New Horizons Computer Learning Centers, Inc.

50400A:

Designing, Optimizing, and Maintaining a Database Administrative Solution for Microsoft® SQL Server® 2008

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Information in this document, including the URL and other Internet Web site references, is subject to change without notice. Unless otherwise noted, the example companies, organizations, products, domain names, e-mail addresses, logos, people, places, and events depicted herein are fictitious, and no association with any real company, organization, product, domain name, e-mail address, logo, person, place, or event is intended or should be inferred. Complying with all applicable copyright laws is the responsibility of the user. Without limiting the rights under copyright, no part of this document may be reproduced, stored in or introduced into a retrieval system, or transmitted in any form or by any means (electronic, mechanical, photocopying, recording, or otherwise), or for any purpose, without the express written permission of Sify Software Limited.

Sify may have patents, patent applications, trademarks, copyrights, or other intellectual property rights covering the subject matter in this document. Except as expressly provided in any written license agreement from Sify, the furnishing of this document does not give you any license to these patents, trademarks, copyrights, or other intellectual property.

The names of manufacturers, products, or URLs are provided for informational purposes only and Sify makes no representations and warranties, either expressed, implied, or statutory, regarding these manufacturers. The inclusion of a manufacturer or product does not imply endorsement of Sify on the manufacturer or product. Links may be provided to third-party sites. Such sites are not under the control of Sify and Sify is not responsible for the contents of any linked site or any link contained in a linked site, or any changes or updates to such sites. Sify is not responsible for webcasting or any other form of transmission received from any linked site. Sify is providing these links to you only as a convenience, and the inclusion of any link does not imply endorsement of Sify of the site or the products contained therein.

© 2010 Sify Software Limited. All rights reserved.

Microsoft and Windows are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

All other trademarks are the property of their respective owners.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

Acknowledgement

Sify Software Limited would like to acknowledge and thank the following for their contribution towards developing this title. Their effort at various stages in the development has ensured that you have a good classroom experience.

Shweta Baijal Mukherjee – Lead Content Developer

Shweta is a Senior Instructional Designer at Sify eLearning Services. She comes with more than five years of experience in the field of Instructional Designing. Currently, she is leading the content development team for the Microsoft ILT courses. She earned her Masters in Computer Applications in 2005 and her Bachelors in Computer Applications in 2002. She has extensive experience working on both technical and non-technical courses.

Kalpana Sharma – Content Developer

Kalpana is an Instructional Designer at Sify eLearning Services. She has three years of experience in Instructional Designing. She completed her Bachelors in Science in 2006. She has successfully completed her ABAP\4 certification in 2007. Her expertise is in designing and developing technical courses.

James Yip - Subject Matter Expert

James has been leading the content development for Eventus and Microsoft Learning since 2007. He has a background as a Project Manager and Solution Architect for a number of large scale projects using Microsoft technologies, such as Windows Server Systems and .NET Development. Before joining Eventus, he was into designing and deploying Microsoft technologies for over eight years.

Mark Fitzgerald – Technical Reviewer

Mark is the Business Intelligence Principal Technologist for QA Ltd in the UK. He holds an honors degree in Computer Studies and multiple certifications in SQL, PerformancePoint, and Windows. With more than 20 years experience in IT, he has become a lead instructor for all Microsoft SQL and BI products within the UK.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Contents

Module 1: Designing an Administrative Solution for SQL Server 2008

Lesson 1: Overview of SQL Server 2008 1-3

Lesson 2: Designing an Administrative Solution 1-28

Lesson 3: Developing and Deploying an Administrative Solution for SQL Server 2008 1-38

Lesson 4: Implementing an Administrative Solution in SSMS 1-44

Lab 1: Designing an Administrative Solution for SQL Server 2008 1-71

Module Review and Takeaways 1-79

Module 2: Deploying SQL Server 2008

Lesson 1: System Requirements for SQL Server 2008 2-3

Lesson 2: Upgrading and Migrating to SQL Server 2008 2-18

Lesson 3: Configuring Instances in SQL Server 2008 2-40

Lab 2: Deploying SQL Server 2008 2-58

Module Review and Takeaways 2-61

Module 3: Designing the Physical Structure of SQL Server 2008

Lesson 1: Planning for Database Files and Transaction Logs 3-3

Lesson 2: Planning for Partitioning 3-28

Lesson 3: Designing a Strategy for Data Compression 3-48

Lab 3: Designing the Physical Structure of SQL Server 2008 3-62

Module Review and Takeaways 3-74

Module 4: Designing a Strategy for Maintaining a Database in SQL Server 2008

Lesson 1: Designing a Strategy for Maintaining Statistics for a Database 4-3

Lesson 2: Designing a Strategy for Maintaining Indexes 4-15

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

Lesson 3: Planning for Full-Text Indexing 4-46

Lab 4: Designing a Strategy for Maintaining a Database in SQL Server 2008 4-69

Module Review and takeaways 4-82

Module 5: Designing Solutions for Managing SQL Server 2008

Lesson 1: Designing a Policy-Based Management Strategy 5-3

Lesson 2: Controlling Resource Usage by Using Resource Governor 5-18

Lesson 3: Auditing Database Changes 5-37

Lab 5: Designing Solutions for Managing SQL Server 2008 5-65

Module Review and Takeaways 5-79

Module 6: Automating the Database Management Strategy for SQL Server 2008

Lesson 1: Managing SQL Server Agent 6-3

Lesson 2: Automating Administrative Tasks in SQL Server 6-29

Lesson 3: Monitoring Events in SQL Server 6-51

Lab 6: Automating the Database Management Strategy for SQL Server 2008 6-62

Module Review and Takeaways 6-72

Module 7: Designing a Strategy for Securing SQL Server 2008

Lesson 1: Securing SQL Server 2008 Instances 7-3

Lesson 2: Planning Database Security 7-40

Lesson 3: Managing Certificates and Keys in SQL Server 2008 7-63

Lesson 4: Implementing Database Encryption 7-80

Lab 7: Designing a Strategy for Securing SQL Server 2008 7-91

Module Review and Takeaways 7-101

Module 8: Designing a Strategy for Monitoring SQL Server 2008

Lesson 1: Overview of Monitoring SQL Server 2008 8-3

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Lesson 2: Monitoring SQL Server Instances 8-23

Lesson 3: Monitoring SQL Server Databases 8-36

Lesson 4: Monitoring Multiple SQL Server by Using Data Collection 8-74

Lab 8: Designing a Strategy for Monitoring SQL Server 2008 8-97

Module Review and Takeaways 8-106

Module 9: Designing a Strategy for Content Distribution in SQL Server 2008

Lesson 1: Distributing Data in SQL Server 2008 9-3

Lesson 2: Working with Distributed Queries 9-12

Lesson 3: Managing Distributed Content by Using Distributed Transactions 9-27

Lesson 4: Managing SSIS Packages 9-38

Lesson 5: Setting Up Scalable Shared Databases 9-60

Lab 9: Designing a Strategy for Content Distribution in SQL Server 2008 9-74

Module Review and Takeaways 9-82

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

About This Course

This section provides information about the course, the audience, student prerequisites, and the course objectives.

Course Description

This five-day instructor-led course is for IT professionals who design and maintain SQL Server databases. This course will help students to design, optimize, and maintain a database administrative solution for Microsoft SQL Server™ 2008.

Audience for this course

Students who take up this course should have three or more years of experience working on databases for two or more of the following phases in the product lifecycle—design, development, deployment, optimization, maintenance, and support:

- Administering databases
- Designing logical database schema solutions
- Defining high-availability solutions
- Automating administrative tasks
- Defining security solutions
- Monitoring and troubleshooting the database server
- Designing and executing deployments
- Defining the infrastructure (storage, hardware, and number of servers or instances, and so on)

Students can also be technical architects and consultants who design and implement SQL Server solutions.

Student Prerequisites

This course expects students to:

- Understand the tradeoffs among the different redundant storage types. For example, what raid levels mean, and how they differ from storage area networks (SANs).

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

- Understand how replication works and how replication is implemented.
- Be familiar with reading user requirements and business-need documents. For example, development project vision/mission statements or business analysis reports.
- Have some knowledge of how queries execute. Must be able to read a query execution plan.
- Have basic knowledge of the dependencies between system components.
- Be able to design a database to third normal form (3NF) and know the tradeoffs when backing out of the fully normalized design (denormalization).
- Be able to design for performance and business requirements in addition to being familiar with design models, such as Star and Snowflake schemas.
- Have monitoring and troubleshooting skills.
- Have knowledge of the operating system and platform, including how the operating system integrates with the database, what the platform or operating system can do, and how the interaction between the operating system and the database works. For example, how integrated authentication interacts with Active Directory® directory service.
- Have knowledge of application architecture; that is, how applications can be designed in three layers, what applications can do, interaction between applications and the database, interaction between the database, and the platform or operating system.
- Must already know how to use:
  - A data modeling tool
  - Microsoft Office Visio® (to create infrastructure diagrams)
  - Be familiar with SQL Server 2008 features, tools, and technologies.
  - Have a Microsoft Certified Technology Specialist: Microsoft SQL Server 2008 credential or equivalent experience.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Objectives of this Course

After completing this course, students will be able to:

- Design an administrative solution for SQL Server 2008.
- Deploy SQL Server 2008.
- Design the physical structure of SQL Server 2008.
- Design a strategy for maintaining SQL Server 2008.
- Design solutions for managing SQL Server 2008.
- Automate database managing strategy for SQL Server 2008.
- Design a strategy for securing databases in SQL Server 2008.
- Design a strategy for monitoring SQL Server 2008.
- Design a strategy for content distribution in SQL Server 2008.
- Design a strategy for replication in SQL Server 2008.
- Design a high-availability solution for SQL Server 2008.
- Design a strategy for backup and recovery solution in SQL Server 2008.

Course Outline

This section provides an outline of the course:

- **Module 1 - Designing an Administrative Solution for SQL Server 2008**: This module provides an overview of an administrative solution and how it helps the administrator.
- **Module 2** - **Deploying SQL Server 2008**: This module describes the considerations for implementing a new version of SQL Server 2008.
- **Module 3** - **Designing the Physical Structure of SQL Server 2008**: This module describes file placement considerations for different functionalities in SQL Server 2008 that help to improve performance and availability.
- **Module 4** - **Designing a Strategy for Maintaining a Database in SQL Server 2008**: This module describes the importance of keeping the database up-to-date and the database performing in optimal state.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

- **Module 5 - Designing Solutions for Managing SQL Server 2008**: This module describes the new approach of managing SQL Servers that has been introduced in SQL Server 2008.
- **Module 6 - Automating the Database Management Strategy for SQL Server 2008**: This module describes the use of SQL Server Agent and scripts, such as VBScripts and SQL PowerShell, to automate some routine jobs that administrators often perform.
- **Module 7 - Designing a Strategy for Securing SQL Server 2008**: This module presents information related to database access control and how SQL Server manages keys and certificates.
- **Module 8 - Designing a Monitoring Strategy for SQL Server 2008**: This module describes the information on monitoring a SQL Server 2008 server to ensure that it is performing properly and is in a healthy state.
- **Module 9 - Designing a Strategy for Content Distribution in SQL Server 2008**: This module describes the different methods of distributing data to multiple servers. The module also describes how to use remote query with linked server and the considerations for using distributed transactions on query span across multiple servers.
- **Module 10 - Designing a Strategy for Replication in SQL Server 2008**: This module provides more in-depth discussions on using replication. The module also describes how to use replication, specifically snapshot and transactional replication, which is one-way replication, to replicate data to other servers.
- **Module 11 - Designing a High-Availability Solution for SQL Server 2008**: This module describes the different high-availability features in SQL Server 2008. The module also covers the scenarios of using them in an organization.
- **Module 12 - Designing a Backup and Recovery Strategy for SQL Server 2008**: This module describes the importance of having a backup of the database. In case of a disaster, backing up the database helps to recover data that cannot be protected by using high-availability solutions mentioned in the previous module.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Course Materials

The following materials are included with the course kit:

- **Course Handbook**. Presents technical information in a simple, direct manner. Only relevant information that enhances the learning experience in class has been presented.
- **PowerPoint Slide Decks**: Provide Key Points, relevant graphics, and illustrations to enhance learning. Also provides instructor notes and additional links for the instructor to effectively prepare for the class.
- **Labs**: Provide a real-world, hands-on experience in the application to help you apply the knowledge and skills learned in the content section of the module.
- **Lab Answer Keys**: Provide the step-by-step detailed steps on labs exercises.
- **Virtual Machine Build Guide**: Provides the detailed steps required to re-create the Virtual Machine and Server images with the necessary configuration.
- **Student Course Files**: Provide a self-extracting executable file named, **Allfiles.exe**, which contains the supporting files for the labs and demonstrations to work as required.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Virtual Machine Environment

This section provides information for setting up the classroom environment to support the business scenario of the course.

Virtual Machine Configuration

In this course, you will use Microsoft Virtual Server 2005 R2 SP1 to perform the labs.

Important: At the end of each lab, you must close the virtual machine and must not save any changes. To close a virtual machine without saving the changes, perform the following steps:

1. In the virtual machine, on the **Action** menu, click **Close**.
2. In the **Close** dialog box, in the **What do you want the virtual machine to do?** list, click **Turn off and delete changes**, and then click **OK**.

The following table shows the role of each virtual machine used in this course:

Virtual machine	Role
50400A-NYC-SQL1	SQL Server 2008

Software Configuration

The following software is installed on each virtual machine (VM):

- Microsoft SQL Server 2008 Developer Edition
- Microsoft Visual Studio 2008 Professional Edition
- Microsoft Office Excel 2007

Course Files

There are files associated with the labs in this course. The lab files are located in the folder, D:\Labfiles\ModXX, on the student computers.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Classroom Setup

Each classroom computer will have the same virtual machine configured in the same way.

Note: The content coverage of this course is about Microsoft products and technologies. Therefore, the content in this course has been paraphrased from Microsoft Web sites and other Web sites that discuss Microsoft products and technologies. Wherever it has been inappropriate to rewrite or paraphrase the content, the original content has been retained. All content sources have been acknowledged in the **Source** section of each topic.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

END USER LICENSE AGREEMENT

This End-User License Agreement ("Agreement") is a legal agreement between Sify Software Limited ("Sify") and you ("Licensee"). By using the licensed content, you accept these terms. If you do not accept them, do not use the licensed content. If you comply with these license terms, you have the rights below.

1. Sify owns or licenses all intellectual property rights (including, without limitation, copyright, trademark, domain name, trade secret, and patent rights) in all materials provided by Sify in connection with this Agreement, including without limitation any and all materials or other works created by Sify relating to the delivery of Course, including any structure or code licensed, developed, or created by Sify or at the direction of Sify. The licensed content is licensed on a per copy per device basis.

2. Sify may change any of the terms in this Agreement. When it does so, it will update the date of the agreement at the top of the agreement.

3. Sify hereby grants a limited, non-exclusive license to Licensee to use the Course Materials for the sole and exclusive purpose of learning the subject matter. This license is personal to Licensee and non-transferable. Licensee may not sublicense, alter, or disclose the Course Materials to third parties in any form for any reason. Licensee is responsible for all uses of the Course Materials through Licensee's Login Name and Password, and for protecting and securing such Login Name and Password from unauthorized use and disclosure. Licensee will not share Licensee's Login Name and Password with any third party. In the event that Licensee becomes aware of or believes there has been any breach of the Course Materials, such as theft or unauthorized use of Licensee's Login Name and Password, Licensee will notify Sify immediately. This license will terminate immediately and without notice in the event Sify determines, at its sole discretion that Licensee has failed to comply with the terms of the Agreement.

4. License commences the day Licensee registers for Course and runs for the specific term purchased, whether or not the Course is used.

5. INSTALLATION AND USE RIGHTS.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

- **Licensed Device.** The licensed device is the device on which you use the licensed content. You may install and use one copy of the licensed content on the licensed device.
- **Portable Device.** You may install another copy on a portable device for use by the single primary user of the licensed device.
- **Separation of Components.** The components of the licensed content are licensed as a single unit. You may not separate the components and install them on different devices.
- **Third Party Programs.** The licensed content may contain third party programs. These license terms will apply to your use of those third party programs, unless other terms accompany those programs.
- **Feedback.** If you agree to give feedback about the licensed content to Sify, you give to Sify, without charge, the right to use, share and commercialize your feedback in any way and for any purpose. You will not give feedback that is subject to a license that requires Sify to license its software or documentation to third parties because we include your feedback in them. These rights survive this agreement.
- **Confidential Information.** The licensed content, including any viewer, user interface, features and documentation that may be included with the licensed content, is confidential and proprietary to Sify and its suppliers.
  i. **Use.** For five years after installation of the licensed content or its commercial release, whichever is first, you may not disclose confidential information to third parties. You may disclose confidential information only to your employees and consultants who need to know the information. You must have written agreements with them that protect the confidential information at least as much as this agreement.
  ii. **Survival.** Your duty to protect confidential information survives this agreement.
  iii. **Exclusions**. You may disclose confidential information in response to a judicial or governmental order. You must first give written notice to Sify to allow it to seek a protective order or otherwise protect the information. Confidential information does not include information that
    - becomes publicly known through no wrongful act;
    - you received from a third party who did not breach confidentiality obligations to Sify or its suppliers; or
    - you developed independently.

6. ADDITIONAL LICENSING REQUIREMENTS AND/OR USE RIGHTS.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

- **Media Elements and Templates.** You may use images, clip art, animations, sounds, music, shapes, video clips and templates provided with the licensed content solely for your personal training use.
- **Academic Materials.** If the licensed content contains academic materials (such as white papers, labs, tests, datasheets and FAQs), you may copy and use the academic materials. You may not make any modifications to the academic materials and you may not print any book (either electronic or print version) in its entirety. If you reproduce any academic materials, you agree that:
  - The use of the academic materials will be only for your personal reference or training use
  - You will not republish or post the academic materials on any network computer or broadcast in any media;
  - You will include the academic material's original copyright notice, or a copyright notice to Sify's benefit in the format provided below:

    Form of Notice:

    © 2010 Reprinted for personal reference use only with permission by Sify Software Limited. All rights reserved.

7. **INTERNET-BASED SERVICES.** Sify may provide Internet-based services with the licensed content. It may change or cancel them at any time. You may not use these services in any way that could harm them or impair anyone else's use of them. You may not use the services to try to gain unauthorized access to any service, data, account or network by any means.

8. **SCOPE OF LICENSE.** The licensed content is licensed, not sold. This agreement only gives you some rights to use the licensed content. Sify reserves all other rights. Unless applicable law gives you more rights despite this limitation, you may use the licensed content only as expressly permitted in this agreement. In doing so, you must comply with any technical limitations in the licensed content that only allow you to use it in certain ways. You may not
   - work around any technical limitations in the licensed content;
   - reverse engineer, decompile or disassemble the licensed content, except and only to the extent that applicable law expressly permits, despite this limitation;
   - make more copies of the licensed content than specified in this agreement or allowed by applicable law, despite this limitation;
   - publish the licensed content for others to copy;

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

- transfer the licensed content;
- allow others to access or use the licensed content;
- rent, lease or lend the licensed content; or
- use the licensed content for commercial licensed content hosting services.
- have Rights to change or modify the server software that may be included with the Licensed Content, including the Virtual Hard Disks does not give you any right to implement Sify patents or other Sify intellectual property in software or devices that may access the server.

9. **TRANSFER TO ANOTHER DEVICE.** You may uninstall the licensed content and install it on another device for your personal training use. You may not do so to share this license between devices.

10. **TRANSFER TO A THIRD PARTY.** You may not transfer those versions marked as 'Beta Content' or 'Pre-release' to a third party. For final versions, these terms apply: The first user of the licensed content may transfer it and this agreement directly to a third party. Before the transfer, that party must agree that this agreement applies to the transfer and use of the licensed content. The first user must uninstall the licensed content before transferring it separately from the device. The first user may not retain any copies.

11. **NOT FOR RESALE SOFTWARE /LICENSED CONTENT.** You may not sell software or licensed content marked as "NFR" or "Not for Resale."

12. As a participant in Course, Licensee owns any original material that Licensee creates and enters online as part of Course, including Licensee's comments and testimonials about the quality of Course and the instructor, and messages posted in chat and threaded discussion or sent via email to Course instructor ("Licensee's Content"). Licensee hereby gives Sify an unlimited, worldwide and royalty-free license for the use and exploitation of Licensee's Content.

13. Licensee agrees that Sify may collect, use, sell, license and otherwise distribute demographic data about users in aggregated form (so that no Licensee is personally identifiable) to assist in market evaluation and for other purposes as reasonably determined by Sify.

14. Some activities within Course may request users to submit information that is tabulated and stored in a database. The data obtained is used to produce statistics

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

that are integral to the educational value of these activities. Anywhere the requested data is personal in nature, a user's specific responses remain confidential and are not associated with name or e-mail address in any way.

15. Sify will not disclose personally identifiable information associated with Licensee's use of the Website (e.g., name, address, access code) to any third parties not affiliated with Sify except to duly authorized investigative agencies for the limited purpose of investigating a good faith claim that a particular user has violated the restrictions on use of or access to Course.

16. **DISCLAIMER OF WARRANTY: THE LICENSED CONTENT IS LICENSED "AS-IS." YOU BEAR THE RISK OF USING IT. SIFY GIVES NO EXPRESS WARRANTIES, GUARANTEES OR CONDITIONS. YOU MAY HAVE ADDITIONAL CONSUMER RIGHTS UNDER YOUR LOCAL LAWS WHICH THIS AGREEMENT CANNOT CHANGE. TO THE EXTENT PERMITTED UNDER YOUR LOCAL LAWS, SIFY EXCLUDES THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NON-INFRINGEMENT. SIFY DOES NOT WARRANT OR GUARANTEE THAT COURSE WILL SATISFY THE INTENDED USE BY LICENSEE. SIFY IS NOT LIABLE FOR ANY INDIRECT, INCIDENTAL OR CONSEQUENTIAL DAMAGES, OR FOR LOSS OF PROFITS, REVENUE, OR DATA, WHETHER IN AN ACTION IN CONTRACT, TORT, PRODUCT LIABILITY, STATUTE OR OTHERWISE, EVEN IF ADVISED OF THE POSSIBILITY OF THOSE DAMAGES. SIFY'S LIABILITY WITH RESPECT TO COURSE PROVIDED HEREUNDER WILL NOT EXCEED THE AMOUNT LICENSEE PAID FOR COURSE, EVEN IF ANY TERM OF THIS AGREEMENT FAILS OF ITS ESSENTIAL PURPOSE.**

17. **ENTIRE AGREEMENT.** This agreement, and the terms for supplements, updates, Internet-based services and support services that you use, are the entire agreement for the licensed content and support services.

18. It also applies even if Sify knew or should have known about the possibility of the damages. The above limitation or exclusion may not apply to you because your country may not allow the exclusion or limitation of incidental, consequential or other damages.

19. Licensee hereby indemnifies and will defend and hold Sify harmless against all claims, damages, injuries or deaths, judgments, liabilities and expenses (including legal costs and attorney fees) arising from or in connection with Licensee's breach of the conditions to this License, or Licensee's negligence or violation or alleged

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

violation of Sify's proprietary information, trademark, copyright, or patent. The obligation to indemnify will survive the termination of this Agreement.

20. This is the entire Agreement between the parties as to the subject matter hereof and supersedes all other agreements, express or implied, written or oral. No modification of, or amendment to, this Agreement shall be binding unless in writing and signed by both parties. Invalidity or unenforceability of one or more provisions of this Agreement shall not affect any other provision of this Agreement. Neither failure nor delay on the part of any party to exercise any right, remedy, power, or privilege hereunder nor course of dealing between the parties shall operate as a waiver thereof, or of the exercise of any other right, remedy, power, or privilege.

21. This Agreement, and any disputes or disagreements concerning Course or Course Materials, will be governed by and construed according to the laws of India, without regard to its choice of law rules, and specifically excluding the U.N. Convention on Contracts for the International Sale of Goods. Any lawsuit filed regarding this Agreement must be filed in courts in Chennai, India. Any notices required or authorized to be given hereunder shall be deemed to be given when mailed by certified or registered mail, postage prepaid, as follows:

To Sify:
Legal Department
Sify Software Limited
2nd Floor, Tidel Park,
No.4, Rajiv Gandhi Salai,
Taramani, Chennai 600113, INDIA

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Module 1

Designing an Administrative Solution for SQL Server 2008

Contents:

Lesson 1: Overview of SQL Server 2008 1-3

Lesson 2: Designing an Administrative Solution 1-28

Lesson 3: Developing and Deploying an Administrative Solution for SQL Server 2008 1-38

Lesson 4: Implementing an Administrative Solution in SSMS 1-44

Lab 1: Designing an Administrative Solution for SQL Server 2008 1-71

Module Review and Takeaways 1-79

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Module Overview

- Overview of SQL Server 2008
- Designing an Administrative Solution
- Developing and Deploying an Administrative Solution for SQL Server 2008
- Implementing an Administrative Solution in SQL Server Management Studio (SSMS)

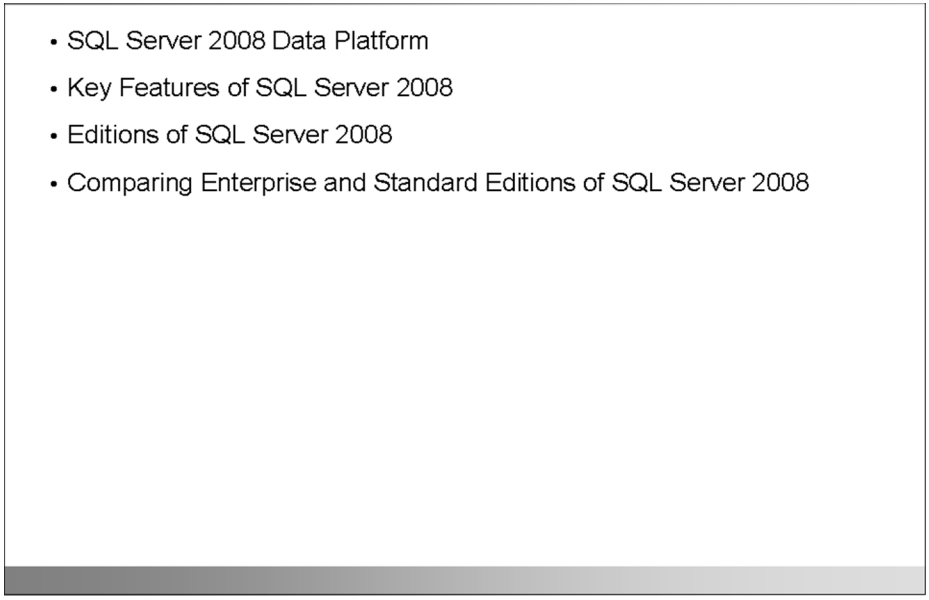
Managing multiple servers would be a complicated and time-consuming task. Microsoft™ SQL Server™ 2008 provides a number of features that help you to reduce the time for administering SQL Servers, and simplify database administration by reducing the workload.

In this module, you will learn the considerations for implementing administrative solutions for SQL Server 2008. You will also learn the process to design an administrative solution for SQL Server 2008 to reduce the administrative cost for managing SQL Server 2008.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Lesson 1
Overview of SQL Server 2008



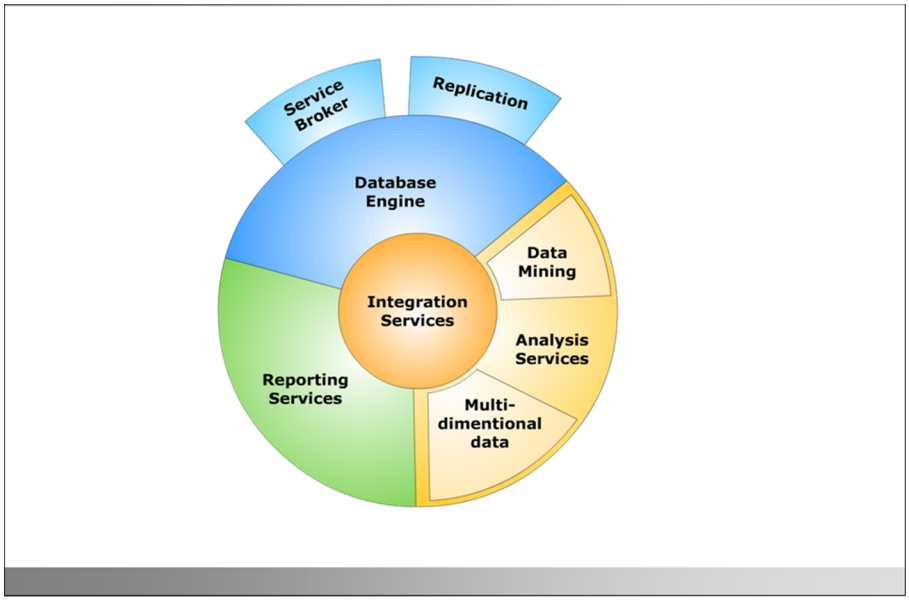
SQL Server 2008 provides a data platform that enables you to run the critical applications. Productivity enhancements help you reduce the cost of managing your data infrastructure while streamlining development of applications. Platform enhancements help you ensure the data security and high availability with enterprise performance and scalability.

In this lesson, you will learn about the basics of SQL Server platform, the editions, and the difference between Standard and Enterprise Edition of SQL Server 2008.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

SQL Server Architecture



Key Points

SQL Server 2008 provides a set of components that enable you to help large enterprise data processing systems to effectively manage data and to produce robust database applications. The various components of the SQL Server 2008 are database engine, SQL Server 2008 Analysis Services (SSAS), integration services, replication, SQL Server 2008 Reporting Services (SSRS), and SQL Server Service Broker.

Database engine

The database engine is the core service for storing, processing, and securing data. The database engine provides controlled access and rapid transaction processing to meet the requirements of the data consuming applications within your enterprise. Use the database engine to create relational databases for online transaction processing or online analytical processing data. This includes creating tables for storing data, and database objects such as indexes, views, and stored procedures for viewing, managing, and securing data. You

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

can use SQL Server Management Studio to manage the database objects, and SQL Server Profiler for capturing server events.

SSAS

SSAS helps you design, create, and manage multidimensional structures that contain detail and aggregated data from multiple data sources, such as relational databases, in a single unified logical model supported by built-in calculations. It provides fast, intuitive, top-down analysis of large quantities of data built on this unified data model, which can be delivered to users in multiple languages and currencies. SSAS works with data warehouses, data marts, production databases, and operational data stores, supporting analysis of both historical and real time data.

SSAS contains the following features and tools that you need to create complex data mining solutions:

- A set of industry-standard data mining algorithms.
- The Data Mining Designer that you can use to create, manage, and explore data mining models, and to create predictions by using those models.
- The Data Mining Extensions (DMX) language that you can use to manage mining models and to create complex prediction queries.

You can use a combination of these features and tools to discover trends and patterns that exist in the data, and then use the trends and patterns to make intelligent decisions about difficult business problems.

Integration services

Integration services is a platform for building enterprise-level data integration and data transformations solutions. You can use Integration Services to solve complex business problems by copying or downloading files, sending e-mail messages in response to events, updating data warehouses, cleaning and mining data, and managing SQL Server objects and data. The packages can work alone or in concert with other packages to address complex business needs. Integration services can extract and transform data from a wide variety of sources, such as XML data files, flat files, and relational data sources, and then load the data into one or more destinations.

Integration services includes a rich set of built-in tasks and transformations, tools for constructing packages, and services for

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

running and managing packages. You can use the graphical integration services tools to create solutions without writing a single line of code, to program the extensive integration services object model, to create packages programmatically, and to code custom tasks and package objects.

Replication

Replication is a set of technologies for copying and distributing data and database objects from one database to another and then synchronizing between databases to maintain consistency. Using replication, you can distribute data to different locations and to remote or mobile users over local and wide area networks, dial-up connections, wireless connections, and the Internet. SQL Server 2008 supports two types of replication—Transactional and Merge.

Transactional replication is used in following server-to-server scenarios that require:

- High throughput
- High scalability and availability
- Data warehousing and reporting
- Integrating data from multiple sites
- Integrating heterogeneous data
- Offloading batch processing

Merge replication is designed for mobile applications or distributed server applications that have possible data conflicts such as, exchanging data with mobile users, consumer point of sale (POS) applications, and integration of data from multiple sites. Snapshot replication is used to provide the initial data set for transactional and merge replication. This replication can also be used when complete refreshes of data are appropriate.

With these two types of replication, SQL Server 2008 provides an efficient and flexible system for synchronizing data across your enterprise. In addition to replication, you can sychronize databases by using Microsoft Sync Framework and Sync Services for ADO.NET. Sync Services for ADO.NET provides an intuitive and flexible API that you can use to build applications that target offline and collaboration scenarios.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

SSRS

SSRS provides various tools and services to help you create, deploy, and manage reports for your organization, as well as programming features that enable you to extend and customize your reporting functionality.

SSRS is a server-based reporting platform that provides comprehensive reporting functionality for a variety of data sources. Reporting Services includes a complete set of tools for you to create, manage, and deliver reports, and APIs that enable developers to integrate or extend data and report processing in custom applications. Reporting Services tools work within the Microsoft Visual Studio environment and are fully integrated with SQL Server tools and components.

With SSRS, you can create interactive, tabular, graphical, or free-form reports from relational, multidimensional, or XML-based data sources. You can publish reports, schedule report processing, or access reports on-demand. SSRS also enables you to create ad hoc reports based on predefined models, and to interactively explore data within the model. You can select from a variety of viewing formats, export reports to other applications, and subscribe to published reports. The reports that you create can be viewed over a Web-based connection or as part of a Microsoft Windows application or SharePoint site. Reporting Services provides the key to your business data.

SQL Server Service Broker

SQL Server Service Broker provides the SQL Server Database Engine native support for messaging and queuing applications. This makes it easier for developers to create sophisticated applications that use the Database Engine components to communicate between disparate databases. Developers can use Service Broker to easily build distributed and reliable applications.

Application developers who use Service Broker can distribute data workloads across several databases without programming complex communication and messaging internals. This reduces development and test work because Service Broker handles the communication paths in the context of a conversation. It also improves performance. For example, front-end databases supporting Web sites can record information and send process intensive tasks to queue in back-end

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

databases. Service Broker ensures that all tasks are managed in the context of transactions to assure reliability and technical consistency.

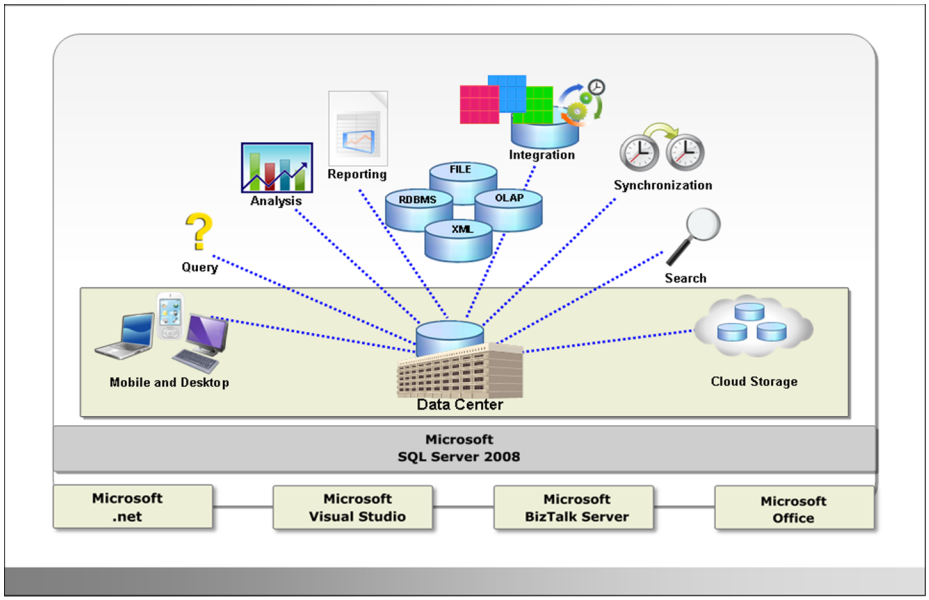
Source:

- http://msdn.microsoft.com/en-us/library/ms130214.aspx
- http://msdn.microsoft.com/en-us/library/ms172359.aspx
- http://msdn.microsoft.com/en-us/library/ms166352(SQL.90).aspx
- ttp://msdn.microsoft.com/en-us/library/ms159106.aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

SQL Server 2008 Data Platform



Key Points

SQL Server provides a trusted, productive, and intelligent data platform.

The data platform enhancements help you ensure the data security and high availability with enterprise performance and scalability. Productivity enhancements in SQL Server 2008 help you reduce the cost of managing your data infrastructure while streamlining development of applications. SQL Server 2008 consists of a business intelligence platform for data integration, analysis, and reporting solutions.

Features of SQL Server 2008 Data Platform

- **Mobile and Desktop**. SQL Server 2008 helps you have control over data stored on the servers within the data center or desktops or mobile devices.
- **Query, Analysis, Reporting, Synchronization, Search, and**

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Integration. SQL Server 2008 delivers a rich set of integrated services that you can use to query, search, synchronize, report, and analyze data.

- **Cloud Storage and Data Center**. Microsoft SQL Data Services (SDS) is a cloud-based relational database platform built on SQL Server technologies. With SDS, you can easily provision and deploy relational database solutions to the cloud, and take advantage of a globally distributed data center. This data center provides enterprise-class availability, scalability, and security with the benefits of built-in data protection and disaster recovery. You can use data stored in other RDBMS, OLAP, XML and even FILE with SQL Server 2008.
- **External applications**. You can use SQL Server 2008 data within custom applications developed by using Microsoft .NET and Visual Studio®. You can also use the data within a service-oriented architecture (SOA) and a business process by using Microsoft BizTalk® Server. Users can access data directly by using various tools, such as the 2007 Microsoft Office system.

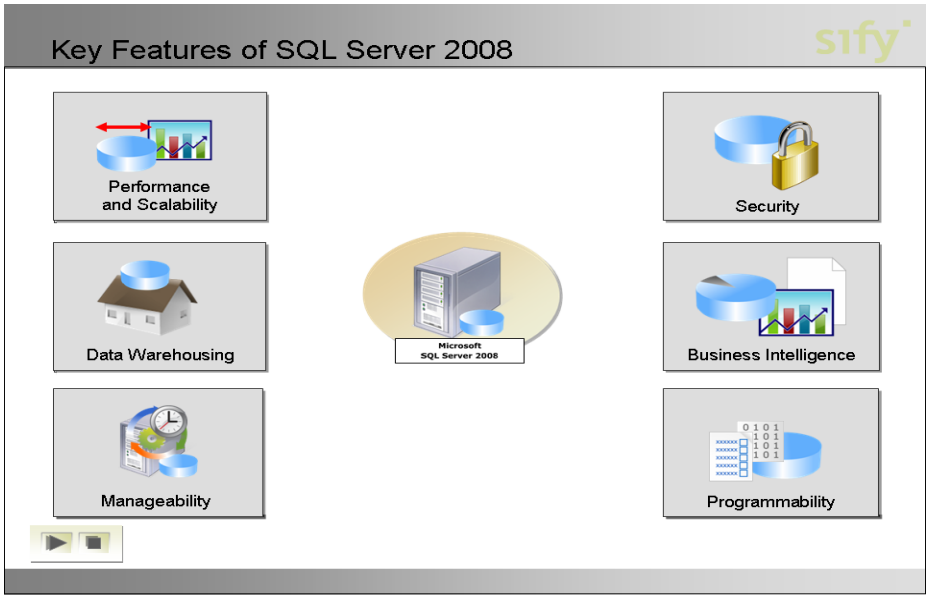
Source:

- http://www.microsoft.com/sqlserver/2008/en/us/overview.aspx
- http://www.microsoft.com/belux/msdn/nl/chopsticks/default.aspx?id=1262
- http://www.microsoft.com/presspass/presskits/sqlserver/docs/SQLSERVER2008FAQ.doc

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Key Features of SQL Server 2008



Key Points

SQL Server 2008 provides features that include performance and scalability, security, data warehousing, business intelligence (BI), manageability, and programmability.

Performance and Scalability

The various tools for better performance and scalability of SQL Server 2008 are:

- **Resource Governor**. SQL Server 2008 provides a new monitoring tool called Resource Governor. By using Resource Governor, you can define resource limits and set priority for different workloads. Resource Governor also helps you to identify which users are running what queries and how long the queries are running, enabling you to optimize server utilization.
- **Data Collection**. SQL Server 2008 provides a new central repository known as Data Collection for storing performance data. SQL Server extends the feature of adding memory resources

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

online with the help of the Hot Add CPU feature. This feature allows a database to be scaled out on demand. You can also add CPU resources to SQL Server 2008 on supported hardware platforms without requiring the application to close.

- **SQL Server Analysis Services (SSAS) and SQL Server Reporting Services (SSRS)**. SQL Server helps you build high-performance analysis solutions with scalability and performance enhancements in SSAS. You can take advantage of on-demand processing and obtain better performance with the re-engineered reporting engine for SSRS.

Data Warehousing

The data warehousing feature of SQL Server 2008 helps you build faster warehouses, manage growing data volumes, integrate with Microsoft business intelligence products, and scale a data warehouse to improve performance by using clustering. The improvements in this feature include:

- **Data Compression**. Data compression built into SQL Server 2008 helps organizations store data efficiently by reducing the amount of storage space needed to store tables and indexes, and reducing the I/O request required to access the data. You need not make any changes to applications to enable this feature.
- **Partitioned Table Parallelism**. Partitions help organizations to manage large, growing tables effectively by transparently breaking them into manageable blocks of data. SQL Server 2008 enhances the partitioning feature of SQL Server 2005 by improving performance on large, partitioned tables.
- **Star Join Query Optimizations**. SQL Server 2008 provides improved query performance for common data warehouse scenarios. Star join query optimizations reduce the query response time by recognizing data warehouse join patterns.
- **Grouping sets**. Grouping sets is an extension to the GROUP BY clause that helps users define multiple groupings in the same query. Grouping sets produce a single result set that is equivalent to a UNION ALL of differently grouped rows, making aggregation, querying, and reporting easier and faster.
- **Change Data Capture (CDC)**. With CDC, you can capture changes and place them in change tables. CDC captures the complete content of changes, maintains cross-table consistency,

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

and even works across schema changes. This enables organizations to integrate the latest information into the data warehouse.

- **MERGE SQL statement**. In a common data warehousing scenario, you can use the MERGE SQL statement to first determine if a row exists and then run an INSERT or UPDATE statement.
- **Scalable Integration Services**. The key improvements in scalability of SQL Services Integration Services (SSIS) are:
  - **SSIS pipeline improvements**. Data Integration packages can now scale more effectively, making use of available resources and managing the largest enterprise-scale workloads. The new design improves the scalability of run time into multiple processors.
  - **SSIS persistent lookups**. The need to perform lookups is one of the most common extraction, transformation, and loading (ETL) operations. This is especially prevalent in data warehousing, where fact records must use lookups to transform business keys to their corresponding surrogates. SSIS improves the performance of lookups to support large tables.
  - **SSIS high-performance connectors**. SSIS has new connectors for SAP Netweaver BI, Oracle, and Teradata, especially designed for high-performance loading of data into an enterprise data warehouse.

Manageability

The following are some of the new features of SQL Server 2008 to manage the servers:

- **Policy-based management**. Microsoft SQL Server 2008 provides a policy-based framework for managing one or more instances of SQL Server, along with tools for performance monitoring, troubleshooting, and tuning, which enable you to more efficiently manage the databases and SQL Server instances. To use this framework, SQL Server policy administrators use SQL Server Management Studio (SSMS) to create policies that manage entities on the server, such as the instance of SQL Server, databases, and other SQL Server objects.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

You can use the built-in framework facets and policies to manage surface area configuration and apply best practices. Policy administrators enable automated policy execution by using one of the Enforce, Check on Changes, and Check on Schedule execution modes.

- **FILESTREAM**. FILESTREAM is a new feature in SQL Server 2008. It allows structured data to be stored in the database and associated unstructured, for example, Binary Large Objects (BLOB), data to be stored directly in the Windows NT® file system (NTFS). You can then access the BLOB data through the high-performance Win32® streaming APIs, rather than having to pay the performance penalty of accessing BLOB data through SQL Server.

Security

The following are some of the security features of SQL Server 2008:

- **Transparent Data Encryption (TDE)**. By using the TDE feature, SQL Server 2008 enables the encryption of entire databases, data files, and log files. Encryption enables organizations to meet the demands for regulatory compliance and the overall concern for data privacy. The benefits of TDE include securing data from unauthorized users and encrypting backups. You can enable these features without changing existing applications.
- **Hardware Security Modules (HSM)**. SQL Server 2008 also provides a comprehensive solution for encryption and key management. To meet the growing need for information security in data centers, some organizations have invested in vendors to manage security keys within an enterprise. SQL Server 2008 provides support for this growing need of information security by supporting third-party key management and Hardware Security Modules products.
- **Database Audit**. In SQL Server 2008, you can improve compliance and security by performing an audit on the data. Auditing can include information about data modifications and the time when data has been read. SQL Server 2008 has features such as enhanced configuration and management of audits, which enable organizations to meet varied compliance needs. For example, a country may require a log to monitor individuals accessing sensitive information. SQL Server 2008 can also define

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

audit specifications in each database, so audit configuration can be ported with databases. By filtering audits to specific objects, you can get a better performance in audit generation and flexibility in configuration.

Business Intelligence (BI)

SQL Server 2008 helps organizations to deliver intelligence by providing a scalable BI platform optimized for data integration, reporting, and analysis. Some of the features of BI include:

- **Design to Scale**. SQL Server 2008 performs a broader analysis with enhanced analytical capabilities and more complex computations and aggregations. New cube design tools help users streamline the development of the analysis infrastructure, enabling them to build solutions for optimized performance. Best Practice Design Alerts have been built into the design, enabling developers to integrate real-time alerts at design time, which optimizes design. The Dimension Designer enables easy viewing and editing of attribute relationships and also provides multiple built-in validations for support of ideal dimension design.
- **Block Computations**. Block computations provide a significant improvement in processing performance, enabling users to increase the depth of their hierarchies and the complexity of the computations.
- **Write-back to multidimensional online analytical processing (MOLAP)**. New MOLAP-enabled write-back capabilities in SSAS do not require querying relational online analytical processing (ROLAP) partitions. This provides users with enhanced write-back scenarios from within analytical applications, without sacrificing online analytical processing (OLAP) performance.
- **Predictive Analysis**. SSAS continues to deliver advanced data mining technologies. Better Time Series support extends forecasting capabilities. Enhanced mining structures deliver more flexibility to perform focused analysis through filtering, in addition to delivering complete information beyond the scope of the mining model in reports. New cross-validation enables confirmation of both accuracy and stability for results that you can trust.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Programmability

In SQL Server 2008, you can easily create the next generation of robust database solutions by using the powerful SQL Server Transact-SQL (T-SQL) programming language and rich data types. These data types can manage virtually any type of data. SQL Server 2008 supports the following development technologies:

- **Microsoft ADO.NET Data Services**. Microsoft ADO.NET Data Services enable Web applications to display data as REpresentational State Transfer (REST)-based data services that client applications can consume in corporate networks and across the Internet.
- **Microsoft ADO.NET Entity Framework**. The Microsoft ADO.NET Entity Framework simplifies data access code in applications by enabling developers to create a customizable conceptual model for data from any source, and easily use it for business requirements.
- **Language-Integrated Query (LINQ)**. LINQ enables data programming developers to build queries over data directly from any .NET-based programming language.
- **Enhanced SQL Server 2008 Native Client**. Enhanced SQL Server 2008 Native Client provides a high-performance, versatile, data access library that enables developers to take advantage of the new capabilities of SQL Server 2008.
- **Microsoft Sync Framework**. The Microsoft Sync Framework is a comprehensive synchronization platform that enables collaboration and offline scenarios for applications, services, and devices. Developers can build synchronized ecosystems that virtually integrate any application and any type of data by using most protocols over virtually any network.

The administrative features of SQL Server 2008 include:

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

- **Database auditing**. Database auditing involves tracking and logging events that occur on the system. There are several levels of auditing for SQL Server, depending on government or standards requirements for your installation. In SQL Server 2008 Enterprise Edition, you can set up automatic auditing by using SQL Server Audit. SQL Server Audit provides the tools and processes necessary to enable, store, and view audits on various server and database objects.
- **Centralized policy-based management**. Policy-based management is a system for managing one or more instances of SQL Server 2008. SQL Server policy administrators create policies to manage entities on the server, such as the instance of SQL Server, databases, or other SQL Server objects by using SSMS. In SQL Server 2008, you can administer multiple servers by assigning Central Management Servers and creating server groups.
- **Enhanced data protection**. Data protection in SQL Server 2008 involves many methods. Database encryption, which is one of the methods to protect data, is the process of encoding data by using a key or password. This can make the data difficult to understand without the corresponding decryption key or password. Encryption does not solve access control problems. However, it enhances security by limiting data loss even if access controls are bypassed. For example, if the database host computer is not configured and an unauthorized user obtains sensitive data, the accessed information might be impossible to understand if it is encrypted.
- **Resources management**. Resource Governor enables you to specify limits on the amount of CPU and memory that incoming application requests can use.

Source:

- http://www.microsoft.com/sqlserver/2008/en/us/white-papers.aspx
- http://www.microsoft.com/sqlserver/2008/en/us/performance-scale.aspx
- https://partner.microsoft.com/download/global/40047533

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

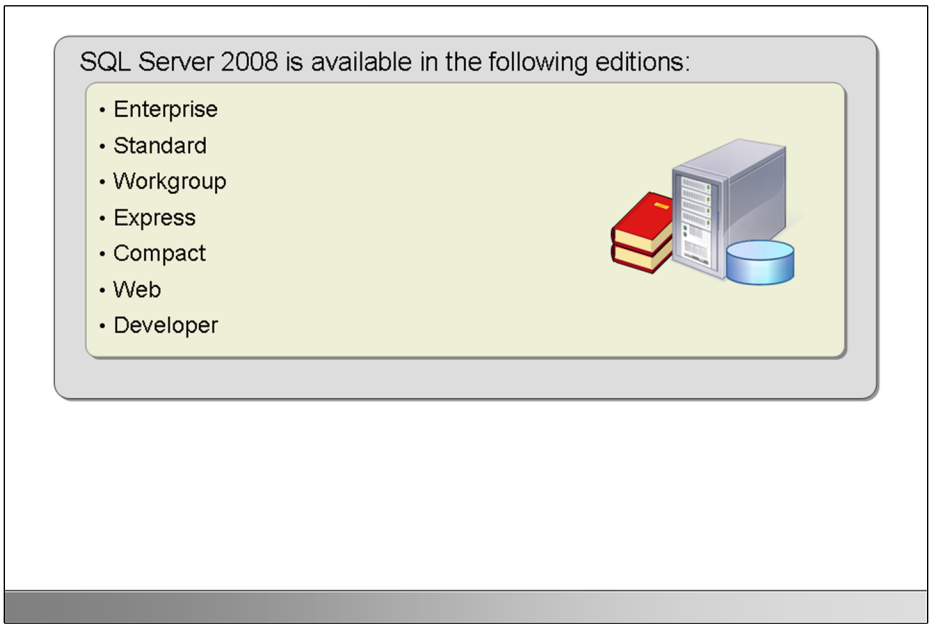
- http://www.microsoft.com/sqlserver/2008/en/us/performance-scale.aspx
- http://download.microsoft.com/download/6/9/d/69d1fea7-5b42-437a-b3ba-a4ad13e34ef6/SQL2008_ProductOverview.docx

For more information, see What's New in SQL Server 2008.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Editions of SQL Server 2008



Key Points

SQL Server 2008 is available in many different editions to suit the needs of the organization and usage. The different editions of SQL Server accommodate the unique performance, runtime, and price requirements of organizations and individuals. The SQL Server components that you install also depend on your specific requirements.

Enterprise Edition

Microsoft SQL Server 2008 Enterprise Edition provides a trusted, productive, and intelligent data platform that enables you to run your most demanding mission-critical applications.

The following are the key features of the Enterprise Edition. The Enterprise Edition:

- Helps reduce time and cost of development.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

- Facilitates easier management of applications.
- Meets the high demands of enterprise online transaction processing (OLTP) and data warehousing applications.

Standard Edition

Standard Edition is a complete data management and Business Intelligence (BI) platform for departmental applications. The Standard Edition:

- Helps reduce operational and development overhead.
- Provides an easy-to-use database platform for small scale-to-medium scale database solutions.

Workgroup Edition

Workgroup Edition is a reliable data management and reporting platform to run branch applications. It includes the core database features of the SQL Server product line and is easy to upgrade to the Standard Edition or the Enterprise Edition.

Express Edition

Express Edition is a free edition of SQL Server, ideal for learning and building desktop and small server applications, and for redistribution by independent software vendors (ISVs).

Express Edition has powerful features that include:

- SSRS, a server-based platform for creating and delivering traditional and interactive reports.
- A graphical management tool for managing the database.

Compact Edition

Compact Edition is a free SQL Server embedded database ideal for building stand-alone and occasionally connected applications for mobile devices, desktops, and clients. It runs on all Windows® operating systems, including Windows XP and Windows Vista®, and on Pocket PCs and smart phones.

Web Edition

Web Edition is specifically designed for Web-serving environments. It provides low cost, large scale, highly available Web applications or hosting solutions to your customers. Microsoft SQL Server 2008 Web

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Edition delivers an optimal balance of scalability, manageability, and cost for hosted Web sites and applications.

Developer Edition

Developer Edition enables developers to build and test applications that run on SQL Server on 32-bit, IA64, and x64 platforms. SQL Server 2008 Developer includes all the functionality of Enterprise Edition, but is licensed only for use during development, testing, and demonstration. This license entitles a developer to use the software on multiple systems.

The following table compares the core editions of SQL Server 2008.

Comparison point	Enterprise	Standard
Target Scenarios	It can be used for enterprise workloads that need redundancy and built-in Business Intelligence.	It can be used in shared data scenarios in departments and small-to-large businesses.
CPU	OS Maximum	4 CPUs
Memory	OS Maximum	OS Maximum
DB Size	Unlimited	Unlimited

The following table compares the specialized editions of SQL Server 2008.

Workgroup	Web	Developer
It can be used in remote offices that need local instances of company data	It can be used for web application hosting	It is a full featured edition but can only be used for development and testing
2 CPUs	4 CPUs	OS Maximum
4 GB	OS Maximum	OS Maximum
Unlimited	Unlimited	Unlimited

The following table compares the free editions of SQL Server 2008.

Comparison point	Express	Compact

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Comparison point	Express	Compact
Target Scenarios	It consists of an entry level database that is ideal for learning and ISV redistribution.	It consists of an embedded database for developing desktop and mobile applications.
CPU	1 CPUs	OS Maximum
Memory	1 GB	OS Maximum
DB Size	4 GB	4 GB

Question: Which edition can be used by a developer to develop applications and test the functionalities on SQL Server 2008?

Question: Which edition provides a free and cost effective database engine for companies with less than 4GB of data and less than 10 users?

Source:

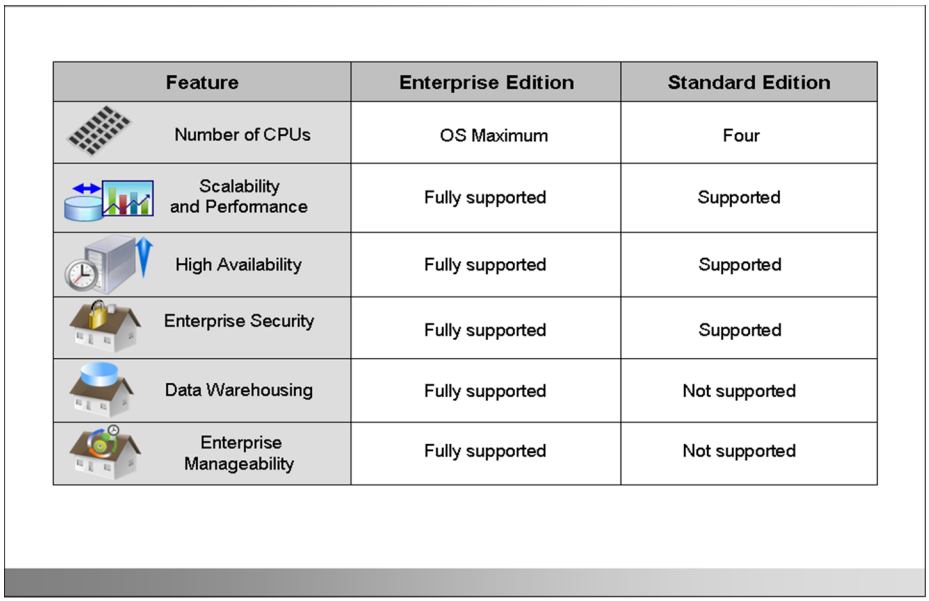
- http://download.microsoft.com/download/1/E/5/1E57389D-3623-4AEE-A109-374F5D41B9D3/SQL%20Server%202008%20For%20Hosters%20Whitepaper.docx
- http://www.microsoft.com/sqlserver/2008/en/us/workgroup.aspx
- http://www.microsoft.com/uk/sqlserver/itpro.mspx
- http://www.microsoft.com/sqlserver/2008/en/us/compact.aspx
- http://www.microsoft.com/sqlserver/2008/en/us/Developer.aspx

For more information, see Editions.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Comparing Enterprise and Standard Editions of SQL Server 2008



Feature	Enterprise Edition	Standard Edition
Number of CPUs	OS Maximum	Four
Scalability and Performance	Fully supported	Supported
High Availability	Fully supported	Supported
Enterprise Security	Fully supported	Supported
Data Warehousing	Fully supported	Not supported
Enterprise Manageability	Fully supported	Not supported

Key Points

SQL Server 2008 Enterprise and Standard Editions offer different capabilities and the choice of either edition depends solely on the needs of the business.

Enterprise Edition

Enterprise Edition is a comprehensive data platform for running mission-critical OLAP, data warehousing, and BI applications. This edition provides enterprise class scalability, high availability, and security to run secure, business-critical applications.

You can use the Enterprise Edition in the following scenarios:

- Mission-critical deployments requiring high availability and uptime
- Existing large-scale OLTP deployments
- OLTP deployments that expect to extend rapidly in the future

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

- Large-scale reporting and analysis deployments
- Data Warehousing
- Server consolidation

Standard Edition

SQL Server 2008 Standard Edition is a data platform for running departmental OLAP and BI applications. This edition is a complete data platform that provides ease of use and manageability for running applications.

You can use the Standard Edition in the following scenarios:

- Departmental deployments
- Small-to-medium scale OLTP deployments
- OLTP deployments that are not expected to rapidly grow in the future
- Reporting and analysis deployments

Comparing SQL Server 2008 Enterprise and Standard Editions

The following table compares the number of CPUs of SQL Server 2008 Enterprise Edition and Standard Edition.

CPU feature	Enterprise edition	Standard edition
Number of CPUs	OS Maximum	Four

The following table compares the scalability and performance feature of SQL Server 2008 Enterprise Edition and Standard Edition.

Scalability and performance features	Enterprise edition	Standard edition
x32 Hardware Support	Fully supported	Fully supported
x64 Hardware Support	Fully supported	Fully supported
I64 Hardware Support	Fully supported	Not supported
Multi-instance Support	50 instances	16 instances
Hypervisor Support	Fully supported	Not supported

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Scalability and performance features	Enterprise edition	Standard edition
Dynamic Address Windows Extensions Memory Management	Fully supported	Not supported
Table and Index Partitioning	Fully supported	Not supported
Parallel Index Operations	Fully supported	Not supported
Parallel Consistency Checks (DBCC)	Fully supported	Not supported
Database snapshots	Fully supported	Not supported
Scalable Shared Databases	Fully supported	Not supported
Indexed views	Fully supported	Supported except for the automatic substitution feature
Enhanced Read-ahead and Scan	Fully supported	Not supported

The following table compares the availability features of SQL Server 2008 Enterprise Edition and Standard Edition.

High-availability (Always On) features	Enterprise edition	Standard edition
Log Stream Compression	Fully supported	Fully supported
Clustering	16-node failover clustering	2-node failover clustering
Database Mirroring	Full	Single-thread
Online operations	Fully supported	Not supported
Resource Governor	Fully supported	Not supported
Backup compression	Fully supported	Not supported
Hot-add Memory and	Fully supported	Not supported

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

High-availability (Always On) features	Enterprise edition	Standard edition
CPU Support		
Database Snapshots	Fully supported	Not supported
Fast Recovery	Fully supported	Not supported
Online Page and File Restore	Fully supported	Not supported
Database Snapshots	Fully supported	Not supported
Fast Recovery	Fully supported	Not supported
Online Page and File Restore	Fully supported	Not supported
Parallel Index Operations	Fully supported	Not supported
Updated Distributed Partitioned Views	Fully supported	Not supported
Table and Index Partitioning	Fully supported	Not supported

The following table compares the Enterprise Security feature of SQL Server 2008 Enterprise Edition and Standard Edition.

Enterprise security features	Enterprise edition	Standard edition
Data encryption	TDE	Limited
Data encryption	Advanced third-party key management	Limited
Auditing	Security auditing	Limited

The following table compares the Enterprise Manageability features of SQL Server 2008 Enterprise Edition and Standard Edition

Enterprise manageability features	Enterprise edition	Standard edition

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Enterprise manageability features	Enterprise edition	Standard edition
Integration Services	Fully supported	Fully supported
SQL Server Import and Export Wizard	Fully supported	Fully supported
Log Providers and Logging	Fully supported	Fully supported
XML Source	Fully supported	Fully supported
SSIS Run-time	Fully supported	Fully supported
Basic Data Profiling Tools	Fully supported	Fully supported
SSIS Package Designer & Service	Fully supported	Fully supported
Data Mining Query Transformation	Fully supported	Not supported
Data Mining Model Training Destination Adapter	Fully supported	Not supported
Fuzzy Grouping Transformation	Fully supported	Not supported
Fuzzy Lookup Transformation	Fully supported	Not supported
Term Extraction Transformation	Fully supported	Not supported
Dimension Processing Destination Adapter	Fully supported	Not supported
Partition Processing Destination Adapter	Fully supported	Not supported

The following table compares the Data Warehousing feature of SQL Server 2008 Enterprise Edition and Standard Edition.

Data warehousing features	Enterprise edition	Standard edition
Partitioned table parallelism	Fully supported	Not supported
Data compression	Fully supported	Not supported
Star join query optimizations	Fully supported	Not supported
CDC	Fully supported	Not supported

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Question: If you need to setup a two-node failover cluster, which edition would you use if you have only 4GB RAM on each node?

Question: If you need to use online restore capability and CDC support, which edition would you use?

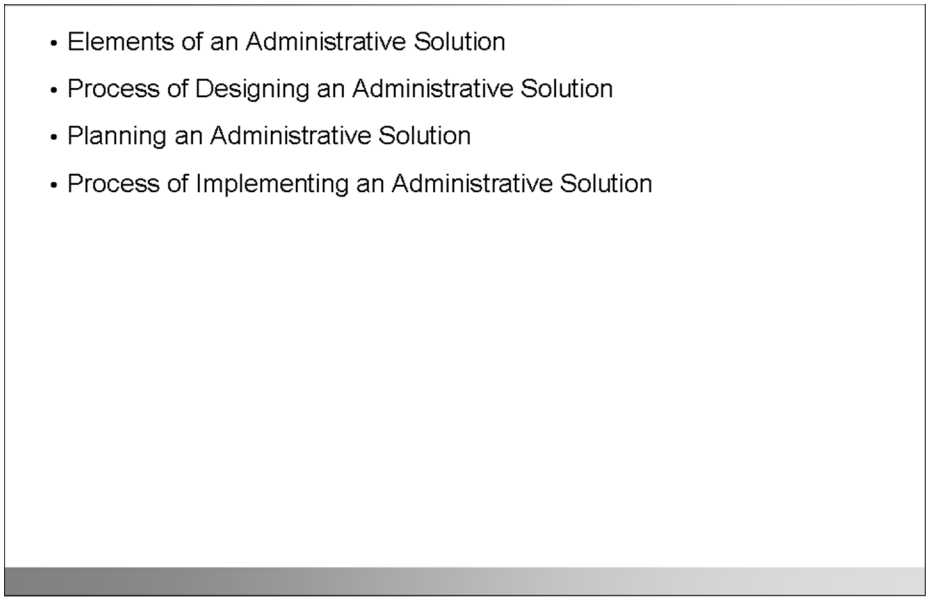
Source:

- http://download.microsoft.com/download/2/D/F/2DF66C0C-FFF2-4F2E-B739-BF4581CEE533/SQLServer2008EnterpriseandStandardFeature%2520Compare.pdf

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Lesson 2
Designing an Administrative Solution



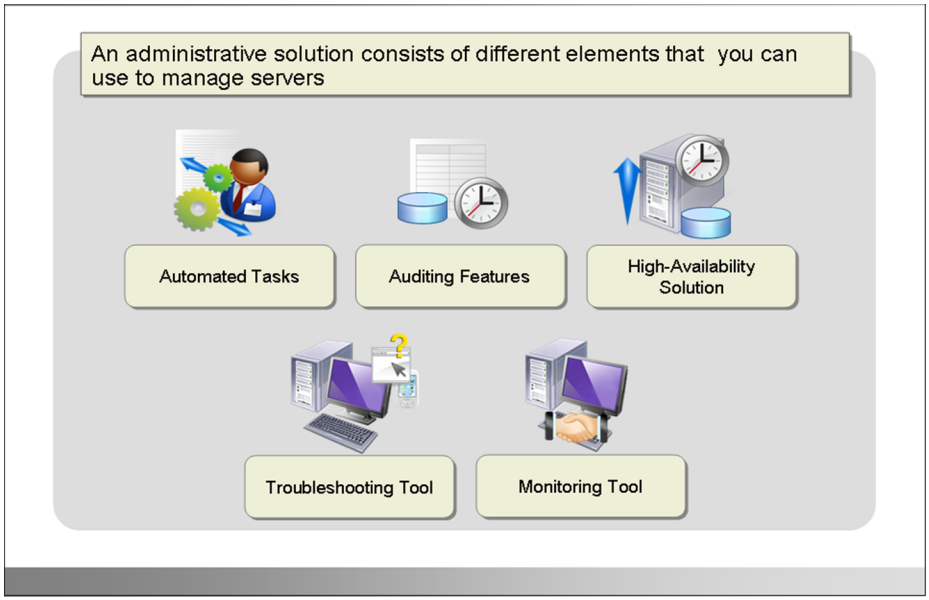
An administrative solution provides a set of features and functionalities in SQL Server that help you simplify the administrative tasks and reduce the workload.

In this lesson, you will learn about the elements of an administrative solution in SQL Server 2008 and the process of designing and implementing an administrative solution.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Elements of an Administrative Solution



Key Points

A SQL Server 2008 database administrative solution is a single package in which you can:

- Manage the administration and support of SQL Server 2008, easily and centrally.
- Reduce operation overhead.
- Reduce total cost of ownership in the long run.
- Reduce administrative overhead by using the high-performance Database Engine in Microsoft SQL Server that consists of several self-tuning configuration options, such as automatic statistics update.
- An administrative solution provides you with an easy operational environment.
- You must also configure an effective availability, backup, and restore strategy.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Elements of an Administrative Solution

A typical administrative solution includes the following elements:

- **Automated Tasks**. An administrative solution helps to automate administrative tasks and reduces the time needed to perform the task manually.
- **Auditing Features**. The auditing features helps to track changes made to a database and log who is making the changes. By tracking the history of changes, you can identify the source of a problem.
- **High-availability Solution**. Implementing the high availability feature ensures that critical business operations can be continued in the event of system failures or disasters that result in system outage.
- **Troubleshooting Tool**. An administrative solution helps to troubleshoot issues such as slow query performance.
- **Monitoring Tool**. Systems should be running in optimal status. An administrative solution should include ways to ensure that the system is running properly and can meet the required Service Level Agreement (SLA).

The reasons to implement an administrative solution include:

- **Reduced downtime**. By using an administrative solution, you can optimize the server for management or to provide service. This will reduce downtime and improve performance.
- **Improved Performance**. You can use the administrative solution tools to investigate and tune the performance of the database. This will improve the productivity of a business.
- **Prevention of unauthorized change and access**. SQL Server 2008 helps to control user access and implement change control policies by using access control and change management provisions. This prevents unauthorized access and traces the history of system deployment, thus reducing the time necessary to troubleshoot problems.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

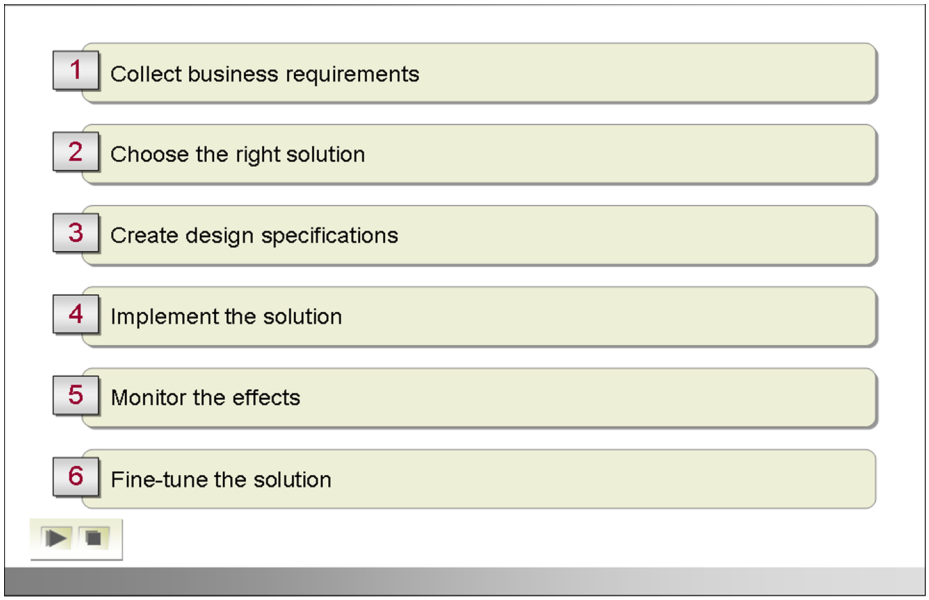
New Horizons Computer Learning Centers, Inc.

- **Automation of administrative tasks**. By automating administrative tasks to manage routine tasks and events, you can perform other administrative functions quickly and in a more controllable manner. This helps to reduce the overall cost of ownership, while ensuring that all required functions are maintained.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Process of Designing an Administrative Solution



Key Points

You need to use a formal process to deploy an administrative solution. After you deploy the solution, you need to monitor it continuously and tune it to suit the business operation environment.

The following are the steps to design an administrative solution:

1. **Collect business requirements**. As a first step for a successful deployment, you need to collect and understand the business requirements. In doing so, you can also identify the scope of the chosen solution.
2. **Choose the right solution**. After you establish the business requirements, you can choose a solution that will fulfill these requirements. To do this, you need to evaluate the requirements and choose a solution that will fulfill either all the requirements or most of them.
3. **Create design specifications**. A design specification is a document that defines the solution that you need to deploy. After

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

you have deployed the solution, you can use this document to validate if the solution deployed is completed as specified in the document.

4. **Implement the solution**. You can implement the solution as specified in the design specification and ensure that the implemented solution meets the business needs.
5. **Monitor the effects**. After you implement the solution, you need to monitor its effectiveness and ensure that the deployed solutions are functioning as expected and meeting the SLA terms.
6. **Fine-tune the solution**. To use the features of SQL Server 2008 effectively, you need to fine-tune the implementation after running the system or implementing the changes. This is a continuous activity, and fine-tuning can help the system to adapt to changing business needs.

Question: Why do you need to create a specification for your administrative solution?

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Discussion: Planning an Administrative Solution

- What are the tasks that need to be performed before creating the technical design of the project?
- What would you ask your business users about their expectations on the services to be provided by SQL Server after an upgrade?
- Why should you perform a check on the existing environment?
- Why should you create a functional specification?

In this discussion, you are presented with a scenario in which you are asked to plan an administrative solution. You and your classmates will discuss possible solutions to the scenario.

Scenario

You are the project manager on a SQL Server upgrade project. You are also responsible for creating the project plan for this project. You lead the technical team to support this project.

Question: What are the tasks that need to be performed before creating the technical design of the project?

Question: What would you ask your business users about their expectations on the services to be provided by SQL Server after an upgrade?

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Question: Why should you perform a check on the existing environment?

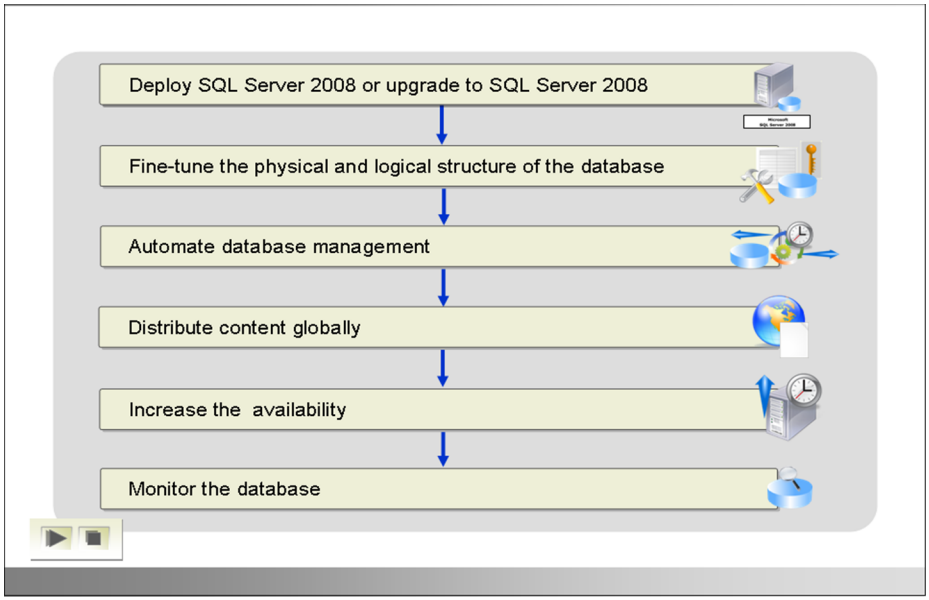
Question: Would you request for the application team to be involved at this stage of the project?

Question: Why should you create a functional specification?

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Process of Implementing an Administrative Solution



Key Points

To implement an administrative solution, you need to review the installation requirements, system configuration checks, and security considerations for a SQL Server 2008 installation.

Steps to implement an administrative solution:

1. You need to run the SQL Server Setup to install SQL Server 2008 or upgrade to SQL Server 2008 from SQL Server 2000 or SQL Server 2005. You need to upgrade all versions before SQL Server 2000 to a supported version, and then upgrade to SQL Server 2008.
2. After installing SQL Server 2008, you need to fine-tune the physical and logical structure where you need to configure the database, instances, and indexes.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

3. Next, you need to automate administrative tasks either by configuring SQL Server Agent or by using scripts. You can automate administrative tasks by using the sqlcmd tool, SSMS, PowerShell, or Windows Management Instrumentation (WMI).
4. You need to monitor your database regularly. You can configure event notification to monitor events in your database. You can use SQL Server Audit to perform an audit at the instance level. You can secure your databases by using database encryption and keys.
5. You need to design a high availability solution for SQL Server. You can use tools like log shipping to implement high availability.
6. Finally, you can distribute secured data to multiple servers by using SSIS. SQL Server supports distributed queries by using OLE DB. You can reference heterogeneous OLE DB data sources in T-SQL statements by using the linked server name or embedding the computer name into the query.

Question: Should you always follow the process in the order specified?

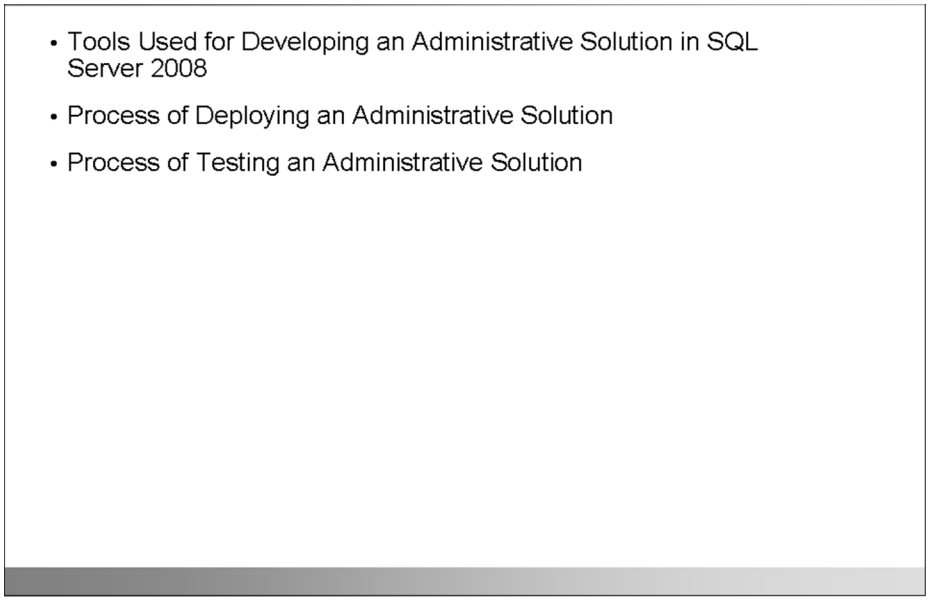
Source:

- http://msdn.microsoft.com/en-us/library/ms188313.aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Lesson 3
Developing and Deploying an Administrative Solution for SQL Server 2008



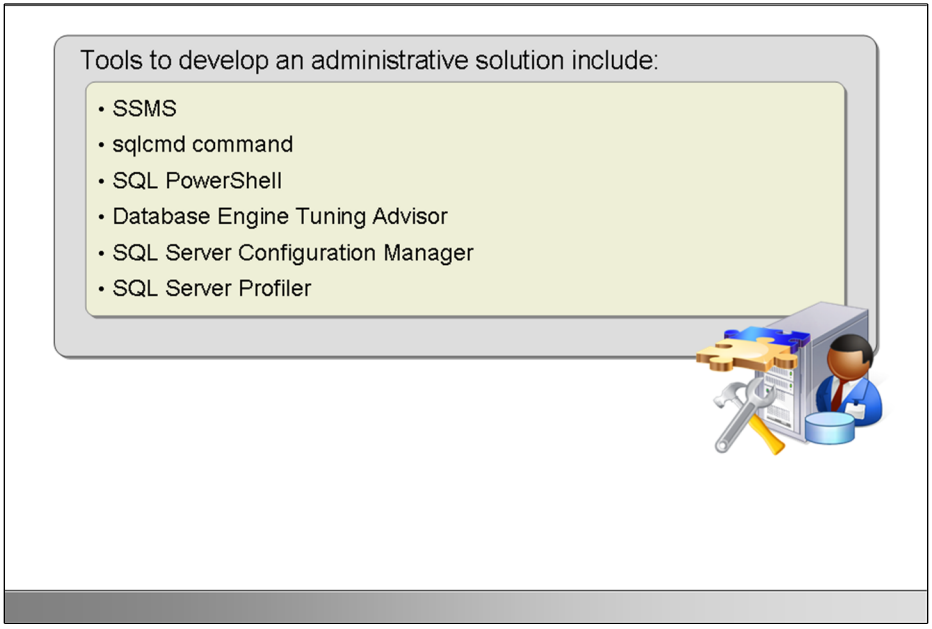
After you define an administrative solution, you need to deploy it. While deploying the solution, you need to ensure that you deploy it without causing problems to other applications or existing databases. Administrative solutions usually consist of functionalities like replication, which cover multiple servers. To ensure that there is no impact on other applications, you should have well-defined testing methodologies and a test lab, so that you can test every change before implementing the solution.

In this lesson, you will learn about the tools used for developing an administrative solution. You will also learn how to deploy and test an administrative solution.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Tools Used for Developing an Administrative Solution in SQL Server 2008



Key Points

The various tools used to develop an administrative solution in SQL Server include:

- **SSMS**. SSMS is an integrated environment that helps you access, configure, manage, administer, and develop all components of SQL Server. Compared to SSMS 2005, SSMS 2008 contains the following improvements:
  - Activity Monitor
  - Object Search
  - Object Explorer Details
  - Multi-Server Queries
  - T-SQL Debugger added to the Query Editor
  - Intellisense® added to the Query Editor

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

- **sqlcmd command-line tool**. The sqlcmd tool helps you enter T-SQL statements, system procedures, and script files at the command prompt. The results of the command are displayed in the Command Prompt window. You can run a sqlcmd job either by specifying a single T-SQL statement or by pointing the utility to a script file that contains T-SQL statements.
- **SQL PowerShell**. PowerShell is a powerful scripting shell that helps administrators and developers automate server administration and application deployment. The PowerShell language supports more complex logic than T-SQL scripts, giving SQL Server administrators the ability to build robust administration scripts. PowerShell scripts can also be used to administer other Microsoft server products. The SQL PowerShell provider exposes the hierarchy of SQL Server objects in paths, similar to file system paths, to locate a database object.
- **Database Engine Tuning Advisor (DTA)**. SQL Server DTA helps you select and create an optimal set of indexes, indexed views, and partitions without requiring an expert understanding of the structure of the database or the internals of SQL Server. DTA helps you analyze a workload and the physical implementation of one or more databases. DTA uses trace files, trace tables, or T-SQL scripts as workload inputs when tuning databases. You can create T-SQL script workloads with the Query Editor in SSMS and trace file and trace table workloads by using the Tuning Template in SQL Server Profiler.
- **SQL Server Configuration Manager**. SQL Server Configuration Manager is a tool to manage the services associated with SQL Server, to configure the network protocols used by SQL Server, and to manage the network connectivity configuration from SQL Server client computers. SQL Server Configuration Manager is a Microsoft Management Console (MMC) snap-in that is available from the Start menu, or can be added to any other MMC display. Microsoft Management Console (mmc.exe) uses the SQLServerManager10.msc file in the Windows System32 folder to open SQL Server Configuration Manager.
- **SQL Server Profiler**. Microsoft SQL Server Profiler is a graphical user interface to SQL Trace for monitoring an instance of the Database Engine or Analysis Services. You can capture and save data about each event to a file or table to analyze later.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Question: Which tool should you use to capture the workload of SQL Server?

Question: Which tools should you use to create scripts for managing different aspects of your SQL Server?

Source:

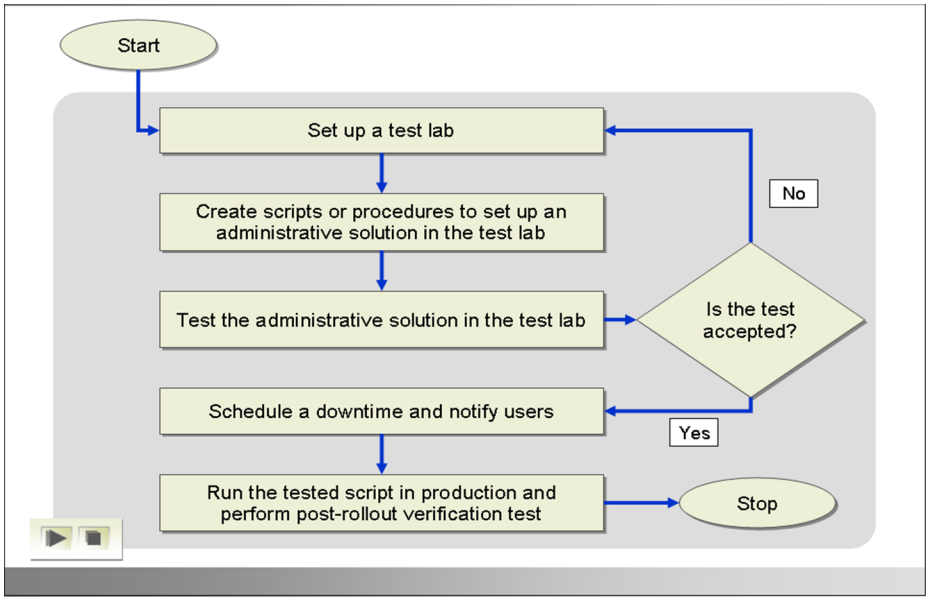
- http://msdn.microsoft.com/en-us/library/ms180944.aspx
- http://msdn.microsoft.com/en-us/library/cc281954.aspx
- http://msdn.microsoft.com/en-us/library/ms173494.aspx
- http://msdn.microsoft.com/en-us/library/ms181091.aspx

For more information, see Features and Tools Overview.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Process of Deploying an Administrative Solution



Key Points

You need to perform the following steps to deploy an administrative solution in a production environment:

1. **Set up a test lab**. You first need to set up a testing environment.
2. **Implement in a test environment**. After the testing environment is set, you need to implement the administrative solution in the test environment by using SQL scripts and commands. SQL scripts are the safest way to ensure that the solution can be consistently applied in both the testing and the production environment.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

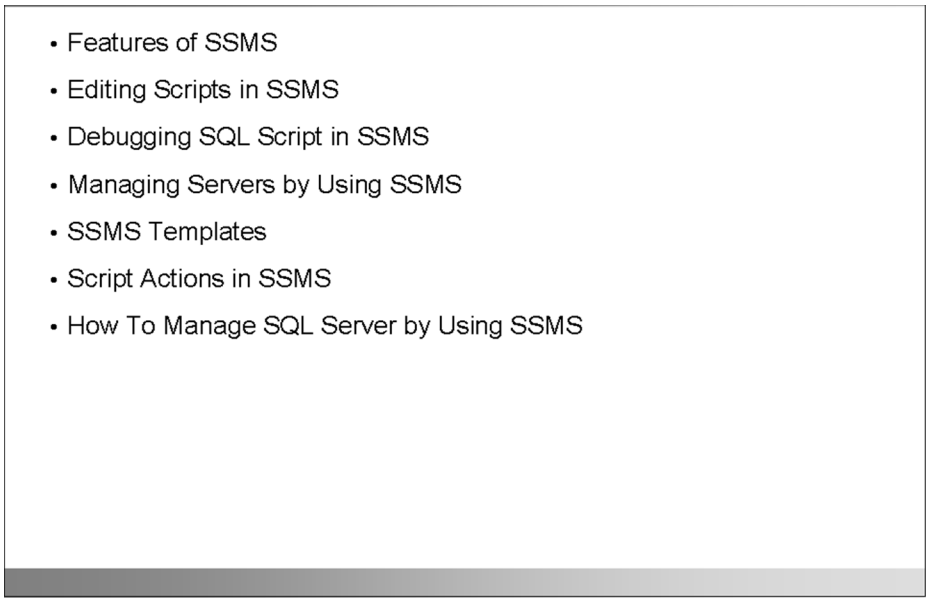
New Horizons Computer Learning Centers, Inc.

3. **Test the performance**. After setting up the solution in the test environment, you need to test its performance and ensure that the results are consistent with the results listed in the test case. If the solution is consistent with the expectation, the change can be deployed to the production environment; otherwise, you need to revaluate and retest it.
4. **Schedule a downtime.** After the solution is deployed and tested in the test environment, the solution is ready to be deployed to production. You need to arrange a downtime for the deployment and notify users of the event. When scheduling downtime, we should consider selecting the downtime which would impact the least number of users. If your application cannot affect an downtime, like Internet portal, consider upgrading server by server or switch over users over to passive site, if any, while upgrading the production site. No matter what, users should avoid making changes in database during the downtime as that would cause issues in catching up the data after deploying the changes and increase the complexity of the rollout.
5. **Rollout to production environment.** You can then roll out the solution to the production environment by using the SQL scripts that you used in the test environment. This ensures the steps you used to deploy the changes in the testing and the production environment are exactly the same.
6. **Conduct post-rollout verification tests.** After the roll out is complete, you need to conduct post-rollout verification tests, similar to those used in the test environment to ensure that the solution is working as per the specifications. The verification test should include the application team and the business users so they can choose some key scenarios to test in the post-rollout verification test to ensure that the applications are functioning properly after the deployment of the change.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Lesson 4
Implementing an Administrative Solution in SSMS



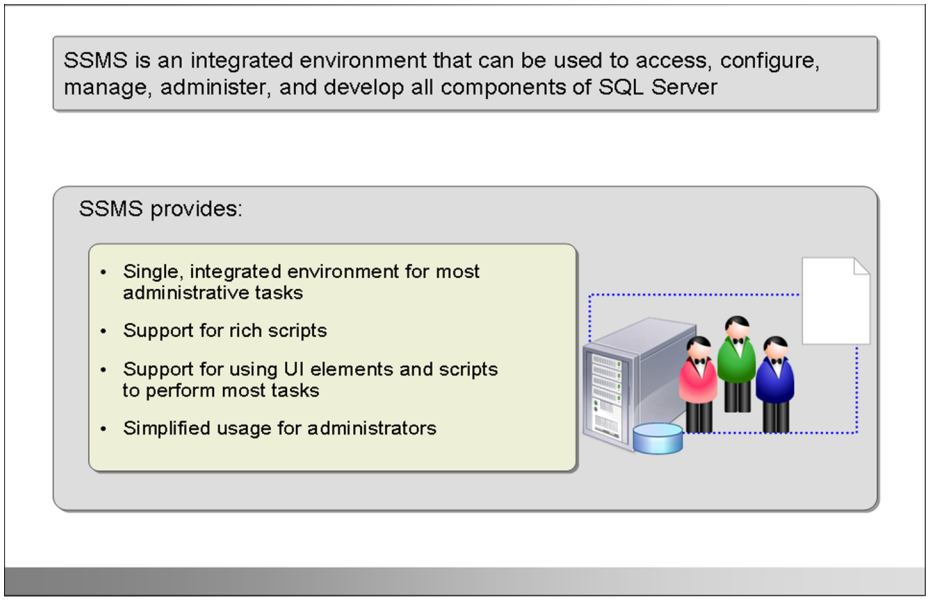
SSMS is an integrated tool for managing different components of SQL Server. You can use it to manage SSIS, database engine, SSRS, and SSAS. SSMS also provides comprehensive support for scripting. Therefore, you can generate scripts based on the action that you want to perform, without having the change really applied to the system.

In this lesson, you will learn about the features and templates of SSMS. You will also learn how to edit and debug scripts in SSMS, and how to manage SQL Server by using SSMS.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Features of SSMS



Key Points

You can access SQL Server by using the broad group of graphical tools combined with a number of rich script editors in SSMS.

SSMS combines the following features of SQL Server 2000 in a single, integrated environment:

- Enterprise Manager
- Query Analyzer
- Analysis Manager

You can view SSMS as a single console that works with all components of SQL Server, such as SSRS, SSIS, and SQL Server Compact 3.5 SP1.

Features of SSMS

The features of SSMS include:

- **Single, integrated environment for most administrative**

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

tasks. With a single graphical user interface, you can work with different aspects of SQL Server without writing a single line of script. By having a consistent interface and familiar layout for the tasks, you can perform all administrative functions efficiently. For example, you can perform tasks such as scheduling backup and managing SQL Server logon.

- **Support for rich scripts**. The Code Editor component of SSMS contains integrated script editors for authoring T-SQL, MDX, DMX, XML/A, and XML scripts. The Code Editor provides:
  - Dynamic Help for immediate access to relevant information while you are working.
  - A rich set of templates with the ability to create custom templates.
  - Support for writing and editing queries or scripts without requiring a connection to a server.
  - Support for scripting for SQLCMD queries and scripts.
  - A new interface for viewing XML results.
  - Integrated source control for solution and script projects, storing, and maintaining copies of scripts as they evolve over time.
  - Microsoft IntelliSense support for MDX statements
- **Easy usage for administrators**. SSMS supports registering multiple SQL servers into a single SSMS console, which helps you to centrally manage multiple SQL Servers. SSMS also supports SQL 2000, SQL 2005, and SQL 2008. This helps you to use a single console for managing different SQL Servers.

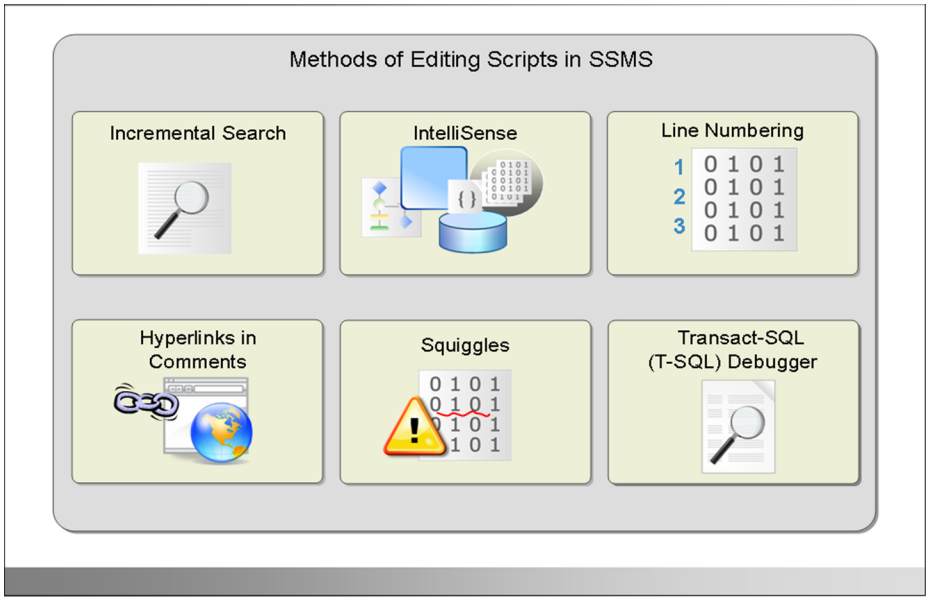
Source:

- http://msdn.microsoft.com/en-us/library/ms174173.aspx
- http://msdn.microsoft.com/en-us/library/ms174219.aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Editing Scripts in SSMS



Key Points

The various methods of editing scripts in SSMS include:

- **Incremental search**. In the incremental search method, you perform a search as you type. This search is a letter-by-letter basis. To activate an incremental search, press CTRL+I.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

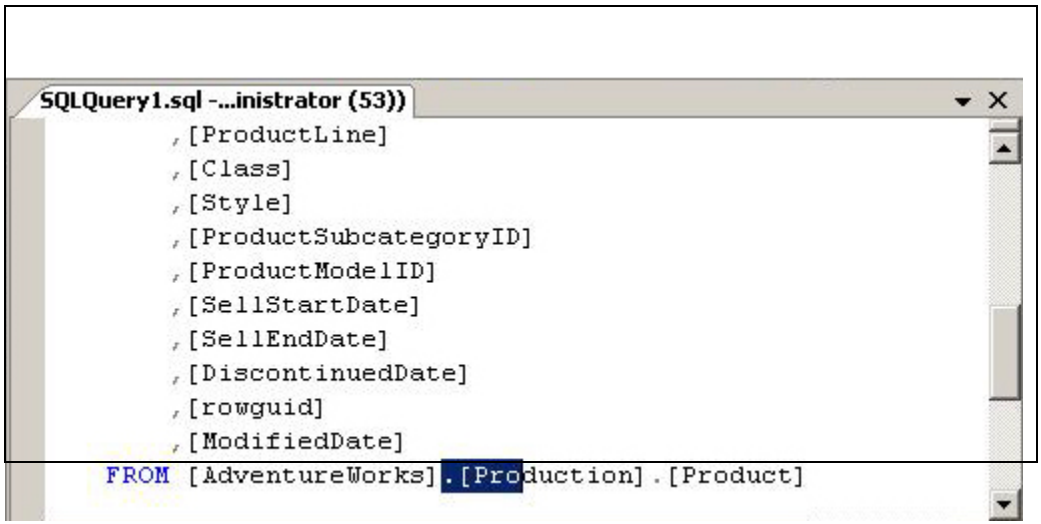


Figure 1.1. Incremental Search

- **IntelliSense**. By using Intellisense, you can enter a function or stored procedure into a query. This method will display the entered complete syntax of a query and arguments as a ToolTip.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

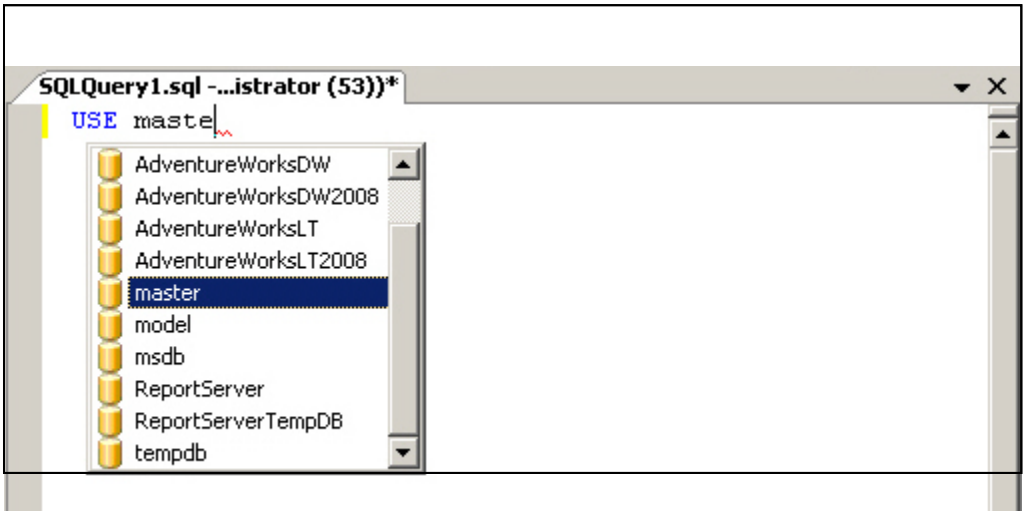


Figure 1.2. Intellisense

- **Line numbering**. By using the line numbering method, you can display line numbers in a code. You can do this on a per-language basis or for all languages. After you have activated line numbering, you can go to a particular line by clicking Go To on the Edit menu, or by pressing CTRL+G.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

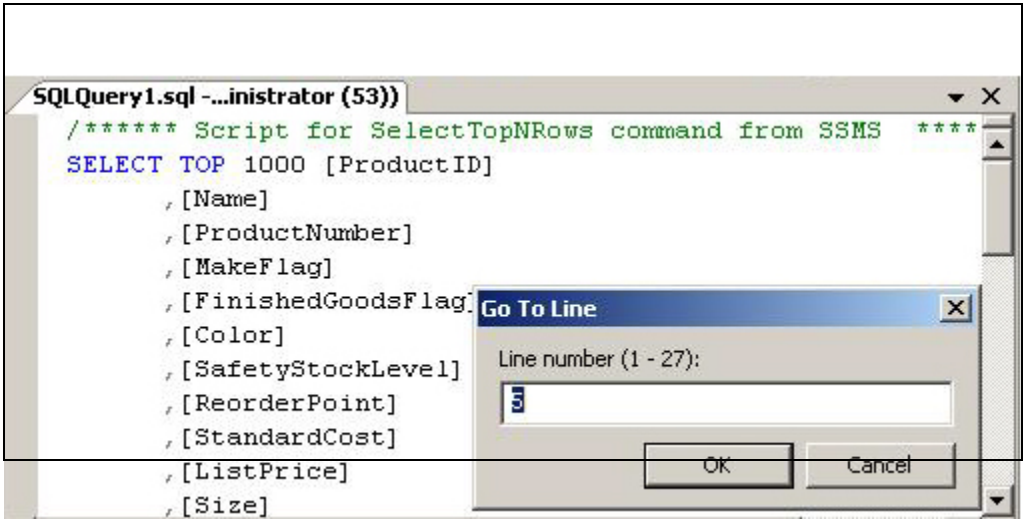


Figure 1.3. Line Numbering

- **Hyperlinks in comments**. You can embed URLs in your comment lines. For example, you can add a link to a Web page that contains a code specification or Help topic.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

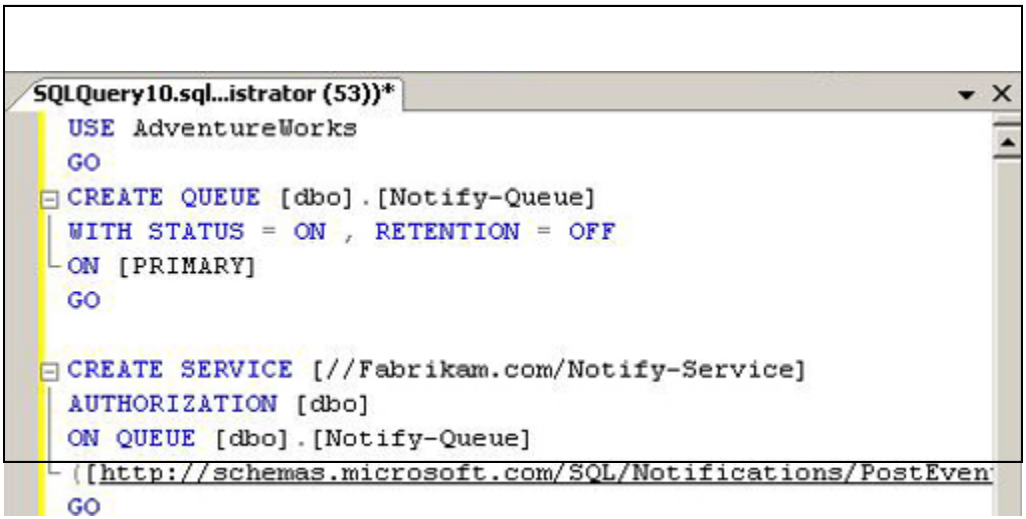


Figure 1.4. Hyperlinks in Comments

- **Squiggles**. Squiggles are wavy lines that are similar to the squiggles in Microsoft Office Word. You will be able to view the wavy lines beneath possible code errors. The wavy lines disappear when the code errors are corrected.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

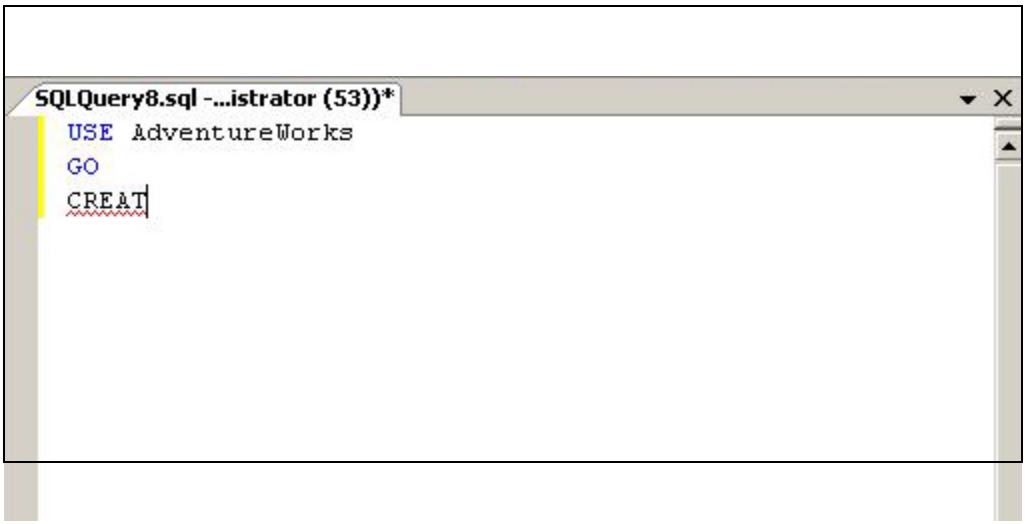


Figure 1.5. Squiggles

- **T-SQL debugger**. You can debug SQL script in SSMS by using the T-SQL debugger.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

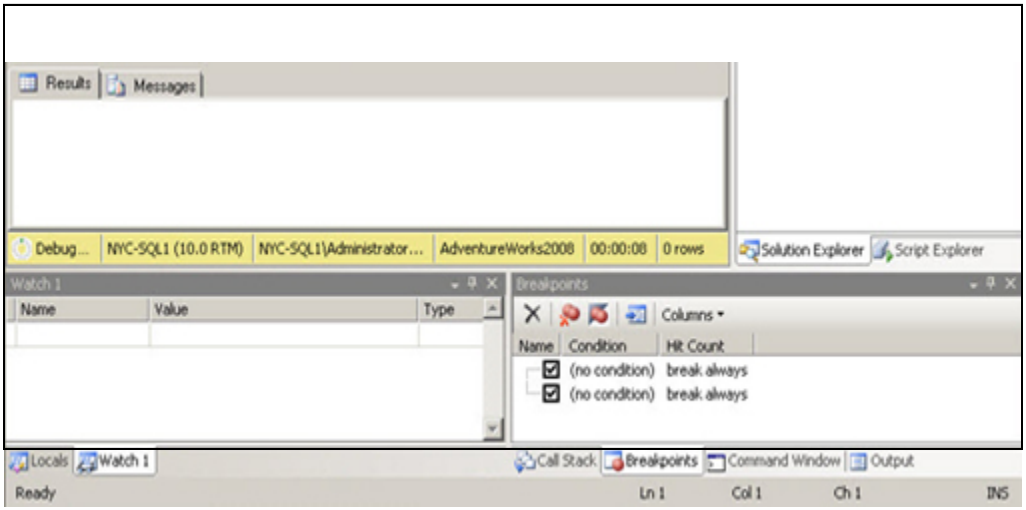


Figure 1.6. T-SQL Debugger

Note: IntelliSense is not available for encrypted database objects such as encrypted stored procedures or user-defined functions. Parameter help and Quick Info are not available for the parameters of extended stored procedures and Common Language Run-time (CLR) Integration user-defined types.

Statements Supported by Intellisense

In SQL Server 2008, IntelliSense supports only the most commonly used T-SQL statements. Some general Database Engine Query Editor conditions might prevent IntelliSense from functioning.

The following table describes how Database Engine Query Editor provides IntelliSense support for the syntax elements in the SELECT statement.

SELECT	WHERE
FROM	ORDER BY
HAVING	UNION
FOR	GROUP BY

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

SELECT	WHERE
TOP	OPTION (hint)

Question: Describe Intellisense support in SSMS?

Question: What are Squiggles in SSMS?

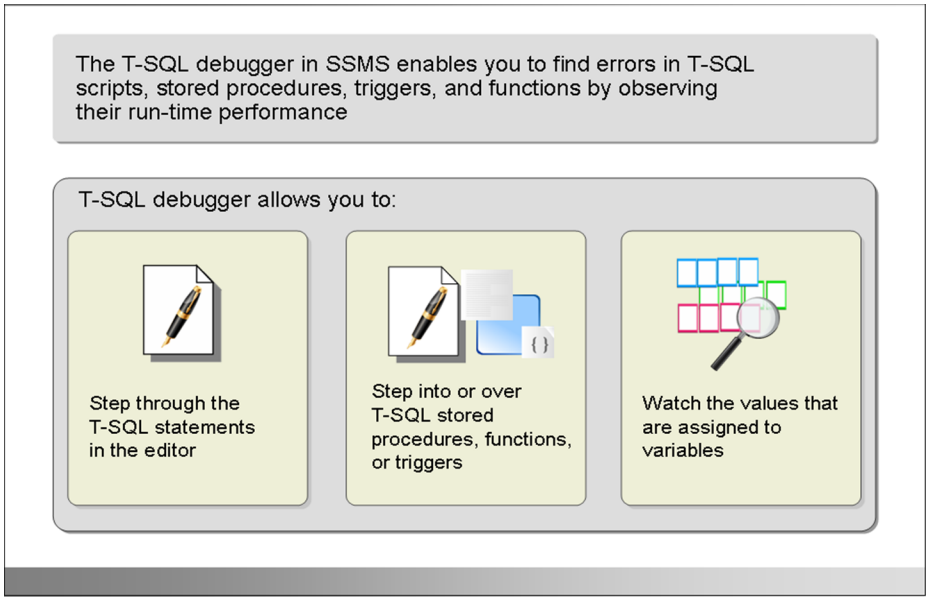
Source:

- http://msdn.microsoft.com/en-us/library/18ct837x%28VS.71%29.aspx
- http://msdn.microsoft.com/en-us/library/bb934481.aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Debugging SQL Script in SSMS



Key Points

You can use different options that are available in T-SQL to debug SQL scripts. These include:

- **Step through the T-SQL statements in the editor line-by-line**. This option enables you to navigate through a set of statements one by one, and to observe their performance while navigating.
- **Set breakpoints to stop at specific lines**. When you start the debugger, it pauses on the first line of code in the Query Editor window. To run to the first breakpoint that you have set, you can use the **Continue** feature. You can also use the **Continue** feature to run to the next breakpoint from any location at which the debugger is currently paused.
- **Step into or over T-SQL stored procedures, functions, or triggers that are run by the code in the editor window**. In the Database Engine Query Editor window, if you are sure there are no errors in a stored procedure, you can step over the

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

procedure. The procedure is run, and the results are returned to the code. If you want to debug a stored procedure or function, you can step into the module. SSMS opens a new Database Engine Query Editor window that displays the source code for the module, places the window into debug mode, and then pauses the execution at the first statement in the module. You can then navigate through the module code by setting breakpoints or by stepping through the code.

Note: Watch the values that are assigned to variables, and observe system objects such as call stack and threads.

Viewing Debugger Information

Each time the debugger pauses execution on a specific T-SQL statement, you can use the following debugger windows to view the current execution state:

- **Locals and Watch**. This window displays currently allocated T-SQL expressions, which are T-SQL clauses that evaluate to a single, scalar expression. This window also displays the data values that are currently assigned to the expressions.
- **QuickWatch**. This window displays the value of a T-SQL expression and helps you to save that expression to a Watch window.
- **Breakpoints**. This window displays the currently set breakpoints and helps you to manage them.
- **Call Stack**. This window displays the current execution location. It also provides information about how execution passed from the original Query Editor window through any functions, stored procedures, or triggers to reach the current execution location.
- **Output**. This window displays various messages and program data, such as system messages from the debugger.
- **Results and Messages**. These tabs in the Query Editor window display the result of previously run T-SQL statements.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

You can start the T-SQL debugger after you open a Database Engine Query Editor window. Then, you can run your T-SQL code in debug mode until you stop the debugger.

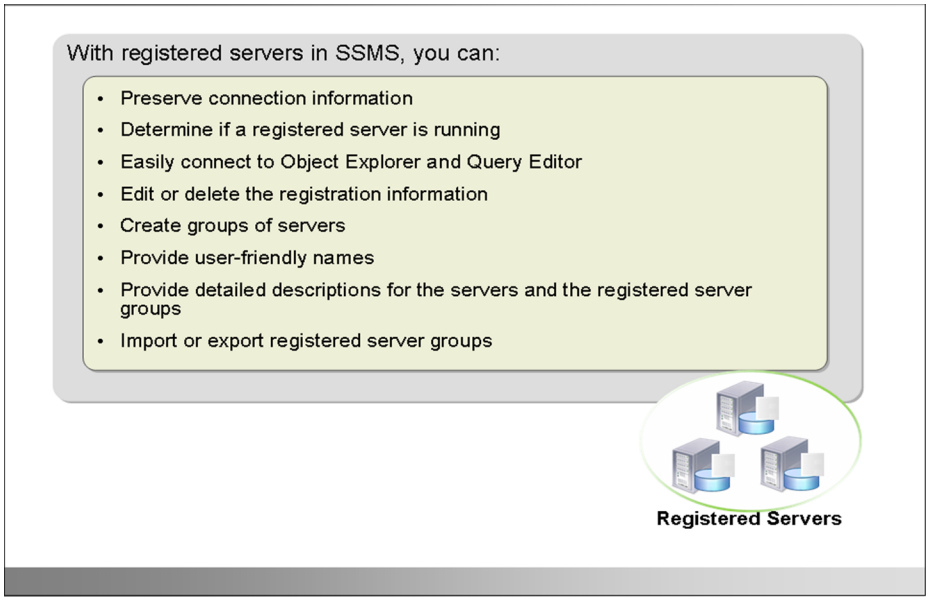
Source:

- http://msdn.microsoft.com/en-us/library/cc645997.aspx
- http://msdn.microsoft.com/en-us/library/cc646024.aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Managing Servers by Using SSMS



Key Points

SSMS helps you to register and manage SQL Servers. This is especially useful in environments where large numbers of SQL Servers are deployed. By registering SQL Server to SSMS, you can quickly review the status of the servers and access the servers, if necessary. You can use registered servers to organize servers that you access frequently. By using registered servers in SSMS, you can:

- Preserve the connection information.
- Determine if a registered server is running.
- Easily connect Object Explorer and Query Editor to a registered server.
- Edit or delete the registration information for a registered server.
- Create groups of servers.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

- Provide user-friendly names for registered servers by providing a value in the Registered server name box which is different from the server names stored in the Server name list.
- Provide detailed descriptions for registered servers.
- Provide detailed descriptions of registered server groups.
- Export registered server groups.
- Import registered server groups.

To view Registered Servers in SSMS:

1. In the Microsoft SQL Server Management Studio window, on the **View** menu, click **Registered Servers**. By default, the Registered Servers pane appears on top of the Object Explorer pane of the main window.

To register a connected server:

1. In the Microsoft SQL Server Management Studio window, in the Object Explorer pane, right-click the connected server, and then click **Register**. The **New Server Registration** dialog box appears.
2. In the **New Server Registration** dialog box, in the **Server name** box, type your server name.
3. In the Authentication box, type **Windows Authentication**.
4. In the Registered server name box, type the name of the second instance, and click **Save**.

To run statements against multiple configuration targets, simultaneously:

1. In the Microsoft SQL Server Management Studio window, on the **View** menu, click **Registered Servers**. The Registered Servers pane appears.
2. In the Registered Servers pane, expand **Central Management Servers**, right-click a server group, and then click **New Query**. The Query Editor pane appears.
3. In the Query Editor pane, type the following T-SQL code.

```
USE master

GO
```

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.


```
SELECT * FROM sysdatabases;
GO
```

4. Click **Execute**.

Note: By default, the results pane will combine the query results from all the servers in the server group.

To change the multiserver results options:

1. In the Microsoft SQL Server Management Studio window, on the **Tools** menu, click **Options**. The **Options** dialog box appears.
2. In the **Options** dialog box, in the navigation pane, expand **Query Results**, expand **SQL Server**, and then click **Multiserver Results**.
3. On the **Multiserver Results** page, specify the option settings that you want, and then click **OK**.

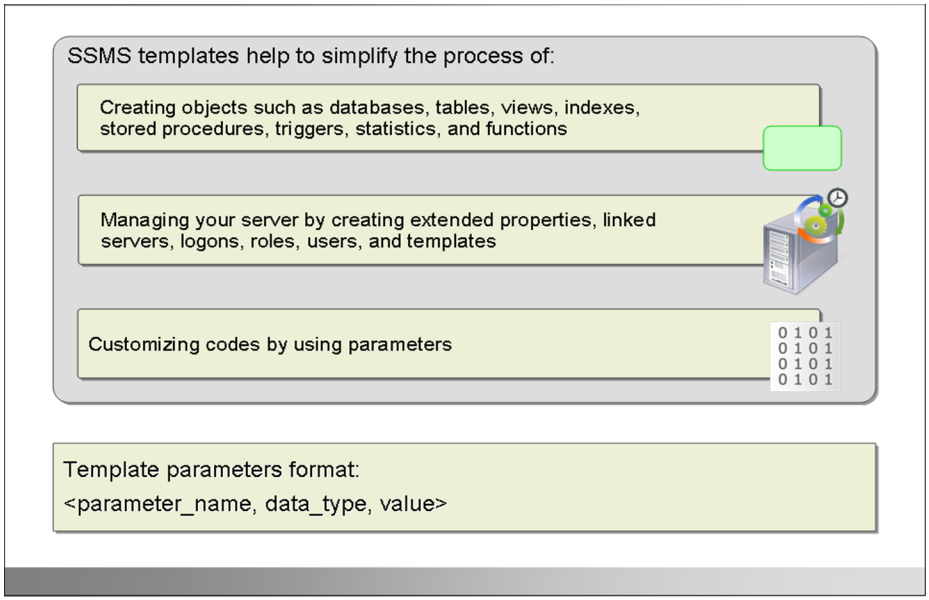
Source:

- http://msdn.microsoft.com/en-us/library/ms173574.aspx
- http://msdn.microsoft.com/en-us/library/bb964743.aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

SSMS Templates



Key Points

SSMS templates are boilerplate files that contain SQL scripts that help you create objects in the database. Templates are available for solutions, projects, and various types of code editors. By using templates, you can create objects such as databases, tables, views, indexes, stored procedures, triggers, statistics, and functions. In addition, there are templates that help you to manage your server by creating extended properties, linked servers, logins, roles, users, and templates for SSAS and SQL Server Compact 3.5 SP1.

SQL Server provides a variety of templates. The first time the template explorer is opened, a copy of the template is placed in the users Documents and Settings folder under Application Data\Microsoft\Microsoft SQL Server\100\Tools\Shell\Templates.

The template scripts provided with SSMS contain parameters that help you customize the code. Template parameter definitions use the following syntax.

```
<parameter_name, data_type, value>
```

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

In the above syntax:

- **Parameter_name**. The name of the parameter in the script.
- **Data_type**. The data type of the parameter.
- **Value**. The value that is to replace every occurrence of the parameter in the script.

Advantages of Using Templates

You can use templates to:

- Reduce time to roll out repetitive objects.
- Reduce human errors.
- Streamline efforts.

To open a template from Template Explorer:

1. In the Microsoft SQL Server Management Studio window, on the **View** menu, click **Template Explorer**. The Template Explorer pane appears.
2. In the Template Explorer pane, in the list of template categories, under SQL Server Templates, expand **Database**, and then double-click **Create Database** to open the template in the appropriate Query Editor pane. Alternatively, you can drag a template from the Template Explorer pane and drop it into a Query Editor pane, to add the template code.
3. In the **Connect to Database Engine** dialog box, complete the connection information, and then click **Connect** to open a new query editor window populated with the Create Database template.

To replace template parameters:

1. In the Microsoft SQL Server Management Studio window, on the Query menu, click **Specify Values for Template Parameters**. The **Specify Values for Template Parameters** dialog box appears.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

2. In the **Specify Values for Template Parameters** dialog box, the **Values** column contains suggested value for the parameter. You can accept the suggested value or replace the value with a new value, and then click **OK** to close the Specify Values for Template Parameters dialog box and modify the script in the Query Editor pane.

To create a custom template:

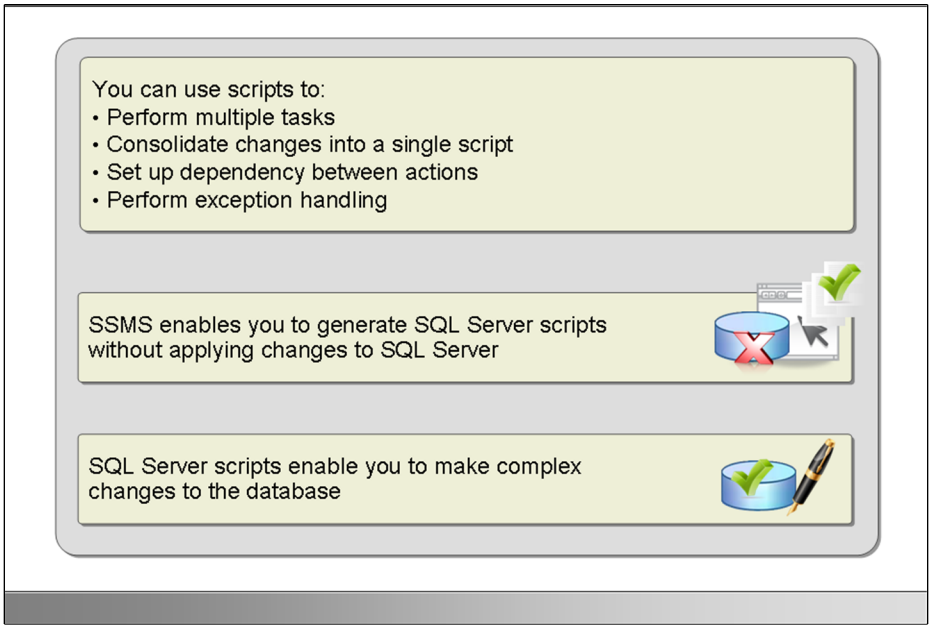
1. In the Microsoft SQL Server Management Studio window, in the Template Explorer pane, browse to the node where you want to store the new template.
2. Right-click the node, point to **New**, and then click **Template**.
3. Type the name for your new template, and then press **ENTER**.
4. Right-click the new template, and then click **Edit**.
5. Create a script in Query Editor pane. Insert parameters in your script in the format <parameter_name, data_type, value>. The data type and value areas must be present, but can be blank.
6. On the toolbar, click **Save** to save your new template.

Source:

- http://msdn.microsoft.com/en-us/library/ms174169.aspx
- http://msdn.microsoft.com/en-us/library/ms179334.aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

Script Actions in SSMS



Key Points

You can use a single script to do multiple tasks, such as creating a database and performing a SQL Server Agent job. When you use a script to apply database changes, the script also allows consistent action to be performed on multiple servers. When compared to a GUI, the number of errors in a script is fewer.

You can also use scripts to consolidate multiple changes into a single script and have all changes applied at once. You can also set up dependency between actions and perform exception handling to ensure the result of the scripts.

In SSMS, there is a script button at the top of the window. You can use this script button to perform an action, such as backup of a database, or make some changes in SQL Server. You can view an SQL script in the SSMS Query Editor, or by using any text editor.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

A schema generated as an SQL script can be used to do the following:

- Maintain a backup script that helps the user to re-create all users, groups, logins, and permissions.
- Create or update database development code.
- Create a test or development environment from an existing schema.

SQL scripts contain descriptions of statements that are used to create a database and its objects. You can generate scripts from the objects in an existing database and then add these objects to another database by running the scripts against that database. In effect, this re-creates the complete database structure and any individual database objects.

You can save the schema for generated objects in a single SQL Script file, or in several files with each file containing the schema of just one object. You can also save the schema generated for a single object, or a group of objects, into one or more SQL script files. You can create T-SQL scripts by using the Generate SQL Server Scripts Wizard, or by using Object Explorer.

Generate SQL Server Scripts Wizard

The Generate Scripts Wizard helps you create T-SQL scripts, and allows you to select a variety of objects to be scripted at once. The wizard has a variety of options for your scripts, such as whether to include permissions, collation, or constraints. You can use the wizard to create scripts for a large number of objects.

Object Explorer

You can use Object Explorer to quickly create scripts of an entire database or a single database object by using the default options. You can create a script in a Query Editor window, to a file, or to the clipboard. The script is created in Unicode format.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

You can create scripts to create or delete the objects. Some object types have additional scripting options, such as ALTER, SELECT, INSERT, UPDATE, DELETE, and EXECUTE.

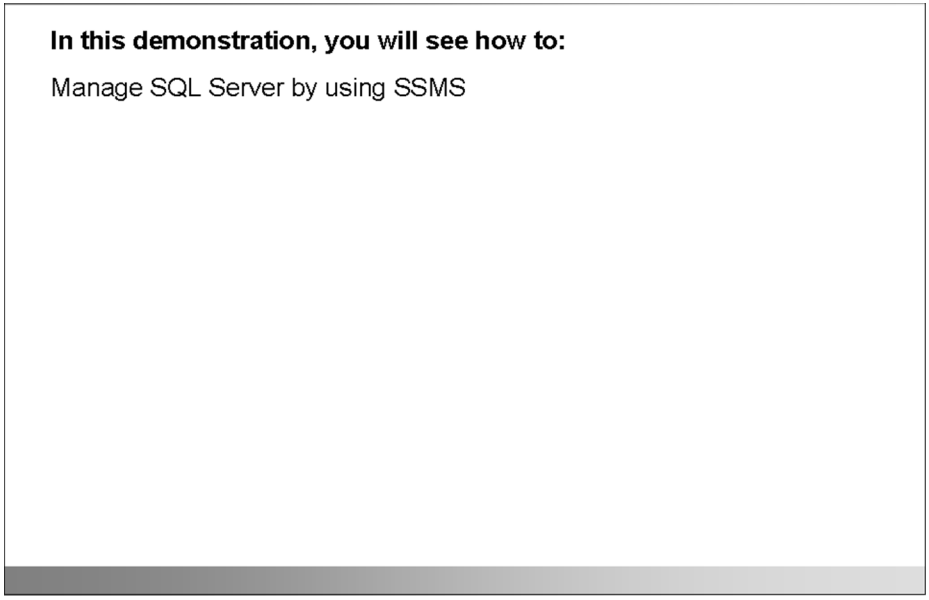
Source:

- http://msdn.microsoft.com/en-us/library/ms191299.aspx
- http://msdn.microsoft.com/en-us/library/ms178078.aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Demonstration: How To Manage SQL Server by Using SSMS



Key Points

The steps to manage SQL Server by using SSMS are:

1. In Windows Explorer, browse to the **D:\Demofiles\Mod01** folder, and then double-click the **M1_Demo1.sql** file.
2. In the **Connect to Database Engine** dialog box, click **Connect**.
3. In the **Microsoft SQL Server Management Studio** window, click **Execute**.
4. On the **View** menu of the Microsoft SQL Server Management Studio window, click **Object Explorer**.
5. In the Object Explorer pane, click **Connect Object Explorer**. The **Connect to Server** dialog box appears.
6. In the **Connect to Server** dialog box, click **Connect**.
7. In the Object Explorer pane, expand **NYC-SQL1**, expand

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Databases, expand **QuantamCorp**, expand **Tables**, right-click a table, and then point to **Script Table as**.

8. The shortcut menu has six scripting options: **CREATE To**, **DROP To**, **SELECT To**, **INSERT To**, **UPDATE To**, and **DELETE To**. Point to **UPDATE To**, and then click **New Query Editor Window**.
9. A new Query Editor window opens, makes a connection, and presents the entire update statement.
10. In the Object Explorer pane, under **Databases**, right-click the **QuantamCorp** database, and then click **Properties**. The **Database Properties – QuantamCorp** dialog box appears.
7. In the **Database Properties – QuantamCorp** dialog box, in the **Select a page** pane, click **Options**.
8. The current compatibility level is displayed in the **Compatibility level** list, and then click **OK**.

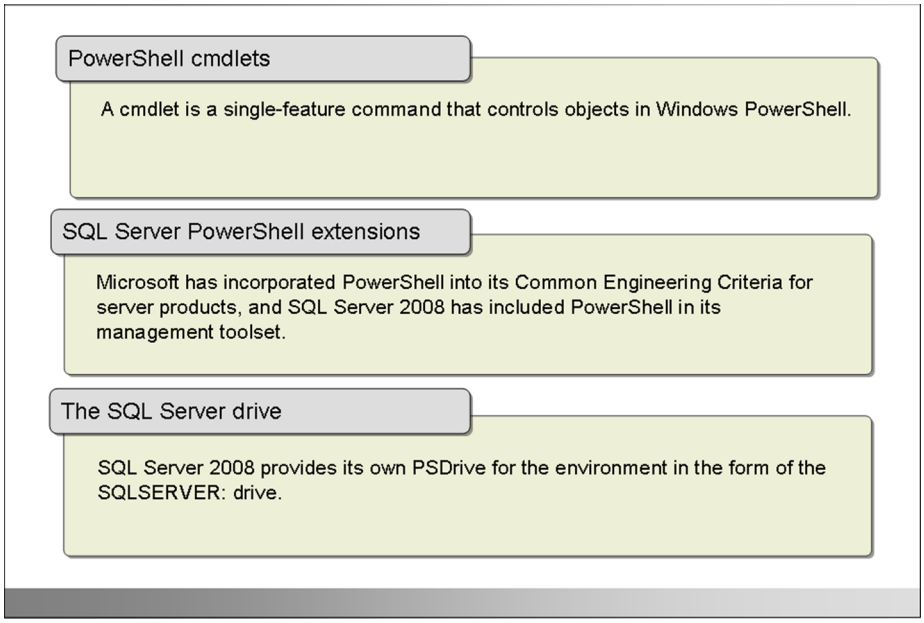
Note: You can change the compatibility level, by selecting the appropriate compatibility level, and then click **OK**.

Question: Is it possible to generate a script to perform an action and store the script for later use?

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Introduction to SQL PowerShell



Key Points

Windows PowerShell is a Windows command-line shell designed for system administrators. The Windows PowerShell includes an interactive prompt and a scripting environment that can be used independently or in combination.

Unlike other command-line shells, which accept and return text, Windows PowerShell is built on top of the .NET Framework common language runtime (CLR) and the .NET Framework to accept and return .NET Framework objects. Windows PowerShell provides the concept of a cmdlet, a single-function command-line tool built into the shell. You can use each cmdlet separately, but their power is realized when you use these simple tools in combination to perform complex tasks. Windows PowerShell contains more than one hundred built-in cmdlets, and you can write your own cmdlets and share them with other users.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Windows PowerShell provides you access to the file system on the computer. In addition, Windows PowerShell providers enable you to access other data stores, such as the registry and the digital signature certificate stores, as you access the file system. Command-line shells, such as cmd.exe, sh, ksh, csh, and bash, support the interaction—operate by executing a command or utility in a new process, and presenting the results to the user as a text message. Other text processing shells, such as sed, awk, and perl also support this interaction. Windows PowerShell uses the same command parser for all shell commands. This makes it much easier to learn how to use each command. You can still use the traditional Windows tools, such as Net, SC, and Reg.exe in Windows PowerShell.

PowerShell cmdlets

A cmdlet is a single-feature command that controls objects in Windows PowerShell. You can recognize cmdlets by their name format—a verb and noun separated by a dash (-), for example, Get-Help, Get-Process, and Start-Service.

In Windows PowerShell, cmdlets are simple, designed to be used in combination with other cmdlets. For example, the get cmdlets are used to retrieve data, the set cmdlets are used to establish or change data, the format cmdlets are used to format data, and the out cmdlets are used to direct the output to a specified destination. Each cmdlet has a help file, and you can access the help by typing the following command.

```
get-help <cmdlet-name> -detailed
```

The help file includes a description of the cmdlet, the command syntax, descriptions of the parameters, and an example to demonstrate the use of the cmdlet.

SQL Server PowerShell extensions

Microsoft has incorporated PowerShell into its Common Engineering Criteria for server products, and SQL Server 2008 has included PowerShell in its management toolset. Windows PowerShell by design allows the inclusion of snap-ins. These snap-ins provide access to the application that wants to allow a PowerShell interface into the application. For example, Microsoft Exchange 2007, has rebuilt the application from the ground up with administrative interfaces built for PowerShell, including hundreds of cmdlets for various functions, and

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

the Exchange administrative application executes the PowerShell cmdlets.

Shell scripting environments address everything on a server as though it were a file system. PowerShell Drives provide this capability in PowerShell, and the Get-PSDrive cmdlet will list the available drives in your session. PSDrives included basic PowerShell, such as environment variables (ENV), HKEY_CURRENT_USER Registry tree (HKCU), and LOCAL_MACHINE (HKLM). You can actually navigate the Windows Registry by using filesystem commands such as cd (Set-Location cmdlet) and dir (Get-ChildItem).

The SQL Server drive

SQL Server 2008 provides its own PSDrive for the environment in the form of the SQLSERVER: drive. With the SQL Server snap-in, you can navigate SQL Server objects through one of four directories in SQL Server 2008:

- The SQL folder accesses the database engine, SQL Server Agent, Service Broker, and Database Mail.
- The SQLPolicy folder accesses Policy-Based Management.
- The SQLRegistration folder accesses the Registered Servers and the new Central Management Server.
- The DataCollection folder accesses the Data Collector objects provided with Management Data Warehouse.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

The screenshot illustrates the four directories in SQL Server 2008.

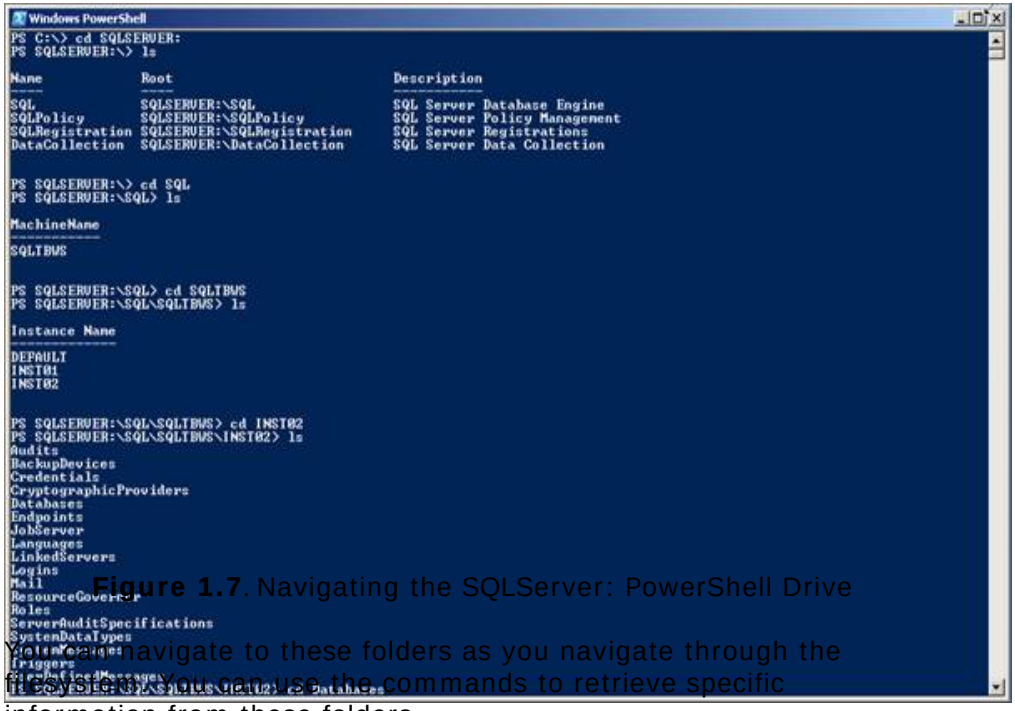


Figure 1.7. Navigating the SQLServer: PowerShell Drive

You can navigate to these folders as you navigate through the file system. You can use the commands to retrieve specific information from these folders.

Scenario 1

By navigating to the Databases folder under an instance of SQL Server, you can use the following command to send the list of databases to a Web page for anyone in the organization.

```
PS> dir | select name, CompatibilityLevel, RecoveryModel,
Size, SpaceAvailable | convertto-html >
c:\inetpub\wwwroot\databases.html
```

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Scenario 2

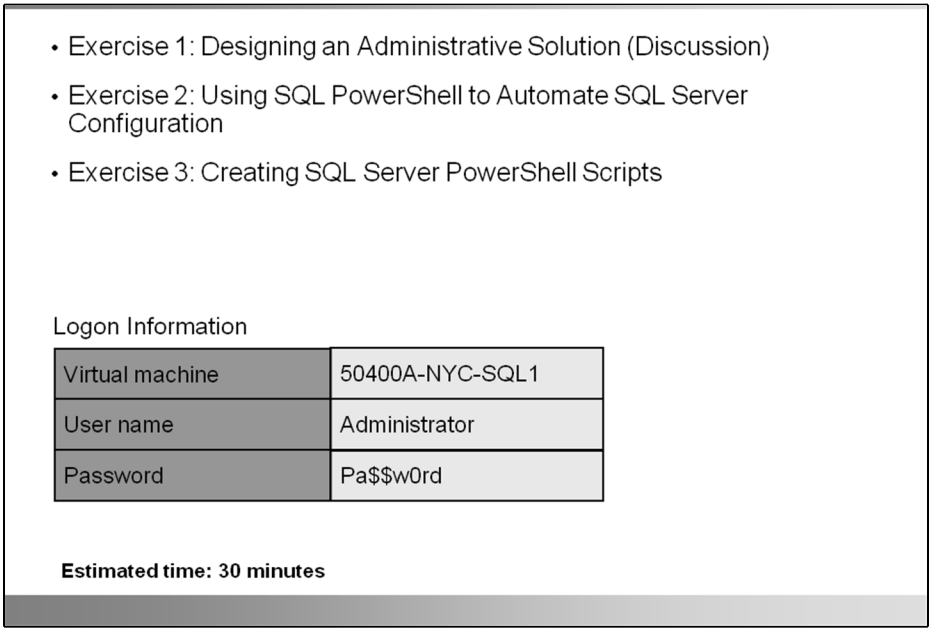
If you then navigate to the Tables folder under a specific database, you can use the following command to send the list of tables in that database, in descending order by the number of rows in the table, showing the largest tables with more rows first, to another Web page for the management tool set.

```
PS> dir | select schema, name, rowcount, filegroup | sort
rowcount -descending | convertto-html >
c:\inetpub\wwwroot\DBTables.html
```

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Lab 1: Designing an Administrative Solution for SQL Server 2008



Objectives

After completing this lab, you will be able to:

- Design an administrative solution.
- Use SQL PowerShell to automate SQL Server configuration.
- Create SQL PowerShell scripts.
- Configure FILESTREAM support in SQL Server 2008.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Scenario

QuantamCorp has deployed two SQL Server 2008 servers in their organization and will continue to deploy more servers that willreplace their existing installation of SQL Server 2005. After this, SQLServer 2008 will be the primary database engine used in theorganization. You need to familiarize yourself with the different administrative features in SQL Server 2008 and understand how they help simplify your administration jobs.

You plan to use a backup job as a pilot to test drive the new and enhanced support for administrative features in SQL Server 2008. In this test, you plan to perform the following:

- Design a project plan.
- Determine the advantages of the new administrative tool—SQL PowerShell.
- Create SQL Server PowerShell scripts.
- Use SQL Server Agent to create scheduled tasks.
- Familiarize yourself with enhanced SSMS.
- Determine how SQL Server 2008 is enhanced to provide better support for administrators.

Exercise 1: Designing an Administrative Solution (Discussion)

Scenario

You are managing over 100 SQL Server servers in your organization. You and your team members perform a number of tasks daily. These tasks take about two to three hours daily.

You have received an assignment to revamp the administrative tasks for managing these SQL Servers. All these SQL Servers will be upgraded to SQL Server 2008 soon. Hence, you have decided to use the new administrative features to perform the administrative tasks. You need to automate all the administrative tasks that you propose.

While carrying out your activities, you face the following problems:

- Colleagues are unable to implement configuration changes correctly and uniformly across multiple servers.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

- You have instructed all administrators to use scripts to deploy changes, but you want to reduce the need to execute the scripts on every server. You want to be able to execute the script once and then apply the changes to all servers.
- A number of administrators stay behind overnight every Friday, to perform recurring administrative jobs or occasional deployment tasks.

Recently, a colleague made a mistake while deploying a change to the production environment. That mistake caused all SQL Servers to fail. As a result, all Internet websites supported by those SQL Servers were unable to function.

Your new application, which is designed to work with SQL Server 2008, will release very soon. This application would use the FILESTREAM feature to store binary data. You need to configure SQL Server that hosts the database to support this feature. In addition, you also need to configure SQL Server to store all the files on the Network Attached Storage (NAS).

Based on the above scenario, answer the following questions:

1. Your management has asked you to prepare a project plan for implementing these changes. What tasks would you include in the project plan?
2. What feature should you use to ensure consistency and uniformity in implementing configuration changes?
3. How can you execute a script once, and then apply the changes to all servers? The suggested method must require minimal effort for managing all SQL Servers.
4. How can you configure your SQL Server environment to minimize the overtime that your colleagues need to work?
5. How can you reduce the possibility of human errors? The suggested method must support both configuration for SQL and Windows.

Exercise 2: Using SQL PowerShell to Automate SQL Server Configuration

Scenario

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

You have decided to use SQL PowerShell to perform administrative tasks as it supports both SQL Server and other configurations. You are assigned to create and configure a database for an application team. You have decided to use this project as a trial to decide if SQL PowerShell can meet your requirements.

The main tasks for this exercise are as follows:

1. Create a test database by using SQL PowerShell.
2. Configure the database compatibility level.
3. Verify the database configuration.
4. Drop the test database.

Task 1: Create a test database by using SQL PowerShell.

1. On the virtual machine, to start **SQL PowerShell**, at the command prompt, run the **sqlps** command.
2. At the database prompt, display all child items of the database.
3. At the database prompt, run SQL PowerShell code to create a test database with the name **PS_Test** by using the **Create()** function.

Task 2: Configure the database compatibility level.

1. At the database prompt, run **cd ps_test** to go to the ps_test prompt.
2. Using the ps_test prompt, set the compatibility level of the database to **80**.

Task 3: Verify the database configuration.

1. Using the ps_test prompt, run the **$MyDB.CompatibilityLevel** command to verify the database configuration.

Task 4: Drop the test database.

1. Using the ps_test prompt, run SQL PowerShell commands to delete the test database along with all its child items.

Results: After completing this exercise, you should have created a test database by using SQL PowerShell, configured the database compatibility level, verified the database configuration, and deleted the test database.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Exercise 3: Managing SQL Server Through Central Management Server

Scenario

You have around 50 SQL Server servers to manage and you constantly need to execute scripts on those servers at around the same time. You understand the Central Management server can be used to centrally manage all those servers. You are now running a test on your test environment to try out the feature.

The main tasks for this exercise are as follows:

1. Create a central management server.
2. Execute multiple server queries.
3. Change the result display format for multi-server query.

Task 1: Create a central management server.

1. In SQL Server Management Studio, on the **View** menu, click **Registered Servers**.
2. In Registered Servers, expand **Database Engine**, right-click **Central Management Servers**, point to **New**, and then click **Central Management Servers**.
3. In the **New Server Registration** dialog box, specify NYC-SQL1, and then click **OK**.
4. In Registered Servers, right-click the Central Management Server, point to **New**, and then click **New Server Group**. Type **Development Servers** as group name, and then click **OK**.
5. In Registered Servers, right-click the Central Management Server group, and then click **New Server Registration**.
6. In the **New Server Registration** dialog box, specify NYC-SQL/INSTANCE1 as instance name, and then click **OK**.
7. In the **New Server Registration** dialog box, specify NYC-SQL/INSTANCE2 as instance name, and then click **OK**.

Task 2: Execute multiple server queries.

1. In SQL Server Management Studio, on the **View** menu, click **Registered Servers**.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

2. Expand a Central Management Server, right-click a server group, point to **Connect**, and then click **New Query**.
3. In Query Editor, type and execute a Transact-SQL statement

```
USE master
GO
SELECT * FROM sysdatabases;
GO
```

Note: All results are combined into a single result.

Task 3: Change the result display format for multi-server query.

1. In Management Studio, on the **Tools** menu, click **Options**.
2. Expand **Query Results**, expand **SQL Server**, and then click **Multiserver Results**.
3. On the **Multiserver Results** page, clear the checkbox **Merge results** and check **Add server name to the results** checkbox, and then click **OK**.
4. In Query Editor, type and execute a Transact-SQL statement

```
USE master
GO
SELECT * FROM sysdatabases;
GO
```

Note: Verify the results that are displayed in the separate result set.

Results: After completing this exercise, you should have created a central management server and executed multiple server queries. In addition, you have changed the result display format for multi-server query.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Exercise 4: Creating SQL Server PowerShell Scripts

Scenario

After evaluating the capability of SQL PowerShell, you have decided to use it to perform most administrative tasks. You are now creating an SQL PowerShell script to perform database backups for your system databases. You will also use the script with SQL Server Agent, to automate the task so that it runs automatically every night.

The main tasks for this exercise are as follows:

1. Add a SQL Server PowerShell script to create a SQL Server Agent job.
2. Add the SQL Server PowerShell script to create a SQL Server Agent job step.
3. Save the SQL PowerShell script.
4. Start SQL PowerShell and run the script.
5. Verify the SQL Server Agent configuration.

Task 1: Add a SQL Server PowerShell script to create a SQL Server Agent job.

1. Open the Notepad window.
2. In the Notepad window, add the SQL Server PowerShell script to create a SQL Server Agent job with the **Backup System Database** job name.

Task 2: Add the SQL Server PowerShell script to create a SQL Server Agent job step.

1. In the Notepad window, add the SQL Server PowerShell code to create two SQL Server Agent job steps named **backup master** and **backup model** by using the **Create()** function as follows, at the end of the existing code.

Job step name	Settings
backup master	Command="BACKUP DATABASE master TO DISK='D:\Labfiles\Mod01\master.bak'" SubSystem = TransactSQL

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Job step name	Settings
	OnSuccessAction = "GoToNextStep"
backup_model	Command="BACKUP DATABASE model TO DISK='D:\Labfiles\Mod01\model.bak'" SubSystem = TransactSQL OnSuccessAction = "GoToNextStep"

Task 3: Save the SQL PowerShell script.

1. Save the Notepad file as **D:\Labfiles\Mod01\create_job.ps1**.

Task 4: Start SQL PowerShell and run the script.

1. Start the command prompt window.
2. At the command prompt, change to the SQL Server PowerShell prompt, and run **D:\Labfiles\Mod01\create_job.ps1**.

Task 5: Verify the SQL Server Agent configuration.

1. Open the **Microsoft SQL Server Management Studio** window.
2. To connect to Microsoft SQL Server Management Studio, in the Connect to Server dialog box, type the following:

   Server type: **Database Engine**

   Server name: **NYC-SQL1**

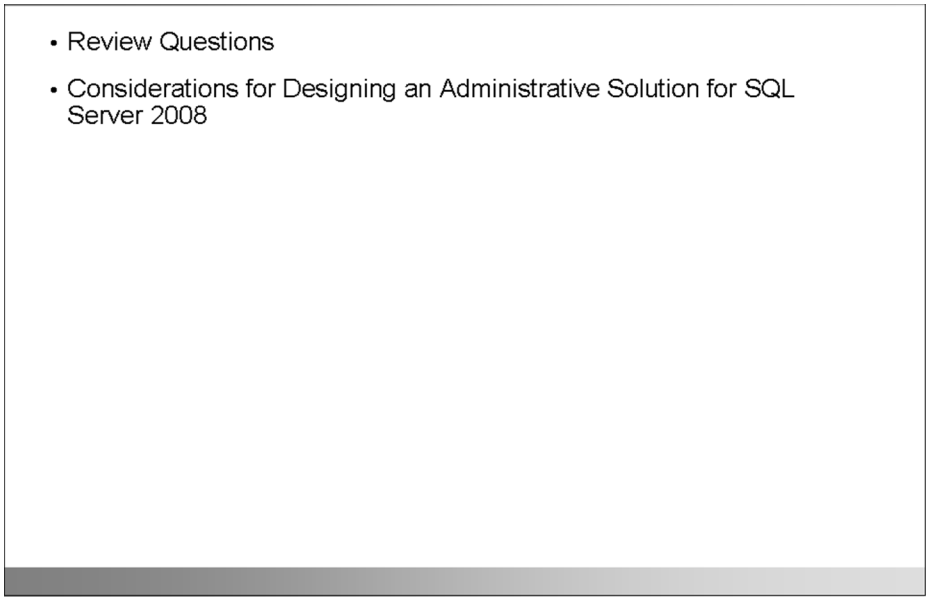
   Authentication: **Windows Authentication**
3. Using the Job Properties – Backup System Database job window, verify the **backup master** and **backup model** steps under **Job step list**.
4. Close the Microsoft SQL Server Management Studio window.

Results: After completing this exercise, you should have created a SQL Server Agent job and a SQL Server Agent job step.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Module Review and Takeaways



Review Questions

1. You are assigned the job of revamping the administrative work for SQL Servers in your organization. What are the benefits of introducing an administrative solution?
2. Your application team informed you that their application requires SQL Server 2008 to work as it requires features only supported in SQL Server 2008, such as the date datatype. You are considering which edition to deploy in your environment. Which edition would you deploy in your environment?
3. You are wrapping up development work on an administrative solution. You want to verify if the deployment scripts developed works in your environment. How will you verify it?
4. While deploying an administrative solution, you suggest to your team to use scripts instead of using the GUI. Why should you always consider using scripts for deploying changes?

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

5. Your application team wants to know which tool is better to create scripts, SQL PowerShell or sqlcmd. You suggest that they use SQL PowerShell instead. What are the benefits of using SQL PowerShell over sqlcmd?
6. Your team wants to know if there is a better way to create scripts than creating everything from the beginning. They feel that writing scripts from the beginning is very time consuming. You ask them to use SSMS to help. Which feature in SSMS helps to generate scripts for deploying changes without really making a change to the production environment?

Considerations for Designing an Administrative Solution for SQL Server 2008

When designing an administrative solution, consider the following:

- Use SQL Server 2008 for improved performance, scalability, security, manageability, and programmability for richer functionalities to enterprises.
- Use the right edition for your solution to minimize its cost. SQL Use Enterprise Edition for all the features and use Standard Edition for smaller deployments.
- Use SQL Server Express Edition to store limited amount of data in a database.
- Deploy an administrative solution that combines different features of SQL Server.
- Use tools like Powershell, sqlcmd, and SSMS to manage servers. Use PowerShell and sqlcmd for command-line support and SSMS for GUI administration.
- Ensure that a well-defined process for testing a solution is established.
- Set up an isolated test lab to test the solution before deploying it in the production environment.
- Use scripts to deploy changes to your environment because scripts help to reduce human errors and provide the ability to schedule tasks.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Module 2

Deploying SQL Server 2008

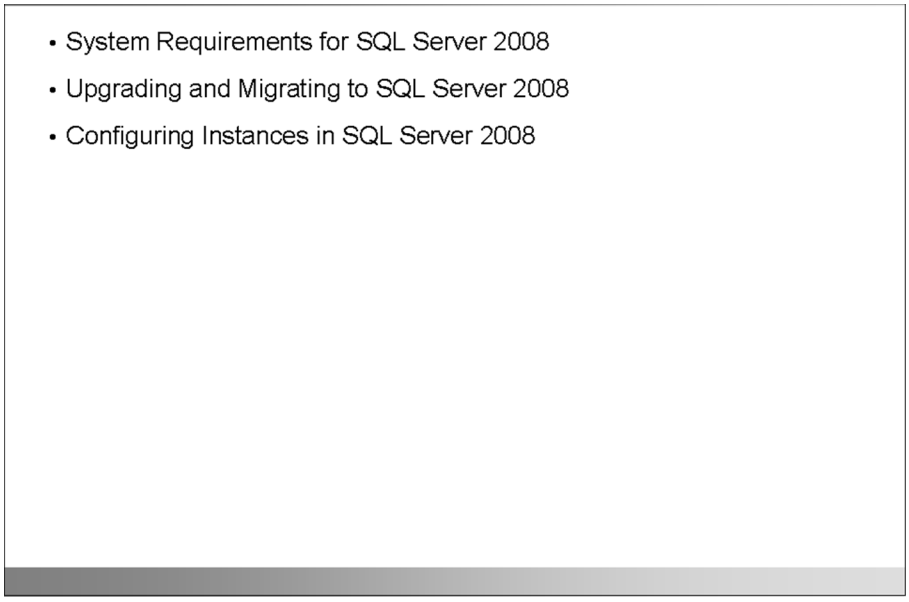
Contents:

Lesson 1: System Requirements for SQL Server 2008 2-3
Lesson 2: Upgrading and Migrating to SQL Server 2008 2-18
Lesson 3: Configuring Instances in SQL Server 2008 2-40
Lab 2: Deploying SQL Server 2008 2-58
Module Review and Takeaways 2-61

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Module Overview



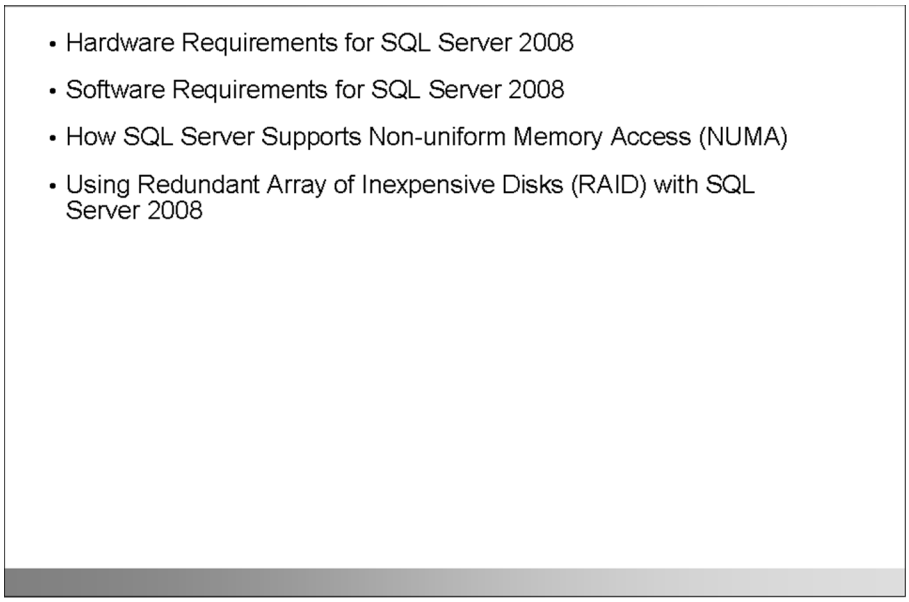
Microsoft™ SQL Server™ 2008 is a comprehensive platform that helps you manage your database. It provides security, reliability, performance, and scalability.

In this module, you will learn about the hardware required for SQL Server 2008 and the different methods of deploying SQL Server 2008. You will also learn about the considerations to upgrade, or migrate to SQL Server 2008 from existing platforms. You will learn how to optimize a SQL Server 2008 installation after deployment.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Lesson 1
System Requirements for SQL Server 2008



You can set up SQL Server 2008 on systems with different hardware configuration. The software requirements for SQL Server 2008 may vary depending on the component that you need to address your business requirement.

In this lesson, you will learn about the various system requirements for SQL Server 2008.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Hardware Requirements for SQL Server 2008

SQL Server Edition	Processor Type	Minimum Processor	Processor Speed	RAM
SQL Server 2008 Standard	x86 (32bit)	Pentium III-compatible processor or faster	• Minimum: 1.0 GHz • Recommended: 2.0 GHz or faster	• Minimum: 512 MB • Recommended: 2 GB or more • Maximum: Operating system maximum
	x64 (64bit)	AMD Opteron, AMD Athlon 64, Intel Xeon with Intel EM64T support, Intel Pentium IV with EM64T support	• Minimum: 1.4 GHz • Recommended: 2.0 GHz or faster	• Minimum: 512 MB • Recommended: 2 GB or more • Maximum: Operating system maximum
SQL Server 2008 Enterprise	x86 (32bit)	Pentium III-compatible processor or faster	• Minimum: 1.0 GHz • Recommended: 2.0 GHz or faster	• Minimum: 512 MB • Recommended: 2 GB or more • Maximum: Operating system maximum
	x64 (64bit)	AMD Opteron, AMD Athlon 64, Intel Xeon with Intel EM64T support, Intel Pentium IV with EM64T support	• Minimum: 1.4 GHz • Recommended: 2.0 GHz or faster	• Minimum: 512 MB • Recommended: 2 GB or more • Maximum: Operating system maximum

Key Points

The hard disk requirements for SQL Server 2008 could vary depending on the data volume that the database is expected to manage.

You can install SQL Server 2008 on Windows Server® 2003 SP2 and Windows Server 2008. For development purposes, you can also install the Developer edition on Windows XP Professional and Windows Vista® Ultimate.

Windows Server 2008 Server Core installations do not support the Enterprise edition and the Professional edition of SQL Server 2008. Moreover, you can install only the Express edition on the Web edition of Windows Server 2008.

64-bit environment for SQL Server

SQL Server 2008 supports both the 32-bit and 64-bit versions for a variety of business applications with innovative offerings from AMD (AMD64) and Intel (Extended Memory 64 Technology or EM64T).

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

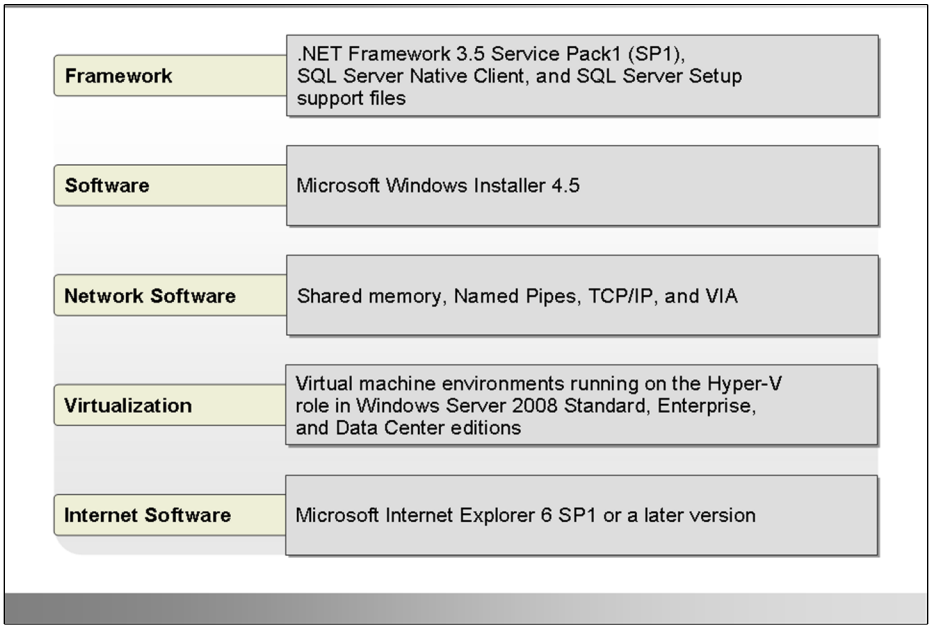
SQL Server 2008 supports the new 64-bit architecture based on the x86 instruction set with 64-bit extensions.

In the 64 bit environment, SQL Server 2008 (64-bit), enterprise applications, and databases can accumulate information gathered by business systems, support a large number of users, and efficiently process and analyze huge amounts of data in complex ways.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Software Requirements for SQL Server 2008



Key Points

The following table lists the software requirements for SQL Server 2008.

Component	Requirement
Framework	SQL Server Setup installs the following software components: .NET Framework 3.5 SP1. You require the following .NET Framework versions: SQL Server 2008 on Windows Server 2003 (64-bit) IA64. .NET Framework 2.0 SP2 SQL Server Express. .NET Framework 2.0 SP2 All other editions of SQL Server 2008. .NET Framework 3.5 SP1 SQL Server Native Client SQL Server Setup support files

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Component	Requirement
	Note: NET Framework 2.0 SP2 is not available as a separate download. You will need to install .NET Framework 3.5 SP1 which includes .NET Framework 2.0 SP2.
Software	SQL Server Setup requires Microsoft Windows Installer 4.5. After installing the required software components, SQL Server Setup will verify that the computer where SQL Server 2008 will be installed meets all the other requirements for a successful installation.
Network Software	Network software requirements for the 32-bit versions and 64-bit versions of SQL Server 2008 are the same. Supported operating systems have built-in network software. Stand-alone named and default instances support the following network protocols: Shared memory Named Pipes TCP/IP VIA
Virtualization	SQL Server 2008 is supported in virtual machine environments running on the Hyper-V™ role in Windows Server 2008 Standard, Enterprise, and Data Center editions. Within the Hyper-V role on Windows Server 2008, you can allocate a maximum of four virtual processors to virtual machines running Windows Server 2008 32-bit or 64-bit editions. You can also allocate a maximum of two virtual processors to virtual computers that are running Windows Server 2003 32-bit editions. For virtual computers that host other operating systems, you can allocate a maximum of one virtual processor to virtual computers. Note: You must shut down SQL Server 2008 before shutting down the virtual machine.
Internet Software	You require Microsoft Internet Explorer® 6 SP1 version for all installations of SQL Server 2008. You

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Component	Requirement
	require Internet Explorer 6 SP1 version for Microsoft Management Console (MMC), SQL Server Management Studio (SSMS), Business Intelligence Development Studio (BIDS), the Report Designer component of SQL Server Reporting Services (SSRS), and HTML Help.
Operating system	SQL Server Setup requires any one of the following operating systems: • Windows Server 2003 SP2 64-bit x64 Standard • Windows Server 2003 SP2 64-bit x64 Datacenter • Windows Server 2003 SP2 64-bit x64 Enterprise • Windows Server 2008 64-bit x64 Standard • Windows Server 2008 64-bit x64 Standard without Hyper-V • Windows Server 2008 64-bit x64 Datacenter • Windows Server 2008 64-bit x64 Datacenter without Hyper-V • Windows Server 2008 64-bit x64 Enterprise, • Windows Server 2008 64-bit x64 Enterprise without Hyper-V • Windows Server 2008 R2 64-bit x64 Web • Windows Server 2008 R2 64-bit x64 Standard • Windows Server 2008 R2 64-bit x64 Enterprise • Windows Server 2008 R2 64-bit x64 Datacenter

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

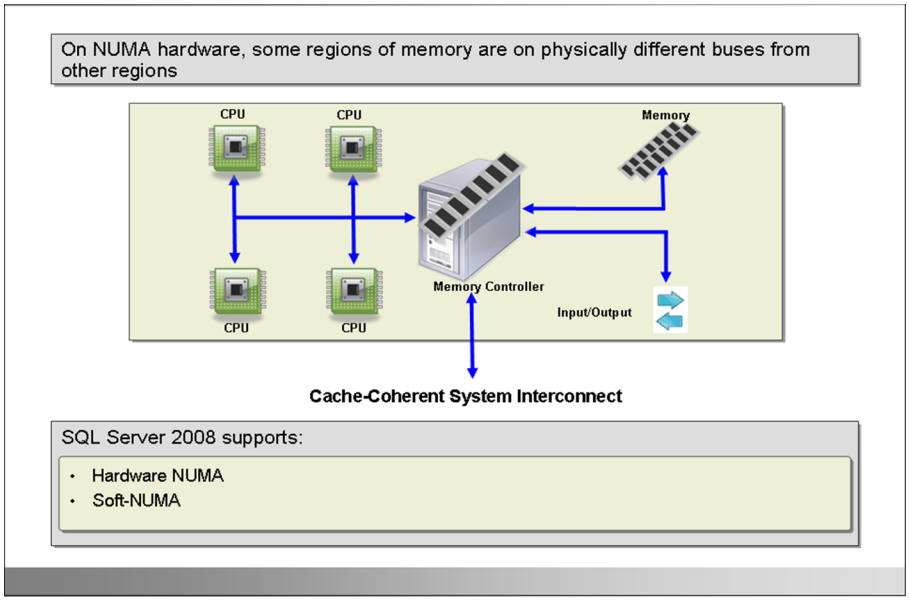
SQL Server Setup will not install some components for different editions. The following table lists the components you must install manually before you run SQL Server Setup:

Edition	Components
SQL Server Express	.NET Framework 2.0 SP2 and Windows installer 4.5. On Windows Vista, use .NET Framework 3.5 SP1.
SQL Server Express with Advanced Services	.NET Framework 3.5 SP1, Windows Installer 4.5, and Windows PowerShell™ 1.0.
SQL Server Express with Tools	.NET Framework 3.5 SP1, Windows Installer 4.5, and Windows PowerShell 1.0.

For more information, see Hardware and Software Requirements for Installing SQL Server 2008.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

How SQL Server 2008 Supports NUMA



Key Points

SQL Server 2008 supports nonuniform memory access (NUMA) and works on NUMA hardware without special configuration. As the number of processors and clock speed increases for hardware, it becomes difficult to reduce the memory latency required to use this additional processing power. NUMA architecture provides a scalable solution to this problem. SQL Server has been designed to take advantage of NUMA-based computers without requiring any application changes.

NUMA Architecture

The trend in hardware is to have more than one system bus, each serving a small set of processors. Each group of processors has its own memory and possibly its own I/O channels. However, each CPU can access memory associated with the other groups in a coherent way. Each group is called a NUMA node. The number of CPUs within a NUMA node depends on the hardware vendor. It is faster to access local memory than the memory associated with other NUMA nodes.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

On NUMA hardware, some regions of memory are on physically different buses from other regions. It will take longer to access some regions of memory because NUMA uses local and foreign memory. Local memory and foreign memory are used in reference to a currently running thread. Local memory is the memory that is on the same node as that of the CPU currently running the thread. Any memory that does not belong to the node on which the thread is currently running is foreign. Foreign memory is also known as remote memory. The ratio of the cost to access foreign memory over local memory is called the NUMA ratio. If the NUMA ratio is 1, it is symmetric multiprocessing (SMP). The more the ratio, the more it costs to access the memory of other nodes. Windows® applications that do not support NUMA (including SQL Server 2000 SP3 and earlier) sometimes perform badly on NUMA hardware.

Hardware NUMA and Soft-NUMA

NUMA can match memory with CPUs through specialized hardware known as hardware NUMA or by configuring SQL Server memory known as soft-NUMA. When SQL Server starts, it is automatically configured based on the underlying operating system and hardware configuration or the soft-NUMA setting. For both hardware and soft-NUMA, when SQL Server starts in a NUMA configuration, the SQL Server log records a multimode configuration message for each node, along with the CPU mask. NUMA can be divided into:

- **Hardware NUMA.** Computers with hardware NUMA have more than one system bus, each serving a small set of processors. Each group of processors has its own memory and possibly its own I/O channels, but each CPU can access memory associated with other groups in a coherent way. Each group is called a NUMA node. The number of CPUs within a NUMA node depends on the hardware vendor. If you have hardware NUMA, you may configure it to use interleaved memory instead of NUMA. In that case, Windows and SQL Server will not recognize it as NUMA. Run the following query to find the number of memory nodes available to SQL Server:

```
SELECT DISTINCT memory_node_id
FROM sys.dm_os_memory_clerks
```

- **Soft-NUMA.** SQL Server allows you to group CPUs into nodes referred to as soft-NUMA. You can configure soft-NUMA when there are a lot of CPUs and when hardware NUMA is unavailable. You can also use soft-NUMA to subdivide hardware NUMA nodes into

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

smaller groups. Only the SQL Server scheduler and SQL Server Network Interface (SNI) support soft-NUMA. You can create memory nodes based on hardware NUMA and soft-NUMA does not affect them. For example, if you have an SMP computer with eight CPUs and you create four soft-NUMA nodes with two CPUs each, you will only have one memory node serving all four NUMA nodes. Soft-NUMA does not provide memory to CPU affinity.

How to Configure SQL Server to Use Soft-NUMA

To configure SQL Server to use soft-NUMA, you must edit the registry to add a node configuration affinity mask. The soft-NUMA mask can be stated as a binary or DWORD (hexadecimal or decimal) registry entry. You must restart the Database Engine to configure soft-NUMA.

If you edit the registry incorrectly, it can severely damage your system. Before making changes to the registry, you must back up any valued data on the computer.

Grouping Common CPUs

For example, consider a computer with eight CPUs that does not have hardware NUMA. Three soft-NUMA nodes are configured. Database Engine instance A is configured to use CPUs 1 through 4. A second instance of the Database Engine, instance B, is installed and configured to use CPUs 5 through 8. You can represent this as:

CPUs 1 2 3 4 5 6 7 8

Soft-NUMA <-N0--><-N1-><----N2---->

SQL Server <instance A ><instance B>

Instance A, which experiences significant I/O, has two I/O threads and two lazy writer threads, while instance B, which performs processor-intensive operations, has only one I/O thread and one lazy writer thread. Differing amounts of memory can be assigned to the instances, but unlike hardware NUMA, they both receive memory from the same operating system memory block and there is no memory-to-processor affinity.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Set the CPU affinity mask

Run the following statement on instance A to configure it to use CPUs 1, 2, 3, and 4 by setting the CPU affinity mask:

```
sp_configure 'show advanced options', 1;
RECONFIGURE;
GO
sp_configure 'affinity mask', 1;
RECONFIGURE;
GO
```

Run the following statement on instance B to configure it to use CPUs 5, 6, 7, and 8 by setting the CPU affinity mask:

```
sp_configure 'show advanced options', 1;
RECONFIGURE;
GO
sp_configure 'affinity mask', 1;
RECONFIGURE;
GO
```

Note: The affinity mask value in the example applies to a computer with a 32-bit operating system.

Map soft-NUMA nodes to CPUs

By using the Registry Editor program, regedit.exe, add the following registry keys to map soft-NUMA node 0 to CPUs 1 and 2, soft-NUMA node 1 to CPUs 3 and 4, and soft-NUMA node 2 to CPUs 5, 6, 7, and 8.

SQL Server 2008	Type	Value name	Value data
HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Microsoft SQL Server\100\NodeConfiguration\Node0	DWORD	CPUMask	0x03
HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Microsoft	DWORD	CPUMask	0x0c

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

SQL Server 2008	Type	Value name	Value data
SQL Server\100\NodeConfiguration\Node1			
HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Microsoft SQL Server\100\NodeConfiguration\Node2	DWORD	CPUMask	0xf0

Source:

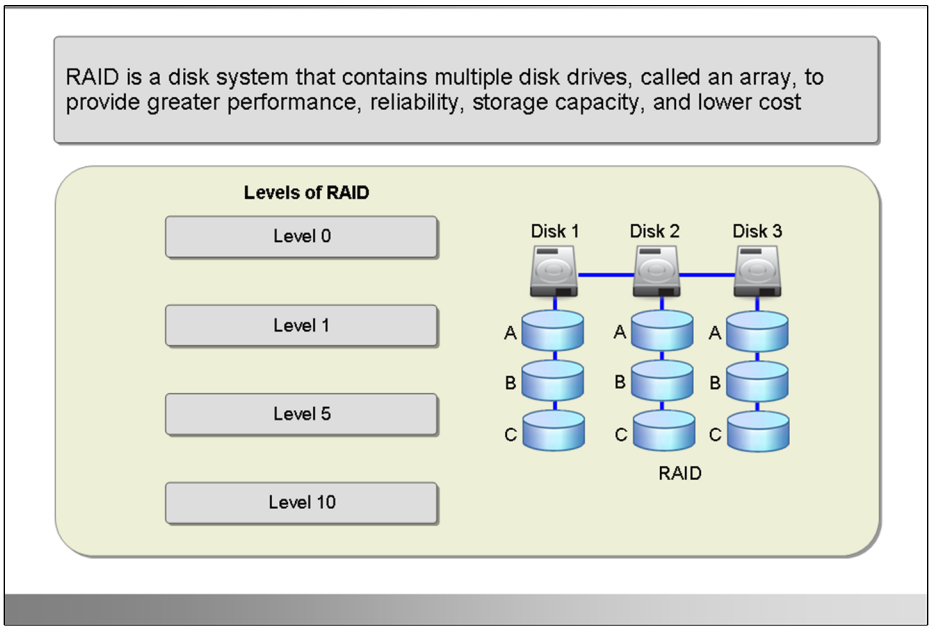
- http://msdn.microsoft.com/en-us/library/ms178144.aspx
- http://technet.microsoft.com/en-us/library/ms345357.aspx
- http://msdn.microsoft.com/en-us/library/ms345357.aspx

For more information, see Understanding Non-uniform Memory Access.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Using RAID with SQL Server 2008



Key Points

Redundant array of inexpensive disks (RAID) is a disk system that comprises multiple disk drives known as an array to provide higher performance, reliability, storage capacity, and lower cost.

Although RAID is not a part of SQL Server, implementing RAID can directly affect the way SQL Server performs.

There are several RAID levels, but levels 0, 1, 5, and 10 are commonly implemented with SQL Server 2008. The descriptions of these RAID levels are as follows:

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

- **Level 0.** This level is also known as disk striping because it uses a disk file system called a stripe set. Data is divided into blocks and spread in a fixed order among all the disks in an array. RAID 0 improves read and write performance by spreading operations across multiple disks. You can then perform operations independently and at the same time. RAID 0 is similar to RAID 5, but RAID 5 also provides fault tolerance. You need at least two disks to implement RAID 0.
- **Level 1.** This level is also known as disk mirroring because it uses a disk file system called a mirror set. Disk mirroring provides a redundant, identical copy of a selected disk. All data written to the primary disk is written to the mirror disk. RAID 1 provides fault tolerance and generally improves read performance but write performance may degrade. You need at least two disks to implement RAID 1.
- **Level 5.** This level is also known as striping with parity and is the most popular strategy for new designs. RAID 5 is similar to RAID 4, because this level stripes the data in large blocks across the disks in an array. However, RAID 5 differs because it writes the parity across all the disks. The parity information provides data redundancy. The data and parity information are arranged on the disk array so that the two types of information are always on different disks. Striping with parity offers better performance than disk mirroring (RAID 1). However, when a stripe member is missing, read performance is decreased, like, when a disk fails. You need at least three disks to implement RAID 5.
- **Level 10 (1+0).** This level is also known as mirroring with striping. RAID 10 uses a striped array of disks that are then mirrored to another identical set of striped disks. Because it is a RAID 0 with all disks mirrored, it is called RAID 10. For example, you can create a striped array by using five disks. You can then mirror the striped array of disks by using another set of five striped disks. RAID 10 provides the performance benefits of disk striping with the disk redundancy of mirroring. RAID 10 provides the highest read-and-write performance of any one of the other RAID levels, but by using many disks. You need at least four disks to implement RAID 10. RAID levels higher than 10 (1 + 0) may offer additional fault tolerance or performance enhancements. These levels, generally, are proprietary systems.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Source:

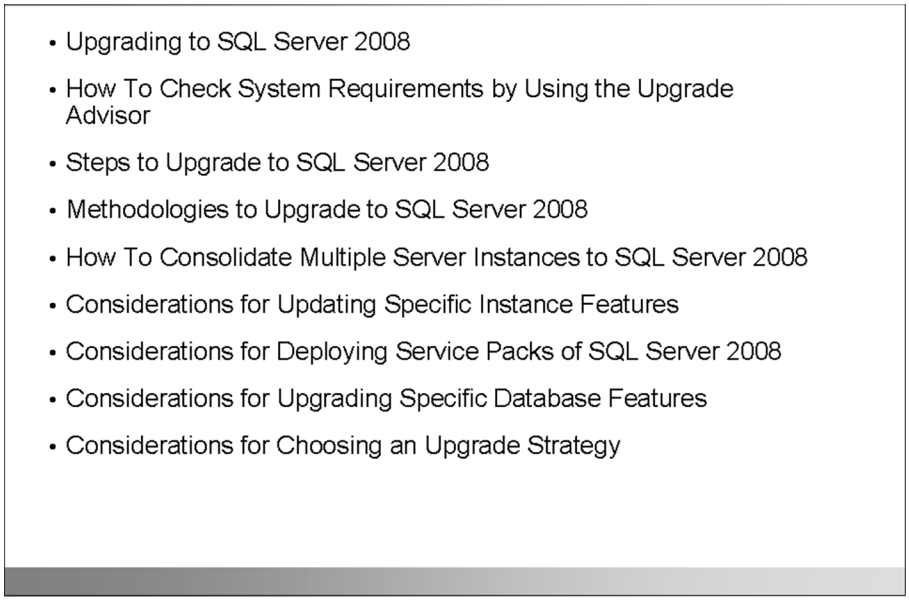
- http://technet.microsoft.com/en-us/library/cc917578.aspx
- http://msdn.microsoft.com/en-us/library/ms190764(SQL.105).aspx
- http://msdn.microsoft.com/en-us/library/ms190764(SQL.105).aspx
- http://msdn.microsoft.com/en-us/library/ms190764(SQL.105).aspx
- http://msdn.microsoft.com/en-us/library/ms190764(SQL.105).aspx

For more information, see RAID Levels and SQL Server.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Lesson 2
Upgrading and Migrating to SQL Server 2008



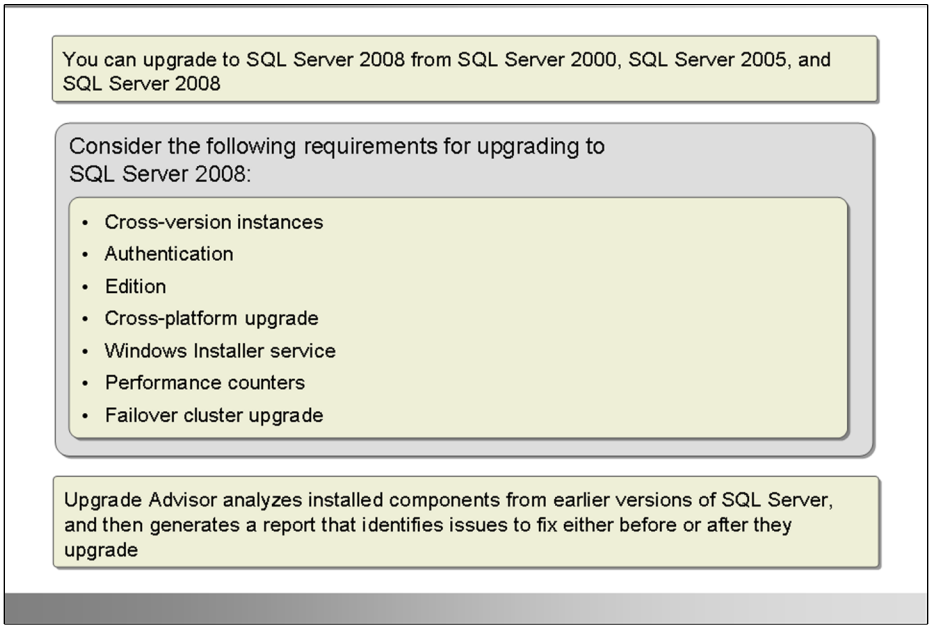
For an organization that is using SQL Server, upgrading to SQL Server 2008 would bring additional business value. There are many methods to upgrade to SQL Server 2008. However, you should use the correct methodology to manage the upgrade process to ensure the success of the upgrade.

In this lesson, you will learn about the various methods to upgrade to SQL Server 2008 and the steps to perform the upgrade. You will also learn about the considerations for updating specific instance features, deploying service packs (SPs) of SQL Server 2008, upgrading specific database features, and choosing an upgrade strategy.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Upgrading to SQL Server 2008



Key Points

The SQL Server Setup program supports upgrading to SQL Server 2008 from SQL Server 2000, SQL Server 2005, and SQL Server 2008. You can also migrate databases from earlier SQL Server versions. You can migrate one SQL Server instance to another on the same computer, or to a SQL Server instance on another computer. You can migrate databases by using the Copy Database Wizard, Backup and restore functionality, SQL Server Integration Services (SSIS) Import and Export Wizard, and bulk export/bulk import methods. Before upgrading to SQL Server 2008, you need to consider the following requirements:

- **Cross-version instances.** SQL Server 2008 supports cross-version instances. However, you need to ensure that the version numbers of the Database Engine, Analysis Services, and Reporting Services components are the same in an instance of SQL Server 2008.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

- **Authentication.** Before upgrading to SQL Server 2008, you should enable Windows Authentication for SQL Server Agent and verify the default configuration.
- **Edition.** Before upgrading from the previous edition of SQL Server to SQL Server 2008, verify that the edition to which you are upgrading to supports the functionality you are using currently.
- **Cross-platform upgrade.** SQL Server 2008 does not support cross-platform upgrade. You cannot upgrade a 32-bit instance of SQL Server to native 64-bit. However you can upgrade a 32-bit instance of SQL Server to a 32-bit instance of SQL Server running Microsoft Windows-32-on-Windows-64 (WoW64) platform on x64 version of Windows. *Operating System.* To upgrade to SQL Server 2008, you must be running a supported operating system. Upgrade will be blocked if there is a pending restart.
- **Windows Installer service.** Upgrade will be blocked if the Windows Installer service is not running.
- **Performance counters.** Upgrade will be blocked if performance counters are corrupt.
- **Failover cluster upgrade.** To upgrade an instance of SQL Server to a SQL Server failover cluster, the instance being upgraded must be a failover cluster. To upgrade a stand-alone instance of SQL Server to a SQL Server failover cluster, install a new SQL Server failover cluster and then migrate user databases from the stand-alone instance by using the Copy Database Wizard.

Issues to Check in the Current Installation of SQL Server

Before upgrading to SQL Server 2008, you need to review the following issues in the current installation of SQL Server:

- Upgrade Analysis Services before you upgrade the Database Engine, when upgrading from a 64-bit edition of SQL Server to a 64-bit edition of SQL Server 2008.
- Back up all SQL Server database files from the instance to be upgraded, so that you can restore them, if required.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

- Run the appropriate Database Console Commands (DBCC) on databases to be upgraded to ensure that you are in a consistent state.
- Estimate the disk space that is required to upgrade SQL Server components, in addition to user databases.
- Ensure that existing SQL Server system databases—master, model, msdb, and tempdb—are configured to auto-grow, and ensure that they have sufficient hard disk space.
- Ensure that all database servers have logon information in the master database. This is important for restoring a database, as system logon information resides in the master database.
- Disable all startup stored procedures because the upgrade process will stop and start the services on the SQL Server instance being upgraded. Stored procedures processed at startup time might block the upgrade process.
- Stop replication and ensure that the replication log is empty.
- Change the max worker threads value to 0 before upgrading to let the Database Engine calculate the optimal number of threads.
- Exit all applications, including all services that have SQL Server dependencies. The upgrade might fail if local applications are connected to the instance being upgraded.
- Minimize downtime for mirrored databases when upgrading server instances, if you are using Database Mirroring, because the passive server can continue to serve user request while you work on upgrading the active server.

Upgrade Advisor

You can prepare for SQL Server 2008 upgrade by using SQL Server Upgrade Advisor. Upgrade Advisor analyzes installed components from earlier versions of SQL Server, and then generates a report that identifies issues to fix either before or after they upgrade. When you run Upgrade Advisor, the Upgrade Advisor Home page appears. From the Home page, you can run the Upgrade Advisor Analysis Wizard, Upgrade Advisor Report Viewer, and Upgrade Advisor Help tools. When you use Upgrade Advisor for the first time, run the Upgrade Advisor Analysis Wizard to analyze SQL Server components. When the wizard completes the analysis, view the resulting reports in the

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Upgrade Advisor Report Viewer. Each report provides links to information in Upgrade Advisor Help that will help you fix or reduce the effect of the known issues. Upgrade Advisor analyzes the Database Engine, Analysis Services, Reporting Services, Integration Services, and Data Transformation Services SQL Server components.

Note: You can upgrade from the retired versions of SQL Server 6.5 and SQL Server 7 to SQL Server 2008.

Source:

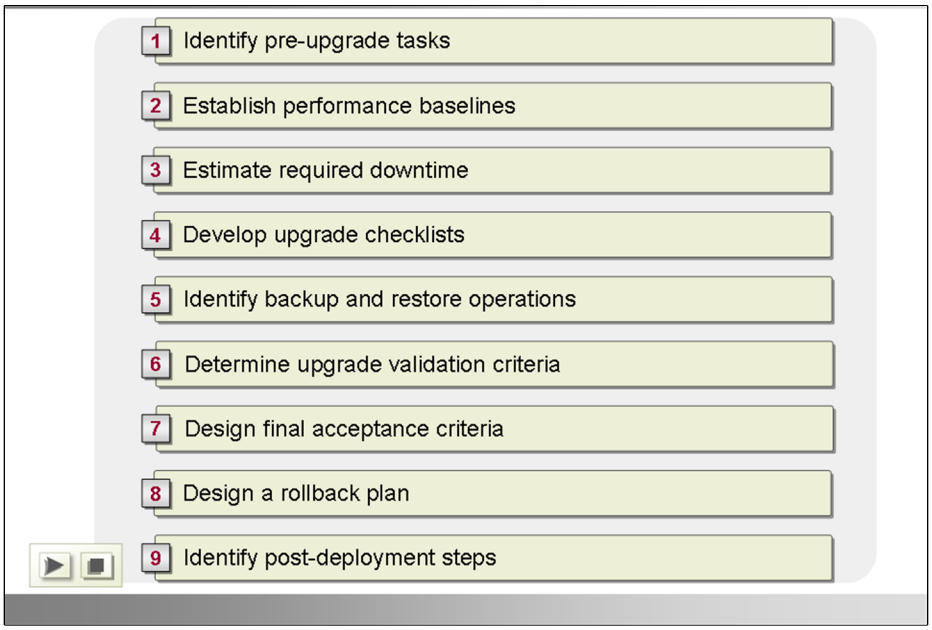
- http://msdn.microsoft.com/en-us/library/ms143393(SQL.105).aspx
- http://msdn.microsoft.com/en-us/library/ms143393(SQL.105).aspx
- http://msdn.microsoft.com/en-us/library/bb933942.aspx

For more information, see Upgrading to SQL Server 2008.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Steps to Upgrade to SQL Server 2008



Key Points

To upgrade to SQL Server 2008 from previous versions of SQL, you need to:

1. **Identify preupgrade task**s. This step includes installing Microsoft Windows Installer (MSI) 4.5, .NET Framework 3.5 SP1, and the SQL Server Native Client on the target instances. You can install these components before the actual upgrade deployment begins. However, you need to restart the server after you install MSI 4.5.
2. **Establish performance baselines**. Collect data that indicates typical performance measurements for important and common queries. If reports show that the performance has not improved after the upgrade, refer to these guidelines. Users might be mistaken, and you might find through the baselines that the new system performs equally well or better.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

3. **Estimate required downtime**. The deployment of an upgrade will involve some downtime for the targeted database servers. When you perform the actual upgrade, ensure you notify users to allow for enough downtime so that the processes can be completed successfully. Try to give yourself time to roll back if an unexpected issue arises that you cannot resolve in the downtime window. This might mean that you must decide whether to finish the upgrade or roll back within the downtime window.
4. **Develop upgrade checklists**. The server environment for targeted database servers might have their own infrastructure complexities. Detail the steps that you must follow for taking the systems offline for a while and bringing them back online. Also detail the steps to take during the upgrade processes. The upgrade steps, in particular, might be more complex. Regardless of the method you use, you might have to apply scripts at certain important points to resolve issues that are identified by Upgrade Advisor.
5. **Identify backup and restore operations**. One of the first steps in the deployment plan should be to back up the targeted databases. Also verify the backups, and decide a strategy for restoring them if necessary.
6. **Determine upgrade validation criteria**. Clearly state what criteria your organization will use to validate that the upgrade was successful. This might consist of scripts that are run to inspect the instances of SQL Server 2008, to verify that issues are resolved, and to ensure that configuration settings, like security settings, are as retained. It might also include bringing applications online selectively, and processing some transactions to confirm successful operation.
7. **Design final acceptance criteria**. The upgrade might succeed at the SQL Server 2008 instance level, but some other unaccounted-for variable in the server infrastructure might still prevent applications from running correctly. Whatever the case, determine how the organization will verify the upgrade, and how it will make the fallback decision if problems arise. This goes beyond validating the upgrade result. It focuses on whether the applications that use the targeted database servers run as expected and required. Usually, the application team would execute a series of tests against a test case defined before the

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

upgrade to ensure that the application control functions properly. It might be appropriate to enlist the support of the QA team to develop appropriate acceptance tests.

8. **Design a rollback plan**. If the upgrade is unsuccessful or if the acceptance tests are unsuccessful, be prepared to back out of the process and restore the original conditions. This is much easier to do with a side-by-side strategy than with an in-place upgrade. For example, with an in-place upgrade, a rollback plan might require restoring a disk image, which is also known as a "ghost" image, of the computer that is running SQL Server 2000 or SQL Server 2005, and then restoring the SQL Server 2000 or SQL Server 2005 databases from the deployment backup.
9. **Identify post-deployment steps**. Even after you validate and accept the upgrade, you might have some remaining tasks to perform, such as updating statistics in the relational database or rebuilding cubes in SQL Server Analysis Services (SSAS). You might also have to reconfigure log shipping, reconfigure database mirroring, reestablish replication, test a failover cluster, or verify that certain SQL Server Agent jobs run correctly.

After you upgrade to SQL Server 2008, you need to complete the following tasks:

- **Register your servers.** Upgrade removes registry settings for the previous instance of SQL Server. After you upgrade, you must re-register your servers.
- **Update statistics.** To help optimize query performance, you should update statistics on all databases after the upgrade. Use the sp_updatestats stored procedure to update statistics in user-defined tables in SQL Server databases.
- **Configure your new SQL Server installation**. To reduce the attackable surface area of a system, SQL Server selectively installs and enables key services and features.

Question: What should you check after upgrading to SQL Server 2008?

Source:

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

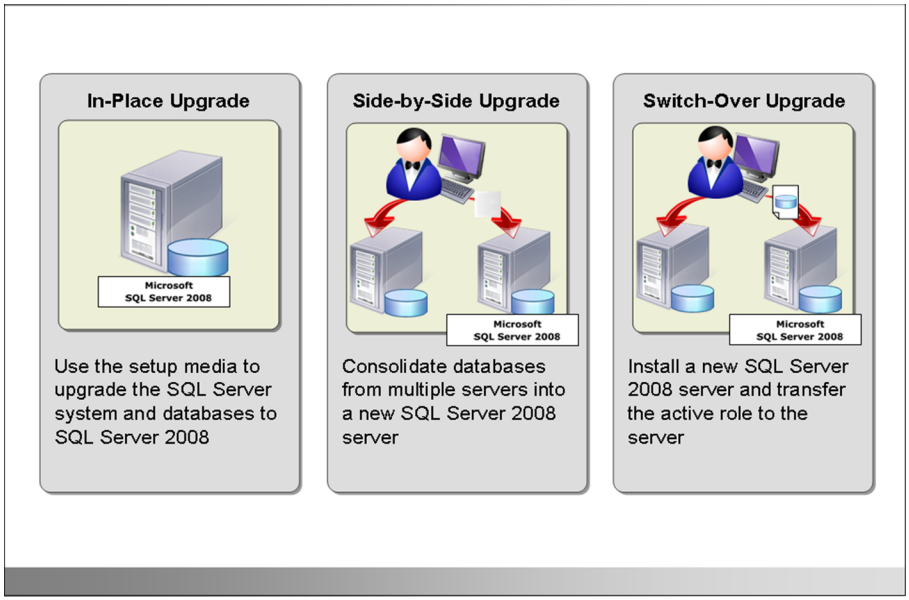
New Horizons Computer Learning Centers, Inc.

- http://msdn.microsoft.com/en-us/library/bb933942.aspx

For more information, see Upgrade How-to Topics.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

Methodologies to Upgrade to SQL Server 2008



Key Points

The various methods to upgrade to SQL Server 2008 include in-place upgrade, side-by-side upgrade, and switch-over upgrade.

In-place upgrade

In the in-place upgrade approach, you use the setup media to upgrade the SQL Server system and databases to SQL Server 2008 by inserting the installation media into the server. You can perform the upgrade by using the following high-level steps:

1. Backup all databases.
2. Perform in-place upgrade by using SQL 2008 setup.
3. Verify the database functionalities.
4. After the upgrade, perform post upgrade tests, such as changing compatibility level or re-enabling full-text index.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Side-by-side upgrade

In the side-by-side upgrade approach, you consolidate databases from multiple servers into a new and more powerful SQL Server 2008 server. The benefit of this approach is that it is easier to fall back or revert to the original version, if required. You can perform the upgrade by using the following high-level steps:

1. Install a new SQL Server 2008 server.
2. Script all the required system configurations objects, such as user and database role, from existing servers to the new server.
3. Use Database Copy Wizard to transfer the database to the new server.
4. Redirect users to the new SQL Server 2008.
5. Decommission the old SQL Server, after testing.

Note: In a switch-over upgrade, instead of using Detach/Attach approach, you can also use backup/restore. However, you get better performance in most case by using detach /attach.

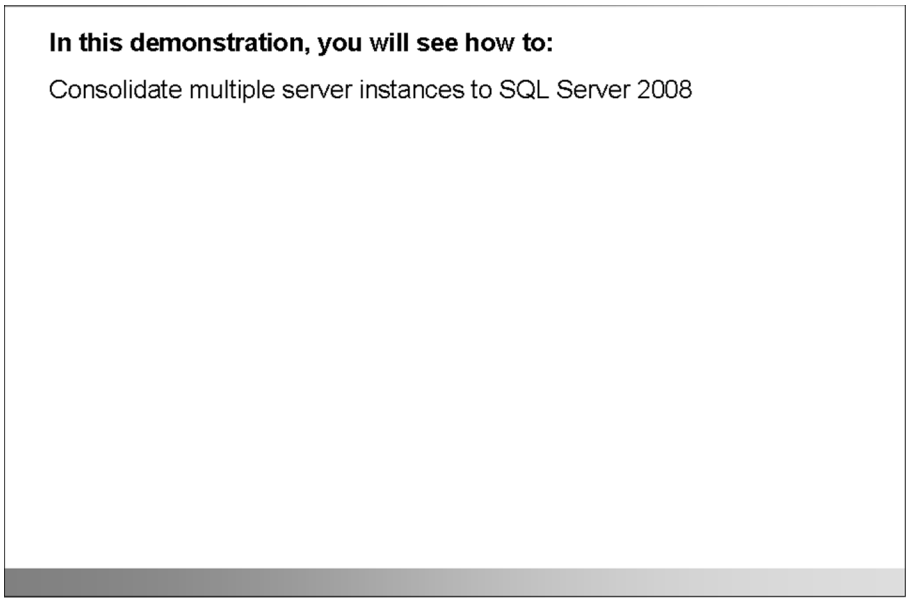
Question: Which methodology should you use to reduce the risk associated with an upgrade?

Question: When should you consider performing an in-place upgrade?

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Demonstration: How To Consolidate Multiple Server Instances to SQL Server 2008



Key Points

The steps to consolidate multiple server instances to SQL Server 2008 are:

1. On the **Start** menu, point to **All Programs**, click **Microsoft SQL Server 2008**, and then click **SQL Server Management Studio**. The Microsoft SQL Server Management Studio window appears.
2. To log on to SQL Server Management Studio, in the **Connect to Server** dialog box, type the following:
   - Server type: **Database Engine**
   - Server name: **NYC-SQL1**
   - Authentication: **Windows Authentication**
3. In the **Connect to Server** dialog box, click **Connect.**

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

4. In the Microsoft SQL Server Management Studio window, in the Object Explorer pane, under **NYC-SQL1(SQL Server 10.0.1600-NYC-SQL1\Administrator)**, expand **Databases**, right-click **QuantamCorpDW**, point to **Tasks**, and then click **Copy Database**. The Copy Database Wizard launches.
5. In the Copy Database Wizard, on the **Welcome to Copy Data Wizard** page, click **Next**.
6. On the **Select a Source Server** page, in the **Source server** box, type **NYC-SQL1**, ensure that the **Use Windows Authentication** option is selected, and then click **Next**.
7. On the **Select a Destination Server** page, in the **Destination server** box, click **browse** button. The **Browse for Servers** dialog box appears.
8. In the **Browse for Servers** dialog box, expand **Database Engine**, select **NYC-SQL1\DEVELOPMENT** and then click **OK**.
9. On the **Select a Destination Server** page, click **Next**.
10. On the **Select the Transfer Method** page, click **Use the SQL Management Object method**, and then click **Next**.
11. On the **Select Databases** page, under the **Copy** column, ensure that the **QuantamCorpDW** check box is selected, and then click **Next**.
12. On the **Configure Destination Database (1 of 1)** page, in the **Destination database** box, type **QuantamCorpDW_new**, and then click **Next**.
13. On the **Select Server Objects** page, under **Available related objects**, click an object, and click the forward arrow button, and then click **Next**.
14. On the **Configure the Package** page, accept default SSIS package properties, and then click **Next**.
15. On the **Schedule the Package** page, ensure that the Run immediately option is selected, and then click **Next**.
16. On the **Complete the Wizard** page, click **Finish**.
17. On the **Performing operation** page, click **Close**.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

Question: What are the benefits of using Copy Database Wizard over the direct upgrade?

Question: Why should you consider consolidating multiple SQL Servers?

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Considerations for Updating Specific Instance Features

Specific Instance Feature	Consideration
 Cluster	• Back up both server instances • Perform the upgrade on the active node • After upgrade, shutdown the passive node, and then restart the active node
 Log Shipping	• Always update the secondary server instance first • Consider the amount of time the database will be unavailable
 SQL Server Agent	• Verify that the job is working before upgrade • After upgrade, change to domain/local user account as the service account
 SQL Server Integration Services (SSIS)	• While upgrading, you must reconfigure an existing installation of SSIS

Key Points

Before upgrading, always backup the system and the user databases. You should also schedule a downtime with users so that they do not change data during the upgrade. If users continue to make changes while you perform the upgrade, you may lose data or your upgrade would fail.

While updating specific instance features, you should consider the following best practices**:**

- **Cluster.** In a cluster, always upgrade both the server instances on the active node. After the upgrade is complete, shutdown the passive node and restart the active node. This helps to reduce the downtime that would occur during the reboot or upgrade.
- **Log shipping.** The upgrade process involves upgrading the secondary server instance of a SQL Server 2005 log shipping configuration to SQL Server 2008, before upgrading the primary

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

server instance. You must also consider the availability of the database while planning an upgrade. You can maximize your database availability by failing over a SQL Server 2005 primary server to a SQL Server 2008 secondary server, before upgrading the original primary server.

- **SQL Server Agent.** Before the upgrade, check if any job is working. After the upgrade, use a local user account as the SQL Server service account. You should consider backing up the msdb database and all scripts used by SQL Server Agent. If required, you can generate a set of scripts for all SQL Server agent jobs for backup.
- **SQL Server Integration Services (SSIS).** While upgrading, if an installation of SSIS exists, you must reconfigure it. You can upgrade from a 32-bit SQL Server version to a 64-bit SQL Server version but not from a 64-bit SQL Server version to a 32-bit SQL Server version. You can also move from one localized version of SQL Server to another localized version.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

Considerations for Deploying Service Packs of SQL Server 2008

Specific Instance Feature	Consideration
 Cluster	• Back up both server instances • Perform the upgrade on the active node • After upgrade, shutdown the passive node, and then restart the active node
 Log Shipping	• Always update the secondary server instance first • Consider the amount of time the database will be unavailable
 SQL Server Agent	• Verify that the job is working before upgrade • After upgrade, change to domain/local user account as the service account
 SQL Server Integration Services (SSIS)	• While upgrading, you must reconfigure an existing installation of SSIS

Key Points

While deploying service packs of SQL Server 2008, you should consider the following guidelines:

- Schedule the deployment of service packs during non-office hours to minimize potential service outage to business services.
- Perform a rolling upgrade for database server with database mirroring to reduce downtime for each mirrored database.
- Test the application against updates before massive rollout.
- Consider installing hotfix and service packs to a pilot group before massive rollout.
- Always deploy updates to non-mission critical systems before mission critical systems.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

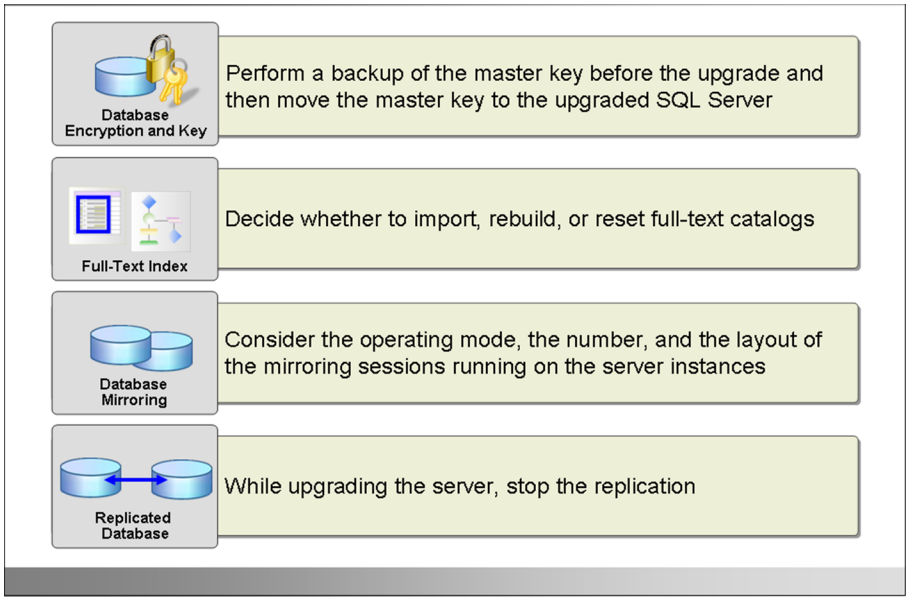
- For database server with database mirroring, the upgrade process involves sequentially upgrading the different instances to Microsoft SQL Server 2008. This form of upgrade, known as a rolling upgrade, reduces downtime to only a single failover.
- Before performing a rolling upgrade, you should consider the following best practices:
- Perform a full-system backup to protect your data. This will help ensure that you have a valid copy of the system backup to help rollback the system if the upgrade fails.
- Run the DBCC CHECKDB command on every principal database. Corrupted databases can cause a failure in the upgrade. This command will ensure that the databases are not corrupted.
- Remove the witness server, if a mirroring session includes a witness server. This would help prevent a switch over happening during the upgrade which would result in data loss.
- Change the operating mode to the high-safety mode, if a mirroring session is running in the high-performance mode. This is to ensure that the changes are replicated to the passive server and SQL Server would ensure that the changes are committed to the other machine before being committed to the active server. This change helps to maintain the data integrity during the upgrade.
- Perform a practice manual failover on at least one of your mirroring sessions. This is to ensure that the team that will perform the upgrade is familiar with the failover procedure as that is required during the upgrade. Failure to perform the failover properly would cause the database to corrupt during the upgrade.

Source:

- http://msdn.microsoft.com/en-us/library/bb497962.aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

Considerations for Upgrading Specific Database Features



Key Points

While upgrading specific database features, you should consider the following best practices:

- **Database encryption and keys**. Before the upgrade, create a backup of the database master key and the service master key. In addition, you need to have the keys available and accessible at all times. Without the keys, you would not be able to access the information even if you want to recover the data.
- **Full-text index**. When you upgrade a server instance to SQL Server 2008, you can import, rebuild, or reset full-text catalogs. Because the full-text indexes are not upgraded during the installation of SQL Server, you have to recreate the index after upgrade. Importing full-text catalogs is faster than rebuilding.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

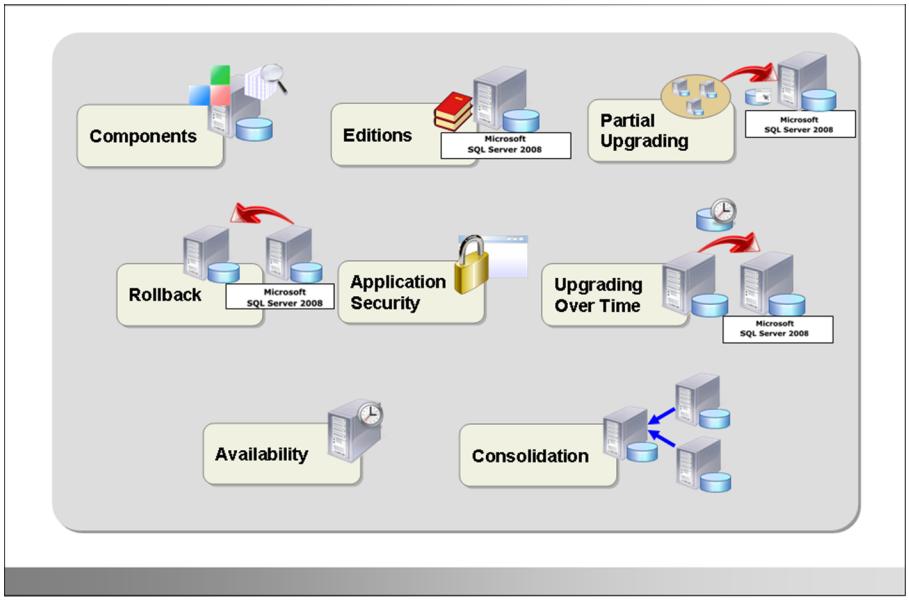
- **Database mirroring**. When you upgrade server instances from SQL Server 2005 to SQL Server 2008, perform a rolling upgrade to reduce the downtime for each mirrored database to only a single manual failover.
- **Replicated database**. While upgrading to SQL Server 2008, stop the replication. This is because the upgrade will modify the table schema which may get replicated to the subscriber, if you are replicating while upgrading. This would cause the subscriber system to fail as the subscriber system may not get upgraded to support the changes applied during the upgrade. After the upgrade is complete on both nodes, restart the replication.

Question: If you have full-text indexing enabled, what should you perform after the upgrade?

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Considerations for Choosing an Upgrade Strategy



Key Points

While selecting an appropriate upgrade method for SQL Server, you need to consider the following factors:

- **Components.** You cannot use a certain upgrade strategy if a component does not support it. For example, you cannot perform an in-place upgrade for SSIS from SQL Server 2000. Also, you should transfer most SSAS components if the source is SQL Server 2000.
- **Editions.** The in-place upgrade strategy does not support all paths between editions. For example, to upgrade a SQL Server 2000 or SQL Server 2005 Enterprise instance to SQL Server 2008 Standard, you must perform a side-by-side upgrade because SQL Server Setup does not support an in-place upgrade path.
- **Partial Upgrading.** To transition only a few databases on a server to SQL Server 2008, perform a side-by-side upgrade.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

- **Rollback.** For many database systems in production, it is impossible to validate a change without a rollback strategy, in case the results are unacceptable. The side-by-side upgrades strategy supports rollback at the time of acceptance testing because the legacy instance is available. However, after users update the databases in the new instance, rollback might not work. Regardless of the strategy you select, you need to test and validate the process. Even if you choose an in-place upgrade strategy, you should test the upgrade process and results on a separate server first.
- **Application Security.** If your organization requires minimal disturbance to the existing applications and users, you can perform an in-place upgrade.
- **Upgrading Over Time.** To transition databases gradually from a legacy instance to SQL Server 2008, you must perform a side-by-side upgrade.
- **Availability.** Databases will become unavailable for some time for both an in-place upgrade and a side-by-side upgrade. The downtime required depends primarily on the size of the data sets. At first, it might seem that an in-place upgrade would be faster than a side-by-side upgrade because the data is not transferred from one server to another. However, an in-place upgrade also requires time for the installation of SQL Server 2008. In a side-by-side upgrade, SQL Server 2008 is already installed on another instance. If the data transfer proceeds quickly and few changes are needed on the new instance, a side-by-side upgrade might be faster than an in-place upgrade.
- **Consolidation.** You should check the performance of the new hardware used to install the SQL Server. If the performance of the new server is more powerful, you should consider consolidating the multiple SQL servers to a single server to leverage better hardware. You may also consider performing consolidation if the resources consumed by the database are low. You can then consolidate the databases, which consume less system resources into a single system.

Question: Why should you consider rollback time when planning for an upgrade?

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

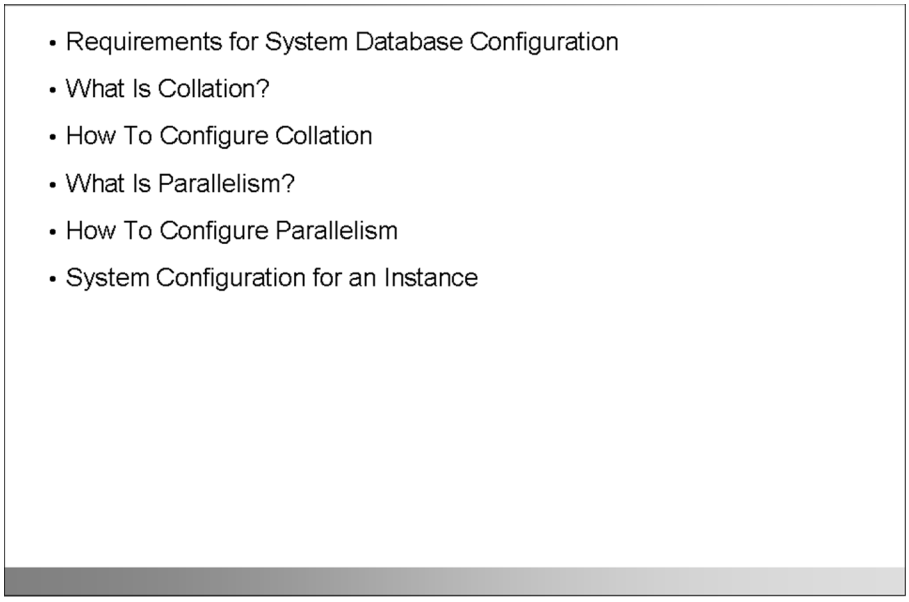
Source:

- http://download.microsoft.com/download/5/3/D/ 53D72434-7BD5-41C6-A806-8212C1B0DCA1/ SQL%20Server%202005%20to%202008%20Upgrade%20White%20Paper.docx
- http://download.microsoft.com/download/D/2/2/ D22D16C3-7637-41D3-99DA-10E7CEBAD290/SQL2008UpgradeTechnicalReferenceGuide.docx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Lesson 3
Configuring Instances in SQL Server 2008



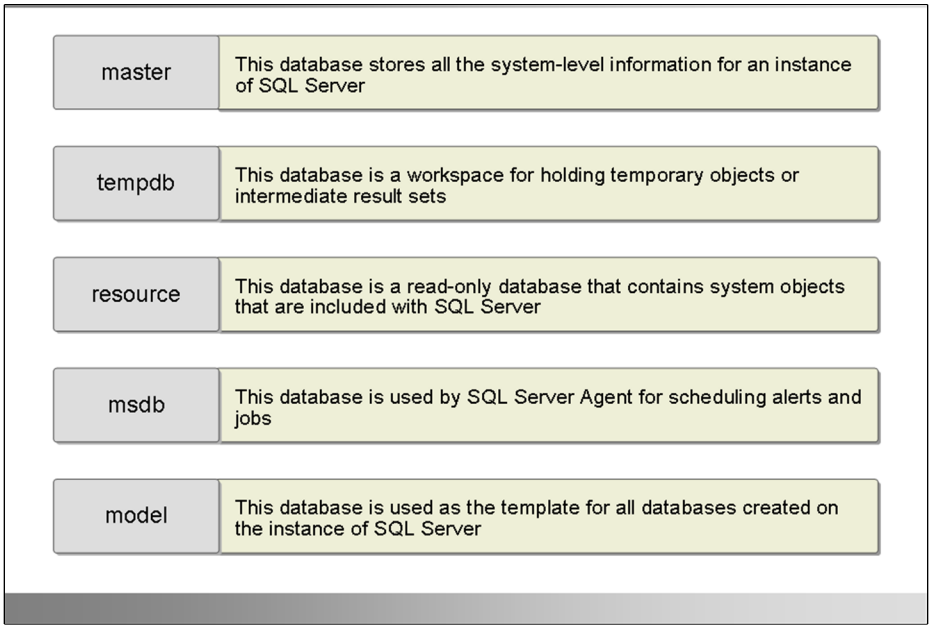
Server hardware usually comes with multiple processors or even multiple processing units within a CPU. SQL Server supports leveraging these server resources to provide better performance and response to users. But sometimes, you would want to restrict the usage of these resources to avoid overuse of system resources. You would also need to configure the behavior of the SQL Server system so that it performs as required by applications.

In this lesson, you will learn how to configure instances in SQL Server 2008. Although SQL Server 2008 is a highly scalable program that can run on various platforms, a proper configuration is important to gain optimized performance.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Requirements for System Database Configuration



Key Points

When you initially setup SQL Server 2008, you need to install system databases such as msdb, tempdb, resource, master, and model. These databases have different usage and configuration requirements.

The following table lists the requirements for configuring system databases.

System database	Requirement
master	The master database stores all the system-level information for an instance of SQL Server. While configuring the master database, you should consider the following best practices to gain optimized performance: • Always have a current backup of the available master database. • Back up the master database immediately after

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

System database	Requirement
	you create, modify, or drop any database; change server or database configuration values; and modify or add logon accounts. • Do not create user objects in the master database. If you create user objects, back up the master database frequently. • Do not enable the TRUSTWORTHY option for the master database.
tempdb	The tempdb database is a workspace for holding temporary objects or intermediate result sets. While configuring the tempdb database, you should consider the following best practices to gain optimized performance: • Cache the temporary tables and table variables. • Improve the allocation page latching protocol. • Reduce the logging overhead for the tempdb database. • Improve the algorithm for allocating mixed pages in the tempdb database.
resource	The resource database is a read-only database that contains system objects that are included with SQL Server. While configuring the resource database, use a drive with fast read access as the resource database is not written. In SQL Server 2008, the default location of the Resource database is <drive>:\Program Files\Microsoft SQL Server\MSSQL10.<instance_name>\Binn\. The database can be moved. However, you must not be move it for the following two reasons: • Applying SQL Server service packs and hotfixes restores the database to the \Binn location. • Moving the Resource database in a failover cluster environment to a nonclustered location will cause failover cluster failure.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

System database	Requirement
msdb	The msdb database is used by SQL Server Agent for scheduling alerts and jobs. While configuring the msdb database, use a faster access drive if SQL Server Agent is enabled.
model	The model database is used as the template for all databases created on the instance of SQL Server. In the model database, you should specify the configuration that should be there when you create a database. For example, you can have the database configured with a schema every time a database is created. When you make changes to model database, ensure that the changes are compatible with the tempdb database as the tempdb database is created from model database every time SQL Server starts.

Considerations for Changing System Databases

- You should always create a backup for your system databases. You should have a backup after you make user defined changes to the databases.
- You should also document the changes applied to these databases to ensure that you and your team are aware of the changes.
- You should ensure that the changes are consistent in your environment. For example, you should make the same changes to model database for all SQL Server in your environment. Otherwise, a database created on one server may be different from a database created on another server.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Source:

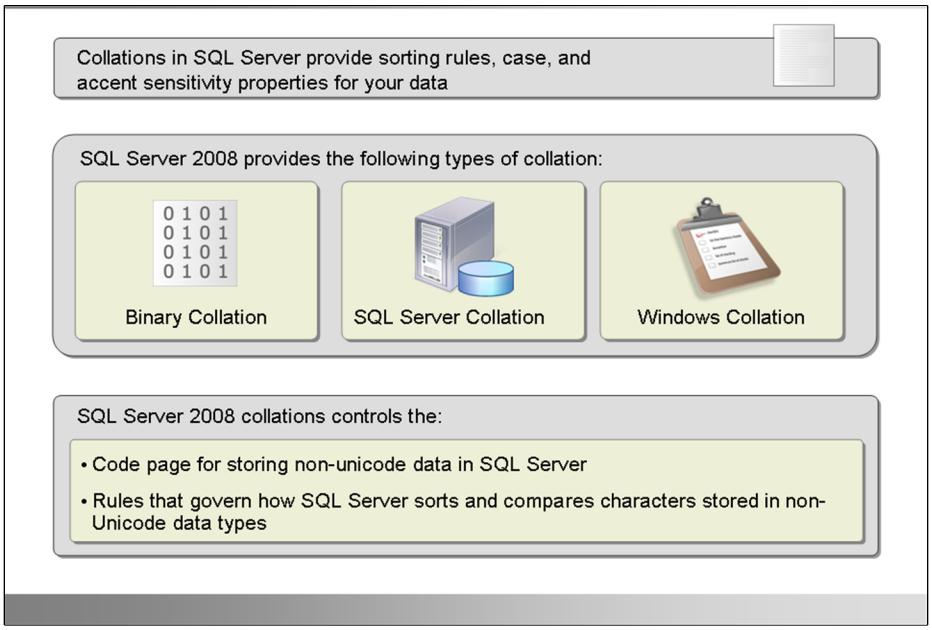
- http://social.msdn.microsoft.com/forums/en-US/sqldatabaseengine/thread/f7885b76-e66b-4dc7-9656-ea4cffa80468

For more information, see Moving System Databases.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

What Is Collation?



Key Points

Collation refers to a set of rules that determine how data is sorted and compared. Collations are a critical part of creating a database and manipulating data.

Types of Collations

SQL Server 2008 provides the following types of collation:

- **Binary collations.** Binary collations sort and compare data based on the bit pattern for each character. Each binary collation in SQL Server maps to a specific language locale and an American National Standards Institute (ANSI) code page. Binary collations perform case-sensitive and accent-sensitive data sorts.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

- **SQL Server collations.** SQL Server collations provide sort order compatibility with SQL Server 2000 and 7.0. SQL Server 2008 collation settings depend on the type of installation. You should select a SQL Server collation that supports the Windows system locale most commonly used in your organization.
- **Windows collations.** Windows collations define rules for storing character data based on an associated Windows locale. When you apply dictionary sorting, the base Windows collation rules specify the alphabet or language in use, and the code page to store non-unicode character data.

If you change a collation type, you may have to rebuild master database or application database. Incorrect collations will lead to incorrect ways of sorting and filtering the content. For example, filtering will be different for case-insensitive collation and case-sensitive collation. Also, there will be difficulties in moving data between servers with incompatible collations.

SQL Server 2008 collations control the:

- Code page for storing non-unicode data in SQL Server.
- Rules that govern how SQL Server sorts and compares characters stored in nonunicode data types. SQL Server Setup will detect the Windows collation on the computer where SQL Server is being installed. Then, it will automatically select the collation for your instance of SQL Server. The results of the sort and comparison might be different for Unicode and nonunicode characters.

SQL Server Collation Settings

SQL Server 2008 collation settings depend on the type of installation. Typically, you should choose a SQL Server collation that supports the Windows system locale most commonly used at your organization. To identify your Windows system locale:

- On computers that are running Windows 2000, click Regional Settings in Control Panel, and then click Advanced.
- On computers running Windows Vista, use the Formats tab.
- On computers that are running Windows Server 2003, click Regional and Language Options in Control Panel, and then click the Regional Options tab.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

In most cases, a computer runs the Windows system locale that matches the language requirements of the user. Therefore, SQL Server Setup automatically detects the Windows system locale and selects the appropriate SQL Server collation. For backward compatibility, the default English-language (US) collation is SQL_Latin1_General*."

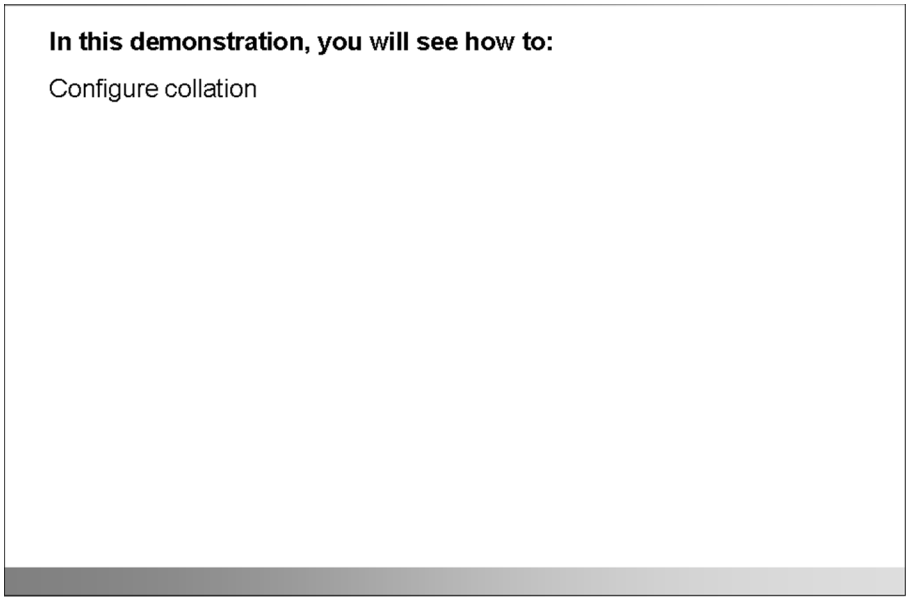
Source:

- http://msdn.microsoft.com/en-us/library/ms143503.aspx
- http://msdn.microsoft.com/en-us/library/ms144260.aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Demonstration: How To Configure Collation



Key Points

The steps to configure collation are:

1. In Windows Explorer, browse to the **D:\Demofiles\Mod02** folder, and then double-click the **M2_Demo3.sql** file.

   Note: If **Connect to Database Engine** dialog box appears, click **Connect**.

2. In the Microsoft SQL Server Management Studio window, in the Query Editor pane, select the **Modify the collation of the HumanResources.Department table to Danish_Norwegian_CS_AS** code, and then click **Execute**.

3. In the Microsoft SQL Server Management Studio window, in the Query Editor pane, select the **Modify the collation of the HumanResources.Department table to Latin1_General_CI_AS** code, and then click **Execute**.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Question: Why should you have different collation for a table or database?

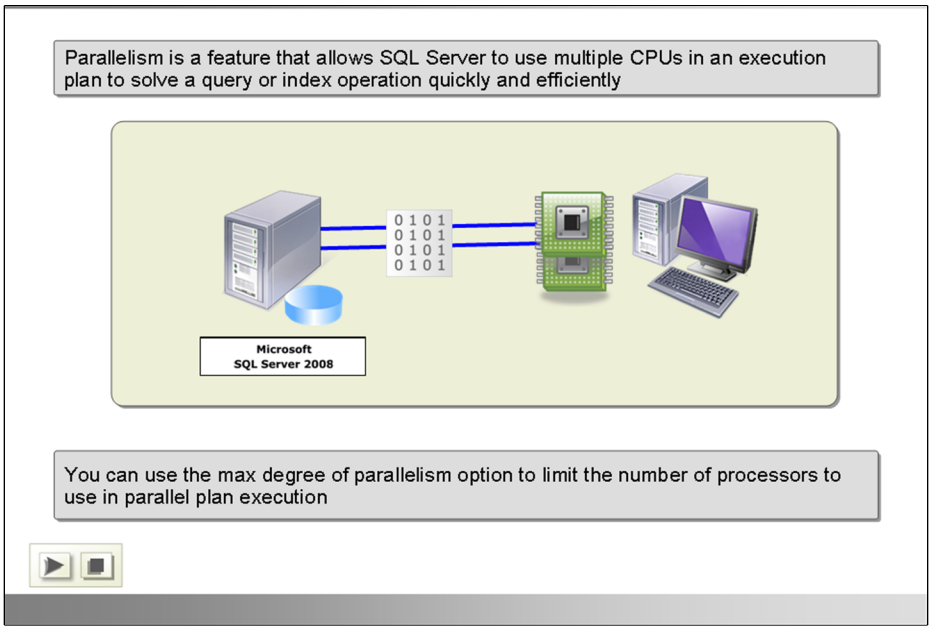
Question: What is the impact of changing the collation?

Question: Can you change the collation for a server?

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

What Is Parallelism?



Key Points

Using parallelism, SQL Server can use multiple CPUs in an execution plan to solve a query or index operation quickly and efficiently.

SQL Server automatically detects the best degree of parallelism for each instance of a parallel query execution or index data definition language (DDL) operation.

SQL server adopts the best degree of parallelism based on the following criteria:

- Is SQL Server running on a computer with more than one microprocessor or CPU, such as a SMP? Only computers with more than one CPU can use parallel queries.
- Are there sufficient threads available? Each query or index operation requires certain number of threads to run. Executing a parallel plan requires more threads than a serial plan. In addition, the number of required threads increases with the

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

degree of parallelism. The degree of parallelism also depends on the type of query or index operation executed.

- Is there sufficient number of rows available to process? If the query optimizer determines that the number of rows is too low, it does not introduce exchange operators to distribute the rows. Consequently, SQL Server executes the operators serially.
- Are there distribution statistics available? If the highest degree of parallelism is not possible, SQL Server considers lower degrees before abandoning the parallel plan.

Scenarios when parallel execution is not considered

- SQL Server query optimizer does not use a parallel execution plan for a query, if any of the following condition is true:
- The serial execution cost of the query is not high enough to consider an alternative, parallel execution plan.
- A serial execution plan is considered faster than any possible parallel execution plan for the particular query.
- The query contains scalar or relational operators that can not be run in parallel. Certain operators can cause a section of the query plan to run in the serial mode, or the whole plan to run in the serial mode.

max degree of parallelism server option

When SQL Server runs on a computer with more than one microprocessor or CPU, it detects the best degree of parallelism, that is, the number of processors employed to run a single statement, for each parallel plan execution.

You can use the max degree of parallelism option to limit the number of processors to use in parallel plan execution. The default value of 0 uses all available processors. Set the max degree of parallelism value to 1 to suppress parallel plan generation. Set the value to a number greater than 1 or up to a maximum of 64 to restrict the maximum number of processors used by a single query execution.

If you specify a value greater than the number of available processors, the actual number of available processors is used. If the computer has only one processor, the max degree of parallelism value is ignored.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Overriding max degree of parallelism with MAXDOP query hints

You can override the max degree of parallelism value in queries by specifying the MAXDOP query hint in the query. You can override the max degree of parallelism value for index operations by specifying the MAXDOP index option in the index statement.

In addition to queries and index operations, the max degree of parallelism option controls the parallelism of DBCC CHECKTABLE, DBCC CHECKDB, and DBCC CHECKFILEGROUP statements. You can disable parallel execution plans for these statements by using trace flag 2528.

Question: You have a query that is using up all system resources. How can max degree of parallelism configuration help to improve performance?

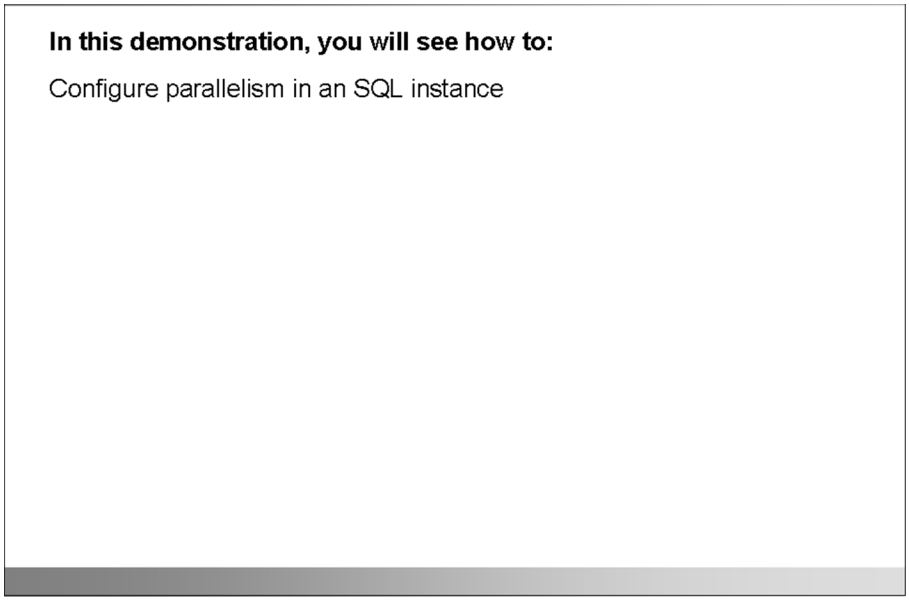
Source:

- http://msdn.microsoft.com/en-us/library/ms188611%28SQL.90%29.aspx
- http://msdn.microsoft.com/en-us/library/ms178065%28SQL.105%29.aspx
- http://msdn.microsoft.com/en-us/library/ms181007%28SQL.105%29.aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Demonstration: How To Configure Parallelism



Key Points

The steps to configure parallelism are:

1. In Windows Explorer, browse to the **D:\Demofiles\Mod02** folder, and then double-click the **M2_Demo4.sql** file.

Note: If **Connect to Database Engine** dialog box appears, click **Connect**.

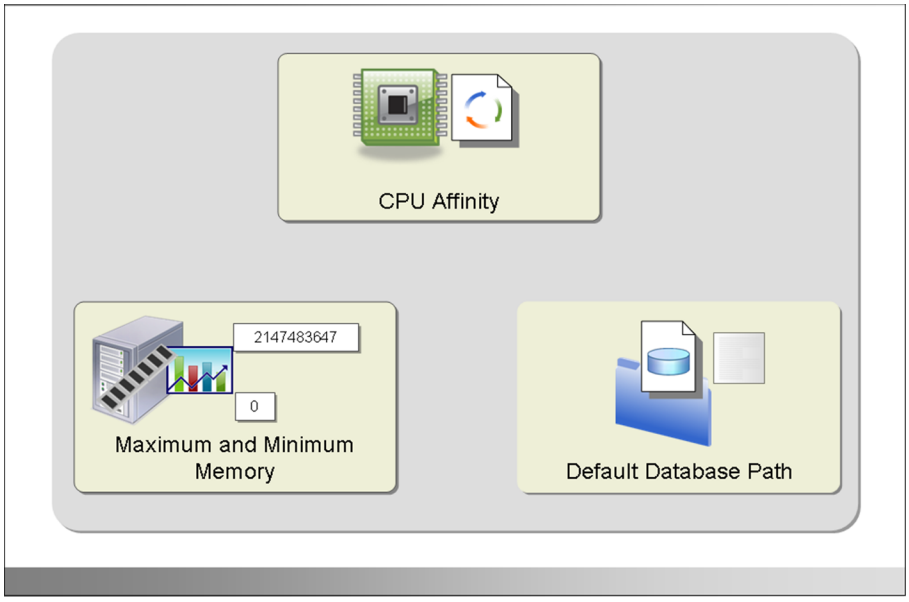
2. In the **Microsoft SQL Server Management Studio** window, click **Execute**.

Question: When would you consider changing the parallelism configuration?

Question: What is the best way to control the parallelism of a query?

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

System Configuration for an Instance



Key Points

To configure an instance of SQL Server 2008 successfully, you need to consider the following parameters:

- **CPU affinity.** During multitasking, Windows server distributes the process threads among various processors. This reduces SQL Server performance under heavy system loads. When you assign processors to specific threads, you can improve the processor performance. SQL Server supports processor affinity by providing two affinity masks options: CPU affinity and affinity I/O. The CPU and I/O affinities support for servers with 33 to 64 processors requires the additional use of the affinity 64 mask and affinity 64 I/O mask options, respectively.

  To configure the affinity mask option:

  1. Open the Microsoft SQL Server Management Studio window.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

2. In the **Microsoft SQL Server Management Studio** window, in the Query Editor pane, delete the existing code and type the following code, and then click **Execute**.

```
sp_configure 'show advanced options', 1;
RECONFIGURE;
GO
sp_configure 'affinity mask', 1;
RECONFIGURE;
GO
```

- **Maximum and minimum memory.** You can use the min server memory and max server memory options to reconfigure the amount of memory in the buffer pool used by an SQL Server instance. The default value for the min server memory option is 0 and for the max server memory option is 2,147,483,647. The minimum amount of memory you can specify for the max server memory option is 16 megabytes (MB).

  To set a fixed amount of memory:

  1. Open the Microsoft SQL Server Management Studio window.
  2. In the Microsoft SQL Server Management Studio window, in the Object Explorer, right-click the SQL Server instance, and then click **Properties**. The **Server Properties –** window appears.
  3. In the **Server Properties –** window, under **Select a page**, click **Memory**.
  4. On the **Memory** page, under **Server Memory Options**, in the **Minimum server memory (in MB)** list, type or select the minimum amount of memory for the server.
  5. On the **Memory** page, under **Server Memory Options**,, in the **Maximum server memory (in MB)** list, type or select the maximum amount of memory for the server, and then click **OK**.

- **Default database path**. You can change the default data and log directories by using SSMS. You can find the default backup directory in the registry.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

To view or change the default locations for database files:

1. Open the Microsoft SQL Server Management Studio window.
2. In the Microsoft SQL Server Management Studio window, in the Object Explorer, right-click the SQL Server instance, and then click **Properties**. The **Server Properties –** window appears.
3. In the **Server Properties –** window, under **Select a page**, click **Database Settings**.
4. Under **Database default locations**, view the current default locations for new data files and new log files. To change a default location, specify a new default pathname in the **Data** box or the **Log** box. Alternatively, you can specify the default path for the files by clicking the **ellipsis (...)** button to find and select a pathname, and then click **OK**.

Why Use CPU Affinity Mask

By segregating SQL Server threads from running on particular processors, Microsoft Windows 2000 or Windows Server 2003 can better evaluate the handling of processes specific to Windows. For example, on an 8-CPU server running two instances of SQL Server (instance A and B), the system administrator could use the affinity mask option to assign the first set of 4 CPUs to instance A and the second set of 4 CPUs to instance B. To configure more than 32 processors, set both the affinity mask and the affinity64 mask. The values for affinity mask are as follows:

- A one-byte affinity mask covers up to 8 CPUs in a multiprocessor computer.
- A two-byte affinity mask covers up to 16 CPUs in a multiprocessor computer.
- A three-byte affinity mask covers up to 24 CPUs in a multiprocessor computer.
- A four-byte affinity mask covers up to 32 CPUs in a multiprocessor computer.
- To cover more than 32 CPUs, configure a four-byte affinity mask for the first 32 CPUs and up to a four-byte affinity64 mask for the remaining CPUs.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Because setting SQL Server processor affinity is a specialized operation, you must use it only when necessary. In most cases, the Microsoft Windows 2000 or Windows Server 2003 default affinity provides the best performance. You should also consider the CPU requirements for other applications when setting the affinity masks.

Question: When should you consider changing affinity settings?

Question: Why should you consider changing the default database folder?

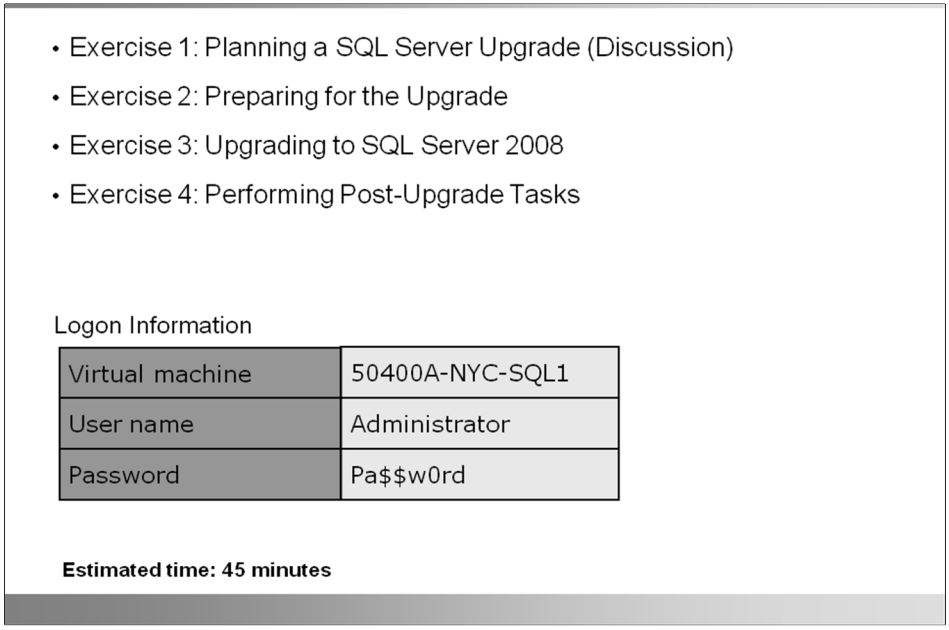
Source:

- http://msdn.microsoft.com/en-us/library/ms187104%28SQL.105%29.aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Lab 2: Deploying SQL Server 2008



Objectives:

After completing this lab, you will be able to:

- Planning a SQL Server upgrade.
- Preparing for the upgrade.
- Upgrade to SQL Server 2008.
- Performing post-upgrade tasks.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Scenario

QuantamCorp has deployed two SQL Server 2008 servers in their organization and will continue to deploy more that will replace their existing installation of SQL Server 2005. After this, SQL Server 2008 will be the primary database engine used in their organization. Now, you are going to perform an upgrade of an existing installation ofSQL Server 2005 in such a way that the application that uses the databases hosted on the SQL Server 2005 server would experience minimum downtime.

You are a part of the project team which is involved in planning and performing the upgrade. As required by the management, you need to achieve the following:

- Decide the best upgrade strategy to use for a scenario.
- Decide how to perform the upgrade with minimum risk and downtime.
- Decide whether consolidation between servers has to be performed and the rationale for doing so.
- Check if an in-place upgrade of a SQL Server 2005 server is possible by using Upgrade Advisor.
- Perform an in-place upgrade on a designated server.
- Perform a post-upgrade check and tasks to ensure that the upgrade is successful.

Exercise 1: Planning a SQL Server Upgrade (Discussion)

Scenario

QuantamCorp has deployed SQL Server as its primary SQL Server database across the enterprise. The following are the distribution of SQL Servers across different departments.

Number of SQL Servers	SQL Server 7.0	SQL Server 2000	SQL Server 2005
IT	0	5	10
Account	5	0	0

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Number of SQL Servers	SQL Server 7.0	SQL Server 2000	SQL Server 2005
HR	1	1	1
Sales	0	0	20
Customer Service	4	10	10

These departments use database servers to support their line-of-business (LOB) applications. The customer services team uses a proprietary custom application and the database is hosted on these SQL servers. The management requires minimizing their downtime during upgrade as well as minimizing the cost associated with the upgrade.

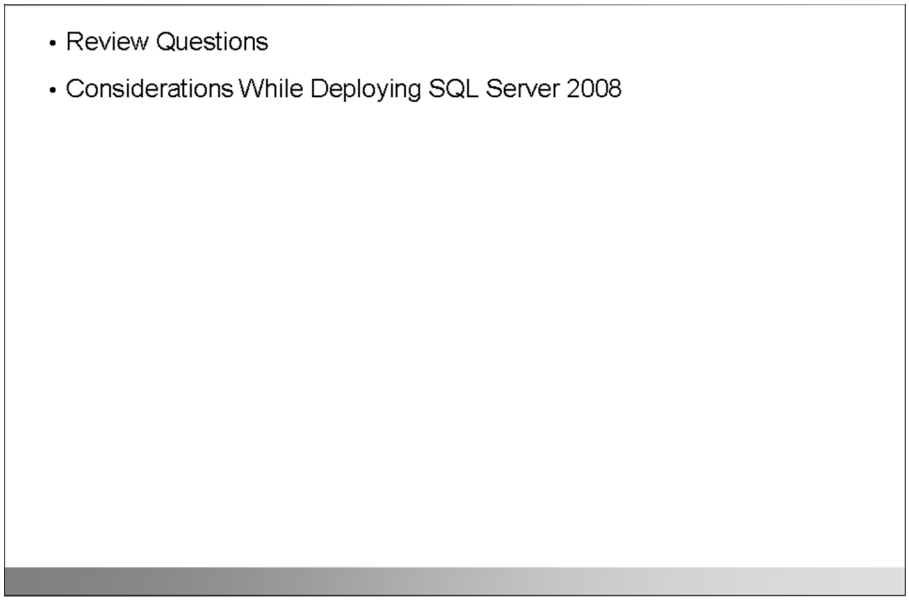
Based on the above scenario, answer the following questions:

1. Which upgrade strategy should you consider while upgrading SQL Servers in the Account department?
2. What should you consider while planning the upgrade strategy for the customer service team?
3. In which order should you upgrade the departments?
4. How can you ensure minimum risk and downtime during upgrade?

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Module Review and Takeaways



Review Questions

1. You are assigned the job of deploying SQL Server 2008. You are asked to deploy 32 bit version instead of 64 bit version. But you want to convince your client to use 64 bit instead. What are the benefits of using a 64 bit version?
2. Would there be any possible compatibility issue when using 64 bit instead of 32 bit?
3. When would you use SQL Server Upgrade Advisor?
4. What should you consider when selecting collation setting for your server? Can you change the collation that is selected after SQL Server setup is done?
5. What is the configuration that you would usually configure after you have deployed SQL Server?
6. When you would consider changing the parallelism configuration?

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Considerations While Deploying SQL Server 2008

- Upgrade to SQL Server 2008 from SQL Server 2000 and SQL Server 2005.
- Perform consolidation during upgrade as new hardware is usually more capable to handle more requests.
- Use the Upgrade Advisor which helps to identify potential issues associated with upgrading the server with in-place upgrade.
- Use Side-by-side upgrade as it allows for easier fallback if the upgrade fails but it requires additional hardware.
- Protect the system databases which are important. Each system database serves a different purpose. For example, the master database is for storing system configuration and the msdb database is for storing SQL Server Agent configuration
- Choose the collation setting carefully during SQL Server setup as you cannot change unless you reinstall it. But you can have different collation at database/table level of the server level
- Change the affinity settings to control how SQL Server consume CPU or I/O resources
- Change the parallelism behavior if you are having problems with queries consuming too many resources.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Module 3

Designing the Physical Structure of SQL Server 2008

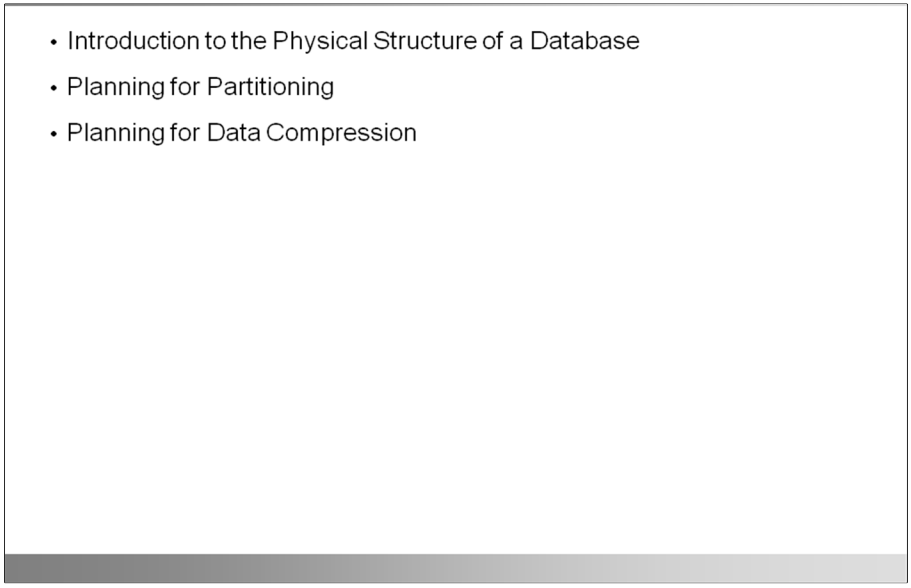
Contents:

Lesson 1: Planning for Database Files and Transaction Logs 3-3
Lesson 2: Planning for Partitioning 3-28
Lesson 3: Designing a Strategy for Data Compression 3-48
Lab 3: Designing the Physical Structure of SQL Server 2008 3-62
Module Review and Takeaways 3-74

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Module Overview



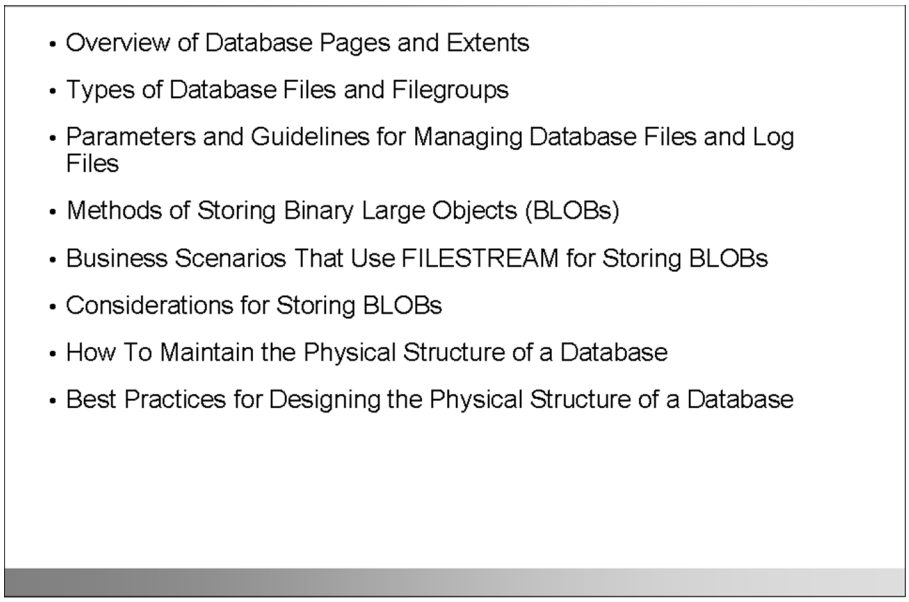
In SQL Server™ 2008, you need to plan the physical structure of the database to ensure that the databases are running in an optimal state. This will help you improve the scalability, performance, and administration of the database. You may also need to configure full-text indexing to speed up some of the queries used in applications.

In this module, you will learn about the physical file structure of the database including partitioned tables and indexes. In addition, you will learn about the full-text indexing feature and how to configure it.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Lesson 1
Introduction to the Physical Structure of a Database



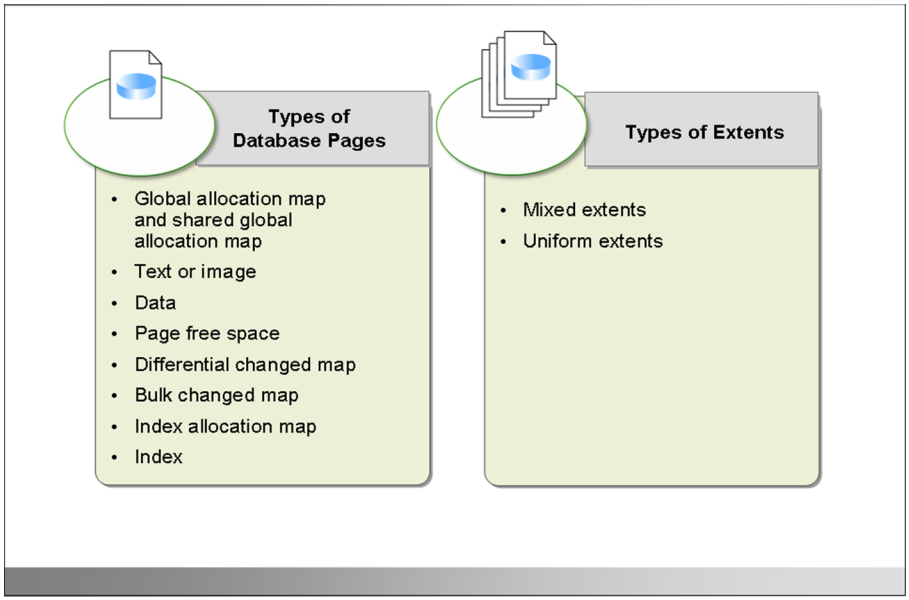
You can design databases in different ways with optimal performance. To get the most out of your database system design, you should understand how database files and log files are organized in SQL Server 2008. Database files consists of user defined data and the database engine consists of a default data file. You can use log files to store any modifications made to data. You can add additional data files for storage. You will need to plan the storage of the database files and transaction logs while deploying SQL Server 2008 to ensure that enough space is available for the database.

In this lesson, you will learn about the fundamentals of the physical structure of the database, with regard to database files, filegroups, and storing binary large objects (BLOBs).In addition, you will learn about parameters and considerations for database files. You will also learn how to maintain the physical structure of the database and the best practices for designing a physical structure of the database.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Overview of Database Pages and Extents



Key Points

Tables and indexes consist of database pages. You can allocate database pages to tables and indexes through extents. You need to understand the architecture of pages and extents for designing and developing databases.

Database Pages

A page is the fundamental unit of data storage in SQL Server. The disk space allocated to a data file (.mdf or .ndf) in a database is logically divided into pages that are numbered contiguously from 0 to n. Disk I/O operations are performed at the page level.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

The various types of pages used in the data files of a SQL Server database include:

- **Global allocation map and shared global allocation map**. The global allocation map and the shared global allocation map store information about extent allocation.
- **Text or image**. The text or image page stores large object data types such as text, ntext, image, nvarchar(max), varchar(max), varbinary(max), and XML. When the size of the data row exceeds 8 KB, the data object uses variable column lengths such as varchar, nvarchar, varbinary, and sql_variant.
- **Data**. The data page stores data rows for all data types except text, ntext, image, nvarchar(max), varchar(max), varbinary(max), and XML.
- **Page free space**. The page free space page contains information about the free space available on pages and page allocations.
- **Differential changed map**. The differential changed map page stores information about changed extents since the last BACKUP DATABASE statement was executed for each allocation unit.
- **Bulk changed map**. The bulk changed map page stores information about extents that you can modify by bulk operations since the last BACKUP LOG statement that was run for each allocation unit.
- **Index allocation map**. The index allocation map page includes information about the extents used by a table or an index per allocation unit.
- **Index**. The index page stores index entries.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Structure of SQL Server Data Page

You can arrange data rows on a page serially, starting immediately after the header. A row offset table starts at the end of the page, and each row offset table contains one entry for each row on the page. Each entry records how far the first byte of the row is from the start of the page. The entries in the row offset table are in reverse sequence from the sequence of the rows on the page.

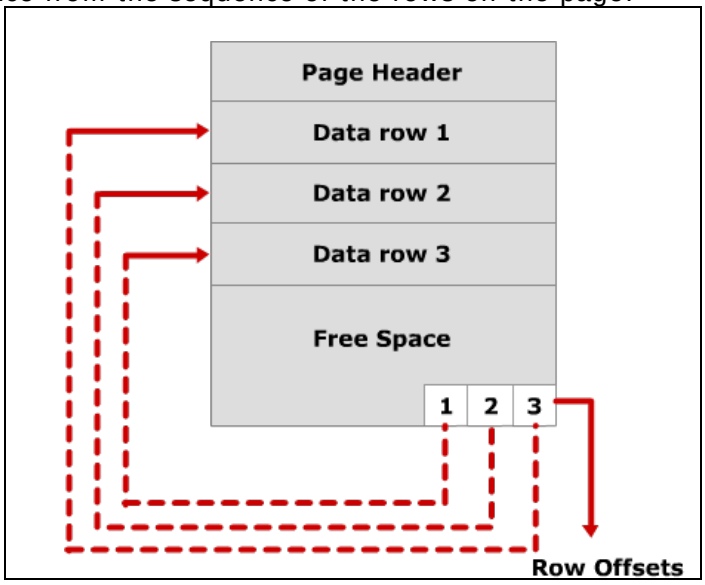


Figure 4.1. Structure of SQL Server Page

Rows cannot span across pages, however you may move out portions of the row from the page so that the row can actually be very large. The maximum amount of data and overhead that a single row on a page contains is 8,060 bytes or 8 KB. However, this amount does not include the data stored in the Text/Image page type. The maximum limit for tables that contains varchar, nvarchar, varbinary, or sql_variant columns varies. When the total row size of all fixed and variable columns in a table exceeds the 8 KB limitation, SQL Server dynamically moves one or more variable length columns to pages in the ROW_OVERFLOW_DATA allocation unit, starting with the column that has the largest width. This movement occurs whenever an INSERT or UPDATE operation increases the total size of

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

the row beyond 8 KB. When SQL Server moves a column to a page in the ROW_OVERFLOW_DATA allocation unit, it maintains a 24-byte pointer on the original page in the IN_ROW_DATA allocation unit. If a subsequent operation reduces the row size, SQL Server dynamically moves the columns back to the original data page.

Extents

Extents are the basic units in which SQL Server manages space. An extent is a collection of eight physically contiguous pages. You can efficiently manage pages by using extents. All pages are stored in extents. To make the space allocation of extents efficient, SQL Server does not allocate whole extents to tables with small amounts of data. SQL Server has two types of extents:

- **Mixed extents**. A maximum of eight objects can share a mixed extent. Each of the eight extent pages can be possessed by separate objects.
- **Uniform extents**. A single object can own all uniform extents. However, only the owning object can access all eight pages of extents.

SQL Server generally allocates pages from mixed extents to a new table or index. When the table or index grows to the point that it has eight pages, it switches to use uniform extents for subsequent allocations. If you create an index on an existing table that has enough rows to generate eight pages in the index, all allocations to the index are in uniform extents.

Question: What is the relationship between a data page and an extent?

Source:

- http://msdn.microsoft.com/en-us/library/ms190969.aspx

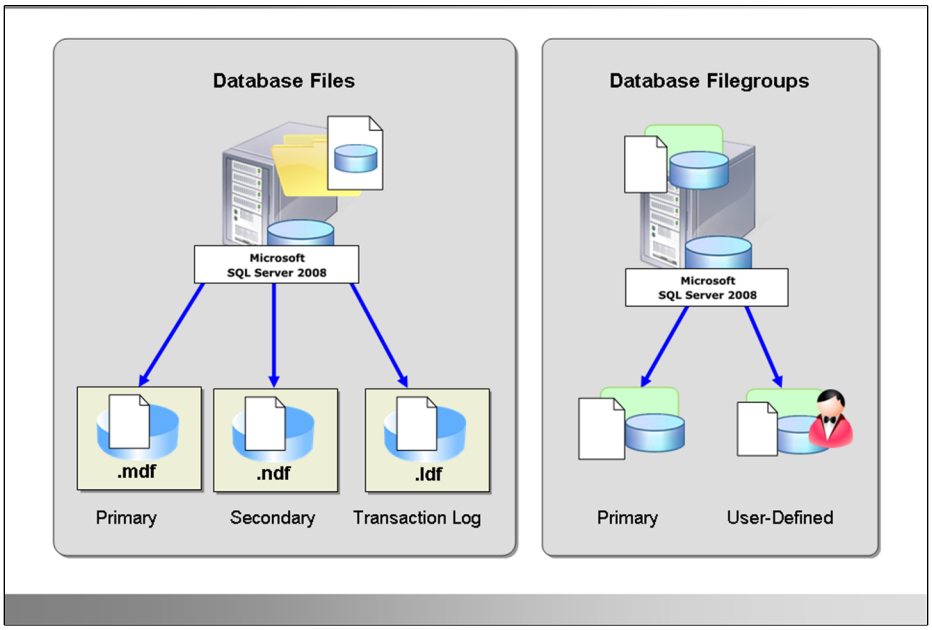
For More Information, see Understanding Pages and Extents.

For More Information, see Managing Extent Allocation and Free Space.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Types of Database Files and Filegroups



Key Points

A SQL Server database contains at least two files, one data file and one log file. Filegroups are named collections of files that can be used for data placement and for administrative tasks such as backup and restore operations. You can group database objects and files together in filegroups.

Database Files

SQL Server databases have three types of files:

- **Primary data files**. The primary data file is the starting point of the database and points to the other files in the database. Every database has one primary data file which usually has the extension .mdf.
- **Secondary data files**. Secondary data files refer to all data files, other than the primary data file. The secondary data files usually have the extension .ndf.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

- **Transactional log files**. Log files contain all the log information that is used to recover a database. A database must have at least one log file for each database, although there can be more than one log file. The log files usually have the extension .ldf.

Note: SQL Server does not enforce the .mdf, .ndf, and .ldf file name extensions, but these extensions help you identify the different kinds of files and their use.

Database Filegroups

The two types of filegroups are:

- **Primary**. The primary filegroup contains the primary data file and any other files that are not assigned to another filegroup. All pages for the system tables are allocated to the primary filegroup.
- **User-defined**. You can specify user-defined filegroups by using the FILEGROUP keyword in a CREATE DATABASE or ALTER DATABASE statement.
  Log files are never part of a filegroup. You can manage log space separately from data space.

  A file cannot be a member of more than one filegroup. Tables, indexes, and large object data can be associated with a specified filegroup. In this case, all their pages will be allocated to that filegroup, or the tables and indexes can be partitioned. The data of partitioned tables and indexes is divided into units, each of which can be placed in a separate filegroup in a database.

  You can designate one filegroup in each database as the default filegroup. When you create a table or index without specifying a filegroup, it is assumed all pages will be allocated from the default filegroup. Only one filegroup at a time can be the default filegroup. Members of the db_owner fixed database role can switch the default filegroup from one filegroup to another. If no default filegroup is specified, the primary filegroup is the default filegroup.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Question: What is the purpose of a primary filegroup?

Source:

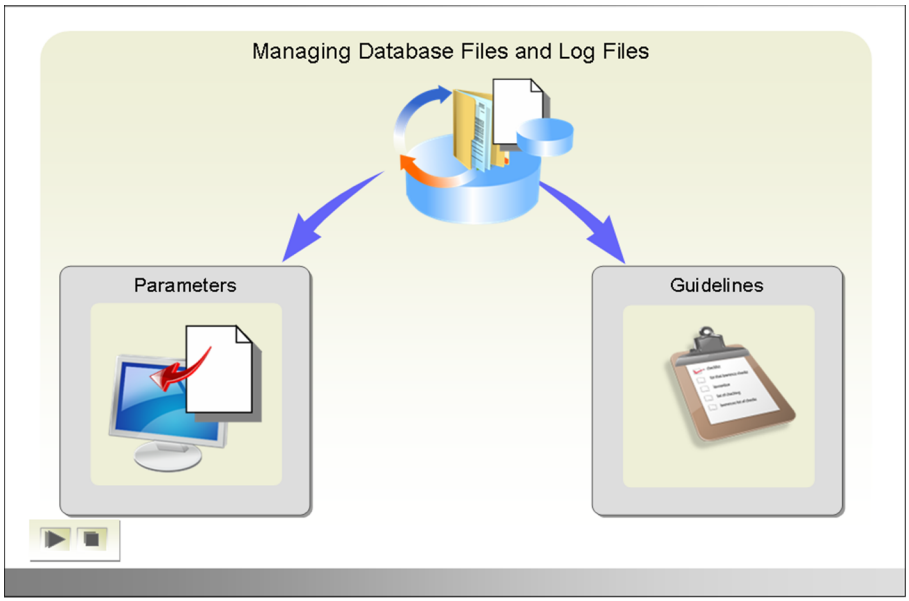
- msdn.microsoft.com/en-us/library/ms179316.aspx
- msdn.microsoft.com/en-us/library/aa174545(SQL.80).aspx

For more information, see Database Files and Filegroups.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Parameters and Guidelines for Managing Database Files and Log Files



Key Points

SQL Server stores the database content into one or more database files. Each database consists at least one database file and log file.

Parameters for Managing Database Files and Log Files

The parameters that a user can specify while managing a database file and a log file include:

- **Current size**. This is the predetermined disk space allocated to a database file. You can specify the current size by using the *size* command. You cannot apply this command to a FILESTREAM filegroup.
- **Maximum size**. This is the upper limit that SQL Server allows to allocate to a database file. You can specify the maximum size by using the *max_size* command. You cannot apply this command to a FILESTREAM filegroup.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

- **Autogrowth size**. This is the disk size that SQL Server extends for a database file. However, consider using a fixed size for auto-growth. If you set a very large size for autogrowth value, it affects the performance of the database at the time when SQL Server is adding additional space to the database file. You can specify the autogrowth value by using the *growth_increment* command. The auto-growth value that you specify cannot exceed the maximum size value. You cannot apply the auto-growth size command to a FILESTREAM filegroup.

Guidelines for Managing Database Files and Log Files

While managing database files and log files, consider the following:

- **Use redundant array of inexpensive disks (RAID) for database files**. You should use RAID 0, 5 or 10 for database files because it provides faster disk read speed and redundancy. Though RAID 0 provides faster disk read speed, it does not provide redundancy. The type of RAID to use depends on the support of hardware and number of disks included in the system.
- **Store log files in RAID 1 drive**. You need to store log files in RAID 1 because it has improved write performance and provides complete redundancy. You cannot ignore the tempdb database file because it grows extensively. You must plan the file structure of the tempdb database file just like the other databases.
- **Specify the size of the database file while creating the file**. When you create a database, you must either specify an initial size for the data and log files, or accept the default size. As data is added to the database, data will be saved to these files. However, you must consider how the database will grow beyond the initial space by specifying the autogrowth amount and a maximum file size.
  By default, the data files grow as much as required until there is no disk space available. Therefore, if you do not want the database files to grow any larger than when they were first created, you must specify this at the time when you create the database by using SQL Server Management Studio (SSMS) or the CREATE DATABASE statement.
  Alternatively, SQL Server lets you create data files that can grow automatically when they fill with data, but only to a predefined maximum size. This can prevent the disk drives from running out of disk space completely.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

- **Specify a limit on the growth of data files**. Permit the data files to grow automatically, but put a limit on the growth by specifying a maximum data file growth size that leaves some available space on the hard disk. Specifying a limit on the data files permits database growth if more data is added than what is expected, but does not fill up the disk drive.
- **Re-evaluate the expected maximum database size**. If the initial data file size is exceeded and the file starts to grow automatically, alter the file size properties by adding more disk space, creating and adding more files or filegroups to the database. Set the maximum size of the database to the same value as the initial size and specify a filegrowth of zero, if the database is not supposed to expand beyond its initial size. This limit prevents the database files from growing. If the database files fill with data, you cannot add more data until more data files are added to the database or until the existing files are expanded manually using the ALTER DATABASE statement.
- **Create the files or filegroups on different local physical disks**. Letting files grow automatically can cause fragmentation of those files, if several of them share the same disk. Therefore, you should create the files or filegroups on as many different local physical disks as you can. Also put objects that compete heavily for space in different filegroups.

  In the example, you create a database with two filegroups, primary and mydb_fg1. mydb_fg1 has 2 files and is assigned as the default filegroup. Each database file is 1 MB in size and the maximum size equals to 10 MB with auto-growth supported. Each auto-growth will increase the database file by 1 MB.

```
USE master;
GO
-- Create the database with the default data
-- filegroup and a log file. Specify the
-- growth increment and the max size for the
-- primary data file.
CREATE DATABASE MyDB
ON PRIMARY
```

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.


```
  ( NAME='MyDB_Primary',
    FILENAME=
       'D:\Data\MyDB_Prm.mdf',
    SIZE=4MB,
    MAXSIZE=10MB,
    FILEGROWTH=1MB),
FILEGROUP MyDB_FG1
  ( NAME = 'MyDB_FG1_Dat1',
    FILENAME =
       'D:\Data\MyDB_FG1_1.ndf',
    SIZE = 1MB,
    MAXSIZE=10MB,
    FILEGROWTH=1MB),
  ( NAME = 'MyDB_FG1_Dat2',
    FILENAME =
       'D:\Data\MyDB_FG1_2.ndf',
    SIZE = 1MB,
    MAXSIZE=10MB,
    FILEGROWTH=1MB)
LOG ON
  ( NAME='MyDB_log',
    FILENAME =
       'D:\Data\MyDB.ldf',
    SIZE=1MB,
    MAXSIZE=10MB,
    FILEGROWTH=1MB);
GO
ALTER DATABASE MyDB
```

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.


```
  MODIFY FILEGROUP MyDB_FG1 DEFAULT;
GO

-- Create a table in the user-defined filegroup.
USE MyDB;
CREATE TABLE MyTable
  ( cola int PRIMARY KEY,
    colb char(8) )
ON MyDB_FG1;
GO
```

Question: What should you consider when configuring the autogrowth setting in SQL Server?

Source:

- msdn.microsoft.com/en-us/library/ms190970.aspx
- msdn.microsoft.com/en-us/library/ms175198.aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Methods of Storing BLOBs

Comparison	Storage Method		
	File server or File System	In-Row SQL Server	FILESTREAM
Maximum BLOB size	NTFS volume size	2 gigabyte (GB) – 1 byte	NTFS volume size
Streaming performance of large BLOBs	Excellent	Poor	Excellent
Security	Manual access control lists (ACLs)	Integrated	Integrated and automatic ACLs
Cost per GB	Low	High	Low
Manageability	Difficult	Integrated	Integrated
Integration with structured data	Difficult	Data-level consistency	Data-level consistency
Application development and deployment	More complex	More simple	More simple
Recovery from data fragmentation	Excellent	Poor	Excellent
Performance of frequent small updates	Excellent	Moderate	Poor

Key Points

SQL Server provides a special data type for storing large amount of data, such as documents, raw files, and XML documents, in a table, known as BLOB. You can write a BLOB to a database, either as binary data or as character data. The various methods to store BLOBs include:

- **File Server or File System**. In this method, the data is stored on the disk directly and only a file name is stored in the database. You can save the BLOB data at file system and create a custom mapping, such as Customer ID as a key for the file name, which allow applications to read the saved BLOB object from file system. However, this method requires manual work for housekeeping of the unused files and difficult to update when a database record, such as the ID, is updated.
- **In-row SQL Server**. In this method, the data is stored inside a SQL database file. SQL Server 2008 supports storing BLOB data with text, ntext, image, varbinary(max) data type.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

- **FILESTREAM**. It is a new feature in the SQL Server 2008 release. It allows structured data to be stored in the database, and associated unstructured or BLOB data to be stored directly in the NTFS file system. It combines the benefits of storing BLOB in file system, for easier management, and in-row SQL server BLOB storage, for simplified application development.
  There are several benefits to FILESTREAM storage:
  - Performance matches the streaming performance of the file system.
  - The size of the BLOB is limited only by the file system volume size.
  - You do not need to manage the FILESTREAM columns separately. They are fully integrated with maintenance operations such as backup and restore.
  - SQL Server ensures data-level consistency between the relational data in the database and the unstructured data that is physically stored in the file system by providing full transactional support.
  - FILESTREAM data is completely integrated with the SQL Server security model.

Question: Where are the BLOB objects stored when you configure FILESTREAM?

Question: If you have a 25MB file to store in each row, which storage option should you use?

Source:

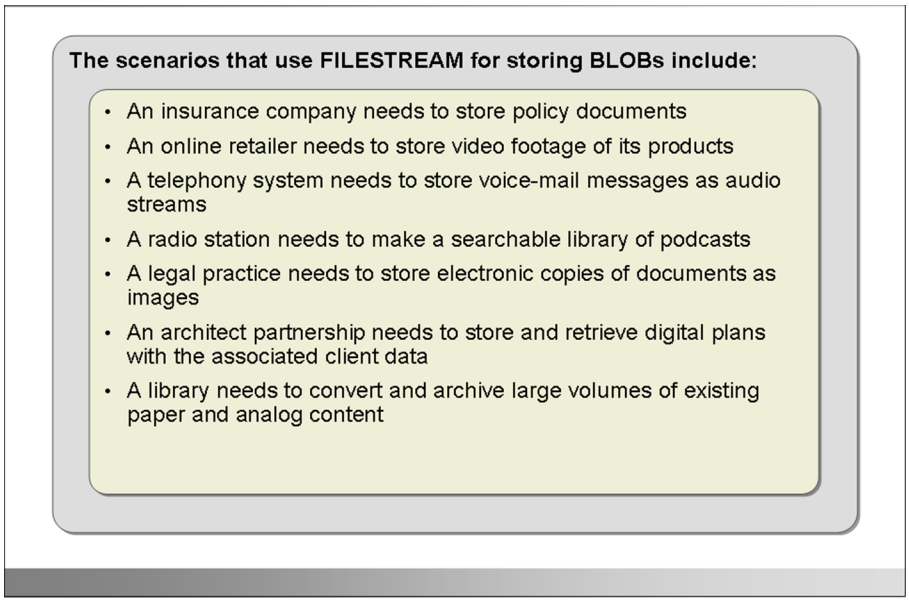
- msdn.microsoft.com/en-us/library/cc949109.aspx
- msdn.microsoft.com/en-us/library/cc716724.aspx

For more information, see FILESTREAM storage in SQL Server 2008.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Business Scenarios That Use FILESTREAM for Storing BLOBs



Key Points

Businesses that contain large volumes of non-relational, non-structured data in the form of digital images, documents, videos, or other multimedia formats can be considered for FILESTREAM data storage.

The following scenarios can use the FILESTREAM functionality to store BLOBs:

- An insurance company needs to store policy documents and retrieve them for claims processing.
- An online retailer needs to store video footage of its products for display in an e-commerce site.
- A telephony system needs to store voice-mail messages as audio streams so that they can be retrieved remotely.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

- A radio station needs to make a searchable library of podcasts available for download from its Web site.
- A legal practice needs to store electronic copies of documents as images and easily retrieve the documents relating to an individual client or case.
- An architect partnership needs to store and retrieve digital plans with the associated client data.
- A library needs to convert and archive large volumes of existing paper and analog content for indexing and use in a digital research tool.

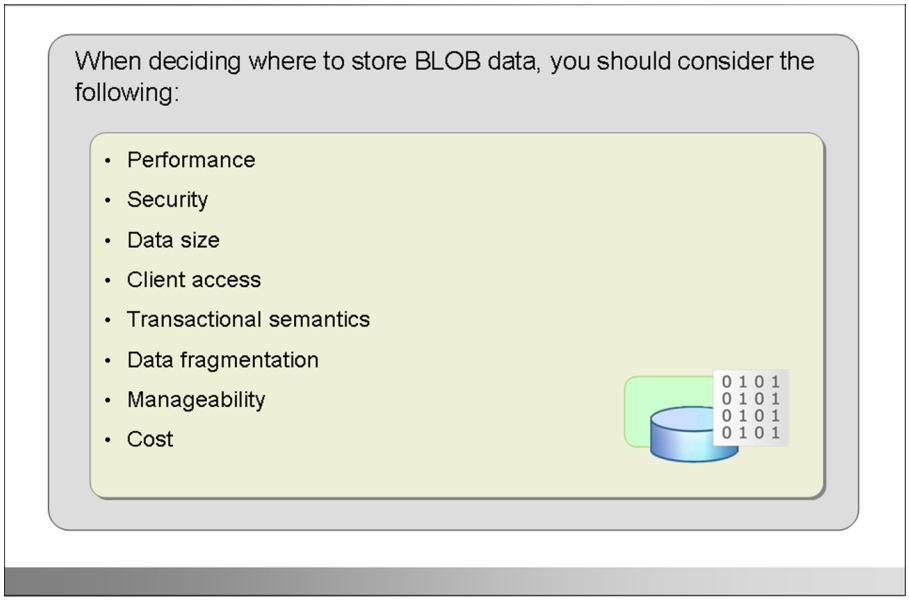
Source:

- msdn.microsoft.com/en-us/library/cc716724.aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Considerations for Storing BLOBs



Key Points

You can store structured and semi-structured data in a relational database. You need to consider the following requirements when deciding where to store BLOB data:

- **Performance**. If you need streaming access, storing data inside a SQL Server database may be slower than storing it externally in a location such as the NTFS file system. By using file system storage, you can read the data from the file and pass it to the client application, either directly or with additional buffering. When you store BLOB in a SQL Server database, you first need to read the data into the memory of SQL Server that is the buffer pool and then you need to pass back the data through a client connection to the client application. This passing back of data means that the memory of SQL Server contains BLOB data, which can cause further performance problems for SQL Server operations.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

- **Security**. You can store sensitive data having tightly-managed access in a database and control its security by using the SQL Server access controls. If the same data is stored in the file system, you need to implement different methods of security such as access control lists (ACLs).
- **Data size**. You can store BLOBs inside or outside a database. You can store BLOBs smaller than 256 kilobytes (KB) such as widget icons inside a database. You can store BLOBs larger than 1 megabyte (MB) outside the database. For storing BLOBs sized between 256 KB and 1 MB, the storage solution depends on the read versus. write ratio of the data, and the rate of "replace". To store BLOB data within the database such as the varbinary(max) data type, you need 2 gigabytes (GB) per BLOB.
- **Client access**. You need to store data in the file system when the protocol that the client uses to access SQL Server data, such as Open Database Connectivity (ODBC), is not suited to applications such as streaming large video files.
- **Transactional semantic**. If the BLOB data has associated structured data that will be stored in the database, you need to ensure that the changes to the BLOB data adhere to transactional semantics. By doing so, the two sets of data will be synchronized. For example, if a transaction creates BLOB data and a row in a database table, but then rolls back, you need to roll back the creation of the BLOB data and the creation of the table row. However, this roll back can become very complex if the BLOB data is stored in the file system with no link to the database.
- **Data fragmentation**. Frequent updates and overwrites will cause the BLOBs to move, either within the SQL Server database files or within the file system, depending on where the data is stored. In this case, if the BLOBs are large, they may become fragmented and will not be stored in one contiguous part of the disk. You can resolve this fragmentation issue easily by using the file system, rather than by using SQL Server.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

- **Manageability**. A solution that uses multiple technologies that are not integrated will be more complex and costly to manage than an integrated solution.
- **Cost**. The cost of the storage solution will vary, depending on the technology you are using.

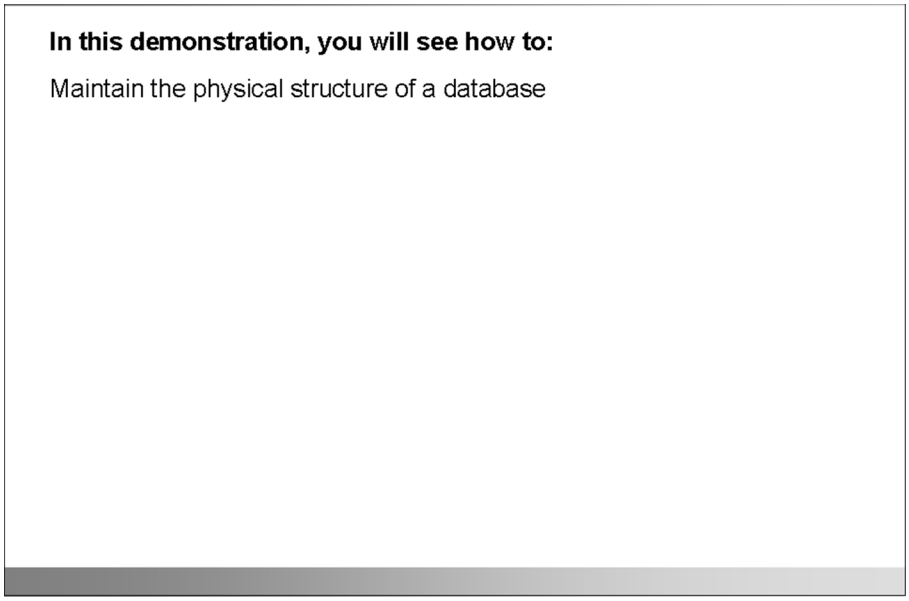
Source:

- msdn.microsoft.com/en-us/library/cc949109.aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Demonstration: How To Maintain the Physical Structure of a Database



Key Points

The steps to maintain the physical structure of a database are:

1. On the Start menu, point to **All Programs**, click **Microsoft SQL Server 2008**, and then click **SQL Server Management Studio**. The Microsoft SQL Server Management Studio window appears.
2. To log on to SQL Server Management Studio, in the **Connect to Server** dialog box, type the following:
   - Server type: **Database Engine**
   - Server name: **NYC-SQL1**
   - Authentication: **Windows Authentication**
3. In the **Connect to Server** dialog box, click **Connect.**
4. In Windows Explorer, browse to the **D:\Demofiles\Mod03** folder, and then double-click the **M3_Demo1.sql** file.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

5. In the Microsoft SQL Server Management Studio window, in the Query Editor pane, select the **check the database structure** code, and then click **Execute**.
6. In the Microsoft SQL Server Management Studio window, in the Query Editor pane, select the **shrink a file and a database** code, and then click **Execute**.

Question: What is the advantage of creating multiple database files in SQL Server?

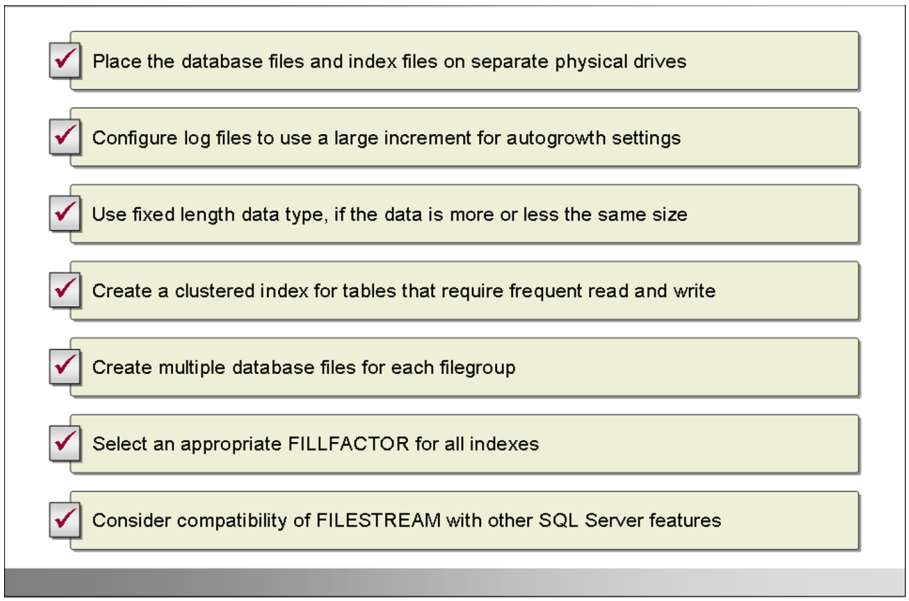
Question: Why do you need to change the default filegroup?

Question: What is the benefit of having multiple filegroups?

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Best Practices for Designing the Physical Structure of a Database



Key Points

To design the physical structure of a SQL Server 2008 database, you need to:

- **Place the database files and index files on separate physical drives**. This helps to reduce the chances for SQL Server to compete for the system resources when SQL Server needs to access both the table and the index at the same time.
- **Configure log files to use a large increment for autogrowth settings**. This helps you to avoid fragmentation on database files as small increments, like few MB, for log files would more likely to cause fragmentation problems on the file.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

- **Use fixed length data type, if the data is more or less the same size**. You should consider using fixed length data type instead of variable length data type, if the data is more or less the same size and similar to the size defined in the table. This helps to reduce the overhead of reading data.
- **Create a clustered index for tables that require frequent read and write**. Adding clustered index to tables helps organize the data in a more structured way which helps SQL Server to easily identify the data page to modify or to read.
- **Create multiple database files for each filegroup**. You should create multiple database files for each filegroup to split the input/output (I/O) process. This also enables parallel operations on the index to avoid contention.
- **Select an appropriate FILLFACTOR for all indexes.** A correctly chosen FILLFACTOR value can reduce potential page splits by providing enough space for index expansion as data is added to the underlying table. A page split makes room for new records, but can take time to perform and is a resource-intensive operation. Also, it can cause fragmentation that causes increased I/O operations. When frequent page splits occur, you can rebuild the index by using a new or existing FILLFACTOR value to redistribute the data.

  You need to specify different values for the FILLFACTOR setting. Different values have different impact on the performance of the table. Some values increase the efficiency of the I/O process while some values degrade the performance.
  The following table describes the FILLFACTOR value and its impact on I/O performance.

FILLFACTOR value	Impact on I/O performance
50-90	This value helps to reduce the I/O request caused by the operation. If you update existing rows with the data that lengthens the size of the rows, use a FILLFACTOR value between 50 -90. The extra byte on each page helps to minimize page splits caused by the extra length in the rows.
< 50	This value would cause database read performance

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

FILLFACTOR value	Impact on I/O performance
	to decrease by two times. Read performance is decreased because the index contains more pages, thereby increasing the disk I/O operations required to retrieve the data. You should avoid using this value as this might cause serious performance issues.
100 or 0	This value causes SQL Server to fully fill each page. Any change to the database would cause data page split to happen which would impact the write performance. But if no data would change happens, the read performance would be optimal as the number of pages to read is kept minimum.

- **Consider compatibility of FILESTREAM with other SQL Server features**. There are some features that will not function if a database has a FILESTREAM column, for instance, you will get an error when you set up snapshot on a database with a FILESTREAM column or you are trying to encrypt a FILESTREAM column.

Question: If you have a database table that requires frequent read-write access, what should you specify as the FILLFACTOR for this table?

Question: If you have a column that has on an average 10-12 characters of information stored, what is the data type you should use?

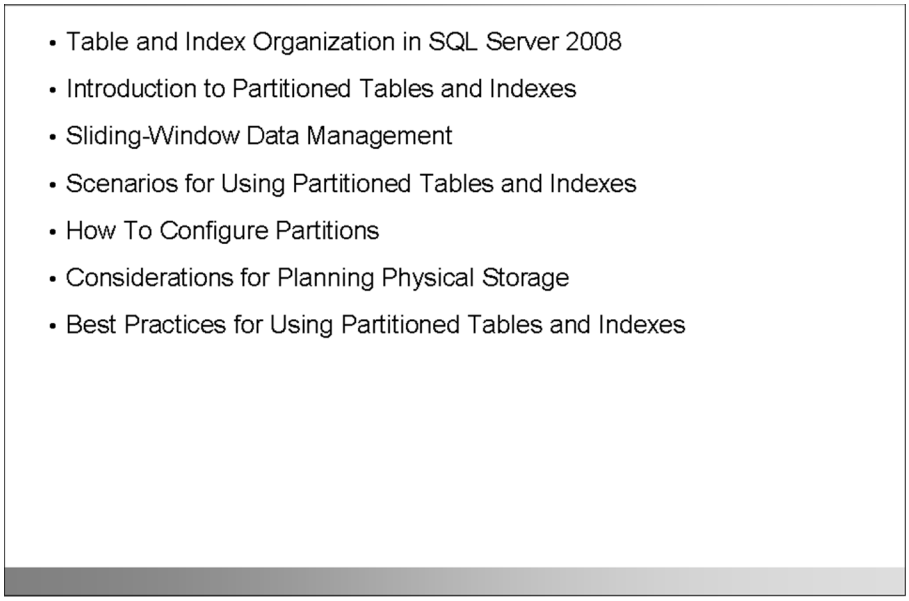
Source:

- msdn.microsoft.com/en-us/library/ms177459.aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Lesson 2
Planning for Partitioning



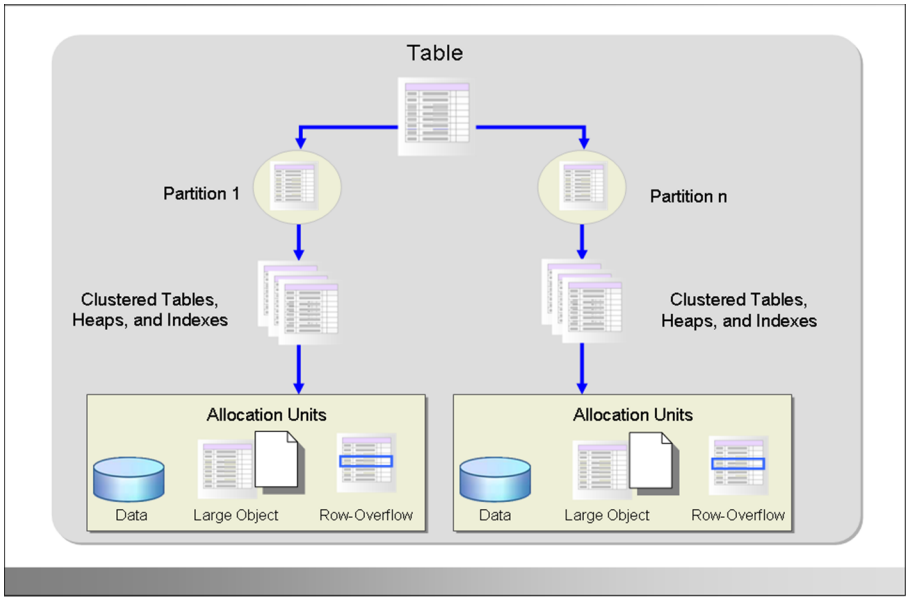
You can use the partitioning method to separate large tables and indexes in smaller parts. You can use SQL Server to partition tables based on the column and defined ranges. By partitioning a table or an index, you can move subsets of data quickly and efficiently. You can improve query performance by considering the type of queries and hardware configuration. You should consider the guidelines and best practices while implementing partitioned indexes. This leads to high performance gain.

In this lesson, you will learn how to improve the performance of database by changing the storage configuration and by table partitioning. You will learn about partitioned tables, indexes, and scenarios for using them. In addition, you will learn about the sliding-window concept. You will also learn about configuring partitions and planning physical storage for the partitioned data.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Table and Index Organization in SQL Server 2008



Key Points

Objects, like table, in the SQL Server database are stored as a collection of 8-KB pages. One or more partitions contain a table and each partition contains data rows in either a heap or a clustered index structure. You can manage the pages of the heap or clustered index in one or more allocation units, depending on the column types in the data rows.

Partitions

Table and index pages are contained in one or more partitions. A partition is a user-defined unit of data organization. By default, a table or index has only one partition that contains all the table or index pages. The partition resides in a single filegroup. A table or index with a single partition is equivalent to the organizational structure of tables and indexes in SQL Server 2000.

When a table or index uses multiple partitions, the data is partitioned horizontally so that groups of rows are mapped into individual partitions, based on a specified column. The partitions can be put on

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

one or more filegroups in the database. The table or index is treated as a single logical entity when queries or updates are performed on the data.

Clustered Tables, Heaps, and Indexes

Clustered tables are tables that have a clustered index. The data rows are stored based on the clustered index key. The clustered index is implemented as a B-tree index structure that supports fast retrieval of the rows, based on their clustered index key values. The pages in each level of the index, including the data pages in the leaf level, are linked in a doubly-linked list. However, navigation from one level to another is performed by using key values.

Heaps are tables that do not have a clustered index. The data rows are not stored in any particular order, and there is no particular order to the sequence of the data pages. The data pages are not linked in a linked list.

Indexed views have the same storage structure as clustered tables. When a heap or a clustered table has multiple partitions, each partition has a heap or B-tree structure that contains the group of rows for that specific partition. For example, if a clustered table has four partitions, there are four B-trees; one in each partition.

Allocation Units

An allocation unit is a collection of pages within a heap or B-tree used to manage data based on their page type. The following table lists the types of allocation units used to manage data in tables and indexes.

Allocation unit type	Is used to manage
IN_ROW_DATA	Data or index rows that contain all data, except large object (LOB) data. Pages are of type Data or Index.
LOB_DATA	Large object data stored in one or more of these data types: text, ntext, image, xml, varchar(max), nvarchar(max), varbinary(max), or CLR user-defined types (CLR UDT). Pages are of type Text/Image.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Allocation unit type	Is used to manage
ROW_OVERFLOW_DATA	Variable length data stored in varchar, nvarchar, varbinary, or sql_variant columns that exceed the 8,060 byte row size limit. Pages are of type Text/Image.

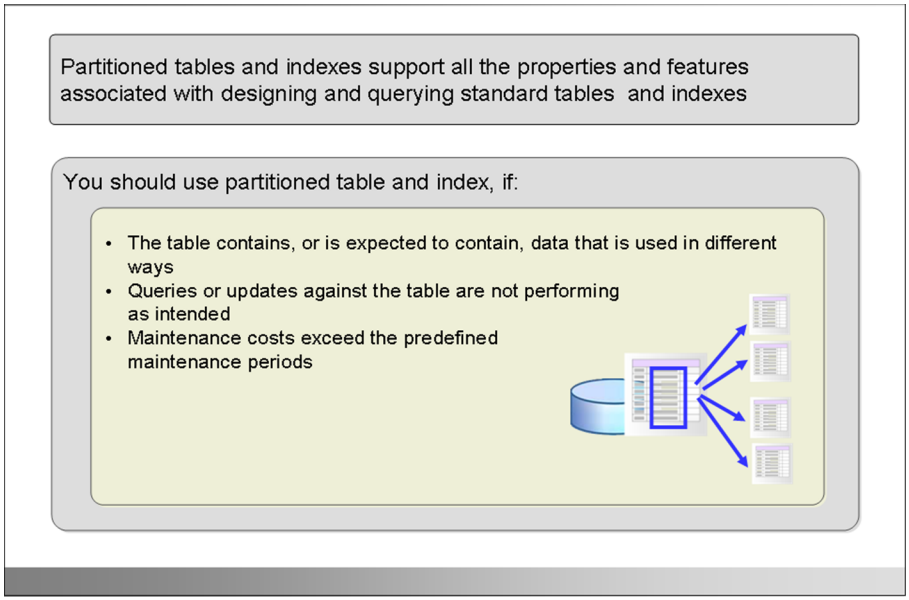
Question: Which type of table will use heaps to arrange the storage of data?

Source:

- http://msdn.microsoft.com/en-us/library/ms189051.aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

Introduction to Partitioned Tables and Indexes



Key Points

Partitioning enables you to manage and access subsets of data quickly and efficiently, while maintaining the integrity of a data collection.

The benefits of partitioning include:

- Data loading from an online transactional processing (OLTP) to an online analytical processing (OLAP) system takes less time to complete.
- Maintenance operations can be performed on subsets of data more efficiently with the help of partitioning because these operations target only the data that is required, instead of the whole table.

Partitioned Table and Index

A partition is a user-defined unit of data organization. By default, a table or index has only one partition that contains all the table or index pages. The partition resides in a single filegroup. A table or

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

index with a single partition is equivalent to the organizational structure of tables and indexes in SQL Server 2000.

You can divide the data of partitioned tables and indexes into units that can be spread across more than one filegroup in a database. You should partition the data horizontally, so that groups of rows are saved into individual partitions. You can treat the table or index as a single logical entity when queries or updates are performed on the data. All partitions of a single index or table must reside in the same database.

Partitioned tables and indexes support all the properties and features associated with designing and querying standard tables and indexes, including constraints, defaults, identity and timestamp values, and triggers. Therefore, if you want to implement a partitioned view that is local to one server, you might want to replace it with a partitioned table instead.

When to use Partitioned Table and Index

Deciding whether to implement partitioning depends primarily on how large the table is or how large it will become, how it is being used, and how well it is performing against user queries and maintenance operations.
You should consider using partitioned table and index, if:

- The table contains or is expected to contain a large amount of data that is used regularly in a segmented fashion.
- Queries or updates against the table are not performing as defined during the setting of the system.
- The maintenance costs exceed predefined maintenance periods.

For example, if you use a current month of data primarily for INSERT, UPDATE, DELETE, and MERGE operations, and use the previous month's data primarily for SELECT queries, managing this table may be easier if it is partitioned by month. This benefit can be especially true if regular maintenance operations on the table only have to target a subset of the data. If you do not partition the table, performing these types of operation can consume a large quantity of resources working on an entire data set. With partitioning, you can perform maintenance operations, such as index rebuilds and defragmentation, on a single month of read/write data, while the read-only data is still available for online access.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

To expand on this example, suppose you want to move one month of read-only data from this table to a data warehouse table for analysis. With partitioning, you can separate subsets of data quickly into staging areas for offline maintenance, and then add these subsets as partitions to existing partitioned tables, assuming these tables are all in the same database instance.

Partitioning a table or index might improve query performance, if you design the partitions correctly, based on the types of queries you frequently run and on your hardware configuration.

Question: What are the benefits of using partition?

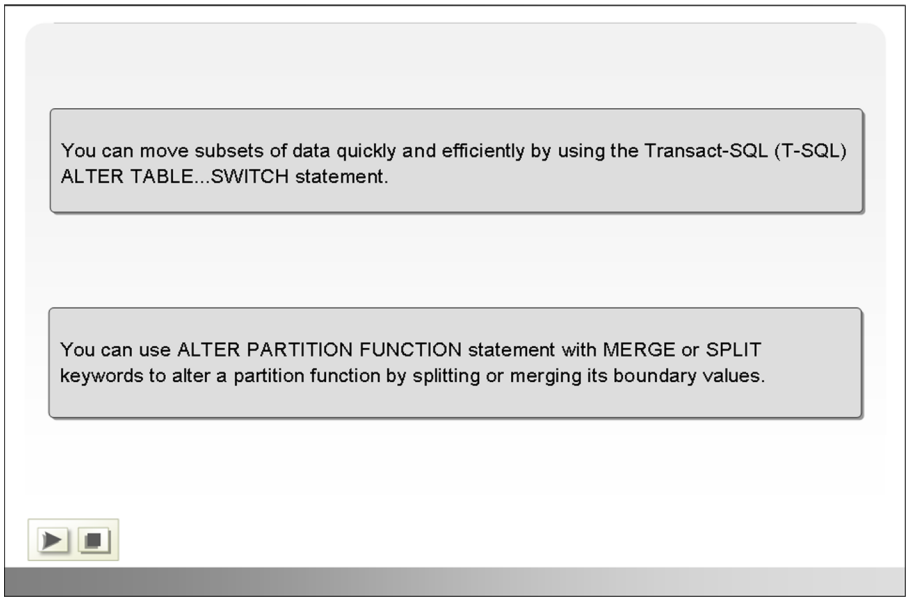
Source:

- http://msdn.microsoft.com/en-us/library/ms189051.aspx
- msdn.microsoft.com/en-us/library/ms190787.aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Sliding-Window Data Management



Key Points

Sliding-window data management is a technique to quickly move data between partitions for storage optimization and simplified management. To move data, you need to follow some rules such as the record date of the data.

By partitioning a table or index, you can move subsets of data quickly and efficiently by using the Transact-SQL (T-SQL) ALTER TABLE...SWITCH statement in the following ways:

- Adding a table as a partition to an existing partitioned table
- Switching a partition from one partitioned table to another
- Removing a partition to form a single table

These scenarios can be useful when you want to add new data to a partitioned table and remove old data from the partitioned table on a regular basis. This operation can involve large or small amounts of data in a variety of scenarios. If new data you are adding has to be loaded, cleaned, or transformed, the data can be treated as a

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

separate entity before adding it as a partition. You can archive or warehouse old data. Regardless of how large or small the collection, the transfer is quick and efficient because, unlike an INSERT INTO SELECT FROM statement, SQL Server does not move data physically.

You can use ALTER PARTITION FUNCTION statement with MERGE or SPLIT keywords to alter a partition function by splitting or merging its boundary values. By executing ALTER PARTITION FUNCTION, a partition of any table or index that uses the partition function can be split into two partitions, or two partitions can be merged into a single partition.

Note: ALTER PARTITION FUNCTION fails when there is a disabled clustered index on any tables that use the partition function.

To change the way a table is partitioned, you can implement any of the following options.

Note: Depending on the configuration of your system, these options can vary in resource consumption.

- Create a new partitioned table with the desired partition function, and insert the data from the old table into the new table by using an INSERT INTO...SELECT FROM statement.
- Create a partitioned clustered index on a heap.

Note: Dropping a partitioned clustered index results in a partitioned heap.

- Drop and rebuild an existing partitioned index by using the Transact-SQL CREATE INDEX statement with the DROP EXISTING = ON clause.
- Perform a sequence of ALTER PARTITION FUNCTION statements.

ALTER PARTITION FUNCTION repartitions any tables and indexes that use the function in a single atomic operation. All filegroups that are affected by ALTER PARITITION FUNCTION must be online. However, the ALTER PARTITION FUNCTION atomic operation is performed offline. Sometimes, the offline operation can be resource-intensive, depending on the extent of repartitioning.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Requirements for Switching Partitions

When a partition is transferred, the data is not physically moved; only the metadata about the location of the data changes. Before you can switch partitions, several general requirements must be met:

- **Both tables must exist before the SWITCH operation**. The table from which the partition is being moved (the source table) and the table that is receiving the partition (the target table) must exist in the database before you perform the switch operation.
- **The receiving partition must exist and it must be empty**. Whether you are adding a table as a partition to an already existing partitioned table, or moving a partition from one partitioned table to another, the partition that receives the new partition must exist and it must be an empty partition.
- **The receiving nonpartitioned table must exist and it must be empty**. If you are reassigning a partition to form one nonpartitioned table, the table that receives the new partition must exist and it must be an empty nonpartitioned table.
- **Partitions must be on the same column**. If you are switching a partition from one partitioned table to another, both tables must be partitioned on the same column.
- **Source and target partitions must share the same filegroup**. The source and the target partitions of the ALTER TABLE...SWITCH statement must reside in the same filegroup, and their large-value columns must be stored in the same filegroup. Any corresponding indexes, index partitions, or indexed view partitions must also reside in the same filegroup. However, the filegroup can be different from that of the corresponding tables or other corresponding indexes.

Example for Partitioning by Using Sliding-Window Data Management Technique

In a sales scenario, consider that the current month data available in the filegroup is October 2002. When the next month's data October 2004 is available, you need to follow a specific sequence of steps to use the existing filegroups and to switch the data in and out. The data currently in the filegroup is data from October 2002. When the

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

data for October 2004 is available, you have two options to add the data to the filegroup depending on available space and archiving requirements. In order to switch a partition in or out of a table quickly, the switch must change metadata only. Specifically, the new data in the table must be created on the same filegroup that is being switched in or out. If you plan to continue to use the existing filegroups, then you need to determine how to handle the space and archiving requirements as follows:

- **Time constraint**. To minimize the amount of time where the table does not have two complete years of data, and if you have the space, you can load the current data of October 2004 into the filegroup without removing the October 2002 data to be archived.
- **Space constraint**. If you do not have enough space to keep both the current month and the month to be archived, then you will need to switch the old partition out first.

The archiving of data is quick both in the time and space constraint cases. It is a good practice to back up the filegroup immediately after you load the new partition and switch it in. You should follow this practice over backing up the filegroup just before you plan to switch it out. For example, if a failure occurs in the RAID array, you should restore the filegroup rather than rebuild or reload the data. In the sales scenario, where the database was recently partitioned, you should consider performing a full backup after stabilizing the partitioning structure. However, performing a full database backup is not the only option. You can select from a number of backup strategies in SQL Server 2008. These backup strategies offer better granularities for backup and restore. In a case where a huge volume of data is not frequently changing, back up the individual filegroups after they are loaded. You should consider this plan as a part of your rolling partition strategy.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

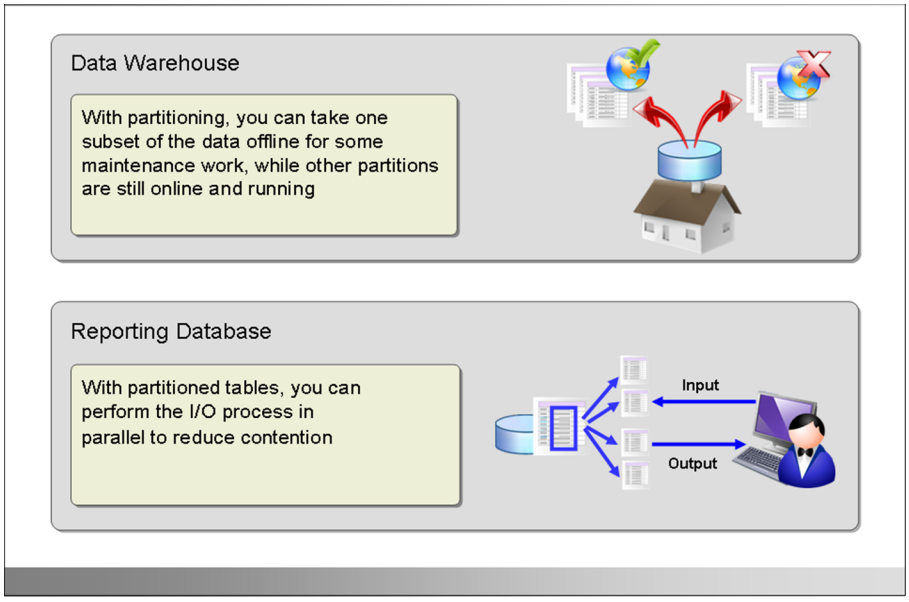
Question: What are the benefits of using the ALTER TABLE ...SWITCH statement over other ways to move data?

Source:

- http://msdn.microsoft.com/en-us/library/ms191174.aspx
- http://technet.microsoft.com/en-us/library/ms191160.aspx
- msdn.microsoft.com/en-us/library/ms345146(SQL.90).aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

Scenarios for Using Partitioned Tables and Indexes



Key Points

The scenarios for using partitioned tables and indexes are:

- **Data warehouse**. Data warehouse can store huge amounts of data, which makes administration difficult in terms of backup and recovery. Dividing data into multiple files and partitions can help shorten the backup and restore window. With partitioning, you can take one subset of the data offline for maintenance work while other partitions are still online and running. For example, for a data warehouse with lots of data, you can use partitioning to have the data stored in multiple database filegroups and have partitioning save the data to the correct filegroup based on the logic defined in a partitioning function.
- **Reporting Database**. By using partitioned tables, you can run the I/O process in parallel. This reduces contention.
  For example, if you have to read the data of 12 months, you can divide data into 12 partitions, and according to the data, the record is created. Therefore, when users are querying data for a

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

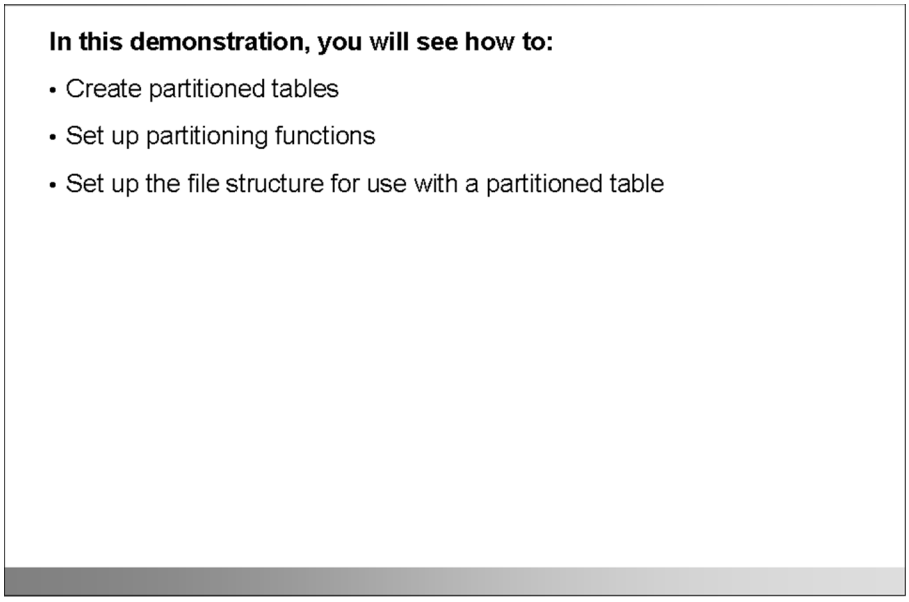
particular year, they are reading data from 12 files simultaneously, which helps the parallel processing of the I/O. This also benefits users who just query data for one month because other users will not be blocked while they are querying data for another month.

Question: How does partitioning help you to manage the data storage of a reporting database?

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Demonstration: How To Configure Partitions



Key Points

The steps to configure partitions are:

1. In Windows Explorer, browse to the **D:\Demofiles\Mod03** folder, and then double-click the **M3_Demo2.sql** file.
2. In the Microsoft SQL Server Management Studio window, in the Query Editor pane, select the **Add database files and filegroups for partitioning** code, and then click **Execute**.
3. In the Microsoft SQL Server Management Studio window, in the Query Editor pane, select the **Create a partition function** code, and then click **Execute**.
4. In the Microsoft SQL Server Management Studio window, in the Query Editor pane, select the **Create a partition scheme** code, and then click **Execute**.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

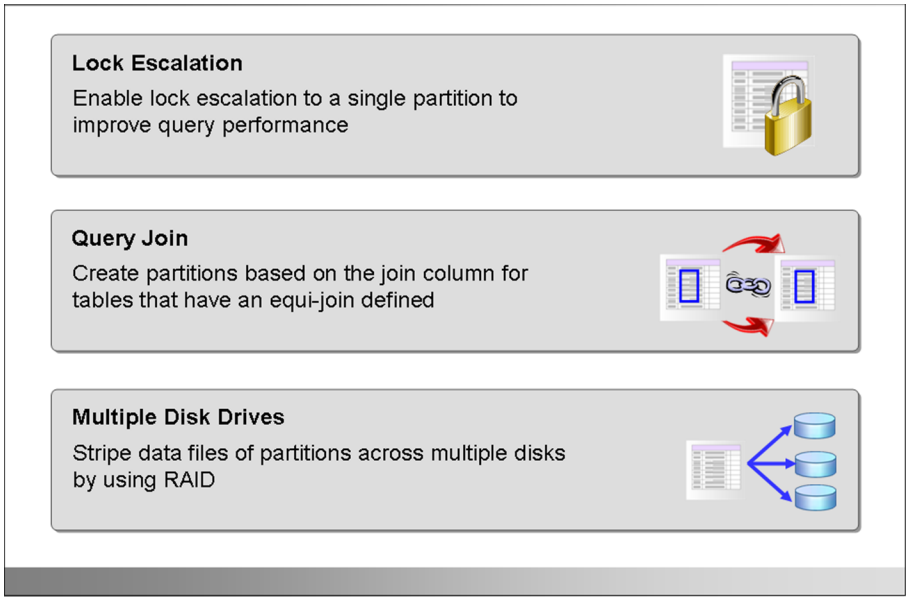
5. In the Microsoft SQL Server Management Studio window, in the Query Editor pane, select the **Apply the partition scheme to the table** code, and then click **Execute**.

Question: What is the use of a partitioning function?

Question: What is the purpose of a partitioning scheme?

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

Considerations for Planning Physical Storage



Key Points

Partitioning a table or index may improve query performance, based on the types of queries you frequently run and on your hardware configuration. You can plan physical storage for the partitioned data depending upon various factors.

These factors are:

- **Lock Escalation Behavior**. Partitioning tables can improve query performance by enabling lock escalation to a single partition, instead of a whole table. To reduce lock contention by allowing lock escalation to the partition, use the LOCK_ESCALATION option of the ALTER TABLE statement.
- **Query Join**. If you frequently run queries that involve an equi-join between two or more partitioned tables, their partitioning columns should be the same as the columns on which the tables are joined. Additionally, you should collate the tables, or their indexes. The partitioning function should:

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

- Have the same number of parameters that are used for partitioning, and the corresponding parameters should be of the same data type.
- Define the same number of partitions.
- Define the same boundary values for partitions.

The SQL Server query optimizer can process the join faster, because the partitions themselves can be joined. If a query joins two tables that are not collocated or are not partitioned on the join field, the presence of partitions may actually slow down query processing, instead of accelerating it.

- **Multiple Disk Drives**. You might map your partitions to filegroups, each accessing a different physical disk drive, to improve I/O performance. When SQL Server performs data sorting for I/O operations, it sorts the data first by partition. Under this scenario, SQL Server accesses one drive at a time, and this might decrease performance. A better solution in terms of performance is to stripe the data files of your partitions across more than one disk by setting up a RAID. In this way, although SQL Server still sorts data by partition, it can access all the drives of each partition at the same time. You can design this configuration regardless of whether all partitions are in one filegroup or multiple filegroups.

Question: How does partitioning help to reduce the lock contention problem?

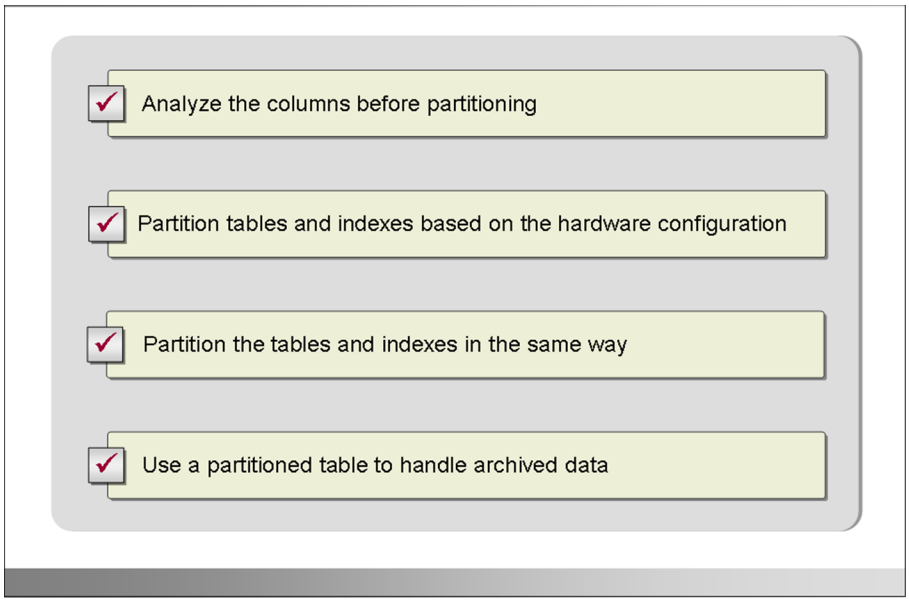
Source:

- http://msdn.microsoft.com/en-us/library/ms177411.aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Best Practices for Using Partitioned Tables and Indexes



Key Points

You should apply the following best practices while implementing a partitioned table and index:

- **Analyze the columns before partitioning**. This helps to improve performance. For example, you can create a partitioning function based on the order date of order tables. This helps in performing the I/O processing in parallel and benefits most users.
- **Partition tables and indexes based on the hardware configuration**. If you have multiple RAID drives, you may want to put partitioned data into multiple drives to perform the I/O processing in parallel.
- **Partition the tables and the indexes in the same way**. Aligning an index with a partitioned table is particularly important if you anticipate that the table will expand by taking on additional partitions, or if will be frequent partition switches. When a table

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

and its indexes are in alignment, SQL Server can switch partitions quickly and efficiently while maintaining the partition structure of both the table and its indexes.

- **Use a partitioned table to handle archived data**. Instead of moving data away from the physical server for archiving, you can use a partitioned table for archiving the data. Partitioning helps keep live data and historical data in the same box, which allows the application to retrieve them at the same time without modifying the program.

Question: What will be the best way to configure partitioning for both a table and an index?

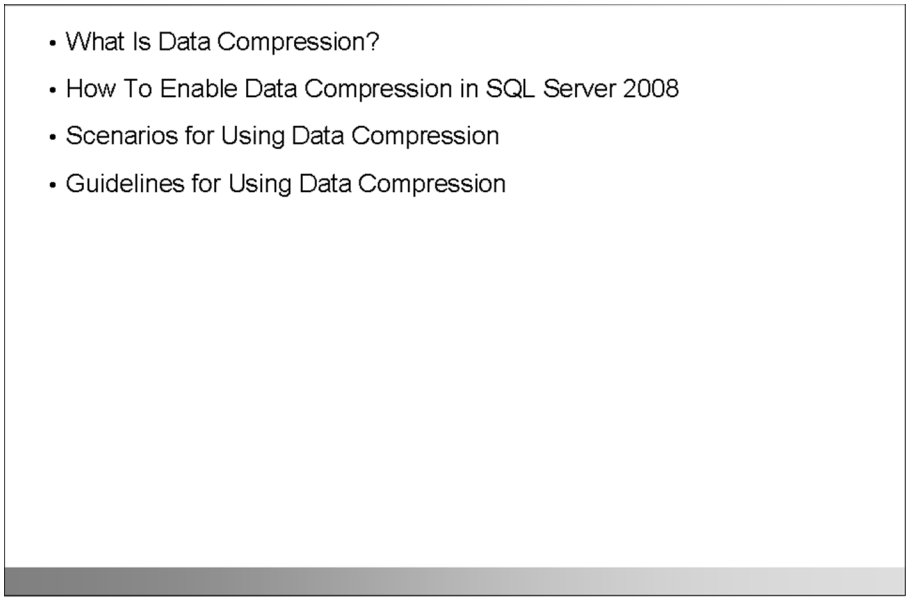
Source:

- msdn.microsoft.com/en-us/library/ms187526.aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Lesson 3
Designing a Strategy for Data Compression



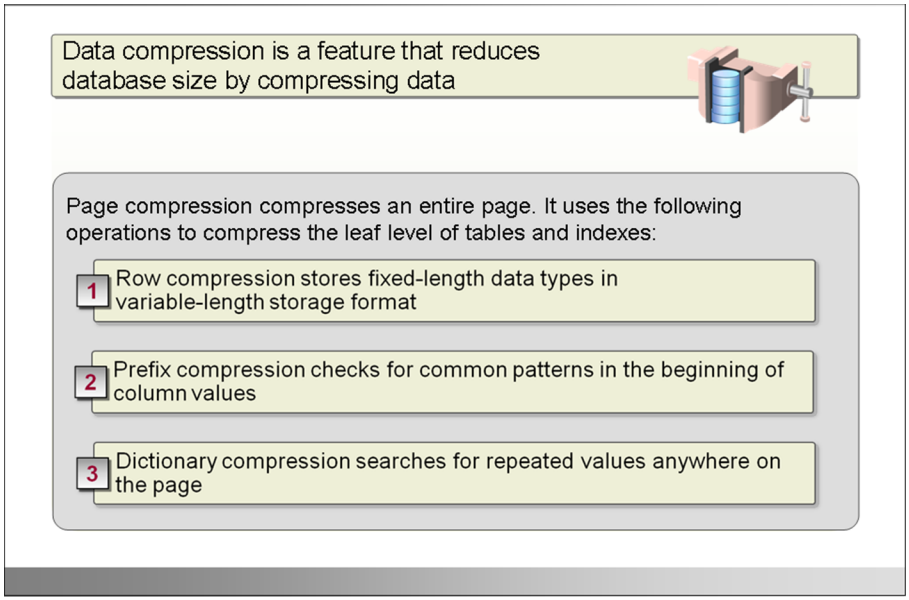
SQL Server allows you to manage the disk space that you need for your data by using compression. SQL Server supports page, row, prefix, and dictionary compression. Compression also helps to reduce the query execution time by reducing the amount of data SQL Server needs to process when it retrieves the information from a database file.

In this lesson, you will learn about page and row compression. You will also learn about the scenarios where compression is suitable. Finally, you will learn the guidelines to use data compression.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

What Is Data Compression?



Key Points

The data compression feature in SQL Server 2008 helps to compress the data in a database. Compressing a database reduces the size of the database. Data compression also helps to improve the performance of input/output (I/O) intensive workloads. This is because compressed data is stored in fewer pages. Therefore, queries need to read fewer pages from the disk.

However, when using data compression, extra CPU resources are required on the database server to compress and decompress the data while data is retrieved from the data page.

You can configure data compression for the following database objects:

- A whole table that is stored as a heap
- A whole table that is stored as a clustered index
- A whole nonclustered index

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

- A whole indexed view

SQL Server 2008 supports row and page compression for both tables and indexes. For partitioned tables and indexes, you can configure compression options for each partition, and it is not mandatory for the various partitions of an object to have the same compression setting.

The compression setting of a table is not automatically applied to its nonclustered indexes. You need to set the compression setting for each index individually. Compression is not available for system tables. You can compress tables and indexes when you create them by using the CREATE TABLE and CREATE INDEX statements. To change the compression state of a table, index, or partition, use the ALTER TABLE or ALTER INDEX statements. SQL Server 2008 provides two types of compression:

Page Compression

Page compression changes the way in which SQL Server stores data into a data page. This feature helps to reduce the storage space required. In addition to storing data efficiently inside a row, page compression minimizes data redundancy in columns in one or more rows on a given page. Minimizing data redundancy optimizes the storage of multiple rows in a page. Page compression is only available in the Enterprise and Developer editions of SQL Server 2008.

Page compression is similar for tables, table partitions, indexes, and index partitions. When you use page compression, nonleaf-level pages of indexes are compressed by using only row compression.

Compressing the leaf level of tables and indexes with page compression consists of three operations in the following order:

1. **Row Compression**

   Row compression helps store data more efficiently in a row by storing fixed-length data types in variable-length storage format. A compressed row uses 4 bits per compressed column to store the length of the data in the column. NULL and 0 values across all data types do not require additional space other than these 4 bits.

   Enabling compression changes only the physical storage format of the data that is associated with a data type, but not its syntax or

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

semantics. You do not require application changes when one or more tables are enabled for compression. The new record storage format has the following main changes:

It reduces the metadata overhead that is associated with the record. This metadata is information about columns, and their lengths and offsets. In some cases, the metadata overhead might be larger than the old storage format.

It uses variable-length storage format for numeric types such as integer, decimal, and float, and for numeric-based types such as datetime and money.

It stores fixed character strings by using variable-length format by not storing the blank characters.

2. **Prefix Compression**

Prefix compression checks for common patterns in the beginning of the column values on a given column across all rows on each page. For each page that is being compressed, prefix compression uses the following steps:

1. For each column, a value is identified that can be used to reduce the storage space for the values in each column.
2. A row that represents the prefix values for each column is created and stored in the compression information (CI) structure that immediately follows the page header.
3. The repeated prefix values in the column are replaced by a reference to the corresponding prefix. If the value in a row does not exactly match the selected prefix value, a partial match can still be indicated.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

The following illustration shows a sample page of a table before prefix compression.

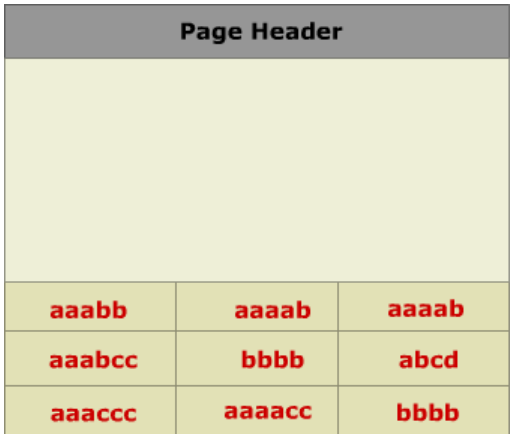


Figure 4.2. A sample page of a table 1

The following illustration shows the same page after prefix compression. The prefix is moved to the header, and the column values are changed to reference the prefix value extracted.

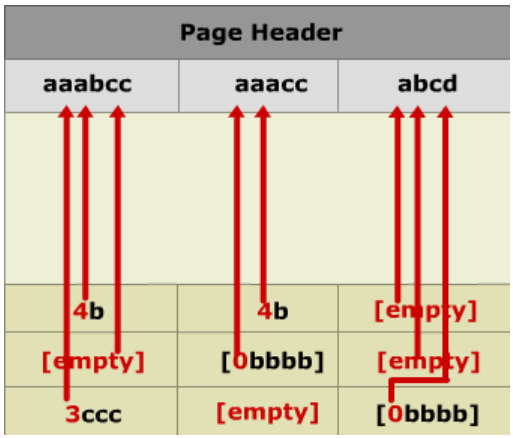


Figure 4.3. The same page after prefix c 1

In the first column of the first row, the value 4b indicates that the first four characters of the prefix (aaab) are present for that row, and

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

also the character b. The resultant value becomes aaabb, which is the original value.

3. **Dictionary Compression**

   After you complete prefix compression, you need to apply dictionary compression. Dictionary compression searches for repeated values anywhere on the page, and stores them in the CI area. Unlike prefix compression, dictionary compression is not restricted to one column. Dictionary compression can replace repeated values that occur anywhere on a page.

   The following illustration shows the same page after dictionary compression.

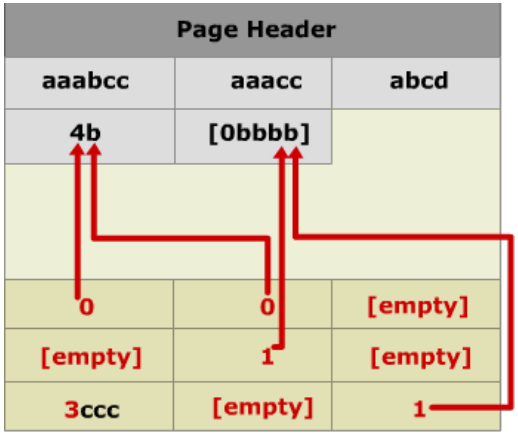


Figure 4.4. The same page after dictionary 4b

Note: The value 4b has been referenced from different columns of the page.

When Page Compression Occurs

When SQL Server creates a new table that has page compression, compression does not occur. However, the metadata for the table indicates that page compression should be used. As SQL Server adds data to the first data page, data is row-compressed. Because the page is not full, no benefit is gained from page compression. When the row is full, the next row that SQL Server adds initiates the page compression operation. It reviews the whole page, evaluates each

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

column for prefix compression, and then evaluates all columns for dictionary compression. If page compression has created enough room on the page for an additional row, SQL Server adds the row, and the data is both row-compressed and page-compressed. If the space gained by page compression minus the space that is required for the CI structure is not significant, SQL Server cannot use page compression for that page. Future rows either fit onto the new page or, if they do not fit, SQL Server adds a new page to the table. Similar to the first page, the new page is not at first page-compressed.

When SQL Server converts an existing table that contains data to page compression, it rebuilds and evaluates each page. Rebuilding all the pages causes the rebuilding of the table, index, or partition.

NTFS Compression

SQL Server 2008 supports a new data type, FILESTREAM that stores the BLOB data in the local or the shared folder directly. Because data is stored in NTFS drives, you can use NTFS compression to optimize the disk space consumption on the server. You can compress the data stored to save disk space. However this would require an extra CPU to compress and decompress the data when it is written or read, respectively. You cannot use compression if the data is essentially uncompressible.

Hence, you should enable compression only when the data can be heavily compressed, and when the extra CPU required will not decrease the workload performance. In addition, compression can only be enabled when the NTFS cluster size is 4,096 bytes or less.

Question: What is the difference between page compression and row compression?

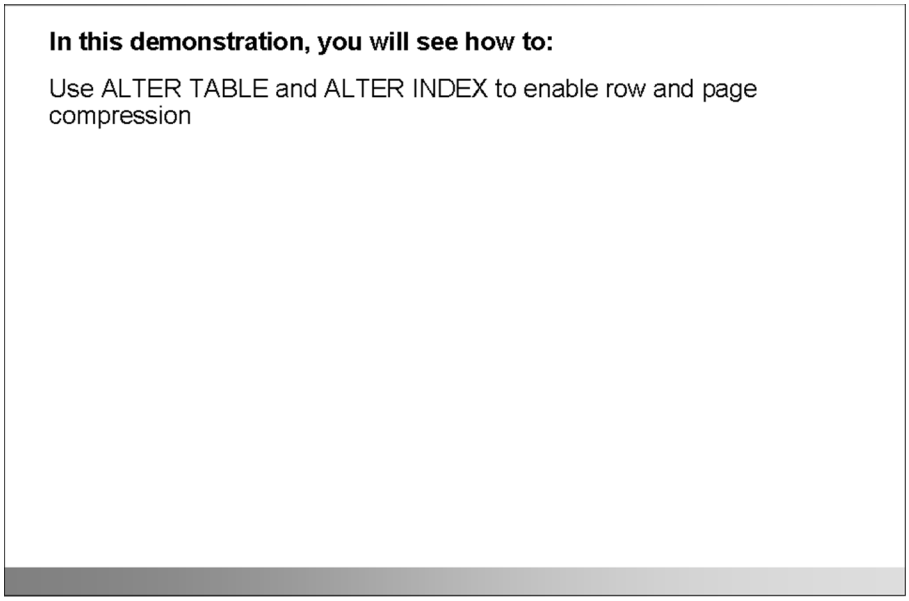
Source:

- http://technet.microsoft.com/en-us/library/cc280464.aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Demonstration: How To Enable Data Compression in SQL Server 2008



Key Points

The steps to use ALTER TABLE and ALTER INDEX to enable row and page compression are:

1. In Windows Explorer, browse to the **D:\Demofiles\Mod03** folder, and then double-click the **M3_Demo4.sql** file.
2. In the Microsoft SQL Server Management Studio window, in the Query Editor pane, select the **Check the disk consumption with different compression** code, and then click **Execute**.

   Note: Ensure whether QuantamCorp is selected in the Available Database list.
3. In the Microsoft SQL Server Management Studio window, in the Query Editor pane, select the **Enable row compression on the Production.WorkOrderRouting table** code, and then click **Execute**.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

4. In the Microsoft SQL Server Management Studio window, in the Query Editor pane, select the **Enable page compression on the Production.WorkOrderRouting table** code, and then click **Execute**.
5. In the Microsoft SQL Server Management Studio window, in the Query Editor pane, select the **Check the status of the compressed table and index** code, and then click **Execute**.

Question: What is the benefit of compressing tables and indexes?

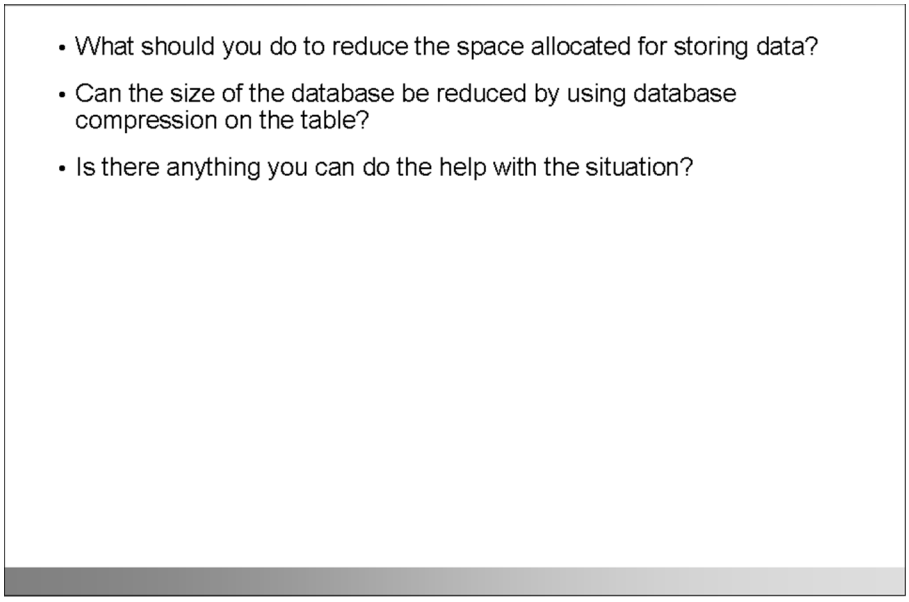
Question: Is data compression a database-level or table-level configuration?

Question: Can a table have different compression settings?

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Discussion: Scenarios for Using Data Compression



Key Points

In this discussion, you are presented scenarios in which you are asked to determine the scenarios for data compression. You and your classmates will discuss possible solutions to the scenario.

Scenario 1

You are reviewing the database design of an application. You notice that a number of fields in a table are using long character data type, such as char(2000). However, most of the fields are not fully filled and contains only five characters or are NULL.

Question: What should you do to reduce the space allocated for storing data?

Scenario 2

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Your application team uses an application for storing the product image. The team wants to reduce the size of the database.

Question: Can the size of the database be reduced by using database compression on the table?

Scenario 3

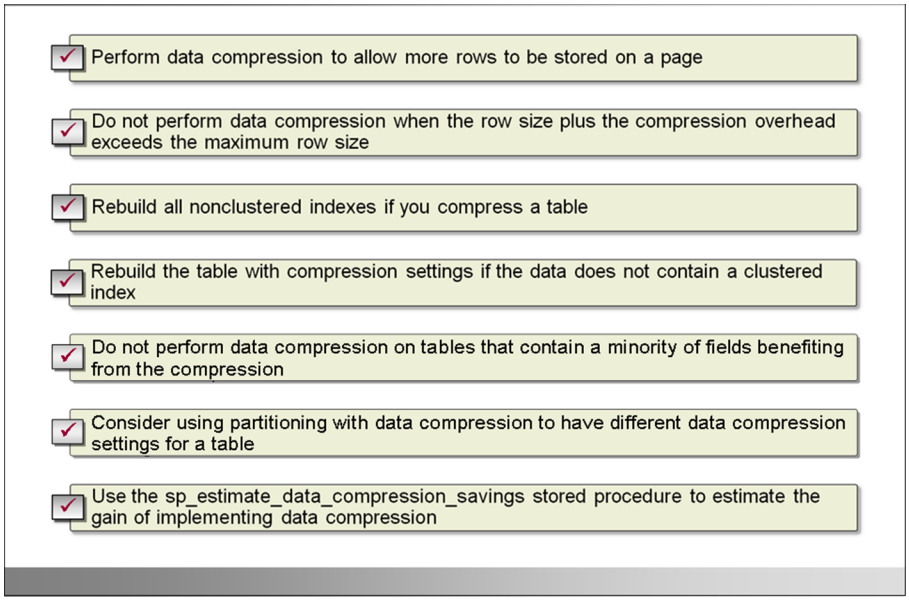
You have an application database which always experience poor performance because of the fragmentation of the indexes. You have configured a SQL Server Agent job to rebuild an index every six hours to ensure that the application is performing properly.

Question: Is there anything you can do the help with the situation?

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Guidelines for Using Data Compression



Key Points

When designing and implementing data compression, refer to the following guidelines:

- **Use compression to increase the number of rows to be stored on a page**. However, compression does not change the maximum row size of a table or index. This is to allow more efficient use of storage space in SQL Server and improve the efficiency for SQL Server to retrieve data.
- **Do not perform data compression when the maximum row size plus the compression overhead exceeds the maximum row size**. This would cause an error as SQL Server does not allow a row to exceed the maximum row size.
- **Rebuild all nonclustered indexes if you compress a table**. The reason being that the new pages that are allocated in a heap as part of DML operations will not use PAGE compression until the heap is rebuilt.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

- **Rebuild the table with compression settings if the data does not contain a clustered index to point to new row locations in the heap**. This is to ensure that the index are pointing to the right location of the data, which reduces the need for SQL Server to further locate the file after getting a pointer from the heap.
- **Do not perform data compression on tables that contain majority of fields benefiting from compression**. Large-value data types are sometimes stored separately from the normal row data on special purpose pages because of their size. Data compression is unavailable for such data.
- **Consider using partitioning with data compression to have different data compression settings for a table**. This variation in settings helps to provide optimal query performance for your queries.
- **Use the sp_estimate_data_compression_savings stored procedure to estimate the gain of implementing data compression**. This procedure returns the current size of tables and estimates the table size for the requested compression state. Compression can be evaluated for whole tables or parts of tables. This includes heaps, clustered indexes, nonclustered indexes, indexed views, and table and index partitions. Table structures can be compressed by using row compression or page compression. If the table, index, or partition is already compressed, you can use this procedure to estimate the size of the table, index, or partition, if it was uncompressed

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

The following example estimates the size of the Production.WorkOrderRouting table in the QuantamCorp database, which is compressed by using ROW compression.

```
USE QuantamCorp2008

GO

EXEC sp_estimate_data_compression_savings 'Production',

'WorkOrderRouting', NULL, NULL, 'ROW' ;

GO
```

Source:

- http://msdn.microsoft.com/en-us/library/cc280574%28SQL.105%29.aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Lab 3: Designing the Physical Structure of SQL Server 2008

- Exercise 1: Planning the Physical Structure of SQL Server 2008 (Discussion)
- Exercise 2: Implementing the Physical Structure of a Database
- Exercise 3: Configuring Partitions
- Exercise 4: Implementing Data Compression
- Exercise 5: Configuring FILESTREAM Support in SQL Server 2008

Logon Information

Virtual machine	50400A-NYC-SQL1
User name	Administrator
Password	Pa$$w0rd

Estimated time: 60 minutes

Objectives:

After completing this lab, you will be able to:

- Plan the physical structure of SQL Server 2008.
- Implement the physical structure of a database.
- Configure partitions.
- Implement full-text indexing.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Scenario

QuantamCorp has deployed two SQL Server 2008 servers in their organization and will add more to replace their existing installation of SQL Server 2005. After this, SQL Server 2008 will be the primary database engine used in their organization. After upgrading from SQL Server 2005 to SQL Server 2008, you need to fine-tune the physical structure of the databases hosted on the server to improve the performance and ensure that the structure can allow for future growth. Also, the application team has requested you to enable the full-text indexing feature because they will be implementing a search feature on an application that will use the full-text query in SQL Server 2008. The application team also wants to exclude a few words from full-text search for compliance reasons.

To optimize the performance of the databases and fulfill the request of the application team so that the internal application will function properly and with reasonable speed, you need to implement the following considerations and tasks:

- Decide whether you should have multiple database files and filegroups for a database.
- Decide whether you should use partitioning in your database.
- Decide how you should partition your data.
- Decide how you should configure database files for optimal performance.
- Configure the database with multiple filegroups and have multiple files for each filegroup.
- Configure partitioning for a table and related index.
- Enable and configure full-text indexing.
- Determine the impact of using a stoplist.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Exercise 1: Planning the Physical Structure of SQL Server 2008 (Discussion)

Scenario

You have received a request to create a database for a new application that is used to store details of all calls made into your customer service hotline. The IVR application that is developed by the application team to automate the enquiry process of thecustomer service hotline will also use the database. The estimatedvolume of data on the database would increase by 10 GB everyweek.

You have reviewed the design of the tables. After communicatingwith the application team, you have grouped the tables into thefollowing categories.

Category	Description
Customer-related tables	These tables are used to store customer information, such as customer profiles.
Product-related tables	These tables are used to store product information, such as product names and product category.
Call Details table	This table stores details of calls.

The application team is concerned about the volume of data in the Call Details table because most of the data is related to the call details. The others tables are quite static in nature and do not expand greatly during normal operations.

The call logging application has features that users can use to search for call details stored as a text field in the call log table. During initial testing, the application team found that the query performance was poor. While you help them investigate the issue, you realize they are using the LIKE operator for querying.

Based on the above scenario, answer the following questions:

1. How would you design the physical structure of the database sto provide optimal query performance?

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

2. How can you resolve the concerns of the application team related to the performance when querying the call details table that contains a large volume of data?
3. What would you suggest to the team to help improve performance of the queries that access the call log table?
4. The application team wants to prevent users from searching for call information based on some restricted words. What would you suggest?
5. After you have configured your filegroups, what should you add to the database documentation that is available to the application design team when they are working on the table design?

Exercise 2: Implementing the Physical Structure of a Database

Scenario

You are creating a new database for your application team. Due to the storage and performance requirements, you want to distribute the I/O workloads and storage across multiple disks for better scalability and performance. You have also decided to add additional filegroups and database files for implementing partitioned table and index.

The main tasks for this exercise are as follows:

1. Create a test database.
2. Add filegroups to the database.
3. Modify the physical structure of the database.
4. Verify the physical file structure of the database.

Task 1: Create a test database.

1. Open the **Microsoft SQL Server Management Studio** window.
2. In the **Connect to Server** dialog box, type the following details to connect to Microsoft SQL Server Management Studio:
   - Server type: **Database Engine**
   - Server name: **NYC-SQL1**

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

- Authentication: **Windows Authentication**

3. In the Microsoft SQL Server Management Studio window, in the Query Editor pane, run a T-SQL code create a test database named **Mod03**, with the following structure.

Database file name	Settings
Mod03_Primary.mdf	Current size: 10 Maximum size: 50 Filegrowth: 5
Mod03_Log.ldf	Current size: 10 Maximum size: 50 Filegrowth: 5

Note: Browse to the **D:\Labfiles\Mod03** folder and select the **E2_T1.sql** file to execute the required T-SQL code.

Task 2: Add filegroups to the database.

1. In the Microsoft SQL Server Management Studio window, in the Query Editor pane, run a T-SQL code to add three filegroups to the database: **HistorialData**, **PresentData**, and **TableData**.

Note: Browse to the **D:\Labfiles\Mod03** folder and select the **E2_T2.sql** file to execute the required T-SQL code.

Task 3: Modify the physical file structure of the database.

1. In the Microsoft SQL Server Management Studio window, in the Query Editor pane, add two files, **Table_Data01** and **Table_Data02** to the **TableData** filegroup, with the following structure.

Database file name	Settings
Mod03_Data01.ndf	Logical name: Table_Data01 Current size: 10 Maximum size: 50

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Database file name	Settings
	Filegrowth: 5
Mod03_Data02.ndf	Logical name: Table_Data02 Current size: 10 Maximum size: 50 Filegrowth: 5

2. Add the **History_Data01** file and the **History_Data02** file to the **HistorialData** filegroup, with the following structure.

Database file name	Settings
Mod03_History01.ndf	Logical name: History_Data01 Current size: 10 Maximum size: 50 Filegrowth: 5
Mod03_History02.ndf	Logical name: History_Data02 Current size: 10 Maximum size: 50 Filegrowth: 5

3. Add the **Present_Data01** file and the **Present_Data02** file to the **PresentData** filegroup.

File name	Settings
Mod03_Present01.ndf	Logical name: Present_Data01 Current size: 10 Maximum size: 50 Filegrowth: 5
Mod03_Present02.ndf	Logical name: Present_Data02 Current size: 10 Maximum size: 50 Filegrowth: 5

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

4. Add a log file named **Mod03_Log02** to the **Mod03** database.

Note: Browse to the **D:\Labfiles\Mod03** folder and select the **E2_T3.sql** file to execute the required T-SQL code.

Task 4: Verify the physical file structure of the database.

1. Verify the file structure of the database by using the **sp_helpdb** stored procedure. sp_helpdb is a system-stored procedure to display the physical file structure of a database.

Results: After completing this exercise, you should have created a database, added filegroups to the database, added files to the filegroup, and verified the physical file structure of the database by using the **sp_helpdb** stored procedure.

Exercise 3: Configuring Partitions

Scenario

For the application database that you implemented earlier, you want to implement partitioned tables so that you can distribute the I/O of the same table across multiple disks for better performance. You also want to implement partitioned table in order to reduce the management overhead for moving data between physical files or disks.

The main tasks for this exercise are as follows:

1. Create a partition function and a partition scheme.
2. Create a test table by using the partition scheme.
3. Verify the partition scheme.

Task 1: Create a partition function and a partition scheme.

1. In the Microsoft SQL Server Management Studio window, create a partition function named **pfCallDetails**. Create the partition function on the date of call. The function should assign data to different partitions by using the date range from **01/01/2006** to **01/01/2009**.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

2. Create a partition scheme named **psCallDetails** for the **pfCallDetails** partition.

Task 2: Create a test table by using the partition scheme.

1. Using the Microsoft SQL Server Management Studio window, create a table, **CallsDetails** with the following fields:
   - Call_date: **datetime**
   - Caller: **varchar**
   - Duration: **int**
   - Additional_info: **nchar**
   - Allow NOT NULL values only in the **Call_date** and the **Caller** fields
2. Create the table by using the partition scheme **psCallDetails**. Create a constraint named **PK_CallDetails** on the primary key with the following options:
   - PAD_INDEX = OFF
   - STATISTICS_NORECOMPUTE = OFF
   - IGNORE_DUP_KEY = OFF
   - ALLOW_ROW_LOCKS = ON
   - ALLOW_PAGE_LOCKS = ON

Note: Browse to the **D:\Labfiles\Mod03** folder and select the **E2_T2.sql** file to execute the required T-SQL code.

Task 3: Verify the partition scheme.

1. In the Microsoft SQL Server Management Studio window, in the **Disk Space Used by Data Files** report, verify the use of partition scheme.
2. Insert records in the CallDetails table as follows.

Field name	Record
call_date	A date between 2004-01-01 and 2009-12-31

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Field name	Record
caller	Peter
duration	30
additional_info	A random number between one and 100000.

3. In the Microsoft SQL Server Management Studio window, in the **Disk Space Used by Data Files** report, verify the use of partition scheme.

Note: Browse to the **D:\Labfiles\Mod03** folder and select the **E2_T3.sql** file to execute the required T-SQL code.

Results: After completing this exercise, you should have created a partition function, created a partition scheme, created a test table function by using the partition scheme, and used the **Disk Space Used by Data Files** report to verify the partition scheme.

Exercise 4: Implementing Data Compression

Scenario

Your application database has a large number of records that consume a large amount of disk space. You are planning to implement database compression to optimize the storage to reduce the disk usage.

The main tasks for this exercise are as follows:

1. Estimate the data compression savings when it is enabled.
2. Enable row compression and page compression.

Task 1: Estimate the data compression savings when it is enabled.

1. Using the Microsoft SQL Server Management Studio window, estimate the data compression savings for a row on the **CallDetailsHistory** table.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

2. Estimate data compression savings for a page on the **CallDetailsHistory** table.

Task 2: Enable row compression and page compression.

1. Using the Microsoft SQL Server Management Studio window, open the Disk Usage By Table report and verify the disk space and query performance of the **CallDetailsHistory** table.
2. Enable row compression on the **CallDetailsHistory** table.
3. Enable page compression on the **IX_CallDetailsHistory** index.
4. After enabling data compression, update the Disk Usage By Table report and verify the disk space and query performance of the **CallDetailsHistory** table.

Results: After completing this exercise, you should have estimated the data compression savings, enabled row compression on a table, and enabled page compression on an index.

Exercise 5: Configuring FILESTREAM Support in SQL Server 2008

Scenario

One of your application teams needs to store BLOB on the application database. For better storage management, you have decided to use FILESTREAM support instead of using image objects for storing the product image. You are now configuring FILESTREAM support in your database and SQL Server, while your application team configures their database to support the use of the FILESTREAM feature.

The main tasks for this exercise are as follows:

1. Enable FILESTREAM support.
2. Change FILESTREAM access level by using Microsoft SQL Server Management Studio.
3. Create a database with FILESTREAM support.
4. Create a table with FILESTREAM support.
5. Verify the FILESTREAM configuration.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Task 1: Enable FILESTREAM support.

1. Open the Sql Server Configuration Manager window.
2. Using the Sql Server Configuration Manager window, open the **SQL Server (MSSQLSERVER) Properties** dialog box.
3. Using the **SQL Server (MSSQLSERVER) Properties** dialog box enable FILESTEAM support by:
   - Selecting the **Enable FILESTREAM for Transact-SQL access** check box.
   - Selecting the **Enable FILESTREAM for file I/O streaming access** check box.
   - Ensure that the **Allow remote clients to have streaming access to FILESTREAM data** check box is selected.

Task 2: Change FILESTREAM access level by using Microsoft SQL Server Management Studio.

1. Open the Microsoft SQL Server Management Studio window.
2. To connect to Microsoft SQL Server Management Studio, in the Connect to Server dialog box, type the following:
   - Server type: **Database Engine**
   - Server name: **NYC-SQL1**
   - Authentication: **Windows Authentication**
3. Using the Microsoft SQL Server Management Studio window, change FILESTREAM access level to two.

Task 3: Create a database with FILESTREAM support.

1. Using the Microsoft SQL Server Management Studio window, create a database named **FILESTREAM_Test** with FILESTREAM support with the following structure.

File	Name	Path
Primary	Data1	D:\Data\archdat1.mdf
Log	Log1	D:\Data\archlog1.ldf

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Also, configure the **D:\data\filestream1** folder as the filestream folder with the **Arch3** name.

Task 4: Create a table with FILESTREAM support.

1. Using the Microsoft SQL Server Management Studio window, create a table named **FILESTREAM_Test.dbo.Records** with the **FILESTREAM_Test** FILESTREAM support with the following structure:

   Id - ROWGUIDCOL (Uniqueidentifier)

   SerialNumber - integer (Unique)

   Chart - Varbinary (Max)

   Allow the **Id** field to have NOT NULL values and the **Chart** field to have NULL values.

Task 5: Verify FILESTREAM configuration.

1. To verify the configuration of the table, insert three records in the **FILESTREAM_Test.dbo.Records** table as follows. Then, select all records from the table.

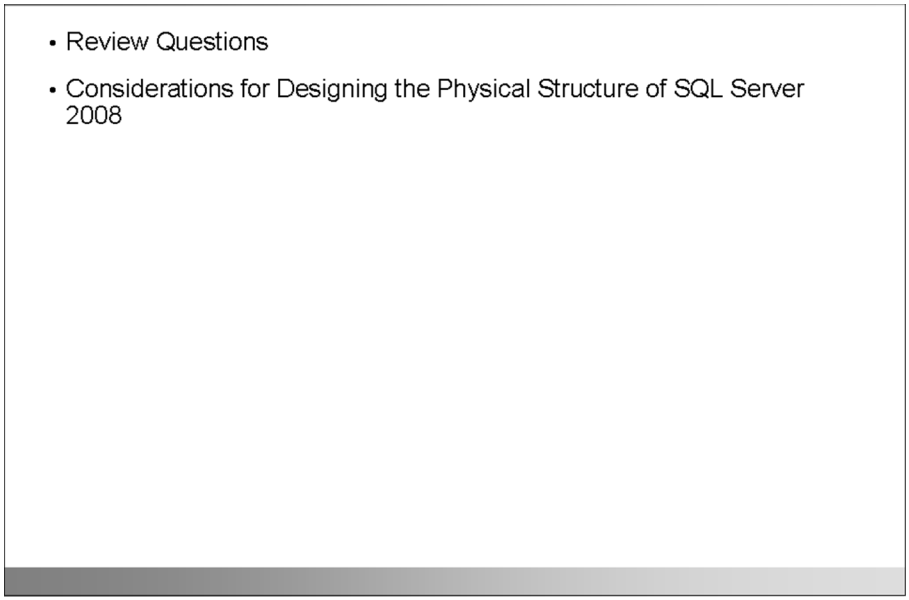
Field name	Record 1	Record 2	Record 3
Id	newid()	newid()	newid()
SerialNumber	1	2	3
Chart	NULL	CAST ('' as varbinary(max))	CAST (1 as varbinary(max))

Results: After completing this exercise, you should have enabled FILESTREAM support, changed FILESTREAM access level by using SSMS, created a database with FILESTREAM support, created a table with FILESTREAM support, and verified FILESTREAM configuration.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Module Review and Takeaways



Review Questions

1. You are configuring the physical storage of an application database. You need to assign multiple database files to a table because of the size of the table. Can you assign a database table to a database file directly?
2. You are evaluating different options for storing BLOB for your application team, which is developing an application for storing the image of company's products in a SQL Server 2008 database. What are the benefits of using FILESTREAM, instead of in-line, for storing BLOB data?
3. You are configuring database partitioning for an OLAP database in your organization. When using partitioning, how can SQL Server determine which partition the data should be assigned to?
4. Your application team informs you that users are complaining about the slow query performance while doing a keyword search. You are considering using full-text indexing to improve performance. Currently, the application is using the LIKE operator

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

to perform the search. What are the benefits of using full-text index over the LIKE operator?

5. You are configuring full-text indexing. You need to provide a stoplist in the configuration window. What is the purpose of using a stoplist?
6. You are configuring the physical architecture of your database. Your senior database architect has asked you to create multiple database files for a table with 1 terabytes of data on multiple physical drives. What are the benefits of using multiple database files?

Considerations for Designing the Physical Structure of SQL Server 2008

When designing the physical structure of SQL Server 2008, consider the following:

- Configure multiple filegroups and have multiple files associated with each filegroup.
- Configure database files located on different physical drives for better performance.
- Configure the default filegroup to any filegroup other than PRIMARY. Store only system tables and data on the PRIMARY filegroup.
- Place table data and indexes on separate database files for better performance.
- Use partitioning to have data of a table and index span into multiple database files for better performance. Always distribute the data based on the column that helps to perform parallel processing of I/O requests.
- Use FILESTREAM to store BLOB data instead of ntext, text and image data type for better manageability.
- Add a clustered index for a table which requires read and write access. Once a clustered index is added, consider adding non-clustered indexes to optimize the performance for a particular query.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

- Create a full-index index on separate drives for better performance and ensure that there is enough space to store the full-text index
- Create a stoplist to prevent users from searching information using specific words.
- Change an application to use full-text predicate instead of LIKE operator for better performance.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Module 4

Designing a Strategy for Maintaining a Database in SQL Server 2008

Contents:

Lesson 1: Designing a Strategy for Maintaining Statistics for a Database 4-3

Lesson 2: Designing a Strategy for Maintaining Indexes 4-15

Lesson 3: Planning for Full-Text Indexing 4-46

Lab 4: Designing a Strategy for Maintaining a Database in SQL Server 2008 4-69

Module Review and takeaways 4-82

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Module Overview

- Designing a Strategy for Maintaining Statistics for a Database
- Designing a Strategy for Maintaining Indexes
- Planning for Full-Text Indexing

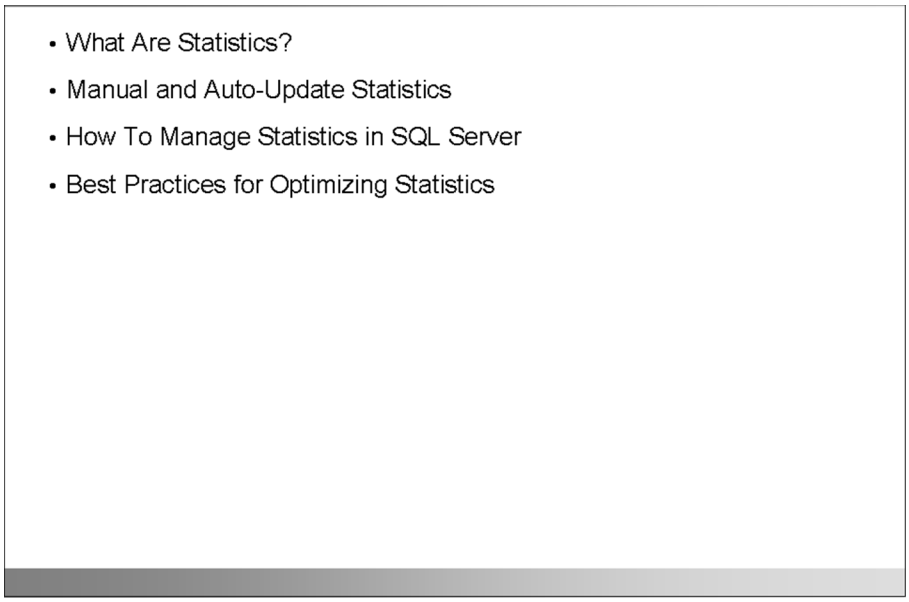
Microsoft™ SQL Server™ 2008 provides tools to govern the performance and usability of the database. These tools include statistics, indexes, and compression. Statistics are important for an optimizer to identify an optimal query plan. SQL Server can automatically create and update statistics. Indexes reduce the data retrieval time and help the database to locate records in a faster and more efficient way. Compression further enhances the operations of the database by optimizing space.

In this module, you will learn how to design a strategy for maintaining statistics and indexes. You will also learn to design a strategy to compress tables and indexes.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Lesson 1
Designing a Strategy for Maintaining Statistics for a Database



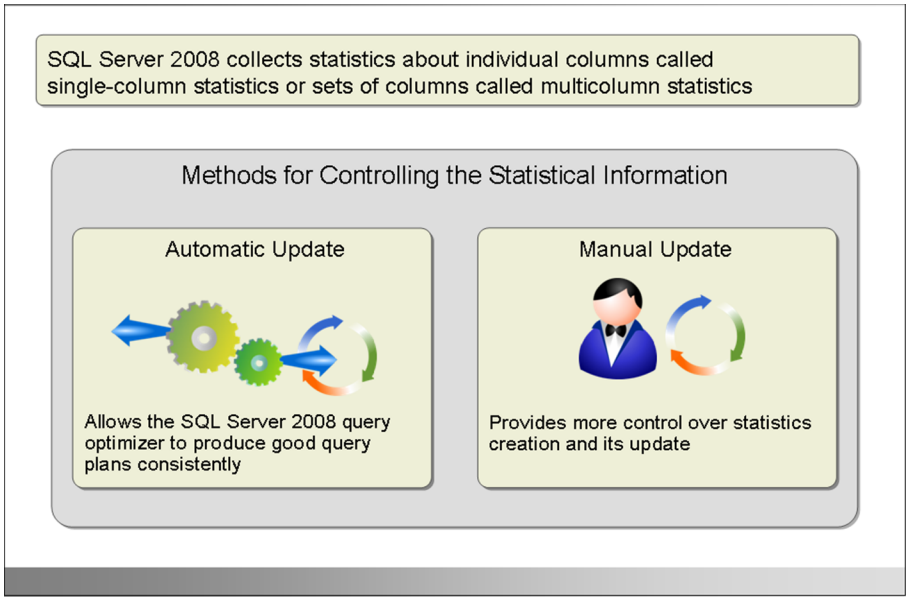
SQL Server 2008 collects statistical information about indexes and column data stored in the database. The SQL Server query optimizer uses this statistical information to choose the most efficient plan for retrieving or updating data. Outdated statistics can cause SQL Server to generate an inefficient execution plan for executing the query.

In this lesson, you will learn about the data that is collected, where it is stored, and how to create, update, and delete statistics in SQL Server. By default, SQL Server 2008 creates and updates statistics automatically, wherever necessary. You will also learn how certain query language features, such as Transact-SQL (T-SQL) variables, interact with the query optimizer, with the help of statistics.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

What Are Statistics?



Key Points

SQL Server 2008 collects statistics about individual columns called single-column statistics or sets of columns called multicolumn statistics. The query optimizer uses statistics to estimate the selectivity of expressions, and thus the size of the intermediate and final query results. Good statistics allow the optimizer to accurately assess the cost of different query plans and then choose a high-quality plan. The sysindexes table stores all the information about a single statistics object in several columns of a single row and in a statistics binary large object (statblob) kept in an internal-only table. In addition, you can find information about statistics in the new metadata views, sys.stats and sys.indexes.

Methods To Control Statistical Outcome

In SQL Server 2008, you can use following methods to control statistical outcome:

- **Automatic Update**. SQL Server 2008 has a number of features for maintaining statistics. The most important one is the ability to

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

automatically create and update statistics. This feature, also called autostats, is enabled by default, starting from SQL Server 2000. Approximately 98 percent of SQL Server 2008 installations leave this feature enabled, which is a best practice. For the majority of databases and applications, you can rely on the automatic creation and update of statistics to provide comprehensive and accurate statistics about their data. This ability allows the SQL Server 2008 query optimizer to produce good query plans consistently, while keeping development and administration costs low.
An important feature for high-throughput online transaction processing application environments is the ability to asynchronously update statistics, automatically. This feature can improve the predictability of query response time in such environments.

- **Manual Update**. If you need more control over statistics creation and update to get good query plans and manage the overhead of gathering statistics, you can use the manual statistics creation and update capabilities. You can use manual update in the following scenarios:
  - A large amount of data is loaded into the table
  - A query frequently experiences a timeout
  - A query is run based on some columns

Advantageous of SQL Server 2008 Statistics

SQL Server 2008 statistics features allow you to:

- Implicitly create and update statistics with the default sampling rate. In the SELECT, INSERT, DELETE, UPDATE, and MERGE commands, use of a column in a query condition, such as a WHERE or JOIN clause, enables you to create or update statistics, if necessary when automatic statistics creation and update is enabled.
- Manually create and update statistics with any desired sampling rate, and drop statistics by using commands such as CREATE STATISTICS, UPDATE STATISTICS, DROP STATISTICS, CREATE INDEX, and DROP INDEX.
- Create and maintain filtered statistics, which are defined over a subset of the rows in the table by using a WHERE clause.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

- Automatically create filtered statistics whenever a filtered index is created.
- Manually create statistics in bulk for all columns of all tables in a database by using sp_createstats.
- Manually update all existing statistics in the database by using sp_updatestats.
- List statistics objects that exist for a table or database by using sp_helpstats, catalog views sys.stats, and sys.stats_columns.
- Display descriptive information about statistics objects by using DBCC SHOW_STATISTICS.
- Enable and disable automatic creation and update of statistics database-wide or for a specific table or statistics object by using ALTER DATABASE options such as AUTO_CREATE_STATISTICS and AUTO_UPDATE_STATISTICS, sp_autostats, and the NORECOMPUTE option on CREATE STATISTICS and UPDATE STATISTICS.
- Enable and disable asynchronous automatic update of statistics by using ALTER DATABASE option, AUTO_UPDATE_STATISTICS_ASYNC.
- SQL Server Management Studio (SSMS) allows you to graphically browse and control statistics objects within its Object Explorer view. Statistics are listed in Object Explorer in a folder under each table object.

Question: How statistics helps in query processing?

Question: Which feature in SQL Server helps to reduce the need for manually editing and managing statistics?

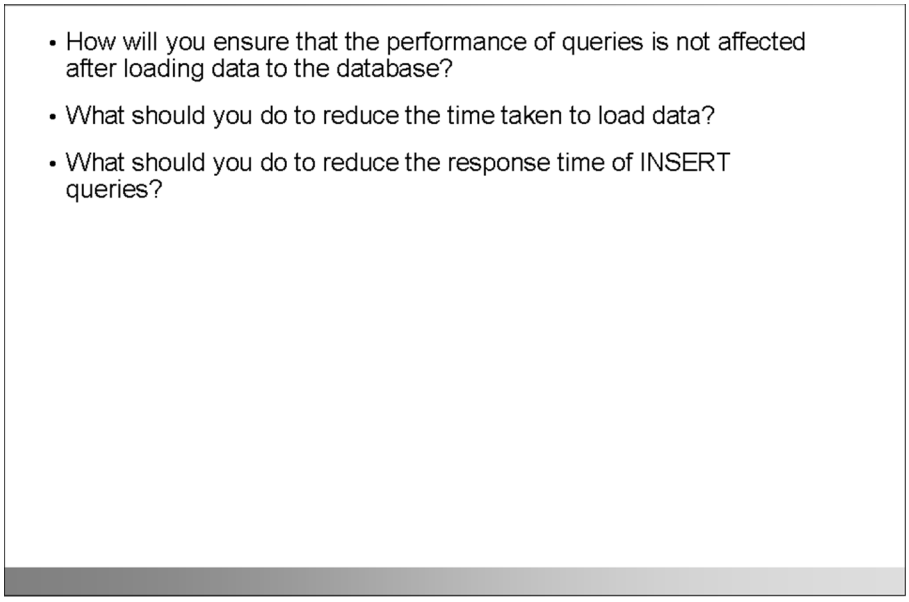
Source:

- http://msdn.microsoft.com/en-us/library/dd535534.aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Discussion: Manual and Auto-Update Statistics



Key Points

In this discussion, you are presented with the following scenarios in which you are asked to compare manual and auto update statistics. You and your classmates will discuss possible solutions to the scenarios.

Scenario 1

You notice that the performance of some queries decreases after you load a large amount of data to the database. While looking at the query plan, you determine that the query is not using the right index.

Question: How will you ensure that the performance of queries is not affected after loading data to the database?

Scenario 2

After enabling automatic statistics update, you notice that the time taken to load data has increased.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Question: What should you do to reduce the time taken to load data?

Scenario 3

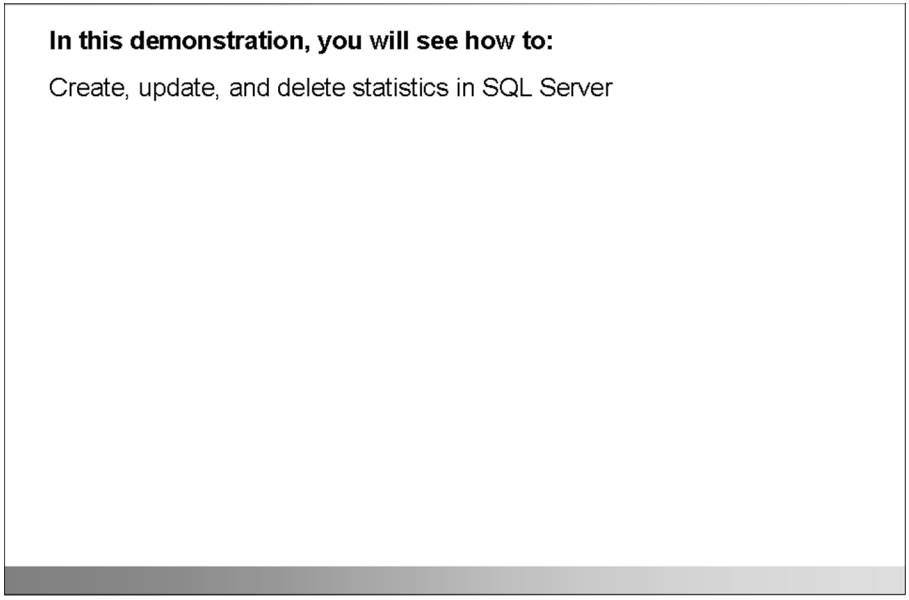
Your application team informs you that the response time for INSERT queries is slow. After some troubleshooting, you determine that the performance of the INSERT query can be increased if you disable the auto statistics update function. However, you cannot disable the auto statistic update function because you might experience performance issues if that option is disabled.

Question: What should you do to reduce the response time of INSERT queries?

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Demonstration: How To Manage Statistics in SQL Server



Key Points

The steps to use statistics in SQL Server are:

1. On the **Start** menu, point to **All Programs**, click **Microsoft SQL Server 2008**, and then click **SQL Server Management Studio**. The Microsoft SQL Server Management Studio window appears.
2. In the Microsoft SQL Server Management Studio window, by using the **Connect to Server** dialog box, connect to the default instance of SQL Server Database Engine.
3. In Windows Explorer, browse to the **D:\Demofiles\Mod04** folder, and then double-click the **M4_Demo1.sql** file.
4. In the Microsoft SQL Server Management Studio window, in the Query Editor pane, select the **Set the statistics xml option on and retrieve records from the HumanResources.Department table** code, and then click **Execute**.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

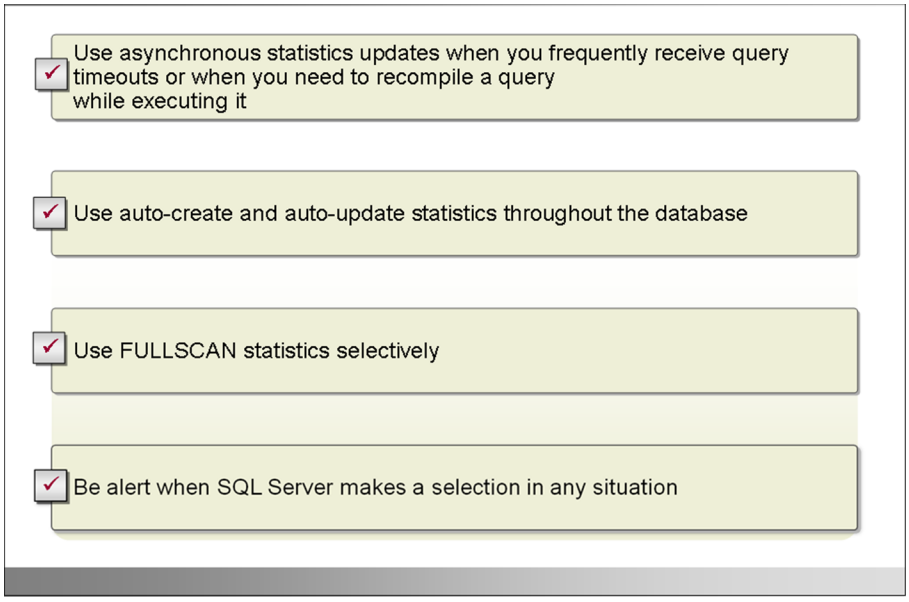
New Horizons Computer Learning Centers, Inc.

5. In the Microsoft SQL Server Management Studio window, in the Query Editor pane, select the **Disable auto update statistics** code, and then click **Execute**.
6. In the Microsoft SQL Server Management Studio window, in the Query Editor pane, select the **Create manual statistics** code, and then click **Execute**.
7. In the Microsoft SQL Server Management Studio window, in the Query Editor pane, select the **Compare the result obtained by running the SET STATISTICS command with the HRDepartment statistics** code, and then click **Execute**.
8. In the Microsoft SQL Server Management Studio window, in the Query Editor pane, select the **Set the statistics xml option on and retrieve records from the HumanResources.Department table** code, and then click **Execute**.
9. In the Microsoft SQL Server Management Studio window, in the Query Editor pane, select the **Insert data into the HumanResources.Department table** code, and then click **Execute**.
10. In the Microsoft SQL Server Management Studio window, in the Query Editor pane, select the **Set the statistics xml option on and retrieve records from the HumanResources.Department table** code, and then click **Execute**.
11. In the Microsoft SQL Server Management Studio window, in the Query Editor pane, select the **Enable auto update statistics, auto create statistics, and auto update statistics async** code, and then click **Execute**.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Best Practices for Optimizing Statistics



Key Points

The goal of using statistics in SQL Server is to allow the optimizer to get good cardinality estimates so that it can find good query plans, while at the same time keeping any overhead or delay associated with gathering statistics reasonable.

Some best practices for managing statistics in SQL Server include:

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

- **Use asynchronous statistics updates when you frequently receive query timeouts or when you need to recompile a query while executing it**. If you have a large database and an online transaction processing (OLTP) workload, and if you enable AUTO_UPDATE_STATISTICS, some transactions that run in a fraction of a second may sometimes take several seconds or more because they cause statistics to be updated. If you want to avoid the possibility of this delay, enable AUTO_UPDATE_STATISTICS_ASYNC. For workloads with long-running queries, getting the best plan is more important than an infrequent delay in compilation. In such cases, use synchronous rather than asynchronous auto update statistics.
- **Use auto-create and auto-update statistics throughout the database**. For a majority of SQL Server installations, the most important best practice is to use auto-create and auto-update statistics database-wide. Auto-create and auto-update statistics are enabled by default. If you observe a query execution plan that is not effective, which may be because of missing or out-of-date statistics are at fault, verify whether you have enabled the auto-create and auto-update statistics.

  When SQL Server creates statistics automatically, the server creates statistics based on a sample of rows from the table by default, instead of performing a full scan. An exception is when SQL Server uses full scan to create statistics as a by-product of index creation. When you update statistics with a different sampling rate from full scan to sampled, or from sampled to full scan, it can cause query plan changes. If query plan stability is a high priority, drop and re-create any statistics with sampling that are the by-product of index creation.
- **Use FULLSCAN statistics selectively**. If you are using auto-create and auto-update statistics and you are getting a query execution plan that is not effective because the statistics are not accurate or current, you can:

  1. Leave auto-create and auto-update statistics on.
  2. Use CREATE STATISTICS … WITH FULLSCAN, NORECOMPUTE and a batch job that does
     UPDATE STATISTICS … WITH FULLSCAN, NORECOMPUTE
     periodically for those statistics that are not accurate or current.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

The frequency with which you should update statistics depends on the application. If the table on which you are performing an update has a high update rate, run a fullscan statistics update every night. If the table has a low update rate, run a fullscan statistics update weekly.

- **Be alert when SQL Server determines how to use statistics in any situation**. SQL Server makes a selection on how to use statistics in a number of situations. Often, the selections are reasonable when the data sizes are small enough, or when the selection does not lead to a query plan that is not effective. When SQL Server must make a selection of a query predicate, sometimes you get a query plan that is not optimal. If a query is not performing as expected, and you suspect that SQL Server has selected a suboptimal query plan, observe the query and its resulting plan for signs that selectivity was made rather than estimate based on statistics. In many cases, you can modify the query or a small part of your application to avoid the selection.

The following table lists some situations that can lead to selections, and possible workarounds.

Situation	Possible workaround
Missing statistics	Verify whether auto create statistics is enabled, or manually create statistics by using CREATE STATISTICS or sp_createstats. Check to see whether the database is read-only. Auto create statistics do not work on read-only databases.
Use of local variables in the query condition	If the query predicate uses a local variable, consider rewriting the query to use a parameter, instead of a local variable. The value of a local variable is not known when the query optimizer creates the query execution plan. When a query uses a parameter, the query optimizer uses the cardinality estimate for the first actual parameter value that is passed to the stored procedure.
Nonconstant-foldable expressions in the	For example, consider the expression, T.x+1 = 0. Rewrite the query to eliminate

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Situation	Possible workaround
query condition	the expression, or evaluate the expression prior to running the query and pass the result into the query as a parameter, not a local variable. In the case of T.x+1 = 0, rewrite the expression as T.x = -1. This is an equivalent expression, and it allows an accurate estimate, instead of a guess.
Complex expressions	Consider the following expression, "Price + Tax > 100" or "Price * (1+TaxRate) > 100". If you see a slow query performance in such a case, you can create a computed column with the equivalent expression, and create statistics or an index on the computed column. Auto create statistics also creates statistics for the computed column if it exists, so you need not create the computed column statistics manually if auto create statistics is enabled.

Question: When you should consider using asynchronous statistics update over automatic statistics update?

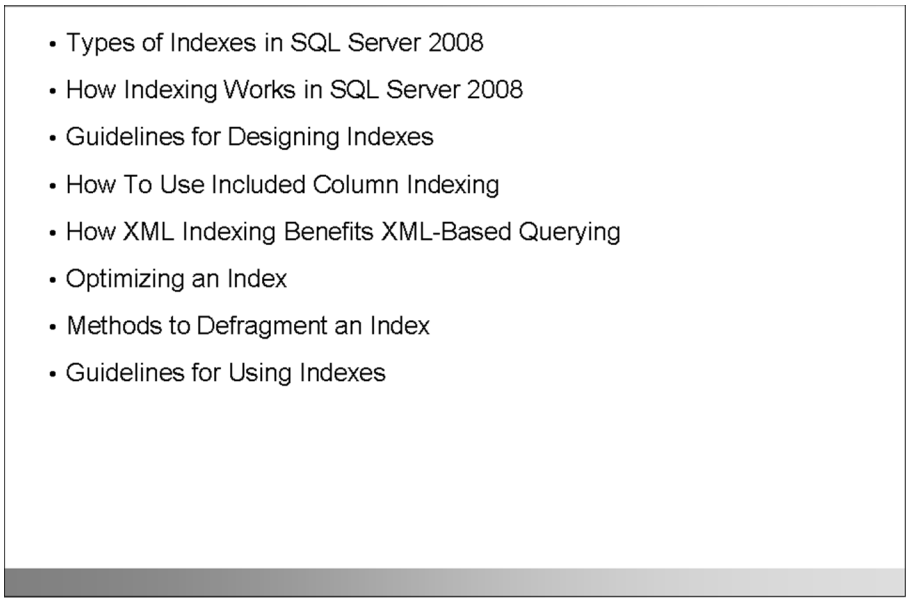
Source:

- http://msdn.microsoft.com/en-us/library/dd535534.aspx
- http://technet.microsoft.com/en-us/library/cc966419.aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Lesson 2
Designing a Strategy for Maintaining Indexes



SQL Server uses indexes to find data quickly when a query is processed, thereby improving the overall performance of the database. You can use indexes to search, sort, and group records. However, indexes require disk space and have to be updated whenever a query modifies the data in a table.

In this lesson, you will learn about various types of indexes, how they should be implemented and used, and the guidelines for using indexing.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Types of Indexes in SQL Server 2008

Type of Index	Description
Clustered Index	Sorts and stores the data rows of the table or view in order, based on the clustered index key
Nonclustered Index	Helps to sort and search information based on values in columns, instead of using a full table scan to locate the information
Spatial Index	Provides the ability to perform certain operations more efficiently on spatial objects in a column of the geometry or geography data type
XML Index	Allows SQL Server to locate information based on values stored inside the XML document
Full-Text Index	Provides ability to perform efficient keyword-based queries in text data
Unique Index	Ensures that the index key contains no duplicate values
Index with Included Columns	Is a nonclustered index that is extended to include nonkey columns in addition to the key columns
Filtered Index	Is an optimized nonclustered index, especially suited to cover queries that select from a well-defined subset of data

Key Points

SQL Server uses different indexes to help improve the query performance. These indexes are designed to cater to different query needs and some are designed to help querying of different data types.

The types of indexes available in SQL Server 2008 include:

- **Clustered Index**. The data rows are stored in order, based on the clustered index key. The clustered index is implemented as a B-tree index structure that supports fast retrieval of the rows, based on their clustered index key values. The pages in each level of the index, including the data pages in the leaf level, are linked in a doubly-linked list. However, navigation from one level to another is performed by using key values.
- **Nonclustered Index**. Nonclustered indexes have a B-tree index structure similar to the one in clustered indexes. The difference is that nonclustered indexes do not affect the order of the data rows. The leaf level contains index rows. Each index row contains

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

the nonclustered key value, a row locator, and any included or nonkey columns. For tables without a clustered index, the locator points to the data row that has the key value. For tables with clustered index, the locator points to the key value in the clustered index. The index helps to sort and search information based on values in columns, instead of using a full table scan to locate the information.

- **Spatial Index**. Spatial index provides the ability to perform certain operations more efficiently on spatial objects (spatial data) in a column of the geometry or geography data type. It reduces the number of objects on which relatively costly spatial operations need to be applied. The geography data type represents geographic objects on an area on the Earth's surface, such as a spread of land. A spatial index on a geography column maps the geographic data to a two-dimensional, non-Euclidean space.

  You can define a spatial index on a table column that contains spatial data. Each spatial index refers to a finite space. For example, an index for a geometry column refers to a user-specified rectangular area on a plane.

- **XML Index**. You can create XML indexes on XML data type columns. XML index indexes all tags, values, and paths over the XML instances in the columns, which benefit query performance. It allows SQL Server to locate information based on values stored inside the XML document.
  XML indexes fall into the following categories:

- **Primary XML index**. The first index on the XML type column must be the primary XML index. By using the primary XML index, SQL Server supports the PATH, VALUE, and PROPERTY secondary indexes. Depending on the type of queries, these secondary indexes might help improve query performance.

- **Secondary XML index**. XML instances are stored in XML type columns as large binary objects (BLOBs). These XML instances can be large, and the stored binary representation of XML data type instances can be up to 2 GB. Without an index, SQL Server would need to process the whole XML stored in BLOB, which can be time consuming. For example, consider the following query.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.


```
WITH XMLNAMESPACES
('http://schemas.microsoft.com/sqlserver/2004/07/Quanta
mCorp/ProductModelDescription' AS "PD")

SELECT CatalogDescription.query('

  /PD:ProductDescription/PD:Summary

') as Result

FROM Production.ProductModel

WHERE CatalogDescription.exist
('/PD:ProductDescription/@ProductModelID[.="19"]') = 1
```

To select the XML instances that satisfy the condition in the WHERE clause, the XML BLOB in each row of table Production.ProductModel is shredded at run time. Then, the expression (/PD:ProductDescription/@ProductModelID[.="19"]) in the exist() method is evaluated. This run-time shredding can be costly, depending on the size and number of instances stored in the column.

If querying XML BLOBs is common in your application environment, it helps to index the XML type columns. However, there is a cost associated with maintaining the index during data modification.

- **Full-Text Index**. A full-text index is a special type of token-based functional index that is built and maintained by the Microsoft Full-Text Engine for SQL Server. It provides efficient support for sophisticated word searches in character string data.

  Creating and maintaining a full-text index involves populating the index by using a process called a population, also known as a crawl. SQL Server supports the following types of population: full population, change tracking-based automatic or manual population, and incremental timestamp-based population.

  A table or a view can have only one clustered index, whereas a table or a view can have multiple nonclustered indexes, spatial indexes, and XML indexes, and full-text indexes.

- **Unique Index**. A unique index ensures that the index key contains no duplicate values. Therefore, every row in the table or view is in some way unique. Both clustered and nonclustered indexes can be unique.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

You can specify a unique index when the data itself is unique. For example, if you want to ensure that the values in the NationalIDNumber column in the HumanResources.Employee table are unique, when the primary key is EmployeeID, create a UNIQUE constraint on the NationalIDNumber column. If the user tries to type the same value in that column for more than one employee, an error message is displayed and the duplicate value is not entered.

With multicolumn unique indexes, the index ensures that each combination of values in the index key is unique. For example, if you create a unique index on a combination of LastName, FirstName, and MiddleName columns, no two rows in the table could have the same combination of values for these columns.

The benefits of unique indexes include the following:

- Data integrity of the defined columns is ensured.
- Additional information helpful to the query optimizer is provided.

When you create a PRIMARY KEY or UNIQUE constraint, a unique index is automatically created on the specified columns. There are no significant differences between creating a UNIQUE constraint and creating a unique index independent of a constraint. Data validation occurs in the same manner and the query optimizer does not differentiate between a unique index created by a constraint or manually created. However, you should create a UNIQUE or PRIMARY KEY constraint on the column when data integrity is the objective.

Index with Included Column

Index with included column is a nonclustered index that is extended to include nonkey columns in addition to the key columns. You can extend the functionality of nonclustered indexes by adding nonkey columns to the leaf level of the nonclustered index. By including nonkey columns, you can create nonclustered indexes that cover more queries. This is because the nonkey columns have the following benefits:

- They can be data types not allowed as index key columns.
- They are not considered by the Database Engine when calculating the number of index key columns or index key size.

An index with included nonkey columns can significantly improve query performance when all columns in the query are included in the

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

index, either as key or nonkey columns. Performance gains are achieved because the query optimizer can locate all the column values within the index and the table or clustered index data is not accessed, resulting in fewer disks I/O operations.

Filtered Index

A filtered index is an optimized nonclustered index, especially suited to cover queries that select from a well-defined subset of data. It uses a filter predicate to index a portion of rows in the table. A well-designed filtered index can improve query performance, reduce index maintenance costs, and reduce index storage costs compared with full-table indexes.

Filtered indexes can provide the following advantages over full-table indexes:

- **Improved query performance and plan quality**. A well-designed filtered index improves query performance and execution plan quality because it is smaller than a full-table nonclustered index and has filtered statistics. The filtered statistics are more accurate than full-table statistics because they cover only the rows in the filtered index.
- **Reduced index maintenance costs**. An index is maintained only when data manipulation language (DML) statements affect the data in the index. A filtered index reduces index maintenance costs compared with a full-table nonclustered index because it is smaller and is only maintained when the data in the index is affected. It is possible to have a large number of filtered indexes, especially when they contain data that is affected infrequently. Similarly, if a filtered index contains only the frequently affected data, the smaller size of the index reduces the cost of updating the statistics.
- **Reduced index storage costs**. Creating a filtered index can reduce disk storage for nonclustered indexes when a full-table index is not necessary. You can replace a full-table nonclustered index with multiple filtered indexes without significantly increasing the storage requirements.

Relationship Between Indexes

The clustered index controls how data is actually organized in a table and how it is stored. The other indexes help to speed up the process of looking up the key of clustered index for a particular type of data.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Situations for Using Clustered and Nonclustered Index

Before you create clustered indexes, you need to understand how your data will be accessed. Consider using a clustered index for queries that do the following:

- Return a range of values by using operators such as BETWEEN, >, >=, <, and <=.
  After the row with the first value is found by using the clustered index, rows with subsequent indexed values are guaranteed to be physically adjacent.
- Return large result sets.
- Use JOIN clauses; typically these are foreign key columns.
- Use ORDER BY, or GROUP BY clauses.

Before you create nonclustered indexes, you should understand how your data will be accessed. Consider using a nonclustered index for queries that have the following attributes:

- Use JOIN or GROUP BY clauses.
- Queries that do not return large result sets.
- Contain columns frequently involved in search conditions of a query, such as the WHERE clause, that returns exact matches.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Question: What is the difference between clustered index and nonclustered index?

Question: When should you use XML index?

Question: When should you consider using included column?

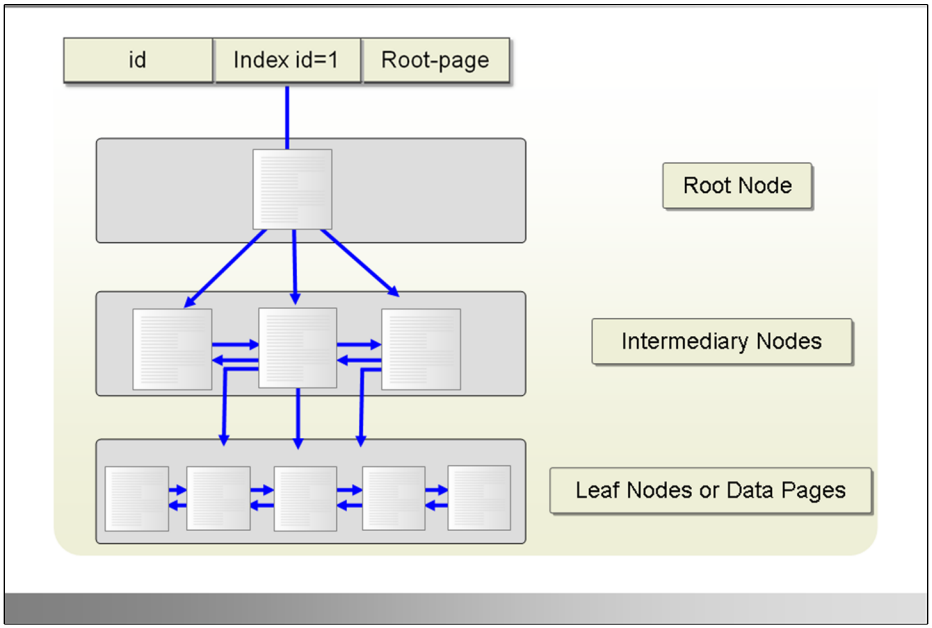
Source:

- http://msdn.microsoft.com/en-us/library/ms191497.aspx
- http://msdn.microsoft.com/en-us/library/ms142575.aspx
- http://msdn.microsoft.com/en-us/library/ms175049.aspx
- http://technet.microsoft.com/en-us/library/ms190806.aspx
- http://technet.microsoft.com/en-us/library/cc280372.aspx
- http://msdn.microsoft.com/en-us/library/ms190639.aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

How Indexing Works in SQL Server 2008



Key Points

SQL Server uses indexes to find data when a query is processed. Indexes speed up the querying process by providing quick access to rows in the data tables. You should a clear understanding of how SQL Server organizes indexes. This will help you determine what kind of indexes can be useful for different kind of queries.

B-tree structure of a SQL Server index

An index is made up of a set of pages or index nodes that are organized in a B-tree structure. A B-tree is constructed with:

- **Root_page or Index Node**. Each page in an index B-tree is called an index node.
- **Root node**. The top node of the B-tree is called the root node.
- **Leaf Node or Data Pages**. The bottom level of nodes in the index is called the leaf nodes or data pages.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

- **Intermediary Levels**. Index levels between the root and the leaf nodes are collectively known as intermediate levels.

Clustered Index Structure

In a clustered index, the leaf nodes contain the data pages of the underlying table. The root and intermediate level nodes contain index pages holding index rows. Each index row contains a key value and a pointer to either an intermediate level page in the B-tree, or a data row in the leaf level of the index.

Clustered indexes have one row in sys.partitions, with index_id = 1 for each partition used by the index. By default, a clustered index has a single partition. When a clustered index has multiple partitions, each partition has a B-tree structure that contains the data for that specific partition. For example, if a clustered index has four partitions, there are four B-tree structures; one in each partition.

Depending on the data types in the clustered index, each clustered index structure will have one or more allocation units in which to store and manage the data for a specific partition. At a minimum, each clustered index will have one IN_ROW_DATA allocation unit per partition. The clustered index will also have one LOB_DATA allocation unit per partition if it contains large object (LOB) columns. It will also have one ROW_OVERFLOW_DATA allocation unit per partition if it contains variable length columns that exceed the 8,060-byte row size limit.

The pages in the data chain and the rows in them are ordered on the value of the clustered index key. All inserts are made at the point where the key value in the inserted row fits in the ordering sequence among existing rows. The page collections for the B-tree are anchored by page pointers in the sys.system_internals_allocation_units system view.

For a clustered index, the root_page column in sys.system_internals_allocation_units points to the top of the clustered index for a specific partition. SQL Server moves down the index to find the row corresponding to a clustered index key. To find a range of keys, SQL Server moves through the index to find the starting key value in the range and then scans through the data pages, using the previous or next pointers. To find the first page in the chain of data pages, SQL Server follows the leftmost pointers from the root node of the index.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Nonclustered Index Structure

Nonclustered indexes have the same B-tree structure as clustered indexes, except for the following significant differences:

- The data rows of the underlying table are not sorted and stored in order based on their nonclustered keys.
- The leaf layer of a nonclustered index is made up of index pages, instead of data pages.

You can define nonclustered indexes on a table or view with a clustered index or a heap. Each index row in the nonclustered index contains the nonclustered key value and a row locator. This locator points to the data row in the clustered index or heap having the key value.

The row locators in nonclustered index rows are either a pointer to a row or are a clustered index key for a row, as described in the following:

- If the table is a heap, which means it does not have a clustered index, the row locator is a pointer to the row. The pointer is built from the file identifier (ID), page number, and number of the row on the page. The pointer is known as a row ID (RID).
- If the table has a clustered index, or the index is on an indexed view, the row locator is the clustered index key for the row. If the clustered index is not a unique index, SQL Server makes any duplicate keys unique by adding an internally generated value called an uniqueifier. This four-byte value is not visible to users. It is only added, when required, to make the clustered key unique for use in nonclustered indexes. SQL Server retrieves the data row by searching the clustered index by using the clustered index key stored in the leaf row of the nonclustered index.

Nonclustered indexes have one row in sys.partitions with index_id >0 for each partition used by the index. By default, a nonclustered index has a single partition. When a nonclustered index has multiple partitions, each partition has a B-tree structure that contains the index rows for that specific partition. For example, if a nonclustered index has four partitions, there are four B-tree structures, with one in each partition. Depending on the data types in the nonclustered index, each nonclustered index structure will have one or more allocation units in which to store and manage the data for a specific

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

partition. At a minimum, each nonclustered index will have one IN_ROW_DATA allocation unit per partition that stores the index B-tree pages. The nonclustered index will also have one LOB_DATA allocation unit per partition if it contains large object (LOB) columns. In addition, it will have one ROW_OVERFLOW_DATA allocation unit per partition if it contains variable length columns that exceed the 8,060-byte row size limit. The page collections for the B-tree are anchored by root_page pointers in the sys.system_internals_allocation_units system view.

Question: What is the difference in the architecture of clustered index and nonclustered index?

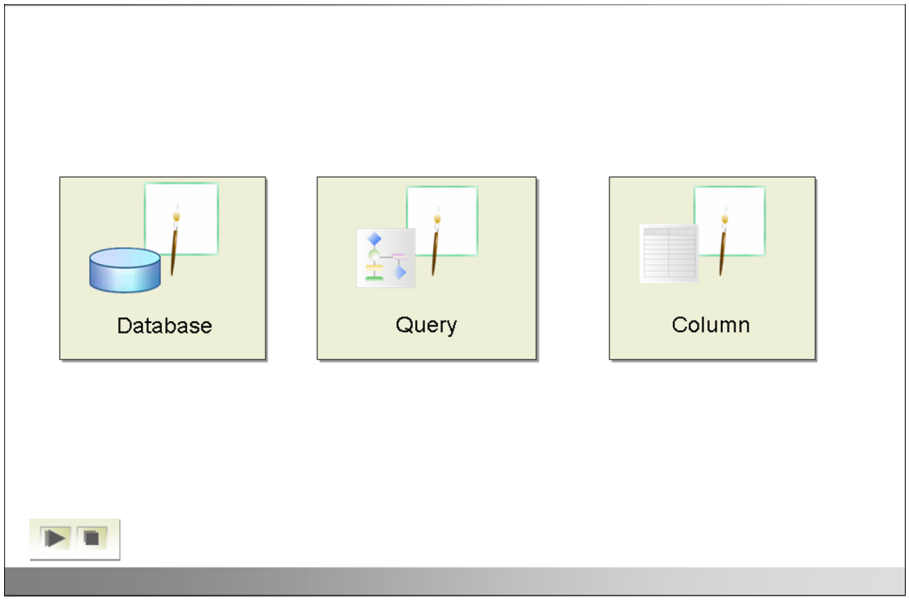
Source:

- http://msdn.microsoft.com/en-us/library/ms177443.aspx
- http://msdn.microsoft.com/en-us/library/ms177484.aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Guidelines for Designing Indexes



Key Points

Well-designed indexes increase query performance and response time without affecting the design of the underlying tables. However, designing an index is very complex, time-consuming, and error-prone process, even for moderately complex databases and workloads. As a database designer, you should design optimal indexes by understanding the characteristics of your database, queries, and data columns.

Database Guidelines

Index helps to improve query performance but may introduce additional storage requirements and management overhead to the database. In order to ensure the index designed would not affect the database; you should consider the following database guidelines:

- **Avoid large numbers of indexes on a table**. Large number of indexes in a table affects the performance of INSERT, UPDATE, and DELETES statements. In addition, all indexes must be adjusted appropriately as data in the table changes.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

- **Avoid over-indexing frequently updated tables**. You should avoid over-indexing frequently updated tables and keep indexes narrow, that is, with as few columns as possible. If you include more columns in an index more time is required to make an adjustment on the index when data changes.
- **Use large number of indexes to improve query performance on tables with low update requirements, but large volumes of data**. Large numbers of indexes can help to improve the performance of queries that do not modify data, such as SELECT statements. This is because the query optimizer has more indexes to choose from to determine the fastest access method.
- **Avoid using indexes on small tables**. Indexing small tables may not be optimal because it can take the query optimizer longer to traverse the index searching for data than to perform a simple table scan. Therefore, indexes on small tables might never be used, but must still be maintained as data in the table changes.
- **Use indexes on views**. Indexes on views can provide significant performance gains when the view contains aggregations, table joins, or a combination of aggregations and joins. The view does not have to be explicitly referenced in the query for the query optimizer to use it.
- **Use the Database Engine Tuning Advisor (DTA) to analyze your database**. The Database Engine Tuning Advisor tool helps you to tune databases and improve query processing. Database Engine Tuning Advisor examines how queries are processed in the databases. In addition, it suggests how you can improve query processing performance by modifying physical design structures such as indexes, indexed views, and partitioning

Query Guidelines

One of the key design objectives for implementing index in a database is to improve the query performance. In order to ensure optimal performance for the query, you should consider the following query guidelines when designing indexes:

- **Create nonclustered indexes on columns that are frequently used in predicates and join conditions in**

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

queries. This improves the performance of the query when performing the join operation by reducing the time required to locate the required records in each table.

- **Cover indexes to improve query performance**. Covering indexes can improve query performance because all the data needed to meet the requirements of the query exists within the index itself; that is, only the index pages, and not the data pages of the table or clustered index, are required to retrieve the requested data, therefore reducing overall disk I/O.
- **Write queries that insert or modify maximum number of rows in a single statement**. You should write queries that insert or modify as many rows as possible in a single statement instead of using multiple queries to update the same rows. By using only one statement, you can reduce the chances of a fragmented index.
- **Evaluate the query type and use of columns in the query**. For example, a column used in an exact-match query type can be a good candidate for a nonclustered or clustered index, or the use of LIKE operator usually mean the normal index may not able to help in locating the information

Column Guidelines

Index is created by storing values of the data into a column which helps SQL Server to quickly locate the information. When you design an index, consider the following guidelines:

- **Keep the length of the index key short for clustered indexes**. This is to ensure that the indexes created are optimal. Also, the keys of clustered index would be used by non-clustered index as well, thus the longer the data size, the large the index would be. Additionally, clustered indexes benefit from being created on unique or non-null columns.
- **Do not specify columns that are of the ntext, text, image, varchar(max), nvarchar(max), and varbinary(max) data types as index key columns**. The data that are not stored inside the data row and the values obtained by SQL server when creating index are actually the pointer of the data which makes the index unable to use. However, varchar(max), nvarchar(max), varbinary(max), and XML data types can participate in a nonclustered index as nonkey index columns.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

- **Use an XML data type as a key column only in an XML index**. Standards indexes such as clustered or nonclustered are designed to handle standard data type, XML data type is a special data type which requires special index structure and thus should consider using XML index instead of standard index
- **Examine column uniqueness**. A unique index instead of a nonunique index on the same combination of columns provides additional information for the query optimizer that makes the index more useful.
- **Examine data distribution in the column**. Frequently, a long-running query is caused by indexing a column with few unique values, or by performing a join on such a column. This is a fundamental problem with the data and query, and generally cannot be resolved without identifying this situation. For example, a physical telephone directory sorted alphabetically on last name will not expedite locating a person if all people in the city are named Smith or Jones.
- **Use filtered indexes on columns that have well-defined subsets**. For example sparse columns, columns with mostly NULL values, columns with categories of values, and columns with distinct ranges of values. A well-designed filtered index can improve query performance, reduce index maintenance costs, and reduce storage costs.
- **Consider the order of the columns if the index contains multiple columns**. The column that is used in the WHERE clause in an equal to (=), greater than (>), less than (<), or BETWEEN search condition, or participates in a join, should be placed first. Additional columns should be ordered based on their level of distinctness, that is, from the most distinct to the least distinct. For example, if the index is defined as LastName, FirstName, the index will be useful when the search criterion is WHERE LastName = 'Smith' or WHERE LastName = Smith AND FirstName LIKE 'J%'. However, the query optimizer would not use the index for a query that searched only on FirstName (WHERE FirstName = 'Jane').

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Question: Your application database has a number of text-based columns and a integer primary key column. On which column should you set up the clustered index?

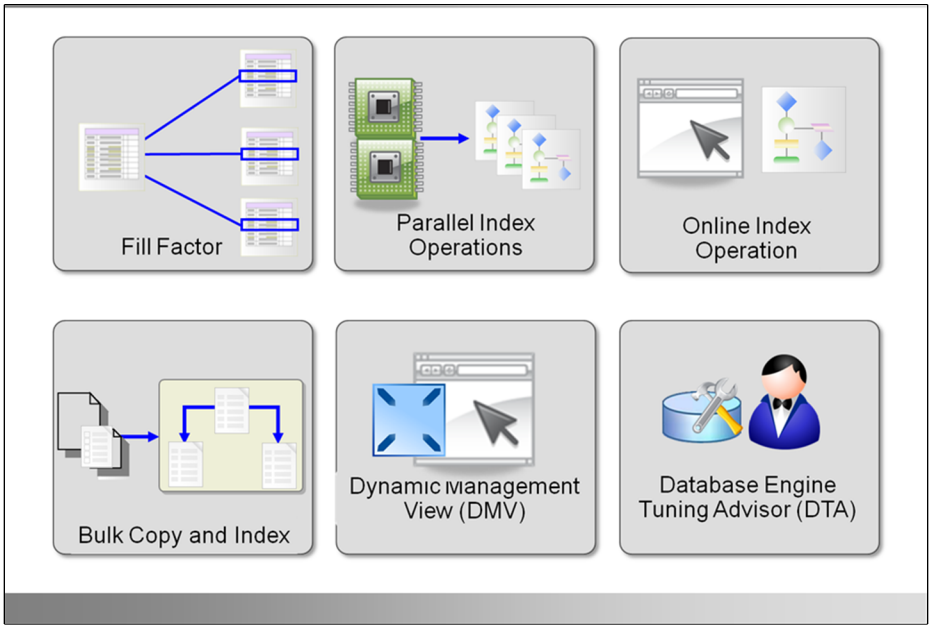
Question: What should you consider when performing a search with multiple values?

Source:

- http://msdn.microsoft.com/en-us/library/ms166575%28SQL.90%29.aspx
- http://msdn.microsoft.com/en-us/library/ms191195.aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

Optimizing an Index



Key Points

A key factor in achieving optimal query performance is by ensuring that you create and then maintain indexes. In addition, an index helps you to reduce the disk I/O and CPU resources consumption.

Some of the factors to optimize an index are:

- **Fill Factor**. You can use the fill-factor option for fine-tuning the index data storage and performance. When you create or rebuild an index, the fill-factor value determines the percentage of space on each leaf-level page to be filled with data, reserving the remainder on each page as free space for future growth. For example, specifying a fill-factor value of 80 means that 20 percent of each leaf-level page will be left empty, providing space for index expansion as data is added to the underlying table. The empty space is available on each index page, rather than at the end of the index. The fill-factor value is a percentage from 1 to 100, and the server-wide default is 0, which means that the leaf-level pages are filled to capacity. PAD_INDEX specifies the open

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

space on each page in the intermediate levels of the index. It is useful only when FILLFACTOR is specified, because PAD_INDEX uses the percentage specified by FILLFACTOR. By default, SQL Server ensures that each index page has enough empty space to accommodate at least one row of the maximum index size, given the set of keys on the intermediate pages. However, if the percentage specified for FILLFACTOR is not large enough to accommodate one row, SQL Server internally overrides the percentage to allow the minimum.

- **Parallel Index Operations**. On multiprocessor computers that are running SQL Server Enterprise, index statements may use multiple processors to perform the scan, sort, and index operations associated with the index statement just like other queries do. The number of processors used to run a single index statement is determined by the configuration option max degree of parallelism, the current workload, and the index statistics. The max degree of parallelism option determines the maximum number of processors to use in parallel plan execution. If the SQL Server Database Engine detects that the system is busy, the degree of parallelism of the index operation is automatically reduced before statement execution starts. The Database Engine can also reduce the degree of parallelism if the leading key column of a nonpartitioned index has a limited number of distinct values, or if the frequency of each distinct value varies significantly.

  The number of processors that are used by the query optimizer typically provide optimal performance. However, operations such as creating, rebuilding, or dropping very large indexes are resource intensive and can cause insufficient resources for other applications and database operations for the duration of the index operation. When this problem occurs, you can manually configure the maximum number of processors that are used to run the index statement by specifying the MAXDOP index option and limiting the number of processors to use for the index operation.

  The MAXDOP index option overrides the max degree of parallelism configuration option only for the query specifying this option. The following table lists the valid integer values that can be specified with the max degree of parallelism configuration option and the MAXDOP index option.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Value	Description
0	Uses the actual number of available CPUs, depending on the current system workload. This is the default value and the setting must be used.
1	Suppresses parallel plan generation. The operation will be run serially.
2-64	Limits the number of processors to the specified value. Fewer processors may be used, depending on the current workload. If a value larger than the number of available CPUs is specified, the actual number of available CPUs is used.

Parallel index execution and the MAXDOP index option apply to the following
T-SQL statements:

- CREATE INDEX
- ALTER INDEX REBUILD
- DROP INDEX (This applies to clustered indexes only.)
- ALTER TABLE ADD (index) CONSTRAINT
- ALTER TABLE DROP (clustered index) CONSTRAINT

All semantic rules that are used with max degree of parallelism configuration option are applicable when the MAXDOP index option is used.

When you run ALTER INDEX REORGANIZE with or without LOB_COMPACTION, the max degree of parallelism value is a single threaded operation. You cannot specify the MAXDOP index option in the ALTER INDEX REORGANIZE statement.

- **Index Operations Online**. You can create, rebuild, or drop indexes online. The ONLINE option allows concurrent user access to the underlying table or clustered index data and any associated nonclustered indexes during these index operations. For example, while a clustered index is being rebuilt by one user, that user and others can continue to update and query the underlying data. When you perform dynamic-link library (DDL) operations offline, such as building or rebuilding a clustered

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

index, these operations hold exclusive locks on the underlying data and associated indexes.

You should perform online index operations for business environments that operate 24 hours a day, seven days a week, in which the need for concurrent user activity during index operations is vital.

The ONLINE option is available in the following T-SQL statements :

- CREATE INDEX
- ALTER INDEX
- DROP INDEX
- ALTER TABLE (To add or drop UNIQUE or PRIMARY KEY constraints with CLUSTERED index option)

- **Bulk Copy and Indexes**. If your table has a clustered index and the data to be inserted is ordered to match the clustered index key, bulk copy the data into the table with the clustered index already in place and specify the ORDER hint. This approach is significantly faster than creating the clustered index after the data is copied into the table.

  If nonclustered indexes are also present in the table, drop these before copying data into the table. It is generally faster to bulk copy data into a table without nonclustered indexes and then to re-create the nonclustered indexes, instead of bulk copying data into a table with the nonclustered indexes in place.

- **Dynamic management view (DMV)**. The DMVs sys.dm_db_missing_index_details and sys.dm_db_missing_index_columns return results that indicate whether an index key column is an equality, inequality, or included column. The result set of sys.dm_db_missing_index_details returns this information in the equality_columns, inequality_columns, and included_columns columns. The result set returned by sys.dm_db_missing_index_columns returns this information in its column_usage column.

  Run the following query against the sys.dm_db_missing_index_details dynamic management view.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.


```
SELECT *

FROM sys.dm_db_missing_index_details
```

- **Database Engine Tuning Advisor**. Microsoft SQL Server Database Engine Tuning Advisor helps you select and create an optimal set of indexes, indexed views, and partitions without requiring an expert understanding of the structure of the database or the internals of Microsoft SQL Server.

  Using the Database Engine Tuning Advisor, you can:

  - Identify the best mix of indexes for databases by using the query optimizer to analyze queries in a workload.
  - Identify indexed views for databases referenced in a workload.
  - Analyze the effects of the proposed changes, including index usage, query distribution among tables, and query performance in the workload.
  - Determine ways to tune the database for a small set of problem queries.

Disabling an Index

Disabling an index prevents user access to the index, and for clustered indexes, to the underlying table data. The index definition remains in metadata and index statistics are kept on nonclustered indexes. Disabling a nonclustered index or clustered index on a view physically deletes the index data. Disabling a clustered index on a table prevents access to the data; the data still remains in the table, but is unavailable for data manipulation language (DML) operations until the index is dropped or rebuilt. To rebuild and enable a disabled index, use the ALTER INDEX REBUILD statement or the CREATE INDEX WITH DROP_EXISTING statement.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Question: For OLTP databases, what should be the commonly used fill-factor values for indexes?

Question: When should you disable an index temporarily?

Source:

- http://msdn.microsoft.com/en-us/library/ms177459.aspx
- http://technet.microsoft.com/en-us/library/ms189329.aspx
- http://msdn.microsoft.com/en-us/library/aa258260%28SQL.80%29.aspx
- http://msdn.microsoft.com/en-us/library/ms177442.aspx
- http://msdn.microsoft.com/en-us/library/ms345405.aspx

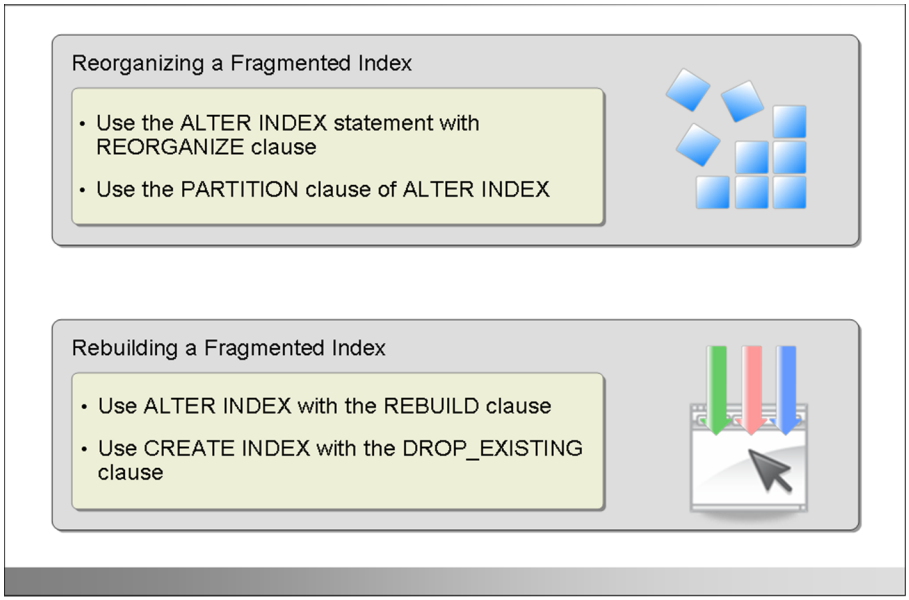
For more information, see Guidelines for Disabling Indexes.

For more information, see Guidelines for Enabling Indexes and Constraints.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Methods to Defragment an Index



Key Points

The SQL Server Database Engine automatically maintains indexes whenever you perform the UPDATE or DELETE operations to the underlying data. Over time, these modifications can cause the information in the index to become fragmented in the database. Fragmentation exists when indexes have pages in which the logical ordering, based on the key value, does not match the physical ordering inside the data file. Heavily fragmented indexes can degrade query performance and cause your application to respond slowly.

Methods to Defragment an Index

You can use following methods to remove index fragmentation:

- **Reorganizing a fragmented index**. The reorganize Index task reorganizes indexes in SQL Server database tables and views. To reorganize an index you can:
- **Use the ALTER INDEX statement with REORGANIZE clause**. To reorganize one or more indexes, use the ALTER INDEX

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

statement with the REORGANIZE clause. This statement replaces the DBCC INDEXDEFRAG statement.

- **Use the PARTITION clause of ALTER INDEX**. To reorganize a single partition of a partitioned index, use the PARTITION clause of ALTER INDEX.
- **Rebuilding a fragmented index**. Rebuilding an index drops the index and creates a new one. To rebuild clustered and nonclustered index you can:
- **Use ALTER INDEX with the REBUILD clause**. This statement replaces the DBCC DBREINDEX statement.
- **Use CREATE INDEX with the DROP_EXISTING clause**. This statement removes the existing index and creates a new one.

Advantages of Index Defragmentation Methods

Each index fragmentation method has its own advantages. Reorganizing a fragmented index helps to:

- Improve index-scanning performance by reordering the leaf-level pages to match the logical order from left to right of the leaf nodes. The index is reorganized within the existing pages allocated to it, and new pages are not allocated. If an index spans more than one file, the files are reorganized one at a time. Pages do not migrate between files.
- Reduce the size of the index pages. Any empty pages created by this compaction are removed, providing additional available disk space. Compaction is based on the fill-factor value in the sys.indexes catalogue view. The reorganizing process uses minimal system resources. Also, reorganizing is automatically performed online. The process does not hold long-term blocking locks; therefore, it will not block running queries or updates.

Rebuilding a fragmented index helps to:

- Remove fragmentation
- Reclaim disk space by reducing the size of the pages using the specified or existing fill-factor setting,
- Reorder index rows in contiguous pages by allocating new pages as needed. This allocation can improve disk performance by reducing the number of page reads required to obtain the requested data.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Determining the Index Defragmentation Methods

You can reorganize an index when the index is not heavily fragmented. However, if the index is heavily fragmented, you will achieve better results by rebuilding the index. You can rebuild an index when the index is heavily fragmented. For partitioned indexes built on a partition scheme, you can use either of these methods on a complete index or on a single partition of an index.

The first step in deciding which defragmentation method to use is to analyze the index to determine the degree of fragmentation. By using the system function sys.dm_db_index_physical_stats, you can detect fragmentation in a specific index, all indexes in a table or indexed view, all indexes in a database, or all indexes in all databases.

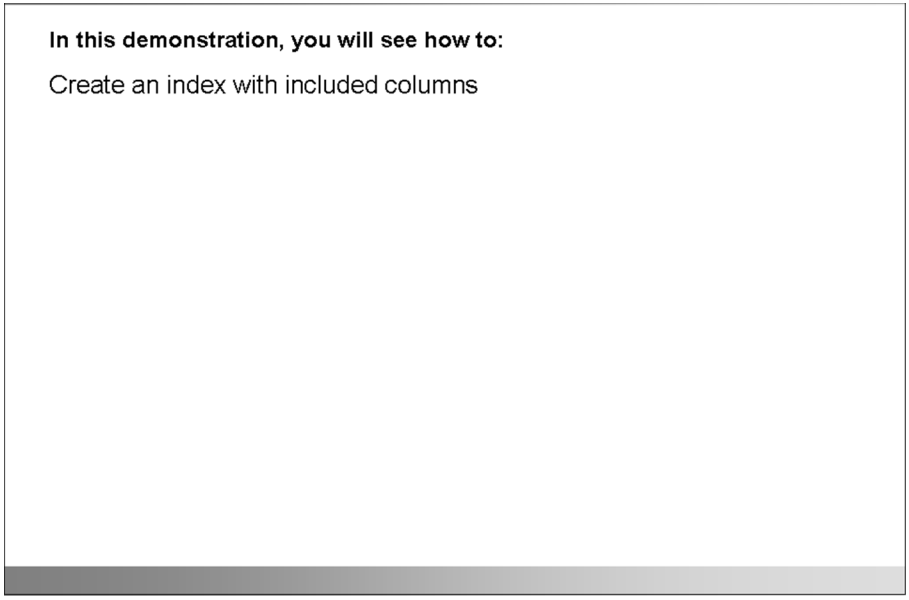
Source:

- http://technet.microsoft.com/en-us/library/ms189858.aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Demonstration: How To Use Included Column Indexing



Key Points

The steps to create an index with included columns are:

1. In the Microsoft SQL Server Management Studio window, in the Query Editor pane, click **Display Estimated Execution Plan**, and then click on the **Execution plan** tab, click **Include Actual Execution Plan**.
2. In Windows Explorer, browse to the **D:\Demofiles\Mod04** folder, and then double-click the **M4_Demo2.sql** file.
3. In the Microsoft SQL Server Management Studio window, in the Query Editor pane, select the **Retrieve the address fields from the Person.Address table** code, and then click **Execute**.
4. In the Microsoft SQL Server Management Studio window, in the Query Editor pane, select the **Drop the IX_Address_PostalCode index, if it exists and then create it** code, and then click **Execute**.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

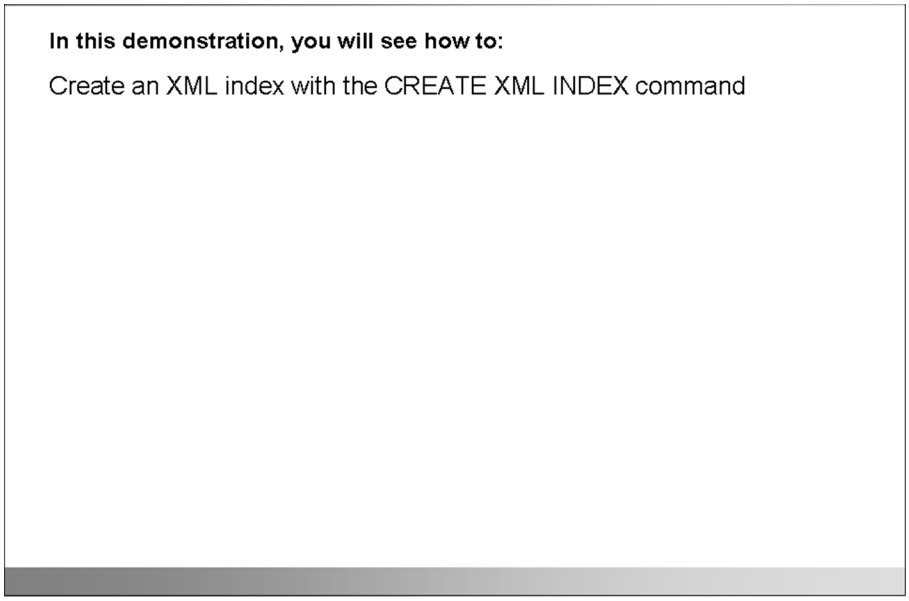
New Horizons Computer Learning Centers, Inc.

5. In the Microsoft SQL Server Management Studio window, in the Query Editor pane, select the **Retrieve the address fields from the Person.Address table** code, and then click **Execute**.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Demonstration: How XML Indexing Benefits XML-Based Querying



Key Points

The steps to create an XML index with the CREATE XML INDEX command are:

1. In the Microsoft SQL Server Management Studio window, in the Query Editor pane, click **Display Estimated Execution Plan**, and then click on the **Execution plan** tab, click **Include Actual Execution Plan**.
2. In Windows Explorer, browse to the **D:\Demofiles\Mod04** folder, and then double-click the **M4_Demo3.sql** file.
3. In the Microsoft SQL Server Management Studio window, in the Query Editor pane, select the **Retrieve product id, product description, and product model id from the Production.ProductModel table** code, and then click **Execute**.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

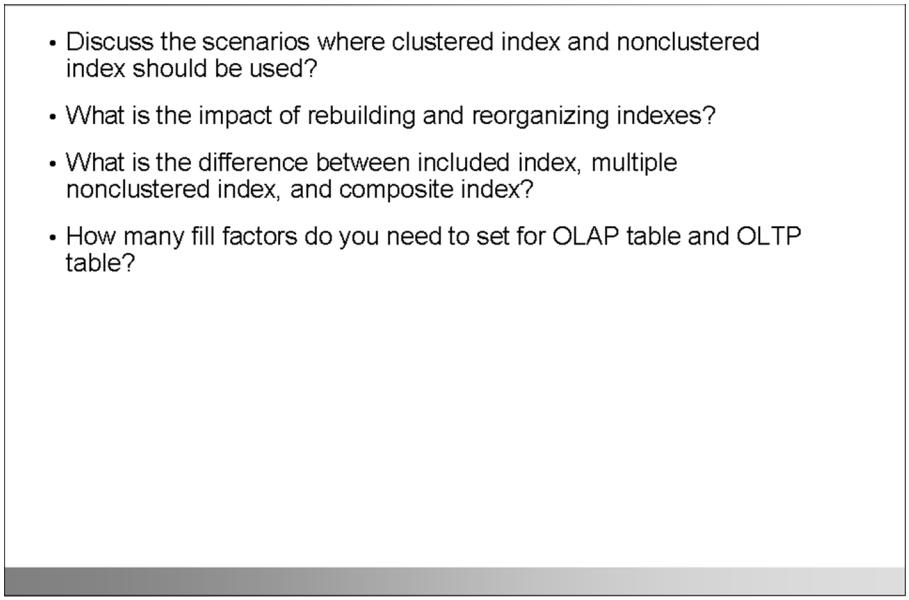
New Horizons Computer Learning Centers, Inc.

4. In the Microsoft SQL Server Management Studio window, in the Query Editor pane, select the **Create a primary XML index** code, and then click **Execute**.
5. In the Microsoft SQL Server Management Studio window, in the Query Editor pane, select the **Create a secondary XML index** code, and then click **Execute**.
6. In the Microsoft SQL Server Management Studio window, in the Query Editor pane, select the **Retrieve product id, product description, and product model id from the Production.ProductModel table** code, and then click **Execute**.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Discussion: Guidelines for Using Indexes



Key Points

In this discussion, you are asked to discuss the guidelines for using indexes. You and your classmates will discuss possible solutions to the scenarios.

Question: Discuss the scenarios where clustered index and nonclustered index should be used?

Question: What is the impact of rebuilding and reorganizing indexes?

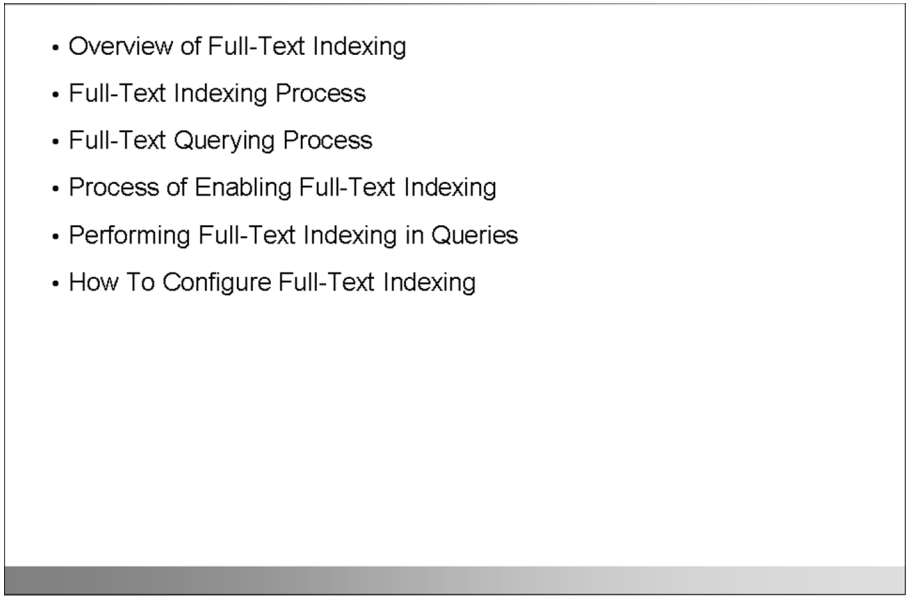
Question: What is the difference between included index, multiple nonclustered indexes, and composite index?

Question: How many fill factors do you need to set for an OLAP table and an OLTP table?

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Lesson 3
Planning for Full-Text Indexing



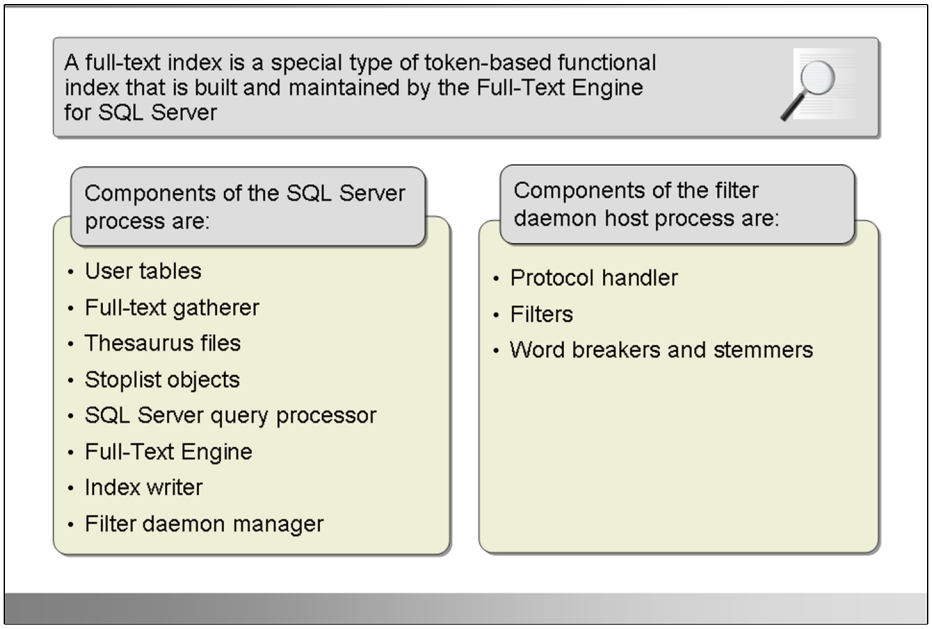
SQL Server 2008 helps applications and users run full-text queries against character-based data in SQL Server tables. Full-text search is applicable in a wide range of business scenarios. Before you run full-text queries on a given table, you must create a full-text index on the table. The full-text index includes one or more character-based columns in the table. Each full-text index, indexes one or more columns from the base table, and each column can have a specific language.

In this lesson, you will learn about the full-text indexing. In addition, you will learn about the process of full-text querying and performing full-text indexing in queries. Finally, you will also learn how to configure full-text indexing.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Overview of Full-Text Indexing



Key Points

A full-text index stores information about significant words and their locations within one or more columns of a database table. A full-text index is a special type of token-based functional index that is built and maintained by the Full-Text Engine for SQL Server. The process of building a full-text index differs from building other types of indexes. Instead of constructing a B-tree structure based on a value stored in a particular row, the Full-Text Engine builds an inverted, stacked, compressed index structure based on individual tokens from the text being indexed.
You can use full-text indexing on columns with the char, varchar, nchar, nvarchar, text, ntext, image, XML, varbinary, or varbinary(max) data types.

In SQL Server 2008, the size of a full-text index is limited only by the available memory resources of the computer on which the instance of SQL Server is running.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

In SQL Server 2008, the full-text indexes are integrated with the Database Engine, instead of being included in the file system, as in SQL Server 2000 and SQL Server 2005. For a new database, the full-text catalog is a virtual object that does not belong to any filegroup. It is a logical concept that refers to a group of full-text indexes.

Architecture of Full-Text Search

The full-text search architecture consists of the following processes:

- **The SQL Server process**. The SQL Server process is also known as the sqlservr.exe process.
  The following table lists the components of the SQL Server process.

SQL Server process	Component
User tables	These tables contain the data to be full-text indexed.
Full-text gatherer	The full-text gatherer works with the full-text crawl threads. It is responsible for scheduling and driving the population of full-text indexes, and for monitoring full-text catalogs. **Note**: In SQL Server 2008, a full-text catalog is a virtual object and does not belong to any filegroup. A full-text catalog is a logical concept that refers to a group of full-text indexes.
Thesaurus files	These files contain synonyms of search terms.
Stoplist objects	Stoplist objects contain a list of common words that are not useful for the search. **Note**: Stoplist objects replace the noise word files of SQL Server 2005 and earlier versions.
SQL Server query processor	The query processor compiles and executes SQL queries. If a SQL query includes a full-text search query, the query is sent to the Full-Text Engine, both during compilation and during execution. The query result is matched against the full-text index.
Full-Text Engine	The Full-Text Engine in SQL Server is now fully integrated with the query processor. The Full-

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

SQL Server process	Component
	Text Engine compiles and runs full-text queries. As part of the query execution, the Full-Text Engine might receive inputs from the thesaurus and stoplist. In SQL Server 2008, the Full-Text Engine for SQL Server runs inside the SQL Server query processor.
Index writer (indexer)	The index writer builds the structure that is used to store the indexed tokens.
Filter daemon manager	The filter daemon manager is responsible for monitoring the status of the Full-Text Engine filter daemon host.

Note: In SQL Server 2008, the MSFTESQL service does not exist. Full-text tasks that were performed by the MSFTESQL service in SQL Server 2005 and earlier versions are now performed by the SQL Server process.

- **The Filter daemon host process**. The filter daemon host process is also known as the fdhost.exe process. In SQL Server 2008, separate processes known as filter daemon hosts loads filters. A server instance uses a multithreaded process for all multithreaded filters and a single-threaded process for all single-threaded filters.

Filter daemon host process	Component
Protocol handler	This component pulls the data from memory for further processing and accesses data from a user table in a specified database. One of its responsibilities is to gather data from the columns being full-text indexed and pass it to the filter daemon host, which will apply filtering and word breaker as required.
Filters	Some data types require filtering before the data in a document can be full-text indexed, including data in varbinary, varbinary(max), image, or XML columns. The filter used for a given document depends on its document type.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Filter daemon host process	Component
	For example, different filters are used for Microsoft™ Word (.doc) documents, Microsoft Excel® (.xls) documents, and XML (.xml) documents. Then the filter extracts chunks of text from the document, removing embedded formatting and retaining the text, and potentially, information about the position of the text.
Word breakers and stemmers	A word breaker is a language-specific component that finds word boundaries based on the lexical rules of a given language (word breaking). Each word breaker is associated with a language-specific stemmer component that conjugates verbs and performs inflectional expansions. At indexing time, the filter daemon host uses a word breaker and stemmer to perform linguistic analysis on the textual data from a given table column. The language that is associated with a table column in the full-text index determines which word breaker and stemmer are used for indexing the column.

Note: fdhost.exe replaces the Full-Text Engine filter daemon (msftefd.exe) of SQL Server 2005 and earlier versions.

An FDHOST launcher service (MSSQLFDLauncher) creates the fdhost.exe processes that are run under the security credentials of the FDHOST launcher service account. This service must be running for full-text indexing and full-text querying to work.

The following table lists the components of the filter daemon host process.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Question: What are the benefits of using full-text indexing?

Question: Which object can you use to prevent certain words getting indexed by full-text search engine?

Source:

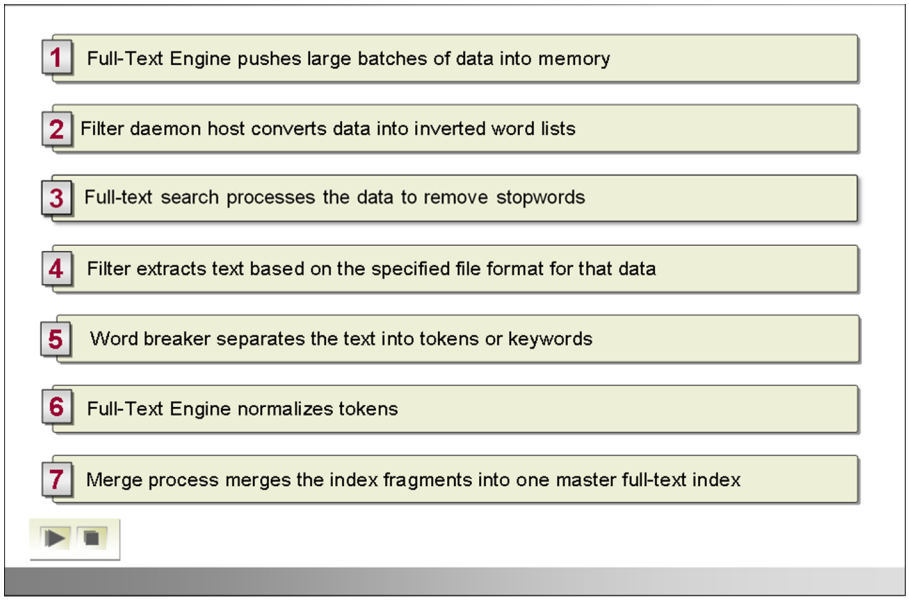
- http://technet.microsoft.com/en-us/library/cc879306.aspx
- http://technet.microsoft.com/en-us/library/ms142541.aspx

For more information, see Full-Text Search Architecture.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Full-Text Indexing Process



Key Points

SQL Server uses an indexing component of full-text search to create the full-text index and the subsequently update this index when modifications to the data occur in the full-text indexed tables.

The process of full-text indexing involves the following steps:

1. **Full-Text Engine pushes large batches of data into memory**. When a full-text population initiated, the Full-Text Engine pushes large batches of data into memory and notifies the filter daemon host.
2. **Filter daemon host converts data into inverted word lists**. The Filter daemon host filters and word breaks the data and converts data into inverted word lists.
3. **Full-text search processes the data to remove stopwords**. The full-text search then pulls the converted data from the word lists, processes the data to remove stopwords, and saves the word lists for a batch into one or more inverted indexes.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

4. **Filter extracts text based on the specified file format for that data**. Full-Text Engine retrieves the data stored as inverted indexes and passes it to the filter. The filter, which implements the IFilter interface, extracts text based on the specified file format for that data when you store indexing data in a varbinary(max) or image column, such as Microsoft Word. In some cases, the filter components require the varbinary (max), or image data to be written to the filter data folder, instead of being pushed into memory.
5. **Word breaker separates the text into tokens or keywords**. As part of processing, the gathered text data is passed through a word breaker. The word breaker separates the text into individual tokens, or keywords. The language used for tokenization is specified at the column level, or can be identified within **varbinary(max)**, **image**, or **xml** data by the filter component.
6. **Full-Text Engine normalizes tokens**. Additional processing may be performed to remove stopwords, and to normalize tokens before they are stored in the full-text index or an index fragment.
7. **Merge process merges the index fragments into one master full-text index**. When a population has completed, a final merge process is started that merges the index fragments into one master full-text index. This merging results in improved query performance because only the master index needs to be queried, rather than a number of index fragments and better scoring statistics may be used for relevance ranking.

To ensure optimal performance ensure that your transaction log contains sufficient space for a long-running transaction, before reorganizing a large full-text index in a database that uses the full recovery model. The master merge can be I/O intensive because large amounts of data must be written and read when index fragments are merged, though it does not block incoming queries. Also, master merging a large amount of data can create a long running transaction, delaying truncation of the transaction log during checkpoint. In this case, the transaction log might grow significantly under the full recovery model.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Question: What is the consideration for configuring full-text index processing?

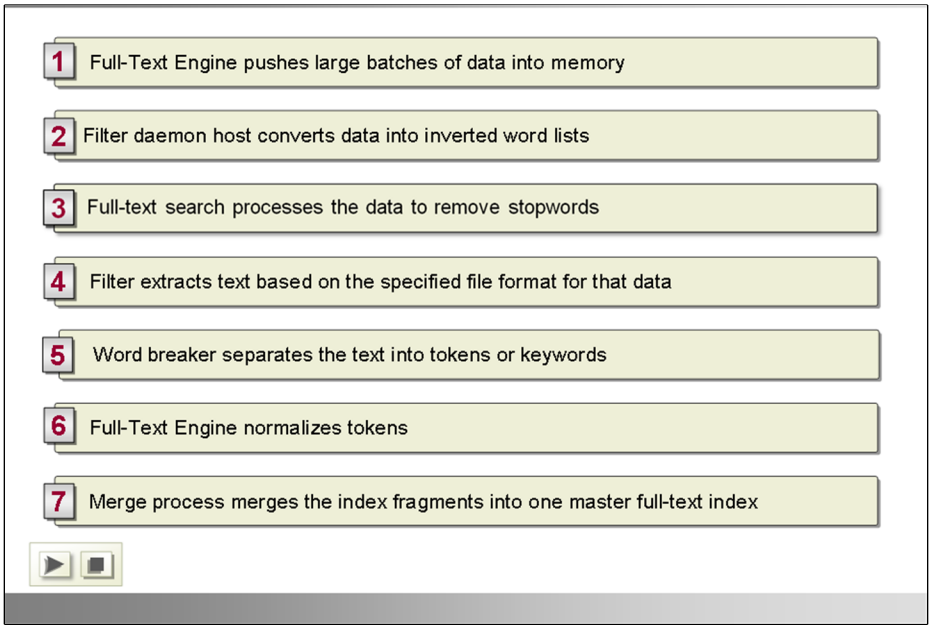
Source:

- http://msdn.microsoft.com/en-us/library/ms142591.aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Full-Text Querying Process



Key Points

After the columns have been added to a full-text index, applications and users can run full-text queries on the text in the columns. These queries can search for any of the following:

- **Simple term**. One or more specific words or phrases.
- **Prefix term**. A word or a phrase where the words begin with specified text.
- **Generation term**. Inflectional forms of a specific word.
- **Proximity term**. A word or phrase close to another word or phrase.
- **Thesaurus**. Synonymous forms of a specific word.
- **Weighted term**. Words or phrases by using weighted values.

Full-text queries use a small set of T-SQL predicates (CONTAINS and FREETEXT) and functions (CONTAINSTABLE and FREETEXTTABLE).

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Process of Full-Text Query

The process of full-text querying involves the following steps:

1. The query processor passes the full-text portions of a query to the Full-Text Engine for processing.
2. The Full-Text Engine performs word breaking, and optionally, thesaurus expansions, stemming, and stopword processing.
3. The Full-Text Engine then represents the full-text portions of the query in the form of SQL operators, primarily as streaming table-valued functions (STVFs). During query execution, these STVFs access the inverted index to retrieve the correct results.
4. The results are either returned to the client at this point, or they are further processed before being returned to the client.

Comparison of LIKE T-SQL Predicate to Full-Text Search

Unlike full-text search, the LIKE T-SQL predicate works on character patterns only. Also, you cannot use the LIKE predicate to query formatted binary data. Furthermore, a LIKE query against a large amount of unstructured text data is slower than an equivalent full-text query against the same data. A LIKE query against millions of rows of text data can take minutes to return; whereas a full-text query can take only seconds or less against the same data, depending on the number of rows that are returned.

Example

The following query return the rows containing the word "performance" from the ProductDescription table.

```
USE QuantamCorp;

GO

CREATE FULLTEXT CATALOG Mod03_Demo;

GO

CREATE FULLTEXT INDEX ON
Production.ProductDescription(Description) KEY INDEX
PK_ProductDescription_ProductDescriptionID ON Mod03_Demo;

GO
```

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.


```
CREATE FULLTEXT INDEX ON Production.Product (Name) KEY
INDEX PK_Product_ProductID ON Mod03_Demo;
GO
SELECT Description
FROM Production.ProductDescription
WHERE CONTAINS(Description, 'aluminum');
```

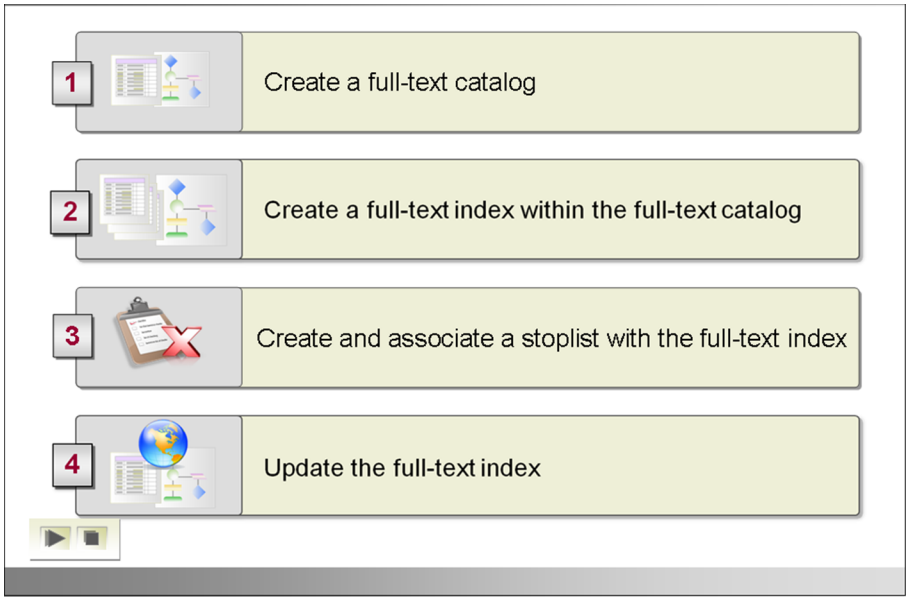
Question: Which functions and predicates are used to perform the full-text query process?

Source:

- http://technet.microsoft.com/en-us/library/ms142547.aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

Process of Enabling Full-Text Indexing



Key Points

SQL Server 2008 user databases are full-text enabled and cannot be disabled. Before you can use full-text index queries on a table, full-text indexing must be enabled on the columns that you want to access with the help of the Full-Text Engine. A full-text indexing catalog must be created before creating full-text indexes. Each full-text index must belong to a full-text catalog. You can create a separate text catalog for each full-text index, or you can associate multiple full-text indexes with a given catalog. The key column specified to be full-text indexed must be a unique, single-column, and non-nullable index, usually a primary key. The Full-Text Engine requires this unique index to map each row in the table to a unique, compressible key.

The process for placing a full-text indexing catalog involves the following steps:

1. **Create a full-text catalog**. For a SQL Server 2008 database, a full-text catalog is a virtual object that does not belong to any filegroup; it is a logical concept that refers to a group of full-text

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

indexes. When creating a full-text index, you need to specify a column-level language for the indexed column. The word breaker and stemmers of the specified language will be used by full-text queries on the column. There are a couple of things to consider for choosing the column language when creating a full-text index. These considerations relate to how your text is tokenized and then indexed by Full-Text Engine.

2. **Create a full-text index within the full-text catalog**. The process of building a full-text index consists of reading data from SQL Server, and then propagating the filtered data to the full-text index. As a best practice, locate a full-text index in the database filegroup that is best for maximizing I/O performance or locate the full-text indexes in a different filegroup on another volume. When ease of management is important to you, you must store table data and any affiliated full-text catalogs in the same filegroup. Sometimes, for performance reasons, you might want to have the table data and the full-text index in different filegroups that are stored on different volumes to maximize I/O parallelism. You must associate tables with the same update characteristics such as a small number of changes versus a large number of changes, or tables that change frequently during a particular time of day together under the same full-text catalog. By setting up full-text catalog population schedules, full-text indexes stay synchronous with the tables without adversely affecting the resource usage of the database server during periods of high database activity.

3. **Create and associate a stoplist with the full-text index**. To prevent a full-text index from becoming bloated, SQL Server has a mechanism that discards commonly occurring strings that do not help the search. These discarded strings are called stopwords. During index creation, the Full-Text Engine omits stopwords from the full-text index. This means that full-text queries will not search on stopwords. A stoplist is created and associated with each full-text index, and the words in that stoplist are applied to full-text queries on that index. By default, the system stoplist is associated with a new full-text index. However, you can create and use your own stoplist.

   For example, to create a new full-text stoplist named myStoplist3 by copying from the system stoplist, use the following CREATE FULLTEXT STOPLIST code.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.


```
USE QuantamCorp;

GO

CREATE FULLTEXT STOPLIST myStoplist FROM SYSTEM
STOPLIST;

GO
```

However, creating a stoplist is not a must because SQL Server 2008 will take the system default stoplist, if you do not define one. To alter a stoplist named myStoplist, add the word 'en', first for Spanish and then for French. Use the following ALTER FULLTEXT STOPLIST code.

```
USE QuantamCorp;

GO

ALTER FULLTEXT STOPLIST MyStoplist ADD 'demo' LANGUAGE
'Spanish';

ALTER FULLTEXT STOPLIST MyStoplist ADD 'demo' LANGUAGE
'French';

GO
```

4. **Update the full-text index**. As data is modified in the associated tables, full-text indexes are automatically updated. This is the default behavior. You can also keep your full-text indexes up-to-date manually or at specified scheduled intervals. Index updating is usually performed as an asynchronous process that runs in the background and keeps the full-text index up-to-date after modifications in the base table.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Question: How many stoplists can you associate with a full-text index?

Question: How can you manage the data storage of full-text indexing?

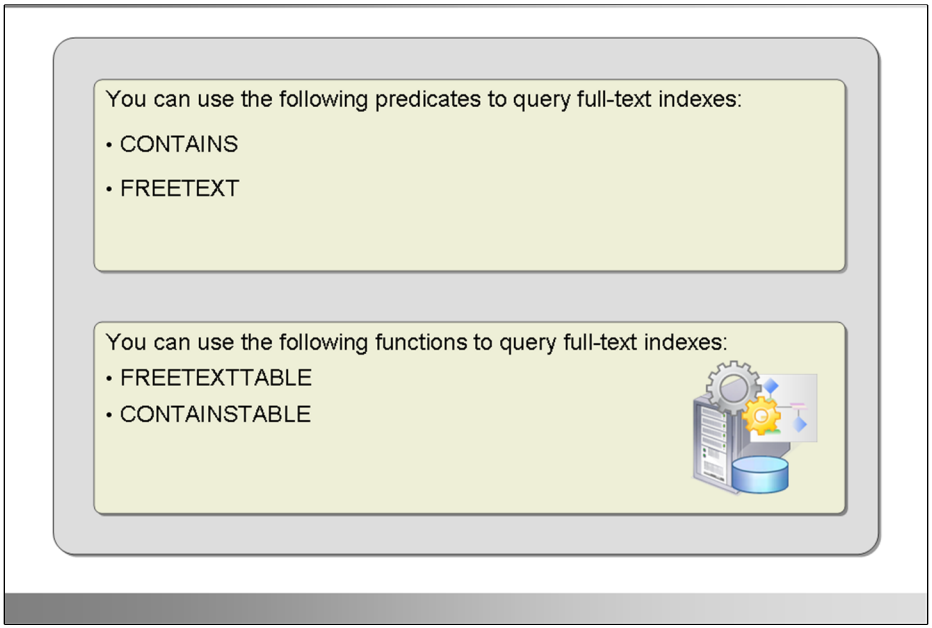
Source:

- technet.microsoft.com/en-us/library/ms142497.aspx
- http://technet.microsoft.com/en-us/library/ms142497.aspx
- http://msdn.microsoft.com/en-us/library/ms142551.aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Performing Full-Text Indexing in Queries



Key Points

Full-text queries use full-text predicates and functions to query full-text indexes. These predicates and functions support rich T-SQL syntax that supports a variety of forms of query terms. Full-text queries run in the same way with both the CONTAINS predicate and the CONTAINSTABLE function.

Full-Text Predicates

Full-text queries use full-text predicates such as CONTAINS and FREETEXT. You can specify CONTAINS and FREETEXT in the WHERE or HAVING clause of a SELECT statement. You can also combine them with any of the other T-SQL predicates, such as LIKE and BETWEEN.

The CONTAINS and FREETEXT predicates return a TRUE or FALSE value. You can use them only to specify selection criteria for determining whether a given row matches the full-text query. Matching rows are returned in the result set.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

When using CONTAINS or FREETEXT, you can specify either a single column, a list of columns, or all columns in the table to be searched. You can also specify the language whose resources will be used by given full-text query for word breaking and stemming, thesaurus lookups, and noise-word removal.

Full-Text Function

Full-text queries use full-text functions such as CONTAINSTABLE and FREETEXTTABLE. You can refer the CONTAINSTABLE and FREETEXTTABLE functions like a regular table name in the FROM clause of a SELECT statement. They return a table of zero, one, or more rows that match the full-text query. The returned table contains only rows of the base table that match the selection criteria specified in the full-text search condition of the function.

The following table lists the keywords used to change the behavior of the full-text function to return a relevant ranking value (RANK) and full-text key (KEY) for each row.

Query	Description
KEY column	The KEY column returns unique values of the returned rows. You can use the KEY column to specify selection criteria.
RANK column	The RANK column returns a rank value for each row that indicates how well the row matched the selection criteria. The higher the rank value of the text or document in a row, the more relevant the row is for the given full-text query. Note that different rows can be ranked identically. You can limit the number of matches to be returned by specifying the optional top_n_by_rank parameter.

When using either of these functions, you must specify the base table that is to be full-text searched. As with the predicates, you can specify a single column, a list of columns, or all columns in the table to be searched, and optionally, the language whose resources will be used by given full-text query.

You can use predicates and functions for the different kind of matches, as follows:

- **Use CONTAINS or CONTAINSTABLE for precise or less precise matches** to single words and phrases, words with

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

proximity within a certain distance of one another, or weighted matches. When using CONTAINS, you must specify at least one search condition that specifies the text that you are searching for and the conditions that determine the matches. You can use logical operation between search conditions. The following query finds all products with a price of $80.99 that contain the word "Mountain".

```
USE QuantamCorp;
GO
SELECT Name, ListPrice
FROM Production.Product
WHERE ListPrice = 80.99
   AND CONTAINS(Name, 'Mountain');
GO
```

The following query finds all products description that contains the text, performance.

```
SELECT *
FROM Production.ProductDescription
WHERE CONTAINS(Description, N'performance');
```

The following query finds all CONTAINSTABLE that contains product description and performance.

```
SELECT * FROM CONTAINSTABLE
(Production.ProductDescription, Description,
N'performance');
```

- **Use FREETEXT or FREETEXTTABLE for matching the**

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

meaning, but not the exact wording, of specified words, phrases or sentences (the freetext string). Matches are generated if any term or form of any term is found in the full-text index of a specified column. In the following example, the @SearchWord variable, which is defined as varchar(30), causes an implicit conversion in the FREETEXT predicate and returns rows with the word "performance" in the ProductDescription table.

```
USE QuantamCorp;

GO

SELECT Description

FROM Production.ProductDescription

WHERE FREETEXT(Description, 'performance');
```

The following query finds all products description that contains the text, performance.

```
SELECT *

FROM Production.ProductDescription

WHERE FREETEXT(Description, N'performance');
```

The following query finds all products description from FREETEXTTABLE that contains the product description, description, and performance.

```
SELECT * FROM FREETEXTTABLE
Production.ProductDescription, Description,
N'performance');
```

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Question: If you want to perform additional processing on one of the result returned by the full-text index engine, should you use a predicate or function of full-text index to perform the query?

Question: What is the difference between FREETEXT and CONTAINS?

Source:

- msdn.microsoft.com/en-us/library/ms142583.aspx

For more information, see How Search Query Results Are Ranked (Full-Text Search).

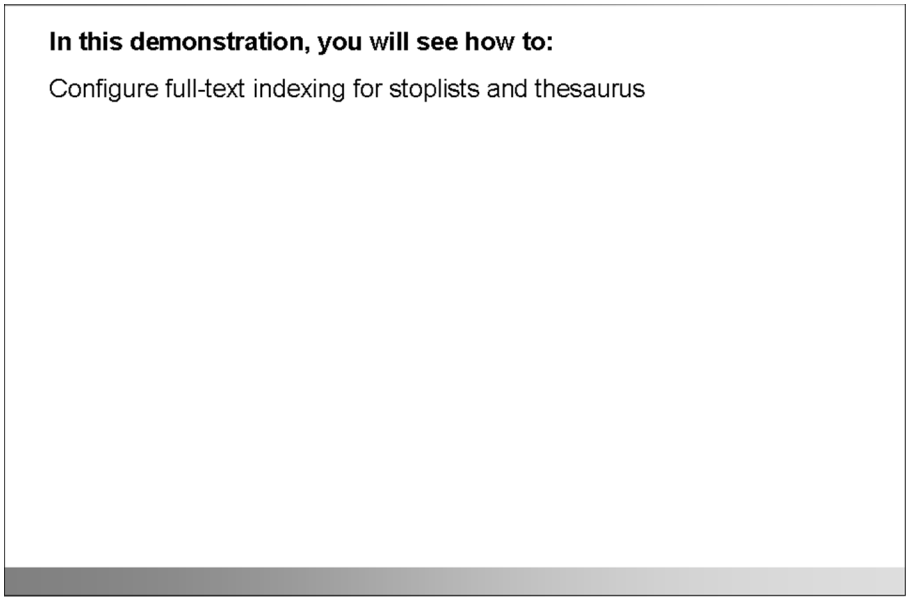
For more information, see CONTAINSTABLE (Transact-SQL).

For more information, see FREETEXTTABLE (Transact-SQL).

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Demonstration: How To Configure Full-Text Indexing



Key Points

The steps to configure full-text indexing are:

1. In Windows Explorer, browse to the **D:\Demofiles\Mod04** folder, and then double-click the **M4_Demo5.sql** file.
2. In the Microsoft SQL Server Management Studio window, in the Query Editor pane, select the **Create a full-text catalog, a full-text stoplist, a full-text index, and configure full-text indexing for stoplists** code, and then click **Execute**.
3. On the Start menu, point to **All Programs**, then click **Accessories**, and then click **Notepad**. The Notepad window appears.
4. In the Notepad window, on the **File** menu, click **Open**. The Open dialog box appears.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

5. In the **Open** dialog box, in the **File name** box, type **C:\Program Files\Microsoft SQL Server\MSSQL10.MSSQLSERVER\MSSQL\FTData\tsglobal.xml**, and then click **Open**.
6. Remove the following comment lines at the beginning and end of the file, respectively:

   <!--Commented out-->
7. Click **File**, and then click **Save**.
8. In the Microsoft SQL Server Management Studio window, in the Query Editor pane, select the **Configure full-text indexing for thesaurus** code, and then click Execute.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Lab 4: Designing a Strategy for Maintaining a Database in SQL Server 2008

- Exercise 1: Designing a Database Maintenance Strategy (Discussion)
- Exercise 2: Maintaining Statistics
- Exercise 3: Maintaining Indexes
- Exercise 4: Implementing Full-Text Indexing

Logon Information

Virtual machine	50400A-NYC-SQL1
User name	Administrator
Password	Pa$$w0rd

Estimated time: 45 minutes

Objectives:

After completing this lab, you will be able to:

- Design a strategy to maintain a database in SQL Server 2008.
- Maintain statistics.
- Maintain indexes.
- Implement data compression.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Scenario

QuantamCorp has deployed two SQL Server 2008 servers in their organization and will continue to deploy more to replace theirexisting installation of SQL Server 2005. SQL Server 2008 will be theprimary database engine used in their organization. After fine-tuning the database physical structures, you realize that the application still does not perform well. You have a meeting with your applicationteam and understand that the performance of the application is poorwhen a large record set is updated on a table, or when they selectdata from large tables. You then propose to do some further fine-tuning on the database to help resolve this issue.

You plan to perform the following to optimize the performance of the database:

- Determine when statistics has to be updated.
- Determine how applications query the database and then determine how to create a certain type of index.
- Determine the impact of data compression.
- Determine if statistics have to be manually updated.
- Create a clustered index to optimize the performance and storage.
- Create an additional nonclustered index to optimize the performance.
- Create an XML index to optimize the performance of query that retrieves XML data.
- Manually update the statistics.
- Enable data compression.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Exercise 1: Designing a Database Maintenance Strategy (Discussion)

Scenario

Your application team recently reports to you that their application performance is very poor. You are assigned to help them investigate and resolve the issues. The application team informs you that they encounter performance degrade only once in a day, usually when there is a large volume of data being uploaded. After initial investigation, you realize that the auto-create and auto-update statistics options are disabled. The application team disabled these options because they were told that these features adversely affect application performance.

In addition, you observe that there is no index created on the table. The application team provides you the following query sample.

```
SELECT * FROM CallDetailsHistory WHERE source_countrycode
BETWEEN 60 AND 100 AND  dest_countrycode BETWEEN 100 AND
300
```

```
SELECT additional_info.query('
  /call_info/source
') as source_code,additional_info.query('
  /call_info/dest
') as dest_code
FROM dbo.CallDetailsHistory
WHERE additional_info.exist ('/call_info/@to[.="1234"]') =
1
```

The following is the schema of a table that experiences poor performance.

	Column Name	Data Type	Allow Nulls
▶	call_date	datetime	☐
	caller	varchar(20)	☐
	dest	varchar(20)	☐
	source_countrycode	int	☐
	dest_countrycode	int	☐
	duration	int	☑
	additional_info	xml	☑

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

The table stores information on two partitions. One partition stores information for the current month, whereas all other months' data is stored in the other partition.

The application team informs you that they are going to load a large amount of historical data into the application. You want to optimize the disk space storage and query performance without modifying the database schema.

Based on the given scenario, answer the following questions:

1. What should you suggest to the application team to resolve the performance issue due to the impact of loading a large amount of data into the database?
2. What would you suggest the application team for modifying the table schema?
3. The application has team raised a concern about the reduced database performance while loading data due to the impact of creating an index. What would you suggest them?
4. After reviewing the table schema you realize that it has many fields created with the nchar(4000) data type. However, the nchar(4000) data type stores only 10-200 characters. What would you suggest the application team to help them optimize their disk usage? In addition, the application team wants to improve the overall query performance without modifying the database schema and application?

Exercise 2: Maintaining Statistics

Scenario

You have loaded a large amount of data into the database. After sometime, you notice that a query is performing badly and you want to check if the statistics are up-to-date. You use the execution plan to check and confirm that the statistics are updated. You also need to optimize the statistic update setting to ensure that the update of statistics happens automatically and would not affect the query performance.

The main tasks for this exercise are as follows:

1. Create a test database.
2. Create a test table and add test records to the table.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

3. Run a query with updated statistics.
4. Disable auto-create and auto-update statistics.
5. Modify the data in the table.
6. Run the query with outdated statistics.
7. Manually update statistics.
8. Enable auto-create and auto-update statistics.

Task 1: Create a test database.

1. Open the Microsoft SQL Server Management Studio window.
2. In the Connect to Server dialog box, type the following to connect to Microsoft SQL Server Management Studio:

- Server type: **Database Engine**
- Server name: **NYC-SQL1**
- Authentication: **Windows Authentication**

3. Using the Microsoft SQL Server Management Studio window, in the Query Editor pane, run a Transact SQL (T-SQL) script to create a test database **Mod04**, by using the following design.

File name	Settings
Mod04_Primary.mdf	Current size: 10 Maximum size: 50 Filegrowth: 5
Mod04_Log.ldf	Current size: 10 Maximum size: 50 Filegrowth: 5

4. Verify that the database file and the transaction log files are created.

Note: Browse to the **D:\Labfiles\Mod04** folder and select the **E2_T1.sql** file to execute the required T-SQL code.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Task 2: Create a test table and add test records to the table.

1. In the Query Editor pane of the Microsoft SQL Server Management Studio window, run a T-SQL script to create a test table named **[dbo].[CallDetails]** in the **Mod04** database, with the following structure:
   - Call_date: **datetime**
   - Caller: **varchar (20)**
   - Duration: **int**
   - Additional_info: **nchar (4000)**
   - Allow only the **Duration** and the **Additional_info** columns to have NULL values.
   - Primary Key: **call_date**
2. In the Query Editor pane, run a T-SQL script to add multiple records with the call_date from 1/1/1999 to 12/31/2004 to the **CallDetails** table.

Note: Browse to the **D:\Labfiles\Mod04** folder and select the **E2_T2.sql** file to execute the required T-SQL code.

Task 3: Run a query with updated statistics.

1. In the Query Editor pane of the Microsoft SQL Server Management Studio window, run the T-SQL command to update the statistics. Use the stored procedure **sp_updatestats** for this.
2. Verify that the updated statistics are updated by selecting all records from the **CallDetails** table where the **Duration** is between **10** and **30**.

Task 4: Disable auto-create and auto-update statistics.

1. In the Query Editor pane of the Microsoft SQL Server Management Studio window, run a T-SQL script to disable the auto-create and auto-update statistics option.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Task 5: Modify the data in the table.

1. In the Query Editor pane of the Microsoft SQL ServerManagement Studio window, delete all records from the **CallDetails** table where **call_date** is greater than **2000-1-1**.
2. Insert records in the **[dbo].[CallDetailsHistory]** table with call dates between 2004-01-01 and 2009-12-31.
3. In the Query Editor pane of the Microsoft SQL Server Management Studio window, delete all records from the **CallDetails** table, where **call_date** is less than **2004-1-1**.

Note: Browse to the **D:\Labfiles\Mod04** folder and select the **E2_T5.sql** file to execute the required T-SQL code.

Task 6: Run the query with outdated statistics.

1. In the Query Editor pane of the Microsoft SQL Server Management Studio window, run a query to view records from the **CallDetails** table, where the call duration is between **10** and **20** and the calling year is above **2005**. Note the execution plan.
2. In the Query Editor pane of the Microsoft SQL Server Management Studio window, run a query to view the statistic information of the **calldetails** table to verify if the statistics is updated.

Task 7: Manually update statistics.

1. In the Query Editor pane of the Microsoft SQL Server Management Studio window, run a query to perform a manual update of statistics.
2. In the Query Editor pane of the Microsoft SQL Server Management Studio window, run a query to view the statistic information and verify if the statistics is updated.

Task 8: Enable auto-create and auto-update statistics.

1. Using the Microsoft SQL Server Management Studio window, enable the auto-create and auto-update statistics.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Results: After completing this exercise, you should have created a test database, created a test table and added records to the table, updated the statistics, disabled the auto-create and the auto-update statistics, modified the data in the table, run a query with outdated statistics, and manually updated the statistics, enabled auto-create and auto-update statistics.

Exercise 3: Maintaining Indexes

Scenario

In one of your application, some of the queries are performing badly. After checking the database schema, you realize that there are no indexes created on that database. You decide to implement indexes on that database to help improve the performance of queries.

The main tasks for this exercise are as follows:

1. Create a test table and add test records to the table.
2. Run a test query without index.
3. Create a clustered index.
4. Create a nonclustered index with included column.
5. Create an XML index.

Task 1: Create a test table and add test records to the table.

1. Using the Microsoft SQL Server Management Studio window, create the **[dbo].[CallDetailsHistory]** table with the following structure:
   - Call_date: **datetime**
   - Caller: **varchar(20)**
   - Dest: **varchar(20)**
   - source_countrycode: **int**
   - dest_countrycode: **int**
   - Duration: **int**
   - Additional_info: **xml**
2. Allow only the **Duration** and the **Additional_info** columns to have NULL values.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

3. In the additional_info column of xml data type, insert xml data that includes call details.
4. Insert records in the [dbo].[CallDetailsHistory] table with call dates between 2004-01-01 and 2009-12-31.

Note: Browse to the **D:\Labfiles\Mod04** folder and select the **E3_T1.sql** file to execute the required T-SQL code.

Task 2: Run a test query without index.

1. Using the Microsoft SQL Server Management Studio window, enable the display of Actual Execution Plan.
2. In the Query Editor pane, run a test query without an index to select all records from the **CallDetailsHistory** table, with the source country code between 60 and 100 and the destination country code between 100 and 300.
3. In the Query Editor pane, on the Execution plan tab, verify the use of table scan.

Task 3: Create a clustered index.

1. Enable the display of Actual Execution Plan.
2. Using the Microsoft SQL Server Management Studio window, add a clustered index **PK_CallDetailsHistory** on the primary key, call_date, of the **dbo.CallDetailsHistory** table.
3. Verify all the records from the **CallDetailsHistory** table with the destination country code between 100 and 300 and the source country code between 60 and 100.

Note: Browse to the **D:\Labfiles\Mod04** folder and select the **E3_T3.sql** file to execute the required T-SQL code.

Task 4: Create a nonclustered index with included column.

1. Enable the display of Actual Execution Plan.
2. Using the Microsoft SQL Server Management Studio window, create a nonclustered index **IX_CallDetailsHistory** on the **dbo.CallDetailsHistory** table.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

3. Verify use of nonclustered index by querying a SELECT statement with a WHERE clause to filter the result by using the **source_countrycode** column.
4. Using the Microsoft SQL Server Management Studio window, delete the nonclustered index **IX_CallDetailsHistory**. Create a nonclustered index **IX_CallDetailsHistory** with included column, dest_country_code on the **dbo.CallDetailsHistory** table of the **Mod04** database.
5. Verify use of nonclustered index by querying a SELECT statement with a WHERE clause to filter the result by using the **source_countrycode** column.

Note: Browse to the **D:\Labfiles\Mod04** folder and select the **E3_T4.sql** file to execute the required T-SQL code.

Task 5: Create an XML index.

1. Enable the display of Actual Execution Plan.
2. Using the Microsoft SQL Server Management Studio window, create a primary XML index and secondary XML indexes for PATH, VALUE, and PROPERTY.
3. Verify use of the index by querying a SELECT statement for the source and the destination country code from the **dbo.CallDetailsHistory** table where the attribute in the **call_info** xml field equals to 1234.

Note: Browse to the **D:\Labfiles\Mod04** folder and select the **E3_T5.sql** file to execute the required T-SQL code.

Results: After completing this exercise, you should have created a test table and added a record in the table, run a test query without an index, created a clustered index, created a nonclustered index, and created an XML index.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Exercise 4: Implementing Full-Text Indexing

Scenario

One of your business applications requires using text-based queriesto locate records. Currently, the application is using the LIKEoperator but the users are complaining about the query performance.You want to implement full-text indexing with custom stoplist tosupport text-based query. This will ensure that the user would notable to search information they are not supposed to search.

The main tasks for this exercise are as follows:

1. Create a test table and add records to the table.
2. Modify the database file structure for full-text index.
3. Create a full-text catalog and a full-text index.
4. Verify the full-text index setup.
5. Manage a full-text index stoplist.
6. Verify the full-text index setup.

Task 1: Create a test table and add records to the table.

1. In the Microsoft SQL Server Management Studio window, create a test table named **CallDetails_FT** with the following structure:
   - id: uniqueidenfitier
   - call_date: datetime
   - caller: varchar(20)
   - duration: int
   - additional_info: nchar(4000)
- Allow all the fields except the **id** field to have NULL values.
- Also, create a constraint named **PK_CallDetails_FT** on the primary key **id** with the following settings:

  PAD_INDEX = OFF

  STATISTICS_NORECOMPUTE = OFF

  IGNORE_DUP_KEY = OFF

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

ALLOW_ROW_LOCKS = ON

ALLOW_PAGE_LOCKS = ON

2. In the Microsoft SQL Server Management Studio window, insert records as follows in the **CallDetails_FT** table.

Field name	Record
id	Assign a unique value by using the NEWID() function.
call_date	A date between 2004-01-01 and 2009-12-31.
caller	Peter
duration	30
additional_info	Call to 'x'. Where 'x' is a random number between one and 100000.

Note: Browse to the **D:\Labfiles\Mod03** folder and select the **E4_T1.sql** file to execute the required T-SQL code.

Task 2: Modify the database file structure for full-text index.

2. In the Microsoft SQL Server Management Studio window, add a filegroup named FullTextData to the database.
3. Add the **Mod03_FullText01.mdf** database file, to the **FullTextData** filegroup with the following settings:
   - Logical file name: FullText_Data01
   - Current size: 10
   - Maxsize: 50
   - Filegrowth: 5
4. Verify the additional database file that is created in the **D:\Data** folder.

Note: Browse to the **D:\Labfiles\Mod03** folder and select the **E4_T2.sql** file to execute the required T-SQL code.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Task 3: Create a full-text catalog and a full-text index.

1. In the Microsoft SQL Server Management Studio window, create a full-text catalog, **Mod03_Catalog** as default.
2. Create a full-text index named **PK_CallDetails** on the **additional_info** column of **dbo.CallDetails_FT** on the **Mod04_Catalog** catalog and the **FullTextData** filegroup.

Note: Browse to the **D:\Labfiles\Mod03** folder and select the **E4_T3.sql** file to execute the required T-SQL code.

Task 4: Verify the full-text index setup.

2. In the Microsoft SQL Server Management Studio window, verify all records from the table **CallDetails_FT** that contains the word **"call"** in the **additional_info** column.

Task 5: Manage a full-text index stoplist.

1. In the Microsoft SQL Server Management Studio window, create a full-text index stoplist named **CallDetailsStopList**.
2. In the **CallDetailsStopList** stoplist, add **English** to the list of call languages.
3. Add the stoplist to the full-text index on the table, **dbo.CallDetails_FT**.

Task 6: Verify the full-text index setup.

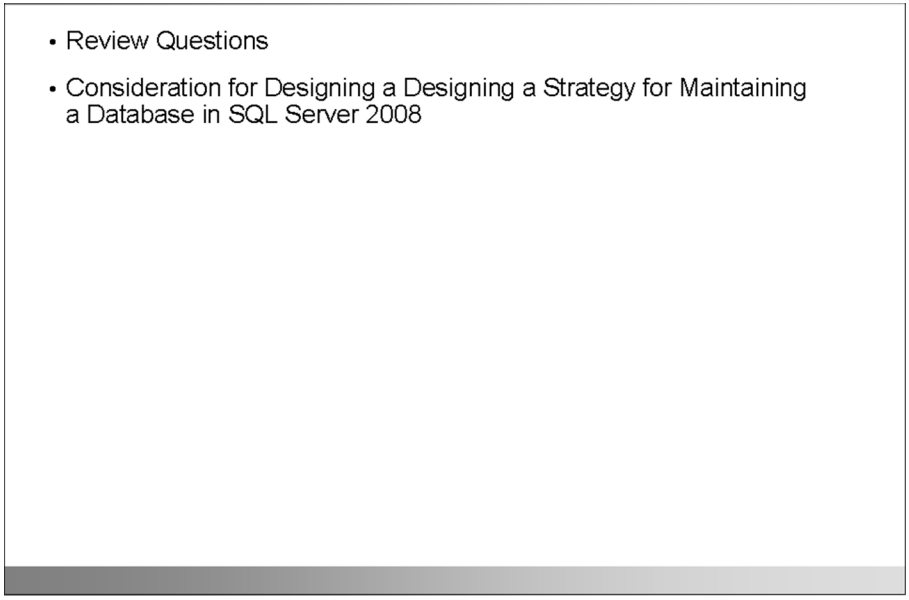
2. In the Microsoft SQL Server Management Studio window, verify all records from the table, **CallDetails_FT** that contains the word **"call"** in the **additional_info** column.

Results: After completing this exercise, you should have created a test table and added records to it, added a filegroup to the database, attached a file to the filegroup, created a full-text catalog, created a full-text index, verified the full-text index setup, created a stoplist on the full-text index, and verified the full-text index setup of the full-text index with the stoplist.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Module Review and Takeaways



Review Questions

1. In your organization, you have set the update mode of the SQL Server to automatic. Recently, you realize that the performance of an application decreases while uploading large amounts of data into the system. To improve the performance of the application, you decide to switch to the manual update mode. How can you resolve this problem by using the manual statistics mode?
2. After you disable automatic statistics update, you realize that the performance of the SQL Server is not optimal. How can you improve the performance of the SQL Server while this option is still disabled?
3. Now, when you review the performance of a poor performance query, you realize that some tables do not have an index defined. You decide to add an index to these tables but you are not sure which one to add. What is difference between a clustered and nonclustered index?

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

4. You have 1 terabyte (TB) of data stored in a database but you notice that most of the data stored is of varchar data type. Also, it is not of the maximum capacity of the field and is only a portion of the data length defined in the table schema. You decide to use data compression on the database. How does data compression helps improve the query performance?
5. You have revised an index design for a table and you find that most of the time SQL Server queries both clustered and nonclusterd indexes to provide results of a query. You are thinking to reduce the need for SQL Server to query both clustered and nonclustered index in the query operation. To do this, you decide to use included column nonclustered index? How does included column nonclustered index help to improve query performance?
6. Recently, your database is running out of space. You decide to use partitioning and compression to optimize the disk space of the database. How will partitioning and compression help you optimize the disk space of the database?

Consideration for Designing a Strategy for Maintaining a Database in SQL Server 2008

When planning a strategy for maintaining a database in SQL Server, consider the following:

- Use clustered index for most tables and add non-clustered index to helps cover specific queries.
- Use specialized index, such as XML and spatial, to optimize the read access of specific information.
- Set FILLFACTOR to around 50-70% for OLTP system and consider using 100% for OLAP.
- Rebuild an index regularly to reorganize data for better performance.
- Disable an index when loading large amount of data to help improve the performance.
- Keep statistics update by using automatic statistics update.
- Use asynchronous automatic statistics update, if automatic statistic update causes performance issues.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

- Enable compressed table and index to reduce the storage requirements for database.
- Use compressed table and index to reduce the data page required to load the data and to improve performance.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Module 5

Designing Solutions for Managing SQL Server 2008

Contents:

Lesson 1: Designing a Policy-Based Management Strategy 5-3
Lesson 2: Controlling Resource Usage by Using Resource Governor 5-18
Lesson 3: Auditing Database Changes 5-37
Lab 5: Designing Solutions for Managing SQL Server 2008 5-65
Module Review and Takeaways 5-79

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Module Overview

- Designing a Policy-Based Management Strategy
- Controlling Resource Usage by Using Resource Governor
- Auditing Database Changes

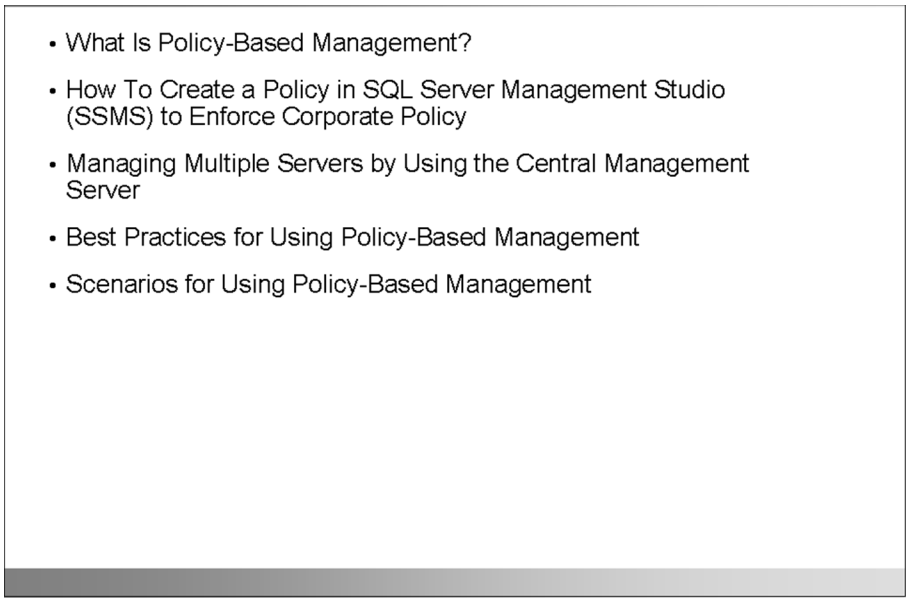
The enhanced scalability of Microsoft™ SQL Server™ 2008 enables business enterprises to run their critical applications while lowering the cost of managing the data infrastructure and delivering insights and information to all users.

In this module, you will learn about the various tools for managing servers in SQL Server 2008. You will then learn about designing policy-based management. You will also learn to control resource usage by using Resource Governor. In addition, you will learn about auditing database changes.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Lesson 1
Designing a Policy-Based Management Strategy



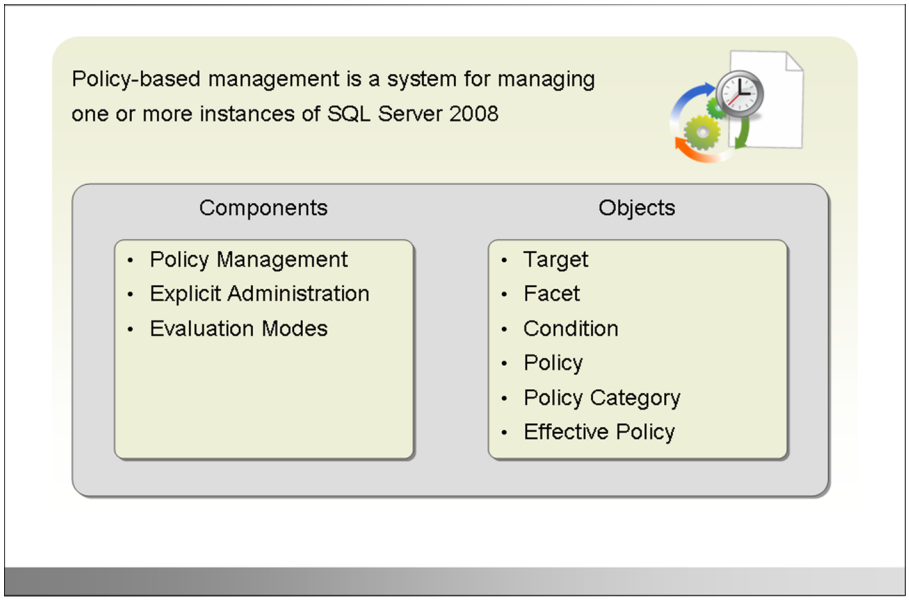
Policy-based management is a system that helps to manage one or more instances of SQL Server 2008. It allows you to manage SQL Server instances through clearly defined policies. You can apply these policies against a single server or a group of servers, to improve the scalability of monitoring and administration.

In this lesson, you will learn about policy-based management. You will learn how to create a policy in SSMS to enforce corporate policy and how to manage multiple servers by using the central management server. You will also learn about the scenarios and best practices for using policy-based management.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

What Is Policy-Based Management?



Key Points

You can use the policy-based management system to manage one or more instances of SQL Server. You can use SSMS to manage entities such as instances of SQL Server, databases, and other SQL Server objects.

Components in a Policy-based Management

The policy-based management system comprises three components:

- **Policy management**. You can create policies to define what to monitor.
- **Explicit administration**. You can select one or more managed targets and check whether the targets comply with a specific policy, or you can explicitly make the targets comply with a policy.
- **Evaluation modes**. There are four evaluation modes, out of which three are automated:

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

- **On demand**. This mode evaluates the policy when it is directly specified by the user.
- **On change: prevent**. This automated mode uses data definition language (DDL) triggers to prevent policy violations.
- **On change: log only**. This automated mode uses event notification to evaluate a policy when a relevant change is made.
- **On schedule**. This automated mode uses a SQL Server Agent job to periodically evaluate a policy.

If you do not enable the automated evaluation modes, the policy-based management system will not perform the policy check at real time and will not affect system performance. Depending on the type of evaluation you choose, you will have a different result. For example, if you use a policy to evaluate on-change, but prevent the change, an error message will appear. A log entry will be created if you configure to use on-change evaluation, but perform logging instead of prevent.

Objects in Policy-Based Management

Policies in SQL Server consist of the following objects:

- **Target**. The managed targets are entities that are managed by policy-based management, such as an instance of the SQL Server Database Engine, a database, a table, or an index. All targets in a server instance form a target hierarchy. A target set is the set of targets that results from applying a set of target filters to the target hierarchy. For example, all the tables in the database owned by the HumanResources schema will form a target set.
- **Facet**. A facet is a set of logical properties that defines the behavior or configuration for certain types of managed targets. The number and characteristics of the properties are built into the facet and can be added or removed only by the creator of the facet, in this case, Microsoft. A target type can implement one or more facets, and a facet can be implemented by one or more target types. Some properties of a facet can only apply to a specific version. For example, the Database Mail property of the Surface™ Area Configuration facet applies only to SQL Server 2005 and SQL Server 2008.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

- **Condition**. A condition is a Boolean expression that specifies a set of allowed states of a policy-based management managed target with regard to a facet.
- **Policy**. A policy refers to a policy-based management condition and the expected behavior, such as evaluation mode, target filters, and schedule. A policy can contain only one condition. Policies can be enabled or disabled. You can choose to schedule a time to evaluate a policy. The scheduled policy will create a SQL Server Agent job, which initiates the evaluation. You can change the method of policy evaluation from the property of the window policy. You need to set the value of the check condition to On Schedule to have a scheduled evaluation of the policy.
- **Policy category**. A policy category is a user-defined category to help manage policies. Users can classify policies into different policy categories. A policy can belong to only one policy category. Policy categories apply to databases and servers. At the database level, the following conditions apply:
  - Database owners can subscribe a database to a set of policy categories.
  - Only policies from its subscribed categories can govern a database.
  - All databases implicitly subscribe to the default policy category.
  - At the server level, policy categories can be applied to all databases.
- **Effective policy**. The effective policies of a target are those policies that govern the target. A policy is effective with regard to a target only, if all the following conditions are true:
  - The policy is enabled.
  - The target belongs to the target set of the policy.
  - The target or one of the target's ancestors subscribes to the policy group that contains this policy.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Question: What is the difference between a scheduled evaluation and real-time evaluation?

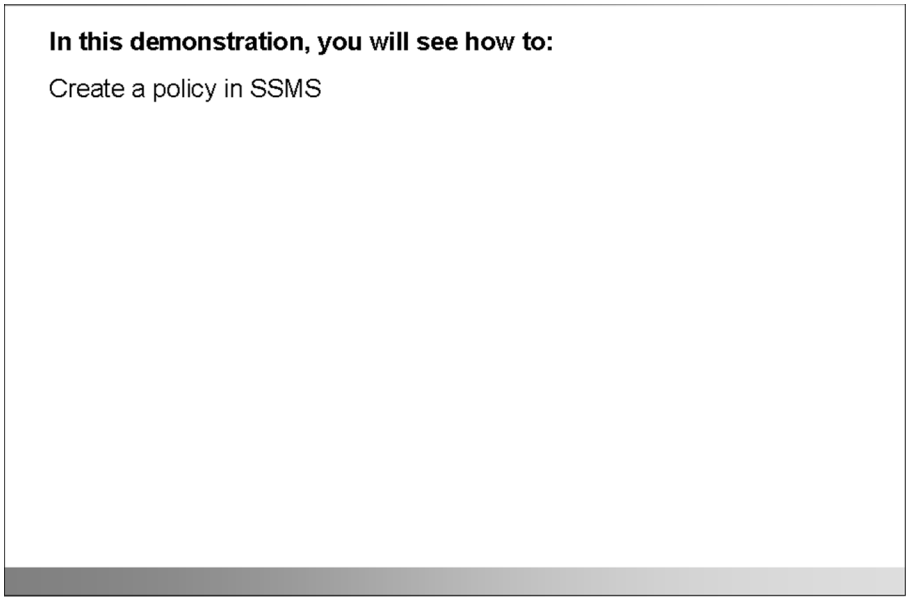
Source:

- http://technet.microsoft.com/en-us/library/bb510667.aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Demonstration: How To Create a Policy in SSMS to Enforce Corporate Policy



Key Points

The steps to create a policy in SSMS are:

1. On the Start menu, point to **All Programs**, click **Microsoft SQL Server 2008**, and then click **SQL Server Management Studio**. The Microsoft SQL Server Management Studio window appears.
2. To log on to SQL Server Management Studio, in the **Connect to Server** dialog box, type the following:
   - Server type: **Database Engine**
   - Server name: **NYC-SQL1**
   - Authentication: **Windows Authentication**
3. In the **Connect to Server** dialog box, click **Connect.**
4. In the Microsoft SQL Server Management Studio window, in the Object Explorer pane, expand **Management**, expand **Policy Management**, right-click **Conditions**, and then click **New**

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Condition. The **Create New Condition** - dialog box appears.

5. In the **Create New Condition** dialog box, in the **Name** box, type **Mod05 Tables**.
6. In the **Facet** list, click **Multipart Name**.
7. In the **Expression** area, for the first row, in the **Field** list, click **@Name**, and in the **Operator** list, click **LIKE**.
8. In the **Expression** area, click the ellipses button besides the Value list. The **Advanced Edit** dialog box appears.
9. In the **Advanced Edit** dialog box, type **'tbl%'** to force all table names to start with the letters **tbl**, and then click **OK** to close the Advanced Edit dialog box.
10. On the **Description** page, type a condition such as**, Database table names must begin with tbl**, and then click **OK** to create the condition.
11. In the Microsoft SQL Server Management Studio window, in the **Object Explorer** pane, right-click **Policies**, and then click **New Policy**. The **Create New Policy –**dialog box appears.
12. In the **Create New Policy –** dialog box, in the **Name** box, type **Mod05 Policy**.
13. In the **Check condition** list, under **Multipart Name**, click **Mod05 Tables.**
14. In the **Against targets** area, select the **Every Table in Every Database** check box to apply the policy.
15. Under the **Every Table** check box, in the **Every Database** area, click the **Every** arrow, and then click **New condition**. The **Create New Condition -** dialog box appears.
16. In the **Create New Condition** dialog box, in the **Name** box, type **Mod05 Database.**
17. In the **Expression** area, for the first row, in the **Field** list, click **@Name**.
18. In the **Expression** area, click the ellipses button besides the **Value** list. The **Advanced Edit** dialog box appears.
19. In the **Advanced Edit** dialog box, type **'Mod05'** and then click **OK.**

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

20. In the **Create New Condition — Mod05 Database** dialog box, click **OK**.
21. In the **Evaluation Mode** list, click **On change: prevent**.
22. In the **Server restriction** box, ensure **None** is selected. Click **OK** to start the policy creation process.
23. In the Microsoft SQL Server Management Studio window, expand **Policies**, right-click the Mod05 created policy, and then click **Enable**.

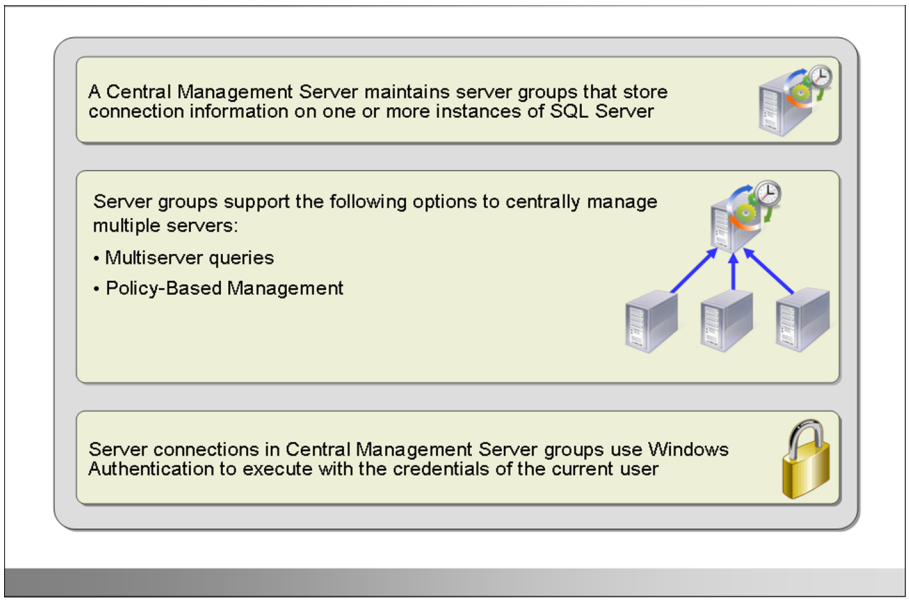
Question: How would you use a policy to enforce corporate standards?

Question: What is the benefit of scheduling a policy execution at regular intervals?

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Managing Multiple Servers by Using the Central Management Server



Key Points

In SQL Server 2008, you can administer multiple servers by using Central Management Servers and by creating server groups. An instance of SQL Server that is designated as a Central Management Server maintains server groups, which maintains the connection information for one or more instances of SQL Server. You can execute Transact-SQL (T-SQL) statements and policy-based management policies at the same time against server groups. However, you cannot designate SQL Server 2000 and SQL Server 2005 as a Central Management Server.

The Central Management Server includes connection information about the configuration targets. It supports only Windows® Authentication; so, authentication information is not stored. This enables the Central Management Server to run T-SQL statements on multiple servers at the same time. To create a Central Management Server and server groups, use the Registered Servers window in SQL

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Server Management Studio. The Central Management Server cannot be a member of a group that it maintains.

Multiserver Queries

To run T-SQL statements against all servers in a server group at the same time, open Query Editor from the server group in the Registered Servers window. When you run T-SQL statements in Query Editor, the statements will be executed against all the servers in the group. You can merge the results of the query into a single results pane, or return the results in separate result panes. When you are merging results, the first server to respond sets the schema for the result set. To merge the result sets, the query must return the same number of columns with the same column names from each server. When you are merging results, a message is displayed for each server that does not match the schema. This message contains the column count and the column names and is returned by the first server to return results. When you do not merge results, the results set from each server will be displayed in its own grid with its own schema.

As an option, Query Editor can include columns that provide the name of the server that produced each row, and the logon that was used to connect to the server that provided each row.

Policy-Based Management

You can evaluate a policy against a registered server or a whole server group. The on-demand policy execution mode enables the evaluate mode and the apply mode for policies. If you cannot establish a connection to one or more servers, those servers are ignored and evaluation against the other servers continues independently.

Security

The connections to servers in Central Management Server groups execute with the credentials of the current user, by using Windows Authentication, and so the effective permissions on the servers in the server groups might vary. For example, the user might be a member of the sysadmin fixed server role on the instance of SQL Server A, but have limited permissions on the instance of SQL Server B.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

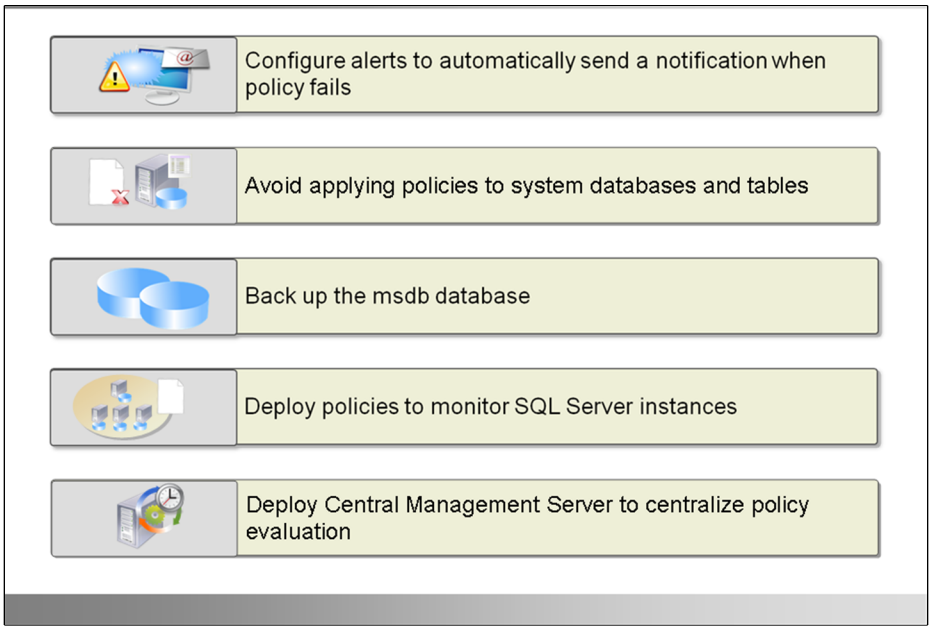
New Horizons Computer Learning Centers, Inc.

Question: What are the benefits of using Central Management Server over policy-based management?

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Best Practices for Using Policy-Based Management



Key Points

You can use policy-based management to automate the deployment of the configuration to multiple servers. This will prevent unwanted changes to databases.

The following are the best practices for using policy-based management:

- **Configure alerts to get notified automatically when a policy fails**. When policy-based management policies are executed in one of the three automated evaluation modes, if a policy violation occurs, a message is written to the event log. You can create an alert to detect the message and correct the problem.

  The following table describes how an alert should detect the messages.

Execution mode	Message number

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Execution mode	Message number
On change: prevent (if automatic)	34050
On change: prevent (if On demand)	34051
On schedule	34052
On change	34053

- **Avoid applying policies to system databases and tables**. Policies can affect the functioning of some SQL Server features if you apply them to system database and tables. For example, the Change Data Capture (CDC) and transactional replication features use the systranschemas table that does not have an index. If you enable a policy in a way such that all tables must have an index, enforcing compliance of the policy will cause these features to fail.
- **Back up the msdb database**. You should back up the msdb database because all policy information is stored on the msdb database.
- **Deploy policies to monitor SQL Server instances**. Policy-based management lets you monitor best practices for the SQL Server Database Engine. SQL Server 2008 provides a set of policy files that you can import as best practice policies, and then evaluate the policies against a target set that includes instances, instance objects, databases, or database objects. You can evaluate policies manually, set policies to evaluate a target set according to a schedule, or set policies to evaluate a target set according to an event.

  SQL Server 2008 includes policies that are not installed on the Database Engine. However, you can import them from the default installation location of C:\Program Files\Microsoft SQL Server\100\Tools\Policies\DatabaseEngine\1033. You can use these policies to monitor an instance of SQL Server.
- **Deploy Central Management Server to centralize policy evaluation**. Deploy Central Management Server, if you have multiple SQL Servers and you plan to use policy-based management. This allows centralized policy deployment and configuration.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Question: How should you get notified if there are objects detected, which are not following the rules defined in a policy?

Source:

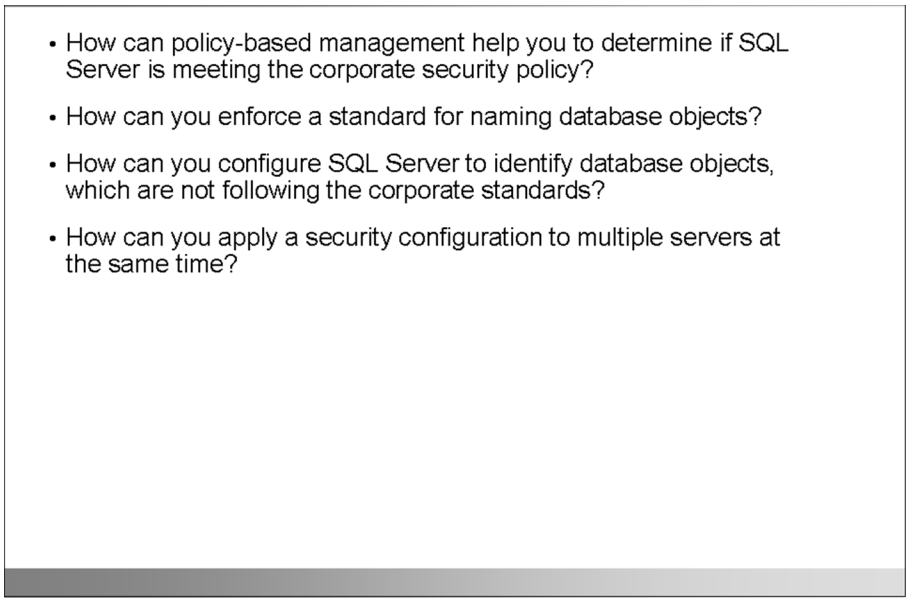
- http://msdn.microsoft.com/en-us/library/cc645723%28SQL.105%29.aspx

For more information, see Monitoring and Enforcing Best Practices by Using Policy-Based Management.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Discussion: Scenarios for Using Policy-Based Management



Key Points

In this discussion, you are presented scenarios in which you are asked to determine the usage of policy-based management. You and your classmates will discuss possible solutions to the scenarios.

Scenario 1

Your security team has asked you to check regularly if the SQL Server that you are managing is meeting the corporate security policy. You want to perform this task regularly with minimum administrative effort.

Question: How can policy-based management help you to determine if SQL Server is meeting the corporate security policy?

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Scenario 2

Recently you realize that in your application there are database objects with different naming conventions. You want to enforce a standard for naming the database objects.

Question: How can you enforce a standard for naming database objects?

Scenario 3

Your development lead wants to know if any developers are creating database objects, which are not following the corporate standards.

Question: How can you configure SQL Server to identify database objects, which are not following the corporate standards?

Scenario 4

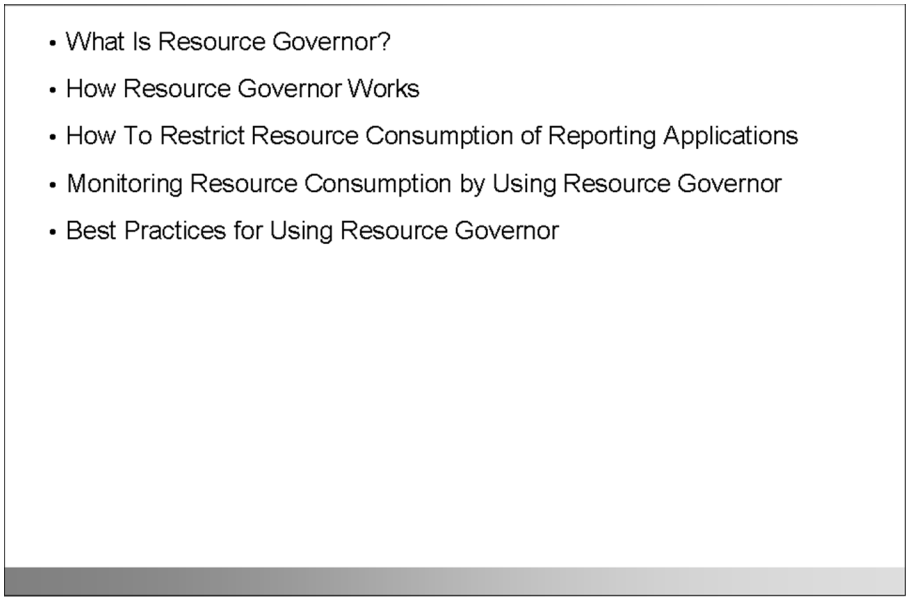
You have several SQL Servers and you want a way to apply a security configuration to all servers at the same time. How can you do this?

Question: How can you apply a security configuration to multiple servers at the same time?

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Lesson 2
Controlling Resource Usage by Using Resource Governor

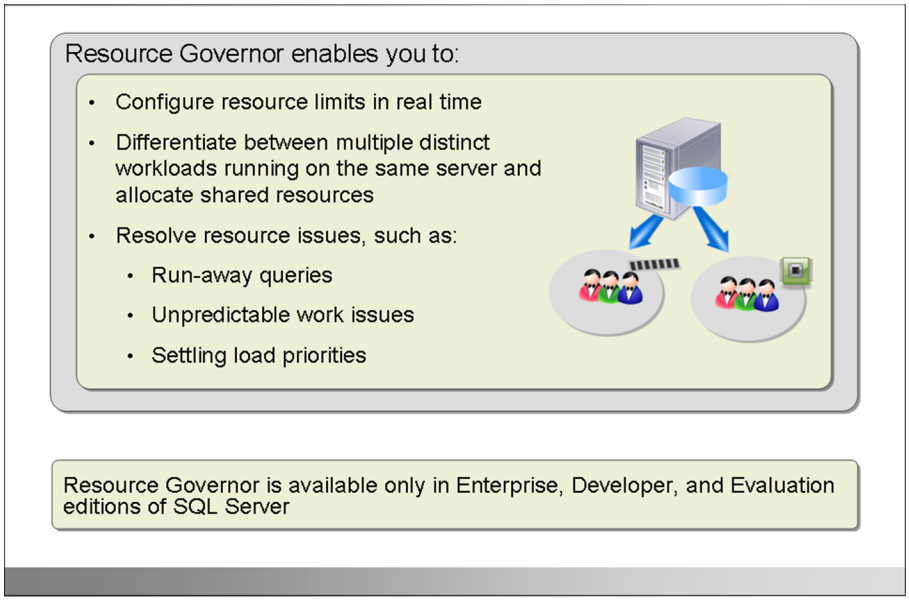


Resource Governor is a new update in SQL Server 2008 that enables you to manage SQL Server workload and resources by specifying limits on resource consumption by incoming requests. You can use Resource Governor to make SQL Server more efficient and to meet Service Level Agreements (SLAs).

In this lesson, you will learn about Resource Governor and how it works. You will also learn how to restrict resource consumption of reporting applications. You will learn to meet SLAs by prioritizing resource groups and the best practices for using Resource Governor.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

What Is Resource Governor?



Key Points

Resource Governor allows organizations to define limits and priorities for different workloads. Resource Governor provides consistent performance to users by enabling concurrent workloads. A workload is a query or request that you can treat as a single entity. You can configure resource limits in real time with minimal impact on executing workloads.

In an environment where multiple distinct workloads are present on the same server, Resource Governor enables you to differentiate these workloads and allocate shared resources as they are requested, based on the limits that you specify. The shared resources refer to CPU and memory.

Types of Resource Issues

You can use Resource Governor to resolve the following types of resource issues, which are common in a database environment:

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

- **Run-away queries on the server**. In this scenario, a resource intensive query can take up most or all of the server resources.
- **Unpredictable workload execution**. In this scenario, concurrent applications on the same server have workloads of different size and type.
- **Setting workload priority**. In this scenario, one workload is allowed to proceed faster than another, or is guaranteed to complete if there is resource contention. Resource Governor enables you to assign relative importance to the workloads.

These issues require the ability to differentiate workloads in some way. To resolve the resource issues, Resource Governor can:

- Classify incoming connections and route their workloads to a specific group.
- Monitor resource usage for each workload in a group.
- Pool resources and set pool-specific limits on CPU usage and memory allocation, to prevent or minimize the probability of run-away queries.
- Associate grouped workloads with a specific pool of resources.
- Identify and set priorities for workloads.

Resource Governor is available only in the Enterprise, Developer, and Evaluation editions of SQL Server.

Disadvantages of Resource Governor

Resource Governor has the following constraints:

- It limits resource management to the SQL Server Database Engine. Resource Governor cannot be used for Analysis Services, Integration Services, and Reporting Services.
- It does not provide workload monitoring or workload management between SQL Server instances.
- It provides limited specification only to CPU bandwidth and memory managed by SQL Server.
- It can manage OLTP workloads but these types of queries, which are typically very short in duration, are not always on the CPU long enough to apply bandwidth controls. This may skew in the statistics returned for CPU usage percentage.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

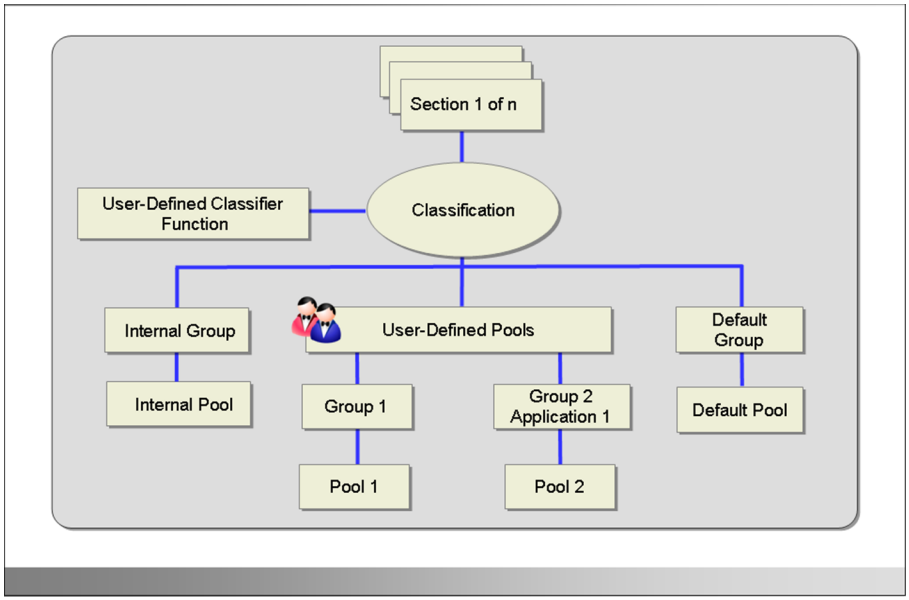
Question: How can you bypass the resources control in Resource Governor?

Source:

- http://technet.microsoft.com/en-us/library/bb895232.aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

How Resource Governor Works



Key Points

The main components of Resource Governor are resource pools, workload groups, and classification.

A resource pool, or pool, represents the physical resources of the server. SQL Server automatically creates and maps two workload groups, internal and default, to their corresponding resource pools when you install SQL Server 2008. Resource Governor also supports user-defined resource pools.

A workload group serves as a container for session requests that are similar according to the classification criteria that are applied to each request. Resource Governor also supports user-defined workload groups.

Classifications are internal rules that classify incoming requests and route them to a workload group. Resource Governor also supports a classifier user-defined function for implementing classification rules.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Resource Pools

The resource pools predefined by Resource Governor are:

- **Internal Pool**. The internal pool represents the resources utilized by SQL Server. This pool always contains only the internal group, and you cannot alter the pool in any way. Resource consumption by the internal pool is not restricted. Any workload in the pool is considered critical for server function. In case of excessive workloads, Resource Governor allocates priority to the internal pool so that it has more resources allocated than other pools.
- **Default Pool**. The default pool is the first predefined user pool. Prior to any configuration, the default pool only contains the default group. You cannot create or drop the default pool, but you can alter it. The default pool can contain user-defined groups in addition to the default group.
- **User-Defined Resource Pools**. Resource Governor provides DDL statements for creating, changing, and dropping resource pools.

A pool has two parts. One part does not overlap with other pools, which enables minimum resource reservation. The other part is shared with other pools, which supports maximum possible resource consumption. You can set the pool resources by specifying one of the following options for each resource:

- MIN or MAX for CPU
- MIN or MAX for memory

MIN and MAX represent the minimum guaranteed resource availability of the pool and the maximum size of the pool, respectively, for each of these resources.

The sum of MIN values across all pools cannot exceed 100 percent of the server resources. You can set the MAX value anywhere in the range between MIN and 100 percent, inclusive.

If you define a nonzero MIN value for a pool, SQL Server readjusts the effective MAX value of other pools as the minimum of the configured MAX value of a pool and subtracts the sum of the MIN values of other pools from 100 percent.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

The shared part of the pool indicates where you can use available resources, if any. However, when resources are allocated they go to the specified pool and you cannot share them. This may improve resource utilization in cases where there are no requests in a given pool and you can free resources configured to the pool for other pools.

Considerations for creating user defined workload groups and resources pools

Resource Governor provides seven DDL commands that you can use to configure all components of the Resource Governor. You can use CREATE, ALTER, and DROP commands for workload groups and resource pools. You can use ALTER RESOURCE GOVERNOR RECONFIGURE statement to apply configuration changes.

The following are the key features of Resource Governor DDL commands:

- The CREATE, ALTER, and DROP statements work on the stored metadata and are transactional. However, completion of these statements does not make the changes effective.
- The preceding statements can be a part of a transaction to prevent simultaneous modifications to the same configuration by multiple users. For example, you can use BEGIN TRANSACTION and COMMIT TRANSACTION with the statements.
- Any locks created on a transaction are visible in sys.dm_tran_locks and the transaction can be stopped if necessary.

Note: You should use transactional syntax when using Resource Governor DDL statements. However, ALTER RESOURCE GOVERNOR DISABLE, ALTER RESOURCE GOVERNOR RECONFIGURE and ALTER RESOURCE GOVERNOR RESET STATISTICS statements cannot be a part of a user transaction.

You can use the ALTER…RECONFIGURE statement to load the metadata into memory. If the ALTER…RECONFIGURE statement fails, the changes to the metadata are not copied to memory and you cannot roll back the statement.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Resource Governor DDL provides various parameters for configuring a resource pool—a minimum and maximum percentage for CPU use, and a minimum and maximum percentage for memory use.

Workload Groups

A workload group allows the aggregate monitoring of resource consumption and the application of a uniform policy to all the requests in the group. A group defines the policies for its members.

Resource Governor predefines two workload groups: the internal group and the default group. A user cannot change anything classified as an internal group, but can monitor it. Requests are classified into the default group when the following conditions exist:

- There are no criteria to classify a request.
- There is an attempt to classify the request into a non-existent group.
- There is a general classification failure.

Resource Governor also provides DDL statements for creating, changing, and dropping workload groups.

Classification

Resource Governor supports the classification of incoming sessions. The classification is based on a set of user-written criteria contained in a function. The results of the function logic enable Resource Governor to classify sessions into existing workload groups.

You can write a scalar function that contains the logic that is used to assign incoming sessions to a workload group. Before you can use this function, you must complete the following actions:

- Create and register the function by using the ALTER RESOURCE GOVERNOR statement.
- Update the Resource Governor configuration by using the ALTER RESOURCE GOVERNOR statement with the RECONFIGURE parameter.

After you create the function and apply the configuration changes, the Resource Governor classifier will use the workload group name returned by the function to send a new request to the appropriate workload group.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

The user-defined function has the following characteristics and functions:

- It is evaluated for every new session, even when connection pooling is enabled.
- It gives workload group context for the session. After group membership is determined, the session is bound to the workload group for the lifetime of the session.
- If the user-defined function returns NULL, default, or the name of a non-existent group, the session is given the default workload group context. The session is also given the default context if the function fails for any reason.
- The function should be defined with server scope (master database).
- The classifier user-defined function designation only takes effect after you run ALTER RESOURCE GOVERNOR RECONFIGURE.
- You can designate only one user-defined function as a classifier at a time.
- You cannot drop or alter the classifier user-defined function unless you remove its classifier status.
- In the absence of a classifier user-defined function, you can classify all sessions into the default group.
- The workload group returned by the classifier function is outside the scope of the schema-binding restriction. For example, you cannot drop a table, but you can drop a workload group.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Working of Resource Governor

- There is an incoming connection for a session (Session 1 of n).
- The session is classified (Classification).
- The session workload is routed to a workload group, for example, Group 1.
- The workload group uses the resource pool it is associated with, for example, Pool 2.
- The resource pool provides and limits the resources required by the application, for example, Application 1.

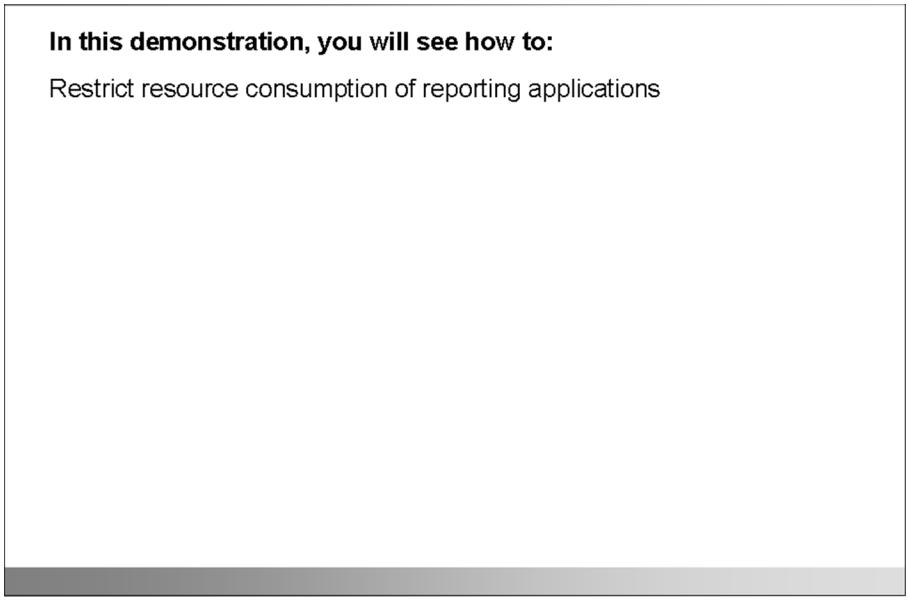
Source:

- http://technet.microsoft.com/en-us/library/bb934084.aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Demonstration: How To Restrict Resource Consumption of Reporting Applications



Key Points

The steps to restrict resource consumption of reporting applications are:

1. In Windows Explorer, browse to the **D:\Demofiles\Mod05** folder, and then double-click the **M5_Demo2.sql** file.

   Note: If the **Connect to Database Engine** dialog box appears, click **Connect.**

2. In the Microsoft SQL Server Management Studio window, in the Query Editor pane, select the **Create two resource pools with maximum CPU percent = 100 and minimum CPU percent = 50** code, and then click **Execute**.

3. In the Microsoft SQL Server Management Studio window, in the Query Editor pane, select the **Create two workload groups with high and low importance** code, and then click **Execute**.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

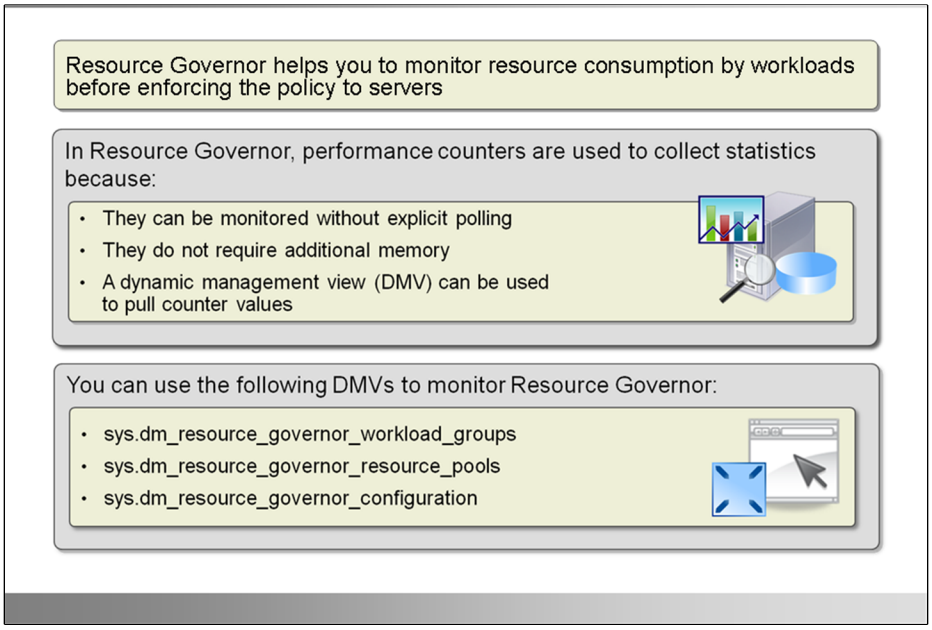
4. In the Microsoft SQL Server Management Studio window, in the Query Editor pane, select the **Load the configuration into memory** code, and then click **Execute**.
5. In the Microsoft SQL Server Management Studio window, in the Query Editor pane, select the **Create a classification function** code, and then click **Execute**.
6. In the Microsoft SQL Server Management Studio window, in the Query Editor pane, select the **Enable resource governor to use the Use the fnGroupClassifier function and apply the configuration changes** code, and then click **Execute**.

Question: How would you create workload groups?

Question: Why should you set the Resource Governor to monitor mode?

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

Monitoring Resource Consumption by Using Resource Governor



Key Points

To monitor resource consumption, you can use Resource Governor. You can configure Resources Governor to generate events for analyzing the impact of policies without applying the policy. The ad-hoc configuration to a long query avoids the impact of an executing query on other queries.

You can use classification functions to assign resources to different resource groups, so that the resources are assigned different priorities.

You can use the sys.dm_resource_governor_workload_groups, sys.dm_resource_governor_resource_pools, sys.dm_resource_governor_configuration dynamic management views (DMV) to monitor the working of Resource Governor.

Monitoring for the SLAs

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

You need to obtain the execution statistics for each workload group and resource pool. Execution statistics enable you to identify workload characteristics that can aid in configuration and detect potential problems early enough to take corrective action.

Performance Counters

In Resource Governor, the majority of statistics collection is based on performance counters for the following reasons:

- Performance counters can be monitored without explicit T-SQL-style polling.
- Performance counters do not require much memory when the total number of expected active workload groups and resource pools is small.

You can use the sys.dm_os_performance_counters dynamic management view to pull counter values. Because it shows raw values, you can use them to calculate average values. You can use existing tools such as Perfmon without any modifications.

The following table shows the performance objects for collecting workload group and resource pool statistics that Resource Governor uses.

Performance counter name	Description
SQLServer:Workload Group Stats	This performance counter reports statistics for each active workload group, such as the number of active requests and the number of blocked requests.
SQLServer:Resource Pool Stats	This performance counter reports statistics for each active resource pool, such as the number of memory grants that are occurring in the resource pool per second and the amount of memory that is used by the resource pool.

Events

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

In addition to using SQL Trace events, Resource Governor introduces new events for monitoring. The following table describes the new events.

Event name	Description
CPU Threshold Exceeded	This event indicates when Resource Governor detects that a query has exceeded the CPU threshold value in REQUEST_MAX_CPU_TIME_SEC.
PreConnect:Starting	This event indicates when a LOGON trigger or Resource Governor classifier function starts execution.
PreConnect:Completed	This event indicates when a LOGON trigger or Resource Governor classifier function finishes execution.

sys.dm_resource_governor_workload_groups View

The sys.dm_resource_governor_workload_groups view returns workload group statistics and the current in-memory configuration of the workload group. This view can be joined with sys.dm_resource_governor_resource_pools to get the resource pool name.

This dynamic management view shows the in-memory configuration. To view the stored configuration metadata, use the sys.resource_governor_workload_groups catalog view.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

When you successfully run the ALTER RESOURCE GOVERNOR RESET STATISTICS, the statistics_start_time, total_request_count, total_queued_request_count, total_cpu_limit_violation_count, total_cpu_usage_ms, max_request_cpu_time_ms, total_lock_wait_count, total_lock_wait_time_ms, total_query_optimization_count, total_suboptimal_plan_generation_count, total_reduced_memgrant_count, out_of_memory_count, and max_request_grant_memory_kb counters are reset. The statistics_start_time counter is set to the current system date and time; the other counters are set to zero.

sys.dm_resource_governor_resource_pools View

The sys.dm_resource_governor_resource_pools view returns information about the current resource pool state, the current configuration of resource pools, and resource pool statistics.

The Resource Governor workload groups and Resource Governor resource pools have a many-to-one relationship. As a result, many of the resource pool statistics are derived from the workload group statistics.

This dynamic management view shows the in-memory configuration. To view the stored configuration metadata, use the sys.resource_governor_resource_pools catalog view.

sys.dm_resource_governor_configuration View

The sys.dm_resource_governor_configuration view returns a row that contains the current in-memory configuration state of Resource Governor.

This dynamic management view shows the in-memory configuration. To view the stored configuration metadata, use the corresponding catalog view.

The following example shows how to get and compare the stored metadata values and the in-memory values of the Resource Governor configuration.

```
Use master

GO

-- Get the stored metadata.
```

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.


```
select
object_schema_name(classifier_function_id) as 'Classifier
UDF schema in metadata',
object_name(classifier_function_id) as 'Classifier UDF
name in metadata'
from
sys.resource_governor_configuration
GO
-- Get the in-memory configuration.
select
object_schema_name(classifier_function_id) as 'Active
classifier UDF schema',
object_name(classifier_function_id) as 'Active classifier
UDF name'
from
sys.dm_resource_governor_configuration
GO
```

Question: How can you use Resource Governor to monitor resources usage without enforcing the resources limitation?

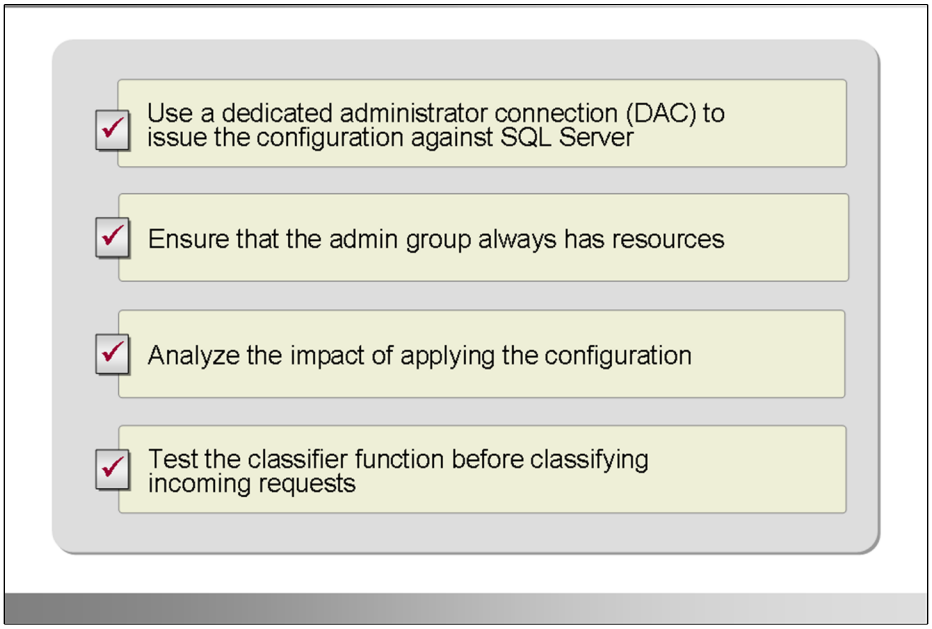
Source:

- http://technet.microsoft.com/en-us/library/bb934023.aspx
- http://technet.microsoft.com/en-us/library/bb895331.aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Best Practices for Using Resource Governor



Key Points

The best practices for using Resource Governor are:

- **Use dedicated administrator connection (DAC) to issue the configuration against SQL Server**. This is because the Resource Governor configuration does not affect DAC. This ensures that the connection used to make the configuration changes will not be affected, even if you make a mistake during the configuration.
- **Ensure that the admin group always has resources**. You can perform administrative tasks even when a system is running out of resources.
- **Analyze the impact of applying the configuration**. Configure the Resource Governor so that it only generates events. Then, monitor the configuration and analyze the impact of applying the configuration.
- **Test the classifier function before classifying incoming**

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

requests. This is because a poorly written function can render the system unusable by timing out, which in turn, can expose configuration information. You can use DAC to troubleshoot a poorly written classifier function because a DAC is not subject to classification and can be used while Resource Governor is running and classifying incoming sessions. You can use existing system functions in your classifier functions and use the values these functions return for classifying sessions. The following system functions can be used for classification:

- HOST_NAME()
- APP_NAME()
- SUSER_NAME()
- SUSER_SNAME()
- IS_SRVROLEMEMBER()
- IS_MEMBER()

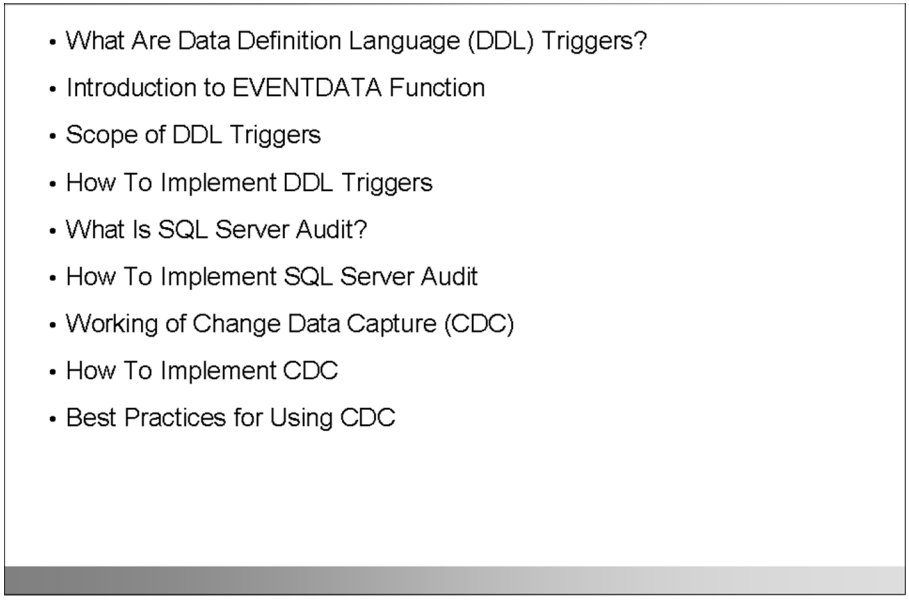
Resource Governor also implements the CONNECTIONPROPERTY function to provide additional properties that you can use for classification. This function returns the same values for the properties for the unique connection that a request came on as the corresponding columns in the sys.dm_exec_connections dynamic management view.

Question: How can you ensure that a classifier function is able to correctly identify an application and assign it to the correct resources?

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Lesson 3
Auditing Database Changes

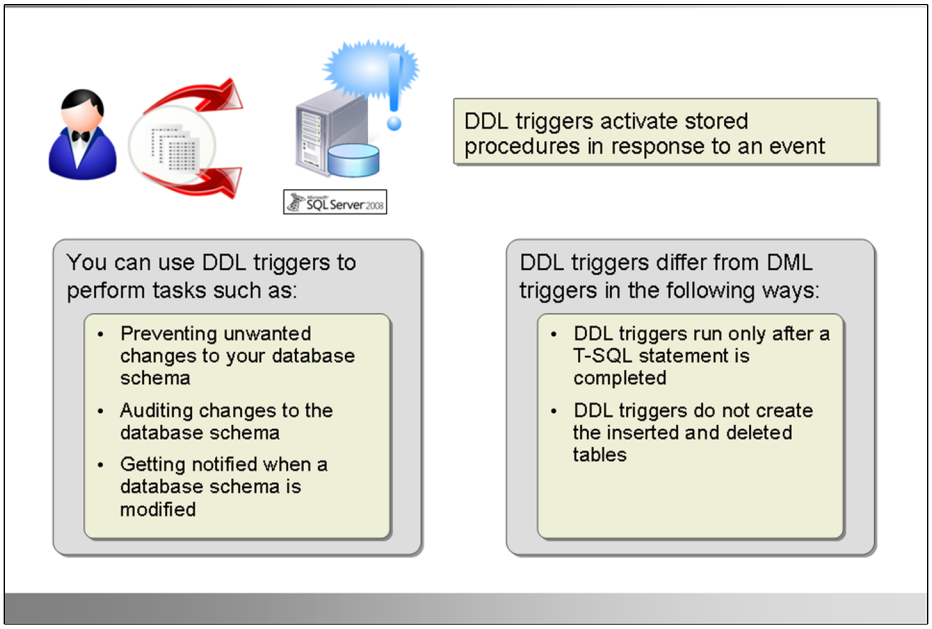


In most organizations unauthorized modification of data can have critical effects. Auditing data is important for processes other than just the retention of data. Auditing an SQL Server database or an instance of SQL Server involves tracking and logging events that occur on the system. You can implement auditing in SQL Server by using DDL triggers. You can track the traditional data changes by using Change Data Capture (CDC).

In this lesson, you will learn about the audit process. You will learn how to implement DDL triggers and SQL Server audit. You will also learn about CDC, ways to implement CDC, and best practices for using CDC.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

What Are DDL Triggers?



Key Points

DDL triggers are similar to regular triggers. They trigger stored procedures in response to an event. However, unlike DML triggers, they do not run in response to UPDATE, INSERT, or DELETE statements on a table or view. Instead, you can use these triggers to run in response to a variety of DDL events. These events primarily correspond to T-SQL statements beginning with keywords CREATE, ALTER, and DROP. You can also use certain system-stored procedures performing operations similar to DDL, to run DDL triggers.

Benefits of DDL triggers

You can use DDL triggers to perform administrative tasks such as auditing and regulating database operations. You can use DDL triggers to:

- Prevent unwanted changes to your database schema.
- Audit changes to the database schema

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

- Get notified when a database schema is modified.

You can run DDL triggers only after the DDL statements that are used to trigger it are run.

The following example shows how you can use a DDL trigger to prevent any table from being modified or deleted.

```
USE QuantamCorp2008

GO

CREATE TABLE NewTable (Column1 int);

GO

CREATE TRIGGER dropTableTrigger

ON DATABASE

FOR DROP_TABLE

AS

   PRINT 'You must disable Trigger "safety" to drop
tables!'

   ROLLBACK;

GO

DROP TABLE NewTable
```

Understanding DDL Triggers and DML Triggers

DDL triggers and DML triggers are used for different purposes.

DDL triggers operate on CREATE, ALTER, DROP, and other DDL statements and stored procedures that perform DDL-like operations. You can use them to perform administrative tasks and enforce rules that affect databases. They apply to all commands of a single type across a database or across a server.

You can create, modify, and drop DML triggers and DDL triggers by using similar T-SQL syntax. Also, like DML triggers, DDL triggers can be nested.

Like DML triggers, DDL triggers can run managed code packaged in an assembly that was created in the Microsoft .NET Framework and uploaded in SQL Server.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Like DML triggers, you can create more than one DDL trigger on the same T-SQL statement. Also, a DDL trigger and the statement that triggers it are run within the same transaction. You can roll back this transaction from within the trigger. A transaction can automatically roll back when an error occurs. DDL triggers that are run from a batch explicitly include the ROLLBACK TRANSACTION statement that cancels the whole batch.

Note: You cannot roll back an ALTER DATABASE event that occurs within a DDL trigger.

When you are designing DDL triggers, consider how they differ from DML triggers in the following ways:

- DDL triggers run only after a T-SQL statement is completed. DDL triggers cannot be used as INSTEAD OF triggers.
- DDL triggers do not create the inserted and deleted tables. You can capture the information about an event that triggers a DDL trigger, and the subsequent changes caused by the trigger by using the EVENTDATA function.

DDL triggers or event notifications are created as a response to the following events when they occur in the database in which the trigger or event notification is created, or anywhere in the server instance. There are three types of DDL events—CREATE_APPLICATION_ROLE, ALTER_APPLICATION_ROLE, and DROP_APPLICATION_ROLE.

- **CREATE_APPLICATION_ROLE**. Applies to the CREATE APPLICATION ROLE statement and sp_addapprole. If a new schema is created, this event also triggers a CREATE_SCHEMA event.

  The following are the examples of the CREATE_APPLICATION_ROLE event:

  - CREATE_ASSEMBLY
  - CREATE_DATABASE
  - CREATE_ASYMMETRIC_KEY
  - ALTER_AUTHORIZATION
  - CREATE_CERTIFICATE
  - CREATE_CONTRACT

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

- ADD_COUNTER_SIGNATURE
  - CREATE_CREDENTIAL
  - GRANT_DATABASE
  - CREATE_DEFAULT
  - CREATE_EVENT_NOTIFICATION
  - CREATE_FUNCTION
  - CREATE_INDEX
  - CREATE_MASTER_KEY
  - CREATE_MESSAGE_TYPE
  - CREATE_PARTITION_FUNCTION
- **ALTER_APPLICATION_ROLE**. Applies to the ALTER APPLICATION ROLE statement and **sp_approlepassword**.

  The following are the examples of the ALTER_APPLICATION_ROLE event:
  - ALTER_ASSEMBLY
  - ALTER _DATABASE
  - ALTER_ASYMMETRIC_KEY
  - ALTER_CERTIFICATE
  - DROP_CONTRACT
  - DROP_COUNTER_SIGNATURE
  - ALTER_CREDENTIAL
  - DENY_DATABASE
  - DROP_DEFAULT
  - DROP_EVENT_NOTIFICATION
  - ALTER_FUNCTION
  - ALTER_MASTER_KEY
  - ALTER_MESSAGE_TYPE
  - ALTER_PARTITION_FUNCTION

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

- **DROP_APPLICATION_ROLE**. Applies to the DROP APPLICATION ROLE statement and **sp_dropapprole**.

  The following are the examples of the DROP_APPLICATION_ROLE event:

  - DROP_ASSEMBLY
  - DROP_DATABASE
  - DROP_ASYMMETRIC_KEY
  - DROP_CERTIFICATE
  - DROP_CREDENTIAL
  - REVOKE_DATABASE
  - DROP_EXTENDED_PROPERTY
  - DROP_FULLTEXT_CATALOG
  - DROP_FULLTEXT_INDEX
  - DROP_FUNCTION
  - DROP_INDEX
  - DROP_MASTER_KEY
  - DROP_MESSAGE_TYPE

  Question: What is the key difference between a DDL trigger and a DML trigger in terms of obtaining information about the change in the database?

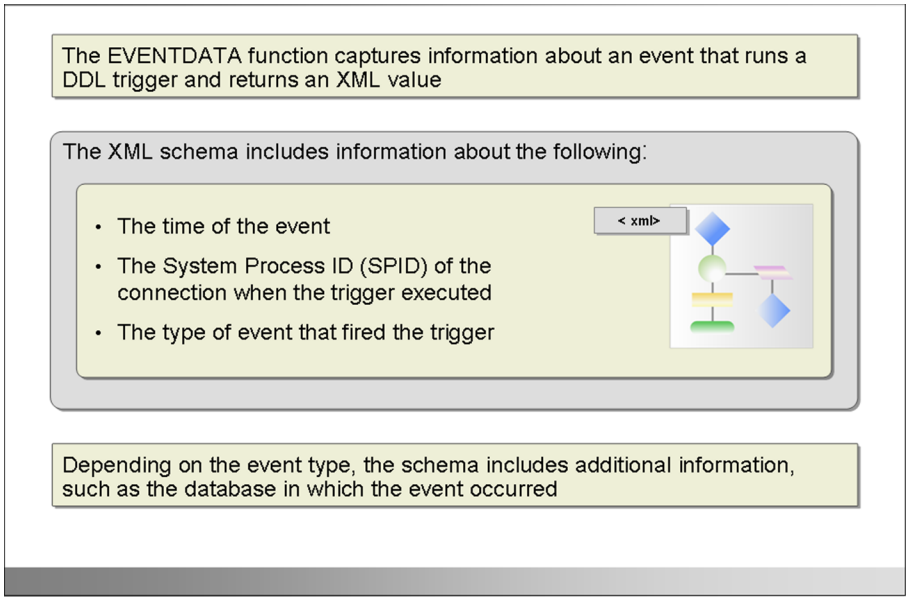
  Source:

  - http://msdn.microsoft.com/en-us/library/ms189599%28SQL.105%29.aspx
  - http://msdn.microsoft.com/en-us/library/ms175941.aspx
    http://msdn.microsoft.com/en-us/library/ms189599.aspx
  - http://msdn.microsoft.com/en-us/library/bb522542.aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Introduction to EVENTDATA Function



Key Points

You can use the EVENTDATA function to capture details of an event that is run by a DDL trigger. The EVENTDATA function returns an XML value. The XML schema includes the time of the event, the System Process ID (SPID) of the connection when the trigger was run, and the type of event that fired the trigger.

The schema includes additional information based on the event such as the database where it occurred, the object against which it occurred, and its T-SQL statement. The following example shows the DDL trigger created in the *QuantamCorp* database.

```
USE QuantamCorp2008

GO

CREATE TRIGGER createTableTrigger

ON DATABASE

FOR CREATE_TABLE
```

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.


```
AS
    PRINT 'CREATE TABLE Issued.'
    SELECT
EVENTDATA().value('(/EVENT_INSTANCE/TSQLCommand/CommandTex
t)[1]','nvarchar(max)')
   RAISERROR ('New tables cannot be created in this
database.', 16, 1)
   ROLLBACK
;
GO
```

The following CREATE TABLE statement is run, after you create the trigger.

```
CREATE TABLE NewTable (Column1 int);
GO
```

The EVENTDATA() statement in the DDL trigger captures the restricted text of the CREATE TABLE statement. You can capture the text by running an XQuery statement against the XML data that is generated by EVENTDATA and retrieving the <CommandText> element.

A key consideration is that EVENTDATA captures data of the CREATE_SCHEMA events along with <schema_element> of the corresponding CREATE SCHEMA definition, if any exists. In addition, EVENTDATA distinguishes <schema_element> from other events. Therefore, a DDL trigger created on both a CREATE_SCHEMA event, and an event represented by the <schema_element> of the CREATE SCHEMA definition, may return the same event data twice, such as the TSQLCommand data. For example, consider a DDL trigger that is created on both the CREATE_SCHEMA and CREATE_TABLE events and the following batch is run:

```
CREATE SCHEMA DemoSchema
GO
CREATE TABLE DemoTable1 (col1 int)
GO
```

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

You should be aware when the application retrieves the TSQLCommand data of the CREATE_TABLE event, because the TSQLCommand data may appear two times. It may appear once when the CREATE_SCHEMA event occurs, and again when the CREATE_TABLE event occurs. You should avoid creating DDL triggers on both the CREATE_SCHEMA events and the <schema_element> texts of any corresponding CREATE SCHEMA definitions. You should also avoid building application logic to prevent the same event from being processed twice.

Using XQuery Path Expression to query xml data returned from EVENTDATA()

XQuery path expressions locate nodes, such as element, attribute, and text nodes, in a document. The result of a path expression always occurs in the document order without duplicate nodes in the result sequence. By specifying a path, you can use either unabbreviated or abbreviated syntax.

A path expression can be relative or absolute.

- **Relative path expression**. A relative path expression is made up of one or more steps separated by one or two slash marks (/ or //). For example, child::Features, is a relative path expression, where child refers only to child nodes of the context node that is currently under process. The expression retrieves the <Features> element node children of the context node.

- **Absolute path expression** An absolute path expression starts with one or two slash marks (/ or //), followed by an optional, relative path. For example, the initial slash marks in the expression, /child::ProductDescription, indicates that it is an absolute path expression. Because a slash mark at the start of an expression returns the document root node of the context node, the expression returns all the <ProductDescription> element node children of the document root.
  If an absolute path starts with a single slash mark, it may or may not be followed by a relative path. If you specify only a single slash mark, the expression returns the root node of the context node.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

You can create a path expression in two steps. For example, to create an absolute path expression, /child::ProductDescription/child::Summary, you need to:

- Retrieve the <ProductDescription> element node children of the document root.
- Retrieve the <Summary> element node children for each retrieved <ProductDescription> element node, which in turn becomes the context node.

To return event data, you should use the XQuery **value()** method instead of the **query()** method. The **query()** method returns XML and ampersand-escaped carriage return and line-feed (CRLF) instances in the output, while the **value()** method renders CRLF instances invisible in the output.

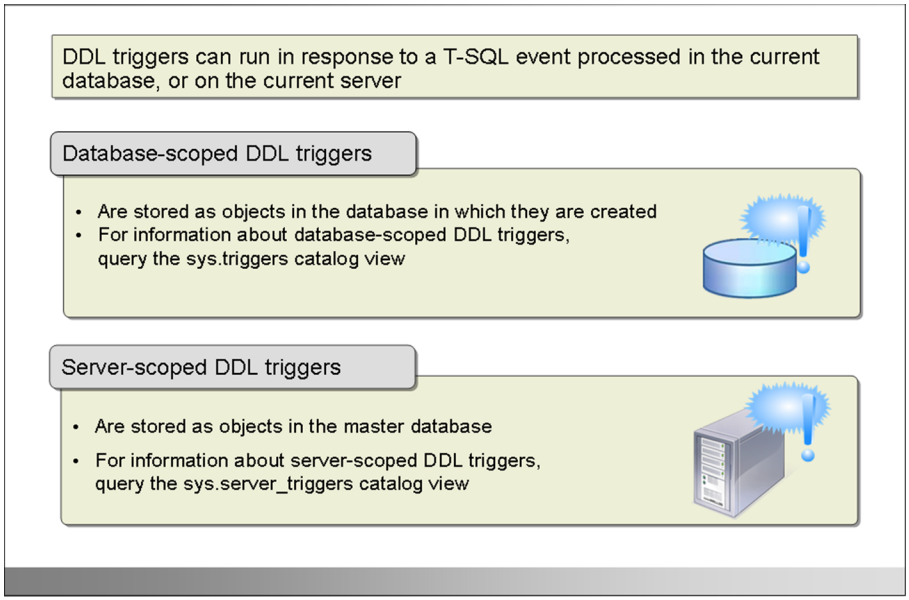
Question: How can you obtain information from the XML data returned by the EVENTDATA function?

Source:

- http://msdn.microsoft.com/en-us/library/ms173781.aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

Scope of DDL Triggers



Key Points

DDL triggers can run in response to a T-SQL event processed in the current database, or on the current server. The scope of the trigger depends on the event. For example, a DDL trigger created to run in response to a CREATE_LOGIN event can do so only when a CREATE_LOGIN event occurs in the server.

The following example shows how a DDL trigger will run whenever a DROP_TABLE or ALTER_TABLE event occurs in the database.

```
--Scope of DDL Triggers

USE QuantamCorp2008

GO

CREATE TRIGGER preventChangeTrigger ON DATABASE FOR
DROP_TABLE, ALTER_TABLE

AS
```

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.


```
   PRINT 'You must disable Trigger "safety" to drop or
alter tables!'
   ROLLBACK
;
```

The following example shows how a DDL trigger prints a message, if any CREATE_DATABASE event occurs on the current server instance. The example uses the EVENTDATA function to retrieve the text of the corresponding T-SQL statement.

```
USE QuantamCorp2008
GO
IF EXISTS (SELECT * FROM sys.server_triggers
    WHERE name = 'ddl_trig_database')
DROP TRIGGER ddl_trig_database
ON ALL SERVER;
GO
CREATE TRIGGER ddl_trig_database
ON ALL SERVER
FOR CREATE_DATABASE
AS
    PRINT 'Database Created.'
    SELECT
EVENTDATA().value('(/EVENT_INSTANCE/TSQLCommand/CommandTex
t)[1]','nvarchar(max)')
GO
DROP TRIGGER ddl_trig_database
ON ALL SERVER;
GO
```

Database-scoped DDL triggers

Database-scoped DDL triggers are stored as objects in the database in which they are created. You can create DDL triggers in the master

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

database and make them perform just like those created in user-designed databases. You can obtain information about DDL triggers by querying the sys.triggers catalog view. You can query sys.triggers within the database context in which you create the triggers, or you can specify the database name as an identifier, such as master.sys.triggers.

Server-scoped DDL triggers

Server-scoped DDL triggers are stored as objects in the master database. However, you can obtain information about server-scoped DDL triggers by querying the sys.server_triggers catalog view in any database context.

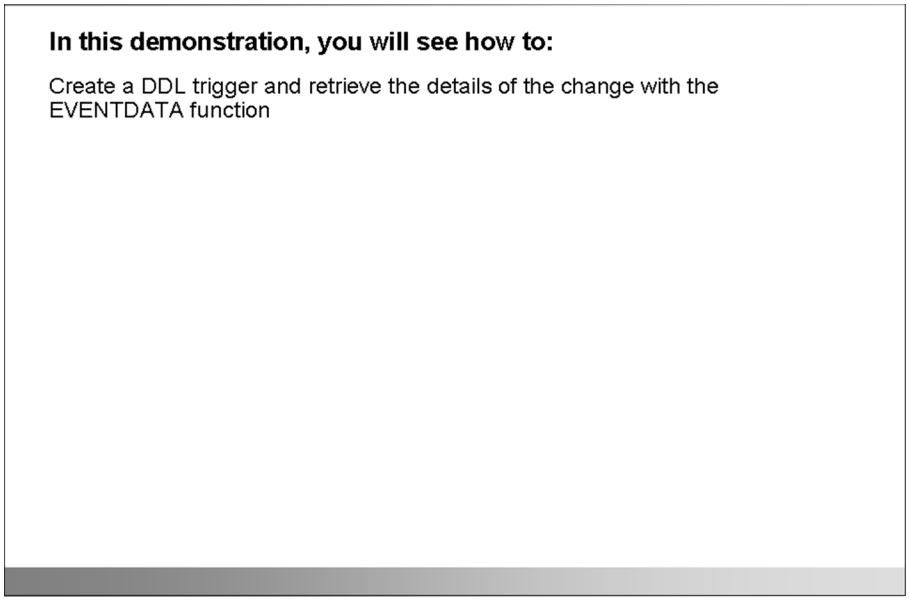
Source:

- http://msdn.microsoft.com/en-us/library/ms186406%28SQL.105%29.aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Demonstration: How To Implement DDL Triggers



Key Points

The steps to create a DDL trigger and retrieve the details of the change with the EVENTDATA function are:

1. In Windows Explorer, browse to the **D:\Demofiles\Mod05** folder, and then double-click the **M5_Demo3.sql** file.

   Note: If the **Connect to Database Engine** dialog box appears, click **Connect.**

2. In the Microsoft SQL Server Management Studio window, in the Query Editor pane, select the **Create a table DDL_Log, a trigger Triggerlog, and then log the event information to the DDL_Log table** code, and then click **Execute**.

3. In the Microsoft SQL Server Management Studio window, in the Query Editor pane, select the **Create and drop the TestTable** table code, and then click **Execute**.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

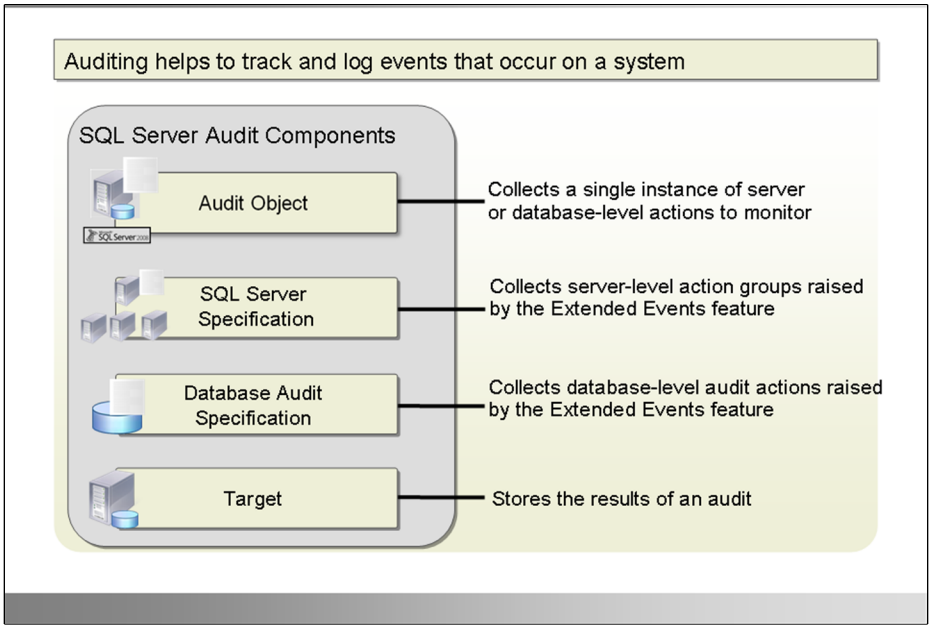
4. In the Microsoft SQL Server Management Studio window, in the Query Editor pane, select the **Retrieve data stored in the ddl_log table** code, and then click **Execute**.

Question: What is the difference between a DDL trigger and a DML trigger?

Question: How will you obtain event information about an event that triggers the DDL Trigger?

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

What Is SQL Server Audit?



Key Points

Auditing an instance of SQL Server or a SQL Server database involves tracking and logging events that occur on the system. You can also set up automatic auditing by using SQL Server Audit in the Enterprise Edition of SQL Server.

There are several levels of auditing for SQL Server. The level of auditing depends on the regulatory or standards requirements for the installation. SQL Server Audit provides the tools and processes to enable, store, and view audits on various server and database objects.

You can record server audit action groups per instance and either database audit action groups or database audit actions per database. The audit event will occur every time that the auditable action is encountered.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

SQL Server Audit Components

An audit is the combination of several elements into a single package for a specific group of server actions or database actions. The components of SQL Server Audit combine to produce an output that is called an audit. There are four components of an audit. They are:

- **Audit Object**. The SQL Server audit object collects a single instance of server or database-level actions and groups of actions to monitor. The audit is at the SQL Server instance level. You can have multiple audits per SQL Server instance. When you define an audit, you specify the location for the output of the results. This is the audit destination. When you create an audit, it is disabled and does not automatically audit any actions. After you enable the audit, the audit destination receives data from the audit.

- **Server Audit Specification**. You can create one server audit specification per SQL Server Audit because it is created at the SQL Server instance scope level along with the audit. The server audit specification collects many server-level action groups raised by the Extended Events feature. You can include audit action groups in a server audit specification. Audit action groups are predefined groups of actions, which are atomic events occurring in the Database Engine. These actions are sent to the audit, which records them in the target.

- **Database Audit Specification**. You can create one database audit specification per SQL Server database per audit. The database audit specification collects database-level audit actions raised by the Extended Events feature. You can add either audit action groups or audit events to a database audit specification. Audit events are the atomic actions that can be audited by the SQL Server engine. Audit action groups are predefined groups of actions. Both are defined at the SQL Server database scope. These actions are sent to the audit, which records them in the target.

- **Target**. The results of an audit are sent to a target, which can be a file, the Windows Security event log, or the Windows Application event log. Logs must be reviewed and archived periodically to ensure that the target has sufficient space to write

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

additional records. Writing to the Windows Security log requires the SQL Server service account to be added to the Generate security audits policy. By default, the Local System, Local Service, and Network Service are part of this policy. Additionally, the Audit object access security policy must be enabled for both Success and Failure. In Windows Vista® or Windows Server® 2008, you can set the more granular application-generated policy from the command line by using the audit policy program (AuditPol.exe). The Windows event logs are global to the Windows operating system. If you need more precise permissions on the audit, use the binary file target.

When you are saving audit information to a file, to help prevent tampering, you can restrict access to the file location in the following ways:

- The SQL Server Service Account must have both the Read and Write permission.
- Audit Administrators typically require the Read and Write permission.
- Audit Readers that are authorized to read audit files must have the Read permission.

Even when the Database Engine is writing to a file, other Windows users can read the audit file if they have permission. The Database Engine does not take an exclusive lock that prevents read operations.

Because the Database Engine can access the file, SQL Server logons that have CONTROL SERVER permission can use the Database Engine to access the audit files. To record any user that is reading the audit file, define an audit on master.sys.fn_get_audit_file. This record the logons with CONTROL SERVER permission that have accessed the audit file through SQL Server.

If an Audit Administrator copies the file to a different location, the Audit Control Lists on the new location should be reduced to the following permissions:

- Audit Administrator—Read/Write
- Audit Reader—Read

You must generate audit reports from a separate instance of SQL Server, such as an instance of SQL Server Express, to which only Audit Administrators or Audit Readers have access. By using a

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

separate instance of the Database Engine for reporting, you can help prevent unauthorized users from obtaining access to the audit record.

You can offer additional protection against unauthorized access by encrypting the folder in which the audit file is stored by using Windows BitLocker™ Drive Encryption or Windows Encrypting File System.

Source:

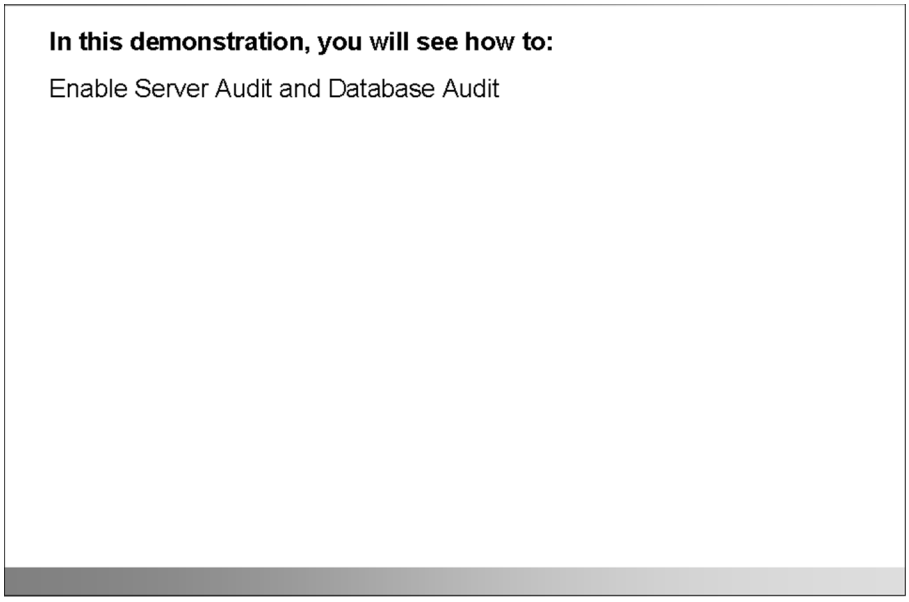
- http://msdn.microsoft.com/en-us/library/cc280386.aspx

For more information, see SQL Server Audit Action Groups and Actions.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Demonstration: How To Implement SQL Server Audit



Key Points

The steps to enable Server Audit are:

1. In Windows Explorer, browse to the **D:\Demofiles\Mod05** folder, and then double-click the **M5_Demo4.sql** file.

   Note: If the **Connect to Database Engine** dialog box appears, click **Connect.**

2. In the Microsoft SQL Server Management Studio window, in the Query Editor pane, select the **Create audit entries to the Windows application log, log the failed logon, and enable server** code, and then click **Execute**.
3. On the **Start** menu, point to **Administrative Tools**, and then click **Event Viewer**. The Event Viewer window appears.
4. In the Event Viewer window, in the **navigator** pane, under **Windows Logs**, in the **Applications** window, verify the audit entries.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

The steps to enable Database Audit are:

1. In Windows Explorer, browse to the **D:\Demofiles\Mod05** folder, and then double-click the **M5_Demo4.sql** file.
2. In the Microsoft SQL Server Management Studio window, in the Query Editor pane, select the **Add audit entries to the Windows application log, log INSERT operations, and enable server auditing** code, and then click **Execute**.
3. In the Event Viewer window, in the **navigator** pane, under **Windows Logs** then on **Applications window**, and then verify the audit entries.

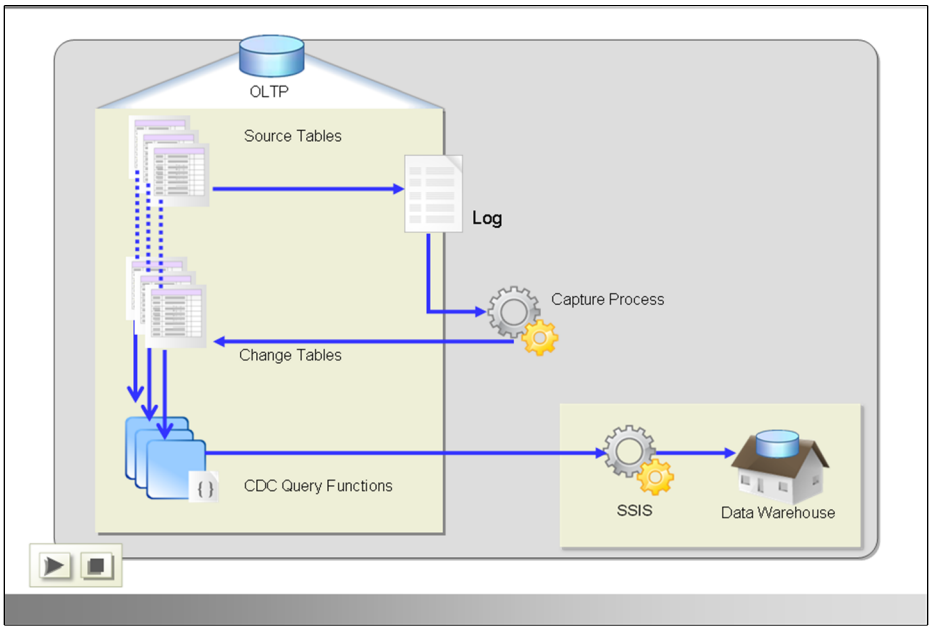
Question: What are the benefits of using SQL Server Audit?

Question: How can you read the audit information generated by SQL Server Audit?

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Working of CDC



Key Points

CDC records insert, update, and delete activity applied to a SQL Server table and stores the changes in a relational format.

CDC Data Flow

1. The source of change data for CDC is the SQL Server transaction log.
2. When you insert, update, or delete records from tracked source tables, entries that describe those changes are added to the log.
3. This log is the input for the CDC capture process. This reads the log and adds information about changes to the tracked table's associated change table.

   Column information and the metadata that is required to apply the changes to a target environment is captured for the modified rows and stored in change tables that mirror the column structure of the tracked source tables.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

4. Functions are provided to enumerate the changes that appear in the change tables over a specified range, returning the information in the form of a filtered result set.
5. The filtered result set is typically used by an application process to update a representation of the source in some external environment.

CDC is available only on Enterprise, Developer, and Evaluation Editions of SQL Server.

For example, consider the extraction, transformation, and loading (ETL) application. An ETL application incrementally loads change data from SQL Server source tables to a data warehouse or data mart. Although the representation of the source tables within the data warehouse must reflect changes in the source tables, an end-to-end technology that refreshes a replica of the source is not appropriate. Instead, you need a reliable stream of change data that is structured so that consumers can apply it to dissimilar target representations of the data. SQL Server change data capture provides this technology.

Understanding CDC and the Capture Instance

Before you track changes to individual tables within a database, you need to enable CDC for the database. You can use the stored procedure sys.sp_cdc_enable_db to enable CDC. When you enable the database, you can identify source tables as tracked tables by using the stored procedure sys.sp_cdc_enable_table. When you enable CDC for a table, an associated capture instance is created to support the dissemination of the change data in the source table. The capture instance consists of a change table and up to two query functions. Metadata that describes the configuration details of the capture instance is retained in the change data capture metadata tables cdc.change_tables, cdc.index_columns, and cdc.captured_columns. You can retrieve this information by using the stored procedure sys.sp_cdc_help_change_data_capture.

All objects that are associated with a capture instance are created in the change data capture schema of the enabled database. The capture instance name should be a valid object name, and should be unique across the database capture instances. By default, the name is <schema name_table name> of the source table. You can name its associated change table by appending _CT to the capture instance name. You can name the function that is used to query for all changes by prepending fn_cdc_get_all_changes_ to the capture

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

instance name. If you configure the capture instance to support net changes, you also create and name the net_changes query function by prepending fn_cdc_get_net_changes_ to the capture instance name.

Data Type Considerations for CDC

CDC supports all base column types. The following table lists the changes captured in change tables and limitations for several column types.

Type of Column	Changes captured in change table	Limitation
Sparse Columns	Yes	Does not support capturing changes when using a columnset.
Computed Columns	No	Changes to computed columns are not tracked. The column will appear in the change table with the appropriate type, but will have a value of NULL.
XML	Yes	Changes to individual XML elements are not tracked.
Timestamp	Yes	The data type in the change table is converted to binary.
BLOB data types	Yes	The previous image of the BLOB column is stored only if the column itself is changed.

Question: What is the difference between CDC and SQL Server audit?

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Question: How can you use the data captured in CDC?

Source:

- http://msdn.microsoft.com/en-us/library/cc645937.aspx

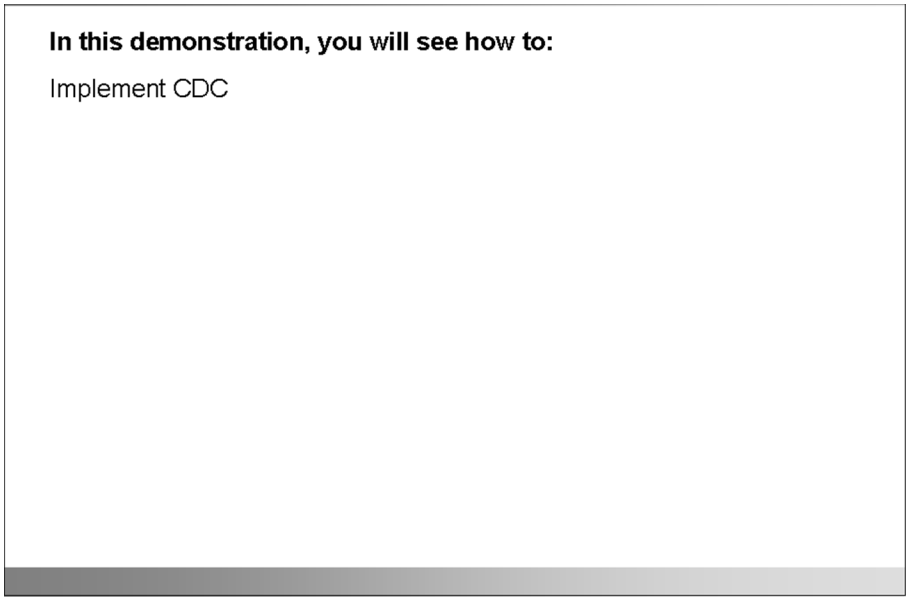
For more information, see sys.dm_cdc_log_scan_sessions (Transact-SQL).

For more information, see sys.dm_cdc_errors (Transact-SQL).

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Demonstration: How To Implement CDC



Key Points

The steps to implement CDC are:

1. In Windows Explorer, browse to the **D:\Demofiles\Mod05** folder, and then double-click the **M5_Demo5.sql** file.

   Note: If the **Connect to Database Engine** dialog box appears, click **Connect.**

2. In the Microsoft SQL Server Management Studio window, in the Query Editor pane, select the **Create TestTable and enable change data capture on it** code, and then click **Execute**.
3. In the Microsoft SQL Server Management Studio window, in the Query Editor pane, select the **Make changes to the TestTable** code, and then click **Execute**.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

4. In the Microsoft SQL Server Management Studio window, in the Query Editor pane, select the **Get all changes captured in cdc.fn_cdc_get_all_changes_dbo_TestTable** code, and then click **Execute**.

Note: There are four rows returned and each row corresponds to a change that was implemented previously.

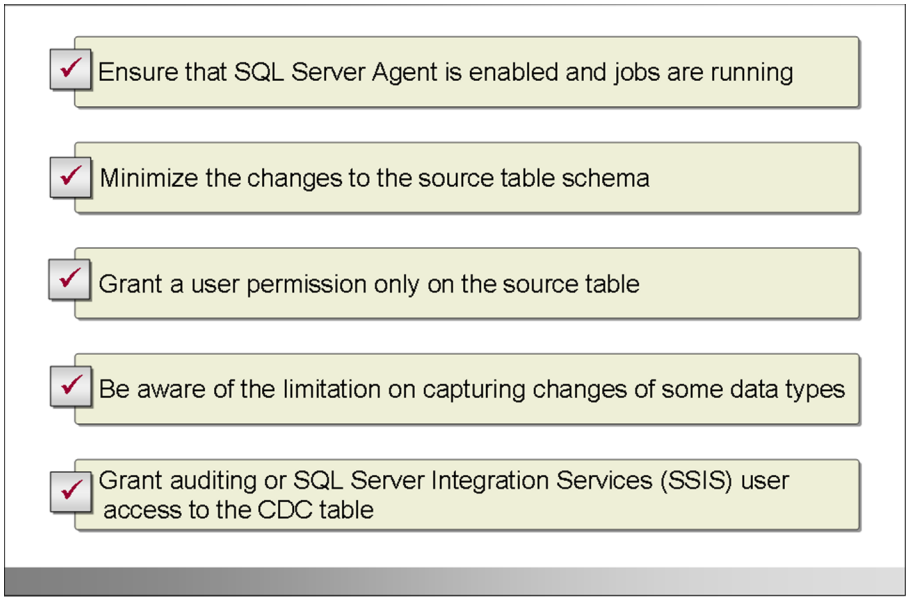
Question: What is the difference between the captured data and the original data?

Question: Why should SQL Server Agent be run before you can use CDC?

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Best Practices for Using CDC



Key Points

When using CDC, consider the following best practices:

- **Ensure that SQL Server agents are enabled and jobs are running**. Two SQL Server Agent jobs are typically associated with a CDC-enabled database: one that is used to populate the database change tables, and the other that is responsible for change table cleanup. Both jobs consist of a single step that runs a T-SQL command. The T-SQL command that SQL Server invokes is a CDC-defined stored procedure that implements the logic of the job. SQL Server creates the jobs when you enable the first table of the database for CDC. SQL Server always creates the cleanup job. SQL Servers creates the capture job only if there are no defined transactional publications for the database. It also creates the capture job when you enable both CDC and transactional replication for a database, and removes the transactional log reader job because the database no longer has defined publications.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

- **Minimize the changes to the source table schema**. CDC helps to mitigate the effect on consumers by allowing the delivered result sets that are returned through the application programming interface (API) to remain unchanged, even as the column structure of the underlying source table changes. This fixed column structure is also reflected in the underlying change table that the defined query functions access.
- **Grant user permission on the source table**. To gain access to the change data that is associated with a capture instance, the user must be granted select access to all captured columns of the associated source table.
- **Be aware of the limitation on capturing changes of some data type**. CDC supports all base column types.
- **Grant auditing or SQL Server Integration Services (SSIS) user access to the CDC table**. Permissions granted to monitored table are not transferred to the CDC table automatically. You need to grant permissions to the CDC table, separately, if you want to allow user access to the data.

Question: How would you manage the permission in CDC-enabled table?

Question: When should you use SSIS with CDC?

Source:

- http://msdn.microsoft.com/en-us/library/cc645937.aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Lab 5: Designing Solutions for Managing SQL Server 2008

- Exercise 1: Designing a Management Strategy for SQL Server 2008 (Discussion)
- Exercise 2: Configuring Policies by Using Policy-Based Management
- Exercise 3: Governing Resources Consumption
- Exercise 4: Implementing DDL Triggers
- Exercise 5: Implementing CDC

Logon Information

Virtual machine	50400A-NYC-SQL1
User name	Administrator
Password	Pa$$w0rd

Estimated time: 90 minutes

Objectives:

After completing this lab, you will be able to:

- Design a management strategy for SQL Server 2008.
- Configure policies by using policy-based management.
- Govern consumption of resources.
- Implement DDL triggers.
- Implement CDC.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Scenario

QuantamCorp has deployed two SQL Server 2008 servers in their organization and will deploy more to replace their existing installation of SQL Server 2005. SQL Server 2008 will be the primary database engine used in the organization. The application team has decided to let more users use their application internally. This may cause problems if a single application database consumes additional resources of the SQL Server 2008 server. In addition, any minor changes made to the database may cause a long downtime beforethe database is up and running again. You need to monitor thedatabase structure to prevent accidental changes, while protectingsensitive information on the database. You need to create a plan toperform the following:

- Determine when to implement CDC.
- Determine how CDC can help in meeting compliance requirements.
- Determine how Resources Governor helps to ensure database resources allocation.
- Determine when to use CDC or policy or DDL triggers to enforce compliance requirements.
- Determine the benefits of implementing centralized database management with policy-based administration.
- Implement CDC to monitor database schema changes.
- Implement a policy to prevent unauthorized database schema changes.
- Implement a policy to enforce corporate standards.
- Implement DDL triggers to enforce custom logging and enforcements on corporate standards.
- Implement Resources Governor to ensure resources allocation.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Exercise 1: Designing a Management Strategy for SQL Server 2008 (Discussion)

Scenario

You are a SQL Server administrator at QuantamCorp. Recently, a system failure has occurred and you spent a lot of time in recovering the system. After a lot of time spent in investigating the cause, you realize that the failure is because developers run long queries against the production system. The query consumes over 90 percent of resources when it runs. This prevents other applications from updating the data in the database.

You company hires a security consultant to study the security of the data stored in SQL Servers managed by you and your team. The security consultant recommends that you should ensure that you log all changes that occur on certain databases, such as the database that stores customer information. This is essential as you must know the modifications made to the database. This will ensure that you have a backup of your data, in the event of a data loss. Keeping track of the modifications made to the database will also help you maintain a regulatory requirement that was recently updated.

In addition, the security consultant recommends that your organization should implement restrictions on database actions for selected staff members in your organization. This will help you minimize the risk associated with database objects that may get accidentally deleted by users. Moreover, you should ensure that all databases are created in a standard manner.

Based on the above scenario, answer the following questions:

1. What should you do to ensure that all changes in the database tables are logged without changing the existing table schema and applications?
2. Your management needs an immediate solution to prevent the system failure from occurring again. In addition, the existing application should not be updated. What solution should you propose?
3. The security consultant wants to prevent developers from making unauthorized changes to the database schema. What would you suggest to resolve this concern?

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

4. How would you ensure that all databases are created in a standardized way?

Exercise 2: Configuring Policies by Using Policy-Based Management

Scenario

You have over 100 servers in your environment that you need to manage. Recently, the company directors have announced a new corporate policy. Based on this policy, you need to ensure that all database objects created follow the defined naming standards. You want to use policy-based management to enforce this policy to all servers simultaneously.

The main tasks for this exercise are as follows:

1. Create a test database.
2. Create a condition.
3. Create a policy category and subscribe it to a database.
4. Create a policy within the policy category.
5. Test the policy.
6. Delete the test database.

Task 1: Create a test database.

1. Open the Microsoft SQL Server Management Studio window and connect to the server, with the following details:
   - Server type: **Database Engine**
   - Server name: **NYC-SQL1**
   - Authentication: **Windows Authentication**
2. In the Query Editor pane of the Microsoft SQL Server Management Studio window, run a Transact-SQL (T-SQL) script to create a test database named **Mod05** with the following structure.

File name	Settings
Mod05_Primary.mdf	Current size: 10 Maximum size: 50

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

File name	Settings
	Filegrowth: 5
Mod05_Log.ldf	Current size: 10 Maximum size: 50 Filegrowth: 5

Note: Browse to the **D:\Labfiles\Mod05** folder and select the **E2_T1.sql** file to execute the required T-SQL code.

Task 2: Create a condition.

1. Using the Microsoft SQL Server Management Studio window, create a new condition, **Mod05 Tables**, with the following details:
   - Facet: **Multipart Name**
   - Field: **@Name**
   - Operator: **LIKE**
   - Force all table names to begin with the letters, **tb1**.

Task 3: Create a policy category and subscribe it to a database.

1. Using the Microsoft SQL Server Management Studio window, open the **Manage Policy Categories** dialog box.
2. Using the **Manage Policy Categories** dialog box, create a policy category named, **Mod05 Category**.
3. Open the **Categories** dialog box and subscribe the **Mod05 Category** of the **Mod05** database.
4. By using the Microsoft SQL Server Management Studio window, create a new policy for every table with the following details:

   Name: **Mod05 Policy**

   Check condition: **Mod05 Tables**

   Create a new condition for every database named Mod05 Database with the expression, **@Name = 'Mod05'**

   Evaluation Mode: **On change: prevent**

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

5. Enable the Mod05 Policy.

Task 4: Create a policy within the policy category.

1. By using the Microsoft SQL Server Management Studio window, open the **Create New Policy –** dialog box.
2. Using the **Create New Policy –** dialog box, create a new policy for every table with the following details:
   - Name: **Mod05 Policy**
   - Check condition: **Mod05 Tables**
   - Create a new condition for every database named Mod05 Database with the expression, **@Name = 'Mod05'**
   - Evaluation Mode: **On change: prevent**
3. Ensure that **Mod05 Policy** is enabled.
4. Using the Microsoft SQL Server Management Studio window, assign the **Mod05 Category** to **Mod05 Policy**.

Task 5: Test the policy.

1. Using the Microsoft SQL Server Management Studio window, apply **Mod05 Policy** to the **CallDetails** table.
2. Apply **Mod05 Policy** to the **tb1CallDetails** table.

Note: Browse to the **D:\Labfiles\Mod05** folder and select the **E2_T5.sql** file to execute the required T-SQL code.

Task 6: Delete the test database.

1. In the Query Editor pane of the Microsoft SQL Server Management Studio window, run a T-SQL script to delete the **Mod05** database.
2. Disable the **Mod05 Policy**.

Results: After completing this exercise, you should have created a test database, created a condition, created a policy category and subscribed it to a database, created a policy within the policy category, tested the policy, and deleted the test database.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Exercise 3: Governing Resources Consumption

Scenario

Recently, a server in your environment has run out of resources. This has caused other applications to stop responding to user requests. Even the administrator is unable to access the server. You realize that a query from an application is responsible for this problem. You want to take preventive action to ensure that such a problem does not occur again. You also want to ensure that an administrator can always access the system even when the system resources are running out.

The main tasks for this exercise are as follows:

1. Create resource pools.
2. Create a workload group.
3. Enable Resource Governor.
4. Create and register a classifier function.
5. Verify Resource Governor.
6. Delete Resource Governor.

Task 1: Create resource pools.

1. In the Query Editor pane of the Microsoft SQL Server Management Studio window, run a T-SQL script to create two resource pools as follows.

Resource pool name	Settings
rpProduction	Maximum CPU percentage: 100 Minimum CPU percentage: 50
rpDevelopment	Maximum CPU percentage: 50 Minimum CPU percentage: 0

Note: Browse to the **D:\Labfiles\Mod05** folder and select the **E3_T1.sql** file to execute the required T-SQL code.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Task 2: Create a workload group.

1. In the Query Editor pane of the Microsoft SQL ServerManagement Studio window, run a T-SQL script to create twoworkload groups named, **rgProduction** and **rgDevelopment** forthe **rpDevelopment** resource group, with **HIGH** and **LOW** importance respectively.

Note: Browse to the **D:\Labfiles\Mod05** folder and select the **E3_T2.sql** file to execute the required T-SQL code.

Task 3: Enable Resource Governor.

1. In the Query Editor pane of the Microsoft SQL Server Management Studio window, run a T-SQL script to enable Resource Governor.

Task 4: Create and register a classifier function.

1. In the Query Editor pane of the Microsoft SQL Server Management Studio window, run a T-SQL script to create a classifier function, **fnGroupClassifier**.
2. Register the **fnGroupClassifier** classifier function.

Note: Browse to the **D:\Labfiles\Mod05** folder and select the **E3_T4.sql** file to execute the required T-SQL code.

Task 5: Verify Resource Governor.

1. In the Query Editor pane of the Microsoft SQL Server Management Studio window, run a T-SQL script to select the resources pool dynamic management view to verify the Resource Governor setup.

Note: Browse to the **D:\Labfiles\Mod05** folder and select the **E3_T5.sql** file to execute the required T-SQL code.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Task 6: Delete Resource Governor.

1. In the Query Editor pane of the Microsoft SQL ServerManagement Studio window, run a T-SQL script to delete theworkload groups: **rgProduction** and **rgDevelopment**.
2. Start the SQL prompt by using Windows Authentication and retrieve a list of all database tables.

Note: Browse to the **D:\Labfiles\Mod05** folder and select the **E3_T6.sql** file to execute the required T-SQL code.

Results: After completing this exercise, you should have created a classifier function, created a workload group, enabled Resource Governor, registered a classifier function, verified Resource Governor, and deleted Resource Governor.

Exercise 4: Implementing DDL Triggers

Scenario

Business requirements require you to have audit log entries on all database schema changes. You have decided to use a DDL trigger to create audit entries in a database table.

The main tasks for this exercise are as follows:

1. Create a test database.
2. Create a log table.
3. Create and test a DDL trigger.
4. Delete the test database.

Task 1: Create a test database.

1. In the Query Editor pane of the Microsoft SQL Server Management Studio window, run a T-SQL script to create a test database named, **Mod05**, with the following structure.

File name	Settings

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

File name	Settings
Mod05_Primary.mdf	Current size: 10 Maximum size: 50 Filegrowth: 5
Mod05_Log.ldf	Current size: 10 Maximum size: 50 Filegrowth: 5

Note: Browse to the **D:\Labfiles\Mod05** folder and select the **E4_T1.sql** file to execute the required T-SQL code.

Task 2: Create a log table.

1. In the Query Editor pane of the Microsoft SQL Server Management Studio window, run a T-SQL script to create a log table named, **DDL_Log** with the structure:
   - PostTime: **datetime**
   - DB_User: **nvarchar (100)**
   - Event: **nvarchar (100)**
   - TSQL: **nvarchar (2000)**

Note: Browse to the **D:\Labfiles\Mod05** folder and select the **E4_T2.sql** file to execute the required T-SQL code.

Task 3: Create and test a DDL trigger.

1. In the Query Editor pane of the Microsoft SQL Server Management Studio window, run a T-SQL script to create a DDL trigger named **log** for the database event **DDL_DATABASE_LEVEL_EVENTS**. This trigger will extract event type and the T-SQL statement from the eventdata() function. This data should be saved to the **DDL_Log** table along with the user name and the current time.
2. In the Query Editor pane, run a T-SQL script to create a test table, and then drop the table to verify the setup of the DDL trigger.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

3. Display the records in the **DDL_Log** table.

Note: Browse to the **D:\Labfiles\Mod05** folder and select the **E4_T3.sql** file to execute the required T-SQL code.

Task 4: Delete the test database.

1. In the Query Editor pane of the Microsoft SQL Server Management Studio window, run a T-SQL script to delete the **Mod05** database.

Results: After completing this exercise, you should have created a test database, created a log table, created a DDL trigger, tested the DDL trigger, and deleted the test database.

Exercise 5: Implementing CDC

Scenario

The corporate policy requires logging of data changes in database tables for sensitive information. To reduce the maintenance and implementation effort, you decide to use CDC to create logs for data changed on designated tables.

The main tasks for this exercise are as follows:

1. Create a test database.
2. Create a test table.
3. Enable CDC for the test table.
4. Verify the CDC configuration.
5. Disable the CDC configuration.

Task 1: Create a test database.

1. In the Query Editor pane of the Microsoft SQL Server Management Studio window, run a T-SQL script to create a test database named, **Mod05**, with the following structure.

File name	Settings
Mod05_Primary.mdf	Current size: 10 Maximum size: 50

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

File name	Settings
	Filegrowth: 5
Mod05_Log.ldf	Current size: 10 Maximum size: 50 Filegrowth: 5

Note: Browse to the **D:\Labfiles\Mod05** folder and select the **E5_T1.sql** file to execute the required T-SQL code.

Task 2: Create a test table.

1. In the Query Editor pane of the Microsoft SQL Server Management Studio window, run a T-SQL script to create a test table named, **TestTable** in the **Mod05** table, with the following structure:
   - id: **int**
   - content: **varchar (200)**
   - create_date: **datetime**

Note: Browse to the **D:\Labfiles\Mod05** folder and select the **E5_T2.sql** file to execute the required T-SQL code.

Task 3: Enable CDC for the test table.

1. In the Query Editor pane of the Microsoft SQL Server Management Studio window, run a T-SQL script to enable CDC on the **Mod05** database and **TestTable**.

Task 4: Verify the CDC configuration.

1. In the Query Editor pane of the Microsoft SQL Server Management Studio window, run a T-SQL script to insert two records, update a record, and delete a record from the **TestTable** table.
2. Verify the CDC configuration by displaying all the actions performed on the **dbo.TestTable** table.
3. View the package.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Note: Browse to the **D:\Labfiles\Mod05** folder and select the **E5_T4.sql** file to execute the required T-SQL code.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Task 5: Disable the CDC configuration.

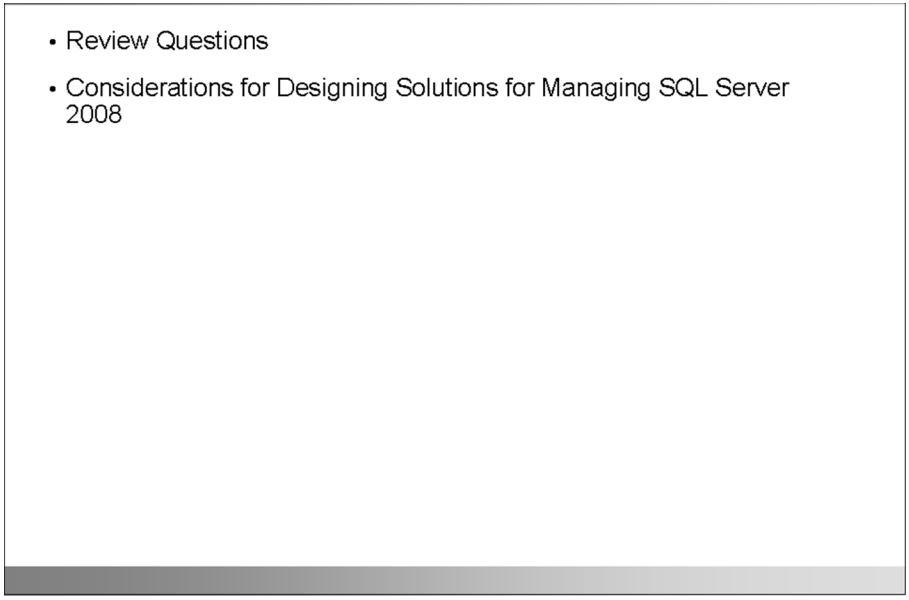
1. In the Query Editor pane of the Microsoft SQL Server Management Studio window, run a T-SQL script to disable the CDC configuration of the **Mod05** database.

Results: After completing this exercise, you should have created a test database, created a test table, enabled CDC for the test table, verified the CDC configuration, and disabled the CDC configuration.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Module Review and Takeaways



Review Questions

1. Your application development policy defines that all database objects should be prefixed with a string used to represent their object type. For example, the stored procedure should be prefixed with usp_. How can you find out if there are any objects that are not following this policy?
2. In your organization, you have to configure a policy to enforce database naming standard. After you configure the policy, you realize that are facing issues in using some SQL Server 2008 features, such as replication. What would be the possible cause of the issue?
3. Your database is always running out of resources when a user runs a report. You want to prevent this shortage of resources because it may affect another application, which shares the same database server. What should you do?
4. Your database is running out of resources now rejecting user connections. You also have trouble connecting to the database by

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

using SSMS. You realize that to overcome these problems, you need to change the configuration of Resource Governor. You try to reduce the number of resources assigned to the query but you a fail to do it. What should you do?

5. You are planning to track access to a database in SQL Server. Which feature of SQL Server should you use?
6. You are creating a custom DDL trigger that provides information about the user who creates additional objects on a database. Which function should you use?

Considerations for Designing Solutions for Managing SQL Server 2008

While designing solutions for managing SQL Server 2008, consider the following:

- Configure a Central Management Server to support centralized administration and management of all SQL Servers in your environment.
- Configure a policy to enforce corporate standards on database configuration and database schema.
- Exclude system databases and objects from policy evaluation.
- Configure a policy to execute at real-time to prevent a user from applying configuration that will violate policy standards.
- Configure a policy for scheduled execution for identifying configurations, which do not comply with the policy standards.
- Configure Resource Governor to control the allocation of system resources assigned to an application or a user.
- Consider using system functions, such as APP_NAME(), to provide additional information to correctly identify the application.
- Use SQL Server Audit to capture changes at server and database schema.
- Use CDC to capture changes at the data-level.
- Use CDC instead of triggers to capture changes in a database table.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Module 6

Automating the Database Management Strategy for SQL Server 2008

Contents:

Lesson 1: Managing SQL Server Agent 6-3
Lesson 2: Automating Administrative Tasks in SQL Server 6-29
Lesson 3: Monitoring Events in SQL Server 6-51
Lab 6: Automating the Database Management Strategy for SQL Server 2008 6-62
Module Review and Takeaways 6-72

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Module Overview

- Managing SQL Server Agent
- Automating Administrative Tasks in SQL Server
- Monitoring Events in SQL Server

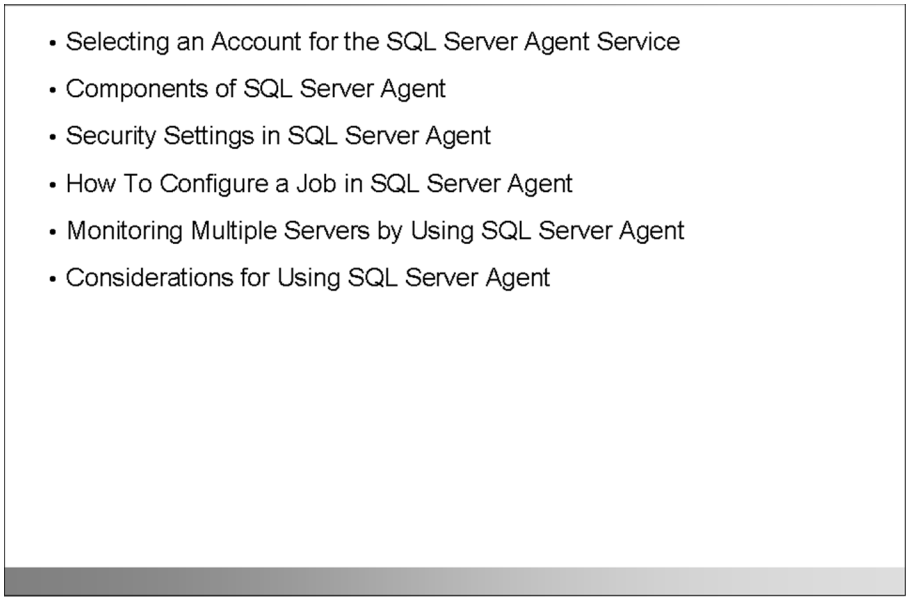
Data management is the process to ensure that diverse data is collected, integrated, labeled, stored, and then retrieved when needed. Business enterprises are concerned with data management applications. SQL provides an array of new features that manages data and protects the server from potential threats. Microsoft ™ SQL Server™ 2008 offers flexibility, improved performance, and scalability of data.

In this module, you will learn about SQL Server Agent and managing SQL Server by using SQL Server Agent. You will also learn to monitor events.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Lesson 1
Managing SQL Server Agent



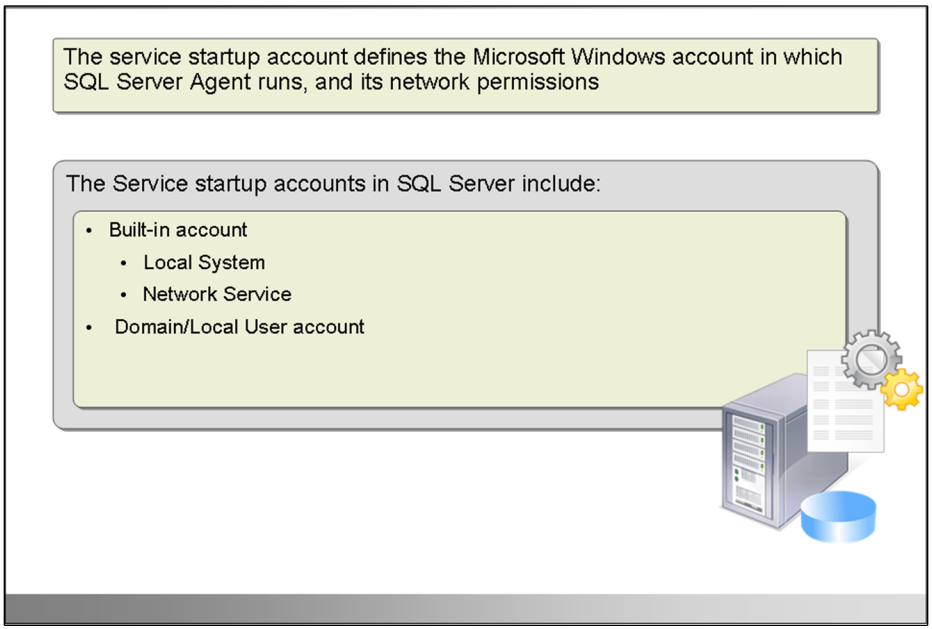
SQL Server Agent is a Microsoft Windows® service that runs scheduled administrative tasks. SQL Server Agent depends on various SQL Server features, such as Change Data Capture (CDC) to provide the infrastructure for performing various tasks. To administer SQL Server, you must ensure that SQL Server Agent is configured properly.

In this lesson, you will learn about the configuration settings and security settings of SQL Server Agent. In addition, you will learn to configure a job in SQL Server Agent. Furthermore, you will learn to monitor multiple servers by using SQL Server Agent and considerations for using SQL Server Agent.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Selecting an Account for the SQL Server Agent Service



Key Points

SQL Server Agent is a Microsoft Windows service that allows you to automate some administrative tasks. SQL Server Agent runs jobs, monitors SQL Server, and processes alerts. The SQL Server Agent service must be running before local or multiserver administrative jobs can run automatically.

The service startup account defines the Microsoft Windows account in which SQL Server Agent runs, and its network permissions. SQL Server Agent runs as a specified user account. For better security, SQL Server Agent should use separate user accounts as compared to SQL Server. For compatibility with earlier versions of SQL Server, SQL Server Agent can also run as the Local System account.

You can select an account for the SQL Server Agent service by using SQL Server Configuration Manager. The two types of accounts are as follows:

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

- **Built-in account**. You can choose from the following list of built-in Windows service accounts.
- **Local System account**. The name of this account is NT AUTHORITY\System. It is a powerful account that has unrestricted access to all local system resources. It is a member of the Windows Administrators group on the local computer, and is, therefore, a member of the SQL Server sysadmin fixed server role.

  Note: The Local System account option is provided for backward compatibility only. The Local System account has permissions that SQL Server Agent does not require. Avoid running SQL Server Agent as the Local System account.

- **Network Service account**. The name of this account is NT AUTHORITY\NetworkService. It is available in Microsoft Windows XP and Microsoft Windows Server® 2003. All services that run under the Network Service account are authenticated to network resources as the local computer.

Note: Because multiple services can use the Network Service account, it is difficult to control the services that have access to network resources, including SQL Server databases. You should not use the Network Service account for the SQL Server Agent service.

Important: Do not select the Local Service account. The SQL Server Agent service cannot be run under this account. The name of this account is NT AUTHORITY\LocalService, and it accesses network resources as a null session with no credentials. It is available in Microsoft Windows XP and Microsoft Windows Server 2003.

- **Domain/Local User Account**. This account helps you specify the Windows domain account in which the SQL Server Agent service runs. You must choose a Windows user account that is not a member of the Windows Administrators group as the local user account. However, there are limitations for using multiserver administration when the SQL Server Agent service account is not a member of the local Administrators group.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Question: Which service account should you choose for SQL Server Agent for most cases?

Source:

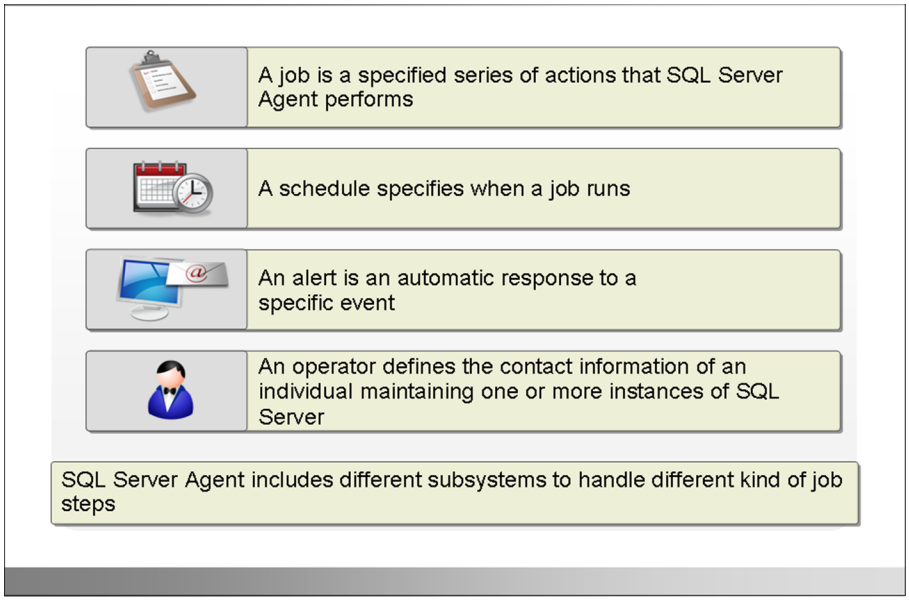
- http://msdn.microsoft.com/en-us/library/aa213810(SQL.80).aspx
- http://msdn.microsoft.com/en-us/library/ms186264.aspx
- http://technet.microsoft.com/en-us/library/ms191543(SQL.105).aspx

For more information, see Service Account Types Supported for SQL Server Agent.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Components of SQL Server Agent



Key Points

You can configure the SQL Server Agent service to start automatically when the operating system starts. You can also start it manually when you need to complete jobs. SQL Server Agent uses components such as jobs, schedules, alerts, and operators to define tasks to be performed; when to perform the tasks; and how to report the success or failure of the tasks.

Jobs

A job is a specified series of actions that SQL Server Agent performs. Use jobs to define an administrative task that can be run one or more times, and monitored for success or failure. A job can run on one local server or on multiple remote servers. You can run jobs in several ways, such as:

- According to one or more schedules.
- In response to one or more alerts.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

By running the **sp_start_job** stored procedure.

Each action in a job is a job step. For example, a job step might consist of running a Transact-SQL (T-SQL) statement, running an SSIS package, or issuing a command to an Analysis Services server. A job step is a part of the job and is managed as part of the job. Each job step runs in a specific security context. For job steps that use T-SQL, use the EXECUTE AS statement to set the security context for the job step. For other types of job steps, use a proxy account to set the security context for the job step.

Schedules

A schedule specifies when a job runs. Multiple jobs can run on the same schedule, and multiple schedules can apply to the same job. A schedule can define the following conditions for the time a job runs:

- When SQL Server Agent starts
- When CPU utilization of the computer is at a level defined as idle
- On a single occasion, such as at a specific date and time
- On a recurring schedule

Alerts

An alert is an automatic response to a specific event. For example, an event can be a job that starts or system resources that reach a specific threshold. An alert can perform actions such as notifying one or more operators, or running a job. You can define the conditions under which an alert occurs. An alert can respond to one of the following conditions:

- SQL Server events
- SQL Server performance conditions
- Microsoft Windows Management Instrumentation (WMI) events on the computer where SQL Server Agent is running

Operator

An operator defines contact information for an administrator responsible for the maintenance of one or more instances of SQL Server. An operator does not contain security information and does not define a security principal. SQL Server 2008 can notify operators of alerts through e-mail, pager, or by using the net send command.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

If you are using SQL Server 2008 on Windows Server 2008, you cannot use net send because this feature is removed from Windows Server 2008.

SQL Server Agent Subsystems

SQL Server Agent includes different subsystems to handle different kind of job steps. You would have to specify what kind of subsystem to use when creating a SQL Server Agent job, which SQL Server would use to execute the specified job steps. The following table lists the subsystems defined by SQL Server.

Subsystem name	Description
Microsoft ActiveX® Script	Microsoft ActiveX Script runs an ActiveX scripting job step. The ActiveX Scripting subsystem will be removed from SQL Server Agent in future versions of SQL Server. Avoid using this feature in new development work, and plan to modify applications that currently use this feature.
Operating System (CmdExec)	The operating system runs an executable program.
PowerShell	The PowerShell runs a PowerShell scripting job step.
Replication Distributor	The Replication Distributor runs a job step that activates the replication Distribution Agent.
Replication Merge	Replication merge runs a job step that activates the replication Merge Agent.
Replication Queue Reader	The Replication Queue Reader runs a job step that activates the replication Queue Reader Agent.
Replication Snapshot	The Replication snapshot runs a job step that activates the replication Snapshot Agent.
Replication Transaction Log Reader	The replication transaction log reader runs a job step that activates the replication Log Reader Agent.
Analysis Services Command	This subsystem runs an Analysis Services command.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Subsystem name	Description
Analysis Services Query	This subsystem runs an Analysis Services query.
SSIS package execution	The SSIS package execution subsystem runs an SSIS package.

Question: Which service account should you choose for SQL Server Agent for most cases?

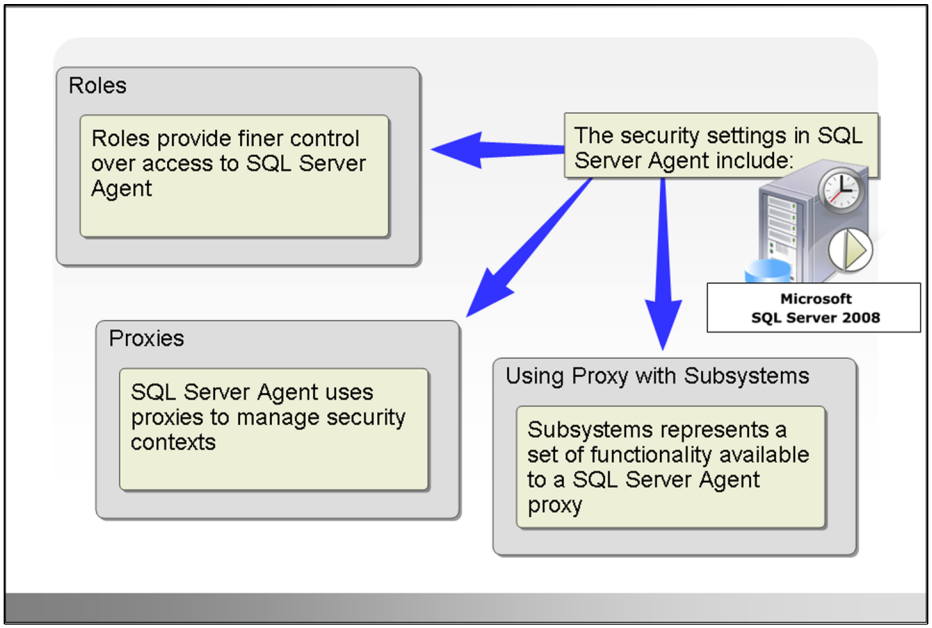
Source:

- http://msdn.microsoft.com/en-us/library/ms190695(SQL.90).aspx
- http://msdn.microsoft.com/en-us/library/aa213814(SQL.80).aspx
- http://msdn.microsoft.com/en-us/library/ms175540(SQL.90).aspx
- http://msdn.microsoft.com/en-us/library/ms187100(SQL.90).aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Security Settings in SQL Server Agent



Key Points

You can use database roles in msdb database to control who have access to configuration and information, such as job history related to SQL Server Agent. SQL Server Agent uses proxies to help you define a job that is executed under another user credential. SQL Server Agent can allow you to control which subsystem the proxy is allowed to use.

Roles

SQL Server 2008 introduced the SQLAgentUserRole, SQLAgentReaderRole, and SQLAgentOperatorRole database roles in msdb database, which provides you with finer control over access to SQL Server Agent. Members of these roles have fixed database roles in msdb, and members of the sysadmin fixed server role have access to SQL Server Agent. Users who are not members of any of these roles cannot use SQL Server Agent.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

The following table describes SQLAgentUserRole, SQLAgentReaderRole, and SQLAgentOperatorRole roles.

Role	Description
SQLAgentUserRole	SQLAgentUserRole is the least privileged of the SQL Server Agent fixed database roles. It has permissions on only operators, local jobs, and job schedules. Members of SQLAgentUserRole have permissions on only local jobs and job schedules that they own. They cannot use multiserver jobs or the master and target server jobs, and they cannot change job ownership to gain access to jobs that they do not already own. SQLAgentUserRole members can view a list of available proxies only in the Job Step Properties dialog box of SQL Server Management Studio. Only the Jobs node in SQL Server Management Studio Object Explorer is visible to members of SQLAgentUserRole.
SQLAgentReaderRole	SQLAgentReaderRole includes all the SQLAgentUserRole permissions as well as permissions to view the list of available multiserver jobs, their properties, and their history. Members of this role can also view the list of all available jobs and job schedules, and their properties, not just those jobs and job schedules that they own. SQLAgentReaderRole members cannot change job ownership to gain access to jobs that they do not already own. Only the Jobs node in SQL Server Management Studio Object Explorer is visible to members of the SQLAgentReaderRole.
SQLAgentOperatorRole	SQLAgentOperatorRole is the most privileged of the SQL Server Agent fixed database roles. It includes all the permissions of SQLAgentUserRole and SQLAgentReaderRole. Members of this role can also view properties for operators and proxies, and enumerate

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Role	Description
	available proxies and alerts on the server.

Proxies

SQL Server Agent uses proxies to manage security contexts. SQL Server Agent can use a proxy in more than one job step. Members of the sysadmin fixed server role can create proxies.

Each proxy corresponds to a security credential. Each proxy can be associated with a set of subsystems and a set of logons. The proxy is used by the SQL Server Agent only for job steps that use a subsystem associated with the proxy. To create a job step that uses a specific proxy, the job owner must either use a logon associated with that proxy or be a member of a role with unrestricted access to proxies. Members of the sysadmin fixed server role have unrestricted access to proxies. Members of SQLAgentUserRole, SQLAgentReaderRole, or SQLAgentOperatorRole can only use proxies to which they have been granted specific access. Each user who is a member of any of these SQL Server Agent fixed database roles must be granted access to specific proxies so that the user can create job steps that use those proxies.

Using Proxy with Subsystems

A subsystem is a predefined object that represents a set of functionality available to a SQL Server Agent proxy. Each proxy has access to one or more subsystems. Subsystems provide security because they delimit access to the functionality that is available to a proxy. Each job step runs in the context of a proxy, except for T-SQL job steps. T-SQL job steps use the EXECUTE AS command to set the security context.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Question: You have a SQL Server Agent job that requires special permission for a step to complete. How can you configure SQL Server Agent to use different credentials to execute the job without changing the service account?

Source:

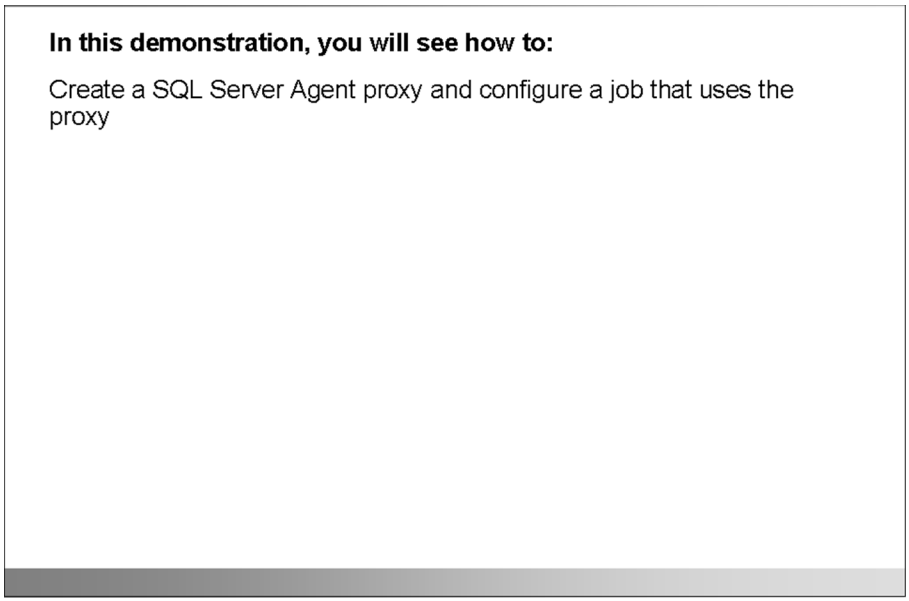
- http://msdn.microsoft.com/en-us/library/ms188283(SQL.90).aspx
- http://msdn.microsoft.com/en-us/library/ms190978(SQL.90).aspx
- http://msdn.microsoft.com/en-us/library/ms187100(SQL.90).aspx

For more information, see SQL Server Agent Fixed Database Roles.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Demonstration: How To Configure a Job in SQL Server Agent



Key Points

The steps to create a SQL Server Agent proxy and configure a job that uses the proxy are:

1. On the **Start** menu, point to **All Programs**, click **Microsoft SQL Server 2008**, and then click **SQL Server Management Studio**. The Microsoft SQL Server Management Studio window appears.
2. In the Microsoft SQL Server Management Studio window, by using the **Connect to Server** dialog box, connect to the default instance of SQL Server Database Engine.
3. On the **Start** menu, point to **All Programs**, point to **Accessories**, and then click **Windows Explorer**.
4. In Windows Explorer, browse to the **D:\Demofiles\Mod06** folder, and then double-click the **M6_Demo1.sql** file.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

5. In the Microsoft SQL Server Management Studio window, in the Query Editor pane, select the **Create an account credential** code, and then click **Execute**.
6. In the Microsoft SQL Server Management Studio window, in the Query Editor pane, select the **Create an account proxy** code, and then click **Execute**.
7. In the Microsoft SQL Server Management Studio window, in the Query Editor pane, select the **Create a job** code, and then click **Execute**.
8. In the Microsoft SQL Server Management Studio window, in the Query Editor pane, select the **Create one or more job steps** code, and then click **Execute**.
9. In the Microsoft SQL Server Management Studio window, in the Query Editor pane, select the **Create a schedule** code, and then click **Execute**.
10. In the Microsoft SQL Server Management Studio window, in the Query Editor pane, select the **Attach a schedule to the job** code, and then click **Execute**.
11. Close all open windows.

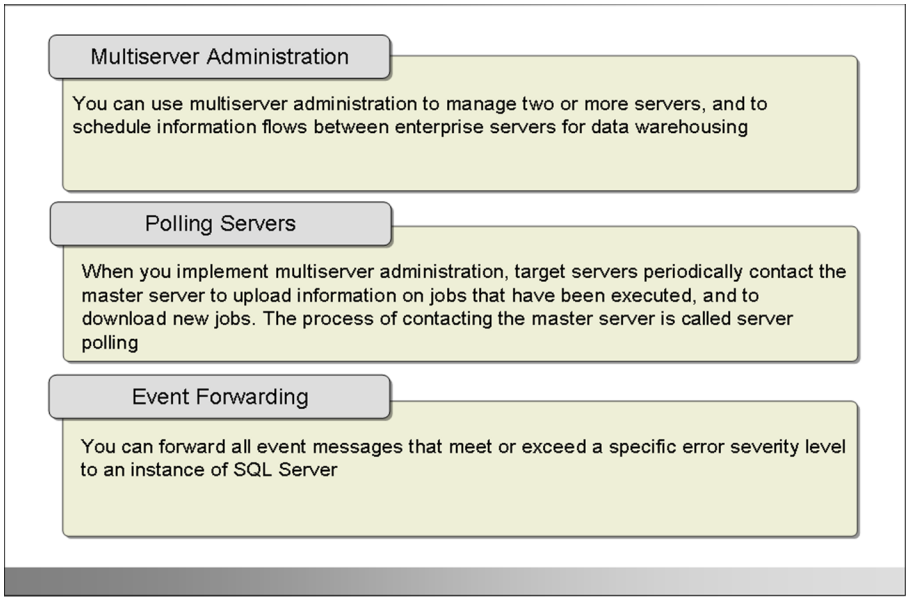
Question: What should you do to schedule an administrative job every night?

Question: Do you need to have the SQL Server Agent service set to automatic when you have created the job?

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Monitoring Multiple Servers by Using SQL Server Agent



Key Points

You can monitor multiple servers by using the multiserver administration, polling servers, and event forwarding features of SQL Server Agent.

Multiserver Administration

Automating administration across multiple instances of SQL Server is called multiserver administration. You can use multiserver administration to manage two or more servers, and to schedule information flows between enterprise servers for data warehousing. To take advantage of multiserver administration, you must have at least one master server and one target server. A master server distributes jobs to, and receives events from, target servers. A master server also stores the central copy of job definitions for jobs that are run on target servers. Target servers connect periodically to the master server to update their schedule of jobs. If a new job exists on the master server, the target server downloads the job.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

After the target server completes the job, it reconnects to the master server and reports the status of the job.

If you administer departmental servers, you can define the following:

- One backup job with job steps
- Operators to notify in case of backup failure
- An execution schedule for the backup job

You can write this backup job a single time on the master server and then enlist each server as a target server. All servers run the same backup job from the time of their enlistment; so you need to define the job only once.

Polling Servers

When you implement multiserver administration, target servers periodically contact the master server to upload information on jobs that have been executed, and to download new jobs. The process of contacting the master server is called server polling. Server polling occurs at regular polling intervals. The polling interval controls how frequently the target server connects to the master server to download instructions and upload the results of job execution. By default, the polling interval is one minute. When a target server polls the master server, it reads the operations assigned to the target server from the sysdownloadlist table in the msdb database. These operations control multiserver jobs and configuration of a target server. Examples of operations include deleting, inserting, and starting a job; and updating the polling interval of a target server. Operations can be posted to the sysdownloadlist table either:

- Explicitly by using the sp_post_msx_operation stored procedure.
- Implicitly by using other job stored procedures.
- You do not have to post operations explicitly if you use:
- Microsoft SQL Server Management Studio (SSMS) to control multiserver jobs.
- Job stored procedures that do not modify job schedules or job steps.

If you use job stored procedures to modify multiserver job schedules or job steps, or SQL Distributed Management Objects (SQL-DMO) to control multiserver jobs, you must keep the target servers

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

synchronized with the current job definition. You can maintain synchronicity by configuring the SQL Server for taking part in multiserver job.

To configure the SQL Server as a master server, perform the following steps:

1. In the Microsoft SQL Server Management Studio window, in the Object Explorer, connect to an instance of the Microsoft SQL Server Database Engine, and then expand that instance.
2. In the Object Explorer pane, right-click **SQL Server Agent**, point to **Multi Server Administration**, and then click **Make this a Master**. The Master Server Wizard – window appears.
3. On the **Welcome to the Master Server Wizard** page, click **Next**.
4. On the **Master Server Operator** page, click **Next**.
5. On the **Target Servers** page, click Add Connection, the Connect to Server dialog box appears.
6. In the **Connect to Server** dialog box, connect to the server to register as a master server, and then click Connect.
7. On the **Target Servers** page, click **Next**. The Checking Server Compatibility window appears.
8. On the **Checking Server Compatibility** page, click **Close**.
9. On the **Master Server Login Credentials** page, click **Next**.
10. In the **Complete the Wizard** page, click **Finish**.

Event Forwarding

You can forward all event messages that meet or exceed a specific error severity level to an instance of SQL Server. This is called event forwarding. The forwarding server is a dedicated server that can also be a master server. You can use event forwarding to centralize alert management for a group of servers, thereby reducing the workload on heavily used servers. When one server receives events for a group of other servers, the server that receives events is called an alerts management server. In a multiserver environment, you designate the master server as the alerts management server.

Consider the following points when creating a multiserver environment:

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

- If you want to enlist a target server with another master server, you must defect a target server from one master server before enlisting with the other one. This is because each target server reports to only one master server.
- If you are changing the name of a target server, you must defect it before changing the name and re-enlist it after the change.
- If you want to dismantle a multiserver configuration, you must defect all target servers from the master server.

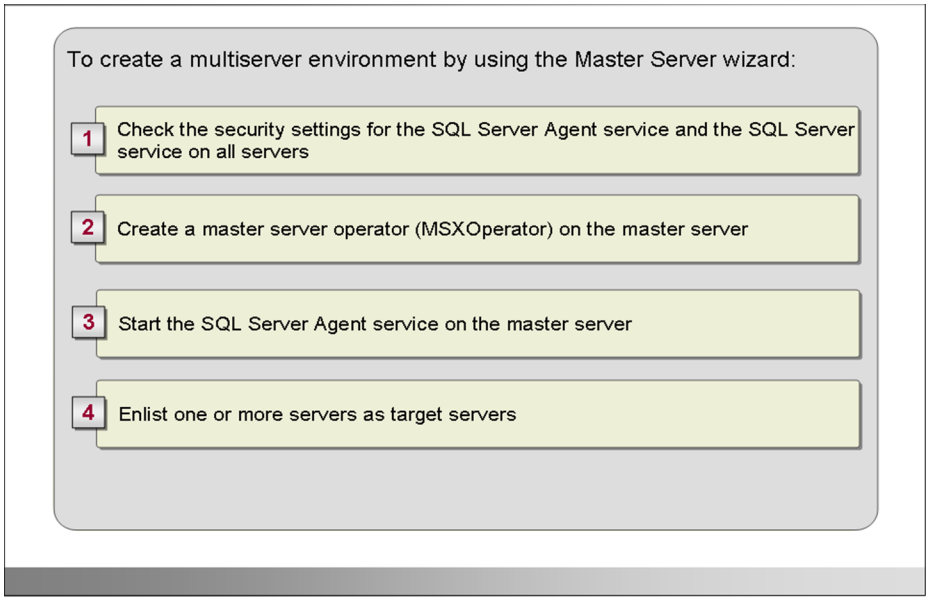
Source:

- http://msdn.microsoft.com/en-us/library/aa213819(SQL.80).aspx
- http://msdn.microsoft.com/en-us/library/ms189633(SQL.90).aspx
- http://msdn.microsoft.com/en-us/library/ms189316(SQL.90).aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Creating a Multiserver Environment



Key Points

In a multiserver administration, you need to set up a master server and one or more target servers. Jobs that are processed on all the target servers are first defined on the master server, and then downloaded to the target servers.

You can create a multiserver environment by using the Master Server Wizard. Full Secure Sockets Layer (SSL) encryption and certificate validation are enabled for connections between master servers and target servers by default.

To create a multiserver environment by using the Master Server wizard, you need to:

1. Check the security settings for the SQL Server Agent service and the SQL Server service on all servers intended to become target servers.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Note: It is recommended that the SQL Server Agent service and the SQL Server service should run on Microsoft Windows domain accounts.

2. Create a master server operator (MSXOperator) on the master server.
   Note: MSXOperator is the only operator that can receive notifications for multiserver jobs.

3. Start the SQL Server Agent service on the master server.

4. Enlist one or more servers as target servers. If you have a large number of target servers, avoid defining your master server on a production server. Otherwise, target server traffic can slow down the performance of the production server. If you can forward events to a dedicated master server, you can centralize administration on one server.

Note: To use multiserver job processing, the SQL Server Agent service account should be a member of the **msdb** database role **TargetServersRole** on the master server. The Master Server Wizard automatically adds the service account to this role as a part of the enlistment process.

Considerations for creating multiserver environments

When creating a multiserver environment, you need to consider the following issues:

- Each target server should report to only one master server. You must defect a target server from one master server before you can enlist it into a different one.

- When changing the name of a target server, you must defect it before changing the name and re-enlist it after the change.

- If you want to dismantle a multiserver configuration, you must defect all the target servers from the master server.

The following are the common tasks performed in creating a multiserver environment:

- **Make changes to multiserver job definitions**. To make changes to multiserver job definitions outside Microsoft SQL

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Server Management Studio, you should post the changes to the download list so that target servers can download the updated job again. To ensure that target servers have current job definitions, post an INSERT instruction after you update the multiserver job, by using the following code.

```
EXECUTE sp_post_msx_operation 'INSERT', 'JOB', '<job id>'
```

- **Set encryption options on target servers**. If you cannot use SSL to encrypt the channel between the master servers and target servers, you can encrypt the channel by configuring the target server to use the level of security required. To configure the appropriate level of security required for a specific master server or target server communication channel, you need to set the SQL Server Agent registry subkey, \HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Microsoft SQL Server\<instance_name>\SQLServerAgent\MsxEncryptChannelOptions(REG_DWORD) on the target server to one of the encryption value. The following table illustrates the encryption values with its description.

Value	Description
0	Disables encryption between the target server and the master server.
1	Enables encryption only between the target server and the master server, but no certificate validation is required
2	Enables full SSL encryption and certificate validation between this target server and the master server. This setting is the default. Unless you have specific reason to choose a different value, we recommend not changing it.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

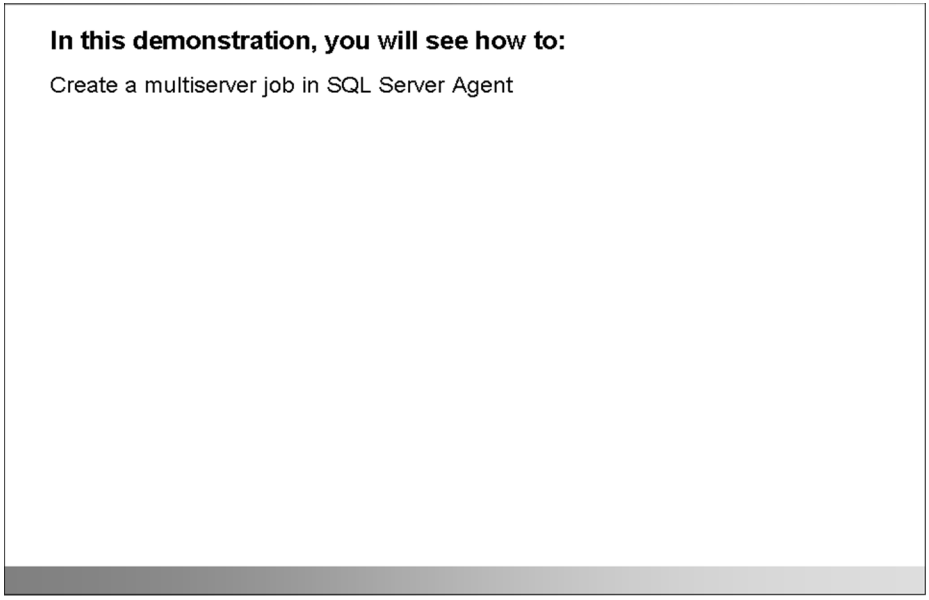
New Horizons Computer Learning Centers, Inc.

Note: The value of <instance_name> is MSSQL.n. For example, MSSQL.1 or MSSQL.3.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Demonstration: How To Create a Multiserver Job in SQL Server Agent



Key Points

The steps to create a multiserver job in SQL Server Agent are:

1. On the **Start** menu of NYC-SQL1, click **Run**.
2. In the **Open** box of the **Run** dialog box, type **regedit**, and then click **OK**.
3. In the Registry Editor window, navigate to **Computer\HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Microsoft SQL Server\MSSQL 10.MSSQLSERVER\SQLServerAgent**.
4. In the result pane, under **Name** list, double-click **MsxEncryptChannelOptions**.
5. In the **Value data** box of the **Edit DWORD(32-bit) Value** dialog box, type **0**, and then click **OK**.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

6. In the Registry Editor window, navigate to **Computer\HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Microsoft SQL Server\MSSQL 10.DEVELOPMENT\SQLServerAgent**.
7. In the result pane, under **Name** list, double-click **MsxEncryptChannelOptions**.
8. In the **Value data** box of the **Edit DWORD(32-bit) Value** dialog box, type **0**, and then click **OK**.
9. In the Registry Editor window, navigate to **Computer\HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Microsoft SQL Server\MSSQL 10.INSTANCE3\SQLServerAgent**.
10. In the result pane, under **Name** list, double-click **MsxEncryptChannelOptions**.
11. In the **Value data** box of the **Edit DWORD(32-bit) Value** dialog box, type **0**, and then click **OK**.
12. On the **Start** menu, point to **All Programs**, point to **Microsoft SQL Server 2008**, and then click **SQL Server Management Studio**. The Microsoft SQL Server Management Studio window appears.
13. In the Microsoft SQL Server Management Studio window, by using the **Connect to Server** dialog box, connect to the default instance of SQL Server Database Engine.
14. In the Microsoft SQL Server Management Studio window, in the Object Explorer pane, right-click **SQL Server Agent**, point to **Multi Server Administration**, and then click **Make this a Master**.
15. On the **Welcome to the Master Server Wizard** page of the Master Server Wizard – NYC-SQL1 wizard, click **Next**.
16. On the **Master Server Operator** page of the Master Server Wizard – NYC-SQL1 wizard, click **Next**.
17. On the **Target Servers** page, expand **Database Engine**, and then click **Add Connection**.
18. In the **Connect to Server** dialog box, in the **Server name** list, click **NYC-SQL1\INSTANCE3**, and then click **Connect**.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

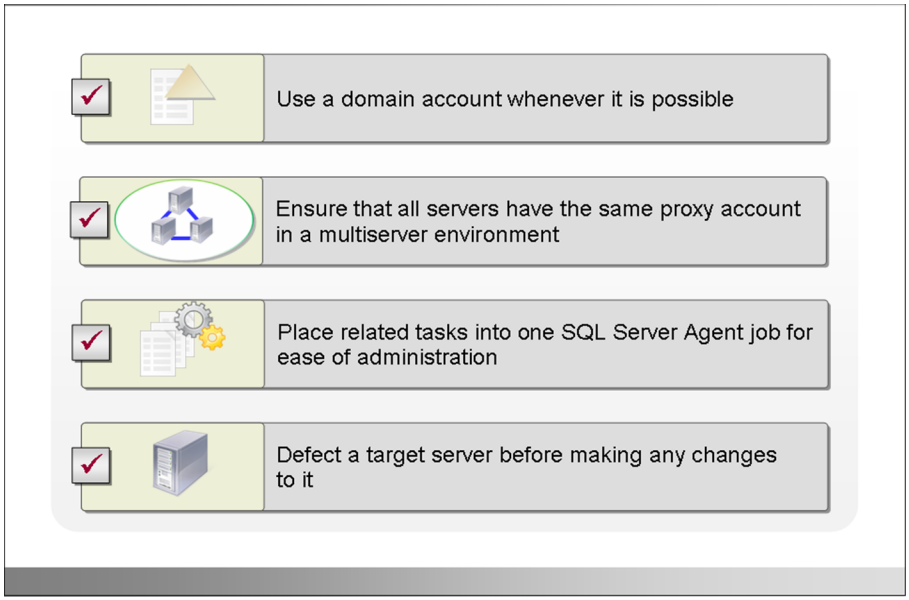
New Horizons Computer Learning Centers, Inc.

19. On the **Target Servers** page, click **Add Connection**.
20. In the **Connect to Server** dialog box, in the **Server name** list, click **NYC-SQL1\DEVELOPMENT**, and then click **Connect**.
21. On the **Target Servers** page, click **Next**.
22. In the **Checking Server Compatibility** dialog box, click **Close**.
23. On the **Master Server Login Credentials** page, click to clear the **Create a new login if necessary and assign it rights to the MSX** check box, and then click **Next**.
24. On the **Complete the Wizard** page, click **Finish**.
25. On the Enlist TSX Progress page, after the MSXOperator is created, click **Close**.
26. On the **Start** menu, point to **All Programs**, point to **Accessories**, and then click **Windows Explorer**.
27. In Windows Explorer, browse to the **D:\Demofiles\Mod06** folder, double-click the **M6_Demo4.sql** file.
28. In the Microsoft SQL Server Management Studio window, on the toolbar of the M6_Demo4.sql page, then click **Execute**.
29. Close all open windows.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Considerations for Using SQL Server Agent



Key Points

SQL Server Agent is allows you to automate administrative tasks or monitor the SQL Servers in your environment. Consider the following points when using SQL Server Agent:

- **Use a domain account wherever it is possible**. Define the domain account as the service account if access to resources outside the server hosting SQL Server is required. If access to such resources is not required, use a local user account on the server as the service account. If access is not required, use the designated user account as the service account. The Local System account is provided only for backward compatibility. It has various permissions that are applicable for the SQL Server Agent. Therefore, you should avoid running SQL Server Agent as the Local System account.
- **Ensure the all servers have the same proxy account in a multiserver environment**. The name of the proxy account on the target server should be the same as the master server proxy

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

account name under which the job step runs otherwise the job would fails to run.

- **Place related tasks into one SQL Server Agent job for ease of administration**. This helps you to define conditions and control the execution order of tasks. It also helps to avoid problems caused by some tasks overlapping with others. For example, you should perform a backup after a batch job is completed. If you start the backup before the batch job is completed, the data might be inconsistent.
- **Defect a target server before making any changes to it**. Consider defecting a target server from a master server before enlisting the target server into a different one. Failure to do so may cause the information to be sent to the wrong target server.

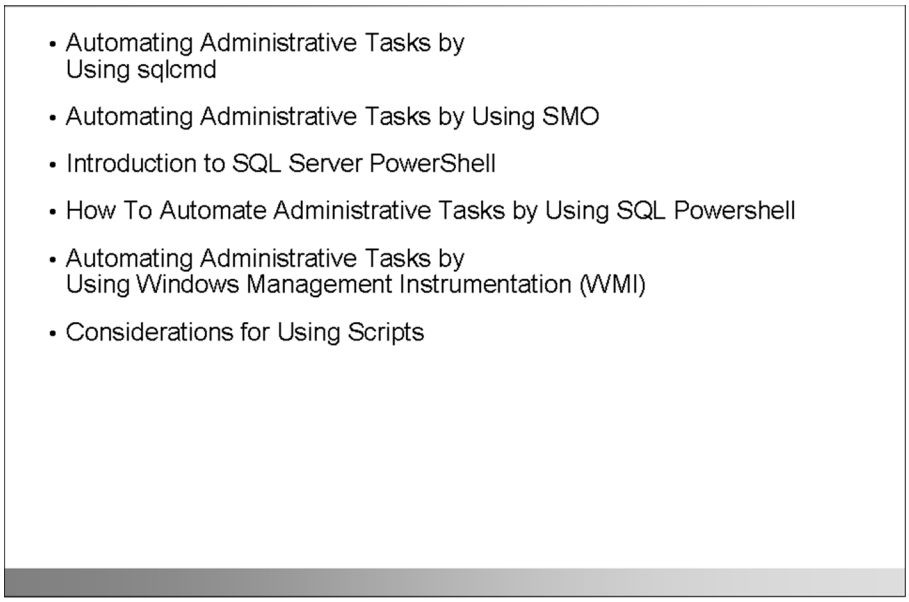
Source:

- http://msdn.microsoft.com/en-us/library/ms191543(SQL.90).aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Lesson 2
Automating Administrative Tasks in SQL Server



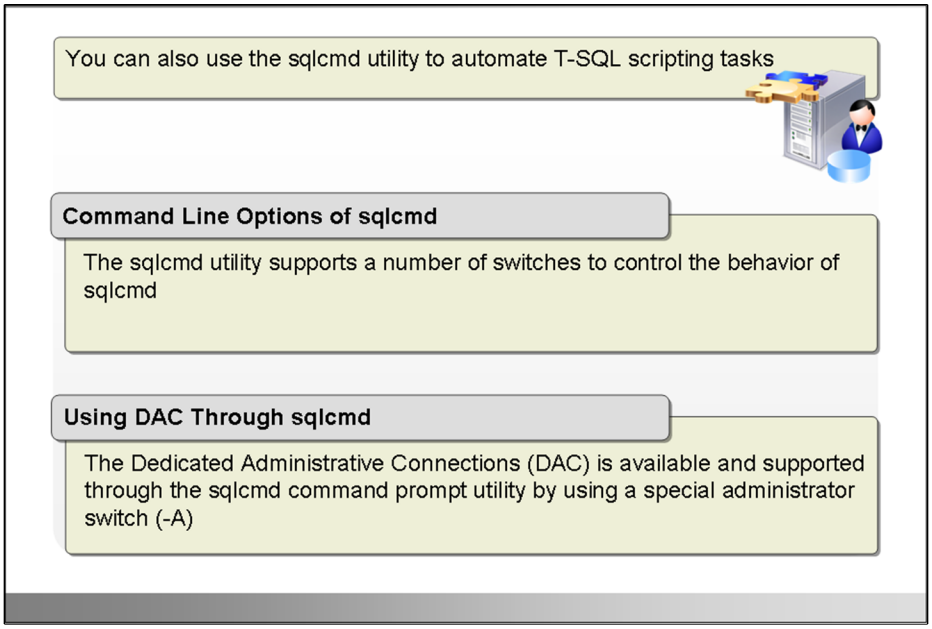
You can use SQL Server Agent to automate administrative tasks. To automate administrative tasks, you define predictable administrative tasks and then specify the conditions under which each task occurs. The tools used for automating the tasks are sqlcmd tool, SQL Server PowerShell, and WMI.

In this lesson, you will learn about automating administrative tasks by using the sqlcmd tool, Server Management Objects (SMO), SQL Server PowerShell, WMI. You will also learn about the considerations for using scripts for managing SQL Server.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Automating Administrative Tasks by Using sqlcmd



Key Points

The sqlcmd utility is a Microsoft Win32® command prompt utility for interactive execution of T-SQL statements and scripts. You can also use it to automate T-SQL scripting tasks. You must understand T-SQL to use sqlcmd interactively and to build script files.

You can use the sqlcmd utility in the following ways:

- Interactively entering T-SQL statements in a manner similar to working at the command prompt. The results are displayed at the command prompt. At the command prompt, type sqlcmd, followed by a list of options that you want.
- Submitting a sqlcmd job either by specifying a single T-SQL statement to run, or by pointing the utility to a text file that contains T-SQL statements to run. The output is usually directed to a text file, but it can also be displayed at the command prompt by using one of the following methods:
    - The SQLCMD mode in SSMS Query Editor

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

- SQL Server Management Objects (SMOs)
- SQL Server Agent CmdExec jobs

sqlcmd Command Line Options

The sqlcmd utility supports a number of command line options to control the behavior of sqlcmd or to specify the parameters of the connection used to connect to the SQL Server.

The following table lists the options used in sqlcmd.

Switch	Description
-E	This switch uses a trusted connection instead of using a user name and password to log on to SQL Server. By default, without -E specified, sqlcmd uses the trusted connection option. The -E option ignores possible user name and password environment variable settings such as SQLCMDPASSWORD. If the -E option is used together with the -U option or the -P option, an error message is generated.
-P password	This password is a user-specified password. Passwords are case sensitive. If the -U option is used and the -P option is not used, and the SQLCMDPASSWORD environment variable has not been set, sqlcmd prompts the user for a password. If the -P option is used at the end of the command prompt without a password sqlcmd uses the default password (NULL).
-U login_id	It is the user login ID. If neither the -U option nor the -P option is specified, sqlcmd tries to connect by using Microsoft Windows Authentication mode. Authentication is based on the Windows account of the user who is running sqlcmd.
-S server_name [\instance_name]	This switch specifies the instance of SQL Server to which to connect. It sets the sqlcmd scripting variable SQLCMDSERVER. Specify server_name to connect to the default instance of SQL Server on that server computer. Specify server_name [\instance_name] to connect to a named instance of SQL Server on that server computer. If no server computer is specified, sqlcmd connects to the default instance of SQL

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Switch	Description
	Server on the local computer. This option is required when you execute sqlcmd from a remote computer on the network.
-d db_name	This switch issues a USE db_name statement when you start sqlcmd. This option sets the sqlcmd scripting variable SQLCMDDBNAME. This specifies the initial database. The default is your login's default-database property. If the database does not exist, an error message is generated and sqlcmd exits.
-A dedicated administrator connection	This switch logs on to SQL Server with a Dedicated Administrator Connection (DAC). This kind of connection is used to troubleshoot a server. This will only work with server computers that support DAC. If DAC is not available, sqlcmd generates an error message and then exits.
-I input_file[,input_file2...]	This switch identifies the file that contains a batch of SQL statements or stored procedures. Multiple files may be specified that will be read and processed in order.
-q " cmdline query "	This switch executes a query when sqlcmd starts, but does not exit sqlcmd when the query has finished running. Multiple-semicolon-delimited queries can be executed.

Using Dedicated Administrative Connections (DAC) through sqlcmd

You can use the DAC of SQL Server when standard connections to the server are not possible. The DAC allows you to access SQL Server, run diagnostic queries, and troubleshoot problems even when SQL Server is not responding to standard connection requests.

The DAC supports encryption and other security features of SQL Server. The DAC only allows changing the user context to another admin user.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

SQL Server makes every attempt to make DAC connect successfully, but under extreme situations, it may not be successful.

By default, the connection is only allowed from a client running on the server. Network connections are not permitted unless they are configured by using the sp_configure stored procedure with the remote admin connections option.

Only members of the SQL Server sysadmin role can connect by using the DAC.

The DAC is available and supported through the sqlcmd command prompt utility by using a special administrator switch (-A). You can also connect prefixing admin: to the instance name in the format sqlcmd -Sadmin:<instance_name>. You can also initiate a DAC from a SQL Server Management Studio Query Editor by connecting to admin:<instance_name>.

For example, the following sqlcmd command shows how to connect to a server that has a blocking problem by using DAC.

```
C:\>sqlcmd -S ServerName -A

1> SELECT blocked FROM sys.dm_exec_requests WHERE blocked
<> 0;

2> GO

Result set:

spid blocked

------ -------

62 64

(1 rows affected)

sqlcmd command to end the blocking process:

1> KILL 64;

2> GO
```

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Question: Why should you use DAC?

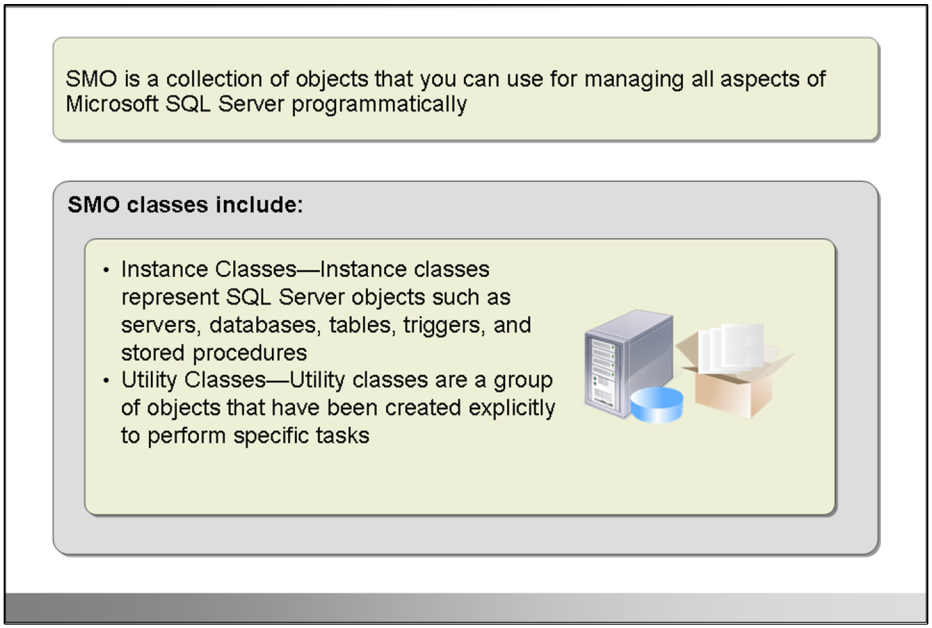
Source:

- http://msdn.microsoft.com/en-us/library/ms180944.aspx
- http://msdn.microsoft.com/en-us/library/bb934450.aspx
- http://msdn.microsoft.com/en-us/library/aa214005(SQL.80).aspx
- http://msdn.microsoft.com/en-us/library/ms162773.aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Automating Administrative Tasks by Using SMO



Key Points

SMO is a collection of objects that you can use for managing all aspects of Microsoft SQL Server programmatically. You can use SMO to build customized SQL Server management applications. Although SSMS is a powerful and extensive application for managing SQL Server, in some situations, using an SMO is a better option.

Some examples include:

- To simplify the user applications that control the SQL Server management tasks to meet the needs of new users and to reduce training costs.
- To re-create the index for improving efficiency of indexes.
- To include third-party hardware or software seamlessly into the database management application.

The SMO object model extends and supersedes the Distributed Management Objects (SQL-DMO) object model. Compared to SQL-

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

DMO, SMO increases performance, control, and ease of use. Most SQL-DMO functionality is included in SMO, and there are various new classes that support new features in SQL Server. The object model is intuitive and uses SQL-DMO terminology, where it is possible, to help transfer your skills.

Because SMO is compatible with SQL Server 2000, SQL Server 2005, and SQL Server 2008, you easily manage a multiversion environment.

SMO Classes

SMO includes a number of .NET classes which are used to perform different operations in SQL Server or represent an object in SQL Server configuration. SMO classes include two categories: instance classes and utility classes.

Instance Classes

Instance classes represent SQL Server objects such as servers, databases, tables, triggers, and stored procedures. The ServerConnection class is used to establish a connection to the instance of SQL Server and control the capture mode of commands sent to it.

The SMO instance objects form a hierarchy that represents the hierarchy of a database server. At the top, you have the instances of SQL Server, under which are the databases, and then the tables, columns, and triggers. If it is logical that there is a one parent to many children relationship, such as a table having one or more columns, then the child is represented by a collection of objects; otherwise, the child is just represented by an object.

Utility Classes

Utility classes are a group of objects that have been created explicitly to perform specific tasks. They have been divided into different object hierarchies based on function:

- **Transfer class**. This is used to transfer the schema and data to another database.
- **Backup and Restore classes**. These are used to back up and restore databases.
- **Scripter class**. This is used to create script files for the regeneration of objects and their dependencies.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Steps to Create SMO Applications

The steps to create SMO applications involve:

1. Create a Visual Studio® project and add reference to the following libraries:
   - Microsoft.SqlServer.ConnectionInfo.dll
   - Microsoft.SqlServer.Smo.dll
   - Microsoft.SqlServer.SmoEnum.dll
   - Microsoft.SqlServer.SqlEnum.dll
2. Create an instance of the Server object and establish a connection to the instance of SQL Server. You can do this in the following three ways:
   - Use a server connection object variable to provide the connection information.
   - Provide the connection information by explicitly setting the Server object properties.
   - Pass the name of the SQL Server instance in the Server object constructor.
3. Add a code to call a method to perform specific tasks related to the object, such as issuing a checkpoint on a database or requesting an enumerated list of logons for the instance of Microsoft SQL Server.
4. Add a code to handle exception. In managed code, exceptions are thrown when an error occurs. SMO methods and properties do not report success or failure in the return value. Instead, exceptions can be caught and handled by an exception handler.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Source:

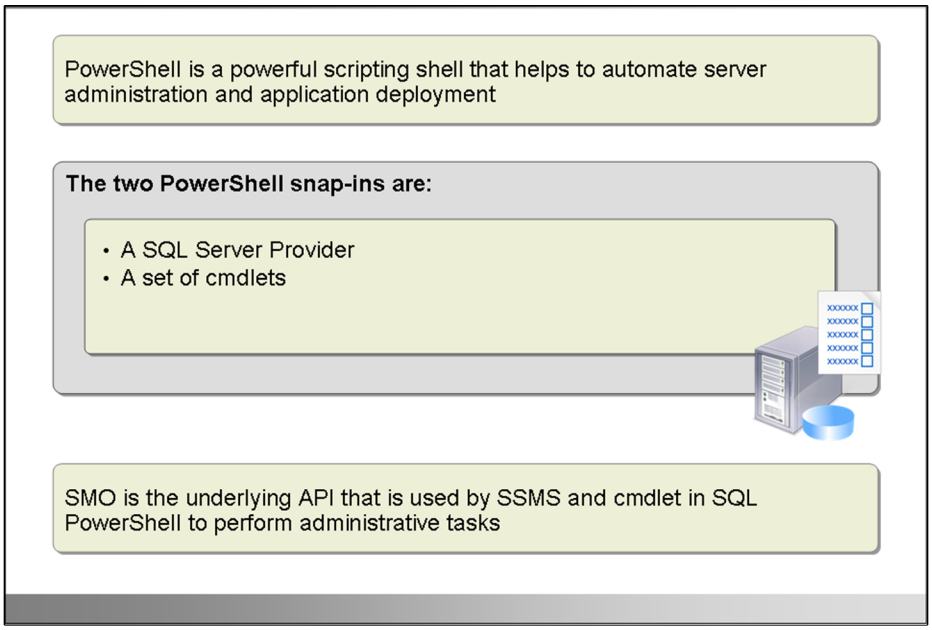
- http://msdn.microsoft.com/en-us/library/ms180944.aspx
- http://msdn.microsoft.com/en-us/library/bb934450.aspx
- http://msdn.microsoft.com/en-us/library/aa214005(SQL.80).aspx
- http://msdn.microsoft.com/en-us/library/ms189595.aspx

For more information, see SMO Object Model.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Introduction to SQL Server PowerShell



Key Points

SQL Server 2008 provides support for Microsoft PowerShell, which is a powerful scripting shell that lets you and developers automate server administration and application deployment. The PowerShell language supports more complex logic than T-SQL scripts, giving you the ability to build robust administration scripts. You can also use PowerShell scripts to administer other Microsoft server products. This provides a common scripting language across servers.

SQL Server provides the following PowerShell implementations:

- **A SQL Server provider**. This implementation which enables a simple navigation mechanism similar to file system paths. You can build paths similar to file system paths, where the drive is associated with an SMO model, and the nodes are based on the object model classes. You can then use commands, such as cd and dir, to navigate the paths similar to the way you navigate folders in a Command Prompt window. You can use other

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

commands, such as ren or del, to perform actions on the nodes in the path.

- **A set of cmdlets**. These are commands used in PowerShell scripts to specify a SQL Server action. The SQL Server cmdlets support actions such as running a sqlcmd script containing Transact-SQL or XQuery statements.

Note: SQL Server 2008 only supports PowerShell scripts for the technologies supported by the SMO, the Database Engine, and Service Broker.

Relationship Between SMO and PowerShell

SMO is the underlying application programming interface (API) that is used by SSMS and cmdlet in SQL PowerShell to perform administrative tasks. By using SSMS, you can create your own script or application to perform administrative tasks. The SQL Server PowerShell provider exposes the hierarchy of SQL Server objects in paths similar to file system paths. You can use the paths to locate an object, and then use methods from the SMO models to perform actions on the objects.

Products whose data or object models can be represented in a hierarchy use Windows PowerShell™ providers to expose the hierarchies. The hierarchy is exposed by using a drive and path structure similar to what the Windows file system uses.

Each PowerShell provider implements one or more drives. Each drive is the root node of a hierarchy of related objects. The SQL Server PowerShell provider implements a SQLSERVER: drive. The SQLSERVER: drive has four primary folders. Each folder and its subfolders represent the set of objects that can be accessed by using a SQL Server management object model. When you are focused on a subfolder in a path that starts with one of these primary folders, you can use the methods from the associated object model to perform actions on the object that is represented by the node.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

The following table lists the PowerShell folders implemented by the SQL Server 2008 PowerShell provider.

Folder	SQL Server object model namespace	Objects
SQLSERVER:\SQL	Microsoft.SqlServer.Management.Smo	Database objects
SQLSERVER:\SQLPolicy	Microsoft.SqlServer.Management.Smo.Agent	Policy-based management objects
SQLSERVER:\SQLRegistration	Microsoft.SqlServer.Management.Smo.Broker	Registered server objects
SQLSERVER:\DataCollection	Microsoft.SqlServer.Management.Smo.Mail	Data collector objects

For example, you can use the SQLSERVER:\SQL folder to start paths that can represent any object that SMO object model supports. The leading part of a SQLSERVER:\SQL path is SQLSERVER:\SQL\ComputerName\InstanceName. You must specify a computer name. You can specify either localhost or `(local`) for the local computer. You must always specify the instance name, even for default instances. For default instances, specify DEFAULT. The nodes after the instance name alternate between object classes such as Database or View, and object names such as QuantamCorp. You cannot represent schemas as object classes. When you specify the node for a top-level object in a schema, such as a table or view, you must specify the object name in the format SchemaName.ObjectName.

The following is the path of the Vendor table in the Purchasing schema of the QuantamCorp database in a default instance of the Database Engine on the local computer.

```
cd SQL\localhost\DEFAULT\Databases\QuantamCorp\TablesGet-ChildItem
```

Object class nodes in a path are associated with a collection class in the associated object model.

The following table lists the object name nodes that are associated with an object class in the associated object model.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Path	SMO class
SQLSERVER:\SQL\locahost\DEFAULT\Databases	DatabaseCollection
SQLSERVER:\SQL\localhost\DEFAULT\Databases\QuantamCorp	Database

Whenever you reference an instance of the Database Engine in a path, the SQL Server PowerShell provider uses SMO to open a Windows Authentication connection to the instance. The connection is made by using the credentials of the Windows account running the PowerShell session. The SQL Server PowerShell provider does not use SQL Server Authentication.

Question: What is the difference between SQL PowerShell and sqlcmd?

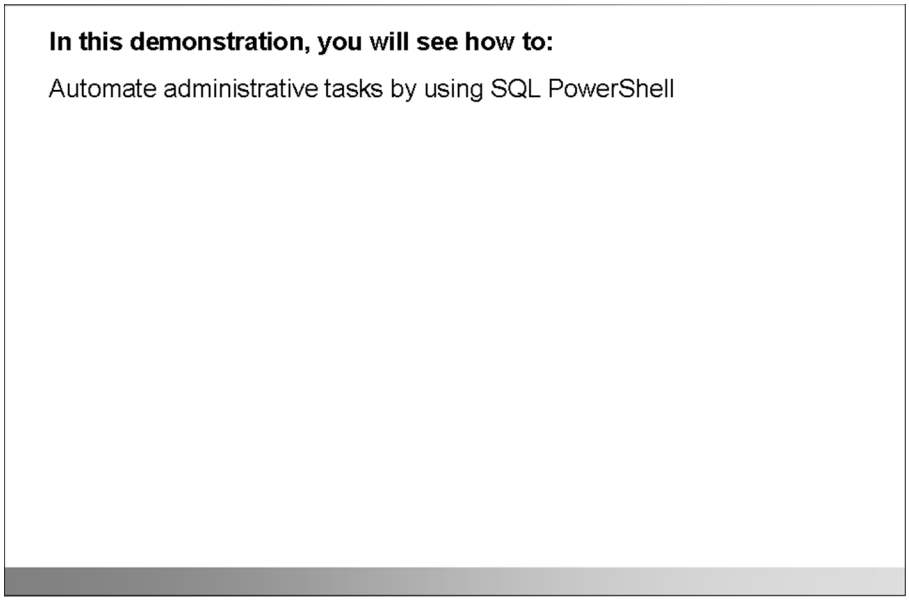
Source:

- http://msdn.microsoft.com/en-us/library/cc281954(SQL.105).aspx
- http://msdn.microsoft.com/en-us/library/cc281947(SQL.105).aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Demonstration: How To Automate Administrative Tasks by Using SQL PowerShell



Key Points

The steps to automate administrative tasks by using SQL Powershell are:

1. On the **Start** menu, point to **All Programs**, point to **Accessories**, and then click **Windows Explorer**.
2. In Windows Explorer, browse to the **D:\Demofiles\Mod06** folder, double-click the **M6_Demo2.sql** file, and then click **Execute**.
3. On the **Start** menu, click **Run**, and then run the **sqlps** command. The SQL PowerShell prompt appears.
4. At the sqlcmd prompt, run the following commands.

```
Encode-SqlName "Table:Test"
Decode-SqlName "Table%3ATest"
```

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

5. At the sqlcmd prompt, run the following command.

```
Invoke-Sqlcmd -Query "SELECT GETDATE() AS TimeOfQuery;"
   -ServerInstance "localhost"
```

6. At the sqlcmd prompt, run the following commands.

```
cd "SQLSERVER:\SQLPolicy\localhost\DEFAULT\Policies"


Get-Item "Database Status" | Invoke-PolicyEvaluation -
   TargetServerName "localhost"
```

7. In the sqlcmd prompt, run the following commands.

```
Set-Location SQLSERVER:\SQL\localhost\DEFAULT\Databases

Get-ChildItem -force
```

8. In the sqlcmd prompt, run the following commands.

```
Set-Location
   SQLSERVER:\SQL\localhost\DEFAULT\Databases\QuantamCo
   rp\Tables

Get-ChildItem | where {$_.Schema -eq "<schema name>"}
```

9. In the sqlcmd prompt, run the following commands.

```
Set-Location SQLSERVER:\SQL\localhost\DEFAULT\Databases

$MyDBVar = New-Object
   Microsoft.SqlServer.Management.SMO.Database

$MyDBVar.Parent = (Get-Item ..)

$MyDBVar.Name = "PS_Demo"

$MyDBVar.Create()

$MyDBVar.State
```

10. In the C:\Program Files\Microsoft SQL Server\100\Tools\Binn\SQLPS.exe window, click the **Close** button.

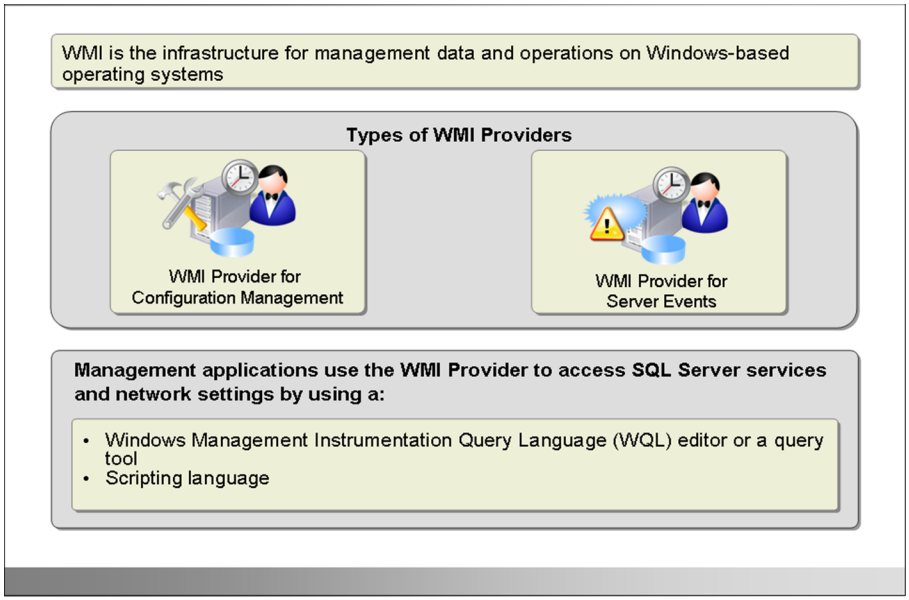
11. Close all open windows.

Question: What is the difference between SQL PowerShell and Windows PowerShell?

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Automating Administrative Tasks by Using WMI



Key Points

WMI is the interface for applications to obtains management data and performs operations on Windows-based operating systems.

The WMI provider is a published layer that is used by the Microsoft SQL Server Configuration Manager. It provides a unified way for interfacing with the API calls that manage the registry operations requested by SQL Server Configuration Manager and/or management applications or scripts.

The SQL Server WMI provider is a dynamic-link library (DLL) and a Managed Object Format (.mof) file, which are compiled automatically by SQL Server Setup. It contains a set of object classes used to control the SQL Server services by using the following methods:

- A script language, such as VBScript, JScript, or Perl, in which Windows Management Instrumentation Query Language (WQL) can be embedded
- The ManagedComputer object in an SMO-managed code program

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

WMI Providers in SQL Server

The various WMI providers in SQL Server are:

- **WMI Provider for Configuration Management**. This provider lets you use WMI to manage SQL Server services, SQL Server client and server network settings, and server aliases. SQL Server services, network settings, and aliases are represented by WMI objects in the root\Microsoft\SqlServer\ComputerManagement namespace of the computer. After a connection is established with the WMI provider on the specified computer, you can query the services, network settings, and aliases by using WQL or a scripting language.
- **WMI Provider for Server Events**. This provider lets you use WMI to monitor events in SQL Server. Any event that can generate an event notification in SQL Server can be leveraged by the WMI by using this provider. Additionally, as a management application that interacts with the WMI, SQL Server Agent can respond to these events, increasing the scope of events covered by SQL Server Agent over earlier releases.

Using WQL and Scripting Languages with WMI

Management applications use the WMI Provider to access SQL Server services and network settings for Configuration Management objects by using a:

- WQL editor or query tool, such as WBEMTest.exe, to query the object set with the WQL.
- Scripting language, such as VBScript.

Management applications/scripts can access WMI information by issuing WMI Query Language (WQL) statements. WQL is a simplified subset of structured query language (SQL), with some WMI-specific extensions. In using WQL, an application retrieves information against a specific database or database object.

You can create a VBScript program that lists the version of installed instances of SQL Server that are running on a computer.

The following code example lists the instances of SQL Server running on the computer and its version:

- Open a new document in a text editor, such as Microsoft Notepad, and copy the following code.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.


```
set wmi =
GetObject("WINMGMTS:\\.\root\Microsoft\SqlServer\Comput
erManagement")

for each prop in wmi.ExecQuery("select * from
SqlServiceAdvancedProperty where SQLServiceType = 1 AND
PropertyName = 'VERSION'")

WScript.Echo prop.ServiceName & " " & prop.PropertyName
& ": " & prop.PropertyStrValue

next
```

The first line in the code, connects to an instance of the WMI Provider for Computer Management at the local computer by using the VBScript GetObject function. The InstancesOf method enumerates a list of services. You can also enumerate these services by using a simple WQL query and an ExecQuery method, instead of the InstancesOf method. The ExecQuery method and a WQL query helps to retrieve the name and the version of the installed instances of SQL Server.

2. Save the file as **test.vbs**.
3. At the command prompt, type **cscript test.vbs** to run the file.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Question: What is the difference between SMO and WMI?

Question: How can you perform WMI queries and trigger actions?

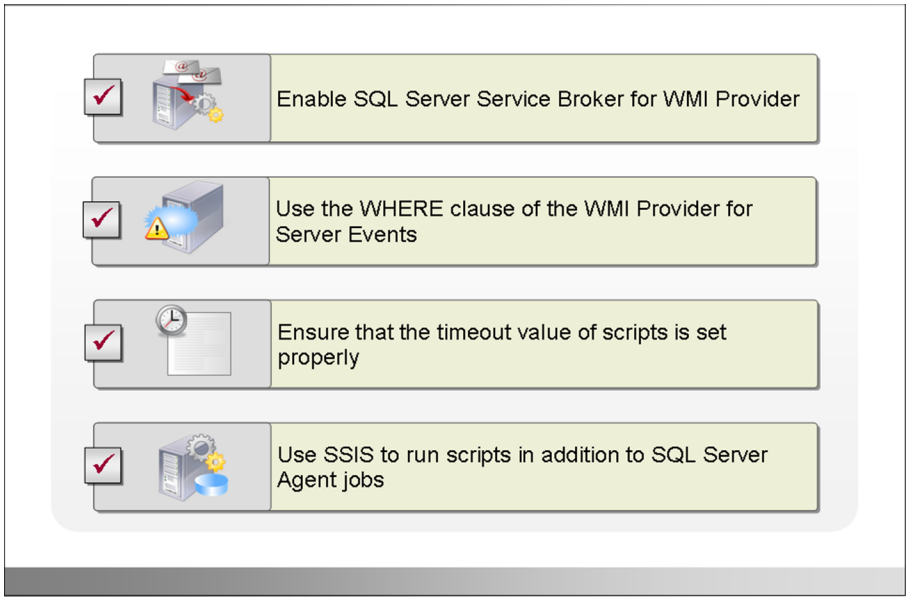
Source:

- http://technet.microsoft.com/en-us/library/ms180499(SQL.90).aspx
- http://msdn.microsoft.com/en-us/library/ms182491(SQL.90).aspx
- http://msdn.microsoft.com/en-us/library/ms181893(SQL.90).aspx
- http://technet.microsoft.com/zh-cn/library/ms186353(SQL.90).aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Considerations for Using Scripts



Key Points

You can use scripts to automate administrative tasks or perform tasks multiple times in exactly the same way. When using scripts to work with SQL Server, you should keep in mind the following considerations:

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

- **Enable SQL Server Service Broker for WMI Provider**. You need to enable Service Broker because the event notifications created by the WMI Provider uses SQL Server to send messages about server events. If a program queries events on a server instance, you must enable Service Broker in msdb of that instance, because that is the location of the target Service Broker service that is created by the provider. You must enable Service Broker in the target database, if your program queries events in a database or on a particular database object. If the corresponding Service Broker is not enabled after your application is deployed, any events generated by the underlying event notification are sent to the queue of the service used by the event notification. However, the underlying event notification is not returned to the WMI management application until you enable Service Broker.
- **Use the WHERE clause of the WMI Provider for Server Events**. WMI Provider for Server Events uses the WHERE clause to control the registration. The where_condition condition of the WMI Provider for Server Events syntax determines:
  - The scope by which the provider tries to retrieve the specified event_type, the server level, the database level, or the object level, of which the only object currently supported is queue. This scope determines the type of event notification created in the target database. This process is called event notification registration.
  - The database, schema, and object, where appropriate, on which to register.
- **Ensure that the timeout value of scripts is set properly**. Scripts may take time to run. Configure the timeout value of scripts properly so that the script can run until it completes. If you run a script by using cscript, the script will run indefinitely without user intervention.
- **Use SSIS to run scripts in addition to SQL Server Agent jobs**. Using SQL Server Agent job provides basic flow control to administrative jobs. With the help of SSIS, the administrator can define a more complex execution flow and condition for administrative tasks. The administrator can also process the result of the administrative tasks and use the result to determine the task to perform next in the SSIS package.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

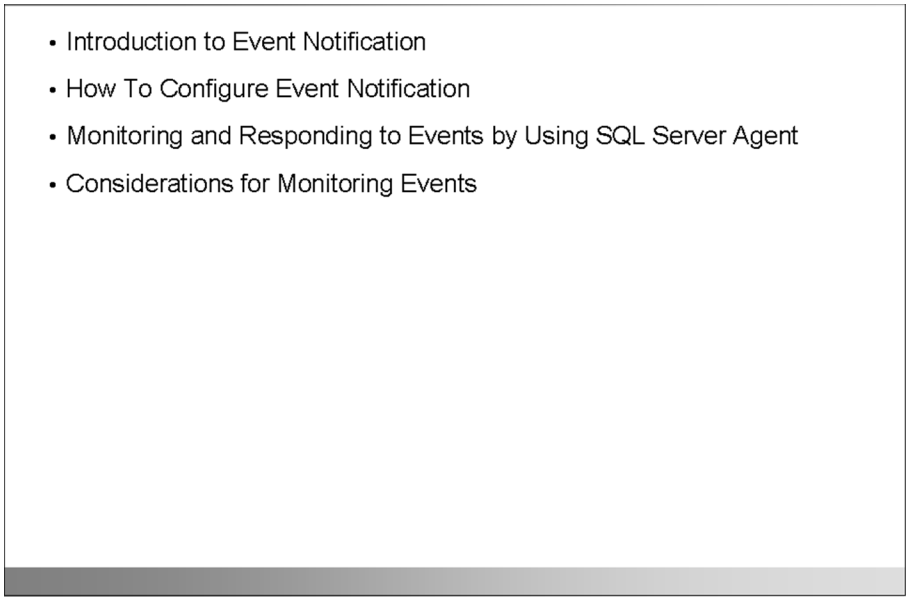
Source:

- http://technet.microsoft.com/es-es/library/ms186371(SQL.90).aspx
- http://msdn.microsoft.com/en-us/library/ms180524(SQL.90).aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Lesson 3
Monitoring Events in SQL Server



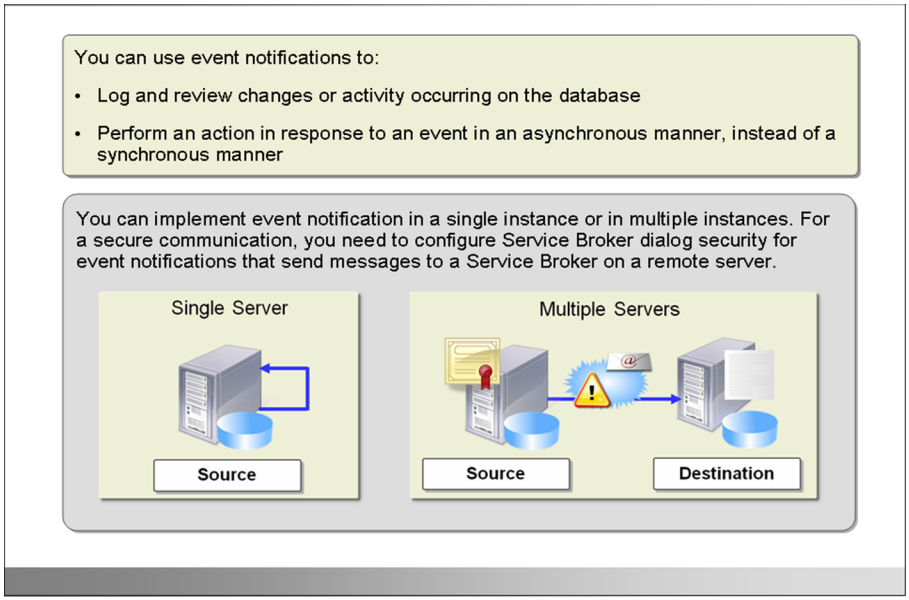
In SQL Server 2008, event monitoring solutions are used to collect information about the database and any related applications when certain events occur. To monitor SQL Server, you can deploy some comprehensive monitoring solutions, such as SCOM. In the absence of a monitoring solution in place, you can configure SQL Server to monitor the events by using event notifications and SQL Server Agent.

In this lesson, you will learn to configure event notification. You will also learn how to monitor events by using SQL Server Agent. In addition, you will also learn the various considerations for monitoring events.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Introduction to Event Notification



Key Points

Event notifications are triggered in response to a variety of T-SQL data definition language (DDL) statements and SQL Trace events, by sending information about these events to a Service Broker service.

You can use event notifications to:

- Log and review changes or activity occurring on the database.
- Perform an action in response to an event in an asynchronous manner, instead of a synchronous manner.

When you create an event notification, one or more Service Broker conversations between an instance of SQL Server and the target service you specify are created. The conversations usually remain open as long as the event notification exists as an object on the server instance. In certain instances, the conversations can close before the event notification is dropped. These conversations are never shared between event notifications. Every event notification has its own exclusive conversation. If you end a conversation

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

explicitly, the action prevents the target service from receiving more messages. Therefore, the conversation does not reopen the next time the event happens.

Event information is delivered to the Service Broker as a variable of type XML that provides information about when an event occurs, about the database object affected, and the Transact-SQL batch statement involved.

You can use the event data to manage SQL Server to track server events, and use that information as the basis to make decisions.

For example, the following event notification sends a notice to Service Broker services every time an ALTER TABLE statement is issued in the QuantamCorp sample database:

```
USE QuantamCorp

GO

CREATE EVENT NOTIFICATION NotifyALTER_T1

ON DATABASE

FOR ALTER_TABLE

TO SERVICE '//Quantam-Corp.com/ArchiveService' ,

    '8140a771-3c4b-4479-8ac0-81008ab17984';
```

Steps to follow while configuring an event notification are:

- Enable Service Broker
- Create a response to a particular event
- Configure routing objects like Queue, services, and route
- Enable event notification for your desired event

Comparison Between Event Notifications and Triggers

Event notifications can offer a programming alternative to DDL triggers and SQL Trace, and run asynchronously, outside the scope of a transaction. Therefore, unlike DDL triggers, you can use event notifications inside a database application to respond to events without using any resources defined by the immediate transaction.

Comparison Between Event Notifications and SQL Trace

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Unlike SQL Trace, you can use event notifications to perform an action inside an instance of SQL Server in response to a SQL Trace event.

Secure Communications between Servers

You can implement event notification in a single instance or in multiple instances. In a single instance, the source and destination database are on the same computer, so dialog security is not required; but if the source and the destination server are not on the same computer, you need to configure dialog security by using certificates. You can use certificates create a secure channel for exchanging information between the servers.

You need to configure Service Broker dialog security for event notifications that send messages to a Service Broker on a remote server. The full-security model enables encryption and decryption of messages that are sent to and from remote servers. Although event notifications are sent in one direction, other messages, such as errors, are also returned to the source of the event.

While you use event notification, you need to ensure that it does not impact the performance of the server hosting production data. The event notification does not occur in real time. This ensures the database transaction is complete before the event information is sent to the destination database.

Question: What is difference between event forwarding in SQL Server Agent and event notification?

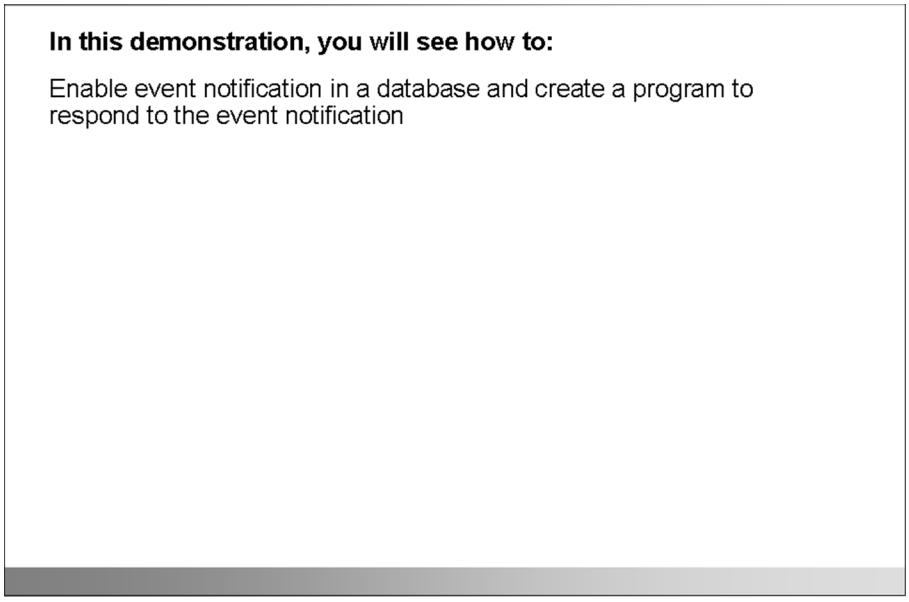
Source:

- http://msdn.microsoft.com/en-us/library/ms190427(SQL.90).aspx
- http://technet.microsoft.com/en-us/library/ms191543(SQL.105).aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Demonstration: How To Configure Event Notification



Key Points

The steps to configure event notification are:

1. On the **Start** menu, point to **All Programs**, point to **Accessories**, and then click **Windows Explorer**.
2. In Windows Explorer, browse to the D:\Demofiles\Mod06 folder, and then double-click the **M6_Demo3.sql** file.
3. In the Microsoft SQL Server Management Studio window, in the Query Editor pane, select the **Create a test table MsgLog** code, and then click **Execute**.
4. In the Microsoft SQL Server Management Studio window, in the Query Editor pane, select the **Enable service broker** code, and then click **Execute**.
5. In the Microsoft SQL Server Management Studio window, in the Query Editor pane, select the **Create a service broker queue** code, and then click **Execute**.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

6. In the Microsoft SQL Server Management Studio window, in the Query Editor pane, select the **Create a stored procedure to log information from event notifications** code, and then click **Execute**.
7. In the Microsoft SQL Server Management Studio window, in the Query Editor pane, select the **Modify the queue object to support automatic activation** code, and then click **Execute**.
8. In the Microsoft SQL Server Management Studio window, in the Query Editor pane, select the **Create a service between the queue and the event notification messages** code, and then click **Execute**.
9. In the Microsoft SQL Server Management Studio window, in the Query Editor pane, select the **Route messages to the service** code, and then click **Execute**.
10. In the Microsoft SQL Server Management Studio window, in the Query Editor pane, select the **Create an event notification** code, and then click **Execute**.
11. In the Microsoft SQL Server Management Studio window, in the Query Editor pane, select the **Create the table1 table** code, and then click **Execute**.
12. In the Microsoft SQL Server Management Studio window, in the Query Editor pane, select the **Check if any information is listed in the MsgLog table** code, and then click **Execute**.
13. Close all open windows.

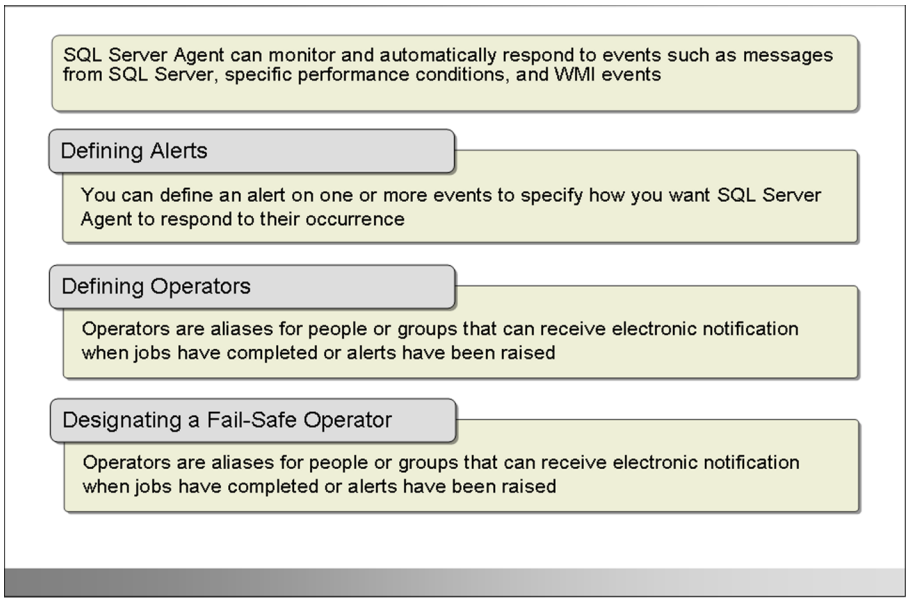
Question: Do you need to configure Service Broker before you can use event notification?

Question: Can you send events to other machines in event notification?

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Monitoring and Responding to Events by Using SQL Server Agent



Key Points

SQL Server Agent can monitor and automatically respond to events such as messages from SQL Server, specific performance conditions, and WMI events. Events are generated by SQL Server and entered into the Microsoft Windows application log.

Defining Alerts

Automated responses to events are called alerts. You can define an alert on one or more events to specify how you want SQL Server Agent to respond to their occurrence. An alert can respond to an event by notifying you or by running a job, or both. An alert can also forward an event to the Microsoft Windows application log on a different computer. For example, you can specify that an operator be notified immediately if an event of severity 19 occurs. By defining alerts, you can more effectively monitor and manage SQL Server.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

When a SQL Server Agent alert is defined for a performance counter, SQL Server Agent directly monitors the performance counter. For a WMI event, SQL Server Agent registers an event query.

To respond to messages from SQL Server, SQL Server Agent monitors the Windows application log. SQL Server Agent can only respond to messages that appear in this log.

When SQL Server logs a message, the SQL Server Agent service compares the message against the alerts defined by you. Regardless of the source of the event, the SQL Server Agent service responds to the event by performing the tasks specified in the alert for the event.

To monitor events, you need to perform the following steps in SQL Server Agent:

1. Configure SQL Server Agent.
2. Define an operator.
3. Create alerts to monitor events.
4. Link an operator to an alert to receive a notification when the alert is triggered.

To define an alert, you need to specify the following:

- The name of the alert
- The event or performance condition that triggers the alert
- The action that SQL Server Agent takes in response to the event or performance condition
- The type of the event determines the parameters that you use to specify the precise event. An alert responds to an event of a specific type. Alerts respond to the following event types:
- SQL Server events
- SQL Server performance conditions
- WMI events

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Defining Operators

The SQL Server Agent service supports the notification of administrators through operators. Operators are aliases for people or groups that can receive electronic notification when jobs have completed or alerts have been raised. The primary attributes of an operator are:

- Operator name
- Contact information

While defining operators, you should remember that the SQL Server Agent service supports the notification of administrators through operators. You must define operators before you define alerts.

Designating a Fail-Safe Operator

The fail-safe operator receives an alert notification after all pager notifications to the designated operators have failed. The fail-safe operator is notified when:

- The operators responsible for the alert could not be paged. The reasons for failure to reach primary operators include incorrect pager addresses and off-duty operators.
- SQL Server Agent cannot access system tables in the msdb database.
  The sysnotifications system table specifies operator responsibilities for alerts.

The fail-safe operator is a security feature. You cannot delete the operator assigned to the fail-safe duty without reassigning the fail-safe duty to another operator, or deleting the fail-safe assignment altogether.

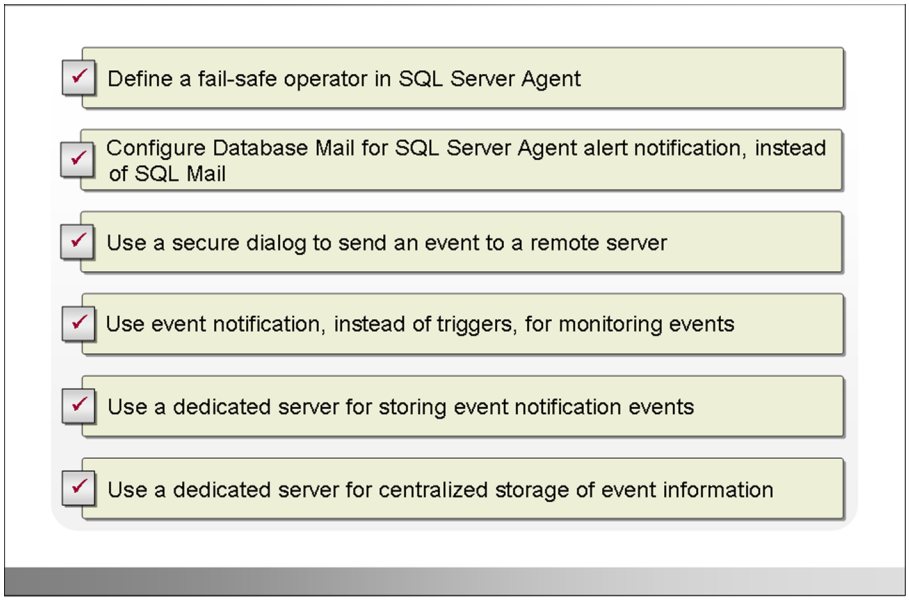
Source:

- http://msdn.microsoft.com/en-us/library/ms191508(SQL.105).aspx
- http://msdn.microsoft.com/en-us/library/aa213820(SQL.80).aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Considerations for Monitoring Events



Key Points

SQL Server Agent provides a number of features for you to perform proactive management.

To monitor events, you need to consider the following:

- **Define fail-safe operators in SQL Server Agent**. The fail-safe operator receives an alert notification after all pager notifications to the designated operators have failed.
- **Configure Database Mail for SQL Server Agent alert notification, instead of SQL Mail**. The SQL Mail feature is only for backward compatibility.
- **Use a secure dialog to send an event to a remote server**. This helps to protect the information exchanged between the remote server and the server sending the events.
- **Use event notification, instead of triggers, for monitoring**

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

events. This helps to reduce the time of monitoring the setup and also avoid the instances of using incorrect coding that may cause issues in database.

- **Use a dedicated server for storing event notification events**. This provides a consolidated view of what is happening in the environment.
- **Use a dedicated server for centralized storage of event information**. This helps reduce the overhead of event information being processed on the server. This also helps to enable centralized reporting.

You should remember that event notification can be used to maintain a log on the system or database events. This helps to minimize the performance impact that would be caused by a trigger.

Question: Why should you use Database Mail over SQL mail?

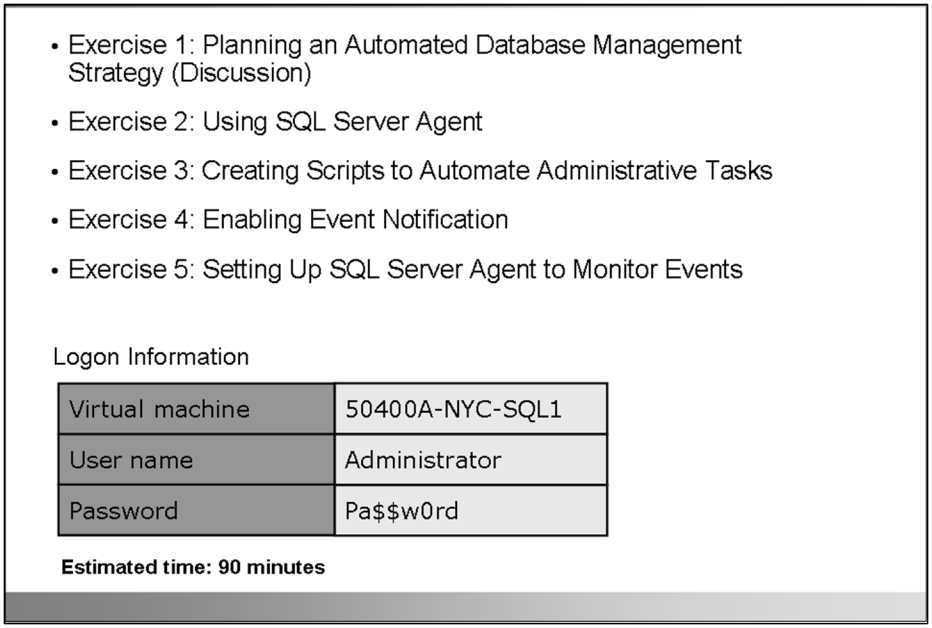
Source:

- http://msdn.microsoft.com/en-us/library/aa213820(SQL.80).aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Lab 6: Automating the Database Management Strategy for SQL Server 2008



Objectives:

After completing this lab, you will be able to:

- Plan an automated database management strategy.
- Use SQL Server Agent.
- Create scripts to automate administrative tasks.
- Enable Event Notification.
- Set up SQL Server Agent to monitor events.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Scenario

QuantamCorp has deployed two SQL Server 2008 servers in their organization and will continue to deploy more to replace theirexisting installation of SQL Server 2005. After this, SQL Server 2008would be the primary database engine used in their organization. Asthe application that is supported by SQL Server 2008 is a veryimportant application, you need to know if any unusual eventhappens on your server. You would also like to automate some of theroutine tasks so that you can be free to upgrade other SQL Server2005 instances to SQL Server 2008.

To achieve this, you need to perform the following:

- Determine the steps to implement automated administrative tasks with SQL Server Agent.
- Determine when to use alerts.
- Decide if event forwarding or event notification has to be used for a scenario.
- Determine the usage of scripts.
- Implement Event Notification in SQL Server.
- Configure SQL Server Agent for automation.
- Configure SQL Server for event monitoring.
- Implement scripts to perform administrative tasks.
- Configure SQL Server Agent to use the scripts.

Exercise 1: Planning an Automated Database Management Strategy (Discussion)

Scenario

The management in your organization initiated a project to reduce the administrative cost for SQL Servers. You submitted a proposal to automate a number of administrative tasks to reduce the time spent every day to perform these tasks.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Your management has asked you to implement this proposal. Before you implement the solution in the production environment, your management wants you to perform a test on automating the following tasks:

- Backing up system databases
- Responding to system events
- Consolidating events information from multiple servers
- Logging and monitoring events that happen at the database level and perform responsive actions on events occurring on the database that support the applications.

Based on the above scenario, answer the following questions:

1. Which strategy would you propose to reduce the operating cost and the administrative effort in managing backup for your SQL Server?
2. Which strategy should you suggest to automate the monitoring of system events?
3. If you want to create an automated response or log events centrally when an error occurs on any of the managed servers, would you consider using event forwarding or event monitoring to monitor events on the servers?
4. When should you consider using event notification over event monitoring?
5. Why should you use scripts rather than performing tasks manually?

Exercise 2: Using SQL Server Agent

Scenario

Currently, you have around 50 servers to manage and you are asking an operator to perform system database backup manually. As the number of servers could grow upto around 100 servers, you decide to configure SQL Server Agent to perform the system database backup automatically in order to save time for the operator.

The main tasks for this exercise are as follows:

1. Start SQL Server Agent.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

2. Create a job schedule.
3. Create a SQL Server Agent job.
4. Assign a schedule to a job.
5. Test the job.

Task 1: Start SQL Server Agent.

1. Open the SQL Server Configuration Manager window.
2. Using the SQL Server Configuration Manager window, configure **SQL Server Agent** to start automatically.

Task 2: Create a job schedule.

1. Start Microsoft SQL Server Management Studio, by using the following credentials in the Connect to Server dialog box:
   - Server type: **Database Engine**
   - Server name: **NYC-SQL1**
   - Authentication: **Windows Authentication**
2. In the Microsoft SQL Server Management Studio window, open the **Manage Schedules** dialog box for **SQL Server Agent**.
3. Using the **Manage Schedules** dialog box, open the New Job Schedule dialog box.
4. Using the **New Job** dialog box, create a schedule with the following details:
   - Name: **Daily 12:00AM**
   - Schedule type: **Recurring**
   - Occurs: **Daily**
   - Recurs every: **1 day**
   - Occurs once at: **12:00AM**

Task 3: Create a SQL Server Agent job.

1. In the Microsoft SQL Server Management Studio window, open the **New Job** dialog box for **SQL Server Agent** to backup master database.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

2. Using the **New Job** dialog box, create a SQL Server Agent job for the**dbo** user. The **dbo** is a member of the sysadmin server role.

Task 4: Assign a schedule to a job.

1. Using the Microsoft SQL Server Management Studio window, open the **Job Properties – Backup System Databases** dialog box for **SQL Server Agent**.
2. Using the **Job Properties – Backup System Databases** dialog box, open the **Pick Schedule for Job – Backup System Databases** dialog box.
3. Using the **Pick Schedule for Job – Backup System Databases** dialog box, create a schedule for the **Daily 12:00AM** job.

Task 5: Test the job.

1. Using the Microsoft SQL Server Management Studio window, start the **Backup System Database** for the SQL Server Agent job.

Results: After completing this exercise, you should have configured the SQL Server Agent to start automatically, created a job schedule, created a SQL Server Agent job, assigned a schedule to a job, and tested the job.

Exercise 3: Creating Scripts to Automate Administrative Tasks

Scenario

The operators have a number of tasks to perform manually. Currently, operators are following steps listed in a manual. Because the steps are performed manually human errors might occur. In order to reduce errors, you decide to create scripts so that operators can run the scripts instead of performing the steps manually.

The main tasks for this exercise are as follows:

1. Create a VBScript for creating a test database.
2. Create a VBScript to modify the database property.
3. Create a VBScript to query WMI information.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Task 1: Create a VBScript for creating a test database.

1. In the Notepad window, create a script for creating a test database **Mod06** with a log file named, **Mod06_Log01.ldf**. Create a primary file named, **Mod06_Data01.mdf** with the following settings and save the Notepad file at **D:\Data\create_db.vbs**.
   - FileGrowthType: **SQLDMOGrowth_MB**
   - FileGrowth: **10**
2. Run the cscript at the command prompt to verify that the test database is created.

Task 2: Create a VBScript to modify the database property.

1. In the Notepad window, create a script to create two database files, for the **Mod06** database and save the file at

   D:\Data\add_db.vbs, with the following structure.

File name	Settings
Mod06_Lookup01	Current size: 10 Type of filegrowth: SQLDMOGrowth_MB Filegrowth: 5
Mod06_Lookup02	Current size: 10 Type of filegrowth: SQLDMOGrowth_MB Filegrowth: 5

   Attach the **Mod06_LookupFG** filegroup to the database.
2. Run the cscript at the command prompt to verify that four database files exist.

Task 3: Create a VBScript to query WMI information.

1. In the Notepad window, create a script to use WMI to set the startup type of SQL Server Browser services to automatic and then start the service. Save the script file at

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

d:\data\sql_service.vbs.

2. Run the cscript at the command prompt.
3. Using the Services window, verify that the SQL Server Browser service is configured to start automatically.

Results: After completing this exercise, you should have created VBScripts for creating a test database, modified the database property, and queried the WMI information.

Exercise 4: Enabling Event Notification

Scenario

A developer deletes a table by mistake. It takes a long time for you to figure out the cause of the issue. In order to avoid this, you decide to implement event notification to monitor the changes in the database schema. Since a large number of servers need to be managed, you want to centralize the monitoring and reduce the management effort.

The main tasks for this exercise are as follows:

1. Configure Service Broker objects for Event Notification.
2. Enable Event Notification.
3. Test Event Notification.

Task 1: Configure Service Broker objects for Event Notification.

1. Using the Microsoft SQL Server Management Studio window, create a queue, service and a route for the **Mod06** database. Configure three Service Broker objects as follows.
   - Queue

     Name: [dbo].[Notify-Queue]

     Status: ON

     Retention: OFF
   - Service
     - **Name**: //QuantamCorp.com/Notify-Service

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

- **Schema**: http://schemas.microsoft.com/SQL/Notifications/PostEventNotification
- Route
  - **Name**: Notify-Route
  - **Service name**: N'//QuantamCorp.com/Notify-Service'
  - **Address**: N'LOCAL'

Note: Browse to the **D:\Labfiles\Mod06** folder and select the **E4_T1.sql** file to execute the required T-SQL code.

Task 2: Enable Event Notification.

1. In the Query Editor pane of the Microsoft SQL Server Management Studio window, run a T-SQL script to enable the event notification, **CreateDatabaseNotification** on the server to create and drop database events.
2. Create an event notification, **AuditLoginLogoutNotification** on the server for the logon, logout, and log failed events related to audit.

Note: Browse to the **D:\Labfiles\Mod06** folder and select the **E4_T2.sql** file to execute the required T-SQL code.

Task 3: Test Event Notification.

1. In the Query Editor pane of the Microsoft SQL Server Management Studio window, run a T-SQL script to test event notification to create and drop the **Mod06_Event** database event.
2. At the command prompt, open a trusted SQL connection.
3. At the sqlcmd prompt, run a query to view all records from the event notification queues.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

4. In the Query Editor pane of the Microsoft SQL Server Management Studio window, run a T-SQL script to create an event notification from the **[dbo].[Notify-Queue]** queue with a timeout value of **10000**.

   Note: Browse to the **D:\Labfiles\Mod06** folder and select the **E4_T3.sql** file to execute the required T-SQL code.

Results: After completing this exercise, you should have prepared your system for event notification, enabled event notification, and tested the event notification.

Exercise 5: Setting Up SQL Server Agent to Monitor Events

Scenario

Recently a database server goes down but you are not aware of the issue until a user calls in and reports the error. You decide to set up SQL Server Agent to monitor some events and have designated operators notified right after an error occurs.

The main tasks for this exercise are as follows:

1. Create a SQL Server Agent job.
2. Create an alert.
3. Create a response for an alert.
4. Verify SQL Server Agent Monitoring.

Task 1: Create a SQL Server Agent job.

1. Using the Microsoft SQL Server Management Studio window, open the **New Job** dialog box for **SQL Server Agent** to create an Alert Response Job.
2. By using the **New Job** dialog box, create a SQL Server Agent job and configure the job for the **[backup_user]** user.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Task 2: Create an alert.

1. Using the Microsoft SQL Server Management Studio window,open the New Alert dialog box, and create an alert for severity 15with the following settings:
   - Type: **SQL Server event alert**
   - Severity: **015 – Syntax Error in SQL Statements**

Task 3: Create a response for an alert.

1. Using the Microsoft SQL Server Management Studio window, open the **'Alert for Severity 15' alert properties** dialog box.
2. Using the **'Alert for Severity 15' alert properties** dialog box, create a response alert for **Alert for Severity 15**.

Task 4: Verify SQL Server Agent Monitoring.

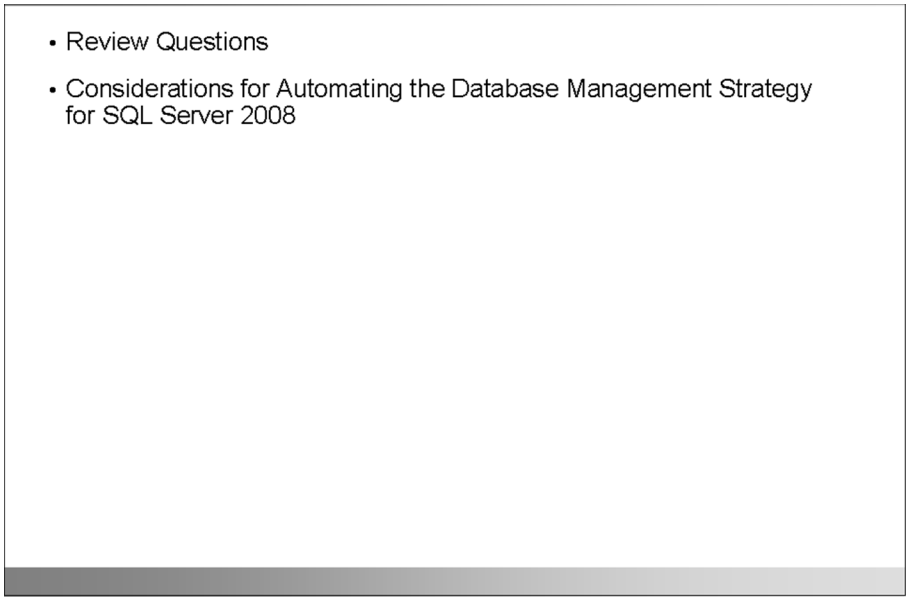
1. In the Query Editor pane of the Microsoft SQL Server Management Studio window, run a T-SQL script to raise an error with the following parameters:
   - Severity level: **15**
   - Error number: **1**
   - Error message: **"This is message %s %d."**
2. View the history of **Alert Response Job** to verify that the job is run.

Results: After completing this exercise, you should have created a SQL Server Agent job, created an alert, created a response for an alert, and verified monitoring of SQL Server Agent.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Module Reviews and Takeaways



Review Questions

1. You are trying to schedule backups for your application databases. You want to ensure that the application databases are backed up in a sequence that is defined by the application team. How would you configure this?
2. You want to review the permissions assigned to SQL Server Agent service account. You only want the service granted required permission but not the service account as the account is shared by multiple SQL Servers. How can you do that?
3. You want to monitor the events happening on a group of SQL Servers and ensure that all events are recorded in a centralized SQL Server database for reporting. What should you configure?
4. You want to get notified by e-mail when an error happens on a SQL Server. You have an operator defined in SQL Server Agent. What will be your next step?

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

5. You want to schedule an SQL script to run on a daily basis. But due to the corporate policy, you need to use another user account to execute the script. What should you configure in SQL Server Agent for this?
6. You want to automate administrative tasks that you are doing manually every day. But you realize that the operation is not supported in T-SQL. What other alternative you can consider to automate the tasks?

Considerations for Automating the Database Management Strategy for SQL Server 2008

When you are automating the Database Management Strategy for SQL Server 2008, consider the following:

- Configure SQL Server Agent to start automatically every time the server starts as jobs would not run if SQL Server Agent is stopped.
- Configure SQL Server Agent to use service SID so you have better control on granting permissions to that particular instance of SQL Server Agent.
- Consider using a proxy for jobs that needs to execute under another user's credentials.
- Consider using VBScript or SQL PowerShell script for administrative tasks that involve actions in addition to SQL Server tasks.
- Consider using WMI to obtain information about SQL Server and other system resources.
- Consider creating scripts for most management jobs so you can reduce the chances of having human errors
- Consider using event notification for monitoring events in one or more SQL Server as this allows the event information to get logged in a centralized database.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Module 7

Designing a Strategy for Securing SQL Server 2008

Contents:

Lesson 1: Securing SQL Server 2008 Instances 7-3
Lesson 2: Planning Database Security 7-40
Lesson 3: Managing Certificates and Keys in SQL Server 2008 7-63
Lesson 4: Implementing Database Encryption 7-80
Lab 7: Designing a Strategy for Securing SQL Server 2008 7-91
Module Review and Takeaways 7-101

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Module Overview

- Securing SQL Server 2008 Instances
- Planning Database Security
- Managing Certificates and Keys in SQL Server 2008
- Implementing Database Encryption

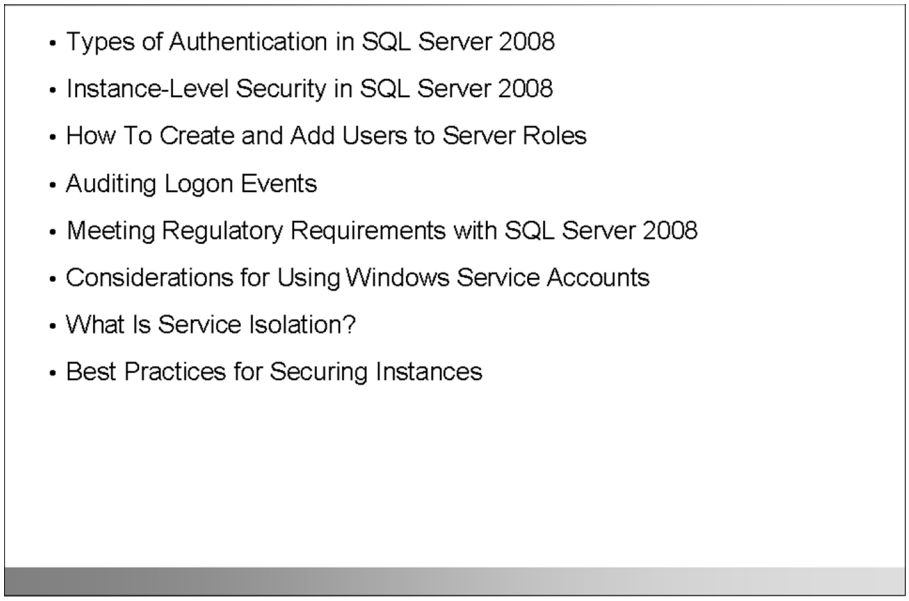
Security is one of the critical features required for a database engine to protect the enterprise data against possible threats. Microsoft™ SQL Server™ 2008 databases store a lot of sensitive information, such as customer information, which must not be accessible to unauthorized users. To enforce security, SQL Server 2008 uses the principle of granting least privileges so that users are not granted more permissions than those necessary to perform their jobs.

In this module, you will learn about the various configurations for securing SQL Server 2008. You will also learn how to implement the security options in SQL Server, such as certificate management and key management. In addition, you will learn how to implement database encryption in SQL Server 2008.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Lesson 1
Securing SQL Server 2008 Instances



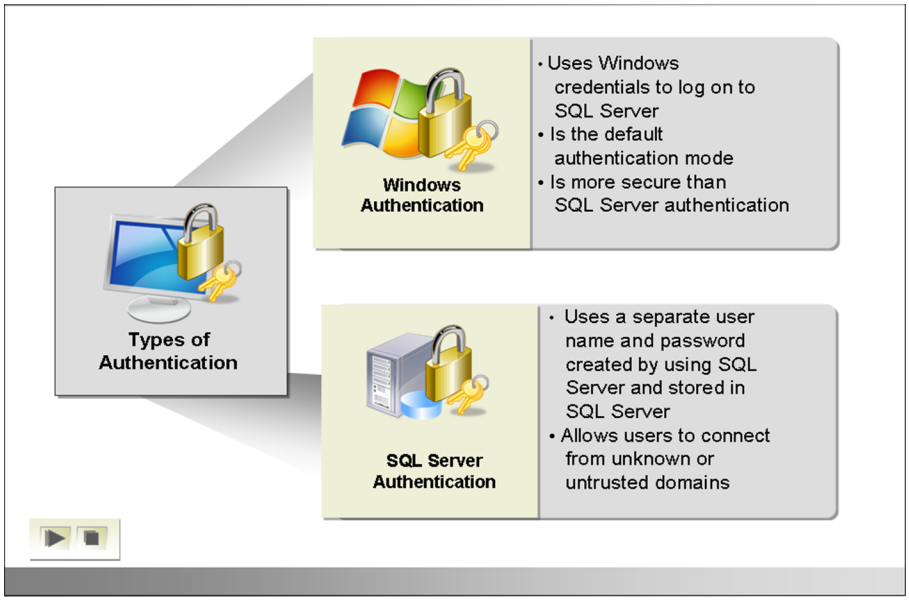
You can secure SQL Server 2008 at the instance level and at the database level. At the instance level, you need to ensure that the server settings are secure. You should ensure that access to the server is controlled and unauthorized access is prevented.

In this lesson, you will learn how to configure SQL Server 2008 instances to prevent unauthorized access to the server. You will learn about principals and securables that provide instance level security. You will also learn about auditing and compliance requirements to secure SQL Server 2008. You will learn about the different service accounts that help manage permissions assigned to the SQL Server instances. You will also learn about service security identifiers (SIDs) that help to secure a service.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Types of Authentication in SQL Server 2008



Key Points

In SQL Server, there are two authentication modes for the database engine that you can select during setup, Windows® authentication and SQL Server authentication. A combination of these two modes is called mixed authentication. In the Windows Authentication mode, Windows Authentication is enabled and SQL Server Authentication is disabled. In the SQL Server and Windows Authentication mode, both SQL Server and Windows Authentication are enabled. Windows Authentication is always available and you cannot disable it.

Windows Authentication

When users connect to SQL Server by using a Windows user account, it validates the account name and password by using the Windows principal token in the operating system. This means that Windows validates the user information and not SQL Server. Windows Authentication is the default authentication mode, and is more secure than SQL Server Authentication. Windows Authentication uses Kerberos security protocol.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Windows Authentication provides:

- Password policy enforcement to validate strong passwords.
- Support for account lockout.
- Support for password expiration.

A connection made by using Windows Authentication is sometimes called a trusted connection, because SQL Server trusts the credentials provided by Windows.

SQL Server Authentication

When using SQL Server Authentication, SQL Server creates the logon names that are not based on Windows user accounts. SQL Server creates and stores both the user name and the password in the master database. When users connect to SQL Server by using SQL Server Authentication, they must provide the logon name and password every time they connect to the server. When you use SQL Server Authentication, you should set strong passwords for all SQL Server accounts. SQL Server Authentication also provides password policy and support for account lockout.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

The following table describes the three optional password policies that are available for SQL Server logons.

Password policy	Description
Change password at next login	This requires users to change the password the next time that the user connects to SQL Server. They can change the password in SSMS. Third-party software developers should provide this feature if this option is used.
Enforce password expiration	The maximum password age policy of the computer is enforced for SQL Server logons.
Enforce password policy	The Windows password policies of the computer are enforced for SQL Server logons. This includes password length and complexity. This functionality depends on the NetValidatePasswordPolicy API, which is only available in Windows Server® 2003 and Windows Server 2008.

Advantages of SQL Server Authentication

SQL Server authentication allows:

- SQL Server to support older applications and applications provided by third-parties that require SQL Server Authentication.
- SQL Server to support environments with mixed operating systems, where all users are not authenticated by a Windows domain.
- Users to connect from unknown or untrusted domains. For example, SQL Server Authentication allows established customers of an application to connect to SQL Server with assigned SQL Server logons to receive the status of their orders.
- SQL Server to support Web-based applications where users create their own identities.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

- Software developers to distribute their applications by using a complex permission hierarchy based on known, preset SQL Server logons.

Note: Using SQL Server Authentication does not limit the permissions of local administrators on the computer where SQL Server is installed.

Disadvantages of SQL Server Authentication

The disadvantages of SQL Server Authentication are:

- SQL Server Authentication requires that Windows domain users who have a login and password for Windows still provide another login and password for SQL Server. Keeping track of multiple names and passwords is difficult for many users.
- SQL Server Authentication cannot use Kerberos security protocol.
- Windows offers additional password policies that are not available for SQL Server logons.

Note: SQL Server Authentication does not limit the permissions of local administrators on the computer where SQL Server is installed.

Configuring the Authentication

To enable SQL Server Authentication during setup, you must provide a strong password for the built-in SQL Server system administrator account named, sa. The sa account connects to SQL Server by using SQL Server Authentication.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

If you select Windows Authentication during setup, the sa account for SQL Server Authentication is created, but disabled. If you later change to Windows and SQL Server Authentication and you want to use the sa account, you must enable the account. You can configure any Windows or SQL Server account. For securing the sa account from unauthorized access, do not enable the account unless your application requires it. Never set a blank or weak password for the sa account.

Question: What is the advantage of Windows Authentication over SQL Server Authentication?

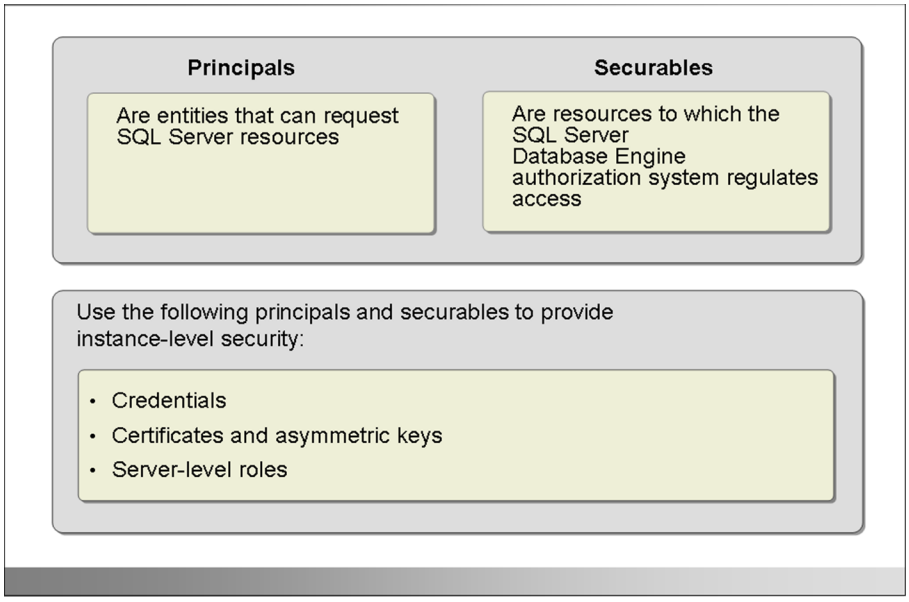
Source:

- msdn.microsoft.com/en-us/library/ms144284.aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Instance-Level Security in SQL Server 2008



Key Points

Instance-level security settings for SQL Server 2008 help to control access to server configuration. This configuration also helps to control which user will have access to the server.

Principals

Principals are entities that can request SQL Server resources. Like other components of the SQL Server authorization model, you can arrange principals in a hierarchy. The scope of influence of a principal depends on the scope of the definition of the principal: Windows, server, and database; and whether the principal is indivisible or a collection. Windows Logon is an example of an indivisible principal, and Windows Group is an example of a principal that is a collection. Every principal has a security identifier (SID).

The following table lists the types of various levels of principals.

Principal	Types
	• Windows Domain Login

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Principal	Types
Windows-level	• Windows Local Login
SQL Server-level	SQL Server Login
Database-level	• Database User • Database Role • Application Role

Securables

Securables are resources to which the SQL Server Database Engine authorization system regulates access. Some securables can be contained within others, creating nested hierarchies called "scopes" that can themselves be secured.

You can provide instance-level security by using the following principals and securables:

Login credentials. A login credential is a record that contains the authentication information required to connect to a resource outside SQL Server. SQL Server can use this information internally. Most credentials contain a Windows-based user name and password. The information stored in a credential enables a user, who has connected to SQL Server by using SQL Server Authentication, to access resources outside the server instance. When the external resource is Windows, the user is authenticated as the Windows user specified in the credential. You can assign a single credential to multiple SQL Server logons. However, you can assign an SQL Server logon to only one credential. System credentials are created automatically and are associated with specific endpoints. Names for system credentials start with two hash signs (##).

Certificates and asymmetric keys. You can use asymmetric encryption by using certificates and asymmetric keys. Certificates are often used as containers for asymmetric keys because they can contain more information, such as expiry dates and issuers. There is no difference between the two mechanisms for the cryptographic algorithm; and for the same key length, there is no difference in strength. You can use a certificate to encrypt other types of encryption keys in a database, or to sign code modules. Certificates and asymmetric keys can decrypt data that the other mechanism encrypts. Generally, you can use asymmetric encryption to encrypt a

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

symmetric key for storage in a database. A public key does not have a particular format that a certificate would have, and you cannot export it to a file.

Server-level roles. To easily manage the permissions on the server, SQL Server provides several roles, which are security principals that group other principals. Roles are like groups in the Microsoft Windows operating system. Server-level roles are also named fixed server roles because you cannot create new server-level roles. Server-level roles are server-wide in their permissions scope. You can add SQL Server logons, Windows accounts, and Windows groups into server-level roles. Each member of a fixed server role can add other logons to that same role.

You can add SQL Server logons, Windows accounts, and Windows groups into server-level roles. Each member of a fixed server role can add other logons to that same role.

The following table shows the server-level roles and their capabilities.

Server-level role	Permission	Description
sysadmin	Granted with GRANT option: CONTROL SERVER	Members of the sysadmin fixed server role can perform any activity on the server.
serveradmin	Granted: ALTER ANY ENDPOINT, ALTER RESOURCES, ALTER SERVER STATE, ALTER SETTINGS, SHUTDOWN, VIEW SERVER STATE	Members of the serveradmin fixed server role can change server-wide configuration options and shut down the server.
securityadmin	Granted: ALTER ANY LOGIN	Members of the securityadmin fixed server role manage logons and their properties. They can GRANT, DENY, and REVOKE server-level permissions. They can also GRANT, DENY, and REVOKE database-level permissions. Additionally,

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Server-level role	Permission	Description
		they can reset passwords for SQL Server logons.
processadmin	Granted: ALTER ANY CONNECTION, ALTER SERVER STATE	Members of the processadmin fixed server role can end processes that are running in an instance of SQL Server.
setupadmin	Granted: ALTER ANY LINKED SERVER	Members of the setupadmin fixed server role can add and remove linked servers.
bulkadmin	Granted: ADMINISTER BULK OPERATIONS	Members of the bulkadmin fixed server role can run the BULK INSERT statement.
diskadmin	Granted: ALTER RESOURCES	The diskadmin fixed server role is used for managing disk files.
dbcreator	Granted: ALTER ANY DATABASE	Members of the dbcreator fixed server role can create, alter, drop, and restore any database.
public	Default	Every SQL Server logon belongs to the public server role. When a server principal has not been granted or denied specific permissions on a securable object, the user inherits the permissions granted to public on that object. Only assign public permissions on any object when you want

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Server-level role	Permission	Description
		the object to be available to all users.

Question: Which server role is used to assign the highest level of permission to a user who needs to manage all aspects of SQL Server?

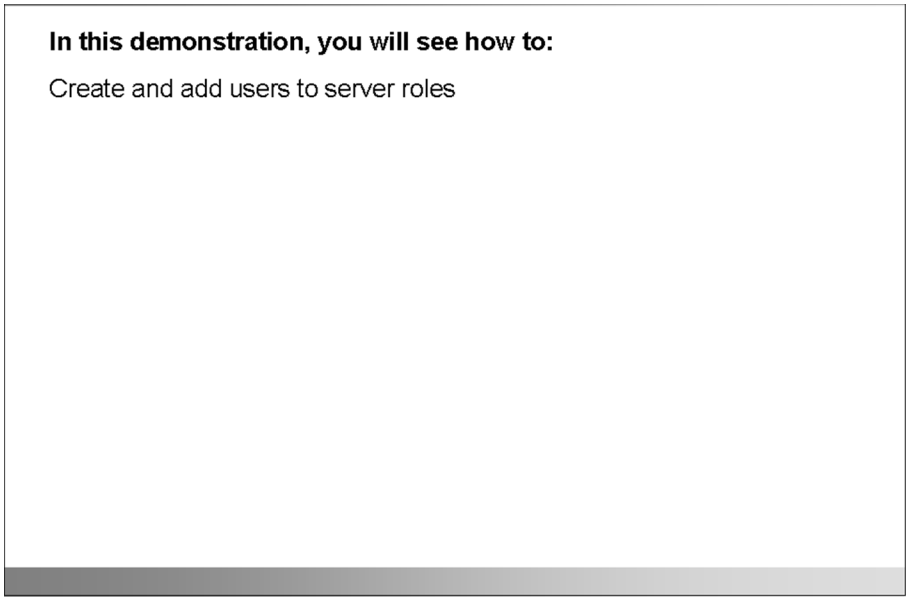
Source:

- http://technet.microsoft.com/en-us/library/ms189522.aspx
- http://msdn.microsoft.com/en-us/library/ms181127.aspx
- msdn.microsoft.com/en-us/library/ms190401.aspx
- msdn.microsoft.com/en-us/library/ms161950.aspx
- msdn.microsoft.com/en-us/library/bb895327.aspx
- http://msdn.microsoft.com/en-us/library/bb895327.aspx
- msdn.microsoft.com/en-us/library/ms188659.aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Demonstration: How To Create and Add Logins to Server Roles



Key Points

The steps to create and add users to server roles are:

1. On the **Start** menu, point to **All Programs**, click **Microsoft SQL Server 2008**, and then click **SQL Server Management Studio**. The Microsoft SQL Server Management Studio window appears.
2. To log on to SQL Server Management Studio, in the Connect to Server dialog box, type the following:
   - Server type: **Database Engine**
   - Server name: **NYC-SQL1**
   - Authentication: **Windows Authentication**
3. Click **Connect**.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

4. In the Microsoft SQL Server Management Studio window, in the Object Explorer pane, under your server, right-click the **Security** folder, point to **New**, and then click **Login**. The **Login – New** dialog box appears.
5. In the **Login – New** dialog box, on the **General** page, in the **Login name** box, type **Ben Darton**.
6. Click **SQL Server Authentication**.
7. In the **Password** box, type **Pa$$w0rd**. Then, in the **Confirm password** box, retype the same password.
8. Ensure that the **Enforce password policy** and the **Enforce password expiration** check boxes are selected.
9. Clear the **User must change password at next login** check box.
10. In the **Default database** list, click **QuantamCorp**, and then click **OK**.
11. On the **Start** menu, point to **All Programs**, point to **Accessories**, and then click **Windows Explorer**.
12. In Windows Explorer, browse to the **D:\Demofiles\Mod07** folder, and then double-click the **M7_Demo1.sql** file.
13. In the **Microsoft SQL Server Management Studio** window, click **Execute**.

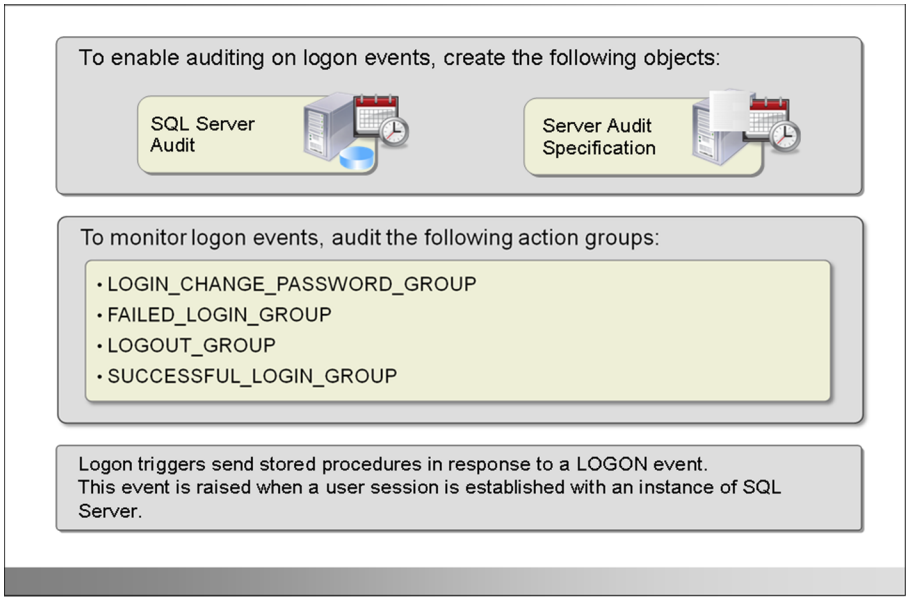
Question: What is the benefit of using server roles?

Question: When should you use SQL Server authentication?

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Auditing Logon Events



Key Points

Auditing an instance of SQL Server or an SQL Server database involves tracking and logging events that occur on the system. There are several levels of auditing for SQL Server, depending on government or standards requirements for your installation. SQL Server Audit provides the tools and processes that you can use to enable, store, and view audits on various server and database objects.

You can record a server audit per instance of action groups, and either database audit action groups or database audit actions per database. The audit event will occur every time you encounter auditable action.

SQL Server Audit Components

An audit is the combination of several elements into a single package for a specific group of server actions or database actions.

Following are the SQL Server Audit Components.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

SQL Server Audit. You can use SQL Server Audit to monitor logon events.

Server Audit Specification. You need to configure server audit specifications to monitor SQL Server Audit.

Audit Action Groups

The following table describes the various audit action groups and provides the equivalent SQL Server Event Class where applicable.

Action group	Description
LOGIN_CHANGE_PASSWORD_GROUP	This event is raised whenever a logon password is changed by way of ALTER LOGIN statement or sp_password stored procedure. It is equivalent to the Audit Login Change Password Event Class.
FAILED_LOGIN_GROUP	This event indicates that a principal tried to log on to SQL Server and failed. Events in this class are raised by new connections or by connections that are reused from a connection pool. It is equivalent to the Audit Login Failed Event Class.
LOGOUT_GROUP	This event indicates that a principal has logged out of SQL Server. Events in this class are raised by new connections or by connections that are reused from a connection pool. It is equivalent to the Audit Logout Event Class.
SUCCESSFUL_LOGIN_GROUP	This event indicates that a principal has successfully logged on to SQL Server. Events in this class are raised by new connections or by connections that are reused from a connection pool. It is equivalent

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Action group	Description
	to the Audit Login Event Class.

Restricting Access to File Locations

When you are saving audit information to a file, to prevent tampering, you can restrict access to the file location in the following ways:

The SQL Server Service Account must have both read and write permission.

You require read and write permission. This assumes that the Audit Administrators are Windows accounts for administration of audit files such as copying them to different shares, and backing up.

Audit readers that are authorized to read audit files must have read permission.

Even when the Database Engine is updating a file, other Windows users can read the audit file, if they have permission. The Database Engine does not take an exclusive lock that prevents read operations.

Because the Database Engine can access the file, SQL Server logons that have the CONTROL SERVER permission can use the Database Engine to access the audit files. To record any user that is reading the audit file, define an audit on master.sys.fn_get_audit_file. This audit records the logons with CONTROL SERVER permission that have accessed the audit file through SQL Server.

If an Audit Administrator copies the file to a different location, such as for archival purposes, the Audit Control Lists on the new location should be reduced to the following permissions:

- Audit Administrator: Read/Write
- Audit Reader: Read

You must generate audit reports from a separate instance of SQL Server, such as an instance of SQL Server Express, to which only Audit Administrators or Audit Readers have access. By using a separate instance of the Database Engine for reporting, you can help prevent unauthorized users from obtaining access to the audit record.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

You can offer additional protection against unauthorized access by encrypting the folder in which the audit file is stored, by using Windows BitLocker™ Drive Encryption or Windows Encrypting File System.

Logon Triggers

Logon triggers run stored procedures in response to a LOGON event. This event is raised when an instance of SQL Server establishes a user session. Logon triggers run after the authentication phase of logon finishes, but before the user session is actually established. Therefore, all messages originating inside the trigger that would typically reach the user, such as error messages and messages from the PRINT statement, are diverted to the SQL Server error log. Logon triggers do not run if authentication fails.

You can use logon triggers to audit and control server sessions such as tracking logon activity, restricting logons to SQL Server, or limiting the number of sessions for a specific logon.

For example, the following code shows how the logon trigger denies logon attempts to SQL Server initiated by login_test if there are already three user sessions created by that logon.

```
USE master;
GO
CREATE LOGIN login_test WITH PASSWORD = '3KHJ6dhx(0xVYsdf'
MUST_CHANGE,
    CHECK_EXPIRATION = ON;
GO
GRANT VIEW SERVER STATE TO login_test;
GO
CREATE TRIGGER connection_limit_trigger
ON ALL SERVER WITH EXECUTE AS 'login_test'
FOR LOGON
AS
BEGIN
IF ORIGINAL_LOGIN()= 'login_test' AND
```

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.


```
    (SELECT COUNT(*) FROM sys.dm_exec_sessions
            WHERE is_user_process = 1 AND
                original_login_name = 'login_test') > 3
    ROLLBACK;
END;
```

Note: The LOGON event corresponds to the AUDIT_LOGIN SQL Trace event, which can be used in event notifications. The primary difference between triggers and event notifications is that triggers are raised synchronously with events, whereas event notifications are asynchronous. For example, if you want to stop a session from being established, you must use a logon trigger. An event notification on an AUDIT_LOGIN event cannot be used for this purpose. To capture XML data about LOGON events for use inside logon triggers, use the EVENTDATA function.

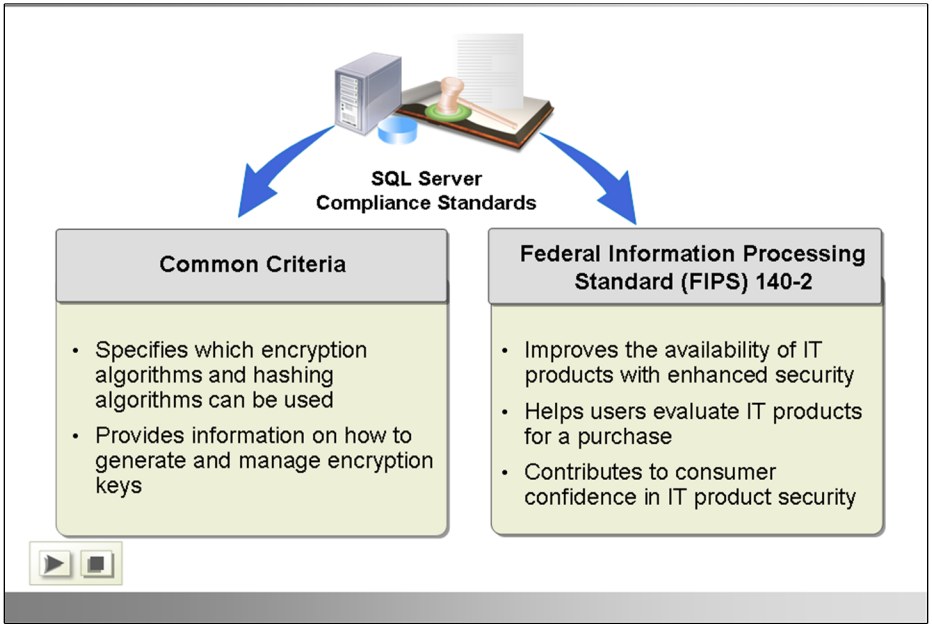
Source:

- http://msdn.microsoft.com/en-us/library/cc280386.aspx
- http://technet.microsoft.com/en-us/library/cc280663.aspx
- msdn.microsoft.com/en-us/library/cc280386.aspx
- http://msdn.microsoft.com/en-us/library/ms189799.aspx
- msdn.microsoft.com/en-us/library/bb326598.aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Meeting Regulatory Requirements for SQL Server 2008



Key Points

Microsoft SQL Server 2008 is a secure and reliable platform for protecting data and for building compliance solutions. Conforming to regulatory legislation often impacts the design of the security policies and related configuration, including what security information needs to be stored in SQL Server. SQL Server 2008 provides Common Criteria and Federal Information Processing Standard (FIPS) 140-2 regulatory requirements to handle organizational needs. You can achieve additional compliance requirements by manually configuring SQL Server Audit.

Common Criteria

The Common Criteria was rated as an international standard in 1999. It supersedes several older evaluation schemes, including the U.S. Trusted Computer Systems Evaluation Criteria (TCSEC), which specified the well-known Class C2 rating, the European Information

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Technology Security Evaluation Criteria (ITSEC), and the Canadian Trusted Computer Product Evaluation Criteria (CTCPEC). A group of nations designed Common Criteria to improve the availability of security-enhanced IT products, help users evaluate IT products for purchase, and contribute to consumer confidence in IT product security. An international body composed of more than 20 nations maintains Common Criteria, which is recognized by the International Standards Organization (ISO) as ISO standard 15408.

The following table describes the elements that are required for the Common Criteria compliance-enabled option.

Criteria	Description
Residual Information Protection (RIP)	RIP requires a memory allocation to be overwritten with a known pattern of bits before memory is reallocated to a new resource. Meeting the RIP standard can contribute to improved security. However, overwriting the memory allocation can slow performance. After the Common Criteria compliance-enabled option is activated, the overwriting occurs.
The ability to view logon statistics	After the Common Criteria compliance-enabled option is activated, logon auditing is enabled. Each time a user successfully logs onto SQL Server, information about the last successful logon time, the last unsuccessful logon time, and the number of attempts between the last successful and current logon times is made available.
Column-level GRANT permission should not override the table-level DENY permission	After the Common Criteria compliance-enabled option is activated, a table-level DENY takes precedence over a column-level GRANT. When the option is not enabled, a column-level GRANT takes precedence over a

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Criteria	Description
	table-level DENY.

The Common Criteria compliance-enabled option is an advanced option and is available on the Enterprise, Evaluation, and Developer Editions of SQL Server. SQL Server 2008 has been evaluated against the Common Criteria at Evaluation Assurance Level 1 augmented (EAL1+). SQL Server 2008 is currently in evaluation at EAL4+ and to prove compliance with the NSA DBMS Protection Profile.

In addition to activating the Common Criteria compliance-enabled option, you also must download and run a script that finishes configuring SQL Server, to comply with Common Criteria Evaluation Assurance Level 4+ (EAL4+).

The following example enables Common Criteria compliance.

```
sp_configure 'show advanced options', 1;
GO
RECONFIGURE;
GO
sp_configure 'common criteria compliance enabled', 1;
GO
RECONFIGURE
GO
```

FIPS 140-2 Compliance

The Federal Information Processing Standardization (FIPS) 140 are series of publications numbered as 140. FIPS is the U.S. government computer security standards that specify requirements for cryptography modules. As of December 2006, the current version of the standard is FIPS 140-2 that was issued on May 25, 2001.

The National Institute of Standards and Technology (NIST) issued the 140 Publication Series to coordinate the requirements and standards for cryptographic modules that include both hardware and software components for use by departments and agencies of the United States federal government. FIPS 140 does not provide sufficient conditions to guarantee that a module conforming to

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

security standards requirements is secure, still less than a system built by using such modules is secure. The requirements cover the cryptographic modules, and some aspects of the documentation in the source code at the highest security level.

User agencies desiring to implement cryptographic modules should confirm that the module they are using is covered by an existing validation certificate. FIPS 140-1 and FIPS 140-2 validation certificates specify the exact module name, hardware, software, firmware, and applet version numbers. For Levels 2 and higher, the operating platform upon which the validation is applicable is also listed. Vendors do not always maintain their baseline validations.

The Cryptographic Module Validation Program (CMVP) is operated jointly by the National Institute of Standards and Technology (NIST) Computer Security Division of the United States government and the Communications Security Establishment (CSE) of the government of Canada. The use of validated cryptographic modules is required by the United States government for all unclassified uses of cryptography. The Government of Canada also recommends the use of FIPS 140 validated cryptographic modules in unclassified applications of its departments.

Security levels

FIPS 140-2 defines four levels of security, Level 1 through Level 4. It does not specify in detail what level of security is required by any particular application.

FIPS 140-2 Level 1 provides the lowest level of security than the other three levels, and imposes very limited requirements.

FIPS 140-2 Level 2 adds requirements for physical tamper-evidence and role-based authentication.

FIPS 140-2 Level 3 adds requirements for a physical tamper-resistance and identity-based authentication, physical or logical separation between the interfaces by which critical security parameters enter and leave the module, and its other interfaces.

FIPS 140-2 Level 4 makes the physical security requirements more stringent, and requires robustness against environmental attacks.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Scope of requirements

FIPS 140 imposes requirements in the following areas:

- Cryptographic module specification
- Cryptographic module ports and interfaces
- Roles, services and authentication
- Finite state model
- Physical security
- Operational environment
- Cryptographic key management
- EMI or EMC
- Self-tests
- Design assurance
- Mitigation of other attacks

You can configure SQL Server 2008 to be compliant with FIPS 140-2. FIPS 140-2 is a statement of the "Security Requirements for Cryptographic Modules." It specifies which encryption algorithms and hashing algorithms you can use, and how you can generate and manage encryption keys. To configure SQL Server 2008 to be compliant with FIPS 140-2, it must be running on a version of Microsoft Windows that is FIPS 140-2 certified or that provides a certified cryptographic module.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

FIPS is a standard developed by two government bodies, NIST, the National Institute of Standards and Technology in the United States, and CSE, the Communications Security Establishment in Canada.

You need to configure the following requirements to ensure that SQL Server 2008 is FIPS 140-2 compliant:

- **Operating system requirement**. You must install SQL Server 2008 on a Windows Server 2003-based server or on a Windows XP-based server.
- **Windows system administration requirement**. You must set the FIPS mode before starting SQL Server.
- **SQL Server administrator requirement**. When the SQL Server service detects that the FIPS mode is enabled when you start SQL Server, SQL Server logs the following message in the SQL Server error log:
  - Service Broker transport is running in FIPS compliance mode

  You may also find that the following message is logged in the Windows Event log:
  - Database Mirroring transport is running in FIPS compliance mode

  You can verify that the server is running in the FIPS mode by these messages.

  For the security of the information exchange between services, the encryption will use the FIPS-certified instance of AES, if the FIPS mode is enabled. If the FIPS mode is disabled, the encryption will use RC4. When you configure a Service Broker endpoint in the FIPS mode, the administrator must specify "AES" for the Service Broker. If you configure endpoint to RC4, SQL Server will generate an error. Therefore, the transport layer will not start.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

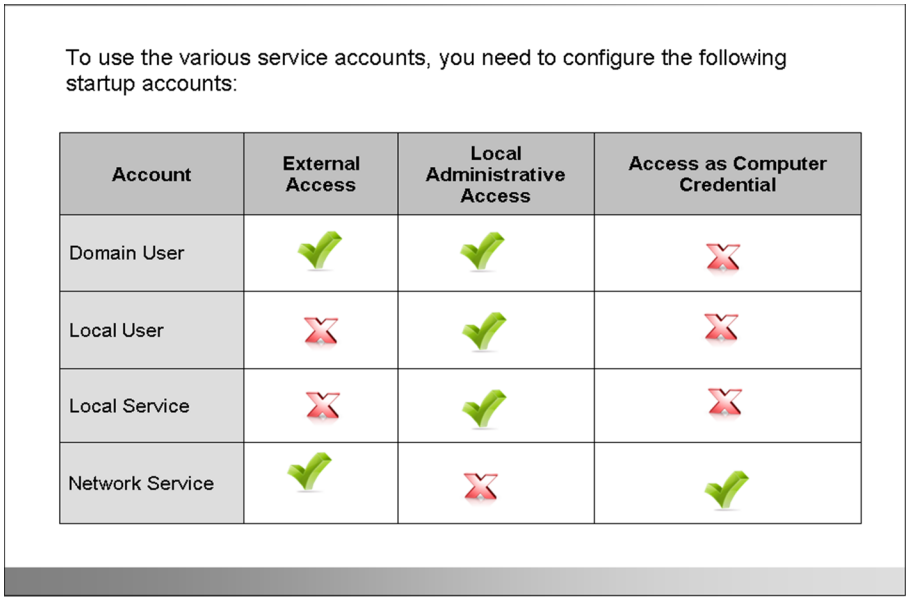
Source:

- www.microsoft.com/sqlserver/2008/en/us/compliance.aspx
- msdn.microsoft.com/en-us/library/bb153837.aspx
- msdn.microsoft.com/en-us/library/bb153837(SQL.90).aspx
- msdn.microsoft.com/en-us/library/bb326650.aspx
- msdn.microsoft.com/en-us/library/bb326611.aspx
- http://social.msdn.microsoft.com/forums/en-US/sqlservicebroker/thread/ 939d8606-ebf4-42df-acb4-8fbdf6f309f0/
- http://social.msdn.microsoft.com/forums/en-US/sqlservicebroker/thread/ 939d8606-ebf4-42df-acb4-8fbdf6f309f0/
- http://social.msdn.microsoft.com/forums/en-US/sqlservicebroker/thread/ 939d8606-ebf4-42df-acb4-8fbdf6f309f0/
- http://social.msdn.microsoft.com/forums/en-US/sqlservicebroker/thread/ 939d8606-ebf4-42df-acb4-8fbdf6f309f0/

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Considerations for Using Windows Service Accounts



Key Points

To start and run each service, you need to configure an account for that service during installation. Startup accounts used to start and run SQL Server can be built-in system accounts, local user accounts, or domain user accounts.

Domain User Account

If a service must interact with network services and access domain resources like file shares, or if it uses linked server connections to other computers running SQL Server, you may use a minimally-privileged domain account. You can perform many server-to-server activities only with a domain user account. This account type must be created by domain administration in your environment before use. The domain user account provides external access and local administrative access but cannot be used as a credential to access a computer.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Local User Account

If the computer is not part of a domain, you must assign a local user account without Windows administrator permissions. The local user account does not provide external access but provide local administrative access. This local user account also cannot be used as a credential to access a computer.

Local Service Account

The Local Service account is a built-in account that has the same level of access to resources and objects as members of the Users group. This limited access helps safeguard the system if the security of individual services or processes is compromised. Services that run as the Local Service account access network resources as a null session without credentials. The Local Service account is not supported for SQL Server or SQL Server Agent services. The actual name of the account is Windows NT® AUTHORITY\Local Service account.

Network Service Account

The Network Service account is a built-in account that has more access to resources and objects than members of the Users group. Services that run as the Network Service account access network resources by using the credentials of the computer account. The actual name of the account is NT AUTHORITY\NetworkService.

The use of the selected logon account should be limited to protecting the encryption of connection strings and passwords. You can assign the same logon account to all SQL Server instances and services, or you can configure each service account individually. However, you must use different service accounts for different instances and services.

In addition to having user accounts, every service has the following three possible startup states that users can control:

- **Disabled**. The service is installed, but not currently running.
- **Manual**. The service is installed, but will start only when another service or application needs its functionality.
- **Automatic**. The service is automatically started by the operating system.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

The following table shows optional accounts for each SQL Server service, and the startup states for each service.

SQL Server service	Optional accounts	Default startup type
SQL Server	In the SQL Server Express edition: Domain User, Local System, Network Service All other editions: Domain User, Local System, Network Service	Automatic
SQL Server Agent	Domain User, Local System, Network Service	Manual Automatic only if the user chooses to autostart
Analysis Services	Domain User, Network Service, Local Service, Local System	Automatic
Reporting Services	Domain User, Local System, Network Service, Local Service	Automatic
SSIS	Domain User, Local System, Network Service, Local Service	Automatic
Full-Text Search	Use an account different than the account for the SQL Server service. The account will change to Local Service by default on Windows Server 2008 and Windows Vista®.	Automatic
SQL Server Browser	Local Service	Disabled Automatic only if user chooses to autostart
SQL Server Active Directory®	Local System, Network Service	Disabled

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

SQL Server service	Optional accounts	Default startup type
Helper		
SQL Writer	Local System	Automatic

Note: Always run SQL Server services by using the lowest possible user rights. Use a specific low-privilege user account or domain account, instead of a shared account for SQL Server services. Do not grant additional permissions to the SQL Server service account or the service groups. Permissions will be granted through group membership or granted directly to a service SID.

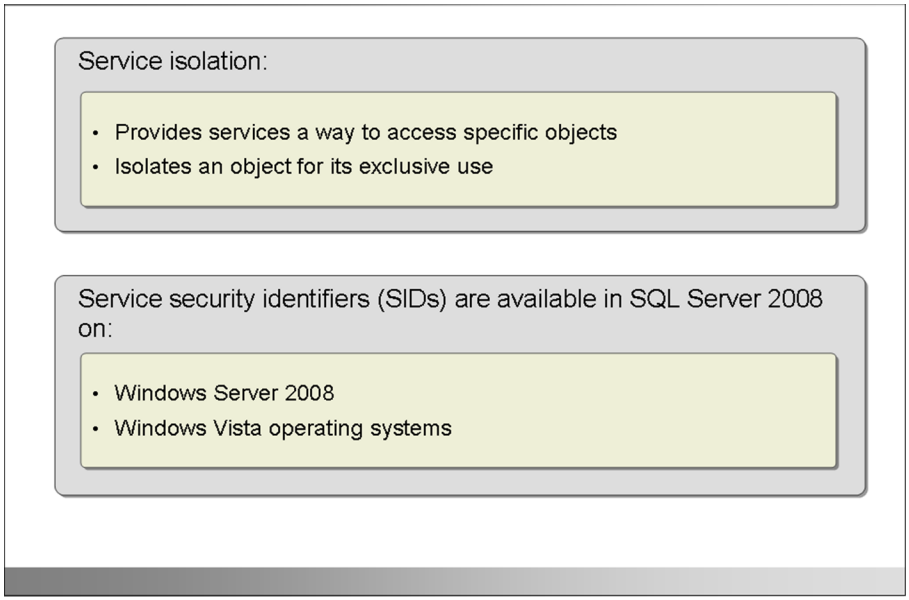
Source:

- http://msdn.microsoft.com/en-us/library/ms143504.aspx
- msdn.microsoft.com/en-us/library/ms143504.aspx
- msdn.microsoft.com/en-us/library/ms143170(SQL.90).aspx
- msdn.microsoft.com/en-us/library/ms143504.aspx
- msdn.microsoft.com/en-us/library/cc281953.aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

What Is Service Isolation?



Key Points

Service SIDs available in SQL Server 2008 on Windows Server 2003, Windows Server 2008, and Windows Vista operating systems enable service isolation. Service isolation provides services a way to access specific objects without having to either run in a high-privilege account or weaken the security protection of the object. An SQL Server service can use this identity to restrict access to its resources by other services or applications.

A service isolates an object for its exclusive use by securing the resource by using an access control entry that contains a service SID. You can derive this perservice SID from the service name and it is unique to that service. For example, a service SID name for an SQL Server service might be NT Service\MSSQL$<InstanceName>. After you enable SID for a service, you need to add the service SID to the local security group in Windows. The service owner can modify the access control list for an object to give access to the service SID. For example, a registry key in HKEY_LOCAL_MACHINE\SOFTWARE would usually be available only to services that have administrative

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

credentials. By adding the perservice SID to the access control list of the key, the service can run in a lower-privilege account, but still have access to the key.

During installation, SQL Server Setup creates a service group for each component of SQL Server. On Windows Server 2003 and Windows XP, service accounts for SQL Server services are added to the local service groups and function on these operating systems as they have in previous SQL Server releases; SQL Server services connect to an instance of SQL Server as members of the local group. On Windows Server 2008 and Windows Vista, you can add the SID to the local security group, instead of the SQL Server service account.

Supported Service Configurations

To install SQL Server 2008 on a Windows Server 2008 domain controller, you need to provide SQL Server services with a service SID. SQL Server 2008 is not supported where failover cluster nodes are domain controllers. Installing SQL Server 2008 on a Windows Server 2003 domain controller requires the domain for SQL Server services to be provisioned with a domain account.

The following table shows the supported service configurations for new and upgraded installations on Windows Server 2008 or Windows Vista.

New installation	Upgrade
• Stand-alone instance—Service group with SID • Failover cluster instance—SID • Domain Controller—SID	• Stand-alone instance—Service group with SID • Failover cluster instance—Service group with service account

The following table shows the service configurations for new and upgraded installations on Windows Server 2003 or Windows XP.

New installation	Upgrade
• Stand-alone instance—Service group with service account • Failover cluster instance—Domain service group with service account • Domain Controller—Service	• Stand-alone instance—Service group with service account • Failover cluster instance—Service group with service account

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

New installation	Upgrade
group with service account	

Question: What is the benefit of using service SIDs?

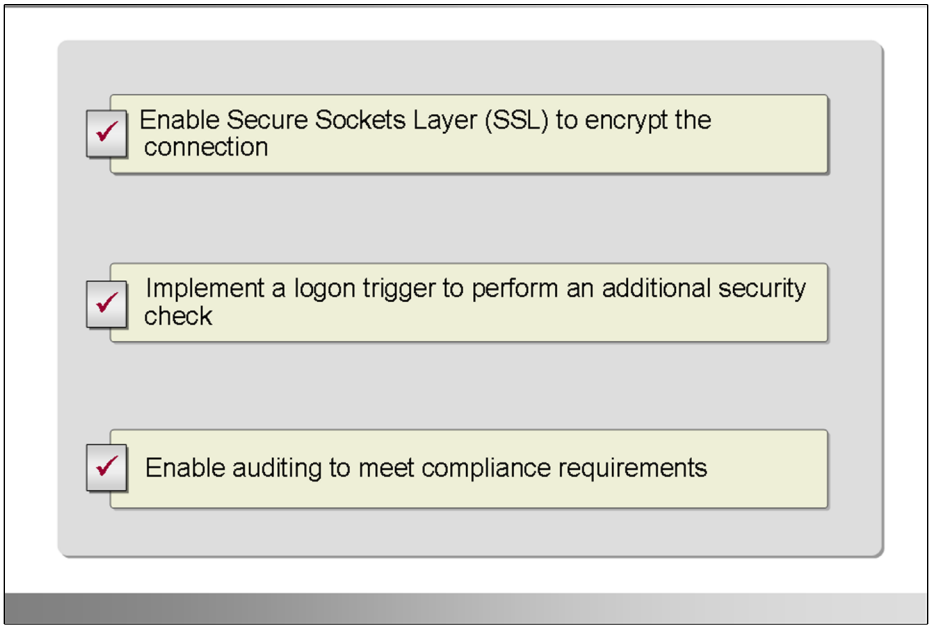
Source:

- technet.microsoft.com/en-us/library/ms143504(SQL.105).aspx
- http://technet.microsoft.com/en-us/library/ms143504%28SQL.105%29.aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Best Practices for Securing Instances



Key Points

You need to secure SQL Server 2008 instances to prevent unauthorized configuration changes and access to the server instance. Secure Sockets Layer (SSL) encrypts data that is transmitted across a network between an instance of SQL Server and a client application. Logon triggers apply additional filters to perform an additional security check on the SQL Server instance. SQL Server includes a number of features to help protect SQL Server for authorized access.

Following are the best practices for securing instances In SQL Server 2008:

- **Enable Secure Sockets Layer (SSL) to encrypt the connection**. The SSL encryption is performed within the protocol layer and is available to all SQL Server clients, except DB Library and MDAC 2.53 clients.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

You can use SSL for server validation when a client connection requests encryption. If the instance of SQL Server is running on a computer that has been assigned a certificate from a public certification authority (CA), the identity of the computer and the instance of SQL Server is guaranteed by the chain of certificates that lead to the trusted root authority. To validate the server, you must configure the computer on which the client application is running to trust the root authority of the certificate that is used by the server. Encryption with a self-signed certificate is possible, but offers only limited protection.

The level of encryption used by SSL, 40-bit or 128-bit, depends on the version of the Microsoft Windows operating system that is running on the application and database computers.

Enabling SSL encryption increases the security of data transmitted across networks between instances of SQL Server and applications. However, enabling encryption does slow the performance. When all traffic between SQL Server and a client application is encrypted by using SSL, the following additional processing is required:

- An extra network roundtrip is required at connect time.
- Packets sent from the application to the instance of SQL Server must be encrypted by the client Net-Library and decrypted by the server Net-Library.
- Packets sent from the instance of SQL Server to the application must be encrypted by the server Net-Library and decrypted by the client Net-Library.

Note: SSL connections that are encrypted by using a self-signed certificate do not provide strong security. You should not rely on SSL by using self-signed certificates in a production environment or on servers that are connected to the Internet.

- **Implement a logon trigger to perform an additional security check**. You can implement a logon trigger to perform an additional security check on the SQL Server instance logon. You can use logon triggers to audit and control server sessions by:

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

- Tracking login activity
- Restricting logons to SQL Server
- Limiting the number of sessions for a specific logon

For example, in the following code, the logon trigger denies log in attempts to SQL Server initiated by login login_test if there are already three user sessions created by that login.

```
USE master;
GO
CREATE LOGIN login_test WITH PASSWORD =
'3KHJ6dhx(0xVYsdf' MUST_CHANGE,
    CHECK_EXPIRATION = ON;
GO
GRANT VIEW SERVER STATE TO login_test;
GO
CREATE TRIGGER connection_limit_trigger
ON ALL SERVER WITH EXECUTE AS 'login_test'
FOR LOGON
AS
BEGIN
IF ORIGINAL_LOGIN()= 'login_test' AND
    (SELECT COUNT(*) FROM sys.dm_exec_sessions
            WHERE is_user_process = 1 AND
                original_login_name = 'login_test') > 3
    ROLLBACK;
END;
```

- **Enable auditing to meet compliance requirements**. You can meet security requirements by using access control features in SQL Server as follows:
- **Enable the password policy enforcement for SQL Server user account**. The enforcement of password policy can be

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

configured separately for each SQL Server login. Use ALTER LOGIN (Transact-SQL) to configure the password policy options of a SQL Server login.

The following rules apply to the configuration of password policy enforcement:

1. When CHECK_POLICY is changed to ON, the following behaviors occur:

   CHECK_EXPIRATION is also set to ON unless it is explicitly set to OFF.

   The password history is initialized with the value of the current password hash.

2. When CHECK_POLICY is changed to OFF, the following behaviors occur:

   CHECK_EXPIRATION is also set to OFF.

   The password history is cleared.

   The value of lockout_time is reset.

- **Set a strong password on the sa account**. Passwords can be the weakest link in a server security deployment. You should always take great care when you select a password.

  A strong password has the following characteristics:

  - Is at least 8 characters long.
  - Combines letters, numbers, and symbol characters within the password.
  - Is not found in a dictionary.
  - Is not the name of a command.
  - Is not the name of a person.
  - Is not the name of a user.
  - Is not the name of a computer.
  - Is changed regularly.
  - Is significantly different from previous passwords.

  Microsoft SQL Server passwords can contain up to 128 characters, including letters, symbols, and digits. Because

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

logins, user names, roles, and passwords are frequently used in Transact-SQL statements, certain symbols must be enclosed by double quotation marks (") or square brackets ([ ]). Use these delimiters in Transact-SQL statements when the SQL Server login, user, role, or password has the following characteristics:

Contains or starts with a space character.

Starts with the $ or @ character.

If used in an OLE DB or ODBC connection string, a login or password must not contain the following characters: [] {}() , ; ? * ! @. These characters are used to either initialize a connection or separate connection values.

- **Use Windows authentication, wherever possible**. During setup, you must select an authentication mode for the Database Engine. If you select Windows Authentication during setup, Setup creates the **sa** account for SQL Server Authentication but it is disabled. If you later change to Mixed Mode Authentication and you want to use the **sa** account, you must enable the account. Any Windows or SQL Server account can be configured as a system administrator. Because the **sa** account is well known and often targeted by malicious users, do not enable the **sa** account unless your application requires it. Never set a blank or weak password for the **sa** account.
- **Avoid granting the CONTROL permission**. Do not grant the CONTROL permission to users because it confers ownership capabilities on the user. A principal that has been granted the CONTROL permission can also grant permissions on the securable.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Users with CONTROL SERVER permission, such as members of the sysadmin fixed server role, can grant any permission on any securable in the server. Users with CONTROL permission on a database, such as members of the db_owner fixed database role, can grant any permission on any securable in the database. Users with CONTROL permission on a schema can grant any permission on any object within the schema.

Question: What should you consider when using SQL Server authentication?

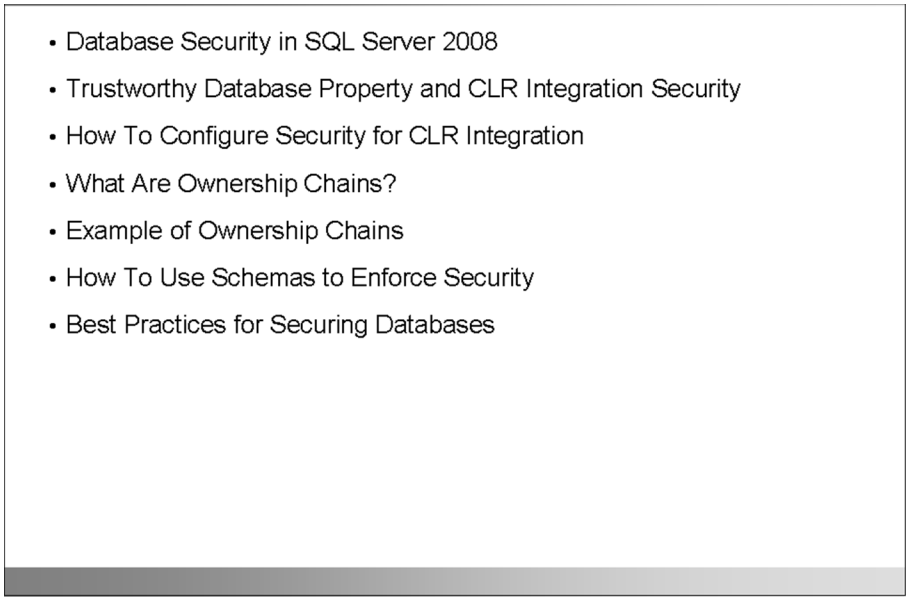
Source:

- msdn.microsoft.com/en-us/library/ms189067.aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Lesson 2
Planning Database Security



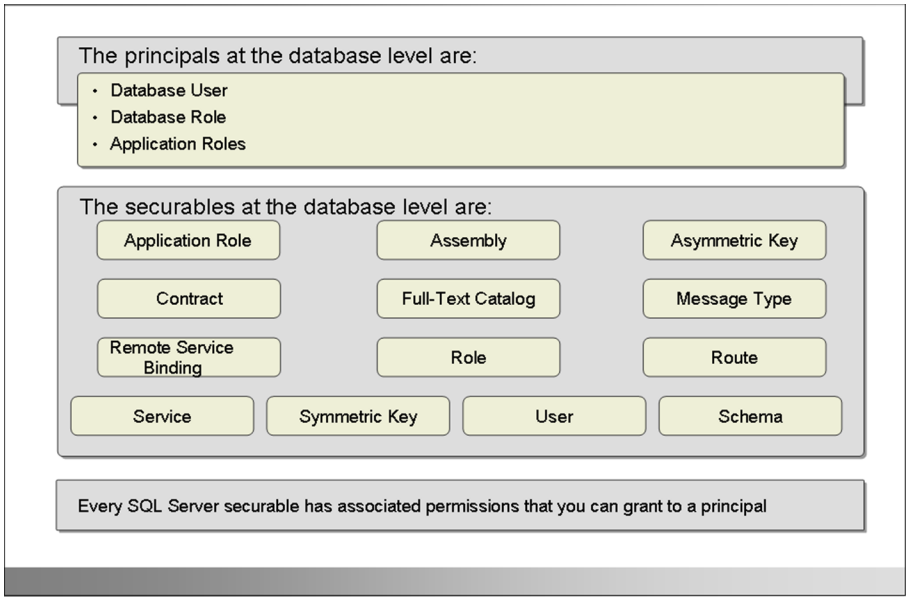
The database is a very important part of the production environment and it is very important that it be secured. Access™ control is the simplest way to control unauthorized access to the database. Sometimes applications may use stored procedure or common language runtime (CLR) components to perform complex business logic. Security for these stored procedures and CLR assemblies are a concern because the programs may require external access to SQL Server, which may become a security risk.

In this lesson, you will learn about the different security configurations at the database level and considerations for securing the database to prevent unauthorized access. You will also learn how to secure the CLR integration component and manage the security risk associated with ownership chain.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Database Security in SQL Server 2008



Key Points

You can control access to a database by assigning permission to database level principles, such as users, application roles. You can also use roles to group permission and assign to a multiple users at once.

Principals

The principals at database level are:

- **Database User**. A database user is a principal at the database level. Every database user is a member of the public role. When you create a database, the database includes a guest user. Users inherit the permissions granted to the guest user who does not have a user account in the database.

  You cannot drop the guest user, but you can disable it by revoking the CONNECT permission in any database apart from master or tempdb.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Use the following T-SQL command to create a user:

```
Use QuantamCorp
CREATE USER Ben Darton FOR LOGIN Ben Darton ;
GO
```

- **Database Role**. Every database user belongs to the public database role. When a user has not been granted or denied specific permissions on a securable, the user inherits the permissions granted to public on that securable.

  Every database includes two entities that appear as users in catalog views: INFORMATION_SCHEMA and sys. These entities are required by SQL Server. They are not principals, and they cannot be modified or dropped.

- **Application Roles**. An application role is a database principal that enables an application to run with its own, user-like permissions. You can use application roles to enable access to specific data to those users who connect through a particular application. Unlike database roles, application roles do not contain members and are inactive by default. Application roles work with both authentication modes. You can enable application roles by using sp_setapprole, which requires a password. Because application roles are a database-level principal, they can access other databases only through permissions granted in those databases to the guest. Therefore, any database in which the guest has been disabled will be inaccessible to application roles in other databases.

  In SQL Server, application roles cannot access server-level metadata because they are not associated with a server-level principal. To disable this restriction and allow application roles to access server-level metadata, set the SQL Server trace flag 4616.

  The following steps describe how an application role switches security contexts:

  1. A user runs a client application.
  2. The client application connects to an instance of SQL Server as the user.
  3. The application then runs the sp_setapprole stored procedure with a password known only to the application.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

4. If the application role name and password are valid, the application role is enabled.
5. At this point the connection loses the permissions of the user and assumes the permissions of the application role.
6. The permissions acquired through the application role remain in effect for the duration of the connection.

In SQL Server 2000, the only way for a user to reacquire its original security context after starting an application role is to disconnect and reconnect to SQL Server. Beginning with SQL Server 2005, sp_setapprole has an option that creates a cookie. The cookie contains context information before the application role is enabled. sp_unsetapprole can use the cookie to revert the session to its original context.

The following example activates the Sales11 application role with the password, fdsd896#gfdbfdkjgh700mM, and creates a cookie. The example returns the name of the current user, and then reverts to the original context by running sp_unsetapprole.

```
EXEC sp_addapprole 'SalesApp', 'Pa$$w0rd' ;

GO

SELECT USER_NAME();

GO

DECLARE @cookie varbinary(8000);

EXEC sp_setapprole 'SalesApp', 'Pa$$w0rd',
@fCreateCookie = true, @cookie = @cookie OUTPUT;

SELECT USER_NAME()

EXEC sp_unsetapprole @cookie;

GO

SELECT USER_NAME();

GO
```

The following principals are created from certificates when SQL Server is installed, and should not be deleted.

- ##MS_SQLResourceSigningCertificate##
- ##MS_SQLReplicationSigningCertificate##

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

- ##MS_SQLAuthenticatorCertificate##
- ##MS_AgentSigningCertificate##
- ##MS_PolicyEventProcessingLogin##
- ##MS_PolicySigningCertificate##
- ##MS_PolicyTsqlExecutionLogin##

Note: Server principals with names enclosed by double hash marks (##) are for internal system use only.

- **Securables**

  The following table lists the types of securables at the database level.

  Application role

  Assembly

  Asymmetric key

  Contract

  Full-Text Catalog

  Message Type

  Remote Service Binding

  Role

  Route

  Service

  Symmetric Key

  User

  Schema

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

The following table describes the permissions that you can grant to a securable.

Permission	Applies to
SELECT	• Synonyms • Tables and columns • Table-valued functions, Transact-SQL (T-SQL) and common language runtime (CLR), and columns • Views and columns
VIEW CHANGE TRACKING	• Tables • Schemas
UPDATE	• Synonyms • Tables and columns • Views and columns
REFERENCES	• Scalar and aggregate functions (T-SQL and CLR) • Service Broker queues • Tables and columns • Table-valued functions (T-SQL and CLR), and columns • Views and columns
INSERT	• Synonyms • Tables and columns • Views and columns
DELETE	• Synonyms • Tables and columns • Views and columns
EXECUTE	• Procedures (T-SQL and CLR) • Scalar and aggregate functions (T-SQL and CLR) • Synonyms • CLR types
RECEIVE	Service Broker queues

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Permission	Applies to
VIEW DEFINITION	• Procedures (T-SQL and CLR) • Service Broker queues • Scalar and aggregate functions (T-SQL and CLR) • Synonyms • Tables • Table-valued functions (T-SQL and CLR) • Views
ALTER	• Procedures (T-SQL and CLR) • Scalar and aggregate functions (T-SQL and CLR) • Service Broker queues • Tables • Table-valued functions (T-SQL and CLR) • Views
TAKE OWNERSHIP	• Procedures (T-SQL and CLR) • Scalar and aggregate functions (T-SQL and CLR) • Synonyms • Tables • Table-valued functions (T-SQL and CLR) • Views
CONTROL	• Procedures (T-SQL and CLR) • Scalar and aggregate functions (T-SQL and CLR) • Service Broker queues • Synonyms • Tables • Table-valued functions (T-SQL and CLR) • Views

Source:

- http://msdn.microsoft.com/en-us/library/ms189121.aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

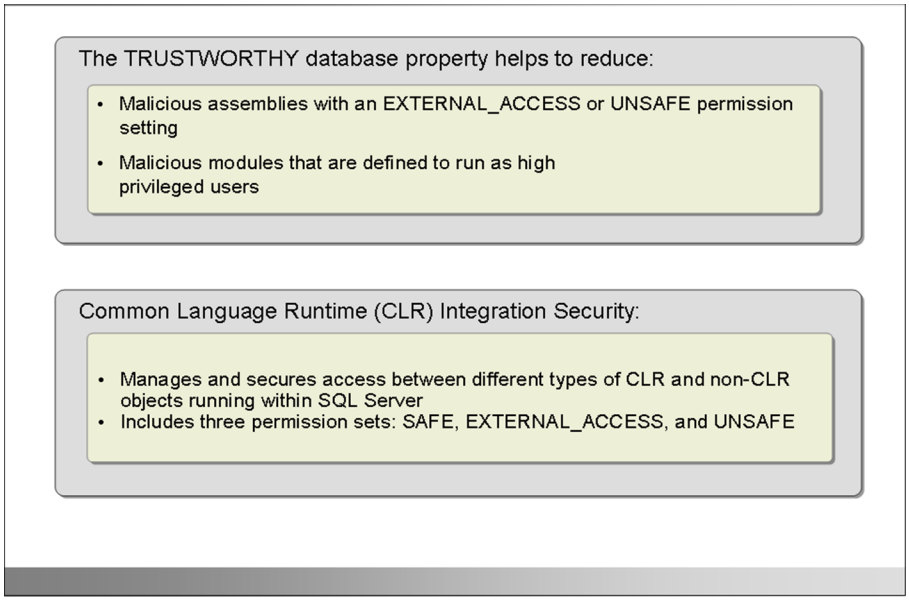
New Horizons Computer Learning Centers, Inc.

- msdn.microsoft.com/en-us/library/ms181127.aspx
- http://msdn.microsoft.com/en-us/library/ms190998.aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Trustworthy Database Property and CLR Integration Security



Key Points

You can use the TRUSTWORTHY database property to indicate whether the programs used in the database are safe to run. By default, this setting is disabled, but you can enable it by using the ALTER DATABASE statement.

Note: To enable the TRUSTWORTHY database property, you need to be a member of the sysadmin fixed server role.

You can use this property to reduce certain threats that can exist as a result of attaching a database that contains one of the following objects:

- Malicious assemblies with an EXTERNAL_ACCESS or UNSAFE permission setting
- Malicious modules that are defined to run as high privileged users

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Both these situations require a specific degree of privilege and are protected by appropriate mechanisms when they are used in the context of a database that is already attached to an instance of SQL Server. However, if the database is taken offline, users that have access to the database file can potentially attach it to an instance of SQL Server of their choice and add malicious content to the database. When databases are detached and attached in SQL Server 2008, certain permissions are set on the data and log files that restrict access to the database files.

Because an instance of SQL Server cannot immediately trust a database, SQL Server does not allow the database to access resources beyond the scope of the database until the database is explicitly marked trustworthy. Also, modules that are designed to access resources outside the database, and assemblies with either the EXTERNAL_ACCESS or UNSAFE permission setting, have additional requirements to run successfully.

The following code sets the TRUSTWORTHY database property ON for QuantamCorp.

```
USE master

GO

ALTER DATABASE QuantamCorp SET TRUSTWORTHY ON
```

The following code sets the TRUSTWORTHY database property OFF for QuantamCorp.

```
USE master

GO

ALTER DATABASE QuantamCorp SET TRUSTWORTHY OFF
```

Common Language Runtime (CLR) Integration Security

The security model of the SQL Server integration with the .NET Framework CLR manages and secures access between different types of CLR and non-CLR objects running within SQL Server. A T-SQL statement or another CLR object running on the server may call these objects. The calls between objects are referred to as links. The types of security checks performed on these objects depend on the types of links involved.

The permission set specified when creating the assembly determines the set of code access security permissions granted to assemblies by

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

the SQL Server host policy level. There are three permission sets: SAFE, EXTERNAL_ACCESS, and UNSAFE.

SAFE

This permission set allows only internal computation and local data access. SAFE is the most restrictive permission set. If you run a code on an assembly with SAFE permissions, it cannot access external system resources such as files, the network, environment variables, or the registry.

The following table lists the permissions and values of SAFE assemblies.

Permission	Value(s) / Description
SecurityPermission	Execution: Permission to run managed code
SqlClientPermission	Context connection = true, context connection = yes. Only the context-connection can be used and the connection string can only specify a value of "context connection=true" or "context connection=yes". AllowBlankPassword = false: Blank passwords are not permitted.

EXTERNAL_ACCESS

EXTERNAL_ACCESS assemblies have the same permissions as SAFE assemblies, with the additional ability to access external system resources such as files, networks, environmental variables, and the registry.

The following table lists the permissions and values of the EXTERNAL_ACCESS assemblies.

Permission	Value(s) / Description
DistributedTransactionPermission	Unrestricted: Distributed transactions are allowed.
DNSPermission	Unrestricted: Permission to request information from Domain Name Servers.
EnvironmentPermission	Unrestricted: Full access to system

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Permission	Value(s) / Description
	and user environment variables is allowed.
EventLogPermission	Administer: The following actions are allowed: creating an event source, reading existing logs, deleting event sources or logs, responding to entries, clearing an event log, listening to events, and accessing a collection of all event logs.
FileIOPermission	Unrestricted: Full access to files and folders is allowed.
KeyContainerPermission	Unrestricted: Full access to key containers is allowed.
NetworkInformationPermission	Access: Pinging is permitted.
RegistryPermission	Allows read rights to HKEY_CLASSES_ROOT, HKEY_LOCAL_MACHINE, HKEY_CURRENT_USER, HKEY_CURRENT_CONFIG, and HKEY_USERS.
SecurityPermission	You can have the following security permissions: Assertion: Ability to assert that all the callers of this code have the requisite permission for the operation. ControlPrincipal: Ability to manipulate the principal object. Execution: Permission to run managed code. SerializationFormatter: Ability to provide serialization services.
SmtpPermission	Access: Outbound connections to SMTP host port 25 are allowed.
SocketPermission	Connect: Outbound connections

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Permission	Value(s)/Description
	(all ports, all protocols) on a transport address are allowed.
SqlClientPermission	Unrestricted: Full access to the datasource is allowed.
StorePermission	Unrestricted: Full access to X.509 certificate stores is allowed.
WebPermission	Connect: Outbound connections to Web resources are allowed.

UNSAFE

UNSAFE allows assemblies unrestricted access to resources, both within and outside SQL Server. Code running from within an UNSAFE assembly can also call unmanaged code. UNSAFE assemblies are given FullTrust.

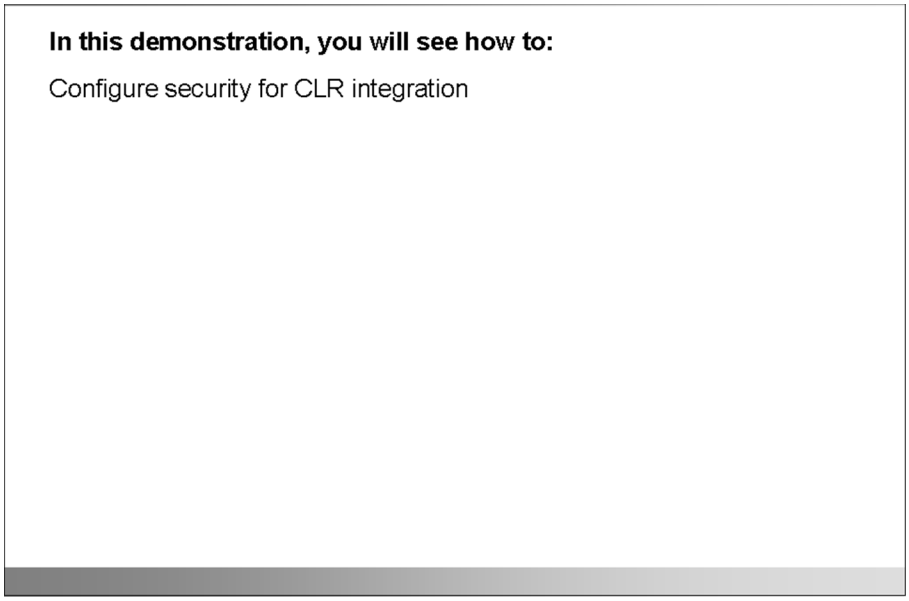
Source:

- msdn.microsoft.com/en-us/library/ms187861.aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Demonstration: How To Configure Security for CLR Integration



Note: Before performing this demonstration, you need to download Mod07_Samples.iso file from the link, http://202.144.75.100/microsoft/SQL_COURSES/50400A/Projects.iso, to the Virtual Machine CD Drive, and then copy the Samples folder from CD Drive to C:\root folder.

Key Points

The steps to configure security for CLR integration are:

1. In the Microsoft SQL Server Management Studio window, in the Query Editor pane, delete the existing code and type the following T-SQL code, and then click **Execute**.

```
USE master
GO
ALTER DATABASE QuantamCorp SET TRUSTWORTHY ON
GO
```

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.


```
sp_configure 'show advanced options', 1;
GO
RECONFIGURE;
GO
sp_configure 'clr enabled', 1;
GO
RECONFIGURE;
GO
```

2. In the Microsoft SQL Server Management Studio window, in the Query Editor pane, delete the existing code, and type the following T-SQL code, and then click **Execute**.

```
USE QuantamCorp
GO
CREATE ASSEMBLY HelloWorld
FROM
'C:\Samples\Engine\Programmability\CLR\HelloWorld\CS\He
   lloWorld\bin\debug\HelloWorld.dll'
WITH PERMISSION_SET = SAFE;
GO
CREATE PROCEDURE HelloWorld
(
   @Greeting nvarchar(12) OUTPUT
)
AS EXTERNAL NAME
   HelloWorld.[Microsoft.Samples.SqlServer.Hello].Hello
   World;
GO
```

3. In the Microsoft SQL Server Management Studio window, in the Query Editor pane, delete the existing code, and type the following T-SQL code, and then click **Execute**.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.


```
USE QuantamCorp
GO
DECLARE @Greeting nvarchar(12)
EXECUTE HelloWorld @Greeting OUTPUT
SELECT @Greeting AS [OutputParameter]
```

Question: What is CLR integration in SQL Server 2008?

Question: What is the advantage of using different security settings in CLR integration?

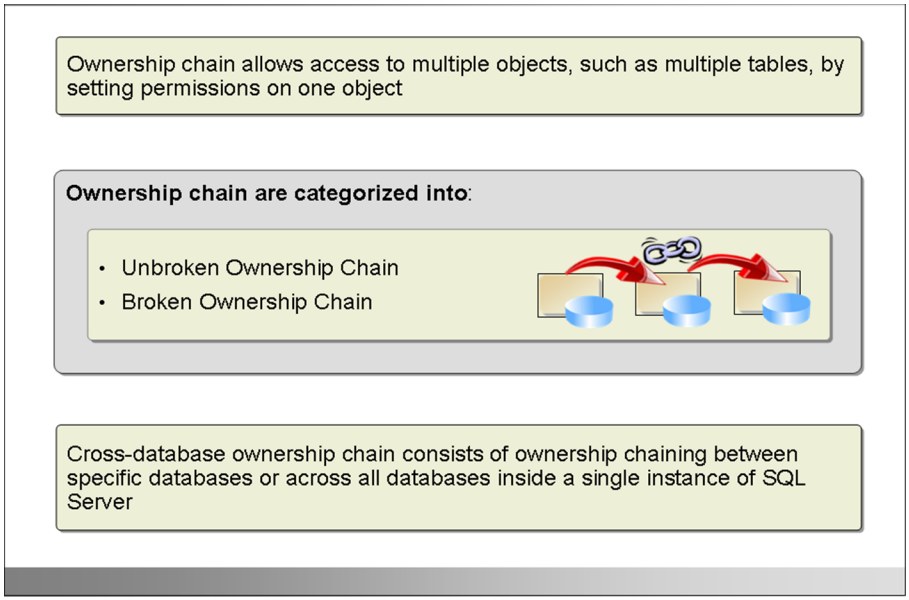
Source:

- msdn.microsoft.com/en-us/library/ms189631.aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

What Are Ownership Chains?



Key Points

When multiple database objects access each other sequentially, the sequence is known as a chain. Although such chains do not exist independently, when SQL Server traverses the links in a chain, SQL Server evaluates permissions on the constituent objects differently than it would if it were accessing the objects separately. These differences have important implications for managing security.

Ownership Chaining

Ownership chaining enables managing access to multiple objects, such as multiple tables, by setting permissions on one object, such as a view.

You should use the following guidelines when working with ownership chains:

- Avoid using cross-database ownership chain.
- Group-related objects under the same schema to avoid overhead caused by additional security checks.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

When an object is accessed through a chain, SQL Server first compares the owner of the object to the owner of the calling object. This is the previous link in the chain. If both objects have the same owner, SQL Server does not evaluate permissions on the referenced object.

Often, the owner of a source object such as a view, stored procedure, or user-defined function also owns all target objects or the objects it depends on. When one user owns the source object and all target objects, the ownership chain is said to be unbroken. If different users own the target objects, the ownership chain is broken. SQL Server relies on the state of the ownership chain to determine when to check permissions.

Ownership chains affect control over permissions, based on which they are categorized as follows:

Unbroken Ownership Chain. When an ownership chain is unbroken, SQL Server checks permission on the source object only. This enables a user to grant permissions on views, stored procedures, and user-defined functions, instead of on each database object. This can greatly simplify permissions management.

Broken Ownership Chain. If the ownership chain is broken, SQL Server checks permissions on each branch of the chain that is owned by a different user. Only those statements where the user has the necessary permissions will be run, and the remaining statements will get an "Insufficient Permissions" error. In this way, SQL Server allows object owners to retain control over permissions.

Cross-Database Ownership Chaining

Cross-database ownership chaining occurs when a source object depends on objects in other databases. Cross-database ownership chaining works in the same way as ownership chaining within a database works. The only distinction is that an unbroken ownership chain is based on all objects being owned by the same logon account, not the same database user name. This is because one logon account can have different user names in different databases.

In a cross-database ownership chain, if the source object in the source database and the target objects in the target databases are owned by the same logon account, SQL Server does not check permissions on the target objects.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Question: What is the benefit of using ownership chains?

Question: What is the risk of enabling cross-database ownership chain?

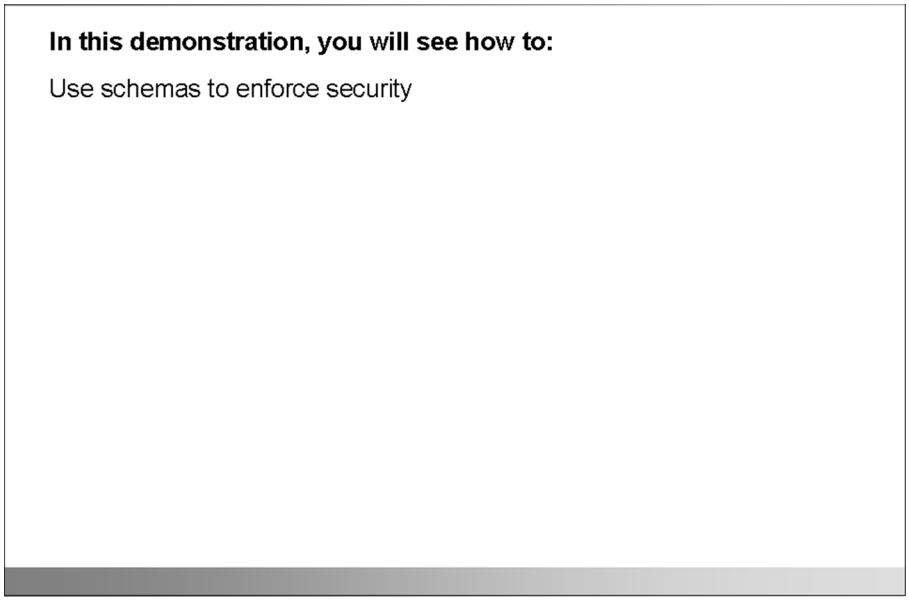
Source:

- http://msdn.microsoft.com/en-us/library/ms188676.aspx
- technet.microsoft.com/en-us/library/cc966453.aspx
- msdn.microsoft.com/en-us/library/aa905173(SQL.80).aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Demonstration: How To Use Schemas to Enforce Security



Key Points

The steps to enforce security by using schemas are:

1. In Windows Explorer, browse to the **D:\Demofiles\Mod07** folder, and then double-click the **M7_Demo2.sql** file.
2. In the Microsoft SQL Server Management Studio window, in the Query Editor pane, select the **Create a schema** code, and then click **Execute**.
3. In the Microsoft SQL Server Management Studio window, in the Query Editor pane, select the **Create a table and assign it to the new schema that you created** code, and then click **Execute**.
4. In the Microsoft SQL Server Management Studio window, in the Object Explorer pane, expand your SQL server instance, and expand the **QuantamCorp** database.
5. Right-click the **Security** folder, point to **New**, and then click

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

User. The **Database User - New** dialog box appears.

6. In the **Database User - New** dialog box, in the **User name** box, type **Ben Darton**, on the **General** page, in the **Login name** box, type the same user name.
7. In the **Default Schema** list, click the ellipses button, type **dbo**, click **OK**, and then click **OK**.
8. In the Microsoft SQL Server Management Studio window, in the Query Editor pane, select the **Retrieve all records from DemoSchema.DemoTable** code, and then click **Execute**.
9. In the Microsoft SQL Server Management Studio window, in the Object Explorer pane under **QuantamCorp**, expand **security** and then expand the **users** folder.
10. Right-click the user that you created and click **Properties**.
11. In the **Database User-Ben** Darton dialog box, in the **Default Schema** list, click the ellipses button, type **DemoSchema**, and then click **OK** two times.
12. In the Microsoft SQL Server Management Studio window, in the Query Editor pane, select the **Retrieve all records from DemoSchema.DemoTable** code, and then click **Execute**.

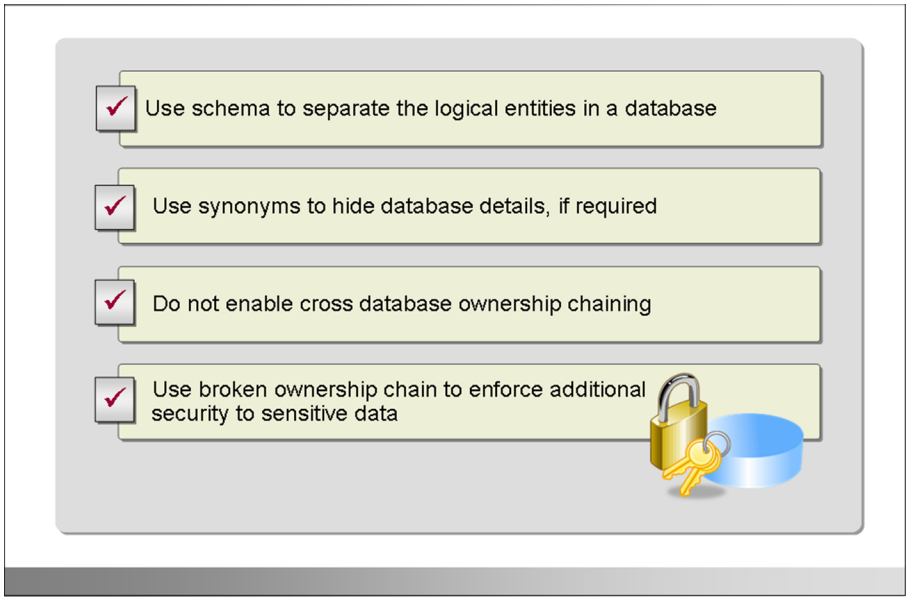
Question: How can a schema enhance the security of a database?

Question: What should you consider when using a schema?

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Best Practices for Securing Databases



Key Points

The best practices for using ownership chains include:

- **Use a schema to separate the logical entities in a database**. You can use schemas to simplify administration of security, backup and restore, and database management. You can do this by grouping database objects or entities logically. Schemas help to manage the logical entity security separately from each other, but still allow objects to work together where required. For example, in the QuantamCorp database, there are conceptually five logical entities: Human Resources, Person, Production, Purchasing, and Sales. These entities could be stored as five separate physical databases; or by using schemas, be combined as five logical entities in a physical database. This reduces the administrative complexity of managing five databases, in preference to one physical database.
- **Use synonyms to hide database details, if required**. A synonym is a database object that allows a user to create another name for some database objects. A synonym provides two advantages:

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

It simplifies the process of naming a fully qualified database object by using an alternative, usually simpler, name.

It abstracts the implementation of tables and schemas.

You can use synonyms to hide the schema name or the fully qualified object name from database consumers. By hiding the schema name or the fully qualified object, you can simplify the implementation of schemas considerably. It also allows objects to be moved between schemas without impacting the application, by using synonyms.

- **Do not enable cross-database ownership chaining**. Members of highly privileged database roles can use cross-database ownership chaining to access objects in databases external to their own. For example, if cross-database ownership chaining is enabled between database A and database B, a member of the db_owner fixed database role of either database can access data of the other database. The process is simple. Diane, a member of db_owner in database A creates user Stuart in database A. Stuart already exists as a user in database B. Diane then creates an object that is owned by Stuart in database A that calls any object owned by Stuart in database B. Because the calling and called objects have a common owner, permissions on the object in database B will not be checked when Diane accesses it through the object she has created.
- **Use broken ownership chain to enforce additional security to sensitive data**. An unbroken ownership chain is the preferred implementation for a production system. In reality, production systems may host sensitive data, so the permissions model implemented in a production system would be significantly more complex and warrant more extensive permission checking.

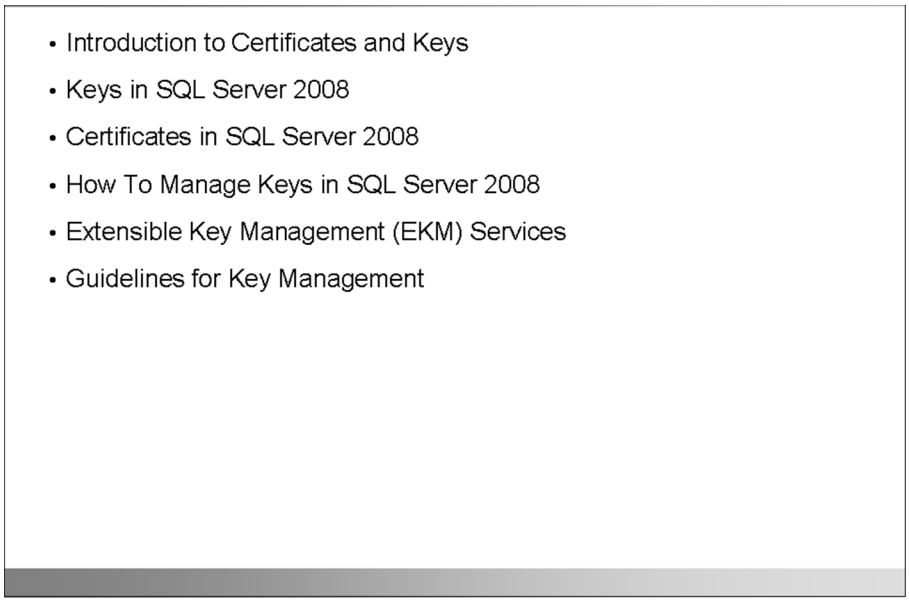
Source:

- msdn.microsoft.com/en-us/library/dd283095.aspx
- msdn.microsoft.com/en-us/library/ms188676.aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Lesson 3
Managing Certificates and Keys in SQL Server 2008

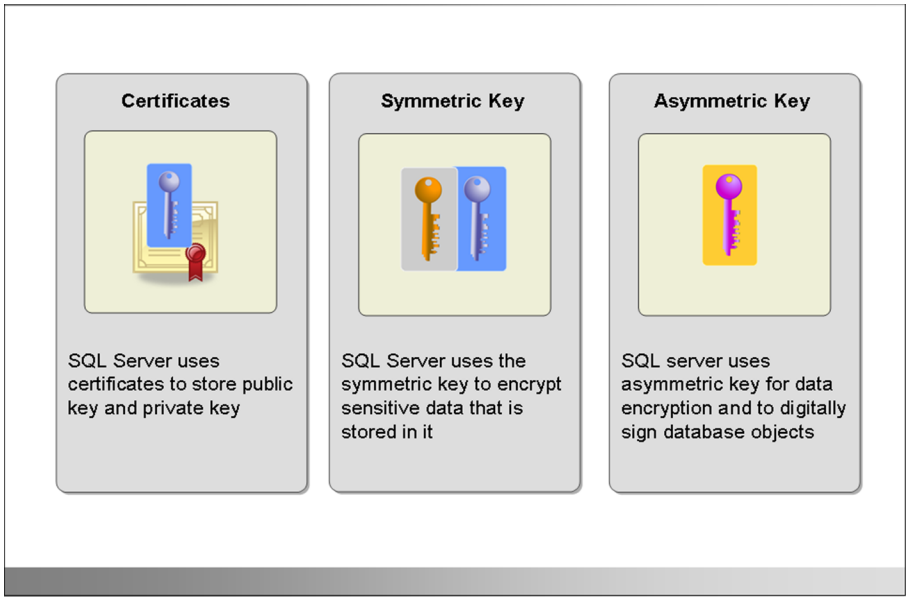


In SQL Server 2008, certificates and keys help to authenticate and safeguard information. Authentication helps to validate the credentials of applications and users. Apart from using SQL Server logon and Windows authentication logon, SQL Server also uses encryption keys to secure data. SQL Server 2008 provides different means to manage these keys in a more secure manner. Certificates in SQL Server help to enhance object and connection security.

In this lesson, you will learn about certificates and keys in SQL Server. You will also learn about the Extensible Key Management (EKM) feature of SQL Server.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

Introduction to Certificates and Keys



Key Points

SQL Server uses encryption keys to help secure data, credentials, and connection information that is stored in a server database. SQL Server has two kinds of keys: symmetric and asymmetric. Symmetric keys use the same password to encrypt and decrypt data. Asymmetric keys use one password to encrypt data called the public key and another to decrypt data called the private key.

Certificates

A certificate is a digitally signed security object that contains a public and optionally a private key for SQL Server. You can use externally generated certificates or SQL Server can generate certificates.

Note: SQL Server certificates comply with the IETF X.509v3 certificate standard.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Certificates are useful because of the option of both exporting and importing keys to X.509 certificate files. The syntax for creating certificates allows for creation options for certificates such as an expiry date. You can use certificates to help secure connections, in database mirroring, to sign packages and other objects, or to encrypt data or connections.

Symmetric Key

The symmetric key is created during SQL Server initialization when you first start the SQL Server instance. SQL Server uses the key to encrypt sensitive data that is stored on it. The operating system creates public and private keys and uses these keys to protect the symmetric key. A public and private key is created for each SQL Server instance that stores sensitive data in a database.

Asymmetric Key

Public Key Infrastructure (PKI) is a form of data security in which a user creates a public key and a private key. The users should keep the private key secret, whereas they can distribute the public key to others. Although the keys are mathematically related, the private key cannot be easily derived by using the public key. The public key is used to encrypt data and the private key is used to decrypt data. The user can decrypt a message that is encrypted by using the public key by using the correct private key. Because there are two different keys, these keys are asymmetric.

You can use asymmetric keys for securing symmetric keys. You can also use them for limited data encryption and to digitally sign database objects. An asymmetric key consists of a private key and a corresponding public key. You can import asymmetric keys from strong name key files, but you cannot export them. They also do not have expiry options. Asymmetric keys cannot encrypt connections.

Certificates and asymmetric keys are both ways to use asymmetric encryption. You can often use certificates as containers for asymmetric keys because they can contain more information such as expiry dates and issuers. Generally, you use a certificate to encrypt other types of encryption keys in a database, or to sign code modules.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Certificates and asymmetric keys can decrypt data that the other encrypts. Generally, you use asymmetric encryption to encrypt a symmetric key for storage in a database.

A public key does not have a particular format like a certificate, and you cannot export it to a file.

Question: Why is a symmetric key preferred over an asymmetric key to encrypt data directly?

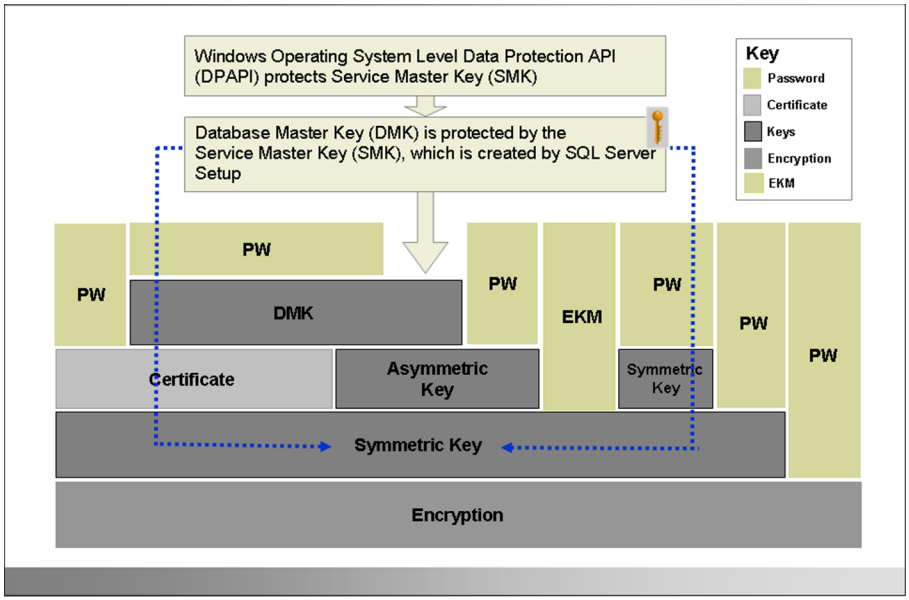
Source:

- technet.microsoft.com/en-us/library/bb964742.aspx
- msdn.microsoft.com/en-us/library/bb895327.aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Keys in SQL Server 2008



Key Points

SQL Server has two primary applications for keys: Service Master Key (SMK) generated on and for an SQL Server instance and Database Master Key (DMK) used for a database.

SMK

The SMK is automatically generated the first time the SQL Server instance is started, and it is used to encrypt a linked-server password, credentials, and the DMK. You can encrypt the SMK by using the local computer key and the Windows Data Protection API (DPAPI). The DPAPI uses a key that is derived from the Windows credentials of the SQL Server service account and the computer's credentials. You can only decrypt the SMK by the service account under which it was created or by a principal that has access to the server's credentials, such as a local administrator account.

To access objects secured by the SMK, you need either the SQL Server Service account that was used to create the key or the computer account. You can change the SQL Server Service account

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

or the computer account without losing access to the key. However, if you change both, you will lose access to the SMK. If you lose access to the SMK without one of these two elements, you will be unable to decrypt data and objects encrypted by using the original key.

DMK

The DMK is a symmetric key that helps to protect the private keys of certificates and asymmetric keys that are present in the database. You can also use it to encrypt data, but it has length limitations that make it less practical for data than using a symmetric key.

When you create the DMK, the master key is encrypted by using the Triple Data Encryption Standard (DES) algorithm and a user-supplied password. To enable the automatic decryption of the master key, a copy of the key is encrypted by using the SMK. SMK is stored in both the database where it is used and in the master system database.

The copy of the DMK that is stored in the master system database is automatically updated whenever you change DMK. However, you can change this default property by using the DROP ENCRYPTION BY SERVICE MASTER KEY option of the ALTER MASTER KEY statement. The DMK that the service master key does not encrypt must be opened by using the OPEN MASTER KEY statement and a password.

To manage encryption keys, you need to create new database keys and a backup of the server and database keys. You also need to know when and how to restore, delete, or change the keys.

Access to objects and data secured with the DMK require only the password that is used to help secure the key.

To manage symmetric keys, you can use SQL Server tools to do the following:

Back up a copy of the server and database keys so that you can use them to recover a server installation, or as part of a planned migration.

Restore a previously saved key to a database. This restoring enables a new server instance to access existing data that it did not originally encrypt.

Delete the encrypted data in a database in the unlikely event that you can no longer access encrypted data.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Re-create keys and re-encrypt data in the unlikely event that the key is compromised. As a security best practice, you should re-create the keys periodically, such as every few months, to protect the server from attacks that try to decipher the keys.

Add or delete a server instance from a server scaleout deployment where multiple servers share both a single database and the key that provides reversible encryption for that database.

The following example creates a database master key for the QuantamCorp database. The key is encrypted by using the password, Pa$$w0rd.

```
USE QuantamCorp;

CREATE MASTER KEY ENCRYPTION BY PASSWORD = 'Pa$$w0rd';

GO
```

The following example removes the master key of the QuantamCorp database.

```
USE QuantamCorp;

DROP MASTER KEY;

GO
```

Question: Why should you back up the SMK?

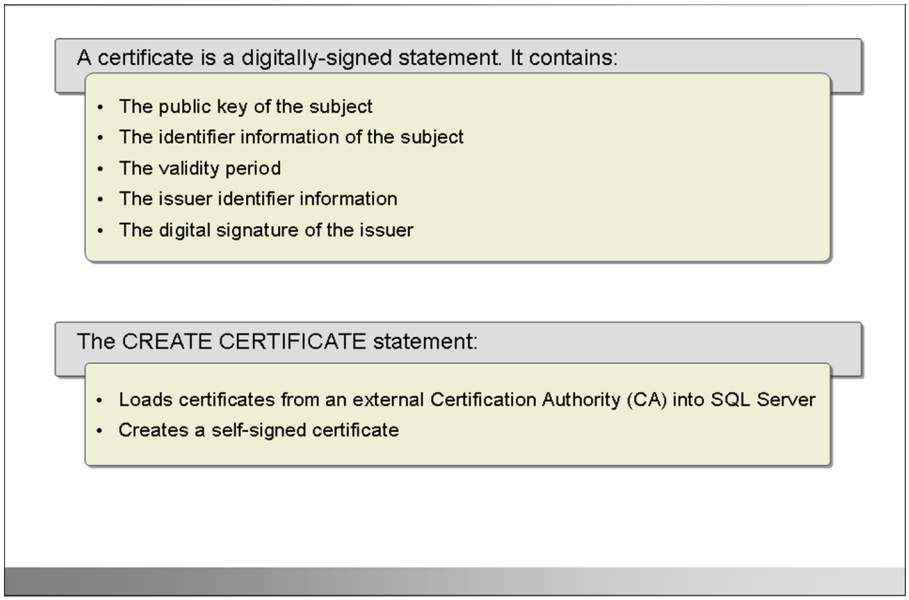
Source:

- technet.microsoft.com/en-us/library/bb964742.aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Certificates in SQL Server 2008



Key Points

A public key certificate, also known as a certificate, is a digitally-signed statement that binds the value of a public key to the identity of the person, device, or service that holds the corresponding private key. It is a database-level securable that follows the X.509 standard and supports X.509 V1 fields. Private keys generated by SQL Server are 1,024 bits long. Private keys imported from an external source have a minimum length of 384 bits and a maximum length of 3,456 bits. The length of an imported private key must be an integer multiple of 64 bits. Certificates are issued and signed by a CA. The entity that receives a certificate from a CA is the subject of that certificate.

Certificates contain the following information:

The public key of the subject.

- The identifier information of the subject, such as the name and e-mail address.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

- The validity period. This is the length of time that the certificate is considered valid. A certificate is valid only for the period of time specified within it.

The issuer identifier information.

- The digital signature of the issuer. This signature attests to the validity of the binding between the public key and the identifier information of the subject.

Using the CREATE CERTIFICATE statement

Some features in SQL Server, such as replication, may require the use of a certificate generated by an external CA. In this case, you can use the CREATE CERTIFICATE statement to load the certificate files provided by the external CA, and load the public or private key or both into SQL Server. Then, you can refer to that certificate in other features.

The CREATE CERTIFICATE statement can also generate a key pair and create a self-signed certificate. You can refer to that certificate in other features.

When you create a certificate from a container, loading the private key is optional. But when SQL Server generates a self-signed certificate, the private key is always created. By default, the private key is encrypted by using the database master key. If the database master key does not exist and if you do not specify the password, the statement will fail.

You do not require the ENCRYPTION BY PASSWORD option in the CREATE CERTIFICATE statement when the DMK encrypts the private key. You can use this option only when you need to encrypt the private key with a password. If you do not specify the password, DMK encrypts the private key of the certificate. Omitting this clause will cause an error if you cannot open the master key of the database. You do not have to specify a decryption password when the private key is encrypted with DMK.

The following example creates a certificate called Cert1. The private key of this certificate is protected by using a password.

```
USE QuantamCorp

GO

CREATE CERTIFICATE Cert1
```

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.


```
    ENCRYPTION BY PASSWORD = 'pGFD4bb925DGvbd2439587y'
    WITH SUBJECT = 'Testing Cert',
    EXPIRY_DATE = '10/31/2009';
GO
```

The following example creates a certificate in the database, loading the key pair from files.

```
USE QuantamCorp
GO
CREATE MASTER KEY ENCRYPTION BY PASSWORD = 'Pa$$w0rd';
GO
CREATE CERTIFICATE Cert_File1
    FROM FILE = 'D:\Demofiles\Cert1.cer'
    WITH PRIVATE KEY (FILE = 'D:\Demofiles\Cert1.pvk',
    DECRYPTION BY PASSWORD = 'Pa$$w0rd');
GO
```

Question: Why should you specify a password when creating a certificate?

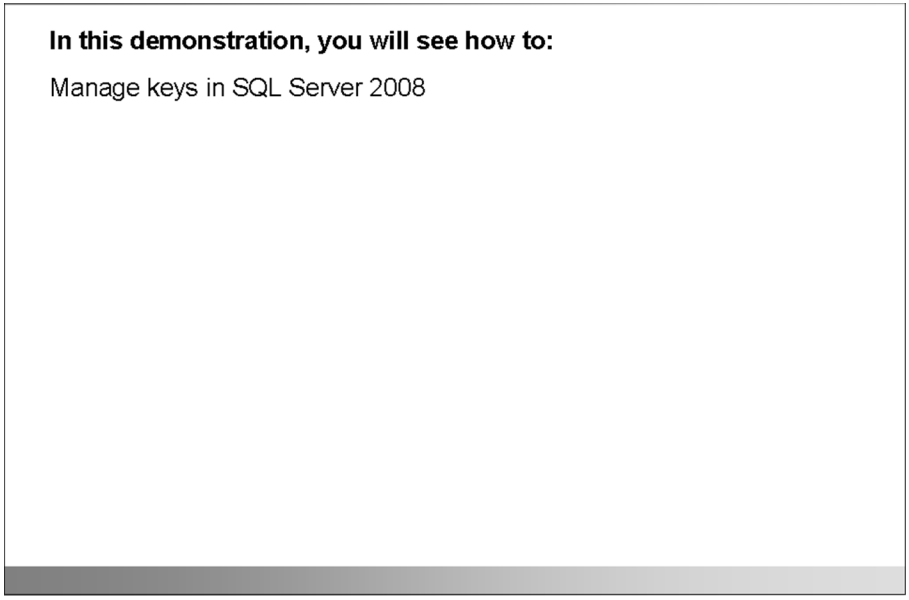
Source:

- technet.microsoft.com/en-us/library/ms189586.aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Demonstration: How To Manage Keys in SQL Server 2008



Key Points

The steps to manage keys in SQL Server 2008 are:

1. In Windows Explorer, browse to the **D:\Demofiles\Mod07** folder, and then double-click the **M7_Demo3.sql** file.
2. In the **Microsoft SQL Server Management Studio** window, click **Execute**.

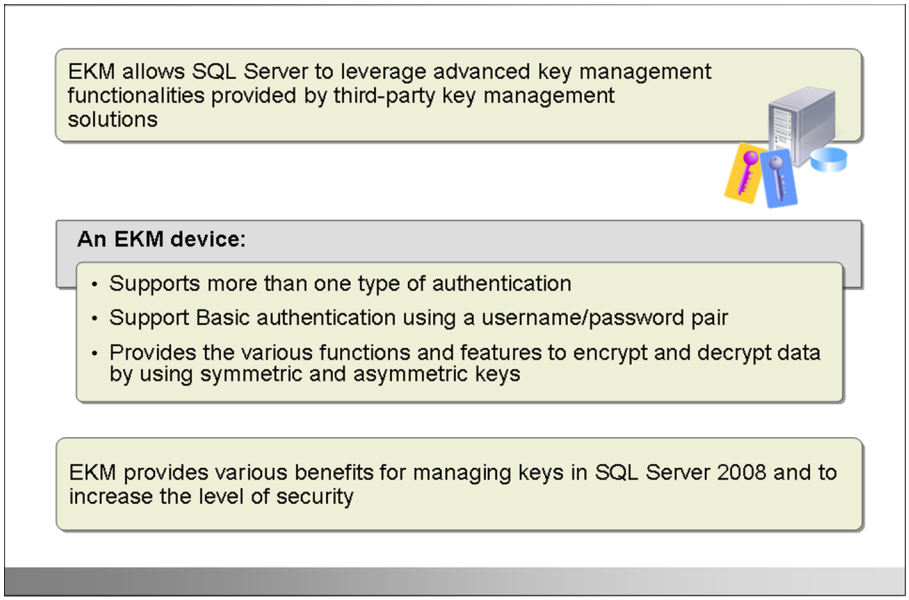
Question: What is the relationship between a DMK and a certificate?

Question: What is the purpose of having keys in SQL Server?

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

EKM Services



Key Points

SQL Server provides data encryption capabilities together with EKM by using the Microsoft Cryptographic API (MSCAPI) provider for encryption and key generation. EKM is a feature that allows SQL Server to leverage advanced key management functionalities provided by third-party key management solution. Encryption keys for data encryption and key encryption are created in transient key containers, and you must export them from a provider before you store them in the database. This approach enables key management that includes an encryption key hierarchy and key backup handled by SQL Server.

With the growing demand for regulatory compliance and concern for data privacy, organizations are taking advantage of encryption as a way to provide a "defense in depth" solution. This approach is often impractical when only database encryption management tools are used. Hardware vendors provide products that address EKM by using HSM. HSM devices store encryption keys on hardware or software

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

modules. This is a more secure solution because the encryption keys do not reside with encryption data.

A number of vendors offer HSM for both key management and encryption acceleration. HSM devices use hardware interfaces with a server process as an intermediary between an application and an HSM. Vendors also implement MSCAPI providers over their modules, which might be hardware or software. MSCAPI often offers only a subset of the functionality that HSM offers. Vendors can also provide management software for HSM, key configuration, and key access.

HSM implementations vary from vendor to vendor, and to use them with SQL Server requires a common interface. Although the MSCAPI provides this interface, it supports only a subset of the HSM features. It also has other limitations, such as the inability to natively persist symmetric keys, and a lack of session-oriented support.

SQL Server 2008 EKM enables third-party EKM and HSM vendors to register their modules in SQL Server. When registered, SQL Server users can use the encryption keys stored on EKM modules. This enables SQL Server to access the advanced encryption features these modules support, such as bulk encryption and decryption, and key management functions such as key aging and key rotation.

SQL Server 2008 EKM enables the encryption keys that protect the database files to be stored in an off-box device, such as a smart card, USB device, or EKM or HSM module. You can encrypt data by using encryption keys that only the database user has access to on the external EKM or HSM module.

Authentication with an EKM Device

An EKM module can support more than one type of authentication. Each provider exposes only one type of authentication to SQL Server, that is, if the module supports Basic or Other authentication types, it exposes one or the other, but not both.

EKM Device-Specific Basic Authentication Using username/password

For those EKM modules that support Basic authentication by using a username and password pair, SQL Server provides transparent authentication by using credentials.

You can create a credential for an EKM provider and link it to a logon for both Windows and SQL Server accounts to access an EKM module

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

on every logon. The Identify field of the credential contains the username; the secret field contains a password to connect to an EKM module.

If there is no logon linked credential for the EKM provider, you can use the credential linked to the SQL Server service account.

You can have multiple credentials linked to a logon, as long as they are used for distinctive EKM providers. You must link only one credential per EKM provider per logon. You can also link the same credential to other logons.

For EKM modules that have authentication other than Windows or user and password combinations, you must perform authentication independently from SQL Server.

Encryption and Decryption by an EKM Device

You can use the following functions and features to encrypt and decrypt data by using symmetric and asymmetric keys:

- Symmetric key encryption
- Asymmetric Key encryption
- EncryptByKey(key_guid, 'cleartext', …)
- DecryptByKey(ciphertext, …)
- EncryptByAsmKey(key_guid, 'cleartext')
- DecryptByAsmKey(ciphertext)

SQL Server can use EKM keys to encrypt other keys in a database. You can create and use both symmetric and asymmetric keys on an EKM device. You can encrypt native or non-EKM symmetric keys with EKM asymmetric keys.

The following example creates a database symmetric key and encrypts it by using a key on an EKM module.

```
CREATE SYMMETRIC KEY Key1

WITH ALGORITHM = AES_256

ENCRYPTION BY EKM_AKey1;

GO

--Open database key
```

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.


```
OPEN SYMMETRIC KEY Key1
```

Note: You cannot encrypt one EKM key with another EKM key.

You can use the keys stored in EKM for transparent data encryption (TDE) and cell-level encryption.

Benefits of Using EKM

EKM provides the following benefits:

- Additional authorization checks enabling separation of duties.
- Higher performance for hardware-based encryption/decryption.
- External encryption key generation.
- External encryption key storage enables physical separation of data and keys.
- Encryption key retrieval.
- External encryption key retention enables encryption key rotation.
- Easier encryption key recovery.
- Manageable encryption key distribution.
- Secure encryption key disposal.

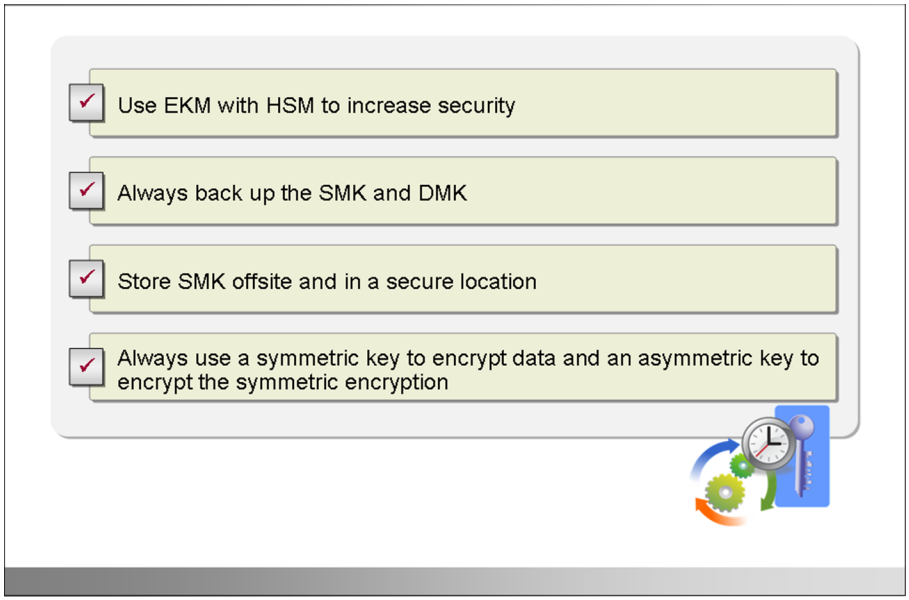
You can use EKM for a user name and password combination or other methods defined by the EKM driver.

Source:

- technet.microsoft.com/en-us/library/bb895340.aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

Guidelines for Key Management



Key Points

When employing Key Management, you need to consider the following guidelines:

- **Use EKM with HSM to increase security**. Storing keys and certificates or either in HSM increases the overall security because HSM provides an additional level of protection for the key. HSM devices use hardware interfaces with a server process as an intermediary between an application and an HSM. MSCAPI often offers only a subset of the functionality that is offered by an HSM. Vendors can also provide management software for HSM, key configuration, and key access.
- **Always back up SMK and DMK**. SMK is the root of the SQL Server encryption hierarchy. It is generated automatically the first time it is needed to encrypt another key. By default, you can encrypt the SMK by using the Windows data protection API and by using the local machine key. Only the Windows service account can open SMK under which it was created, or by a principal with

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

access to both the service account name and its password. To regenerate or restore the SMK, you need to decrypt and re-encrypt the complete encryption hierarchy. Unless the key has been compromised, you should schedule this resource-intensive operation during a period of low demand. You can use the database master key to encrypt other keys and certificates inside a database. If you delete the key or it gets corrupted, SQL Server may be unable to decrypt those keys, and the data encrypted by using these keys will be effectively lost. For this reason, you should back up the database master key, and store the backup in a secure off-site location.

- **Store SMK offsite and in a secure location**. To access objects that SMK secured, you require either the SQL Server Service account that was used to create the key or the computer account; that is, the computer is tied to the system where the key was created. You can change the SQL Server Service account or the computer account without losing access to the key. However, if you change both, you will lose access to the SMK. If you lose access to the service master key without one of these two elements, you will be unable to decrypt data and objects encrypted by using the original key. You should, therefore, store the SMK off-site to restore all the content of the database if a site disaster occurs.
- **Always use a symmetric key to encrypt data and an asymmetric key to encrypt the symmetric encryption for better performance**. Symmetric keys always encrypt data in faster speed. Therefore, to increase the performance, use symmetric keys to protect the data and asymmetric keys to protect the symmetric encryption, to avoid compromising security.

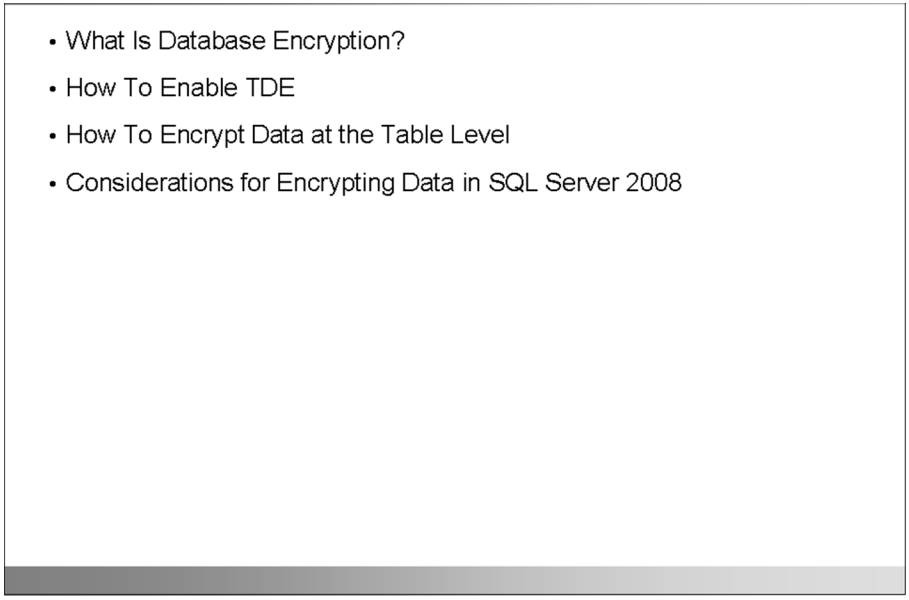
Source:

- technet.microsoft.com/en-us/library/bb895340.aspx
- technet.microsoft.com/en-us/library/bb964742.aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Lesson 4
Implementing Database Encryption



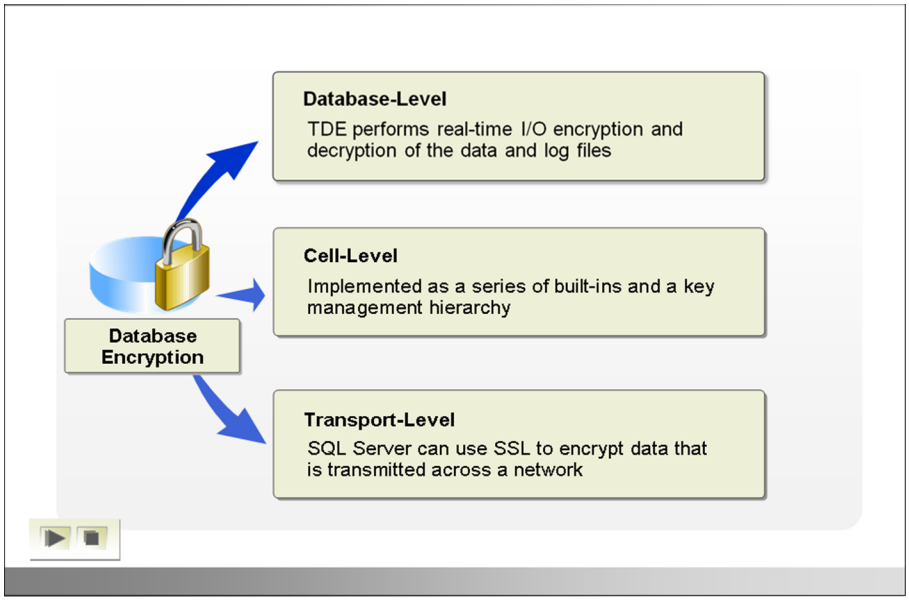
In many organizations, access control may not be able to prevent information leakages or unauthorized access to information. To provide additional security, SQL Server provides database encryption at the cell-level to protect data so that even administrators cannot access the information stored in the databases without permission. With the introduction of TDE in SQL Server 2008, users now have the choice between cell-level encryption as in SQL Server 2005, full database-level encryption by using TDE, or the file-level encryption options provided by Windows.

In this lesson, you will learn how to apply TDE to improve data security. You will also learn how to encrypt data at the table level and the considerations for encrypting data in SQL Server 2008.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

What Is Database Encryption?



Key Points

To secure the database, you can take several precautions, such as designing a secure system, encrypting confidential assets, and building a firewall around the database servers.

For example, if physical media, such as drives or backup tapes, is stolen, an unauthorized user can restore or attach the database and access the data. You can prevent this unauthorized access by encrypting sensitive data in the database and protecting the keys that are used to encrypt the data with a certificate. This restricts unauthorized access to the data.

Securing SQL Server can be viewed as a series of steps involving four areas: the platform; authentication; objects, including data; and applications that access the system.

Types of Database Encryption

In SQL Server 2008, you can perform database encryption at the database-level, cell-level, and transport-level as given below.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

- **Database-Level Encryption**. Transparent Data Encryption (TDE) performs real-time I/O encryption and decryption of the data and log files. Thus, any data that is written into the database file (*.mdf) is encrypted. You can also design snapshots and backups to take advantage of the encryption provided by the TDE so these are encrypted on disk, also. Data that is in use, however, is not encrypted because TDE does not provide protection at the memory level or transit level. The transaction log is also protected, but additional caveats apply.

  The encryption uses a database encryption key (DEK), which is stored in the database boot record for availability during recovery. The DEK is a symmetric key secured by using a certificate stored in the master database of the server or an asymmetric key protected by an EKM module. TDE protects the data and log files, and provides the ability to comply with many laws, regulations, and guidelines established in various industries. This enables software developers to encrypt data by using AES and 3DES encryption algorithms without changing existing applications.

  For data that is in use, you can decrypt all pages as they are read and stored into the buffer pool and are in clear text in memory. The operating system may page data out of memory as part of memory management. In this process, decrypted data may be written to disk. You can configure Windows and SQL Server to prevent memory from being paged to disk, although the performance cost of this can be high. Other operating system actions such as hibernation can cause memory to be written to disk. Data in transit is not protected because the information is already decrypted before it reaches this point; SSL should be enabled to protect communication between the server and any clients.

  When the database page file is written to disk, the headers are not encrypted with the rest of the data because this information is necessary for the page to be reloaded. The header contains status details such as the database compatibility level, database version, and mirroring status. You can encrypt any important data such as DEK before it is inserted into the header. The header also includes a data corruption checksum (CRC). Users can have both a checksum on the plaintext and a checksum on the encrypted text. This is not a cryptographic checksum; it detects only data corruption, to see if the data is readable; and not data integrity, to see if the data was modified. You can encrypt all other user data stored in the database

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

page, including any unused or previously deleted sections of data, to avoid information leakage. When you enable TDE on any user database, encryption is also automatically enabled for the temporary database (tempdb). This prevents temporary objects that are used by the user database from leaking to disk. System databases other than tempdb cannot currently be encrypted by using TDE.

TDE is available only in the Enterprise and Developer Editions of SQL Server 2008. Databases used in other Editions cannot be encrypted by using TDE, and TDE-encrypted databases cannot be used in other Editions.

- **Cell-level encryption**. SQL Server offers encryption at the cell level. Cell-level encryption was introduced in Microsoft SQL Server 2005. You can implement cell-level encryption as a series of built-ins and a key management hierarchy. This encryption follows a manual process that requires a redesign of the application to call the encryption and decryption functions. In addition, you must modify the schema to store the data as varbinary and then recast back to the appropriate data type when read. The traditional limitations of encryption are inherent in this method because you cannot use any of the automatic query optimization techniques.

  Cell-level encryption has a number of advantages over database-level encryption. It offers a more granular level of encryption. In addition, data is not decrypted until it is used (when a decryption built-in function is called) so that even if a page is loaded into memory, sensitive data is not in clear text. Cell-level encryption also allows for explicit key management. Keys can be assigned to users and protected by passwords to prevent automatic decryption. This offers another degree of control (users can, for example, have individual keys for their own data); however, the administrator is further burdened with maintaining the keys. Because cell-level encryption is highly configurable, it may be a good fit for applications that have targeted security requirements.

- **Transport-level encryption**. Microsoft SQL Server can use Secure Sockets Layer (SSL) to encrypt data that is transmitted across a network between an instance of SQL Server and a client application. The SSL encryption is performed within the protocol layer and is available to all SQL Server clients, except DB Library and MDAC 2.53 clients.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Encryption Algorithm

Encryption algorithms define data transformations that unauthorized users cannot easily reverse. You can choose from among several algorithms, including DES, Triple DES, TRIPLE_DES_3KEY, RC2, RC4, 128-bit RC4, DESX, 128-bit AES, 192-bit AES, and 256-bit AES.

The following general principles apply for algorithms:

- Strong encryption generally consumes more CPU resources than weak encryption.
- Long keys generally yield stronger encryption than short keys.
- Asymmetric encryption is stronger than symmetric encryption by using the same key length, but it is relatively slow.
- Block ciphers with long keys are stronger than stream ciphers.
- Long, complex passwords are stronger than short passwords.
- If you are encrypting lots of data, you should encrypt the data by using a symmetric key, and encrypt the symmetric key with an asymmetric key.
- Encrypted data cannot be compressed, but compressed data can be encrypted. If you use compression, you should compress data before encrypting it.

Question: What is the difference between cell-level encryption and database-level encryption?

Question: What is the disadvantage of using cell-level encryption?

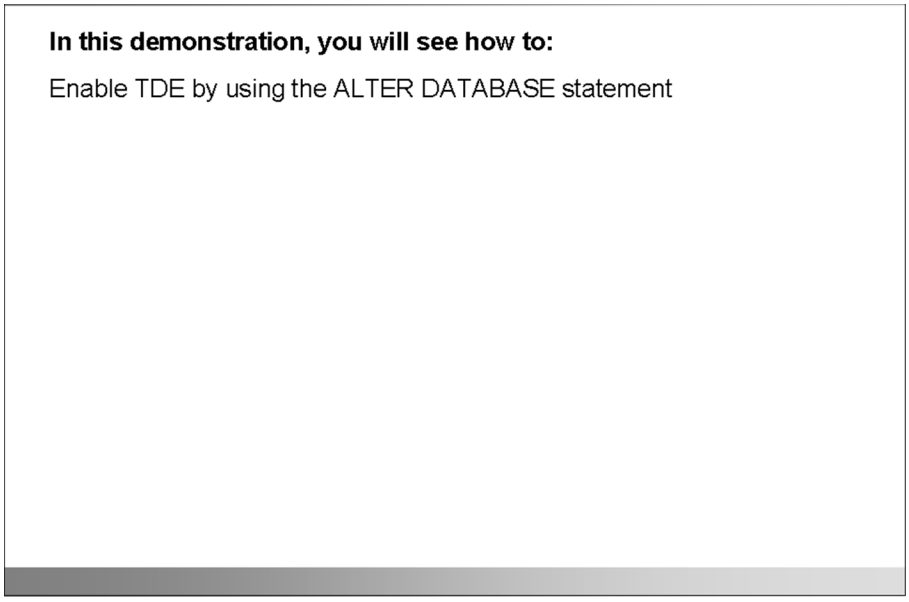
Source:

- http://technet.microsoft.com/en-us/library/bb283235.aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Demonstration: How To Enable TDE



Key Points

The steps to enable TDE with the ALTER DATABASE statement:

1. In Windows Explorer, browse to the **D:\Demofiles\Mod07** folder, and then double-click the **M7_Demo4.sql** file.
2. In the **Microsoft SQL Server Management Studio** window, click **Execute**.

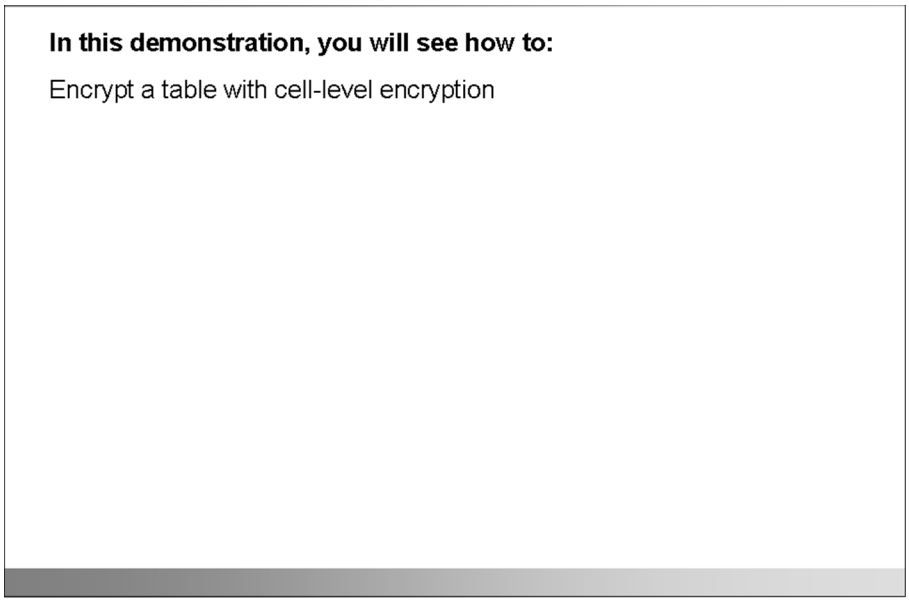
Question: What is the disadvantage of enabling database-level encryption?

Question: What is the difference between database level-encryption and transport-level encryption?

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Demonstration: How To Encrypt Data at the Table Level



Key Points

The steps to encrypt a table with cell-level encryption are:

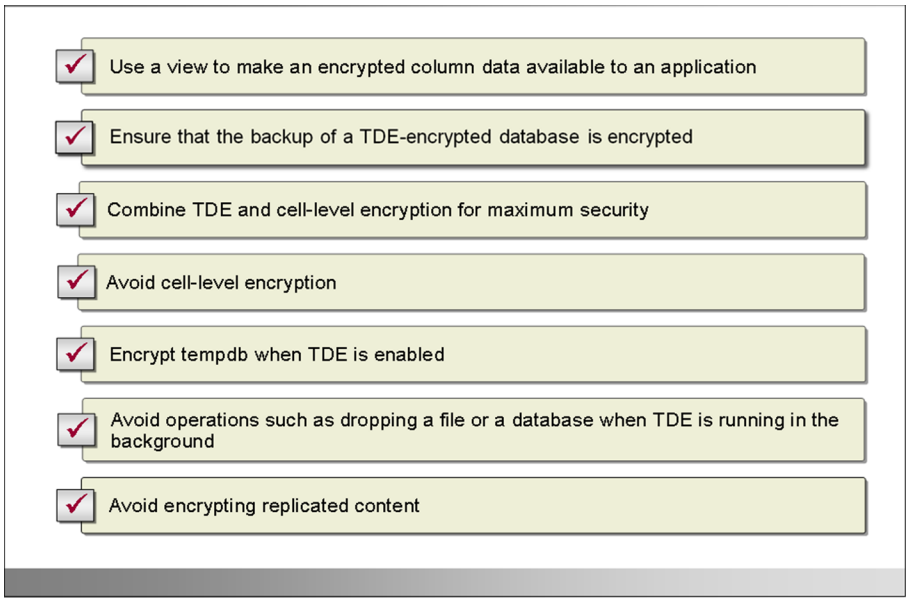
1. In Windows Explorer, browse to the **D:\Demofiles\Mod07** folder, and then double-click the **M7_Demo5.sql** file.
2. In the **Microsoft SQL Server Management Studio** window, click **Execute**.

Question: What is the difference between TDE and cell-level encryption?

Question: Which encryption keys are used in cell-level encryption?

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

Considerations for Encrypting Data in SQL Server 2008



Key Points

When encrypting data in SQL Server, consider the following.

- **Always use a view to make an encrypted column data available to an application**. If you use a view to decrypt a column, each row in the table will be decrypted before any where clause is run. You could also use a table-valued function to use the decryption only on the rows required in the result set. You need not change the application to retrieve the encrypted content stored in the underlying table.
- **Ensure that the backup of a TDE-encrypted database is also encrypted**. When you enable TDE on a database, you need to encrypt all backups. So, you must ensure that you back up and maintain the certificate that was used to protect DEK, along with the database backup. If you lose this certificate or certificates, the data will be unreadable. Back up the certificate along with the database. Each certificate backup should have two files; both of these files

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

should be archived separately from the database backup file for security. Alternatively, consider using the EKM feature for storage and maintenance of keys used for TDE.

- **Combine TDE and cell-level encryption for maximum security**. Although you can mitigate these performance concerns for cell-level encryption by explicit application design, you must take more care to prevent the accidental leakage of data. For example, consider a quick scheme to enable fast equality searches by using hashes of the sensitive data. If these hashes are stored in a column along with the encrypted data, it quickly becomes obvious if two rows have identical values because the hashes will be the same. You must use extra security reviews to ensure that unintended data disclosures do not occur; so, both the database and application must be security aware. TDE prevents these data leak scenarios by encrypting at the broadest scope. In both cell-level and database-level encryption, information is decrypted on the server, and decrypted data is sent to clients in plaintext. You must use SSL to protect this channel.
- **Avoid cell-level encryption because it might degrade performance of a system**. As a rough comparison, performance for a very basic query that selects and decrypts a single encrypted column when using cell-level encryption tends to be around 20 percent worse. This inversely scales with workload size resulting in performance degradations that are several magnitudes worse when attempting to encrypt an entire database. One sample application with 10,000 rows was four times worse with one column encrypted and 20 times worse with nine columns encrypted. Because cell-level encryption is custom to each application, performance degradation will vary depending on application and workload specifics.
- **Encrypt tempdb when TDE is enabled**. The tempdb system database will be encrypted if any other database on the instance of SQL Server is encrypted by using the TDE. This might have a performance effect for unencrypted databases on the same instance of SQL Server. Some operations are not allowed when the TDE is run in the background.
- **Avoid operations such as dropping a file or a database when TDE is running in the background**. The following operations are not allowed during initial database encryption, key change, or database decryption:

  Dropping a file from a filegroup in the database

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Dropping the database

Taking the database offline

Detaching a database

Transitioning a database or filegroup into a READ ONLY state

The following operations are not allowed during the CREATE DATABASE ENCRYPTION KEY, ALTER DATABASE ENCRYPTION KEY, DROP DATABASE ENCRYPTION KEY, or ALTER DATABASE...SET ENCRYPTION statements:

Dropping a file from a filegroup in the database

Dropping the database

Taking the database offline

Detaching a database

Transitioning a database or filegroup into a READ ONLY state

Using an ALTER DATABASE command

Starting a database or database file backup

Starting a database or database file restore

Creating a snapshot

The following operations or conditions will prevent the CREATE DATABASE ENCRYPTION KEY, ALTER DATABASE ENCRYPTION KEY, DROP DATABASE ENCRYPTION KEY, or ALTER DATABASE...SET ENCRYPTION statements:

The database is read-only or has any read-only file groups.

An ALTER DATABASE command is executing.

Any data backup is running.

The database is in an offline or a restore condition.

A snapshot is in progress.

Database maintenance tasks.

- **Avoid encrypting replicated content**. Replication does not automatically replicate data from a TDE-enabled database in an encrypted form. You must separately enable TDE if you want to protect the distribution and subscriber databases. Snapshot

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

replication, and the initial distribution of data for transactional and merge replication, can store data in unencrypted intermediate files such as the bcp files. During transactional or merge replication, encryption can be enabled to protect the communication channel.

Source:

- msdn.microsoft.com/en-us/library/cc278098.aspx
- msdn.microsoft.com/en-us/library/bb934049.aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Lab 7: Designing a Strategy for Securing SQL Server 2008

- Exercise 1: Designing a Strategy for Database Security (Discussion)
- Exercise 2: Configuring Security in SQL Server
- Exercise 3: Enabling Database Encryption

Logon Information

Virtual machine	50400A-NYC-SQL1
User name	Administrator
Password	Pa$$w0rd

Estimated time: 60 minutes

Objectives:

After completing this lab, you will be able to:

Design a strategy for database security.

Configure security in SQL Server.

Enable database encryption.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Scenario

QuantamCorp has deployed two SQL Server 2008 servers in their organization and will continue to deploy more to replace theirexisting installation of SQL Server 2005. After this, SQL Server 2008will be the primary database engine used in their organization.Recently, the company performed a technical audit. The auditsuggested that the organization had to take some measures toensure sensitive data is not accessible to unauthorized users,
including administrators. To achieve this, they have decided to reconfigure the security settings for the database and also to implement database encryption. You have been asked to help in securing the databases. You plan to perform the following in your database:

- Identify the security requirements.
- Determine the access control level.
- Determine the best SQL Server authentication to use for your organization.
- Determine whether to use schema to simplify permission administration.
- Determine whether to use cell-level encryption or TDE.
- Implement database security configuration.
- Implement schema.
- Implement cell-level database encryption.
- Implement TDE.

Exercise 1: Designing a Strategy for Database Security (Discussion)

Scenario

Your organization’s sensitive information has been accessed by unauthorized users. This has led to a serious damage to the organization’s financial and corporate image. An employee's credit card details stored on an application have been accessed by an unauthorized IT staff. The IT staff has misused the information.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net’s unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed ″as-is.″ Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

You are assigned a task to review the security configuration to ensure that only authorized users are allowed to access the organizational data. Also ensure that the sensitive information is protected even if the database file is stolen or when you directly open the database table.

Based on the above scenario, answer the following questions:

1. How can you ensure that only authorized people have access to the data stored on a SQL Server?
2. At which level should you apply access control, the database-level, the server-level, or both the levels?
3. Your organization has more than 100 tables in some of its databases. One of your colleagues has raised a concern about not wanting to grant individual permissions for each table. How would you address the concern?
4. How would you ensure that an unauthorized user cannot view any sensitive information? The user should not be able to view the data even after opening the database by using SSMS. In addition, the user should not be allowed to view the data even after mounting the database file obtained from the server.
5. How would you implement database encryption to protect the database keys?

Exercise 2: Configuring Security in SQL Server

Scenario

While configuring the database you realize that the database is not secure and users can access sensitive information. To make the database more secure, you need to enable authentication and configure the security of the database. Instead of applying security restrictions at the table level, you decide to configure the security of a group of tables by defining a schema.

The main tasks for this exercise are as follows:

1. Enable SQL authentication.
2. Create a test database.
3. Create SQL Server logon.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

4. Create a database user.
5. Create schemas.
6. Create test tables.
7. Grant a permission on the test table.
8. Test the permission settings.

Task 1: Enable SQL authentication.

1. Open the Microsoft SQL Server Management Studio window.
2. In the Connect to Server dialog box, type the following to connect to Microsoft SQL Server Management Studio.
   - Server type: **Database Engine**
   - Server name: **NYC-SQL1**
   - Authentication: **Windows Authentication**
3. Using the Microsoft SQL Server Management Studio window, open the **Server Properties – SQL2008** dialog box.
4. Set the server authentication mode to **SQL Server and Windows Authentication mode**.
5. Restart the NYC-SQL1 SQL server instance.

Task 2: Create a test database.

1. In the Query Editor pane of the Microsoft SQL Server Management Studio window, run a T-SQL script to create a test database named **Mod07** with the following physical structure.

File name	Settings
Mod07_Primary.mdf	Current size: 10 Maximum size: 50 Filegrowth: 5
Mod07_Log.ldf	Current size: 10 Maximum size: 50 Filegrowth: 5

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Note: Browse to the **D:\Labfiles\Mod07** folder and select the **E2_T2.sql** file to execute the required T-SQL code.

Task 3: Create SQL Server logon.

1. Using the Microsoft SQL Server Management Studio window, open the **Login – New** dialog box for the **Security** node.
2. Create a new SQL Server authentication logon for **peter** with the following settings:
   - Default Database: **Mod07**
   - Password: **Pa$$w0rd**
   - Ensure that the **Enforce password policy** and the **Enforce password expiration** check boxes are selected.
   - Clear the **User must change password at next login** check box.
3. Similarly, create a new SQL Server authentication logon for **michael** with the following settings:
   - Default Database: **Mod07**
   - Password: **Pa$$w0rd**
   - Ensure that the **Enforce password policy** and the **Enforce password expiration** check boxes are selected.
   - Clear the **User must change password at next login** check box.

Task 4: Create a database user.

1. Using the Microsoft SQL Server Management Studio window, add **peter** and **michael** in the list of database users for **Mod07**.

Task 5: Create schemas.

1. Using the Microsoft SQL Server Management Studio window, create two new security schema named **Calls** and **Customers** for the **dbo** owner.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Task 6: Create test tables.

1. In the Query Editor pane of the Microsoft SQL ServerManagement Studio window, run a T-SQL script to create threetest tables in the **Mod07** database as follows.

Table name	Column name	Data type of column
Calls.CallDetails	key_col	int
Customers.CustomerInfo	customer_id	int
Customers.CustomerAddress	customer_id	int

Note: Browse to the **D:\Labfiles\Mod07** folder and select the **E2_T6.sql** file to execute the required T-SQL code.

Task 7: Grant a permission on the test table.

1. In the Query Editor pane of the Microsoft SQL Server Management Studio window, run a T-SQL script to grant the SELECT permission on the **Customers** and the **Calls** schema to peter. Also, grant the SELECT permission on **Calls.CallDetails** to **michael**.

Note: Browse to the **D:\Labfiles\Mod07** folder and select the **E2_T7.sql** file to execute the required T-SQL code.

Task 8: Test the permission settings.

1. At the command prompt, connect to SQL Server by using the following credentials:
   - User name: **peter**
   - Password: **Pa$$w0rd**
2. At the sqlcmd prompt, run a command to display all records from the **Calls.CallDetails** table. Note that no error message is displayed.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

3. At the sqlcmd prompt, run a command to display all records from the **Customers.CustomerInfo** table. Note that no error message will be displayed.
4. Restart the command prompt.
5. At the command prompt, connect to SQL Server by using the following credentials:
   - User name: **michael**
   - Password: **Pa$$w0rd**
6. At the sqlcmd prompt, run a command to display all records from the **Calls.CallDetails** table. Note that no error message will be displayed.
7. At the sqlcmd prompt, run a command to display all records from the **Customers.CustomerInfo** table. Note that an error message will be displayed.

Results: After completing this exercise, you should have enabled SQL authentication, created a test database, created a SQL Server logon, created a database user, created a schema, created three test tables, granted the SELECT and the CONNECT SQL permissions to users, and tested the permission settings.

Exercise 3: Enabling Database Encryption

Scenario

QuantamCorp recently released a new application which would help eCommerce Web sites to process credit card payments. In order to comply with regulations, you need to ensure that unauthorized users do not get access to the credit card information. You decide to implement database encryption to prevent unauthorized users from accessing credit card information.

The main tasks for this exercise are as follows:

1. Create a test table.
2. Create a database master key.
3. Create a certificate.
4. Create a symmetric key.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

5. Enable encryption on the test table.
6. Verify configuration.
7. Create a server certificate.
8. Create a database encryption key.
9. Enable transparent database encryption.

Task 1: Create a test table.

1. In the Query Editor pane of the Microsoft SQL Server Management Studio window, run a T-SQL script to create a test table named **TestTable** with the following structure:
   - id: int
   - content: varchar(20)
   - create_date: datetime
2. Insert two records into the **TestTable** table as follows.

id	content	create_date
1	Test	Current date and time
2	Test	Current date and time

Note: Browse to the **D:\Labfiles\Mod07** folder and select the **E3_T1.sql** file to execute the required T-SQL code.

Task 2: Create a database master key.

1. In the Query Editor pane of the Microsoft SQL Server Management Studio window, run a T-SQL script to create a database master key, with the password, **P@ssw0rd**.

Task 3: Create a certificate.

1. In the Query Editor pane of the Microsoft SQL Server Management Studio window, run a T-SQL script to create a certificate named **Mod07_Cert** for the **Mod07** database.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Task 4: Create a symmetric key.

1. In the Query Editor pane of the Microsoft SQL ServerManagement Studio window, run a T-SQL script to create asymmetric key named **Mod07_Key**, by using the algorithm **AES_256** for the **Mod07_Cert** certificate.

Task 5: Enable encryption on the test table.

1. In the Query Editor pane of the Microsoft SQL Server Management Studio window, run a T-SQL script to enable cell-level encryption on the **content** column of the **TestTable** table and to store the encrypted content in the **encrypted_content** column.

Note: Browse to the **D:\Labfiles\Mod07** folder and select the **E3_T5.sql** file to execute the required T-SQL code.

Task 6: Verify configuration.

1. In the Query Editor pane of the Microsoft SQL Server Management Studio window, run a T-SQL script to verify the database configuration by selecting the records from the **TestTable** table, before and after decrypting the table.

Note: Browse to the **D:\Labfiles\Mod07** folder and select the **E3_T6.sql** file to execute the required T-SQL code.

Task 7: Create a server certificate

1. In the Query Editor pane of the Microsoft SQL Server Management Studio window, run a T-SQL script to create a server certificate named **Mod07_ServerCert** that uses an encryption key **P@ssw0rd**.

Note: Browse to the **D:\Labfiles\Mod07** folder and select the **E3_T7.sql** file to execute the required T-SQL code.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Task 8: Create a database encryption key.

1. In the Query Editor pane of the Microsoft SQL ServerManagement Studio window, run a T-SQL script to create adatabase encryption key by using the **AES_128** algorithm andthe **Mod07_ServerCert** certificate.

Note: Browse to the **D:\Labfiles\Mod07** folder and select the **E3_T8.sql** file to execute the required T-SQL code.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Task 9: Enable transparent database encryption.

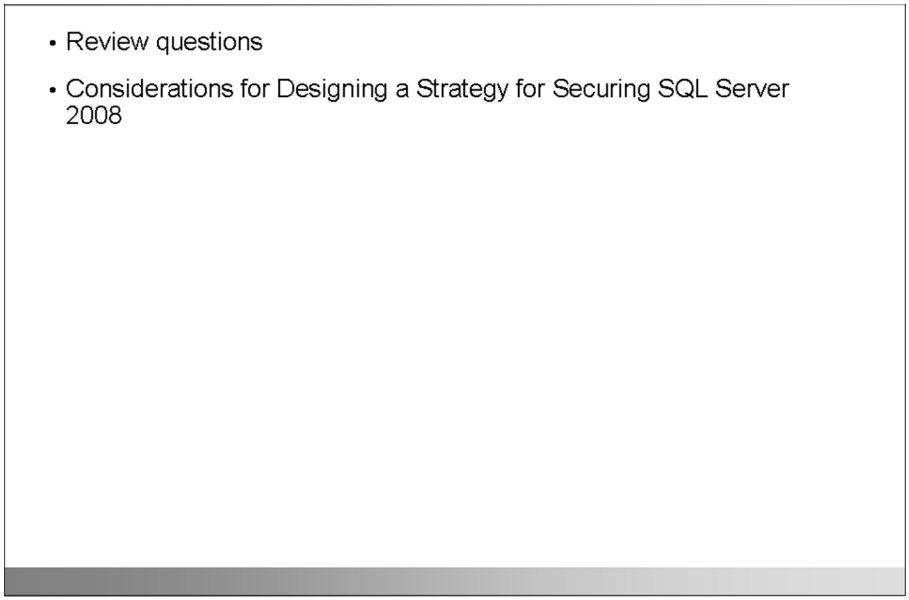
1. In the Query Editor pane of the Microsoft SQL Server Management Studio window, run a T-SQL script to enable transparent database encryption on the **Mod07** database.

Results: After completing this exercise, you should have created a test table, created a database master key, created a certificate, created a symmetric key, enabled encryption on the test table, verified the database configuration, created a server certificate, created a database encryption key, and enabled transparent database encryption.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Module Review and Takeaways



Review Questions

1. You have users who need to access various tables in your database. How can you simplify the permission assignment to a group of tables?
2. In your organization, you need to enable SQL Server authentication for an application. According to the policy of your organization, you should enable password policy for all logon accounts defined. What should you do to enable password policy check for those accounts?
3. In your organization, you have SQL Server that requires you to use a CLR stored procedure for calling a Web service to obtain information for a business process. You do not want to grant excess permission to the stored procedure. What service account should be used when you have CLR integration component which requires access to external resources?
4. Your company has recently launched a service to process credit card payments. To avoid unauthorized access, you need to

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

ensure that the credit card information stored in the database is secure. Which encryption will you use and what should you consider while encrypting the data?

5. You have implemented database encryption in one of your SQL Servers that stores credit card information. What should you consider while implementing database encryption?
6. You are working on an application and are considering whether to use TDE or cell-level encryption to encrypt data in the application. When should you consider using TDE over cell-level encryption?
7. Your corporate policy requires HSM to protect all keys used by servers for encrypting data. You recently enabled EKM and database encryption on a SQL Server. How can EKM help to protect keys?

Considerations for Designing a Strategy for Securing SQL Server 2008

When designing a strategy for securing SQL Server 2008, consider the following:

- Use server roles and database roles to simplify the administration of the permission.
- Separate database tables into different groups by using schema. These groups help to simplify assigning permissions because you can assign permission to a schema instead of assigning permissions to individual tables.
- Use Windows Authentication first as it is more secure. Consider using SQL Server Authentication only if necessary.
- Apply additional security check for each user trying to connect to SQL by using logon triggers. You should also consider using password policy on SQL Server Authentication as well.
- Run SQL Services to avoid unnecessary permissions by using a local user account.
- Apply TDE to a database that would store sensitive information, like credit card information.
- Protect data at the table level by applying cell-level encryption. Ensure that you configure keys to protect the data before enabling cell-level encryption.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

- Apply EKM for protecting keys used to encrypt and decrypt the database content.
- Apply transport-level encryption to protect data exchanged between the server and the client.
- Ensure you backup both SMK and DMK as you cannot restore a database if you do not have the correct keys.
- Configure an operator who is responsible for receiving e-mail notification from SQL Server Agent. Ensure that you have a fail safe operator defined in case you don't have an operator on-duty according to the schedule and e-mail notification should still get sent out
- Consider using secure communication between servers to secure event information transferred between SQL Servers.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Module 8

Designing a Strategy for Monitoring SQL Server 2008

Contents:

Lesson 1: Overview of Monitoring SQL Server 2008 8-3
Lesson 2: Monitoring SQL Server Instances 8-23
Lesson 3: Monitoring SQL Server Databases 8-36
Lesson 4: Monitoring Multiple SQL Server by Using Data Collection 8-74
Lab 8: Designing a Strategy for Monitoring SQL Server 2008 8-97
Module Review and Takeaways 8-106

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Module Overview

- Overview of Monitoring SQL Server 2008
- Monitoring SQL Server Instances
- Monitoring SQL Server Databases
- Monitoring Multiple SQL Servers by Using Data Collection

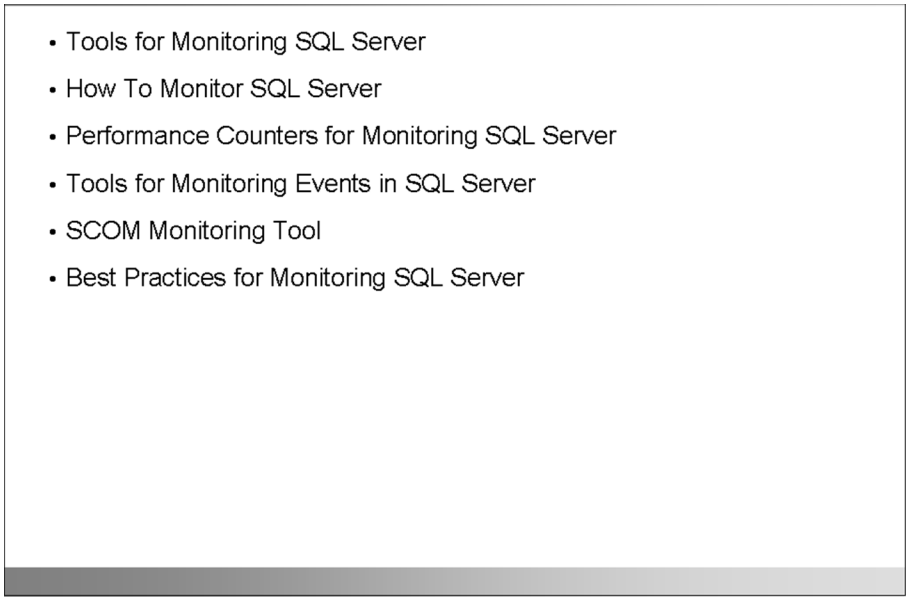
Microsoft™ SQL Server™ offers a variety of tools to monitor the database performance and the events happening on the database server. Some of the tools for monitoring include event notifications, Log File Viewer, System Monitor, SQL Server Profiler, SQL Trace, data collection, and SQL Server events.

In this module, you will learn how to use these tools to monitor SQL Server instances, databases, and multiple instances of SQL Servers.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Lesson 1
Overview of Monitoring SQL Server 2008



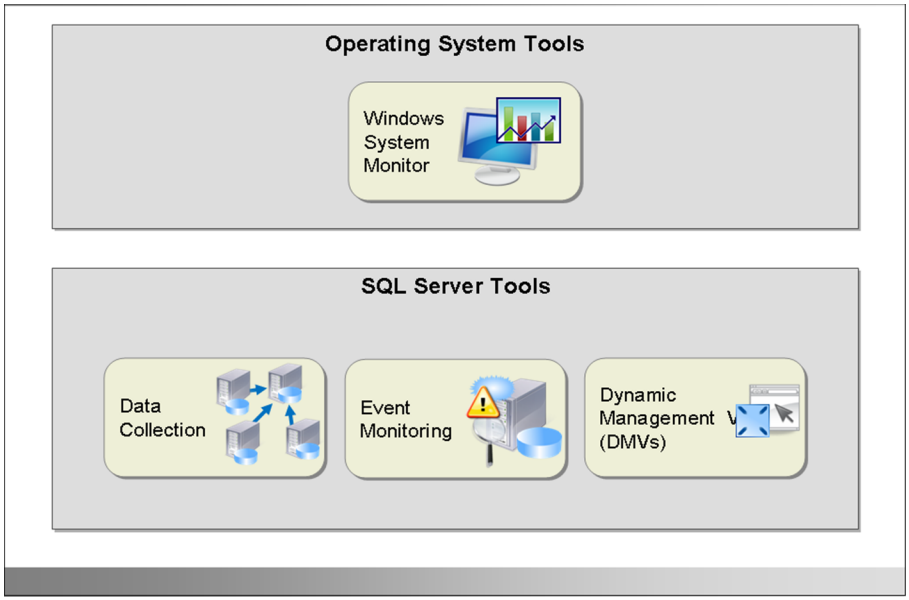
The goal of monitoring databases is to assess how a server is performing. To monitor a system effectively, you need to take periodic snapshots of the current performance to isolate processes that are causing problems, and gather data continuously over time to track performance trends. SQL Server and the Windows® operating system provide utilities that let you view the current state of the database and to track performance.

In this lesson, you will learn about the tools provided by SQL Server and the Windows operating system to monitor the database. In addition, you will learn about the performance counters for monitoring SQL Server. You will also learn about the tools to monitor events. You will learn the System Center Operations Manager (SCOM) monitoring tool for end-to-end monitoring.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Tools for Monitoring SQL Server



Key Points

You can use the tools provided by the Windows operating system and SQL Server to view the current condition of the database and to track its performance as conditions change.

Operating System Tools

The monitoring tool available with the Windows operating system is:

- **Windows System Monitor**. You can use Windows System Monitor to track resource usage in SQL Server. System Monitor uses remote procedure calls (RPCs) to collect information from SQL Server. Any user who has Microsoft Windows permissions to run the System Monitor can use it to monitor SQL Server.

SQL Server Tools

The monitoring tools available with SQL Server are:

- **Data Collection**. SQL Server 2008 provides a data collector to obtain and save data that is gathered from multiple SQL Server

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

servers. The data collector enables you to use data collection containers. These containers help you to determine the scope and frequency of data collection on a computer that is running SQL Server 2008.

- **Event Monitoring**. You can monitor events that occur in the Database Engine to maintain the performance of databases. Tools used to monitor events are Activity Monitor, Event Notification, SQL Server Log Viewer, and Windows Event Viewer.
- **Dynamic Management Views (DMVs)**. DMVs and functions return server state information that you can use to monitor the state of a server instance, diagnose problems, and optimize performance.

Question: To understand the execution of SQL scripts, which tool would you use to monitor SQL Server?

Question: Which tool would you use to monitor the real time performance of the server running SQL Server?

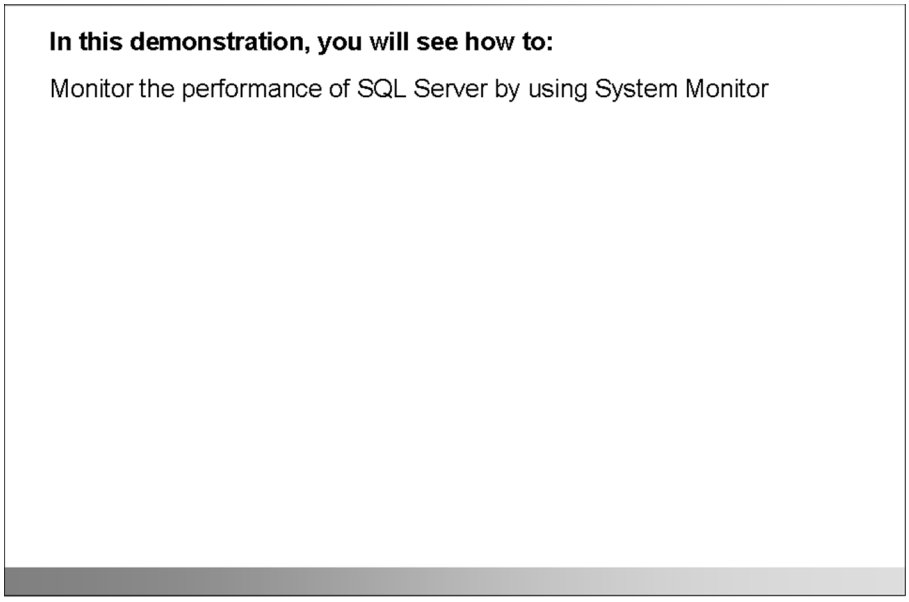
Source:

- http://technet.microsoft.com/en-us/library/bb677179.aspx
- http://msdn.microsoft.com/en-us/library/ms188754.aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Demonstration: How To Monitor SQL Server



Key Points

The steps to monitor the performance by using SQL Server are:

1. On the **Start** menu, point to **Administrative Tools**, and then click **Reliability and Performance Monitor**. The Reliability and Performance Monitor window appears.
2. In the **Reliability and Performance Monitor** window, in the navigation pane, under **Monitoring Tools**, click **Performance Monitor**.
3. In the result pane, click the **add** button. The **Add Counters** dialog box appears.
4. In the **Add Counters** dialog box, in the **Available counters** list for **<Local computer>**, expand **SQL Server:Buffer Manager**, click **Page reads/sec**, hold Ctrl, click **Page writes/sec**, and then click **Buffer cache hit ratio**. Then click **Add**.
5. In the **Available counters** list for **<Local computer>**, expand **SQL Server:Locks**, click **Lock Waits/sec**, and then click **Add**.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

6. Click **OK** to close the Add Counters dialog box.

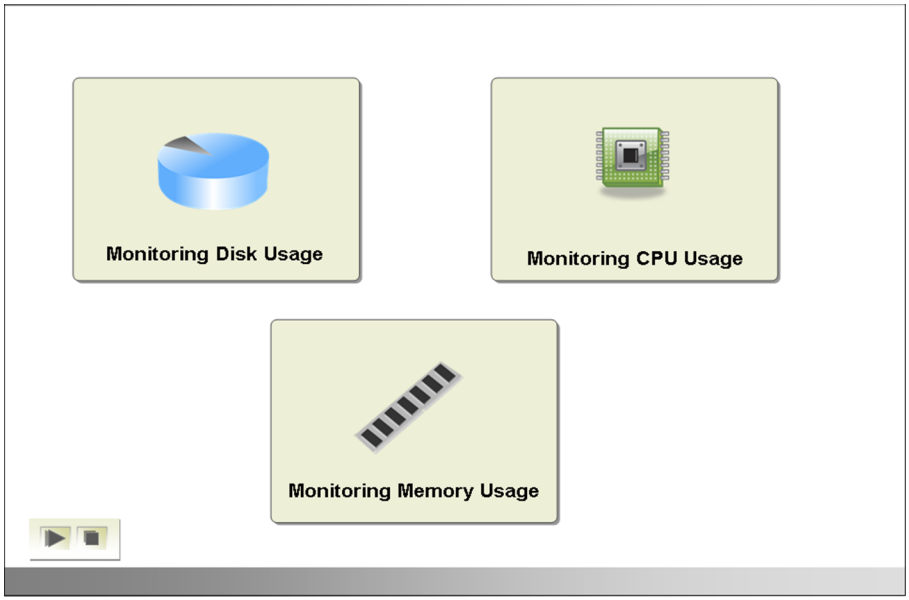
Question: Which performance counter should you monitor to determine if SQL Server is using the buffer cache?

Question: Which tools should you use to monitor the real-time performance of SQL Server and Operating System?

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Performance Counters for Monitoring SQL Server



Key Points

To monitor the performance of SQL Server 2008 effectively, you need to monitor the performance counters for disk usage, CPU usage, and memory usage.

Monitoring Disk Usage

SQL Server uses Microsoft Windows operating system input/output (I/O) calls to perform read and write operations on a disk. SQL Server manages when and how disk I/O is performed, but the Windows operating system performs the underlying I/O operations. The I/O subsystem includes the system bus, disk controller cards, disks, tape drives, CD-ROM drive, and many other I/O devices. Disk I/O is frequently the cause of bottlenecks in a system.

Monitoring disk activity involves the following two areas of focus: monitoring disk I/O and detecting excess paging isolating disk activity created by SQL Server.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Monitoring Disk I/O and Detecting Excess Paging

There are three counters that you need to monitor to determine disk activity. They are:

- **PhysicalDisk: % Disk Time. The PhysicalDisk**: % Disk Time counter monitors the percentage of time that the disk is busy with the read/write activity.
- **PhysicalDisk: Avg. Disk Queue Length**. The PhysicalDisk: Avg. Disk Queue Length counter returns the average number of requests stored in the disk queue. Disk queue is a list of pending tasks that are waiting for disk subsystem/physical disk to process.
- **PhysicalDisk: Current. Disk Queue Length**. The PhysicalDisk: Current. Disk Queue Length counter returns the current number of requests stored in the disk queue.

If the PhysicalDisk: % Disk Time counter is high or more than 90 percent, check the PhysicalDisk: Current Disk Queue Length counter to see how many system requests are waiting for disk access. You should sustain the number of waiting I/O requests at not more than 1.5 to 2 times the number of spindles that make up the physical disk. Most disks have one spindle, although redundant array of independent disk (RAID) devices usually have more than one spindle.

Use the values of the Current Disk Queue Length and % Disk Time counters to detect bottlenecks within the disk subsystem. If Current Disk Queue Length and % Disk Time counter values are consistently high, you can:

- Use a faster disk drive.
- Move some files to an additional disk or server.
- Add disks to a RAID array, if one is being used.

If you are using a RAID device, the % Disk Time counter can indicate a value greater than 100 percent. If it does, use the PhysicalDisk: Avg. Disk Queue Length counter to determine how many system requests, on average, are waiting for disk access.

Applications and systems that are I/O-bound may keep the disk constantly active.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Monitor the Memory: Page Faults/sec counter to ensure that the disk activity is not caused by paging. In Windows, paging is caused by processes configured to use too much memory and file system activity.

If you have more than one logical partition on the same hard disk, use the Logical Disk counters rather than the Physical Disk counters. You can determine which files are frequently accessed by using the logical disk counters. After you have found the disks with high-levels of read/write activity, evaluate the read-specific and write-specific counters to identify the type of disk activity that causes the load on each logical volume, such as Logical Disk: Disk Write Bytes/sec.

Isolating Disk Activity Created by SQL Server

Two counters that you need to monitor to determine the amount of I/O generated by SQL Server components are:

- **SQL Server:Buffer Manager:Page reads/sec**. This counter reports the number of physical database page reads that are issued per second. This statistic displays the total number of physical page reads across all databases. However, to reduce the cost of physical I/O, you can either use a larger data cache, intelligent indexes, and more efficient queries, or by change the database design.
- **SQL Server:Buffer Manager:Page writes/sec**. This counter reports the number of physical database page writes issued per second.

In System Monitor, these counters monitor the amount of I/O generated by SQL Server components by examining the following performance areas:

- Writing pages to disk
- Reading pages from disk

If the values for these counters reach the capacity limit of the hardware I/O subsystem, try to reduce the values by tuning your application or database to reduce I/O operations, such as index coverage, better indexes, or normalization, increasing the I/O capacity of the hardware, or adding memory. For example, you can use Database Engine Tuning Advisor (DTA) to analyze SQL Server workloads for indexes, indexed views, and partitioning to improve server performance.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Monitoring CPU Usage

Monitor an instance of SQL Server periodically to determine whether CPU usage rates are within normal ranges. A continually high rate of CPU usage may indicate the need to upgrade the CPU or add multiple processors. Alternatively, a high CPU usage rate may indicate a poorly tuned or designed application. Optimizing the application can lower CPU utilization.

An efficient way to determine CPU usage is to use the Processor: % Processor Time counter in System Monitor. This counter monitors the amount of time the CPU spends executing a thread that is not idle. A consistent state of 80 percent to 90 percent may indicate the need to upgrade the CPU or add more processors. For multiprocessor systems, monitor a separate instance of this counter for each processor. This value represents the sum of processor time on a specific processor. To determine the average for all processors, use the System: %Total Processor Time counter.

Optionally, you can also use the following counters to monitor processor usage:

- **Processor: % Privileged Time**. This counter corresponds to the percentage of time the processor spends on execution of Microsoft Windows kernel commands, such as processing of SQL Server I/O requests. If this counter is consistently high when the Physical Disk counters are high, consider installing a faster or more efficient disk subsystem.

Note: Different disk controllers and drivers use different amounts of kernel processing time. Efficient controllers and drivers use less privileged time, leaving more processing time available for user applications, increasing overall throughput.

- **Processor: %User Time**. This counter corresponds to the percentage of time that the processor spends on executing user processes such as SQL Server.
- **System: Processor Queue Length**. This counter corresponds to the number of threads waiting for processor time. A processor bottleneck develops when threads of a process require more processor cycles than that are available. If more than a few processes attempt to utilize the time of processors time, you

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

might need to install a faster processor. If you have a multiprocessor system, you could add a processor.

When you examine processor usage, consider the type of work that the instance of SQL Server performs. If SQL Server performs many calculations, such as queries involving aggregates or memory-bound queries that require no disk I/O, 100 percent of the time of a processor time can be used. If this utilization affects the performance of other applications, try changing the workload. For example, dedicate the computer to running the instance of SQL Server.

Usage rates around 100 percent, where many client requests are being processed, may indicate that processes are queuing up, waiting for processor time, and causing a bottleneck. Resolve the problem by adding faster processors.

Monitoring Memory Usage

Monitor an instance of SQL Server periodically to confirm that the memory usage is within typical ranges.

To monitor for a low-memory condition, use the following performance counters:

- **Memory: Available Bytes**. This counter indicates the number of bytes of memory that is currently available for use by processes.
- **Memory: Pages/sec**. This counter indicates the number of pages that either were retrieved from disk due to hard page faults or written to disk to free space in the working set due to page faults.

Low values for the *Memory:* Available Bytes counter can indicate that there is an overall shortage of memory on the computer or that an application is not releasing memory. A high rate for the Pages/sec counter could indicate excessive paging. You need to monitor the Memory: Page Faults/sec counter to ensure that the disk activity is not caused by paging.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

A low rate of paging and hence page faults are typical, even if the computer has plenty of available memory. The Microsoft Windows Virtual Memory Manager (VMM) takes pages from SQL Server and other processes as it trims the working-set sizes of those processes. This VMM activity tends to cause page faults. To determine whether SQL Server or another process is the cause of excessive paging, monitor the Process: Page Faults/sec counter for the SQL Server process instance.

By default, SQL Server changes its memory requirements dynamically, on the basis of available system resources. If SQL Server needs more memory, it queries the operating system to determine whether free physical memory is available and uses the available memory. If SQL Server does not need the memory currently allocated to it, it releases the memory to the operating system. However, you can override the option to dynamically use memory by using the min server memory, and max server memory server configuration options.

To monitor the amount of memory that SQL Server uses, examine the following performance counters:

- **Process: Working Set**. The Working Set counter shows the amount of memory that is used by a process. If this number is consistently below the amount of memory that is set by the min server memory and max server memory server options, SQL Server is configured to use too much memory.
- **SQL Server: Buffer Manager: Buffer Cache Hit Ratio**. The Buffer Cache Hit Ratio counter is specific to an application. However, a rate of 90 percent or higher is desirable. Add more memory until the value is consistently greater than 90 percent. A value greater than 90 percent indicates that more than 90 percent of all requests for data were served from the data cache.
- **SQL Server: Buffer Manager: Total Pages**. The SQL Server: Buffer Manager: Total Pages counter returns the number of pages in the buffer pool, including database pages, free pages, and stolen pages.
- **SQL Server: Memory Manager: Total Server Memory (KB)**. If the Total Server Memory (KB) counter is consistently high compared to the amount of physical memory in the computer, it may indicate that more memory is required.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Question: Which performance object would you use to monitor the memory usage of your SQL Server process?

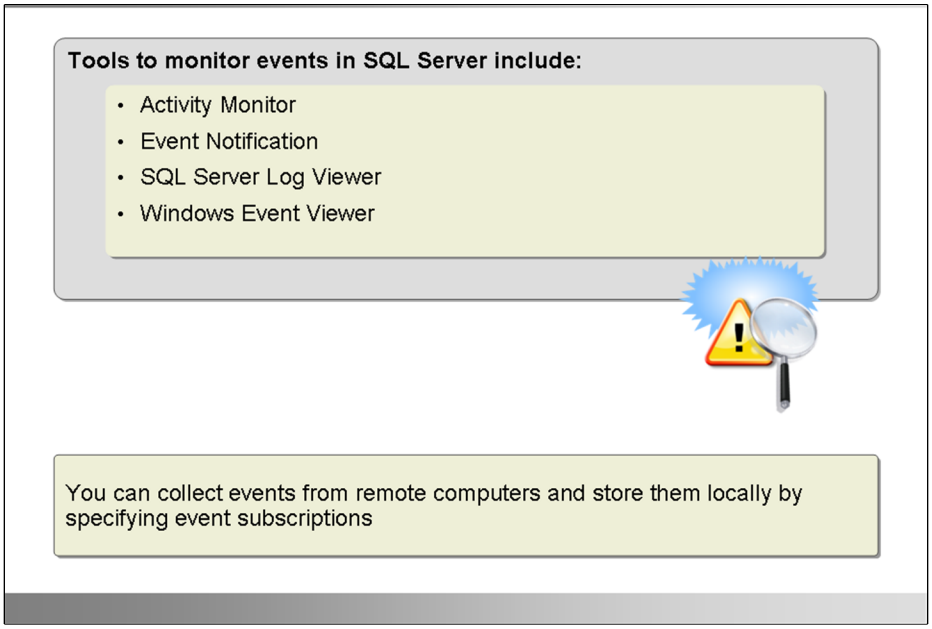
Source:

- http://msdn.microsoft.com/en-us/library/ms175903.aspx
- http://msdn.microsoft.com/en-us/library/ms178072(SQL.90).aspx
- http://technet.microsoft.com/en-us/library/cc917570.aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Tools for Monitoring Events in SQL Server



Key Points

SQL Server provides SQL Trace and event notifications to monitor events that occur in the Database Engine.

While running, SQL Trace captures events in the Database Engine as they occur. The captured events are event class instances that you choose to monitor in a process called a trace. Each event class consists of an event name and data columns that describe attributes of a particular occurrence of an event class. When an event occurs that belongs to an event class specified in a trace, SQL Trace records the event name and the associated data in a trace file. You can access SQL Trace either through SQL Server Profiler or through T-SQL system-stored procedures entered at the command line.

Tools to Monitor Events

The following tools in SQL Server help you to monitor events:

- **Activity Monitor**. Use SSMS to view the following information about SQL Server activity:

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

- Current user connections and locks.
- Process number, status, locks, and commands that active users are running.
- Objects that are locked, and the kinds of locks that are present.

If you are a system administrator, you can view additional information about a selected process, or close a selected process.

Use the Activity Monitor in SSMS to perform specific monitoring of an instance of SQL Server. This allows you to determine, at a glance, the volume and general types of activity on the system. For example, you can view the current blocked and blocking transactions, currently connected users on an instance of SQL Server and the last statement executed, and locks that are in effect.

You can monitor SQL Server activity by using the sp_who and sp_lock system stored procedures.

- **Event Notification**. Event notifications execute in response to a variety of T-SQL data definition language (DDL) statements and SQL Trace events by sending information about these events to a Service Broker service.

  Event notifications send information about the same events that are captured by SQL Trace to a Service Broker service. But unlike traces, you can use event notifications to perform an action inside SQL Server in response to events. Because event notifications execute asynchronously, these actions do not consume any resources defined by the immediate transaction.

- **SQL Server Log Viewer**. SQL Server logs certain system events and user-defined events to the SQL Server error log and the Microsoft Windows application log. Both logs automatically timestamp all recorded events. Use the information in the SQL Server error log to troubleshoot problems related to SQL Server.

  You can view the SQL Server error log by using SSMS or any text editor, such as notepad. By default, the error log is located at: Program Files\Microsoft SQL Server\MSSQL.n\MSSQL\LOG\ERRORLOG and ERRORLOG.n files

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

- **Windows Event Viewer**. When SQL Server is configured to use the Microsoft Windows application log, each SQL Server session writes new events to that log. The Windows application log records events that are logged by applications. For example, a database application might record a file error in the application log. Unlike the SQL Server error log, a new application log is not created each time you start an instance of SQL Server.

  The Windows application log provides an overall picture of events that occur on the Windows operating system, as well as events in SQL Server, and SQL Server Agent. You can use the Windows Event Viewer to view the Windows application log and to filter the information. For example, you can filter events, such as information, warning, error, success audit, and failure audit.

Event Subscriptions in Event Viewer

Event Viewer enables you to view events on a single remote computer. However, troubleshooting an issue might require you to examine a set of events stored in multiple logs on multiple computers.

Computers that run on the following operating systems can be event sources: Windows XP with Service Pack 2 (SP2), Windows Server® 2003 with Service Pack 1 (SP1), Windows Server 2003 with Service Pack 2 (SP2), Windows Server 2003 R2, Windows Vista®, Windows Vista with SP1, or Windows Server 2008.

You can collect copies of events from multiple remote computers and store them locally. To specify which events to collect, you create an event subscription. Among other details, the subscription specifies exactly which events will be collected and in which log they will be stored locally. Once a subscription is active and events are being collected, you can view and manipulate these forwarded events as you would any other locally stored events.

If you use the event collecting feature, you must configure both the forwarding and the collecting computers. The functionality depends on the Windows Remote Management (WinRM) service and the Windows Event Collector (Wecsvc) service. Both of these services must be running on computers participating in the forwarding and collecting process.

To create a new subscription:

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

1. On the collector computer, on the Start menu, point to **Administrative Tools**, click **Event Viewer**. The Event Viewer window appears.
2. In the console tree, under **Event Viewer (Local)**, right-click **Subscriptions**.and then click **Create Subscription**. The **Subscription Properties** dialog box appears.
3. In the **Subscription Properties -** dialog box, in the **Subscription name** box, type a name for the subscription.
4. Under **Subscription type and source computers**, click **Select Computers**. The **Computers** dialog box appears.
5. In the **Computers** dialog box, click **Add Domain Computers**. The **Select Computer** dialog box appears.
6. In the **Enter the object name to select (examples)** box, type the name of the computer from which events are to be collected, and then click **OK**.
7. In the **Computers** dialog box, click **OK**.
8. In the **Subscription Properties –** dialog box, click **Select Events**. The **Query Filter** dialog box appears.
9. In the **Query Filter** dialog box, use the controls to specify the criteria that events must meet to be collected, and then click **OK**.
10. In the **Subscription Properties -** dialog box, click **OK**. The subscription will be added to the Subscriptions pane and, if the operation was successful, the Status of the subscription will be Active.

Question: Which tool would you use to check the error logs of your SQL Server?

Question: What is the difference between Event Notification and Event Subscription?

Source:

- http://msdn.microsoft.com/en-us/library/ms190378.aspx
- http://msdn.microsoft.com/en-us/library/ms178094.aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

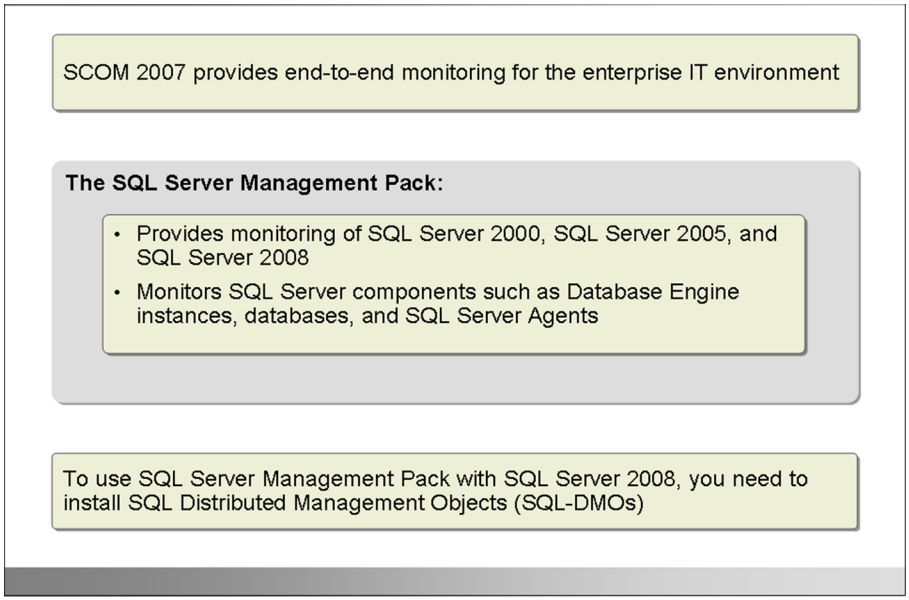
New Horizons Computer Learning Centers, Inc.

- http://msdn.microsoft.com/en-us/library/ms191202.aspx
- http://technet.microsoft.com/en-us/library/cc749183(WS.10).aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

SCOM Monitoring Tool



Key Points

System Center Operations Manager (SCOM) 2007 provides end-to-end monitoring for the organization IT environment. An Operations manager can monitor thousands of servers, applications, and clients and provide comprehensive views of their health states.

SQL Server Management Pack

The SQL Server Management Pack provides both proactive and reactive monitoring of the following:

- SQL Server 2000
- SQL Server 2005
- SQL Server 2008

SQL Server Management Pack monitors SQL Server components such as Database Engine instances, databases, and SQL Server Agents.

The monitoring provided by the management pack includes availability and configuration monitoring, performance data

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

collection, and default thresholds. You can integrate the monitoring of SQL Server components into your service-oriented monitoring scenarios.

In addition to performance monitoring capabilities, this management pack includes dashboard views, extensive knowledge with embedded in-line tasks, and views that enable near real-time diagnosis and resolution of detected issues.

Requirements of SQL Server 2008 to Support the SQL Server Management Pack

A number of workflows in the SQL Server 2008 management pack run scripts, which rely on SQL Distributed Management Object (SQL-DMO) to query information from the SQL Server. SQL-DMO is not provided as a part of SQL Server 2008. Every system with SQL Server 2008 that will be monitored must have SQL-DMO installed from the Microsoft SQL Server backward compatibility components.

Question: What are the key benefits of using SCOM over the other monitoring tools?

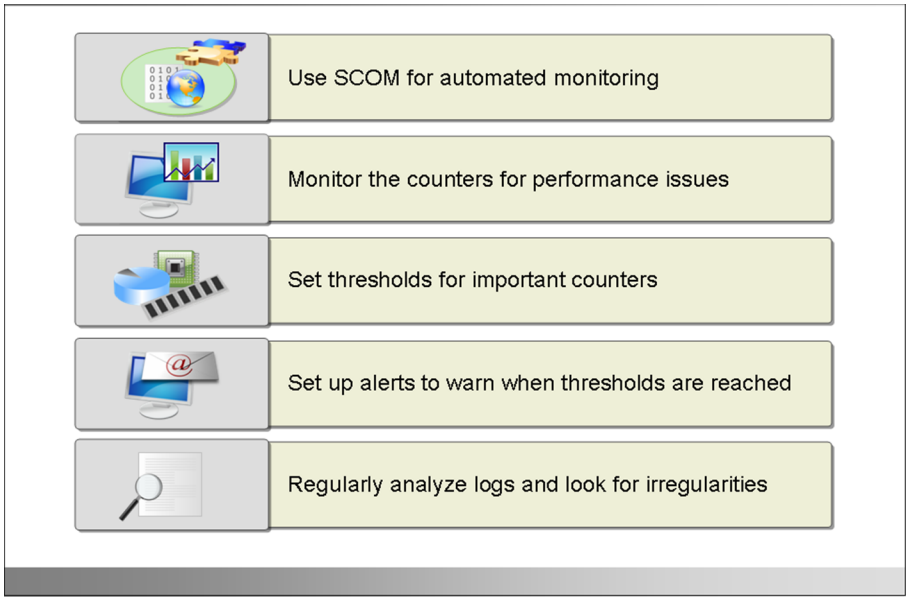
Source:

- http://technet.microsoft.com/en-us/library/dd767429.aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Best Practices for Monitoring SQL Server



Key Points

SQL Server plays an important role in an organization's IT environment as it support line-of-business applications and used as data store for business data. You need to monitor it regularly to ensure they it is running properly and to minimize chances of server failure. There are several tools that help you monitor the server. The guidelines to monitor SQL Server include:

- **Use SCOM for automated monitoring**. You can perform automated and proactive monitoring of SQL Servers in your environment by using SCOM and SQL Server Management Pack.
- **Monitor the counters for performance issues**. Monitor the performance counters for potential CPU, disk, and memory issues so that you can react before the problem arises.
- **Set thresholds for important counters**. With this, you have a baseline to compare to and have a clear picture on when the server is about to reach the limit. You should always monitor important counters and set thresholds on those counters. You

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

should collect performance data for a period of time and analyze it to find out the normal level of the reported value.

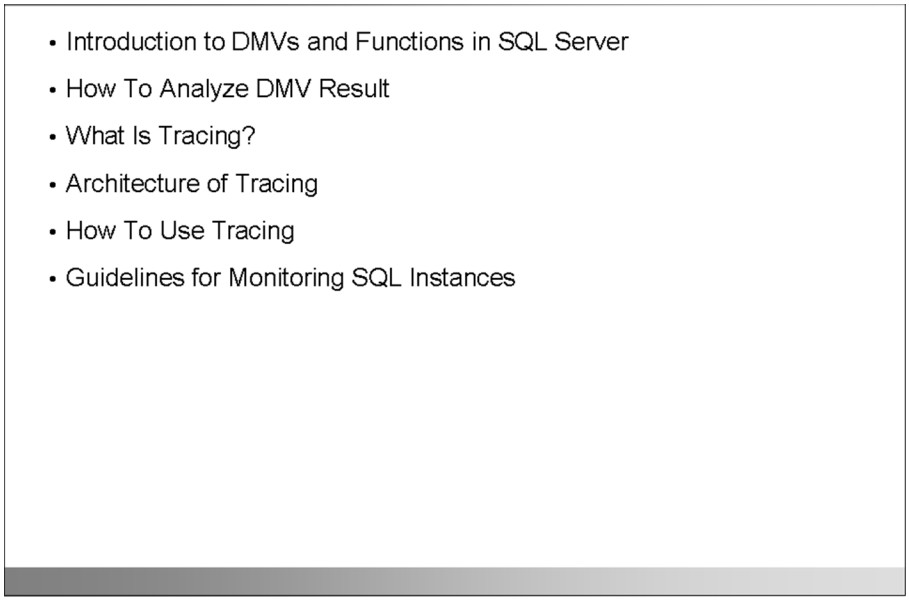
- **Set up alerts to warn when thresholds are reached**. Set up alerts to warn when thresholds are reached. Configure SQL Server or Windows to generate alerts when the reported values are over the threshold, so that you will get alerts when there are issues.
- **Regularly analyze logs and look for irregularities.** Check Windows application log and SQL Server error log for errors. There could be some information in these logs about some unusual behavior of SQL Servers.

Question: Why should you set up alerts to monitor events?

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Lesson 2
Monitoring SQL Server Instances



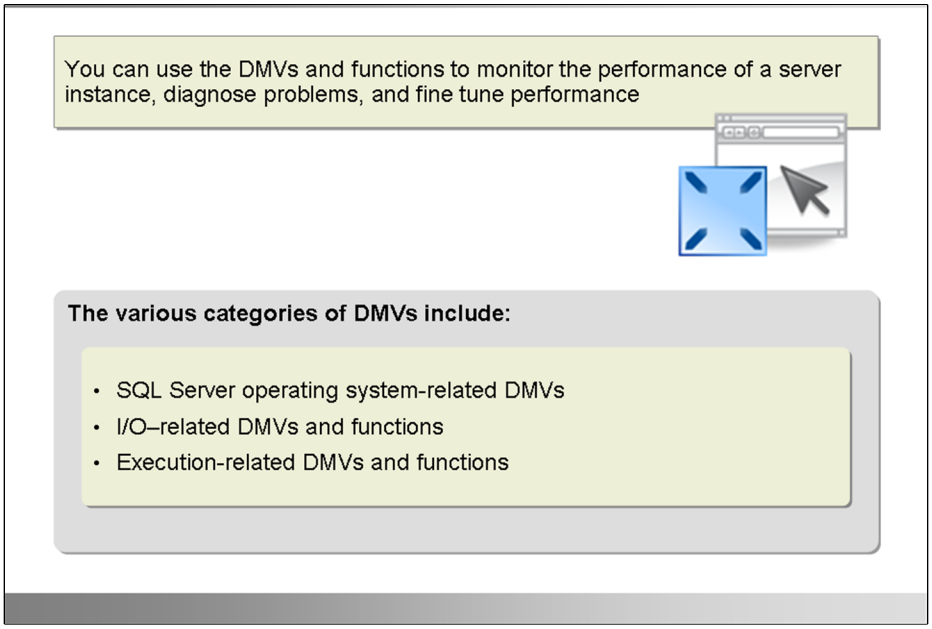
You can use DMVs and functions to return server state information that can be used to monitor the health of a server instance, diagnose problems, and tune performance. SQL Server provides SQL Trace to monitor events that occur in the Database Engine. By recording specified events, SQL Trace helps you troubleshoot performance, audit database activity, gather sample data for a test environment, debug T-SQL statements and stored procedures, and gather data for performance analysis tools. You can access SQL Trace either through SQL Server Profiler or through T-SQL system stored procedures.

In this lesson, you will learn about DMVs and functions in SQL Server. In addition, you will learn how to analyze DMV result. You will also learn about SQL Server Trace. You will then learn the architecture of tracing and how to use tracing. Finally, you will learn about the guidelines for monitoring SQL Server instances.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Introduction to DMVs and Functions in SQL Server



Key Points

DMVs and functions provide real-time information about the performance and state of SQL Server. They return the system state information.

Categories of DMVs

The various categories of DMVs include:

- **SQL Server operating system-related DMVs**. The SQL Server Operating System (SQLOS) is responsible for managing operating system resources that are specific to SQL Server.
- **I/O–related DMVs and functions**. This DMV provides information about the I/O operations of SQL Server. The various I/O related DMVs are:
  - sys.dm_io_backup_tapes
  - sys.dm_io_cluster_shared_drives

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

- sys.dm_io_pending_io_requests
- sys.dm_io_virtual_file_stats

- **Execution-related DMVs and functions**. This DMV provides information about query execution in SQL Server. The various execution-related DMVs and functions include:
  - sys.dm_tran_active_snapshot_database_transactions
  - sys.dm_tran_active_transactions
  - sys.dm_tran_current_snapshot
  - sys.dm_tran_current_transaction
  - sys.dm_tran_database_transactions
  - sys.dm_tran_locks
  - sys.dm_tran_session_transactions
  - sys.dm_tran_top_version_generators
  - sys.dm_tran_transactions_snapshot
  - sys.dm_tran_version_store

For more information, see I/O-related DMVs.

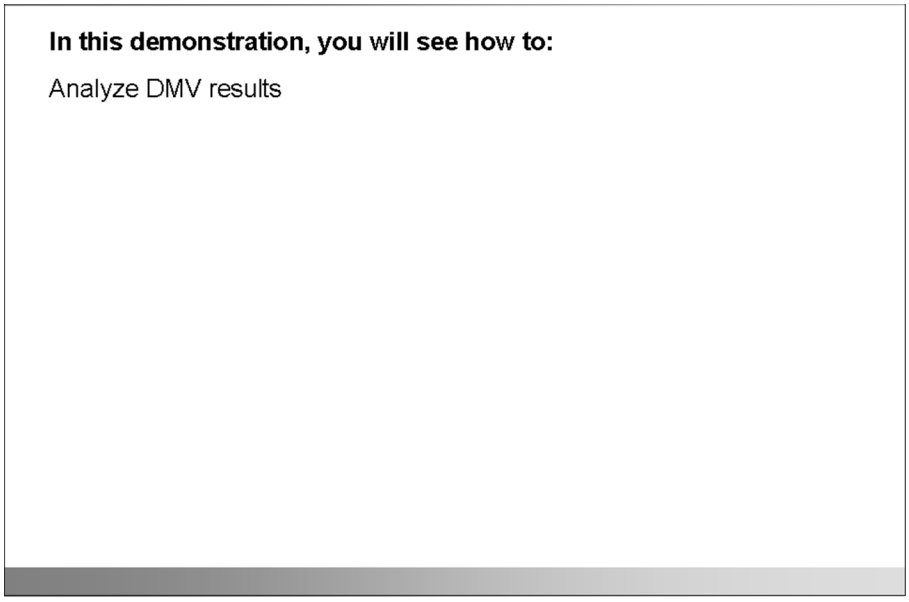
For more information, see Execution-related DMVs.

For more information, see SQL Server Operating System-related DMVs.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Demonstration: How To Analyze DMV Results



Key Points

The steps to analyze DMV results are:

1. On the **Start** menu, point to **All Programs**, click **Microsoft SQL Server 2008**, and then click **SQL Server Management Studio**. The Microsoft SQL Server Management Studio window appears.
2. In the Microsoft SQL Server Management Studio window, by using the **Connect to Server** dialog box, connect to the default instance of SQL Server Database Engine.
3. In Windows Explorer, browse to the **D:\Demofiles\Mod08** folder, and then double-click the **M8_Demo2.sql** file.
4. In the **Microsoft SQL Server Management Studio** window, click **Execute**.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

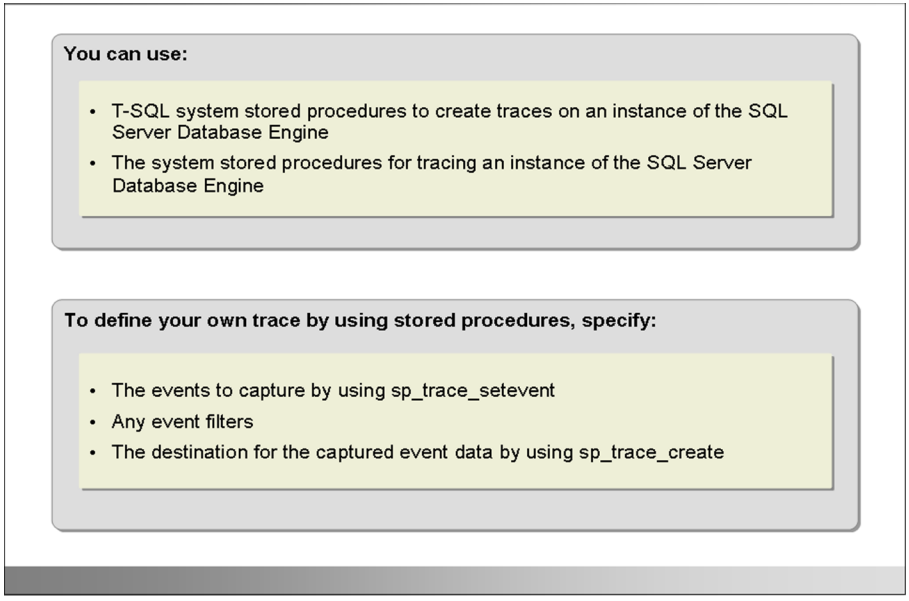
Question: What is the advantage of using the sys.dm_os_waiting_tasks DMV over sp_who2?

Question: What is the consideration of using DMV?

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

What Is Tracing?



Key Points

SQL Server provides T-SQL system stored procedures to create traces on an instance of the SQL Server Database Engine. These system stored procedures can be used from within your own applications to create traces manually, instead of using SQL Server Profiler. This manual creation allows you to write custom applications specific to the needs of your enterprise.

The following table lists the system stored procedures for tracing an instance of the SQL Server Database Engine.

Stored procedure	Task performed
fn_trace_geteventinfo	This procedure returns information about events included in a trace.
fn_trace_getinfo	This procedure returns information about a specified trace or all existing traces.
sp_trace_create	This procedure creates a trace definition. The

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Stored procedure	Task performed
	new trace will be in a stopped state.
sp_trace_generateevent	This procedure creates a user-defined event.
sp_trace_setevent	This procedure adds an event class or data column to a trace, or removes one from it.
sp_trace_setstatus	This procedure starts, stops, or closes a trace.
fn_trace_getfilterinfo	This procedure returns information about filters applied to a trace.
sp_trace_setfilter	This procedure applies a new or modified filter to a trace.

To define your own trace by using stored procedures, specify:

- The events to capture by using sp_trace_setevent.
- Any event filters.
- The destination for the captured event data by using sp_trace_create.

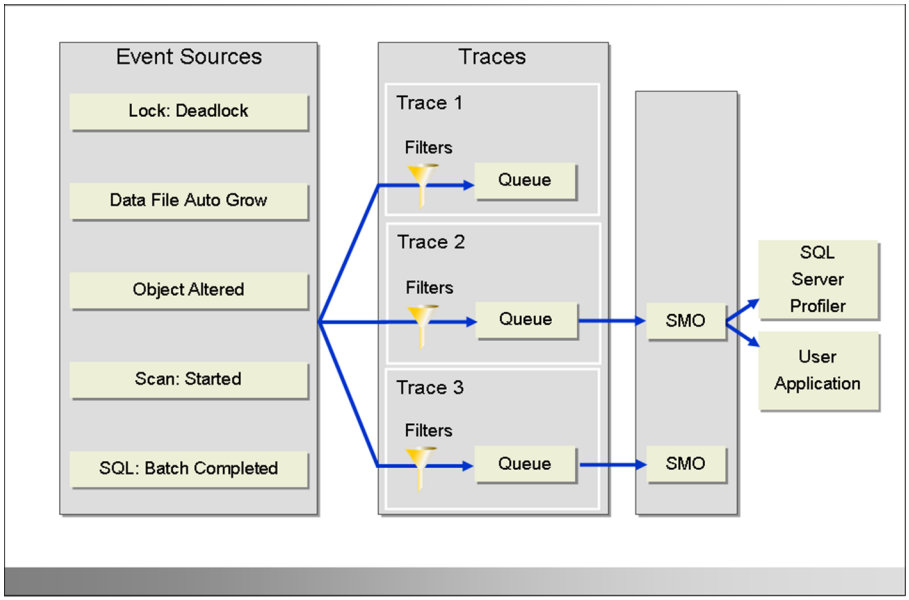
Source:

- http://msdn.microsoft.com/en-us/library/ms191006.aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Architecture of Tracing



Key Points

In SQL Trace, events that are instances of event classes listed in the trace definition are gathered. These events can be filtered out of the trace or queued for their destination. The destination can be a file or SQL Server Management Objects (SMO), which can use the trace information in applications that manage SQL Server.

The Event Source can be any source that produces the trace event, such as T-SQL batches or SQL Server events, such as deadlocks. After an event occurs, if the event class has been included in a trace definition, the trace gathers the event information. If filters have been defined for the event class in the trace definition, the filters are applied and the trace event information is passed to a queue. From the queue, the trace information is either written to a file, or is used by SMO in applications, such as SQL Server Profiler.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Question: What is the difference between using SQL Server Profile and SQL Tracing?

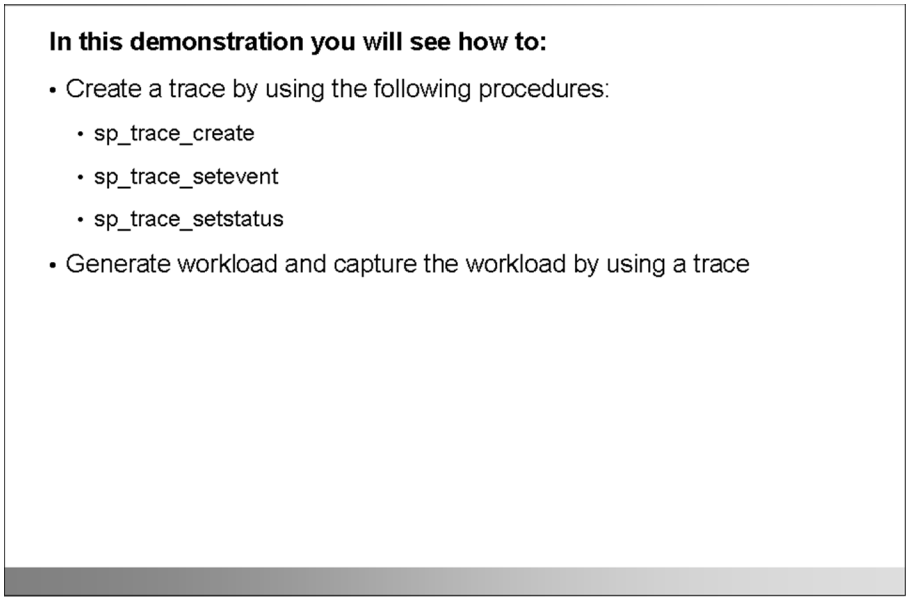
Source:

- http://msdn.microsoft.com/en-us/library/ms187933.aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Demonstration: How To Use Tracing



Key Points

The steps to enable tracing and monitor the result are:

1. In Windows Explorer, browse to the **D:\Demofiles\Mod08** folder, and then double-click the **M8_Demo3.sql** file.
2. In the Microsoft SQL Server Management Studio window, in the Query Editor pane, select the **Create and configure a trace** code, and then click **Execute**.
3. In the Microsoft SQL Server Management Studio window, in the Query Editor pane, select the **Generate a workload** code, and then click **Execute**.
4. In the Microsoft SQL Server Management Studio window, in the Query Editor pane, select the **Capture the workload** code, and then click **Execute**.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

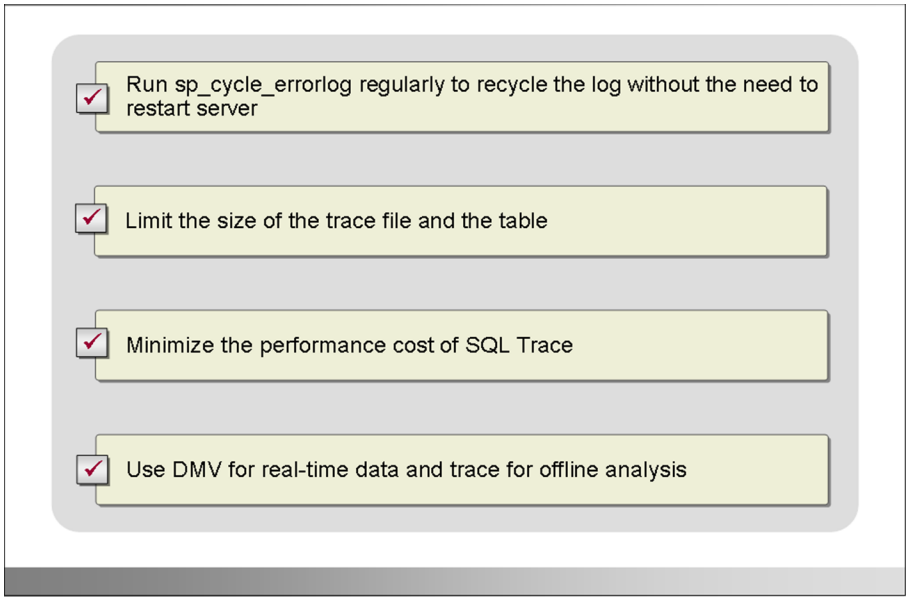
Question: What should you consider when using SQL Server Tracing?

Question: Where is the trace saved?

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Guidelines for Monitoring SQL Instances



Key Points

You should monitor SQL Server instances for improved performance. The guidelines to monitor SQL Server instances are:

- **Run sp_cycle_errorlog regularly to recycle the log without the need to restart the server**. A new error log is created each time an instance of SQL Server is started. You can use the sp_cycle_errorlog system stored procedure to cycle the error log files without having to restart the instance of SQL Server. Generally, SQL Server retains backups of the previous six logs and gives extensions depending upon the sequence of the log backups. For example, SQL Server gives the most recent log backup, the extension .1 and the second most recent the extension .2. The current error log has no extension.
- **Limit the trace file and the table sizes**. The results of SQL Trace vary in size depending on the event classes that are included in the trace and the way in which the Database Engine is used. If you trace event classes that occur frequently, you can

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

minimize the amount of data that the trace collects by setting the maximum file size or the maximum number of rows. By specifying the maximum file size or rows, you can ensure that the trace file or table will not increase beyond the specified limit.

- **Optimize SQL**. Running SQL Trace affects system performance because it uses system resources to gather data. You can minimize the performance cost incurred by a trace by implementing the following guidelines:
  - Use the command prompt to run traces.
  - Avoid including events that occur frequently.
  - Focus on the trace to collect only events that provide relevant data.
  - Avoid collecting duplicate data
  - Use filters in the trace definition
- **Use DMV for real-time data and trace for offline analysis**. Use DMV to retrieve real-time statistics of SQL Server, and use SQL Trace for the detailed offline analysis of the performance of the SQL Server.

Question: Why should SQL Trace or SQL Server Profiler not be enabled on a production system during production hours?

Source:

- http://msdn.microsoft.com/en-us/library/ms187885.aspx
- http://msdn.microsoft.com/en-us/library/ms189034.aspx
- http://msdn.microsoft.com/en-us/library/ms187023.aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Lesson 3
Monitoring SQL Server Databases

- DMVs to Monitor SQL Server Databases
- What Is SQL Server Profiler?
- SQL Server Profiler Templates
- How To Use SQL Server Profiler
- Locking in Database Engine
- Monitoring Deadlocks
- How To Detect Deadlocks by Using SQL Server Profiler
- What Is Row Versioning?
- Database Engine Isolation Levels
- Optimizing Query Performance
- Guidelines for Monitoring Databases

You can monitor database to identify potential problems and possible threats that may affect the services of the databases. To safeguard the database against malicious attacks, you need to ensure that the data is stored in an optimal way in the database. You can use tools such as DMVs and SQL Server Profiler to monitor the database. You can also restrict changes to a database by using locking and row versioning. SQL Server uses this feature to prevent multiple users from modifying the same data at the same time. You can monitor queries to optimize performance. You can use the execution plan feature of the SQL Server Database Engine to view and analyze query execution plans.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

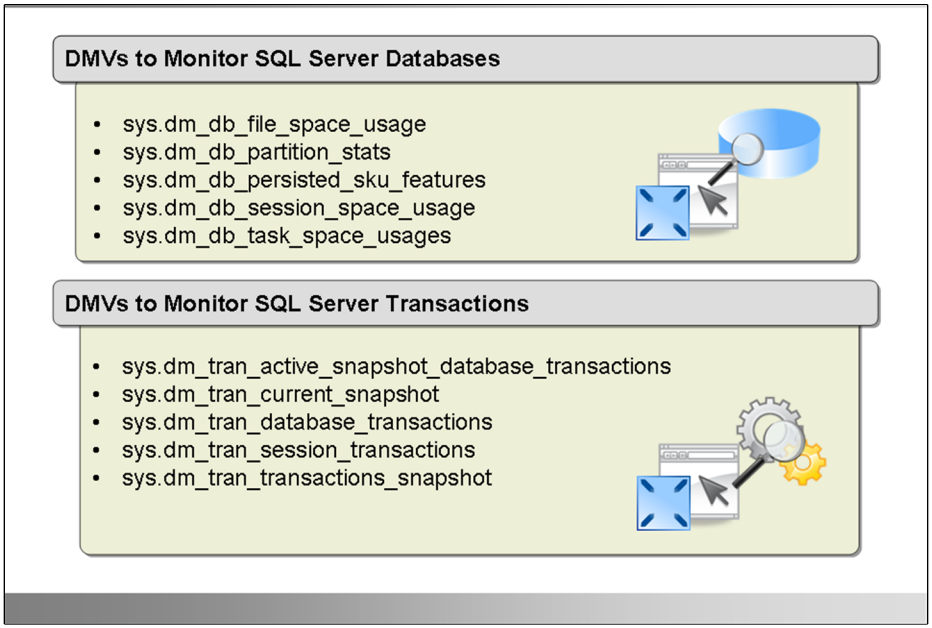
New Horizons Computer Learning Centers, Inc.

In this lesson, you will learn about the DMVs to monitor SQL Server databases. In addition, you will learn about SQL Server Profiler. You will also learn about deadlocks and how to detect deadlocks by using SQL Server Profiler. You will learn the locking and row versioning mechanism of Database Engine. You will also learn about the Database Engine isolation-levels.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

DMVs to Monitor SQL Server Databases



Key Points

There are some DMVs that provide detailed information about SQL Server databases.

DMVs to Monitor SQL Server Databases

You can use the following DMVs to monitor SQL Server databases:

- **sys.dm_db_file_space_usage**. This DMV returns space usage information for each file in the database.
- **sys.dm_db_partition_stats**. This DMV returns page and row-count information for every partition in the current database.
- **sys.dm_db_persisted_sku_features**. Some features of the SQL Server Database Engine change the way the Database Engine stores information in the database files. These features are restricted to specific editions of SQL Server. A database that contains these features cannot be moved to an edition of SQL Server that does not support them Use the

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

sys.dm_db_persisted_sku_features dynamic management view to list all edition-specific features that are enabled in the current database.

To determine whether a database uses any features that are restricted to specific edition, you can use the following command.

```
SELECT feature_name FROM
sys.dm_db_persisted_sku_features ;

GO
```

- **sys.dm_db_session_space_usage**. This DMV returns the number of pages allocated and deallocated by each session for the database.

  To determine the Amount of Free Space in tempdb, you can use the following command.

```
SELECT SUM(unallocated_extent_page_count) AS [free
pages],

(SUM(unallocated_extent_page_count)*1.0/128) AS [free
space in MB]

FROM sys.dm_db_file_space_usage;
```

- **sys.dm_db_task_space_usages**. This DMV returns page allocation and deallocation activity by task for the database.

DMVs to Monitor SQL Server Transactions

You can use the following DMVs to monitor SQL Server transactions:

- **sys.dm_tran_active_snapshot_database_transactions**. This DMV returns a virtual table for all active transactions that generate or potentially access row versions. Transactions are included for the following conditions:
  - When either or both ALLOW_SNAPSHOT_ISOLATION and READ_COMMITTED_SNAPSHOT database options are enabled:
    - There is one row for each transaction that is running under snapshot isolation level, or read-committed isolation level that is using row versioning.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

- There is one row for each transaction that causes a row version to be created in the current database. For example, the transaction generates a row version by updating or deleting a row in the current database.

- When a trigger is fired, there is one row for the transaction under which the trigger is executing.
- When an online indexing procedure is running, there is one row for the transaction that is creating the index.
- When Multiple Active Results Sets (MARS) session is enabled, there is one row for each transaction that is accessing row versions.
- This DMV does not include system transactions. The syntax for this DMV is sys.dm_tran_active_snapshot_database_transactions.

The following example uses a test scenario in which four concurrent transactions, each identified by a transaction sequence number (XSN), are running in a database that has the ALLOW_SNAPSHOT_ISOLATION and READ_COMMITTED_SNAPSHOT options set to ON.

```
SELECT transaction_id, transaction_sequence_num,
commit_sequence_num, is_snapshot session_id,
first_snapshot_sequence_num, max_version_chain_traversed,
average_version_chain_traversed, elapsed_time_seconds FROM
sys.dm_tran_active_snapshot_database_transactions;
```

- **sys.dm_tran_current_snapshot**. This DMV returns a virtual table that displays all active transactions at the time when the current snapshot transaction starts. If the current transaction is not a snapshot transaction, this function does not return any row. sys.dm_tran_current_snapshot is similar to sys.dm_tran_transactions_snapshot, except that sys.dm_tran_current_snapshot returns only the active transactions for the current snapshot transaction. The syntax for this function is sys.dm_tran_current_snapshot.
- **sys.dm_tran_database_transactions**. This DMV returns information about transactions at the database level.
- **sys.dm_tran_session_transactions**. This DMV returns correlation information for associated transactions and sessions.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

- **sys.dm_tran_transactions_snapshot**. This DMV returns a virtual table for the sequence_number of transactions that are active when each snapshot transaction starts. You can use the information that is returned by this DMV to:
  - Find the number of currently active snapshot transactions.
  - Identify data modifications that are ignored by a particular snapshot transaction. For a transaction that is active when a snapshot transaction starts, all data modifications by that transaction, even after that transaction commits, are ignored by the snapshot transaction. The syntax for this DMV is dm_tran_transactions_snapshot.

Query plans for various types of T-SQL batches, such as ad hoc batches, stored procedures, and user-defined functions, are cached in an area of memory called the plan cache. Each cached query plan is identified by a unique identifier called a plan handle. You can specify this plan handle with the sys.dm_exec_query_plan DMV to retrieve the execution plan for a particular Transact-SQL query or batch.

If a T-SQL query or batch runs a long time on a particular connection to SQL Server, retrieve the execution plan for that query or batch to discover what is causing the delay. The following example shows how to retrieve the XML Showplan for a slow-running query or batch.

Note: To run this example, replace the values for session_id and plan_handle with values specific to your server.

First, retrieve the server process ID (SPID) for the process that is executing the query or batch by using the sp_who stored procedure:

```
USE master;
GO
exec sp_who;
GO
```

The result set that is returned by sp_who indicates that the SPID is 54. You can use the SPID with the sys.dm_exec_requests dynamic management view to retrieve the plan handle by using the following query:

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.


```
USE master;
GO
SELECT * FROM sys.dm_exec_requests
WHERE session_id = 1;
GO
```

The table that is returned by sys.dm_exec_requests indicates that the plan handle for the slow-running query or batch is 0x06000100A27E7C1FA821B10600, which you can specify as the plan_handle argument with sys.dm_exec_query_plan to retrieve the execution plan in XML format. The execution plan in XML format for the slow-running query or batch is contained in the query_plan column of the table returned by sys.dm_exec_query_plan.

```
USE master;
GO
SELECT * FROM sys.dm_exec_query_plan
(0x06000100A27E7C1FA821B10600);
GO
```

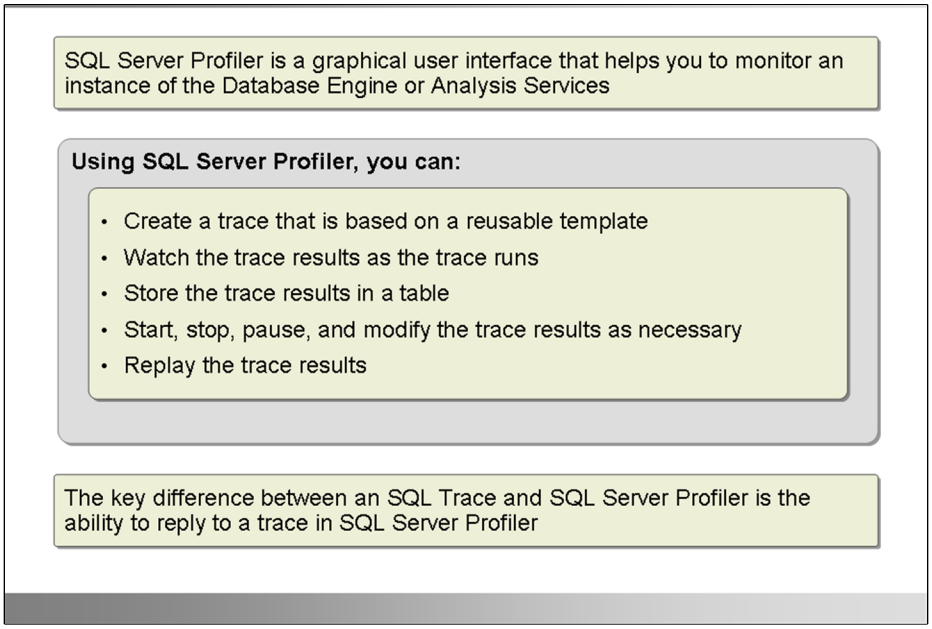
Source:

- http://msdn.microsoft.com/en-us/library/cc280724.aspx
- http://msdn.microsoft.com/en-us/library/ms187737.aspx
- http://technet.microsoft.com/en-us/library/ms180023.aspx
- http://technet.microsoft.com/en-us/library/ms184390.aspx
- http://technet.microsoft.com/en-us/library/ms186957.aspx
- http://technet.microsoft.com/en-us/library/ms173540.aspx
- http://msdn.microsoft.com/en-us/library/ms189747.aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

What Is SQL Server Profiler?



Key Points

SQL Server Profiler is a graphical user interface that helps you to monitor an instance of the Database Engine or Analysis Services. You can capture and save data about each event in a file or a table that you can analyze later. For example, you can monitor a production environment to see which stored procedures are affecting performance by executing too slowly.

Using SQL Server Profiler

SQL Server Profiler shows how SQL Server Database Engine and SSAS resolve queries internally. This allows you to see exactly what T-SQL statements or multi-dimensional expressions (MDX) are submitted to the server and how the server accesses the database or cube to return result sets.

Using SQL Server Profiler, you can:

- Create a trace that is based on a reusable template.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

- Watch the trace results as the trace runs.
- Store the trace results in a table.
- Start, stop, pause, and modify the trace results as necessary.
- Replay the trace results.

Comparison Between SQL Trace and SQL Server Profiler

The key difference between an SQL Trace and SQL Server Profiler is the ability to reply to a trace in SQL Server Profiler. SQL Server Profiler helps you to monitor the server in a graphic way, such as support of a deadlock graph in SQL Server Profiler.

You can use SQL Server Profiler to monitor only the events in which you are interested. If traces become too large, you can filter them based on the information you want, so that only a subset of the event data is collected. Monitoring too many events adds overhead to the server and the monitoring process, and this can increase the size of the trace file, XML file, or the trace table, especially when the monitoring process takes place over a long period of time. Using a trace file and an XML file, you can reduce the overhead of enabling SQL Server Trace, which may affect the performance of the monitored server.

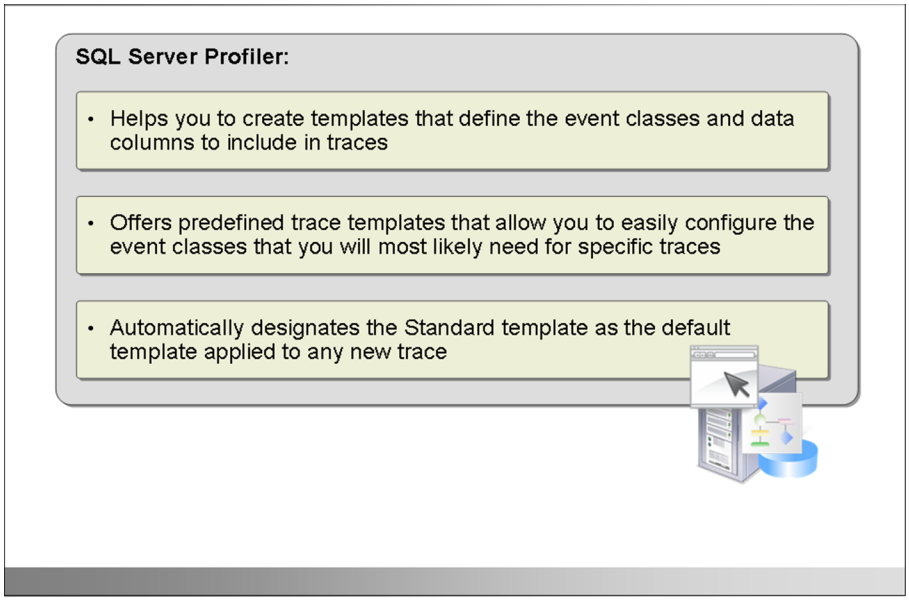
Source:

- http://www.microsoft.com/sql/prodinfo/compare/oracle/oracledba/quickguide/04_manageability/s8_SQL_server_manageability.swf
- http://msdn.microsoft.com/en-us/library/ms187929.aspx
- http://msdn.microsoft.com/en-us/library/cc296782%28BTS.10%29.aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

SQL Server Profiler Templates



Key Points

You can use SQL Server Profiler to create templates that define the event classes and data columns to include in traces. After you define and save the template, you can run a trace that records the data for each event class you selected. You can use a template on many traces; the template is not itself executed.

Predefined Templates

SQL Server Profiler offers predefined trace templates that allow you to easily configure the event classes that you will most likely need for specific traces. For example, the Standard template helps you create a generic trace for recording logons, logoffs, batches completed, and connection information. You can use this template to run traces without modification, or as a starting point for additional templates with different event configurations. You can use SQL Server Profiler to create templates that define the event classes and data columns, and these templates can be included in traces. After you define and save the template, you can run a trace that records the data for each

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

event class you selected. You can use a template in many traces; the template itself is not executed.

The following table lists the predefined templates, their purpose, and the event classes for which they capture information.

Template name	Template purpose	Event classes
SP_Counts	This template captures stored procedure execution behavior over time.	SP:Starting
Standard	This template is a generic starting point for creating a trace. It captures all stored procedures and T-SQL batches that are run. You can use this template to monitor general database server activity.	Audit Login Audit Logout ExistingConnection RPC:Completed SQL:BatchCompleted SQL:BatchStarting
TSQL	This template captures all T-SQL statements that are submitted to SQL Server by clients and the time issued. You can use this template to trublesgoot client applications.	Audit Login Audit Logout ExistingConnection RPC:Starting SQL:BatchStarting
TSQL_Duration	This template captures all T-SQL statements submitted to SQL Server by clients, their execution time in milliseconds, and groups them by duration. You can use this template to identify slow queries.	RPC:Completed SQL:BatchCompleted
TSQL_Grouped	This template captures all T-SQL statements submitted to SQL Server and the time they were issued. It groups information by user or client that submitted the statement. You can use this template to investigate queries from a particular client or user.	Audit Login Audit Logout ExistingConnection RPC:Starting SQL:BatchStarting
TSQL_Locks	This template captures all of the	Blocked Process

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Templatename	Template purpose	Event classes
	T-SQL statements that are submitted to SQL Server by clients along with exceptional lock events. You can use this template to troubleshoot deadlocks, lock time-out, and lock escalation events.	Report SP:StmtCompleted SP:StmtStarting SQL:StmtCompleted SQL:StmtStarting Deadlock Graph (Use against SQL Server 2005 or SQL Server 2008 instance.) Lock:Cancel Lock:Deadlock Lock:Deadlock Chain Lock:Escalation Lock:Timeout (Use against SQL Server 2000 instance.) Lock:Timeout (timeout>0) (Use against SQL Server 2005 or SQL Server 2008 instances.)
TSQL_Replay	This template captures detailed information about T-SQL statements that is required if the trace will be replayed. You can use this template to perform iterative tuning, such as benchmark testing.	CursorClose CursorExecute CursorOpen CursorPrepare CursorUnprepare Audit Login Audit Logout Existing Connection RPC Output Parameter RPC:Completed RPC:Starting Exec Prepared SQL Prepare SQL SQL:BatchCompleted

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Templatename	Template purpose	Event classes
		SQL:BatchStarting
TSQL_SPs	This template captures detailed information about all executing stored procedures. You can use this template to analyze the component steps of stored procedures. You need to add the SP:Recompile event if you suspect that procedures are being recompiled.	Audit Login Audit Logout ExistingConnection RPC:Starting SP:Completed SP:Starting SP:StmtStarting SQL:BatchStarting
Tuning	This template captures information about stored procedures and T-SQL batch execution. Use to produce trace output that DTA can use as a workload to tune databases.	RPC:Completed SP:StmtCompleted SQL:BatchCompleted

Default Template

SQL Server Profiler automatically designates the Standard template as the default template applied to any new trace. However, you can change the default template to any other predefined or user-defined template.

Question: If you want to optimize the performance of your application, which template would you use?

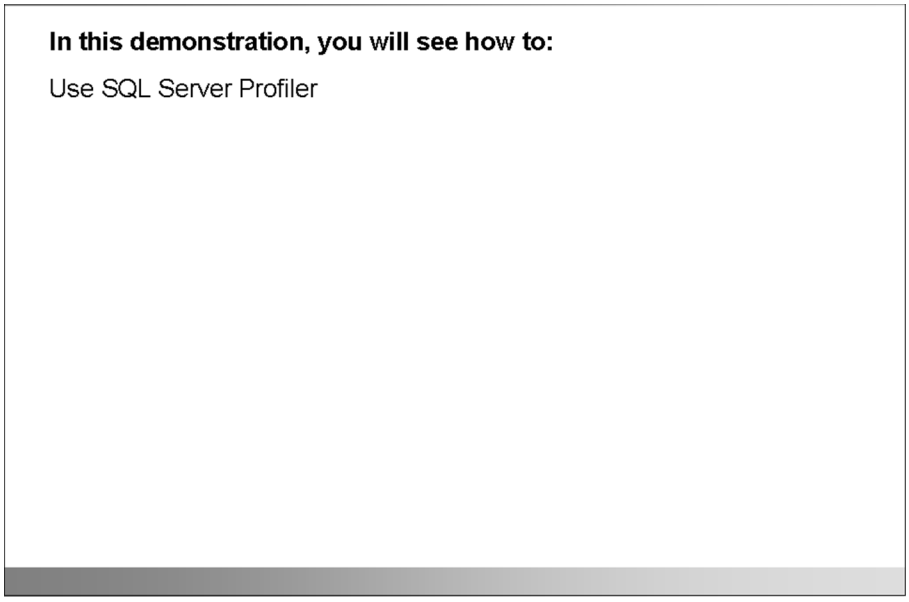
Source:

- http://msdn.microsoft.com/en-us/library/ms190176.aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Demonstration: How To Use SQL Server Profiler



Key Points

The steps to use SQL Server Profiler are:

1. On the **Start** menu, click **Microsoft SQL Server 2008**, click **Performance Tools**, and then click **SQL Server Profiler**. The SQL Server Profiler window appears.
2. In the SQL Server Profiler window, on the **File** menu, click **New Trace**. The **Connect to Server** dialog box appears.
3. In the **Connect to Server** dialog box, connect to the default instance of SQL Server by using **Windows Authentication**. The **Trace Properties** dialog box appears.
4. In the **Trace Properties** dialog box, click **Run** to start the trace with default settings.
5. In Windows Explorer, browse to the **D:\Demofiles\Mod08** folder, and then double-click the **M8_Demo4.sql** file.
6. In the Microsoft SQL Server Management Studio window, click

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Execute.

7. In the **SQL Server Profiler** window, click the **Stop** button to stop the trace and then observe the result.

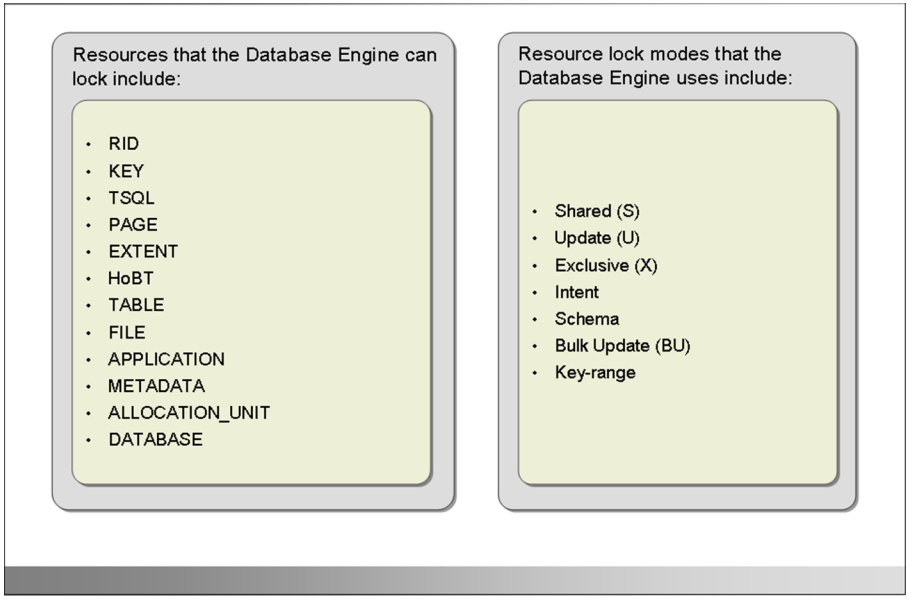
Question: What is difference between SQL Trace and SQL Profiler?

Question: What is a template in SQL Profiler?

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Locking in Database Engine



Key Points

SQL Server Database Engine has multigranular locking that allows different types of resources to be locked by a transaction. To minimize the cost of locking, the Database Engine locks resources automatically at a level appropriate to the task. Locking at a smaller granularity, such as rows, increases concurrency but has a higher overhead because more locks must be held if many rows are locked. Locking at a larger granularity, such as tables, are expensive in terms of concurrency because locking an entire table restricts access to any part of the table by other transactions. However, it has a lower overhead because fewer locks are being maintained.

Lock Resources

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

The Database Engine often has to acquire locks at multiple levels of granularity to fully protect a resource. This group of locks at multiple levels of granularity is called a lock hierarchy. For example, to fully protect a read of an index, an instance of the Database Engine may have to acquire share locks on rows and intent share locks on the pages and tables.

The following table shows the resources that the Database Engine can lock.

Resource	Description
RID	A lock is applied to the row identifier that is used to lock a single row within a heap.
KEY	A row lock is applied within an index used to protect key ranges in serializable transactions within a table with clustered index.
PAGE	A lock is applied to a 8-kilobyte (KB) page in a database, such as data or index pages.
EXTENT	A lock is applied to a contiguous group of eight pages, such as data or index pages.
HoBT	HoBT refers to a heap or B-tree. A lock is applied that protects a B-tree (index) or the heap data pages in a table that does not have a clustered index.
TABLE	A lock is applied to the entire table, including all data and indexes.
FILE	A lock is applied to a database file.
APPLICATION	A lock is applied to an application-specified resource.
METADATA	A lock is applied to metadata locks.
ALLOCATION_UNIT	A lock is applied to an allocation unit.
DATABASE	A lock is applied to the entire database.

Resources Lock Modes

SQL Server Database Engine locks resources by using different lock modes that determine how the resources can be accessed by concurrent transactions.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

The following table shows the resource lock modes that the Database Engine uses.

Lock mode	Description
Shared (S)	This lock mode is used for read operations that do not change or update data, such as a SELECT statement.
Update (U)	This lock mode is used on resources that can be updated. It helps prevent a common form of deadlock that occurs when multiple sessions are reading, locking, and potentially updating resources later.
Exclusive (X)	This lock mode is used for data-modification operations, such as INSERT, UPDATE, or DELETE. It ensures that multiple updates cannot be made to the same resource at the same time.
Intent	This lock mode is used to establish a lock hierarchy. The types of intent locks are: intent shared (IS), intent exclusive (IX), and shared with intent exclusive (SIX).
Schema	This lock mode is used when an operation dependent on the schema of a table is executing. The types of schema locks are: schema modification (Sch-M) and schema stability (Sch-S).
Bulk Update (BU)	This lock mode is used when bulk copying data into a table and the TABLOCK hint is specified.
Key-range	This lock mode protects the range of rows read by a query when using the serializable transaction isolation level. It ensures that other transactions cannot insert rows that would qualify for the queries of the serializable transaction if the queries were run again.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Question: How does SQL Server use locks to control concurrency?

Question: When designing a query, why should you not add query hints to modify the default locking behavior of your query?

Source:

- http://msdn.microsoft.com/en-us/library/ms189849.aspx
- http://technet.microsoft.com/en-us/library/ms175519.aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

Monitoring Deadlocks

A deadlock occurs when two or more tasks permanently block each other with each task having a lock on a resource which the other tasks are trying to lock

Resources that can cause deadlocks are:

- Locks
- Worker threads
- Memory
- Parallel query execution-related resources
- Multiple Active Result Sets (MARS) resources

Deadlock detection is performed by a lock monitor thread that periodically initiates a search through all of the tasks in an instance of the Database Engine

Key Points

A deadlock occurs when two or more tasks permanently block each other with each task having a lock on a resource which the other tasks are trying to lock.

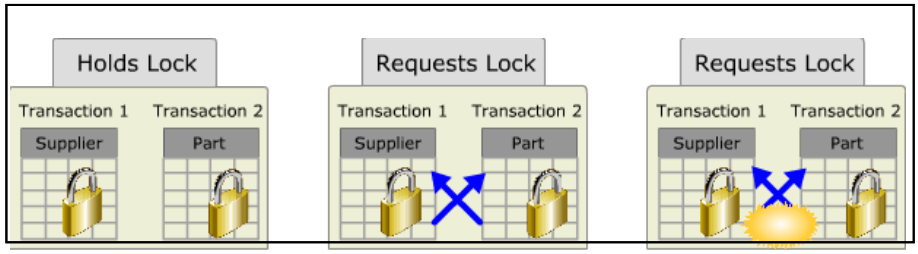


Figure 8.1. Deadlocking

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

The SQL Server Database Engine automatically detects deadlock cycles within SQL Server. The Database Engine chooses one of the sessions as a deadlock victim and the current transaction is terminated with an error to break the deadlock.

Resources that can Cause Deadlock

Each user session might have one or more tasks running on its behalf where each task might acquire or wait to acquire a variety of resources. The following types of resources can cause blocking that could result in a deadlock:

- **Locks**. Waiting to acquire locks on resources, such as objects, pages, rows, metadata, and applications can cause deadlock. For example, transaction T1 has a shared (S) lock on row r1 and is waiting to get an exclusive (X) lock on r2. Transaction T2 has a shared (S) lock on r2 and is waiting to get an exclusive (X) lock on row r1. This results in a lock cycle in which T1 and T2 wait for each other to release the locked resources.
- **Worker threads**. A queued task waiting for an available worker thread can cause deadlock. If the queued task owns resources that are blocking all worker threads, a deadlock will result. For example, session S1 starts a transaction and acquires a shared (S) lock on row r1 and then goes to sleep. Active sessions running on all available worker threads are trying to acquire exclusive (X) locks on row r1. Because session S1 cannot acquire a worker thread, it cannot commit the transaction and release the lock on row r1. This results in a deadlock.
- **Memory**. When concurrent requests are waiting for memory grants that cannot be satisfied with the available memory, a deadlock can occur. For example, two concurrent queries, Q1 and Q2, execute as user-defined functions that acquire 10MB and 20MB of memory respectively. If each query needs 30MB and the total available memory is 20MB, then Q1 and Q2 must wait for each other to release memory, and this result in a deadlock.
- **Parallel query execution-related resources**. Coordinator, producer, or consumer threads associated with an exchange port may block each other causing a deadlock usually when including at least one other process that is not a part of the parallel query. Also, when a parallel query starts execution, SQL Server determines the degree of parallelism, or the number of worker

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

threads, based upon the current workload. If the system workload unexpectedly changes, for example, where new queries start running on the server or the system runs out of worker threads, a deadlock could occur.

- **Multiple Active Result Sets (MARS) resources**. The following resources are used to control interleaving of multiple active requests under MARS:
  - **User resource**. When a thread is waiting for a resource that is potentially controlled by a user application, the resource is considered to be an external or user resource and is treated like a lock.
  - **Session mutex**. The tasks running in one session are interleaved, which means that only one task can run under the session at a given time. Before the task can run, it must have exclusive access to the session mutex.
  - **Transaction mutex**. All tasks running in one transaction are interleaved, which means that only one task can run under the transaction at a given time. Before the task can run, it must have exclusive access to the transaction mutex.

Deadlock Detection

All the resources that can cause deadlock participate in the Database Engine deadlock detection scheme. Deadlock detection is performed by a lock monitor thread that periodically initiates a search through all of the tasks in an instance of the Database Engine. The deadlock search interval is 5 seconds.

The following points describe the search process:

- If the lock monitor thread finds deadlocks, the deadlock detection interval will drop from 5 seconds to as low as 100 milliseconds depending on the frequency of deadlocks.
- If the lock monitors thread stops finding deadlocks, the Database Engine increases the intervals between searches to 5 seconds.
- If a deadlock has just been detected, it is assumed that the next threads that must wait for a lock are entering the deadlock cycle.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

The first couple of locks that wait after a deadlock has been detected will immediately trigger a deadlock search rather than wait for the next deadlock detection interval. For example, if the current interval is 5 seconds, and a deadlock was just detected, the next lock wait will kick off the deadlock detector immediately. If this lock wait is part of a deadlock, it will be detected right away rather than during next deadlock search.

The Database Engine typically performs periodic deadlock detection only. Because the number of deadlocks encountered in the system is usually small, periodic deadlock detection helps to reduce the overhead of deadlock detection in the system.

When the lock monitor initiates deadlock search for a particular thread, it identifies the resource on which the thread is waiting. The lock monitor then finds the owner(s) for that particular resource and recursively continues the deadlock search for those threads until it finds a cycle. A cycle identified in this manner forms a deadlock.

After a deadlock is detected, the Database Engine ends a deadlock by choosing one of the threads as a deadlock victim. The Database Engine terminates the current batch being executed for the thread, rolls back the transaction of the deadlock victim, and returns a 1205 error to the application. Rolling back the transaction for the deadlock victim releases all locks held by the transaction. This allows the transactions of the other threads to become unblocked and continue. The 1205 deadlock victim error records information about the threads and resources involved in a deadlock in the error log.

By default, the Database Engine chooses as the deadlock victim the session running the transaction that is least expensive to roll back. Alternatively, a user can specify the priority of sessions in a deadlock situation using the SET DEADLOCK_PRIORITY statement. DEADLOCK_PRIORITY can be set to LOW, NORMAL, or HIGH, or alternatively can be set to any integer value in the range (-10 to 10). The deadlock priority defaults to NORMAL. If two sessions have different deadlock priorities, the session with the lower priority is chosen as the deadlock victim. If both sessions have the same deadlock priority, the session with the transaction that is least expensive to roll back is chosen. If sessions involved in the deadlock cycle have the same deadlock priority and the same cost, a victim is chosen randomly.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Question: How can you minimize the chances for a deadlock to happen?

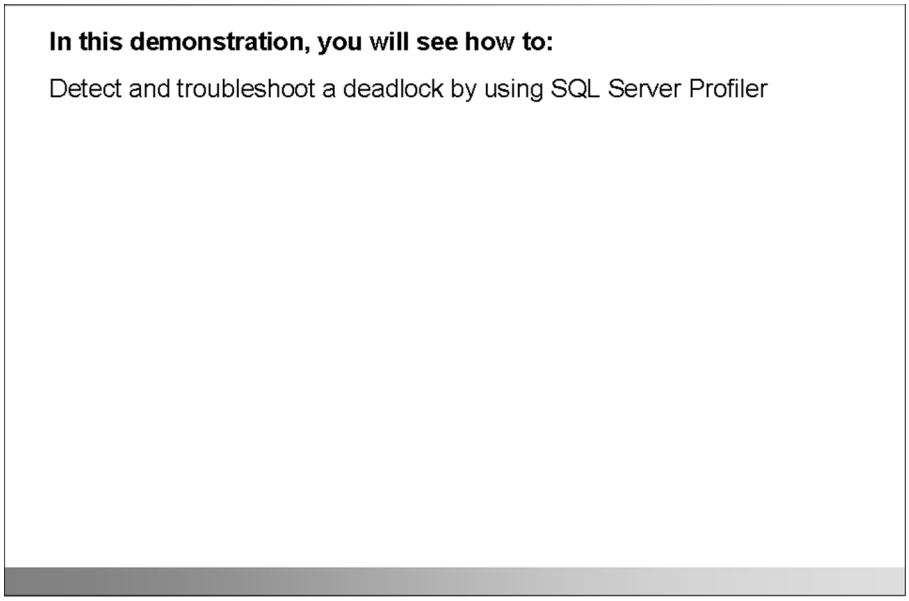
Source:

- http://msdn.microsoft.com/en-us/library/ms178104.aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Demonstration: How To Detect Deadlocks by Using SQL Server Profiler



Note: Before performing this demonstration, you need to execute the Prereq_mod08.sql file from the path, Demo files the D:\Demofiles\Mod08 folder.

Key Points

The steps to detect deadlocks by using SQL Server Profiler are as follows:

1. In the SQL Server Profiler window, on the **File** menu, click **New Trace**. The **Connect to Server** dialog box appears.
2. In the **Connect to Server** dialog box, connect to the default instance of SQL Server by using **Windows Authentication**. The **Trace Properties** dialog box appears.
3. In the Trace Properties dialog box, in the **Templates** list, click **TSQL_Locks** as the template, and then click **Run** to start the trace with the default settings

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

4. On the Start menu, click **Run**, type **cmd**, and then click **OK**. The Command Prompt window appears.
5. At the command prompt, type **sqlcmd**, and then press ENTER.
6. At the sqlcmd prompt, run the following command.

```
USE QuantamCorp
GO
CREATE TABLE dbo.DLTable
(col1 int)
GO
CREATE TABLE dbo.DLLock
(col1 int)
GO
INSERT INTO DLTable (1)
INSERT INTO DLLock (1)
```

7. At the sqlcmd prompt, run the following command.

```
BEGIN TRAN
UPDATE Mod08.dbo.DLTable SET col1 = 1
```

8. At the sqlcmd prompt, run the following command.

```
BEGIN TRAN
UPDATE Mod08.dbo.DLLock SET col1 = 1
UPDATE Mod08.dbo.DLTable SET col1 = 1
```

9. At the sqlcmd prompt, run the following command.

```
UPDATE Mod08.dbo.DLLock SET col1 = 1
```

10. In the **SQL Server Profiler** window, click the **Stop** button to stop the trace and then observe the result.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

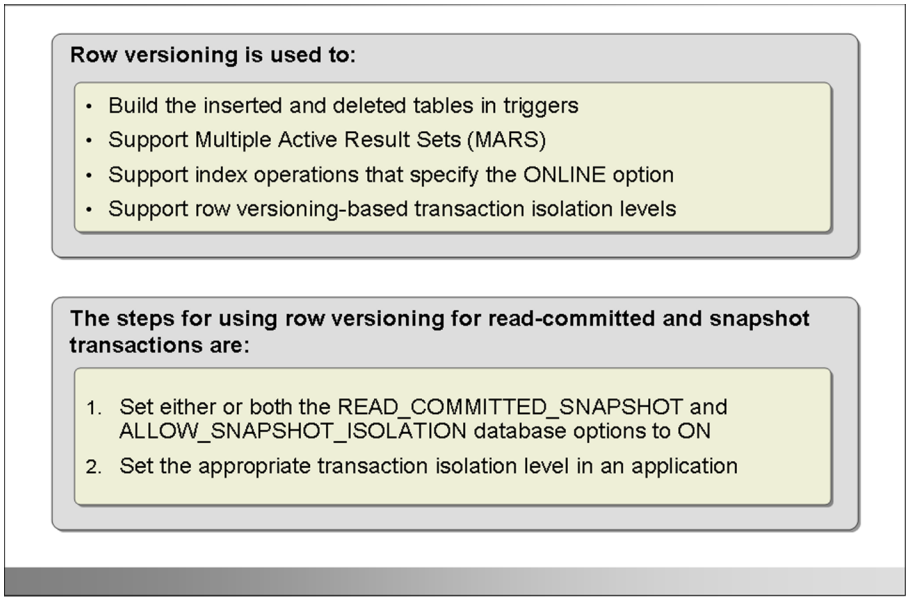
Question: What is the benefit of using SQL Server Profiler to monitor deadlock?

Question: How can you prevent deadlocks?

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

What Is Row Versioning?



Key Points

Row versioning is a general framework in SQL Server that is used to:

- Build the inserted and deleted tables in triggers. Any rows modified by the trigger are versioned. This includes the rows modified by the statement that launched the trigger and any data modifications made by the trigger.
- Support Multiple Active Result Sets (MARS). If a MARS session issues a data modification statement, such as INSERT, UPDATE, or DELETE, at a time there is an active result set, the rows affected by the modification statement are versioned.
- Support index operations that specify the ONLINE option.
- Support row versioning-based transaction isolation levels, such as:
  - A new implementation of read committed isolation level that uses row versioning to provide statement-level read consistency.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

- A new isolation level, snapshot, to provide transaction-level read consistency.

The tempdb database must have enough space for the version store. When tempdb database is full, UPDATE operations will stop generating versions and continue to succeed, but READ operations might fail because a particular row version that is needed no longer exists. This affects operations like triggers, MARS, and online indexing.

Process of Row Versioning

Using row versioning for read-committed and snapshot transactions is a two-step process:

1. Set either or both the READ_COMMITTED_SNAPSHOT and ALLOW_SNAPSHOT_ISOLATION database options to ON.
2. Set the appropriate transaction isolation level in an application.
    - When the READ_COMMITTED_SNAPSHOT database option is set to ON, transactions can set thee row versioning.
    - When the ALLOW_SNAPSHOT_ISOLATION database option is set to ON, transactions can set the snapshot isolation level.

When either READ_COMMITTED_SNAPSHOT or ALLOW_SNAPSHOT_ISOLATION database option is set to ON, the SQL Server Database Engine assigns a transaction sequence number (XSN) to each transaction that manipulates data by using row versioning. Transactions start at the time a BEGIN TRANSACTION statement is executed. However, the XSN starts with the first read or write operation after the BEGIN TRANSACTION statement. The XSN is incremented by one each time it is assigned.

When you set either the READ_COMMITTED_SNAPSHOT or ALLOW_SNAPSHOT_ISOLATION database options to ON, versions are maintained for all data modifications performed in the database. Every time a row is modified by a specific transaction, the instance of the Database Engine stores a version of the previously committed image of the row in tempdb. Each version is marked with the XSN of the transaction that is used to identify the version the transaction is looking for. The versions of modified rows are chained by using a link list. The newest row value is always stored in the current database and connected to the versioned rows stored in tempdb.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Note: For modification of large objects (LOBs), only the changed fragment is copied to the version store in tempdb.

Row versions are held long enough to satisfy the requirements of transactions running under row versioning-based isolation levels. The Database Engine tracks the earliest useful XSN and periodically deletes all row versions stamped with XSN that are lower than the earliest useful sequence number.

When both the READ_COMMITTED_SNAPSHOT and ALLOW_SNAPSHOT_ISOLATION database options are set to OFF, only rows that are modified by triggers or MARS sessions, or read by ONLINE index operations, are versioned. These row versions are removed when they are no longer needed. A background thread periodically executes to remove old row versions.

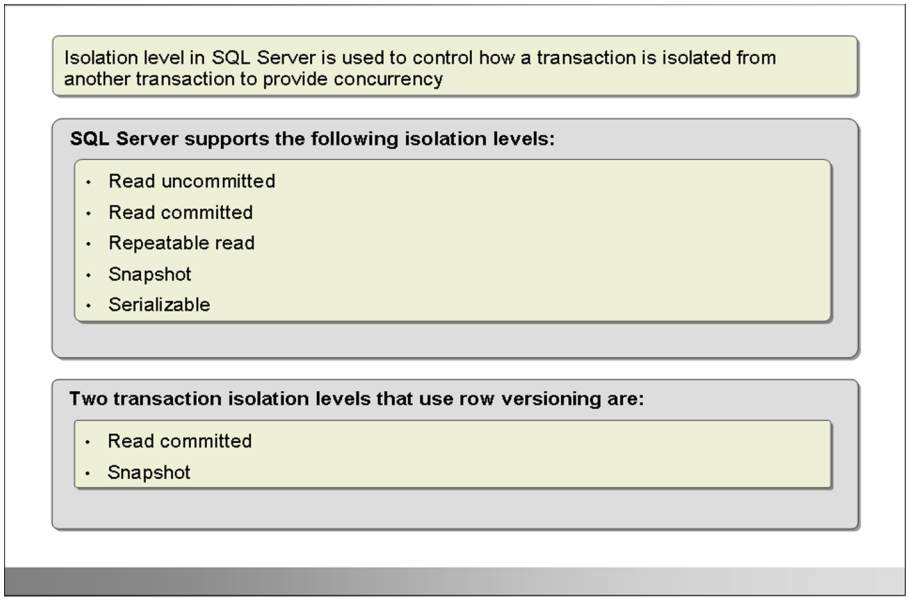
Source:

- http://msdn.microsoft.com/en-us/library/ms189050.aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Database Engine Isolation Levels



Key Points

Transactions specify an isolation level that defines the degree to which one transaction must be isolated from a resource or data modifications made by other transactions. Isolation levels are described in terms of which concurrency side-effects, such as dirty reads or phantom reads, are allowed.

- Transaction isolation levels control:
- The use of locks when data is read, and the type of lock requested.
- The time period to hold the read locks.
- The level of protection for a read operation referencing rows modified by another transaction by applying the following protection:
  - Blocking until the exclusive lock on the row is freed.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

- Retrieving the committed version of the row that existed at the time the statement or the transaction started.
- Reading the uncommitted data modification.

Choosing a transaction isolation level does not affect the locks acquired to protect data modifications. A transaction always gets an exclusive lock on any data it modifies, and holds that lock until the transaction completes, regardless of the isolation level set for that transaction. For read operations, transaction isolation levels primarily define the level of protection from the effects of modifications made by other transactions.

A lower isolation level increases the ability of many users to access data at the same time, but increases the number of concurrency effects, such as dirty reads or lost updates, that users might encounter. Conversely, a higher isolation level reduces the types of concurrency effects that users may encounter, but requires more system resources and increases the chances that one transaction will block another. Choosing the appropriate isolation level depends on balancing the data integrity requirements of the application against the overhead of each isolation level. The highest isolation level, serializable, guarantees that a transaction will retrieve exactly the same data every time it repeats a read operation, but it does this by performing a level of locking that is likely to impact other users in multi-user systems. The lowest isolation level, read uncommitted, may retrieve data that has been modified but not committed by other transactions. All the concurrency side effects can happen in read uncommitted, but there is no read locking or versioning, so overhead is minimized.

Database Engine Isolation Levels

The International Organization for Standardization (ISO) standard defines the following isolation levels, all of which are supported by the SQL Server Database Engine:

- **Read uncommitted**. In this level, statements can read rows that have been modified by other transactions, but not yet committed. This lowest isolation level not only increases the ability of users to access data simultaneously, but also increases the number of concurrency effects users might encounter.
- **Read committed**. In this level, statements cannot read data that has been modified, but not committed by other transactions.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

This option prevents dirty reads. It is the SQL Server default option.

- **Repeatable read**. In this level, statements cannot read data that has been modified, but not yet committed by other transactions and no other transaction can modify data that has been read by the current transaction until the current transaction completes.
- **Snapshot**. In this level, data read by any statement in a transaction will be the data that existed at the start of the transaction. The transaction can only recognize data modifications that were committed before the start of the transaction. Data modifications made by other transactions after the start of the current transaction are not visible to statements running in the current transaction. The effect is as if the statements in a transaction get a snapshot of the committed data as it existed at the start of the transaction.
- **Serializable**. In this level, statements cannot read data that has been modified, but not yet committed by other transactions. Other transactions cannot modify data that has been read by the current transaction until the current transaction completes. Other transactions cannot insert new rows with key values that would fall in the range of keys read by any statements in the current transaction until the current transaction completes. This level guarantees that a transaction will retrieve exactly the same data every time it repeats a read operation; but it does this by performing a level of locking that is likely to impact other users in multiuser systems.

SQL Server also supports two transaction isolation levels that use row versioning. They are:

- **Read committed**. When the READ_COMMITTED_SNAPSHOT database option is set to ON, read committed isolation level uses row versioning to provide statement-level read consistency. Read operations require only SCH-S table level locks and no page or row locks. When the READ_COMMITTED_SNAPSHOT database option is set to OFF, which is the default setting, read committed isolation function as it did in earlier versions of SQL Server. Both implementations meet the American National Standards Institute (ANSI) standard of read committed isolation.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

- **Snapshot**. The snapshot isolation level uses row versioning to provide transaction-level read consistency. Read operations acquire no page or row locks; only SCH-S table locks are acquired. When another transaction modifies the reading rows, they retrieve the version of the row that existed when the transaction started. Snapshot isolation is enabled when you set ON the ALLOW_SNAPSHOT_ISOLATION database option. By default, this option is set to OFF for user databases.

SQL Server does not support versioning of metadata. The following table shows the concurrency side effects enabled by the different isolation levels.

Isolation level	Dirty read	Nonrepeatable read	Phantom
Read uncommitted	Yes	Yes	Yes
Read committed	No	Yes	Yes
Repeatable read	No	No	Yes
Snapshot	No	No	No
Serializable	No	No	No

Question: What is the impact of using Serializable isolation level?

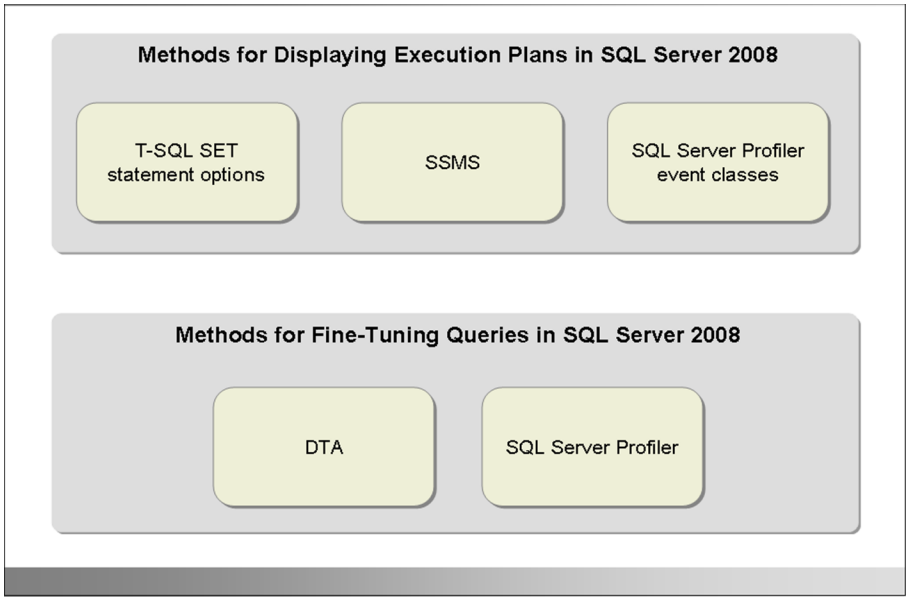
Source:

- http://msdn.microsoft.com/en-us/library/ms378149.aspx
- http://msdn.microsoft.com/en-us/library/ms189122.aspx
- http://msdn.microsoft.com/en-us/library/ms173763.aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Optimizing Query Performance



Key Points

The SQL Server Database Engine can display how it navigates tables and uses indexes to access or process the data for a query or other DML statement, such as UPDATE. This is called an execution plan. To analyze a slow-running query, it is useful to examine the query execution plan to determine what is causing the problem.

Methods for Displaying Execution Plans in SQL Server 2008

You can display execution plans by using the following methods:

- **T-SQL SET statement options**. When you use the T-SQL SET statement options, you can produce estimated and actual execution plans in XML or text. You can produce execution plans with SET SHOWPLAN_ALL or SET SHOWPLAN_TEXT for tabular representation, or the XML for representation produced by SET SHOWPLAN_XML.
- **SSMS**. SSMS displays an estimated graphical execution plan for statements that do not execute and an actual graphical execution

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

plan for executed statements, which you can save and view in Management Studio.

- **SQL Server Profiler event classes**. You can select SQL Server Profiler event classes to include traces that produce estimated and actual execution plans in XML or text in the trace results.

When you use one of these methods to display execution plans, the best execution plan used by the Database Engine for individual data manipulation language (DML) and T-SQL statements is displayed. The plan reveals compile-time information about stored procedures and called stored procedures that are invoked to an arbitrary number of calling levels. For example, executing a SELECT statement may show that the Database Engine uses a table scan to obtain the data. Execution of the SELECT statement may also show that an index scan will be used if the Database Engine determines that an index scan is a faster method of retrieving the data from the table.

Methods for Fine-Tuning Queries in SQL Server 2008

The methods for fine-tuning queries in SQL Server are:

- **DTA**. SQL Server DTA helps to select and create an optimal set of indexes, indexed views, and partitions without requiring an expert understanding of the structure of the database or the internals of Microsoft SQL Server.

  DTA analyzes a workload and the physical implementation of one or more databases. DTA uses trace files, trace tables, or T-SQL scripts as workload input when tuning databases. You can create T-SQL script workloads with the Query Editor in SSMS. You can create trace file and trace table workloads by using the Tuning Template in SQL Server Profiler.

  After analyzing a workload, DTA can recommend that you add, remove, or modify physical design structures in your databases. The advisor can also recommend what statistics should be collected to back up physical design structures. The physical design structures include clustered indexes, nonclustered indexes, indexed views, and partitioning. DTA recommends a set of physical design structures that reduces the query optimizers estimated cost of the workload.

- **SQL Server Profiler**. SQL Server Profiler shows how SQL Server resolves queries internally. This allows you to see exactly what T-SQL statements or Multi-dimensional expressions are submitted

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

to the server and how the server accesses the database or cube to return result sets.

Using SQL Server Profiler, you can create a trace, run the trace and watch the results as it runs, and store the results in a table. You can start, stop, pause, and modify results and you can also replay the trace results.

Question: You are working on fine tuning queries provided by your application developer. You want to fine tune the table structure and schema and see if update is required. Which performance tuning tools should you use?

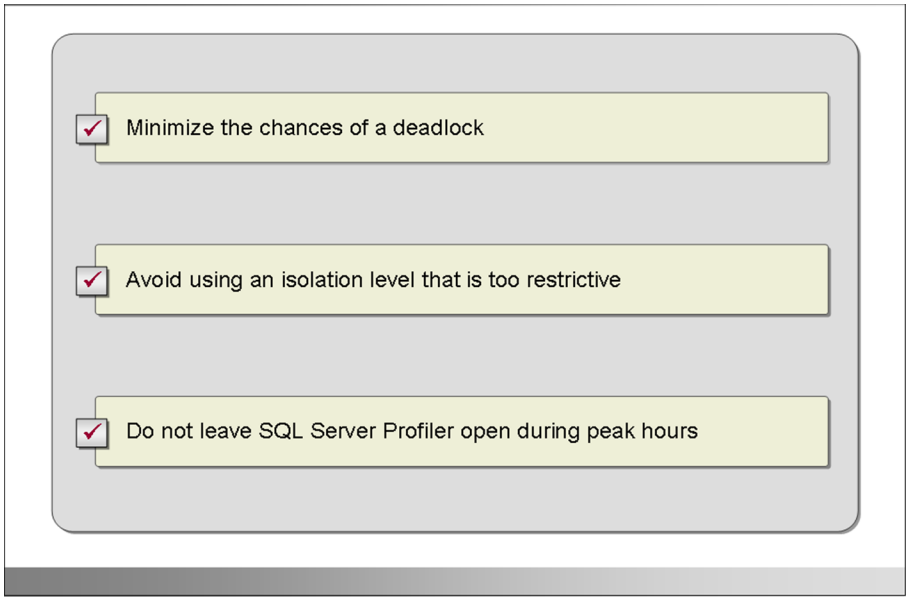
Source:

- http://technet.microsoft.com/en-us/library/ms191227.aspx
- http://msdn.microsoft.com/en-us/library/ms173494.aspx
- http://msdn.microsoft.com/en-us/library/ms187929.aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Guidelines for Monitoring Databases



Key Points

The guidelines to monitor databases are:

- **Minimize the chances of a deadlock**. Although deadlocks cannot be completely avoided, following certain coding conventions can minimize the chance of generating a deadlock. Minimizing deadlocks can increase transaction throughput and reduce system overhead because fewer transactions are:
- Rolled back, undoing all the work performed by the transaction.
- Resubmitted by applications because they were rolled back when deadlocked.

  To help minimize deadlocks:

  - Access objects in the same order.
  - Avoid user interaction in transactions.
  - Keep transactions short and in one batch.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

- Use a row versioning-based isolation level.
- Set READ_COMMITTED_SNAPSHOT database option ON to enable read-committed transactions to use row versioning.
- Use snapshot isolation.
- Avoid using too restrictive isolation level.

- **Avoid using an isolation level that is too restrictive**. Overly restrictive isolation level will cause excessive wait during SQL server query execution. Try to use the default isolation level and change it only if needed.
- **Do not leave SQL Server Profiler open during peak hours**. Use SQL Trace instead for detailed monitoring of SQL performance. Having SQL Server Profiler would consume more resources than others, which increase the wait time and lock duration of the system resources causing performance issues on the application database.

  SQL Profiler can cause huge performance impact on SQL server performance. If long time monitoring, like capturing SQL statement executed against the server, is need, use SQL Trace instead.

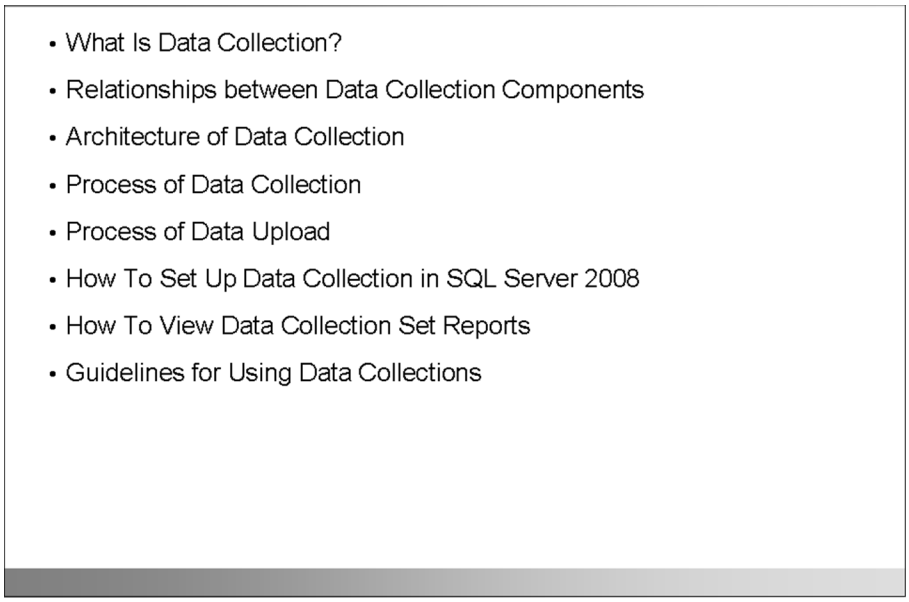
Source:

- http://msdn.microsoft.com/en-us/library/ms191242.aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Lesson 4
Monitoring Multiple SQL Server by Using Data Collection



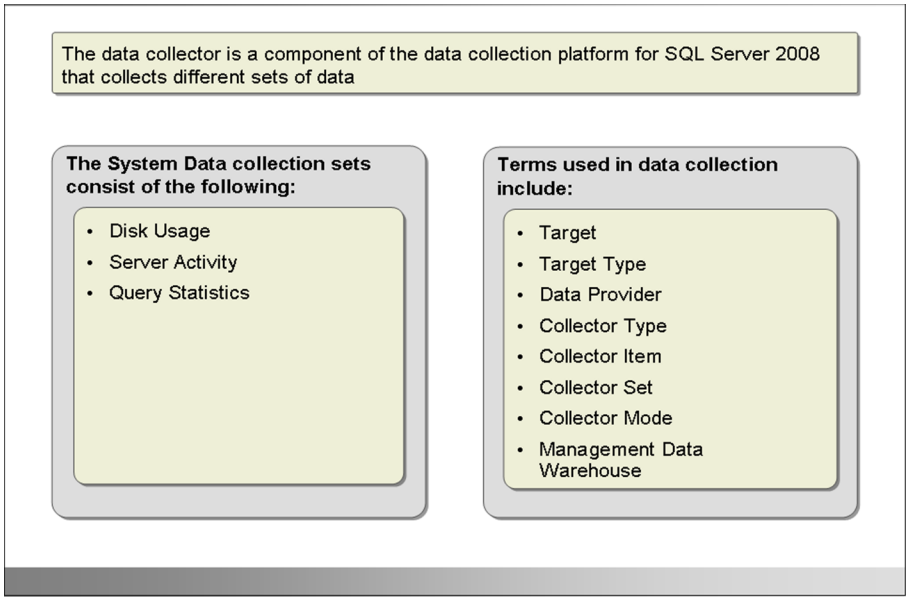
SQL Server 2008 provides a data collector that you can use to obtain and save data that is gathered from multiple SQL Servers. The data collector enables you to use data collection containers with which you can determine the scope and frequency of data collection on a computer that is running SQL Server 2008.

In this lesson, you will learn about data collection and the relationship between its components. In addition, you will learn the architecture of data collection. You will also learn about the processes of data collection and data upload. You will learn how to set up data collection in SQL Server 2008 and how to view collected data with data collection reports. You will finally learn about the guidelines for using data collectors.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

What Is Data Collection?



Key Points

The data collector is a core component of the data collection platform for SQL Server 2008. The data collector provides one central point for data collection across your database servers and applications. This collection point can obtain data from a variety of sources and is not limited to performance data, unlike SQL Trace.

The data collector enables you to adjust the scope of data collection to suit your test and production environments. The data collector also uses a data warehouse, a relational database that enables you to manage the data that you collect by setting different retention periods for your data.

The data collector supports dynamic tuning for data collection and is extensible through its application programming interface (API).

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

System Data Collection Sets

The data collector installs three System Data collection sets during the SQL Server 2008 Setup process. You can configure these collection sets to suit your monitoring requirements but you cannot delete them. The System Data collection sets consist of the following:

- **Disk Usage**. This collection set collects data about disk and log usage for all the databases installed on the system.
- **Server Activity**. This collection set collects resource usage statistics and performance data from the server and SQL Server.
- **Query Statistics**. This collection set collects query statistics, individual query text, query plans, and specific queries.

Data Collection Terminology

Some terms associated with data collection include:

- **Target**. A target is an instance of the Database Engine in an edition of SQL Server that supports Data Collection.
- **Target root**. A target root defines a subtree in the target hierarchy. A target set is the group of targets that results from applying a filter to a subtree defined by a target root. A target root can be a database, an instance of SQL Server, or a computer instance.
- **Target type**. It refers to the type of target, which has certain characteristics and behavior. For example, a SQL Server instance target has different characteristics than a SQL Server database target.
- **Data provider**. A data provider is a known data source, specific to a target type, which provides data to a collector type.
- **Collector type**. A collector type is a logical wrapper around the SSIS packages that provide the actual mechanism for collecting data and uploading it to the management data warehouse.
- **Collection item**. A collection item is an instance of a collector type. A collection item is created with a specific set of input properties and a collection frequency.
- **Collection set**. A collection set is a group of collection items. A collection set is a unit of data collection that a user can interact with through the user interface.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

- **Collection mode**. Collection mode refers to the manner in which the data is collected and stored. Collection mode can be cached or noncached. Cached mode supports continuous collection, whereas noncached mode is intended for on-demand collection or a collection snapshot.
- **Management data warehouse**. The management data warehouse is a relational database used to store collected data.

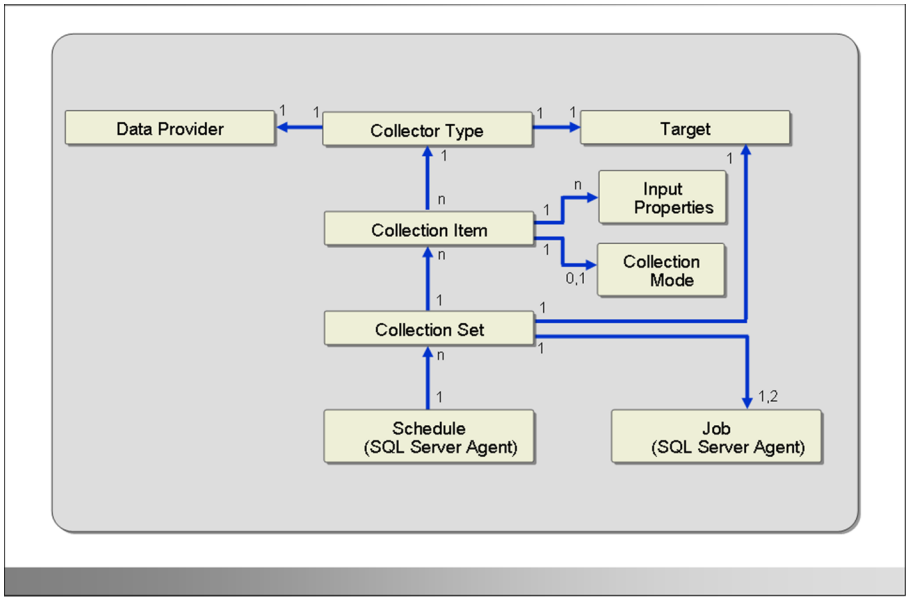
Source:

- http://msdn.microsoft.com/en-us/library/bb677248(SQL.105).aspx
- http://msdn.microsoft.com/en-us/library/bb964725.aspx
- http://msdn.microsoft.com/en-us/library/bb677279.aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Relationships between Data Collection Components



Key Points

In data collection, the data provider is external to the data collector and has an implicit relationship with the target. The data provider is specific to a particular target and provides data such as system views in SQL Server, Performance Monitor counters, and Windows Management Instrumentation (WMI) providers, which can be consumed by the data collector.

The collector type is specific to a target type, based on the logical association of a data provider to a target type. The collector type defines how data is collected from a specific data provider and specifies the data storage schema. The data provider schema and the storage schema are required to store the data that is collected. The collector type also provides the location of the management data warehouse, which can reside on the computer running data collection or on a different computer.

A collection item is an instance of a specific collector type, parameterized with input parameters, such as the XML schema for

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

the collector type. All collection items must operate on the same target root or on an empty target root. This enables the data collector to combine collector types from the operating system or from a specific target root, but not from different target roots.

A collection item has a collection frequency defined that determines how often snapshots of values are taken. Although it is a building block for a collection set, a collection item cannot exist on its own.

Collection sets are defined and deployed on a server instance and can be run independently of each other. Each collection set can be applied to a target that matches the target types of all the collector types that are part of a collection set. The collection set is run by a SQL Server Agent job or jobs, and data is uploaded to the management data warehouse on a predefined schedule.

All the data collected by different instances within the collection set is uploaded to the management data warehouse on the same schedule. This schedule is defined as a shared SQL Server Agent schedule and more than one collection set can use it. A collection set is turned on or turned off as a single entity; collection items cannot be turned on or turned off individually.

When you create or update a collection set, you can configure the collection mode for collecting data and uploading it to the management data warehouse. The type of scheduling is determined by the type of collection: cached or noncached. If the collection is cached, data collection and upload each run on a separate job. Collection runs on a schedule that starts when the SQL Server Agent starts and it runs on the frequency specified in the collection item. Upload runs according to the schedule specified by the user.

Under non-ached collection, data collection and upload both run on a single job, but in two steps. Step one is collection, step two is upload. No schedule is required for on-demand collection.

After a collection set is enabled, data collection can start, either according to a schedule or on demand. When data collection starts, SQL Server Agent spawns a process for the data collector, which in turn loads the Integration Services packages for the collection set. The collection items, which represent collection types, gather data from the appropriate data providers on the specified targets. When the collection cycle ends, this data is uploaded to the management data warehouse.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Question: What is the difference between data collection and event notification?

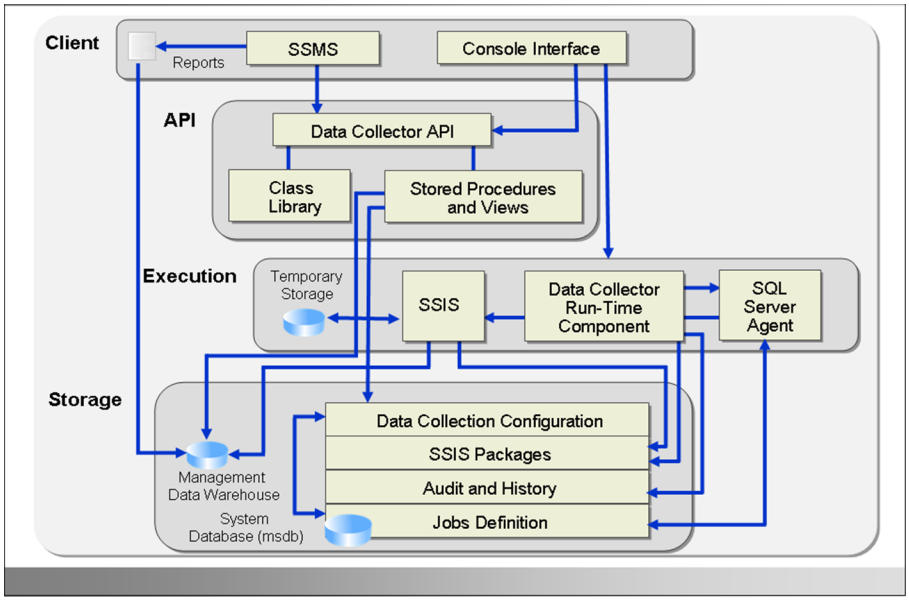
Source:

- http://msdn.microsoft.com/en-us/library/bb677279.aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Architecture of Data Collection



Key Points

The data collector is integrated with SQL Server Agent and SSIS, and uses both extensively. You can use SSIS to execute packages that collect data from individual data providers.

Data Collection Architecture

The components of the data collection architecture can be grouped into four categories:

- **Storage components**. They are the databases that contain configuration information and the collected data.
- **Execution components**. They are used for data collection and storage.
- **API components**. They are used to enable interaction between the user interfaces and the data collector.
- **Client components**. They are the user interfaces for the data collector.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Storage Components

The data collector stores data in the msdb database and the management data warehouse.

- **msdb System Database**. The data collector uses the msdb database for storing configuration information, run-time information, auditing, and collection history information. This database must be present on the instance of SQL Server that runs the data collector. By using msdb, you get the following benefits:
  - All the data needed to configure and run the data collector is in one location.
  - Data collection configuration can be deployed across several servers without having to use the file system.
  - The data collector can use existing SQL Server security mechanisms to protect the data. In addition, database roles can provide granular security, and there is no need to implement cross-database chaining.
  - It is possible to ensure the referential integrity of the configuration and run-time data because msdb is a relational database.

    In addition to storing collector-specific information, you can also store SQL Server Agent job information and SSIS package information.
- **Data Collection Configuration**. Data collection configuration-related tables and views include information about collection sets, collector types, and collection items. The tables and views also contain global data collector parameters, such as the location of the management data warehouse. Database roles for the data collector and management data warehouse are also stored here.
- **SQL Server Agent Jobs**. SQL Server Agent job information is stored in msdb by using a separate schema. This job information is created or updated whenever collection sets are added, removed, or changed. When a collection set is started, SQL Server Agent uses this information to execute data collector jobs.
- **SSIS Packages**. SSIS packages are stored in msdb by using existing SSIS storage tables. The data collector uses these

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

packages to collect data and upload it to the management data warehouse.

- **Auditing and History**. Dedicated storage is provided for saving detailed audit and historical records of data collection runs. The data collector run-time component and the SSIS packages generate events during data collection. This information is used to monitor and troubleshoot the collection process.
- **Management Data Warehouse**. The management data warehouse is a relational database that contains all the data that is retained. You can have this database on the same system as the data collector, or it can be on another computer.

Execution Components

Components in the execution category run in-memory and are responsible for collecting and storing data.

- **SQL Server Agent**. SQL Server Agent hosts the data collector run-time component. Data collection is scheduled as a SQL Server Agent job. When the job is run, it invokes the data collector run-time component using the parameters required to create and execute a collection set.
- **Data Collector Run-Time Component**. The data collector run-time component is hosted by a stand-alone process called Dcexec.exe. This component manages data collection based on the definitions provided in a collection set, and can accept any collection set as input. The data collector run-time component is responsible for loading and executing the SSIS packages that are part of a collection set.

  A collection set is run in one of the following collection and upload modes:

- **Non-cached mode**. Data collection and upload are on the same schedule. The packages start, collect and upload data at their configured frequency, and run until they are finished. After the packages finish, they are unloaded from memory.
- **Cached mode**. Data collection and upload are on different schedules. The packages collect and cache data until they receive a signal to exit from a loop control-flow task. This ensures that the data flow can be executed repeatedly, which enables continuous data collection.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

- **SSIS Run-Time Component**. The data collector run-time component invokes the SSIS packages that the data collector uses. The packages are run in-process.

  The two most important tasks for the SSIS packages are data collection and data upload, which are carried out by separate packages.

  - The collection package gathers data from a data provider and keeps it in temporary storage. If configured to do so, this package also appends additional information to the data it collects, such as the snapshot time or information about the source of the data.
  - The upload package reads the data in temporary storage, processes the data as required and then uploads the data to the management data warehouse. The upload is finished as a bulk insert to minimize the impact on server performance.

  Separating data collection and data upload provides more flexibility and efficiency. This design supports scenarios where snapshots of the data are captured at frequent intervals, but the collected data only needs to be uploaded every hour. Data collection and upload frequency should be determined by the monitoring requirements of a particular SQL Server installation.

- **Temporary Storage**. Temporary storage is used during data collection to store property values exposed by collection sets or calculated dynamically. These values can be raw data obtained during the collection process, or processed data, such as aggregated data.

API Components

The API category includes the class library and object model used to configure, maintain, and extend the data collector. Client access to the system database must come through the API.

- **Class Library**. The data collector has its own namespace and is independent of SQL Server Management Objects (SMOs).
- **Stored Procedures and Views**. The stored procedures and views provide Create, Read, Update, and Delete access to data collector configuration data and the management data warehouse.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Client components

The data collector uses SSMS and a Windows Powershell™ command-line interface as clients for configuring and using the data collector.

- **SSMS**. You can use SSMS to do the following:
  - Configure data collection using the Configure Data Collection Wizard in Object Explorer.
  - Enable or disable data collection by using Object Explorer.
  - Start or stop a specific collection set by using Object Explorer.
  - View or configure data collection or data collection set properties by using Object Explorer.
  - View reports based on historical data stored in the management data warehouse. This database can be located on the local instance or on another server. The latter is recommended.
- **Windows Powershell**. You can use the Windows Powershell command-line interface to do the following:
  - Start or stop a collection set.
  - Add or remove a collection set.
  - Change the editable properties of a collection set.

Question: You find that data collection is not functioning as expected. What is the first thing you should check?

Question: What is the purpose of having a centralized management database?

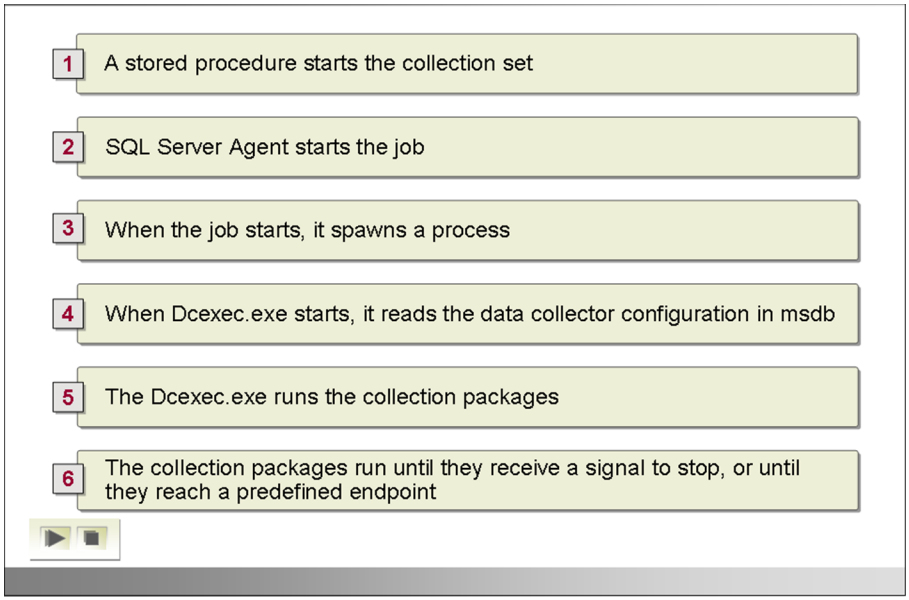
Source:

- http://msdn.microsoft.com/en-us/library/bb677355.aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Process of Data Collection



Key Points

The data collection process is as follows:

1. **A stored procedure starts the collection set**. The state of the collection set in msdb is changed to Running. This state change enables the SQL Server Agent jobs that are associated with the collection set.
2. **SQL Server Agent starts the job**. The job is started as follows:
   - If the collection set mode is set to cached, the collection job is started immediately.
   - If the collection set mode is set to noncached, the collection job is started according to the schedule that is configured.
3. **When a job starts, it spawns a process**. The process is to run the data collector run-time component (Dcexec.exe). The Dcexec.exe runs the collection package associated with the collection set.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

4. **When Dcexec.exe starts, it reads the data collector configuration in msdb**. Dcexec.exe helps to find out which collection mode, cached or noncached, to use.
5. **The data collector run-time component runs the collection packages**. The data collector run-time component uses the SSIS run-time component to run the collection packages defined for the collection set.
6. **The collection packages run until they receive a signal to stop, or until they reach a predefined endpoint**. The collection packages save the data that is collected in a staging area until it is uploaded.

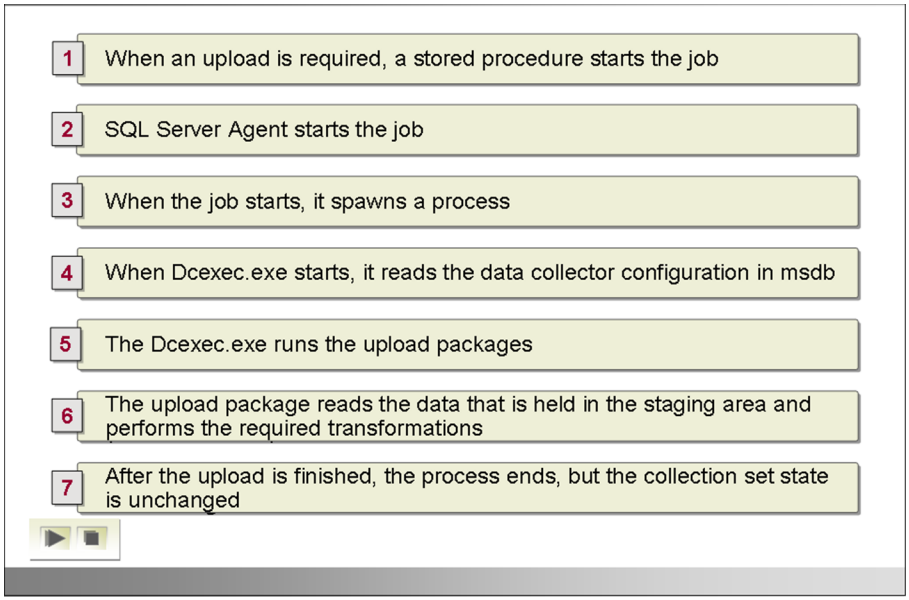
Source:

- http://msdn.microsoft.com/en-us/library/bb677355.aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Process of Data Upload



Key Points

The data upload process is as follows:

1. **When an upload is required, a stored procedure starts the job**. This job is handled by SQL Server Agent.
2. **SQL Server Agent starts the job**. The job is started as follows:
   - If an on-demand upload request is submitted by the user, the upload job is started immediately.
   - If there is no on-demand upload request, the upload job is started according to the schedule that is configured for the collection set.
3. **When a job starts, it spawns a process**. The process runs the data collector run-time component (Dcexec.exe). The Dcexec.exe runs the upload package associated with the collection set.
4. **When Dcexec.exe starts, it reads the data collector configuration in msdb**. Dcexec.exe helps to find out how to

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

handle the upload, such as the connection string to the management data warehouse, or any data processing that is required.

5. **The Dcexec.exe runs the upload packages**. The data collector run-time component uses the SSIS run-time component to run the upload packages that are defined for the collection set.
6. **The upload package reads the data that is held in the staging area and performs the required transformations**. The transformations can be normalization or sampling. When the transformations are finished, the upload package connects to the management data warehouse and inserts the data.
7. **After the upload is finished, the process ends, but the collection set state remains unchanged**. The next upload starts according to the collection set schedule.

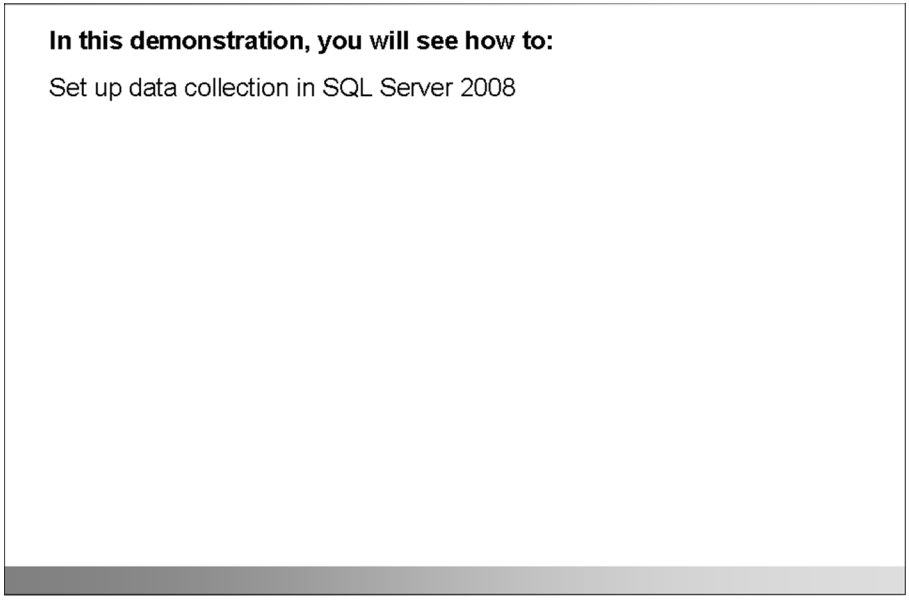
Source:

- http://msdn.microsoft.com/en-us/library/bb677355.aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Demonstration: How To Set Up Data Collection in SQL Server 2008



Key Points

The steps to set up data collection in SQL Server 2008 are:

1. In the Microsoft SQL Server Management Studio window, on the Object Explorer pane, expand **Management**, right-click **Data Collection**, and then click **Configure Management Data Warehouse**. The Configure Management Data Warehouse Wizard appears.
2. On the **Welcome to the Configure Management Data Warehouse Wizard** page, click **Next**.
3. On the **Select configuration task** page, ensure that the **Create or upgrade a management data warehouse** check box is selected, and then click **Next**.
4. On the **Configure Management Data Warehouse Storage** page, click **New**. The **New Database** dialog box appears.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

5. In the **New Database** dialog box, in the **Database name** box, type a name for the database, and then click **OK**.
6. In the **Configure Management Data Warehouse Wizard**, click **Next**.
7. On the **Map Logins and Users** page, under **Users mapped to this login**, select the check box of the login.
8. Under **Database role membership for:**, select the appropriate role, and then click **Next**.
9. On the **Complete the Wizard** page, click **Finish**.
10. On the **Configure Data Collection Wizard Progress** page, click **Close**.
11. In the Microsoft SQL Server Management Studio window, on the Object Explorer pane, under **Management**, right-click **Data Collection**, and then click **Configure Management Data Warehouse**. The Configure Management Data Warehouse Wizard appears.
12. On the **Welcome to the Configure Management Data Warehouse Wizard** page, click **Next**.
13. On the **Select configuration task** page, select the **Set up data collection** check box, and then click **Next**.
14. On the **Configure Management Data Warehouse Storage** page, click the **ellipsis** button besides Server name. The **Connect to Server** dialog box appears.
15. In the **Connect to Server** dialog box, connect to your default SQL Server instance with **Windows Authentication**
16. In the **Database name** list, click the database name that you created, and then click **Next**.
17. On the **Complete the Wizard** page, click **Finish**.
18. On the **Configure Data Collection Wizard Progress** page, click **Close**.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

19. In the Microsoft SQL Server Management Studio window, on the Object Explorer pane, under Data Collection, expand **System Data Collection Sets**, and right-click **Disk Usage**, and then click **Collect and Upload Now**. The **Collect and upload Data Collection Sets –** message box appears.
20. In the **Collect and upload Data Collection Sets –** message box, after the data collection process is complete, click **Close**.
21. In the Microsoft SQL Server Management Studio window, in the Object Explorer under **Data Collection**, right-click **Query Statistics**, and then click **Collect and Upload Now**.
22. In the **Collect and upload Data Collection Sets –** message box, after the data collection process is complete, click **Close**.
23. In the Microsoft SQL Server Management Studio window, in the Object Explorer under **Data Collection**, right-click **Server Activity**, and then click **Collect and Upload Now**.
24. In the **Collect and upload Data Collection Sets –** message box, after the data collection process is complete, click **Close**.

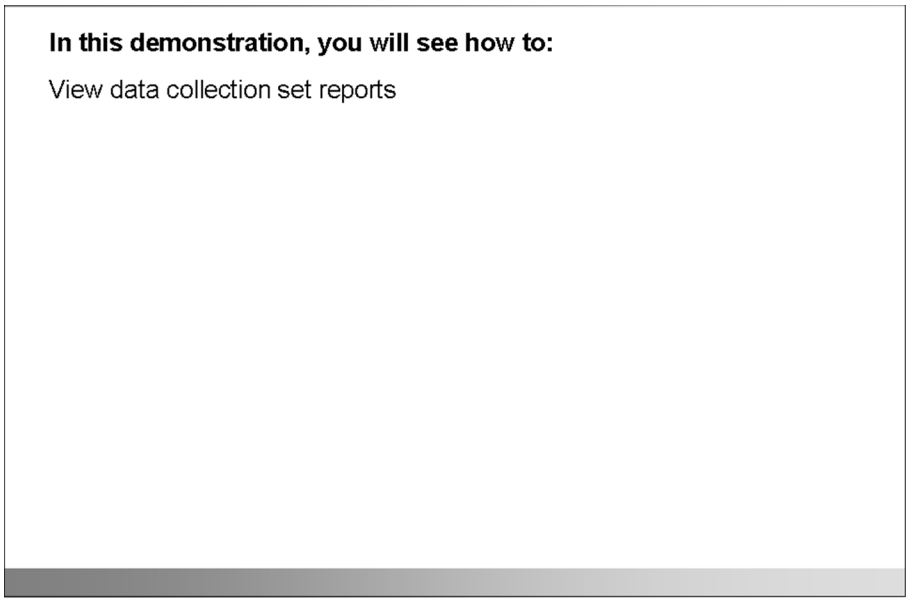
Question: What is the purpose of having a management data warehouse?

Question: Is it possible to have a single management data warehouse for multiple servers?

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Demonstration: How To View Data Collection Set Reports



Key Points

The steps to view data collection set reports are:

1. In the Microsoft SQL Server Management Studio window, in the Object Explorer pane, right-click Data Collection, point to **Reports**, click **Management Data Warehouse**, and then click **Server Activity History**. The Server Activity History report is displayed in the work pane.
2. In the Object Explorer pane, right-click Data Collection, point to **Reports**, click **Management Data Warehouse**, and then click **Disk Usage Summary**. The Disk Usage Collection Set report is displayed in the work pane.
3. In the Object Explorer pane, right-click Data Collection, point to **Reports**, click **Management Data Warehouse**, and then click **Query Statistics History**. The Query Statistics History report is displayed in the work pane.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Question: What information is collected by the data collector?

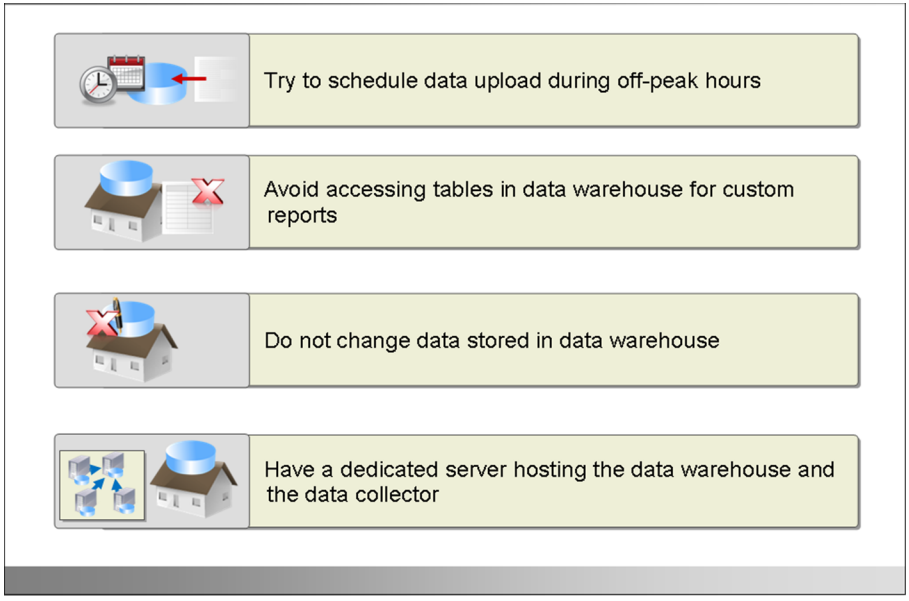
Question: Is it possible to create custom reports for the data collector?

Question: Is it possible to collect additional information from a target server?

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Guidelines for Using Data Collections



Key Points

The guidelines for using data collection are:

- **Try to schedule the data upload during off-peak hours**. Uploading a data job can be a resource-intensive process. Consider scheduling the upload of data to occur during off-peak hours to avoid impacting other applications.
- **Avoid accessing tables in data warehouse directly for custom reports**. Instead of directly using the tables, use the documented stored procedures and functions that are provided with the data collector to access instance and application data. The table names and definitions can change when the application is updated, in future releases.
- **Do not change data stored in data warehouse, directly**. Changing the data that you have collected invalidates the legitimacy of the collected data.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

- **Have a dedicated server hosting the data warehouse and data collector**. Data collection can be a resource-intensive process. Consider hosting the data collector and data warehouse on a dedicated computer to avoid performance impact to other computers.

Source:

- http://msdn.microsoft.com/en-us/library/bb677306.aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Lab 8: Designing a Monitoring Strategy for SQL Server 2008

- Exercise 1: Designing a Data Collection Architecture (Discussion)
- Exercise 2: Monitoring SQL Server Instance
- Exercise 3: Implementing Tracing in SQL Server 2008
- Exercise 4: Using DTA
- Exercise 5: Monitoring Performance by Using Data Collection

Logon Information

Virtual machine	50400A-NYC-SQL1
User name	Administrator
Password	Pa$$w0rd

Estimated time: 60 minutes

Objectives:

After completing this lab, you will be able to:

- Design a Data Collection architecture.
- Monitor SQL Server instance.
- Implement trace in SQL Server 2008.
- Use DTA.
- Monitor performance by using Data Collection.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Scenario

QuantamCorp has deployed two SQL Server 2008 servers in their organization and will continue to deploy more to replace theirexisting installations of SQL Server 2005. QuantamCorp wants tomake SQL Server 2008 the primary database engine for their organization. Recently, the application team reported that one application performed badly after they released a new version to the production environment. The application team has asked for yourhelp in diagnosing the problem. You will be checking differentaspects of SQL Server 2008 server and databases to see if they arecausing the problem.

You plan to monitor and fine tune your database performance. In this test, you plan to perform the following:

- Configure data collection.
- Determine the difference between Data Collection and Data Collector Set in Windows.
- Familiarize yourself with database tuning tools provided in SQL Server.
- Implement data collection.
- Use SQL Server Profiler to monitor performance and detect deadlocks.
- Use DTA to fine tune the database performance.

Exercise 1: Designing a Data Collection Architecture (Discussion)

Scenario

QuantamCorp is concerned about the performance and availability of its application. Your task is to collect performance information from all SQL Server 2008 servers in your organization for assistance in further analysis.

You organization has about 200 SQL Server 2008 servers worldwide. Your organization wants a solution that will help them collect basic information about each SQL Server. In addition, it should be feasible to collect additional information on each server.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

You have decided to use the Data Collector features to perform this task.

Based on the above scenario, answer the following questions:

1. How would you configure the data collection architecture?
2. How would you configure data collection such that the data upload process will have minimum impact to business?
3. Your colleague wants to know the difference between using Data Collection in SQL Server and the Data Collector Set in Windows. How would you respond to help him understand the difference?

Exercise 2: Monitoring SQL Server Instance

Scenario

You have a number of SQL Server servers whose performance is very slow. You want to monitor their performance and determine the cause for the issue. You decide to use the Performance Monitor to identify the possible cause for the performance issue and check whether the operating system or SQL Server is causing the problem.

The main tasks for this exercise are as follows:

1. Start and configure the Reliability and Performance Monitor tool.
2. Generate SQL workload.
3. Monitor the result.

Task 1: Start and configure the Reliability and Performance Monitor tool.

1. Open the Reliability and Performance Monitor window.
2. Using the Reliability and Performance Monitor window, add the following performance counters:
   - PhysicalDisk: **% Disk Time**
   - PhysicalDisk: **Avg. Disk Queue Length**
   - SQL Server:Buffer Manager: **Page writes/sec**
   - SQL Server:Buffer Manager: **Buffer Cache Hit Ratio**
   - SQL Server:Buffer Manager: **Total Pages**

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

- Processor: **% Privileged Time**
- Processor: **%User Time**
- System: **Processor Queue Length**
- Memory: **Available Bytes**
- Memory: **Pages/sec**
- Process: **Working Set**
- SQL Server:Memory Manager: **Total Server Memory (KB)**

Task 2: Generate SQL workload.

1. Open the Microsoft SQL Server Management Studio window.
2. In the **Connect to Server** dialog box, type the following to connect to Microsoft SQL Server Management Studio:
   - Server type: **Database Engine**
   - Server name: **NYC-SQL1**
   - Authentication: **Windows Authentication**
3. In the Query Editor pane of the Microsoft SQL Server Management Studio window, run a T-SQL code to perform the following:
4. Create a database **Mod08** with the following structure:
   - File Name: 'D:\data\Mod08_Primary.mdf'
   - Size: 10
   - MaxSize: 50
   - FileGrowth: 5
5. Create a table **[dbo].[CallDetails]** with the following structure:
   - Call_date: datetime
   - Caller: varchar
   - Duration: int
   - Additional_info: nchar
6. Add multiple records into the **CallDetails** table.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Note: You can browse to the **D:\Labfiles\Mod08** folder and select the **E2_T2.sql** file to refer to the required T-SQL code.

Task 3: Monitor the result.

1. Using the **Reliability and Performance Monitor** window, monitor the change in values reported to represent the SQL workload.

Results: After completing this exercise, you should have started and configure the Reliability and Performance Monitor tool generated SQL workload, and monitored the SQL workload result.

Exercise 3: Implementing Tracing in SQL Server 2008

Scenario

After checking the value of the performance counters, you realize that the application's queries could be the cause for the performance issue. You now want to monitor the queries executed against SQL Server and identify the query that is causing the problem.

The main tasks for this exercise are as follows:

1. Enable SQL trace.
2. Generate workload on the database.
3. Stop SQL Trace.
4. Analyze the captured trace.

Task 1: Enable SQL trace.

1. In the Query Editor pane of the Microsoft SQL Server Management Studio window, run a T-SQL script to create a trace named **Mod08_Trace** and select the return code to save to

   @TraceID. Then, set the events and data columns that you need to capture with **sp_trace_setevent** as per the following:

Event ID	Column ID
10	1

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Event ID	Column ID
12	15
13	1
13	11
13	14

Also, enable the trace by using the **sp_trace_setstatus** command.

Note: You can browse to the **D:\Labfiles\Mod08** folder and select the **E3_T1.sql** file to refer to the required T-SQL code.

Task 2: Generate workload on the database.

1. In the Query Editor pane of the Microsoft SQL Server Management Studio window, run a T-SQL script to generate workload on the database by clearing the data in the **CallDetails** table and then inserting the following record in the table:
   - call_date: **2004-01-01**
   - caller,duration: **Peter**
   - duration: **30**
   - additional_info: 'Call to' + Product of a random number and 100000.
2. Run a script to generate a table schema for testing the deadlock. Create two tables **dbo.DLTable** and **dbo.DLLock** with an integer column named **col1**.
3. Open the command prompt window, and start the SQL prompt.
4. At the SQL prompt, perform a transaction to update the **col1** column of the **dbo. DLTable** table to **1**.
5. At the SQLCMD prompt, perform a transaction to lock the **col1** column of the **dbo. DLTable** table with value **1**.
6. At the SQLCMD prompt, perform a transaction to lock the **col1** column of the **dbo. DLTable** table with value **1**.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Note: You can browse to the **D:\Labfiles\Mod08** folder and select the **E3_T2.sql** file to refer to the required T-SQL code.

Task 3: Stop SQL trace.

1. In the Query Editor pane of the Microsoft SQL Server Management Studio window, run a T-SQL script to disable the trace located at **D:\Data\Mod08_Trace.trc**.

Note: You can browse to the **D:\Labfiles\Mod08** folder and select the **E3_T3.sql** file to refer to the required T-SQL code.

Task 4: Analyze the captured trace.

1. Using the SQL Server Profiler window, open the SQL Server Profiler trace and monitor the recorded SQL statements.

Results: After completing this exercise, you should have enabled SQL trace, generated workload on the database, stopped SQL trace, and analyzed the captured trace.

Exercise 4: Using DTA

Scenario

After resolving the performance issue, you still notice that some of the queries are not providing optimal performance. You want to use DTA to analyze the structure of the database and check if you can do anything to further improve the performance.

The main tasks for this exercise are as follows:

1. Use DTA to analyze the captured trace.
2. Analyze the result.

Task 1: Use DTA to analyze the captured trace.

1. Open the Database Engine Tuning Advisor window.
2. In the Connect to Server dialog box, type the following to connect to Microsoft SQL Server 2008:
   - Server name: **NYC-SQL1**

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

- Authentication: **Windows Authentication**

3. Using the Database Engine Tuning Advisor window, perform tuning for the **Mod08_Trace.trc** captured workload against the **Mod08** database by specifying a time limit for the tuning. Also, let the physical design structure of the trace in the database include only clustered indexes.

Task 2: Analyze the result.

1. Using the Database Engine Tuning Advisor window, analyze Index Recommendations and Partition Recommendations.
2. Analyze the **index details (current)** and the **index details (recommended)** report.

Results: After completing this exercise, you should have generated a report for the captured trace by using DTA and analyzed the report.

Exercise 5: Monitoring Performance by Using Data Collection

Scenario

To help to monitor the performance of SQL Server, you decide to implement data collection to collect performance information about SQL Server for centralized monitoring.

The main tasks for this exercise are as follows:

1. Configure Management Data Warehouse.
2. Enable Data Collection.
3. Manually trigger the Data Collection process.
4. Generate Data Collection reports.

Task 1: Configure Management Data Warehouse.

1. Using the Microsoft SQL Server Management Studio window, start the Configure Management Data Warehouse Wizard and configure Management Data Warehouse for the **MDW** database with the following settings:

- Create a database **MDW**.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

- Map Logins and Users.
- Set Database role membership for the database **MDW**.

Task 2: Enable Data Collection.

1. Using the Microsoft SQL Server Management Studio window, start the Configure Management Data Warehouse Wizard.
2. Run the Configure Management Data Warehouse Wizard, and enable Data Collection for the **MDW** database by connecting to the server with the following credentials:
   - Server name: **NYC-SQL1**
   - Authentication: **Windows Authentication**

Task 3: Manually trigger the Data Collection process.

1. Using the Microsoft SQL Server Management Studio window, manually start the Data Collection process for the following system Data Collection sets:
   - Disk Usage
   - Query Statistics
   - Server Activity

Task 4: Generate Data Collection reports.

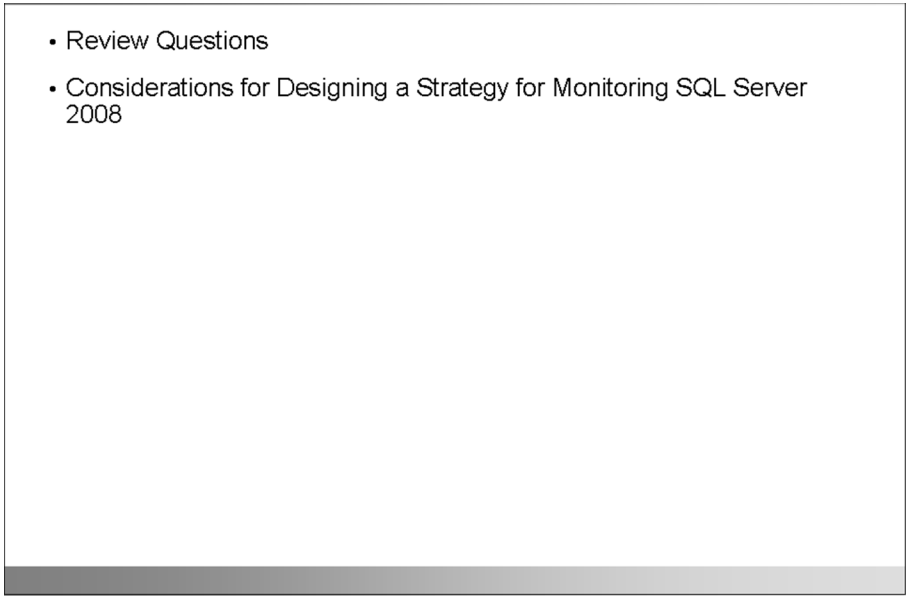
1. Using the Microsoft SQL Server Management Studio window, generate the following Data Collection reports:
   - Server Activity History
   - Disk Usage Summary
   - Query Statistics History

Results: After completing this exercise, you should have started SQL Server Agent, configured Management Data Warehouse, enabled Data Collection, manually triggered the Data Collection process, and generated Data Collection reports.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Module Review and Takeaways



Review Questions

1. What should you use to monitor statements executed against your SQL Server database?
2. What is difference between system monitor and DMVs?
3. What is the consideration for using row-versioning?
4. What is the purpose of isolation level?
5. What is the best way to avoid deadlocks?
6. What is the purpose of management data warehouse?
7. Can you configure event collection to collect additional information?

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Considerations for Designing a Strategy for Monitoring SQL Server 2008

When designing a strategy for monitoring SQL Server 2008, consider the following:

- Use System Monitor to monitor the real-time performance of SQL Server.
- Use SQL Trace to collect detail workload of your SQL Server if you experience a performance problem on the server. You can then analyze the collected workload with the help of SQL Server Profiler.
- Configure filtering when using SQL Trace or SQL Profiler as there will be a lot of information collected.
- Use SQL Server Profiler to troubleshoot a deadlock because it would provide a graphical way to present information.
- Use DMVs to monitor performance.
- Avoid modifying the query execution plan with query hints unless it is really necessary.
- Use row versioning isolation-levels for better performance and isolations on multiple transactions.
- Ensure tempdb is allocated enough disk space and located in a fast disk when using row versioning.
- Change isolation level only if you want to change the behavior of how SQL server isolates different transactions.
- Use data collection for disk, CPU, and query statistics.
- Configure a dedicated management data warehouse, if you have large number of servers to manage.
- Create custom reports if information included in standard reports cannot meet your needs.
- Ensure SQL Server Agent is started when using event collection.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Module 9

Designing a Strategy for Content Distribution in SQL Server 2008

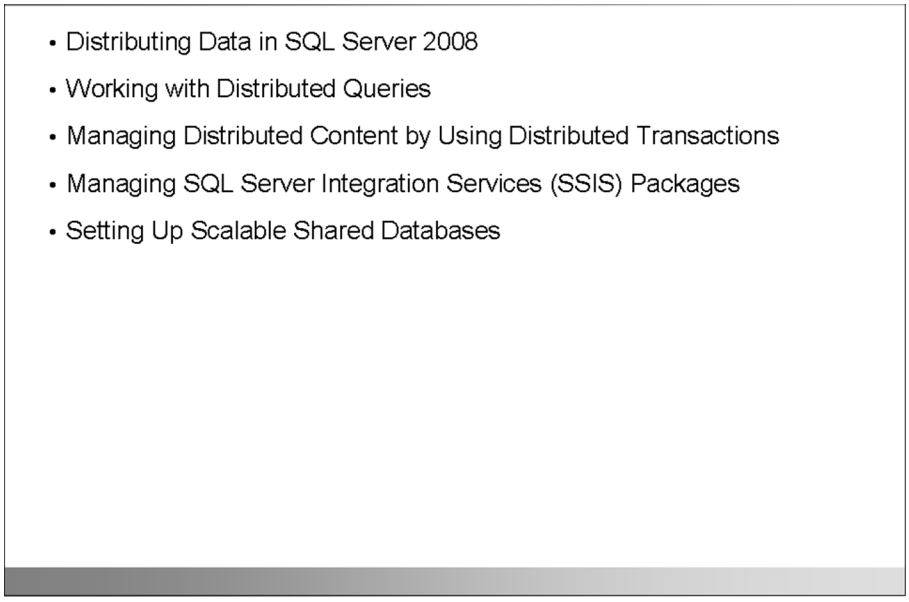
Contents:

Lesson 1: Distributing Data in SQL Server 2008 9-3
Lesson 2: Working with Distributed Queries 9-12
Lesson 3: Managing Distributed Content by Using Distributed Transactions 9-27
Lesson 4: Managing SSIS Packages 9-38
Lesson 5: Setting Up Scalable Shared Databases 9-60
Lab 9: Designing a Strategy for Content Distribution in SQL Server 2008 9-74
Module Review and Takeaways 9-82

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Module Overview



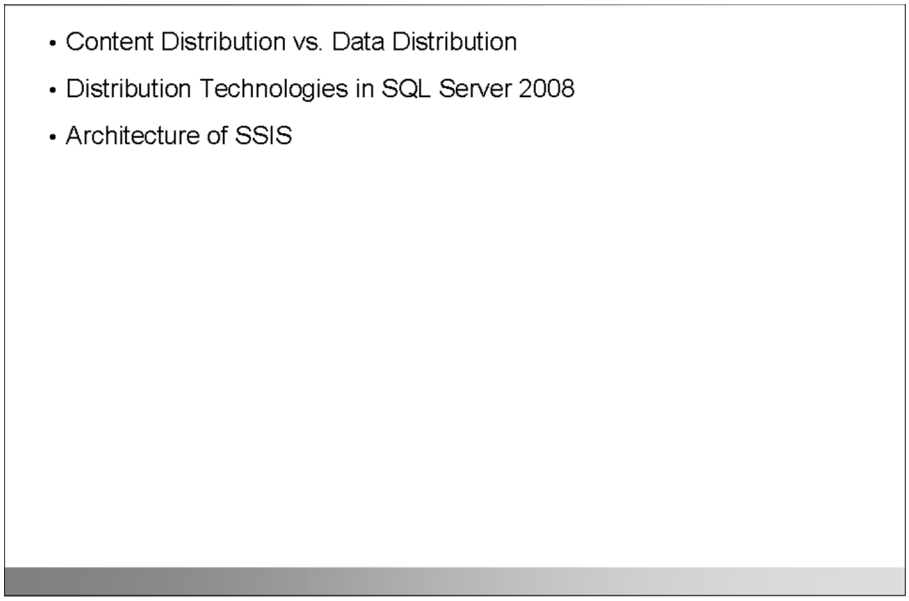
Organizations use many applications to manage business processes and provide new services. Consequently, the number of application servers and data storage servers that are required to support those applications has increased. Therefore it becomes very important to manage the content distribution among multiple servers. Microsoft™ SQL Server™ 2008 provides a number of technologies that help an organization to distribute information between multiple servers. SQL Server also includes different technologies for managing transactions across multiple servers and for consolidating the results into a single server.

In this module, you will learn about data distribution and about how to query data from heterogeneous data sources by using distributed queries. You will also learn about how to manage distributed transactions. You will learn about how to consolidate data by using an SSIS package. Finally, you will learn about scalable shared databases which help multiple servers to have the same copy of a database.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Lesson 1
Distributing Data in SQL Server 2008



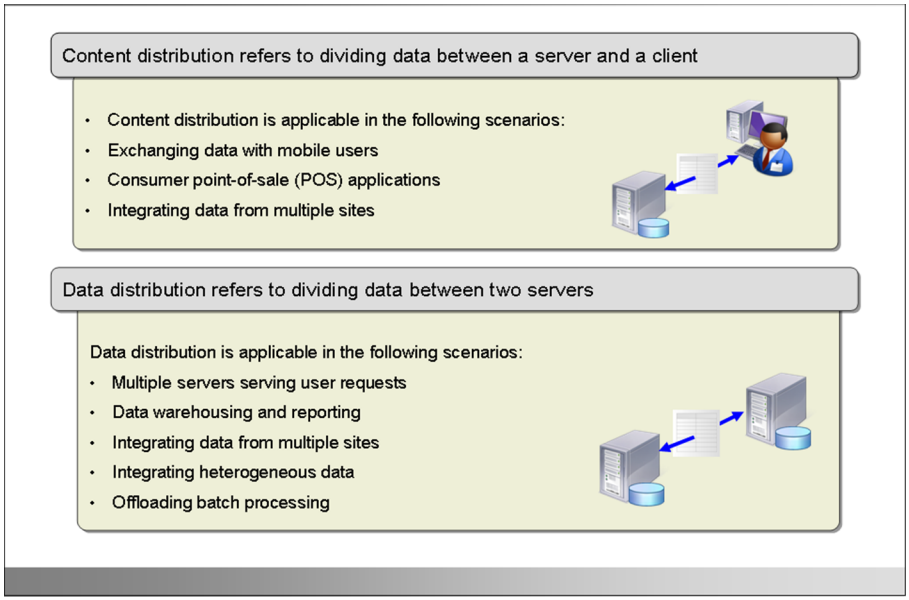
Data distribution involves moving or copying data between different servers. You can distribute data to different locations and to remote or mobile users over local and wide area networks, dial-up connections, wireless connections, and the Internet. Usually these tasks are integrated into the overall business process. SQL Server Integration Services (SSIS) helps perform data distribution across various servers.

In this lesson, you will learn about the various data distribution methods and the technologies that help in data distribution in SQL Server.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Content Distribution vs. Data Distribution



Key Points

Content distribution refers to dividing data between a server and a client. Data distribution refers to dividing data between two servers.

Content Distribution

In content distribution, data is typically replicated between servers and clients to support the following applications:

- **Exchanging data with mobile users**. Many applications, such as customer relationship management (CRM), sales force automation (SFA), and field force automation (FFA), require data to be available to remote users, including sales people and delivery drivers. These applications require data to be replicated between the server and the client.
- **Consumer point-of-sale (POS) applications**. POS applications, such as checkout terminals and ATM machines, require data to be replicated from remote sites to a central site.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

- **Integrating data from multiple sites**. Applications often integrate data from multiple sites. For example, an application that is used in both the business office and regional offices might require data to flow in one or both directions between the business office and the regional offices.

Data Distribution

In data distribution, data is typically replicated between servers to support the following applications and requirements:

- **Multiple servers serving user requests**. You can maintain updated copies of data so that a read activity scales across multiple servers. The redundancy resulting from maintaining multiple copies of the same data is crucial during planned and unplanned system maintenance.
- **Data warehousing and reporting**. Data warehouse and reporting servers often use data from online transaction processing (OLTP) servers. Use replication to move data between OLTP servers, and reporting and decision support systems.
- **Integrating data from multiple sites**. Data is often consolidated from remote offices into a single database at a central office. Similarly, data can be replicated out to remote offices.
- **Integrating heterogeneous data**. Some applications depend on data being sent to or from databases other than SQL Server. Use replication to integrate data from non-SQL Server databases.
- **Offloading batch processing**. Batch operations are often too resource-intensive to run on an OLTP server. Use replication to offload processing to a dedicated batch processing server.

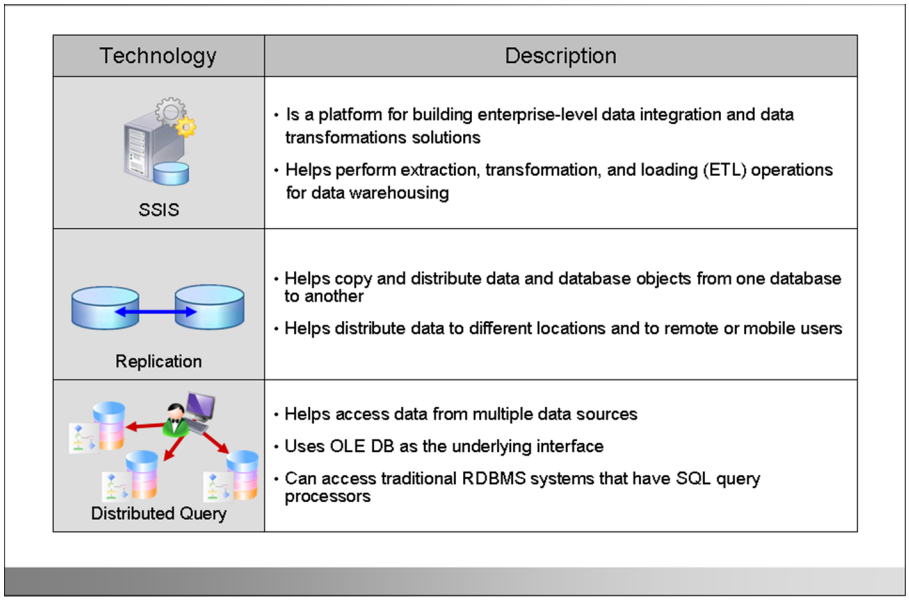
Source:

- http://msdn.microsoft.com/en-us/library/ms152565.aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Distribution Technologies in SQL Server 2008



Key Points

SQL Server 2008 provides technologies such as SSIS, replication, and distributed query to help you perform data distribution tasks.

SSIS

The SSIS service is a management service used for administering SSIS packages. It consists of application programming interfaces (APIs) that help program the SSIS object model. SSIS is a platform for building enterprise-level data integration and data transformations solutions. It helps perform extraction, transformation, and loading (ETL) operations for data warehousing in business intelligence (BI) scenarios.

You can use SSIS to solve complex business problems by copying or downloading files, sending e-mail messages in response to events, updating data warehouses, cleaning and mining data, and managing SQL Server objects and data. The packages can work alone or in collaboration with other packages to meet complex business needs. SSIS can extract and transform data from a wide variety of sources

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

such as XML data files, flat files, and relational data sources, and then load the data into one or more destinations.

SSIS includes a rich set of built-in tasks and transformation tools for developers to build logic in packages; and the Integration Services service is used to run and manage packages. You can use the graphical Integration Services tools to create solutions without writing a single line of code; or you can program the extensive Integration Services object model to create packages programmatically and code custom tasks and other package objects.

You can use SSIS to build high-performance data integration and workflow solutions. It helps perform ETL operations for data warehousing in BI scenarios.

SSIS includes graphical tools and wizards for building and debugging packages. Tasks perform workflow functions such as FTP operations, SQL statement execution, and e-mail messaging. SSIS consists of data sources and destinations for extracting and loading data. It also consists of transformations that clean, aggregate, merge, and copy data.

Replication

Replication helps copy and distributes data and database objects from one database to another and then synchronize between databases to maintain consistency. Using replication, you can distribute data to different locations and to remote or mobile users over local and wide area networks, dial-up connections, wireless connections, and the Internet.

The types of replication for use in distributed applications by SQL Server include:

- Transactional replication
- Merge replication
- Snapshot replication

With these three types of replication, SQL Server provides a powerful and flexible system for synchronizing data across an enterprise.

The replication type that you choose for an application depends on many factors, such as the physical environment, the type and quantity of data to be replicated, and whether the data is updated at the Subscriber. The physical environment includes the number and

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

location of computers that are involved in replication and whether these computers are clients or servers.

Distributed Query

You can use distributed queries to access data from multiple data sources. SQL Server supports distributed queries by using Object Linking and Embedding Database (OLE DB).

The OLE DB provider displays data in tabular objects called rowsets. SQL Server allows rowsets from OLE DB providers to be referenced in T-SQL statements as if they were SQL Server tables.

You can reference tables and views in external data sources directly in SELECT, INSERT, UPDATE, and DELETE T-SQL statements. Because distributed queries use OLE DB as the underlying interface, distributed queries can access traditional RDBMS systems that have SQL query processors, and can also access data managed by data sources of varying capabilities. As long as the software that owns the data exposes the data in a tabular rowset through an OLE DB provider, the data can be used in distributed queries.

Note: Using distributed queries in SQL Server is similar to the linked table functionality through Open Database Connectivity (ODBC). This functionality, which was supported previously by Microsoft Access™, is now built into SQL Server by using OLE DB as the interface to external data.

Question: What is the difference between distributed query and replication for distributing data across multiple servers?

Question: What is the difference between SSIS and replication for distributing data across multiple servers?

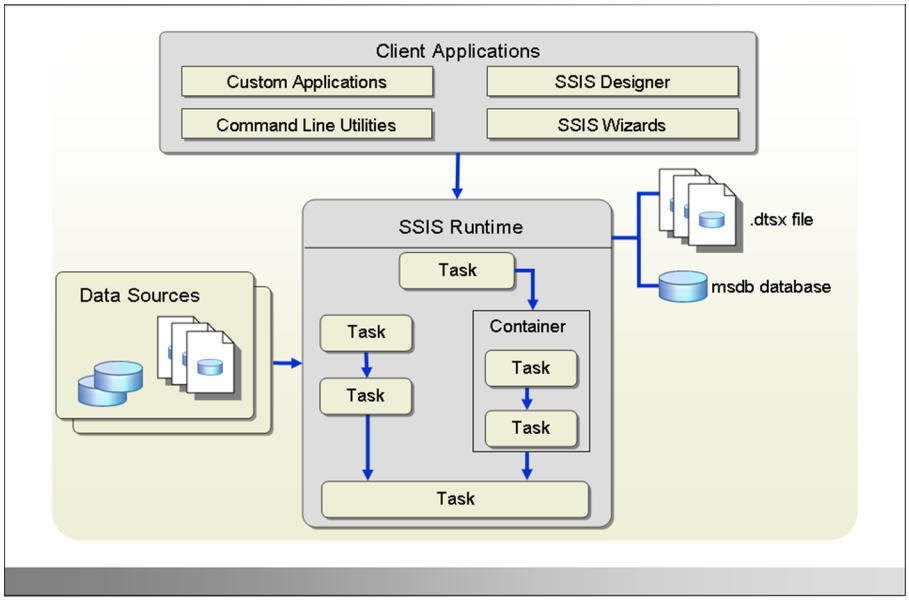
Source:

- http://msdn.microsoft.com/en-us/library/ms166352(SQL.90).aspx
- http://msdn.microsoft.com/en-us/library/ms188721.aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Architecture of SSIS



Key Points

SSIS is a platform for building high performance data integration and workflow solutions, including extraction, transformation, and loading (ETL) operations for data warehousing.

SSIS includes graphical tools and wizards for building and debugging packages; tasks for performing workflow functions such as FTP operations, SQL statement execution, and e-mail messaging; data sources and destinations for extracting and loading data; transformations for cleaning, aggregating, merging, and copying data; a management service, the SSIS service, for administering SSIS packages; and application APIs for programming the SSIS object model.

The architecture of SSIS consists of various components, such as the SSIS Designer, containers and tasks.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

SSIS Designer

SSIS Designer is a graphical tool that you can use to create and maintain Integration Services packages. SSIS Designer is available in Business Intelligence Development Studio as part of an Integration Services project.

You can use SSIS Designer to:

- Construct the control flow in a package.
- Construct the data flow in a package.
- Add event handlers to the package and package objects.
- View the package content.
- View the execution progress of the package at run time.

Run-time Engine

The SSIS run-time engine controls the management and execution of packages by implementing the infrastructure that enables the execution order, logon, variables, and event handling. Programming the SSIS run-time engine allows developers to automate the creation, configuration, and execution of packages, and create custom tasks and other extensions.

Run-time Executables

The SSIS run-time executables are packages, containers, tasks, and event handlers that are part of SSIS. Run-time executables also include custom tasks that you develop.

Container

Containers are objects in SSIS that provide structure to packages, and services to tasks. Containers support repeating control flows in packages, and group tasks and containers into meaningful units of work. Containers can include other containers, in addition to tasks.

Task

Tasks are control flow elements that define units of work that a package control flow performs. An SSIS package is made up of one or more tasks. If the package contains more than one task, they are connected and sequenced in the control flow by precedence constraints.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

SSIS Object Model

The SSIS object model includes managed APIs for creating custom components to be used in packages or custom applications that create, load, run, and manage packages. The developer can write custom applications or custom tasks or transformations by using any common language runtime (CLR) compliant language.

SSIS Service

The SSIS run-time engine controls the management and execution of packages, by implementing the infrastructure that enables execution order, logon, variables, and event handling. Programming the SSIS run-time engine lets developers automate the creation, configuration, and execution of packages and create custom tasks and other extensions.

The SSIS service allows you to use SQL Server Management Studio (SSMS) to monitor running Integration Services packages and manage the storage of packages.

SQL Server Import and Export Wizard

The SQL Server Import and Export Wizard can copy data to and from any data source for which a managed .NET Framework data provider or a native OLE DB provider is available. This wizard also offers the simplest method to create an SSIS package that copies data from a source to a destination.

Command Prompt Utilities

SSIS includes additional tools, wizards, and Command Prompt utilities for running and managing Integration Services packages.

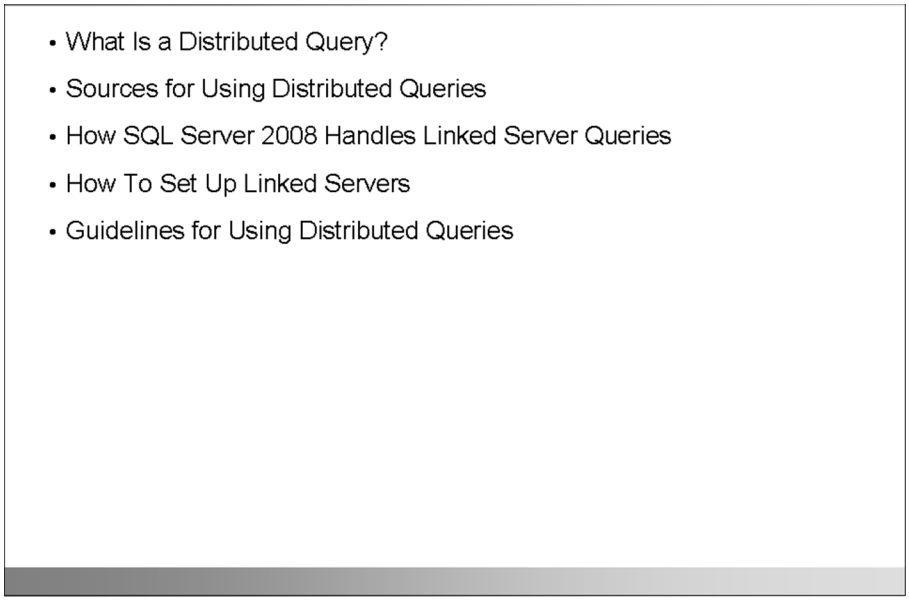
Source:

- http://technet.microsoft.com/en-us/library/ms166363(SQL.90).aspx
- http://msdn.microsoft.com/en-us/library/ms137973.aspx
- msdn.microsoft.com/en-us/library/bb522498.aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Lesson 2
Working with Distributed Queries



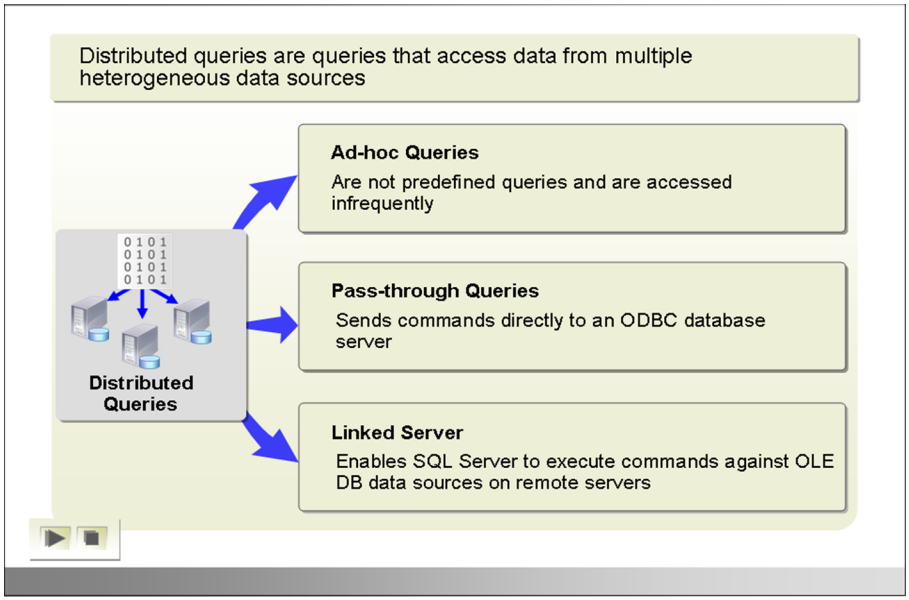
With multiple data sources, querying data has been a challenge. Distributed queries access data from multiple heterogeneous data sources. You can store these data sources on the same computer or different computers. SQL Server 2008 supports distributed queries through OLE Database (OLE DB).

In this lesson, you will learn about distributed queries, and how to query data by using a linked server. You will also learn about the sources for a distributed query and to set up a linked server. Finally, you will learn the guidelines for using distributed queries.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

What Is a Distributed Query?



Key Points

SQL Server users can use distributed queries to access:

- Distributed data stored in multiple instances of SQL Server
- Heterogeneous data stored in various relational and nonrelational data sources accessed by using an OLE DB provider

The following are the types of distributed queries:

Ad-hoc Query

By default, SQL Server does not allow you to use ad-hoc distributed queries by using the OPENROWSET and the OPENDATASOURCE functions. When this option is set to 1, SQL Server allows ad hoc access. When this option is not set or is set to 0, SQL Server does not allow ad hoc access.

```
sp_configure 'show advanced options', 1
reconfigure
```

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.


```
go

sp_configure 'Ad Hoc Distributed Queries', 1

reconfigure
```

Ad-hoc distributed queries use the OPENROWSET and OPENDATASOURCE functions to connect to remote data sources that use OLE DB. You should consider using the OPENROWSET and the OPENDATASOURCE functions only to reference OLE DB data sources that are accessed infrequently. When data sources are frequently accessed, you should consider defining a linked server because the OPENDATASOURCE and the OPENROWSET functions do not provide all the functionality of a linked server. This includes security management and the ability to query catalog information. You can specify the OPENROWSET or OPENDATASOURCE functions with the information needed to connect to the linked server. You can then reference the rowset the same way you reference a table in T-SQL statements.

```
USE master

GO

SELECT *

FROM OPENROWSET('Microsoft.ACE.OLEDB.12.0',

        'D:\Demofiles\DemoDB.accdb' ;

      'admin';'' ,

        Person)
```

You can use OPENROWSET with any OLE DB provider that returns a rowset and can be used where a table or view is referenced in a T-SQL statement. You can specify OPENROWSET with the following:

- All the information required to connect to the OLE DB data source
- The name of an object that will generate a rowset or a query that will generate a rowset

OPENDATASOURCE provides connection information as part of a four-part object name. This function supports only OLE DB providers that expose multiple rowsets by using the catalog.schema.object notation. You can use OPENDATASOURCE in the same location in the T-SQL syntax where you use a linked server name. You can specify OPENDATASOURCE with the following:

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

- The name registered as the PROGID of the OLE DB provider that is used to access the data source
- A connection string that specifies various connection properties to be passed to the OLE DB provider

By default, SQL Server does not allow Distributed Ad-hoc queries by using OPENROWSET and OPENDATASOURCE. When this option is set to 1, SQL Server allows improvised access. When this option is set to 0, SQL Server does not allow ad-hoc access.

Pass-Through Query

SQL Server sends pass-through queries as uninterpreted query strings to an OLE DB data source. The query must be in a syntax that the OLE DB data source will accept. A T-SQL statement uses the results from a pass-through query as a regular table reference.

You can use the OPENQUERY statement to run a pass-through query on the specified linked server. This server is an OLE DB data source. You can reference OPENQUERY in the FROM clause of a query as a table name. You can also reference OPENQUERY as the target table of an INSERT, UPDATE, or DELETE statement. This is subject to the capabilities of the OLE DB provider. Although the query may return multiple result sets, OPENQUERY returns the first one only.

The following are ways to generate a rowset from an OLE DB provider:

- Reference an object in the data source that the provider can expose as a tabular rowset. All providers support this capability.
- Send the provider a command that the provider can process and expose the results of the command as a rowset. This capability requires that the provider support the OLE DB command object and all its mandatory interfaces.
- When an OLE DB provider supports the command object, you can use the following T-SQL functions to send the provider commands, called pass-through queries:
- OPENQUERY sends a command string to an OLE DB data source by using a linked server name.
- OPENROWSET and OPENDATASOURCE support sending a command string to an OLE DB data source. The resulting rowset can be referenced by using an ad-hoc name.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

- OPENROWSET and OPENQUERY support pass-through queries in which the query string contains a stored procedure call only when the query has one of the following characteristics:

  Executes a single, static SELECT statement

  Does not use any dynamic EXECUTE statements

  Does not perform any data modification operations

The OLE DB specification does not define one command language that must be used by all OLE DB providers. OLE DB providers can support any command language that is related to the data they expose. OLE DB providers that expose the data in relational databases generally support the SQL language. Other types of providers, such as those exposing the data in an e-mail file or network directory, generally support a different language.

Linked Server

A linked server configuration enables SQL Server to run commands against OLE DB data sources on remote servers. Linked servers offer the following advantages:

- Ability to access remote servers
- Ability to issue distributed queries, updates, commands, and transactions on heterogeneous data sources across the enterprise
- Ability to address diverse data sources similarly
- A linked server includes an OLE DB provider and an OLE DB data source to specify where and how to connect the server.

An OLE DB provider is a dynamic-link library (DLL) that manages and interacts with a specific data source. An OLE DB data source identifies the database, which you want to access through OLE DB. Although data sources queried through linked server definitions are ordinarily databases, OLE DB providers exist for a variety of files and file formats. These include text files, spreadsheet data, and the results of full-text content searches.

For example, if a linked server name of DeptSQLSrvr is defined against another instance of SQL Server, the following statement references a table on that server.

```
USE master;
```

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.


```
GO
EXEC sp_addlinkedserver
   @server='<VAN-SQL>',
   @srvproduct='',
   @provider='SQLNCLI',
   @datasrc='<NYC-SQL1>'

-- VAN-SQL represents the remote server and NYC-SQL1
represents the host server.
EXEC sp_addlinkedsrvlogin '<VAN-SQL>', 'true'
GO
SELECT * FROM [VAN-SQL].QuantamCorp.sys.tables
```

Question: When should you consider using OPENROWSET over linked server?

Question: Does SQL Server enable support for distributed ad-hoc query by default?

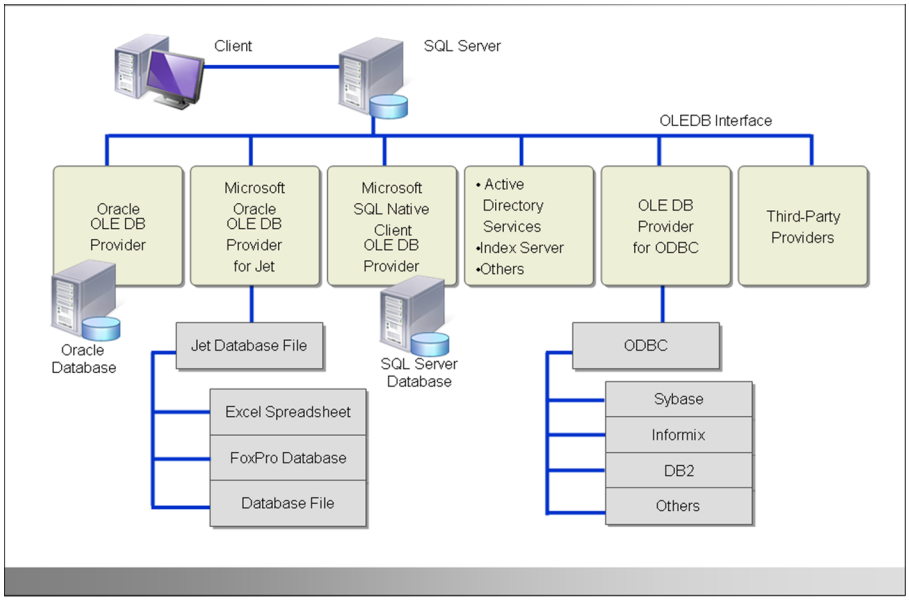
Source:

- http://msdn.microsoft.com/en-us/library/ms190448.aspx
- http://msdn.microsoft.com/en-us/library/ms188279(SQL.90).aspx
- http://msdn.microsoft.com/en-us/library/ms187569.aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Sources for Using Distributed Queries



Key Points

In SQL Server, you can reference heterogeneous OLE DB data sources in T-SQL statements by using the linked server name or the improvised computer name.

To access data from an OLE DB data source, you must provide SQL Server with the following information:

- The name of the OLE DB provider that exposes the data source
- Any information the OLE DB provider must have to locate the source of the data

You can query an OLE DB data source by using an object name directly, such as a table name, or you can use a SQL statement that is compatible with the data source. The objects that you can expose as rowsets are known as remote tables. The queries that generate rowsets are known as pass-through queries.

Optionally, you can supply SQL Server with valid logon IDs for the OLE DB data source. SQL Server supports referencing heterogeneous

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

OLE DB data sources in T-SQL statements by using the linked server name or the improvised computer name.

Distributed queries allows users to access another data source (for example, files, nonrelational data sources such as Active Directory®, and so on) by using the security context of the Microsoft Windows® account under which the SQL Server service is running. SQL Server impersonates the login appropriately for Windows logins; however, that is not possible for SQL Server logons. This can potentially allow a distributed query user to access another data source for which they do not have permissions, but the account under which the SQL Server service is running does have permissions. Use sp_addlinkedsrvlogin to define the specific logons that are authorized to access the corresponding linked server. This control is not available for ad-hoc names.

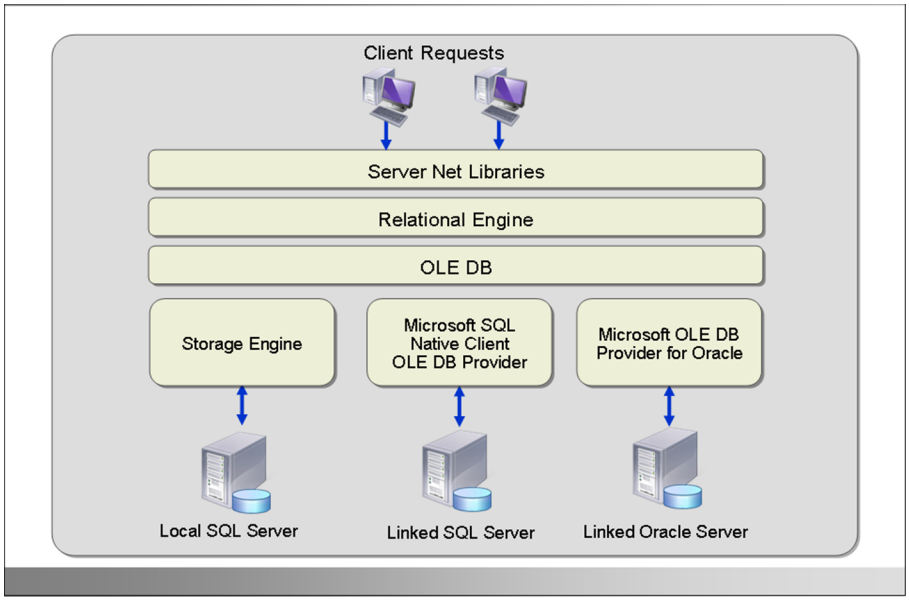
When possible, SQL Server pushes relational operations such as joins, restrictions, projections, sorts, and group by operations to the OLE DB data source. SQL Server does not default to scanning the base table into SQL Server and performing the relational operations itself. SQL Server queries the OLE DB provider to determine the level of SQL grammar it supports, and, based on that information, pushes as many relational operations as possible to the provider.

Source:

- http://technet.microsoft.com/en-us/library/ms191277(SQL.90).aspx
- http://msdn.microsoft.com/en-us/library/ms188313.aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

How SQL Server 2008 Handles Linked Server Queries



Key Points

You can use linked servers to handle distributed queries. When a client application runs a distributed query through a linked server, SQL Server parses the command and sends requests to OLE DB. The rowset request may be in the form of running a query against the OLE DB provider or opening a base table from the provider.

For a data source to return data through a linked server, the OLE DB provider for that data source must be present on the same server as the instance of SQL Server.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

When you use a third-party OLE DB provider, the account under which the SQL Server service runs must have read and execute permissions for the directory, and all subdirectories, in which the provider is installed. SQL Server uses OLE DB to communicate between the relational engine and the storage engine. The relational engine breaks down each Transact-SQL statement into a series of operations on simple OLE DB rowsets opened by the storage engine from the base tables. This means the relational engine can also open simple OLE DB rowsets on any OLE DB data source.

When you are setting up a linked server, you should register the connection information and data source information with SQL Server. After registration, that data source can be referred to with a single logical name.

By using stored procedures and catalog views, you can:

- Create a linked server definition by running sp_addlinkedserver.
- View information about the linked servers defined in a specific instance of SQL Server by running a query against the sys.servers system catalog views.
- Delete a linked server definition by running sp_dropserver. You can also use this stored procedure to remove a remote server.

You can define linked servers by using SQL Server Management Studio (SSMS). In Object Explorer, you should right-click Server Objects, select New, and select Linked Server. You can delete a linked server definition by right-clicking the linked server name and then selecting Delete.

When you run a distributed query against a linked server, you should include a qualified, four-part table name for each data source to query. This four-part name should be in the form, linked_server_name.catalog.schema.object_name.

The following example creates a linked server named SEATTLESales. The product name is SQL Server, and no provider name is used.

```
USE master;

GO

EXEC sp_addlinkedserver

   '<NYC-SQL1>',
```

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.


```
   N'SQL Server'
GO
```

To create a linked server definition by using the OLE DB provider for Jet to access an Excel® spreadsheet in the 1997–2003 formats, first create a named range in Excel by specifying the columns and rows of the Excel worksheet to select. You can then reference the name of the range as a table name in a distributed query.

```
EXEC sp_addlinkedserver 'ExcelSource',
   'Jet 4.0',
   'Microsoft.Jet.OLEDB.4.0',
   'D:\Demofiles\People.xls, NULL, Excel 5.0'
GO
```

To connect to an Excel spreadsheet in the Excel 2007 format, use the ACE provider.

```
EXEC sp_addlinkedserver @server = N'ExcelDataSource',
@srvproduct=N'ExcelData',
@provider=N'Microsoft.ACE.OLEDB.12.0',
@datasrc=N'D:\Demofiles\People.xlsx',
@provstr='EXCEL 12.0' ;
```

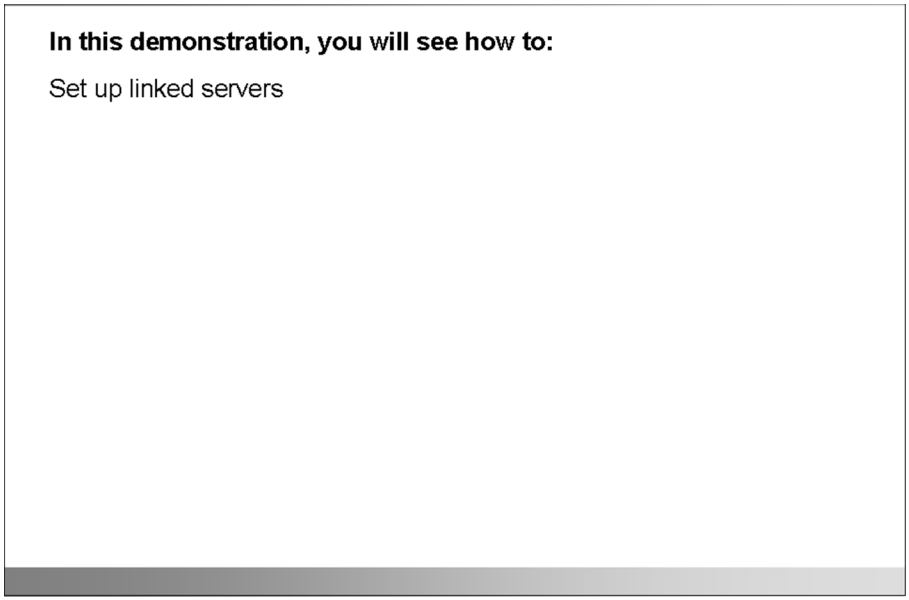
Source:

- http://msdn.microsoft.com/en-us/library/ms188279.aspx
- http://technet.microsoft.com/en-us/library/ms191277(SQL.90).aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Demonstration: How To Set Up Linked Servers



Key Points

The steps to set up linked servers are:

1. On the **Start** menu, point to **All Programs**, click **Microsoft SQL Server 2008**, and then click **SQL Server Management Studio**. The Microsoft SQL Server Management Studio window appears.
2. In the Microsoft SQL Server Management Studio window, by using the **Connect to Server** dialog box, connect to the default instance of SQL Server Database Engine.
3. In Windows Explorer, browse to the **D:\Demofiles\Mod09** folder, and then double-click the **M9_Demo1.sql** file.
4. In the Microsoft SQL Server Management Studio window, select the **Add an instance of SQL Server as a linked server** code, and then click **Execute**.
5. In the Microsoft SQL Server Management Studio window, select the **Enable linked server for self mapping of credentials** code, and then click **Execute**.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

6. In the Microsoft SQL Server Management Studio window, select the **Verify if the linked table is working** code, and then click **Execute**.
7. In the Microsoft SQL Server Management Studio window, select the **Add and link a linked server record to an excel file and retrieve the information of the excel file** code, and then click **Execute**.

Question: What does a linked server support?

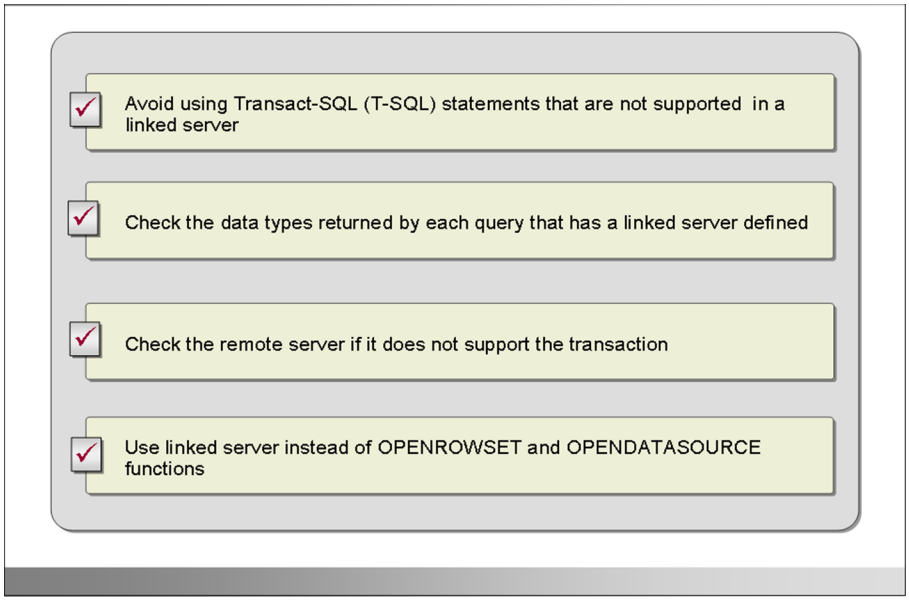
Question: How would you refer to a linked server in a query?

Question: What credential does SQL Server use to connect to a linked server?

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Guidelines for Using Distributed Queries



Key Points

You must consider the following guidelines when using a distributed query:

- **Avoid using some T-SQL functionalities that are not supported in a linked server**. Some of them are a table with an XML column and the READTEXT, WRITETEXT, and UPDATETEXT statements. Stored procedures are only supported if the remote server is SQL Server
- **Check the data types returned by each query that has a linked server defined**. Because data type conversion happens when a remote query runs. Data types are converted between OLE DB and SQL Server system data types by mapping data.
- **Check the remote server if it does not support the transaction**. It would be set as read-only. By using distributed transactions, SQL Server ensures that a transaction involving multiple nodes is committed or rolled back in all nodes. If the

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

provider does not support participating in a distributed transaction, only read-only operations are allowed against that provider when inside a transaction.

- **Use a linked server, instead of OPENROWSET and OPENDATASOURCE functions**. OPENROWSET and OPENDATASOURCE functions may impose reduced security and additional performance overhead.

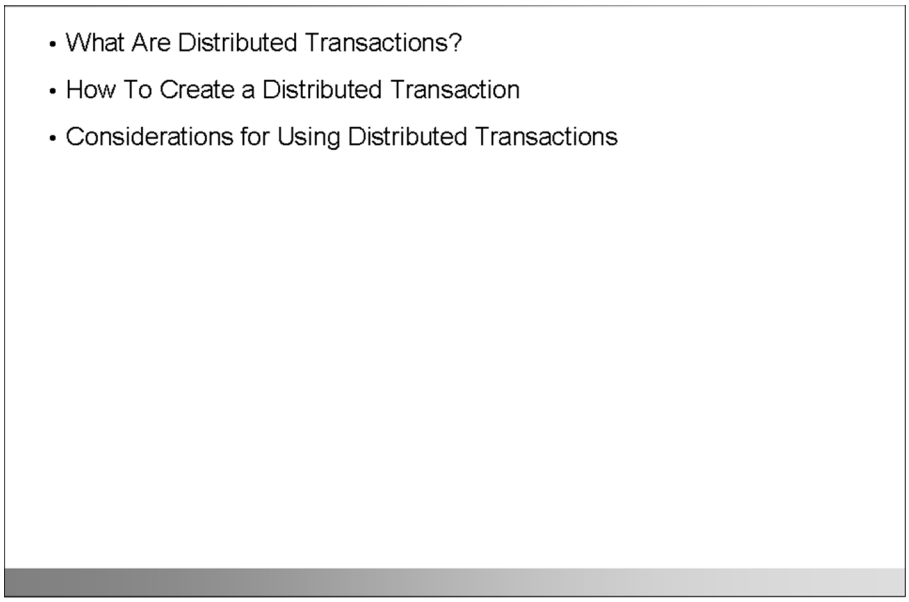
Source:

- http://msdn.microsoft.com/en-us/library/ms189266(SQL.90).aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Lesson 3
Managing Distributed Content by Using Distributed Transactions



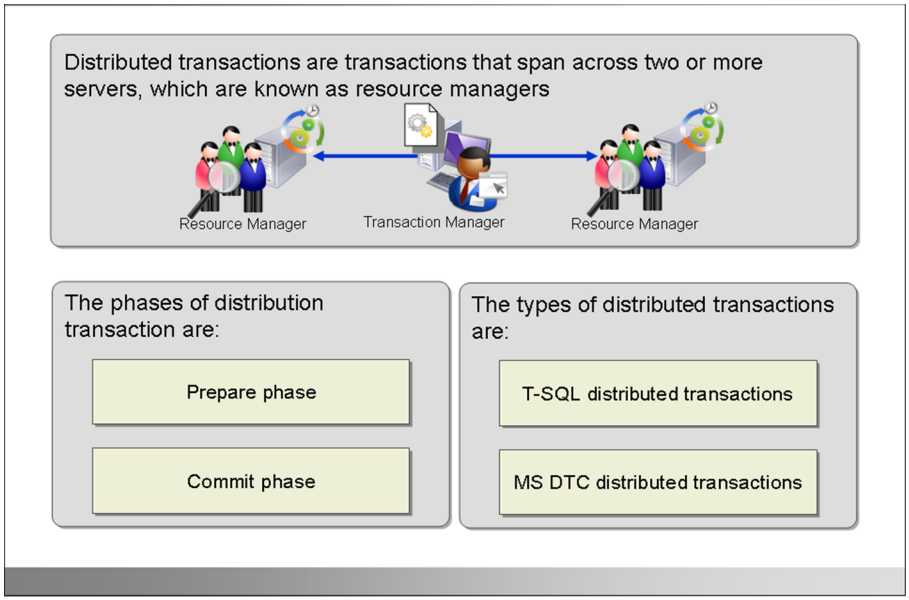
A transaction within a single SQL Server that spans two or more databases is a distributed transaction. SQL Server, however, manages the distributed transaction internally; to the user it operates as a local transaction. Distributed transactions extend the benefits of transactions to applications that update distributed data. Object OLE DB and ODBC also contain support at the API level for managing distributed transactions.

In this lesson, you will learn about distributed transactions and how to perform a distributed transaction.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

What Are Distributed Transactions?



Key Points

A transaction within a single instance of the Database Engine that spans across two or more databases is called a distributed transaction. A distributed transaction updates data on two or more servers. These servers are called resource managers. A server component known as the transaction manager coordinates the management of the transaction between resource managers. Each instance of the SQL Server Database Engine can operate as a resource manager in distributed transactions coordinated by transaction managers, such as MS DTC.

Distributed transactions extend the benefits of transactions to applications that update distributed data. At the application level, a distributed transaction is managed in the same way as the local transaction. At the end of the transaction, the application requests the transaction to be committed or rolled back. The transaction manager must manage a distributed commit differently to minimize the risks associated with a network failure. In the event of a network failure, some resource managers may successfully commit the

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

transaction, while others may roll back the transaction. To prevent such errors, the commit process is managed in two phases: prepare phase and commit phase. This distributed commit process is also known as a two-phase commit (2PC).

Prepare Phase

The following are the tasks of the prepare phase:

- When the transaction manager receives a commit request, it sends a prepare command to all resource managers involved in the transaction.
- Each resource manager makes the transaction durable and all log buffers for the transaction are written back to the disk.
- As each resource manager completes the prepare phase, it returns success or failure of the prepare phase to the transaction manager.

Commit Phase

The following are the tasks of the commit phase:

- When the transaction manager receives a notification of a successful prepare phase from all resource managers, it sends commit commands to each resource manager involved in the transaction.
- The resource managers then complete the commit phase.
- If all resource managers report a successful commit, the transaction manager sends a success notification to the application.
- If any resource manager reports a failure of the prepare phase, the transaction manager sends a rollback command to each resource manager and indicates the failure of the commit phase to the application.

The different types of distributed transactions are:

T-SQL Distributed Transactions

A T-SQL script or application connection runs a T-SQL statement that starts a distributed transaction. In this transaction, SQL Server becomes the controlling server that runs the statement. The script or application then runs either distributed queries against linked servers or remote stored procedures against remote servers. When

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

distributed queries and remote procedure calls are made, the controlling server automatically calls MS DTC to enlist the linked and remote servers in the distributed transaction.

When the script or application issues either a COMMIT or ROLLBACK statement, the controlling SQL Server calls MS DTC to manage the two-phase commit process, or to notify the linked and remote servers to roll back their transactions.

You can start distributed transactions in T-SQL by using the BEGIN DISTRIBUTED TRANSACTION statement or by running a distributed query. Start an explicit distributed transaction using the BEGIN DISTRIBUTED TRANSACTION statement. You can also run a distributed query against a linked server. The instance of the Database Engine you have connected to calls MS DTC to manage the distributed transaction with the linked server. You can also call remote stored procedures on a remote instance of the Database Engine as part of the distributed transaction.

MS DTC Distributed Transactions

Applications written by using OLE DB, ODBC, ActiveX® Data Objects (ADO), or DB-Library can use T-SQL distributed transactions by issuing T-SQL statements to start and stop T-SQL distributed transactions. OLE DB and ODBC also contain support at the API level for managing distributed transactions. OLE DB and ODBC applications can use these API functions to manage distributed transactions that include other Component Object Model (COM) resource managers that support MS DTC transactions other than the SQL Server Database Engine. They can also use the API functions to gain more control over the boundaries of a distributed transaction that includes several computers running instances of the Database Engine.

In MS DTC, the MS DTC transaction manager is responsible for coordinating DTC distributed transactions. Usually, one MS DTC transaction manager process is present on each computer However, the system administrator can choose to configure a computer to depend on an MS DTC transaction manager running on another system.

MS DTC resource management

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

The MS DTC transaction manager maintains an MS DTC log file on disk on its local system. The MS DTC log file captures the following information:

- The identity of the transaction.
- The resource managers that are enlisted with the transaction manager.
- The MS DTC transaction manager that imported the transaction to this node, if the transaction is imported to this node from another node.
- The transaction managers that are sent the exported transaction, if the transaction is exported from this node to other nodes.
- The outcome of the transaction and which resource managers and transaction managers have been told of the transaction outcome.

The information that is maintained in the MS DTC log file is retained for a short period of time. All information regarding the transaction can be discarded after the transaction commits or aborts and after all of the subordinate transaction managers and resource managers have been received the transaction outcome. Therefore, the DTC log file is implemented as a simple circular file: The MS DTC writes to the log file sequentially, and when it reaches the end of the log file, it continues writing at the beginning of the file, overwriting obsolete information.

Some of the information maintained in the MS DTC log file is retained for a long period of time. The MS DTC prevents overwriting and loss of this information by periodically copying the data to the current active part of the log file. The process of copying long-lived information to the log is referred to as check pointing. MS DTC performs check pointing periodically.

ODBC and OLE DB applications manage distributed transactions in the following ways:

ODBC Distributed Transactions

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

You can control local transactions at the ODBC API level by setting the connection attribute SQL_ATTR_AUTOCOMMIT to SQL_AUTOCOMMIT_OFF and then by calling the ODBC SQLEndTran function to commit or roll back each transaction.

To manage a distributed transaction in an ODBC application, you can use the MS DTC COM methods as follows:

- Call DtcGetTransactionManager to connect to MS DTC.
- Call ITransactionDispenser::BeginTransaction to start the distributed transaction and get a transaction object.
- For each ODBC connection participating in the distributed transaction, call the ODBC function SQLSetConnectAttr with fOption set to SQL_COPT_SS_ENLIST_IN_DTC and vParam holding the address of the transaction object from ITransactionDispenser::BeginTransaction.
- When the transaction is completed, instead of calling the ODBC SQLEndTran function, call the ITransaction::Commit or ITransaction::Rollback methods on the transaction object obtained from ITransactionDispenser::BeginTransaction.

OLE DB Distributed Transactions

The method used for controlling a distributed transaction in OLE DB is similar to controlling a local transaction.

To control a local transaction, an OLE DB consumer:

- Uses the ITransactionLocal::StartTransaction method to start the local transaction and get a transaction object.
- Calls the ITransaction::Commit or ITransaction::Rollback methods on the transaction object obtained by ITransactionLocal::StartTransaction.

To control a distributed transaction, the consumer:

- Calls DtcGetTransactionManager to connect to MS DTC.
- Calls ITransactionDispenser::BeginTransaction to start the distributed transaction and get a transaction object.
- Calls the ITransactionJoin interface of the distributed transaction object for each connection participating in the distributed transaction.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

- Calls the ITransaction::Commit or ITransaction::Rollback methods of the distributed transaction object to complete the transaction.

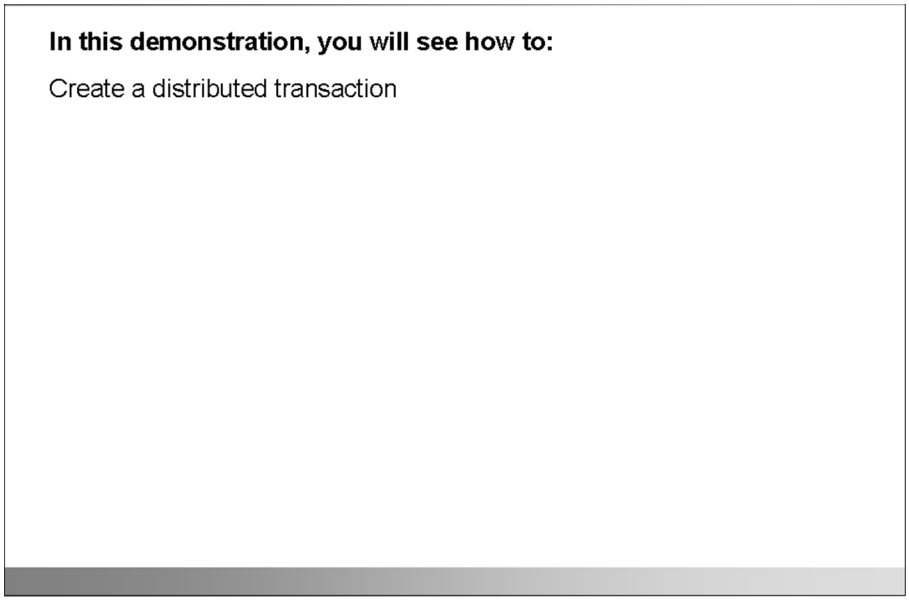
Source:

- http://msdn.microsoft.com/en-us/library/ms191440.aspx
- http://msdn.microsoft.com/en-us/library/aa213081(SQL.80).aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Demonstration: How To Create a Distributed Transaction



Key Points

The steps to create a distributed transaction are:

1. In Windows Explorer, browse to the **D:\Demofiles\Mod09** folder, and then double-click the **M9_Demo2.sql** file.
2. In the Microsoft SQL Server Management Studio window, in the Query Editor pane, select the **Create a distributed transaction** code, and then click **Execute**.
3. In the Microsoft SQL Server Management Studio window, in the Query Editor pane, select the **Retrieve database records from the local server** code, and then click **Execute**.
4. In the Microsoft SQL Server Management Studio window, in the Query Editor pane, select the **Retrieve database records from the remote server** code, and then click **Execute**.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

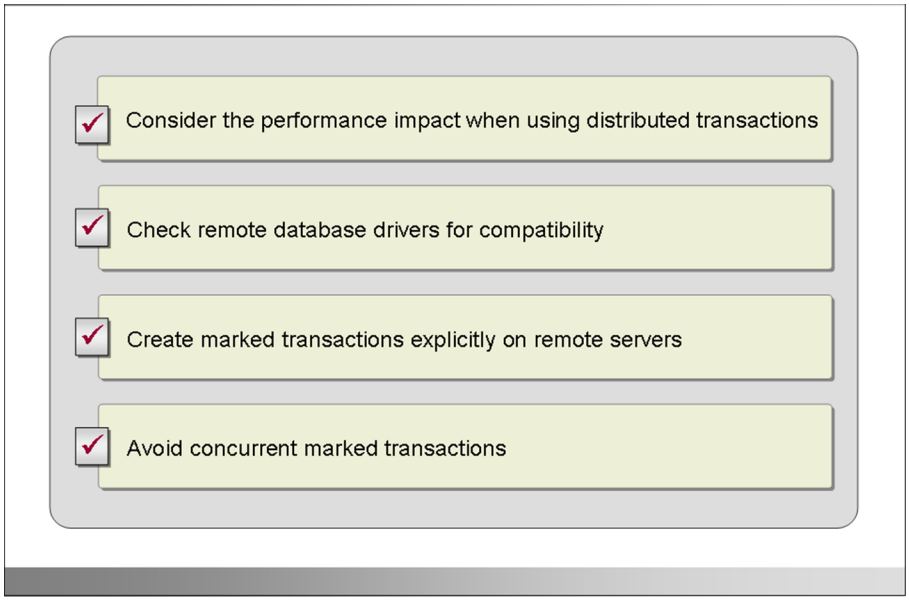
Question: What is the difference between the BEGIN DISTRIBUTED TRANSACTION statement and the BEGIN TRANSACTION statement?

Question: If you do not specify the BEGIN DISTRIBUTED TRANSACTION statement, what will be the outcome of the transaction, when you make changes to the data on a linked server?

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Considerations for Using Distributed Transactions



Key Points

You can perform real-time updates to multiple databases simultaneously by using distributed transactions. Every database should commit its part of the transaction for the entire transaction to be successful. In this operation, multiple communications happen between the originating server and the remote servers. If one of the servers does not respond to the prepare phase, the transaction will be refused, and all servers involved will roll back the transaction.

There are several other factors that might affect the successful completion of a distributed transaction. Therefore, it is important to take into account the following considerations when using distributed transactions:

- **Consider the performance impact when using distributed transaction**. A distributed transaction occurs when the individual statements in a transaction update databases on multiple servers. This involves many network communications between the originating server and the remote servers. In addition,

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

distributed transactions require external services such as MS DTC to help coordinate the transaction. Therefore, you should use distributed transactions to avoid the performance getting affected.

- **Check the remote database drivers for compatibility**. All remote database and OLE DB drivers do not work under distributed transaction. Check the latest Microsoft documentation on compatibility before using such drivers.
- **Create marked transactions explicitly on remote servers**. A transaction mark name is not automatically distributed to another server as the transaction spreads there. For the mark to spread to other servers, you need to write a stored procedure that contains a BEGIN TRAN name WITH MARK statement. You must then run the stored procedure on the remote server under the scope of the transaction in the originating server.
- **Avoid concurrent marked transactions**. You must avoid running concurrent marked transactions. It is rare but possible for the commit of a distributed marked transaction to deadlock with other distributed marked transactions that are committing at the same time. When this happens, the marking transaction is chosen as the deadlock victim and is rolled back. When this error occurs, the application can retry the marked transaction. When multiple marked transactions try to commit concurrently, there is a higher probability of deadlock.

Question: What should you consider when using marked transaction with distributed transactions?

Question: What is the key consideration for using distributed transaction in terms of performance?

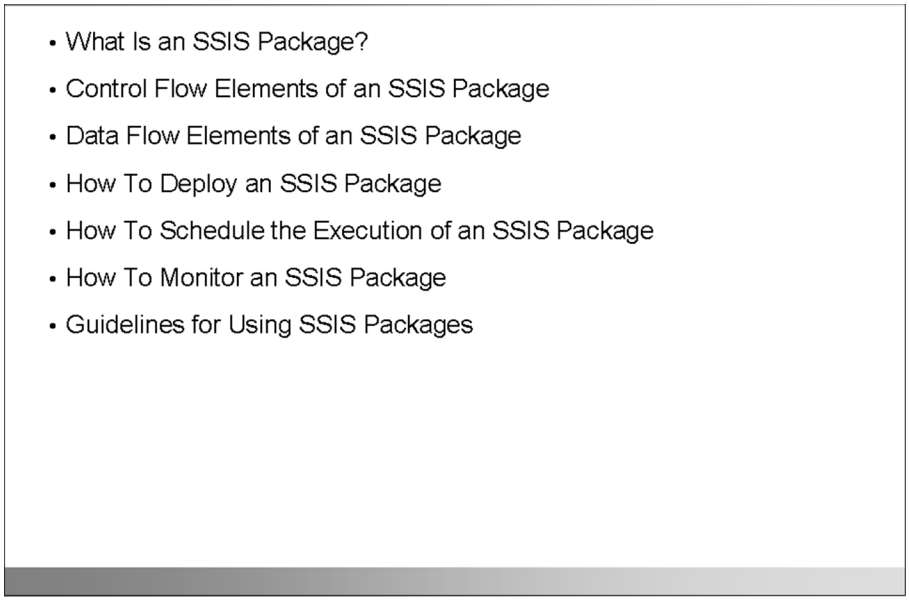
Source:

- http://msdn.microsoft.com/en-us/library/ms187014.aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Lesson 4
Managing SSIS Packages

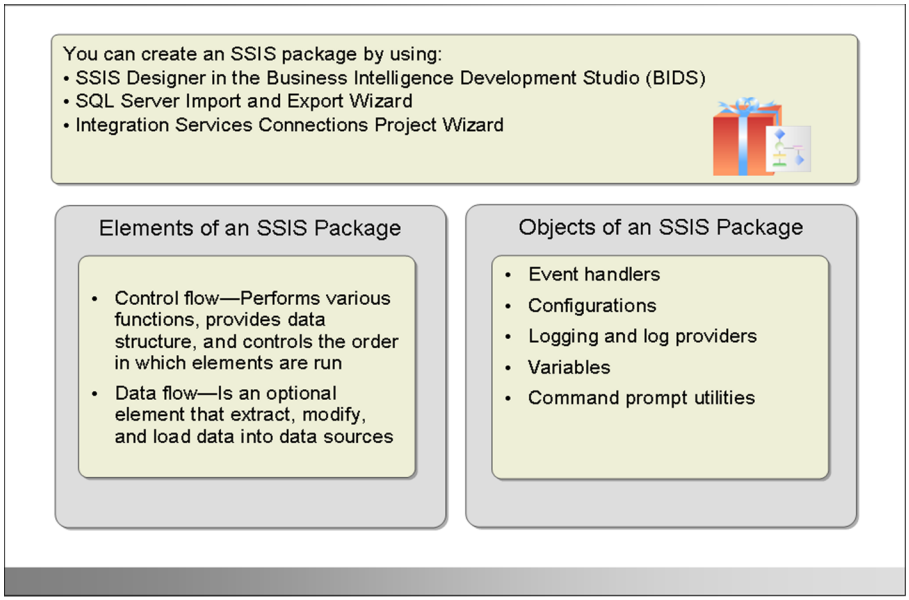


SSIS is a comprehensive data integration platform that can be used to transfer, mine, transform, and consolidate your information from heterogeneous sources and load it to multiple systems. An SSIS package is the object that implements the SSIS functionality to extract, transform, and load data. You can create an SSIS package by using the SSIS Designer in Business Intelligence Development Studio (BIDS).

In this lesson, you will learn about SSIS packages, the elements in the packages, and how to work with SSIS packages.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

What Is an SSIS Package?



Key Points

An SSIS package implements the SSIS functionality to extract, transform, and load data. You can create an SSIS package by using:

- SSIS Designer in the BIDS.
- SQL Server Import and Export Wizard.
- Integration Services Connections Project Wizard.

An SSIS package is created within the context of an Integration Services project. Before creating an SSIS package by using SSIS Designer, you need to open an existing Integration Services project or create a new project. When you create a new Integration Services project in BIDS, a new package is automatically added to the project.

SSIS Designer does not depend on the Integration Services service, and it is not necessary for the service to be running to create or modify packages in SSIS Designer.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Elements of an SSIS Package

A basic SSIS package includes Control flow and Data flow elements. Control flow elements perform various functions, such as providing data structure, and controlling the order in which elements are run. The main control flow elements are tasks, containers, and precedence constraints. There must be at least one control flow element in a package.

Data flow elements are optional elements that extract, modify, and load data into data sources. The main data flow elements are sources, transformations, and destinations. There is no minimum requirement of data flow elements in a package.

In addition to the Control and Data flow elements, SSIS packages can include the following objects that provide advanced features or extend the existing functionality:

- **Event handlers**. An event handler is a workflow that runs in response to the events raised by a package, task, or container.
- **Configuration**. A configuration is a set of property-value pairs that defines the properties of the package and its tasks, containers, variables, connections, and event handlers when the package runs.
- **Logging and log providers**. A log is a collection of information about the package that is collected when the package runs. A log provider defines the destination type and the format that the package and its containers and tasks can use to log run-time information. Integration Services includes various built-in log providers for logging. You can also create and use custom log providers.
- **Variables**. Integration Services supports system variables and user-defined variables. The system variables provide useful information about package objects at run time. User-defined variables support custom scenarios in packages.
- **Command Prompt utilities**. Integration Services includes several Command Prompt utilities for running and managing Integration Services packages, such as:
    - **dtexec**. You can use dtexec to run an existing package. To access the utility, run dtexec and specify any additional command switches that you require.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

- **dtutil**. You can use dtutil to manage the existing packages at the command prompt. You can access packages that are stored in the SQL Server msdb database, the SSIS Package Store, and the file system. You can also perform tasks such as copying, deleting, moving, and signing packages. In addition, you can verify that a specified package exists. However, dtutil does not support the use of command files or redirection. The dtutil utility includes features, such as:
  - Remarks in the command prompt, which makes the command prompt action easier to understand.
  - Overwrite protection, to prompt for a confirmation before overwriting an existing package when copying or moving packages.
  - Console help, to provide information about the command options for dtutil.
  - Security through user names and passwords. If the utility accesses a package that is stored in msdb, the com006Dand prompt may require a user name and a password. Also, if the instance of SQL Server uses SQL Server Authentication, the command prompt requires a user name and a password. If the user name is missing, dtutil tries to log on to SQL Server by using Windows Authentication.

Configuring Windows Firewall for SSIS

Windows Firewall helps prevent unauthorized access to computer resources over a network connection. To access SSIS through this firewall, you need to configure the firewall to enable access.

The SSIS service uses the DCOM protocol.

If the firewall supports application-level filtering, you can use the user interface that Windows provides to specify the exceptions that are allowed through the firewall, such as programs and services. Otherwise, you need to configure DCOM to use a limited set of TCP ports.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

The SSIS service uses port 135, and the port cannot be changed. You need to open TCP port 135 for access to the service control manager (SCM). SCM performs tasks such as starting and stopping Integration Services and transmitting control requests to the running service.

You can configure Windows Firewall by running a command at the command prompt, or by setting properties in the Windows Firewall dialog box.

Source:

- msdn.microsoft.com/en-us/library/ms141134.aspx
- http://msdn.microsoft.com/en-us/library/ms140258.aspx
- http://technet.microsoft.com/en-us/library/ms137861.aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Control Flow Elements of an SSIS Package

Element	Description
Containers	Provide structure in packages and services to tasks in the control flow
Tasks	Perform the work in packages
Precedence Constraints	Connect containers and tasks in packages into an ordered control flow

Key Points

An SSIS package consists of a control flow and one or more data flows. The data flows are optional. SSIS provides the following three types of control flow elements:

Containers

SSIS includes the following container types for grouping tasks and implementing repeating control flows:

- The For Loop container repeats its control flow until a specified expression evaluates to FALSE.
- The Foreach Loop container enumerates a collection and repeats its control flow for each member of the collection.
- The Sequence container lets you define a subset of the control flow within a container and manages tasks and containers as a unit.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

- Containers support repeating control flows in packages. They group tasks and containers into meaningful units of work. In addition to tasks, containers can include other containers.

Packages use containers to:

- Repeat tasks for each element in a collection, such as files in a folder, schemas, or SQL Server Management Objects (SMO) objects.
- Repeat tasks until a specified expression evaluates to false.
- Group tasks and containers that must succeed or fail as a unit.

SSIS architecture supports nested containers. A control flow can include multiple levels of nested containers. For example, a package can contain a container such as a Foreach Loop container, which in turn could contain another Foreach Loop container.

Event handlers also have control flows, which are built using the elements as control flow.

Tasks

SSIS Designer, the graphical tool in SSIS for working with packages, provides the design surface for creating package control flow, and provides custom editors for configuring tasks. You can also program the SSIS object model to create packages programmatically.

Tasks perform the work in packages. SSIS includes tasks for performing various functions, such as copying files and directories, sending and receiving messages between packages, and downloading files and data. The following table describes some of the tasks and their functions.

Task Name	Functionality
Data flow	Defines and runs data flows that extract data, apply transformations, and load data
Data preparation	• Copies files and directories • Downloads files and data • Saves data returned by Web methods • Works with XML documents
Workflow	• Communicates with other processes to run packages or programs

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Task Name	Functionality
	• Sends and receives messages between packages • Sends e-mail messages • Reads Windows Management Instrumentation (WMI) data • Watches for WMI events
SQL Server	Accesses, copies, inserts, deletes, or modifies SQL Server objects and data
Analysis services	Creates, modifies, deletes, or processes Analysis Services objects
Scripting	Extends package functionality by using custom scripts
Maintenance	Performs administrative functions such as backing up and shrinking SQL Server databases, rebuilding and reorganizing indexes, and running SQL Server Agent jobs

Precedence Constraints

Precedence constraints connect containers and tasks in packages into an ordered control flow. You can control the sequence execution for tasks and containers, and specify conditions that determine whether tasks and containers can run.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Precedence constraints link executables, containers, and tasks in packages into a control flow, and specify conditions that determine whether executables run. An executable can be a For loop, Foreach loop, Sequence container, task, or an event handler. Event handlers use precedence constraints to link their executables into a control flow.

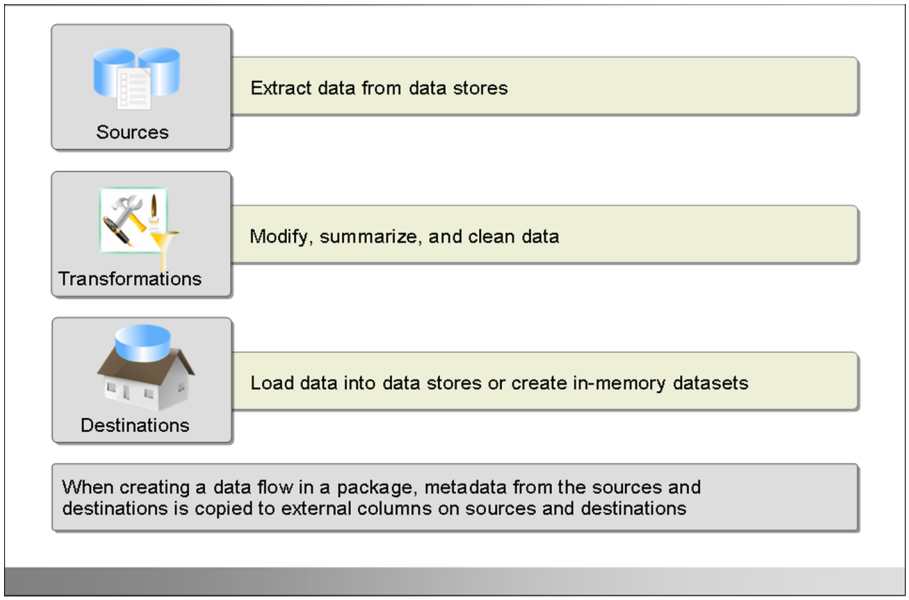
Source:

- http://msdn.microsoft.com/en-us/library/ms137681.aspx
- http://msdn.microsoft.com/en-us/library/ms139892(SQL.90).aspx
- http://technet.microsoft.com/en-us/library/ms141261(SQL.90).aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Data Flow Elements of an SSIS Package



Key Points

SSIS provides the following three types of data flow components:

- **Sources**. It extracts data from data stores such as tables and views in relational databases, files, and Analysis Services databases.
- **Transformations**. It modifies, summarizes, and cleans data.
- **Destinations**. It loads data into data stores or create in-memory datasets.

Additionally, SSIS provides paths that connect the output of one component to the input of another component. Paths define the sequence of components, and let you add annotations to the data flow or view the source of the column.

Sources

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

In SSIS, a source is the data flow component that makes data from different external data sources available to the other components in the data flow.

The source for a data flow typically has one regular output. The regular output contains output columns, which are columns the source adds to the data flow.

An error output for a source contains the same columns as the regular output. It also contains two additional columns that provide information about errors.

There are three types of sources:

- Sources that Setup installs when you install SSIS.
- Sources that you can download from the Microsoft Web site.
- Custom sources that you create if the existing sources do not meet your needs.

Transformations

Transformations can perform tasks such as updating, summarizing, cleaning, merging, and distributing data.

The inputs and outputs of a transformation define the columns of incoming and outgoing data. Depending on the operation performed on the data, some transformations have a single input and multiple outputs, while other transformations have multiple inputs and a single output. Transformations can also include error outputs that provide information about the error that occurred, together with the data that failed.

SSIS transformations are the components in the data flow of a package that aggregate, merge, distribute, and modify data. Transformations can also perform lookup operations and generate sample datasets.

Destinations

A destination is the data flow component that writes the data from a data flow to a specific data store, or creates an in-memory dataset.

An SSIS destination must have at least one input. The input contains input columns, which come from another data flow component. The input columns are mapped to columns in the destination.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Many destinations also have one error output. The error output for a destination contains output columns, which typically contain information about errors that occur when writing data to the destination data store.

There are three types of destinations:

- Destinations that Setup installs when you install SSIS.
- Destinations that you can download from the Microsoft Web site.
- Custom destinations that you create if the existing destinations do not meet your needs.

External Metadata

When you create a data flow in a package by using SSIS Designer, the metadata from the sources and destinations is copied to the external columns on sources and destinations. This metadata provides a snapshot of the schema. When SSIS validates the package, SSIS Designer compares this snapshot against the schema of the source or destination, and generates errors and warnings, depending on the changes.

The SSIS project provides an offline mode. When you work offline, no connections are made to the sources or destinations the package uses, and the metadata of external columns is not updated.

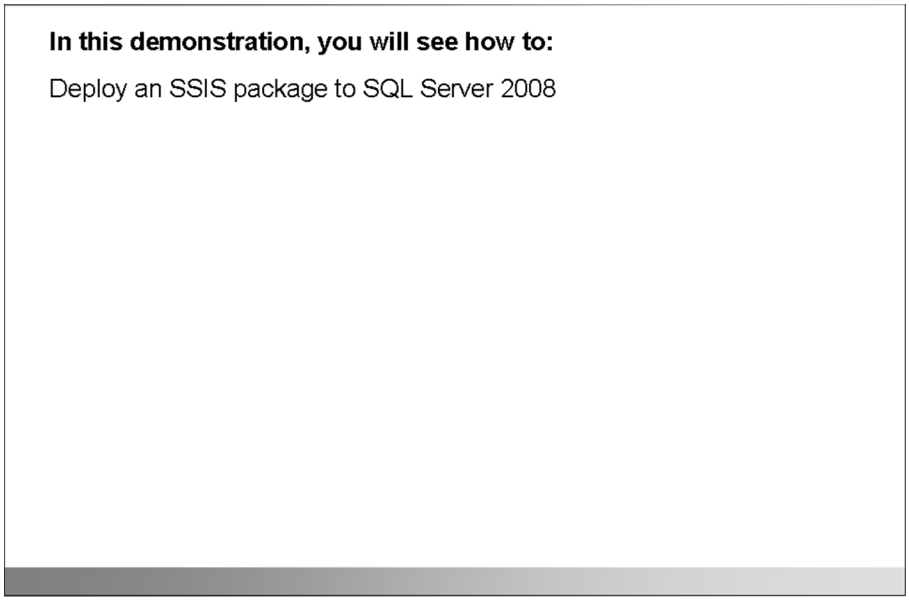
Source:

- http://msdn.microsoft.com/en-us/library/ms140080.aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Demonstration: How To Deploy an SSIS Package



Key Points

The steps to configure windows firewall for SSIS package to allow port 135 by using netsh are:

1. On the Start menu, click **Run**, type **cmd**, and then click **OK**. The command prompt window appears.
2. At the command prompt, run the following command.

   ```
   netsh firewall add portopening protocol=TCP port=135
   name="RPC (TCP/135)" mode=ENABLE scope=SUBNET
   ```

3. At the command prompt, run the following command.

   ```
   netsh firewall add allowedprogram
   program="%ProgramFiles%\Microsoft SQL
   Server\100\DTS\Binn\MsDtsSrvr.exe" name="SSIS Service"
   scope=SUBNET
   ```

4. At the command prompt window, run the **Exit** command.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

The steps to run a SSIS package and apply a XML configuration file are:

1. In Windows Explorer, browse to the **C:\Program Files\Microsoft SQL Server\100\Samples\Integration Services\Package Samples\CalculatedColumns Sample\CalculatedColumns** folder, and then double-click the **CalculatedColumns.dtproj** file.
2. In the CalculatedColumns – Microsoft Visual Studio window, under Solution Explorer pane, double-click CalculatedColumn.dtsx. the **CalculatedColumn.dtsx** page appears.
3. On the **CalculatedColumn.dtsx** page, on the **Control Flow** tab, click the **Calculated Values** box.
4. On the **SSIS** menu, click **Package Configurations**. The Package Configuration Organizer window appears.
5. In the **Package Configuration Organizer** window, click **Add**. The Package Configuration Wizard launches.
6. In the Package Configuration Wizard, on the **Welcome to the Package Configuration Wizard** page, click **Next**.
7. On the **Select Configuration Type** page, in the **Configuration type** list, ensure that XML configuration file is selected.
8. Click the **Browse** button besides the **Configuration file name** text box. The **Select Configuration File Location** dialog box appears.
9. In the **Select Configuration File Location** dialog box, browse to the **D:\Data** folder, then in the **File name** box, type **demo.dtsconfig**, and click **Save**.
10. On the **Select Configuration Type** page, click **Next**.
11. On the **Select Properties to Export** page, under **Objects**, select the **SaveCheckPoints** check box, and then click **Next**.
12. On the **Completing the Wizard** page, click **Finish**, and then click **Close**.
13. In Windows Explorer, browse to the **D:\Data** , and open the **demo.dtsconfig** file in Notepad.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

The steps to import a SSIS package from a file by using SSMS are:

1. In the Microsoft SQL Server Management Studio window, in the Object Explorer pane, click **Connect**, and then click **Integration Services**. The **Connect to Server** dialog box appears.
2. In the **Connect to Server** dialog box, in the Server name combo box, type or select the name of the SQL Server instance, and then click **Connect**.
3. In the Object Explorer pane, under the SQL Server instance that is running, expand the **Stored Packages** folder.
4. Right-click the File System in which you want to import a package, click **Import Package**. The **Import Package** dialog box appears.
5. In the **Import Package** dialog box, in the **Package Location** list, click **File System**.
6. Click the **ellipsis** button (…) next to **Package path** box. The **Load Package** dialog box appears.
7. In the **Load Package** dialog box, browse to **C:\Program Files\Microsoft SQL Server\100\Samples\Integration Services\Package Samples\CalculatedColumns Sample\CalculatedColumns\CalculatedColumns.dtsx**, and then click **Open**.
8. In the **Import Package** dialog box, click **OK** to complete the import.

The steps to export a SSIS package to a file by using SSMS are:

1. In the Object Explorer pane, under the SQL Server instance that is running, under the **Stored Packages** folder, expand the subfolders to locate the **CalculatedColumns** package.
2. Right-click the **CalculatedColumns** package, and then click **Export Package**. The **Export Package - \FileSystem\CalculatedColumns** dialog box appears.
3. In the **Export Package - \FileSystem\CalculatedColumns** dialog box, in the **Package Location** list, click **File System**.
4. Click the **ellipsis** button (…) next to **Package path** box. The **Save Package To File** dialog box appears.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

5. In the **Save Package To File** dialog box, browse to the **C** drive, in the File name box, type a name for the **exported package**, and then click **Save**.
6. In the **Export Package - \FileSystem\CalculatedColumns** dialog box, click **OK.**

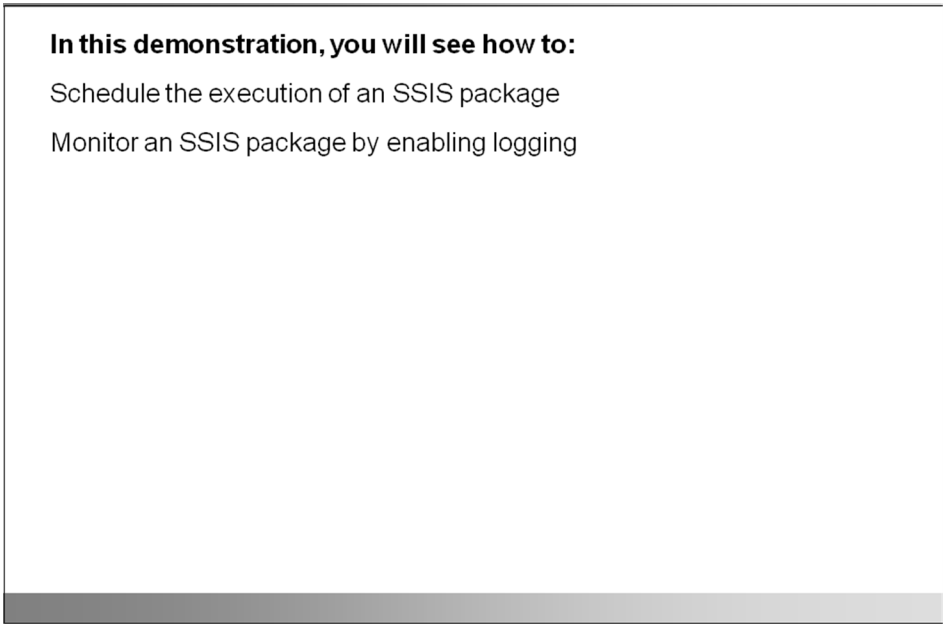
Question: What is the purpose of the XML configuration file?

Question: What is the SSIS deployment manifest file?

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Demonstration: How To Execute and Monitor an SSIS Package



Key Points

The steps to schedule the execution of an SSIS package are:

1. On the Start menu, click **Run**, type **cmd**, and then click **OK**. The command prompt window appears.
2. At the command prompt, run the following command.

```
dtexec/f " C:\Program Files\Microsoft SQL
Server\100\Samples\Integration Services\Package
Samples\CalculatedColumns
Sample\CalculatedColumns\CalculatedColumns.dtsx"
```

3. In the Microsoft SQL Server Management Studio window, open the instance of SQL Server in which you want to create a job, or the instance that contains the job to which you want to add a step.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

4. In the Microsoft SQL Server Management Studio window, in the Object Explorer pane, under the SQL Server instance that is running, expand SQL Server Agent and perform one of the following tasks:
   - To create a new job, right-click **Jobs**, and then click **New**. The New Job window appears
   - To use an existing job, expand **Jobs**, right-click the job, and then click **Properties**. The Job Properties – window appears.
5. In the window, on the **General** page, if you are creating a new job, provide a job name, **New**, select an owner and job category, **Data Collector**.
6. To make the job available for scheduling, ensure that the **Enabled** check box is selected.
7. On the **Steps** page, click New button. The New Job Step window appears.
8. In the New Job Step window, in the **Step name** box provide a step name, Step, in the **Type** list, select a job step type that is based on the version of the dtexec utility (dtexec.exe) that you want to run the job.

Note: To run the job by using the version of the dtexec utility that the system automatically invokes, select **SQL Server Integration Services Package**.

9. On the **General** tab, in the **Package source** list, in the **Package source** list, select **SSIS Package Store**.
10. If the source is SQL Server or the SSIS Package Store, in the **Server** list, select **NYC-SQL1**, and then select the authentication mode to use. If the source is the file system, provide the path of the package.
11. In the **Package** box, provide the name of the package, or click the **ellipsis** button (…). The **Select an SSIS Package** dialog box appears.
12. In the **Select an SSIS Package** dialog box, browse to the **\MSDB\Data Collector\PerfCountersCollect** package.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

13. On the **Configurations** tab, click **Add**. The **Browse for File – NYC-SQL1** dialog box appears.
14. In the **Browse for File – NYC-SQL1** dialog box, under **Demo**, select **demo.dtsconfig**, and then click **OK**.
15. On the **Command files** tab, click **Add.** The **Browse for File – NYC-SQL1** dialog box appears.
16. In the **Browse for File – NYC-SQL1** dialog box, browse to the **C:\Program Files\Microsoft SQL Server\100\Samples\Integration Services\Package Samples\CalculatedColumns Sample\CalculatedColumns** location, and then click **CalculatedColumns.dtsx**.
17. On the **Data sources** tab, select the **ConfigConnection** check box and then update the connection string.
18. On the **Execution options** tab, select the **Fail the package on validation warnings** check box.
19. In the **New Job Step** window, click **OK**.
20. On the **Start** menu, point to **All Programs**, click **Microsoft SQL Server 2008**, and then click **SQL Server Business Intelligence Development Studio**. The Start Page – Microsoft Visual Studio (Administrator) window appears.
21. In the Start Page – Microsoft Visual Studio (Administrator) window, on the **File** menu, point to **Open**, and then click **File**. The **Open File** dialog box appears.
22. In the **Open File** dialog box, in the **File name** box, type **D:\Demofiles\Import Product Info.dtsx**, and then click **Open**.
23. In the Import Product Info.dtsx [Design] – Microsoft Visual Studio (Administrator) window, on the Control Flow tab, right-click anywhere, and click **Logging**. The Configure SSIS Logs: Import Product Info window appears.
24. In the Configure SSIS Logs: Import Product Info window, in the Containers pane, ensure that the **Import Product Info** check box is selected.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

25. In the work pane, on the **Providers and Logs** tab, in the **Provider type** list, click **SSIS log provider for Text files**, and then click the **Add** button.
26. On the **Details** tab, ensure that the **Events** check box is selected.
27. On the **Providers and Logs** tab, under **Select the logs to use for the container**, in the **Configuration** list for SSIS log provider for Text files, double-click **log.txt**, and then click **New Connection**. The **File Connection Manager Editor** dialog box appears.
28. In the **File Connection Manager Editor** dialog box, in the **Usage type** list, click **Create file**.
29. Click the **Browse** button next to the File box. The **Select File** dialog box appears.
30. In the **Select File** dialog box, in the **File name** box, type **D:\Demofiles\log.txt**, and then click **Open**.
31. In the **File Connection Manager Editor** dialog box, click **OK**.
32. In the Configure SSIS Logs: Import Product Info window, click **OK**.

Question: Is there an alternate method to schedule the execution of an SSIS package, instead of using the SQL Server Agent?

Question: Is it possible to run just a task within a package?

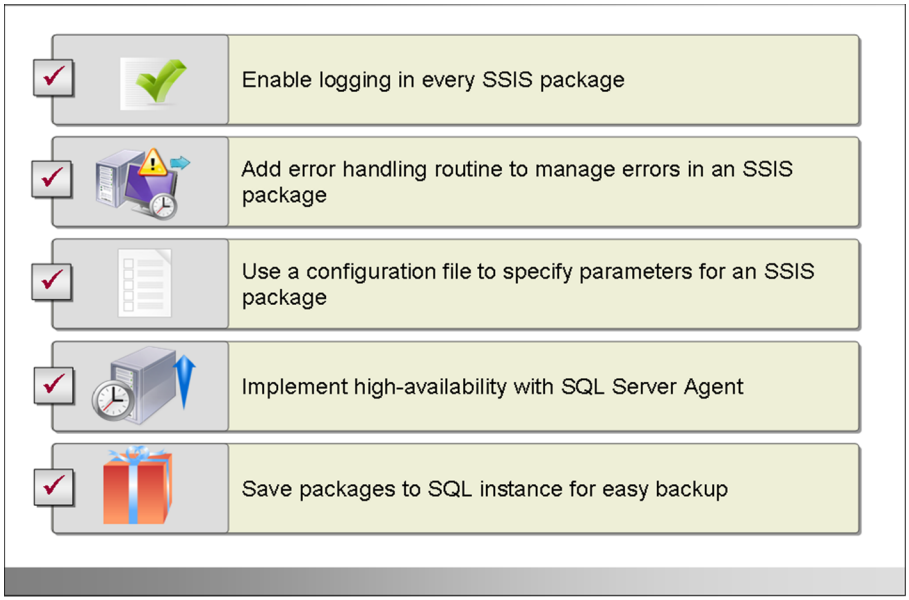
Question: How can you write the error information from an SSIS package?

Question: How can you modify the behavior of a package when you try to run it?

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Guidelines for Using SSIS Packages



Key Points

Consider the following guidelines when using SSIS packages:

- **Enable logging to track running packages**. Log providers capture information about specified events for later analysis. The Log providers save the captured information in a database table, a flat file, an XML file, or another supported output format.
- **Add error-handling routines to manage errors in the SSIS package**. Create an event handler for the OnError event. In the event handler, you can:
  - Use a Send Mail task to notify an administrator of the failure.
  - Use a Script task and custom logic to obtain system information for troubleshooting.
  - Clean up temporary resources or incomplete output.
- **Use a configuration file to specify parameters for the SSIS**

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

package. This simplifies the process of switching the configuration between environments. Using a configuration file also helps avoid mistakes when you modify properties in a package.

- **Implement high-availability with SQL Server Agent**. When you need to run multiple packages, use other available servers. This method of using other servers to run packages when the packages are all under the control of one parent package is called load balancing. In SSIS, load balancing is a manual procedure that the owners of the packages perform. Servers do not automatically perform load balancing. In addition, the packages that are run on the remote servers must be whole packages, not individual tasks in other packages.
- **Save packages to SQL instances for easy backup**. You can save SSIS packages to the file system or msdb, which is a SQL Server system database. You can back up and restore packages saved to msdb by using SQL Server backup and restore features. Any configuration files that the packages include are stored in the file system. These configuration files are not backed up when you back up the msdb database. Therefore, ensure that the configuration files are backed up regularly as part of your plan for securing packages saved to msdb. To include configurations in the backup of the msdb database, consider using the SQL Server configuration type, instead of file-based configurations.

Question: What should you do to enable logging for an SSIS package?

Question: Where should you store the SSIS package if you want simplified backup and consolidated management?

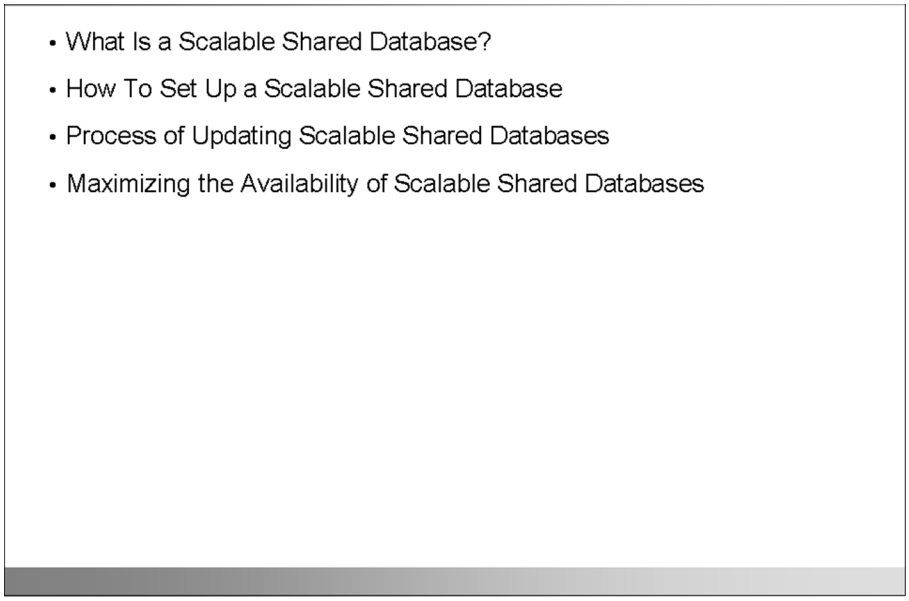
Source:

- http://msdn.microsoft.com/en-us/library/ms345184.aspx
- http://msdn.microsoft.com/en-us/library/ms141699.aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Lesson 5
Setting Up Scalable Shared Databases



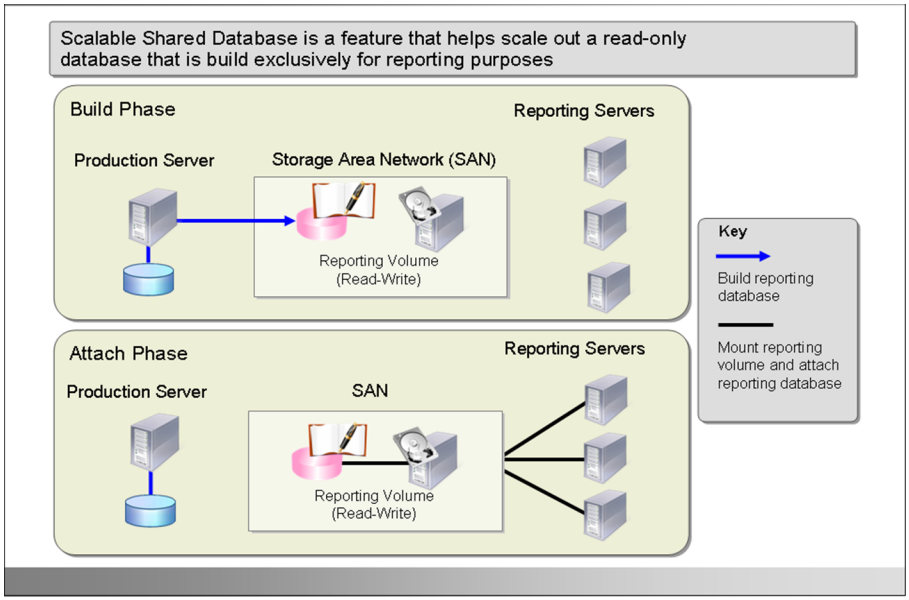
The scalable shared database feature allows database developers and administrators to attach a read-only reporting database to multiple server instances over a storage area network (SAN). This feature is used by database developers and administrators for reporting purposes to create a scalable platform for reporting solutions by gathering identical data as read-only on multiple servers. It helps in reducing storage complexity, simplifying SQL Server scale-out for complex, off-hours reporting workloads.

In this lesson, you will learn about the architecture of a scalable shared database. You will also learn how to work with the scalable shared database.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

What Is a Scalable Shared Database?



Key Points

You can use the Scalable Shared Database feature to scale out a read-only database built exclusively for reporting purposes. The reporting database must reside on a set of dedicated, read-only volumes whose primary purpose is hosting the database. By using commodity hardware for servers and volumes, you can scale out a reporting database that provides an identical view of the reporting data on multiple reporting servers. The Scalable Shared Database feature allows a smooth update path for the reporting database.

After you build the reporting database on a set of reporting volumes, you should mark the volumes as read-only and mount them to multiple reporting servers. You can then attach the reporting database to an instance of Microsoft SQL Server 2005 or SQL Server 2008 on each reporting server. This reporting database then becomes available as a scalable shared database. After it is established as a scalable shared database, a reporting database can be shared by clients by using different reporting servers. To query the database, a user or application can connect to any server

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

instance to which the database is attached. For a given version of a reporting database, clients on different servers obtain an identical view of the reporting data, making query results consistent across servers.

Building and Scaling Out a Reporting Database

Building and scaling out a reporting database involves the following two phases:

- **The Build phase**. This phase is also called the refresh phase. When you configure a new scalable shared database, you should first build a new reporting database on a set of reporting volumes or refresh an outdated version of the reporting database on the reporting volumes.
  To build a reporting database, you first mount the reporting volume on the production server. After you mount the volume onto the production system, the volume is marked as read-write. Then, build a reporting database on the volume to copy data or database. After building the database, you should set each reporting volume to read-only and dismount it.
- **The Attach phase**. After you build the reporting database, you need to scale out the database by configuring it as a scalable shared database on multiple server instances.
  In this phase, first, you mount the read-only reporting volume onto multiple reporting servers over the SAN. Then, on each reporting server, you attach the reporting database to an instance of SQL Server. The database is attached as a read-only database because the volumes are read-only. After the completion of this process on a given reporting server, the reporting database becomes a scalable shared database on that server. However, the attach phase as a whole continues until the database is attached on all the reporting servers.

A copy of a reporting database remains available as a scalable shared database for as long at is remains attached on any of the reporting servers.

Benefits of Scalable Shared Databases

The following are the benefits of scalable shared databases:

- **It provides workload scale-out of reporting databases**. Scalable shared databases help you scale out the workload on the reporting databases by using commodity servers and hardware. A

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

scalable shared database is a cost-effective way of making a read-only data mart or data warehouses accessible to multiple server instances for reporting purposes such as running queries or using Reporting Services.

- **It provides workload isolation**. Each server in a warehouse or a data mart uses its own memory, CPU, and tempdb database. This distribution of workload prevents one poorly-tuned query from consuming all your server resources and affecting other queries.
- **It guarantees the same view of reporting data from all servers**. The scalable shared database provides the same view of reporting data from the servers if all the server instances are configured identically. For example, all server instances need to use the same collation.

Note: You can also update the reporting database on a second reporting volume.

Limitations of Scalable Shared Database

Scalable shared databases have the following limitations:

- The database must be on a read-only volume.
- The data files are accessible over a storage area network (SAN).
- The databases are supported by Windows Storage running only on Windows Server® 2003 SP1 or Windows Server 2008.
- The scalable shared database configurations must be limited to eight server instances per shared database.
- The scalable shared databases do not support database snapshots.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Important: You must ensure that the SAN environment is working properly to configure a scalable shared database.

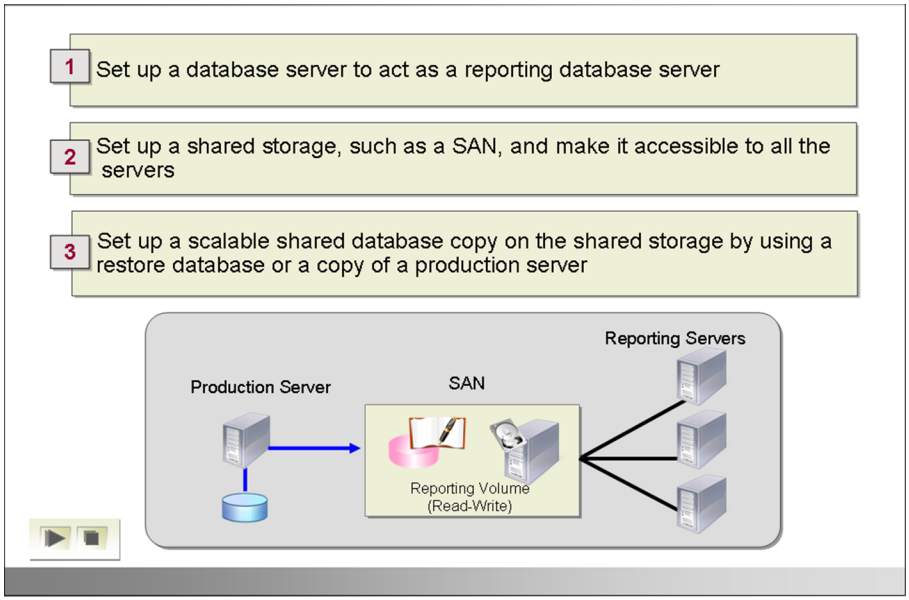
Question: When should you consider using scalable shared database?

Source:

- http://technet.microsoft.com/en-us/library/ms345392.aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

How To Set Up a Scalable Shared Database



Key Points

To use a scalable shared database, you need to have shared disk storage, such as a SAN and two volumes for storing the read-only and read-write versions of the database.

While setting up a scalable shared database, perform the following steps:

1. **Set up a database server to act as a reporting database server**. For a scalable shared database, you require that a reporting server and the associated reporting database run on identical platforms. Whenever you update a reporting server, you need to update any other reporting servers that serve the same scalable shared database or databases. The Scalable Shared Database feature is tested under a concurrent access workload by up to eight server instances. SQL Server does not enforce an instance limit, but you must limit your scalable shared database configurations to eight server instances per shared database.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

2. **Set up a shared storage such as SAN and make it accessible to all the servers**. After you mount the reporting volume on a reporting server, ensure that it is correctly mounted and that the data files are accessible by running the DIR <drive-letter>:\<database-directory> command at the command prompt. In this command, <drive-letter> is the letter assigned to the reporting volume and <database-directory> specifies the location of the database's data files on the volume. You should perform this test for each of the reporting servers to ensure that you get the same results for all of them. You should also ensure that the reporting volume is read-only. To ensure that the reporting volume is read-only, you should attempt to create a file on the volume. The easiest way to do this is to attempt to copy or save a plain text file onto the volume. The attempt should fail.
3. **Set up a scalable shared database copy on the shared storage by either using a restore database, or a copy of a production server**.
    - **Restore database**. Use the same path for the reporting database and the production database to simplify building or updating a reporting database. However, if you build the reporting database by restoring a database backup, consider specifying the full path of the restored data files in your RESTORE DATABASE statement. For example, if the production database is on E:\SQLdata, if possible, use E as the drive letter of the reporting volume and \SQLdata as the directory of the reporting database. However, if they are different, you could write a script using explicit paths to handle the differences.
    - **Copy of a production server**. If you build the reporting database by copying the production database, consider listing each file and specifying the full path when you attach the reporting database in the FOR ATTACH clause of your CREATE DATABASE statement.

You can optimize the reporting database for the anticipated query workload. Such optimizations could include:

- Updating or recreating statistics on affected tables and indexes as needed.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

- Rebuilding indexes with a FILL factor of 100.
- Putting the reporting database into simple recovery mode and performing a checkpoint to truncate the transaction log and save log space.

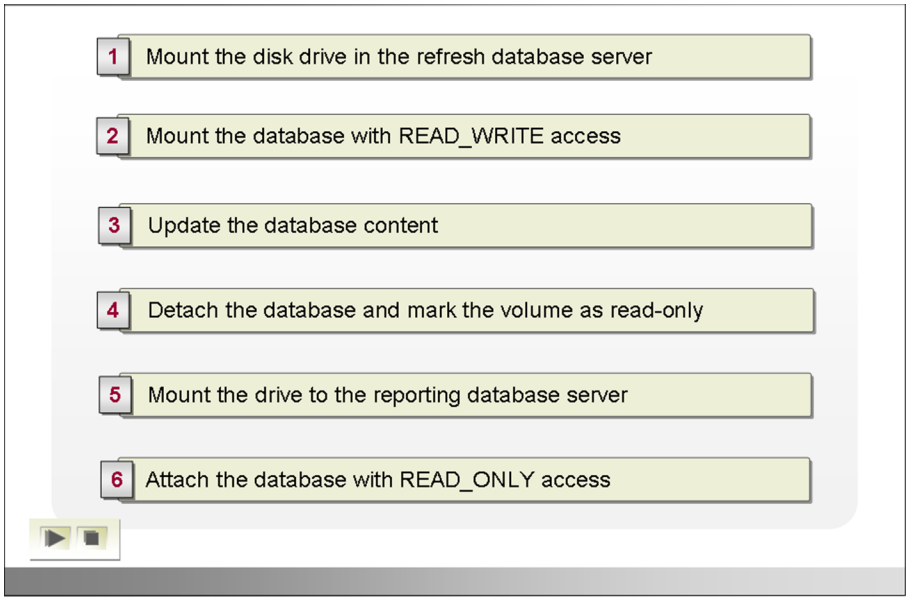
Source:

- http://technet.microsoft.com/en-us/library/ms345523.aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Process of Updating Scalable Shared Databases



Key Points

You can create a scalable shared database by building a new reporting database on a set of reporting volumes. When the reporting data becomes obsolete, the reporting database is said to be outdated. An outdated reporting database is refreshed as part of each update cycle. You can refresh a stale reporting database by updating its data or building a completely new version of the database.

To build or refresh a reporting database, you can use any of the data-copy methods provided by SQL Server for copying the data or databases. These methods include SSIS, backup and restore, and copying the production database.

SSIS

You can create or copy a database by running SSIS packages and by using the Execute SQL task or the Transfer Database task. The Execute SQL task runs SQL statements or stored procedures from a package. By using the Execute SQL task, you can create a database

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

by running a CREATE DATABASE statement and then populating the database by copying in one or more tables or views.

The Transfer Database task can copy a database within the same server instance or between instances.

Backup and Restore
You can restore a backup of a production database onto the reporting volume. This involves restoring and recovering a full database backup on to the reporting volume.

If you are using the same drive letter, mount the reporting volume on to a different host and connect to a server instance there to restore the database.

If the reporting volume uses a different drive letter from the production volume, you must add a WITH MOVE clause to your RESTORE DATABASE statement that specifies the drive letter of the reporting volume in the pathname of the restored database.

Copying the Production Database
Before you can copy a database manually or use the Detach and Attach Method of the Copy Database Wizard, you must take the database offline, and after copying it, bring it back online. The Copy Database Wizard, however, offers an alternative method, the SMO Transfer method, that copies the database while it remains online. Though slower than the Detach and Attach Method, the SMO Transfer method has the advantage of preserving active connections to the database.
While updating a scalable shared database, perform the following steps:

1. **Mount the disk drive in the refresh database server**. You need to unmask the logical unit number (LUN) corresponding to each reporting volume to update a scalable shared database. This unmasking will make the volume accessible to the production server.
2. **Mount the database with READ_WRITE access**. You can use the DiskPart list volumes command to list information about all disks and volumes. The following commands describe how to mount the volume by using the DiskPart tool.

   ```
   DiskPart

   DISKPART> select volume= 1
   ```

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.


```
DISKPART> assign letter= e:

DISKPART> attribute clear readonly

DISKPART> exit
```

3. **Update the database content**. To refresh an existing reporting database, first, you need to attach the database to a server instance. Then, you need to set the read-write access to the database by using the ALTER DATABASE <database_name> SET READ_WRITE statement. Finally, you need to build or refresh the database.
4. **Detach the database and mark the volume as read-only**. You can use the sp_detach_db @dbname='<database_name>' statement to detach the database from the production server instance. In addition, you need to mark each reporting database as read-only.
5. **Mount the drive to the reporting database server**. You need to make the volume accessible to the production server to update a scalable shared database. You can make the volume accessible by unmasking the LUN corresponding to each reporting volume. However, if the database uses multiple reporting volumes, you need to mount the volume corresponding to a LUN on each of the reporting servers.
6. **Attach the database with READ_ONLY access**. After mounting a set of reporting volumes to each of the reporting servers, you should validate that the volume is accessible as a read-only volume over the SAN. Finally, you need to attach the database to one or more server instances on each of the reporting servers.

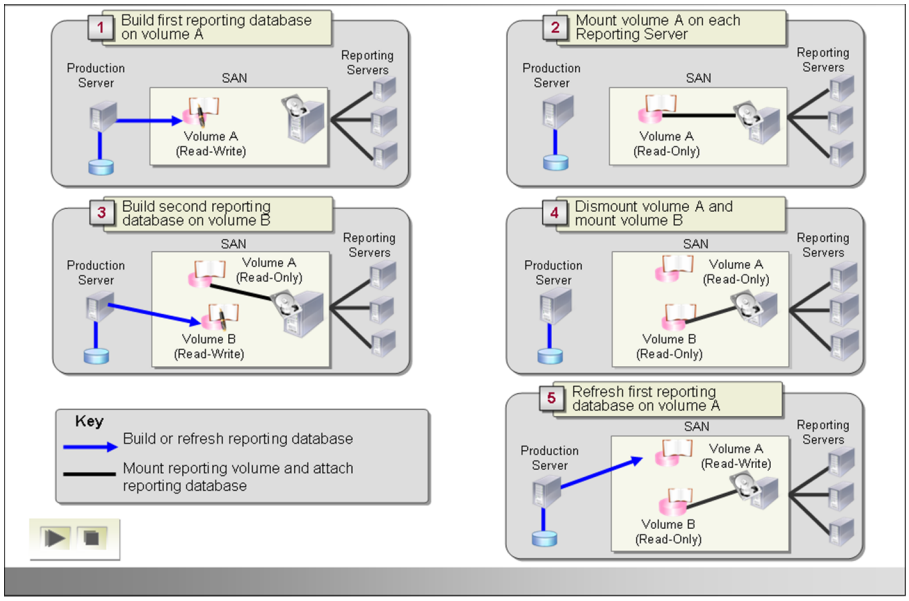
Source:

- http://technet.microsoft.com/en-us/library/ms345390(SQL.90).aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Maximizing the Availability of Scalable Shared Databases



Key Points

You must use two alternate sets of reporting volumes to maximize availability of reporting data by overlapping two update cycles. You can mount one set of reporting volumes to the reporting servers and build an up-to-date version of the reporting database on the alternate set of volumes. After completing the build phase or refresh phase for the fresh reporting database, you can overlap the attach phase and the detach phase of the two sets of volumes.

To overlap the detach phase of the outdated reporting database with the attach phase of the fresh reporting database, on each reporting server, start the attach phase for the fresh database, by mounting the alternative set of volumes onto the reporting server. Then, perform the whole detach phase for the stale database by detaching it, which makes the scalable shared database unavailable on the reporting server, and dismounting the corresponding reporting volume or volumes. Finally, after you detach the scalable shared database, complete the attach phase for the fresh version by

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

attaching the database to the server instance, which makes it available as a scalable shared database.

You can complete the overlapped attach and detach phases on one server at a time or you can perform each step on every server before you continue to the next step. You should select the best approach based on your business requirements. If you want your reporting database to return identical results to clients, you need to synchronize the updates by detaching the old copy from all the server instances before attaching the new copy to any of them. In this case, start the attach phase for the fresh database on all the servers, then complete the detach phase on all the servers, and, finally, attach the database on each of the reporting server as quickly as possible. If you want your reporting database to return different results on different server instances, you can perform a rolling update of the reporting database. You can perform a rolling update by completing the overlapping attach and detach phases on one reporting server at a time.

While synchronizing the update cycle on all server instances, perform the following steps:

1. **Build the first reporting database on volume A**. You can start the synchronized update by mounting the reporting volume A on the production server and marking it as read-write. This builds the initial version of the reporting database. Then, dismount the database only after marking the reporting volume as read-only.
2. **Mount volume A on each reporting server**. Mount volume A as a read-only volume on each of the reporting servers and attach to each of the server instances.
3. **Build the second reporting database on volume B**. While the reporting database is available as a scalable shared database on volume A, mount the other reporting volume, B, on the production server and mark it as read-write to build a fresh, up-to-date version of the database for the second update cycle. Then, you need to dismount the volume B.
4. **Dismount volume A and mount volume B**. After performing step 3, the detach phase for volume A and the attach phase for volume B overlap. To separate the two phases, first, mount the alternate volume B on the reporting server. Then, detach the database on volume A. Next, attach the refreshed reporting

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

database, on volume B, to each of the reporting servers. Finally, dismount volume A from each of the reporting servers.

5. **Refresh the first reporting database on volume A**. While the reporting database is available as a scalable shared database on volume B, mount the reporting volume A on the production server and mark it as read-write. This process might involve updating the existing database or building a completely new database.

Source:

- http://technet.microsoft.com/en-us/library/ms345596(SQL.90).aspx
- http://technet.microsoft.com/en-us/library/ms345596.aspx

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Lab 9: Designing a Strategy for Content Distribution in SQL Server 2008

- Exercise 1: Selecting a Content Distribution Technology (Discussion)
- Exercise 2: Performing a Remote Query
- Exercise 3: Monitoring an SSIS Package

Logon Information

Virtual machine	50400A-NYC-SQL1
User name	Administrator
Password	Pa$$w0rd

Estimated time: 60 minutes

Objective:

After completing this lab, you will be able to:

- Select a content distribution technology.
- Perform a remote query.
- Monitor an SSIS package.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Scenario

QuantamCorp has deployed the first two SQL Server 2008 servers in its organization and will continue to deploy more to replace the existing installations of SQL Server 2005. QuantamCorp wants to make SQL Server 2008 the primary database engine. As an enhancement to the application package, the application team has deployed a new SQL Server Agent job, which runs daily. This SQL Server Agent job first extracts data from another SQL Server 2008 server and merges it with the data in your SQL Server 2008 server. Next, an SSIS package performs the transformation of merged data into a form required by the reporting team. A task that you deployed is not functioning properly. After running a series of diagnostic tests, you realize that you need a linked server. Then, you can continue to diagnose the problem by enabling logging on the SSIS package to help the application team check the error in the package. You will then perform the following during the project:

- Determine the appropriate content distribution approach.
- Determine an appropriate approach for querying data in another SQL Server.
- Determine when to enable logging in an SSIS package.
- Configure a linked server and perform distributed query.
- Enable logging in an SSIS package and test it.

Exercise 1: Selecting a Content Distribution Technology (Discussion)

Scenario

You have recently received a lot of complaints about the performance of one of your business application that is currently using a database located at the corporate headquarter. After conducting a detailed investigation along with your application team, you find out that a performance issue with the application is caused whenever a user attempts to generate a large report. You need to minimize this performance impact on applications and allow users to generate large reports at their convenience.

In addition to solving the performance issue, the application team wants you to build a report that combines data from multiple

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

databases stored on multiple database servers. The application team wants to use the filtering ability in SQL Server to reduce the amount of data that can be downloaded to the client application to generate reports.

You are planning to implement logging on all SQL Server Integration Services (SSIS) packages. When an SSIS package fails, the cause of the failure takes long to identify because there are no logs stored, related to the execution of the package.

Based on the above scenario, answer the following questions:

1. What would you suggest to help resolve the performance issue related to the report generation without modifying the report and the application code?
2. How would you resolve the need of the application team to access data stored on multiple databases?
3. What would you propose to help improve the ability to troubleshoot an SSIS package execution?

Exercise 2: Performing a Remote Query

Scenario

The management of QuantamCorp has asked you to create a set of reports based on the information stored in the servers in the regional offices. Instead of replicating the data over, you choose to connect to those remote servers to retrieve information directly.

In this exercise, you would configure a linked server to query the remote databases by using your local SQL Server.

The main tasks for this exercise are as follows:

1. Create a test database on the NYC-SQL1 instance.
2. Create a test database on the NYC-SQL1\DEVELOPMENT instance.
3. Set up a linked server.
4. Verify the linked server setup.

Task 1: Create a test database on the NYC-SQL1 instance.

1. Open the Microsoft SQL Server Management Studio window.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

2. Using the Microsoft SQL Server Management Studio window, open the **Connect to Server** dialog box.
3. In the **Connect to Server** dialog box, use the following credentials to log on to Microsoft SQL Server Management Studio:
   - Server type: **Database Engine**
   - Server name: **NYC-SQL1**
   - Authentication: **Windows Authentication**
4. Using the Microsoft SQL Server Management Studio window, in the Query Editor pane, run a Transact SQL (T-SQL) code to create a test database, **Mod09_1**, by using the following credentials.

File name	Settings
Mod09_2_Primary.mdf	Current size: 10 Maximum size: 50 Filegrowth: 5
Mod09_2_Log.ldf	Current size: 10 Maximum size: 50 Filegrowth: 5

Also add a table named **CallDetails** to the **Mod09_1** database with the following structure, and also insert the following record.

Field name	Data type	Value	Record
call_date	datetime	NOT NULL	2004-01-01
caller	varchar(20)	NOT NULL	Peter
duration	Int	NULL	30
product	Int	NOT NULL	A random number between one and three.
Cs_rep	Int	NOT NULL	A random number between one and five.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Field name	Data type	Value	Record
additional_info	nchar(4000)	NULL	Call to 'x'. Where 'x' is a random number between one and 100,000.

Note: Browse to the **D:\Labfiles\Mod09** folder and select the **E2_T1.sql** file to execute the required T-SQL code.

Task 2: Create a test database on the NYC-SQL1\DEVELOPMENT instance.

1. In the **Connect to Server** dialog box, use the following credentials to log on to Microsoft SQL Server Management Studio:
   - Server type: **Database Engine**
   - Server name: **NYC-SQL1\DEVELOPMENT**
   - Authentication: **Windows Authentication**
2. Using the Microsoft SQL Server Management Studio window, in the Query Editor pane, run a T-SQL code to create a test database, **Mod09_2**, by using the following credentials.

File name	Settings
Mod09_2_Primary.mdf	Current size: 10 Maximum size: 50 Filegrowth: 5
Mod09_2_Log.ldf	Current size: 10 Maximum size: 50 Filegrowth: 5

Also add a table named, Employee**,** to the **Mod09_2** database. Insert three records in the table, as follows.

Employee ID	Name
1	Peter

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Employee ID	Name
2	Michael
3	John

Note: Browse to the **D:\Labfiles\Mod09** folder and select the **E2_T2.sql** file to execute the required T-SQL code.

Task 3: Set up a linked server.

1. Using the Microsoft SQL Server Management Studio window, open the **Connect to Server** dialog box.
2. In the **Connect to Server** dialog box, use the following credentials to log on to Microsoft SQL Server Management Studio:
   - Server type: **Database Engine**
   - Server name: **NYC-SQL1**
   - Authentication: **Windows Authentication**
3. In the Query Editor pane for NYC-SQL1 of the Microsoft SQL Server Management Studio window, run a T-SQL code to set up a linked server for the **NYC-SQL1\DEVELOPMENT** instance.

Note: Browse to the **D:\Labfiles\Mod09** folder and select the **E2_T3.sql** file to execute the required T-SQL code.

Task 4: Verify the linked server setup.

1. In the Query Editor pane for NYC-SQL1 of the Microsoft SQL Server Management Studio window, run a T-SQL code to verify the setup of the linked server by creating an INNER join with the **call_date**, **caller**, and the **e.name** columns of the **CallDetails** table, and combining them with the **name** column of the **Employee** table in the local instance.

Note: Browse to the **D:\Labfiles\Mod09** folder and select the **E2_T4.sql** file to execute the required T-SQL code.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Results: After completing this exercise, you should have created a test database on the NYC-SQL1 instance and the NYC-SQL1\DEVELOPMENT instance, set up a linked server, and verified the linked server setup.

Exercise 3: Monitoring an SSIS Package

Scenario

You prepare the report by using the databases in the remote offices. You find that the report generation is very slow when you retrieve data from remote offices. Therefore, you decide to switch the model to replicate data from the regional office to your SQL Server.

In this exercise, you will configure an SSIS package to copy data from the remote SQL Server.

The main tasks for this exercise are as follows:

1. Enable logging in an SSIS package.
2. Test the SSIS package.

Task 1: Enable logging in an SSIS Package.

1. Open the Start Page – Microsoft Visual Studio (Administrator) window.
2. Using the Start Page – Microsoft Visual Studio (Administrator) window of Business Intelligence Development Studio, enable logging for the **Import Product Info.dtsx** package, with the following settings:
   - Add the **SSIS log provider for Text files** provider type.
   - Ensure that the **Events** check box is selected to log all events for Import Product Info.
   - Using the **File Connection Manager Editor** dialog box, configure the file connection manager to save the log in **log.txt** at run time.
3. Save the **Import Product Info.dtsx** package configuration.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Task 2: Test the SSIS package.

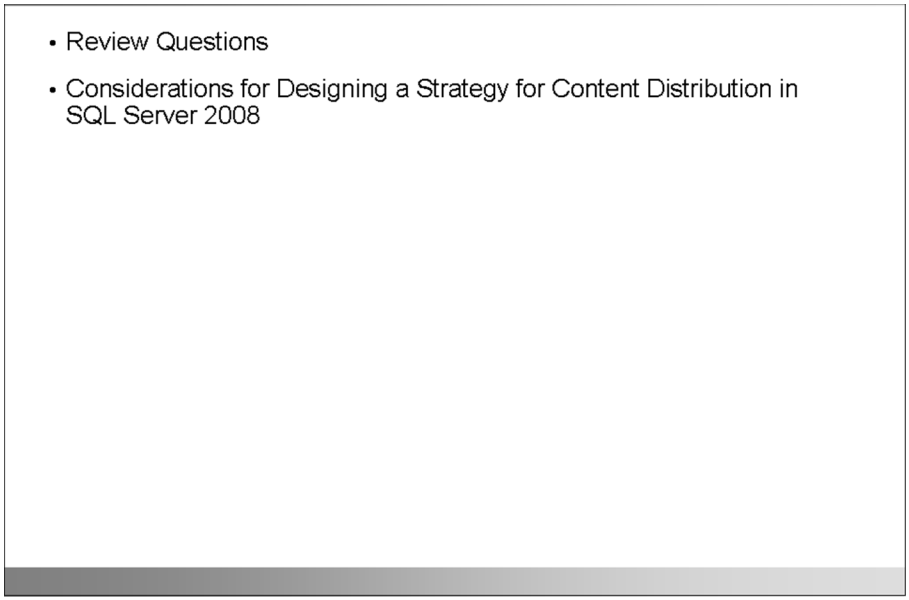
1. In Windows Explorer, open the Execute Package Utility window for the **Import Product Info.dtsx** package.
2. Using the Execute Package Utility window, test the SSIS package by creating a **log.txt** file with **SSIS log provider for Text files**. Verify the content of the **log.txt** file.

Results: After completing this exercise, you should have enabled logging in an SSIS package and tested the SSIS package.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

Module Review and Takeaways



Review Questions

1. Your application team reports that an SSIS package has failed. You want to help the application team to troubleshoot the issue. What should you do?
2. Your company has recently merged with another company. Users need to create a report that contains data from a database managed by the merged company. Security policies prevent you from granting database access to the user directly. How can you provide access to the data in the database managed by the merged company?
3. Your application team is suggesting that you split the database into two for better performance. How can you ensure changes are committed on both the databases at the same time?
4. You are managing databases for an application. Recently you notice that the database backup jobs are not completed successfully. The database backup uses marked transaction backup through linked server. What is the cause of the problem?

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use to
mscwinfo@microsoft.com.

New Horizons Computer Learning Centers, Inc.

5. You are helping the application team to regularly load information into a SQL Server database. You need to perform some cleaning and view data before data gets loaded into the database. Which tool should you use?

Considerations for Designing a Strategy for Content Distribution in SQL Server 2008

- When designing a strategy for content distribution in SQL Server 2008, consider the following:
- Bring data closer to the user, if possible for better performance.
- Use distributed query if data cannot be replicated.
- Use linked server instead of ad-hoc distributed queries for better security control and performance.
- Avoid using distributed transactions, if possible to avoid performance overhead.
- Use separate configuration, such as XML file, database table, to specify the configuration for the SSIS package.
- Enable logging in SSIS package, if possible, to generate logs.
- Install SSIS package on the msdb database for better security and centralized backup.

11/25/2012 10:58:12 PM 48bb073b-a3c2-4961-8af5-be421e6e7158 Travis Runyard
Warning: This is trunyard@sentric.net's unique copy. It is illegal to reprint, redistribute or resell this conte
The Licensed Content is licensed "as-is." Microsoft does not support this Licensed Content in any way
Microsoft gives no express warranties, guarantees or conditions.Please report any unauthorized use t
mscwinfo@microsoft.com.