



Analysis Services Overview

White Paper

Published: December 2007

Summary: Microsoft SQL Server 2008 Analysis Services builds on the value delivered with the significant investments in Analysis Services 2005 around scalability, advanced analytics and Microsoft Office interoperability. Through substantially improved performance and scalability, and developer productivity you can build enterprise scale Online Analytical Processing (OLAP) solutions. The Unified Dimensional Model consolidates data access and provides a wide range of analytical capabilities while deep integration with Microsoft Office and an open, embeddable architecture, allows you to reach every user with familiar tools and drives actionable insight to users across the enterprise.

Contents

Introduction	1
Build Enterprise-Scale Solutions	1
High Developer Productivity	1
Scalable Infrastructure.....	3
Superior Performance	4
Extend Solutions with Comprehensive Analytics	5
Unified Dimensional Model.....	5
Central Manageability of Key Enterprise Metrics	5
Predictive Analysis	5
Drive Actionable Insight through Familiar Tools.....	6
Optimized Office Interoperability.....	6
Rich Partner Extensibility.....	8
Open Embeddable Architecture.....	8
Conclusion	9

Introduction

Analytical solutions are quickly becoming mission critical for many organizations. This has led to an explosion of data stored in these systems and a need to support larger, faster solutions that can be created and developed quickly and effectively.

Build Enterprise-Scale Solutions

Microsoft SQL Server 2008 Analysis Services is designed to provide exceptional performance and scales to support applications with millions of records and thousands of users. Innovative, consolidated tools help improve developer productivity and result in better design and faster implementation.

High Developer Productivity

Developers typically have to learn and use multiple tools to build and deploy a solution. With Analysis Services however, developers can use the SQL Server Business Intelligence Development Studio (BIDS) throughout the entire development cycle from the start of the project through development to deployment. Because Business Intelligence Development Studio is based on the Visual Studio development environment, it is fully integrated with the Visual Studio Team System; which provides design, development, collaboration, optimization, and testing resources. This provides an environment where developers can work faster and more effectively within an integrated, intuitive environment. Furthermore, to even further enhance productivity BIDS also offers sophisticated Business Intelligence Wizards. A set of easy to use wizards will help even the most novice user in modeling some of the more complex business intelligence problems making the developments of BI projects more accessible to a larger number of people and organizations.

Inefficiencies in the design occurring in the early development phase often waste large amounts of development time because work that developers have already completed based on the incorrect design needs to be re-done when design mistakes have been rectified. SQL Server 2008 Analysis Services introduces a set of new, innovative Best Practice Design Alerts that provide automatic notification of potential design issues early in the development process, which reduces wasted time caused by design mistakes and facilitates a faster development process. Figure 1 shows an alert on the **Time** dimension and **Calendar** hierarchy. As you can see from Figure 1, the alerts highlight problem areas. However, they do not in any way affect functionality as the alerts can simply be ignored or dismissed individually, or globally.

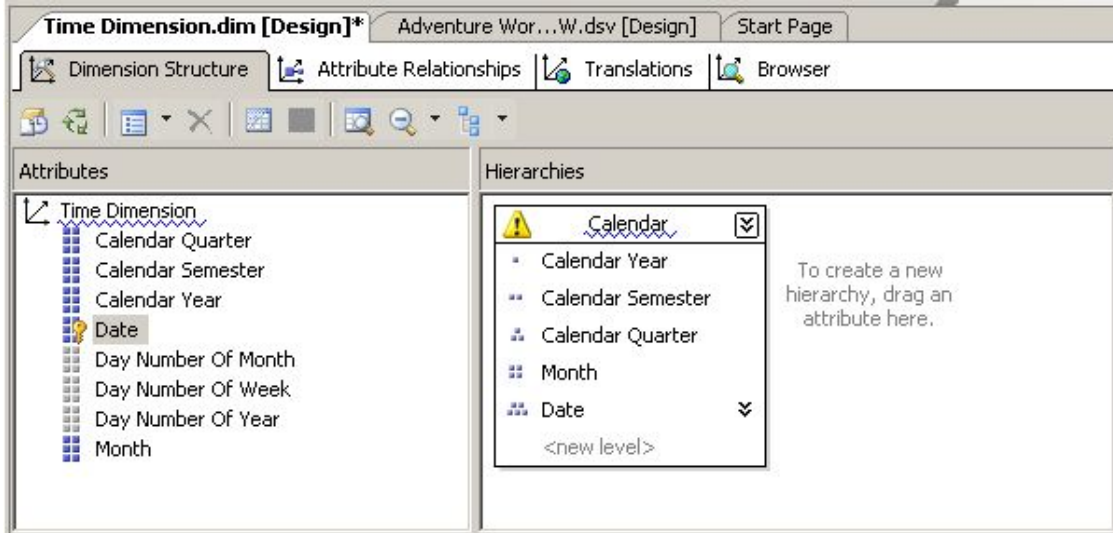


Figure 1

In addition to real time alerts, you can scan your solution design for all alerts. Figure 2 shows the current alerts on a design.

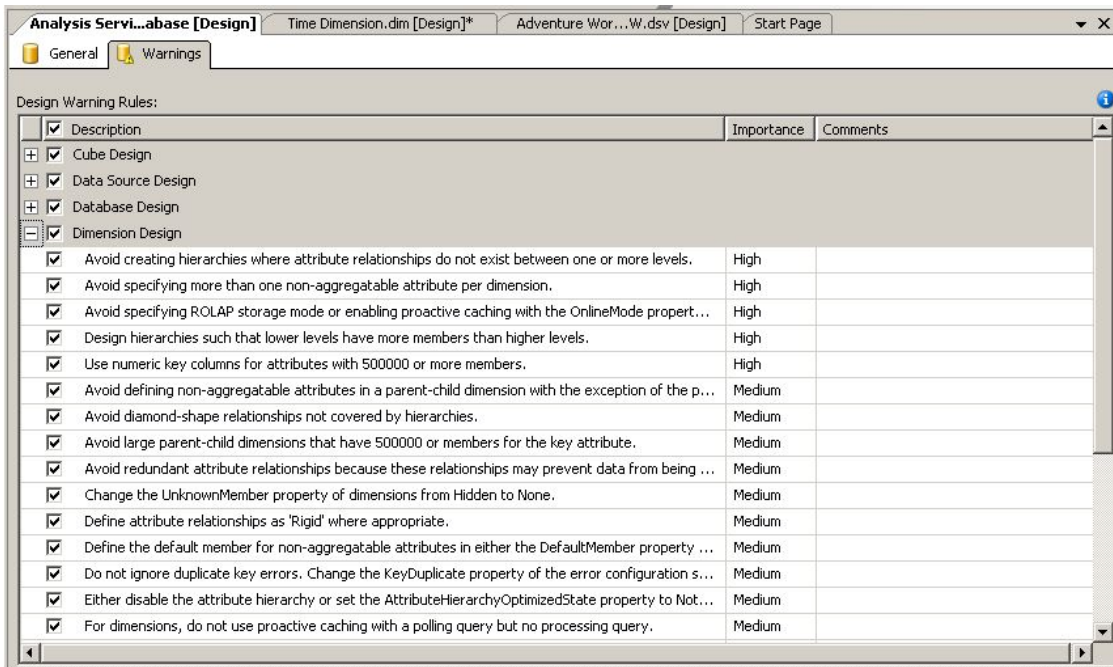


Figure 2

SQL Server 2008 Analysis Services further increases developer productivity with new, enhanced cube, dimension, and attribute designers. Figure 3 shows the new Attribute Relationships designer.

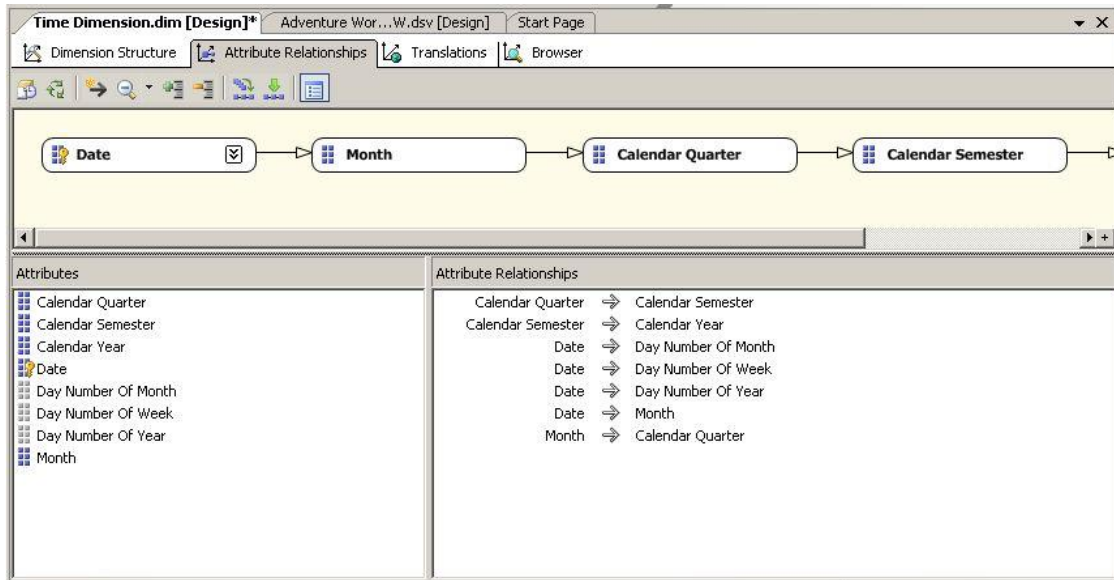


Figure 3

Scalable Infrastructure

Analysis Services can scale to support databases of many terabytes in size with many thousands of users. To support many users, avoid contention, and reduce costs you can scale out an Analysis Services solution. Scaling out an Analysis Services solution typically adds processing and storage overhead to store and synchronize several versions of the data, but SQL Server 2008 Analysis Services can share one read-only Analysis Services database between several Analysis Services servers completely removing this overhead.

Real time resource monitoring becomes essential as systems scale in both size and number of users. SQL Server 2008 Analysis Services provides Dynamic Management Views similar to those available to the database engine. These provide real time enterprise system information for monitoring, analysis, and performance tuning.

As databases increase in size, the time and cost of maintaining backups increases correspondingly. Very often backup time increases exponentially when working with OLAP databases once the databases reach certain sizes, but with SQL Server 2008 Analysis Services a new backup storage subsystem results in backup times that increase linearly with database size. This removes limitations on backup size and therefore removes limitations on database size.

As databases become larger, the information that a user requires can be harder to find. Perspectives provide a filtered view of the UDM giving all of the advantages of data marts while eliminating redundant storage, reducing processing costs, eliminating the synchronization requirement between data marts and removing data consistency and integrity issues caused by storing multiple copies of the same data.

With increasing globalization, solutions need to be presented to a worldwide audience. Data is typically the same for the whole world, but metadata such as cube, measures, dimension names and levels, and Key Performance Indicators (KPI's) will differ for each language required. Translations provide the ability to create different metadata values for each language and globally scale your solution. Financial information will also need to be localized to present results in the correct currency. Offering powerful translation capabilities and automatic currency conversions Analysis Services provides localized analytical data to users in their own language.

Superior Performance

Analysis Services cubes are multidimensional structures that enable fast access to high volumes of pre-aggregated data, empowering end users to gain insight into relevant business data at the speed of thought. Analysis Services stores its data in a highly optimized and compressed format called Multidimensional OLAP (MOLAP). It also allows the flexibility of storing the data (in part or completely) in a relational database as Relational OLAP (ROLAP) or in a hybrid mode called Hybrid OLAP (HOLAP). MOLAP provides significantly better performance than ROLAP and HOLAP.

Multidimensional data is inherently sparse by nature. For example, you do not buy every product in every branch of a retailer on every day. SQL Server, unlike most OLAP systems, does not store these NULL values, resulting in a significant reduction in database size, protection from data explosion, and a resulting performance improvement. Many OLAP systems waste a substantial portion of query processing time in aggregating data from cells with NULL values that will subsequently yield NULL results. SQL Server 2008 Analysis Services uses a technique called Block Computation that exploits cube sparsity to improve query performance by focusing only on the non-NULL data. This can improve query performance by orders of magnitude and therefore allow a finer granularity of analysis.

Another area where SQL Server delivers superior performance is attribute-based hierarchies. Typically, databases contain hierarchies that share common attributes. In most OLAP systems, these common attributes must be duplicated for each hierarchy, but SQL Server provides attribute-based hierarchies that avoid the need for any duplication and improve performance and scalability.

Writeback is a core functionality in Analysis Services that allows the user to modify cell values. It is commonly used in planning, budgeting, and forecasting applications. Previous versions of Analysis Services required writeback data to be stored in ROLAP format. SQL Server 2008 Analysis Services allows writeback data to be stored in MOLAP format resulting in significantly better performance for query and writeback operations.

Proactive caching provides MOLAP performance with real time analytics. This is achieved by keeping an up-to-date copy of data organized for high-speed

access using the UDM structure as its foundation. This prevents users from overloading the relational database by providing a high performance, transparent, synchronized aggregate cache.

Extend Solutions with Comprehensive Analytics

When thinking of OLAP most people think of a storage and aggregation engine. This is also valid for Analysis Services. However, Analysis Services takes the analytical platform to a new level offering more advanced features than those traditionally related to OLAP. This enables organizations to accommodate multiple analytical needs within one solution offering so much more than a traditional OLAP platform. In this effort, the Unified Dimensional Model (UDM) plays a central role, providing extensive analytical capabilities.

Unified Dimensional Model

The UDM was a new concept for Analysis Services that was introduced with the release of SQL Server 2005. The UDM provides an intermediate logical layer between the physical relational database used as the data source and the proprietary cube and dimension structures that are used to resolve user queries. In this way, you can think of the UDM as the centerpiece of the OLAP solution. However, as mentioned above the concept of the UDM impacts multiple aspects of the Analysis Services solution. One of the key benefits of the UDM is the ability to combine the flexibility and richness of the traditional relational reporting model with the powerful analytics and superior performance of the classic OLAP model. In addition, a wide range of advanced business intelligence capabilities have been included in the model to provide best of breed relational and OLAP analysis and to further allow organizations to easily extend solutions leveraging the unique Key Performance Indicator Framework as well as the sophisticated predictive analytic capabilities that are all delivered through one approach: The UDM.

Central Manageability of Key Enterprise Metrics

In SQL Server 2008 Analysis Services enterprise wide Key Performance Indicators (KPI's) can be centrally stored and managed. This provides a central repository for users to access key enterprise metrics through a variety of applications including Microsoft Office PerformancePoint Server 2007, Microsoft Office Excel 2007, Microsoft Office SharePoint Services 2007, and Microsoft SQL Server Reporting Services.

Predictive Analysis

Traditional data analysis looks at historical data and quickly returns results based on this data. However, many questions asked by business users cannot be answered by this sort of analysis as they are not looking for the results of what has happened, but instead they are looking for predictions of what might happen. The ability to predict future trends is potentially one of the most

important factors in the success of any organization, but it is not as simple as extending a trend line. Members need to be grouped to create clusters that behave in a similar way; contributing factors need to be assessed to measure their effect on a particular result; interdependencies need to be identified.

Data mining algorithms in Analysis Services provide this predictive analysis and SQL Server 2008 Analysis Services improves the data mining algorithms to enable analysis that is more extensive.

Microsoft SQL Server Data Mining Add-Ins for Office 2007

The Data Mining Add-Ins for Office 2007 is a set of easy to use data mining capabilities that allows you to access data mining functionality from within Office 2007, thus enabling predictive analysis at every desktop. Being able to harness the highly sophisticated data mining algorithms of Microsoft SQL Server 2008 Analysis Services within the familiar environment of Office, business users can easily gain valuable insight into complex sets of data with just a few mouse clicks. Designed with the end users in mind, the Data Mining Add-Ins for Office 2007 empowers end users to perform advanced analysis directly in Microsoft Excel and Microsoft Visio.

There are three individual components:

- **Data Mining Client for Excel** enables you to create and manage an entire Analysis Services data mining project from within Excel 2007.
- **Table Analysis Tools for Excel** enables you to use the powerful Analysis Services data mining capabilities to analyze data stored in Excel spreadsheets.
- **Data Mining Templates for Visio** enables you to render decision trees, regression trees, cluster diagrams, and dependency nets in Visio diagrams.

Drive Actionable Insight through Familiar Tools

Powerful analytical solutions provide no business benefit if the information is not easily accessible by all users. SQL Server 2008 Analysis Services goes beyond business users and provides analytical information to everyone in the organization using the familiar tools in Microsoft Office. Further client interfaces can be developed using the open architecture of SQL Server 2008 Analysis Services and developers can take advantage of the extensibility of the product to expand its functionality.

Optimized Office Interoperability

The 2007 Microsoft Office system provides optimized interoperability with SQL Server 2008 Analysis Services. Information is provided on the desktop through familiar tools to extend the reach of your analytical information. For example, Excel 2007 is a fully functional, rich Analysis Services client, while Microsoft Office PerformancePoint Server 2007 Analytics provides a thin Analysis Services client. The following 2007 Office system components provide Analysis Services interoperability:

Microsoft Office Excel

Excel 2007 is a fully functional Analysis Services client. Excel 2007 provides functionality in the following areas:

- Excel provides access to data stored in Analysis Services OLAP cubes. Excel provides pivot tables that present multidimensional data to the user and allow the user to slice and dice the data. The server performs the processing, and the results are cached on both the server and the client to enhance performance.
- Excel brings Analysis Services features and analytical capabilities such as KPIs, calculated members, named sets, actions and translations to users.
- Excel can use the Data Mining Add-Ins for Office 2007 to provide rich predictive and statistical analysis to end users.
- Excel can add automatic analysis features, such as highlighting exceptions where data seems to differ from patterns in other areas of the table or data range, forecast future values based on current trends, analyze what if scenarios, and determine what needs to change to meet a specific goal.
- Reporting Services can create reports from Analysis Services data and render them as Excel spreadsheets to increase availability to end users.

Figure 4 shows the Excel PivotTable being used for client access Analysis Services data.

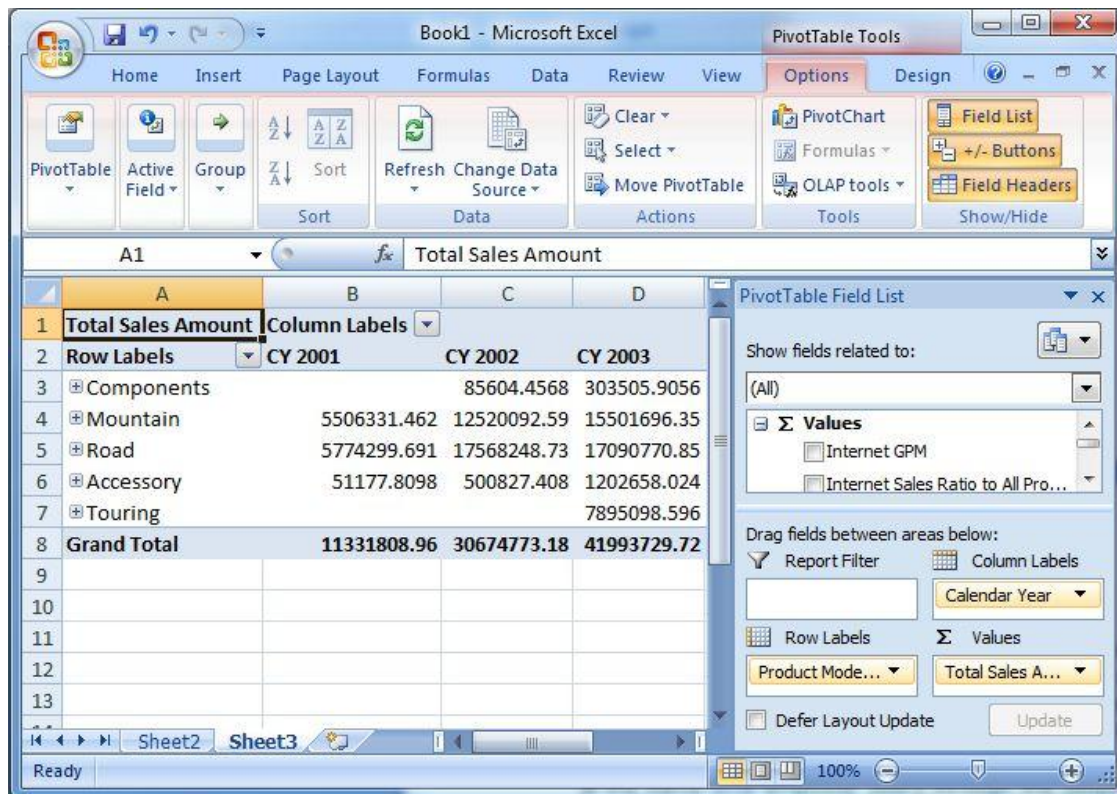


Figure 4

Microsoft Office Word

Reporting Services can create reports from Analysis Services data and render them as Microsoft Office Word documents to increase availability to end users. These reports can then be edited directly in Microsoft Office Word.

Microsoft Office Visio

You can use Microsoft Office Visio to annotate, enhance, and present data mining graphical views. With SQL Server 2008 and Visio 2007, you can:

- Render decision trees, regression trees, cluster diagrams, and dependency nets.
- Save data mining models as Visio documents embedded into other Office documents, or saved as a Web page.

Microsoft Office SharePoint Server 2007

A comprehensive collaboration, publishing, and dashboard solution that you can use as a centerpiece for providing one central location for placing all your enterprise-wide Analysis Services data, so that everyone in your organization can view and interact with relevant and timely analytical views, reports and KPIs.

Microsoft Office PerformancePoint Server 2007

An integrated performance management application that employees can use to monitor, analyze, and plan business activities based on the data provided by SQL Server 2008 Analysis Services. Office PerformancePoint Server 2007 provides scorecards, dashboards, management reporting, analytics, planning, budgeting, forecasting, and consolidation functionality to provide extensive performance management capabilities.

Rich Partner Extensibility

SQL Server 2008 provides an open architecture allowing developers to build solutions on top of Analysis Services and extend its functionality. Analysis Services includes stored procedures to provide straightforward access to Analysis Services functionality to external programming languages. Stored procedures provide cross-language exception handling, versioning and deployment support.

Data mining represents any form of statistical analysis and, as this field is constantly evolving, new data mining algorithms could make an analytical system obsolete. Analysis Services supports plug-in algorithms to extend data mining capabilities and allow the addition of new data mining algorithms by third parties or in-house developers.

Open Embeddable Architecture

Many organizations will require a customized client interface or they will need to consume the Analysis Services data in another service or application.

Analysis Services has long supported OLE DB for OLAP, ADOMD, and ADOMD.Net, but this is extended by SQL Server 2008 Analysis Services to

expose data using the XML for Analysis (XML/A) standard. Each Analysis Services server is now a provider of web services and, as such, this makes it straightforward to integrate analytical data into modern applications.

Conclusion

Microsoft SQL Server 2008 Analysis Services builds on a strong foundation of analytical tools to provide a truly enterprise scale solution. Performance and scalability are substantially improved with faster processing, improved large database backups, and new monitoring functionality. Data is more available to users by combining data marts into a UDM and centralizing access and manageability of key enterprise metrics. Analytical capabilities are extended with the predictive abilities of an enhance data mining toolset.

Access to data is not sufficient to drive this information into the business. Users need familiar tools and application developers need to be able to integrate the data into their applications. Analysis Services provides optimized Office interoperability to provide a familiar interface and an open, embeddable architecture to allow developers to integrate the data.

For more information:

<http://www.microsoft.com/sql/>

Please give us your feedback:

Did this paper help you? Tell us on a scale of 1 (poor) to 5 (excellent), how would you rate this paper and why have you given it this rating? For example:

- Are you giving it a high rating because it has good examples, excellent screenshots, clear writing, or another reason?
- Are you giving it a low rating because it has poor examples, fuzzy screenshots, unclear writing?

This feedback will help us improve the quality of white papers we release.

[Send feedback.](#)

The information contained in this document represents the current view of Microsoft Corporation on the issues discussed as of the date of publication. Because Microsoft must respond to changing market conditions, it should not be interpreted to be a commitment on the part of Microsoft, and Microsoft cannot guarantee the accuracy of any information presented after the date of publication.

This white paper is for informational purposes only. MICROSOFT MAKES NO WARRANTIES, EXPRESS OR IMPLIED, IN THIS DOCUMENT.

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 Microsoft Corporation.

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

© 2008 Microsoft Corporation. All rights reserved.

Microsoft, PowerShell, SharePoint, SQL Server, Visual Basic, Visual C#, Visual Studio, Windows, Windows Server, and the Server Identity Logo are trademarks of the Microsoft group of companies.

All other trademarks are property of their respective owners.