

Microsoft SQL Server Plug-in 6.82

User's Guide

Revision: This manual is updated for Version 6.82
Software Version: 6.82 (March, 2011)

Copyright © 1997-2010. All rights reserved.

The software manufacturer makes no representations or warranties with respect to the contents hereof and specifically disclaims any implied warranties of merchantability or fitness for any particular purpose. Furthermore, the software manufacturer reserves the right to revise this publication and to make changes from time to time in the content hereof without obligation of the software manufacturer to notify any person of such revision of changes. All companies, names and data used in examples herein are fictitious unless otherwise noted.

No part of this document may be reproduced, transmitted, transcribed, stored in a retrieval System or translated into any language including computer language, in any form or by any means electronic, mechanic, magnetic, optical, chemical or otherwise without prior written permission.

All other products or company names mentioned in this document are trademarks or registered trademarks of their respective owners.

Acknowledgements: Two encryption methods, DES and TripleDES, include cryptographic software written by Eric Young. The Windows versions of these algorithms also include software written by Tim Hudson. Bruce Schneier designed Blowfish encryption.

"Part of the software embedded in this product is gSOAP software. Portions created by gSOAP are Copyright (C) 2001-2006 Robert A. van Engelen, Genivia inc. All Rights Reserved. THE SOFTWARE IN THIS PRODUCT WAS IN PART PROVIDED BY GENIVIA INC AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHOR BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE."

The Agent, Agent Console, and Vault applications have the added encryption option of 128/256 bit AES (Advanced Encryption Standard). Advanced Encryption Standard algorithm (named Rijndael, pronounced "Rain Doll") was developed by cryptographers Dr. Joan Daemen and Dr. Vincent Rijmen. This algorithm was chosen by the National Institute of Standards and Technology (NIST) of the U.S. Department of Commerce to be the new Federal Information Processing Standard (FIPS).
See: <http://csrc.nist.gov/encryption/aes/round2/r2report.pdf> for details.

The Agent and Vault applications have the added security feature of an over the wire encryption method.

Contents

1	Introduction to the SQL Plug-In.....	1
1.1	Overview	1
1.2	Features in Version 6.x.....	2
1.3	Clustering.....	3
1.4	Supported Platforms	3
1.5	About this Guide	3
1.6	Release Notes and Online Help.....	4
2	Installing the SQL Plug-In.....	5
2.1	Licensing.....	5
2.2	Cluster Awareness – Cluster Plug-In	6
2.3	Main Features of the Cluster Plug-In.....	6
2.4	Installation Setup Recommendations.....	7
3	Working with Backups.....	8
3.1	Backup Information	8
3.2	Performing Backups.....	9
3.3	Transaction Logs.....	10
3.4	Simultaneous Backups	11
3.5	Creating a New Agent.....	11
3.6	Resolving Failed Server Connections	12
3.7	Creating a Backup Job	12
3.8	Scheduling a Backup	17
3.9	Verifying the Backup	18
3.10	Viewing Job Properties.....	19
3.11	Error Messages.....	20
3.12	Re-Registering an Agent.....	21
3.13	Removing a Database from the Backup	21
4	Working with Restores	22
4.1	Notes About Restores.....	23
4.2	Deleting a Database.....	23

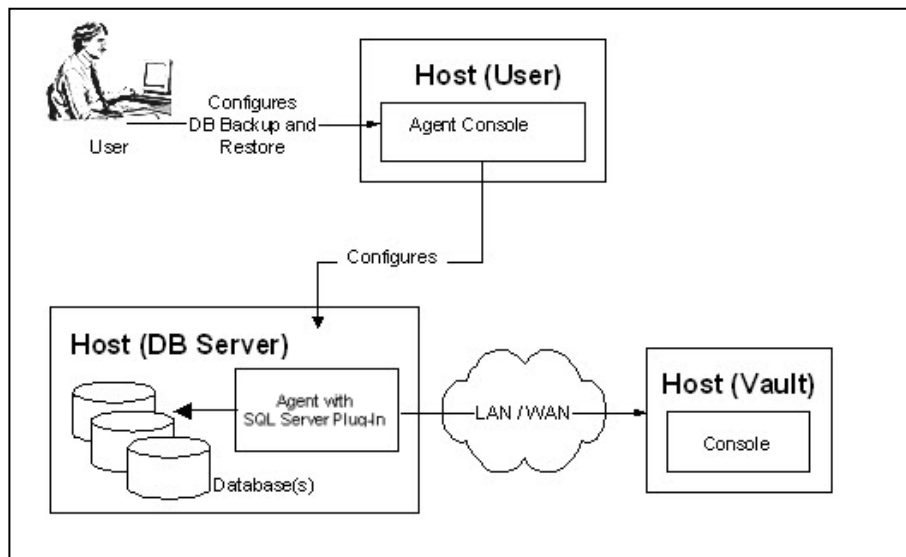
4.3	About the With Replace Option	24
4.4	Restoring Databases	24
4.5	Database types	24
4.6	Recovering System Databases	26
4.7	Recovering the MASTER Database	26
4.8	Recovering the MASTER database on an SQL 2008 Cluster	27
4.9	Recovering a Single User database	28
4.10	Restore Database with Recovery	29
4.11	Microsoft SQL Bare Metal Restore	29
4.12	Recovering from a Worst Case Disaster	30

1 Introduction to the SQL Plug-In

This chapter provides an overview. It also describes new features, release notes and help.

1.1 Overview

The Microsoft SQL Server Plug-In allows a user to perform a database backup on a Microsoft SQL Server. The Plug-In is installed on top of the Windows Agent on the database host to perform the backups, either on demand (ad hoc), or scheduled.



This diagram illustrates the basic product implementation:

A user, typically a DBA (with Administrator rights), configures the backup via the Agent Console application.

Agent Console configures the Agent, which may reside on a different host than Agent Console. However, the Agent and Plug-In must reside together on the SQL Server system.

The user can now schedule a backup of the Database(s), at which time the Agent, with the aid of the Microsoft SQL Plug-In, sends the database information to the Vault.

1.2 Features in Version 6.x

The Plug-In application has a separate license requirement, but is installed along with the Agent. It is operated and configured within the Agent Console program GUI.

- Support for Windows Server 2008 Standard, Datacenter, and Enterprise.
- Support for Microsoft SQL Server 2008 (SP2) and for Microsoft SQL Server 2008 R2.
- Support for Microsoft SQL Server 2005 (SP3). The Plug-In API on Microsoft SQL Server 2005 uses ODBC, and is also backward-compatible with Microsoft SQL 2000.
- The database(s) to be backed up must run on a single host.
- The Windows Agent and the Microsoft SQL Server Plug-In must run together on the same system that runs Microsoft SQL Server. The Windows Agent and the Microsoft SQL Server Plug-In must always have the same version number.
- If a database host is completely lost, the database software may be loaded, and the database completely restored, after a full system restore. See Chapter 4 ("Restores").
- Hot backup - occurs without taking the database down.
- A Restore can be to the same or an alternate database location, or to a flat file.
- Provides full and incremental backups.
- Microsoft SQL Server 2000/2005 support a default instance, and up to 15 named instances (i.e., "computer_name\instance_name").
- The Plug-In can implement larger delta block sizes (32 KB). This improves performance for backups. The larger block sizes result in the overall delta changes that are sent to the Vault being smaller in total size. Thus, less disk storage space is used.
- Version 6.x includes much better support for delta. The Microsoft SQL Server backup stream is now processed logically rather than physically. This will greatly reduce the size of delta backups.
- The Plug-In uses the high-performance Microsoft-defined Virtual Backup Device Interface (VDI).
- The Plug-In will dynamically determine what SQL Server client is installed. With SQL Server 2000/2005 client, the Plug-In will support access to all instances of the server installed (regardless of what version they are).
- Backups and Restores can handle multiple SQL Databases per Job. The Restore destination can be the original database, another database, or a file. You will be prompted if the Restore is about to overwrite an existing database.
- Upgrades of earlier Plug-In versions are supported, and do not require any reconfiguration and/or re-seeding of existing data.
- Licensing (obtaining and validating) is different when connecting to a Vault lower than 5.53. See section 2.1 of this Guide.

1.3 Clustering

On Windows 2003 Enterprise Edition, it is possible to create a two-node cluster for SQL Server 2000/2005.

On Windows Server 2008 Enterprise Edition, it is possible to create a two-node cluster for SQL Server 2005/2008.

Clustering is supported for Windows Agents, with a separately licensed Cluster Support Plug-In (see section 2.2). The main function of the Cluster Support Plug-In is for the Agent on a Microsoft SQL Server, which has a virtual IP address in the cluster, to be able to follow the server when it fails over to another node in a cluster.

The Agent can still access its configuration (on a shared drive), and scheduled backups will occur as usual, without it looking like a "different" backup and causing a reseed.

1.4 Supported Platforms

The Microsoft SQL Server Plug-In is supported on currently supported Windows Agents. See the Windows Agent Release Notes for the latest versions of supported Windows.

Also, the Shipping Products Chart lists all Software Supported Products.

These documents are available from your service provider.

The 6.82 Microsoft SQL Server Plug-In is used in conjunction with Vault versions 5.53, and 6.04 and above.

1.5 About this Guide

This guide should be used in conjunction with other manuals that describe the Windows Agent and Agent Console.

Agent Console Operations Guide

- Installing the main Agent Console software (GUI)
- Using the Agent Console GUI – Workspace, Agents, Agent Configurations, Jobs, Safesets, Catalogs and Log files
- Performing backups – Types, Seeding, Mapped drives and databases, Options, Tape, Retentions, Notification, Expiration, Scheduling and Ad-hoc (on demand) Backups
- Report Logs – Creating and Managing Log files
- Data Security – User Authentication and Encryption
- Open File Backup – Shared files, OTM, and OFM
- Troubleshooting and Command-Line Interface

Windows Agent User Guide

- Agent for Windows Installation
- Using the Agent for backups and restores
- Windows Systems Recovery
- Cluster Support Plug-In

1.6 Release Notes and Online Help

Release notes provide “up to the minute” information about the released product. Release notes contain an overview of new features, any known defect (bug) fixes incorporated since the last release, a description of any known issues, and a section on product support. Release notes are available from your service provider.

Agent Console (Windows Agent Console/Web Agent Console) provides online help, which contains information similar to the contents of this User Guide.

There is also context-sensitive “What’s This” help on each Windows Agent Console GUI screen. You can access the context-sensitive “What’s This” help by clicking the Help icon (question mark) in the Agent Console application. Note: If the Windows Agent Console F1 Help screen is open (even minimized), the “What’s This” help will not be active. The F1 help must be closed for the “What’s This” help to function properly.

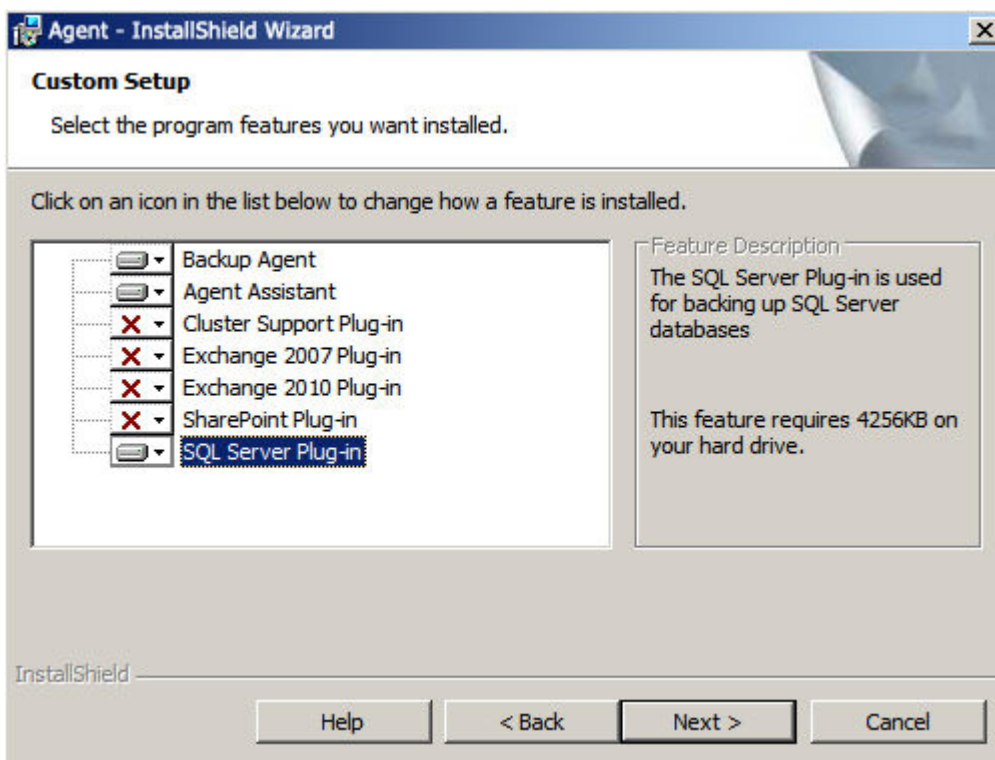
2 Installing the SQL Plug-In

The Microsoft SQL Server Plug-In integrates into the existing Agent architecture, providing the user with the ability to backup/restore Microsoft SQL Server databases to/from the Vault.

The Microsoft SQL Server Plug-In is installed during the Windows Agent installation. See the Agent for Microsoft Windows User's Guide. The Plug-In can be installed when installing the Agent, or it can be installed later, by re-running the installation with the Modify selection.

The Cluster Support Plug-In can be installed the same way.

For supported Vault, Web and Windows Agent Console versions, see the release notes.



Note: For more information about functions such as “Create a New Agent”, “Create a Backup Job”, “Scheduling Backups”, and Disaster Recovery, refer to the Agent Console Operations Guide.

2.1 Licensing

The Microsoft SQL Server Plug-In requires a separate license if you are connecting to a Vault that is version 5.27 or lower. If you are connecting to a Vault that is version 5.53 or more, the license is supplied automatically from the Vault.

To apply a 5.27 license, install the Microsoft SQL Server Plug-In. Next, enter and validate the license from Agent Console through the Agent Configuration > Plug-In tab.

Obtain a license key from your service provider.

1. In Agent Console, under <Agent name> → Agent Configuration → Plug-In, select the Microsoft SQL Server Plug-In, and enter the license key information.
2. Press Test to ensure the license is correct/valid. If not, re-enter the data.
3. Press Set to save the information. Click OK to finish.
4. From the Agent Console program, if you create a new Job, you should now be able to see "Microsoft SQL Server" under Backup Source Type. See section 3.5 of this manual, "Create a Backup Job".

(Note that the "Plug-In Options" button is not used for the Microsoft SQL Server Plug-In.)

2.2 Cluster Awareness – Cluster Plug-In

The Microsoft Server Clustering Services (MSCS) Cluster Support Plug-In is available for the Windows Agent.

The main function of the Cluster Support Plug-In is for the Agent on a Microsoft SQL or Microsoft Exchange Server, which has a virtual IP address in the cluster, to be able to follow the server when it fails over to another node in a cluster. The Agent can still access its configuration (on a shared drive), and scheduled backups will occur as usual, without it looking like a "different" backup and causing a reseed.

The Cluster Support Plug-In is supported on Windows Agents. It requires a separate license.

2.3 Main Features of the Cluster Plug-In

A user can connect to an Agent (with a Plug-In and proper license) on a Virtual Server or Local machine (a node) via IP or host name.

The Virtual Server Agent can back up virtual server shared data without reseeding, or in case of a failover.

Once created, Jobs (on a shared drive belonging to a virtual server) can be used by all Agents on the cluster.

Scheduling of virtual server backups is handled between node Agents without schedule overlapping. The configuration files are located on the drive owned by the virtual server.

Each physical node in a cluster configuration requires a separate installation of the Agent, each with a separate Plug-In and license. You also need to enter the licenses of the Plug-Ins on the Virtual Server. The Cluster Plug-In should not be visible on the Virtual Server.

When you first configure an Agent on a Virtual Server, you will be prompted for a location on a drive that the Virtual Servers see. So, after a failover, the Agent configuration will still be available to all servers owned by the virtual server.

The icons that represent the servers in Agent Console are different (e.g., a "regular" local Agent versus a Virtual Server Agent).

Note: Icons only change after you have selected the Agent. Virtual Server Agents have white icons, and regular Server Agents have blue icons.

2.4 Installation Setup Recommendations

1. Install the Agents and Plug-Ins on the Physical Nodes.
2. Set Cluster, SQL, and Exchange Plug-Ins on the Physical Nodes.
3. Create a new Agent for your Exchange or SQL cluster on the Virtual Node using the IP address or host name.
4. From the newly created Agent, double click the Global file to open the "Virtual server shared area" window. From here you must select a drive letter for your SQL or Exchange Cluster. Click OK. This will launch the Agent Configuration Window.
5. Here you must specify all of your Vault connection information. Click OK when complete. Note: Once the folder for configuration files on the shared disk has been created you will not be prompted again for its location.

Note: You must create the backup Job from the Virtual Node in order to use the Cluster Plug-In failover features even though you can perform backups from the Virtual Node as well as the Physical Node.

3 Working with Backups

This section describes how to create an Agent and Job, schedule the Job and check for completion and errors.

3.1 Backup Information

Before performing an Microsoft SQL Server database backup, make sure that you have all of the information such as names, locations, passwords, etc., that the wizard will require. You can use the following table for reference.

Note: Backups and restores must be run with Administrator rights.

System Requirement	Customer/User supplied value	Comments
New Job name	Name =	Job used to communicate with the Agent that has the Microsoft SQL Server Plug-In
Vault profile	Profile =	Profile of an existing (already created) Vault. Chosen from a pull-down menu.
Backup Source Type	Type =	Select "MS SQL Server" from the pull-down menu.
Microsoft SQL Server Connection Information	Computer = Instance = Authentication = Windows or SQL User Name = Password =	Dependent on user's configuration and DBA setup. Validates the fields, and allows access to the database. For SQL authentication, the maximum length of the password is 31 characters.
Microsoft SQL Server Database selection	Database =	Pull-down list of all available databases
Network Account Information	User Name = Password = Domain =	Supplies Windows credentials to the server. Enables backup process to run across the network.
Quick File Scanning		Set on or off.
Backup Type	Choose "Full" or "Incremental"	First backup is always Full.
Encryption type	Type = Password =	If you select a type, you must supply a password.
Log options	Options = Detail level = Copies =	User defined
Start the backup	Immediate = Schedule =	Can use the Scheduling Wizard

3.2 Performing Backups

To back up your Microsoft SQL Server you will first need to add to your workspace the Agent with the Microsoft SQL Server on it. Then create a new Job in that Agent, using "MS SQL Server" type. It is not recommended to do multiple Jobs, at the same time, on the same database.

The Backup Wizard will direct the user through the process to create a Job. The steps are described in detail later in this chapter, but briefly, the steps are:

1. From the Agent Console GUI program (which communicates with the Agent and the Plug-In), create a new Job.
2. From the "Backup Source Type" screen, select "MS SQL Server".
3. Select a Vault profile where the Job will be targeted.
4. Supply a user-defined name for this new Job.
5. Choose an instance from the "Select Server Instance" list. The list will have the first entry pre-selected as <Default>, designating the default instance. Other entries will be present only if they can be retrieved.
6. Next, choose the Windows or SQL Authentication option. (See the Microsoft SQL Server Help section on Windows/SQL Authentication for more information.)

If you choose Windows Authentication, access is controlled through the user name and password from the Windows logon. In this case, User Name and Password on the screen are not used.

If you choose SQL Authentication, access is controlled through the User Name and Password on this screen.

Note: For SQL authentication, the maximum length of the password is 31 characters.

Select one or more databases for backup. You can select all the databases shown in the list of databases. These only represent the ones shown in the list, rather than all existing databases. If you choose to back up "all existing databases" for this instance, it will use all of the current ones. So if databases are added or removed from the instance, the next backup will automatically choose all of them. There can also be more than one Job doing backups on different databases.

It is recommended that the system (master, model and msdb) databases be backed up in one Job, and the other (user) databases be backed up in one or more Jobs. During a recovery process, the master database must be recovered first, by itself. Then the other system databases can be restored.

7. Enter network account information to enable the backup process to run across the network – user name, password, and domain. This allows the Agent to connect to a local or remote database.
8. Select or de-select the "Quick file scanning" option. Note that the "deferred" option is not available here; the backup must be completed in one step.

9. If desired, select an encryption type, and supply an encryption password. Note that if you lose the password, the data will be unrecoverable.
10. Select logging options and level of detail, and log copies.
11. Start the backup (immediately), schedule it for later, or exit the Wizard without doing a backup (but do save the backup Job).

Log files are created on the Server machine (with the Agent), in directories with the Job name. They are normally viewed from the Agent Console application (GUI) screen.

Backup Notes:

- A Microsoft SQL Server 2000/2005 does not allow incremental backups for databases if the recovery model of database is set to 'Simple'.
- Microsoft SQL Server Plug-In has standard backup/restore permission requirements for the logged user.
- Backup Permissions: BACKUP DATABASE and BACKUP LOG permissions default to members of the db_owner fixed database role, who can transfer permissions to other users, and to members of the db_backupoperator fixed database role.

3.3 Transaction Logs

It is important when running Microsoft SQL to choose whether you are running a Simple Recovery Model, or a Full Recovery Model.

The Simple Recovery Model will provide you with simple and efficient backups and restores. This does not back up the transaction logs with the database. Instead it deletes them when transactions that they contain are committed to the database. This is similar to a circular logging option. You must handle the truncation and reclamation of transaction log space yourself.

The Full Recovery Model uses the Agent to manage the backups and logs. You need to run regular (full) and incremental backups. Incremental backups do not reduce the size of the transaction logs. To manage the transaction log space, the logs must be truncated using SQL maintenance for the space used by the logs to be reclaimed.

When the logs are truncated, the space left by previous obsolete transactions is reused. SQL logs can be truncated after an incremental backup.

If you want to use the Agent to manage the backups and logs, you need to use the Full Recovery model, and run regular full and incremental backups, with transaction log truncation.

When using the Full Recovery Model, these other factors must be considered:

- An incremental backup Job must be scheduled.
- The level of activity on the Database, as this affects the rate at which Microsoft SQL transaction logs will grow.

- Amount of space required for the local transaction logs and how their growth is managed.
- Amount of space required on the Vault.
- The complexity of the restore procedure (a single-pass restore is still supported).

Please consult your Microsoft SQL documentation on how to properly manage your Database and transaction logs.

3.4 Simultaneous Backups

Scheduling simultaneous backups using different backup jobs backing up same database instance may cause your backups to fail with error "NONE-E-6406 Cannot find database to back up".

As a good practice, the “staggering” of scheduling different jobs that back up the same instance is recommended.

3.5 Creating a New Agent

The Agent Console Operations Guide provides details about creating Agents and Jobs. Here is some “quick reference” material about the process:

1. In the Agent Console GUI screen, select a Workspace, and then right-click to display the submenu.
2. Click New Agent. The Agent Properties window opens.
3. Enter a Description (the name of the Agent) and the Network Address of your Microsoft SQL Server. The Network Address can be an IP address or a valid host name.
4. Enter the user name, password and domain. Note that the password is case-sensitive.
5. After the Agent Information and Authentication data has been entered, click Get Status. If the information is validated, your data is displayed in the Agent Status window. If the information is not validated, a message from the Agent Console application appears. Check your information and revise it as required. Once again, click Get Status.
6. Click OK.

When configuring the Agent Console application to connect to a remote Agent, you need to specify a particular user name and password. For details, refer to “Agent Access Privileges” in the Agent Console Guide or Help. If an authentication password changes (e.g., the user listed on the Agent Properties dialog changes his or her system password), the Agent Properties dialog must be updated.

The Authentication information here (user name/password) allows the Agent to access the data. It also allows you to browse to see the data. Typically the authentication can be an Administrator

password. It can also be the same user name/password that the backup Agent needs to "Run As" a particular user. See "Run As" in "Create a Backup Job".

3.6 Resolving Failed Server Connections

If you created a Job using Windows Authentication and you did not supply a Domain name when you created the Job, the Job will work properly, even after a failover.

But after a failover to another virtual node in the cluster, if you try to create a new Job without the Domain name in Agent Properties, it will fail to connect.

Your Job must have a Domain name in Agent Properties in order to use Windows authentication after a failover.

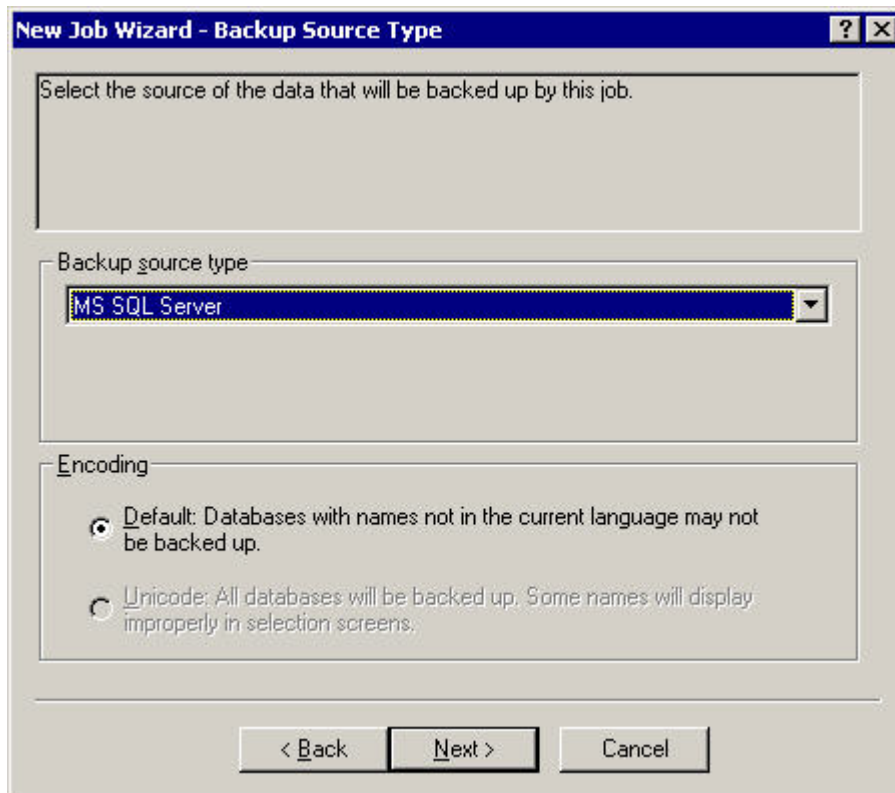


SQL Server authentication should work correctly in these cases.

3.7 Creating a Backup Job

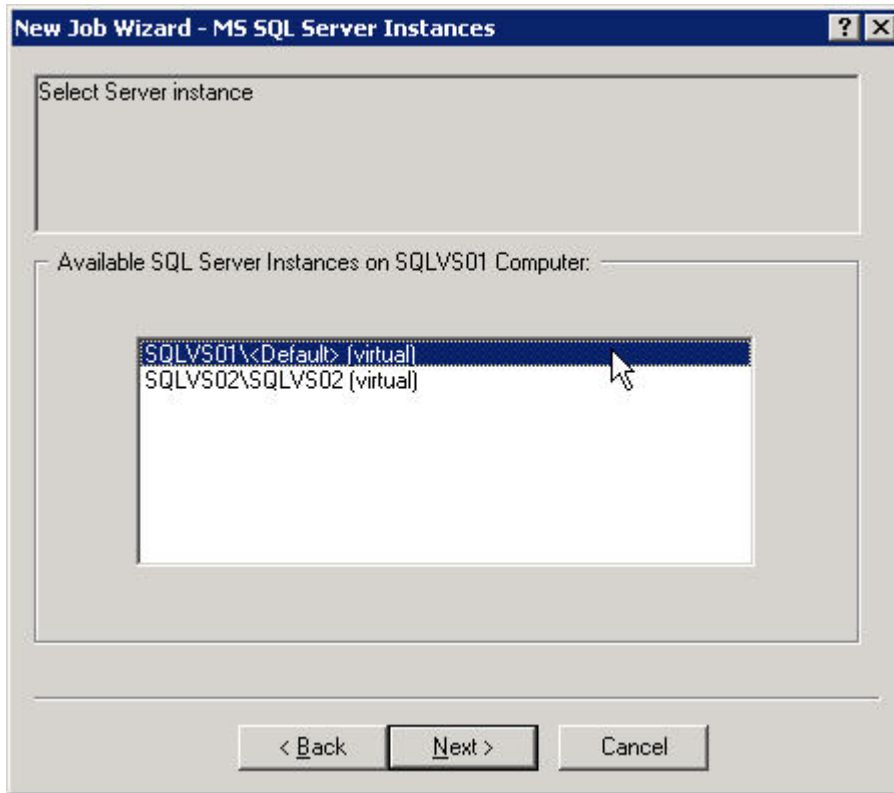
The Agent Console Operations Guide provides a more detailed look at, and understanding of, creating Agents and Jobs. This section of the User's Guide is meant as a "quick reference" for the user.

1. Right-click on the Agent from the previous section, and select "New Job". The New Job Wizard displays. Enter a name for this Job, and click Next.
2. Next, select a Backup Source Type (MS SQL Server) from the dropdown list. (Your list may appear different from this diagram.) The Microsoft SQL Server Plug-In must be installed and licensed in order to appear here. See Section 2 of this manual. Click Next to continue.

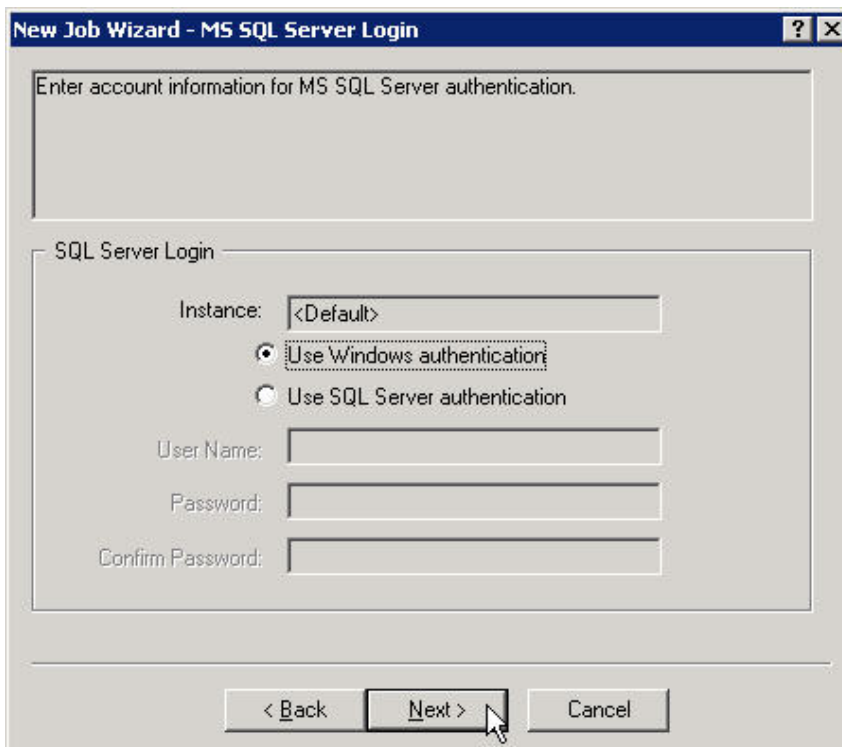


Note that the Encoding type is "Default". Unicode is not supported here.

3. On the next screen, select a Vault where the backup will go, and then click Next.
4. On the next screen, enter a name for this Job.
5. Select a Microsoft SQL Server Instance. This presents a list of the SQL Server instances available for the client installed. The list will always begin with a pre-selected entry named <Default>, designating the default instance. Other entries will be present only if they can be retrieved. If there is only one, you must select <Default>.

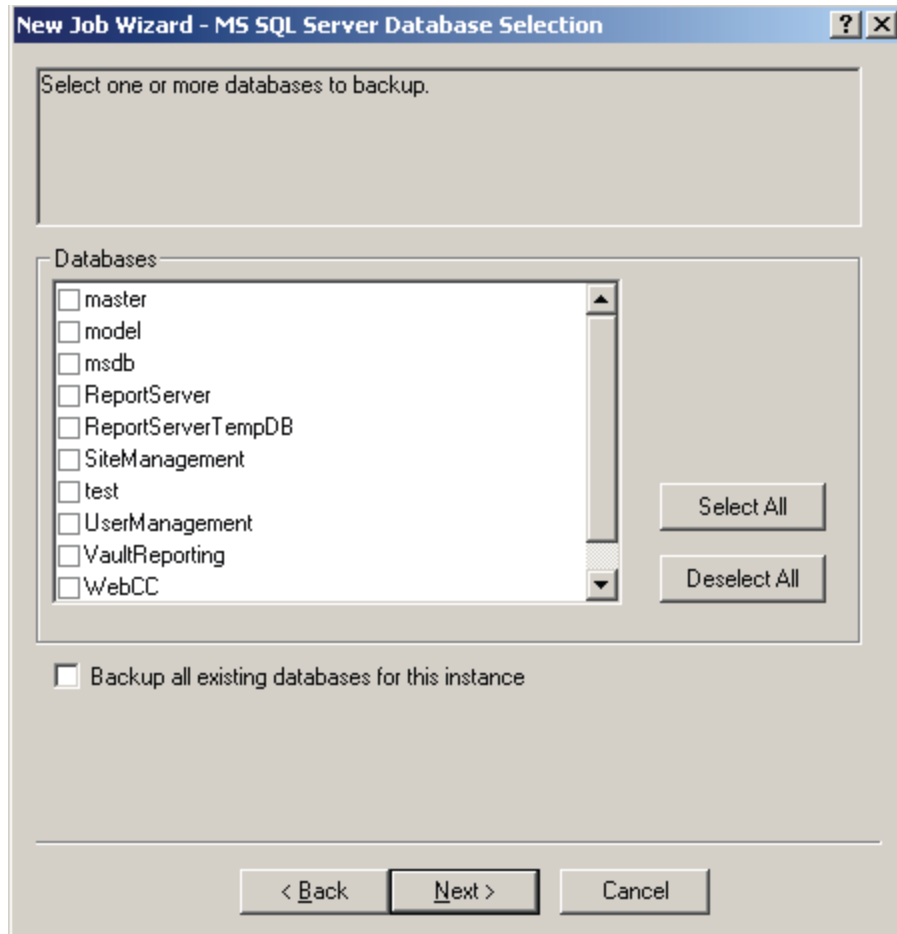


6. SQL Server Login. This screen will allow the user to select the authentication type (Windows or SQL, depending on how it is set on the SQL server) and enter a user name and password to connect to the selected instance. After entering valid credentials, the user will proceed to the database selection dialog.



Note here that Computer and Instance names are presented in read-only format, for information purposes. Also, for SQL authentication, the maximum length of the password is 31 characters.

7. Database selection. Select a specific database to back up. See Section 4.2.1 for specific database types.



It is strongly recommended that the system (master, model and msdb) databases be backed up in one Job, and any other (user) databases be backed up in one or more separate Jobs.

During a recovery process, the master database must be recovered first, by itself. Then the other system databases can be restored.

8. Run As. Enter network account information to enable the Agent Backup process to run as a particular user, to be able to access the data. This would only be needed if the Authentication information (user name/password) in the Agent properties did not have enough rights for the Agent Backup process to access the database.

A Microsoft SQL Job can back up and restore successfully without "Run As" credentials. To run a Backup/Restore Job with the same privileges that the Agent has, leave the user name and password blank in the "Run As" credentials.

New Job Wizard - Run As

Enter network account information that will enable the backup process to run.

Run as

User name: administrator

Password: *****

Confirm password: *****

Domain: adcluster

< Back Next > Cancel

If this is a concern or security risk, run the Agent service with a Windows account that cannot read/write to the database (to prevent accidental access), and supply a user name in the “Run As” parameters.

- Options. Quick File Scanning reduces the amount of data read during the Backup process. Any file streams that are deemed unchanged since the last Backup, are skipped over. Without Quick file scanning turned on, files are read in their entirety.

The first time that you back up an SQL database, it will always be a backup type of “full”. You cannot choose between full or incremental backups in this case.

For subsequently scheduled backups, you will have a choice. (See Section 3.6 of this manual.)

The difference between "Full" and "Incremental" is that "Full" will back up the entire database, while "Incremental" will only back up the transaction log files. During a restoration, the log files will be played back to achieve the most up to date restore since the last (full) backup.

- Encryption. You may choose to use encryption, or not. Select one of the types shown in the list. Choose your own encryption password. This password is not stored anywhere on the system. If you lose the encryption password, your data will be inaccessible.
- Log files. Choose one of these options:
 - Create log file: Check this box to instruct the system to generate log files for each Job executed. These printable log files contain the start-connect-completion and disconnect times, file names (i.e., the name of each file that was copied during a Backup process) and any processing errors.

- Log detail level: You can select the detail of logging to go from least to most: None, Files, Directory, Summary. More detailed logging creates larger log files, and is normally used only for troubleshooting problems.
 - Changing the detail level only affects log files that are created from that point on. It does not affect any previously created log files.
 - Automatically purge expired log files only: You have a choice of either automatically purging expired log files, or keeping a selected number of them before they get deleted (oldest one first).
 - Automatic purging will delete the log file associated with a backup when that backup (safeset) is deleted.
 - Keep the last number of log files: You may select a number here, whereby the system will keep that number of log files, which are associated with that Backup. When that number is reached, the oldest log file will be deleted to make way for the newest one.
12. Run, Schedule or Exit. You can run this Job immediately, or schedule it for later. If you simply exit now, the Job will still be available later. See the next section for information about setting up a schedule.

3.8 Scheduling a Backup

1. Select your Microsoft SQL Server Agent on the left pane of the Agent Console application. The Microsoft SQL Server Job you created plus the Schedule, Global and Inventory files appear in the right pane.
2. Double-click on the Schedule file. The Schedule List appears.
3. Click the New button. The Scheduling Wizard launches.
4. Work through the Schedule Wizard as described in Section 5.1.1 of the Agent Console manual. "Add a New Schedule Entry."
5. On the Schedule wizard – Weekly panel, select the days you want the Job to run. For example, Monday through Friday.
6. Continue working through the Schedule Wizard until finished.
7. Repeat the preceding steps to run a backup of your Microsoft SQL Server Job once per week (Full backup recommended).

You will see a choice of full or incremental backup in the options screen, but the first time that you back up an SQL database, it will always be a backup type of "full". For subsequently scheduled backups, backup types can be full or incremental.

The difference between "Full" and "Incremental" is that "Full" will back up the entire database, while "Incremental" will only back up the transaction log files. During a restoration, the log files will be played back to achieve the most up to date restore since the last (full) backup.

8. Next, you should decide on how to tailor your backup and recovery options based on your specific Microsoft SQL Server (not discussed in this section).

3.9 Verifying the Backup

To verify a backup, right-click the Safeset and select Properties.

Also, in the logs directory you can view the log file that the backup produced. The sample here is the last part of the log of a successful backup. Notice that there were no "errors encountered", and no "warnings encountered".

```

April12 15:32 BKUP-I-0218 Agent Version 6.72.1071 April 8 2010 10:51:21
April12 15:32 BKUP-I-0031 Job started at 12-APRIL-2010 15:32:19.57 -0400
April12 15:32 BKUP-I-0030 Job name /db1/
April12 15:32 MSQL-I-0001 SQL Server Plug-In Version:6.72.1071
April12 15:32 SSET-I-0034 connect to the Vault at server11
April12 15:32 SSET-I-0001 Connection to the Vault is using AES encryption
April12 15:32 SSET-I-0181 login at 12-APRIL-2010 15:39:19.73 -0400 as c/c
-----/ ----- / -----
April12 15:32 BKUP-I-0219 Vault Version 6.21
April12 15:32 SSET-I-0036 synching catalog number is 1
April12 15:32 BKUP-I-0001 Backup Encoding: ANSI
April12 15:32 MSQL-I-0001 VDI driver is loaded ...
April12 15:32 MSQL-I-0001 SQL Command: SET IMPLICIT_TRANSACTIONS OFF IF
@@TRANCOUNT > 0 ROLLBACK TRAN BACKUP DATABASE [CustNorth01] TO
VIRTUAL_DEVICE='vvdv628' WITH STATS=1
April12 15:32 MSQL-I-0001 SQL server message 3014: Backup or restore
operation successfully processed 36793 pages in 22.675 seconds (13.292
MB/sec).
April12 15:32 BKUP-I-0017 copied
$SQL1$\SQL7ONW2K\DEF\CustNorth01.FULL      5,632      704
April12 15:32 BKUP-I-0017 copied $SQL1$\SQL7ONW2K\DEF\CustNorth01.END
          0          0
April12 15:32 SSET-I-0037 committed catalog number is 1
April12 15:32 SSET-I-0035 disconnect from the Vault at 12-APRIL-2010
15:39:44.12 -0400
April12 15:32 CTLG-I-0001 catalog created
April12 15:32 MSQL-I-0001 object processed:
$SQL1$\SQL7ONW2K\DEF\CustNorth01

```

```

April12 15:32 BKUP-I-0000 errors encountered:                0
April12 15:32 BKUP-I-0000 warnings encountered:             0
April12 15:32 BKUP-I-0000 files/directories examined:       2
April12 15:32 BKUP-I-0000 files/directories filtered:       0
April12 15:32 BKUP-I-0000 files/directories deferred:       0
April12 15:32 BKUP-I-0000 files/directories backed-up:      2
April12 15:32 BKUP-I-0000 files backed-up:                  2
April12 15:32 BKUP-I-0000 directories backed-up:            0
April12 15:32 BKUP-I-0000 data stream bytes processed:      5,632
(5.5 KB)
April12 15:32 BKUP-I-0000 all stream bytes processed:        301,707,776
(287.7 MB)
April12 15:32 BKUP-I-0000 pre-delta bytes processed:         301,707,776
(287.7 MB)
April12 15:32 BKUP-I-0000 deltized bytes processed:          301,707,776
(287.7 MB)
April12 15:32 BKUP-I-0000 compressed bytes processed:        71,548,688
(68.2 MB)
April12 15:32 BKUP-I-0000 approximate bytes deferred:        0 (0
bytes)
April12 15:32 BKUP-I-0000 reconnections on recv fail:        0
April12 15:32 BKUP-I-0000 reconnections on send fail:        0
April12 15:32 BKUP-I-0032 Job completed at 12-APRIL-2010 15:32:44.40 -
0400
April12 15:32 BKUP-I-0033 elapsed time 00:00:25

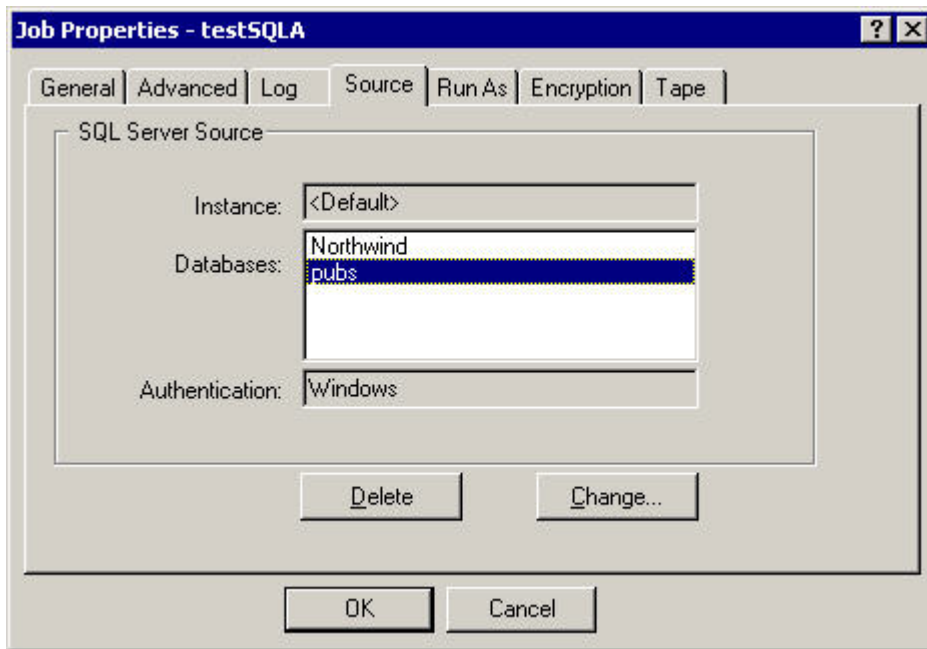
```

You may also have the option set to receive an email notification on a successful (or failed) backup. See section 3.3.5 in the “Agent Console Operations Guide”.

3.10 Viewing Job Properties

This section describes the Job Properties screen for displaying/changing Job properties. Select the properties screen by right-clicking on a Job, or by using F2 with a specific Job selected.

From here, you may view or modify the Job properties.



In the Job Properties, the tabs are:

- General
- Advanced
- Log
- Source
- Run As
- Encryption
- Tape

3.11 Error Messages

When an SQL database is backed up using the incremental (transaction log) method, the logs are sent to the Vault, and are applied against the last full backup when doing a restore. However, if a backup is performed outside of the Plug-In applications (such as Enterprise Manager), this causes the database to produce a transaction log file. This is recorded by the SQL server in the msdb, but the Plug-In does not know about this.

In this case the transaction log files are out of synch for the Plug-In, and a full backup must be performed to ensure the integrity of the data. If incremental backups are continued they will not fail (however, attempts to restore from an incremental SQL Backup may fail with a "REST-F-0014 Job failed to complete" message, and the database may be left in a "loading" state).

3.12 Re-Registering an Agent

If you delete an Agent from a Vault, you are deleting the actual profile on the Agent PC. If you then add that same Agent (it uses the Agent's computer name) to the Vault, the Vault recognizes it and prompts you for a re-registration. This will also happen on a Restore From Another Computer.

The original profile is downloaded from the Vault back to the Agent, but in this case for security reasons, it omits several fields:

- The encryption password, if you used encryption on backups
- The domain, user name and password of the account used to back up the SQL Server ("Run As" tab in Job Properties)
- If "Source" was used with SQL Authentication, you must re-enter that information in the Job Properties > Source tab.

You will receive error-log messages similar to the following when a backup or restore fails because of a re-registration, or a problem with restoring from another computer.

```
PARS-W-0002    Due to a computer registration, configuration file
"Global" is missing the following information:
```

```
PARS-W-0002    Plugin0.Password (Password)
```

```
Please use the Agent Console to re-enter the missing information.
```

In this version of the Agent, the Agent re-registration process creates a "register.log" log file that reports any missing Job file settings. The log file can be viewed via Agent Console once the re-registration is complete.

Any attempt to perform a backup or restore using one of the affected Job files will fail until the Job file has been reconfigured. Should this failure occur, the backup or restore log file will contain information similar to that of the "register.log" (indicating which Job settings are missing).

3.13 Removing a Database from the Backup

A failure in one SQL database can cause the entire SQL backup Job to fail. This may happen if the database is open, or running some process.

To be able to skip the database, you can use either the command line option, or a Registry entry.

For the CLI, use /LOGANDCONTINUE=YES in the vv.exe parameter list.

To set (create) a Registry entry, create a DWORD and set it to 1. For example:

```
HKEY_LOCAL_MACHINE\SOFTWARE\\Agent\LOGANDCONTINUE
```

This works for an SQL Server instance, or a Virtual SQL Server instance in a clustered environment.

4 Working with Restores

For a Microsoft SQL restore, the user will specify a backup (a particular Safeset) from which they would like to restore. Restores can be to the original location, an alternate location, or a file.

Note: Backups and restores must be run with Administrator rights.

If you are restoring to an alternate location, the logical file names of the new database and the transaction files in the database must be the same as the original ones.

Restoring to a file still requires Microsoft SQL Server tools to get the data back into the database. Note that this also requires at least twice the database disk space, in order to keep the entire file on disk, along with the database.

Restores may be necessary in three scenarios:

- Restoring the full database, with any incremental backups, overwriting the existing database.
- With no system backup, restoring the system from the ground up (“bare metal restore”) – installing the OS, applications and then the full database, and any incremental backups, onto a new system.
- If there is an Microsoft SQL Server backup, and a full system backup, install the OS and then restore the System State, and Microsoft SQL Server system.

The Restore Wizard will direct the user through the process. Briefly, the steps to perform this include:

1. From Agent Console, and with a Job highlighted, select the Restore function (Wizard).
2. Select a Source – Vault and Safeset.
3. Select one or more databases to restore.
4. Select the instance where the database should be restored.
5. Confirm account information for Microsoft SQL Server authentication – computer name, instance selection, user name, and password. Choose the Windows or SQL Authentication option.

Note: For SQL authentication, the maximum length of the password is 31 characters.

6. Select the destination where the Microsoft SQL Server database should be restored – original database, other database, or Directory.
7. Enter network account information that will enable the restore process to run – user name, password and domain.
8. Select logging options and level of detail.
9. Start the restore process.

4.1 Notes About Restores

When restoring a single database, the Microsoft SQL Plug-In expects, for Microsoft SQL 2000 and Microsoft SQL 2005 Servers, that the source and destination databases will use identical database and transaction log "file names". Database name and physical file "locations" may be different. The path to the database being restored must exist prior to the restore.

Choosing the REPLACE option allows the data to be restored to a non-original database, having the same logical file name as the original database. You can only do restores to another database, whether on the original server or not, if the REPLACE option is checked.

In such a case, the existing database is deleted. When the REPLACE option is not specified, a safety check occurs (which prevents overwriting a different database by accident). The safety check ensures that the RESTORE operation will not restore the database to the current server if:

- The database named in the RESTORE statement already exists on the current server, or;
- For Microsoft SQL 2000 and Microsoft SQL 2005 Servers, the file names in the destination databases are different from that recorded in the backup set.

With multiple databases restored in the same Job, the destination can only be the original database name, or to a directory (file). If only a single database is restored, you have the choice of original, or directory as well as another (different) database name.

Note: The Plug-In reports errors differently on SQL 2008 versus SQL 2005 when it restores a database from one instance to another.

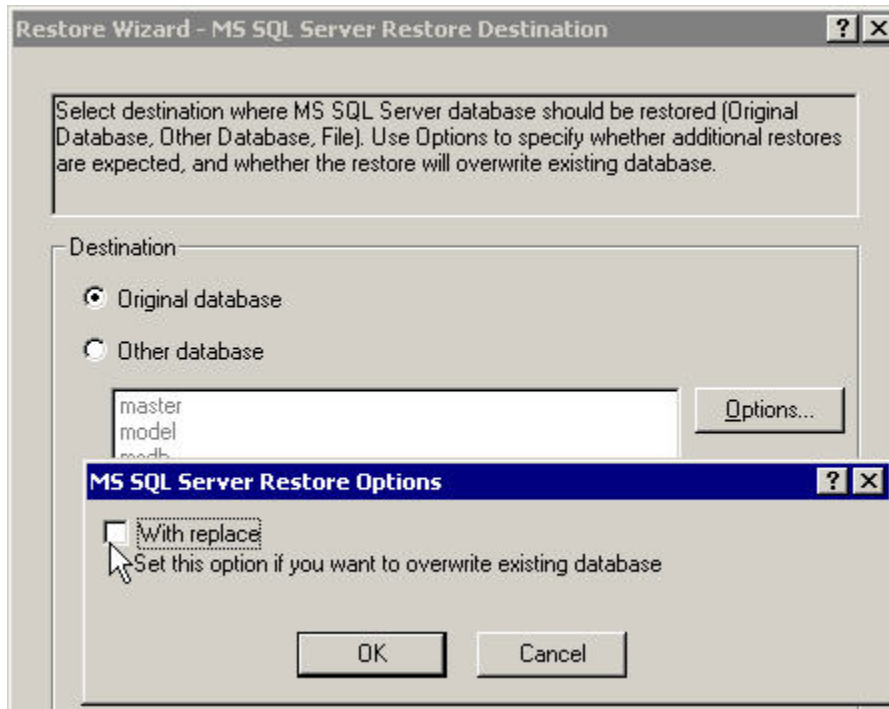
4.2 Deleting a Database

If a database has been taken offline, or removed from the server, it will cause a backup failure. To be able to continue, you can delete the database causing the problem.

In Job Properties -> Source, in the SQL Server Source screen you can select one or more databases and delete them from there, by clicking on the Delete button. You are prompted with "Are you sure?".

4.3 About the With Replace Option

The “With Replace” option allows you to overwrite an existing database. This can be used in two cases:



First, when you restore an existing database (in the Full Recovery model), SQL Server 2005 requires you to back up the tail of the log before you restore the database under the full or bulk-logged recovery model. Trying to restore a database before you back up the tail of the log causes an error, unless the Restore statement contains either a “with replace” or “with stopat” clause.

Secondly, an SQL restore may fail when you try to restore an SQL database from one instance to another instance on the same SQL Server. In this case, make sure the original database (the one that you are restoring) does not exist on the SQL Server, on any of the available instances. Then use the Options -> With Replace selection.

4.4 Restoring Databases

This section provides description of steps to recover Microsoft SQL databases. These methods are applicable to Microsoft SQL 2000, SQL 2005 and SQL 2008 servers.

Microsoft SQL 2000, SQL 2005 and SQL 2008 servers support multiple, named instances of the database.

4.5 Database types

SQL databases fall into two major categories:

- System Databases: - any database created by/for SQL Server for server metadata storage. They are:
 - The MASTER database – As the name suggests this is the main system database and contains information about all other databases as well as SQL Server instance configuration. We assume that this database is being backed up regularly, especially when system changes are made.
 - The MSDB database – This contains information about Backup History, Replication, and Log Shipping. This assumes this database is being protected by regular backups.
 - The MODEL database – This a template database for SQL database creation. If customized database configuration is done this database needs to be protected by periodic backups depending on how regularly customization is changed.
 - The TEMPDB database (Microsoft SQL 2000 only) – This is a temporary working database. No backups are recommended.
- User databases: - any database created for or by users. Microsoft SQL Server 2000/2005 installations come with sample user databases:
 - Northwind
 - Pubs (Microsoft SQL 2000 only)
 - AdventureWorks (Microsoft SQL 2005)

Note: AdventureWorks database for SQL 2008 can be downloaded from Microsoft.

Notes: It is strongly recommended that the system (master, model and msdb) databases be backed up in one Job, and the other (user) databases be backed up in one or more Jobs.

During a recovery process, the master database must be recovered first, by itself. Then the other system databases can be restored.

A restore from another computer, to a database with the original name will fail if the directory structure does not exist on the destination computer. The location and database names are stored as binary data inside the files, and cannot be accessed directly during a restore. The path to the database being restored must exist prior to the restore. You must create the path manually, before the restore.

4.6 Recovering System Databases

- The MASTER database contains key information about SQL Server installation. If it is damaged or removed, the SQL Server instance will not start. The process of recovering the MASTER database is described below.
- The MSDB database can be restored in the same way as user databases.
- The MODEL database can be restored in the same way as user databases.
- The TEMPDB database contains temporary working information, so it should not be backed up or restored.

4.7 Recovering the MASTER Database

To recover the MASTER database:

1. Stop the Microsoft SQL Server service(s).
2. Choose one of these options:
 - (SQL 2000) - Using the REBUILD utility create a blank MASTER database for the SQL Server instance. This utility can be found in the /tools/bin directory of the Microsoft SQL Server. After successful rebuild, user created databases will not be available.

Note: If rebuildm.exe appears to be stuck in an infinite loop, try removing the read-only attributes from the exe file. You may need to copy it from CD to a disk.

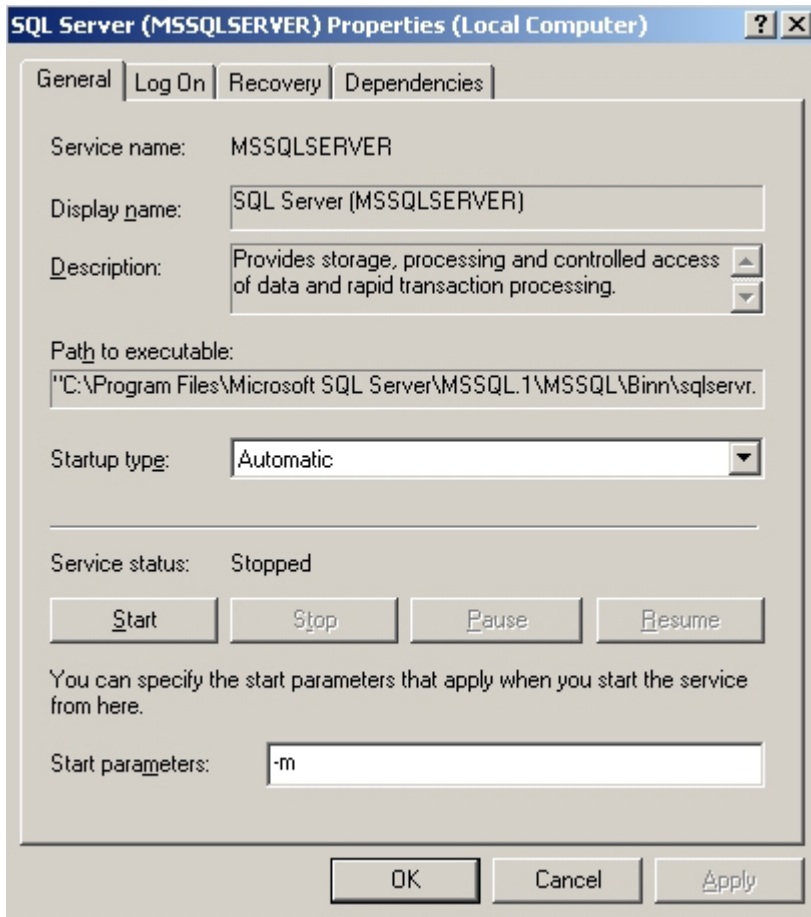
- (SQL 2005/2008) – SQL 2005/2008 does not have the REBUILD utility. The Master database must be rebuilt manually. See “Rebuilding the Master Database” and “How to Install SQL Server 2005” in Microsoft’s documentation.
 - (SQL 2000 MSDE) – SQL 2000 MSDE does not have the REBUILD utility. The user must uninstall and reinstall the database instance. The DB instance name must be the same as it was previously.
3. Re-install the Microsoft SQL Server service pack.
 4. Restart Microsoft SQL Server in single-user mode from Administrative Tools > Services:
 - If the default instance is being used, right-click on the (MS SQL SERVER) > Properties > General:

```
Start parameters: -m
```

Note: You must click the "Start" button on this page now. Clicking the "OK" button only will Not start the service with the -m parameter.
 - If a specific instance is being restored, right-click on the (<instance_name>) > Properties > General:

```
Start parameters: -m
```

Note: You must click the "Start" button on this page now. Clicking the "OK" button only will Not start the service with the -m parameter.



5. Restore the MASTER database using regular restore procedures with the help of Agent/Agent Console.
6. Restart Microsoft SQL Server instance service in normal mode. Verify the data in the MASTER database has been successfully restored. At this point information about user created databases should be visible.
7. Restore other databases as needed.

You may see three files, called Master.full, Master.Log and Master.End that are created by the Plug-In. They are only used by the Plug-In and are not needed by other programs (such as Enterprise Manager). They can be deleted by the user without affecting a restore via EM.

4.8 Recovering the MASTER database on an SQL 2008 Cluster

You may not be able to restore an SQL Server 2008 MASTER database on a Windows Server 2008 Cluster in Disaster Recovery mode. Windows Agent Console may show a "Failed to connect to the server" message, and Web Agent Console may show the Agent as offline. This is because the SQL Server Agent service is in a stopped state.

For example, assume that a two node Server 2008 cluster exists with SQL 2008 (default and named instances) with a separate active directory machine. If the AD, and two nodes are totally destroyed, they would need to be recovered from your backup.

The O/S is reinstalled on the AD machine and the two nodes. Clustering is configured with a default install of SQL 2008. The Agent is installed on both nodes, and reregistered.

When you attempt to restore the SQL instance with Agent Console you will see that the SQL Server Agent service is in a stopped state.

Steps for restoring an SQL Server MASTER database on a Cluster:

1. Make sure the Cluster Service and SQL Services are functioning correctly. If SQL 2008 is being used, make sure that SQL Server 2005 Backward Compatibility is installed.
2. Install the Agent on the owner physical node, and register the Virtual Agent for clustered SQL service. The SQL job is ready to be restored.
- 3.
4. Open the Failover Cluster Management Console, and take the SQL Server resources offline, except for network and disk resources. These two resources need to be online and accessible.
5. Open a CMD.exe window, and go to the program folder of the SQL instance that is going to be restored. (You can find the path from the SQL instance service properties on the service snap-in. For example:
C:\Program Files\Microsoft SQL Server\MSSQL10.MSSQLSERVER\MSSQL\Binn)
6. Run "sqlservr -c -m -s <instance_name>". This command is used to start the SQL instance in single user mode. You can use "sqlservr -c -m" for the default instance.
7. Go to Agent Console, and operating from the virtual Agent, restore the MASTER (only) database. The "sqlservr" command stops when the restore is finished.
8. Check the restore log to make sure the MASTER database was restored successfully.
9. On the Cluster Management Console, bring the SQL instance online.

Note: Not only SQL 2008, but SQL 2005 should follow the same steps to be able to restore a clustered SQL server MASTER database.

4.9 Recovering a Single User database

To recover a single user database:

1. Microsoft SQL Server instance must be running
2. User database metadata information must be available. This can be accomplished in one of two methods:
 - Master database must already contain configuration of the user database to be restored. See the Recovering MASTER database section if the information was lost, or

- User database must be recreated manually using Enterprise Manager. The database must be recreated with the same server name/instance/database name combination. Additionally, the logical file name for the database (Enterprise Manager -> SQL Server instance-> user database properties -> Data Files-> File name) must match that of the original database. Manually recreating the user database is not recommended as it may not provide full database functionality with respect to access privileges, relationships etc.
3. For the database to be available for restore, it must not be placed in single-user mode, and it must not be accessed by any users. User access to the databases can be viewed via the Enterprise Manager ->SQL Server instance-> Management -> Current activity -> Process Info. This information is not refreshed in real time, so it may not be reliable on active systems.
 4. Using the Agent, restore the database from the required safeset as indicated by the Restore Wizard.

In the Restore Destination screen, you will be prompted for:
 - a) the original database name (to be restored on the same database on the same SQL Server),
 - b) another database name (single database restore only - the "Replace" checkbox option must be used),
 - c) a directory on disk for a file (to be restored with Enterprise Manager).
 5. Repeat this process for any other user databases that require restoring.

4.10 Restore Database with Recovery

If you attempt a restore from an incremental backup where an external application (such as Enterprise Manager) was also performing a backup of the same database, the database may hang in a "loading" state.

The log could have "REST-F-0014 job failed to complete."

As a workaround, from SQL Query Analyzer, run `Restore Database <database name> with Recovery` where <database name> is the name of your database.

4.11 Microsoft SQL Bare Metal Restore

To complete a bare-metal restore of a Microsoft SQL server you need these items:

- Replacement hardware, if necessary
- The original operating system disks
- Full drive backups of the system drives, and other logical drives where critical applications or data were installed. A "full backup" consists of the 'System State backup' and the 'full drive or system backup'. A 'System State backup' captures Active Directory, Registry, IIS metabase, and types of data that may not be backed up by some other backup systems.

4.12 Recovering from a Worst Case Disaster

1. Reconfigure hardware that is similar (i.e., equal or better) to the original hardware.
2. Create a logical drive that matches the original configuration. Although hardware does not always need to be identical, be aware that some drivers that are listed in the backup set may be incompatible with hardware on the new systems, and may require you to manually remove or install drivers in Safe mode. Test the system state restoration on replacement hardware before you actually need to perform a system state restoration.
3. Reinstall the operating system. Install the same version of Windows as a stand-alone server to the same drives and paths to which the Windows Server was previously installed.
4. Using Restore, restore full drive backups. The full backup consists of your 'System State backup' and the 'full drive or system backup'. By restoring the System State, you have restored Active Directory, the IIS metabase, etc.

See the Agent for Microsoft Windows User Guide for another example of bare-metal recovery on a Windows system.

Section 4.2 of this Guide describes how to restore databases for Microsoft SQL 2000 and Microsoft SQL 2005.