

Best Practices for Storing the ArcGIS® Workflow Manager Workspace in an Enterprise Geodatabase for SQL Server

An Esri® White Paper
November 2012



Copyright © 2012 Esri
All rights reserved.
Printed in the United States of America.

The information contained in this document is the exclusive property of Esri. This work is protected under United States copyright law and other international copyright treaties and conventions. No part of this work may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording, or by any information storage or retrieval system, except as expressly permitted in writing by Esri. All requests should be sent to Attention: Contracts and Legal Services Manager, Esri, 380 New York Street, Redlands, CA 92373-8100 USA.

The information contained in this document is subject to change without notice.

Esri, the Esri globe logo, ArcGIS, ArcSDE, ArcCatalog, ArcMap, esri.com, and @esri.com are trademarks, service marks, or registered marks of Esri in the United States, the European Community, or certain other jurisdictions. Other companies and products or services mentioned herein may be trademarks, service marks, or registered marks of their respective mark owners.

Best Practices for Storing the ArcGIS Workflow Manager Workspace in an Enterprise Geodatabase for SQL Server

An Esri White Paper

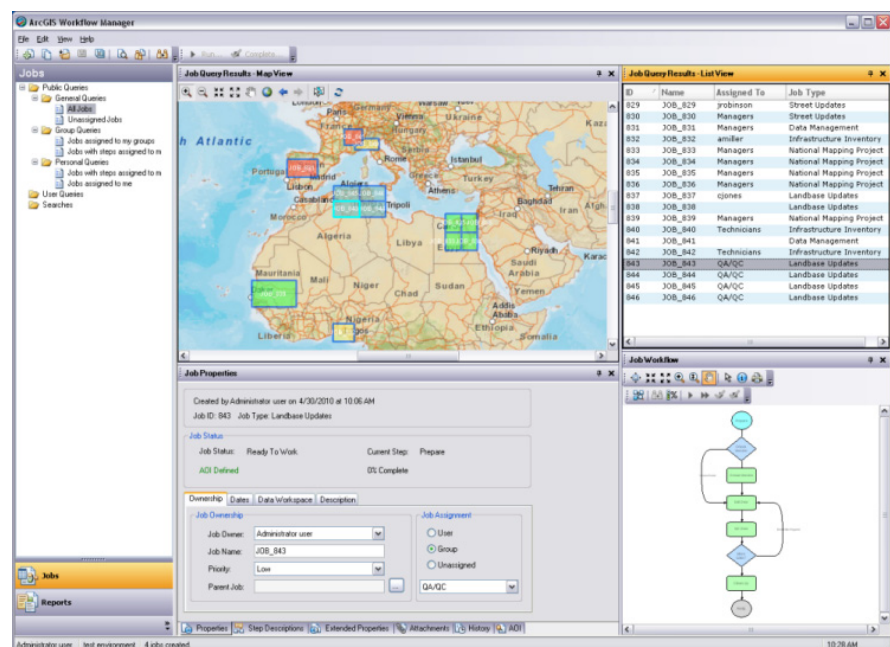
Contents	Page
Introduction.....	1
Postinstallation Utility	2
Preparing the Database	2
ArcSDE DBTUNE.....	2
Disk Configuration.....	3
Reducing Disk I/O Contention.....	3
Transparent Data Encryption	4
Step 1: Create Data Files.....	6
Step 2: Create Workflow Manager User.....	8
Step 3: Modify DBTUNE	9
Step 4: Configure SQL Server Parameters	10
Step 5: Configure ArcSDE Parameters.....	10
Step 6: Create and Configure the Workspace	10
Step 7: Verify Storage.....	15
Step 8: Grant Permissions and Roles	16

Contents	Page
Step 9: Log File Tables	17
Step 10: Create Database Users	18
Editor User	18
Viewer User	18
Replication	18
Conclusion	19

Best Practices for Storing the ArcGIS Workflow Manager Workspace in an Enterprise Geodatabase for SQL Server

Introduction

This white paper is intended to help database administrators establish the ArcGIS® Workflow Manager workspace in an enterprise geodatabase for SQL Server®. The enterprise geodatabase uses ArcSDE® technology as the gateway between geographic information system (GIS) clients and SQL Server. The Workflow Manager workspace uses ArcSDE, and this paper covers the best practices to create the ArcSDE instance. ArcGIS Workflow Manager, an extension to ArcGIS for Desktop and ArcGIS for Server, is an enterprise workflow management application that provides an integration framework for ArcGIS multiuser geodatabase environments. It simplifies many aspects of job management and tracking and streamlines the workflow, resulting in significant time savings for any implementation. Workflow Manager provides tools for allocating resources and tracking the status and progress of jobs.



ArcGIS Workflow Manager Client Application Interface

A detailed history of job actions is automatically recorded for each job to give managers a complete report on how the job was completed. This information can be supplemented

with comments and notes to provide even richer job documentation. Workflow Manager handles complex geodatabase tasks behind the scenes by assisting the user in the creation and management of versions. Integration between Workflow Manager and ArcGIS geodatabase tools provides a way of tracking feature edits made through Workflow Manager jobs using the geodatabase archiving tools.

Postinstallation Utility

For the Workflow Manager applications to function, all Workflow Manager system tables must exist in the database, and of those, the tables that host required configuration information must be populated with values that will drive the application behavior. The following section will detail how to create and start configuring the Workflow Manager system tables.

The Workflow Manager postinstallation utility is used to insert the Workflow Manager system tables into the workspace you've created for the Workflow Manager repository (also referred to as the Workflow Manager database).

This repository contains a set of geodatabase tables (Workflow Manager system tables) used to store the job and configuration information for your work management system and one feature class that is used to store the geometries for the areas of interest (AOIs) for your jobs.

The Workflow Manager system tables added by the postinstallation utility are geodatabase tables registered with your ArcSDE geodatabase. If at any time you are unsure of the state of the database, connect to the repository using ArcCatalog™. Because the Workflow Manager system tables are geodatabase tables, it is strongly recommended that you do not delete or modify these tables using your RDBMS tools.

Preparing the Database

Before you can install the Workflow Manager system tables into an ArcSDE instance, you must properly set up the instance to receive the repository; the steps below will show how to accomplish this task.

ArcSDE DBTUNE

DBTUNE storage parameters let you control how ArcSDE technology creates objects within a SQL Server database. You can determine such things as how to allocate space to a table or index, which FileGroup a table or index is created in, and other SQL Server-specific storage attributes. They also allow you to specify one of the available storage formats for the geometry of a spatial column.

The DBTUNE storage parameters are stored in the DBTUNE table. The DBTUNE table, along with all other metadata tables, is created during the setup phase that follows the installation of ArcSDE. ArcSDE installation creates a dbtune file under the etc directory from which the DBTUNE table is populated. If no dbtune file is present during setup, the DBTUNE table will be populated with default values.

When a large number of database connections access the same files in the same location on disk, database performance is slower because the connections are competing with one another for the same resources. To reduce this competition, you can store database files in different locations on disk.

Thus, DBTUNE can be modified to store the Workflow Manager system tables in separate data files across different locations on disk. This will reduce disk contention and improve database input/output (I/O).

Standard GIS storage recommendations favor keeping index and log files separate from vector and tabular business tables. For performance reasons, it is better to position the business, feature, and spatial index tables separately and position FileGroup data files based on their usage pattern. For a multiversioned, highly active editing geodatabase, database files of the VERSIONS FileGroup may be separated and dispersed across available disks to avoid I/O contention.

Disk Configuration

Large production enterprise geodatabase systems should employ a hardware striping solution. Your best disk and data organization strategies involve spreading your data across multiple disks.

With data spread across multiple disks, more spindles actively search for it. This can increase disk read time and decrease disk contention. However, too many disks can slow down a query. There are two main ways of achieving striping: FileGroups and redundant array of independent disks (RAID). You can also combine the two by creating FileGroups within disk arrays. You can employ data segregation strategies; keeping tables from indexes or certain types of tables from other tables will improve performance and alleviate administrative burdens.

The suggested SQL Server optimal configuration is as follows:

- Disk 0—SQL Server/Application software
- Disk 1—master, model, msdb
- Disk 2—tempdb
- Disk 3—Log files
- Disk 4—Feature data tables
- Disk 5—Spatial index data tables
- Disk 6—Attribute data/Business tables
- Disk 7—SQL Server indexes

Reducing Disk I/O Contention

As a rule, you should create database files as large as possible, based on the maximum amount of data you estimate the database will contain, to accommodate future growth. By creating large files, you can avoid file fragmentation and gain better database performance. In many cases, you can let data files grow automatically; just be sure to limit autogrowth by specifying a maximum growth size that leaves some hard disk space available. By putting different FileGroups on different disks, you can also minimize the physical fragmentation of your files as they grow.

To configure data and log files for best performance, follow these best practices:

- To avoid disk contention, do not put data files on the drive that contains the operating system files.
- Put transaction log files and data files on separate drives. This will give you the best performance by reducing disk contention between data and transaction log files.

- Put the tempdb database on a separate drive if possible, preferably on a RAID 10 or RAID 5 system. For environments where there is intensive use of tempdb databases, you can get better performance by putting tempdb on a separate drive, which lets SQL Server perform tempdb operations in parallel with database operations.
- The RAID configuration that is best for your database files depends on several factors, including performance and recoverability needs. RAID 10 is the recommended RAID system for transaction log, data, and index files. If you have budget restrictions, you can consider keeping the transaction log files in a RAID 10 system and storing the data and index files in a RAID 5 system.

For more information about RAID, see RAID Levels and SQL Server at [http://technet.microsoft.com/en-us/library/ms190764\(SQL.105\).aspx](http://technet.microsoft.com/en-us/library/ms190764(SQL.105).aspx) and chapter 7, "Planning Fault Tolerance and Avoidance," by Charlie Russel and Sharon Crawford, from *Microsoft® Windows® 2000 Server Administrator's Companion* (Microsoft Press) at [http://technet.microsoft.com/pt-br/library/bb742464\(en-us\).aspx](http://technet.microsoft.com/pt-br/library/bb742464(en-us).aspx).

- Use partitioning on large tables. Partitioning lets you split a table across multiple FileGroups; by using partitions, you can place a subset of a table or index on a designated FileGroup. This capability lets you separate specific pieces of a table or index onto individual FileGroups and effectively manage file I/O for volatile tables. Partitions let you easily manage archival routines and data loading operations.

Below is a suggested design to reduce disk I/O contention:

File Type	Database Activity	Move File to Disk With
Transaction log files	Frequent edits	Relatively low I/O
Transaction log files	Few or no edits	Moderate I/O
tempdb	Frequent edits	Low I/O but separate from transaction log files
master, model, msdb	Few edits	Moderate I/O
Data	Frequent edits	Relatively low I/O

Transparent Data Encryption

The precautions you can take to help secure the database include designing a secure system, encrypting confidential assets, and building a firewall around the database servers. However, if the physical media (drives or backup tapes) are stolen, a malicious party can just restore or attach the database and browse the data. One solution is to encrypt the sensitive data in the database and protect the keys that are used to encrypt the data with a certificate. This prevents anyone without the keys from using the data, but this kind of protection must be planned in advance.

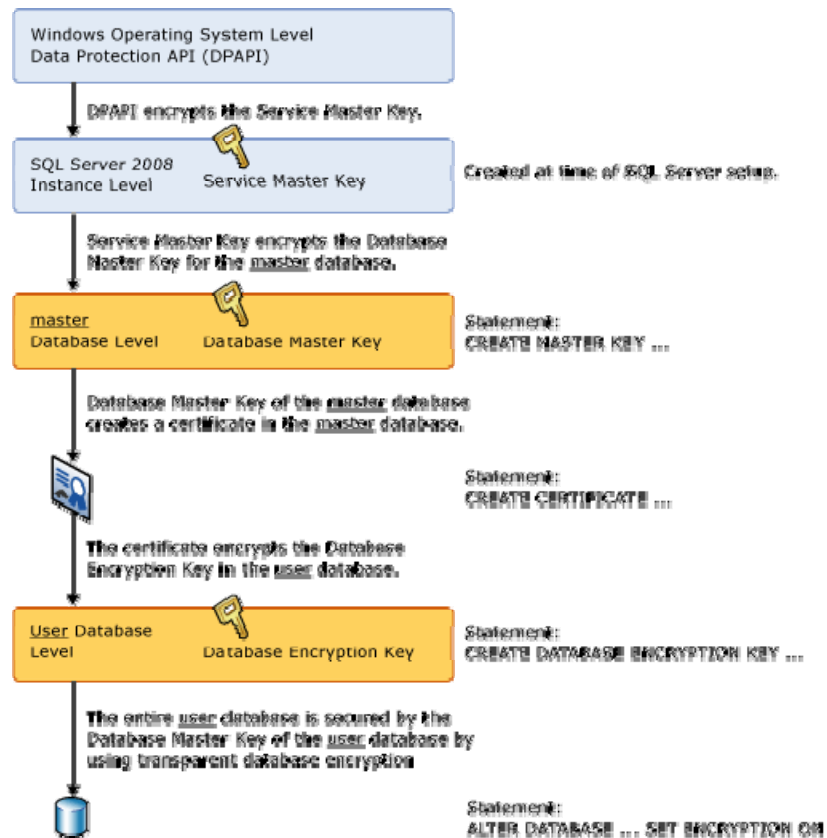
Transparent data encryption (TDE) performs real-time I/O encryption and decryption of the data and log files. The encryption uses a database encryption key (DEK), which is stored in the database boot record for availability during recovery. The DEK is either a symmetric key secured by using a certificate stored in the master database of the server or an asymmetric key protected by an extensible key management (EKM) module. TDE protects data "at rest," meaning the data and log files. It provides the ability to comply with many laws, regulations, and guidelines established in various industries. This enables software developers to encrypt data by using Advanced Encryption Standard

J10043

(AES) and Triple Data Encryption Standard (3DES) encryption algorithms without changing existing applications.

Database files are encrypted at the page level. The pages in an encrypted database are encrypted before they are written to disk and decrypted when read into memory. TDE does not increase the size of the encrypted database.

The following illustration shows the architecture of TDE encryption:



TDE Encryption Architecture

Learn more about TDE at <http://msdn.microsoft.com/en-us/library/bb934049.aspx>.

To use TDE, follow these steps:

- Create a master key.
- Create or obtain a certificate protected by the master key.
- Create a database encryption key and protect it with the certificate.
- Set the database to use encryption.

```
USE master
GO
/* Verify master key */
SELECT * FROM sys.symmetric_keys WHERE name LIKE '%MS_DatabaseMasterKey%'
GO

/* if there are no records found, then it means there was no predefined Master Key.
To create a Master Key, you can execute the below mentioned TSQL code. */

/* Create master key */
CREATE MASTER KEY ENCRYPTION BY PASSWORD = 'wmxadmin';
GO
/* Backup master key */
OPEN MASTER KEY DECRYPTION BY PASSWORD = 'wmxadmin';
GO
BACKUP MASTER KEY TO FILE = 'D:\mssqlbackup\master\masterkey.mk'
    ENCRYPTION BY PASSWORD = 'wmxadmin';
GO

/* Create Certificate */
CREATE CERTIFICATE wmx_cert WITH SUBJECT = 'WMX Server Certificate';
GO

/* Verify Certificate */
SELECT * FROM sys.certificates where [name] = 'wmx_cert'
GO

/* Backup certificate */
BACKUP CERTIFICATE wmx_cert TO FILE = 'D:\mssqlbackup\master\wmx.cer'
    WITH PRIVATE KEY (
        FILE = 'D:\mssqlbackup\master\wmx.pvk',
        ENCRYPTION BY PASSWORD = 'wmxadmin');
GO

USE wmxdb
GO
/* Create Encryption key */
CREATE DATABASE ENCRYPTION KEY
    WITH ALGORITHM = AES 256
    ENCRYPTION BY SERVER CERTIFICATE wmx_cert;
GO

/* Encrypt database */
ALTER DATABASE wmxdb SET ENCRYPTION ON;
GO

/* Verify Encryption */
SELECT
    DB_NAME(database_id) AS DatabaseName
    ,Encryption_State AS EncryptionState
    ,key_algorithm AS Algorithm
    ,key_length AS KeyLength
FROM sys.dm database encryption keys
GO
SELECT
    NAME AS DatabaseName
    ,IS_ENCRYPTED AS IsEncrypted
FROM sys.databases where name = 'wmxdb'
GO
```

Step 1: Create Data Files

Create new FileGroups to store the Workflow Manager system tables.

FILEGROUP	ArcSDE_PARAMETER
WMX_BDATA	Business table
WMX_BINDEX	Business table index
WMX_FDATA	Feature table
WMX_FINDEX	Feature table index
WMX_SDATA	Spatial Index table
WMX_SINDEX	Spatial Index table index
WMX_ADATA	Adds table (versioned)
WMX_AINDEX	Adds table index
WMX_DDATA	Deletes table (versioned)
WMX_DINDEX	Deletes table index

J10043

```
USE MASTER
GO
ALTER DATABASE [WMXDB] ADD FILEGROUP [WMX_BDATA]
GO
ALTER DATABASE [WMXDB] ADD FILE(NAME = N'wmx Bdata01', FILENAME =
N'C:\mssql\data\wmxdb\wmx_Bdata01.NDF' , SIZE = 1, MAXSIZE = 400, FILEGROWTH = 1) TO FILEGROUP
[WMX_BDATA]
GO
ALTER DATABASE [WMXDB] ADD FILEGROUP [WMX_BINDEX]
GO
ALTER DATABASE [WMXDB] ADD FILE(NAME = N'wmx Bindex01', FILENAME =
N'C:\mssql\data\wmxdb\wmx_Bindex01.NDF' , SIZE = 1, MAXSIZE = 400, FILEGROWTH = 1) TO FILEGROUP
[WMX_BINDEX]
GO
ALTER DATABASE [WMXDB] ADD FILEGROUP [WMX_FDATA]
GO
ALTER DATABASE [WMXDB] ADD FILE(NAME = N'wmx Fdata01', FILENAME =
N'C:\mssql\data\wmxdb\wmx_Fdata01.NDF' , SIZE = 1, MAXSIZE = 400, FILEGROWTH = 1) TO FILEGROUP
[WMX_FDATA]
GO
ALTER DATABASE [WMXDB] ADD FILEGROUP [WMX_FINDEX]
GO
ALTER DATABASE [WMXDB] ADD FILE(NAME = N'wmx Findex01', FILENAME =
N'C:\mssql\data\wmxdb\wmx_Findex01.NDF' , SIZE = 1, MAXSIZE = 400, FILEGROWTH = 1) TO FILEGROUP
[WMX_FINDEX]
GO
ALTER DATABASE [WMXDB] ADD FILEGROUP [WMX_SDATA]
GO
ALTER DATABASE [WMXDB] ADD FILE(NAME = N'wmx Sdata01', FILENAME =
N'C:\mssql\data\wmxdb\wmx_Sdata01.NDF' , SIZE = 1, MAXSIZE = 400, FILEGROWTH = 1) TO FILEGROUP
[WMX_SDATA]
GO
ALTER DATABASE [WMXDB] ADD FILEGROUP [WMX_SINDEX]
GO
ALTER DATABASE [WMXDB] ADD FILE(NAME = N'wmx Sindex01', FILENAME =
N'C:\mssql\data\wmxdb\wmx_Sindex01.NDF' , SIZE = 1, MAXSIZE = 400, FILEGROWTH = 1) TO FILEGROUP
[WMX_SINDEX]
GO
ALTER DATABASE [WMXDB] ADD FILEGROUP [WMX_ADATA]
GO
ALTER DATABASE [WMXDB] ADD FILE(NAME = N'wmx Adata01', FILENAME =
N'C:\mssql\data\wmxdb\wmx_Adata01.NDF' , SIZE = 1, MAXSIZE = 400, FILEGROWTH = 1) TO FILEGROUP
[WMX_ADATA]
GO
ALTER DATABASE [WMXDB] ADD FILEGROUP [WMX_AINDEX]
GO
ALTER DATABASE [WMXDB] ADD FILE(NAME = N'wmx Aindex01', FILENAME =
N'C:\mssql\data\wmxdb\wmx_Aindex01.NDF' , SIZE = 1, MAXSIZE = 400, FILEGROWTH = 1) TO FILEGROUP
[WMX_AINDEX]
GO
ALTER DATABASE [WMXDB] ADD FILEGROUP [WMX_DDATA]
GO
ALTER DATABASE [WMXDB] ADD FILE(NAME = N'wmx Ddata01', FILENAME =
N'C:\mssql\data\wmxdb\wmx_Ddata01.NDF' , SIZE = 1, MAXSIZE = 400, FILEGROWTH = 1) TO FILEGROUP
[WMX_DDATA]
GO
ALTER DATABASE [WMXDB] ADD FILEGROUP [WMX_DINDEX]
GO
ALTER DATABASE [WMXDB] ADD FILE(NAME = N'wmx Dindex01', FILENAME =
N'C:\mssql\data\wmxdb\wmx_Dindex01.NDF' , SIZE = 1, MAXSIZE = 400, FILEGROWTH = 1) TO FILEGROUP
[WMX_DINDEX]
GO
```

By setting the data files' initial size to 1 MB, there is no delay in the creation of the FileGroups; to avoid fragmentation, you can resize the data files.

```
ALTER DATABASE [WMXDB] MODIFY FILE ( NAME = N'wmx Bdata01', SIZE = 400MB )
ALTER DATABASE [WMXDB] MODIFY FILE ( NAME = N'wmx Bindex01', SIZE = 400MB )
ALTER DATABASE [WMXDB] MODIFY FILE ( NAME = N'wmx Fdata01', SIZE = 400MB )
ALTER DATABASE [WMXDB] MODIFY FILE ( NAME = N'wmx Findex01', SIZE = 400MB )
ALTER DATABASE [WMXDB] MODIFY FILE ( NAME = N'wmx Sdata01', SIZE = 400MB )
ALTER DATABASE [WMXDB] MODIFY FILE ( NAME = N'wmx Sindex01', SIZE = 400MB )
ALTER DATABASE [WMXDB] MODIFY FILE ( NAME = N'wmx Adata01', SIZE = 400MB )
ALTER DATABASE [WMXDB] MODIFY FILE ( NAME = N'wmx Aindex01', SIZE = 400MB )
ALTER DATABASE [WMXDB] MODIFY FILE ( NAME = N'wmx Ddata01', SIZE = 400MB )
ALTER DATABASE [WMXDB] MODIFY FILE ( NAME = N'wmx Dindex01', SIZE = 400MB )
```

Verify FileGroups and data files:

```
EXEC sp_helpdb wmxdb
GO
```

Step 2: Create Workflow Manager User

Create a new database user to store the Workflow Manager system tables; grant the appropriate permissions.

Create user and schema:

```
USE WMXDB
GO
CREATE USER [wmx] FOR LOGIN [wmx]
GO
CREATE SCHEMA [wmx] AUTHORIZATION [wmx]
GO
ALTER USER [wmx] WITH DEFAULT SCHEMA=[wmx]
GO
```

Grant privileges:

```
USE WMXDB
GO
EXEC sp_droprolemember 'gis_data_creator', 'wmx'
GO
EXEC sp_droprole 'gis data creator'
GO
CREATE ROLE gis data creator AUTHORIZATION dbo
GO
GRANT CREATE TABLE TO gis data creator
GO
GRANT CREATE PROCEDURE TO gis_data_creator
GO
GRANT CREATE VIEW TO gis data creator
GO
EXEC sp_addrolemember 'gis data creator', 'wmx'
GO
```

Verify role:

```
EXEC sp_helprolemember 'gis_data_creator'
GO
```

Verify role permissions:

```
select dp.NAME AS principal name,
dp.type desc AS principal type desc,
o.NAME AS object_name,
p.permission_name,
p.state desc AS permission state desc
from sys.database_permissions p
left OUTER JOIN sys.all_objects o
on p.major_id = o.OBJECT_ID
inner JOIN sys.database_principals dp
on p.grantee_principal_id = dp.principal_id
where dp.NAME = 'gis_data_creator'
GO
```

J10043

Verify user permissions:

```
select USER_NAME(p.grantee_principal_id) AS principal_name,
       dp.type_desc AS principal_type_desc,
       p.class_desc,
       OBJECT_NAME(p.major_id) AS object_name,
       p.permission_name,
       p.state_desc AS permission_state_desc
from sys.database_permissions p
inner JOIN sys.database_principals dp
on p.grantee_principal_id = dp.principal_id
where USER_NAME(p.grantee_principal_id) = 'wmx'
```

Associate login wmx with user wmx:

```
USE WMXDB
GO
EXEC sp_change_users_login 'update_one', 'wmx', 'wmx'
GO
EXEC sp_helpuser 'wmx'
GO
```

Step 3: Modify DBTUNE

Export the dbtune file before making any modification.

```
sdedbtune -o export -f dbtune_exp.sde -u sde -p sde
-i sde:sqlserver:mcsdbsrv -D wmxdb
```

Copy *dbtune_exp.sde* to *dbtune_wmx.sde*.

Create a new DBTUNE keyword following the steps below:

- Copy the ##DEFAULTS keyword and paste it at the end of the *dbtune_wmx.sde* file.
- Rename it ##WMX and change the FileGroup name for the appropriate parameters.

```
dbtune_wmx.sde

##WMX
A_INDEX_RASTER "WITH FILLFACTOR = 75 ON WMX_AINDEX"
A_INDEX_ROWID "WITH FILLFACTOR = 75 ON WMX_AINDEX"
A_INDEX_SHAPE "WITH FILLFACTOR = 75 ON WMX_AINDEX"
A_INDEX_STATEID "WITH FILLFACTOR = 75 ON WMX_AINDEX"
A_INDEX_USER "WITH FILLFACTOR = 75 ON WMX_AINDEX"
A_INDEX_XML "WITH FILLFACTOR = 75 ON WMX_AINDEX"
A_STORAGE "ON WMX_ADATA"
B_INDEX_RASTER "WITH FILLFACTOR = 75 ON WMX_BINDEX"
B_INDEX_ROWID "WITH FILLFACTOR = 75 ON WMX_BINDEX"
B_INDEX_SHAPE "WITH FILLFACTOR = 75 ON WMX_BINDEX"
B_INDEX_TO_DATE "WITH FILLFACTOR = 75 ON WMX_BINDEX"
B_INDEX_USER "WITH FILLFACTOR = 75 ON WMX_BINDEX"
B_INDEX_XML "WITH FILLFACTOR = 75 ON WMX_BINDEX"
B_STORAGE "ON WMX_BDATA"
D_INDEX_ALL "WITH FILLFACTOR = 75 ON WMX_DINDEX"
D_INDEX_DELETED AT "WITH FILLFACTOR = 75 ON WMX_DINDEX"
D_STORAGE "ON WMX_DDATA"
F_INDEX_AREA "WITH FILLFACTOR = 75 ON WMX_FINDEX"
F_INDEX_FID "WITH FILLFACTOR = 75 ON WMX_FINDEX"
F_INDEX_LEN "WITH FILLFACTOR = 75 ON WMX_FINDEX"
F_STORAGE "ON WMX_FDATA"
GEOMETRY_STORAGE "GEOMETRY"
GEOMTAB_PK "WITH FILLFACTOR = 75 ON WMX_FINDEX"
GEOMTAB_STORAGE "ON WMX_FDATA"
I_STORAGE "ON WMX_FDATA"
S_INDEX_ALL "WITH FILLFACTOR = 75 ON WMX_SINDEX"
S_INDEX_SP_FID "WITH FILLFACTOR = 75 ON WMX_SINDEX"
S_STORAGE "ON WMX_SDATA"
END
```

If your database only stores the WMX repository, you can edit ##DEFAULTS; otherwise, create a new configuration keyword as described above.

- Import the modified *dbtune_wmx.sde* file.

```
sdedbtune -o import -f dbtune_wmx.sde -u sde -p sde  
-i sde:sqlserver:mcsdbsrv -D wmxdb
```

Step 4: Configure SQL Server Parameters

It is recommended that you use the following parameter values when creating a SQL Server database:

SQL Server Parameters for Workflow Manager

Parameter name	Value
Server Memory: Use AWE to allocate memory	Enabled
Processors: Boost SQL Server Priority	Enabled
Security SQL Server and Windows Authentication mode	Enabled
Connections: Maximum number of concurrent connections	0 = Unlimited
Connections: Allow remote connections to this server	Enabled

Step 5: Configure ArcSDE Parameters

You need to configure the MAXBLOBSIZE and TCPKEEPALIVE parameters for the ArcSDE geodatabase used as the WMX repository. The MAXBLOBSIZE value is -1 by default. However, if you are using SQL Server or another enterprise DBMS, make sure that this value is set to -1 and the TCPKEEPALIVE value is set to 1. This command should be used from the command prompt of a machine where ArcSDE is installed.

```
sdeconfig -o alter -v MAXBLOBSIZE=-1 -i <service> -u sde -p <sde_password>  
sdeconfig -o alter -v TCPKEEPALIVE=1 -i <service> -u sde -p <sde_password>
```

For more information, see the ArcSDE Administration Command Reference.

Step 6: Create and Configure the Workspace

The postinstallation utility creates all the system tables required for you to use ArcGIS Workflow Manager.

Steps

- From the **Start** menu, navigate to **ArcGIS Workflow Manager** and select **Workflow Manager Post Installation**.

The **ArcGIS Workflow Manager Post-Installation** wizard appears.

J10043

- b. Enter the connection information for the designated workspace, then test the connection information.

The screenshot shows the 'ArcGIS Workflow Manager Post-Installation' dialog box. It has several input fields: 'Server' with 'production', 'Service' with '5151', and 'Database' which is empty with a note '(If supported by your DBMS)'. Under the 'Account' section, 'Database Authentication' is selected, with 'User Name' as 'wmx' and 'Password' as '****'. 'Operating System Authentication' is unselected. A 'Test Connection' button is to the right. Under the 'Version' section, 'SDE.DEFAULT' is selected, with a 'Change...' button. At the bottom are 'Help', '< Back', 'Next >', and 'Cancel' buttons.

Enter connection information for the ArcSDE geodatabase that will be the container for the Workflow Manager repository.

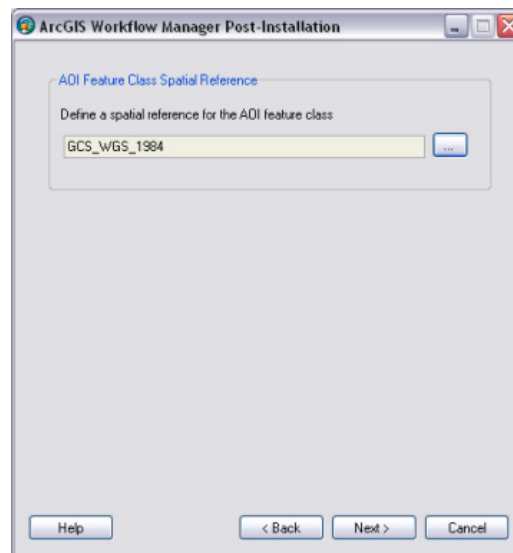
- c. Click **Next**.
- d. Choose a configuration keyword for your Workflow Manager geodatabase; **use the WMX keyword created in step 3.**

The screenshot shows the 'ArcGIS Workflow Manager Post-Installation' dialog box at the 'Configuration Keyword' step. Two radio buttons are present: 'Default' (unselected) and 'Use Configuration Keyword' (selected). Below the selected option is a text box with 'WMX' and a dropdown arrow. At the bottom are 'Help', '< Back', 'Next >', and 'Cancel' buttons.

Choose the configuration keyword.

- e. Click **Next**.
- f. Define the spatial reference for the Area of Interest feature class.

A separate two-page wizard appears allowing you to select the coordinate system and extents.



Define the spatial reference for the Area of Interest feature class.

- g. Click **Next**.

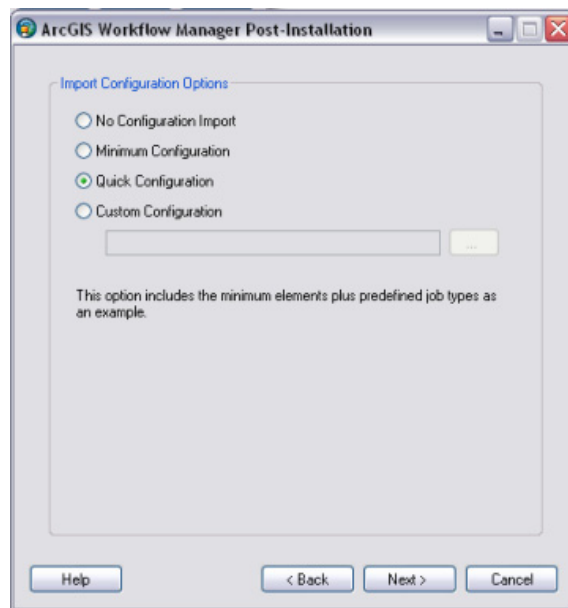
Choose your starting configuration option for the Workflow Manager repository.

- **No configuration** will put the Workflow Manager system tables schema into your designated geodatabase and will not prepopulate any of the configuration elements.
- **Minimum configuration** imports just the basic elements that the Workflow Manager system requires.

Note: If you plan to use all the functionality available in the Workflow Manager client application, it is recommended that you import at least the minimum configuration.

- **Quick configuration** includes the minimum elements plus predefined job types as an example.
- **Custom configuration** allows you to import a preexisting Workflow Manager configuration file.

J10043



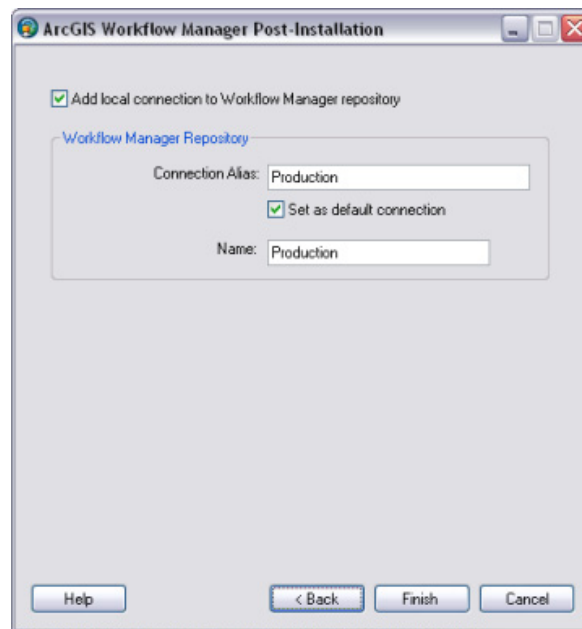
Choose the starting configuration for your repository.

Note: If you plan to import a Workflow Manager configuration exported from a 3.x database, choose the **No Configuration Import** option. Tools are provided in the Workflow Manager Configuration Manager to import the legacy configuration files.

- h. Click **Next**.
- i. Choose whether you would like to create a Workflow Manager database connection to this new repository when the postinstallation is complete.
 - Specify a connection alias.
 - Specify a repository name.

Note: This is used as a unique identifier when defining a cluster for Repository Replication.

- Specify to set the database as default.



Decide if you want Workflow Manager to add a Workflow Manager database connection for you.

- j. Click **Finish**.

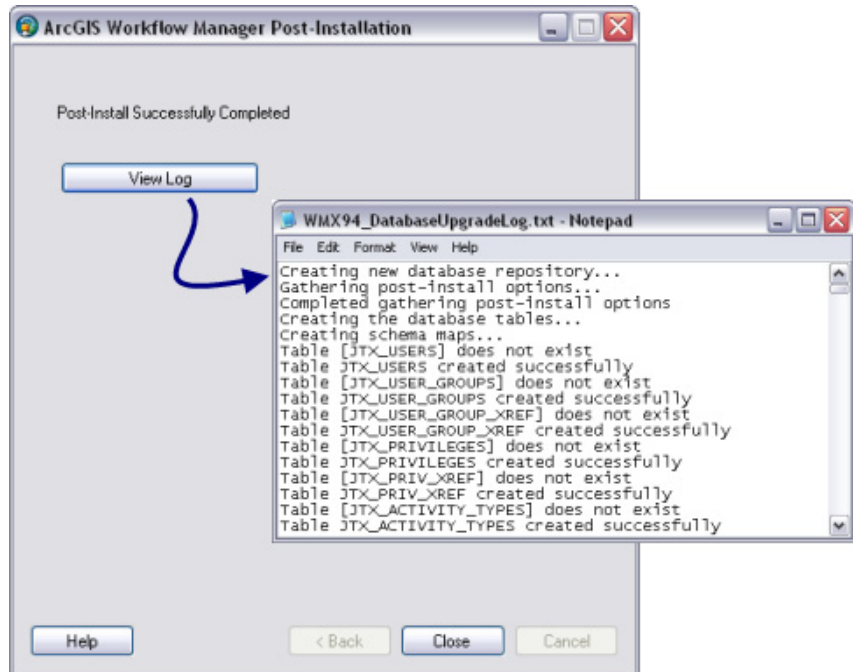
The system tables needed to use Workflow Manager are created in the database. Also, the current user running the postinstallation is automatically added to the list of users and has administrator access.

The final page in the wizard appears, indicating whether the postinstallation was successful.

Note: Administrator access means that the current user can connect to the selected database through the configuration manager to make edits to components of your system.

J10043

- k. Choose whether you would like to view the log.



The final page of the wizard indicates the success of the postinstallation and provides an option to see a log of the process.

- l. Click Close.

Step 7: Verify Storage

Run the SQL queries below to verify that the Workflow Manager workspace was created under the correct FileGroups.

```
USE WMXDB
GO
```

List FileGroups and data files:

```
EXEC sp_helpdb wmxdb
GO
```

List FileGroup data files:

```
EXEC sp_helpfilegroup 'PRIMARY'
GO
```

List tables by FileGroup:

```
SELECT USER_NAME(o.uid) [Owner],  
       OBJECT_NAME(i.id) [Table Name],  
       FILEGROUP_NAME(groupid) AS [Filegroup Name]  
FROM sysindexes i inner join sysobjects o  
  ON i.id = o.id  
WHERE i.indid IN (0, 1) AND OBJECTPROPERTY(i.id, 'IsMSShipped') = 0 AND  
      USER_NAME(o.uid) = 'wmx'  
ORDER BY 1,3,2  
GO
```

List indexes by table and FileGroup:

```
select 'owner'=user name(o.uid)  
, 'table name'=object name(i.id), i.indid  
, 'index name'=i.name , i.groupid  
, 'filegroup'=f.name , 'file_name'=d.physical_name  
, 'dataspace'=s.name from sys.sysindexes i  
, sys.sysobjects o, sys.filegroups f  
, sys.database files d, sys.data spaces s  
where objectproperty(i.id, 'IsUserTable') = 1  
and i.id = o.id  
and f.data_space_id = i.groupid  
and f.data_space_id = d.data_space_id  
and f.data space id = s.data space id  
and user name(o.uid) = 'wmx'  
order by object name(i.id), i.name, f.name  
go
```

If any tables or indexes are stored in the wrong FileGroup, use ALTER TABLE and ALTER INDEX to change the FileGroup (see SQL Server Books Online at <http://msdn.microsoft.com/en-us/library/ms130214.aspx>).

Also, in Management Studio, you can re-create the DDL script of tables and indexes. Then within *create script*, you can modify the FileGroup parameter and re-create the tables and indexes in the correct FileGroups. This is particularly useful when tables are empty and you are allowed to re-create database objects.

Step 8: Grant Permissions and Roles

Grant permissions to the Workflow Manager system tables through the schema.

```
USE [wmxdb]  
GO  
EXEC sp_droprole 'wmx_editor', 'giseditor'  
GO  
EXEC sp_droprole 'wmx_editor'  
GO  
EXEC sp_addrole 'wmx_editor', 'wmx'  
GO  
GRANT DELETE ON SCHEMA::[wmx] TO [wmx_editor]  
GRANT EXECUTE ON SCHEMA::[wmx] TO [wmx_editor]  
GRANT INSERT ON SCHEMA::[wmx] TO [wmx_editor]  
GRANT SELECT ON SCHEMA::[wmx] TO [wmx_editor]  
GRANT UPDATE ON SCHEMA::[wmx] TO [wmx_editor]  
GO  
EXEC sp_droprole 'wmx_viewer', 'gisviewer'  
GO  
EXEC sp_droprole 'wmx_viewer'  
GO  
EXEC sp_addrole 'wmx_viewer', 'wmx'  
GO  
GRANT SELECT ON SCHEMA::[wmx] TO [wmx_viewer]  
GO
```

Verify role:

```
EXEC sp_helprolemember wmx_editor'
GO
EXEC sp_helprolemember wmx_viewer'
GO
```

Verify role permissions:

```
select dp.NAME AS principal name,
       dp.type desc AS principal type desc,
       o.NAME AS object name,
       p.permission_name,
       p.state_desc AS permission_state_desc
from sys.database_permissions p
left OUTER JOIN sys.all_objects o
on p.major_id = o.OBJECT_ID
inner JOIN sys.database_principals dp
on p.grantee_principal_id = dp.principal_id
where dp.NAME in ('wmx_editor','wmx_viewer')
GO
```

Step 9: Log File Tables

Enterprise geodatabases use log file tables to maintain lists of selected records. Records are written to log file tables for later use by the application whenever a selection of a specific size is made, a reconciliation or post on a versioned database is performed, or a disconnected editing checkout is done in a client application. The log file tables store the ObjectIDs of the selected features so they can be redisplayed. This allows faster analysis and processing of information.

In ArcGIS, by default, log file tables are used if the selection set contains 100 or more records. This selection threshold of 100 features is set in the registry. Although it can be changed, Esri does not recommend doing so. There is no proven performance reason for changing it, and doing so could even cause performance problems. Thus, log file tables store feature selections in ArcMap™ that have more than 100 records for each connected ArcSDE editor/viewer user.

Log file options are set using specific parameters in the SERVER_CONFIG and DBTUNE tables (sde_server_config and sde_dbtune in a SQL Server database). Parameters in these tables are altered using the sdeconfig and sdedbtune commands, respectively.

In SQL Server, one table is created in tempdb in the format ##SDE_SESSION<SDE_ID>. This table is truncated when the connecting application deletes its log files, and the table is dropped when the session disconnects. When using the default setting, users do not require CREATE TABLE permission in the database for the session table to be created in tempdb.

The DBTUNE SESSION_TEMP_TABLE parameter must be set to 1 (true) to allow the session-based log file table to be created in tempdb. If you change the SESSION_TEMP_TABLE parameter to 0 (false), the SDE_LOGFILES, SDE_LOGFILE_DATA, and SDE_SESSION<SDE_ID> tables will be created in the connecting user's schema; hence, the user requires CREATE TABLE permission.

Learn more about ArcSDE log file tables at
resources.arcgis.com/en/help/main/10.1/index.html#/What_are_ArcSDE_log_file_tables/002q00000011000000/.

Step 10: Create Database Users

The example below shows how to create an editor and viewer ArcSDE user:

Editor User

```
USE master
GO
EXEC sp_addlogin N'giseditor', 'gis$editor', @logindb, @loginlang
GO
```

Create user:

```
USE [wmxdb]
GO
CREATE USER [giseditor] FOR LOGIN [giseditor]
GO
```

Grant privileges:

```
USE [wmxdb]
GO
EXEC sp_addrolemember N'wmx editor', N'giseditor'
GO
```

Viewer User

```
USE master
GO
EXEC sp_addlogin N'gisviewer', 'gis$viewer', @logindb, @loginlang
GO
```

Create user:

```
USE [wmxdb]
GO
CREATE USER [gisviewer] FOR LOGIN [gisviewer]
GO
```

Grant privileges:

```
USE [wmxdb]
GO
EXEC sp_addrolemember N'wmx viewer', N'gisviewer'
GO
```

Replication

ArcGIS Workflow Manager repository replication is designed to allow organizations to maintain multiple Workflow Manager repositories locally in different geographic regions for performance reasons, while the users at each location continue to work on the same set of jobs.

You will define a collection of Workflow Manager repositories that you want to have identical Workflow Manager contents. This collection of repositories is known as a repository cluster.

One Workflow Manager Repository is designated as the parent repository and is tasked with coordinating the synchronization between each repository. After the parent repository has been specified, each child repository is added. The child repositories can be designated as either connected repositories, where the synchronization will happen automatically, or offline repositories, where the synchronization occurs by creating files that can be manually imported.

Learn more about replicating Workflow Manager repositories at
[resources.arcgis.com/en/help/main/10.1/index.html#
/Replicating_Workflow_Manager_Repositories/003800000077000000/](https://resources.arcgis.com/en/help/main/10.1/index.html#/Replicating_Workflow_Manager_Repositories/003800000077000000/).

Conclusion

You can reduce disk contention and improve database I/O by storing the ArcGIS Workflow Manager workspace in different locations on disk. However, this practice alone does not guarantee optimal database performance, and additional tuning tasks may be needed.

Learn more about the recommended tuning tasks:
[resources.arcgis.com/en/help/main/10.1/index.html#
/Minimize_disk_I_O_contention_in_SQL_Server/002q00000021000000/](https://resources.arcgis.com/en/help/main/10.1/index.html#/Minimize_disk_I_O_contention_in_SQL_Server/002q00000021000000/)

For more information on ArcGIS Workflow Manager, visit esri.com/workflowmanager or e-mail workflowmanager@esri.com.

Access blogs, forums, downloads, and more, from the ArcGIS Workflow Manager resource center at resources.arcgis.com/en/communities/workflow-manager/.

You can access other resources at ArcGIS 10.1 for Desktop Help:
http://resources.arcgis.com/en/help/main/10.1/index.html#/Welcome_to_the_ArcGIS_Help_Library/00qn0000001p000000/ and Esri Support (support.esri.com)



Esri inspires and enables people to positively impact their future through a deeper, geographic understanding of the changing world around them.

Governments, industry leaders, academics, and nongovernmental organizations trust us to connect them with the analytic knowledge they need to make the critical decisions that shape the planet. For more than 40 years, Esri has cultivated collaborative relationships with partners who share our commitment to solving earth's most pressing challenges with geographic expertise and rational resolve. Today, we believe that geography is at the heart of a more resilient and sustainable future. Creating responsible products and solutions drives our passion for improving quality of life everywhere.



Contact Esri

380 New York Street
Redlands, California 92373-8100 USA

1 800 447 9778
T 909 793 2853
F 909 793 5953
info@esri.com
esri.com

Offices worldwide
esri.com/locations