

High Availability and Disaster Recovery

- ▶ The main point of a high-availability strategy is to keep the critical data as available as possible.
- ▶ The different high availability solutions provide ways to keep a database system online even at the time of hardware failure or other unforeseen incidents.
- ▶ Disaster recovery efforts are all about restoring systems and data to a previous acceptable state in the event of partial or complete failure of computers due to natural or technical causes

Recovery Point Objective (RPO)

- ▶ The RPO is the amount of data one can lose, measured by time.
- ▶ It is likely to be different from system to system and application to application.
- ▶ For instance, a critical system may have an RPO measured in minutes while a non-critical one may have it measured in days.

Recovery Time Objective (RTO)

- ▶ How much time is needed to bring systems online in case a disaster happened.
- ▶ DBAs need to work closely with other members of the IT team to establish the RPO and RTO

High availability and Disaster recovery Solutions

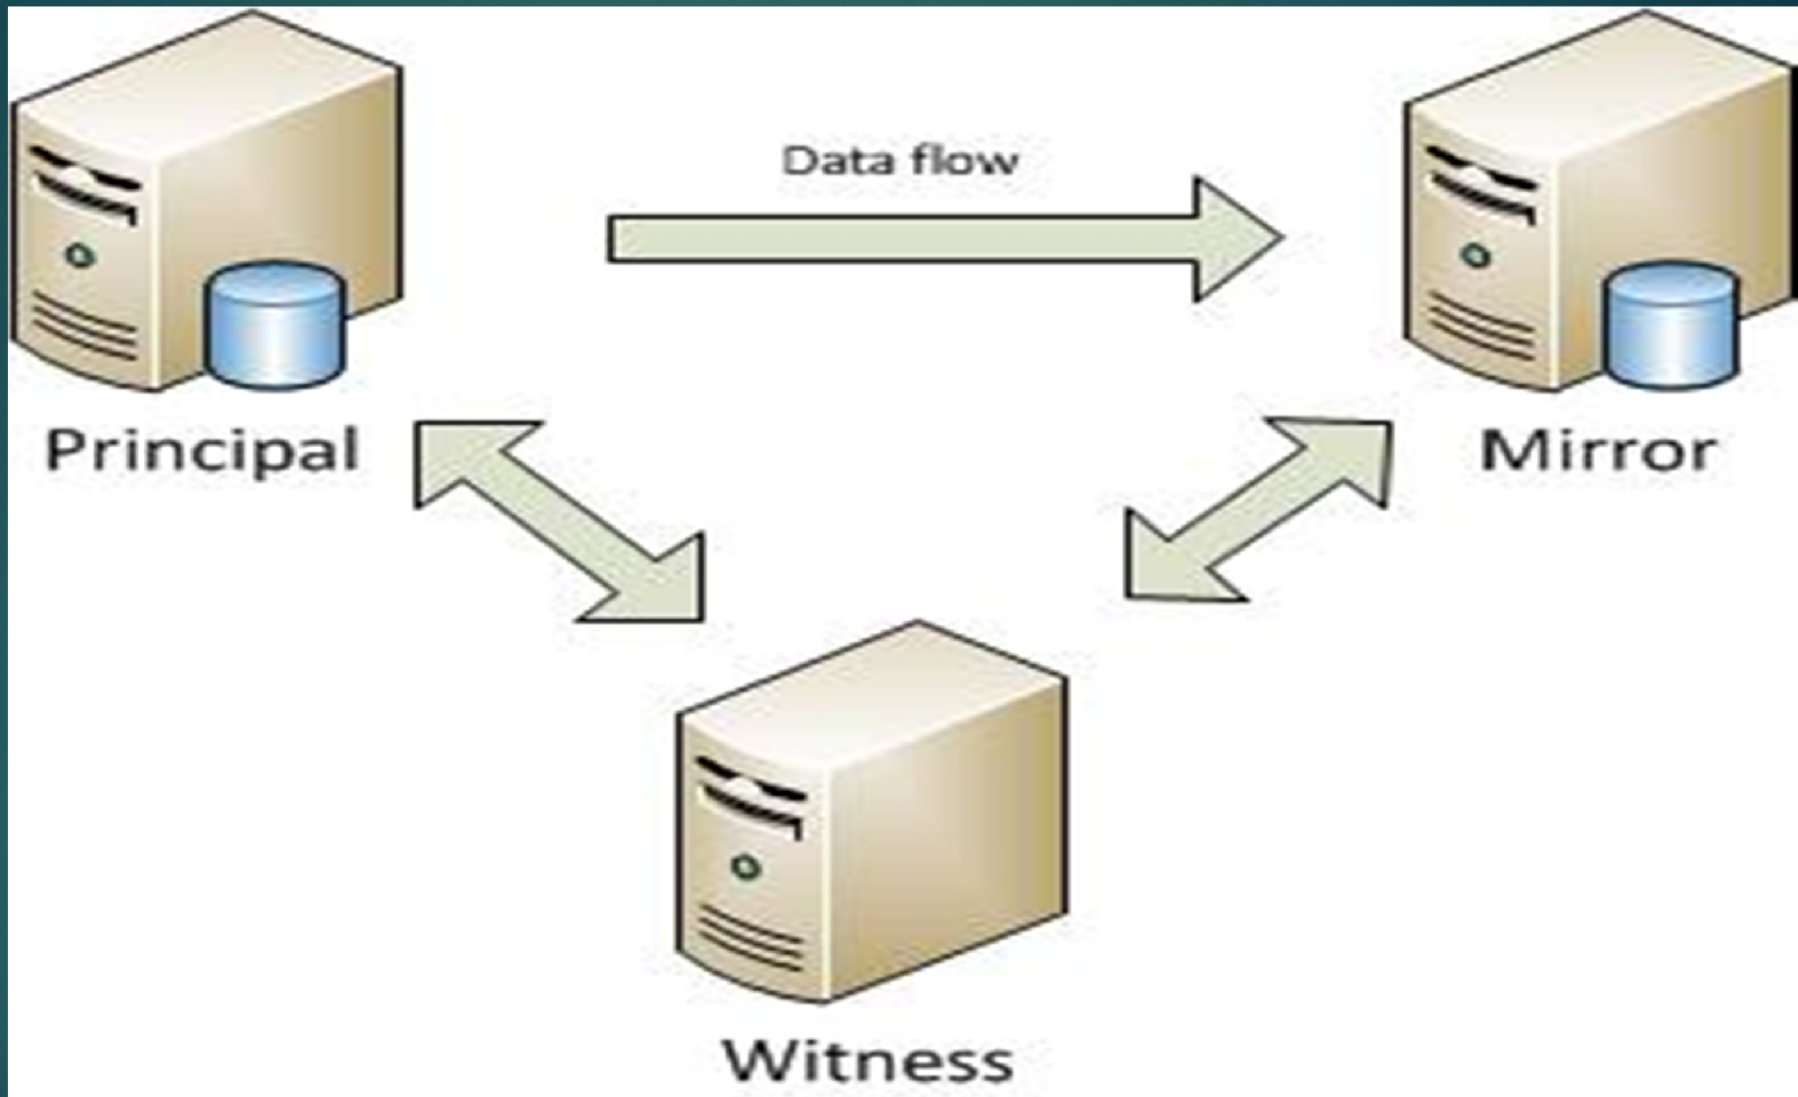
- ▶ Mirroring
- ▶ Log Shipping
- ▶ Replication
- ▶ AlwaysOn availability

Mirroring

- SQL Server database mirroring is a disaster recovery and high availability technique that involves two SQL Server instances on the same or different machines. One SQL Server instance acts as a primary instance called the **principal**, while the other is a mirrored instance called the **mirror**. In special cases, there can be a third SQL Server instance that acts as a witness
 1. Principal Instance - Take a full backup and a log backup as well
 2. Copy the full/log backups from Principal Instance to Mirror instance
 3. Mirror Instance - Restore with NORECOVERY option the full backup
 4. Mirror Instance – Apply the log backup
 5. Principal Instance - Start synchronization

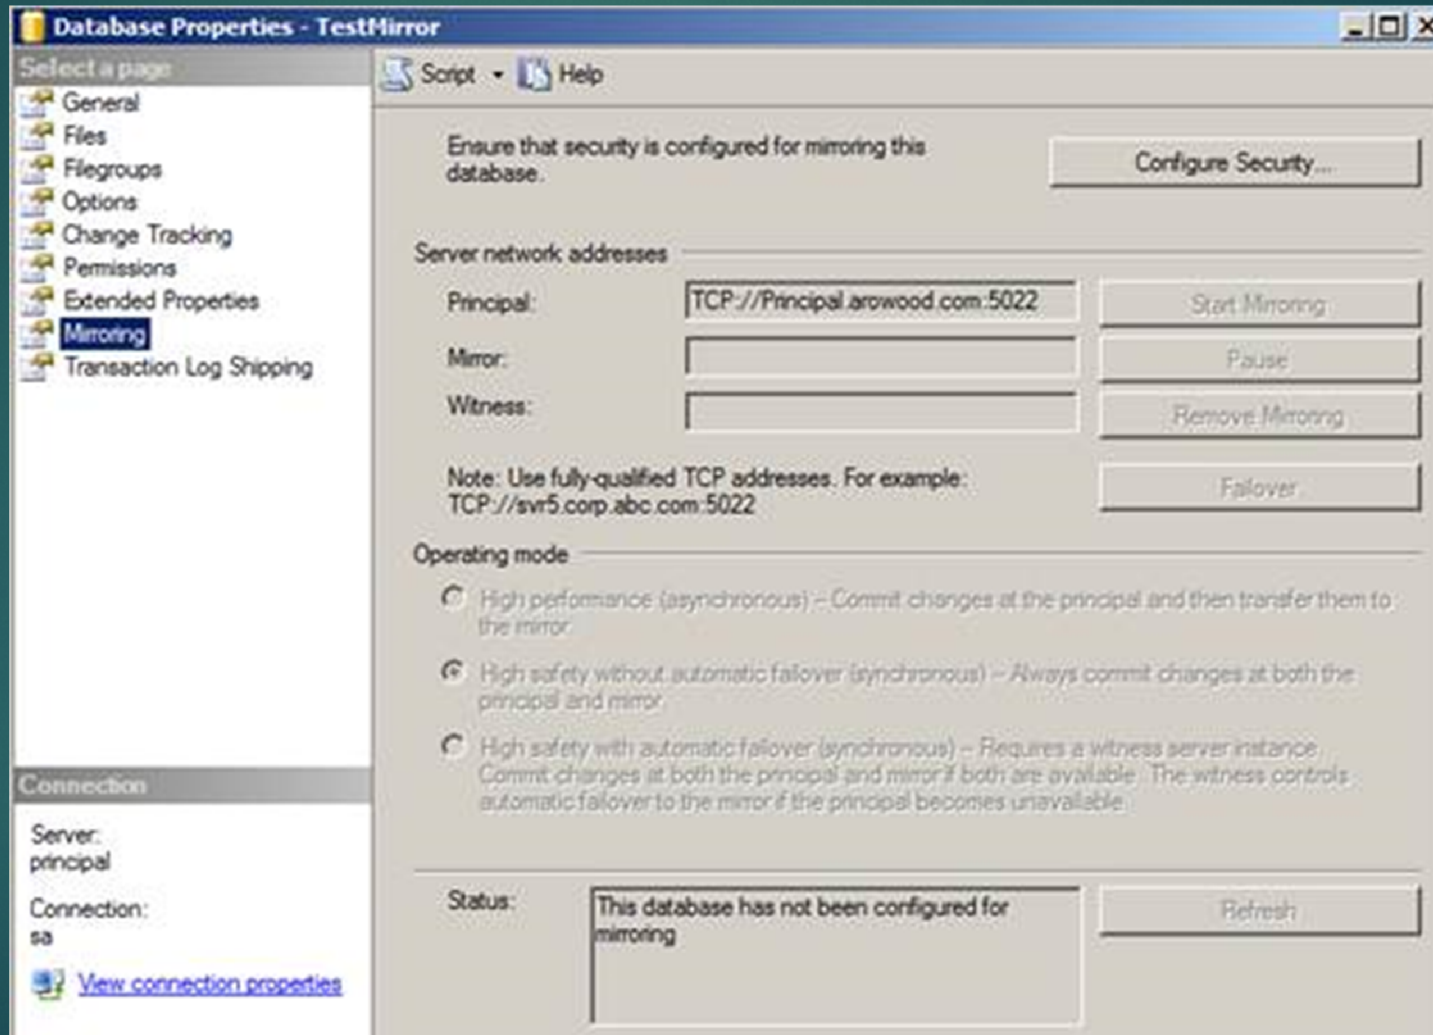
Witness: The role of the witness is to verify whether a given partner server is up and functioning. If the mirror server loses its connection to the principal server but the witness is still connected to the principal server, the mirror server does not initiate a failover

Mirroring



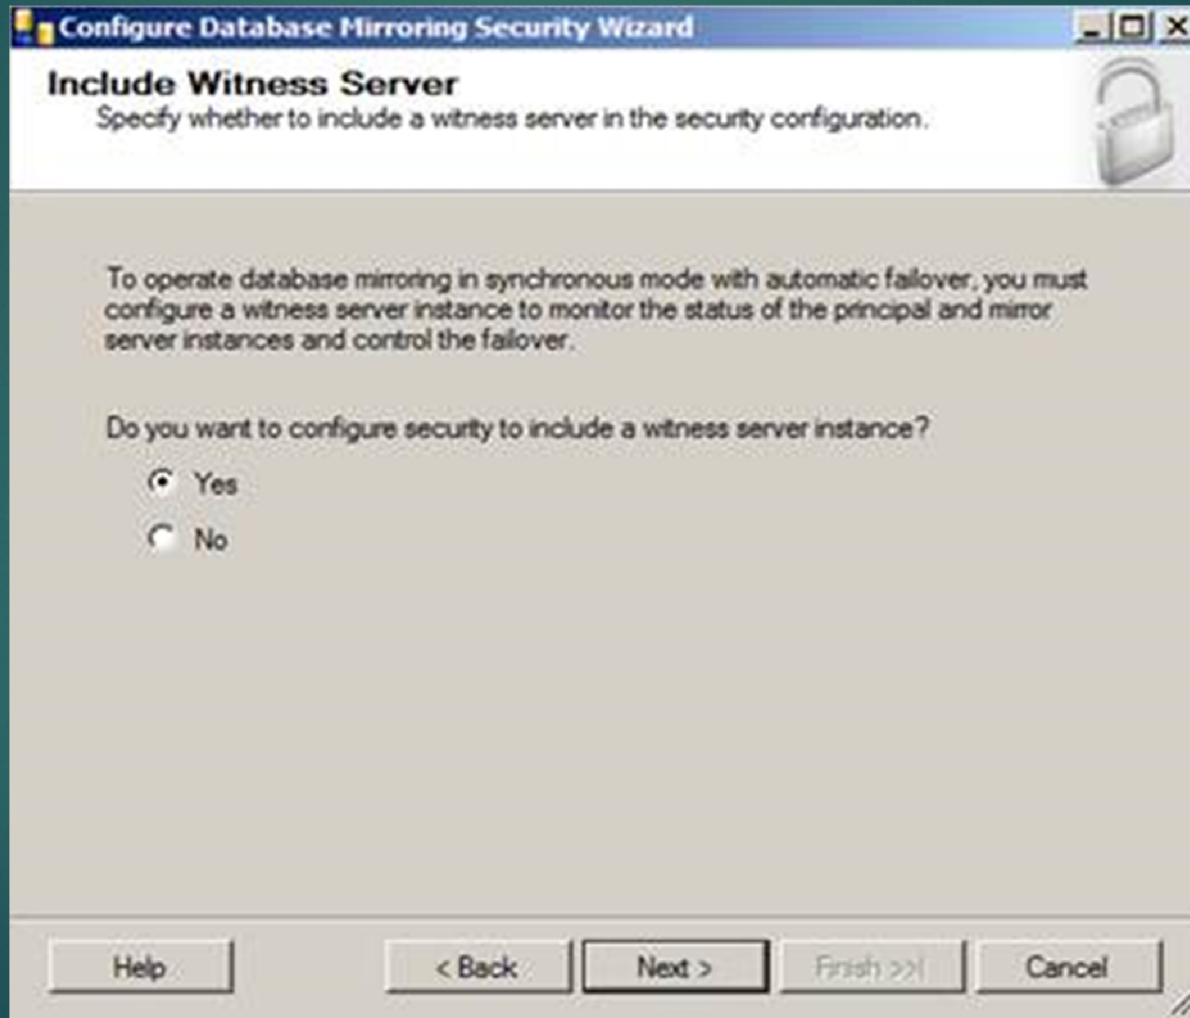
Mirroring

- From the Principal server, right click the database and choose "Tasks" | "Mirror" or choose "Properties" | "Mirroring".



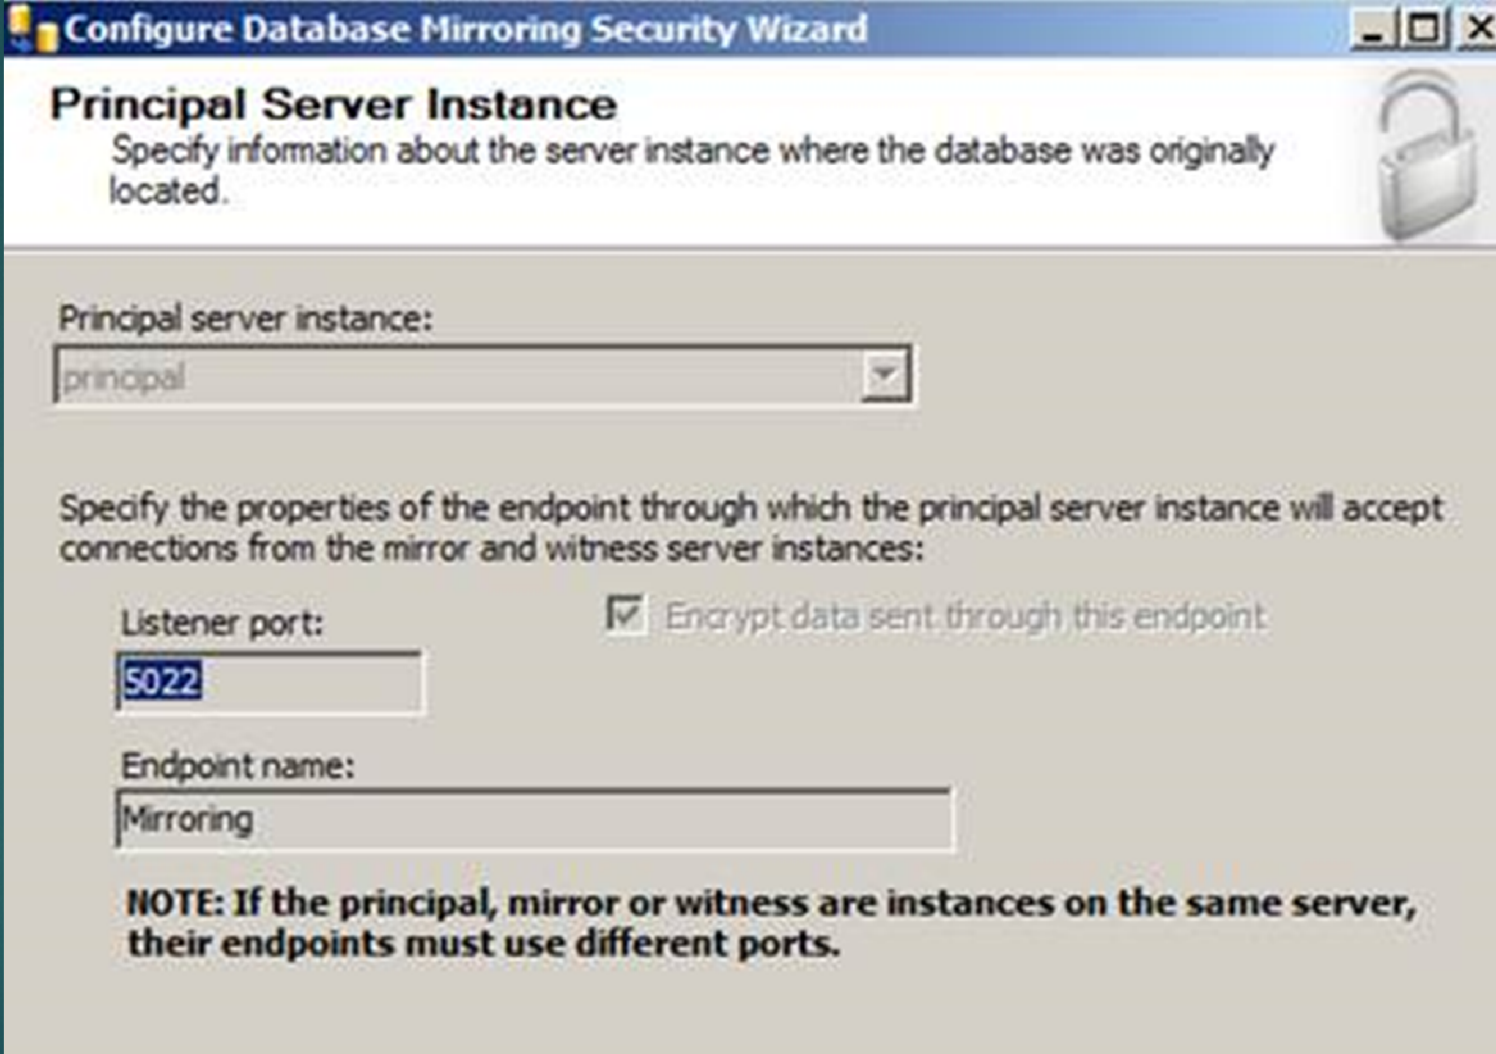
Mirroring

- Click the "Configure Security" button and click "Next >" if the Configure Database Mirroring Security Wizard intro screen appears. The next screen should be the Include Witness Server screen:



Mirroring

- The next screen will give you options to configure the Principal Server Instance



Configure Database Mirroring Security Wizard

Principal Server Instance

Specify information about the server instance where the database was originally located.

Principal server instance:
principal

Specify the properties of the endpoint through which the principal server instance will accept connections from the mirror and witness server instances:

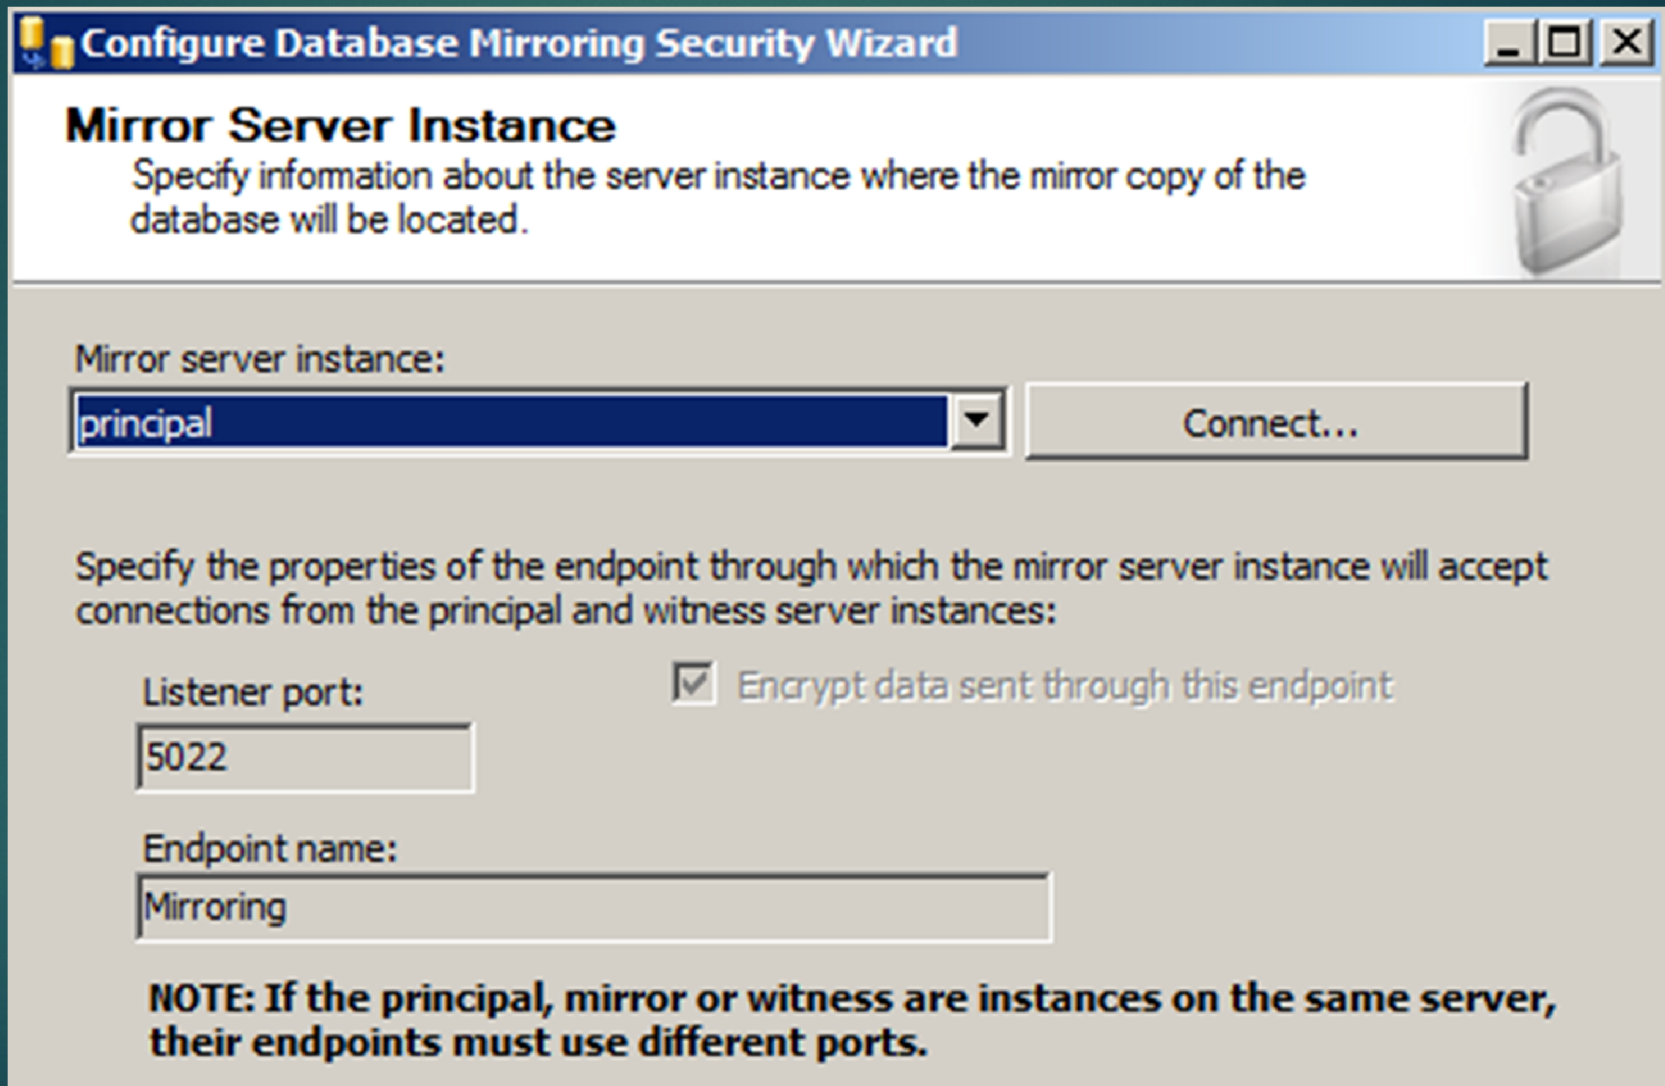
Listener port: Encrypt data sent through this endpoint

Endpoint name:

NOTE: If the principal, mirror or witness are instances on the same server, their endpoints must use different ports.

Mirroring

- The next screen will give you options to configure the Mirror Server Instance



Configure Database Mirroring Security Wizard

Mirror Server Instance

Specify information about the server instance where the mirror copy of the database will be located.

Mirror server instance:

principal

Specify the properties of the endpoint through which the mirror server instance will accept connections from the principal and witness server instances:


Listener port: Encrypt data sent through this endpoint

Endpoint name:

NOTE: If the principal, mirror or witness are instances on the same server, their endpoints must use different ports.

Mirroring

- To connect to the Mirror server instance we will need to click the "Connect..." button then select the mirror server and provide the correct credentials:



The screenshot shows the "Connect to Server" dialog box for Microsoft SQL Server 2008 R2. The dialog is titled "Connect to Server" and features the Microsoft SQL Server logo. The fields are as follows:

- Server type: Database Engine
- Server name: mirror
- Authentication: SQL Server Authentication
- Login: sa
- Password: [Redacted]
- Remember password

At the bottom of the dialog, there are four buttons: "Connect", "Cancel", "Help", and "Options >>".

Mirroring

- Once connected, we also notice our endpoint name is Mirroring and we are listening on port 5022.
Click "Next >" and you'll see the Service Accounts screen.



Configure Database Mirroring Security Wizard

Service Accounts

Specify the service accounts of the server instances.

For SQL Server accounts in the same domain or trusted domains, specify the service accounts below. If the accounts are non-domain accounts or the accounts are in untrusted domains, leave the textboxes empty.

Service accounts for the following instances:

Principal:

Mirror:

After you specify the service accounts, logins will be created for each account, if necessary, and will be granted CONNECT permission on the endpoints.

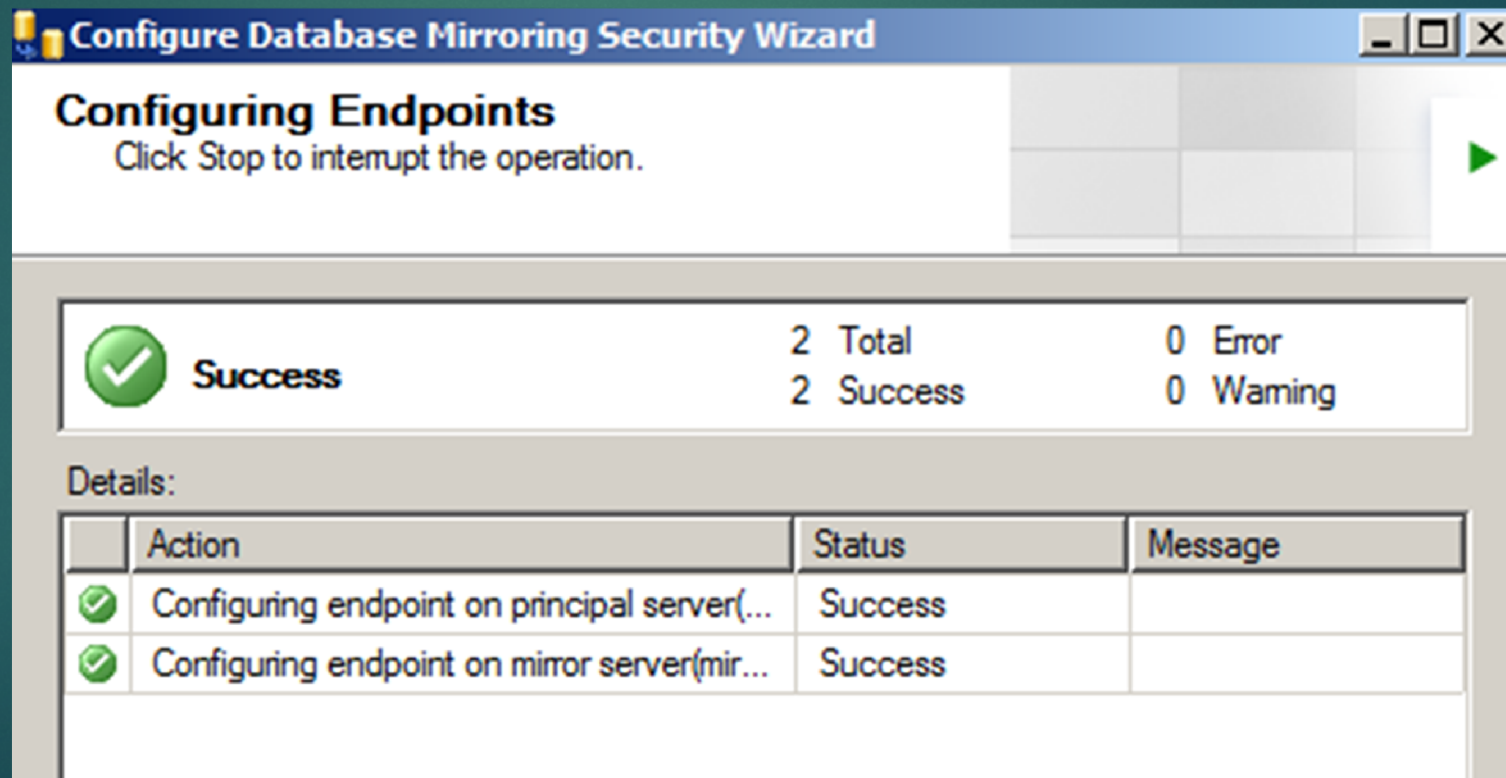
Mirroring

When using Windows Authentication, if the server instances use different accounts, specify the service accounts for SQL Server. These service accounts must all be domain accounts (in the same or trusted domains).

If all the server instances use the same domain account or use certificate-based authentication, leave the fields blank.

Since my service accounts are using the same domain account, I'll leave this blank.

Click "Finish" and you'll see a Complete the Wizard screen that summarizes what we just configured. Click "Finish" one more time.



Mirroring

- If you see the big green check mark that means Database Mirroring has been configured correctly. However, just because it is configured correctly doesn't mean that database mirroring is going to start...

Next screen that pops up should be the Start/Do Not Start Mirroring screen:

Database Properties



Specified database mirroring configuration settings :

Principal network address: TCP://Principal .com:5022

Mirror network address: TCP://MIRROR com:5022

Witness network address: None

Operating mode: High safety without automatic failover (synchronous)

To use the specified network addresses for mirroring this database, click Start Mirroring. To wait to start mirroring, click Do Not Start Mirroring; you can then start mirroring by clicking Start Mirroring on the Mirroring page of the Database Properties dialog box. Alternatively, you can exit the Database Properties dialog box without starting mirroring now, but you will need to configure the operating modes and server network addresses again before you can start mirroring.



Start Mirroring


Do Not Start Mirroring

Log Shipping

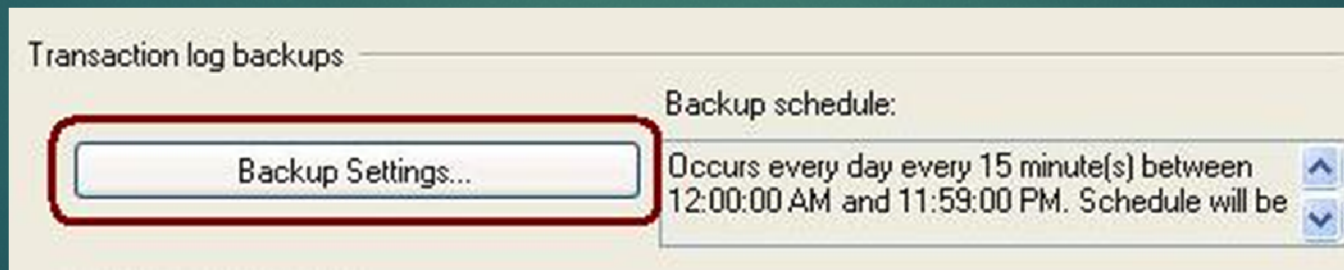
- Log Shipping is a basic level SQL Server high-availability technology that is part of SQL Server. It is an automated backup/restore process that allows you to create another copy of your database for failover.
- Log shipping involves copying a database backup and subsequent transaction log backups from the primary (source) server and restoring the database and transaction log backups on one or more secondary (Stand By / Destination) servers. The Target Database is in a standby or no-recovery mode on the secondary server(s) which allows subsequent transaction logs to be backed up on the primary and shipped (or copied) to the secondary servers and then applied (restored) there.
- **Permissions**
 - To setup a log-shipping you must have sysadmin rights on the server.
- **Minimum Requirements**
 - SQL Server 2005 or later
 - Standard, Workgroup or Enterprise editions must be installed on all server instances involved in log shipping.
 - The servers involved in log shipping should have the same case sensitivity settings.
 - The database must use the full recovery or bulk-logged recovery model
 - A shared folder for copying T-Log backup files
 - SQL Server Agent Service must be configured properly
 - In addition, you should use the same version of SQL Server on both ends. It is possible to Log Ship from SQL 2005 to SQL 2008, but you can not do it the opposite way. Also, since Log Shipping will be primarily used for failover if you have the same versions on each end and there is a need to failover you at least know you are running the same version of SQL Server.

Log Shipping

- On the primary server, right click on the database in SSMS and select Properties. Then select the **Transaction Log Shipping** Page. Check the "**Enable this as primary database in a log shipping configuration**" check box.



- The next step is to configure and schedule a transaction log backup. Click on Backup Settings... to do this.



Log Shipping

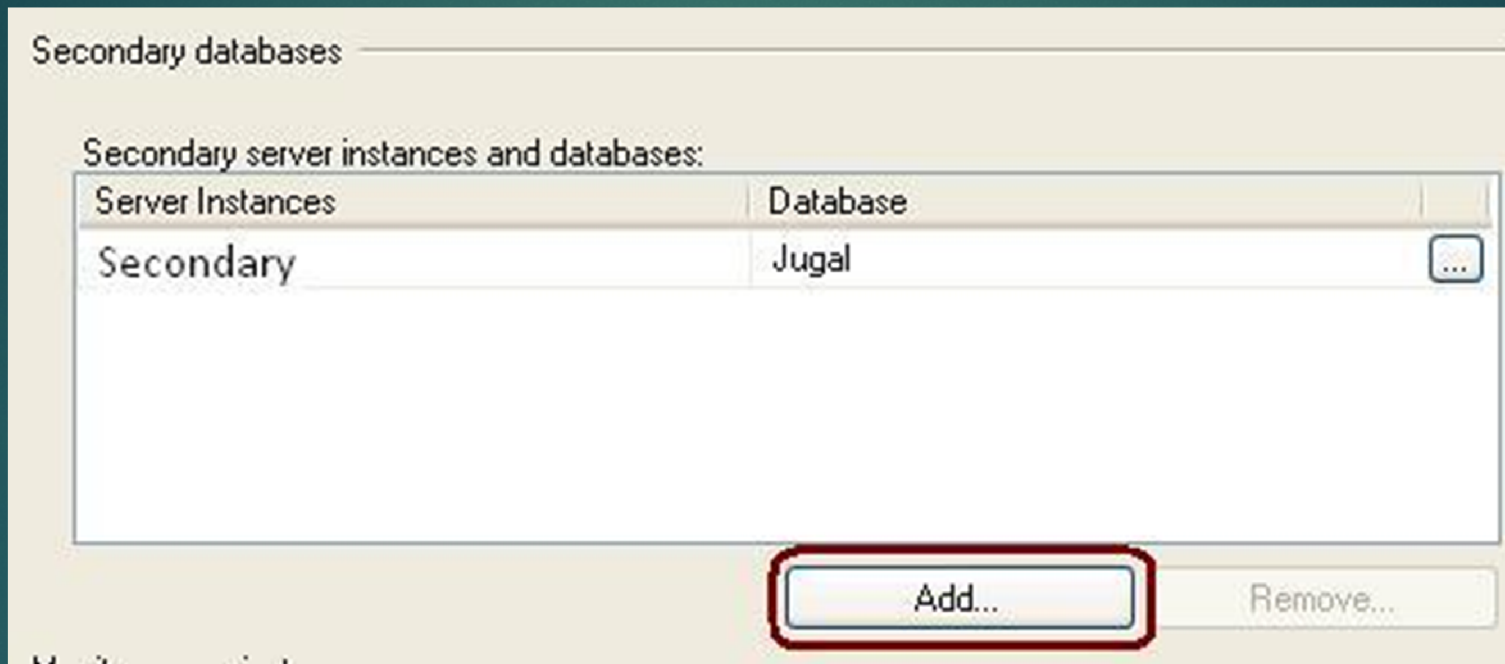
- If you are creating backups on a network share enter the network path or for the local machine you can specify the local folder path.

The screenshot shows the 'Transaction Log Backup Settings' dialog box. It contains the following fields and options:

- Transaction log backups are performed by a SQL Server Agent job running on the primary server instance.**
- Network path to backup folder (example: \\fileserver\backup):** A text box containing '\\backupserver\'.
- If the backup folder is located on the primary server, type a local path to the folder (example: c:\backup):** An empty text box.
- Note:** you must grant read and write permission on this folder to the SQL Server service account of this primary server instance. You must also grant read permission to the proxy account for the copy job (usually the SQL Server Agent service account for the secondary server instance).
- Delete files older than:** A dropdown menu set to '72' and another dropdown menu set to 'Hour(s)'.
- Alert if no backup occurs within:** A dropdown menu set to '1' and another dropdown menu set to 'Hour(s)'.
- Backup job:**
 - Job name:** A text box containing 'LSBackup_Jugal' and a 'Schedule...' button.
 - Schedule:** A text box containing 'Occurs every day every 15 minute(s) between 12:00:00 AM and 11:59:00 PM. Schedule will be used starting on 1/3/2011.' and a 'Disable this job' checkbox.
- Compression:**
 - Set backup compression:** A dropdown menu with options: 'Use the default server setting', 'Use the default server setting', 'Compress backup', and 'Do not compress backup'. The second 'Use the default server setting' option is currently selected.
- Note:** If you backup the transaction logs of this database with any of the following options, you will not be able to restore the backups on the secondary server instances.
- Buttons:** 'Help', 'OK', and 'Cancel'.

Log Shipping

- In this step we will configure the secondary instance and database. Click on the **Add...** button to configure the Secondary Server instance and database. You can add multiple servers if you want to setup one to many server log-shipping.



- When you click the Add... button it will take you to the below screen where you have to configure the Secondary Server and database. Click on the Connect... button to connect to the secondary server. Once you connect to the secondary server you can access the three tabs as shown below.

Log Shipping

- In this step you can specify how to create the data on the secondary server. You have three options: create a backup and restore it, use an existing backup and restore or do nothing because you have manually restored the database and have put it into the correct state to receive additional backups

Secondary Database Settings

Secondary server instance: secondary Connect...

Secondary database: Jugal
Select an existing database or enter the name to create a new database.

Initialize Secondary Database Copy Files Restore Transaction Log

You must restore a full backup of the primary database into secondary database before it can be a log shipping destination.

Do you want the Management Studio to restore a backup into the secondary database?

Yes, generate a full backup of the primary database and restore it into the secondary database (and create the secondary database if it doesn't exist) Restore Options...

Will take the fresh backup of primary database and restore it on secondary server

Yes, restore an existing backup of the primary database into the secondary database (and create the secondary database if it doesn't exist)

Will use the existing backup of the primary database restore it on secondary server

Specify a network path to the backup file that is accessible by the secondary server instance.

Backup file: Restore Options...

No, the secondary database is initialized.

Help OK Cancel

Log Shipping

- In this step you can specify how to create the data on the secondary server. You have three options: create a backup and restore it, use an existing backup and restore or do nothing because you have manually restored the database and have put it into the correct state to receive additional backups

Secondary Database Settings

Secondary server instance: secondary

Secondary database: Jugal
Select an existing database or enter the name to create a new database.

Initialize Secondary Database **Copy Files** Restore Transaction Log

Files are copied from the backup folder to a destination folder by a SQL Server Agent job running on the secondary server instance.

Destination folder for copied files: (This folder is usually located on the secondary server.)
\\secondary\SQLDBPool_Logshipping_Secondary

Note: you must grant read and write permission on this folder to the proxy account for the copy job (usually the SQL Server Agent service account on the secondary server instance).

Delete copied files after: 72 Hour(s)

Copy job

Job name: LSCopy_Secondary_Jugal

Schedule: Occurs every day every 15 minute(s) between 12:00:00 AM and 11:59:00 PM. Schedule will be used starting on 12/29/2010. Disable this job

Log Shipping

- Here you have to specify the database restoring state information and restore schedule. This will create the restore job on the secondary server.

Secondary Database Settings

Secondary server instance: Secondary Connect...

Secondary database: Jugal
Select an existing database or enter the name to create a new database.

Initialize Secondary Database | Copy Files | **Restore Transaction Log**

Files are restored from the destination folder by a SQL Server Agent job running on the secondary server instance.

Database state when restoring backups:

No recovery mode Secondary DB will be in restoring state, no-one can read it until its online.

Standby mode Secondary DB will be in read only mode, user can read the data.

Disconnect users in the database when restoring backups

Delay restoring backups at least: 0 Minute(s)

Alert if no restore occurs within: 45 Minute(s)

Restore job

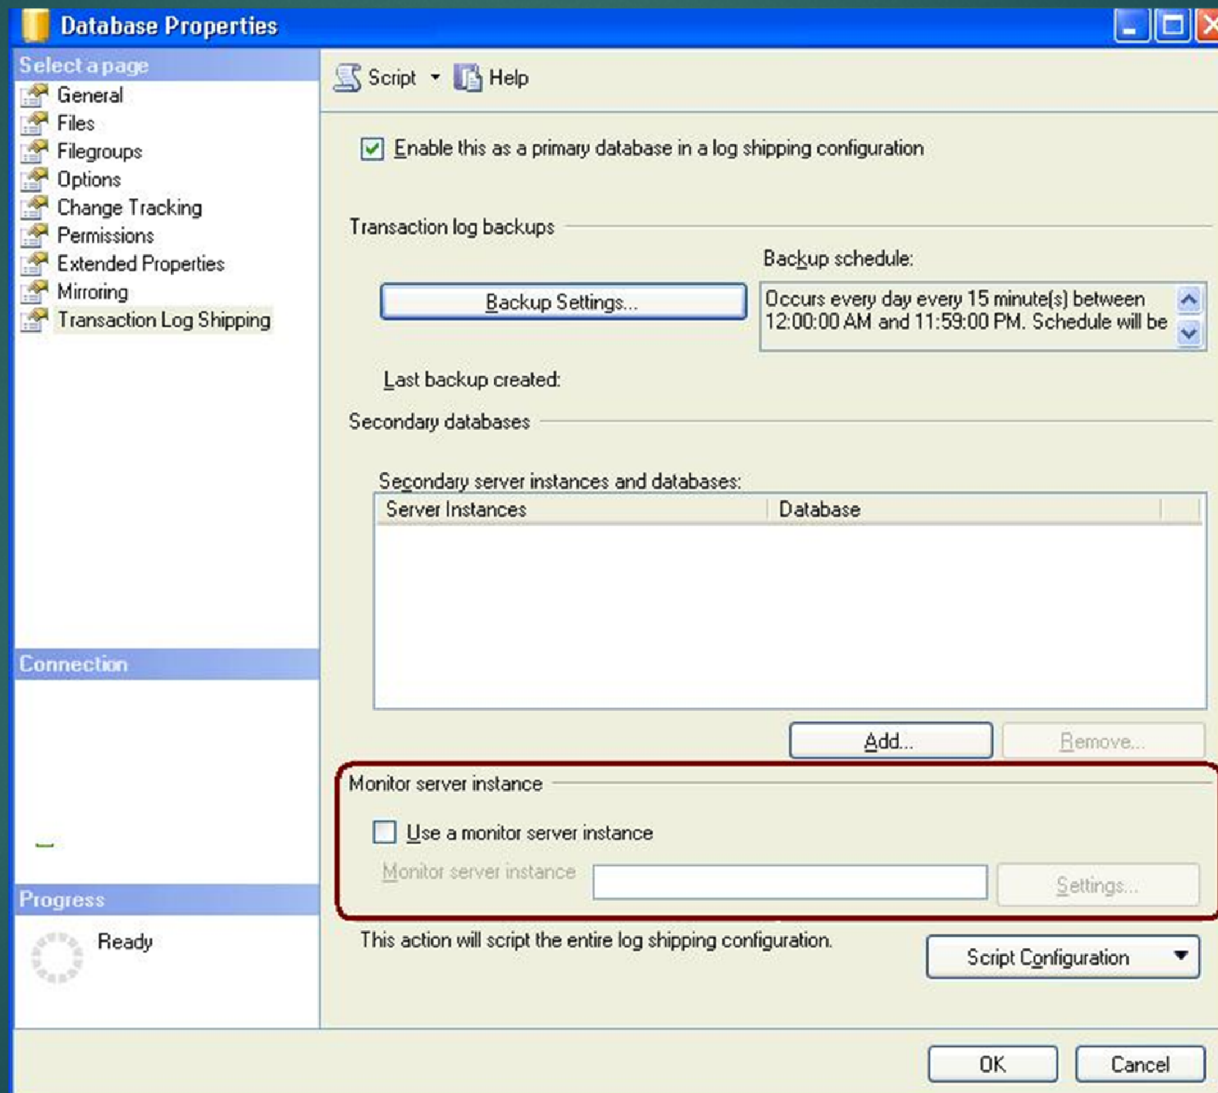
Job name: LSRrestore_primary_Jugal Schedule...

Schedule: Occurs every day every 15 minute(s) between 12:00:00 AM and 11:59:00 PM. Schedule will be used starting on 12/29/2010. Disable this job

Help OK Cancel

Log Shipping

- In this step we will configure Log Shipping Monitoring which will notify us in case of any failure. Please note Log Shipping monitoring configuration is optional.



Log Shipping

Click on **Settings...** button which will take you to the "**Log Shipping Monitor Settings**" screen. Click on **Connect ...** button to setup a monitor server. Monitoring can be done from the source server, target server or a separate SQL Server instance. We can configure alerts on source / destination server if respective jobs fail. Lastly we can also configure how long job history records are retained in the MSDB database. Please note that you cannot add a monitor instance once log shipping is configured.

The screenshot shows the "Log Shipping Monitor Settings" dialog box. The title bar is blue with the text "Log Shipping Monitor Settings" and a close button. The main area is white with a light blue border. It contains several sections: "Monitor server instance" with a text box containing "Secondary" and a "Connect..." button; "Monitor connections" with a heading and a sub-heading "Backup, copy, and restore jobs connect to this server instance:" followed by two radio button options: "By impersonating the proxy account of the job (usually the SQL Server Agent service account of the server instance where the job runs)" (selected) and "Using the following SQL Server login:" (unselected). Below the second option are three text boxes labeled "Login:", "Password:", and "Confirm Password:". "History retention" section has a "Delete history after:" label, a spinner box set to "96", and a dropdown menu set to "Hour(s)". "Alert job" section has a "Job name:" label and a text box containing "LSAlert_secondary", a "Schedule:" label and a text box containing "Start automatically when SQL Server Agent starts", and a "Disable this job" checkbox which is unchecked. At the bottom are three buttons: "Help", "OK", and "Cancel".

Log Shipping Monitor Settings

The monitor server instance is where status and history of log shipping activity for this primary database are recorded. It is also where the log shipping alert job runs.

Monitor server instance:
Secondary

Monitor connections

Backup, copy, and restore jobs connect to this server instance:

By impersonating the proxy account of the job (usually the SQL Server Agent service account of the server instance where the job runs)

Using the following SQL Server login:

Login:

Password:

Confirm Password:

History retention

Delete history after: 96 Hour(s)

Alert job

Job name: LSAAlert_secondary

Schedule: Start automatically when SQL Server Agent starts Disable this job

Log Shipping

- Click on the **OK** button to finish the Log Shipping configuration and it will show you the below screen.

Restoring backup to secondary database

Success 5 Total
5 Success 0 Error
0 Warning

Details:

Action	Status	Message
Backing up primary database [Jugal]	Success	
Restoring backup to secondary databa...	Success	
Saving secondary destination configura...	Success	
Saving primary backup setup	Success	
Saving Monitor configuration	Success	

Filter Close Report