Enterprise Deployment: Failover Clustering Considerations for Laserfiche

White Paper

Laserfiche®

Table of Contents

Overview

You can use Laserfiche with failover clustering technology to maintain high availability; allowing users to access Laserfiche after a hard disk, network, software, or other type of failure occurs. You can achieve high availability by creating a cluster that contains multiple redundant machines (otherwise known as **nodes**). If a failure is detected on the active node, the cluster will automatically failover to a backup node, which takes over the original node's work without causing significant interruptions in service.

Scope of the Paper

Clustering is a broad and complex topic that primarily involves non-Laserfiche technology. Though it is not intended to be a step-by-step guide, this paper will provide a basic foundation for integrating Laserfiche Rio 10 with a Microsoft Server 2012 or 2016 failover cluster. The paper is designed as a starting point, not as a definitive guide, as more research will be required depending on the specifics of your system.

This paper will primarily focus on creating a failover cluster to support the core Laserfiche components:

- Laserfiche Server, including Laserfiche Full-Text Indexing and Search
- Laserfiche Directory Server
- The Microsoft SQL Server instance that hosts the repository database(s)
- Repository volumes
- Repository folder

Though other Laserfiche products and components will not be discussed in detail, the following table briefly explains how they work in a failover cluster.

Product	**Supports failover clustering?**
Workflow	Supported for the Workflow Subscriber and Message Queuing as of Workflow 8.3. The password for WFUSER$ is no longer tied to the local computer on which the Workflow Subscriber is running, so the Subscriber can fail over without needing to be reconfigured. Workflow Message Queuing is based on Microsoft's Message Queuing Services

	(MSMQ). Consult Microsoft's resources on how to build a clustered MSMQ role.
Web Client and WebLink	Supported. The IIS machine can run in a cluster. Depending on the configuration, users may have to sign in again after failover occurs. Microsoft's general recommendation for achieving high-availability of Internet Information Services (IIS) servers is by using Network Load Balancing (NLB). For more information, see the IIS documentation.
Laserfiche Mobile Server	Supported in the same way as Web Client and WebLink.
Forms	Failover clustering for the Forms Web Server is supported the same way that Web Client and WebLink are supported. Failover clustering for the Forms database can be configured following the instructions in this paper for the Laserfiche Server database.
Agenda Manager	Not supported.
Discussions	Not supported.
Import Agent	Not supported.
Quick Fields Agent	Not supported.
Connector	The XML files that store Connector profiles should be housed in shared storage so that all nodes in the cluster can access them.
Audit Trail	Audit Trail can work in a failover cluster as long as all nodes share the same config.xml configuration file. The Audit Trail SQL database can be

	clustered using Microsoft SQL's built-in support for clustering.

In addition, note that this paper will:

- Focus on configuring failover clustering for Laserfiche Rio 10.
- Focus exclusively on Microsoft's clustering technology, and specifically on Windows Server 2012 and later. Outside research is required for information on other clustering options.
- Only discuss software clustering, not hardware clustering, which requires specialized hardware to be installed on one or more machines in a cluster (Windows Server uses software to manage the cluster).
- Not discuss clustering or failover options provided through virtual machine technologies, such as those provided by VMware Distributed Resource Scheduler (DRS) or VMware Fault Tolerance. In some cases, virtualization may be a more robust and cost-efficient approach than physical failover clusters.
- Only focus on failover clusters, and not on other cluster types, such as network load balancing clusters (which help distribute application workload) or compute clusters (which are used for complex computational purposes). See the network load balancing white paper for clustering web products. For distributed computing, see the documentation for the Distributed Computing Cluster.

Overview of Failover Clustering in Microsoft Server 2012 and later

The programs, disks, and other tools used by a cluster are **resources**, which are organized into **resource groups**. When a node becomes unavailable, all resources in a resource group fail over together to a different node. **Dependencies** are relationships between resources that ensure the cluster brings resources online in the correct order during failover.

When clustering the Laserfiche Server itself, or when using Failover Clustering Instances with SQL Server, you will need to make sure that **shared storage** is available. With shared storage, application data (e.g., database files, repository volumes, audit logs) is not stored on the nodes, but rather on a disk that all nodes can access (e.g., a storage area network or SAN). Although the storage disk is shared, only the active node can read from or write to it at any given time. The storage disk can be physical or virtual, as Microsoft's server virtualization technology, Hyper-V, can be used to share virtual hard disks between different computers.

The **quorum** configuration in a cluster determines the number of node failures that the cluster can sustain. Nodes in a Windows Server Failover Cluster (WSFC) regularly communicate to establish their health. Unresponsive nodes are considered to be failed. A majority of the nodes must be healthy in order for the cluster to be healthy.

The guidance provided in this paper is independent of whether or not Hyper-V is used to create the nodes in the cluster. The nodes in this paper may be either physical or virtual.

Microsoft SQL Server

Microsoft SQL Server provides support for failover clustering, which can be utilized by all Laserfiche products that require a database.

Since SQL Server 2012, Microsoft has used the term Always On to refer to both Failover Clustering instances (FCIs) and availability groups. Either kind of high availability solution works with Laserfiche—the choice of which solution to use depends on your organization's business needs. Both FCIs and availability groups require a pre-existing cluster configured using Windows Server Failover Clustering (WSFC). You should install SQL Server after you have created and configured all cluster nodes.

Note: Oracle Database, which is also supported by various Laserfiche components and offers a failover solution, will not be covered in this paper. For more information, refer to Oracle Database documentation.

Note: Depending on your business needs, other high availability solutions for your SQL Server, such as VMWare High Availability, may be more appropriate. This section covers only SQL Always On high availability solutions.

Failover Clustering Instances

FCIs work in the same manner described in Clustering Scenarios. A cluster consists of several nodes on which an instance of SQL Server is installed. At any time, one node is active, meaning that it owns the resources of the SQL Server instance. Shared storage is required. The shared storage constitutes a single point of failure unless it is also backed up in some other way—FCI does not contain within itself a mechanism to recover from shared storage failures. When the active node fails, ownership of the SQL Server instance's resources transfers over to another node, which then becomes the active node.

Availability Groups

Availability groups provide a high availability setup where there is no shared storage. An availability group is a set of databases that fail over together. The database copies on different nodes are known as **replicas**. Each replica is hosted by an instance of SQL Server on a different node of the cluster. There is a **primary replica**, which is the main means of updating a database in the group. The other replicas are known as **secondary replicas**. In contrast with FCIs, each secondary replica contains all the information in the databases in the availability group—there is no shared storage. To keep their copies of the database up-to-date, the secondary replicas regularly communicate with the primary replica.

Availability groups differ from FCIs as follows:

- The lack of shared storage and the multiple copies of the databases mean that availability groups do not have a single point of failure. In contrast, shared storage is a single point of failure in FCIs (unless measures are taken to make the shared storage highly available).
- Availability group replicas live on multiple instances of SQL Server. Each replica requires a different instance of SQL Server. In an FCI cluster, there is only one instance of SQL Server, and the active node owns the resources of the instance. The different nodes do not host separate instances of SQL Server.

- When using availability groups, only the databases fail over, not the entire instance of SQL Server. This means that logins, certificates, and SQL Server Agent jobs, for example, do not fail over in availability groups. In contrast, FCIs fail over at the server level, which means that all data in the server, not just the databases, is protected.

See the Microsoft documentation for more information comparing FCIs with availability groups.

It is possible to combine availability groups with Failover Clustering Instances. However, in this situation, the availability groups can no longer automatically fail over (FCIs can still fail over, in which case the replicas hosted by the node that is failing over will be moved together with the SQL Server instance to another node). Planned manual failovers and forced manual failovers for availability groups are still possible in this combined setup.

Tips for Using SQL Always On with Laserfiche

The following actions may improve your experience of using SQL Always On with Laserfiche:

- If you choose to use Always On availability groups, you will need to configure a listener for each availability group. This requires a unique DNS name and virtual IP addresses. You should use the listener's IP address or DNS name when configuring the database for your Laserfiche server. If you are using Windows authentication to connect to SQL Server, you will also need to register a Service Principal Name for the listener so that your Laserfiche application can use Kerberos to authenticate to the underlying SQL Server instances that fall under the listener.
- You will need ODBC Driver 11 or later for SQL Server to implement availability groups.
- After you have implemented SQL Server Always On, make the Laserfiche Server resource dependent on the SQL Server resource.
- To speed up the failover process and minimize downtime, add the following string to the SQL Connection string for your Laserfiche Server:
    - `MultiSubnetFailover=True;`

  For Laserfiche Directory Server, you can specify a new connection string by navigating to **Settings** tab on the Directory Server configuration site. In the **General** tab under Settings, locate the **Licensing Site** section and click **Modify connection string**. In the **Modify Connection String** dialog box, you can specify a new connection string to replace the default one. You must enter a full

connection string—the dialog box does not let you append to the default connection string.

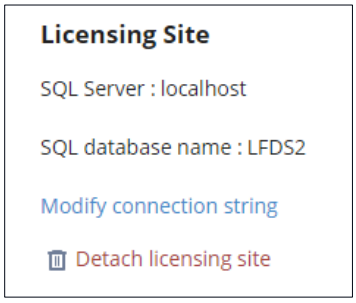

For Laserfiche Forms, the connection string is automatically set when you connect to the database from the Forms Configuration page. This string is written to the files **Web.config** and **RoutingEngineServiceHost.exe.config**, located in the Forms installation folder (by default, the relevant paths are **C:\Program Files\Laserfiche\Laserfiche Forms\Forms** for Web.config and **C:\Program Files\Laserfiche\Laserfiche Forms\Forms\bin** for RoutingEngineServiceHost.exe.config). To specify a different connection string, you can manually edit the connection string in these files. The connection string is located in the `<connectionStrings>` block in both files.

- Use performance counters to monitor your SQL Always On setup. See Microsoft's advice on monitoring Always On availability groups.
- Avoid turning on read-only routing to secondary replicas for Always On availability groups. This feature is not yet supported by Laserfiche.
- When upgrading to a newer version of Forms, you will need to update your Forms database. This action will result in an error if Always On availability groups are enabled. To avoid this, temporarily disable availability groups before you update your database, then turn the feature back on after the update.

The following settings should be chosen based on your business needs. Laserfiche's requirements alone do not dictate any particular choice.

- **Quorum modes**. Microsoft recommends having an odd number of nodes in a quorum. If you are using FCIs with two nodes and shared storage, the shared storage can add a third vote to make the total number of votes odd.
- **Failover modes**. While automatic failover offers the highest availability, it may not be possible in some circumstances, for example if you want to use availability groups in combination with FCIs. Other modes of failover may involve data loss and more downtime.
- **Availability modes for availability groups**. Either synchronous or asynchronous commits are possible. Asynchronous commits minimize transaction latency but

allow secondary databases to lag behind primary databases in their updates, increasing the risk of data loss. Synchronous commits increase transaction latency but protect against data loss. Synchronous commits are necessary for automatic failover. See Microsoft's documentation on how availability modes influence which failover modes are possible.

SQL Always On Resources

The following resources may help you plan for your SQL Always On setup:

- Prerequisites, restrictions, and recommendations for configuring Always On availability groups
- Overview of Always On availability groups
- Microsoft's recommended adjustments to quorum voting
- How the different failover modes work for availability groups

Overview of Clustering for Laserfiche Components

The following sections highlight failover considerations for the various components of a Laserfiche installation

Laserfiche Repository Folder

The Laserfiche repository folder, which is specified during repository creation, contains configuration files the Laserfiche Server needs to interact with (by default, this folder also contains a Laserfiche volume and search index files). Shared storage should be created in order to house the Laserfiche repository folder, and the Laserfiche Server resource should be dependent on this storage disk.

Laserfiche Volumes

Laserfiche volumes contain images, text, electronic files, thumbnails, and word location data. A Laserfiche installation will continue to function after a volume experiences a hardware failure, although the contents of documents will not be accessible. If document access is mission-critical, you should incorporate volumes in a failover cluster. If you do so, we recommend not making the Laserfiche Server dependent on the disk resource that contains volumes, unless the volumes reside on the same disk resource as the Laserfiche repository folder or SQL Server database files.

Laserfiche Server

Though the Laserfiche Server can be used in a failover cluster, it is cluster-unaware, meaning its behavior does not change when running in a cluster. In addition, it operates as an active/passive application, meaning only one Laserfiche Server instance within a cluster should be running at any given time. In other words, while the Laserfiche repository should be attached to all Laserfiche Servers within a cluster, only the Laserfiche Server on the node that is in control of the quorum should be running and actively broadcasting a Laserfiche repository; all other instances should not be running. Before failover, the cluster should direct all Laserfiche Server requests to the cluster's active node. After failover, the cluster should stop the Laserfiche Server on the node that has failed (if necessary), start the Laserfiche Server on the new active node, and direct all Laserfiche Server requests to it.

Laserfiche Directory Server

Laserfiche Directory Server includes a Failover Cluster Support feature that can be included at installation. If this is included, the Directory Server administration site will have a built-in **Enable Cluster** option. With Directory Server 10.3 and later, modify the SQL Connection string to configure Always On Availability Groups for the Directory Server database.

This paper contains instructions on configuring Directory Server to work in a failover cluster.

Laserfiche Forms

Forms servers can work in a network load balancing configuration. The Forms SQL database can be configured to use SQL's failover clustering capabilities. When clustering the Forms database, take note of where you can change the SQL connection strings in the Forms configuration files.

Laserfiche Licensing

Your license type determines the number of Laserfiche Servers you can install and run. For Laserfiche Rio, there is no limit, so you can use Directory Server to create one license for each cluster node that will host a Laserfiche Server. For all other license types (Avante, Team, United), you must purchase one license for each Laserfiche Server you want to install/run.

> **Note:** Laserfiche offers discounted prices for Laserfiche Servers that are intended to run on failover servers. For more information, talk to your Laserfiche value-added reseller or account manager.

Regardless of your license type, you will need one unique license for each machine you want to install the Laserfiche Server on. This is because some Server features defined in the license, like the public portal, are licensed per server, so the Laserfiche Server must be activated for exclusive use on a particular machine (using a hardware ID defined in the license). If you had a cluster with two Laserfiche Server nodes (one active and the other for use during failover), you would need to purchase two public portal features in order for both Laserfiche Servers to support a public portal.

Upgrades for Clustered Applications

If you have Laserfiche Server or Directory Server configured in a failover cluster, you will have to take the entire cluster offline in order to upgrade to a newer version of Laserfiche Server or Laserfiche Directory Server. Turn off all cluster nodes, upgrade each node, then turn the nodes back on.

Sample Clustering Scenarios

In this section, we will explore some real-world clustering scenarios and configurations. In the following diagrams, shapes with red borders represent physical devices, while shapes with blue borders represent virtual servers. Windows supports failover clusters with either physical or virtual nodes. While the diagrams that follow include virtual nodes created on the physical nodes, scenarios 1 and 2 apply just as well to clusters that have only physical nodes. If you have only physical nodes, you may treat the Laserfiche clients in Scenarios 1 and 2 as connecting directly to one physical node, bypassing the virtual server layer. Scenario 3 applies only when you have virtual servers.

Scenario 1: Before Failover

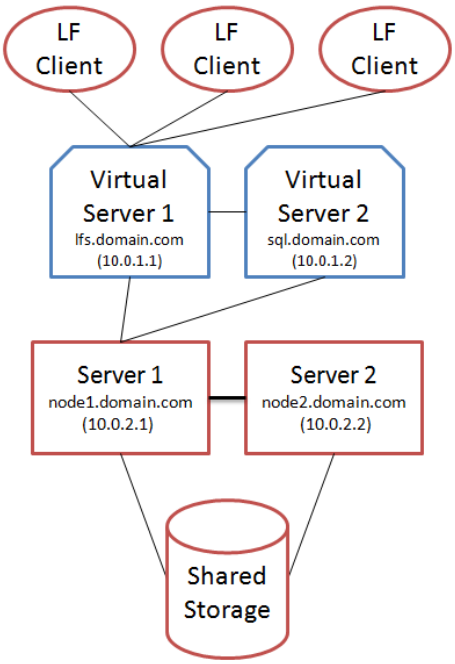

In this scenario, a Laserfiche Server and a Microsoft SQL Server instance are installed on both **Server 1** and **Server 2**. These two machines are clustered together to create:

- **Virtual Server 1:** Hosts the Laserfiche Server and can be accessed at **lfs.domain.com** or **10.0.1.1**
- **Virtual Server 2:** Hosts the SQL Server and can be accessed at **sql.domain.com** or **10.0.1.2**

Server 1 is the active node for both virtual servers, while **Server 2** waits idly on hot standby in case **Server 1** becomes unavailable.

When a Laserfiche client request is made to **lfs.domain.com**, it is forwarded to **node1.domain.com**. If necessary, **node1.domain.com** makes a database call to **sql.domain.com**, which is forwarded back to **node1.domain.com**.

Regardless of which cluster node is active, the repository folder and the SQL Server database files are both stored on **Shared Storage**.

> **Note:** To keep things simple, assume the repository volumes are stored on the same virtual server as the Laserfiche Server. You could separate these onto different virtual servers, if necessary.

Scenario 2: Failover Occurs

As a follow up to Scenario 1, consider the following diagram, which represents what will happen if **Server 1** becomes unavailable (e.g., motherboard failure):

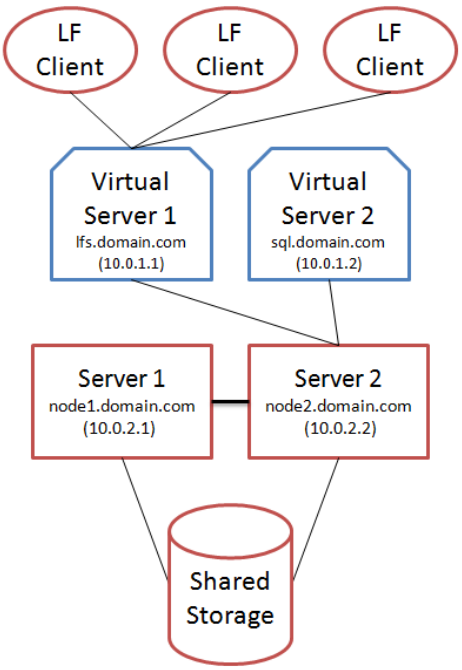

When the cluster determines that **Server 1** is unavailable, the quorum sets **Server 2** as the active node. All requests to **lfs.domain.com** and **sql.domain.com** will now be forwarded to **node2.domain.com**. Since all application data (repository folder, database files, and volumes) are stored on **Shared Storage**, no data loss will occur. In addition, there will be zero to minor interruption in service.

Scenario 3: Alternate Configuration

The previous two scenarios use a configuration where both the Laserfiche and the SQL servers run on the same physical node. To increase their return on investment (ROI), many organizations strive to avoid idle servers, so they may use a setup like that shown in the diagram below. This setup is possible only with virtual servers, so it applies only if you are using Microsoft's Hyper-V technology. A setup of this type has variously been called a multiple-instance or active/active configuration. However, it is not to be confused with an alternative interpretation of "active/active," used to describe configurations that enable load balancing. In load balancing, a single instance can be distributed over multiple nodes in order to even out network traffic loads. Our scenario involves no load balancing. Each instance (Laserfiche Server or SQL) operates on one node only. It can be thought of as combining a Laserfiche Server instance in active/passive mode with an SQL instance in active/passive mode.

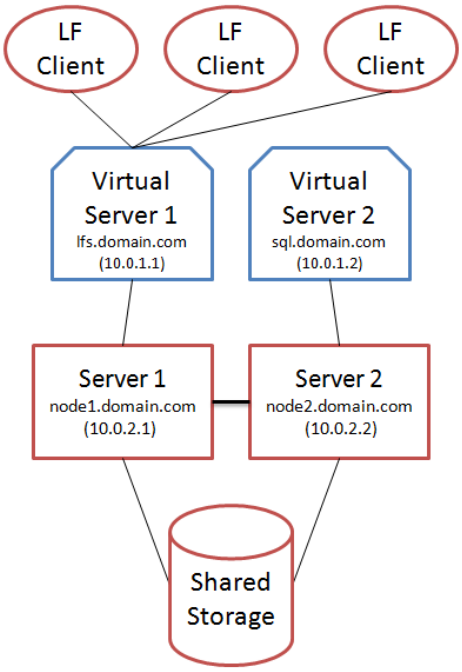

In this scenario, a Laserfiche Server and a Microsoft SQL Server instance are still installed on *both* **Server 1** and **Server 2**. However, before failover, the virtual servers each run on a separate node: **lfs.domain.com** forwards to **node1.domain.com** and **sql.domain.com** forwards to **node2.domain.com**.

If **Server 1** becomes unavailable, the cluster will automatically begin forwarding **lfs.domain.com** to **node2.domain.com**. The same logic applies if **Server 2 becomes** unavailable, in which **sql.domain.com** will be forwarded to **node1.domain.com**. In other words, both physical servers are identical, but, before failover, each is responsible

for only a *single* virtual server. After failover, one of the physical servers becomes responsible for *both* virtual servers.

As mentioned earlier, this scenario also allows you to spread out the workload between two machines. In addition, you can take a server down for maintenance without causing any interruption in service, as the cluster will automatically failover to the adjacent node.

Configuration Details

The following sections document step-by-step configuration details for select Laserfiche components.

Configuring Laserfiche Server to Run in a Failover Cluster

After you have created and configured your WSFC cluster, install the Laserfiche Server and the Laserfiche Full-Text Indexing and Search service on each node.

To configure the Laserfiche Server to run in a server cluster:

1. From the Server Manager, select the **Failover Cluster Manager** from **Tools**.
2. Select your previously-created server cluster.
3. From the **Action** menu, select **Configure a Role**. This will open the **High Availability** wizard and display the **Before You Begin** step. Click **Next**.
4. In the **Select Service or Application** step, choose **Generic Service** from the list of services and applications. Click **Next**.
5. In the **Select Service** step, select the **Laserfiche Server** service. Click **Next**.
6. In the **Client Access Point** step, enter a name for the role. This will be displayed in the Failover Cluster Manager as the Client Access Name or Server Name, and cannot be identical to the cluster name. Click **Next**.
7. In the **Select Storage** step, select the disk you are using for shared storage. Click **Next**.
8. In the **Replicate Registry Settings** step, add **HKEY_LOCAL_MACHINE\SOFTWARE\Laserfiche\Engine** as the registry key to be replicated to all nodes in the cluster. Click **Next**.
9. In the **Confirmation** step, confirm that you are ready to configure the service for high availability. Once you have done so, the cluster will be configured to work with the Laserfiche Server.
10. In the **Summary** step, you can view the report created by the wizard or close the wizard by clicking **Finish**.
11. To add your SQL server, select the Server Name in the **Roles** pane.
12. From the **Actions** menu, select **Add Resource** and then **Generic Service**. This will open the **New Resource** wizard.

13. In the **Select Service** step, select your SQL server. Click **Next**, and confirm you want to add the service by clicking **Next** again.
14. In the **Resources** pane, right-click on **Laserfiche Server** and select **Properties**.
15. In the **Dependencies** tab, add the IP address of the cluster and your SQL Server as dependencies. Click **OK** to save changes and close the dialog.
16. To view your dependency report, right-click on **Laserfiche Server**, open the **More Actions** menu, and select **Show Dependency Report**. The dependency report will open in your browser, and should show that Laserfiche Server depends on the storage disk, the client access name, and the IP address you selected. Close the browser to close the report.
17. To add your search service, select the client access name in the **Roles** pane. In the **Action** menu, select **Add Resource** and then **Generic Service**. This will open the **New Resource** wizard again.
18. In the **Select Service** step, select the **Laserfiche Full-Text Indexing and Search Engine**. Click **Next**, and confirm you want to add the service by clicking **Next** again.
19. With the client access name selected, the **Laserfiche Full-Text Indexing and Search Engine** should appear in the **Resources** pane. Right-click on this resource and then select **Properties**.
20. In the **Dependencies** tab for the Laserfiche Full-Text Indexing and Search Engine service, designate **Laserfiche Server** as a dependent service.
21. In the Registry Replication tab, add **HKEY_LOCAL_MACHINE\SOFTWARE\Laserfiche\Engine**.
22. Click **OK** to save your changes.

Additional Settings

- **SQL Connection String:** On each node, make sure there is a **ConnectionString** string value in the Windows registry under **HKEY_LOCAL_MACHINE\SOFTWARE\Laserfiche\Engine\8.0\Repositories*SampleRepositoryName***. The value is an ODBC connection string to the appropriate SQL database. For a Microsoft SQL Server database, the values should be as follows:
    - For Microsoft SQL Server with Windows Authentication:

        - Driver={*Name of ODBC Driver 11 for SQL Server*};SERVER={*MSSQLServerName*};Trusted_Connection=yes;DATABASE={*DATABASENAME*};
    - For Microsoft SQL Server with SQL Server Authentication:
        - Driver={*Name of ODBC Driver 11 for SQL Server*};SERVER={*MSSQLServerName*};UID={*SqlLoginName*};PWD={*SqlPassword*};DATABASE={*DATABASENAME*};

  Replace *Name of ODBC Driver 11 for SQL Server*, *MSSQLServerName*, *DATABASENAME*, *SqlLoginName*, and *SqlPassword* with the values appropriate to your setup.

  With an Oracle database, the value will look similar to:

    - Driver={Oracle in OraClient10g_home1};DBQ=*ORACLESERVERNAME*;UID=*schemaName*;PWD=*schemaPassword*;

  The connection string is also first specified when you create a repository in the repository creation wizard, under ODBC driver options.
- **Failover Policy:** The failover policy for the Laserfiche Server determines the number of times that the cluster service will attempt to restart the Laserfiche Server when it detects that it is no longer started. If it is not able to start after the specified number of times, it will failover the active node. It is important to limit the number of restart attempts, since each attempt will increase the amount of downtime for a Laserfiche repository. You can adjust the policy for Laserfiche Server by right-clicking on the server name in the **Roles** panel, then selecting **Properties** and **Failover**.

Configuring Laserfiche Directory Server

These instructions assume that the cluster has already been created and exists in the Failover Cluster Manager in Windows Server.

To configure Windows Server for Laserfiche Directory Server:

1. Install Laserfiche Directory Server with the Laserfiche Directory Server Failover Cluster Support component on all nodes in the cluster. Complete the rest of the Windows Server Configuration steps here before proceeding to the **To configure Laserfiche Directory Server for failover clustering** section below.
2. Select the cluster you are using in the Failover Cluster Manager. In **Roles**, create and configure a new role.
    a. In the wizard, on the **Select Role** step, select **Other Server**.

b. On the **Client Access Point** step, specify a cluster role name. Note that this cannot be the same as the cluster name.

c. When prompted to select a Resource Type, choose **Laserfiche Directory Service Resource**.

3. Back in the Failover Cluster Manager, right-click on the newly created role and view the **Properties**.

   a. On the **General** tab, in the **Preferred Owners** section, select all the nodes.

   b. On the **Failover** tab, adjust failover settings as appropriate.

4. Check that Resource properties for the role are configured.

   a. Select the Role.

   b. At the bottom of **Failover Cluster Manager**, click the **Resources** tab.

   c. Under **Server Name**, right-click on the Server name (the "Client Access Point") and click **Properties**.

   d. Under Other Resources, right-click on **Laserfiche Directory Service Resource**, and click **Properties**.

   e. Start the cluster. Right-click on **Laserfiche Directory Service Resource** and click **Bring Online**.

Note: Please stop all Directory Server services before starting the cluster.

To configure Laserfiche Directory Server for failover clustering:

1. View the Laserfiche Directory Server Administration site.
2. On the **Settings** tab, enable the **Enable Cluster** (this action cannot be undone) option.
3. Next to “Cluster Role Name”, specify the cluster role name created in the previous section.
4. Next to “Cluster Fingerprint”, specify the hardware fingerprint of the first cluster node hosting Directory Server. You can retrieve your hardware fingerprint using the Hardware Fingerprint Utility, by navigating to the Laserfiche Server installation directory and double-clicking **showhwfp.exe**.
5. Click **Add new host to cluster** and add a cluster node. Repeat as necessary for all nodes in the cluster. Use the fully qualified domain name for each host name.
6. After configuring Directory for all the nodes, delete the file **C:\ProgramData\Laserfiche\LFDS\usersettings.config** on each node.

Note: The licensing database connections must be configured separately for each node. Administrators should create and activate the Directory Server database on one node, then attach the database to the other Directory Server nodes.

Creating/Registering Laserfiche Repositories

Once you have properly configured your server cluster, SQL Server, and Laserfiche, you are ready to create a repository or to register an existing one. When creating or registering a Laserfiche repository, keep the following in mind:

- The name of the server cluster should be used to identify the computer hosting the Laserfiche Server. This name was specified when creating the server cluster.
- If you are using Laserfiche Directory Server, you will need to generate licenses for the various Laserfiche products that will be registered under the Directory Server, such as the web client, Laserfiche Server, or Forms. You should generate licenses for these in the Laserfiche Directory Server using the cluster role name.
- If you are using Laserfiche Directory Server and configuring single sign-on, use the cluster role name as the Directory Server host name.
- The SQL Server instance is the name of the virtual SQL Server, followed by a backslash, and the instance name. If an instance name was not specified, you can use the name of the virtual SQL Server.
- The Laserfiche repository location should be set to a folder on shared storage.

Note: Once you have created or registered a Laserfiche repository, it will be immediately broadcast across your network. When attaching the Laserfiche repository to a Laserfiche client, or otherwise pointing clients at the repository, ensure you use the name of the server cluster, as this will ensure clients will be able to access the Laserfiche repository after a failover occurs.

Laserfiche®

Enterprise Deployment: Failover Clustering Considerations for Laserfiche
July 2018

Author: Leif Hancox-Li, Jonathan Powers, and Constance Anderson
Editor: Roger Wu

Description:
This paper explains how to use Laserfiche with failover clustering technology to maintain high availability. This enables users to access Laserfiche after a hard disk, network, software, or other type of failure occurs.

Laserfiche
3545 Long Beach Blvd.
Long Beach, CA 90807
U.S.A

Phone: +1.562.988.1688
www.laserfiche.com