Enterprise Deployment: Failover Clustering Considerations for Laserfiche

White Paper

March 2012

Laserfiche®

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

This chapter is for informational purposes only. COMPULINK MAKES NO WARRANTIES, EXPRESS OR IMPLIED, AS TO THE INFORMATION IN THIS DOCUMENT.

Table of Contents

Overview

You can 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. 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 interruption 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 with a Microsoft Server 2008 cluster using Microsoft Cluster Server (MSCS). 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
- 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 (only relevant products are listed; current released versions as of May 2011).

Product	**Supports failover clustering?**
Workflow	Not supported.
WebAccess and WebLink	Supported. The IIS machine can run in a cluster. Depending on the configuration, users may have to log in again after failover occurs. Microsoft's general recommendation for achieving high-availability of Internet Information Services (IIS)

	servers (which includes Web Access and WebLink) is by using Network Load Balancing (NLB). For more information, see IIS documentation.
Agenda Manager	Not supported.
Import Agent	Not supported.
Quick Fields Agent	Not supported.

In addition, note that this paper will:

- Focus exclusively on Microsoft's clustering technology, and specifically on Windows Server 2008. 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 (Microsoft Cluster 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. For more information, see the **Enterprise Deployment - Virtualization Considerations for Laserfiche** white paper on the Laserfiche Support Site.

  Note: The paper will discuss "virtual servers," though the term "virtual" here does not refer to a virtual machine (which runs on top of one or more physical machines). Rather the "virtual servers" discussed in this paper refer to virtual interfaces created by clustering technology.

- 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).

Overview of Microsoft Cluster Server (MSCS)

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 failover together to a different node. **Dependencies** are relationships between resources that ensure the cluster brings resources online in the correct order during failover.

Clusters require **shared storage**, meaning application data (e.g., database files, repository volumes, audit logs) is not stored on the nodes, but rather on an external storage device that all nodes can access (e.g., a storage area network or SAN). Although the storage device is shared, only the active node can read from or write to it at any given time. Each storage device is referred to as a **physical disk resource**.

The **quorum** is the cluster's configuration database, which, among other things, tells the cluster which node is currently active and which nodes are in standby. The quorum is stored on a shared storage device.

MSCS enables the creation of **virtual servers**, which are virtual interfaces comprised of configuration information and resources, such as physical devices, IP addresses, network names, and applications. Unlike physical servers, virtual servers are not tied to a specific physical machine, as they can failover from one to another.

Microsoft SQL Server

Microsoft SQL Server provides support for failover clustering, which can be utilized by all Laserfiche products that require a database. You should install SQL Server after you have created and configured all cluster nodes. For detailed information on planning and installing SQL Server support for failover clustering, refer to SQL Server documentation.

Note: Support for failover clusters is dependent on the version and edition of SQL Server you are using. In addition, some editions can use more nodes than others. Refer to SQL Server documentation for more information.

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.

After you have implemented SQL Server support for failover clustering, make the Laserfiche Server resource dependent on the SQL Server resource.

> **Note:** SQL Server has various other technologies that provide high availability, such as database mirroring, log shipping, and replication. For more information, refer to SQL Server documentation.

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). A physical disk resource should be created for the hard disk on which the Laserfiche repository folder will reside, and the Laserfiche Server resource should be dependent on this 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. It is therefore not strictly necessary to include the hard disks on which volume data resides as a part of a failover cluster. However, if document access is mission-critical, you should incorporate volumes in a failover cluster. If you do so, we recommend that you do not make 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 Licensing

Your license type determines the number of Laserfiche Servers you can install and run. For Rio, there is no limit, so you can use Laserfiche License Manager 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 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).

> **Note:** Some Server features, like public portal, must be purchased for each individual Laserfiche Server, meaning you may need to buy multiple instances of the feature to support all nodes in a cluster. For example, though you can create an unlimited number of Laserfiche Servers with Rio, you must purchase each public portal feature individually. 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 public portal.

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.

Scenario 1: Before Failover

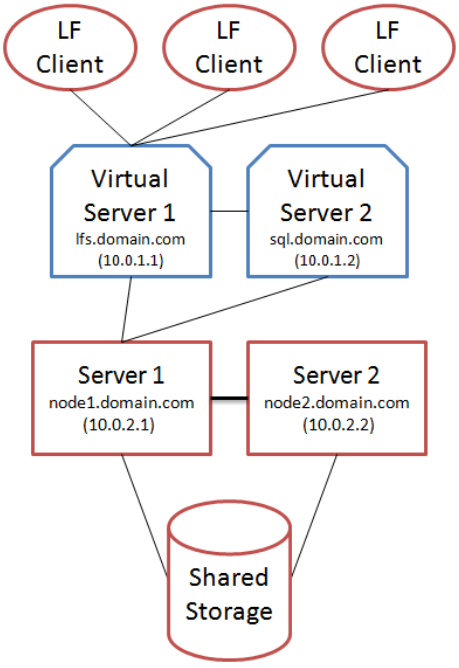

In this scenario, a Laserfiche Server and a Microsoft SQL Server instance are installed on both **Server 1** and **Server 2**. These two machines are clustered together to create:

- **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., its motherboard dies):

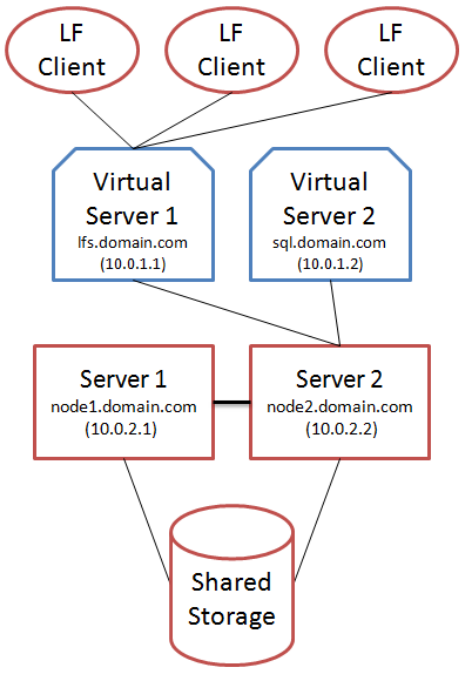

When the cluster determines that **Server 1** is unavailable, the quorum sets **Server 2** as the active node. All requests to **lfs.domain.com** and **sql.domain.com** will now be forwarded to **node2.domain.com**. Since all application data (repository folder, database files, and volumes) are stored 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 an active/passive configuration, where one node is always idle. To increase return on investment (ROI), many

organizations strive to avoid idle servers, in which case the diagram below applies.

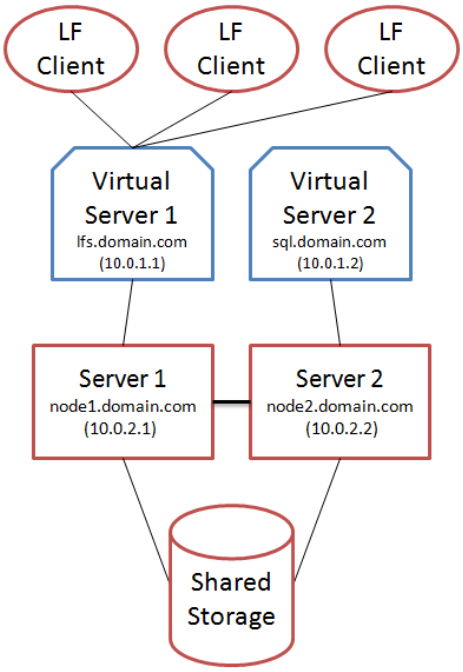

In this scenario, a Laserfiche Server and a Microsoft SQL Server instance are still installed on *both* **Server 1** and **Server 2**. However, before failover, the virtual servers each run on a separate node: **lfs.domain.com** forwards to **node1.domain.com** and **sql.domain.com** forwards to **node2.domain.com**.

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.

Note: 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.

Configuring your Laserfiche Server

After you have created and configured your cluster, install the Laserfiche Server and the Laserfiche Full-Text Indexing and Search (LFFTS) service on each node. Using MSCS' Cluster Administrator, configure both the Laserfiche Server and the LFFTS to run in the cluster using the following settings:

- **Virtual server**: When configuring the Laserfiche Server, choose to use the existing virtual server called **Cluster Group**.
- **Application cluster resource**: An application cluster resource is used to manage the settings of the Laserfiche Server. A **Generic Service** resource should be created for the Laserfiche Server and the LFFTS.
- **Service name**: The name of the service used by the Laserfiche Server is **LFS**. The service used by the search engine is **Laserfiche Full-Text Indexing and Search Service**.
- **SQL Connection String:** On each node, make sure there is a **ConnectionString** string value 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 value should be something similar to either:
  - **Driver={SQL Native Client};SERVER={*MSSQLServerName*};Trusted_Connection=yes;DATABASE={*DATABASENAME*};**
  - **Driver={SQL Native Client};SERVER={*MSSQLServerName*};UID={*SqlLoginName*};PWD={*SqlPassword*};DATABASE={*DATABASENAME*};**

  With an Oracle database, the value will look similar to:
  - **Driver={Oracle in OraClient10g_home1};DBQ=*ORACLESERVERNAME*;UID=*schemaName*;PWD=*schemaPassword*;**
- **Registry replication key**: The registry branch that keeps track of Laserfiche repository settings must be replicated on each node. This branch is the following: **HKEY_LOCAL_MACHINE\SOFTWARE\Laserfiche\Engine**.

- **Resource dependencies**: Establishing a dependency relationship ensures that the Laserfiche Server will always have access to the resources that it requires. The Laserfiche Server should be dependent on the following resources: SQL Server, physical disk resource(s) for the hard disks on which the repository and SQL database files reside, cluster IP, and cluster name. Additionally, the LFFTS must be dependent on the Laserfiche Server.
- **Restart policy**: The restart 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.

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

1. From Windows Server 2008, start **Failover Cluster Management**.
2. Select your previously-created server cluster.
3. From the **Action** menu, select **Configure a Service or Application**. 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 the cluster name for the service that the clients will use to access the service. Enter the appropriate IP address for the network. Click **Next**.
7. In the **Select Storage** step, choose the cluster disk you have created. 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 create a physical disk resource for the Laserfiche repository, expand **Services and Applications** and select **LFS**.
12. Right-click **LFS** and select **Add Storage**.
13. Choose the disk you want to add and click **OK**.
14. Click **OK**.
15. To add your SQL server, again select **LFS**.
16. From the **Action** menu, select **Add a resource** and then **Generic Service**. This will open the **New Resource** wizard.
17. In the **Select Service** step, select your SQL server. Click **Next**, and confirm you want to add the service by clicking **Next** again.
18. In the main **Failover Cluster Management** window, again select **LFS** under **Services and Applications**. The LFFTS should appear under the **Other Resources** section. Select this resource and select **Properties** from the **Action** menu.
19. In the **Dependencies** tab, add the IP address and your SQL Server as dependencies. Click **OK** to save changes and close the dialog.
20. To view your dependency report, select **LFS**, open the **Action** menu, and select **Show Dependency Report**. The dependency report will open in your browser, and should show that **LFS** depends on the physical disk, the name, and the IP address you selected. Close the browser to close the report.
21. To add your search service, select **LFS**, open the **Action** menu and select **Add a resource** and then **Generic Service**. This will open the **New Resource** wizard.
22. In the **Select Service** step, select the LFFTS. Click **Next**, and confirm you want to add the service by clicking **Next** again.

23. With **LFS** selected, the LFFTS should appear under the **Other Resources** section. Select this resource and then select **Properties** from the **Action** menu.
24. In the **Dependencies** tab for the LFFTS service, designate **Laserfiche Server 8.x** as a dependent service.
25. In the **Registry Replication** tab, add **HKEY_LOCAL_MACHINE\SOFTWARE\Laserfiche\Engine**.
26. Click **OK** to save your changes.

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.
- 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 one of the cluster's physical disk resources.

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

March 2012

Author: Jonathan Powers and Constance Anderson

Technical Editor: Jeanie Conner

Editor: Tammy Kaehler

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