the professional labeling software

www.nicelabel.com, info@nicelabel.com

NiceLabel

Clustering NiceLabel

White Paper

Version 20090219-01

www.nicelabel.com

Table of Contents

1 Executive Summary

From small local production companies to global enterprises, all companies today rely on running business services on computers. The modern computers of the era are much more reliable than they were a few years ago, but the increasing demands on computer equipment adversely affect their availability. If a business-critical computer fails, the application is offline and this directly translates into lost revenue.

This White Paper explains how clustering can help decrease or eliminate service downtime and the benefits of running NiceLabel products on clustered servers. It provides step-by-step configuration instructions for high-availability printers, NiceLabel Suite Network edition, and NiceWatch (or NiceWatch Server).

2 The Reasons for Clustering

Today, an increasing number of companies depend on the business processes running on their servers. International companies operate with customers all over the world, accessing their services at any time, on a twenty-four-hour, seven-days-a-week basis. If such a server fails, the service is unavailable and every second of its downtime means lost revenue.

One way to minimize or avoid downtime is to run business-critical services on clustered computers. Clusters cannot prevent computer faults, but they limit the service downtime by providing a host on which this service may recover.

Large enterprises are not the only ones experiencing downtime problems; the demand for business-critical computing is equally present in smaller companies. If, for example, a production procedure running on a server stops, the planned amount of products can not be produced, and the company fails to deliver the order to its customer on time.

Another important area is system management. There is increasing demand for shorter downtime during the performance of routine computer maintenance tasks. The users need and demand access to the system's services even while administrators perform maintenance tasks.

Apart from improving availability, clustering technologies can also improve the performance of the server. A typical example is a company server with a business application running out of system resources (such as processor cycles, available memory and disk space). By forming a clustered group of servers and running the application on this cluster, you can significantly improve the application's performance.

The financial aspect of a clustering solution is clear: if the potential revenue loss due to downtime or inadequate performance is larger than the equipment cost increase represented by the clustering solution, the clustering solution pays off.

3 What Is Clustering?

Clustering is a method of grouping two or more servers with the goal of increasing availability or performance of such a group, compared to a single server. Usually, the cost of such a group is lower than the cost of a single server offering same level of reliability and performance.

A cluster is a group of servers that are working in unison and are function as one virtual server. Usually, the servers are connected to each other with fast network connections.

3.1 Term Definitions

Certain aspects of clustering are not common computer knowledge. For the sake of understanding this document, please familiarize yourself with the following terms:

- **Node:** a server working as a part of the cluster. The Microsoft Cluster Service on Windows 2003 Server operating systems supports two to eight nodes in a cluster (depending of the Windows 2003 Server edition).
- **Cluster:** a group of nodes working together as a virtual server.
- **Active/Passive Cluster:** in this configuration, the active node offers resources or services to the clients, while the passive node remains in standby mode, waiting for the active node to fail. Nodes communicate through a dedicated network connection. If the communication is stopped, the passive node assumes that the active node has failed and takes control of cluster resources or services, provided to the clients.
- **Active/Active Cluster:** in this configuration, all nodes in the cluster can run applications and offer resources to clients. At the same time, each node is acting as a backup to other nodes and can dynamically assume the other node's role.
- **Heartbeat:** node-to-node communication through a dedicated networking interface is used by the nodes to query the status of other connected nodes. Based on the communication status, the passive nodes may decide to take control of the cluster.
- **Failover:** the process where the standby node assumes control of resources.
- **Failback:** the process where the original node retakes control of resources from the passive node after a failover.
- **Cluster Quorum:** Is a common resource for all nodes in a cluster. A quorum provides consistent information about the cluster state. It holds the cluster configuration information such as which servers are part of the cluster, which resources are installed in the cluster, and the status of those resources (for example, online or offline).

3.2 Clustering Types

There are two main reasons for clustering:

- To provide failover capability and increase availability of services and applications
- To provide parallel calculating power and better application performance in cases of high-performance computing (HPC) clusters.

Based on these two reasons, the two most common clustering implementations are Failover Clustering and High-Performance Computing (HPC) Clustering.

3.2.1 Failover Clustering

In this type of clustering implementation, the cluster offers failover services, which means that if a failure occurs on any cluster member server (cluster node), the resources that the

failing server was hosting will automatically restart themselves on another server that is a member of the same cluster. This process of resource moving from one server to another is called *failover*. A typical example is the Microsoft Cluster Service (MSCS), available on Microsoft Windows Server operating systems. Please refer to Microsoft Web page for a detailed overview of clustering support available in different Windows Server operating systems and their editions.

3.2.2 Load Balancing Clustering

Load-balancing clusters are used to distribute the client workload (for example, access to Web pages) to several servers forming a cluster. Such a cluster of computers is sometimes referred to as a server farm. The workload is divided by one or several computers forming the load balancing at the front end. While the primary goal of this clustering technology is to increase performance, such a cluster also offers the high availability feature. There are many commercial load balancing clusters available today, for example Network Load Balancing service in the Microsoft Windows Server family of operating systems or Linux Virtual Server, the free software package for the Linux operating system.

3.2.3 High-Performance Computing (HPC) Clustering

High-performance computing (HPC) cluster technology is used in supercomputers. Multiple computers are working in parallel and acting as one computer to achieve very high calculating power. Unlike failover clusters, HPC clusters do not require shared disks as they only do calculations and do not need to read or write to a common storage. If one of the machines in a HPC-cluster fails, the only consequence is reduced cluster calculating performance, but none of the functionality is lost. Microsoft Windows Compute Cluster Server 2003 (CCS) is an example of such cluster technology software.

3.3 Microsoft Cluster Service

There are several failover clustering implementations offered by different vendors. Currently, NiceLabel software is using the Microsoft implementation of clustering, called Microsoft Cluster Service (MSCS). Therefore, this document will focus solely on the MSCS clustering architecture.

In a typical Microsoft Cluster Service example, you can configure two servers in a cluster, and present them to the outside world as a Virtual Server. Both nodes have access to external shared storage, which can be SCSI- or Fiber Channel-attached storage. If one of the servers fails, the remaining server is used to ensure that the virtual server and the services it offers are still available to clients, thereby providing users transparent access to the cluster.

Every service or group of services running in the cluster uses cluster node resources. Every service or group of services has its own assigned disk (which is shared with the other failover cluster nodes), it has its own IP Address and its own Network name. All of the resources that a clustered service uses are called a Resource Group. The Resource group contains the basic resources that every service needs: a Disk Drive, an IP Address, a Network Name, and the service itself. These resources together form a virtual server which can be moved from one server to another in a matter of seconds (Failover) without any dependence on a specific server. Users access this virtual server like any other physical server, using its host name or IP address.

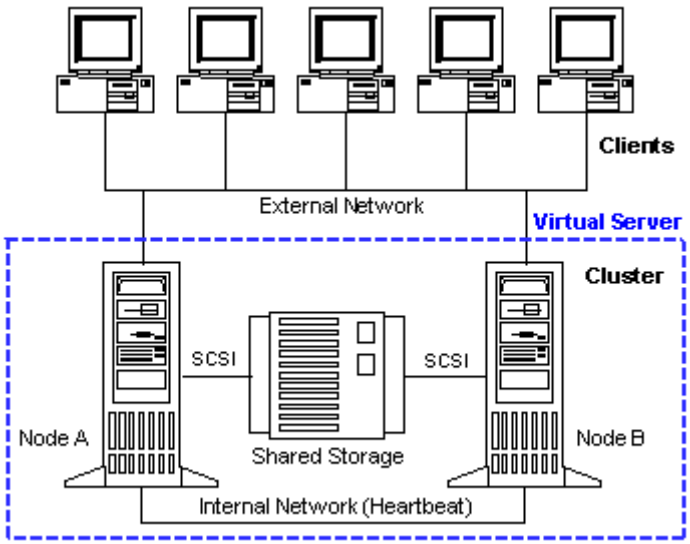

Figure 1: Microsoft Cluster diagram

Note: To learn more about the Microsoft Cluster Service, please see the literature referenced in section Literature of this document or visit the Microsoft Web site.

4 NiceDrivers and NiceLabel Applications on Cluster

The following NiceLabel resources and products can be installed on a cluster to increase their availability:

- NiceDrivers
- NiceLabel Suite (Network edition) labeling application
- NiceWatch and NiceWatch Server middleware application

4.1 NiceDriver on a Cluster

To achieve a highly available printer, you can use Microsoft Cluster Service functionality and make sure the printer is always accessible.

If you connect the printer on one node in the cluster, the printer will not be available, if that node fails. When using the printer in a cluster environment, make sure that the printer is accessible directly on the network using its own network interface. It does not matter which node in the cluster prints labels, the printer is always available through its network interface, and is not shared from a node.

A clustered printer solution setup consists of the following main phases:

1. Make sure that you have an installed and fully functional cluster (in our example called ClusterTT)
2. Install the Port Monitor (Microsoft does not support the use of third-party port monitors and recommends that you use the Standard TCP/IP port monitor)
3. In Microsoft Cluster Administrator create a Group with a Disk Resource, using the wizard.

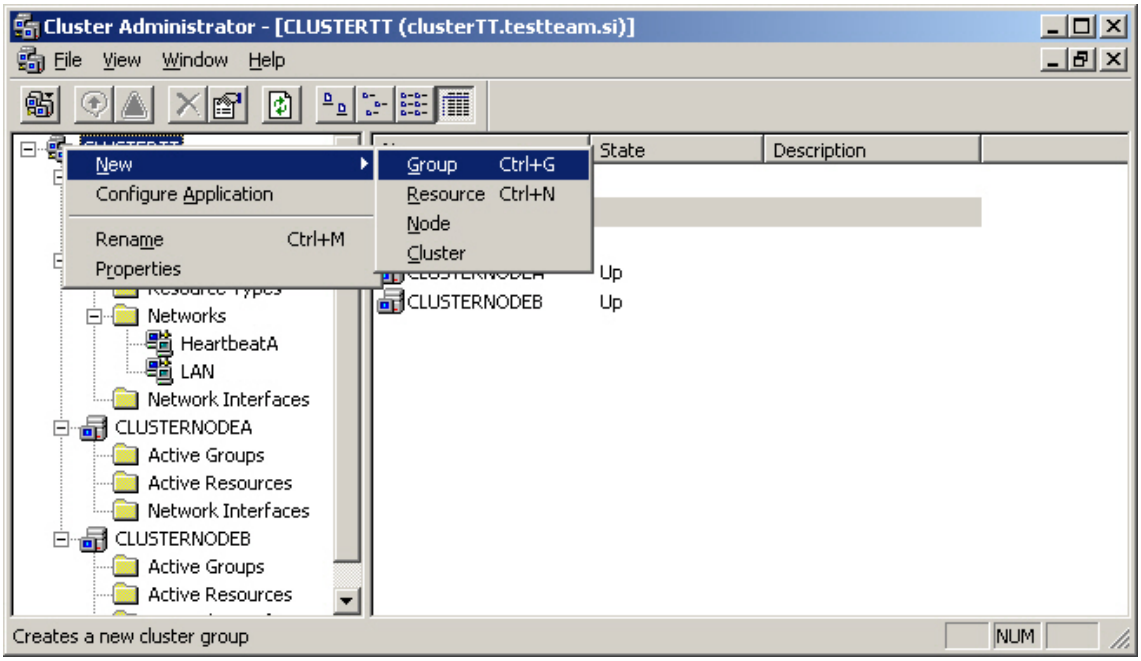

Figure 2: Adding a new group

4. Configure the Print Spooler using the *Cluster Application Wizard*.

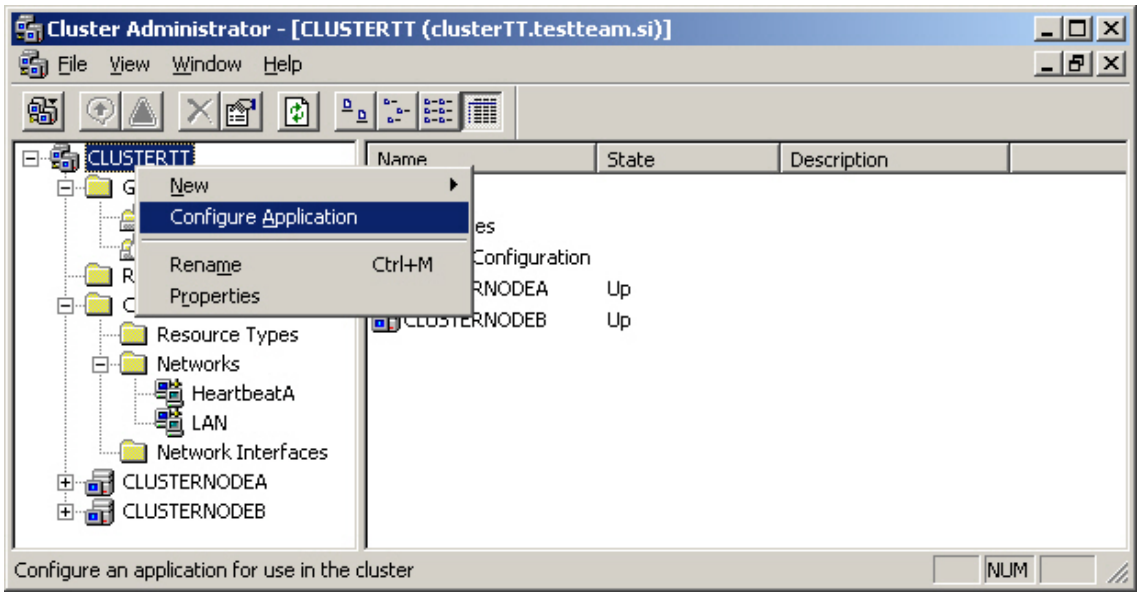

Figure 3: Configure the Print Spooler as an application in the cluster

5. Click **Use an existing resource group** and then select the group you created previously.
6. For the resource group name, provide a name that accurately represents the group, such as "Group Spooler". This name is used in Cluster Administrator for administrative purposes.
7. At the *Virtual Server Access Information* window enter:
    - The NetBIOS name that clients can use to connect to this virtual server (for example, \\ClusterTT\CLPRINT).
    - The IP address that clients can use to connect to this virtual print server.

    And click **Next**.
8. At the *Create a Resource for My Application* window, leave the default selection and click **Next**.
9. At the *Application Resource Type* window, select the **Print Spooler** resource type and, when prompted, name the Spooler resource and enter a meaningful description. This name is used only in the Cluster Administrator for administrative purposes.
10. Click the **Advanced Properties** button to open the resource properties window. On the *Dependencies* tab, click **Modify**. Double-click the **Physical Disk** and the **Network Name** resources that you just created.

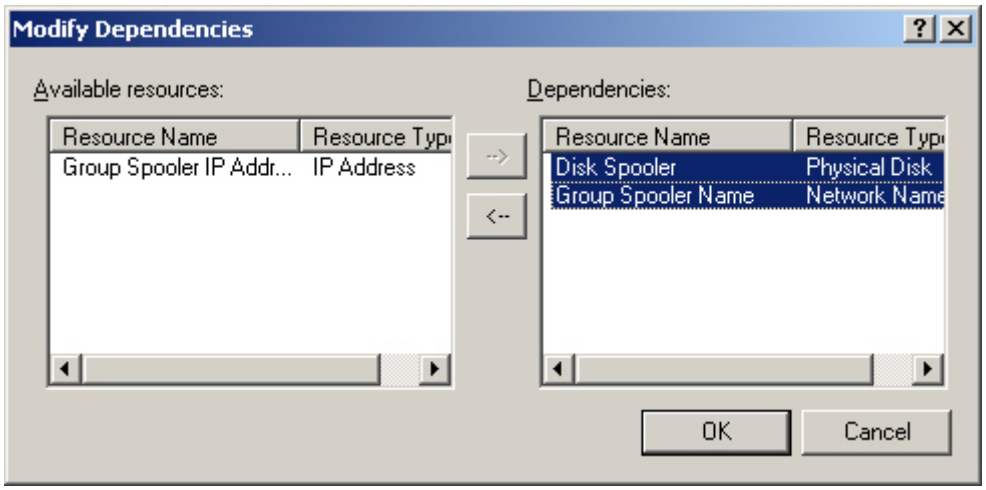

Figure 4: Adding dependencies

11. On the *Print Spooler Parameters* window, type the path and folder name for the spooler files location. This is where the SHD and SPL files are created when the Spooler service runs. Unless necessary, leave the Job completion timeout as 160 seconds and click **Next**.
12. At the Cluster Application Wizard final window, you can review the configured virtual server and its resources. Click **Finish** to complete the wizard.
13. Now bring the virtual server online by right-clicking the new group and selecting **Bring Online**.

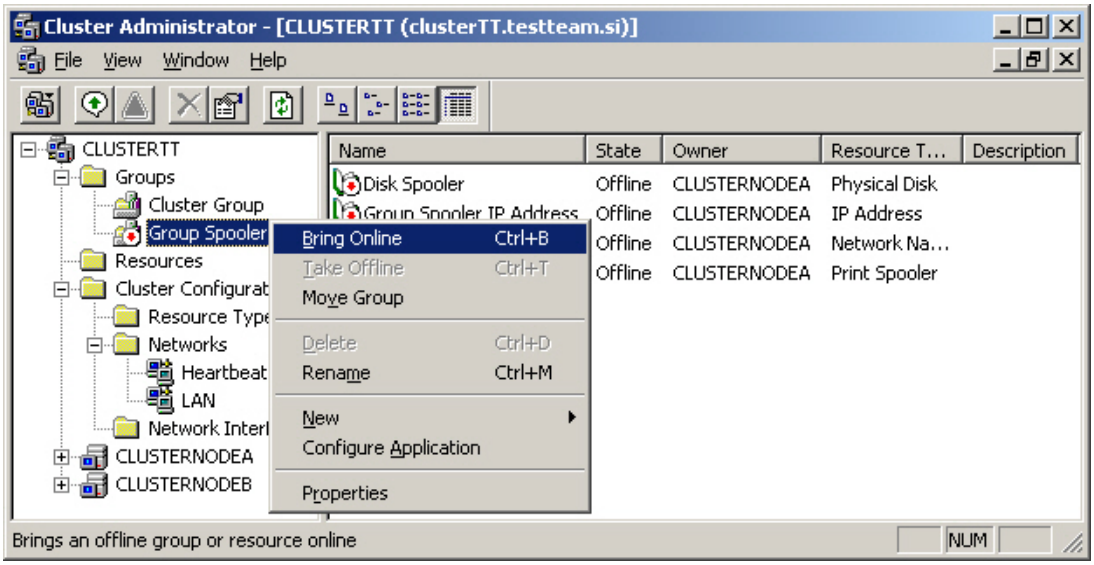

Figure 5: Bringing the group online

14. Now right-click the first node and select **Start Cluster Service**. Wait for a few seconds and click the **Active resources** folder. All resources should be returned to this node (provided you configured the Failback setting in the group properties dialog).

Note: If you are setting up an active/active print server, you need to create one group for each node and you want to set each spooler group to have a different preferred owner. You cannot have multiple Spooler resources in the same group. An active/active print server configuration is one in which there are multiple nodes in the cluster that are processing print jobs for clients with multiple spoolers. This could include as many as two to four nodes that are actively handling requests.

15. To install the NiceDriver driver on your virtual server, on any cluster node, click **Start** and **Run** and in the Open box, type \\NetworkName of the virtual server on which the print spooler resource depends.

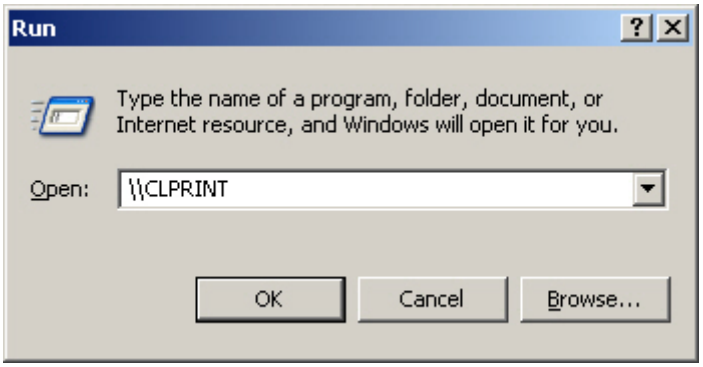

Figure 6: Connect to virtual server

16. Proceed through the steps and complete the Add Printer Wizard. The information on the *Windows Printer Test Page* describes the printer configuration.

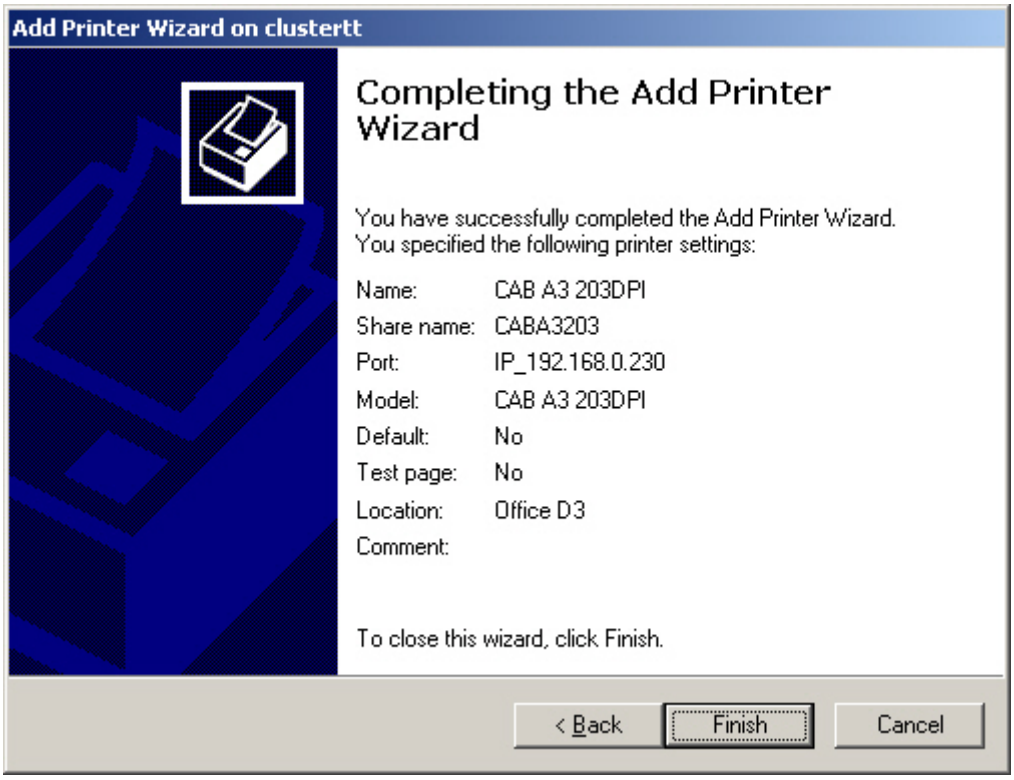

Figure 7: Completing the Add Printer Wizard

4.2 NiceLabel Network on a Cluster

This solution is using a file share on a functional MS failover cluster to store the System folder of a NiceLabel Suite Network installation. Using this approach, we provide NiceLabel workstations with uninterrupted access to the NL Primary server by storing important NL files (licensing, logging, security settings, etc.) on the clustered file share.

In this case it is not necessary to have a secondary NiceLabel server in your network.

The main steps for implementing a NiceLabel Network Suite on a set of clustered servers are:

1. Make sure that you have an installed and fully functional cluster (in our example called ClusterTT).
2. Create a folder on a common disk (for example, F:\Network).
3. Install NiceLabel Suite Network (server part) edition into this folder. For NiceLabel Suite Network installation details, see the NiceLabel Quick Start Guide.
4. Create a new File Share resource in the new folder.
5. Install NiceLabel Suite Network on the workstation. During the installation procedure, provide the network address of the cluster File Share, created in the previous step. Make sure you use the UNC naming convention, providing the Cluster name, followed by the File Share name (for example \\ClusterTT\File_Share)
6. Activate NiceLabel Suite Network edition using the standard procedure.

Known issue:

By enabling logging on the NiceLabel workstation, the log will be written to the File Share on the cluster. This means that the NiceLabel log is always up to date regardless of the status of NiceLabel Suite Network server. However, there might be a slight interruption in

logging during failover (and consequently, loss of logging data), depending on the hardware specifics of your clustering solution (for example, failover of Fiber optical shared storage in general takes less time than failover on its SCSI counterpart).

4.3 NiceWatch Server on a Cluster

In this solution, Microsoft failover clustering is monitoring the NiceWatch service, installed on both nodes. Additionally, files important for NiceWatch Server operation (triggers configuration, label files, etc.), are stored on a clustered File Share and thus are monitored by the cluster. Each node is using its own NiceWatch Server license. Cluster is operating in an Active-Passive failover mode. If NiceWatch Server on one fails, clustering service starts NiceWatch service on the other node, which in turn connects to the NiceWatch files on the clustered File Share and continues the operation.

4.3.1 Clustered NiceWatch using Enterprise Print Manager (EPM) – Recommended method

Main steps are listed below for implementing a fault-tolerant NiceWatch solution using the functionality of NiceLabel Enterprise Print Manager:

1. Make sure that you have an installed and fully functional cluster (in our example called ClusterTT)
2. Create a folder on a common disk (for example, F:\Network)
3. On the new folder, create a new File Share resource.
4. Install NiceWatch on both nodes in the usual way. License them by pointing to the EPM server.
5. Configure NiceWatch on both nodes to use global settings (Tools -> Options -> Settings: Global).
6. Start NiceWatch on the first node and register it as a service.
7. Set the NiceWatch service for manual start.
8. Repeat the steps 6. and 7. on the second node.
9. Stop the NiceWatch service on the second node.
10. In the NiceWatch installation on the first node, create your desired triggers and save them into the trigger configuration file (for example TEST.MIS) to the clustered File Share.
11. Using Microsoft Cluster Administrator, create a new generic service resource using the Resource Wizard. In the process, provide the necessary command and location of the .MIS file. For example: net start EPNWS5 / \\ClusterTT\File_Share\TEST.MIS)
12. Enable the *Set Cluster Usage* (Tools menu) function in order to share the program settings of all clustered instances of NiceWatch. When a shared path is chosen by the first instance, the settings from that instance will be copied to the location. When it is selected on other instances, the settings will be detected and adopted by all subsequent clustered instances of NiceWatch.
13. For a completely fault-tolerant solution, you can point your NiceWatch installation to a clustered printer driver.

Important: Using the above implementation, all settings are synchronized automatically between both NiceWatch instances, thus effectively eliminating the limitation from the previous case.

Known issues:

1. Two NiceWatch Enterprise licenses are needed to activate the NiceWatch software. Only one instance of the NiceWatch Server software is in use at any given time. However, when the failover happens, the second NiceWatch Server needs its own

license as the license of the first NiceWatch Server remains in use for some time after the node failed.

2. Two NiceWatch Enterprise licenses are needed.
3. COM Port Trigger cannot be used because of physical limitations (each node has its own COM Port connector.

4.3.2 Clustered NiceWatch Enterprise using two separate NiceLabel software key licenses

Main steps are listed below for implementing a fault-tolerant NiceWatch Server solution using separate licenses:

1. Make sure that you have an installed and fully functional cluster (in our example called ClusterTT)
2. Create a folder on a common disk (for example, F:\Network)
3. On the new folder, create a new File Share resource.
4. Install and license NiceWatch Server on both nodes in the usual way, using a software key license for each node.
5. Start NiceWatch Server on the first node and register it as a service.
6. Set the NiceWatch service for manual start.
7. Repeat the steps 5. and 6. on the second node.
8. Stop the NiceWatch service on the second node.
9. In the NiceWatch Server installation on the first node, create your desired triggers and save them into the trigger configuration file (for example TEST.MIS) to the clustered File Share.
10. Using Microsoft Cluster Administrator, create a new generic service resource using the Resource Wizard. In the process, provide the necessary command and location of the .MIS file. For example: net start EPNWS5 / \\ClusterTT\File_Share\TEST.MIS)
11. For completely fault tolerant solution, you can point your NiceWatch installation to a clustered printer driver.

Known issues:

1. Using the above implementation, no settings are synchronized automatically between both NiceWatch instances. If synchronization is needed, it has to be done manually.
2. Two NiceWatch licenses are needed.
3. COM Port Trigger cannot be used because of physical limitations (each node has its own COM Port connector.

4.3.3 Clustered Enterprise Print Manager (EPM) – Recommended Method

The Enterprise Print Manager is a web application, which often finds its way onto network load balancing clusters along side pre-existing web services. The installation of EPM on a NLB cluster is fairly straight-forward, but does carry some specific recommendations regarding the setup process.

Installation

1. Install EPM on cluster Node A (primary node)
2. License EPM on the primary node
3. Open the page http://<cluster name>/EPM
4. Install EPM on Node B (secondary node)

Note: If EPM is installed on the secondary node before the primary node installation is licensed, the secondary installation may be unable to detect its license, even though the primary node was installed. This is solved by restarting the IIS service on the secondary node.

Version upgrade

1. Uninstall EPM from Node B (secondary node)
2. When prompted whether you wish to keep the existing database, make sure that the choice “Yes” is selected.
3. Perform the normal upgrade procedure on Node A (primary node)
4. The upgrade process will automatically remove the old version and install the new version of EPM
5. Install EPM on node B (secondary node) and subsequent nodes, if more are present.

Note: None of the nodes require a reboot and the performance of other content on the server cluster should not be affected in any way.

License upgrade

1. Upgrades of licenses are performed on the primary node in the same fashion as in non-cluster EPM installations.

Known upgrade issues:

1. If the labeling client is printing during the **version upgrade** process, XML files with print job information that are sent to the EPM dynamically will queue. Once one of the EPM nodes is accessible, the XML files will be sent to the EPM.
2. Labels and other files stored in the Storage Server on the EPM may be inaccessible during the **version upgrade** process.

5 NiceLabel Literature

To learn more about NiceLabel software or Microsoft Cluster Service (MSCS), see the following list of literature and documentation. You will find the documents on your NiceLabel CD ROM, at **www.nicelabel.com** or at **www.microsoft.com**:

NiceLabel Pro:	▪ NiceLabel Quick Start Guide ▪ NiceLabel Pro User Guide ▪ NiceLabel Online Help
NiceLabel Print Center:	▪ NiceLabel Print Center User Guide
Microsoft Cluster Service (MSCS)	▪ Guide to Creating and Configuring a Server Cluster under Windows Server 2003 ▪ Creating and Configuring a Highly Available Print Server under Microsoft Windows Server 2003 Using a Server Cluster White Paper

6 Glossary

Node	A server working as a part of a cluster. Microsoft Cluster Service in the Windows 2003 Server operating systems supports from two to eight nodes as part of a cluster (depending of the Windows 2003 Server edition).
Cluster	A group of nodes working together as a virtual server.
Active/Passive Cluster	In this configuration, the active node is offering resources or services to the clients, while the passive node remains in standby mode, waiting for the active node to fail. Nodes are communicating through a dedicated network connection. If the communication is stopped, the passive node assumes the active has failed and takes control of cluster resources or services, offered to the clients.
Active/Active Cluster	In this configuration, all nodes in the cluster can run applications and offer resources to clients. At the same time, each node is acting as a backup to other nodes and can dynamically assume the other's role.
Heartbeat	Node to node message, sent through a dedicated networking interface. It is used by the nodes to query the status of remaining node(s). If this message is not received in regular interval, the passive node(s) decide to take control of the cluster.
Failover	The process of changing the control of resources to the standby node.
Failback	The process of changing the control of resources back to the original node.
Quorum	Is a common resource, available to all nodes in a cluster. Quorum provides consistent information about the cluster state. It holds the quorum log and cluster configuration information such as which servers are part of the cluster, what resources are installed in the cluster, and what state those resources are in (for example, online or offline).
Downtime	Period of time that a computer (or clustered group of computers) is unavailable, usually as a result of a failure or routine maintenance.

7 Online Support

You can find the latest builds, updates, workarounds for problems and Frequently Asked Questions (FAQ) under the Support section on our Web site at www.nicelabel.com. If you cannot solve the problem on your own, please contact your local vendor or representative offices listed in the topic Contact Information.

For more information please refer to:

- Support FAQ: http://www.nicelabel.com/Support/FAQ
- NiceLabel Learning Center: http://www.nicelabel.com/Learning-center
- NiceLabel Forums: http://forums.nicelabel.com/

8 Contact Information

Head Office
Euro Plus d.o.o.
Ulica Lojzeta Hrovata 4c
SI-4000 Kranj, Slovenia
+386 4 280 50 00 Tel
+386 4 233 11 48 Fax

www.europlus.si
info@europlus.si

North American Office
Niceware International, LLC
10437 Innovation Drive, Suite 147
Milwaukee, WI 53226
+1 414 476 6423 Tel
+1 414 476 7955 Fax

www.nicewareintl.com
sales@nicewareintl.com
support@nicewareintl.com

NiceLabel France
Cobarsoft SARL
6 rue Salengro
ZI hippodrome
32000 Auch
France
+33 (0) 562 709 201 Tel
+33 (0) 562 708 004 Fax

support@nicelabel.fr
www.nicelabel.fr

German Office
NiceLabel Germany GmbH
Liebknechtstr. 29
63179 Obertshausen
Germany
+49 (0)6104 405 400 Tel
+49 (0)6104 405 4020 Fax

info@nicelabel.de
www.nicelabel.de