Load Balancing
VisNetic® MailServer

©2003 Deerfield Communications, Inc. Portions © Multiple Sources. All rights reserved.

Load Balancing VisNetic® MailServer

Why implement load balancing in VisNetic MailServer?

There are a number of reasons that a company may wish to load balance traffic across multiple VisNetic MailServers, these reasons include, but are not limited to:

- Traffic volume exceeds the capacity of a single server
- Reducing failure recovery times
- Improving the reliability of connections
- Reducing server response delays by spreading the load

Implementing a simple solution which is reliable and which provides the same level of performance as more complex solutions is preferred. However there are times when a more complex solution offers distinct advantages in reliability, scalability and performance and in these cases, the more complex solution should be deployed.

The overall assessment of complexity is subjective but should include implementation costs, standards reliance, other affected systems and the resources required to perform operations.

The following chart may be helpful in determining whether or not a load balanced VisNetic MailServer cluster is warranted in your environment. This diagram addresses server traffic volume and server response delays; it does not address failure recovery times or the reliability of connections.

VisNetic MailServer Version & Hardware Recommendations

Users	Dom	Ver	Minimum SW/HW
6 - 50	1	Plus	2K, XP CPU PIII 300Mhz, RAM 256MB
6 - 50	1+	Pro	2K, XP, 2K3 CPU PIII 300Mhz, RAM 256MB
50 - 2k	1+	Pro	2K, XP, 2K3 CPU PIII 300Mhz, RAM 512MB
2k – 5k	1+	Pro	2K, XP, 2K3 CPU PIII, RAM 512 MB, ODBC or HSMM
5k – 30k	1+	Pro	2K, XP, 2K3 Dual CPU PIII, RAM 512 MB, ODBC
30k+	1+	Pro	Used Load Balanced Servers

What is load balancing?

Load Balancing is the dynamic sharing of load between servers, possibly on a per session basis. Almost always using 'sticky' paths, where traffic will continue to flow via a particular route for the duration of a session.

There is little question that load balanced solutions will be more likely to provide a more even spread of load between servers. However, understanding if this is useful, i.e. actually adds value, is a little more pertinent, and should be assessed by understanding if the solution is technically more robust and whether increased performance is realized. There is also a case to be made for scalability; what fits now should ideally be able to grow bigger without radical change.

©2003 Deerfield Communications, Inc. Portions © Multiple Sources. All rights reserved.

Note: It is important to note that the Professional versions of VisNetic MailServer support an unlimited number of Domains, however the physical capability of each server varies from machine to machine. There is a processing limit that is reached with some machines around 3000 Domains (Users are not affected). At this point you might wish to load balance VisNetic MailServer regardless of the total number of configured Users.

Implementing Load Balanced VisNetic MailServers

There are two important aspects of load balancing as it applies to VisNetic MailServer, the first is the routing of traffic to the server cluster and the second is the configuration of VisNetic MailSever within the cluster. We will first address the issue of routing traffic.

There are a number of ways to load balance traffic. For the purpose of this document, we will limit ourselves to the three most common methods:

- Round-robin Domain Name System (DNS)
- Load balancing Switches and Routers
- Windows 2000 Advanced Server Network Load Balancing

Round-robin DNS involves setting up your site's DNS server to return the set of all the IP addresses of the servers in the cluster in a different order on each successive request. The client typically forwards the request to the first IP address in the list of IP addresses returned. Consequently, the request is directed to a different server in the cluster, and the traffic is distributed across the servers. We'll discuss this later in more depth.

Load balancing switches redirect TCP requests to servers in a cluster. These implementations provide a highly scalable, interoperable solution that is also very reliable. They typically present a common virtual IP address to the requesting clients, and then forward the requests to an available server.

Windows® 2000 Advanced Server Network Load Balancing (NLB) distributes the IP traffic across multiple servers that provide TCP/IP services. Network Load Balancing presents a common virtual IP address for the entire cluster and transparently partitions client requests across the servers in the cluster. Network Load Balancing provides high availability and high scalability to the Internet applications running on the servers.

We will now look at each of these load-balancing solutions in detail.

Round-Robin DNS

Before understanding round-robin DNS in detail, let's take a brief look at DNS and how it functions.

How DNS works

When an application or a user requests a resource on the Internet, they can use the Internet name of the resource. For example, an application may use the name “www.deerfield.com” instead of an IP address. However, to locate the resource on the

©2003 Deerfield Communications, Inc. Portions © Multiple Sources. All rights reserved.

Internet, the application needs to know the IP address of the resource. Domain Name System (DNS) is the standard method of translating the Internet name of a resource to its IP address. The method of translating names to their corresponding IP addresses is called *name resolution*.

For example, a query for "www.deerfield.com" will look for an "address resource record" for the specified name. The address resource record that is returned contains an ordered set of IP addresses for the resource. When the query is issued:

The DNS client tries to resolve the query locally using information cached from a previous query.

If the query cannot be resolved locally, the query is passed to the DNS server. The DNS server checks if the query can be resolved locally using the information in the zone configured on the server.

If the DNS server cannot find a corresponding resource record, it queries other DNS servers.

There are three important things to note here. First, names can be resolved in pieces. For example, a server with the authoritative entries for the ".com" domain could contact an authoritative server for the "deerfield.com" domain. (This process could repeat.) Second, individual organizations can maintain DNS servers that are available to anyone on the net. For instance, Deerfield.com itself maintains the authoritative servers for resolving names (such as "deerfield.com", "mail.deerfield.com", and "www.deerfield.com") within their domain. Third, DNS records can be cached in any number of places on the net—they contain information about how long they're valid (like the freshness date on your milk).

How Round-robin DNS works

Given those three traits, it's possible then for a domain's DNS servers to give a different IP address (actually, a different ordering of the set of possible IP addresses—A/B/C, B/C/A, C/A/B, A/B/C, . . .) each time it's queried. Each of the IP addresses points to a logically identical server that's equally capable of handling the request. And because different clients are routed to different machines (with different IP addresses), this gives us a primitive form of load balancing.

The main advantage of round-robin DNS is that it requires no additional hardware—you just set up your DNS server properly, and it works. However, there are several disadvantages that prevent many sites from using round-robin DNS for load balancing:

The caching feature of DNS prevents complete load balancing because not every request that comes in will get its address directly from our DNS server.

You can solve the above issue by disabling caching, but doing so means that every resolution will have to be resolved by your servers, which is expensive and potentially slower for your users.

The DNS server has no way of knowing if one or more of the servers in your cluster is overloaded or out of service. So, the round-robin scheme will send traffic to all servers in

©2003 Deerfield Communications, Inc. Portions © Multiple Sources. All rights reserved.

turn, even if some are overburdened or offline. This is not exactly a wonderful user experience, although a browser user can hit the reload button to try again other applications may not have a user interface and may not retry to connect.

Because of this last issue, round-robin DNS isn't used much (at least not by itself) for large or mission-critical server clusters. But you can use round-robin DNS for load balancing in a server cluster with two or three servers or you can use it to balance load across two or three server clusters, each of which is load-balanced with one of the methods below. (The chances that an entire cluster will fail are quite small.)

Load Balancing Switches

Load balancing switches are hardware Internet scalability solutions that distribute TCP requests across multiple servers.

These switches sit between the connection to the Internet and the server cluster. All requests come to the switch using the same IP address, and then the switch forwards each request to a different server based on various algorithms implemented in the switch. Switches will frequently be able to ping the servers in the cluster to make sure they're still up, and to get an estimate of how busy they are so they can be relatively intelligent about load balancing.

Another common algorithm is to load balance based on the content of the request. Perhaps the IP address of the requestor, or some other information in the request. Using the IP address alone doesn't work well since some ISPs, such as AOL, and companies use proxy servers that change the IP address of all of the requestors that go through the proxy to the same address.

Using a load-balancing switch is much better and more scalable than using round-robin DNS, but switches can be quite expensive—and you'll need multiple switches to avoid making the switch the single point of failure for your entire server cluster. The next solution, Microsoft® Windows® 2000 Advanced Server Network Load Balancing (NLB), is often less expensive than a load-balancing switch and avoids having a single point of failure.

Windows 2000 Advanced Server Network Load Balancing

In Windows 2000 Advanced Server Network Load Balancing, a single cluster IP address addresses all the servers in the cluster. All of the machines are connected in parallel, so each machine sees all of the requests for the cluster. It is relatively easy to put machines into and remove machines from the cluster both manually and programmatically through a series of command-line utilities.

In addition to the common external IP address for the cluster, each server in the cluster will respond to a *dedicated network address* as well. So, each machine responds to two network addresses: a *cluster network address* and a dedicated network address.

Network Load Balancing is implemented using a network driver that is logically placed between the higher-level protocol TCP/IP, and the network adapter of the host. All the cluster hosts receive the incoming traffic. The Network Load Balancing network driver acts as a filter and allows the host to process only a part of the incoming traffic. The

©2003 Deerfield Communications, Inc. Portions © Multiple Sources. All rights reserved.

incoming requests are accepted according to the Network Load Balancing settings for the host.

The benefits of utilizing Network Load Balancing are:

Fault tolerance. Network Load Balancing automatically detects a nonfunctional host and can recover from it. In case of an offline host, the incoming traffic is distributed across the remaining online servers.

Higher scalability. Network Load Balancing supports up to 32 computers in a cluster and can load balance client requests for individual TCP services across the cluster. If you need a cluster of more than 32 computers, you can load balance across NLB clusters of 32 machines or fewer using one of the previous two methods.

Manageability. You can specify load balancing for a single TCP port to customize processing load for a host. You can also block access to certain ports.

Ease of use. Network Load Balancing installs a standard networking driver component and does not need any special hardware to run.

To summarize, it may be necessary, from a reliability and scalability standpoint, to implement a VisNetic MailServer cluster. In order to make effective use of the server cluster, you have to have some sort of traffic load balancing.

The simplest form is round-robin DNS, but it has numerous technical limitations, the worst of which is that it has no way to detect a slow or offline server and will therefore send requests to servers regardless of whether that server can handle them or not. However, round-robin DNS is sometimes successfully used for small, non-critical server clusters and for load balancing across a set of clusters.

Many sites, particularly those that aren't using Microsoft® Windows NT® Server or Windows 2000, use a load-balancing switch to distribute requests. Load-balancing switches do a great job of distributing the load, but they're expensive and represent a single point of failure, so they can reduce reliability and availability of your servers.

If you're running Windows 2000, you'll want to consider Windows 2000 Advanced Server Network Load Balancing. It doesn't add a single point of failure, doesn't require expensive hardware, and it is often less expensive than a set of load-balancing switches. (There is also a version, called Windows Load Balancing Service (WLBS), for those still running Windows NT.)

The VisNetic MailServer Cluster

VisNetic MailServer not Cluster aware, meaning the different servers in the Cluster are unaware of one another’s operating state or even if the other servers are available. Therefore it is recommended that an external process be put into place to programmatically move servers in and out of the Cluster based on the state of the services they are providing and/or their availability.

©2003 Deerfield Communications, Inc. Portions © Multiple Sources. All rights reserved.

As with load balancing network traffic, there are a number or ways to configure the VisNetic MailServer Cluster. For the purposes of this document, we will outline the most common configurations starting with the simplest (single non-balanced server) and ending with the most complex.

(1) Stand Alone Server (no load balancing or failure protection)

MX5
x.x.x.x

Stand Alone Server

This is a typical VisNetic MailServer implementation in the small to medium sized business environment. This configuration provides no failure protection or load balancing. If the server goes offline, the remote sending server retains mail until the retry cycle times out, at which time the sender will receive a delivery failure notification.

Typical DNS configuration:

mail.domain.com = x.x.x.x
mx 5 = mail.domain.com

(2) Primary with Backup Server (no load balancing with failure protection)

MX5 MX10
x.x.x.x y.y.y.y

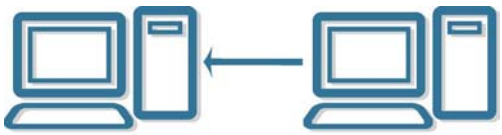

Primary with Backup Server

This is the recommended minimum configuration, which provides failure protection for an offline primary server. The backup server can be hosted locally or offsite by some other entity (usually an ISP).

Delivery to the primary and backup servers is controlled by Mail Exchange (MX) DNS records. The primary server typically has a lower mail server priority, which is defined in the MX DNS entry. Lower priority servers are tried first, if delivery is not successful, the sending server will typically connect to the next lower priority server to attempt delivery and so on until it exhausts the list of MX entries returned by the DNS server.

©2003 Deerfield Communications, Inc. Portions © Multiple Sources. All rights reserved.

In the diagram above, the sending server will first attempt a connection to the primary server, if a connection cannot be made (primary offline) or if delivery fails, the sending server will connect to the backup or secondary server.

The backup server is configured to accept mail for the primary server and to deliver that mail when the primary comes back online. It is recommended that you host your backup or secondary server offsite.

Typical DNS configuration:

mail.domain.com = x.x.x.x
backup.domain.com = y.y.y.y
mx 5 = mail.domain.com
mx 10 = backup.domain.com

(3) Dual Primary Servers (load balancing with no failure protection)

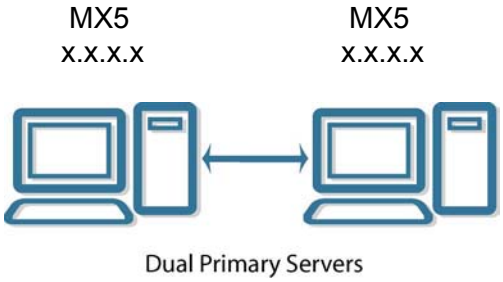

This configuration provides the lowest level of load balancing achievable with VisNetic MailServer but lacks an offsite backup or secondary server to protect against power outages or other failures that could take both servers offline.

The servers share a common publicly addressable IP Address provided by the selected load balancing software or hardware.

In this configuration, the MailServer Configuration, Mail Store and Domain / User Configuration are stored on one of the two primary servers. The two servers share a mapped drive (not the c:\ drive) with both servers running the Mail Services (SMTP, POP3/IMAP, WebAdmin and WebMail). The inbound load to the servers is balanced utilizing one of the methods described earlier in this document. This solution provides the lowest level of scalability.

Typical DNS configuration:

mail.domain.com = x.x.x.x
mx 5 = mail.domain.com

(4) Primary with Data and Backup Servers (no load balancing with failure protection)

©2003 Deerfield Communications, Inc. Portions © Multiple Sources. All rights reserved.

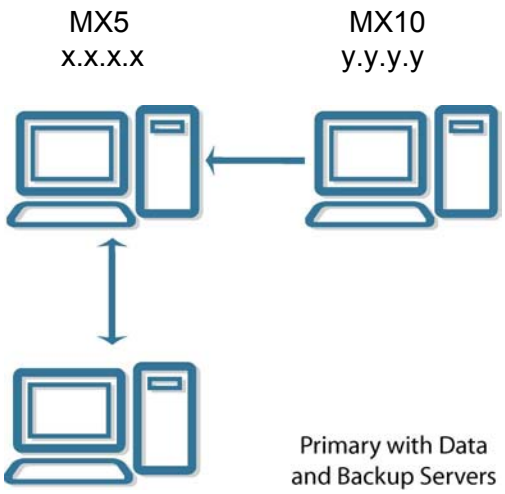

Data Server
z.z.z.z

This is the recommended configuration if you anticipate adding load balanced servers in the future. The data server holds the MailServer Configuration, Mail Store and Domain / User Configuration. The primary server runs the Mail Services and the backup (hosted offsite) provides offline mail receipt. The primary advantage of this configuration is that it provides the foundation for scalability, as multiple load balanced primary servers can be added as the domain and/or user database grows.

Typical DNS configuration:

mail.domain.com = x.x.x.x
backup.domain.com = y.y.y.y
mx 5 = mail.domain.com
mx 10 = backup.domain.com

©2003 Deerfield Communications, Inc. Portions © Multiple Sources. All rights reserved.

(5) Multiple Primaries with Data Server

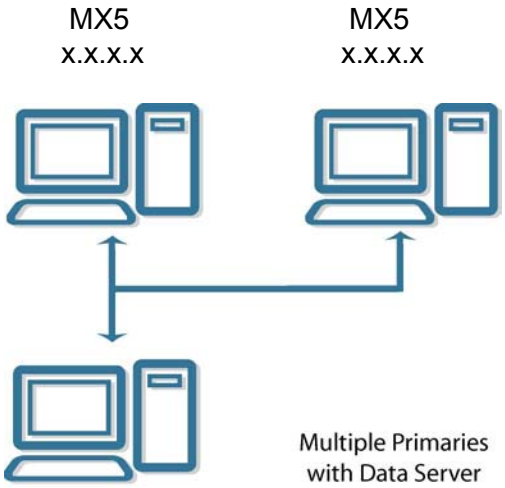

Data Server
z.z.z.z

As outlined above, this configuration can scale upwards to support multiple load balanced primary servers.

Typical DNS configuration:

mail.domain.com = x.x.x.x
mx 5 = mail.domain.com

(6) Service Specific Primaries with Data Server

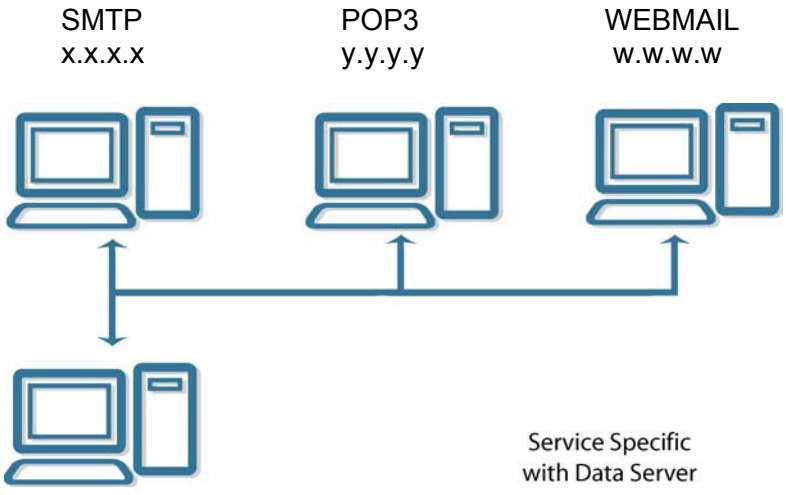

Data Server
z.z.z.z

©2003 Deerfield Communications, Inc. Portions © Multiple Sources. All rights reserved.

This configuration is typically implemented as a load sharing as opposed to a load balancing solution. In this configuration, the primary servers are not load balanced, as each is responsible for one or more Mail Services. As an example Primary 1 may be responsible for running SMTP Services for the domain(s) and Primary 2 may be running POP3 and IMAP. A third server could be added to host WebAdmin and/or WebMail.

Routing to the individual servers is accomplished via DNS host record entries. In the example outlined above, Primary 1 may resolve as smtp.domain.com while Primary 2 may resolve as pop3.domain.com, with Primary 3 resolving as webmail.domain.com.

The primary advantage of this configuration is that the service load is shared among multiple servers.

Typical DNS configuration:

smtp.domain.com = x.x.x.x
pop3.domain.com = y.y.y.y
webmail.domain.com = w.w.w.w
mx 5 = smtp.domain.com

Note: Data Servers house three primary components; the Mail Store, Server Configuration and Domain and User Configuration data. In all cases, the Domain and User Configuration data may be written to an ODBC compliant database on the Data Server or another local network accessible Database Server. For the purposes of this document, the Data Server will be represented as a single machine. The migration of the cluster to the ODBC database is accomplished post installation by configuring the datasource on each server via the Control Panel, Data Sources applet. The conversion from the file system to the database will be run on only one of the clustered servers.

©2003 Deerfield Communications, Inc. Portions © Multiple Sources. All rights reserved.

Installing a VisNetic MailServer Cluster

Requirements

Non Load Balanced

Configurations 1 and 2 – Utilize minimum system requirements as outlined in the diagram above.

Load Balanced

Configurations 3, 4, 5 and 6 – All load balanced configurations require a mapped drive to a resource accessible to all cluster servers. The same mapped drive letter must be designated for all servers (they must all be mapped to drive V:\ as an example). The same windows user must be used, whether installing in a domain environment or not, and must have full rights to the mapped drive. Utilize the minimum system requirements outlined in the diagram above for the individual clustered servers.

Note: These steps are only required for configurations 3-6 outlined above, configurations 1 and 2 are stand alone VMS servers and do not require special cluster configuration.

Note: For the purpose of these instructions, we will configure the Data Server on the shared V:\ drive.

Note: Download the latest VisNetic MailServer Version 6 setup file from: http://www.deerfield.com/download/visnetic_mailserver/

Note: Ensure your servers meet the recommended system requirements as outlined above and that a shared drive (utilizing the same drive letter) is mapped from all servers to the common data server.

VisNetic MailServer configuration settings that are utilized in the implementation of load-balanced clusters include:

Required

a. Mail Path setting (Options, Directories)
b. Config Path (Options, Other Options, System Path Settings)

Optional

a. Temp Path setting (Options, Directories)
b. Log Path setting (Options, Directories)
c. Web Path (Options, Other Options, System Path Settings)
d. Web Temp Path (Options, Other Options, System Path Settings)
e. Local Service ID (Options, Other Options, System Path Settings)

©2003 Deerfield Communications, Inc. Portions © Multiple Sources. All rights reserved.

Configuring the First Server

1. Run the downloaded setup file on the first server (this will vary depending on the type of cluster you are installing, but will typically be the lowest level MX server). Installation to the default drive and directory is highly recommended. (C:\Program Files\Deerfield.com\VisNetic MailServer)

2. Create a directory structure on your shared data server to house the VMS Configuration and Mail Store (V:\VMS\ as an example).

3. Copy the following folders and any sub-folders from C:\Program Files\Deerfield.com\VisNetic MailServer to V:\VMS.

    a. ..\mail
    b. ..\config

4. Access the VMS configuration utility, select the “Options” tab and then select the “Other Options” button.

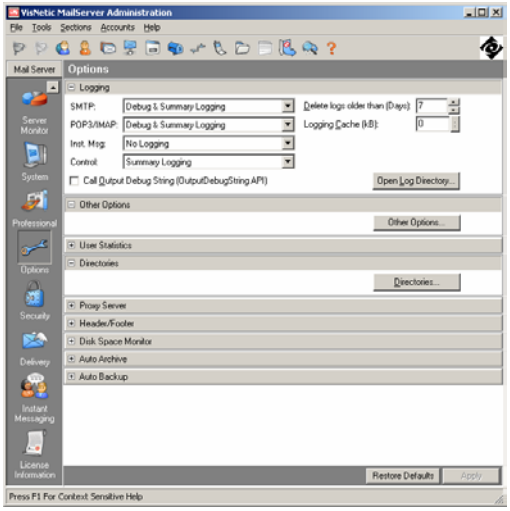

5. Select the “System Path Settings” button and in the resulting dialog, enter V:\VMS on the first line then save the file (File, Save or Ctrl-S)

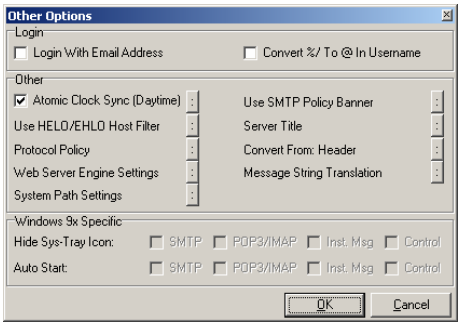

6. Close the “Other Options” dialog by clicking “OK”.

©2003 Deerfield Communications, Inc. Portions © Multiple Sources. All rights reserved.

7. While still on the “Options” tab, select the “Directories” button.

8. In the resulting dialog, enter V:\VMS\mail\ as the “Mail Path” setting, and then click “OK”.

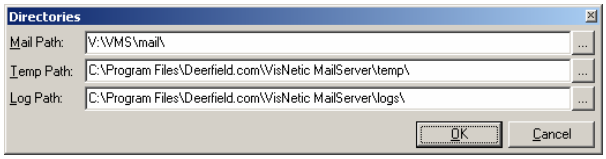

9. Access the “System” tab and stop all services (“SMTP”, “POP3/IMAP”, “Instant Messaging” and Control).

10. Open the Windows Services Control panel and set the VMS services (all four) to Log On with an account that has access to the shared drive that was previously mapped (V:\VMS).

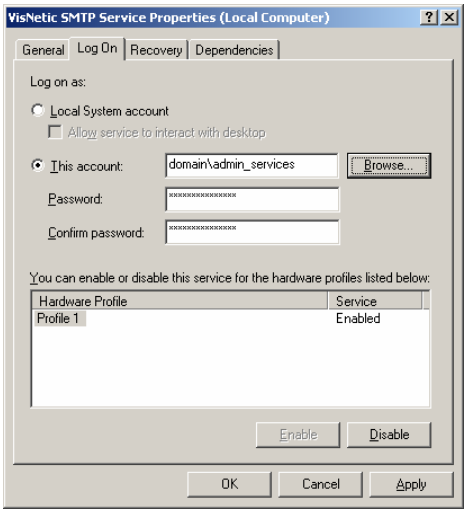

Note: By default the services Log On with a Local System Account, which does not have the appropriate rights to access the mapped drive.

11. Back on the VMS “System” tab, start the services.

VMS is now using the configuration and mail store previously copied to V:\VMS\config and V:\VMS\mail.

Configuring Other Servers

12. Run the downloaded setup file on the next server(s). Installation to the default drive and directory is highly recommended. (C:\Program Files\Deerfield.com\VisNetic MailServer)

13. Access the VMS configuration utility, select the “Options” tab and then select the “Other Options” button.

©2003 Deerfield Communications, Inc. Portions © Multiple Sources. All rights reserved.

14. In the resulting dialog, enter V:\VMS on the first line then save the file (File, Save or Ctrl-S)

15. Close the “Other Options” dialog by clicking “OK”.

16. Access the “System” tab and stop all services (“SMTP”, “POP3/IMAP”, “Instant Messaging” and Control).

17. Open the Windows Services Control panel and set the VMS services (all four) to Log On with an account that has access to the shared drive that was previously mapped (V:\VMS).

    Note: By default the services Log On with a Local System Account, which does not have the appropriate rights to access the mapped drive.

18. Back on the VMS “System” tab, start the services.

VMS is now using the configuration and mail store on the shared data server.

Configure DNS and Load-Balancing

Once the clustered VMS servers are in place, it will be necessary to configure DNS to route incoming requests to the appropriate server(s), depending on the load-balance method you have implemented.

It is important to note that VMS does not provide the load-balancing component of the cluster, as outlined above that functionality is provided by third party software and/or hardware.

Important Considerations

1. The configuration utility can be accessed locally or remotely from any of the clustered servers, however because the configuration is shared it is recommended that you reload (File, Reload Config or F5) the configuration before making any changes. In most cases, you will be warned that the configuration has changed and will be prompted to reload.

2. When the servers boot up, it is important that the shared drive be connected before the VMS services start. This behavior varies widely from installation to installation, dependent on your unique configuration. To ensure the services start in the correct order, you will need to set them to be dependent on the last system services to start. For detailed instructions on how to do this, please see this Microsoft article: http://support.microsoft.com/default.aspx?scid=kb;en-us;193888&Product=win2000

©2003 Deerfield Communications, Inc. Portions © Multiple Sources. All rights reserved.