



Multi-Homing and Network Interface Card Aggregation



Revision History

Revision	Change Description	Updated By	Date
1	Original version	Scott Thompson	9/20/2000
2	Modified to current template	Drew Robbins	3/7/2001
3			
4			
5			
6			



Table of Contents

1	Overview	1
2	Multi-Homing.....	2
2.1	Definition.....	2
2.2	CCS Client Examples.....	2
2.3	CCS Client Issues	2
2.3.1	Company C	2
3	Network Interface Card Aggregation.....	4
3.1	Definition.....	4
3.2	Benefits.....	4
3.2.1	Additional Bandwidth	4
3.2.2	Redundancy	4
3.3	Server Configuration.....	5
3.4	Switch Implications and Configuration.....	5
3.5	Example Topologies and Scenarios	5
4	Best Practices.....	7
4.1	Recommendations.....	7



1 Overview

This white paper provides an overview of Multi-Homing and Network Interface Card (NIC) Aggregation technologies in a MetaFrame environment. It gives examples of each technology as well as recommended best practices for implementation of these technologies in a MetaFrame environment.



2 Multi-Homing

2.1 Definition

Multi-Homing is defined as servers that have more than one network interface card (NIC). Each interface can be connected to one or more networks. Multi-homing is typically used to provide load balancing or redundancy. In the below examples each interface is connected to a different network subnet.

2.2 CCS Client Examples

Company A

Company A has multi-homed servers in their hosted MetaFrame environment. Each server has two interfaces that are on different network subnets. One interface is dedicated for public Internet traffic. The other interface is for access to backend SQL Servers and Company A's internal network. This internal network is not accessible from the public Internet.

Company B

Company B is using multi-homed MetaFrame servers. One interface is on a public network segment and the other is on a private network segment. One of these MetaFrame servers also contains the NFuse extensions. There is a dedicated ICA Master Browser with only one interface that is located on the public segment. The dedicated ICA Master Browser is also a standalone server and not a member of a NT Domain. Additionally, there is a web server that has interfaces on both the public and private network segments.

Company C

Company C has multi-homed MetaFrame servers in their internal MetaFrame environment. One network interface is used for ICA connections from users. The other network interface is dedicated for back-end services such as communication to Domain Controllers or databases.

2.3 CCS Client Issues

Using multi-homed servers must be very well planned to implement a proper architecture. There have been several issues related to the use of multi-homed servers. Below are examples of these issues.

2.3.1 Company C

The issues described below were seen at Company C. Company C's environment included Windows 2000 Terminal Services SP1 and MF 1.8 SP2. As described above, each MetaFrame server had two network interfaces on different subnets.

Binding Order: In the Windows 2000 GUI you are able to change the binding order of your NIC adapter. It has been seen consistently that after changing the binding order a few times, the change no longer takes affect. This is important in a MetaFrame environment because Citrix relies on the adapter binding order when sending and receiving ICA packets. The first bound adapter will have priority over all subsequent adapters. The first bound adapter will be listed first in the server's routing table and will be used as the default route. In this scenario, after changing the binding order a few times, it will no longer work. The GUI will display that it has changed when in fact it has really not. The "QSERVER" command should return the default interface, the first adapter in the binding order. After this function breaks, performing the "QSERVER" command will always return the same interface regardless of the modified binding order. This is also demonstrated by the use of the "NBTSTAT" command. This command should also return the default interface. Running this command showed the same results as the



"QSERVER" command. Additionally, when populating the server list, Published Application Manager uses similar functionality as the "QSERVER" command. Therefore, if the client's network adapter is not the first bound adapter, only server addresses on the inaccessible network are returned.

Resolution: When installing the server verify the adapter binding order on all multi-homed servers. Make sure the network that users will be connected to is bound to the first adapter. Note: Data can be pulled from each interface individually by running QSERVER with additional parameters. The proper syntax is the following: QSERVER /tcpserver: <ip address>

Applications published on the W2K Servers are not added to the ICA Master

Browser's application list: In Windows 2000, a key is added to the registry for each network adapter in the following path: HKLM\System\CurrentControlSet\Services. MetaFrame servers send their published application information to the ICA Master Browser using the first adapter key found in the registry. If the network interface that cannot communicate with the ICA Master Browser is configured as the first adapter key, application information will be stored locally and will not reach the ICA Master Browser.

Resolution: Verify that the first adapter in the registry points to the network that can communicate with the Mater Browser.

Inability for users to establish connections to W2K Servers or published applications:

When using Load Balancing, a Load Level is determined for each MetaFrame Server. Load for each factor is computed on a 0-10000 basis. Each factor contributes a percentage of the final load, which is also between 0 and 10,000. When calculating the load value, the MetaFrame server verifies that listener ports are available on all adapters. If a listener port is not configured for each of the adapters, a load level of "10002" will be generated. A load level of "10002" indicates that no ICA connections are available on that server. If this is the case, users will not be able to establish connections to the W2K server or its published applications.

Resolution: Change the TCP/IP listener in Citrix Connection Configuration to include all adapters or add a new listener for each adapter that did not have one. Load levels can be found by running the QSERVER /APP command. Other load levels are:

9999 = No load balancing installed

0 to 9998 = "normal" load level

0000! = Application is disabled for this server

10000 = Load is at 100%

10001 = Out of licenses

ICA Master Browser retains IP address even if changed: The ICA Master Browser stores the IP addresses of servers within the farm until it receives updates from one the servers. If an adapter is added with a new IP address or if an IP address is changed, the ICA Master Browser's server list retains the previous IP address. In this instance clients will be returned the wrong IP address.

Resolution: Stop and start the ICA Browser Service on the server with the new adapter or IP address.

3 Network Interface Card Aggregation

NIC aggregation using multiple NICs provides additional bandwidth, fail-over and reduced complexity over a standard multi-homed server. The aggregation feature of the NIC driver enables the OS to view multiple NICs as one interface. As a single interface, many of the multi-homed environment issues are resolved. NIC aggregation is also known as Multi-Link Trunking, Fast-Etherchannel, and NIC bonding. NIC aggregation requires support from the NIC vendor and switch vendor.

3.1 Definition

NIC aggregation, also referred to as multi-link trunking and fast ether-channel, is a technique that allows parallel physical links between a switch and a server to be used simultaneously. This multiplies the bandwidth and provides additional redundancy between the devices.

3.2 Benefits

NIC aggregation provides two major benefits. As described above, these two major benefits include additional bandwidth and redundancy. These benefits can be gained over Ethernet infrastructures.

3.2.1 Additional Bandwidth

NIC aggregation co-effectively provides incremental bandwidth between servers and switches. NIC aggregation offers the potential to multiply aggregate bandwidth. Using aggregation can resolve bandwidth bottlenecks that can be found at the server. By providing scalable network bandwidth, NIC aggregation prevents server I/O bottlenecks at the network interface.

3.2.2 Redundancy

NIC aggregation provides automatic point-to-point redundancy between servers and switches. If a link in a trunk fails, the network traffic flows mapped to that link are dynamically reassigned to the remaining links of the aggregated link. The process of redirecting traffic from the failed link to the remaining links is less than a second. This process is transparent to end-users. Since no host protocol timers expire, no sessions are dropped.

Figure 1 gives an example of a NIC failure. In this example two links form an aggregated link. Network traffic is distributed across the links in the event of a failure. In this example, if Physical link 1 fails, the network traffic flow between Server A and Server B is re-directed to Physical link 2. The re-direction occurs as soon as the switch learns that an address has moved from one port to another. As a result NIC aggregation is extremely fault-tolerant.




Figure 1. *Fault Tolerance with Link Aggregation*

3.3 Server Configuration

NIC aggregation is achieved on the server by the use of network interface cards that support aggregation. NIC software drivers coordinate distribution of loads across multiple network interfaces. The NICs assume one IP address and one MAC address and they appear as one device to the Operating System.

3.4 Switch Implications and Configuration

The switch vendor must support NIC Aggregation for it to work. Aggregation is configured through identification and definition of the number of ports that make up a channel using a command line interface or with SNMP management applications. Once the ports have been configured as one channel, devices can be connected to it. Aggregation can be implemented across different switches in a back plane or different chassis. Switch chassis will need to be physically connected by ATM, fiber-channel, or some other media.

3.5 Example Topologies and Scenarios

Aggregated Links: Figure 2 shown below shows an example of aggregated links between multiple servers. Servers are shown with 100, 200, and 400 Mbps connections.




Figure 2. *Aggregated Links Between Multiple Servers*

Aggregation Between Switches: Figure 3 shown below shows an example of an aggregated link between switches.




Figure 3. *Aggregated Link Between Switches*



4 Best Practices

The use of multi-homing and NIC Aggregation depends on what is trying to be accomplished. Below are some recommendations on when to use each method in a MetaFrame implementation.

4.1 Recommendations

✓ **Multi-Homing**

Multi-Homed MetaFrame servers pose design considerations for ICA browser functionality. Multi-homing a MetaFrame server requires extensive planning and testing. Due to the complexity of multi-homed MetaFrame servers, it is typically not recommended. In instances where it is necessary to have an additional layer of security between clients and back-end servers, multi-homed servers, if well planned and properly configured, can aid in this scenario. Below is a diagram of such a scenario.

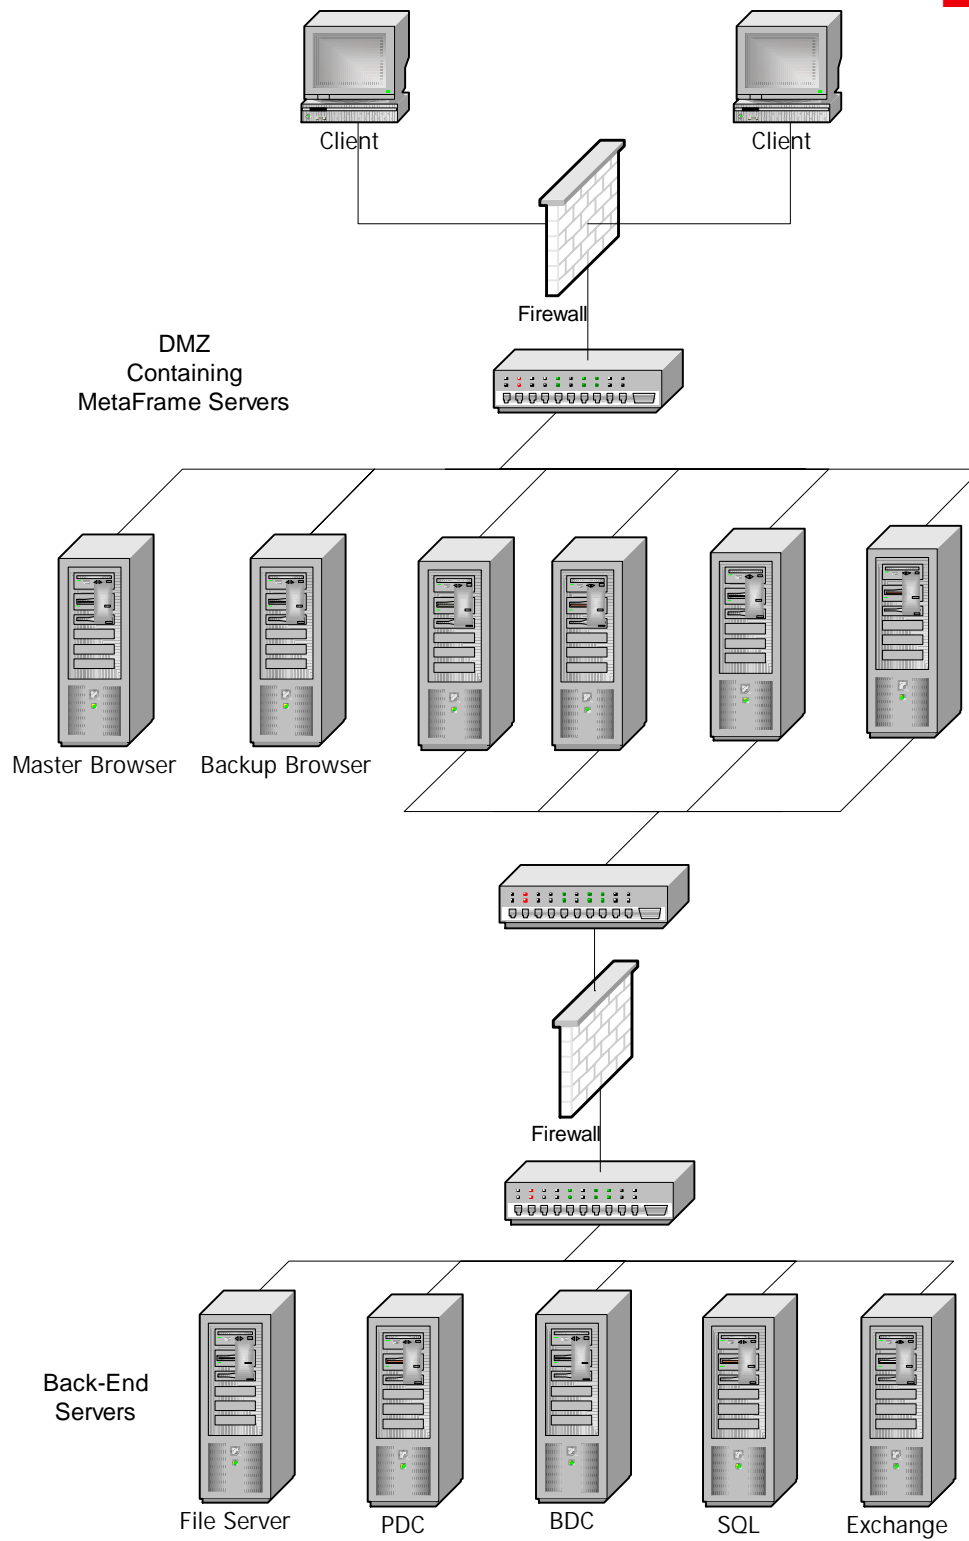


Figure 4. Use of Multihomed Servers



This is a high-level general scenario. It does not include other servers that would typically be located in a DMZ. In this simplified scenario clients are connecting through a firewall. All appropriate TCP Ports are open to allow traffic flow. All the MetaFrame servers are multi-homed with the exception of the ICA Master and Backup Browser. One network interface is external and is accessible to the clients. Network address translation (NAT) is set up for these interfaces. The second interface is an internal network that is not accessible to the clients. The multi-homed servers have been configured appropriately in terms of binding order to ensure that load information is appropriately routed through the external interface to the dedicated ICA Master Browser. TCP/IP listeners have been set up for ICA on each interface using Citrix Connection Configuration. There is an additional firewall that resides between the MetaFrame servers and back-end data servers. In this instance these back-end servers consist of a Primary Domain Controller (PDC), Backup Domain Controller (BDC), File Server, SQL Server, and Exchange. The ICA Master Browser and Backup Browser are dedicated and do not host any applications. Within Citrix Server Administration the dedicated ICA Master Browser is configured to "Always attempt to become Master Browser" and the dedicated ICA Backup Browser is configured to "No Preference". All other MetaFrame servers are set to "Do not attempt to become Master Browser". The server designated as the ICA Master Browser will become the ICA Master Browser on the external network. If all the multi-homed servers are configured properly and tested they should send load information through the external interface to the dedicated ICA Master Browser.

Although this configuration will work there are some caveats.

1. The single-homed ICA Browsers cannot be part of an NT Domain. Given this it will make them more difficult to manage.
2. There will also be an election for an ICA Master Browser on the internal network. Although clients will not use this ICA Master Browser, it is important to note. This could be controlled by also having a dedicated ICA Master Browser that resides on the internal network. This server could be less expensive and only serve this purpose.

✓ **NIC Aggregation**

If redundancy and the additional benefit of increased throughput are desired, the use of a NIC aggregation method is recommended over multi-homing. NIC aggregation eliminates all of the issues associated with multi-homing.