



Cisco's Windows Networking Design Implementation Guide
Revised for Windows Server 2008r2 and Windows 7

Paul Willy CCSI/MCITP-EA



Contents

Introduction	4
Before You Begin	4
Conventions	
Prerequisites	
Components Used	
What Is Windows Networking?	4
Domains versus Workgroups	
What Protocol Did It Use?	5
What Protocol Does It Use Now?	6
IPv6 Primer	
Back to Windows	
Dynamic IP Addressing	8
What Is DHCP?	
DHCP Scopes	
DHCP Relay	
DHCP Options	
Cisco DHCP Servers	
Name Resolution	10
Host Name Cache	
DNS	
LLMNR	
Hosts File	
WINS	
Name Lookup Order	
NetBIOS Node Type Name Search Order	
The Microsoft LAN Services Browser	11
Scaling to Larger Networks	12
Domain Trusts	
Single Domain	
Transitive Kerberos Trusts	
Answers to New Challenges	12
Etherchannel	
Virtual PortChannel (vPC)	
Load Balancing and Forwarding Methods	14
Microsoft Network Load Balancing	
Hyper-V and Networking	14
New Horizons Class Offerings	16



Introduction

The term “networking” covers a broad range of technologies, which, combined together, allow computers to share information. Networking components can be segmented into end-system applications, network operating systems, and networking equipment.

A network operating system is software run on all interconnected systems. When Cisco originally wrote this design guide, they included Novell NetWare, Sun’s NFS (Network File System), AppleShare, and Microsoft’s implementation of a network operating system commonly called Windows Networking. Windows Networking is now extensively deployed with millions of nodes. Fundamental changes were made with Windows Vista and Server 2008 (NT 6) and Windows 7 and Server 2008 R2 (NT 6.1). This updated document will reference NT 6 and NT 6.1 instead of Vista, Windows 7, Server 2008 and Server 2008r2 to minimize the verbosity.

This design guide explains the basic concepts of Windows Networking and provides insight on how to design networks (LANs and WANs) to best utilize this operating system. The guide also explains protocols, naming, and scaling issues associated with Windows Networking.

In the decade since the original guide (1) was published Microsoft has made significant changes to their network architecture however some of the more troublesome processes that this document was written for remain in production. The tenets of the original document remain valid because of this.

Before You Begin

Conventions

For more information on document conventions, see the Cisco Technical Tips Conventions.

Prerequisites

There are no specific prerequisites for this document.

Components Used

This document is not restricted to specific software and hardware versions. The information presented in this document was created from devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. **If you are working in a live network, ensure that you understand the potential impact of any command before using it.**

What is Windows Networking?

Windows Networking refers to the networking system shared by the software that comes with all the following Microsoft operating systems or servers:

- Windows 95, 98, and ME (Legacy)
- Windows NT 3.1/3.50/3.51/4.0 (Legacy)
- Windows 2000 and XP/2003 also known as NT 5 and NT 5.1
- Windows NT 6 Windows Vista and Server 2008
- Windows NT 6.1 Windows 7 and Server 2008 r2

Microsoft LAN Manager, the LAN Manager client for MS-DOS, Windows for Workgroups, Windows 9X and Windows NT 3.1 through NT 4 are not discussed in this document except in a historical context.

Domains versus Workgroups

Windows Networking has three concepts of a group of related computers — workgroups, a domain and a domain hierarchy. Workgroups can be any logical collection of computers; any computer on the network can join an existing workgroup or create a new one.

More formal entities, Active Directory Domains are created and managed by a set of domain controllers (DC). They run a process that runs on a Windows server. A domain has security and administrative properties that a workgroup does not. Each Active Directory Domain must have at least one 2000 thru 2008r2 server, which is responsible for the PDC emulator role. User account and computer account information in the domain and security within the domain are controlled by one or more DCs. Only the Accounts Data, Scripts and Group Policy settings are replicated between the Domain Controllers. This creates a “loosely consistent” replicated database.

Windows Networking domains are not the same as Internet domain names as used by the Domain Name System (DNS), although they are dependent on the DNS name space. Using a special type of DNS record called a SRV to locate DCs and their services. A domain hierarchy or Active Directory Hierarchy is a collection of domains organized into parent-child relationships. This convention, introduced with Windows 2000, enables easier searching through multiple domains in a single query (among other things). This hierarchy maps exactly to a DNS namespace.

Microsoft’s choice to use DNS to locate DCs through SRV records has caused most organizations to use “Split Brain” DNS. DNS was intended to be a single database that did not exist in a single place. On the Internet domains are located by walking the tree. The root servers are configured with referrals to the com, org, edu, us etc. Those servers are configured with referrals to Cisco, Microsoft, PBS etc. Most organizations do not want referrals to the domain containing their internal servers workstations and devices. This has led to the “Split Brain” where organizations have a different set of DNS records for public and private networks.

What Protocol Did It Use?

Prior to Windows 2000, Windows Networking used the NetBIOS protocol for file sharing, printer sharing, messaging, authentication, and name resolution. Microsoft claimed a pure Windows 2000 installation would require NetBIOS only for interoperability with earlier versions of Windows Networking using the flat NetBIOS namespace. NetBIOS is a session-layer protocol that can run on any of the following transport protocols:

- NetBEUI (NetBIOS over LLC2)
- NWLink (NetBIOS over Internetwork Packet Exchange [IPX])
- NetBIOS over TCP/IP (NBT)

Through Windows Server 2003, Microsoft recommended that clients use only one transport protocol for maximum performance. You should pick a protocol to use for your entire network, preferably TCP/IP, and then turn the other protocols off because the NetBIOS name service maintains information about computer names (a name space) separately for each transport. Name spaces do not interact with each other; each transport operates as a separate network.

NetBEUI (NetBIOS over LLC2) is the least scalable of the three protocols because it must be bridged. NetBEUI is included only to support very old services (for example, old versions of LAN Manager). NetBEUI does not require any client address configuration. There is no fixed limit to the number of Windows clients can have with NetBEUI, but it is common for this solution to run into performance problems as the number of clients



in a single bridge group goes above 50 to 100 users. The flat topology and reliance on broadcasts did not scale, especially when traffic had to traverse a WAN link.

NWLink was recommended only for networks already running IPX that cannot be upgraded to use TCP/IP. Similar to NetBEUI, NWLink requires no client address configuration. NWLink used IPX type-20 packets to exchange registration and browsing information. To forward type-20 IPX packets across Cisco routers, you had to configure ipx type-20 propagation on each interface on every router on your network.

For scalability reasons, it was recommended to utilize NetBIOS over TCP (NBT) for most networks, or anytime the network includes a WAN. Since NBT uses TCP/IP, each computer must be configured to use a static IP address, or to fetch an IP address dynamically with the Dynamic Host Configuration Protocol (DHCP). For ease of network administration, it was highly recommended to use DHCP; for optimum network performance, it was highly recommended to use a (Windows Internet Name Service) WINS server as well. A WINS server allows clients to get browsing information without having to broadcast requests every time. There is a direct correlation between the number of broadcasts in a network and network performance; broadcasts are necessary for a network to function, but minimizing them can be critical.

Cisco recommends that most customers use TCP/IP for Windows Networking. The bulk of the old design guide focused on designs using NBT. Microsoft used the same IP stack from Windows for Workgroups through Windows Server 2003 r2. Most networkers grew to dislike the Browser service this method depended on because of browser elections that consumed the network during periods of peak use. With Windows NT 6 all this has changed. Windows no longer exhibits the conflicted name resolution methods of the old stack and "browser elections" can become a thing of the past.

What Protocol Does it Use Now?

With NT 6 the issues have changed. Now we have a mix of legacy applications that may have dependencies on old protocols with new servers that introduce IPv6 in addition to IPv4. To the extent that it is possible the old protocols should be disabled, however in manufacturing some of the processes on the shop floor may still require them. Also ancient accounting systems may have NetBIOS hooks that fail if the old protocols are disabled.

- TCP/IP v6
- Pure TCP/IP v4
- NetBIOS over TCP/IP (NBT)

Ipv6 is the future, how long will it take and how does Windows fit into an Ipv6 discussion. Simply, the Ipv4 address space is exhausted. Network address translation (NAT) has extended the use of Ipv4 for over a decade, but it has issues. NAT breaks protocols that embed IP addresses. For instance, with VoIP, the client computer says to the server, "Please send incoming calls to this address." Obviously this doesn't work if the address in question is a private address. Working around this requires a significant amount of special case logic in the NAT device, the communication protocol, and/or the application. For this reason and a few others, most of the people who participate in the Internet Engineering Task Force (IETF) don't care much for NAT (2).

IPv6 Primer

IPv6 addresses are written down as eight 16-bit values with colons between them, and each 16-bit value is displayed in hexadecimal, using numbers 0-9 and the letters A - F. For example, 2001:db8:31:1:20a:95ff:fe5:246e. It's not uncommon for IPv6 addresses to have a sequence of consecutive zeroes. In these cases, exactly one of those sequences can be left out. So 2001:db8:31:0:0:0:0:1 becomes 2001:db8:31::1 and the IPv6 loopback address 0:0:0:0:0:0:0:1 becomes ::1. There are a total of 340,282,366,920,938,463,463,374,607,431,768,211,456 IPv6 addresses.

IPv6 hosts use stateless auto configuration by default. Much like IPX; a router advertisement informs the client what the network address is, and the client does some math on its MAC address and adds the host ID to the network ID it got from the router. There is DHCPv6, but it is not as important as it was in IPv4. The experimental RFC 5006 enables a router to configure DNS for the client as well as other parameters.

There is a flaw in this router discovery process, Cisco and Linux have patched for it, but Juniper and Microsoft have not. It does require link local access, but when processor utilization goes to 100% and the device must be power cycled, that would be bad, the Bugtraq ID:45760.

IPv6 was supposed to have significant built in security, it does and it doesn't. IPsec is built in, but its implementation is not. Deploying certificates is easier now but not widely used. Many software firewalls only support IPv4 and simply pass IPv6 traffic, the Windows firewall in NT 6 and 6.1 is an exception. This is one point where Microsoft is ahead of MAC OS, Linux and BSD (2 p3).

Back to Windows

How do you know if NetBIOS is in use? The best method to discover if NetBIOS names are being used is the Windows command line. The command **nbtstat -c** will show what names are being resolved as NetBIOS names. This command can be executed on both servers and workstations to determine what names are using NetBIOS. XP and Server 2003 (NT 5.1) still have some dependencies, Windows Server 2008, and R2 as well as Windows Vista and 7 (NT 6 and 6.1) seem to have finally put NetBIOS to rest, however that does not mean all processes installed on them will function without it.

Windows NT 6 and 6.1 have added IPv6 dependencies for some services. The most likely service you may run into that requires v6 is DirectAccess. There were some poorly received services like The Network Meeting Place in NT 6 that were deprecated in NT 6.1. The vast majority of services available on NT 6 and 6.1 work over either IPv4 or IPv6. Both NT 6 and NT6.1 prefer IPv6 over IPv4, to view this the command **netsh interface ipv6 show prefixpolicy** shows the prefix policy table and the command **netsh interface ipv6 add|set|delete prefixpolicy** allows its modification. IPv6 can be disabled or modified through the registry, see KB 929852. Group Policy Preferences can be used to deploy registry settings to entire organizational units, use them carefully!

From Microsoft's perspective, IPv6 is a mandatory part of the Windows operating system and it is enabled and included in standard Windows service and application testing during the operating system development process. Because Windows was designed specifically with IPv6 present, Microsoft does not perform any testing to determine the effects of disabling IPv6. If IPv6 is disabled on Windows 7, Windows Vista, Windows Server 2008 R2, or Windows Server 2008, or later versions, some components will not function. Moreover, applications that you might not think are using IPv6 — such as Remote Assistance, HomeGroup, DirectAccess, and Windows Mail — could be (7).



Microsoft has implemented IPv6 over IPv4 Tunneling (3) where the The IPv4 Protocol field is set to 41 to indicate an encapsulated IPv6 packet, the source and destination addresses are Ipv4. It is possible to configure tunnels, but the ones of most concern are the automatic ones. ISATAP is used for unicast communication across an IPv4 intranet and is enabled by default. 6to4 is used for unicast communication across the IPv4 Internet and is enabled by default. Teredo is used for unicast communication across the IPv4 Internet over network address translators (NATs). Teredo support is included and is enabled but inactive by default.

ISATAP has received the most press as a potential security flaw. The reason for this is that it's possible for a rogue device to dynamically register this name in DNS (4). If that happens, the client systems will auto-configure themselves to use the rogue device as their web proxy, or configure their ISATAP adapters to use the rogue device as their ISATAP gateway. Both of these scenarios are enabled by the fact that Internet Explorer uses auto-discovery by default to configure the web proxy, and the ISATAP adapter is enabled by default if the name ISATAP can be resolved and the client can contact an ISATAP router.

6to4 may also be an issue, because the tunnel endpoint is not preconfigured but is determined dynamically based on the destination of the packet. Netsh can expose these connections and either netsh or group policy can be used to disable them. If routes to 2002::/16 exist in the output of **netsh interface ipv6 6to4 routing** command there is a 6to4 tunnel. The command **netsh interface ipv6 6to4 set state disabled** or group policy can disable them (5).

Dynamic IP Addressing

What Is DHCP?

Manually addressing TCP/IP clients has always been both time consuming and error prone. To solve this problem, the Internet Engineering Task Force (IETF) developed DHCP, the Dynamic Host Configuration Protocol. DHCP is designed to automatically provide clients with a valid IP address and related configuration information (see the section DHCP Options below). Each range of addresses that a DHCP server manages is called a scope.

DHCP Scopes

You must configure a range of addresses for every IP subnet where clients will request a DHCP address; each range of addresses is called a DHCP scope. You can configure a DHCP server to serve several scopes since the DHCP server or servers do not need to be physically connected to the same network as the client. If the DHCP server is on a different IP subnet from the client, then you need to use DHCP relay to forward DHCP requests to your DHCP server.

DHCP Relay

DHCP relay typically runs on a router and the relay support is available on Windows NT Server versions. On Cisco 700 series routers, you can turn on DHCP relay with the **set dhcp relay** command. You can turn on DHCP relay on a Cisco IOS router by configuring **ip helper-address** with the address of the DHCP server on each interface that will have DHCP clients. The **ip helper-address** command forwards many other IP broadcasts, including DNS, Trivial File Transfer Protocol (TFTP), and NetBIOS name service packets. To forward

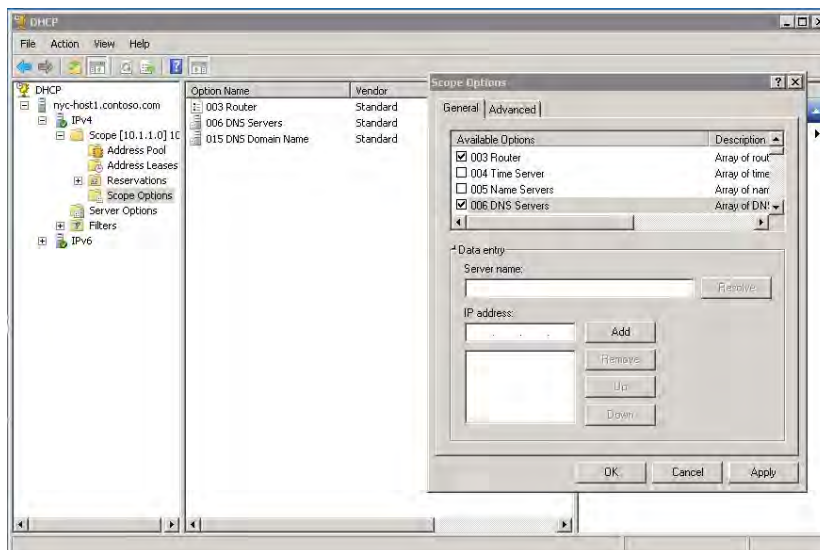
only DHCP requests, see the following example configuration. For more information, see the “Configuring Broadcast Handling” section in the Network Protocols Configuration Guide, Part I.

```
no ip forward-protocol udp tftp
no ip forward-protocol udp dns
no ip forward-protocol udp time
no ip forward-protocol udp netbios-ns
no ip forward-protocol udp netbios-dgm
no ip forward-protocol udp tacacs
ip forward-protocol udp bootpc
interface ethernet 0
ip helper-address 172.16.12.15
interface ethernet 1
ip helper-address 172.16.12.15
```

Where the address 172.16.12.15 is the DHCP server.

DHCP Options

In addition to its IP address, a DHCP client can get other TCP/IP configuration information from a DHCP server, including the subnet mask, default gateway, and DNS information. These pieces of information, called DHCP options, can be configured in the DHCP Manager on your Windows DHCP server.



If your clients are using Windows Internet Name Service (WINS) for name resolution (discussed later), you should configure the address of the WINS server (option 44) and the WINS node type (option 46). A brief list of node types is included in the “Name Resolution” section. The node type p-node (0x2) is strongly recommended when a WINS server is available.

Cisco DHCP Servers

Cisco has an integrated DHCP and DNS server for Windows, Apple and UNIX; the server has a graphical interface, support for secondary addressing, and many other enterprise features. The Cisco 700 series routers (in Release 4.1 and later) and Cisco IOS routers (in Release 11.2(7)F and later) also include a DHCP server that can assign addresses on local network segments. Cisco routers also include network and port-level address translation.



Name Resolution

Name resolution is the process of associating a convenient name, such as FRED or fred.domain.com, with a network address (often an IP address). In legacy systems, this applied to the way that Windows Networking resolved a NetBIOS unique workstation names (described as WORKSTATION<00> in a later section) to an IP address. This process should not be confused with the related but different process of browsing (which uses other types of NetBIOS names). As of the release of Microsoft Windows NT 6, Windows Networking clients use several methods of name resolution:

- Host name cache
- DNS
- LLMNR
- HOSTS file
- WINS

Host Name Cache

NT 6 uses host name cache to store names resolved by DNS for the time to live (TTL) assigned by the primary DNS server. NT 6 like previous versions also stores errors for 15 minutes this can be problematic and can be solved by an "ipconfig /flushdns", a reboot, or create a DWORD value in HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\DNSCache\Parameters and name it MaxNegativeCacheTtl Set the value to "0" (zero).

DNS

Any DNS server can be configured statically to answer queries for computers with static IP addresses. This scenario is useful if computers in your network have static IP addresses. When Windows systems use DHCP to get an IP address, netmask, default gateway, DNS and WINS to register a NetBIOS name, you can set up a Windows DNS server to query a WINS server for names or addresses that were not entered statically. In both cases, Windows and non–Windows systems can resolve IP addresses correctly. This configuration is no longer recommended.

If an administrator configures each Windows Networking server with a static IP address, it may be convenient to enter each server in the DNS system and use DNS for name resolution. With Windows, DNS servers can also be dynamically configured with address to name mappings. DHCP clients, DHCP Servers and Dynamic DNS servers work together to update name to address mappings in the DNS server. The DHCP server can perform this update for non–Windows NT DHCP clients. Dynamic DNS has significant security issues.

LLMNR

Link-local Multicast Name Resolution (LLMNR) is a new protocol that provides an additional method to resolve the names of neighboring computers; it is not enabled by default. LLMNR is especially useful for networks that do not have a Domain Name System (DNS) server like a home or very small business network. LLMNR uses a simple exchange of request and reply messages to resolve computer names to Internet Protocol version 6 (IPv6) or IP version 4 (IPv4) addresses. It does not scale, it uses a ttl of 1 so routers will not forward it. It is considerably less hostile than the NetBIOS Browser Service and the browser elections it caused.

Hosts file

The hosts file found in C:\windows\system32\drivers\etc works as host files have since the beginning of IPv4. It should be avoided whenever possible unless documented due to the scalability problem of touching every host.

WINS

WINS or the NetBIOS Name Service is almost done. For most organizations as soon as the last Windows XP/2003 is turned off, NT 6 and 6.1 have no NetBIOS dependencies this author has seen. In fact when reviewing event logs, no browser elections in months produces convincing evidence that what Microsoft promised in 1999 at TechEd, they have finally delivered. Do not forget the shop floor, or accounting, there still may be legacy programs or devices that require NetBIOS name resolution. The command **nbtstat -c** on those machines will show if they are using NetBIOS names.

Name Lookup Order

Windows networking components sends name resolution queries in a different order, depending on the NetBIOS node type. If the system is Windows NT 6 or 6.1 and the name is longer than 15 characters, then Windows sends only a DNS query. Other networking components and services may also use a different order depending upon the API called to perform name resolution. This is where legacy systems running '90s operating systems or accounting systems with NetBIOS hooks may still use the old order. Name lookup in NT 6 is performed in the following order:

- Check for local host name.
- Check the Host name cache (which is populated with entries in hosts).
- Send a DNS query (use the list until one responds).
- Send a broadcast query or a WINS query, depending on the current NetBIOS node type.
- Check the LMHOSTS file if configured.

NetBIOS Node Type Name Search Order

- b-node (0x1) Broadcast only
- p-node (0x2) WINS only
- m-node (0x4) Broadcast, then WINS
- h-node (0x8) WINS, then broadcast

The Microsoft LAN Services Browser

Windows Networking was originally designed to run on a single LAN segment or a bridged (flat) network. At that time, only the NetBEUI protocol was supported. Microsoft developed the LAN Services Browser to enable the user to browse a list of all computer services available on the network. Each Windows Networking client registered its NetBIOS names periodically by sending broadcasts.

Every computer also had to send broadcasts to elect a browse master for the network. The browse master (and several backup browse masters) maintained the list of computers and their addresses. When a user browsed the network, the client sent a broadcast request and the browse master responded. If the browse master was busy, an election packet would cause every potential browser to broadcast 3 times over every configured protocol.

Eventually Microsoft added support for NetBIOS over IPX and NetBIOS over TCP/IP, but Windows Networking still assumed that all clients and servers were on the same logical IPX network or IP subnet — they still sent broadcasts to register and find computers on the network.

This architecture, although simple to implement, generated an enormous burden on the network and on the CPU of each client on the network. Because of these scalability problems, Microsoft began to offer other methods of browsing and name resolution — ways for clients to map a name to the IP address of other computers on the network. Eventually Microsoft also provided a way to browse and resolve names without broadcasts.



The rest of this section explains how browsing works in various environments. The previous section explained how individual NetBIOS names are resolved. These two activities are similar but distinct. Users browse the network when opening the network neighborhood, using the net view command, or logging into a Windows NT domain at startup. Name resolution is the process of resolving names previously known, or found when browsing. Please note that this discussion is unrelated to Web browsers.

In Windows NT, it is not necessary to turn off browsing in most cases, although it may be desirable. In Windows NT, set the Hkey_local_machine\system\CurrentControlSet\Services\Browser\Parameters\MaintainServerList registry key to No. Administrators can control broadcasts sent by DHCP clients by selecting the appropriate WINS node type (p-node: 0x2).

Scaling to Larger Networks

Domain Trusts

When planning a Windows network, consideration of what domain model to use is important. The following paragraphs discuss the benefits and drawbacks of several domain models. If you have several domains, you probably want to exchange data with other domains in your network. Trust relationships are a way to gain or grant access to a domain without having to manage each user individually, because Global and Universal security groups can be assigned permissions, or nested in Domain Local groups in any domain that is trusting.

Single Domain

This domain model is the simplest; the network has only one domain. This setup works for small medium or large sized installations. Only the largest organizations require more than one production domain. The chemical company in Midland would be an example of an organization that because of 275,000 copies of Windows and a global presence determined that more than one domain in a forest was appropriate. Microsoft and HP did a technology proof of concept where they created a single domain with 100,000 users and it worked fine.

Transitive Kerberos Trusts

Designed for companies without a central administrative or IS organization, the transitive Kerberos trust model is the easiest to implement. Every domain is installed as a child of a Forest Root, or as a new tree in an existing forest. This model adds administrative complexity because an Administrator can only be an admin in one domain, and there are many functions only possible by an Enterprise Admin which only exists in the forest root.

Answers to New Challenges

There are many new challenges to Networkers in an environment implementing NT6 and 6.1 beyond the challenge of IPv6. Virtualization and Failover clustering present issues unanticipated by Cisco's original design guide. Maximizing availability and security simultaneously is non-trivial. Hyper-V and Failover Clusters are the next big challenge.

Servers are increasingly being tasked with higher and higher utilization. Bulldozer and Xeon are up to the task; soon hundreds of CPU cores on a single server with a quarter terabyte of Ram will not seem unusual. This means when a single server loses network connectivity, many services go offline. One possible answer is the Etherchannel another is Virtual Port Channel.

Etherchannel

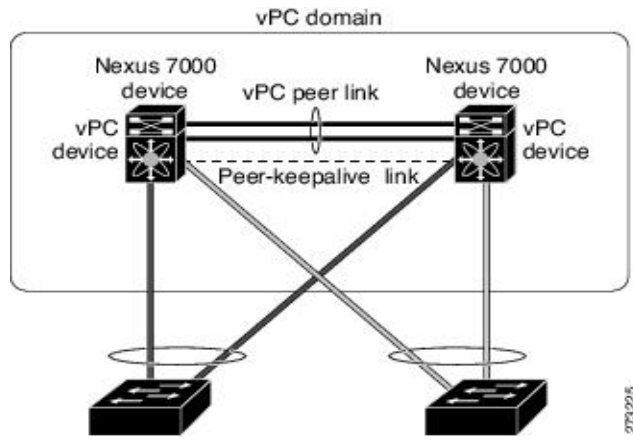
Creating highly available services has led to interesting networking challenges. Network teams to Etherchannels have not been easy and Microsoft, HP and Cisco have not helped. Issues with specific component installation order for Proliant's (6) and non-standard configuration parameters;

```
interface GigabitEthernet0/1
switchport trunk encapsulation dot1q
switchport mode trunk
no ip address
channel-group 1 mode desirable
spanning-tree bpdudfilter enable
!
interface GigabitEthernet0/2
switchport trunk encapsulation dot1q
switchport mode trunk
no ip address
channel-group 1 mode desirable
spanning-tree bpdudfilter enable
```

Do not try to configure more than 6 EtherChannels on the switch. Configure all ports in an EtherChannel to operate at the same speeds, duplex modes, trunking modes and spanning-tree modes. Do not configure a port to be a member of more than one EtherChannel group. Use good documentation and labeling practices! LACP and PaGP are not supported for trunking.

Virtual PortChannel (vPC)

A virtual port channel (vPC) allows links that are physically connected to two different Cisco Nexus 7000 Series devices to appear as a single port channel by a third device (see Figure). The third device can be a switch, server, or any other networking device that supports port channels. Beginning with Cisco NX-OS Release 4.1(4), you can configure up to 256 vPCs per device. A vPC can provide Layer 2 multipathing, which allows you to create redundancy and increase bisectional bandwidth by enabling multiple parallel paths between nodes and allowing load balancing traffic.





You can use only Layer 2 port channels in the vPC. A vPC domain is associated to a single VDC, so all vPC interfaces belonging to a given vPC domain must be defined in the same VDC. You must have a separate vPC peer-link and peer-keepalive link infrastructure for each VDC deployed. Consolidating a vPC pair (two vPC peer devices of the same domain) in two VDCs of the same physical device is not supported. The vPC peer link must use 10-Gigabit Ethernet ports for both ends of the link or the link will not form (8).

Load Balancing and Forwarding Methods

EtherChannel balances the traffic load across the links in a channel by reducing part of the binary pattern formed from the addresses in the frame to a numerical value that selects one of the links in the channel. EtherChannel load balancing can use MAC addresses or IP addresses, source or destination addresses, or both source and destination addresses. The selected mode applies to all EtherChannels configured on the switch. You configure the load balancing and forwarding method by using the port-channel load-balance global configuration command.

With source-MAC address forwarding, when packets are forwarded to an EtherChannel, they are distributed across the ports in the channel based on the source-MAC address of the incoming packet. Therefore, to provide load balancing, packets from different hosts use different ports in the channel, but packets from the same host use the same port in the channel.

With destination-MAC address forwarding, when packets are forwarded to an EtherChannel, they are distributed across the ports in the channel based on the destination host's MAC address of the incoming packet. Therefore, packets to the same destination are forwarded over the same port, and packets to a different destination are sent on a different port in the channel.

Microsoft Network Load Balancing

Network Load Balancing (NLB) is a scalability method included in NT 6 Servers that allows scaling out. It works for stateless applications like Web servers. Cisco has excellent documentation that can help prevent NLB from breaking your network (9). Network Load Balancing is a clustering technology offered by Microsoft as part of all Windows 2000 Server Windows Server 2003 and Windows Server 2008 family operating systems. NLB uses a distributed algorithm to load balance network traffic across a number of servers.

Hyper-V and Networking

Etherchannel is an effective method to add both bandwidth and availability to a solution incorporating server virtualization. For Hyper-V in particular, the order of installation is important. Back on page 15 and in the references number 6 are guides to getting an Etherchannel to work with a network team that will also allow dot1q trunking. Microsoft has implemented interesting network flexibility in Hyper-V.

Hyper-V allows 3 types of network switches to be created. In typical Microsoft fashion their names are not the same as other vendors. The type External uses a NIC on the host machine connected to the outside network. The type Internal creates a virtual NIC on the host machine that cannot be connected to any external interface. The type Private can only have other virtual machines connected and cannot communicate with the host or external networks. The following graphic shows the virtual network manager.

As you can see in the graphic, creating an unmanaged switch on Hyper-V is straightforward. If the Host server has dot1q and trunking configured, a single interface or a network team can allow multiple vlans with dot1q tags to share a trunk to the Cisco switch. Once configured for dot1q a ProLiant will expose a separate NIC for each Vlan to the virtual network manager. Be careful with the “allow management operating system to share this network adapter” check box. Even with r2 it has been the source of some Hyper-V host network stability problems.

One caveat, if you ghost a Hyper-V host, and drop multiple copies they all have the same MAC address they do not use the MAC address of the NIC, they use one stored in the registry. When you create virtual machines on those hosts they will have duplicate MAC addresses.

References

(1) Windows Networking Design Implementation Guide. by Cisco http://www.cisco.com/en/US/tech/tk870/tk877/tk880/technologies_tech_note09186a00801aa01f.shtml

(2) Everything you need to know about IPv6 By Iljitsch van Beijnum <http://arstechnica.com/hardware/news/2007/03/IPv6.ars>

(3) Ipv4 to Ipv6 Transition Technologies by Microsoft <http://download.microsoft.com/download/1/2/4/124331bf-7970-4315-ad18-0c3948bdd2c4/IPv6Trans.doc>

(4) Does Removing ISATAP for the DNS Block List Impact Security? by Thomas W Shinder <http://blogs.technet.com/b/tomshinder/archive/2011/04/19/does-removing-isatap-for-the-dns-block-list-impact-security.aspx>

(5) Ipv6 Security by Scott Hogg and Eric Vyncke <http://books.google.com/books?id=kwOv0Aw2IIUC&pg=PT424&lpg=PT424&dq=6to4+security+Windows&source=bl&ots=Qkn77LCVYi&sig=TRTNOXGEN0VpGHnYO L2qA6io6LE&hl=en&sa=X&ei=gb4FT96KleomosQKW9JyQCg&ved=0CE0Q6AEwAQ#v=onepage&q=6to4%20security%20Windows&f=false>

(6) Using HP ProLiant Network Teaming with Microsoft Hyper-V <http://h20000.www2.hp.com/bc/docs/support/SupportManual/c01663264/c01663264.pdf>

(7) IPv6 for Microsoft Windows: Frequently Asked Questions <http://technet.microsoft.com/en-us/network/cc987595.aspx>

(8) Configuring vPCs by Cisco http://www.cisco.com/en/US/docs/switches/datacenter/sw/4_2/nx-os/interfaces/configuration/guide/if_vPC.html

(9) Catalyst Switches for Microsoft Network Load Balancing Configuration Example http://www.cisco.com/en/US/products/hw/switches/ps708/products_configuration_example09186a0080a07203.shtml



New Horizons Course Offerings

(CCNAX) Interconnecting Cisco Networking Devices: Accelerated

CCNA

- (ICND1) Interconnecting Cisco Network
- (ICND2) Interconnecting Cisco Network
- (IINS) Implementing Cisco IOS Network Security
- (IIUC) Implementing Cisco IOS Unified Communications
- (IUWNE) Implementing Cisco Unified Wireless Networking Essentials
- (IUC) Implementing Cisco Unity Connection v8

CCNP

- (ROUTE) Implementing cisco Ip routing
- (SWITCH) implementing cisco switched networks
- (TSHOOT) troubleshooting and maintaining cisco ip networks
- (TVOICE) troubleshooting cisco Unified

CCNP Voice

- (CAPP5) Integrating Cisco Unified communications applications
- (CIPT1) Implementing Cisco Unified Communications IP Telephony Part 1 v7.0/6.0
- (CIPT2) Implementing Cisco Unified Communications IP Telephony Part 2 v7.0/6.0
- (CVOICE) Cisco Voice over IP v6.0
- (QOS) Implementing Cisco Quality of Service v2.3

CCNP Security

- (ASA) Upgrade v8.3 Cisco ASA Firewall
- (ASA) UPGRaDe v8.3 – Cisco ASA 8.x to 8.3+ Migration
- (Firewall) Deploying Cisco ASA Firewall solutions v1.0
- (IPS) – Implementing Cisco Intrusion Prevention System
- (SECURE) securing the cisco Routers and Switches
- (VPN) Deploying cisco ASA VPN solutions v1.0

Additional Courses

(ACUCM) v8.0 – Administering Cisco Unified Communications Manager
(ACEAP) v1.0 – Cisco Application Control Engine Appliance
(ACUCM w/auc) Administering Cisco Unified Communications Manager v7.0
(AUC) – Administering Cisco Unity Connection
(AUM) Administering Unified
(BGP) Configuring BGP on Cisco Routers
(CANAC) Implementing Cisco NAC Appliance
(CUWN) Cisco Unified wirelessNetwork
(CVPI) Cisco Unified Customer Voice Portal Implementation v7.0
(CWAAS) Cisco Wide-area Application Services
(CWLMS) v4.0 – Implementing CiscoWorks LMS
(DCNI2) Implementing cisco Data Center Network Infrastructure
(DCUCI) Data Center Unified Computing Implementation 3.0 (Exam N/A)
DCUCD v4.0 – Data Center Unified Computing Design
(IASNS) v4.2 – Implementing Cisco
(ICMBC) Intelligent Content Manager bootcamp
(ICOMM) Administering cisco voice and Unified communications
(ICSNS) v4.2 – Implementing Cisco
(IP6FD) – IPv6 Fundamentals and Deployment
(IPS) v.7 – Implementing Cisco Intrusion Prevention System
(MARS) – Cisco MARS
(MPLS) Implementing Cisco MPLS v2.3
(SNAA) Securing Networks with ASA Advanced v1.0
(SNAF) Securing Networks with ASA Fundamentals v1.0
(TUC) Troubleshooting Cisco Unified Communications Systems v1.0
(UCCE/ICM w/cvp) Unified contact Center Enterprise/Intelligent content Manager Administration
(UCCE/ICM w/ipivr) Unified contact Center Enterprise/Intelligent content Manager Administration
(UCCxa) Unified contact Center Express advanced
(UCCXD) Deploying Unified Contact Center Express v3.0
(UC-UCS) v3.0 – UCS for UC Engineers