




Networking and Double-Take Failover



Networking and Double-Take Failover published May 2002

NSI and Double-Take are registered trademarks of Network Specialists, Inc. All other products are trademarks of their respective companies. © 1996–2002 NSI Software

Double-Take Support for Application Failover

Double-Take's file system replication process is application independent and replicates any file system changes (including permissions and attributes) written to NTFS, FAT or FAT32 file systems by any application or process, subject to specific exceptions called out in the *User's Guide* or readme.txt file. Maintaining point-in-time consistent file system replicas and providing server monitoring and automatic or manual failover of the server name and IP address are the primary functions of the Double-Take software and we offer support to qualified customers should these functions fail to operate in accordance with our published documentation, regardless of what application or process is manipulating the data.

NSI Software may provide application notes and other documents that provide implementation guidelines on how to use Double-Take functions and replicas to manually or automatically failover or recover many popular third party applications and a general process to accomplish failover or recovery of many other third party applications. While these steps are believed to be accurate for the specific configuration, Double-Take version, and application versions originally tested, due to the number of possible configurations and variables, NSI Software can only test selected combinations and may provide only limited support for the operation and configuration of third party applications or the behavior of those applications before, during, or after failover, in its discretion. In cases where NSI Software has no direct access to or experience with a particular application or configuration, NSI Software support may also be limited to only the actual replication of the file system data and failover (name and IP address) of the server.

For assistance in validating, implementing or troubleshooting these or other possible configurations with third party applications, NSI Software and its partners may offer professional services on a fee basis to apply best practices for assisting with third party applications to recover automatically or manually using replicated data.

This, and any other, application note is provided solely for the convenience of our customers and is not intended to bind NSI Software to any obligation.

Table of Contents

Introduction	1
SMB and NetBT	1
Failover and WINS	3
WINS Management Scripting	4
Creating WINSCL Scripts	5
Failover and DNS	7
Name Caching	8
Double-Take Failover and ARP	9
Confirming Double-Take Failover	9
Replace Option	11
Double-Take Failover and Domain Controllers	11
Windows 2000	11
Windows NT	12
IP Address Failover to a Remote Target	12

Introduction

It is helpful to understand how Windows systems communicate when implementing a high availability solution using Double-Take's failover capabilities. Double-Take failover can be implemented with very little configuration necessary in small or simple networks, but additional configuration may be required in large or complex environments to ensure that all clients will resolve the source name to the correct IP address.

Understanding how Server Message Block (SMB) communications occur in Windows NT 4.0, Windows 2000, and mixed environments is the key to understanding what actions must be taken at failover in regard to DNS and WINS services. Due to the way SMB is implemented in Windows 2000, pure Windows 2000 environments should require no name resolution adjustments after failover. However, in Windows NT 4.0 or mixed environments, successful SMB communications require specific name resolution requirements to be met.


This document also discusses other networking topics related to failover, including how to view name caches and the ARP cache to troubleshoot. Additionally, common questions such as how failover affects domain controllers and how to fail over IP addresses to remote targets are addressed.

SMB and NetBT


Windows file and print sharing uses the SMB protocol, which has historically relied on NetBIOS. NetBIOS, in turn, required NetBIOS over TCP/IP (NetBT) to function on IP networks. NetBT uses TCP port 139 and has a limitation of binding only to the primary IP address of each NIC. This is explained in Microsoft Knowledge Base article Q131641, and can be seen by using a port scanner to probe TCP port 139 (the "nbssession" port) on an adapter with multiple addresses. This will show that NetBT is listening on TCP port 139 only on the primary address.

Windows 2000 and later versions do not require the NetBT layer and use SMB directly on top of TCP/IP using port 445 (TCP and UDP). This implementation does not have the aforementioned binding limitation and allows clients to establish SMB sessions to any IP address on the server using port 445. In order to be backward compatible with legacy clients and servers, Windows 2000 also supports SMB on NetBT using port 139, which inherits the primary IP address limitation. If NetBT is disabled, a Windows 2000 system will use only port 445 for SMB session.. See the following three diagrams.


Windows NT 4.0 and Earlier



Windows 2000 and Later with NetBT Enabled



Windows 2000 and Later with NetBT Disabled



The command `netstat -a -n` will show a Windows 2000 system listening on UDP 0.0.0.0:445 and TCP 0.0.0.0:445. An IP address of 0.0.0.0 indicates a global binding, so it is listening on TCP and UDP port 445 on all IP addresses for SMB communications. The NetBT binding will be shown as TCP 172.16.137.5:139 for the primary IP address on each NIC, indicating that it is listening only on those IP addresses.

Within Microsoft there is not a standardized way to refer to the implementation of SMB on port 445. Knowledge Base article Q150543 refers to it as "Direct Hosting of SMB Over TCP/IP." Article Q204279 refers to it as the "NetBIOS-less" transport, even while mentioning that the `NET CONFIG SERVER` command will report the binding as `NetbiosSmb (000000000000)`. Regardless, this document will refer to it as SMB/IP for the sake of convenience.

The NetBT bindings are displayed with the network adapter's GUID and MAC address as:

```
NetBT_Tcpip_{030319AA-84D2-481D-8E28-1EAC6D4AF2A8} (000102d03b7d)
```

The `netstat -a` command returns display names for ports 139 and 445 as `netbios-ssn` and `microsoft-ds` respectively.

Due to the differences between operating system versions, behavior of client/server communications after Double-Take failover will vary depending on the operating system version of the clients and server. When Double-Take failover occurs, the source IP address is added as a secondary IP address by default. Accordingly, if either the target server or clients are pre-Windows 2000 systems, one of the following is required in order for clients to access files on the target using the source's NetBIOS name:

- ◆ WINS and DNS servers must be updated so the source NetBIOS and host name resolve to a primary IP address on the target
- ◆ The IP address of the source must be made a primary IP address on the target (the "replace" option)

If all clients and the target server are Windows 2000 or later, WINS and DNS do not need to be updated if the source IP address is failed over. The clients, using SMB/IP, can connect to the target server with either the source or target IP address, so it does not matter which IP address the source name resolves to.

Microsoft Knowledge Base article Q142309 lists the order of name resolution methods used by clients when resolving a NetBIOS name:

1. NetBIOS name cache
2. WINS server
3. B-node broadcast
4. LMHOSTS file
5. HOSTS file
6. DNS server

Failover and WINS

When Double-Take failover occurs, Windows becomes aware of the new NetBIOS name that the server now owns and initiates a WINS registration with the target's primary WINS server. This registration associates the source NetBIOS name with the target's primary IP addresses and will be distributed to other WINS servers on the network via WINS replication. The length of time required for all WINS servers to obtain the new registration will depend on the number of WINS servers, the architecture of WINS replication, and the interval of replication. If there is only one WINS server or if the target and all clients that need access are configured with the same primary WINS server, then no additional configuration is necessary.

However, if some or all clients are configured with primary WINS servers other than the target's primary WINS server, failover scripts can be used to make the necessary WINS registrations on all WINS servers, or to initiate WINS replication. The first method, making the WINS registrations, incurs less network overhead, but will require the appropriate permissions (administrator group membership) on all WINS servers. The second method, initiating WINS replication, will rely on the WINS infrastructure to distribute the new records, but will require the system and network resources required to complete WINS replication. The impact of WINS replication will depend on the size of the network and the WINS architecture.

WINS Management Scripting

WINS registrations can be made via scripts as part of the failover process. This can be accomplished by using the Windows Resource Kit WINS Administration Tool (winscl.exe) or the Windows 2000 NETSH command in the failover script to either initiate WINS replication or make the appropriate registration with each WINS server.

WINS scripting is required when some or all clients have a primary WINS server other than the target's primary WINS server, the target server or any clients are pre-Windows 2000 systems, and the "replace" failover option is not used.

WINS scripting is not required when the target and all clients have the same primary WINS server (regardless of whether clients and target are in a LAN or WAN environment), the target server and all clients are Windows 2000 or later versions and the IP address is failed over, or the "replace" failover option is used.

Since Windows 2000 clients can use SMB/IP with any IP address on a Windows 2000 server, it does not matter if they resolve the source name to the source IP address or the target IP address. Both will work as long as the source IP address is failed over to the target. However, if the source IP address is not failed over (typically because the source and target are on different subnets), WINS servers must be updated at failover so that clients will resolve the source name to the target IP address.

Management of Windows NT 4.0 WINS servers can only be scripted with WINSCL, which includes two different methods of registration. The first way is to import an LMHOSTS file that is located on the WINS server. This method requires less scripting, but more files to manage since a separate LMHOSTS file must be maintained on each WINS server for each source server. If there are ten source servers and ten WINS servers, 100 LMHOSTS files, ten on each WINS server, will be required.

The second method is to script each registration individually. Although this requires only one script per source and does not require any files to be stored on the WINS servers, the scripts are much longer since a complete WINS registration for a member server actually has three registrations: SERVER[0x0], SERVER[0x3], and SERVER[0x20]. See Microsoft Knowledge Base article Q119495 for descriptions of names registered with the WINS service. (The SERVER[0x3] entry is a registration for the Messenger service, and can typically be omitted unless there is an application dependency.) See the following section for sample WINSCL scripts.

Windows 2000 WINS servers can be managed with the WINSCL utility or the Windows 2000 NETSH command, which supports numerous Windows 2000 IP management functions. NETSH can be used interactively or in scripts. The "add name" command in the WINS context will register SERVER[0x0], SERVER[0x3], and SERVER[0x20] with the specified WINS server. See Microsoft Knowledge Base article Q233375 for more information on adding WINS registrations with the NETSH command. Following is an example that adds a dynamic (RecType=1) registration for a server named EUROPA with an IP address of 172.16.137.5 to a WINS server with an IP address of 172.16.137.1.

```
netsh wins server 172.16.137.1 add name Name=EUROPA RecType=1 IP={172.16.137.5}
```

Using the NETSH command is the preferable method in Windows 2000 environments since each registration is simply a one-line command.

Creating WINSCL Scripts

Updating WINS servers with WINSCL scripts will require two LMHOSTS files on each WINS server that need to be updated, as well as text files containing the WINSCL commands that will be executed at failover and failback. An LMHOSTS file (TARGET_HOST) mapping the source's NetBIOS name to the target's IP address will be imported into the WINS database at failover, and another file (SOURCE_HOST) mapping the source's NetBIOS name to the source's IP address will be imported into the WINS database at failback.

The following sample LMHOSTS files use a source computer named PRODSVR with an IP address of 10.5.0.2, while the target server's IP address is 10.4.0.4. As you can see, the SOURCE_HOST file associates PRODSVR with its own 10.5.0.2 IP address, while the TARGET_HOST associates PRODSVR with the target's IP address of 10.4.0.4.

SOURCE_HOST

10.5.0.2	PRODSVR
----------	---------

TARGET_HOST

10.4.0.4	PRODSVR
----------	---------

There are seven WINSCL commands that must be scripted to import an LMHOSTS file into a WINS database. The following example, TARGET.DAT, includes the commands necessary to connect to a WINS server with an IP address of 10.5.0.11 and import the TARGET_HOST file that is located on the WINS server.

TARGET.DAT

I 10.5.0.11 SI I D:\TARGET_HOST 0 EX
--

For this example, the following command would be placed in the post-failover script to run the TARGET.DAT script.

POST_OVER.BAT

WINSCL < D:\SCRIPTS\TARGET.DAT

The SOURCE.DAT script would be used at failback to import the SOURCE_HOST file.

SOURCE.DAT

```
I
10.5.0.11
SI
I
D:\SOURCE_HOST
0
EX
```

The post-failback script would contain the following command to execute the script:

POST_BACK.BAT

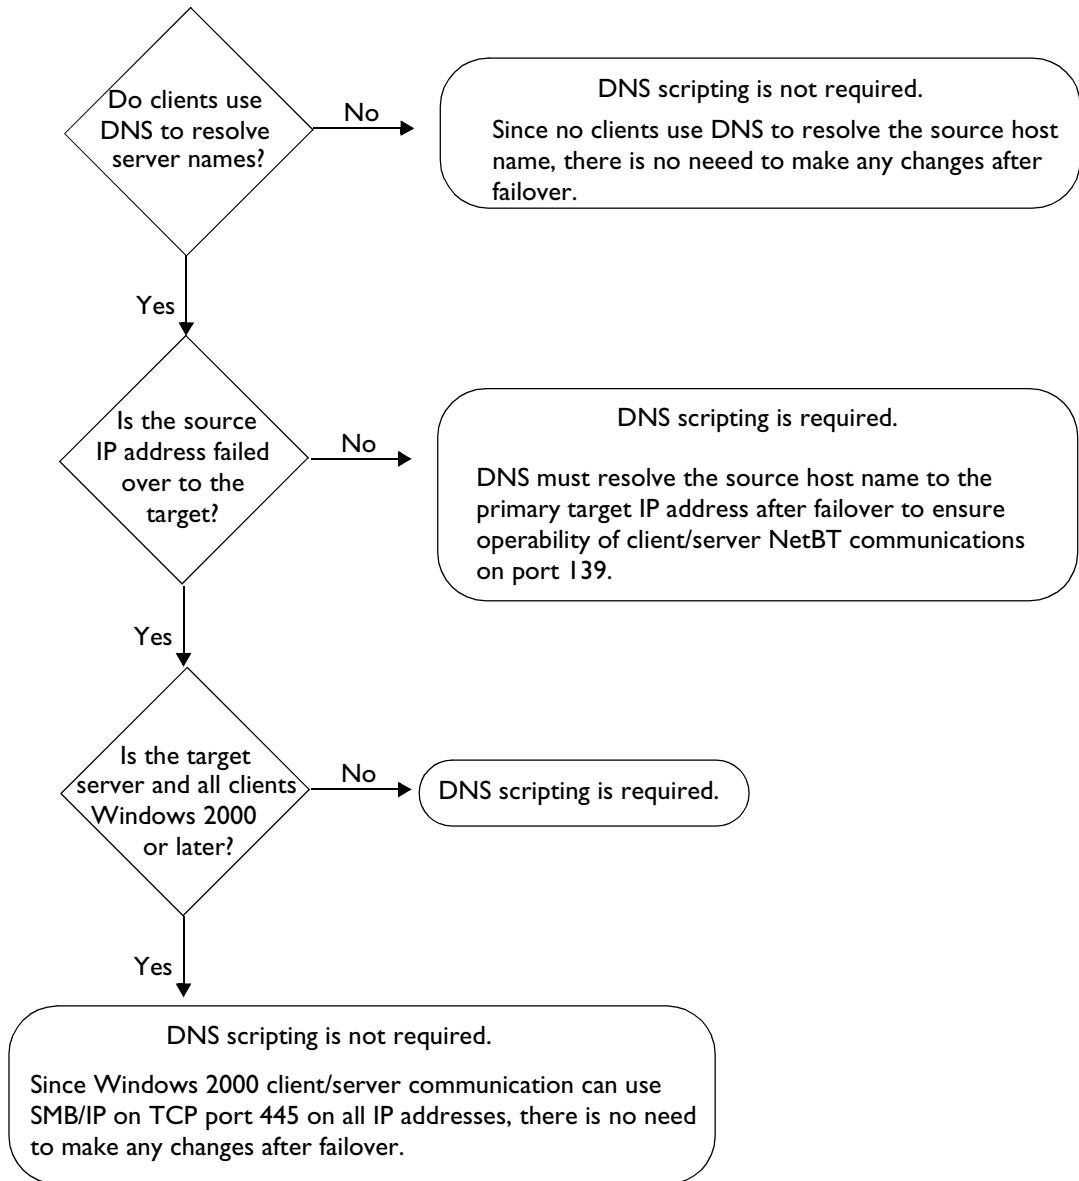
```
WINSCL < D:\SCRIPTS\SOURCE.DAT
```

NOTE: The LMHOSTS files must be located on the WINS server. The WINSCL utility does not actually import the LMHOSTS. It simply tells the WINS service to import the file. Accordingly, the WINS service interprets the path to the file (D:\SOURCE_HOST, for example), and the file must be at that location on the WINS server.

See the Failover chapter of the *Double-Take User's Guide* for information on how to configure failover scripts and the Windows NT 2000/4.0 Resource Kit documentation for WINSCL documentation.

Failover and DNS

DNS resolution is also a consideration after failover, especially when the source IP address is not failed over to the target. Use the following decision tree to determine when DNS host entries need to be changed after failover.



The DNS Server Troubleshooting Tool utility (DNSSCMD) from the Windows 2000 Support Tools can be used in the failover and failback scripts to delete and add host and reverse lookup entries so that the source host name will resolve to the target IP address. The following example commands delete the host and reverse lookup entries associating the host name "europa" with 192.168.15.14 in the nsissoftware.com zone on the DNS server (dnssvr.nsissoftware.com), and add the entries to associate europa with 192.168.15.18.

```
dnscmd dnssvr.nsisw.com /RecordDelete nsissoftware.com europa A 192.168.15.14 /f
dnscmd dnssvr.nsisw.com /RecordDelete 168.192.in-addr.arpa 14.15 PTR europa.nsisw.com /f
dnscmd dnssvr.nsisw.com /RecordAdd nsisw.com europa A 192.168.15.18
dnscmd dnssvr.nsisw.com /RecordAdd 168.192.in-addr.arpa 18.15 PTR europa.nsisw.com
```

DNSSCMD commands will only work if dynamic updates are enabled on the DNS zone. This is configured on the DNS zone's Properties dialog in the Windows 2000 Microsoft Management Console DNS snap-in. If only secure updates is enabled (this option is available only on Active Directory-integrated zones), the DNSSCMD utility must be used in the context of a user who is in the domain DnsAdmins group (i.e., the Double-Take service logon account must be in the DnsAdmins group if the commands are in failover/failback scripts). Failover/failback scripts do not run in the security context of the user specified in the failover monitor Account option despite the dialog's implication (in Double-Take 4.1) to the contrary.

The Windows 2000 Dynamic DNS (DDNS) client does not initiate a registration reflecting the failed-over name and IP address when failover occurs, and the ipconfig /registerdns command will not cause the failed-over name and IP address to be registered. Accordingly, host records for the source will remain intact after failover, and any required changes must be made on all DNS servers used by relevant clients.

Changes to non-Windows 2000 DNS servers and Windows 2000 DNS servers with dynamic updates disabled must be implemented by some other means. At this time, NSI Software does not have any specific documentation on how to script changes to DNS records other than the Windows 2000 DDNS solution. However, since DNS zone files are text-based they can be manipulated with any scripting language that can open, parse, and write to a text file.

Name Caching

By default, Windows systems cache DNS and NetBIOS name resolutions in the DNS Resolver Cache and NetBIOS Remote Cache Name Table respectively. This functionality does not impede the ability of clients to access the source name after failover as long as the appropriate WINS and DNS changes are made. Even in name-only failover scenarios (where the IP address is not failed over), clients will use WINS or DNS to re-resolve the name if an attempt to initialize a session with the cached entry fails.

Although name caching does not present any issues if WINS and DNS are updated after failover, information about the entries in the name cache may be useful in troubleshooting if difficulties are experienced. The cached entries can quickly show whether the client has resolved the source name to the correct IP address.

A system's NetBIOS Remote Cache Name Table can be viewed by using the nbtstat -c command, and the ipconfig /displaydns command displays the DNS Resolver Cache. The NetBIOS and DNS name caches can be purged with the nbtstat -R (the -R is case-sensitive) and ipconfig /flushdns commands respectively.

Microsoft Knowledge Base articles Q120642, Q245437, and Q187709 contain information regarding the configuration of timeout values for the NetBIOS and DNS name caches. However, keep in mind that the name resolution cache settings do not need to be adjusted for successful failover. If WINS and DNS entries are updated properly, previously cached name resolutions will not impede the ability of clients to establish SMB sessions with the target server.

Double-Take Failover and ARP

After TCP/IP has resolved the NetBIOS or host name to an IP address, it passes the IP packet to ARP (Address Resolution Protocol). ARP resolves the IP address to an adapter MAC (media access control) address and then passes the packet to the data-link layer (Ethernet, Token Ring, ATM, etc.). ARP maintains a cache of IP address-to-MAC address resolutions on each system. This cache must be updated at failover so that clients with cached entries will not attempt to send packets to the source's MAC address.

When Double-Take is installed, the NSI ARP Responder device driver is installed and set to a startup type of demand (Windows 2000) or manual (Windows NT 4.0). Double-Take uses the NSI ARP Responder to broadcast an unsolicited (gratuitous) ARP when failover occurs. The unsolicited ARP forces the systems on the same physical network to update their ARP caches with an entry associating the source IP address with the target adapter's MAC address. The following event will be created in the target's Application Event Log when the unsolicited ARP is broadcast:

Event ID: 5400

Source: Double-Take

Type: Information

Description: Broadcasted new MAC address [target adapter MAC address] for IP address [source IP address].

The `arp -a` command can be used to see the ARP cache on a system for troubleshooting purposes.

Confirming Double-Take Failover

When Double-Take executes a failover, it makes entries in the Double-Take logs, the Windows Event Log, and will generate an SNMP trap (if the SNMP component is installed). Operating system commands can also be used to show that the appropriate name and IP address changes have been made.

The following messages are logged in the target's Double-Take log file when failover occurs:

```
05/01/2002 13:34:14.5230 101 Failover in progress!!!
```

```
05/01/2002 13:34:44.2860 700001 Failover complete for SERVER1
```

A number of entries will be made in the Application Event Log as well:

Event ID: 5100

Source: Double-Take

Type: Information

Description: Failover completed for [source computername]

Event ID: 5101

Source: Double-Take

Type: Information

Description: IP address [source IP address] with subnet mask [source subnet mask] was added to target machine's [target adapter name] adapter.

If the SNMP service and the Double-Take SNMP component are installed, the target will generate the following SNMP trap when failover occurs:

Trap: DtttrapFailoverInProgress

Description: Failover is occurring.

The `nbtstat -n` command can be used to verify that the source's NetBIOS name is successfully added to the target after failover. This command shows the NetBIOS names registered on the local system. Following is a sample that shows `nbtstat -n` output on the target server CALLISTO before and after failover of the source server GANYMEDE (in the JUPITER domain).

Before Failover

```
D:\>nbtstat -n
TS:
Node IpAddress: [172.16.137.31] Scope Id: []

                NetBIOS Local Name Table

    Name                Type                Status
    -----
CALLISTO                <00>    UNIQUE            Registered
JUPITER                 <00>    GROUP             Registered
CALLISTO                <20>    UNIQUE            Registered
JUPITER                 <1E>    GROUP             Registered
INet~Services          <1C>    GROUP             Registered
IS-callisto...<00>    UNIQUE            Registered
```

After Failover

```
D:\>nbtstat -n
TS:
Node IpAddress: [172.16.137.31] Scope Id: []

                NetBIOS Local Name Table

    Name                Type                Status
    -----
CALLISTO                <00>    UNIQUE            Registered
JUPITER                 <00>    GROUP             Registered
CALLISTO                <20>    UNIQUE            Registered
JUPITER                 <1E>    GROUP             Registered
INet~Services          <1C>    GROUP             Registered
IS-callisto...<00>    UNIQUE            Registered
GANYMEDE                <00>    UNIQUE            Registered
GANYMEDE                <03>    UNIQUE            Registered
GANYMEDE                <20>    UNIQUE            Registered
```

Additionally, `nbtstat -a [target IP address]` can be run from a remote system to show the target's NetBIOS name table to retrieve the same information.

The `ipconfig` command can be used to verify that the source's IP address has failed over to the target. Failed-over IP addresses will be listed in the `ipconfig` output as well as in the network adapter's TCP/IP Properties Advanced dialog.

Replace Option

Double-Take includes a replace option that enables the target to replace its computer name and IP address with the source's computer name and IP address at failover. This option is supported on Windows NT 4.0 targets in the released versions of Double-Take, and in a hotfix version for Windows 2000.

When the replace option is used on a Windows NT 4.0 target, the Computer Browser, Net Logon, and Server services are stopped and started after the name and IP address is changed when failover occurs. Windows 2000 systems must be rebooted after the failover occurs if the replace option is used.

Since the source name and IP address remain the same when the replace option is used, there are no name resolution issues to consider as long as the source's IP address is the target's primary IP address after failover occurs.

Double-Take Failover and Domain Controllers

NSI recommends that solutions using Double-Take failover be implemented on member servers whenever possible. However, there are certain environments where the use of a domain controller as the source or target is unavoidable. Domain controllers can be successfully failed over with Double-Take when necessary, but the domain controller functionality is not failed over. Following is a discussion of the issues that should be considered when implementing a Double-Take failover solution with domain controllers.

Windows 2000

Windows 2000 Active Directory domain controllers uses a pull-based replication architecture, so there is no risk of Active Directory updates being sent to the wrong server due to Double-Take failover adding the source's computer name to the target. The only items to consider are the effects of a given domain controller being unavailable. A brief outline of some of these issues follows, but a complete understanding can only be gained by an in-depth knowledge of Active Directory. Active Directory documentation is included in the Windows 2000 Resource Kit.

Active Directory has five FSMO (Flexible Single Master Operation) roles, and each role is assigned to one domain controller in the domain or forest. Role ownerships can be easily moved between domain controllers to facilitate changes to the domain controller infrastructure. The five FSMO roles are:

1. Schema master (one per forest)
2. Domain naming master (one per forest)
3. PDC emulator (one per domain)
4. RID master (one per domain)
5. Infrastructure master (one per domain)

See Microsoft Knowledge Base article Q197132 for a concise description of these roles. Unavailability of some of the FSMO roles can cause immediate effects, such as a Windows NT 4.0 user not being able to change his password (PDC emulator), inability to extend the AD schema (schema master), and inability to add a domain to a forest (domain naming master).

Global Catalog (GC) servers are also critical for proper domain functionality, particularly the logon process in a multi-domain forest (see Microsoft Knowledge Base article Q216970). A properly designed Active Directory infrastructure will have multiple GC servers placed strategically throughout the network to ensure that the failure of a given GC server will not impact users.

Windows NT

Windows NT domains also use a pull-based directory replication architecture, so a failed-over BDC will not cause any inconsistencies in directory replication. The PDC prompts BDCs to request replication on a scheduled interval, and the BDCs then request updates from the PDC. The BDC's request informs the PDC of the last change it received, and the PDC sends the subsequent updates.

Again, there are no issues related to Double-Take, and the only concerns are those related to a PDC or BDC being unavailable for a period of time. A thorough understanding of Windows NT domains will be necessary to be aware of all possible issues, and the Windows NT 4.0 Resource Kit contains the relevant documentation. Some of the most important issues to consider include:

- ◆ Clients may be required to authenticate across a WAN link if a local domain controller is not available, which may cause a delay based on the available bandwidth.
- ◆ Extended downtime of the BDC may result in a full directory synchronization, which, depending on the size of the directory, can utilize significant bandwidth. By default, the PDC keeps a change log of 2000 entries, and a BDC will require a full synchronization if more than 2000 changes are made during its downtime. This is particularly a concern if the BDC is separated from the PDC by a WAN link.
- ◆ Some applications, such as Microsoft Exchange Server 5.5, require access to the PDC at start time.
- ◆ User and computer accounts, user rights, and other directory objects cannot be created, modified, or deleted if the PDC is unavailable.
- ◆ Trusts cannot be created if the PDC is unavailable.

IP Address Failover to a Remote Target

In most cases, failing over to a remote target can be accomplished by failing over the computer name only and updating the name resolution servers to associate the source name with the target's IP address. However, some solutions may require an IP address to be failed over to a remote target. Failing over IP addresses to remote targets can be accomplished a number of ways. If a VPN infrastructure exists so that the source and target can be on the same subnet, IP address failover will work exactly as it does in a LAN environment.

If a VPN does not exist, routers can be automatically reconfigured with a failover script after the IP address is failed over so that the failed over address will be routed to the target. This would entail configuring the routers to move the source's subnet from the source's physical network to the target's physical network. There are a number of issues to consider when designing a solution that requires router configuration to achieve IP address failover. Since the route to the source's subnet will be changed at failover, the source server must be the only system on that subnet, which in turn requires all server communications to pass through a router. Additionally, it may take several minutes or even hours for routing tables on other routers throughout the network to converge.

Depending on the router's capabilities, other options may also exist. Some routers can be configured to provide a routing infrastructure that can accommodate IP address failover to another segment. Additional discussion on this topic is beyond the scope of this paper due to the number of router manufacturers and various capabilities of router operating systems. If additional information is required, NSI Software's Professional Services division can be engaged to provide fee-based engineering services to analyse, design, and implement solutions.