

UNCLASSIFIED

Report Number: C4-051R-00

Microsoft Windows 2000® Network Architecture Guide

Systems and Network Attack Center (SNAC)

Author:
Paul F. Bartock, Jr.
Paul L. Donahue
Daniel J. Duesterhaus
Julie M. Haney
Prentice S. Hayes
Trent H. Pitsenbarger
1Lt Robin G. Stephens, USAF
Neil L. Ziring



Updated: April 19, 2001
Version 1.0

National Security Agency
9800 Savage Rd. Suite 6704
Ft. Meade, MD 20755-6704

410-854-6015
securew2k@dewnet.ncsc.mil

UNCLASSIFIED

UNCLASSIFIED

This Page Intentionally Left Blank

UNCLASSIFIED

Warnings

- **Do not attempt to implement any of the settings in this guide without first testing in a non-operational environment.**
- This document is only a guide containing recommended security settings. It is not meant to replace well-structured policy or sound judgment. Furthermore this guide does not address site-specific configuration issues. Care must be taken when implementing this guide to address local operational and policy concerns.
- The security changes described in this document only apply to Microsoft Windows 2000 systems and should not be applied to any other Windows 2000 versions or operating systems.
- SOFTWARE IS PROVIDED "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE EXPRESSLY DISCLAIMED. IN NO EVENT SHALL THE CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.
- This document is current as of April 19, 2001. See [Microsoft's web page](#) for the latest changes or modifications to the Windows 2000 operating system.

UNCLASSIFIED

This Page Intentionally Left Blank

UNCLASSIFIED

Acknowledgements

The authors would like to acknowledge the authors of the “*Guide to Implementing Windows NT in Secure Network Environments*” and the “*Guide to Securing Microsoft Windows NT Networks*” versions 2.0, 2.1, 3.0, 4.0, and 4.1.

Some parts of this document were drawn from Microsoft copyright materials with their permission.

Trademark Information

Microsoft, MS-DOS, Windows, Windows 2000, Windows NT, Windows 98, Windows 95, Windows for Workgroups, and Windows 3.1 are either registered trademarks or trademarks of Microsoft Corporation in the U.S.A. and other countries.

All other names are registered trademarks or trademarks of their respective companies.

Table of Contents

Warnings..... iii

Acknowledgements v

Trademark Information vi

Table of Contents..... vii

Table of Figures viii

Introduction 1

Getting the Most from this Guide 1

About the Microsoft Windows 2000 Network Architecture Guide 1

Chapter 1 Recommended Network Architecture Configuration 3

Network Description 3

Perimeter Security..... 5

Domain Controllers..... 8

Mail Servers 9

Web Servers..... 9

Network Support Services..... 10

Corporate Servers and Services..... 11

Appendix A Contained Domain Model Diagram 13

Appendix B Extended Domain Model Diagram..... 15

Appendix C Further Information..... 17

Appendix D References..... 19

Table of Figures

Figure 1 Contained Domain Model Diagram	13
Figure 2 Extended Domain Model Diagram	15

Introduction

The purpose of this guide is to inform the reader about the services that are available in the Microsoft Windows 2000 environment and how to integrate these services into their network architecture. This guide is not intended to provide individual security settings for the network devices. Instead, it is designed to provide the reader an idea of what functionality is recommended in each part of the network.

The ***Microsoft Windows 2000 Network Architecture Guide*** presents a general overview of the network and recommended network services. This overview is designed to show the recommended functionality in various locations within a network. The authors intend this guide to be used to help the planning phase of a network and it should not be used on its own as an all-encompassing blueprint for a network.



NOTE: This guide does not address specific security issues for the Microsoft Windows 2000 operating system or any of the other network operating systems or services mentioned.

This document is intended for Microsoft Windows 2000 network administrators and network designers. However, it should be useful for anyone involved with designing a network that includes Microsoft Windows 2000 hosts and/or servers.

Getting the Most from this Guide

The following list contains suggestions for successfully using the Microsoft Windows 2000 Network Architecture Guide:

- ❑ Read the guide in its entirety. Subsequent sections can build on information and recommendations discussed in prior sections.
- ❑ If applicable, compare the recommendations in this guide to the existing network architecture.
- ❑ Use a reasonable man theory when planning what a network needs:
 - Small networks will not need a separate device for each recommended function while large networks may need more than one host per function.
 - Implementing network devices without properly configuring them could lead to a more vulnerable network.
 - Network planning should include the necessary personnel to configure, manage, and monitor all the devices and hosts on the network.

About the Microsoft Windows 2000 Network Architecture Guide

This document consists of the one chapter and several appendices:

Chapter 1, “Recommended Network Architecture Configuration,” contains descriptions of the various hosts and network devices found in the diagrams in the appendices.

Appendix A, “Contained Domain Model Diagram,” contains a diagram that demonstrates the recommendations for a single domain that could be connected to the Internet.

UNCLASSIFIED

Appendix B, “Extended Domain Model Diagram,” contains a diagram that demonstrates several additional considerations for domains that spans one external network.

Appendix C, “Further Information,” contains a list of the hyperlinks used throughout this guide.

Appendix D, “References,” contains a list of resources cited.

Recommended Network Architecture Configuration

All the sections below reference the figures found in Appendix A and Appendix B.

Network Description

Contained Domain Model

A Contained Domain is a Microsoft Windows 2000 domain that does not extend across any networks not controlled by the organization. For example, a Protected Network could be considered a LAN protected by a firewall and a router. Any domain that does not extend past this firewall or router would be considered contained. Microsoft defines a site as any part of a network that is connected by a WAN link. Based on this definition, a Contained Domain is any domain that does not cover more than one site as long as the entire site is controlled by the organization defining the network architecture.

Extended Domain Model

An Extended Domain is a Microsoft Windows 2000 domain that extends beyond the Protected Network. Using Microsoft's site definition, an Extended Domain model represents any network that extends beyond the boundaries of one site. This type of network will need more functionality including domain controllers from higher in the domain tree located at lower branches' sites.

The Internet

Almost everyone today is aware of the Internet, but it is often confused with the World Wide Web (WWW, or "the web"). The web is part of the Internet, as are USENET (newsgroups) and ISPs (Internet Service Providers) such as AOL (America On Line), CompuServe and MSN (Microsoft Network). As the "World Wide Web" implies, the Internet spans the world. Anyone can connect to the Internet. Hackers, whether they are domestic or foreign, hobbyist, script-kiddies or professionals, scan systems looking for vulnerable computers. Once they find them, they take advantage of the vulnerabilities and can then deface web pages, take out servers (Denial of Service attacks), retrieve sensitive data or use the system to attack other systems. In short, the Internet is an environment that cannot be blindly trusted. The Internet provides a convenient medium to connect to other networks, but it does not provide reliable security features, such as user authentication and validation, server validation, or protection from hostile users or software.

In **Figure 1**, the Internet portion represents any network where the above activity can occur and not be controlled by the Protected Network's administrators.

Wide Area Network (WAN) Link

The WAN Link connects the external router to a backbone network. For example, the backbone network could be an ISP network on the Internet. The link is the boundary between the Protected Network and all external networks. It is also the point where the Protected Network administrator must share and coordinate control of network traffic with the ISP administrators. The choice of technology used to supply connectivity between geographically disjoint facilities is a function of distance, bandwidth, availability, and cost. With the emergence of affordable high-speed carrier technology, the choice needs to be tailored to operational requirements.

Protected Network

The Protected Network is connected to an interface of the firewall. This network includes the Internal Router and all of the subnets behind the Internal Router. These subnets comprise the internal backbone, the User Legs, the Legacy Leg and the Corporate Leg.

Demilitarized Zone (DMZ)

A DMZ is a small isolated network logically positioned between the Protected Network and an external network (e.g., the Internet). In the recommended configuration, the DMZ consists of a DMZ Router, which connects to an interface on the firewall, and the public servers (e.g., HTTP, Mail, DNS, FTP). Isolating the DMZ from the Protected Network provides a layer of defense from attacks launched from compromised public servers. However, the configuration may change if the administrators of the Protected Network do not control the firewall or the External Router. The External Router and the firewall are responsible for managing Internet access to the DMZ.

User Leg 1 and User Leg 2

The User Leg portions of the network contain user workstations. The workstations in these legs should be Windows 2000 Professional clients only. Multiple user legs will likely be established to account for logical separations in such areas as physical workstation/network locations, departments, user functions, or roles. In some environments it may be necessary to have a DHCP server, Domain Controller, and DNS server local to the user leg and not on the corporate backbone.

Legacy Leg

The Legacy Leg contains non-Windows 2000 client machines. These legacy machines could include Windows 9x and Windows NT 4.0 systems. Separation of legacy operating systems allows greater administrative control over non-Windows 2000 boxes and helps to maintain a clear picture of systems that may one day need to be upgraded. Within this leg, administrators can also maintain legacy services, protocols, and applications only needed for legacy Windows operating systems.

Corporate Leg

The Corporate Leg contains a number of Windows 2000 servers used internally by the corporate network. These servers may include an internal web server, certificate server, database server, mail server, domain controllers, and a network administrative host. In an extended domain model, other site domain controllers would be found in this area.

Perimeter Security

External Router

The External Router connects the Protected Network and DMZ to the WAN Link. The router provides the first opportunity to actively permit or deny access for clients and servers and for network services. For network traffic, the router can perform packet filtering and may be able to do some stateful inspection. Typically, this router acts as the screening filter. The screening filter provides a basic set of controls that do not change on a regular basis. These controls may include protection against the following types of attacks: IP address spoofing, denial of service attacks, and connections to unauthorized services. It is also important to log and analyze traffic that is denied access to the Protected Network and DMZ. This information can be used to better protect a network by identifying potential weaknesses.

Firewall

The best way to describe a firewall is to state what a firewall is not: a firewall is not simply a router, host system, or a collection of systems that provides security to a system. A firewall is not a single security solution, but it should be implemented as part of a defense in depth strategy. A properly configured firewall can reduce the risk of exposure to inherently insecure services. A firewall is simply an enforcer of a security policy. A firewall should reside at the perimeter of your network and protect your data from malicious entities. Firewalls can also control the availability of outside resources to the “trusted user”. Typically, a firewall resides on a separate machine, called a bastion host, and should be installed on a hardened operating system. Three interfaces are recommended: one for incoming traffic, one for access to the demilitarized zone, and one that connects to the Protected Network. The market place offers a wide variety of firewalls all claiming superiority over the other. But there are three traditional categories: packet filter, stateful packet filter, and application proxy. Vendors also supply hybrid firewalls providing the flexibility to deny packets at the network layer, IP address of an unwanted host or network, or the ability to check for malformed packets at the application layer. Firewalls should provide a degree of privacy. An example of privacy is the use of Network Address Translation (NAT) to protect an internal IP address scheme from the outside.

Most firewalls provide Virtual Private Networks (VPN) further enhancing privacy via encryption. Also, most provide comprehensive logging and auditing functions. These functions allow the administrator to determine if the firewall is withstanding probes or attacks, and if the rule base is configured properly. When implementing a firewall there is a risk involved with any service permitted into the protected network. A correctly configured firewall can help manage that risk.

DMZ Router

The DMZ router provides redundant protection. The only scenario where this router is necessary is if the firewall is not part of the Protected Network. The DMZ router filters on the services the DMZ provides, and denies all other traffic. If the firewall provides network address translation (NAT), the DMZ router will verify that all connections originate from the firewall.

Internal Router

The Internal Router provides an additional level of security against attack. The biggest threat to the Protected Network is from an insider. The internal router can be configured to protect the firewall and DMZ from internal attacks. The filters on the Internal Router should reflect the firewall policy providing an additional level of security. The rule base on this router should be configured to perform sanity checking (i.e., IP spoofing, etc) and all nonessential services on the router should be disabled. All recommendations for the external router should be incorporated into the Internal Router's rule base.

Intrusion Detection System

An Intrusion Detection System (IDS) is a security tool that gathers and analyzes information from various areas within a computer or a network to identify possible security breaches, assess system and file integrity, provide real time intrusion detection transparent to legitimate users, and provide response to unauthorized activity that blocks an intruder's access the network. The minimum installation recommendation is for a network based IDS placed inside the firewall on the internal backbone. However, as allowed by budget, IDSs are useful outside of the External Router, in the DMZ, and on each of the network legs. An IDS can be configured to detect internal and external threats.

An advantage of network based IDSs is that they do not affect the speed of the network or add load to the monitored system. Most network IDSs have comprehensive attack signature databases that can recognize known attack signatures. However, IDSs are similar to anti-virus products in that they cannot detect attacks until the database is updated. Also, IDSs struggle against low profile, time consuming attacks.

Another form of IDS is a host-based IDS. Host systems monitor their network connections and file system status. Host-based IDSs provide the ability to detect back doors in local systems, perform some degree of vulnerability assessment via inspection of configuration files, check password files for weak passwords and inspect other areas to detect policy violations. Host-based IDSs can be time and resource demanding. However, IDS logs can be stored and parsed at the local host or a centralized location. Since a host-based IDS only protects the host it is installed on, it is recommended to place host-based IDS on core hosts: DNS, Mail, Web servers etc.

IPSec Server

The IPSec Server provides the ability to protect traffic between hosts and sets of hosts within the network and identified hosts on other networks. For example, the IPSec Server could provide a secured channel for active directory replication between two domain controllers on different networks at different sites. Depending on network size and capacity requirements, the IPSec Server may be a separate network component or its functions may be provided by the Internal Router or the firewall. Because the IPSec Server protects traffic sent by other hosts, it must have the ability to support and maintain multiple distinct IPSec tunnels using a variety of configuration options. The security of an IPSec Server depends on strict authentication and authorization. Only specifically identified administrators should be able to log in to the IPSec Server or change its configuration. The IPSec Server should maintain audit logs of its operations.

Junction Devices

The physical Junction Devices can consist of a wide variety of physical network devices used to connect a local leg to the internal backbone. These physical network devices can be unintelligent port hubs, Ethernet switches, routers or firewalls.

The primary purpose of these junction devices is to provide the connectivity between the internal corporate network, the network backbone, and the local leg. The network device selection depends on the local leg size and type, the level of security required, and the amount of network traffic to traverse the junction box.

The Junction Connections consist of three basic types: User Leg Junction, Legacy Leg Junction and Corporate Leg Junction. The User Leg Junction is defined as the network device connecting a pure Windows 2000 Professional and Server LAN segment (Leg) to the internal network backbone. The Legacy Leg Junction is defined as the network device connecting a LAN segment containing a mixture of legacy Windows Operating Systems and other non-Microsoft Operating Systems to the internal network backbone. The Corporate Leg Junction is defined as the network device connecting corporate network servers to the internal network backbone.

Junction Connection properties

Authentication forwarding should be Kerberos V5 if available. If not using Kerberos for authentication, a server supporting the appropriate form of authentication should be placed on the local leg. Authentication verifies the identity of a user against a provided account and password. The Windows 2000 environment supports several different authentication methods. These authentication methods are Kerberos V5 certificate based authentications, NTLM, and NTLMv2.

The Junction Connection should allow all three types of Domain Name System (DNS) queries that a client or secondary DNS server can make from the local leg to an upstream DNS server. These DNS queries are recursive, iterative, and inverse. This functionality is necessary for Windows 2000 clients because they use DNS for name resolution and service location, including locating the domain controllers for logon.

While Windows 2000 uses DNS as the primary method for matching host name to IP address, Windows NT 4.0 and earlier operating systems use Windows Internet Name Service (WINS) for matching NetBIOS name to its IP address. These operating systems should be located on the Legacy Leg. Therefore, it is recommended to place a WINS server on the legacy leg and stop all NetBIOS traffic at the junction connection. However, if a separate WINS server is unavailable and interoperability is required with legacy operating systems that require NetBIOS naming to identify network resources, the Junction Connection must be able to pass WINS queries and responses or have configured LMHOSTS files on the legacy operating systems.

If using Dynamic Host Configuration Protocol (DHCP) to automatically configure TCP/IP addresses in the network, the Junction Connection should forward the DHCP clients' broadcast messages to the DHCP server and allow the DHCP server to respond back to the DHCP client. Note all DHCP communication is done over User Datagram Protocol (UDP) ports 67 and 68.

The following items are recommended if using a router, a firewall, or another network traffic blocking capable device such as a Junction Connection:

- Limit the general broadcast traversing the Junction Connection device in both directions.
- Block unnecessary protocols through Junction Connection.

- Block unnecessary services from traversing through the Junction Connection.

Domain Controllers

Windows 2000 uses a multi-master Domain Controller (DC) system rather than the NT 4.0 Primary Domain Controller (PDC) and Backup Domain Controller (BDC) system. A Domain Controller holds part (or all) of the Active Directory for the domain. All user information, including account information, group memberships, etc., is contained in the Active Directory. Domain Controllers are different from stand-alone servers since stand-alone servers do not contain any part of the Active Directory. Providing multiple Domain Controllers on a network allows redundancy. In **Figure 1**, there are three different Domain Controllers shown which are described below.

Corporate Domain Controller

The Corporate Domain Controller represents domain controllers and services from domains not located within the Protected Network. Any network that is not part of a forest or tree and has no external trust relationships will not have this server on the network. These servers provide the following functionality:

- Users from other domains can log in to their local Domain Controller without sending network traffic outside the Protected Network.
- Replication between the Corporate Domain Controllers can be scheduled for off-peak hours.

The Corporate Domain Controllers should only come from trusted domains. Furthermore, traffic from these Domain Controllers outside the Protected Network should be sent through an encrypted tunnel as shown in **Figure 2**.

Domain Controller 1

Domain Controller 1 represents the first domain controller installed in a domain. By default this domain controller has special functionality including schema master, PDC emulator, etc. Therefore, it is recommended that this host have more protection and not be used for simple user authentication, but for domain administration purposes. If the Protected Network is a site in a larger domain, this domain controller will be replaced with a Corporate Domain Controller. If any of the master tokens, like schema master, have been moved from this domain controller then extra protection should be placed on that domain controller as well.

Domain Controller 2

Domain Controller 2 represents any domain controller that is not the first one installed in a domain. This domain controller could be a member of any domain that has a domain controller on the corporate leg. Domain Controller 2 should be placed so clients can authenticate to this server. However, it should also be configured to only replicate with the domain controllers found on the corporate leg. This setting will limit the replication traffic between protected networks.

Mail Servers

Mail Forwarder

There are several risks associated with the receipt of e-mail from potentially untrusted entities outside the site. Chief among these concerns are attacks against the recipient e-mail server itself. Examples of this include attempts to exploit buffer overflows and content driven attacks in the form of malicious code.

The Mail Forwarder is simply a mail server that forwards e-mail messages intended for internal users to the internal mail server and accepts mail destined for the external network for delivery. As the Mail Forwarder is the only mail server that is exposed to the external network, the risks associated with e-mail are reduced by precluding direct access to the internal mail server. This Mail Forwarder, as should be the case for all servers in the DMZ, must not be a member of any internal Windows 2000 domains. This will limit the damage that could result from its compromise. Content checking, initial virus scanning, and filtering ideally should also be performed here to guard against malicious code.

Internal Mail Server

The Internal Mail Server is utilized by users within the site to both send and receive mail. It should be configured to send all outgoing mail to the mail forwarder. The risks associated with the internal mail server are similar to that described for mail forwarder. Again, content checking and filtering capabilities are critical countermeasures. Additionally, a large concern is preserving data confidentially and integrity by utilizing user authentication and access control mechanisms to limit users to content for which they are authorized access (e.g., their own mailbox). Data encryption can also be utilized for sensitive data with the common options including the S/MIME standard for reader-to-writer data protection or SSL for protecting data in transmission between the client and server (e.g. SMTP over SSL and IMAP over SSL).

Web Servers

External Web Server

Organizations frequently have data they want to publish to the external network via a web server. A web server in the DMZ should be utilized to provide this functionality in lieu of allowing direct connectivity to an internal server. However, there remains a security concern regarding unauthorized data access. In some applications it is necessary to partition data between users, and in all applications it is necessary to protect the web site from unauthorized modification. Access control mechanisms ranging from simple passwords to certificate based authentication can be used to segment user data. Using secure protocols like SSL can also be used to secure communications to the server.

If a FTP server is used in conjunction with the web server, it is recommended that it be setup as read-only. If it is necessary to allow users to write to the FTP server, then the use of a drop-box is recommended. A drop-box is a directory on the FTP server that is configured for write-only access. An administrator then reviews each submission to ensure that uploaded files do not contain inappropriate content and utilizes a virus scanner as a malicious code countermeasure. Only then are the files moved to other directories where other users can access them.

A database server is also commonly used in conjunction with the web server. Here again the primary security concerns relate to ensuring that individuals only have access to the data for which they are authorized in order to preserve data confidentiality and integrity.

Internal Web Server

The Internal Web Server is used to publish data via HTTP to users within the site. External users should never be granted access to the internal web server. Instead any data intended for external consumption should be published on the external web server. Despite the fact that external users do not connect, the concerns related to the internal web server are very similar to that for the external web server. For instance, it is still important to prevent unauthorized access from users within the site. The internal web server has the potential advantage of being in a more homogenous environment, which would allow the use of Windows authentication mechanisms. FTP servers and database servers exposed via HTTP, if utilized, should follow the guidelines offered for the External Web Server.

Network Support Services

External DNS Server

The External DNS Server resides in the DMZ and is for public use. The only information that should be on the External DNS Server is the information that needs to be advertised to Internet clients. No internal information should be available to the External DNS Server. Furthermore, only the External DNS Server can communicate with the outside. The Internal DNS Server should handle resolution of external DNS information for the firewall and all hosts on the Protected Network. Although the Internal DNS Server is unable to communicate directly with the external network, it should be configured to send queries and receive the responses via the External DNS Server.

For proper functionality, DNS UDP connections should be allowed through the firewall to the External DNS Server. These connections allow external clients to query the DNS. However, DNS TCP connections should only be allowed between the External DNS server and trusted DNS servers located beyond the External Router. All other DNS TCP connections should be denied. This setting reduces an attacker's ability to map the Protected Network via zone transfers. Also, note that even if DNS information is hidden there are other sources that will provide information about the internal naming scheme. For instance, email headers, NetBIOS, and other services could still supply host information to an attacker.

Internal DNS Server

The Internal DNS Server provides the functionality required by Windows 2000 for registering domain services and clients in the domain. This server is extremely important because it holds a complete map of the domain. Therefore, this server should not share any information with the External DNS Server. Since there may be requirements for reverse lookup zones to include clients in the Protected Network, the reverse lookup zone on the internal server may need to be passed to the External DNS Server. The functionality of this server may be included on a domain controller. However, if DNS data is required externally, a specific, non-domain controller based DNS should be established to communicate the information to the external server. Like domain controllers, the better connection clients have to the internal DNS server, the better the network will perform.

Corporate Servers and Services

Admin Host

The Admin Host resides in the corporate leg of the network, and it acts as a central workstation for many network security administration tasks. In a large network, there may be several Admin Host workstations. The Admin Host should be configured to allow administrators to log in, conduct audit reviews, configure DNS and DHCP, administer the web and e-mail servers, and maintain group policy. Depending on the mechanisms in use, the Admin Host may also act as a central control point for performing backups. Where possible, the Admin Host should hold centralized audit logs for important parts of the network, such as domain controllers, routers, switches, intrusion detection systems, and critical servers (e.g., DNS, DHCP servers). Typically, the Admin Host will be one or more dedicated Windows 2000 hosts, but in a small network the functions of the Admin Host may be supported by one of the other hosts on the corporate leg. The Admin Host must be configured to restrict access to authorized administrators.

Certificate Server

The Certificate Server is used to issue public key certificates to users and computer accounts. Certificates are commonly used as an authentication mechanism (e.g., public key based authentication) and to support data confidentiality (e.g., SSL). The primary security concerns relate to certificate issuance policies and protecting the integrity of the certificate authorities. A certificate used for authentication is not any more secure than the procedures used to verify the certificate was issued to the correct person. Windows 2000 offers mechanisms that tie certificate issuance to account information contained in the Active Directory, helping to ensure that certificates are appropriately bound to user accounts. Protection of the integrity of the certificate authority is primarily a function of appropriately configuring the Certificate Server and limiting physical access.

Remote Access Server

The Remote Access Service (RAS) allows users to dial into your network. RAS provides external users the ability to remotely access networks and should not be used unless absolutely required by operations. However, there is information available on securely connecting to networks using the Remote Access Security Program (RASP), which can use devices such as FORTEZZA modems for dialing into networks. A commercial site for more information on this technology is <http://ias.itse realm.com/rasp>. It is especially critical to properly secure Remote Access Servers as they provide a backdoor to a secure network.

Corporate Database Server

The Corporate Database Server functions as a data repository that may be accessed by users within the site. The primary security concern related to database servers is protecting data from unauthorized access. Database servers typically support user authentication and the assignment of users to specific roles based upon their needs, ranging from a simple user role to a database administration role. Furthermore, one can typically assign access rights to database objects based upon these roles (e.g., read access for all users to a table) or for specific users (grant user x write access to a table). A finer level of access control, typically down to the table column, is possible with some

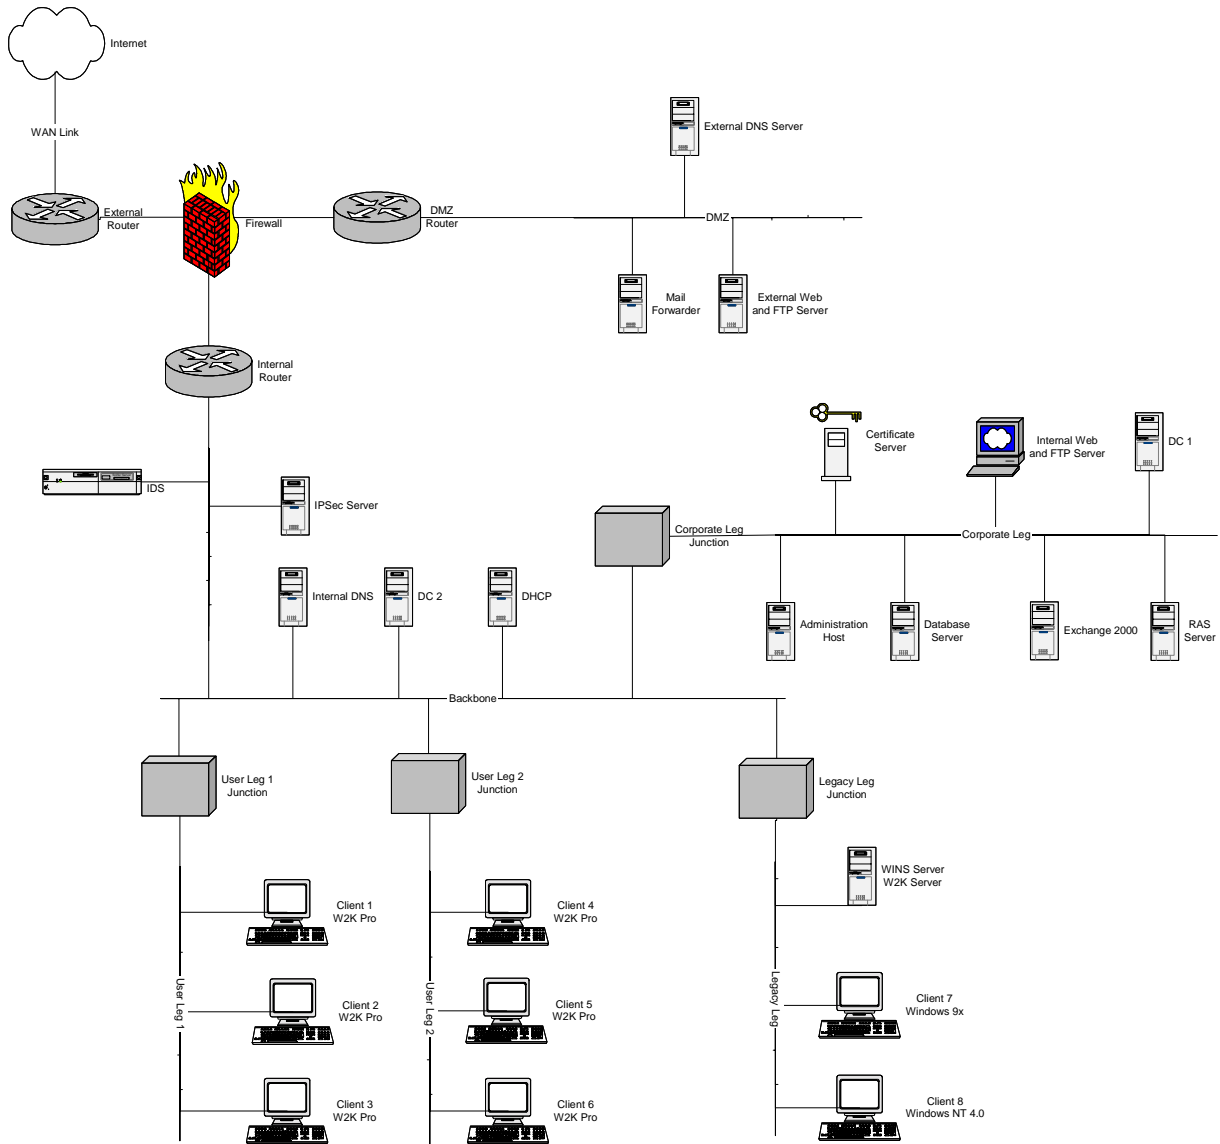
UNCLASSIFIED

applications. Encryption can be used to protect data in transit between the server and client.

If publishing database information via HTTP, the configuration of the web server is critical to the security of the database information. For further information reference the External Web Server section.

Appendix
A

Contained Domain Model Diagram



Appendix A – Contained Domain Model Diagram

Figure 1 Contained Domain Model Diagram

UNCLASSIFIED

This Page Intentionally Left Blank

Appendix A – Contained Domain
Model Diagram

UNCLASSIFIED

Extended Domain Model Diagram

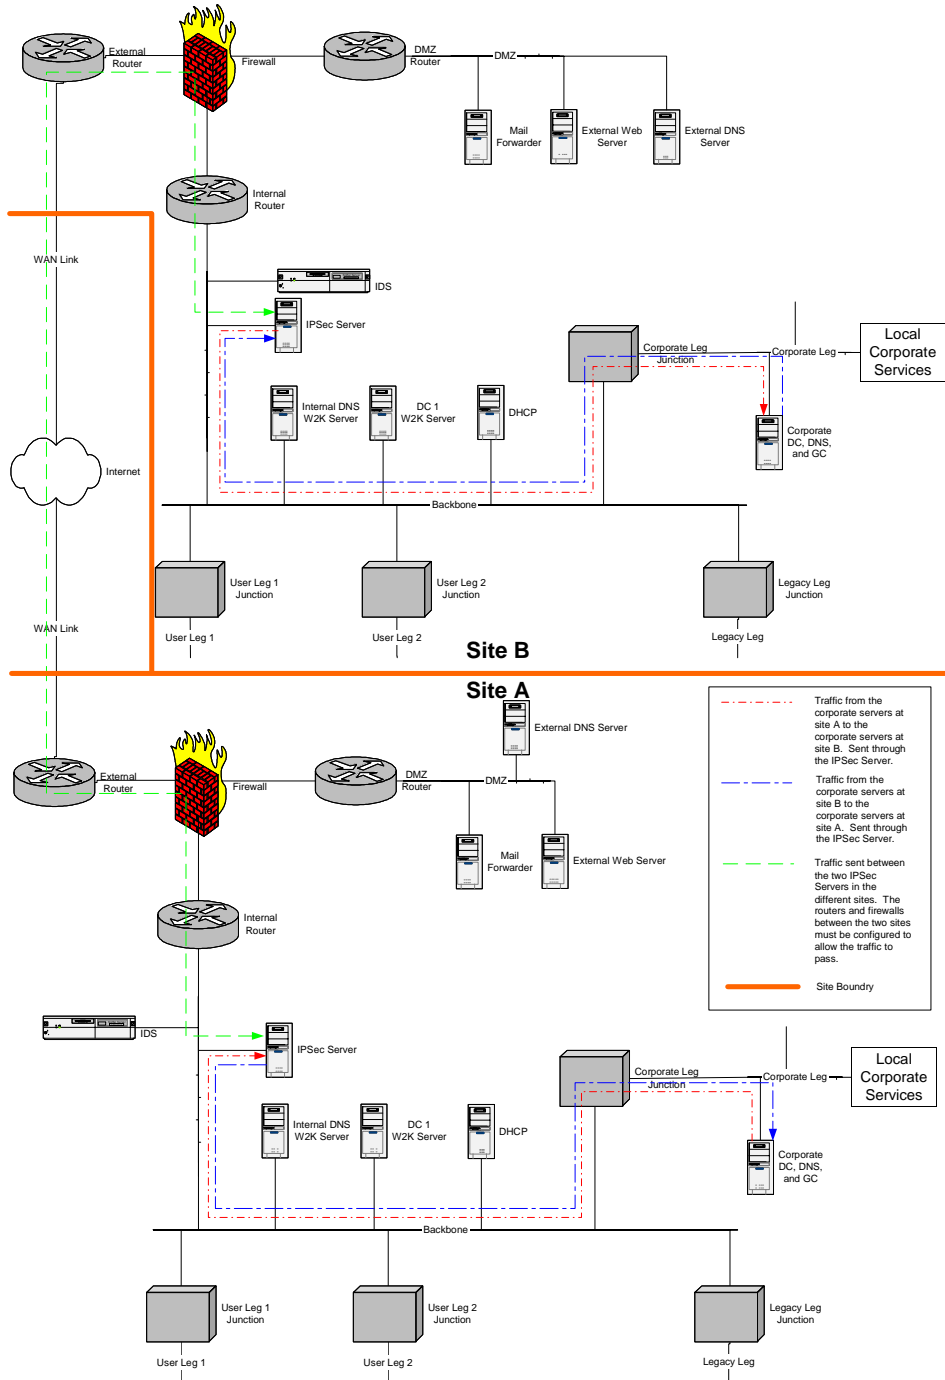


Figure 2 Extended Domain Model Diagram

UNCLASSIFIED

This Page Intentionally Left Blank

Appendix B – Extended Domain
Model Diagram

UNCLASSIFIED

Further Information

Microsoft Corporation. "Designing a Site Topology for Active Directory Replication."

http://www.microsoft.com/windows2000/library/resources/reskit/deploymentscenarios/scenarios/repl_design_sitetopology_active_directory_repl.asp.

Microsoft Corporation. "Deployment Lab Scenarios."

<http://www.microsoft.com/windows2000/library/resources/reskit/deploymentscenarios/default.asp>.

Microsoft Corporation. "DNS Requirements for Deploying Active Directory."

<http://www.microsoft.com/technet/win2000/win2ksrv/dnsreq.asp>.

Microsoft Corporation. "Enabling Resource Sharing Between Windows 2000 Forests."

http://www.microsoft.com/windows2000/library/resources/reskit/deploymentscenarios/scenarios/trust_enable_resshare_w2000_forests.asp.

Microsoft Corporation. "Windows 2000 Advanced Server Help."

<http://windows.microsoft.com/windows2000/en/advanced/help/>.

NSA Windows 2000 Guides pertaining to specific security settings and options.

Riley, Steve. "Active Directory Replication over Firewalls."

<http://www.microsoft.com/TechNet/win2000/win2ksrv/adrepfir.asp#c>.

UNCLASSIFIED

This Page Intentionally Left Blank

UNCLASSIFIED

References

[Microsoft's Web Site.](#)

Microsoft Corporation. Designing a Microsoft Windows 2000 Directory Services Structure: Class Pack. February 2000. Part Number: X05-39392.