



**DRAFT**

**DOMAIN NAME SYSTEM  
SECURITY TECHNICAL IMPLEMENTATION GUIDE**

Version 3, Release 0

23 January 2006

**Developed by DISA for the DOD**

UNCLASSIFIED

This page is intentionally left blank.

## TABLE OF CONTENTS

	Page
TABLE OF CONTENTS.....	iii
SUMMARY OF CHANGES.....	viii
1. INTRODUCTION .....	1
1.1 Background.....	1
1.2 Authority.....	1
1.3 Scope .....	2
1.4 Writing Conventions.....	2
1.5 DISA Information Assurance Vulnerability Management (IAVM).....	3
1.6 STIG Distribution .....	3
1.7 Document Revisions.....	3
1.8 Vulnerability Severity Code Definitions .....	4
2. DNS TERMINOLOGY AND CONCEPTS .....	5
2.1 Domains.....	5
2.2 Zones .....	6
2.3 Name Servers.....	7
2.3.1 Authoritative Name Servers .....	7
2.3.2 Caching Name Servers .....	8
2.4 Name Server Support Roles .....	8
2.5 Resolvers .....	9
3. DNS ARCHITECTURE AND GENERAL SECURITY REQUIREMENTS .....	11
3.1 Name Server Software.....	11
3.2 Name Server Operating System Platforms .....	12
3.3 Redundancy, Dispersal, and Availability .....	13
3.3.1 Network related availability .....	14
3.3.2 Stub Zones.....	15
3.4 Authentication and Access Control .....	15
3.4.1 Zone Updates.....	16
3.4.1.1 Updates of Individual Records .....	16
3.4.1.2 Zone Transfers.....	17
3.4.2 Query Restrictions.....	18
3.4.2.1 Split DNS .....	18
3.4.2.2 Operational Security.....	19
3.4.2.3 Restrictions on Recursion.....	19
3.4.3 Firewalls and DNS .....	20
3.5 Logging.....	20
3.6 Zone Files .....	21
3.6.1 Change and Ownership Documentation.....	21
3.6.2 Zone-Spanning Records and Glue.....	22
3.6.3 Improper NS Records and Lame Delegation .....	23

3.6.4 Root Hints.....	24
4. CONFIGURATION OF BIND.....	26
4.1 Obtaining Software.....	26
4.2 Configuration of the named.conf File.....	27
4.2.1 key Statement .....	28
4.2.2 controls Statement and the rndc.conf File .....	29
4.2.3 server Statement .....	31
4.2.4 acl Statement .....	32
4.2.5 logging Statement.....	32
4.2.6 options Statement .....	36
4.2.6.1 version Option .....	37
4.2.6.2 Options Related to Query and Recursion Restrictions .....	38
4.2.6.3 Options Related to Zone Update Notification .....	39
4.2.6.4 IQUERY Simulation .....	40
4.2.6.5 RFC 2308 .....	40
4.2.7 view Statement and Split DNS .....	41
4.2.8 zone Statement .....	42
5. CONFIGURATION OF WINDOWS 2000/2003 DNS.....	45
5.1 Comparison of Windows 2000/2003 DNS and BIND .....	45
5.1.1 BIND 9 Security Features Not Present in Windows 2000 DNS .....	45
5.1.2 Windows 2000/2003 Security Features Not Present in BIND 9 .....	46
5.1.3 Implications of BIND 9/Windows 2000/2003 DNS Differences.....	47
5.2 Network Restrictions .....	48
5.3 Secure Dynamic Updates and Active Directory .....	48
5.4 Zone Transfers.....	51
5.5 Forwarders and Recursion .....	53
5.6 WINS Integration .....	55
5.7 Logging.....	58
6. CISCO CSS DNS.....	60
6.1 CSS DNS Background.....	60
6.2 CSS DNS Zone Delegation .....	61
6.2.1 CSS DNS Childhood .....	61
6.2.2 CSS DNS Parenthood.....	63
6.3 Simulation of Zone Transfer Restrictions in CSS DNS .....	64
6.4 Simulation of Query Restrictions in CSS DNS .....	65
6.5 Recursion and CSS DNS Forwarders.....	65
7. STANDARD OPERATING PROCEDURES FOR DNS .....	68
7.1 Security Management Responsibilities .....	68
7.1.1 Personnel .....	68
7.1.2 Physical Access Control.....	68
7.1.3 Business Continuity.....	69
7.2 System Administration Responsibilities.....	69
7.2.1 Software Patches .....	69

7.2.2 Backup.....	70
7.3 DNS Software Administration Responsibilities .....	70
7.3.1 DNS Configuration File Changes.....	70
7.3.2 Cryptographic Key Supersession .....	71
7.3.3 Log Review .....	71
7.4 DNS Database Administration Responsibilities.....	72
8. DNSSEC .....	74
8.1 DNSSEC Mechanisms and Operations .....	75
8.1.1 Sign and Serve.....	75
8.1.2 Verify Signature .....	75
8.2 Generation of Public Key-Private Key Pair (DNSSEC-OP1) .....	76
8.2.1 Key Pair Generation—Illustrative Example.....	78
APPENDIX A. RELATED PUBLICATIONS.....	82
APPENDIX B. INFORMATION ASSURANCE VULNERABILITY MANAGEMENT (IAVM) .....	84
APPENDIX C. GENERATION AND STORAGE OF TSIG KEYS.....	86
APPENDIX D. MODEL STANDARD OPERATING PROCEDURES FOR DNS DATABASE ADMINISTRATION .....	90
APPENDIX E. UNIX OPERATING SYSTEM CONFIGURATION TO SUPPORT BIND .....	93
E.1 Network Services .....	93
E.2 Non-privileged User for BIND .....	94
E.3 File Permissions .....	95
E.3.1 BIND Support Files .....	95
E.3.2 TSIG Keys .....	96
E.3.3 Log Files .....	96
E.4 Running BIND in chroot(ed) Directory Structure .....	96
APPENDIX F. MICROSOFT WINDOWS OPERATING SYSTEM CONFIGURATION TO SUPPORT BIND .....	99
F.1 Network Services .....	99
F.2 Non-Privileged User for BIND .....	100
F.3 File Permissions .....	102
APPENDIX G. ROOT HINTS .....	105
APPENDIX H. LIST OF ACRONYMS.....	107

## LIST OF TABLES

Table 1-1. Vulnerability Severity Code Definitions .....	4
Table F-2. Windows File Permissions .....	102
Table G-1. Root Hints .....	105

## TABLE OF FIGURES

Figure 2-1. Domain Tree .....	5
Figure 2-2. Domains and Zones .....	6
Figure 5-1. Change Zone Type .....	49
Figure 5-2. Dynamic Updates .....	50
Figure 5-3. Zone Transfers Tab .....	52
Figure 5-4. Forwarders Tab .....	54
Figure 5-5. Advanced Tab .....	55
Figure 5-6. WINS tab .....	57
Figure 5-7. Logging Tab .....	58
Figure F-1. Service Account .....	101

This page is intentionally left blank.

## SUMMARY OF CHANGES

Version 3, Release 0.1 of the Domain Name System (DNS) STIG includes numerous changes relative to the last release Version 2, Release 2, dated March 04, 2005.

### General:

All references to BIND versions have been replaced with BIND 9.3.1 or later or BIND 9.3.x.  
All sections in the document were reworded and edited.  
Added Table and Figure Labels.

### Section 1:

Applied general changes to Section 1 in an attempt to update for consistency with other STIGs.  
Changed the acceptable BIND version levels.  
Section 1.3: Changed Scope to include all DOD authoritative and recursive name servers.  
Section 1.8: Added DNS applicability to the Vulnerability Severity Code Definitions

### Section 2:

Provided additional clarification on caching and authoritative name servers.

### Section 3:

Rewording and grammatical changes made throughout the section.  
Changed the requirement for OSs to coincide with the appropriate OS STIG.  
Removed Windows NT operating system as a viable platform for DNS server software.  
Added DNS0160, CAT III on documentation requirements for the DNS architecture  
Added DNS0170, CAT II requirement to ensure underlying operating system is STIG compliant  
Added verbiage to refer to the appropriate OS STIG for currently supported OS types.  
Added DNS0175: CAT I: DNS server software will not reside on an OS that is no longer supported by the vendor.  
Added DNS0180, CAT II, requirement for HIDS on a DNS server.  
Added a new section, 3.2.2, on Stub Zones provided by the DOD DNS Working Group  
Section 3.4.2.1 Split DNS. Substantial change. Split DNS is no longer a requirement due to operational considerations, unless using RFC 1918 or private IP addresses.  
Section 3.4.1.2 Changed DNS0460 from a CAT III to a CAT II.  
Section 3.4.2.2: New section: Operational Security and added DNS0185: CAT III requirement for the DNS administrator to review the contents of the zones they are responsible for, at a minimum, annually.  
Section 3.4.3: Added information on firewalls and the use of TCP port 53  
Section 3.6.1 Changed DNS0225 from a CAT III to a CAT II finding.  
Section 3.6.2: Changed DNS0235 from a CAT IV to a CAT III finding and reworded to include "zones with lesser security".

### Section 4:



Added a DNS0190, CAT III requirement to subscribe to “bind-announce” if BIND is utilized.  
Added a DNS0195 CAT III requirement for TSIG minimum key length.  
Added a DNS0250, CAT III requirement for new TSIG key generation for each type of transaction (zone transfer, dynamic updates, etc)

## **Section 5:**

Added additional information/clarification on Windows Server 2003 throughout the section.  
Added information on vulnerabilities associated with DHCP. Added two new requirements to section 5.3: DNS0255, CAT II – the disablement of DHCP on any Windows Domain Controller; and DNS0260 CAT II, ensuring computer accounts for DHCP servers are not members of the DNSUpdateProxy group.

## **Section 8:**

This section was completely rewritten by the DOD DNS Working Group in collaboration with the NIST Draft Special Publication 800-81, Secure Domain Name System (DNS) Deployment Guide.

## **Appendix A:**

Added the Draft NIST Special Publication  
Added the CIS DNS BIND Benchmark

## **Appendix B:**

Updated IAVAs to include 2005-A-0005: Multiple Vulnerabilities In BIND and 2003-B-0001 Multiple Buffer Overflow Vulnerabilities in Various DNS Resolver Libraries

## **Appendix E:**

Added information on the CIS DNS BIND Benchmark.  
Changed DNS4450 to the following: The SA will ensure all network services not required for operations are disabled. Any network service that is utilized is documented with the IAO with complete mission justification and vulnerability mitigation.  
Added “BIND will not run as a root user” to DNS4440.  
E.3.1: Added that the *named* process will not have write access to any configuration files, key files, etc.  
E.4 , DNS4445 CAT III, Added a new requirement for BIND to be configured in a chroot(ed) directory structure.

## **Appendix F:**

F.1: Added the word “required” to the disabling statement for unnecessary services, per the Windows STIGs.

F.2: Changed DNS4540 from a CAT III to a CAT II requirement.

## **Appendix H:**

Edited the Acronym List.

This page is intentionally left blank.

## 1. INTRODUCTION

*The Domain Name System (DNS)* is a hierarchically structured, distributed database that provides name resolution services to Internet Protocol (IP) networks such as the Internet. The most common use of DNS is to map a host name to its network address. For example, the host name `www.cert.mil` is mapped via DNS to `199.211.123.12`, the dotted decimal notation for its 32-bit IP address.

DNS also supports reverse lookup, in which a network address is mapped to a host name. Other supported directory capabilities include lookup of name servers and mail exchangers. Newer technologies such as IPv6, Microsoft's Active Directory, and wireless Internet also depend on DNS for basic functionality.

The security posture of DNS within the Department of Defense (DOD) is of concern for several reasons. The data it provides must be highly reliable and available to support both peacetime and wartime requirements; and the dissemination of DNS information must not itself introduce security risks into the DOD computing environment. Specifically, mechanisms must exist for name servers to verify the accuracy and origin of DNS information received.

As DNS has many security weaknesses, yet is also a critical infrastructure component, it is a major potential target for adversaries. Early implementations of DNS did not provide for transmission authentication or guarantee the integrity or accuracy of the global database, which left the DNS service open to attacks. More recent versions of DNS software provide greater information assurance (IA), but it must be appropriately configured to achieve this assurance. Without reliable DNS, nearly all DOD information systems are left vulnerable, including mission critical applications.

This *DNS Security Technical Implementation Guide (STIG)* is designed to assist administrators with the configuration of DNS server software and related portions of the underlying operating system. This STIG also provides guidance for standard operating procedures related to configuration management, business continuity, and other topics.

### 1.1 Background

This document supersedes previous DNS guidance found in the Network Infrastructure and UNIX STIGs. It also has implications for organizations within the scope of the Enclave and Web Server STIGs. The DOD DNS Working group is a major contributing factor to the requirements in this and all DOD DNS related initiatives.

It should be noted that FSO Support and Technical Support for the STIGs, Checklists, and Tools are only available to DOD Customers.

### 1.2 Authority

DOD Directive 8500.1 requires that "all IA and IA-enabled IT products incorporated into DOD information systems shall be configured in accordance with DOD-approved security configuration guidelines" and tasks DISA to "develop and provide security configuration

guidance for IA and IA-enabled IT products in coordination with the Director, NSA.” This document is provided under the authority of DOD Directive 8500.1.

The use of the principles and guidelines in this STIG will provide an environment that meets or exceeds the security requirements of DOD systems operating at the Mission Assurance Category (MAC) II Sensitive level, containing unclassified but sensitive information.

### 1.3 Scope

This document details DOD DNS security practices and procedures applicable to all DOD name servers including authoritative and recursive servers. The requirements set forth in this document are designed to assist Information Assurance Officers (IAOs) and DNS administrators. The policy portions of this STIG are relevant to all name servers connected to either the DOD Un-Classified (But Sensitive) Internet Protocol Router Network (NIPRNet) or Secret Internet Protocol Router Network (SIPRNet).

This document supports the following implementations of DNS:

- BIND 9.3.1 and above
- BIND 9.3.2 for, Windows 2000, and Windows 2003
- Windows 2000/2003 DNS
- CSS DNS

The field use of Berkeley Internet Name Domain (BIND) releases are 9.2.6 and above, 8.4.7 and above, and 9.3.1 and above are acceptable, but the document does not detail the syntax of BIND 8 configuration statements. In most cases, BIND 8 and BIND 9 statements are identical. When they are not, organizations deploying BIND 8 must make the appropriate changes to the BIND 9 syntax in this document to achieve the desired effect.

The field use of Windows 2003 DNS is also acceptable, but the document does not include the configuration of Windows 2003 dialog boxes. Nevertheless, security relevant configuration options for Windows 2000 and Windows 2003 DNS are very similar. Any references to Windows 2000 in this document also apply to Windows 2003.

This STIG does not address the DNS configuration of DNS clients (i.e., the workstations, servers, and network devices that query name servers). Each of these clients runs DNS resolver software. Any requirements concerning those resolvers would be addressed in the STIG corresponding to the underlying technology such as the Desktop Services or Operating Systems STIGs(e.g., Desktop).

### 1.4 Writing Conventions

Throughout this document, statements are written using words such as “**will**” and “**should**.” The following paragraphs are intended to clarify how these STIG statements are to be interpreted.

A reference that uses “**will**” implies mandatory compliance. All requirements of this kind will also be documented in the italicized policy statements in bullet format, which follow the topic paragraph. This will make all “**will**” statements easier to locate and interpret from the context of the topic. IAO will adhere to the instruction as written. Only an extension issued by the Designated Approving Authority (DAA) will table this requirement. The extension will normally have an expiration date, and does not relieve the IAO from continuing their efforts to satisfy the requirement.

A reference to “**should**” is considered a recommendation that further enhances the security posture of the site. These recommended actions will be documented in the text paragraphs but not in the italicized policy bullets. Nevertheless, all reasonable attempts to meet this criterion will be made.

For each italicized policy bullet, the text will be preceded by parentheses containing the italicized Short Description Identifier (SDID), which corresponds to an item on the checklist and the severity code of the bulleted item. An example of this will be as follows "(DNS0100: CAT II)". If the item presently has no Potential Discrepancy Item (PDI), or the PDI is being developed, it will contain a preliminary severity code and "N/A" for the PDI (i.e., "[N/A: CAT III]").

### **1.5 DISA Information Assurance Vulnerability Management (IAVM)**

The DOD has mandated that all IAVMs are received and acted on by all commands, agencies, and organizations within the DOD. The IAVM process provides notification of these vulnerability alerts and requires that each of these organizations take appropriate actions in accordance with the issued alert. IAVM notifications can be accessed at the Joint Task Force - Global Network Operations (JTF-GNO) web site: <http://www.cert.mil>.

### **1.6 STIG Distribution**

Parties within the DOD and Federal Government's computing environments may obtain the applicable STIG from the Information Assurance Support Environment (IASE) web site. This site contains the latest copies of all STIGs, as well as checklists, scripts, and other related security information and guidance. The NIPRNet URL for the IASE site is: <http://iase.disa.mil/>.

### **1.7 Document Revisions**

Comments or proposed revisions to this document should be sent via e-mail to **fso\_spt@disa.mil**. DISA FSO will coordinate all change requests with the relevant DOD organizations before inclusion in this document.

## 1.8 Vulnerability Severity Code Definitions

<b>Category I</b>	<p>Vulnerabilities that allow an attacker immediate access into a machine, allow superuser access, or bypass a firewall.</p> <p>DNS Application: Vulnerabilities that may result in maliciously fabricated information being returned to a client. Vulnerabilities that allow an attacker to bypass access controls on the underlying system or network.</p>
<b>Category II</b>	<p>Vulnerabilities that provide information that has a high potential of giving access to an intruder.</p> <p>DNS Application: Vulnerabilities that may result in a customer-perceived denial of service persisting longer than 30 seconds. Vulnerabilities that provide information that have a high potential of giving access to an intruder.</p>
<b>Category III</b>	<p>Vulnerabilities that provide information that potentially could lead to compromise.</p> <p>DNS Application. Vulnerabilities that may result in administratively controlled DNS information (e.g., via split DNS) being released to improper entities as well as vulnerabilities that may result in a customer perceived denial of DNS service persisting less than 30 seconds</p>
<b>Category IV</b>	<p>Vulnerabilities, when resolved, will prevent the possibility of degraded security.</p> <p>DNS Application: Other Vulnerabilities or configurations, which tend to reduce the survivability and robustness of the DNS system.</p>

**Table 1-1. Vulnerability Severity Code Definitions**

## 2. DNS TERMINOLOGY AND CONCEPTS


DNS is a hierarchically structured, distributed database that provides name services to IP networks such as the Internet. This database contains hostnames and their corresponding IP addresses. Name servers support hostname to IP address translation (forward resolution) as well as the reverse, IP address to hostname translation (inverse resolution). The DNS protocol establishes rules for the following:

- Querying and updating the names database
- Database replication
- Database schema

DNS has evolved into a critical component of the Internet infrastructure by making network navigation less cumbersome for the user and applications on large IP networks. The following sections describe major DNS concepts, components, and security objectives.

### 2.1 Domains

The data stored in the DNS is identified by *domain names* that are organized as a hierarchical tree according to organizational or administrative boundaries. A hypothetical domain tree might be as follows:



**Figure 2-1. Domain Tree**

Each branch (or node) on the tree is given a *label*. A fully qualified domain name is the concatenation of all these labels on a given path from the domain to the root. Typically, domain names are written as a string of labels listed from right to left and separated by periods. For example, consider the following fully qualified domain name(FQDN):


example.disa.mil.



In the illustration, “example” is a third or lower level domain, “disa” is a second level domain, and “mil” is a Top Level Domain (TLD). Note that mil is followed by a period. This represents the root but is usually omitted in most applications today because all Internet domains share the same root.

## 2.2 Zones

Most users of IP-based applications are very familiar with domains and their use, but zones are actually the more relevant building blocks of the naming structure in DNS. A zone is a logical group of network devices and may be either an entire domain, a domain with all of its sub-domains, or a portion of a domain. In our hypothetical example, the relationship of domains and zones might be as follows:



**Figure 2-2. Domains and Zones**

In this illustration, the example and cert domains are associated with their own zones. The zones for *disa* and *mil*, on the other hand, are combined with the sub-domains *lab* and *osd*, respectively. There are many different ways to partition any domain tree into various zones.

The rationale for the existence of zones is that they make it easier to distribute the name database. Just as it would not be practical for the entire Internet DNS to reside on a single server, it also does not make sense to assign a unique name server for each lower-level domain. Zones allow network architects and administrators to combine domains in ways that optimize the management of a given portion of the name space.

## 2.3 Name Servers

A *name server*'s primary function is to respond to client queries for information from the DNS. Although name servers can be configured in a wide variety of ways, there are essentially two types of name servers:

- **Authoritative:** Those that store zone records for one or more DNS zones.
- **Caching:** Those that provide responses based on a series of queries for records stored on other name servers or from cache of responses built by using previous queries.

The first type of name server is called an *authoritative* name server because it can answer authoritatively when asked for information about records in one of its zone files. It consults its internal data to answer incoming requests rather than trying to pass them on to another server. In some references, authoritative name servers are also called *advertising* name servers.

The second type of name server is called a *caching* name server because it caches the queries it resolves on behalf of clients. Most computers on the Internet do not search for and query authoritative name servers directly, but instead are configured to query from one to three caching name servers that do this work on their behalf. Caching name servers are also referred to as *recursive* or *resolving* name servers.

Importantly, a name server can be simultaneously both an authoritative and a caching name server; however, this is not considered a secure configuration. When an authoritative name server is also a caching server, both authoritative and external records are stored in the cache. This configuration means that if an adversary is able to corrupt or poison the cache, through external queries, the adversary may be able to modify the name server's responses to legitimate queries for authoritative records. When the authoritative function is separated from the resolving function, as best practice dictates, this risk is eliminated.

### 2.3.1 Authoritative Name Servers

For fault tolerance and continuity of operations, most zones have two or more authoritative servers. The authoritative server, where the master copy of the zone data is maintained is called the master (primary) name server, hereafter simply called the master name server. It loads the zone contents from a local zone file that the DNS database administrator creates and edits.

Each master for a zone may have one or more associated slave (secondary) servers, hereafter called slaves. A slave server is also authoritative for that zone but gets its updates from its master server using a replication process known as a *zone transfer*. Typically, each zone has only one master from which all slave name servers obtain updates. There are, however, multiple configuration options for any DNS architecture.

### 2.3.2 Caching Name Servers

A caching name server can resolve client queries using one of two mechanisms:

- Forwarding
- Recursion

The forwarding feature provides the ability for the name server to simply forward the request or query on to another caching server without any additional effort. With *recursion*, the name server recursively searches for another name server that has authoritative name information to get the records that the client (resolver) requested.

In this document, servers that support forwarding are referred to as *forwarding* servers. Similarly, servers that support recursion are called *recursive* servers.

Caching servers store the results of the lookups they perform in memory. This process, known as *caching*, improves performance because frequently queried records are readily available. Caching also reduces network traffic, since the zone's authoritative name servers do not need to be queried as often.

When a caching name server attempts to resolve a host name, it first searches its cache for the lowest level information in the DNS hierarchy it can find. For example, if trying to resolve example.disa.mil, the name server first checks if there is a Resource Record (RR) known as an Address (A) record for that name in the cache. If there is not an A record, then the name server checks whether there is an Name Server (NS) resource record for example.disa.mil in the cache. If there is an NS record, then it sends a query to the authoritative server for example.disa.mil. If there is not an NS record, then it searches for an authoritative name server for disa.mil and so on.

Upon system initialization, the caching name server has no previously stored information in its cache and thus must start with the root authoritative name servers. Knowledge of the root name servers is located in the "root hints", a file that is included in the installation of the name server software and can be updated if IP addresses of the root servers change.

### 2.4 Name Server Support Roles

Supporting a name server involves administration of three separate components:

- The underlying operating system
- The name server software
- The DNS zone database

In this document, the term System Administrator (SA) refers to the individual responsible for operating system functions, DNS software administrator refers to the individual responsible for administering the DNS software (e.g., modifying configuration files), and DNS database administrator refers to the individual responsible for entering, modifying, and deleting records in

the DNS zone database. When the term DNS administrator appears, it refers to both DNS software and DNS database administrators.

At some sites, an individual will be responsible for all these functions, while at other sites responsibility for these functions will be divided among several individuals or teams. Each site or component determines the manner in which it separates or assigns duties, however, regardless of the division of labor, all supporting personnel must comply with this STIG. To facilitate certification and accreditation requirements, the IAM/IAO will be expected to explain how DNS support functions are delineated among personnel at that location.

## **2.5 Resolvers**

Another key component of DNS is the client resolver, which formulates and sends DNS queries to name servers. Most resolvers are rather simple. They merely send host names to caching name servers and wait for the response. The resolvers are typically configured with a list of two or three caching name servers, thereby ensuring there is a backup in case the master name server is unavailable.

In the future, resolvers may perform more complex functions. For example, they may sign queries if the caching name server requires authentication of requests. They may also request that replies be signed and the signatures be validated. This functionality is not available in common commercial resolvers today.

This page is intentionally left blank.

### 3. DNS ARCHITECTURE AND GENERAL SECURITY REQUIREMENTS

DNS security can be addressed at many different levels ranging from enterprise architecture to operating system parameters. Some aspects of the DNS protocols require strict consistency across all DNS products and operating systems, otherwise the wide variety of hosts on the Internet and within enterprise networks could not discover the host names and IP addresses they need to function properly. Other portions of the DNS, particularly with regard to security configuration, allow for differences among DNS software products.

This section addresses issues that impact all DNS implementations regardless of the product selected or the operating system on which it operates.<sup>1</sup> It also provides guidance on what products are acceptable in the DOD computing environment. Later sections in the STIG provide additional guidance for specific implementations of DNS (i.e., BIND, Windows DNS, and Cisco CSS DNS).

- *(DNS0160: CAT III) The IAO will ensure the DNS architecture is documented to include specific roles for each DNS server, the security controls in place, and what networks are able to query each server (CIS DNS BIND Benchmark).*

#### 3.1 Name Server Software

At the present time, BIND 9.3.1 or later is the only DNS implementation that supports all of the current and planned DOD required security functionality. DOD has a large investment in the development of BIND, for which DISA rigorously tests new releases. BIND is free and, therefore, does not impose licensing costs on organizations that deploy it. Given these advantages, organizations must have a compelling reason for use of a DNS alternative.

Some situations exist in which a BIND alternative may be acceptable. For example, the use of Windows 2000/2003 DNS greatly reduces the complexity of administering Windows networks. In addition, DNS is a fundamental component of the Active Directory environment. Guidelines for the use of this alternative are delineated in *Section 5, Configuration of Windows 2000/2003 DNS*, of this document.

There are also requirements associated with particular high-availability server implementations that are not met with BIND 9 versions. For example, some enterprise implementations require that DNS automatically remove records for a member of a server cluster on a temporary basis when that machine is offline and then return the record when operations resume. When servers are distributed over a wide area, the preference is to direct a client computer (resolver) to the server that provides the fastest response time for that client, which could vary depending upon geographic proximity, network traffic, or server loads. The Cisco CSS DNS offers this functionality, which is not present in BIND 9. Guidelines for the use of this alternative are delineated in the section on *Configuration of Cisco CSS DNS*, of this document.

---

<sup>1</sup> The exception to this rule is Cisco CSS DNS, which does not support all the requirements in this section but nonetheless has been approved for specialized applications.

If an organization believes that it requires name server software other than those discussed above, it should contact DISA FSO to determine an appropriate course of action. In most cases, use of an alternative will generate a certification finding, although extensions may be granted when either it meets the general requirements listed in this section or mitigating controls provide equivalent levels of assurance as BIND.

- *(DNS0400: CAT II) The IAO will ensure name server software on production name servers is BIND or Windows 2000 (or later) DNS unless operational requirements cannot be met with these options, in which case the alternative must still be configured in a manner to satisfy the general security requirements listed in this STIG.*

The only alternative currently permitted is Cisco CSS DNS, which is currently limited to only those hosts defined in the csd.disa.mil domain. CSS DNS is both subject to the some of these general security requirements, where applicable, and the specific STIG guidance for this product.

### 3.2 Name Server Operating System Platforms

Another critical component of DNS security is the security of the operating system (OS) platforms on which the DNS software runs. If it is not possible to secure the OS, then DNS itself cannot be secure. Accordingly, organizations must select an appropriate OS for its name servers, one that has a well documented, secure configuration.

Even a securely configured operating system is vulnerable to the flaws of the programs and applications that run on it. To prevent DNS software from being subject to the vulnerabilities of other programs, it is best not to run other programs at all – or at a minimum, run only those programs that are necessary for either OS or DNS support. In other words, in a properly secured environment, a name server would not run on the same device that also provides users web, e-mail, firewall, or database services.

- *(DNS0410: CAT II) The IAO will ensure DNS server software runs only on approved operating systems as defined by the appropriate OS STIG.*
- *(DNS0170: CAT II) The IAO will ensure the underlying operating system of the DNS server is in compliance with the appropriate OS STIG.*
- *(DNS0175: CAT I) The IAO will ensure the DNS server software is not installed and operational on operating systems that are no longer supported by the vendor.*
- *(DNS0415: CAT II) The IAO will ensure DNS software runs only on dedicated hardware, with the exception of Windows 2000/3 DNS, which must run on a domain controller that is integrated with Active Directory services.*
- *(DNS0420: CAT II) The SA will ensure only the DNS software process ID (PID) has read access to the files containing DNS encryption keys.*

- *(DNS0425: CAT II) The SA will ensure only the DNS software Process ID (PID) and/or the DNS database administrator have edit access to the zone database files.*
- *(DNS0430: CAT II) The SA will ensure only the DNS software administrator and the DNS software PID have read access to the DNS software configuration files and only the DNS software administrator has write access to these files.*
- *(DNS0435: CAT II) The SA will ensure the name server's IP address is statically defined.*
- *(DNS0440: CAT II) The SA will ensure an integrity checking tool is installed and monitoring for any modifications to the root.hints and named.conf files.*
- *(DNS0180: CAT II) The SA will ensure a Host Based Intrusion Detection System (HIDS) is installed on the DNS server.*

Refer to the appropriate Operating System STIG for the currently supported versions of operating systems.

Cisco CSS runs on its own Internetworking Operating System (IOS) and therefore is not subject to many of these requirements. Updates to this list prior to the next publication of the STIG may appear in the *DNS Security Checklist*, a companion document to this STIG that provides instructions on how to check for vulnerabilities and the appropriate mitigations (fixes) for those vulnerabilities.

The technical implementation guides, security recommendations, and associated tools for supported operating systems can be obtained from the Information Assurance Support Environment website (<http://iase.disa.mil> website).

### **3.3 Redundancy, Dispersal, and Availability**

A critical component of securing an information system is ensuring its availability. The best way to ensure availability is to eliminate any single point of failure in the system itself and in the network architecture that supports it.

Fortunately, the inherent design of DNS supports a high-availability environment. Master and slave servers regularly communicate zone information, so if any name server is disabled at any time, another can immediately provide the same service. The task for the network architect is to ensure that a disaster or outage cannot simultaneously impact both the master and all of its slave servers. If a disaster occurs, the DNS protocols cannot prevent total loss of name resolution services for hosts within affected zones.

The solution is to disperse name servers in such a way as to avoid single points of failure. At minimum, authoritative name servers for the same zone should be on different network segments in order that at least one name server is available in the event that a router or switch fails. This fault tolerance should also extend to wide area data communications lines. For example, if a site has multiple leased lines connecting the network on which the name server resides to a larger



network such as the NIPRNet, routing protocols should be configured such that if one of the lines fails, another one will still be available to support the name server.

Organizations should also be prepared for greater disasters, such as the destruction of a building, an entire campus, or in the case of a hurricane, an entire city. In situations in which all the hosts defined on an authoritative name server are located in the same building as the name server, then loss of DNS will not impact availability of service simply because the computing infrastructure is already down. On the other hand, if all the authoritative name servers for a zone reside in a single building, but hosts defined within the zone are located elsewhere, then the loss of the DNS will impact service. The loss of service occurs because users (and other infrastructure devices and servers) will not be able to resolve host names for servers/services that are otherwise still operational at an unaffected site.

Given that name servers can be dispersed across a network and still be physically located near one another, the DNS architecture should require reasonable geographic dispersal as well. Understandably, many small office sites will not have local name servers. However, if the host records for a particular site reside on a name server at a remote location, there should be a backup for that name server at an alternate location. If an organization does not have the resources for this level of dispersal, it can partner with another organization to have each organization's master name server to serve as a slave for the other organization's zone. In this configuration, the name servers are both masters for some zones and slaves for others.

- *(DNS0200: CAT I) The IAO will ensure each authoritative master name server has at least one and preferably two or more active slave servers for each of its zones.*
- *(DNS0205: CAT I) The IAO will ensure name servers authoritative for a zone are located on separate network segments.*
- *(DNS0210: CAT II) If a zone includes hosts located in more than one building or site, the IAO will ensure at least one of the authoritative name servers supporting the zone resides in a different facility from the other name server(s).*

### **3.3.1 Network related availability**

One of the most critical vulnerabilities of the DNS is that local queries for information can be affected by non-local communications failures. For example, an attempt to look up "local.disa.mil" could be prevented by an inability to reach a root DNS or a .MIL server caused by a broken link, a routing failure, a mis-configured firewall, or other network related issues. This can prevent even physically co-located systems from communicating with one another due to an inability to convert a configured DNS name into the IP address(es) they require to connect.

For the vast majority of applications, normal configurations of communication and server redundancy are sufficient to provide the level of service required. However, there may be circumstances where additional DNS survivability or continuity of operations will be required to meet mission critical requirements.

### 3.3.2 Stub Zones

One technique to enhance survivability is "Stub Zones". Recursive resolvers can be configured with information about a specific zone, which allows the resolver to bypass normal hierarchical lookup and go directly to the zone. For example, a resolver configured with stub zone information about ".MIL" does not need to query any of the root servers about the location of .MIL name servers as that information will be maintained at/by the resolver.

The recursive resolver is configured with the name of the zone, the location of the master, and optionally a few slaves. The resolver then downloads the complete set of NS records for the zone as well as the A records for names in the NS records. These are updated and maintained on the same schedule as if the recursive resolver were a slave server for the zone.

For the purposes of this STIG, resolver operators should consider the installation of up to three classes of stub zones. First, a .MIL stub zone as described above to insulate against issues external to the DOD. Second, a local stub zone covering the local network or the organizational enterprise network. Third, one or more stub zones covering zones which point to systems that contain resources critical to the local mission.

Stub zones are not a panacea and do require additional administrative oversight to ensure configuration information does not become stale. Operators should balance the additional operational burden against the mission needs when deciding on the extent of stub zone implementation.

### 3.4 Authentication and Access Control

The general security objectives for all information systems are confidentiality, integrity, and availability. The primary objective of DNS authentication and access control is the integrity of DNS records; only authorized personnel must be able create and modify resource records, and name servers should only accept updates from authoritative master servers for the relevant zones. Integrity is best assured through authentication and access control features within the name server software, though firewalls also play a significant role in controlling DNS transactions on a network.

A secondary objective of DNS authentication and access control is the confidentiality of DNS records; only those with a valid business purpose should be able to obtain the host names and IP addresses defined within a zone. Both name server configuration and firewalls can support this objective as well. Nevertheless, an organization should never rely on restrictions to DNS host records to provide a significant safeguard to the hosts themselves, as this would constitute a very thin veil of protection. The operating assumption should be that a determined adversary would be able to obtain an IP address of a host within the zone. With a defense-in-depth posture, other controls will prevent the adversary from being able to do much with this address if it were obtained.

### 3.4.1 Zone Updates

DNS resource records are created and modified through *zone updates*. These occur either through update of individual records on a master name server or through zone transfers from a master name server to one of its slaves.

#### 3.4.1.1 Updates of Individual Records

Individual resource records can be updated manually, where a DNS database administrator either edits the zone file directly or uses a tool to do so, or dynamically, where an automated process enters changes to the zone file without the intervention of the DNS database administrator.

The *dynamic update* capability has considerable appeal in an environment in which IP addresses change so frequently that it would be unacceptably burdensome or expensive to dedicate the time of a DNS database administrator to this function. This condition would likely be met at sites that rely on the Dynamic Host Configuration Protocol (DHCP) to assign IP addresses to client devices such as workstations, laptops, and IP telephones. It would also apply to sites that utilize frequently changing service (SRV) records.

On the other hand, dynamic updates can pose a security risk if the proper security controls are not implemented. When dynamic updates are permitted without any mitigating controls, a host with network access to the name server can modify any zone record with an appropriately crafted dynamic update request.

The solution is to require cryptographic authentication of all dynamic update requests, but not all DNS software supports this functionality. When it does not, dynamic updates must be prohibited. BIND 9 name servers support cryptographic authentication of dynamic update requests using Transaction Signature (TSIG) (also referred to as Transaction Authentication), a symmetric key technology, but it is relatively rare for DNS clients to support TSIG. As a consequence, dynamic updates are not likely to be feasible for most implementations of BIND. Microsoft's Windows 2000 (and above) DNS supports what it terms *secure dynamic update*. This technology utilizes GSS-TSIG, which leverages a Kerberos-based cryptographic authentication infrastructure shared among both servers and clients. Therefore, GSS-TSIG is required for secure dynamic updates to a Microsoft Windows DNS server.

- (DNS0450: CAT I) *The DNS software administrator will disable dynamic updates unless the DNS software is configured to require all dynamic updates be cryptographically authenticated.*

For detailed configuration guidance for dynamic updates within a Windows 2000/2003 DNS environment, refer to the section on *Secure Dynamic Updates and Active Directory*.

For detailed configuration guidance for dynamic updates within BIND, refer to the *zone Statement section*.

### 3.4.1.2 Zone Transfers

A slave updates its zone information by requesting a zone transfer from its master server. In this transaction, the risk for the slave is that the response to its request is not in fact from its authorized master but rather from an adversary posing as the master. In this scenario, such an adversary would be able to modify and insert records into the slave's zone at will. To protect against this occurrence, the slave must be able to authenticate the master to provide assurance that any zone updates are valid.

The risk to the master server in this situation is that it would honor a request from a host that is not an authorized slave, but rather an adversary seeking information about the zone. To protect against this possibility, the master must first have knowledge of what machines are authorized slaves. Then the master must authenticate each slave when that server requests a zone transfer.

One way an organization can help prevent bogus requests for zone transfers from the master server is to hide the existence of the name server. When this technique is employed, the master name server is called a "stealth master". Hiding the stealth master involves removing its NS record from all related zone files and also ensuring that it does not appear as the master name server in any Start of Authority (SOA) records. Slaves are still configured to request zones from the master, but no hosts other than the slaves can learn the master's IP address without access to the software configuration files, which should be adequately protected. The stealth master architecture also has the advantage of improving performance because the master server is not burdened by client queries and exists only for zone updates and transfers.

- *(DNS0455: CAT I-II<sup>2</sup>) The DNS software administrator will configure each slave supporting a zone to cryptographically authenticate its master before accepting zone updates.*
- *(DNS0460: CAT II) The DNS software administrator will configure each zone master name server to limit zone transfers to a list of active slave name servers authoritative for that zone.*
- *(DNS0465: CAT III) The DNS software administrator will configure each zone master server to cryptographically authenticate any slave requesting a zone transfer.*
- *(DNS0470: CAT II) The DNS software administrator will configure a name server to only accept notifications of zone changes from a host authoritative for that zone.*

For detailed configuration guidance for zone updates and transfers within BIND 9, refer to the *zone Statement* section.

---

<sup>2</sup> A violation of this requirement can have one of two severity levels depending upon the extent of the violation. If slaves do not authenticate masters in any manner, then the discrepancy would be a Category I finding. If some form of authentication exists (i.e., based on IP address), but it is not based on cryptography, then the discrepancy would be a Category II finding.

### 3.4.2 Query Restrictions

The default configuration in nearly all DNS server software is to permit any client to access any record. There are two reasons why this configuration would be considered weak security. First, some of the records may reveal information about how an internal network is configured and therefore, should not be shared with external clients. Adversaries could potentially use this information to plan an attack on the internal network. Second, processing queries for any client could allow an adversary with access to the name server to exploit known or unknown flaws in the query process. The exploits might disable the name server, degrade its performance, modify the record cache, or provide the adversary with elevated privileges.

#### 3.4.2.1 Split DNS

The following section represents a substantial departure from previous DNS STIG guidance on Split DNS.

Split DNS or (split horizon) is a deployment of DNS where the same DNS question is answered differently depending on who is asking the question. For example, the query for A records associated with `www.example.mil` could return one set of addresses if the query comes from NIPRNet connected IP addresses, and could return another set of addresses if the query originates external to the NIPRNet.

Previous versions of this STIG recommended the use of Split DNS as a form of operational security. The expectation was that by denying an attacker the ability to associate names with addresses, some measure of security for internal hosts could be gained. Analysis of the use of Split DNS both within the DOD, and in the Internet at large, strongly suggests that any supposed security gain is more than offset by the additional operational complexity of a Split DNS configured zone.

Finally, Split DNS interacts poorly with the deployment of DNSSEC. Split DNS results in different answers being given to different questioners, which is at odds with the DNSSEC model. The DNSSEC model generally provides for (and can result in the authentication of) only one set of answers for any given DNS query at any given time. While it is, with great difficulty, possible to configure DNSSEC to support Split DNS, this would add substantially to the operational complexity of DNS and reduce the robustness of the DNS infrastructure.

DNS is a public information utility. DNS operators should assume that any data placed in the DNS would be available to anyone connected to the Internet. Split DNS shall not be considered a mitigating factor or technique to deny DNS information to an attacker. Split DNS will continue to be required in one situation only: to prevent private address space (e.g. 10.0.0.0/24) from leaking into the public DNS system.

- *(DNS0215: CAT III) The IAO will ensure, when using private IP address space within an Enclave, a split DNS configuration is implemented to prevent the private address space from leaking into the public DNS system.*

### 3.4.2.2 Operational Security

DNS administrators must review the contents of their zones at least as often as annually for content or aggregation of content that may provide an adversary information that can potentially compromise operational security. This specifically includes names that provide an outsider some indication as to the function of the referenced system. The DNS administrator is the final adjudicator of the sensitivity of DNS information, but should make a conscious decision to include such information based on operational need (e.g., naming a DNS server something other than dns.zone.mil is unlikely to provide any additional security given these servers are easily identifiable during DNS queries).

- *(DNS0185: CAT III) The IAO will ensure the DNS administrator reviews the contents of the zones they are responsible for, at a minimum, annually.*

### 3.4.2.3 Restrictions on Recursion

A potential vulnerability of DNS is that an attacker can poison a name server's cache by sending queries that will cause the server to obtain host-to-IP address mappings from bogus name servers that respond with incorrect information. Once a name server has been poisoned, legitimate clients may be directed to non-existent hosts (which constitutes a denial of service) or, worse, hosts that masquerade as legitimate ones to obtain sensitive data or passwords.

To guard against poisoning, name servers authoritative for .mil domains should be separated functionally from name servers that resolve queries on behalf of internal clients. Organizations may achieve this separation by dedicating machines to each function or, if possible, by running two instances of the name server software on the same machine; one for the authoritative function and the other for the resolving function. In this design, each name server process may be bound to a different IP address or network interface to implement the required segregation.

In addition to enforcing this separation, organizations must ensure that the caching servers only accept queries from known supported clients, as they are much less likely to attempt poisoning attacks than unknown external hosts. In most cases, the hosts inside an enclave constitute the list of known supported clients. However, in some cases, the caching name server may be expected to serve clients distributed over a wide area, which is acceptable, as long as the name server's audience is limited in some fashion. In no case, should a name server accept queries from any Internet host.

- *(DNS0475: CAT III) The DNS software administrator will prohibit recursion on an authoritative name server.*
- *(DNS0480: CAT III) The DNS software administrator will configure a caching name server to accept recursive queries only from the IP addresses and address ranges of known supported clients.*

### 3.4.3 Firewalls and DNS

Firewalls are an important component of a defense-in-depth protection strategy for DNS. One design consideration is the location of name servers relative to the location of firewalls within a network topology. An authoritative name server must never be on the external (or untrusted) side of a firewall. This would make the DNS server an untrusted device by definition. Furthermore, placing the server externally to the firewall would render it vulnerable to attack because it would not have the benefit of firewall protection.

Similarly, a name server that handles queries from external hosts must never be on an internal network. Externally accessible name servers must reside within a perimeter network, DMZ, where firewalls can monitor and block traffic to and from the name servers and the outside world, while also ensuring that external hosts do not directly communicate with internal hosts. On the other hand, if the name server provides authoritative information exclusively for internal hosts, then it should reside on the internal side of the firewall thus it cannot be reached from the outside.

A firewall administrator will need to configure the firewall to support DNS transactions (UDP and TCP 53). The administrator should ensure that this traffic is limited to authorized name servers; in particular, that inbound Port 53 requests to other hosts must be prohibited. The firewall administrator may also need to configure the firewall to support secure shell (TCP 22) or an acceptable alternative means of remote administration of the name server. The exact specification of these rules is beyond the scope of this STIG. However, zone and firewall administrators should only implement limits of access to TCP/53 if the DNS server software supports no other methods of controlling zone transfers, in which case the server software does not meet zone transfer restriction requirements.

Some organizations use firewalls to limit zone transfers, particularly in cases in which the DNS software does not have this capability. Zone transfers usually involve a TCP connection, while queries typically utilize the connectionless UDP. If, for example, an organization placed an authoritative server in its DMZ but limited zone transfers to servers behind the firewall, then it could allow inbound and outbound UDP 53 to and from the DMZ name server to allow queries, but deny TCP 53 in both directions to prohibit zone transfers. This configuration, however, is strongly discouraged because it may prevent legitimate DNS transactions that involve large responses (e.g., a DNSSEC signature). In general, limitations on queries, zone updates and transfers should be implemented in the name server's configuration rather than through configuration of firewalls, routers, or other communications devices.

- *(DNS0405: CAT II) The IAO will ensure hosts outside an enclave cannot directly query or request a zone transfer from a name server that resides on the internal network (i.e., not in a DMZ).*

### 3.5 Logging

DNS software administrators require DNS transaction logs for a wide variety of reasons including troubleshooting, intrusion detection, and forensics. These logs should be appropriately

secured, having file permissions that restrict unauthorized changes or viewing, and archived, being appropriately backed-up and stored in order for them to be examined at a future time. Numerous software products are available to aid the DNS software administrator in examining these transaction logs.<sup>3</sup>

- *(DNS0485: CAT I-II<sup>4</sup>) The DNS software administrator will configure the DNS software to log, at a minimum, success and failure of the following events:*
  - *start and stop of the name server service or daemon*
  - *zone transfers*
  - *zone update notifications*
  - *queries*
  - *dynamic updates*
- *(DNS0490: CAT II) The DNS software administrator will configure the DNS software to send all log data to either the system logging facility (e.g., UNIX syslog or Windows Application Event Log) or an alternative logging facility with security configuration equivalent to or more restrictive than the system logging facility.*
- *(DNS0495: CAT III) The DNS software administrator will configure the DNS software to add timestamps and severity information to each entry in all logs.*

### 3.6 Zone Files

Many implementations of DNS, including BIND, store zone information in text files, while others store this information in databases primarily accessed through GUI utilities. As both methods are acceptable, regardless of the format, there are basic practices that DNS database administrators should follow when managing the zones for which they have been assigned responsibility.

#### 3.6.1 Change and Ownership Documentation

A zone file should contain adequate documentation that would allow an IAO or newly assigned administrator to quickly learn the scope and structure of that zone. In particular, each record (or related set of records, such as a group of LAN workstations) should be accompanied by a notation of the date the record was created, modified, or validated and record the owner's name, title, and organizational affiliation. The owner of a record is an individual with the authority to request that the record be modified or deleted.

---

<sup>3</sup> For a list of tools focused on DNS, visit the ISC web site. The URL for ISC is <http://www.isc.org/>.

<sup>4</sup> A violation of this requirement can have one of two severity levels depending upon the extent of the violation. If no logging exists, then the discrepancy would be a Category I finding. If some logging exists, but not for all of the events listed, then the discrepancy would be a Category II finding.



This information will help administrators and auditors verify that the zone records are current and that only authorized personnel modify them. If records are not current, there is the potential that an adversary could simulate the activities of a retired host in order to capture logon credentials and other information. For example, presume a user has a bookmark for a retired web server in his browser. If the record for the server is not removed from DNS, a perpetrator could stand-up another server to mimic the behavior of the retired server, which users may still attempt to access because they may not have deleted or updated the bookmark for that server.

- *(DNS0220: CAT III) The DNS database administrator will document the owner of each zone record (or group of related records) and the date the record was created, last modified, or verified. This documentation will preferably reside in the zone file itself through comments, but if this is not feasible, the DNS database administrator will maintain a separate database for this purpose.*
- *(DNS0225: CAT III) Record owners will validate their records no less than annually. The DNS database administrator will remove all zone records that have not been validated in over a year.*

### 3.6.2 Zone-Spanning Records and Glue

If a name server were able to claim authority for a resource record in a domain for which it was not authoritative, this would pose a security risk. In this environment, an adversary could use illicit control of a name server to impact IP address resolution beyond the scope of that name server (i.e., by claiming authority for records outside of that server's zones). Fortunately, all but the oldest versions of BIND and most other DNS implementations do not allow this behavior. Nevertheless, the best way to eliminate this risk is to eliminate from the zone files any records for hosts in another zone. The two key exceptions to this rule involve glue for NS records and CNAME records for legacy resolution support.

*Glue* is a term used for a situation in which A records for a delegated zone's name servers appear in the zone file of the parent zone. This is illustrated in the following hypothetical excerpt from the zone file for disa.mil:

```
example    IN    NS    ns1.example.disa.mil
           IN    NS    ns2.example.disa.mil

ns1.example.disa.mil.    IN    A    132.40.11.2
ns2.example.disa.mil.    IN    A    132.40.15.3
```

In this instance, the zone file for disa.mil is authoritative for ns1 and ns2 in example.disa.mil, a different zone. Yet if the zone file did not include A records for ns1 and ns2, any client seeking records in example.disa.mil would be unable to reach that zone. Hence, the glue connecting example.disa.mil and disa.mil is a necessary exception to the idea that one zone should not contain authoritative records for another zone.

Another situation in which canonical names from one zone might appear in the zone file of a different zone is in the case of aliases. Imagine the fictitious “example” branch office closed and all of its resources were migrated to the Systems Management Center (SMC) in Oklahoma City. In this situation, the DNS database administrator might retire the example.disa.mil domain but still keep its zone operational until the user community has learned of the changes. The DNS database administrator can replace the A records in example.disa.mil domain with CNAME records for the new servers in okc.disa.mil as is shown in the following hypothetical excerpt from the example.disa.mil zone file:

```
$ORIGIN example.disa.mil
jupiter          IN      CNAME      jupiter.okc.disa.mil.
saturn           IN      CNAME      saturn.okc.disa.mil.
```

In general, zone-spanning aliases should be temporary (e.g., to facilitate a migration). When a host name is an alias for a record in another zone, an adversary has two points of attack; the zone in which the alias is defined, and the zone authoritative for the alias’s canonical name. This configuration also reduces the speed of client resolution because it requires a second lookup after obtaining the canonical name.

**Note:** There are certain situations where longer-lived CNAME records are appropriate (e.g., when Akamai (akadns) or similar services are utilized). Specifically, these services have higher than normal availability requirements and/or are using caching services.

- *(DNS0230: CAT III) The DNS database administrator will not include any resource records for a host in a zone file if its fully qualified domain name resides in another zone, unless the record is a glue record or CNAME record supporting a system migration.*
- *(DNS0235: CAT III) The DNS database administrator will ensure that no zone-spanning CNAME record, that points to a zone with lesser security is active for more than six months.*

### 3.6.3 Improper NS Records and Lame Delegation

An NS record should map a domain name to an active name server authoritative for that domain. Unfortunately, in poorly configured zone files, these NS records may refer to machines that are no longer in operation or ones that do not provide name services. In some cases, they may provide authoritative name service, but for different zones than the one intended in the NS record. This latter case is called *lame delegation*.

Poorly constructed NS records pose a security risk because they create conditions under which an adversary might be able to provide the missing authoritative name services that are improperly specified in the zone file. The adversary could issue bogus responses to queries that clients would accept because they learned of the adversary’s name server from a valid authoritative name server, one that need not be compromised for this attack to be successful.

The list of slave servers must remain current within 72 hours of any changes to the zone architecture that would affect the list of slaves. If a slave server has been retired or is not

operational but remains on the list, then an adversary might have a greater opportunity to impersonate that slave without detection, rather than if the slave were actually online. For example, the adversary may be able to spoof the retired slave's IP address without an IP address conflict, which would likely not occur if the true slave were active.

- *(DNS0240: CAT I) The DNS database administrator will ensure each NS record in a zone file points to an active name server authoritative for the domain specified in that record.*

### 3.6.4 Root Hints

Caching name servers require a set of servers to begin the query to look up data for a resolver, and then ultimately to cache that data. In order to bootstrap caching servers, information referring them to the root servers is provided since the root zone is the starting point for all zones. Modern versions of most DNS server software have this information hard coded into the program. However, to ensure that the data is current and correct, and for possible future custom configurations, it is best to configure a "root.hints" zone into a caching name server.

When authoritative servers are sent queries for zones that they are not authoritative for, and they are configured as a non-caching server (as recommended), they can either be configured to return a referral to the root servers or they can be configured to refuse to answer the query. The recommendation is to configure authoritative servers to refuse to answer queries for any zones for which they are not authoritative. This is more efficient for the server, and allows it to spend more of its resources doing what its intended purpose is; answering authoritatively for its zone.

The security risk is that an adversary could change the root hints and direct the caching name server to a bogus root server. At that point, every query response from that name server is suspect, which would give the adversary substantial control over the network communication of the name server's clients. The DNS software administrator may use *Appendix G, Root Hints*, to validate the root hints file entries. In nearly all cases, the Internet root hints are part of the installation package of the DNS software. Fortunately, the Internet root servers rarely change their addresses, so DNS software administrators will infrequently need to modify or update the root hints for servers meant to be resolving from the NIPRNet/Internet.

A DNS software administrator can obtain root hints from the following locations:

- <ftp://ftp.rs.internic.net> (198.41.0.6) (for the Internet)
- <ftp://www.ssc.smil.mil/ftp/domain/root.cache> (for SIPRNET)<sup>5</sup>

DNS software administrators should update or verify the root hints periodically (annually should suffice) to ensure that they have current records. There are several methods for obtaining the most current root hints. This file may be obtained from <ftp://ftp.rs.internic.net> (198.41.0.6) using

---

<sup>5</sup> An IP address for this SIPRNET server is not included in the STIG. Retrieving root hints from this URL assumes that one already has a valid root hints file and therefore there is some risk that the user could be directed to a bogus FTP server if this assumption is incorrect.

anonymous File Transfer Protocol (FTP), as mentioned above, to retrieve the *named.root* from the domain directory and then rename this file to the root hints file.

Bootstrapping a new caching server on the SIPRNet is problematic, as the bundled root hints file will not provide appropriate information to allow the server to resolve the URL above properly. The preference would be to retrieve a copy from another known SIPRNet DNS server, using ftp or secure copy (scp) to its IP address, or an out-of-band transfer method. If this is not possible, The SIPRNet Support Center (SSC) can provide the location or IP address, from which to retrieve it, upon request. Once bootstrapped, the updates to the SIPRNET root hints should be accessible from the ftp site above, which should be resolvable with the installed SIPRNET root hints file.

The Domain Information Groper (dig) utility is an additional method for obtaining an update of the Internet root hints file. A DNS software administrator can run the following command using the dig utility bundled with BIND and other DNS software:

```
dig @198.41.0.4 . ns > root.hints
```

- *(DNS0500: CAT I) The DNS administrator will ensure only valid root name servers appear in the local root zone file.*
- *(DNS0505: CAT III) The DNS software administrator will remove the root hints file on an authoritative name server in order for it to resolve only those records for which it is authoritative, all other queries are refused.*

## 4. CONFIGURATION OF BIND

The most common implementation of DNS protocols on the Internet and large enterprise name servers is Berkeley Internet Name Domain (BIND). BIND is most often found on UNIX based operating systems, where it runs as the *named* daemon. There is also an implementation of BIND for Microsoft Windows 2000 and 2003, where it runs as a service.

BIND 9 is the most recent and preferred version of BIND. Some name servers running BIND 8, as well as other alternatives, are still in operation, however, they are expected to migrate to BIND 9. There are several advantages of BIND 9 relative to its predecessors:

- Support for DNSSEC
- The ability to differentiate responses to queries using views
- Multiprocessor support
- IPv6 compatibility, including support for IPv6 ports and records
- Improved portability architecture

This section will first explain how to obtain authorized BIND software and then will review security related configuration of the *named.conf* file, which controls the behavior of BIND.

### 4.1 Obtaining Software

System Administrators should download source code from the Internet Software Consortium (ISC) (<http://www.isc.org>). The ISC site is an excellent source for documentation on DNS and BIND, and its various releases. The administrator should compile the source code using a trustworthy compiler running on a trustworthy computer.

If administrators need to compile code, they should do so on a trustworthy computer other than the name server and then transfer the resulting binaries to the name server for execution. If BIND is compiled on the same machine on which it will eventually run, this provides the potential for an adversary to either modify the compiler to insert malicious code or use the compiler to create unauthorized programs. In general, no compiler should be resident on a production name server at any time, yet it is understood that some versions of UNIX operating systems require local compilers for specific administrative functions.

In some cases, organizations may wish to use operating system vendor supplied versions of BIND in order to comply with or leverage a support agreement with the operating system vendor. This is acceptable so long as the vendor updates its code to be consistent with the most recent versions of BIND in order to avoid security vulnerabilities.

When an SA applies that vendor's patches, the software might be functionally equivalent to a more recent version of BIND than is stated in response to a query. For example, a fully patched version of Solaris 9's BIND might be functionally equivalent to BIND 8.2.7 although a query would respond that it is running BIND 8.2.4. To avoid any confusion, organizations that use vendor-supplied versions of BIND should work closely with the vendor to ascertain consistency with ISC versions of BIND.

An effort is underway to provide a DOD-authorized source for BIND binaries. The network address of this source and when it will become available has not yet been determined. Using a single authorized source will protect DOD against scenarios in which a malicious entity introduces faulty BIND code into the DOD environment to breach security.

- *(DNS0700: CAT II) The IAO will ensure version BIND 8.4.7 and above, 9.2.6 and above, or 9.3.1 and above is installed. If subsequent IAVA guidance recommends a BIND upgrade, that guidance will supersede this requirement.*
- *(DNS0190: CAT III) If BIND is utilized, the SA will subscribe to ISC's mailing list called "bind-announce" for vulnerabilities and software notifications.*

## 4.2 Configuration of the named.conf File

Once BIND has been downloaded from the ISC web site and installed, the next step is to configure it. Configuration information is contained in the *named.conf* file, a text file that has special statements and syntax unique to BIND. This section will review the statements that impact security and list requirements for how they should be configured. The relevant *named.conf* statements are as follows:

- key
- controls
- server
- acl
- logging
- options
- view
- zone

For additional information on statements and their syntax, see the *BIND 9 Administrator Reference Manual*. The guide may be obtained from the ISC's web site. On the BIND 9 page (<http://www.isc.org/products/BIND/bind9.html>), scroll down to the Documentation section and select the "PDF format" link.

### 4.2.1 key Statement

The *key* statement defines a key that can be used in Transaction Signatures (TSIG)<sup>6</sup>, which is a technology that authenticates information shared between two or more hosts by signing a hash of transferred information with a private key. Today, TSIG keys will most commonly be used to authenticate that transferred zone data come from an authorized source. Other potential uses include authentication of queries, dynamic updates, and *rndc* control clients. Importantly, the *key* statement does not secure anything by itself. Rather, it labels the keys so they can be used in other statements.

The syntax of the *key* statement is:

```
key key_id {  
    algorithm algorithm_name;  
    secret string;  
};
```

An example of *key* statement in practice might be:

```
key ns1.example.disa.mil_ns2.example.disa.mil {  
    algorithm hmac-md5;  
    include "/etc/dns/keys/ns1_ns2.key";  
};
```

And the file */etc/dns/keys/ns1\_ns2.key* might contain:

```
secret "2njlQNnzn6HTwKLCjStUXg==" ;
```

In this example the name of the key consists of the two hosts that are going to share the key, that is, the name servers ns1 and ns2 in the fictitious domain example.disa.mil. The algorithm is the hmac-md5 hashing algorithm, the only one currently supported in BIND 9. The use of the *include* statement nested within the *key* statement allows the DNS software administrator to separate the *named.conf* file from the actual key strings. This is helpful when there is a requirement to hide the keys from those who have the ability to view the *named.conf* file.

The generation of the required keys and their placement in the correct files involves several steps. Administrators should use the key generation tools supplied with the DNS software and not third party utilities. Administrators who are not familiar with this process may consult *Appendix C, Generation and Storage of TSIG Keys*, for further information.

- (DNS0195: CAT II) The SA will ensure the TSIG key is a minimum of 128 bits in length.
- (DNS0250: CAT III) The SA will ensure a new TSIG key is generated and utilized for each type of transaction (zone transfer, dynamic updates, etc)

---

<sup>6</sup> Transaction Signatures is also referred to as Transaction Authentication (abbreviated as TSIG).

A future release of BIND 9 is expected to include support for SHA-1, which is the algorithm currently specified in the National Institute of Standards and Technology's (NIST's) Secure Hashing Standard (FIPS180-1) and required throughout DOD. When it is available, hmac-md5 should be replaced with hmac-sha1 for DNS TSIG applications. In general, only NIST or National Security Agency (NSA) approved algorithms should be utilized in the DOD computing infrastructure.

- *(DNS0705: CAT III) The DNS software administrator will utilize 128-bit HMAC-SHA1 keys when available and 128-bit HMAC-MD5 keys until that time.*
- *(DNS0710: CAT II) The DNS software administrator will place each required key in its own dedicated file.*

#### 4.2.2 controls Statement and the rndc.conf File

The *controls* statement defines how BIND will listen for control messages (e.g., commands that reload zones, stop the name server, and unload a copy of the name server's internal database to a file). A DNS software administrator can use the BIND 9 *rndc* utility to generate the control messages. The BIND 8 equivalent is the *ndc* utility.

The risk posed by the *controls* statement and the associated use of the *rndc* or *ndc* commands is that an adversary could use them to remotely control the name server without having to authenticate to the operating system on which the name server resides. To prevent this occurrence, one can limit control messages to the local host, thereby prohibiting the use of the control utilities from remote hosts. This restriction should not hinder administration because authorized DNS software administrators should have user accounts on the machine and be able to remotely access it. Therefore, they can log on remotely and then run the control utility locally.

If a DNS software administrator uses the BIND 9 *rndc* statement under the restrictions just described, he or she can execute the command ***rndc-confgen -a*** and apply appropriate permissions (i.e., 440) to the resulting *rndc-key* file that should appear in the */etc* directory. The default configuration that results from this procedure will meet all security requirements. The remainder of this section provides additional information in case more complex configurations are required.

The syntax of the *controls* statement is:

```
controls {  
    inet ip_address | *  
        [ port ip_port ]  
        allow { address_match_list }  
        keys { key_list };  
    [ inet ... ]  
};
```



An *inet* phrase will exist for the network interface(s) on the name server that listens for control messages. One would enter “\*” to enable control messages to be heard on any interface. The listening port is configurable but defaults to TCP 953. The *allow* and *keys* statements provide security – the listener will accept messages only from those IP addresses listed in the *allow* clause and only when accompanied by a key listed in the key list.

A corresponding *key* statement in the *named.conf* file must accompany the *keys* phrase. The previous section reviewed the syntax for the *controls* statement and related security requirements.

An example of the *controls* statement in practice is:

```
controls {  
    inet 127.0.0.1  
    allow 127.0.0.1  
    keys { "rndc_key" };  
};
```

To complete the configuration, the *named.conf* file would also include a *key* statement for the *rndc\_key*.

The statement above restricts the use of *rndc* to the local host, which is the desired behavior. If there is no *controls* statement in the *named.conf* file, *named* will implement this behavior by default and look for *rndc\_key* in */etc/rndc.key* (or the path defined using the *sysconfdir* parameter when the BIND build was configured).

*rndc* also has its own configuration file, *rndc.conf*, that has a similar syntax to the *named.conf* file, but is limited to the *options*, *key*, *server*, and *include* statements. An example of a minimal configuration is as follows:

```
key rndc_key {  
    algorithm hmac-md5;  
    secret "2njlQNzn6HTwKLcjStUXg==" ;  
};  
options {  
    default-server localhost;  
    default-key rndc_key;  
};
```

A DNS software administrator can run ***rndc-confgen*** to create this file and a corresponding suggested *controls* statement for the *named.conf* file.

- (DNS0715: CAT II) The DNS software administrator will configure BIND to accept control messages only when the control messages are cryptographically authenticated and sent from an explicitly defined list of DNS administrator workstations.

### 4.2.3 server Statement

If a *key* statement defines a key intended to support TSIG, then it must be combined with a corresponding *server* statement to produce the desired behavior. The relevant syntax of the *server* statement is:

```
server ip_address {  
    keys {string; [string; [...]]};  
};
```

In our example, suppose ns1.example.disa.mil has IP address 10.1.1.1 and ns2.example.disa.mil has IP address 10.2.2.2. In this case, the following statement referencing ns2 would appear on ns1:

```
server 10.2.2.2 {  
    keys {ns1.example.disa.mil_ns2.example.disa.mil};  
};
```

Likewise, the reciprocal statement referencing ns1 would appear on ns2:

```
server 10.1.1.1 {  
    keys {ns1.example.disa.mil_ns2.example.disa.mil};  
};
```

Each name server will sign responses with the specified key, but will not require that requests be signed. Hence, both servers must include the relevant server statement to ensure that both ends of a transaction are signed. Note that both servers use the same key for communication between them.

- *(DNS0720: CAT II) The IAO will either ensure any two name servers sharing zone information utilize a unique TSIG key for communication between them or, in cases in which more than four servers support a zone, create a written key management plan that will document how keys are shared and replaced in a manner to reduce residual risk to an acceptable level.*

#### 4.2.4 *acl* Statement

The *acl* statement implements Access Control Lists (ACLs) within the *named.conf* file. It assigns a name to a list of IP addresses, networks, or keys to simplify references in other statements. For example, it is easier to read the name *trusted\_hosts* or *internal* rather than a long list of subnets. BIND ACLs should not be confused with access control lists on routers or server files, which must also be modified to secure DNS. In the case of BIND, the ACLs merely label various hosts and networks for use in later statements.

The syntax of the *acl* statement is:

```
acl acl-name {  
    address_match_list  
};
```

An example of *acl* for IP hosts inside an enclave might be:

```
acl trusted_hosts {  
    225.10.230.0/24;  
    239.10.241.128/25;  
    239.10.241.15;  
};
```

In this example, the list of trusted hosts consists of the networks 225.10.230.0/24 and 239.10.241.128/25, and a host with IP address 239.10.241.15.

ACLs may also include TSIG keys, as in the following example in which *ns1* communicates with *ns2* and *ns3*:

```
acl zone_partners {  
    "ns1.example.disa.mil_ns2.example.disa.mil";  
    "ns1.example.disa.mil_ns3.example.disa.mil";  
};
```

In this example, *ns1* only accepts responses from *ns2* and *ns3* when they are authenticated with the appropriate TSIG key.

As a best practice, DNS software administrators should create BIND ACLs for all supported hosts or networks that will be permitted to use the name server on which the configuration file resides. The DNS software administrator should also create an ACL for the name servers with which it will transfer zone information. Organizing the relevant hosts and networks in this manner will simplify the administration and audit of the *named.conf* file.

#### 4.2.5 *logging* Statement

The default behavior for BIND 9 complies with the general logging requirements as specified in the section on *Logging*. In many cases, the default configuration may not be appropriate for

operational reasons. For example, separation of duties between SAs and DNS software administrators may require them to maintain separate logs, perhaps one capturing events that the other does not. The precise nature of the logging configuration is at the discretion of the site as long as it meets the general requirements. Organizations that desire non-default logging behavior will need a strong understanding of the *logging* statement and its *categories* and *channels* phrases, which are explained in the remainder of this section.

The *logging* statement controls what BIND messages will be sent and to where. The *categories* sub-statement (or phrase) defines the “what” part of logging – i.e., the scope and content of log entries. The *channels* phrase defines the “where” part of logging – i.e., the file or process to which the log entries are written.

The syntax of the *logging* statement is:

```
logging {  
    [ channel channel_name  
        ( file path_name  
            [ versions ( number | unlimited ) ]  
            [ size size_spec ]  
            | syslog syslog_facility  
            | stderr  
            | null );  
        [ severity (critical | error | warning | notice |  
                    info | debug [level]| dynamic); ]  
        [ print-category yes/no; ]  
        [ print-severity yes/no; ]  
        [ print-time yes/no; ]  
    }; ]  
    [ category category_name {  
        channel_name ; [ channel_name ; ... ]  
    }; ]  
};
```

An example of the *logging* statement in practice might be:

```
logging {
    channel "query_log" {
        file "/var/dns/named.query"
        versions unlimited
        size 30 m;
        severity info;
        print-category no;
        print-severity yes;
        print-time yes;
    };
    category "queries" {
        "query_log";
    };
};
```

This statement will log DNS queries to a file */var/dns/named.query*. Entries will include the severity of the message and the time and date it occurred, but will not include the category name (we already know that it is a query). When the log size reaches 30 MB, a new log will be opened. Query logs can grow very quickly, so setting limits on their size ensures that they can be managed appropriately (e.g., moved to CD-ROM for archival, viewed in a text editor, etc.). In addition, if a size limit is set, unusually rapid growth in the number of query logs is a good indication that there is something wrong on the network and perhaps a denial of service attack is in progress. When selecting the log size for a particular name server, DNS and System Administrators should consider the server's disk space utilization, the typical volume of DNS transactions, and the expected log rotation and archival scheme. Lack of disk space can be problematic for server performance and availability.

The rationale for logging query data to a separate log is that in most environments, the number of query entries is likely to greatly exceed the entries of all other categories combined. Important messages are more likely to be overlooked during daily log reviews if they are easily lost in the noise. Moreover, archival requirements for system information will exceed those of query information in nearly all cases. Combining both types of information into a single log would require a great deal of data to be stored much longer than necessary, which increases administrative burden and overhead.

The *logging* statement will also log most other behavior via a category named *default*, which does not need to be explicit in the *named.conf* file. There are several predefined phrases that the DNS software administrator should know to fully understand the logging behavior of BIND. The most relevant predefined phrases are:

```
channel "default_syslog" {
    syslog daemon;
    severity notice;
};
channel "default_debug" {
    file "named.run";
    severity dynamic;
};
category "default" {
    "default_syslog";
    "default_debug";
};
```

If the DNS software administrator needs to supplement these predefined phrases to meet a business or critical mission requirement, then the complete phrase needs to be included in the *named.conf* file. The DNS software administrator should consult the *BIND 9 Administrator Reference Manual* for more detailed information on the *logging* statement and its *channel* and *category* phrases. The information here constitutes only a minimally acceptable configuration for BIND logging.

On UNIX systems, the DNS software administrator should also read the syslog man pages to determine how *syslog.conf* statements may interact with the logging directives in the *named.conf* file (type **man syslog** at the command prompt to obtain this information). In general, the *syslog.conf* file settings will supersede the *named.conf* file settings if they are set at a higher severity level. For example, if BIND specifies that the logging channel records everything of severity notice and above, but syslog is configured to log only critical messages, then only critical messages will be captured. The DNS software administrator may need to modify the *named.conf* and/or *syslog.conf* files to ensure that STIG requirements are met.

In the example above, there is a special *channel* for the queries *category* and all other logging and categories log to the default *channel* and default *category*. Therefore, entries of severity level notice and above will be sent to the syslog daemon for all categories other than the query category. The *named.run* file will capture all entries other than queries at the global severity level of the operating system (which results from using the value *dynamic* for the severity parameter). The DNS software administrator should maintain a separate log for the queries category to avoid crowding critical application messages with extensive query records. In addition, the DNS software administrator should set query log size restrictions appropriate to the environment, but never to exceed 80 percent of available disk space.

## 4.2.6 options Statement

The *options* statement in the *named.conf* file defines global options for BIND. Some options can also be defined for each zone. Not all options impact security. The relevant BIND options for this STIG are as follows:

- version
- recursion
- forward
- forwarders
- allow-query
- allow-transfer
- allow-recursion
- notify
- allow-notify
- also-notify
- fake-iquery
- rfc2308-type1

The general syntax of the *options* statement is:

```
options {  
    [ option1 option_value1; ]  
    [ ... ]  
    [ optionN option_valueN; ]  
};
```

Importantly, the *options* statement on an authoritative name server will look quite different than one on a caching name server because of the different roles these servers have in the DNS infrastructure. The following is an example of an *options* statement that includes security-relevant parameters for a DNS STIG compliant configuration on an authoritative server:

```
options {  
    version " ";  
    recursion no;  
    allow-query {none};  
    allow-transfer {none};  
};
```

Again, there will be other options listed in a typical *options* statement. Only security-relevant phrases are shown above. Each of these options will be discussed in more detail later in the section.

In the example above, both queries and transfers are prohibited. At first, these prohibitions might seem strange since obviously DNS requires both queries and transfers to function properly. The *allow-query* and *allow-transfer* phrases, however, can appear in both the *options*

statement and the *zone* statement. Any settings in the *zone* statement override any included in the *options* statement. If the *zone* statement contains no such settings, then the settings in the *options* statement are utilized. Given the interaction between the two statements with respect to the *allow-query* and *allow-transfer* phrases, the most secure configuration is to deny everything in the *options* statement and then permit only what is necessary in the *zone* statements. The desired behavior when a zone has been incompletely configured is to prevent any transactions from occurring.

On a caching server, a DNS STIG compliant configuration might include the following:

```
options {  
    version " ";  
    recursion yes;  
    forward "first";  
    forwarders {forwarding_servers;};  
    allow-query {trusted_hosts;};  
    allow-recursion {trusted_hosts;};  
};
```

The major differences between the authoritative and caching name server configurations are as follows:

- Caching servers allow recursion and forwarding.
- Caching servers only accept queries from defined hosts.

Note that the BIND ACLs "forwarding servers" and "trusted hosts" are utilized here. The alternative is to list the IP addresses and address ranges in the *option* phrase. Using ACLs is the preferred approach because it helps prevent against an attack in which an errant host is inserted into a potentially long lists without detection. With ACLs, auditors and administrators need only examine the *acl* statements to verify the integrity of referenced hosts within the *named.conf* file. The use of ACLs also makes the options statements easier to read.

The remainder of this section discusses each of the options in more detail.

#### 4.2.6.1 version Option

The *version* phrase in the *options* statement determines how the name server will respond to queries regarding which version of BIND is running. Responding with this information could notify the attacker of what potential vulnerabilities might exist on the name server. To protect against this scenario, BIND can be configured to provide a null response using the following phrase:



```
options {  
    version " ";  
    [ ... ]  
};
```

Alternative text is acceptable as long as it does not include version information, as in the following example:

```
version "Version not disclosed";
```

- (DNS0725: CAT III) *The DNS software administrator will configure BIND not to disclose its version number to queries.*

#### 4.2.6.2 Options Related to Query and Recursion Restrictions

A potential vulnerability of DNS is that an attacker can poison a name server's cache by sending queries that will cause the server to obtain host-to-IP address mappings from bogus name servers that respond with incorrect information. Once a name server has been poisoned, legitimate clients may be directed to non-existent hosts (which constitutes a denial of service) or, worse, hosts that masquerade as legitimate ones to obtain sensitive data or passwords.

To guard against such a scenario, name servers authoritative for .mil domains must be separated functionally from name servers that resolve queries on behalf of internal clients. In addition to enforcing this separation, organizations must ensure that the caching servers only accept queries from trusted clients under the theory that they are much less likely to attempt poisoning attacks than unknown external hosts. These general requirements were first stated in a previous section on Restrictions on Recursion.

The following BIND options can implement the desired DNS architecture:

- recursion
- forward
- forwarders
- allow-query
- allow-recursion

On an authoritative server, if the DNS software administrator sets *recursion* to "no" and leaves *forward* undefined, then poisoning will be prevented because the name server will never obtain host-to-IP address mappings from external hosts. As discussed above, the *allow-query* option appears in both the *options* statement and the *zone* statement. To keep the default behavior as restrictive as possible, the DNS software administrator sets the *allow-query* option to "none" and permits queries as required in the *zone* statement. With regards to the *options* statement, these objectives are accomplished with:

```
options {
```

```
    [ ... ]  
    recursion no;  
    allow-query {none};  
    [ ... ]  
};
```

The “recursion no” phrase implements the first general requirement listed in the Restrictions on Recursion section, which prohibits recursion on an authoritative name server.

On a caching server, the DNS software administrator wants to permit recursion and forwarding, but limit it to trusted hosts. In the *named.conf* file, this is accomplished with:

```
options {  
    [...]  
    recursion yes;  
    forward "first";  
    forwarders {forwarding_servers};  
    allow-query {trusted_hosts};  
    allow-recursion {trusted_hosts};  
    [...]  
};
```

The “allow recursion {trusted\_hosts;}” phrase implements the second general requirement listed in the section on Restrictions on Recursion, which allows caching servers to accept recursive queries only from the IP addresses and address ranges of known supported clients.

The terms “forwarding\_servers” and “trusted\_hosts” in these options are BIND ACLs defined with the *acl* statement. As mentioned, the use of these ACLs makes the *named.conf* file easier to read, manage, and audit.

Another approach to limiting queries is to use the *option* statement’s *blackhole* phrase, which forces the name server to drop any query from a host whose source address matches the phrases address match list. In a sense, the *blackhole* phrase is the inverse of the *allow-query* phrase, specifying a list of hosts from which queries are prohibited rather than accepted, as is the case with *allow-query*. For example, if a site does not use RFC 1918 private address space (e.g., 10.0.0.0 network addresses), then any request coming from such an address would be illegitimate and should be dropped. In other cases, there may be known address ranges from which DNS attacks are likely. Despite these situations, use of the *blackhole* phrase is not recommended because these types of address restrictions are more efficiently implemented and easier to maintain in firewall rules and router ACLs. DNS software administrators should not be expected to frequently update the *blackhole* list based on emerging and rapidly changing threats.

#### 4.2.6.3 Options Related to Zone Update Notification

The Zone Transfers section lists a general requirement that name servers only accept notifications of zone changes from the hosts authoritative for that zone. If all of a zone’s NS

records are valid, then the default behavior in BIND complies with this requirement and does not require the DNS software administrator to take any additional action.

In some cases, the DNS software administrator must implement a non-default configuration to comply with operational or mission critical requirements. If this is the case, the DNS software administrator must have an understanding of the *named.conf* options that govern how master name servers notify other hosts of zone changes and when slave servers will accept notifications. If none of these options are selected, the resulting behavior represents an acceptable security risk.

The three phrases within the *options* statement that govern this behavior are:

- *notify* – which turns notification on or off (defaults to on)
- *also-notify* – which defines servers other than those listed in NS records that will be sent notifications (defaults to none)
- *allow-notify* – which defines from which servers a slave will accept notifications (defaults to the master name server only)

In summary, the default behavior is as follows:

- A master name server sends notifications to all other authoritative servers in a given zone (as defined by NS records for that zone) whenever the zone changes.
- A slave name server only accepts notification of zone changes from its master.

A requirement for non-default behavior is uncommon. Some sites may have a *stealth slave*, i.e., a name server for which there is no NS record in the zone, but which nonetheless is intended to be authoritative for that zone. If a stealth slave does not appear in an *also-notify* phrase in the *options* statement, it would not receive notifications but it still could periodically pull zones without notifications.

#### 4.2.6.4 IQUERY Simulation

In earlier versions of BIND, there is a security vulnerability associated with simulating the obsolete query type IQUERY, which was implemented with the *fake-iquery* phrase in the options statement. BIND 9 never performs IQUERY simulation.

- (DNS0730: CAT III) The DNS software administrator will ensure that the *named.conf* options statement does not include the phrase "fake-iquery yes;".

#### 4.2.6.5 RFC 2308

A potential vulnerability in DNS is that a name server will provide a list of name servers (i.e., NS records) along with negative responses to queries. Presumably, this response is intended to provide additional assistance to a client trying to resolve a host name to an IP addresses. The

general idea is, "I cannot help you but here are lists of servers I know that might be able to." Unfortunately, if the query were not sent from a legitimate lost client but an attacker attempting to infiltrate a network, then the name server would provide valuable information to that attacker.

To ensure that this undesirable behavior is not supported on a given name server, the DNS software administrator must set the option *rfc2308-type1* to "no", which is the default behavior in BIND 9.

- (DNS0735: CAT II) The DNS software administrator will ensure the *named.conf* options statement does not include the phrase "rfc2308-type1 yes;".

#### 4.2.7 view Statement and Split DNS

To maintain the confidentiality of information about a network, a DNS software administrator may want to restrict untrusted outsiders from learning the names and IP addresses of hosts within the internal network while also allowing external hosts to reach public web sites, e-mail servers, and other services supporting the organization's mission. The solution is to *split* the DNS into two: a version accessible to outsiders that contains limited information and another version for internal use that contains all information.

In the past a split DNS was achieved by hosting different zone data on different servers – one residing on a publicly accessible DMZ and another on the internal network. Another complex but acceptable alternative is to run two instances of BIND on the same server, each bound to a unique IP address and loading a distinct zone file. This mechanism for achieving split DNS is sometimes referred to as *split brain* DNS.

BIND 9 introduces the ability to generate *views* on zones that can direct clients to different records based on the source IP address of the client generating the query. This greatly simplifies the administration of a split DNS and can also reduce the number of name servers required to support the DNS infrastructure.

Views are implemented using the *view* statement in the *named.conf* file. The syntax of the *view* statement is as follows:

```
view view_name {  
    match-clients { address_match_list };  
    match-destinations { address_match_list };  
    recursion yes/no;  
    [ view_option; ... ]  
    [ zone-statistics yes/no; ]  
    [ zone_statement; ... ]  
};
```

An example of *view* statements in practice might be as follows:

```
view "internal" {
```

```
match-clients {trusted_hosts;};  
recursion no;  
zone example.disa.mil {  
    type master;  
    file "example-internal.db";  
};  
};  
view "external" {  
    match-clients {any;};  
    match-destinations {dmz;};  
    recursion no;  
    zone example.disa.mil {  
        type master;  
        file "example-external.db";  
    };  
};
```

First, note that the *zone* statements in the example above only contain the minimum information to support views and, therefore, are incomplete security configurations. The next section provides a more complete discussion of the *zone* statement.

The ordering of the views is often critical because the source and destination IP addresses of a query are matched against each view in the order they appear in the *named.conf* file. Once a query matches one of the views, no further views are checked. As a result, the DNS software administrator must be cognizant of potential problems that might result from view sequencing.

In this example, the internal view precedes the external view. Consequently, internal clients are always served by the internal view because the phrase "match-clients {trusted\_hosts;};" would always be true in this circumstance. Alternatively, had the external view preceded the internal view, internal clients would be captured in the phrase "match-clients {any;};" included in the external view. Thus, the internal clients would utilize the external view to resolve names of hosts in the DMZ. In this case, the ordering does not impact DNS behavior, but in more complex cases, it might.

The site may implement a split DNS in whichever way it deems appropriate so long as it meets the requirements delineated in this STIG.

#### 4.2.8 zone Statement

Zone configuration is achieved with the *named.conf* *zone* statement. The *zone* statements dictate query handling, permissible zone transfers, the location of zone data and more. They tie together and make sense of the *acl*, *key*, *server*, and *options* statements that precede them.

The *zone* statements come with their *zone* options, which are similar to the global *named.conf* options. Recall that zone options supersede global options. The general strategy is to make the global options as restrictive as possible and then loosen these restrictions in *zone* statements based on business requirements. This strategy follows the security principle of least privilege.

The *zone* statements can either appear on their own within the *named.conf* file or be nested within *view* statements. However, in the presence of a *view* statement, a standalone *zone* statement would be rendered superfluous – it would never be utilized by any view and, consequently, by any client. As a result, if any *view* statements are explicit in the *named.conf* file, then all *zone* statements must be nested within *view* statements. Most *zone* statements in the DOD computing environment will be located within *view* statements.

The syntax of security-relevant components of the *zone* statement is as follows:

```
zone zone_name {  
    type ( master | slave | hint | stub | forward );  
    allow-query { address_match_list };  
    allow-transfer {  
        address_match_list;  
        key key_ID  
    };  
    allow-update {address_match_list };  
    file string;  
};
```

The general security requirements related to the *zone* statement, which are listed in the Zone Transfers section; state that zone transfers must be limited to authorized name servers and that these transfers be cryptographically authenticated. An example STIG-compliant *zone* statement in practice might be as follows:

```
zone example.disa.mil {  
    type master;  
    allow-query {any;};  
    allow-transfer {  
        key ns1.example.disa.mil_ns2.example.disa.mil;  
    };  
    allow-update {none;};  
    file "example.db";  
};
```

In the example, the phrase "allow-query {any;};" permits DNS queries from any host, which is the default behavior and the appropriate configuration for an authoritative server. Recall, however, that in the *named.conf options* statement the default was changed to "allow-query {none;};" (i.e., deny all queries unless explicitly told otherwise). Accordingly, queries must be activated in the *zone* statement for it to function properly.

In the example, the use of a key in the *allow-transfer* phrase activates transaction signatures that ensure that transactions are authenticated. The slave server ns2.example.disa.mil is the only other server that has this key, so this phrase ensures that zone transfers are restricted to the slave only.

In the example, the phrase "allow-update {none;};" disables dynamic updates. Dynamic updates are useful when hosts frequently change their IP address (e.g., workstations on a local area network utilizing DHCP). Unfortunately, dynamic updates represent a serious security vulnerability because they allow an attacker to modify zone records using bogus updates from spoofed addresses. Therefore, dynamic updates must be cryptographically authenticated if they are permitted.

In the example provided above, the *zone* statement restricted zone transfers to those signed by a particular TSIG key. While cryptographic authentication satisfies the general requirements in this STIG, in a maximum-security posture zone transfers and dynamic updates are restricted both by IP address and by TSIG key. Such a configuration would provide additional protection against a scenario in which someone is able to obtain the key (e.g., via theft) and then use it from an unauthorized host, because the adversary both would have to obtain the TSIG key and spoof the IP address of the slave name server.

The implementation of this configuration is somewhat tricky because it utilizes the negated match feature twice – first in an *acl* statement and a second time within the *allow-transfer* statement. Nevertheless, this configuration is recommended. An example is as follows:

```
acl not_ns2.example.disa.mil {!239.10.241.129; any; }

zone example.disa.mil {
    ...
    allow-transfer { !not_ns2.example.disa.mil;
        key ns1.example.disa.mil_ns2.example.disa.mil;
    };
    ...
};
```

In this example, the IP address of the ns2.example.disa.mil is 239.10.241.129. The *acl* statement defines an access list that includes every host except the ns2. The “!” character before any list entry means not – in this example, not 239.10.241.129. If a negation element is true, BIND moves to the next element in the list, in this example, “any”.

The negation operator in the *allow-transfer* statement, the first element is effectively a double negative. It is true when the requesting host is 239.10.241.129 and false when it is not. Recall, however, that BIND continues to check other items in the list when a negation statement is true. Therefore, through the use of the double negative, the requesting host must both have IP address 239.10.241.129 and meet at least one other condition in the access list. In this case, the other condition is that the host signed its request with the correct TSIG key. If either condition is not met, then access is denied. DNS software administrators should secure zone transfers and dynamic updates using this technique when it is feasible to do so.

## **5. CONFIGURATION OF WINDOWS 2000/2003 DNS**

Many DOD organizations utilize Windows 2000 or 2003 Server<sup>7</sup> technology to support their Windows networks. Windows 2000 is heavily dependent on DNS and comes bundled with its own proprietary DNS software. Unfortunately, Windows 2000 DNS does not support several important features. Therefore, personnel who support Windows DNS should prepare for a forthcoming mandate to migrate to BIND, or other DOD-mandated alternative.

This section provides interim guidance for Windows 2000/2003 DNS implementations in use prior to the migration to BIND or other DOD-mandated alternative.

### **5.1 Comparison of Windows 2000/2003 DNS and BIND**

To understand the security posture of Windows DNS relative to BIND 9.3.x, one must know the features that each of the DNS software suites support. There are some features supported in BIND 9.3.x that are not supported in Windows 2000/2003 DNS. Similarly, there are features in Windows 2000/2003 DNS that are not currently supported in BIND 9.3.x. For Windows Server 2003, Microsoft made minor changes to the DNS software such as the addition of Stub Zones and Conditional Forwarding features. The changes to the DNS software from Windows 2000 to 2003 are subtle and therefore any references to Windows DNS or Windows 2000 also refers to Windows 2003.

#### **5.1.1 BIND 9 Security Features Not Present in Windows 2000 DNS**

BIND 9.3.x and above, has the following features that Windows DNS does not have:

- TSIG support for zone transfers and queries
- DNSSEC support, especially the ability to cryptographically sign zones
- Ability to restrict queries and recursion by IP address
- Ability to explicitly define split DNS configurations on an authoritative name server

TSIG is required for zone transfers, but not for queries. Zone transfers that occur without the cryptographic authentication that TSIG provides are less secure than zone transfers that rely on other forms of authentication such as source IP address. Reliance on weak authentication methods for zone transfers increases the possibility that adversaries can alter authoritative records, which has significant implications for the integrity of the DOD DNS infrastructure. Moreover, if an adversary is able to perform a zone transfer, he or she can map many network resources with a single transaction, which is much less likely to be detected than an authorized scan of the network.

DNSSEC provides cryptographic authentication of query responses through the use of signed zones. BIND 9 name servers can check the validity of signatures on behalf of clients, while Windows DNS servers have no ability to do this. This is a major shortcoming of Windows DNS

---

<sup>7</sup> As mentioned in the introduction, Windows 2000 and 2003 DNS are very similar and have identical security configuration requirements. Throughout this document, references to Windows will apply to both Windows 2000 and 2003.



because DNSSEC signed zones are a cornerstone of future efforts to secure DOD DNS transactions.

Windows DNS is able to restrict zone transfers by IP address; however, unlike BIND 9, cannot restrict queries and recursion by IP address. As mentioned, IP address restrictions are a weak form of authentication for zone transfers, but they currently are the only practical form of restricting transactions that involve direct interaction between name servers and DNS client resolvers. Without IP address restrictions, sensitive host records may be sent to an unauthorized user. In addition, the potential exists for adversaries to poison the name server's cache if it will accept their queries.

Address restrictions are particularly important in the implementation of split DNS, in which the objective is to provide different host information to different clients depending on their network location. With Windows, the only means to create a split DNS is to use two independent name servers authoritative for the same zone and then use firewall rules or router ACLs to restrict access to those name servers by network address. This is a highly artificial and undesirable architecture. In BIND 9, one can use views to distinguish what information will be given to whom.

### **5.1.2 Windows 2000/2003 Security Features Not Present in BIND 9**

The notable advantages of Windows DNS relative to BIND 9 are as follows:

- Support for (Generic Security Service) GSS-TSIG
- Granular access rights to DNS zones and objects

The key feature in Windows DNS that is absent in BIND 9 is GSS-TSIG, which is the basis for the Windows DNS "Secure Dynamic Updates" feature. BIND 9 can be configured to accept dynamic updates, but can only authenticate them through TSIG shared secret keys. Very few clients support TSIG today, and even if they did, supporting such clients in a large environment would impose an unacceptable administrative burden given the large number of secret keys the system would require.

GSS-TSIG differs from ordinary TSIG in that it leverages an existing security relationship for authentication and session keying. In the case of Windows, DNS can be integrated with Microsoft's Active Directory (AD), a distributed database of user and machine information and other records. AD uses the accounts associated with computers defined within AD to provide the existing security relationship for GSS-TSIG. This eliminates the administrative burden of managing separate TSIG keys.

With the use of GSS-TSIG, Windows offers reasonable (but not absolute) assurance that dynamic updates are legitimate – a powerful advantage when a network supports DHCP clients, who's IP addresses may change frequently. If an organization has a requirement for DHCP clients to serve data or an application to other users, then secure dynamic updates offer the only feasible means to operate under those conditions. The alternative is for the DNS database

administrator to manually update the zone records whenever a DHCP client obtains a new IP address, which is an unreasonably burdensome approach.

The second advantage of Windows 2000 DNS relative to BIND 9 is that AD integration allows an administrator to grant user groups AD permissions to DNS zones and objects. This feature offers far more granular control of DNS objects than BIND does. For example, an SA can assign various other DNS software administrators different levels of access to each zone depending upon their roles and responsibilities. In BIND, on the other hand, an SA cannot grant other administrators write permissions to particular sections of the *named.conf* file, where options for all zones are defined. In BIND, an administrator either has full control or no control.

Finally, Windows requires a number of SRV records in the DNS for the LDAP, Kerberos, and catalog services that are critical to the operation of Windows. Without these services, users could not log on, obtain directory information or perform other critical functions. When Windows servers utilize Windows 2000/2003 DNS, these records are automatically entered into the zone upon installation, and automatically modified later when there are changes. This feature is desirable because the SRV entries can be intimidating for administrators new to either Windows 2000 or DNS in general. In fact, the transfer of these records from Windows 2000 DNS to BIND 9 is not difficult. They can be simply cut and pasted from a Windows DNS zone into a BIND 9 zone. Changes to the SRV records would require manual intervention, but this is not expected to be a frequent occurrence on most networks.

### **5.1.3 Implications of BIND 9/Windows 2000/2003 DNS Differences**

The DOD has a strong preference that organizations deploy BIND 9 to support DNS for the following reasons:

- BIND 9 supports several key required security features that Windows 2000 DNS and most other DNS software products do not.
- DISA examines, compiles, and rigorously tests BIND 9 code but does not have similar access to Windows DNS code nor does it perform similar evaluation of that code.
- DISA has a large investment in BIND 9 software development and a significant voice in the future direction of BIND, while it has less influence over Microsoft's future DNS development effort.

At the same time, Windows 2000/2003 has certain features that make it an acceptable solution for many environments. The use of Windows DNS should be limited in such a manner to reduce the cost and complications of the future migration to BIND. For instance, there is no reason to use Windows 2000 DNS as a caching name server or for any application that requires recursion. Microsoft Windows clients can still send queries to BIND 9 name servers, which can in turn query authoritative Windows 2000/2003 DNS servers when needed.

Once GSS-TSIG is available in BIND 9, there should no longer be a need to deploy Windows DNS servers. At that point, BIND 9 will provide all the services and features that make Windows DNS an attractive alternative for some applications today. Of course, both software functionality and security requirements can change rapidly so DOD guidance on this issue may change as well.

The remainder of this section provides specific implementation guidance for organizations that still choose to deploy Windows 2000/2003 DNS.

## **5.2 Network Restrictions**

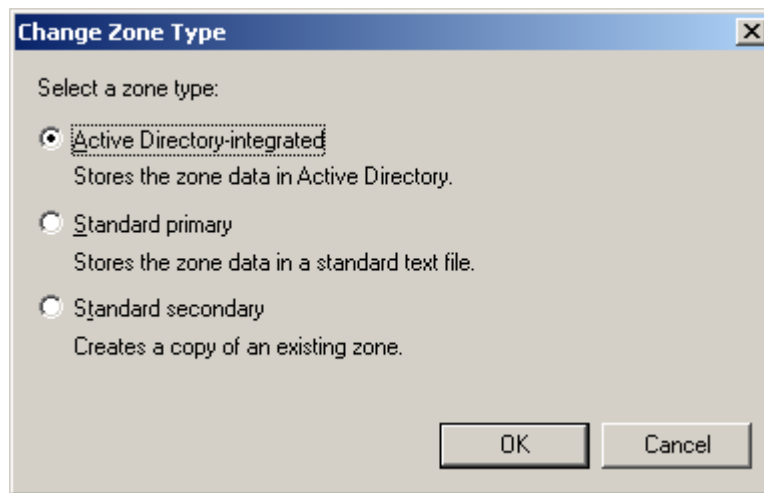
As mentioned, Windows DNS cannot restrict transactions other than zone transfers by IP address. Instead, organizations must use firewalls or router ACLs to enforce those restrictions. These additional protections are typically a component of defense-in-depth security architecture, but in this case they comprise the first and likely only line of defense against unauthorized access to zone records.

- *(DNS0800: CAT II) The IAO will ensure that firewall rules or router ACLs prevent unauthorized hosts from outside the enclave from querying Windows DNS servers.*

## **5.3 Secure Dynamic Updates and Active Directory**

The most compelling reason for using Windows DNS occurs if there is a requirement for dynamic updates, which must be cryptographically authenticated as specified previously in this STIG. Windows DNS provides cryptographic authentication through the Secure Dynamic Updates feature. Appropriate implementation of this capability requires a number of supporting security settings. For example, the feature is only available when DNS is integrated with AD. Moreover, DNS can only be integrated with AD if the DNS service runs on a domain controller.

AD integration can be configured by zone, as is shown in the dialog box below and as the requirement states, the zone must be Active Directory-integrated:

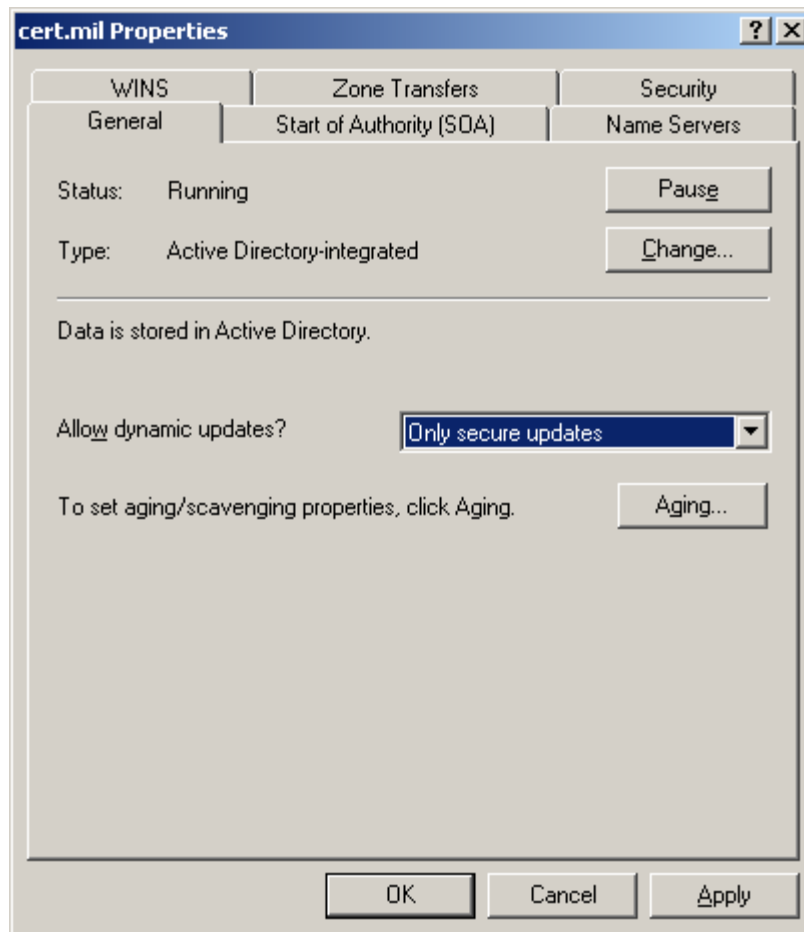


**Figure 5-1. Change Zone Type**

Once a zone is configured for AD integration, an administrator can require secure dynamic updates on the “General” tab of the “Properties” dialog box for the zone, as is shown in the following dialog box<sup>8</sup>: The Windows 2003 dialog box looks very similar, however it shows a warning about vulnerabilities associated with non-secure updates.

---

<sup>8</sup> This example and many that follow show a fictitious cert.mil domain. They do not reflect the settings of an actual cert.mil name server.



**Figure 5-2. Dynamic Updates**

To prevent unauthorized dynamic updates of records, the ability to modify the records should be limited to the individuals or objects with authority over that record. In many cases, this may be a machine account. Fortunately, SAs can implement these restrictions using AD permissions. A SA would grant a host change permissions to the AD object corresponding to its zone record only and no other objects. This permissions scheme is typically created by default, but SAs should verify that it has been implemented as expected.

Importantly, AD read permissions usually must be granted to all users because they will need them to get responses to queries. Furthermore, there should be no restrictions on auto-generated SRV records because users typically need to issue queries for these records prior to authenticating. For example, a background process must locate a domain controller to determine which machine the user must communicate with in order to log on. This creates a potential "catch-22." If a user must first log on to query DNS and must query DNS to log on, then the user can accomplish nothing. To resolve this dilemma, the EVERYONE group should be granted read permissions to the LDAP and Kerberos SRV records, which is the default in AD-integrated Windows 2000 DNS. Most SAs are conditioned through STIG guidance and elsewhere to remove the EVERYONE group from all permissions assignments, but this is one case in which that advice should not be followed.

There is a significant vulnerability potential when the DHCP service runs using the computer account of a Windows Domain Controller, as is the default Windows configuration. This account has full control over all DNS objects stored in Active Directory. In this case the DHCP server has access to modify the SRV (and other) records for all the Domain Controllers. When these records were replicated to other domain controllers (when AD Integrated DNS is used as required by the STIG), all the Windows DNS servers could potentially be compromised.

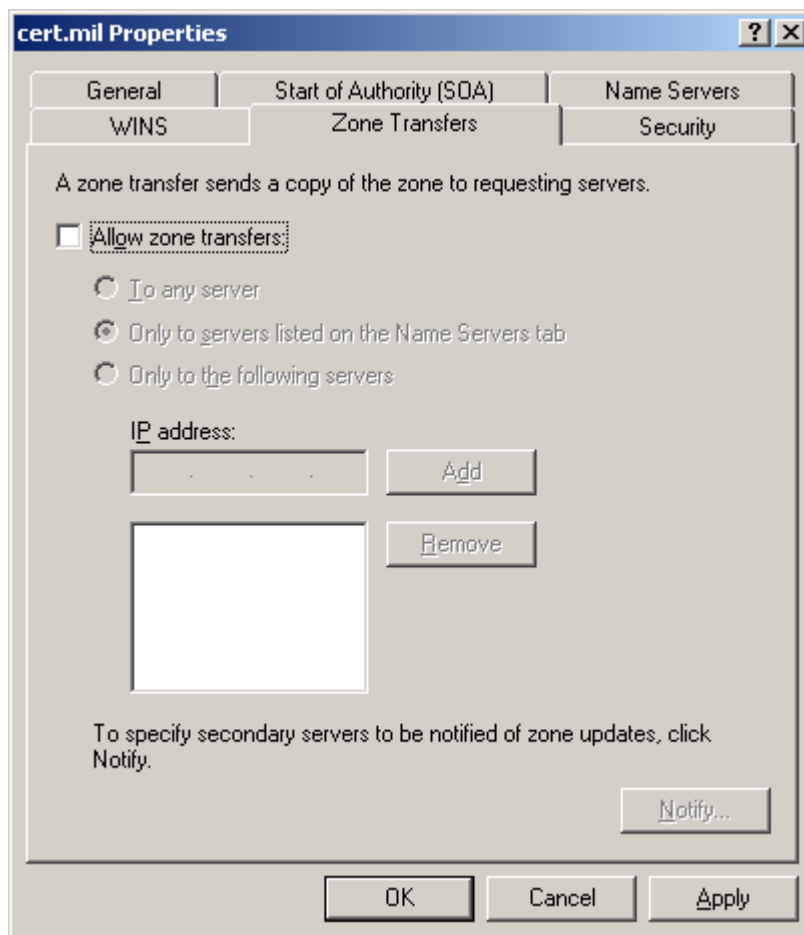
A built-in security group, DNSUpdateProxy, is provided as of Windows 2000. This group can update DNS records for clients without becoming the owner of the records. When DHCP servers are added as members of this group, any of the (member) DHCP servers can update the records. The first user that is not a member of the DNSUpdateProxy group to modify the records associated with a client; becomes the owner. There is a limited vulnerability for all servers (even non-domain controllers) on which a DHCP service runs. The DNS records associated with the DHCP server host could be modified by other DHCP servers that are members of the DNSUpdateProxy group. In order to prevent this from occurring, DHCP should not be installed on a domain controller if the group DNSUpdateProxy is utilized.

- *(DNS0805: CAT I) The SA will disable the DHCP server service on any Windows 2000/2003 DNS server that supports dynamic updates.*
- *(DNS0255: CAT II) The SA will disable the use of DHCP on any Windows Domain Controller.*
- *(DNS0260: CAT II) The IAO will ensure computer accounts for DHCP servers are not members of the DNSUpdateProxy group.*

## 5.4 Zone Transfers

Windows allows for two ways of synchronizing zone data across name servers: (1) traditional RFC-compliant DNS zone transfers; and (2) AD-replication. The latter only works when Windows DNS is integrated with AD, which makes each of the DNS records an AD object. The Windows 2000/2003 DNS implementation of traditional zone transfers does not meet the STIG requirement that the transfers be cryptographically authenticated using a technology such as TSIG. Fortunately, AD-replication is cryptographically authenticated. Therefore, the solution in

a pure Windows 2000/2003 DNS implementation is to integrate DNS with AD and disable zone transfers. This configuration –shown in the example on the following page – will force all zone updates to rely on AD replication exclusively.



**Figure 5-3. Zone Transfers Tab**

In a heterogeneous environment, the problem is more complex. For instance, if both BIND and Windows 2000 DNS name servers are authoritative for the same zone, they will need to execute zone transfers, but they cannot employ TSIG because Windows 2000 DNS does not offer this capability.

For this reason, mixed environments are highly discouraged. Instead, DNS software administrators should create subordinate zones for Microsoft services. For example if BIND servers are authoritative for hosts located in example.disa.mil, then hosts that require secure dynamic update can be placed in a lower-level domain named ms.example.disa.mil or ad.example.disa.mil. Windows 2000 DNS servers will be authoritative for the lower-level domain. The name servers authoritative for example.disa.mil will reference the lower-level Windows 2000 DNS servers through NS records.


In cases in which domain segregation is not feasible, zone transfers must still be cryptographically authenticated with an alternative to TSIG. The most practical alternative is to create host-to-host virtual private network (VPN) connections between the name servers that need to perform zone transfers. The keys used to encrypt communication between the name servers serve the role that would have been filled by the TSIG keys had both ends of the transaction been able to support TSIG.

- *(DNS0810: CAT I) The SA will either configure Windows 2000/2003 DNS to prohibit zone transfers or will implement a VPN solution that requires cryptographic authentication of communicating devices and is used exclusively by name servers authoritative for the zone.*

## **5.5 Forwarders and Recursion**

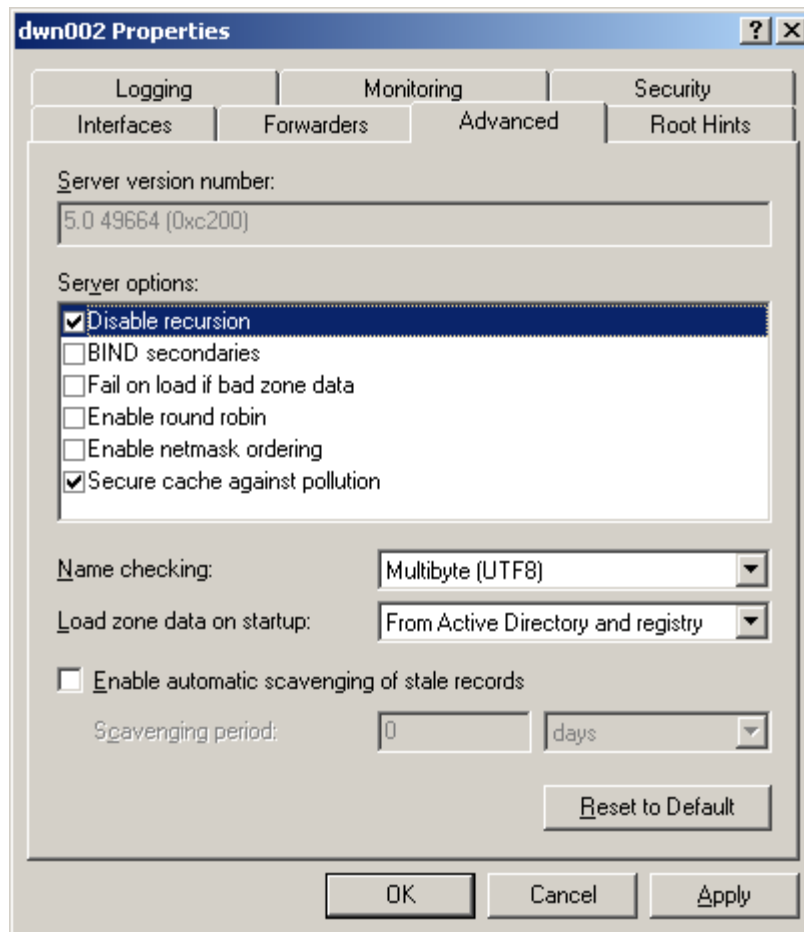
Windows DNS has historically been more vulnerable to cache poisoning attacks than BIND as the algorithm used for answering recursive queries also makes it more prone to self-imposed denial of service attacks and as an amplification device for attacks on other DNS servers. Additionally Windows DNS does not allow for the fine-grained access control restrictions that are allowed by BIND and other recursive DNS appliances. Therefore, Windows 2000/2003 DNS should not be deployed as a caching name server. Consequently, the use of forwarders and recursion is prohibited on Windows 2000/2003 DNS servers. To preclude forwarding, an administrator can make sure that the "Enable Forwarders" check box is not marked on the "Forwarders" tab of the name server properties, as shown below:





**Figure 5-4. Forwarders Tab**

To prohibit recursion, the administrator can select the server option "Disable recursion" on the "Advanced" tab of the name server properties. The administrator can also select "Secure cache against pollution." Recall that the rationale for prohibiting recursion is to do just that. These settings are shown in the following illustration. Secure cache against pollution is on by default in Windows Server 2003.



**Figure 5-5. Advanced Tab**

- (DNS0815: CAT II) The SA will disable forwarders on an authoritative Windows 2000 DNS server.
- (DNS0820: CAT II) The SA will disable recursion on an authoritative Windows 2000 DNS server.

The configuration shown above is also necessary to comply with the general requirement in the section on *Restrictions on Recursion*, that the DNS software administrator prohibit recursion on authoritative name servers.

## 5.6 WINS Integration

Prior to Windows 2000, Windows operating systems utilized NetBIOS rather than DNS names for most internal network communication, including file, print and messaging services. Microsoft developed the Windows Internet Naming Service (WINS) to provide a central

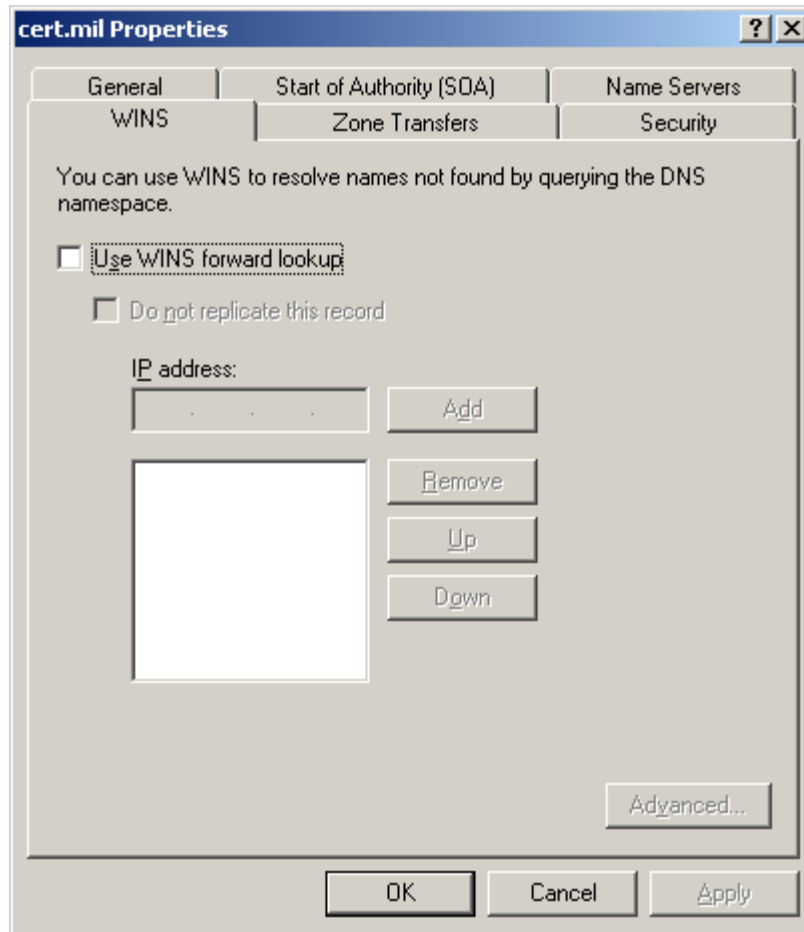
host-name-to-IP-address resolution capability and other services similar to those provided by DNS. WINS, however, supports Windows NetBIOS clients exclusively.

When Microsoft developed its own DNS server software, it included options for DNS/WINS integration. This feature allows a client to first query DNS, which forwards the request to WINS if the record is not found. Importantly, this integration is not required to resolve names because clients can also be configured to send a query to a WINS server when DNS replies that a host name or IP address cannot be resolved.

WINS is not secure because, among other reasons, it will accept dynamic updates without authentication, thereby allowing adversaries to easily modify WINS records. Integrating DNS and WINS creates the potential for DNS to be poisoned with bogus entries in WINS. Therefore, administrators must never integrate WINS with DNS.

With the introduction of Windows 2000, Microsoft embraced DNS as the primary mechanism for host name definition and address resolution. For networks that deploy only Windows 2000 or later versions of Windows, there is no need for WINS. On these networks, DNS can support all desired functionality. WINS is still required for some legacy Windows NT and Windows 9x applications, but pre-2000 Windows computers can be configured to query DNS and WINS separately, eliminating any need for DNS/WINS integration.

To ensure that DNS is not integrated with WINS on a Windows 2000 DNS server, an administrator should make sure the "Use WINS forward lookup" is not checked on the "WINS" tab on the properties dialog of each zone, as shown on the following page.



**Figure 5-6. WINS tab**

- (DNS0825: CAT I) The SA will configure Windows 2000 DNS to prohibit WINS lookup.

## 5.7 Logging

The general requirements for logging (as specified in *Section 3.5, Logging*) are that the DNS service should log success and failure of the following:

- Start and stop of the DNS service
- Zone transfers
- Zone update notifications
- Dynamic updates
- Queries

Events related to the start and stop of the DNS service appear in the Windows 2000 System Event Log. Other events are logged to a file named `%systemroot%\system32\dns\dns.log`. Windows 2000 DNS has its own logging facility, which is primarily designed to debug DNS problems rather than maintain a record of DNS transactions. The Windows 2003 logging tab has been split into two; debugging and event logging. To implement the general logging requirements specified in the section on *Logging*, the DNS software administrator must select the “Query,” “Notify,” and “Update” debug logging options as shown in the dialog box below.

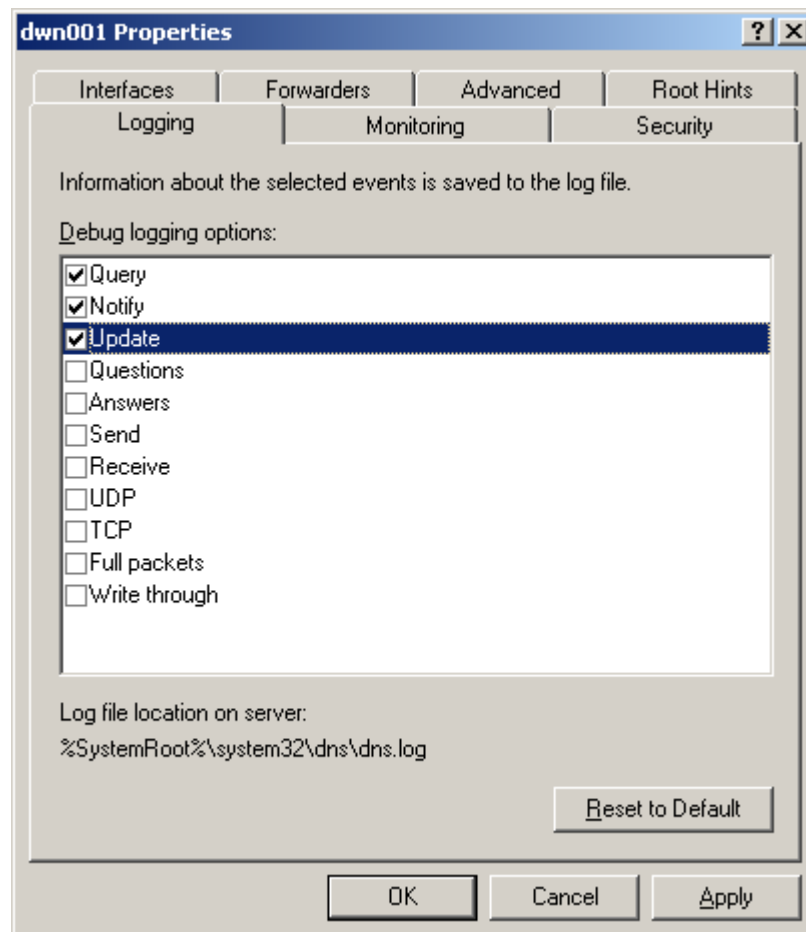


Figure 5-7. Logging Tab

This page is intentionally left blank.

## 6. CISCO CSS DNS

This section provides guidance for the configuration of Cisco Content Services Switch (CSS) DNS.

### 6.1 CSS DNS Background

Cisco CSS is a network device that optimizes a user's Internet experience by directing the user to the server best able to respond to the user's request. At a basic level, this can involve load balancing, in which requests are distributed among several servers to optimize the processing load on each. More sophisticated CSS architectures utilize content routing, in which CSS forwards a user to the *closest* server based on the location of the client (i.e., the one with the fastest current response time given that location). CSS DNS can also sense when a server is down and temporarily remove host records for that server until it is back on line.

CSS DNS is a core technology behind content routing. For example, suppose identically configured web servers around the world supported a web site with the URL `www.test.mil`. CSS DNS might provide an Air Force officer in the Pentagon with the IP address of a `www.test.mil` server in Quantico, Virginia while CSS DNS might provide a Marine at Pearl Harbor the IP address of a server with identical content in Hawaii. In the absence of content routing, both the warfighter in the Pentagon and at Pearl Harbor would be able to traverse the NIPRNet to a server on the other side of the globe, but content routing provides faster response times by directing users to servers closer by. Moreover, it utilizes the network more efficiently because traffic does not have to travel over long-haul routes.

Traditional implementations of DNS, such as BIND, cannot support this functionality. At best, these traditional implementations provide a *round robin* capability that would alternate between providing a user the IP address of the server at Quantico server and the one at Pearl Harbor, but this crude form of load balancing does not optimize the user experience in the same manner as content routing. In addition, the traditional implementations of DNS do modify response behavior when a server is temporarily unavailable. Thus, good reasons exist for an organization that has geographically distributed web services to want to deploy a DNS product with the features present in Cisco CSS DNS.

On the other hand, Cisco CSS DNS lacks many of the key security features of BIND, including TSIG authentication and access control lists. Therefore, without mitigating controls – which will be discussed later in this section – Cisco CSS DNS' default configuration provides considerably less protection against DNS attacks than a BIND name server that complies with the requirements listed in this STIG.

Cisco CSS DNS also lacks support for DNSSEC, which means that DNSSEC-enabled BIND servers will not be able to cryptographically validate any host-to-IP address mappings issued from a CSS name server. Arguably, the value of DNSSEC zone record validation rises with the frequency that a record changes because an adversary that spoofs or changes an IP address is more likely to go unnoticed when the address is already changing often under normal conditions. Indeed, CSS DNS' ability to rapidly change the DNS response to a query based on current

conditions is precisely the environment under which a user would most want the assurance that DNSSEC provides. Consequently, at this time, the deployment of CSS DNS (and similar technologies) is a significant impediment for the successful adoption of DNSSEC.

## 6.2 CSS DNS Zone Delegation

CSS DNS plays a special role in improving the performance of key applications, but in most large enterprises, BIND and Windows name servers dominate the DNS infrastructure. CSS DNS must fit into the enterprise DNS architecture to be a functional component of that architecture. However, CSS DNS does not currently support zone transfers with non-CSS DNS name servers. Accordingly, its interaction with other name servers must involve either hosting child zones of higher-level zones (likely supported by BIND name servers), or serving as a parent to lower-level zones.

### 6.2.1 CSS DNS Childhood

CSS DNS *childhood* refers to cases in which the CSS DNS authoritative records are in a child zone of a higher-level parent zone. So that readers can better appreciate security concerns and requirements, we will use the `www.test.mil` example introduced earlier to explain on a basic level how CSS DNS zone delegation works in practice. CSS DNS is highly complex and a complete understanding of its options and behavior is beyond the scope of this STIG.

In a typical DNS implementation (i.e., one not utilizing CSS DNS), `www.test.mil` would be defined as an A record in a zone file. If `www.test.mil` were actually several servers supporting a virtual host, there might be several A records. The records would appear in a zone file shown in Partial Zone File Example 1:

#### **Partial Zone File Example 1: Traditional DNS**

```
$ORIGIN test.mil
www      IN      A      150.10.34.200      ; Quantico server
          IN      A      254.35.39.5       ; Pearl Harbor server
```



This configuration provides a round robin effect in which the name server alternates between providing the address of the Quantico and Pearl Harbor servers in response to queries for www.test.mil. As discussed, the round robin load balancing behavior is not equivalent to content routing, in which DNS clients would be provided either the Quantico or Pearl Harbor IP address depending upon the location of the client. To implement a content routing architecture using CSS DNS, the zone file supporting the test.mil zone might appear as follows:

### **Partial Zone File Example 2: Delegation to CSS DNS with NS and Glue Records**

```
$ORIGIN test.mil
www      IN      NS      css101.csd.disa.mil.
          IN      NS      css201.csd.disa.mil.
css101.csd.disa.mil.  IN      A        150.10.34.2
css201.csd.disa.mil.  IN      A        254.35.39.2
```

In Example 2, the name servers css101 and css201 are CSS devices in the csd.disa.mil domain. Importantly, this zone file might reside on a BIND name server, whose function might be merely to provide the DNS glue to the content routing CSS devices. Given the configuration in the zone file above, the CSS devices are authoritative for www.test.mil and, therefore, can determine whether to respond to www.test.mil queries with the Quantico or Pearl Harbor addresses based on internal configuration and algorithms.

In the first example, the server hosting the test.mil zone file maintains authority for www.test.mil through its A records in that file, a traditional approach that limits functionality. In the second example, authority for www.test.mil is delegated to CSS name servers, which can use their content routing rules to determine how to best respond to queries for www.test.mil.

The CSS commands to support this configuration on css101 might be as follows (using the specialized CSS DNS configuration syntax):

```
(config)# dns-server zone 0 tier1 "US-East" leastloaded
(config)# dns-record a www.test.mil 150.10.34.200 single kal-icmp sticky-enabled proximity
```

Similarly, the commands on css201 might be as follows:

```
(config)# dns-server zone 1 tier1 "Pacific" script
(config)# dns-record a www.test.mil 254.35.39.5 single kal-icmp sticky-enabled script
```

As mentioned, the full scope of CSS command syntax and functionality is beyond the scope of this STIG. Nonetheless, several elements of these commands illustrate some of the features of the CSS DNS implementation. For instance, note that in CSS DNS zone 0, the A record for www.test.mil server resolves to the Quantico IP address, while in zone 1 it resolves to the Pearl Harbor IP address. In addition, the descriptions "US-East" and "Pacific" suggest that the zones

correspond to geographic regions, which can assist with content routing decisions. Indeed, the final term in the *dns-record* command controls that behavior. In zone 0, responses to client queries for *www.test.mil* are based on a proximity database. In zone 1, the responses are based on the *scrip* method, which returns zone records based on a source IP address hash. Both methods use an algorithm to return an IP address to the DNS client that should improve the user's experience.

The primary security concern with regard to the type of delegation discussed above is that to implement this approach, an organization would have to migrate its authoritative records from a well-known DNS implementation with proven, tested security controls to a relatively new DNS implementation without similar controls. Therefore, this migration should only occur when the performance and availability advantages of CSS significantly outweigh the increased residual security risk of using a less mature technology.

- *(DNS0905: CAT II) The IAO will ensure Cisco CSS DNS is not utilized to host the organization's authoritative records unless DISA Computing Services supports that host in its csd.disa.mil domain and associated high-availability server infrastructure.*

There are no security requirements related to specific CSS DNS configuration options related to method of content routing or content rules in general, but the DOD may issue guidance at a future date if security vulnerabilities with a particular method or algorithm arise.

Cisco CSS utilizes load distribution algorithms that have the capability to route client requests based on server availability, network congestion levels, and the relative proximity of the client to each server, while BIND does not have these capabilities. Both BIND and Cisco CSS DNS support round-robin resolution in which the name server rotates through a list of IP addresses associated with a host name. If an organization only needs the round-robin capability and not the more sophisticated load distribution algorithms, then it should implement BIND because it is more mature and has greater security functionality than Cisco CSS DNS.

### 6.2.2 CSS DNS Parenthood

The previous section reviewed how a traditional DNS implementation (perhaps utilizing BIND) could delegate authority to CSS DNS so that an organization can take advantage of features that CSS DNS offers but BIND and many other DNS implementations do not. In this section, we review the inverse situation in which a CSS DNS server delegates authority from one of its zones to a non-CSS DNS server that supports subordinate zones.

Although it is technically possible to delegate zones within CSS DNS, there is almost never a rationale to do so because such delegation could be achieved as easily with BIND, which offers security features not present in CSS DNS. Moreover, the performance enhancing features of CSS typically would not apply to name server records because these records are obtained easily and quickly across the wide area without significant impact on a user's experience.

To eliminate unnecessary NS records and related zone delegations from CSS DNS, an administrator can use the following command:

```
(config)# no dns-record ns domain_name
```

- *(DNS0910: CAT III) The Systems Administrator will not delegate zones in Cisco CSS DNS.*

### 6.3 Simulation of Zone Transfer Restrictions in CSS DNS

CSS DNS does not support RFC-compliant zone transfers. Instead, it shares information related to DNS records with other CSS devices using the Application Peering Protocol (APP). Although APP sessions are substantially different from zone transfers in practice, they still must meet the same general security requirements because their underlying purpose (to share information on zone records) is similar.

The core requirements related to zone transfers are that an authoritative name server transfers zone information only to designated zone partners and that name servers only accept zone data when it is cryptographically authenticated.

CSS APP provides means to designate with which devices it can share zone data and to authenticate those transactions. CSS devices can define their peers using IP addresses and authenticate them using Challenge Handshake Authentication Protocol (CHAP) with a shared-secret. This configuration can be supplemented with MD5 hashing encryption. While this does not provide the equivalent strength of cryptographic authentication as BIND's TSIG HMAC-MD5, it does provide a satisfactory level of information assurance when CSS DNS operates within a trusted network environment.

To activate these features, the CSS Systems Administrator can issue the following command for each designated peer:

```
(config)# app session ip_address authChallenge shared_secret  
encryptMd5hash
```

In this command, *ip\_address* refers to the IP address of the designated peer and the *shared\_secret* is a text string up to 32 characters in length.

Other security controls also exist outside of CSS to enhance information assurance. For instance as an alternative to using the built-in authentication features of CSS APP, one can place IPSEC VPN routers in front of each CSS device that will serve to both authenticate end points and provide confidentiality to transferred data. One can also limit APP communication to an out-of-band network, which would make it considerably more difficult for adversaries to spoof the addresses of peers or hijack APP sessions.

- *(DNS0915: CAT I) The SA will activate APP CHAP authentication and MD5 hashing features for APP sessions with each APP peer with which a CSS device communicates. Alternatively, the SA will achieve authentication and encryption using an IPSEC VPN between each peer couple.*

- *(DNS0900: CAT III) The SA will ensure the shared secret in any APP session is a randomly generated 32-character text string.*
- *(DNS0920: CAT III) The SA will configure CSS to transmit APP session data over an out-of-band network if one is available.*

## **6.4 Simulation of Query Restrictions in CSS DNS**

CSS DNS does not support query restrictions, which represents a significant vulnerability in environments in which these restrictions are necessary to assure the confidentiality of zone data. The core DNS query restriction applicable to all DNS implementations is that internal records only be shared with relevant trustworthy clients. This requirement leads to split-DNS implementations in cases in which a DNS zone must support both DMZ hosts that are accessed from outside the enclave and internal hosts that are accessed inside the enclave.

Fortunately, the absence of query restrictions in CSS DNS should not pose a security problem for most CSS applications. CSS technology is focused on optimizing the performance of high-volume, highly-distributed application systems. It is likely that these application systems are either available to the public or a wide range of DOD users located in a number of different geographic locations. In this environment, a split-DNS is inappropriate because CSS DNS does not provide clients with IP addresses of hosts inside a protected enclave – most likely all of its authoritative records should be available to a broad audience.

In cases in which queries must be limited for a valid business reason, the IAO should implement one of the following alternatives to query restrictions within CSS DNS to mitigate the risk of their absence:

- Query restrictions within the relevant zone statement in the DNS configuration file on the name servers supporting the parent zone to the CSS DNS records (which prevents clients from learning the NS records required to reach the CSS devices).
- Router or firewall access control lists that only allow authorized hosts to access the CSS devices or the servers that they support.
- Security modules within the CSS-supported application that authenticate clients through PKI certificates, source IP address, or .mil reverse lookup.

## **6.5 Recursion and CSS DNS Forwarders**

CSS DNS is not vulnerable to attacks associated with recursion because it does not support recursion, but does offer a forwarder feature that sends un-resolvable or unsupported requests to another name server. This feature poses a risk because the forwarder feature merely redirects potential attacks to another name server.

If zone delegation is implemented as described in *Section 6.2.1, CSS DNS Childhood*, no reason exists for CSS DNS to receive an un-resolvable or unsupported query unless the query is indeed un-resolvable or unsupported. Although these queries may occur unintentionally, many sophisticated DNS attacks rely on improper handling of malformed requests to achieve their desired impact, whether that is to crash the server, poison the record cache, or obtain the privileges of the DNS process. In this environment, CSS DNS forwarding may be a back door means of reaching a name server than would have otherwise been inaccessible to the attacker. Given the near certainty that the forwarder will not provide requested records to a legitimate client if a STIG-compliant CSS DNS cannot do so, the best course of action is to drop these requests rather than forward them.

To disable forwarders in CSS DNS, the Systems Administrator enters the following commands:

```
(config)# no dns-server forwarder primary  
(config)# no dns-server forwarder secondary
```

- *(DNS0925: CAT II) The SA will disable forwarders in Cisco CSS DNS.*

This page is intentionally left blank.

## 7. STANDARD OPERATING PROCEDURES FOR DNS

To secure DNS, a secure baseline configuration must be implemented and must be managed securely from that point forward. Monitoring DNS transactions and planning for contingencies is also crucial to a secure DNS infrastructure. This section covers these operational aspects of DNS security.

A recommended reference is RFC 2870 (Root Name Server Operational Requirements). This RFC lists mandates for root servers, but its guidance is of interest to anyone with responsibility for a name server at any level in the DNS hierarchy. The IAO at each facility should review this document and ensure that all DNS administrators at the site review it as well.

Standard Operating Procedures impact individuals in a variety of DNS support roles, including security management, systems administration, DNS software administration, and DNS database administration. Each of these roles is addressed in the subsections that follow.

### 7.1 Security Management Responsibilities

Security management involves the tracking of security-relevant personnel assignments, physical access control, and business continuity. The IAO will either have direct responsibilities for these areas or work closely with those that do.

#### 7.1.1 Personnel

The security of the DOD DNS infrastructure depends upon the availability and integrity of the professionals that support it.

- *(DNS0120: CAT IV) The IAO will maintain a list of personnel authorized to administer each zone and name server over which he or she has responsibility.*
- *(DNS0125: CAT II) The IAO will ensure there is at least one backup DNS database administrator identified for each supported zone and at least one backup DNS software administrator for each name server.*

#### 7.1.2 Physical Access Control

RFC 2870 states that the specific area in which the server is located must have positive access control. In other words, the number of individuals permitted access to the area must be limited, controlled, and recorded. At a minimum, control measures should include mechanical or electronic locks.

These protection mechanisms are highly recommended but may not be feasible at all sites. Name servers, however, should be among the most secured computers/assets at a location because compromise of a name server can directly impact the security of the services it supports. For example, a facility should not house web servers in a locked cage within a secured data center, but allow a name server that resolves the IP addresses for those web servers to reside on an administrator's desk.

- *(DNS0100: CAT II) The IAO will ensure a caching name server is protected by equivalent or better physical access controls than the clients it supports.*
- *(DNS0105: CAT II) The IAO will ensure an authoritative name server is protected by equivalent or better physical access controls than each of the hosts in its zone.*

### 7.1.3 Business Continuity

Business continuity involves planning for contingencies in which an entire site is lost due to a terrorist attack, fire, or severe natural disaster. In general, the strategy must be to move production-computing services to another location prepared to handle them. To accomplish this task, hardware, software, and data must be present or easily obtainable at the other location. Any such planning must include DNS if it is to be effective.

In the case of DNS, a well-designed architecture might have already placed an operating name server (e.g., a zone slave) at the disaster recovery site. If cost or other constraints preclude this, the IAO must, at a minimum, ensure that an off-site copy of zone information exists (preferably in a digital format but in a hard copy format otherwise) to prevent complete loss of records in the event of a disaster.

A key component of business continuity is continuity of electrical power. For root servers, the RFC requires power continuity of at least 48 hours, which essentially necessitates on-site electric generators. Of course, not all sites have this capability or are able procure it. Nevertheless, if all the name servers supporting a zone lose power, users may be unable to reach the hosts defined in the zone because they cannot resolve the host's name to its IP address. Therefore, name servers should have Uninterruptible Power Supply (UPS) or alternative power source similar to the hosts that they support.

## 7.2 System Administration Responsibilities

The Systems Administrator of a computer running name server software has the responsibility to ensure that security patches on the computer are current and that critical files are backed up.

### 7.2.1 Software Patches

DNS software has a history of vulnerabilities and new ones may be discovered at any time. To ensure that attackers cannot take advantage of known DNS flaws or vulnerabilities, applicable patches and patch documentation must be maintained. Guidance on appropriate patches is included in IAVM notices issued by the DOD Computer Emergency Response Team (CERT). *Appendix B* contains a list of IAVM notices relevant to DNS.

- *(DNS0130: CAT II) The SA will maintain a log of the date and time each patch or upgrade to DNS software was implemented.*



## 7.2.2 Backup

Fortunately, the DNS architecture is such that there should always be a hot backup of zone information present whenever the master name server is unavailable for any reason (i.e., the authoritative slave server maintains a copy of the zone files on the master). This built-in redundancy, however, does not extend to configuration files and logs. Therefore, name servers should be backed up to external media (e.g., tape, optical disk, etc.) on a regular basis.

At some locations, an automated enterprise backup system supports many servers. In this case, name servers can simply be added to the enterprise system. At other locations, backups must be performed manually, placing a considerably higher burden on administrators. In circumstances in which zone and configuration information is very static, remaining the same for several months at a time, it would make little sense to conduct full daily backups. Backups should occur as frequently as needed to capture changes on the name server.

- *(DNS0135: CAT II) The IAO will establish operating procedures that will ensure, at a minimum, DNS configuration and resource record data is backed up on any day on which there are changes.*

## 7.3 DNS Software Administration Responsibilities

The DNS software administrator maintains DNS configuration files and DNS-related cryptographic keys, and reviews DNS logs for security-related events.

### 7.3.1 DNS Configuration File Changes

Much of this document explains how configuration file settings affect system security. Consequently, changes to configuration files are likely to impact security and, therefore, must be tracked.

- *(DNS0140: CAT II) The IAO will maintain a log of the date and time any DNS configuration file was modified and the business justification for that modification.*

### 7.3.2 Cryptographic Key Supersession

Similar to user account passwords, cryptographic keys such as TSIG keys must be changed periodically to minimize the probability that they will be compromised. If there is a known compromise of a TSIG key, then it needs to be replaced immediately. One of the most important aspects of *key supersession* is the method that will be used to transfer newly generated keys. Possibilities, in rough order of preference, are as follows:

- SSH
- Encrypted e-mail using DOD PKI certificates
- Secure fax (STU-III)
- Regular mail (using the expedited mailing service holding the current GSA contract for "small package overnight delivery service")
- Hand courier

The operational details of key supersession are beyond the scope of this STIG.

- *(DNS0445: CAT II) The IAO will ensure a cryptographic key used to secure DNS transactions are changed at least annually.*
- *(DNS0145: CAT II) The IAO will establish written procedures for the replacement of cryptographic keys used to secure DNS transactions that will cover, at a minimum:*
  - *Frequency of key supersession*
  - *Criteria for triggering emergency supersession events*
  - *Notification of relevant personnel during emergency and non-emergency supersession*
  - *Methods for securely transferring newly generated keys*

### 7.3.3 Log Review

This STIG specifies configuration requirements for logging, but if administrators never review the logs that the configuration generates, they are of little value. Enterprises whose personnel resources or DNS transaction logs are such that frequent log review is not feasible should procure log management tools that assist administrators with identifying the most important log entries.

- *(DNS0110: CAT II) The IAO will ensure DNS log archival requirements meet or exceed the log archival requirements of the operating system on which the DNS software resides and in accordance with DODI 8500.2 requirements.*
- *(DNS0115: CAT II) The DNS software administrator will review DNS logs daily or employ a real-time log analysis or network management tool that immediately alerts an administrator of critical DNS system messages.*

## 7.4 DNS Database Administration Responsibilities

One of the simplest changes to DNS may also be one of its greatest potential vulnerabilities, specifically, adding a host name and its associated IP address to the zone file. Without a rigorous process for adding and modifying resource records, an attacker can social engineer the system (i.e., manipulate human processes rather than circumvent technical controls). For example, an attacker might be able to simply call the DNS database administrator on the phone and successfully request that the IP address of an e-mail server be changed to a rogue server under the attacker's control. Resource records can be modified to compromise security in countless ways.

To best assure the integrity of zone files, requests to change the DNS records should be carefully managed and the records should be checked periodically to ensure their validity. For example, when equipment is retired, SAs often fail to remove the associated host from the DNS. Without periodic checks, an attacker may use a retired host IP address to obtain valuable information from another user who was unaware of the change.

The details of the procedures to add or modify resource records are beyond the scope of this STIG. The requirement is that there be a written procedure in place that meets certain basic criteria.

- *(DNS0150: CAT II) The IAO will establish written procedures for the following:*
  - *The process for updating zone records*
  - *Who is authorized to submit and approve update requests*
  - *How the DNS database administrator verifies the identity of the person from whom he or she received the request*
  - *How the DNS database administrator documents any changes made*

This page is intentionally left blank.

## 8. DNSSEC

The DNS Security protocol (DNSSEC) is a promising technology for authentication of DNS transactions. It relies on public key cryptography to sign transactions, which, if implemented properly, greatly reduces the administrative burden associated with managing and transferring keys. A powerful feature of DNSSEC is the ability to sign records to ensure their integrity and authenticity throughout the DNS infrastructure and not just between the authoritative name server and its zone partner or local client. The advantages of this feature become apparent when DOD users wish to securely validate records from other organizations, including commercial vendors, business partners, and other Government agencies. For example, suppose a user wants assurance that information obtained from what it believes to be cert.mil servers are, in fact, actually cert.mil servers and not an adversary's servers masquerading as cert.mil servers. Presumably, the user or its local caching server could share a TSIG key with the name servers authoritative for the cert.mil zone, but this would be a very cumbersome approach. The cert.mil administrator would have to generate and distribute keys to a wide variety of users. Those users, in turn, would probably want keys for numerous zones, not just cert.mil. For a complete implementation, caching servers would also need TSIG keys for root servers and each top-level domain. When more than a handful of hosts and domains are involved, the key sharing agreements become so complex that they become impractical. Indeed, this problem is inherent in any symmetric key cryptosystem, of which TSIG is only one instance. With DNSSEC and its accompanying public key infrastructure (PKI), organizations only need to trust a common signing authority. In the above example, if cert.mil has its zone signed by an authority that the osd.mil name server also trusts, then the desired assurance is achieved.

The objective is to establish signing authorities broadly recognized throughout DOD and elsewhere. Although creating the DNSSEC key infrastructure is a challenge, it is much more feasible than acquiring equivalent capabilities with TSIG. There are no requirements for DNSSEC at this time because the technology is not mature enough to prevent or delay Certification and Accreditation (C&A) of systems that have not yet implemented it. Once DNSSEC has been formally tested in the environment, personnel have been trained, an appropriate public key infrastructure is available, and operating procedures are in place, DNSSEC is likely to be the preferred technology for securing DNS. Organizations can implement DNSSEC today if they choose to do so, but are advised to perform the necessary planning, training and infrastructure upgrades prior to such an implementation.

This section describes the mechanisms involved in the DNSSEC approach, the operations those mechanisms involve, and a secure way of performing those operations by using checklists. In other words, this section provides guidelines for secure deployment of DNSSEC features through DNSSEC operations that are supported in name server software. These guidelines are covered in Sections 8.2 through 8.6. To ensure end-to-end protection of DNS query/response transactions, additional protection measures (apart from what DNSSEC specification provides)—such as securing the communication path between local DNSSEC-aware caching/resolving name servers and stub resolvers—are needed. These measures are discussed in Section 8.7. Zone administrators need to have an understanding of the logic of dynamic updates in the presence of additional RRsets (especially NSEC RRset) introduced into the zone file by DNSSEC specification. This logic is explained in Section 8.8.

## 8.1 DNSSEC Mechanisms and Operations

DNSSEC mechanisms involve two main processes: sign and serve, and verify signature. These processes are described below.

### 8.1.1 Sign and Serve

The first task in this process is to generate digital signatures associated with every RR in the zone file. Instead of generating a signature for every RR, DNSSEC specifies generation of a signature for an RRset (a set of RRs with the same owner name, class, and RRTYPE). The digital signature and its associated information (ID of the key used, start and expiry dates for the signature, etc.) are encapsulated in a special RR of RRTYPE RRSIG (Resource Record Signature). The actual key string and associated information about the public key that is to be used to verify the signature (in RRSIG) are given in a special RRTYPE DNSKEY. Another RRTYPE, NSEC (Next Secure), is used to list RRTYPES (in canonical order) available for a given domain (owner name), and a signature (RRSIG RR) for that RRTYPE is generated to provide authenticated proof of nonexistence to queries for any nonexisting RRTYPE in that domain. In addition, there is an optional RRTYPE DS (Delegation Signer) in case a zone wants to vouchsafe the authenticity of the public key of its child zone. In other words, the DS RR carries the signature for the RR that contains the (hash) public key of the child zone. The detailed syntax for each of these additional RRTYPES introduced by DNSSEC specification is specified in RFC 4034. The most important of these is the RRSIG RR because it contains the actual signature string. The RRSIG RR, like any other RR, contains the owner name, TTL, class, RRTYPE, and RDATA fields. The digital signature and all its associated information are contained within the RDATA field. The layout of the RDATA field in the RRSIG RR, with all subfields, is shown in Figure 8-1. A brief description of each subfield follows.

The “RRTYPE Covered” field is the type of RRset for which this RRSIG holds the signature. The “Algorithm Code” field is a code integer assigned to represent a given cryptographic algorithm used to generate the signature. The “Labels” and “Original TTL” fields are the number of labels (actual number of RRs in the signature) and the TTL value of the RRset this signature covers. The “Signature Expiration” and “Signature Inception” values are absolute time values that span the signature validity period—the time period for which this RRSIG is considered valid for the zone. The “Key Tag” and “Signer’s Name” fields are the hash and FQDN of the DNSKEY RR that the client needs to validate the signature, and the final field is the encoded signature itself. A zone that contains the additional RRs along with the regular RRs is called a signed zone. A name server that hosts these signed zones and includes the appropriate signatures (i.e., the corresponding RRSIGs) along with requested RRs in its response is called a DNSSEC-aware name server.

### 8.1.2 Verify Signature

The response coming from a signed zone is called a signed response. A resolver that has the capability to verify signatures in a signed response is called a DNSSEC-aware validating resolver. Before a resolver can verify the signature associated with RRsets (in the response) of a zone using the public key of the zone (sent along with the response), it has to establish trust in that public key. In DNSSEC this requirement is addressed by having the resolver go through a

sub-process called building a trusted chain. In this sub-process, the resolver starts from a known list of trusted public keys (called trust anchors) and establishes trust in the public key of the zone in a given response by traversing through a chain of public keys, building the chain by using the DNS name space hierarchy. The trust anchors in a resolver ideally consist of the root public keys (if root servers are DNSSEC-aware) or public keys of zones lower in the hierarchy. The trust anchor list in a resolver is not built through a DNS transaction; it uses an out-of-band mechanism. The DNSSEC processes described above involve several name server operations and a few resolver operations.

The name server operations are as follows:

- DNSSEC-OP1: Generation of public key-private key pair
- DNSSEC-OP2: Secure storage of private keys
- DNSSEC-OP3: Public key distribution
- DNSSEC-OP4: Zone signing
- DNSSEC-OP5: Key rollover (changing of keys)
- DNSSEC-OP6: Zone re-signing.

The resolver operations are DNSSEC-OP7: Trust anchors' configuration; and DNSSEC-OP8: Establishing trusted chain and signature verification.

Name server operations DNSSEC-OP1 through DNSSEC-OP4 and resolver operations DNSSEC-OP7 and DNSSEC-OP8 are performed either prior to DNSSEC deployment or for secure operation (serving signed responses and verification of signatures) and hence are covered in this section. The remaining name server operations (key rollovers and zone re-signing – DNSSEC-OP5 and DNSSEC-OP6 respectively) are performed on a periodic basis after a fully operational DNSSEC deployment.

## **8.2 Generation of Public Key-Private Key Pair (DNSSEC-OP1)**

DNSSEC specifies generation and verification of digital signatures using asymmetric keys. This requires generation of a public key-private key pair. Although the DNSSEC specification does not call for different keys (just one key pair), experience from pilot implementations suggests that for easier periodical security administration operations such as key rollover (changing of keys) and zone re-signing, at least two different types of keys are needed. One set is called Key Signing Key (KSK). This key (specifically, the private part of the key pair, called KSK-private) will be used only for signing the key set (i.e., DNSKEY RRset) in the zone file. The other key type is called the Zone Signing Key (ZSK) (whose private part is called ZSK-private) and will be used to sign all RRsets in the zone (including DNSKEY RRset). An administrative distinction is made between the KSK and ZSK keys by setting the Secure Entry Point (SEP) flag bit in the DNSKEY RR that represents the public part of those keys (called KSKpublic and ZSK-public, respectively). The logic behind creation of two types of key pairs is to provide separate set of functions for each key type and thus reduce the overall complexity of tasks involved in key rollovers and zone re-signing. Accordingly, the KSK (KSK-private) is used to sign the key set (i.e., DNSKEY RRset) and is the key type that is sent to the parent to be used for authenticated delegation (by generating a DS RR, using the hash of the child's KSK-public key and generating

a corresponding a signature (RRSIG RR) using its own ZSK-public). The KSK (KSK-public) is also meant to be used as a trust anchor (SEP keys) in validating resolvers to establish trust chains for verification of signatures. The ZSK (ZSK-private) is to be used for signing the entire zone file (all RRsets). The public portion of this key (ZSK-public) will not be sent to the parent and will always remain in the zone. The decision parameters involved in KSK and ZSK key pair generation are as follows: Choice of digital signature algorithm, choice of key sizes, and choice of crypto period (duration for which the key will be used).

The choice of digital signature algorithm will be based on recommended algorithms in well-known standards. NIST's Digital Signature Standard (DSS) (FIPS 186-2) provides three algorithm choices: Digital Signature Algorithm (DSA), RSA, and Elliptic Curve DSA (ECDSA).

Of these three algorithms, RSA and DSA are more widely available and hence are considered candidates of choice for DNSSEC. In terms of performance, both RSA and DSA have comparable signature generation speeds, but DSA is much slower for signature verification. Hence, RSA is the recommended algorithm as far as this guideline is concerned. RSA with SHA-1 is the only cryptographic algorithm required to be implemented with DNSSEC. It can be expected that name servers and clients will be able to use the RSA algorithm at the minimum. It is suggested that at least one ZSK for a zone use the RSA algorithm.

NIST's Secure Hash Standard (SHS) (FIPS 180-2) specifies SHA-1, SHA-224, SHA-256, SHA-384, and SHA-512 as approved hash algorithms to be used as part of the algorithm suite for generating digital signatures using the digital signature algorithms in the NIST's DSS. Considering ease of implementation and availability, SHA-1 is recommended as a good candidate, in spite of a recent reported hash collision attack that took away about 11 bits from its theoretical bit strength.

The choice of key size is a tradeoff between the risk of key compromise and performance. The performance variables are signature generation and verification times. The size of the DNS response packet also is a factor because DNSKEY RRs are sent in the additional section of the DNS response. Because the KSK is used only for signing the key set (DNSKEY RRset), performance is not much of an issue. Compromise of a KSK could have a great impact, however, because the KSK is the entry point key for a zone. Compromise of a KSK in a zone high in the DNS hierarchy could expose a great portion of the DNS subtree (hence a large number of zones) to spoofing attacks. In addition, rollover of a KSK in the event of a compromise involves potential update of trust anchors in many validating resolvers. Hence, a large key size is recommended for the KSK: it has negligible performance impact but a large impact relating to key compromise.

As far as the choice of key size for the ZSK is concerned, performance certainly will be a factor because the ZSK is used for signing all RRsets in the zone. In terms of impact, however, it is restricted to just a single zone because the ZSK's usage is limited to signing RRsets only for that zone but not for providing authenticated delegation for a child zone. Hence, a key size smaller than that for the KSK can be used for the ZSK.



The choice of crypto period (rollover period) is dictated by the risk of key exposure. In the case of the KSK, the volume of signature sets generated is not large (because the KSK signs only the DNSKEY RRset and the frequency of change of this RRset also is small). The minimal key exposure (small volume of data available to guess KSK-private) combined with the large size of the key implies that the crypto period for the KSK can be long (usually a year or two).

In the case of ZSK, the risk of key guessing is higher because of larger key exposure. The larger key exposure is a result of the fact that the number of signature sets generated is very large (because the ZSK signs all RRsets in a zone and other RRsets change much more frequently than DNSKEY RRsets, so the number of fresh signatures generated is high). This factor, combined with the relatively smaller size of the key, implies that the crypto period for ZSKs must be much smaller than those for KSKs (usually a month or two).

The key size for the KSK should be sufficiently large because of the greater impact on DNS due to KSK key compromise. The crypto period (rollover period) for the ZSK should be sufficiently short because of the greater risk of key guessing resulting from larger ZSK key exposure.

In terms of the number of keys of each type (KSK and ZSK) to be generated, a good practice is to generate an extra ZSK in addition to the one that will be used for signing. Hence, the zone administrator should use the key generation program to generate one KSK and two ZSKs during initial deployment of DNSSEC. One ZSK is treated as the active key, and its private part (ZSK-private) will be used for signature generation. The other ZSK (ZSK-public) will be made part of DNSKEY RRset, but its associated private part (ZSK-private) will not be used for signing RRsets. This additional ZSK will provide a readily available ZSK for immediate rollover in emergency situations such as key compromise and a form of advance notification to validating resolvers that this key is to be the one into which the zone is going to roll over after the current crypto period expires. The mere presence of the key in the DNSKEY RRset enables validating resolvers to cache and establish trust in the new key so that they can immediately use the key for signature verification as soon as rollover occurs.

The recommended digital signature algorithm suite, key sizes, and crypto periods for the KSK and ZSK keys are given in Table 8-1.

### 8.2.1 Key Pair Generation—Illustrative Example

Every DNSSEC-aware name server software should provide a utility program for generating asymmetric key pairs (a public-private key pair). The use of one such program, `dnssec-keygen` (provided by BIND 8.3.x), is illustrated below: `dnssec-keygen -a algorithm -b bits -n type [options] name` where algorithm (under `-a` parameter) can be one of the following:

RSASHA1  
DSA

bits (key size) for the `-b` parameter has the following ranges:

[512..4086] for RSASHA1

[512..1024] for DSA (must be divisible by 64)

type for the `-n` parameter can be one of `ZONE` | `HOST`

name is the owner of the key (usually the domain name in the zone apex).

This command generates two files—one containing the public key and the other the associated private key. The generic names for these files are as follows:

`K<domain_name>+algorithm_id+Key_id.key` `K<domain_name>+algorithm_id+Key_id.private`

The `domain_name` is the value of the name parameter specified on the command line. The algorithm IDs are:

- 3 – DSA
- 5 – RSASHA1

The `key_id` is the unique identifier for the key generated by the program. For example, to generate a ZSK keyset of length 1024 bits that uses the RSASHA1 algorithm suite for signing the zone `example.com`, the following command would be issued:

```
dnssec-keygen -a RSASHA1 -b 1024 -n ZONE example.com
```

For this command, the following files containing the private and public keys, respectively, are generated:

```
Kexample.com.+005+28345.private  
Kexample.com.+005+28345.key
```

In these file names, 005 stands for the `algorithm_id` and 28345 is the unique key ID. In the `*.key` file, the public key information is expressed in the same syntax as that of a zone file RR. The content of the file `Kexample.com.+005+28345.key` will be: `example.com IN DNSKEY 256 3 5 BQFG+KGJ7.....` (Base64 encoded key string) Hence, the contents of the file containing the public key can be readily added to the zone file (`zonedb.example.com`) contents by using the following command: `cat *.key >> /var/named/zonedb.example.com`. After the DNSKEY RR containing the public key is added to the zone file, the zone's serial number in the SOA RR must be increased prior to actually signing the zone.

Windows Server 2003 DNS provides basic support of the DNS Security Extensions (DNSSEC) protocol as defined in RFC 2535. The current feature support allows DNS servers to perform as secondary DNS servers for existing DNSSEC-compliant, secure zones. DNS supports the storing and loading of the DNSSEC-specific resource records (RRs). Currently, a DNS server is not capable of signing zones and resource records (creating cryptographic digital signatures) or validating the SIG RRs. The DNSSEC resource records are KEY, SIG, and NXT.

To obtain additional information and suggested requirements for the implementation of DNSSEC, refer to the NIST Special Publication 800-81 (Draft), Secure Domain Name System (DNS) Deployment Guide.

This page is intentionally left blank.

## **APPENDIX A. RELATED PUBLICATIONS**

### **Publications:**

Comer, Douglas. *Internetworking with TCP/IP*. Prentice Hall, 4<sup>th</sup> Edition, 2000.

Albitz, Paul and Liu, Cricket. *DNS and BIND*. O'Reilly & Associates, Inc., 4<sup>th</sup> Edition, 2001.

DISA API. Department of Defense Implementation Guidelines for Domain Name System (DNS) Security (DRAFT), 7 September 2000.

DISA API. Action Plan for Department of Defense (DOD) Domain Name System (DNS) Hardening (DRAFT), 16 August 2002.

DISA API. Discussion of Advanced Domain Name System (DNS) Features, Windows 2000 Active Directory, and Their Impact Within DOD (DRAFT), 26 August 2002.

ISC, BIND 9 Administration Reference Manual, 2001.

NIST Special Publication 800-81 (Draft) Secure Domain Name System (DNS) Deployment Guide, August 2005. Authors: Ramaswamy Chandramouli and Scott Rose.

The Center for Internet Security DNS BIND Benchmark, Version 1.0, January 2006.

### **Request For Comment (RFC):**

RFCs are currently managed by the Internet Engineering Task Force (IETF) and can be found at <http://www.rfc-editor.org>.

### **Recommended RFCs:**

RFC 1034, Domain Names – Concepts and Facilities, November 1987

RFC 1035, Domain Names – Implementation and Specification, November 1987

RFC 2181, Clarifications to the DNS Specification, July 1997

RFC 2182, Selection and Operation of Secondary DNS Servers, July 1997

RFC 2535, Domain Name System Security Extensions, March 1999

RFC 2870, Root Name Server Operational Requirements, June 2000

## **Applicable RFCs:**

1713 Tools for DNS Debugging  
1912 Common DNS Operational and Configuration Errors  
2536 DSA KEYs and SIGs in the Domain Name System (DNS)  
2537 RSA/MD5 KEYs and SIGs in the Domain Name System (DNS)  
2538 Storing Certificates in the Domain Name System (DNS)  
2539 Storage of Diffie-Hellman Keys in the Domain Name System (DNS)  
2540 Detached Domain Name System Information  
2541 DNS Security Operational Considerations  
2845 Secret Key Transaction Authentication for DNS (TSIG)  
2930 Secret Key Establishment for DNS (TKEY RR)  
2931 DNS Request and Transaction Signatures (TSIG)  
3007 Secure Domain Name System (DNS) Dynamic Update  
3008 Domain Name System Security Signing Authority

## **Web Sites:**

<http://www.isc.org/>

Internet Software Consortium

<http://www.nominum.com/>

Nominum, a for-profit company, is now performing the principal maintenance and development of BIND on behalf of the ISC. Hosts the BIND 9 Administration Reference Manual.

(<http://www.nominum.com/content/documents/bind9arm.pdf>)

<http://www.dns.net/dnsrd/>

The DNS Resource Directory

<http://www.nwfusion.com/research/dns.html>

Network World Fusion DNS Resources

<http://www.cist-east.saic.com/>

SAIC - Center for Information Security. This site is specifically maintained in support of the DOD DNS Security effort.

<Http://www.microsoft.com/>

Microsoft Corporation.

[www.dnssec.net](http://www.dnssec.net)

Information on DNSSEC

## **APPENDIX B. INFORMATION ASSURANCE VULNERABILITY MANAGEMENT (IAVM)**

The IAVM notices listed below relate to DNS. Administrators should read these notices and comply with the guidance therein.

2000-B-0001: BIND NXT Buffer Overflow

<ftp://www.cert.mil/pub/bulletins/dodcert2000/2000-b-0001.htm>

2000-B-0008: BIND 8.2.2-P6 Denial of Service Vulnerabilities

<ftp://www.cert.mil/pub/bulletins/dodcert2000/2000-b-0008.htm>

2001-A-0001: Multiple Vulnerabilities in BIND

<ftp://www.cert.mil/pub/bulletins/dodcert2001/2001-a-0001.htm>

2002-A-0006 (v1): Multiple Vulnerabilities in ISC BIND 4 and 8

<ftp://www.cert.mil/pub/bulletins/dodcert2002/2002-a-0006.htm>

2002-T-0010: Denial of Service Vulnerability in ISC-BIND 9

<ftp://www.cert.mil/pub/bulletins/dodcert2002/2002-t-0010.htm>

2003-B-0001: Multiple Buffer Overflow Vulnerabilities in Various DNS Resolver Libraries

<ftp://ftp.cert.mil/pub/bulletins/dodcert2003/2003-b-0001.htm>

2005-A-0005: Multiple Vulnerabilities in BIND

<ftp://www.cert.mil/pub/bulletins/dodcert2005/2005-a-0005.htm>

This page is intentionally left blank.



## APPENDIX C. GENERATION AND STORAGE OF TSIG KEYS

TSIG Keys are required to authenticate zone transfers and updates. This appendix provides instructions on how to create and install a key in a manner that meets the requirements of this STIG. The steps are as follows:

### 1. Generate the key

From a command prompt on one of the name servers that will utilize the key, change to a directory where the user would like the keys to be created. Ideally, the keys should be located in a designated area of the file system for that purpose (e.g., */etc/dns/keys/*). The command and its results will look similar to the following:

```
% dnssec-keygen -a HMAC-MD5 -b 128 -n HOST example_key.  
Kexample_key.+157+38675
```

Use the HMAC-MD5 algorithm only when the SHA-1 algorithm is not available. BIND 9.2.1 only supports MD5 but future releases should support SHA-1.

This will create two new files in the current directory:

```
Kexample_key.+157+38675.key  
Kexample_key.+157+38675.private
```

Note that the string that is returned by **dnssec-keygen** is the base name for the two files. The first is followed by “.key” and the second by “.private.” The string itself is the name chosen for the key preceded by a “K” and followed by a numerical string related to the key.

An example of the text in the *.key* file is:

```
example_key. IN KEY 512 3 157 2njlQNnzn6HTwKLcjStUXg==
```

An example of the text in the *.private* file is:

```
Private-key-format: v1.2  
Algorithm: 157 (HMAC_MD5)  
Key: 2njlQNnzn6HTwKLcjStUXg==
```

In this example, the generated key is:

```
2njlQNnzn6HTwKLcjStUXg==
```

Since both files contain the secret key, they should both be protected accordingly. Anyone who obtains one of these files could use the key to perform unauthorized changes to DNS zone records.

## **2. Create a new designated file for that key**

Using a text editor, create a file with the following content:

```
secret "generated_key" ;
```

In our example, the contents would be:

```
secret "2nj1QNzn6HTwKLCjStUXg==" ;
```

The syntax of the statement is critical. Ensure that:

- The word "secret" appears at the beginning of the line followed by a space
- The key is included in quotes with no extra spaces before or after the key
- A semi-colon (;) follows the quotation mark after the key
- There are no extra characters, lines, or carriage returns before or after the statement

Importantly, any key longer than approximately 320 bits will contain a space within the key field of the original .key file. This space can be left within the string, as long as it is enclosed within double quotes (") in the new file created to house the key.

## **3. Transfer the key to the other name server that will utilize the key**

Two name servers will need to use the same key to authenticate transactions between them. Therefore, the generated key must be transferred to the other system. If this transfer occurs over a network, then it must be encrypted so that no other party can view the key. If any paper copy of the key is created, then the paper must be shredded after the key is entered on the other system.

## **4. Set appropriate permissions on key files**

As specified in *Appendix E.3.2, TSIG Keys*, TSIG key files should have permissions of 400 and be owned by the user account used to run the name server software.

## 5. Define the key within the *named.conf* file on each relevant name server

For a key file named *tsig-example.key* that resides in */etc/dns/keys/*, an appropriate *key* statement would be:

```
key example_key {  
    algorithm hmac-md5;  
    include "/etc/dns/keys/tsig-example.key";  
};
```

When *named* is started, it will interpret the include statement and build the following key definition (assuming the generated key in our example):

```
key example_key {  
    algorithm hmac-md5;  
    secret "2nj1QNzn6HTwKLcjStUXg==" ;  
};
```

This page is intentionally left blank.

## **APPENDIX D. MODEL STANDARD OPERATING PROCEDURES FOR DNS DATABASE ADMINISTRATION**

The model procedures are not mandatory and designed to provide organizations with suggested language for local written procedures. Nevertheless, adoption of the model procedures will ensure that the organization is fully compliant with the administrative requirements related to resource record updates found in *Section 7, Standard Operating Procedures for DNS*, of this STIG. In review, the requirement is that the IAO will establish written procedures regarding the following:

- The process for updating zone records
- Who is authorized to submit and approve update requests
- How the DNS database administrator verifies the identity of the person from whom he or she received the request
- How the DNS database administrator documents any changes made

A key objective of the procedures is separation of duties, which limits the potential for a single insider to misuse the system. There are three key components of any changes to a DNS zone:

- Request
- Approval
- Implementation

Ideally, a different person should perform each of these activities. For example, an application team deploying a new database server would request a host name and IP address for that server. A network security officer or operations manager would approve that request. Finally, a DNS database administrator would implement the approved request.

In some environments, two of the three components may be combined. For example, designated personnel may be authorized to submit requests that require no further approval. In smaller environments, technical staff may perform multiple roles (e.g., IAO and DNS database administrator), in which case the approval and implementation steps may be combined in practice. Combining all three components is ill advised in all but small LAN environments in which one or two IT support professionals perform all IT functions, which makes separation of duties impossible or impractical.

The model procedures are listed below. Optional items in the procedures are placed in brackets ([...]). Comments or guidance within the model procedures are in italics.

1. The authority to request changes to a zone database (e.g., create, modify or delete a record) is limited to (select one of the options below)  
[an IAO-maintained list of individuals stored at \_\_\_\_\_ location]  
[the following organizational roles: \_\_\_\_\_]

2. The authority to approve change requests is limited to (select one of the options below)  
[the IAO]  
[an IAO-maintained list of individuals stored at \_\_\_\_\_ location]  
[the owner of the domain as specified in the zone database]  
[the DNS database administrator]
3. The authority to implement approved change requests is limited to an IAO-maintained list of DNS database administrators stored at \_\_\_\_\_ location.
4. The approved method(s) for submitting a requests is/are as follows: (select one or more of the options below)  
[send a [digitally signed] e-mail message to the designated approver]  
[send a signed written request to the designated approver]  
[open a ticket with the help desk]
5. Before implementing approved change requests, the DNS database administrator verifies the identity of the individual from whom he or she receives the request: (select one or more of the options below):  
[through voice or face recognition, when the requestor or approver is well known]  
[using the "sent from" field when requests are delivered via e-mail]  
[validating the digital signature in the e-mail client software]  
[through call-back to the requestors desk telephone number]  
[through a comparison of the approving signature on the request to one on file for that individual].
6. The DNS database administrator saves approved requests to \_\_\_\_\_ location (e.g., network file folder, e-mail archive, file cabinet, etc.)
7. When implementing approved change requests, the DNS database administrator documents the following information related the request:
  - the name[, title][, and organizational affiliation] of the requestor
  - the name[, title][, and organizational affiliation] of the approver
  - the name of the DNS database administration implementing the change
  - the date of the change
  - [- who is permitted to make subsequent changes to the affected records]
8. Records of updates and changes are located in (select one of the options below)  
[zone files]  
[a database named \_\_\_\_\_ located at \_\_\_\_\_]

This page is intentionally left blank.

## APPENDIX E. UNIX OPERATING SYSTEM CONFIGURATION TO SUPPORT BIND

The UNIX STIG provides guidance on how to secure a UNIX computer so that it may operate at an acceptable level of risk. That STIG also includes guidance specific to BIND, particularly with respect to file permissions on key BIND files and patches and upgrades needed to comply with DNS-related IAVM notices. This appendix provides additional guidance on OS-level concerns related to BIND. Eventually, the issues will be addressed in the UNIX STIG and, when that occurs, this appendix will be removed.

For general operating system requirements and guidance, refer to the section on *Name Server Operating System Platforms*.

For specifics on configuring BIND to run on various OSs such as Red Hat Linux and Solaris, refer to the Center for Internet Security (CIS) DNS BIND Benchmark.

### E.1 Network Services

Only network services required to ensure the availability, confidentiality, and integrity of DNS information must run on a name server in a DOD DNS infrastructure. Running unnecessary services leaves BIND vulnerable to potential exploits related to software not critical to DNS. Therefore, a maximum-security posture prohibits these services.

Listing exactly what services should be available in this STIG is not practical because different sites deploy different hardware, intrusion detection agents, network management tools, and other items required to maintain operations. Nevertheless, the IAO must adhere to the spirit of service minimization to the extent operationally feasible. Clearly, BIND should not run on a server that also hosts a web application, handles e-mail, routes network traffic, or houses a non-DNS database. Both */etc/services* and */etc/inetd.conf* should be configured to reflect this.

Some SAs like to install BIND on application servers to quicken lookups. For example, an e-mail server that forwards messages to e-mail servers in other domains must issue queries for the MX and A records associated with those other mail servers. If these could be obtained from a locally installed name server, the presence of that name server would speed up query response and reduce network traffic. Nevertheless, the security risks of this configuration outweigh the speed advantages. In the case of e-mail, the few extra milliseconds it takes to query remote name servers will have no appreciable impact on end-users' e-mail experience, while exposing both e-mail and DNS to each other's security vulnerabilities could have a significant impact on an organizations computing infrastructure.



Beyond *named* and *sshd*, there are a few core services that are still required. For example, name servers must comply with all applicable host intrusion detection and system integrity requirements. Consequently, it is likely that products such as Tripwire, and Real Secure Server Sensor are resident on the system. In addition, there may be agent software to support network monitoring or management and time synchronization services. If an organization believes it can best ensure 24x7 DNS availability through integration of the name server into site-wide monitoring services, then it is appropriate that software run on the name server to support that service. Similar cases can be made for backup and uninterruptible power supply (UPS) if the site supports these services across the infrastructure with agent software. If a network service such as SNMP is utilized, it must be secured in accordance with the Unix and Network Infrastructure STIGs.

The only permitted services to be running on a DNS UNIX BIND server are those implementing:

- DNS
  - Secure shell
  - Host intrusion detection
  - Host file integrity
  - Network management or monitoring
  - Anti-virus
  - Backup
  - UPS
- *(DNS4450: CAT II) The SA will ensure all network services not required for operations are disabled. Any network service that is utilized is documented with the IAO with complete mission justification and vulnerability mitigation.*

## **E.2 Non-privileged User for BIND**

The BIND *named* daemon runs as root by default, but this practice represents a security vulnerability because an adversary that managed to gain control of *named* could then control the entire system. To prevent this from happening, BIND can run as a non-privileged user with a UID other than 0 (which belongs to root). The SA, however, must create such an account to implement this approach. For the remainder of this document, it will be assumed that the name of the user is *dnsuser* and that it is the sole member of a group named *dnsgroup*.

Unfortunately, the existence of a new non-privileged user introduces a point of entry for a potential attacker. If the attacker can log on as *dnsuser*, then the attacker would have full control of DNS services. To provide assurance that *dnsuser* can only be used to run BIND, the SA can, at a minimum, define its shell as */bin/false* (i.e., no shell at all). This prevents an adversary from using the non-privileged user account to execute commands.

- (DNS4440: CAT III) The SA will create a special non-privileged DNS user account and configure BIND to run as that user account. BIND will not be configured to run as a-root user.
- (DNS4460: CAT III) The SA will ensure no user can obtain a command shell using the DNS user account (e.g., by setting its shell to /bin/false in /etc/passwd or a similar alternative).

### E.3 File Permissions

OS STIGs provide guidance on how SAs can secure a file system, all of which is applicable to name servers. In addition, there are additional files specific to BIND that the SA must also lock down. These include BIND support files and TSIG keys.

#### E.3.1 BIND Support Files

BIND support files include the configuration files, zone databases, and logs. There is no reason why anyone other than the *named* process and the DNS or System Administrator needs access to these files. The *named* process will need to write to log and zone files, but only the DNS software administrator needs the ability to modify configuration files. The *named* process will not have write access to any configuration files, key files, or any directories in which the files are contained.

- (DNS4470: CAT II) The SA will set ownership and permissions on all BIND name servers at least as restrictive as listed in the table below:

<b>FILE NAME</b>	<b>OWNER</b>	<b>GROUP</b>	<b>PERMISSIONS</b>
<i>named.conf</i>	<i>root</i>	<i>dnsgroup</i>	<i>640</i>
<i>named.pid</i>	<i>root</i>	<i>dnsgroup</i>	<i>600</i>
<i>Root hints file</i>	<i>root</i>	<i>dnsgroup</i>	<i>640</i>
<i>master zone file</i>	<i>root</i>	<i>dnsgroup</i>	<i>640</i>
<i>slave zone file</i>	<i>root</i>	<i>dnsgroup</i>	<i>660</i>

Common names for the root hints file are *root.hints*, *named.cache*, or *db.cache*. The name is configurable within the *named.conf* file.

- (DNS4470: CAT III) On BIND 8 name servers, the SA will set ownership and permissions on all BIND name servers at least as restrictive as listed in the table below (if the listed files exist):

<b>FILE NAME</b>	<b>OWNER</b>	<b>GROUP</b>	<b>PERMISSIONS</b>
<i>named.run</i>	<i>root</i>	<i>dnsgroup</i>	<i>660</i>
<i>named_dump.db</i>	<i>root</i>	<i>dnsgroup</i>	<i>660</i>
<i>ndc (FIFO)</i>	<i>root</i>	<i>dnsgroup</i>	<i>600</i>
<i>ndc.d (directory containing ndc)</i>	<i>root</i>	<i>dnsgroup</i>	<i>700</i>

Note that *dnsgroup* (or an alternative name) must be created before its use.

### E.3.2 TSIG Keys

Required access to TSIG keys is extremely limited, nothing more than allowing the *named* process the ability to read the keys to sign transactions. Even the DNS and System Administrators should not be able to read or modify the keys. The exception is supersession events – when the keys are replaced. In these cases, the SA should temporarily change file ownership to root and assign write permissions for the duration of the switch only.

<b>FILE NAME</b>	<b>OWNER</b>	<b>GROUP</b>	<b>PERMISSIONS</b>
<i>unique to each key</i>	<i>dnsuser</i>	<i>dnsgroup</i>	<i>400</i>

For detailed requirements and guidance for encryption key file permissions, refer to *Section 3.2, Name Server Operating System Platforms*.

### E.3.3 Log Files

Log files should be appropriately secured because they could possibly reveal information about DNS configuration that potentially could assist an adversary attempting to breach the security of DNS. Nevertheless, allowing authenticated users to read the files is still acceptable. If a name server is configured to write DNS messages to the system logging facility, then permissions on these files should follow the relevant OS STIG guidance for this access. If BIND is configured to log events to other files, these logs should be appropriately secured, having file permissions that restrict unauthorized changes.

<b>FILE NAME</b>	<b>OWNER</b>	<b>GROUP</b>	<b>PERMISSIONS</b>
<i>any log file</i>	<i>dnsuser</i>	<i>dnsgroup</i>	<i>660</i>

For detailed requirements guidance for log file permissions, refer to the section on *Logging*.

## E.4 Running BIND in chroot(ed) Directory Structure

With any network service, there is the potential that an attacker can exploit a vulnerability within the program that allows the attacker to gain control of the process and even run system commands with that control. One possible defense against this attack is to limit the software to particular quarantined areas of the file system, memory or both. This effectively restricts the service so that it will not have access to the full file system. If such a defense were in place, then even if an attacker gained control of the process, the attacker would be unable to reach other commands or files on the system. This approach often is referred to as a *padded cell*, *jail*, or *sandbox*. All of these terms allude to the fact that the software is contained in an area where it cannot harm either itself or others. A more technical term is a *chroot(ed)* directory structure.

It is highly recommended that the SA configure BIND to run in a padded cell or chroot(ed) directory structure.

- *(DNS4445: CAT III) The SA will ensure BIND is configured in a chroot (ed) directory structure.*

This page is intentionally left blank.

## APPENDIX F. MICROSOFT WINDOWS OPERATING SYSTEM CONFIGURATION TO SUPPORT BIND

BIND was written to run on UNIX and other operating systems prior to Windows, but the current source code has been modified to also run on Microsoft Windows NT and Windows 2000. UNIX is the preferred platform for BIND because the behavior of its secure configuration with BIND is better known than the one for Windows. Nevertheless, some sites may select Windows because their staff may be much more familiar with Windows than UNIX or because it may be more cost-effective for them to deploy it.

If a site chooses to install BIND on Windows, then the DNS software administrator should still carefully read the appendix on UNIX configuration that precedes this one. Windows configuration settings related to BIND are derived from corresponding UNIX settings. Thus, to fully understand the motivation behind Windows settings and how they may differ from those in UNIX, the DNS software administrator must have a working knowledge of the secure UNIX configuration.

Eventually, the material in this appendix will be transferred to the Windows STIG. When that transition is complete, this appendix will be removed. For general DNS requirements and guidance related to operating systems, refer to *Section 3.2, Name Server Operating System Platforms*.

### F.1 Network Services

The general principle behind acceptable services on Windows is the same as that for UNIX. Only network services absolutely required, to ensure the availability, confidentiality, and integrity of DNS information must run on a name server in the DOD DNS infrastructure.

The implementation on Windows, however, is substantially different than the implementation on UNIX. Startup settings for UNIX daemons are defined in *inetd.conf*. In Windows, startup settings for server software are defined in the services applet of the Control Panel or Microsoft Management Console. A UNIX SA can prevent a daemon from running at startup by removing the daemon from *inetd.conf*. A Windows SA can prevent a service from running at startup by clicking the "disable" options under the configuration dialog box for that service. Using the disable option is required for unnecessary services because it prevents unauthorized or accidental use of those services at a later time.

In UNIX, secure remote administration is typically achieved with secure shell (SSH) implemented using the *sshd* daemon. Although there are implementations of SSH for Windows, they are not commonly used and usually require separate licenses. Perhaps the closest equivalent to SSH in Windows is Terminal Services, which provides users with a remote graphical user interface (GUI) as if they were on the machine itself, but without encrypted sessions. Use of Terminal Services is not recommended for remote management of name servers because of potential security vulnerabilities associated with this service.

On DOD local area networks (LANs), a DNS software administrator can remotely administer BIND on Windows via use of standard Windows remote administration tools. The use of the **net start** and **net stop** commands can be used to start and stop the *ISC BIND* service. Text editors such as notepad can be utilized to access remote configuration files. With these tools, a DNS software administrator will be able to perform all required tasks. Relative to SSH, the disadvantage of this type of in-band management is that most communication is not encrypted. Nevertheless, the challenge-response method of authentication at least guarantees that passwords are not sent in clear text.

Remote administration of this sort must be limited to the LAN. If there is a requirement to remotely administer a Windows-based BIND name server from a host on an untrusted network or on the external side of an organization's firewall, then third-party network encryption tools must be utilized to ensure the confidentiality and integrity of all transferred data.

Use of the Windows net commands and remote file access requires that the Server service be running on the name server. To make sure that the Server service is not utilized unnecessarily, there must be no file shares beyond the default administrative shares such as C\$.

The only permitted services to be running on a Windows ISC BIND DNS server are those implementing:

- DNS (*i.e., the ISC BIND service*)
  - Host intrusion detection
  - Host file integrity
  - Network management or monitoring
  - Anti-virus
  - Backup
  - UPS
- *(DNS4570: CAT II) The IAO will ensure that appropriate encryption software is correctly installed and configured on Windows BIND name servers if it is required that in-band remote management be performed from hosts outside the enclave in which the name server resides.*
  - *(DNS4580: CAT II) The SA will ensure that no other shares other than the default administrative shares are enabled on a name server.*

## **F.2 Non-Privileged User for BIND**

As in UNIX, the Windows SA can create a special user account dedicated to running BIND. This ensures that if an external entity somehow gains control of BIND, its access to other system resources is restricted to the greatest extent feasible. If this technique were implemented, then the ISC BIND service dialog box would appear as follows:



**Figure F-1. Service Account**

In Windows, a user does not have to be a member of any group other than the implicit groups "Everyone" and "Authenticated Users." Thus, to best ensure security, dnsuser must be removed from all explicit groups, including the "Users" group, into which all users are placed by default. There should not be a dnsgroup group as is recommended for UNIX.

To log on a service, dnsuser must have the user right "Log on as service." No other rights are necessary. In Windows 2000, users can also be explicitly denied some user rights so that they do not inherit them from group membership. Accordingly, dnsuser must be assigned the user rights "Deny Access to This Computer from the Network" and "Deny Logon as a Batch Job." On properly configured Windows systems dnsuser would never inherit such rights in any case, but the action is a defense-in-depth measure that offers protection when there are authorized changes to the system that could impact dnsuser.

- (DNS4530: CAT II) The SA will create an ISC BIND service user account and configure ISC BIND to run as that user.
- (DNS4540: CAT III) The SA will ensure that the ISC BIND service user is not a member of any group whose membership can be configured.
- (DNS4550: CAT III) The SA will grant the ISC BIND service the user rights of "log on as service," "Deny Access to This Computer from the Network," and "Deny Logon as a Batch Job," which are required for the proper configuration and security of ISC BIND.



### F.3 File Permissions

On UNIX systems, DNS software administrators will sometimes know the root password to perform additional administrative duties based on roles established at each site and will have the ability to switch (or **su**) to root when performing these duties. In Windows, this type of privilege elevation is uncommon. Instead, DNS administrators should be added to a newly created group named dns-admins. When they log on to their regular user accounts, these users will obtain the permissions needed through their membership in the dns-admins group. An alternative is to work through the administrators group, but this would give DNS administrators far more rights and permissions than is necessary to perform their job functions. Consequently, that approach should be avoided.

The Windows SA must ensure that Windows BIND name servers provide comparable file system protections as UNIX BIND name servers.

- (DNS4590: CAT II) The Windows SA will set ownership and permissions on all Windows BIND name servers at least as restrictive as listed in the table below (if a user or group is not listed in the table, then it should not have any permissions to the listed file or folder):

<b>FOLDER/FILE NAME</b>	<b>OWNER</b>	<b>USER/GROUP</b>	<b>PERMISSIONS</b>
%systemroot%\system32\dns\bin	Administrators	Administrators	Full control
		dns-admins	Read
		dnsuser	Read
%systemroot%\system32\dns\etc	Administrators	Administrators	Full control
		dns-admins	Change
		dnsuser	Change
named.conf	Administrators	Administrators	Full control
		dns-admins	Change
		dnsuser	Read
named.pid	Administrators	Administrators	Full control
		dns-admins	Read
		dnsuser	Change
named.stat	Administrators	Administrators	Full control
		dns-admins	Read
		dnsuser	Change
root hints file	Administrators	Administrators	Full control
		dns-admins	Change
		dnsuser	Read
Any zone file	Administrators	Administrators	Full control
		dns-admins	Change
		dnsuser	Change
Any TSIG key file	Administrators	dnsuser	Read

**Table F-2. Windows File Permissions**

Note that `dns-admins` (or an alternative name) must be created before its use. The folders listed are defaults, but the actual location may vary.

Common names for the root hints file are *root.hints*, *named.cache*, or *db.cache*. The name is configurable within the *named.conf* file.

This page is intentionally left blank.

## APPENDIX G. ROOT HINTS

The following table lists the IP address for each of the Internet root servers. DNS software administrators can use this list to validate the entries in their local root hints files. These addresses rarely change.

Root Server	IP Address
A	198.41.0.4
B	192.228.79.201
C	192.33.4.12
D	128.8.10.90
E	192.203.230.10
F	192.5.5.241
G	192.112.36.4
H	128.63.2.53
I	192.36.148.17
J	192.58.128.30
K	193.0.14.129
L	198.32.64.12
M	202.12.27.33

**Table G-1. Root Hints**

This page is intentionally left blank.

## APPENDIX H. LIST OF ACRONYMS

A	Address (a type of DNS resource record)
ACL	Access Control List
APP	Application Peering Protocol
BIND	Berkeley Internet Name Domain
C&A	Certification and Accreditation
CDE	Common Desktop Environment
CERT	Computer Emergency Response Team
CHAP	Challenge Handshake Authentication Protocol
CNAME	Canonical Name (a type of DNS resource record)
CS	Computing Services (a division of DISA)
CSS	Content Services Switch
DECC	Defense Enterprise Computing Center
DECC-D	Defense Enterprise Computing Center-Detachment
DHCP	Dynamic Host Configuration Protocol
DIG	Domain Information Groper
DISA	Defense Information Systems Agency
DMS	Defense Message System
DMZ	Demilitarized Zone (literal meaning). (In the context of network security, this refers to a perimeter network connected to a firewall that is not part of the external or internal network but is accessible by external hosts.)
DNS	Domain Name System
DNSSEC	Domain Name System Security
DOD	Department of Defense
DSN	Defense Switched Network
FIPS	Federal Information Processing Standard
FSO	Field Security Operations
FTP	File Transfer Protocol
FQDN	Fully Qualified Domain Name
GCCS	Global Command and Control System
GSS	Generic Security Service
GUI	Graphical User Interface
HMAC	Keyed-Hashing for Message Authentication

IA	Information Assurance
IAM	Information Assurance Manager
IAO	Information Assurance Officer
IASE	Information Assurance Support Environment
IAVA	Information Assurance Vulnerability Alert
IAVM	Information Assurance Vulnerability Management
IOS	Internetworking Operating System
IP	Internet Protocol
IPSEC	Internet Protocol Security
ISC	Internet Software Consortium
JTF	Joint Task Force
LAN	Local Area Network
LDAP	Lightweight Directory Access Protocol
MAC	Mission Assurance Category
MD5	“Message Digest” 5 Algorithm
MX	Mail Destination / Mail Forwarder (a type of DNS resource record)
NIPRNet	Non-classified (but Sensitive) Internet Protocol Router Network
NIST	National Institute of Standards and Technology
NS	Name Server (a type of DNS resource record)
NSA	National Security Agency
OS	Operating System
PDI	Potential Discrepancy Item
PDL	Potential Discrepancy List
PID	Process ID
PKI	Public Key Infrastructure
PTR	Pointer Record
RFC	Request for Comment
RNOSC	Regional Network Operations and Security Center
RNDC	Remote Name Daemon Control
RR	Resource Record
SA	System Administrator
SHA-1	Secure Hash Algorithm 1
SIPRNet	Secret Internet Protocol Router Network
SMC	Systems Management Center
SOA	Start of Authority
SRR	Security Readiness Review
SRRDB	Security Readiness Review Database
SRV	Service (a type of DNS resource record)
SSC	SIPRNet Support Center
SSH	Secure Shell
SSO	Systems Support Office
STIG	Security Technical Implementation Guide

TCP	Transmission Control Protocol
TLD	Top Level Domain
TS	Terminal Services
TSIG	Transaction Authentication
TXT	Text (a type of DNS resource record)
UDP	User Datagram Protocol
UID	User Identificaiton
UPS	Uninterruptible Power Supply
URL	Uniform Resource Locator
VCTS	Vulnerability Compliance Tracking System
VMS	Vulnerability Management System
VPN	Virtual Private Network



This page is intentionally left blank.