



Global Information Assurance Certification Paper

Copyright SANS Institute
Author Retains Full Rights

This paper is taken from the GIAC directory of certified professionals. Reposting is not permitted without express written permission.

Interested in learning more?

Check out the list of upcoming events offering
"Security Essentials: Network, Endpoint, and Cloud (Security 401)"
at <http://www.giac.org/registration/gsec>

Securely implementing a BIND DNS server on RedHat Linux

Jody Steadman
GSEC V1.4b

Introduction/Summary:

Domain Name Service (DNS) is a vital piece of a network's infrastructure and that same vitality demands a secure implementation. Today, more and more companies are connecting their Local Area Networks (LAN) to the Internet, hosting their own web pages, and utilizing their network to provide services like DNS.

This paper will step through the process of redeveloping and implementing a secure DNS infrastructure for a large organization that spans the continental United States and the world. The organization, which will be called Acme, had a DNS structure developed at a time when computer security was looked at as more of a burden and not an enabler. The configuration and reconfiguration will show what to do and what not to do when configuring DNS using Berkeley Internet Name Daemon (BIND) running on RedHat Linux.

About the Organization:

Acme has a main location, a smaller location across the United States, and smaller satellite offices throughout the world. These will be called IT East, IT West, and satellites respectively. IT East, which has the largest infrastructure, houses both primary and secondary Internic registered DNS servers. IT East has a Class B of Internet routable IP addresses and maintains a firewall at the border to their network.

IT West has a sub-domain furnished by IT East's DNS servers. When a DNS query for IT West's sub-domain is received, the DNS servers in IT East forward the request to the DNS servers in IT West. IT West has limited IP space, so space is conserved using a Network Address Translation (NAT) solution. A PIX firewall sits at the border to their network routable address space. This location utilizes 4 DNS servers, with two DNS servers that reside in the Internet Routable IP space and two residing inside the private address space to resolve private IP addresses on the local network. A point-to-point connection between the two locations does exist and is used for certain administrative functions.

Acme has a major dependence on a properly working DNS infrastructure. The organization's web site is crucial to business operations through marketing and internal and external customers. Acme participates in many research projects and is looked upon as a world leader in its industry. Therefore, continuous day and night operation is essential to its business and standing in the international community.

Old Configuration and its issues:

The DNS servers in IT East are running BIND version 9 on RedHat 7.3. These servers are patched on a regular basis using the RedHat Network and up2date. The servers do reside on separate subnets and in 2 different physical locations within the organization's infrastructure. The OS on the servers were not installed using an IT department wide OS template for RedHat Linux, so the OS

installations are configured differently. The firewall blocks all access to the machines except requests to port 53 (Domain).

The DNS servers in IT west are running BIND version 8 on RedHat 6.2. These servers are rarely patched or upgraded and are not on the RedHat Network. The servers reside on separate subnets and in 2 different physical locations. The OS on the servers were not installed using a department wide OS template. The firewall does block all access to the machines except those requests to port 53.

BIND 8 has a number of bugs and should be patched or upgraded to the latest version, being 8.3.3 as of the date of this paper. One problem, a buffer overflow in transaction signature handling code is serious and can allow arbitrary code to be executed on the overflowed system. This bug affects versions of BIND 8 prior to 8.2.3. (Carnegie)

BIND 9 has a few security concerns, but not as many as BIND 8. For example, BIND 9, versions prior to 9.2.1, has a bug that allows a specifically designed packet to trigger a consistency check in BIND, thus shutting the service down. (Carnegie)

The configuration of BIND for both servers is the same with the exception of a few differences because of the usage of NAT in IT West. When looking through the config files, one can see many features not being used, some not used properly, and some enabled options are insecure. The following issues were discovered when reviewing the old configuration file.

1. Recursion Yes

- This option allows any computer to use the DNS server to resolve an address. When used in conjunction with views, setting this to "no" will disallow unwanted hosts from using the server. If the hosts cannot use the machine, fewer resources are used and another layer of defense is in place.

2. Logging

- The security category is logging to other syslog channels and only provides informative data and not debug. It is a good idea to log as much as possible when it comes to security and use a separate file for logging.
- Queries are not logged. The use of DNS query logs is often essential in tracking down break-ins or break-in attempts.

3. No BIND views

- BIND views allow an administrator to group permissions based on IP address. BIND Views are arguably the most important option for securing a BIND installation.

4. CHAOS class records

- By default, an attacker can gain information regarding the version of BIND by querying the versions.bind and authors.bind chaos records. At the least, these should be changed to something other than the correct version.

BIND process specific configurations, configurations other than the config file, are the same across both IT East and IT West. There are some notable issues with the way BIND was installed and configured.

1. Default installation directory used

- By default, RedHat Linux keeps map files in /var/named. A different directory structure is more secure because it does not reside in a default place.

2. BIND is not chroot jailed

- By chroot jailing a process to a directory, the process may not traverse outside of the specified directory, thus diminishing the success of an intruder/attacker.

3. File Permissions are too loose

- The file permissions are not restrictive enough to stop the reading of map files, config files, and log files. When using chroot jailing, setting permissions within the chroot jail is of the utmost importance.

The servers, where BIND resides, were not adequately secured. While securing BIND does obviously decrease the likelihood of an intrusion, improperly securing the OS leaves the machine open to intruders because of default “out-of-box” configurations.

1. Non-essential services running

- Some services like Sendmail, Telnet, and Apache are running and do not need to be. These programs have their exploits and provide another door for an intruder to enter through.

2. No regular upgrades/patches

- Some of the servers do not have a patch schedule and are not members of the RedHat Network.

Testing and Implementation:

A test environment was setup to accommodate 3 machines running RedHat Linux 7.3 and BIND 9.2.1. One machine represents a primary server residing in IT west and the other 2 represent a primary and secondary server located in IT East. The primary server for IT West houses domains authoritative to that name server. The primary server for IT East contains all domains authoritative for its name server. The secondary server will contain data from the primary in IT East to test the primary-slave relationship.

Goals were established based on problems found in the old config and new action items decided by meetings with those involved in the reconfiguration. They include full redundancy between locations, a true primary slave relationship between servers per location, correcting detected problems with the previous configuration, developing a way to immediately propagate changes without using notify, and changing the primary and secondary DNS server assignments with Internic.

In doing research, a Secure BIND template was found that would provide the framework for the reconfiguration. This template employs many of the features needed in the reconfiguration (Thomas). The template will be "tweaked" as needed to accommodate specifics needed by this reconfiguration.

In the following, I will step through each of the issues discovered and the implemented solution. The master config file and slave config file is located in appendix A and appendix B respectively.

1. No BIND views

- Named.conf file entries:

```
view "internal-hosts" in {
    match-clients {
        trusted-hosts;
    };
    recursion yes;
    ...
};
view "external-hosts" in {
    match-clients {
        any;
    };
    recursion no;
    ...
};
```

};

- Views are divided by internal and external views. The internal view identifies internal hosts using match-clients and the trusted-hosts acl (access control list). Recursion is set to yes and will be discussed in the next bullet. The external view is more restrictive and places any host that is not included in the internal view in the external view. The “any” setting is used with the “match-clients” option to achieve this. Recursion is set to no (Internet).

2. Recursion

- Named.conf file entries:

```
view "internal-hosts" in {  
    ...  
    recursion yes;  
    ...  
};  
view "internal-hosts" in {  
    ...  
    recursion no;  
    ...  
};
```

- Recursion allows a host to resolve addresses not authoritative for the DNS server. This is not advantageous for 2 reasons. Good security is the tightest security possible without hindering normal operations, so denying recursion for hosts outside the network is acceptable. Secondly, allowing outside hosts to use recursion ties up resources and could give false data when planning for additional resources on the machine. (Internet)

3. Logging

- Named.conf file entries:

```
logging {  
    channel "default_syslog" {  
        syslog local2;  
        severity info;  
    };  
    channel secure_log {  
        file "bind/named.log";  
        severity debug;  
        print-time yes;  
    };  
};
```

```
};

category default { default_syslog; };
category general { default_syslog; };
category security { secure_log; };
category config { default_syslog; };
category resolver { default_syslog; };
category xfer-in { secure_log; };
category xfer-out { secure_log; };
category client { secure_log; };
category network { secure_log; };
category queries { secure_log; };
category lame-servers { default_syslog; };

};
```

- The template comes with logging split into 2 separate channels by using the channel option. "secure_log" is for security and "default_syslog" is for general information from BIND. BIND allows categorization of log entries using the category option. For more information on these categories, see the BIND 9 Admin guide located in the documentation section on ISC's web site.

4. CHAOS class records

- Named.conf file entries:

```
view "external-chaos" chaos {
    match-clients {
        any;
    };
    recursion no;
    zone "." {
        type hint;
        file "/dev/null";
    };
    zone "bind" {
        type master;
        file "db.bind";
        allow-query {
            trusted-hosts;
        };
        allow-transfer {
            none;
        };
    };
};
```


- By using this method, chaos record queries can be logged for obvious security reasons. It is important to note that a version statement in the options section of the named.conf will block version info, but will not log the event. The db.bind file is used in conjunction with the chaos view in the config file and can contain a customizable response to the chaos version query. (Internet)

5. Default directory used and BIND not chrooted

- /etc/init.d/named file entries:

```
[ -f /usr/sbin/named ] || exit 0
[ -f /named/etc/named.conf ] || exit 0
case "$1" in
    start)
        # Start daemons.
        echo -n "Starting named: "
        daemon /usr/sbin/named -u named -t /named/ -
            c /etc/named.conf
        echo
        touch /var/lock/subsys/named
    ;;
```

- The named init file was modified with the above changes to start BIND in chroot-jailed environment every time the server is booted. In the above excerpt from the init script, the daemon line is the most critical portion of this configuration. The “/usr/sbin/named” is the binary to be executed, “-u named” is the user to run BIND as “-t /named” tells named to chroot to the specified directory, and “-c /etc/named.conf” tells the daemon which config file to use when starting the daemon. The “/named” directory is where the zone files, configs, and some logs will reside.

6. File permissions are too loose:

/named directory structure:

- All files and directories have the strictest permissions on them only allowing what it needs. The default directory permission, including /named/, is 770, because named is the owner. File permissions will be discussed in the next section.
- /named/bind/: This directory contains the security log called named.log. It can include other logs, as we deem necessary.

- `/named/etc/`: This directory contains the `named.conf` and could contain other files at our discretion.
- `/named/master/`: This directory contains all zone files that are master to the server and shell/PERL scripts used to administrate the server.
- `/named/slave/`: This directory contains all zone files that are slave zones for the server. It contains no scripts, because changes in zone files should occur on the master for the zone. These files are generated by BIND.

Specific file structure and permissions:

- This section discusses the files that reside in the `/var/named/` directories that are not usual components of BIND. Default file permissions for zone related files are 664 and owned by root unless otherwise specified.
- `/named/db.bind`: This file stores the CHAOS class records for version and authors. Having a file in place, instead of declaration in `named.conf`, allows us to log CHAOS class queries and set the information we want others to view.
- `/named/bind/named.log`: This is an audit log file for security related logging. Default mask is 770 and owned by named.
- `/named/dev/*`: This contains the “null” and “random” needed by BIND. Default permission is 666. “`mknod /named/dev/null c 1 3`” and “`mknod /named/dev/random c 1 8`” are used for it’s creation respectively.
- `/named/etc/named.conf`: This is the config file used by BIND. It cannot reside in `/etc` because of the chrooted environment. The file is owned by root and has permissions of 664.
- `/named/master/named.*`: These are the master zone files for the server. They are owned by root with 664 permissions.
- Scripts in `/named/master`: The script ‘`updateslaves.pl`’ are owned by root with 770 permissions because we only want the root user to execute these scripts.
- `/named/slave`: These are slave zone files from the slave server/s. Permissions are default from BIND with 664 owned by named. It

must stay this way so BIND can update and/or replace them when needed.

- /etc/init.d/named: This file has been modified to accommodate the chrooted environment.

7. Non-essential services running:

- All services were shutdown except BIND and SSH. This was accomplished by running `setup` and deselecting services from startup and using `chkconfig` in Linux. The `xinet.d` directory located in `/etc` contains files used by `xinetd` to start certain services. The person installing the machine should look at these files and verify what should and should not be running. These files are named the name of the service they start. An option to disable the service is in each of those files located in the `xinet.d` directory. (Guardian)
- A Nessus scan was done against all the machines before and after it was patched on the RedHat Network. All services were verified turned off.

8. No regular upgrades/patches:

- All servers were placed on the RedHat Network. The RedHat Network options were set to notify the administrator of any upgrades/patches available for the machines.

9. Provide full redundancy between locations:

- The zones authoritative in IT East were slaved on IT West's DNS and vice-versa. The transfers occur over the PTP connection between locations. The zone transfers were secured using an `acl` in the `named.conf` file.

10. Implement Primary-Secondary relationship between servers per location:

- In IT East, one secondary server is used. The secondary slaves all zones from the primary. This includes the zones the primary server slaves for IT West. The same setup exists in IT West, except more DNS servers are used to accommodate the private address space. One Primary and Secondary resides on the routable IP space, and 2 exist inside the private network. The servers in the private network slave all authoritative zones from the Primary.

11. Immediate and centralized updates without using notify:

- A script was created that reloads the zone on the primary server and logs into the secondary servers via SSH. After login, the script issues an "rndc reload" of the zone. Multiple servers can exist in the script. The script also notifies on a failure to update. See Appendix C.

12. Change primary and secondary DNS server assignments with Internic:

- A request was made with Internic to have the Root Servers changed so the domain primary server is the primary in IT East and the domain secondary server is the primary in IT West.

Mission Complete

The servers were configured using the tested configuration and all problems were addressed. When a contrast is made between how the DNS infrastructure was originally setup and how it is presently, one can see how important and easy it is to properly configure DNS infrastructure to provide some redundancy and security in an increasingly hostile Internet environment.

Appendices:

Appendix A (Master Config):

```
//  
// Master server config file.  
// Modified Date  
//  
  
acl "transfer" {  
    192.168.100.1;  
    192.168.100.2;  
    172.16.0.1;  
};  
  
acl "trusted-hosts" {  
    192.168/16;  
    172.16.0.1;  
};  
  
acl "bogon" {  
    See Rob Thomas's Secure BIND version 3.7 for list of IP addresses  
};
```

```

key "rndc-key" {
    algorithm hmac-md5;
    secret "****insert key here****";
};

controls {
    inet 127.0.0.1 port 953
    allow {
        127.0.0.1;
    }
    keys {
        "rndc-key";
    };
};

logging {

    channel "default_syslog" {
        syslog local2;
        severity debug;
    };

    channel secure_log {
        file "bind/named.log";
        severity debug;
        print-time yes;
    };

    category default { default_syslog; };
    category general { default_syslog; };
    category security { secure_log; };
    category config { default_syslog; };
    category resolver { default_syslog; };
    category xfer-in { secure_log; };
    category xfer-out { secure_log; };
    category client { secure_log; };
    category network { secure_log; };
    category queries { secure_log; };
    category lame-servers { default_syslog; };

};

options {
    pid-file "named.pid";
};

```

```
statistics-file "named.stats";
dump-file "named.dump";
zone-statistics yes;
notify no;
transfer-format many-answers;
max-transfer-time-in 60;
interface-interval 0;

allow-transfer {
    localhost;
    transfer;
};

allow-query {
    trusted-hosts;
};

blackhole {
    bogon;
};
};

view "internal-hosts" in {
    match-clients {
        trusted-hosts;
    };
    recursion yes;
    additional-from-auth yes;
    additional-from-cache yes;
};

zone "." in {
    type hint;
    file "named.root";
};

zone "0.0.127.in-addr.arpa" in {
    type master;
    file "named.local";
    allow-query {
        any;
    };
    allow-transfer {
        none;
    };
};
```

```

//
// IT East's zone files
//

zone "168.192.in-addr.arpa" {
    type master;
    file "master/named.rev";
    allow-query {
        any;
    };
};

zone "acme.org" {
    type master;
    file "master/acme_org";
    allow-query {
        any;
    };
};

//
//IT West's Zone files
//

zone "16.172.in-addr.arpa" in {
    type slave;
    file "slave/named.16_172_in-addr_arpa";
    allow-query {
        any;
    };
    masters {
        172.16.0.1;
    };
};

zone "west.acme.org" in {
    type slave;
    file "slave/named.west_acme_org";
    allow-query {
        any;
    };
    masters {
        172.16.0.1;
    };
};

```

```

view "external-hosts" in {
    match-clients {
        any;
    };
    recursion no;
    additional-from-auth yes;
    additional-from-cache no;

zone "." in {
    type hint;
    file "named.root";
};

//
// IT East's zone files
//

zone "168.192.in-addr.arpa" {
    type master;
    file "master/named.rev";
    allow-query {
        any;
    };
};

zone "acme.org" {
    type master;
    file "master/acme_org";
    allow-query {
        any;
    };
};

//
//IT West's Zone files
//

zone "16.172.in-addr.arpa" in {
    type slave;
    file "slave/named.16_172_in-addr_arpa";
    allow-query {
        any;
    };
    masters {
        172.16.0.1;

```



```
};  
};  
  
zone "west.acme.org" in {  
    type slave;  
    file "slave/named.west_acme_org";  
    allow-query {  
        any;  
    };  
    masters {  
        172.16.0.1;  
    };  
};  
  
view "external-chaos" chaos {  
    match-clients {  
        any;  
    };  
    recursion no;  
  
    zone "." {  
        type hint;  
        file "/dev/null";  
    };  
  
    zone "bind" {  
        type master;  
        file "db.bind";  
  
        allow-query {  
            trusted-hosts;  
        };  
        allow-transfer {  
            none;  
        };  
    };  
};  
};
```

Appendix B (Slave Config):

```
//  
// Master server config file.  
// Modified Date  
//  
  
acl "transfer" {  
    192.168.100.1;  
    192.168.100.2;  
    172.16.0.1;  
};  
  
acl "trusted-hosts" {  
    192.168/16;  
    172.16.0.1;  
};  
  
acl "bogon" {  
    See Rob Thomas's Secure BIND version 3.7 for list of IP addresses  
};  
  
key "rndc-key" {  
    algorithm hmac-md5;  
    secret "****insert key here****";  
};  
  
controls {  
    inet 127.0.0.1 port 953  
    allow {  
        127.0.0.1;  
    }  
keys {  
    "rndc-key";  
};  
};  
  
logging {  
  
    channel "default_syslog" {  
        syslog local2;  
        severity debug;  
  
    };  
};
```

```

channel secure_log {
file "bind/named.log";
severity debug;
print-time yes;

};

category default { default_syslog; };
category general { default_syslog; };
category security { secure_log; };
category config { default_syslog; };
category resolver { default_syslog; };
category xfer-in { secure_log; };
category xfer-out { secure_log; };
category client { secure_log; };
category network { secure_log; };
category queries { secure_log; };
category lame-servers { default_syslog; };

};
options {
pid-file "named.pid";
statistics-file "named.stats";
dump-file "named.dump";
zone-statistics yes;
notify no;
transfer-format many-answers;
max-transfer-time-in 60;
interface-interval 0;

allow-transfer {
localhost;
transfer;
};

allow-query {
trusted-hosts;
};

blackhole {
bogon;
};
};

```

```

view "internal-hosts" in {
    match-clients {
        trusted-hosts;
    };
    recursion yes;
    additional-from-auth yes;
    additional-from-cache yes;

zone "." in {
    type hint;
    file "named.root";
};

zone "0.0.127.in-addr.arpa" in {
    type master;
    file "named.local";
    allow-query {
        any;
    };
    allow-transfer {
        none;
    };
};

//
// IT East's zone files
//

zone "168.192.in-addr.arpa" {
    type slave;
    file "slave/named.rev";
    allow-query {
        any;
    };
    masters {
        192.168.100.1;
    };
};

zone "acme.org" {
    type slave;
    file "master/acme_org";
    allow-query {
        any;
    };
    masters {

```

```

        192.168.100.1;
    };
};

//
//IT West's Zone files
//

zone "16.172.in-addr.arpa" in {
    type slave;
    file "slave/named.16_172_in-addr_arpa";
    allow-query {
        any;
    };
    masters {
        192.168.100.1;
    };
};

zone "west.acme.org" in {
    type slave;
    file "slave/named.west_acme_org";
    allow-query {
        any;
    };
    masters {
        192.168.100.1;
    };
};

view "external-hosts" in {
    match-clients {
        any;
    };
    recursion no;
    additional-from-auth yes;
    additional-from-cache no;
};

zone "." in {
    type hint;
    file "named.root";
};

//
// IT East's zone files
//

```

```
zone "168.192.in-addr.arpa" {
    type slave;
    file "slave/named.rev";
    allow-query {
        any;
    };
    masters {
        192.168.100.1;
    };
};

zone "acme.org" {
    type slave;
    file "slave/acme_org";
    allow-query {
        any;
    };
    masters {
        192.168.100.1;
    };
};

//
//IT West's Zone files
//

zone "16.172.in-addr.arpa" in {
    type slave;
    file "slave/named.16_172_in-addr_arpa";
    allow-query {
        any;
    };
    masters {
        192.168.100.1;
    };
};

zone "west.acme.org" in {
    type slave;
    file "slave/named.west_acme_org";
    allow-query {
        any;
    };
    masters {
        192.168.100.1;
    };
};
```

```

};
};

view "external-chaos" chaos {
    match-clients {
        any;
    };
    recursion no;

    zone "." {
        type hint;
        file "/dev/null";
    };

    zone "bind" {
        type master;
        file "db.bind";

        allow-query {
            trusted-hosts;
        };
        allow-transfer {
            none;
        };
    };
};

```

Appendix C (Update Slaves Script):

```

#!/usr/bin/perl

use Net::SSH::Perl;
require '/etc/passwd.pl';

$zone = @ARGV[0];

if ($zone eq "") {
    print "usage: ./updateslaves.pl <zone-name>\n";
    exit;
}

# Reload the primary

```

```
system("rndc reload $zone IN internal-in");
system("rndc reload $zone IN external-in");
```

```
# Slave servers
```

```
$slave1 = '192.168.100.2';
```

```
# Commands to be ran
```

```
$cmd1 = "/usr/sbin/rndc reload $zone in internal-in";
$cmd2 = "/usr/sbin/rndc reload $zone in external-in";
```

```
print "Logging in to $slave1\n";
```

```
my $ssh1 = Net::SSH::Perl->new($slave1);
$ssh1->login($user, $pass);
```

```
print "Updating $slave1\n";
$foutput1 = $ssh1->cmd($cmd1);
$soutput1 = $ssh1->cmd($cmd2);
```

```
if (($foutput1 > 0) or ($soutput1 > 0)) {
    print "Reload of $slave1 failed\n";
}
```

```
else {
    print "Reload of $slave1 successful\n";
    print "Please allow 10 minutes for propagation\n";
}
```

© SANS Institute 2003, Author retains full rights.

Works Cited

Carnegie Mellon University. "CERT Advisory CA-2001-02 Multiple Vulnerabilities in BIND." 29 Jan. 2001. URL: <http://www.cert.org/advisories/CA-2001-02.html> (10 Jan. 2003).

Carnegie Mellon University. "CERT Advisory CA-2002-15 DOS Vulnerability in ISC BIND 9." 4 Jun. 2002. URL: <http://www.cert.org/advisories/CA-2002-15.html> (11 Jan. 2003).

Guardian Digital Inc. "Linux Security Quick Reference Card." 2000. URL: <http://www.linuxsecurity.com/docs/QuickRefCard.pdf> (27 Feb. 2003)

Internet Software Consortium. "BIND 9 administrators guide." 2001. URL: <http://www.nominum.com/content/documents/bind9arm.pdf>. (3 Feb. 2003).

Thomas, Rob. "Secure BIND Template Version 3.7." 13 Feb 2003. URL: <http://www.cymru.com/Documents/secure-bind-template.html>. (20 Feb. 2003).

© SANS Institute 2003, Author retains full rights.