

Domain Name Service

Configuration



PSINet *Setting up a Name Server*

- Hardware Resources**
 - Dedicated or shared machine
- Time**
 - Setup and administration
- Information**
 - Host names and IP address
 - Mail exchangers
 - Any aliases

When setting up a name server, you need to consider many things.

First and foremost are the resources necessary. Appropriate hardware must be allocated to the name server. Depending how busy you expect it to be, you may even need to dedicate a machine for name service.

If your name server is going to be the primary source of domain information for your site, then the name administrator at your site will be responsible for insuring the information is correct and available. This includes keeping the machine running as well as insuring that the information is accessible to other name servers when queried.

Finally, the host information that is going to be stored on the server must be collected and meet the needs of your users.

- Checked when the name server starts up.
- Directives
 - directory
 - cache
 - primary
 - secondary

```
directory      /etc/
cache          .          named.ca
primary        company.com  company.com
primary        125.7.204.in-addr.arpa  company.rev
primary        0.0.127.in-addr.arpa  named.local
secondary      foo.com     38.8.75.3 foo.com
secondary      193.2.206.in-addr.arpa  38.8.75.3 foo.rev
```

The **boot file** is read when the name server starts up. It defines the location of the zone files that the name server is responsible for as well as the root-servers file. Above is the sample boot file configuration on the ns.company.com. name server.

directory

This line defines the working directory of zone information.

cache

Denotes the location of the root-servers file. This file contains a list of all the root name servers and their IP addresses. It must be updated periodically on each name server. A copy of the file can be FTPed from rs.internic.net.

primary

This directive tells the name server the zone(s) that it is primary for and what the name of each zone file is. It is also used for the reverse zone as well as the local host zone.

secondary

This name server is also secondary for foo.com and the reverse of foo.com, 193.2.206.in-addr.arpa. Ns.company.com does a zone transfer from a primary name server for foo.com at 38.8.75.3.



Cache File

- Current list of root name servers
- Must be kept current!

```
.           3600000 IN      NS      A.ROOT-SERVERS.NET.
A.ROOT-SERVERS.NET.  A      198.41.0.4
.           3600000     NS      B.ROOT-SERVERS.NET.
B.ROOT-SERVERS.NET.  A      128.9.0.107
.           3600000     NS      C.ROOT-SERVERS.NET.
C.ROOT-SERVERS.NET.  A      192.33.4.12
.           3600000     NS      D.ROOT-SERVERS.NET.
D.ROOT-SERVERS.NET.  A      128.8.10.90
.           3600000     NS      E.ROOT-SERVERS.NET.
E.ROOT-SERVERS.NET.  A      192.203.230.10
.           3600000     NS      F.ROOT-SERVERS.NET.
F.ROOT-SERVERS.NET.  A      39.13.229.241
.           3600000     NS      G.ROOT-SERVERS.NET.
G.ROOT-SERVERS.NET.  A      192.112.36.4
.           3600000     NS      H.ROOT-SERVERS.NET.
H.ROOT-SERVERS.NET.  A      128.63.2.53
.           3600000     NS      I.ROOT-SERVERS.NET.
I.ROOT-SERVERS.NET.  A      192.36.148.17
```

Copyright © 1997 PSINet Inc.

Confidential and Proprietary Information

T10.4

This file needs to be updated manually in order to be kept current. This file is necessary for a name server to successfully query for any information outside of its zone because all zone delegations start at the root name servers.

In the previous boot file configuration, this file was named named.ca. A copy of this file can be FTPed from rs.internic.net.

Resource Records (RR)

- Stored on name servers
- Queried for by resolvers
- Two main types: structural and node data
- All queries are based on:
 - label**
The domain name that the RR is associated with
 - class**
Almost always “IN” for Internet
 - type**
The RR type, e.g., “A” for “address”
- The general form for a resource record is:
<name> <TTL> <class> <type> <data>

Copyright © 1997 PSINet Inc.

Confidential and Proprietary Information

T10.5

The **resource records** are the information stored in the zone files on the name servers. There are two main types of information: structural and node data. Structural data is mostly name service overhead that facilitates the distribution of name information. Node data is the information that is usable by the end user of the service.

All queries are based on the label, class and type of a resource record. The **label** is the domain name that the RR is associated with. The **class** for our purpose is IN (Internet) and the **type** is the kind of DNS record that is listed. These will each be covered in more detail later.

<name>

Domain name of the host that the record is associated with.

<TTL>

Time to Live. This is the length of time that a name server is allowed to cache the record. When the time expires the name server can no longer use the data and must query an authoritative name server again for the data.

<class>

IN Internet. The scope of our discussion will be limited to the Internet.

<type>

This is the type of record. For example, A, PTR, CNAME, NS, etc.

<data>

The appropriate data for the record. It can be an IP address, or name or a brief sentence about the server.



Resource Records: Structural Data

- Structural data relates to the entire zone.
- Structural data specifies:
 - Zone boundaries
 - Delegation information
i.e., what name servers handle psi.net?
- Includes DNS overhead records
 - SOA - Start of Authority
 - NS - Name Server

The **structural data** specifies where the name server's authority starts and stops for each domain it is authoritative for. The structural records are the **start of authority record (SOA)** and the **Name Server (NS)**.



Resource Records: Node Data

- Node data specifies characteristics of each node in the zone.
- Includes host specific records.
 - A Name-to-Address Mapping
 - PTR Address-to-name mapping
 - CNAME Canonical Name
 - MX Mail Exchanger

Node data specifies characteristics of each node in the zone. For example, the host ns.psi.net has an IP address associated with it. The scope of our discussion will include A, PTR, CNAME, MX but there are other types of node data such as HINFO, TXT and WKS.



Resource Records: SOA Record

- Structural Resource Record.
- Identify the start of a new DNS zone.
- Required for secondary zone transfers.

```
company.com. IN      SOA      ns.company.com.    jane.comany.com. (
                    96062101    ; serial
                    10800    ; refresh (3 hours)
                    3600    ; retry (1 hour)
                    1209600   ; expire (14 days)
                    86400    ; minimum (1 day)
                    )
```

Copyright © 1997 PSINet Inc.

Confidential and Proprietary Information

T10.8

The **SOA record** identifies the start of a new DNS zone. There is always only one SOA record for each zone. The record outlines information which is mostly useful to a secondary name server. Therefore, if the SOA record does not exist, then the zone transfer from the secondary name server will fail and any information that it had will eventually become stale.

origin company.com. The domain name of the root of this zone.

SOA This is the record type. Start of Authority.

source ns.company.com. The name server this record was created on.

person jane.comany.com. Zone administrator. jane@company.com

serial

The version number of the current zone data in the file. This needs to be incremented whenever any change to the data is made so that any secondary name servers know to do a zone transfer.

refresh

A directive to the secondary name server to check if a zone transfer is necessary by checking the value of the serial number.

retry

If a refresh attempt fails, the secondary should retry every 3600 seconds.

expire If a retry is never successful, then expire all zone information.

minimum The default TTL for records with no specific setting.



Resource Records: NS Records

- Structural Resource Record
- NS records identify the authoritative name servers for the zone.
- Returned at each stage of an iterative query.

```
company.com. IN      SOA      ns.company.com.    jane.company.com. (  
96062101           ; serial  
10800              ; refresh (3 hours)  
3600               ; retry (1 hour)  
1209600            ; expire (14 days)  
86400              ; minimum (1 day)  
)  
company.com. IN      NS       ns.company.com.  
company.com. IN      NS       ns2.psi.net.
```

Copyright © 1997 PSINet Inc.

Confidential and Proprietary Information

T10.9

NS records are structural records used to identify the boundaries of the zone. There are two sets of NS records for each zone. One is kept on the delegating name server and the other is on the name server being delegated to. When name servers delegate authority for a zone away, they must retain data on who has the information so when they are queried they can at least point towards the appropriate name server. This is what makes the distributed name service work. If the NS records are not configured correctly, then query results will be unpredictable.

NS

Name Server. An authoritative name server for this zone. Ns.company.com and ns2.psi.net are the domain names of authoritative name servers for company.com.



Resource Records: A Records

- A records map names to IP addresses.
- There should be one A record for each IP address of a host.
- Accordingly, multi-homed hosts may have more than one A record.

```
company.com. IN SOA ns.company.com. jane.company.com. (  
[serial, refresh, retry, expire, TTL]  
)  
company.com. IN NS ns.company.com.  
company.com. IN NS ns2.psi.net.  
smtp.company.com. IN A 204.7.125.2  
www.company.com. IN A 204.7.125.3  
mac1.company.com. IN A 204.7.125.10  
rock.company.com. IN A 204.7.125.11  
ns.company.com. IN A 204.7.125.15
```

Copyright © 1997 PSINet Inc.

Confidential and Proprietary Information

T10.10

A records are node data that map host names to IP addresses.

There should be one A record for each IP address of a host. Therefore, it is acceptable to have one host name map to multiple IP addresses.

On the same token, it is also acceptable to have multiple host names map to the same IP address. However, this does not imply that you can have multiple machines with the same IP address.

The CNAME record can also be used to alias a name to an existing host in this fashion.



Resource Records: PTR Records

- Map IP addresses to a hostname.
- PTR (“reverse lookup”) records are part of a different zone than the forward records.
- Also handled with NS and SOA records.

```
125.7.204.in-addr.arpa. IN      SOA   ns.company.com. jane.company.com. (
                        [serial, refresh, retry, expire, TTL]
                        )
125.7.204.in-addr.arpa. IN      NS    ns.company.com.
125.7.204.in-addr.arpa. IN      NS    ns2.psi.net.

2.125.7.204.in-addr.arpa. IN     PTR   smtp.company.com.
3.125.7.204.in-addr.arpa. IN     PTR   www.company.com.
10.125.7.204.in-addr.arpa. IN    PTR   mac1.company.com.
11.125.7.204.in-addr.arpa. IN    PTR   rock.company.com.
15.125.7.204.in-addr.arpa. IN    PTR   ns.company.com.
```

Copyright © 1997 PSINet Inc.

Confidential and Proprietary Information

T10.11

PTR records are used to map IP addresses back to host names.

The reverse, as it is commonly called, has its own zone which is part of the in-addr.arpa domain. The in-addr (inverse address) is written in the reverse.

In the case of company.com which uses the 204.7.125 network, there would be two zone files. The first would be used for forward lookups and would be called company.com. The zone for the reverse lookups would be called 125.7.204.in-addr.arpa.

As with the company.com zone, ns.company.com and ns2.psi.net will be authoritative for the 125.7.204.in-addr.arpa zone.



Resource Records: CNAME Records

- Specifies host name aliases.
- Preferred method of aliasing.
- CNAME carries the data not the alias.

```
company.com. IN      SOA      ns.company.com. jane.company.com. (  
                  [serial, refresh, retry, expire, TTL]  
                  )  
www.company.com.  A          204.7.125.3  
ftp.company.com.  CNAME     www.company.com.
```

The canonical (official) name of a host can also have aliases.

In the above example, www has the alias of ftp. So a query on ftp.company.com will return the IP of www and note that ftp is an alias for www.



Resource Records: MX Records

- Specify where mail should be sent.
 - Final delivery
 - Forwarding
 - Outages
- Prioritize mail destinations.
 - "Final" destination has lowest preference value.

Copyright © 1997 PSINet Inc.

Confidential and Proprietary Information

T10.13

MX records are used to help route electronic mail to its final destination. They are particularly useful when hosts don't process their own mail but rely on some other machine to accept it for them.

Such customers are UUPSI customers who UUCP their mail from a PSINet UUPSI server. Correctly configured MX records are used to route mail destined for uupsicust.com to a UUPSI server. Customers connect to the UUPSI machine periodically to retrieve their mail.

TCP connected sites running a POP3 server also find MX records useful. Mail destined for chris@securehost.company.com can be routed to a general POP3 server such as pop3.company.com. Chris can then retrieve his mail using a POP3 client running on securehost.

Another feature of the MX record is that alternative hosts that are able to process mail or store it temporarily can be added to the zone file. These sites can act as a backup if the primary mail server is unreachable. The MX records use a prioritization scheme to determine which host should be contacted first to deliver mail. The scheme is priority based; the MX record with the lowest preference value has the highest priority.



Resource Records: MX Records

company.com. IN	SOA	ns.company.com. jane.company.com. ([serial, refresh, retry, expire, TTL])
company.com.	MX	10 smtp.company.com.
company.com. MX	5000	mx.smtp.psi.net.
company.com.	MX	5000 mx2.smtp.psi.net.
smtp.company.com.	MX	10 smtp.company.com.
smtp.company.com.	MX	5000 mx.smtp.psi.net.
smtp.company.com.	MX	5000 mx2.smtp.psi.net.
*.company.com.	MX	10 smtp.company.com.
*.company.com.	MX	5000 mx.smtp.psi.net.
*.company.com.	MX	5000 mx2.smtp.psi.net.

Copyright © 1997 PSINet Inc.

Confidential and Proprietary Information

T10.14

smtp Mailhost for the domain.

Any mail sent to someone@company.com will be directed to the mail server at smtp.company.com.

PSINet offers backup Mail Exchangers if your site has an outage. In this case if smtp.company.com is down, the mail will be rerouted to a PSINet mail exchanger. Since the PSINet mx machines are set to the same priority, the mail will be sent to one or the other. Once it is on one PSINet exchanger, it will only be forwarded on to someone or sent back to the sender after 7 days of failed delivery attempts.

The wildcard * is used to denote any other match. So mail sent to someone@gogly-moogly.company.com. will be forwarded on to smtp.company.com and be backed up by mx and mx2.



Configuration: Zone File

```
company.com. IN SOA ns.company.com. jane.company.com. (  
[serial, refresh, retry, expire, TTL])  
company.com. IN NS ns.company.com.  
company.com. IN NS ns2.psi.net.  
smtp.company.com. IN A 204.7.125.2  
www.company.com. IN A 204.7.125.3  
ftp.company.com. IN CNAME www.company.com.  
mac1.company.com. IN A 204.7.125.10  
rock.company.com. IN A 204.7.125.11  
ns.company.com. IN A 204.7.125.15  
company.com. IN MX 10 smtp.company.com.  
company.com. IN MX 5000 mx.smtp.psi.net.  
company.com. IN MX 5000 mx2.smtp.psi.net.  
smtp.company.com. IN MX 10 smtp.company.com.  
smtp.company.com. IN MX 5000 mx.smtp.psi.net.  
smtp.company.com. IN MX 5000 mx2.smtp.psi.net.  
*.company.com. IN MX 10 smtp.company.com.  
*.company.com. IN MX 5000 mx.smtp.psi.net.  
*.company.com. IN MX 5000 mx2.smtp.psi.net.
```

Copyright © 1997 PSINet Inc.

Confidential and Proprietary Information

T10.15

This is the final zone file for company.com.

Remember that there is a second zone file for the 125.7.204.in-addr.arpa zone that is not shown here.

Note that the SOA records are not written in full. Refer to the previous slide on the SOA record for the correct configuration.



Resource Records: Zone File Shorthand

@	IN	SOA	ns.company.com.	jane.company.com. (
			[serial, refresh, retry, expire, TTL])	
		NS	ns.company.com.	
		NS	ns2.psi.net.	
		MX	10	smtp.company.com.
		MX	5000	mx.smtp.psi.net.
		MX	5000	mx2.smtp.psi.net.
smtp		A	204.7.125.2	
		MX	10	smtp.company.com.
		MX	5000	mx.smtp.psi.net.
		MX	5000	mx2.smtp.psi.net.
www		A	204.7.125.3	
ftp		CNAME	www	
mac1		A	204.7.125.10	
*		MX	10	smtp.company.com.
		MX	5000	mx.smtp.psi.net.
		MX	5000	mx2.smtp.psi.net.
rock		A	204.7.125.11	
ns		A	204.7.125.15	

Copyright © 1997 PSINet Inc.

Confidential and Proprietary Information

T10.16

Building a zone file can be quite tedious. It is very important that all information is logically and syntactically correct. Even forgetting a “.” will cause problems.

Zone files can become relatively large, complex and difficult to read. It is possible to use “shorthand” when configuring the file to make it easier to read and reduce the risk or error. Here is the same company.com zone file rewritten using some shorthand.

@ Defined as the zone of origin, or the domain at the root of the zone. It is appended to all names in the zone files that do not already end in a dot.

The records for each host can be divided into sections. When the label field is left blank, the last label mentioned is used.

