NetDB and DNS at Stanford

Drew W. Saunders
UIT Networking

NetDB at Stanford

- NetDB was developed at Stanford in the late ’80’s, converted to a Web application in the late ’90’s, and will likely never have fancy Javascript or other very modern “features” that just slow it down.
    - It is open source, but relies on Oracle, so isn’t used by many other organizations.
- NetDB is the source for our DNS servers
- NetDB is the source for our DHCP servers

- NetDB can be used for other purposes beyond just DNS and DHCP. There are lots of fields that are potentially of use by many organizations or automated systems, and an extensive search function.

- Access within NetDB is by “Group” and “Record Type.”

- The “Administrator” field does NOT grant any access to modify a record. It is used extensively by other systems (such as the SSL Certification group), but within NetDB, it’s just data.

- If you can’t modify a record, most likely you don’t have the correct Group access, or it’s a Record Type you can’t modify.

NetDB and DNS

- NetDB is the source for all of Stanford's DNS information (plus a few custom files for really strange stuff)
- .Stanford.EDU is appended by default, so it's not displayed, but it's there.
- If you use 3rd or 4th level domains, you have to type them all out. If you use a domain other than Stanford.EDU, you'll have to type it all out as well.
- We also use ".Nodomain" if you want to use a name that never gets to DNS, and ".sunet" if you want a name to exist only on campus, but never get sent outside.

- DNS is updated every half hour from NetDB, and the update takes 5-20 minutes. It's generally done at 5 and 35 past the hour.
- The NetDB Name field creates the A record in DNS (and, if appropriate, the AAAA record for IPv6). A is for "Authoritative."
- The NetDB Alias field creates the CNAME ("Canonical," which is lesser than "Authoritative"), but the term "Alias" is commonly used as well. An A name can have many Aliases/CNAMEs.
- The MX field is used to redirect mail using the MX (Mail Alias) feature of DNS. If you want mail sent to host1 to end up at host2, you enter "host1" into the "Receive mail for" field for host2.
- An "Advanced" node can have multiple Names (each with their own list of Aliases) as well as interface Names (but no interface Aliases), which can be handy for complicated servers.

NetDB and DHCP

- Stanford uses two centralized DHCP servers (Dusk and Dawn) to provide campus-wide DHCP. Running your own departmental DHCP server is strongly discouraged, but can be accommodated for special projects. UIT configures the switches we run to reject DHCP replies from user ports.

- Although DHCP is not routed by default, we have put in “helper” addresses on the router interfaces for nearly every network to forward DHCP requests to dusk and dawn.

- DHCP is designed to pick from a pool of numbers, but you can limit that pool to one number, so that you always get the same IP number. That’s very commonly used at Stanford.

- The more traditional large pool of IP numbers is called “Roaming” DHCP at Stanford, and is the primary way that DHCP works on wireless, but can be set up for wired as well.

- To have NetDB send the desired information to Dusk and Dawn, the MAC address has to be entered and at the least “DHCP” has to be checked. If it’s not in NetDB, it won’t get a DHCP response (with some exceptions).

- If you need one IP number to be handed out every time, enter that IP number and, if the DHCP request comes from that subnet, Dusk and Dawn will reply with only that number.

- If you need to be able to roam, check the “roaming” checkbox. For wireless connections, not checking “roam” will make your entry useless.

- For wired connections, you can have a fixed IP for one (or more) subnet(s) and roaming for all others: just check both DHCP and Roaming and enter the fixed IP number(s) you need appropriate to the subnet(s) where you will regularly connect to the wired network.

Data fields of NetDB

- Department: If your department is not listed, or has changed names, please put in a help ticket. There is no automated update.
- Location: If your building is not listed, or has changed names, please put in a ticket. You need to enter a room as well, but that field is just text. Please use the “Cyberspace” location for “Cloud” computers.
- Make and Model: Please don’t just use “Unknown” when your Make doesn’t exist, just put in a help ticket. Users can create new Models, but it would be best not to create too many.
- Operating System: You can have multiple operating systems, and users can create new ones, but please check to see if one that’s close enough exists.
- Comments: Just plain text, and lots of it!
- Expiration date: Nothing will happen on this date, use it for your records and searching.

Data fields of NetDB that aren't just text

- Node Type: Most users can only make regular nodes, Templates, or Virtual nodes. Select your non-regular node with a checkbox.
  - A "Template" node type lets you make many nodes with similar information easily, without having to use up a real IP number.
  - A "Virtual" node is identical to a regular node, but you can now search on that node type if that's useful for your organization.
- Administrator: This must be either one or more SUNet IDs or, preferably, an Admin Team. NetDB treats this as plain text, but many other departments or campus systems make use of the Administrator field. Admin teams are very nice!
- User: This must be a SUNet ID, but is not a required field. If the user doesn't have a SUNet ID, you can add appropriate information into the comments field.

The “Group” Field

- “Group” grants most of the power in NetDB.
- A user has to share a Group with a record in order to be able to modify it. Every user is a member of the Group “Stanford,” so the Stanford Group can be used to “pass” records around (but remove it when you’re done!)
- A user has to share a Group with an IP range to assign IP numbers in that range.
- A user has to share a Group with a domain to assign names in that domain. The Group “Stanford” can assign names in the “Stanford.EDU” domain, so most users never notice this.

The “State” field

- The State field is relatively new. Anything other than “Good” may have restrictions.
- “Dubious” is added by ISO based on their policy compliance. You have to work with ISO, not UIT Networking, to figure out why and fix it.
- “Unknown” is usually an incomplete self-reg record. It’s best to just delete it and force the host to go through SNSR again.
- “Stale” is used to clean up old nodes. A “Stale” node will be deleted automatically after 90 days, giving conscientious administrators plenty of time to search quarterly for “Stale” nodes and fix them as needed.
- “Unknown,” “Dubious,” and “Vile” nodes will not receive DHCP requests. If you need to temporarily kick a user off and get their attention, setting their node to “Vile” will cause DHCP to fail and, presumably, get them to contact their administrators.

Custom Fields

- There are four custom fields in the form of Label:Value
- You can use them any way you want, but should probably be consistent in your labels so as to make searching easier.
- The NetDB firewall automation tool scrubs information from the custom fields.
- SNSR uses special Template nodes with lots of information in the custom fields to work.

Command-Line NetDB and RMI

- There is a CLI version of NetDB that can make some tasks much easier and faster, especially bulk changes. See http://netdb-cli.stanford.edu
- Searches are limited, and there are a few functions that still have to be done via the web. You can use a complex search from the web version to get a list of hosts, and then apply changes to that list through the CLI.
- Using the "clone" feature, with some preparation, you could create hundreds of records in minutes.
- If you need more capabilities, create your own CLI using the Java RMI! See http://netdb-rmi.stanford.edu

DNS at Stanford

- DNS at Stanford has become more complicated with off-campus hosting services.
- 3rd (and 4th and 5th) level domains used to be prohibited, but a few years ago the office of the President and Provost changed their mind, so you can have any *.Stanford.EDU domain you want, it's just more work for you.
- Domains other than Stanford.EDU can be hosted by Stanford, but they must meet the teaching and research needs of the University, so you can't just buy one and expect it to be hosted by Stanford. It's generally much cheaper to have your custom domains hosted off campus.

DNS: Internal vs. External

- Internal users should use Ice and Iron as their DNS servers. DHCP will hand these out. Ice and Iron use anycast DNS to front for up to 5 other servers, so don't worry about a server going offline.
- External DNS servers will get their information from Atalante, Avallone, Argus and 3 servers from "dnsmadeeasy.com."
- Any name or alias with the .sunet domain will only be seen internally, and will not be sent to the external DNS servers. Use .sunet for non-routed IP numbers to be clean.
- A name in the .nodomain domain won't be sent to any DNS servers, but will be searchable in NetDB. Use .nodomain for complicated Advanced nodes with lots of interface names.

How does DNS work?

- Computers need IP numbers, not names. Humans like names.
- The name that points to an IP number (or numbers) is called an A record (or AAAA for IPv6). A is for "Address record"
- Aliases (CNAMES) point a name to another name. Eventually, your computer still needs a number. C is for "Canonical" and it means the CNAME record points to the canonical record.
- A computer will always ask its own DNS servers for all DNS lookups, no matter the domain. If those DNS servers don't own that domain, they'll find out who does it and pass the DNS request along and then send back the reply to the initial requestor.
    - All of your DNS conversations are with your own DNS servers.

- If you are told that the name you're inquiring about is an Alias, your DNS servers begin a new DNS lookup (which could point to another Alias) and repeat until they finally get an A name and a real IP number.
  - You can't have "cascading" Aliases within NetDB, but you could have a [name].stanford.edu alias pointing to a non-Stanford name, which is itself an alias of another name.
- MX information has to be associated with the A record, it can't be associated with an Alias/CNAME. This can be problematic for redirecting Stanford names to off-campus providers.
- A name can have multiple IP numbers associated with it. You'll be handed them in random order by the DNS servers and your OS will usually pick the first. At Stanford, if one is of those IP numbers on your local subnet, it will be handed out first.
- An IP number can have multiple names, but that's less likely to be a problem.

Domains in NetDB

- In order to create names within a domain (even Stanford.EDU), that domain has to be defined in NetDB.
- For every domain entry, “Create names in” defines which groups can create names within that domain (or subdomain).
- “Use as name” defines which groups can use that domain as a node name. i.e: cs.stanford.edu is both a node name and a domain, and the group “Computer Sci” has the rights to use that domain as a node name.
- Domains that don’t belong to Stanford but still need to be used for redirecting a Stanford alias to an off-site A record (name) still need to be defined in NetDB. There are a lot of *.amazonaws.com sub-domains in NetDB just for this purpose.

3 DNS examples

- Let's look at the complete process for DNS for anything.stanford.edu, gift.stanford.edu and next.stanford.edu (as of November 2017).

- All of these are aliases, but the process is different for each.

- We will look at the process for an off-campus user (using Comcast DNS as an example), it's similar for on-campus.

- The Comcast DNS servers are 75.75.75.75 (cdns01.comcast.net) and 75.75.76.76 (cdns02.comcast.net), referred to as "cdns01 and 02" in subsequent slides.

DNS process for anything.stanford.edu

- Our fictional Comcast user will send a DNS request to either cdns01 or 02 servers asking "Could you give me the IP for anything.stanford.edu?" It will also ask "Could you give me the IPv6 information for anything.stanford.edu?" if the host uses IPv6. Your host will likely choose to only use one DNS server.
- cdns01 or cdns02, recognizing that they don't own Stanford.EDU, will look up the DNS servers for Stanford and get a list of 6 of them. (dig +noall +answer -t ns stanford.edu will get you that same list)
- cdns01 or cdns02 will ask all 6 "What is the IP for anything.stanford.edu"? (Twice, v4 and v6) All 6 will reply "anything.stanford.edu is an alias(CNAME) for networking.stanford.edu" (Again, twice)
- Cdns01/02 will now check to see if they know the DNS servers for Stanford.EDU (they likely didn't forget them) and ask all 6 "What is the IP number for networking.stanford.edu?" and "What is the IPv6…?" All 6 will reply "171.64.20.23" in the first response and "2607:f6d0:0:a13f::f0ad" in the second response.
- Cdns01/02 will tell you "anything.stanford.edu is an alias for networking.stanford.edu. networking.stanford.edu has the IP number of 171.64.20.23" and another response that says "anything.stanford.edu is an alias for networking.stanford.edu. networking.stanford.edu has the IPv6 address of 2607:f6d0:0:a13f::f0ad"

DNS process for gift.stanford.edu

- Our Comcast user will ask their cdns01/02 “what is the IP number for gift.stanford.edu?” They may ask again for IPv6 (we’ll skip IPv6 for the next two examples)
- Cdns01/02 will find out the name servers for Stanford (same list of 6) and ask all of them. They’ll get the reply that “gift.stanford.edu is an alias for dc-63166-17256102.us-east-1.elb.amazonaws.com”
- Cdns01/02 will look up the name servers for us-east-1.elb.amazonaws.com (there are 4 of them) and ask all four “what is the IP number of dc-63166-17256102…?”
- All four will give two replies: “dc-63166-17256102… has address 107.21.119.220” and “dc-63166-17256102…. has address 54.243.238.45”
- Cdns01/02 will reply to you “gift.stanford.edu is an alias for dc-63166-17256102.us-east-1.elb.amazonaws.com and dc-63166-17256102.us-east-1.elb.amazonaws.com has 107.21.119.220 and 54.243.238.45”
- To make this work in NetDB, we had to create 3 domains and a node entry with the name of dc-63166-17256102.us-east-1.elb.amazonaws.com and many aliases, including “gift.”

DNS process for next.stanford.edu

- Our Comcast user will ask their cdns01/02 “what is the IP number for next.stanford.edu?” They may ask again for IPv6.
- Cdns01/02 will look up the name servers for Stanford and ask them. They’ll get the reply that “next.stanford.edu is an alias for ext.squarespace.COM”
- Cdns01/02 will look up the name servers for squarespace.com (there are 8 of them) and ask all eight “what is the IP number of ext.squarespace.COM?”
- All four will give the reply that “ext.squarespace.com is an alias for ext-cust.squarespace.com”
- Cdns01/02 will remember the 8 DNS servers for squarespace.com and ask them “what is the IP number for ext-cust.squarespace.com?
- All 8 will reply with the four IP numbers (198.185.159.145, 198.49.23.144, 198.185.159.144 and 198.49.23.145)
- Cdns01/02 will tell you “next.stanford.edu is an alias for ext.squarespace.COM, ext.squarespace.com is an alias for ext-cust.squarespace.com, and ext-cust.squarespace.com has address …”
- Apparently squarespace used to use “ext” but changed their minds to use “ext-cust” for external customers, and we could clean up that record in NetDB, but it still works, so we don’t.

DNS gory details from Wireshark (packet analyzer) for jrbp.stanford.edu

$ host jrbp.stanford.edu
jrbp.stanford.edu is an alias for live-jrbp.pantheonsite.io.
live-jrbp.pantheonsite.io is an alias for fe2.edge.pantheon.io.
fe2.edge.pantheon.io has address 23.185.0.2
fe2.edge.pantheon.io has IPv6 address 2620:12a:8000::2
fe2.edge.pantheon.io has IPv6 address 2620:12a:8001::2

Source	Destination	Protocol	Length	Info
171.66.72.11	171.67.1.234	DNS	77	Standard query 0xa2ba AAAA jrbp.stanford.edu
171.66.72.11	171.67.1.234	DNS	77	Standard query 0x9b76 A jrbp.stanford.edu
171.67.1.234	171.66.72.11	DNS	492	Standard query response 0xa2ba AAAA jrbp.stanford.edu CNAME live-jrbp.pantheonsite.io CNAME fe2.edge.pantheon.io AAAA 26…
171.67.1.234	171.66.72.11	DNS	452	Standard query response 0x9b76 A jrbp.stanford.edu CNAME live-jrbp.pantheonsite.io CNAME fe2.edge.pantheon.io A 23.185.0…

Let's open the web page:

Source	Destination	Protocol	Length	Info
171.66.72.11	23.185.0.2	HTTP	1313	GET / HTTP/1.1
23.185.0.2	171.66.72.11	TCP	66	80 → 54245 [ACK] Seq=1 Ack=1449 Win=32256 Len=0 TSval
23.185.0.2	171.66.72.11	TCP	66	80 → 54245 [ACK] Seq=1 Ack=2696 Win=34816 Len=0 TSval
23.185.0.2	171.66.72.11	TCP	953	80 → 54245 [PSH, ACK] Seq=1 Ack=2696 Win=34816 Len=88
23.185.0.2	171.66.72.11	TCP	1514	80 → 54245 [ACK] Seq=888 Ack=2696 Win=34816 Len=1448
23.185.0.2	171.66.72.11	TCP	1514	80 → 54245 [ACK] Seq=2336 Ack=2696 Win=34816 Len=1448
23.185.0.2	171.66.72.11	TCP	1514	80 → 54245 [ACK] Seq=3784 Ack=2696 Win=34816 Len=1448
23.185.0.2	171.66.72.11	TCP	1514	80 → 54245 [ACK] Seq=5232 Ack=2696 Win=34816 Len=1448
23.185.0.2	171.66.72.11	TCP	1514	80 → 54245 [ACK] Seq=6680 Ack=2696 Win=34816 Len=1448
23.185.0.2	171.66.72.11	TCP	1514	80 → 54245 [ACK] Seq=8128 Ack=2696 Win=34816 Len=1448
23.185.0.2	171.66.72.11	TCP	1147	80 → 54245 [PSH, ACK] Seq=9576 Ack=2696 Win=34816 Len
23.185.0.2	171.66.72.11	HTTP	860	HTTP/1.1 200 OK (text/html)

Notice: The IP number is for pantheon, the HTTP GET still refers to jrbp.stanford.edu, so pantheon's web server knows what you want.

```
▶ Internet Protocol Version 4, Src: 171.66.72.11, Dst: 23.185.0.2
▶ Transmission Control Protocol, Src Port: 54245, Dst Port: 80, Seq: 1449, Ack: 1, Len: 1247
▶ [2 Reassembled TCP Segments (2695 bytes): #207(1448), #208(1247)]
▼ Hypertext Transfer Protocol
  ▶ GET / HTTP/1.1\r\n
    Host: jrbp.stanford.edu\r\n
```

Stanford DNS pointing to an external IP number.

- Sometimes, we can't use an alias (if there's an MX record) or the off-site vendor prefers we point a *.Stanford.edu host directly to an IP number instead of a name.
    - An MX has to point to the A record only, it can't point to an alias. "ASSU.stanford.edu" is an example: they use squarespace for their web hosting and google for their mail forwarding.
- In NetDB, UIT Networking will have to define that IP range and give that range appropriate group permissions. We will give it a useful name, list the location as "Off Campus" and usually put a comment that it's used for IP redirects.
- For example, if you ask your DNS servers "what is the IP for awesome.stanford.edu" they will ask the Stanford DNS servers, which will reply "awesome.stanford.edu has address 52.8.220.191" and that's it. Even though Stanford doesn't route that IP range, the DNS servers don't care.
- If Amazon AWS changes that IP number, this will break until we fix our entry in NetDB.

What could possibly go wrong?

- The "Cloud" is not made up of unicorns and rainbows. It's someone else's computers and someone else's network, and those can fail, and Stanford can't fix them.
- If there's a problem with any of the DNS servers that don't belong to Stanford, there will be a failure of DNS, and Stanford will not be able to fix it.
- The name server response for a domain will have a time to live (48 hours for Stanford), and if that organization changes its name servers, our DNS servers may hold onto old information until that TTL expires. This is very common when organizations move hosting providers and the previous provider has a long TTL.
- The reply for a DNS query will have a TTL (generally fairly short, 30 minutes at Stanford), so if the name or IP number changes, there may be a delay until the correct information propagates out. DNS servers will hold onto information until that information times out to reduce queries.