

Improving DNS Security and Resiliency

Carlos Vicente
NSRC-University of Oregon



UNIVERSITY OF OREGON



Threats to DNS

- Server crashes
- Server compromise
- Denial of service attacks
- Amplification attacks
- Cache poisoning
- Targeted host attacks using zone information
- More
 - <http://www.dnssec.net/dns-threats>

DoS attacks

- Saturating the target machine with external requests, such that it cannot respond to legitimate traffic
- When your DNS servers are the target of a Denial of Service attack:
 - Your customers can't resolve other domains
 - The world can't resolve your own domains
 - Might as well not be connected to the Internet

Amplification Attacks

- Also known as “Reflection Attacks”
- DNS servers being used as tools in the attack
 - Sending responses to queries whose source addresses have been spoofed
- The actual node that owns the spoofed address is the victim

NEWS TECHNOLOGY

[Home](#) | [US & Canada](#) | [Latin America](#) | [UK](#) | [Africa](#) | [Asia-Pac](#) | [Europe](#) | [Mid-East](#) | [South Asia](#) | [Business](#) | [H](#)

4 November 2010 Last updated at 11:33 ET



Burma hit by massive net attack ahead of election

An ongoing computer attack has knocked Burma off the internet, just days ahead of its first election in 20 years.

The attack started in late October but has grown in the last few days to overwhelm the nation's link to the net, said security firm Arbor Networks.

Reports from Burma say the disruption is ongoing.

The attack, which is believed to have started on 25 October, comes ahead of closely-watched national elections on 7 November.

International observers and foreign journalists are not being allowed into the country to cover the polls.

It will raise suspicions that Burma's military authorities could be trying to restrict the flow of information over the election period.

The ruling generals say the polls will mark a transition to democratic



Huge amounts of traffic easily overwhelmed Burma's links to the net

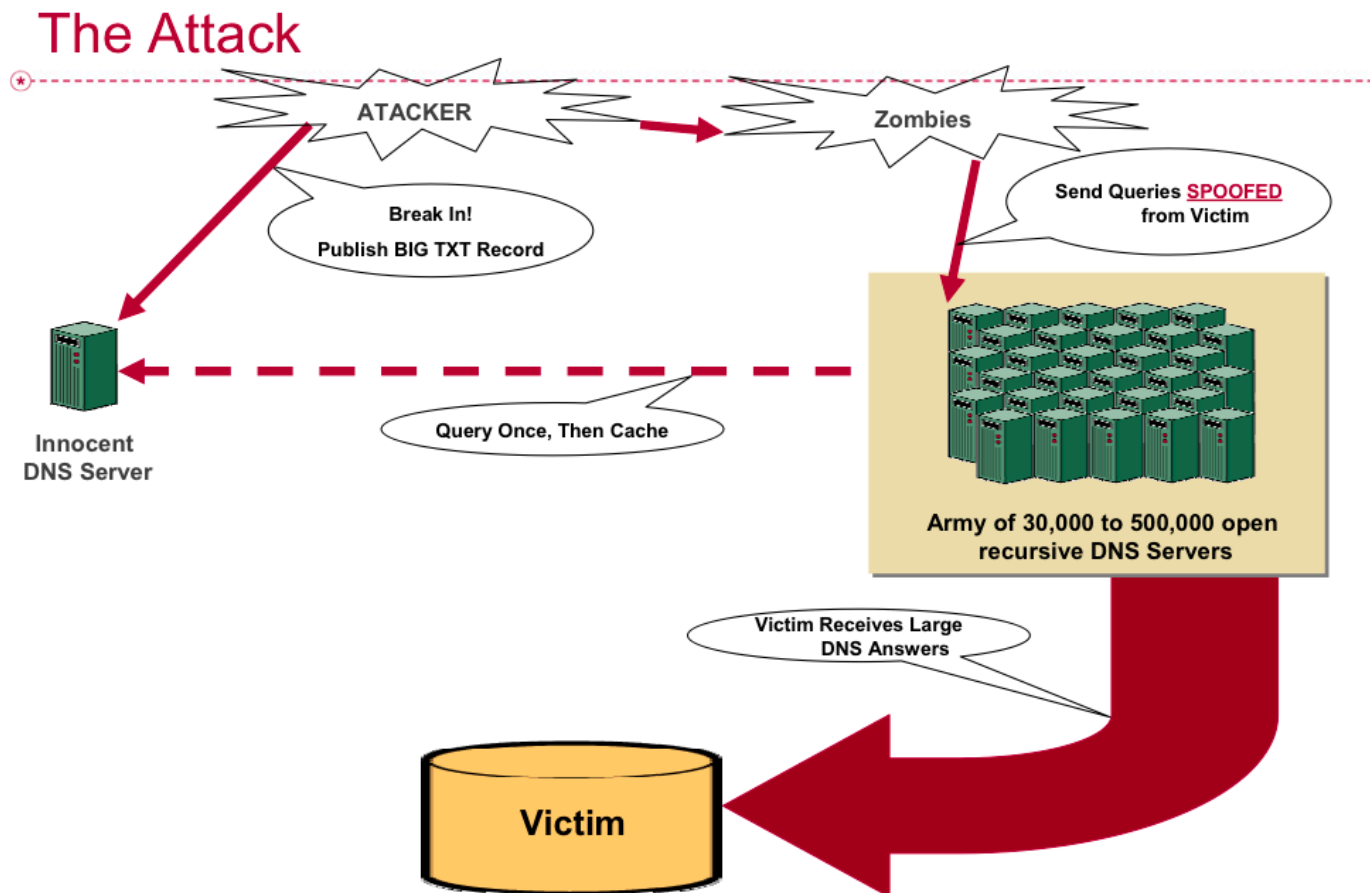
Related stories

[Burma election: Q&A](#)[Spanish police smash huge botnet](#)[Cyber wars in Iran](#)

UNIVERSITY OF OREGON



Amplification Attacks



Source: <http://www.nanog.org/meetings/nanog37/presentations/frank-scalzo.pdf>



UNIVERSITY OF OREGON




Amplification Attacks


- Difficult to protect our users against
 - Impossible to filter thousands of servers
- What we can do is avoid taking part
 - Ingress filtering (IETF BCP 38)
 - Restrict access to recursive DNS servers
 - However, authoritatives can still be used in attacks
- What we should NOT do
 - Limit the size of DNS packets (breaks DNSSEC)

Cache Poisoning

- Attacker tricks a caching server to store a wrong answer
 - www.mybank.com -> 1.2.3.4
 - 1.2.3.4 is the attacker's web server, disguised as your bank!
 - One successful attack affects many (if not all) users


The Memory Experts™
TECH TIPS

Run the Crucial™ System Scanner to find out what type of memory you have installed.



 Print  Retweet  Facebook

Alert 

Cache-poisoning attack snares top Brazilian bank Google AdSense spoofed

By [Dan Goodin in San Francisco](#) • [Get more from this author](#)

Posted in [Crime](#), 22nd April 2009 00:32 GMT

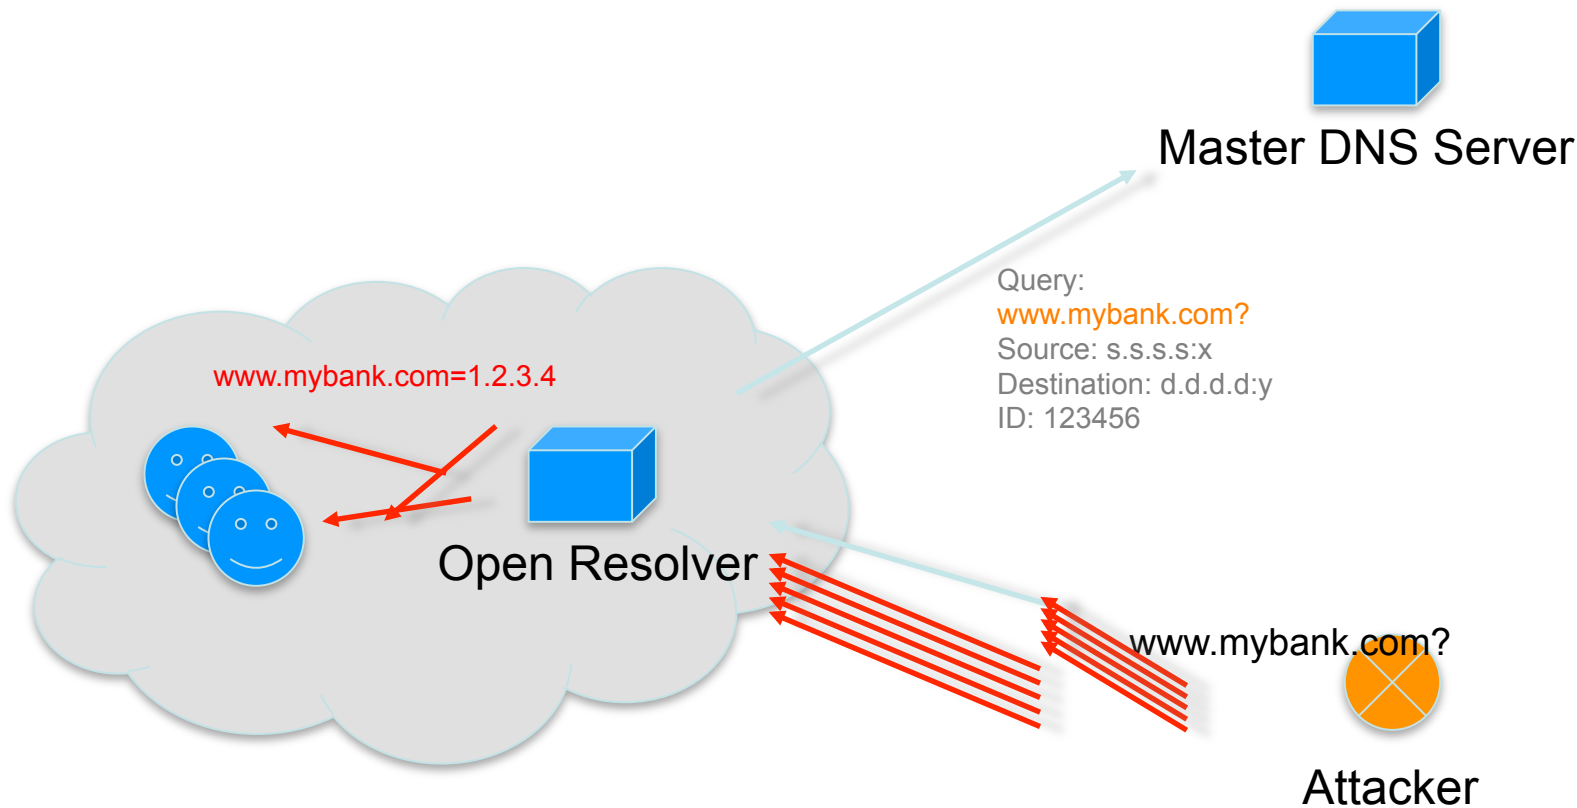
[Free whitepaper – The 10 myths of safe web browsing](#)

One of Brazil's biggest banks has suffered an attack that redirected its customers to fraudulent websites that attempted to steal passwords and install malware, according to an unconfirmed report.

According to [this Google translation](#) of an article penned in Portuguese, the redirection of [Bradesco](#) was the result of what's known as a cache poisoning attack on Brazilian internet service provider [NET Virtua](#).



Cache Poisoning



Reply:
www.mybank.com=1.2.3.4
Source: d.d.d.d:y
Destination: s.s.s.s:x
ID: 123456



UNIVERSITY OF OREGON



IANA — Cross-Pollination Scan

http://recursive.iana.org/

Most Visited ▾ Netflix Everywhere: ...

IANA — Cross-Pollination Scan

iana
Internet Assigned Numbers Authority

[Domains](#) [Numbers](#) [Protocols](#) [About IANA](#)

Cross-Pollination Check

The discovery of a [highly-effective cache poisoning attack](#) that can affect name servers providing recursive name service has made it important that such servers be patched to mitigate against the problem. Furthermore, the risk of cache poisoning for servers that share recursive and authoritative functions can cross-pollinate the authoritative function with incorrect data. This tool is designed to assess the authorities for a given domain and determine whether they provide vulnerable recursive service.

Provide a **domain name** to analyse [Submit Query](#)

Safe.

The servers tested for UOREGON.EDU appear not to be vulnerable to cache poisoning.
Note that not all authoritative name servers could be reached, so there may be additional issues that were not discovered.

Name server	IP Address	Results
ARIZONA.EDU	128.196.128.233	Not recursive
BIGDOG.LSU.EDU	192.16.176.1	Not recursive
	2620:0:da0:f000::1	⚠ Could not scan
DNS.CS.UOREGON.EDU	128.223.6.9	Not recursive
	2001:468:d01:6::80df:609	Not recursive
PHLOEM.UOREGON.EDU	128.223.32.35	Not recursive
	2001:468:d01:20::80df:2023	Not recursive
RUMINANT.UOREGON.EDU	128.223.60.22	Not recursive
	2001:468:d01:3c::80df:3c16	Not recursive
SNS-PB.ISC.ORG	192.5.4.1	Not recursive
	2001:500:2e::1	Not recursive



Dangers of zone transfers

- Zone transfers meant to be used to distribute zones among authoritative servers
- Transfers are expensive operations in terms of resources
 - Could be used for DoS attack
- Having your whole zone makes hacker's life easier:
 - No need to scan your address space
 - Better understanding of your network

Authoritative vs. Recursive

Server Function	Information	Target audience
Authoritative	Your domains	The Internet
Recursive	All other domains	Your users



Separation of Duties

- Physically separating authoritative and recursive servers gives you:
 - Easier control
 - Apply restrictions to what the servers can be used for, and by whom
 - Easier troubleshooting
 - Consider what happens when a DNS-hosted customer moves their domain to another provider without telling you

Authoritative – BIND options

```
options {  
    version "9999.9.9";  
    allow-transfer { peers; };  
    blackhole { attackers; };  
    recursion no;  
    allow-query { any; };  
    ...  
};
```



Authoritative – IP filters

- Can't really filter much here
 - Ports udp/53 and tcp/53 should be open to the world.
- Just don't run any other services
 - No web server, mail server, etc.
 - Keep it really simple

Authoritative - Location

- Locate your servers topologically and geographically dispersed
 - Establish a relationship with another operator, or
 - There are companies that provide secondary service
 - Ask for anycast, DNSSEC and IPv6 support!
 - See RFC 2182

Recursive – BIND options

```
options {  
    version "9999.9.9";  
    recursive-clients 5000;  
    allow-transfer { none; };  
    blackhole { attackers; };  
    allow-recursion { customers; };  
    allow-query { customers; };  
    dnssec-enable yes;  
    dnssec-validation yes;  
    ...  
};
```



Recursive – IP filters

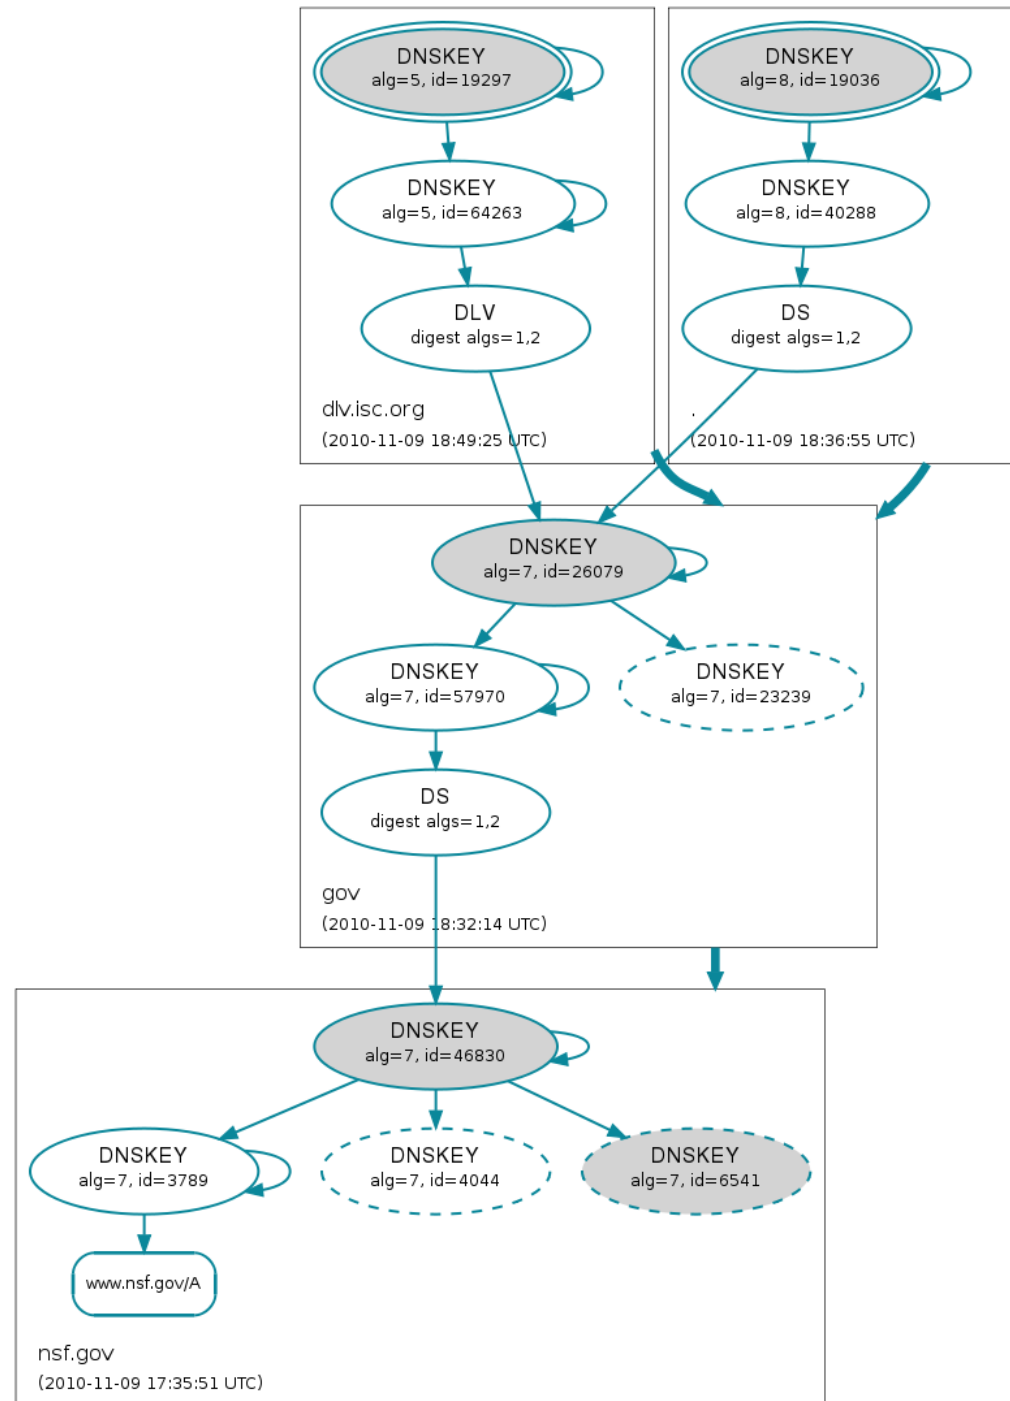
- udp/53 and tcp/53 open only to customers
 - Drop the packets early, don't bother the DNS daemon
 - Remember to filter IPv6 as well if you have v6 connectivity
 - Can be done simply with iptables on Linux.



DNSSEC Validation

- The root is now signed!
- Only true way to avoid cache poisoning
- Started with universities and research organizations, now large ISPs are joining:
 - <http://www.dnssec.comcast.net/>

Source:
dnsviz.net



UNIVERSITY OF OREGON



DNSSEC Validation

```
options {  
    dnssec-enable yes;  
    dnssec-validation yes;  
}
```

```
managed-keys {  
    "." initial-key 257 3 8 "AwEAAagAIKIVZrpC6la7gEzahOR  
+9W29euxhJhVVLOyQbSEW0O8gcCjFFVQUTf6v58fLjwBd0YI0EzrAcQqBGCzh/  
RStloO8g0NfnfL2MTJRkxoXbfDaUeVPQuYEhg37NZWAJQ9VnMVDxP/  
VHL496M/QZxkjf5/  
Efucp2gaDX6RS6CXpoY68LsvPVjR0ZSwzz1apAzvN9dlzEheX7ICJBBtuA6G3LQ  
pzW5hOA2hzCTMjJPJ8LbqF6dsV6DoBQzgul0sGlcGOYI7OyQdXfZ57relSQageu  
+ipAdTTJ25AsRTAoub8ONGcLmqrAmRLKBP1dfwhYB4N7knNnulqQxA  
+Uk1ihz0=";  
};
```



DNSSEC packet size implications

- Responses can easily exceed previous max. of 512 bytes over UDP
- Two solutions:
 - Use EDNS0: The client signals that it can support larger UDP packets
 - Use TCP
- In both cases, need to make sure that the path between your customers and your name servers is capable
 - Especially, check out firewalls

Client failover behavior

- Clients of authoritative servers (other recursive servers)
 - Fail over well using different NS records
- Clients of recursive servers (stub resolvers)
 - Do a very poor job at failing over
 - Users complain immediately
 - Other services break due to timeouts

Anycast


- Routing trick in which the same IP address is announced by multiple routers so that a particular sender reaches the topologically nearest node that responds to that address
- Excellent solution to enhance DNS:
 - Load-balancing
 - Failover
 - DoS attack isolation
 - Cache poisoning isolation

Anycast DNS

- Two approaches
 - Running a routing daemon on the DNS server
 - Zebra, Quagga, etc.
 - Must tie the prefix announcements to DNS service start/stop and... daemon crashes
 - Using IP SLA with Cisco routers
 - Check that the service is operational before injecting prefix in the routing domain
 - No need to trust your sysadmins injecting routes into your routing domain ;-)
 - Server configuration much simpler




Anycast Topology



UNIVERSITY OF OREGON



Anycast Topology



Anycast DNS – Cisco IP SLA

```
ip sla 1
  dns www.mydomain.com name-server 192.0.2.202
  timeout 500
  frequency 10
ip sla schedule 1 life forever start-time now

track 1 ip sla 1
ip route 192.0.2.100 255.255.255.255 192.0.2.200 track 1 tag 999

route-map V4-STATIC permit 10
  match tag 999

router isis mynet
  redistribute static ip metric 100 route-map V4-STATIC level-1
```



Anycast – Server Interfaces

eth0 Link encap:Ethernet HWaddr F0:4D:A2:01:65:42
 inet addr:192.0.2.202 Bcast:192.0.2.203 Mask:255.255.255.252

lo Link encap:Local Loopback
 inet addr:127.0.0.1 Mask:255.0.0.0

lo:1 Link encap:Local Loopback
 inet addr:192.0.2.100 Mask:255.255.255.255



Configuration Management

- Keep configurations and zone files under revision control
 - *SVN, Git, CVS*
- Generate, don't edit, zone files
 - *<http://netdot.uoregon.edu>*
 - *<http://www.nictool.com/info/>*
 - *<http://www.debianadmin.com/bind-dns-server-web-interfacefrontend-or-gui-tools.html>*
- Use CM tool to distribute these files, reload services, etc.
 - *Puppet, CFEngine, etc.*
 - Run a syntax check before loading

```
named-checkzone mydomain.com zonefile
```



Diversify OS and DNS software

- Consider running different DNS software (Bind, Unbound, NSD, etc.) on different OSs
 - Saves you from total disaster when you hit a bug, but...
 - Makes configuration management a lot more challenging

Periodic zone checks

- Periodically run checks for
 - Consistent, missing or bad data
 - Catching common misconfigurations
 - RFC 1912
- Check out dnscheck
 - <https://github.com/dotse/dnscheck>

Watch those logs

- Use a tool to analyze your DNS logs and alarm on important messages
 - Swatch, Tenshi, etc.

Monitoring Availability – Nagios


- Use the `check_dns` to make sure that the server is actually resolving
 - Don't just ping the server
- You can also use this to make sure that very important A records are there:
 - `www`, `smtp`, `imap`,...
- Make sure that your alarms will work despite DNS being down!

Monitoring Availability - Nagios

Service 'DNS' On Host 'ns1'

01-01-2010 00:00:00 to 11-07-2010 21:08:40
Duration: 310d 21h 8m 40s

[Availability report completed in 0 min 16 sec]

Service State Breakdowns:

State	Type / Reason	Time	% Total Time	% Known Time
OK	Unscheduled	90d 22h 8m 40s	29.247%	100.000%
	Scheduled	0d 0h 0m 0s	0.000%	0.000%
	Total	90d 22h 8m 40s	29.247%	100.000%
WARNING	Unscheduled	0d 0h 0m 0s	0.000%	0.000%
	Scheduled	0d 0h 0m 0s	0.000%	0.000%
	Total	0d 0h 0m 0s	0.000%	0.000%
UNKNOWN	Unscheduled	0d 0h 0m 0s	0.000%	0.000%
	Scheduled	0d 0h 0m 0s	0.000%	0.000%
	Total	0d 0h 0m 0s	0.000%	0.000%
CRITICAL	Unscheduled	0d 0h 0m 0s	0.000%	0.000%
	Scheduled	0d 0h 0m 0s	0.000%	0.000%
	Total	0d 0h 0m 0s	0.000%	0.000%
Undetermined	Nagios Not Running	0d 0h 0m 0s	0.000%	
	Insufficient Data	219d 23h 0m 0s	70.753%	
	Total	219d 23h 0m 0s	70.753%	
All	Total	310d 21h 8m 40s	100.000%	100.000%



UNIVERSITY OF OREGON

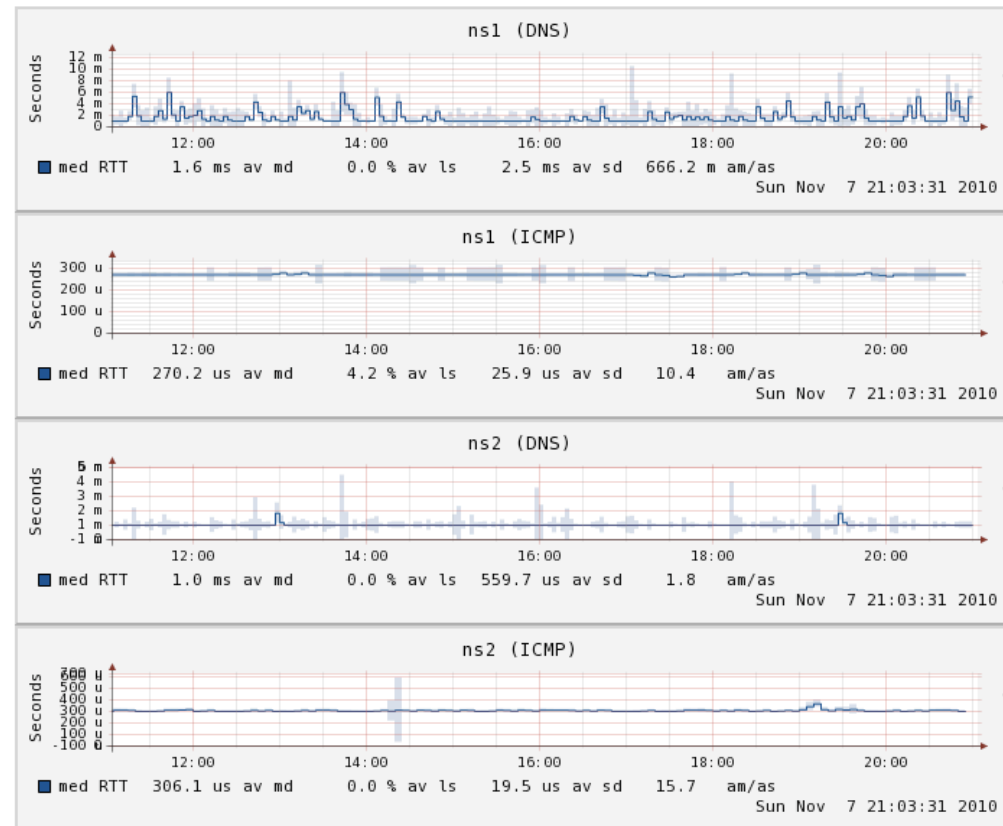


Monitoring Delay


- Important to look at both
 - Network delay
 - DNS service delay

Monitoring Delay - Smokeping

Recursive



DNS Statistics - DSC



UNIVERSITY OF OREGON



Questions?

- Thank you



UNIVERSITY OF OREGON

