DNS Best Practices

Mike Jager
Network Startup Resource Center
mike@nsrc.org

Authoritative server

- Gives answers for specific zones
  - "authoritative" for these zones
- Only respond to queries for these zones
- Never ask other DNS servers anything
- A server can be authoritative for >1 zone
- A zone should have >1 authoritative server

Recursive server

- Receives queries from clients
  - CPE, user's PCs, mail servers, etc
- Send queries to authoritative servers
- Follow referrals down from the root servers until an answer is found
- Answer stored in local cache

Authoritative vs recursive

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

Threats to DNS

- Denial of service attacks
- Reflection/amplification attacks
- Cache poisoning
- Information disclosure
- Human error
- Hardware/software failure

DoS attacks

- Saturating the target with 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

DoS attacks

- Your authoritative servers may be attacked
- Mitigate by having multiple servers
  - Well distributed globally
- Anycast a good technique to absorb DoS
- Many commercial anycast services
  - May act as secondary servers for your zones
- Some services available for ccTLDs, etc.

Amplification attacks

- Specific class of “reflection attack”
- DNS servers used as tools in the attack
- Queries with spoofed source addresses sent to DNS servers
- Server replies to the “source” with packet many times larger than the request
- The node legitimately using the spoofed address is the victim

Amplification attacks

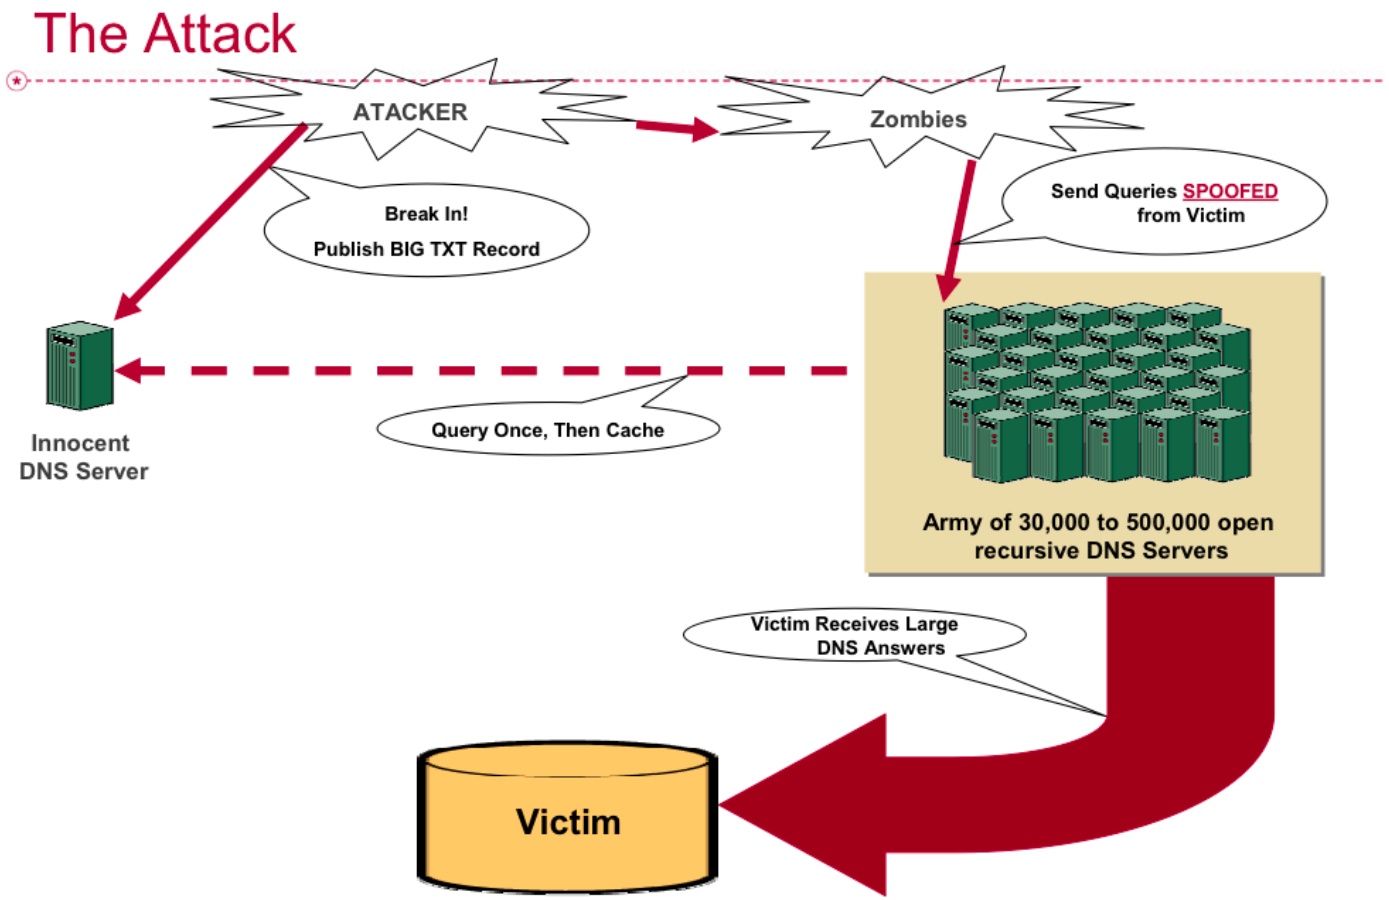

Source: http://www.nanog.org/meetings/nanog37/presentations/frank-scalzo.pdf

Amplification attacks

- Victims see lots of UDP source 53 traffic
- Many different source addresses
- Standard DDoS mitigation technique
- Tempting to limit DNS packets by size
  - But this breaks DNSSEC!
- Often open recursive DNS servers
- Important to not be part of the problem!

Amplification attacks

- Don’t run open recursive servers
  - Drop queries that are not from customers
  - Authoritative servers used in attacks too
- Ensure BCP 38 adherence
  - https://tools.ietf.org/html/bcp38

Cache poisoning

- Attacker fools recursive server into caching an incorrect answer
- www.mybank.com -> 192.0.2.1
  - 192.0.2.1 is under attacker control
  - Looks like your bank, but isn't!
- Successful cache poisoning attack affects many (if not all) users

Cache poisoning

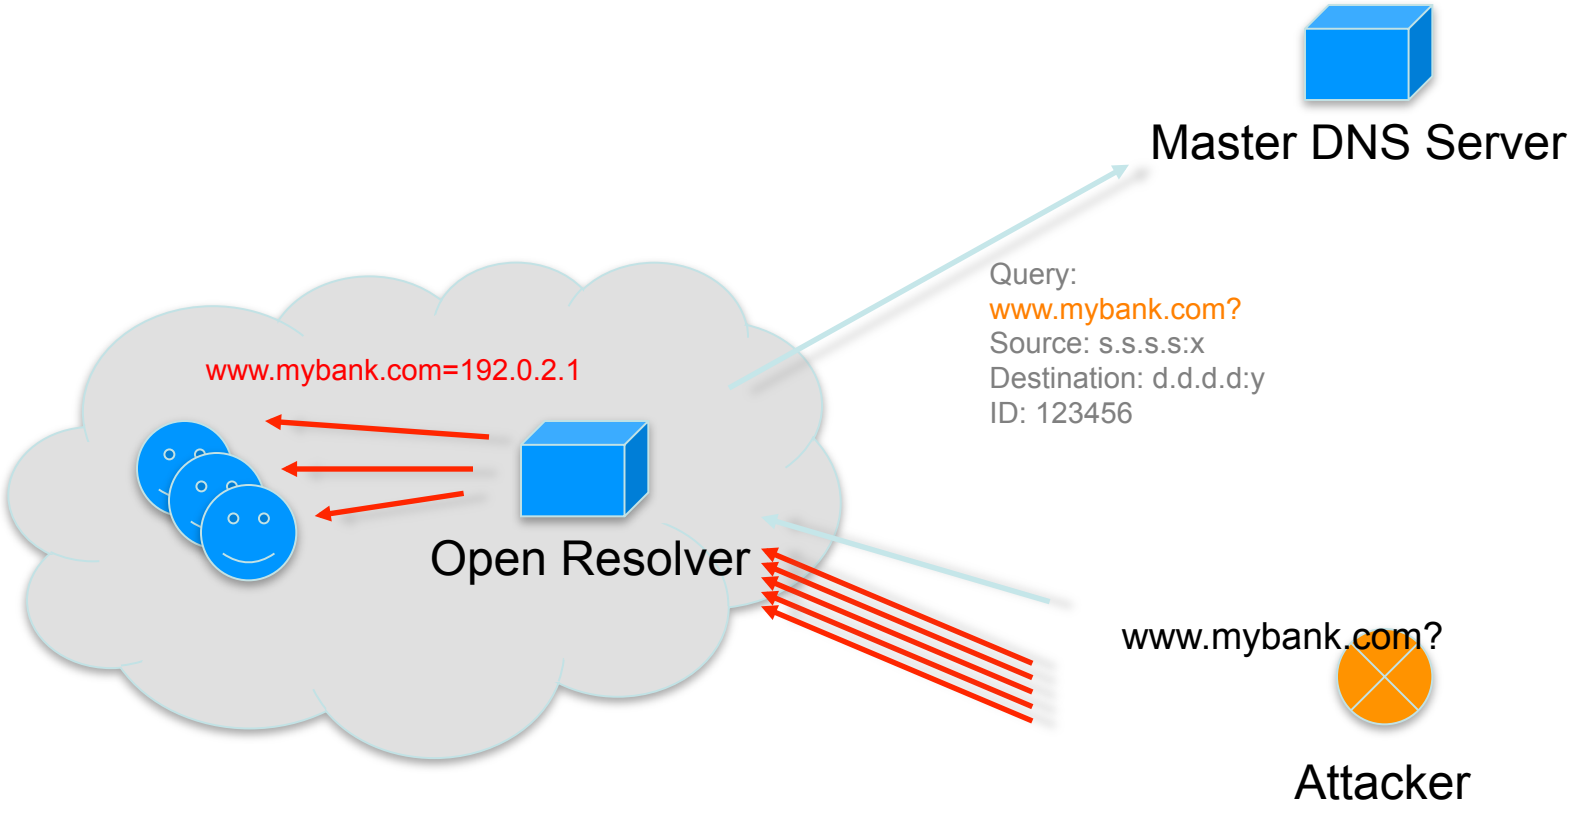

Cache poisoning

- Many tweaks to make poisoning harder
  - Being careful about processing responses
  - Transaction ID randomisation
  - Source port randomisation
    - NAT can undo this
- DNSSEC is the only true way to avoid it

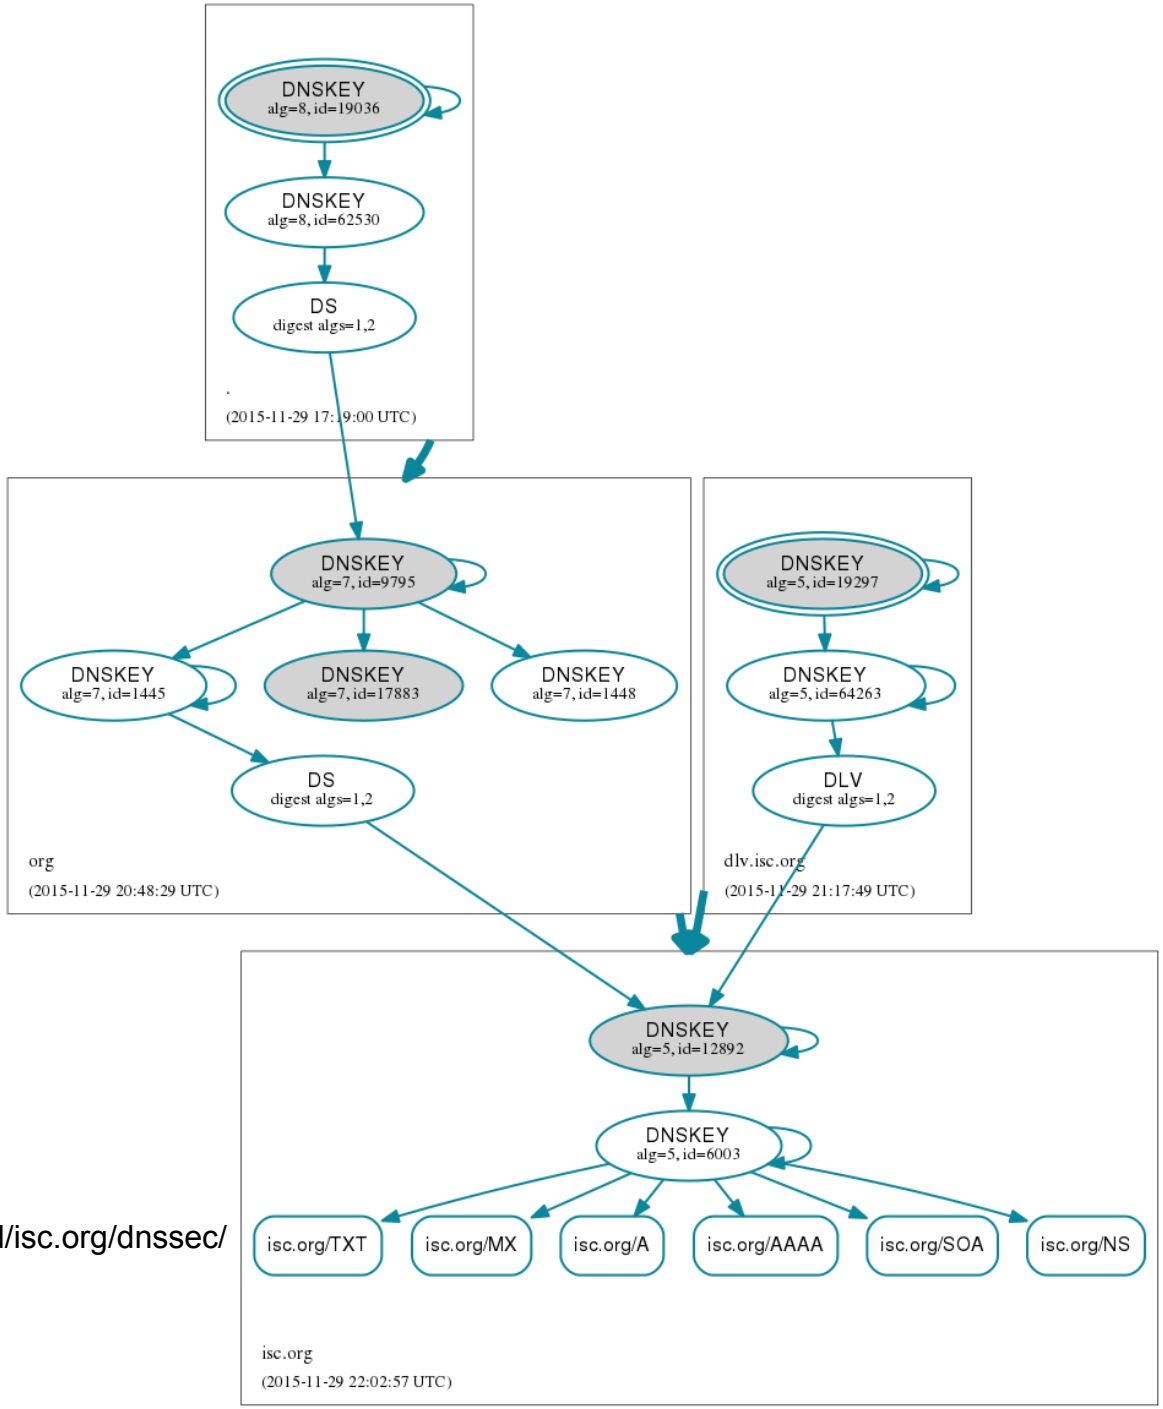

Source: http://dnsviz.net/d/isc.org/dnssec/

Information disclosure

- DNS is clear text
  - DNSSEC provides authentication
  - Not confidentiality
- Zone transfers
  - Allow the entire contents of a zone to be read
  - Easier for an attacker to find targets

Separation of duties

- Authoritative and recursive separated
- Scale each service independently
- Failure of one does not affect the other
- Easier control
  - Restrict what each can be used for by whom
- Easier troubleshooting
  - Not confusing authoritative and cached data

Protecting authoritative servers

- Disable recursion!
- UDP/TCP dest port 53 from everywhere
- No other services on the same servers
- Run multiple authoritative servers
  - RFC 2182
  - Including some outside of your network
  - Trade secondary service with another operator
  - Commercial DNS hosting services

Protecting recursive servers

- Only permit queries from your customers
  - Otherwise you **will** be used for amplification
- Stateless packet filter
  - Permit UDP/TCP dest port 53 from customers
  - On-server firewall (iptables/ipfw)
  - ACL deployed to router/switch
  - Do not keep packet state!

Client failover

- Clients of authoritative servers
  - Recursive servers
  - Fail over well using different NS records
- Clients of recursive servers
  - Stub resolvers in CPE, PCs, servers, etc
  - Do a very poor job at failing over
  - Users complain immediately
  - Services break due to timeouts

Anycast

- Routing trick
- Same prefix announced from >1 location
- Client reaches “nearest” instance
  - based on network topology
  - BGP path selection
- Works well with short-lived sessions
  - like DNS!

Anycast topology

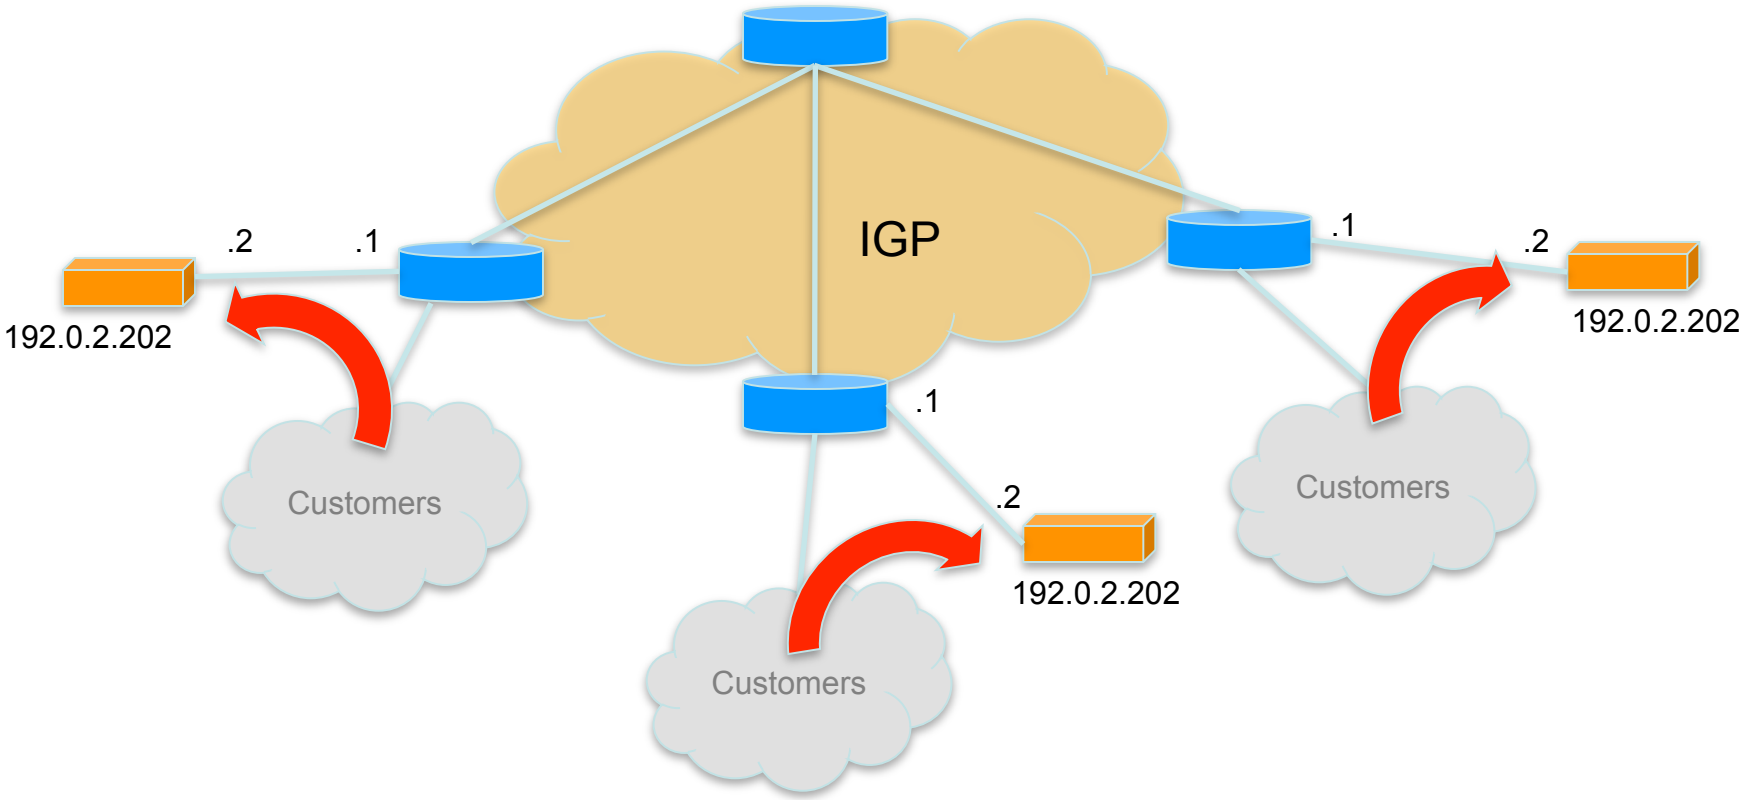

Anycast topology

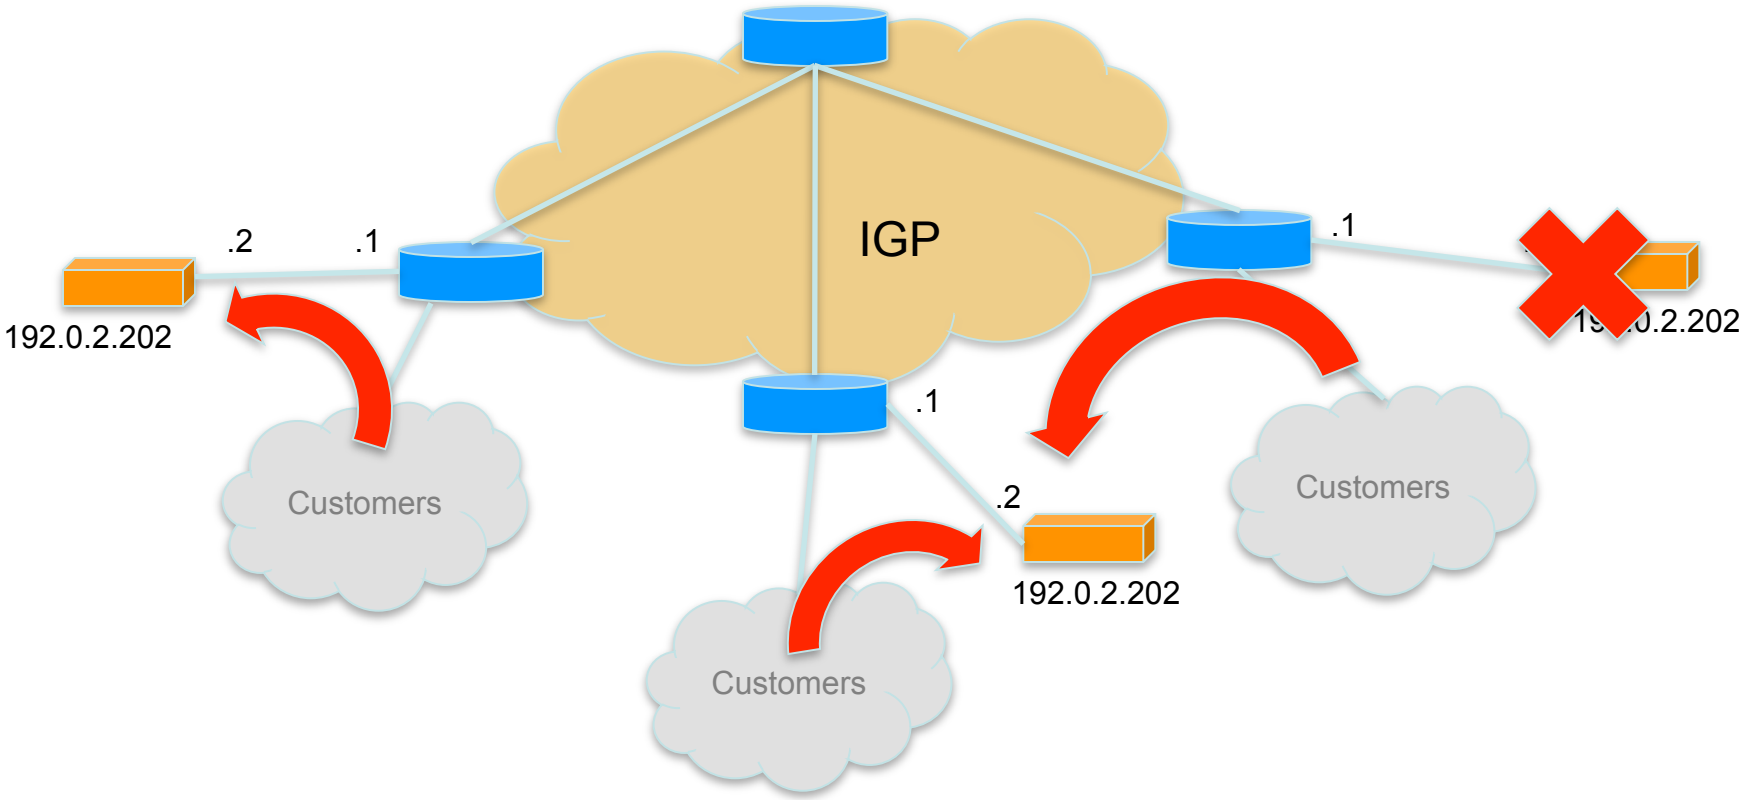

Anycast DNS

- Load balancing
- Failover
- Distributed sinking of DDoS traffic
- Minimise impact of cache poisoning

Anycast DNS

- Run a routing daemon on DNS server
  - BIRD, Quagga, etc
  - Must withdraw prefix if DNS service stops
  - More complex server configuration
- IP SLA monitors DNS service
  - Advertises prefix if service is operational
  - No routing protocol on server
  - More complex router configuration

Diversification

- Different location
- Different network
- Different hardware
- Different OS
- Different DNS software
- Reduced chance of total service failure
- Increased configuration complexity

Configuration management

- Use a tool for configuration/zones revisions
  - Git, Subversion, etc
- Use a tool to generate zone files
  - Avoid error-prone manual edits
  - Netdot
- Use a tool to deploy configuration files
  - Ansible, Puppet, Chef, etc
- Use a tool...

Sanity checking

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

Monitoring availability

- Don’t just ping the DNS server address!
- Check that server responds to queries
- Check that important records still exist
  - www, smtp, imap, etc
- DNS failure may impact alarming
  - Out-of-band alerting mechanism required

Monitoring delay

- Network delay
- DNS service response time
- Smokeping can do both

Monitoring delay

Recursive

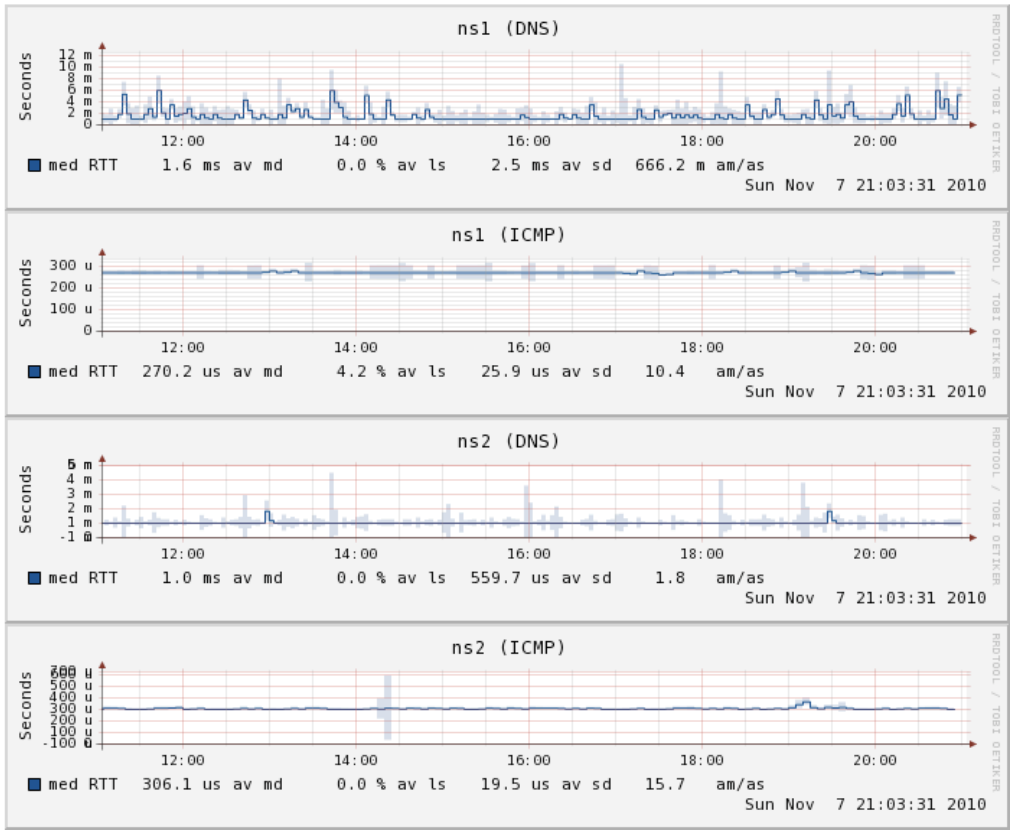

Query statistics - dsc

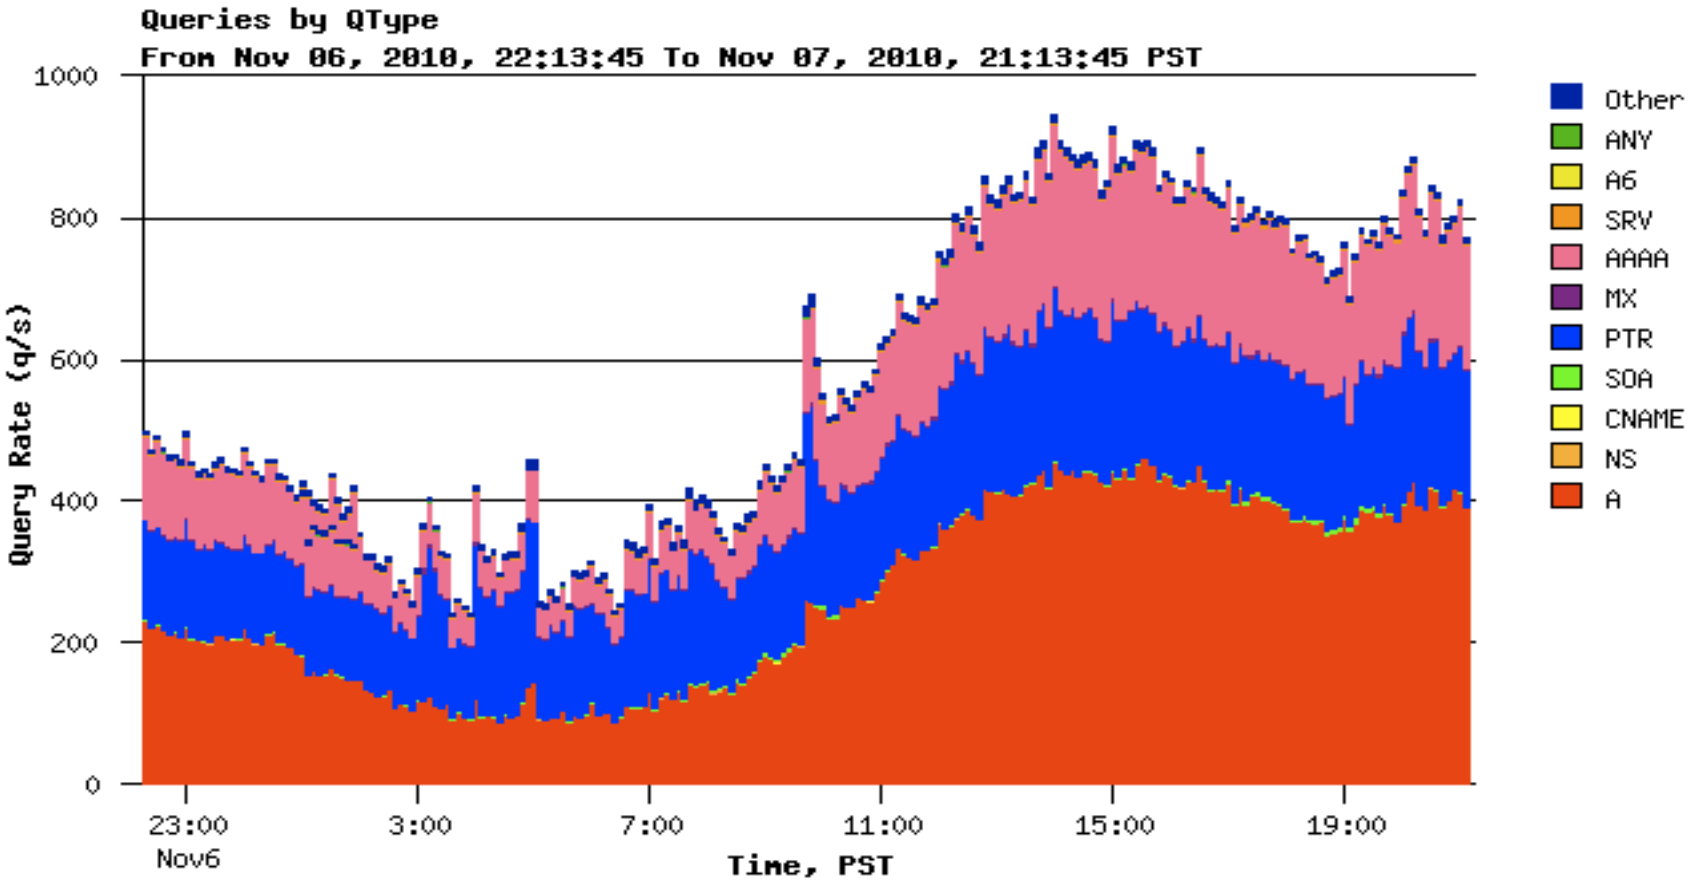

Monitoring logs

- Use a tool to analyse DNS logs
  - Simple Log Watcher
  - tenshi
- Alarm on important messages
  - zone syntax errors
  - zone transfer errors
  - DNSSEC validation errors

Questions?