Exceptional service in the national interest

Sandia National Laboratories

"From 0 to 6 in…": IPv6 deployment experiences at Sandia National Laboratories

Casey Deccio, Rich Gay

Sandia National Laboratories

Winter 2012 ESCC/Internet2 Joint Techs

January 23, 2012, Baton Rouge, LA

U.S. DEPARTMENT OF ENERGY

National Nuclear Security Administration

Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

Outline

Sandia National Laboratories

- Addressing
- Architecture
- Host/application observations
- DNSViz – testing DNS consistency with IPv6

IPv6 Addressing

Sandia National Laboratories

- 52-bit network environments
  - 4-bit environment identifier
- 64-bit prefixes – all non-PTP subnets
  - 12-bit network identifier
  - Based on VLAN value
- 126-bit prefixes – PTP subnets
  - Last network in an environment (all network bits set) reserved for PTP addressing within that environment
  - PTP subnets assigned sequentially from reserved network
  - PTPs – 126 vs. 127 vs. 64

IPv6 Host Addressing

- Static addressing (e.g., servers)
  - 64-bit host identifier uses decimal-encoded value of IPv4 last octet, padded by zeroes, for facilitated identification
    - 192.0.2.**13** => 2001:db8:1234:abcd::**13**

- Dynamic addressing
  - DHCPv6
    - Enterprise IP management
    - IPv6 DNS server advertisement
    - DDNS updates (to forward/reverse DNS zones) via DHCP server
  - 96-bit prefix from each subnet network used for dynamic pool
    - 32 bits for non-temporary address assignment
    - Doesn't conflict with static addressing scheme

Outline

Sandia National Laboratories

- Addressing
- Architecture
- Host/application observations
- DNSViz – testing DNS consistency with IPv6

Architecture

Sandia National Laboratories

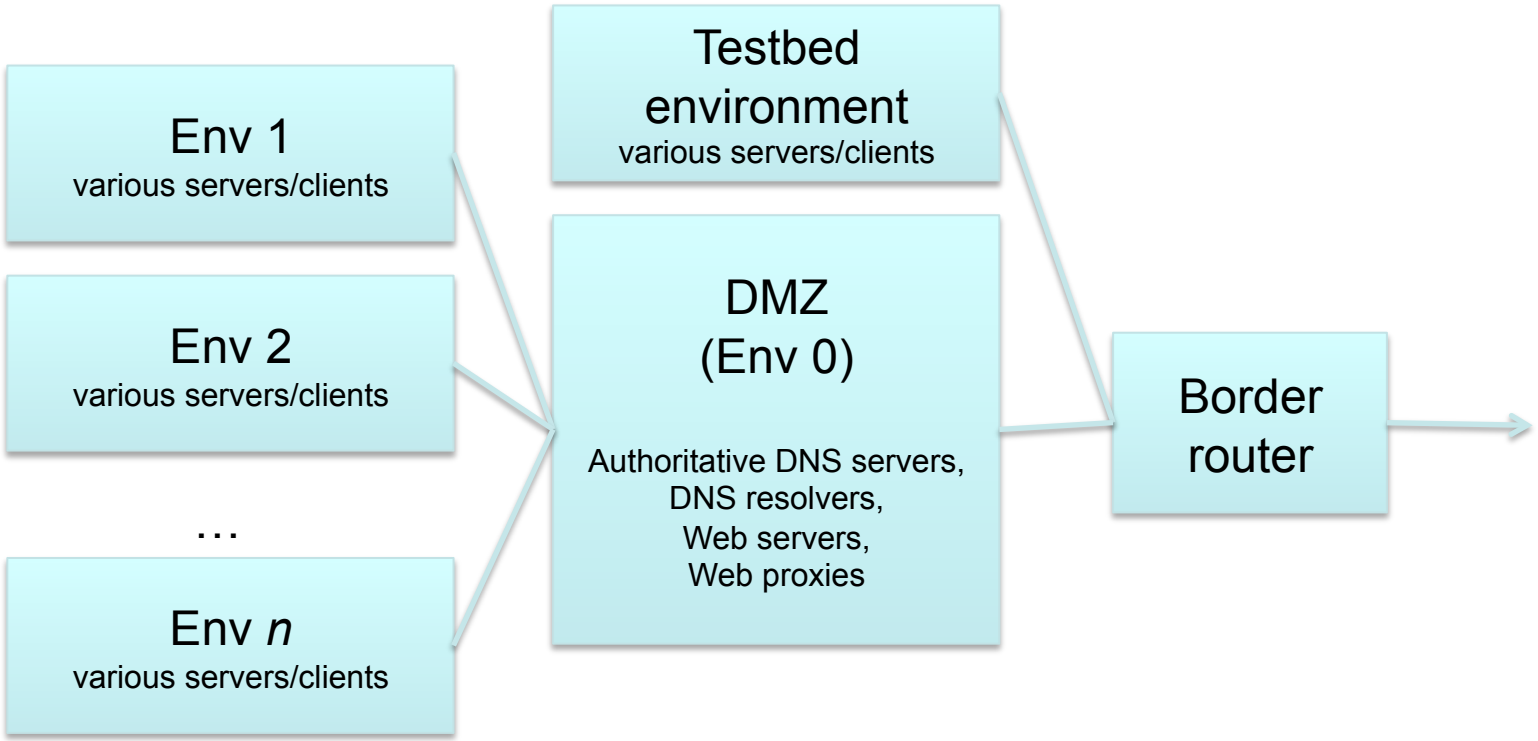

Deployment Plan

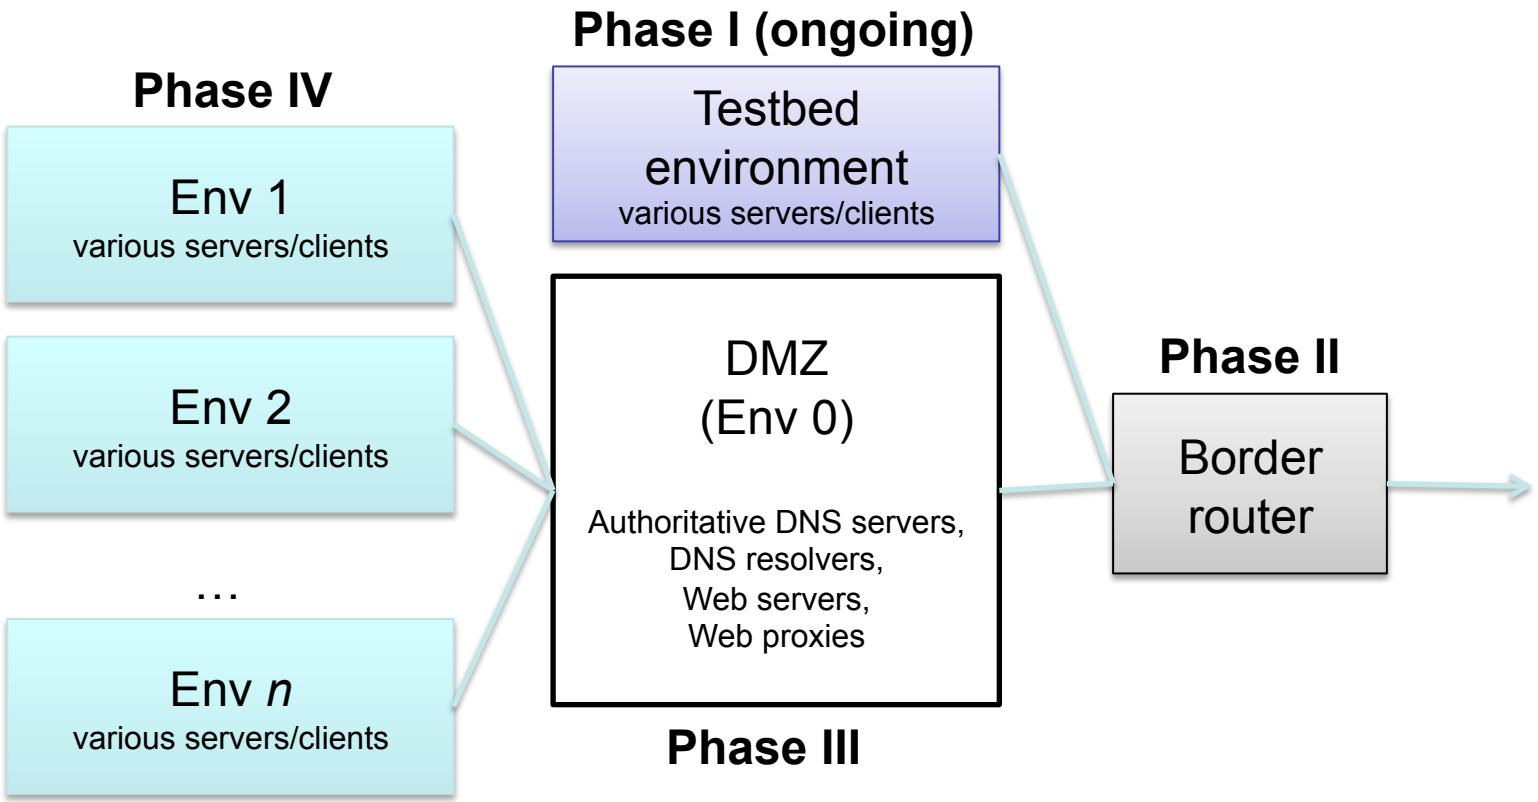

IPv6 Security

Sandia National Laboratories

- Usual stuff blocked at the border:
  - Protocol 41
  - Teredo
  - Unnecessary ICMPv6
  - Reserved IPv6 addresses
  - Obsolete IPv6 addresses
- Observations
  - Zero IPv6 host scans
  - Zero port scans of live IPv6 hosts

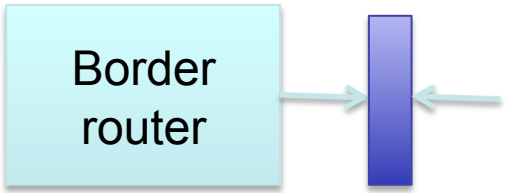

Firewall Woes

Sandia National Laboratories

- Application-level Gateway (ALG)
  - Some implementations have problems handling fragmented packets
- RHEL5
  - Linux kernel 2.6.18 doesn't filter properly; unable to re-assemble packet fragments
- RHEL6 (and RHEL5?)
  - Default firewall rules don't allow return DHCPv6 responses

- Fragmentation
  - Mostly affects DNS/ DNSSEC
  - Use large DNS responses to test IPv6 connectivity

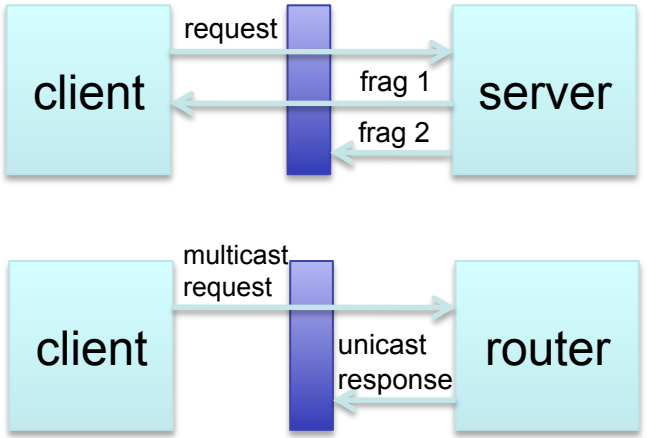

DHCPv6 RA Configuration

Sandia National Laboratories

- Router Advertisements (RAs) for DHCPv6
  - Managed (**M**) address configuration bit **set**
    - Indicates that addresses are available via DHCPv6
  - Autonomous (**A**) address-configuration bit **cleared** from prefix
    - Indicates that prefix cannot be used for stateless address configuration
- Results from initial testing
  - WinXP doesn't support DHCPv6
  - Mac OS X pre-Lion doesn't support DHCPv6
  - Tested OSes respect cleared A-bit on prefix (i.e., don't use SLAAC)

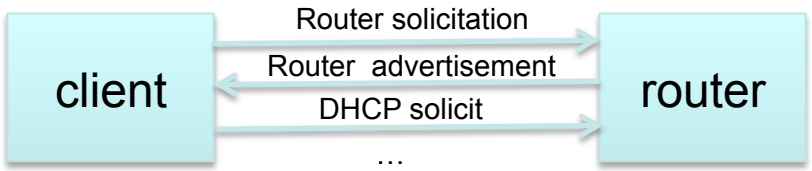

Challenges with ISC dhcp for DHCPv6

Sandia National Laboratories

- Features not yet fully developed as for IPv4
- "host" statements use DHCP Unique Identifier (DUID), rather than MAC address
  - IPAM must have client DIUDs to populate hosts for dhcpd6.conf
  - ISC dhcp 4.2 includes retrofit that allows old-style MACs for dhcpd6.conf hosts
  - RHEL6 ships with ISC dhcp 4.1, but will backport functionality
- "pool" statements unusable within subnet6
  - Allow/deny clients, based on existence of "host" statement
- DDNS
  - updates can't update both A and AAAA records
    - Current update algorithm doesn't allow updating AAAA when A already exists for name
    - Reverse doesn't get updated either
    - Work-arounds exist, but aren't clean
  - Only Windows 7 clients are sending FQDN option (with default settings)

Outline

Sandia National Laboratories

- Addressing
- Architecture
- Host/application observations
- DNSViz – testing DNS consistency with IPv6

Major OSes

Sandia National Laboratories

- Windows 7
  - DHCPv6 works as expected, out of box
- Mac OS X Lion
  - DHCPv6 works as expected, out of box
  - Uses IPv4 DNS servers **before** IPv6 servers
- RHEL6 (NetworkManager)
  - IPv6 must be explicitly enabled on network interface (default: “ignore”)
  - DHCPv6 works as expected
  - Uses IPv4 DNS servers **before** IPv6 servers
- Ubuntu 11.10 (NetworkManager)
  - IPv6 must be explicitly enabled on network interface (default: “ignore”)
  - DHCPv6 requires “priming” – change from “Automatic” to “Automatic, DHCP Only” and back
  - Uses IPv4 DNS servers **before** IPv6 servers

Other IPv6 Applications

Sandia National Laboratories

- BlueCoat Secure Gateway (Web proxy)
  - Doesn't fail over to IPv4 in the case of IPv6 connectivity issues
    - Works well for identifying others' IPv6 issues
    - Requires manually whitelisting troubled domains
- World IPv6 Day
  - June 8, 2011 – 10% HTTP traffic used IPv6
  - Oct 5, 2011 – 3.6% HTTP traffic used IPv6

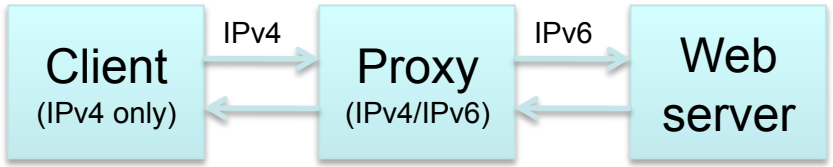

Other Challenges

Sandia National Laboratories

- No mechanism for inserting AAAA glue into .gov
- Monitoring
  - Our current monitoring tools don't fully support IPv6
  - We're setting up Nagios to supplement existing toolset
- Current corporate protection suite for Windows 7 doesn't support IPv6

Outline

Sandia National Laboratories

- Addressing
- Architecture
- Host/application observations
- DNSViz – testing DNS consistency with IPv6

DNS Consistency with IPv6

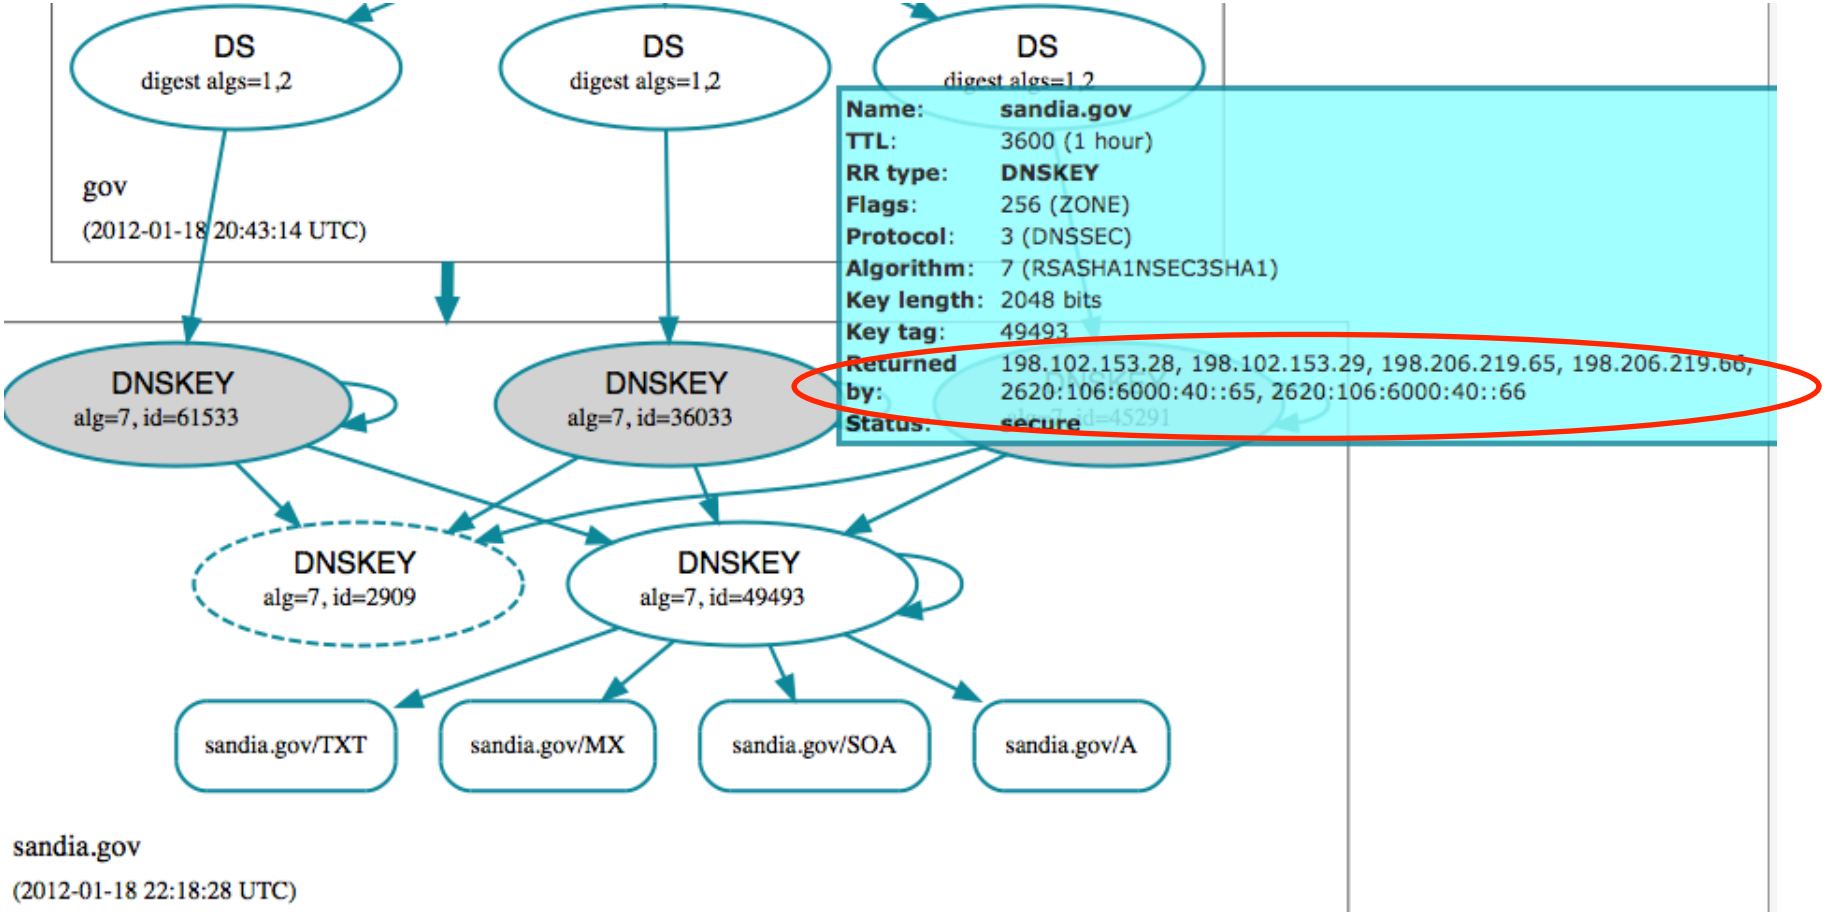

http://dnsviz.net/

DNS Consistency with IPv6

Sandia National Laboratories

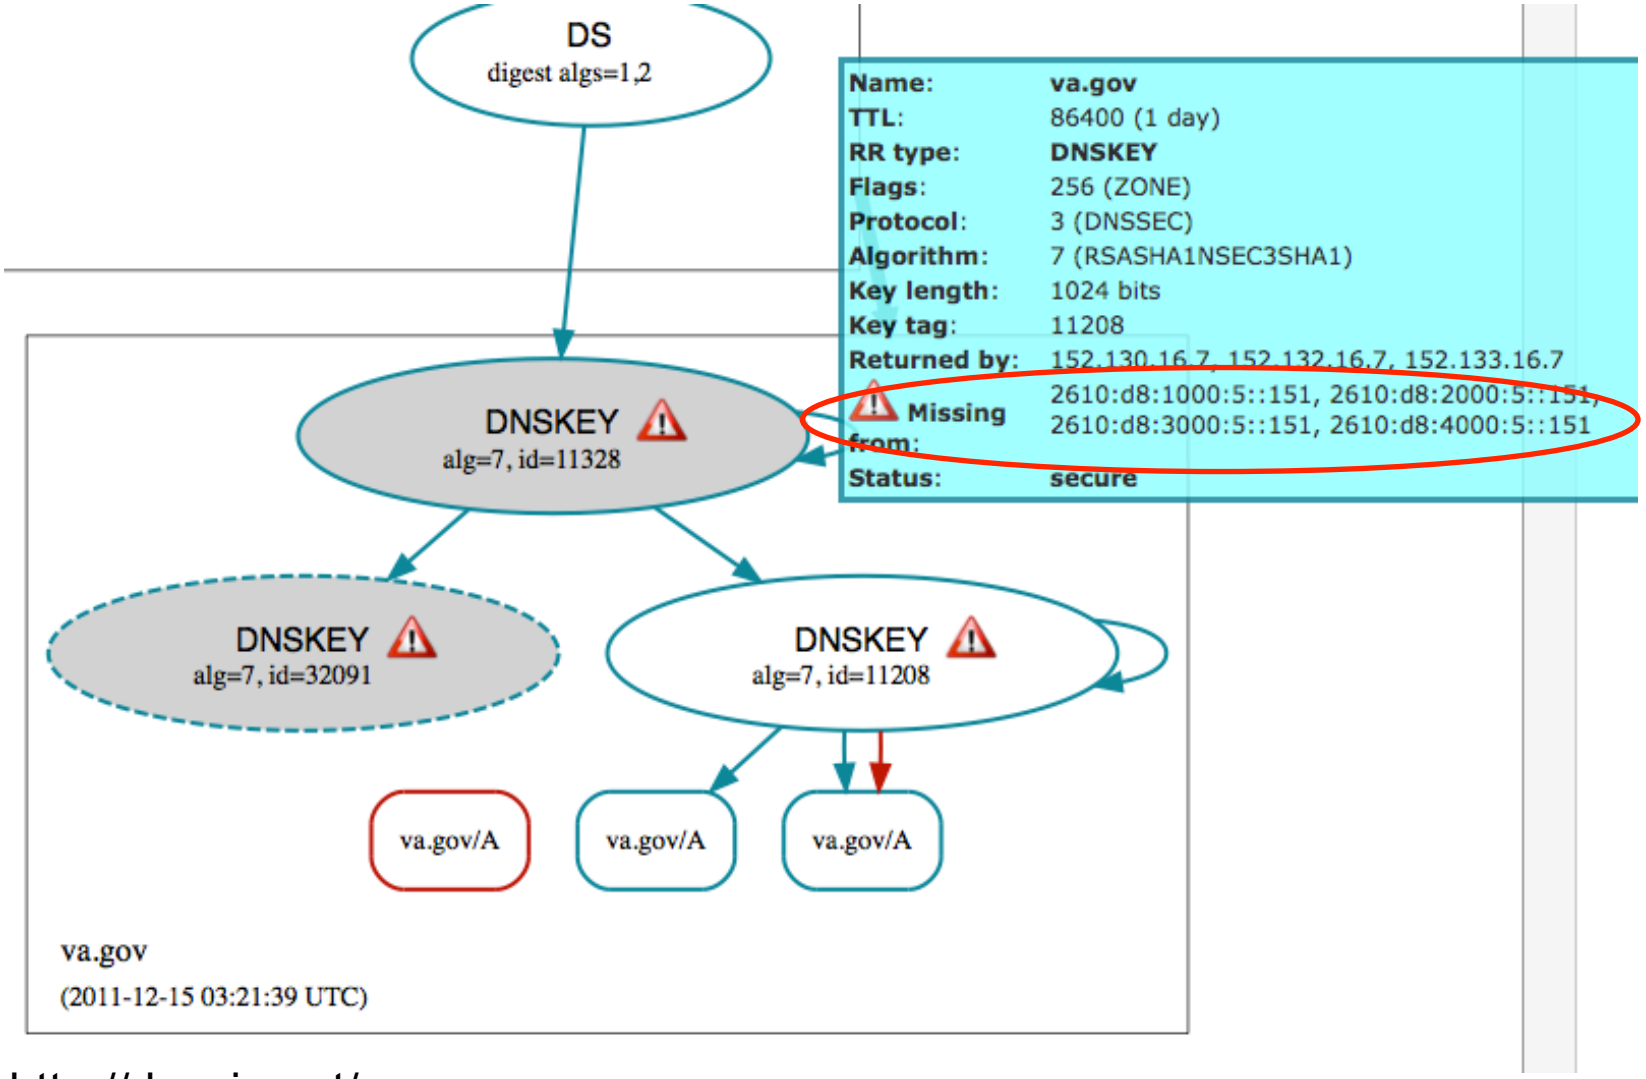

http://dnsviz.net/

DNS Consistency with IPv6

Sandia National Laboratories

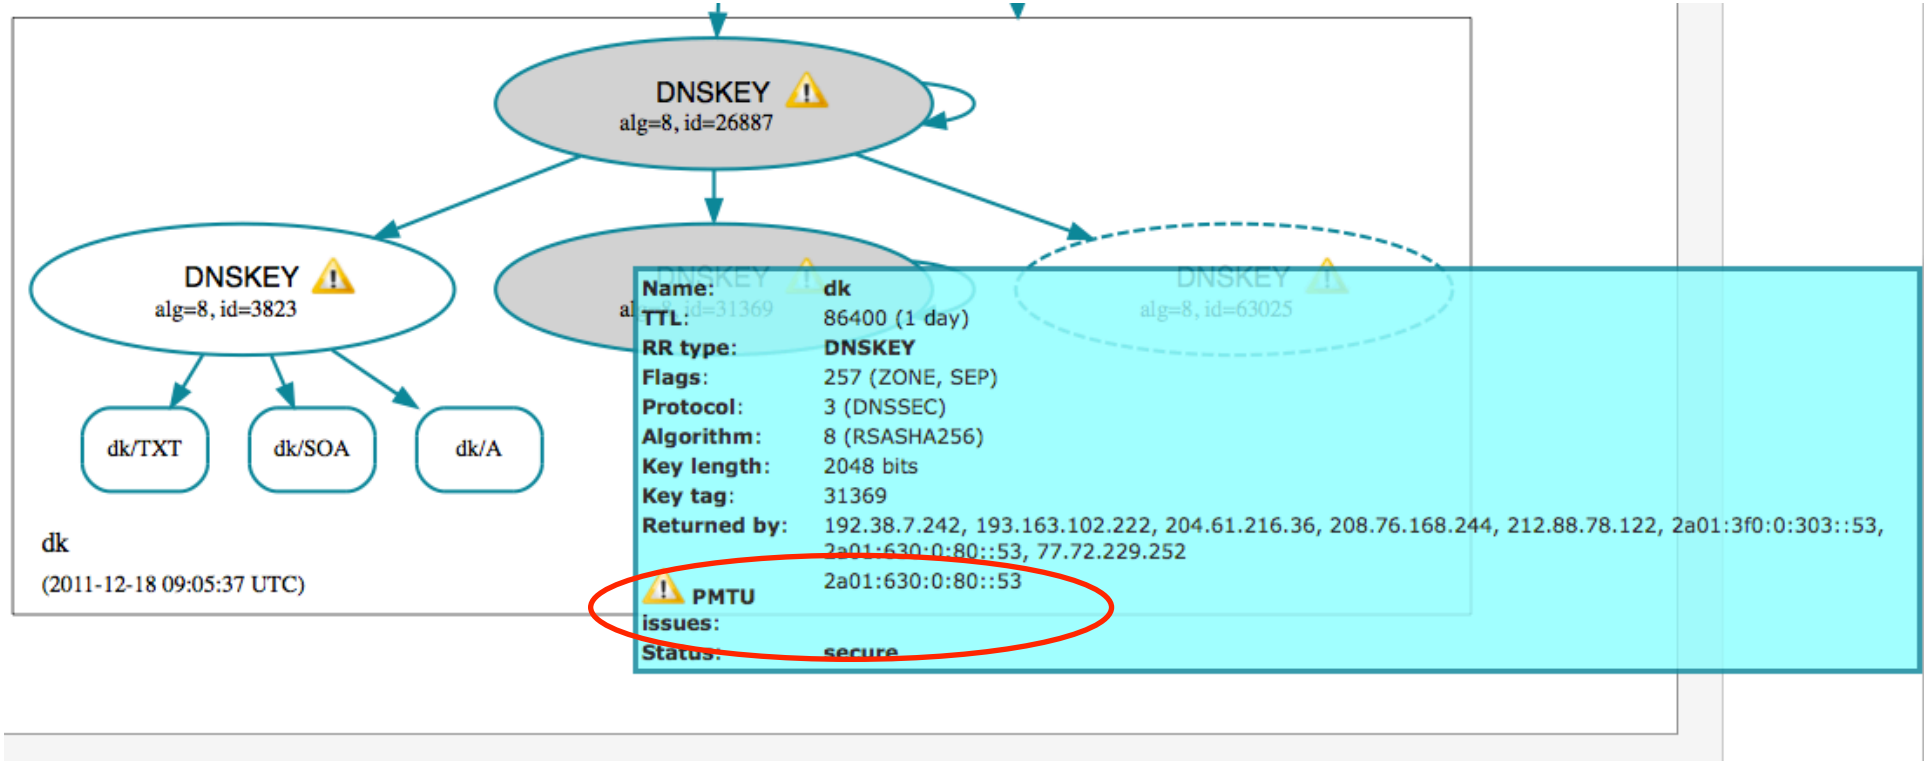

http://dnsviz.net/

Future Plans

Sandia National Laboratories

- Continue testing and deployment to clients
- Enable IPv6 on all Internet-facing servers
- Improved monitoring and IPAM solutions for IPv6

Questions?

Sandia National Laboratories

Other Observations

Sandia National Laboratories

- Safari Web browser
  - Fails over to IPv4 almost immediately after IPv6 attempt
  - Timeout for other browsers/applications 20 – 75 seconds
- OpenSSH_5.2p1 (distributed with Mac OS X Snow Leopard)
  - Fails over from IPv4 to IPv6, even with only link-local address
  - Seems to be fixed in later version