

*Exceptional service in the national interest*



# Deploying IPv6 in the Enterprise

Casey Deccio

Sandia National Laboratories

ARIN XXIX

April 23, 2012, Vancouver, B.C.



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.

# A Few Points on IPv6 Deployment



- Enterprises will typically be working in an existing IPv4 architecture
- Many enterprises will need to largely work with existing IP address configuration management processes
  - Automatic configuration
  - Managed IP addresses allocation
- IPv6 is not expected to be deployed everywhere at once in an enterprise

# Outline

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

\*Addressing and architectural guidelines are based on actual deployment at Sandia, though they are generally consistent with Best Current Operational Practices.

# IPv6 Addressing



- 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 (skip first)
  - PTP prefix length – 126 vs. 127 vs. 64
    - Determined by both network availability and hardware limitations

# IPv6 Host Addressing



- Fixed addressing

- 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

- 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

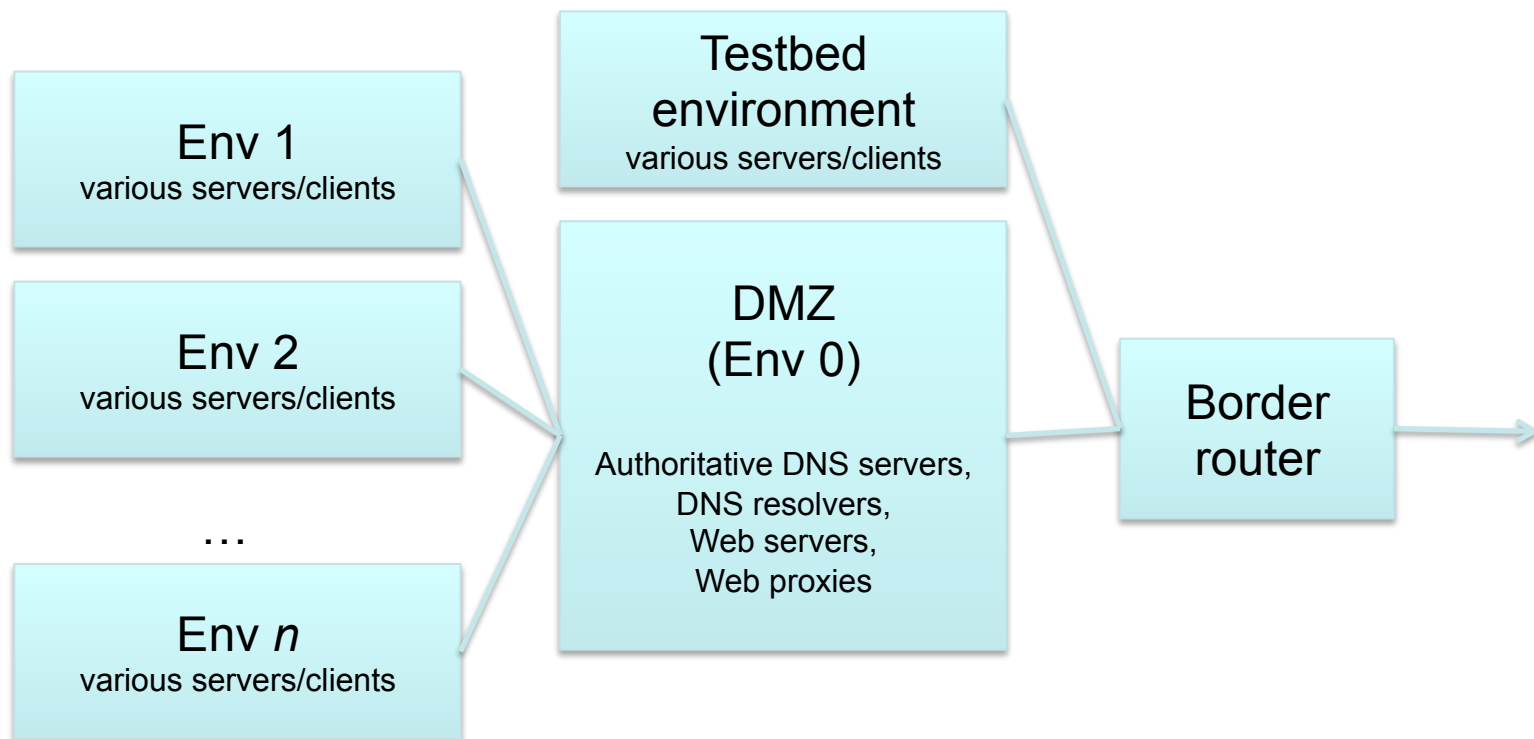
# Address Configuration

- Manual host configuration
  - Fixed addressing only
  - Servers
- Automatic configuration
  - DHCPv6 (no SLAAC)
    - Supports both fixed addressing (pre-assigned addresses) and dynamic addressing (from /96 pools)
    - Enterprise host/IP management
    - IPv6 DNS server advertisement
    - DDNS updates (to forward/reverse DNS zones) via DHCP server

# Outline

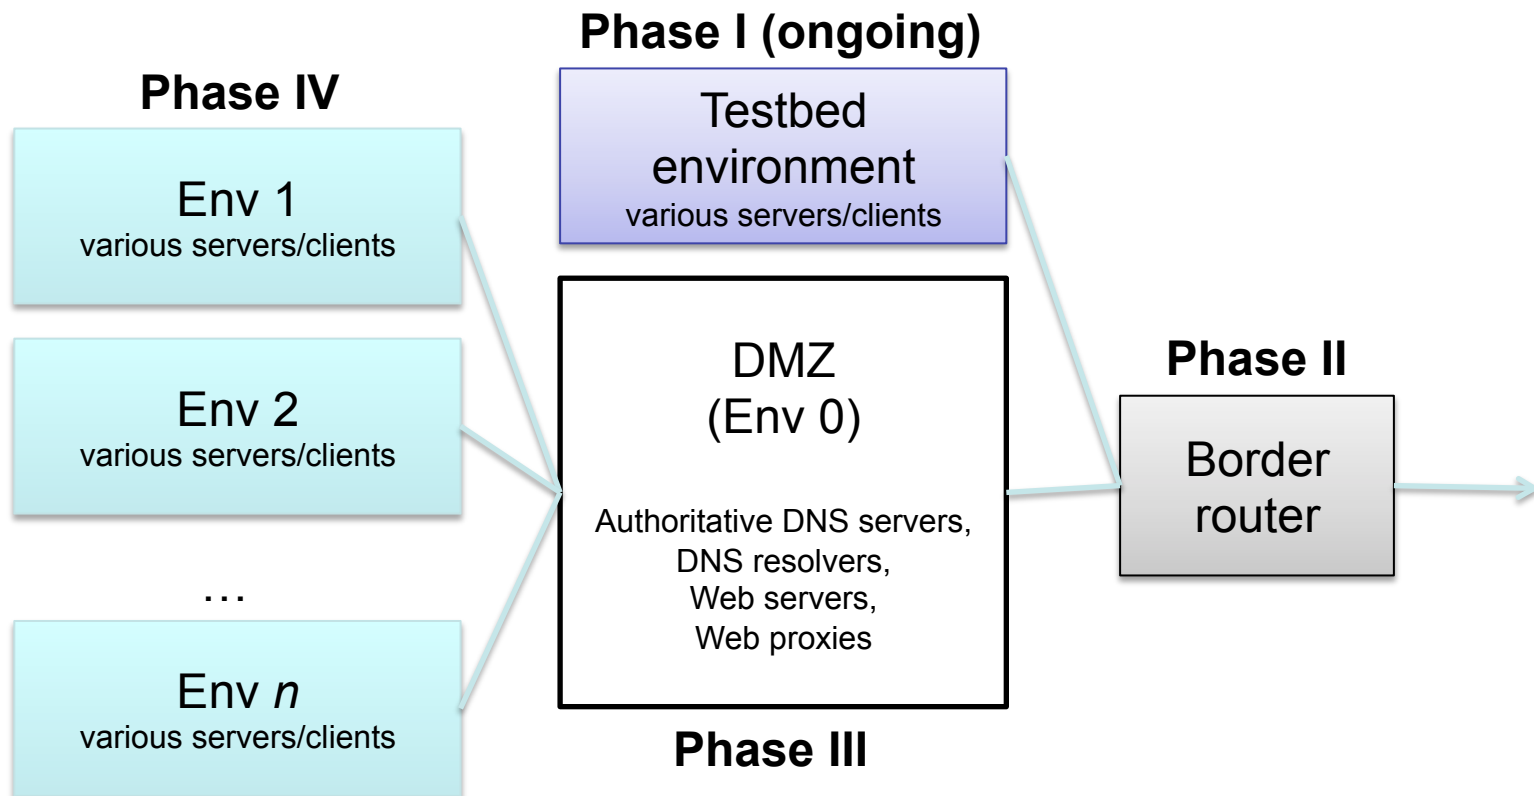
- Addressing
- Architecture
- Host/application observations
- DNSViz – testing DNS consistency with IPv6

# Architecture






# Deployment Plan




# IPv6 Security

- Usual stuff blocked at the border:
  - Protocol 41
  - Teredo
  - Unnecessary ICMPv6
  - Reserved IPv6 addresses
  - Obsolete IPv6 addresses




# Firewall Woes

- 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

- 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

- 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 backported functionality
- “pool” statements unusable within subnet6
  - Unable to 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


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

# Major OSes

- 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

- BlueCoat Secure Gateway (Web proxy)
  - Allows IPv4-only client to access IPv6 Web servers
  - 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
  - Loses its own IPv6 route with bounce of physical interface!
- World IPv6 Day
  - June 8, 2011 – 10% HTTP traffic used IPv6
  - Oct 5, 2011 – 3.6% HTTP traffic used IPv6



```

$ dig +short www.bluecoat.com aaaa
2001:418:9804:111::9
$ curl -I6 http://www.bluecoat.com/
curl: (7) couldn't connect to host
$ curl -I4 http://www.bluecoat.com/
HTTP/1.1 200 OK
  
```





# Other Challenges

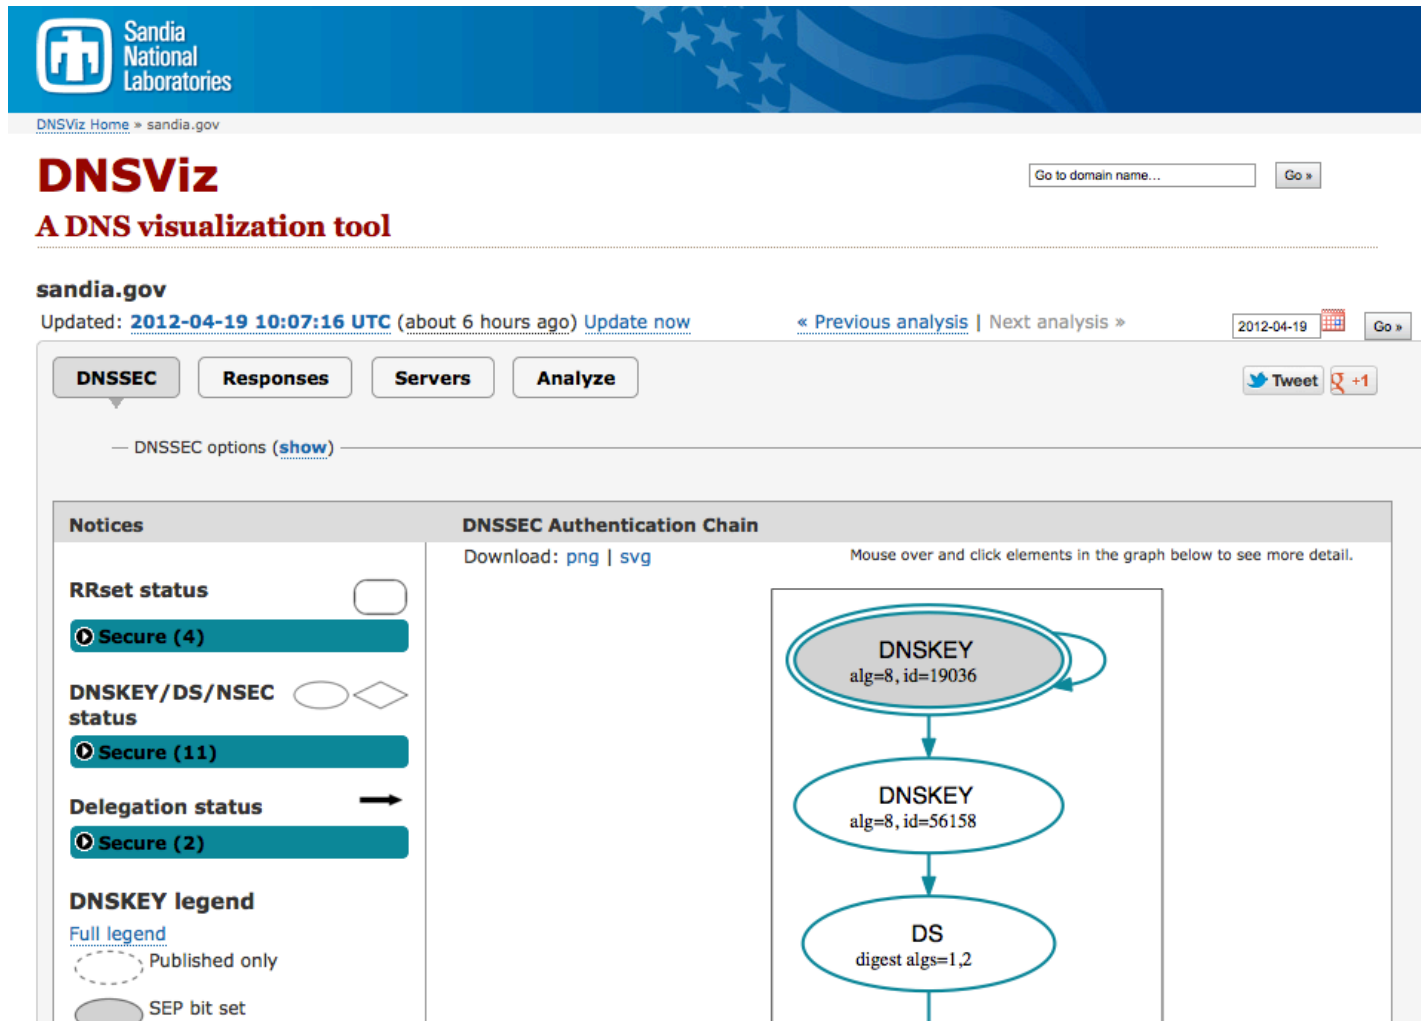
- No mechanism for inserting AAAA glue into .gov
- Monitoring
  - Our current monitoring tools don't fully support IPv6
  - We're setting up Nagios to supplement existing toolset
- Current corporate protection suite for Windows 7 doesn't support IPv6

# Outline

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

# DNS Visualization – and IPv6

<http://dnsviz.net/>



The screenshot displays the DNSviz web application interface. At the top, there is a blue header with the Sandia National Laboratories logo and the text "DNSviz Home » sandia.gov". Below the header, the main title "DNSviz" is prominently displayed in red, followed by the subtitle "A DNS visualization tool". A search bar with the placeholder "Go to domain name..." and a "Go »" button is located to the right of the title.

The main content area shows the analysis for "sandia.gov". It includes a timestamp "Updated: 2012-04-19 10:07:16 UTC (about 6 hours ago) Update now" and navigation links "« Previous analysis | Next analysis »". A date selector shows "2012-04-19" and a "Go »" button. Below this, there are four tabs: "DNSSEC" (selected), "Responses", "Servers", and "Analyze". Social media sharing buttons for "Tweet" and "+1" are also present.

The "DNSSEC options (show)" section is visible. The main visualization area is titled "DNSSEC Authentication Chain" and includes a "Download: png | svg" link and a note "Mouse over and click elements in the graph below to see more detail." The authentication chain is represented by a flowchart with three nodes:

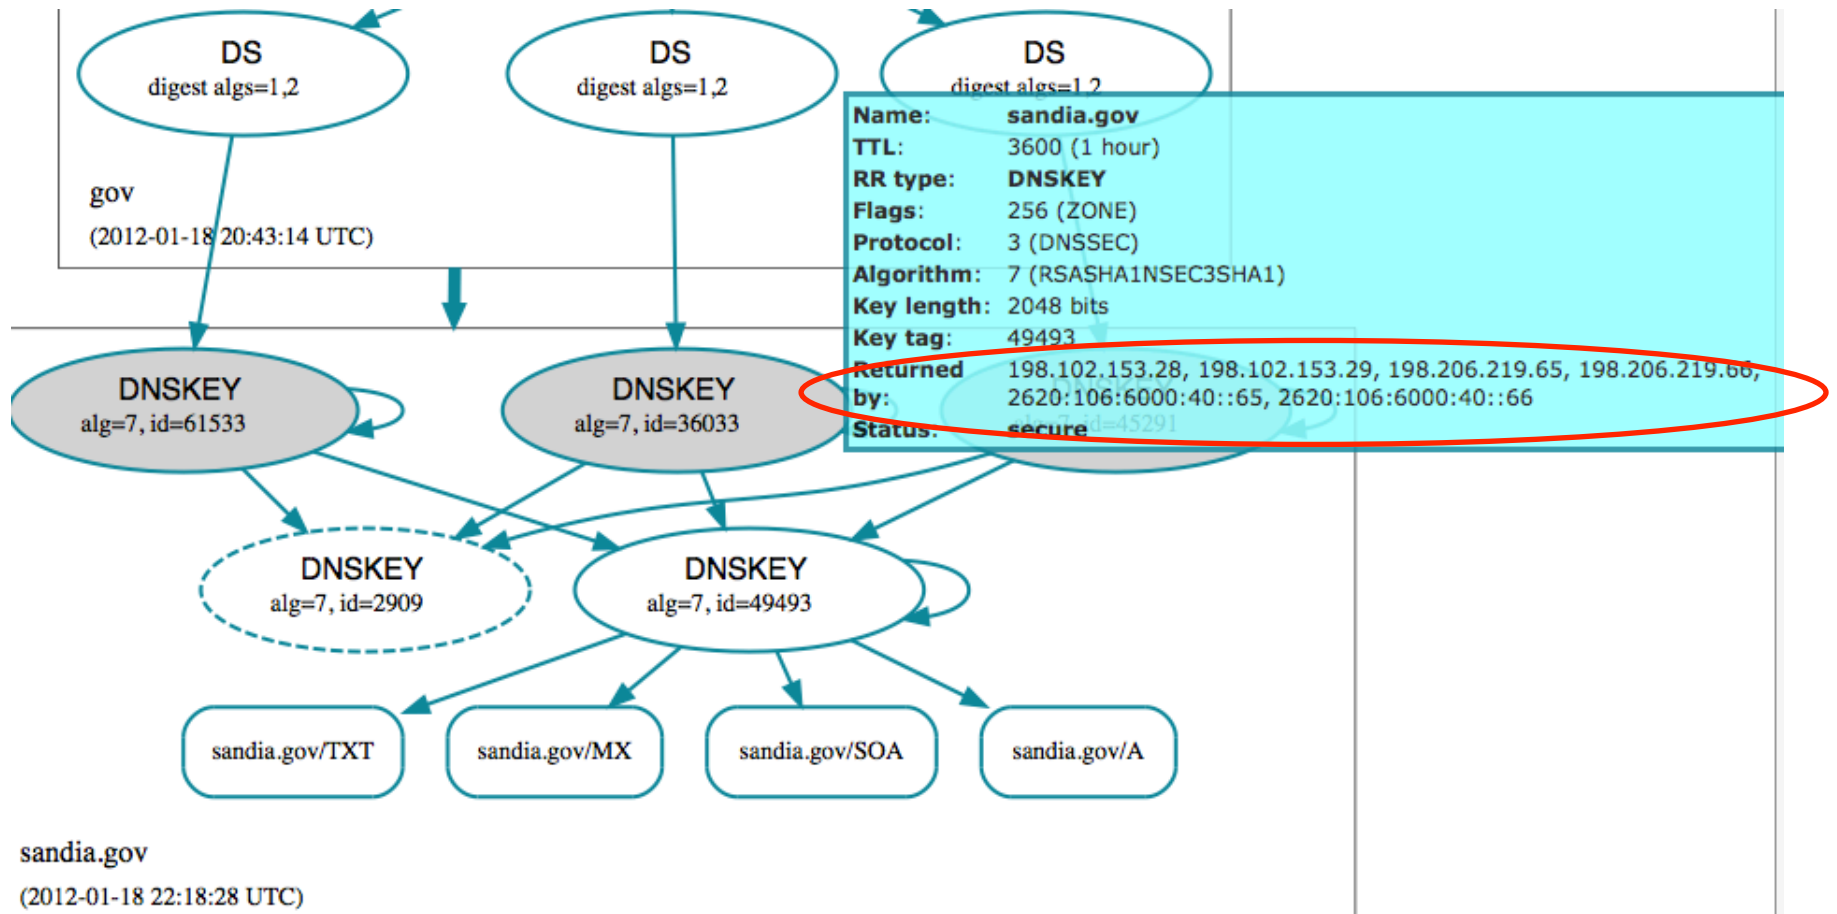
- Top node: DNSKEY (alg=8, id=19036) with a self-loop arrow.
- Middle node: DNSKEY (alg=8, id=56158)
- Bottom node: DS (digest algs=1,2)

Arrows indicate the flow from the top DNSKEY node to the middle DNSKEY node, and from the middle DNSKEY node to the bottom DS node. A vertical line extends downwards from the DS node.

On the left side of the interface, there are several status indicators:

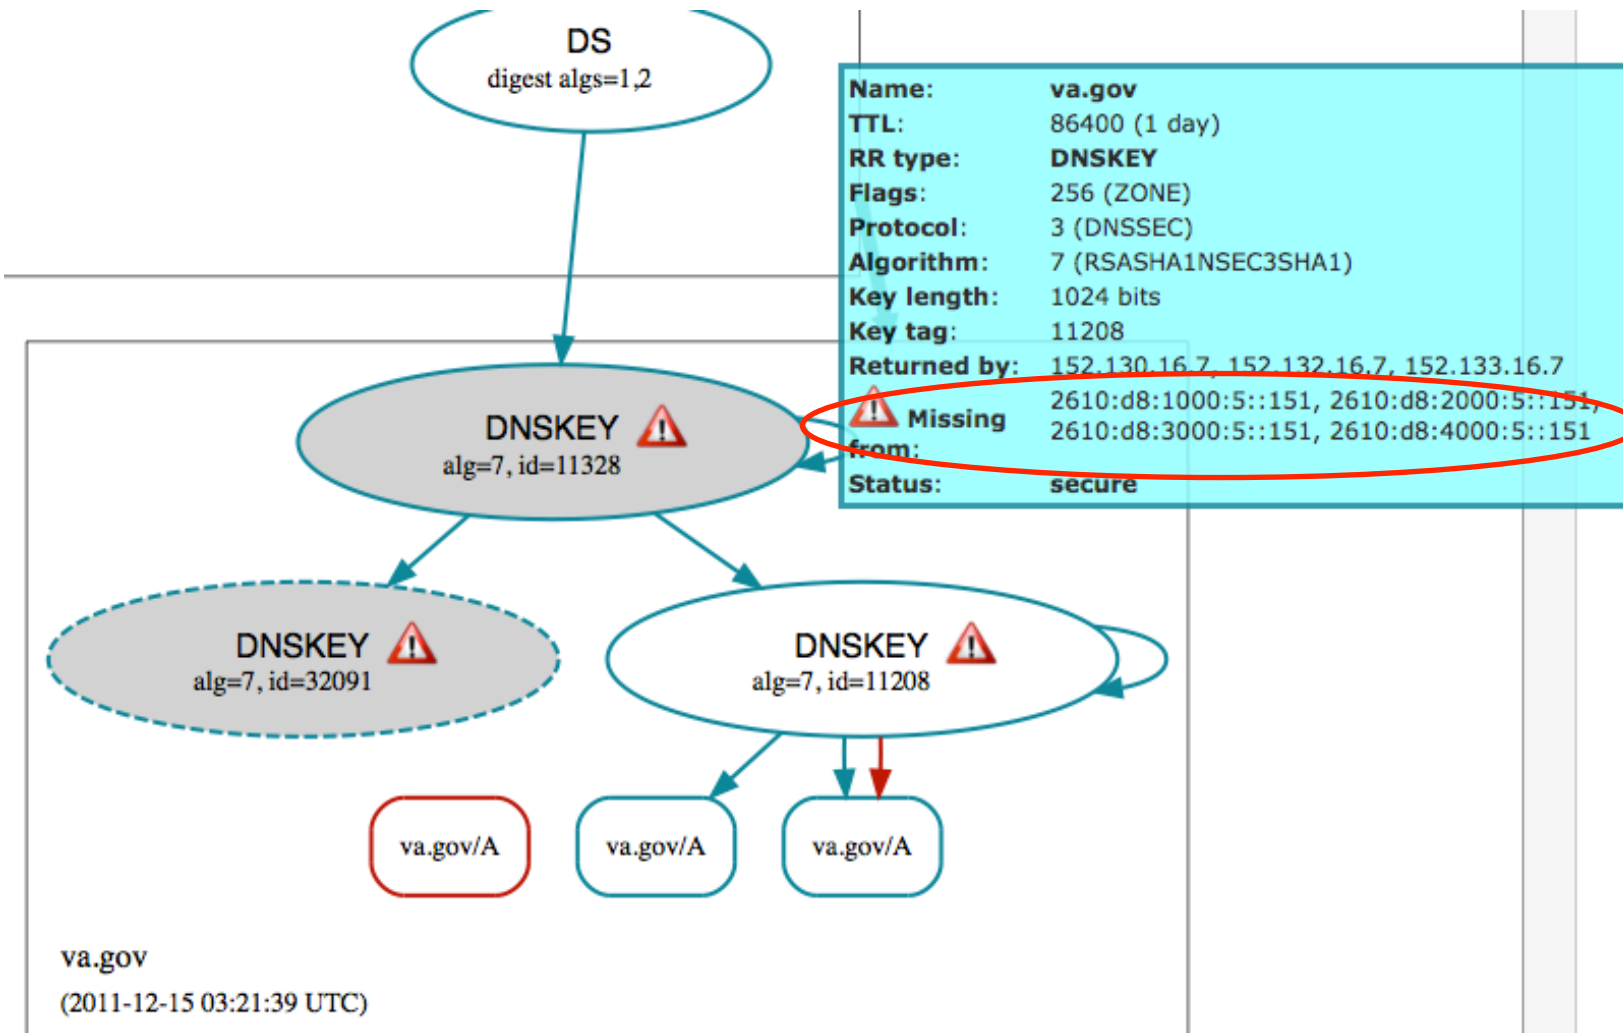
- RRset status:** Secure (4)
- DNSKEY/DS/NSEC status:** Secure (11)
- Delegation status:** Secure (2)
- DNSKEY legend:** Full legend, Published only (dashed circle), SEP bit set (solid circle)

# DNS Consistency with IPv6



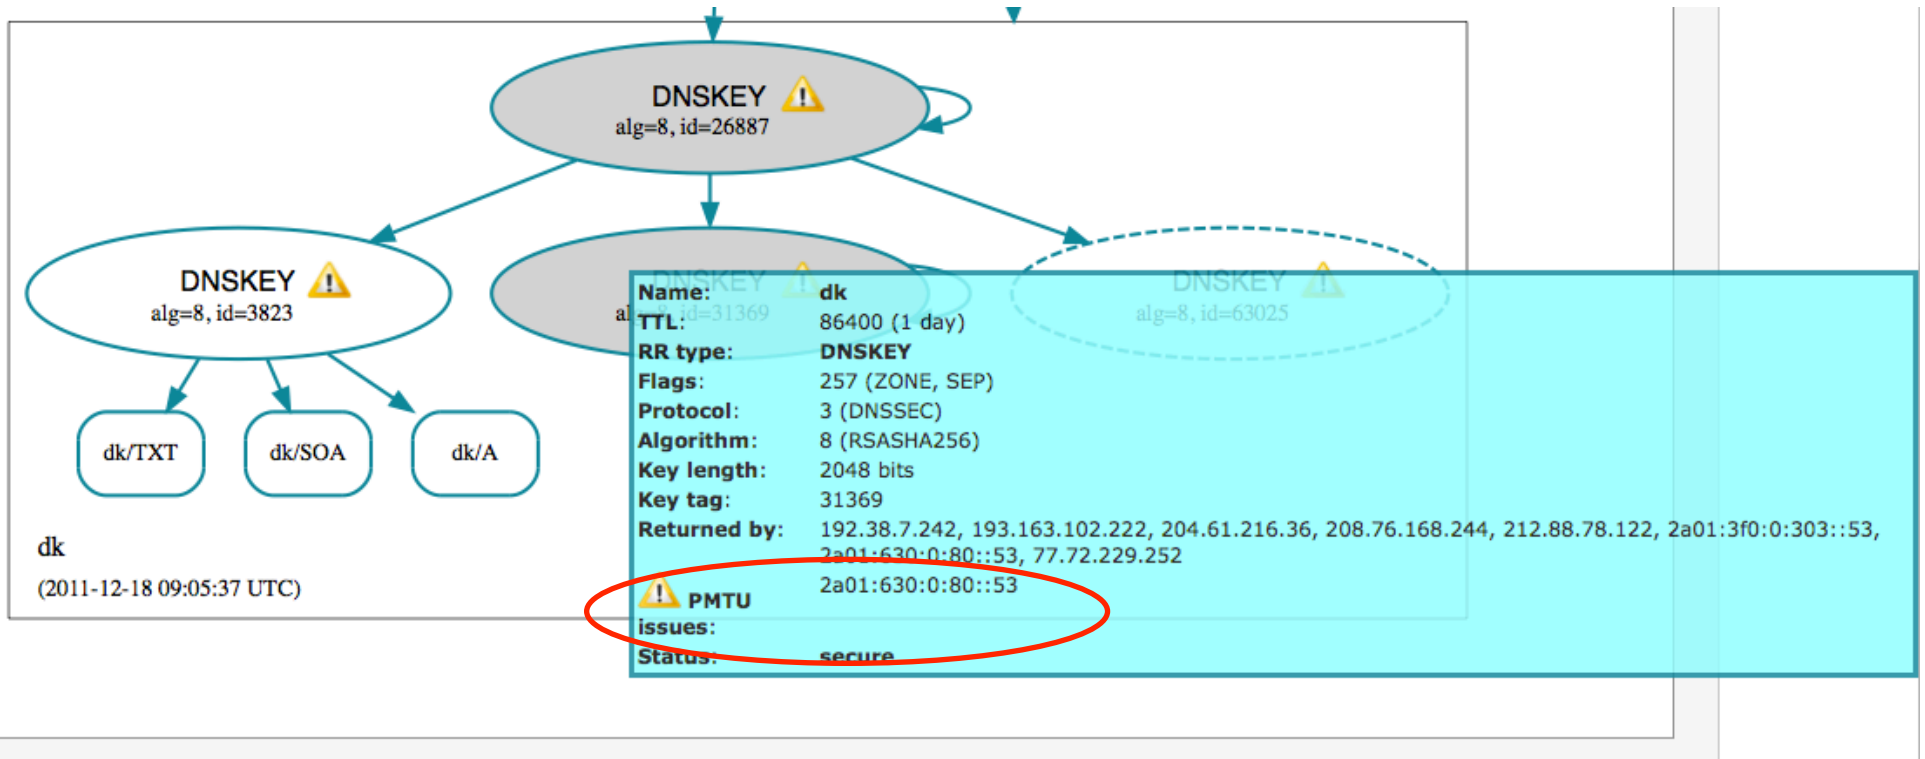
<http://dnsviz.net/>

# DNS(SEC) Consistency with IPv6



<http://dnsviz.net/>

# DNS Consistency with IPv6 – PMTU issues



<http://dnsviz.net/>