NATTED – A Field Report

Troopers 2016, 14th of March 2016

Gabriel Müller, Senior Consultant

Facts and Figures

Activity	Consulting, engineering und project management for information technology from a single source
Owner	The share capital is wholly owned by the partners
Founded in	1986
Employees	Over 170 staff
Clients	Over 400
Projects	Over 4‘000
Site Locations	Zurich, Berne, Basle, Lausanne

Qualification of our Consultants

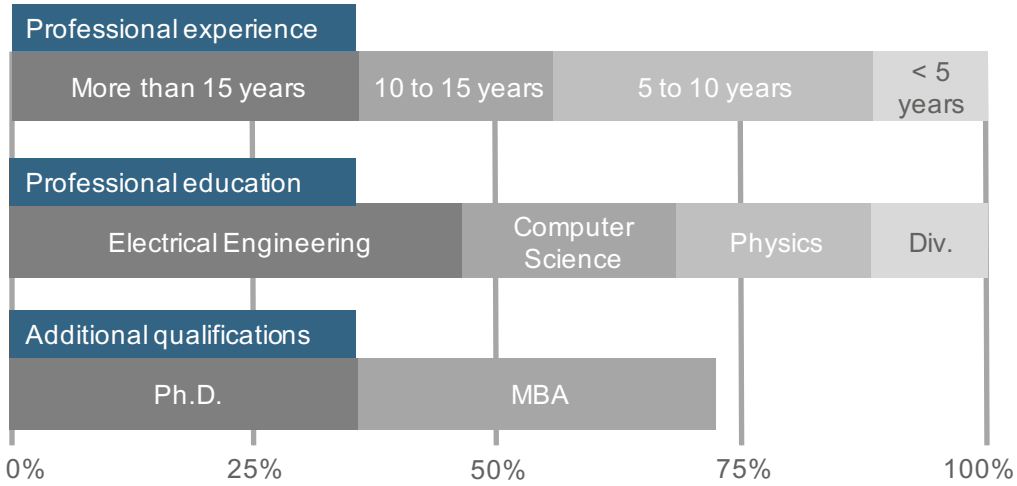

Turnover

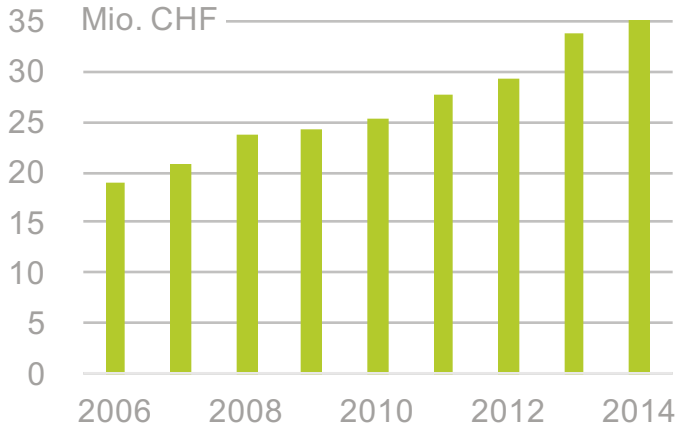

Partners of AWK

From left to right:
Ralph Tonezzer,
Peter Gabriel,
Kurt Biri,
Christian Mauz,
Oliver Vaterlaus
(Managing Partner),
Ueli Sandmeier,
André Arrigoni,

Contents

Why are we doing this?

- … our statement for the need of monitoring in general
  - Complexity
  - Redundancy
- … curiosity
  - Does it work?
  - Performance impact
- … a valid starting point for IPv6 deployment
  - Lifecycle
  - Today: IPv6 islands – tomorrow: IPv4 islands

NAT and its Reputation

- One-to-Many — permanent
  - IPv4: Public IP address shortage
  - IPv6: n/a
- One-to-One — permanent
  - IPv4: Merging networks of multiple organisations / entities
  - IPv6: Hiding (?)
- NAT64 — temporary
  - Transition mechanism
- The big difference

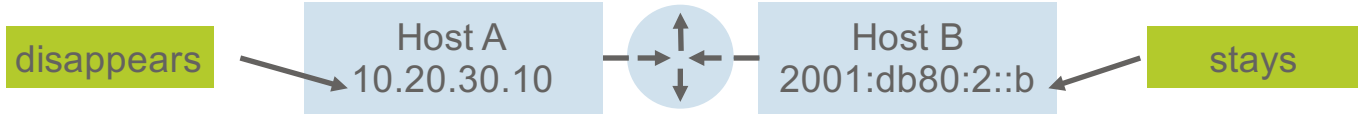

Contents

Transition Mechanism - Recap

- When do you use a transition mechanism?
  - If your transport is IPv4-only or IPv6-only.
  - If your communication endpoints IP protocols miss-match.

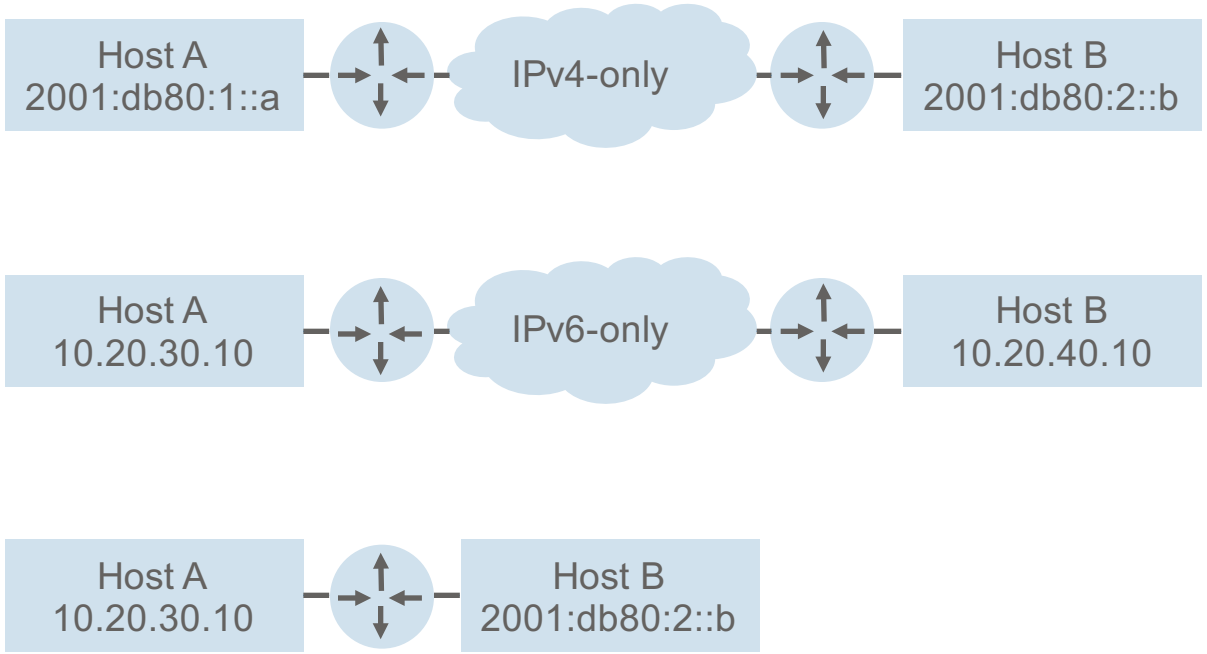

Transition Mechanisms - Overview

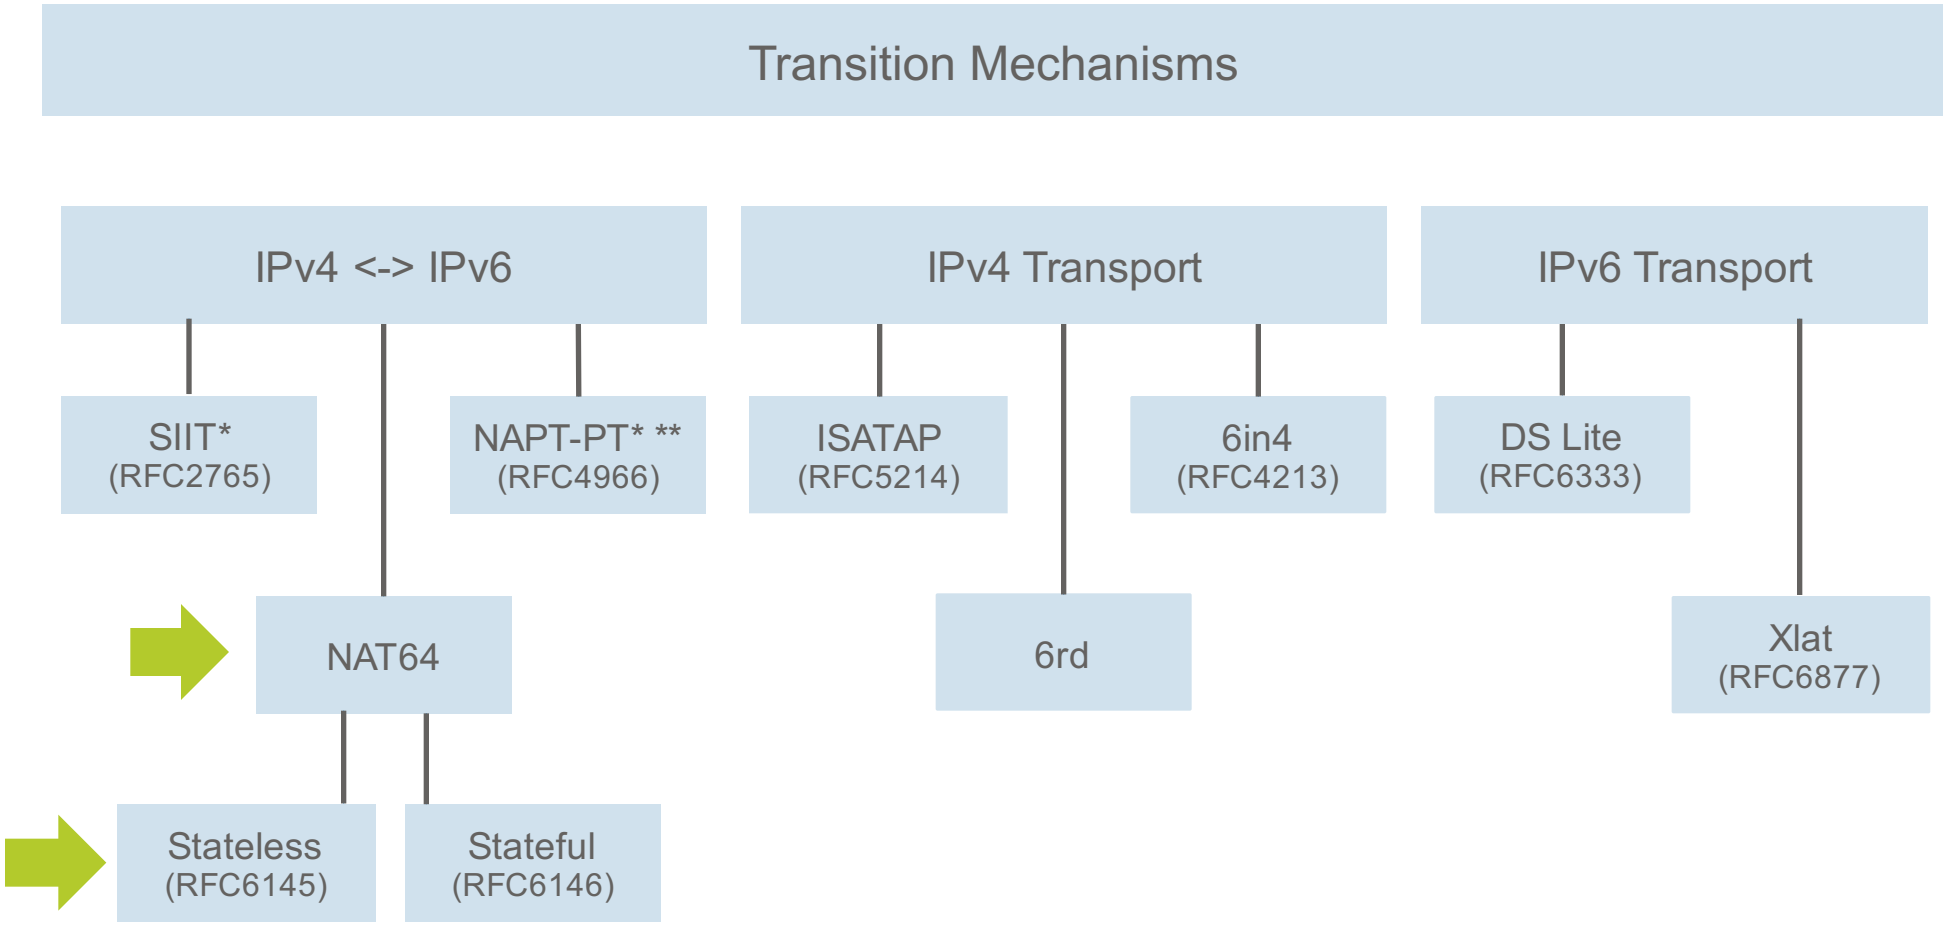

* Deprecated RFCs / ** Also NAT-PT (RFC2766)

Contents

Starting Position

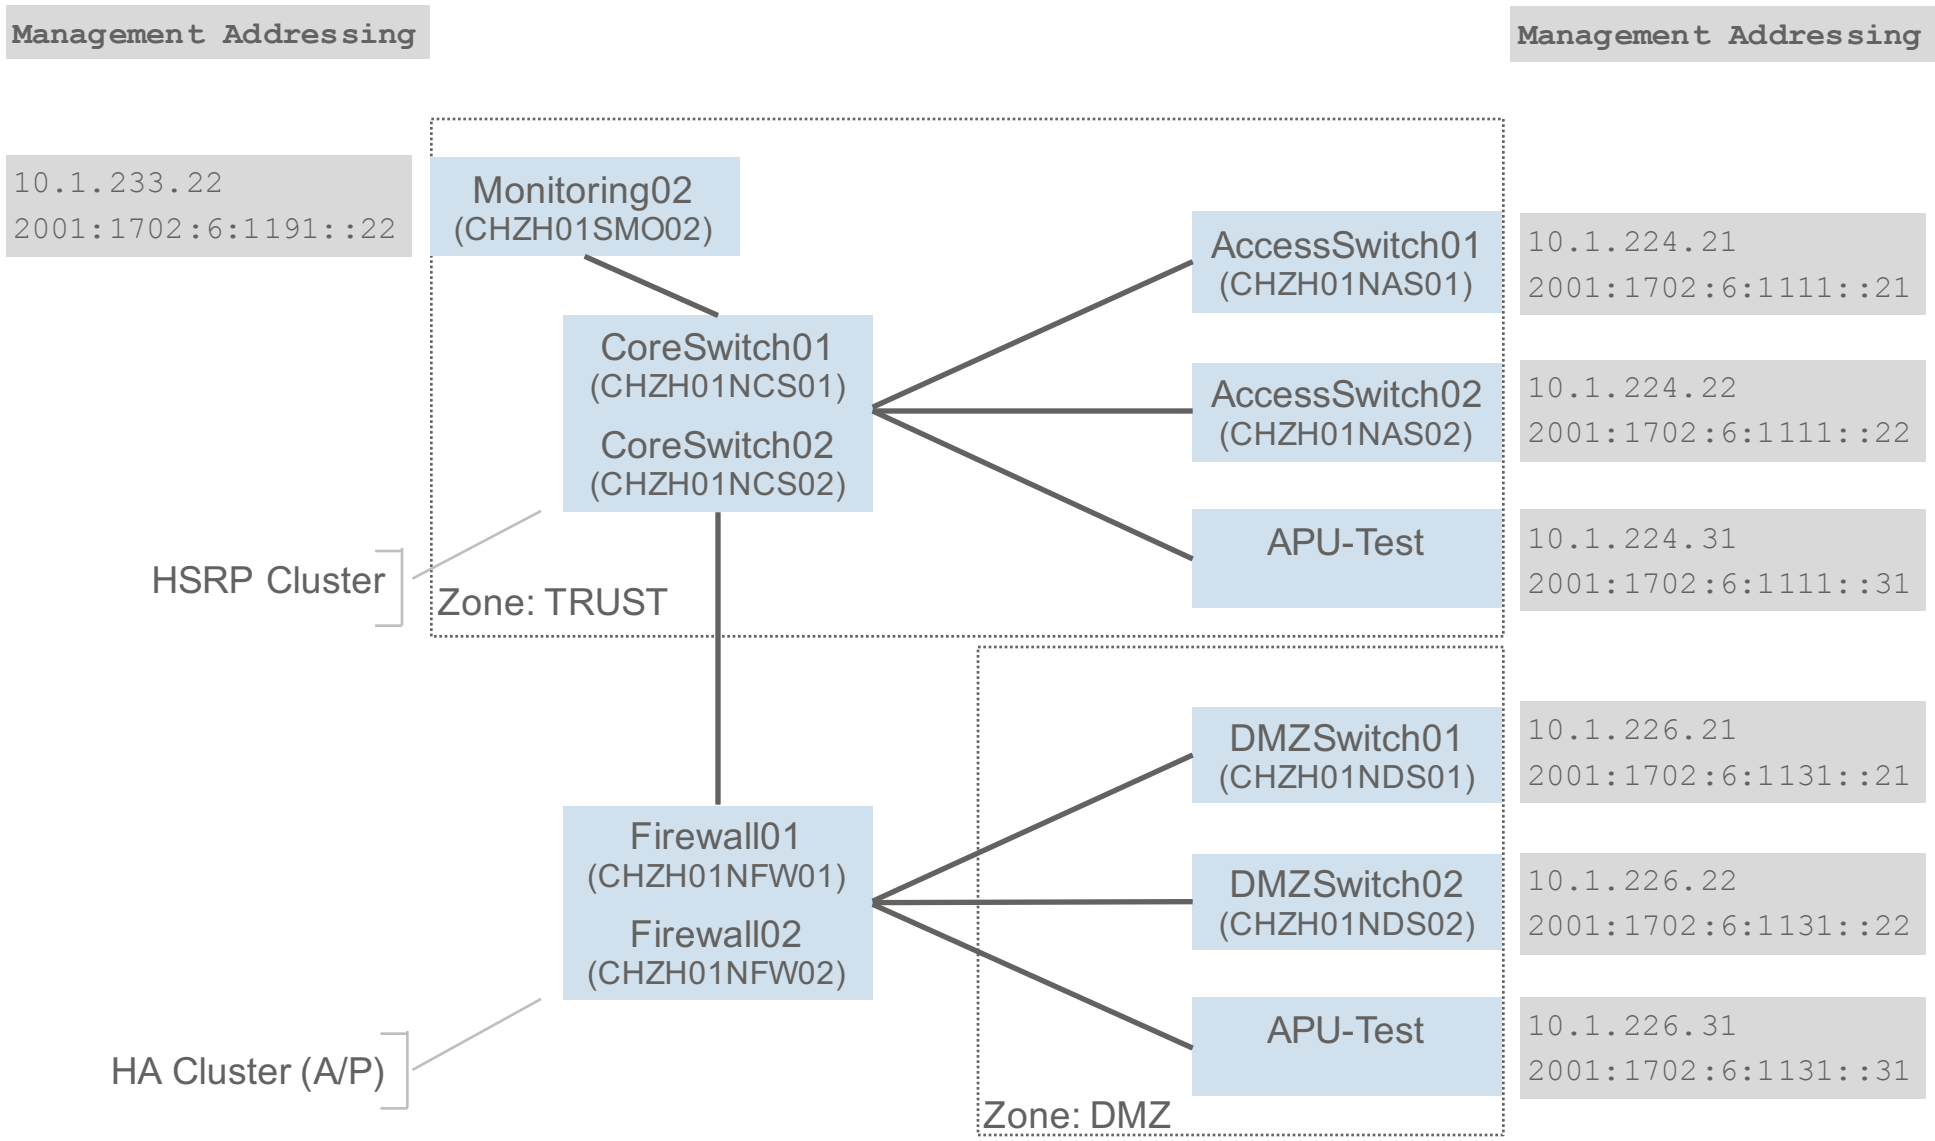

Target

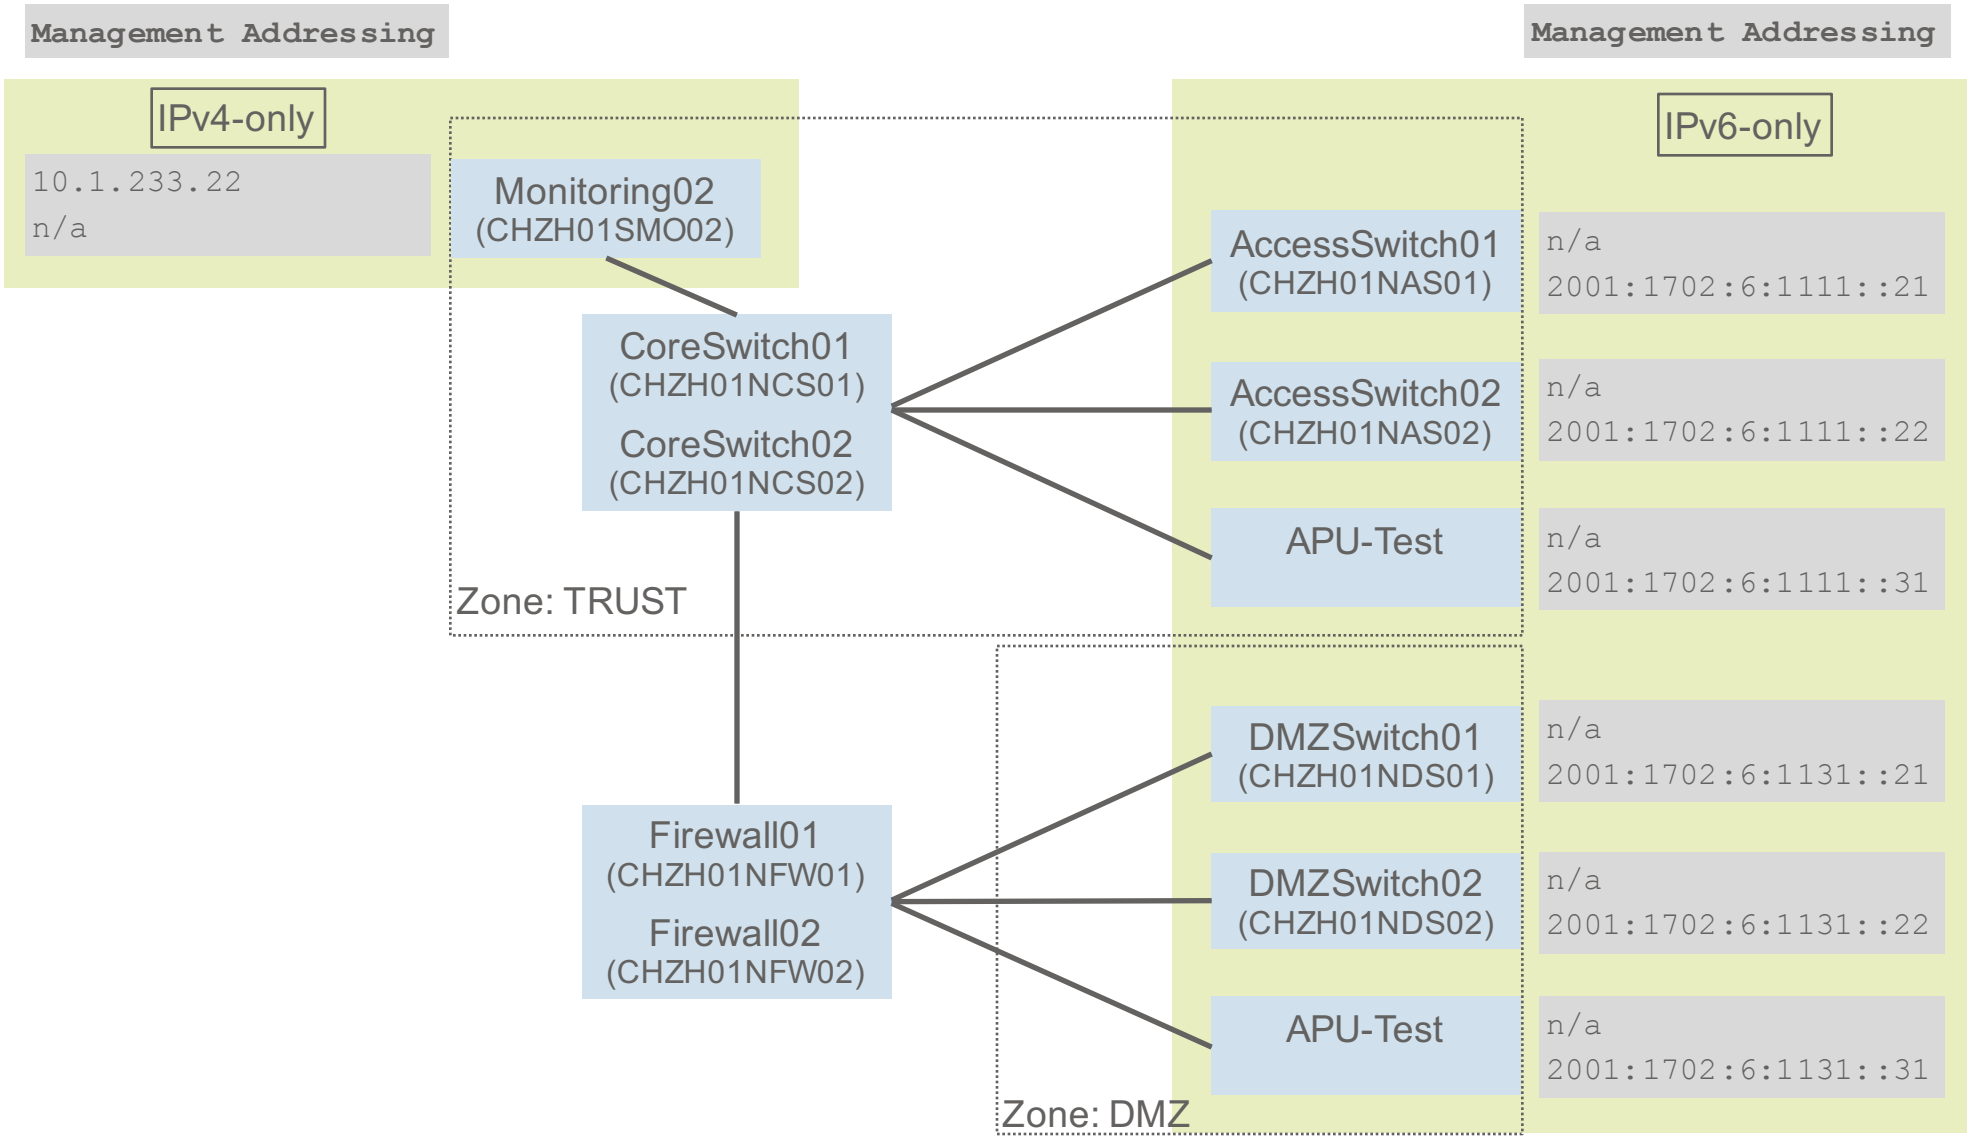
Management Addressing
Management Addressing
IPv4-only
10.1.233.22
n/a
IPv6-only
Monitoring02
(CHZH01SMO02)
CoreSwitch01
(CHZH01NCS01)
CoreSwitch02
(CHZH01NCS02)
AccessSwitch01
(CHZH01NAS01)
n/a
2001:1702:6:1111::21
AccessSwitch02
(CHZH01NAS02)
n/a
2001:1702:6:1111::22
APU-Test
n/a
2001:1702:6:1111::31
Zone: TRUST
Firewall01
(CHZH01NFW01)
Firewall02
(CHZH01NFW02)
DMZSwitch01
(CHZH01NDS01)
n/a
2001:1702:6:1131::21
DMZSwitch02
(CHZH01NDS02)
n/a
2001:1702:6:1131::22
APU-Test
n/a
2001:1702:6:1131::31
Zone: DMZ

Design Considerations (1/4)

- What do we need?
  - Some device doing NAT64 translation
  - Available devices
    - CoreSwitch01/02
    - Firewall01/02
- Checking device support for NAT64
  - Core switch
    - Cisco Catalyst 4500e
    - Sup 6L-E / IOS 15.2(2)E3
  - Firewall
    - Juniper SRX240 H2
    - Junos 12.3X48-D15.4

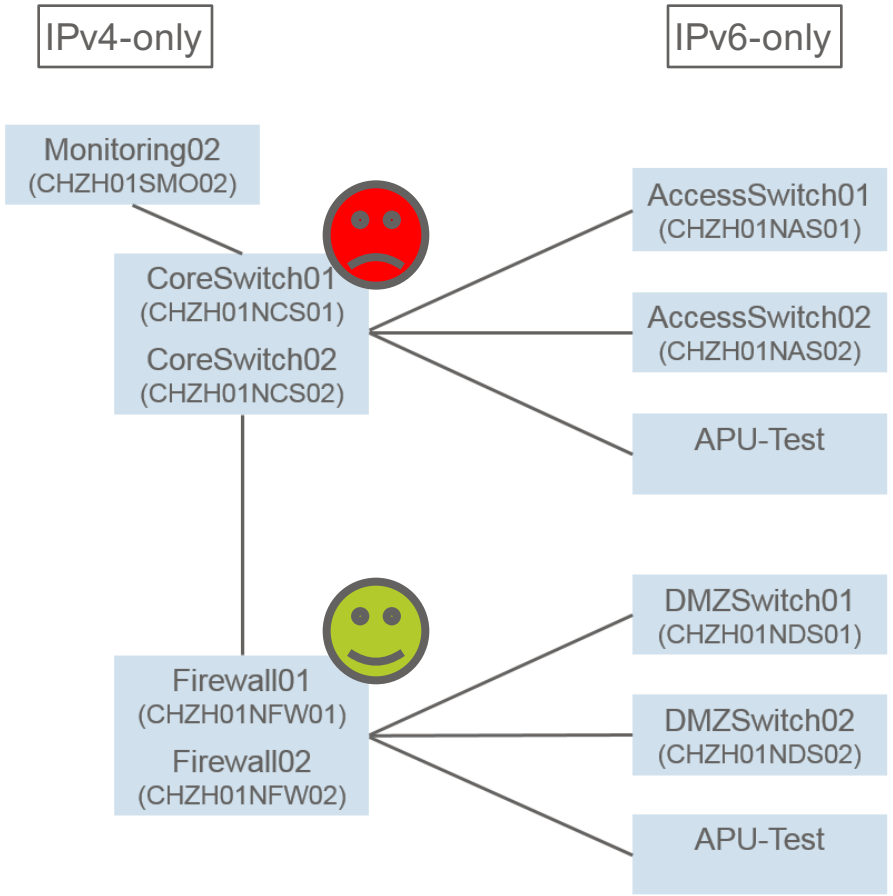

Design Considerations (2/4)

- Monitoring02 ⇔ DMZ Devices
  - No issues, SRX supports NAT64
- Monitoring02 ⇔ Trust Devices
  - Problematic
  - Solution: Route through Firewall

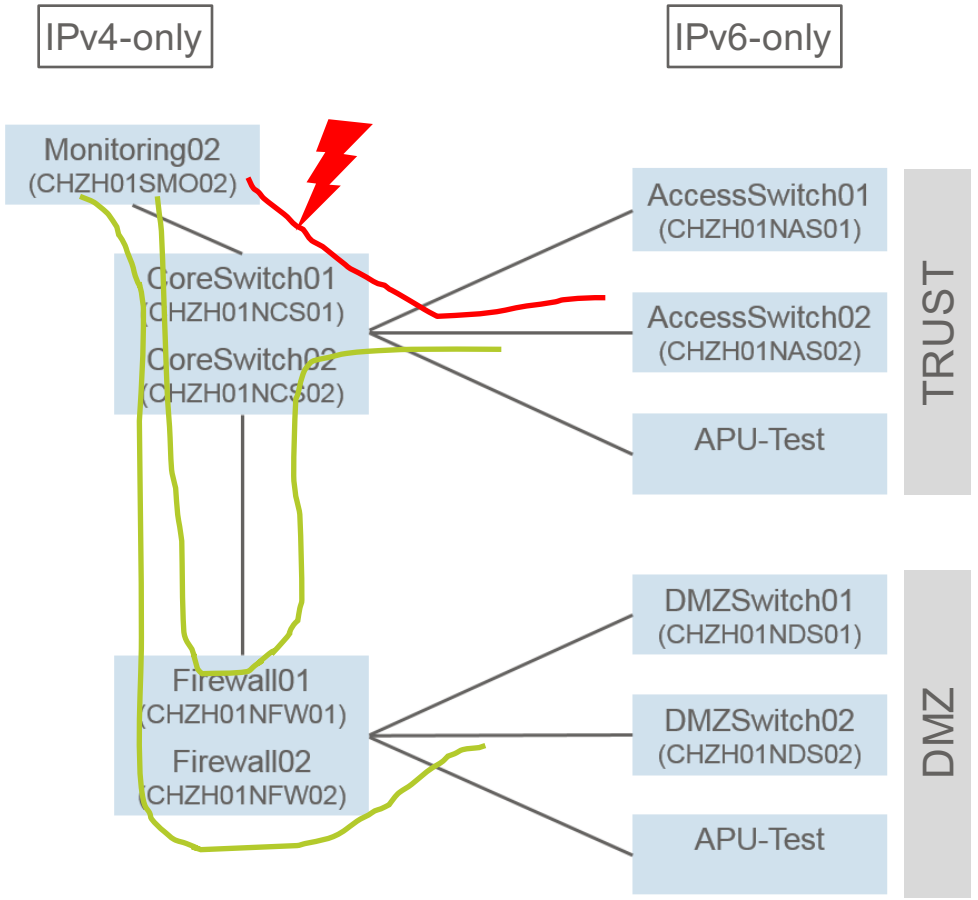

Design Considerations (3/4)

- Addressing – Static 1to1 NAT

Device	IPv4 Address	IPv6 Address
AccessSwitch01	**10.1.234.121**	2001:1702:6:1111::21
AccessSwitch02	**10.1.234.122**	2001:1702:6:1111::22
APU-Test	**10.1.234.131**	2001:1702:6:1111::31
DMZSwitch01	**10.1.226.121**	2001:1702:6:1131::21
DMZSwitch02	**10.1.226.122**	2001:1702:6:1131::22
APU-Test	**10.1.226.131**	2001:1702:6:1131::31
Monitoring02	10.1.233.22	**2001:1702:6:1191::1122** **2001:1702:6:11a1::1122**

Bold means virtual IP, not assigned to a device

Design Considerations (4/4)

- Something missing?
  - What about network 10.1.234.0/24?
  - Purely virtual, used to forward packets from monotoring02 to firewall
  - It is new…
  - ⇨ Routing entry on CoreSwitch01/02 required, next hop Firewall01/02

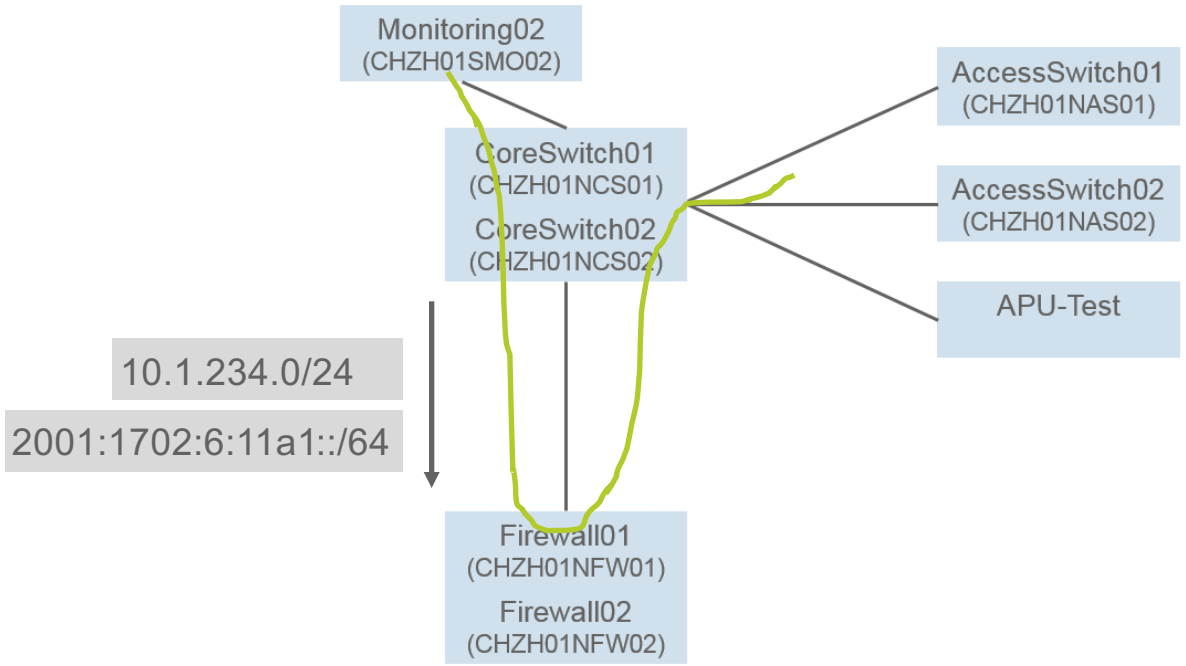

NAT64 Configuration (SRX 240) – to Zone DMZ

- Static NAT for destination address

```
mug@CHZH01NFWCL01# show security nat static
rule-set RS_NAT6446_from_reth21013 {
    from interface reth2.1013;
    rule R_001_DMZ {
        match {
            destination-address-name 10.1.226.121/32;
        }
        then {
            static-nat {
                prefix-name {
                    2001:1702:6:1131::21/128;
                }
            }
        }
    }
}
```

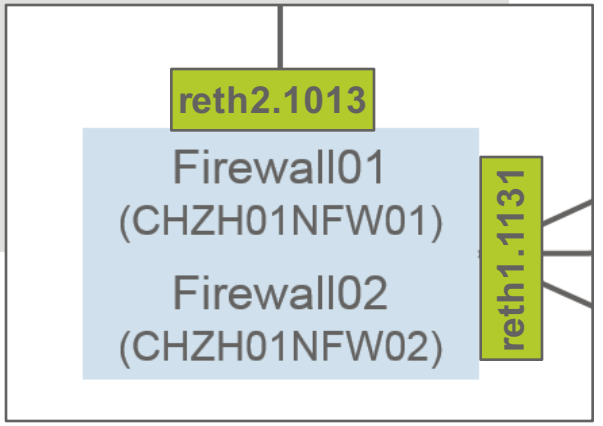

NAT64 Configuration (SRX 240) – to Zone DMZ

- Changing source IP

```
mug@CHZH01NFWCL01# show security nat source
rule-set RS_NAT6446_from_reth11131 {
    from interface reth2.1013;
    to interface reth1.1131;
    rule R_001_DMZ {
        match {
            source-address-name CHZH01_MGMT_10.1.233.0/24;
            destination-address-name CHZH01_MGMT_2001:1702:6:1131/64;
        }
        then {
            source-nat {
                interface;
            }
        }
    }
}
```

```
// interface configuration
family inet6 {
                address 2001:1702:6:1131::10/64;
                address fe80::1131:0:0:10/64;
            }
```

NAT64 Configuration (SRX 240) – from Zone DMZ

- Static NAT for destination address

```
mug@CHZH01NFWCL01# show security nat static
rule-set RS_NAT6446_from_reth11131 {
    from interface reth1.1131;
    rule R_001_ToSyslog {
        match {
            destination-address-name 2001:1702:6:1191::1122/128;
        }
        then {
            static-nat {
                prefix-name {
                    CHZH01_SMO02_v4-10.1.233.22/32;
                }
            }
        }
    }
}
```

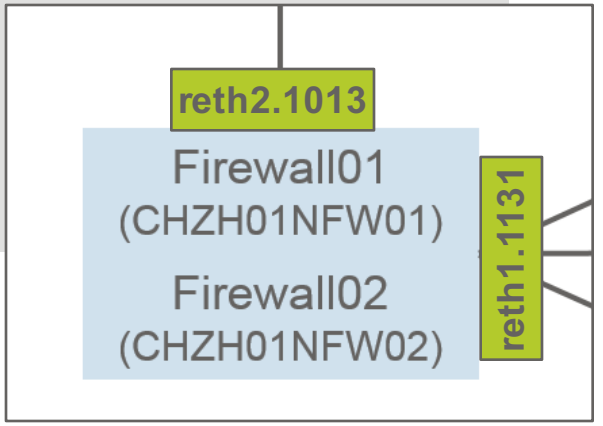

NAT64 Configuration (SRX 240) – from Zone DMZ

- Changing Source IP
  - Not required
  - Static NAT entry is used (applied in reversed order)

```
mug@CHZH01NFWCL01# show security nat static
rule-set RS_NAT6446_from_reth21013 {
    from interface reth2.1013;
    rule R_001_DMZ {
        match {
            destination-address-name 10.1.226.121/32;
        }
        then {
            static-nat {
                prefix-name {
                    2001:1702:6:1131::21/128;
                }
            }
        }
    }
}
```

NAT64 Configuration (SRX 240) – to ZONE DMZ

- From IPv4 to IPv6

```
mug@CHZH01NFWCL01# show security policies from-zone MGMT to-zone MGMT
...
policy MGMT_MGMT_007-TEMP {
    match {
        source-address [ CHZH01_SMO02_v4-10.1.233.22/32 … ];
        destination-address [ 2001:1702:6:1131::21/128 … ];
        application [ snmp junos-ssh junos-ping junos-https ];
    }
    then {
        permit;
    }
}
```

NAT64 Configuration (SRX 240) – from ZONE DMZ

- From IPv6 to IPv4

```
mug@CHZH01NFWCL01# show security policies from-zone MGMT to-zone MGMT
...
policy MGMT_MGMT_009-TEMP {
    match {
        source-address [ 2001:1702:6:1131::21/128 … ];
        destination-address CHZH01_SMO02_v4-10.1.233.22/32;
        application [ junos-ssh junos-syslog junos-pingv6 junos-https ];
    }
    then {
        permit;
    }
}
```

NAT64 Configuration (SRX 240) – Policy Handling

- By the way…
  - Why do we specify rules with Src IPv4 – Dst IPv6 or vice versa?
  - Any idea?

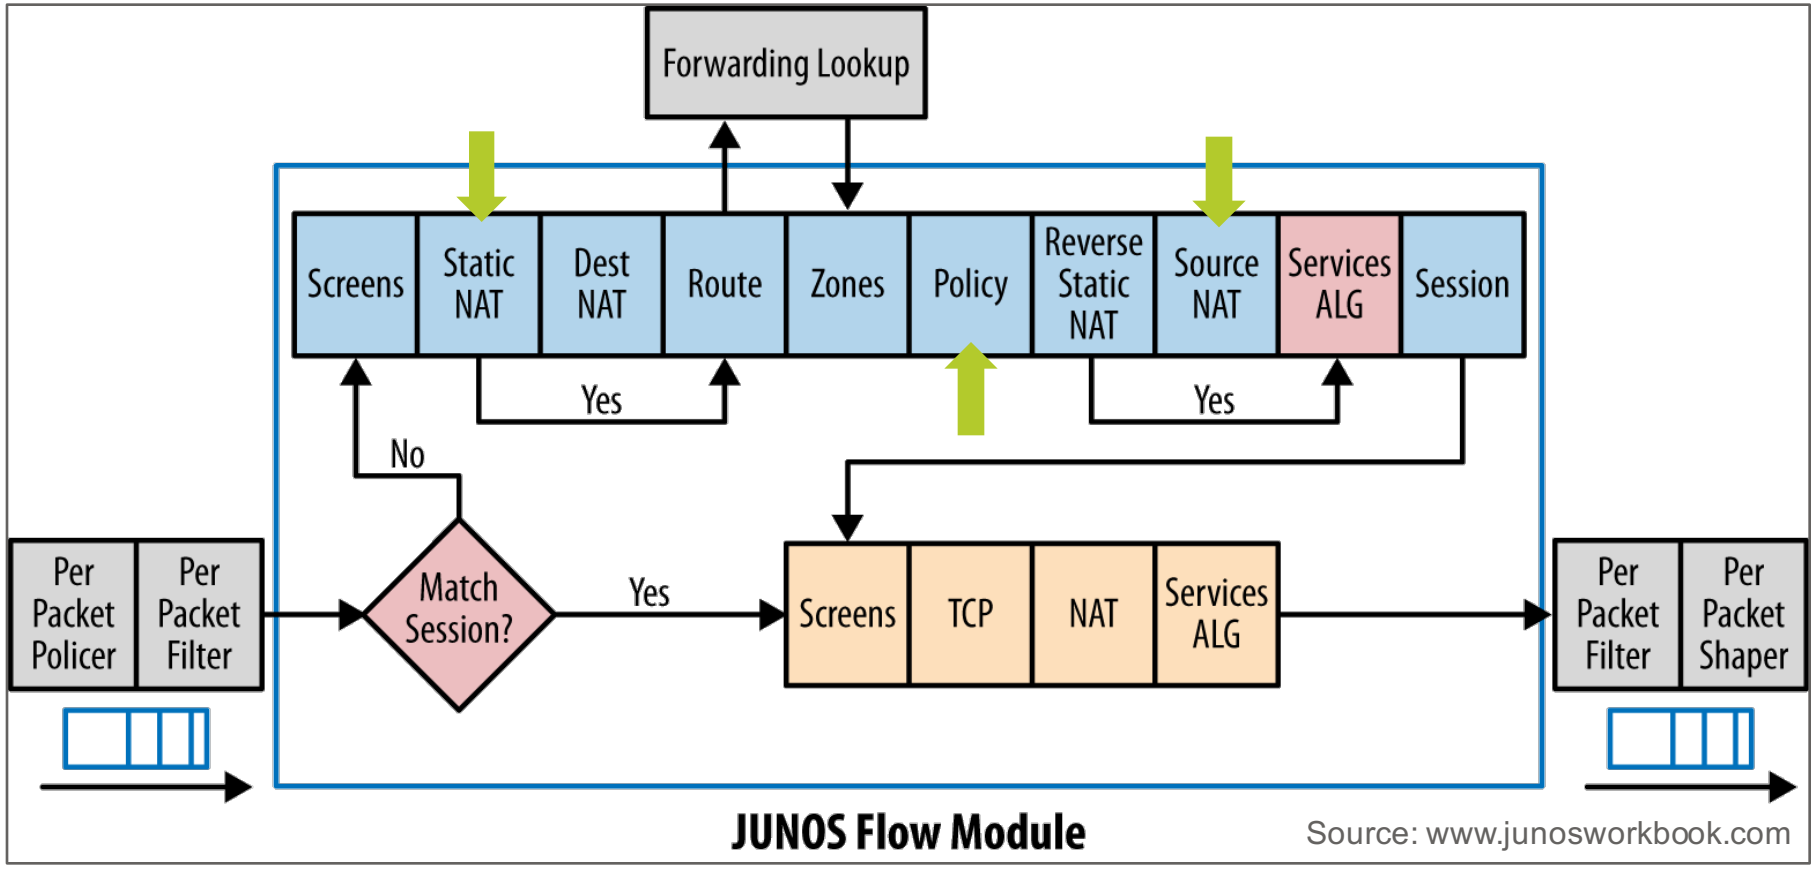

Contents

A very first packet (1/2)

- Do you remember?

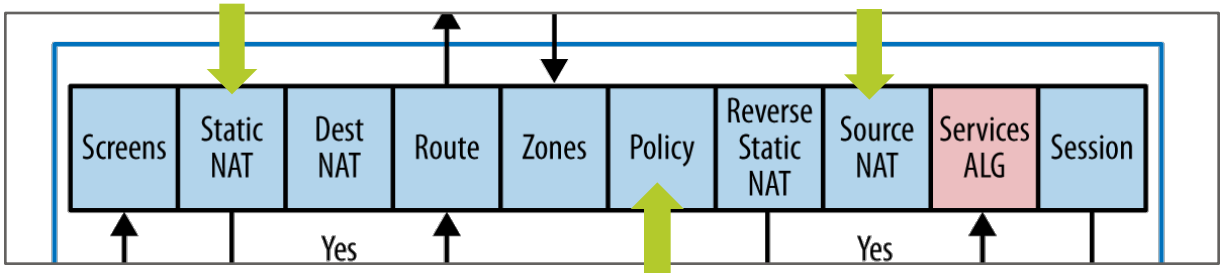


```
// Packet arriving (v4 -> v4)
Jan 27 08:27:06 08:27:06.173280:CID-1:RT:  reth2.1013:10.1.233.22->10.1.226.101, icmp, (8/0)
Jan 27 08:27:06 08:27:06.173421:CID-1:RT:  chose interface reth2.1013 as incoming nat if.

// Static NAT & Route Lookup
Jan 27 08:27:06 08:27:06.173421:CID-1:RT:flow_first_rule_dst_xlate: packet 10.1.233.22->10.1.226.101
nsp2 change to 2001:1702:6:1131:0:0:0:21.
Jan 27 08:27:06 08:27:06.173421:CID-1:RT:  Doing DESTINATION addr route-lookup
Jan 27 08:27:06 08:27:06.173421:CID-1:RT:  routed (x_dst_ip 2001:1702:6:1131:0:0:0:21) from MGMT
(reth2.1013 in 1) to reth1.1131, Next-hop: 2001:1702:6:1131:0:0:0:21

// Searching for Security Policy
Jan 27 08:27:06 08:27:06.173421:CID-1:RT:  flow_first_policy_search: policy search from zone MGMT->
zone MGMT (0x114,0x1378a,0x378a)
Jan 27 08:27:06 08:27:06.173421:CID-1:RT:  Policy lkup: vsys 0 zone(6:MGMT) -> zone(6:MGMT) scope:0
Jan 27 08:27:06 08:27:06.173421:CID-1:RT:  10.1.233.22/2048 -> 2001:1702:6:1131:0:0:0:21/17974 proto
Jan 27 08:27:06 08:27:06.173913:CID-1:RT:  permitted by policy MGMT_MGMT_006_TEMP(176)
```

A very first packet (2/2)

- Do you remember?

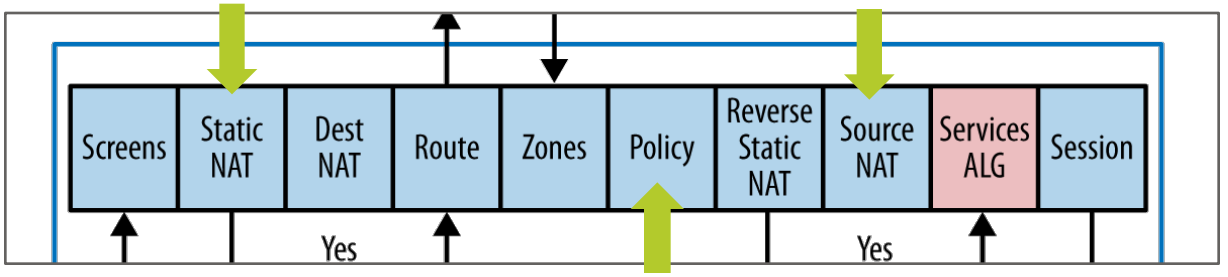


```
// Source NAT
Jan 27 08:27:06 08:27:06.173913:CID-1:RT:flow_first_src_xlate: src nat returns status: 1, rule/pool
id: 5/2, pst_nat: False.
Jan 27 08:27:06 08:27:06.173913:CID-1:RT:  dip id = 2/0, 10.1.233.22/1-
>2001:1702:6:1131:0:0:0:10/38506
Jan 27 08:27:06 08:27:06.174727:CID-1:RT:handle icmp xlate v4 to v6
Jan 27 08:27:06 08:27:06.174786:CID-1:RT:  post addr xlation: 2001:1702:6:1131:0:0:0:10-
>2001:1702:6:1131:0:0:0:21.

// Packet leaving (v6 -> v6)
Jan 27 08:27:06 08:27:06.174812:CID-1:RT:**** jump to packet after xlate:2001:1702:6:1131:0:0:0:10-
>2001:1702:6:1131:0:0:0:21
131
```

- Questions ?
  - Why is :1131::10 used as source IP?

SSH Login & Syslog

- SSH from monitoring02 to DMZSwitch01

```
mug@monitoring02:~$ ssh 10.1.226.121
Password:
CHZH01NDS01>en
Password:
CHZH01NDS01#show users
    Line       User       Host(s)              Idle       Location
*  1 vty 0     mug        idle                 00:00:00 2001:1702:6:1191::1122
CHZH01NDS01#
```

- Syslog from DMZSwitch01 to monitoring02

```
root@monitoring02:/home/mug# tail -f /var/log/syslog
...
Feb 10 21:57:18 chzh01nds01t.awkgroup.com 11444:
  .Feb 10 20:57:17.853: %SYS-6-LOGGINGHOST_STARTSTOP: Logging to host
  2001:1702:6:1191::1122 port 514 started - CLI initiated
```

10.1.226.121

```
CHZH01NDS01#show run | include 1122
logging host ipv6 2001:1702:6:1191::1122
```

Monitoring of DMZSwitch01 (native via IPv6)

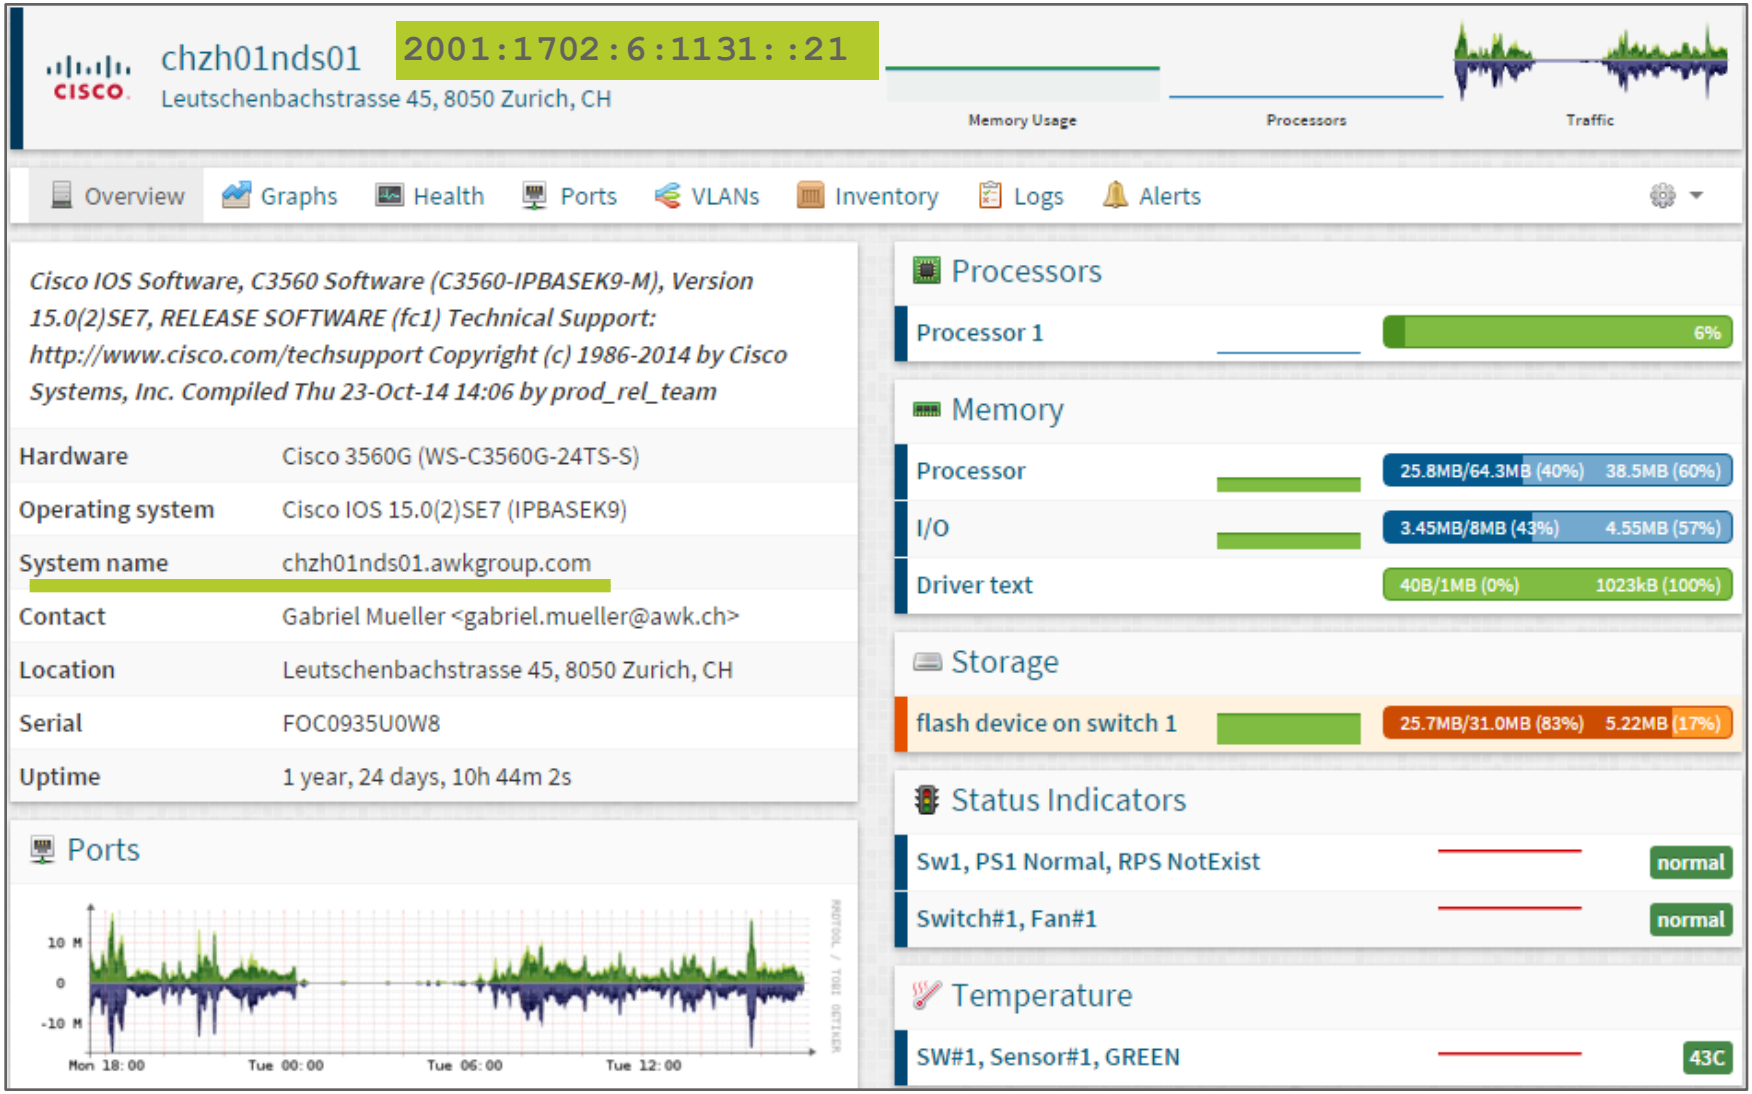

Monitoring of DMZSwitch01 (NAT64)

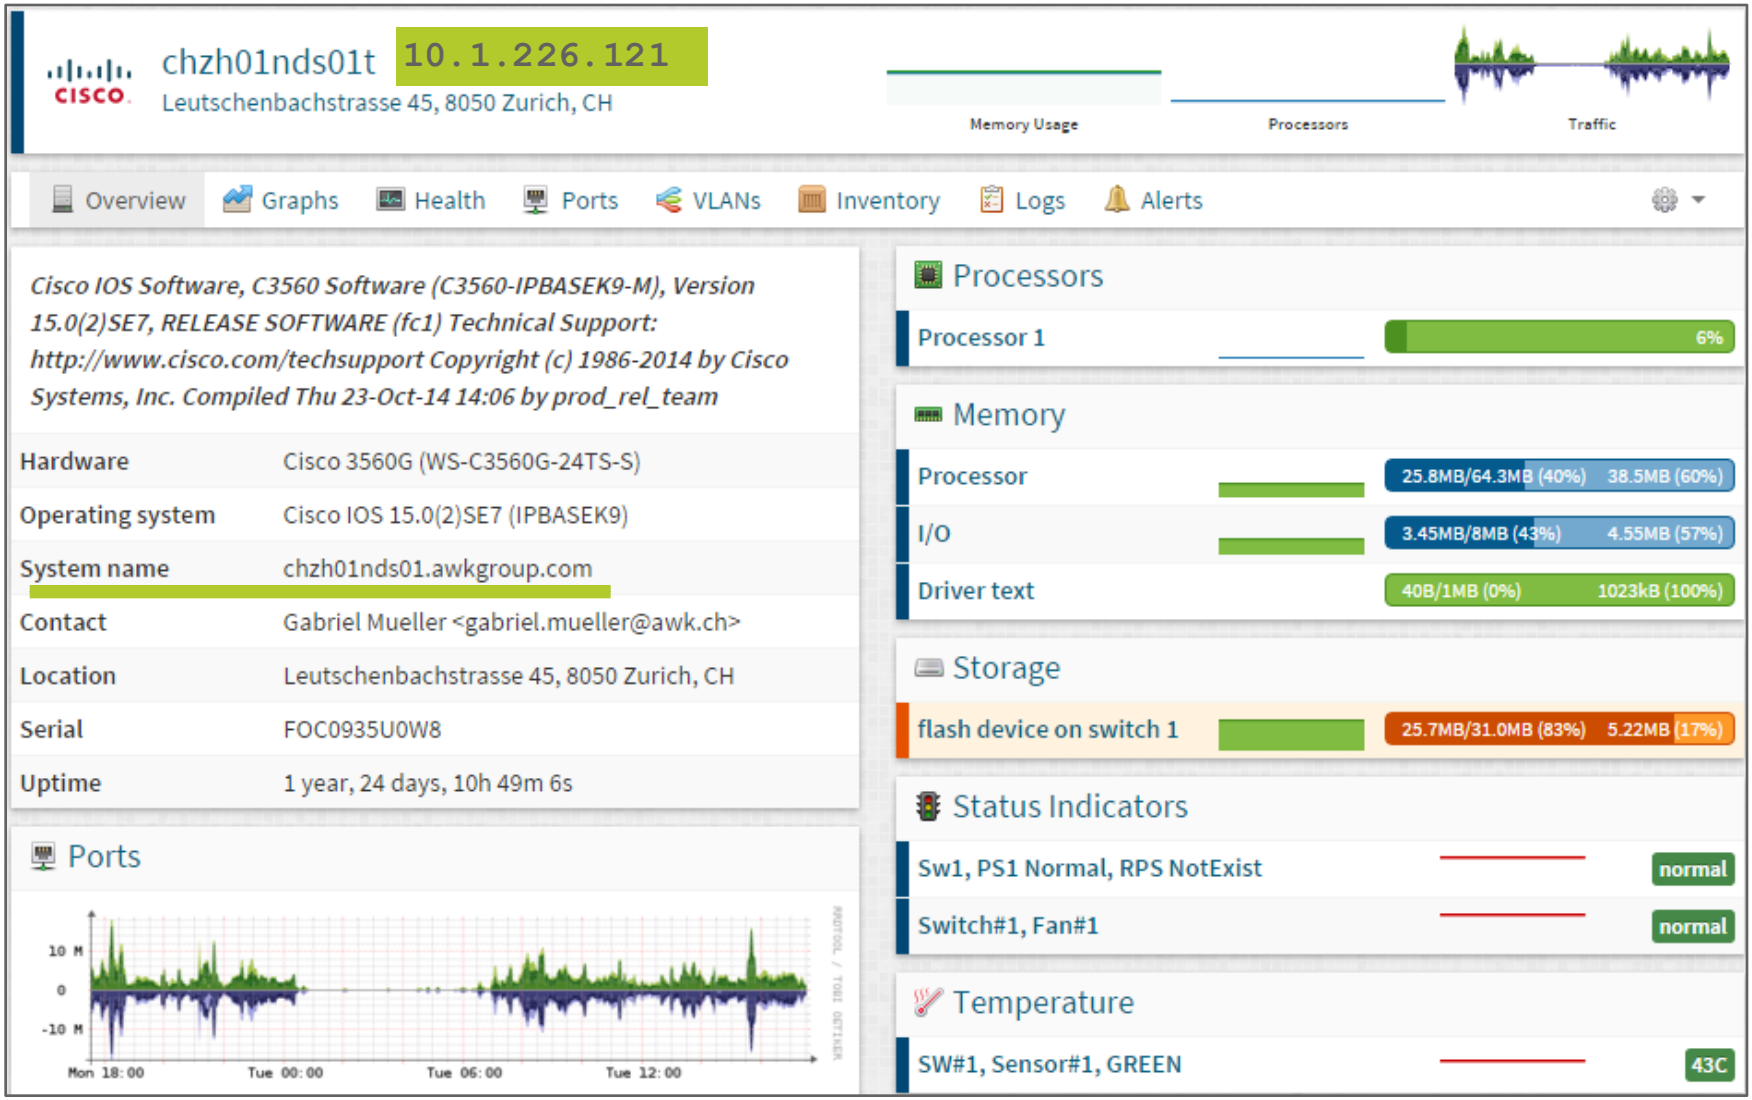

Performance – Ping Flooding – to Zone DMZ (1/2)

- Native IPv4

```
root@Monitoring02:/home/mug# ping -f 10.1.226.31 -c 10000
PING 10.1.226.31 (10.1.226.31) 56(84) bytes of data.

--- 10.1.226.31 ping statistics ---
10000 packets transmitted, 10000 received, 0% packet loss, time 9023ms
rtt min/avg/max/mdev = 0.626/0.852/14.050/0.486 ms, pipe 2, ipg/ewma 0.902/0.820 ms
```

- Native IPv6

```
root@monitoring02:/home/mug# ping6 -f 2001:1702:6:1131::31 -c 10000
PING 2001:1702:6:1131::31(2001:1702:6:1131::31) 56 data bytes

--- 2001:1702:6:1131::31 ping statistics ---
10000 packets transmitted, 10000 received, 0% packet loss, time 7644ms
rtt min/avg/max/mdev = 0.420/0.707/13.250/0.380 ms, pipe 2, ipg/ewma 0.764/1.189 ms
```

- ⇨ IPv6 flow processing faster (-;

Performance – Ping Flooding – to Zone DMZ (2/2)

- NAT64

```
root@monitoring02:/home/mug# ping -f 10.1.226.131 -c 10000
PING 10.1.226.131 (10.1.226.131) 56(84) bytes of data.

--- 10.1.226.131 ping statistics ---
10000 packets transmitted, 10000 received, 0% packet loss, time 9391ms
rtt min/avg/max/mdev = 0.633/0.895/44.000/0.818 ms, pipe 4, ipg/ewma 0.939/0.850 ms
```

- Comparission
  - IPv4 Native: 9023ms
  - IPv6 Native: 7644ms
  - NAT64: 9391ms (+ 4% / +23%)

Performance – Ping Flooding – to Zone TRUST (1/2)

- Native IPv4

```
root@monitoring02:/home/mug# ping -f 10.1.224.31 -c 10000
PING 10.1.224.31 (10.1.224.31) 56(84) bytes of data.

--- 10.1.224.31 ping statistics ---
10000 packets transmitted, 10000 received, 0% packet loss, time 8600ms
rtt min/avg/max/mdev = 0.573/0.800/14.027/0.482 ms, pipe 2, ipg/ewma 0.860/0.804 ms
```

- Native IPv6

```
root@monitoring02:/home/mug# ping6 -f 2001:1702:6:1111::31 -c 10000
PING 2001:1702:6:1111::31(2001:1702:6:1111::31) 56 data bytes

--- 2001:1702:6:1111::31 ping statistics ---
10000 packets transmitted, 10000 received, 0% packet loss, time 5291ms
rtt min/avg/max/mdev = 0.240/0.479/13.292/0.364 ms, pipe 2, ipg/ewma 0.529/0.404 m
```

- ⇨ Again, IPv6 flow processing faster (-;

Performance – Ping Flooding – to Zone TRUST (2/2)

- NAT64

```
root@monitoring02:/home/mug# ping -f 10.1.234.131 -c 10000
PING 10.1.234.131 (10.1.234.131) 56(84) bytes of data.

--- 10.1.234.131 ping statistics ---
10000 packets transmitted, 10000 received, 0% packet loss, time 11928ms
rtt min/avg/max/mdev = 0.858/1.143/14.059/0.549 ms, pipe 2, ipg/ewma 1.192/1.552 ms
```

- Comparison
  - IPv4 Native: 8600ms
  - IPv6 Native: 5291ms
  - NAT64: 11928ms (+ 39% / +125%)

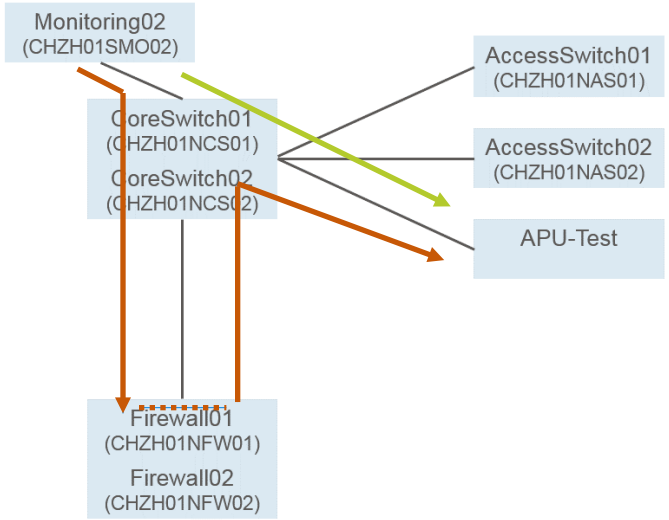

Performance – Observing SRX Load (1/4)

- Juniper SRX Branch Architecture
  - No dedicated hardware for control and data plan (all on same processor)
    - Dedicated core(s) for control plane
    - Dedicated core(s) for data plane
  - Control plane
    - Show chassis routing-engine
  - Data plane
    - Show chassis forwarding

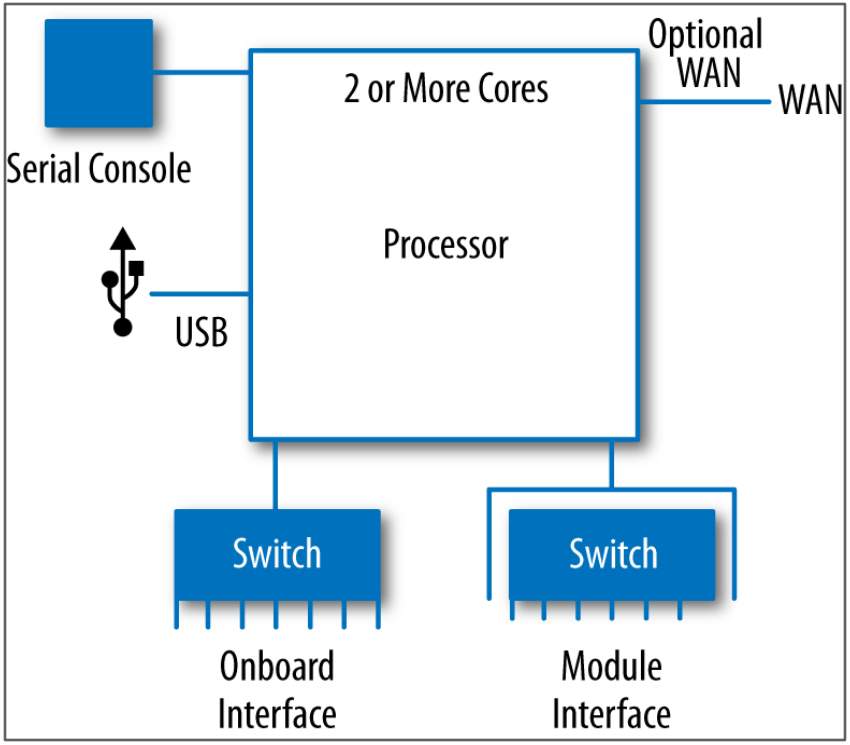

Source: http://chimera.labs.oreilly.com

Performance – Observing SRX Load (2/4)

- Control plane

```
mug@CHZH01NFWCL01> show chassis routing-engine
node0:
--------------------------------------------------------------------------
Routing Engine status:
    Temperature                 43 degrees C / 109 degrees F
    CPU temperature             43 degrees C / 109 degrees F
    Total memory              2048 MB Max  1311 MB used ( 64 percent)
      Control plane memory    1072 MB Max   600 MB used ( 56 percent)
      Data plane memory        976 MB Max   703 MB used ( 72 percent)
    CPU utilization:
      User                      18 percent
      Background                 0 percent
      Kernel                     9 percent
      Interrupt                  0 percent
      Idle                      73 percent
    Model                        RE-SRX240H2
    Serial ID                    ACMK1932
    Start time                   2016-03-01 08:15:25 CET
    Uptime                       2 days, 1 hour, 33 minutes, 44 seconds
    Last reboot reason           Router rebooted after a normal shutdown.
    Load averages:               1 minute   5 minute  15 minute
                                     0.22       0.34       0.32
```

Performance – Observing SRX Load (3/4)

- Data plane

```
mug@CHZH01NFWCL01> show chassis forwarding
node0:
--------------------------------------------------------------------------
FWDD status:
  State                                 Online
  Microkernel CPU utilization          12 percent
  Real-time threads CPU utilization     1 percent
  Heap utilization                     73 percent
  Buffer utilization                    1 percent
  Uptime:                               16 days, 7 minutes, 16 seconds
```

Performance – Observing SRX Load (4/4)

- Security related processing (sub process / thread of data plane)

```
mug@CHZH01NFWCL01> show security monitoring fpc 0
node0:
--------------------------------------------------------------------------
FPC 0
  PIC 0
    CPU utilization           :   16 %
    Memory utilization        :   73 %
    Current flow session      : 6877
    Current flow session IPv4: 1133538
    Current flow session IPv6: 4293840635
    Max flow session          : 409600
Total Session Creation Per Second (for last 96 seconds on average): 1015
IPv4  Session Creation Per Second (for last 96 seconds on average): 1015
IPv6  Session Creation Per Second (for last 96 seconds on average):    0
```

Performance – Observing SRX Load – to Zone DMZ (1/2)

- Firewall load
  - Native IPv4

```
root@monitoring02:/home/mug# ping -f 10.1.226.131 -c 100000
```

  - Native IPv6

```
root@monitoring02:/home/mug# ping6 -f 2001:1702:6:1131::31 -c 100000
```

  - NAT64

```
root@monitoring02:/home/mug# ping -f 10.1.226.131 -c 100000
```

Performance – Observing SRX Load – to Zone DMZ (2/2)

- Forwarding plane
  - Native IPv4

```
Microkernel CPU utilization        +2-3%
CPU utilization                    +16%  // sub process
Session Creation Per Second (for last 96 seconds on average): ~1000
```

  - Native IPv6

```
Microkernel CPU utilization        +2-3%
CPU utilization                    +22%  // sub process
Session Creation Per Second (for last 96 seconds on average): ~1000
```

  - NAT64

```
Microkernel CPU utilization        +3-4%
CPU utilization                    +22%  // sub process
Session Creation Per Second (for last 96 seconds on average): ~850
```

  - Comment: SRX240h2 should be able to handle 9600 new sessions / second

Performance – iperf Testing – to Zone DMZ

- Control Plane
  - No significant increase

- Overall

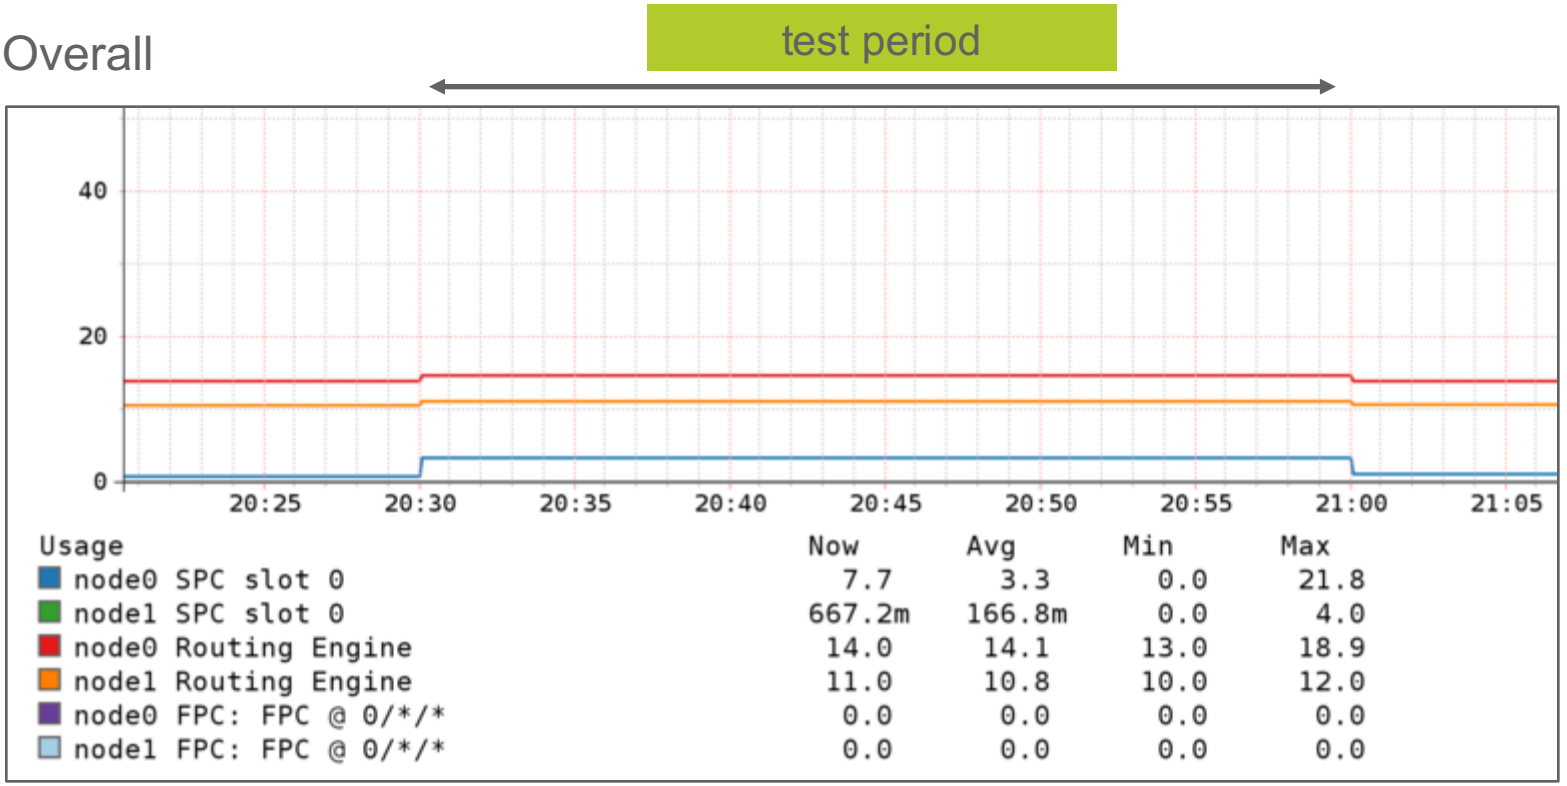

Comparing monitored values (CLI and GUI)

- Monitoring issue (?)

```
mug@CHZH01NFWCL01> show chassis routing-engine
node0:
--------------------------------------------------------------------------
Routing Engine status:
    Temperature                 43 degrees C / 109 degrees F
    CPU temperature             44 degrees C / 111 degrees F
    Total memory              2048 MB Max  1311 MB used ( 64 percent)
      Control plane memory    1072 MB Max   590 MB used ( 55 percent)
      Data plane memory        976 MB Max   712 MB used ( 73 percent)
    CPU utilization:
      User                      10 percent
      Background                 0 percent
      Kernel                    16 percent
      Interrupt                  0 percent
      Idle                      73 percent
    Model                       RE-SRX240H2
    Serial ID                   ACMK1932
    Start time                  2016-02-07 20:47:00 CET
    Uptime                      14 days, 23 hours, 57 minutes, 41 seconds
    Last reboot reason          Router rebooted after a normal shutdown.
    Load averages:              1 minute   5 minute  15 minute
                                    0.16       0.23       0.24
```

Performance – iperf Measurement – to Zone DMZ (1/2)

- Native IPv4

```
mug@Monitoring01:~$ iperf -c 10.1.226.31 -p 443 -P 10
------------------------------------------------------------
Client connecting to 10.1.226.31, TCP port 443
TCP window size: 85.0 KByte (default)
------------------------------------------------------------
…
[SUM]  0.0-10.3 sec   121 MBytes  98.6 Mbits/sec
mug@Monitoring01:~$
```

- Native IPv6

```
mug@Monitoring01:~$ iperf -c 2001:1702:6:1131::31 -p 443 -P 10 -V
------------------------------------------------------------
Client connecting to 2001:1702:6:1131::31, TCP port 443
TCP window size: 85.0 KByte (default)
------------------------------------------------------------
…
[SUM]  0.0-10.0 sec   989 MBytes   828 Mbits/sec
mug@Monitoring01:~$
```

Performance – iperf Measurement – to Zone DMZ (2/2)

- NAT64

```
mug@Monitoring01:~$ iperf -c 10.1.226.131 -p 443 -P 10
------------------------------------------------------------
Client connecting to 10.1.226.131, TCP port 443
TCP window size: 85.0 KByte (default)
------------------------------------------------------------
…
[SUM]  0.0-10.4 sec   118 MBytes  95.8 Mbits/sec
mug@Monitoring01:~$
```

Performance – iperf Measurement – to Zone TRUST (1/2)

- Native IPv4

```
mug@monitoring02:/home/mug# iperf -c 10.1.224.31 -p 443 -P 10
------------------------------------------------------------
Client connecting to 10.1.224.31, TCP port 443
TCP window size: 85.0 KByte (default)
------------------------------------------------------------
…
[SUM]  0.0-11.1 sec   104 MBytes  79.2 Mbits/sec
mug@chbe01nmp01:/home/mug# iperf -s -p 443 -V
```

- Native IPv6

```
…
```

Performance – iperf Measurement – to Zone TRUST (2/2)

- NAT64

```
mug@monitoring02:/home/mug# iperf -c 10.1.234.131 -p 443 -P 10
------------------------------------------------------------
Client connecting to 10.1.234.131, TCP port 443
TCP window size: 85.0 KByte (default)
------------------------------------------------------------
…
[SUM]  0.0-10.5 sec   124 MBytes  99.1 Mbits/sec
mug@monitoring02:/home/mug#
```

Contents

Performance – Observing SRX Load – to Zone TRUST (1/2)

- Firewall load
  - Testing

```
root@monitoring02:/home/mug# ping -f 10.1.224.31 -c 100000
root@monitoring02:/home/mug# ping6 -f 2001:1702:6:1111::31 -c 100000
root@monitoring02:/home/mug# ping -f 10.1.234.131 -c 100000
```

  - Observing forwarding and control plane

```
mug@CHZH01NFWCL01> show chassis forwarding
mug@CHZH01NFWCL01> show security monitoring fpc 0
```

```
mug@CHZH01NFWCL01> show chassis routing-engine
```

Performance – Observing SRX Load – to Zone TRUST (2/2)

- Firewall load (continued)
  - Native IPv4
    - n/a
  - Native IPv6
    - n/a
  - NAT64

```
Microkernel CPU utilization          +7%
CPU utilization                      +20%  // sub process
Session Creation Per Second (for last 96 seconds on average): ~800
```

Demo

Juniper Space / Security Director

Contents

Our Summary

- What we tested
  - Management related traffic / applications
  - TCP / UDP / ICMP
- What we did not test
  - DNS
  - Multicast
- Does it work?
  - Yes
- Would we recommend to use it?
  - Yes
    - Allows you to introduce IPv6 islands
    - Enabler for IPv6 experiences

Your Opinion

- What do you think?
- Would you give it a try?
- Any other ideas where / how to make use of NAT64?

In case of further questions

AWK GROUP

Gabriel Müller
Dipl. El.-Ing. ETH
Senior Consultant

gabriel.mueller@awk.ch

Contents

Setup – NAT64 Configuration (SRX 240) – to Zone TRUST

- Static NAT for destination address

```
mug@CHZH01NFWCL01# show security nat static
rule-set RS_NAT6446_from_reth21013 {
    from interface reth2.1013;
     rule R_004_TRUST {
        match {
            destination-address-name 10.1.234.121/32;
        }
        then {
            static-nat {
                prefix-name {
                    2001:1702:6:1111::21/128;
                }
            }
        }
    }
}
```

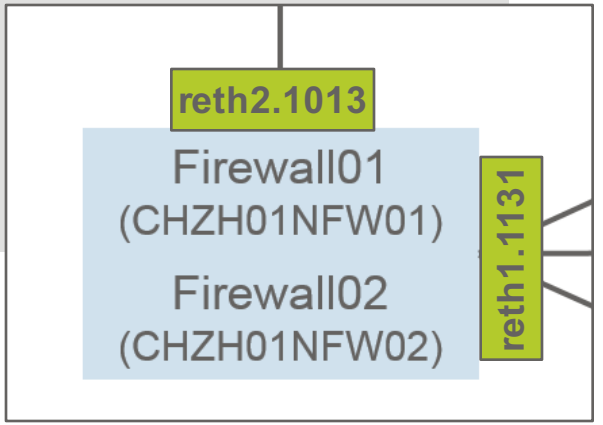

Setup – NAT64 Configuration (SRX 240) – to Zone TRUST

- Changing source IP

```
mug@CHZH01NFWCL01# show security nat source
rule-set RS_NAT6446_from_reth21013 {
    from interface reth2.1013;
    to interface reth2.1013;
    rule R_004_TRUST {
        match {
            source-address-name CHZH01_MGMT_10.1.233.0/24;
            destination-address-name CHZH01_MGMT_2001:1702:6:1111/64;
        }
        then {
            source-nat {
                interface;
            }
        }
    }
}
```

Incoming interface equals outgoing interface since we looping through firewall

```
family inet6 {
                address 2001:1702:6:1131::10/64;
                address fe80::1131:0:0:10/64;
            }
```

Setup – NAT64 Configuration (SRX 240) – from Zone TRUST

- Static NAT for destination address

```
mug@CHZH01NFWCL01# show security nat static
rule-set RS_NAT6446_from_reth21013 {
    from interface reth2.1013;
    rule R_099_ToSyslog {
        match {
            destination-address-name 2001:1702:6:11a1::1122/128;
        }
        then {
            static-nat {
                prefix-name {
                    CHZH01_SMO02_v4-10.1.233.22/32;
                }
            }
        }
    }
}
```

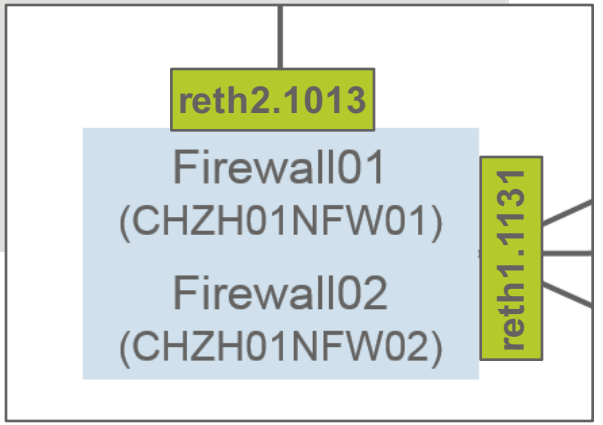

Setup – NAT64 Configuration (SRX 240) – from Zone TRUST

- Changing Source IP
  - Not required
  - Static NAT entry is used (applied in reversed order)

```
mug@CHZH01NFWCL01# show security nat static
rule-set RS_NAT6446_from_reth21013 {
    from interface reth2.1013;
    rule R_001_DMZ {
        match {
            destination-address-name 10.1.226.121/32;
        }
        then {
            static-nat {
                prefix-name {
                    2001:1702:6:1131::21/128;
                }
            }
        }
    }
}
```

NAT64 Config – Junos Security Director View

- Firewall Policies

MGMT_MGMT_007-TEMP	MGMT	CHZH01_SMO02_v4-10.1.233.22/32 CHZH01_SMO01_v4-10.1.233.21/32	MGMT	2001:1702:6:1131::21/128 2001:1702:6:1131::22/128 2001:1702:6:1131::31/128	snmp ssh ping https	Permit
MGMT_MGMT_008-TEMP	MGMT	CHZH01_SMO02_v4-10.1.233.22/32 CHZH01_SMO01_v4-10.1.233.21/32	MGMT	2001:1702:6:1111::21/128 2001:1702:6:1111::22/128 2001:1702:6:1111::31/128	snmp ssh ping https	Permit
MGMT_MGMT_009-TEMP	MGMT	2001:1702:6:1131::21/128 2001:1702:6:1131::22/128 2001:1702:6:1131::31/128	MGMT	CHZH01_SMO02_v4-10.1.233.22/32	ssh syslog pingv6 https	Permit
MGMT_MGMT_010-TEMP	MGMT	2001:1702:6:1111::21/128 2001:1702:6:1111::22/128 2001:1702:6:1111::31/128	MGMT	CHZH01_SMO02_v4-10.1.233.22/32	ssh syslog pingv6 https	Permit

NAT64 Config – Junos Security Director View

- NAT Policies

RS_NAT46_MGMT_DMZ_Source (18 - 18)										
R_001_DMZ	SOURCE	**Interfaces:** reth2.1013	CHZH01_MGMT_10.1.233.0/24		**Interfaces:** reth1.1131	CHZH01_MGMT_2001:1702:6:1131/64			**Interface**	Not Applicable
RS_NAT46_MGMT_TRUST_Source (19 - 19)										
R_004_TRUST	SOURCE	**Interfaces:** reth2.1013	CHZH01_MGMT_10.1.233.0/24		**Interfaces:** reth2.1013	CHZH01_MGMT_2001:1702:6:1111/64			**Interface**	Not Applicable
RS_NAT6446_from_reth11131 (20 - 20)										
R_001_ToSyslog	STATIC	**Interfaces:** reth1.1131	empty		Not Applicable	2001:1702:6:1191::1122/128		Not A	Not Ap...	CHZH01_SMO02_v4-10.1.233.22/32
RS_NAT6446_from_reth21013 (21 - 27)										
R_001_DMZ	STATIC	**Interfaces:** reth2.1013	empty		Not Applicable	10.1.226.121/32		Not A	Not Ap...	2001:1702:6:1131::21/128
R_002_DMZ	STATIC	**Interfaces:** reth2.1013	empty		Not Applicable	10.1.226.122/32		Not A	Not Ap...	2001:1702:6:1131::22/128
R_003_DMZ	STATIC	**Interfaces:** reth2.1013	empty		Not Applicable	10.1.226.131/32		Not A	Not Ap...	2001:1702:6:1131::31/128
R_004_TRUST	STATIC	**Interfaces:** reth2.1013	empty		Not Applicable	10.1.234.121/32		Not A	Not Ap...	2001:1702:6:1111::21/128
R_005_TRUST	STATIC	**Interfaces:** reth2.1013	empty		Not Applicable	10.1.234.122/32		Not A	Not Ap...	2001:1702:6:1111::22/128
R_006_TRUST	STATIC	**Interfaces:** reth2.1013	empty		Not Applicable	10.1.234.131/32		Not A	Not Ap...	2001:1702:6:1111::31/128
R_099_ToSyslog	STATIC	**Interfaces:** reth2.1013	empty		Not Applicable	2001:1702:6:11a1::1122/128		Not A	Not Ap...	CHZH01_SMO02_v4-10.1.233.22/32