LONAP™

Network Automation at an Internet Exchange Point

Ian Chilton

UKNOF 52

28th September 2023 | London

Recently...

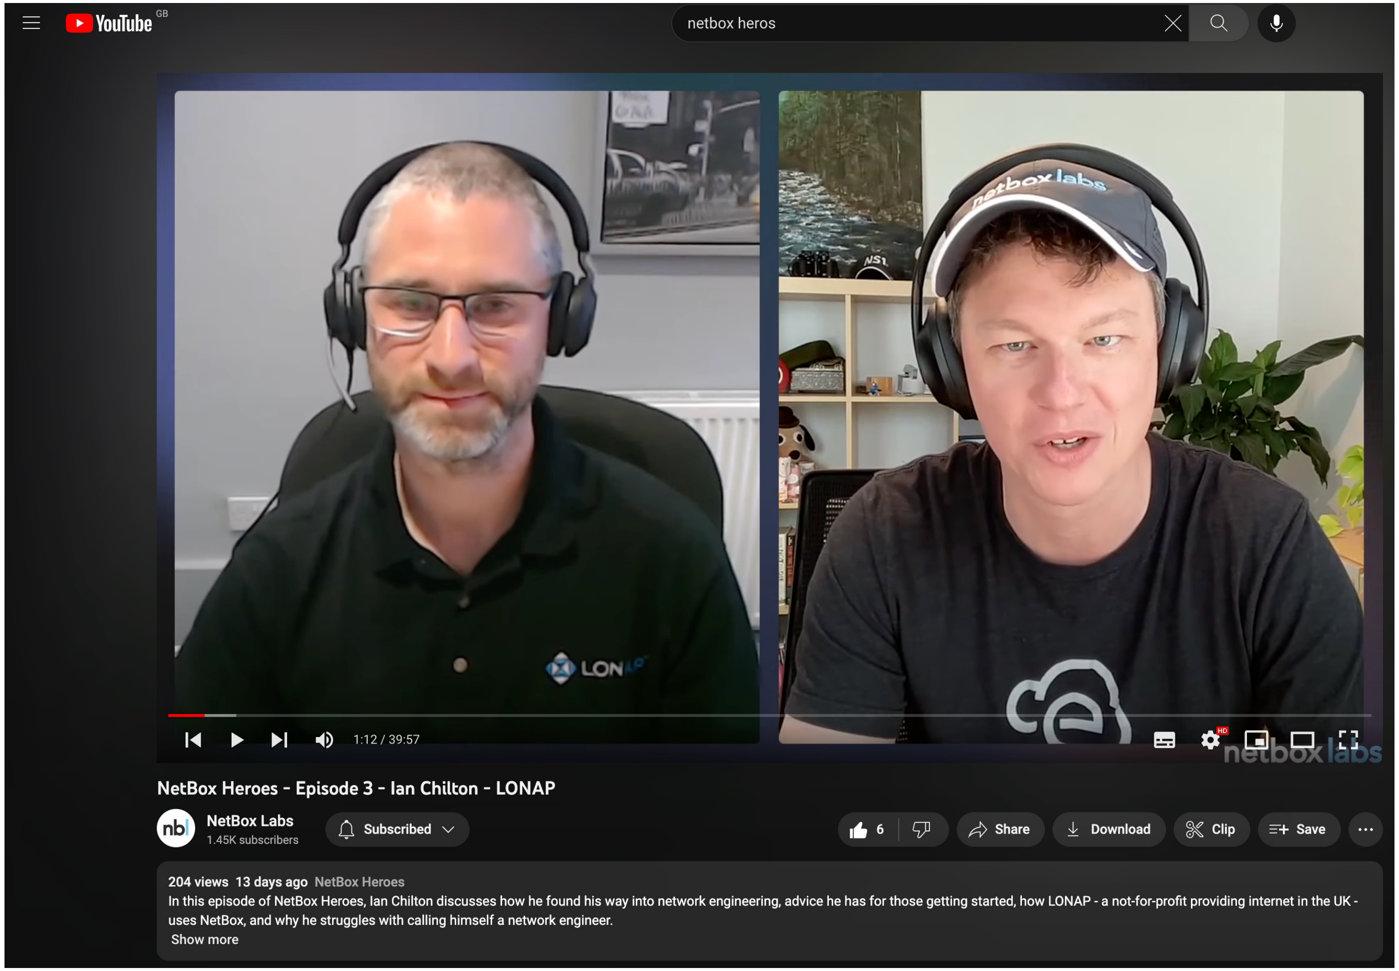

NetBox Heroes - Episode 3 - Ian Chilton - LONAP

204 views 13 days ago NetBox Heroes
In this episode of NetBox Heroes, Ian Chilton discusses how he found his way into network engineering, advice he has for those getting started, how LONAP - a not-for-profit providing internet in the UK - uses NetBox, and why he struggles with calling himself a network engineer.
Show more

LONAP

- Not-for-profit, neutral Internet Exchange Point (IXP).
- Established in 1997.
- Members are all equal stakeholders.
- 8 Points of Presence across 7 of the major London data centres.
- ~250 members.
- ~350 connected ports.
- 1Tbps, ~130Mpps peak traffic.
- 4x full time staff.
- Board of elected non-exec directors.

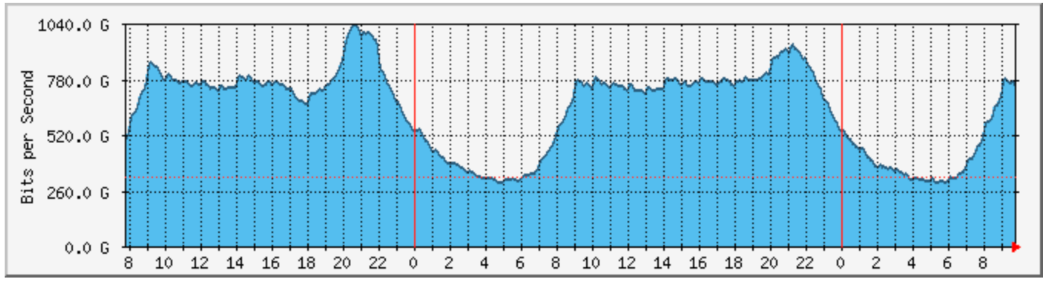

	Max	Average	Current
In	1029.6 Gb/s	634.8 Gb/s	750.6 Gb/s
Out	1030.5 Gb/s	634.9 Gb/s	750.5 Gb/s

- Connected networks are a mix of Internet Service Providers, content/hosting networks and CDNs. They include: Google, Netflix, Microsoft, BBC, Apple, Twitter, TikTok, eBay, Sky, Vodafone.
- Full list: https://lonap.net/members

Points of Presence

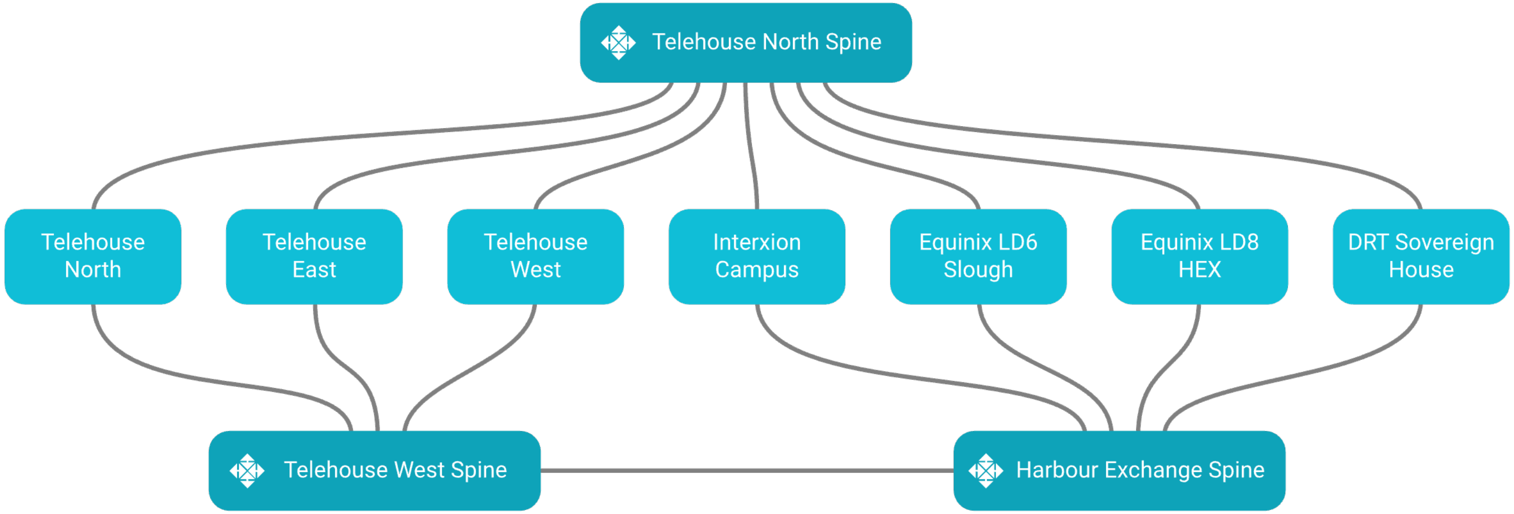

- **TELEHOUSE NORTH BUILDING**
  - **TELEHOUSE NORTH 2**
- **TELEHOUSE EAST BUILDING**
- **TELEHOUSE WEST BUILDING**

- **EQUINIX LD8, 6-7 AND 8-9 HARBOUR EXCHANGE**
- **INTERXION LONDON (SOVEREIGN HOUSE)**
  - **INTERXION LONDON (CLOUD HOUSE)**
- **EQUINIX LD6, SLOUGH**
  - **EQUINIX LD4, LD5, LD7 & LD10, SLOUGH**
- **INTERXION LON1 CAMPUS HANBURY STREET**
  - **INTERXION LON2, LON3**

Points of Presence

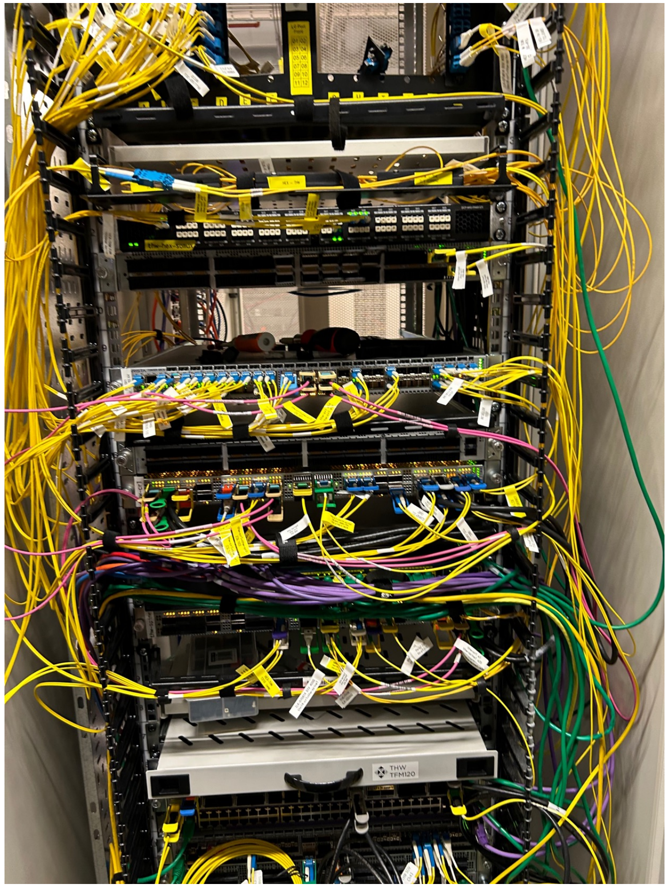

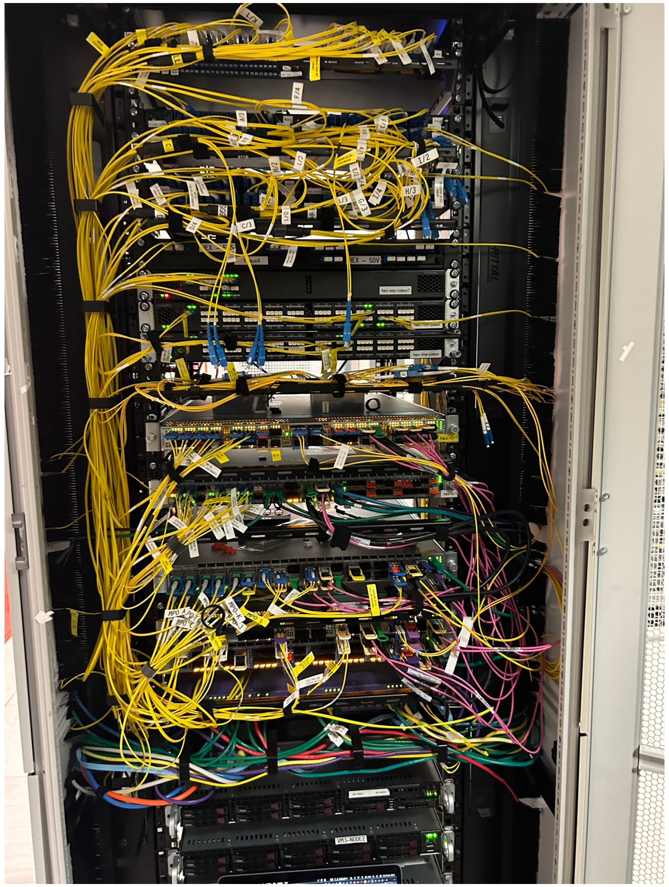

LONAP

IXP Network
20230609 WH

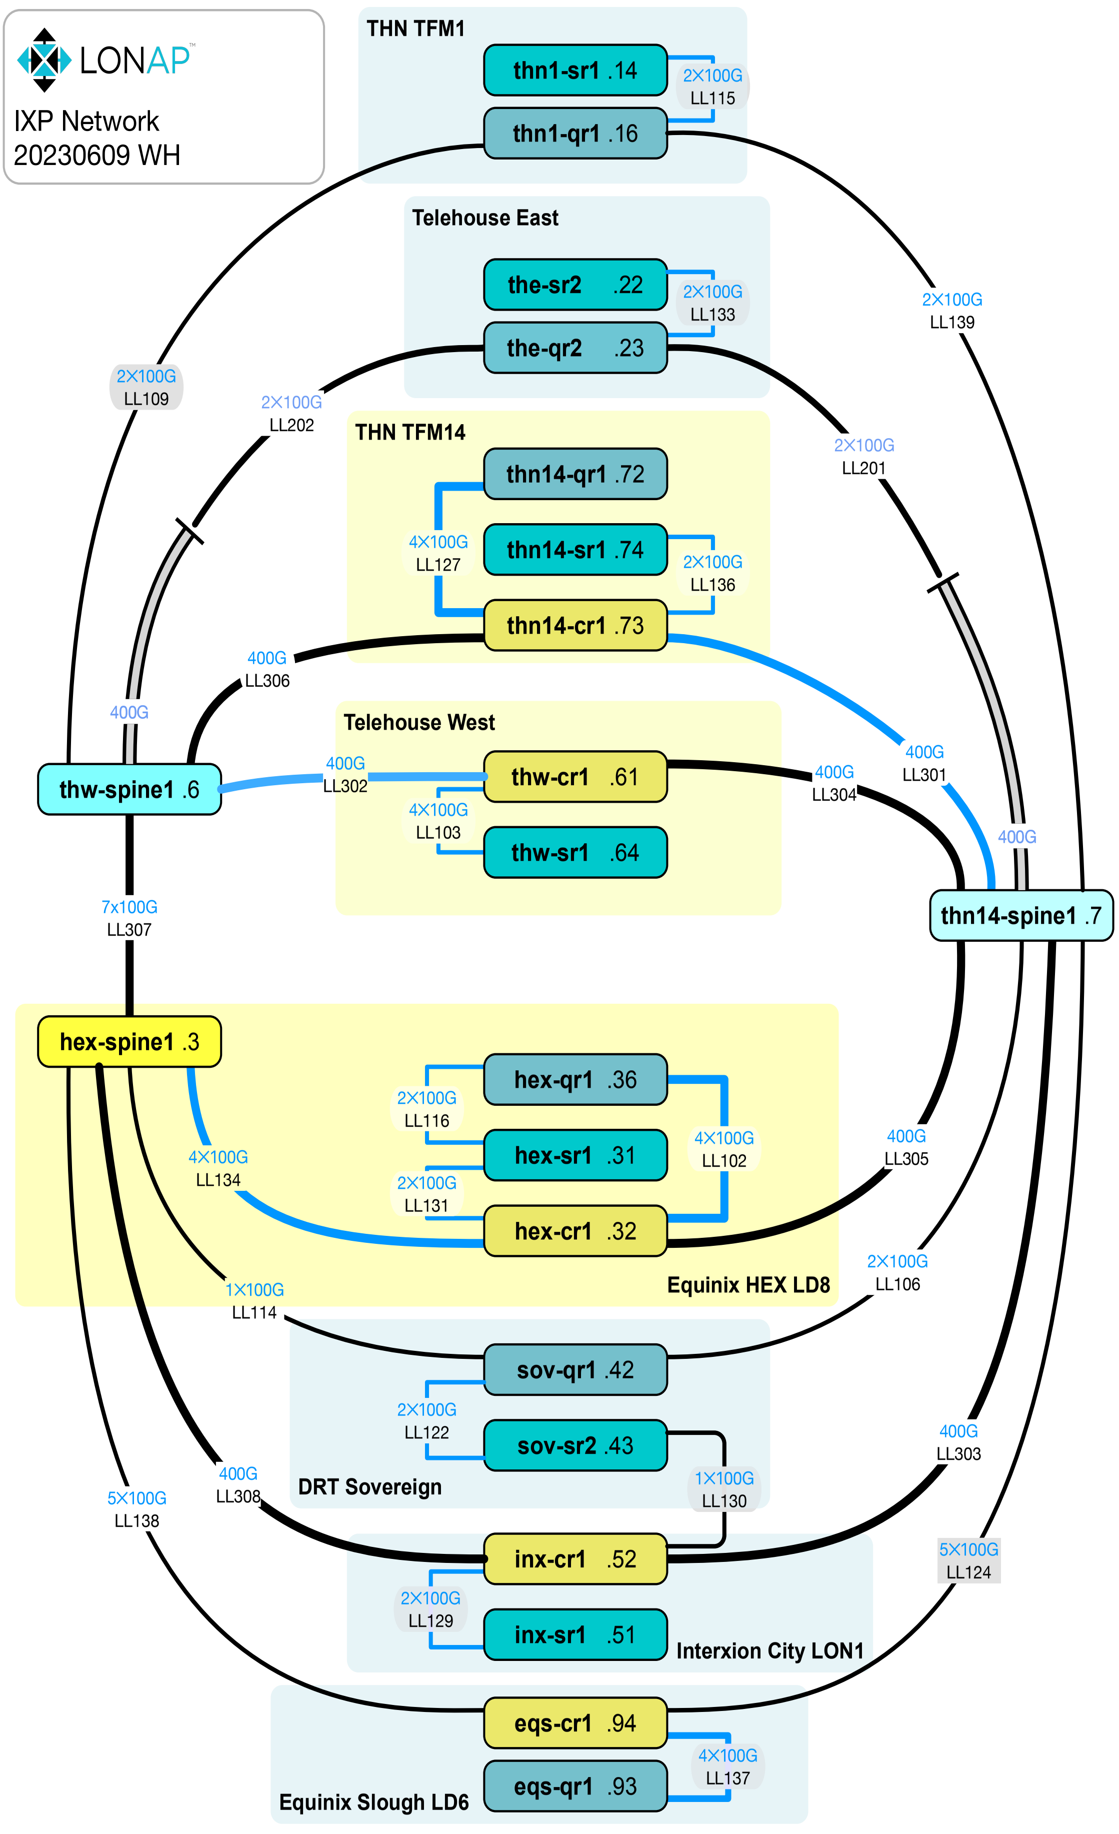

LONAP

LONAP

Management LAN
20230609

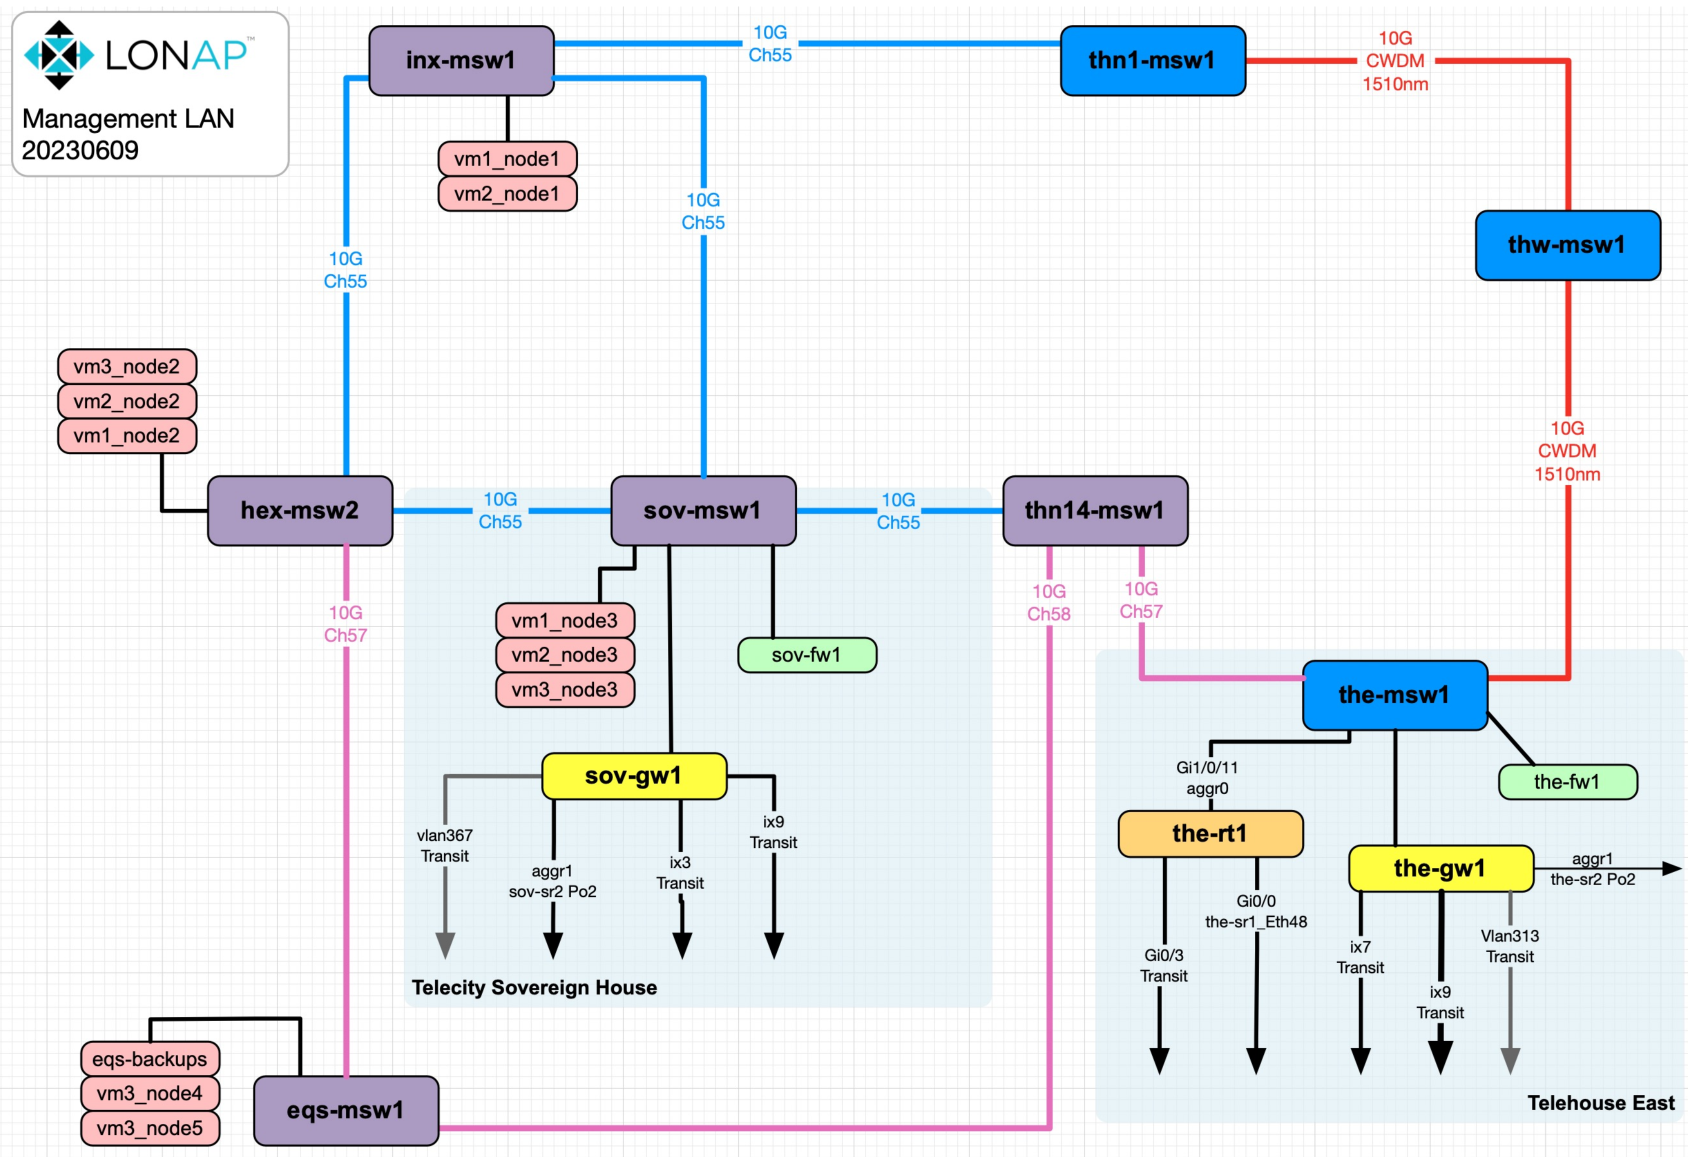

inx-msw1
vm1_node1
vm2_node1
10G
Ch55
thn1-msw1
10G
CWDM
1510nm
10G
Ch55
10G
Ch55
thw-msw1
vm3_node2
vm2_node2
vm1_node2
10G
CWDM
1510nm
hex-msw2
10G
Ch55
sov-msw1
10G
Ch55
thn14-msw1
10G
Ch57
vm1_node3
vm2_node3
vm3_node3
sov-fw1
10G
Ch58
10G
Ch57
the-msw1
sov-gw1
vlan367
Transit
aggr1
sov-sr2 Po2
ix3
Transit
ix9
Transit
Telecity Sovereign House
Gi1/0/11
aggr0
the-rt1
Gi0/3
Transit
Gi0/0
the-sr1_Eth48
the-fw1
the-gw1
aggr1
the-sr2 Po2
ix7
Transit
ix9
Transit
Vlan313
Transit
Telehouse East
eqs-backups
vm3_node4
vm3_node5
eqs-msw1

LONAP

“DevOps” Engineer

- Operations.
- Software development.
- Network Engineering.
- System administration (10x servers, >100 VMs).
- Management of many internal and member facing systems.
- Maintenance - physical infrastructure changes, software upgrades, configuration rollout.
- Deployment.
- Configuration management.
- Automation.
- Monitoring.
- Data centre engineering.
- Handling alerts – 24/7/365.
- Member support (where needed).

Network Automation

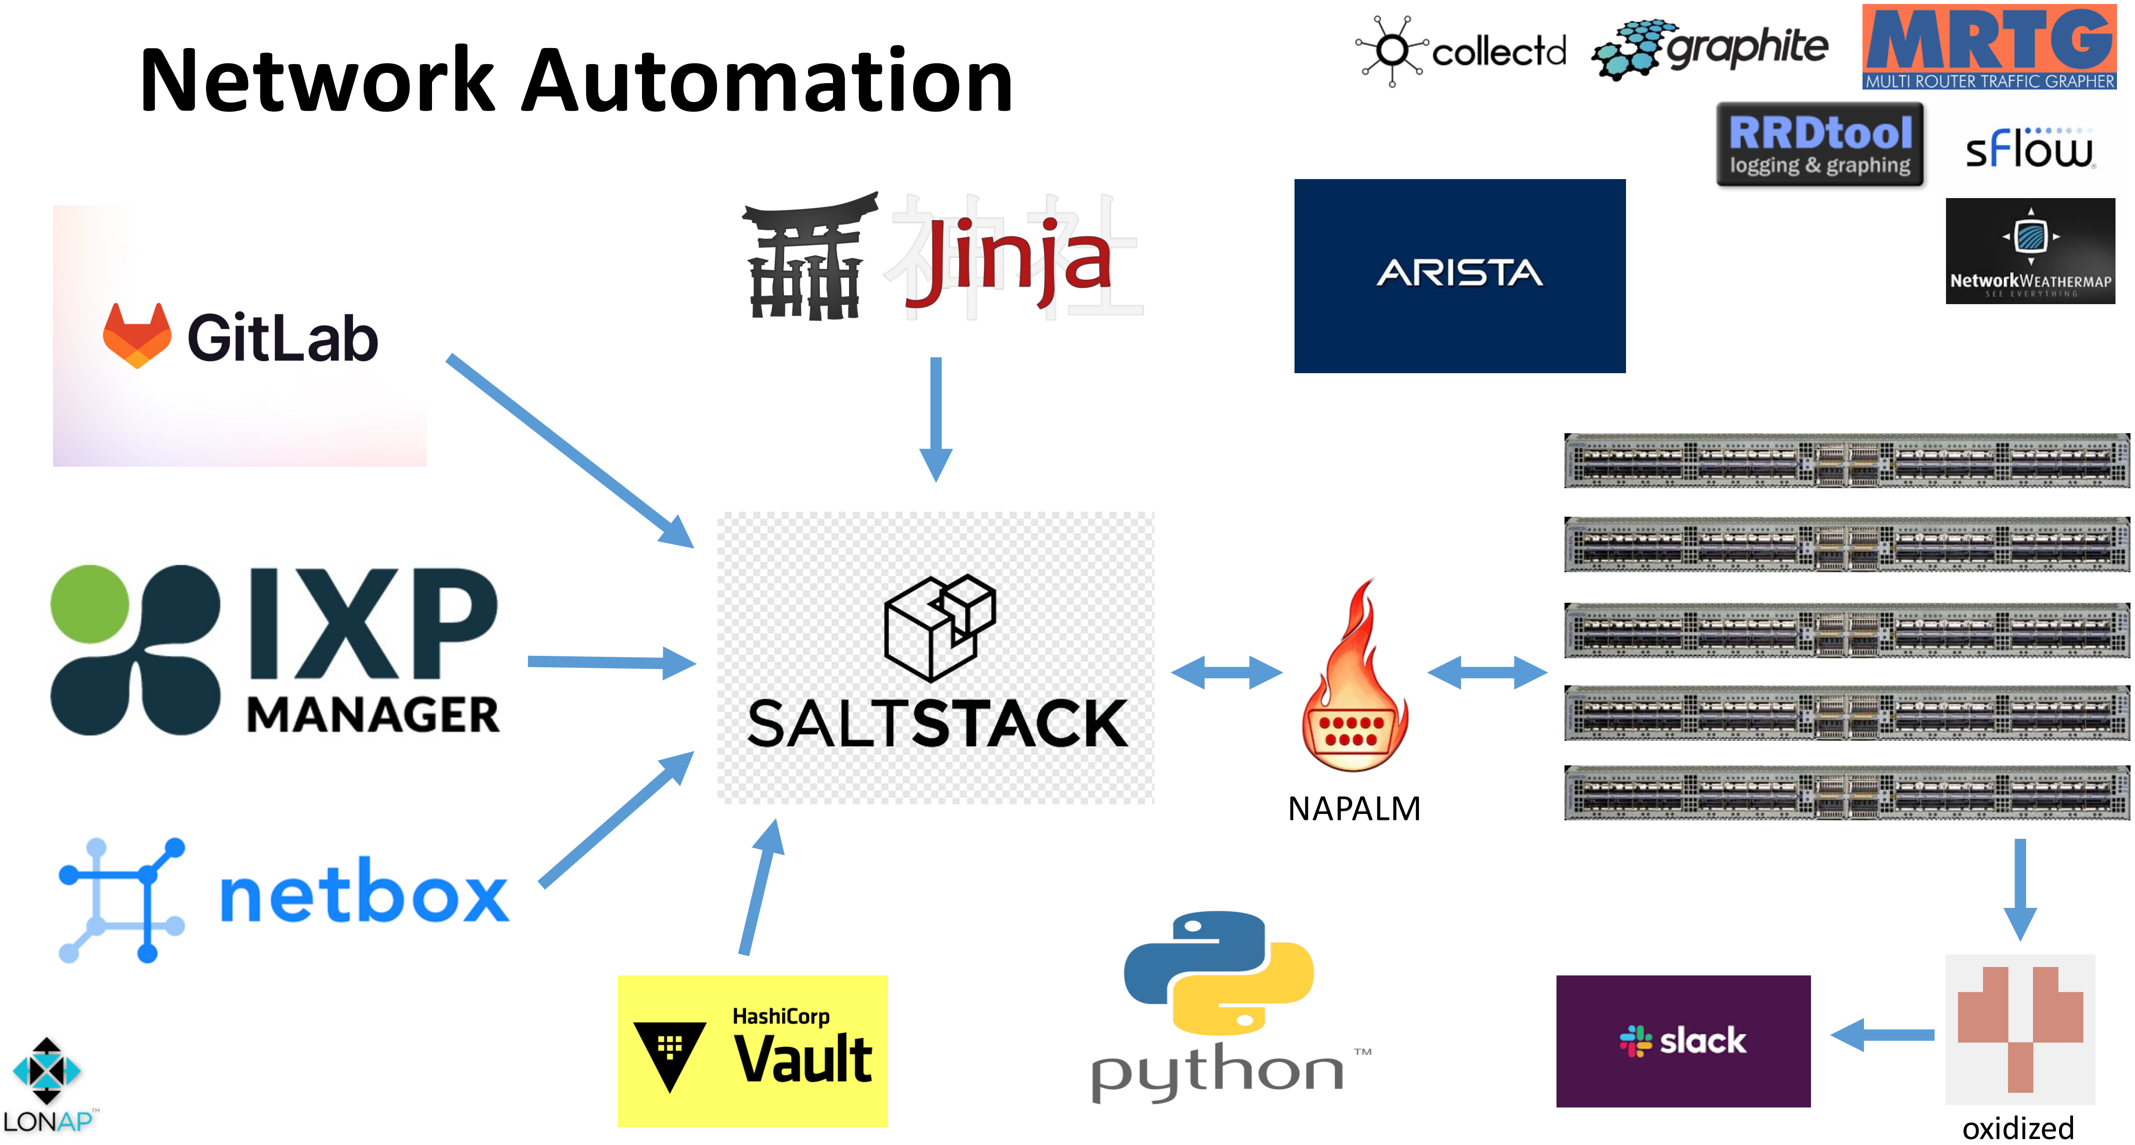

NAPALM

oxidized

Existing Network or New Deployment..?

Which is best?

Existing:

- Lot of work to (hopefully) generate the same configuration that’s already in place.
- Good time to review configuration.
- Potential inconsistencies from years of manual configuration.
- Scary deployment – can’t break network!

New network (“green field”):

- Can deploy without worrying about breaking anything.
- Testing/debugging/fixing configuration at the same time as working on automation.

SaltStack

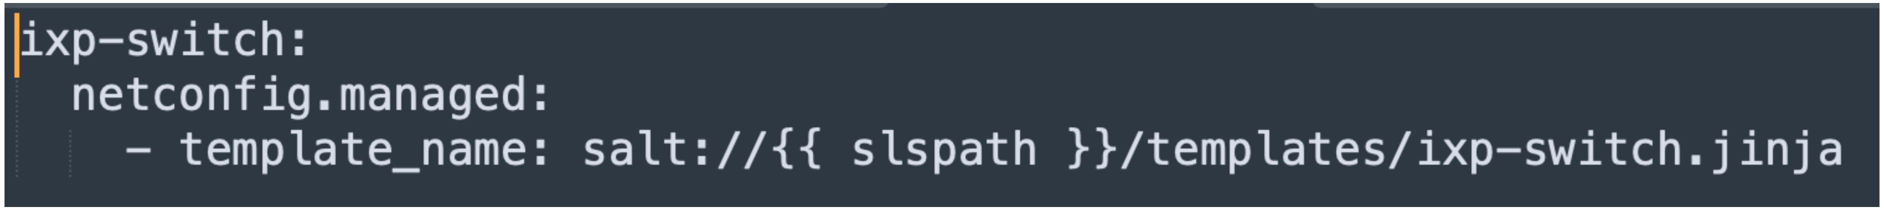

```
ixp-switch:
  netconfig.managed:
    - template_name: salt://{{ slspath }}/templates/ixp-switch.jinja
```

All logic built in to the templates.

This became beneficial later on!

SaltStack: Grains

“facts”

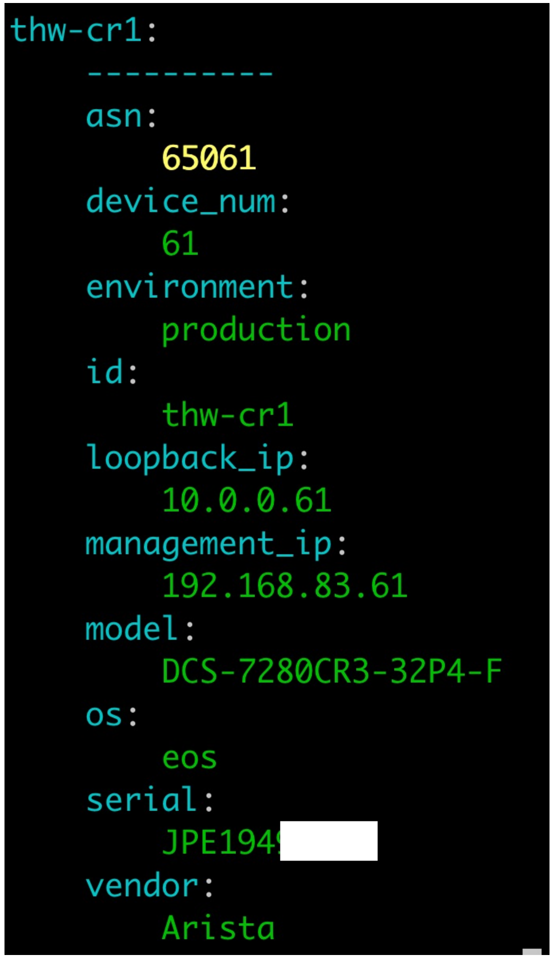

```
thw-cr1:
    ----------
    asn:
        65061
    device_num:
        61
    environment:
        production
    id:
        thw-cr1
    loopback_ip:
        10.0.0.61
    management_ip:
        192.168.83.61
    model:
        DCS-7280CR3-32P4-F
    os:
        eos
    serial:
        JPE194[redacted]
    vendor:
        Arista
```

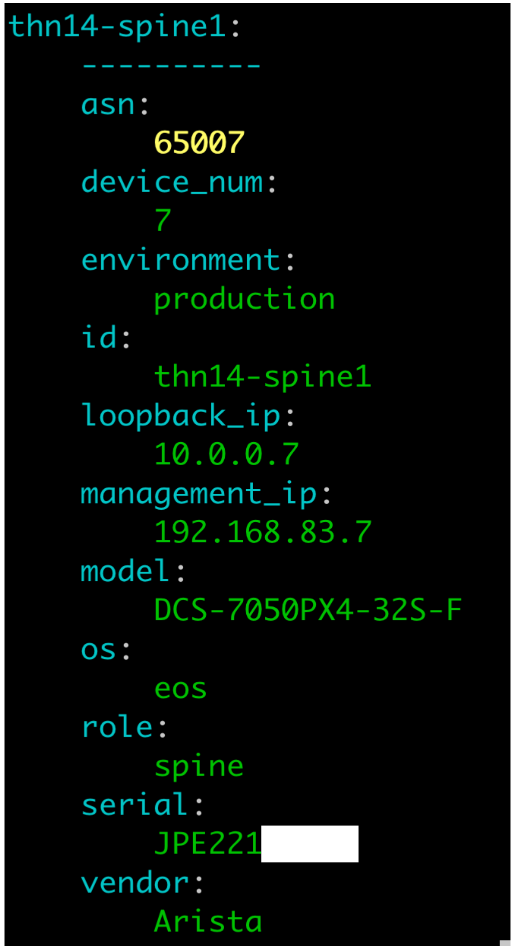

```
thn14-spine1:
    ----------
    asn:
        65007
    device_num:
        7
    environment:
        production
    id:
        thn14-spine1
    loopback_ip:
        10.0.0.7
    management_ip:
        192.168.83.7
    model:
        DCS-7050PX4-32S-F
    os:
        eos
    role:
        spine
    serial:
        JPE221[redacted]
    vendor:
        Arista
```

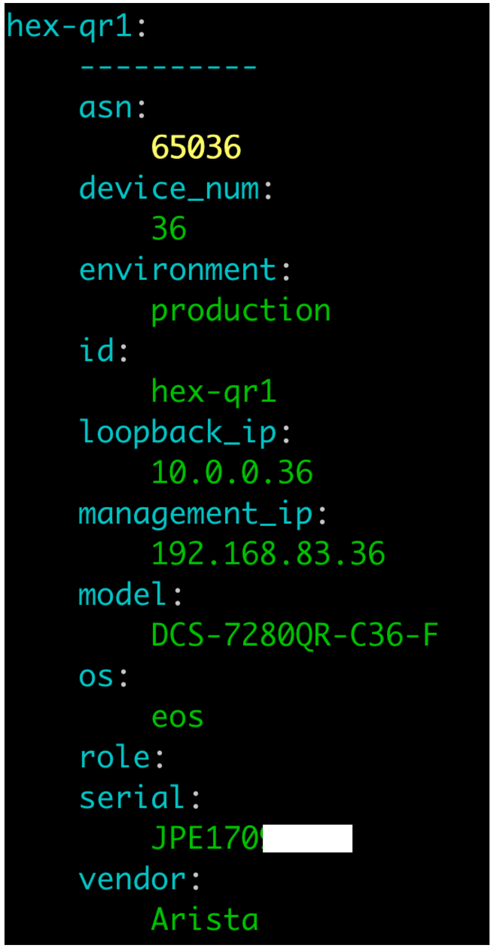

```
hex-qr1:
    ----------
    asn:
        65036
    device_num:
        36
    environment:
        production
    id:
        hex-qr1
    loopback_ip:
        10.0.0.36
    management_ip:
        192.168.83.36
    model:
        DCS-7280QR-C36-F
    os:
        eos
    role:
    serial:
        JPE170[redacted]
    vendor:
        Arista
```

SaltStack: Pillar

“configuration”

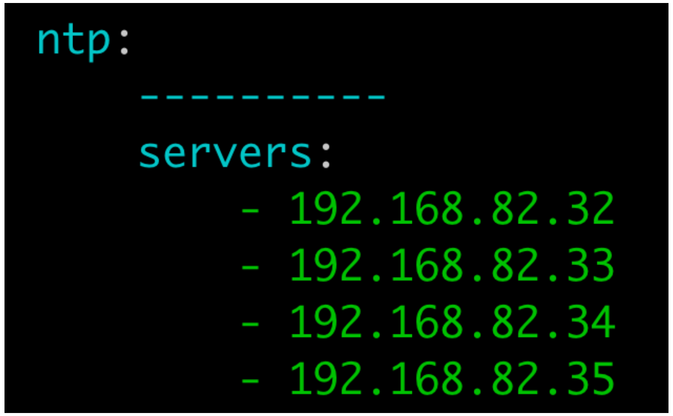

```
ntp:
    ----------
    servers:
        - 192.168.82.32
        - 192.168.82.33
        - 192.168.82.34
        - 192.168.82.35
```

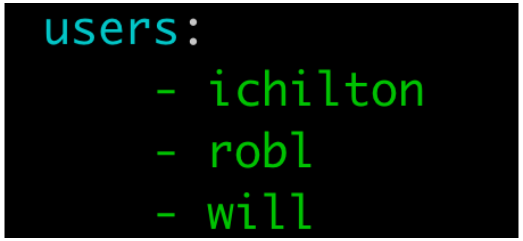

```
users:
        - ichilton
        - robl
        - will
```

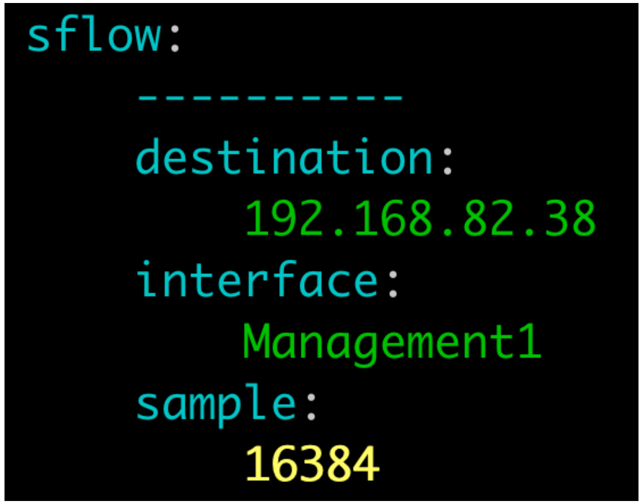

```
sflow:
    ----------
    destination:
        192.168.82.38
    interface:
        Management1
    sample:
        16384
```

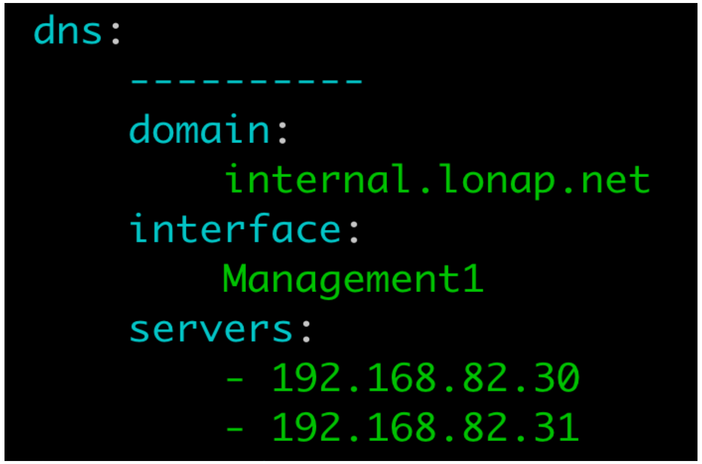

```
dns:
    ----------
    domain:
        internal.lonap.net
    interface:
        Management1
    servers:
        - 192.168.82.30
        - 192.168.82.31
```

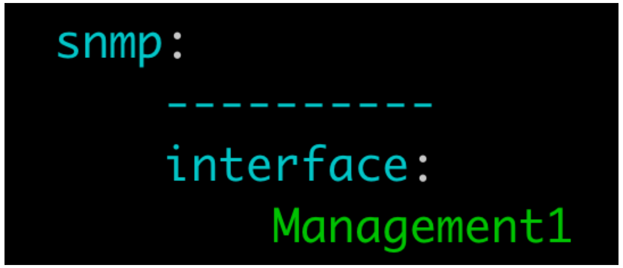

```
snmp:
    ----------
    interface:
        Management1
```

SaltStack: Pillar

This is merged…

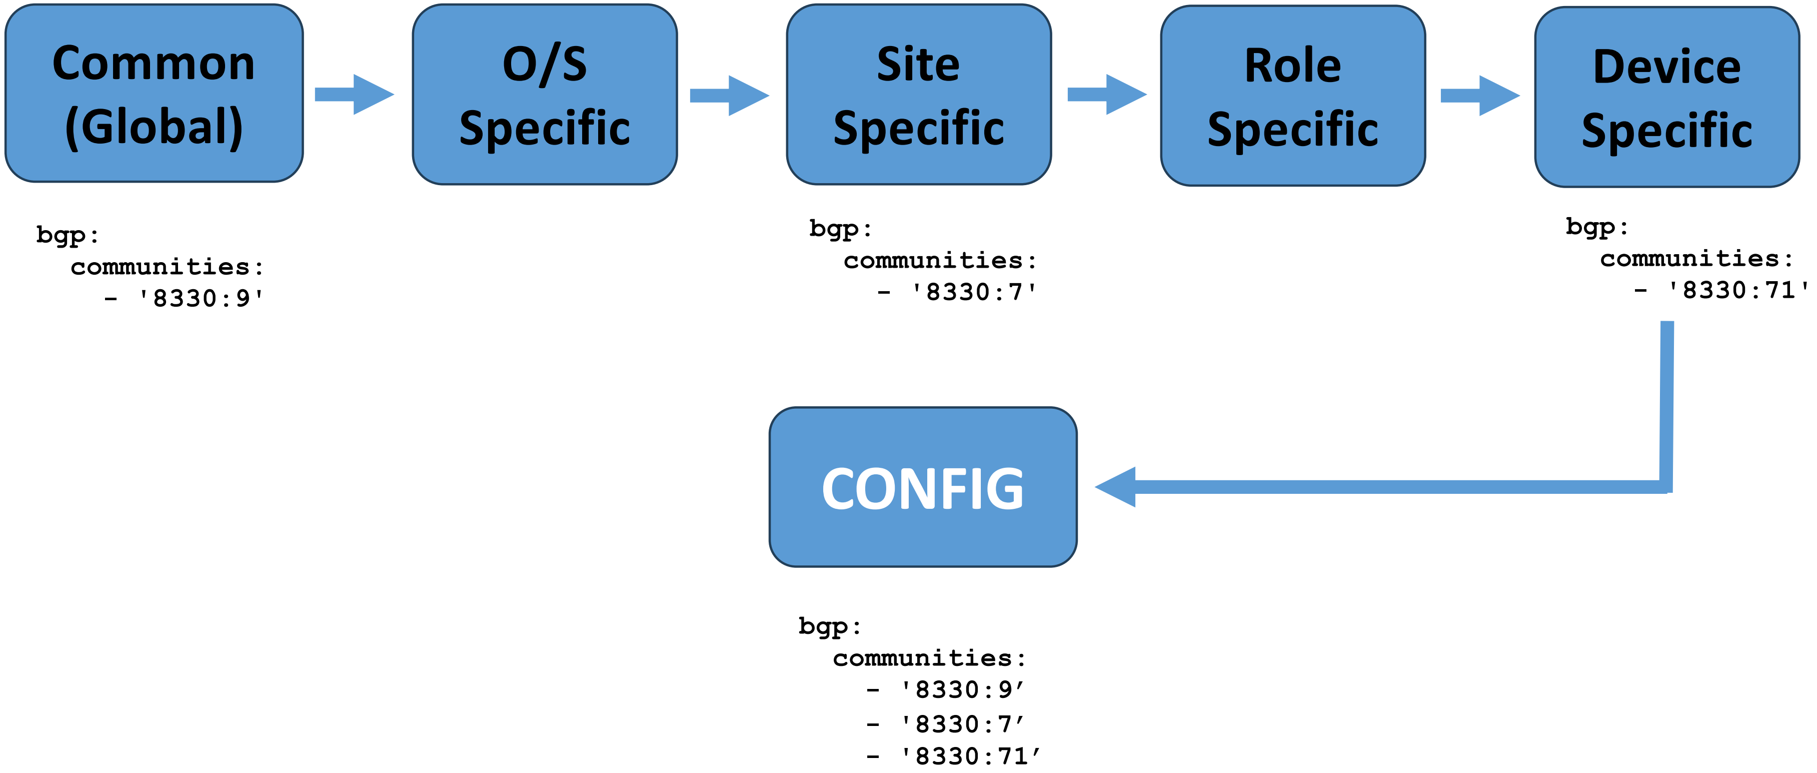

Common (Global) → O/S Specific → Site Specific → Role Specific → Device Specific

```
bgp:
  communities:
    - '8330:9'
```

```
bgp:
  communities:
    - '8330:7'
```

```
bgp:
  communities:
    - '8330:71'
```

CONFIG

```
bgp:
  communities:
    - '8330:9’
    - '8330:7’
    - '8330:71’
```

SaltStack: Salt Mine

Dynamic information about other devices.

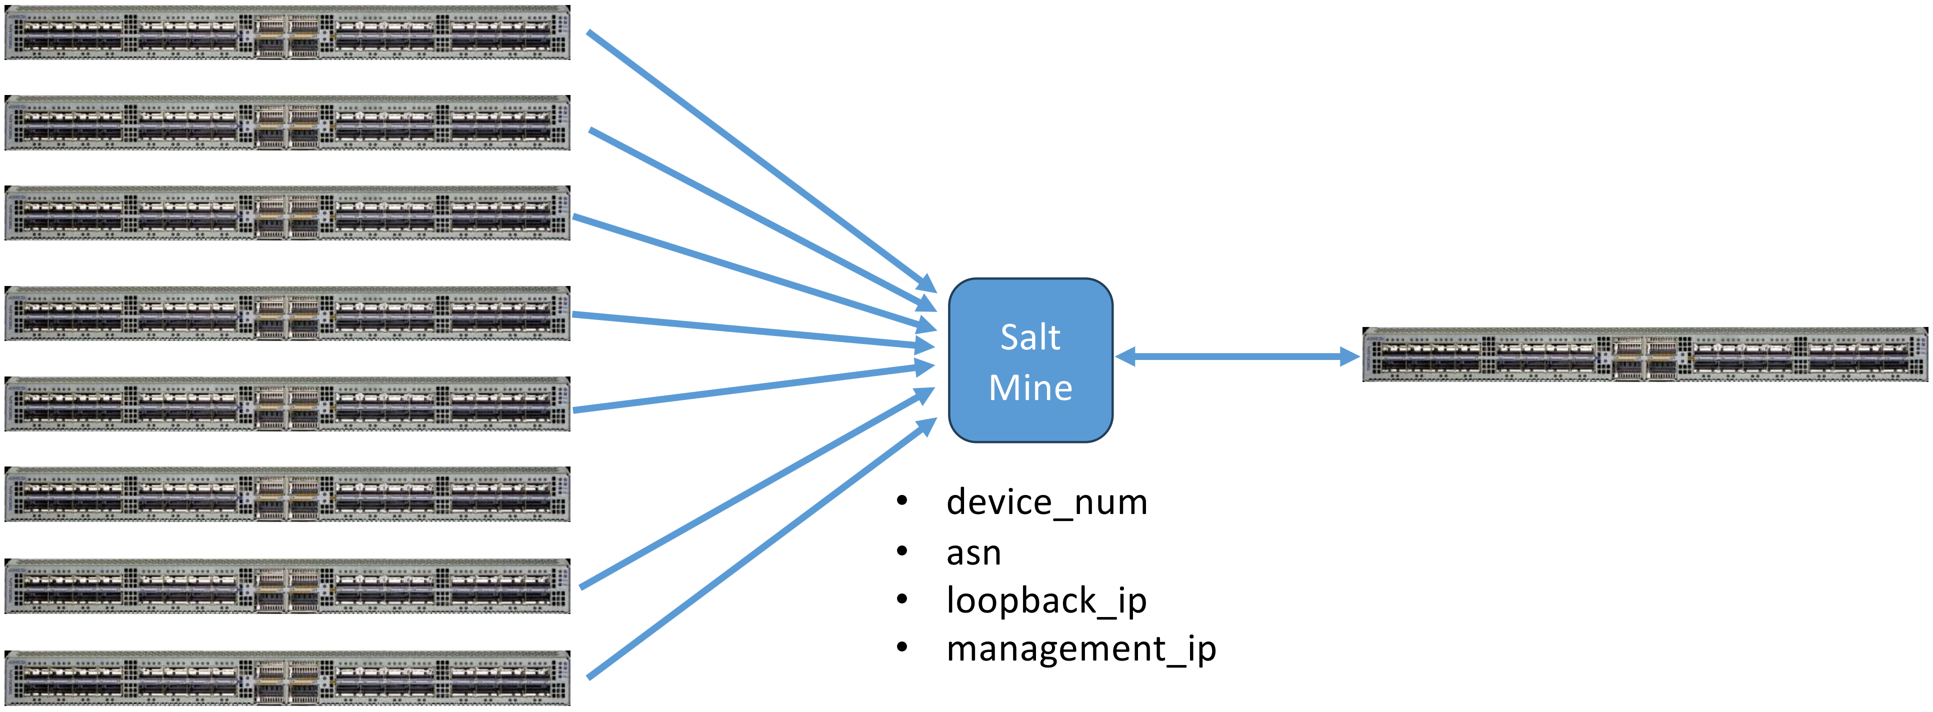

- device_num
- asn
- loopback_ip
- management_ip

SaltStack: Jinja2 Templates

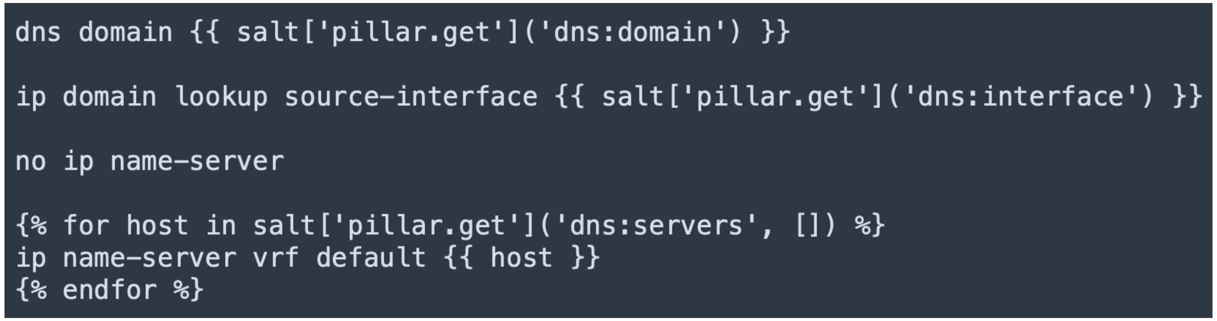

```
dns domain {{ salt['pillar.get']('dns:domain') }}

ip domain lookup source-interface {{ salt['pillar.get']('dns:interface') }}

no ip name-server

{% for host in salt['pillar.get']('dns:servers', []) %}
ip name-server vrf default {{ host }}
{% endfor %}
```

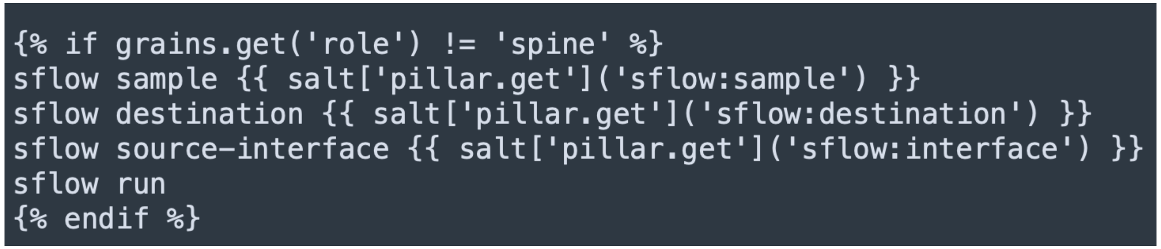

```
{% if grains.get('role') != 'spine' %}
sflow sample {{ salt['pillar.get']('sflow:sample') }}
sflow destination {{ salt['pillar.get']('sflow:destination') }}
sflow source-interface {{ salt['pillar.get']('sflow:interface') }}
sflow run
{% endif %}
```

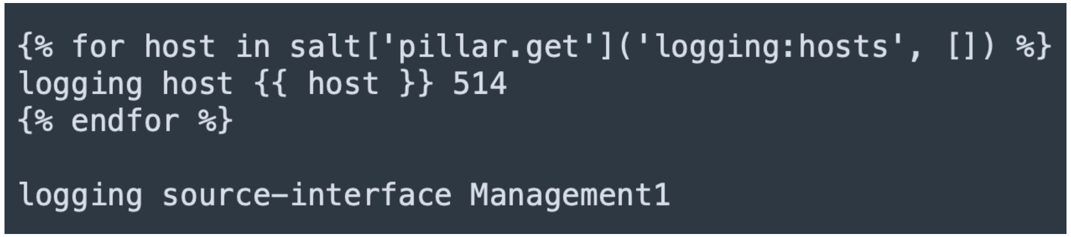

```
{% for host in salt['pillar.get']('logging:hosts', []) %}
logging host {{ host }} 514
{% endfor %}

logging source-interface Management1
```

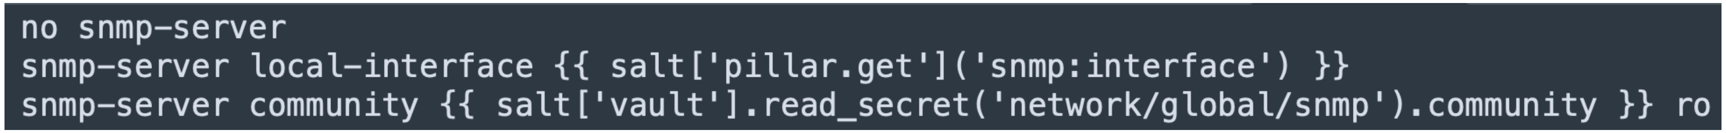

```
no snmp-server
snmp-server local-interface {{ salt['pillar.get']('snmp:interface') }}
snmp-server community {{ salt['vault'].read_secret('network/global/snmp').community }} ro
```

LONAP™

Merging (vs Replace)

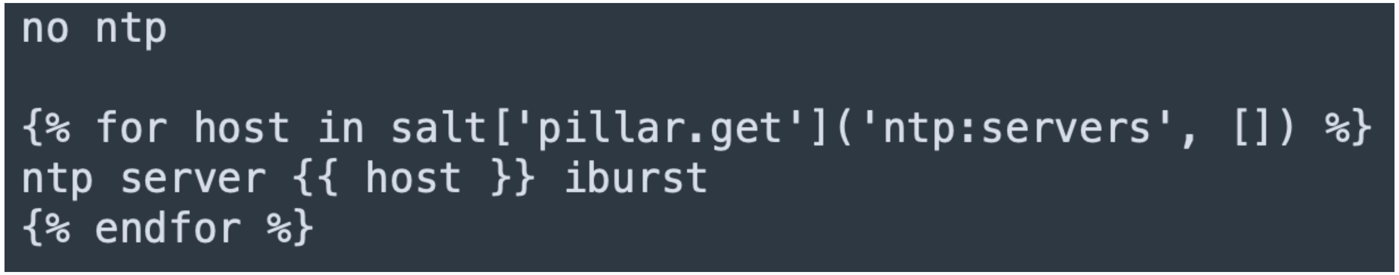

```
no ntp

{% for host in salt['pillar.get']('ntp:servers', []) %}
ntp server {{ host }} iburst
{% endfor %}
```

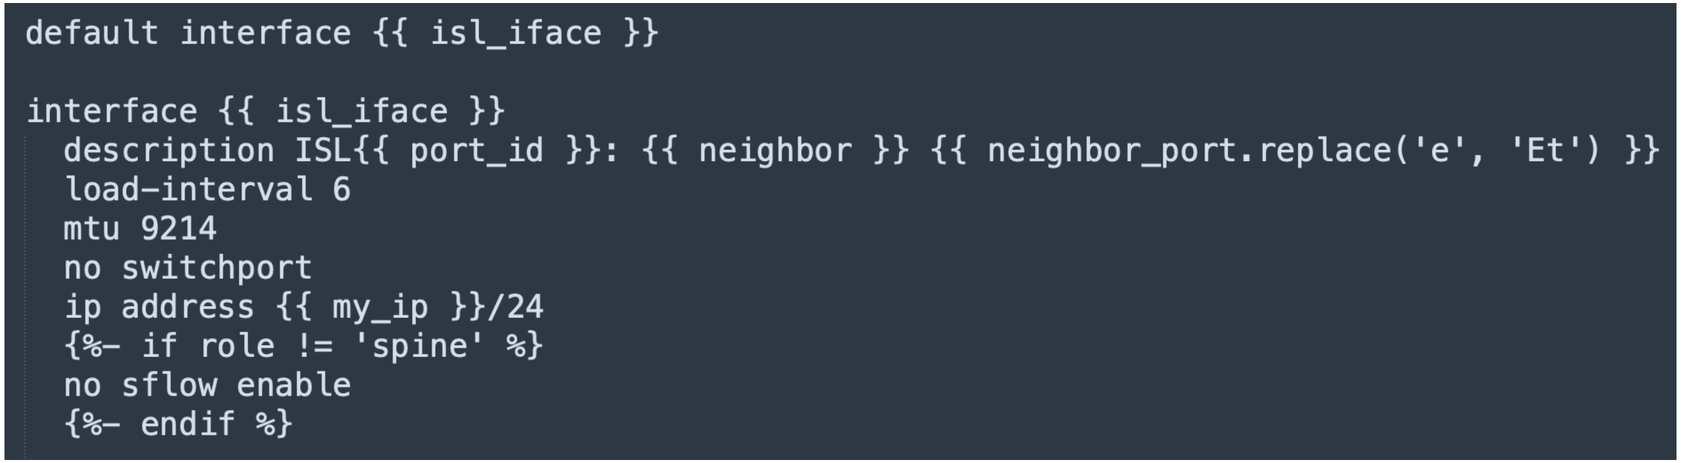

```
default interface {{ isl_iface }}

interface {{ isl_iface }}
  description ISL{{ port_id }}: {{ neighbor }} {{ neighbor_port.replace('e', 'Et') }}
  load-interval 6
  mtu 9214
  no switchport
  ip address {{ my_ip }}/24
  {%- if role != 'spine' %}
  no sflow enable
  {%- endif %}
```

Configuration Deployment

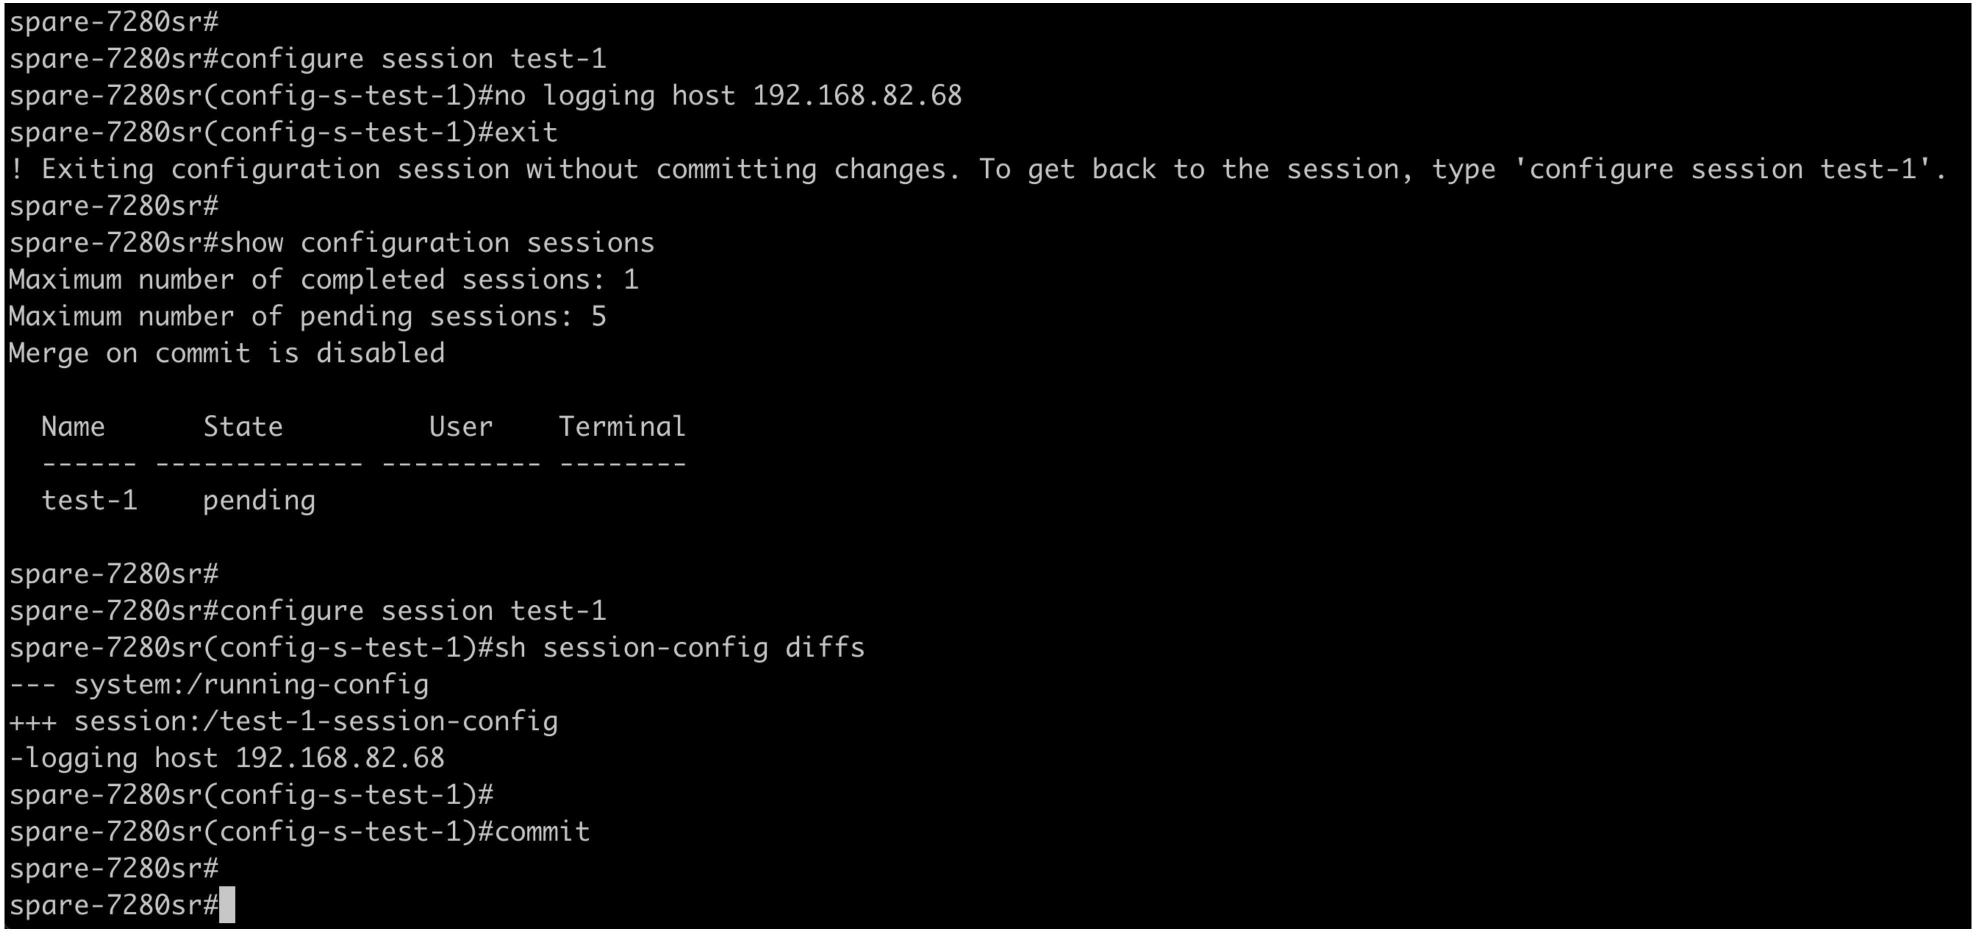

```
spare-7280sr#
spare-7280sr#configure session test-1
spare-7280sr(config-s-test-1)#no logging host 192.168.82.68
spare-7280sr(config-s-test-1)#exit
! Exiting configuration session without committing changes. To get back to the session, type 'configure session test-1'.
spare-7280sr#
spare-7280sr#show configuration sessions
Maximum number of completed sessions: 1
Maximum number of pending sessions: 5
Merge on commit is disabled

  Name      State         User     Terminal
  ------ ------------- ---------- --------
  test-1    pending

spare-7280sr#
spare-7280sr#configure session test-1
spare-7280sr(config-s-test-1)#sh session-config diffs
--- system:/running-config
+++ session:/test-1-session-config
-logging host 192.168.82.68
spare-7280sr(config-s-test-1)#
spare-7280sr(config-s-test-1)#commit
spare-7280sr#
spare-7280sr#
```

‘run-salt’

Humans vs cron…

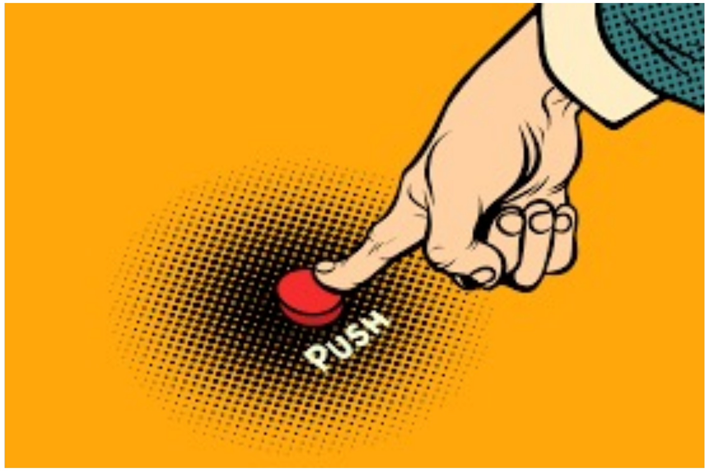

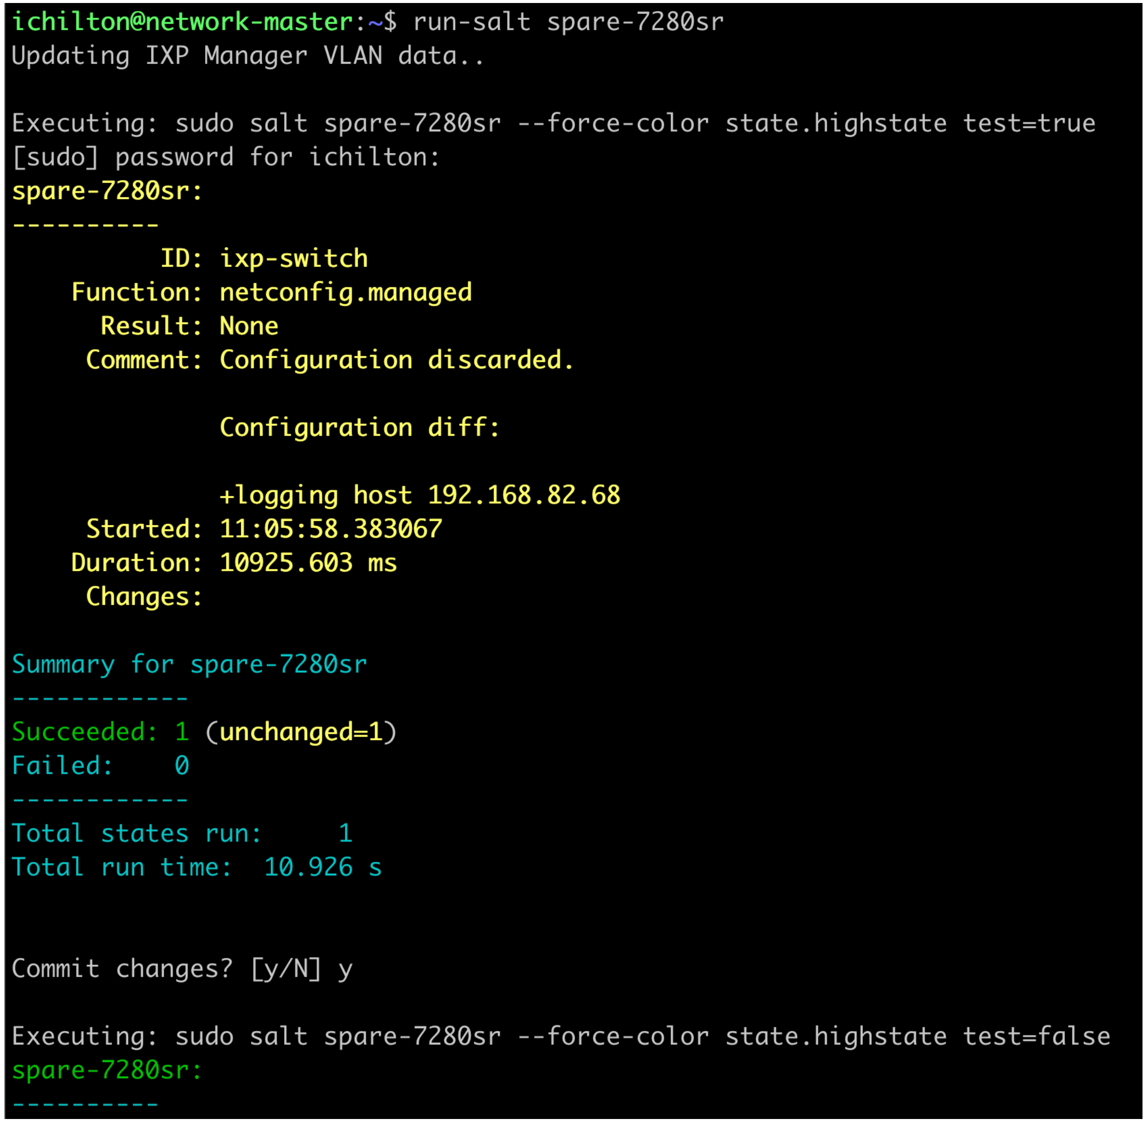

```
ichilton@network-master:~$ run-salt spare-7280sr
Updating IXP Manager VLAN data..

Executing: sudo salt spare-7280sr --force-color state.highstate test=true
[sudo] password for ichilton:
spare-7280sr:
----------
          ID: ixp-switch
    Function: netconfig.managed
      Result: None
     Comment: Configuration discarded.

              Configuration diff:

              +logging host 192.168.82.68
     Started: 11:05:58.383067
    Duration: 10925.603 ms
     Changes:

Summary for spare-7280sr
------------
Succeeded: 1 (unchanged=1)
Failed:    0
------------
Total states run:     1
Total run time:  10.926 s


Commit changes? [y/N] y

Executing: sudo salt spare-7280sr --force-color state.highstate test=false
spare-7280sr:
----------
```

LONAP™

Changes you forgot to apply!

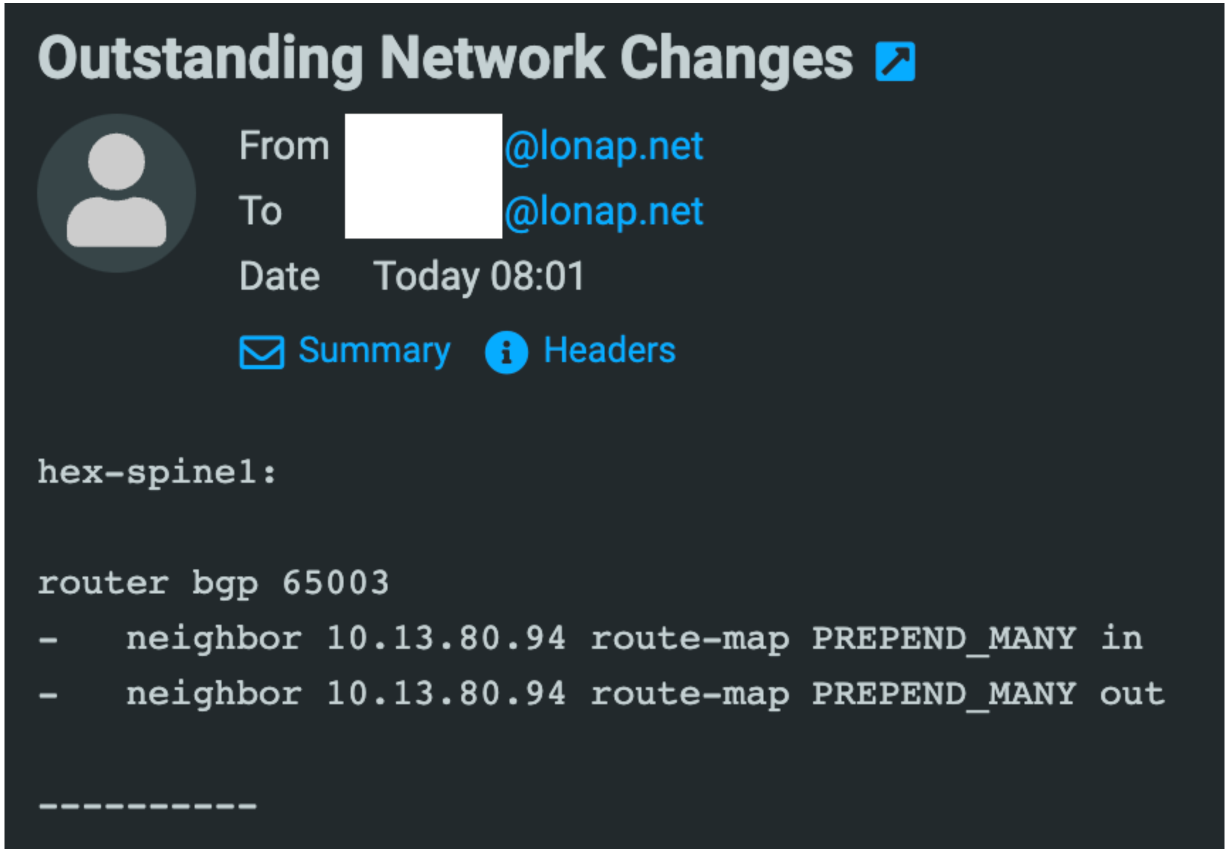

Outstanding Network Changes

From [redacted]@lonap.net

To [redacted]@lonap.net

Date Today 08:01

Summary Headers

```
hex-spine1:

router bgp 65003
-   neighbor 10.13.80.94 route-map PREPEND_MANY in
-   neighbor 10.13.80.94 route-map PREPEND_MANY out

----------
```

The End ?????

Let’s revisit this...

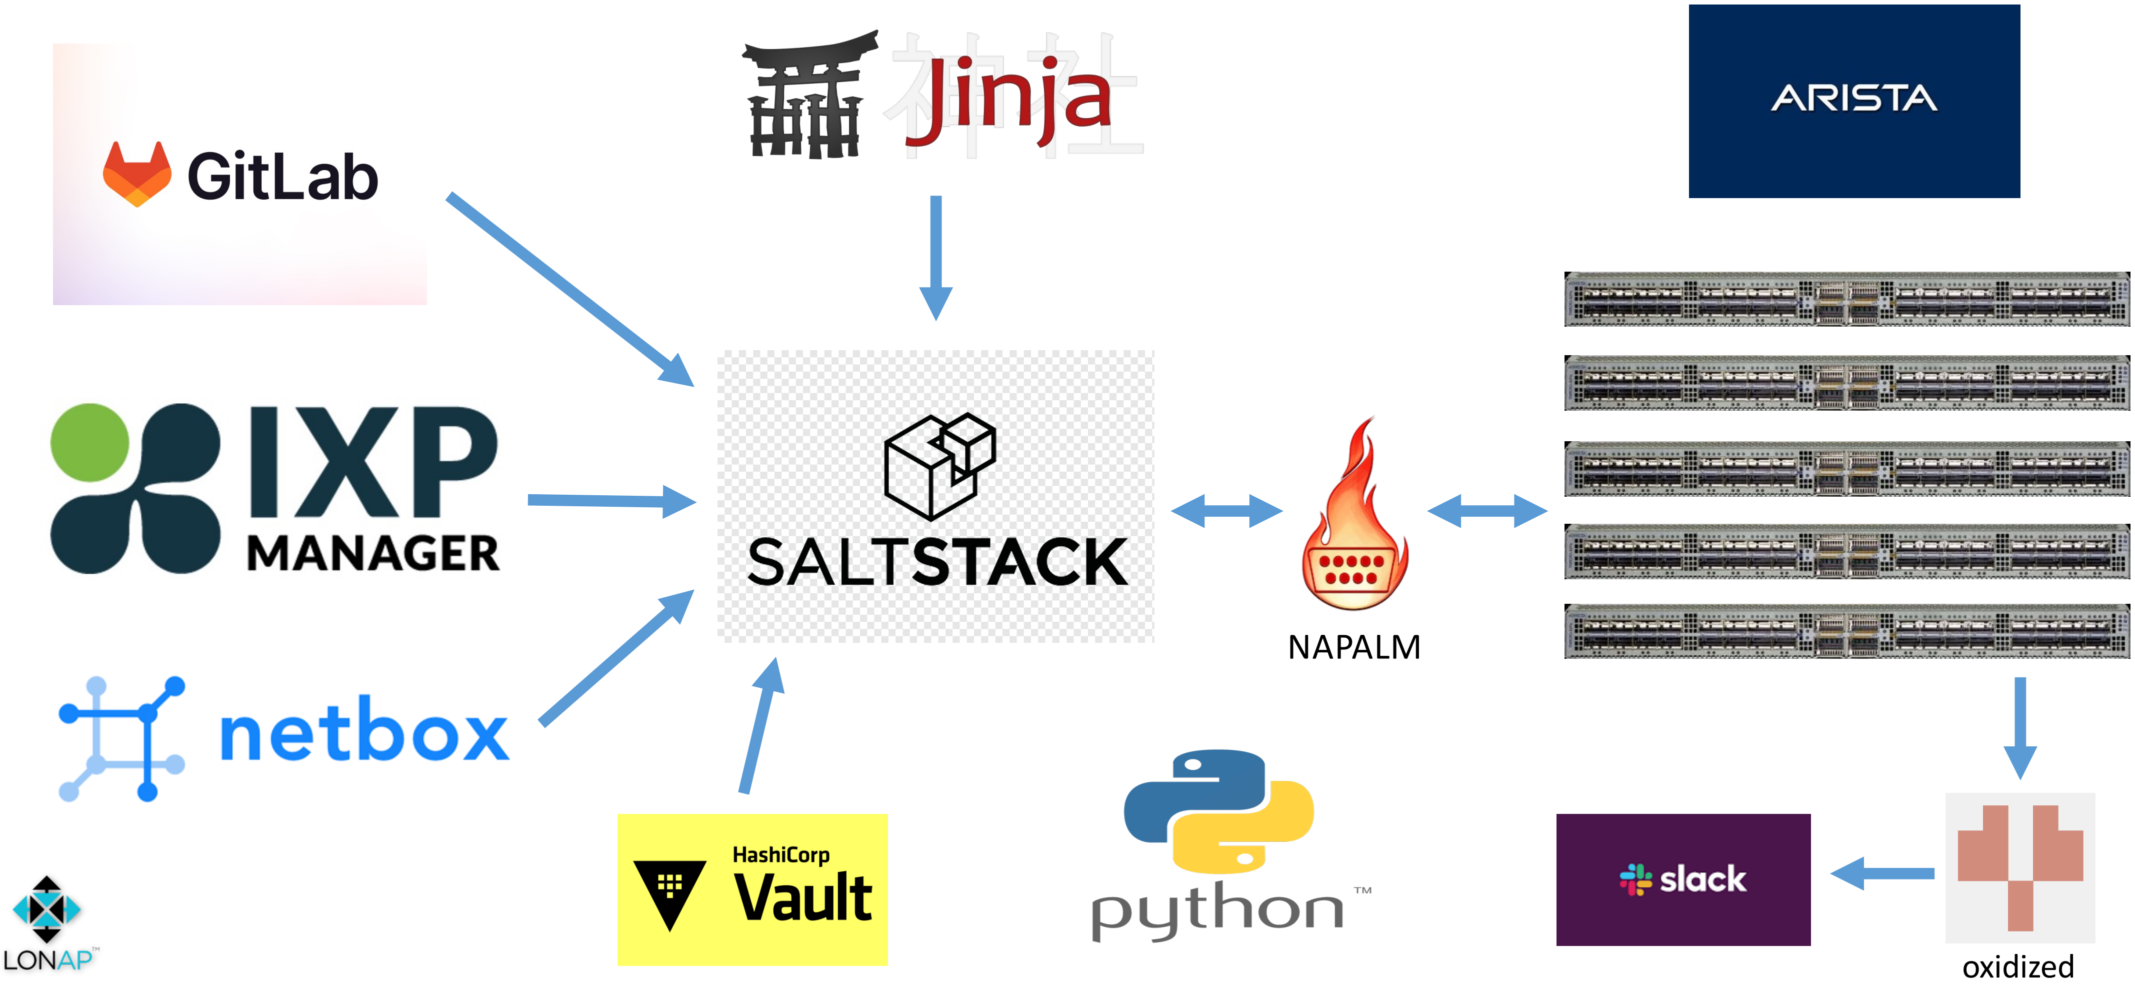

Hmm...do we actually need SaltStack? 🤔

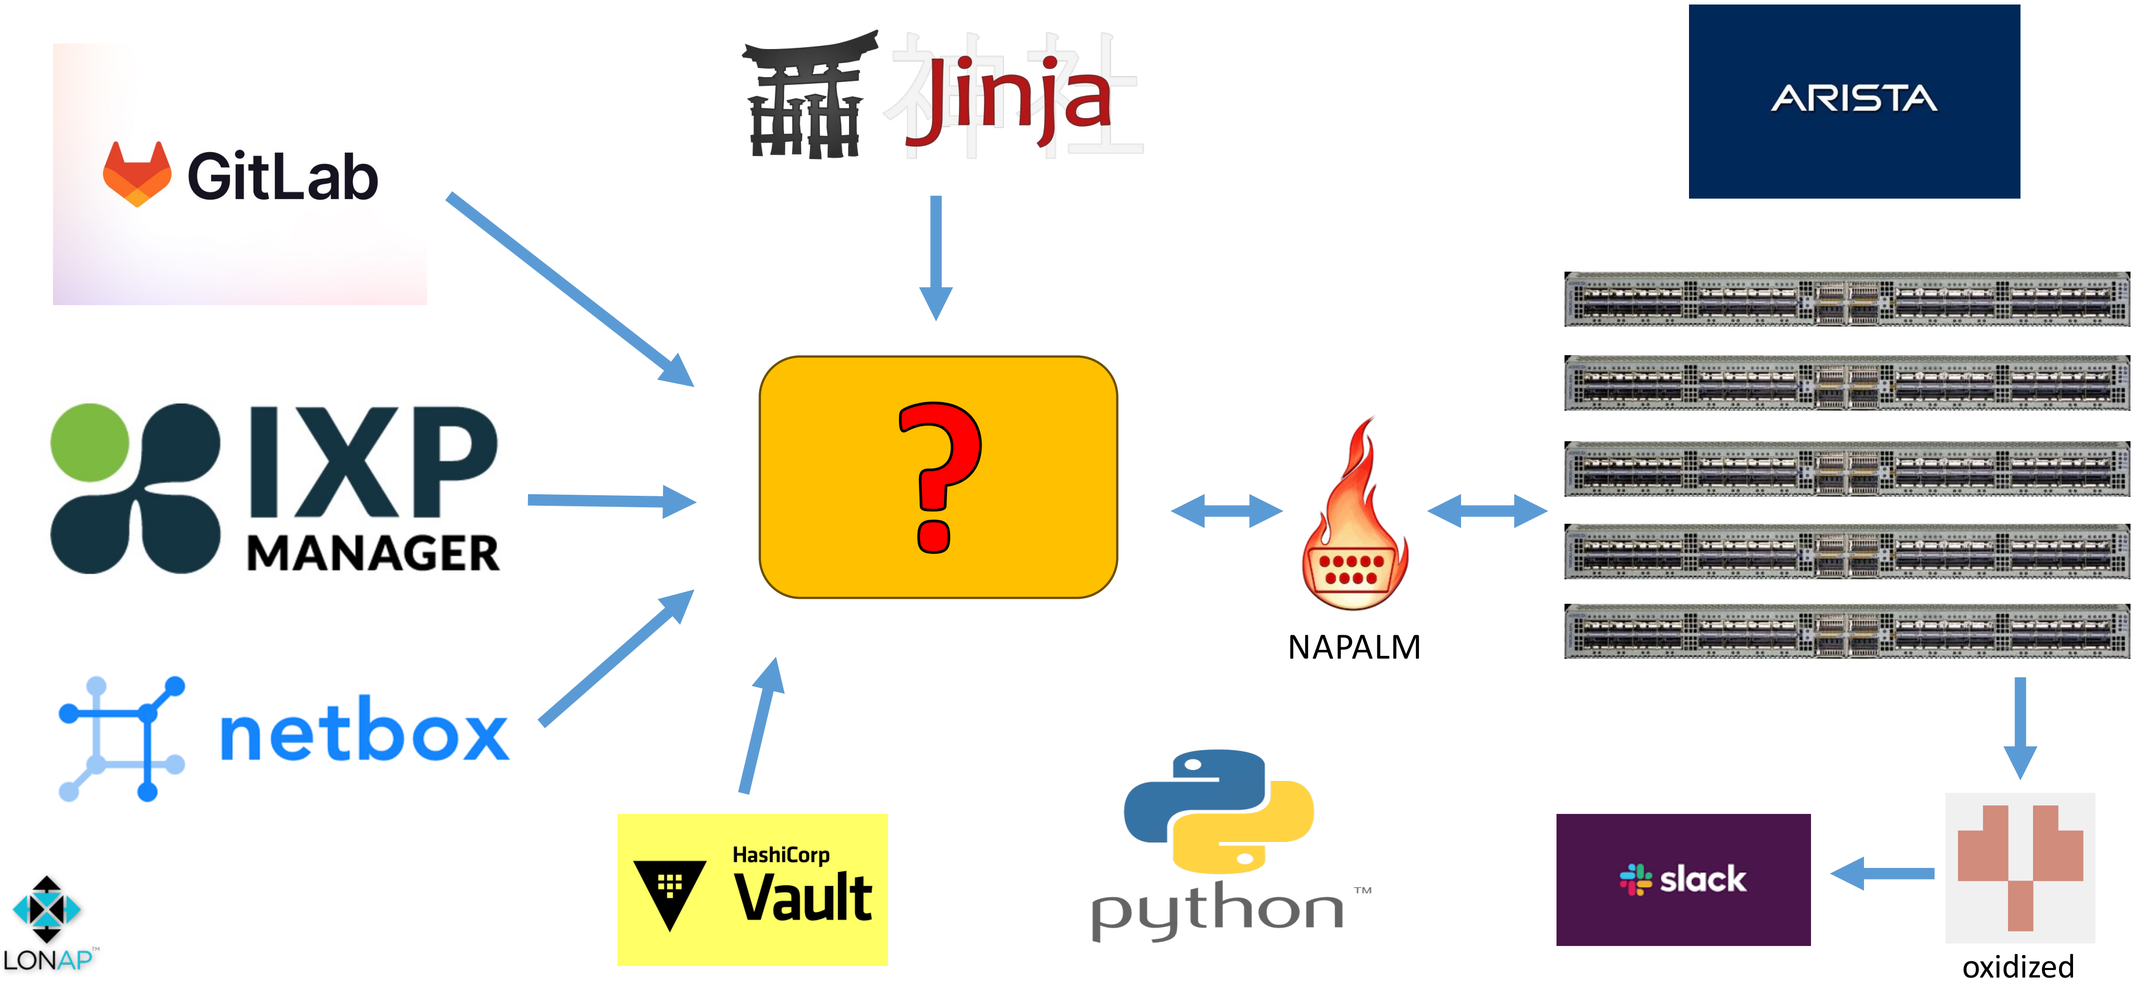

SaltStack

- SaltStack has worked great, but it’s a big codebase maintained by others.
  - Breakages are a real problem in network operations.
  - Awaiting packages.
- Only using a small part of it.
- Managing a minion proxy process per-device.
- Template / data errors can be difficult to debug.

Custom Tooling

- Own code – all within our control.
- Can run from anywhere.
  - Minimal dependencies.
  - Though does require all of the data sources to be in place.
- Faster.
- More flexible.
- Can design interface how we want it to work.
- Can easily expand scope of tooling.

Starting afresh… ish.

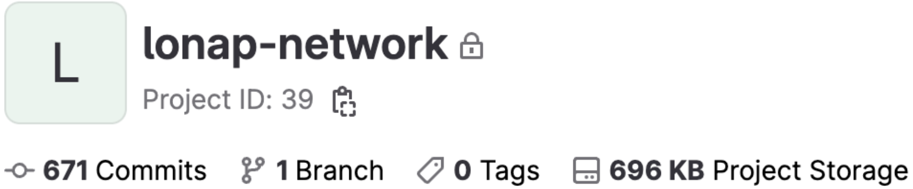

L

lonap-network 🔒

Project ID: 39

671 Commits   1 Branch   0 Tags   696 KB Project Storage

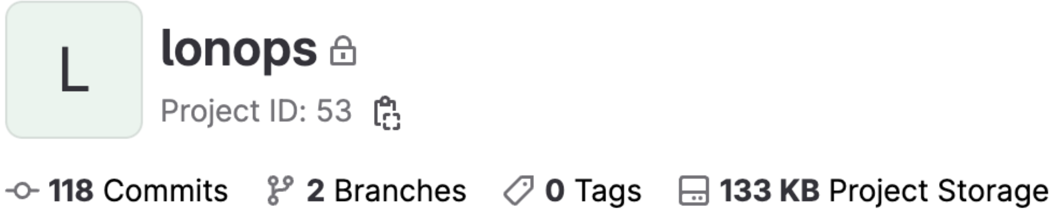

L

lonops 🔒

Project ID: 53

118 Commits   2 Branches   0 Tags   133 KB Project Storage

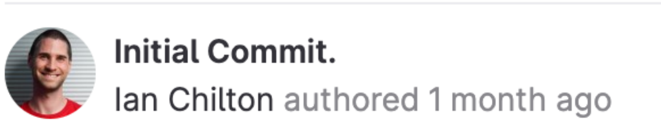

Initial Commit.

Ian Chilton authored 1 month ago

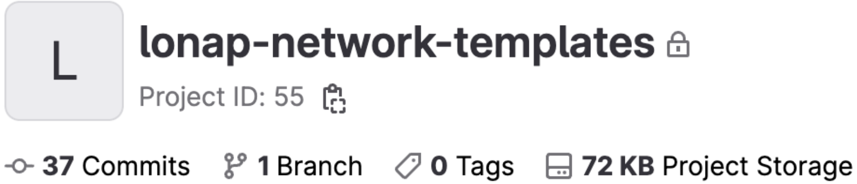

L

lonap-network-templates 🔒

Project ID: 55

37 Commits   1 Branch   0 Tags   72 KB Project Storage

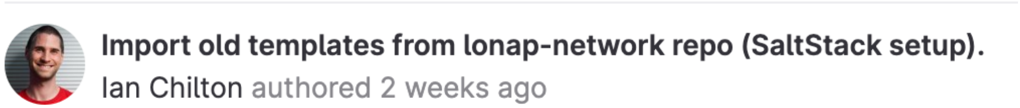

Import old templates from lonap-network repo (SaltStack setup).

Ian Chilton authored 2 weeks ago

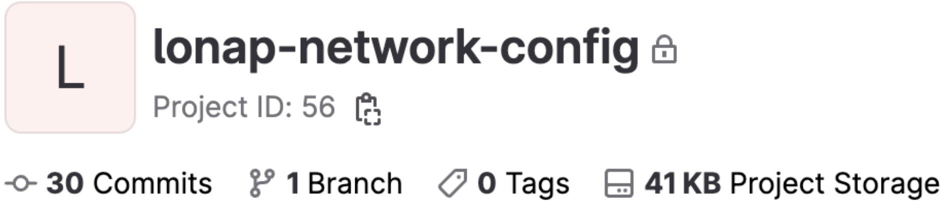

L

lonap-network-config 🔒

Project ID: 56

30 Commits   1 Branch   0 Tags   41 KB Project Storage

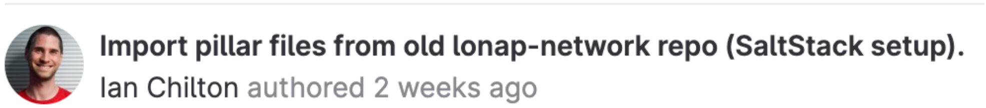

Import pillar files from old lonap-network repo (SaltStack setup).

Ian Chilton authored 2 weeks ago

Safety Concerns..

- The tooling breaking could be a nightmare during maintenance or DC work.
- However…. primary concern is to generate the **correct** configuration. We must not break the network.
- Better to fail than to generate the wrong configuration, because some data is not available.
- Reduce dependencies – we don’t want the tool we need to fix the network to be broken with the network.
- Think about race conditions when committing configuration (see change -> data changes -> commit change). Ideally we want to be sure we are committing the changes we’ve seen/approved.
- Cache external data. Instead of calling a HTTP API directly, write it to a file.
- Keep good logs / history – useful for operational reasons (history), essential for debugging.
- Beware of vendor bugs.

‘LONOPS’

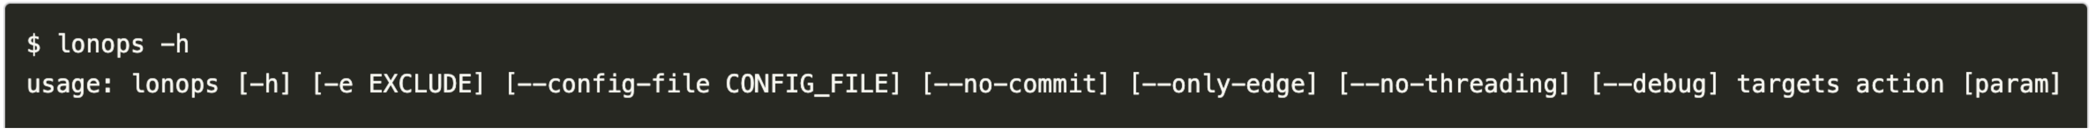

```
$ lonops -h
usage: lonops [-h] [-e EXCLUDE] [--config-file CONFIG_FILE] [--no-commit] [--only-edge] [--no-threading] [--debug] targets action [param]
```

$ lonops <targets> <action> [param]

Targets (1)

Single device:

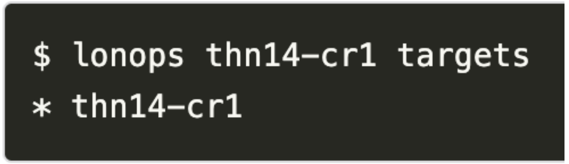

```
$ lonops thn14-cr1 targets
* thn14-cr1
```

Multiple devices:

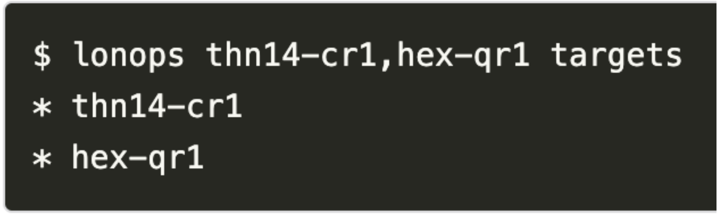

```
$ lonops thn14-cr1,hex-qr1 targets
* thn14-cr1
* hex-qr1
```

Wildcard:

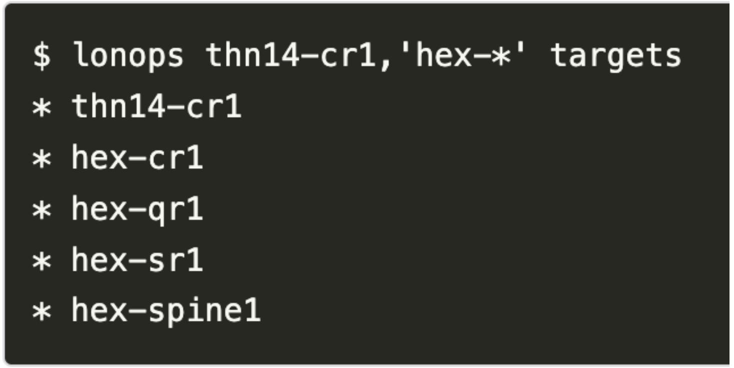

```
$ lonops thn14-cr1,'hex-*' targets
* thn14-cr1
* hex-cr1
* hex-qr1
* hex-sr1
* hex-spine1
```

Targets (2)

Only edge boxes (no spines):

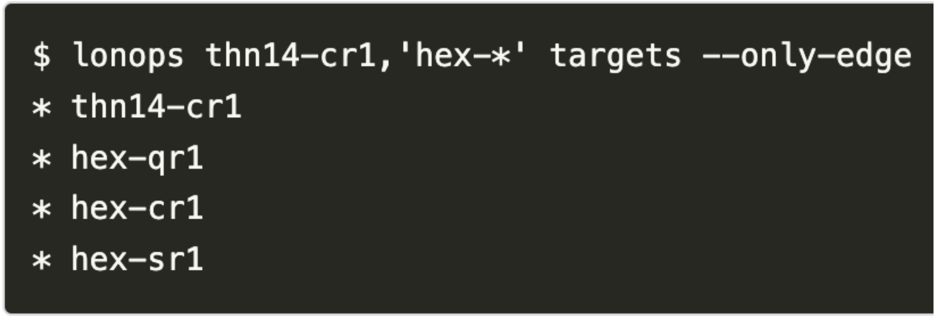

```
$ lonops thn14-cr1,'hex-*' targets --only-edge
* thn14-cr1
* hex-qr1
* hex-cr1
* hex-sr1
```

Exclude some matches:

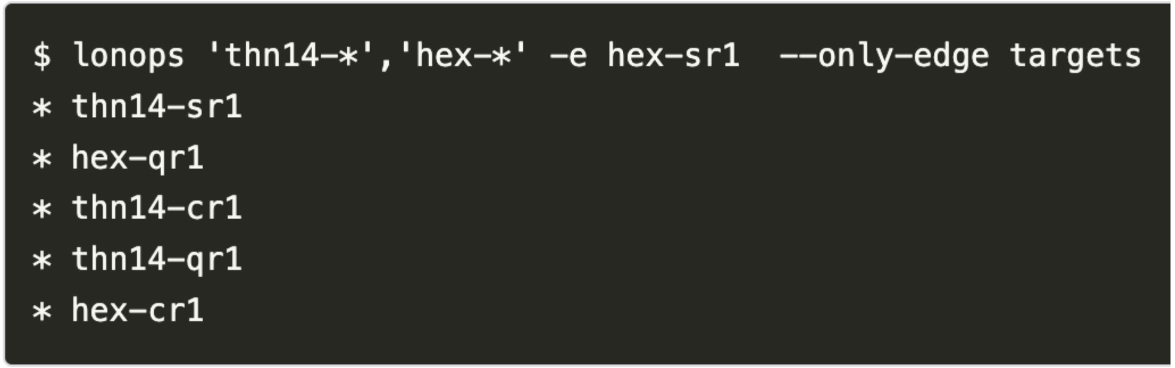

```
$ lonops 'thn14-*','hex-*' -e hex-sr1  --only-edge targets
* thn14-sr1
* hex-qr1
* thn14-cr1
* thn14-qr1
* hex-cr1
```

Exclude with wildcard:

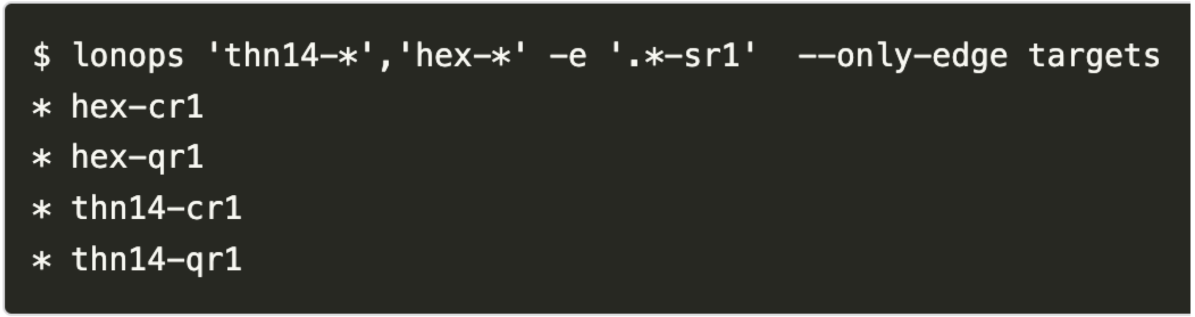

```
$ lonops 'thn14-*','hex-*' -e '.*-sr1'  --only-edge targets
* hex-cr1
* hex-qr1
* thn14-cr1
* thn14-qr1
```

Targets (3)

Just do everything:

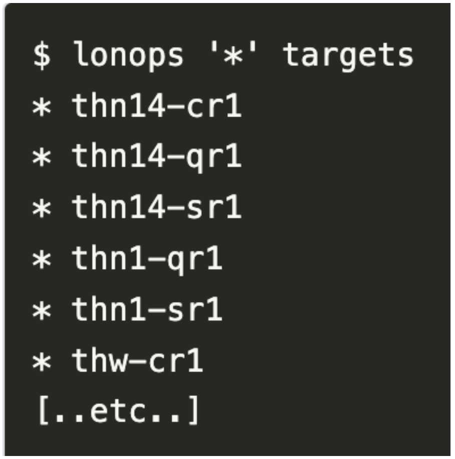

```
$ lonops '*' targets
* thn14-cr1
* thn14-qr1
* thn14-sr1
* thn1-qr1
* thn1-sr1
* thw-cr1
[..etc..]
```

..but all at the same time:
(using threads)

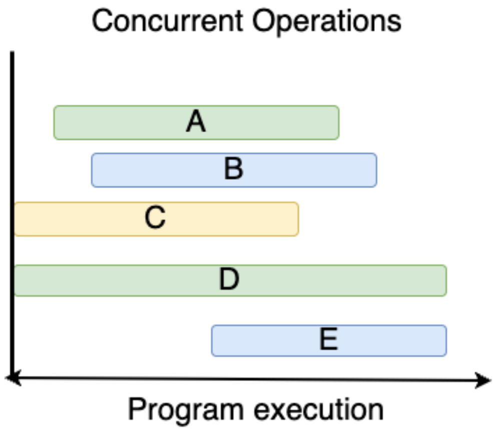

Coming soon:

- By environment.
- By role.

Action: facts

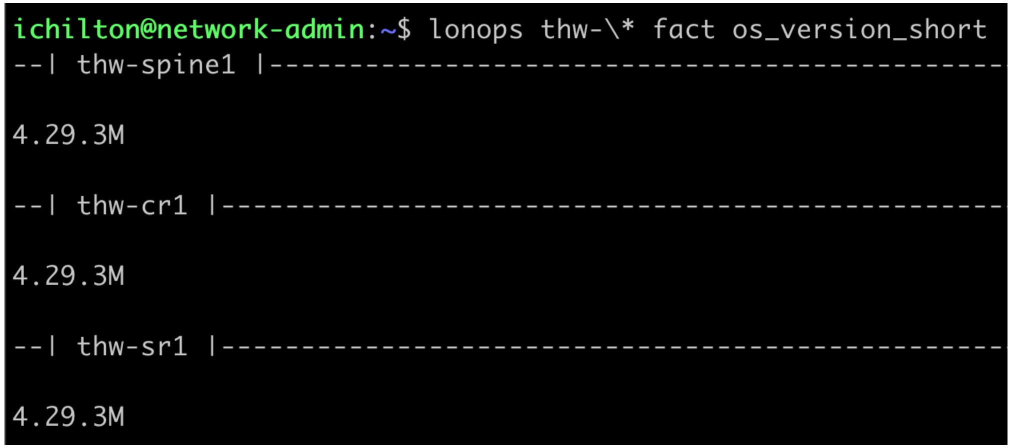

```
ichilton@network-admin:~$ lonops thw-\* fact os_version_short
--| thw-spine1 |----------------------------------------------

4.29.3M

--| thw-cr1 |-------------------------------------------------

4.29.3M

--| thw-sr1 |-------------------------------------------------

4.29.3M
```

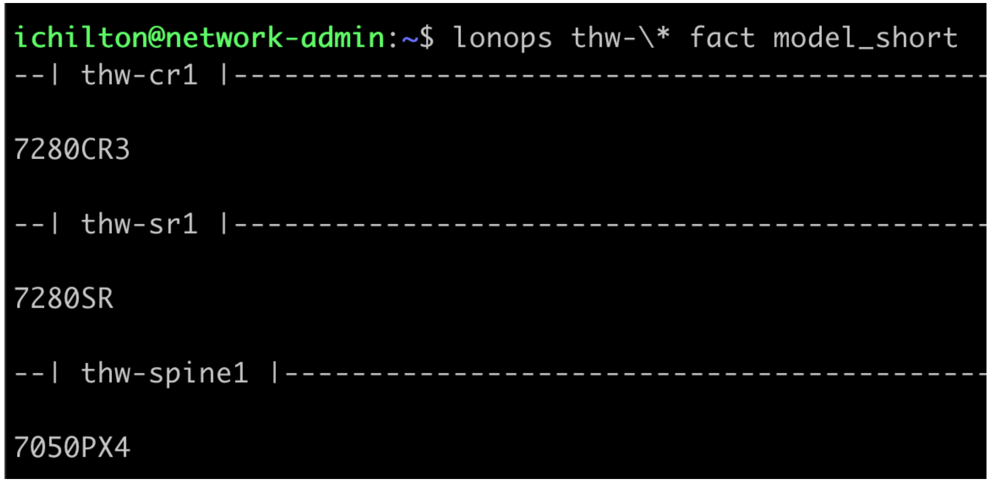

```
ichilton@network-admin:~$ lonops thw-\* fact model_short
--| thw-cr1 |---------------------------------------------

7280CR3

--| thw-sr1 |---------------------------------------------

7280SR

--| thw-spine1 |------------------------------------------

7050PX4
```

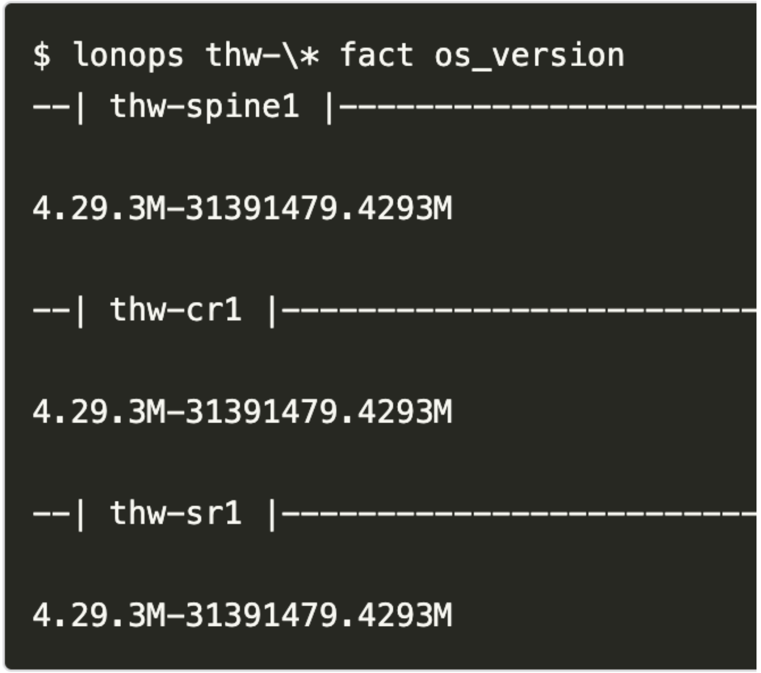

```
$ lonops thw-\* fact os_version
--| thw-spine1 |------------------------

4.29.3M-31391479.4293M

--| thw-cr1 |---------------------------

4.29.3M-31391479.4293M

--| thw-sr1 |---------------------------

4.29.3M-31391479.4293M
```

Action: set_config

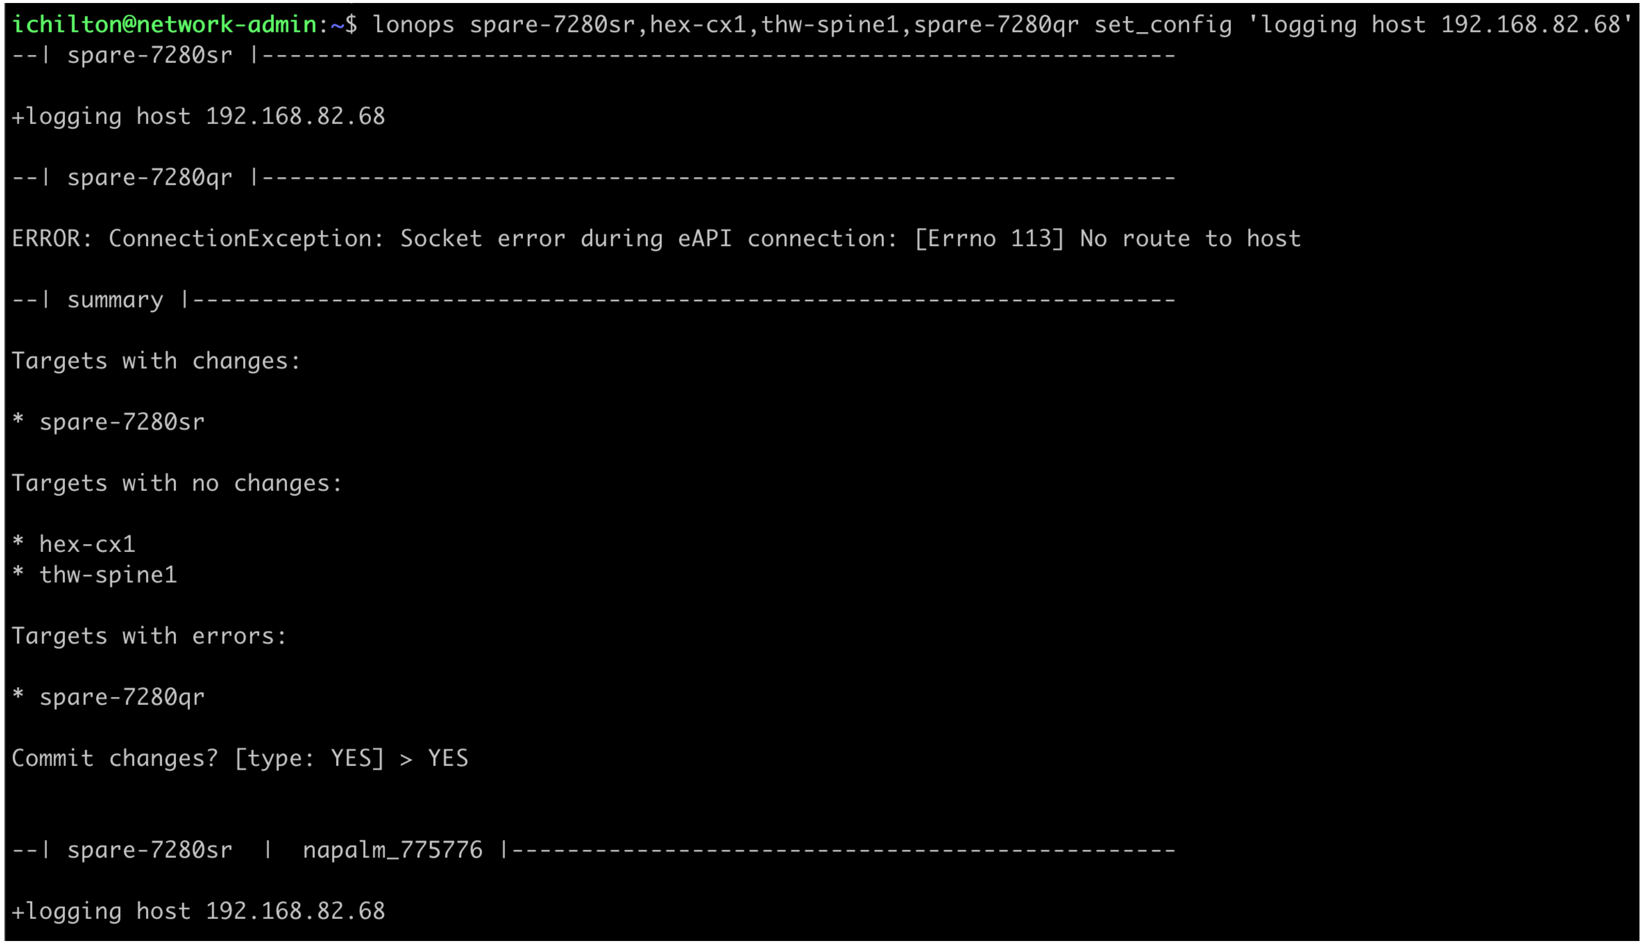

```
ichilton@network-admin:~$ lonops spare-7280sr,hex-cx1,thw-spine1,spare-7280qr set_config 'logging host 192.168.82.68'
--| spare-7280sr |------------------------------------------------------------------

+logging host 192.168.82.68

--| spare-7280qr |------------------------------------------------------------------

ERROR: ConnectionException: Socket error during eAPI connection: [Errno 113] No route to host

--| summary |-----------------------------------------------------------------------

Targets with changes:

* spare-7280sr

Targets with no changes:

* hex-cx1
* thw-spine1

Targets with errors:

* spare-7280qr

Commit changes? [type: YES] > YES


--| spare-7280sr  |  napalm_775776 |------------------------------------------------

+logging host 192.168.82.68
```

Instant backups/notification with Oxidized

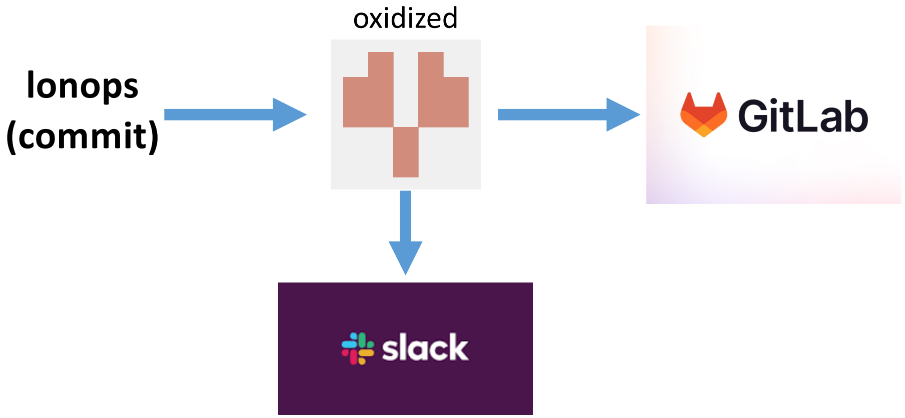

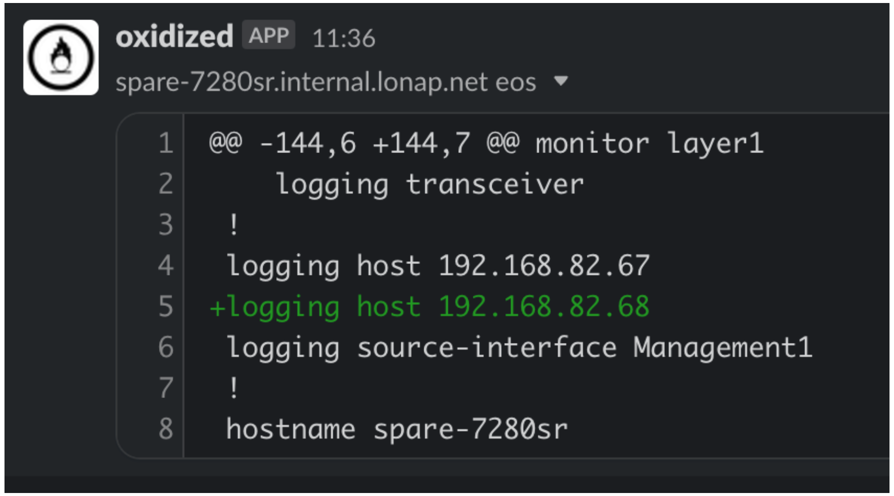

oxidized APP 11:36

spare-7280sr.internal.lonap.net eos

```
@@ -144,6 +144,7 @@ monitor layer1
    logging transceiver
 !
 logging host 192.168.82.67
+logging host 192.168.82.68
 logging source-interface Management1
 !
 hostname spare-7280sr
```

Action: print, load, replace

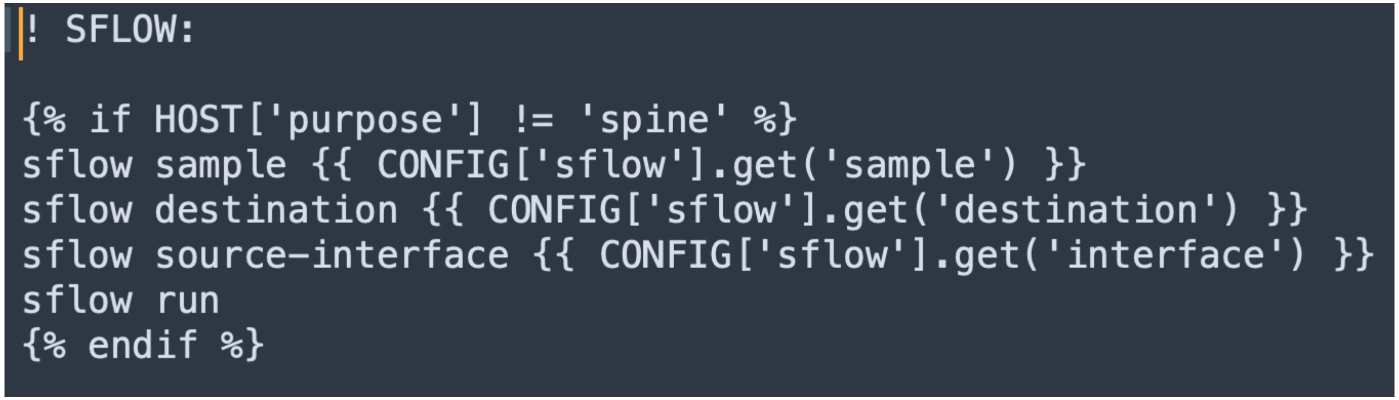

```
! SFLOW:

{% if HOST['purpose'] != 'spine' %}
sflow sample {{ CONFIG['sflow'].get('sample') }}
sflow destination {{ CONFIG['sflow'].get('destination') }}
sflow source-interface {{ CONFIG['sflow'].get('interface') }}
sflow run
{% endif %}
```

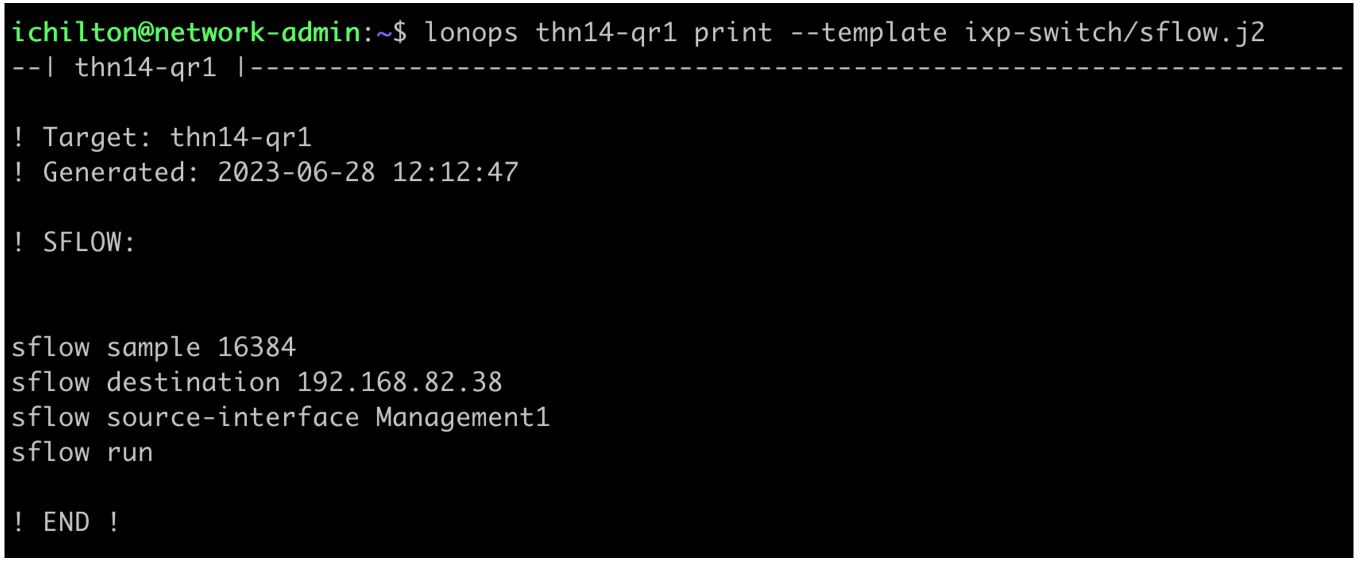

```
ichilton@network-admin:~$ lonops thn14-qr1 print --template ixp-switch/sflow.j2
--| thn14-qr1 |-------------------------------------------------------------

! Target: thn14-qr1
! Generated: 2023-06-28 12:12:47

! SFLOW:


sflow sample 16384
sflow destination 192.168.82.38
sflow source-interface Management1
sflow run

! END !
```

LONAP™

Action: print, load, replace

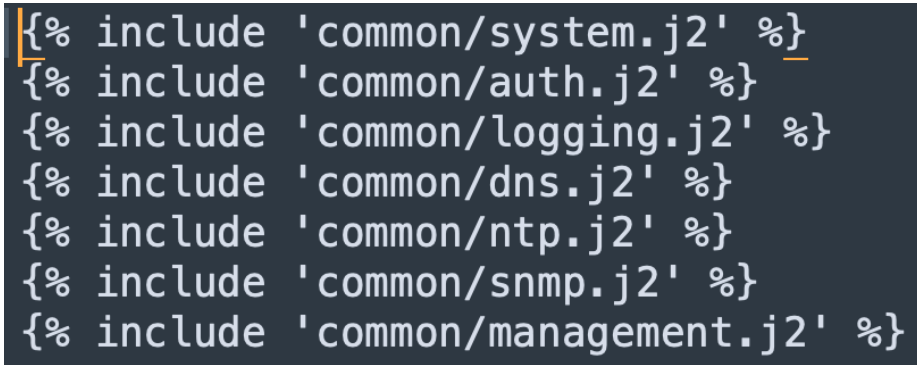

```
{% include 'common/system.j2' %}
{% include 'common/auth.j2' %}
{% include 'common/logging.j2' %}
{% include 'common/dns.j2' %}
{% include 'common/ntp.j2' %}
{% include 'common/snmp.j2' %}
{% include 'common/management.j2' %}
```

common.j2

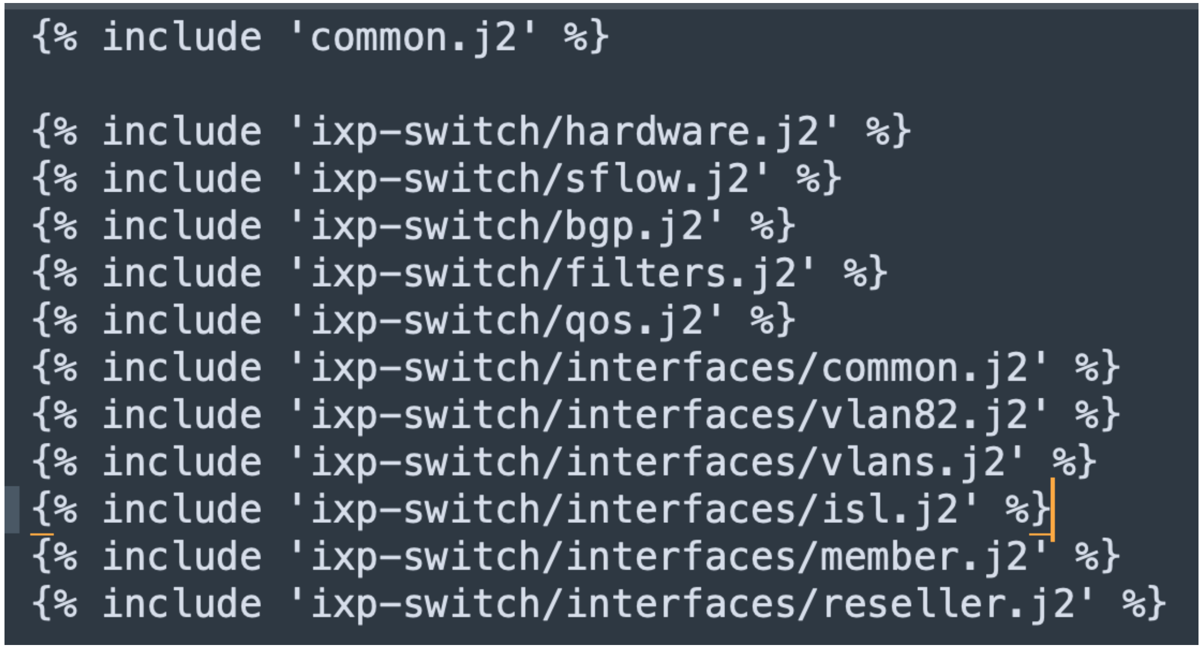

```
{% include 'common.j2' %}

{% include 'ixp-switch/hardware.j2' %}
{% include 'ixp-switch/sflow.j2' %}
{% include 'ixp-switch/bgp.j2' %}
{% include 'ixp-switch/filters.j2' %}
{% include 'ixp-switch/qos.j2' %}
{% include 'ixp-switch/interfaces/common.j2' %}
{% include 'ixp-switch/interfaces/vlan82.j2' %}
{% include 'ixp-switch/interfaces/vlans.j2' %}
{% include 'ixp-switch/interfaces/isl.j2' %}
{% include 'ixp-switch/interfaces/member.j2' %}
{% include 'ixp-switch/interfaces/reseller.j2' %}
```

ixp-switch.j2

Action: maintenance (member affecting)

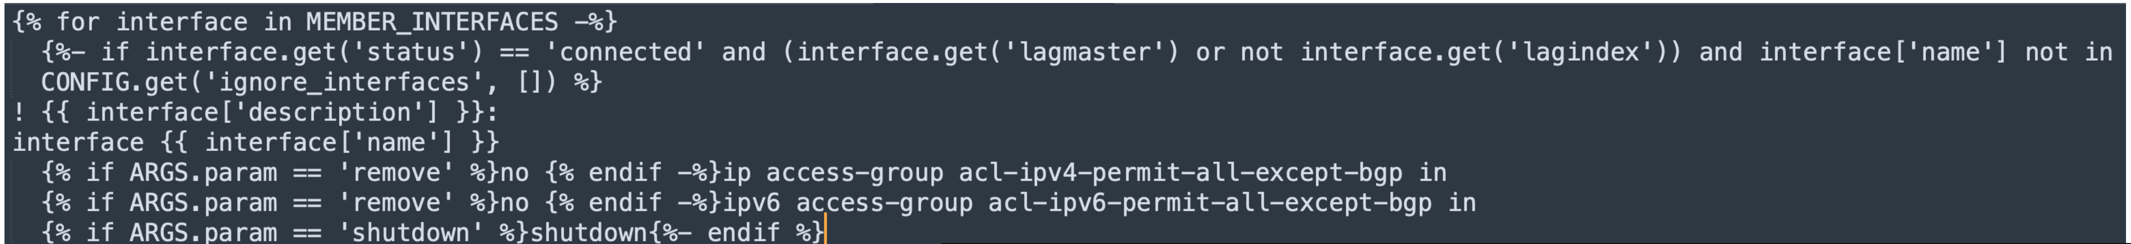

```
{% for interface in MEMBER_INTERFACES -%}
  {%- if interface.get('status') == 'connected' and (interface.get('lagmaster') or not interface.get('lagindex')) and interface['name'] not in
  CONFIG.get('ignore_interfaces', []) %}
! {{ interface['description'] }}:
interface {{ interface['name'] }}
  {% if ARGS.param == 'remove' %}no {% endif -%}ip access-group acl-ipv4-permit-all-except-bgp in
  {% if ARGS.param == 'remove' %}no {% endif -%}ipv6 access-group acl-ipv6-permit-all-except-bgp in
  {% if ARGS.param == 'shutdown' %}shutdown{%- endif %}
  {% endif -%}
{% endfor -%}
```

BCP 214

RFC 8327

Mitigating the Negative Impact of Maintenance through BGP Session Culling.

https://www.rfc-editor.org/info/bcp214

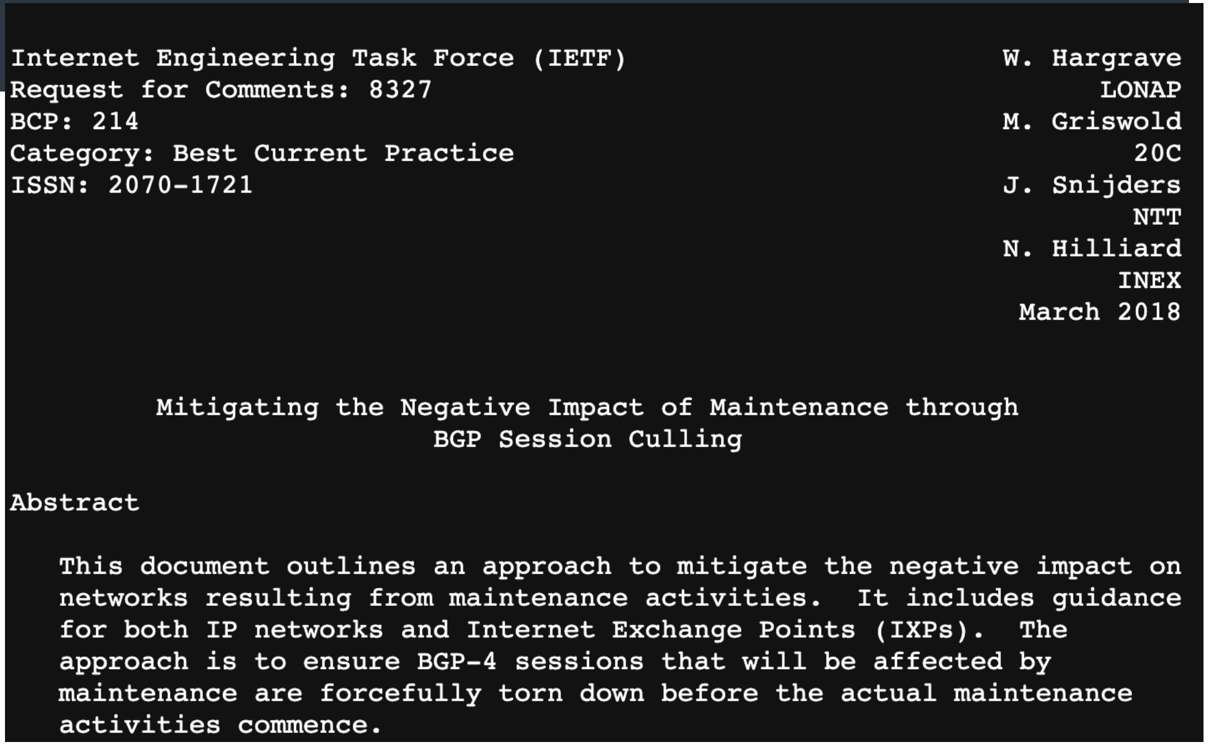

```
Internet Engineering Task Force (IETF)                      W. Hargrave
Request for Comments: 8327                                        LONAP
BCP: 214                                                    M. Griswold
Category: Best Current Practice                                     20C
ISSN: 2070-1721                                             J. Snijders
                                                                    NTT
                                                            N. Hilliard
                                                                   INEX
                                                             March 2018


         Mitigating the Negative Impact of Maintenance through
                          BGP Session Culling

Abstract

   This document outlines an approach to mitigate the negative impact on
   networks resulting from maintenance activities.  It includes guidance
   for both IP networks and Internet Exchange Points (IXPs).  The
   approach is to ensure BGP-4 sessions that will be affected by
   maintenance are forcefully torn down before the actual maintenance
   activities commence.
```

Action: maintenance

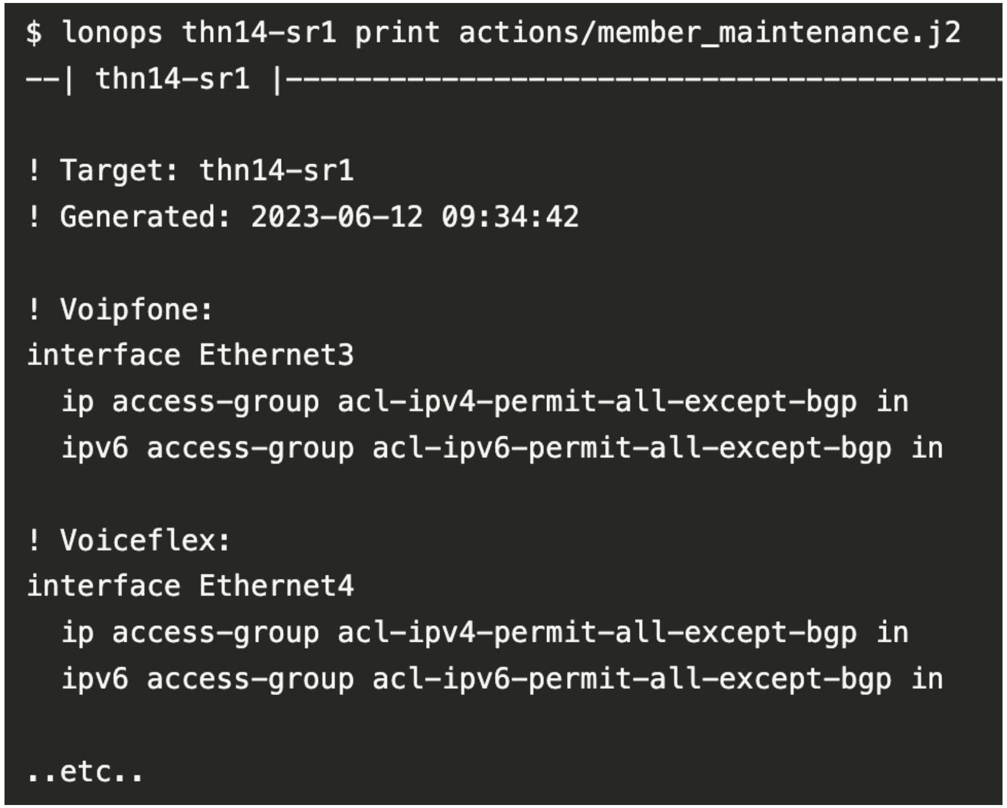

```
$ lonops thn14-sr1 print actions/member_maintenance.j2
--| thn14-sr1 |-----------------------------------------

! Target: thn14-sr1
! Generated: 2023-06-12 09:34:42

! Voipfone:
interface Ethernet3
  ip access-group acl-ipv4-permit-all-except-bgp in
  ipv6 access-group acl-ipv6-permit-all-except-bgp in

! Voiceflex:
interface Ethernet4
  ip access-group acl-ipv4-permit-all-except-bgp in
  ipv6 access-group acl-ipv6-permit-all-except-bgp in

..etc..
```

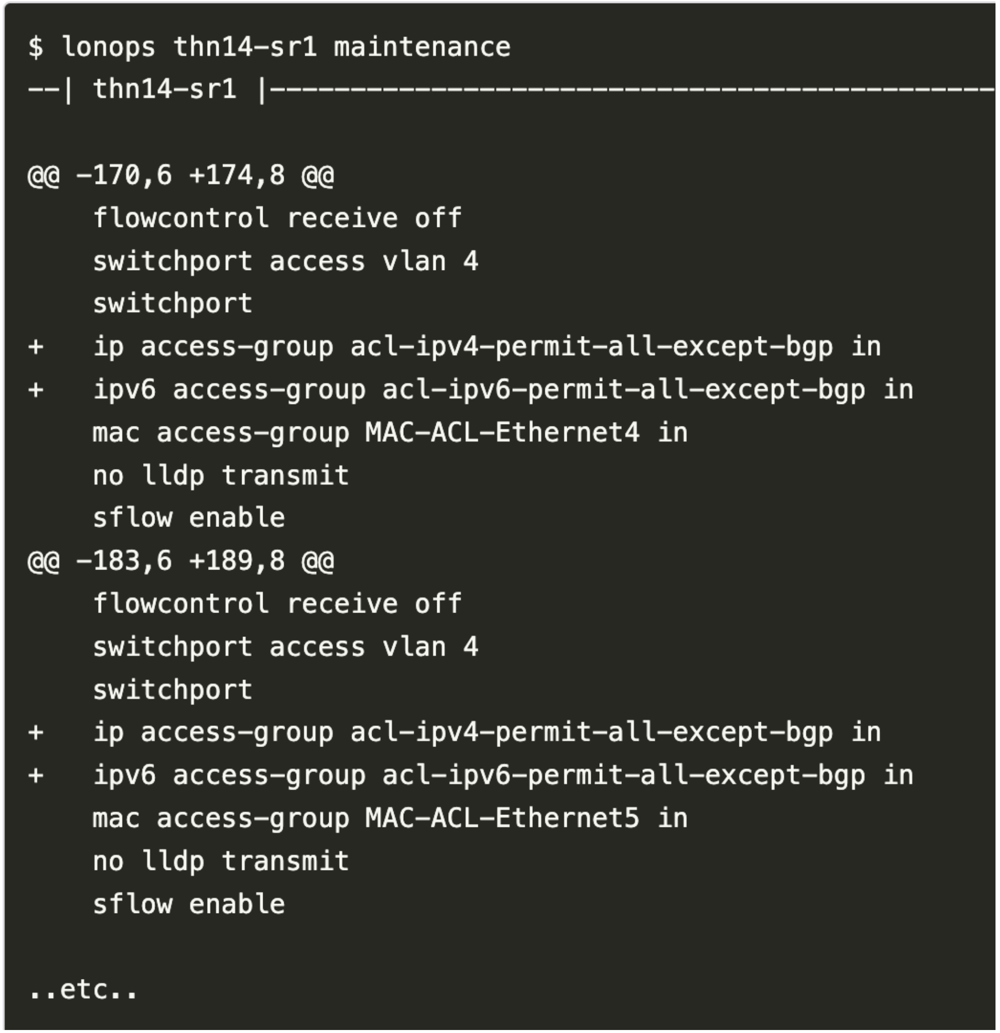

```
$ lonops thn14-sr1 maintenance
--| thn14-sr1 |---------------------------------------------

@@ -170,6 +174,8 @@
    flowcontrol receive off
    switchport access vlan 4
    switchport
+   ip access-group acl-ipv4-permit-all-except-bgp in
+   ipv6 access-group acl-ipv6-permit-all-except-bgp in
    mac access-group MAC-ACL-Ethernet4 in
    no lldp transmit
    sflow enable
@@ -183,6 +189,8 @@
    flowcontrol receive off
    switchport access vlan 4
    switchport
+   ip access-group acl-ipv4-permit-all-except-bgp in
+   ipv6 access-group acl-ipv6-permit-all-except-bgp in
    mac access-group MAC-ACL-Ethernet5 in
    no lldp transmit
    sflow enable

..etc..
```

Development / Troubleshooting

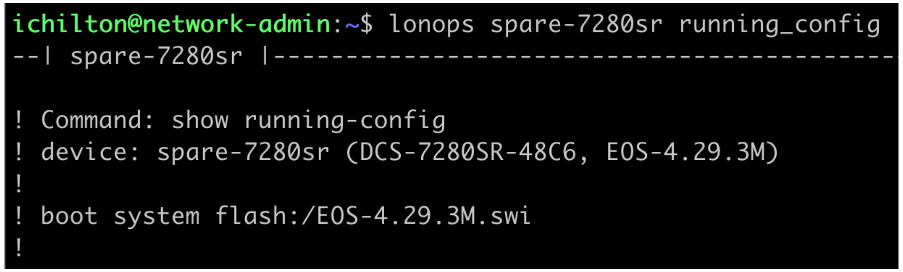

```
ichilton@network-admin:~$ lonops spare-7280sr running_config
--| spare-7280sr |----------------------------------------

! Command: show running-config
! device: spare-7280sr (DCS-7280SR-48C6, EOS-4.29.3M)
!
! boot system flash:/EOS-4.29.3M.swi
!
```

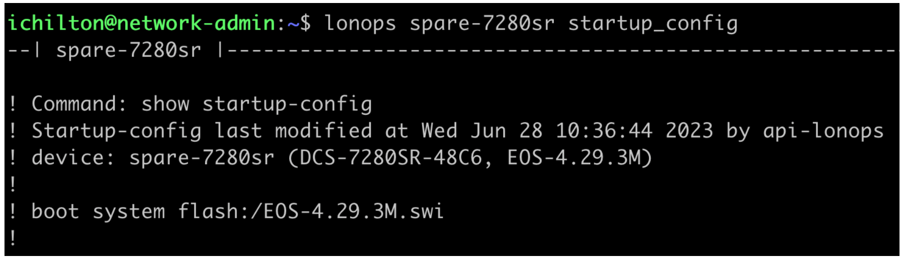

```
ichilton@network-admin:~$ lonops spare-7280sr startup_config
--| spare-7280sr |------------------------------------------------------

! Command: show startup-config
! Startup-config last modified at Wed Jun 28 10:36:44 2023 by api-lonops
! device: spare-7280sr (DCS-7280SR-48C6, EOS-4.29.3M)
!
! boot system flash:/EOS-4.29.3M.swi
!
```

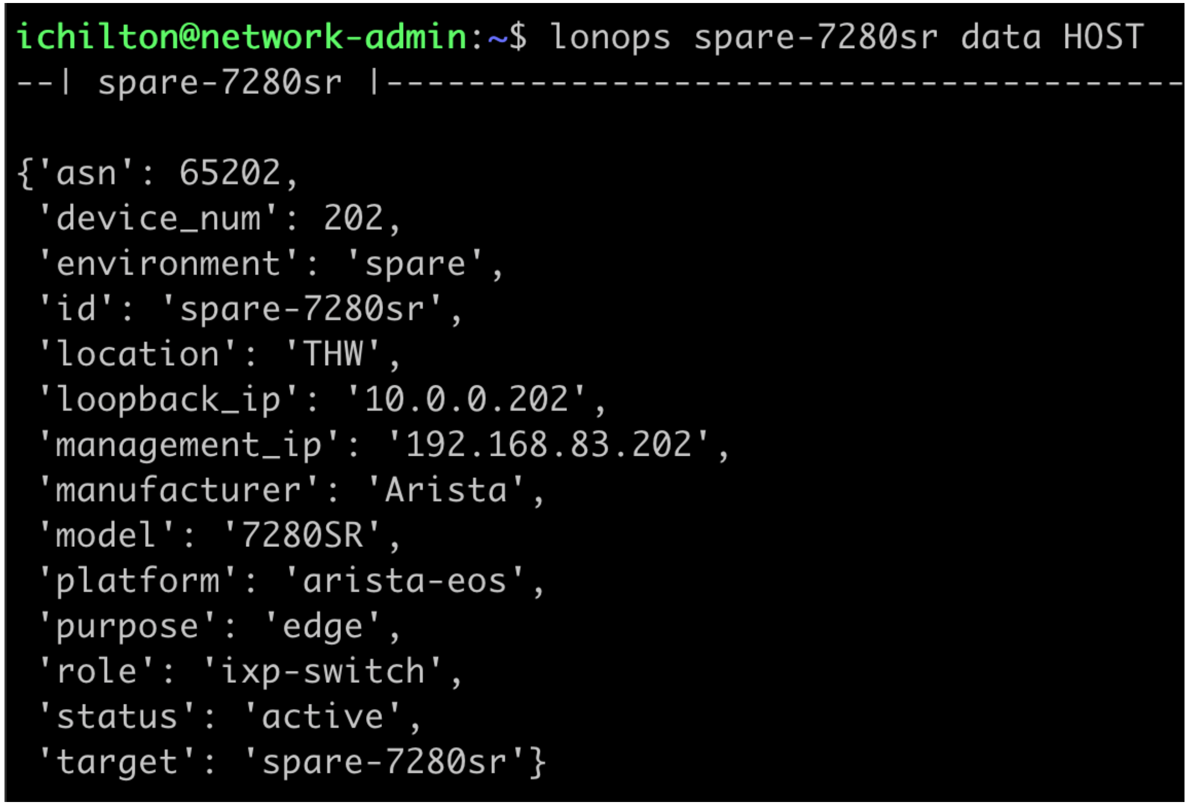

```
ichilton@network-admin:~$ lonops spare-7280sr data HOST
--| spare-7280sr |------------------------------------------

{'asn': 65202,
 'device_num': 202,
 'environment': 'spare',
 'id': 'spare-7280sr',
 'location': 'THW',
 'loopback_ip': '10.0.0.202',
 'management_ip': '192.168.83.202',
 'manufacturer': 'Arista',
 'model': '7280SR',
 'platform': 'arista-eos',
 'purpose': 'edge',
 'role': 'ixp-switch',
 'status': 'active',
 'target': 'spare-7280sr'}
```

Coming soon..

One tool for all day to day activities…

- ISLs – depref, repref.

- Look up member ports, by name, ASN etc:
  - Status.
  - Traffic.
  - Check light levels.

- Optic inventory - where are the spares?

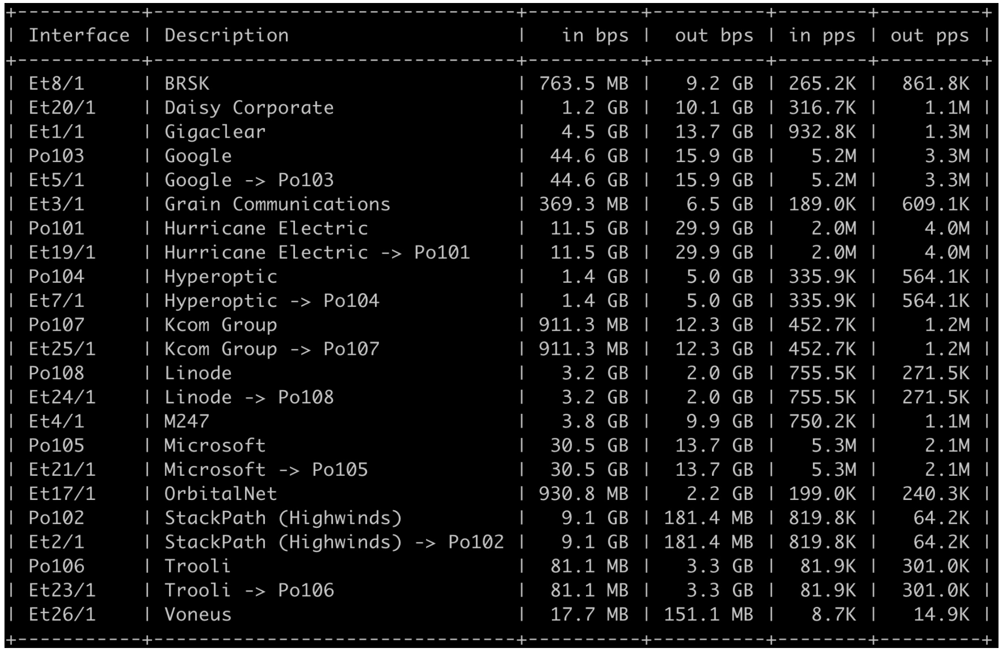

```
+-----------+--------------------------------+----------+----------+--------+---------+
| Interface | Description                    |   in bps |  out bps | in pps | out pps |
+-----------+--------------------------------+----------+----------+--------+---------+
| Et8/1     | BRSK                           | 763.5 MB |   9.2 GB | 265.2K |  861.8K |
| Et20/1    | Daisy Corporate                |   1.2 GB |  10.1 GB | 316.7K |    1.1M |
| Et1/1     | Gigaclear                      |   4.5 GB |  13.7 GB | 932.8K |    1.3M |
| Po103     | Google                         |  44.6 GB |  15.9 GB |   5.2M |    3.3M |
| Et5/1     | Google -> Po103                |  44.6 GB |  15.9 GB |   5.2M |    3.3M |
| Et3/1     | Grain Communications           | 369.3 MB |   6.5 GB | 189.0K |  609.1K |
| Po101     | Hurricane Electric             |  11.5 GB |  29.9 GB |   2.0M |    4.0M |
| Et19/1    | Hurricane Electric -> Po101    |  11.5 GB |  29.9 GB |   2.0M |    4.0M |
| Po104     | Hyperoptic                     |   1.4 GB |   5.0 GB | 335.9K |  564.1K |
| Et7/1     | Hyperoptic -> Po104            |   1.4 GB |   5.0 GB | 335.9K |  564.1K |
| Po107     | Kcom Group                     | 911.3 MB |  12.3 GB | 452.7K |    1.2M |
| Et25/1    | Kcom Group -> Po107            | 911.3 MB |  12.3 GB | 452.7K |    1.2M |
| Po108     | Linode                         |   3.2 GB |   2.0 GB | 755.5K |  271.5K |
| Et24/1    | Linode -> Po108                |   3.2 GB |   2.0 GB | 755.5K |  271.5K |
| Et4/1     | M247                           |   3.8 GB |   9.9 GB | 750.2K |    1.1M |
| Po105     | Microsoft                      |  30.5 GB |  13.7 GB |   5.3M |    2.1M |
| Et21/1    | Microsoft -> Po105             |  30.5 GB |  13.7 GB |   5.3M |    2.1M |
| Et17/1    | OrbitalNet                     | 930.8 MB |   2.2 GB | 199.0K |  240.3K |
| Po102     | StackPath (Highwinds)          |   9.1 GB | 181.4 MB | 819.8K |   64.2K |
| Et2/1     | StackPath (Highwinds) -> Po102 |   9.1 GB | 181.4 MB | 819.8K |   64.2K |
| Po106     | Trooli                         |  81.1 MB |   3.3 GB |  81.9K |  301.0K |
| Et23/1    | Trooli -> Po106                |  81.1 MB |   3.3 GB |  81.9K |  301.0K |
| Et26/1    | Voneus                         |  17.7 MB | 151.1 MB |   8.7K |   14.9K |
+-----------+--------------------------------+----------+----------+--------+---------+
```

LONAP™

Comments? Questions?