VERA C. RUBIN OBSERVATORY

The Network Saga

and where we are now

DM All Hands Meeting

Hernan Stockebrand
Julio Constanzo

7 August 2023

NSF AURA U.S. DEPARTMENT OF ENERGY SLAC CHARLES AND LISA SIMONYI FUND FOR ARTS AND SCIENCES LSST CORPORATION

Less Layers

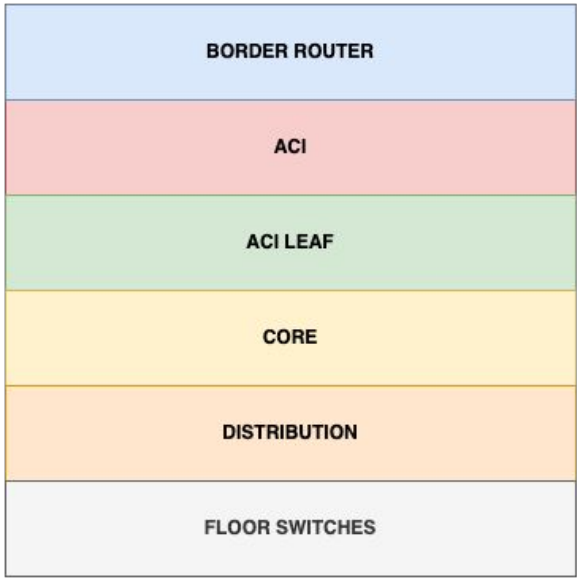

ACI WORLD

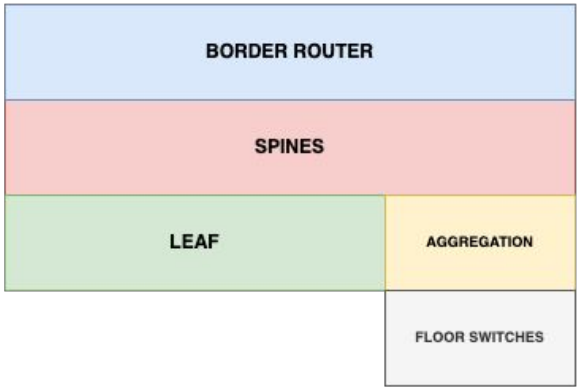

NOW

Improvements

- Less hops
- No Spanning Tree storms
- No Multicast Issues
- No Downtime

- Physical and Logical Redundancy
- One Routing Protocol

Actual Topology

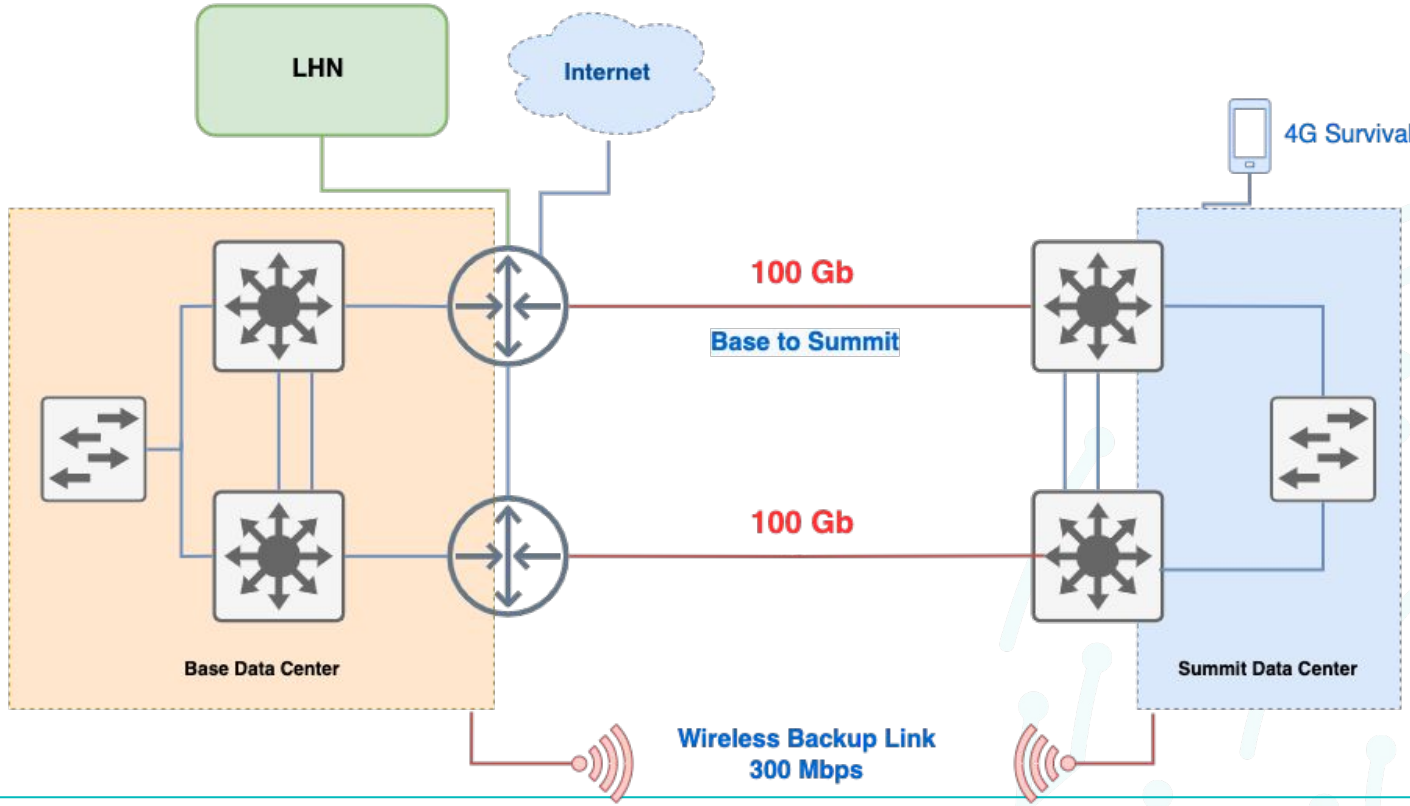

Auto-Deployment and Orchestration

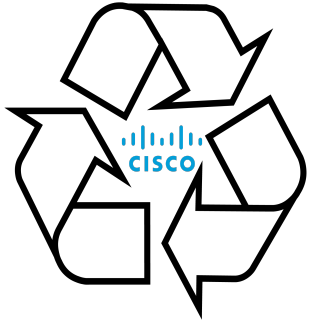

Migration Process

- 19 devices converted from ACI to NEXUS
- 34 devices configured with ZTP/POAP

Infrastructure as a Code

- 87 devices managed with Ansible
- 94 devices on Twice daily backup
- Software Upgrade
- Massive common configuration

Monitoring

- 183 Devices Monitored
- Call, email and instant messaging notifications

Monitoring


General / PING LS-CP
Ping LS-CP
Ping LS-CP
1.40 ms
10.150.1.2
1.24 ms
10.150.1.6
4.37 ms
10.150.7.6
Uptime


LHN Monitoring
Telconor
1.181 kb/s
TELCONOR - IN
1.579 kb/s
TELCONOR - OUT
LHN Monitoring
448.4 kb/s
LHN-01-IN
4.5 Mb/s
LHN-01-OUT
3.9 kb/s
LHN-02-IN
2.8 kb/s
LHN-02-OUT


Base to Summit Link - IN
Summit Data Center Temperature
17.7 °C
Base to Summit Link - OUT
Total Base to Summit - IN
31.2 Mb/s
Total Base to Summit - OUT
21.6 Mb/s
Telconor IN
1.784 kb/s
Telconor OUT
2.985 kb/s
100G Base to Summit
100G-01-IN
100G-01-OUT
100G-02-IN
100G-02-OUT
Telconor
771 b/s
TELCONOR - IN
1.608 kb/s
TELCONOR - OUT


LibreNMS
Arista Networks EOS version 4.28.3M running on an Arista Networks DCS-7260CX3-64
Device Group Membership
BDC-NET
Overall Traffic
Processors
Memory

Auxtel Run over Wireless Link

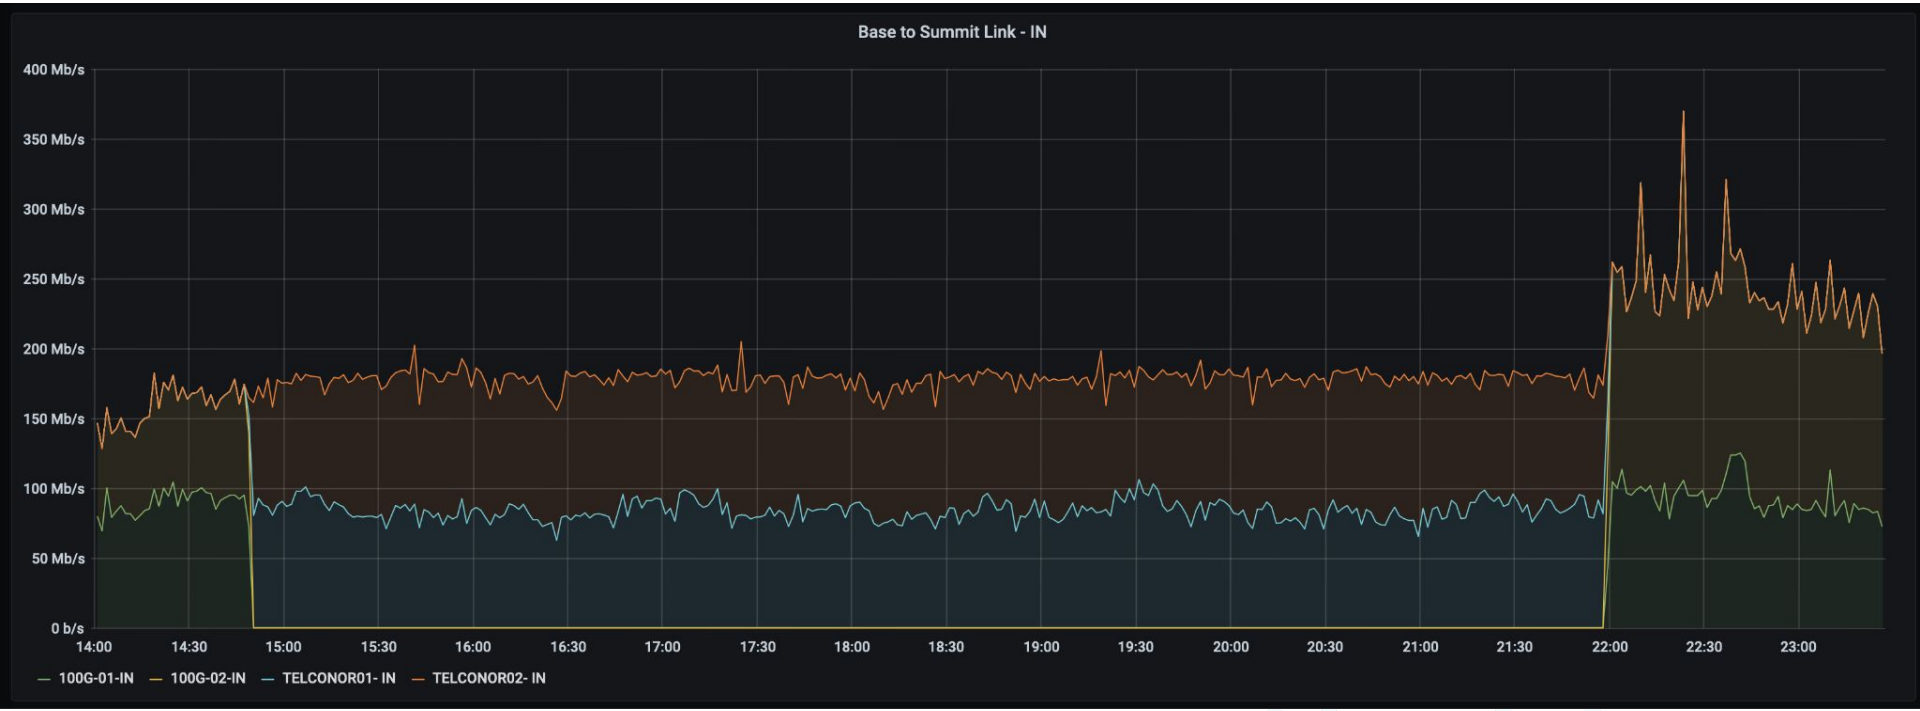

Summit - Base Link Improvement

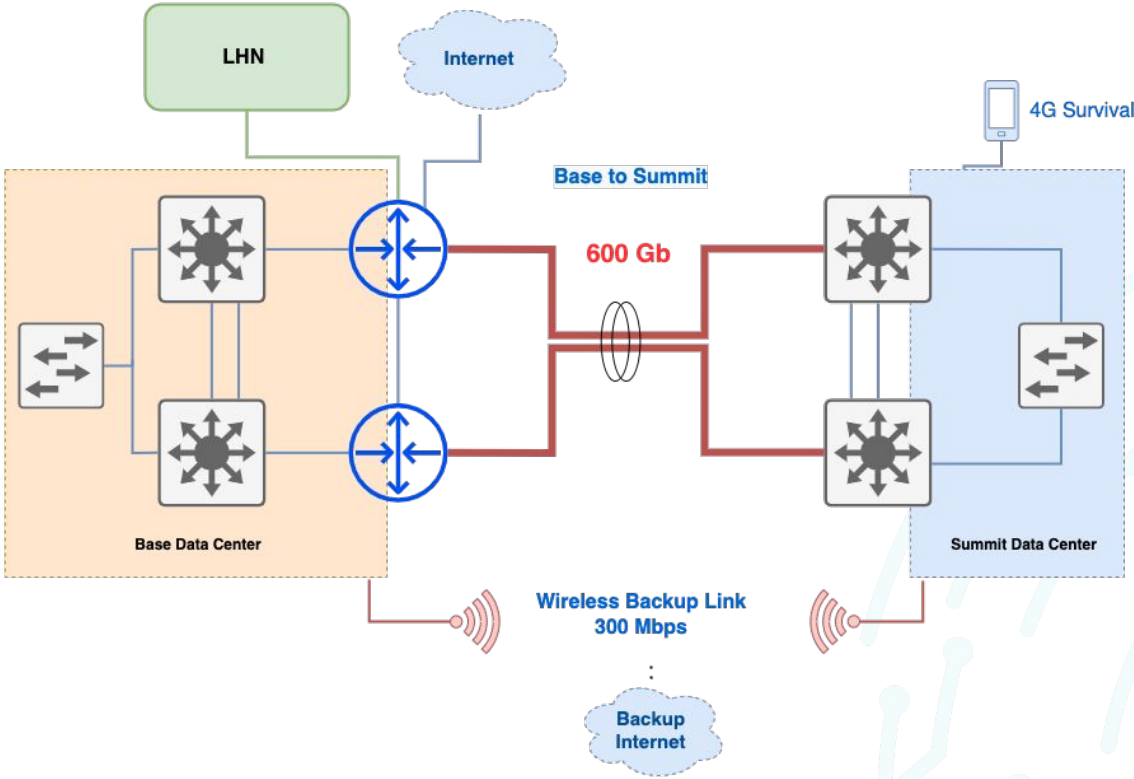

Where we are with LHN?

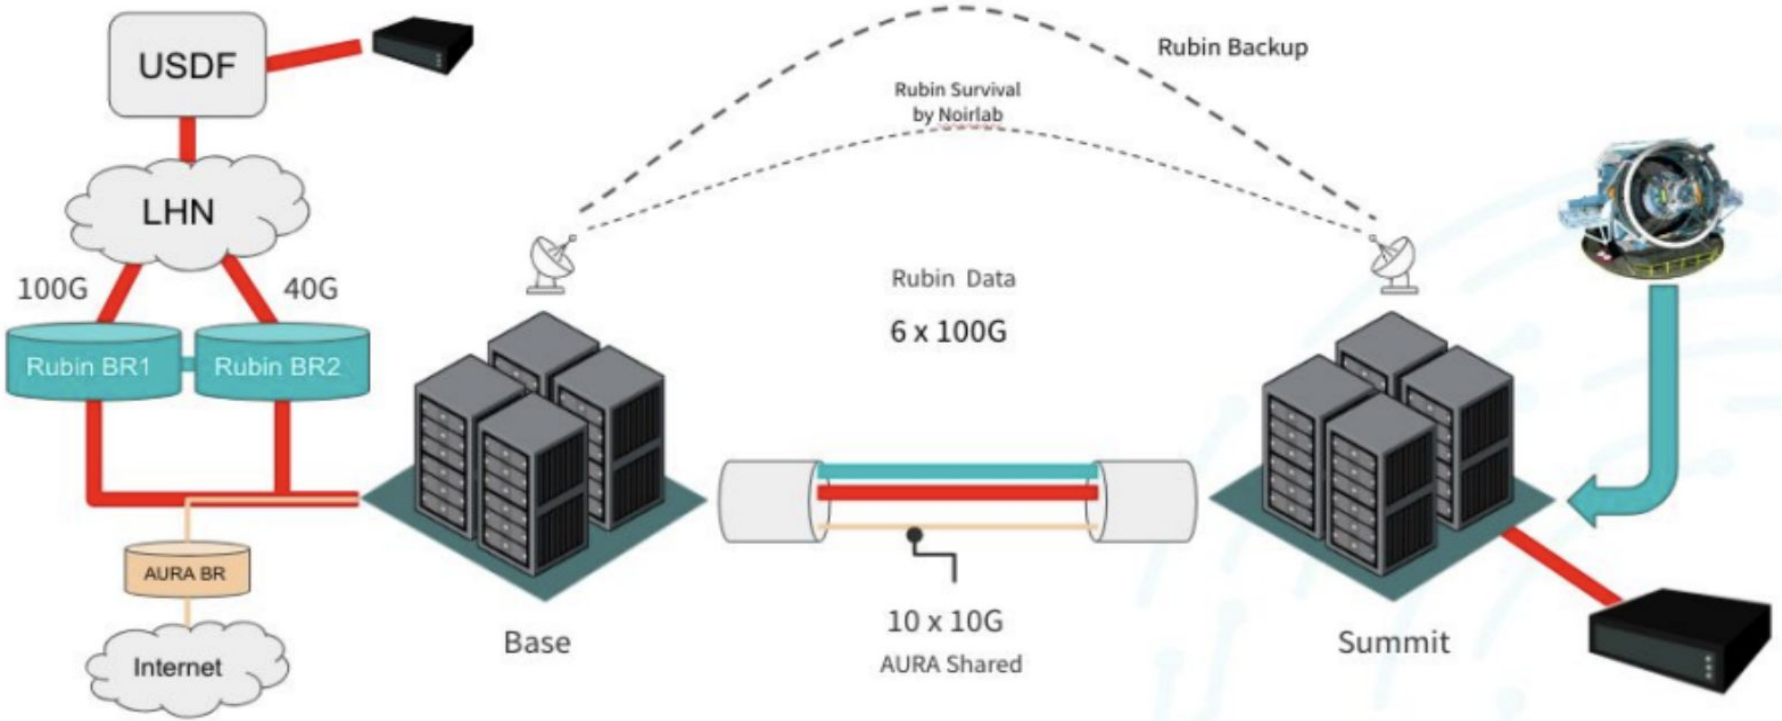

LHN Access Data Flow

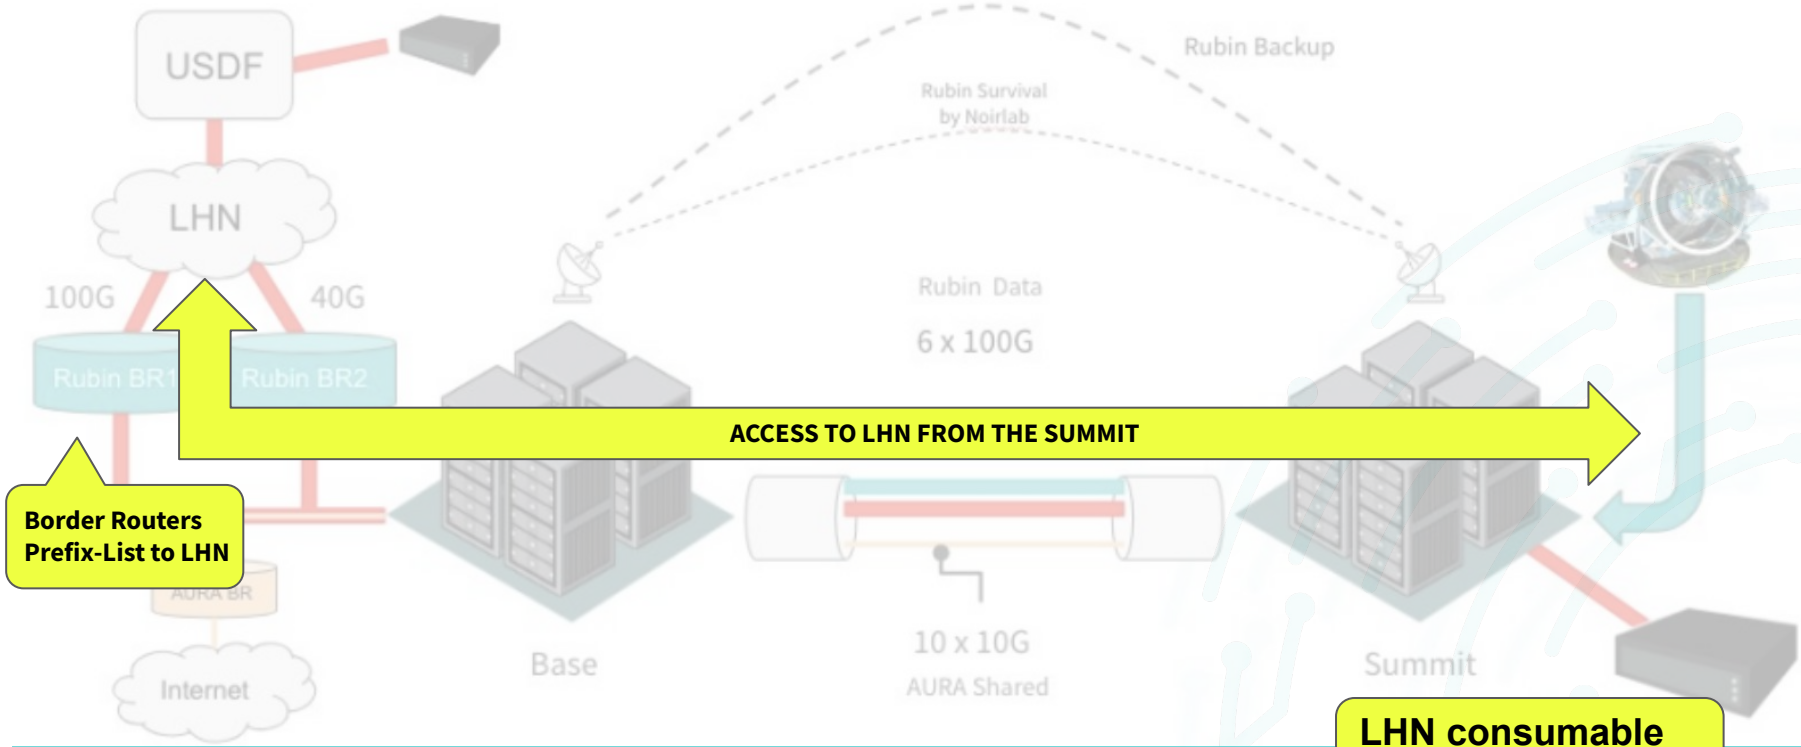

LHN Upgrades

DWDM line cards upgrade from:

- 40x10 Gbps to 4x100 Gbps. Total of 600 Gbps from summit to base.
- ATM only 2x100Gbps deployed.
- 4x100 Gbps line cards are awaiting for the new Border and Encryptions nodes to arrive to La Serena (Late 2023)

Looking if you deployments are using the LHN access: https://confluence.lsstcorp.org/display/IT/LHN+Prefixes – Take a look here!

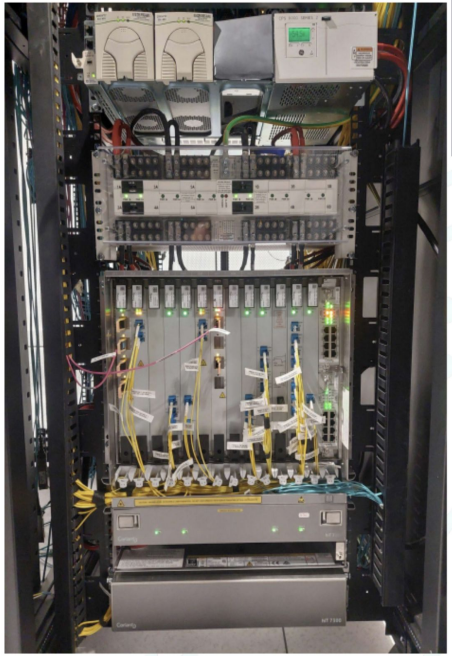

Advertised by SLAC

prefix	purpose	affected servers	added on
134.79.20.0/23	Pixel Data	134.79.20.58 - sdf-dtn01.slac.stanford.edu 134.79.20.60 - sdf-dtn10.slac.stanford.edu	01 Sep 2022
134.79.23.0/24	Pixel Data	134.79.23.41 - sdfdtn001.slac.stanford.edu	01 Sep 2022
134.79.235.224/28	Perfsonar	134.79.235.226 - Psnr-oct-v40.slac.stanford.edu	01 Sep 2022
134.79.235.240/28	Personar	134.79.235.242 - Psnr-oct-v41.slac.stanford.edu	01 Sep 2022

Advertised by Rubin

prefix	purpose	affected servers	added on
139.229.140.0/27	Forwarders	not used, could be removed.	01 Sep 2022
139.229.140.130/31	DTN	139.229.140.131	01 Sep 2022
139.229.140.132/31	DTN	not used	01 Sep 2022
139.229.140.134/31	Base Perfsonar	139.229.140.135 - perfsonar1-360.ls.lsst.org	01 Sep 2022
139.229.140.136/31	Base Perfsonar	139.229.140.137 - perfsonar1-370.ls.lsst.org	01 Sep 2022
139.229.153.0/24	BTS		27 Jan 2023
139.229.164.0/24	Wifi (summit only)	ssid: Rubinobs-LHN	27 Jan 2023
139.229.165.0/24	All Sky	139.229.165.6 - auxtel-archiver.cp.lsst.org	01 Sep 2022
139.229.175.0/26	Comcam	139.229.175.3 - comcam-fp01.cp.lsst.org	15 Sep 2022
139.229.175.64/26	LSSTcam	139.229.175.65-74 - lsstcam-dc01-1-.cp.lsst.org	15 Sep 2022
139.229.175.128/25	Auxtel	139.229.175.131 - auxtel-fp01.cp.lsst.org	15 Sep 2022
139.229.180.0/24	Yagan		27 Jan 2023

IPsec & Measurement-Monitoring Infrastructure

What's Next?

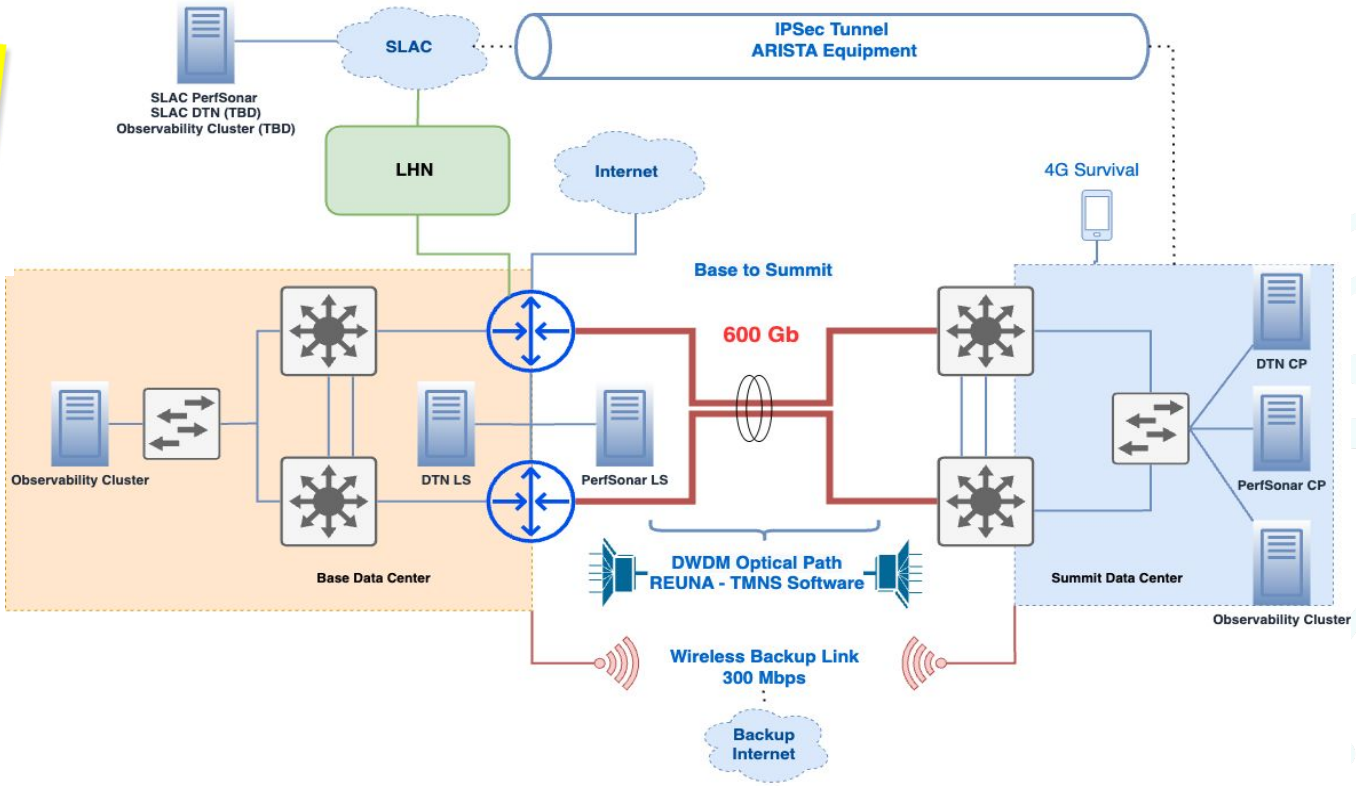

Measurement & Monitoring Tools

Monitoring and measurement tools deployed at summit - work in progress:

- Rubin CP&LS perfSONAR connected at 10/100 Gbps to Rubin Border Routers for testing measurement and monitoring of LHN access from summit-to-base & summit-to-USDF.

- Preliminary network testing using iperf3 between summit-to-base, base-to-USDF and summit-to-USDF.

DTN LS and CP nodes connected at 40/100GbE — scheduled testing to avoid competing with science data.

Arista built-in monitoring feature to explore.

Observability Cluster framework in progress.

SNMP/Monitoring proxy for VNOC access

perfsonar/maddash

The Monitoring and Debugging Dashboard (MaDDash) is a tool for collecting large amounts of inherently two-dimensional data and presenting it...

Prometheus

Grafana

slack

ARISTA

LibreNMS

squadcast

USDF perfSONAR Infrastructure

SLAC Border LHN perfSONAR at 100GbE, 10GbE

Rubin USDF internal perfSONAR TBD

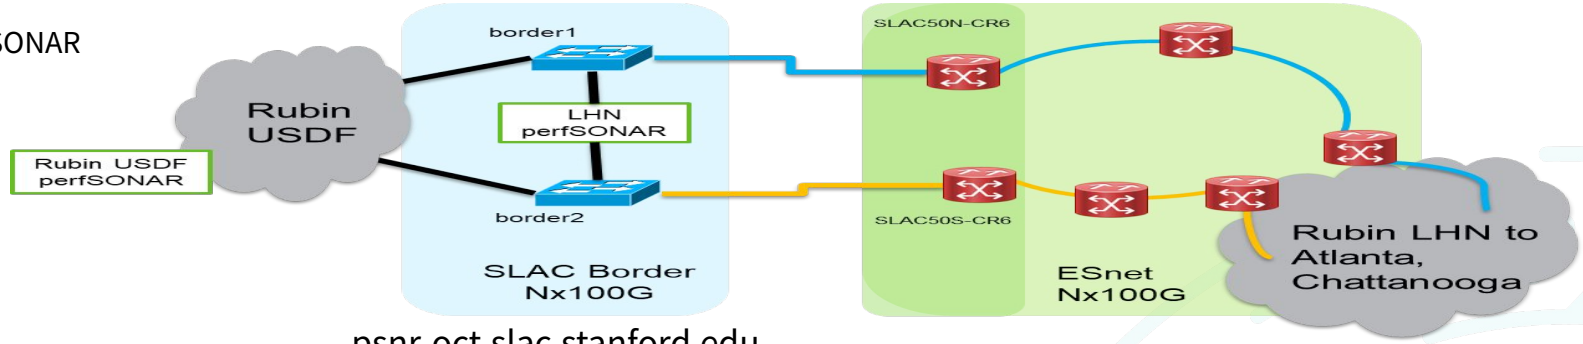

psnr-oct.slac.stanford.edu

```
[SUM]   0.00-7.00   sec  7.68 GBytes  9.43 Gbits/sec  43271             sender
[SUM]   0.00-7.00   sec  7.58 GBytes  9.30 Gbits/sec                  receiver
```

From just one pod inside Yagan cluster to Base — iperf3 preliminary testing in 7 sec using 100 TCP streams

```
[SUM]   0.00-7.01   sec  5.25 GBytes  6.43 Gbits/sec  117               sender
[SUM]   0.00-7.01   sec  5.12 GBytes  6.27 Gbits/sec                    receiver
```

From just one pod inside Yagan cluster to USDF — iperf3 preliminary testing in 7 sec using 100 TCP streams

Questions?