

**Internet2 AL2S
Triannual Report
12-1-14 through 4-1-14**

Introduction

This is the fourth in a series of periodic reviews of the Internet2 AL2S infrastructure. The goal is to examine what is working and what isn't with an eye toward improvements and a focus on information important to the community for their decision-making process.

The raw data that was used to create this report is contained in the AL2S Weekly reports, compiled in to spreadsheet form.

Report Overview:

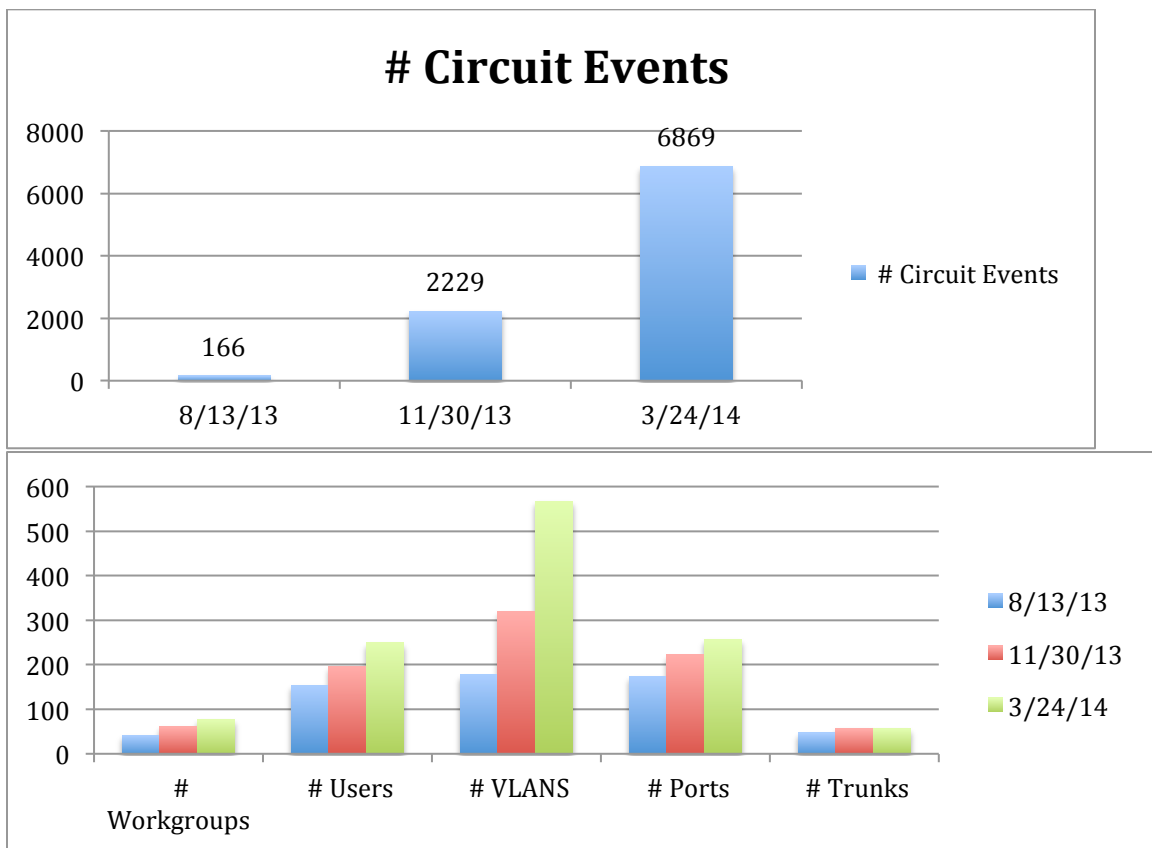
- OESS Utilization*
- Bandwidth Utilization*
- Upgrade Details*
- Future Plans for Layer 2*
- Unscheduled Outages/Incident Management*
- Scheduled Outages/Change Management*
- Post-Mortem Improvements*
- Availability Information*
- Brocade*
- Progress from Last Report Deliverables*
- Timeline of Installs/Upgrades*

OESS Utilization

OESS usage, the manner in which users configure VLANs across the infrastructure, has seen a decent uptick over the reporting period. We continue to add workgroups, users, and managed ports at a steady pace. We're up to 567 VLANs up/up.

There were 6,869 circuit events in this reporting period, about 3 times more than the last period. The bulk continues to come through the OSCARS IDC. There were 24 Workgroups, most representing major members, that configured more than 15 VLANs over the last reporting period, and five with more than 50 VLANs configured.

OESS usage is the major metric by which the success of the service is being tracked. Both the Static and Dynamic VLAN services seem to be getting good uptake. Based on Usage Statistics, or lack thereof, we will be reaching out to one Workgroup to offer additional training.



Bandwidth Utilization

The last report predicted the busy EAST-WEST corridor from Chicago to DC would need an upgrade over the last fourth months. The trend proved accurate and the circuits are in the process of being upgraded.

Trending the bandwidth four months out reveals one circuit which may need to be addressed. PHIL-WASH is currently only running at 25% utilization however it's seeing a large increase in traffic resulting in a steep slope of the trending line. If the slope remains the circuit will need an upgrade in about ten weeks. We're anticipating the upgrade in this next 16 week period and are doing an analysis of NEWY-PHIL to determine if it would be better to bypass PHIL with the upgrade, turning up a cut-through between NEWY-WASH to take the pressure off of PHIL-WASH.

We're going to expand our Capacity Planning further in the next period to provide better integration of port and slot usages, inventory, and the rapid deployment pool. We're also going to begin to integrate sflow records from AL2S in to our Netflow Analysis system, assuming the vendor sflow support arrives when we believe it should.

Upgrade Details

New nodes were added in Charlotte, Philadelphia, Baton Rouge, Jackson MS and Dallas based on user demand. Along with these new nodes eighteen new user connections were added. This includes a brief flurry of special activity in late February/Early March due to changes in the R&E networking environment.

There was one vendor software load (Brocade or Juniper code) during the past four months. This took place in the closing weeks of the reporting period and allowed for the Juniper TAC to move to a mainline support program for the features. The previous reporting period had no code loads. This points to the relatively stable nature of the Openflow code coming from the vendors.

There were two major upgrades to the OESS software during the reporting period: 1.1.1 and 1.1.2. 1.1.1 primarily added support for multi-point circuits, a major new feature, while 1.1.2 fixed several small non-service impacting bugs.

Future Software Plans for Layer 2

Vendor Software

Brocade: Internet2 plans to move to a GA release of the Brocade 5.6 code train (5.6c or later) in Q2 2014, once feature evaluation and testing is fully complete. Promised drivers include sFlow, converged layer2 and 3 matching, bigger table sizes, and improved flow mod performance.

Virtualization

Internet2 successfully demonstrated FlowSpace Firewall (FSF 1.0.1) at ONS 2014 on NDDI. Internet2 will deploy FSF 1.0.1 as part of the AL2S software stack in Q2 2014.

Core Functionality

Internet2 will deploy OESS 1.1.5 in late Q2 2014 to support ___. (ERIC TO UPDATE THIS AFTER 2 PM ROADMAP MEETING.)

Internet2 continues to evaluate support for OpenFlow 1.3, including controllers. Future support for many features (e.g. QnQ, QoS, Multipoint via MAC Learning) are gated on OpenFlow 1.3 support in the controller and the vendor software. We expect a have a fully developed plan in the next four months.

Unscheduled Outages/Incident Management

The largest impact to availability for AL2S remains maintenance or fiber cuts on the underlying L1 circuit infrastructure. Of the fifty Incidents over the last four months only about 18% of them were initiated by L2.

The five hardware failures over the last four months has been in line with the previous reporting periods, with seven in the last period and again five in the period preceding that. This would be seventeen in the last year, or about 1 per year for every two nodes, with about 50% of the failures being potentially service impacting. More analysis needs to be given to the role of backbone vs. client failures and the real-world impact of various fails on the community.

We did see an OESS bug in this reporting period. A new version of OSCARS issued a more thorough teardown request. This triggered OESS to issue a teardown with the “undefined” VLAN id, which the controller and switches interpreted as a wildcard vlan ID. Node Downtime was 1h47m.

There was also an instance of the Juniper RE's not failover correctly, which resulted in 1h18m of node downtime while an engineer manually failed over the RE.

Brocade details can be found in that break-out section.

Scheduled Outages/Change Management

During this reporting period the improvement in our ability to enact change continued: fewer changes were rescheduled. The only reschedule was a fiber event performed by Zayo. In addition there were fewer changes that delivered unexpected results, although this may be related to the generally lower rate of change. The Juniper code upgrade resulted in some differing behavior of their XML data. This had the potential of impacting only a single alarm, but we are following up by questioning our pre-upgrade testing.

Our process for tracking software changes doesn't seem quite correct. While in the past a change could be considered a loss of service, albeit one not counted by industry best practices, the separation of the control plane from the data plane means that many software changes are not actually impacting to the data plane. An analogy might be swapping out the primary or backup supervisor/RE card: you lose some control functionality briefly however there's no impact to data transport services. It's unclear if the loss of provisioning during these times could be counted separately, or even if it should be separated as a separate service. What to do, what to do?

The active management and mature software continues to show in the figures as numbers hold stable or move in better directions.

Post-Mortem Improvements

Two incidents during this reporting period have resulted in post-mortems being conducted, a goal from the last quarter. The following is a list of improvements implemented in the AL2S cycle:

There was an incident in Houston on February 5th. During the incident packets transiting through the switch were dropped without the node alerting. The incident was detected through the SmokePing infrastructure put in place after the Sept 3rd incident. As a result of this incident number of changes have been implemented in operational process. There's now an emphasis on collecting diagnostic data quickly, with about a 1-hour window given, prior to service restoration efforts. This is a compromise position between Brocade's desire to explore a root-cause live and the need for service restoration. To supplement this goal we worked with Brocade to improve their diagnostic collection scripts to shorten the time to run and to target certain error conditions with specific subset of commands. We have also modified the Incident process to bring in Software & traditional network engineer jointly and improved alarming on syslog events related to fabric switch modules.

There was an incident in Denver on February 20th. SmokePing and ISIS adjacency alerts were raised for the DENV-SALT backbone, with traffic not passing from DENV to SALT although SALT to DENV was having no trouble. The node did not raise any alarms. A reboot restored forwarding and the line card was RMA'd. As a result of this incident a BFD-like feature is being added to OESS so that it can detect one-way link and fail over circuits automatically. The data collection portion of the Workflow was also reemphasized, to assist in root-cause identification.

Availability

After twelve months of looking at Availability figures and graphs I am convinced that the metric, as currently measured, is useless. What we measure is not helpful and it is difficult if not impossible to derive actionable data from it, and yet the metric is de rigueur. The circuits that interconnect the nodes make up the majority of the data and yet those figures are troubled because of the impact of Layer 1 issues and the ability of the infrastructure to fail over to other paths. We're paying the vast majority of attention to figures that we have no control over and, because we're smart and know we have no control over them, we engineer around the problem with failover paths. This is absurd.

When we think about availability the real question is: can the user reach the services they need through our infrastructure? Does it really matter if there is a Scheduled Maintenance, the removal of which from metrics is a well-established industry practice? Perhaps in a less transparent backbone the figures would provide some indication, if they could be trusted to not be manipulated, of issues on the backbone. Internet2 attempts to very transparent, disclosing details on every issue, scheduled or unscheduled. In this environment what is the purpose of an availability figure whose primary variable is fiber cuts?

We're committed to have a better understanding of Availability metrics/measurement in the next quarter and having a plan ready to implement by August with something that is meaningful to the community and useful.


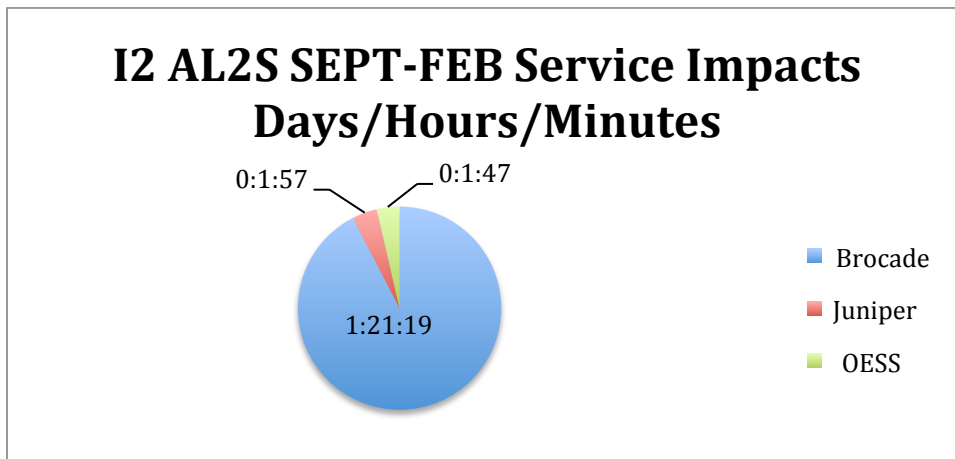
Brocade

Our interactions with Brocade have increased significantly over the last four months. We now have executive level contacts every three weeks as well as twice-weekly improvement calls in addition to our normal weekly vendor management call.

Our engagements at the executive level have centered around the root causes of our troubles and efforts to mitigate them through better failover of the Brocade hardware. As a result of these conversations Brocade has a timeline of new features to release over the next year that greatly improve the ability of the nodes to detect error conditions and fail the hardware, mitigating the Severely Degraded Data Paths issues I2 has encountered.

Our case reviews and twice-weekly calls with Brocade have centered on joint efforts to improve MTTR after each event. This involves improvements to process, error detection, and restoration. As the graph below shows, there has been a significant reduction of MTTR due to these efforts, and we're now hitting root causes in about 30%-40% of our cases.

We continue to work with Brocade on what appears to be a higher than expected level of hardware failures. We have a physical layer review scheduled for the next reporting period to help further mitigate MTTR issues on hardware fails.



Goals: April 1 2014 – August 1, 2014

- **Figure out Availability.** Over the next four months we're going to work out a plan for Availability reporting that is meaningful and useful to both Internet2 and the member community.
- **More Capacity Planning.** As usage ramps up we will need to be aware of potential service bottle necks. We should integrate our various capacity planning efforts and in the interests of transparency publish it weekly to our website. Bandwidth trending, slots, ports, CPU. Our goal is to react proactively to capacity issues and not be in a position where we are unable to meet commitments.
- **Member Outreach.** Contact Workgroups with no usage and offer additional training & resources. IE: Ensure our documentation efforts are not the reason for low usage.
- **Integrate AL2S in to our Network Analytics system.** Work with our vendors to get sflow working to our Deepfield Netflow system. Our goal is to make L2 stats available to the members and to be able to better work with traffic trends for capacity planning purposes.
- **Document fully our efforts to monitor the network.** Detail what we do and how we do it. The goal is to become more methodical about our alerting and improvements to it.
- **Document fully our Provisioning process.** Develop a process for the Workflow system for the technical portion of the AL2S provisioning process. Our goal is to better serve the members through our better understanding of the provisioning process.

Progress on Previous Report Deliverables (December 2013-March 2014)

- We plan to improve our capacity planning in the next four months. Bandwidth trending should be a routine part of reporting and we should have explicit reports noting low port & slot capacity on the AL2S infrastructure.
 - Complete. Weekly bandwidth trending numbers are now present, which allow us to plan backbone upgrades. We also have scripts that provide chassis/port capacity data.
- We plan to improve the tracking of our users usage of the infrastructure. The goal here to identify sites that could use more outreach and education, and ultimately reduce the barriers to usage.
 - Complete. We track who is using the infrastructure and weekly compare it to who is *able* to use the infrastructure, without additional outreach targeted to gaps.
- We plan to improve the descriptions of our services. Descriptions are lacking for AL2S support of OESS, dynamic, circuits, static circuits, inter-domain circuits, etc. Further the bootstrap infrastructure, NDDI, is not well documented, nor is its capabilities.
 - Complete. We've put a significant amount of effort in to the service descriptions on the new page. This includes differentiating the various services, focusing on the core VLAN service, documenting the users, etc.
- We should do a better job institutionalizing our data. The current process generally revolves around sending e-mail to various groups. Proposals and plans should be archived and available through the Internet2 website on a routine basis.
 - Complete. The process for doing this is now well-know. The distribution structure is still a little rough around the edges but the core is there to improve upon.
- The Layer 3 backbone will begin using Layer 2 (via vlans and SDN-signalled circuits) for 100g transport.
 - Complete.

Progress on previous report deliverables

Headroom

- A number of uncompleted previous goals fall in to the Headroom category. These are tied to the community efforts to better understand our joint headroom goals. An updated policy is expected to deliver in April.

Availability

- A number of uncompleted goals are tied to Availability. The patchwork effort gets partially completed by several tertiary projects, such as a Performance Assurance Systems, being completed. This was almost certainly overscoped and the new Availability goal is a nod to the inability to complete the overscoped goal.

AL2S Timeline

- 3/7/14 PacWave (SEAT) added
- 3/4/14 TR-CPS (NEWY) added
- 2/25/15 NEREN (UVM, UMS, UNH, DART) added in ALBA
- 2/15/14 MREN added in STAR
- 2/14/14 LEARN added in ELPA
- 2/14/14 UNM added in ELPA
- 2/12/14 SOX added in ATLA & DALL
- 2/6/14 FLR/Bato added
- 2/6/14 NOAA McLean added
- 2/6/14 NOAA/Denver added
- 2/6/14 UEN/SLC added
- 2/6/14 LEARN/Dallas added
- 2/6/14 MARIA/Ashburn added
- 2/2/14 OESS 1.1.2a upgrade
- 1/24/14 Dallas inserted
- 1/21/14 OESS 1.1.2 Upgrade
- 1/17/14 LEARN (HOUS)
- 1/10/14 AL2S now supporting AL3S
- 12/20/13 BATO Inserted
- 12/19/13 OESS 1.1.1.1a Deployed
- 12/16/13 OESS 1.1.1.1 Deployed
- 11/16/13 Jackson Inserted
- 12/11/13 PHIL inserted
- 12/6/13 Smithsonian, 10g, Ashburn
- 12/6/13 CAAREN 100g, McLean
- 12/3/13 Charlotte Inserted