Update from US ATLAS and AGLT2

Shawn McKee / University of Michigan

WLCG SOC Hackathon 2023

August 14, 2023

Coseners House, Abingdon, UK

Why WLCG SOC for USATLAS/AGLT2?

AGLT2 has been concerned about operational security for a long time

- Limited manpower at a Tier-2 and overloaded
- Would like to benefit from the broader community actively involved in operational security

USATLAS has also discussed how best to implement security monitoring for its distributed facilities.

WLCG Security Operations Center effort seemed like a good opportunity

- Provides example best practices, tools and docs

2018 Optical Splitter / Bro / MISP

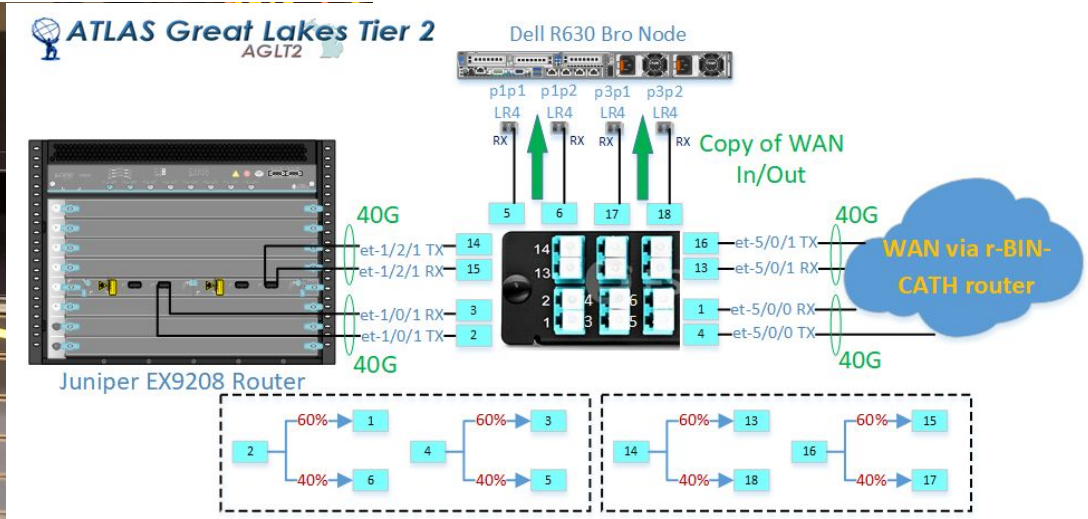

bro.aglt2.org

misp.aglt2.org

Was inexpensive to enable (~$1.2K). Splitter and shelf was $300, Intel XL710-Q2 40G nics $400 x2, $100 in cables (reused worker node for server)

Bro has been running at AGLT2 since August 10, 2018

Monthly avg of **63.1 billion** packets captured and 266 million packets lost (0.4%)

Fiber Optical Splitter Connections

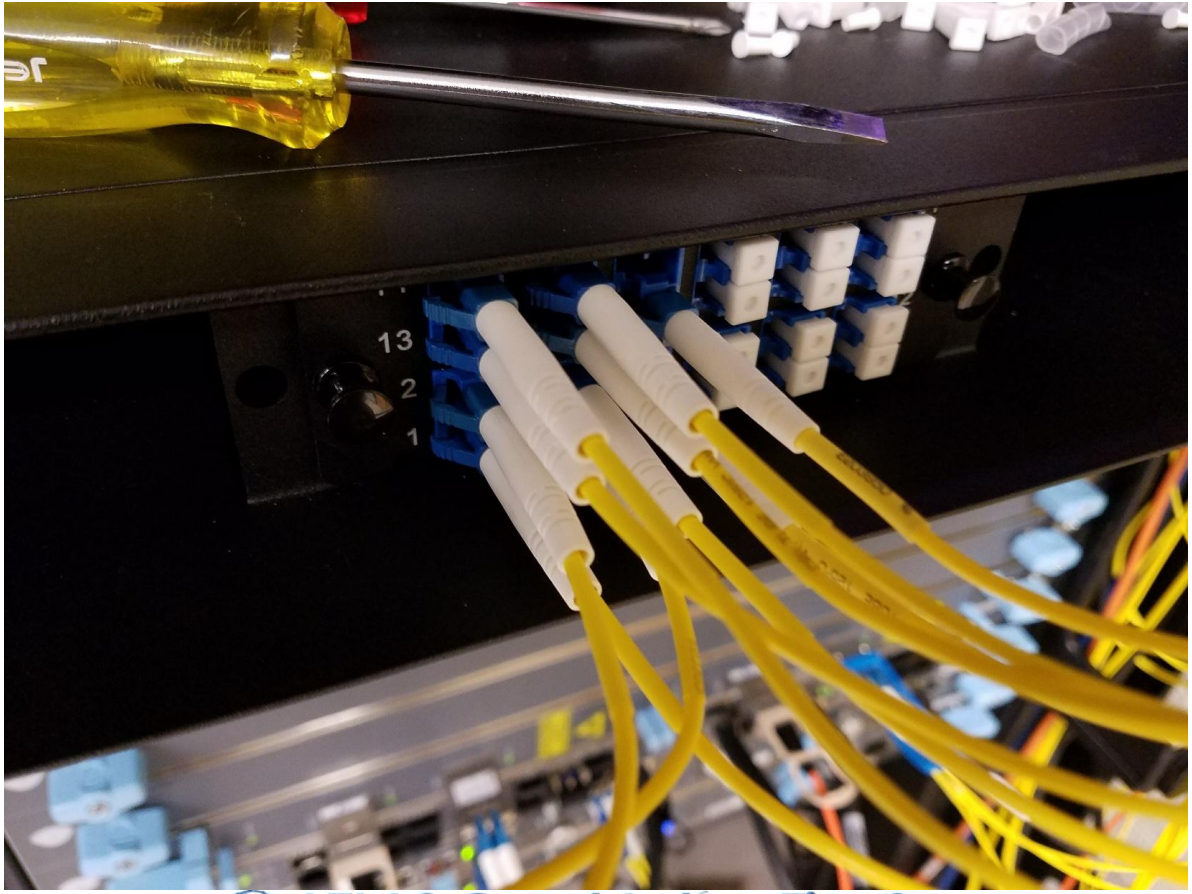

2021 Network Upgrade at UM site

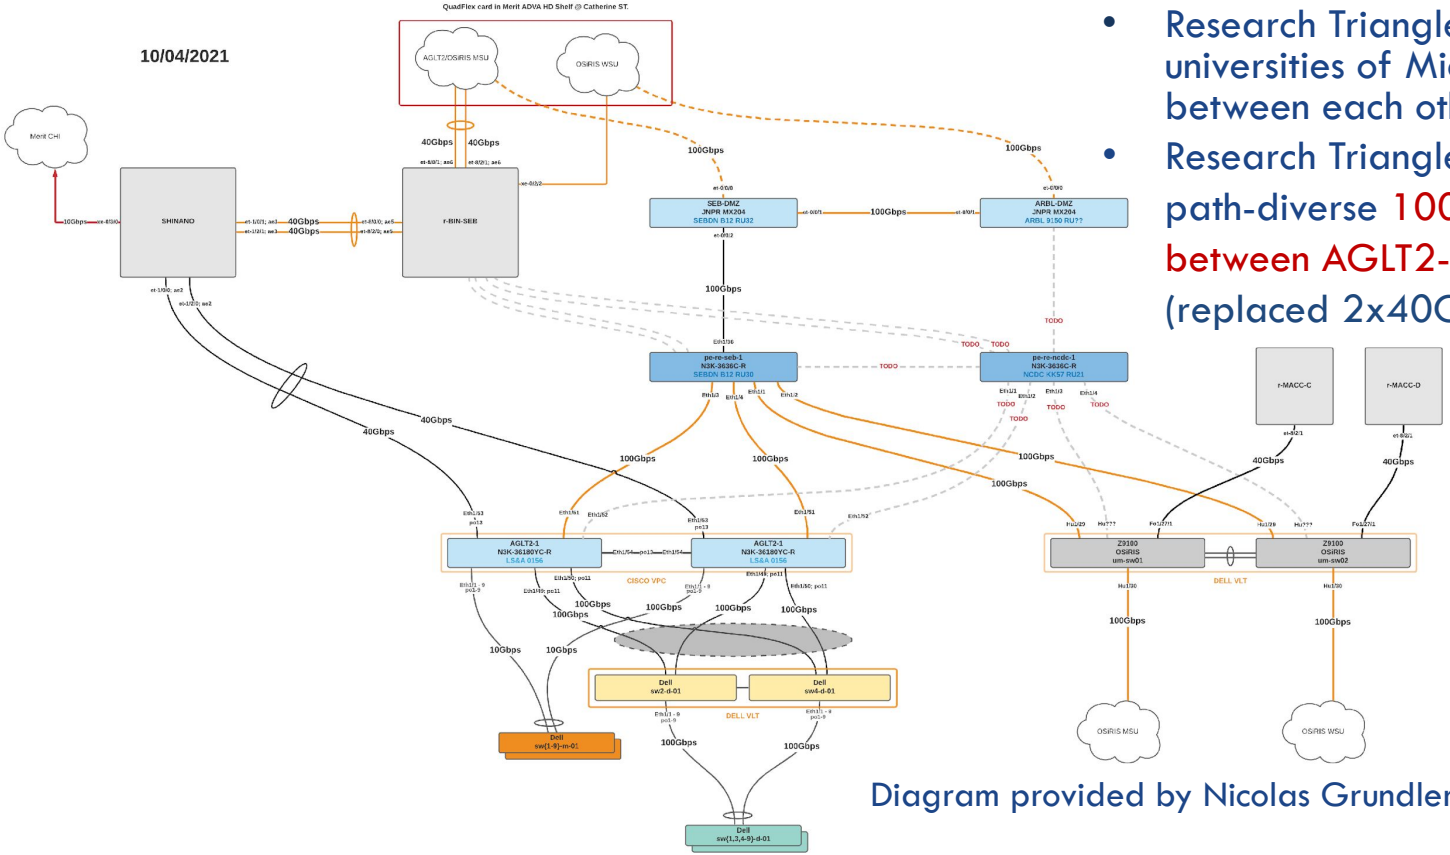

- Research Triangle Network among 3 universities of Michigan. 100Gbps between each other.
- Research Triangle adds 2x path-diverse 100G connections between AGLT2-MSU and AGLT2-UM (replaced 2x40G)

Diagram provided by Nicolas Grundler from LSA network.

2021 Network Upgrade at UM site

- Ansible and GitHub are used for configuration management
- 2 border switches from Cisco (C36180YC), management switch Dell S3048-ON, data switch Dell S5232-ON
- Each Rack has its own management switch and data switch(es), positioned at the middle of the Rack to reduce cross rack cabling.
- Management switch (Gbps) is for iDRAC and provision with different color of cables, data switch(100Gbps/40Gbps) provides dual links for data transfers.
- A dedicated switch/VLAN for switches' oob connections and a dedicated switch/VLAN for PDU/ARU management network
- Each row of Racks has one KVM switch and console
- Use slim cables (Slim RJ45, AOC breakout cables: 100Gb->25Gb*4 and 40Gb->10Gb*4 and SFP+/SFP28 transceivers)
- New PDUs have socket level meter and control
- Replaced all power cords with colors (red for single PSU, blue/black for dual PSU)
- All cables are labelled with printed labels.

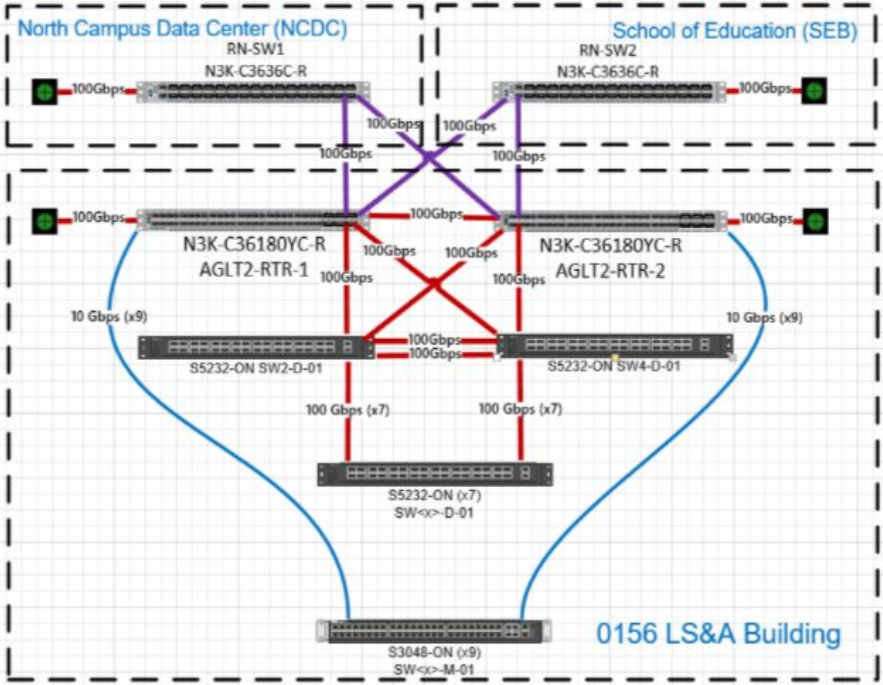

2021 New network capture nodes

- AGLT2 was running WLCG SOC network capture on our 2x40G links 2018-2021.
- Network upgrade to multiple 100G connections required us to purchase new hardware
- We bought **two** Dell R7525 systems, each with two Mellanox/NVidia Bluefield-2, dual port 100G NICs and deployed Zeek on CentOS 8 base OS. Each monitors **2** 100G fibers
  - Initial deployment and testing by summer students
- On the to-do list:

- Final in-line fiber taps in place once 100G connections are operational.
- Update to current configs based upon WLCG SOC.
- Explore the Bluefield-2 cards, with their multiple cores, for Zeek use.

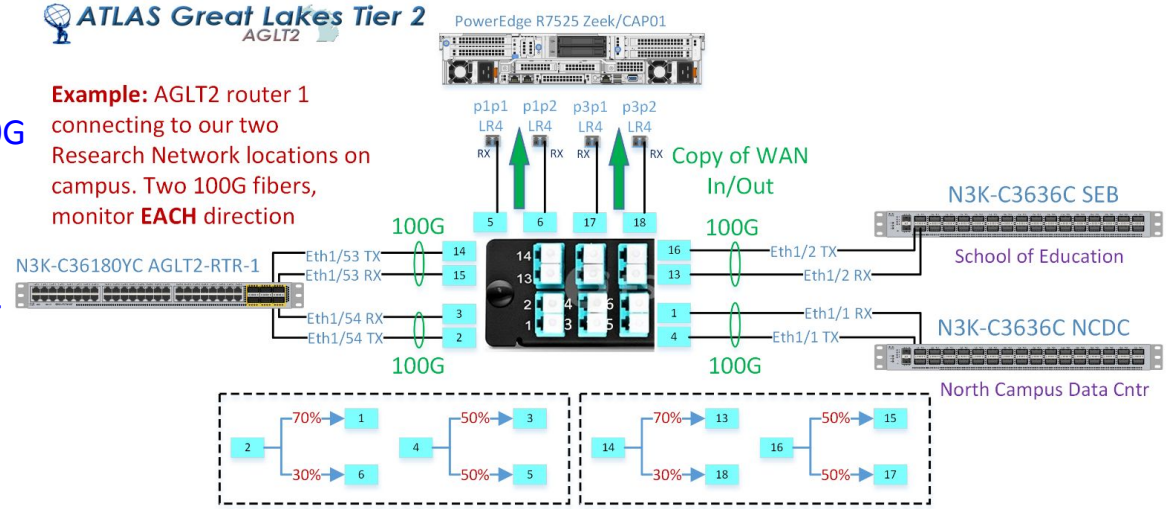

ELK at AGLT2

- AGLT2 has been using Elasticsearch, Logstash and Kibana for a few years, primarily to host a central syslogging service
- Currently we have a 4 node cluster running Elasticsearch 7.17
  - The primary storage nodes are NVMe based, 128G ram, 36 2.3Ghz HT-cores
  - Total space available is 24 TB (ZFS mirror pools on NVMe)
- Elastic search has **3560** indices, **64.8** billion documents and **1632** primary and replica shards as of today (Aug 14, 2023)
- The main data sources are 1) **syslogging** from all our devices, 2) **dCache** logs, 3) **Netflow/Sflow** and (soon) 4) **Bro log files**
- **Want to convert to Opensearch (soon)**

New Efforts at CAIDA

Frank Wurthwein / USCMS / OSG Executive Director / Director of the UCSD Supercomputing Center contacted a few of us last week about a new CAIDA NSF award:

https://www.caida.org/funding/cici-starnova/

Scalable Technology to Accelerate Research Network Operations Vulnerability Alerts

From Frank:

"It builds on the UCSD network telescope activities:

https://www.caida.org/projects/network_telescope/

The difference being, that it does so on the production network, rather than the /10 of the amateur radio folks.

I talked with K. Claffy, the CAIDA group leader, and she would be interested in collaboration with a larger community effort to share threat-intelligence gained from both. There is obvious academic and operational interest in correlating thread intelligence from the /10 of the amateur radio folks with the production network.

I'd be happy to facilitate discussions."

We should discuss how this might be worked into the WLCG SOC concepts and how we should engage.

Summary

- The AGLT2 goals is network monitoring with Zeek/MISP providing us with new info; **we need to incorporate it into our operations**
- For USATLAS, we are interested in that can help monitor and secure our osites, for teams that don't have significant security expertise or much available effort.
- Our main interest is in configuring some level of **alerting when attacks are occurring**.
  - Some way to create a report summarizing identified attacks would also be a great addition
  - Automated response is also of interest

We also need to consider how to integrate activities ongoing elsewhere (like CAIDA)

Questions ?

Additional / Backup Slides Follow

2018 Netflow/Sflow Monitoring via ELK

- In addition to Bro monitoring we wanted to have better visibility into our network traffic.
- Because we already had an ELK stack, when we heard about ElastiFlow we were intrigued
  - https://github.com/robcowart/elastiflow
  - Install https://github.com/robcowart/elastiflow/blob/master/INSTALL.md
- It was pretty easy to setup. Some challenges getting the sflow-codec and the Kibana elastiflow index imported (maybe better now?)
  - Contact me if you want details!
- Once it was setup we just needed to point our Juniper router to it

ATLAS Alarms & Alerts Service

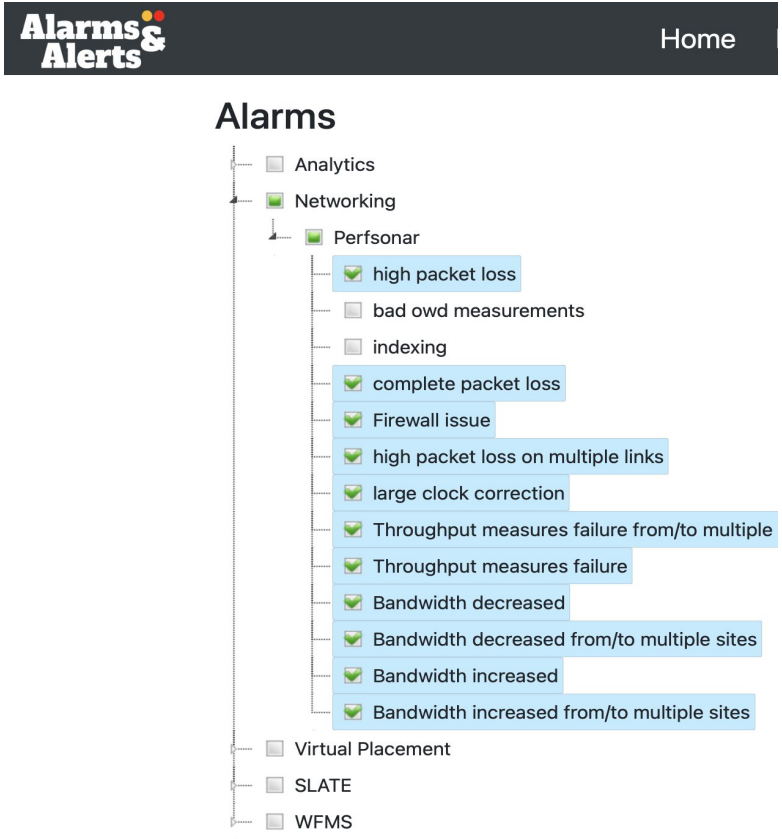

Currently available:

- Main packet loss issues
- Main throughput issues

Future plans:

- Add traceroute alarms:
  - Destination never reached
  - Path changes too often
  - Node causes issues with multiple sites

https://aaas.atlas-ml.org/

2018 Fiber Splitter Connection Details