



# AGLT2 Status Update

Shawn McKee / University of Michigan  
Chip Brock, Philippe Laurens, Mike Nila, Wenjing Wu  
USATLAS Facilities Meeting, ANL  
December 3, 2018

# AGLT2 Numbers

- The ATLAS Great Lake Tier-2 (AGLT2) is a distributed LHC Tier-2 for ATLAS spanning between UM/Ann Arbor and MSU/East Lansing. Roughly 50% of storage and compute at each site
  - 10680 logical cores
    - MCORE slots 1134 (dynamic) + 10 (static)
  - **Additional** 936 Tier-3 job slots usable by Tier-2
  - Average 10.69 HS06/slot
  - **6.9 Petabytes** of storage
  - Total of **117 kHS06**
  - Tier-2 services virtualized in VMware 5.5 (upgrading to 6.5, really!!)
- 2x40 Gb inter-site connectivity, UM has 100G to WAN, MSU has 10G to WAN, lots of 10Gb internal ports and 20 x 40Gb ports, 32x100G/40G or 64x50G/25G ports
- High capacity storage systems have 2 x 50Gb or 2 x 100Gb bonded links
- 40Gb link between Tier-2 and Tier-3 physical locations

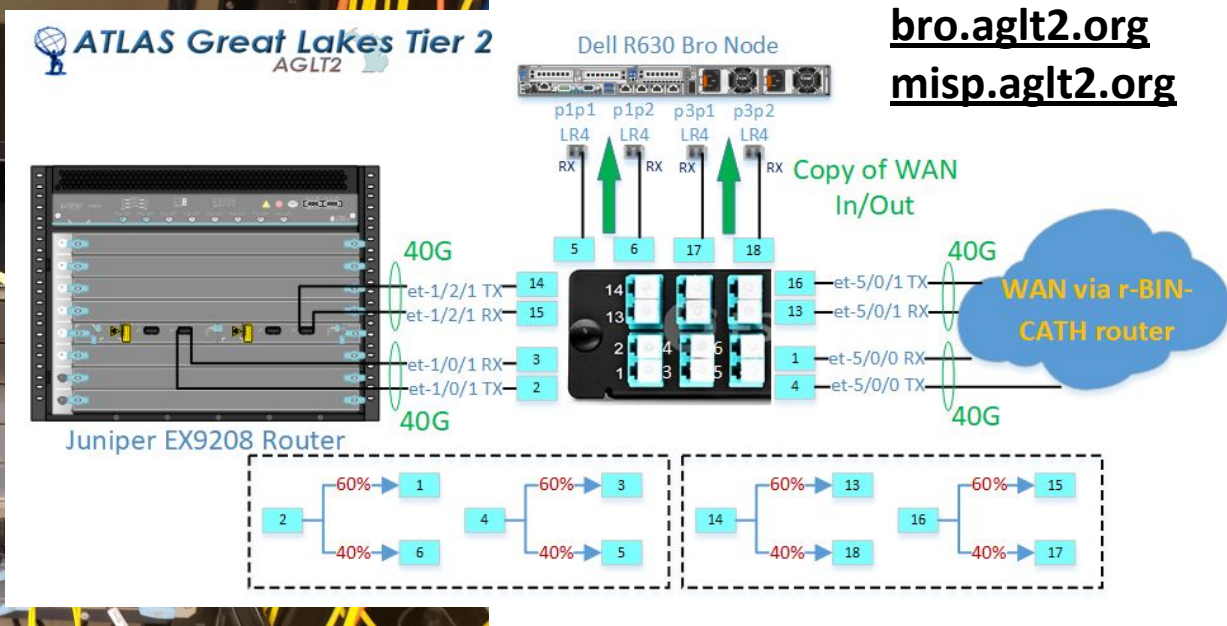
# Personnel Updates

- As most of you know, **Bob Ball** our Tier-2 Manager since we started in 2006, has retired 😞
  - He will be missed!
- **Wenjing Wu** has taken over for much of Bob's work at AGLT2 at UM..
- **Philippe Laurens** will be increasing his involvement at AGLT2 MSU and will be ramping up to 85% on AGLT2 (from 40% currently) by February 2019
- Both Wenjing and Philippe have been added as the AGLT2 collaboration board members, replacing Bob

# 2018 Hardware Additions

- For **AGLT2** we have gone with the *C6420* configuration from line 20 in [https://docs.google.com/spreadsheets/d/1vQ8lc3dTolnu--8m9HGy5tiOHr26eK-tNjOOaSV\\_EuQ/edit#gid=0](https://docs.google.com/spreadsheets/d/1vQ8lc3dTolnu--8m9HGy5tiOHr26eK-tNjOOaSV_EuQ/edit#gid=0) (\$10.78/HS06)
  - **UM** getting two C6420 enclosures and 9 servers, 1 ML3 tape drive, three R740 server to replace existing ESXi hosts (+7081 HS06)
    - Old ESXi hosts (2xE5-2670) will add 1450 HS06
  - **MSU** getting one C6420 enclosure, 7 servers, 2 PDUs (+5507 HS06)
- Earlier this year we purchase a SLATE edge node and are about to put in place an S4248-ON switch to serve the SLATE and a few AGLT2 nodes
  - Switch provides 40x10G, 2x40G & 6x100G ports with deep buffers

# SOC Work: Optical Splitter / Bro / MISP



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)  
Currently Bro has been running continuously since August 10<sup>th</sup>

**63,115,662,766** packets captured and **266,317,103** packets lost (**0.4%**) in ~45 days

# 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>
- It was pretty easy to setup. Some challenges getting the sflow-codec and the Kibana elastiflow index imported.
  - Contact me if you want details!
- Once it was setup we just needed to point our Juniper router to it

# Netflow/Sflow Monitoring via ELK (2)

- Setting up our Juniper EX9208 was pretty simple
- The configuration on the right is the bulk of what is needed
- Add additional interfaces as need (those interfaces that connect to the WAN)

```
sflow {  
  agent-id 10.10.1.2;  
  polling-interval 1;  
  sample-rate {  
    ingress 100;  
    egress 100;  
  }  
  source-ip 10.10.1.2;  
  collector 10.10.1.9 {  
    udp-port 6343;  
  }  
  interfaces xe-0/0/3.0 {  
    polling-interval 1;  
    sample-rate {  
      ingress 100;  
      egress 100;  
    }  
  }  
}
```

# ElastiFlow @ AGLT2 Examples

kibana

- Discover
- Visualize
- Dashboard
- Timelion
- APM
- Dev Tools
- Monitoring
- Management

Overview | Top-N | Flow | Geo IP | AS Traffic | Exporters | Traffic Details | Flow Records




Flow Exporter: Select... Client: Select... Server: Select... Service: Select...

Servers and Clients (bytes)



- 192.41.230.36
- 192.41.230.26
- 192.41.231.130
- 192.41.230.29
- 10.10.1.51
- 192.41.230.23
- 192.41.230.33
- 192.41.236.120
- 192.41.231.134
- 192.41.230.250
- 192.41.236.54
- 192.41.231.135
- 192.41.231.130


Services (bytes)



- exp1 (TCP/1021)
- TCP/37764
- TCP/13078
- TCP/26586
- TCP/8588
- TCP/39826
- complex-link (TCP/...
- TCP/27624
- TCP/15798
- TCP/25338
- TCP/20477
- TCP/36813
- TCP/36813

RST ACK ECE  
PSH SYN FIN CWR

Autonomous Systems (bytes)



- Merit Network Inc. (2...
- University of Illinois (...)
- University of Michiga...
- Simon Fraser Univers...
- University of Californi...
- Boston University (111)
- European Organizati...
- University of Wiscons...
- University of Texas at...
- University of Chicago ...
- Brookhaven National ...
- CHINANET SiChuan T...
- Ma...

IP Versions and Protocols (bytes)

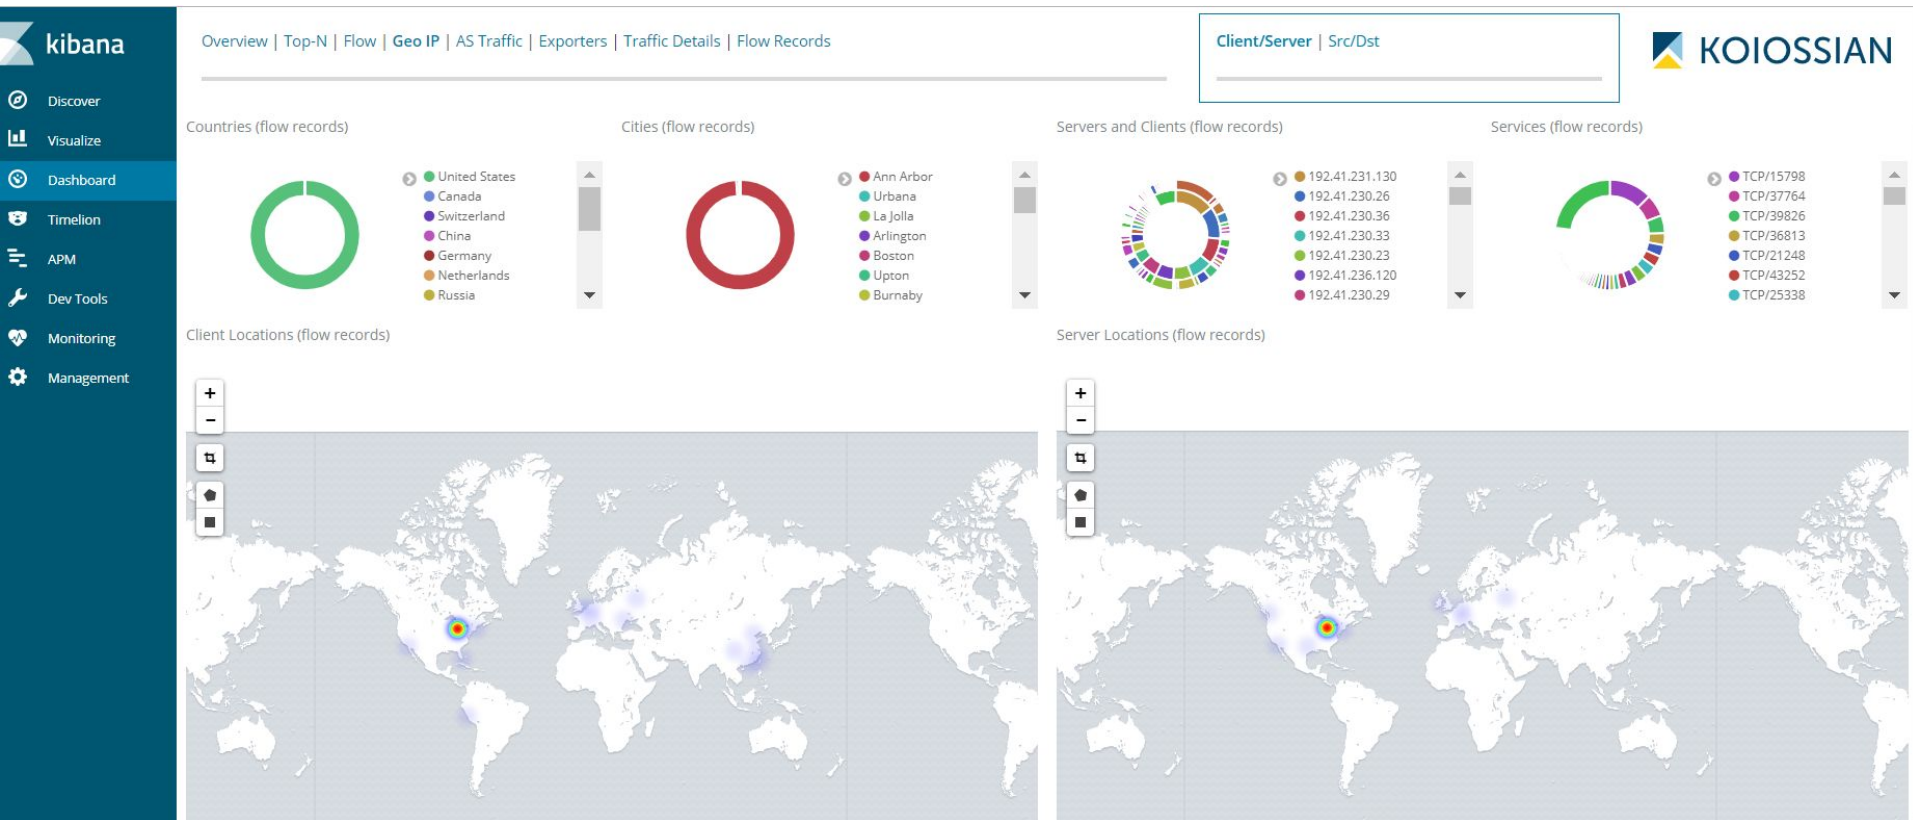


- IPv4
- TCP
- UDP
- ICMP

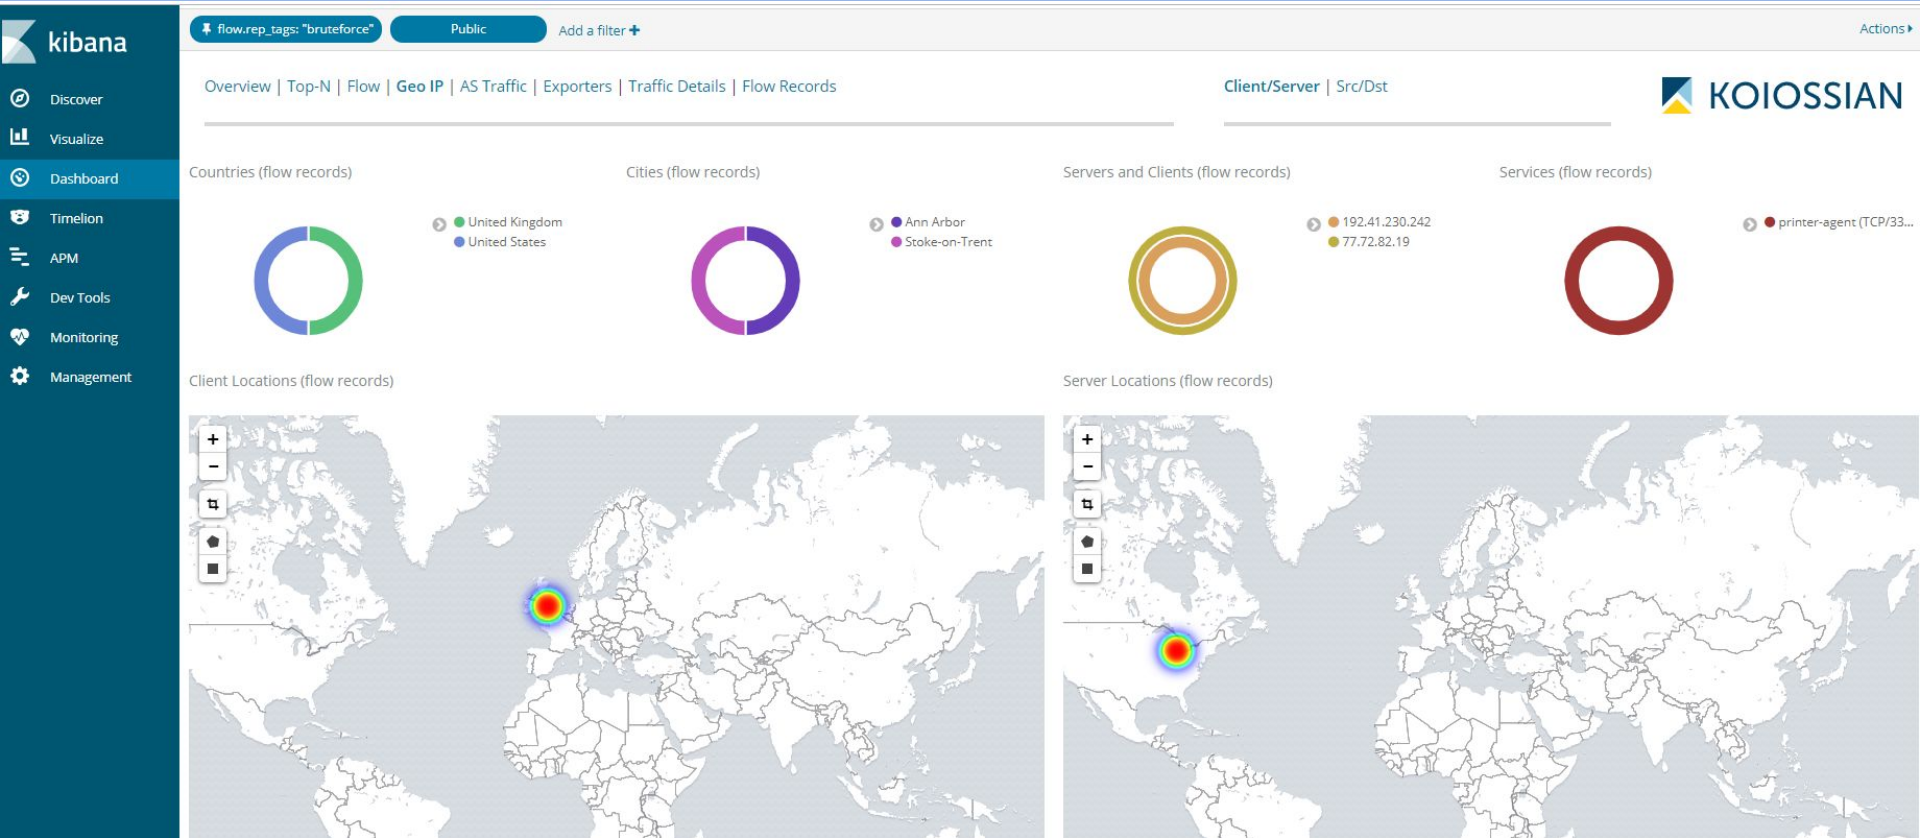
screenshare  
ddos  
bruteforce  
rdp bot ssh  
suspicious



# ElastiFlow @ AGLT2 Examples



# ElastiFlow @ AGLT2 Examples



# Near Term Plans

- SLATE containerization deployments
- Experimenting with SDN/NFV/OVS in our Tier-2 and as part of LHCONe point-to-point testbed.
  - This has been delayed but still planning to work with KIT and MWT2 once they have OVS on their production storage
- Update to VMware soon: new version (5.5->6.5, then to 6.7), new configuration for HA, new ESXi host at UM deployed
- Will begin testing dCache on Ceph using AGLT2 and OSiRIS
- Participating in the WLCG SOC (Security Operations Center) work
- Have students engaged in creating new dashboards for AGLT2 that leverage ELK for syslogs and dCache