

MetraNOVA

A Consortium for Advancing Network Observation and Analysis

Edward Balas, ESnet
Ivana Golub, PSNC




Nordunet Conference NDN2024, 11 September 2024, Bergen Norway





What is Measurement and Monitoring?

- Measuring the network over time, and creating useful things with the observations
 - maps, dashboards, capacity plans, operational alarms, annual reports
 - Whatever helps improve situational awareness, and lets you tell the story
- Typically involve a handful of data types
 - Ports stats, Network Flow Summaries, Optical Performance, Routing Tables, End to End Perf
- Includes systems many of you use today
 - perfSONAR, Nagios, Prometheus, TICK, Stardust, Netsage, Kentik, Arbor Networks, Deepfield





Measurement is about Storytelling

We all have stories to tell

- **Operational Stories:**
 - **What** just happened to the network in the northeast?
 - **Who** is the source of that giant surge in new traffic?
- **Planning/Engineering Stories:**
 - **When** will we run out of capacity between EU and US?
 - **Where** are we seeing performance degradation?
 - **What** do we need to budget for in the next cycle?
- **Community Engagement Stories:**
 - **How** has a project's use of the network changed in last year, and is there anything we can do facilitate better data transfer?
- **Value Proposition Stories:**
 - **Why** is R&E awesome? *full technicolor.*




Beyond points on a graph, a good story needs to relate back to the real world things people care about: *who, what, where, when, why, how*

- **Metadata in measurements is essential to provide this relation**



What we dream of

Networks will employ data-driven design and operations



Insight Cons-1e

https://console.internet2.edu/#?command=show+version&nodes=aggl.dal13

Community Interfaces Looking Glass

Filter by node name or location

Aggregation Router

- aggl.ashb AnBurm, VA
- aggl.dal13 Dallas, TX
- aggl.eoeh Chicago, IL
- aggl.los42 Los Angeles, CA
- aggl.newy2 New York, NY
- aggl.sanj San Jose, CA
- aggl.seat Seattle, WA
- aggl.ashb AnBurm, VA
- aggl.dal13 Dallas, TX

(2 nodes selected) show version

```
aggl.dal13 > show version
Cisco IOS Software, Version 7.3.3 (OT)
Copyright (c) 2013-2022 by Cisco Systems, Inc.

Build Information:
  Built By      : ingunwa
  Built On     : Sat Jan 20 10:33:37 UTC 2022
  Build Host   : iox-uc1-004
  Workspace    : /auto/src/archive16/prod/7.3.3/8888/vs
  Version      : 7.3.3
  Label        : 7.3.3

Cisco 8888 (Intel(R) Xeon(R) CPU D-1538 @ 2.40GHz)
Cisco 8888-025 (Intel(R) Xeon(R) CPU D-1538 @ 2.40GHz) processor with 32GB of memory
aggl.dal13 uptime is 7 weeks, 1 day, 19 hours, 6 minutes
Cisco 8882 2RU System w/ 12x48AGE QSFP56-DD & 68x180GE QSFP28

aggl.dal13 > show bfd
IPv4 Sessions Up: 2, Down: 0, Unknown/Retry: 0, Total: 2
```


show lacp
show l2vpn
show lisp neighbors
show route
show version
show vrf all
traceroute
uncheck

show version core1.denv
show bfd core1.denv



What we dream of

Combining additional types of information will provide a fuller context as to what is influencing the network





Is this a solved problem?

Multi-layer visibility with generalized infrastructure


Detecting nationwide ATM network disruption with SNMP CPU stats

Georeferencing + Data Fusion for situational awareness

Hurricane Katrina inspires georeferenced weathermap to show actual weather.

Geographic capacity planning

Metadata used to aggregate capacity and usage at continental level and combined with forecasting



We have the concepts, but lack ubiquitous technical capabilities and policy frameworks

- Scaling still difficult at intranetwork level
- Metadata tends to be variable
- Interdomain sharing is limited
- Required domain knowledge a barrier



Imagine the following scenario

- An Engineer gets a report that a customer is seeing poor performance moving data as part of a scientific pipeline
 - checks their measurement systems and local PerfSONAR results and find all clear
 - Presumes issue is likely at far end of the path close to last mile

- To support the customer effectively, the Engineer needs access to more data



Where does the engineer do next?

- **Best Case:**
 - **External networks have established measurement collections**
 - **federated auth and enough structure and documentation exist to support self service**
- More often:
 - External network has established measurement collections
 - Low fidelity data is externally accessible
 - The data needed is private, the engineer however knows his peers and can make a wetware request
- Worst case:
 - External network might itself be decentralized with no ubiquitous measurement approach
 - Multiple human interactions required to find the right engineer
 - The external engineer has to log into the router or other device to debug with you



Where does the engineer do next?

- Best Case:
 - External networks have established measurement collections
 - federated auth and enough structure and documentation exist to support self service
- **More often:**
 - **External network has established measurement collections**
 - **Low fidelity data is externally accessible**
 - **The data needed is private, the engineer however knows his peers and can make a wetware request and has to manually correlate information shared**
- Worst case:
 - External network might itself be decentralized with no ubiquitous measurement approach
 - Multiple human interactions required to find the right engineer
 - The external engineer has to log into the router or other device to debug with you



Where does the engineer do next?

- Best Case:
 - External networks have established measurement collections
 - federated auth and enough structure and documentation exist to support self service
- More often:
 - External network has established measurement collections
 - Low fidelity data is externally accessible
 - The data needed is private, the engineer however knows his peers and can make a wetware request
- **Worst case:**
 - **External network might itself be decentralized with no ubiquitous measurement approach**
 - **Multiple human interactions required to find the right engineer**
 - **The external engineer has to log into the router or other device to debug with you**



MetranOVA is here to help

- **Advocate for quality ubiquitous collections with appropriate access within all of R&E**
 - Provide training and policy guidance
 - Create knowledge base articles and howtos
- **Lower the barriers through technical and policy collaboration**
 - Reduce need for bespoke solutions
 - Amortize software sustainment costs through collaboration.
- **Retain Network Measurement as a core competency through the next generation**
 - Requires ongoing care and feeding
 - Deep domain knowledge in networking, systems, and to an extent stats
 - Support next generation of R&E engineers



The Secret Sauce of Research and Education



- **Timeless design constructs**

- Ubiquitous Access
- Loose Coupling
- Vendor Neutrality
- Open Standards
- Rough consensus and working code

- Technology != differentiator

- Same software and hardware used in R&E and in Commodity Internet
- It's how you use it, not what you use
- like an artist and a paintbrush

- **Combined with community focus**

- We are a not for profit community
- Our values differentiate us
- Its how we apply these technologies to address needs and facilitate scientific and educational endeavors.

- Additional Considerations

- Collaboration and trust are key
- Ubiquitous access tempered with appropriate access control
- We need: Design Patterns, Service Definitions, and Policy Guidance



Consortium Details

Goals

- Tools, Tactics and Techniques
- Develop and Share
 - Open Architectures
 - Technical Components
 - Design Patterns
 - Best Practices
 - Policy Recommendations.

Vision

- <https://github.com/MetrANOVA/.github/blob/main/profile/vision.md>
- A collaboratively developed ecosystem exists
- Open Source, loosely coupled, without cloud service dependence
- Solid foundation for production services and innovation
- Facilitate data driven design in engineering and operations

Executive Committee

- Provides governance and oversight.
- Decides on new membership organizations.
- Representatives from each member.
 - Inder Monga - ESNet
 - Ivana Golub - PSNC/GÉANT
 - James Deaton - Internet2
 - Luke Fowler - Indiana University GlobalNOC
 - Nathaniel Mendoza - TACC
 - Ed Balas - Consortium Lead

Participation Model

- Member Organizations
 - Requires ≥ 1 Full Time Staff Equivalent
 - Participates in governance process
- Affiliates
 - Any organization or individual able to contribute.
 - Lower bar to participate, more flexibility



What have we been up to this year?

- Community survey completed
- Established near term roadmap
 - Vetted technical stack
 - Policy guidance for appropriate data sharing
- Technical work in progress
 - Elasticsearch Time Series Data Stream evaluation
 - SNMP vs Streaming assessment
 - Science Registry refactor / up keep (TACC)




2024 State of Community Survey

- 19 responses from 18 organizations
- 42% NREN, 26% Regional, 16% Campus, 10% Lab or Facility
- Majority identify as Network Engineers, 36% as leaders and 36% as Syseng

What type of organization?

19 responses






Survey: How Measurement and Monitoring is organized

Do you have separate monitoring and measurement systems? In this context measurement is defined as the act of collecting, ingesting, storing... deviation often to network operators and engineers.


19 responses



- No, our monitoring functions are part of our measurement platform
- Yes, though that's not deliberate
- Yes, we do that for resiliency and separation
- Yes, we do it deliberately for non-technical reasons




Survey: Data in use





Survey: What people need

% who found this valuable vs. Development Area





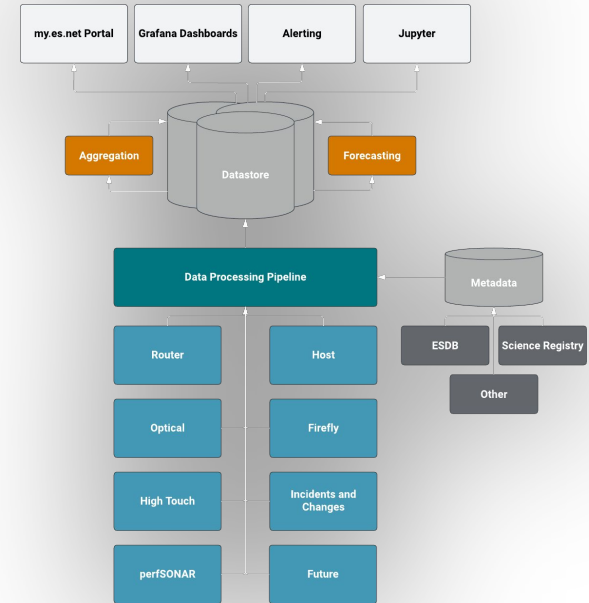
Value Engineering of Elasticsearch Datastore

Broad Elastic Adoption:

- Elasticsearch and OpenSearch used within most member networks and with in PerfSONAR
- Flow, SNMP, Optical, Open Telemetry, Streaming

A few members motivated to explore improved scaling

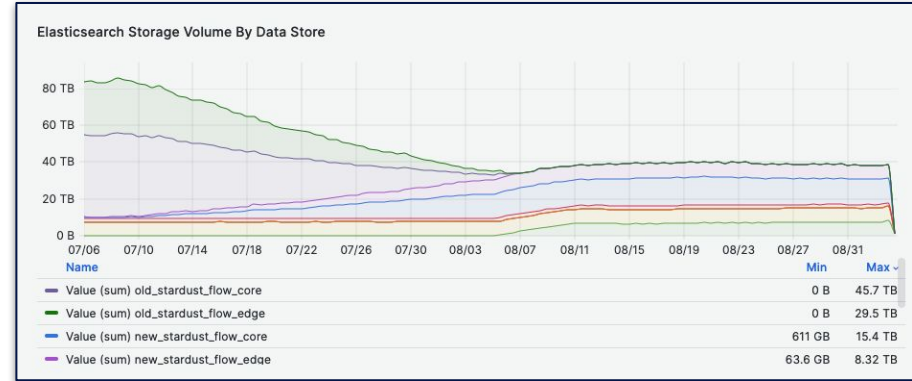
- scale of > 50 nodes
- new features since adoption to improve costs / scaling





Elastic Time Series Data Stream (TSDS)







- <https://www.elastic.co/guide/en/elasticsearch/reference/current/tsds.html>
- As of version 8.9
- Reported savings of up to 70%
 - <https://medium.com/squareshift/up-to-70-metrics-storage-savings-with-tsds-enabled-integrations-in-elastic-observability-4cf8b6217c1>
- We are evaluating this in particular for both Flow and Port metrics with encouraging results
- ESnet has deployed this at scale for a few months
 - Observed **63% reduction** for single packet flows
 - Observed **72% reduction** for multi packet flows





Impact of SNMP timing variance on measurement quality

- Trying to create map of max instantaneous usage we noticed links with impossibly high values
- Exhaustive investigation found:
 - Telegraph SNMP Poller using same timestamp for all values in a getBulk sequence (based on goSNMP)
 - The variance in Elapsed Time is higher than anticipated: $O(4)$ sec
 - **We estimate 30 second rate calculations have a 22% average error rate**
- Implications:
 - With our ET variance it does not make a great deal of sense to poll at 30 seconds
 - **To get < 2% we would need 5 min polling**
 - Moving to streaming telemetry we hope will provide qualitative improvements
 - There are ways to through rewriting pollers to improve SNMP
- Data and Scripts
 - <https://github.com/MetrANOVA/SNMP-ET-Jitter-exploration/tree/main>





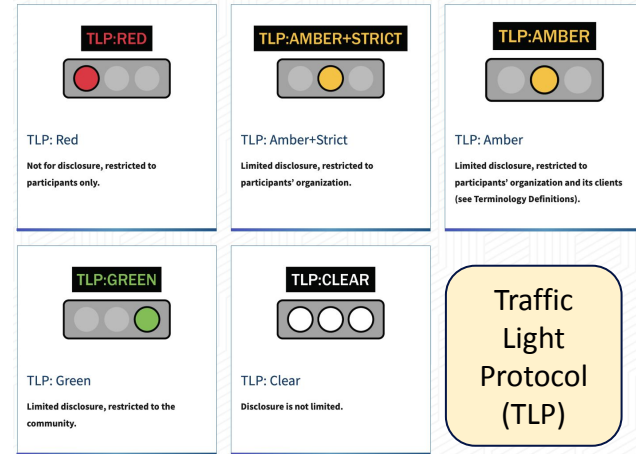
Data Sharing, Federation, Anonymization, Policy

- Federated services rely on data from multiple systems and domains
- Appropriate controls that respect each domains policies and constraints are a must for data sharing
- Having well defined policies is a precursor which today does not always exist
- Example of constraints you are facing:
 - GDPR, FERPA, HIPAA.
 - NDAs and customers who wish to remain low profile
 - Institutional policies, funding bodies, etc.



Different kinds of Sharing

- Different data products have different levels of sensitivity
 - Raw measurements
 - API access to measurement repository with query language
 - Access to online dashboards and reports
- Clear policies helps set expectations
 - What is and what is not shared and at what level of access
- An example:
 - perfSONAR measurement dashboards
 - Performance monitoring data
 - Geopositioning data
 - Display in a single pane of glass.
 - Less policy constraints what is collected and what is shared





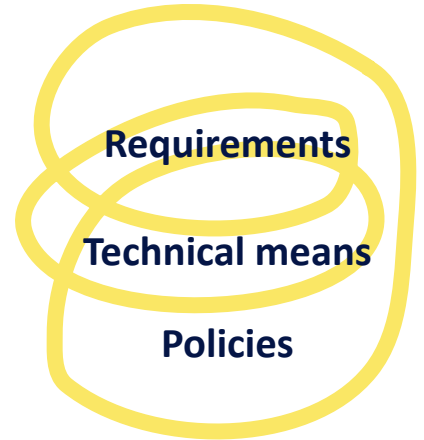
Laws, Rules, Policies, Guidelines, Best Practices

- MetrANOVA is not a substitute for legal advice.
- Just as we do with the technical elements, we aim to look at use cases and provide documents/processes that people can use, or at least start with.
- Covering issues like...
 - How the data is collected and transported to the storage infrastructure?
 - How is it stored, and where? Is it encrypted at rest?
 - Anonymization - how, at what stage in the process?
 - Storage, collection, display?
 - Retention - how long will the data be kept?
 - Sharing - who is the data shared with, on what terms? For what purposes?



There is no “best” policy

- Technical, development, and engineering work (sometimes) have relatively clear-cut ways to define what the “right” solution is
 - Performance metrics, etc.
- There’s no “best” set of policies
 - We’re not trying to set a standard, more give people a starting point
- What we *can* do is document a set of policies that *can* work.
 - Based on the actual knowledge and experience of participants
- Reduce the amount of effort that people need to put into getting started.





Output

- Documenting performed analysis and best practices
 - Technical work e.g. dataspace solutions comparisons
 - Data management approaches and policies, re. Collection, transport, etc. etc.
- Consider specific use cases
 - Rather than generic exploration of data management
 - Focus on Network monitoring and management
- Use as-is or adapt to specific legal/technical/institutional environment
 - No reinventing the wheel
 - From the community to the community



Many Thanks!

Edward Balas

MetrANOVA Consortium Lead

ebalas@es.net

Ivana Golub

Executive Committee rep for PSNC/GÉANT

igolub@man.poznan.pl

For more information:

- Github: <https://github.com/MetrANOVA>
- Web: <http://www.metranova.org/>



Questions?

