MetrANOVA

A Consortium for Advancing Network Observation and Analysis

Introduction and updates

ESnet GÉANT INDIANA UNIVERSITY INTERNET2 TACC

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

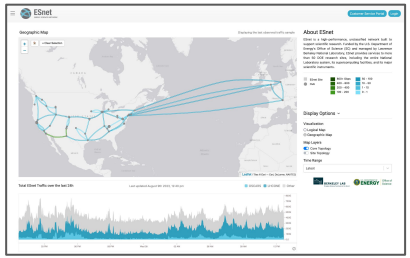

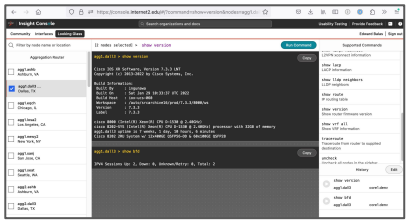

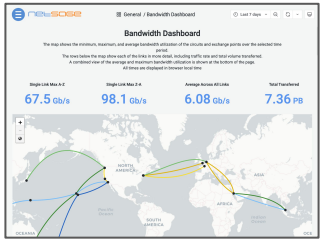

In the future…

1. **Neworks will employ data-driven design and operations**
2. Diverse Measurements will be combined with metadata to provide a more detailed understanding of the network
3. Composite / End to End views of the R&E infrastructure will give stakeholders appropriate awareness

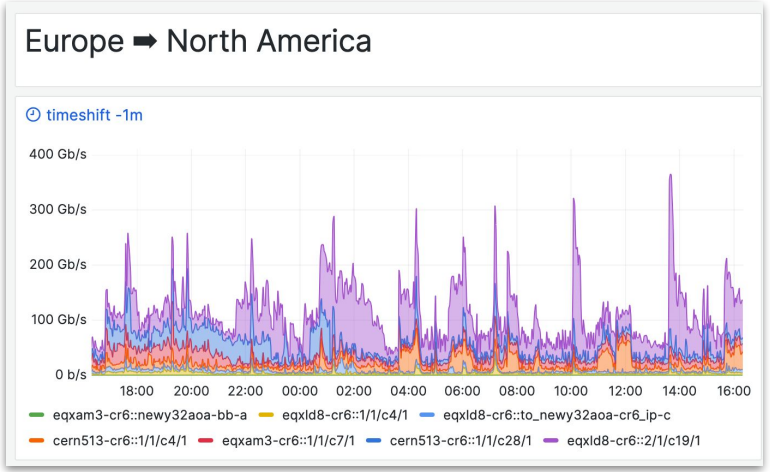

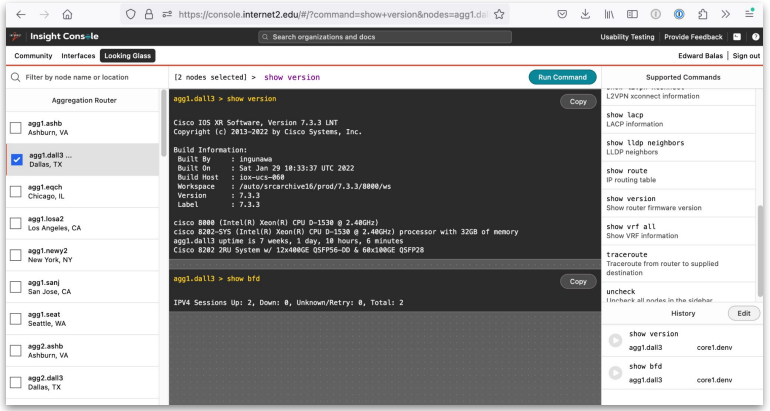

In the future…

1. Neworks will employ data-driven design and operations
2. **Diverse measurements will be combined with metadata to provide a more detailed understanding of the network**
3. Composite / End to End views of the R&E infrastructure will give stakeholders appropriate awareness

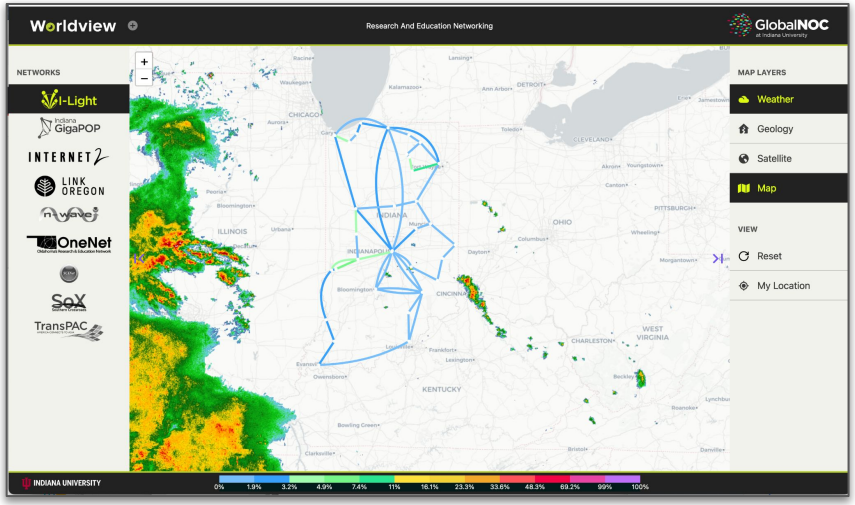

In the future…

1. Neworks will employ data-driven design and operations
2. Diverse Measurements will be combined with metadata to provide a more detailed understanding of the network
3. **Composite / end to end views of the R&E infrastructure will give stakeholders appropriate awareness**
    - Recent examples include:
        - International Networks @ Indiana University, NetSAGE, TACC and EPOC
        - GlobalNOC and Global Research Map
        - PerfSONAR

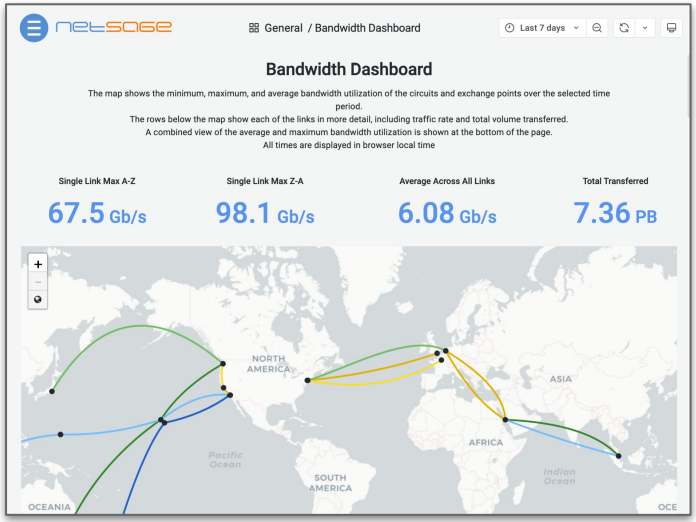

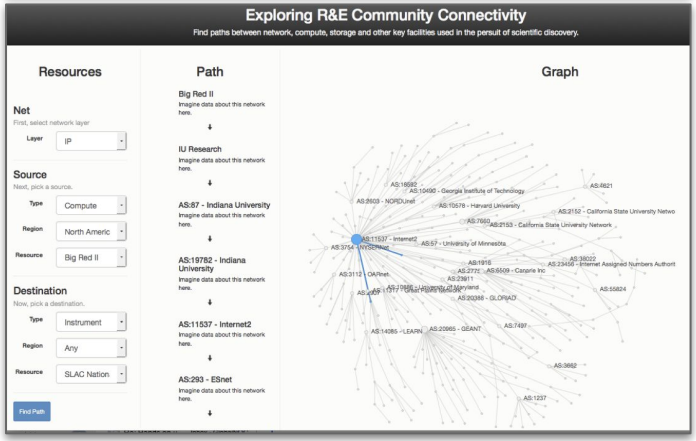

Is this a solved problem?

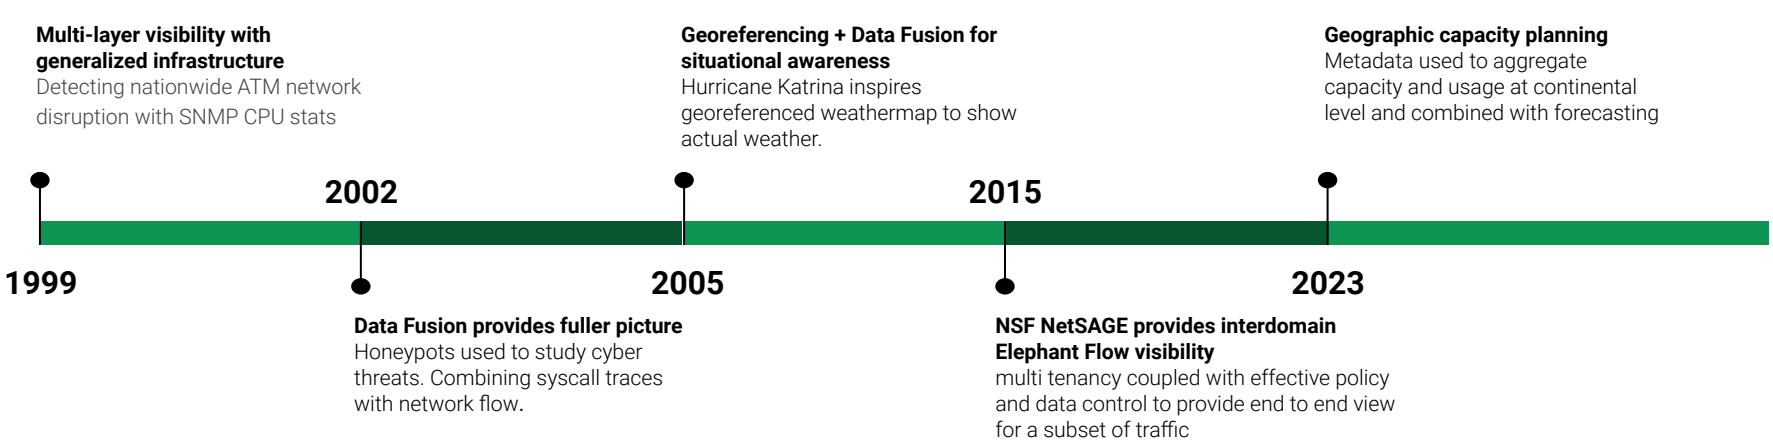

Most of the key ingredients are available however…

- 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 exit 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 exit 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 exit 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

- **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.

- 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

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

ESnet GÉANT INDIANA UNIVERSITY INTERNET2 TACC

Consortium Details

Goals

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

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

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

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?

- Internal member survey
- community survey launched
- 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)

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 adoptions to improve costs / scaling

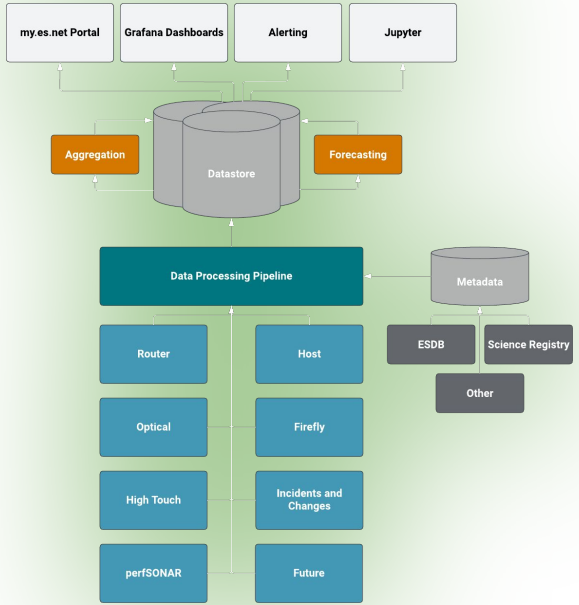

Elastic Time Series Data Stream (TSDS)

- 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 a few weeks ago
- We have observed **~63% data reduction** for single packet flows

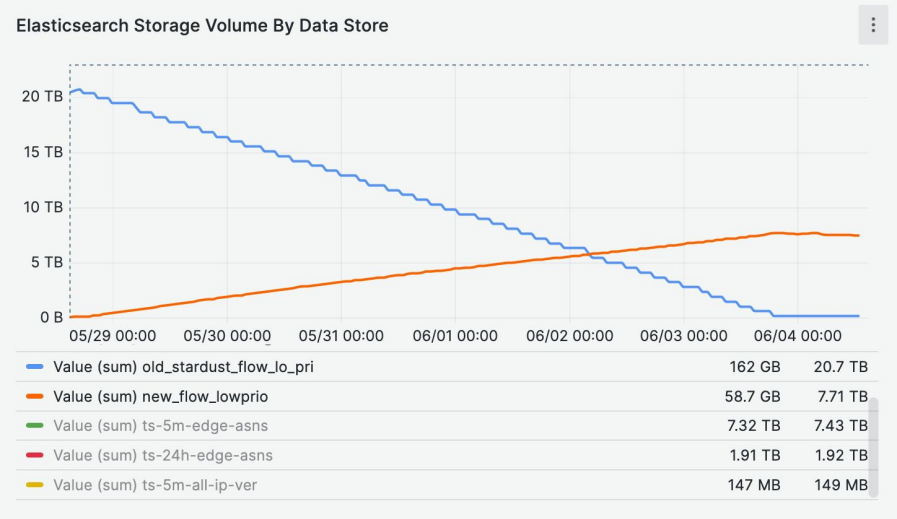

Impact of SNMP timing variance on measurement quality

- 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 that 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-explore/tree/main

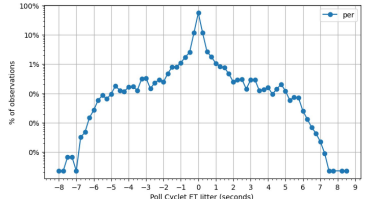

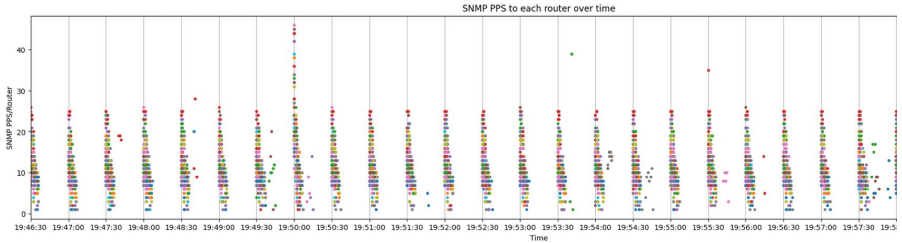

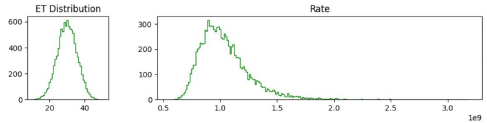

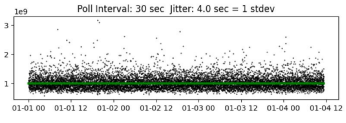

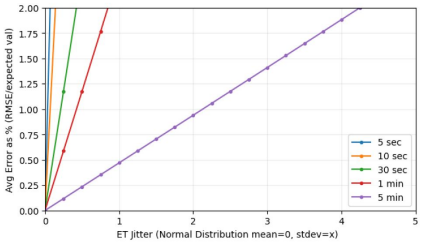

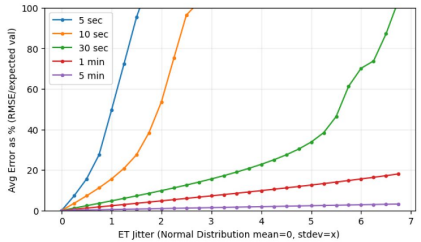

Data Sharing, Federation, Anonymization, Policy

- 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
- We will be developing recommendations for the creation of functional policy
  - With technical means to enforce
- Example of constraints you are facing:
  - GPDR, FERPA, HIPPA.
  - NDAs and customers who which to remain low profile
  - Institutional policies, funding bodies, etc.

Sharing is Caring

- One of MetrANOVA's key goals is enabling sharing of data
- We want the mechanisms to be easily enabled
  - That's the technical part
- We want to design data models, APIs etc that support sharing of all kinds
- *We want defining the supporting policies etc to be easy, or at least easier than it would be if one was starting from scratch*.

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 let consumers and providers helps set expectations
  - What is and what is not shared and at what level of access
- An example:
  - International Networks at IU is an example
  - Supporter of IU's contribution to the consortium.
  - Vital for them to be able to collect/accept data from multiple sources
    - Display in a single pane of glass.
  - 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 give people a set of documents/processes that they 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 not a direct equivalent - 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

- What does the policy side want to produce?
- A set of documents that define policies re. Collection, transport, etc. etc.
- Align with a MetrANOVA-style architecture
  - But doesn't assume it
- Use as-is or adapt to specific legal/technical/institutional environment

How you can help

Share your thoughts to guide us

- MetrANOVA 2024 Community Survey
  - Help influence our strategy and roadmap
  - https://forms.gle/zeYAu8Hp1bZFU8hF7

Do it meow!

Many Thanks!

Edward Balas

MetrANOVA Consortium Lead

ebalas@es.net

David Ripley

Head of Policy Workstream

daripley@iu.edu

Andrew Lake

Head of Technical Workstream

andy@es.net

For more information:

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

Questions?