

Internet2 Performance Working Group: Community Metrics and Telemetry Project Meeting Notes

10-13-2020

Project info:

https://docs.google.com/document/d/1A_6Ooowm33MZR3A-fuTnhGwIUX6UEeaPumRfGUSBEgl/edit#heading=h.b5cbzrmmu6l

Agenda:

- Agenda bash
- Project status updates
 - Respective Pilots
 - Memorandum of Understanding → Letter of Intent to share
- Prototype Tooling
 - Parallel Traceroute Visualization
- Action items from last month
- Next steps and action items for next month
- Upcoming Meetings of interest
- Other

Attendees (Please sign in)

Joe Breen, Dale Finkleson, Jeff Schwab, Michael Lambert, Larry Blunk, Celeste Anderson, Corey Eichelberger, Andrew GoLightly, Matt Zeukaskas, Paul Love, Andrew Gallo, George Uhl, Alan Whinery, Ventsi Gotov, Mark Feit

Apologies
Tim Chown

Project Status updates

Pilots:

Clemson / C-Light

- Initial meeting and pilot scope completed
 - Pilot will pick up C-Light routers and Clemson router path to the Clemson Science DMZ
- Next Steps: Implement pilot at C-Light and Clemson routers

Imperial College / JISC pilot (UK)

- Notes from Tim Chown: I have emailed David (Imperial), Niall (GÉANT) and Paul (Jisc) to remind them of where we were earlier in the summer, with the desired list of data items we'd like to extract from devices on the path, and the sample traces from Utah to Imperial's internal and external DTNs for examples of specific devices we'd like to access.
-
- Is there any new development on the proxy that we might use between the queriers and the back end NREN databases? The good news from our last meeting was that GÉANT, Jisc and Imperial at least all use influxdb, so there are some commonalities. I'm not sure what the status of the proxy is though; I doubt direct access would be possible.
-
- Finally, the GÉANT GN4-3 project (my work package) is hosting a telemetry workshop on 10th Nov, which I think would be the date of the next I2 perf WG meeting, but they won't clash due to the TZ difference. There's still room in the agenda to fit in a 10-15 min talk about this pilot, the rationale behind it, the issues it throws up, and the desire to share telemetry data. There's some info at <https://wiki.geant.org/display/PUB/Telemetry+and+Big+Data+Workshop> and registration open at <https://events.geant.org/event/207/>. Please let me know if you're interested to present. We have a couple of US presenters already on P4/INT and NetSage.
 - JB: Karl N pointed out the GEANT workshop yesterday and I have added it to our notes for today. Karl mentioned he will already be there and would be willing to present so we are planning to put some material together.

LinkOregon pilot

- On hold through end of year due to CARES act projects

NCSA

- Has been on hold due to project load - trying to re-engage now
- Next steps: Additional follow up and scoping of initial pilot

- Corey hoping to setup something over the next couple of weeks

University of Hawaii pilot

- UH and Astronomy have been on hold due to project load - trying to re-engage now
- Next steps: Scoping the initial pilot

University of Michigan prototype with container

- Nick Grundler has been able to implement the project telegraf container so that it is polling core UMich core interfaces
- Next steps: Nick is coordinating with Shawn McKee to add router interfaces for his infrastructure and for pertinent Data Transfer nodes

University of Utah / Utah Education and Telehealth network

- Successfully completed initial prototype with direct SNMP polling
- Have started developing tools with early data set results
- Telemetry data collectors purchased through CARES act funding in order to create local repositories of data
 - UETN has over 1600 routers
 - UofU Science DMZ infrastructure and Testbed infrastructure will contribute to data
- UoU rolling out 6 new 100G DTNs by November in addition to the 6 existing 40G DTNs
 - UofU is instrumenting all DTNs will have perfSONAR, SNMP, and some PCP instrumentation
- Sloan Digital Sky Survey rolling out DTNs
 - Asking if they would be willing to participate in the project
 - Ask them on Science Registry
- UofU migrating to the project Telegraf container to pick up the data and send it to I2
- UETN hiring new person focused on measurement and automation who will help on UETN/UofU pilot as a portion of duties
-

ESNet

- Andy Lake's team has modified NetBeam so tools can query paths via IPv4 addresses
 - This modification allows for traceroute tools and other tools to obtain information across the ESnet backbone using the IPv4 addresses in addition to interfaces

- This work has allowed prototype tools to do interdomain visualization from university, across regional, across I2 to router in ESnet
- IPv6 query capability will be a future endeavor

Letter of Intent to Share

- Rough draft:
<https://docs.google.com/document/d/1Z8OTssi5K9cQG118KYSr1hA0UKPZ4jqAKtvIUrbLglQ/edit#heading=h.5c13see8bx63>
- Suggestions, feedback, and willingness to help welcome
- Intention to do something deliberately -- Not a contract or legal document
- Receiving review from Quilt members / CEO Roundtable -- James Deaton is following up
- Q: other institutions to give feedback A: Have received feedback from all but one or 2 institutions
 - << >> one more pass and then go for signing
- Other example documents:
 - Measurement Lab Acceptable Use Policy: <https://www.measurementlab.net/aup/>
 - RIPE NCC Terms and Conditions for Data Access: <https://labs.ripe.net/datarepository/conditions/basic>

Tools

- Internet2 and GRNOC Telegraf container for collecting data:
 - Latest: <https://github.internet2.edu/internet2/grnoc-telegraf-docker>
 - Work happening to port to Kubernetes on the SLATE (<http://slateci.io>) platform
 - Work happening to expand the configuration to use a new version of these
- IU GRNOC Augmented Traceroute:
<https://snapp-portal.grnoc.iu.edu/grafana/dashboard/script/atr.js?orgId=2>
- Parallel Traceroute Visualizer status

- Successful initial integration ESnet NetBeam API data using separate scripts:
- Ability to now support a traceroute from a university across a regional, across the Internet2 backbone and across the ESnet routed backbone
 - For the I2 portion, the tool is currently leveraging embedded Grafana visualization from Augmented Traceroute
 - For the NetBeam portion, some work is happening to visualize -- for now, simple web tables with the data are in place
 -
- Tool description and architecture:
 - <https://docs.google.com/document/d/11TJL7QYcu5roEwls92CcHVc856mrSfDOIiMJxulrSHg/edit?usp=sharing>
- NetBeam API:
 - <https://documenter.getpostman.com/view/4148513/SzKZrvTv?version=latest#intro>
- Internet2 GRNOC database and TimeSeries Data Base

Q: are we trying to pull from individual GRNOC and NetBeam or some sort of standard A: pulling from 2 different data sources;

Observation:

UH: DTNs are running in a maddash - running scheduled tests -- "pscheduler pause" when transfers are happening

Q: how to ensure capacity on network? AW: focus on getting as much data as can vs getting within a scheduled time window

Internet2 Performance Working Group wiki

- <https://spaces.at.internet2.edu/display/PerformanceWG/Home>
- Some progress continuing on revamping site

Review of Action Items from last month:

- <<>> Finish UMich pilot within its scope and finish testing of container -- **IN PROCESS**

Next Steps and Action Items for next month:

- <<>> Follow up on pilots - specifically NCSA, JISC/Imperial, Utah, UMich
-

Other

- DoE SBIR grant - Technologies for Sharing Network Performance Data
 - https://science.osti.gov/-/media/sbir/pdf/TechnicalTopics/FY2021_Phase_I_Release_1_Topics.pdf?la=en&hash=D7C553C3316BF15F3A73AEA0CAED21A161692F2B
 - Anonymization Tools and Services
 - Correlate Log Data and/or Host Sensor data with Network Trace Data
 - Questions –Contact: Richard Carlson, Richard.Carlson<at>Science.doe.gov
 - **References:**
 - 1. Dijkhuizen, N. V., Der Ham, J. V. (2018). A Survey of Network Traffic Anonymisation Techniques and Implementations, Article No.: 52. <https://dl.acm.org/doi/10.1145/3182660>
 - 2. Liu, Z., Kettimuthu, R., Leyffer, S., Palkar, P., and Foster, I. (2018). A Mathematical Programming-and Simulation-Based Framework to Evaluate Cyberinfrastructure Design Choices, IEEE 13th International Conference on e-Science (e-Science), Auckland, 2017, pp. 148-157. <https://ieeexplore.ieee.org/document/8109132>
 - 3. Chung, J., Liu,Z., Kettimuthu, R.,andFoster, I. (2019). Elastic Data Transfer Infrastructure for a Dynamic Science DMZ, 15th IEEE International Conference on eScience (eScience 2019), San Diego, CA, Sep. 2019<https://www.mcs.anl.gov/~kettimut/publications/elasticDTN.pdf>
 - 4. Liu, Z., Kettimuthu, R., Balaprakash, P., Rao, N. S., Foster, I. (2018). Building a Wide-area File Transfer Performance Predictor: An Empirical Study, International Conference on Machine Learning for Networking, (MLN2018), November 27-29, Paris, France. <https://www.ornl.gov/publication/building-wide-area-file-transfer-performance-predictor-empirical-study>

Upcoming Meetings of Interest

- perfSONAR TechExtra day -- Nov 2

- <https://www.internet2.edu/news-events/events/techextra-perfsonar-day/>
- Please sign up -- registration closes end of month
 - No charge
 - Invite those who are playing with perfSONAR and not able to usually attend due to fees
 - If questions, please drop a note
- Update on perfSONAR -- hoping for 4.3 announcement
- Measurement in cloud
- GEANT talks
- Measuring wireless quality
- Other topics
-
- GEANT - Nov 10
 - <https://wiki.geant.org/display/PUB/Telemetry+and+Big+Data+Workshop>
 - What data do researchers want?
 - What operational data needed to expose?
 - How to best present the data?
 - Other questions
- perfClub next Monday (3rd Mon of the month)

OSHEAN

- Ventsi - would like to start a pilot from OSHEAN's perspective -- will reach out

#####