Monitoring Mini-Workshop

Sarah Edwards
July 27, 2011
www.geni.net

Sponsored by the National Science Foundation

Agenda

- Introduction [10 min]
  - Sarah Edwards, BBN/GPO
- Point of View Talks [10 min each]
  - Lessons learned from monitoring Plastic Slices
    - Chaos Golubitsky, BBN/GPO
  - Sharing Data via GMOC
    - Camilo Viecco, Indiana University/GMOC
  - OpenFlow Monitoring
    - Nick Bastin, Stanford/OpenFlow
  - Campus and Experimenter Needs and Resources
    - Sarah Edwards, BBN/GPO
- Open Discussion - All [30 min]
- Wrap Up/Conclusions [10 min]

Monitoring: Introduction

Sarah Edwards – sedwards@bbn.com

July 27, 2011

www.geni.net

Sponsored by the National Science Foundation

GEC 11 Monitoring Related Meetings

- Yesterday in the Campus track
  - OpenFlow Campus Deployment
  - Operations Update
  - Plastic Slices Report Out
- Later today **Monitoring BOF Dinner**
  - Meet in the hotel lobby at 6:30pm
- Tomorrow in the Software track
  - I&M Working Session from 1pm-3pm

Purpose & Outcome

- Operations are supporting aggregates like they are production
  - Need monitoring now
- Purpose of the Mini-Workshop is to:
  - Share lessons learned, available tools, needs
- The Outcome is:
  - a list of important monitoring topics to inform priorities in the next few months
- Full Description:
  - http://groups.geni.net/geni/wiki/GEC11MonitoringMiniWorkshop

Scope of today's discussion

- Focus on areas of shared interest
  - Tools to aid people with responsibilities
    - e.g. campuses who have signed the aggregate agreement
  - Lower burden on monitoring done multiple places
    - e.g. monitoring MyPLC happens at lots of campuses
  - Coordination when we rely on each other
    - e.g. debugging a VLAN path across multiple networks
  - Monitoring that is truly GENI-wide
    - e.g. GMOC DB is currently acting as a shared repo of data

Questions for today's discussion

- What resources do campuses and experimenters need monitored?
- What resources are currently being monitored?
- What monitoring data would be easy and useful to share?
- What tools are missing?
- What agreements do we need? APIs? data formats? etc.?
- Are there concerns about sharing data?
- **Anything else?**

Agenda

- Introduction [10 min]
  - Sarah Edwards, BBN/GPO
- Point of View Talks [10 min each]
  - Lessons learned from monitoring Plastic Slices
    - Chaos Golubitsky, BBN/GPO
  - Sharing Data via GMOC
    - Camilo Viecco, Indiana University/GMOC
  - OpenFlow Monitoring
    - Nick Bastin, Stanford/OpenFlow
  - Campus and Experimenter Needs and Resources
    - Sarah Edwards, BBN/GPO
- Open Discussion - All [30 min]
- Wrap Up/Conclusions [10 min]

geni
Exploring Networks of the Future

GENI Monitoring

Lessons learned from Plastic Slices

Chaos Golubitsky, GPO

GEC11 July 27, 2011

www.geni.net

Sponsored by the National Science Foundation

Plastic Slices Monitoring Overview

- Features of Plastic Slices project:
  - A few types of aggregates (OpenFlow, MyPLC)
  - Many instances of aggregates (nine campuses)
- Goals of monitoring:
  - Diverse sites submit data to a central location
    - GMOC, via data submission API
  - Collected data is available
    - GMOC web UI: http://gmoc-db.grnoc.iu.edu/api-demo/
    - Recent data can be downloaded by interested parties
  - Collected data is useful
    - Diagnostics at remote aggregates
    - Resource status data for experimenters to use

Plastic Slices Monitoring Status (1)

- Diverse campuses submitting data:
  - Eight campuses and two backbones submitted data to GMOC's staging database via the API
  - Now migrating to (authenticated) production database
- Data availability:
  - Recent data (past 10 minutes) is publically available: http://gmoc-db.grnoc.iu.edu/web-services/gen_api.pl
    - To see older data, use an existing long-term collection (e.g. at GMOC or GPO), or download recent data to make your own
  - GPO downloaded recent data every five minutes, starting early May, with roughly 96% success rate
  - Repackaged data into local graphs and Nagios alerts

Plastic Slices Monitoring Status (2)

- How data was used during plastic slices:
  - Operator debugging: monitoring failure may imply aggregate is down
  - Experimenter debugging: has my OpenFlow sliver been opted in?
  - Experimenter results: MyPLC per-sliver traffic counters

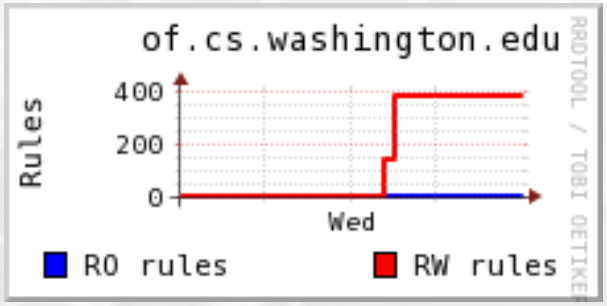

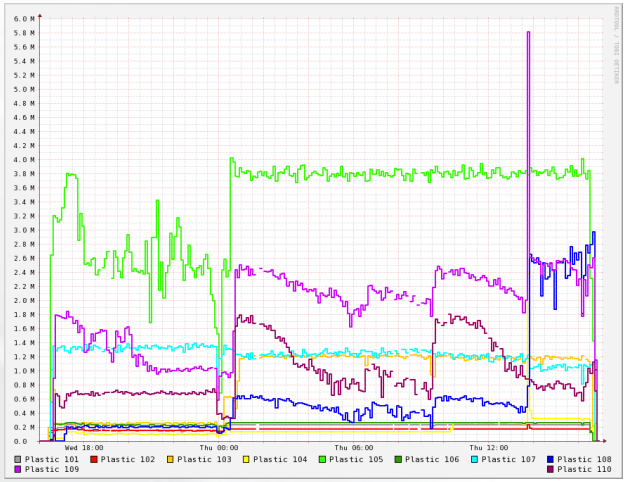

Lessons Learned and Next Steps

- Lessons:
  - Clock synchronization is important for time-series data with self-reported timestamps
  - Collecting, downloading, and processing data is a high-performance application (needs sufficient hardware)
  - There's never enough data to debug all the problems
- Next steps:
  - Finish move to production, work on performance issues and code robustness
  - Improve flowvisor message enumeration
  - Improve slice-to-sliver mapping tools
  - Help other data submitters use the API

Other mesoscale network monitoring at GPO

- Public real-time monitoring at GPO to help detect network problems in the OpenFlow backbone:
  - Can dedicated monitoring nodes ping each other? http://monitor.gpolab.bbn.com/connectivity/core.html
  - Do we see traffic leaks or broadcast storms in the core topology?

    http://groups.geni.net/geni/wiki/NetworkCore/TrafficLeaks
- Data sharing helped operators work together
- Advertise your site's public monitoring data on your aggregate pages

Agenda

- Introduction [10 min]
  - Sarah Edwards, BBN/GPO
- Point of View Talks [10 min each]
  - Lessons learned from monitoring Plastic Slices
    - Chaos Golubitsky, BBN/GPO
  - Sharing Data via GMOC
    - Camilo Viecco, Indiana University/GMOC
  - OpenFlow Monitoring
    - Nick Bastin, Stanford/OpenFlow
  - Campus and Experimenter Needs and Resources
    - Sarah Edwards, BBN/GPO
- Open Discussion - All [30 min]
- Wrap Up/Conclusions [10 min]

Sharing Data with GMOC

What can we do for your measurements

Camilo Viecco

GlobalNOC
Global Research Network Operations Center

gMOC
GENI META-OPERATIONS CENTER

GMOC Overview

- GENI Meta-Operations
  - Provide unified view of GENI
    - (url: http://gmoc-db.grnoc.iu.edu/)
  - Provide a initial point of contact for GENI related operations
  - **Provide Monitoring visualizations and API**
    - (url: http://gmoc-db.grnoc.iu.edu/api-demo/)
  - Provide emergency shutdown services

GlobalNOC
Global Research Network Operations Center

gMOC
GENI META-OPERATIONS CENTER

GMOC Monitoring

- 'The place for unified data'
  - In particular infrastructure data
- We do two types of data monitoring: measurement and alerting.
  - Only measurement is considered Production ready.
  - Alerting (nagios)
    - Running but alerts only are local
    - Url: http://gmoc-db.grnoc.iu.edu/nagios/
    - Challenge is to make sure notifications are both targeted (no false negatives/positives) and actionable.

GlobalNOC
Global Research Network Operations Center

GMOC
GENI META-OPERATIONS CENTER

Measurement

- Long time trend Analysis
  - Status of hosts/interfaces/slices
  - Time series based
- Tools (SNAPP,Ganglia,Measurement Manager)
- We collect data directly and provide API for submission and recollection of Data.

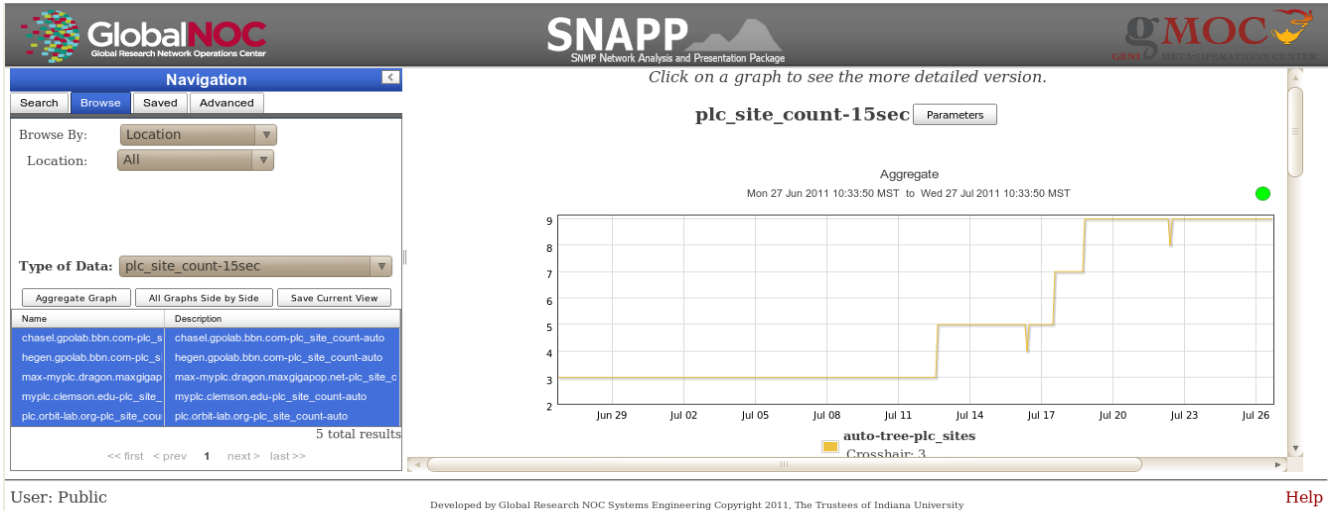

GlobalNOC
Global Research Network Operations Center

gMOC
GENI META-OPERATIONS CENTER

SNAPP

- A Suite for SNMP collection and visualization
  - Snapp-collector
    - Very scalable, up to 130000 interfaces monitored in a single host
    - Used to monitor Internet2,NLR, IU,ProtoGENI BB switches
  - Snapp frontend
    - Reused the SNAPP front-end to display the time series data.
    - Javascript, csv and png output
    - Aggregation
    - Data tagging

Data API

- Simple format for time-series data
  - Simple XML file
    - You can write your own generator or use tools developed by GMOC for RRD files.
  - Can be used for both input and output.
- Data is sent via a post to the gmoc servers
  - Production server requires auth.

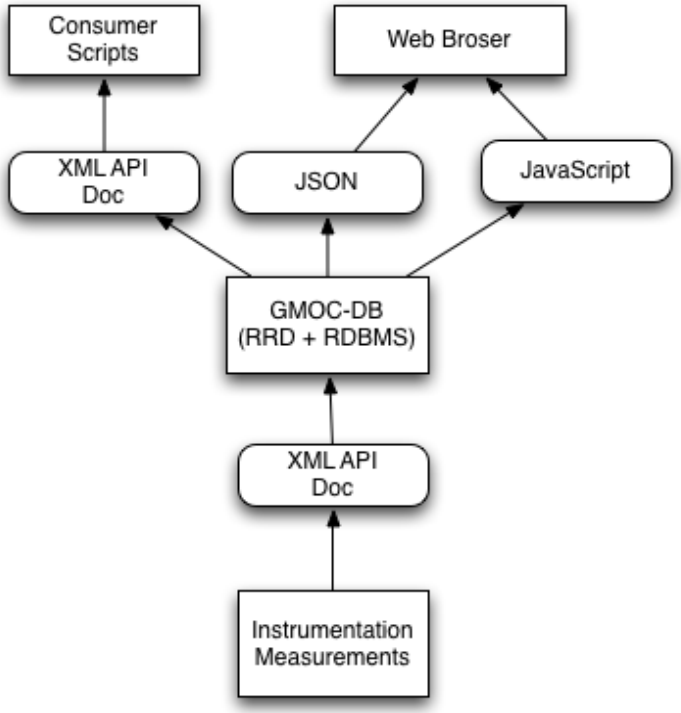

GlobalNOC
Global Research Network Operations Center

gMOC
GENI META-OPERATIONS CENTER

Want to be reported at GMOC?

- Give us direct SNMP access
  - Used by Internet2, NLR, ProtoGENI (and soon CRON)
- Send us data via the Measurement API
  - Currently only used by the Plastic Slices effort (INSTOOLS coming next).
  - Requires shared secret for production

Agenda

- Introduction [10 min]
  - Sarah Edwards, BBN/GPO
- Point of View Talks [10 min each]
  - Lessons learned from monitoring Plastic Slices
    - Chaos Golubitsky, BBN/GPO
  - Sharing Data via GMOC
    - Camilo Viecco, Indiana University/GMOC
  - OpenFlow Monitoring
    - Nick Bastin, Stanford/OpenFlow
  - Campus and Experimenter Needs and Resources
    - Sarah Edwards, BBN/GPO
- Open Discussion - All [30 min]
- Wrap Up/Conclusions [10 min]

OpenFlow Network Monitoring

- SNMP
  - Existing MIBs already have extensive layer 2/3 protocol coverage
  - Existing monitoring SW already speaks SNMP
  - Traps
- Custom
  - Hooks into data that SNMP doesn't handle well
  - "Quick-and-Dirty" way to expose data before exposing to SNMP

OpenFlow Component Stack

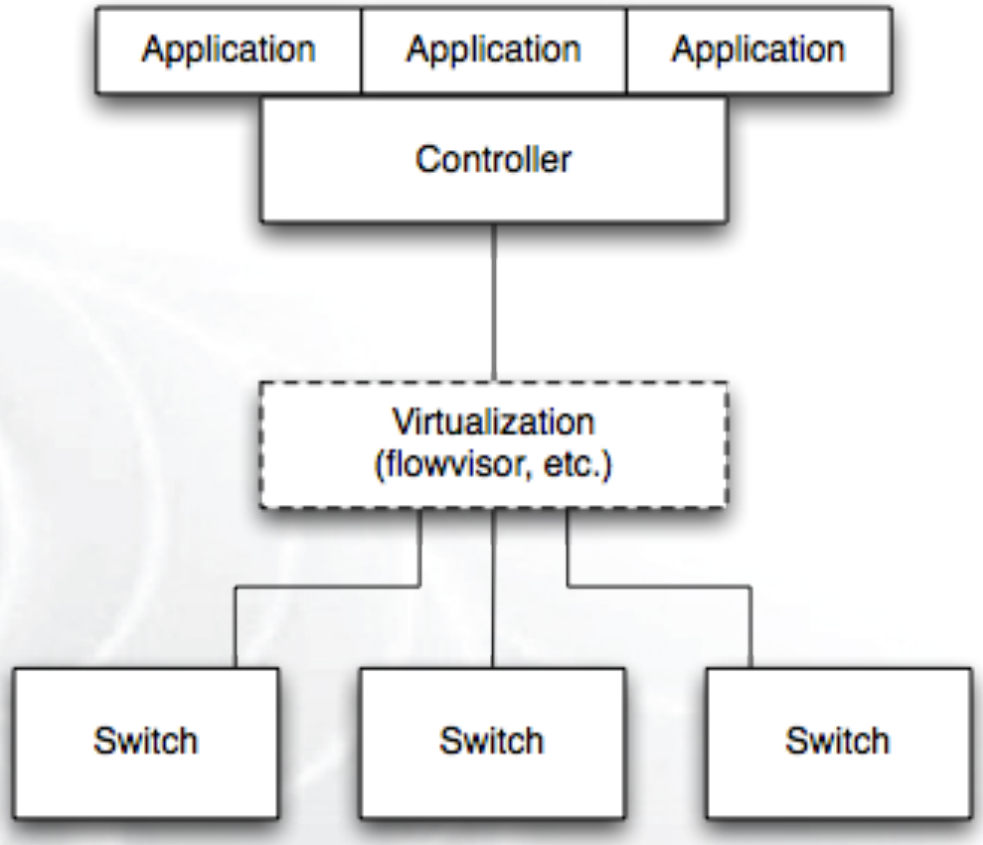

How you can help

- Researchers
  - Assume people will use your software after you publish your paper
  - Build in monitoring hooks from the start (you'll be glad you did)
- Network Operators
  - Insist on vendor support for SNMP (sw and hw)
  - Vote with your dollars
- Protocol Designers
  - Don't try to re-invent SNMP

Agenda

- Introduction [10 min]
  - Sarah Edwards, BBN/GPO
- Point of View Talks [10 min each]
  - Lessons learned from monitoring Plastic Slices
    - Chaos Golubitsky, BBN/GPO
  - Sharing Data via GMOC
    - Camilo Viecco, Indiana University/GMOC
  - OpenFlow Monitoring
    - Nick Bastin, Stanford/OpenFlow
  - Campus and Experimenter Needs and Resources
    - Sarah Edwards, BBN/GPO
- Open Discussion - All [30 min]
- Wrap Up/Conclusions [10 min]

Campus and Experimenter Needs and Resources

Sarah Edwards – sedwards@bbn.com

July 27, 2011

www.geni.net

Sponsored by the National Science Foundation

- Spoke with a small number of folks
  - Some campuses
  - Some aggregates
  - Some experimenter-representatives
- The following are a sampling of possibly interesting things they said

GEC 10 Monitoring BOF Breakfast

- Tool: “slice top”
- Want to be able to say: “this is not my fault”
- State future date when instantiate policy regarding privacy

Russ Clark – Georgia Tech.

- OpenFlow
  - OpenFlow switch manufacturers will eventually need to provide per-flow stats in the switch
  - Interested in OpenFlow monitoring for experiments
  - Dynamically setting up mirrored ports on OpenFlow switches has great potential for doing security analysis and for IDS
- General Monitoring
  - Interested in sharing live (or recent) data without giving direct SNMP access to switches
- Broader Impact
  - As always, need data to argue for continued support from administration
    - e.g. How many nationwide experiments did we support?
    - e.g. How many papers published?

Chaos Golubitsky – GPO Lab

- Currently, doing the following among others:
  - Standard host monitoring; E2E network reachability
  - For GENI Core: alerting on stray packets on wrong VLAN
  - Sharing everything at: http://monitor.gpolab.bbn.com
  - Could tailor pointers to bits of interest to people.
- Questions
  - Useful network stats; traffic stats
  - Is the integrity of the topology good? Packet storm?
  - Currently building up graphs involves guessing on sliver names (especially for OpenFlow)
- Advice
  - For GENI trending is more important than alerting. Because anyone can look at it even if they weren't interested yesterday. Helps interested parties not just operators.
  - Performance matters. Monitoring server is heavily loaded.

Chris Small & John Meylor – Indiana U.

- Currently, doing the following among others:
  - Measurement Manager: monitors OpenFlow including topology
  - Capture OpenFlow stats, send to GMOC DB
- Needs/Suggestions
  - View of how GENI resources are integrated together
  - Would like a per campus/per aggregate view
  - Topology info
  - Would be nice to get monitoring for free when install nodes
  - OpenFlow Rspecs contain spaces for descriptions which aren't used

Experimenter Needs

- Spoke with: Mark Berman & Niky Riga
- What resources and connectivity exist and are available
  - Existing aggregates. Available resources in aggregate. Connectivity.
  - Currently 4 types of connectivity: commodity Internet, NLR/I2, PG nodes in I2 pops, TangoGENI via OpenFlow. Becoming hard to keep in our heads.
  - OpenFlow topology
- Troubleshooting aids
  - Map configuration to monitoring data (a.k.a. per slice view)
    - port 47 traffic to port 48 only matters if know my configuration includes those ports
    - Possible tool: Input manifest Rspec, and return the 7 graphs that matter
- GENI utilization statistics
- Standard Monitoring Information
  - Usual host and networks stats (cpu, memory, up/down).
  - Aggregate-specific stats.
  - Storage and visualization of the above.
  - Specify granularity of: stats collection frequency, how long to store.

ProtoGENI aggregate monitoring

- NEW! Flack & INSTOOLS now integrated
- PG considers nodes to be the experimenter's responsibility
- Idle monitor to tell when reserved nodes are NOT in use
- Emulab (not PG) nodes have a serial console

PlanetLab aggregate monitoring

- PlanetLab has three main monitoring tools
  - CoMon
    - Per node statistics available to all users
    - Info at: http://comon.cs.princeton.edu/
  - PlanetFlow
    - Logs flows in and out of a node
    - Provides a publicly accessible database to identify sources of IP packets
    - Running at: http://planetflow.planet-lab.org/#planetlab
  - MyOps
    - Automates bringing nodes up
    - Prods maintainers when attention is needed
    - Running at: http://monitor.planet-lab.org/monitor/

Agenda

- Introduction [10 min]
  - Sarah Edwards, BBN/GPO
- Point of View Talks [10 min each]
  - Lessons learned from monitoring Plastic Slices
    - Chaos Golubitsky, BBN/GPO
  - Sharing Data via GMOC
    - Camilo Viecco, Indiana University/GMOC
  - OpenFlow Monitoring
    - Nick Bastin, Stanford/OpenFlow
  - Campus and Experimenter Needs and Resources
    - Sarah Edwards, BBN/GPO
- Open Discussion - All [30 min]
- Wrap Up/Conclusions [10 min]

Questions for today's discussion

- What resources do campuses and experimenters need monitored?
- What resources are currently being monitored?
- What monitoring data would be easy and useful to share?
- What tools are missing?
- What agreements do we need? APIs? data formats? etc.?
- Are there concerns about sharing data?
- **Anything else?**

Possible Monitoring Topics to Address

How to participate

- If you are interested in working these topics:
  - e-mail Sarah Edwards (sedwards@bbn.com)
- Come to the **Monitoring BOF Dinner**
  - Let Sarah know if you are coming
  - Meet in the hotel lobby at 6:30pm