2nd GENI Instrumentation and Measurement Workshop
Chicago, IL, June 8-9, 2010

Contents

1. Technical Goals

The first GENI measurement workshop was held on June 26, 2009[1]. It brought together measurement experts to review the following topics: 1) measurement architecture, 2) instrumentation, 3) experiment specification, and 4) data management. The speakers suggested approaches to each topic that would allow GENI to meet its goals. The consensus was that the design of an effective GENI measurement architecture had just begun.

To continue the effort, six GENI Instrumentation and Measurement (I&M) prototyping projects were established following Solicitation 2, joining three I&M prototyping projects continuing from Solicitation 1. Also, a GENI I&M Working Group (WG) was formed at the beginning of Spiral 2.

The WG has affirmed that GENI I&M systems will provide broad data gathering, analysis and archival capability, sufficient for GENI's research mission and sufficient for operations. Furthermore, the GENI I&M WG will create and document a GENI I&M architecture in Spiral 2, and coordinate the design and deployment of a first GENI I&M system in Spiral 3.

The objective of this 2nd GENI I&M workshop is to gather contributors from key I&M prototyping projects to define priority pieces of the I&M architecture by consensus, assemble teams to complete the documentation, and draft a roadmap for implementations during Spirals 2 and 3.

A GENI I&M Capabilities Catalog[2] has been drafted, which reviews each current GENI I&M project, and other selected projects, and lists: architecture components addressed or implemented; implementations in GENI or elsewhere; and uses in GENI or elsewhere. This document identifies projects that can best contribute to architecture topics and to identify projects that can implement enhanced GENI I&M capabilities in Spirals 2 and 3.

Based upon this catalog, several key projects were identified (see invitee list below) that could contribute to a GENI I&M architecture because they have already implemented pieces of I&M functionality in a manner consistent with GENI goals. Four of these projects were invited to give presentations at the GEC7 WG meeting, highlighting how their work mapped into the evolving GENI I&M architecture. Since then, the organizing committee has gathered technical references from and had

[1] See http://groups.geni.net/geni/wiki/GENIMeasWS

[2] See http://groups.geni.net/geni/wiki/GENIIandMCAPCAT

extended discussions with these projects, and gained a better understanding of how they can best contribute to the GENI I&M architecture.

An early draft of a GENI I&M Architecture document was completed[3] and reviewed at the GEC7 WG meeting. Although there was general agreement on the draft of the architecture, the following priority topics were identified as needing to be defined first:
GENI I&M use cases
GENI measurement plane
GENI I&M services
Interfaces, protocols and schema for measurement data in GENI

This workshop will:

- Gather contributors from the key projects (see invitee list below).
- For each priority topic, the organizers will outline a suggested approach or solution, including how certain key projects might contribute functionality.
- Then, a representative from these key projects will review how they could best contribute the suggested functionality.
- Finally, each priority topic will be discussed in a structured manner, with the goal of achieving a consensus on a proposed solution or approach, plus identifying gaps that need further work.
- Assemble teams for each priority topic, identify the action items required to close any identified gaps, complete the proposed solution or approach, and write a revised section(s) for the architecture document
- Draft a roadmap for implementations in Spiral 2 and 3 by the key projects to realize the proposed solutions.

The revised architecture document will then be reviewed by the WG. It and the roadmap will be used for guiding future work on GENI I&M systems.

[3] See http://groups.geni.net/geni/wiki/GeniInstrumentationandMeasurementsArchitecture

2. Organization

Dates: Tuesday, June 8, 2010, 1:00 pm – Wednesday, June 9, 2010, 2:00 pm

Location: Hilton Chicago O'Hare Airport, Chicago, IL,
http://www1.hilton.com/en_US/hi/hotel/CHIOHHH-Hilton-Chicago-O-Hare-Airport-Illinois/index.do

By Monday, May 24, please contact the hotel at 773-686-8000 , and book a room using the code "BBN", to receive a discounted rate.

Number of attendees: 19

Agenda for June 8:

1:00 pm	Welcome and introductions
1:15 pm	Suggest a basic set of GENI I&M use cases, and review contributions from key projects
1:45 pm	Discuss basic set of GENI I&M use cases, summarize consensus and identify gaps
2:30 pm	Break
2:45 pm	Suggest definition of GENI I&M measurement plane, services, interfaces and protocols (APIs), and review possible contributions from key projects
4:15 pm	Break
4:30 pm	Discuss GENI I&M measurement plane, services, interfaces and protocols (APIs), summarize consensus and identify gaps
6:00 pm	Adjourn
7:00 pm	Dinner

Agenda for June 9:

8:00 am	Suggest contents and structure of GENI measurement data schema, and review possible contributions from key projects

9:30 am	Break
9:45 am	Discuss contents and structure of GENI measurement data schema, summarize consensus and identify gaps
11:15 am	Break
11:30 am	Identify teams for each priority topic, draft action items to close identified gaps, and make writing assignments for revised sections of the architecture document
12:30 pm	Lunch
1:00 pm	Review consensus of GENI I&M use cases; GENI I&M measurement plane, services, interfaces and protocols (APIs); and contents and structure of GENI measurement data schema; and draft roadmap for how key projects could implement them in Spirals 2 and 3
2:00 pm	Adjourn

Participation: Attendance will be limited to invitees. The capacity of the current room is 16+ attendees. Currently, we expect at least 18 attendees to be present, and we are looking to accommodate even a few more if at all possible,

It is important that each attendee come prepared with possible contributions from their project for certain priority topics, as requested by the organizers (coming soon), and be willing to help write revised sections for the architecture document.

Sponsorship: Funding for the workshop will be provided by the NSF through the GPO. This will cover all expenses associated with the workshop itself, including travel expenses for all participants.

Within 30 days upon returning from your trip, please submit a brief invoice, along with receipts for all travel expenses incurred (*including meals*), to BBN Technologies. (See the attached instructions)

Please email your invoice and receipts as a single PDF file to: krich@bbn.com

Organizing Committee:

Paul Barford - University of Wisconsin – Madison (no)
Bruce Maggs – Duke University and Akamai (yes)
Harry Mussman – BBN/GPO (yes)

Vic Thomas - BBN/GPO (yes)
Evan Zhang – BBN/GPO (yes)

Invitee List:
OML (ORBIT Measurement Library) OMF (ORBIT Management Framework)
Max Ott – NICTA (yes, by phone)
Ivan Seskar – Rutgers WINLAB (yes)
Instrumentation Tools
Jim Griffioen - Univ Kentucky (yes)
perfSONAR
Matt Zekauskas - Internet2 (no)
Jason Zurawski – Internet2 (yes)
Martin Swany - Univ Delaware (yes)
Guilherme Fernandes – Univ Delaware (yes)
Ezra Kissel – Univ Delaware (yes)
Scalable Sensing Service (S3)
Sonia Fahmy – Purdue (yes)
Puneet Sharma - HP Labs (yes)
OnTimeMeasure for network measurements
Prasad Calyam - Ohio Supercomputing Ctr (yes)
GENI Meta-Operations Center and NetKArma
Jon-Paul Herron - Indiana Univ (no)
Camilo Viecco - Indiana Univ (yes)
Chris Small - Indiana Univ (yes)
Virtual Machine Introspection (VMI)
Brian Hay – Univ Alaska (yes)
Data-Intensive Cloud Control for GENI
Michael Zink (yes)
Experiment Management Service – Digital Object Registry
Jim French - CNRI (yes)
Giridhar Manepalli - CNRI (yes)

3. GENI I&M Use Cases and Architecture Requirements

Agenda for June 8:

1:15 pm	Suggest a basic set of GENI I&M use cases, and review contributions from key projects
1:45 pm	Discuss basic set of GENI I&M use cases, summarize consensus and identify gaps
2:30 pm	Break

3.1 Suggest a basic set of GENI I&M use cases and architecture requirements

Initial view of I&M vision, requirements, strawman, and WG Objectives for Spiral 2:
Paul Barford, University of Wisconsin, WG Co-Chair
November 18, 2009

Vision for GENI I&M
instrumentation and measurement systems provide broad data gathering, analysis and archieval capability
sufficient for scientific mission
sufficient for operations
key for success of the infrastructure
Requirements
measure details of GENI behavior with high precision and accuracy
no impact on experiments
ubiquitous
extensible
large capacity
high availability
resilient
strong access control
tight integration with CFs
Conceptual strawman
instrumentation - taps in the metwork and systems that provide basic signals
collection and synthesis - programmable systems that collect, combine and transform basic signals
archieve - measurement data index and repository
Instrumentation

link sensors - deployed on network links via taps, provide basic link signals
node sensors - deployed on all systems , provide basic utilization/state/configuration data
time sensors - deployed at all sites, provide fine-grained, synchronized timestamps

Collection and synthesis

programmable systems connected to sensors
transform basic signals into data suitable for more standard analysis
transformations can be more sophisticated
select/transfer protocol moves data from node sensors
short term storage capability

Data archieve

high capacity data repository deployed
data catalog

Security and access control

only accessible by authorized users
different views depending on authorization level
secure
private
some mechanisms defined by CF

NOTE: Requirements for data collected for experiments vs data needed for operations may be different.

An overview of basic GENI I&M use cases:
Ref GIMS_Design_UseCases: "Use-cases for GENI Instrumentation and Measurement Architecture Design", Prasad Calyam - Ohio Supercomputing Ctr

- **Use-cases for GENI Instrumentation and Measurement Architecture Design**
- **Prasad Calyam, Ph.D.**
- ***(PI – OnTimeMeasure, Project #1764)***
- **pcalyam@osc.edu**
- ***March 31st 2010***

- **What is different in GENI facility measurements?**
- **GENI supports testbeds aimed at "clean-slate" re-design of the Internet to overcome limitations of current Internet**
- **Users have greater options/control on measurements**
  - **Measurement server software/hardware**
  - **Advanced open-source/commercial instrumentation**
  - **Measurement service providers (who may customize)**
  - **Measurements across wired/wireless aggregates**
  - **Internet-scale measurements with "interesting" cross-traffic**

- **Goals for GENI Instrumentation and Measurement Architecture (GIMA) Design**
  - **Provide drill-down performance transparency of system and network resources at *hop*, *link*, *path* and *slice* levels**
  - **Allow and make-it-easy for users (NOC staff, experimenter) to access and control instrumentation and measurement functions involving interactions between GIMS sub-services**
  - **Remove burden on researcher to become a network measurement infrastructure expert so that researcher can better focus on the science in the experiments**
  - **Provide performance transparency of the status of the individual GIMS sub-service components and their interfaces with other sub-services**
- **For each sub-service (e.g., MP, MC, MAP, MO, MDA) in GIMA, following information could be specified:**
  - **Capabilities**
  - **Input, Output**
  - **Instrumentation components**
  - **Software components**
  - **Schemas**

- **Use cases from User point-of-view**
  - **Interfaces: Web-pages, Command-line options**
  - **Classification**
    - **NOC monitoring**
    - **Experiment monitoring**
    - **Measurement utilities**

- **NOC Monitoring**
  - **Capabilities: Availability, Heath Status, Diagnosis of perceived or impending problems ---- *context of the entire physical infrastructure***
    - **Availability: Up/Down, Up-Good, Up-Acceptable, Up-Poor**
    - **Health Status: Metrics and their levels for Hop, Link, Path and Slice**
  - **Use cases:**
    - **For a physical topology of Nodes {A, ... Z} show me if any slice is mis-behaving so that I can invoke "emergency shutdown" to swap it out**
    - **Experimenter called NOC about non-responsiveness of resources or unexpected behavior in a slice, notify status of user slice resources**
    - **We would like to keep meta-data of all the experiments, send us experiment meta-data after each slice expires**

- **Experiment Monitoring**

- **Capabilities: Availability, Heath Status, Diagnosis of perceived or impending problems ---- *context of the experiment slice***
  - **Availability: Up/Down, Up-Good, Up-Acceptable, Up-Poor**
  - **Health Status: Metrics and their levels for Hop, Link, Path and Slice**
- **Use cases:**
  - **A slice has been setup for me, have I got all the resources I asked for**
  - **Show me a dashboard of some or all of the resource performance measurements as I run my experiments**
  - **My experiment data shows inconsistencies, let me query the status of user slice resources so that I can notify GMOC about it**
  - **Provide me with an archive of some or all of the slice resource performance measurements so that I can reference them during offline analysis of the data collected in my experiment after the slice expires**

- **Measurement Utilities**
  - **Capabilities: Active measurements and passive measurements ---- *context of the experiment slice pertaining to research needs***
    - **Support tools that researchers of different problem domains will want to use (e.g., traffic engineering researcher will want SNMP, TCP protocol researcher would like throughput measurements from Iperf, video quality researcher would like PSNR measurements from Evalvid)**
  - **Use cases:**
    - **Setup up passive measurement taps at hops a, b, c**
    - **Setup up active measurements on paths x, y, z using p, q, r tools**
      - **On-demand or On-going (sampling patterns of periodic, random, stratified random, adaptive)**
    - **I am writing an event-driven simulation, at certain time points, I would like to be notified of anomalies and forecasts of system and network performance at hops a, b, c on paths x, y, z pertaining to tools p, q, r**
    - **I am running an experiment to deploy a novel IPTV system protocol, provide me with PSNR measurements of video quality between paths x, y, z (e.g., Evalvid tool that will need source and destination packet captures)**
    - **Provide me with an archive of some or all of the slice resource performance measurements that I requested as part of my experiment**

- **Use cases from measurement-services designer point-of-view**

- **How will we authenticate NOC staff versus researcher and what measurement privileges can we assign to users based on roles**
- **What is the workflow for a user to interface with a measurement service that manipulates the user's slice resources**
- **What is the schema we will use to exchange various "messages" between the measurement sub-services**
- **What is the schema we will support for users to query measurement data using web-service clients**
- **What are the sorts of examples/templates of measurement service usage that should be made available**
- **?**

Jim : 6 user groups:

NOC – cross CF
CF –management of their control framework + Aggregate providers
Archive providers
Experimenters
Experiment users (Opt-in users?)
Archived data for researchers (user looking for archived data)

Consider archived data for researchers (user looking for archived data):

Brian on 6/9 via email: If we want to make this happen, we need to make it "easy for data providers (e.g., instrumentation device designers, experiments who implement custom instrumentation devices, etc) to supply data to other entities."

He suggests a transformation service to accomplish this:

Being able to define an I&M schema would be great, but there are some practical problems with this approach that limit it (particularly as time goes on). A basic goal of the I&M framework should be to make it easy for data providers (e.g., instrumentation device designers, experiments who implement custom instrumentation devices, etc) to supply data to other entities. Supplying data for their own needs is generally fairly easy, but I believe we also want to encourage them to make as much of their data as possible (within legal and ethical limits) to be public – other users may find value in the data from an experiment (or some set of experiments in combination) that the original data providers did not envision, and this is only possible if the data is available (and can be found, of course, but that is mostly a different problem). Based on past experience, data providers tend to make data available to others when it is easy to do so, and the requirement to meet standard formats is seen by many data providers as a burden that they choose to ignore (i.e., they don't make the data available to others if it requires any effort on their part).

This approach has been used successfully in other areas, including a variety of scientific domains. For example, in the early 90s there was an international effort to provide data centers to which scientist could submit their data from individual studies of environmental contamination in the Arctic region. By the mid to late 90s the United States has contributed no datasets, primarily due to data format issues (the data centers had data formats, and the scientists in the US had no reason to spend the time to format the data in the required formats). A transformation library approach was applied at two of the data centers, resulting in a substantially simplified data submission process, with the result that US scientist began contributing their datasets (as it required almost no additional effort on their part). We face a similar challenge in GENI, where we have lots of data providers, and some data aggregation/collection systems (both in real time and for archival purposes).

Adding transformations to the GENI I&M infrastructure

- Likely to have lots of different data providers, not all of which will have same schema. This is particularly true when experimenters implement new instrumentation methods.
- Likely to have lots of data consumers, not all of which will have same schema
- Even if we agree on a schema today, some new requirement (device, data, frequency, format, statistical analysis, etc) tomorrow may (probably will) require changes to the schema

Approach to this is to embed transformation capability into the framework (probably at the collection points).

- Acceptable format gets null transformation
- Transformable format gets transformed
- Ability to add new transformations when necessary

Some advantages of this approach

- Schema changes/upgrades (don't all have to upgrade at one moment, as the transformation can handle this)
- New data providers (don't have to have them fit our format – they give us what they have, and we manage the transformation to what we need). We therefore reduce the burden placed on data providers and encourage them to make their data available.
- Data providers sending data to multiple consumers (collection points). Provider sends its native data format and collection points transform to their desired format.
- Can also be applied to metadata (at least to provide some automated metadata in the case of missing

- We can reduce some duplication of effort – once someone in the GENI world writes an A→ B transformation, there is a good chance no-one else has to (if we implement the transformation capability carefully)

To some extent we're doing this informally at some locations (GMOC, for example), but if we build this capability into the I&M infrastructure we can benefit across the GENI project (e.g., transformations written by GMOC can be applied at DOR, or vice versa)

Basic types of I&M services, and requirements:

1) Experiment I&M services

    Structure part of Researcher's slice
    Configured by Researcher
    Measurement data "owned" by Researcher, and they decide who can use it

Chris S.: Some of this data is owned by opt-in uses. Privacy and anonymity are major issues.

Ivan: This is a thorny problem---users don't even know what data is being collected about them. Rutgers has policy that cannot backup data beyond 60 days. What implications does this have on GENI archives?

Agreed: This workshop focuses on mechanisms, but not policies, so generally out of scope for this workshop.

Bruce: We need to provide mechanisms that will allow expts to implement different policies.

Martin: Mechanisms for controlling access to data should be same as mechanisms used to control access to other GENI resources.

    Can Operator ever see this data?
    Needs to be easy to use
    Like OMF/OML
    Like Instrumentation Tools

2) Network/Testbed I&M services

    Structure part of Operator's slice
    Structure and basic measurements configured by Operator

2a) Common set of measurement data

    In public domain
    Anonymization

Advertised

Can be shared with other Operators and Researchers, when authorized by Operator

Like perfSONAR

When used by a Researcher, they are receiving data from multiple slices

Jim : Look at PlanetLab. Nobody knows how the network as a whole is doing.

Jim: How often does an experimenter need to go outside the expt. to look at "common measurements".
If they need it, they can go to GMOC to get this common info. Or, they can use the control framework to get this data?
Lack of access to common data is a problem w. PL. Can't tell if an expt. isn't behaving as expected because of problems with the substrate.

Max: PL does collect common data (esp. reliability of nodes and links).
Other communities (medical, physics) are doing a lot of work in the area of archiving expt. data and sharing expt data. Even sharing data from running expts is being looked at by other disciplines.

Brian: Isn't this simply an access control issue. There are entities that own resources, entities that own data collected about these resources, and entities that can access this data.

Martin: Agree. Same mechanisms used to describe access to resources should be used for measurement data.

2b) Customized sets of measurement data

Each set of data measurement created by a sliver that is part of a slice, typically a slice belonging to a Researcher

Each sliver provides customized data using, for example, distinct filters.

Each set of measurement data "owned" by corresponding Researcher, and they decide who can use it

Like ShadowNet

3) Interoperability of I&M services

Essential for efficient development of I&M structure

Essential for efficient use of I&M structure, including mix of measurement data from both experiment and network/testbed I&M services

Requires services within essentially all Aggregates to exchange data, even when an Aggregate uses private IP space

Interoperability between services needs to be authorized and established via Control Framework mechanisms

One method to authorize communications between services: CF drops keys or credentials into both services

Jim: What is the relationship between the CF and measurement system. His project spends more time talking to Rob of PGENI than writing code.
Is the measurement resource a 1st class object?
Perhaps an authorization service used by the control plane and the measurement plane.
Issue: Permissions on measurement data will outlive slice.

Ivan: Measurement support infrastructure should be treated as a resource?
Archival data set is a resource just as a substrate is a resource.
Issue: in one case experimenter owns resource and in another case the resource provider owns the resource.

Harry: There are layers of ownership: owners of the data, owners of measurement resources, owners of servers.
Should CF provide mediation mechanisms between users and data?

Prasad: Think of I&M as a service that users can access. Now it looks like a resource access to which is mediated by the CF.

Martin: CF should be involved with slice set up but not mediating access to data.

Jim: Widgets pulling in data may be instantiated and accessed using CF mechanisms. However using CF mechanism to access data may not be appropriate.

Ivan: Issue is where is data stored? If held by CF, then CF auth may be appropriate. If somebody else holds data then CF may not be appropriate.

Jim: Have had similar discussions w/ Rob. Difficult to do (technical reasons)

Harry: Measurement data can drive experiment. Movement & storage of measurement data is itself part of the experiment.
See, for example, DICloud project

Common data: does it always go to GMOC and they publish this? Not necessarily---resource owner may keep common data.

Issue: Data collected outside the slice (common data) may have to be cleaned up before it is published to protect privacy. This is a hard problem.

If there is a device that collects data, the organization that owns the device must be trusted to protect the data.

Common data: How does an experimenter get "common data" that isn't already being collected? Go to GMOC and have them collect this additional data? Or, go to substrate owner (e.g. to get ping times) and ask to use their resource to collect ping time.

Martin: If you have permissions to subscribe to substrate resources (e.g. ping times) you should get it.

Legal issues associated with substrate/aggregate handing over data to GMOC. Must trust the entity to which data is being handed over.

3.2 Review contributions from key projects

Brief contributions from:
OML (ORBIT Measurement Library) OMF (ORBIT Management Framework)
Instrumentation Tools
perfSONAR
Scalable Sensing Service (S3)
OnTimeMeasure for network measurements
GENI Meta-Operations Center and NetKArma
Virtual Machine Introspection (VMI)
Data-Intensive Cloud Control for GENI
Experiment Management Service – Digital Object Registry

3.3 Discuss basic set of GENI I&M use cases and architecture requirements, summarize consensus and identify gaps

What use cases should we document?
We have user groups. Use cases for user groups.

Can we agree on basic types of I&M services, and requirements?

What gaps have been identified?

4. GENI I&M Measurement Plane, Services, Interfaces and Protocols

Agenda for June 8:

2:45 pm	Suggest definition of GENI I&M measurement plane, services, interfaces and protocols (APIs), and review possible contributions from key projects
4:15 pm	Break
4:30 pm	Discuss GENI I&M measurement plane, services, interfaces and protocols (APIs), summarize consensus and identify gaps
6:00 pm	Adjourn

4.1 Suggest definition of GENI I&M services

Fig 1-1: I&M Services for Researchers

Fig 1-2: I&M Services for Operators

Fig 1-3: I&M Services

Including: Lookup Service. Also Topology Service?

4.2 Review possible contributions from key projects, and discuss

Review how each project maps to the suggested I&M Services:

Fig 2-1: OML (ORBIT Measurement Library) OMF (ORBIT Management Framework)

Max: For OMF, measurement is considered to be inside a slice.
Same framework needs to be used to instrument user provider resources.
OMF treats measurement resources just like any other resources.
Instrumentation layer itself can be instrumented to steer it (just as expt. is steered based on measurements).

Fig 2-2: Instrumentation Tools

Fig 2-3: perfSONAR

LookupSvc has meta-data but not data.
Topology service: Used in PerfSonar to figure out what measurements mean. What is its role in GENI? Is this info held by the CF (in the clearinghouse?).

Fig 2-4: Scalable Sensing Service (S3)

Puneet: S3 Experimenter interacts with orchestration svc that then talks to the meas pt services.

Fig 2-5: OnTimeMeasure for network measurements

Fig 2-6: Data-Intensive Cloud Control for GENI

Note: Flow of measurement data is central to Data-Intensive Cloud Control project. Does this enlarge our view of the I&M architecture?

DI Cloud: How much is the CH involved since resources are needed to do expts?
The measurement collection service is in the cloud.
The presentation, analysis and archiving of data is in the cloud----archiving missing from the diagram.

Fig 2-7: Experiment Management Service – Digital Object Registry

4.3 Summarize consensus and identify gaps

Our consensus is summarized in the following figure:

Fig 1-3b: I&M Services

What gaps have been identified?

4.4 Suggest definition of GENI I&M measurement plane and interfaces

GENI basics:

Includes infrastructure from a wide variety of Aggregates

Resources from any of these aggregates can be included in a Researcher's slice.

GENI backbone networking resources:

Currently provided by Internet 2 and NLR.

Including both IP backbones, and Layer 2 (VLAN) services.

Addresses on the IP backbones are not always reachable from the Internet.

GENI control and experiment traffic:

Control traffic is carried by the Internet and/or Internet2 and NLR backbones, so that a Researcher can setup experiments on GENI, while located at any site, without the need for special network access.

Question: Does this mean that Aggregates and other GENI resources connected to Internet2 or NLR must have publically reachable addresses? Or that the Researcher in this case must have access to the Internet2 or NLR backbones?

Experiment Traffic may be carried on a Layer 2 (VLAN) connections, setup as part of a Researcher's slice, to carry traffic between the included Aggregates.

Layer 2 (VLAN) connections are carried by arrangements of Ethernet switches and/or tunnels.

Some Experiment Traffic may flow to or from the Internet.

This is consistent with current ProtoGENI practice.

GENI Aggregates:

Some (or possibly most) of the Aggregates will have their resources (hosts, etc.) connected via private address space. They will not be directly reachable from the Internet, or Internet2 or NLR backbones.

Experiment Traffic carried by a Layer 2 (VLAN) network connection into the Aggregate will be able to connect with hosts, etc., in their private address space.

The Aggregate Manager is expected to have a public (or reachable) IP address, so that the Researcher can send messages to reserve resources, etc. In turn, it will interact with the hosts.

Some arrangement will be necessary for the Researcher to login to an assigned host in the private address space to load code, etc.

Assume: Researcher can login to a host in the private address space using an SSH Proxy, provided as part of the Aggregate, which has a public (or reachable) IP address.

Assume: Researcher, once logged-in to a host, can use SCP to download code from a repository with a public (or reachable) IP address to a host.

GENI Measurement Traffic:

The flow of GENI Measurement Traffic (the “measurement plane”) has not yet been defined.

We have two obvious choices: Internet, or Internet2 or NLR backbones, like Control Traffic; or an assigned Layer 2 (VLAN) connection like Experiment Traffic.

Max: We don’t need to differentiate this from how we connect resources across aggregates.

Ivan: Must be able to guarantee some QoS for measurement data (e.g for steerable expts where measurements steer expts)

Jim: Measurement traffic must be carried reliably even if things such as routing go wrong with the experiment.

But, measurement traffic can overwhelm a path – this needs to be prevented, if possible.

Max: Still no major difference between experiment resources and measurement resources.

Some Measurement Traffic will flow between the Researcher and the I&M services, like Control Traffic.

Some Measurement Traffic will flow between I&M services, like Experiment Traffic.

If we require a Layer 2 (VLAN) network connection for Measurement Traffic, it is another complication in setting up I&M services.

Assume: Carry most (if not all) Measurement Traffic like we carry Control Traffic, via Internet, or Internet2 or NLR backbones

Assume: When an I&M service is in an Aggregate with private IP addresses, include proxies (or other access servers) to allow the necessary access.

Assume: Some measurement traffic may be carried via a Layer 2 (VLAN) network connections, but preferably implemented by a tunnel arrangement, to avoid the need for Ethernet switches

This approach is summarized in Fig 3-1.

Ivan: The expt console in OMF is an ssh or http proxy.

Max: Connectivity of measurement points, collectors, etc. is a CF problem. Not an I&M problem.

Jim: Most of what is in the figure is already being provided by CFs: e.g. ssh proxy, http proxy, etc.

Fig 3-1: Measurement Traffic Flows

GENI Measurement Traffic Proxies:

Proxies are required when the desired I&M services are located in an Aggregate that uses private IP addresses.

Several proxies are required, depending on the underlying protocol.

Authentication and authorization must be managed by the GENI Control Framework (CF). Is is assumed that this is done by "dropping keys or credentials" into appropriate services.

The suggested proxies are shown in Fig 3-2.

SSH Proxy

HTTP Proxy

VPN Access Server, to provide a tunnel between two Aggregates.

Fig 3-2: Measurement Traffic Proxies

Reference that explains two approaches to HTTP-based web services:
Ref MeasPlane-1: RESTful Web Services vs. "Big" Web Services: Making the Right Architectural Decision

4.5 Review possible contributions from key projects, and discuss

OML (ORBIT Measurement Library) OMF (ORBIT Management Framework):
Carries measurement traffic between hosts and services within a site via a dedicated VLAN.

Instrumentation Tools
Follows protoGENI arrangements, and carries measurement traffic like control traffic, using Internet2 backbone.

perfSONAR
Carries measurement traffic over backbone IP network

Scalable Sensing Service (S3)
Carries measurement traffic over backbone IP network

OnTimeMeasure for network measurements
Carries measurement traffic over backbone IP network

GENI Meta-Operations Center and NetKarma

Virtual Machine Introspection (VMI)

Data-Intensive Cloud Control for GENI
Carries measurement data over VLAN connection

Experiment Management Service – Digital Object Registry

Suggested:
Continue discussion after we consider interfaces and protocols.

4.6 Suggest definition of GENI I&M interfaces and protocols (APIs)

Referring to Fig 1-3, we define these I&M interfaces and messages/flows/APIs:

1) Discover Resources and Assign Slivers: EC srvc uses CF to discover resources, and then assign slivers to slice/researcher for I&M srvc's

2) Configure and Program Slivers: EC srvc uses CF and/or ssh to load std or customized software images for I&M srvc's

Note: 1) and 2) are not specific to I&M services

3) Manage Services: EC srvc and MO srvc use CF and/or https to check status of I&M srvcs, receive event notifications, and execute functions such as start, stop, reset, reboot, and checkpoint

4) Measurement Data Flows: Measurement data flows between I&M srvcs. Two options: Pull and Push.

5) Measurement Data File Transfers: Measurement data file transfers between I&M srvcs. Expect to Pull from and Push to Repository

6) Register I&M Service: Operator configures I&M srvc to register with Lookup Srvc, advertising name, location, and available metadata

7) Discover I&M Service and Establish Meas Data Flow: ECS or I&M srvc discovers I&M srvc advertisement, and establishes data flow

8) Conduct and Observe Experiment: Researcher uses browser to interact with and observe services via web portals

Fig 1-3: I&M Services

4.7 Review possible contributions from key projects

4.7.1 OML (ORBIT Measurement Library) OMF (ORBIT Management Framework)

Summary:
Fig 4-1: OMF/OML Services and Messages

Also these references:
Fig 4-2: OML Component Architecture
Fig 4-3: OMF/OML Overview
Fig 4-4: ORBIT Network Diagram

Ref OMF_OML-1: "XDR: External Data Representation Standard"
Ref OMF_OML-2: "ORBIT Measurements Framework and Library (OML): Motivations, Design, Implementation, and Features"
Ref OMF-OML-3: "OML Overview" slides
Ref OMF-OML-4: "Measurement Architectures for Network Experiments with Disconnected Mobile Nodes"

Interfaces and protocols:

3) Manage Services

Via HTTP to all srvc's, with APIs based on REST.

~~Via HTTP to OML Client srvc, to config files specifying filtering and streaming, which are then compiled into code~~

Max: Everything in OMF is organized as services (http).

4) Measurement Data Flows

Researcher defines measurement streams, gathering data samples and averaging, etc.

Meas data is series of typed vectors, XDR coded, and then streamed from client to collection server using proprietary OML protocol, on top of TCP, over dedicated Control VLAN

Considering using IPFIX instead of prop OML protocol; IPFIX has many extensions, and typically uses SCTP for transport

If path becomes disconnected from time-to-time. data is cached in Proxy Server FIFO, and then forwarded when path is reestablished

Max: Seriously considering IPFIX for transport.

5) Measurement Data File Transfers

Meas Analysis Present Srvc running outside of OMF/OML.
Can import directly from SQL DB
EC can arrange to convert tables into graphs
Portal service to view

Max: OMF has a portal service that allows visualization of measurement database. Looking into moving to streaming DBs.

6) Register I&M Service

Max: No different from general resource discovery.

7) Discover I&M Service and Establish Meas Data Flow

Max: Does not map to OMF.

8) Conduct and Observe Experiment
Experiment Portal early prototype
Each experiment results in a separate page containing all the experiment related information (script, parameter, resources used, time) as well as a pointer to the measurement database.
Where?

Max: OMF has a portal, which sits behind the firewall, talks to different services, shows what is going on. Interprets xml from the different services.

Max on 6/9 via email:

I still believe that we should clearly separate what we expect to be provided by the aggregates and control frameworks and what is specific to measurements. I agree with you that a lot of hard questions have not been answered, but we can't take on all of the issues.

The capability to acquire and instantiate resources, being it compute nodes, links, storage, radios, ... is a control framework task as is the establishment of links between entities with specific QoS and security properties if available.

What we should concentrate on is the architecture and models of an instrumentation capability or framework. To that extent we (or maybe just me :) look at the world a bit differently, instead of a service oriented architecture, we primarily concentrate on 'streams'.

We have 'measurement (instrumentation) points' which can be instantiated and configured to produce a stream of tuples of a specific schema. These streams are directed towards processing nodes which consume streams (possibly multiple) and either store them or originate streams which are processed by other nodes. This is essentially what the streaming database community has developed (no points for originality).

This has a couple of implications:

* The processing nodes and links are resources which need to be acquired and accounted for. Something which requires proper authorization (note, this is different from authentication, which needs to be an existing system or control framework service). This includes instruments, computing and storage resources, as well as 'analysis' services which produce "higher level" information (e.g. perfSONAR's traceroute).

* Measurement streams are uni-directional between a sender and a receiver with no feedback (in the stream)

* Control and configuration is done through the same mechanisms the control framework provides (which also enforces access and policy)

* As we can model a tuple stream as an unbound table, we can support real-time (sequence of tuples) as well as offline (table where each row is a tuple). This also allows us to 'join & process' streams (e.g. joining on the packet id and calculating the difference of timestamps of packet observations at the source and sink to arrive at a latency stream) as well as 'batch process' them (e.g. visualise results, deeper statistical analysis through tools like R)

* As the streams are defined and created by the sending entity, the instrumentation system can instrument itself and report that downstream. That's especially important if it becomes overloaded and needs to shed load (throwing things away, or simplify processing)

* What we need to agree on is the protocol (or format) of the streams and we propose to use IPFIX with some extensions (which are within the standard) to make the streams more self-explanatory (essentially carry meta-data). (An alternative XML representation of an IPFIX dump is straight forward, if not already defined in an IETF draft, or we could adopt OGF's NM-WG).

* IPFIX can be transported over many channels as long as they are reliable.

* This model still encompasses a service oriented view, where the 'processing node' (there should be a better word for it) can be requested to return, or produce data - again, viewing this as a table with a specific scheme. (Now this last point may look a bit like a cop out :) but I want the request to go through the normal authorisation framework we already have for operations on resources).

4.7.2 Instrumentation Tools

Summary:
Fig 5-1: Instrumentation Tools Services and Messages

Also these references:
Fig 5-2: Instrumentation Tools Components
Fig 5-3: Instrumentation Tools Topology

Ref InsTools-1: "Architectural Design and Specification of the INSTOOLS Measurement System"

Jim: Much of what instr. tools has done relates to items 1 & 2.
Control thru RSpec.

Interfaces and protocols:

3) Manage Services
   MC Srvc has collection control software
   MP Srvc includes remote access daemon to execute capture software

Jim: Was in Emulab. Now re-implementing in protoGENI

4) Measurement Data Flows
   MC Srvc gets data from MP Srvc via SSH/SCP
   MC Srvc gets data from MP Srvc via SNMP
   Emulab (ssh) key distribution mechanism used to authorize MC to get data from MPs
5) Measurement Data File Transfers

Jim: Using ssh and scp to move files.

6) Register I&M Service
7) Discover I&M Service and Establish Meas Data Flow
8) Conduct and Observe Experiment
   Portal to MCs, then GUI in MCs displays data as table or graph

Jim: Done on demand based on what people want to see (using Drupal).

Only control: Instrumentation on or off.
View collected data using content management system.

Portal to various measurement collectors
Every aggregate has a measuremet controller.
Portal has links to each of the MCs.
Can click to any MC.

How can new collection tools be added easily?
How do you add new instruments to the system?
Have added a way of adding new tools to the content management system.

4.7.3 perfSONAR

Summary:
Fig 6-1: perfSONAR Services and Messages

Also these references:
Fig 6-2: perfSONAR Measurement data Schema

Ref perfSONAR-1: "Scalable Framework for Representation and Exchange of Network Measurements"
Ref perfSONAR-2: "An Extensible Schema for Network Measurement and Performance Data"
Ref perfSONAR-3: "NM-WG/perfSONAR Topology Schema"

Interfaces and protocols:

3) Manage Services

Manage Services GUI: perfAdmin (CGI script to locate and manage perfSONAR services and data)

How does this work?

perAdmin is more of a GUI to see what is out there. (No authentication).

4) Measurement Data Flows

Pulls data from MA srvc, with these messages:

Echo Request – Not specific to measurement. All services need to be able to respond to echo req

Metadata Key Request

Setup Data Request

Note:

All perfSONAR Messages
addressed to each service at a service URL
formatted in XML using SOAP over HTTP
always a Request and then a Response

Also an interface that is pub-sub (push) [uses WS Notify].
Looking at other data formats (e.g NetLogger, a compressed xml format, ..)
Metadata is regular and extensible. How it is encoded is a different issue.
Looking at AMQ---advanced message queuing protocol---high perf pub-sub system.

Q: Pub-sub or not? How to do authorization in pub-sub?

Ivan: It is possible to do so. Can constrain who can access the channel---broker can deny access to channel.
Or can encrypt data and give keys only to auth subscribers. Each channel has an id (channel name). Whoever creates the channel owns it (likely the publisher). Most systems use xmpp for pub-sub.

pub-sub will continue to grow in PerfSonar world.

(currently) no encryption
(currently) no authentication and authorization
Not completely true. EU developent uses edugain based on Shib.

(since Authentitcation Srvc (AS) not yet built or deployed)

Note:

Each message follows perfSONAR schema, and contains;
message container
one or more one metadata elements
zero or more data elements

Perfsonar metadata format is extensible.

5) Measurement Data File Transfers

Used to ship whole sql-lite files. No std mechanism. Add to document?

6) Register I&M Service

Each MA service registers with LS service
homeLS registers with globalLS
globalLS updates other globalSL
LS Register Request
LS Deregister Request
LS Keepalive Request
Also, operator (?) registers topology information, with these messages:
TS Query Request
TS Add Request
TS Update Request
TS Replace Request

Being changed to make it better.
Introducing another level in hierarchy. Going from home LS to global LS wasn't part of original design. Now can accommodate arbitrary levels of hierarcy.
Moving to REST.
Also changing topology schema.
Unification of lookup service & topology service.
An API shields discovery from clients (don't have to know how many LSs to talk to before finding service).
Similar to DNS. Can add new services (define new name space).

Jim: Authentication and Auth---user management system. Lots of way of doing it. Who is going to do it?

7) Discover I&M Service and Establish Meas Data Flow
   Client can discover service and gain access to data using these messages:
   LS Query Request - XQuery
   LS Query Request – Discovery
   LS Key Request
8) Conduct and Observe Experiment
   GUI types on MAP srvc includes:
   active Service
   GMAPS
   acad (Java-based visualization)
   E2EMon (link monitoring)
   ESNet (domain utilization)
   trace (traceroute visualization)
   perfAdmin (CGI script to locate and manage perfSONAR services and data)
   perfER GUI (displys the results of pingER testing)
   perfSONAR-BUOY (displays the results of latency and bandwidth testing)

4.7.4 Scalable Sensing Service (S3)

Summary:

Fig 7-1: Scalable Sensing Services (S3) Services and Messages

Interfaces and protocols:

3) Manage Services
   Via HTTP to GUI on Sensing Info Mgmt Backplane
   Via HTTP to GUI on Sensor Pods
   How?

On demand measurement. Start/stop are main controls.

4) Measurement Data Flows
   Pull via HTTP from Sensor Pod web intfc
   Query, specified by URL parameters
   Control?
   Notification?

HTTP, with all plain text—not even xml (no encoding).

5) Measurement Data File Transfers

Measurements stored on local hosts. Pulled using ssh/scp.
Measurements moved to central location.

6) Register I&M Service

No registration.

7) Discover I&M Service and Establish Meas Data Flow
8) Conduct and Observe Experiment
    GUI on Sensing Info Mgmt Backplane

Portal for experimenters (demo at GEC8)
Queries on sql dbase.
Scripts to extract information from plain text.

Plan to do admission control on measurements.

Have estimates on load introduced by different kinds of measurements.
Can feed into what measurements to admit.
Also can infer some measurements.

Can request measurements to be conducted periodically.

Measuring end-to-end path properties.
Like perfSONAR, used mostly for infrastructure monitoring.

4.7.5 OnTimeMeasure for network measurements

Summary:
Fig 8-1: OnTimeMeasure Services and Messages

Interfaces and protocols:
3) Manage Services

Central scheduler for admitting measurement requests. Coordinates measurements---detect conflicting measurements, schedule requests. Implements policies including how much resources can be used for measurement.

Measuremments in a dbase (proprietary schema)

Configuration: Similar to Instrumentation Tools

4) Measurement Data Flows
    Pull via HTTP from Node Beacon and Root Beacon web interfaces?

Using scp
Trying to do http

5) Measurement Data File Transfers

Raw files and processed files (e.g. time series data).
Use Graphite to see different kinds of dashboards.
sql dump for archives

6) Register I&M Service

7) Discover I&M Service and Establish Meas Data Flow

Will probably use perfSONAR schema.

8) Conduct and Observe Experiment

Via HTTP from GUI on Policy/Publish Authority?

web portal based. Future: command line interface also.

4.7.6 Data-Intensive Cloud Control for GENI

Summary:

Fig 9-1: Data Intensive Cloud Services and Messages

Interfaces and protocols:

3) Manage Services
4) Measurement Data Flows

A large amount of radar data flows "in real time" from radar system, through ViSE server, to Amazon EC2 and S3 resources, where it is collected and analyzed

Radar data follows NetCDF format.

Radar data flows to Amazon public IP address. How is this done?

~~Push or pull? Always as a file? How? Streamed in chunks? How?~~

~~One option: File transferred with ftp (or equivalent)~~

~~One option: File transferred with OPenDAP, that uses http to transfer data that can be in NetCDF format.~~

Data transfer done using LDM (local data manager) --- basically pub-sub over TCP.
Data chunked (a few hundred MBs).
Starts with LDM server and queue at radar system
Also LDM manager and queue on vise server.
Final LDM server and queue in cloud
Also push data from final queue to archive in cloud

Each LDM server configured to connect to other server(s) via channel.
Publisher LDM pubs chunks.
Reliable.
High data rate.

(CERN is using gridFTP to move terabytes of data between Tier 1 sites.)

5) Measurement Data File Transfers
6) Register I&M Service
7) Discover I&M Service and Establish Meas Data Flow
8) Conduct and Observe Experiment

4.7.7 Digital Object Registry

Summary:

Fig 10-1: DOR MDA Service Services and Messages

Also these references:

Fig 10-2: DOR MDA Service File Organization

Interfaces and protocols:

3) Manage Services
4) Measurement Data Flows
5) Measurement Data File Transfers

Interfaces to the MDA srvc include: https; scp or sftp

From another I&M srvc, MDA srvc can provide these basic functions: put/update file; get file; delete file

When file is first introduced, it is assumed that file contains type info (extension), metadata, and "file self description" info. A wide range of files and associated metadata is permitted by the MDA srvc.

Also persistent object identifier, e.g., handle.

Each file is "owned" by a GENI slice and one or more users (operators/researchers)

MDA srvc allows the owner to specify who has read and/or write access to the file.

MDA srvc utilizes the mechanisms provided by the CF to authenticate and authorize users.

Assume: CF drops public keys of authorized users into MDA srvc, so that: presence of key indicates an "account" on the MDA srvc; additional info indicates nature of access (CNRI)

6) Register I&M Service
7) Discover I&M Service and Establish Meas Data Flow
8) Conduct and Observe Experiment

DOR – for measurement archive service.

Can manage meas data

provide storage

enable discovery

provide access control on data

(has been done in other contexts)

-

4.8 Discuss GENI I&M interfaces and protocols (APIs)

Summary of interfaces and protocols from all projects:

3) Manage Services

OMF/OML:

Via HTTP to all srvc's, with APIs based on REST.

Instrumentation Tools:

MC Srvc has collection control software

MP Srvc includes remote access daemon to execute capture software

perfSONAR:

Manage Services GUI: perfAdmin (CGI script to locate and manage perfSONAR services and data)

How does this work?

Scalable Sensing Service:

Via HTTP to GUI on Sensing Info Mgmt Backplane

Via HTTP to GUI on Sensor Pods

How?

4) Measurement Data Flows

OMF/OML:

Meas data is series of typed vectors, XDR coded, and then streamed from client to collection server using proprietary OML protocol, on top of TCP, over dedicated Control VLAN

Considering using IPFIX instead of prop OML protocol; IPFXI typically uses SCTP for transport

If path becomes disconnected from time-to-time. data is cached in Proxy Server FIFO, and then forwarded when path is reestablished

Instrumentation Tools:

MC Srvc gets data from MP Srvc via SSH/SCP

MC Srvc gets data from MP Srvc via SNMP

Emulab (ssh) key distribution mechanism used to authorize MC to get data from MPs

perfSONAR:

Pulls data from MA srvc, with these messages:

Echo Request

Metadata Key Request

Setup Data Request

Note:

All perfSONAR Messages

addressed to each service at a service URL

formatted in XML using SOAP over HTTP

always a Request and then a Response

(currently) no encryption
(currently) no authentication and authorization
(since Authentitcation Srvc (AS) not yet built or deployed)
Note:
Each message follows perfSONAR schema, and contains;
message container
one or more one metadata elements
zero or more data elements
Scalable Sensing Service:
Pull via HTTP from Sensor Pod web intfc
Query, specified by URL parameters
Control?
Notification?
OnTimeMeasure:
Pull via HTTP from Node Beacon and Root Beacon web interfaces?
Data-Intensive Cloud Control;

A large amount of radar data flows "in real time" from radar system, through ViSE server, to Amazon EC2 and S3 resources, where it is collected and analyzed

Radar data follows NetCDF format.

Radar data flows to Amazon public IP address. How is this done? Push or pull? Always as a file? How? Streamed in chunks? How?

One option: File transferred with ftp (or equivalent)

One option: File transferred with OPenDAP, that uses http to transfer data that can be in NetCDF format.

5) Measurement Data File Transfers

OMF/OML:

Meas Analysis Present Srvc running outside of OMF/OML can import directly from SQL DB

DOR:

Interfaces to the MDA srvc include: https; scp or sftp

From another I&M srvc, MDA srvc can provide these basic functions: put/update file; get file; delete file

When file is first introduced, it is assumed that file contains type info (extension), metadata, and "file self description" info. A wide range of files and associated metadata is permitted by the MDA srvc.

Each file is "owned" by a GENI slice and one or more users (operators/researchers)

MDA srvc allows the owner to specify who has read and/or write access to the file.

MDA srvc utilizes the mechanisms provided by the CF to authenticate and authorize users.

Assume: CF drops public keys of authorized users into MDA srvc, so that: presence of key indicates an "account" on the MDA srvc; additional info indicates nature of access (CNRI)

6) Register I&M Service
   perfSONAR:
   Each MA service registers with LS service
   homeLS registers with globalLS
   globalLS updates other globalSL
   LS Register Request
   LS Deregister Request
   LS Keepalive Request
   Also, operator (?) registers topology information, with these messages:
   TS Query Request
   TS Add Request
   TS Update Request
   TS Replace Request

7) Discover I&M Service and Establish Meas Data Flow
   perfSONAR:
   Client can discover service and gain access to data using these messages:
   LS Query Request - XQuery
   LS Query Request – Discovery
   LS Key Request

8) Conduct and Observe Experiment
   OMF/OML:
   Experiment Portal
   Instrumentation Tools:
   Portal to MCs, then GUI in MCs displays data as table or graph
   perfSONAR:
   Many GUI types on MAP srvc
   Scalable Sensing Service:
   GUI on Sensing Info Mgmt Backplane
   OnTimeMeasure:
   Via HTTP from GUI on Policy/Publish Authority?

4.9 Summarize consensus and identify gaps

Possible consensus set of GENI interfaces and protocols:

3) Manage Services

Consider all contributions, but particularly OMF/OML and perfSONAR

a) Via HTTP (or HTTPS) to all services, standardized API, following SOAP to allow use of credentials that flow through HTTP proxy

SOAP can be heavy weight.
Giridhar: Can use id in https cert to look up credentials.

Answer depends on how much we want graphics
(http good) vs. scripting (xml-rpc).
Max: xml-rpc is synchronous.
http for transport is good.
What is the nature of the API?
What about revocation of credentials?
Max: We should have an asynch model.
Must be able to pass around self-contained xml strings.

4) Measurement Data Flows

Define a variety of supported flows
Both pull and push
Both repeated transactions and streaming

a) Pull with repeated transactions via SSH/SCP

Like **Instrumentation Tools:**
MC Srvc gets data from MP Srvc via SSH/SCP
Emulab (ssh) key distribution mechanism used to authorize MC to get data from MPs
Restricted to one site?

Why not same data flow service to write data to the archive?
Max: Issues should be naming, finding archived data using meta-data, access controls, etc.

Two issues: the transport mechanisms (e.g. scp). Don't restrict to any one data flow mechanism now. Think about variety of xport mechanisms since we don't know enough.
More important to define what is in the flow (schema).

Are we reinventing what the CF has done?

Naming is very important. Can we agree that everything has a URI (services, flows, etc).

Need to consider all options:

Protocol	Data Flows	File Transfers
SNMP	Pull	
SCP	Push	
HTTP	Pub/sub	
XMPP		
SCTP		
…		

b) Pull with repeated transactions via SNMP

Like **Instrumentation Tools:**

MC Srvc gets data from MP Srvc via SNMP

Emulab (ssh) key distribution mechanism used to authorize MC to get data from MPs

Restricted to one site?

Same as (a).

c) Pull with repeated transactions via HTTP (HTTPS), can use HTTP proxy to traverse some site boundaries

Like **perfSONAR:**

Pulls data from MA srvc

All perfSONAR Messages

addressed to each service at a service URL

formatted in XML using SOAP over HTTP

always a Request and then a Response

(currently) no encryption

(currently) no authentication and authorization

Each message follows perfSONAR schema, and contains;

message container

one or more one metadata elements

zero or more data elements

But consider also **Scalable Sensing Service** and **OnTimeMeasure:**

Same as (a)?

d) Consider (?) push with repeated transactions via HTTP (HTTPS); can use HTTP proxy to traverse site boundaries

e) Push stream via TCP or SCTP; needs VPN Access Server (or VLAN connection) to traverse site boundaries

Like **OMF/OML:**

Meas data is series of typed vectors, XDR coded, and then streamed from client to collection server using proprietary OML protocol, on top of TCP, over dedicated Control VLAN

Considering using IPFIX instead of prop OML protocol; IPFXI typically uses SCTP for transport

If path becomes disconnected from time-to-time. data is cached in Proxy Server FIFO, and then forwarded when path is reestablished

f) Push high-bandwidth stream via TCP or SCTP, needs VPN Access Server (or VLAN connection) to traverse some site boundaries

Like: **Data-Intensive Cloud Control:**

A large amount of radar data flows "in real time" from radar system, through ViSE server, to Amazon EC2 and S3 resources, where it is collected and analyzed

Radar data follows NetCDF format.

Radar data flows to Amazon public IP address. How is this done? Push or pull? Always as a file? How? Streamed in chunks? How?

One option: File transferred with ftp (or equivalent)

One option: File transferred with OPenDAP, that uses http to transfer data that can be in NetCDF format.

5) Measurement Data File Transfers

a) Define a basic method to push or pull files, using HTTPS, SCP or perhaps SFTP

Like **DOR:**

Interfaces to the MDA srvc include: https; scp or sftp

From another I&M srvc, MDA srvc can provide these basic functions: put/update file; get file; delete file

When file is first introduced, it is assumed that file contains type info (extension), metadata, and "file self description" info. A wide range of files and associated metadata is permitted by the MDA srvc.

Each file is "owned" by a GENI slice and one or more users (operators/researchers)

MDA srvc allows the owner to specify who has read and/or write access to the file.

MDA srvc utilizes the mechanisms provided by the CF to authenticate and authorize users.

Assume: CF drops public keys of authorized users into MDA srvc, so that: presence of key indicates an "account" on the MDA srvc; additional info indicates nature of access (CNRI)

6) Register I&M Service

a) Messages formatted in XML using SOAP over HTTP to Lookup Service

Like **perfSONAR:**

Which GENI services should be registered? Even those dedicated to experiments?

Lookup services: Do we need a separate lookup svc for I&M? PerfSonar has one if we do.

7) Discover I&M Service and Establish Meas Data Flow

a) Messages formatted in XML using SOAP over HTTP to Lookup Service and other services

Like **perfSONAR:**

8) Conduct and Observe Experiment

a) HTTP (HTTPS) to GUIs at web services.

Following SOAP so that credentials can pass through HTTP proxies
Many GUIs have been defined and can be used

Is there anything standardized about these GUIs?

Can we reach a consensus on defining these interfaces and protocols?

Which gaps have been identified?

4.10 Discuss GENI I&M measurement plane and interfaces, summarize consensus and identify gaps

Based upon our discussion of interfaces and protocols, can we agree on modified measurement traffic flows and proxies?

Fig 3-1b: Measurement Traffic Flows

Fig 3-2b: Measurement Traffic Proxies

What gaps have been identified?

5. GENI Measurement Data Schema

Agenda for June 9:

8:00 am	Suggest contents and structure of GENI measurement data schema, and review possible contributions from key projects
9:30 am	Break
9:45 am	Discuss contents and structure of GENI measurement data schema, summarize consensus and identify gaps
11:15 am	Break

5.1 Suggest contents and structure of GENI measurement data schema

GENI measurement data schemas:

Multiple schemas will be defined and used within GENI, reflecting many different I&M arrangements

Examples from current projects:

a) perfSONAR defines certain network measurements, identifies source and/or point of measurement, and stores time, value pairs (not arrays, that I can tell)

b) OMF/OML has researchers define experiments, and collect data. No metadata, since researcher knows what it is. Probably time, value pairs, but I suppose it could be arrays.

c) ViSE radar collects arrays of data, following one of the NetCDF formats; later, can be used as input to well-defined visualization programs.

Definition of a GENI measurement data schema:

1) Mechanism:
    - Packets of chunks, in messages or flows
    - Files or records, in transfers or storage
2) Format of data:
    - Text
    - Bytes
    - Binary

3) Metadata:
    - How attached
    - Format of contents

How is GENI measurement data schema defined?

Contents of metadata:

1) Flow or record identifier (1 or more):
    - Index
    - Local
    - Globally unique
2) Annotation, for identification and/or searching (1 or more):
    - Slice, experiment, run
    - Researcher's notes
3) Provenance:
    - Owner, contact
    - Access rules
    - Encryption
4) Privacy:
    - Type of private info
    - Access rules
    - Anonymity applied
5) Processing of data (zero or more):
    - Where, when, what
    - Filters applied
6) Collection of data
    - Where, when, what
    - Filters applied
7) Description of data:
    - Time and value, pairs or tuples
    - Logs or events, with timestamps
    - File(s)
    - Binary images

Meta-data may apply to a single data set or a collection of data sets.
What meta-data is required by GENI?
Some meta-data fields may be invariant and others may change?

5.2 Review possible contributions from key projects

5.2.1 OML (ORBIT Measurement Library) OMF (ORBIT Management Framework)

Summary:
Fig 4-1: OMF/OML Services and Messages

References:
Ref OMF_OML-1: RFC4506 – XDR: External data Representation Standard
Ref OMF_OML-2: "ORBIT Measurements Framework and Library (OML): Motivations, Design, Implementation, and Features"
Ref OMF-OML-3: "OML Overview" slides
Ref OMF-OML-4: "Measurement Architectures for Network Experiments with Disconnected Mobile Nodes"

Interfaces and protocols:

4) Measurement Data Flows

Researcher defines measurement streams, gathering data samples and averaging, etc.

Meas data is series of typed vectors, XDR coded, and then streamed from client to collection server using proprietary OML protocol, on top of TCP, over dedicated Control VLAN

Considering using IPFIX instead of prop OML protocol; IPFXI typically uses SCTP for transport

If path becomes disconnected from time-to-time. data is cached in Proxy Server FIFO, and then forwarded when path is reestablished

5) Measurement Data File Transfers

Meas Analysis Present Srvc running outside of OMF/OML.
Can import directly from SQL DB
EC can arrange to convert tables into graphs

Measurement data schema:

4b) Measurement data flow schema:

Meas data follows schema defined by researcher, including: measurement-point id's, metric id's, etc.

A sensor (or application, or service) define a set of measurement points, with each measurement point defined by a name and a typed vector (sensor schema).

At runtime, the experimenter (or operator) provides a streams spec which defines what measurement points are going to be activated and what initial processing is going to be performed - that defines the actual schema going over the wire and/or ending up in the collection database

Ivan & Max: Strongly advocate use of IPFIX templates (data model).
Already has basic data types defined including data types relevant for networking.
Also has an extension mechanism (which is what spectrum sensing folks have done).
We may want to develop recommendations on how we can use IPFIX in GENI.
Don't want to put everything in the meta-data carried with data.

5b) Measurement data storage schema:

Application definition is used to create DB schema for experiment, using XSLT.

DB table is created for each measurement point, names based on id attribute of the group element.

Includes mandatory fields for name/id, timestamp, sequence number

Protocol is self describing

Server automatically creates a table for every distinct stream (distinct in terms of schema not source).

Streams carry their own name which is translated into a database using a simple naming convention.

5.2.2 Instrumentation Tools

Summary:

Fig 5-1: Instrumentation Tools Services and Messages

Interfaces and protocols:

4) Measurement Data Flows

MC Srvc gets data from MP Srvc via SSH/SCP

MC Srvc gets data from MP Srvc via SNMP

Emulab (ssh) key distribution mechanism used to authorize MC to get data from MPs

Measurement data schema:

4b) Measurement data flow schema:

N/A?

5b) Measurement data storage schema:

Internal to Measurement Controller?

What?

5.2.3 perfSONAR

Summary:

Fig 6-1: perfSONAR Services and Messages

Also these references:

Fig 6-2: perfSONAR Measurement data Schema

Ref perfSONAR-1: "Scalable Framework for Representation and Exchange of Network Measurements"
Ref perfSONAR-2: "An Extensible Schema for Network Measurement and Performance Data"
Ref perfSONAR-3: "NM-WG/perfSONAR Topology Schema"

Interfaces and protocols:

4) Measurement Data Flows

Pulls data from MA srvc, with these messages:
Echo Request
Metadata Key Request
Setup Data Request

Note:

All perfSONAR Messages
addressed to each service at a <u>service URL</u>
formatted in XML using SOAP over HTTP
always a Request and then a Response
(currently) no encryption
(currently) no authentication and authorization
(since Authentitcation Srvc (AS) not yet built or deployed)

Note:

Each message follows perfSONAR schema, and contains;
message container
one or more one metadata elements
zero or more data elements

Measurement data schema:

4b) Measurement data flow schema:

5b) Measurement data storage schema:

Follows perfSONAR schema, and contains;
message container
one or more one metadata elements
zero or more data elements

message or store file

From Ref perfSONAR-1

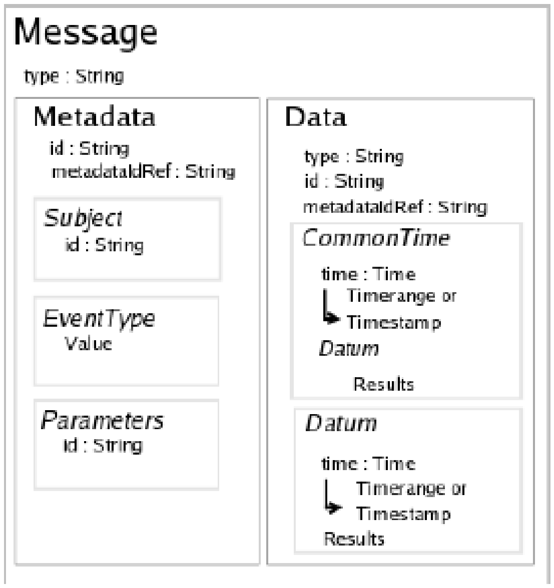

Figure 1. NM-WG Base Schema

The metadata section is subdivided into three parts, only the first of which is required:

- Subject – The physical or logical entity being described. For example, a host pair or router address. Like the subject of the sentence: *Host A to Host B* measured ICMP latency is 100ms.
- EventType – The canonical name of the aspect of the subject being measured, or the actual event (i.e. "characteristic") being sought. Like the object of the sentence: Host A to Host B *measured ICMP latency* is 100ms.
- Parameters – The way in which the description is being gathered or performed. For example, command-line arguments to *traceroute* or whether the round-trip delay packet used ICMP or UDP. Like the descriptive clause of the sentence: *When you use* 100 *byte packets*, Host A to Host B ICMP latency is 100ms.

CommonTime is valid for a group of data
Prefer NTP

MetaData blocks can be chained. Subject then points to prev. metadata block.

EvenType: the type of operation performed on data (e.g. average over multiple data items).

XML name spaces used to tell what subjects and parameters are valid for an eventtype.

Schema is extensible: Create new name spaces.

There is dependency on topology information for using this data.

No notion of owner in current PerfSonar metadata.

Topology schema:

Presently used for:
Internet2
UNINET (Norway)

From Ref perfSONAR-3:
Structured by layers and the same elements recurring there
- Varied by namespaces
- Reuse visualization logic, etc.
- Validate layer- or technology-specific attributes
- 4 Layers: Base (both abstract and L1), L2, L3, L4

(No longer valid. Can name specific technologies such as ethernet, IP, etc).
Identifiers use URN notation
- Prefixed with "urn:ogf:network:"
- Consists of name/value pairs separated by colons
- Possible field names: domain, node, port, link, path, network
- Set of rules defined for each field to keep identifiers compact and finite

Use Case:
- A client would use a topology service to look up the identifier for a network element and then would query a lookup service using the identifier to find the measurements associated with that element.

5.2.4 Scalable Sensing Service (S3)

Summary:
Fig 7-1: Scalable Sensing Services (S3) Services and Messages

Interfaces and protocols:
4) Measurement Data Flows
Pull via HTTP from Sensor Pod web intfc
Query, specified by URL parameters
Control?
Notification?

Measurement data schema:

4b) Measurement data flow schema:

What?

5b) Measurement data storage schema:

What?

5.2.5 OnTimeMeasure for network measurements

Summary:

Fig 8-1: OnTimeMeasure Services and Messages

Interfaces and protocols:

4) Measurement Data Flows

Pull via HTTP from Node Beacon and Root Beacon web interfaces?

Measurement data schema:

4b) Measurement data flow schema:

What?

5b) Measurement data storage schema:

What?

5.2.6 GENI Meta-Operations Center and NetKarma

Reference:

Ref GMOC-1: "GMOC Topology-Entity Data Exchange Format Specification"
Ref GMOC-2: "Proposal: Use of URN's as GENI Identifiers"

Measurement data schema:

4b) Measurement data flow schema:

5b) Measurement data storage schema:

From Ref GMOC-1:

Identifiers, encodings, and field sizes

We assume that at slices and devices are uniquely identified in GENI-wide by a human readable name. Names character set can be any unicode representable character set, but they must be encoded using UTF-8. Names are limited in size to 128 octets (bytes).

Principals are uniquely identified GENI-wide by a primary email address. Principals given names and last names are limited to 40 bytes. The email address is limited to 128 bytes.

Locations are uniquely identified (in the scope of each exchange document) by a human readable name.The minimal specification for a location is either the tuple (city, state_province, mail_code, country) or by the tuple (longitude and latitude).

Organizations are uniquely identified (in the scope of each exchange document) by a human readably name. These names are limited to 60 bytes.

Device's interfaces can be uniquely identified within a device by a device specific unique name. It is assumed that this name binding will be remain unique for subsequent documents as long as there are no changes in configuration of either the interface or the device.

Data model

In our data model every network device is considered a device. Devices can have a single parent device. The graph of the parent-hood for devices is a forest (a set of trees) .Slivers are modeled as virtual devices, that is a device with a parent device. Slices can be associated with both slivers and full devices. Each device can be associated with one sliver at most, this the graph of the relationship of slivers and devices is another forest.

A circuit is any network connection, between two or more devices. Circuits refer to any layer in the network stack and can be connected to any interface. Circuits can be build a multiplicity of other circuits. A circuit can be part of multiple circuits. The graph of circuit relationships is a disjoint set of directed acyclic graphs.

Data format

The date exchange format is defined using the relax-NG compact syntax as follows:

```
datatypes xsd = "http://www.w3.org/2001/XMLSchema-datatypes"
grammar {

start = element geni_aggregate {geni_aggregate-content}

geni_aggregate-content =
  attribute name {text},
  attribute public_key{text}?,
  element location {location-content}+,
  element contact  {contact-content}+,
  element organization {organization-content}+,
  element point_of_presence {pop-content}+,
  element device {device-content}+,
  element slice {slice-content}*,
  element net_topology {net_topology-content}

  location-content =
     attribute name {text},
     ( element address {address-content} |
      element geo_location {geo_location-content} |
       ( element address {address-content} ,

        element geo_location {geo_location-content} ) )

address-content =
   attribute address {text}?,
   attribute city    {text},
   attribute province {text},
   attribute country  {text}

geo_location-content =
   attribute latitude {xsd:double},
   attribute longitude {xsd:double}

contact-content=
   attribute email_address {text},
   attribute last_name {text},
   attribute given_names {text},
   attribute phone {text}?,
   attribute organization_name {text}?

organization-content=
   attribute name{text},
   element   primary_contact_email {text},
   element   location_name {text},
   element   parent_organization_name {text}?,
   element   url {text}?
```

```
pop-content =
    attribute name{text},
    attribute location_name{text},
    element  operator_org_name{text}?,
    element  admin_org_name{text}?

administrative_state-content =
    attribute state {"Planning"|"Provisioning"|"Available"|"NormalOperation"|
Maintenance" |"Unknown" | "Decomissioned"}

operational_state-content =
    attribute state {"Up" | "Degraded" | "Down" | "Unknown"}

device-content =
   attribute name {text},
   element  device_location {device_location-content},
   element  operator_org_name {text},
   element  admin_org_name {text}?,
   element  device_type {text},
   element  sw_version {text}?,
   element  hw_version {text}?,
   element  operational_state {operational_state-content}?,
   element  administrative_state {administrative_state-content}?,
   element  interface {interface-content}*

device_location-content =
    element pop_name {text} |
    element parent_device_name {text}

interface-content =
    attribute name {text},
    element contracted_bw {xsd:double}?,
    element max_bps       {xsd:double}?,
    element administrative_state {administrative_state-content}?,
    element net_addr {net_addr-content}*

net_addr-content =
    element net_addr_type {text},
    element addr{text},
    element netmask{text}

slice-content =
    attribute name {text},
    element   operator_org_name {text},
    element   primary_contact_email {text},
    element   device_names {text}+

net_topology-content =
    element  network {network-content}+,
    element  circuit {circuit-content}+,
    element  circuit_hierarchy {circuit_hierarchy-content}*
```

```
 network-content =
     attribute name {text},
     element   operator_org_name {text}?,
     element   admin_org_name {text}?

 circuit-content =
     attribute name {text},
     attribute circuit_type {text},
     element   channel {xsd:integer}?,
     element   reserved_bw {xsd:integer}?,
     element   vlan {xsd:integer}?,
     element   circuit_endpoint {circuit_endpoint-content}*

 circuit_endpoint-content =
     attribute   device_name {text},
     attribute   interface_name {text}

 circuit_hierarchy-content =
     element   upper_circuit_name {text},
     element   lower_circuit_name {text}
}
*
```

Currently getting all this data for PlanetLab (periodically) and ProtoGENI very occasionally.
GMOC can be queried for this information.
Problem: topology isn't defined for wireless networks.

5.2.7 Virtual Machine Introspection (VMI)

Measurement data schema:

4b) Measurement data flow schema:
Suggestions by Brian Hay?

5b) Measurement data storage schema:
Suggestions by Brian Hay?

5.2.8 Data-Intensive Cloud Control for GENI

Summary:
Fig 9-1: Data Intensive Cloud Services and Messages

Interfaces and protocols:

4) Measurement Data Flows

A large amount of radar data flows "in real time" from radar system, through ViSE server, to Amazon EC2 and S3 resources, where it is collected and analyzed

Radar data follows NetCDF format.

Radar data flows to Amazon public IP address. How is this done? Push or pull? Always as a file? How? Streamed in chunks? How?

One option: File transferred with ftp (or equivalent)

One option: File transferred with OPenDAP, that uses http to transfer data that can be in NetCDF format.

5) Measurement Data File Transfers

Measurement data schema:

4b) Measurement data flow schema:

5b) Measurement data storage schema:

Radar data follows NetCDF format.

Radar data flows to Amazon public IP address. How is this done? Push or pull? Always as a file? How? Streamed in chunks? How?

One option: File transferred with ftp (or equivalent)

One option: File transferred with OPenDAP, that uses http to transfer data that can be in NetCDF format.

5.2.9 Digital Object Registry

Summary:

Fig 10-1: DOR MDA Service Services and Messages

Also these references:

Fig 10-2: DOR MDA Service File Organization

Interfaces and protocols:

5) Measurement Data File Transfers

Interfaces to the MDA srvc include: https; scp or sftp

From another I&M srvc, MDA srvc can provide these basic functions: put/update file; get file; delete file

When file is first introduced, it is assumed that file contains type info (extension), metadata, and "file self description" info. A wide range of files and associated metadata is permitted by the MDA srvc.

Each file is "owned" by a GENI slice and one or more users (operators/researchers)

MDA srvc allows the owner to specify who has read and/or write access to the file.

MDA srvc utilizes the mechanisms provided by the CF to authenticate and authorize users.

Assume: CF drops public keys of authorized users into MDA srvc, so that: presence of key indicates an "account" on the MDA srvc; additional info indicates nature of access (CNRI)

Measurement data schema:

5b) Measurement data storage schema:

From another I&M srvc, MDA srvc can provide these basic functions: put/update file; get file; delete file

When file is first introduced, it is assumed that file contains type info (extension), metadata, and "file self description" info. A wide range of files and associated metadata is permitted by the MDA srvc.

5.3 Discuss contents and structure of GENI measurement data schema, summarize consensus and identify gaps

Can we agree:

Multiple schemas will be defined and used within GENI, reflecting different I&M arrangements

Can we agree to include schemas that describe these different I&M arrangements:

a) perfSONAR defines certain network measurements, identifies source and/or point of measurement, and stores time, value pairs (not arrays, that I can tell)

b) OMF/OML has researchers define experiments, and collect data. No metadata, since researcher knows what it is. Probably time, value pairs, but I suppose it could be arrays.

c) ViSE radar collects arrays of data, following one of the NetCDF formats; later, can be used as input to well-defined visualization programs.

d) Are there additional I&M arrangements that need to be described?

Can we agree on the following approach to defining a schema?

1) Mechanism:
    - Packets of chunks, in messages or flows
    - Files or records, in transfers or storage
2) Format of data:
    - Text
    - Bytes
    - Binary
3) Metadata:
    - How attached
    - Format of contents

Can we describe the schema for perfSONAR-like data?
Can we describe the schema for OMF/OML-like data?
Can we describe the schema for ViSE-like data?

Can we agree on the following contents in metadata?

1) Flow or record identifier (1 or more):
    - Index
    - Local
    - Globally unique
2) Annotation, for identification and/or searching (1 or more):
    - Slice, experiment, run
    - Researcher's notes
3) Provenance:
    - Owner, contact
    - Access rules

Encryption

4) Privacy:
    Type of private info
    Access rules
    Anonymity applied
5) Processing of data (zero or more):
    Where, when, what
    Filters applied
6) Collection of data
    Where, when, what
    Filters applied
7) Description of data:
    Time and value, pairs or tuples
    Logs or events, with timestamps
    File(s)
    Binary images
8) Are there other items that may need to be included?

Can we map the metadata contents to GMOC-requested data?
Can we map the metadata contents to perfSONAR-like data?
Can we map the metadata contents to OMF/OML-like data?
Can we map the metadata contents to ViSE-like data?

Can we agree which metadata is required for all schemas?

6. Identify Teams for Each Priority Topic

Agenda for June 9:

11:30 am	Identify teams for each priority topic, draft action items to close identified gaps, and make writing assignments for revised sections of the architecture document
12:30 pm	Lunch

Potential team members: (attended workshop: yes or no)

Paul Barford - University of Wisconsin – Madison (no)
Bruce Maggs – Duke University and Akamai (yes)
Harry Mussman – BBN/GPO (yes)
Vic Thomas - BBN/GPO (yes)
Evan Zhang – BBN/GPO (yes)

OML (ORBIT Measurement Library) OMF (ORBIT Management Framework)

Max Ott – NICTA (yes, by phone)
Ivan Seskar – Rutgers WINLAB (yes)

Instrumentation Tools

Jim Griffioen - Univ Kentucky (yes)

perfSONAR

Matt Zekauskas - Internet2 (no)
Jason Zurawski – Internet2 (yes)
Martin Swany - Univ Delaware (yes)
Guilherme Fernandes – Univ Delaware (yes)
Ezra Kissel – Univ Delaware (yes)

Scalable Sensing Service (S3)

Sonia Fahmy – Purdue (yes)
Puneet Sharma - HP Labs (yes)

OnTimeMeasure for network measurements

Prasad Calyam - Ohio Supercomputing Ctr (yes)

GENI Meta-Operations Center and NetKArma

Jon-Paul Herron - Indiana Univ
Camilo Viecco - Indiana Univ (yes)
Chris Small - Indiana Univ (yes)
Beth Plale - Indiana Univ (yes)

Virtual Machine Introspection (VMI)

Brian Hay – Univ Alaska (yes)

Data-Intensive Cloud Control for GENI

Michael Zink - UMass Amherst (yes)

Experiment Management Service – Digital Object Registry

Jim French - CNRI (yes)
Giridhar Manepalli - CNRI (yes)
Larry Lannom – CNRI (no)

6.1 GENI I&M use cases

Team members:
Paul Barford - University of Wisconsin – Madison (no)
Jim Griffioen - Univ Kentucky (yes)
* Prasad Calyam - Ohio Supercomputing Ctr (yes)
Camilo Viecco - Indiana Univ (yes)
Brian Hay – Univ Alaska (yes)
* agreed to organize first discussion and writing

Identify all user groups, and provide basic use cases:
1) GMOC operations group (NOC)
2) Cluster and aggregate operations groups (NOCs)
3) Archive service providers
4) Experiment researchers
5) Experiment (opt-in) users (see http://groups.geni.net/geni/attachment/wiki/041409NYCOptInWGAgenda/071509%20%20GENI-SE-OI-Overview-01.4.pdf for listing of opt-in issues, such as privacy)
6) Researchers that use measurement data archived by other researchers (DatCat model)

6.2 GENI I&M services

Team members:
* Harry Mussman – BBN/GPO (yes)
Evan Zhang – BBN/GPO
Giridhar Manepalli - CNRI (yes)
Chris Small - Indiana Univ (yes)
Beth Plale - Indiana Univ (yes)
* agreed to organize first discussion and writing

Summarize current view

Identify different types of services:
Completely dedicated to an experiment
Common portion, plus parts associated with different experiments
Common service, with data provided to multiple experiments

Need: Basic definition of an archive service

6.3 GENI I&M resources

Team members:

Vic Thomas - BBN/GPO (yes)
* Jim Griffioen - Univ Kentucky (yes)
Martin Swany - Univ Delaware (yes)
Camilo Viecco - Indiana Univ (yes)
Brian Hay – Univ Alaska (yes)
Giridhar Manepalli - CNRI (yes)
* agreed to organize first discussion and writing

Significant question uncovered at workshop!
Jim on 6/25 via email: We should involve Rob Ricci in the discussion.

What are resources:

1) Hosts, VMs, etc.
2) Network connectivity
3) Software, e.g., I&M software that can be included in an experiment
3) I&M services
4) I&M data flows and file transfers
5) I&M data files stored in archives

How are each of these discovered, specified, authorized and assigned:

a) Always by mechanisms provided by the CF?
b) With CF plus additional mechanisms?

Consider example of LS in perfSONAR
Consider example of data file stored in archive, owned by an experimenter

Goals:

Need to define and then compare these options
Need to understand interop with CF for each option
Does CF setup secondary authorization mechanisms in some cases? If so, how?

Does each item have:

Unique and persistent name?
Unique and persistent identifier?
Need to carefully consider this for all of GENI

For each item, consider how to:

Create
Name

Register and discover
Authorize and assign

For each item, consider:
Ownership
What sort of policies the owner may want to apply

6.4 GENI I&M measurement plane and interfaces

Team members:
* Harry Mussman – BBN/GPO (yes)
Ezra Kissel – Univ Delaware (yes)
Chris Small - Indiana Univ (yes)
* agreed to organize first discussion and writing

Consider:
IP network
Layer 2 (VLAN) connections

Discuss
Which protocols are active
Access to resources in aggregates, even when resources are in private address space, via GWs or proxies
How to provide authentication and authorization
How to provide QoS to protect measurement traffic
How to provide QoS to protect other traffic when measurement traffic is large.
Reserve bandwidth?

Martin on 6/28: Consider XSP (extensible session protocol) to provide transport layer GW functions.

6.5 GENI I&M interfaces and protocols (APIs): manage services

Vic Thomas - BBN/GPO (yes)
Ivan Seskar – Rutgers WINLAB (yes)
Max Ott – NICTA (yes, by phone)
* Sonia Fahmy – Purdue (yes)

Giridhar Manepalli - CNRI   (yes)
* agreed to organize first discussion and writing

Define an approach based on OMF/OML and S3:

HTTP(S)
REST vs SOAP
Authorization by credentials or ? If credentials, how to revoke?
Pass XML fragments
Define basic API

6.6 GENI I&M interfaces and protocols (APIs): data flows and data file transfers

* Harry Mussman – BBN/GPO (yes)
Ivan Seskar – Rutgers WINLAB (yes)
Max Ott – NICTA (yes, by phone)
Ezra Kissel – Univ Delaware (yes)
Prasad Calyam - Ohio Supercomputing Ctr (yes)
Michael Zink - UMass Amherst (yes)
* agreed to organize first discussion and writing

Consider data flows and data file transfers between all services

Define range of options:

What:

Data flows
Data files transfers

Type:

Pull
Push
Pub/Sub

Protocol:

SNMP
SCP
FTP and gridFTP
HTTP
XMPP
TCP
SCTP

Consider:

Naming
Discovery
Connectivity
Authentication and authorization mechanisms

Map to current projects, giving examples:
Consider: Minimum set required for GENI

6. 7 GENI I&M interfaces and protocols (APIs): service registration and discovery

Team members:
* Jason Zurawski – Internet2 (yes)
Prasad Calyam - Ohio Supercomputing Ctr (yes)
* agreed to organize first discussion and writing

Consider approach used in perfSONAR

Summarize for:
Services with data flows
Also sources of file transfers?
Also GUIs?

6.8 GENI I&M interfaces and protocols (APIs): GUIs

Team members:
Jeremy Reed - Univ Kentucky (yes)
* Guilherme Fernandes – Univ Delaware (yes)
Puneet Sharma - HP Labs (yes)
* Agreed to organize team
Define overall goals for GENI GUIs

Types of GUIs:
Control experiments
Display I&M results
Report status
View archive service

Consider portal, for access to multiple GUIs

Consider need for authentication and authroization

6.9 GENI measurement data schema

Team members:

Bruce Maggs – Duke University and Akamai (yes)
Max Ott – NICTA (yes, by phone)
Ivan Seskar – Rutgers WINLAB (yes)
* Martin Swany - Univ Delaware (yes)
Camilo Viecco - Indiana Univ (yes)
Michael Zink - UMass Amherst (yes)
Jim French - CNRI (yes)
* agreed to organize first discussion and writing

Consider:

Measurement data schema
Metadata schema
Metadata contents

Consider measurement data schema and/or metadata schema from:

perfSONAR
GMOC-provided
Current OML
Proposed using IPFIX
NetCDF (as used by DI Cloud)

Consider: Minimum set required for GENI

Provide overall template for GENI metadata, considering above.
Which items in GENI metadata template are:

Required?
Invariant?

7. Review Consensus and Draft Roadmap

Agenda for June 9:

1:00 pm	Review consensus of GENI I&M use cases; GENI I&M measurement plane, services, interfaces and protocols (APIs); and contents and structure of GENI measurement data schema; and draft roadmap for how key projects could implement them in Spirals 2 and 3
2:00 pm	Adjourn

7.1 Review consensus of GENI I&M use cases; GENI I&M measurement plane, services, interfaces and protocols (APIs); and contents and structure of GENI measurement data schema

How close are we to consensus on each of these priority topics:

1) GENI I&M use cases

2) GENI I&M services

3) GENI I&M measurement plane and interfaces

4) GENI I&M interfaces and protocols (APIs)

5) GENI measurement data schema

7.2 Draft roadmap for how key projects could implement them in Spirals 2 and 3

How can each of the following projects move towards a standard GENI approach? With how much effort?

OML (ORBIT Measurement Library) OMF (ORBIT Management Framework)
Max Ott – NICTA (yes, by phone)
Ivan Seskar – Rutgers WINLAB (yes)
Instrumentation Tools
Jim Griffioen - Univ Kentucky (yes)
perfSONAR

Matt Zekauskas - Internet2 (no)
Jason Zurawski – Internet2 (yes)
Martin Swany - Univ Delaware (yes)
Guilherme Fernandes – Univ Delaware (yes)
Ezra Kissel – Univ Delaware (yes)

Scalable Sensing Service (S3)

Sonia Fahmy – Purdue (yes)
Puneet Sharma - HP Labs (yes)

OnTimeMeasure for network measurements

Prasad Calyam - Ohio Supercomputing Ctr (yes)

GENI Meta-Operations Center and NetKArma

Jon-Paul Herron - Indiana Univ
Camilo Viecco - Indiana Univ (yes)
Chris Small - Indiana Univ (yes)

Virtual Machine Introspection (VMI)

Brian Hay – Univ Alaska (yes)

Data-Intensive Cloud Control for GENI

Michael Zink (yes)

Experiment Management Service – Digital Object Registry

Jim French - CNRI (yes)
Giridhar Manepalli - CNRI (yes)

8. References

Ref GIMS_Design_UseCases: "Use-cases for GENI Instrumentation and Measurement Architecture Design"

Ref MeasPlane-1: "RESTful Web Services vs. "Big" Web Services: Making the Right Architectural Decision"

Ref OMF_OML-1: "XDR: External Data Representation Standard"

Ref OMF_OML-2: "ORBIT Measurements Framework and Library (OML): Motivations, Design, Implementation, and Features"

Ref OMF-OML-3: "OML Overview" slides

Ref OMF-OML-4: "Measurement Architectures for Network Experiments with Disconnected Mobile Nodes"

Ref InsTools-1: "Architectural Design and Specification of the INSTOOLS Measurement System"

Ref perfSONAR-1: "Scalable Framework for Representation and Exchange of Network Measurements"

Ref perfSONAR-2: "An Extensible Schema for Network Measurement and Performance Data"

Ref perfSONAR-3: "NM-WG/perfSONAR Topology Schema"

Ref GMOC-1: "GMOC Topology-Entity Data Exchange Format Specification"

Ref GMOC-2: "Proposal: Use of URN's as GENI Identifiers"