GENI Instrumentation and Measurement WG: GENI I&M Architecture

GENI Engineering Conference 7
Duke University, Durham, NC

GPO System Engineer: Harry Mussman

March 17, 2010

www.geni.net

Sponsored by the National Science Foundation

Instrumentation and Measurement WG

- Scope:
  - Discuss, develop and build consensus around the architectural framework for the instrumentation and measurement infrastructure that will be deployed and used in GENI
  - Deploy basic instrumentation and measurement capabilities in GENI Spiral 2

GENI I&M Architecture: Document

- Purpose:
  - Provide a comprehensive and ordered list of topics that must be addressed for a complete architecture
  - Identify the priority topics that the WG needs to address first
  - Pull together contributions by the WG though Spiral 2

- Plan:
  - Now : v0.1 DRAFT completed, by GPO; see http://groups.geni.net/geni/wiki/GeniInstrumentationandMeasurementsArchitecture
  - By GEC8: v0.5 draft, by GPO, with contributions from WG
  - By GEC9: v1.0 draft, reviewed by WG

GENI I&M Architecture: Document Outline

- 1. Document Scope
- 2. Introduction
- 3. Definition and configuration of I&M services
- 4. Interfaces, protocols and schema for Measurement Data (MD)
- 5. Ownership of MD and privacy of owners
- 6. Interfaces, protocols and APIs for using I&M services
- 7. Basic GENI I&M use cases
- 8. MD transport via the GENI Measurement Plane
- 9. Discovery, authorization, assignment and binding of GENI I&M services
- 10. Measurement Orchestration (MO) service
- 11. Measurement Point (MP)
- 12. Time-stamping MD
- 13. Measurement Collection (MC) service
- 14. Measurement Analysis and Presentation (MAP) service
- 15. Measurement Data Archive (MDA) service
- 16. Additional GENI I&M use cases

GENI I&M Architecture: I&M Services

- 1. Document Scope
- 2. Introduction
- 3. Definition and configuration of I&M services
- 4. Interfaces, protocols and schema for Measurement Data (MD)
- 5. Ownership of MD and privacy of owners
- 6. Interfaces, protocols and APIs for using I&M services
- 7. Basic GENI I&M use cases
- 8. MD transport via the GENI Measurement Plane
- 9. Discovery, authorization, assignment and binding of GENI I&M services
- 10. Measurement Orchestration (MO) service
- 11. Measurement Point (MP)
- 12. Time-stamping MD
- 13. Measurement Collection (MC) service
- 14. Measurement Analysis and Presentation (MAP) service
- 15. Measurement Data Archive (MDA) service
- 16. Additional GENI I&M use cases

GENI I&M Architecture: Priority Topics

- 1. Document Scope
- 2. Introduction
- **3. Definition and configuration of I&M services**
- **4. Interfaces, protocols and schema for Measurement Data (MD)**
- 5. Ownership of MD and privacy of owners
- 6. Interfaces, protocols and APIs for using I&M services
- 7. Basic GENI I&M use cases
- **8. MD transport via the GENI Measurement Plane**
- 9. Discovery, authorization, assignment and binding of GENI I&M services
- 10. Measurement Orchestration (MO) service
- 11. Measurement Point (MP)
- 12. Time-stamping MD
- 13. Measurement Collection (MC) service
- 14. Measurement Analysis and Presentation (MAP) service
- 15. Measurement Data Archive (MDA) service
- 16. Additional GENI I&M use cases

3. Configuration of I&M Services: Issues

- Issues:
  - What is a complete group of I&M services?
  - How are they configured in GENI implementations, from small to large?

3. Configuration of I&M Services: Approach

- Based on GENI I&M Capabilities Catalog (v0.1), these GENI projects have comprehensive, end-to-end capabilities:
  - OML (ORBIT Measure Library) in OMF (ORBIT Mgmt Framework)
    (Ott, NICTA and Gruteser, WINLAB/Rutgers, 1660)
  - Instrumentation Tools
    (Griffioen, Univ Kentucky, 1642)
  - perfSONAR for network measurements
    (Zekauskas, I2 and Swany, Univ Delaware, 1788)
  - Scalable Sensing Service
    (Fahmy, Purdue and Sharma, HP Labs, 1723)
  - OnTimeMeasure
    (Calyam, Ohio Super Ctr, 1764)
  - others?

- What measurement services do these projects have in common?

OML in OMF (Ott and Gruteser, 1660): Architecture

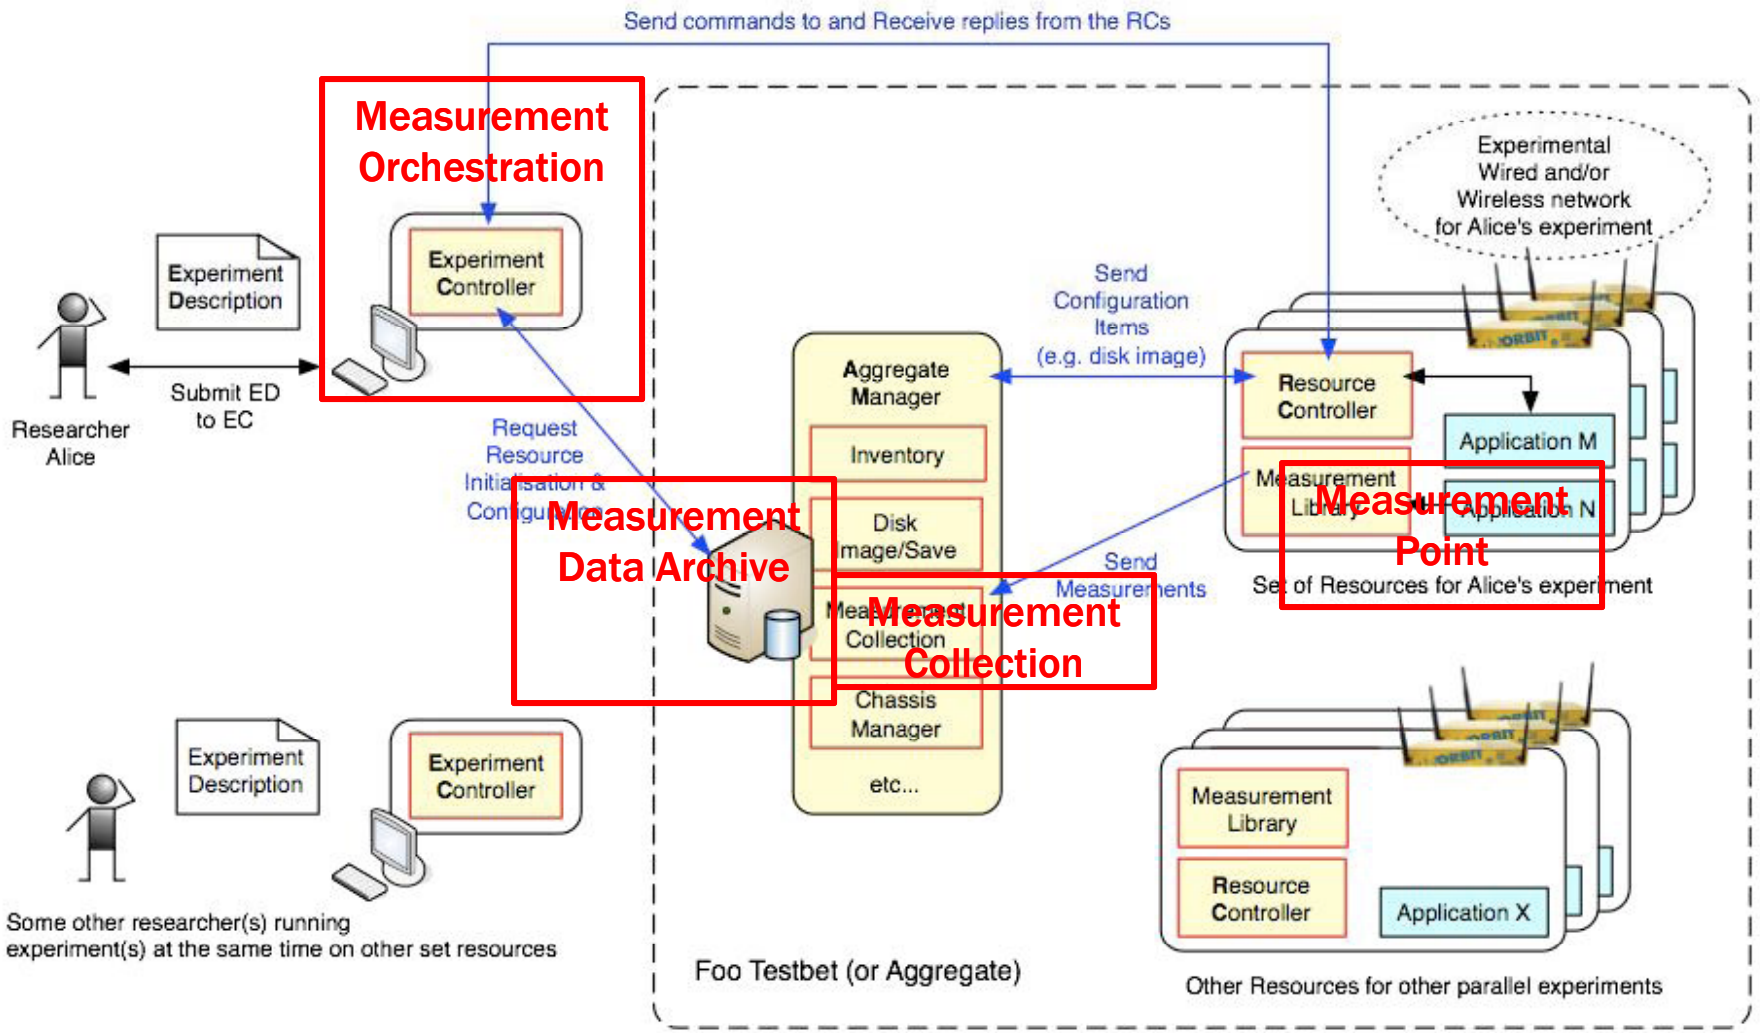

Instrumentation Tools (Griffioen, 1642): Architecture

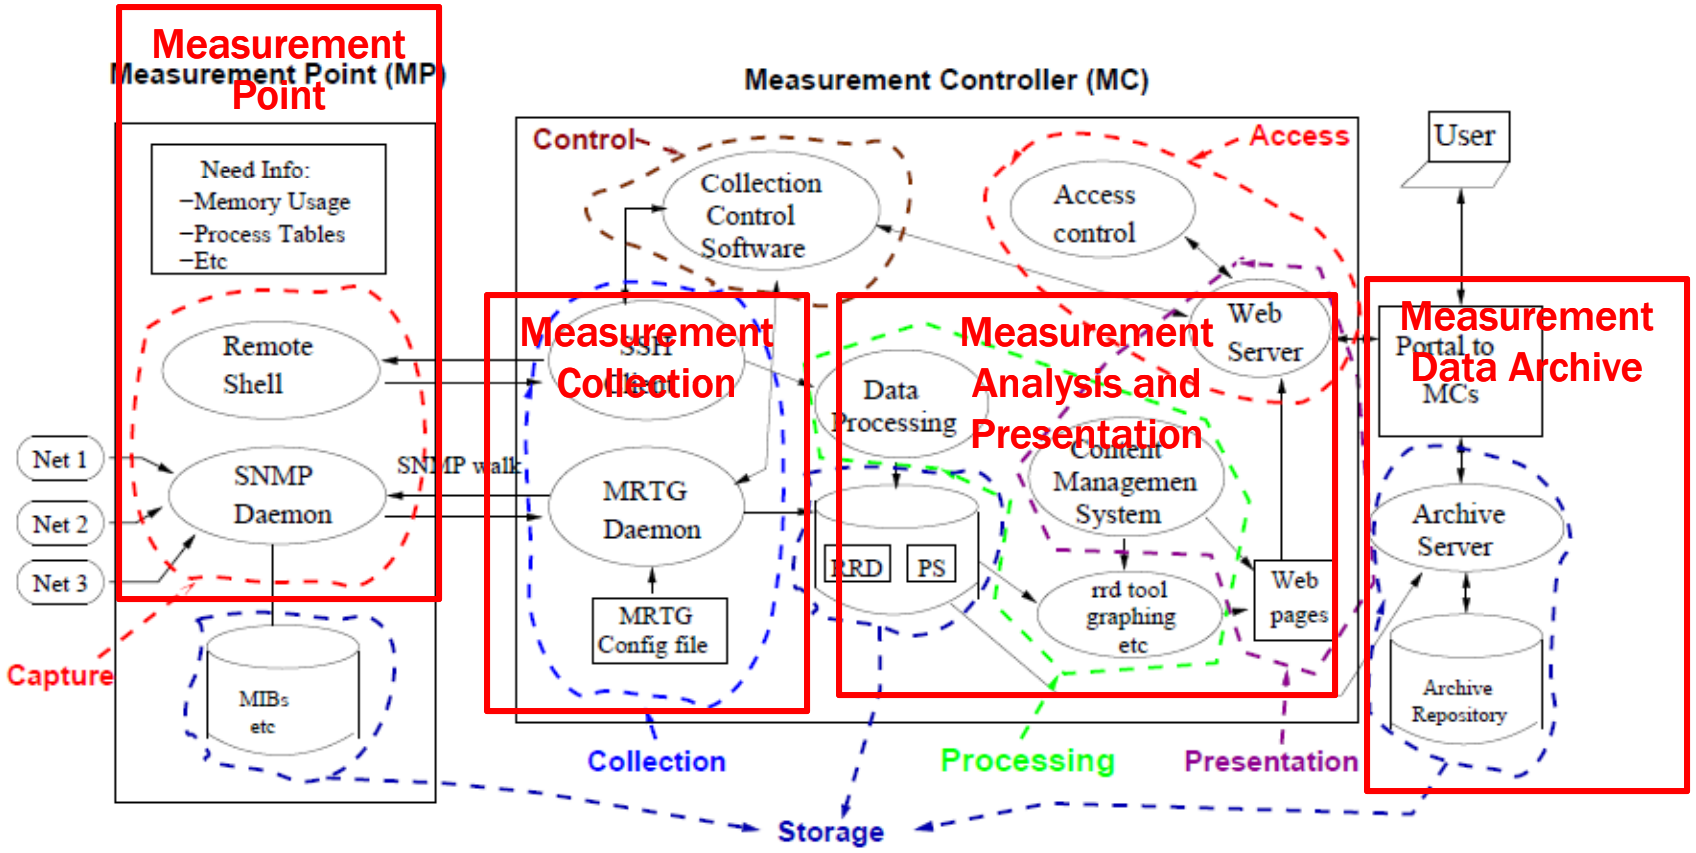

Figure 3: The Architectural Components of the INSTOOLS Toolset

perfSONAR Architecture

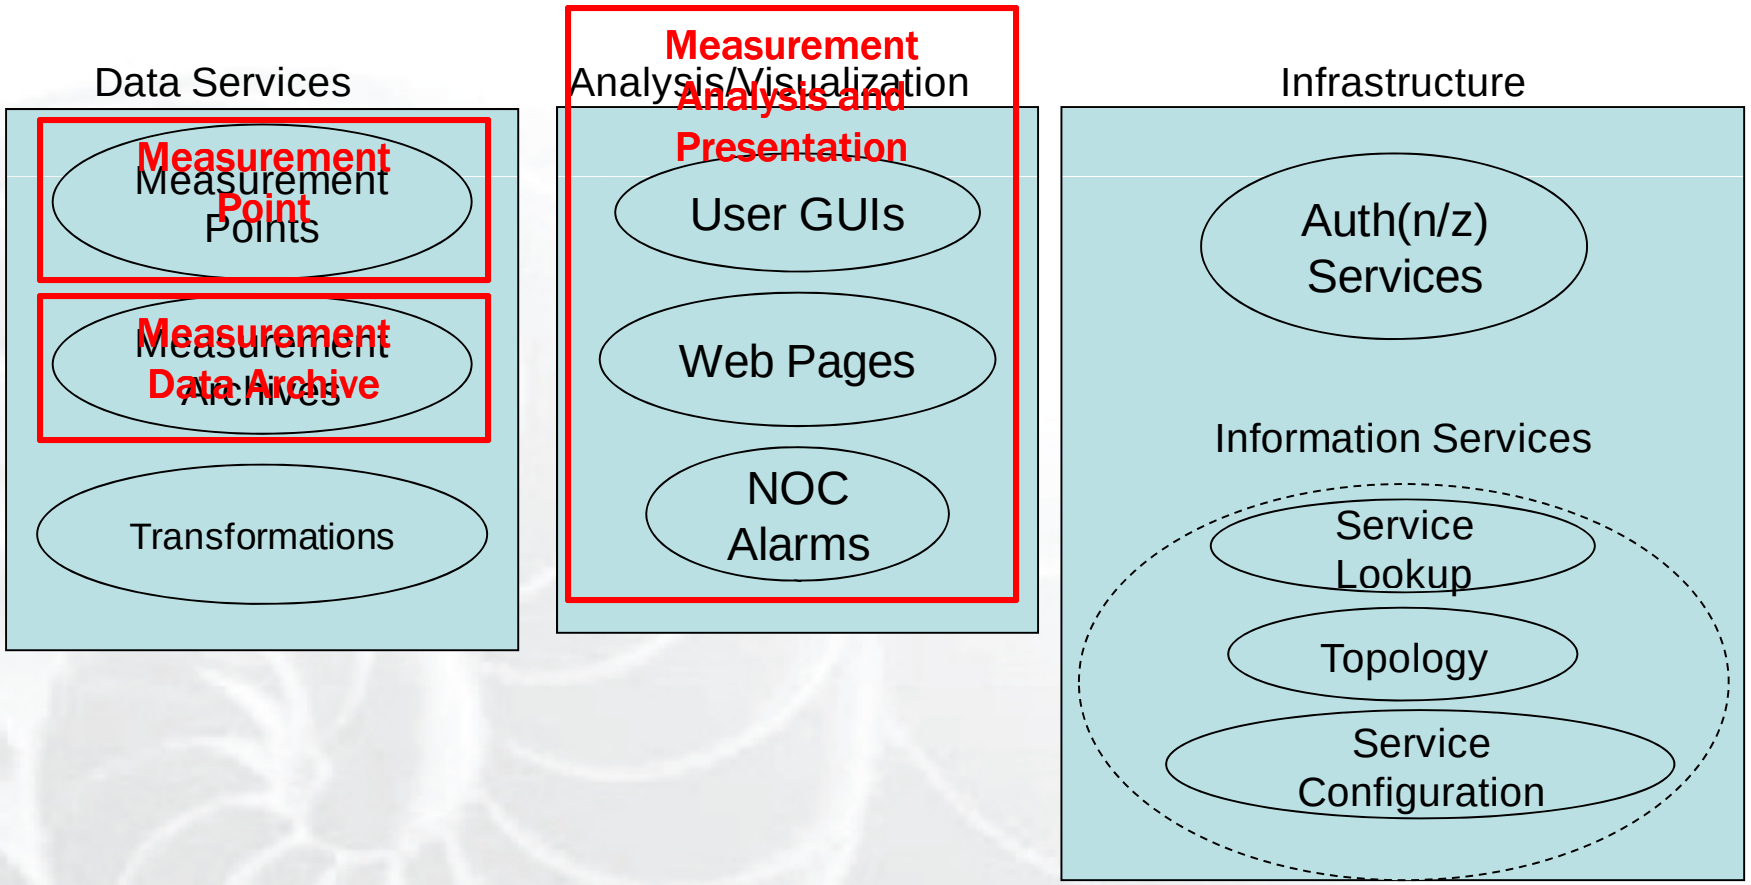

Scalable Sensing Service (Fahmy and Sharma, 1723): Architecture

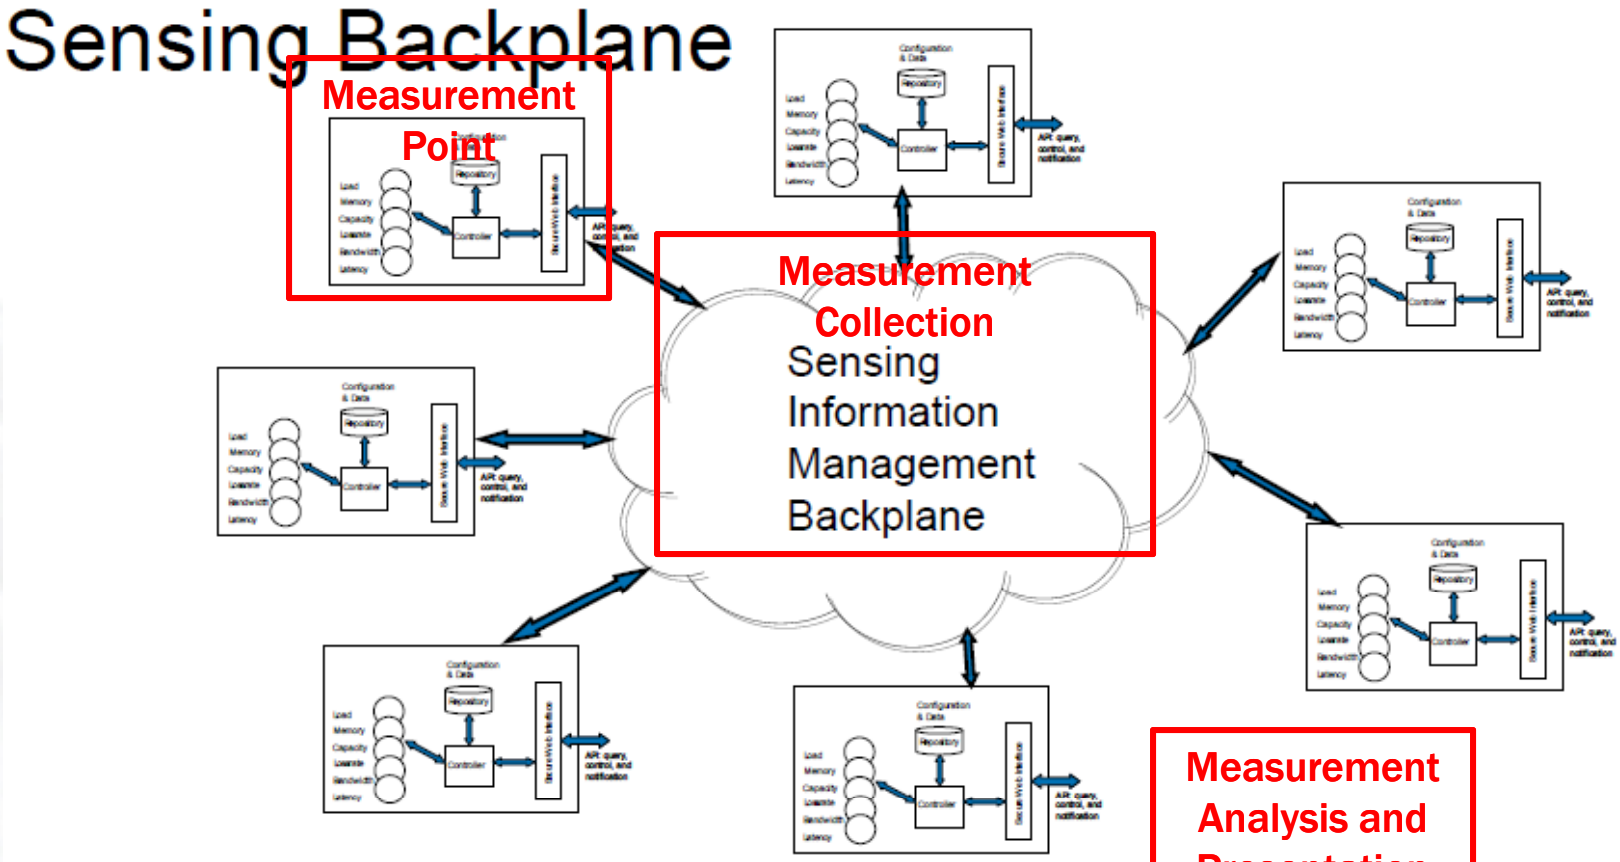

- Aggregate data from end-points
- Configurable and self-managing
- E.g., SDIMS [SIGCOMM 2004]

OnTimeMeasure (Calyam, 1764): Architecture

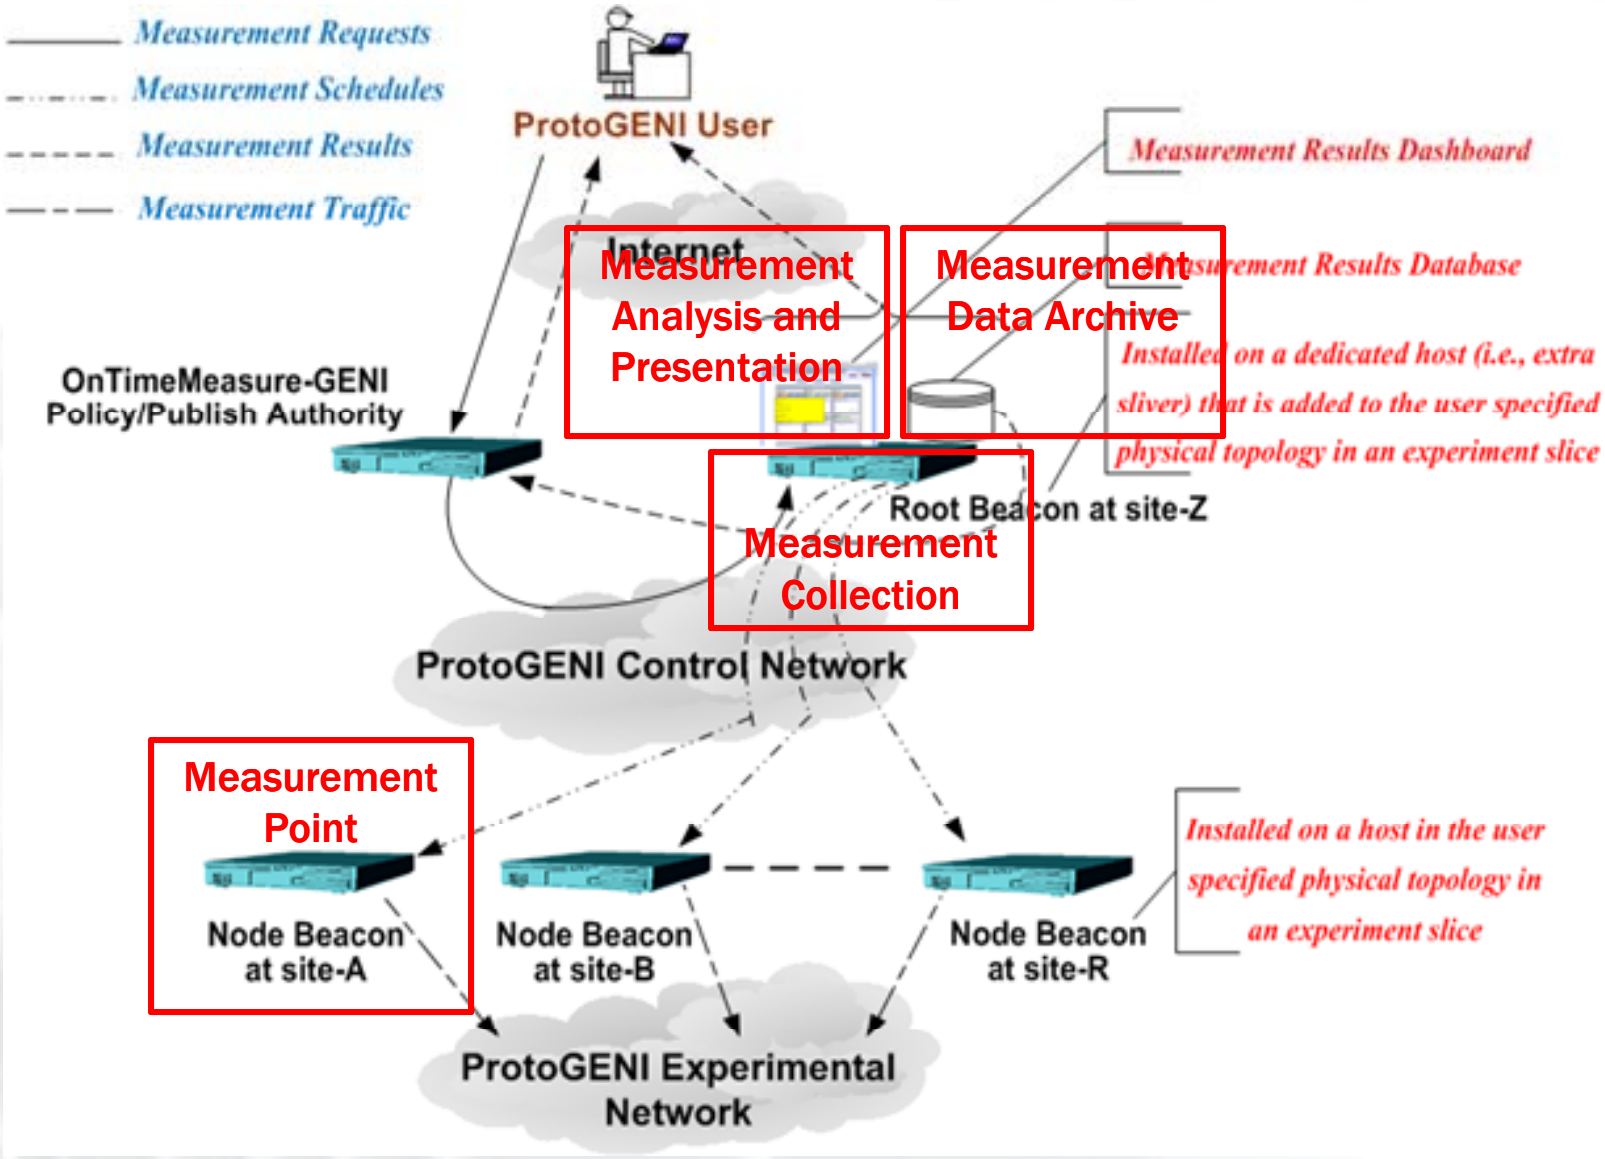

3. Configuration of I&M services: Proposed definition

- After considering projects with comprehensive, end-to-end capabilities, here are five services they have in common:

- Measurement Orchestration (MO) service
  (p/o Experiment Control service, uses a language to orchestrate I&M services)
- Measurement Point (MP) service
  (instrumentation that taps into a network and/or systems, links and/or nodes, to capture measurement data and format it using a standardized schema)
- Measurement Collection (MC) service
  (programmable systems that collect, combine, transform and cache measurement data)
- Measurement Analysis and Presentation (MAP) service
  (programmable systems that analyze and then present measurement data)
- Measurement Data Archive (MDA) service
  (measurement data repository, index and portal)

3. Configuration of I&M services: Proposed configuration

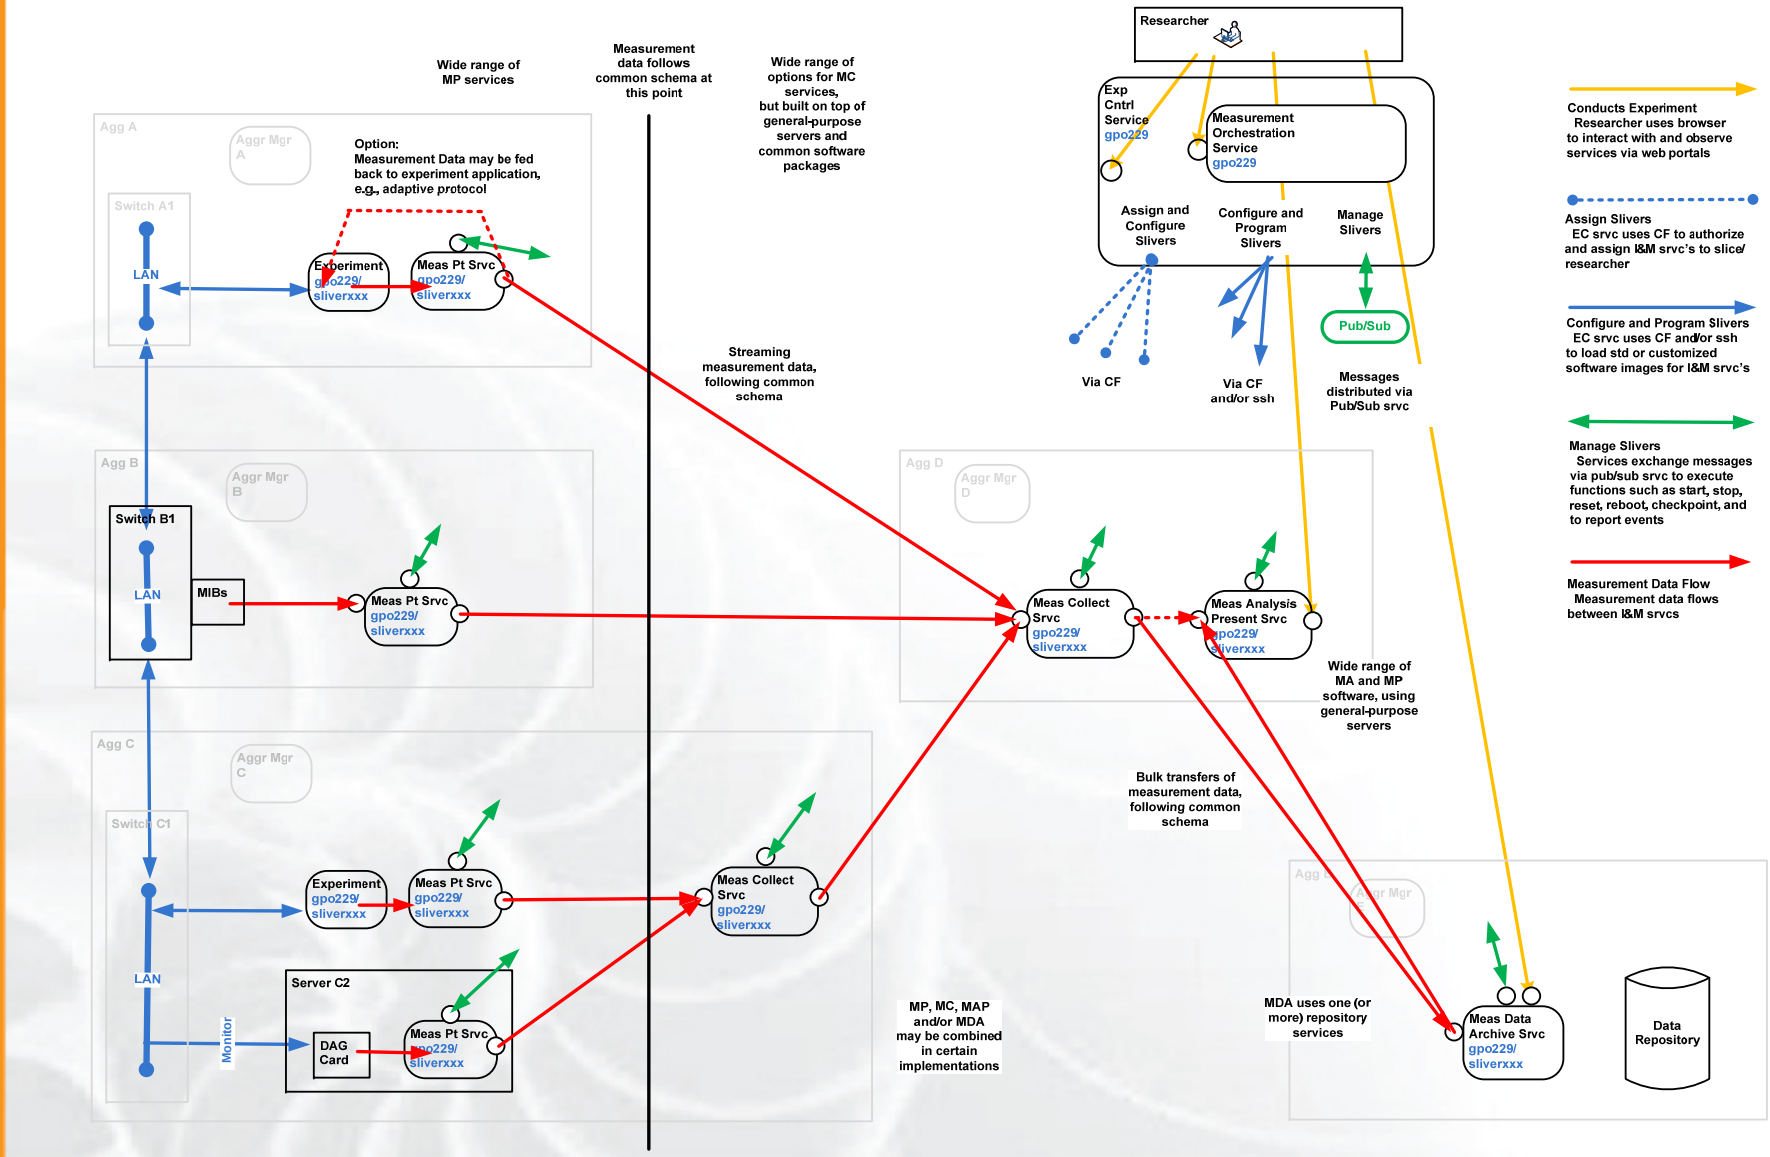

3. Configuration of I&M services: Expected range of implementations

- Small-scale implementations might put all I&M services within one aggregate, and even in one server
  - interfaces between services would be internal to the aggregate, or even internal to the server
- Large-scale implementations might have I&M services distributed over many aggregates
  - with measurement data flowing between services
  - with orchestration mechanisms based upon message exchanges

3. Configuration of I&M services: WG discussion topics

- Are these five services a complete group of I&M services?
- Are these good names for the five I&M services?
- Is this five the right granularity for I&M services?
- Is this a complete and flexible configuration for I&M services?
- Can this configuration accommodate the range from small-scale to large-scale implementations?
- How can we obtain a consensus, so that we can set a firm foundation for the other topics?

GENI I&M Architecture: Priority Topics

- 1. Document Scope
- 2. Introduction
- 3. Definition and configuration of I&M services
- **4. Interfaces, protocols and schema for Measurement Data (MD)**
- 5. Ownership of MD and privacy of owners
- 6. Interfaces, protocols and APIs for using I&M services
- 7. Basic GENI I&M use cases
- 8. MD transport via the GENI Measurement Plane
- 9. Discovery, authorization, assignment and binding of GENI I&M services
- 10. Measurement Orchestration (MO) service
- 11. Measurement Point (MP)
- 12. Time-stamping MD
- 13. Measurement Collection (MC) service
- 14. Measurement Analysis and Presentation (MAP) service
- 15. Measurement Data Archive (MDA) service
- 16. Additional GENI I&M use cases

4. Interfaces, protocols and schema for MD: Issues and Approach

- Issues:
  - This topic suggested at GEC6 meeting: Common schema for MD
  - Can we identify a common set of interfaces, protocols and schema for MD, or at least a limited number of types?
  - What needs to be included in the MD schema?

- Approach:
  - Assume all MD after MPs follows this common set of interfaces, protocols and schema
  - Start with definition of MD schema
  - Next, understand [8. MD Transport via GENI Measurement Plane]
  - Then, complete first set of interfaces and protocols

4. Interfaces, protocols and schema for MD: WG discussion topics

- From GENI I&M Capabilities Catalog (v0.1), these GENI projects (and others) are working on data schema and/or data archives:
  - perfSONAR for network measurements (Swany, Univ Delaware, 1788)
  - IMF project (Dutta, NC State, 1718)
  - Embedded Real-Time Measurements (Bergman, Columbia, 1631)
  - GENI Meta-Operations Center (Herron, Indiana Univ, 1604)
  - netKarma: GENI Provenance Registry (Pale and Small, Indiana Univ, 1706)
  - DatCat project at http://www.datcat.org/ (Klaffy, CAIDA)
  - Crawdad project at http://crawdad.cs.dartmouth.edu/ (Kotz, Dartmouth)
  - Amazon Simple Storage Service
  - Data-Intensive Cloud Control (Zink and Cecchet, UMass Amherst, 1709)
  - Experiment Mgmt System (Lannom and Manepalli, CNRI, 1663)
  - others?
- What can we learn from these projects?

4. Interfaces, protocols and schema for MD: WG discussion topics

- Standardized interfaces between measurement services
  - Pt-to-pt vs pt-to-multipoint (e.g., pub/sub)
  - Stream vs bulk transfer
  - Disconnection operation expected, or not.
- Protocols for moving measurement data
  - Streaming data
  - Bulk-transfer of data
- Schema for measurement data
  - Data record identifier
  - Annotation, or meta data
  - Data types and values, with timestamps
- How can we obtain a consensus on first set of intfc's/protocols/schema for MD?
- What is the process for extending the set?

GENI I&M Architecture: Priority Topics

1. Document Scope
2. Introduction
3. Definition and configuration of I&M services
4. Interfaces, protocols and schema for Measurement Data (MD)
5. Ownership of MD and privacy of owners
6. Interfaces, protocols and APIs for using I&M services
7. Basic GENI I&M use cases
8. **MD transport via the GENI Measurement Plane**
9. Discovery, authorization, assignment and binding of GENI I&M services
10. Measurement Orchestration (MO) service
11. Measurement Point (MP)
12. Time-stamping MD
13. Measurement Collection (MC) service
14. Measurement Analysis and Presentation (MAP) service
15. Measurement Data Archive (MDA) service
16. Additional GENI I&M use cases

8. MD Transport via the GENI Measurement Plane: Issues and Approach

- Issue:
  - Need to understand how MD traffic flows are transported by the GENI Measurement Plane before the interfaces and protocols for MD can be fully defined

- Approach:
  - Understand current view of GENI Control Plane and Experiment Plane
  - Consider options for GENI Measurement Plane to transport MD flows, using networks that implement GENI Control and Experiment Planes

- Current view of GENI Control Plane:
  - Implemented using an IP network, carried via a GENI backbone implemented on I2 and/or NLR, with public IP addresses, that should be reachable from the Internet.
- Current view of GENI Experiment Plane:
  - Implemented using a switched Layer 2 network, with Layer 2 connections switched within and between aggregates, and carried via a Layer 2 GENI backbone implemented on I2 and/or NLR.

8. MD Transport via the GENI Measurement Plane: WG discussion topics

- Assume: MD traffic flows are transported between I&M services using the GENI Measurement Plane
  - Expected to include both incremental blocks of data, and occasional large file transfers

- Option 1: Carry all MD traffic flows using a dedicated measurement VLAN
  - As in OMF, where the dedicated measurement VLAN is called "Control"
  - This is ideal, but requires nodes/servers to have three NICs, a significant complication
- Option 2: Carry all MD traffic flows using the same IP network that supports the Control Plane.
  - How can the MD traffic flows be controlled to avoid disrupting the control traffic?
  - This would be consistent with servers having two NICs
- Option 3: Carry most MD traffic flows per Option 2, but define a dedicated measurement VLAN to the slice/experiment for high-rate MD traffic flows, using the same Layer 2 network that supports the Experiment Plane.
  - Perhaps this needs to be done in only exceptional cases
  - This could still be consistent with servers having two NICs

8. MD Transport via the GENI Measurement Plane: WG discussion topics

- How can the WG reach consensus?

- How can there be a consensus with all of GENI on Measurement Plane, and relationship with Control Plane and Experiment Plane?

GENI I&M Architecture: Next Steps

- Purpose:
  - Provide a comprehensive and ordered list of topics that must be addressed for a complete architecture
  - Identify the priority topics that the WG needs to address first
  - **Pull together contributions by WG though Spiral 2**

- Plan:
  - Now : v0.1 DRAFT completed, by GPO; see http://groups.geni.net/geni/wiki/GeniInstrumentationandMeasurementsArchitecture
  - **By GEC8: v0.5 draft, by GPO, with contributions from WG**
  - By GEC9: v1.0 draft, reviewed by WG

GENI I&M Architecture: Priority Topics and Proposed I&M Services

- 1. Document Scope
- 2. Introduction
- **3. Configuration of I&M services (includes five services)**
- **4. Interfaces, protocols and schema for Measurement Data (MD)**
- 5. Ownership of MD and privacy of owners
- 6. Interfaces, protocols and APIs for using I&M services
- 7. Basic GENI I&M use cases
- **8. MD transport via the GENI Measurement Plane**
- 9. Discovery, authorization, assignment and binding of GENI I&M services
- **10. Measurement Orchestration (MO) service**
- **11. Measurement Point (MP)**
- 12. Time-stamping MD
- **13. Measurement Collection (MC) service**
- **14. Measurement Analysis and Presentation (MAP) service**
- **15. Measurement Data Archive (MDA) service**
- 16. Additional GENI I&M use cases

GENI I&M Architecture: Proposed Configuration