Breakout Summary: Optical, Measurement, and Real-time

GENI Planning Group Resource: Larry Peterson

NSF Resource: Darleen Fisher

Chair: Rene Cruz

July 13: 15 people present

July 14: 10 people present

“Detailed version”

Optical Networks

- Two levels or research
  - Physical layer, hardware oriented
  - Control plane for fast circuit switching
- Former may be less viable (cost) for GENI, but is important to support with research dollars, as well as future integration witn GENI
  - federation with optical testbeds suggested
    - E.g. MONET, BOSSNET, and CHEETAH, Dragon (GMPLS testbed), Enlighten
    - Applies to other technology testbeds (non-optical, e.g. wireless)

Optical Networking Control Plane Research

- Control of rapidly configurable circuits
  - Could be lambdas or sub-lambda (TDM)
  - Intersects heavily with GENI facility design
    - planning group is encouraged to seek input from research community, and provide flexibility for, slice allocation, topology embedding algorithms
  - Intersects heavily with control plane experiments on GENI facility
    - Within a “slice”
- Important to provide capability for dynamic provisioning of circuits which terminate at ***edge*** of GENI
  - Experiments with networking architectures where bulk of switching happens at network edge
  - Support for optical access networks

Measurement, Summary

- GENI can be viewed as an embedded sensor network
  - network events are what is to be sensed
  - cross layer measurement support is critical
  - Collect, measure, aggregate, analyze, archive, replay
- Complexity and cost of measurement could easily dominate cost of GENI
  - Flexibility, slicing of measurement infrastructure is crucial
  - Heisenberg (Bandiwdth and processing requirements for supporting measurements and dissemination compete with network resources)
  - “Tap” versus embedded instrumentation
- Need to pay attention to privacy issues

Measurement, detailed (1)

- In a sense, the true scientific grand challenge that GENI addresses is to understand at a fundamental level the behavior of the most complex artifact known to humankind -- global-scale globe-spanning networked systems. It is fair to say that there are many aspects of the detailed behavior of the Internet that are simply not understood today, let alone how it will behave in the near future as the number and diversity of nodes increases by several orders of magnitude and ever more critical applications migrate to it over the next few years. As this represents a major societal risk, Congress should be concerned!

- Therefore, it is very important to consider the design and deployment of an "observatory" capability, commensurate with a research facility of the size and scale of GENI, to provide the foundation for a more fundamental understanding of the behavior of such complex engineered systems.

- Essentially this is an embedded sensor network where that being sensed are network events themselves. This is intended as a more general concept than that connoted by the term "measurement":
  - Collect, measure, aggregate, analyze, archive, replay;

Measurement, detailed (2)

- An aggressive data collection capability is viewed as essential for any research facility in the physical sciences, and these often stretch the state-of-the-engineering-art in terms of bandwidth, capacity, and data storage/archive abilities (e.g., large hadron collider, nuclear initiation facility, etc.).
  - (Extreme) technical challenge for GENI observatory: collect full packet traces (or equivalent network level fragment) at 10 Gbps or faster at O(1000) places in the network simultaneously.
- In reality, the data collection capability will support a design space from collecting (1) just headers to full packets, (2) every packet to a variety of statistical reductions, (3) from many simultaneous places to just a few, (4) for one slice to every slice (whole network), (5) retain for a short time to months, and (6) from network level data up through the protocol stack to application events. The collection capability will have some maximum limitations in terms of collection bandwidth and capacity, and this will need to be configured and allocated to experiments just like any other resource in the systems.
- The volume of data collection cannot be collected and analyzed in one place. The ability to slice the analyses by geography, and to combine statistical abstracts or other methods of aggregation of results is likely to be a significant research challenge. Therefore, it is important that such research be able to be enabled by the native capability provided by the GENI facility architecture.

Measurement, detailed (3)

- Where and how will the measurements be taken?
  - Taps via specialized hardware at various network pipes?
  - Assume this is a capability of the underlying routers?
  - Implemented by specialized blades with direct connectivity to disk arrays?
  - Other approach?

- Must also consider:
  - (1) Support for cross-layer instrumentation, tracing, and measurement
  - (2) Privacy policies when user data is collected and archived.
  - (3) How much of facility budget is to be dedicated to the observatory capability?
- The extreme technical challenge described above could dominate the cost of the whole facility!

Real-time applications

- focus on latency control, rather than high throughput
  - End-to-end
  - “performance is bad, who is to blame?”
- high fidelity and easy to use measurement infrastructure is critical
- need for programmable impairments
- need to define abstraction for slice requirements carefully in order to support real-time requirements
  - For slicing a link via TDM this is straightforward, but in general, where slice boundaries are more soft, this is quite a complex issue that deserves considerable thought and deliberation

Real-time applications (2)

- Focus on timing precision (more important than high BW)
  - Should be able to ask any GENI component - what time is it? And get answer with known, bounded precision
  - Need to effect actions at specific times (e.g sending a packet)
- Endpoints may not be PCs, support important new apps where endpoints may not have heavyweight OS
- Might be supported within a single slice
  - Important to have such a capability within GENI facility to support research of this type