

User Opt-in within GENI: *Broad Perspective*

Prasad Calyam, Ph.D.
Assistant Professor, Department of Computer Science
University of Missouri-Columbia

GEC23, June 2015

GEC23 User Opt-in Session Agenda

Session Goal:

- “This session is an opportunity for an exchange of information on user opt-in issues that include user recruitment, experiment design, technologies for user testbeds and application development. It will feature short talks and a panel discussion between GENI experimenters focused on application development and infrastructure operators”
- All speakers can give brief update on their project-specific issues/best practices, raise open questions and potential solution approaches; discussions will build on their talk content, and also moderators will engage audience for questions or use prepared questions

Session Agenda:

- 11am - 11:15am - User Opt-in within GENI: Broad Perspective (Prasad Calyam)

Part-1: Campus User Traffic Opt-in (Moderated by Prasad Calyam)

- 11:15 am - 11:35am - Short talks by Russ Clark and KC Wang
- 11:35am - 11:50am - Discussion following Short talks

Part-2: Experiment User Opt-in (Moderated by Russ Clark)

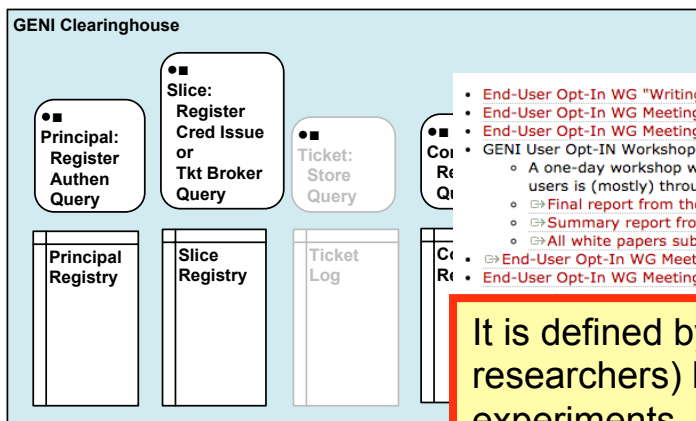
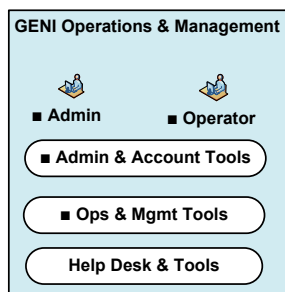
- 11:50am - 12:10pm - Short talks by Ezra Kissel and Parmesh Ramanathan
- 12:10pm - 12:30pm - Discussion following Short talks

Topics of Discussion

- User Opt-in Definition and Relevance
 - A bit of GENI history!
- Lessons from prior GENI projects
 - Opt-in Actors, Use Cases
- Issues from my current GENI/US Ignite projects
 - Simulation-as-a-Service
 - PhysicalTherapy-as-a-Service

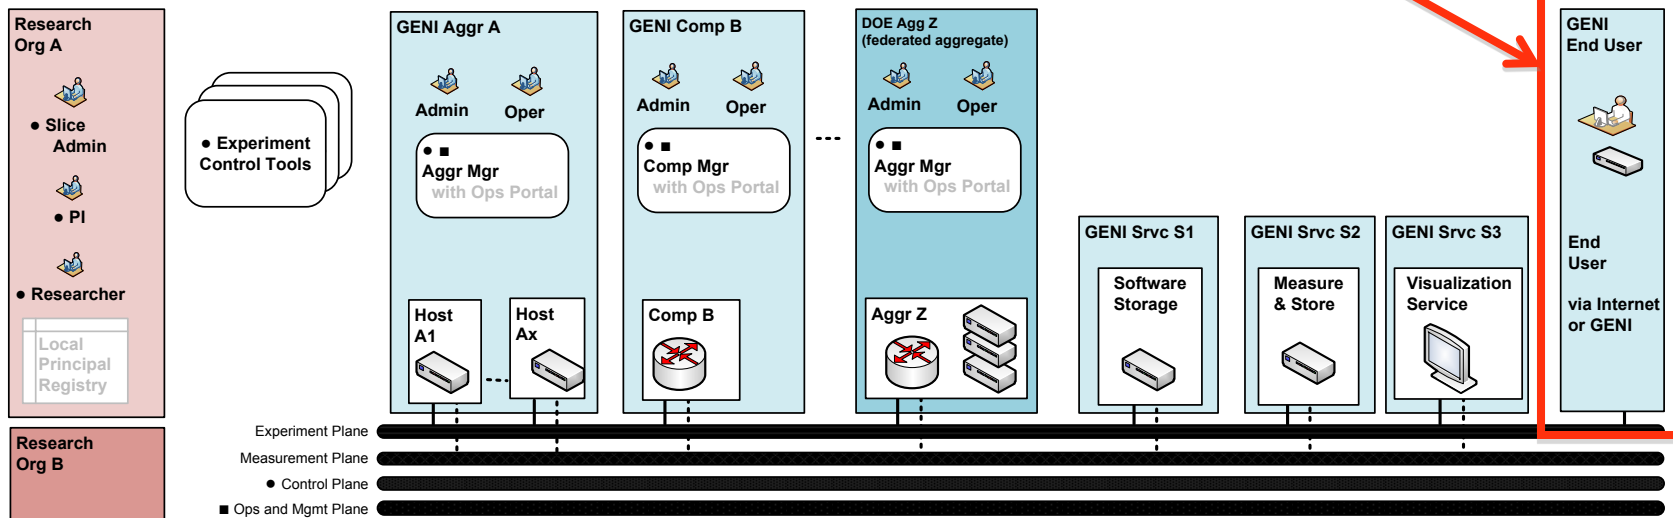
Opt-in Working Group Definition

End user in the Experiment Plane



- End-User Opt-In WG "Writing Meeting" in NYC: Agenda, summary and references (April 14, 2009)
- End-User Opt-In WG Meeting at the GEC-4: Meeting minutes, including all presentations (April 2, 2009)
- End-User Opt-In WG Meeting at the GEC-3: Meeting minutes, including all presentations (October 2008)
- GENI User Opt-IN Workshop (July 22, 2008)
 - A one-day workshop was held on July 22, 2008 in Cambridge, MA, that was structured to solicit ideas from users is (mostly) through getting great applications into GENI slices. Thus, we need to think about how to get great applications into GENI slices.
 - Final report from the GENI User Opt-In Workshop (July 22, 2008)
 - Summary report from the GENI User Opt-In Workshop (July 22, 2008)
 - All white papers submitted by the GENI User Opt-In Workshop participants.
- End-User Opt-In WG Meeting at the GEC-2: Presentation (March, 2008)
- End-User Opt-In WG Meeting at the GEC-1: Meeting minutes (October 2007)

It is defined by use cases where end users (not researchers) become involved with GENI experiments, and by the services and capabilities necessary to support them.



Credit: Harry Mussman & Henning Schulzrinne, 2007 - 2009 - <http://groups.geni.net/geni/wiki/GeniOptIn>

Opt-in Working Group Charter

Credit: Harry Mussman & Henning Schulzrinne, 2007 - 2009 - <http://groups.geni.net/geni/wiki/GeniOptIn>

- Charter Scope:
 - How do end-users (including Internet users) participate in GENI experiments?
 - What are the various aspects including user interfaces, scheduling, debugging, measurement, archiving data, sandboxes, etc.?
 - What are the privacy and legal issues involved in user opt-in?

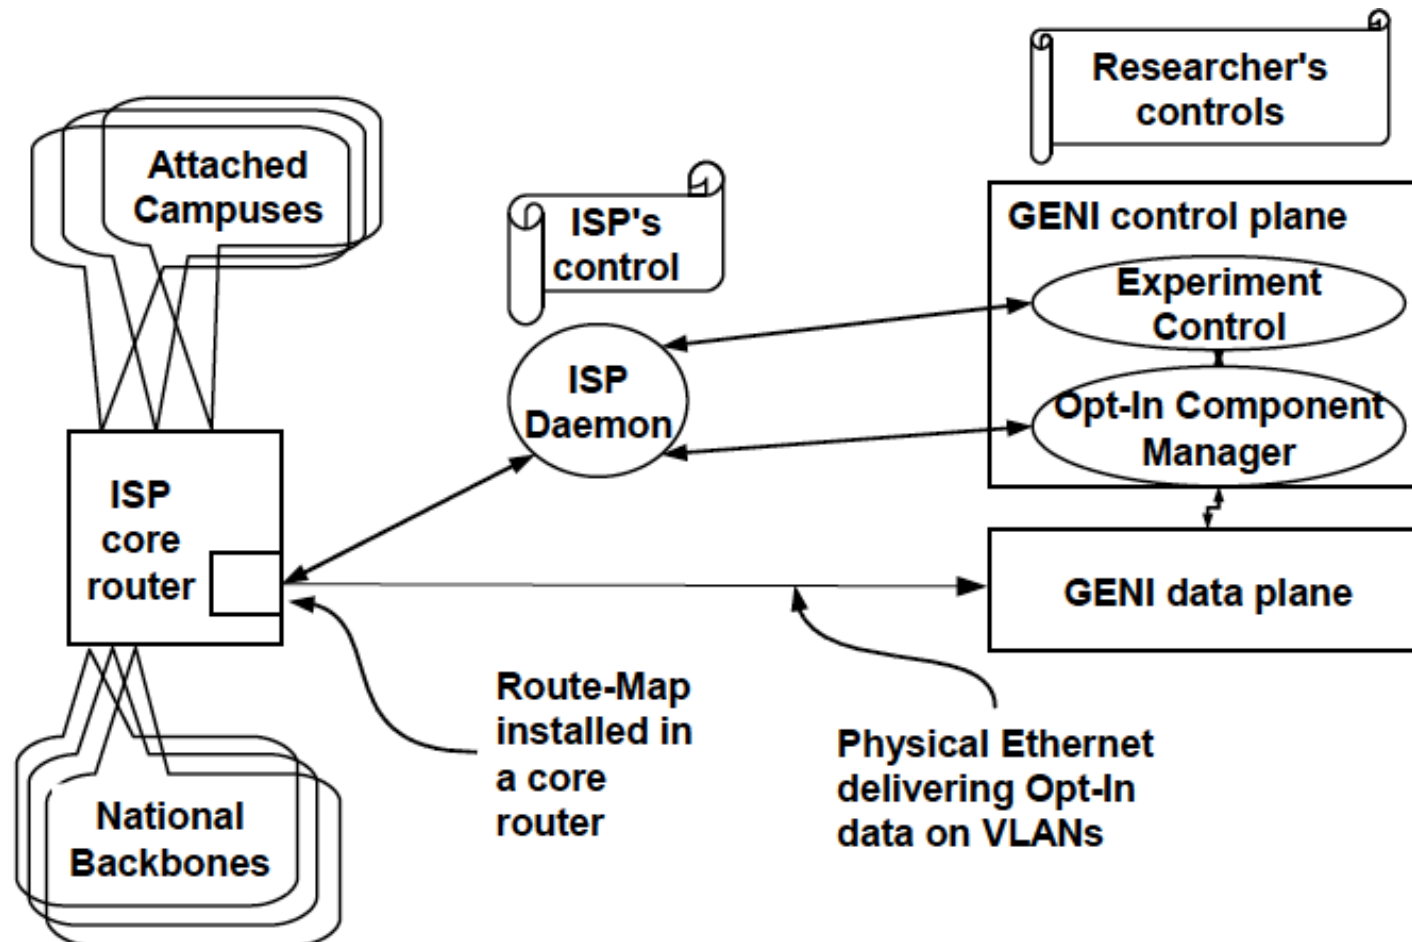
Why bother about user opt-in?

Credit: Harry Mussman & Henning Schulzrinne, 2007 - 2009 - <http://groups.geni.net/geni/wiki/GeniOptIn>

- Opt-in will add realism to the GENI experiments and make the experimental results relevant to real-world networks
 - Users bring e.g., mobility, compute/storage cycles, notion of scale to App
 - User opt-in is an essential step in building a large-scale long-running GENI experiment that has tangible user benefits
 - In turn, it helps launching a new network structure or service that could be commercially viable
- Opt-in user recruitment challenges
 - *bribe'em, woo'em, force'em (Henning in GEC2 Talk)*
 - Attract users through 'Value' offered to user; set pertinent expectations from user to participate and help the research
 - GENI provides - incentive \$, cheaper, faster, freedom/flexibility, reliable
- Issues:
 - What is "real" data or "realism" in user opt-in for an experiment?
 - How much opt-in do we need (number of users) for realism?

Regional Opt-in Project View

ISP-controller user traffic in the Experiment Plane



Credit: Matt Mathis, 2008 - <http://groups.geni.net/geni/wiki/RegionalOptIn>

Why re-visit GENI User Opt-in today?

- Earlier focus was GENI-project (larger experiment) Opt-in
 - What should GPO do to involve more users in prototyping?
 - Researchers and educators, industry, ...
 - How to incentivize good experimenters?
 - Focus on Apps and users will come
(GENI User Opt-In Workshop Report, July 2008)
 - Users may not even know GENI is underneath!
 - Users choose to opt into experiments of interest to them
 - Build a App marketplace; need active marketing to opt-in users
 - Training and support for App builders
 - Rely on employees/students as users
 - Similar to product beta testing (e.g., @ Microsoft, Google)
 - APIs and testbeds for App builders
 - Needs: resource brokering, performance measurements, security/privacy, ...

Why re-visit GENI User Opt-in today? (2)

- Today's focus has to be GENI-experiment user Opt-in
 - How to capitalize on GENI's success leveraging:
e.g., growing user community, infrastructure with GENI Racks, tools and documentation, GEC knowledge base, best practices?
 - (Still) Focus on Apps and users will come
 - GENI Shakedown Experiments, US Ignite Apps, Mozilla Ignite Apps, CC-* (Domain Science based) Campus Apps, ...
 - Residential broadband growth is in an interesting phase
 - Growing foot print of Google Fiber, DublinK, ...
 - Path more evident for experiment life cycle in GENI
 - Short-time experiment slices to long-standing services in slices
 - Graduated experiments have a path to cloud migration and even commercialization

Topics of Discussion

- User Opt-in Definition and Relevance
 - A bit of GENI history!
- Lessons from prior GENI projects
 - Opt-in Actors, Use Cases
- Issues from my current GENI/US Ignite projects
 - Simulation-as-a-Service
 - ElderCare-as-a-Service

User Opt-in Actors

- Operations group of CIO and staff
- Operations group of PI and staff
- Provider group
- Research group
 - including Slice Administrators, PIs and Researchers
- Experiment review board
 - (GENI-specific: suggested by PlanetLab experience, and important to screen opt-in experiments)
- Institutional Review Board
 - Institution/Site-specific procedures, legal matters around user risk from experiment participation

Source: <http://groups.geni.net/geni/wiki/GeniOptIn>

Exemplar prior GENI Projects on User Opt-in

- Seattle:
 - (Justin Cappos) <http://groups.geni.net/geni/wiki/MillionNodeGENI>
Educational uses, Research experiment demos
 - Issues: How to really build large-scale ‘million-node’ testbeds?
- Regional Opt-in:
 - (Matt Mathis) <http://groups.geni.net/geni/wiki/RegionalOptIn>
 - Proposal for wholesale opt-in via ISP Daemon
 - Permit GENI experimenters to request that ISPs redirect traffic from a huge population of innocent users to GENI infrastructure
 - ‘*Active Users*’ (actively choose to participate in an experiment) versus ‘*Innocent Users*’ (do nothing to participate in the experiment) -- many experiments may choose hybrid users
 - Issues: How to give ISP’s control? How to design means to opt-out?
- Virtual Tunnels
 - (Nick Feamster) <http://groups.geni.net/geni/wiki/DTunnels>
 - Bringing Experimenters and External Connectivity to GENI (BGP Mux)
 - Issues: How to seamlessly integrate external connectivity with virtual networks of experiments?

Topics of Discussion

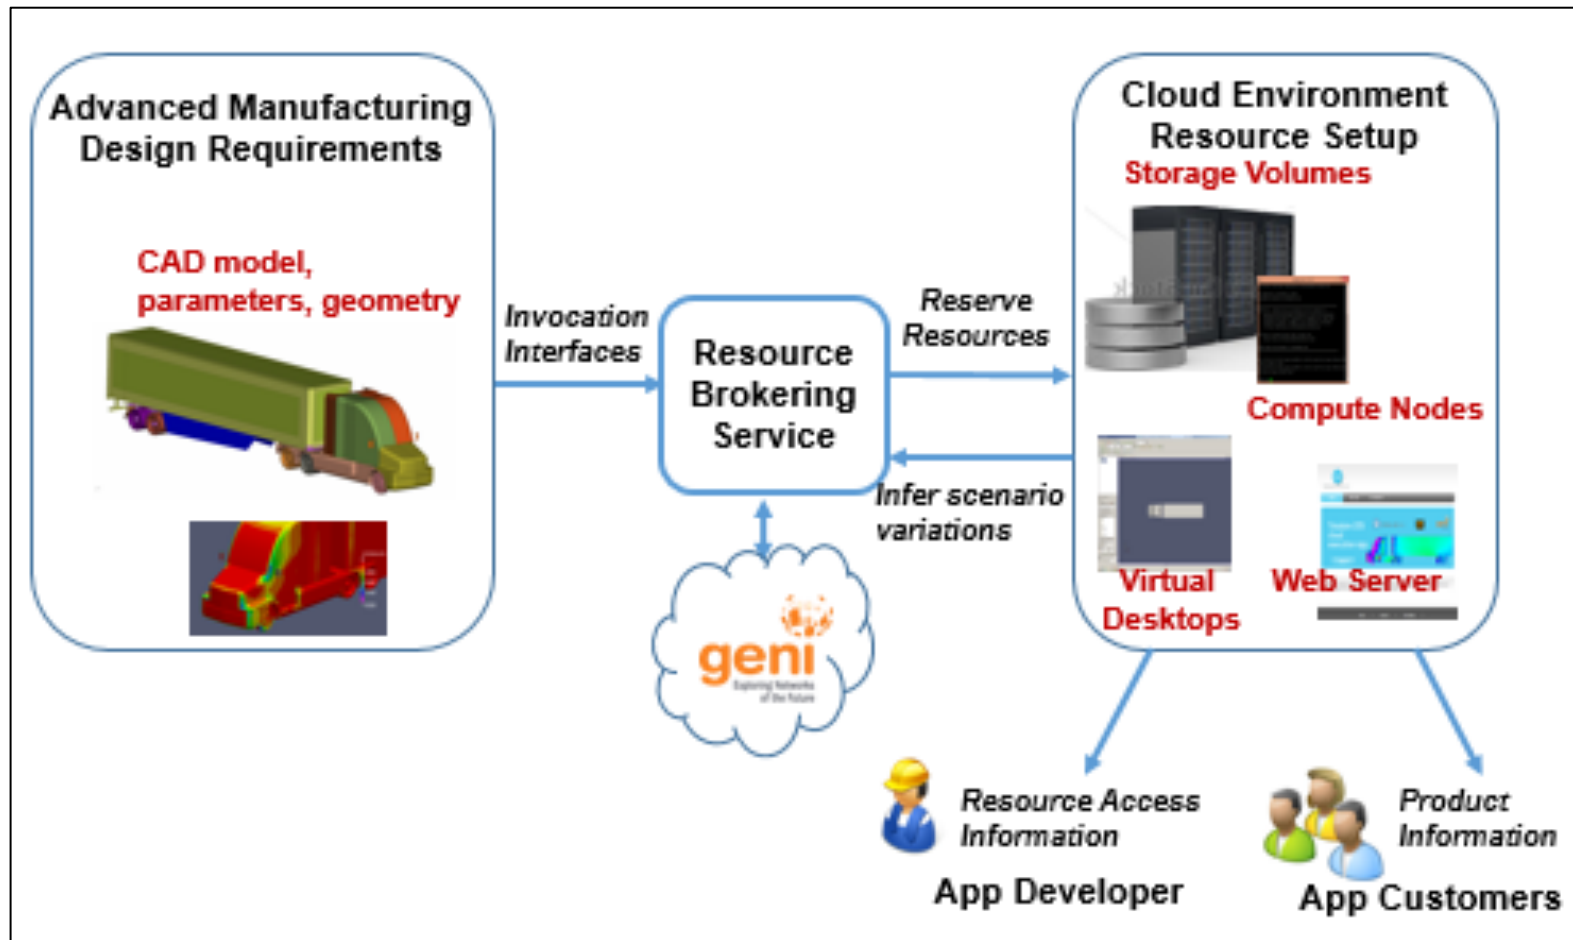
- User Opt-in Definition and Relevance
 - A bit of GENI history!
- Lessons from prior GENI projects
 - Opt-in Actors, Use Cases
- Issues from my current GENI/US Ignite projects
 - Simulation-as-a-Service
 - ElderCare-as-a-Service

Simulation-as-a-Service

Hybrid Cloud Experiments with GENI for Multi-site Opt-in Use Cases

GENI for SaaS Resource Brokering

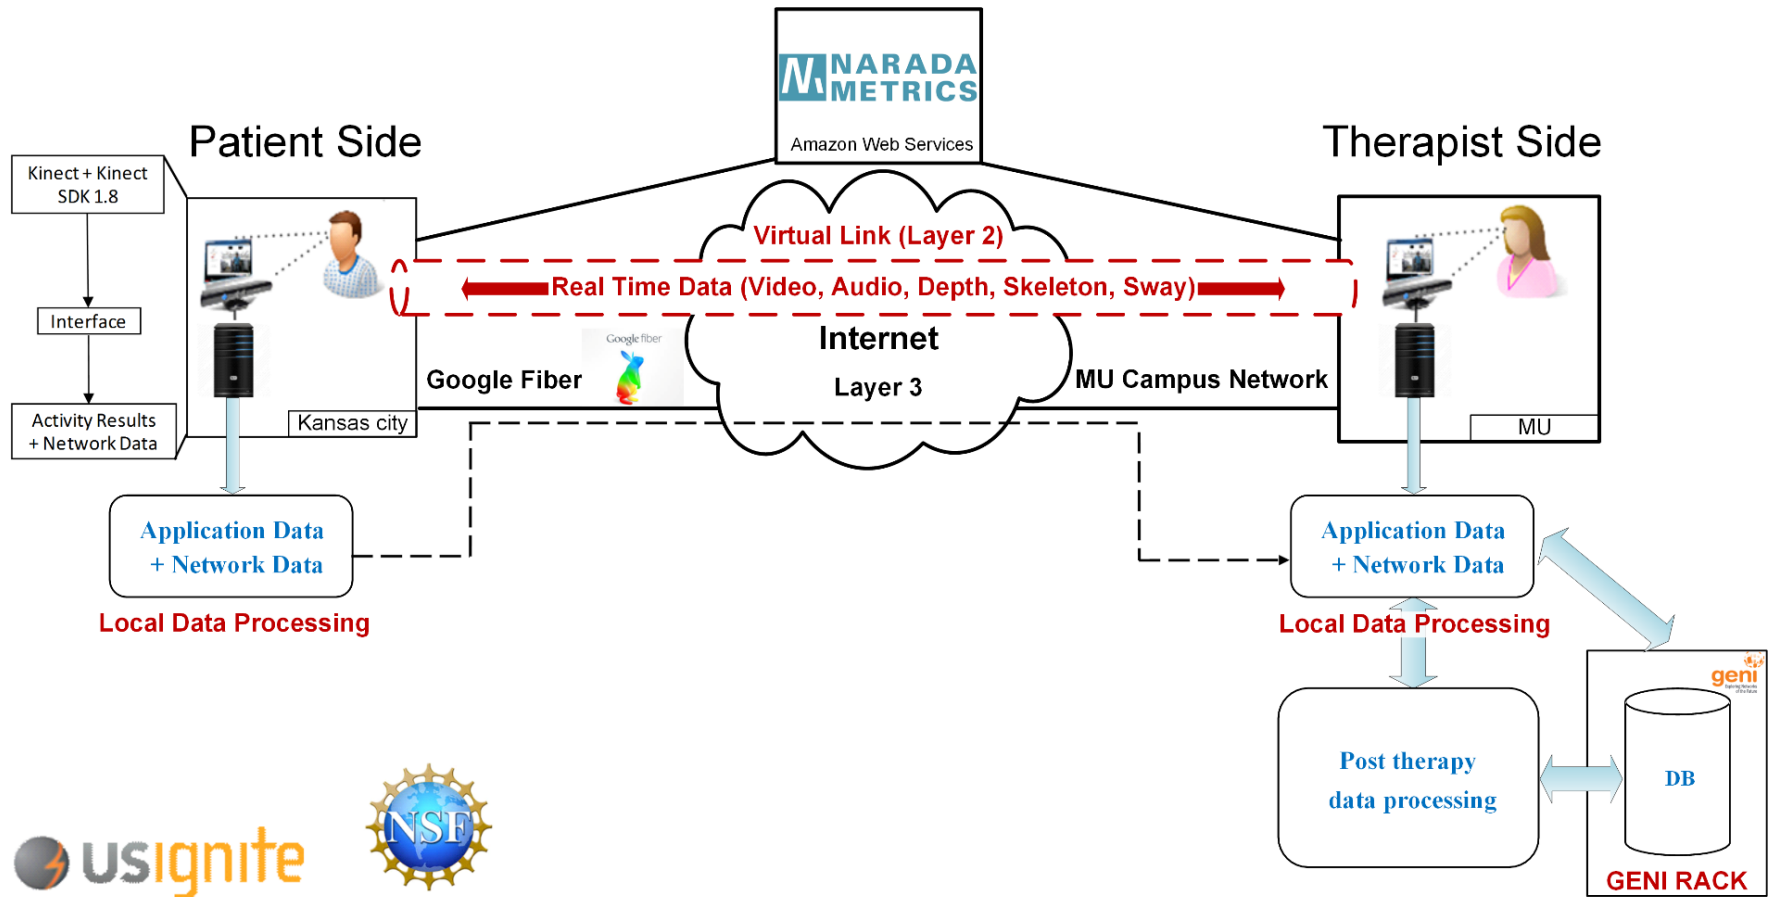
Our environment uses GENI to provision resources to both App developer and App customers based on the individual requirements



PhysicalTherapy-as-a-Service

GENI-Enabled In-Home, Personalized Health Monitoring and Coaching

PhysicalTherapy-as-a-Service



Summary of experiences in supporting GENI Services for User Opt-in

- Last-mile networks are a challenge
 - High-speed connections (e.g., Google Fiber in Kansas City, OARnet & MDC in Dublin, OH) are needed, but not enough...
 - Home networks do not have public IPs
 - Users are novices in networking technologies
- Getting a GENI Rack into a city data center/experiment slice was necessary
 - Challenges with hosting, networking and platform setup
 - Need for compatibility: what ever works in GENI, would it work in AWS?
- Virtualization complicates performance troubleshooting
 - Is it the application, network or virtualization?
 - App performance visibility is only possible if application-and-network-monitoring are closely integrated
- Issues such as cost of the App/Service become important to users
 - More effort on economic models in GENI can help Apps in their path to commercialization

Thank you for your attention!



IRB Issues

- User opt-in within GENI brings unfamiliar territory
 - PI group is hosting data or experiment outside the institution - how is it different from survey monkey?
 - What are the risks to users for experiment opt-in?
 - E.g., loss of reputation, loss of data privacy, loss of data or corruption of applications through spyware, loss of service or resources, (accidental/intentional) DDoS, legal exposure, ...
- The ultimate goal of the IRB can be paraphrased as balancing the risk to the subjects against the gains from the research
 - IRB is explicitly responsible for considering the rights and interests of the experimental subjects, and acts as a proxy for them

Use Cases (Moderated by Prasad)

- **Use Case: A Group of Users are Pulled into a GENI Experiment**
 - Issue: If group opt-in is arranged by a service provider, how do we provide for informed consent to “opt-out”?
 - Issue: Can a group “innocent” users be included in an experiment without their agreement? If so, what process must be followed? What if they are customers of an ISP?
 - Issue: How can we provide for “fast restoration” when the experiment fails, or must be brought down?
 - Issue: How do we deal with a user who has been pulled into an experiment, notices the difference, feels there has been a security breach, and complains?

Source: <http://groups.geni.net/geni/wiki/GeniOptIn>

Use Cases (2) (Moderated by Prasad)

- **Use Case: A 3rd Party User is Undesirably Affected by a Disruptive GENI Experiment (i.e., un-intended opt-in)**
 - Issue: Need capability to detect disruptive traffic where possible.
 - Issue: Need capability to trace back the disruptive traffic, through the aggregate that handled it, to the experiment that caused it.
 - Issue: Need capability to stop disruptive experiment.
 - Issue: Experiments that probe the Internet (or another network) should have capability to implement a “blocked call list” so that probing of a user’s machine can be blocked.
 - Issue: Should minimize overhead on operators.

Source: <http://groups.geni.net/geni/wiki/GeniOptIn>

Use Cases (3) (Moderated by Russ)

- **Use Case: A User Chooses to Participate in a GENI Experiment to Receive a Service**
 - Issue: When users make a choice to opt-in, what is the mechanism used to make the choice? Typically a web interface? If there is no web interface as part of the service, what can be done?
 - Issue: When users make a choice to opt-in, how can we insure that we have informed consent? What information must be provided? Is there a way to standardize on the approach to assure that it is done properly every time? Is there a standardized software package that can be utilized?
 - Issue: Privacy of individual opt-in users is essential, to meet both ethical and legal requirements.

Source: <http://groups.geni.net/geni/wiki/GeniOptIn>

Use Cases (4) (Moderated by Russ)

- **Use Case: A User Chooses to Participate in a GENI Experiment to Contribute Resources for Others**
 - Issue: Contributing resources entails significant risks for the user since they are typically running code on their own PC or node that has been provided by the researcher and/or other users. These include stability, security and privacy risks. Every effort must be made to mitigate these risks.
 - Issue: This user must fully understand the nature of the experiment, thus have provided informed consent.
 - Issue: If others are going to be using the shared resources, can the user who is contributing the resources limit what can be done with them? Perhaps just agree to a policy before they are contributed?

Source: <http://groups.geni.net/geni/wiki/GeniOptIn>