

# Intelligent Data Movement Service

## *Shakedown Experiments*

Ezra Kissel, Frank Diaz, Martin Swany  
Indiana University

October 29<sup>th</sup> 2013

GEC18



# Goals

---


- Create a distributed, dynamic data storage and movement service as a long-running GENI experiment
  - Aim high, then deal with reality!
- Distributed across many GENI aggregates, dynamism through responsiveness to network measurements
- Use existing, proven services as core components
- Allow for opt-in users to use IDMS as a functional system
  - Collect “real-world” data and adjust/optimize as necessary
- Explore the limits of GENI resource allocation capabilities

# Components

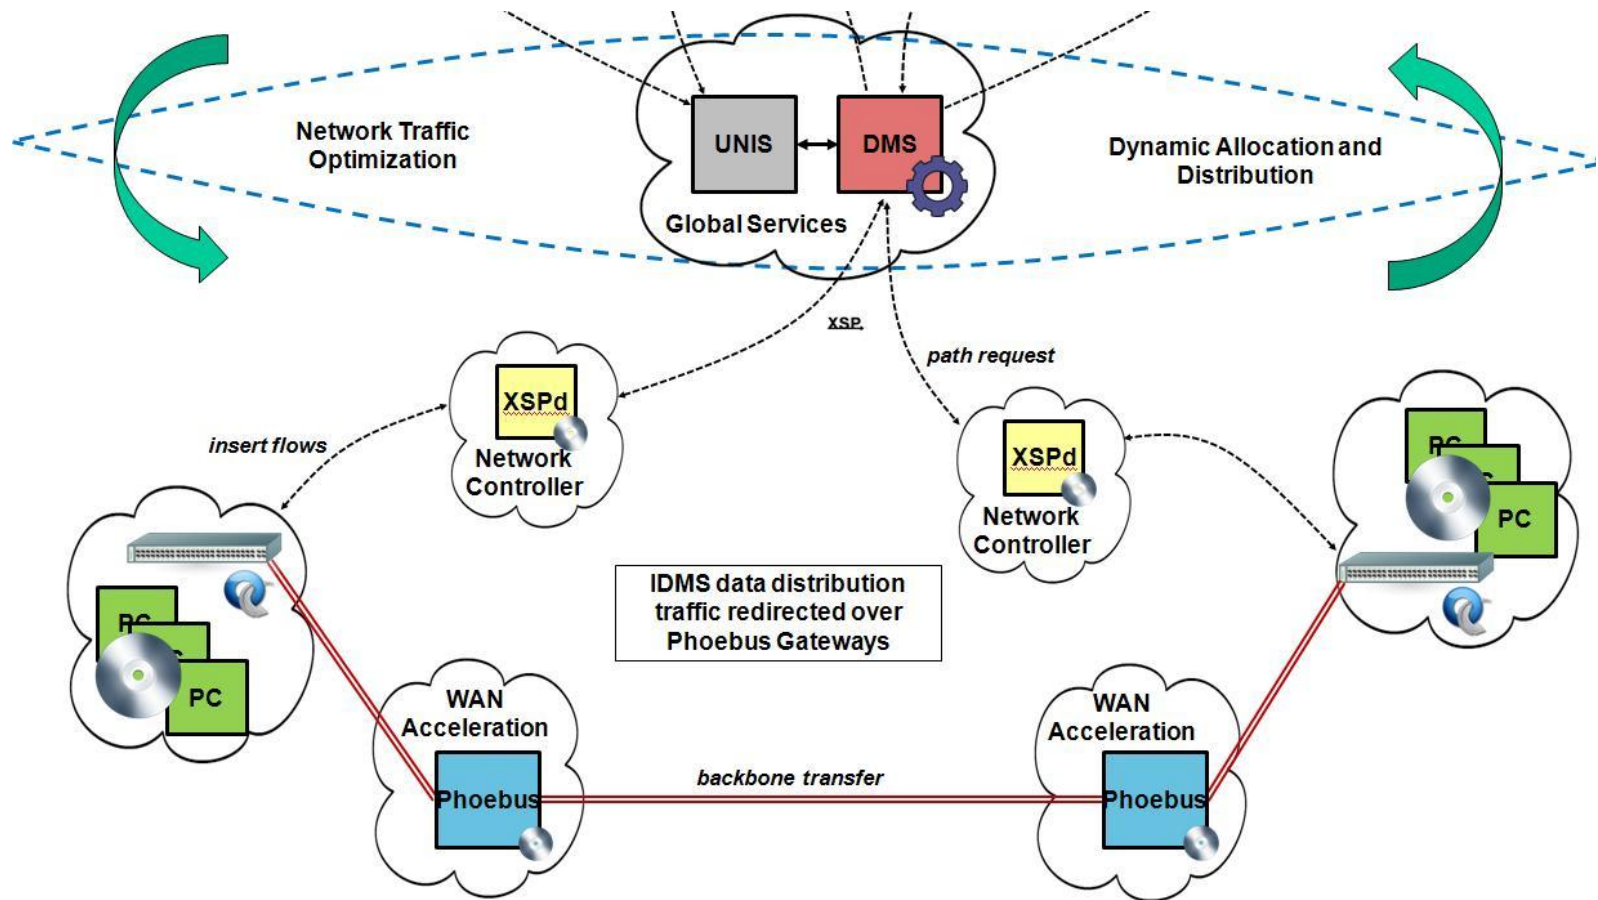
---

- **Storage depot**
  - IBP implementation from our colleagues at LoCI and ACCRE
  - Manages physical storage, allows for distributed “allocations”
  - Build upon existing client tools
- **Network Controller**
  - XSPd, session protocol interface for dynamic network control
  - Communicates with ION and OpenFlow
- **WAN Acceleration**
  - Phoebus software router, uses XSP for signaling
  - Deploy near backbone paths
- **Experiment I&M**
  - UNIS for service registration and discovery, plus tracking experiment resources
  - Active and passive measurements over experiment lifetime
- **Data Manager Service (new)**
  - Dynamic resource allocation and service placement
  - Algorithms for effective data distribution

# Resource allocation and instantiation



# Network optimization



# GENI Resources

---

- Racks for storage and control appliances
  - Leverage both InstaGENI and ExoGENI deployments
- “sliverable” storage
  - Recent addition to ExoGENI
- Backbone PCs
  - ProtoGENI and elsewhere, suitable for WAN acceleration
- FOAM aggregates for traffic flow optimization and redirection
- GENI I&M
- Existing GENI CF interfaces
  - Omni and stitching tools

# Challenges

---

- Availability of resources
  - Experiment design requires resiliency in the face of failures
- Aggregate storage capacity
  - Limits on user data may need to be enforced
- Understanding and reacting to network conditions
  - Building a non-trivial data management service
  - Do intra-slice measurements give us the right level of visibility?
- Compatible AuthN/AuthZ for users
  - Existing services already use GENI-compatible PKI mechanisms
  - Clients, scripts, and potentially a webGUI

# Reporting and availability

---

- Appliance images
  - IDMS services plus skeleton images
- Code repositories
- Experiment log
  - Experiences and specific GENI troubleshooting advice
- GENI Wiki
  - Progress reports
  - Links and references to everything above



# Thank you!

---

- Contact:
  - [ezkissel@indiana.edu](mailto:ezkissel@indiana.edu)
  - [fediaz@indiana.edu](mailto:fediaz@indiana.edu)
  - [swany@iu.edu](mailto:swany@iu.edu)
  
- DAMSL Group
  - <http://damsl.cs.indiana.edu>