GEC 10 Demonstrations
TransCloud: A Distributed Environment Based On Dynamic Networking

Rick McGeer, HP Labs
Joe Mambretti, Northwestern
Paul Müller , TU Kaiserslautern
Chris Matthews, Chris Pearson, Yvonne Coady, Victoria
Jim Chen, Fei Yeh, Northwestern
Andy Bavier, PlanetWorks
Marco Yuen, Princeton
Jessica Blaine, Alvin Au Young, HP Labs
Alex Snoeren, UC San Diego

March 16, 2010

http://www.icair.org

http://www.geni.net

Sponsored by the National Science Foundation

Introduction – TransCloud

- TransCloud = A Cloud Where Services Migrate, Anytime, Anywhere In a World Where Distance Is Eliminated
  - Joint Project Between GENICloud, iGENI, et al
  - GENICloud Provides Seamless Interoperation of Cloud Resources Across N-Sites, N-Administrative Domains
  - iGENI Optimizes Private Networks of Intelligent Devices Capable of Dynamically Provisioned Low-Latency, High-Performance Communications Among Multiple Physically-Distributed Infrastructures and Federated Domains

Introduction – TransCloud

- General Premise = Transition From Legacy Architecture, Technologies, Protocols, Implementations
  - Legacy Environments Reflect Obsolete Concepts of Facility-based Services
  - Legacy Environments Do Not Leverage The Potential Of A Wide Range Of Exceptionally Powerful Emerging Models, Architecture, Technologies, and Services
  - The GENI -- “Clean Slate” Approach -- Enables Creating and Implementing Totally New Types of Distributed Environments, Including Those Integrating New Cloud and Networking Architecture and Technology

Context 1: Seamless Computation Services Available Anytime, Anywhere

- Everybody Wants “the Cloud”…BUT…
  - Performance of Cloud services Highly Dependent On Location
    - Of End-User, Applications, Middle Processes, Network Topology
    - Of Cloud Data, Compute Processes, Storage, etc
- Why?
  - Performance of Legacy Protocols and Other Elements, Highly Dependent on Latency Issues
- Therefore:
  - If the Clouds Are Too Far Away, Performance Will Be Severely Restricted
  - If Clouds Have Large Internal Latencies AND Uses Legacy Protocols, Performance Will Be Very Severely Restricted
- Ergo
  - Clouds Needs To Be Close To Experience Sites OR
  - Networks (And Clouds) Can Be Designed To Eliminate Distance

Context 2: Living With Legacy Protocols Over Commodity Internet vs Creating Alternatives

- Legacy Is There For a Reason
  - Compatibility
  - Fairness
  - Congestion Avoidance
  - Other Considerations
- Therefore: Distributed Cloud
  - Minimal Latencies Over Legacy Internet To Anywhere/Everywhere
- Therefore: Private Internal Networks
  - Eliminate Latency Dependence Internally
  - Use Aggressive Internal Transport/Application Protocols
    - TIA-1039, Reliable Blast UDP, Lambda RAM
    - Flow Control Enabled….And Other Techniques/Technologies

Context 3: General Considerations

- Major Cloud Use Case: Big Data, Distributed Collection, Must Live With Available Networks
  - Smart Cities
  - Sensor Nets
  - Enterprise
- Current World: May Have To Rely On Provider Service
- Best Case: Create Private Network
  - Owning Optical Fiber
  - Create High Performance Wireless Point-to-Point Links
- Many Data Intensive Science Projects, Including
  - High Energy Physics (e.g. LHCNet, Science Data Network, I-WIRE)
  - Atmospheric Sensing Apparatus
  - Ocean Observing (e.g., Project Neptune)
  - Distributed Radio and Optical Telescopes, etc

Premise: Compute Where Data Lives!

- Computation is Ubiquitous and Easy To Obtain
- Programs Are Small and Easy to Transmit
- Most Programs Can *Reduce* Data
- Often Data Is Large and Challenging To Transmit
  - E.g., Jim Gray distributing SDSS by sending computers by FedEx!
- *Solution -- Send Programs to Data*

Solution – TransCloud

- Introducing TransCloud Prototype
  - An Early Instantiation of the Proposed Architecture
  - A Distributed Environment That Enables Component and Interoperability Evaluation
  - A Distributed Environment That Can Enhance GENI's Potential As a Research Instrument
  - A Testbed On Which Early Experimental Research Can Be Conducted
  - An Environment That Can Be Used To Explain/Showcase New Innovative Architecture/Concepts Through Demonstrations, e.g., Three Demonstrations at GEC 10

Solution: TransCloud

- TransCloud = A Highly Geographically Distributed Environment That Can Support Services Based On
  - Highly Distributed Processes – Any Process, Anywhere World Wide – Freed From Physical Dependencies
  - Using Multiple Independently Administered and Distributed Resources, Including Compute Processing, Dynamic Networking, Storage, Data, Analytics, etc.
  - Using a Ubiquitous Environmental "Stack" ~ Conceptually, Creating a Type of "TCP" Stack Oriented For This Much Larger Blend of Resources

Introduction – TransCloud

- Several Basic TransCloud Concepts
  - High Performance Highly Distributed Cloud Architecture Allowing Processes Across Multiple Administrative Domains Integrated With Dynamic Networking (GENI)
  - Utilizing Scalable Lightweight Federation Processes
  - Services Are Based On Processes That Can Be Executed Anywhere World-Wide (Location Independent)
  - Top Level Services Can Be Accessed Via Public Internet
  - Core Processes and Data Streams Leverage Sophisticated Communication Services Not Merely "Best Effort" Commodity Internet

Introduction – TransCloud

- TransCloud Architectural Components
  - High Level APIs
  - A High Performance General Programming Environment
  - A Wide Area Programming Environment Integrated With Query Systems And High Performance Data Access Services
  - Resource Management Frameworks, Including Cluster, VM and Network Resource Management
  - High Levels of Virtualization Based on VMs and Network Abstractions

TransCloud Architecture

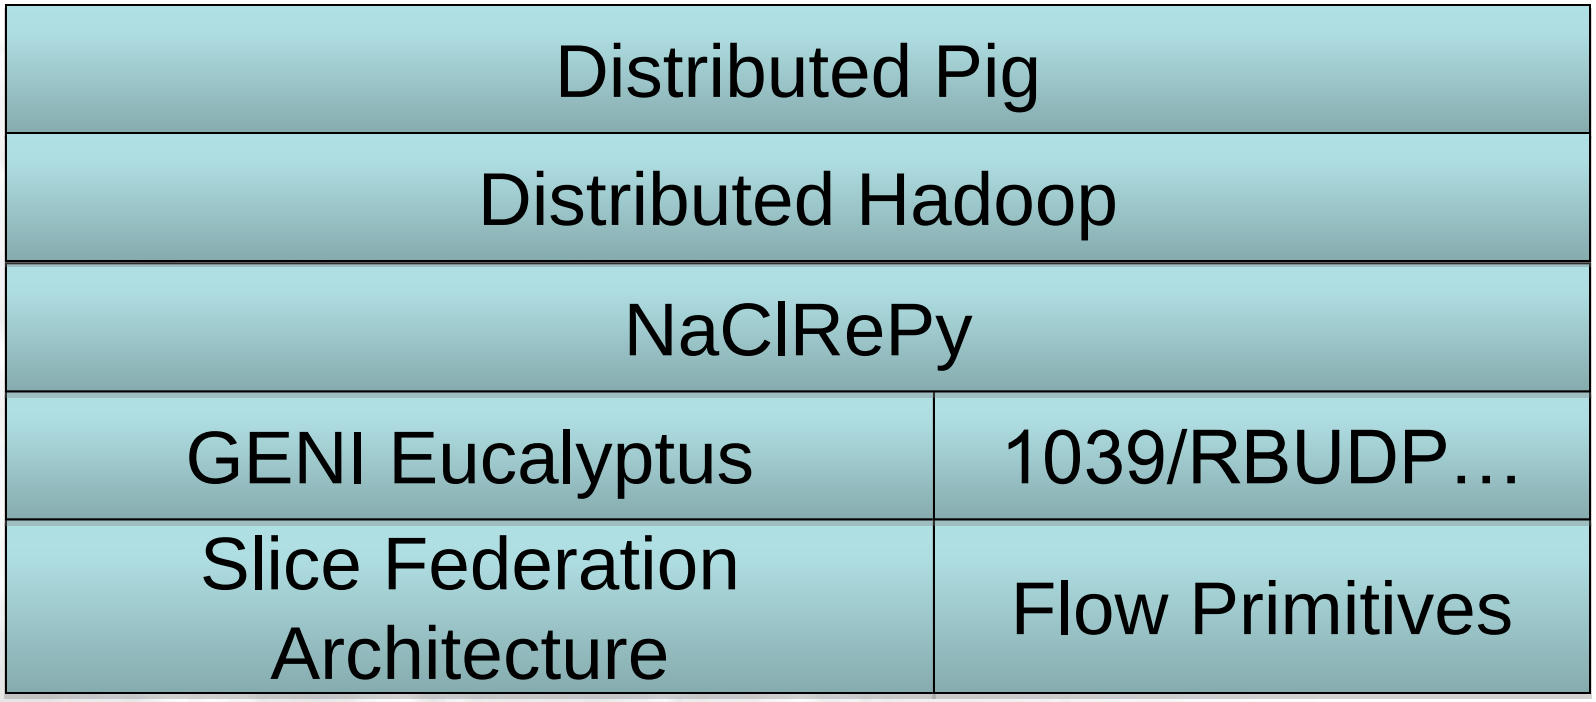

Demonstrations – TransCloud Sites and Mesh –Using the Global Lambda Integrated Facility (GLIF)

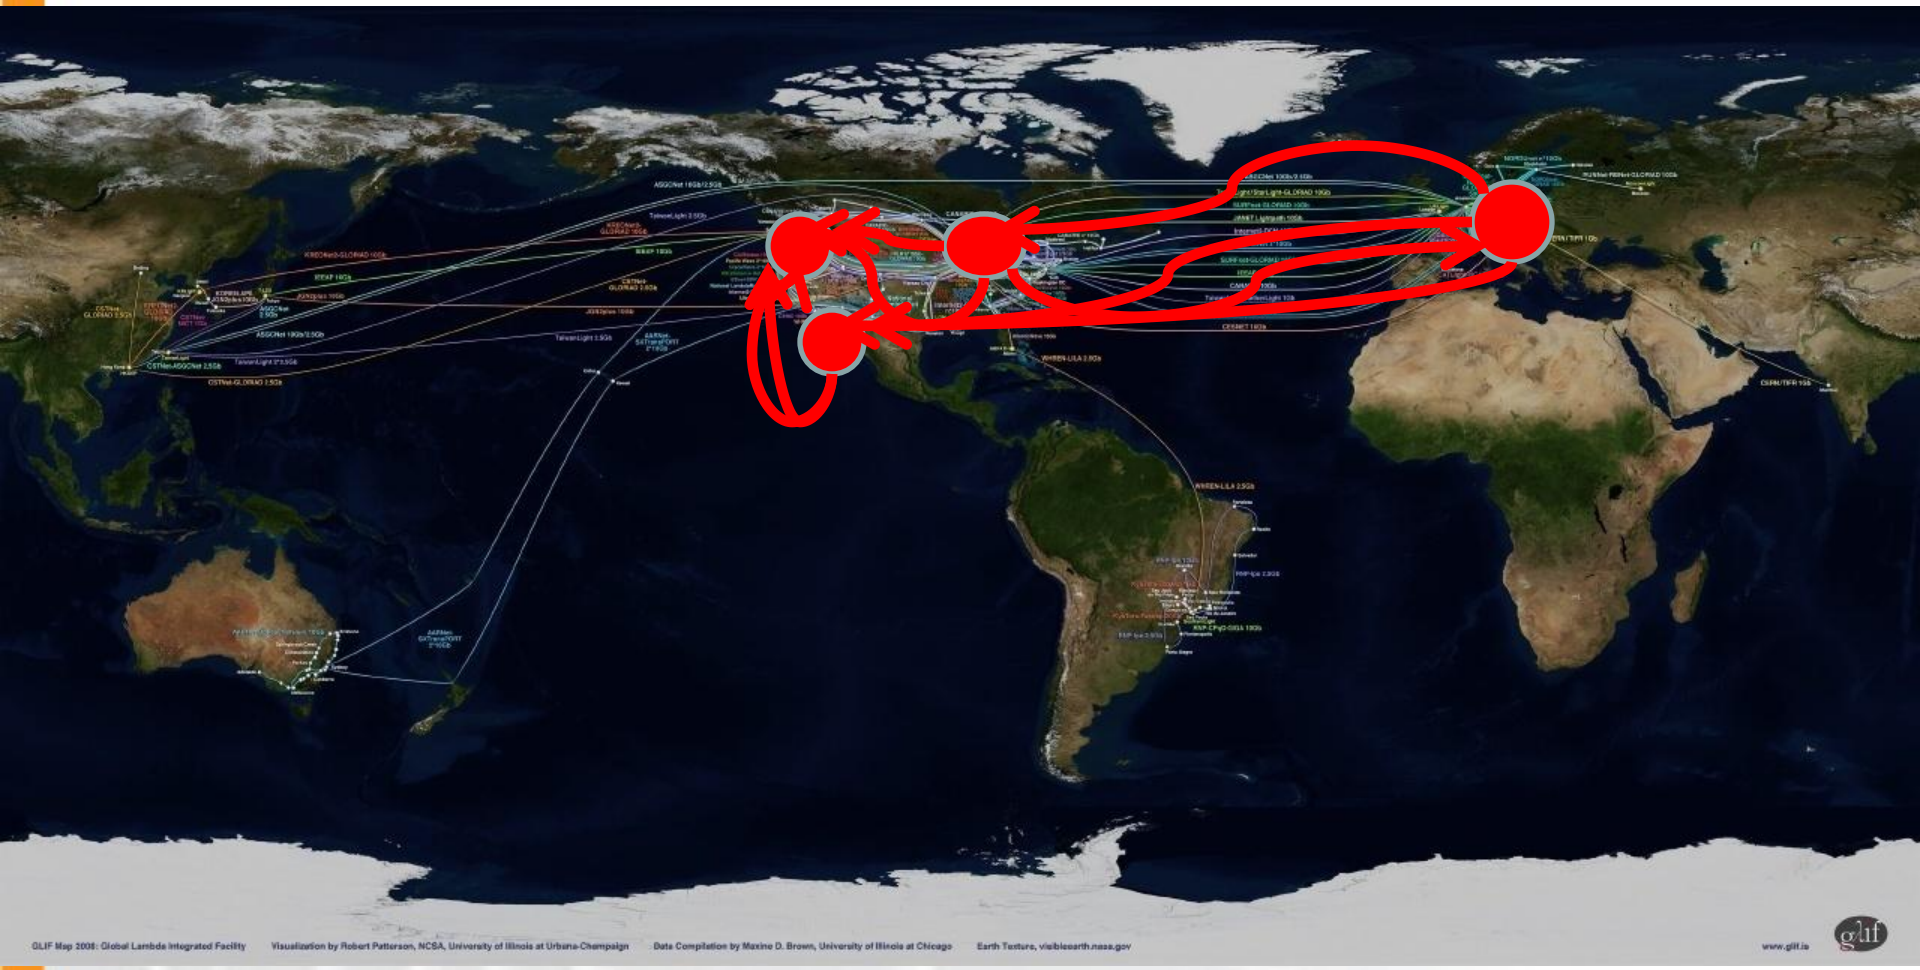

TransCloud Today

- Sites at
  - HP Labs, Palo Alto
  - UC San Diego
  - Northwestern
  - Kaiserslautern
- Tomorrow (*literally*!)
  - Amsterdam
- Connectivity provided by:
  - CAVEWave, StarLight, NetherLight, DFN, National Lambda Rail, Global Lambda Integrated Facility

Demonstration 1A – TransCloud

- TransCloud Prototype Demonstration 1A
  - A Complex Query Is Initiated
  - Process Discovers Packages/Integrates Required Resources Resident At Multiple Sites, Across Multiple Domains, Including International (US⇔Germany)
  - Query Is Executed Using Dynamically Instantiated Fabric
  - Result Is Produced

Demonstration 1B – TransCloud

- TransCloud Prototype Demonstration 1B
  - Complex Query Number 1 Is Initiated
  - Process Discovers Packages/Integrates Required Resources Resident At Multiple Sites, Across Multiple Domains, Including International (US⇔Germany)
  - Query 1 Is Executed
  - Result 1 Is Produced
  - Complex Query Number 2 Is Initiated
  - Process Discovers Packages/Integrates Required Resources Resident At Multiple Sites, Across Multiple Domains
  - Query 2 Is Executed
  - Result 2 Is Produced
  - Results 1 and 2 are Combined Are Delivered

Demonstration 2 – TransCloud

- TransCloud Transcoding
  - Demonstration of Converting New Technology Capability Into Practical Service (Using In Part Control Based on XML RPC)
  - Traditionally Digital Media Has Used Different Infrastructure For Different Edge Delivery Platforms
  - This Demonstration Shows the TransCloud Provides a Capability for Using One High Performance Distributed Environment for Transcoding For Multiple Platforms
    - Mobile Phones
    - Computers
    - Tablets
    - Tile Displays
    - Et Al -- Unlimited

How TransCoding Demonstration Works

- TransCloud: Set of Protocols, Standards, Management Software That Enables Interoperation of Distinct Cloud Resources
- TransCloud: Advanced Distributed Global Environment That Enables Dynamic Creation of Communication Services, Including Those Based On Rapid Migration of Virtual Network and Cloud Resources

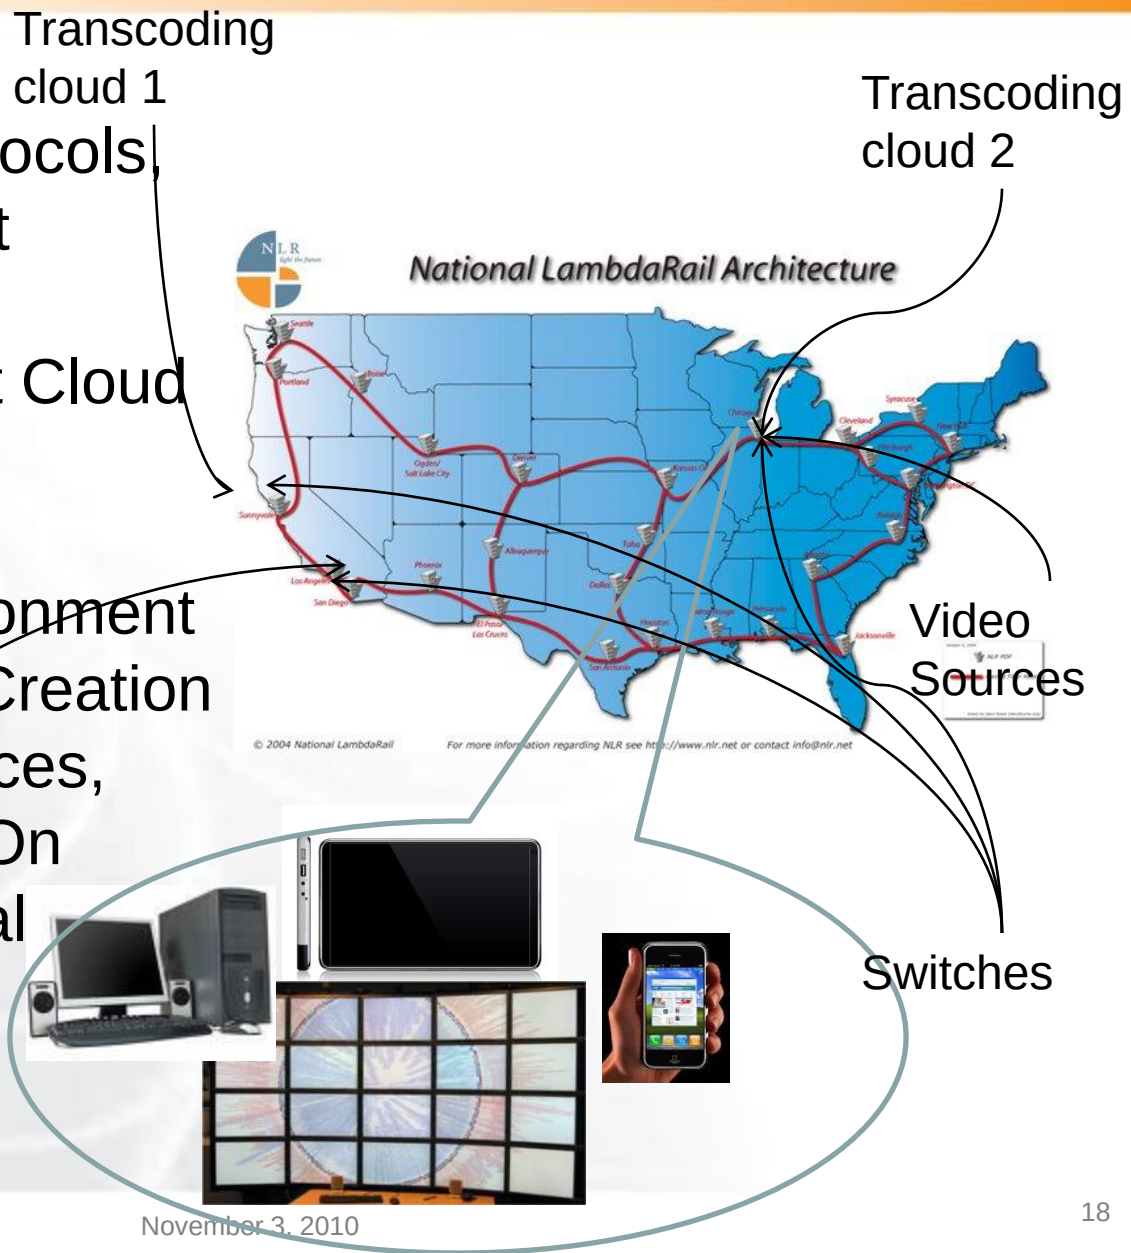

Demonstration 2

UCSD TransCloud

HP Labs OpenCirrus TransCloud

iCAIR Trans Cloud

Kaiserslautern TransCloud

The TransCloud on GENI

- GENI = An Indispensible Resource For Transcloud
  - GENI Standard (Slice-Based Federation Architecture) Is Key To Interoperation of Multiple Domain Clouds and Forms the Critical Federation Standard
  - Deep Network Programmability and Long-Distance L2 Networking Are Vital to Seamless Core Process and Data Migration
- Key Future Goals
  - Hosting Researchers On the TransCloud Platform
  - Providing Key Resources for e-Science, Network Science
  - Continuing Intercontinental Expansion of the TransCloud

Advancing TransCloud

- If You Are Interested In Using This Environment, Contact Us
- If You Would Like To Contribute Resources, Contact Us

TransCloud at GEC 10

- **THANKS!**
- **Questions????**