



D u k e S y s t e m s

NSF Workshop on Designing Tools and Curricula for Undergraduate Courses in Distributed Systems, Boston, July 8, 2012

ExoGENI: A Multi-Domain IaaS Testbed

Jeff Chase

Duke University / RENCI

Ilia Baldine
Yufeng Xin
Anirban Mandal



Paul Ruth
Aydan Yumerefendi
Chris Heerman




Exploring networks of the future

Evolving technological and social networks, intertwined and worldwide in scope, are rapidly transforming societies and economies. The Global Environment for Network Innovations (GENI), a project sponsored by the National Science Foundation, is open and broadly inclusive, providing collaborative and exploratory environments for academia, industry and the public to catalyze groundbreaking discoveries and innovation in these emerging global networks.

GENI is a virtual laboratory at the frontiers of network science and engineering for exploring future internets at scale. GENI creates major opportunities to understand, innovate and transform global networks and their interactions with society.

Understand. Innovate. Transform – Make an Impact on Global Networks.

Federated substrate with end-to-end virtualized “slices”



RENCI, Duke to build experimental networking infrastructure

Published: Friday, December 9, 2011

CHAPEL HILL, Dec. 9, 2011—RENCI at the University of North Carolina, Chapel Hill and Duke University in partnership with IBM will lead a new project to build a nationwide test bed for networking and networked cloud computing.

The project is part of NSF's Global Environment for Network Innovation (GENI) initiative, which enables researchers to explore networks of the future.

The National Science Foundation (NSF) awarded just over \$2 million to the three-year ExoGENI project, led by Ilia Baldine, director of RENCi's networking research group and Jeff Chase, a Duke University computer science professor.



<http://wiki.exogeni.net>



Ilia
Baldine

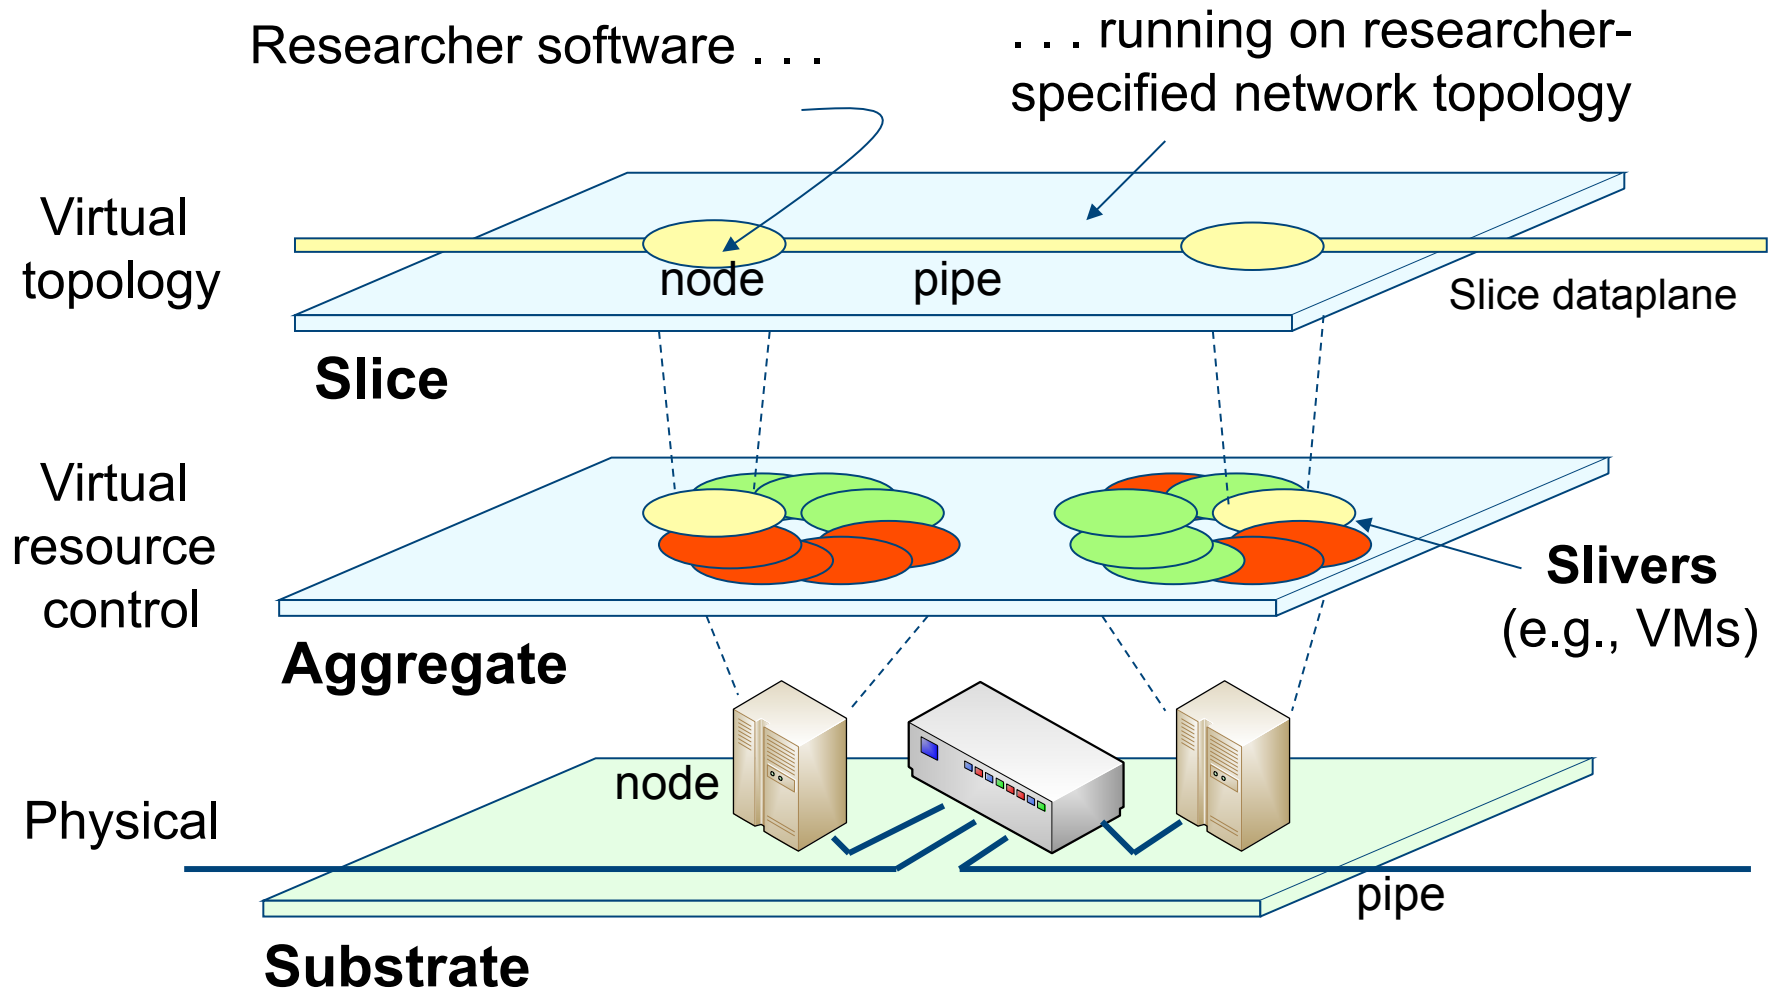


ExoGENI Testbed

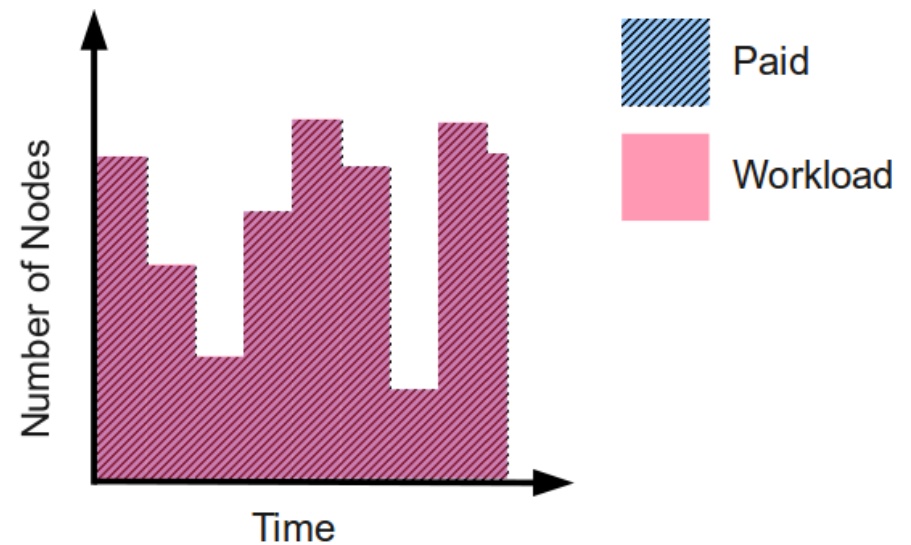

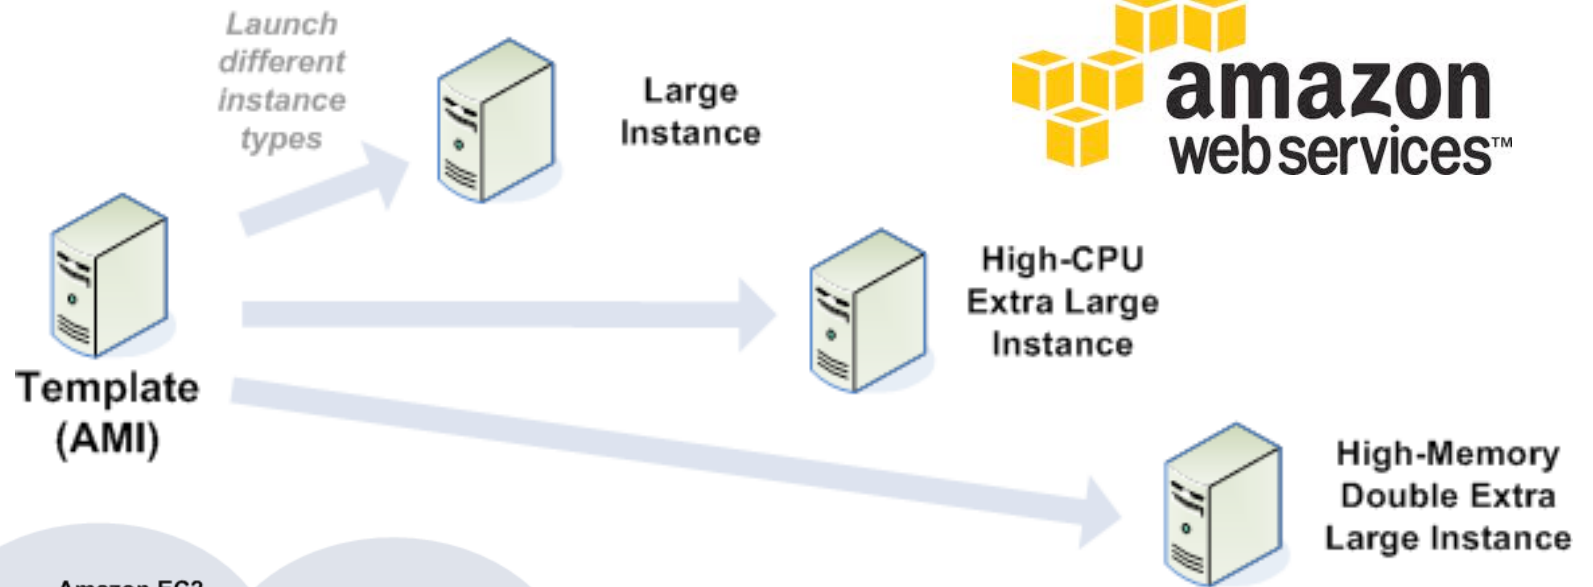
- ▶ 14 GENI-funded “Racks”
 - IBM x3650 M3/M4 servers
 - 2 local drives, 48G RAM
 - Dual-socket 8-core CPU w/ Sandy Bridge
 - 10G dual-port Chelso adapter
 - BNT 8264 10G/40G OpenFlow switch
 - DS3512 6TB sliverable storage
 - iSCSI interface for head node image storage as well as experimenter slivering
- ▶ Each rack is a small networked cloud
 - Interconnected by combination of dynamic and static L2 circuits: regionals + national fabrics
 - OpenFlow-enabled
- ▶ Also accommodates “volunteer” sites
- ▶ <http://wiki.exogeni.net>



GENI resource model



EC2: The Canonical Cloud



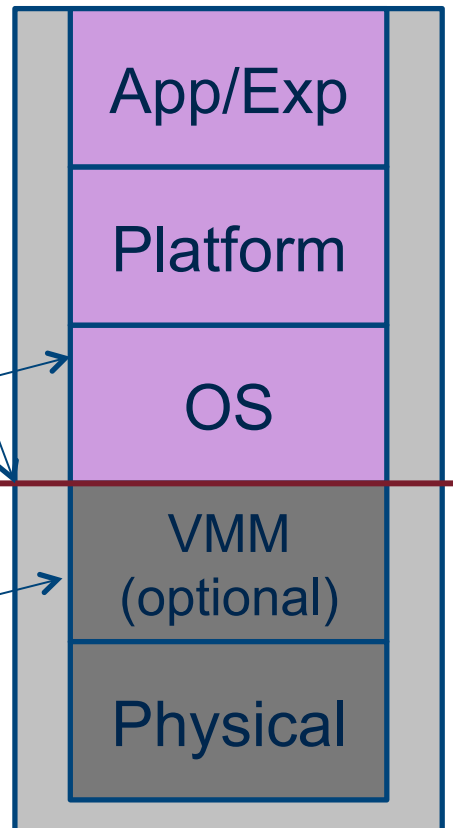
IaaS: infrastructure services

NIST Cloud Computing Reference Architecture


“Consumers of **Infrastructure-as-a-Service** have access to virtual computers, network-accessible storage, network infrastructure components, and other **fundamental computing resources...**”

Example: virtual server

1. Host offers instances of a machine abstraction.
2. Guest (slice) requests and programs virtual resource instances.
3. Host contract defines performance and isolation properties.




Deployment of IaaS clouds is growing rapidly w/ open cloud software.



IaaS: clouds and network virtualization

Virtual Compute and Storage Infrastructure




Cloud APIs (Amazon EC2 ..)




Cloud Providers

Virtual Network Infrastructure



Dynamic circuit APIs (NLR Sherpa, DOE OSCARS, I2 ION, OGF NSI ...)



Transport Network Providers

ORCA

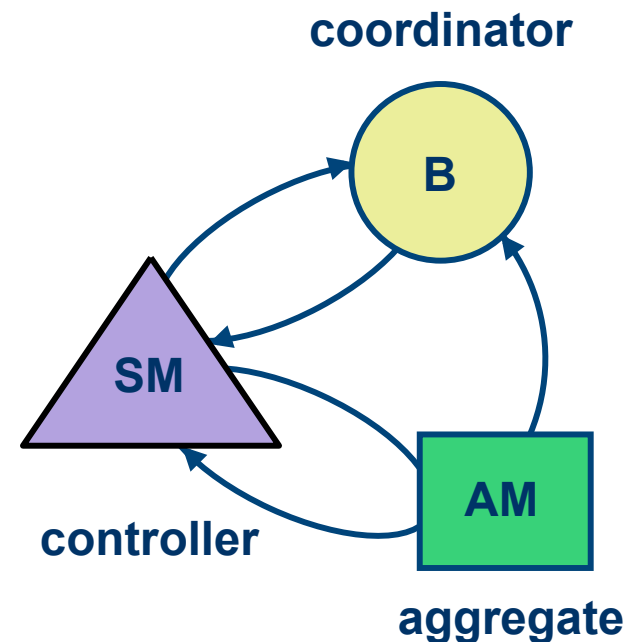
Open Resource Control Architecture



- **ORCA is a “wrapper” for off-the-shelf cloud and circuit nets etc., enabling federated orchestration:**

- + Resource brokering
- + VM image distribution
- + Topology embedding
- + Stitching
- + Authorization

- **GENI, DOE, NSF SDCI+TC**
- <http://networkedclouds.org>
- <http://geni-orca.renci.org>

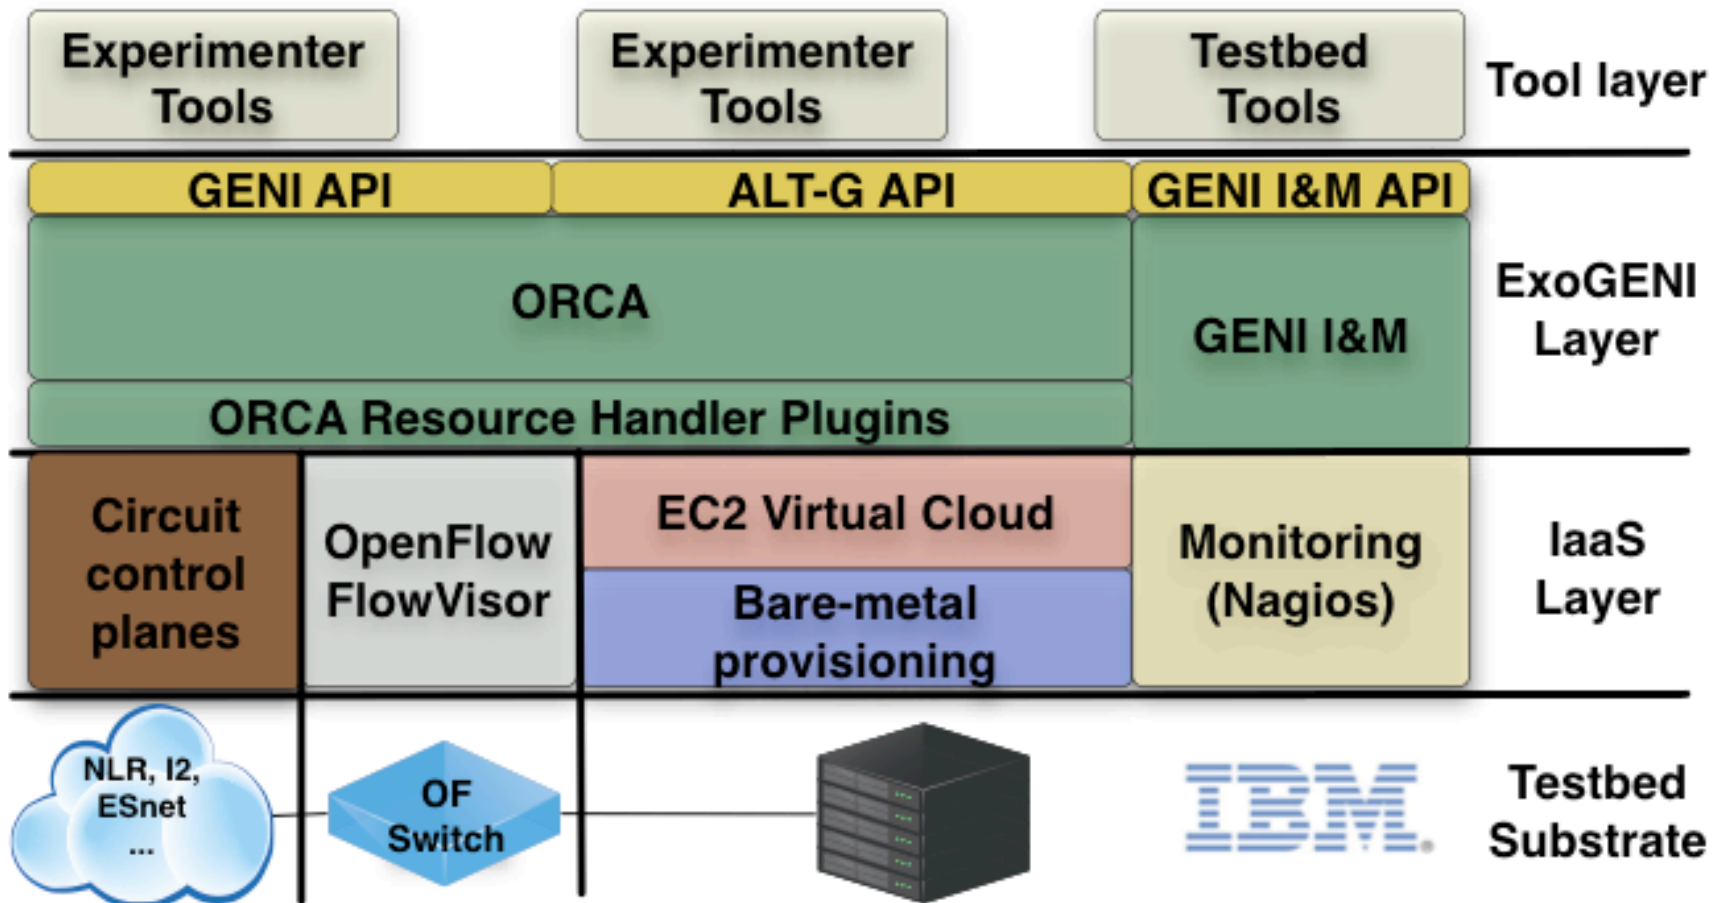


ExoGENI

- **Every Infrastructure as a Service, All Connected.**
 - Substrate may be volunteered or rented.
 - E.g., public or private clouds and transport providers
- **ExoGENI Principles:**
 - Open substrate
 - Off-the-shelf back-ends
 - Provider autonomy
 - Federated coordination
 - Dynamic contracts
 - Resource visibility

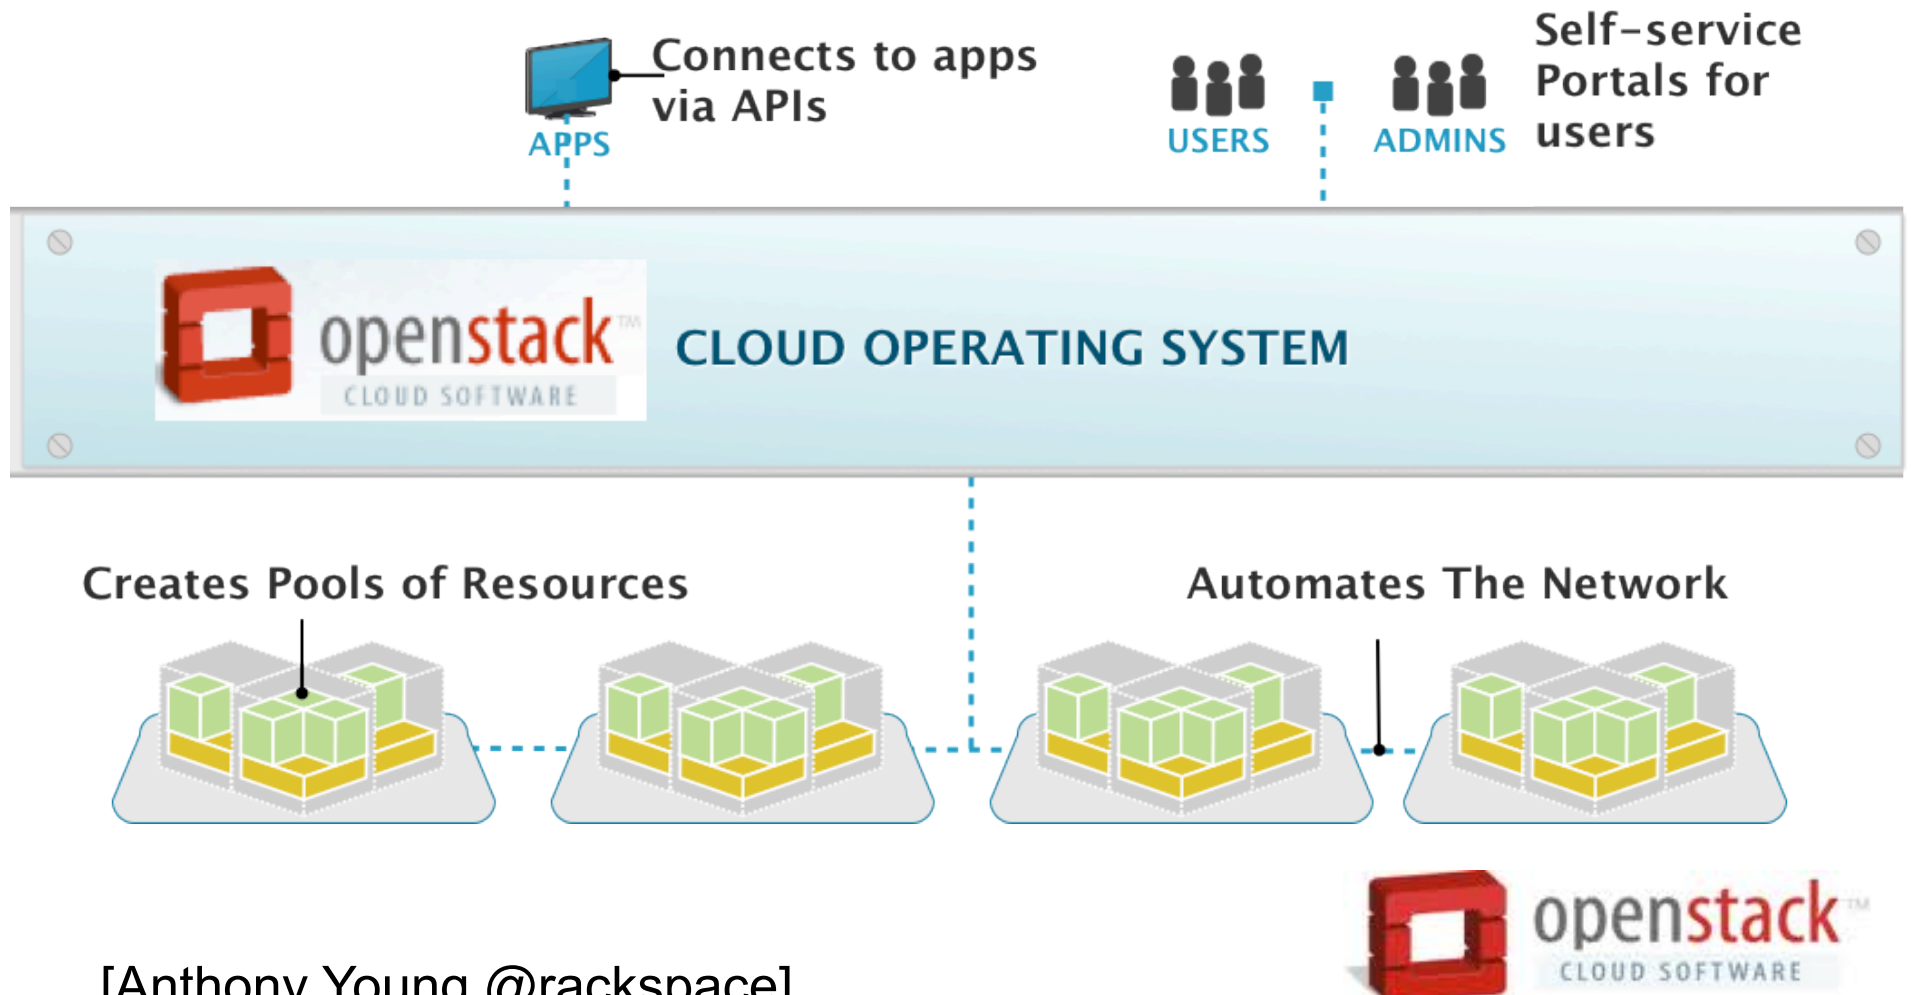


ExoGENI software structure



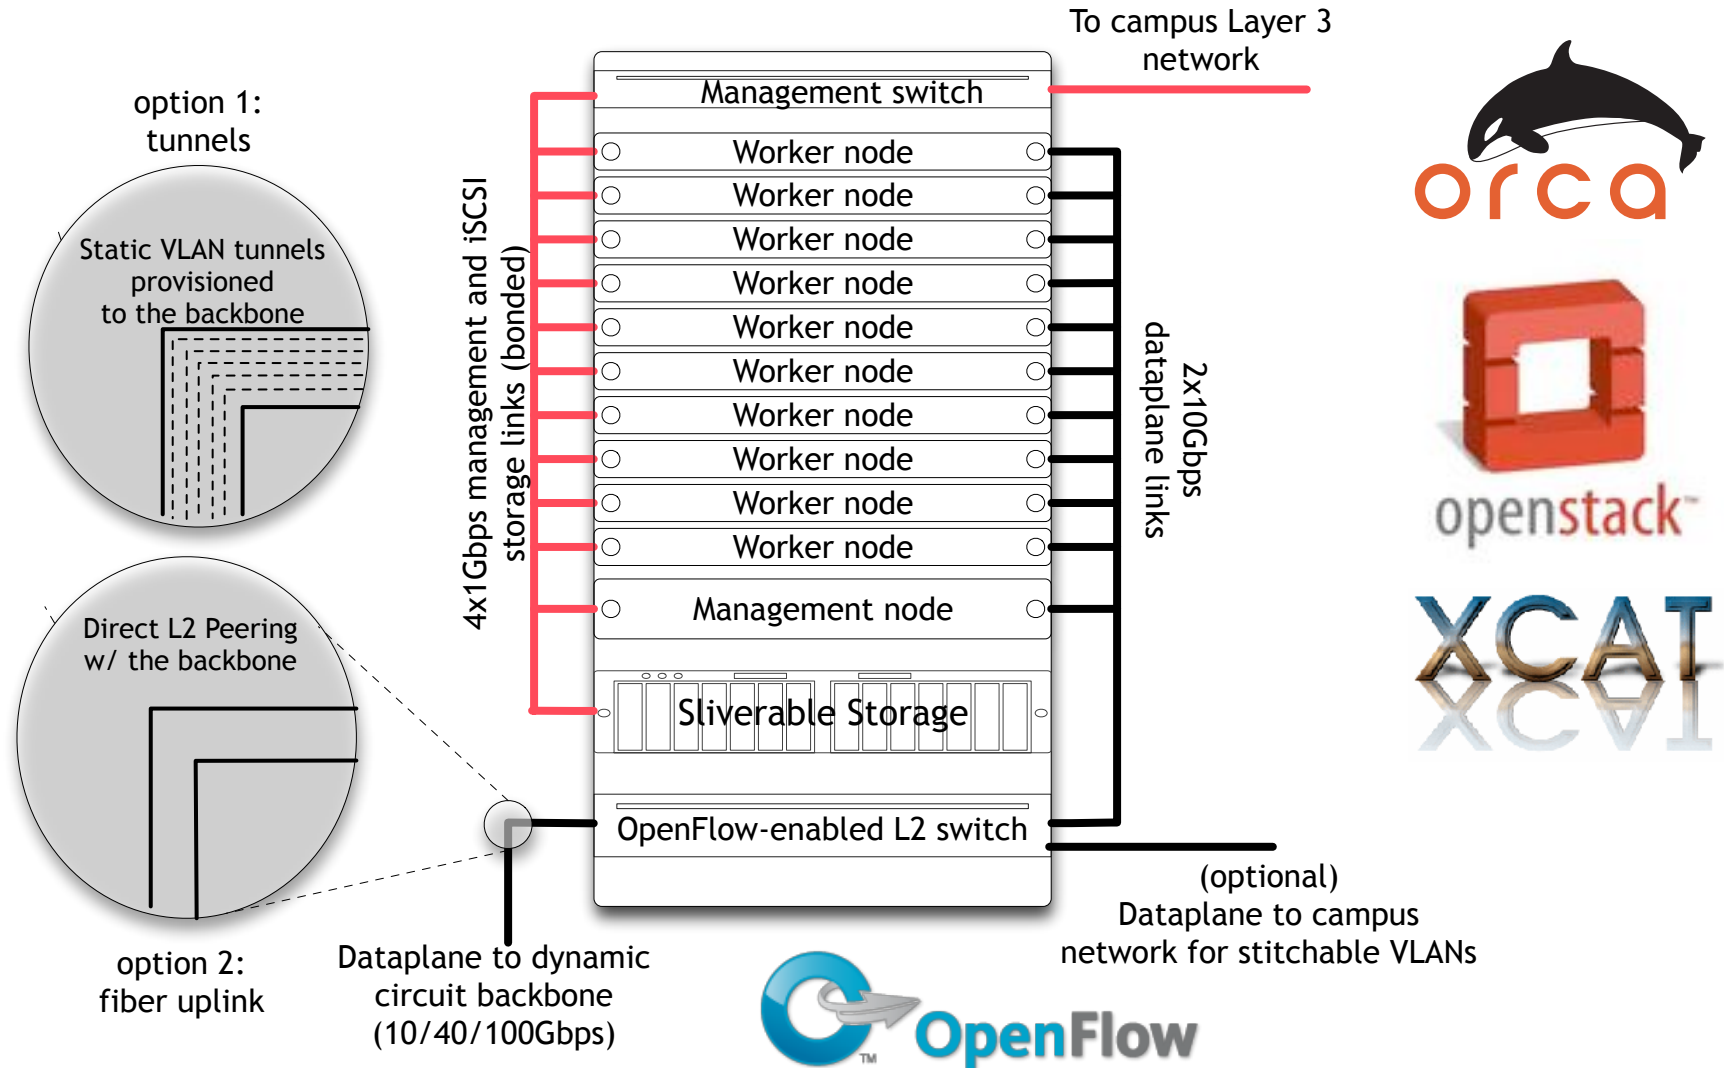
OpenStack, the Cloud Operating System

Management Layer That Adds Automation & Control



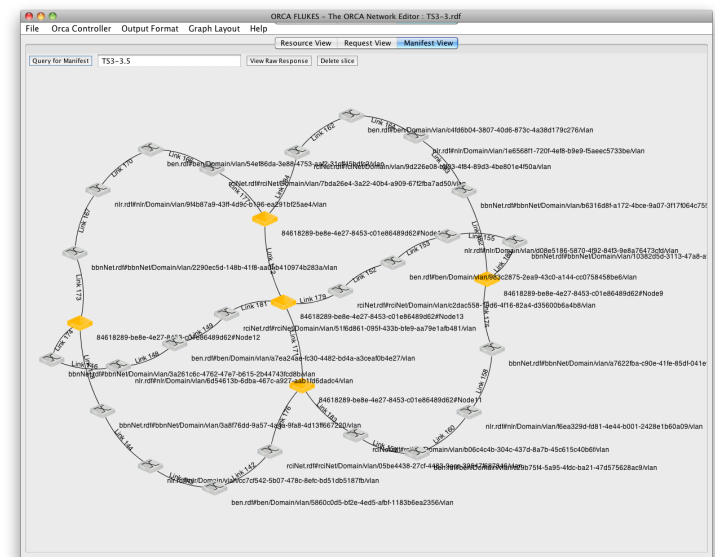

[Anthony Young @rackspace]

An ExoGENI cloud “rack site”

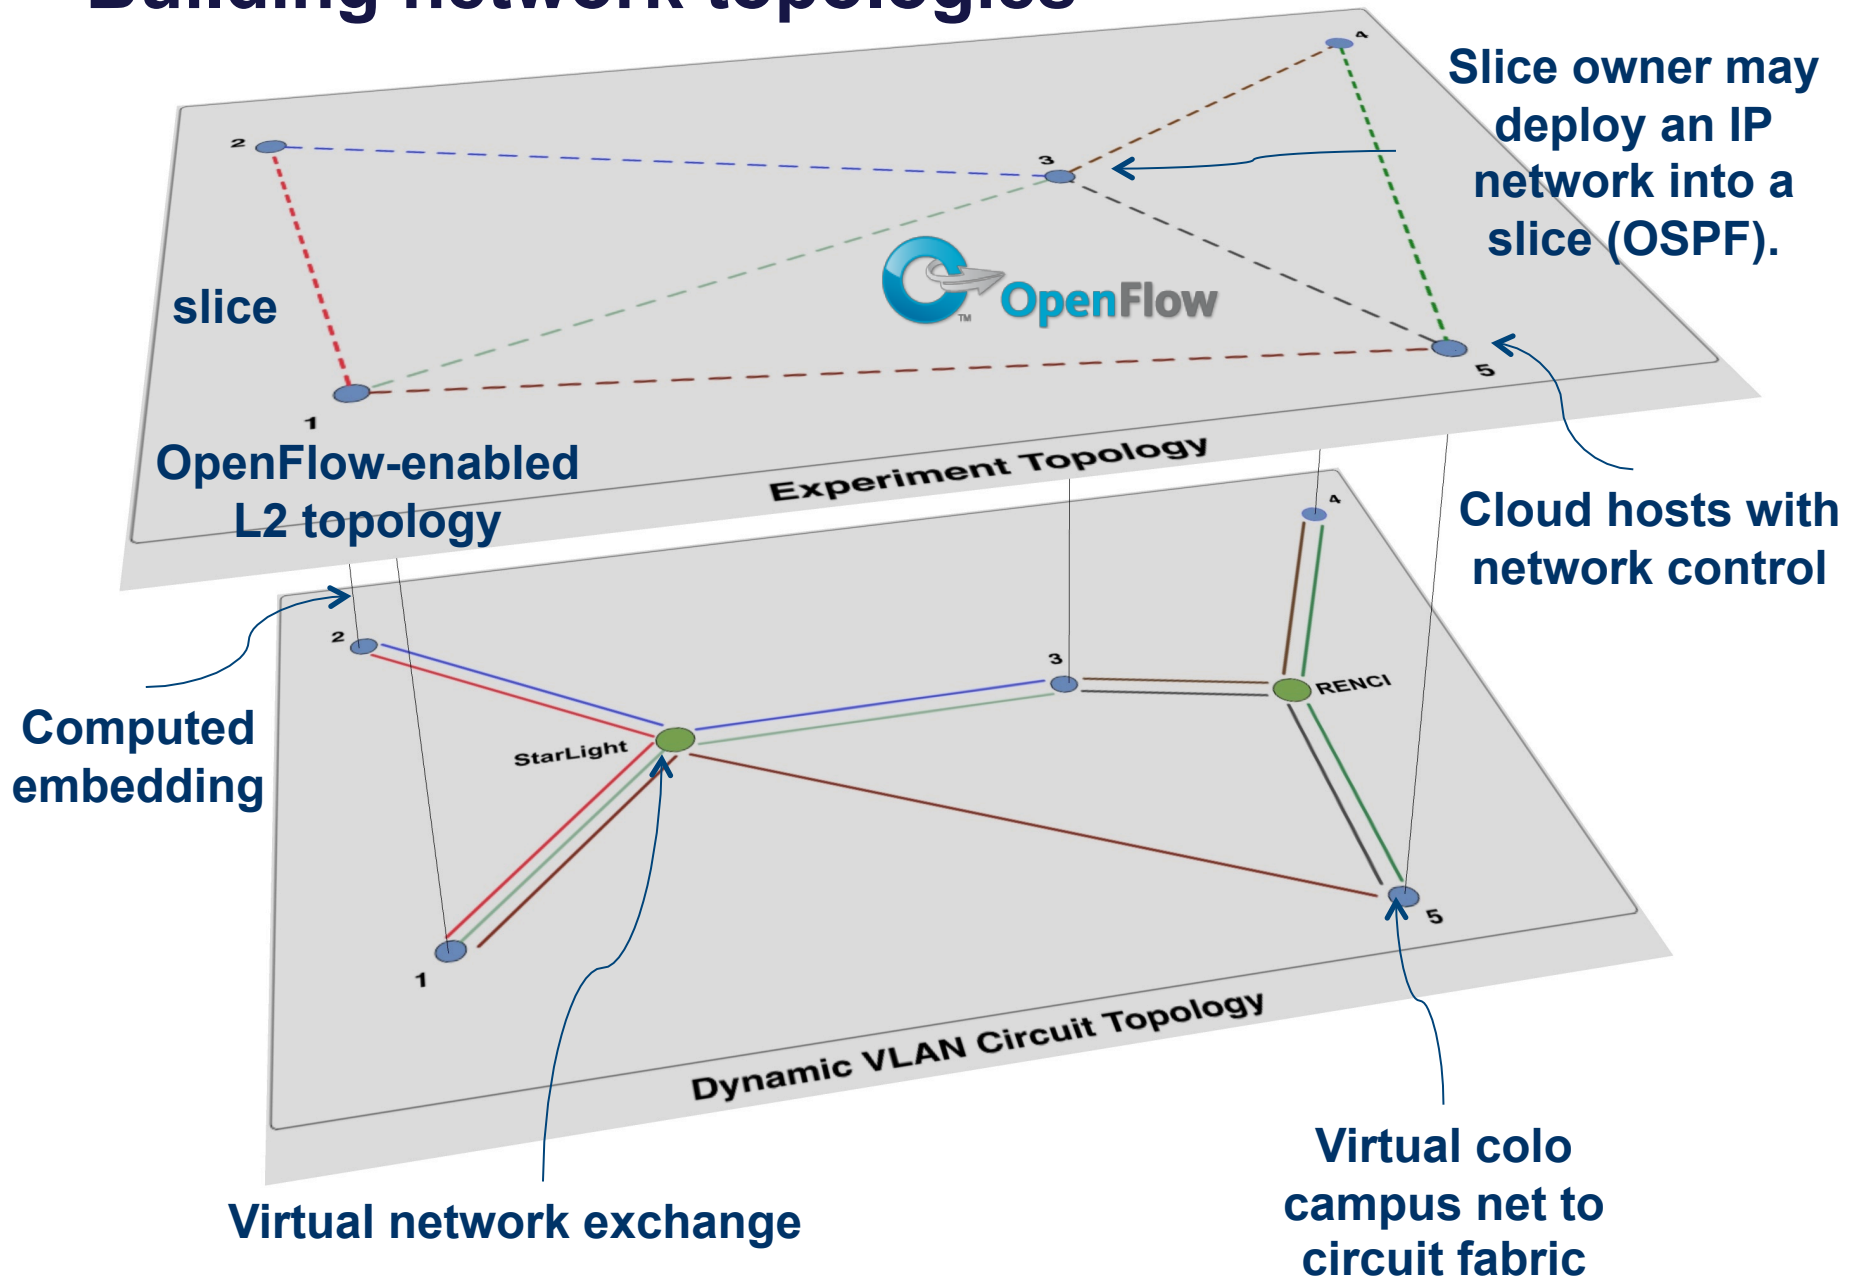


ExoGENI slice isolation

- ▶ Strong isolation is the goal
- ▶ Compute instances are KVM based and get a dedicated number of cores
- ▶ VLANs are the basis of connectivity
- ▶ Register OpenFlow controllers for your VLANs



Building network topologies



Flukes GUI

The image displays the ORCA FLUKES - The ORCA Network Editor GUI. It features three overlapping windows:

- Top Left Window:** Shows a network graph with three yellow nodes connected by links labeled LINK1, LINK2, and LINK3. A text field contains the query: `ebdf01a7-10bb-4e5d-acb1-d7177705874#Node`.
- Top Right Window:** Displays a map of the United States with various regions highlighted in different colors (red, blue, green, yellow). Red cube icons are placed on the map, with labels such as "Straight site authority" and "Pittsburgh".
- Bottom Window:** Shows a network graph with a central yellow node labeled "Node0" connected to three blue server rack icons labeled "NodeGroup0", "NodeGroup1", and "NodeGroup2" via links "Link2", "Link1", and "Link0" respectively. A "Node Properties" dialog box is open, showing configuration options for "NodeGroup2":
 - Name: NodeGroup2
 - Select node type: Site default (selected), Euca c1.medium, Euca m1.small
 - Select image: None
 - Select domain: System select (selected), Duke CS, NERSC
 - Splittable between domains:
 - Select dependencies: NodeGroup0, Node0, NodeGroup1
 - Link0 IP Address: [] . [] . [] . [] / []
 - Internal VLAN IP Address: [] . [] . [] . [] / []
 - Number of servers: 1
 - Additional properties: PostBoot Script



ExoGENI Status

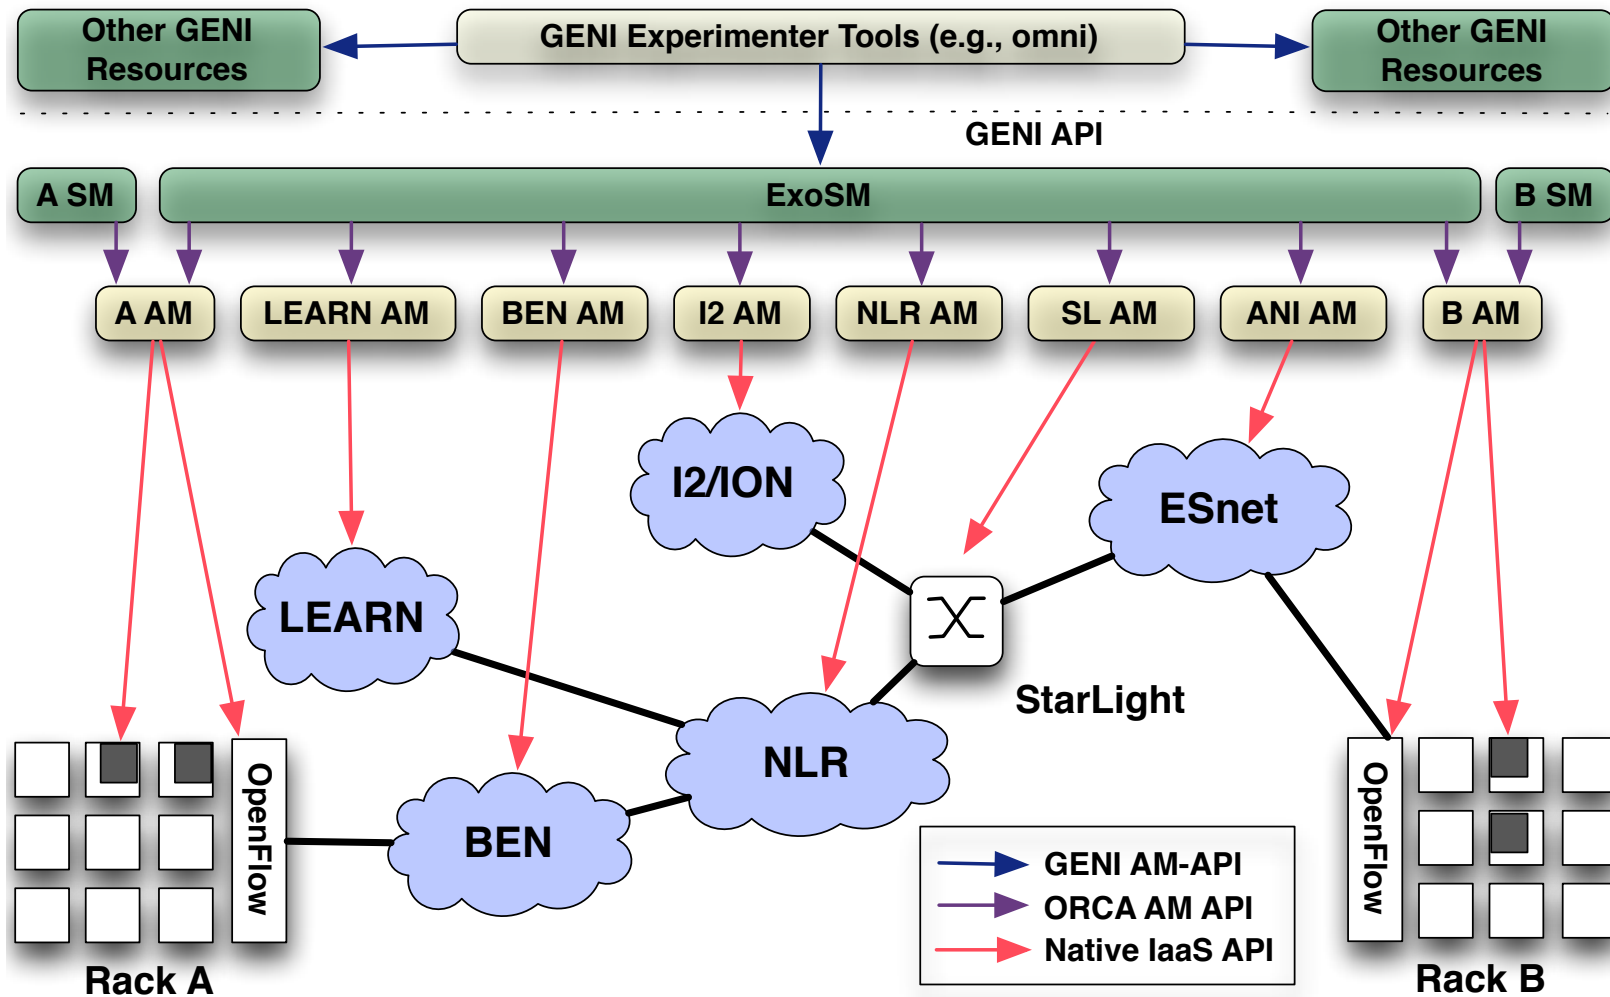
- ▶ 3 racks deployed
 - RENCI, GPO and NICTA
- ▶ 2 existing racks
 - Duke and UNC
- ▶ 2 more racks coming
 - FIU and UH
- ▶ Connected via BEN
(<http://ben.renci.org>),
LEARN and NLR
FrameNet, I2



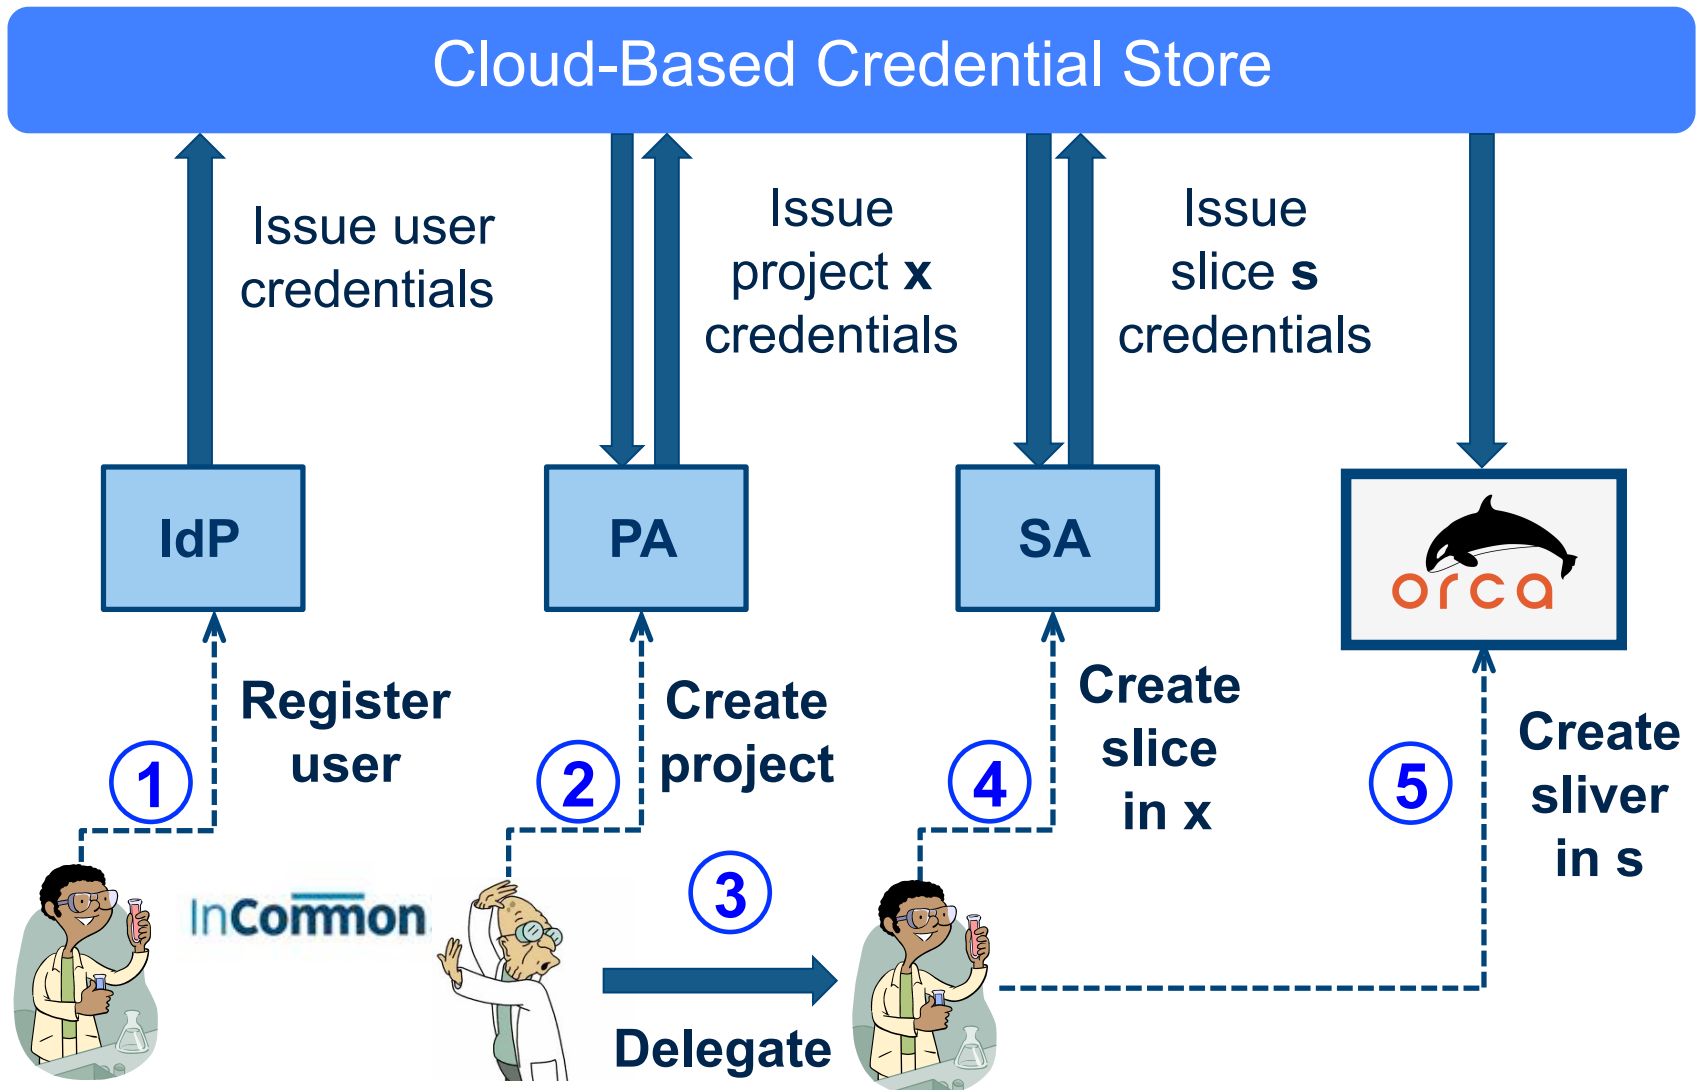
From ORCA to ExoGENI

- **ExoGENI is a facility deployment based on the ORCA framework/toolkit.**
- **ExoGENI instantiates ORCA actors with specific configurations and relationships.**
 - **One AM is configured for each aggregate/provider.**
 - **SMs proxy the GENI API, check identity/access.**
 - **GENI users may access site racks as independent aggregates.**
 - **per-rack SM sees only resources within the rack**
 - **Testbed-wide SMs (ExoSMs) perform topology embedding across all of ExoGENI, acting as a single GENI aggregate.**
 - **Testbed-wide broker assists ExoSMs.**

Deployment structure



Getting GENI credentials



Summary

- **ExoGENI is a network of standard OpenStack cloud sites deployed (deploying) at campuses.**
 - Initial sites centrally managed from RENCI, other providers may join and advertise portions of their resources.
- **Layered orchestration software (ORCA) manages multi-cloud slices and integrates with GENI.**
 - Proxies GENI APIs, checks identity/authorization.
- **Circuit backplane for L2 network connectivity.**
 - By agreement with circuit providers....
- **Configurable/flexible L3 connectivity.**
 - “Easy button” to configure IP network within slice.
 - Host campuses may offer L3 connectivity to slices.