

# Intelligent Data Movement Service

## *Shakedown Experiments*

Ezra Kissel, Akshay Dorwat  
Indiana University

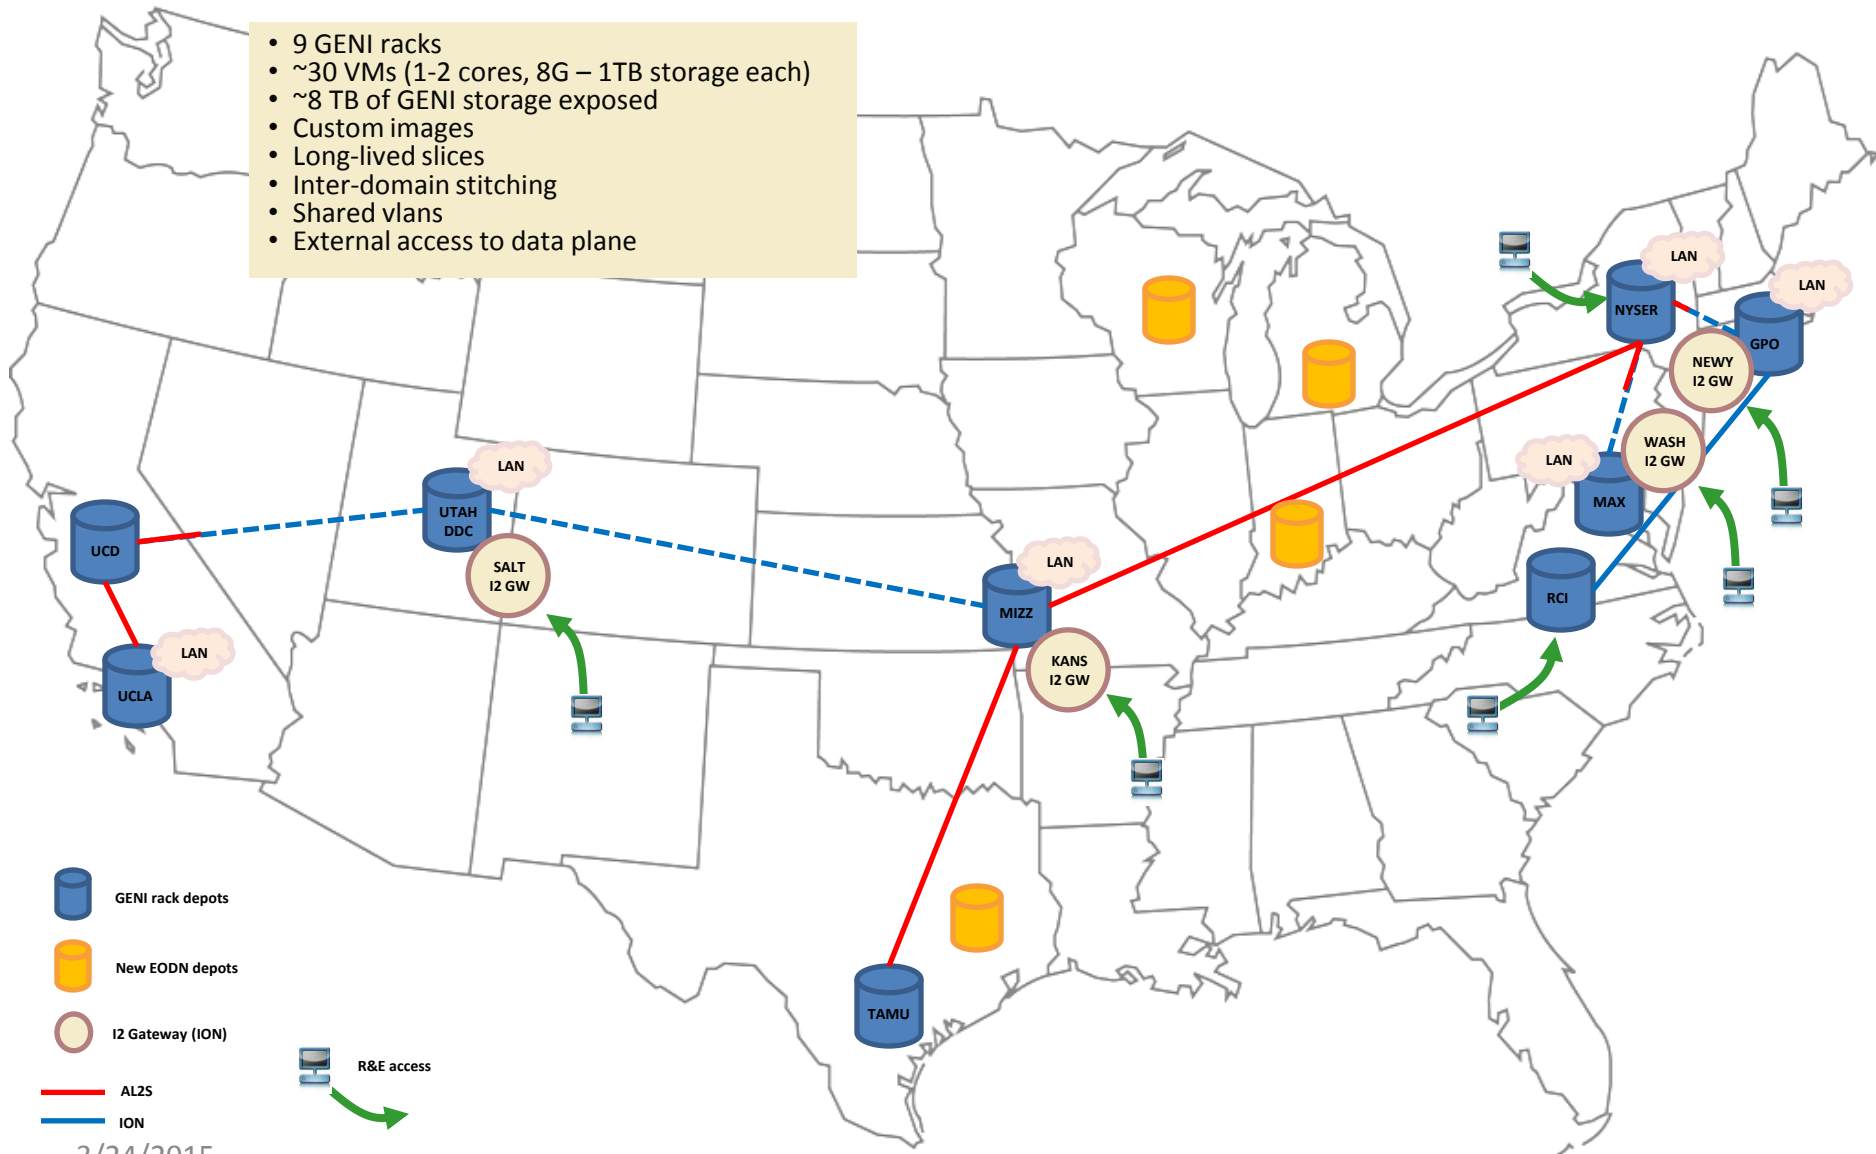
March 24<sup>th</sup> 2015

GEC22 – Experimenter/Developer Roundtable



# GEC22 core IDMS topology

- 9 GENI racks
- ~30 VMs (1-2 cores, 8G – 1TB storage each)
- ~8 TB of GENI storage exposed
- Custom images
- Long-lived slices
- Inter-domain stitching
- Shared vlans
- External access to data plane



# The good/bad list

---

- **Stitching**
  - stitcher.py keeps getting better
  - AL2S and ION working well in 2.8
  - Adding external, non-GENI endpoints is now a clear process
  - Order of operations can sometimes be frustrating with larger requests
    - For example - one hop fails because of resource limitation, must wait for all previous allocated VMs to release
- **Shared vlan POA (PG)**
  - Finding the right nodes to bind to is difficult without supporting lib
    - Maybe genilib does this now?
  - Works reliably once the right combination of steps is followed (the ShareALan HowTo has most of the info now)
  - Issues with VM isolation and policy for dealing with usage could use some thought

# The good/bad list (cont.)

---

- IG/EG interaction
  - Common image process remains unchanged
    - Still an involved and lengthy procedure fraught with danger
  - Increased AL2S connectivity between many racks is promising
  - Would be nice to resolve NDL->RSpec conversion issues
- Stale AM state
  - AM response doesn't reflect reality
  - Occurs frequently on both EG/IG racks (mostly interface/link components)
- The “fighting with GENI” vs. “doing your real work” ratio
  - Improvement in my case: 80/20 → 50/50
  - Still acceptable for shakedown experiments
  - Not so good for quickly evaluating an idea on a non-trivial topology

# End

---

- Contact:
  - ezkissel@indiana.edu
  - adorwat@indiana.edu
- Project Wiki Page
  - <http://groups.geni.net/geni/wiki/sol4/IDMS>
- Thanks to the GPO and the rest of the GENI community