Ubiquitous Edge Platform

Lincoln Bryant
Rob Gardner

US ATLAS Distributed Facilities Meeting, 14 March 2016

ATLAS EXPERIMENT

Ubiquitous & Easy “CI Substrate”

- Pioneer a new phase of advanced cyberinfrastructure deployment, allowing sites to flexibly evolve and sustain both on-premise and commercial cloud-based infrastructure
- Hosted services, such as CEs, data caches, squid, etc., could be centrally deployed onto “CI substrates” within a trusted CI zones and remotely operated, upgraded, and optimized for performance
- Extend to shared, opportunistic university clusters and cloud resources

Distributed Virtualized Data Centers

- Reduce IT footprint and ops burden
  - Centralize deployment & ops; reduce local admin cost
- Explore virtualized data center frameworks
  - E.g. container management over bare metal or VMs
- “Blue sky” goal
  - Establish a “trusted pattern” for a “CI substrate” on sites
  - Create distributed virtualized data center(s) overlaying the fabric substrate

Canonical SciDMZ

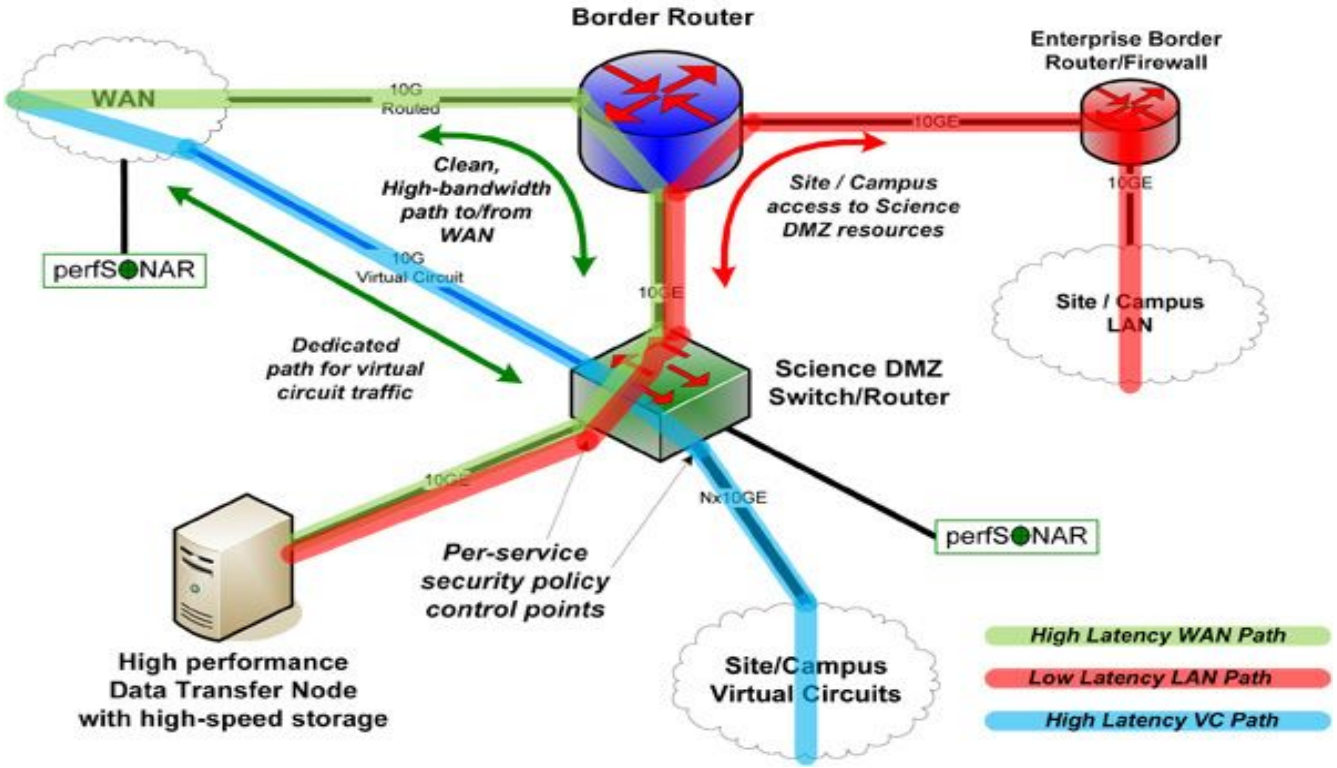

SciDMZ with CI Substrate

US ATLAS Facility central ops console:

$ slate install osg-squid.3.3 --sites MWT2 ALGT2 SWT2 WT2 NET2

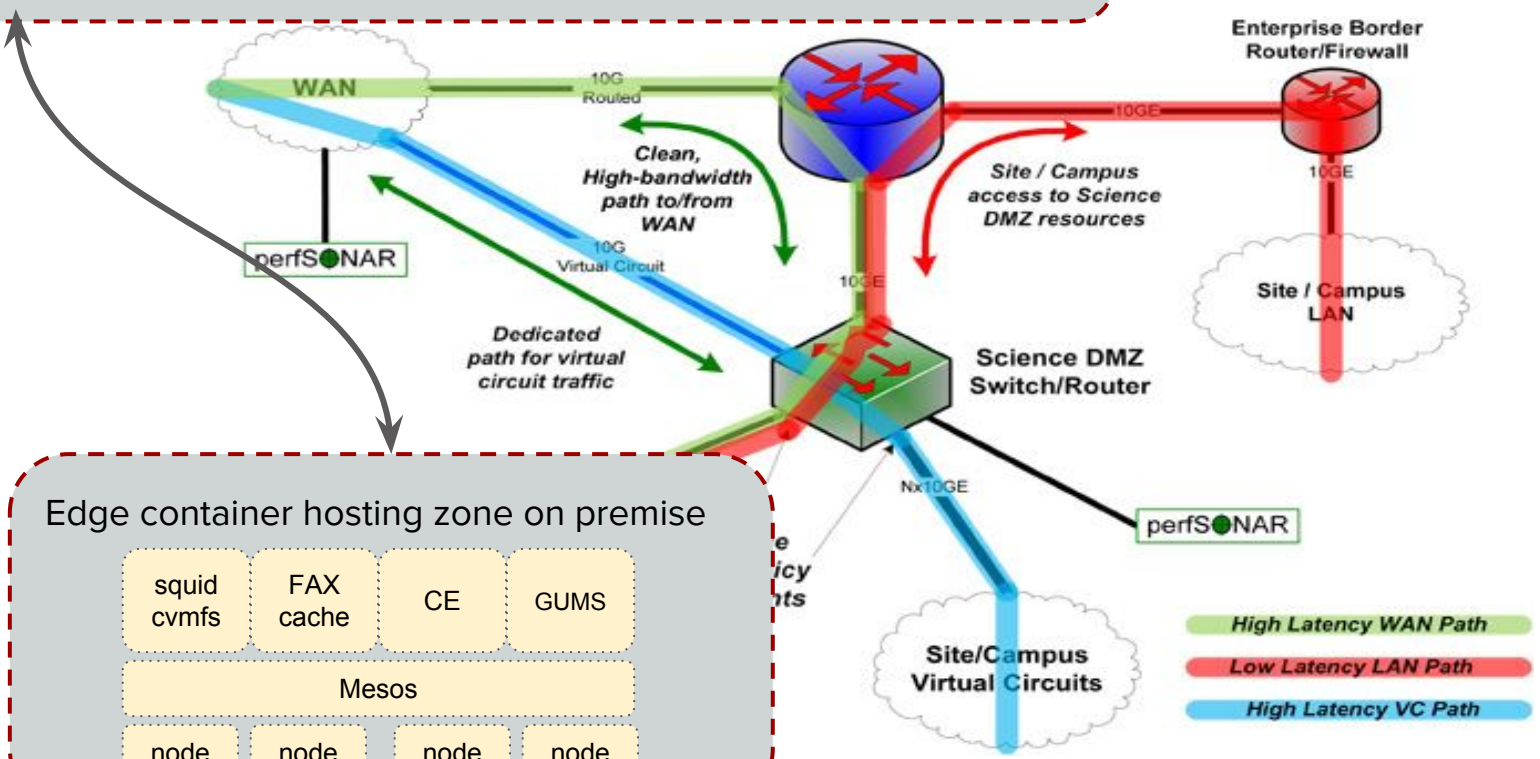

Edge container hosting zone on premise

squid cvmfs	FAX cache	CE	GUMS
Mesos			
node	node	node	node

Deploying research software at the edge

Open Science Grid

perfSONAR

Xrootd Cache

globus

Hardware

- Produce reference specification for supportability reasons
  - No more than 2-3 vendor options.
- Cloud providers like Joyent have done a really good job in this space. Something similar to:
  - https://docs.joyent.com/private-cloud/hardware/specs

Operating system

- Many choices to evaluate in this area
- Traditional distributions:
  - EL, Ubuntu, etc
- Upcoming projects building around containers:
  - CoreOS, Boot2Docker, RancherOS, Project Atomic
- Exotic alternatives:
  - SmartOS (Solaris-based, emulating Linux kernel ABI)

Software

- Microservices-y architecture
  - Follow the Docker model of 1 application per container
- Service discovery and configuration tools
  - Consul, etcd, etc
- Scheduling
  - Kubernetes, Docker Swarm, Fleet, Mesos
  - (HTCondor?)

Software

- Dockerized applications created, vetted, maintained by central operations team.
  - Pushed by operators down to subscribed sites
  - Or, pulled by local site admins without interaction with central.
- Built-in monitoring
  - Graphite, ELK, etc

Containerizing Services

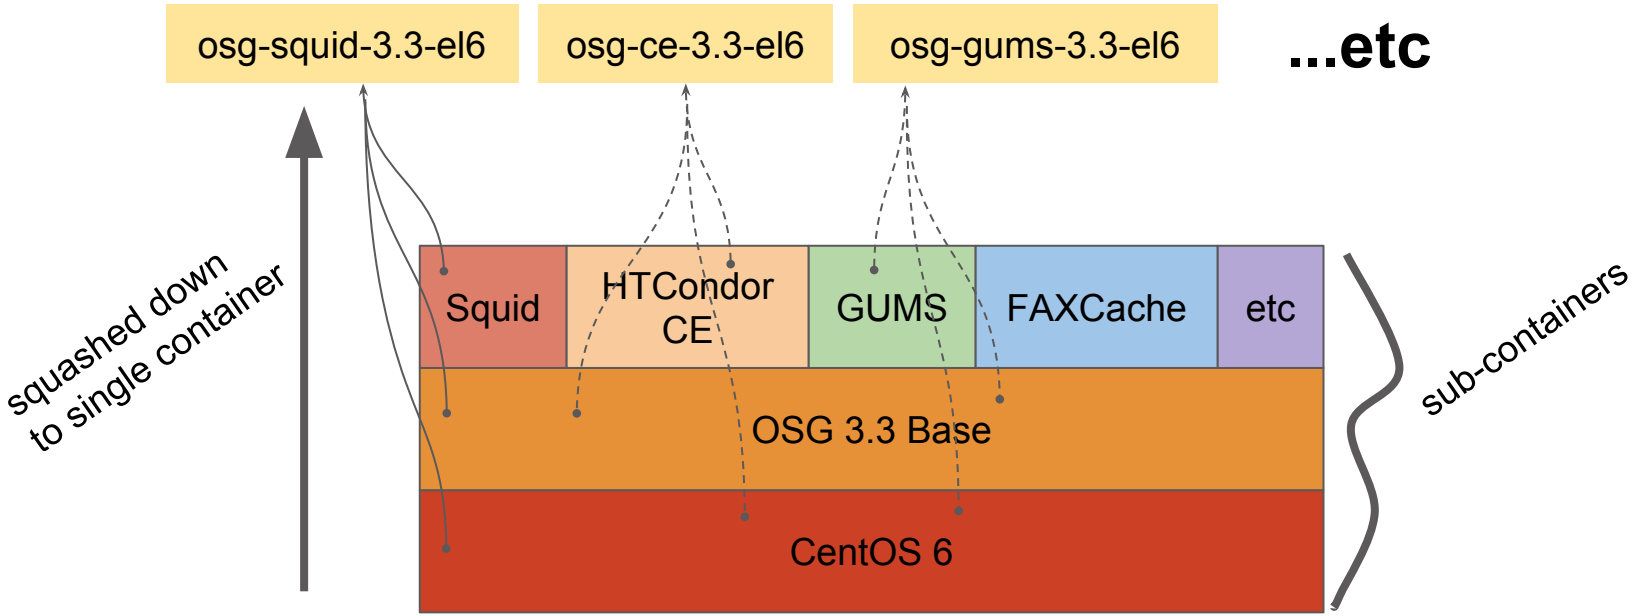

Frontier-Squid Containerized

- Deployed in a hybrid cloud @ Midwest Tier 2:

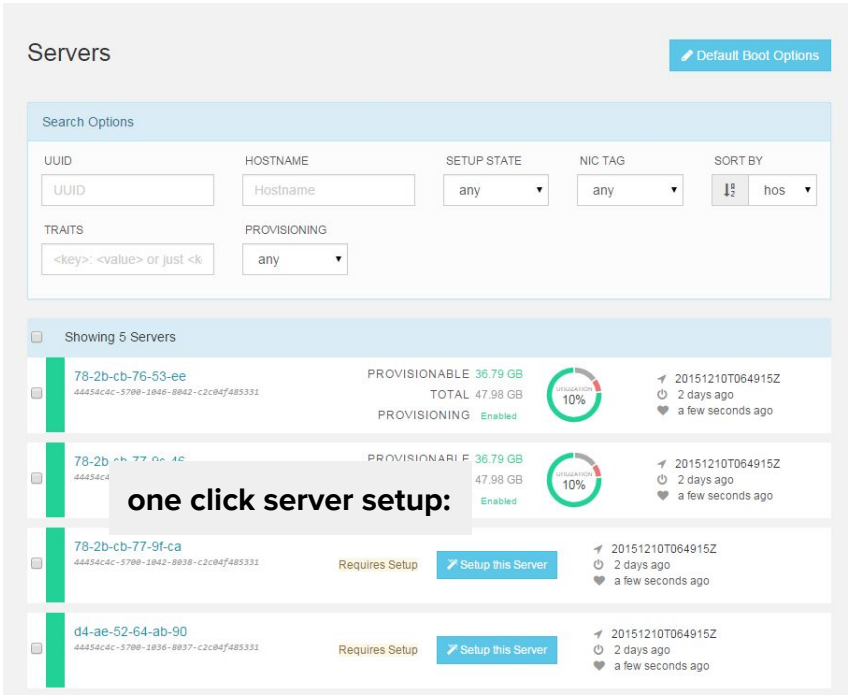

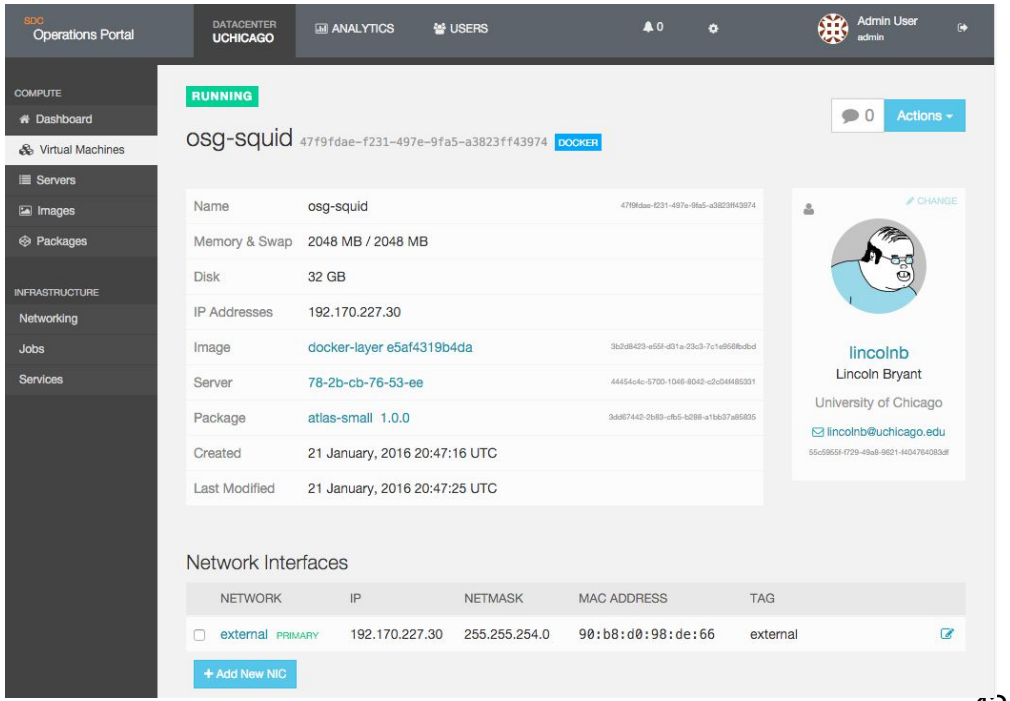

Benefits for ATLAS

- Easily deploy Tier 2 and Tier 3 services
  - PROOF on Demand
  - Remote desktop / NX
  - FAX doors
  - XRootD caches
  - etc

Current pain points

- Many points where human interaction is currently needed
  - Can we automate here?
- Is it possible for me to stand up, then destroy an entire ATLAS site in an automated way?
  - CE, SE, all interactions with AGIS, etc.

Security considerations

- Who has root on the machine?
- Can trusted users allocate resources and start containers remotely without having root?
  - Unprivileged containers are semi-working in newer kernels, but here be dragons..
- Ultimately: What is the correct privilege separation between owner and operator?

Other considerations

- Should there be a VPN / control channel setup such that these nodes are all accessible via the same private IP space?
- Can we use this platform as a testbed for things like SDN?
- What does it look like when we have multiple nodes per site?

Summary

- Platform for edge services on Science DMZs
- Container-based applications, maintained by a central team
- Built-in service discovery, configuration, and monitoring
- Flexible, adaptable to the needs of other projects.

Thank you!
Questions?