

Virtualization on Comet

XSEDE Service Provider Forum – October 13, 2017

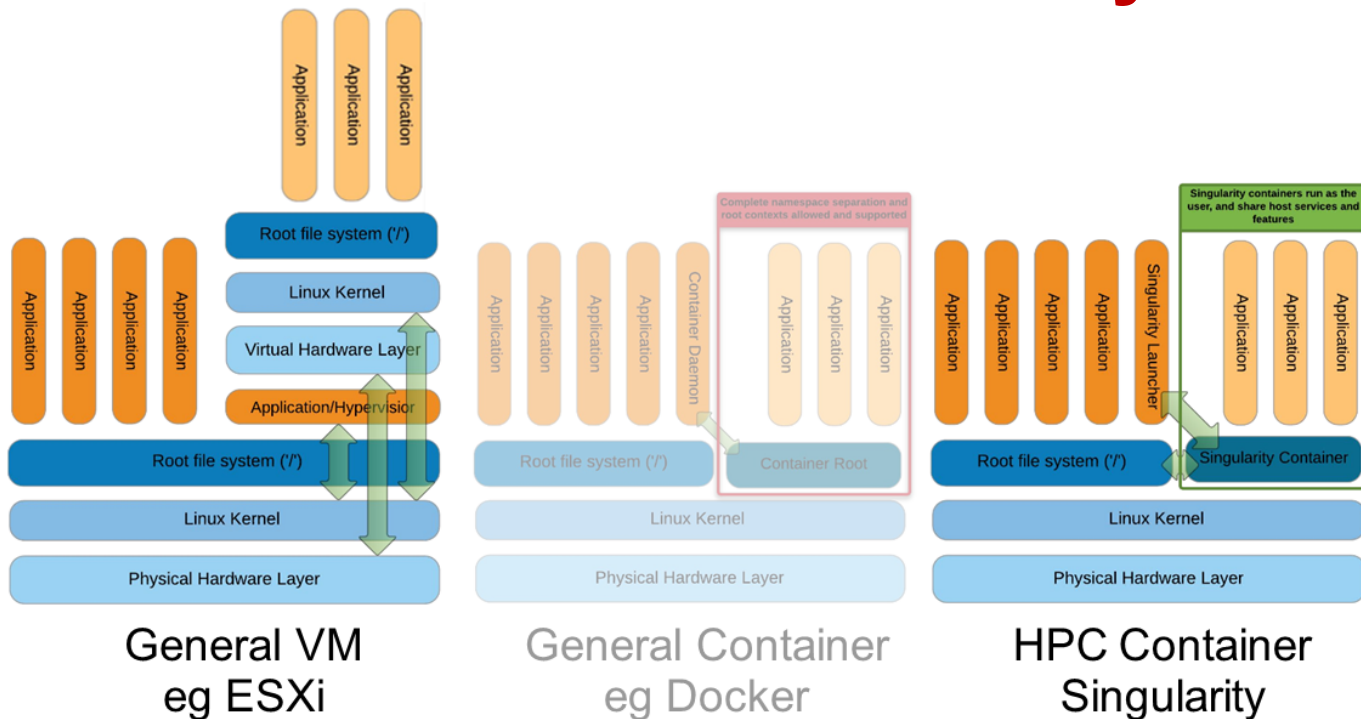
Trevor Cooper, tcooper@sdsc.edu

Mahidhar Tatineni, mahidhar@sdsc.edu

Dmitry Mishin, dmishin@sdsc.edu



Virtualization Three Ways



Source: [Greg Kurtzer keynote at HPC Advisory Council 2017 @ Stanford](#)

Comet Virtual Clusters

- **KVM based Full Virtualization**
- **Full root access, PXE install, persistent disk images, near-native Infiniband performance**
- **Nucleus Rest API and Cloudmesh MGMT**
- **Consume SUs from XSEDE allocations**

Comet VC Use Cases

- **CAIDA Hackathon**

- Root access to nodes for custom OS and software stack.
- Full control of network stack inside virtual compute nodes by attendees and easy 'repair' by CAIDA admins
- Full isolation of virtual cluster from production resources and filesystems

- **Open Science Grid**

- Simple install using existing management infrastructure (PXE, Foreman, Puppet)
- Multiple XSEDE allocations consuming SUs via OSG VC with no effort from allocated projects
- Largest OSG provider of resources (> 2x) for last LIGO run

Comet Gateway Frontend Hosting

- **KVM based Full Virtualization**
- **Full root access, PXE install, persistent disk images, modest resource requirements**
- **Nucleus Rest API and Cloudmesh management**
- **Appropriate for Gateways running compute jobs on Comet**
- **Possibilities for SDSC onsite storage outside of Comet**

Comet Singularity

- **HPC Application Containerization**
 - Userspace only (with namespaces and minimal SetUID)
- **Singularity v.2.3.2 is available now**
 - support for 'pull' from Docker and Singularity Hub
- **Singularity v.2.4 is being tested**
- **Utilization growing quickly...**

Comet Singularity Use Cases

- **Special Requirements**

- Frequently updated or requested application installs that are difficult in Comet OS (TensorFlow, Torch, Caffe, Keras)
- Commercial application binaries incompatible with Comet base OS

- **Neuroscience Gateway**

- Custom software stack difficult to install or support in Comet OS (TensorFlow, GENESIS, Neuron)
- NSG allows some custom user codes... constrain user codes from modifications of NSG environment.

Comet Kubernetes

- **Generalized Container Orchestration**
 - Currently in development testing
 - Not (currently) for users
- **Use Cases Under Investigation**
 - Backend Services (MySQL, RabbitMQ)
 - User facing Services (Slurm, GridFTP)

Comet Allocations for Virtualization

	Trial	Startup	XRAC
Virtual Cluster	- na -	50 K SUs	50 K – 7,000 K SUs
Singularity	1000 SUs	50 K SUs	50 K – 6,000 K SUs

***Trial requests typically handled in 1 business day.
Startup requests reviewed locally. Can be approved in a few days.
XRAC requests reviewed by standard XSEDE process.***

Comet User Support

- **Virtual Cluster Support**

- Frontend support to Virtual Cluster Admins by Indiana University team which interfaces with HPC Operations for backend issues
- Basic installation assistance, Cloudmesh support, Nucleus API and Cloudmesh feature requests

- **Singularity Support**

- SDSC HPC User Services Staff via XSEDE and local ticketing systems
- Building and testing containers on Comet hardware (ie. P100)
- Integration of containers with Scheduler (ie. example containers)

Resources

- **Cloudmesh Comet CLI for Comet VC**
 - http://cloudmesh.github.io/client/commands/command_comet.html
- **Singularity – Containers for Science**
 - http://www.hpcadvisorycouncil.com/events/2017/stanford-workshop/pdf/GMKurtzer_Singularity_Keynote_Tuesday_02072017.pdf
- **Docker vs Singularity vs Shifter vs UGE Container Edition**
 - https://tin6150.github.io/psg/blogger_container_hpc.html
- **Performance Evaluation of Container-based Virtualization for High Performance Computing Environments (paper)**
 - <https://arxiv.org/abs/1709.10140>
- **Comet Singularity Examples**
 - <https://github.com/mkandes/naked-singularity>