

Containers, Clouds & Combos

Emerging Delivery Models for HPC

Jonathan Muller
Team Lead TPSI

Agenda

- Cluster Workload
- Bright Architecture
- Clouds
- Containers



Life of HPC administrator up until recently

- Manage some compute & storage nodes
- Use batch scheduler to schedule MPI jobs



Life of HPC administrator today

- New types of workload and use-cases for clusters
- Various ways of hosting workload



Cluster Workloads

- Workload types:
 - Traditional compute jobs (e.g. MPI)
 - Hadoop / Spark (i.e. long running services)
 - Deep neural networks (e.g. DIGITS user front end)
 - Long running services that need to scale up/down
- Hosting dimensions:
 - Physical versus virtual machines
 - Containerized versus non-containerized
 - On-premise versus off-premise (i.e. in public cloud)






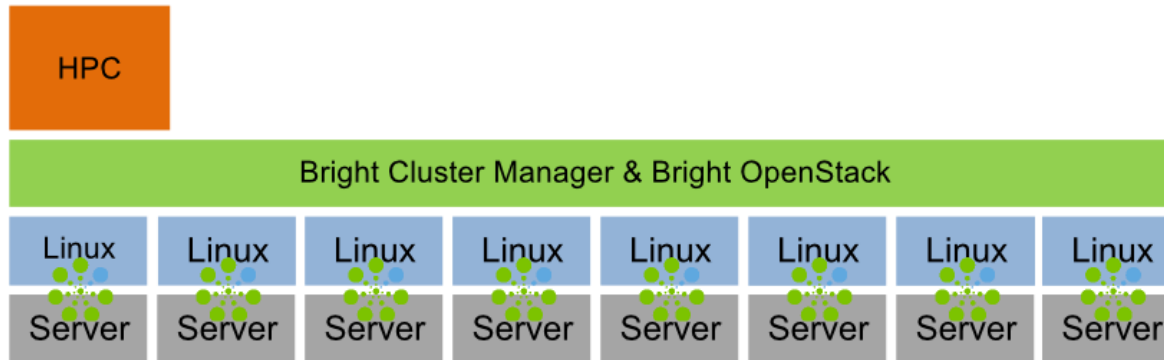
Bright's **mission** is to provide the best software platform for deploying, managing, and monitoring clustered infrastructure, in the datacenter or in the cloud.

Bright Architecture

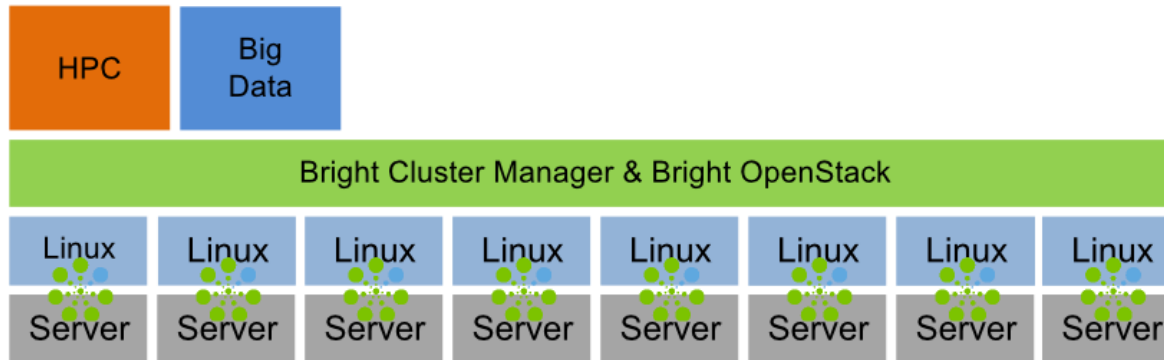
Bright Cluster Manager Architecture




Bright Cluster Manager Architecture



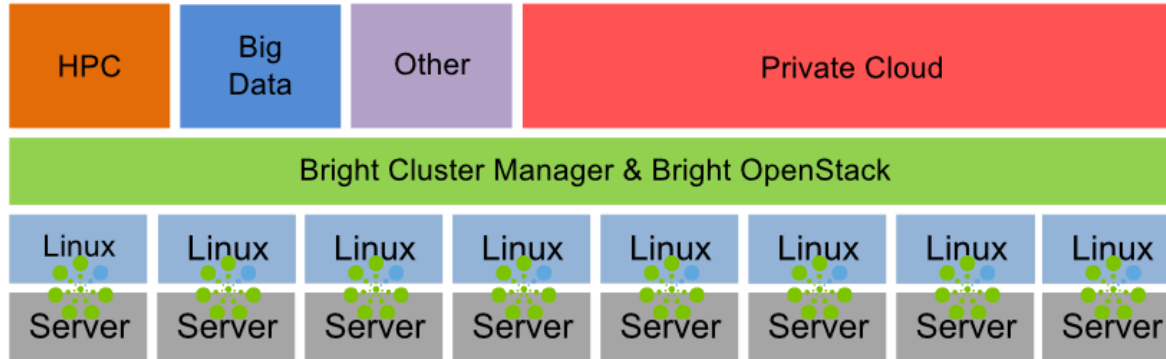
Bright Cluster Manager Architecture




Bright Cluster Manager Architecture



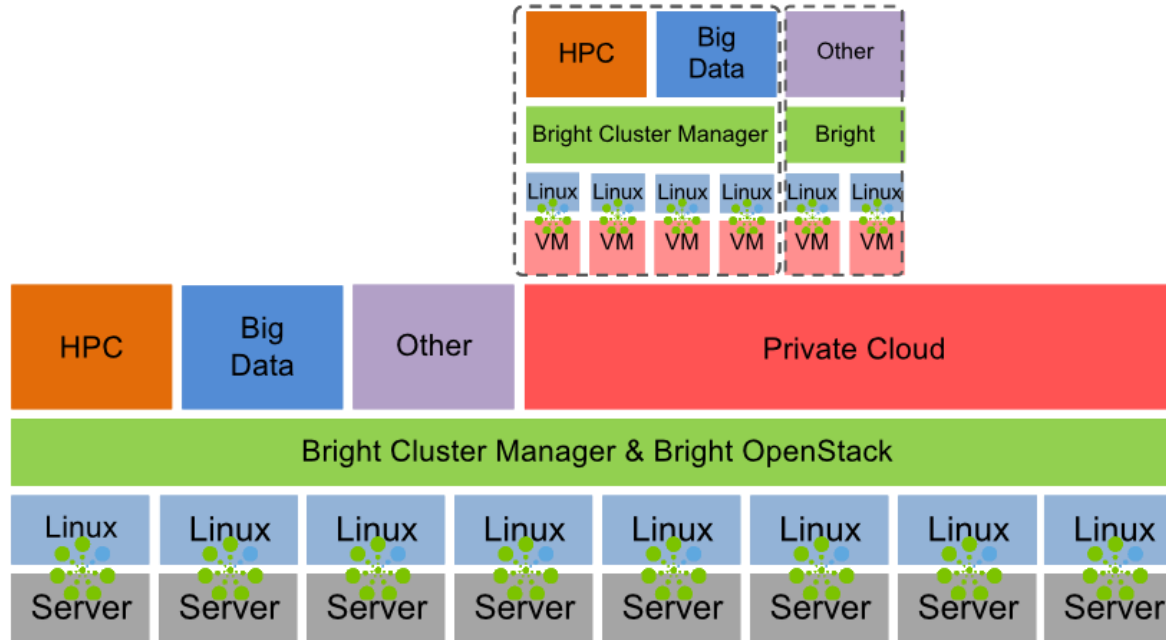
Bright Cluster Manager Architecture




Bright Cluster Manager Architecture




Bright Cluster Manager Architecture



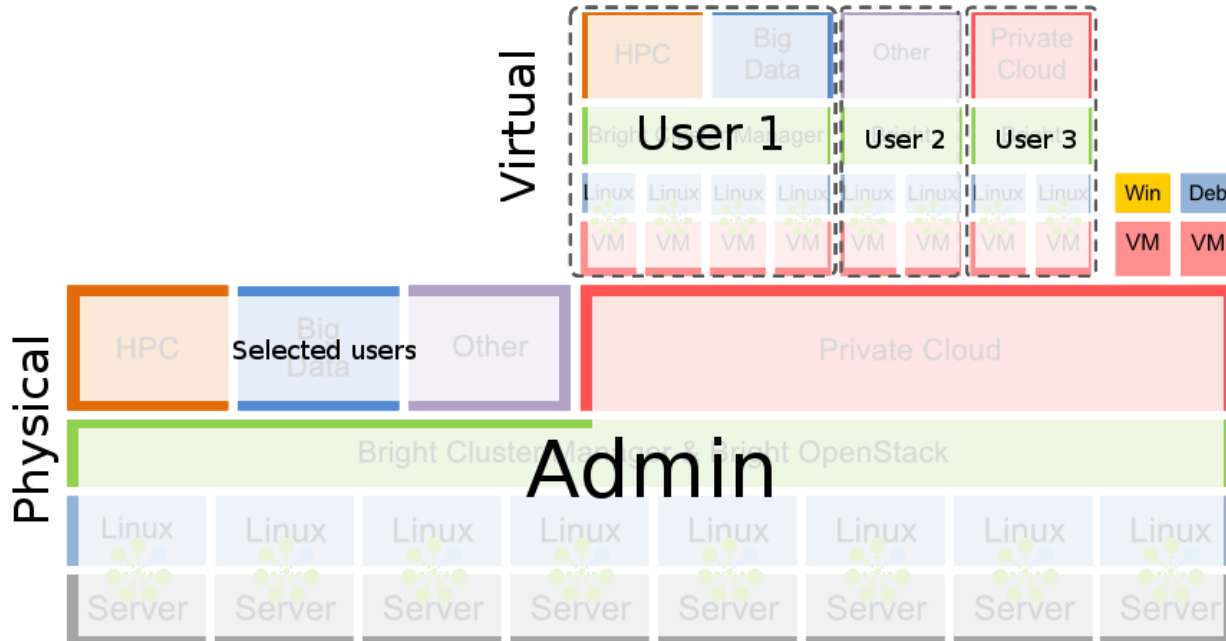
Bright Cluster Manager Architecture



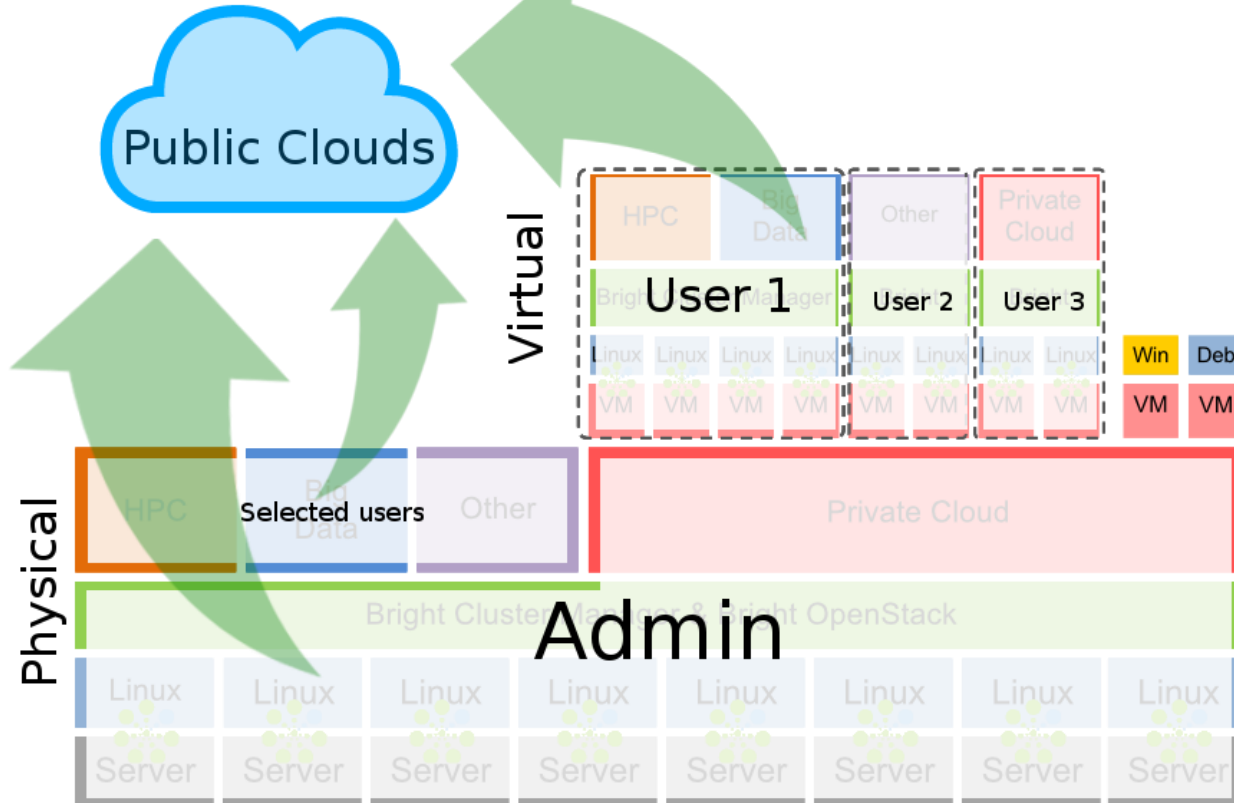
Bright Cluster Manager Architecture



Bright Cluster Manager Architecture

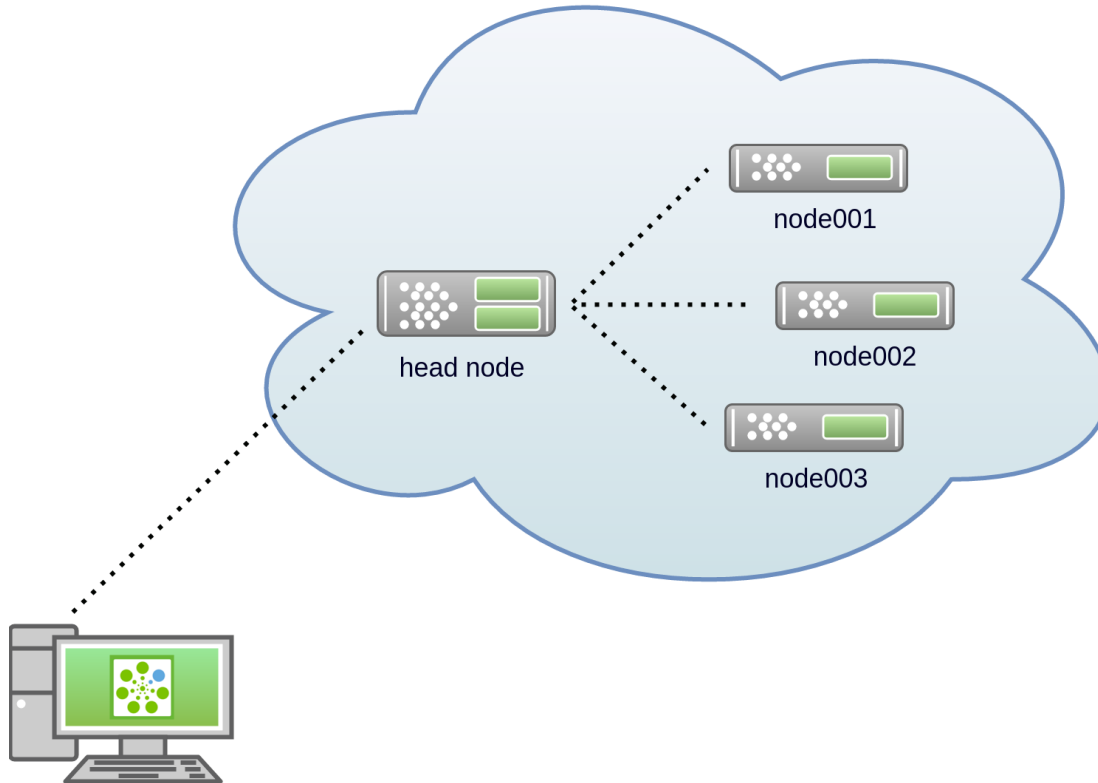


Bright Cluster Manager Architecture

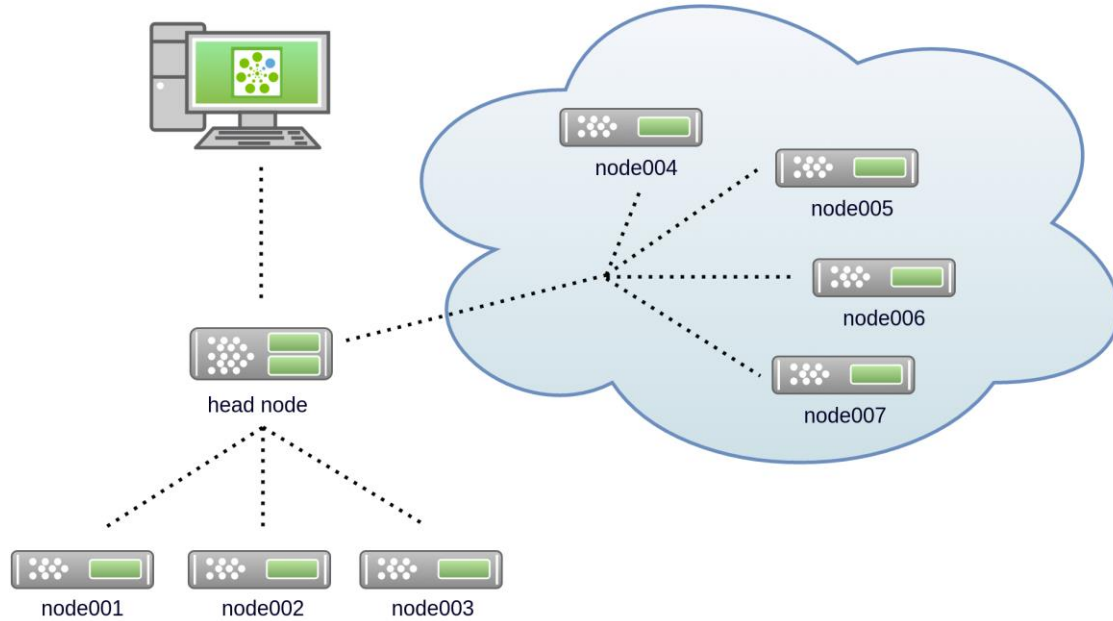


Cloud Bursting

Cluster on Demand



Cluster Extension



Moving Workload to Cloud

Bright Cluster Manager provides:

- Uniformity between cloud and on-premise
- Single workload management system setup
- Auto-scaling of nodes based on workload
- Transparent moving of input/output data



Containerization

Container 101

- Containerization: method of running applications in isolated environment
- Relies on Linux kernel features:
 - Namespaces (PID, network, IPC, ...) for *isolation*
 - cgroups for *restricting* resource usage
 - UnionFS
- Container image: root filesystem snapshot which serves as template
- Container: instance of an image (normally active)
- Docker: tool to manage containers on single host
- Kubernetes: cluster-wide container orchestration tool



Containers in Bright

- Container support since Bright 7.2:
 - Docker v1.9
 - Kubernetes v1.1.2
- New in Bright 7.3:
 - Docker v1.11 (containerd and runC)
 - Kubernetes v1.3.0
 - Flannel networking integration
 - Docker Registry
 - Singularity

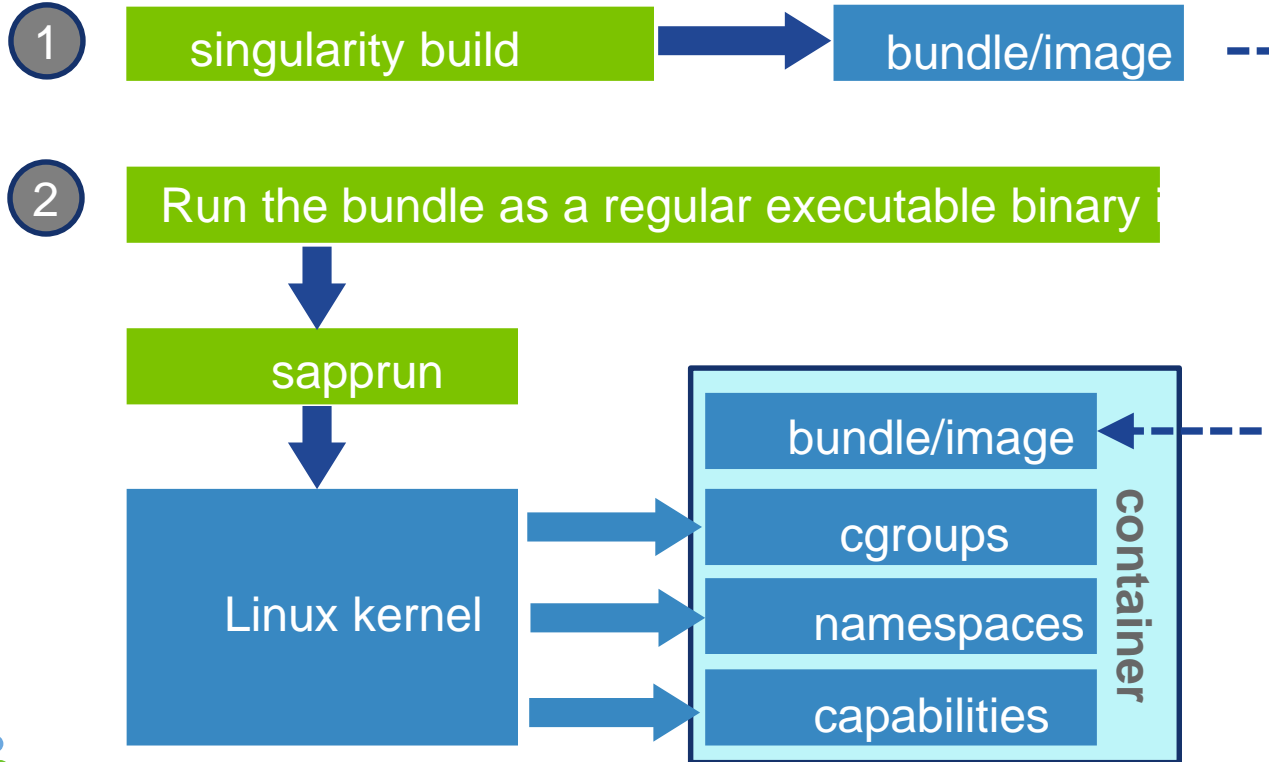


Singularity

- Dangerous to allow ordinary users to start Docker containers:
 - Root on the network
 - Image poisoning
 - Easy to become root on host by mapping host filesystem into container
- Docker not suitable for HPC jobs on multi-user system
- Singularity: allow containers to be used for HPC applications
- Primary goal: application portability
- Packages up application, dependencies and optionally input data
- Provide isolation but not resource restriction
- Possibilities for cloud bursting



Singularity Architecture



Containerization Use Cases

- HPC jobs: Singularity
- Per user web front-ends (e.g. DIGITS): Kubernetes
- Software development/testing: Docker with local Docker registry
- Long running services to be scaled dynamically: Kubernetes
 - Alternative: Mesos with Marathon
- Non-parallel jobs that finish at some point: Kubernetes
 - Alternative: Mesos with Chronos



Conclusion

Bright Cluster Manager allows HPC workload to be run:

- On-premise, off-premise or hybrid
- Inside containers, or outside containers
- On physical or on virtual machines



Q&A and huge Thank you!

Jonathan Muller, Container Integration Lead
jonathan.muller@brightcomputing.com

We tweet at [@BrightDevTeam](https://twitter.com/BrightDevTeam) [@BrightEMEA](https://twitter.com/BrightEMEA) and [@BrightComputing](https://twitter.com/BrightComputing)

Visit us at www.brightcomputing.com/products

Read our new ebook on Data Center Technology for the Life Sciences
<http://bit.ly/ls-ebook>



