TASHI: OPEN-SOURCE CLOUD COMPUTING

Jim Cipar, Greg Ganger, Julio Lopez, Michael Stroucken, Dave O'Hallaron, Michael Kozuch (Intel), Richard Gass (Telefonica)

TASHI AND OPENCIRRUS

Tashi: an Apache Incubator project

- VM-based cluster management
- Architected for Big Data activities
- Flexible, to enable scheduling research

Zoni: bottom layer of Tashi

- Dynamic physical resource provisioning
- Enables experimentation on raw hardware

OpenCirrus: global cloud testbed

- 15 sites worldwide
- Supporting cloud systems research
- ... and many academic cloud users
- PDL, Intel, and MIMOS sites run Tashi

PDL'S OPENCIRRUS HARDWARE

- 78 nodes, 624 cores, 1.6 TF, 0.3 PB, 1.2 TB
- Node configuration:
  - CPU: 2x quad-core Intel Xeon 2.66 GHz
  - Local storage: 1 TB SATA
  - Netapp VM image storage
  - RAM: 16 GB
  - Network: 2 x 1-GE
- 4x 48-port Gigabit Ethernet switches
- 8x 10-GE links to the DCO Backbone

ASSOCIATED SYSTEMS RESEARCH

- Power management
  - Lower power modes for idle machines
  - Dynamic work scheduling to maximize benefits
- Elastic distributed storage
  - Enabling temporary reuse of nodes
  - With full data availability and balanced load
- Collaborative two-level scheduling
  - Supporting heterogenous workload mixes
  - Batch frameworks, elastic services, etc.

- See related posters for more details

TASHI ARCHITECTURE

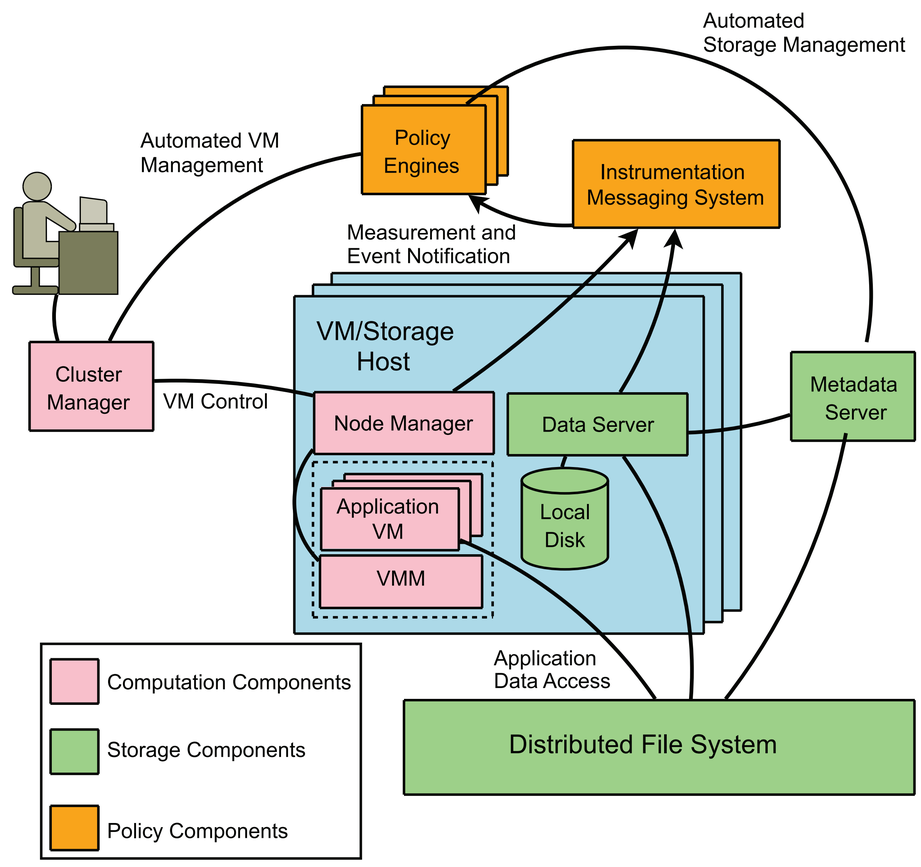

- Cluster manager provides mechanism for managing VMs
  - Keeps track of running VMs, watches for state changes
  - Users, apps., and policy engines all act as clients
  - Client calls CM to start, stop, pause, migrate VMs
- Node manager provides common VMM management
  - Converts calls from CM to native Xen/KVM commands
  - Observes and reports performance metrics

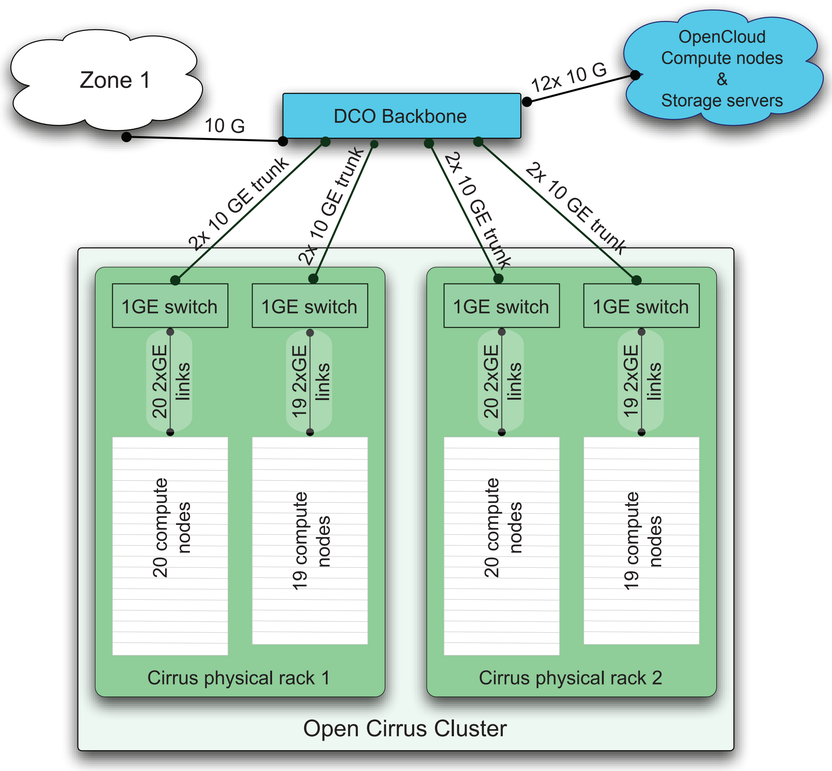

Intel Science & Technology Center for Cloud Computing

Carnegie Mellon University · Georgia Tech · intel · PRINCETON UNIVERSITY · UC Berkeley