Karlsruhe Institute of Technology

Building a Private Cloud with Eucalyptus

5th IEEE International Conference on e-Science – Oxford – December 9th 2009

Christian Baun, Marcel Kunze

Forschungszentrum Karlsruhe
in der Helmholtz-Gemeinschaft

Universität Karlsruhe (TH)
Forschungsuniversität • gegründet 1825

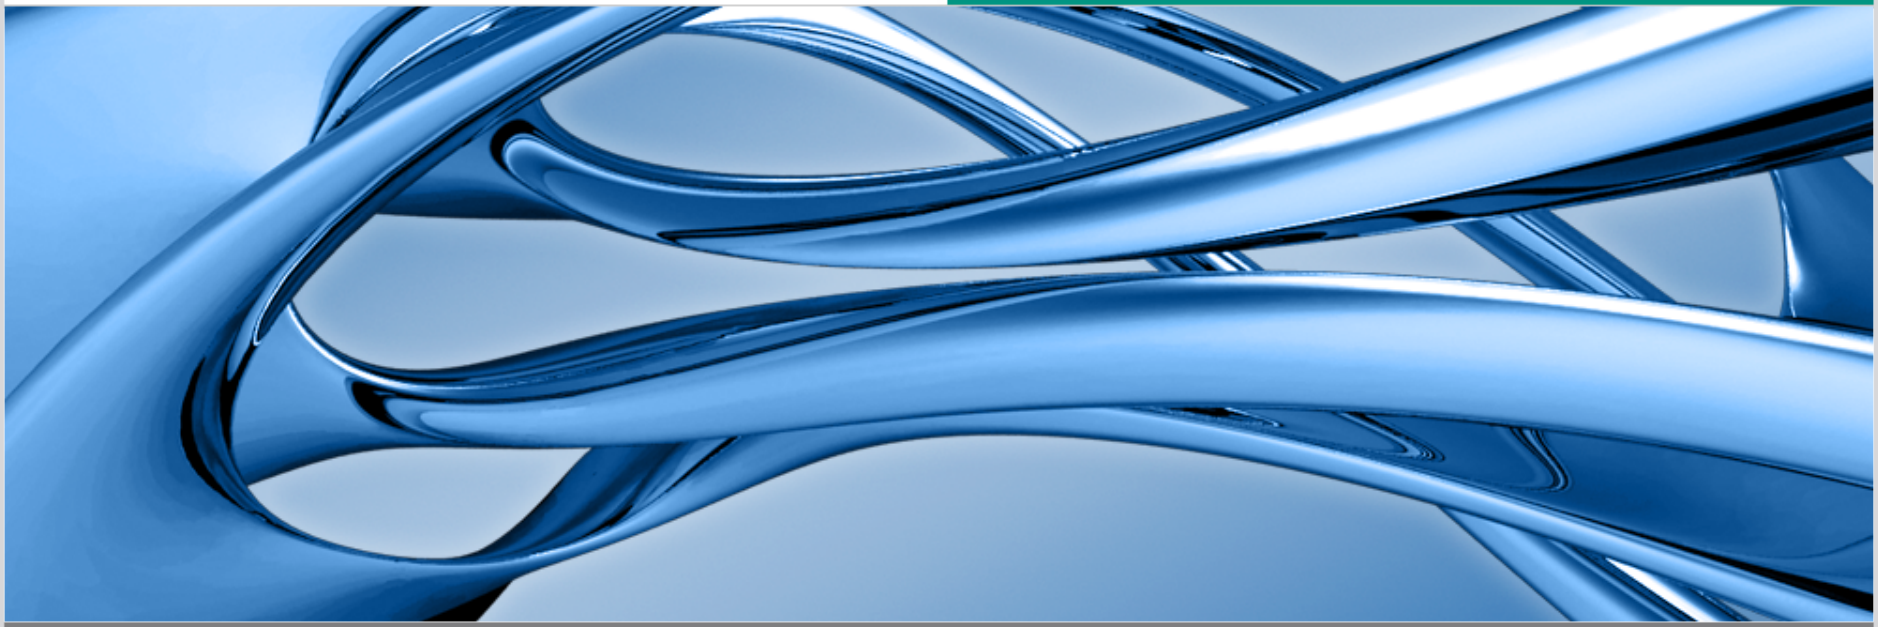

Eucalyptus

http://open.eucalyptus.com

- Open Source software infrastructure for implementing cloud computing on clusters
- Developed at UC Santa Barbara. Now, Eucalyptus Systems, Inc.
- EUCALYPTUS - Elastic Utility Computing Architecture for Linking Your Programs To Useful Systems
- Implements Infrastructure as a Service (IaaS) – gives the user the ability to run and control virtual machine instances (Xen, KVM) deployed across a variety of physical resources
- Interface compatible with Amazon EC2
- Includes „Walrus“, an implementation of Amazon S3 storage service
- Includes „Storage Controller“, an implementation of Amazon EBS (persistent) storage service
- Potential to interact with the same popular tools, known to work with Amazon EC2, S3 and EBS e.g. S3 Curl, ElasticFox, s3cmd, ...
- Eucalyptus is an important step to establish an open cloud computing infrastructure standard

Eucalyptus (v1.5.2) – Components

http://open.eucalyptus.com

- Cloud Controller (CLC)
  - Includes S3 and EBS services
  - Collects resource information from the CC
  - Operates like a meta-scheduler in the Cloud
- Cluster Controller (CC)
  - Schedules the distribution of virtual machines to the NCs
  - Collects (free) resource information from the NCs
- Node Controller (NC)
  - Runs on every node in the Cloud
  - Xen-Hypervisor or KVM running
  - Provides resource information to the CC

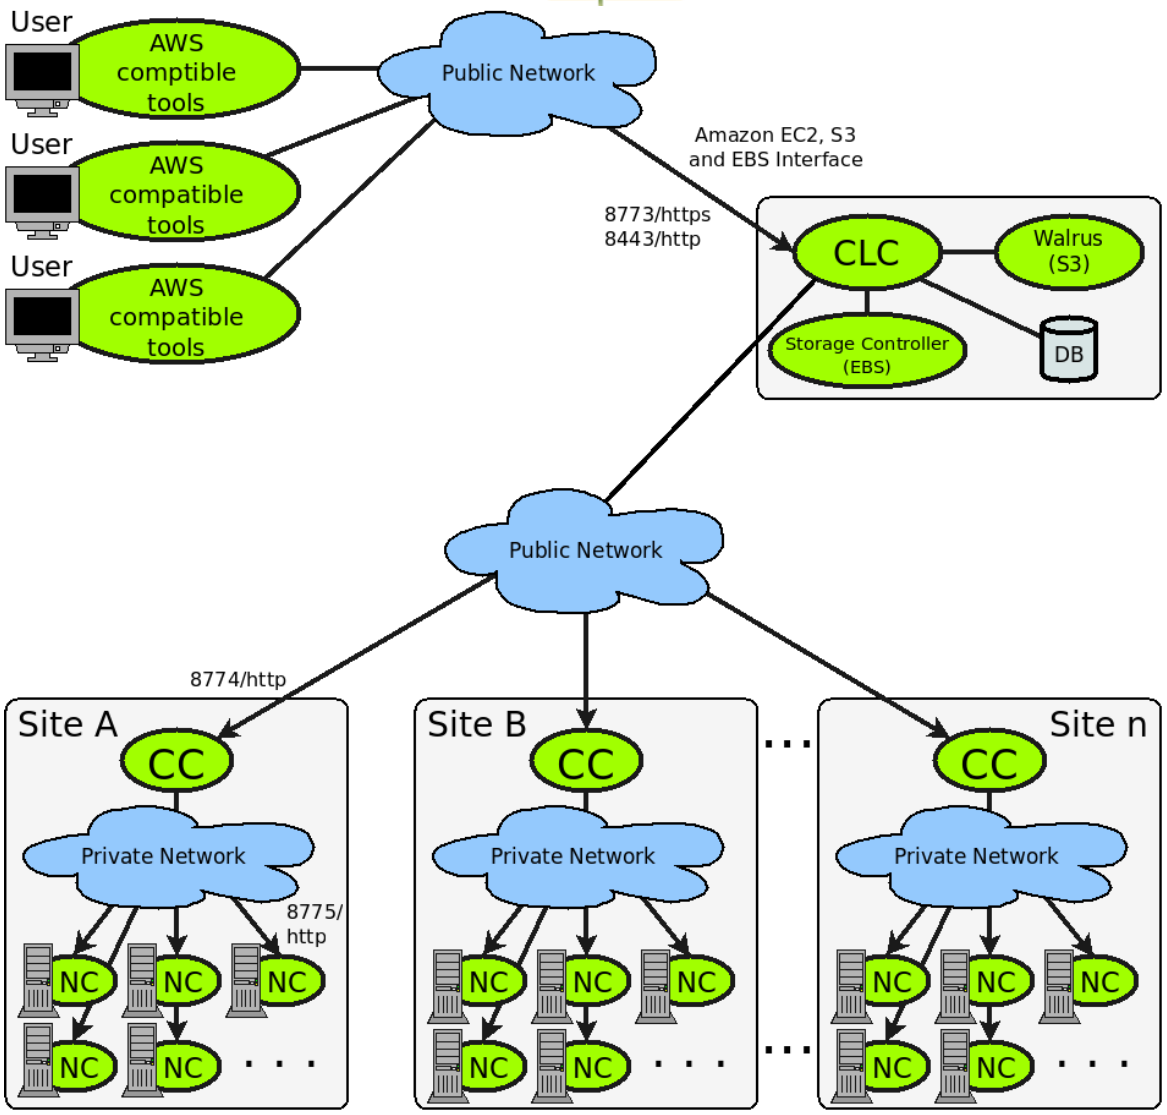

Eucalyptus at KIT

- In 2009 different installations with Eucalyptus 1.4, 1.5.x and 1.6.x have been build up and tested at KIT
- Private Cloud Installation with Eucalyptus 1.5.2 (July – October 2009)
  - Used to gain experience and for performance measurements
  - HP Blade Center c7000
  - 5x HP Blade ProLiant BL2x220c G5
  - 2 Server per blade:
    - 2x Intel Quad-Core Xeon (2,33GHz)
    - 16GB RAM

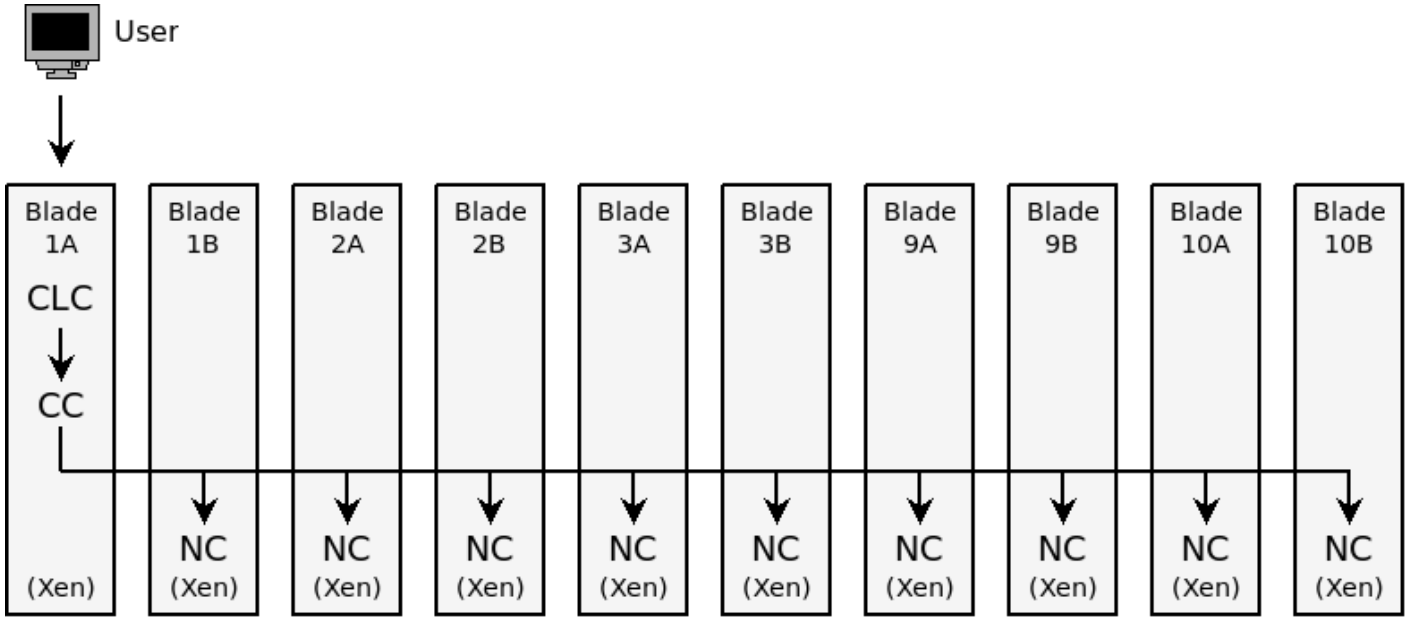

Storage Performance: S3/EBS vs. Eucalyptus

- Measurements made:
  - Sequential output/input (per character and blockwise)
  - Random seeks and file creation/deletion
- EC2 stores the images/instances in S3 possibly distant from EC2
  - Eucalyptus stores the images in S3 (Walrus) too, but the instance data is stored at the NC locally
- Eucalyptus' sequential read/write performance is weaker than EC2/S3
- Eucalyptus' random seeks and file creation/deletion performance is better
  - Caching behavior of the Linux Kernel helps
- EBS performance of Eucalyptus is weak
  - EBS files are stored at the CLC
  - All access goes over the network
- Eucalyptus' blade server contain SATA hard disks with 7200 RPM
- Storage performance in the private cloud can be improved with fast hard disks, RAID, SAN, SSDs ...

CPU Performance: EC2 vs. Eucalyptus

- Performance was measured with Linux Kernel compilation
  - All instance types have been measured
- Additional RAM and CPUs lead to a performance boost when using more threads
- Eucalyptus performs much better for `m1.small`
  - Because the CPU differs
  - All instance types use the same physical servers
  - Instances run separately on the blades to avoid interferences
- For EC2, `m1.small` instances perform awfully bad compared to the other instance types
- We cannot make any assertion about the load of the physical hosts at Amazon during our CPU testing
- Using more threads than virtual/physical CPUs/cores available is not advisable due to the thread context switching overhead

CPU Performance: EC2 vs. Eucalyptus

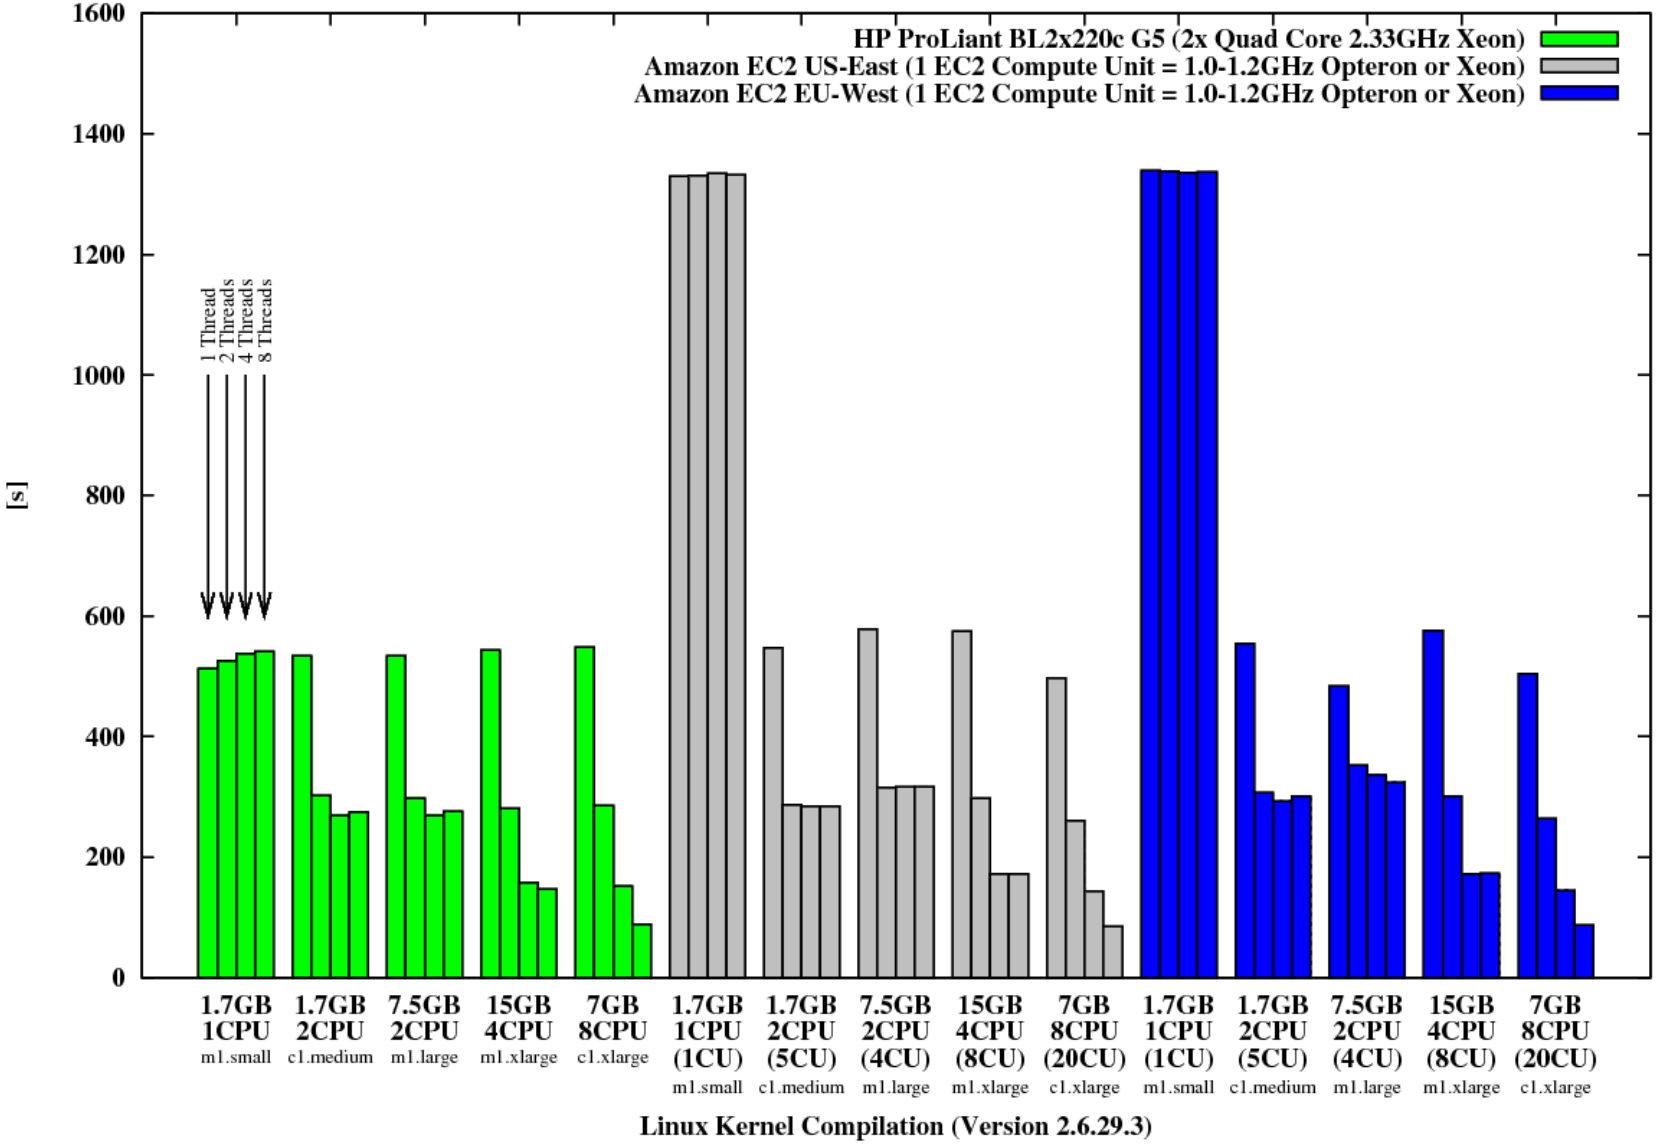

Network Transfer Rate

- The network transfer rate between and inside Amazon EC2 and Eucalyptus was measured
- Transfer rate inside the Eucalyptus infrastructure is best
  - 1000 Mbit/s Ethernet bandwidth between the blade servers
- Transfer rate inside EC2 imply there is 1000 Mbit/s Ethernet used as well, but with a higher workload

Scenario	Transfer Rate [MB/s]
EC2 EU West measured from KIT	11,7
EC2 US East measured from KIT	6
EC2 EU West locally	78,9
EC2 US East locally	84,6
Between EC2 EU West and EC2 US East	10
Eucalyptus at KIT locally	113,1

Network Latency – Round-Trip Time (RTT)

- **Is Cloud ready for HPC application?**
  - **RTT was measured**
  - **For almost all scenarios `m1.small` performs worse compared to the other instace types**
- **For MPI jobs where loosely coupled tasks compute a few seconds it seems possible to use cloud infrastructures**

Scenario	Round-Trip-Time [ms]
Amazon EC2 measured from KIT	102 – 106
EC2 sites locally	0.3 – 0.5
Between both zones in EC2 EU West	1 – 2
Between EC2 EU West and EC2 US East	87 – 92
Eucalyptus at KIT locally	< 0.2

Next steps

- Upgrade to Eucalyptus 1.6 (done)
- Connect the Eucalyptus cluster with SAN and use the benefits of SAN storage where it is useful (done)
- Increase the stability of the infrastructure (in progress)
- Find typical administration scenarios of an productive cloud environment and test them with Eucalyptus 1.6 (in progress)
- Using the infrastructure as a productive environment (in progress)
- Deployment of Eucalyptus components (Cloud Controller, Cluster Controller, S3 and EBS) inside virtual machines (in progress)
- Migration of Grid services in the infrastructure (todo)

Now available...

- Available since November 2009
- First cloud computing book in German language
- Covers the latest topics in cloud computing
- Only € 14,95

Thank you for your attention