Performance Measurement of a Private Cloud in the OpenCirrus^{TM} Testbed

Christian Baun and Marcel Kunze

Steinbuch Centre for Computing,
Karlsruhe Institute of Technology,
Herrmann-von-Helmholtz-Platz 1,
76344 Eggenstein-Leopoldshafen, Germany
baun,kunze@kit.edu

Abstract. Cloud computing realizes the advantages and overcomes the restrictions of the grid computing paradigm. Elastic infrastructures can be easily created and managed by cloud users. In order to accelerate the research on data center management and cloud services the OpenCirrus[1] Research Testbed has been started by HP, Intel and Yahoo!. Although commercial cloud offerings are proprietary, an Open Source solution exists in the field of IaaS with Eucalyptus. This paper examines the I/O and CPU performance as well as the network transfer rate of cloud computing infrastructures implemented with Eucalyptus in contrast to Amazon EC2/S3.

1 Cloud Computing – An Upcoming Trend in IT

During the last years with the support of public funding, grid computing evolved from a computer scientists' field of research to a common working environment for scientific disciplines like physics, medicine and meteorology. A grid definition from Ian Foster and Carl Kesselman, summarizing the focus of grids is:

> A computational grid is a hardware and software infrastructure that provides dependable, consistent, pervasive, and inexpensive access to high-end computational capabilities. [1]

At the same time another trend became imminent in the commercial IT sector: Cloud computing [2] aims at consolidating hardware and software resources in large data centers. Cloud computing realizes the advantages and overcomes the restrictions of the grid computing paradigm [3].

All resources are marketed by providers as a service over the Internet based on a utility model. In the world of cloud computing dynamically scalable (elastic) infrastructures can easily be created and managed by the user. Only the consumed resources are accounted following the pay-as-you-go principle. Many different cloud definitions exist and according to our understanding we want to phrase the following short and concise definition:

[1] Open Cirrus is a trademark of Yahoo! Inc.

H.X. Lin et al. (Eds): Euro-Par 2009 Workshops, LNCS 6043, pp. 434–443, 2010.
© Springer-Verlag Berlin Heidelberg 2010

> Building on compute and storage virtualization, and leveraging the modern Web, cloud computing provides scalable, network-centric, abstracted IT infrastructure, platforms, and applications as on-demand services that are billed by consumption. [4]

Based on this definition we compare grids and clouds. Both technologies focus on IT resources and try to provide an user friendly, inexpensive and pervasive access to these resources over the internet.

The actual situation is that most grid infrastructures consist of geographically distributed, heterogeneous resources without central control, are best suited for special application domains like high energy physics, are publicly funded and use well developed and maintained middleware systems.

In contrast to grids, most cloud infrastructures consist of one or few data centers under central control, are well suited for generic applications, have commercial business models and use proprietary middleware systems. A strong advantage of clouds is the more comfortable usability, as the ownership of resources is granted to the service consumer.

1.1 Everything as a Service

When talking about cloud computing it must be kept in mind that different technical types of cloud services exist.

- **Infrastructure as a Service** (IaaS) implements an abstract view towards the hardware (servers, network,...) and allows to run virtual instances of servers without the need to directly access the bare metal.
- **Platform as a Service** (PaaS) takes the level of abstraction further. PaaS appears as a virtual appliance and makes it simple to scale from a single server to many. Here, the user has no need to worry about the operating system, fundamental software and related application software packages.
- **Software as a Service** (SaaS) provides enterprise quality software (complete applications) to be consumed as a utility.

Cloud computing has the potential to radically change the way IT services are implemented and managed. Project and business funds can be spent to support the core business rather than spending it for IT infrastructure. As they have resource ownership, cloud users are free to run the operating systems, infrastructures, applications and programming languages of their choice. The flexibility of cloud computing has its origin in the combination of virtualization technologies with web services.

2 The OpenCirrus Research Testbed

In July 2008 the OpenCirrus[2] project was announced by HP, Intel and Yahoo!. OpenCirrus aims to build an open, internet-scale global testbed for cloud computing research focusing on data center management, cloud services, systems

[2] `http://opencirrus.org`

and application level research. OpenCirrus is a loose federation of three sponsors which are HP Labs, Intel Research and Yahoo! At the moment there are three also academic partners: The University of Illinois at Urbana-Champaign (UIUC), the Singapore Infocomm Development Authority (IDA) and the Karlsruhe Institute of Technology (KIT).

OpenCirrus consists of six sites initially, hosted by the sponsors and partners, each equipped with 1000–4000 CPU cores and 1 Petabyte of data store.

The basis of the OpenCirrus research testbed is formed by the Physical Resource Sets (PRS) [5]. These provide logical mini-datacenters to the researchers and isolate the experiments from each other. Inside a PRS ensembles of physical nodes are allocated and isolated via virtual local area networks (VLAN) using already existing software like Emulab, a network emulation testbed from the University of Utah and HP Opsware, a tool for server provisioning, configuration and management.

The PRS are the basis to implement Virtual Resource Sets (VRS). The VRS abstract from physical resources by the introduction of a virtualization layer. The virtualization concept applies to all IT aspects like CPU, storage, networks and applications. The main advantage of VRS is the potential to create IT services exactly fitting customers varying needs. IT services can be deployed on demand by automated resource management. Service levels can be easily guaranteed and live migration of services is possible. As a consequence capital expenditures and operational expenditures are both reduced.

The VRS are implementing Infrastructure as a Service (IaaS) because they provide compute, storage and networking services.

OpenCirrus strongly differs from other cloud computing testbeds because it supports both, system and application level research. In contrast to cloud infrastructures like Amazon Elastic Compute Cloud (EC2)[3], Amazon Simple Storage Service (S3)[4] and Google AppEngine[5], all software layers and the hardware itself can be accessed and adapted by the OpenCirrus researchers. Intel platform features like Intel Data Center Management Interface (DCMI) and Node Manager (NM) that support cloud computing could also be utilized.

3 Eucalyptus

A cloud service that will be examined by KIT in the OpenCirrus research testbed is the cloud computing infrastructure service Eucalyptus. Eucalyptus[6] is an Open Source software developed at the University of California, Santa Barbara, implementing cloud computing on university compute clusters. EUCALYPTUS stands for Elastic Utility Computing Architecture for Linking Your Programs To Useful Systems. It implements Infrastructure as a Service (IaaS) while giving the user the ability to run and control virtual machine instances (Xen) deployed

[3] `http://aws.amazon.com/ec2/`
[4] `http://aws.amazon.com/s3/`
[5] `http://code.google.com/appengine/`
[6] `http://open.eucalyptus.com`

across a variety of physical resources [6]. The interface is compatible with EC2 which is the most popular IaaS, and Eucalyptus includes Walrus, a basic implementation of the S3 interface and a block storage service that is interface compatible with Amazon EBS[7].

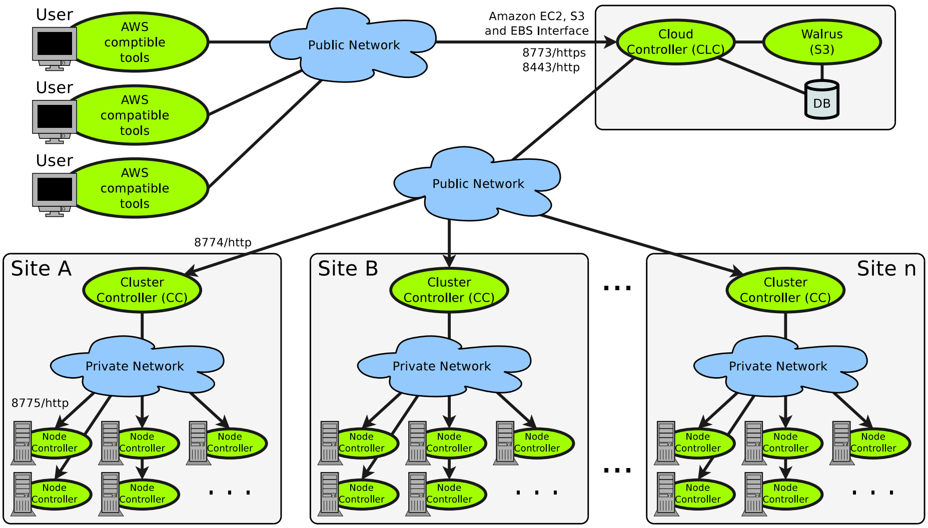

Fig. 1. Structure of Eucalyptus

With Eucalyptus it is possible to build up a private cloud that can be controlled by the same tools known to work with Amazon EC2 and S3. Examples are the ElasticFox EC2 plugin for the Firefox browser or s3cmd utilities for S3 storage management. Eucalyptus has the potential to help establish an open cloud computing infrastructure standard. The main components are the Cloud Controller (CLC), Cluster Controller (CC) and Node Controller (NC) [7]. The NC runs on every node in the cloud as well as a Xen-Hypervisor[8] or KVM[9]. The NC provides information about free resources to the CC. The CC schedules the distribution of virtual machines to the NC and collects (free) resource information. The CLC collects resource information from the CC and operates like a meta-scheduler in the cloud. Figure 1 shows the structure of Eucalyptus including CLC, CC and NC.

For preliminary testing an Eucalyptus R&D cloud installation has been set up at KIT running Eucalyptus 1.4. This installation is used for gaining experience with Eucalyptus and several performance tests. The R&D cloud consists of:

– 2x IBM Blade LS20 (2x Single Core Opteron at 2.4 GHz, 4 GB RAM)
– 2x IBM Blade HS21 (2x Dual Core Xeon at 2.33 GHz, 16 GB RAM)

[7] `http://aws.amazon.com/ebs/`
[8] `http://www.xen.org`
[9] `http://www.linux-kvm.org`

The LS20 Blade is from 2005 and the HS21 Blade from 2006. One LS20 acts as CLC, CC and NC. This consolidation leads not to performance issues because of the small number of NCs in this installation. The other LC20 and the two HS21 are NCs.

While the performance of an Eucalyptus private cloud depends on the hardware available, it is interesting to see the performance key data of this Eucalyptus installation, compared to Amazon EC2/S3. The focus of the performance benchmarks was storage, network and CPU performance.

3.1 Storage Performance of Eucalyptus

In order to compare the storage performance of Amazon S3 and Eucalyptus 1.4 with Walrus, the benchmark tool `Bonnie++`[10] was used. This software measures the rate of sequential output and input. The measurements in Figure 2 are sequential output (per character, per block and rewrite) and sequential input (per character and per block).

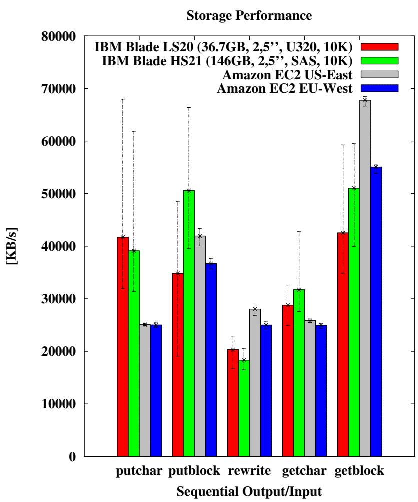

Fig. 2. Storage Performance of Amazon S3 and Eucalyptus

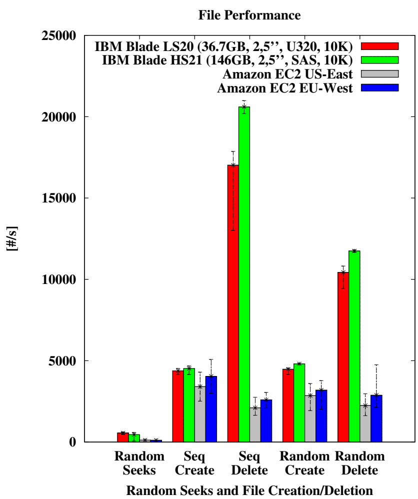

Fig. 3. Performance of Random Seeks and File Creation/Deletion for Amazon S3 and Eucalyptus

The RAM of the Eucalyptus NCs was reduced to overcome memory caching effects. The storage performance of Eucalyptus depends on the features of the storage subsystem. For these tests, the write performance of Eucalyptus, using a modern SAS hard disk with 10000 RPM (revolutions per minute) is faster than Amazon S3. The performance for read in contrast is faster at the Amazon sites.

[10] `http://sourceforge.net/projects/bonnie/`

Bonnie++ also measures the performance for random seeks and especially file creation (Figure 3). Both is faster with Eucalyptus. A possible explanation for these measurements is that Eucalyptus stores the data at the NCs locally. It is unknown whether data stored by Amazon S3 is located near or far the EC2 instances. The reason why file deletion at this Eucalyptus installation performs that much better compared to Amazon S3, remains unclear.

While testing the storage performance of Eucalyptus, the instances were running alone on the blade servers to avoid interferences. We cannot make any assertion about the workload of the Amazon S3 system and the load of the physical server the EC2 instances were hosting at Amazon during our testing.

3.2 Network Transfer Rate

The network transfer rate inside and between the Eucalyptus site and the Amazon EC2 sites was measured at a working day (July 2th 2009) with `iperf`.

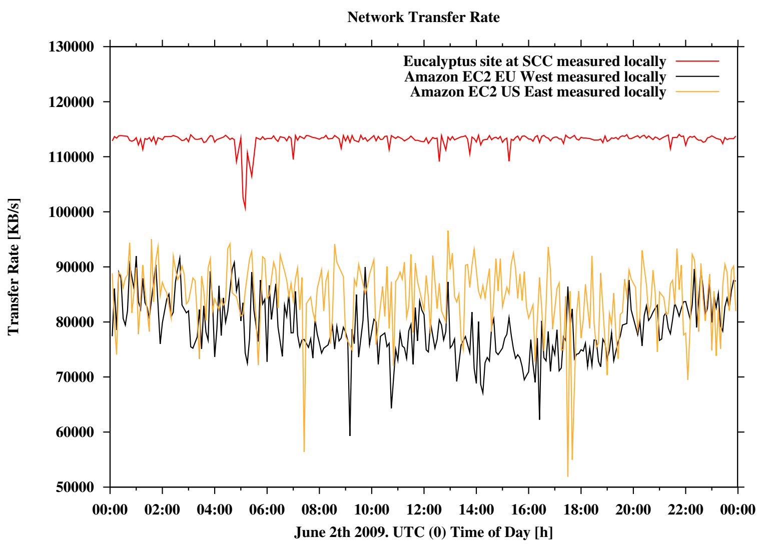

Fig. 4. Network Transfer Rate (1)

The strong in-house network transfer rate (Figure 4) is not surprising because of the 1000 Mbit/s Ethernet. But it is evident that the network transfer rate to the Eucalyptus infrastructure is much more constant in contrast to Amazon EC2. The network transfer rates inside the Amazon EC2 US East and EU West sites imply that there is 1000 Mbit/s Ethernet also used, but with a higher workload.

Figure 5 shows, that the network transfer rate from Karlsruhe to Amazon EC2 EU West is approximately twice as much better compared to Amazon EC2 US East. The network transfer rate to Eucalyptus over the German national research and education network (DFN) is more constant compared to Amazon

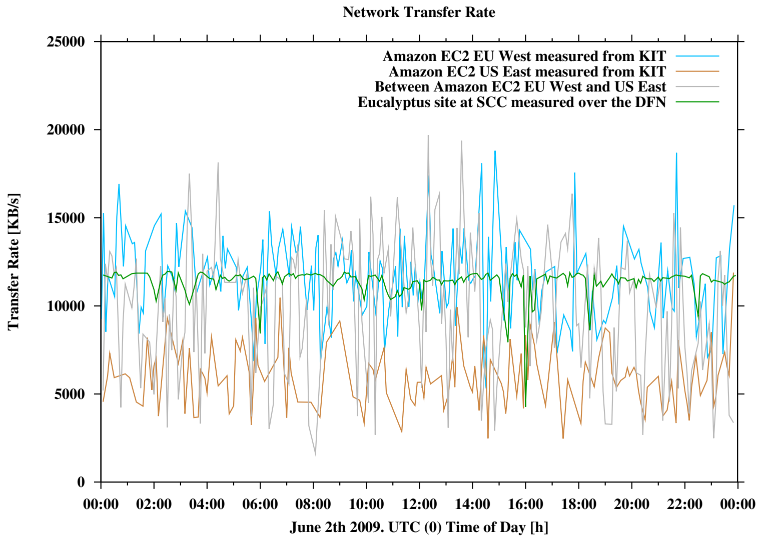

Fig. 5. Network Transfer Rate (2)

EC2. This is also not surprising. An interesting outcome is that peaks of the network transfer rate to Amazon EC2 EU West are much better compared to the connection to the Eucalyptus site in Karlsruhe over the DFN.

3.3 Network Latency

To examine the potential for using Ecalyptus and public cloud infrastructures at all for HPC, the network latency was measured. The results in Table 1 show that HPC in the cloud over institutional/geographical borders is impossible.

For MPI-Jobs where every task computes a few seconds like Monte Carlo methods it is possible to use cloud infrastructures. The network latency inside Amazon EC2 is poor and a surprising outcome is that the network latency between Amazon EC2 EU West and Amazon EC2 US East is better than inside the Amazon sites.

Table 1. Network Latency

`ping ip -f -c 10000`	time [ms]	min. Round-Trip-Time [ms]	avg. Round-Trip-Time [ms]	max. Round-Trip-Time [ms]
EC2 EU West from KIT	138262	27.943	28.192	59.399
EC2 US East from KIT	137014	92.839	93.154	118.853
inside EU West	146447	87.493	90.069	145.109
inside EC2 US East	147380	87.527	92.266	115.461
EC2 EU from EC2 US	138451	88.260	90.776	144.078
Eucalyptus at SCC via DFN	131145	15.093	15.197	29.863
inside Eucalyptus at SCC	2064	0.125	0.146	0.806

3.4 CPU Performance

To compare the CPU performance of Amazon EC2 and Eucalyptus 1.4, the Linux Kernel was compiled for benchmarking. All available Amazon EC2 instance types (see Table 2) were tested.

Table 2. Amazon EC2 Instance Types

`m1.small`	(Small Instance)	1.7 GB RAM	1 virtual Core
`c1.medium`	(High-CPU Medium Instance)	1.7 GB RAM	2 virtual Cores
`m1.large`	(Large Instance)	7.5 GB RAM	2 virtual Cores
`m1.xlarge`	(Extra Large Instance)	15 GB RAM	4 virtual Cores
`c1.xlarge`	(High-CPU Extra Large Instance)	7 GB RAM	8 virtual Cores

For `m1.small` instances, the single virtual core is equivalent to one EC2 Compute Unit. One Amazon EC2 Compute Unit provides the equivalent CPU capacity of a 1.0-1.2 GHz 2007 Opteron or 2007 Xeon processor. This is also the equivalent to an early 2006 1.7 GHz Xeon processor.[11] The virtual cores of the `m1.large` and `m1.xlarge` instances are equivalent to two EC2 Compute Units each. The virtual cores of the `c1.medium` and `c1.xlarge` instances are equivalent to 2.5 EC2 Compute Units each.

Eucalyptus even provides five instance types (see Table 3) following the identical naming scheme than EC2.

Table 3. Eucalyptus Instance Types

`m1.small`	(Small Instance)	128 MB RAM	1 virtual CPU
`c1.medium`	(High-CPU Medium Instance)	256 MB RAM	1 virtual CPU
`m1.large`	(Large Instance)	512 MB RAM	2 virtual CPUs
`m1.xlarge`	(Extra Large Instance)	1 GB RAM	2 virtual CPUs
`c1.xlarge`	(High-CPU Extra Large Instance)	2 GB RAM	4 virtual CPUs

For CPU benchmarking, the time needed to compile Linux Kernel 2.6.29.3 with 1, 2, 4 and 8 threads was measured. The results in Figure 6 show that additional RAM and CPUs are leading to a significant performance boost when using more threads. For none of the instance types more than 8 threads lead to better results.

The measurements in Figure 6 also show our Eucalyptus infrastructure performs approximately twice better for `m1.small` instances compared to Amazon EC2. This is not surprising because the CPU differs and due to the fact that while testing the CPU performance of Eucalyptus the instances were running alone on the blade servers to avoid interferences. We cannot make any assertion about the load of the physical server the EC2 instances were hosting at Amazon

[11] `http://aws.amazon.com/ec2/instance-types/`

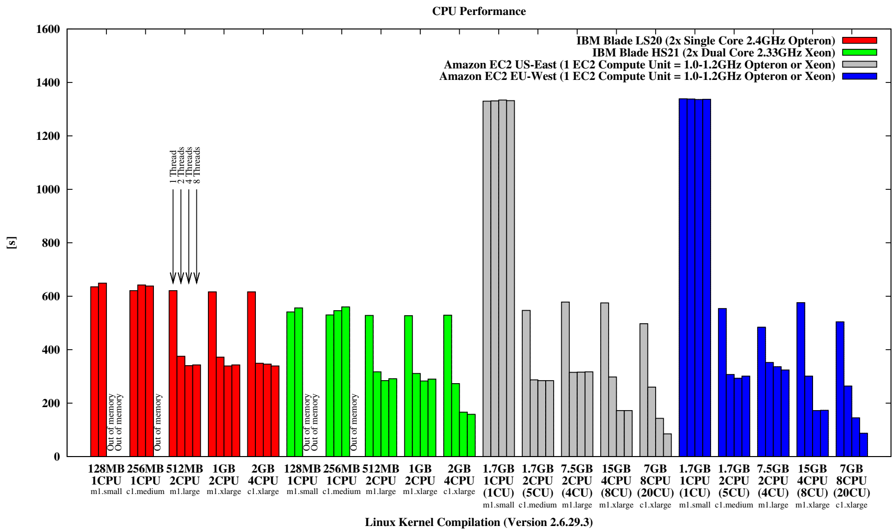

Fig. 6. CPU Performance measurement via Linux Kernel compilation

during our CPU testing. But it is likely that the Amazon servers storing `m1.small` instances have to share their resources between lots of instances thus reducing the CPU performance radically.

Using more threads than virtual/physical CPUs/cores available is not leading to a performance boost because of the thread context switching overhead. The reason why using more than 2 threads at `c1.xlarge` with Eucalyptus at the IBM LS20 is not leading to a significant performance boost is because the LS20 has only two single core CPUs. The IBM HS21 has two dual core CPUs and therefore using 4 threads at `c1.xlarge` leads to an performance enhancement.

The CPU performance measurements strongly depend on the workload of the physical machines. For Eucalyptus, the instances were running alone on the blade servers to avoid interferences but it is impossible to make any assertion about the workload of the physical servers inside the Amazon sites.

4 Further Steps

Currently a second Eucalyptus R&D cloud is set up at KIT with 5x HP Blade ProLiant BL2x220c. Each blade includes two server nodes (2x Intel Quad-Core Xeon at 2.33 GHz, 16 GB RAM). This system is running Eucalyptus 1.5.2 with AppScale 1.1. The purpose of this new installation is to gain experience to operate the OpenCirrus KIT site with virtualization services based on 2656 Nehalem cores in 332 HP T2 servers starting in autumn 2009.

5 Conclusion

Cloud computing allows flexible and elastic resource provisioning. The high degree of automation and the large economies of scale make it attractive for both, academia and business, pathing the way from manufacture towards the industrialization of IT.

OpenCirrus offers interesting R&D opportunities for cloud systems research and application development.

With Eucalyptus, an Open Source implementation of the perhaps most popular IaaS offering of Amazon is available, representing a first step towards the creation of a cloud standard. Although Eucalyptus is a new development the software performs sufficiently stable and when using up to date hardware the users have no need to fear a lower performance as compared to Amazon. The performance that can be achieved with Eucalyptus depends on the physical servers and their workload.

With commodity hardware and Open Source software, a private cloud can be build up providing the same functionality and better performance compared to the most popular public clouds.

References

1. Foster, I., Kesselman, C. (eds.): The Grid: Blueprint for a New Computing Infrastructure. Morgan Kaufmann, San Francisco (1999)
2. Buyya, R., Yeo, C.S., Venugopal, S.: Market-Oriented Cloud Computing: Vision, Hype, and Reality for Delivering IT Services as Computing Utilities. In: HPCC'08: Proceedings of the 2008 10th IEEE International Conference on High Performance Computing and Communications (2008)
3. Bégin, M.E.: An EGEE Comparative Study: Grids and Clouds – Evolution or Revolution, CERN (2008)
4. Baun, C., Kunze, M., Nimis, J., Tai, S.: Cloud Computing: Web-basierte dynamische IT-Services. Springer, Heidelberg (2009)
5. Campbell, R., Gupta, I., Heath, M., Ko, S., Kozuch, M., Kunze, M., Kwan, T., Lai, K., Yan Lee, H., Lyons, M., Milojicic, D., O'Hallaron, D., Soh, Y.C.: Open Cirrus™ Cloud Computing Testbed: Federated Data Centers for Open Source Systems and Services Research. In: HotCloud'09: Workshop in Hot Topics in Cloud Computing (2009)
6. Nurmi, D., Wolski, R., Grzegorczyk, C., Obertelli, G., Soman, S., Youseff, L., Zagorodnov, D.: The Eucalyptus Open-source Cloud-computing System. In: CCA'08: Proceedings of Cloud Computing and Its Applications workshop (2008)
7. Nurmi, D., Wolski, R., Grzegorczyk, C., Obertelli, G., Soman, S., Youseff, L., Zagorodnov, D.: Eucalyptus: A Technical Report on an Elastic Utility Computing Architecture Linking Your Programs to Useful Systems. In: UCSB Computer Science Technical Report Number 2008-10 (2008)