Università degli Studi di Bologna
Facoltà di Ingegneria

Corso di

Reti di Calcolatori L-A

Cloud Computing

Antonio Corradi

Luca Foschini

What is Cloud computing?

"The architecture and terminology of cloud computing is as clearly and precisely ***defined as, well, a cloud****."*

Source:
www.opencloudmanifesto.org

Cloud Computing Problem Space

"It starts with the premise that the **data services and architecture should be on servers**. We call it **cloud computing** – **they should be in a 'cloud' somewhere**. And that if you have the right kind of browser or the right kind of access, it doesn't matter whether **you have a PC or a Mac or a mobile phone or a BlackBerry or what have you – or new devices still to be developed – you can get access to the cloud**…"

– Dr. Eric Schmidt, Google CEO, August 2006

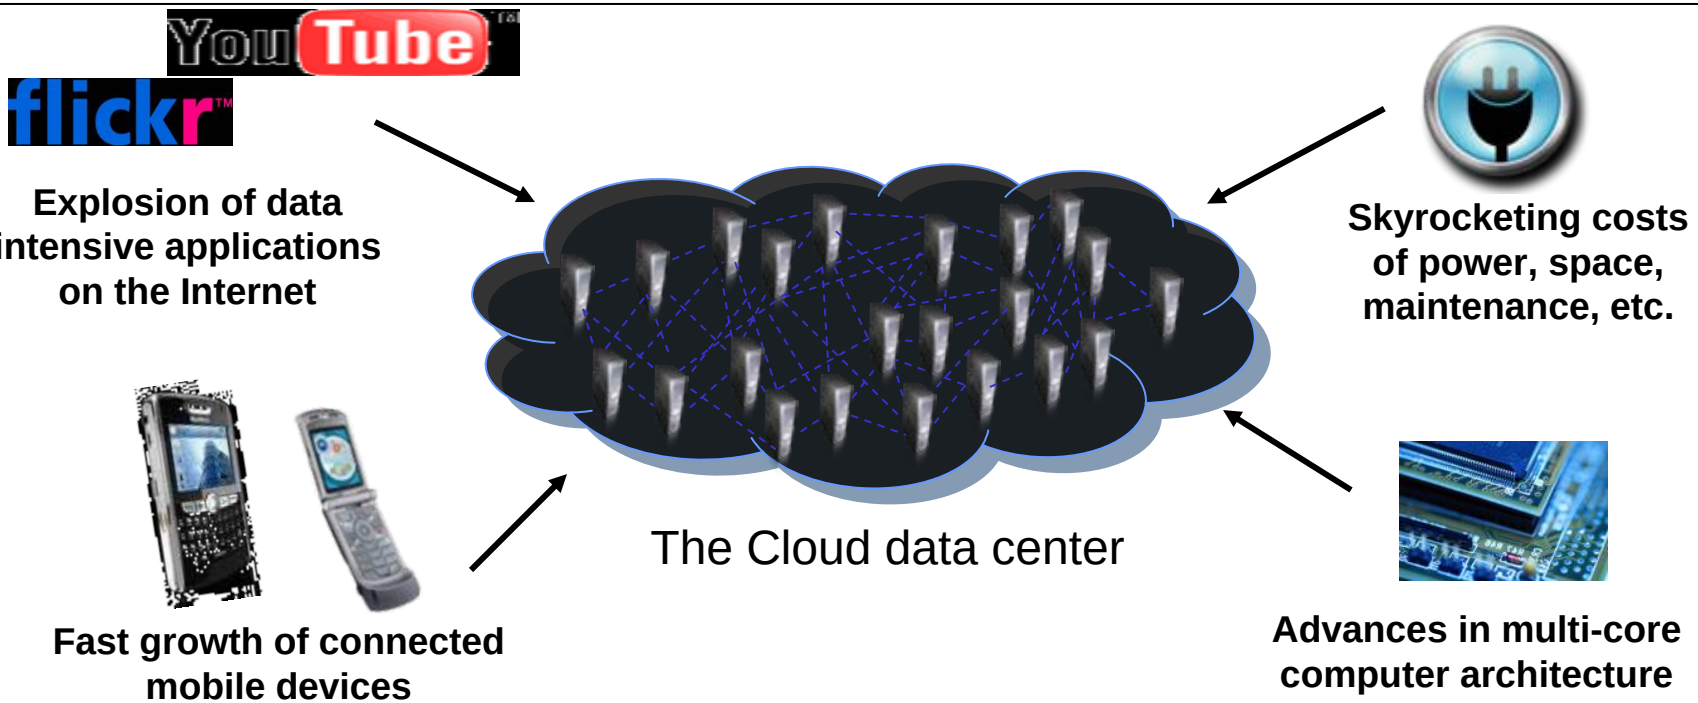

Cloud computing is…

- **Main requirements**
  - **Scalability on demand** (elastic and highly virtualized resources/images, Service Level Agreements – SLA, …)
  - **Automated provisioning and ease-of-use** (utility computing + infrastructure, platform, and software as a service)
  - **Cost efficiency** (minimized startup costs, energy-saving,…)
- **Challenges**
  - **Management** (system resources, **power-saving**, …)
  - **Interoperability and portability** (data, applications, and virtualized images)
  - **Metering and monitoring** (dynamic monitoring of used resources, accounting, …)
  - **Security**

Key Goals - Infrastructure Perspective

- How can we provide **flexible compute resources** quickly to promote **rapid prototyping?**
- How do we **deploy applications** that **scale up** to meet **increasing demands** over time**?**
- How do we **manage 100,000's** of machines with **minimal human intervention?**
- How can we make the **most efficient use** of all **the compute resources** in a data center**?**

What Can You Do With a Cloud?

- **Low-barrier prototyping and development**
  - Setting up and tearing down VMs is extremely fast
  - Promotes innovation
- **Scale-out Web 2.0 applications**
  - Can clone app server VMs in response to user demand
- **Scalable transaction processing**
  - Support for high end UNIX and mainframe servers enables secure, reliable backend database support
- **Data intensive applications**
  - Grid-like computations supported

Cloud computing: reality check

- **Amazon Elastic Computing – EC2**: virtualized images (DB+Software and middleware+OS), Xen, simple SLA console
- **HP/Yahoo/Intel Test Bed**: virtualized images, Xen, simple SLA console
- Research initiatives (***RESERVOIR EU FP7 project***, previous projects on grid computing such as EEGE, ...)
- **Google App Engine** (**S**oftware as **a S**ervice **SAS**, Web applications, Google App Engine, sandbox for management and security)
- **IBM Blue Cloud**: virtualized images (DB+Software and middleware+OS), Xen, Tivoli (monitoring and management), simple SLA console

Others ongoing projects: **Eucalyptus**, **3Tera**, ...

Google App Engine

- **Web Application on Google infrastructures**
- **Application Environment**
  - Sandbox: secure environment that distributes Web requests for the application across multiple servers and starts/stops servers to meet traffic demands
  - Python runtime environment
  - Datastore service
  - Google Accounts Integration
- Preview period, only free accounts are available
  - 500 MB and up to 5 million page views a month
  - Up to 3 applications
  - Scalable quotas
- **What you need is**
  - Google App Engine SDK
  - Google Account
  - Text Editor

Amazon EC2

- Features
  - AMI Amazon Machine Image
    - Use pre-configured, templated images to get up and running immediately.
    - Create image containing applications, libraries, data and associated configuration settings
    - Restriction: Linux-based Images
  - Amazon S3 (Simple Storage Service)
    - Providing safe (?), reliable (?) and fast (?) repository to store the AMIs
  - Amazon EC2 (Elastic Computing Cloud)
    - Web service that lets the user requisition AMIs
- Price
  - Pay only for the resources that are used
    - Different SLAs
      - Small, Large, Extra Instances
      - Data Transfer Levels
    - Different prices
  - SLAs example
    - Small Instance
      - $0,10 per instance-hour
      - 1.7 GB of memory
      - 1 EC2 Compute Unit
      - 160 GB of instance storage
      - 32-bit platform

Amazon EC2

- How does it work?
  - Subscribe account
  - Get Firefox Plug-In
  - Run your image
    - Example
      - Fedora Core 4
      - Apache
      - MySQL
  - Manage it

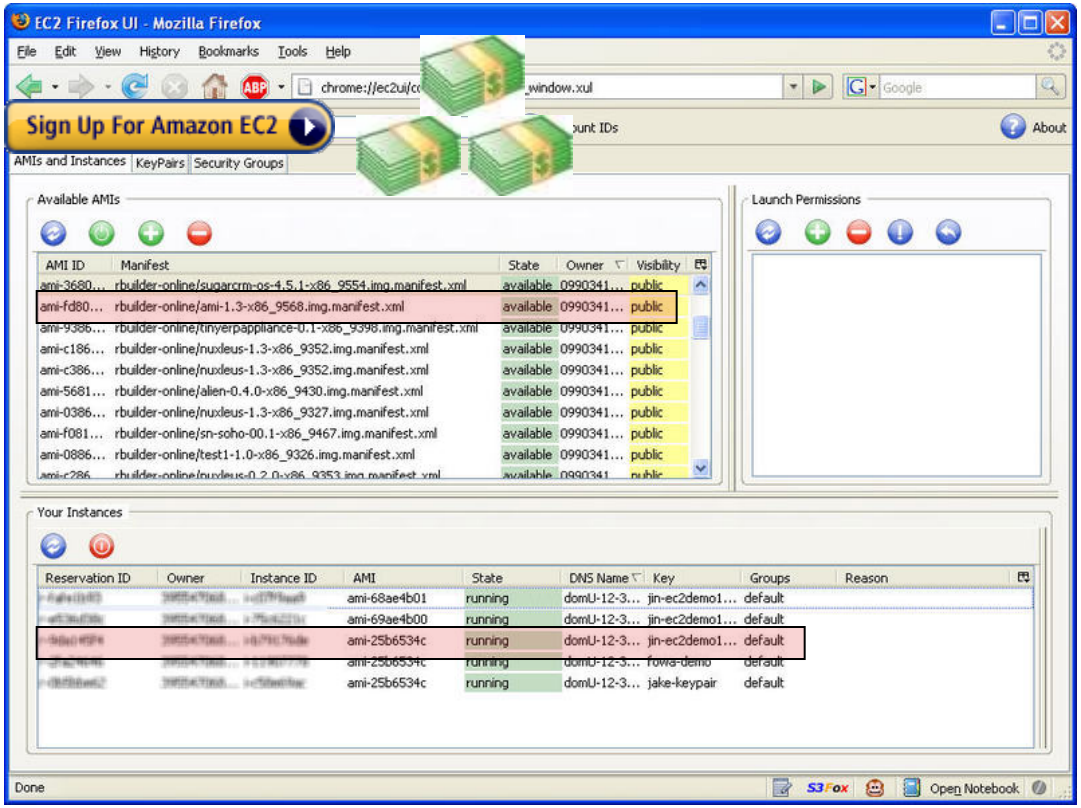

HP/Yahoo/Intel Test Bed

- **Different projects on Cloud Computing**
  - Configure the number of hosts and their memory and disk
  - Specify how much it's willing to pay for the virtualized cluster through a spending rate
  - All of these variables can be changed at any point without interrupting running jobs Increasing the spending rate will immediately increase the CPU share on the cluster nodes
  - Possible Hadoop Integration

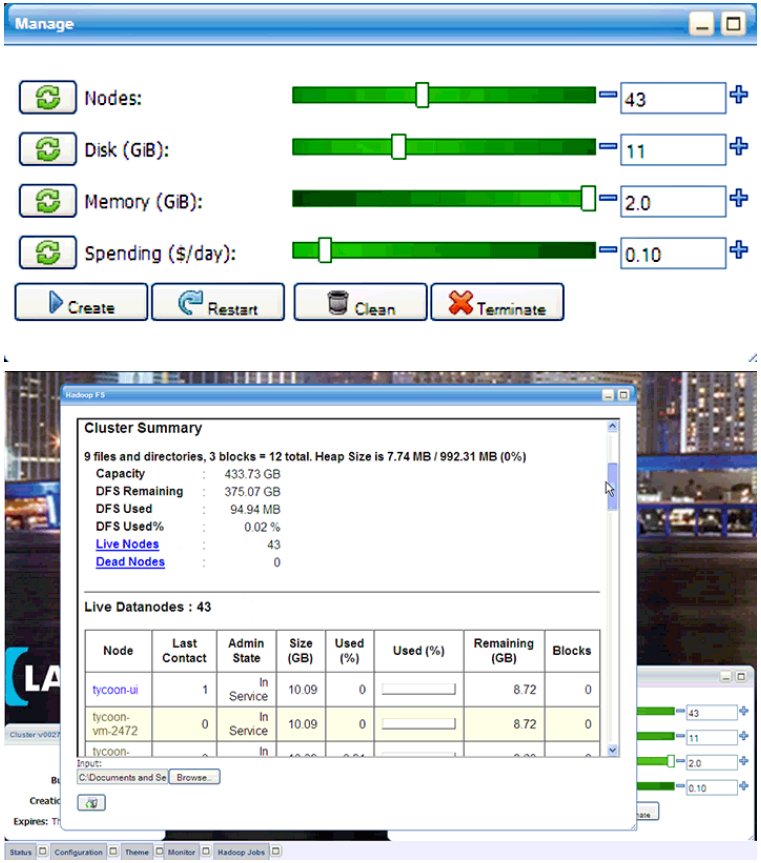

3Tera

Architecture intended to provide an **open framework** allowing the development of a **Cloud computing environment** that's rigorous enough to take Web or Enterprise application

Configuration options

Resource	Min	Max
CPU's	2	1024
RAM, GB	2	2048
Storage, GB*	750	512,000
IP addresses	32	1024

Architecture

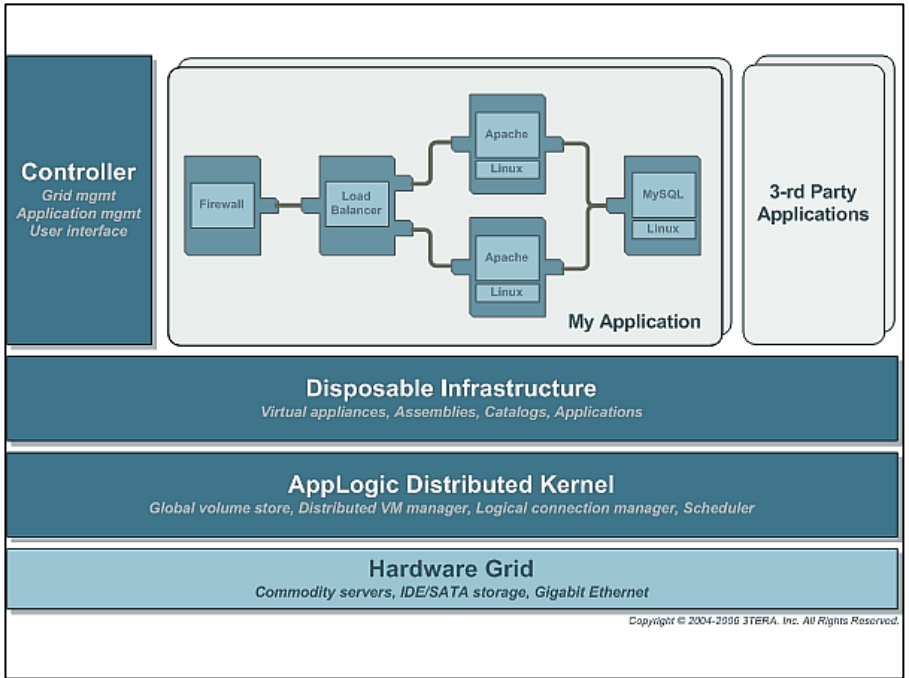

Eucalyptus

- Open-source software infrastructure for implementing Cloud computing on clusters
- Linux systems
- Xen (versions 3.*) for virtualization
- Rocks based (open-source cluster manager)
- Virtual Machines Provisioning

Eucalyptus

- **Eucalyptus Features (1.2 and 1.3)**
  - Installation
    - Rocks-based "almost-one-button" binary install
    - Experts-only "you-are-on-your-own" source install
    - RPM packages for "non-Rocks" RPM based systems
  - Administration
    - Adding/approving/disabling/deleting users (via the Web interface)
    - Adding/listing/disabling images (Web interface with command line use)
    - Adding/deleting nodes and clusters (via edit of configuration files)
  - Amazon's EC2 compatibility:
    - In terms of command-line tools

Xcerion

- Xcerion is an Internet service providing a virtual desktop and OS for free
- Technology
  - XIOS/3 XML Internet Operating System
  - XML Virtual Machine executes the applications locally instead of in the cloud
  - Cloud used for data persistence, storing the users files

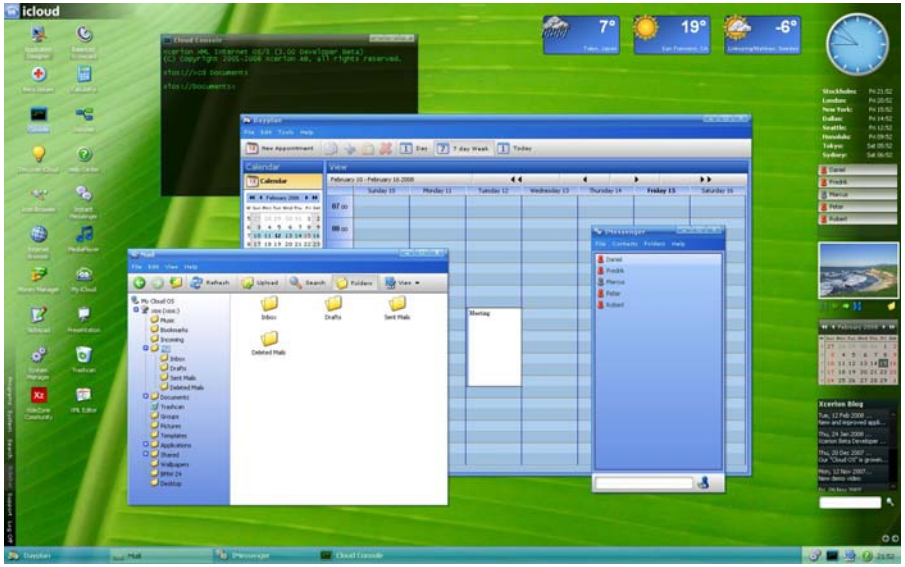

Basic Cloud Computing Architecture

Data Center

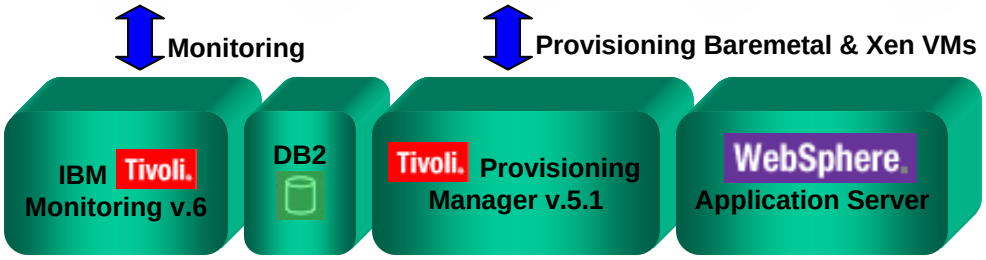

Provisioning Management Stack

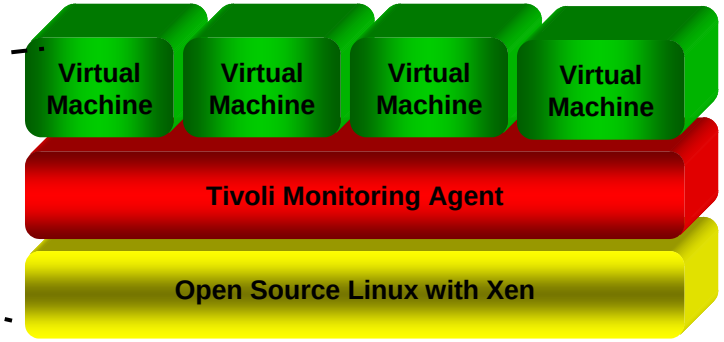

Virtualized Infrastructure Based on Open Source Linux & Xen

Virtualization: all physical machines act as virtual machine hosts; all workloads run on virtual machines

Provisioning: dispense preloaded virtual machines in minutes

Monitoring: ensure systems that go down are recycled quickly

MapReduce Programming Model

- Functional programming that is easily parallelizable
- Split into two phases:

  Map – Perform custom function on all items in an array

  Reduce – Collate map results using custom function
- Scales well – computation separated from processing dataflow
- Illustrative example
  - Map that squares the value of numbers in an array

    {1, 2, 3, 4} -> {1, 4, 9, 16}
  - Reduce that sums the squares : 30

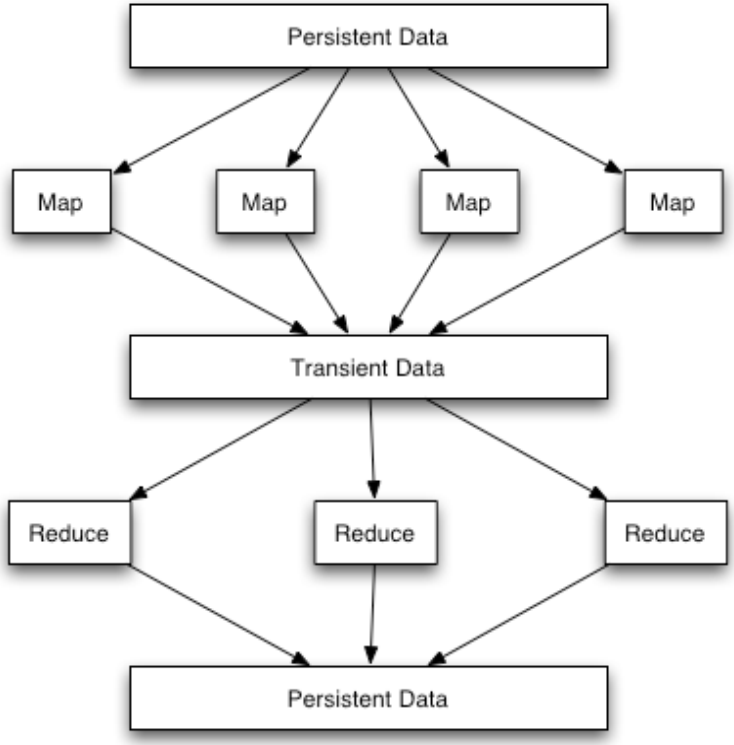

Our experience with the Cloud

We currently work with the Cloud Computing Center at Dublin

- **Virtual machine** (VM) consolidation for power-saving
- Significant for **pervasive computing environment with a high number of VMs**, e.g., one proxy-VM for each node…

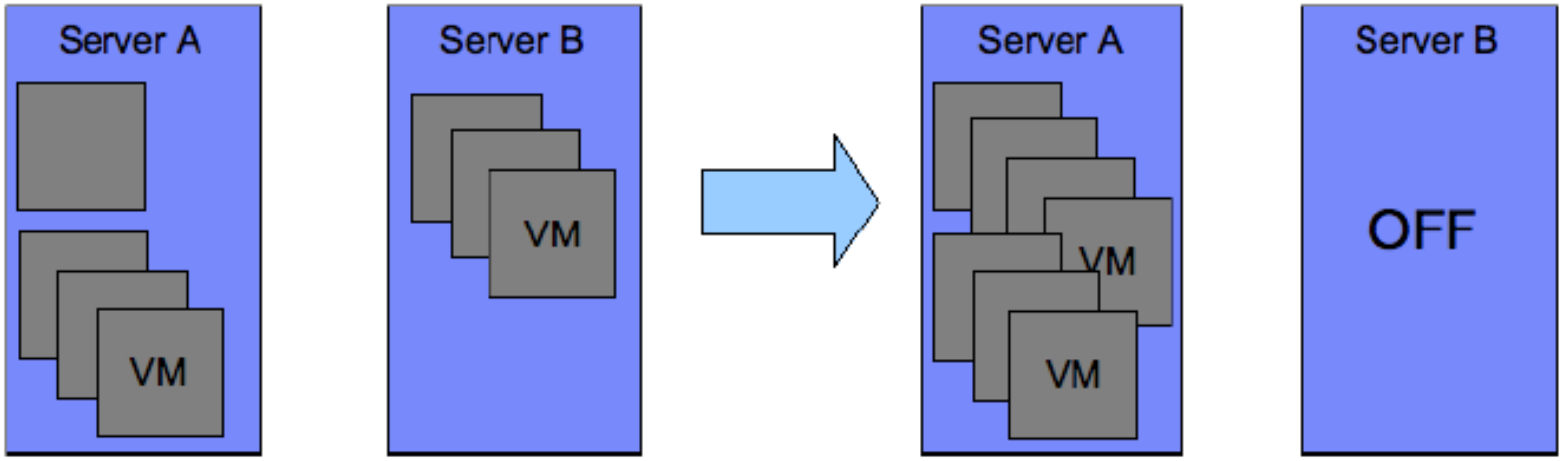

The fog has gone…

… and the Clouds are disclosed into the sky!

Thanks for your attention!

Resources and links

- Introductive YouTube video on Cloud Computing: http://www.youtube.com/watch?v=XdBd14rjcs0
- Introductive YouTube video on Cloud Computing

3Tera

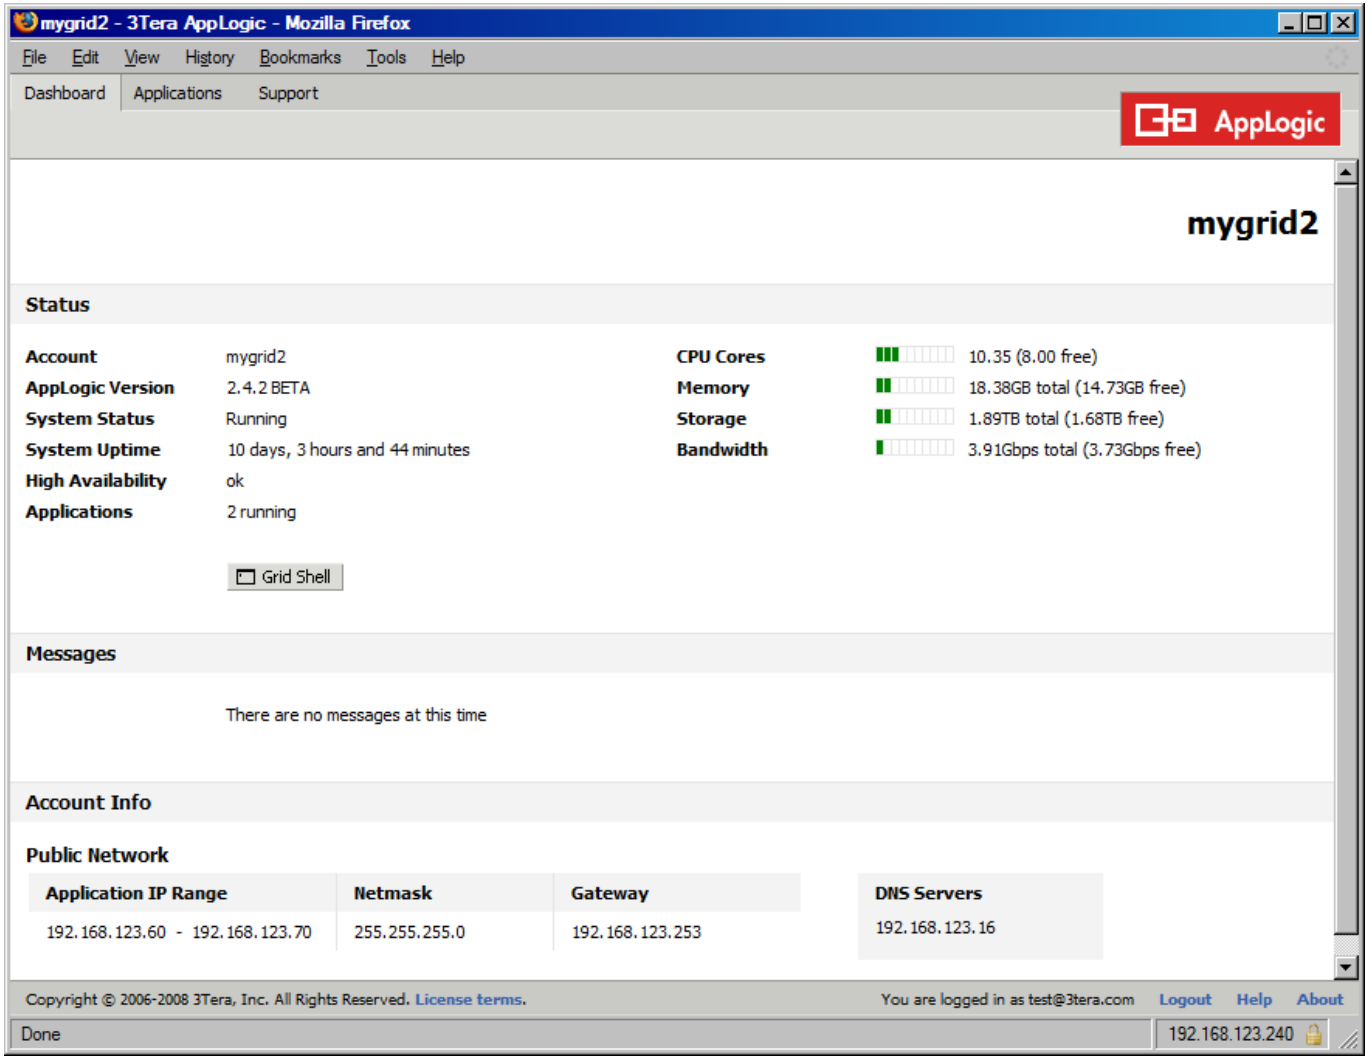
mygrid2 - 3Tera AppLogic - Mozilla Firefox
File Edit View History Bookmarks Tools Help
Dashboard Applications Support
AppLogic
mygrid2
Status
Account mygrid2
AppLogic Version 2.4.2 BETA
System Status Running
System Uptime 10 days, 3 hours and 44 minutes
High Availability ok
Applications 2 running
CPU Cores 10.35 (8.00 free)
Memory 18.38GB total (14.73GB free)
Storage 1.89TB total (1.68TB free)
Bandwidth 3.91Gbps total (3.73Gbps free)
Grid Shell
Messages
There are no messages at this time
Account Info
Public Network
Application IP Range
Netmask
Gateway
DNS Servers
192.168.123.60 - 192.168.123.70
255.255.255.0
192.168.123.253
192.168.123.16
Copyright © 2006-2008 3Tera, Inc. All Rights Reserved. License terms.
You are logged in as test@3tera.com Logout Help About
Done
192.168.123.240

3Tera

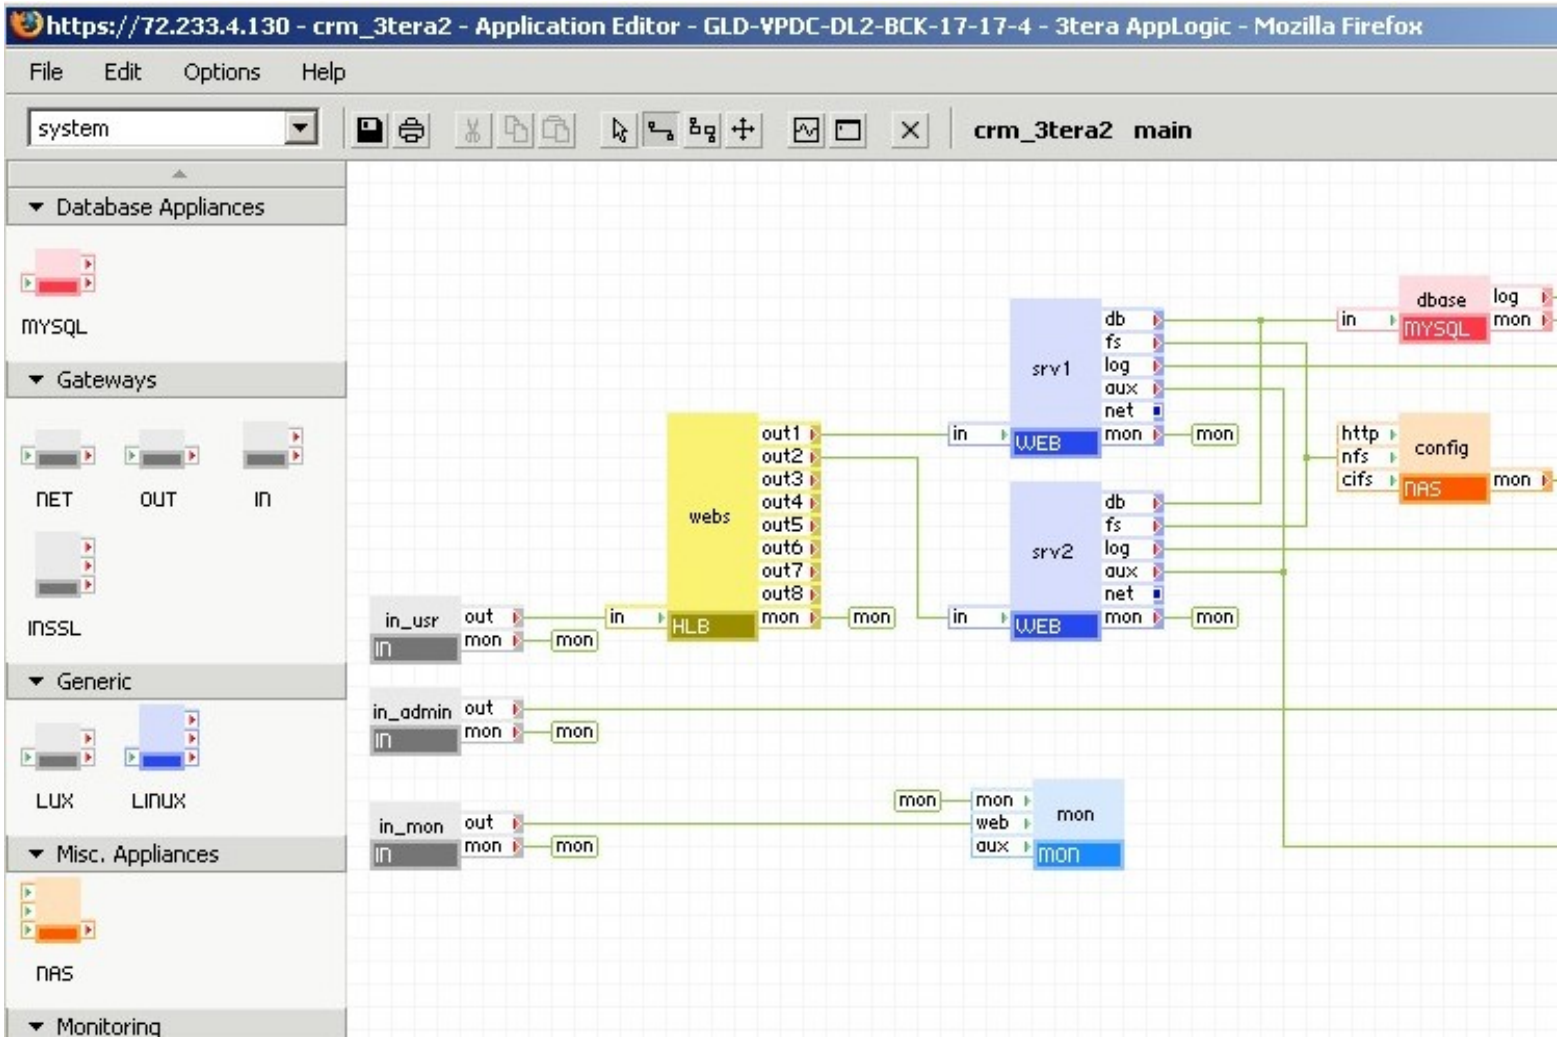
https://72.233.4.130 - crm_3tera2 - Application Editor - GLD-VPDC-DL2-BCK-17-17-4 - 3tera AppLogic - Mozilla Firefox
File
Edit
Options
Help
system
crm_3tera2 main
Database Appliances
MYSQL
Gateways
NET
OUT
IN
INSSL
Generic
LUX
LINUX
Misc. Appliances
NAS
Monitoring
webs
out1
out2
out3
out4
out5
out6
out7
out8
in
HLB
mon
srv1
srv2
db
fs
log
aux
net
WEB
dbase
MYSQL
config
http
nfs
cifs
NAS
in_usr
in_admin
in_mon
out
IN
web

3Tera

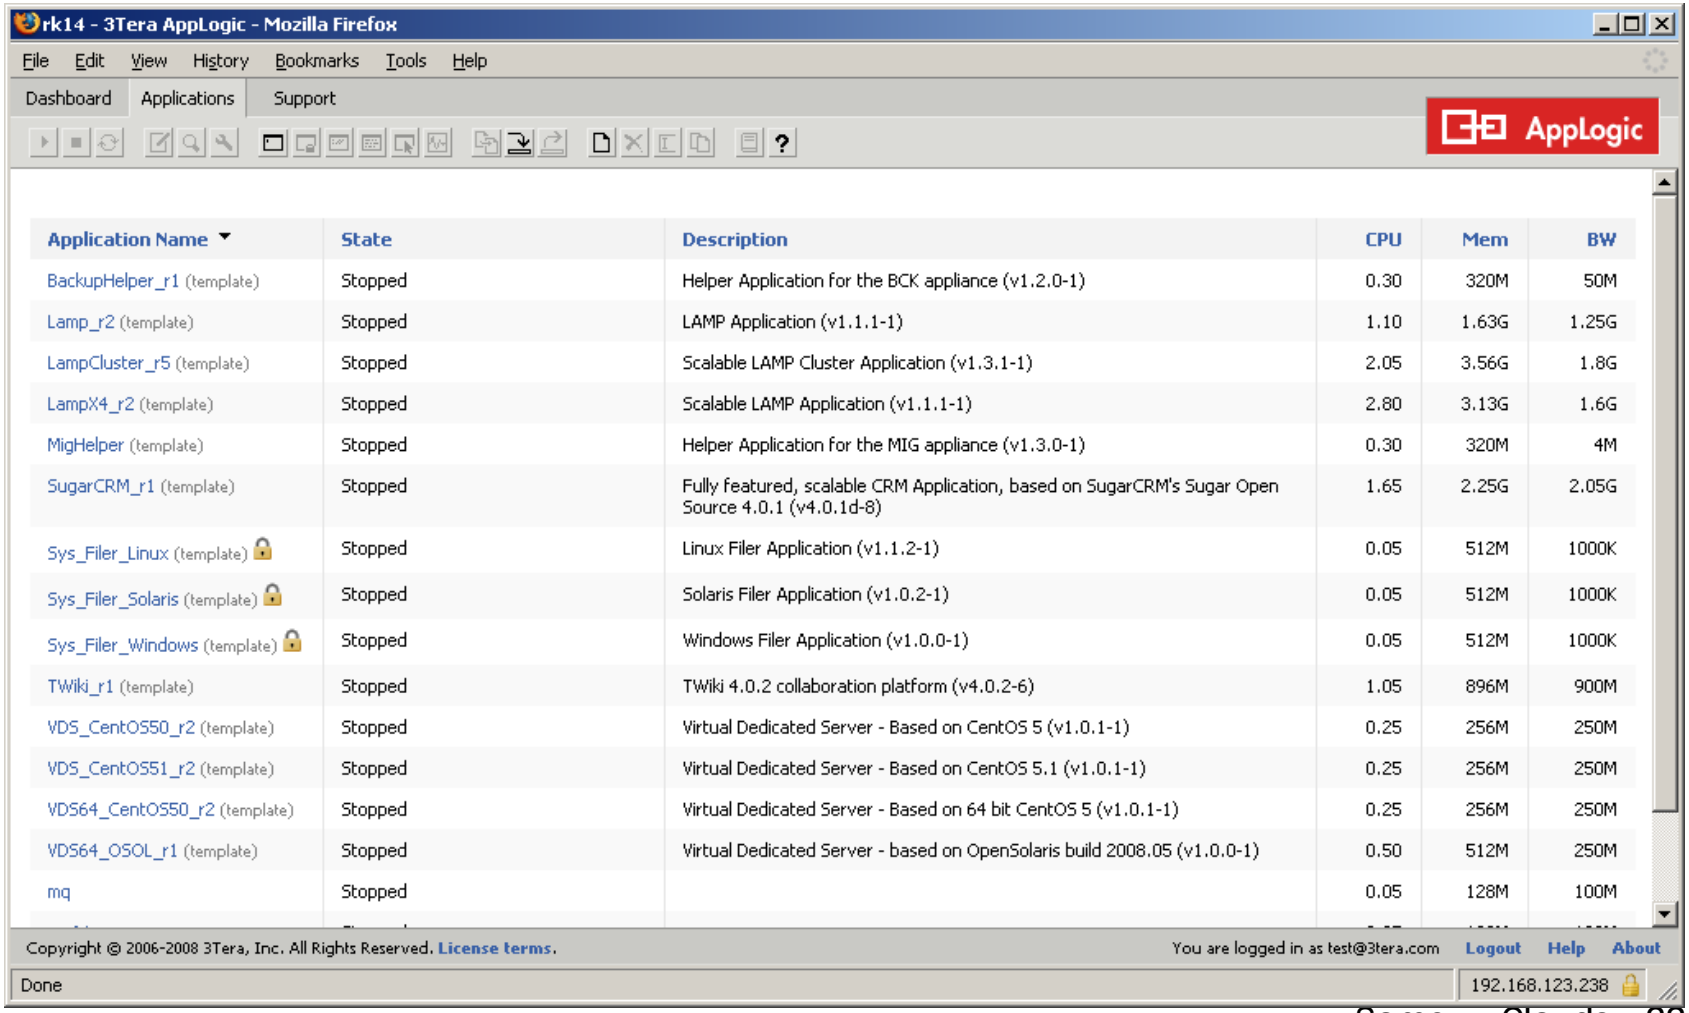
rk14 - 3Tera AppLogic - Mozilla Firefox
File Edit View History Bookmarks Tools Help
Dashboard Applications Support
AppLogic
Application Name | State | Description | CPU | Mem | BW
BackupHelper_r1 (template) | Stopped | Helper Application for the BCK appliance (v1.2.0-1) | 0.30 | 320M | 50M
Lamp_r2 (template) | Stopped | LAMP Application (v1.1.1-1) | 1.10 | 1.63G | 1.25G
LampCluster_r5 (template) | Stopped | Scalable LAMP Cluster Application (v1.3.1-1) | 2.05 | 3.56G | 1.8G
LampX4_r2 (template) | Stopped | Scalable LAMP Application (v1.1.1-1) | 2.80 | 3.13G | 1.6G
MigHelper (template) | Stopped | Helper Application for the MIG appliance (v1.3.0-1) | 0.30 | 320M | 4M
SugarCRM_r1 (template) | Stopped | Fully featured, scalable CRM Application, based on SugarCRM's Sugar Open Source 4.0.1 (v4.0.1d-8) | 1.65 | 2.25G | 2.05G
Sys_Filer_Linux (template) | Stopped | Linux Filer Application (v1.1.2-1) | 0.05 | 512M | 1000K
Sys_Filer_Solaris (template) | Stopped | Solaris Filer Application (v1.0.2-1) | 0.05 | 512M | 1000K
Sys_Filer_Windows (template) | Stopped | Windows Filer Application (v1.0.0-1) | 0.05 | 512M | 1000K
TWiki_r1 (template) | Stopped | TWiki 4.0.2 collaboration platform (v4.0.2-6) | 1.05 | 896M | 900M
VDS_CentOS50_r2 (template) | Stopped | Virtual Dedicated Server - Based on CentOS 5 (v1.0.1-1) | 0.25 | 256M | 250M
VDS_CentOS51_r2 (template) | Stopped | Virtual Dedicated Server - Based on CentOS 5.1 (v1.0.1-1) | 0.25 | 256M | 250M
VDS64_CentOS50_r2 (template) | Stopped | Virtual Dedicated Server - Based on 64 bit CentOS 5 (v1.0.1-1) | 0.25 | 256M | 250M
VDS64_OSOL_r1 (template) | Stopped | Virtual Dedicated Server - based on OpenSolaris build 2008.05 (v1.0.0-1) | 0.50 | 512M | 250M
mq | Stopped | | 0.05 | 128M | 100M
Copyright © 2006-2008 3Tera, Inc. All Rights Reserved. License terms.
You are logged in as test@3tera.com Logout Help About
Done
192.168.123.238