



Implementing Private Cloud Computing Environment with OpenSource Tools

Yim, Sang Hyok
sangyim@hotmail.com



Cloud Computing? (end user)



Cloud Computing? (IT)



How do we convert to CLOUD?



Cloud Computing Attributes

1. Resource Abstraction and Pooling
2. Network-Centric
3. Simple, Fast provisioning of Resources
4. Rapidly and Elastically provisioned Resources
5. Utility Pricing (pay-per-use)

Components

- Computing Unit (CPU/Memory/OS)
- Database Service (RDS,NoSQL)
- Storage Service (Block Storage, File System)
- Middleware Service (WAS, Datapower application)
- Network Service (load-balancing, HA/DR)

Types

- Public
 - Hybrid
 - Private
-
- SaaS – applications, processes (Quicken/Live)
 - PaaS – virtualized middleware (DatSERVICE Farm)
 - IaaS – virtualized servers, storage, networks (EC2,S3)


Hypes

- Not just virtualization or GRID
 - Must be self-serviced/automated facilities to control the resource allocation
- Google/Yahoo are not your typical CLOUDs
- CLOUD is a blackbox to end users
 - Highly available and utility computing environment
- CLOUD is a data center(s) to IT people
 - HA/DR/Scalability
- **Not Revolution – It is Evolution!**

Private Cloud


How to build an enterprise level
cloud infrastructure like Amazon EC2
in your organization

OLD Data Center



Cloud Data Center


Network Farm



Storage Farm



Computing Farm




5 Phases of CLOUDization

1. Virtualization
2. High Availability (HA)
3. Scalability (Elastics)
4. Disaster Recovery (DR)
5. Automation

Phase I - Virtualization

- Virtualized OS - IaaS
 - Virtualized Storage
 - Virtualized Network
-
- Virtualized DB - PaaS
 - Virtualized Middleware - PaaS
 - Virtualized Applications - SaaS

OS Virtualization



OS Virtualization Features


(e.g. Amazon EC2)

- OS Live Migration with Shared Storage
- OS Live Migration without Shared Storage
- OS offline migration with Shared Storage
- OS offline migration without Shared Storage
- OS snapshot copy
- Dynamic Resource Allocation (CPU/Memory/Disk/USB/PCI devices)
- Resource Sharing (CPU/Memory/IO devices)

Hypervisors


1. KVM (RedHat:RHEL/Fedora, CentOS)
2. Xen (Linux – Amazon EC2)
3. VMWARE
4. PowerVM (IBM p4/5/6/7)
5. VBOX (Sun/Oracle)
6. Hyper-V (MS)

Storage Virtualization



Storage Virtualization

(like Amazon S3)



OpenFiler


Open Source Solution for Virtualizing Storage

1. Linux Based Kernel (Under GNU license)
2. Supports SAN& NAS protocols
3. HA cluster failover capable
4. Block-level Remote Replication for DR, support for snapshot
5. Web-based GUI


Consideration for iSCSI instead of SAN

1. Comparable speed with Gigabit speed (1g bps or 10g bps) against Fiber (100m Bps – 1600m Bps)
2. No need for FC hardware
3. Open source tools vs. proprietary tools


Network Virtualization




Migration of Host



Load Balancing




Reverse Proxy




OpenSource Tools for Network Virtualization

- LVS: Piranha
- NAT: iptables (port forwarding)
- Reverse Proxy: squid, apache (mod_proxy)
- Dynamic DNS services:
 - dnshdynamic.org, dyn.com, no-ip.com

DB Virtualization (PaaS)




DB Virtualization (IaaS)




Phase II – High Availability (HA)

- Application (app/middleware/db) level HA – Clustering – heartbeat, redhat cluster (ricci/luci) – hacmp/VCS
- Storage HA – rsync, drbd, RAID, (EMC storage suite – timefinder/srdf bcv)
- Network HA – ethernet bonding, etherchanneling, switch trunk
- Hardware HA – fault-tolerant hardware


DB Active/Passive Cluster




DB Active/Passive Cluster with Standby Server




DB Active/Active Cluster with Shared Storage (e.g. Oracle RAC)




DB Active/Active Cluster without Shared Storage (e.g. MySQL Enterprise Cluster)




Filesystem Replication (rsync)




Shared Filesystem (GFS)



Replicated Filesystem (XtreemFS)



Remote Block Copy (DRBD)




Phase III – Scalability


- Vertical Scaling
 - Dynamic resource allocation
 - Static resource assignment
- Horizontal Scaling
 - Shared Data
 - Load Balancer (Piranha)
- Computing on-demand
 - Grid Computing (RedHat MRG, Condor)

Network Virtualization


RedHat Piranha LVS



From RedHat/CentOS Virtual Server Guide v5.4



From RedHat/CentOS Virtual Server Guide v5.4




From RedHat/CentOS Virtual Server Guide v5.4

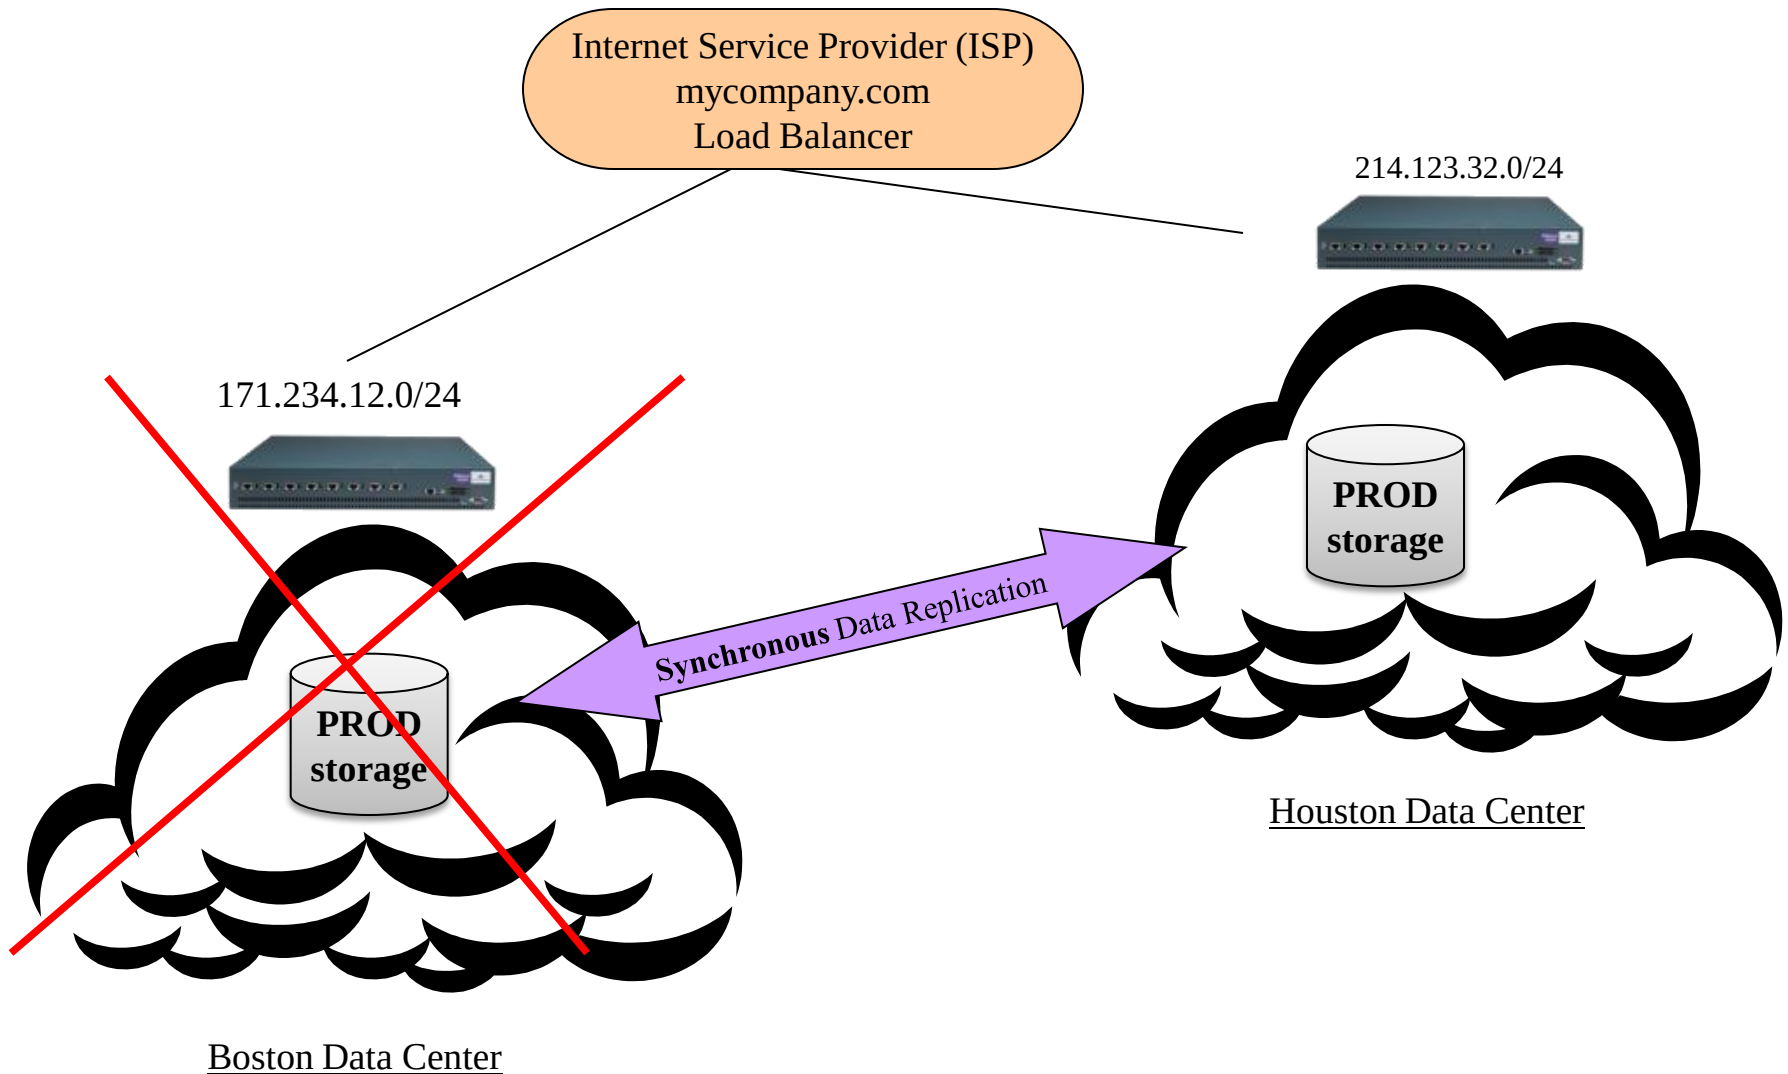
Phase IV – Disaster Recovery

- Remote data center data replication (Standby Data Center)
 - Data replication over slower WAN
 - Data synchronization latency issues
 - Traffic redirection by Communication Provider
 - Disaster Detection
- Real time data replication over multiple data centers
 - Fast WAN communication
 - Multiplexing of many slower communication lines


DR Site (aka: Standby Site)



Multiple Active Data Center



Multiple Active Data Center with Data Bank



Phase V – Automation

- OpenSource Cloud Deployment tools
 - OpenStack
 - OpenNebula


⇒ Provides Cloud Management Tools

⇒ Easy migration to other Cloud Service Providers


⇒ OpenStack by Rackspace

Cloud Service


(e.g. Amazon EC2, Rackspace)




Hybrid Cloud




Demo - Infrastructure




Demo - OpenFiler




Demo - KVM




Demo – MySQL Cluster



Demo - GFS



Demo – Piranha (LVS)



Demo – Putting All Together (okmindmap.com)

