

# Building Storage Service in a Private Cloud

Sateesh Potturu & Deepak Vasudevan

Wipro Technologies

- ❑ Storage in a private cloud is the storage that sits within a particular enterprise security domain and the data is considered more secure. This presentation discusses how we can build a storage solution for a private cloud using off the shelf servers and open source components.
  
- ❑ The presentation explains the suitability of open source components in building private cloud storage with respect to
  - ❑ Types of Storage [ NFS/CIFS/Block storage ]
  - ❑ Backup solutions
  - ❑ Scalability and
  - ❑ Multi-Site Recovery

# Learning Objectives

1. Exploring suitability of different open source components in building storage for a private cloud.
2. Backup Strategies for storage in a private cloud
3. Walkthrough of an implementation of storage service in private cloud.

- ❑ **Cloud Overview**
- ❑ **Public vs Private cloud**
- ❑ **Storage in private cloud**
  - ❑ Background
  - ❑ Requirements
  - ❑ Building an inexpensive storage pool
  - ❑ Block and File storage
    - ❑ iSCSI, NFS & CIFS
  - ❑ Backup in private cloud
- ❑ **Private Cloud – An implementation**
  - ❑ Background
  - ❑ Tiered storage model
  - ❑ Block and File store
  - ❑ Backup
  - ❑ Storage provisioning
  - ❑ Security
  - ❑ Usage billing
- ❑ **Conclusion**
- ❑ **Future**
- ❑ **QA**

# Cloud Overview

# Cloud concepts – Quick glance

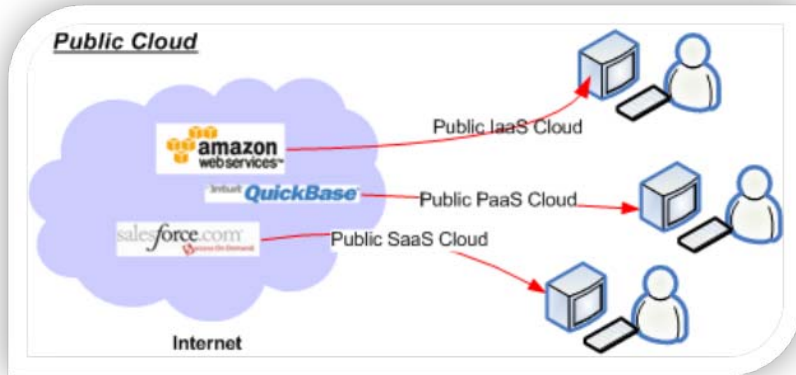
“A systems integrated software solution that enables the orchestration, delivery and management of an on-demand, dynamically scalable, highly available virtualized IT environment that is secure, self-healing, agile and pay-as-you-use”

## Clouds are factories for IT services on an industrial scale



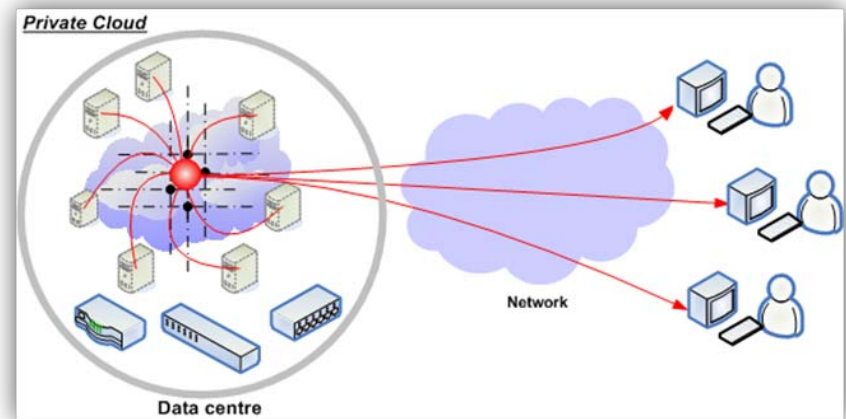
Sharing is achieved through resource virtualization in IaaS, while in PaaS/SaaS the applications and databases support multi-tenancy

# Cloud Models

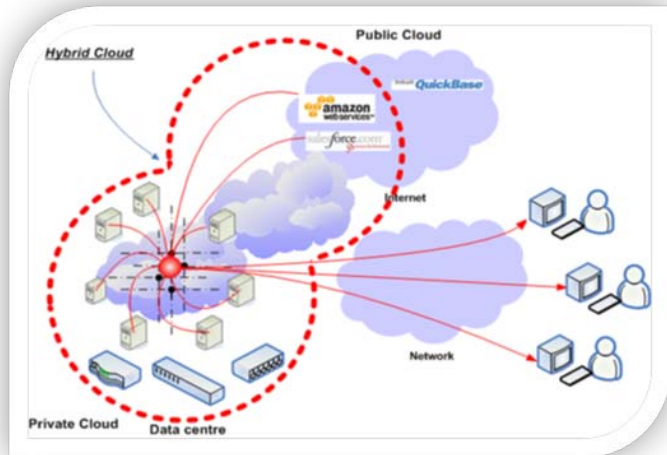


Public Cloud or External Cloud\* - *Resources are dynamically provisioned on a fine-grained, self-service basis over the Internet*

Private Cloud, Internal Cloud, Enterprise Cloud – *Emulation of the Public Cloud on a private network*

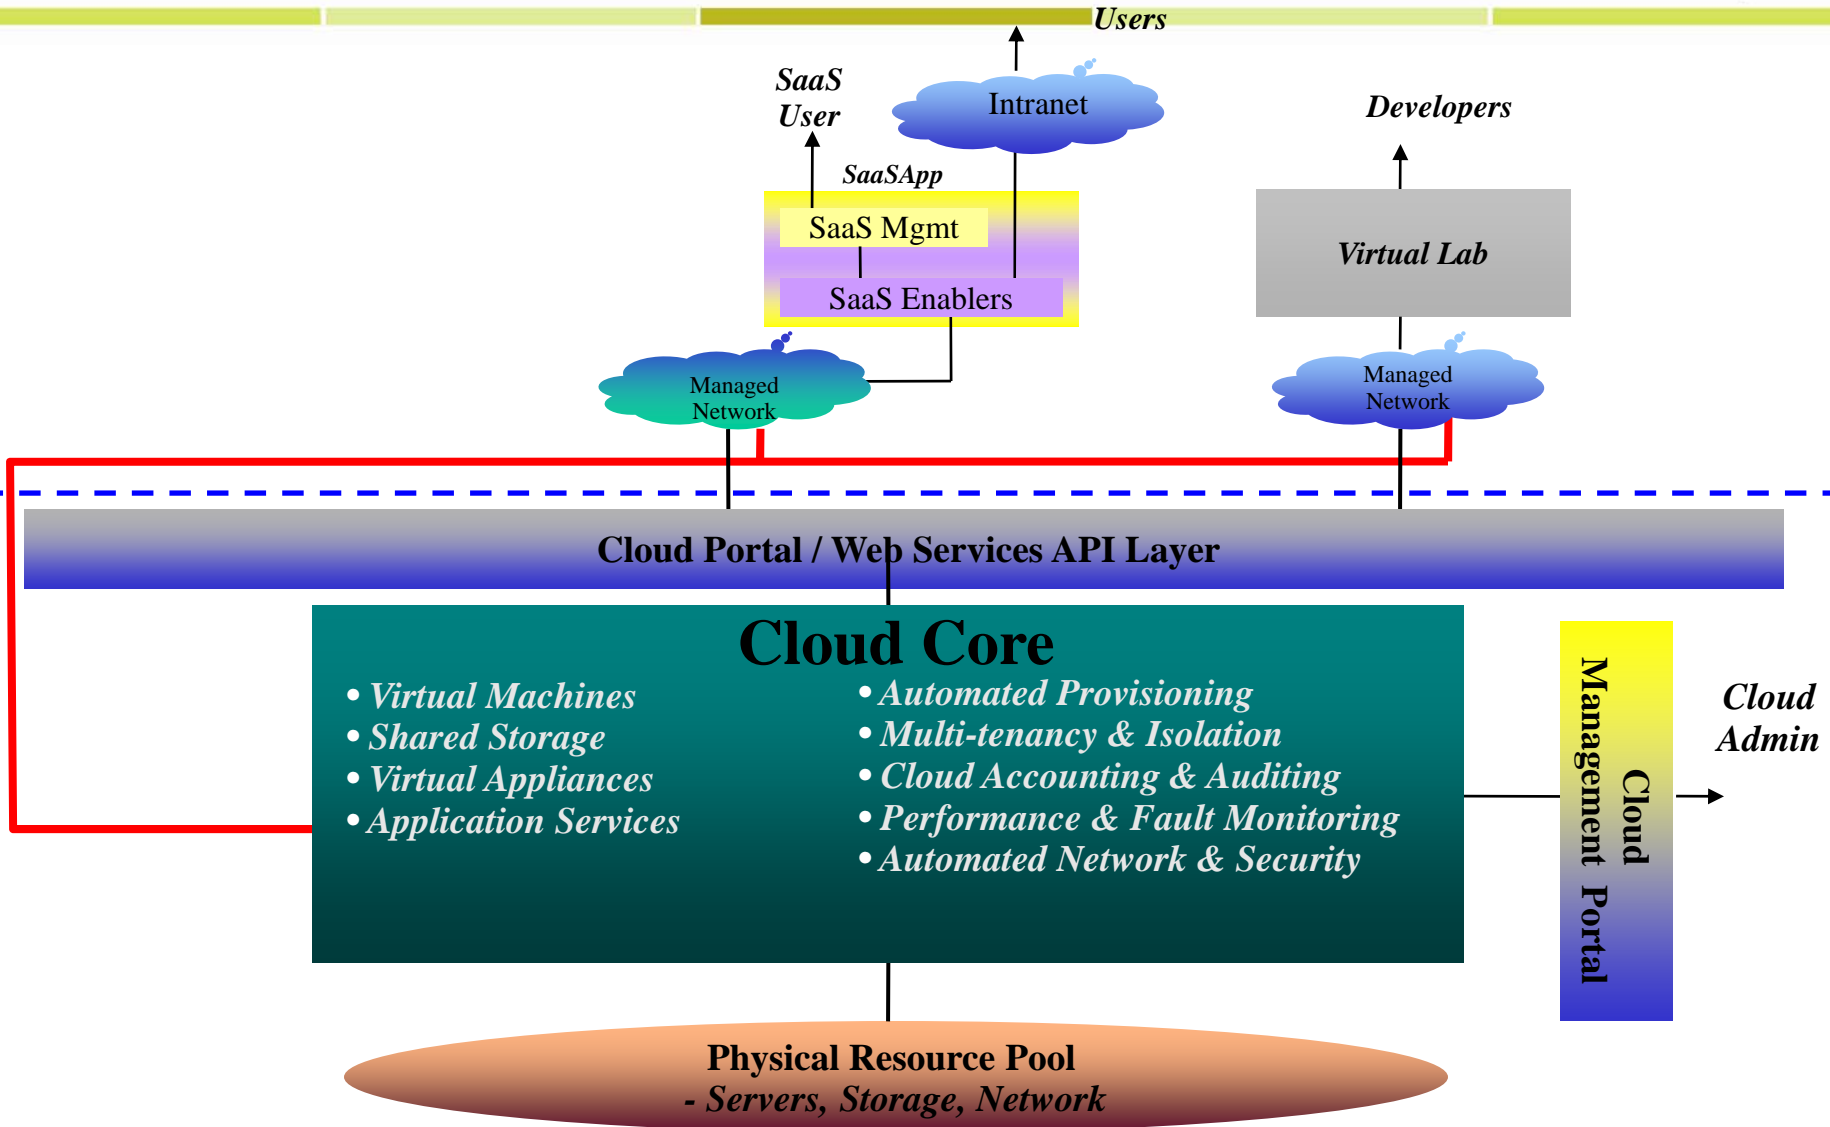


Hybrid Cloud – *A combination of Public and Private cloud*



Source: Wikipedia

# Private Cloud Overview





# Private Cloud Storage

- ❑ Storage pool within an enterprise.
- ❑ Consumed by IT, business domains, R&D, physical labs.
- ❑ Typical usage – hosting databases for organizational functions, Backup, test & deploy solutions.
- ❑ Usage scale lower than public cloud.
- ❑ Yet challenges of automated allocation, dynamic provisioning, scalability, security, fault tolerance, multi tenancy, disaster recovery remain.
- ❑ Cost is a concern – ROI realized via savings accrued.

# Client Storage requirements

- ❑ File system Storage
  - ❑ Provisioning/de-provisioning NFS shares.
  - ❑ Provisioning/de-provisioning CIFS shares.
  
- ❑ Block Storage
  - ❑ Provisioning/de-provisioning of iSCSI luns.

# Building Inexpensive Storage pools

## Drivers:

- ❑ Private Cloud storage places implicit demand for dynamic scalability
- ❑ Charge back – marginal or absent. ROI is via accrued savings.
- ❑ Commercial solutions are expensive.

## Solution:

- ❑ Employ storage aggregation techniques , use cheap disks.
- ❑ TB Sata drives, JBODs can be aggregated to provide block level aggregation – using LVM.
- ❑ File systems can be aggregated using open source Lustre FS
- ❑ Gluster FS – clustered file system, provides FS aggregation.

## Benefits:

- ❑ Using consolidate block and file storage pools, CIFS, NFS, iSCSI solutions can be deployed over it to provide mass storage solution.
- ❑ Other applications – space for Backup, scratch space for app developers.

# Block & File Storage (low cost)

## □ NFS

- Linux server running NFS as a Filer.
- Inexpensive disks/JBOD attached as backend store.
- NFS shares carved out of aggregated storage.
- Filer hooked onto the cloud management network.
- HA configuration of service via Linux Heartbeat.
- Data redundancy using replication (DRBD).
- Dynamic scaling of storage using LVM.

## □ CIFS

- Linux server running Samba.
- Cheap disks for backend store.
- HA, data protection and dynamic provisioning provided by open source components.

## □ iSCSI

- Open source iSCSI target servers – iSCSI Enterprise Target (IET), UNH-iSCSI
- iSCSI luns carved out of cheap JBOD storage.
- Service failover using Linux Heartbeat.
- Data replication using DRBD.

# Backup – an integral part of Storage

## Drivers:

- ❑ Data protection and Disaster recovery.
- ❑ Backup should be quick and reliable.

## Solution:

- ❑ Variety of open source backup solutions available – Amanda, Bacula, Time Vault etc.
- ❑ Use inexpensive storage from Storage pool.

## Benefits:

- ❑ Cheaper solution.
- ❑ Low maintenance, quick deployment.

# Private Cloud Storage – An Implementation



Let us walk through an implementation of Private cloud storage, hereafter referred as PCS.

## Background:

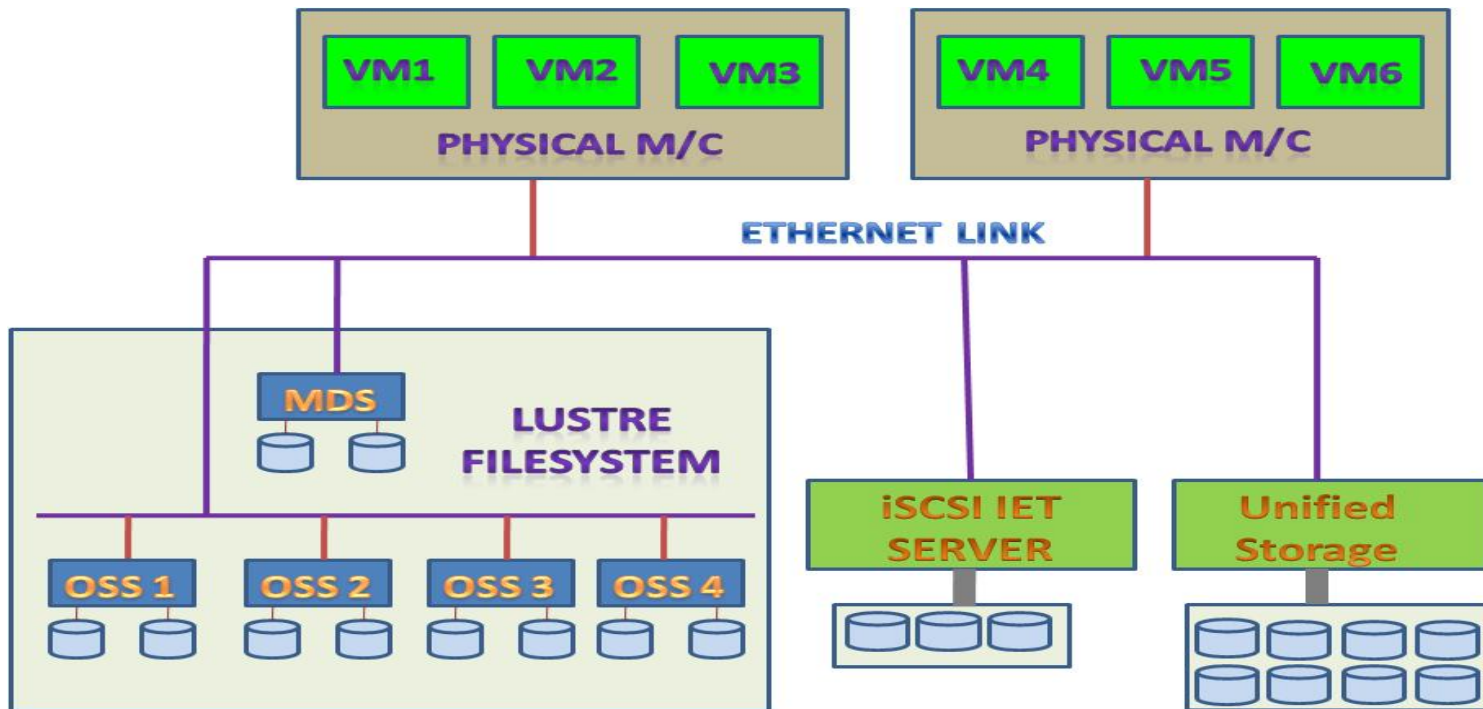
- ❑ Primary consumers of storage in PCS – enterprise applications, virtual desktop users/Servers, COEs that build and tests POC apps.
- ❑ It is envisioned that users would have their additional storage needs sourced from the cloud – Ex: backup
- ❑ Storage Solution built to address above class of usage.
- ❑ Applications access storage using protocols like iSCSI, NFS and CIFS.

## Strategy:

- ❑ Provide a mix of low cost solutions coupled with open source solutions using commodity hardware.

# Tiered Storage Allocation Model

- ❑ Tiered storage allocation model to meet the storage needs of the private cloud.
- ❑ Storage allocated from a pool of inexpensive storage, Unified storage & Aggregated storage.

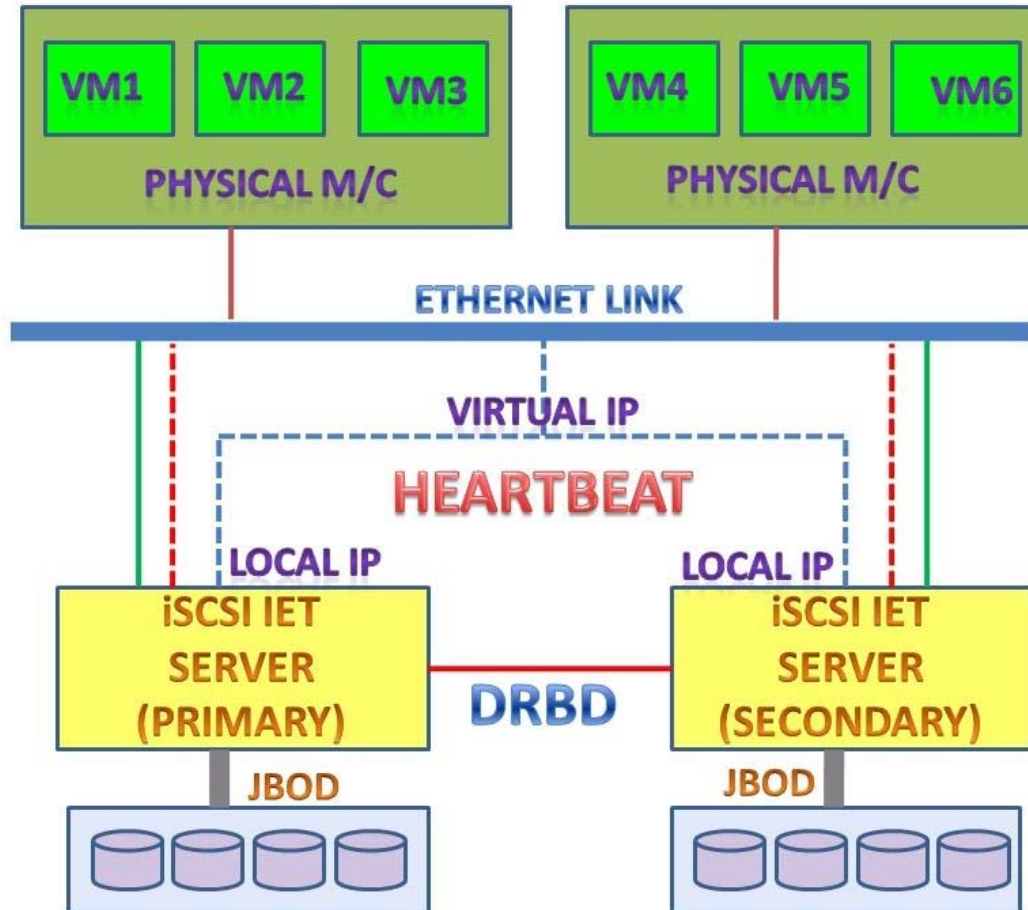


## ❑ Inexpensive storage

### ❑ iSCSI

- ❑ Storage solution built using IET server with DAS.
- ❑ Users provision iSCSI luns over a self-service portal
- ❑ Storage accessed via virtual machines associated with iSCSI luns.
- ❑ Multi-site solution, mirror server in another building.
- ❑ Data replicated over n/w via DRBD.
- ❑ HA using Linux Heartbeat.
- ❑ Lun isolation among VMs provided using ACLs.

# Low Cost Block storage ( IET)



## □ NFS

- NFS server using std distribution on Linux.
- Cheap DAS storage.
- Accessed via VMs.
- HA, DRDB configuration.
- VM isolation using fine grained control.

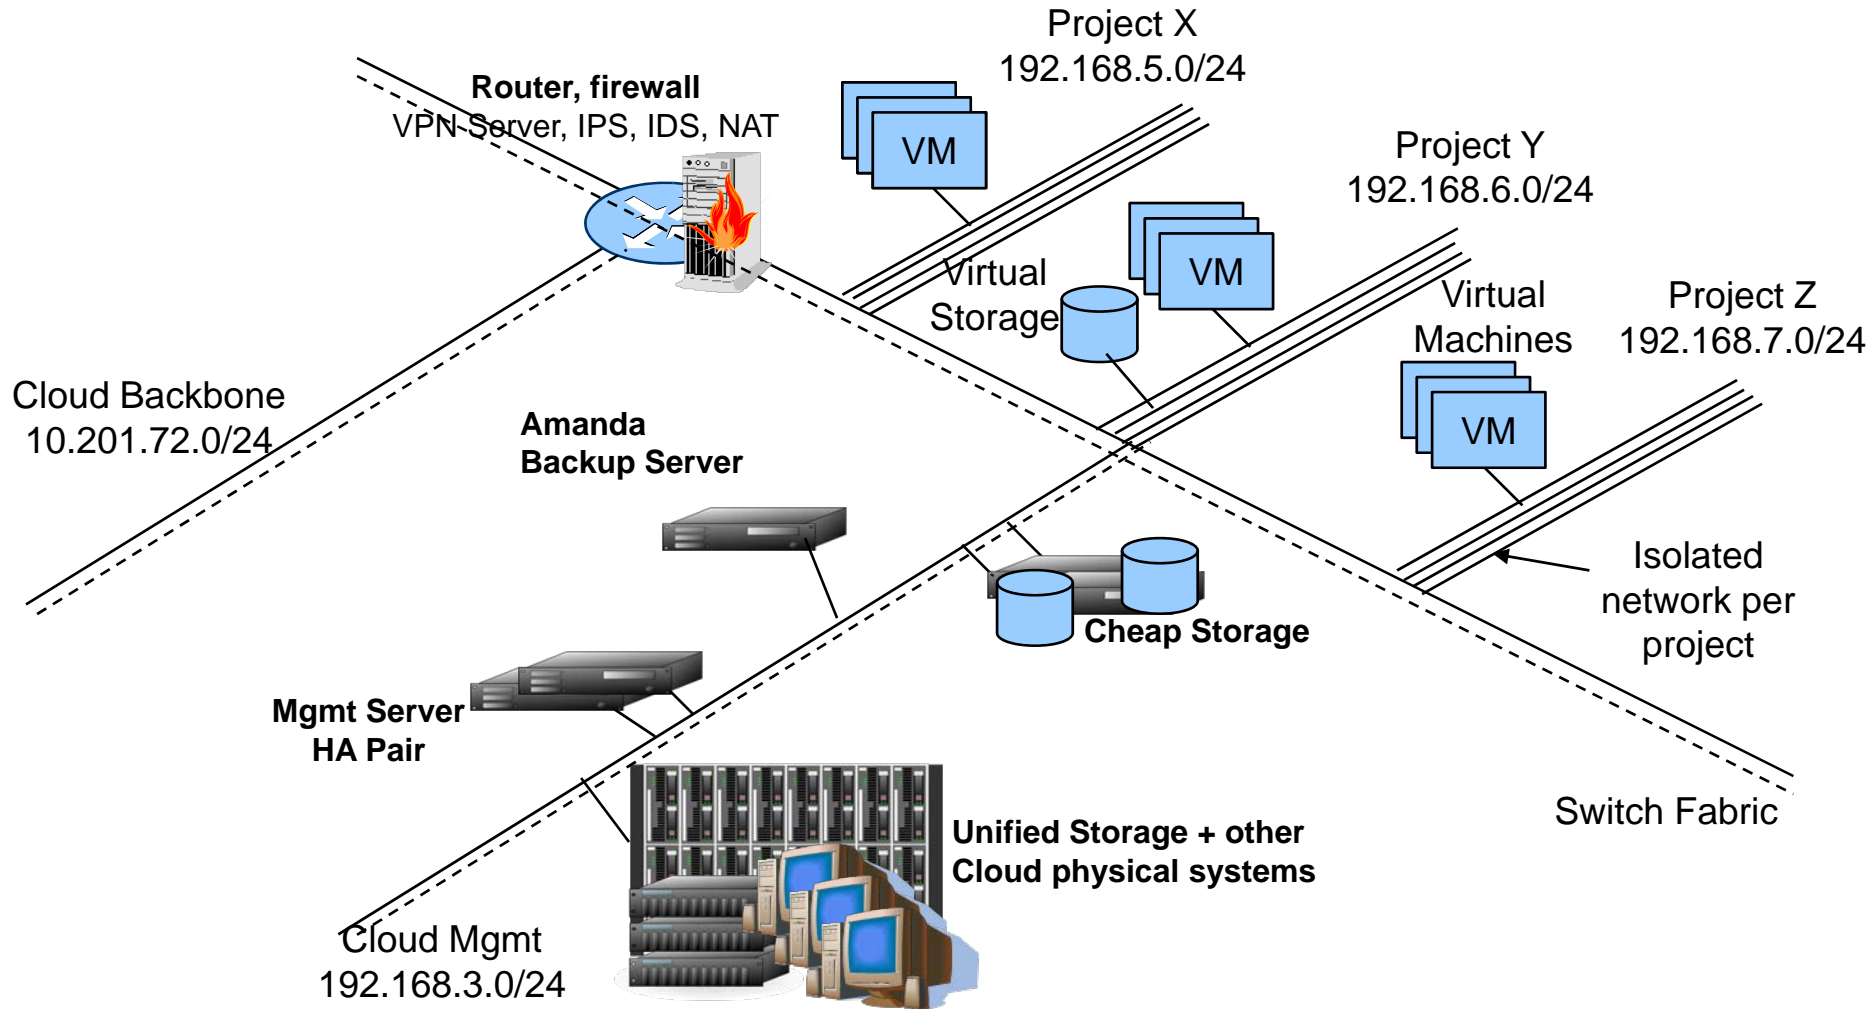
## □ CIFS

- Implemented using Linux servers running Samba with DAS storage.
- Access via VMs.
- HA, DRBD configuration.
- Security based on restricting share access to specific users.
- CIFS storage integrated with enterprise Active directory (AD) for authentication.

# Tiered Storage contd..

- ❑ Unified Storage
  - ❑ Unified Storage Server deployed for alternate storage.
  - ❑ iSCSI, CIFS and NFS shares.
  - ❑ 17TB, thin provisioning, data protection, HA.
  - ❑ GBps link + link aggregation (4Gbps).
- ❑ Aggregated File storage using Lustre
  - ❑ Storage aggregated from cheap desktops (TB Sata drives).
  - ❑ Massively scalable up to peta bytes of storage with hundreds of nodes.
  - ❑ Redundancy was built at meta data server level.
  - ❑ Solution under deployment – used as storage pool for file & block storage + primary target for cloud backup.
- ❑ Advantages of Tiered model
  - ❑ Provides grades of storage – capacity, speed, protection.
  - ❑ SLA based storage allocation in the future.
  - ❑ Heterogeneous storage support.

# Deployment Example



# Cloud Backup



- ❑ Data Distribution
  - ❑ Internal Clients
    - ❑ General purpose image repositories
    - ❑ Cloud management software, tools etc.
  - ❑ External Clients
    - ❑ Images generated by the customers.
    - ❑ Customer data – iSCSI luns , CIFS & NFS shares.
- ❑ Backup strategy
  - ❑ Data saved by internal and external clients need to be backed up.
  - ❑ Solution should scale to backup data from multiple storage pools.
  - ❑ Needs to be fast and reliable.

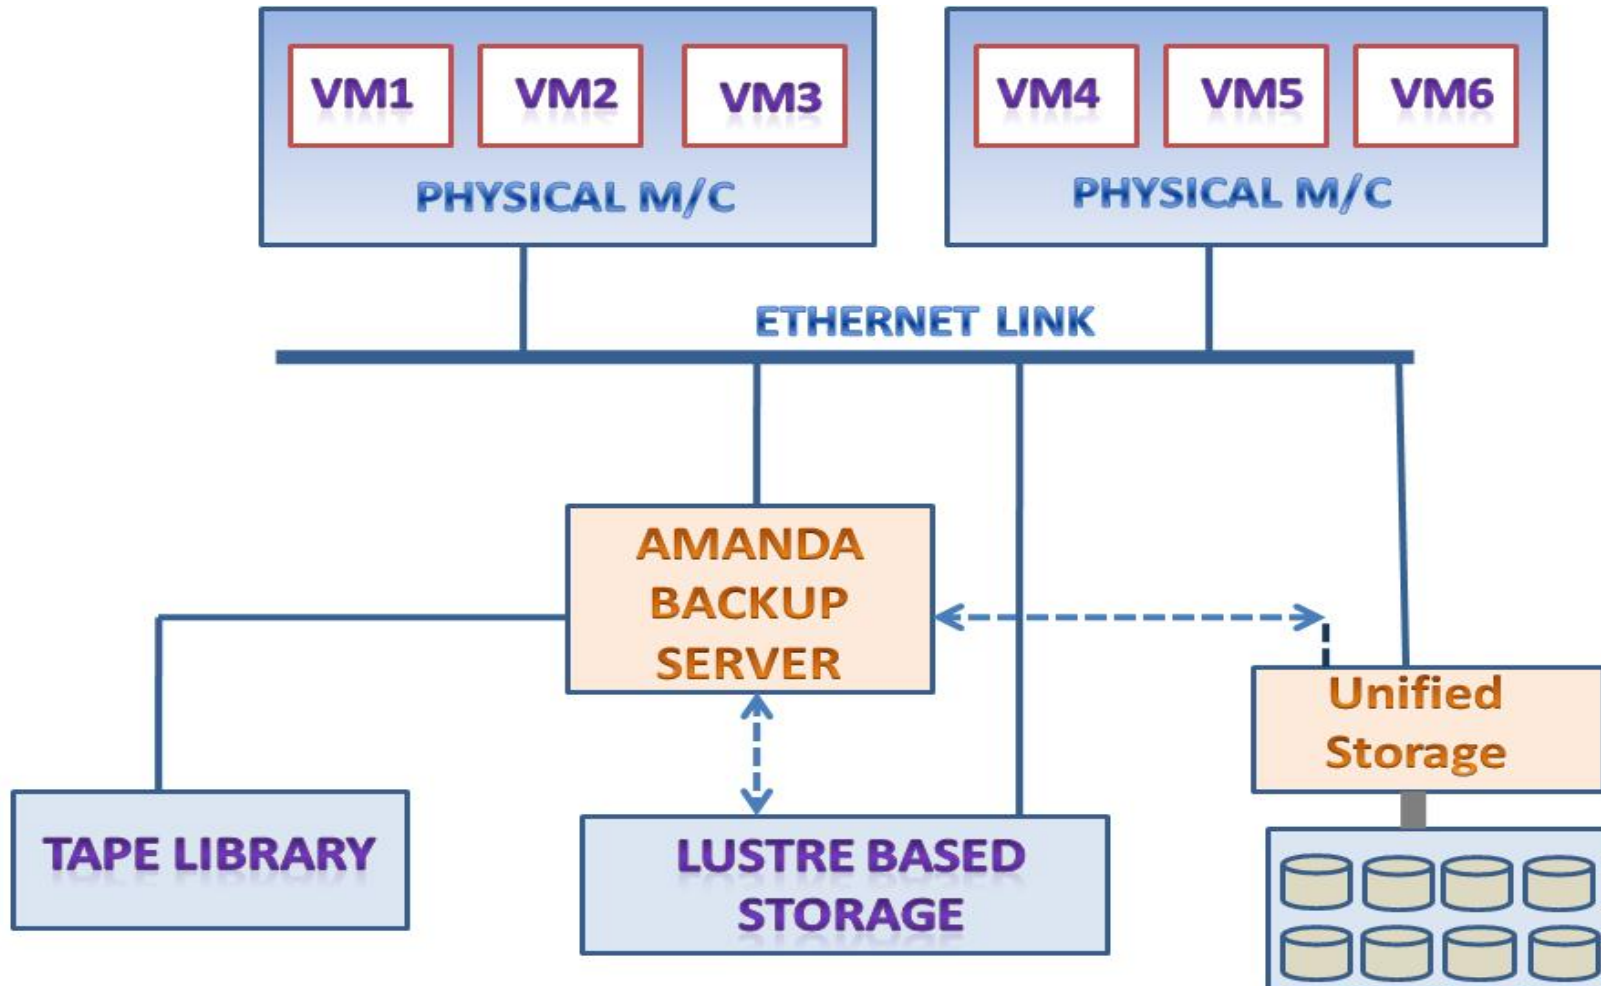
# Backup tenets contd..

- ❑ Options explored:
  - ❑ Use Fast Tape Library.
  - ❑ Use Raid Array
  - ❑ Use existing Dell Power Vault Auto Loader
  - ❑ Alternate low cost option
    - ❑ Used aggregated low cost storage using Lustre FS.
    - ❑ Scales well, dynamic provisioning.
    - ❑ Cheap and efficient

# Backup solution

- ❑ Implemented using open source Amanda backup solution.
  - ❑ Uses standard utils – dump, GNU tar
  - ❑ Well tested and widely used.
- ❑ Lustre FS provides for backend storage.
- ❑ Backup/restore implemented for backing up shares and data on Unified storage.
- ❑ Backup/restores are fast – disk to disk copy.

# Backup layout



- ❑ Service portal
  - ❑ Provisioning via subscription & Self-service
  - ❑ Access over intranet by providing AD credentials.
- ❑ Allocation
  - ❑ Can create Virtual machines & NFS/CIFS/iSCSI storage.
  - ❑ Associate storage with different VMs.
- ❑ Access
  - ❑ The actual access to resource is via public IP.
  - ❑ Ex: A Linux user may 'ssh' into the VM and launch applications that could use the attached storage.

- ❑ Allocated storage accessed through VMs.
- ❑ Each customer on a separate VLANs.
- ❑ Storage access controlled through ACLs.
- ❑ Intranet access, within enterprise security domain.

# Usage billing

- ❑ PCS has tools and mechanisms in place to collect storage usage data for each client.
- ❑ A policy and processing module needs to be developed to use the statistics and compute billing for individual clients.
- ❑ Service is under development..

PCS provides :

- ❑ affordable low cost, high capacity storage.
- ❑ Scalable architecture and supports heterogeneous storage.
- ❑ An integrated backup solution helping in protecting data and enabling disaster recovery management
- ❑ Solution achieved using
  - ❑ commodity hardware
  - ❑ readily available open source components
  - ❑ low cost and effort in designing, implementing, deploying and maintaining the service.



- ❑ HSM – Integrating backup with Tape Library
- ❑ Heterogeneous storage support
- ❑ Customer initiated backup via portal.
- ❑ Storage as a service

**QA**  
**thanks.!**