





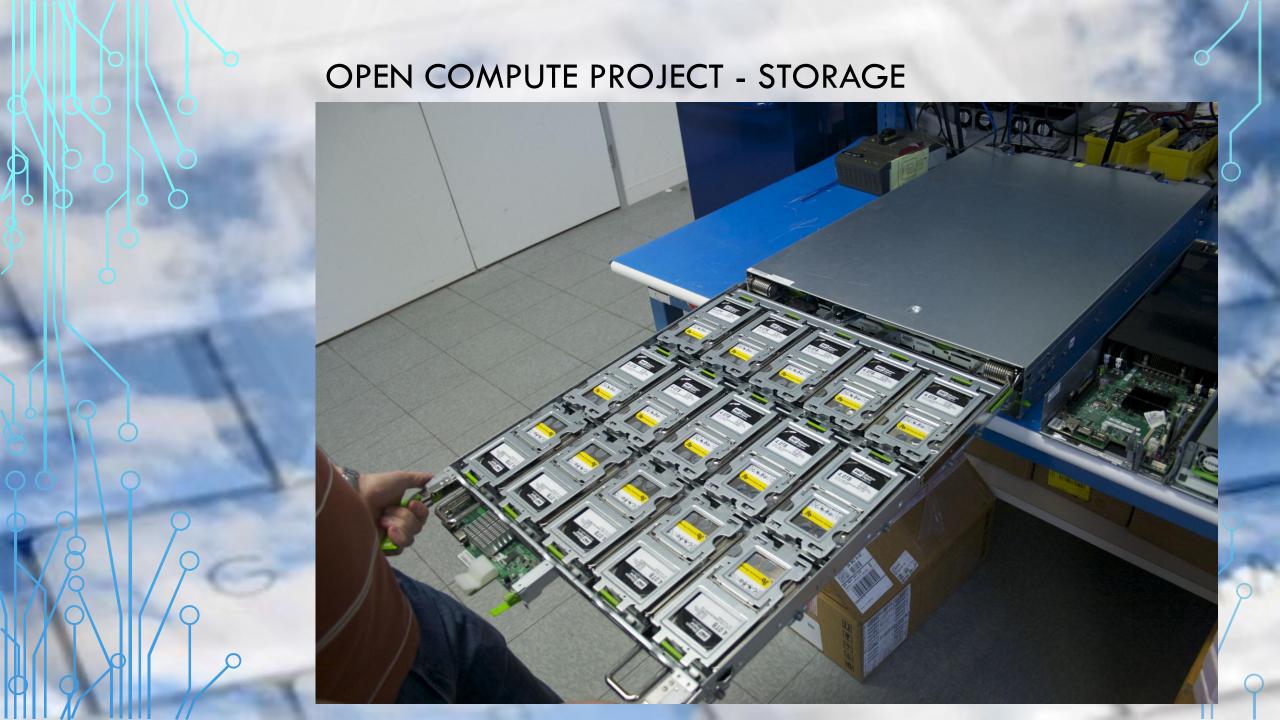


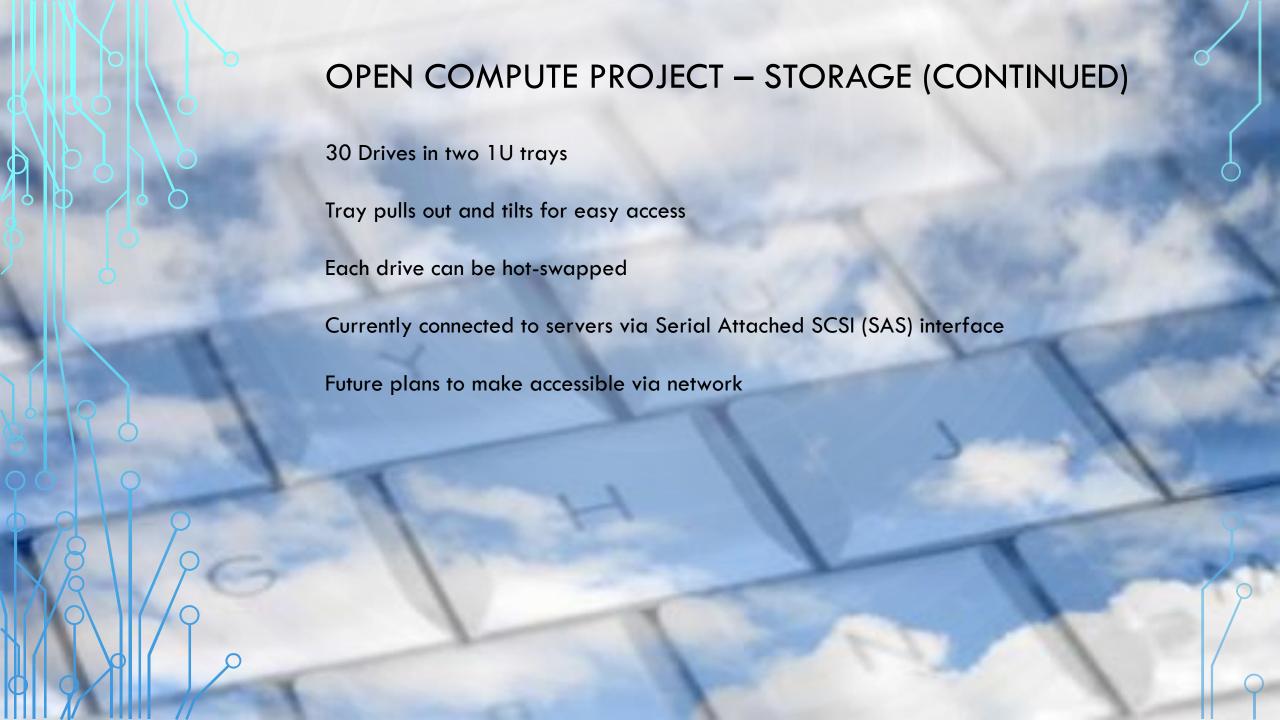
OPEN COMPUTE PROJECT Divided into 6 Focus areas Server Vanity free, lower cost Storage Low cost building blocks Data center design **Energy efficient** Networking Open network technology stack

Hardware management

Certification

Remote management tools





BACKBLAZE

BACKBLAZE - STORAGE POD



BACKBLAZE STORAGE PODS @ RPI

Petascale Storage Using Ceph And The BackBlaze Storage Pod

Tim Wickberg, Christopher D. Carothers, R. Lindsay Todd Rensselaer Polytechnic Institute

Design Goals

- Inexpensive: \$500,000
- Large capacity: Over 1 Petabyte Usable
- Fast over 5 GB/sec write, 10GB/sec read
- Reliable can loose any two storage servers and their disks

RAID Kills Performance, Use Replication Instead

A RAID array using 3TB hard drives will have at least one uncorrectable bit error during reconstruction. This drives the adoption of RAID-6.

But RAID kills performance. A single SATA disk provides 75 random 4k IOPS. A 4+2P RAID array, due to parity and stripe requirements, is limited to the same 75 IOPS write, and 300 IOPS read.

If you add replication to guard against the loss of any storage servers or arrays, you now have 12 disks as a replicated pair of RAID6 4+2P arrays providing only 75 IOPS write and 600 IOPS read, for 4/6 (RAID overhead) * 1/2 (replication overhead) = one third of the raw disk space.

Instead you can replicate data twice (2x replication). Each file written in is stored on three separate disk drives. Same one-third of the raw disks space available as usable space, but the write IOPS is no longer constrained by RAID. The same 12 disks under 2x replication instead can provide 300 IOPS write, and 900 IOPS read.

This is 300% better write performance, and 50% better read performance with the same hardware.

Ceph

Ceph is the only mainstream open-source parallel filesystem that can currently support the replication model required.

The CRUSH layer in Ceph allows us to control the data placement, such that replicas are spread out between storage pods and racks, to prevent any single pod failure from causing data loss.

BackBlaze Storage Pod

The BackBlaze Storage Pod is a commodity hardware approach to large scale storage systems.

There is no redundancy within the pod. But by aggregating a large number of these pods together with Ceph's replication mechaisms we can provide high availability at the system level.

Our pod design varies from the stock BackBlaze Pod:

- Different motherboard, with additional PCIe slots.
- 40Gbps Infiniband adapter, rather than using the Gigabit Ethernet.
- Reduced capacity, but better performance.
 Changing the SATA controllers from 3 4-port models to 4 2-port models gives us an additional 100MB/s path to the disks, at the loss of 1 SATA port (5 disks).
- · Intel SSD added in for metadata storage.



Proposed System

The proposed system includes 64 storage pods, each with 24x 3TB Hitachi Deskstar hard drives, and an Intel 300GB SSD.

Quantity	Component	Price	Total
64	Modified BackBlaze Pods	\$4,500	\$288,000
64 * 24 = 1536	3TB Hitachi Hard Drives	\$120	\$184,320
64	300GB Intel SSDs	\$500	\$32,000
			\$504,320

This results in 4.7-Petabytes of raw disk storage, and 19.2-Terabytes of SSD storage for metadata.

After replication on both the filesystem (2x) and metadata (3x), the resulting system is expected to have 1.5-Petabytes of usable storage, 4.5-Terabytes of metadata storage.

Expected performance from the system including replication overhead is:

- · 8.5 GB/sec write, 25GB/sec read
- 21k IOPS write, 64k IOPS read
- Metadata operations at 300k IOPS write, 2.5million IOPS read.



BackBlaze Pod - Photo by Chris Dag - http://www.flickr.com/photos/chrisdag/6074480108/



AARNET CLOUDSTOR+

100GB storage per researcher

Access controlled by the Australian Access Federation

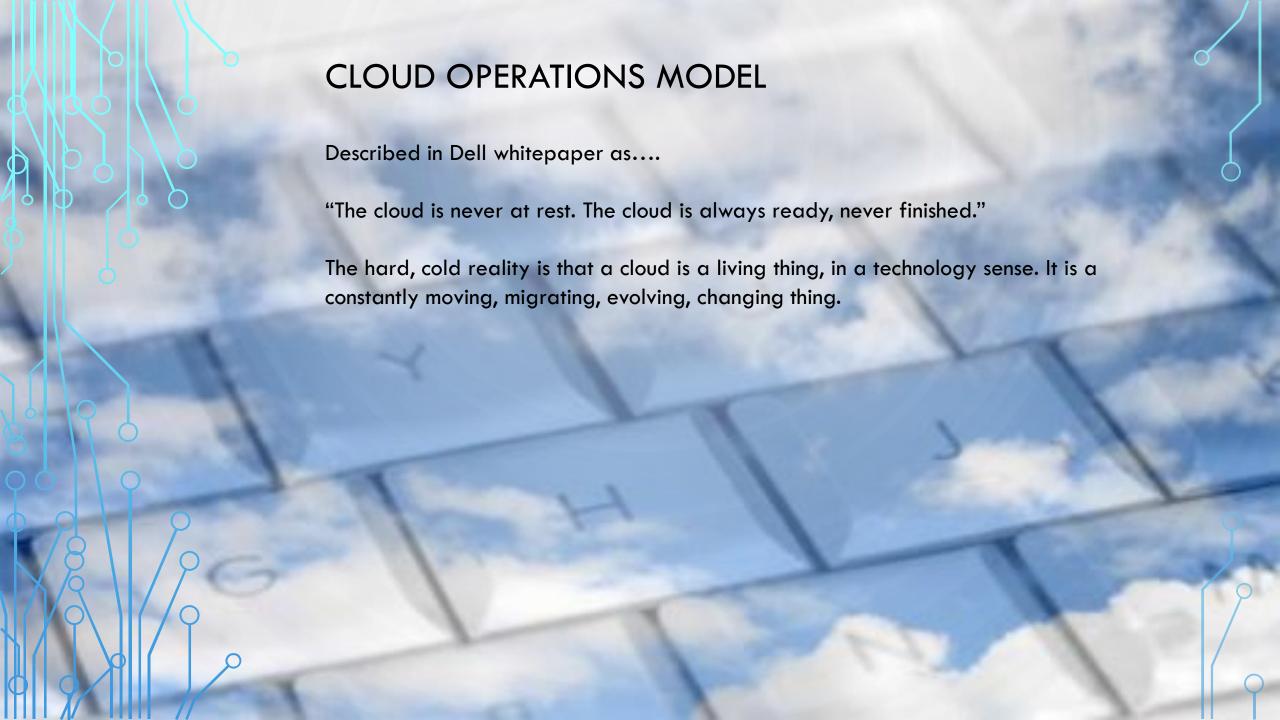
Hadoop backend (HDFS)

Objects replicated geographically, at least 3 copies
Data stored closer to the point of frequent use

Open source software used

Dropbox like functionality, with clients for Windows, Linux, Mac OS plus Android and IOS devices and a web interface

MICROSOFT'S VIEW OF CLOUD COMPUTING Service Delivery Service Information Business Continuity & Service Service Level Financial Demand Capacity Relationship Availability Catalog Lifecycle Security Management Management Management Management Management Management Management Mgmt. Management Provides Services Defines -**Defines** Service Change Reporting Management Software Service Service Asset & Manages Management Configuration System Management Service Health Release & Monitoring Provides Services Deployment Management Configuration Manag Management Operations System Knowledge Management Platform Fabric Management Incident & Manages Problem Defines: Deployment & Management Provisioning Management Request Provides Services Fulfillment Data protection Hypervisor Access Network Management Management Network Compute Storage Facilities Manages Systems Security Administration Management Infrastructure





CATEGORIES (WEBOPEDIA)

Private cloud

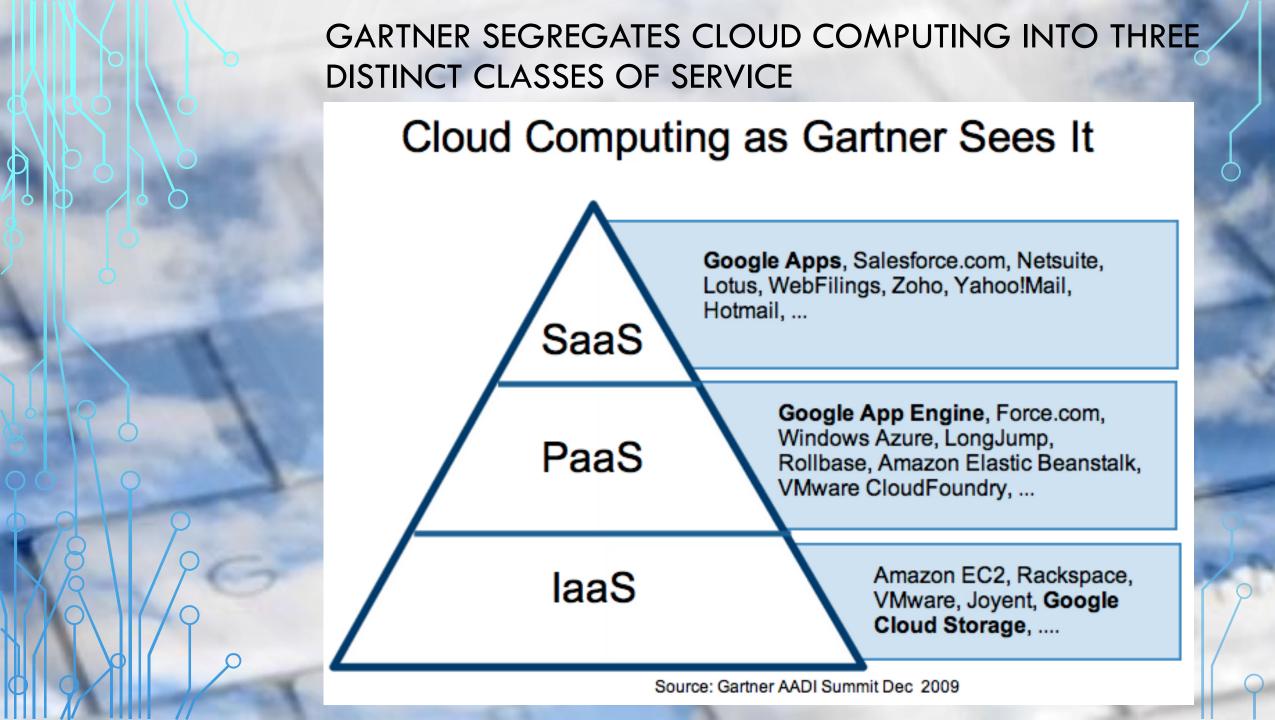
A <u>cloud computing platform</u> that is implemented within the corporate <u>firewall</u>, under the control of the IT department.

Public cloud

A form of <u>sour computing</u> in which a company relies on a third-party cloud service provider for services such as <u>servers</u>, <u>data storage</u> and applications, which are delivered to the company through the Internet.

Hybrid cloud.

A combined form of private clouds and <u>public clouds</u> in which some critical data resides in the enterprise's private cloud while other data is stored in and accessible from a public cloud.





Applications that only available online. Hotmail, Gmail, Salesforce.com etc. Accessed through a web browser or mobile app. Data stored online

Hardware and infrastructure outsourced. User is responsible for the whole stack... from operating system to application

Vendor is responsible for the platform. End user deploys application on top of the platform and is only responsible for that layer



DRIVING FORCES BEHIND CLOUD COMPUTING

- •The **emergence of new types of apps** that are inherently social (like, follow, share as verbs), that use open APIs, and new app frameworks have a notion of intelligence built right in (e.g. autocomplete in every search bar).
- •The explosion of data—both structured and unstructured—from a growing world of sources like sensors and social media is creating huge challenges for managing and processing all that data, but also creates great opportunities for smarter insights.
- •The **consumerization of IT** has resulted in a proliferation of mobile devices that users bring to work and expect IT to support. It's estimated that in 2015, the number of connected devices will be twice the number of people on the planet. Today's users expect their apps to be available on these devices with a seamless experience. However, the need for enterprise security and compliance cannot be sacrificed to meet these expectations.
- •And of course there's tremendous interest in the economic benefits of cloud computing.

DRIVING FORCES BEHIND CLOUD COMPUTING

Shared Resources

- Reduce CAPEX —run more applications on less hardware
- Multi-tenanted to improve utilization by pooling resources
- •Repeatable & reliable deployment provides agility and confidence
- Application owners only pay for allocated capacity

Always Up, Always On

- •Resilient and redundant platform for all applications
- •Reconfigure virtual hardware with zero app downtime
- •Improved business continuity helping you meet your SLA's
- •Deliver high availability with commodity hardware

Scalable and Elastic

- •Deploy apps across multiple clouds for scale and burst loads
- Meet demand without building for peak scenarios
- •A single pane of glass for app management wherever it is running
- •Deliver better experiences to web and mobile app users



COMPARING...

Traditional	Cloud	
Storage Area Network	Distributed / Local attached	
Enterprise class HDD's	Commodity HDD's	
Highly redundant	Redundancy via replication	
Quick turnaround maintenance agreements	Standard warranty	
Rigid network configuration	Software defined networking	

