WORKSHOP GARR 2024

NET MAKERS

The Requirements and the Hidden Costs of the Storage

Sustainable models for Education & Research Data

Enrico Venuto

Politecnico di Torino

Politecnico di Torino

From Flat Cloud Storage for universities to . . .

Google Disk

In the beginning: Google Drive was Flat (unlimited space for universities)

Current Approach: Every university has a small space (pay for surplus)

OneDrive Disk

In the beginning: OneDrive was Flat (unlimited disk space for universities)

Current Approach: Every university has a small disk space (pay for surplus)

Surplus: Listino CRUI 2.651,40 €/Y per 10 TB

Dropbox ……..

Keyword: data-growing

Some History

In the beginning: Every system has its Own storage — **MB-GB – 10-100 Mbps**

Past Approach: One centralized storage for all the needs — **GB-TB - 1-10 Gbps**

Current Approach: One different storage for each different need — **TB-PB – 10-100 Gbps**

Plus: Optimized for specific load/bandwidth
Cost effective

Minus: Huge quantity of data moving between different storage

Next Approach: ?? — **PB-EB – 100 Gbps-Tbps**

Fact: Data is exponentially increasing:
Current approach is going to become unsustainable

Keyword: data-growing

One different storage for each different need

Improvement of the technology system require a process of continuous revision of the technological solutions and their development model combined with the analysis of strategic opportunities

Backup solutions
NAS Storage
HPC
Private Cloud (eg OpenStack, Kubernetes)
Business continuity & Always-On
Cloud Storage

Open Source as enabling factors for sharing and collaboration

Keyword: always-on, flexibility, security, sustainability, data sovereignty

Methodological Notes

Cost to Buy (CtB):

The cost of the acquisition of the storage includes 5 years of SW and on-site NBD HW support

Yearly cost (YC):

One Fifth of the "Cost to Buy" plus the cost of one year of electricity (YE) multiplied for the Power Usage Effectiveness (PUE) of the Datacenter

$$YC = CtB / 5 + YE \times PUE$$

What's included?

Only the storage with "basic" protocol like CIFS (Samba), nfs, *fpt* or iSCSI, optionally with Compression, DeDupe, Remote Replication, Remote incremental Snapshot - No user web interface, document sharing, video streaming: they must be added.

Backup solutions

Over than 10PB NAS in a third Campus Datacenter providing **de-dupe/compressed backup** and security copies for

Databases, VMs, Research Data

Cloud Mail & Cloud Collaboration Suites

Requirements: NAS, High **Capacity**, Hi **Bandwidth**, ZFS, Low **cost**

Cost to buy: 30.000€/PB (22TB SATA disks. Cost of only disks: 22.000€/PB)

6.000€/PB/Year (life: 5years) + 1.000€ Power consumption → **7€/TB/Y**

50Gbps aggregate bandwidth

Plus: sustainability

Minus: only one controller

Keywords: security, sustainability, data sovereignty, capacity, replication

NAS Storage

Over than 4PB of High Availability NAS in two different Campus Datacenter with cross site replication and snapshot providing Always-On Data for

- Research
- Open Science, GitLab, […]

Requirements: NAS, iSCSI, **Resilience**, High **Capacity**, High **Bandwidth**, ZFS, Cost

Cost to buy: 53.000€/PB (16TB NL-SAS disks. Cost of only disks: 30.000€/PB)

10.600€/PB/Year (life: 5years) + 1.200€ Power consumption → **12€/TB/Y**

100Gbps aggregate bandwidth

Plus: sustainability, double controller active/active

Minus: Max 128 disks

HPC

1 PB NVMe BeeGFS
with **InfiniBand** & 100Gbps **RoCE** Storage Node

Requirements: Resilience, Scalability, High Capacity,
High **Bandwidth**, High **Performance**

Cost to buy: 200.000€/PB (disks NVMe)

40.000€/PB/Year (life: 5years) + ? € Power consumption → **50€/TB/Y**

Tbps aggregate bandwidth

100 nodes
5.000 core
50 TB RAM
100 GPU

Plus: performance, BandWidth
Minus: Cost

Keywords: HPC, data sovereignty, OpenSource

Private Cloud – OpenStack & K8S - Ceph

kubernetes

OpenStack: Virtual Datacenter, Kubenetes clusters
100/400 Gbps connectivity

Requirements: Resilence, Scalability, Performance, High Bandwidth

Ceph **1.200 disks** (SAS, 1.8TB -> 2,16 PB)

Cost to buy: ??? Circa 200.000 €/PB

40.000 €/PB/Year (life: 5years) + ? € Power consumption → **? 50€/TB/Y**

? aggregate bandwidth

85 nodes
14.000 core
75 TB RAM
1.200 disks

Plus: ∞ Scalability, Vendor independent, Eterogenous Hardware

Minus: Hi storage replication, Hi CPU-Memory-Network usage,
Low Performance

Keywords: IaaS, federation, data sovereignty, OpenSource

Private Cloud – OpenStack & K8S

From Ethernet Fabric to IP Fabric

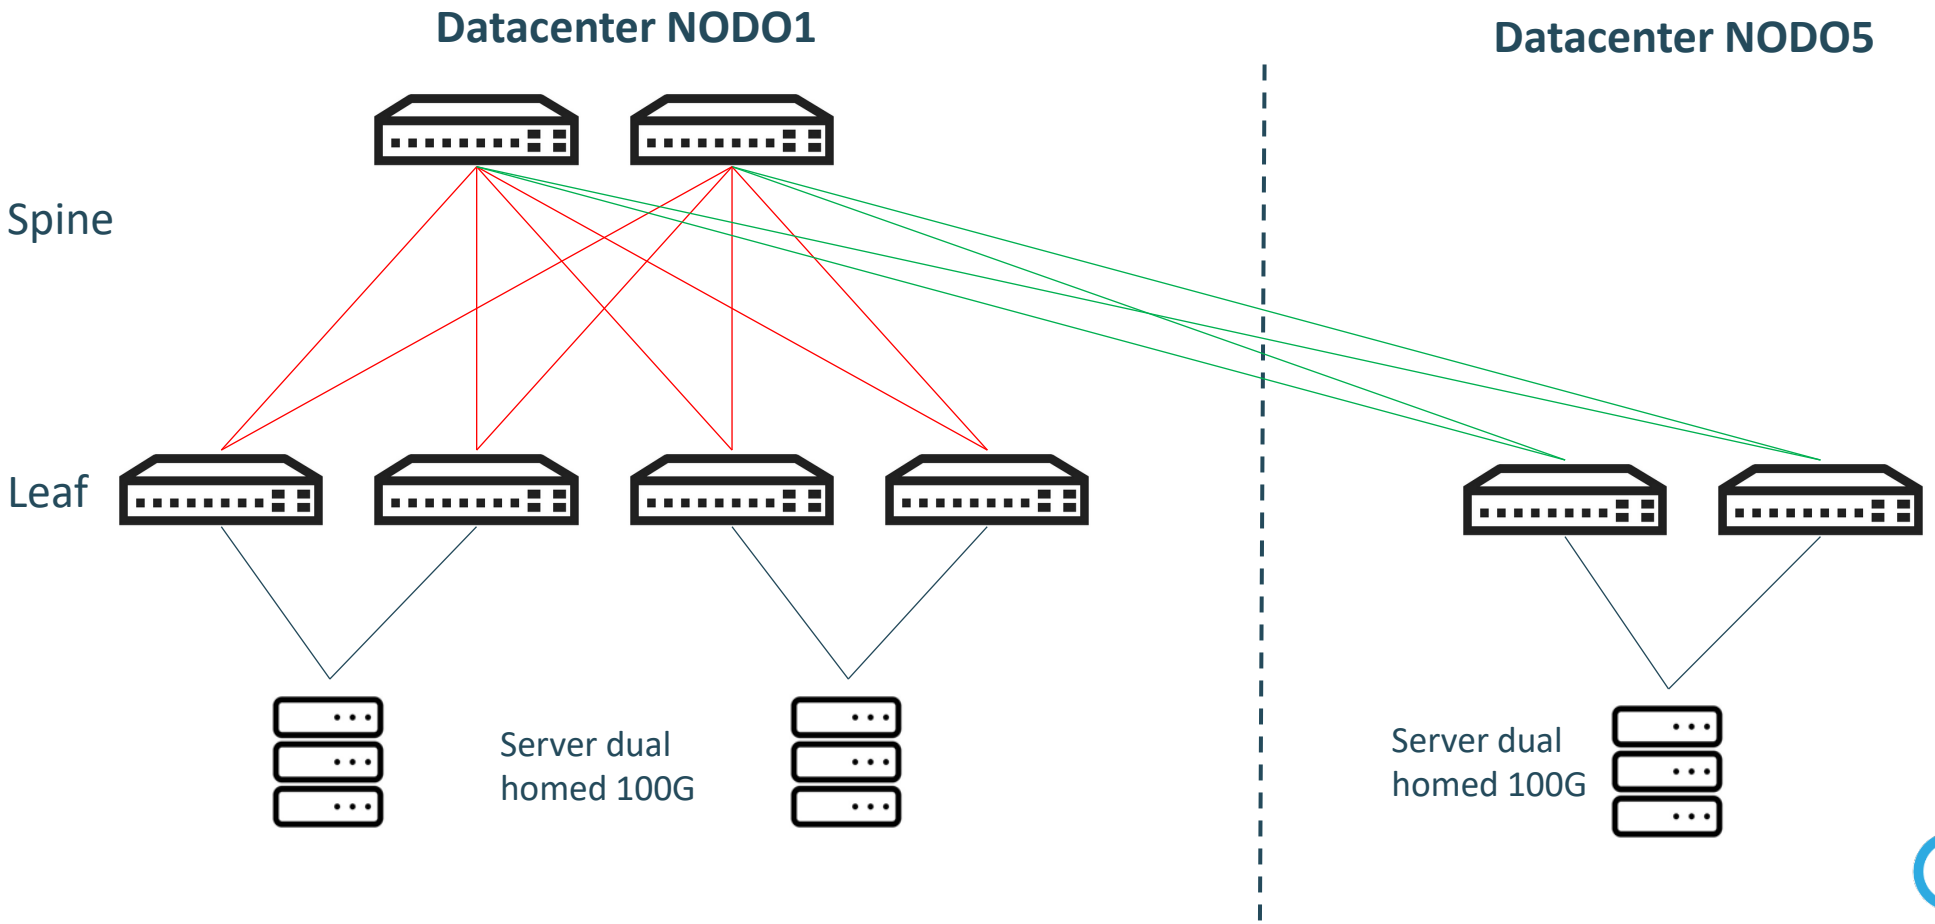

Replica 3

(3 copies of data)

Spine&Leaf - Buy: 140k€

Cloud@PoliTO

400 Gbps AOC 3 mt

400 Gbps con fibra monomodale e ottiche QSFP DD FR (LC max 2KM)

100 Gbps AOC 3-5 mt

Private Cloud – OpenStack & K8S

From Ethernet Fabric to IP Fabric

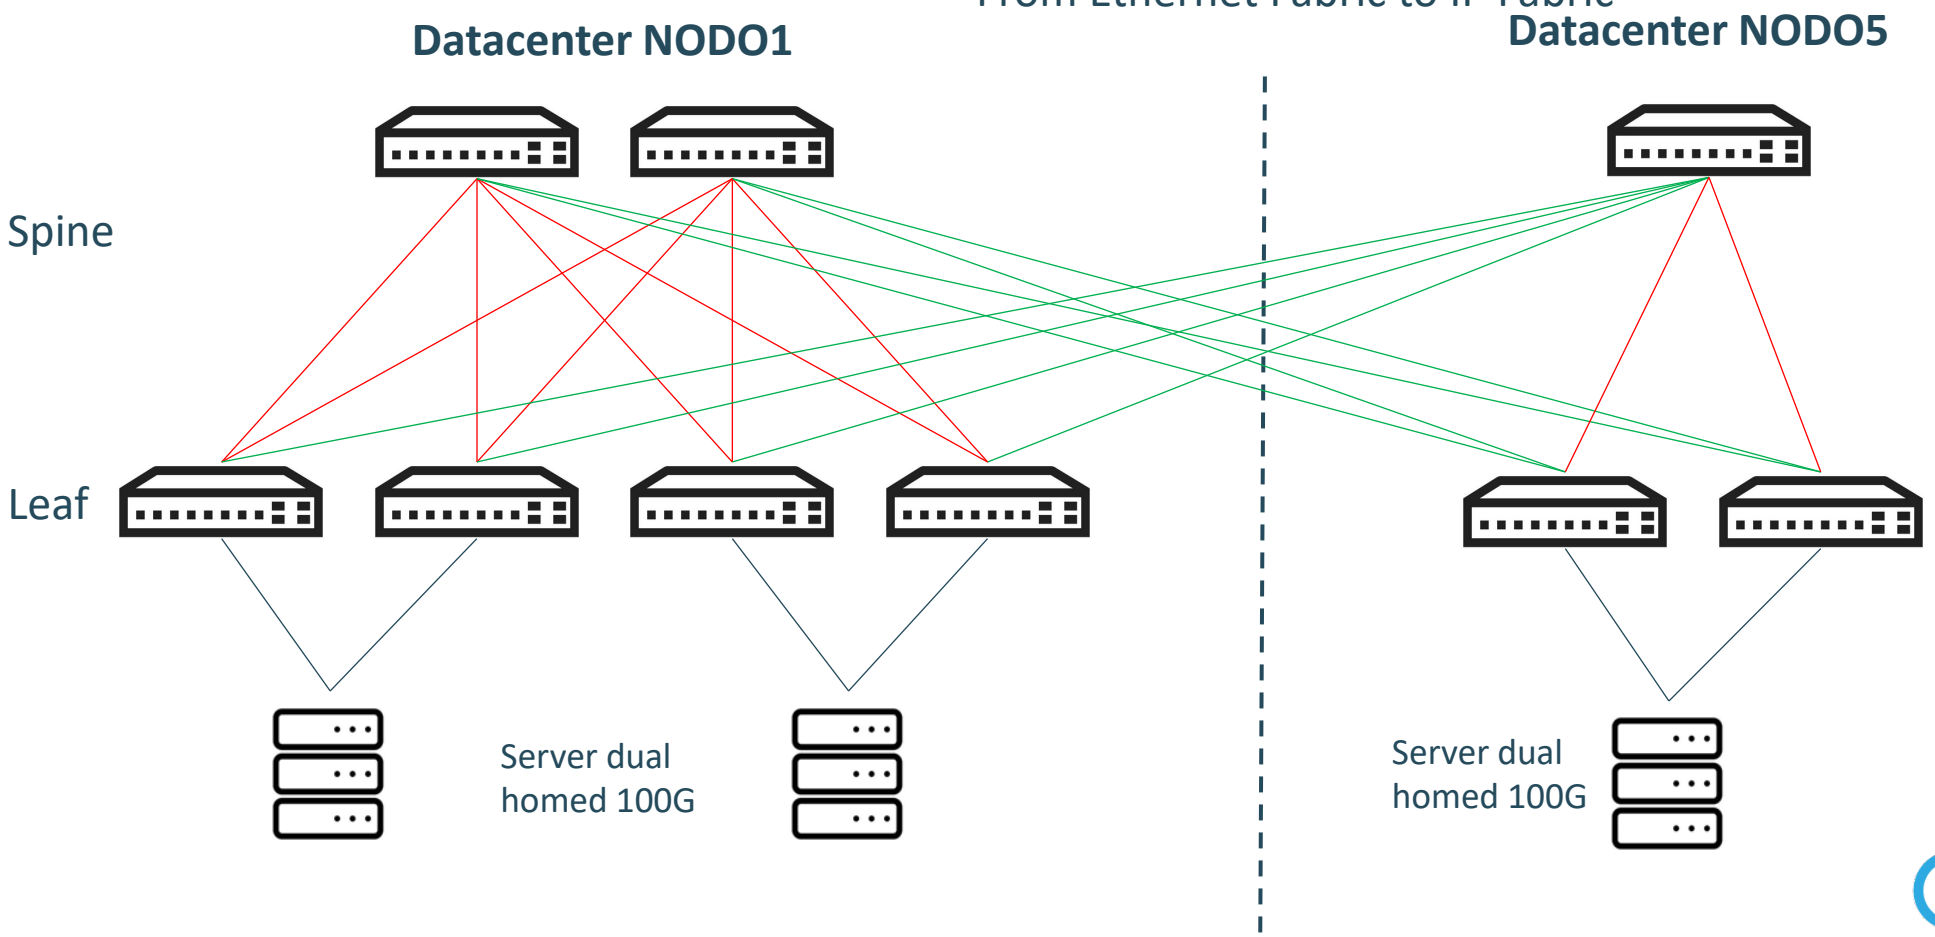

400 Gbps AOC 3 mt

400 Gbps con fibra monomodale e ottiche QSFP DD FR (LC max 2KM)

100 Gbps AOC 3-5 mt

Replica A4

(2x2 copies of data)

Spine&Leaf - Buy: 190k€

Require 3 Datacenter

Cloud@PoliTO

Private Cloud – OpenStack & K8S - Cinder

OpenStack: Virtual Datacenter, Kubenetes clusters (on VM)
100/400 Gbps connectivity

Requirements: Resilence, Scalability, Performance, High Bandwidth

kubernetes

Cinder

Cost to buy: 50.000€/PB

10.000€/PB/Year (life: 5years) + 2.000 € Power consumption → **12€/TB/Y**

500 Gbps aggregate bandwidth

Plus: Sustainability, Hi BandWidth, Multiple controller

Minus:

Keywords: IaaS, federation, data sovereignty, OpenSource

Business Continuity for Mixed Load

A couple of storage located in 2 different Campus Datacenter, Operating in Business continuity, Providing sufficient performance & storage space for mixed load (Hosting of core applications, Databases, VMs, Application Server)

Requirements: Hi **Resilience**, High **Performance**, High Bandwidth, **Scalability**
Cost to buy: 300.000€/PB (disks NL-SAS & SSD-SAS)

60.000€/PB/Year (life: 5years) + 10.000€ Power consumption → **70€/TB/Y**

100Gbps aggregate IP bandwidth

Plus: Business Continuity, multiple replicated controller

Minus: Cost

Keyword: availability, always-on, always-growing

Public Cloud Storage – OneDrive

10 PB Cloud Storage for mixed use

Requirements: NAS, High **Capacity**, Hi **Bandwidth**, Low **cost**

OneDrive:

265.000€/PB/Year (Listino CRUI) → **265€/TB/Y**

?? aggregate IP bandwidth

Plus: nothing to do
Minus: very expensive, low performance, low bandwidth for mixed use

Keyword: Public-Cloud

Public Cloud Storage – AzureFile v1 model

10 PB Hi performance Cloud Storage for mixed use

Requirements: NAS, High **Capacity**, Hi **Bandwidth, Hi Performance**

Azure File Provisioned v1 model

SSD, 3 Year Reservation term, provides storage, IOPS and throughput in a fixed ratio to each other, similar to how storage is purchased in an on-premises storage solution.

Capacity (GiB)	Baseline IOPS	Burst IOPS	Burst credits	Throughput (ingress + egress)
102,400	100,000	Up to 100,000	0	10 GiB/s

11.268€/100TB/mese → 13.520.000€/10PB/Y → **1.352€/TB/Y**

800Gbps IP bandwidth

Plus: nothing to do

Minus: very expensive, require full Cloud Adoption

Keyword: Public-Cloud

From Nominal to Usable Space

Formatted Space → about 90% of Nominal Space

Protected Space → portion of Formatted Space based on RAID protection & HotSpares overhead

Usable Space → portion of the Protected Space based on Replication Factor

Corrective Factor → Multiplier to transform Nominal Cost in Effective Cost

Keyword: Sustainability

From Nominal to Effective Cost – 1 TB

		Nominal cost €/TB/Y	Usable space	Corrective Factor	BandWidth Gbps	Effective Cost €/TB/Y
OneDrive		**265**	100%	1	low	**265**
Azure File v1 model		**1.352**	100%	1	800	**1.352**
Backup		**7**	75%	1,33	50	**9**
NAS	*Unique*	**12**	75%	1,33	100	**16**
	Replicated		37,5%	2,66	100	**32**
HPC		**50**	75%	1,33	1000+	**66**
Openstack Ceph	*Unique*	**-**	30%	**3,33**	-	**166**
	Replicated		22,5%	**4,44**	-	**222**
Openstack Cinder	*Unique*	**12**	75%	1,33	500	**16**
	Replicated		37,5%	2,66	500	**32**
BusinessContinuity Mixed Load		**70**	75%	1,33	100 + FC	**93**

From Nominal to Effective Cost – 10 PB

		Nominal cost €/TB/Y	Usable space	Corrective Factor	BandWidth Gbps	Effective Cost €/TB/Y	Effective Cost €/10PB/Y
OneDrive		**265**	100%	1	low	**265**	**2.650.000**
Azure File v1 model		**1.352**	100%	1	800	**1.352**	**13.520.000**
Backup		**7**	75%	1,33	50	**9**	**90.000**
NAS	*Unique*	**12**	75%	1,33	100	**16**	**160.000**
	Replicated		37,5%	2,66	100	**32**	**320.000**
HPC		**50**	75%	1,33	1000+	**66**	**660.000**
Openstack Ceph	*Unique*	**-**	30%	**3,33**	-	**166**	**1.660.000**
	Replicated		22,5%	**4,44**	-	**222**	**2.220.000**
Openstack Cinder	*Unique*	**12**	75%	1,33	500	**16**	**160.000**
	Replicated		37,5%	2,66	500	**32**	**320.000**
BusinessContinuity Mixed Load		**70**	75%	1,33	100 + FC	**93**	**930.000**

Aphorisma's corner on CONSTANTS

NET MAKERS

The only Constant in Life is Change *[Eraclitus]*

Aphorisma's corner on CONSTANTS

NET MAKERS

The only Constant in Life is Change *[Eraclitus]*

The only Constant in Data is Growing *[Enrico]*

Aphorisma's corner on CONSTANTS

NET MAKERS

The only Constant in Life is Change *[Eraclitus]*

~~***The only Constant in Data is Growing*** *[Enrico]*~~

The only Constant in Data and its Cost is Growing *[Enrico – Two year later]*

Domande?

wooclap.com
Codice: WSGARR24

WORKSHOP GARR 2024

NET MAKERS

Thanks

www.linkedin.com/in/enricovenuto

Enrico Venuto
Politecnico di Torino

Politecnico di Torino