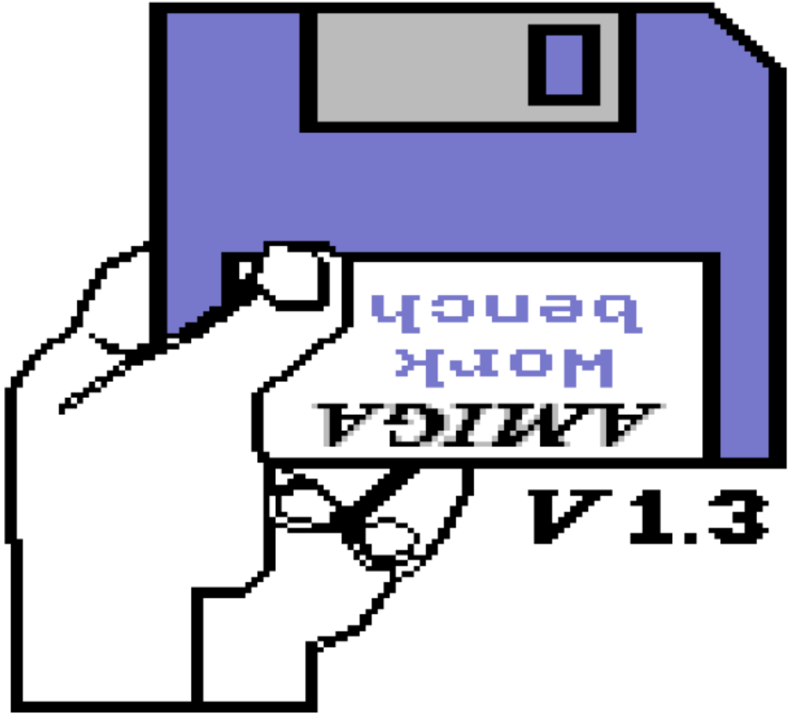
AMIGA
work
bench
V1.3

Storage Cost Modeling

PASIG 2019

Julian M. Morley, Digital Library Systems & Services
jmorley at stanford dot edu
@JulianMMorley

Preservation Storage Cost Modeling

PASIG 2019

Julian M. Morley, Digital Library Systems & Services
jmorley at stanford dot edu
@JulianMMorley

Preservation Storage Cost Modeling at Stanford

PASIG 2019

Julian M. Morley, Digital Library Systems & Services
jmorley at stanford dot edu
@JulianMMorley

Digital Preservation at Stanford

Stanford Libraries runs SDR:
the **Stanford Digital Repository**.

Over 600TB of data projected to double in size in 2019.

Our digital collections are **a strategic resource** for the university **that must be preserved** for future generations.

Ground Rules

Preservation is meaningless without **access**

Preservation means multiple copies

Data sovereignty is a strategic imperative

Continuity planning is vital for long-term preservation

Ground Rules Applied

Access: At least 1 copy on spinning disk.

Sovereignty: At least 1 copy on Stanford hardware.

Preservation: At least 3 copies (we want 4).

Preservation: At least 2 different out-of-state sites.

Preservation: Regular content audits & fixity.

Continuity: Avoid vendor monoculture.

Continuity: Have a plan to add and drop vendors.

SDR Preservation Core

Campus Datacenters

Online copy
provides
Access & Preservation

Stanford, CA

US-West-2, OR

AWS

Archive #1
AWS Glacier

Dallas, TX

IBM Cloud
Object Storage

Archive #2
COS Vault
Archive

Third Vendor

???

Archive #3

Offsite

Archive Storage: the Cloud!

Any cloud provider that has:

- S3-compatible API
- A ‘cold’ or ‘cloud tape’ product
- U.S. Data centers outside California
- Published pricing, no significant ingest costs

Amazon, IBM, Oracle & Wasabi* meet these criteria.

Dell & Iron Mtn have an S3 cloud product, but no cold tier.

Microsoft, Google & Backblaze have no S3 API.

Disk Storage: local SDS

Software Defined Storage (SDS) running on commodity server hardware.

- SDS used by all cloud providers; it works. It scales.
- Allows use of S3 API across local and cloud storage.
- Buy hardware from any vendor, in small batches.
- Commercial and Open Source solutions available.

Local SDS Cost Model

Assuming a hypothetical Ceph cluster.

Need to know:

- A node configuration, price & capacity it provides
- The amount of new storage being added to the cluster per year
- Amount of old storage that must be retired (aged out) each year
- Total number of in-service nodes, and the associated datacenter costs
- Licensing and support costs

Assuming:

- Prices in current (2019) dollars
- Storage costs will decrease over time, and storage density will increase.

Storage Node Configuration

Based on Ceph Storage Recommendations & SMC reference architecture.

SuperMicro Superstorage Server 6028R-E1CR24L
24x 3.5" HDD, 20 cores, 256GB RAM, 2x 1TB NVMe SSD cache.
$25K in 2019 - assuming price remains flat in 2019 dollars.

	Year 1 model	**Year 2 model**	**Year 3 model**	**Year 4 model**	**Year 5 model**
Disk size	10TB	12TB	14TB	16TB	18TB
Total raw TB	240TB	288TB	336TB	384TB	432TB
Net TB	**144TB**	**173TB**	**202TB**	**230TB**	**260TB**

Projected Cluster Growth

Assuming an initial cluster size of 1PB and a 30% YoY data growth rate.
At the end of each year we need at least this much provisioned capacity, plus 10%+ headroom.

	Year 1	Year 2	Year 3	Year 4	Year 5
preserved data	1000 TB	1300 TB	1690 TB	2200 TB	2860 TB
required cluster capacity	> 1100 TB	> 1430 TB	> 1860 TB	> 2420 TB	> 3150 TB

Bring it together

	Year 1	Year 2	Year 3	Year 4	Year 5
preserved data	1000 TB	1300 TB	1690 TB	2200 TB	2860 TB
required SDS capacity	> 1100 TB	> 1430 TB	> 1860 TB	> 2420 TB	> 3150 TB
New cap/yr	1100 TB	330 TB	430 TB	560 TB	730 TB
Replacement cap req'd /yr	0	0	0	1100 TB	330 TB
current year node gives:	144TB each	173 TB each	202 TB each	230 TB each	260 TB each
Need to add:	8 nodes (1152 TB)	2 nodes (346 TB)	5 nodes (1010 TB)	8 nodes (1840 TB)	5 nodes (1300 TB)

Simple Preservation Cost Model

Disclaimers:

- This is a model for budgetary planning purposes only.
- All vendor prices are from publicly available sources.
- Stanford's internal model is a little different; this spreadsheet does not reflect what Stanford is spending.
- Internal staffing costs are left as an exercise for others.

http://bit.ly/PASIGStorageModel

In Conclusion: