


Can Disks Replace Tape (or Cloud) for Long-Term Storage?

The Perspective from SGI: we sell stuff that creates an awful lot of data, but not really a "Storage Company"



sgi


Can Disk Replace Tape?



<http://www.haystack.mit.edu/tech/vlbi/mark5/>

- Yes, of course it can. In some ways it already has. But *should* disk replace tape?
- Many people focus on the costs when having this discussion
- HSM was created as a hybrid solution leveraging the performance of disk and cost advantages of tape

Can Disk Replace Tape?



<http://www.forbes.com/sites/tomcoughlin/2014/12/30/archiving-storage-tiers-part-2/>

- Really the conversation should be about using the best tool for the job
- Tape plays a role, but it is a role that is getting smaller as disk architectures expand capabilities
- Disk already being replaced by Flash at the higher boundaries of these roles – how long before we are asking “Can Flash Replace Disk (or Tape or Cloud)?”

What about Cloud?



<https://thelocalbrand.com/take-advantage-cloud-backup-services-business/>

- Cloud is disk (or tape possibly, but everyone is using disk) only with better Marketing...
- Cloud is an access method, possibly a paradigm shift, but not a storage technology
- Cloud works if you have Small Data, are Very Patient, and/or have Ridiculous Bandwidth on tap
- Cost advantages at Petascales? No.

“Supposed” advantages of each platform



	Cost	Speed	Latency	Power	Density	Disaster Recovery / Business Continuity	Forward Migration	PITA Factor
Tape	😊	✗	✗	😊	😊	😊 Send Offsite	Ugh!	✗
Disk	✗	😊	😊	✗	✗	😊 Replication	Meh...	😊
Cloud	?	✗	✗	N/A	N/A	😊 It's Magic!	Don't care!	😊😊

SGI Cold Storage Elements



- Really dense disk arrays (4U84, 4U96)
- SAS 3.0 JBOD, hang 3PB from each server
- Remember COPAN? It's SGI ZeroWatt™ now!
- Introducing JBFS: a “filesystem” created specifically for mounting drives like tapes, rich capabilities for power management, and active validation/repair of data
- Integrated and Certified with DMF: HSM, meet your new best friend. Make Flash look bigger & cheaper.
- Cost, Power, and Density of Tape. Speed, Latency, and low PITA factor of Disk. Add a cloud-access model and you might be on to something here....

SGI JBFS


XFS Feature Profile	DMF/JBFS Feature Profile
A large number of files	Small number of volumes
File sizes change	Volume sizes fixed
Flexible directory structure	Flat volume structure
Arbitrary organization of files	Fixed organization of volumes
Primarily random access	Primarily sequential access
Bursty access	Sustained access
Mount/dismount infrequently	Mount/dismount frequently

Other Benefits

- Recoverability
- Data Assurance
- “High” Performance
- Flexibility
- Zero-Watt™

Sample performance comparisons

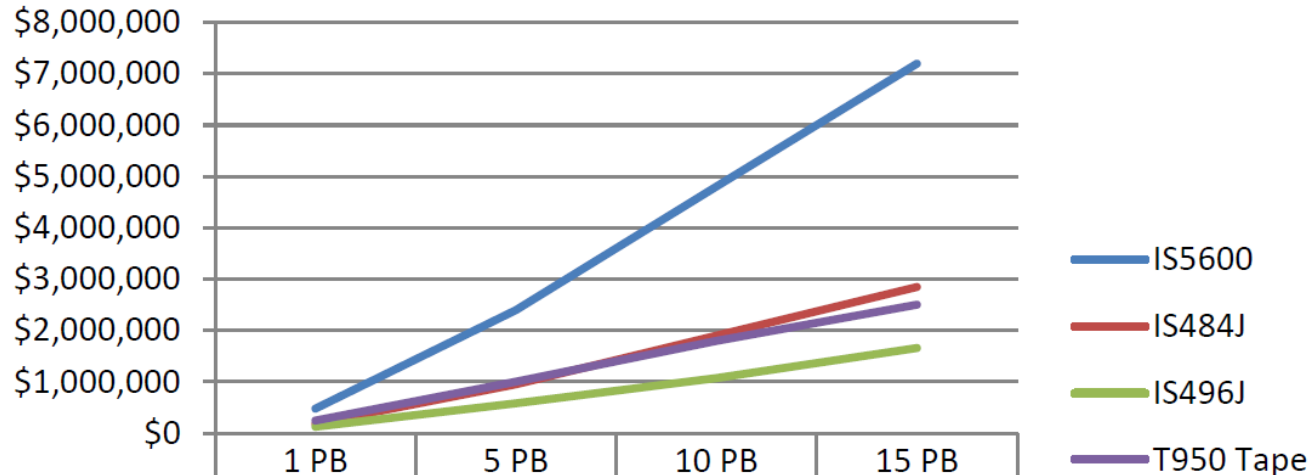
Throughput GB/Sec by Storage Type



- 1PB to 15PB capacity modelled
- No compression, de-dupe, etc.
- SLC T950 with LTO-6 tapes (2.3TB)
- 4 tape drives per PB of capacity
- (IS5600 performance virtually identical to IS484J)

Sample price comparisons

Storage Cost for Total Capacity




	1 PB	5 PB	10 PB	15 PB
IS5600	\$480,000	\$2,400,000	\$4,800,000	\$7,200,000
IS484J	\$190,030	\$950,150	\$1,900,300	\$2,850,450
IS496J	\$130,000	\$579,124	\$1,075,516	\$1,654,640
T950 Tape	\$250,000	\$1,000,000	\$1,800,000	\$2,500,000

- 1PB to 15PB capacity modelled
- No compression, de-dupe, etc.
- SLC T950 with LTO-6 tapes (2.3TB)
- 4 tape drives per PB capacity

Sample price comparisons

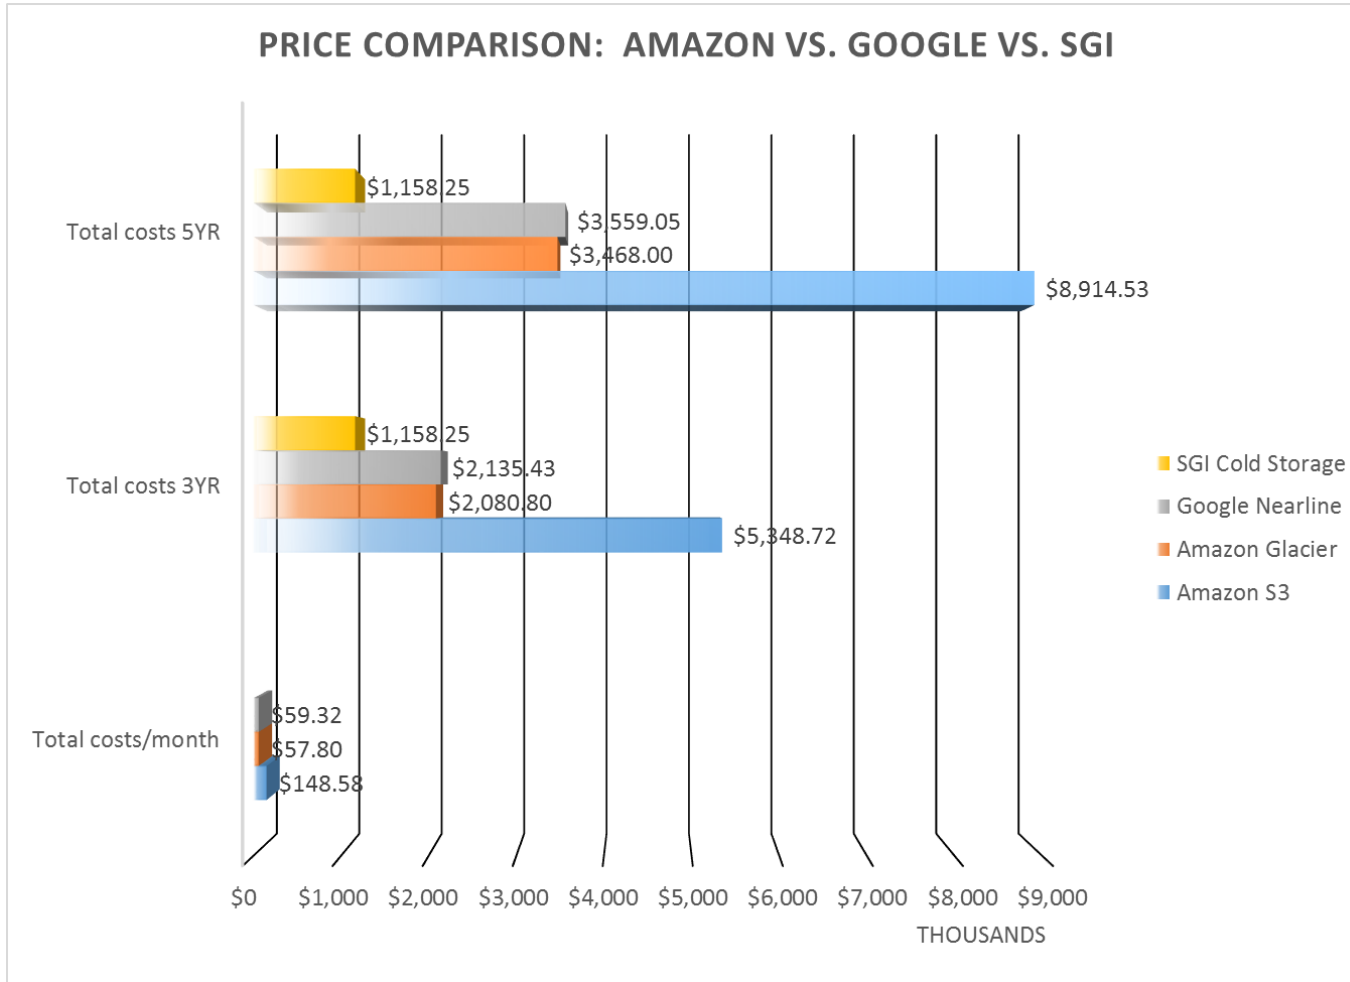
Storage Cost \$/GB



- 1PB to 15PB capacity modelled
- No compression, de-dupe, etc.
- SLC T950 with LTO-6 tapes (2.3TB)
- 4 tape drives per PB capacity

	1PB	5PB	10PB	15PB
IS5600	\$0.48	\$0.48	\$0.48	\$0.48
IS484J	\$0.19	\$0.19	\$0.19	\$0.19
IS496J	\$0.13	\$0.12	\$0.11	\$0.11
T950 Tape	\$0.25	\$0.20	\$0.18	\$0.17

What about cloud?



- 5PB capacity modelled, two copies of DMF data, separate hardware and separate locations (replicated)
- 3- and 5-year lifespan modelled
- Google Nearline, Amazon S3 and Glacier assume some free access/retrieval, basic data request charges, 100TB/month charged retrieval (2% of total), and data transfer from least expensive Americas/EMEA data centers