WHITE PAPER

# *Dramatically Lowering Storage Costs*

*2012*

*Peter McGonigal* petermc@sgi.com
*James Hill* jamesh@sgi.com

Version 7

Analysts and observers of today's digital age are forecasting exponential growth of data volumes and this is especially prevalent in the internet, telecommunications, technical and research areas. The challenge for organisations is both in storing and extracting value from these vast quantities of data. This paper looks at available techniques and solutions to reduce the cost of managing, storing and curating these large data volumes.

Data storage costs are more than the cost of the disk or tape infrastructure holding the information, it includes the electricity to power and cool the equipment, the people to manage and curate the data as well as techniques and software to store the information so that it can be successfully located and reused in the future for business value.

While storage devices are rapidly increasing in capacity and the cost of disks are increasingly commoditised, forward projections of data volumes and the escalating cost of power are far outstripping the gains made by the technology increases.

Nearly all data sources from any field considered, across civilian, research, medical and military, have diverse data sources ranging from consumer devices such as digital and video cameras, medical devices such as genomic scanners, electron microscopes and widespread data sensors for all manner of purposes. These are increasing significantly in both number and resolution while becoming more affordable resulting in a higher uptake across the world and greater storage challenges.

As more and more data is collected, the cost of data curation is becoming one of the major considerations of data custodians. For some organisations the cost of culling the data is more than the cost of storing it. In health sciences, data may need to be kept for 30 years, or for the life of the patient. For researchers, data may need to be kept for seven years past the date of publication of research and scientific papers, while regulatory compliance may require some data to be stored for even longer periods.


In June 2011, IDC released the findings of a study that found:

- The world's information is doubling every two years
- By 2020 the world will generate 50 times the amount of information and 75 times the number of "information containers" while the number of staff to manage this growth will be increased by less than 1.5 times.

Source:
http://www.computerworld.com/s/article/print/9217988/World_s_data_will_grow_by_50X_in_next_decade_IDC_study_predicts

To deal with these increases and effectively manage costs, organisations and data custodians need to have cost effective strategies for:

- Managing long-term data accessibility and retention
- Cost effectively growing storage capacity to keep pace with data growth
- Automating storage management
- Minimising the cost of curating and managing large volumes of data.

Simply adding additional spinning disks for storage capacity is not a viable long term strategy as this requires substantial expenditure on capital items, power and cooling, computer room floor space and data administration resources. Another challenge from growing online storage capacity is the ability to manage excessive backup and recovery times for large data repositories.

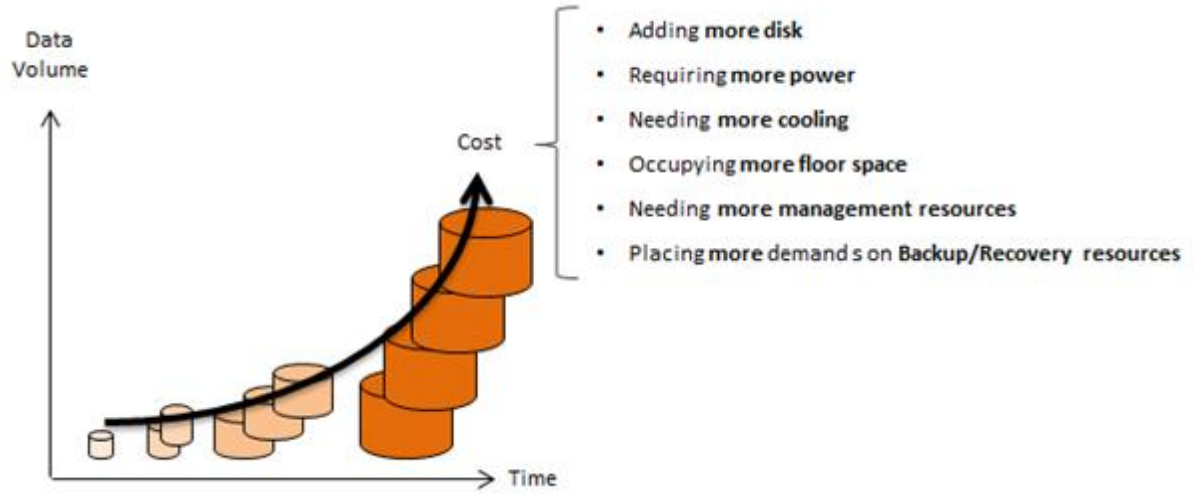


Understanding the value of an organisation's data over time is the key to selecting an appropriate solution.

Typically the value of data reduces over time, as it is accessed frequently when created and less frequently as time goes on. For more information on Time Value of Data, please refer to SGI's white-paper (http://www.sgi.com/products/storage/whitepapers). When the spectrum of data stored is considered against time, new current and high value data is typically only a small portion of the overall data being stored.

Companies and organisations need to question why non-current, long-term and infrequently used data needs to be stored on fast, expensive spinning disk storage along with current and frequently used high value data. Typically it doesn't and by utilising Tiered Storage Virtualisation, data can

be automatically moved between storage tiers that have different performance characteristics and operating costs. Storage tiers can range from high performance (high-power high cost) storage to lower performance (lower power and cost) spinning or idle disk and finally out to a much lower cost tape platforms that provide large storage volumes at very low cost.

When using a Tiered Storage Virtualisation solution, all data stored remains visible and accessible to users and applications, however the bulk of data that is infrequently used and of lower value can be stored on higher capacity, lower performing and lower cost storage. All data is typically copied to tape (and MAID if it is included in the configuration) within 24 hours while leaving the primary copy on high performance disk. As the primary disk storage is required for new data storage, data blocks are freed up leaving the file indexes and directory structures intact on the primary file system. As data ages and is used less frequently, it is gradually displaced from the lower disk tiers leaving the copies on tape. Data migration is automatic and under control of policies and rules set by the system administrators. Because the file catalogues and directories are held online, a user is typically unaware that the data maybe on a slower tier of disk or on tape and accesses files as if they are still online. Determining where the data actually resides requires special commands or command options.

Whenever a user accesses the data stored on these lower tiers it is automatically and transparently recalled onto the primary storage location – the only difference users might notice is a delay in the initial access time for less frequently used files. These techniques ensure that the least frequently used data is stored and managed at the lowest possible storage cost while frequently used and higher value data is stored and managed on fast storage, providing a cost effective and elastic storage capability without sacrificing accessibility for users.

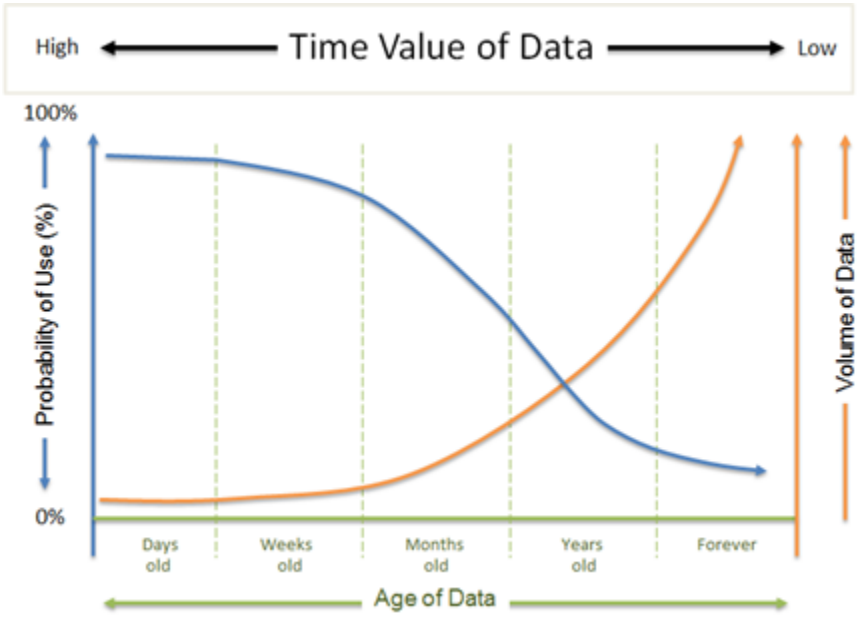


There are three broad categories of storage:

1. **RAID** (Redundant Array of Independent Disks)

   RAID Array Storage is typically a collection of drives that can be made up of SSD (Solid State drives), high performing SAS (Serial Attached SCSI), and/or large capacity SATA (Serial ATA) drives.

2. **MAID** (Massive Array of Idle Disks)

   MAID is a collection of drives that can powered down until needed. Typically only about 25% of the drives are active at any one time, the remaining drives and associated electronics are in a powered off state.

3. **Tape**

   Tape storage silos provide an automated low cost reliable storage solution. A number of different tape media technologies are available, with the most common being LTO (Linear Tape-Open technology). The LTO-6 technology holds 2.5TB of uncompressed data per tape. More information on LTO is available on

   http://www.lto.org/technology/index.html

   Other tape technologies such as Oracle's T10000c hold even greater uncompressed capacities.

The table below provides a broad overview of these three different categories of storage.

| | RAID | MAID | TAPE |
|---|---|---|---|
| **Latency** (Average Access Time to first byte) | 12ms | 15s | 70s |
| **Bandwidth** (Thru-put performance) | 20GB/s | 6.4GB/s | 45TB/hr (12.5GB/s) |
| **Capacity** | 2.5PB in 50U (840x 3TB SATA) | 2.7PB in 48U | 6.7PB/single frame (950x LTO6 tape cartridges) |
| **Power** | ~20kW | 3.9kW operational 2.1kW stand-by | < 1kW max. |
| **Life span of media** | ~4 to 5 years | 20 years | 30 years |

There is a large difference in the cost of power to store the data on the different technologies and this makes a significant difference in operational costs when considered over the extended life span of data in a global environment where power costs are escalating.

This promotes a cost effective storage strategy that incorporates locating high value, frequently accessed data on high performing RAID type storage with the remaining data stored on higher capacity, lower cost MAID and/or Tape storage, depending on the data access and use profile. Migrating data from one tier to the other is carried out automatically and transparently by the data management software. The graph below show how a single pool of data might be virtualised.

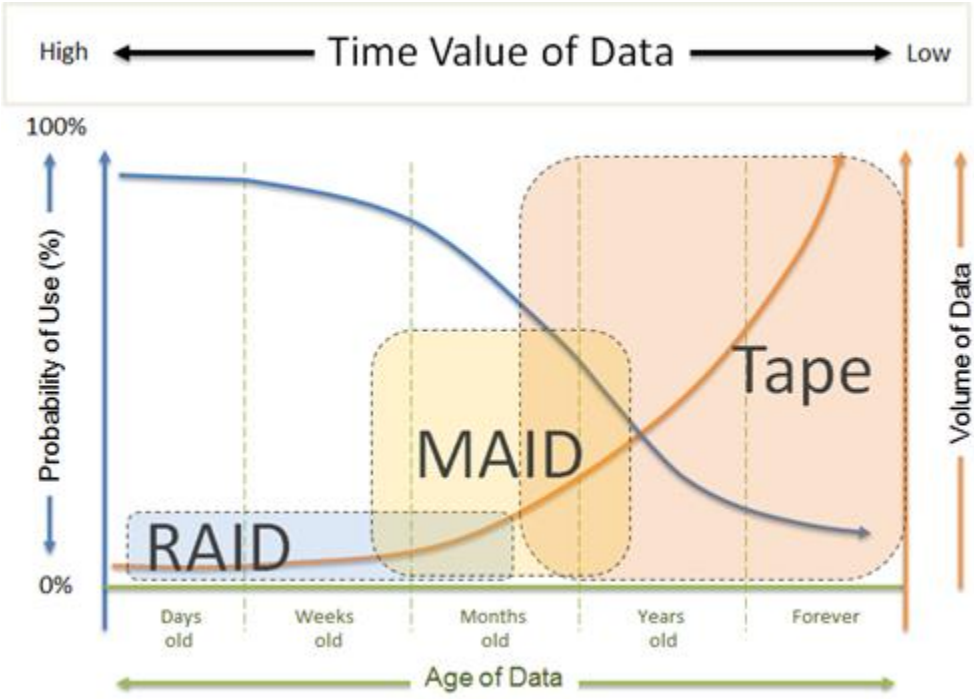


By way of an example, the following table compares 50 Petabytes of data stored using RAID storage compared to a combination of both RAID and Tape storage (ie. Tiered Virtualised storage).

| | Price | Raw Capacity | $/TB Raw | Usable Capacity | $/TB Usable | Number of Racks | Est. Power (max kW) |
|---|---|---|---|---|---|---|---|
| RAID | $17M | 67PB | $251 | 50PB | $338 | 36 | 336kW |
| RAID & Tape | $4.67M | 52PB | $93 | 50PB | $93 | 16 | 38kW |

Both the initial purchase price and the going infrastructure management costs from floor space, power and cooling make the tiered storage solution a compelling choice, with a combination of RAID and Tape for storing the data. Data residing on both RAID and Tape remains immediately accessible and visible to users. The time to complete the initial read of data located on powered-off MAID disk will be approximately 16 seconds (spin up time) and for data on a tape that has to be loaded into the tape slot by a robot taking up to a minute. Data sets that are anticipated being used in the near future can easily be scheduled to be migrated to fast spinning disk prior to user access to avoid any access lag to data.

As the majority on data is always on nearline tape and only very recent data is located exclusively on online RAID storage, the time to back up the data is dramatically reduced. Nearline storage does not need to be backed up as the data is already on multiple tapes. When data is copied to secondary or tertiary tiers, copies of this data can be made automatically to tape silos at remote locations. This provides an automated managed solution for disaster recovery of data combined with cost effective data storage and backup.

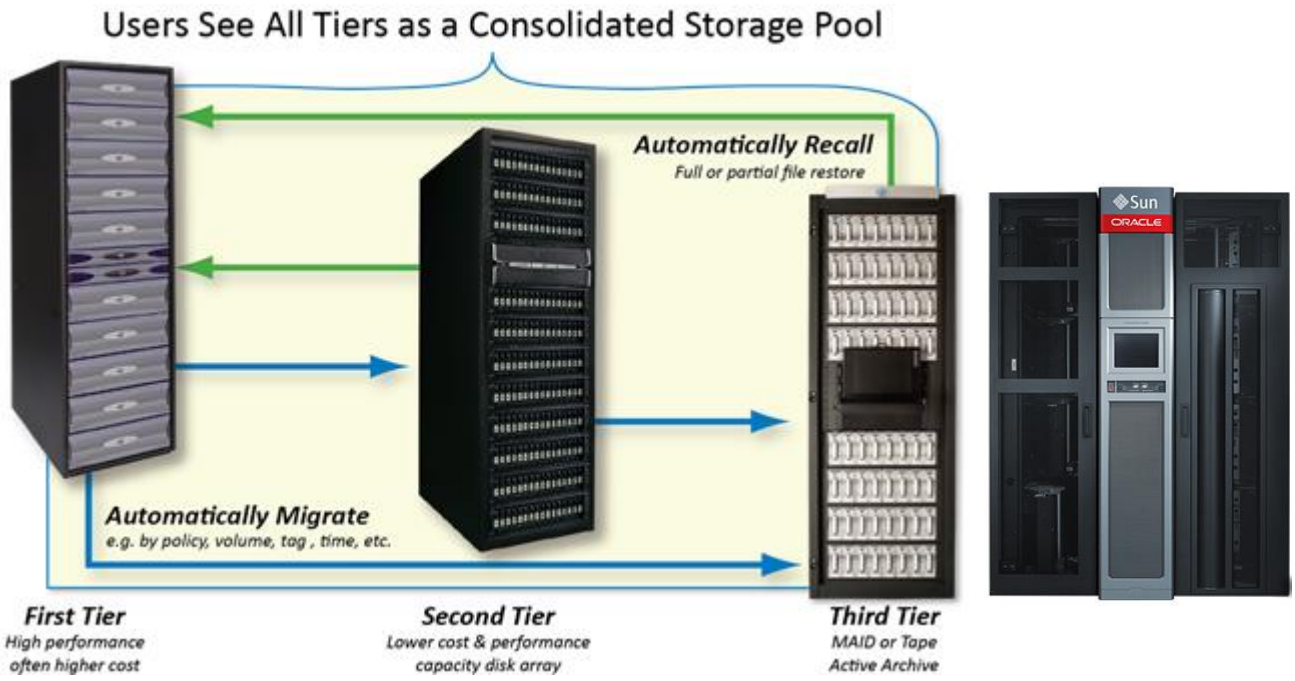


If appropriate, de-duplication or compression can be used to reduce the volume of data stored. As well as automatically migrating data between tiers according to predefined rules and policies, the data management software will regularly check tapes in the silo, merging partially used tapes to use all the available capacity.

Integrating and automating these functions reduces the complexity of managing large volumes of data, delivering efficiency gains for IT operations staff and so freeing time for more strategic tasks or innovative initiatives.

This technology also facilitates the automatic migration of data from one storage technology to the next with minimum administrative overheads. SGI have customers in Australia where data created 21 years ago has been migrated through multiple generations of technology and is available today for immediate access.

Knowledge management and digital asset management software tools that are aware of the tiered storage is regularly used in conjunction with these technologies, providing high level value add capabilities to manage, search and curate the data.

Tiered Storage Virtualisation solutions from SGI  are very cost effective for storing, managing and curating large data volumes, with savings accruing year on year from all areas including initial capital expenditure for a given storage volume, support, power and data administration costs.