

SGI InfiniteStorage

Managing the data explosion...


Floyd Christofferson
Director, Storage Products

SGI® Company Proprietary



SGI: Wrangling Big Data for > 25 Years

- 25+ years of innovation
- 500+ patents
- 6500+ customers worldwide
- 1500 employees
- Headquarters in Fremont, CA
- 2013: \$767M, strong balance sheet
- World-class servers, storage, services
- World's fastest supercomputer
- Petabyte scale
- Open Source leadership



SGI customers create and manage extreme volumes of data

Staying in Front of the Data Gap



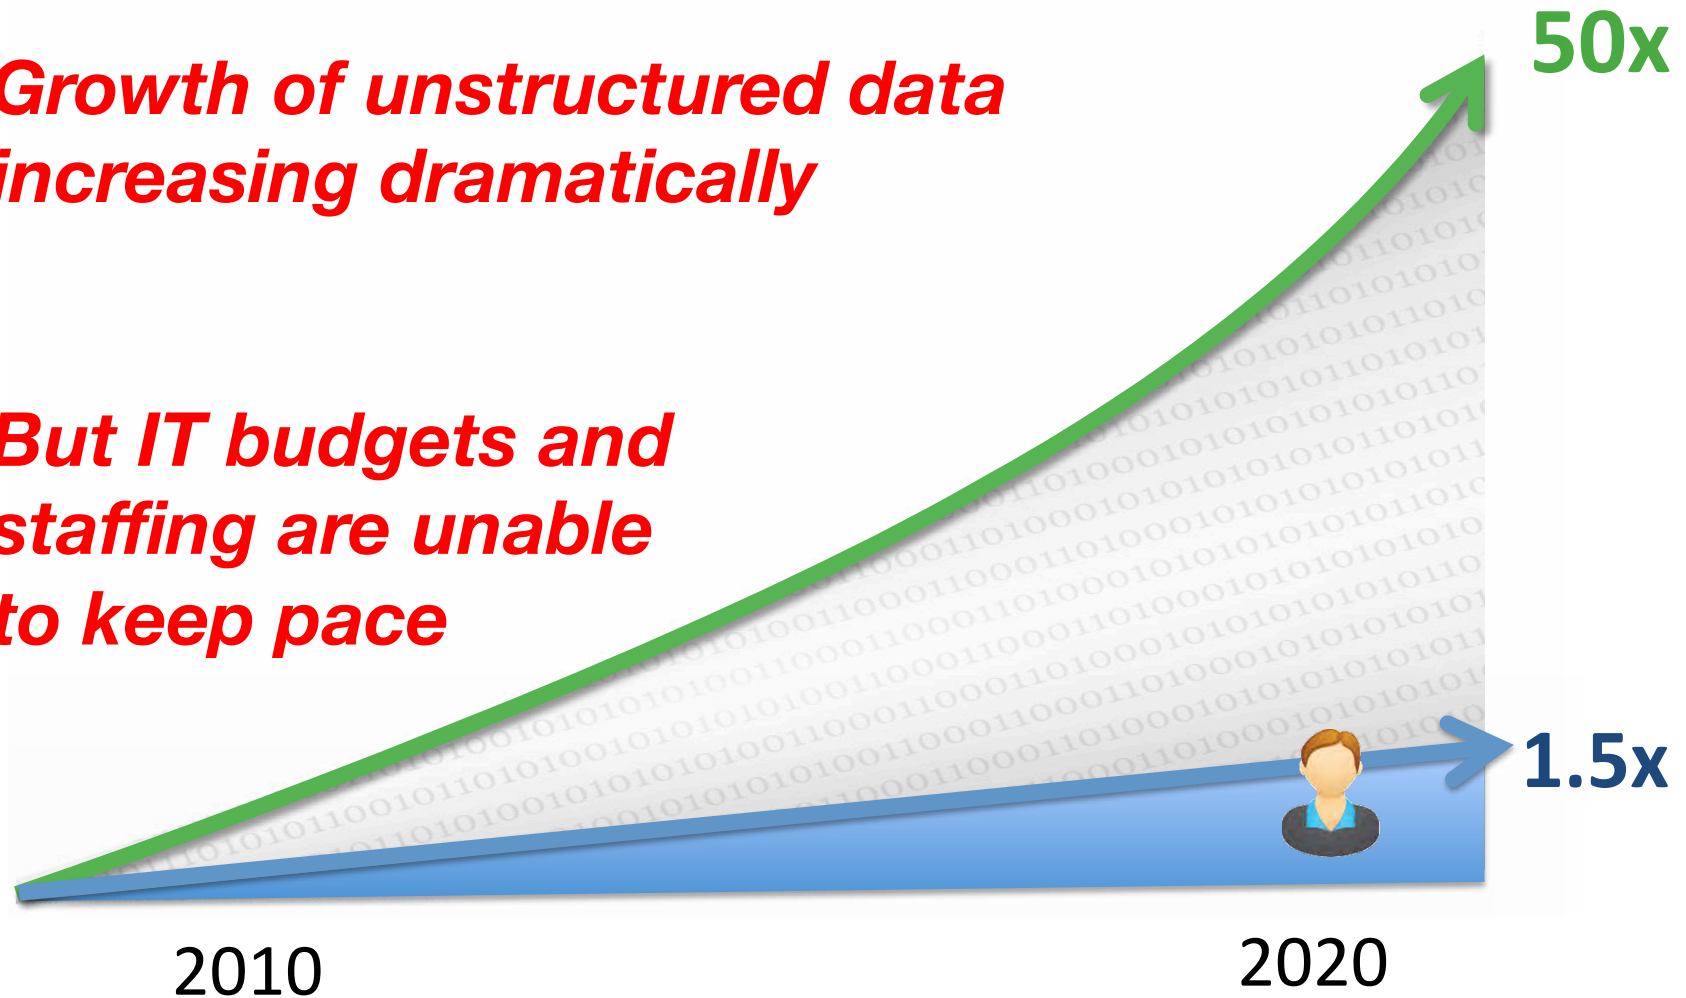
SGI's InfiniteStorage ecosystem has been driven by extreme data requirements of high-performance computing

SGI Shipped 1.2 Exabytes of Disk in 2012-13

Extreme Data Growth in All Industries

***Growth of unstructured data
increasing dramatically***

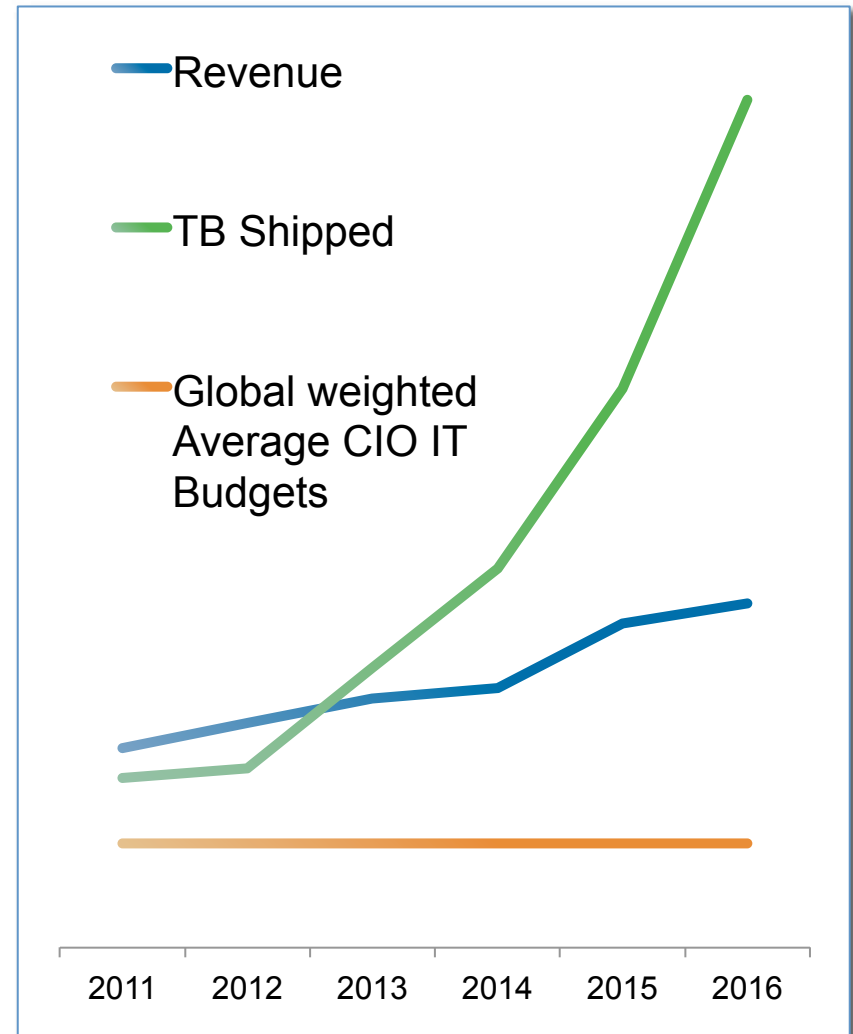
***But IT budgets and
staffing are unable
to keep pace***



Data is growing, storage budgets are not

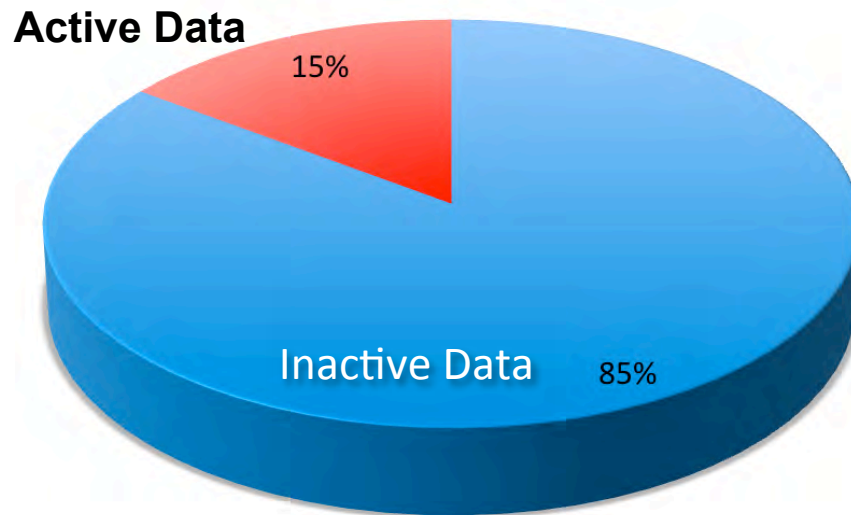
- IT Budgets anticipated to be virtually flat
- Data growing faster than:
 - IT budget
 - Recovery capabilities
 - Staff resources

Unstructured data, is growing
at 60% to 80% year over year



Source: Gartner: Who's Who in Object Storage,
Published: 10 April 2013 ID:G00248527

Facts the Many Storage Vendors Don't Discuss



Eighty-Five percent of production data is Inactive --- 68% not accessed in 90 days

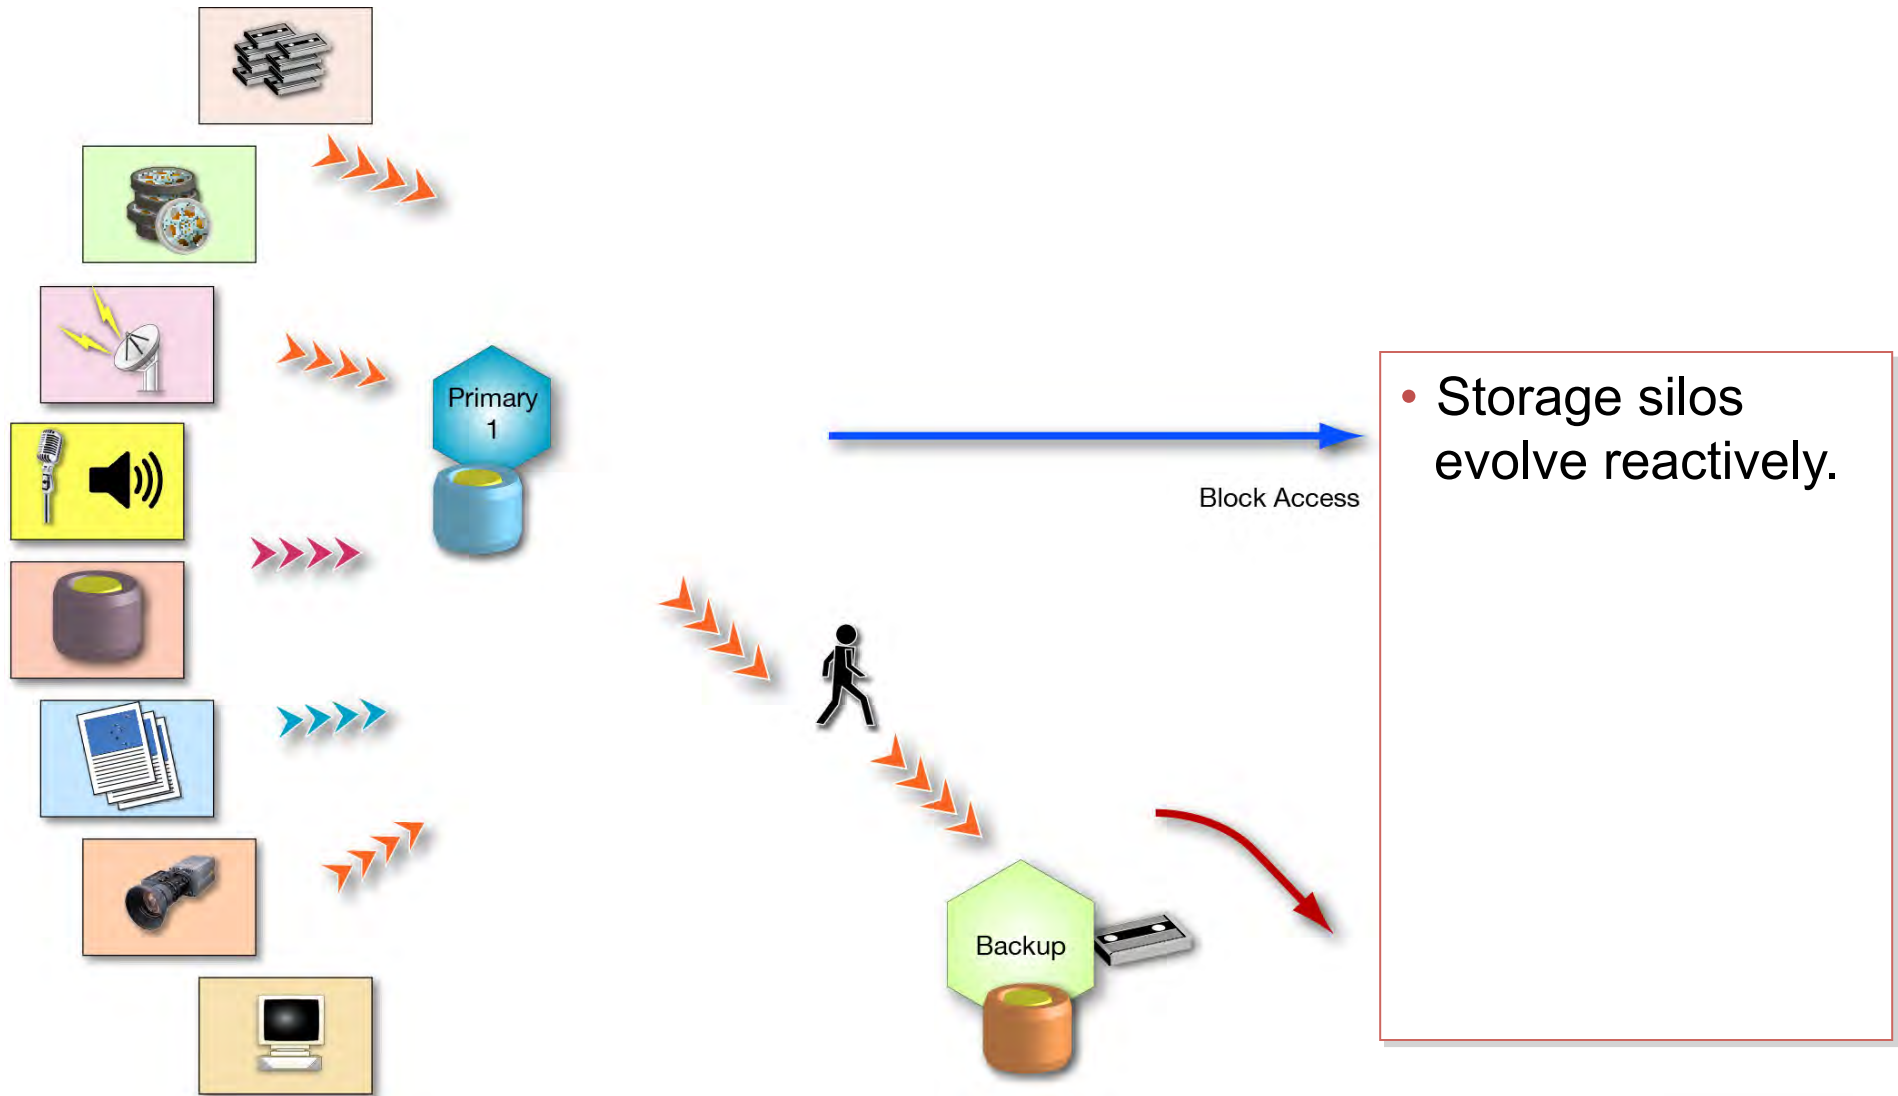
According to Forrester Research

Over 66% of files are re-opened once and 95% fewer than 5 times.

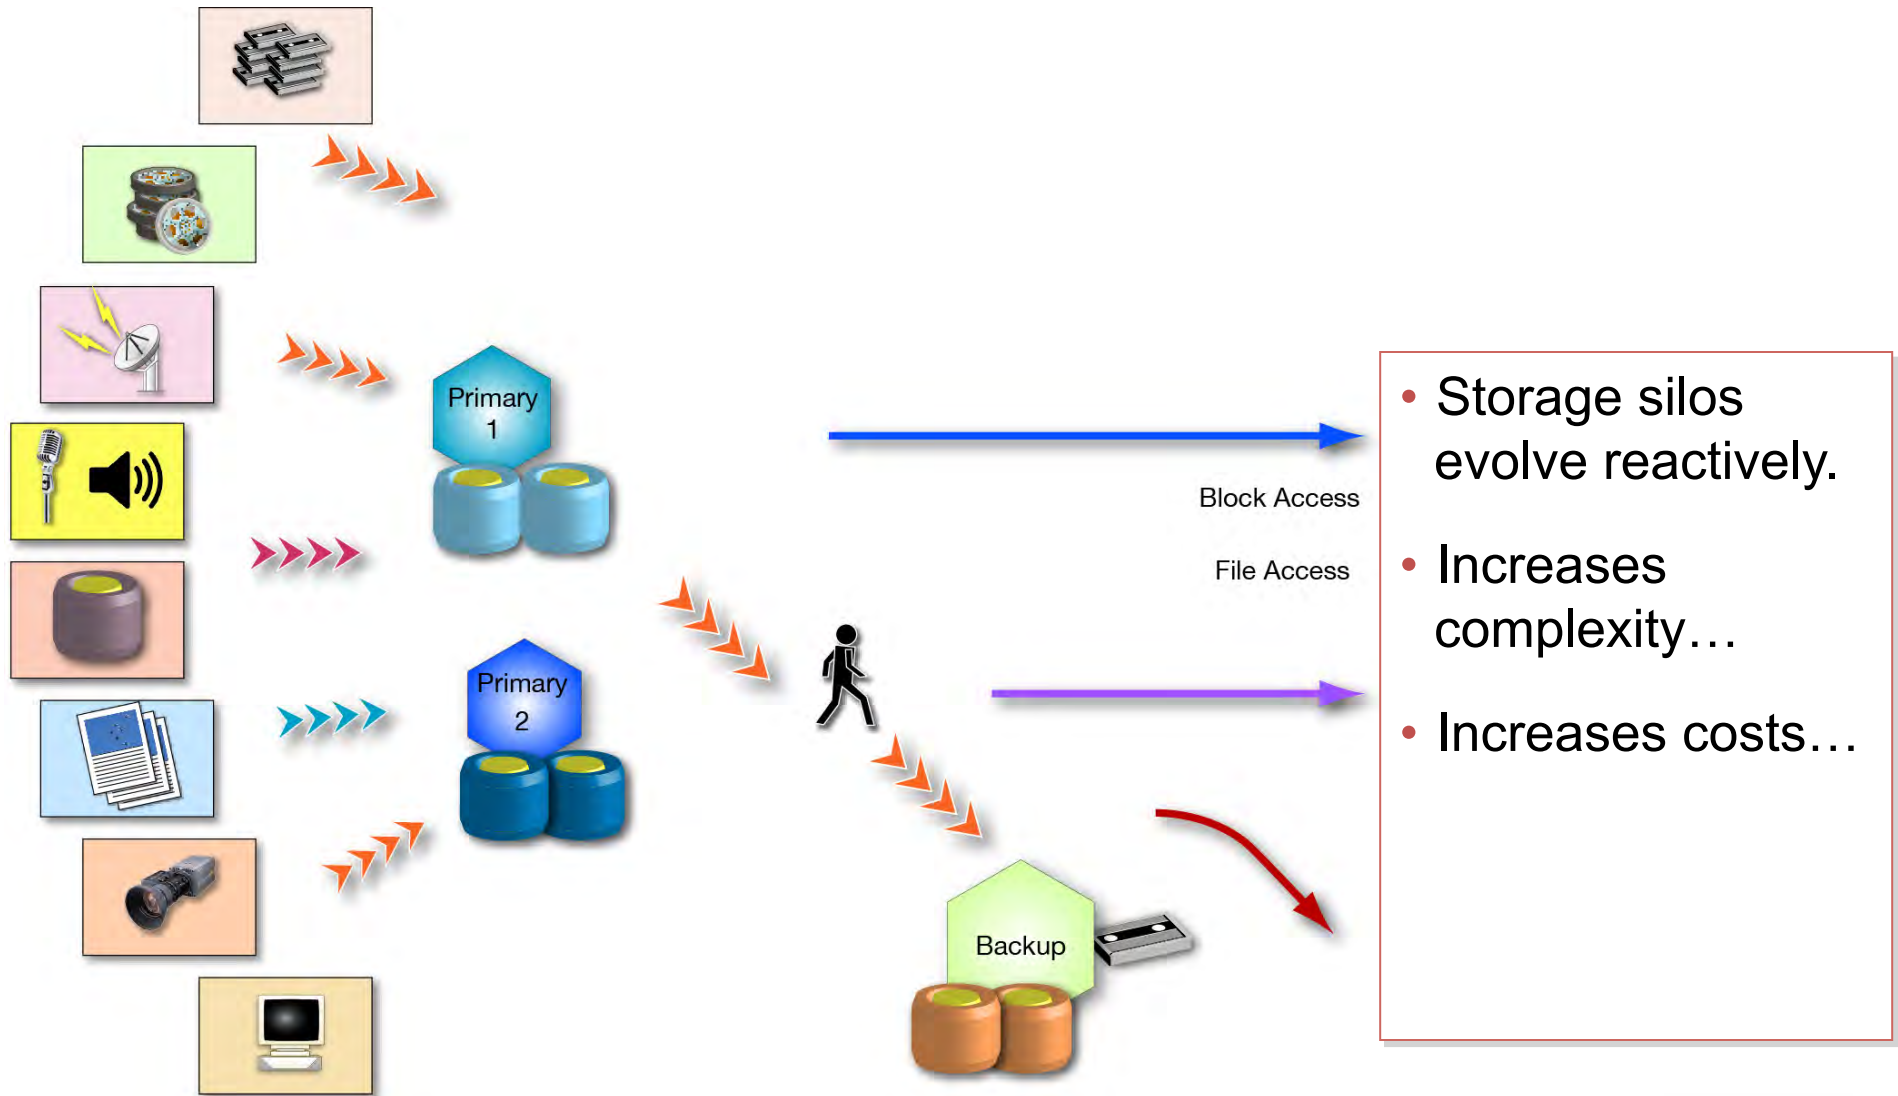
U.C. Santa Cruz Study on Data Access Patterns

Inactive Data Doesn't Need Expensive Primary Storage

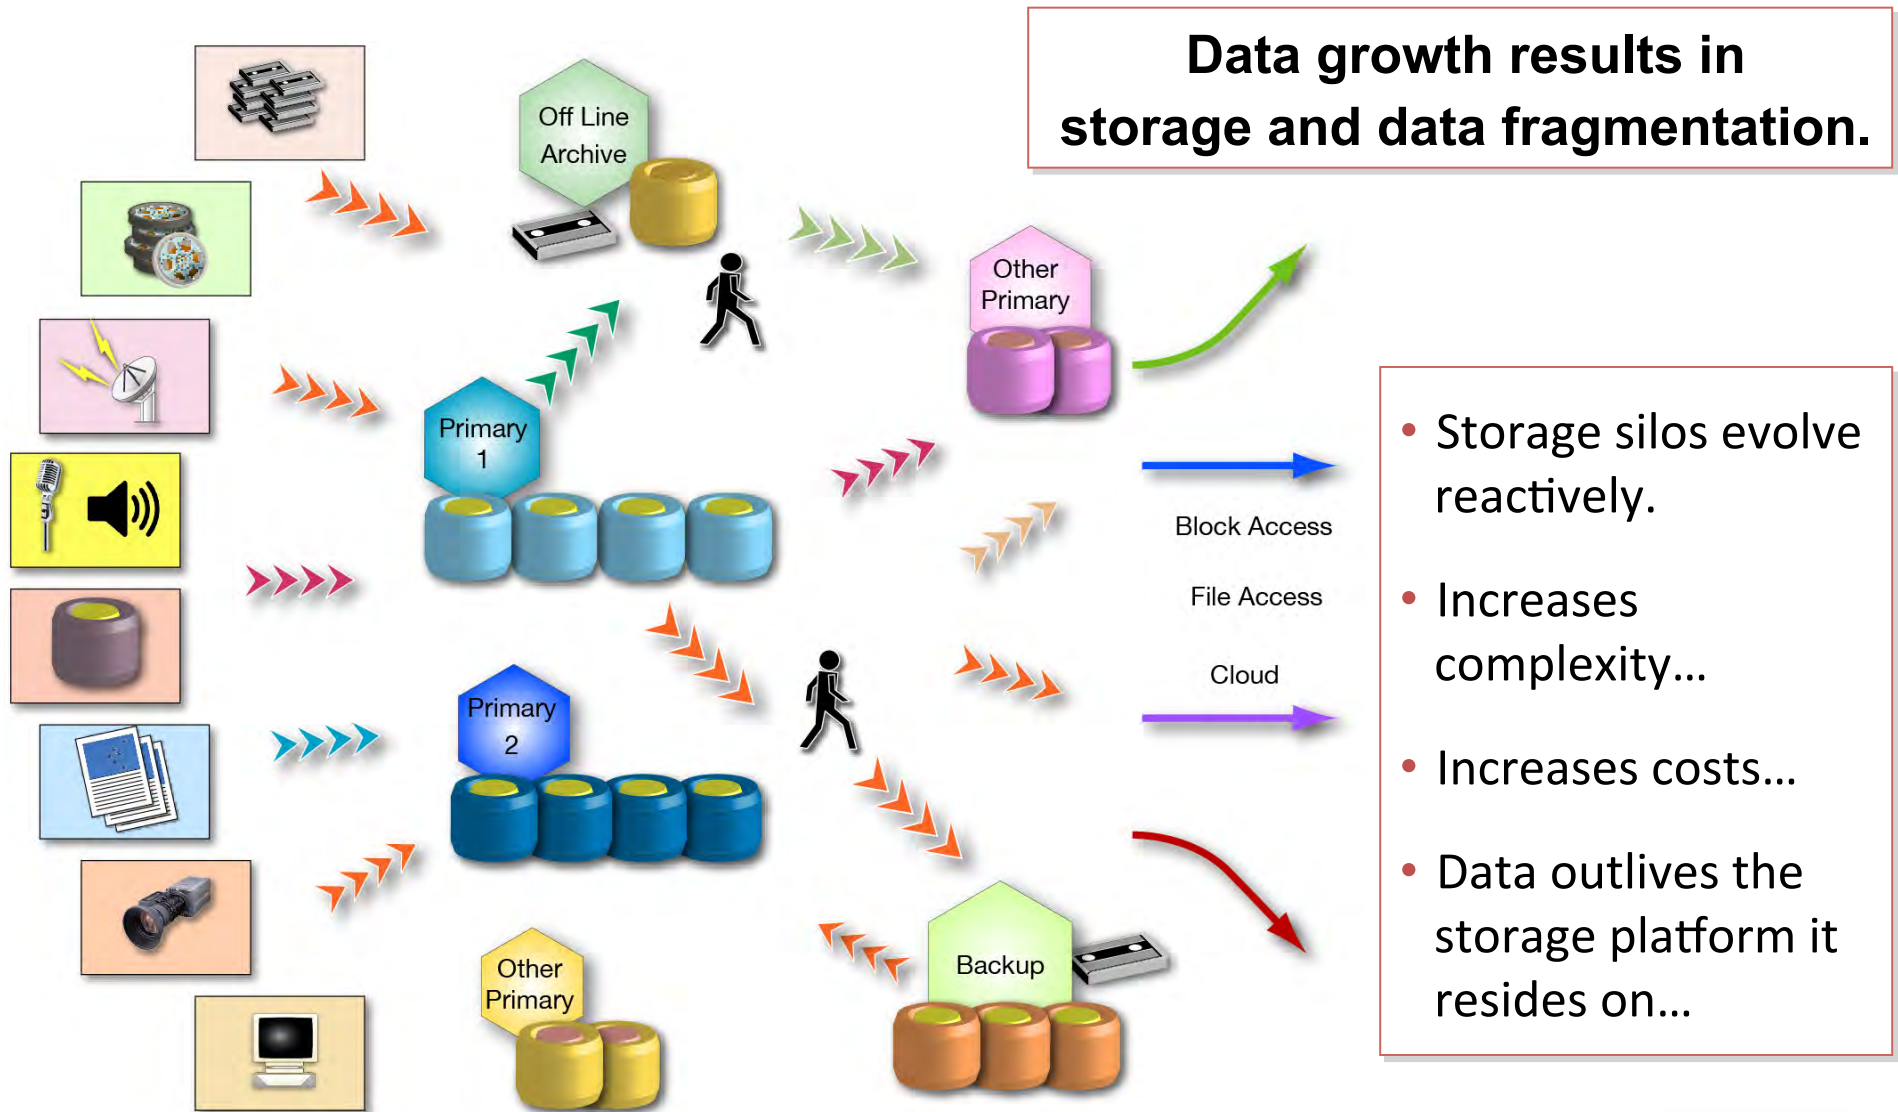
Storage Growth is Relentless...



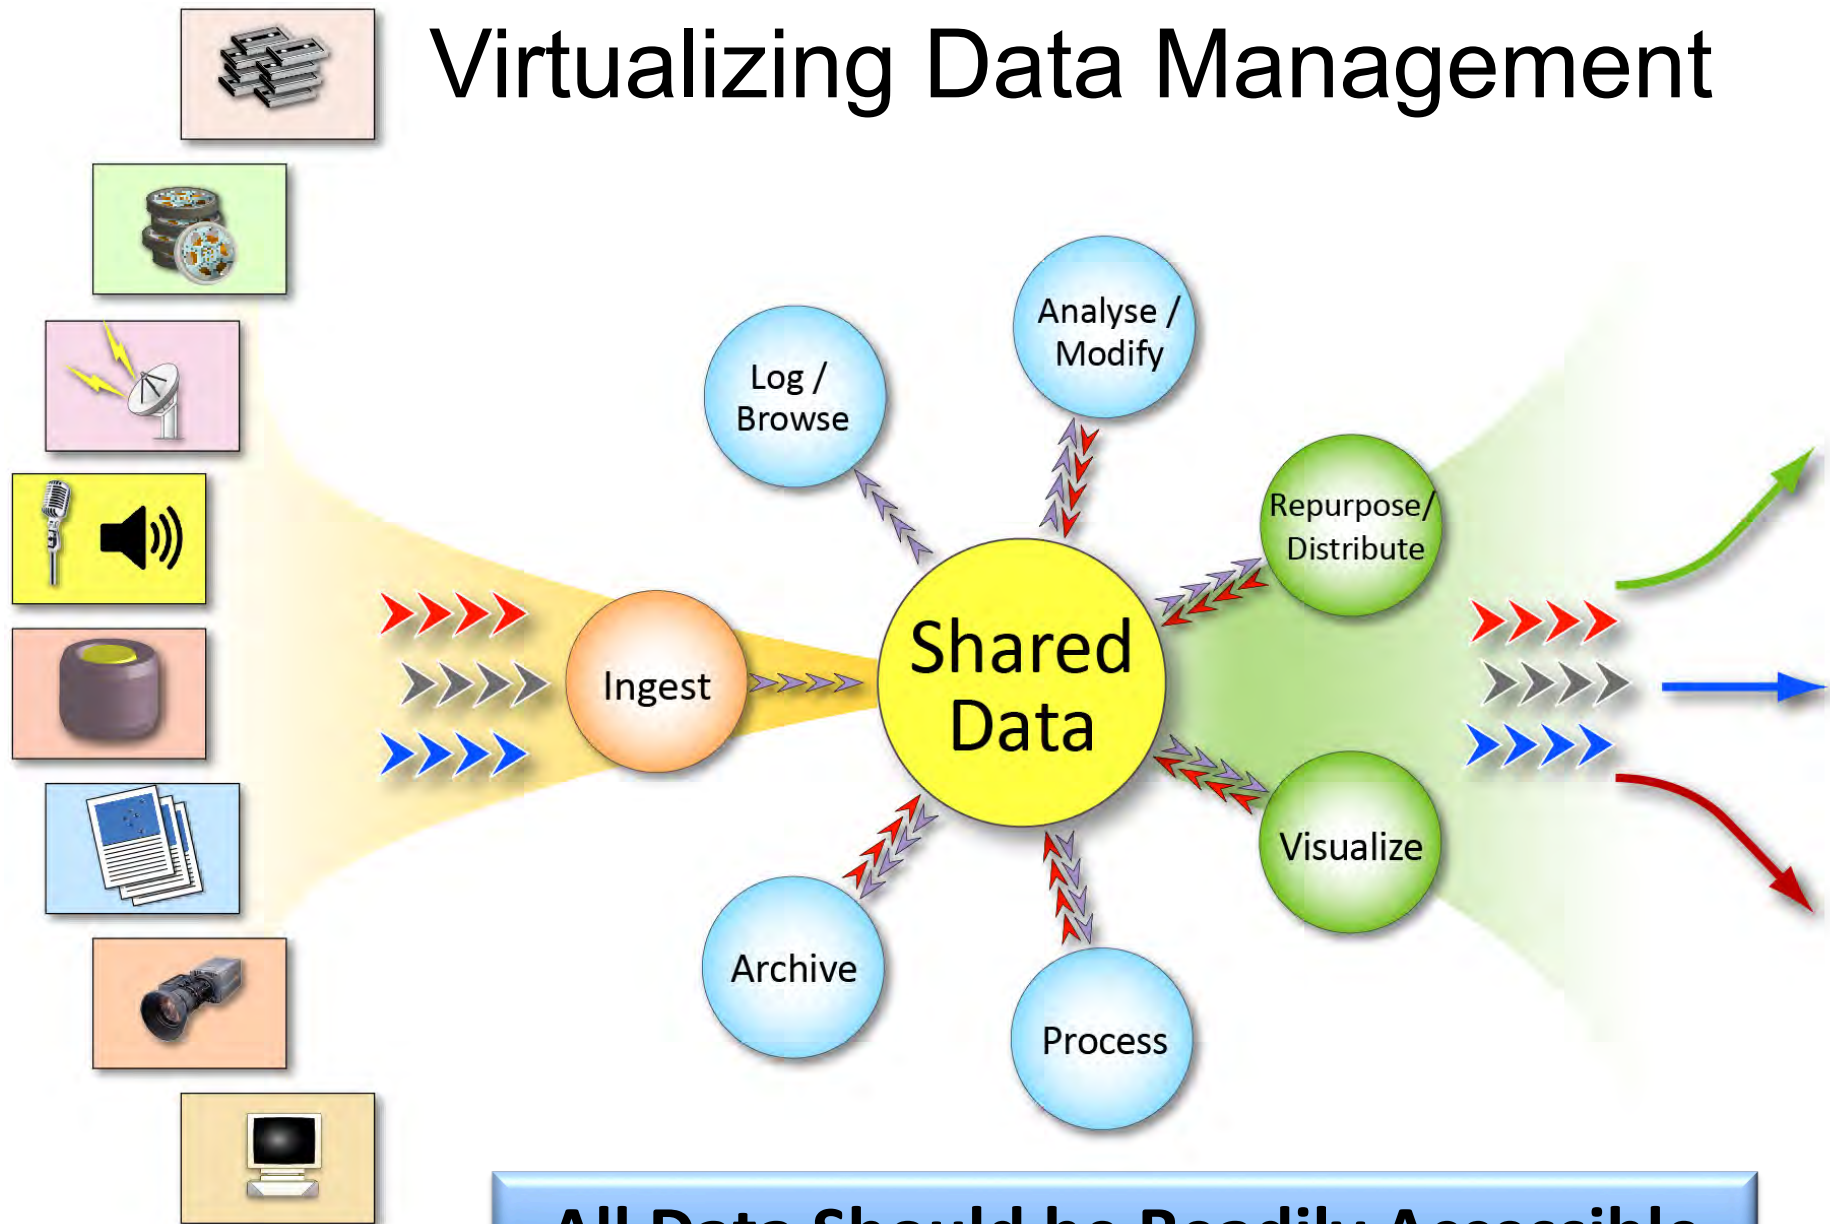
Management Headaches Grow...



Storage Sprawl = Data Problems



Virtualizing Data Management

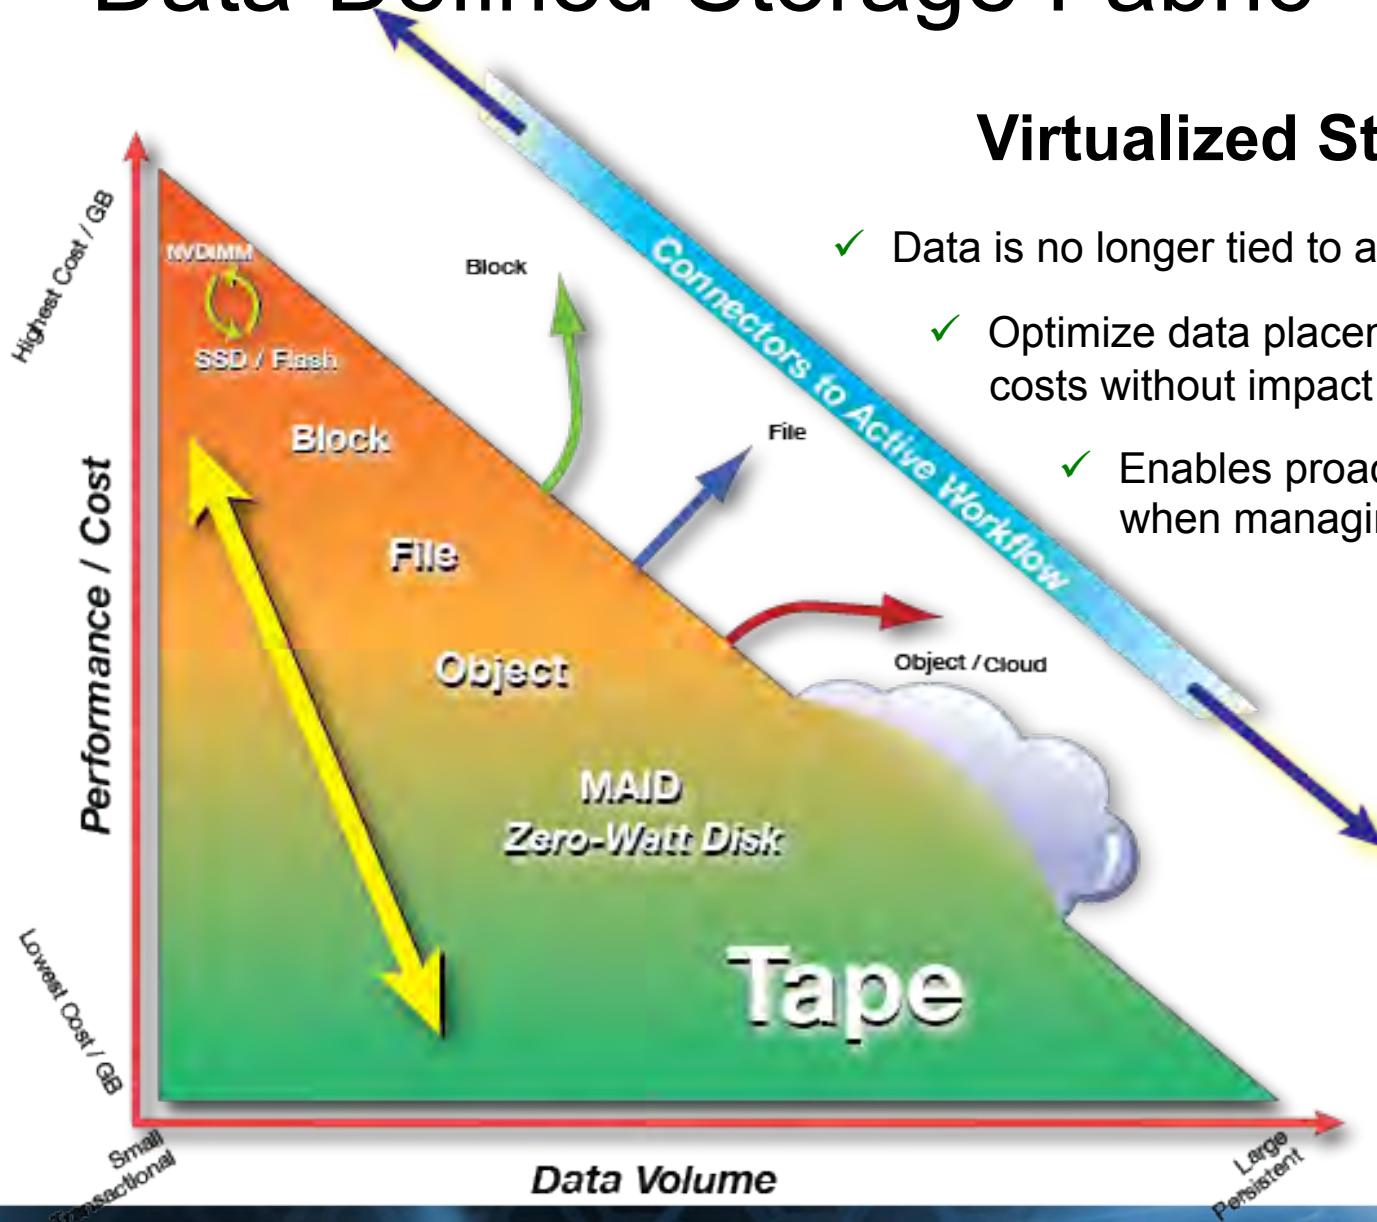


All Data Should be Readily Accessible

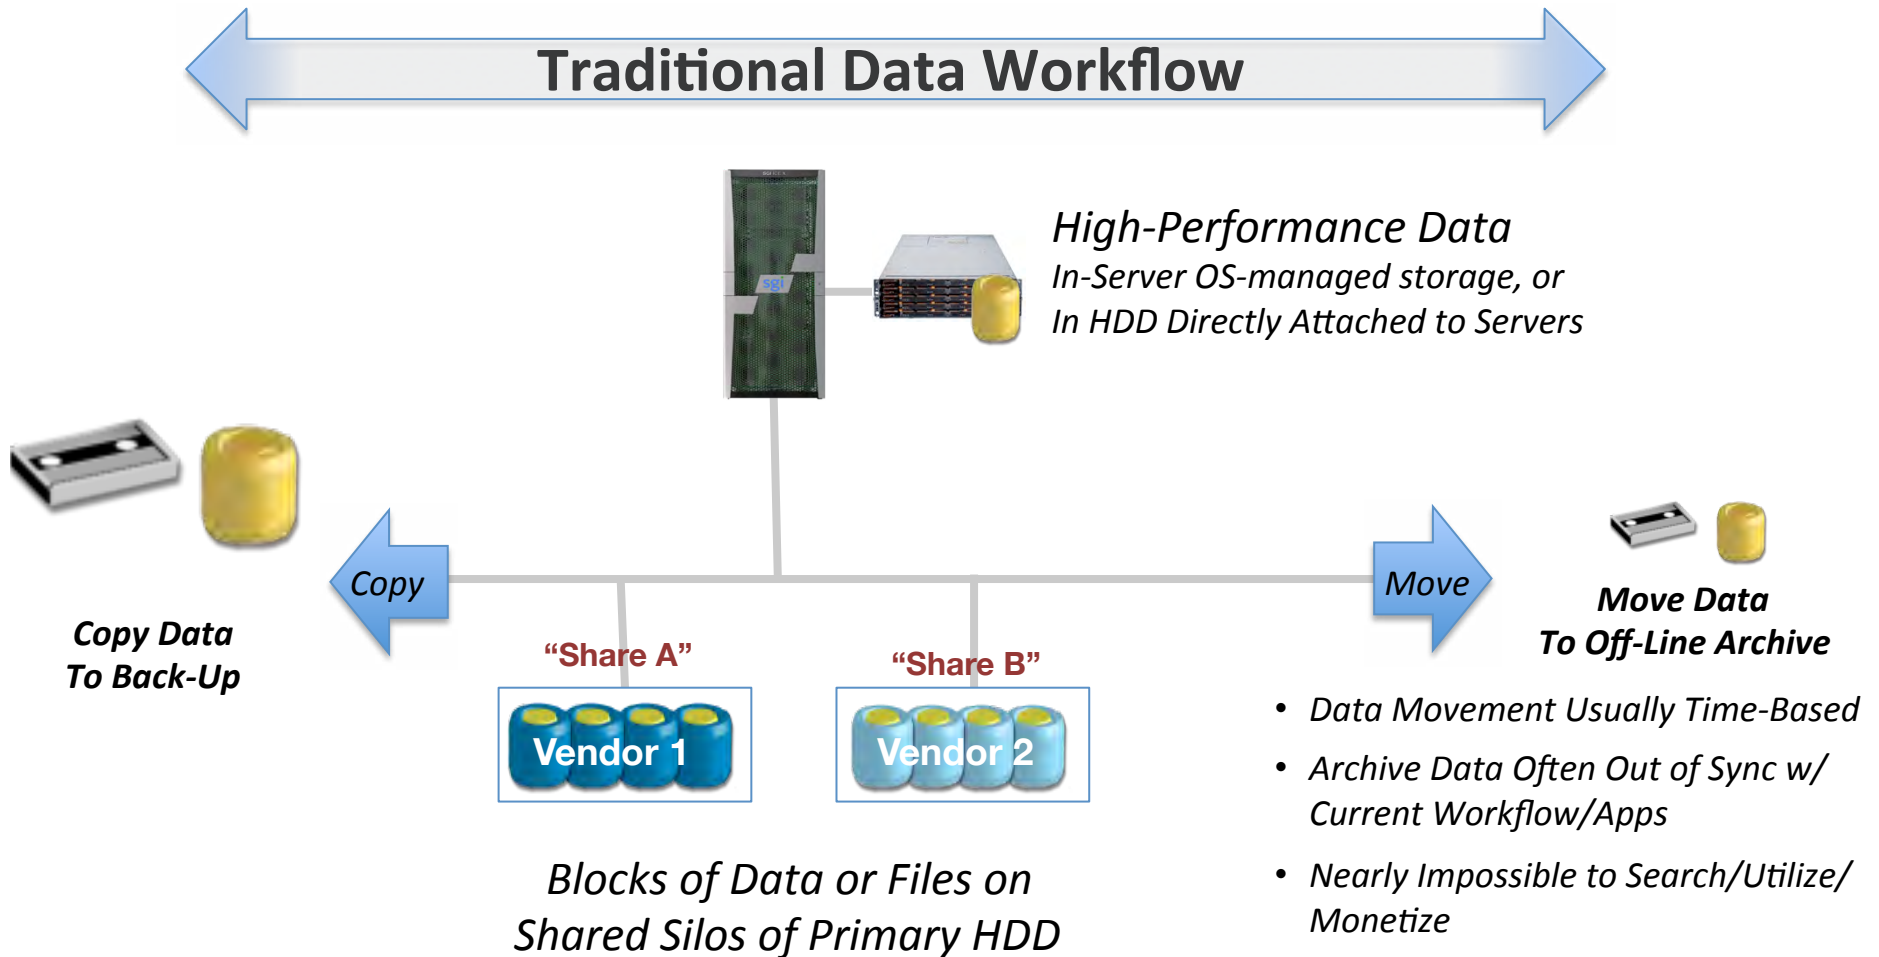
Data-Defined Storage Fabric

Virtualized Storage Tiers:

- ✓ Data is no longer tied to a particular device or tech...
- ✓ Optimize data placement & reduce operational costs without impact to users...
- ✓ Enables proactive IT, not reactive – when managing data growth...

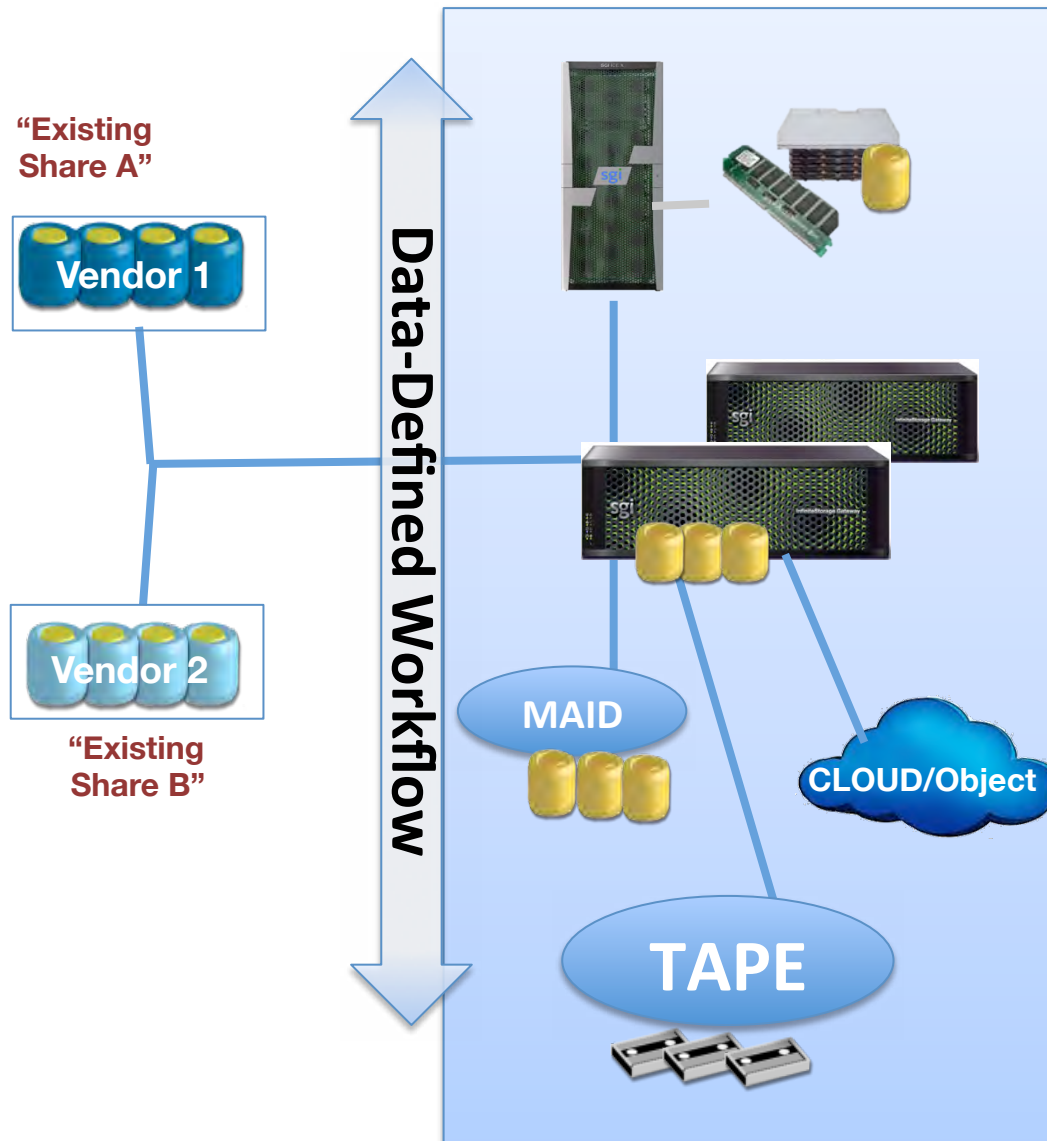


Traditional Data Management Model



Does NOT distinguish between Hot, Warm, Cool or Cold Data

Data-Defined Storage Fabric



Extreme Data – Server-Tethered, Clustered

- On-blade and in-cluster managed caching
- DAS / High-Performance / Hybrid

Hot Data – Shared, Online

- Hi Performance / Utilization
- Business transactions / performance-sensitive applications
- Traditional "Primary" Block & File / Hybrid Arrays
- Full RAS

Warm, Cool, Cold Data – Shared, Active

- Data not for immediate transactions, but need access for monetization / management
- Capacity-sensitive
- Active Archive Fabric – data remains indexed, visible and searchable
- Virtualized intelligently-managed tiers
- Object/Cloud, 2d HDD, MAID, Tape

SGI addressing data explosion for over 20 years

- Cosmological Research - 100.0 PB
- Aerospace agency (40 GB/sec) (21 years online) - 60.0 PB
- US National Weather Analysis (300 TB/day I/O – 105GB/s NFS throughput) - 55.0 PB
- European Weather Analysis - 30.0 PB
- Movie visual effects – 1.8 Billion files) - 24.0 PB
- Aerospace research (21 years online) - 20.0 PB
- University, Australia – Genomics Research - 20.0 PB
- Sports League Active Archive (~40TB/day ingest) - 18.0 PB
- Central Europe – Media library - 15.0 PB
- Scientific Research (21 yrs in prod, always online) - 8.0 PB
- Global Media and Entertainment Film Library - 7.0 PB
- Weather prediction (+13TB/day) – With Lustre - 6.5 PB
- European National Institute for Audio & Video - 4.5 PB
- National Computing and Network Services (Europe) - 4.0 PB
- European National Research Agency - 4.0 PB
- European Space Research - 4.0 PB
- European Oil and Gas Customer - 3.0 PB
- Multinational Oil and Gas Customer - 2.7 PB
- Aircraft Manufacturer - US - 2.0 PB
- Earth Data Research, Europe - 1.7 PB
- Cancer Research center, UK - 1.6 PB

SGI is bringing the experience of managing some of the world's largest HPC data environments to the growing needs of today's GPC market.

Customer Perspectives

"We have **enough primary** storage. Our data is growing, but we **don't need more primary** to store it. We need the ability to **access** it at **any time** in an **online** state, but this should **not** be on **expensive primary** disk."

Kathy Gruzas, CIO of Weta Digital

sgi