



SGI® Complex Data Management Strategy and Solutions



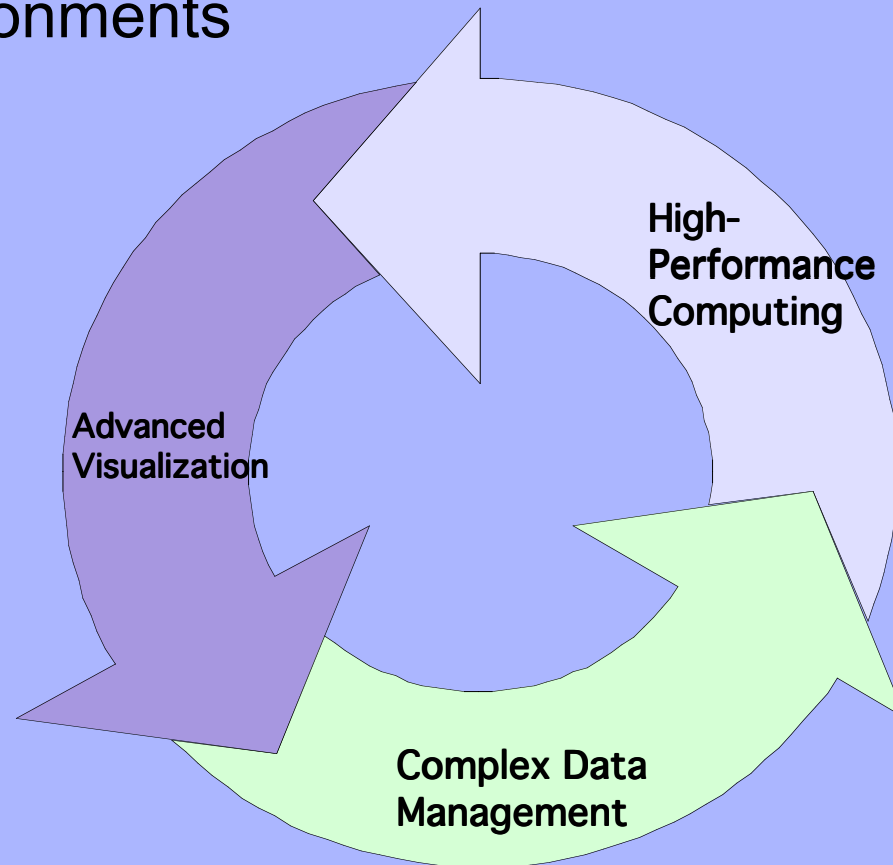
LaNet Merrill
SGI SAN Product Manager

Topics

- Overview of SGI® Complex Data Management
- Architectures and Solutions
- SGI Solutions in Action

SGI® Delivers Solutions with Unmatched Capabilities and Power

SGI is focused on the demanding requirements of technical and creative professionals and extending our technology into heterogeneous environments



Data Explosion - The Problem

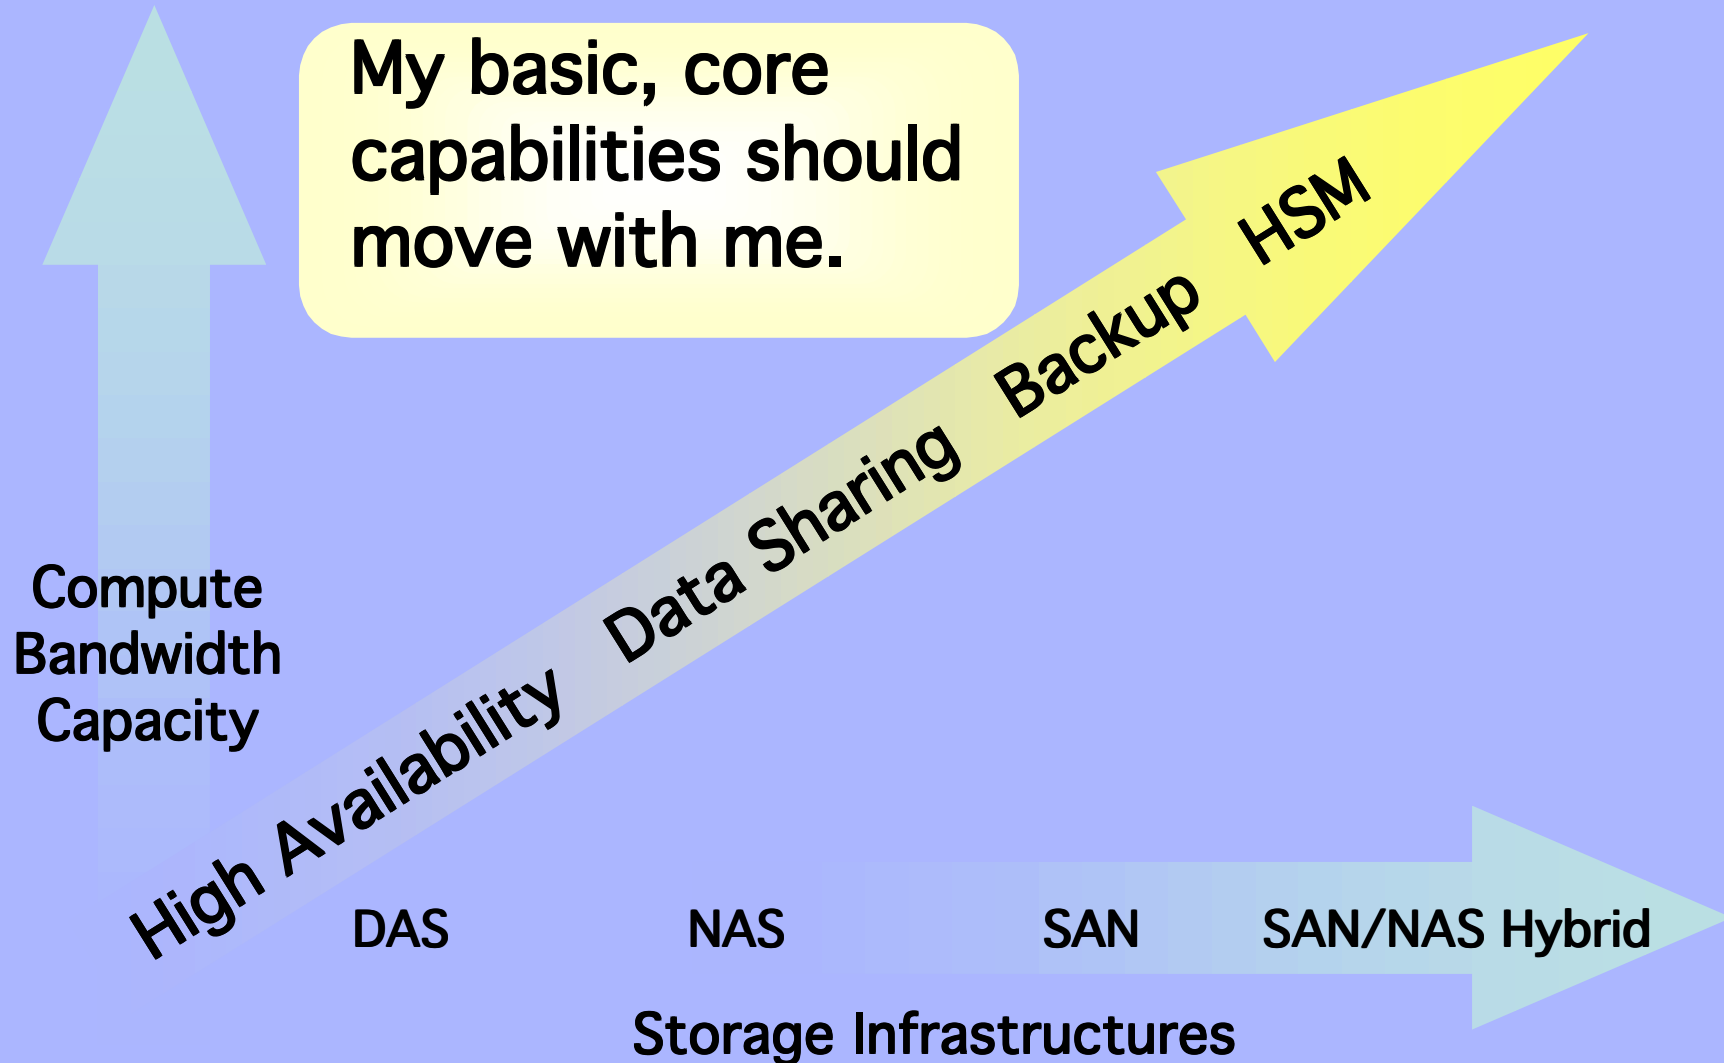
'As my requirements change, I may have to move between DAS, NAS and SANs.

I don't want to get locked in -

Dead-ends, box upgrades, broken applications and lots of data format changes are NOT acceptable'

capacity change'

Problem Statement: The Data Explosion— What WOULD Solve the Problem?



SGI® Complex Data Management Strategy

High-Performance, Integrated Storage Solutions

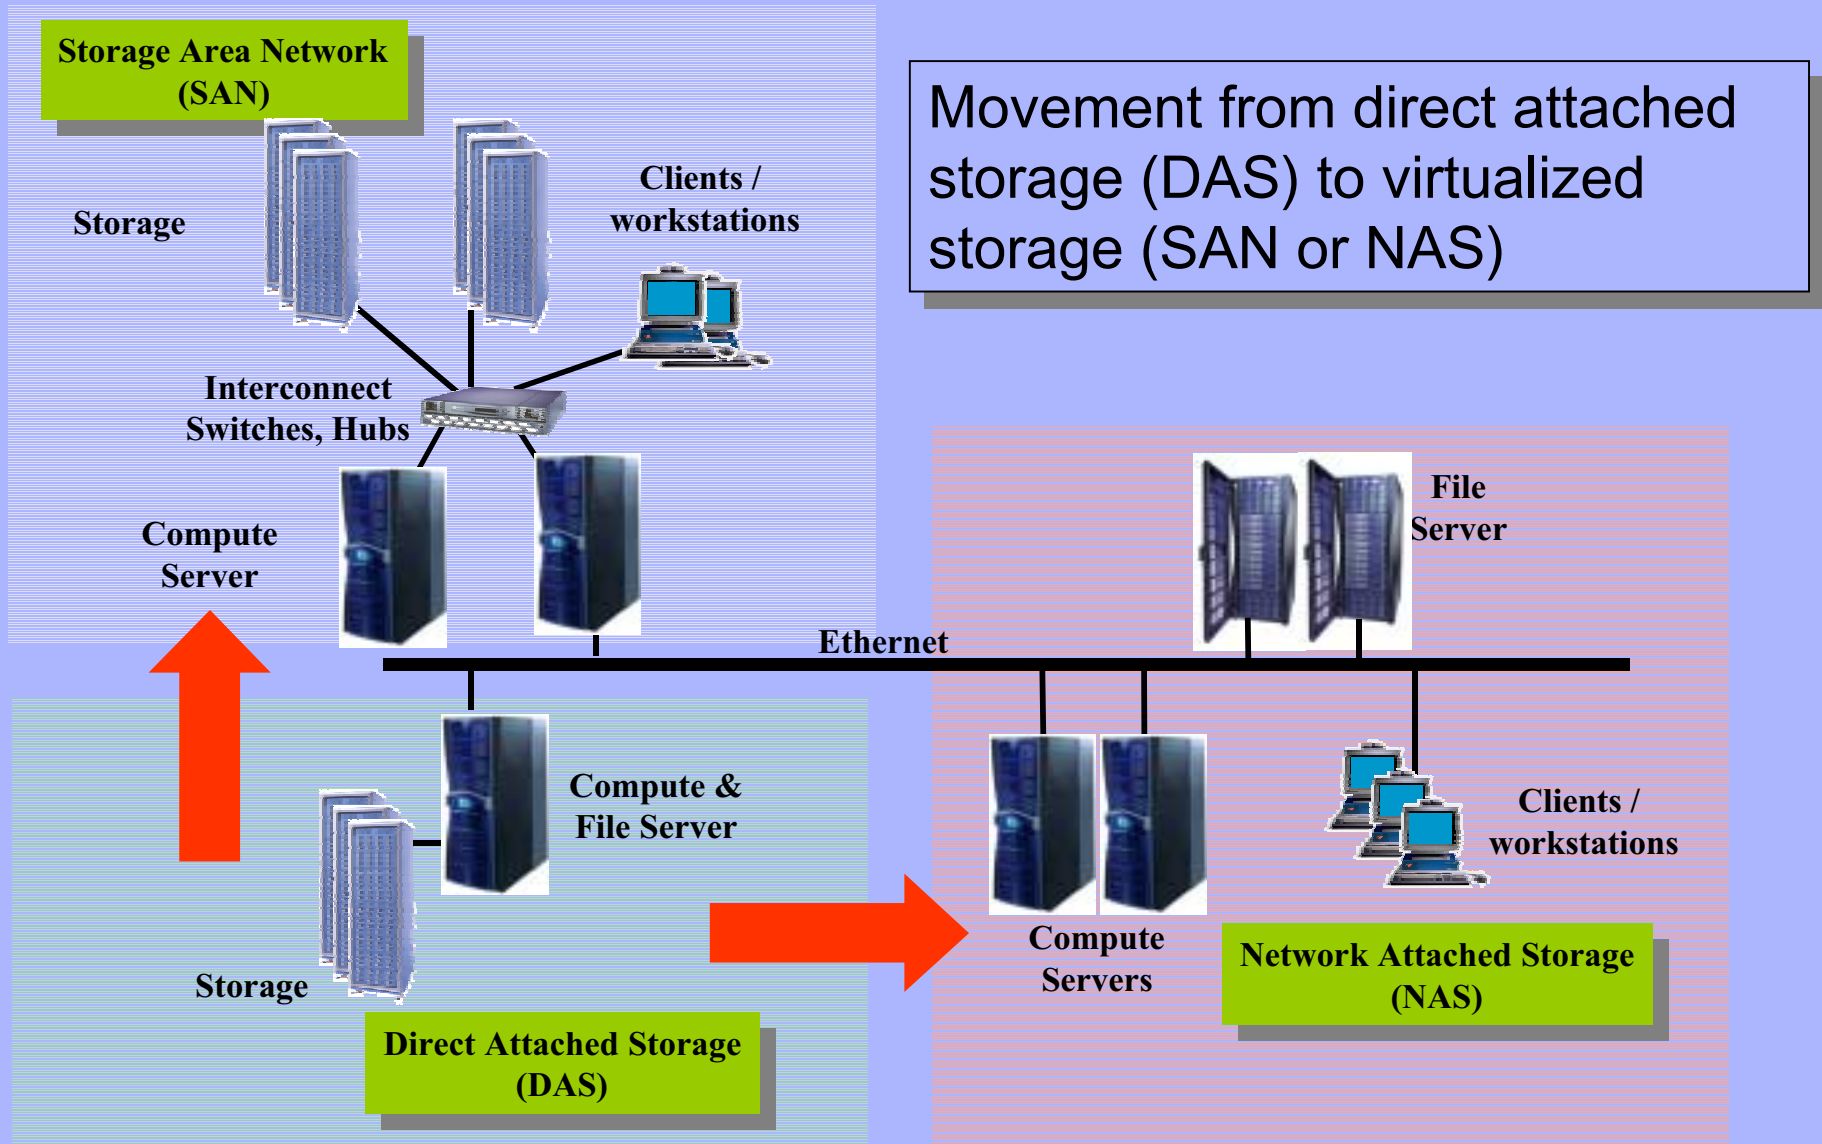
Performance: Integrating the bandwidth and data needs of servers, visualization and storage providing ultimate bandwidth and capacity.

Modular: Scalable software and hardware components which may be deployed separately or together.

Reduced Complexity: Information where you need it when you need it, allowing you to focus on your core business.

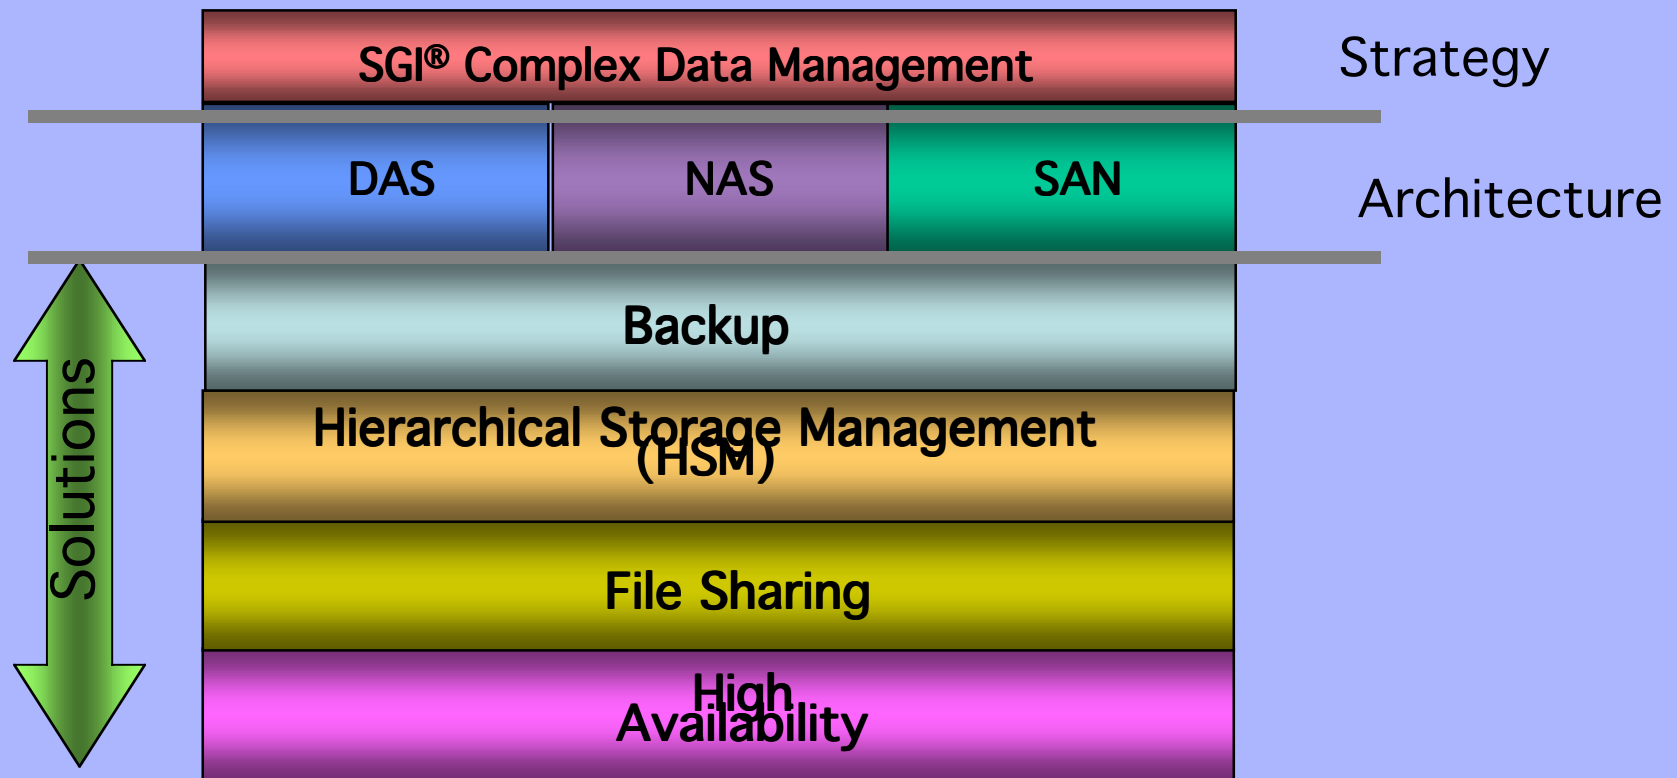


Data Management Architectures

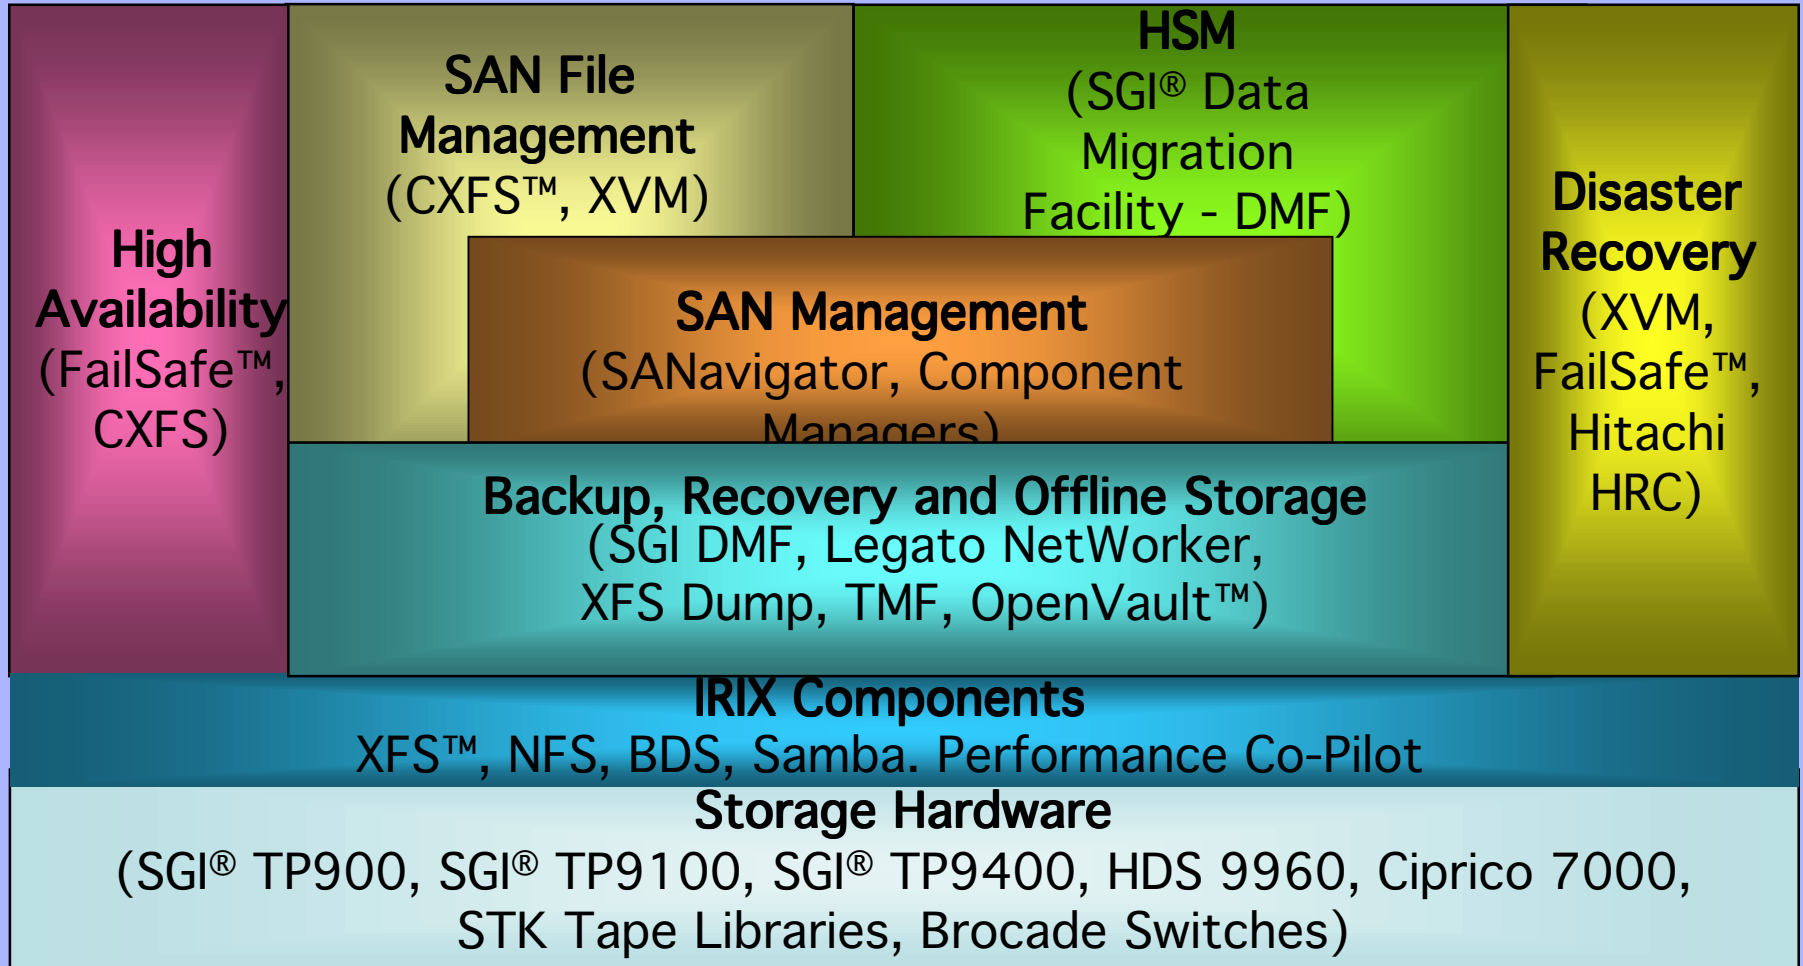


Complex Data Management Solutions

- Mix and match to build unique custom solutions
 - Between architectures (DAS, NAS and SAN)
 - Within architectures



Complex Data Management Modular Technologies



The Value of SGI® Data Management

- Storage solutions that work
 - **Fibre Channel solutions longer than anyone in industry (1997)**
- One stop shopping and one stop support
 - **Layered products for backup, HSM, file sharing and high availability**
 - First HPC Shared File system - CXFS™ (Dec 1999)
 - Pioneered OpenVault™ (1997)
 - Clustered HSM file system - DMF (1999)
- Focus on performance-sensitive solutions
 - **Broke 3 Gbyte/sec SAN barrier (aggregate BW) Jan 2000**
 - **First 2 GB SAN Fabric (October 2001)**
 - **Delivering first 12 GB/sec SAN Solution (aggregate, 15 GB/sec peak)**
- Customers trust SGI with their intellectual property (IP)
 - **Bringing the three major components of computing (serving, visualization and data management) together in one consolidated, complete package**

SGI[®] Data Management Architectures and Solutions

SGI® Data Management Direct Attach Solutions


- SGI DAS Solutions Value:
 - Maximum I/O Bandwidth
 - Unlimited file size and number of files
 - Highly reliable “system”

<h2>Backup</h2> <ul style="list-style-type: none">• SGI DMF• Legato NetWorker• XFS™ Dump• STK Tape Libraries	<h2>HSM</h2> <ul style="list-style-type: none">• Data Migration Facility (DMF)• TMF• OpenVault™• STK Tape Libraries
<h2>File Sharing</h2> <ul style="list-style-type: none">• XFS™• NFS• BDS• Samba• TP900, TP9100, TP9400, HDS 9960, Ciprico 7000	<h2>High Availability</h2> <ul style="list-style-type: none">• FailSafe™• Redundant paths via hardware

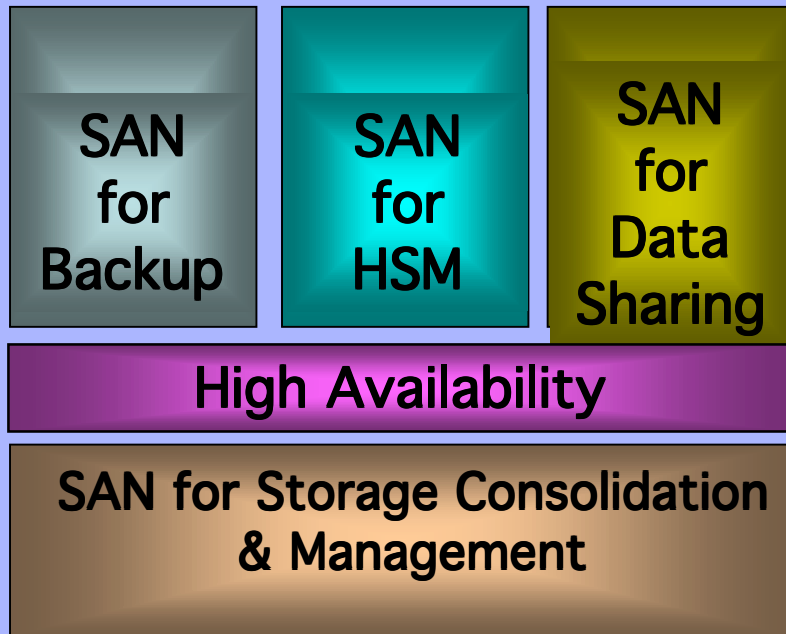
SGI[®] File Server

Network Attach Solutions

- Scalable network-attached storage solutions
- Two configurations
 - Model 830—SCSI JBOD, S/W RAID
 - Model 850—Fibre Channel RAID5
- UNIX[®] clients with NFS; Windows NT[®] clients w/CIFS
- Leverages SGI NUMA for high bandwidth
- Storage management options
 - HSM
 - Incremental I/O bandwidth
 - Application failover (HA)
 - Backup and archive
- Upgrade SAN support



SGI[®] Custom SAN Solutions




- ***SGI SAN Solutions address:***
 - Right-sizing of storage assets
 - Storage administration (complexity & cost)
 - Ineffective or complex backup and recovery
 - Improved file access and workflow

The Key to SGI™ SAN Solutions


- **Same SW elements used in DAS, NAS and SAN solutions**
 - Your investment is protected as you modify architectures.
 - Stability from widely used components
 - The components are designed to work together or separately
- **Highest performance data-sharing capabilities available, via CXFS™**
 - Focus on innovating, not spending time waiting for data
- **DMF is the most reliable, innovative, highest capacity HSM in the industry**
 - Your data is safe
 - Limitless growth of near-line capacity
 - Clustered HSM file system (1999) provides visibility for all files - even those in your HSM
- **SAN Performance Innovation**
 - Broke 3GB/sec SAN barrier (aggregate BW) (Jan. 2000)
 - First 2 Gbit SAN Fabric (October 2001)
 - Delivering first 12GB/sec SAN Fabric (aggregate, 15GB/sec peak) (November 2001)

SAN for Consolidation & Management

Current Environment

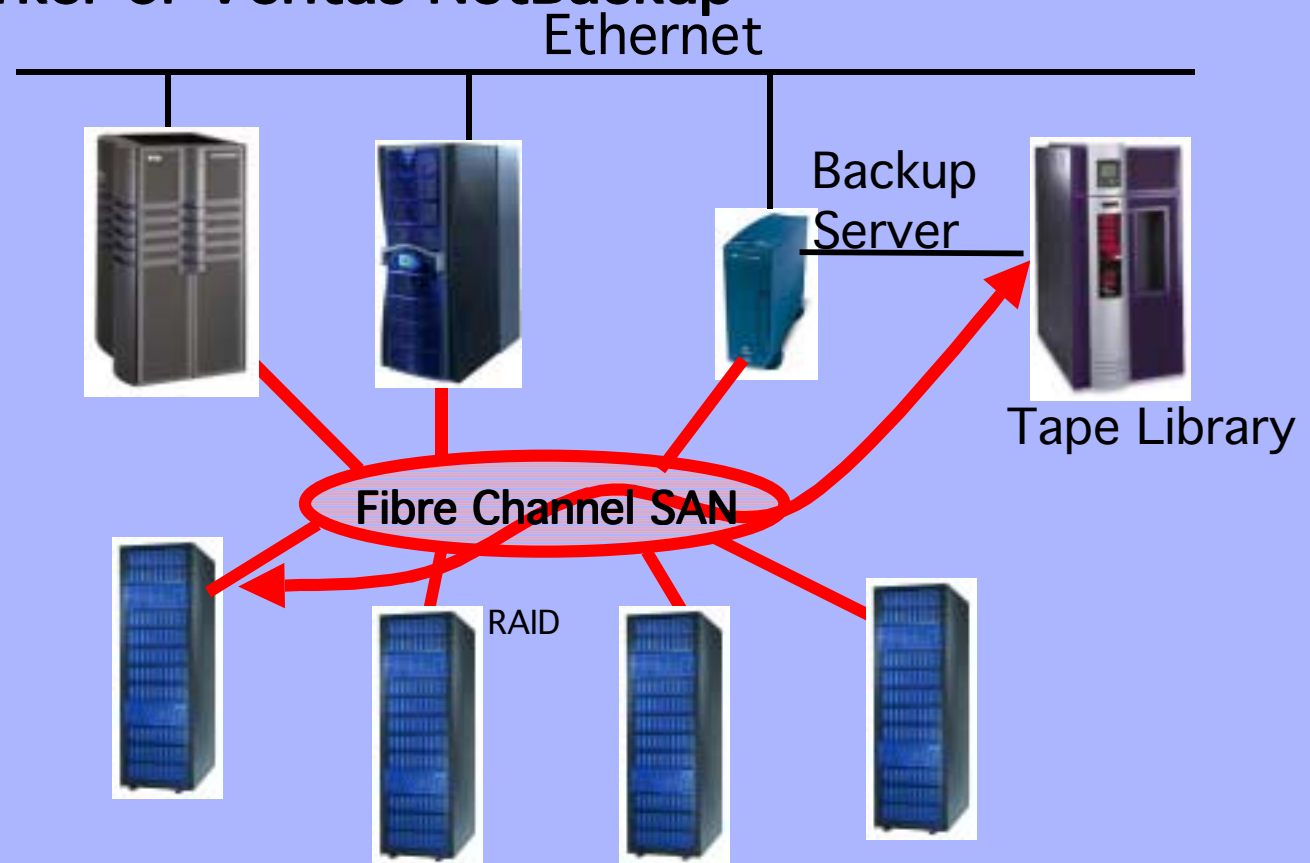


Consolidated Storage



SAN for Backup, Archive and Recovery


- Offloads backup load on the LAN
- Reduce backup window by using SAN
- Fast reliable data recovery
- Legato NetWorker or Veritas NetBackup




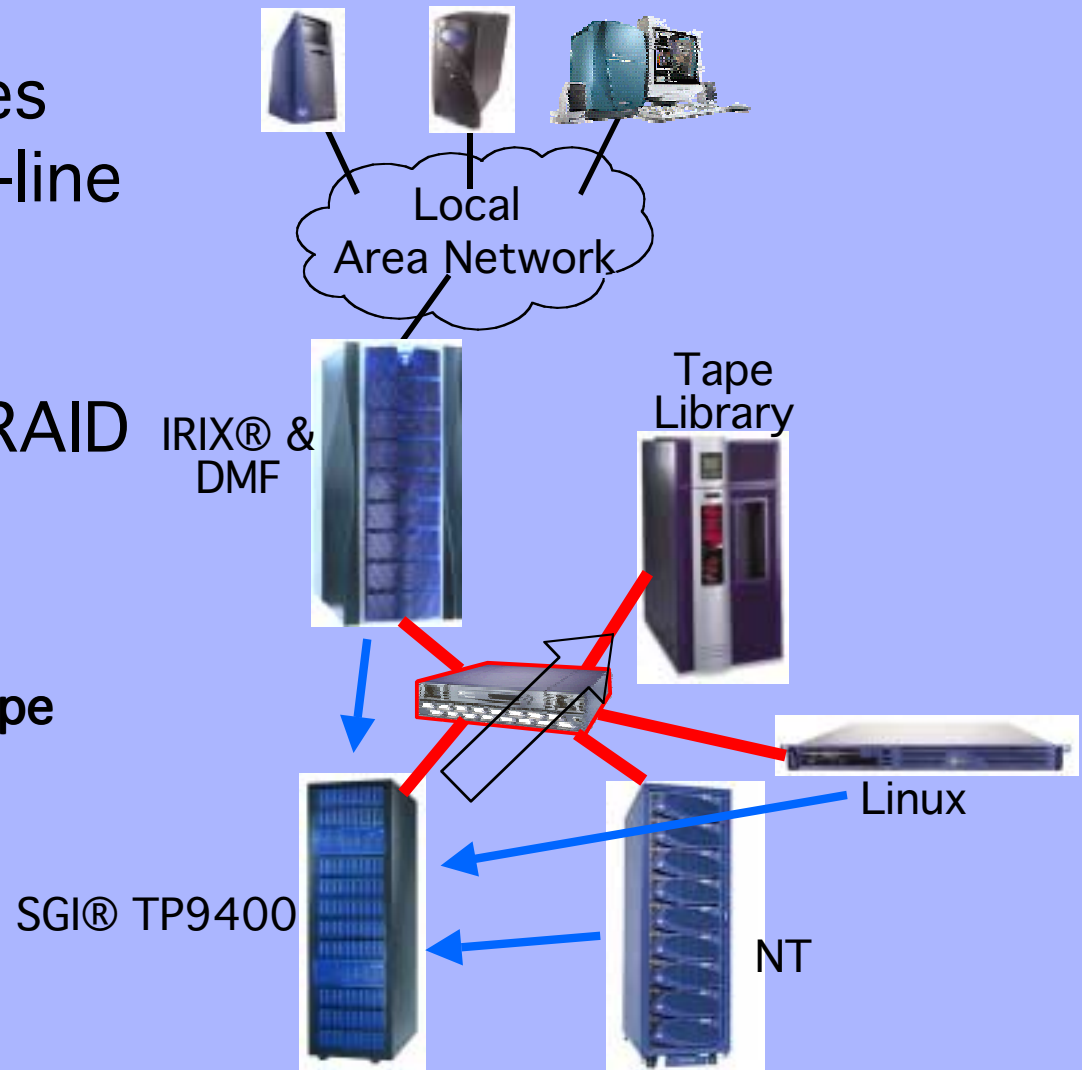
SAN for Hierarchical Storage Management

SGI® Data Migration Facility (DMF)

- Automatically migrates data off RAID to near-line tape library
- Better performance
- Higher utilization of RAID
- Data integrity

 = Data moves to RAID

 = Data moves off to Tape



SAN for Data Sharing


CXFS™ Shared File system

No Shared Data




- Multiple views of data
- Replication of data
- XFS™ performance

Shared Files with CXFS



- Single view of data
- No replication
- XFS performance


IRIX™



Tape Library




IRIX




Tape Library



TP9400




32-bit Linux




NT

SGI® TP9400



Linux®



Windows® NT

SGI® Data Management Solutions in Action

Data Management in Media

Customer Challenges

- Moving mountains of extremely large data files
- 30TB of online storage, 40TB near-line storage to create a feature-length movie

Workflow

- An “army” of postproduction artists using multiple software applications operating on the same data
- Digitization, color correcting, editing, effects, compositing



The Lord of the Rings

Conventional Solutions

- Importing a production sequence take can take hours with conventional NFS
- Replicating data further increases storage costs and impairs manageability

Laboratoires Éclair

- Customer Overview
 - French leader in movie postproduction
 - Moving from Analog to 100% Digital
- Customer Issues
 - Need to move huge datasets from one host to another for the post-production workflow
 - As many as 12 transfer are necessary to post-produce a movie, each transfer can be overnight
- Customer requirements
 - Decrease time to transfer data between hosts
 - Streamline the workflow

Laboratoires Eclair

Number of personnel: 350

Number of post-prods (2001): ~12
including 2 full digital

Average dataset size: 7GB

Curent number of datasets: ~ 500
per production

Online Storage: 6TB

Projected Growth: Investigating
HSM solutions (DMF)

Results

Boosted media workflow by a 1:3 ratio enabled Laboratoires Eclair to work in a 3x8 hours rolling process. Previously one 10 shift.

San for Storage Consolidation and Management + File Sharing using CXFS™

Data Management in Weather

Customer Challenges

- Supply real-time weather data and forecasts to military and civilian agencies
- Provide selective access to classified and nonclassified users
- 8TB of online storage, 80TB near-line storage

Workflow

- 365x24 operation; 6 million observations per day input into two weather models running on 512CPU supercomputers
- 2TB of data created per day
- In emergency situations, must be able to acquire all computational and data resources



Image courtesy of the Laboratory for Atmospheres, NASA Goddard Space Flight Center

Conventional Solutions

- Alternate solutions can't support multilevel security operating environment, selective file-access permission, preemptive operational change, file sharing between weather models, and required transfer speeds

Fleet Numeric Meteorology and Oceanography Center

•Customer Overview

- FNMOC is widely acclaimed as the world leader in operational coupled air –ocean modeling.

•Customer Issues

- Improving the resolution and accuracy of weather forecasts have resulted in increased model complexity and the addition of new models to the FNMOC workload.
- Improve work flow for time and data capacity

•Customer requirements

- 365X24 Secure Operation
- 200 GFlops in 2002
- 1TB data throughput every 12 hours

FNMOC at a Glance:

35 Officers, 65 Enlisted Personnel
160 Civilians

Output: 500,000 graphes, chartes, analysis, forecasts data sets per day

Online Storage: 8TB (SGI®TP9400)

Offline Storage: 80TB

Data Management Software:

- Trusted IRIX™
- CXFS™
- DMF

Results

- **After a careful evaluation process, FNMOC chose SGI® to fulfill these intense operational and security needs**

**San for Storage Consolidation and Management +
Backup + HSM + Shared Files using CXFS™**


Data Management in Sciences

Customer Challenges

- The Scientific community shares large numbers of current and historic brain image files to cure disease, aid in brain surgery, and study brain structure
- 2TB online storage, 40TB near-line storage, each file 100s of gigabytes

Workflow

- Researchers frequently compare current data with historical records; the same data is shared by multiple researchers working on multiple projects
- 8TB of data is added per year



Conventional Solutions

- Shared access to images using NFS is tedious; large image archive means costly laborious management and limited availability.
- Scientists want to do science, not data management

UCLA LONI - Scientific Imaging

- Customer Overview
 - Multi-user HPC center for research
 - Large image files
- Customer Issues
 - Big, complicated, admin-intensive backups
 - Data Access, Data Sharing Bottlenecks
- Customer requirements
 - Increase multi-user data access speeds
 - Streamline cumbersome, labor-intensive backup

UCLA Laboratory of Neuro Imaging

Number of researchers: ~100

Average dataset size: 20GB

Online Storage: 7 TB (SGI® TP9100/SGI® TP9400)

Available Near-line Storage: 40 TB

Migrated data: 8TB

Projected Growth: 8TB per year

Results

"The increased capacity and performance of the new storage architecture allowed... us to carry out multiple projects simultaneously using the same data on different systems, tremendously increasing our productivity while reducing non-productive waiting time and system downtime."

San for Storage Consolidation and Management + High Availability using FailSafe + HSM using DMF + File Sharing using CXFS™

Government- Geospatial Imaging

Customer Challenges

- The acquisition, processing, storage, serving, and exploitation of earth referenced images and their associated data
- Data files average size 300MB - 2Gb
- Need to provide quick access to the data for exploitation

Workflow

- RAID is used to satisfy a requirement for secure storage and fast access. Near-line storage (tape robot or offline tape) is used to supplement the RAID and may be configured as an HSM system.



Conventional Solutions


- Data accessed via NFS. Access by more than one user to the same file means waiting or making a copy. Access to a file stored by a server other than the one the user is NFS mounted on compounds wait times.

SGI® Geospatial Solution - CXFS™/SAN Architecture for TerraPoint

- 2 SAN storage units
 - Dual-ctrl Clariion (1TB)
 - Dual-ctrl SGI®TP9400 (6TB)
- 14 SAN-attach hosts
 - 14 Silicon Graphics®Octane/Silicon Graphics®O2
 - 2 Origin 200's
- 2 Brocade Switches
- Metadata network (not shown)

- 800MB LIDAR file takes 15 - 20 min to collect
- NFS feature extraction required 45 min - 1 hr
- CXFS™ and SAN, extraction takes 5 min
- Extraction process now keeps pace with other post processing operations

The customer is able to take on more projects and saves money on labor.



Feature extraction time was reduced by an order of magnitude using CXFS/SAN!