Hitachi Data Systems

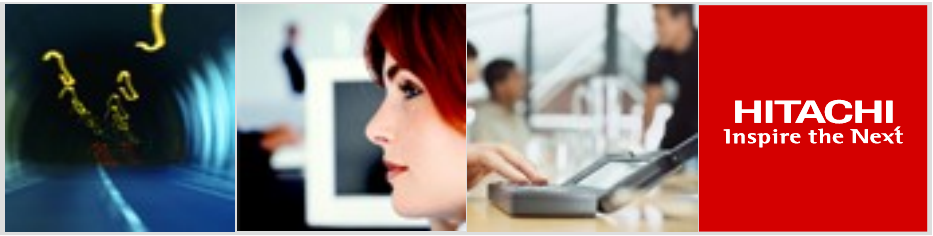

Scaling Storage Up and Out

NSC'08 Storage Track
Linköping 081014

Agenda

HITACHI
Inspire the Next

- Current Trends in the Storage Market
- Some challenges in the storage arena
- Technology Shifts Affecting Storage

Hitachi Data Systems

HITACHI
Inspire the Next

- Wholly owned subsidiary of Hitachi Ltd. based in Santa Clara, California
- Full focus on information management and data storage with hardware, software and services
- A strong Nordic organization with about 100 employees out of which 40 are based in Sweden
- Main storage supplier to many of the world's largest corporations
- Through Hitachi a manufacturer of both storage systems and hard drives

HITACHI
Inspire the Next

Hitachi Data Systems

The Hard Drive Has been Around for 52 Years - We have come a LONG way

HITACHI
Inspire the Next

Model	IBM 350	Hitachi Ultrastar A7K1000
Capacity	4,4MB	1000GB
Usage	Part of IBM RAMAC 305	Server, desktop, storage
Media Type	Fifty 24” disk platters	Five 3,5” disk platters
Rotational Speed	1200RPM	7200RPM
Areal Density	105bit/in^2	148Gbit/in^2
Max Transfer Rate	8,8kB/sec	85MB/sec
Price	$3200 per month for system with storage ($25200 in CPI’08)	$250 per unit

Trend – Exponential Data Growth

HITACHI
Inspire the Next

- The average annual capacity growth in Swedish organizations is about 70%, meaning that the data volume doubles in 18 months
- The lions share of growth stems from file, databases, and data warehousing
- More information is stored about each transaction
- Much of the capacity growth comes from copies and from test/development projects
- Little or no growth in mainframe and traditional databases

- In HPC the challenge is two-fold:
  - Scaling **UP** and scaling **OUT**

Trend - Falling Prices but Rising Cost for Storage

HITACHI
Inspire the Next

- The purchase price per GB for storage is falling very rapidly, and is driven by two main factors
  - Larger capacity drives being shipped to the market independent on underlying technology
  - A technology shift towards more SATA and other low-cost drives
- The total cost for storage is still going up in most organizations
  - Rapid price reduction is offset by even more rapid data growth
  - Cost of labor is increasing very rapidly affecting operations and maintenance prices
  - Deletion and archiving polices are few and far apart, meaning that there is constant influx to the storage infrastructure but no outflux
- This is leading to more focus on consolidation and standardization to drive down the operational cost

Trend – The Demise of Monolithic Storage

HITACHI
Inspire the Next

What is a typical high-end storage system comprised of?

Cheap, Commodity Media (HDDs) Control Unit Cheap, Commodity Media (HDDs)

- High technology central control unit encapsulated with cheap media
- 20-year-old architectures that were designed for
  - Direct attached storage with no device sharing and GBs of capacity
  - Two shift operations with maintenance windows
- To recover control unit R&D costs, more and more disks are needed
- Ever larger disk systems become more difficult to provision and refresh

This growth model is not sustainable !

Trend - Tiered Storage and Virtualization is Going Mainstream

HITACHI
Inspire the Next

- Static tiered storage is a de-facto standard in almost every organization with over 20TB
- Virtualization is finding its way into more and more enterprise customers for a variety of purposes
  - Standardization of storage services across heterogeneous hardware
  - Facilitate migration of old hardware
  - Enabling online mid-range storage maintenance online
  - Building a platform for dynamic tiered storage management
- Hitachi is leading the pack in enterprise virtualization with the controller-based virtualization in USP V and USP VM used by a large number of customers in Sweden and around the world
- Most larger companies have two or three main storage tiers and a ”tier-junk”
  - With the addition of access protocols and DR/BC schemes this generally gives 8-10 service classes

Challenge - The Performance vs Capacity Challenge

HITACHI
Inspire the Next

- Hard drive capacity is increasing much faster than performance, meaning that access density (IOPS/GB) is going down
  - Hard Disk capacity increases by 40-60% per year
  - Hard Disk performance increases by 10% per year
  - Amount of data increases by 70% per year
- Current solution to growth is often simple - Throw more hardware capacity at the problem
- This will lead to future migration and retention challenges, which will require more intelligent storage architectures with broader cost focus

Challenge - Putting an End to Exponential Storage Infrastructure Growth

HITACHI
Inspire the Next

- Exponential growth with finite resources will never be sustainable
- Factors that will eventually re quire work to reduce growth rates are:
  - Disruption in the cost reduction curve of storage capacity when hard drives require a different manufacturing process to increase density (will happen with Bit Patterned Media and/or HAMR)
  - Future cost challenges for finding the gold nuggets in today's data (compare to the cost of maintaining mainframe applications today)
  - Cost of migration becomes paramount if data is never deleted
  - Migration becomes a nightmare when service windows vanish and amount of data increases

Challenge – Addressing the Long-Term Storage of Data

HITACHI
Inspire the Next

- Fixed content is data objects that have a value over a long time, and have content that
- How long is long-term?
  - 3 years?
  - 10 years?
  - 100 years?
- Which compliance requirements or legal regulation is affecting different data?
- What information will be business-critical in the future, and must therefore be kept?
- How do we recreate information from a disk array in 100 years?
  - A hard drive has a life expectance of 10 years
  - Operating systems change major versions every two or three years
- How can we guarantee that all copies are deleted when the delete button is pressed?

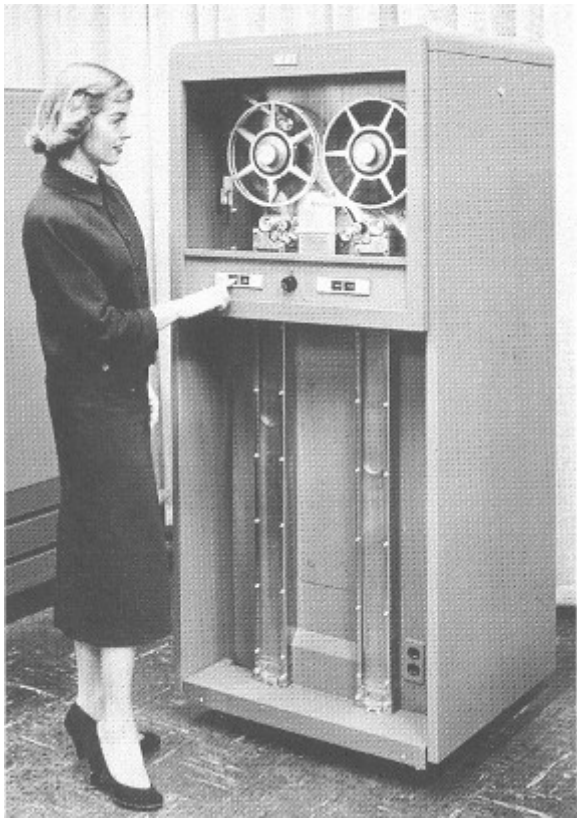

Tech Shift – Virtualization, Virtualization, Virtualization, Virtualization

HITACHI
Inspire the Next

- Virtualization x 4 is sizzling hot at the moment, and we have only seen the beginning of it
  - Server virtualization for partitioning physical servers into multiple virtual servers
  - Storage virtualization to provide shared storage services across heterogeneous storage hardware
  - File Virtualization to create a global name-space over multiple file-systems and over heterogeneous hardware
  - Network virtualization to consolidate previously separate networks (storage, wired LAN, wireless LAN, telecommunications, etc) onto a single network infrastructure with rapid deployment of services
- However virtualization is not a panacea for all problems
  - Example: Huge challenge with power and cooling developing in virtualized blade-server environments
  - Expect to see similar problems for other virtualization technologies

Tech Shift – Storage Migrations Transforming into Process Instead of Project

HITACHI
Inspire the Next

- Yesterday
  - The systems were mainly used 9-5, Monday-Friday
  - The systems were in use until they became obsolete
- Today
  - Normally a 36-48 month lifecycle of systems
  - At best 3-5 maintenance windows per year during inconvenient times
  - An average enterprise migration project with an online replacement of an enterprise infrastructure lasts for 6-12 months
- Tomorrow
  - Storage space is purchased continuously on-demand in a virtualized environment
  - Maintenance windows are out of the question, and 100% availability with minimal RPO and RTO is a requirement
  - Migration is a constant process, where data is continuously moved and new disk systems are gradually integrated into the environment

Tech Shift – The Hard Drive is Transforming

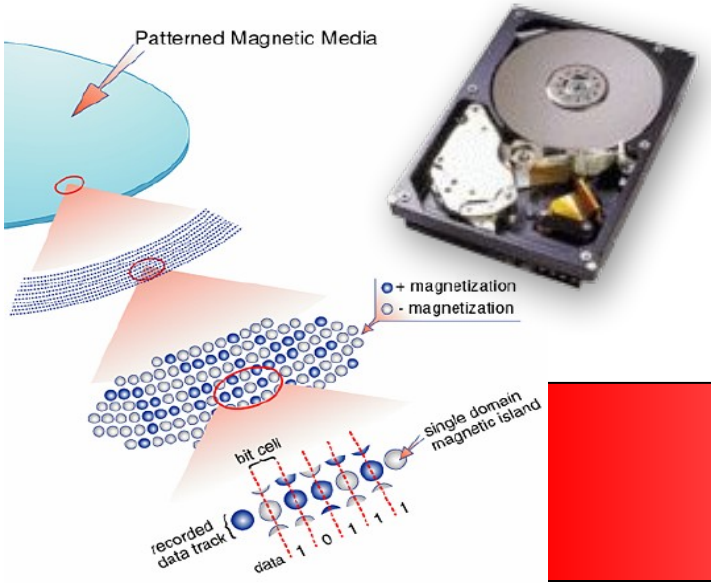

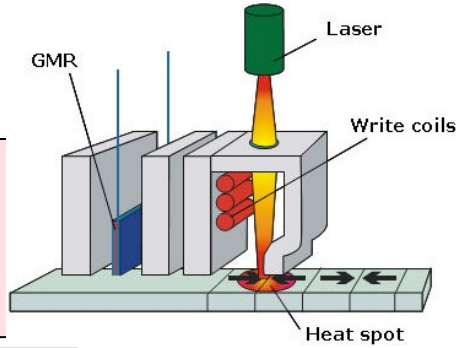

3.5" Disk Capacity

Thermally - Assisted
50TB

Patterned Media
10TB

Perpendicular Recording
3TB

Longitudinal Recording
500GB

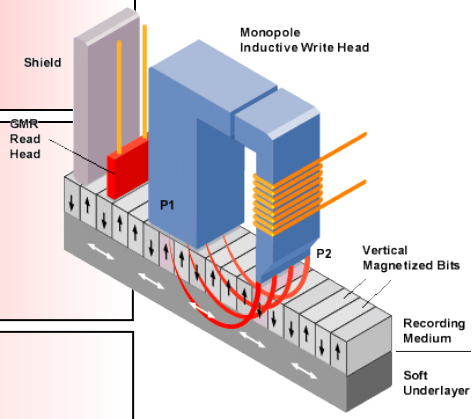

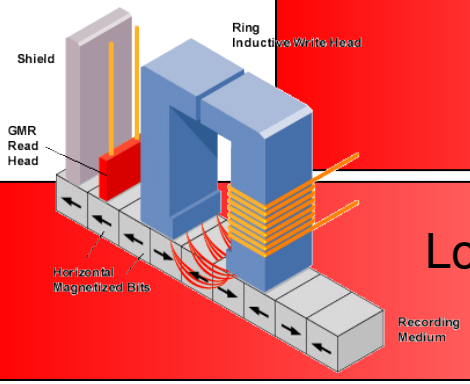

2006 2010 2013 2016

Tech Shift – The Storage Media of the Future

HITACHI
Inspire the Next

- The hard drive will for an overseeable future play an important role for storing data and has a roadmap for future capacity increases
  - HAMR, Heat-Assisted Magnetic Recording, disks with heating of the surface while writing to create increased magnetic stability (up to 50Tbit/in^2)
  - Bit Patterned Media where the magnetic layer is an ordered array of islands storing bits rather than a random mosaic
- Solid-state will increase in importance for high-performance applications and low-power scenarios
- Further into the future, methods like IBM Millipede (atomic force probing/nanotech on chip) and quantum storage (using the atomic spin (qubit) for storing information) will dramatically change the parameters of storage media

Tech Shift – The Storage Media of the Future (2)

- Solid-state for high performance applications

HITACHI
Inspire the Next

- A need for a higher performance tier than disk is being seen with hot files, OLTP and databases that are read/written by a large number of threads
- A limitation of the hard drive is the capability to handle random I/O
  - A 15k RPM Fibre Channel hard drive can handle about 150-200 IOPS
- Solid-state devices allows the highest available performance
  - Two types in the form of Flash based and DRAM based
- Texas Memory Sys. claims to have the world's fastest storage (RamSan)
  - Claims to offer 600 000 random IOPS with a latency of 15 microseconds
  - Scales up to 1TB and 3,2M IOPS in 24U
- Hitachi can offer solid state capability by locking LUNs in the DRAM cache
  - Up to 4,5M IOPS in a fully built out USP V
- Benefits of SSD
  - Constant seek time gives excellent and deterministic random performance
  - Great write-performance (DRAM only)
  - Less power consumption and therefore less heat dissipation
  - Very easy to erase data from the drives
- Drawbacks of SSD
  - Price is still high and identifying which data to place on SSD is hard
  - Poor sequential I/O performance
  - Volatility when power is lost for DRAM, and shorter reliable lifetime for NAND

Future formats of storage – File vs Block vs Object-based

HITACHI
Inspire the Next

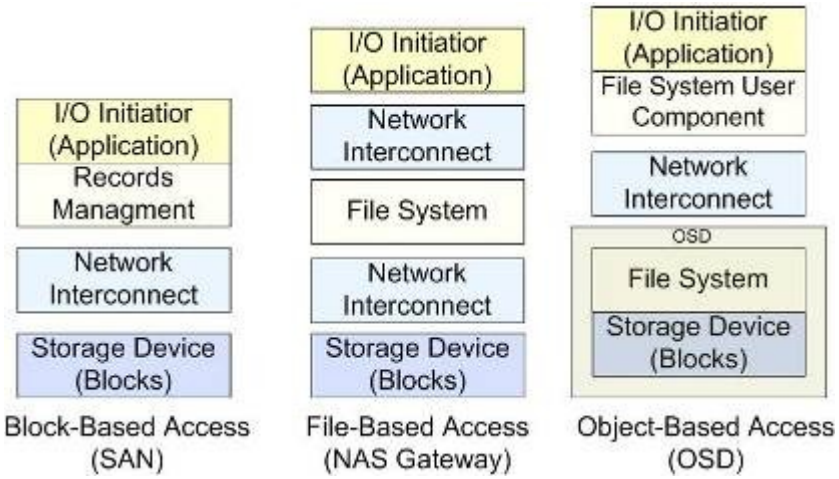

- Block-based access is the foundation for all I/O today
  - Some applications like databases and OLTP require the performance given by low-level access at the block-level, and are likely to do so for an overseeable future
  - Requires intelligence high in the stack to create records that can be understood at the application level
  - Access using iSCSI or Fibre Channel, no multi-tenancy

Future formats of storage – File vs Block vs Object-based (2)

HITACHI
Inspire the Next

- File-based access allows access to an aggregation of blocks (files) through the interface of a file system managing low-level metadata
  - Accessing files is very intuitive and sharing is straight-forward
  - NAS devices decouples the dependence on a single file server, ensuring good availability and performance
  - To improve scalability, a distributed file-system can be implemented but challenges with lock traffic and metadata traffic arise
  - End-user or application access using NFS or CIFS
- Object Based Storage offers a new paradigm of working with data
  - A collection of objects rather than indexing into array of blocks by LBA
  - Metadata and storage resource management is pushed to the very bottom into an Object-Based Storage Device (OSD)
  - Combines the scalability of block-level access with sharability of files at the expense of performance
  - The first deployments occur in archiving environments where there is signficant interest for decoupling metadata and data from applications
  - Hitachi Content Archive Platform builds on the principle of OSD
  - Likely to increase in importance at the expense of files, not an alternative to SAN yet

Scaling up for block-storage - The Universal Storage Platform V & VM

HITACHI
Inspire the Next

- The Hitachi Universal Storage Platform V / VM is not a single server that needs to be clustered for availability
  - It is a cluster of 128 / 32 processors within a physical frame sharing a global cache
- The first enterprise storage system with a switched back-end
- No single point of failure and can be maintained nondisruptively
- 4Gps end to end
- Dynamic HDD Sparing
- HiTrack® “call home” service/remote maintenance tool

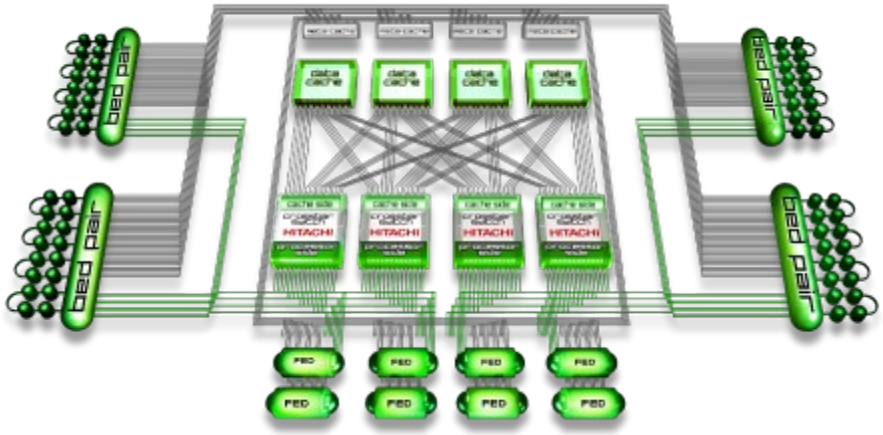

4Gb/sec end-to-end internal architecture including front end, disks, and back end

Design for Scaling up and out for File Services - High-performance NAS Platform Architecture

HITACHI
Inspire the Next

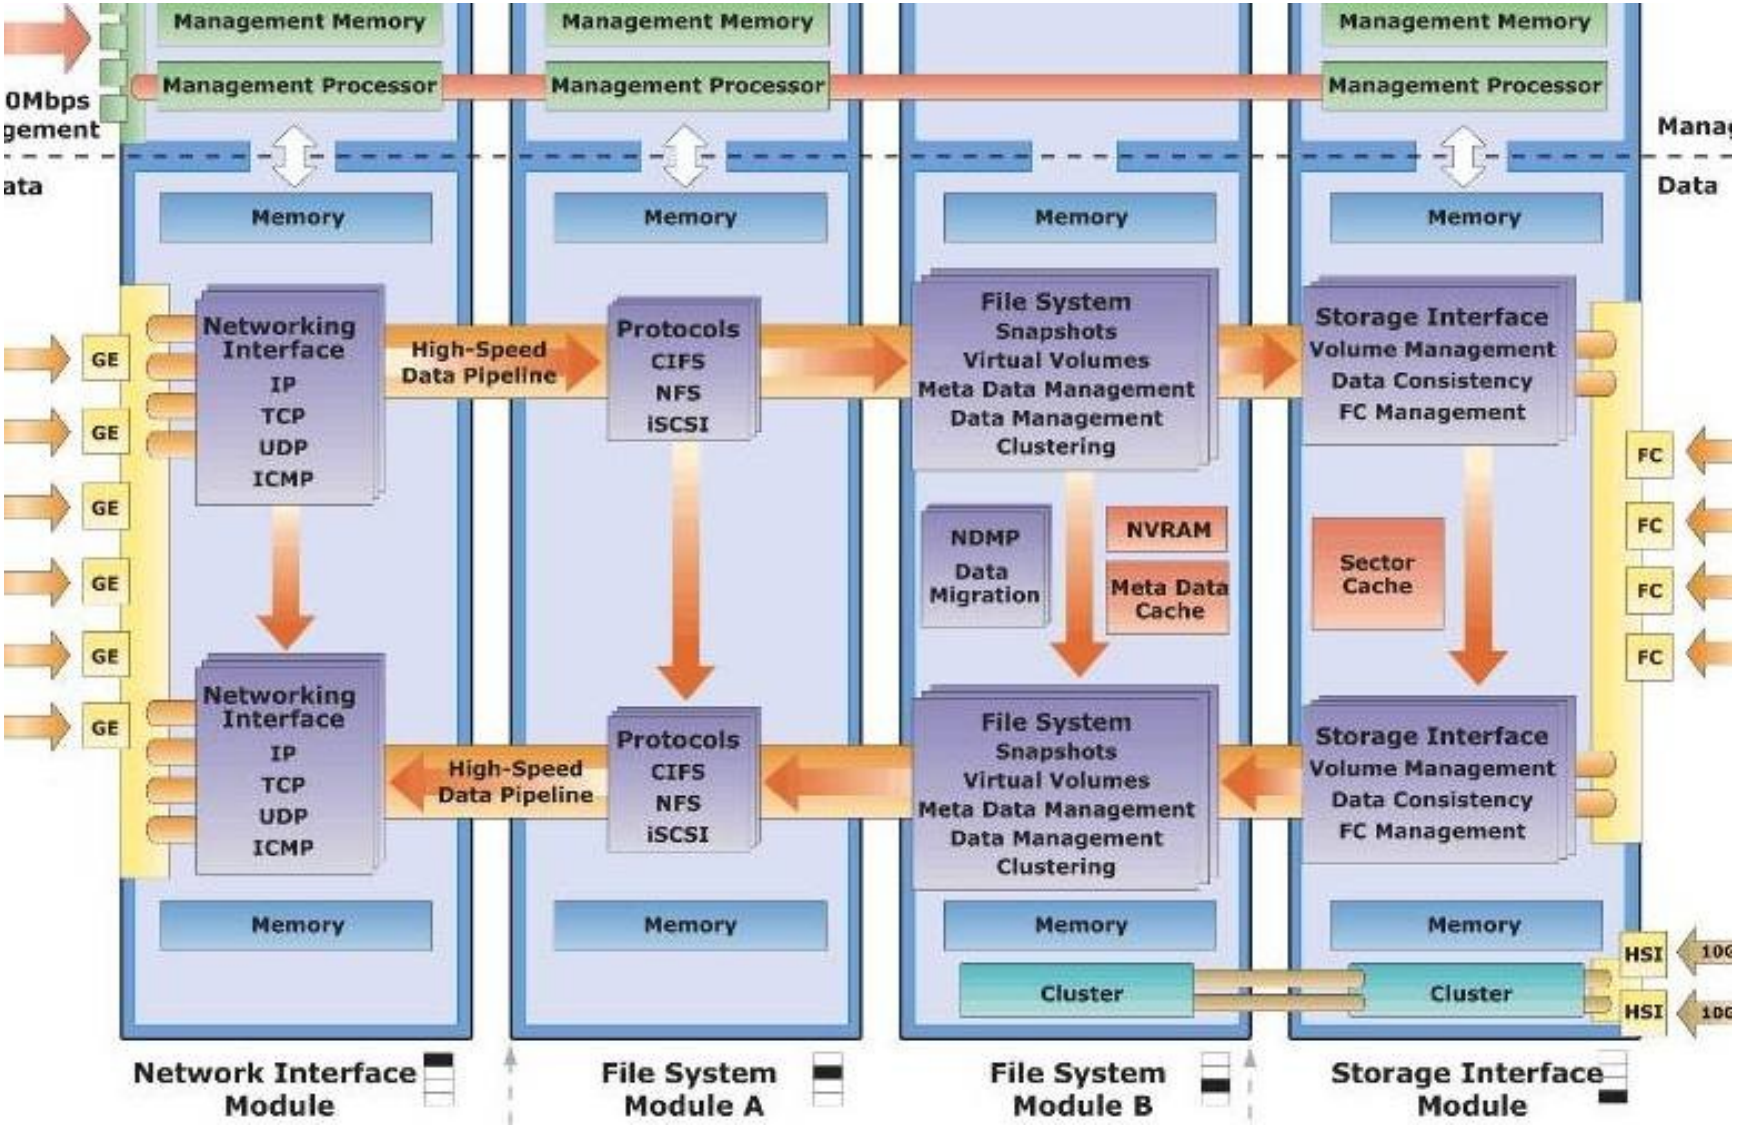

Hitachi Data Systems

Summary

HITACHI
Inspire the Next

- Scale out! – Use low cost components with a shared file system or invest in a more expensive shared storage array behind a powerful NAS
- Scale up! – Consider SSD as a serious alternative provided that data is short lived;

- No matter what – Keep data in motion and purge things that are not needed

Questions and Comments?

HITACHI
Inspire the Next

E-mail: staffan.strand@hds.com

Thank You!

HITACHI
Inspire the Next

Thank You for Your Attention!

E-mail: staffan.strand@hds.com

HITACHI
Inspire the Next