

Storage & Storage Management

Dennis Martin

RMWTUG April 2005 Meetings

Agenda

- Storage Interfaces
- Replication Technologies
- Storage Resource Management
- Information Lifecycle Management

Storage Interfaces

- Hardware Block Storage Interfaces
 - ATA & Serial ATA (SATA)
 - SCSI & Serial-Attached SCSI (SAS)
 - Fiber Channel
 - ESCON & FICON (mainframe)

Storage Interfaces

- Software Block Storage Interfaces
 - iSCSI
- Software File Storage Interfaces
 - NFS & CIFS

ATA and SATA

- ATA also known as IDE
- Used in desktops, laptops and low-end servers
- Very low-cost
- Very high-volume
- Lower reliability than SCSI, SAS, FC
- Used for hard drives, CD-ROM, DVD
- Moving to serial interface for hard drives

SCSI and SAS

- Used in higher-end workstations and servers
- Higher cost than ATA & SATA
- Lower volume than ATA & SATA
- Used for hard drives, CD-ROM, others
- Industry moving to serial interface for hard drives
- Longer warranties than ATA & SATA

Fiber Channel (FC)

- Designed for high-performance, high reliability and long distance
- Hard drive costs, volumes and warranties are similar to SCSI drives
- Can be used in a SAN
- Three configurations
 - Point to point
 - Arbitrated Loop
 - Fabric

ESCON & FICON

- Mainframe only
- ESCON – up to 17 MB/sec
- FICON – ESCON encapsulated over FC
- Can use copper or fiber-optic cabling
- Mainframe storage subsystems command a premium price because of extra hardware and features required for mainframes

Mixing and Matching

- SATA and SAS drives use the same connectors.
- Some disk arrays use FC interface to the outside world, but use SATA drives internally. This is known as FATA.
- Mainframes can share “plumbing” of SAN with open systems, but generally use distinct storage devices.

Hardware Block Interfaces

	ATA	SATA	SCSI	SAS	FC
Number of devices	2	1	8 or 16	16K	16M
Maximum distance	18 inches	1 meter	3 – 25 meters	10 meters	100+ KM
Cable type	Copper	Copper	Copper	Copper	Copper or fiber optic
Interface type	Parallel	Serial	Parallel	Serial	Serial
Transfer Speeds (MB/sec)	Up to 133	150	Up to 320	300	100, 200, 400
Hard drive Rotational Speed (RPM)	Up to 7.2K	Up to 10K	10K, 15K	10K, 15K	10K, 15K

Software Block Interfaces

- iSCSI (Internet SCSI)
 - SCSI protocol over TCP/IP
 - Long distance
 - Slower performance (actual data transfer) than the direct storage hardware interfaces due to the software overhead
 - Speed improvements with accelerator cards
 - Software and hardware targets available
 - Can be used in a SAN

Software File Interfaces

- Used for applications that need file-oriented access
- NFS – Available for UNIX & Linux systems
- CIFS – Available for Windows systems

Trends

- Disk Drives
 - Mainstream disk drives moving to 2.5 inch form factor over the next couple of years
 - Maximum RPM has reached a plateau for now
 - SATA will introduce deeper command queuing, approaching what is available for SCSI & FC
- Interfaces
 - SATA roadmap includes 300, 600 MB/sec
 - SAS roadmap includes 600, 1200 MB/sec
 - FC roadmap includes 800, 1000 MB/sec
- Other
 - RAID-6 (double-parity) will be used for SATA subsystems

Windows Simple SAN

- Goal: to drive robust, serviceable, easily-deployable SAN technology to all customers
- Targeted at the small and medium business market
- Available for Fibre-channel and iSCSI technology
- Affordable components
- Configurable for complete solutions
- Uses Windows Server 2003 storage features
- Listing of solutions on Microsoft web site
- Expect more emphasis in 2005

Replication Technologies


- Synchronous & Asynchronous
- Replication Types
- Location
 - Host software-based
 - Appliance-based
 - Storage subsystem-based

Replication Methods

- Synchronous – I/O must complete to both devices before response is sent to host
- Asynchronous – I/O completes at primary device before response is sent to host
- Either method can be implemented in hardware or software

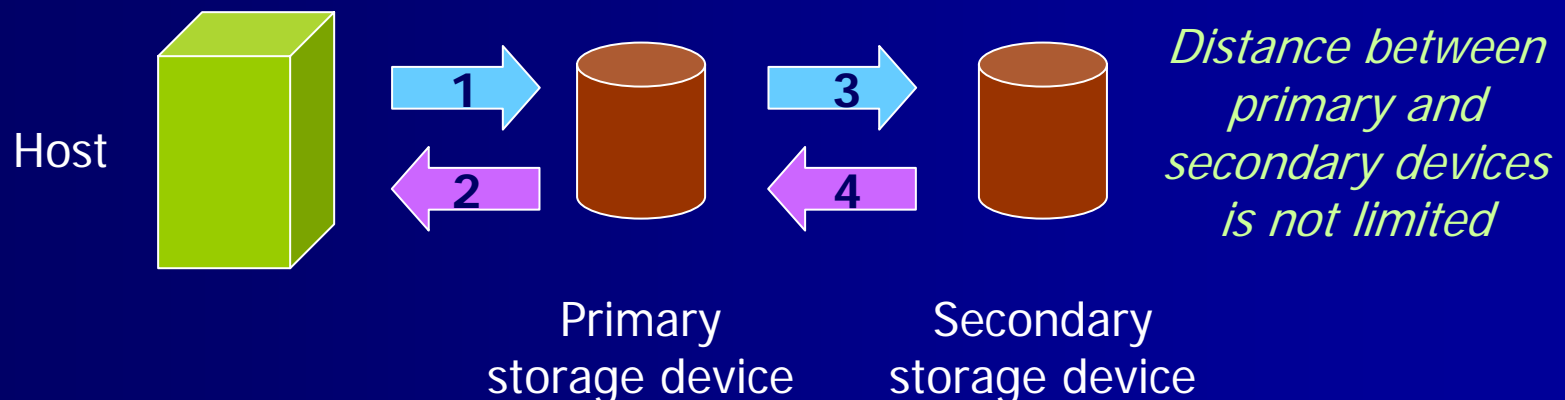
Synchronous Replication

1. Application writes to primary device and waits
2. Primary device sends write to secondary device, either via device or host software
3. Secondary sends response to primary
4. Primary sends response to application, and then application can continue with next write



Asynchronous Replication

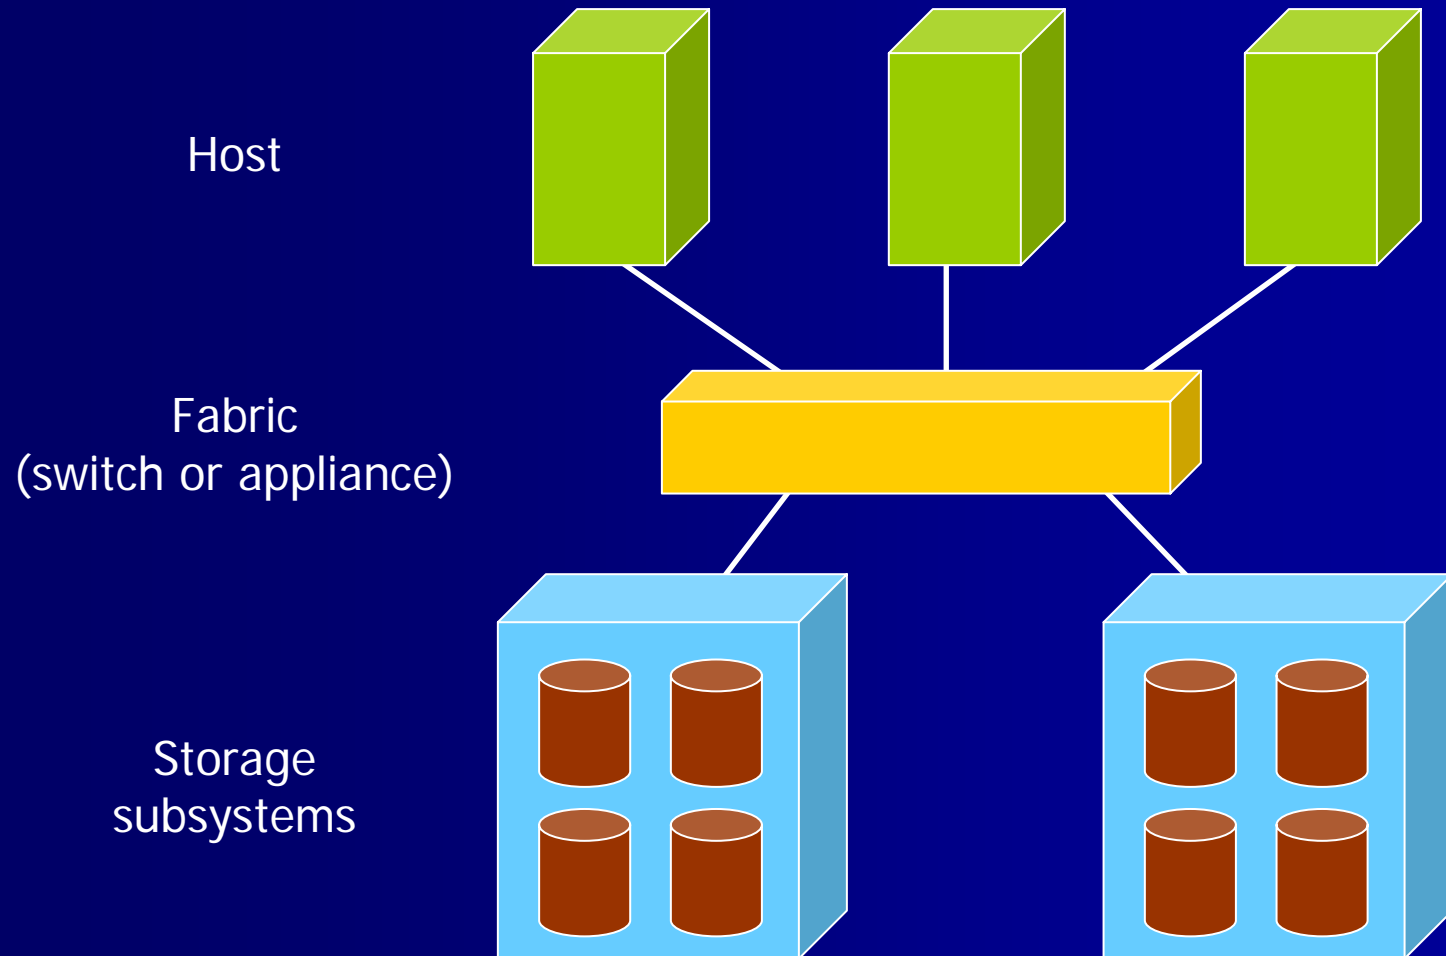
1. Application writes to primary device and waits
2. Primary device sends response to application, then application may continue with next write
3. Primary sends write to secondary device, either via device or host software
4. Secondary sends response to primary



Replication Types

- Point-in-Time (PIT) Copy
 - A copy of data at a specific instant in time
- Remote Copy
 - Copy of data at some geographical distance
- Mirror
 - Local copy of data in synchronization
- Data Migration
 - Move data between different storage systems or different hosts

Replication Location



Replication Location

- Host
 - Works at volume, file or file system level
 - Consumes host resources
 - Independent of storage hardware
- Fabric
 - In the data path and adds latency
 - Specific to that switch or appliance
- Storage Subsystem
 - High performance
 - Expensive
 - Generally specific to that subsystem

Storage Resource Management (SRM)

- Measure and report storage usage
- Activity by file, user, server, etc.
- Active vs. Passive
- Application aware (database, email, etc.)
- Beginning to absorb SAN management solutions
- Pricing: by server, SAN switch port count or by storage capacity managed, some extras

Passive SRM

- File activity: by server, user, file type, etc.
- “Vulnerable” files – no backup copy
- “Wasted Space” – duplicate files
- “Obsolete” files – not accessed in certain time period
- “Orphan” files – creator of files no longer with company
- Charge-back
- Quota reporting (not enforcement)

Active SRM

- Capacity actions: provisioning (add or grow LUNs/volumes/filesystems), file movement or deletion, etc.
- Performance actions: re-routing paths, file movement, etc.
- Business Continuity actions: make extra backup, make remote copies, etc.
- Records Retention/Compliance actions: move data to alternate location or archive

Information Lifecycle Management (ILM)

- Defined time periods for data creation, backup, retention and deletion
- Track movement and location of data throughout its life
- Metadata is important
- Rules-based automated actions & policies
- Builds on SRM software
- Not all data treated the same
- Must traverse all storage management layers

Data Definitions

■ Activity level

- Active – data in use
- Reference – data may be required, many times read only
- Archive – data not expected to be needed

■ Structure

- Structured – has predefined format for access
- Unstructured – no predictable access controls
- Semi-structured – access control added to unstructured data

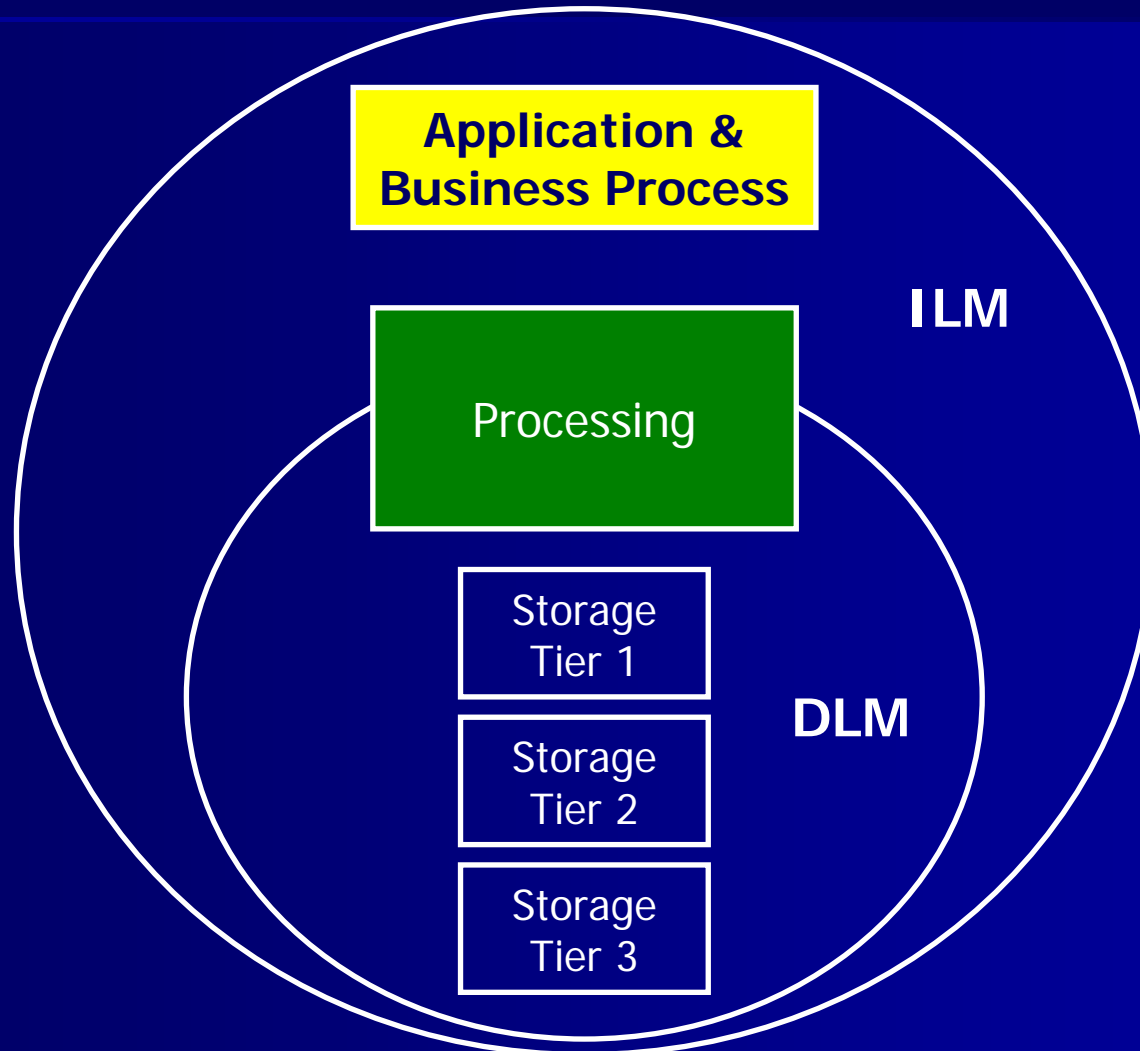
■ Location

- Online – readily available
- Nearline – accessible without operator activity
- Backup – copies of data to disk or tape
- Archive – removable media

ILM & DLM

- Information Lifecycle Management
 - Requires information about business process
 - Holistic view of managing information
 - Uses functions of Data Lifecycle Management
 - Includes all data – structured / unstructured / and semi-structured
- Data Lifecycle Management
 - Focus is to make efficient use of storage resources
 - Based on characteristics of data and data access

ILM & DLM Diagram



Evaluator Group

- Independent Industry Analysts
- Focused exclusively on computer storage
- Web site includes online database of product information
- Advanced storage education: storage infrastructure, management and enterprise
- Free monthly newsletter
- www.evaluatorgroup.com