



THE FUTURE OF DATA STORAGE



WITH NO CLEAR PATH
IN SIGHT,
WILL ALIENS SOLVE
OUR STORAGE
PROBLEMS?



**Closing Keynote
Presented by
Michael Del Rosso, CTO
Quantum**

Quantum.

NETWORLD+INTEROP

NETWORK STORAGE DAY 2004



Trends and Directions for Storage 2004 and Beyond

Michael J. Del Rosso

Chief Technology Officer

Quantum Corporation



Agenda

- A Brief History of Unmet Promises
- Information Management Modeling Language (IM²L)
- Information Lifecycle Management (ILM)
- Continuous Data Protection (CDP)
- Intelligent in the SAN (NAAAY)



A Brief History of Unmet Promises

"Those who cannot remember the past are condemned to repeat it." – George Santayana

- Late 80's
 - » DCE/DFS/OSF, IEEE MSSRM V.4
- Early 90's
 - » ISO protocols, ADA
- Mid 90's
 - » ATM to the desktop
- Late 90's
 - » Client/server
- Early 00's
 - » Virtualization, SSP
- Mid 00's
 - » ILM?

WHAT'S THE COMMON FACTOR?



Information Management Modeling Language (IM²L)



The Case for IM²L

- Define common vocabulary
- Avoid definition-related arguments/discussions/misunderstandings
- Elevate conceptual level beyond physical infrastructure
- Avoid vendor lock-in
- Avoid thinking of Information Management needs in terms of vendor product capabilities
- Help establish enterprise-wide understanding of information management requirements
- Allow definition of Service Level Objectives (SLO) for information management
- Allow determination of organization's maturity level in dealing with information management
- Can be used for design, planning, and documentation



What IM²L Covers

- Data generators (applications), collectors, classifiers, repositories
- Storage classes
- Policies:
 - » retention
 - » security
 - » placement
 - » migration
 - » replication
 - » protection
 - » recovery objectives
- Metadata: required versus available; generation and aggregation points
- Value of data
- Costs of storage operations, and costs of movement operations
- Data volumes, data growth



Various IM²L Considerations

- What is the right level of granularity for describing:
 - » Applications
 - » Information
 - » Volumes
 - » Filenames
 - » Directories
 - » Tables
- Syntax, diagrams and symbols



Information Lifecycle Management (ILM)



What is ILM? – The Cynical View

- ILM is a soup of storage management process and technologies being stirred to reposition products into meaningfulness
 - » HSM
 - » Backup
 - » SRM
 - Virtualization
 - SAN
 - SDM
- ILM is not a product
 - » Beware companies that claim to have ILM today
 - » Possibly a ploy for valuation



ILM Vendor & Analyst Perspectives

- Gartner
 - » “ILM is an approach to information and storage management that recognizes that the **value** of data changes over time.”
- StorageTek
 - » “Align the cost of storage with the **value** of customer information.”
- Enterprise Storage Group
 - » “ILM is a combination of technology and methodology that helps end users manage data from the moment it is created to the time at which it is no longer needed.”
- EMC
 - » “(A lack of ILM) results in overinvestment in the protection and availability of **non-critical** information... ILM is the optimal management of information throughout its life, from creation and use to archiving and disposal.”



It's all about the **Value**

- If it's all about the **value**, who determines the **value** of your data?
 - » Application administrators, CIOs, CFOs, Sys Admins,
 - » It can't be automated, it can't be inferred
 - » What is the taxonomy of **value**?
- Differentiation of the value of data in a single application
 - » It is possible to effectively manage email today
 - » It is possible to sort the attachments from the message, search on message content
 - » It is impossible to determine personal from business, business irrelevant from business critical, so it must all be managed the same
 - » And, if one decides that after three years all email is to be deleted, how does one go about deleting all the backups of email too?
- Imagine **valuing** individual records in a database?
- **Value** drives Recovery Time Objectives, Recovery Point Objectives, security, availability and cost sensitivity
- Vendors' classifications of value map to their products (is your data CLARiiON important or Symmetrix important?)



Impediments to *Comprehensive* ILM

- Outstanding Issues
 - » Data Valuation
 - To support the automatic migration between cost layers
 - » Heterogeneity
 - If you think SANs were tough...
 - Imagine getting one vendor's value generator to work with another's auto-migration tool
 - » Data Protection
 - Still the primary pain point today
 - » Regulatory Compliance
 - They are just getting started!
- According to Gartner, ILM will be 80% slideware and 20% product through 2005



Continuous Data Protection (CDP)



Data Protection Customer Pain Points & Pressures

- Pain Points
 - » Solution cost
 - 30% of storage budget
 - » Performance (at least 2x too slow for critical apps)
 - Backup window
 - Recovery time
 - » Data Loss
 - 52% of Restores fail (unprotected window of vulnerability + media reliability)
 - » Cost of downtime
 - For 54% of businesses > \$50K per hour, for 39% > \$100K per hour
 - » Enterprise critical apps are: Transaction processing, Email, Financials, ERP, CRM, SCP, Decision Support
 - » 30% of Enterprises need <1 hr recovery, 60% <4hr for critical apps

- Data protection pressures
 - » Data growth (70% annually)
 - » Regulatory compliance (HIPAA, SEC 17a-4, Sarbanes-Oxley)
 - » Viruses, hackers
 - » Round the clock availability




Classic Data Protection Technologies

Mirroring: Improves availability, often mis-categorized as a data protection tool

Snapshot : Primarily solve some of the limitations in client-server backup regarding backup window, but aren't safe enough for Data Protection

Client Server Backup : The most widely used and understood data protection practice. Performance, cost and data loss potential are huge drawbacks





Continuous Data Protection

- What is it?
 - » Continuous capture of all write activities in real-time
- Benefits
 - » Speed
 - Backup time: imperceptible
 - Restore time: seconds
 - » Protection
 - No more window of vulnerability
 - Restore to any point in time
 - » Cost Savings
 - Software, hardware, management



CDP Implementations

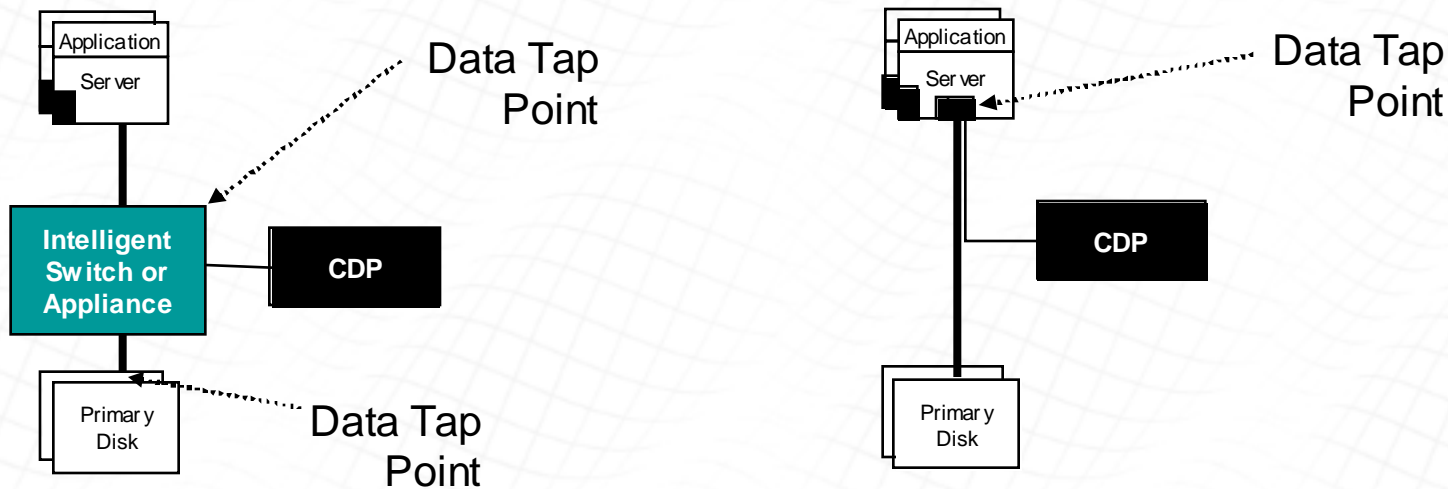
- Host-centric
 - » Driver or volume manager duplicates writes, sends to backup appliance or backup disk array
- Fabric appliance or switch
 - » Device duplicates writes, sends to backup disk array
- Storage-centric
 - » Primary storage array duplicates writes, sends to backup disks or separate backup array

- In-band vs out-of-band
 - » Everything has an in-band component
 - » Hardware or software



CDP data path configurations

- CDP data taps
 - » Host-based with driver shim or volume manager
 - » Fabric-based with switch or appliance
 - » Primary storage





Intelligence in the SAN



Intelligence in the SAN

- Today higher level storage operations reside in:
 - » Hosts
 - » Storage arrays
- Examples
 - » Virtualization / Volume Management / Storage Pooling
 - » Replication
 - » Tape library emulation
 - » Continuous Data Protection (CDP)



Intelligence in the SAN

- Benefits of intelligence in the fabric:
 - » vs Hosts
 - Aggregation/concentration instead of widely dispersed
 - No host performance impact
 - Reduces management burden
 - » vs Primary Storage
 - Services have visibility of multiple storage systems
 - Brings services to lower-end storage
 - Simplifies multiple vendor-specific approaches
 - Reduces management burden
- Unique fabric functionality
 - » Intelligent load management (QoS)



Intelligence in the SAN

- Common choices for SAN-based services
 - » Appliance
 - Often Intel server-based
 - Bandwidth, latency, scalability
 - Availability
 - » Intelligent Switch
 - Custom ASICs
 - Historical availability focus
- New business issues for intelligent switch vendors
 - » Go to market
 - » Support
 - » Multi-vendor application co-residence
 - » Platform versus embedded
 - » Management
 - » Over time broader APIs will resolve all this

NETWORLD+INTEROP

NETWORK STORAGE DAY 2004



NETWORLD+INTEROP

NETWORK STORAGE DAY
Las Vegas, Nevada 12 May 2004

2004

ADDRESSING THE STORAGE SCHISM

IS THERE ANY HOPE OF ALIGNING STORAGE TECHNOLOGY SPENDING WITH BUSINESS OBJECTIVES?

N+I 2004

PROTOCOL WARS

I-WARP
SAS/S-ATA
PEER TO PEER NETWORK STORAGE INTERCONNECT IS ALREADY UNDER WAY

Are you ready to choose sides?

N+I 2004

INVASION OF THE INTELLIGENT STORAGE DEVICES

PEER TO PEER OR JUST ANOTHER PANACEA?

N+I 2004

ENHANCED DATA PROTECTION SOLUTIONS

REAL IMPROVEMENTS OR REPACKAGED TECHNOLOGY

N+I 2004

THE INTELLIGENT STORAGE UTILITY

WE WANT TO BELIEVE... BUT...

N+I 2004

THE FUTURE OF DATA STORAGE

WITH NO CLEAR PATH IN SIGHT... WILL WE EVER SOLVE OUR STORAGE PROBLEMS?

N+I 2004

Featuring speakers and panelists from more than 25 of today's most important storage technology vendors... An opportunity to interact with peers and industry leaders...

Hosted by The Data Management Institute LLC




BE THERE. YOUR DATA MAY DEPEND ON IT.

Thank You
for
Your Participation!



COMING ATTRACTIONS



“ALAS, POOR DISK DRIVE, I KNEW HIM...”