Storage Decisions

Hosted by SearchStorage.com STORAGE

Real-World Tiered Storage

Craig D. Taylor
Associate Director, Open Systems

cme
Chicago Mercantile Exchange

Storage Decisions

Hosted by SearchStorage.com STORAGE

Real-World Tiered Storage

Presented by Craig D. Taylor
Associate Director Open Systems
Chicago Mercantile Exchange
ctaylor@cme.com
www.cme.com

Storage Decisions

Hosted by SearchStorage.com STORAGE

What I assume you know

- **Storage area networks (SAN)**
- **WORM media and regulatory type storage**
- **Disaster recovery (DR)**
- **Sarbanes Oxley, SAS70, etc.**
- **Your own challenges and issues**

Storage Decisions

Hosted by SearchStorage.com STORAGE

By the end of the session, you'll know the following:

- **One practical tiered architecture approach**
- **Storage risk assessment**
- **MAID (Massive Array of Idle Disks)**
- **ATA storage offerings pros and cons**

Storage Decisions

Hosted by SearchStorage.com STORAGE

Real-world tiered architecture

- The CME (Chicago Mercantile Exchange)
- Current environment
- Challenges
- Develop a storage strategy
- Backup to disk
- New technology direction
- MAID architecture vs. standard ATA offerings

Storage Decisions

Hosted by SearchStorage.com STORAGE

Overview of CME

- Founded in *1898*
- North America's largest futures exchange and the world's largest futures clearinghouse
- Notional value of contracts traded daily exceeding $3.7 trillion last month.
- Annual growth of transactions tracking around 100% - 140%
- Global footprint

Storage Decisions

Hosted by SearchStorage.com STORAGE

Increasing workloads

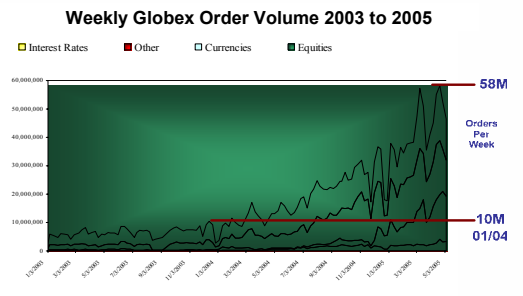

Storage Decisions

Hosted by SearchStorage.com STORAGE

*** Press Release ***

CME reports April volume

- **Average daily volume for April was 4.9 million contracts, up 47% from previous year.**
- **CME's fourth consecutive record month.**
- **Average daily volume on CME® Globex® rose to a record 3.4 million contracts, a 113% increase from April 2004.**

Storage Decisions

Hosted by SearchStorage.com STORAGE

CME's current infrastructure

- **In 2 1/2 years, CME's SAN has grown from 4 TB to over 180 TBs.**
- **In the same time frame, Unix/Linux server count went from 500 to over 2,000 servers**
- **Linux presence has grown from a few servers to over 500 in one year. (Now over 1,000)**
- **3 data centers**

Storage Decisions

Hosted by SearchStorage.com STORAGE

Data growth challenges

- **New technology implementation while growing data volume exponentially**
- **I/O performance critical to application performance**
- **Ever-changing regulatory requirements**
- **Disaster recovery**
- **Capacity planning**
- **Backup environment**

Storage Decisions

Hosted by SearchStorage.com STORAGE

Questions to be answered

- **Where does data belong?**
- **How long do we keep it?**
- **What does legal say about data retention?**
- **What is the service level for data at points throughout the life cycle?**
- **How do we back it up within our window?**
- **How do we plan for unpredictable growth?**

Storage Decisions

Hosted by STORAGE SearchStorage.com

How much of your regulated data do you have clear, "legally" defined retention periods for?

1. **100% - Our legal department has signed off on all of our retention policies.**
2. **75%**
3. **50%**
4. **25%**
5. **0% - I have no idea. . . .**

Audience Response
sponsored by TechTarget

Storage Decisions

Hosted by SearchStorage.com STORAGE

Objective -- Develop a storage strategy

- **Identify success factors**
  - **"Data classification"**
  - **Service-level performance**
  - **Cost reduction (containment) (Value in the market)**
  - **Management capabilities (Ease of management)**
  - **Compatibility**
  - **Etc.**
- **Your success factors may be different**

Storage Decisions

Hosted by SearchStorage.com STORAGE

Data classification

- **Use broad strokes (tiers)**
- **Develop data classes, (tiers), based on:**
  - **Performance criteria**
  - **Service level**
  - **Disaster recovery**
  - **Regulation**
  - **Retention**

Storage Decisions

Hosted by SearchStorage.com STORAGE

Tiers of data

- **Tier I**
  - Critical applications and databases that need high performance
  - and/or replication
- **Tier II**
  - Other production databases, QA, no replication
- **Tier III**
  - Data with long-term, regulated retention. Regulatory reports,
  - SOX records, e-mail, etc.
- **Tier IV**
  - Backup and restores/synthetic full backups
  - NearLine file system storage
- **Tier V**
  - Tape for offsite storage

Storage Decisions

Hosted by SearchStorage.com STORAGE

Tier platforms

- Many tiers already existed in the current environment

TIER I *DMX*	TIER II *Hitachi*	TIER III *WORM*	TIER IV *ATA*	TIER V *Tape*
Critical and Replicated Data	Lower Priority, Non-Replicated Data	Archived Regulatory Data	ATA, Virtual Tape	Disaster Recovery, Off-Site Data
 EMC DMX	 Hitachi 9960	HP 2200MX 	?	STK PowderHorn Silo 

Storage Decisions

Hosted by SearchStorage.com STORAGE

Tier platform issues

- Tier III was inadequate, not scalable, not flexible, and difficult to replicate.
- Tier IV was defined but didn't exist in our environment.
- Many Tier V issues

Storage Decisions

Hosted by STORAGE SearchStorage.com

How satisfied are you with your current backup environment?

1. **Very satisfied.**
2. **It gets the job done.**
3. **It usually gets the job done, but there are plenty of issues.**
4. **I'm in big trouble, lots of backup and restore failures, backup windows are too small, etc.**

Audience Response sponsored by TechTarget

Storage Decisions

Hosted by SearchStorage.com STORAGE

NetBackup environment challenges

- STK Silo arm(s) are controlled by our mainframe. NetBackup communicates through ACSLS.
- Mainframe "tech time", in the middle of weekend backup window.
- 3 distinct "physical", environments and only 2 silos.
- Instances of very long running backups.
- High level of restore failures.
- High occurrence of backup failures
- Back up Windows and Novell in addition to all Unix.

Storage Decisions

Hosted by STORAGE SearchStorage.com

Are you using any type of disk-based backup in your environment?

1. **Yes, I back up mostly to disk and use little tape.**
2. **Yes, I'm staging to disk and still most data to tape.**
3. **No, but thinking about it.**
4. **No, tape only. Tape still does the job for me.**

Audience Response sponsored by TechTarget

Storage Decisions

Hosted by SearchStorage.com STORAGE

Backup-to-disk solutions

Do the research, know your options!

- **Initial take on backup-to-disk solutions was to use straight disk, no need for virtual tape**
  - **Talented staff**
  - **No additional software costs**
  - **Easy to manage**

Storage Decisions

Hosted by SearchStorage.com STORAGE

Backup-to-disk solutions (Continued)

- **Evaluated 40 TB "usable" ATA solutions (x3)**
  - **RFQs and RFPs to several vendors.**
  - **Standard low-cost ATA platforms "very similar".**
  - **Discovered new architecture, MAID**
  - **(Massive Array of Idle Disks)**
  - **Reworked our requirements to better compare MAID architecture and standard low-cost ATA.**

Storage Decisions

Hosted by SearchStorage.com STORAGE

Backup-to-disk solutions (Continued)

- **Direction change from backup to disk, to virtual tape**
  - **No compression with straight backup to disk. (NetBackup compression is client-based)**
  - **Disk space management issues with straight backup to disk.**
  - **Replication challenges and limited options with straight disk.**
  - **What is the purpose of ATA storage in our environment?**

Storage Decisions

Hosted by SearchStorage.com STORAGE

CME's new technology direction

- **New technology evaluation and integration**
  - No "bleeding-edge" technology, vendor financial stability.
  - Evaluate from a "risk", point of view.
- **What changed and why?**
  - Growth was traditionally underestimated.
  - Development cycles too slow to meet sudden volume increases.
  - Risk variables changed.
- **Risks with standard low-cost ATA vs. MAID**
  - Almost all risk factors go against MAID startup vendor.
  - Gartner recommends evaluating new players.
  - Clear back out plan with current infrastructure.

Comparing MAID vs. standard ATA

Cost

- Standard ATA – typically higher than tape.

Scalability

- Standard ATA about 14 to 56 TB (depending on vendor)
- MAID Storage – 224 TB capacity
- Most inexpensive ATA storage systems have a low maximum capacity requiring multiple systems to match the capacity of 1 system with MAID architecture.

Storage Decisions
Hosted by SearchStorage.com STORAGE

Comparing MAID vs. standard ATA

Factor	MAID		Standard ATA	
	56 TB	224 TB	56 TB	224 TB
Floor space	10 Sq. Ft.	10 Sq. Ft.	15.4 Sq. Ft	46.2 Sq. Ft.
Power consumption	1 to 1.6	3.4 to 5.8	6.24	25.76
BTU output	3,650 to 5,667 BTU/hr	11,840 to 20,000 BTU/hr	21,959 BTU/hr	87,837 BTU/hr
FalconStor	1 controller	2 controllers	2 per config	8 per config
Fabric connections	4	4	8	32

Storage Decisions
Hosted by SearchStorage.com STORAGE

Other MAID benefits

- **SATA drives are designed for power cycling, typically less than 50% duty cycle. The MAID architecture uses SATA as intended and spins down drives when not in use.**
- **SATA drive MTBF is 400K hrs vs. 1M hrs for SCSI. The MAID architecture has a maximum-drive duty cycle of 25%. Theoretically, 4X expected service life of always-on SATA drives.**
- **Super dense, scalable frame and simple management.**

Storage Decisions

Hosted by SearchStorage.com STORAGE

Comparing MAID and standard ATA footprint

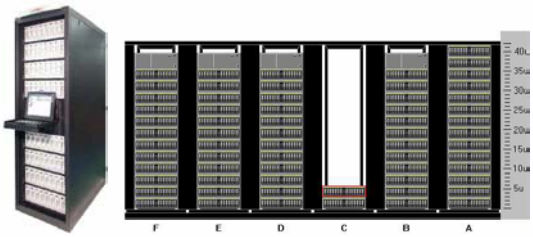

Storage Decisions

Hosted by SearchStorage.com STORAGE

MAID testing and implementation results

- **Backups and restores were up to 228% faster on a MAID architecture system compared to tape.**
- **No backups have failed on the new storage compared to a continuous array of failures on the tape side.**
- **Since installation, there have had no virtual tape issues. We spend an entire FTE chasing down tape backup issues.**

Storage Decisions

Hosted by SearchStorage.com STORAGE

CME open systems SAN storage strategy

TIER I *DMX*	TIER II *Hitachi*	TIER III *Centera*	TIER IV *COPAN - ATA*	TIER V *Tape*
Critical and replicated data	Lower priority, non-replicated data	Archived regulatory data	ATA, virtual tape	Disaster recovery, offsite data
 EMC DMX	 Hitachi 9960	EMC Centera 	COPAN Revolution 200T 	STK PowderHorn Silo 

Storage Decisions

Hosted by SearchStorage.com STORAGE

Vendor scoop

- **COPAN is the leading MAID Vendor**
  - **Startup**
  - **Listens to customers**
- **EMC**
  - **Aggressive pricing (Really. . . .)**
  - **Good support**
- **Enterprise tape**
  - **Only 2 real competitive vendors IBM and STK**
  - **Currently in negotiations for upgrade**

Storage Decisions

Hosted by SearchStorage.com STORAGE

Recommendations / summary

- **Think long term when planning your storage strategy**
- **Identify your success factors**
- **Push for clear regulatory retention definition**
- **Evaluate new technology**
  - **Great learning experience**
  - **Evaluate vendors as well as their technology**

Storage Decisions

Hosted by SearchStorage.com STORAGE

QUESTIONS?

I can be reached at ctaylor@cme.com for any further questions.