

eSentio

THOUGHT LEADERSHIP

PARTN

BC31: A Case Study in the Battle of Storage Management

Scott Kindred, CBCP

eSentio Technologies



▶ Agenda



- Background
- The plan
 - Phase 1: Needs analysis
 - Phase 2: Design
 - Phase 3: Procurement & implementation planning
 - Phase 4: System installation & burn-in
- Results
- Next time.....
- Conclusion

▶ The Starting Point



- Direct attached storage on each server
- Limited number of “critical systems”
- Growth over 15 years to over 90 critical systems with data stored on SAN, NAS, & DASD
- Contracted “first-come, first-served” DR site

▶ Growing Pains



- Shrinking backup window with growing backup time required
- Global growth resulting in smaller window time to conduct backups
- Recovery time objective lengthening because of the amount of data
- Reliance on IT increasing pain level caused by any outage

▷ Specifics



- Grew from 2 critical systems on 2 servers to over 90 critical systems spread over 350 servers
- Recovery point objective increased from 8 hours to well over 40 hours
- Costs of contracted DR facility drastically increasing
- Use of SAN and local replication to hosted data facility became the goal

▶ The Plan



- Phase 1: Needs analysis
- Phase 2: Design
- Phase 3: Procurement & implementation planning
- Phase 4: System installation & burn-in



Phase 1: Storage Management Team



- Managers and supervisors from IT, business continuity, database management and practice support
- Project manager leadership
- Charged with development of this project, and long-term management of data storage



Phase 1:

Where are we? Where are we going?



- Identification & location of ALL data stores, or, a voyage to the last frontier!
- Review recovery time objective and recovery point objective
- Resulted in data classification system



Phase 1: Data Tier Classifications



- Tier 1: Mission critical
- Tier 2: Mission essential
- Tier 3: Deferrable



Tier 1: Mission Critical Data



- Interruption to the access of this data will result in severe financial, regulatory, or safety issues
- RTO = 4 hours
- Examples include email, accounting systems, and document management system



Tier 2: Mission Critical Data



- Occasionally accessed data important to survivability of the firm
- RTO = greater than 4 hours, but less than 24 hours
- Examples include databases in a variety of applications, primarily for litigation support



Tier 3: Deferrable Data



- Data rarely accessed, but may be required
- RTO = as needed after the first 24 hours
- Examples include some image databases, accounting historical records, annual budgeting applications

▶ Phase 1: Analysis Results



- Only Tier 1 applications and data previously restored after a disaster
- Now include all critical and mission essential data (Tier 1 and 2)
- Consolidate servers, employ virtualization where possible, store data on SAN, replicate
- Resulted in a need for 20 TB of Tier 1, and 10 TB Tier 2 storage at each location

▶ Phase 2: Design Questions



- Alternate data center location?
 - How far is far enough?
- Array or software based?
 - Chose array based
 - Array houses both Tier 1 & 2 disks
- Number of servers required?
 - 65 of various sizes
- Equipment availability?
 - Shared or single user?
- Financing
 - Lease or own?

▶ Phase 2: Array Requirements



- 20 TB usable RAID5 fiber channel disk, 300GB, 10k RPM
- 10 TB usable RAID5 SATA, or equivalent, disk, largest and fastest available at delivery date
- All software for AIX and Windows server connectivity
- 24x7x4 hour support

▶ Phase 2: Array Comparisons



	Vendor 1	Vendor 2	Vendor 3
Fiber Channel	90 each 300 GB 10k RPM 27 TB total	94 each 300 GB 10k RPM 28.8 TB total	84 each 300 GB 10k RPM 25.2 TB total
SATA or Equal	43 each 320 GB 5.4k RPM 13.7 TB PATA	58 each 250 GB 7.2k RPM 14.5 TB ATA	52 each 250 GB 7.2k RPM 13 TB SATA
Array Cost	Middle	High	Low

▶ Phase 2: Server Requirements



- Lack of internal experience with blade servers and costs eliminated use of blade servers
- Small, medium and large servers were defined to meet application requirements
- Large servers to be used as hosts for virtual servers

▶ Phase 2: Server Definitions



Small	Medium	Large
Dual processor 2 GB RAM 73 GB disk 10k RPM Single FC connectivity 4 GigE ports	Dual processor 4 GB RAM 73 GB disk 10k RPM Single FC connectivity 4 GigE ports	Quad processor 8 GB RAM 4 each 73 GB disk 10k RPM Dual FC connectivity

▶ Phase 2: Server Comparisons



Server (# required)	Vendor 1	Vendor 2	Vendor 3
Small (x36)	3 GHz (x2) 2 MB cache 2 GB RAM HBA (x1)	3 GHz (x2) 4 MB cache 4 GB RAM HBA (x1)	3 GHz (x2) 2 MB cache 4 GB RAM HBA (x1)
Medium (x24)	3 GHz (x2) 2 MB cache 4 GB RAM HBA (x1)	3 GHz (x2) 4 MB cache 4 GB RAM HBA (x1)	3 GHz (x2) 4 MB cache 4 GB RAM HBA (x1)
Large (x4)	3.3 GHz (x4) 8 MB cache 8 GB RAM HBA (x2)	2.2 GHz (x4) 4 MB cache 8 GB RAM HBA (x2)	3.16 GHz (x4) 4 MB cache 8 GB RAM HBA (x2)
Server Cost	Low	Middle	High

▶ Phase 3: Procurement & Planning



- Team expanded to include technicians
- Sub-teams created for:
 - Storage array & replication
 - Direct attached storage devices
 - Wide-area network (WAN)
 - Security
 - Server build planning & execution
- Develop a test plan

▶ Phase 3: Sub-team Responsibilities



- Storage array & replication:
 - Installation of arrays in both facilities
 - Set up fiber switches
 - Determine how to move data from current locations to new SAN
 - Determine how to replicate data
- Direct attached storage:
 - Identify current locations of all Tier 1 and 2 data
 - Determine order to move data
 - Create process to move data

▶ Phase 3: Sub-team Responsibilities



- **WAN:**
 - Develop WAN design
 - Order required data links
 - Set up and configure WAN
- **Security:**
 - Develop information security design
 - Order needed equipment
 - Set up and configure security equipment
- **Server Build**
 - Rack and configure all servers

▶ Wrap Up!



- Install all equipment
- Burn-in for one week
- Move data
- Begin replication
- Test!

▶ The Results!



- Recovery time for Tier 1 data reduced from over 40 hours to 4 hours
- Tier 2 data replicated and available for use when servers acquired
- Use of virtualization reduced number of physical servers from 90 to 65
- Use of facility within range of metro-area network saved money
- Two years to complete entire project

▶ Next Time...



- Educate users on data storage costs & need to classify data
- Storage management must constantly review data stores and ensure proper hardware is in place
- Require vendors to demonstrate equipment in YOUR environment with YOUR data
- Data tiers must be based on data value to the organization, AND access requirements

▶ Conclusion



- Go slow up front to go fast later
- New technology is a given!
- Serial attached SCSI (SAS)
- IP based SCSI (iSCSI)
- Data deduplication
- Information Lifecycle Management
- Buy what you can afford
- Storage management requires a continuous effort, not a one time shot

▶ Questions?



▶ Thank You!



Thank you for your time and interest!

Scott Kindred, CBCP

Senior Consultant

eSentio Technologies

700 12th Street NW, Suite 700

Washington, DC 20005

Email: Scott.Kindred@esentio.com

Phone: (202)725-0618