Online SAN Re-engineering and Data Migration Case Study

Danny B. Gross
Member of Technical Staff
Motorola, Inc

Atul Patel
Unix Systems Administrator
Motorola, Inc.

Aaron Dahl
Unix Systems Administrator
Motorola, Inc.

HP WORLD 2002
Conference & Expo

Data Migration Case Study
Agenda

- Description of the Project
- The Project Environment
- The Target Architecture
- Implementation Plan and Execution Report
- Engineering Challenges and Solutions
- Questions and Answers

Data Migration Case Study
Project Description

Overall Objective: Reduce costs by consolidating to two data centers

-Significant cost avoidance in storage purchase/lease over three years.
-Applications survivability or rapid recovery in major infrastructure disasters.
-Optimization of resources for Database, Storage, and Applications performance
-Net reduction in capital equipment/service costs.

Data Migration Project Goals:

–3960 GB in optimized usable storage to the framework.
–All allocated disks to be mirrored between two data centers.
–Applications survival/recovery from the loss of an entire data center within minutes.
–Protect all storage to not less than RAID 5 within a single array
–Improve applications performance when framework is fully allocated.
–Consolidate all storage onto two Storage Area Networks (SAN).
–No outage (other than momentary and recoverable) to the application.

Data Migration Case Study
Pre-Project Infrastructure

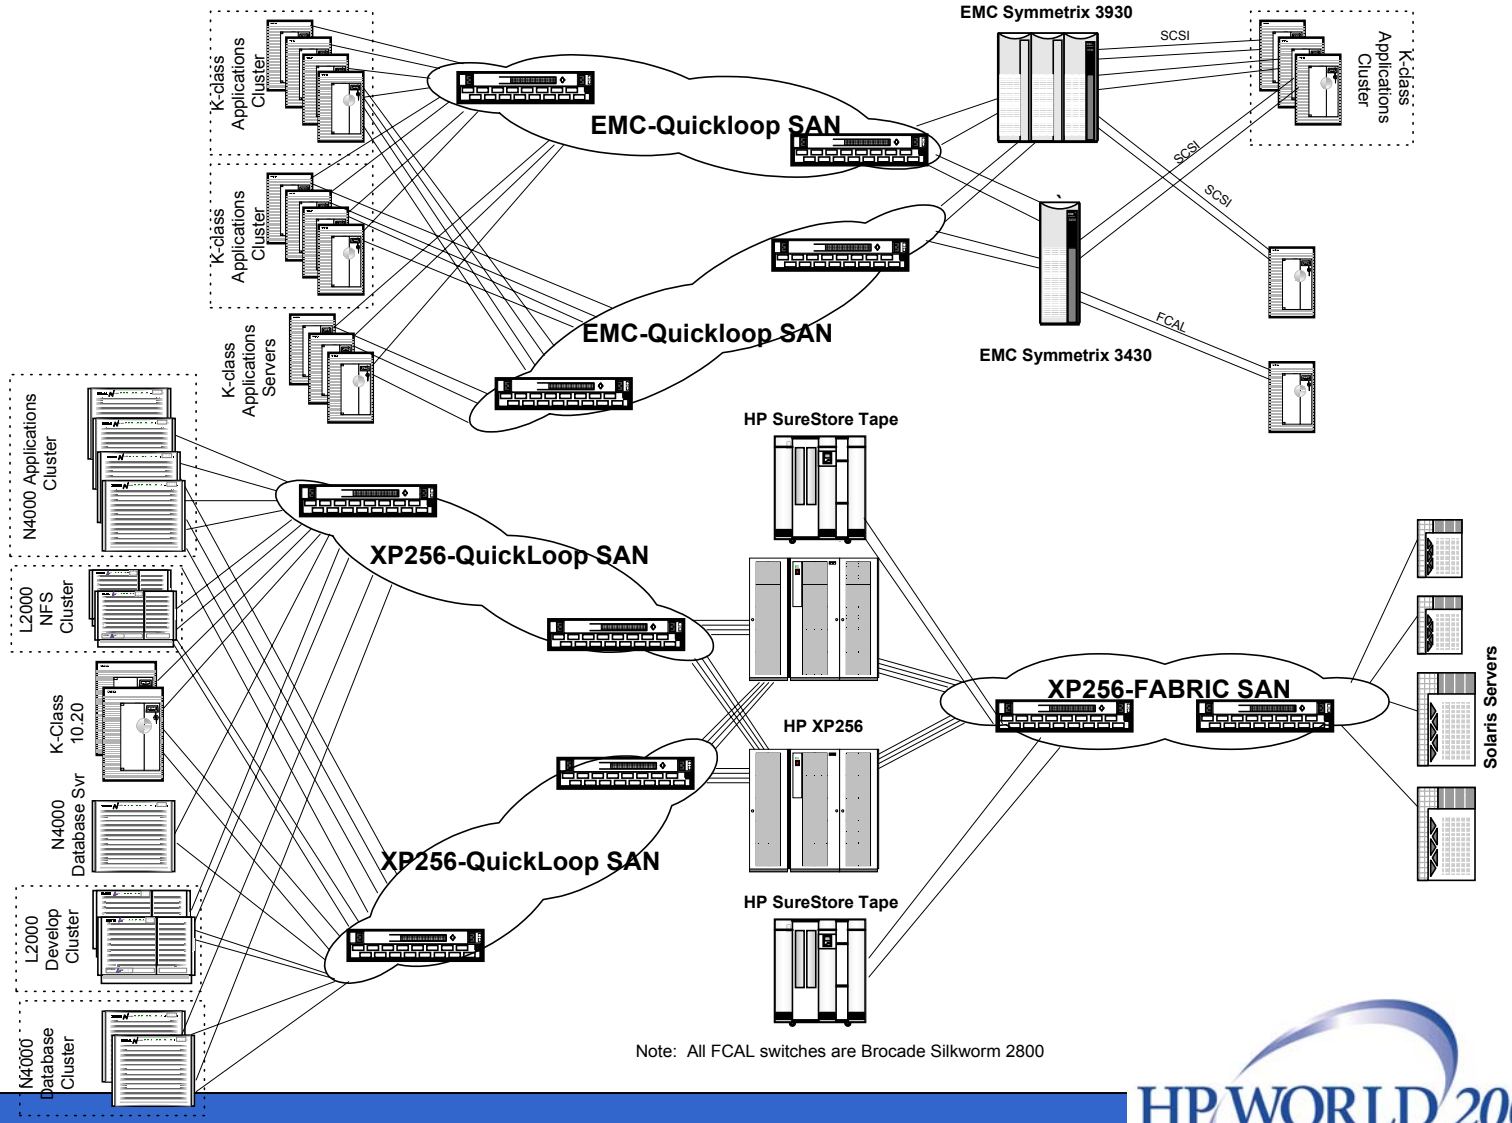

Note: All FCAL switches are Brocade Silkworm 2800

Data Migration Case Study
Pre-Project XP256 Allocation

CL1-J / CL2-J	CL1-K / CL2-K	CL1-J / CL2-J	CL1-K / CL2-K	CL1-F / CL2-F	CL1-E / CL2-E	CL1-B / CL2-B	CL1-A / CL2-A
4-8	**4-7**	**4-6**	**4-5**	**4-4**	**4-3**	**4-2**	**4-1**
3:85	3:72	3:5F	3:4C	3:39	3:26	3:13	3:00
3:86	3:73	3:60	3:4D	3:3A	3:27	3:14	3:01
3:87	3:74	3:61	3:4E	3:3B	3:28	3:15	3:02
3:88	3:75	3:62	3:4F	3:3C	3:29	3:16	3:03
3:89	3:76	3:63	3:50	3:3D	3:2A	3:17	3:04
3:8A	3:77	3:64	3:51	3:3E	3:2B	3:18	3:05
3:8B	3:78	3:65	3:52	3:3F	3:2C	3:19	3:06
3:8C	3:79	3:66	3:53	3:40	3:2D	3:1A	3:07
3:8D	3:7A	3:67	3:54	3:41	3:2E	3:1B	3:08
3:8E	3:7B	3:68	3:55	3:42	3:2F	3:1C	3:09
3:8F	3:7C	3:69	3:56	3:43	3:30	3:1D	3:0A
3:90	3:7D	3:6A	3:57	3:44	3:31	3:1E	3:0B
3:91	3:7E	3:6B	3:58	3:45	3:32	3:1F	3:0C
3:92	3:7F	3:6C	3:59	3:46	3:33	3:20	3:0D
3:93	3:80	3:6D	3:5A	3:47	3:34	3:21	3:0E
3:94	3:81	3:6E	3:5B	3:48	3:35	3:22	3:0F
3:95	3:82	3:6F	3:5C	3:49	3:36	3:23	3:10
3:96	3:83	3:70	3:5D	3:4A	3:37	3:24	3:11
3:97	3:84	3:71	3:5E	3:4B	3:38	3:25	3:12

A	**A**
C	**C**
P	**P**
3	**1**

CL1-A / CL2-A	Cl1-B / CL2-B	CL1-E / CL2-E	CL1-F / CL2-F	CL1-A / CL2-A	Cl1-B / CL2-B	CL1-E / CL2-E	CL1-F / CL2-F
2-1	**2-2**	**2-3**	**2-4**	**2-5**	**2-6**	**2-7**	**2-8**
01:00	01:0F	01:1E	01:2D	01:3C	01:4B	01:5A	01:69
01:01	01:10	01:1F	01:2E	01:3D	01:4C	01:5B	01:6A
01:02	01:11	01:20	01:2F	01:3E	01:4D	01:5C	01:6B
01:03	01:12	01:21	01:30	01:3F	01:4E	01:5D	01:6C
01:04	01:13	01:22	01:31	01:40	01:4F	01:5E	01:6D
01:05	01:14	01:23	01:32	01:41	01:50	01:5F	01:6E
01:06	01:15	01:24	01:33	01:42	01:51	01:60	01:6F
01:07	01:16	01:25	01:34	01:43	01:52	01:61	01:70
01:08	01:17	01:26	01:35	01:44	01:53	01:62	01:71
01:09	01:18	01:27	01:36	01:45	01:54	01:63	01:72
01:0A	01:19	01:28	01:37	01:46	01:55	01:64	01:73
01:0B	01:1A	01:29	01:38	01:47	01:56	01:65	01:74
01:0C	01:1B	01:2A	01:39	01:48	01:57	01:66	01:75
01:0D	01:1C	01:2B	01:3A	01:49	01:58	01:67	01:76
01:0E	01:1D	01:2C	01:3B	01:4A	01:59	01:68	01:77

CL1-K / CL2-K	CL1-J / CL2-J	CL1-K / CL2-K	CL1-F / CL2-F	CL1-E / CL2-E	CL1-B / CL2-B	CL1-A / CL2-A
	3-7	**3-6**	**3-5**	**3-4**	**3-3**	**3-2**	**3-1**
A	2:72	2:5F	2:4C	2:39	2:26	2:13	2:00
	2:73	2:60	2:4D	2:3A	2:27	2:14	2:01
	2:74	2:61	2:4E	2:3B	2:28	2:15	2:02
	2:75	2:62	2:4F	2:3C	2:29	2:16	2:03
	2:76	2:63	2:50	2:3D	2:2A	2:17	2:04
	2:77	2:64	2:51	2:3E	2:2B	2:18	2:05
	2:78	2:65	2:52	2:3F	2:2C	2:19	2:06
	2:79	2:66	2:53	2:40	2:2D	2:1A	2:07
	2:7A	2:67	2:54	2:41	2:2E	2:1B	2:08
	2:7B	2:68	2:55	2:42	2:2F	2:1C	2:09
	2:7C	2:69	2:56	2:43	2:30	2:1D	2:0A
	2:7D	2:6A	2:57	2:44	2:31	2:1e	2:0B
	2:7E	2:6B	2:58	2:45	2:32	2:1f	2:0C
	2:7F	2:6C	2:59	2:46	2:33	2:20	2:0D
	2:80	2:6D	2:5A	2:47	2:34	2:21	2:0E
	2:81	2:6E	2:5B	2:48	2:35	2:22	2:0F
	2:82	2:6F	2:5C	2:49	2:36	2:23	2:10
	2:83	2:70	2:5D	2:4A	2:37	2:24	2:11
	2:84	2:71	2:5E*	2:4B	2:38	2:25	2:12

A	**A**
C	**C**
P	**P**
2	**0**

CL1-A / CL2-A	Cl1-B / CL2-B	CL1-E / CL2-E	CL1-F / CL2-F	CL1-A / CL2-A	Cl1-B / CL2-B	CL1-E / CL2-E	
1-1	**1-2**	**1-3**	**1-4**	**1-5**	**1-6**	**1-7**	
00:00	00:0F	00:1E	00:2D	00:3C	00:4B	00:5A	
00:01	00:10	00:1F	00:2E	00:3D	00:4C	00:5B	
00:02	00:11	00:20	00:2F	00:3E	00:4D	00:5C	
00:03	00:12	00:21	00:30	00:3F	00:4E	00:5D	S
00:04	00:13	00:22	00:31	00:40	00:4F	00:5E	
00:05	00:14	00:23	00:32	00:41	00:50	00:5F	P
00:06	00:15	00:24	00:33	00:42	00:51	00:60	
00:07	00:16	00:25	00:34	00:43	00:52	00:61	A
00:08	00:17	00:26	00:35	00:44	00:53	00:62	
00:09	00:18	00:27	00:36	00:45	00:54	00:63	R
00:0A	00:19	00:28	00:37	00:46	00:55	00:64	
00:0B	00:1A	00:29	00:38	00:47	00:56	00:65	E
00:0C	00:1B	00:2A	00:39	00:48	00:57	00:66	
00:0D	00:1C	00:2B	00:3A	00:49	00:58	00:67	
00:0E	00:1D	00:2C	00:3B	00:4A	00:59	00:68	

Data Migration Case Study
Target Infrastructure

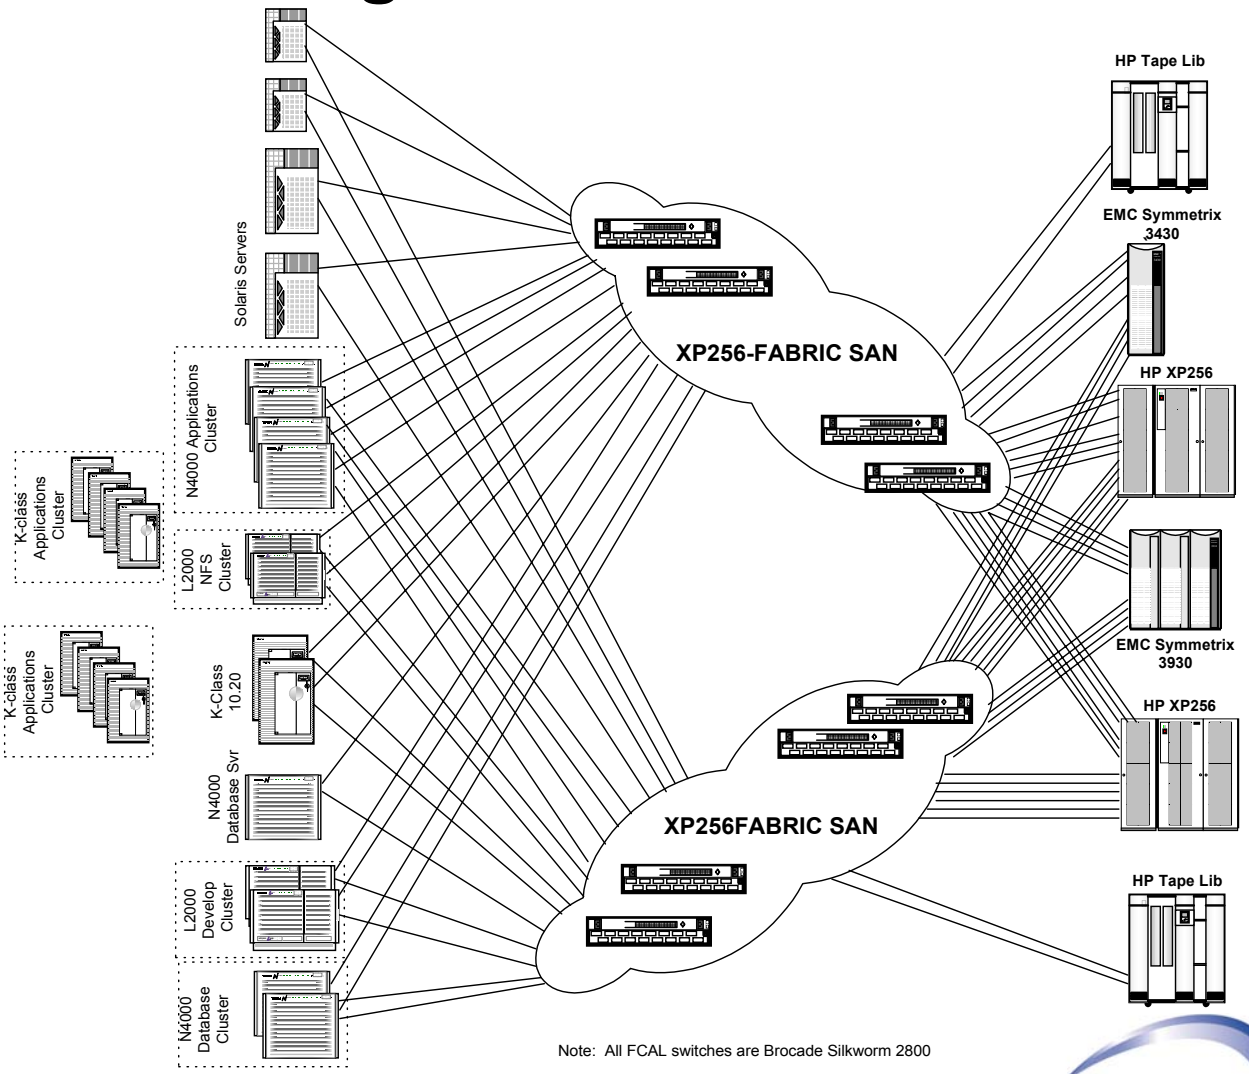

Data Migration Case Study
Target XP256 Allocation

CL1-J / CL2-J	CL1-E / CL2-E	CL1-F / CL2-F	CL1-J / CL2-J	CL1-A / CL2-A	CI1-B / CL2-B	CL1-E / CL2-E	CL1-F / CL2-F			CI1-B / CL2-B	CL1-A / CL2-A	CL1-J / CL2-J	CL1-F / CL2-F	CL1-E / CL2-E	CI1-B / CL2-B	CL1-A / CL2-A	CL1-F / CL2-F
4-8	**4-7**	**4-6**	**4-5**	**4-4**	**4-3**	**4-2**	**4-1**			**2-1**	**2-2**	**2-3**	**2-4**	**2-5**	**2-6**	**2-7**	**2-8**
3:75	3:62	3:4F	3:3C	3:2D	3:1E	3:0F	3:00			1:00	1:0F	1:1E	1:2D	1:3C	1:4F	1:62	1:75
3:76	3:63	3:50	3:3D	3:2E	3:1F	3:10	3:01			1:01	1:10	1:1F	1:2E	1:3D	1:50	1:63	1:76
3:77	3:64	3:51	3:3E	3:2F	3:20	3:11	3:02			1:02	1:11	1:20	1:2F	1:3E	1:51	1:64	1:77
3:78	3:65	3:52	3:3F	3:30	3:21	3:12	3:03			1:03	1:12	1:21	1:30	1:3F	1:52	1:65	1:78
3:79	3:66	3:53	3:40	3:31	3:22	3:13	3:04	**A**	**A**	1:04	1:13	1:22	1:31	1:40	1:53	1:66	1:79
3:7A	3:67	3:54	3:41	3:32	3:23	3:14	3:05	**C**	**C**	1:05	1:14	1:23	1:32	1:41	1:54	1:67	1:7A
3:7B	3:68	3:55	3:42	3:33	3:24	3:15	3:06	**P**	**P**	1:06	1:15	1:24	1:33	1:42	1:55	1:68	1:7B
3:7C	3:69	3:56	3:43	3:34	3:25	3:16	3:07			1:07	1:16	1:25	1:34	1:43	1:56	1:69	1:7C
3:7D	3:6A	3:57	3:44	3:35	3:26	3:17	3:08	**3**	**1**	1:08	1:17	1:26	1:35	1:44	1:57	1:6A	1:7D
3:7E	3:6B	3:58	3:45	3:36	3:27	3:18	3:09			1:09	1:18	1:27	1:36	1:45	1:58	1:6B	1:7E
3:7F	3:6C	3:59	3:46	3:37	3:28	3:19	3:0A			1:0A	1:19	1:28	1:37	1:46	1:59	1:6C	1:7F
3:80	3:6D	3:5A	3:47	3:38	3:29	3:1A	3:0B			1:0B	1:1A	1:29	1:38	1:47	1:5A	1:6D	1:80
3:81	3:6E	3:5B	3:48	3:39	3:2A	3:1B	3:0C			1:0C	1:1B	1:2A	1:39	1:48	1:5B	1:6E	1:81
3:82	3:6F	3:5C	3:49	3:3A	3:2B	3:1C	2:0D			1:0D	1:1C	1:2B	1:3A	1:49	1:5C	1:6F	1:82
3:83	3:70	3:5D	3:4A	3:3B	3:2C	3:1D	3:0E			1:0E	1:1D	1:2C	1:3B	1:4A	1:5D	1:70	1:83
3:84	3:71	3:5E	3:4B											1:4B	1:5E	1:71	1:84
3:85	3:72	3:5F	3:4C											1:4C	1:5F	1:72	1:85
3:86	3:73	3:60	3:4D											1:4D	1:60	1:73	1:86
3:87	3:74	3:61	3:4E											1:4E	1:61	1:74	1:87

	CI1-B / CL2-B	CL1-E / CL2-E	CL1-F / CL2-F	CL1-J / CL2-J	CL1-A / CL2-A	CI1-B / CL2-B	CL1-E / CL2-E			CL1-A / CL2-A	CL1-J / CL2-J	CL1-F / CL2-F	CL1-E / CL2-E	CI1-B / CL2-B	CL1-A / CL2-A	CL1-J / CL2-J	
	3-7	**3-6**	**3-5**	**3-4**	**3-3**	**3-2**	**3-1**			**1-1**	**1-2**	**1-3**	**1-4**	**1-5**	**1-6**	**1-7**	
	2:62	2:4F	2:3C	2:2D	2:1E	2:0F	2:00			0:00	0:0F	0:1E	0:2D	0:3C	0:4F	0:62	
	2:63	2:50	2:3D	2:2E	2:1F	2:10	2:01			00:01	0:10	0:1F	0:2E	0:3D	0:50	0:63	
S	2:64	2:51	2:3E	2:2F	2:20	2:11	2:02			00:02	0:11	0:20	0:2F	0:3E	0:51	0:64	S
	2:65	2:52	2:3F	2:30	2:21	2:12	2:03			00:03	0:12	0:21	0:30	0:3F	0:52	0:65	
P	2:66	2:53	2:40	2:31	2:22	2:13	2:04	**A**	**A**	00:04	0:13	0:22	0:31	0:40	0:53	0:66	P
	2:67	2:54	2:41	2:32	2:23	2:14	2:05	**C**	**C**	00:05	0:14	0:23	0:32	0:41	0:54	0:67	
A	2:68	2:55	2:42	2:33	2:24	2:15	2:06	**P**	**P**	00:06	0:15	0:24	0:33	0:42	0:55	0:68	A
	2:69	2:56	2:43	2:34	2:25	2:16	2:07			00:07	0:16	0:25	0:34	0:43	0:56	0:69	
R	2:6A	2:57	2:44	2:35	2:26	2:17	2:08	**2**	**0**	00:08	0:17	0:26	0:35	0:44	0:57	0:6A	R
	2:6B	2:58	2:45	2:36	2:27	2:18	2:09			00:09	0:18	0:27	0:36	0:45	0:58	0:6B	
E	2:6C	2:59	2:46	2:37	2:28	2:19	2:0A			00:0A	0:19	0:28	0:37	0:46	0:59	0:6C	E
	2:6D	2:5A	2:47	2:38	2:29	2:1A	2:0B			00:0B	0:1A	0:29	0:38	0:47	0:5A	0:6D	
	2:6E	2:5B	2:48	2:39	2:2A	2:1B	2:0C			00:0C	0:1B	0:2A	0:39	0:48	0:5B	0:6E	
	2:6F	2:5C	2:49	2:3A	2:2B	2:1C	2:0D			00:0D	0:1C	0:2B	0:3A	0:49	0:5C	0:6F	
	2:70	2:5D	2:4A	2:3B	2:2C	2:1D	2:0E			00:0E	0:1D	0:2C	0:3B	0:4A	0:5D	0:70	
	2:71	2:5E	2:4B											0:4B	0:5E	0:71	
	2:72	2:5F	2:4C											0:4C	0:5F	0:72	
	2:73	2:60	2:4D											0:4D	0:60	0:73	
	2:74	2:61	2:4E											0:4E	0:61	0:74	

Data Migration Case Study
Project Plan/Execution Report

Phase II – Isolate and Re-engineer 1st Datacenter

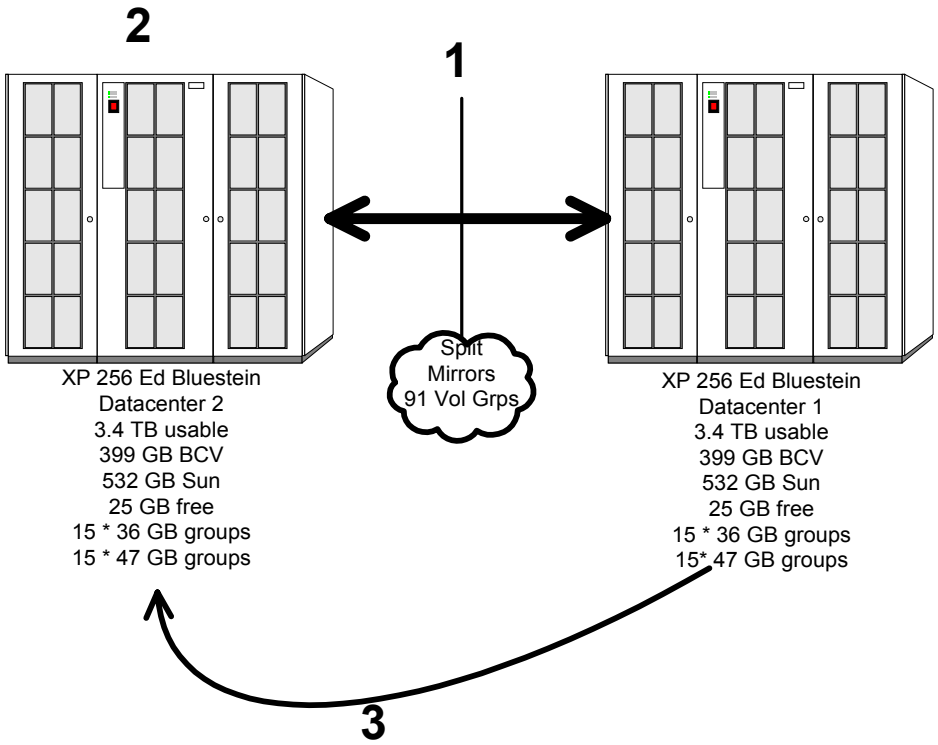

1 - Split Mirrors/reduce pvlinks+mirror primaries	Planned - 8 hrs, Actual - 18 hrs
2 - Re-engineer XP256 and Switch Framework	Planned - 2 d, Actual - 4 d
3 - Re-mirror volumes into desired locations	Planned - 10 d, Actual - 3 d

Data Migration Case Study
Project Plan/Execution Report

Phase III – Isolate and Re-engineer 2nd Datacenter

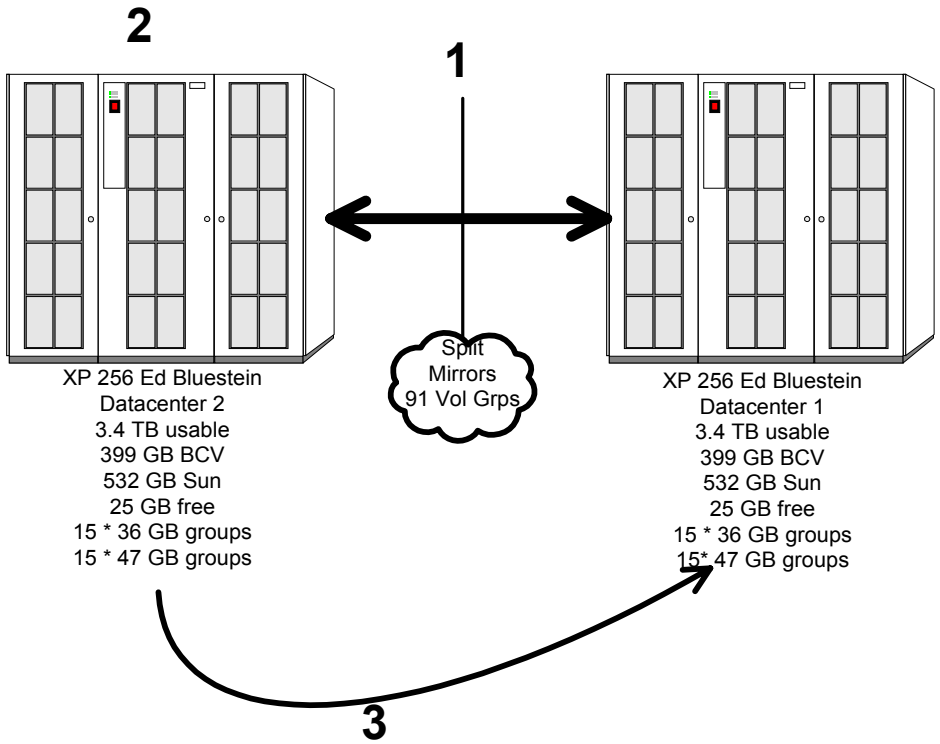

1 - Split Mirrors/reduce pvlinks+mirror primaries	Planned - 8 hrs, Actual - 4hrs
2 - Re-engineer XP256 and Switch Framework	Planned - 2 d, Actual - 2 d
3 - Re-mirror volumes into desired locations	Planned - 10 d, Actual - 3d

Data Migration Case Study
Engineering Challenges

- How to maintain ServiceGuard viability during massive reorganization
- How to convert QuickLoop to Fabric Mode with no outage.
- How to resolve logical volumes allocated between two arrays.
- How to manage and plan for rapidly changing infrastructure.
- How to recover HP and Veritas Volume Managers without rebooting.
- How to rapidly re-import modified volume groups within the SAN.

Data Migration Case Study
ServiceGuard Viability

Issues:

- Modification of clusterlock disk requires cluster to be taken down
- Clusterlock disk would be lost during mirror split operation
- Clusterlock disk would be reorganized on array

Solution:

- Use cmcld reports to syslog to verify loss and recovery of clusterlock.
- Rebuild device file for configured clusterlock to match minor number for target location.
- Extend the rebuilt device back into the clusterlock volume group.
- Use “cminitlock” to re-initialize the clusterlock.

Data Migration Case Study
QuickLoop-Fabric Conversion

Issues:

- Brocade 2800 Switches required firmware upgrade for Fabric Mode.
- No host reboot was possible.

Solution:

- Install Fabric-capable Adapters in hosts.
- Conduct mirror splits, remove primary/pvlink and remaining pvlinks
- Re-engineer arrays, install firmware upgrade, and re-build Fabric.
- Re-enable the host ports. Host adapters will re-negotiate correct topology.

Data Migration Case Study
Logical Volume Split Issues

Issues:

- Unmirrored physical volume resides on two arrays

Solution:

- Identify unallocated extents in volume group.
- Use pvmove to re-locate extents to appropriate array.

Data Migration Case Study
Managing Rapidly Changing Infrastructure

Issues:

- Several systems required redeployment to SAN infrastructure.
- Several applications needed additional disk space
- Project needed 22 days to complete

Solution: Establish management-supported rules for change:

- No changes to clusterlock volume groups until completion of project
- No changes that risk a loss of quorum to a cluster
- No changes that prevent our ability to manipulate the FCAL infrastructure
- New allocations of LUNs can be made after Phase II is complete. Risk is: No mirror.

Data Migration Case Study
LVM and Veritas Recovery

HPUX:

- "vgreduce" pvlinks, mirror disks from volume groups.
- "rmsf" to remove device files.
- "ioscan –fnC disk" to re-scan infrastructure.
- "insf –e" to re-install special files.
- "vgextend" and "lvextend" to rebuild mirrors and pvlinks.

Solaris/Veritas:

- "vxassist remove mirror <vol> <diskname1> <diskname2> …" to remove mirrors
- "vxdg –g <disk group> rmdisk <disk>" to reduce disks.
- "port disable <portnum>" (Brocade) to disable port.
- "port enable <portnum>" (Brocade) to enable port after SAN work done.
- "drvconfig;disks" to rescan infrastructure/rebuild device files.
- "format –f <labelfile> <disk>" to apply label to disks.
- "vxdctl enable" to force Veritas to rescan disks.
- "vxdisksetup –i <disk>" to create a Veritas disk
- "vxdg –g <disk group> adddisk <disk name>" to extend a disk into a disk group
- "vxassist –g <disk group> mirror <vol> <disk name1> <disk name 2> …" to remirror

Data Migration Case Study
Managing and Importing Volume Groups

Issues:

- 624 allocated physical volumes in 90 volume groups
- 18 different systems

Solution:

- Turn off SecureManagerXP to simplify and speed disk management.
- Create a single file that holds critical data:

VG Name, VG Minor Number, PVG assigned, LUN, Information field

EX: vg_admin,0x160000,vg_admingrp5,0:3f,system1
vg_admin,0x160000,vg_admingrp5,1:3f,system1

- Manage all changes via this file and share with all SAN-hosts.
- Expect every LUN to have mirror in identical location on other array.
- Set kernel param “maxvgs” high (default= 10) on all SAN-hosts
- Use “xpinfo” or “inquiry256” to match LUNs and device files.
- Use scripts to parse “vgimport” command.

Online SAN Re-engineering and Data Migration Case Study

Questions?

Danny B. Gross
Member of Technical Staff
Motorola, Inc
6501 William Cannon West
Austin, Texas 78735
danny.gross@motorola.com
(512) 895-4825

Atul Patel
Unix Systems Administrator
Motorola, Inc.
3501 Ed Bluestein Blvd.
Austin, Texas 78721
Atul.C.Patel@motorola.com
(512) 933-5472

Aaron Dahl
Unix Systems Administrator
Motorola, Inc.
3501 Ed Bluestein Blvd.
Austin, Texas 78721
Aaron.Dahl@motorola.com
(512) 933-6310