HP's AutoRAID and the XP256 on the HP3000

Chris Wong
Cerius Technology Group

HP World 2000

Presentation #017

Agenda

- What is a LUN?
- What is the AutoRAID?
- What is the XP256?
- Why would I want to use one?
- How would I use one?
- What about LDEV1?
- What about beyond LDEV1?
- CA & BC with XP256

What is a LUN?

- Logical Unit Number
- LUN looks like single disk to host
- LUN is actually across multiple disks

Disks and their relationship to MPE

HASS, internal, etc.. - 1 PDEV per disk
AutoRAID - 1 PDEV per LUN - LUN is across ALL disks
XP256 - 1 PDEV per LDEV (LUN) - LDEV is across disks in array group

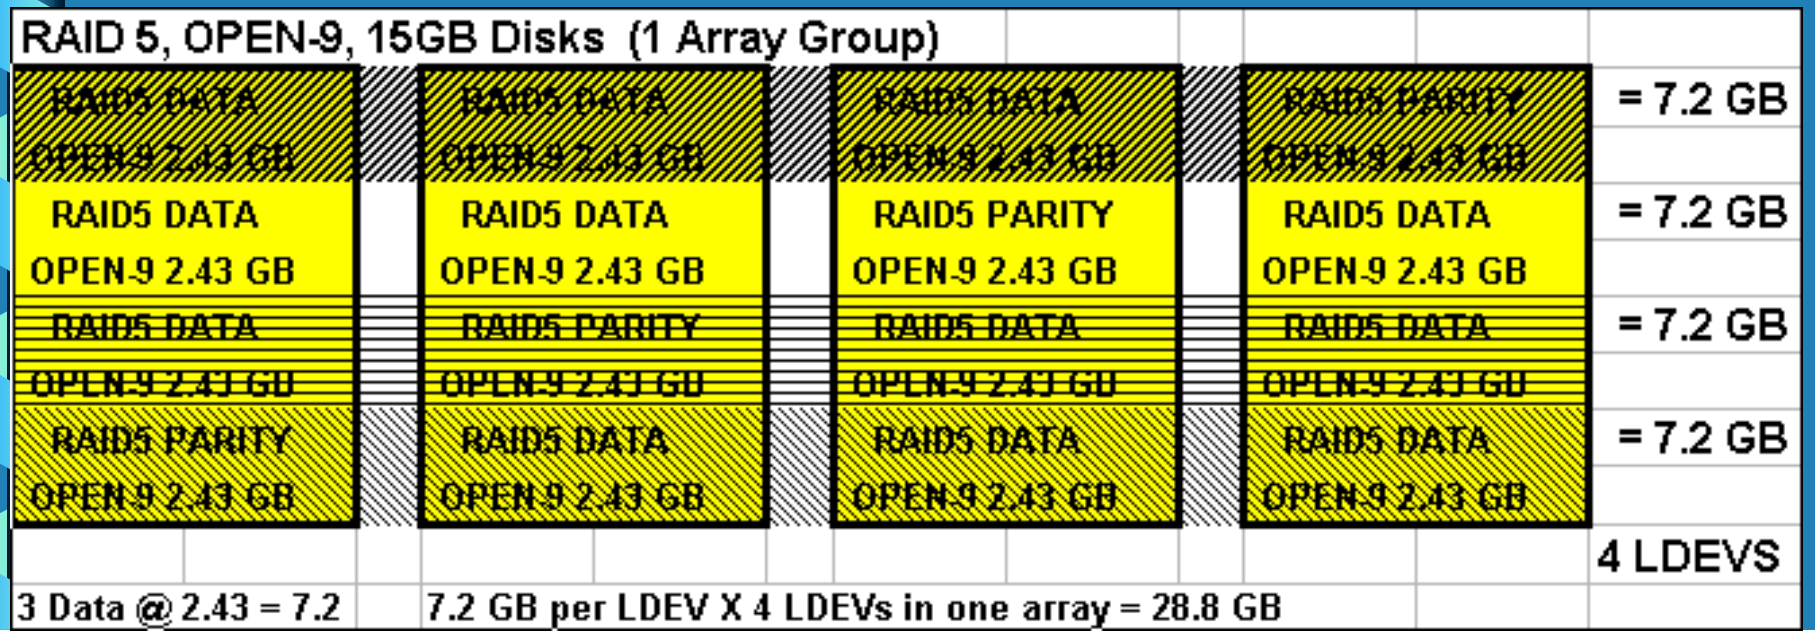

What is the AutoRAID?

- Name: HP SureStore E Disk Array 12H
- Automagically manages data between RAID 1 and RAID 5
  - HSM: Hierarchical Storage Management - automatic
    - New and active data: RAID 1
    - Old and inactive data: RAID 5*
    - Striping & file placment - automatic
- Component Redundancy
- * Only after array starts to fill

AutoRAID

- Data spread across up to 12 drives
- Config up to 8 LDEVs
- RAID 1 and RAID 5 - not dedicated

What is the XP256?

- Cache-centric system
- Redundant, Super HA
- Large & expandable
- Expensive
- Fast & cool

XP256

- Dedicated SCSI bus
- Can’t daisy chain XP256s
- Max 15 LDEVs per SCSI bus
- Max 8 LDEVs for best performance

CHIP Client Host Interface Processor
MPE/iX doesn't support redundant controllers

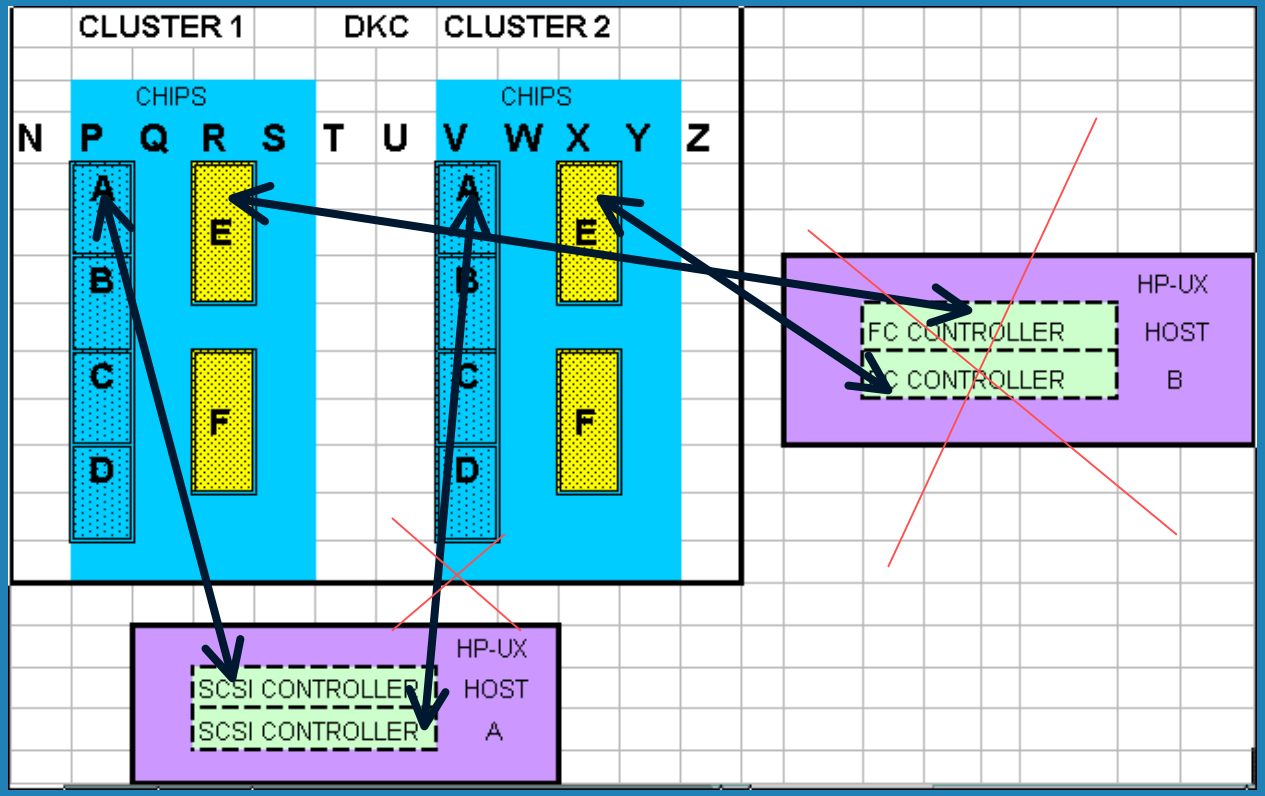

CL1A to c3 on HostA

CL2V to c5 on HostA

CL1E to c7 on HostB

CL2E to c9 on HostB

Host A: 2 paths to XP256 - c3 to Cluster1, slot P, port A and c5 to Cluster2, slot V, port A

Assign LDEV to CHIP scsi

Hi! I'm LDEV 0:01

RAID5 DATA OPEN-9 2.43 GB	RAID5 DATA OPEN-9 2.43 GB	RAID5 PARITY OPEN-9 2.43 GB	RAID5 DATA OPEN-9 2.43 GB	= 7.2 GB LDEV
CONFIGURE HOST ACCESS TO THE LDEV:		Pick CHIP that you want to access the LDEV from.		

Access LDEV #0:01

From CL1-A

Using SCSI ID 1

and LUN 00

And access this same LDEV

From CL2-A

Using SCSI ID 9

and LUN 00

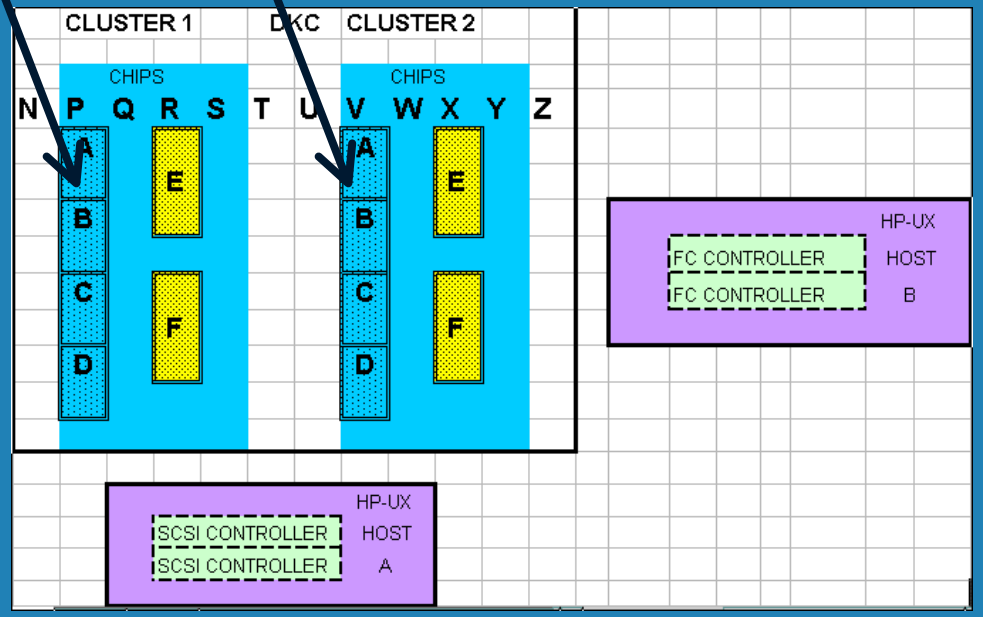

Why? - High Availability

- Mirroring Options
  - Software Mirroring is not supported on LDEV1 or any USER VOLUME
  - Hardware Mirroring is only supported method for LDEV1
- Uptime
  - Add disks on-line

Why? Protect Investment

- AutoRAID could be used for HP-UX or NT
- XP256 can be used for UNIX, NT, NovelI, Linux, MVS/VM and MPE or all at the same time!
- Welcome to the SAN (Storage Area Network)

Why? Capacity Planning

- No downtime for increasing capacity
- Notification when capacity should be increased

Why? Performance Management

- AutoRAID - Completley self tuning (HSM)
  - No or small IT department
- XP256 - Tools to assist
- AutoRAID - don't use if more than 276 I/Os per second - then the single controller becomes an issue
- Arraymgr - High Performance Mode
  - array -J HighPerformance [array-id]

System Requirements (6/16/00)

- AutoRAID
  - MPE/iX 6.5 or MPE/iX 6.0 with Express 2
  - MPE/iX 5.5 PP 7 + MPEKXU3, ARMKXW5, SYLLX17
- XP256
  - MPE/iX 6.5 or MPE/iX 6.0 with Express 2
  - MPE/iX 5.5 PP7 + MPEKXT8

AutoRAID & LDEV 1

- Add additional card(s) if needed
- Configure lun0 on AutoRAID as 4 GB
- Connect to controller X
- ISL> ODE mapper run
  - Get path for SCSI card and LUN0
- sysgen
  - Add card (if added)
  - Modify ldev1 to new PATH and new ID (hpdarray)
  - ho | keep | tape

LDEV1 - continued

- Boot from prim? N
- Boot from alt? Y
- Interact with IPL? Y
- ISL> PRIMPATH 40.1.0 (whatever for LUN0)
- ISL> INSTALL
- Boot from prim? Y
- Interact with IPL? Y
- ISL> START NORECOVERY NOSYSSTART

XP256 & LDEV1

- LUN0 must be created on RAID1 and must be 4 GB (Requires ACP & array group)
- All other steps are the same (even the ID #)

Next steps - above LDEV1

- Private Volume Sets
- HPPXUDC set
- LUNs 0,1,2,3,4,5,6
  - 0 = 40.1.0 (LDEV1, LUN0)
  - 1 = 40.1.1, 2=40.1.2, 3=40.1.3, etc.
  - ad 32 id=hpdarray path=40.1.2
  - VOLUTIIL

LDEVs and XP256

- Open level is important :
  - Limited number of LDEVs on one controller
  - Example: Open-3 uses 12 LDEVs while for virtually the same space Open-8 uses 4
    - 2.4 GB vs. 7.3 GB
  - Write Queue Depth not an issue on MPE

RAID1			RAID5		
OPEN-3	OPEN-8	OPEN-9	OPEN-3	OPEN-8	OPEN-9
28/12	27/4	28/4	41/18	41/6	41/6
Usable Space in GB / LDEVS					

LDEVs and XP256

- Use OPEN 3 / RAID 1 for LDEV1
- Use OPEN 3 / RAID 1 for any other system_volume_set_members
- Use OPEN 3/ RAID 1 for Master Volumes of User Volumes
- Use OPEN 9 / RAID 5 for User Volumes
- SCSI Host Mode: 08

XP256, Business Copy & MPE

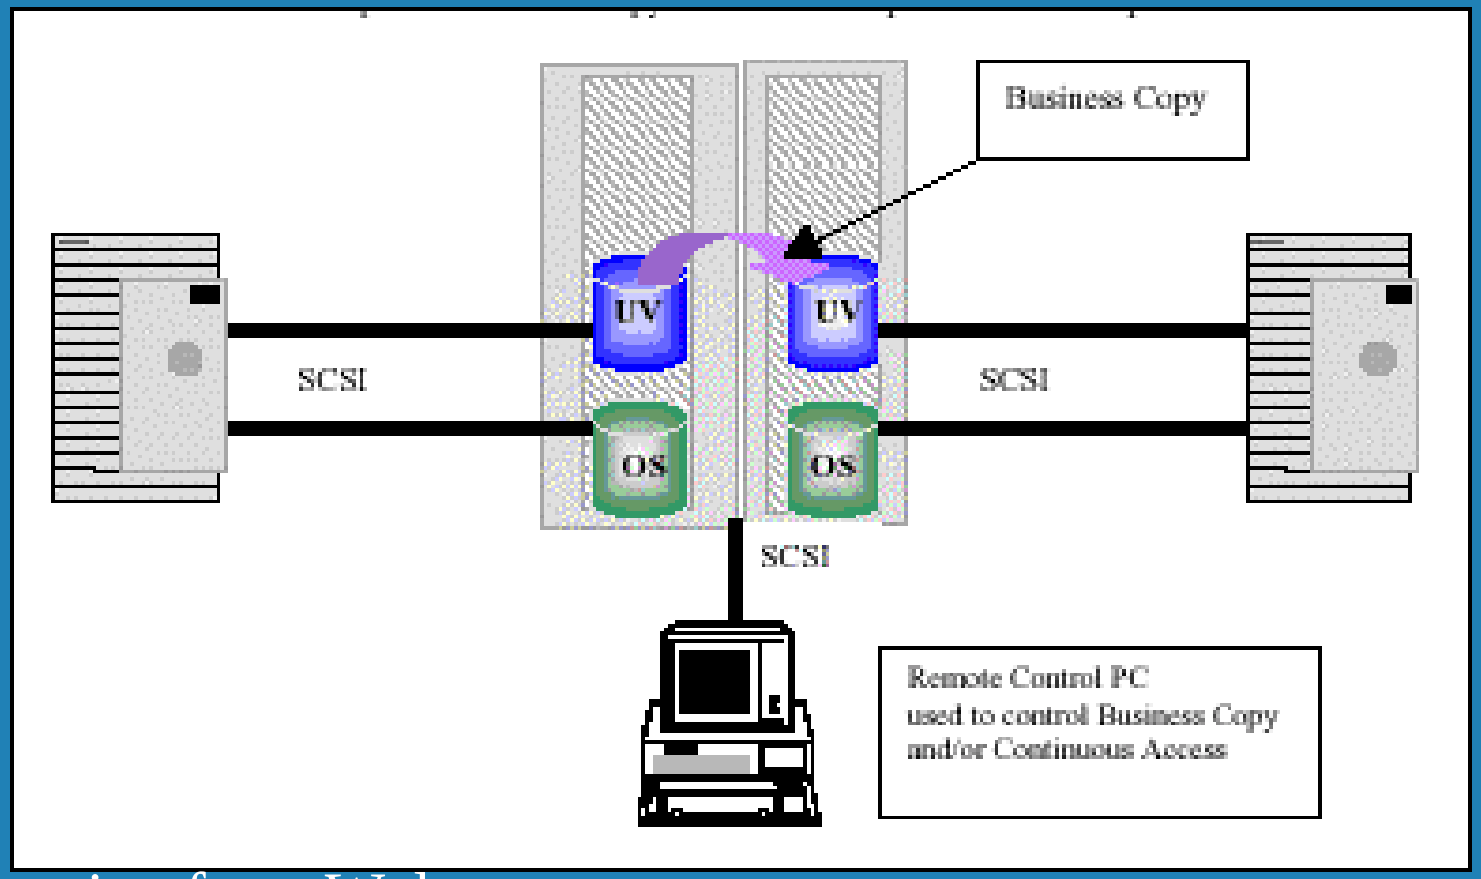

Drawing from Walter McCullough's White Paper

XP256, Continuous Acces & MPE

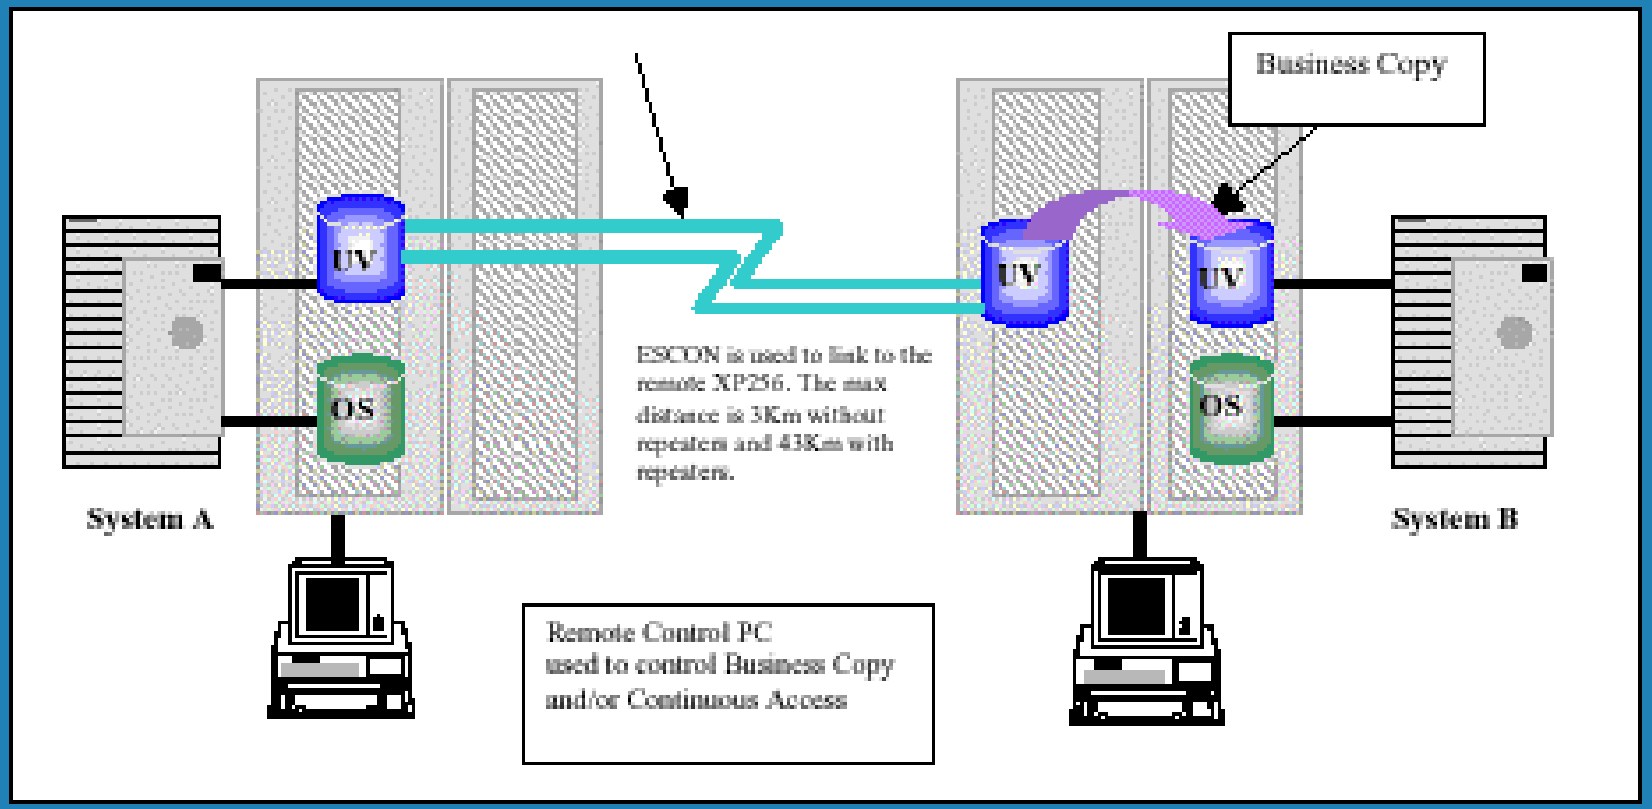

Drawing from Walter McCullough's White Paper

BC or CA? Get this paper

- White Paper by Walter McCullough (HP)
- December 1999

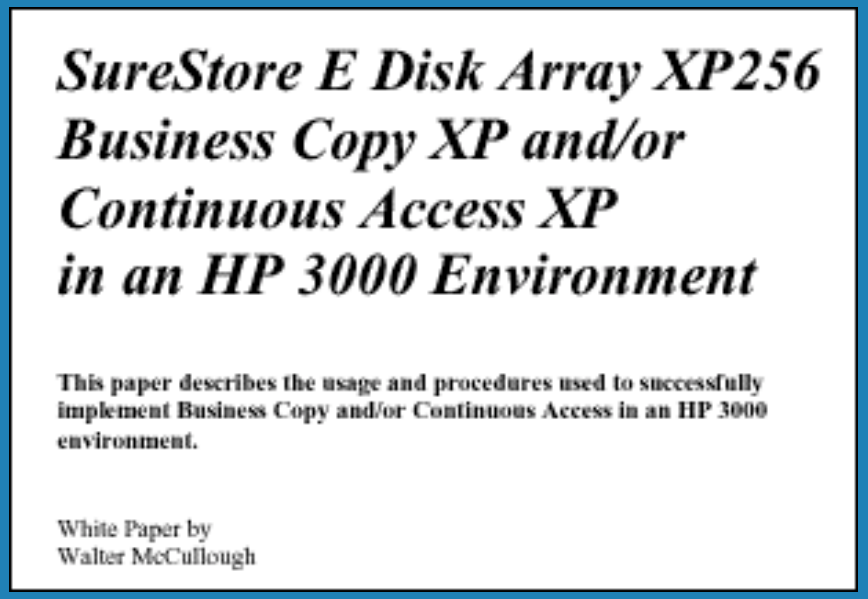

SureStore E Disk Array XP256
Business Copy XP and/or
Continuous Access XP
in an HP 3000 Environment

This paper describes the usage and procedures used to successfully implement Business Copy and/or Continuous Access in an HP 3000 environment.

White Paper by
Walter McCullough