

OpenVMS Filesystem Updates



P Muralidhar Kini
OpenVMS Engineering

© Copyright 2011 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice. Confidentiality label goes here



AGENDA

TOPICS

What's new in File System post V8.4 ?

- XFC Updates
 - ❑ Support for 4GB GH regions

- XQP Updates
 - ❑ Directory validation patch
 - ❑ Guard LBN 0 allocation

- RMS Updates
 - ❑ MBC Enhancement
 - ❑ SYMLINKS Update



XFC UPDATE

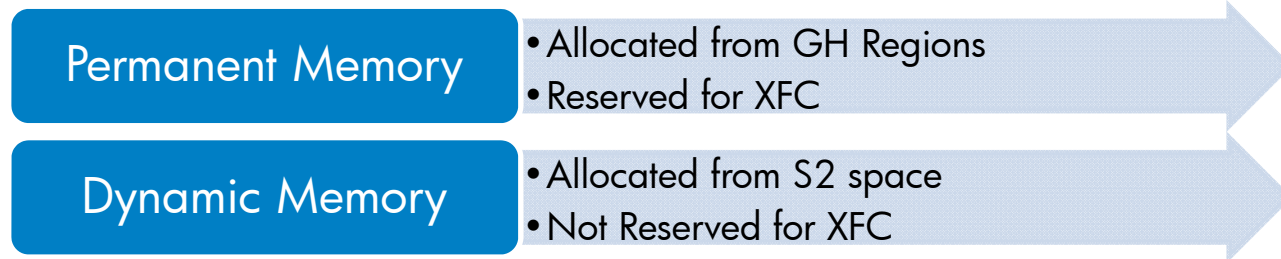
SUPPORT FOR 4GB GH REGIONS



SUPPORT FOR 4GB GH REGION

CURRENT LIMITATION

➤ XFC Memory



➤ XFC supports max GH region size of 4MB – ALPHA/INTEGRITY

- ❑ Request for 8GB XFC Permanent memory gets satisfied as 2048 4MB pages



➤ Integrity Servers support GH regions of sizes up to 4GB

- ❑ Request for 8GB can get satisfied as 2 4GB pages



SUPPORT FOR 4GB GH REGION

SOLUTION

- XFC enhanced to support GH regions of size up to 4GB on Integrity
 - ❑ Request for 8GB XFC Permanent memory can get satisfied as 2 4GB pages
 - ❑ Less TLB entries means less TLB misses and hence faster response times



- Affected images
 - ❑ XFC (SYS\$XFCACHE.EXE)
 - ❑ SYS (SYSTEM_PRIMITIVES.EXE)

- Patch containing this fix
 - ❑ VMS84I_UPDATE-V0600

NEW

XFC now supports GH region sizes as per the ALPHA/INTEGRITY architecture

XQP UPDATES

1) DIRECTORY VALIDATION PATCH



DIRECTORY VALIDATION PATCH

WHY?

- Problem reports
 - Directory corruption problem
 - Paged pool corruption problem
- Analysis results
 - Inadequate consistency checks in XQP for early problem detection
 - Directory corruption problem not yet root caused
- Can XQP take any proactive measure?
 - Add consistency checks for early problem detection
 - Mechanisms to report the problem to the user
 - Provide controls to enable/disable the consistency checks

XQP made more robust by adding basic consistency checks for early problem detection



DIRECTORY VALIDATION PATCH

TARGET – DIRECTORY CORRUPTION PROBLEM

Directory corruption problem involves failure to access file(s) within a directory

Known symptoms –

1) Missing directory entry

- Files that existed in the directory go missing.
- These files aren't accessible and not listed by any form of "DIR" or "DUMP" command

2) Stale directory entry

- Files that were successfully deleted before, reappear in the directory.
- These files aren't accessible but are listed by "DIR" and "DUMP" commands.

3) Filenames out of order

- Files that aren't accessible are listed by "DIR" but not by "DIR <FILESPEC>" command
- Files in the directory are not in alphabetically sorted order



DIRECTORY VALIDATION PATCH

FEATURES

- Additional Consistency checks added, no tracing

- Checks controlled by Dynamic SYSGEN parameter XQPCTLD7 (Default value 0)
 - ❖ Bit 0 -> Controls filling/checking of known pattern in unused blocks of directory
 - Disabled by default

 - ❖ Bit 1 -> Controls action taken (Crash/OPCOM message) when a inconsistency is detected
 - OPCOM message setting by default

 - ❖ Bit 2 -> Controls using internal counter for directory entry validation before/after every file create/delete operation
 - Feature enabled by default



DIRECTORY VALIDATION PATCH

SETTINGS

XQPCTLD7	Directory Entry Validation	Action on Inconsistency	Check for Pattern	
000	ENABLED	OPCOM	DISABLED	Default Mode
001	ENABLED	OPCOM	ENABLED	Min Mode
010	ENABLED	CRASH	DISABLED	
011	ENABLED	CRASH	ENABLED	Max Mode
100	DISABLED	NA	DISABLED	None Mode
101	DISABLED	OPCOM	ENABLED	
110	DISABLED	NA	DISABLED	
111	DISABLED	CRASH	ENABLED	



DIRECTORY VALIDATION PATCH

RELEASE/PERFORMANCE IMPACT

➤ Patches Available

- ❑ VMS83A_F11X-V0200, VMS831H1I_F11X-V0200
- ❑ VMS84I_UPDATE-V0600, VMS84A_UPDATE-V0600



NEW

➤ To disable all new functionality, use none mode

- ❑ Set XQPCTLD7 to 4(100) or 6(110)

➤ Performance Impact

- ❑ Performance tests involved heavy file create/delete operation
- ❑ Features enabled by default has negligible performance impact <1%
- ❑ Overall performance impact is minimal <3%

Once this patch is installed, set XQPCTLD7 based on the functionality desired



XQP UPDATE

2) GUARD LBN0 ALLOCATION



GUARD LBN0 ALLOCATION

WHY THIS FEATURE?

- Problem reports
 - Disk corruption, LBN 0 corrupted
 - LBN 0 contains the boot block of bootable disk
- Analysis results
 - XQP never allocates LBN0, LBN0 always allocated to INDEXF.SYS/GPT.SYS file
 - Some LBN 0 corruption problems, root caused to be with 3rd party application
- Can XQP take any proactive measure?
 - Make XQP more robust to handle 3rd party application errors while updating file headers LBN
 - Handle XQP data structure tampering which can result in adding LBN 0 to free LBN list

Unexpected behavior by 3rd party application can trick XQP to add LBN 0 to free LBN list



GUARD LBN0 ALLOCATION

FEATURE DETAILS

- Existing Limitation
 - No consistency check for LBN 0 allocation while creating files
- Solution
 - Consistency check added to detect LBN 0 allocation to files
 - LBN 0 allocation to file results in a System crash
- Problem that still remains
 - Applications can still do Logical IO to LBN 0 and corrupt the disk
- Patch with the feature
 - ❑ VMS84I_UPDATE-V0500, VMS84A_UPDATE-V0500
 - ❑ Will be back ported to earlier releases

Consistency check added to detect LBN 0 allocation to files



RMS UPDATE

MBC (MULTI BLOCK COUNT) INTRODUCTION
V8.4 & BEFORE



MBC

INTRODUCTION

- MBC stands for Multi block count, used by RMS

- MBC value of X means RMS buffers for the file is of size X blocks
 - ❑ IO's issued by RMS will be of size X blocks
 - ❑ Basic unit of synchronization for the file is X blocks

- MBC – Default value is 32 blocks, Maximum value is 127 blocks

- Applicable only to sequential files

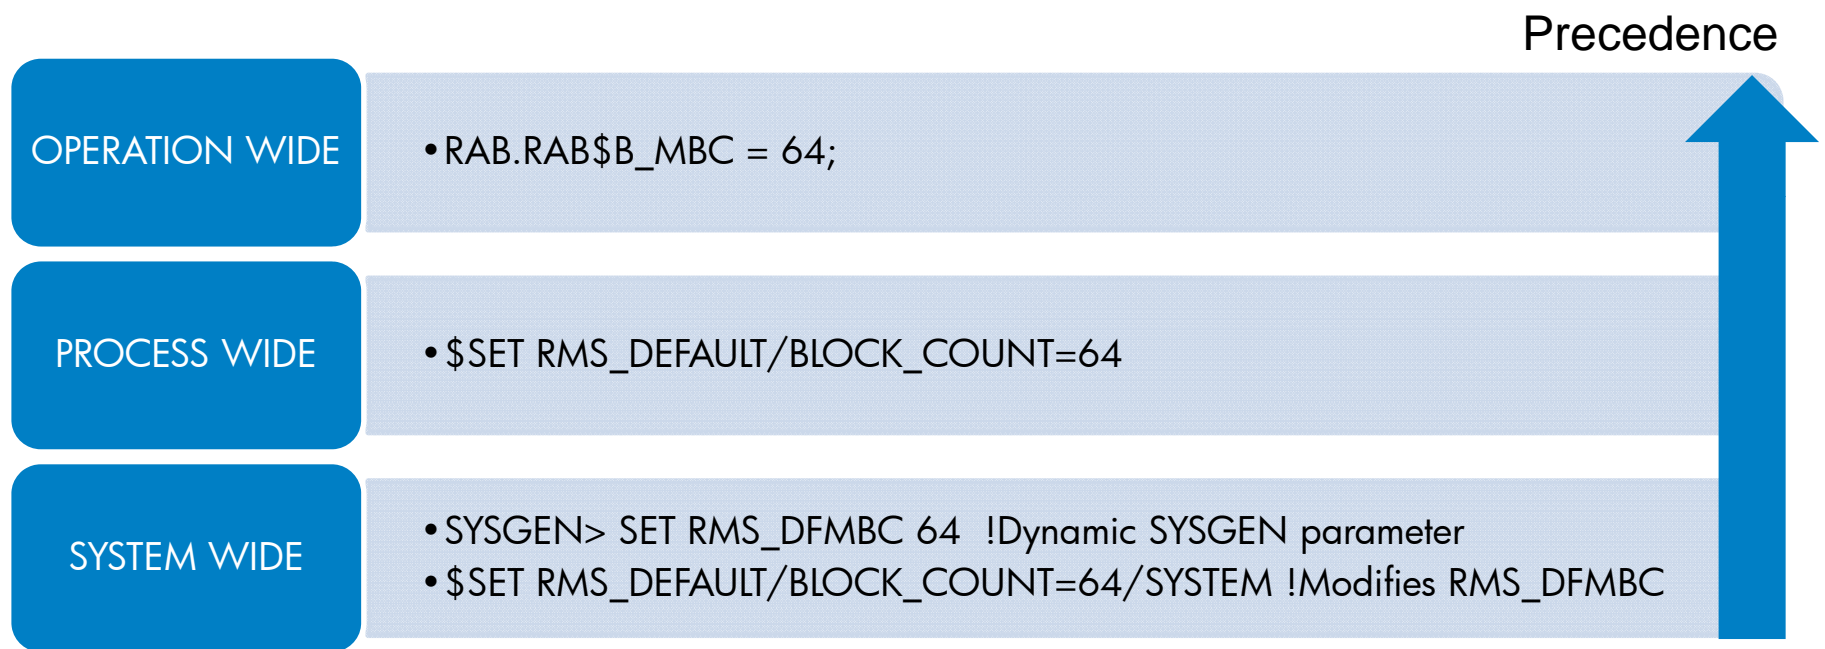
MBC feature existed before OpenVMS V84 Release , No changes in V84 SSB



MBC

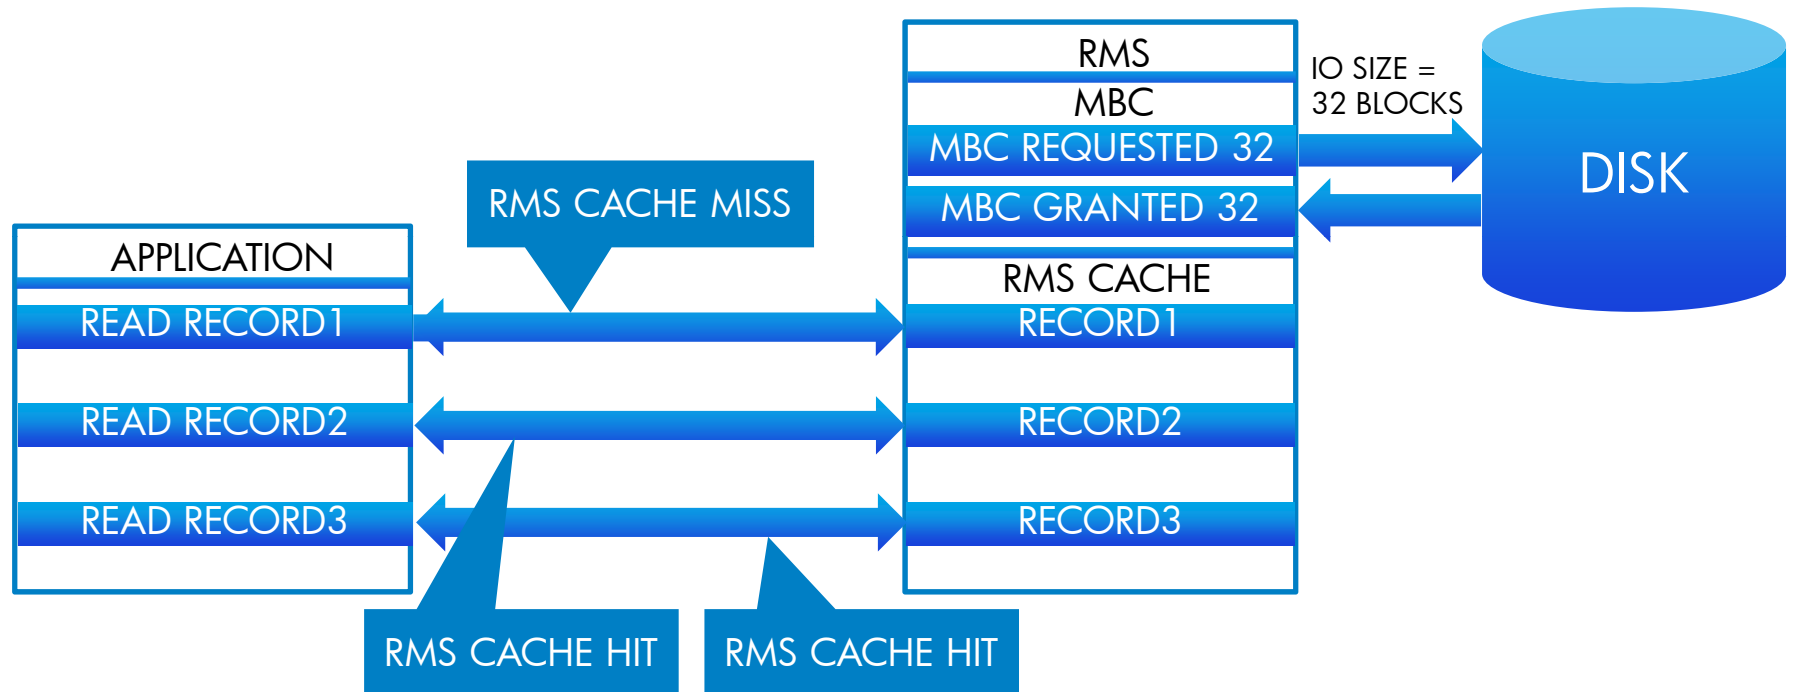
CONTROLS

➤ MBC Controls



MBC

USAGE – NON CLUSTER ENVIRONMENT

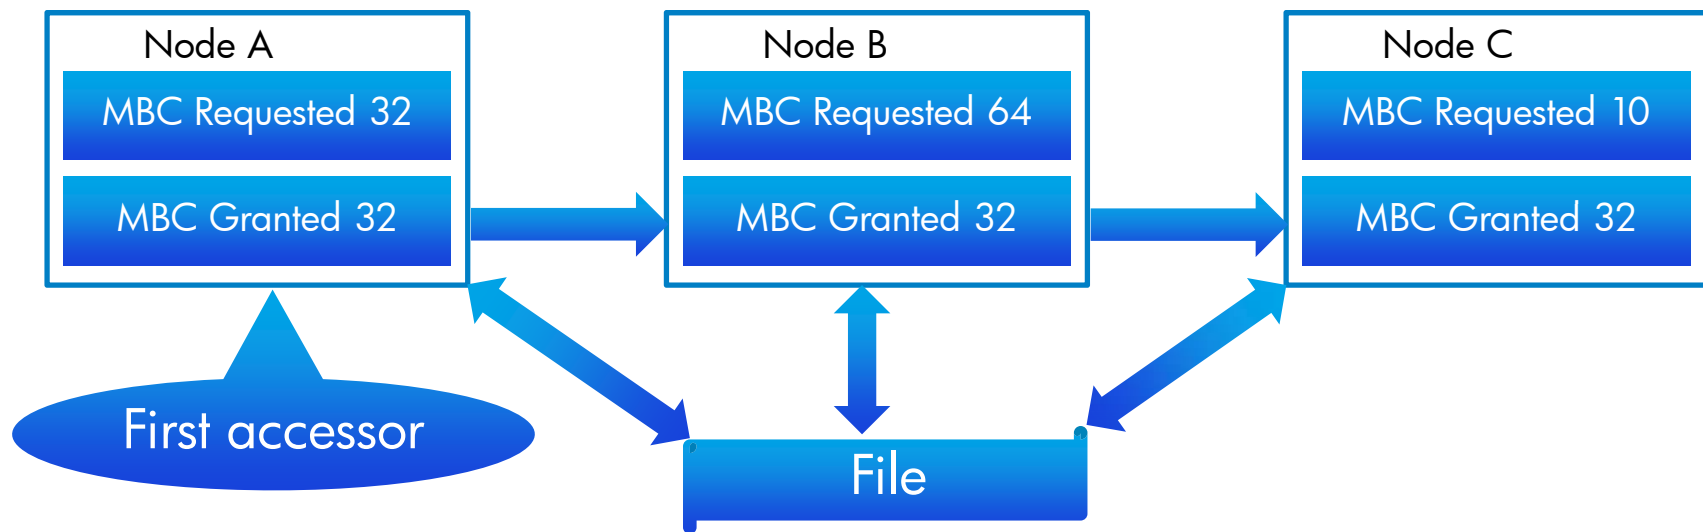


Larger MBC value -> Larger RMS IO's -> More data read from file per IO

MBC

USAGE – CLUSTER ENVIRONMENT

- File can be accessed cluster-wide with different MBC values
 - ❑ First accessor of the file decides the MBC value
 - ❑ Subsequent accessor's of the file settle for MBC value decided by the first accessor



First accessor of the file decides the MBC value for the file to be used cluster-wide

RMS UPDATE

1) MBC (MULTI BLOCK COUNT) ENHANCEMENT
POST V8.4



MBC ENHANCEMENT

INTRODUCTION

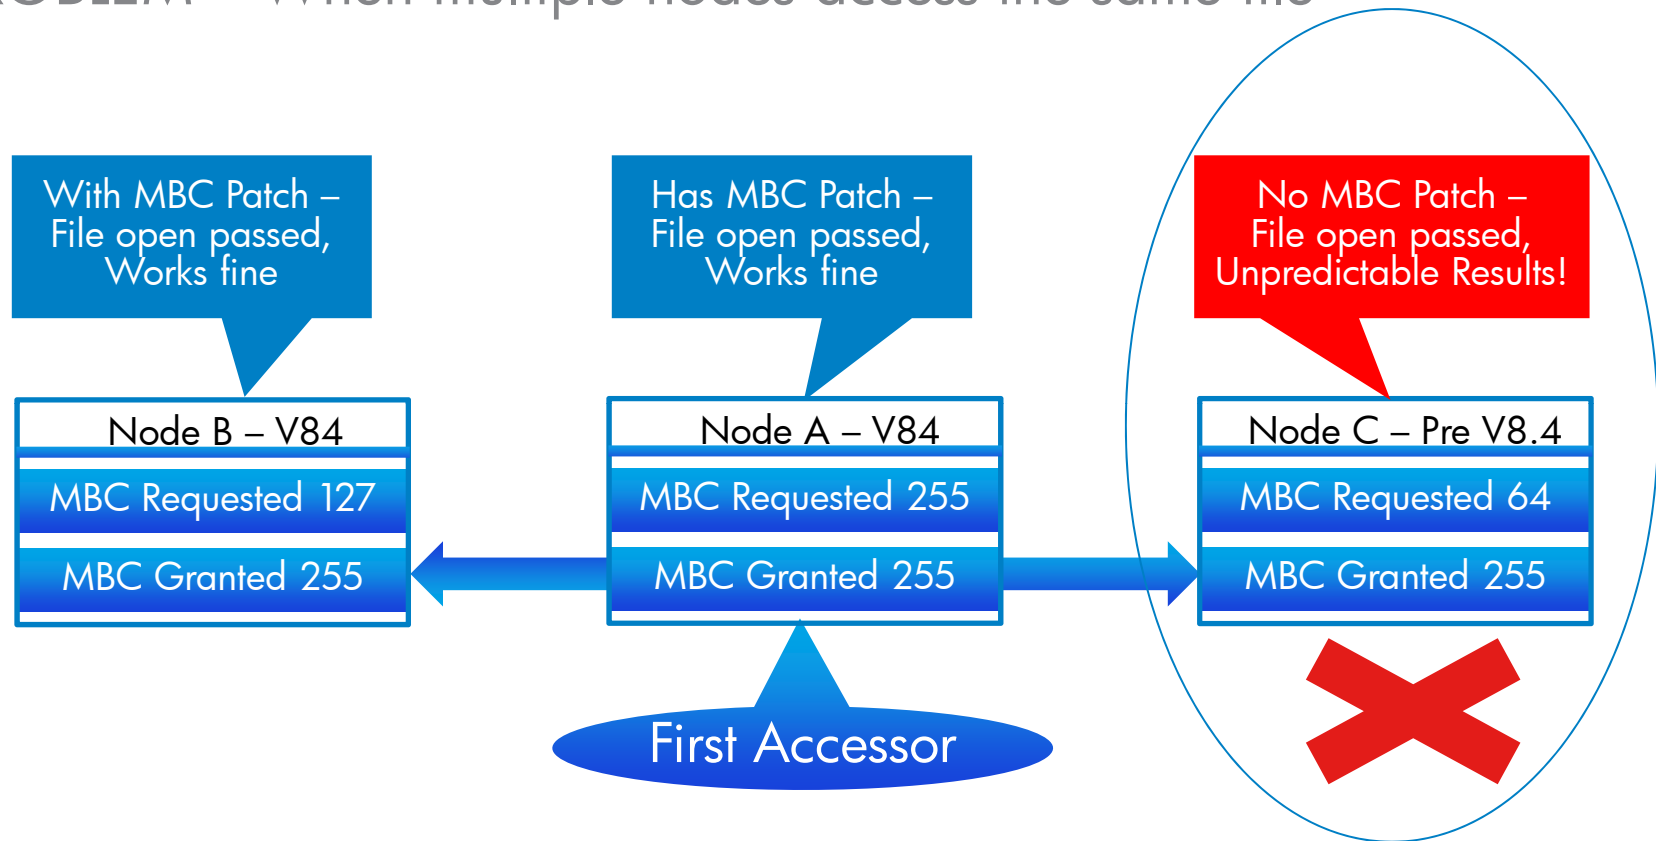
- Maximum value of MBC increased from is 127 to 255 blocks
- VCC_MAX_IO_SIZE Dynamic SYSGEN parameter needs to be increased
 - ❑ Set it as per MBC value, XFC will then cache MBC IO's
- Helps reducing the number of IO's by increasing the IO size
- Feature Release
 - ❑ Introduced in VMS84I_UPDATE-V0500/ VMS84A_UPDATE-V0500

MBC enhancement feature released for OpenVMS V8.4



MBC ENHANCEMENT

PROBLEM – When multiple nodes access the same file



Nodes without MBC patch that access files opened with MBC>127 will cause unpredictable results

MBC ENHANCEMENT

SOLUTION – When multiple nodes access the same file

- On pre V84 nodes, fail access to files already opened with MBC > 127
 - ❑ File access fails, problem avoided

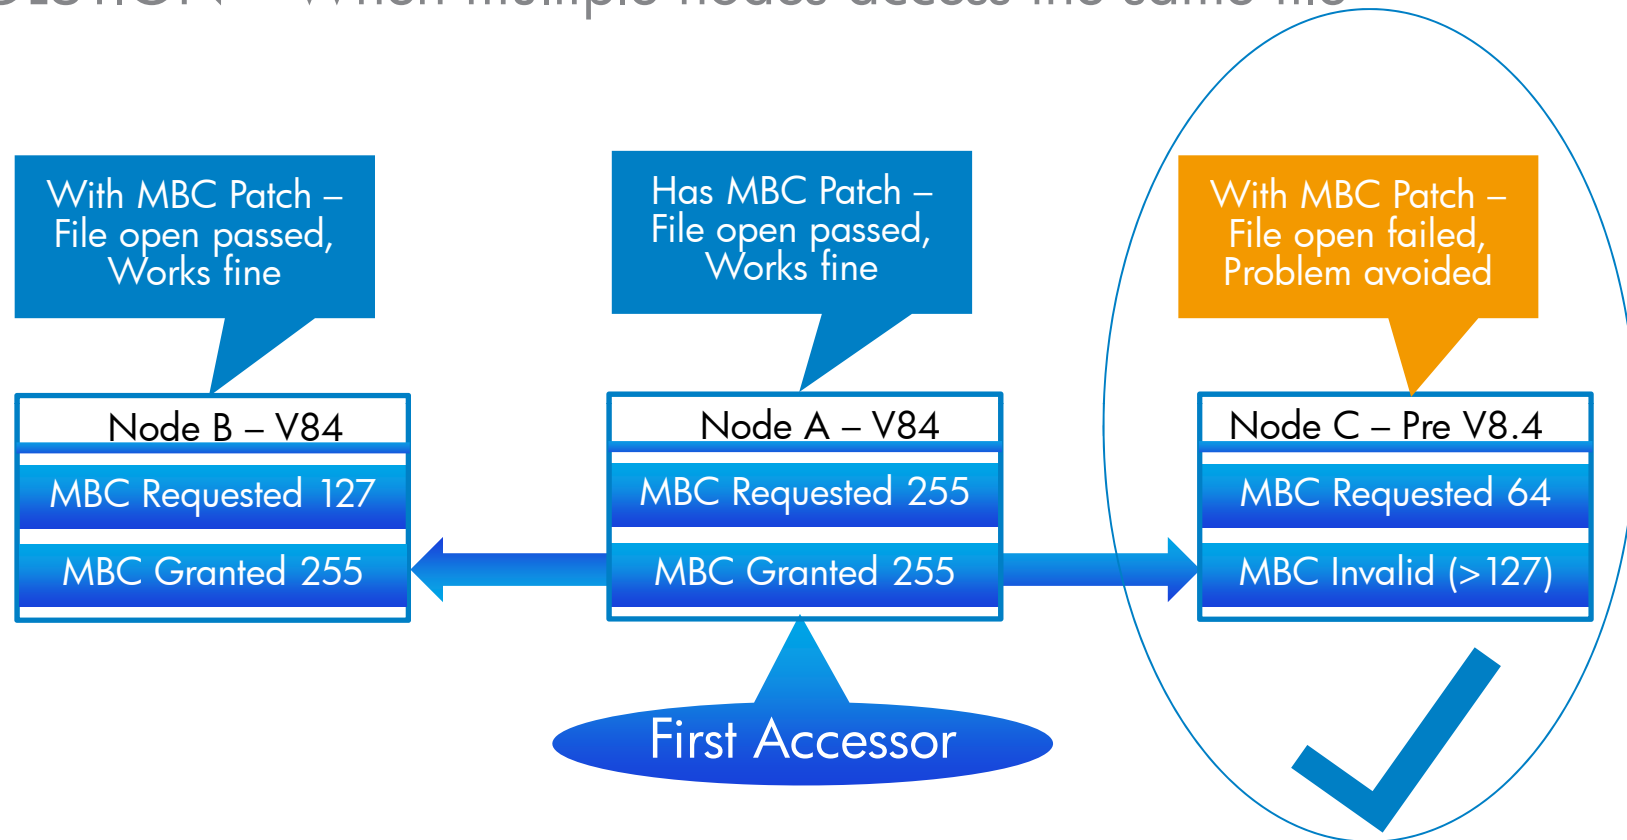
- Patches released for pre V84 versions
 - ❑ VMS732_RMS-V0600
 - ❑ VMS82A_RMS-V0400
 - ❑ VMS821I_RMS-V0500
 - ❑ VMS83I_RMS-V1000/VMS83I_SYS-V1600
 - ❑ VMS83A_RMS-V1200/VMS83A_SYS-V2000
 - ❑ VMS831H1I_RMS-V0400/VMS831H1I_SYS-V1300

MBC enhancement fixes included in latest V8.4 RMS TIMA Kits



MBC ENHANCEMENT

SOLUTION – When multiple nodes access the same file



Install MBC patch on nodes that access files opened with MBC > 127



MBC ENHANCEMENT

Latest Update

- MBC Enhancement – Latest version
 - VMS84I_RMS-V0200/VMS84A_RMS-V0200
 - Will be included in VMS84I_UPDATE-V0600/VMS84A_UPDATE-V0600



NEW

MBC enhancement fixes included in latest V8.4 RMS TIMA Kits





RMS UPDATE

2) SYMLINKS FEATURE



SYMLINKS FEATURE

UPDATE

SYMLINK functionality strengthened with continuous improvement

➤ Fixes in V8.4 UPDATE500

- ❑ Unable to access directories with UNICODE characters in its filename
- ❑ SEARCH/NOWARNING unable to suppress unresolved symbolic link traversal errors

➤ Fixes in V8.4 UPDATE600

NEW

- ❑ ANAL/DISK/REPAIR unable to repair large number of Pre V84 SYMLINK files on the disk
- ❑ COPY/SYMLINK copies the target file instead of the source files
- ❑ BACKUP does not report any message for not traversing the SYMLINK file
 - ✓ BACKUP to report "BACKUP-I-SYMNOTFLW" message
- ❑ Unable to handle directory names ending with . (dot)
- ❑ SYMLINK is supported on ODS5 disk but still able to create SYMLINK on ODS2 disk
 - ✓ SYMLINK creation on ODS2 disk blocked



QUESTIONS/COMMENTS

CONTACT

➤ Business Manager (Rohini Madhavan)

☐ rohini.madhavan@hp.com

➤ Office of Customer Programs

☐ OpenVMS.Programs@hp.com



Thank you

