Patching and Upgrading Tru64 UNIX and TruClusters

Martin Moore

Team Leader, Tru64 UNIX Support
HP Services

Session Topics

- Patching Tru64 UNIX and TruClusters
  - Background info
  - Best practices
- Upgrading Tru64 UNIX and TruClusters
  - Overview
  - Best practices
- New Hardware Delivery (NHD) kits
- Related sessions:
  - 3375 -- HP's UNIX Patch Strategy: Moving from Tru64 to HP-UX
  - 3759 – V5.1B and its Enhancement Releases

Types of patches

- Customer-Specific Patch (CSP)
  - Addresses a reported problem for a specific customer
    - Customized for customer's system and patch level
    - Made available only to a specific customer
  - Unit testing performed
  - Shipped in *dupatch/setld* format
  - Rolled into a future release patch kit

Types of patches (2)

- Early Release Patch (ERP)
  - A single patch available to all customers in response to a critical or widespread problem, such as:
    - Data integrity
    - Security
    - Detrimental to system operation
  - Uses CSP packaging mechanism with slightly different documentation
  - More extensive testing than CSP
  - Accompanied by Product Advisory and/or CERT advisory
  - Rolled into a future release patch kit
  - Available via patch web/FTP site – current ERP list at http://h30097.www3.hp.com/unix/EarlyReleasePatch-download.html

Types of patches (3)

- Release Patch Kit
  - General distribution of patches
  - Available to all customers via FTP/web site:
    - http://www.itrc.hp.com/service/patch/mainPage.do
  - Testing is a scaled version of a full functional release
  - Two types: Initial Patch Kit (IPK) or Aggregate Patch Kit
    - IPK contains fixes to prevent upgrade regression (forward porting of existing patches not in new version of O/S) plus fixes for new problems found in release qualification testing
    - Aggregate kits (subsequent base level kits after IPK) contain cumulative fixes for problems reported by customers or HP
    - Both may also include new functionality

Patch distribution model

- Ship 1 or 2 kits per base level
  - Deliver kits more frequently for releases with the highest volume of incoming problem reports
  - Deliver kits less frequently for releases with low volume
- All release patch kits, aggregate and IPKs, go through an extensive test matrix before being released
  - Scaled version of full quality testing that is performed in base OS/TruCluster development releases

Patch roadmap

- Recent and *forthcoming* patch kits
  - BL24: V5.1B pk3, V5.1A pk6, V5.1B-1
  - BL25: V5.1B pk4, *V5.1B-2*
  - BL26: *V5.1B pk5, V5.1B-3*
- Patch roadmap viewable at:
  - http://www.tru64unix.compaq.com/docs/patch/roadmap.html

    - or -
  - http://h30097.www3.hp.com/docs/patch/roadmap.html

Supported versions

- Currently Supported Releases
  - V4.0F/TCR1.6 (Prior Version Support thru 30-Jun-2007)
  - V4.0G/TCR1.6 (PVS thru 30-Jun-2007)
  - V5.1A/TCR5.1A (PVS thru 30-Jun 2007)
  - V5.1B/TCR5.1B (Standard support through EOL – at least 2011)

http://www.hp.com/hps/os/os_pvs_amap.html

Patch kit paradigm

- Prior to V5.1B pk4, patches are “a la carte”
  - Install/Remove all patches or selected patches (within limits; some patches depend on others)
  - HP strongly recommends installing **all** patches
- V5.1B pk4 and later are “all or nothing” – no more “a la carte” selection of patches!
- Patches can be installed as reversible or nonreversible
  - Reversible
    - Consumes additional disk space
    - Patches can be removed recursively
    - Strongly recommended
  - Nonreversible
    - Additional disk space not needed
    - Patches cannot be removed

Patch kit naming

- Release kits (as of BL18)
  - Example:

    OS Product|Version|Base Level|KitType|Kit#|-MfgDate

Patch kit file T64V51B18AS0003-20011020.tar is read as:

Tru64 UNIX (and TruCluster) software (T64)
Version 5.1 (V51)
Base level 18 (B18)
Aggregate (selective installation) patch kit (AS)
Patch Kit 3 (0003)
Manufactured on October 20, 2001 (-20011020)

Patch kit naming (2)

- ERP's and CSP's – Old style
  - Example: T64V51AB21-C0020100-12345-ES-20030611.tar
    - T64V51AB21: Product, version, base level (Tru64 UNIX V5.1A, BL 21)
    - C0020100: Patch number of one patch included in kit (patch 201.00) – kit may or may not include other patches
    - 12345: Unique kit ID number (unique for this CSP)
    - ES: Kit type identifier
      - M = Manual patch (not installed with dupatch)
      - E = ERP (not present implies CSP)
      - S = Security patch
    - 20030611: Manufacture date (June 11, 2003)
  - The kit number (12345 in this example) is the key field to uniquely identify this kit

Patch kit naming (3)

- ERP's and CSP's – New style (since late 2003)
  - Example: T64KIT0012345-V51AB21-ES-20030611.tar
  - Minor changes from previous style
    - Individual patch number dropped
    - "KIT" and unique kit number moved to first field
  - Kit type identifier may include "C" to identify CSP
    - C and E mutually exclusive

Patch documentation

- Publicly available at http://www.tru64unix.compaq.com/docs/patch - or - http://h30097.www3.hp.com/docs/patch/
- Patch Kit Installation Instructions
  - Installing and removing patches using the *dupatch* utility
  - Describes baselining techniques
  - Instructions for rolling and no-roll patching in clusters
  - Provides other information for working with patches
- Patch Summary and Release Notes
  - Provides information about specific patches
  - Summarizes Tru64 UNIX and TruCluster patches included in the kit
  - New: includes list of CSP's superseded by kit

Best practices for patches

- Patch “Best Practice” document on-line at http://www.tru64unix.compaq.com/docs/best_practices/
  - Not the same as this presentation
- Back up your system first!
  - If a catastrophe occurs while patches are installing, you may have to restore system (or undo rolling upgrade)
  - At a minimum: vdump the /, /usr, and /var filesystems
- Review patch kit release notes
  - Especially if you have CSP’s or ERP’s installed
  - If questions remain about whether a CSP or ERP fix is included in a release patch kit, contact HP support

Best practices for patches (2)

- Check for presence of CSP's / ERP's
  - "Show System Patch History" option from "Patch Tracking" menu
  - (or) *dupatch -track -type kit*
- Possibly remove CSP's/ERP's installed by dupatch
  - Historically, dupatch didn't know if fixes were in the release kit, so it wouldn't install patches on top of them
  - In new kits, dupatch can reconcile most CSP's with release kits -- and install patches over CSP's when fixes are included in the release kits
    - Some CSP's may not reconcile and will need to be removed
    - Release notes now include list of CSP's with fixes in kit

Best practices for patches (3)

- Run “Baseline Analysis/Adjustment” in dupatch
  - Checks for additional missing or inconsistent files that block the installation of specific patches
  - Could be manual (non-dupatch) patches, third-party or freeware utilities, locally modified programs, etc.
  - Enable patch installation IF you determine that’s the right course
- Put a meaningful comment in installation log
  - e.g., “Installing security ERP for CERT advisory XXXX”
  - Dupatch transactions are logged in /var/adm/patch/log/
- Always install patches reversibly
  - Preserve backed-up files in case you need to reverse patches
    - Location selectable; */var/adm/patch/backup* by default

Patching clusters

- Rolling Patch
  - Allows cluster to remain up while patch operations are performed sequentially on cluster members
  - Maintains availability of critical applications
  - Can combine aggregate kit + ERP's/CSP's in one roll
  - Can combine *installupdate* (and/or NHD) + patches in a single roll
- No-Roll Patch
  - Introduced in BL19 patch kits (Spring 2002)
  - Allows patches to be managed on the whole cluster at once rather than using rolling patch
  - Provides the ability to patch a cluster quickly with minimal downtime

Version switch

- Prevents the introduction of potentially incompatible new features until all members have been updated with necessary components
- Applicable mostly in cluster environments, but may be required for a few patches on standalone systems
  - Release notes & dupatch special instructions will indicate if needed
- Command to activate new version varies:
  - Cluster, rolling upgrade: *clu_upgrade switch*
  - Cluster, no-roll install: */var/adm/patch/noroll/noroll_versw*
  - Standalone system: *versw –switch*
- Use of the version switch requires a cluster-wide shutdown to remove those patches requiring the version switch
  - Procedure varies for different versions; consult kit documentation

Phases of the rolling upgrade process

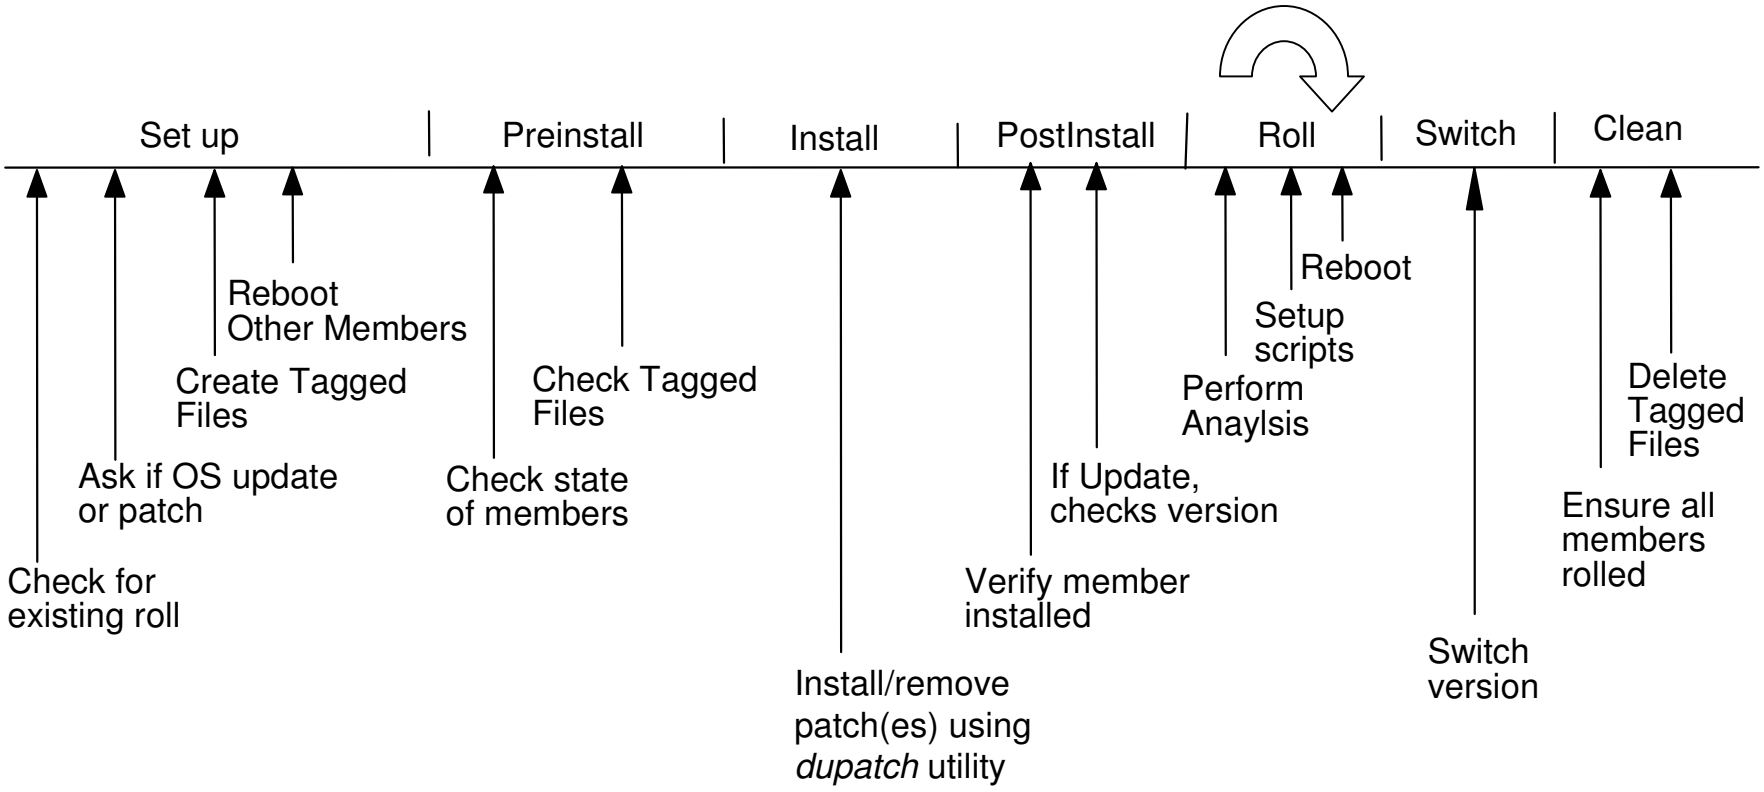

Rolling upgrade: Preparation

- Back up the cluster
- Upgrade firmware if necessary
- Choose a lead member
- *clu_upgrade -v check setup <lead_member_id>*
  - "check" is a keyword used to check readiness for any clu_upgrade step: e.g., *clu_upgrade check install*

clu_upgrade setup

- Verifies:
  - System is part of a cluster
  - No roll is currently in progress
- Asks whether OS upgrade, NHD, patch, or combo
- Creates tagged files (prefixed with ".Old..")
  - For best performance, relocate CFS server for /, /usr, /var to the member running this step
- Instructs user to reboot all cluster members except the lead member
  - Non-lead members now running on tagged files
  - Usually done one member at a time to maintain quorum
- Usually takes 45 minutes to 2 hours

clu_upgrade preinstall

- Executed on the lead member
- Verifies:
  - Tagged files are present
  - Tagged files match the inventory
  - All non-lead cluster members are running on tagged files
- Usually takes 15 to 30 minutes

clu_upgrade install

- Executed on the lead member
- Patches installed/removed using *dupatch* utility
- At this point, only the lead member is running with patches installed (or removed); all other members are running with tagged files and patches uninstalled (or not yet deleted)
- dupatch can be run multiple times
  - Can install from multiple patch kits (aggregate and/or ERP and/or CSP)
  - Multiple patch deletions and/or installations can occur within the same roll

clu_upgrade postinstall

- Executed on the lead member only
- Checks to see if dupatch has been executed before this step is completed

clu_upgrade roll

- All non-lead members must be rolled
- This is an iterative step
  - Repeated on each member until complete
  - Members can be rolled in parallel if cluster quorum is maintained
- Only member-specific files are modified because the shared files are already in place from the roll of the lead member
- Sets up *it(8)* job that will be executed on reboot to perform the roll

clu_upgrade roll (2)

- Runs in single-user mode on each member
- Backs up member's member-specific files
- Reboots the member and upon reboot:
  - Copies and configures member-specific files from *member0* directory
  - Builds new custom kernel for each member
  - Reboots customized kernel
- All cluster members now have patches configured

clu_upgrade switch

- Run once after all members have completed roll
- Any installed patch controlled by a version switch will be enabled on all cluster members
- Prior to this phase, a patch containing a version switch can be deleted
- After this step, the roll itself must be completed; the user is no longer permitted to undo steps

clu_upgrade switch (2)

- After the version switch has been thrown, removing a patch containing a version switch must follow a special procedure:
  1) Complete the roll if still active
  2) Run the associated undo script (supplied with the patch subset)
  3) Remove the patch

clu_upgrade clean

- Verifies that the switch stage has completed
- Removes the tagged files
- Ensures proper cleanup so that future rolls can be performed
- Creates an archive of the logs and status files under */cluster/admin/clu_upgrade/history*
- Usually takes 30 to 90 minutes
  - As with other steps, run on member that is CFS serving /, /usr, and /var filesystems (or relocate them)

No-roll patch process

- Introduced in BL19 (V5.1A PK2/V5.1 PK5) as an alternative to the rolling patch process
  - Some admins considered the cluster to be unavailable while in the midst of a roll
  - May be applying only a single patch (ERP or CSP) and don't need or want to use the lengthy roll process
- Provides the admin with a procedure to maintain patches with minimal downtime and less manual intervention
- Applies patches to all members automatically

No-roll patch process (2)

- Requires less disk space than rolling patch process (no tagged files)
- Cluster services are not available during the no-roll process!
- Cluster restarted after all members have been patched

No-roll patch – Preconditions

- Quorum configuration is important
  - Quorum must be configured such that all cluster members can go down (serially or in tandem) while maintaining quorum
  - Detailed information provided in the *Cluster Installation Guide*
- Cluster members must be in multi-user mode because EVM is used to generate cluster-wide events

No-roll patch timeline

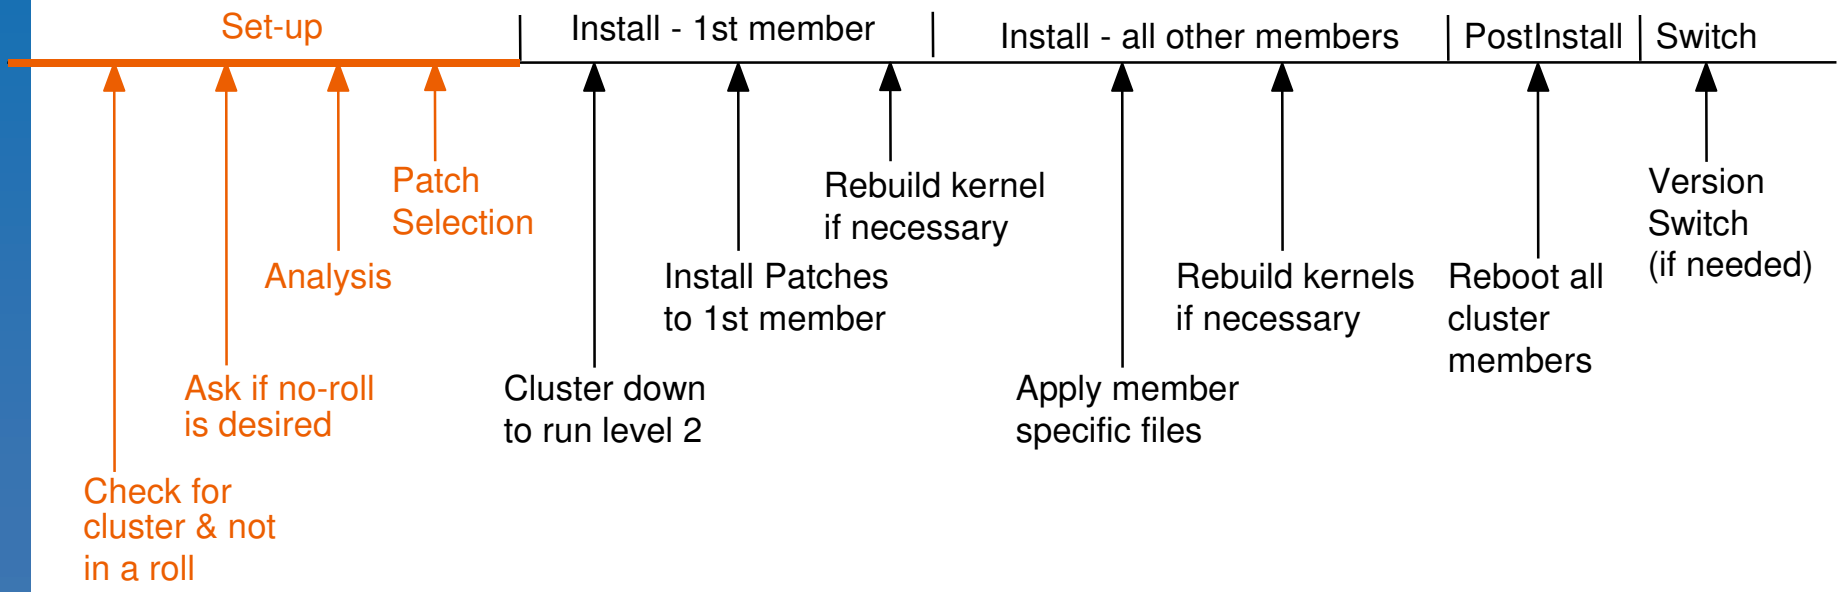

No-roll patch - Set-up

- In *dupatch*, user selects install/delete patches
- Checks for correct state:
  - Configuration is a cluster
  - A cluster roll is not in progress
- Confirms that a no-roll operation is desired
- Verifies that cluster is in multi-user mode
- Performs analysis
  - Disk space
  - Dependencies
- User selects patches to be installed (older kits only)
- Remainder of the process is automated

No-roll patch timeline

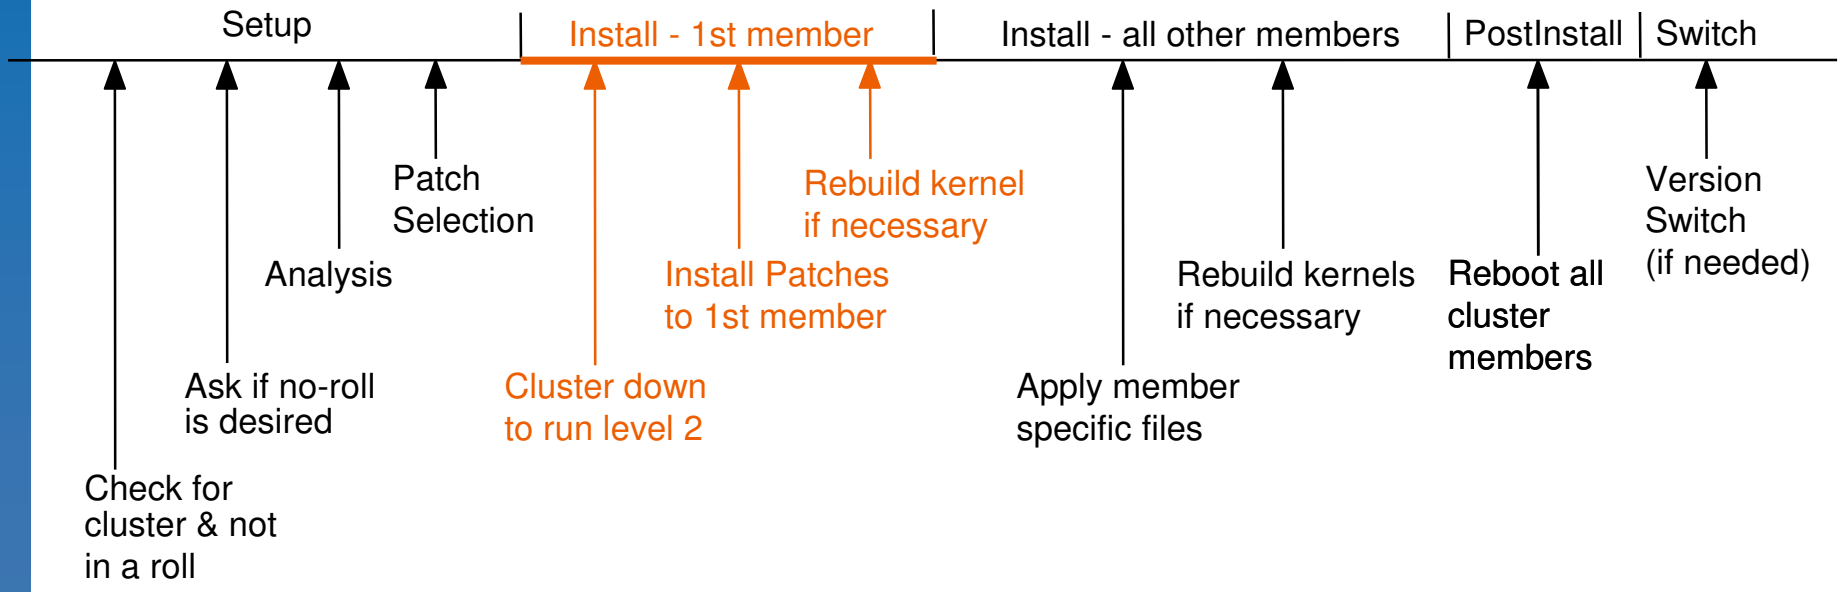

No-roll patch - Install first member

- “First member” is the cluster member running dupatch
- Cluster is brought down to run level 2 via a cluster-wide event
- Patches are loaded and configured on the first cluster member
- If any patches require a kernel build, a new kernel is built automatically

No-roll patch timeline

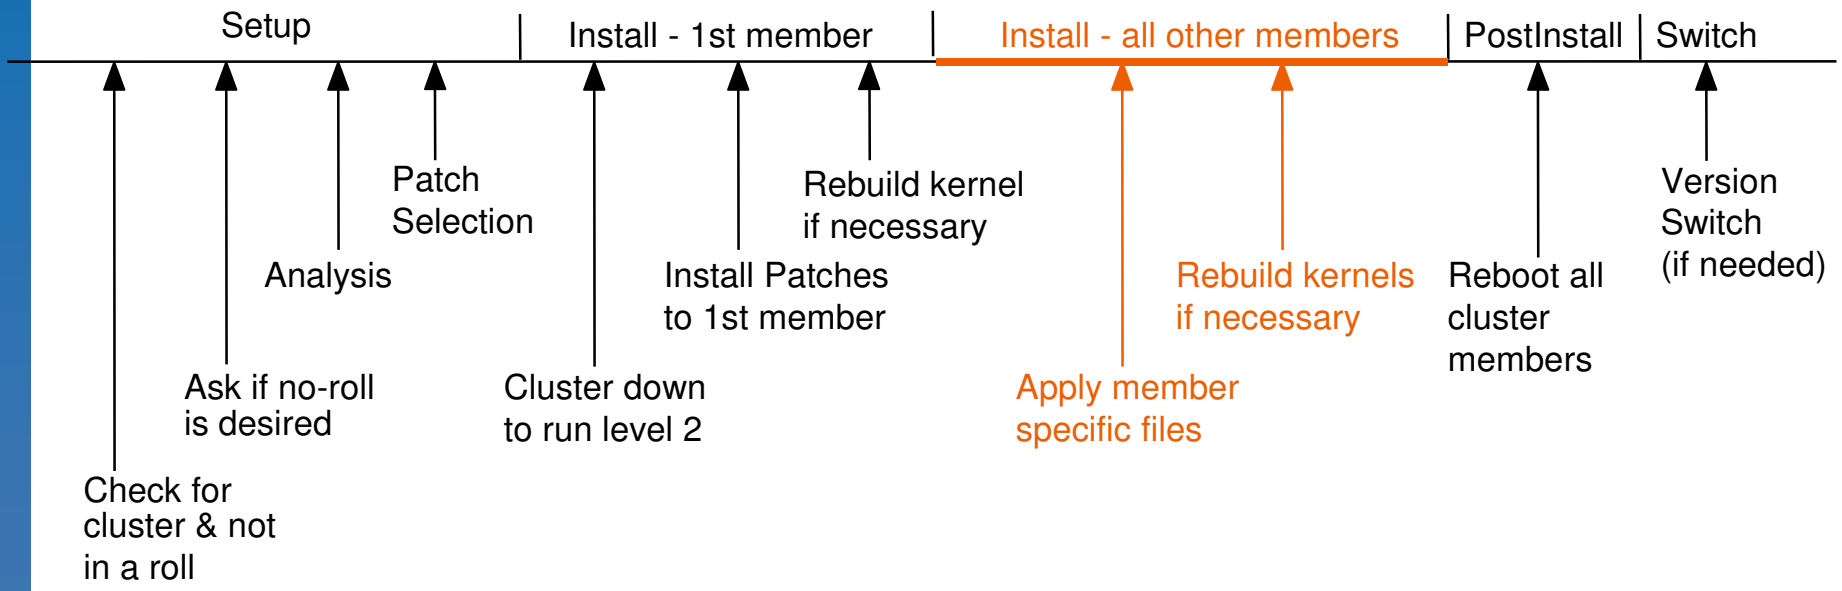

No-roll patch - Install all members

- A second cluster-wide event is issued on all running cluster members:
  - Executes */var/adm/patch/patch_member_noroll* script
  - Operation is done in parallel to minimize downtime
- For a down member:
  - Posts an *it(8)* job which, upon reboot, will run /var/adm/patch/*patch_member_noroll*
  - Populates member-specific files and automatically rebuilds kernel, if necessary

No-roll patch timeline

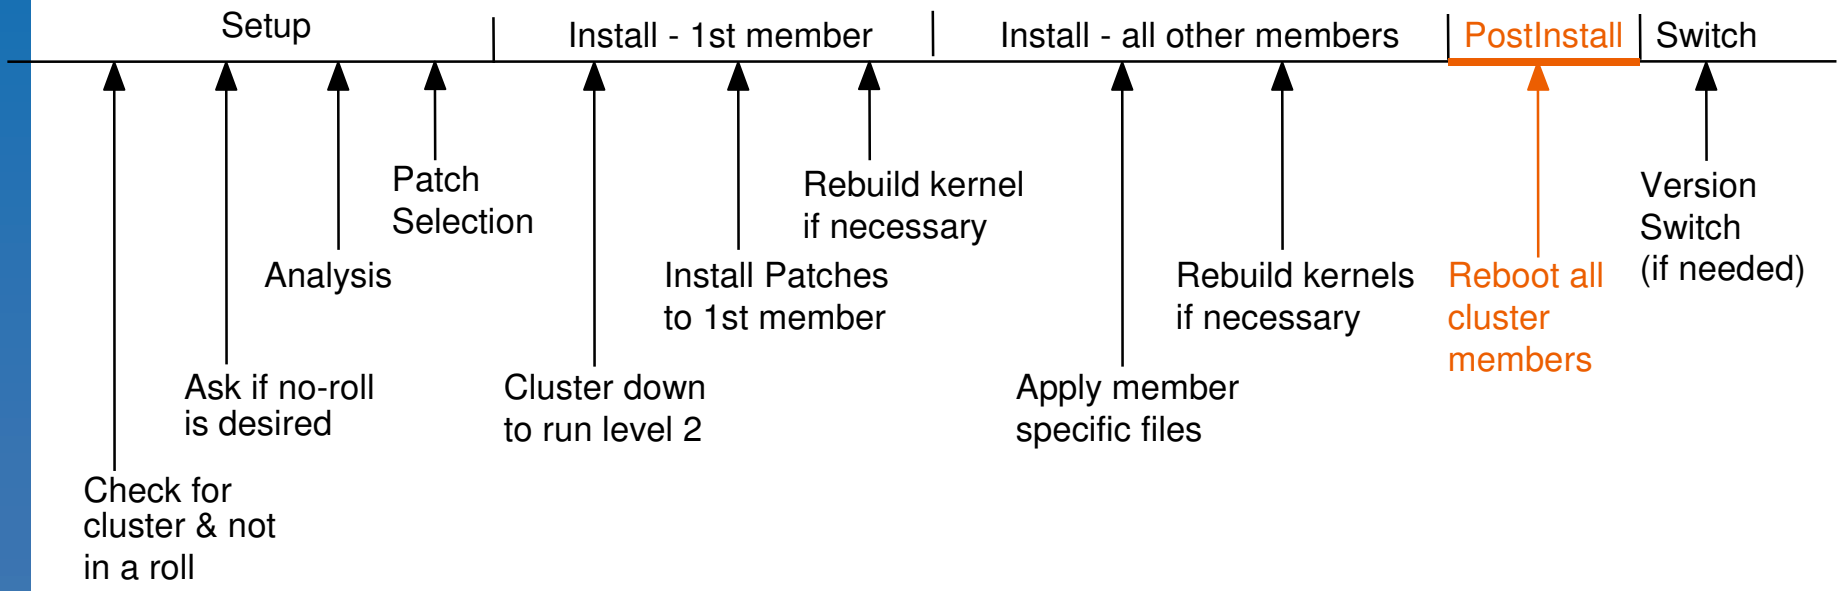

No-roll patch - Postinstall

- Status of all members is checked
- Cluster-wide event is issued that causes all members in the cluster to be rebooted
- Cluster members will reboot in a fashion that will not lose cluster quorum during the reboot

No-roll patch timeline

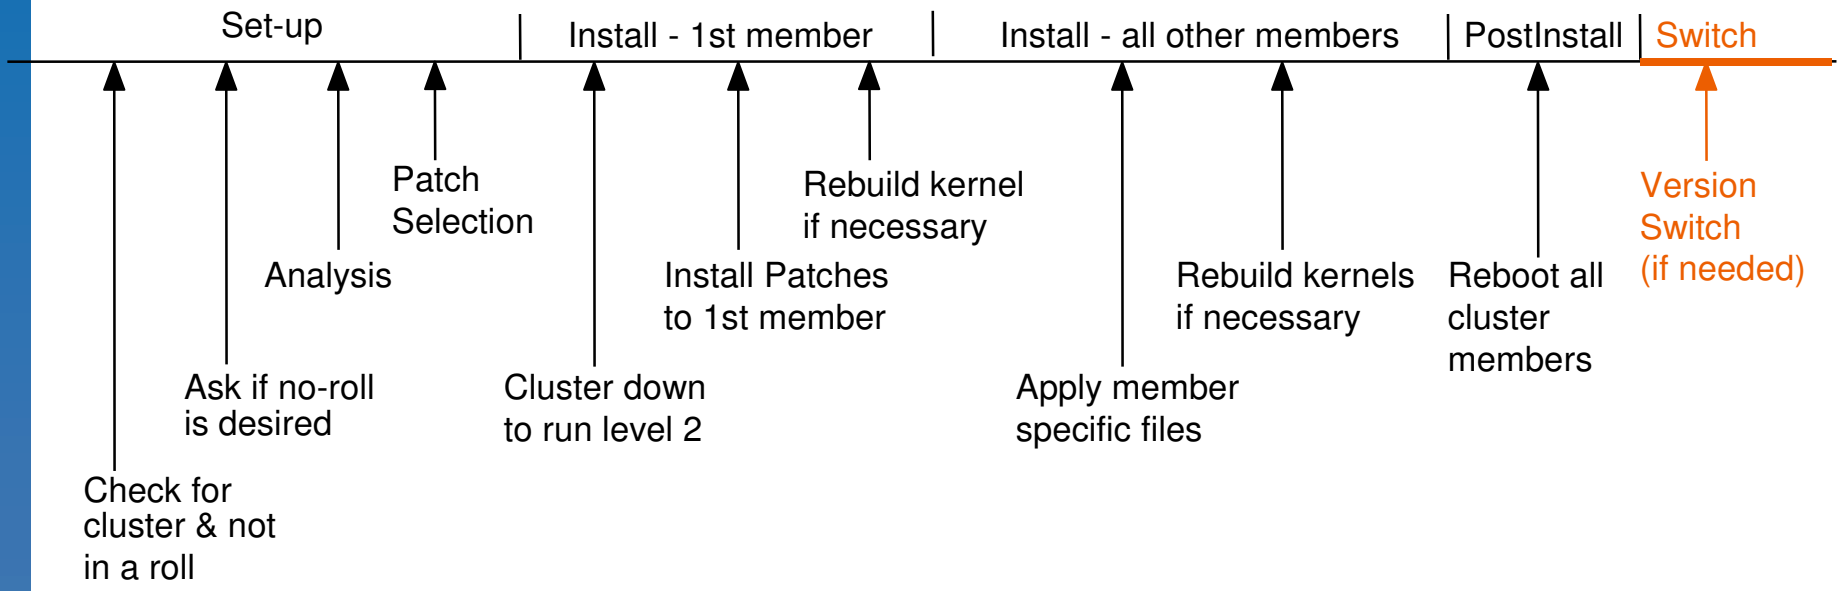

No-roll patch - Switch

- If a patch containing a version switch has been installed, the user must manually throw the version switch for the cluster as a whole, by running the script */var/adm/patch/noroll/noroll_versw*
  - All cluster members must be up at this time
- Once the switch has been thrown, if the version switch patch needs to be removed, the undo script must be run followed by patch deletion (identical to rolling patch process)
- It is recommended that the cluster be rebooted after throwing the switch

Handling errors in no-roll patch

- At the start of a no-roll operation, *dupatch* determines which members are presently up
  - Members that were known to be running at the start of the no-roll operation, but do not respond to later cluster events, will cause the no-roll operation to be suspended until the problem has been corrected by the user
- Status from the script is reported in a file in the shared directory */var/adm/patch/noroll_results*

Rolling vs. no-roll comparison

	Rolling Patch	No-Roll Patch
Cluster Services	Available	Unavailable
Time to Complete	Lengthy	Minimal
Reboots per Member	4	1
Additional Disk Space Needed	Yes (for tagged files)	No
Automated Process	No	Yes
Combine Patch with OS Upgrade	Yes	No
Removal of version-switch patch	Manual	Manual

Upgrading Tru64 UNIX

- *installupdate* utility does the following:
  - Checks for file type conflicts
  - Checks for sufficient disk space
    - Provides several options to free up space
  - Notifies of possibly conflicting layered products
  - Optionally removes blocking layered products
  - Updates base OS, TCR, and WLS subsets to new version

Tru64 UNIX upgrade paths (standalone systems)

To get to this version:	Do one *installupdate* from:
V4.0D	V4.0A, V4.0B, V4.0C
V4.0E	V4.0B, V4.0C, V4.0D
V4.0F	V4.0D, V4.0E
V4.0G	V4.0D, V4.0E, V4.0F
V5.0	V4.0D, V4.0F
V5.0A	V4.0F, V5.0
V5.1	V4.0G, V5.0A
V5.1A	V5.0A, V5.1
V5.1B (, B-1, B-2, B-3, …)	V5.1, V5.1A

TruCluster upgrade paths

- A rolling upgrade can only be performed to the next higher version of TCR software
  - V1.4A to V1.5
  - V1.5 to V1.6
  - V5.0A to V5.1
  - V5.1 to V5.1A
  - V5.1A to V5.1B
- No rolling upgrade from TCR 1.6 to TCR version 5 due to introduction of Single System Image
- So clusters must roll in small steps
  - V5.1 cluster must roll to V5.1A before rolling to V5.1B
  - Standalone V5.1 system could go to V5.1B in one step

The “dash” releases: V5.1B-n

- Ship new fixes and minor functionality within the patch kit
  - V5.1B-1 = V5.1B base + pk3 (+ NHD-7 if needed)
  - V5.1B-2 = V5.1B + 5.1B-2/PK4 (+ NHD-7 if needed)
  - V5.1B-3 = V5.1B + 5.1B-3/PK5 (+ NHD-7 if needed)
- V5.1B-n is new minor update version
- Base version is still V5.1B (ISVs don't want a 5.1C)
  - Is a minor release with all the bits on separate media as opposed to being snapped onto a 5.1C base OS CD-ROM.
  - “sizer -v”, */etc/motd* still report “V5.1B”, not “V5.1B-1”
- Update Associated Products CD and related media
  - Includes updates to many APCD products.
- *installupdate* command not needed
- See session 3759, Friday 08/20/2004 08:00AM

Best practices for upgrades

- Back up your system first
  - If a catastrophe happens while subsets are loading, you'll probably have to restore the system (or undo rolling upgrade)
  - At a minimum: vdump the /, /usr, and /var filesystems
- Check release notes and installation guide
  - Is platform still supported? (Very old ones may not be.)
- Upgrade firmware if needed
  - After upgrade, re-run ECU on systems with EISA buses
- Verify AdvFS domains
- Run *cdslinvchk* to verify CDSL integrity
- No need to install patches on old version prior to upgrade (but do install them on the new version!)

Things to consider

- How many upgrades to do vs. a full install?
  - Time/risk tradeoff vs. effort to recustomize system
- 4.0F to 5.0A upgrade is problematical
  - Avoid if possible by choosing a different path: e.g., 4.0F -> 4.0G -> 5.1 -> 5.1B instead of 4.0F -> 5.0A -> 5.1A -> 5.1B
- If root directory is mirrored via LSM, consider unencapsulating it during upgrade
  - One less layer of complexity
  - Former mirror plex is a snapshot of pre-upgrade root
  - But don't just "break mirror" by disassociating one plex; this can cause problems when reassociating after the upgrade is finished.

Something to watch out for

- This problem exists only if Enhanced Security is in use
- Rolling upgrade of cluster from V5.1A + pk5 or lower, to V5.1B + pk3 or higher
  - prpasswdd inter-node communication changed
  - After roll of first member, all logins will hang (thus, you can't log back in to complete the upgrade!)
  - Workaround: disable prpasswdd before roll phase

```
# cd /sbin/rc3.d
# mv S22prpasswd disable.S22prpasswd
# /sbin/init.d/prpasswd stop
```

← on all members

  - After rolling upgrade finishes, re-enable prpasswdd

```
# cd /sbin/rc3.d
# mv disable.S22prpasswd S22prpasswd
# /sbin/init.d/prpasswd start
```

← on all members

New Hardware Delivery (NHD)

- Periodic kits issued to provide Tru64 UNIX support for new hardware options
  - Seven kits issued to date
  - Recent kits provide support for V5.1A and V5.1B only
- NHD-7 added support for DS15
  - Original NHD-7 included BL24 (V5.1B pk3, V5.1A pk6)
  - Current kit includes BL25 (pk4) for V5.1B
- NHD-6 added support for DS20L & several new devices
- Where to get NHD kits?
  - Factory installed on new systems
  - CD-ROM bundled with Tru64 UNIX media kit or available separately (part number QA-MT4AX-H8)
  - Search for “NHD” and Tru64 UNIX on ITRC patch search page

HP WORLD 2004

Solutions and Technology Conference & Expo

Co-produced by:

interex

shared knowledge • shared power

encompass

AN HP USER GROUP

RECOMMENDED TRAINING VENUE FOR THE

HP Certified Professional