

HP-UX 11.31 in Comparison with Previous HP-UX Releases

DECUS Presentation, 04/17/2007, Nürnberg
Authors: Victor Senderov
Zatina Yordanova



© 2006 Hewlett-Packard Development Company, L.P.
The information contained herein is subject to change without notice

Recapitulation of the new Features we've Discussed in this Presentation

- **Introduction**
- Web-based administration
- System Fault Management
- Interface Expansion/ Longer Hostnames
- Unified File Cache
- File System Enhancements
- vPars
- LVM/ MirrorDisk
- Dump/ Livedump
- Miscellaneous
- Next Generation Mass Storage Stack
- Cell OL*



About the Authors



Victor Senderov

*Studied
Computermathematics at
the Otto-von-Guericke
Universität in Magdeburg*



Zlatina Yordanova

*Student at the Technical
University in Sofia in
Informatics*

- HP-UX Support Team Members at the Global Delivery Support Center Sofia



Global Delivery Center Sofia

- HP-UX SW&HW, Microsoft, Data Protector
- Microsoft, Mass Storage, ProLiant Servers
- + Other teams = Totally 120 HP-employees
- Support teams provide support for DACH, also
- English speaking planned

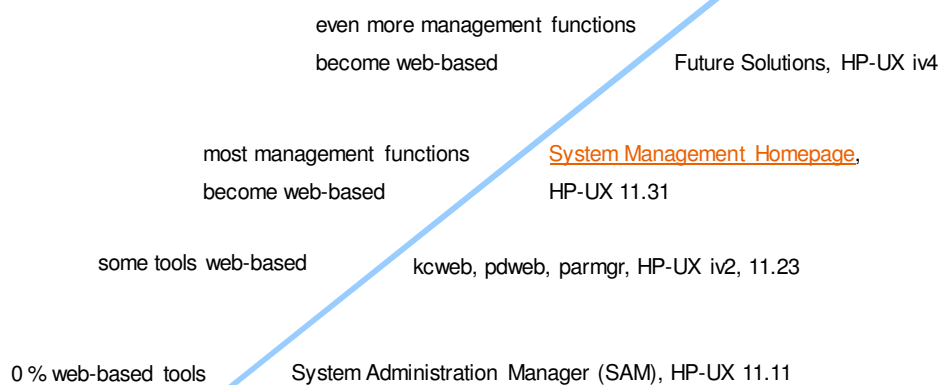


Agenda

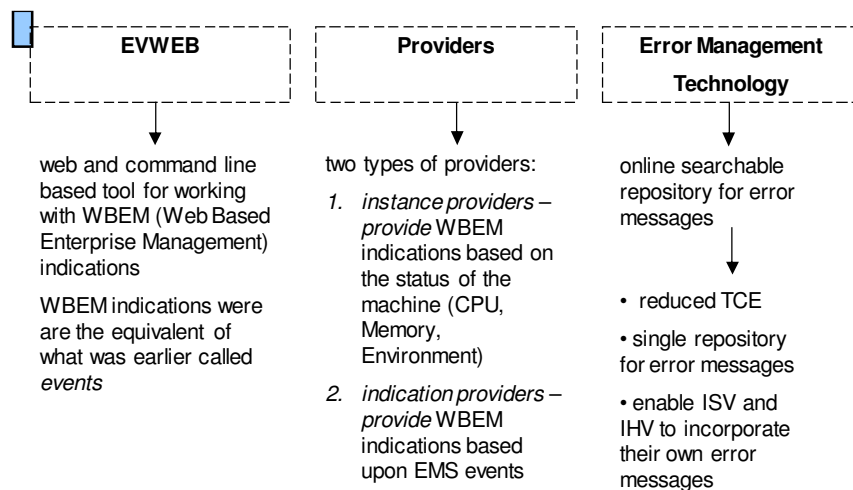
- Introduction
- **Web-based administration**
- **System Fault Management**
- Interface Expansion/ Longer Hostnames
- Unified File Cache
- File System Enhancements
- vPars
- LVM/ MirrorDisk
- Dump/ Livedump
- Miscellaneous
- Next Generation Mass Storage Stack
- Cell OL*



Web Based Administration



System Fault Management



Agenda

- Introduction
- Web-based administration
- System Fault Management
- **Interface Expansion/ Longer Hostnames**
- Unified File Cache
- File System Enhancements
- vPars
- LVM/ MirrorDisk
- Dump/ Livedump
- Miscellaneous
- Next Generation Mass Storage Stack
- Cell OL*



Longer Hostnames

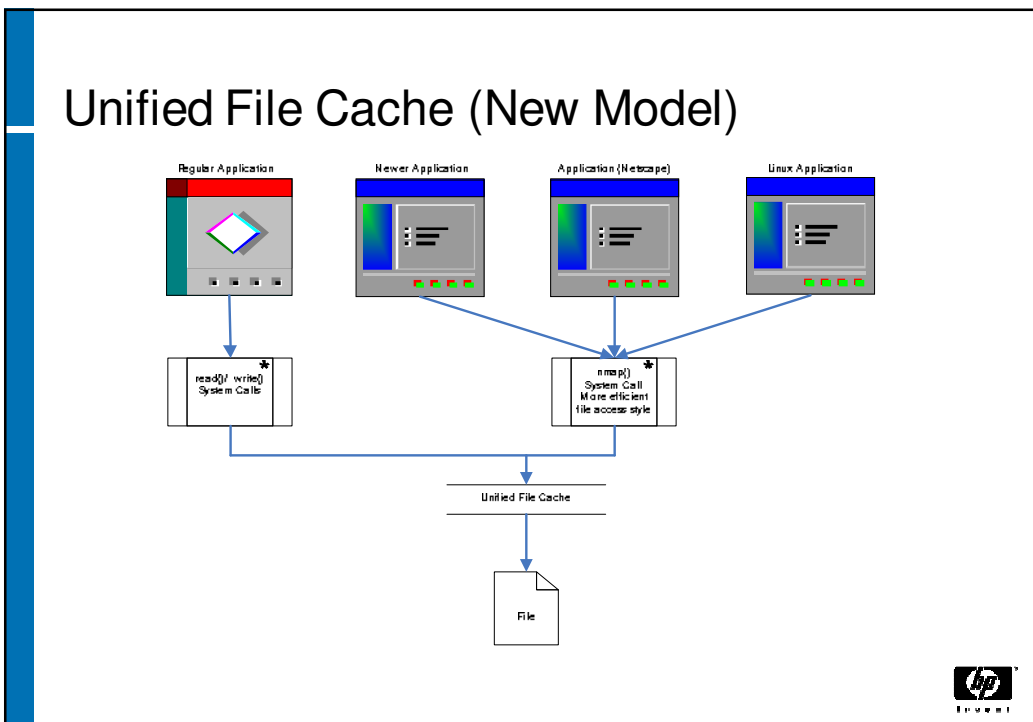

- Hostnames and node names now up to 255 characters long
- Existing software making use of the affected interfaces - `uname()` and `gethostname()` - will not need to be changed, however changes may be necessary to accommodate the new features



Agenda

- Introduction
- Web-based administration
- System Fault Management
- Interface Expansion/ Longer Hostnames
- **Unified File Cache**
- File System Enhancements
- vPars
- LVM/ MirrorDisk
- Dump/ Livedump
- Miscellaneous
- Next Generation Mass Storage Stack
- Cell OL*





Agenda

- Introduction
- Web-based administration
- System Fault Management
- Interface Expansion/ Longer Hostnames
- Unified File Cache
- **File System Enhancements**
- vPars
- LVM/ MirrorDisk
- Dump/ Livedump
- Miscellaneous
- Next Generation Mass Storage Stack
- Cell OL*



General File System Enhancements

- VxFS 4.1 now supports large file systems of up to theoretically 256 TB and large files of up to 16 TB
- VFS Stacking
 - fsdaemon
- VXFS, HFS and CDFS now use Unified File Cache
- UFC-aware and OL*-aware file cache tunables
 - filecache_max
 - filecache_min
- fadvise(2) interface and Posix.1 posix_fadvise(2) system interface



General File System Enhancements 2

- File system calls for TRU64 Migration
- `pselect(2)` provides Unix 2003 compliance for VFS
- Core FS Enhancements in support of Interface Expansion
- Core FS Enhancements in support of ONC+ 2.3
- VFS integration with HP-UX CIFS
- `/etc/mnttab` implemented as a pseudo device driver



General File System Enhancements 3

- Better handling of HFS write errors
- New improved multi-threaded file systems syncer
- Asynchronous I/O performance improvement for applications that use `aio_reap(2)`
- `pax(1)` supports multi-TB files and long filenames
- Long pathnames cached in DNLC
- HFS `fsckclean(1M)` command performs faster



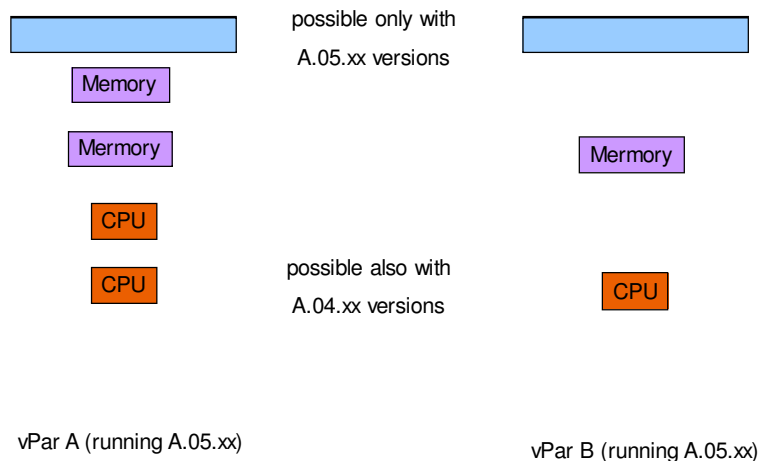
Agenda

- Introduction
- Web-based administration
- System Fault Management
- Interface Expansion/ Longer Hostnames
- Unified File Cache
- File System Enhancements
- **vPars**
- LVM/ MirrorDisk
- Dump/ Livedump
- Miscellaneous
- Next Generation Mass Storage Stack
- Cell OL*



vPars A.05.xx

1/2



vPars A.05.xx

2/2

- Other new features:
 - hyperthreading
 - mixed 11iv2/v3 environment
 - vPars A.04.xx (xx ≥ 02) and vPars A.05.xx can coexist
 - vPars A.05.01 Monitor must be running
 - firmware requirements follow vPar A.05.xx



vPars A.05.xx

3/3

Mixed Environment

Hyperthreading is ...	Monitor Booted is...	Monitor Boot Result	Virt Partition booted	Virt. Partition result
OFF	A.05.01	OK	A.05.01	OK
	A.04.02	OK	A.04.02 A.05.01 A.04.02	OK Fails OK
ON	A.05.01	OK	A.05.01 A.04.02	OK Fails
	A.04.02	FAIL	N/A N/A	N/A N/A



Agenda

- Introduction
- Web-based administration
- System Fault Management
- Interface Expansion/ Longer Hostnames
- Unified File Cache
- File System Enhancements
- vPars
- **LVM/ MirrorDisk**
- Dump/ Livedump
- Miscellaneous
- Next Generation Mass Storage Stack
- Cell OL*



Logical Volume Manager New Features

- New online features:
 - Online resizing of a LUN
 - `vgmodify`
 - LVM Device Online Replacement (OLR)
 - `pvchange -a`
 - Volume Group Quiesce/Resume
 - `vgchange [-Q | -R]`
- Mass Storage Stack Support
 - new style DSF are supported
- SLVM Single Node Online Volume Reconfiguration (SNOR)
 - `vgchange -x`
- Dynamic Volume Group Modification
 - `vgmodify`
- Boot resiliency



Logical Volume Manager Enhancements

- Increased logical volume size
 - lvols up to 16 TB
- Display enhancements
 - `lvdisplay`, `pvdisplay`, `vgdisplay`, and `vgscan` commands all support long hostnames
 - new “-F” option for easier parsability
 - `pvdisplay -l`
- `vgscan [-f | -k]`



LVM Deprecated or Obsolete Features

- Bad block reallocation obsolete
- `maxvgs` deprecated



MirrorDisk UX

- Striped Mirrors with the entire range of stripe sizes



Agenda

- Introduction
- Web-based administration
- System Fault Management
- Interface Expansion/ Longer Hostnames
- Unified File Cache
- File System Enhancements
- vPars
- LVM/ MirrorDisk
- **Dump/ Livedump**
- Miscellaneous
- Next Generation Mass Storage Stack
- Cell OL*



Livedump

- Uses of livedump
 - Analysis of recoverable operating system faults
 - Failures with DLKM that do not destabilize the kernel (for example, I/O drivers)
 - Performance analysis
 - Analysis of a snapshot of a running kernel.
- See the man page `livedump(1M)` for details



New and Improved Crashdump

- Performance Improvements (parallelization)
- Configuration Improvements
- Availability and Manageability Improvements
- Further Information can be found in
 - “HP-UX 11i v3 Crash Dump Improvements,”
 - Residing on <http://docs.hp.com>




Agenda

- Introduction
- Web-based administration
- System Fault Management
- Interface Expansion/ Longer Hostnames
- Unified File Cache
- File System Enhancements
- vPars
- LVM/ MirrorDisk
- Dump/ Livedump
- **Miscellaneous**
- Next Generation Mass Storage Stack
- Cell OL*



Some New Features That We Would Just Like To Mention

- Universal Disk Format
- Dynamic Root Disk
- ProcSM
- Bigger PIDs and NPROC (256,000)
- I/O Infrastructure
- InfiniBand



Effective theoretical throughput in different configurations

	single	double	quad
1X	2 Gbit/s	4 Gbit/s	8 Gbit/s
4X	8 Gbit/s	16 Gbit/s	32 Gbit/s
12X	24 Gbit/s	48 Gbit/s	96 Gbit/s



Agenda

- Introduction
- Web-based administration
- System Fault Management
- Interface Expansion/ Longer Hostnames
- Unified File Cache
- File System Enhancements
- vPars
- LVM/ MirrorDisk
- Dump/ Livedump
- Miscellaneous
- **Next Generation Mass Storage Stack**
- Cell OL*



Legacy and Persistent DSFs

- Until HP-UX 11i v2 - **Legacy DSFs**
- Path dependant
- Format: /dev/(r)dsk/cxydz
- Impact of Hardware Changes on the system
- In HP-UX 11i v3 - **Persistent DSFs**
- Path independent
- DSFs bound to WWID (World Wide Identifier)
- Format: /dev/(r)disk/diskX



11.31 Mass Storage Stack

- Migration from Secure Path to Mass Storage Stack Native Multipathing
- Auto-detection and configuration of new LUNs & LUN Paths.
- Persistent DSFs.
- Load Balancing configured on all paths.
- Zero User configuration required.
- Performance
- Example ioscan – new Options
- scsimgr
- msv2v3check – a script to check whether the storage devices used are supported



Native Multipathing in LVM

- Alternate and Primary Paths
- alternate links -> up to 8, in native multipathing -> up to 32
- Persistent DSFs supported
- Legacy DSFs available for backward compatibility
- Migration variations
- vgdsf
- ioscan – new Options
- iobind
- io_redirect_dsf



Agenda

- Introduction
- Web-based administration
- System Fault Management
- Interface Expansion/ Longer Hostnames
- Unified File Cache
- File System Enhancements
- vPars
- LVM/ MirrorDisk
- Dump/ Livedump
- Miscellaneous
- Next Generation Mass Storage Stack
- **Cell OL***

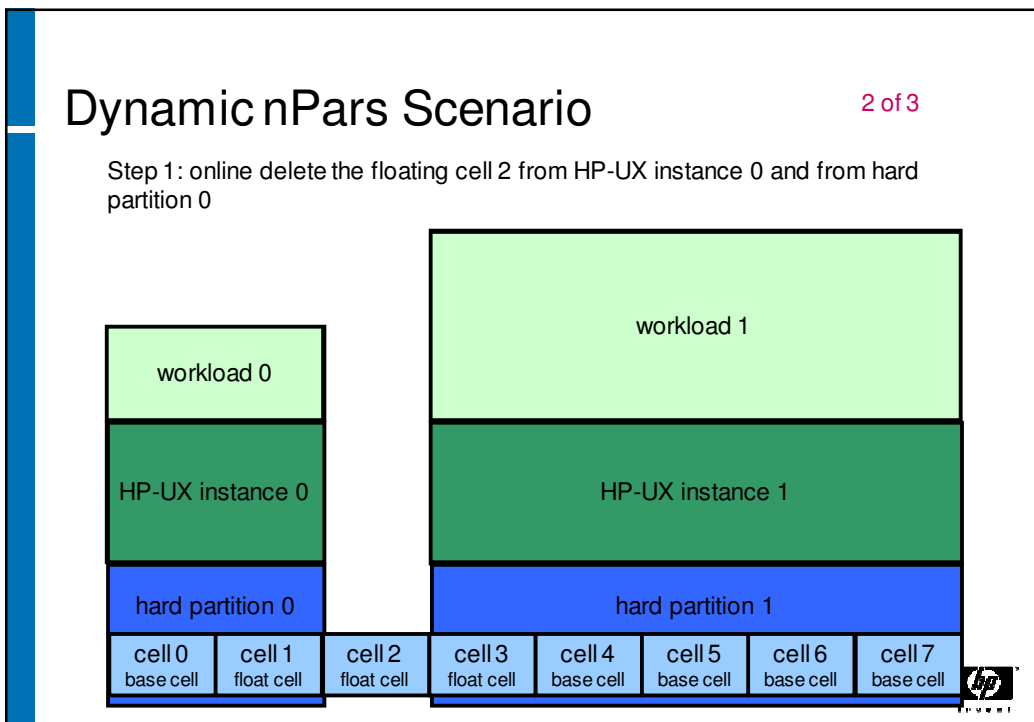
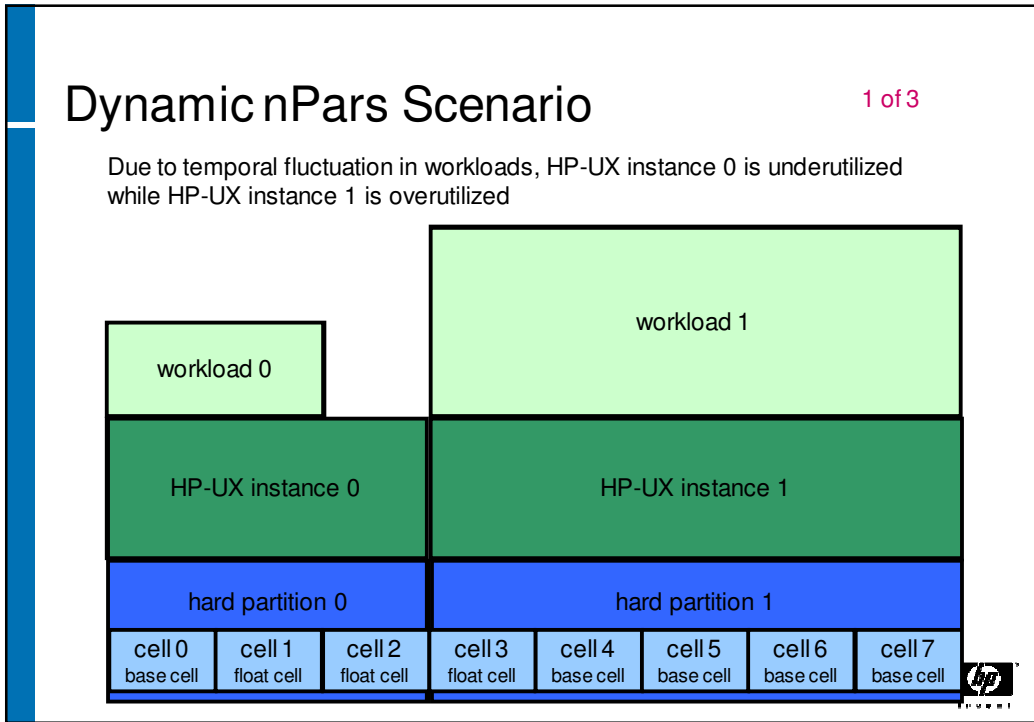


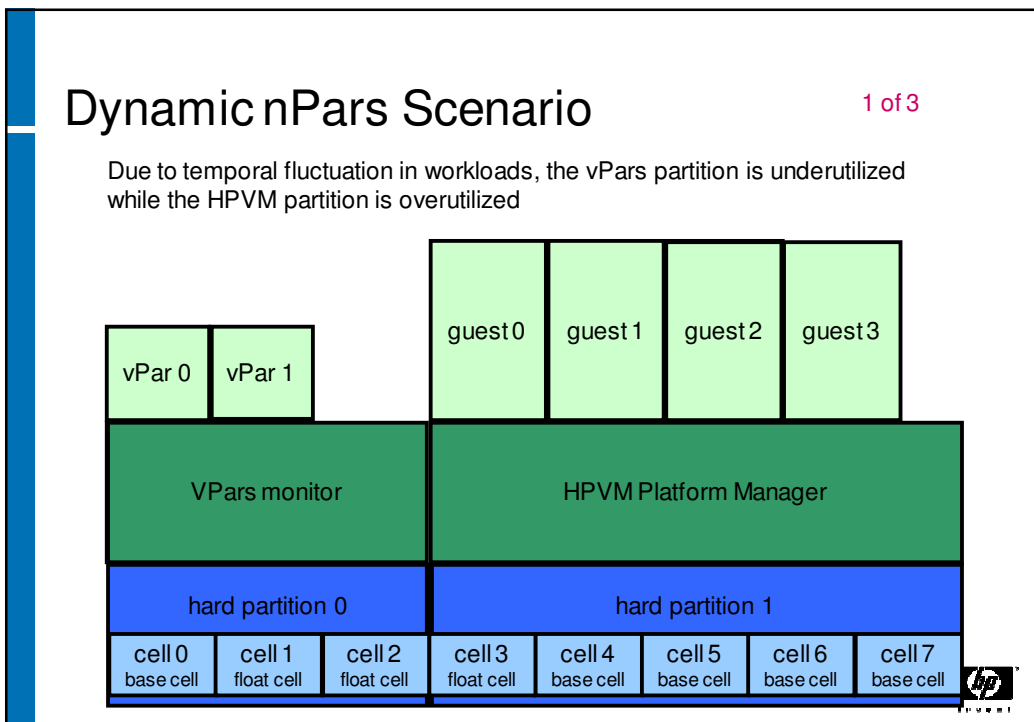
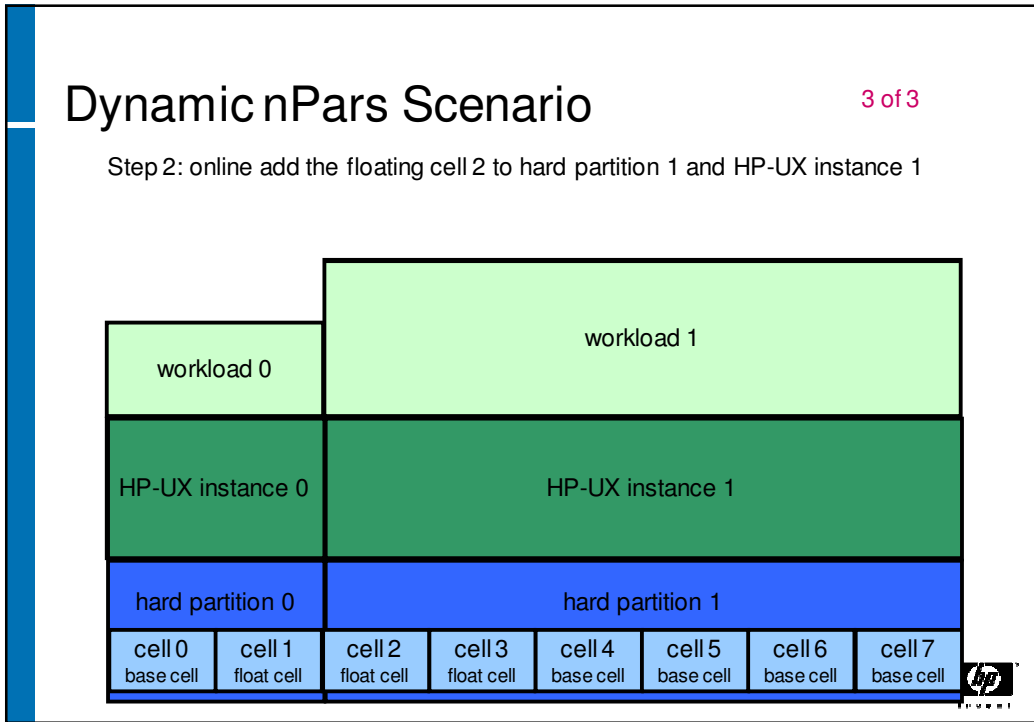
Cell OL* Functionality in 11.31

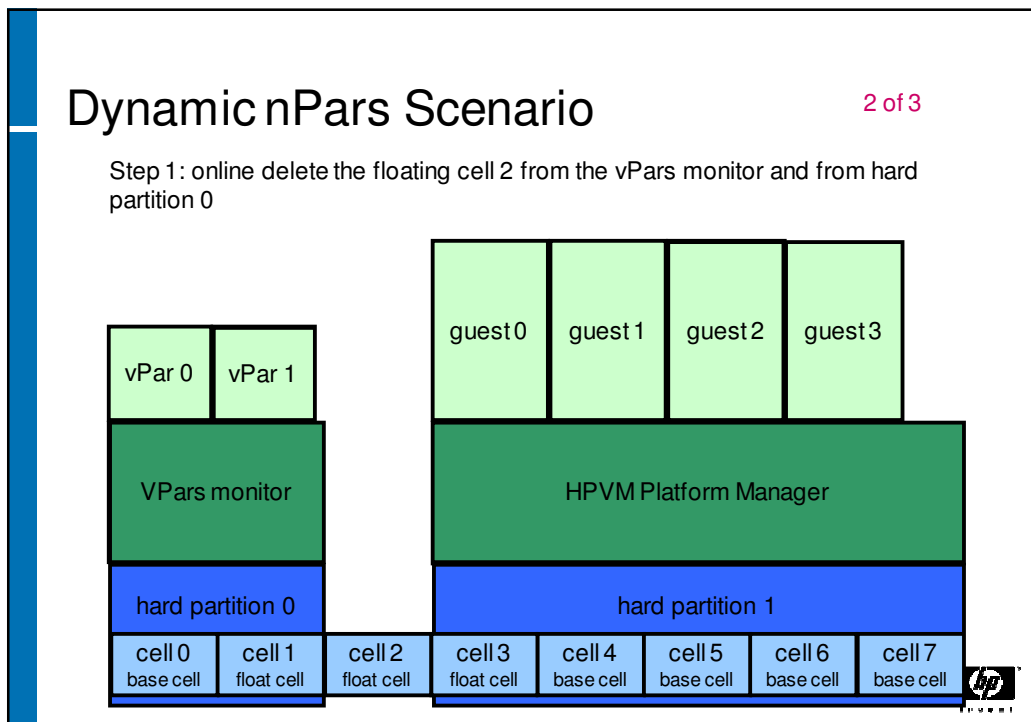
1/2

- Will Come with Sept. 2007 OE
- Cell Online Activation (Cell OLA)
- Cell Online Deactivation (Cell OLD)
 - using “float cells” dynamic migration between nPars possible
- Cell Online Replacement (Cell OLR = OLD + OLA)
 - HA increased via replacing troublesome HW before it fails
 - in case of HW failure second reboot not needed









Recapitulation

- Introduction
- Web-based administration
- System Fault Management
- Interface Expansion/ Longer Hostnames
- Unified File Cache
- File System Enhancements
- vPars
- LVM/ MirrorDisk
- Dump/ Livedump
- Miscellaneous
- Next Generation Mass Storage Stack
- Cell OL*

