

Exadata Patching

or, “How I Learned to Stop Worrying and Love Exadata Patching”

Andy Colvin, Enkitec

<http://blog.oracle-ninja.com>

About Me

- Background in system/network administration
- Working with Oracle for 12 years
- Began working on Exadata 16 months ago
- Involved in 10 separate Exadata implementation projects (24 Exadata racks)
- Patched nearly all of them with different requirements and versions

Agenda

- **Patching Overview**
- Patch Application Process
- Mitigating Risk with Patching
- Patching “Gotchas”

Exadata Patching Overview

- Holy Grail of Exadata notes – MOS #888828.1
 - Lists current and previous versions for 11.2 software
 - General rule is to only apply patches listed in this note
- Patches apply to 3 stacks
 - Database/Clusterware – Bundle Patches
 - Operating System/Firmware – Exadata Storage Server Patches
 - Infiniband – Infiniband Switch Patches
- KVM and Cisco switch
 - Oracle does not provide patches directly

Bundle Patches

- Affect Database and Grid Infrastructure (GI) Homes
- Released Monthly (depending on software version)
- Contain standard Oracle patches and Exadata-specific patches
- Installed using OPatch

```
[oracle@enkdb01 ~]$ opatch lsinventory
Interim patches (3) :

Patch 12332686      : applied on Wed Jul 20 10:33:34 CDT 2011
Unique Patch ID: 13789775
Created on 26 May 2011, 04:53:25 hrs PST8PDT
Bugs fixed:
12332686, 10626132
```

Exadata Storage Server Patches

- One patch includes updates for OS, kernel, IB, ILOM, other firmware, new features
- Includes “minimal pack” for database servers – formerly called “convenience pack”
- Applied using patchmgr

```
[root@enkcel01 ~]# imageinfo

Kernel version: 2.6.18-194.3.1.0.4.el5 #1 SMP Sat Feb 19
03:38:37 EST 2011 x86_64
Cell version: OSS_11.2.0.3.0_LINUX.X64_110520
Cell rpm version: cell-11.2.2.3.2_LINUX.X64_110520-1

Active image version: 11.2.2.3.2.110520
Active image activated: 2011-06-02 13:41:55 -0500
Active image status: success
Active system partition on device: /dev/md6
Active software partition on device: /dev/md8
```

Infiniband Switch Patches

- Only install patches provided by Oracle
- Only install versions mentioned in MOS note #888828.1
- Various installation methods depending on the version

```
[root@enksw-ib2 ~]# version  
SUN DCS 36p version: 1.3.3-2  
Build time: Apr  4 2011 11:15:19
```

Patch Release Cycles

- Bundle Patches
 - 11.2.0.1 – released every 2 months (through Sept. 2011)
 - 11.2.0.2 – released monthly (through Oct. 2011, then every 2 months)
 - 11.2.0.3 (when released) – monthly
- Exadata Storage Server Patches
 - Released quarterly
- Infiniband Patches
 - Semi-annually to annually

Agenda

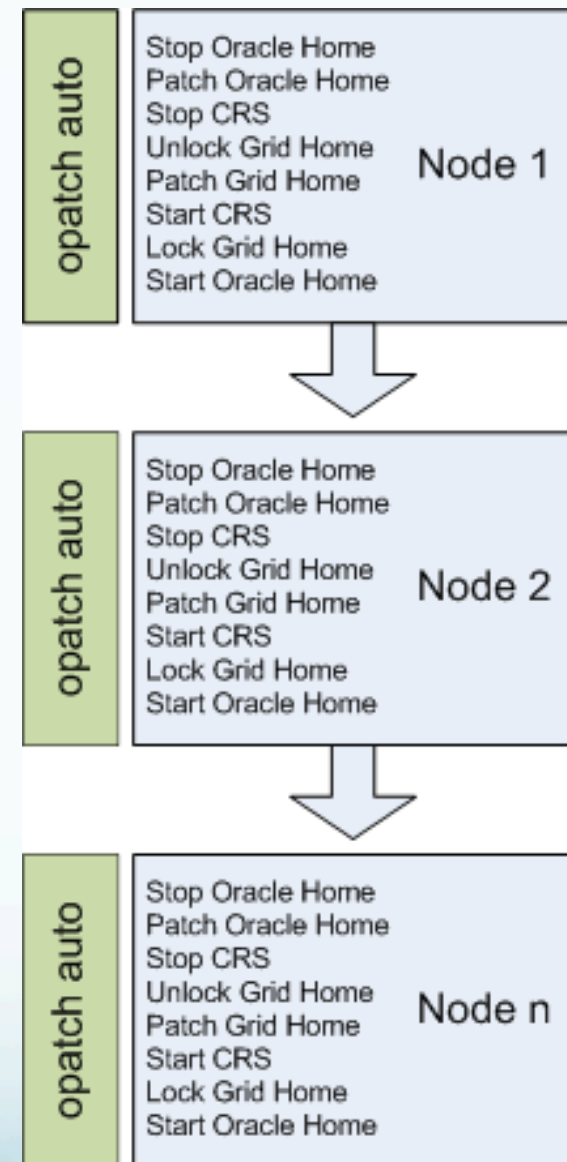
- Patching Overview
- **Patch Application Process**
- Mitigating Risk with Patching
- Patching “Gotchas”

Bundle Patch Application

- Rolling patches
- Installed using OPatch
 - 11.2.0.2 uses `opatch auto` exclusively
 - Allows administrators to issue one `opatch` command to patch all Oracle homes on the server
 - After the first server is patched, repeat the process on remaining servers
 - 11.2.0.1 uses a combination of `opatch apply` and `opatch auto` depending on type of patch

OPatch Auto

- OPatch automation installs the bundle patch with one command per compute node
- Only patches Oracle homes that have a database registered
- To patch specific Oracle homes, add the `-oh <ORACLE_HOME>` flag to the command
- Requires that OPatch run with root privileges



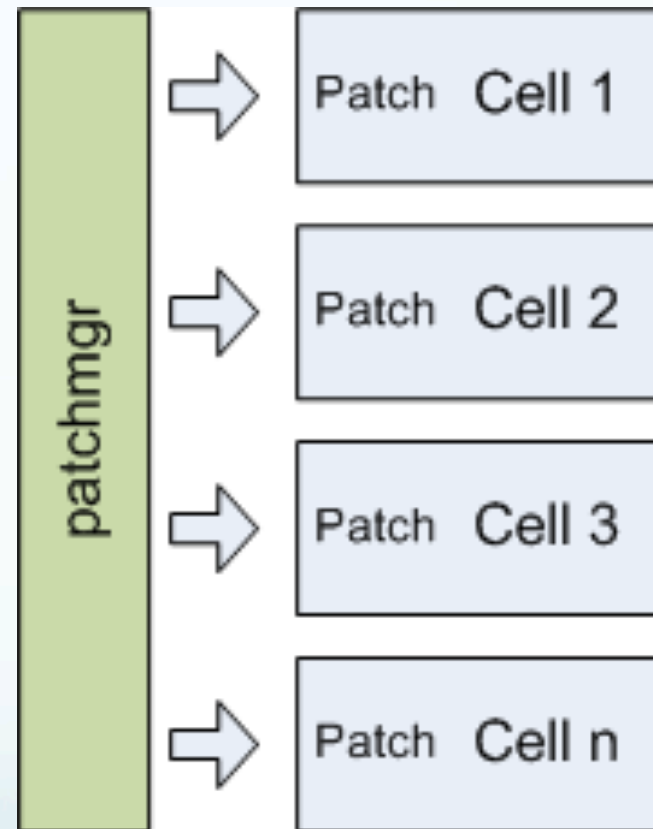
Storage Server Patches

- Installed using patchmgr
- Include operating system updates, firmware updates, and new features
- patchmgr utilizes dcli to push patch software to storage cells
- Can either be installed rolling or non-rolling

Storage Server Patches

Non-rolling patch update

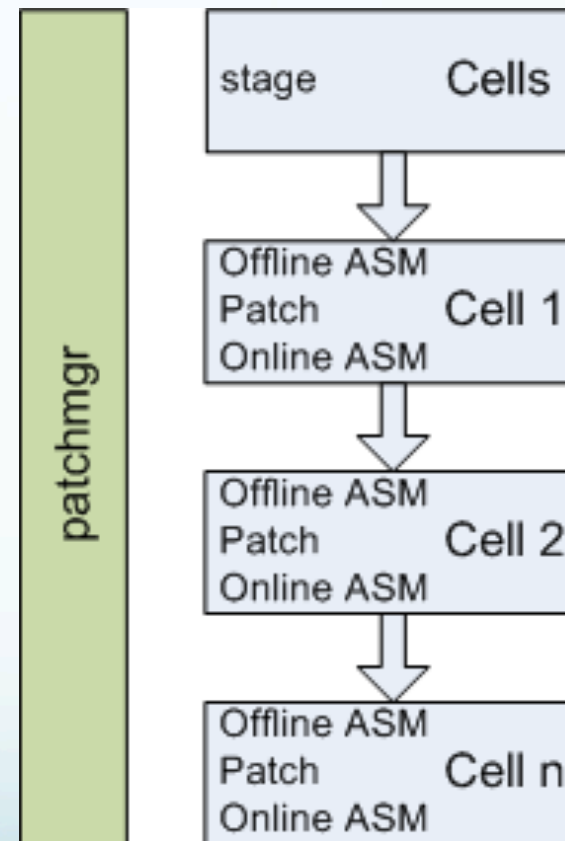
- Benefits
 - Shorter patch window
 - Between 1.5 and 3 hours for entire process
- Considerations
 - Full outage during patch apply process
 - If patch fails, all cells can be affected



Storage Server Patches

Rolling patch update

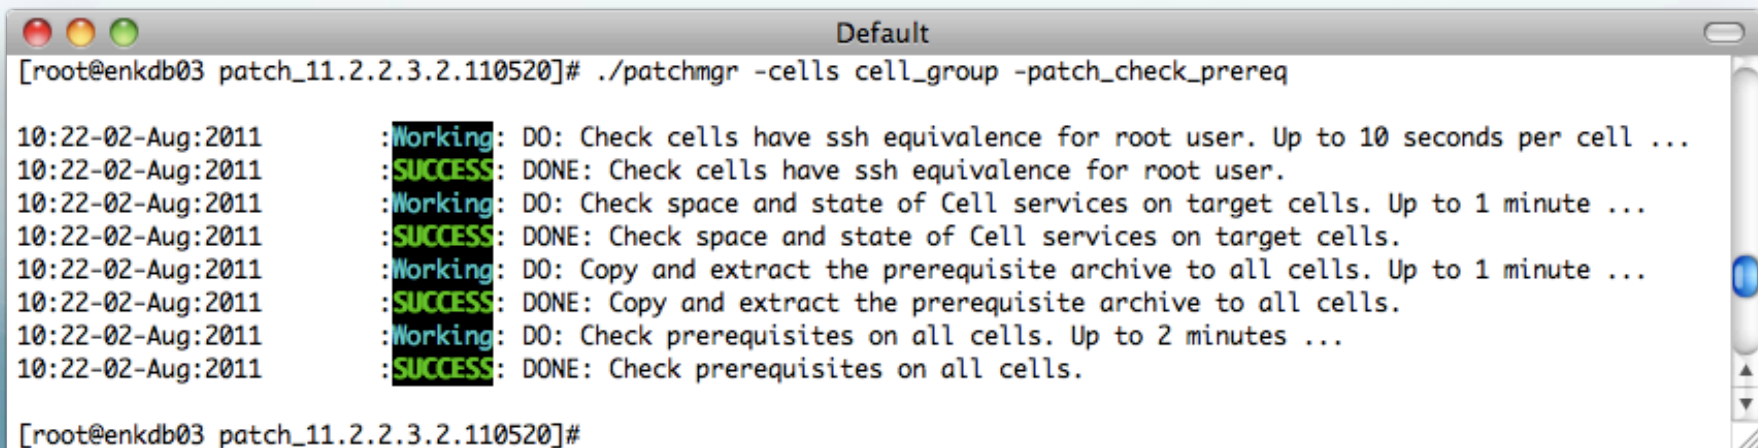
- Benefits
 - No downtime
 - If patch fails, only one cell is affected
- Considerations
 - Longer patch apply time
 - Between 1.5 and 3 hours per cell for patch apply
 - Reduced redundancy during patch apply process
 - Minimum BP requirements



Storage Server Patches

Applying the patch

1. Download and stage patch on compute node
2. Unpack patch and check SSH connectivity `dcli -l root -g cell_group hostname -a`
3. Run patch prerequisite check before applying `./patchmgr -cells cell_group -patch_check_prereq [-rolling]`
4. Apply patch `./patchmgr -cells cell_group -patch [-rolling]`



```
Default
[root@enkdb03 patch_11.2.2.3.2.110520]# ./patchmgr -cells cell_group -patch_check_prereq

10:22-02-Aug:2011      :Working: DO: Check cells have ssh equivalence for root user. Up to 10 seconds per cell ...
10:22-02-Aug:2011      :SUCCESS: DONE: Check cells have ssh equivalence for root user.
10:22-02-Aug:2011      :Working: DO: Check space and state of Cell services on target cells. Up to 1 minute ...
10:22-02-Aug:2011      :SUCCESS: DONE: Check space and state of Cell services on target cells.
10:22-02-Aug:2011      :Working: DO: Copy and extract the prerequisite archive to all cells. Up to 1 minute ...
10:22-02-Aug:2011      :SUCCESS: DONE: Copy and extract the prerequisite archive to all cells.
10:22-02-Aug:2011      :Working: DO: Check prerequisites on all cells. Up to 2 minutes ...
10:22-02-Aug:2011      :SUCCESS: DONE: Check prerequisites on all cells.

[root@enkdb03 patch_11.2.2.3.2.110520]#
```

Storage Server Patches

Patch application details

- Cells utilize software RAID partitions on the first 2 physical disks
- RAID-1 partitions are created for swap, /boot, /, /opt/oracle, and /var/log/oracle.
- / and /opt/oracle have active and inactive partitions
- Patch is pushed to current inactive partition, then active/inactive partitions are switched

/dev/md2 - swap
/dev/md4 - /boot
/dev/md5 - / (active)
/dev/md6 - / (inactive)
/dev/md7 - /opt/oracle (active)
/dev/md8 - /opt/oracle/ (inactive)
/dev/md11 - /var/log/oracle

Storage Server Patches

Patch apply process

1. Patch contents are copied to `/root/_patch_hctap_`
2. New operating system image is pushed to `/` and `/opt/oracle`
3. Cell reboots
4. During cell reboot, `exachkcfg` service determines if any new firmware updates need to be installed
5. Cell reboots several times to finalize firmware updates
6. Post-installation validation checks are run against the cell
7. Cell usb recovery media is updated to the latest version

Storage Server Patches

Compute Node “Minimal” Packs

- Installed on database servers after cells are patched
- Include Infiniband/OFED driver updates, kernel patches, and firmware updates for hardware
- Minimal packs do not install operating system updates for the database servers
- Generally involve only one reboot (storage server patches reboot the cell several times)
- Are applied one node at a time to allow for a rolling patch

Storage Server Patches

Installing Compute Node “Minimal” Packs

- Unzip the db_patch_<version>.zip on each compute node
- Run ./install.sh script on one node, and wait for the node to reboot and come back up
- Repeat the process on the remaining compute nodes



```
root@enkdb03:/u01/stage/patches/patch_11.2.2.3.2.110520/db_patch_11.2.2.3.2.110520
[root@enkdb03 db_patch_11.2.2.3.2.110520]# ls
11.2.2.3.2.110520.tbz  filesort.pl      image_functions      install.sh  README.txt
dopatch.sh          ibdiagtools.tar  image.id.11.2.2.3.2.110520.in  osw.tbz
[root@enkdb03 db_patch_11.2.2.3.2.110520]# ./install.sh
NOTE This db host may reboot during the minimal pack application process.
NOTE Ensure all Oracle software components are shut down.

WARNING Do not start more than one instance of installation.
WARNING Do not interrupt the installation.
WARNING Do not open log files in editor in write mode or try to alter them.

NOTE You may interrupt this run in next 60 seconds with control-c.

[root@enkdb03 db_patch_11.2.2.3.2.110520]#
```

Infiniband Switch Patches

- Infiniband switches are running CentOS Linux 5.2
- Infiniband patches can be installed in a rolling fashion
 - Patch the first switch, and wait for reboot to complete, then repeat patch process on remaining switches
- Older versions (1.1.3) were installed by placing update files on a web or FTP server, and downloading RPM package updates to the switch
- Later versions (1.3.3) were installed by placing the update package on the filesystem of the switch, and updating from the ILOM of the IB switch

Agenda

- Patching Overview
- Patch Application Process
- **Mitigating Risk with Patching**
- Patching “Gotchas”

Mitigating Patching Risks

- Patching was originally very scary, but new features are making it better.

Feature	Minimum Version
out of partition cell upgrades	11.2.1.3.1
DB+GI BP Merge	11.2.0.2
BP OPatch Auto Install	11.2.0.2 BP2/11.2.0.1 GI BP4
BP OEM Install	11.2.0.2 BP1/11.2.0.1 BP7
BP DG Standby-First Install	11.2.0.2 BP1/11.2.0.1 BP8
Oplan	11.2.0.2

Mitigating Patching Risks

Data Guard Standby First Patch Apply

- Patches can be applied to the standby database before the primary database
- Available in 11.2.0.2 BP1+ and 11.2.0.1 BP7+

Mitigating Patching Risks

Data Guard Standby First Patch Apply Process

1. Apply patch to standby server
2. Test patch using snapshot standby or active data guard
3. Switchover to standby
4. Apply patch to primary database
5. Switchover back to original primary
6. Run post-patch install scripts

Mitigating Patching Risks

OPlan

- OPlan is available for 11.2.0.2 bundle patches
- Simplifies the patching process and gives step by step instructions tailored to your environment
- Covers both patch application and rollback, including out of place patching

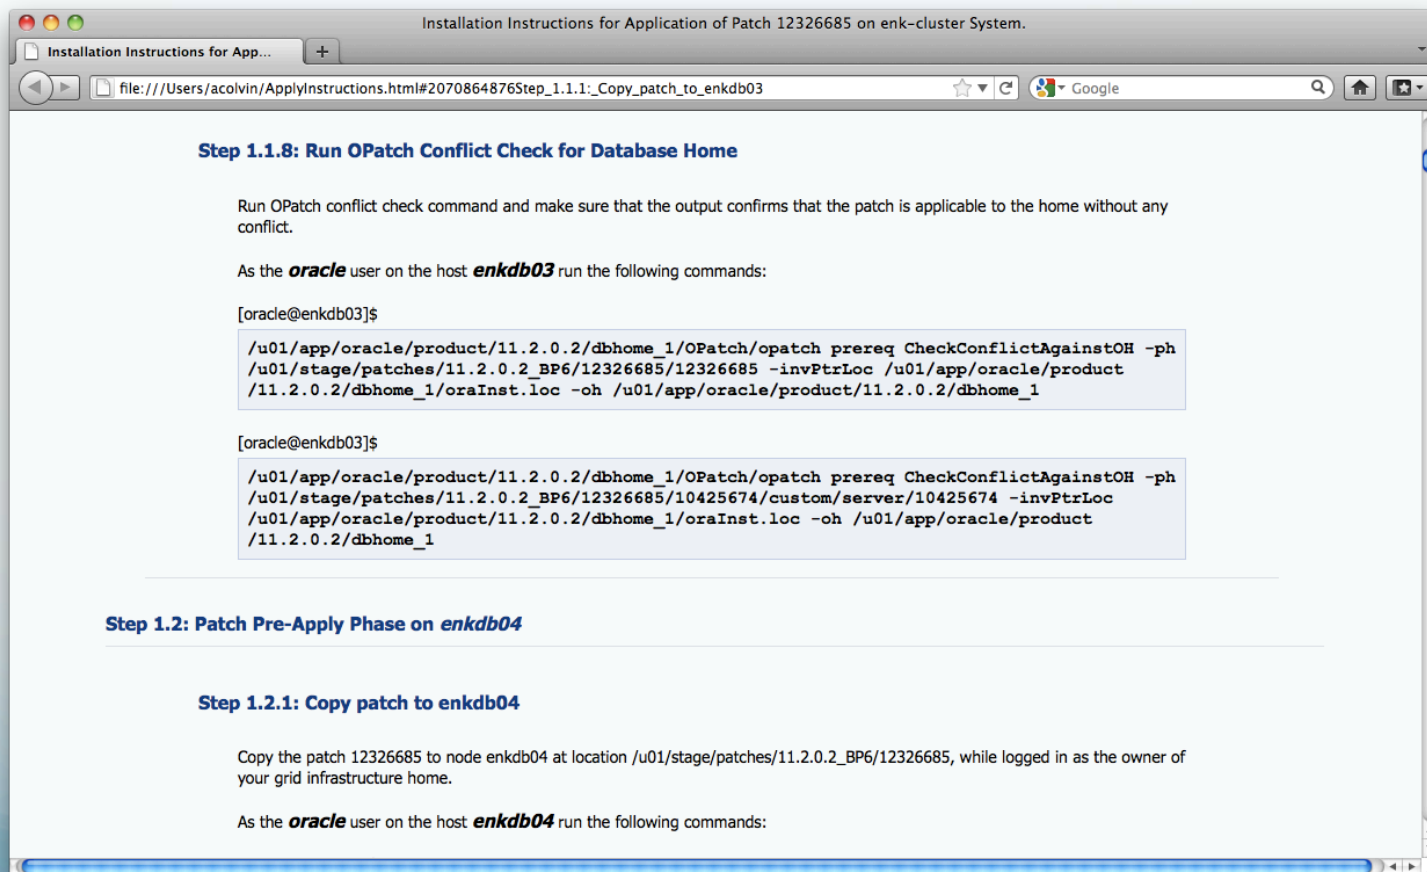
Mitigating Patching Risks

Running OPlan

- Download OPlan to your Oracle home
- Issue the command `oplan generateApplySteps <patch location>`
- Oplan will generate an HTML file with specific installation steps for your system

Mitigating Patching Risks

Oplan Output



Agenda

- Patching Overview
- Patch Application Process
- Mitigating Risk with Patching
- **Patching “Gotchas”**

Patching “Gotchas”

- No 2 patch applications are alike
- Ensure that you have tested before patching production
- For storage server patches, ensure that you have a good connection that will not drop during the entirety of the patch process
 - Use either VNC or screen to connect to the server you’re patching from
- Always read through the README several times
 - Read through the “Known Issues” sections to be prepared for what may go wrong – 11.2.2.3.2 software release has 20 known issues listed in the README
 - Exadata storage server patches have a README and a supplemental note (11.2.2.3.2 has README and note #132958.1)

Patching “Gotchas”

Horror stories

- Minimal pack installation failed on one database server, and it wouldn't boot up completely
 - Messages were received regarding missing kernel modules, and the system was not able to bring up network interfaces
 - Only access was through KVM or ILOM's serial console
 - Discovered that the minimal pack had wiped out the listing of kernel modules in `/lib/modules/<kernel_version>` before backing it up
 - Solved the issue by formatting a USB stick with ext3 and copying modules over from the surviving compute node
- Database servers now back up these files before applying a minimal pack

Patching “Gotchas”

Horror stories

- Cell would not start back up with correct IP address entries
 - After applying the Exadata storage software patch, the cell would start up with the IP address settings from the factory
 - Manually recreating the network interfaces would not persist after a reboot
 - Issue was traced back to a corrupted `/opt/oracle.cellos/cell.conf` file. After recreating this file with the correct values, the network settings stayed correct after reboots
- Oracle now recommends in the patch README to run `ipconf -dry` before issuing a storage server patch

Patching “Gotchas”

Horror stories

- ILOM would not come back online after firmware patch
 - After applying the Exadata storage software patch, one cell did not complete the patch, and the ILOM was unresponsive
 - After connecting directly to the ILOM using the serial connection, it was determined that the ILOM was left in rescue mode
 - The ILOM was flashed up to the newer version, and flashed back down to the original firmware version
 - A cold reset command was sent to the ILOM
 - Patch installation was run again for the single cell that had failed
- One of the bugs being fixed by this version of the firmware was a memory leak in the ILOM software
 - Oracle now sends a cold reset to the ILOM before patching

Questions?

Andy Colvin, Enkitec

Email: andy.colvin@enkitec.com

Blog: <http://blog.oracle-ninja.com>

Twitter: @acolvin