Ksplice – Is Rebooting Your Oracle Linux Database Server Now Obsolete?

Robert Bialek

Principal Consultant

Who Am I

- Principal Consultant, Partner and Trainer at Trivadis GmbH in Munich
  - robert.bialek@trivadis.com
- Main Focus: Oracle Database High Availability
  - High availability architecture design
  - Review, troubleshooting, coaching
  - Backup and recovery
  - Performance tuning
  - Linux administration
- Trainer for the following Trivadis courses
  - Oracle Grid Infrastructure (O-GRINF)
  - Oracle Real Application Cluster (O-RAC)
  - Oracle Data Guard (O-DG)

ORACLE Certified Master
Oracle Database 10*g* Administrator

ORACLE Certified Master
Oracle Database 11*g* Administrator

Brief introduction of Trivadis

Trivadis is a market leader in IT consulting, system integration, solution engineering and the provision of IT services focusing on ORACLE and Microsoft technologies in Switzerland, Germany and Austria.

We provide our services in the following strategic business areas:.

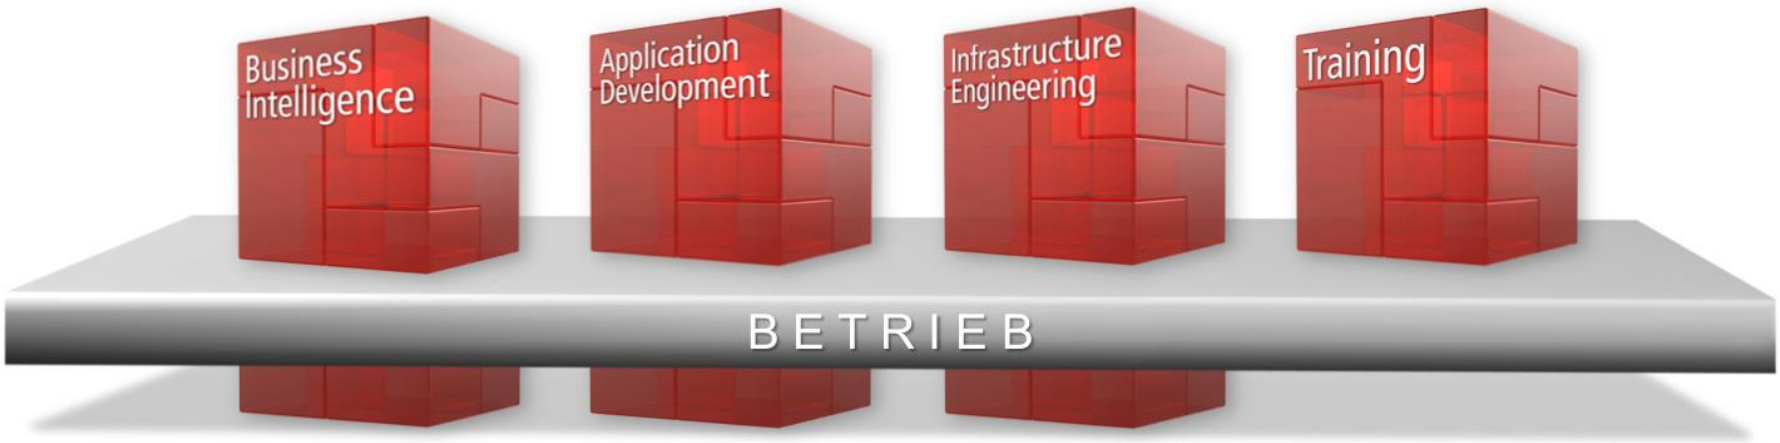

Our training services guarantee transfer of know-how.

With over 600 IT experts and specialists on site for you

11 Trivadis branches with over 600 employees

200 service level agreements

More than 4000 training participants

Research and development budget: CHF 5.0 /EUR 4 million

Financially independent and consistently profitable

Experience from more than 1900 projects per year for over 800 customers

12/2012

AGENDA

1. Introduction
2. How it Works
3. Installation and Configuration
4. Patch Management
5. Oracle Database Environment
6. Core Messages

OS Patch Management

- In most cases a **reactive** OS patching policy will be applied
  - *“Do not touch a running system”*
  - *“We will patch in case we hit an issue”*
  - *“We will upgrade if the application compatibility requires it”*
- Do I need to care about security ?
  - If your OS has been compromised, it might be easy to compromise the applications too (e.g. databases)
  - Most of the attacks are not 0-day and base on common vulnerabilities
- **Better**, switch to **proactive** OS patching and try to keep up your Linux server with the security alerts

OS Patching – User Space Applications

- Most of the Linux maintenance patches can be applied **online**
  - *Graceful switchover* between different program versions
- How is it possible?
  - Different versions (inodes) of a program binary on a disk **and** in memory
  - *New* processes will use the new patched program version (inode)
  - *Old* processes will still use the *old* version (inode)

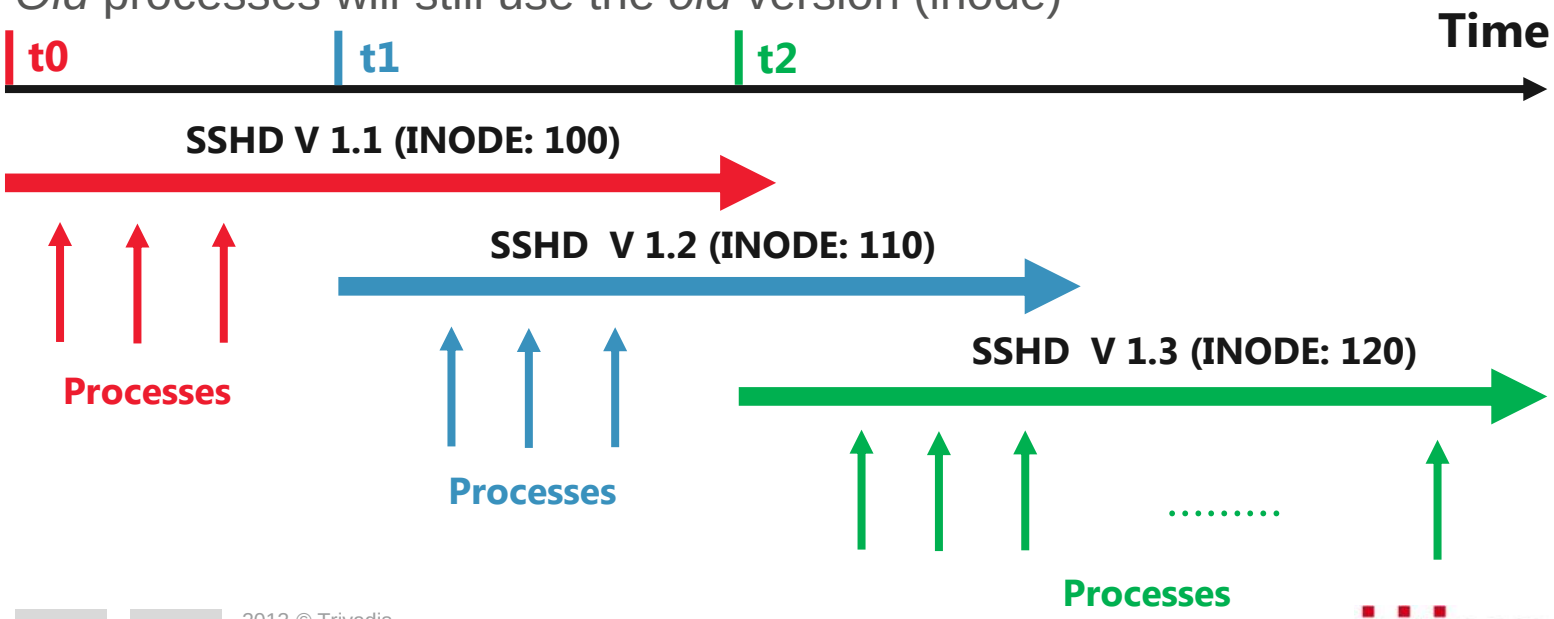

OS Patching – Kernel (The Problem)

- OS kernel is **always** an exception
  - A new patched kernel version will be installed online on a disk but **not loaded** into memory
  - To activate the new patched version a reboot is required
  - Loss of all software states on that server

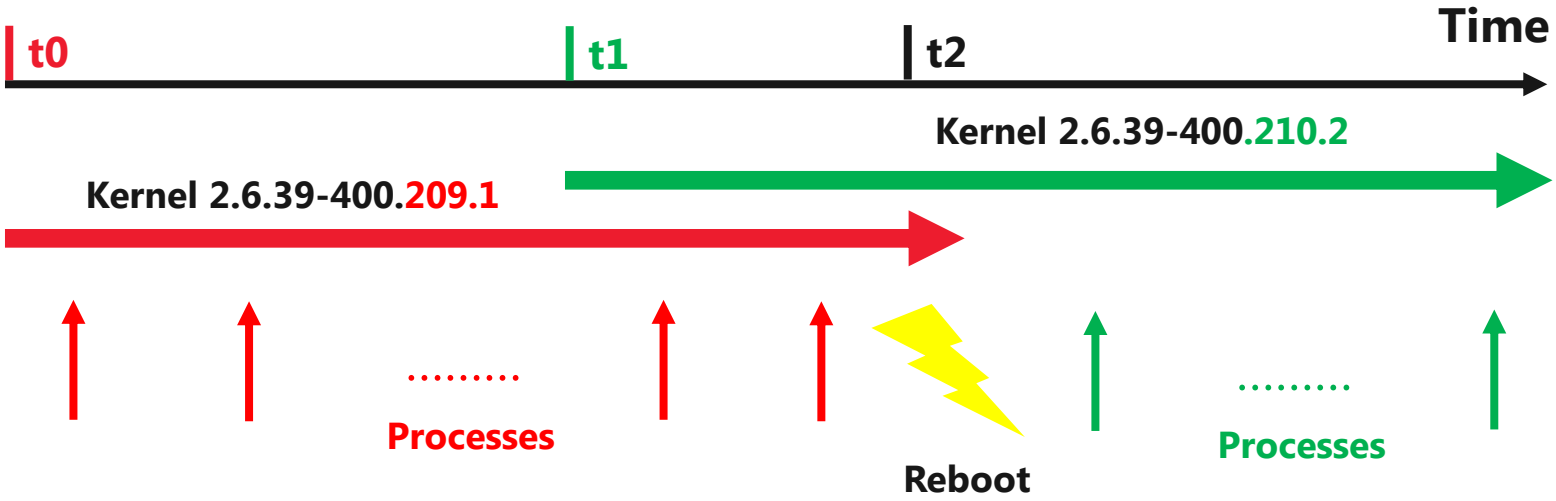

- On average, there is one ELSA/ELBA per month
  - To stay up-to-date you need to reboot on average **once per month**

OS Patching – Kernel (The Solution ?)

- Ksplice, Inc. software company was founded in 2008
  - Goal ➔ **apply important kernel patches online using Ksplice Uptrack**
  - How ➔ patch the running kernel directly in memory
  - Initially prebuilt kernel patches on a subscription basis
  - Free of charge for Ubuntu and Fedora

- On July 21, 2011 Oracle Corporation acquired Ksplice
  - A standard feature of the Oracle Linux Premier Support level subscription
  - Supported all OEL 6 and most of the OEL 5 kernels
  - Not supported for other enterprise Linux distributions (RH 30-days trial)

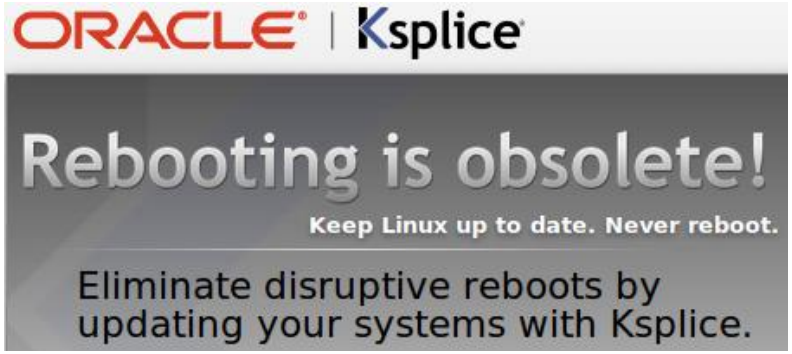

AGENDA

1. Introduction
2. How it Works
3. Installation and Configuration
4. Patch Management
5. Oracle Database Environment
6. Core Messages

Ksplice Patches – Overview

- The core functionality base on dynamically loadable **kernel modules**
  - Very flexible approach
  - Allows not only applying but also rolling back patches online

- Ksplice patches can be used to
  - Efficiently fix security or critical kernel bugs
  - Create an online kernel diagnostic patch
- Most of the patches can be created, without the developer writing any new code!
  - In case a traditional patch do not change the semantic of persistent kernel data structures

Ksplice Patches – Kernel Modules

- Ksplice patches replace always a **whole** kernel function
  - Each patch consist amongst other files of three kernel modules
- Example for patch *oviq2kch* (*CVE-2013-2634*)

ksplice-oviq2kch.ko	Performs safety checks Locates the function to be replaced in memory Applies/rollback the patch
ksplice-oviq2kch_vmlinux-new.ko	Contains the new patched kernel function object code
ksplice-oviq2kch_vmlinux-old.ko	Contains the old kernel function Used to locate the function in memory At the end it will be unloaded

Ksplice Patches – Application of a Patch

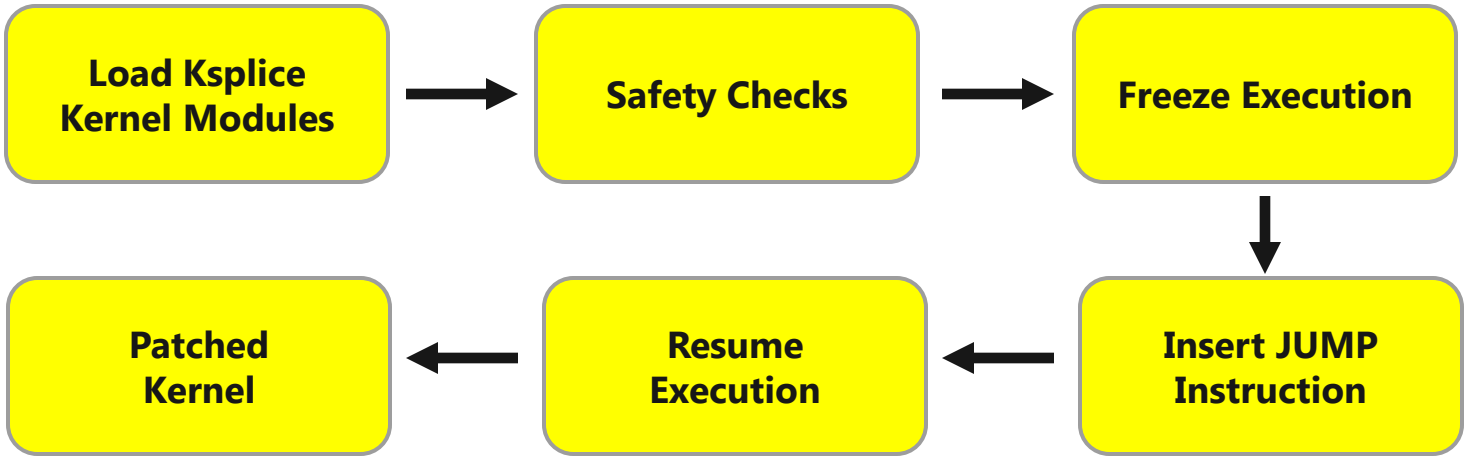

- The function to be replaced cannot be executed by any CPU
  - In other case, after multiple attempts the procedure will abort
- The execution of other programs will be suspended for about 0.7 ms
- Overhead for each applied patch
  - Some more CPU cycles used because of the JUMP instruction
  - On average 200kB memory overhead

AGENDA

1. Introduction
2. How it Works
3. Installation and Configuration
4. Patch Management
5. Oracle Database Environment
6. Core Messages

Ksplice Uptrack Installation – Requirements

- Ksplice Uptrack can be installed/uninstalled **online**
  - No reboot or activation of a specially prepared kernel
- Requirements
  - Oracle Linux Premier Support subscription
  - Ksplice Access Key (can be requested via *http://linux.oracle.com)*
- Each server needs to subscribe to the correct *Ksplice for Oracle Linux* channel at the Oracle ULN
  - During the server registration (*uln_register*)
  - After the installation at *http://linux.oracle.com*

Ksplice Uptrack Installation – RPM Package

- Install the Ksplice Uptrack package

```
# yum install uptrack -y
```

- Configuration file: */etc/uptrack/uptrack.conf*

```
# /etc/uptrack/uptrack.conf
accesskey = 56dac...
https_proxy =
install_on_reboot = yes
#upgrade_on_reboot = yes
autoinstall = no
```

  - Uptrack connects to *https://updates.ksplice.com:443* – make sure your firewall is aware of it

Ksplice Uptrack Offline Client – Local YUM Repository

- In many cases it is not practical or even possible to download and apply patches directly from the Oracle Ksplice server
- For each supported kernel version Oracle bundles all available Ksplice patches to one RPM
- Register one server at the Oracle ULN and define it as a local YUM repository

- Download all patches locally, by using the Oracle supplied script 167283.sh

```
# 167283.sh
...
[ol6_x86_64_ksplice: 7      of 110    ] Downloading getPackage/uptrack-
updates-2.6.32-100.28.11.el6.x86_64-20130721-0.noarch.rpm
```

Ksplice Uptrack Offline Client – Installing Updates

- Install the offline version of the Ksplice Uptrack client

```
# yum install uptrack-offline.noarch
```

- For Oracle Enterprise Linux 6

```
# yum install uptrack-updates-$(uname -r)
Installing: uptrack-updates-2.6.39-400.17.1.el6uek.x86_64-20130808-
0.noarch
The following steps will be taken:
Install [ptj4l5wq] CVE-2013-0268: /dev/cpu/*/msr local privilege
                                  escalation.
...
Your kernel is fully up to date.
Effective kernel version is 2.6.39-400.109.5.el6uek
```

- For Oracle Enterprise Linux 5

```
# yum install uptrack-updates-$(uname -r).$(uname -m)
```

AGENDA

1. Introduction
2. How it Works
3. Installation and Configuration
4. Patch Management
5. Oracle Database Environment
6. Core Messages

Patch Management – Overview

- Ksplice Uptrack is **not a replacement** for your regular package management system
- Ksplice Uptrack **will not**
  - patch any files on disk
  - provide any updates for userspace applications
  - provide all regular kernel RPM packages
  - provide any updates to 3rd party kernel modules
  - update any important shared libraries

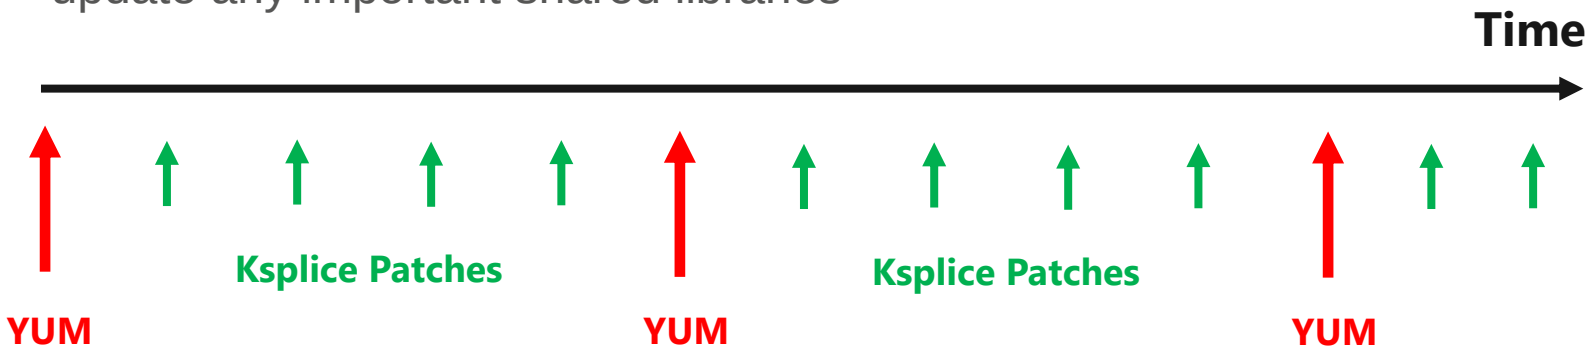

Patch Management – Tools

- Command line tools for Oracle Enterprise Linux (no GUI)
  - uptrack-upgrade, uptrack-install, uptrack-remove, uptrack-show, uptrack-uname
- Standard Linux tools, do not show the in memory effective kernel version

```
# uname -rv
2.6.39-400.17.1.el6uek.x86_64 #1 SMP Fri Feb 22 18:16:18 PST 2013
```

- For this purpose use *uptrack-uname*

```
# uptrack-uname -rv
2.6.39-400.109.5.el6uek.x86_64 #1 SMP Tue Jul 23 16:40:27 PDT 2013
```

- Applied/Available Ksplice patches will be stored in */var/cache/uptrack*

Applying Ksplice Uptrack Patches

- To download and apply all Ksplice patches available for your system

```
# uptrack-upgrade -y
...
Installing [seloxg7f] CVE-2012-6544: Information leak in Bluetooth
                                     L2CAP socket name.
Your kernel is fully up to date.
Effective kernel version is 2.6.32-400.29.3.el6uek
```

- To install only one Ksplice patch
  - Because of dependencies, in many cases a bundle of Ksplice patches will be installed

```
# uptrack-install xxqhbxgk -y
...
Installing [xxqhbxgk] CVE-2013-1929: Buffer overflow in TG3 VPD
                                     firmware parsing.
Effective kernel version is 2.6.18-348.6.1.el5
```

Ksplice Patches – Non-Quiescent Code

- On a heavy loaded system some kernel code might be executed constantly by many processes
- In this case, the apply procedure might fail

```
The following actions failed:
Install [t6yfkr0b] NULL pointer dereference in SCSI device removal.

Ksplice was unable to install the update because one or more programs are
constantly using the kernel functions patched by this update.

  - kworker/0:2 (pid 32)
  - kworker/0:1 (pid 18)
```

- Re-trying manually many times helps in most cases
  - Activating auto-installation for new patches might also be a solution

Listing Applied Ksplice Uptrack Patches

- To show all applied patches, including a short patch description

```
# uptrack-show
Installed updates:
[hlhx6na6] Clear garbage data on the kernel stack when handling signals.
[6b96sdwt] CVE-2012-1568: A predictable base address with shared
                          libraries and ASLR.
```

- To check which patches have been already downloaded, but have not been installed yet, use the option *--available*:

```
# uptrack-show --available
Available updates:
[xq7rv6l7] Use-after-free in USB serial driver probing.
[limx5afw] NULL pointer dereference on futex wakeup.
[koo5a2ja] Deadlock in block device journalling layer JBD.
[ipmp2zo8] Bogus return value in Virtio driver memory allocation.
[drummexy] NULL pointer dereference in ATA driver core.
```

Removing Ksplice Uptrack Patches

- To remove one patch
  - In most cases a patch bundle will be removed

```
# uptrack-remove hcz2odxp

The following steps will be taken:
Remove [qx38sixd] Data loss in filesystems due to delayed writeback.
Remove [hcz2odxp] Data loss in filesystems due to missing writeback.

Go ahead [y/N]? y
Removing [qx38sixd] Data loss in filesystems due to delayed writeback.
Removing [hcz2odxp] Data loss in filesystems due to missing writeback.
Effective kernel version is 2.6.39-400.109.4.el6uek
```

- To remove all applied Ksplice Uptrack patches

```
# uptrack-remove -all

Removing [cjz1git0] Networking failure with BladeEngine network adapter.
Removing [jb8gouo5] CVE-2013-3225: Kernel stack information leak in
                                   Bluetooth rfcomm.
...
```

Removing Ksplice Uptrack Patches – Kernel Modules

- Removing a patched kernel module, without a reboot is possible
  - Example: *[ulct1xyk] Missing return value in IPv6 socket options manipulation*

```
# lsmod | grep ulct1xyk
ksplice_ulct1xyk_ipv6_new    14116  1
ksplice_ulct1xyk            135789  2  ksplice_ulct1xyk_ipv6_new
ipv6                        333733  23 ksplice_ulct1xyk_ipv6_new
```

- But there might be a problem **after** a reboot

```
# lsmod | grep -E 'ulct1xyk|ipv6'
ipv6                   327781  22
ksplice_ulct1xyk       135789  1
```

```
# uptrack-remove ulct1xyk
The following actions failed:
Remove [ulct1xyk] Missing return value in IPv6 socket options
manipulation.
```

Re-Applying Ksplice Uptrack Patches During a Reboot

- A server reboot will not lead to a permanent loss of all Ksplice patches
- The *uptrack* service is responsible for

```
/etc/rc3.d/S09uptrack -> ../init.d/uptrack
```

  - Saving Ksplice Uptrack state during a shutdown or reboot
  - Applying Ksplice patches in a very early boot phase
- To deactivate the automatic application of the patches during boot time
  - Create a file */etc/uptrack/disable* before reboot

```
# touch /etc/uptrack/disable
```

  - Boot the kernel with the option *nouptrack*

```
# cat /proc/cmdline
ro root=/dev/mapper/vg_sys-lv_root ... quiet numa=off nouptrack
```

AGENDA

1. Introduction
2. How it Works
3. Installation and Configuration
4. Patch Management
5. Oracle Database Environment
6. Core Messages

Ksplice Uptrack & Oracle Database Environment

- Ksplice patches can be **applied/uninstalled online** in an Oracle database environment (cluster or a single instance)
  - They do not touch any shared libraries
  - Relinking software or scheduling a downtime is not necessary
  - OS kernel remains **up-to-date** (regarding security/critical bug fixes)!
- Not fully supported for Oracle Engineered Systems

- To stay up-to-date with the whole OS software stack we need **also** to consider all the **regular patches**
  - non-kernel bug fixes – most of them do not require any downtime
  - 3rd party kernel modules / HW driver patches – they might require a downtime
  - some shared libraries – they might require a downtime

Shared Libraries & Oracle Database Environment

- Oracle database and Grid Infrastructure have strong dependencies to some OS shared libraries (e.g. GLIBC, LIBAIO)

```
# ps -ef | grep oracle
oracle   31926     1   0 Aug31 ?        00:01:41 oracleP9981

# cat /proc/31926/maps
3e37200000-3e3721c000 r-xp 00000000 fd:3d8 1933337    /lib64/ld-2.5.so
3e37600000-3e3774f000 r-xp 00000000 fd:3d8 1933339    /lib64/libc-2.5.so
3e37a00000-3e37a82000 r-xp 00000000 fd:3d8 1933483    /lib64/libm-2.5.so
3e38200000-3e38202000 r-xp 00000000 fd:3d8 1933369    /lib64/libdl-2.5.so
```

- Although OS vendors guarantee backward compatibility for shared libraries during maintenance upgrades
  - MOS articles **recommend** to relink the software after upgrading OS (shared OS libraries)
  - For Grid Infrastructure a software relink is according to MOS even **required**

AGENDA

1. Introduction
2. How it Works
3. Installation and Configuration
4. Patch Management
5. Oracle Database Environment
6. Core Messages

Core Messages

- Ksplice Uptrack is a great and powerful technology
- It can be used to apply important kernel patches at once, without a downtime
  - Especially in case security plays a very important role for you
  - Combined with the regular patches
- But, **set the right expectations**!
  - In some cases you cannot avoid a downtime
  - Ksplice cannot be your only HA solution
- Does it make sense to combine Ksplice with RAC (rolling-upgrade)?
  - It depends on the business needs…
- During evaluation consider always the whole software stack

Questions and answers ...

Robert Bialek

Principal Consultant

Tel.: +49 89 99 27 59 30

robert.bialek@trivadis.com

BASEL BERN BRUGG LAUSANNE ZUERICH DUESSELDORF FRANKFURT A.M. FREIBURG I.BR. HAMBURG MUNICH STUTTGART VIENNA

trivadis
makes IT easier.