

# REDUCE PLANNED DATABASE DOWN TIME WITH ORACLE TECHNOLOGY SMART


Kirill Loifman  
dadbm.com  
DOAG 2015

# WHO IS KIRILL LOIFMAN

- Working with Oracle since 1996
- Oracle DBA experience on Unix AIX, OpenVMS, Tru64, HP-UX, Linux
- Senior DB Architect at adidas
- OCP DBA 9i,10g,11g,12c,...
- Blogger, speaker & a father
- Contacts  
www: dadbm.com  
Twitter: @loifmkir  
Google.com/+Dadbm  
Facebook, LinkedIn, Xing ...

The image shows a screenshot of the DaDBm Facebook page and website content. The Facebook page header includes the name 'DaDBm', navigation tabs for 'Page', 'Activity', 'Insights', and 'Settings', and a search bar. The main content area features a red banner with a photo of Kirill Loifman and the text 'DaDBm Community'. Below this is a 'Timeline' section with a photo of Kirill Loifman in front of a SAP sign. The website content on the right includes a header with 'DaDBm' and 'Oracle Consulting and DBMS Blog', a search bar, and navigation tabs for 'HOME', 'About', 'Services', 'Feedback', 'Blog', and 'Contact'. The main content area features a 'DaDBm - your Remote DBA Service with OCP Expertise' section, an 'About Me' section, a 'DBM's Blog' section, an 'Oracle Consulting' section, and a 'Timeline' section. The 'About Me' section includes a photo of Kirill Loifman and text describing his experience. The 'DBM's Blog' section includes a list of articles from the blog. The 'Oracle Consulting' section includes text describing the services offered. The 'Timeline' section includes a photo of Kirill Loifman in front of a SAP sign and text describing his experience. The 'Articles from the Blog' section includes a list of articles with titles and dates. The 'How We Work?' section includes text describing the services offered. The 'USER'S FEEDBACK' section includes a list of user testimonials. The 'DBA NEWS, TIPS, HOW-TOs, BLOG' section includes a list of news items. The 'ORACLE Certified Associate' and 'ORACLE Certified Professional' sections include text describing the certifications. The 'TAGS' section includes a list of tags. The 'MY SHORT PROFILE' section includes text describing Kirill Loifman's experience. The 'CORE SERVICES' section includes a list of services. The 'CONTACT ME' section includes text describing the contact information.

# DOAG ABOUT ME IN CONFERENCE BROCHURE 2014



## Die Datenbank hacken

Kirill Loifman, geboren im russischen Kaliningrad, lebt seit dreizehn Jahren in Deutschland. Er ist Oracle-Certified Database Consultant (OCP DBA) und besitzt mehr als siebzehn Jahre Erfahrung als Oracle-Datenbank-Administrator. Sein Blog ([www.dadbm.com](http://www.dadbm.com)) sowie seine Kommentare auf Twitter (@loifmkir), Google plus und Facebook sind in der Community weithin bekannt. Der Vortrag „live adventure – from my PC to Oracle Remote Database“ hat zum Ziel, sich live ohne Passwort in einen Datenbank-Server einzuhacken, um zu demonstrieren, wie man seine Daten vor unerlaubten Zugriffen schützen kann. Drei Fragen an Kirill Loifman:

*Wie wichtig ist Ihnen die Sicherheit Ihrer Daten?*

Daten sind das Wertvollste, was Unternehmen besitzen. Deshalb sollte die Datensicherheit auch die größte Beachtung erhalten, noch vor deren Verfügbarkeit.

*Was bedeutet Ihnen Ihre Arbeit?*

Mich fasziniert die IT-Technologie im Allgemeinen sowie die Oracle-Technologie im Besonderen. Nichts bleibt bestehen, alles ändert sich schnell und dynamisch. Als Lernbegeisterter konnte ich mich deshalb über mehr als siebzehn Jahre weiterentwickeln und bin somit in der Lage, neue Datenbank-Lösungen zu entwickeln, Datenbank-Probleme zu beheben und die Performance von Anwendungen zu steigern.



*Werden Sie es im Rahmen Ihres Vortrags schaffen, ohne Passwort die Verbindung zu einem Server herzustellen?*

Ich gebe zu, die Beschreibung meines Vortrags ist etwas provozierend. Um sich mit einer Oracle-Datenbank zu verbinden, ist ein Authentifizierungsmechanismus erforderlich, und ich werde einige davon beschreiben. Es gibt jedoch undokumentierte Tricks, mit denen ein erfahrener Oracle-Anwender ein ineffizientes Sicherheitssystem umgehen kann. Ich werde auch einige Techniken vorstellen, die den alltäglichen Datenbank-Zugang für Administratoren und Entwickler vereinfachen – ohne die Sicherheit außer Acht zu lassen.


# ELEMENTS OF SYSTEM AVAILABILITY

- **High availability**
  - Eliminating single points of failure through redundancy
  - Examples: redundant HW, SAN/ASM, RAC databases
- **Disaster recovery**
  - Recovering from failure
  - Examples: Extended clusters, standby databases
- **Oracle Maximum Availability Architecture (MAA)**
  - Combining HA and DR elements
- **Downtime**
  - How much time you can tolerate your system being unavailable
  - Includes unexpected as well as planned down time
  - 99.9% availability ~ 8 hours and 45 minutes of down time per year



# LEVELS OF SYSTEM AVAILABILITY

- **Level 0: Out-of-the-box** (no specific high-availability elements)
- **Level 1: Storage-level protection** (database instance with protected storage)
- **Level 2: CFC or/and Standby**  
(Cold Failover Cluster or/and Standby DB at the same physical location)
- **Level 3: Recovery via redundant components**  
(Hot Failover with multi-node RAC at the same physical location)
- **Level 4: Active and passive recovery**  
(2 site concept, MAA, RAC, Data Guard, Storage mirroring, Flashback DB,...)



Note: HA Level combinations possible; \* percentages shown are for illustrative purposes only.

# AVAILABILITY TARGET AND COST

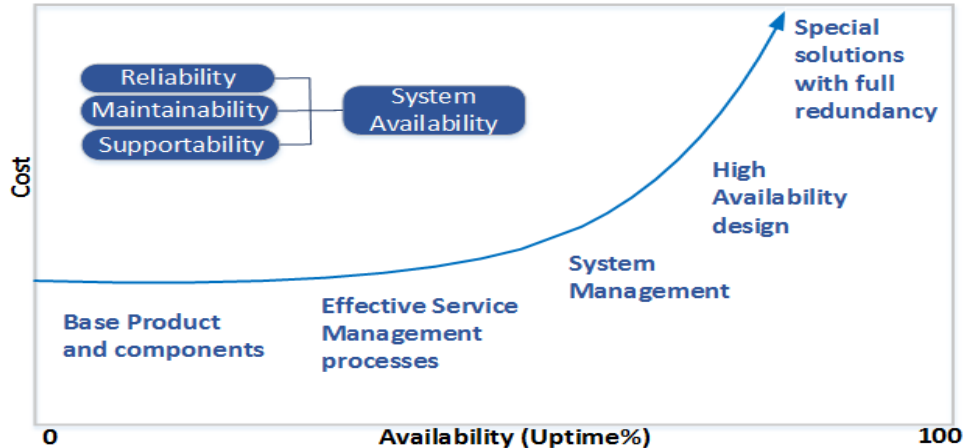
- **How many 9s do you need?**

Sharp decrease in downtime for each additional 9 in the availability target.

Availability Target	Downtime Per Year (Approx.)
90 %	36 days
98 %	7.3 days
99.7 %	26 hours
99.99 %	52 minutes
99.999 %	5 minutes

- **What will it cost?**

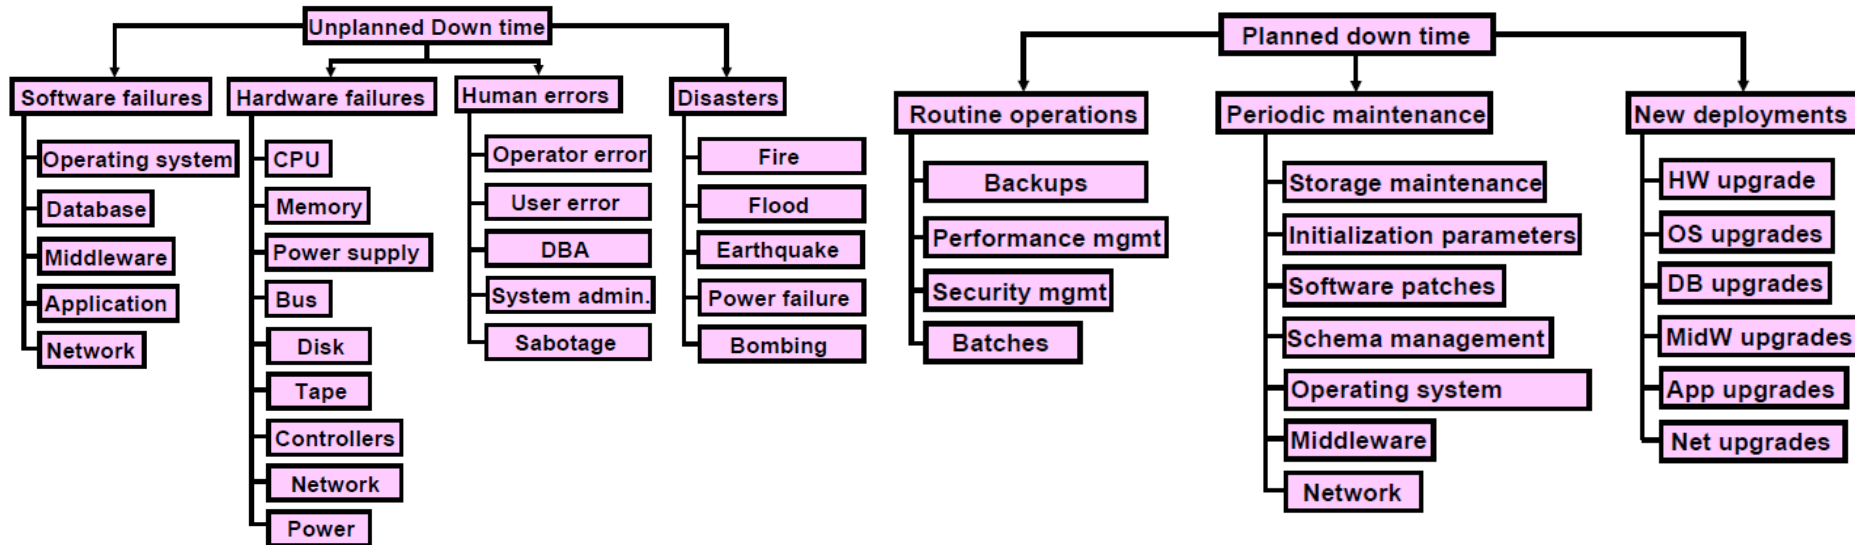
The higher the percentage of overall availability the greater the cost.





# CAUSES OF DOWNTIME – UNPLANNED VS PLANNED

- 2 ways of dealing with the potential for disaster:
  - a) expect and plan for it*
  - ~~*b) do nothing, and hope for the best*~~



# HOW TO MINIMIZE PLANNED DOWNTIME

- Categorize and Plan Maintenance Activities
- Minimize duration of every maintenance occurrence
- Different system architecture for different availability requirements
- Develop/Improve company service and change management processes

Maintenance Type	Examples
Application Maintenance	Application upgrade, code changes, etc.
Schema maintenance	Objects/Segment maintenance, etc.
Database maintenance	DB files/tablespace changes, backup, etc.
Instance maintenance	Instance parameters, SGA changes
DB software maintenance	DB/GRID patching/upgrade/migration
OS Maintenance	OS Patching / Upgrade
Hardware maintenance	Hardware extention; Server/HBA Firmware patching



# PLANNED DOWNTIME – MAINTENANCE CATEGORIES

- Improve what already exists
- Introduce complexity only if required
- Spend more budget only if required.
- Start out-of-the box, from green to red

Maintenance Type	Level 1	Level 2 (CFC)	Level 3 (RAC)	Level 4 (DR/MAA)
Application Maintenance	Green	Green	Green	Green
Schema maintenance	Green	Green	Green	Green
Database maintenance	Green	Green	Green	Green
Instance maintenance	Green	Green	Blue	Blue
DB software maintenance	Green	Blue	Yellow	Red
OS Maintenance	Green	Blue	Yellow	Yellow
Hardware maintenance	Green	Blue	Yellow	Yellow

# PLANNED MAINTENANCE TECHNIQS

- **Out-of-the box**

- Online System Reconfiguration
- Online Linux Patching
- Online database patching and patching best practices
- Online database parameter changes
- Online database file move/compression (12c+)
- Online partition move/compression (12c+)
- Online Table Redefinition including 12c enhancements
- Online Statistics Gathering during Bulk-Loads (12+)
- Online Application Upgrades with Edition-Based Redefinition
- Online / Active Database Duplication (11g+) and 12c enhancements

...

## **High Availability Solutions**

- Cold Failover Cluster – CFC (Oracle RAC 1 Node)
- Hot Failover Cluster (Oracle RAC/VM)
- Disaster Recovery Configuration (Oracle Data Guard)

...

# HARDWARE MAINTENANCE TECHNICS

## Online System Reconfiguration

- Online grow and shrink of shared memory and database cache
  - Reserve enough memory (sga\_max\_size /memory\_max\_size)
  - Linux Huge Pages
- Online addition/extension or removal of disks (ASM/LVM/SAN)
- Online addition or removal of clustered nodes (RAC/GRID)
- Online addition or removal of RMA/CPU's on SMP servers
  - Oracle SPARC M-Series & Oracle Solaris (Dynamic Reconfiguration & System Domains)
  - Hitachi Compute Blade logical partitioning (LPAR) & RHL/Windows
- Oracle 12c Database Smart Flash Cache Enhancements

# HARDWARE MAINTENANCE TECHNICS

## Oracle 12c Database Smart Flash Cache Enhancements

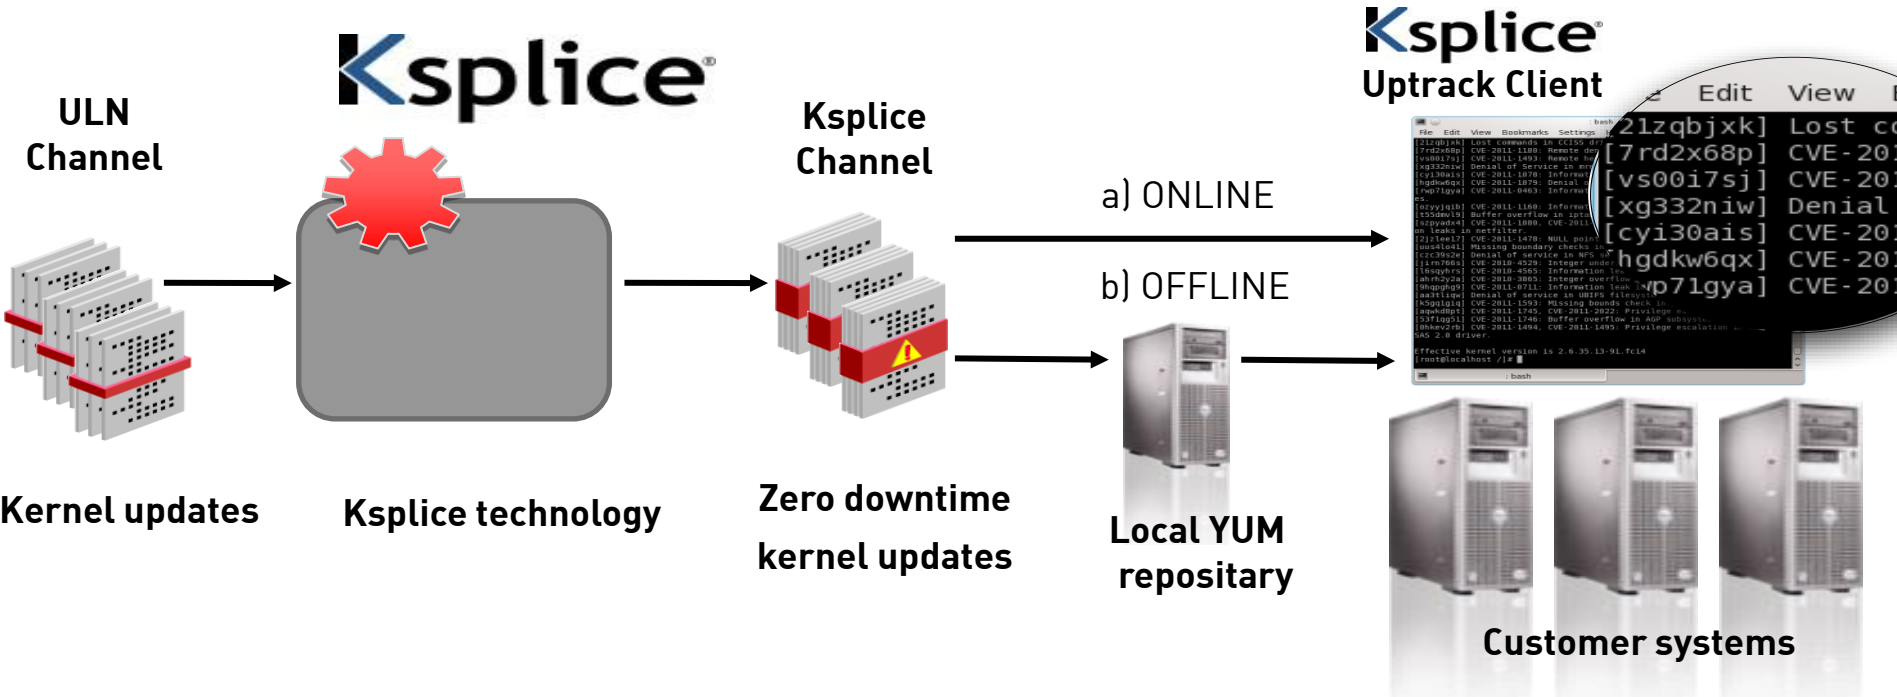
- Multiple flash drives in the flash cache (up to 16 devices in 12c)
- Specify flash devices at instance startup:  
`db_flash_cache_file = /dev/raw/sda, /dev/raw/sdb`  
`db_flash_cache_size = 32G, 64G`
- Dynamic enable/disable for flash cache devices:  
`db_flash_cache_size = 0, 64G`  
`db_flash_cache_size = 32G, 64G`
- HW example: HP Blade BL460c + HP IO Accelerator Flash Card 1,2T

# OS MAINTENANCE TECHNICS

## Online Oracle Enterprise Linux (OEL) Patching

- Enabled by Oracle Ksplice technology (founded in 2008; at Oracle since 2011)
- Types of patches
  - Security updates and other critical fixes
  - Kernel Diagnostic patches
  - Userspace patching: glibc & OpenSSL (last OOW)
- Prerequisites
  - Oracle Enterprise Linux 5.x+ with OU or RH compatible Kernels
  - Oracle Linux Premier Support
  - Get ULN access and Ksplice access key

# OS MAINTENANCE TECHNICS



# OS MAINTENANCE TECHNICS

## Two Ways to Consume Ksplice

- Individual Servers can register with Oracle's Ksplice server directly
  - Each system must be connected to Internet
  - Each system will check for new updates ~every 4 hours
  - Oracle provides an interactive Web portal to monitor users' systems
  - Updates can be auto-installed if desired
- Ksplice YUM repository
  - All internal servers configured to use local Ksplice Repository
  - Cron job updates local patch mirror once per day
  - Patching triggered manually with yum
  - Decrease patch data downloaded from Internet.
  - Useable with: Spacewalk Server or EM12c



# OS MAINTENANCE TECHNICS

## Ksplice Off-Line with Intranet Connection

- Create a local YUM Mirror and register the Ksplice Channel(s)
  - [http://docs.oracle.com/cd/E37670\\_01/E37355/html/ol\\_offlncl\\_ksplice.html](http://docs.oracle.com/cd/E37670_01/E37355/html/ol_offlncl_ksplice.html)
  - Subscribe each machine in `yum.repos.d` to your local YUM Ksplice channel

```
[ol6_x86_64_ksplice]
name=Ksplice for $releasever - $basearch
baseurl=http://local_yum_server/yum/OracleLinux/OL6/ksplice/$basearch/
gpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY
gpgcheck=1
enabled=1
```
- Install Ksplice “Uptrack-offline” RPM onto target machines.

```
# yum install uptrack-offline
```
- Perform kernel patching with your Local YUM Ksplice Repository

```
# yum install uptrack-updates-`uname -r`
```
- Can be integrated with Spacewalk and EM12c

# DB SOFTWARE MAINTENANCE

## Patching Best Practices (Single One-Off, CPU, PSU, Bundle Patch)

- Patch Strategy: at least every 6 months, before DB creation, before upgrade, ...
- Test and automate patching (functional, load, performance testing)
- Avoid combining patches to minimize risk due to untested combinations
- Review patch issues on MOS after 4 weeks of the patch release
- Do prerequisite checks in advance (opatch or MOS Patch Conflict Checker)
- 98% of patches can be installed in a highly available manner
  - Online patch (small amount interim one-off patches )
  - RAC rolling (some interim, all PSU)
  - Data Guard Standby First (some interim, all PSU and Bundle)
  - All patches can be applied Out of Place (11.2.0.2+)

# DB SOFTWARE MAINTENANCE

## Optimizing Patch Application

- Unzip patches in the same patch location and apply them during same downtime
- Apply several patches during the same downtime

- **Execute best practices**

```
opatch napply -skip_subset -skip_duplicate  
opatchauto apply [-nonrolling] [-silent] [-local]  
SQL> startup  
datapatch [-verbose] [-apply / -rollback <patch_id>] [-force]
```

- 12.1) datapatch takes care of applying the post SQL
- 12.1.0.2) No need for catbundle.sql with PSU, BP installs  
    dba\_registry\_sqlpatch instead of dba\_registry\_history  
    datapatch -rollbackall

# DB SOFTWARE MAINTENANCE

## Online database patching

- Available only for small amount interim one-off patches and bug fixes
- Applied and removed from a running instance
- Additional memory consumption and process start time penalty
- Rollback and replace them with “regular” ones during next downtime
- Oracle 12c -> no improvements found
- Use OPatch utility (which uses oradebug to install/uninstall the patch)

```
opatch apply online
```

# DB SOFTWARE MAINTENANCE

## Out of place patching

- Available from 11.2.0.2+ / supported by SAP
- Minimize downtime during patching of NON-RAC databases
- In-Place vs Out-of-Place

11gR2 regular In-Place	12c Out-of-Place
Shutdown DB / Listener	Clone existing ORACLE_HOME online
Binary part: opatch apply	Binary part: opatch apply
Logs check / fixes / conflict resolution patches	Logs check / fixes / conflict resolution patches
Possible rollback of binaries	Possible rollback of binaries
DB Startup	DB Restart in the new cloned \$ORACLE_HOME
DB part: SQL>@catbundle.sql psu apply	DB part: ./datapatch -verbose

# DB SOFTWARE MAINTENANCE

## Out of place patching – apply PSU example

```
# ONLINE clone (as root) $ORACLE_HOME_O (dbhome_1) to $ORACLE_HOME_C (dbhome_2)
mkdir $ORACLE_HOME_C
chown oracle:oinstall $ORACLE_HOME_C
cd $ORACLE_HOME_O
tar cvfp - . --exclude=*.log --exclude=*.trc | ( cd $ORACLE_HOME_C; tar xvf - )
```

### # Clone the installation with OUI (using Perl or OUI itself)

```
./runInstaller -clone -silent -noconfig -defaultHomeName ORACLE_BASE=$ORACLE_BASE
ORACLE_HOME=$ORACLE_HOME_C oracle_install_OSDBA=dba oracle_install_OSOPER=oper
```

### # Run root.sh (as root) from the cloned ORACLE\_HOME

```
/oracle/product/12.1.0/dbhome_2/root.sh
```

### # Patch the cloned ORACLE\_HOME with PSU 12.1.0.2.5

```
export ORACLE_HOME=$ORACLE_HOME_C
cd /oracle/stage/12.1.0.2.5-21359755
$ORACLE_HOME/OPatch/opatch apply Database home
```

# DB SOFTWARE MAINTENANCE

## Out of place patching – apply PSU example

# Switch database services to the cloned ORACLE\_HOME

# Update new ORACLE\_HOME in the /etc/oratab and in any .profile files

# Restart DB and listener from the cloned ORACLE\_HOME

```
export ORACLE_HOME=$ORACLE_HOME_O
```

```
lsnrctl stop
```

```
sql> shutdown immediate
```

```
export ORACLE_HOME=$ORACLE_HOME_C # or better . oraenv
```

```
lsnrctl start
```

```
sql> startup
```

# Complete the Patch post installation steps

```
./datapatch -verbose
```

# Cleanup the old ORACLE\_HOME

```
export ORACLE_HOME=$ORACLE_HOME_O
```

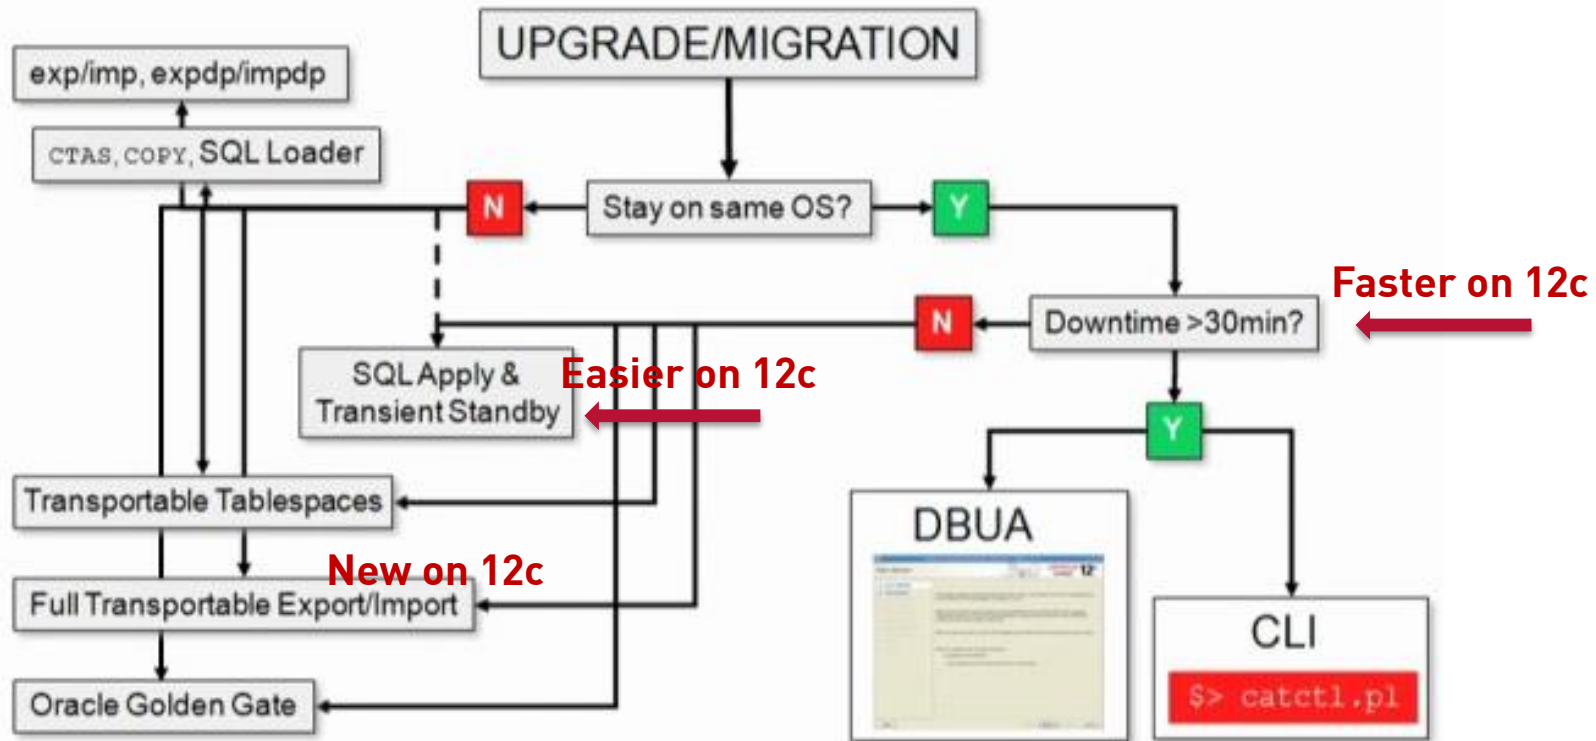
```
$ORACLE_HOME/oui/bin/detachHome.sh
```

```
rm -rf $ORACLE_HOME
```



# DB SOFTWARE MAINTENANCE


Database upgrade / migration enhancements



# DATABASE MAINTENANCE


## Oracle 12c Multitenant Architecture

- Greater and easier **application separation** for DB consolidation environments
  - >added DBA complexity though
- Faster **database provisioning**
  - Pluggable PDB creation from PDB\$SEED
  - Clone PDB into the same or another CDB (still not officially ONLINE -> wait for 12cR2)
  - Plug an unplugged PDB into the same or another CDB (even to higher version)
  - Plug non-CDB in as PDB
- More **patching / upgrade** options
  - Patch / Upgrade all PDBs in a CDB at once
  - More PDBs -> longer patching / upgrade time
  - Plug PDB into the CDB container with higher version



# DATABASE MAINTENANCE

Online database file move (new in 12c)



- Move **data file online** to different location , disk or storage system
- Progress in V\$SESSION\_LONGOPS (One row per file; blocks moved so far)
- SQL Examples:

```
ALTER DATABASE MOVE datafile 5 TO '+DiskGroup3' REUSE;  
ALTER DATABASE MOVE datafile '/disk1/myexample01.dbf'  
TO '/disk2/myexample01.dbf' KEEP ;
```

Compatible	Not compatible
Queries, DML and DDL operations Block media recovery Tablespace made READ ONLY or READ WRITE Data file RESIZE (extension) Online backup, Flashback Database*	Data file OFFLINE Media recovery Data file RESIZE (shrink) operation FLASHBACK DATABASE operation

# DATABASE MAINTENANCE

## Active Database Duplication Enhancements

- Default use of **backup sets** for active duplication (*Pull method* instead of *Push*)
- Choice of **compression**, **section size**, and encryption
- Option to end duplication with DB in mounted state (Move DB to ASM; Upgrading DB)
- Duplication of pluggable databases including all or individual PDBs
- Example

```
RMAN> DUPLICATE DATABASE TO cdb2 PLUGGABLE DATABASE pdb1, pdb3;
```

```
RMAN> DUPLICATE TARGET DATABASE TO orcl2 FROM ACTIVE DATABASE  
[USING BACKUPSET]  
[SECTION SIZE ...]  
[USING COMPRESSED BACKUPSET]  
[NOOPEN] ...;
```

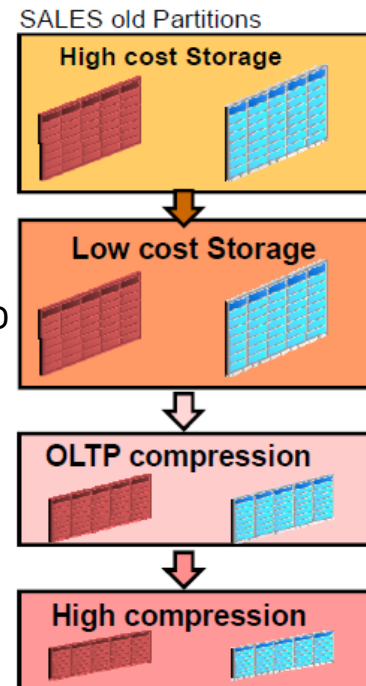
# DB SCHEMA MAINTENANCE

Online partition move/compression (new in 12c)

- Move / split / merge partitions and subpartitions **ONLINE**
- Compress partitions and subpartitions **ONLINE**
- DML allowed, not DDL / Global and local indexes maintained
- Used by ILM to move / compress table partitions according ADO po
- SQL examples:

```
ALTER TABLE orders MOVE PARTITION ord_p1  
TABLESPACE lowtbs UPDATE INDEXES ONLINE;
```

```
ALTER TABLE orders MOVE PARTITION ord_p1  
ROW STORE COMPRESS UPDATE INDEXES ONLINE;
```



# DB SCHEMA MAINTENANCE

## Online Table Redefinition (new in 12c)

- Following enhancements introduced in Oracle 12c:
  - Support for redefinition of **multiple partitions** in a single redefinition session.
  - Support for tables with **VPD policies**
  - New parameter **dml\_lock\_timeout** to handle pending DML in FINISH\_REDEF\_TABLE procedure
  - New procedure **REDEF\_TABLE** allows a one-step operation to redefine a table or partition
  - **Improved performance** of sync\_interim\_table and resilience of finish\_redef\_table
- ```
SQL> EXEC DBMS_REDEFINITION.START_REDEF_TABLE ( uname => 'hr',-  
orig_table => 't1', int_table => 'int_t1',-  
options_flag => DBMS_REDEFINITION.CONST_USE_PK,-  
copy_vpd_opt => DBMS_REDEFINITION.CONST_VPD_AUTO)
```

```
EXEC DBMS_REDEFINITION.FINISH_REDEF_TABLE (uname => 'hr',-  
orig_table => 't1', int_table => 'int_t1'-  
dml_lock_timeout => 100)
```

# DB SCHEMA MAINTENANCE

## Enhanced Online DDL Capabilities (new in 12c)

- Following operations do not block end user DML statements

DROP INDEX **ONLINE** (except cluster index, index on queue table)

DROP CONSTRAINT **ONLINE** (except CASCADEd and referencing constraints)

ALTER INDEX UNUSABLE **ONLINE** (except index on temporary table)

SET COLUMN UNUSED **ONLINE** (except columns with DEFERRABLE constraint)

ALTER INDEX [VISIBLE | INVISIBLE]



# DB SCHEMA MAINTENANCE


## Online Statistics Gathering during Bulk-Loads (new in 12c)

- In Oracle Database 12c, the database automatically gathers table statistics during the following types of bulk loads:
  - `CREATE TABLE AS SELECT`
  - `INSERT INTO ... SELECT` into an empty table by using a direct path insert.
- Statistics are available immediately after load.
- No additional table scan is required to gather statistics.
- All internal maintenance operations that use CTAS (MV refresh) benefit
- Index statistics or histograms are not gathered -> manual `gather_table_stats`

# APPLICATION MAINTENANCE

## Indexes on the Same Set of Columns (**new in 12c**)

- Several different indexes on the same set of columns
  - Unique and non-unique
  - Partitioned and not-partitioned
  - Locally and Globally Partitioned
  - Indexes with differ partitioning type (range or hash)
- Only **one** of the indexes **visible at a time**.
- If `OPTIMIZER_USE_INVISIBLE_INDEXES=TRUE`, optimizer uses the invisible index.



# APPLICATION MAINTENANCE

## Invisible table columns (new in 12c)

- Invisible columns are user-specified hidden columns

```
SQL> CREATE TABLE mytab (col1 VARCHAR2(10), col2 NUMBER INVISIBLE);
```

- Hidden columns do not affect existing applications that access the table.
- A hidden column can be accessed only by referring to its name explicitly

```
SQL> INSERT INTO mytab (col1, col2) values ('A',1);
```

- Can be made visible later

```
SQL> SET COLINVISIBLE ON
```

```
SQL> DESC mytab
```

```
SQL> ALTER TABLE mytab MODIFY (col2 VISIBLE);
```

# APPLICATION MAINTENANCE

## Edition-Based Redefinition (11gR2+)

## Database Development Guide



### Table of Contents

- Oracle Database Development Guide
  - Preface
  - Changes in This Release for Oracle Database Development Guide
  - Design Basics
  - Performance and Scalability
  - Security
  - High Availability
  - Advanced PL/SQL Features
  - SQL Processing for Application Developers
  - Using SQL Data Types in Database Applications
  - Using Database Expressions in Database

## 24 Using Edition-Based Redefinition

**Edition-based redefinition (EBR)** lets you upgrade the database component of an application while it is in use, thereby minimizing or eliminating downtime.

To upgrade an application while it is in use, you copy the database objects that comprise the application and redefine the copied objects in isolation. Your changes do not affect users of the application—they continue to run the unchanged application. When you are sure that your changes are correct, you make the upgraded application available to all users.

Using EBR means using one or more of its component features. The features you use, and the downtime, depend on these factors:

# SUMMARY

- Classify company applications based on the availability requirements
- Categorize and plan maintenance activities
- Minimize duration of every maintenance activity using best practices and new technics.
- Build DB environments based on availability requirements
- Do end-to-end availability and failover testing
- Develop/Improve corporate service and change management processes

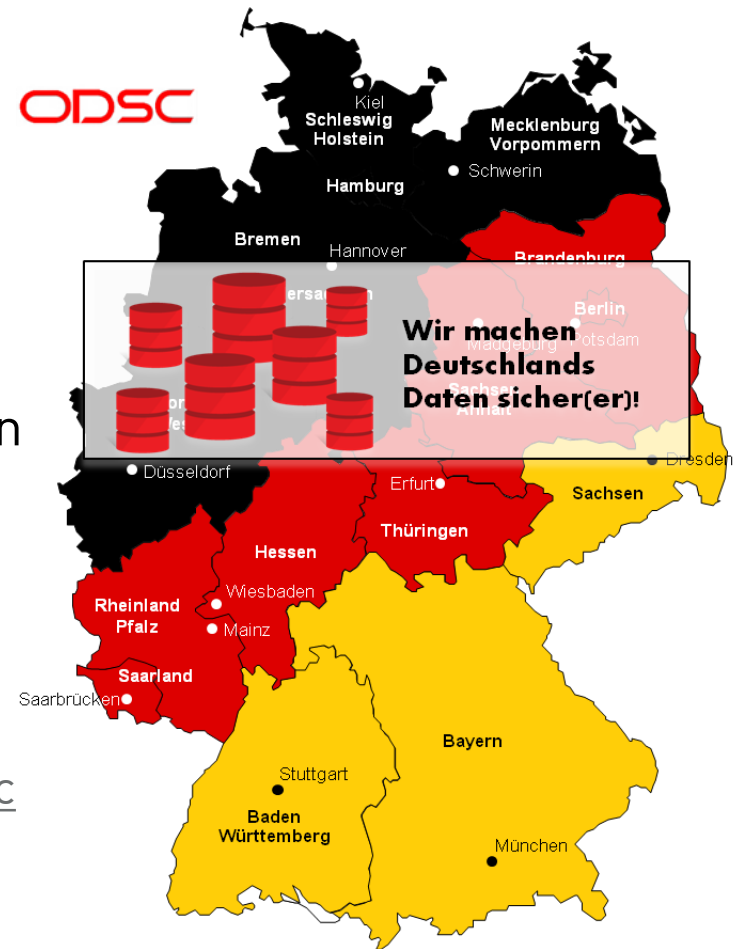
# Wir machen Deutschlands Daten sicher(er)!

Eine Community für DBAs, die aktiv die Daten Ihrer Datenbanksysteme schützen wollen oder müssen.

Wir erarbeiten sinnvolle Konzepte und tauschen unser Wissen aus.

[https://blogs.oracle.com/SecurityDE/entry/db\\_security\\_community\\_wurde\\_gegr%C3%BCndet](https://blogs.oracle.com/SecurityDE/entry/db_security_community_wurde_gegr%C3%BCndet)

Email to join: [carsten.muetzlitz@oracle.com](mailto:carsten.muetzlitz@oracle.com)



THANK YOU!

???

NOW or at [DADBM.COM](http://DADBM.COM)

Kirill Loifman

11-Nov-2015