CoMinder

Harmony 2017 Conference
17-18 of May, 2017
Helsinki, Finland

Shrinking the database

all you should know
to make your database smaller
(v. 1.02)

Andrey Chervonets (Андрей Червонец)

e-mail: a.chervonets@cominder.eu

Company: SIA CoMinder (Latvia)

About speaker

Andrey Chervonets (Андрей Червонец)

- founded company: SIA CoMinder (Latvia) in 2011.
- Experience: 15+ years' experience of supporting systems based on Oracle products and technologies (and all around).
- Certified:
  - Oracle Database OCP/OCE
  - Oracle AS/WebLogic OCP
  - IBM DB2, Guardium
- Contacts:
  - e-mail: a.chervonets@cominder.eu
  - LinkedIn: http://www.linkedin.com/in/andreychervonets

About CoMinder

CoMinder

- Has been founded in 2011:
  - to provide better DBA expert services
  - for customers interested in improvements
- Registered partner of:
  - Oracle, Red Hat, Microsoft, XMind
- Main focus areas:
  - Database/Middleware administration:
    consulting, **project works**, 24*7*365 SLA based **outsourcing**
  - Operating system administration (Linux, Windows) + (AIX, Solaris, HP-UX)
  - Data processing automation
  - Support tools development
  - Specific task upon customers request

Agenda

- Why should/may I need to keep DB size smaller?!
- Space Control concept
- Ideas / Technologies for used space size reduction
- Applied technologies (during a System life cycle)
- How to stay alive... and make life better
- 12+N ways to “lose weight”
- TOOLS
- Summary
- Q + A

Disclaimer

- Will NOT discuss:
  - Exact command syntax, specifications
  - Does it work in …. version ….
- Use at Your own risk!
- Do NOT believe me!
  - Try to understand first how does it really work
  - Test in your specific project / environment(s)
  - Everything may change in new version
  - Invent, design better (and share to community)

Any feedback is welcome!

Do we really need

to keep DB smaller

?

Who cares?

Why me?

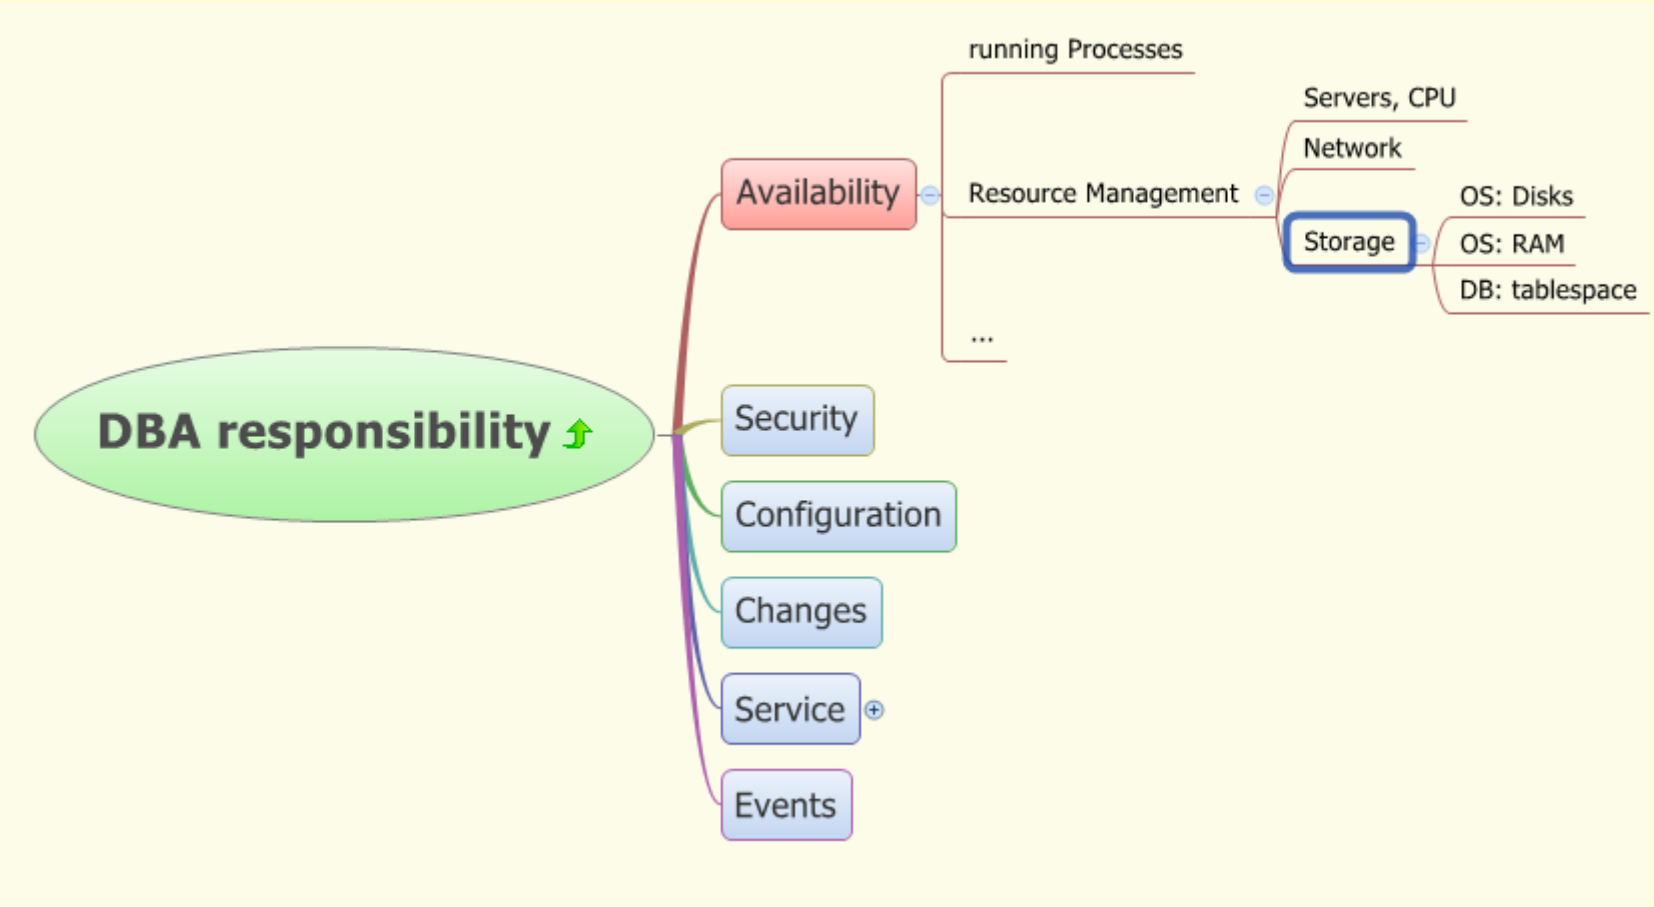

Survey summary

- What IT-experts think about this?!

How much time do You spend for database space management for ALL databases You are responsible for.

52 responses

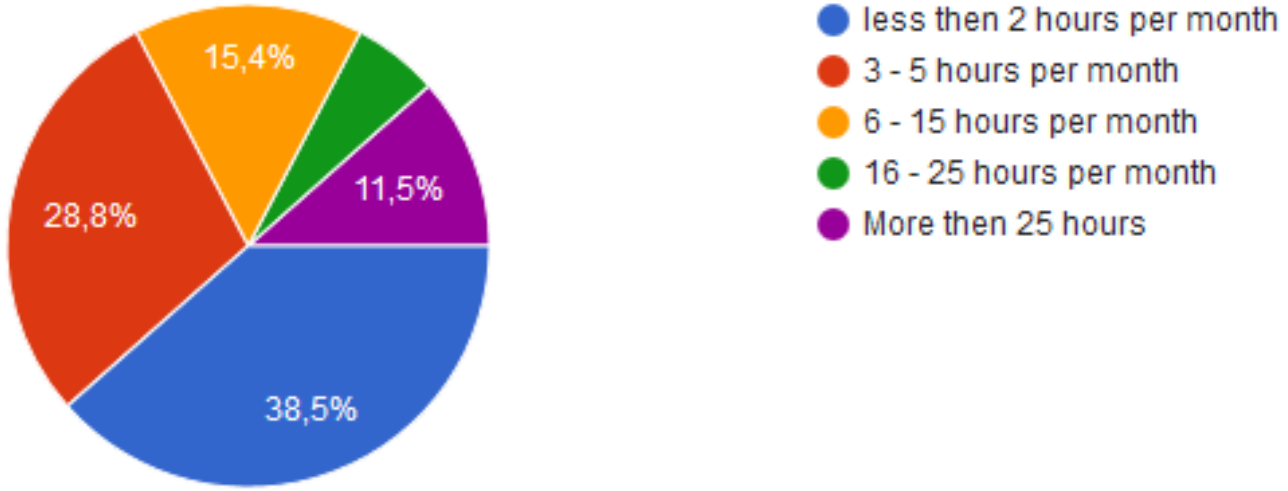

Survey summary

- What IT-experts think about this?!

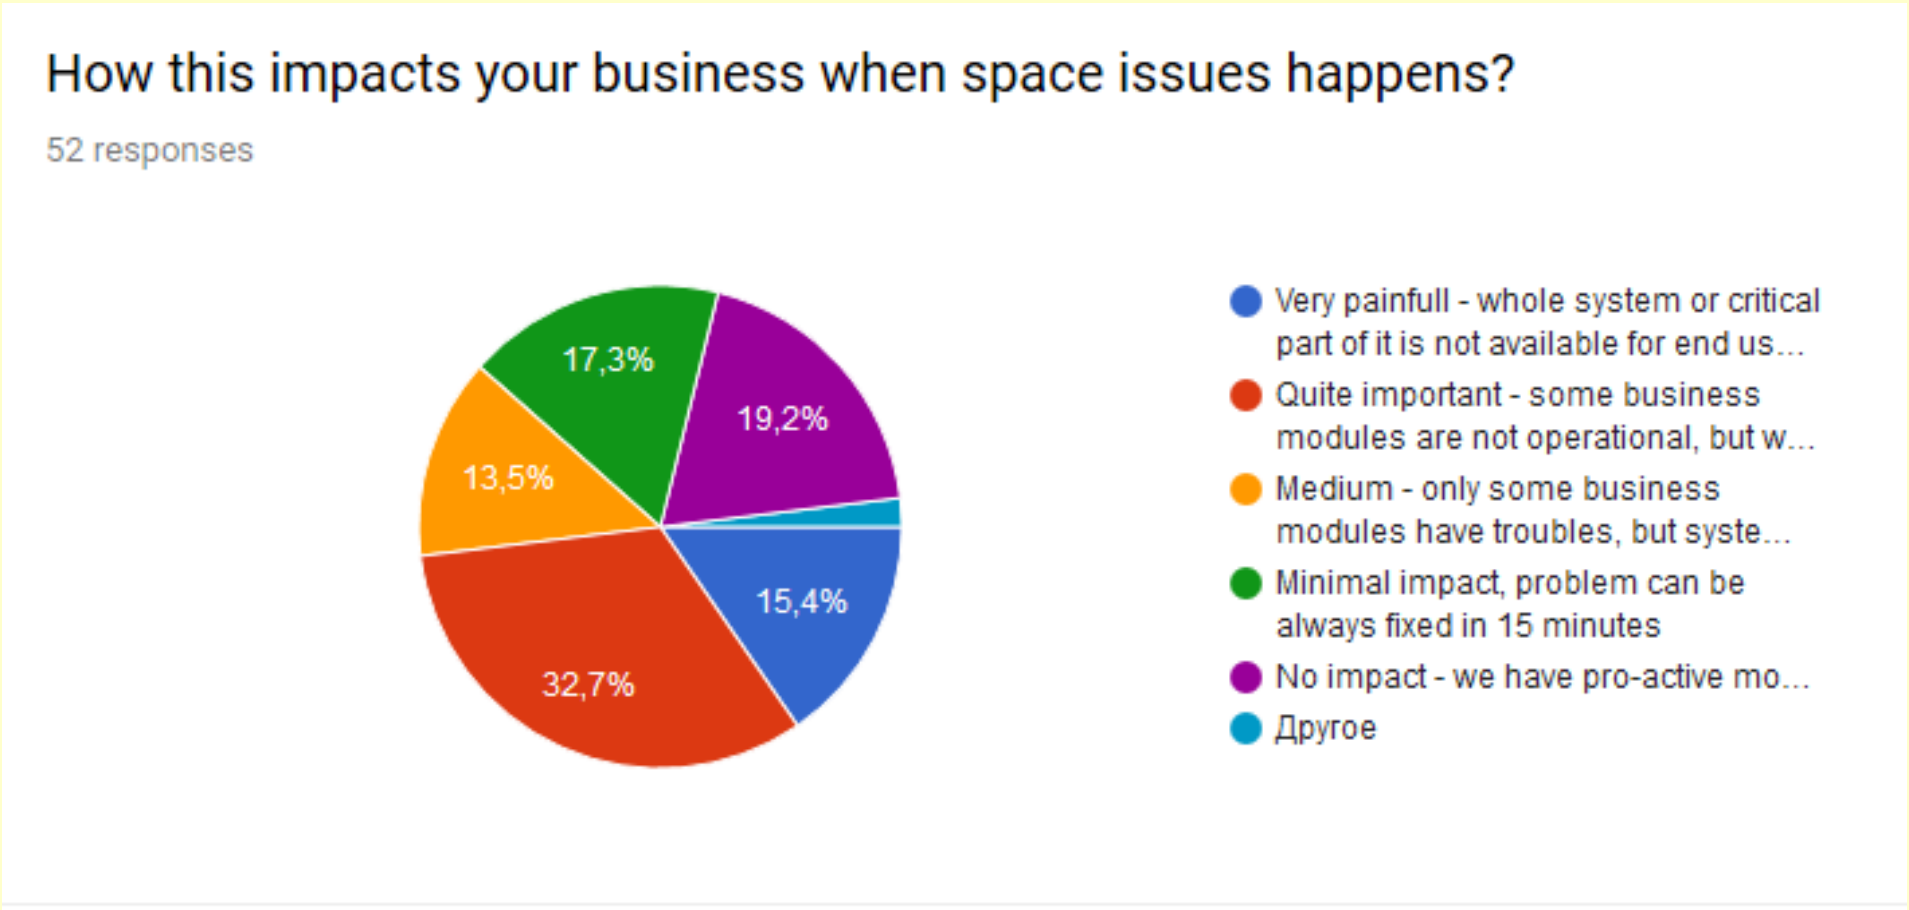

Survey summary

- What IT-experts think about this?!

Database server space management - when it is critical?!

Do You believe that space management in general (regardless of database size and environment) is important or even critical to keep database operational?

52 responses

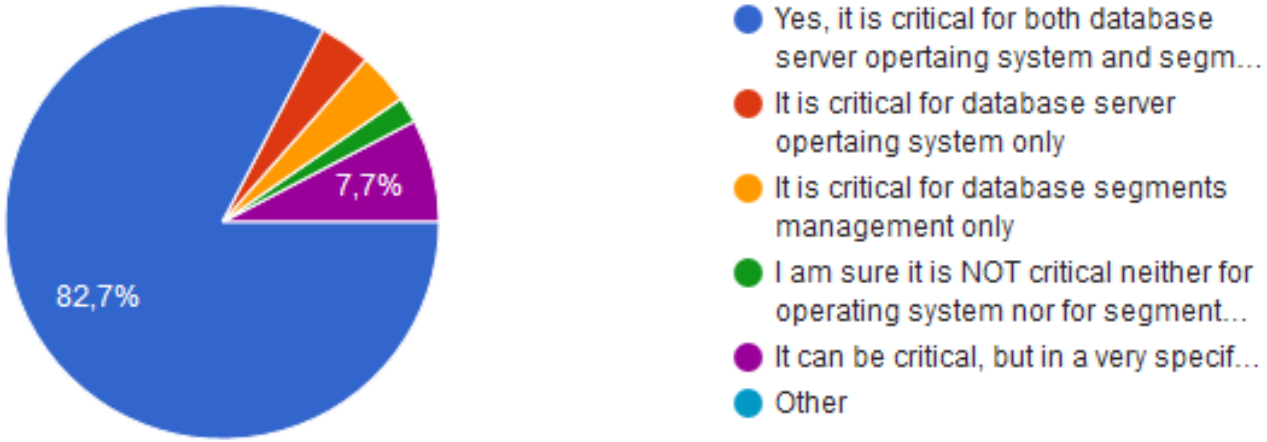

Survey summary

- What IT-experts think about this?!

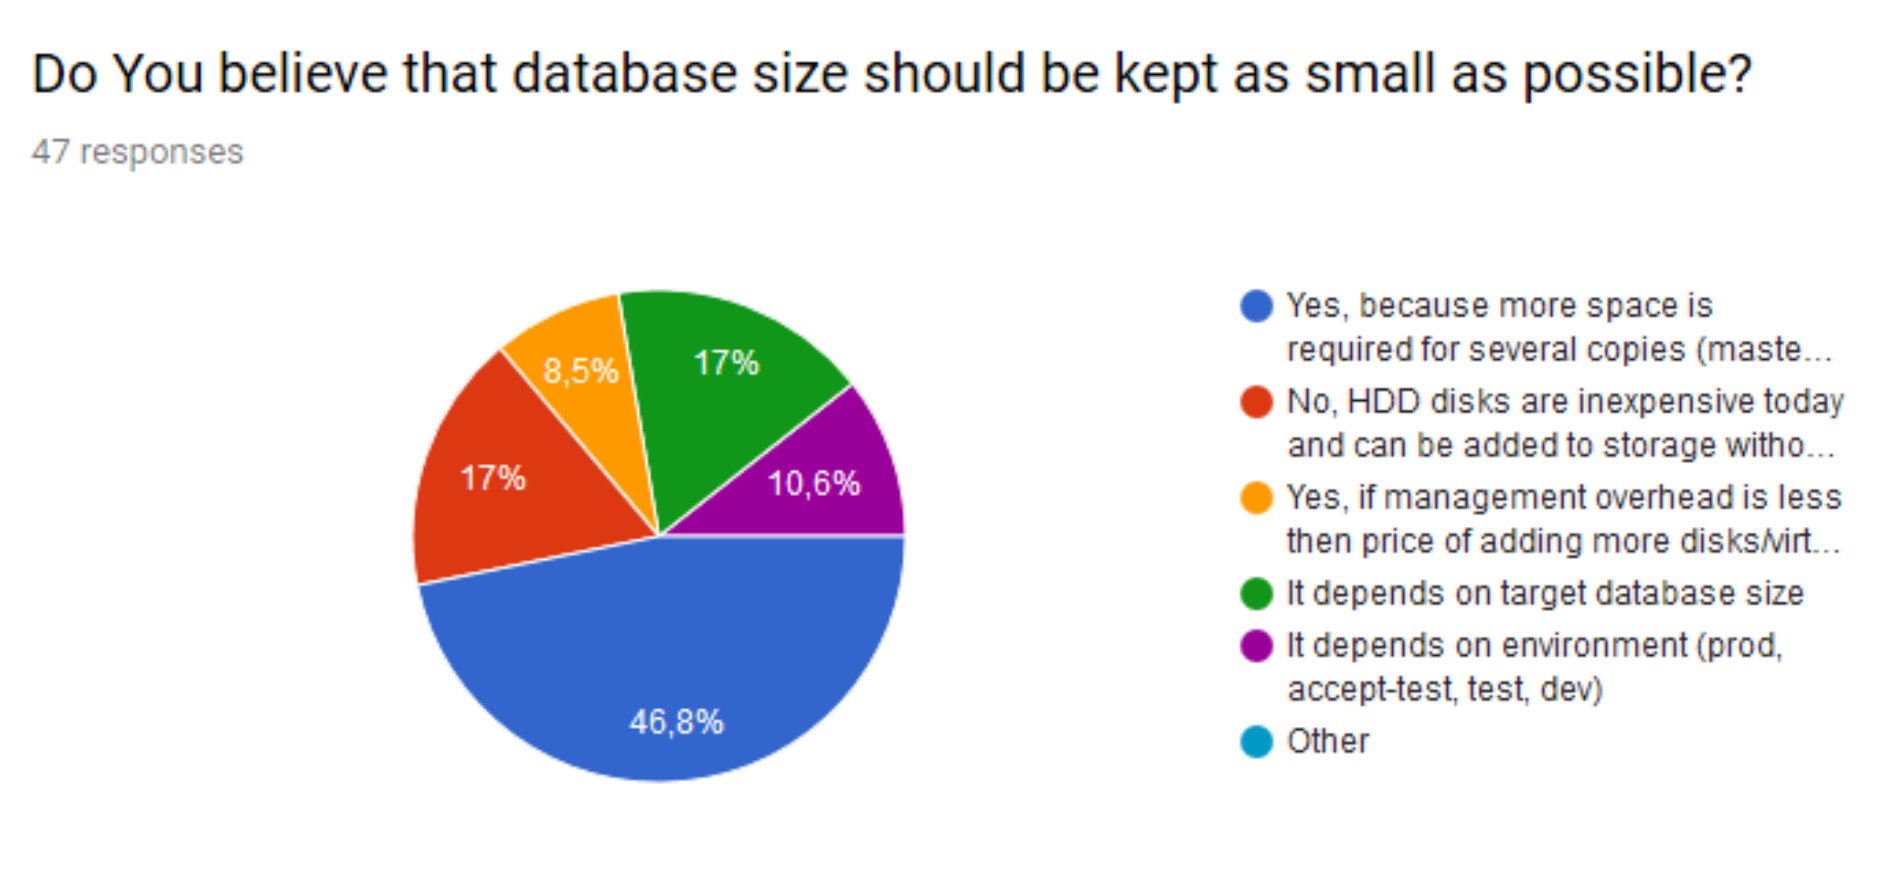

Survey summary

- What IT-experts think about this?!

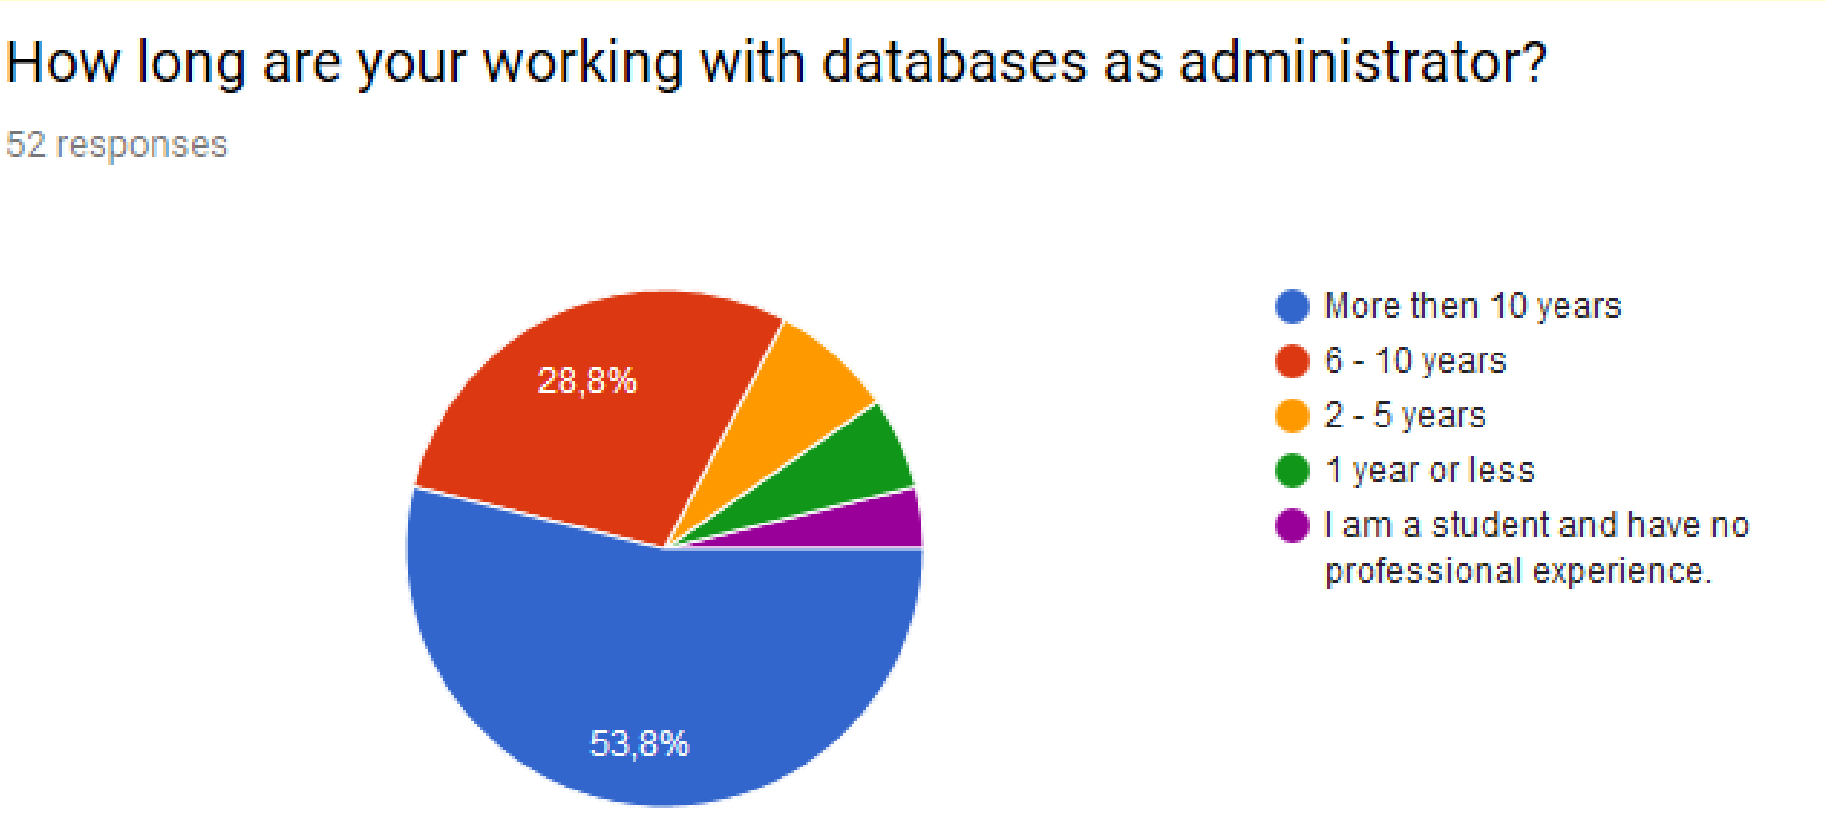

Big deal?

Is it big deal to resolve?

- In most cases short-term solutions can be found
- Kill It With Iron (KIWI) – is still most popular

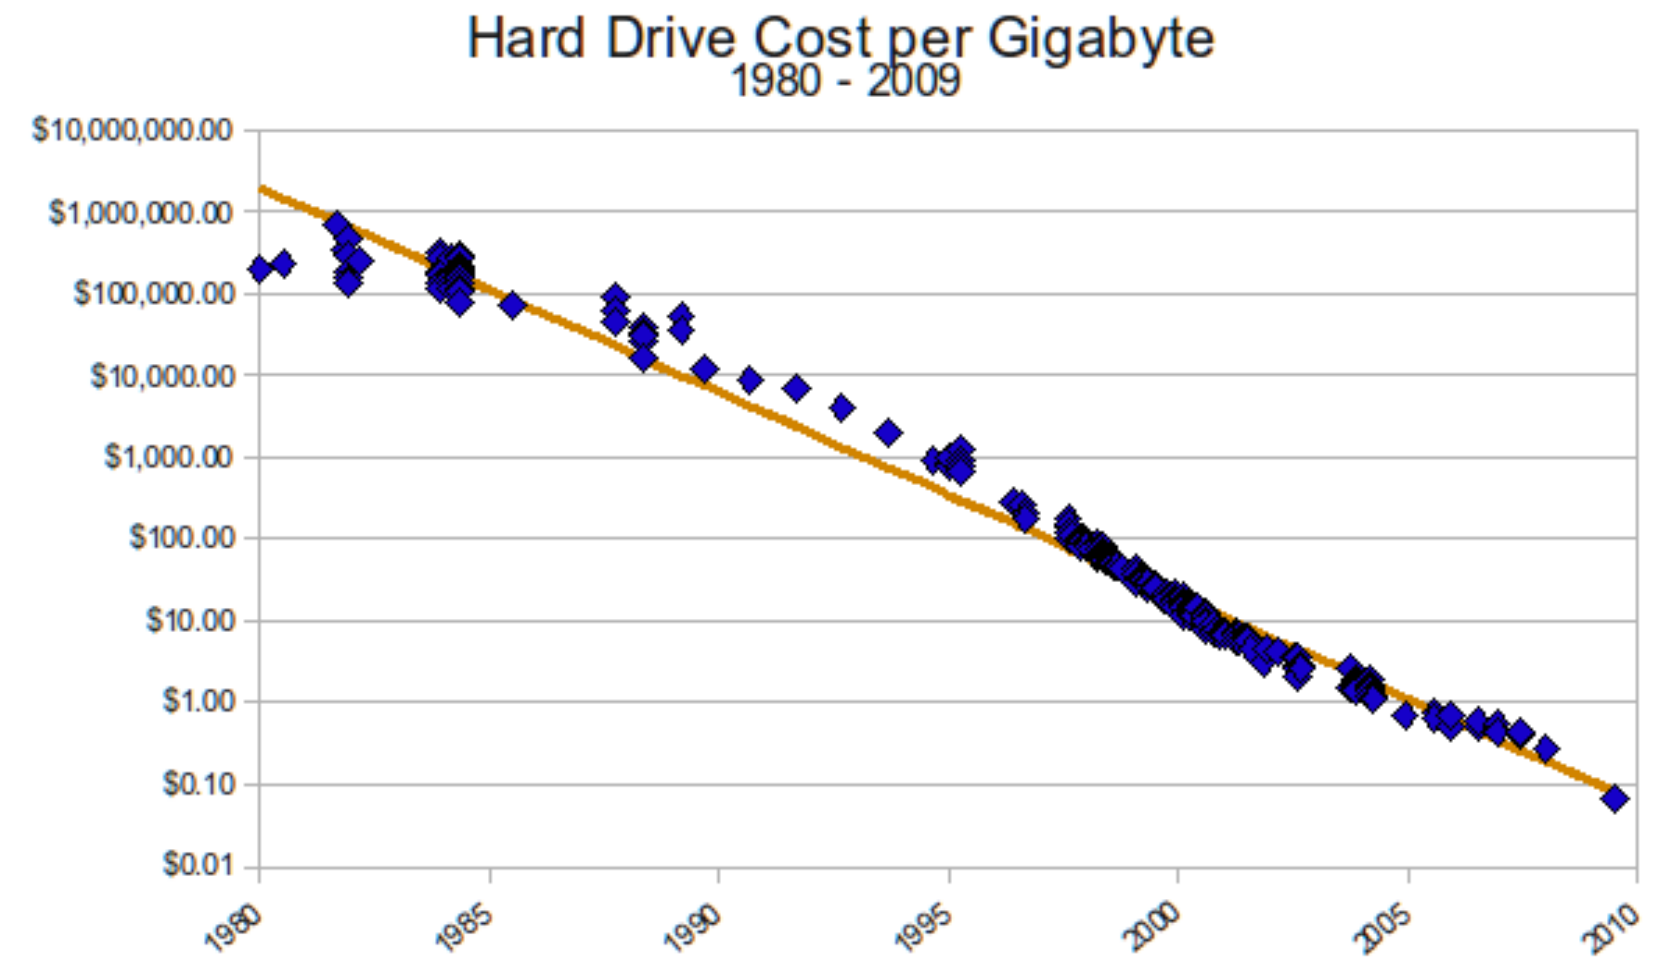

Big deal?!

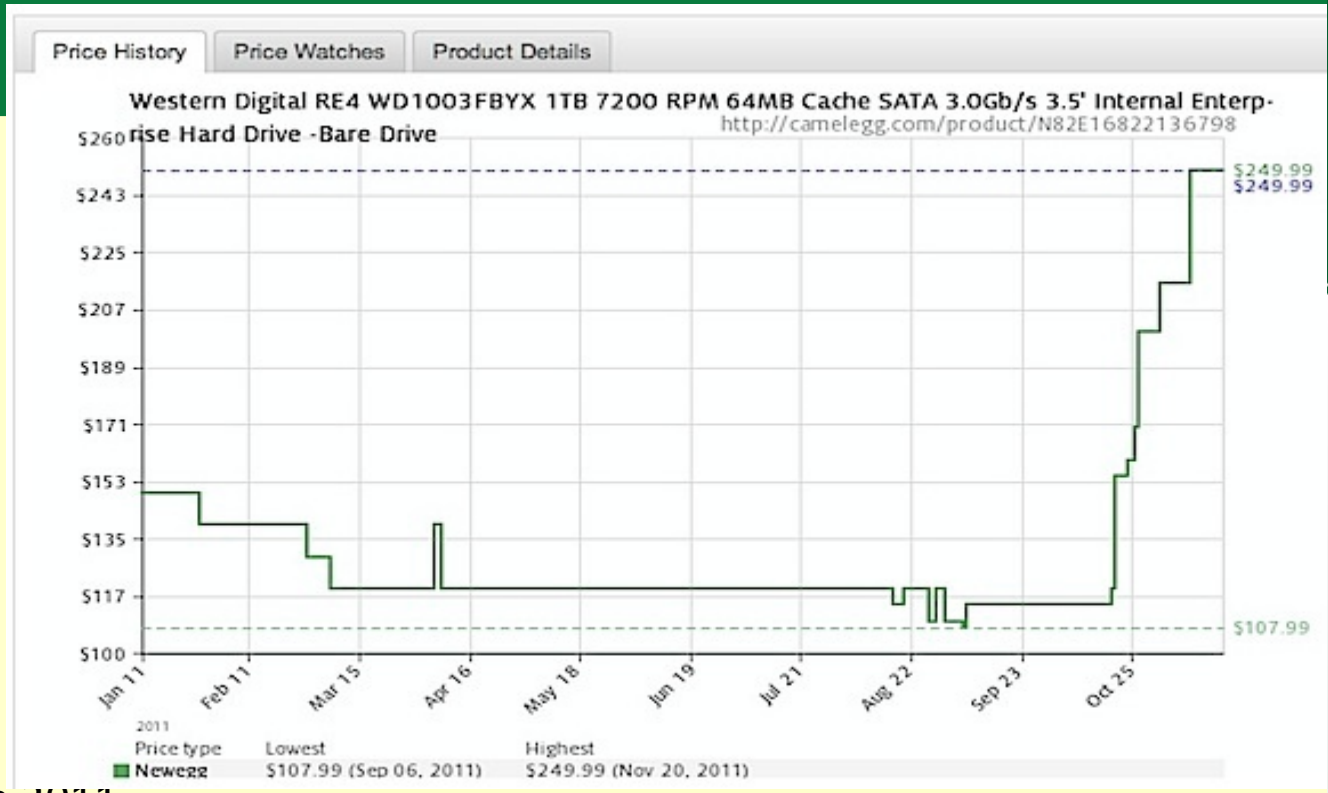

Is it big deal to resolve?

- Not so easy in long-term
  - amount of data is rising Exp

KIWI costs:

- Just HDD prices
  - Cost per GB is falling
  - But can be exceptions like in 2011
- Disk array (has physical limits on HDD count)
- Branded hardware expert service- Special price in remote DC!
- Sysadmins/DBA time - to make new HDD space available for DB/APP
- **Business downtime – lost revenue + penalties.**
  - **Does Your company know the real cost of 1 hour of Downtime? 10 hours?**

Cost of Space - 1

- **What should be counted to know the cost to leave a database “a bit larger”:**
  - Hardware:
    - just HDD prices – per GB
    - New Storage array or extension
  - Service:
    - Expert time to add more space (physical or virtual)
    - Ordering expences (someone should spend time for this)
  - Downtime – may be required for:
    - Add physical/logical device
    - Extend LVM/partition/file-system

Cost of Space - 2

- **What should be counted to know the cost to leave a database “a bit larger”:**
  - How much **GB needed** more **X** number of **Copies** of data:
    - Master, Stand-By, Accept-Test, EDU, Test, DEV
    - Backup (master, stand-by, off-site copy)
  - More RAM:
    - Larger DB may require larger Buffer Cache
  - More Networking:
    - More I/O + longer network load for backups, server-storage
  - Performance degradation
    - Slower I/O during disks rebalancing
    - More I/O due to semi-empty blocks reads

Make the right choice….

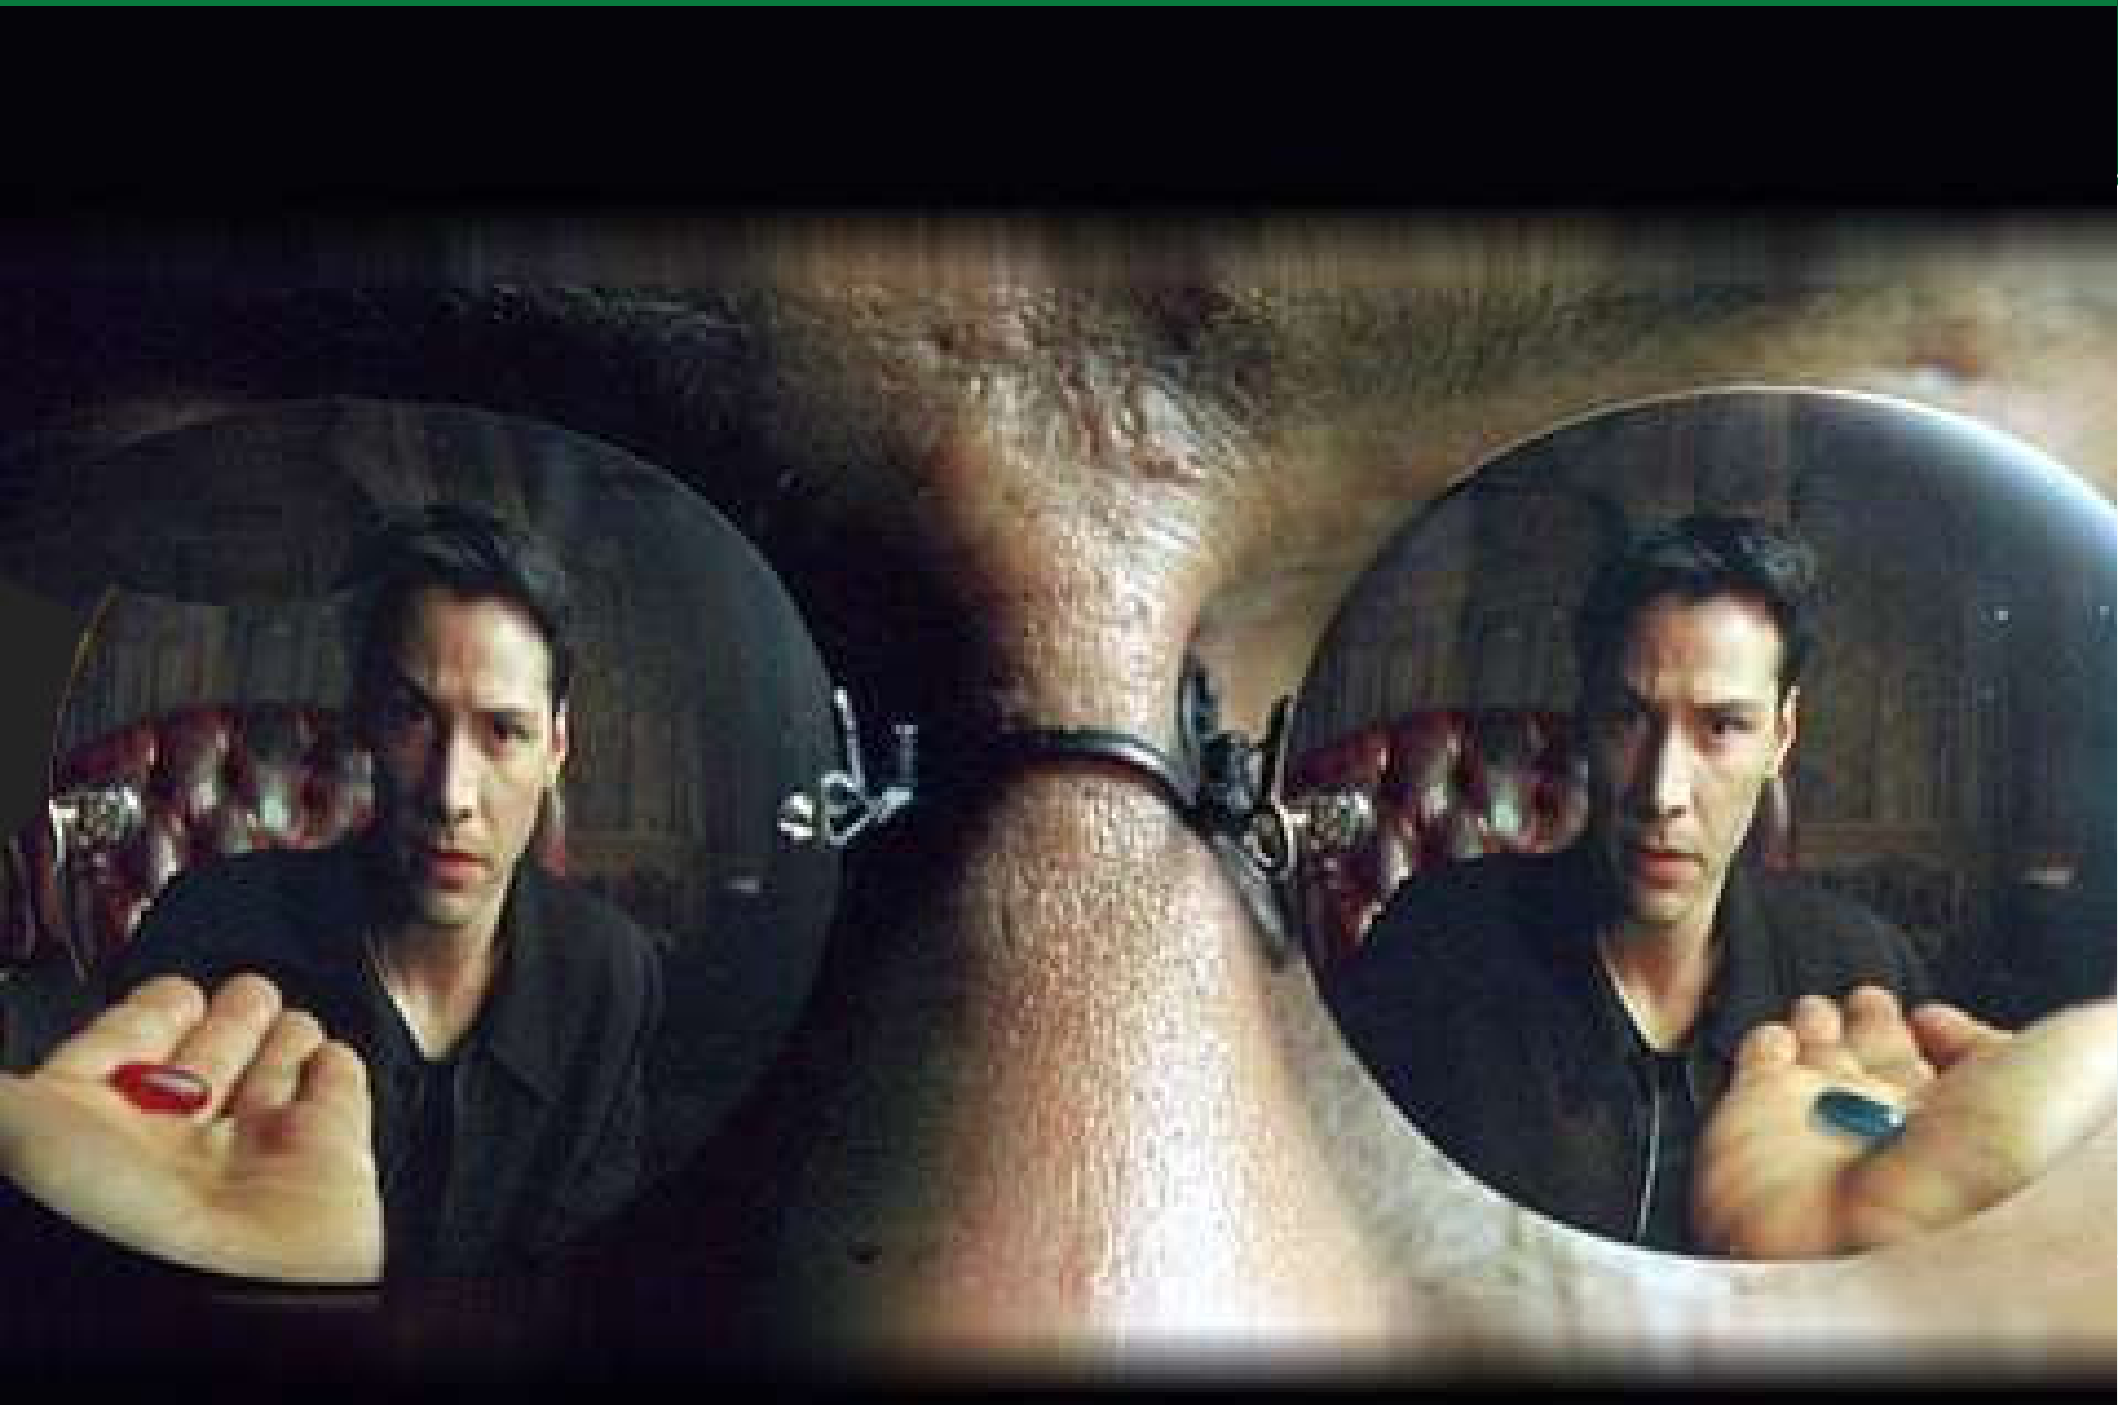

Space Control Concept

Key principles of Space Control (SC)

The Law of Space Dynamic:

In isolated System, without external impact (like adding new HW, etc...) -
disk space does not appear,
and does not disappear
– it just distributed between objects.

```
space provided (Server) = space allocated + free (Client)
```

- There is no "silver bullet" - we should understand how it works (how System use space)
- Holistic approach - we should control **ALL** elements of the system
- System approach - all elements are related and impact each other on many layers
- 1st focus on elements with top usage or/and growth for period
- Do check-outs and implement solutions on regular basis
- Automate what is reasonable

Who cares?

Why me? Once more - “You are the One”! Responsible for the Results!

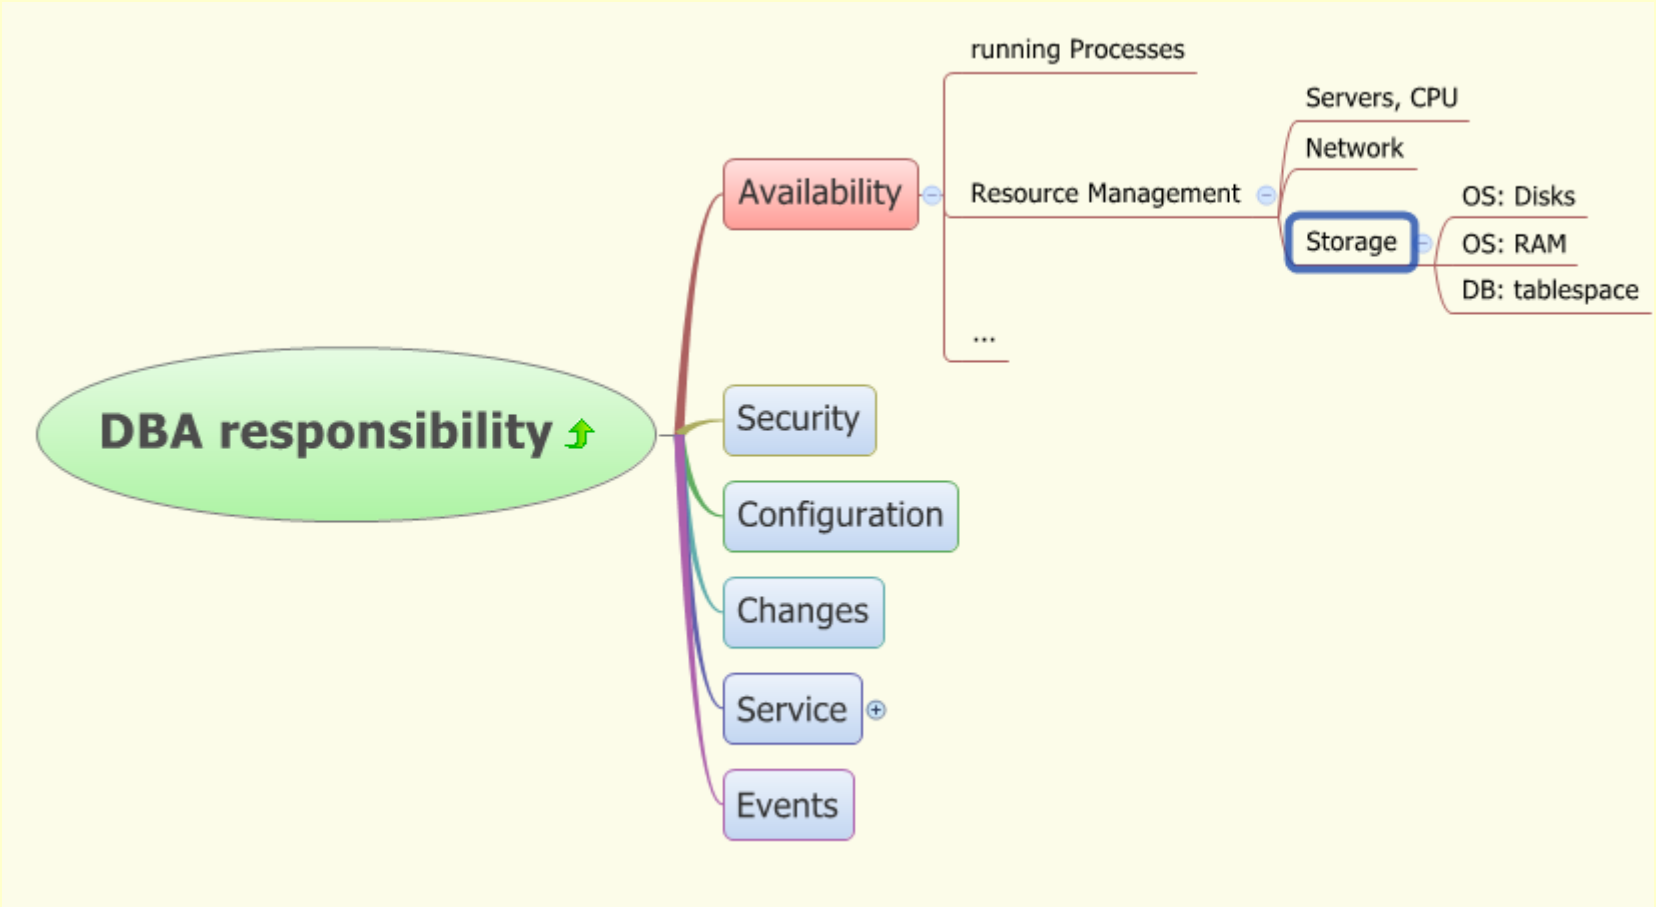

System Life Cycle

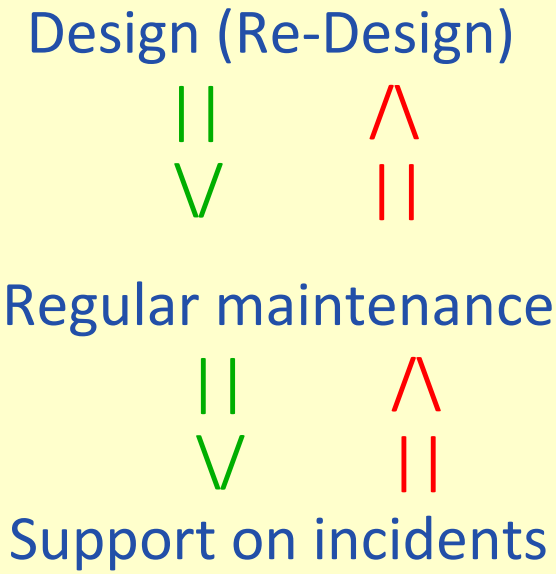

Used space reduction Technologies

How to make DB smaller - techniques

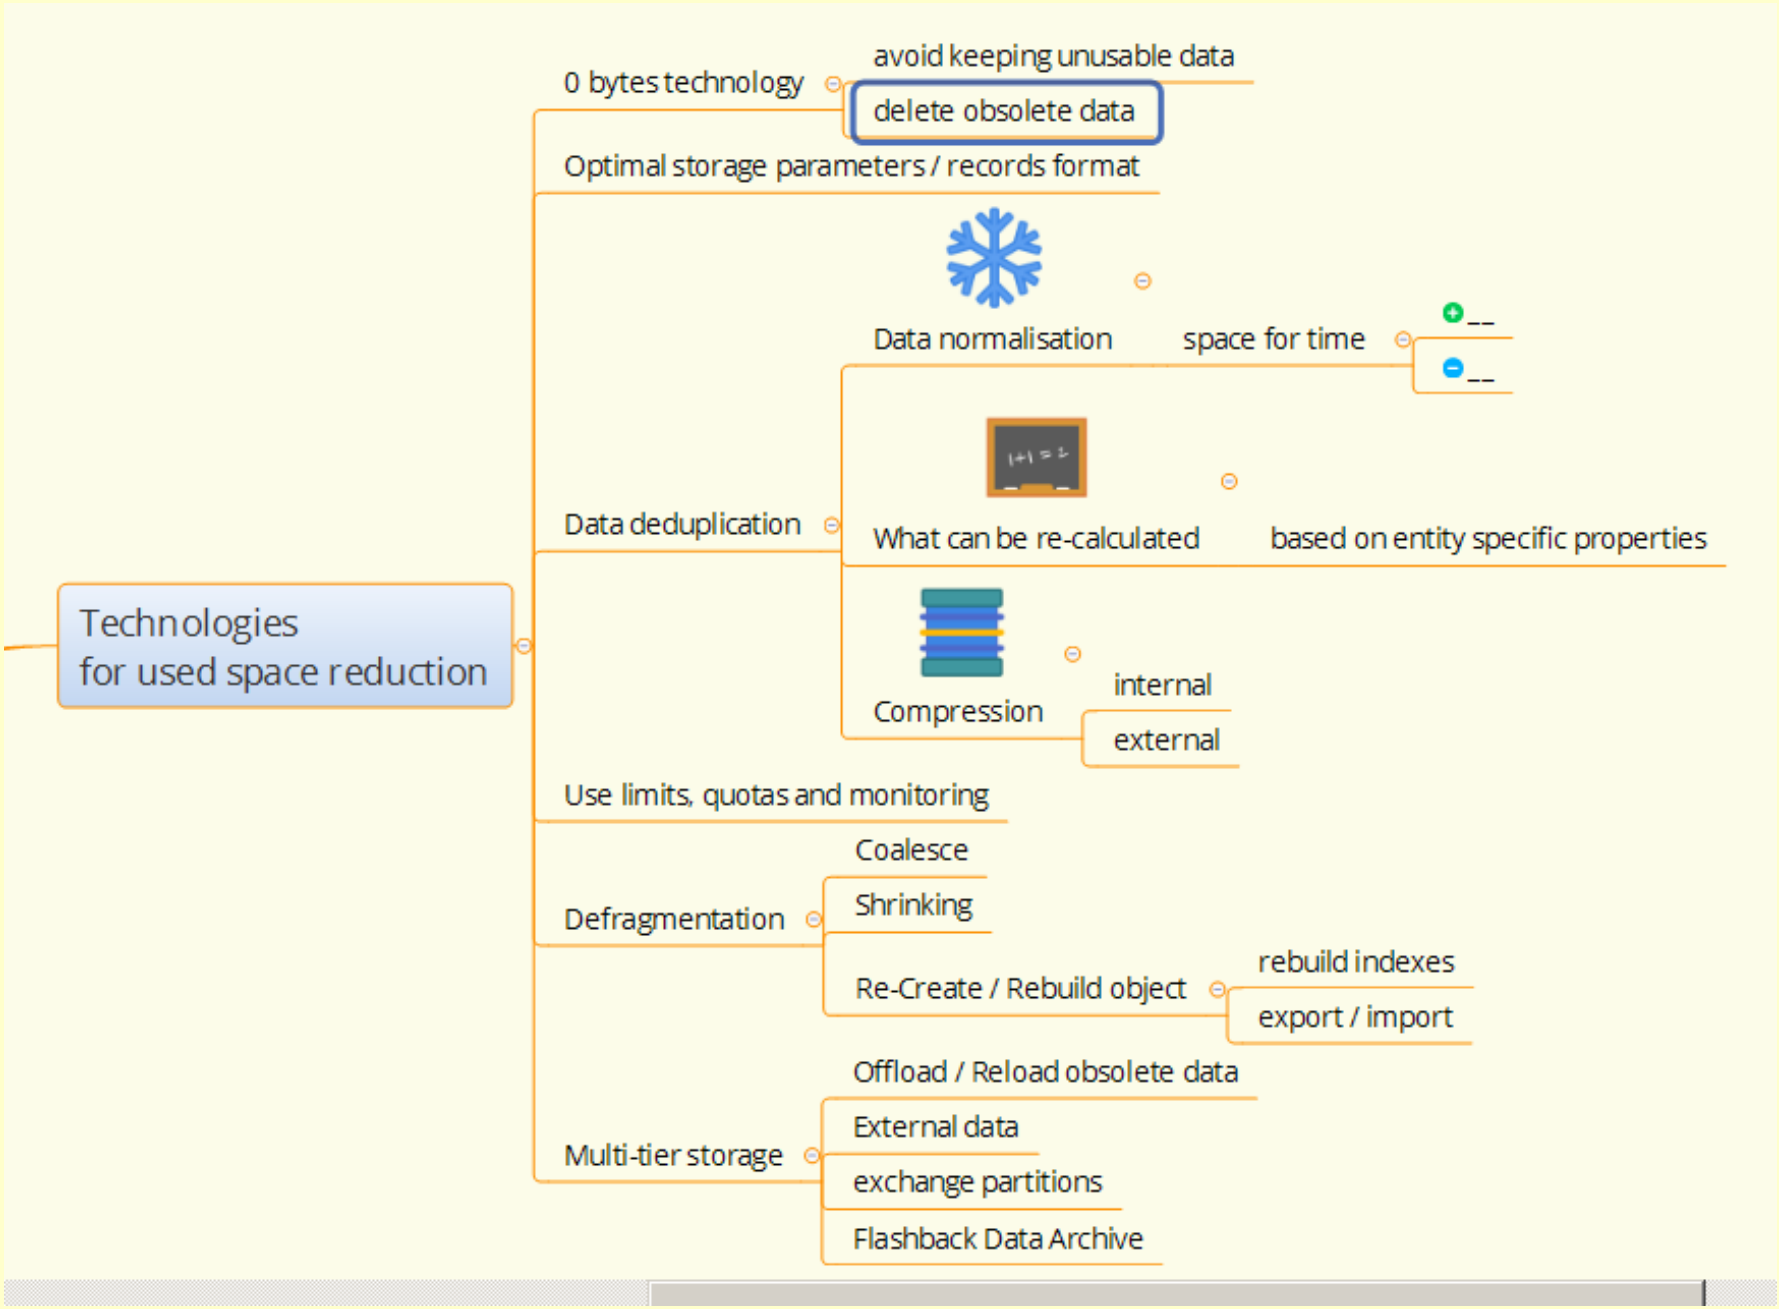

Apply techniques
to every
System
Life Cycle phase:

- design and development
- testing
- maintenance
- support

Design phase impact – like a car

- Car design impact
  - Engine power, max speed
  - Gas consumption
  - Special maintenance required
  - Driver, passengers comfort
- Database design impact:
  - Planned performance, data volumes to process
  - Response time
  - I/O volume, minimal I/O rate required
  - Special maintenance required
  - OLTP/DWH → response time
  - Space required for data, undo, temp, RAM….

Maintenance phase impact – like a car

- Car good/bad maintenance impact:
  - Safety
  - Gas consumption
  - Real speed
  - Unexpected troubles
- Database good/bad maintenance impact:
  - Resources consumption
  - Real response time
  - DRP
  - Response time
  - Space required
  - Unexpected troubles

Support phase impact – like a car

- Having good Car service impact:
  - Fast problems diagnostics
  - Fast parts replacement
  - How long fast “fix” will work until next incident
  - Replacement car service for time of “recovery”
- Having good DB/OS Support tools impact:
  - Time for problems diagnostics
  - Downtime time
  - How often issues will repeat again
  - “Temperature” of Service Desk phone

Cost of Space Reduction

- **What should be counted to know the cost of keeping a database smaller:**
  - Service:
    - Expert time to design, support special maintenance tasks
    - Expert time to execute/control execution
  - Hardware:
    - Temporary space that may be required
    - Specific equipment
  - Performance:
    - More CPU, I/O for time of spec. maintenance tasks
    - Tables, indexes locking

To reduce or not to reduce…?!

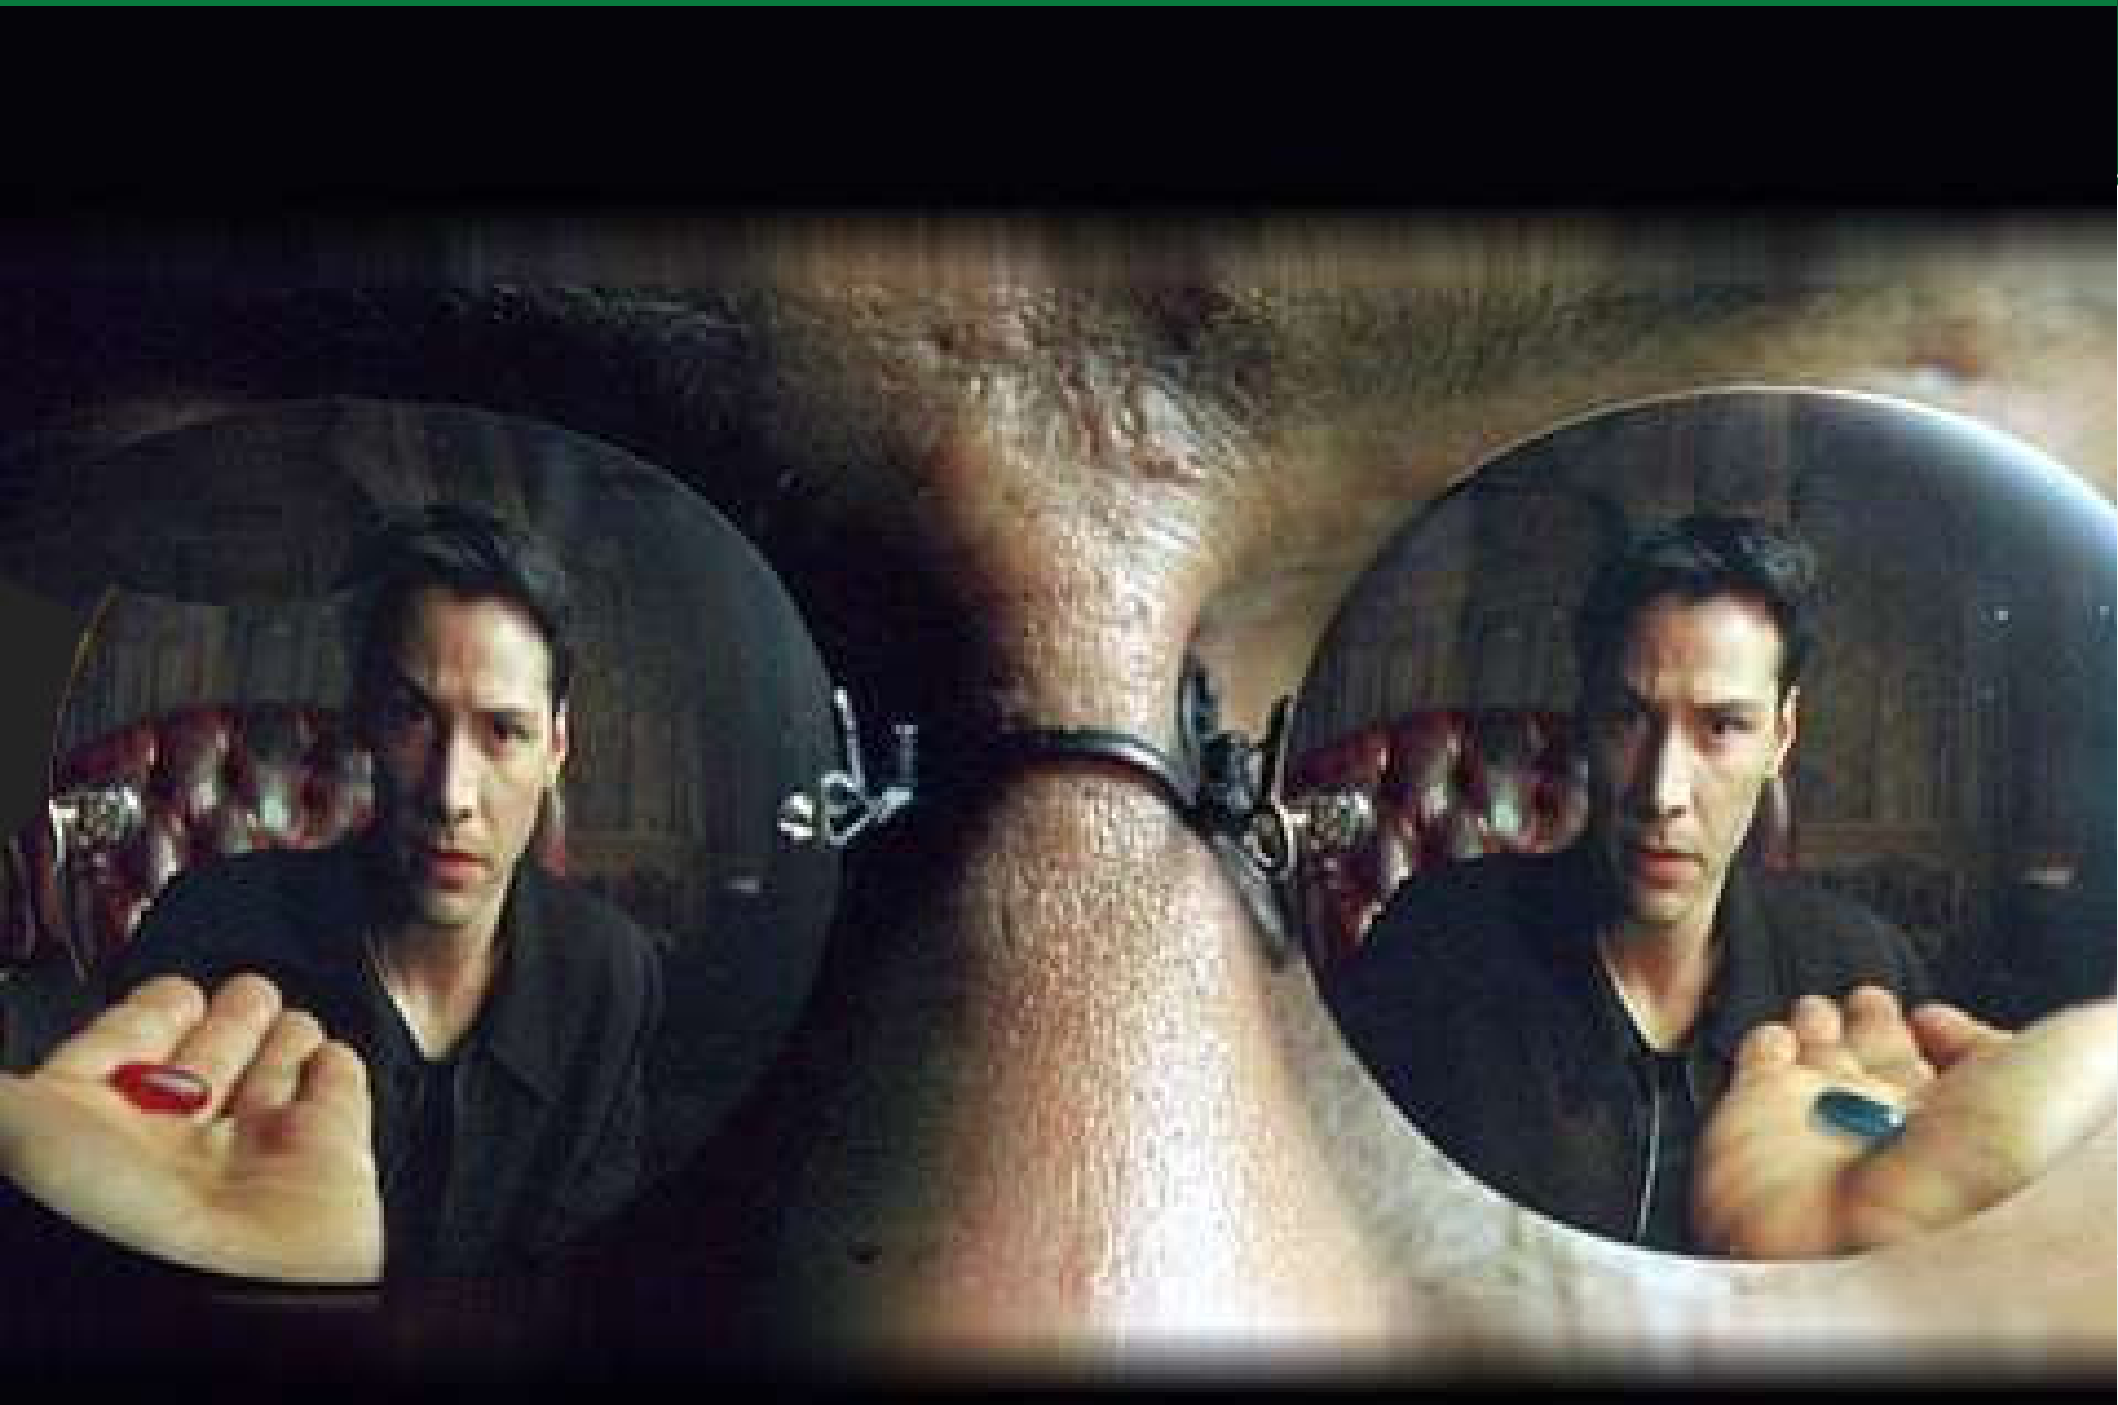

Why this is still an issue?!

- Too seldom incidents
  - This really keeps admins unprepared
  - This cause significant impact on business
  - Until incident fixed, then forgotten until next time (up to 3-5 times)

  because:

  It is always possible to free or find more space in the system!
  (in most cases)

- Too complex (many components, many layers to manage)

Space Architecture

- 7 level vertical model

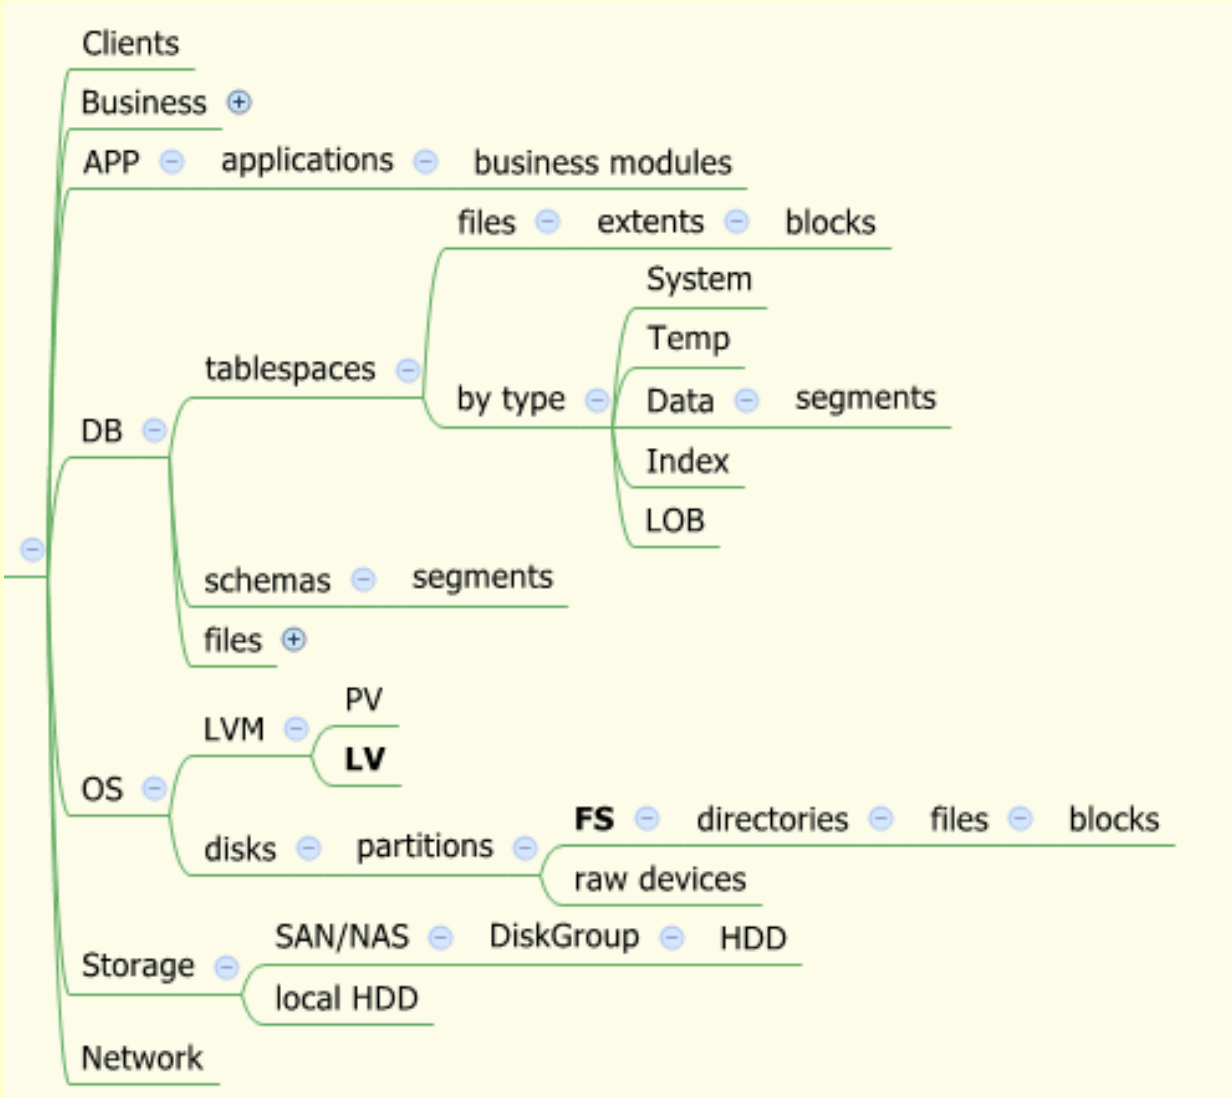

How to stay alive...

and

make life better

Live is possible (after 50+ too)!

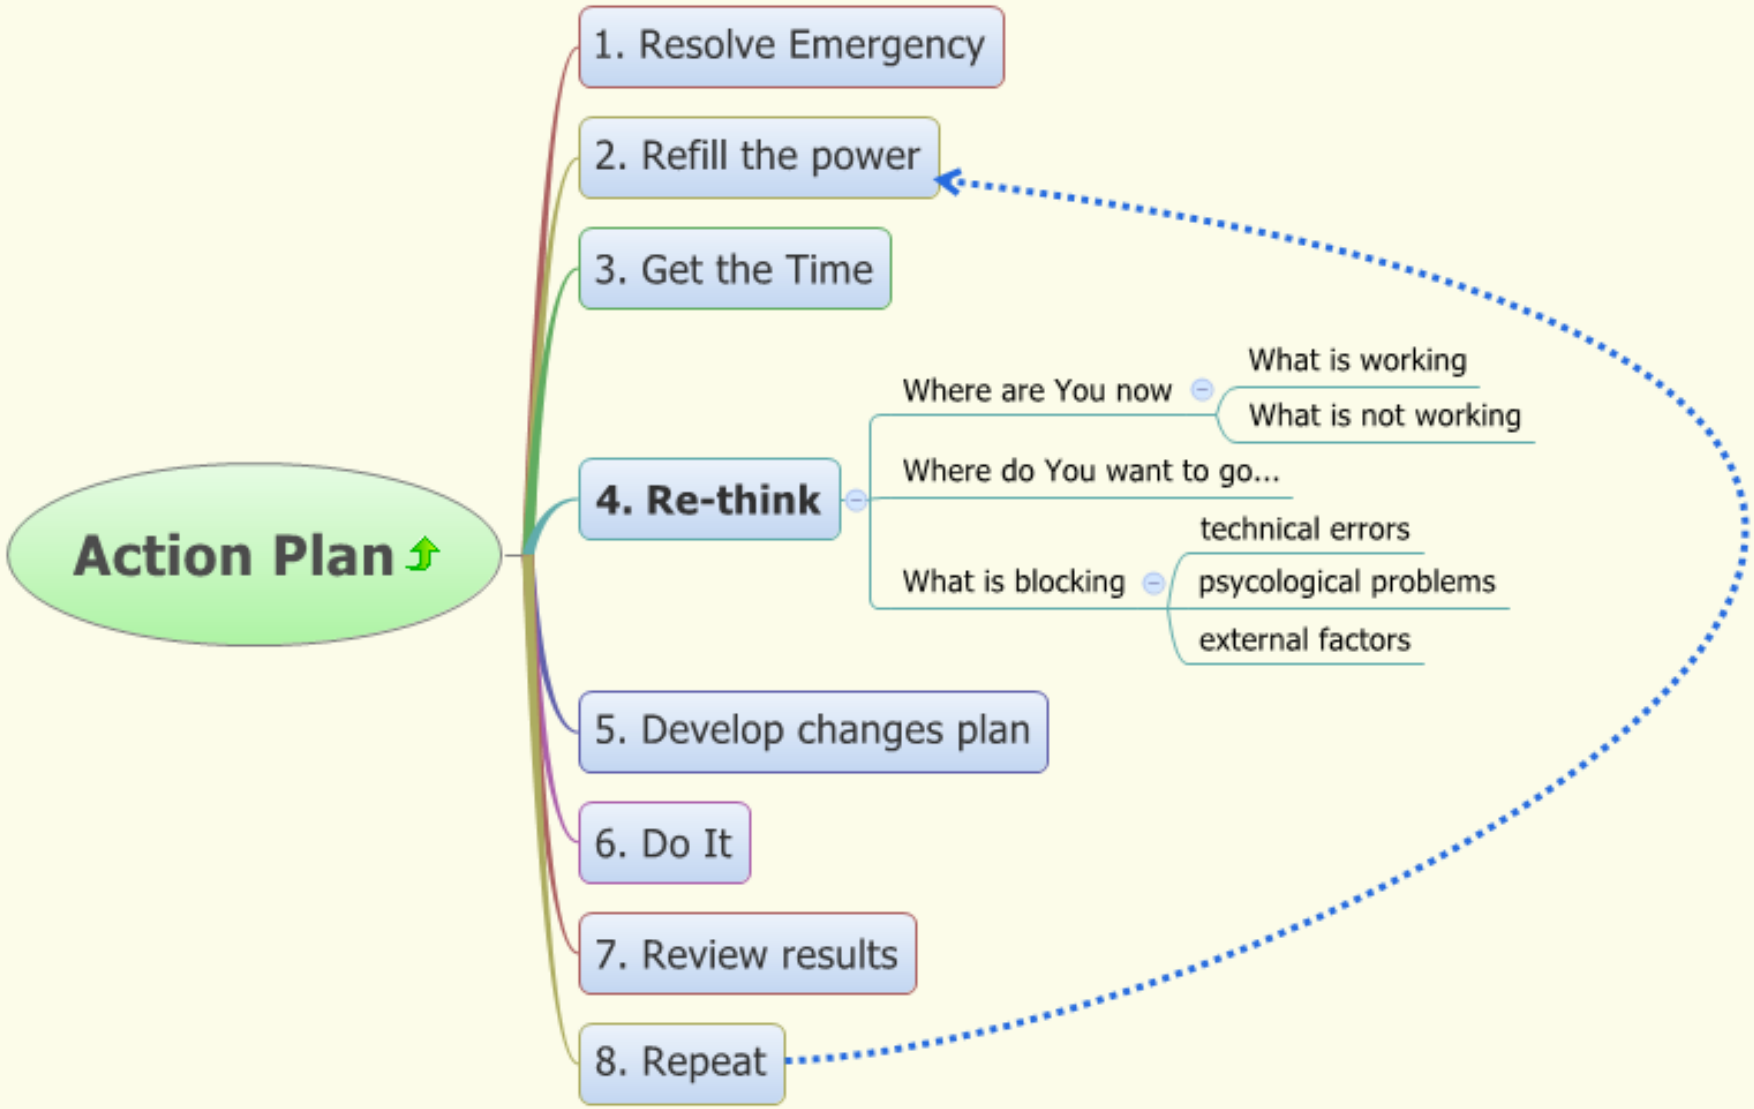

Pro-Active SC - workflow

- Workflow:

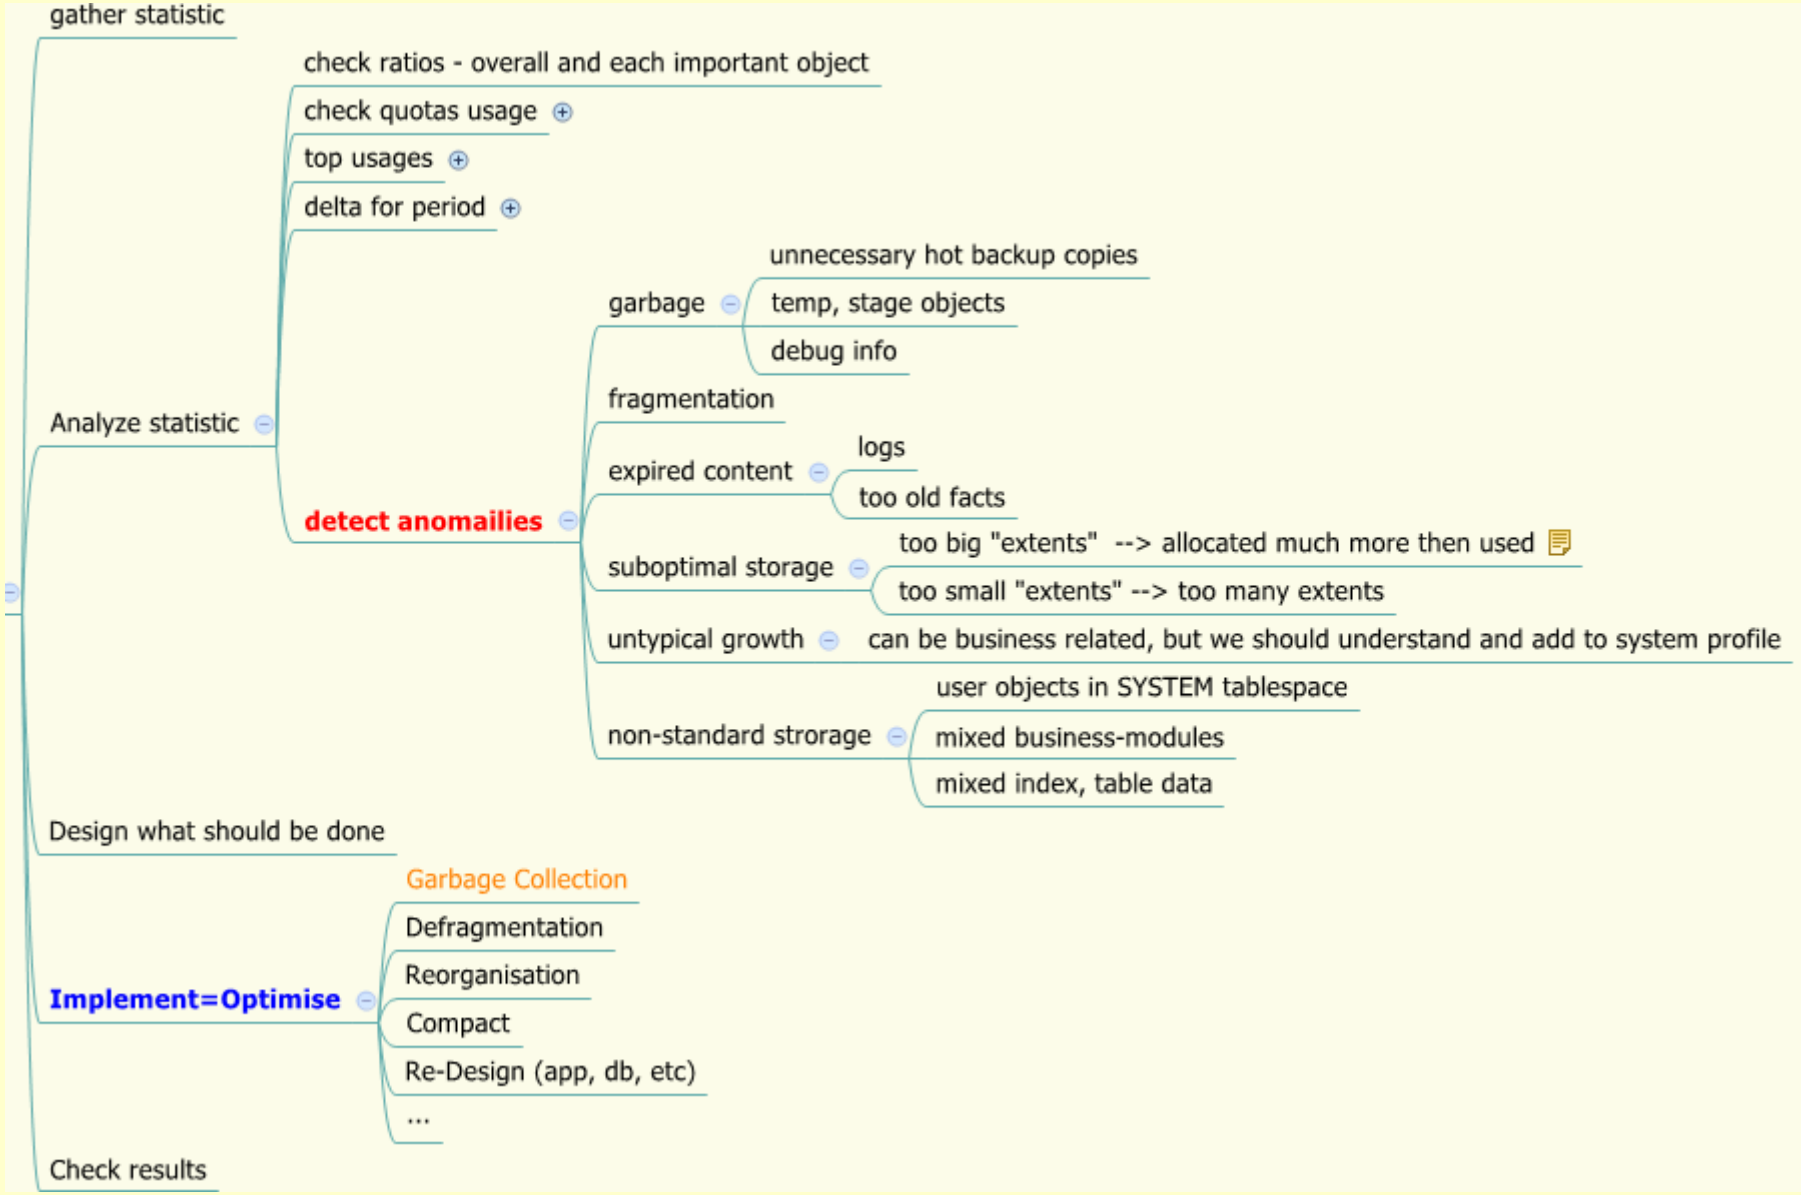

Re-Active SC - workflow

- Workflow:

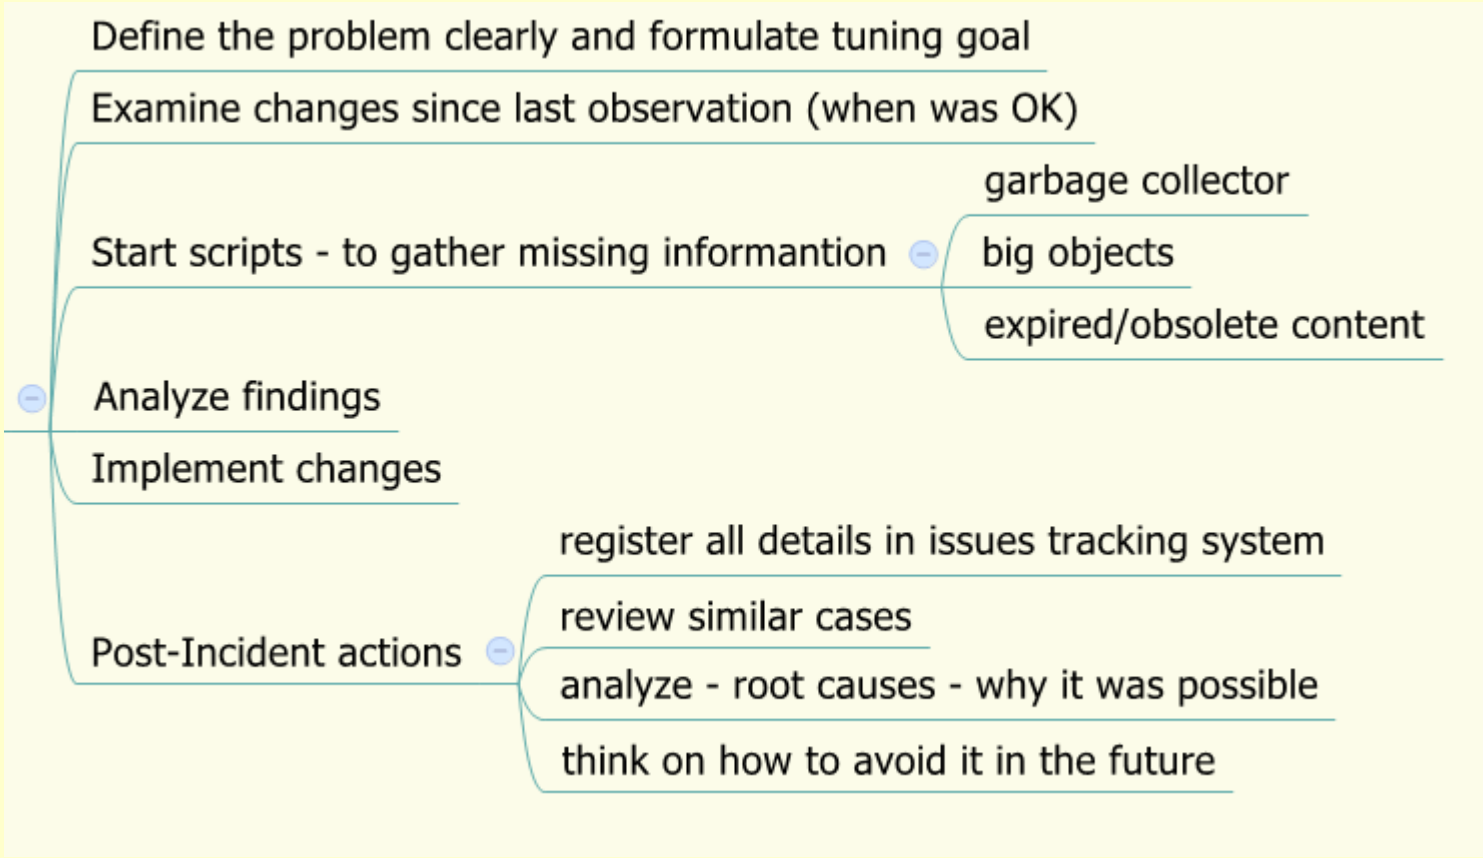

12 + N ways
to “lose weight”

How to make DB smaller - techniques

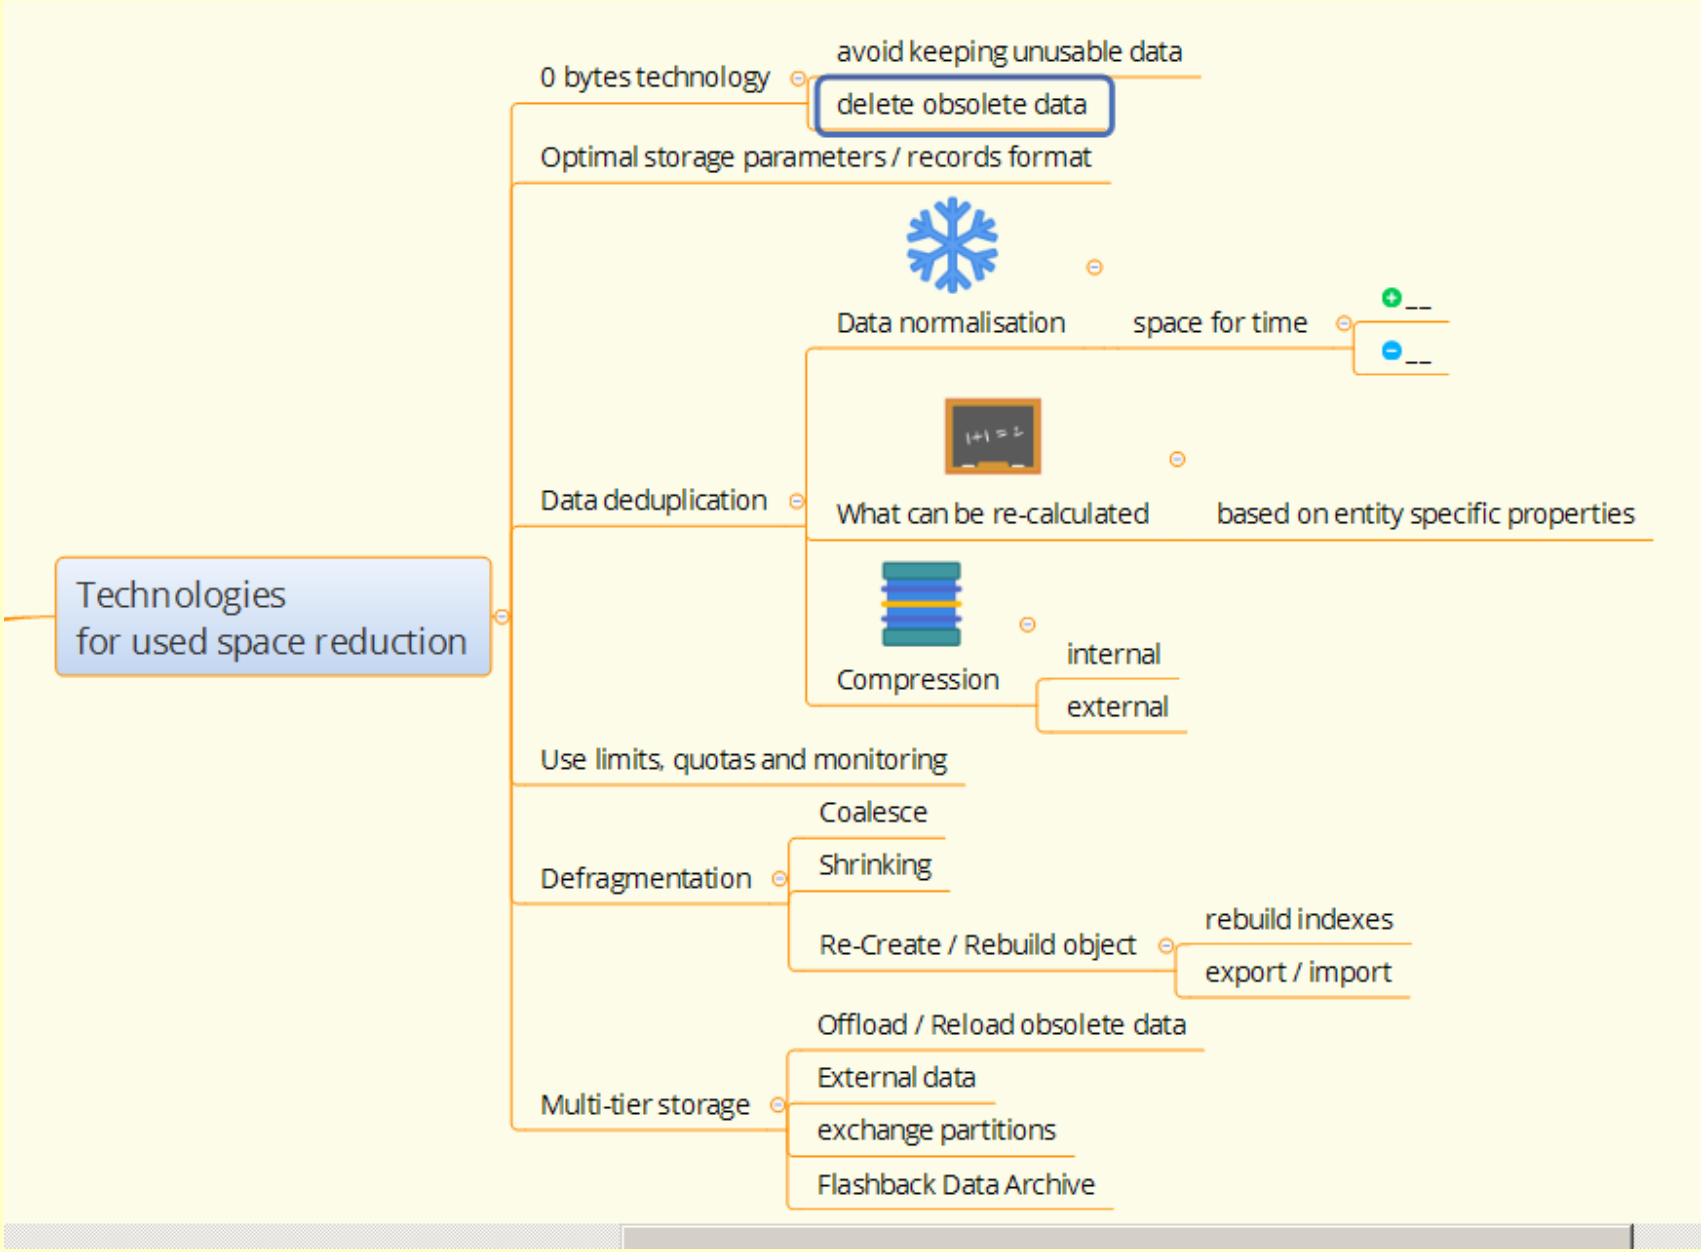

What can be done – database:

Prevent problems:

- set proper DEFAULT TABLESPACE for users (not default=SYSTEM!!!)
- set quotas for users
- set maxsize limits for datafiles
- track space usage, growth
- adjust retention policy for audit data, keep AUDIT data of of SYSTEM tablespace
- reasonable audit (not ALL – 11g default audit can burst SYSTEM tablespace!)
- avoid 1-side balanced indexes

0 bytes technology:

- garbage collection (dba_segments where segment_name like '%TMP%' or..)
- drop unused indexes (enable index usage monitoring)
- purge or offload obsolete data (debug, audit, temp, transition, etc.)

What can be done – database:

Defragment, re-create:

- shrink segments (10g+)
- coalesce + deallocate unused (8i+)
- rebuild indexes, IOT online/offline (with pctfree=1 or even pctfree=0)
- defragment tablespace (move/rebuild to stage tbs, back to original/rename)
- defragment tablespace (exp/imp if have LONG fields)
- re-create/move segments in tablespace with optimal UNIFORM size extents (for example: 4K is OK for StatsPack)
- re-create/move segments to tablespace with compression (! have limits)
- re-create TEMP tablespace
- re-create UNDO tablespace
- re-create REDO logs with smaller size

What can be done – database:

Offload data (prepare to reload):

- export (+ compress dump file + named pipes if required) + delete
- external files feature
- exchange (swap) partitions (partition tablespace files with empty partition)
- database on other server (access over database link)

Other:

- move files to less expensive storage
- find space on other file system partitions
- keep in order all file system partitions (may be used space during emergency)
- compress backups (if possible on the fly – build in or via named pipes for exp)

What can be done – OS:

Just some hints:

- daily purge/compress/offload obsolete data (audit, trace, log-files, temp, etc.)
- check time to time for “lost” stuff by staff (*.gz, *.tar, rpm-s, etc)
- monitor disk space usage (space, inodes) by mountpoint (most monitoring tools can do this)
- track space usage by directory/file
- remove 5% root reservation for data mountpoints (tune2fs)
- clean /mail/spool/username
- user quotas may be employed at OS level too
- compare database data directories size with dba_data|temp_files

Hints

- "Missing" space in Linux:

System 1

```
[root@abc-monika00 ~] df -k
Filesystem           1K-blocks      Used Available Use% Mounted on
/dev/simfs            10485760   2996300   7489460  29% /
none                    287144         4    287140   1% /dev
```

System 2

```
[root@abc-oracle04 ~]$ df -k
Filesystem                 1K-blocks      Used Available Use% Mounted on
/dev/mapper/VolGr40LV00     36832848  13158184  21803612  38% /
/dev/sda1                     295561     19479    260822   7% /boot
none                         4080932         0   4080932   0% /dev/shm
/dev/mapper/VolGr320-s01   309604352 274390644  19486708  94% /s01
/dev/mapper/VolGr9-s02      10288760   8432824   1333288  87% /s02
/dev/mapper/VolGr30-s03     30930940  13741012  15618704  47% /s03
```

Hints

- “Missing” space in Linux (2 different servers):

Filesystem	1K-blocks	Used	Available	Use%	Mounted	Missing	Miss %
/dev/simfs	10485760	2995120	7490640	29%	/	0	0.00%
Filesystem	**1K-blocks**	**Used**	**Available**	**Use%**	**Mounted**	**Missing**	**Miss %**
/dev/mapper/VolGr40-LogVol00	36832848	13158184	21803612	38%	/	1871052	5.08%
/dev/sda1	295561	19479	260822	7%	/boot	15260	5.16%
/dev/mapper/VolGr320-s01	309604352	274390644	19486708	94%	/s01	15727000	5.08%
/dev/mapper/VolGr9-s02	10288760	8432824	1333288	87%	/s02	522648	5.08%
/dev/mapper/VolGr30-s03	30930940	13741012	15618704	47%	/s03	1571224	5.08%

Q: Who is using 5%?

A: reserved by ext3/ext4 file system

Info: man **tunefs** (or **tune2fs**)

Hint: When totally out of space, try reduce percent of disk space reserved by file system for **root**:

```
tune2fs -m pct_reserved /device-path
```

Hints

- OS – Top usage

`du -xb /u01/ | sort -k 1g`

```
[root@abc-oracle04 temp]$ du -xb /u01/ | sort -k 1g | more
4096    /u01/app/cominder/work
4096    /u01/app/oracle/admin/ATEST/cdump/core_2148
4096    /u01/app/oracle/admin/ATEST/cdump/core_8818
4096    /u01/app/oracle/admin/ATEST/cdump/core_9815
4096    /u01/app/oracle/admin/PRODDB/cdump
4096    /u01/app/oracle/product/9.2.0.8/assistants/dbca/logs
...
309856116       /u01/app/oracle/product/9.2.0.8/assistants
487250056       /u01/inst.ora/rdbms/9204_64/9208
1982678664      /u01/inst.ora/rdbms/9204_64
2064689619      /u01/inst.ora/rdbms
2187986576      /u01/inst.ora
2299627375      /u01/app/oracle/product/9.2.0.8
3255055872      /u01/oradata/ATEST
3255064064      /u01/oradata
3396730237      /u01/app/oracle/product
3439307656      /u01/app/oracle
3856168124      /u01/app
9299222860      /u01/
```

Hints

- OS – Who is growing (also right now)

```
   du -xb /u01/ | sort -k 2 > before.out
… wait some time or do something in FS ...
   du -xb /u01/ | sort -k 2 > after.out
… compare output ...
  diff before.out after.out | grep '/' | sort -k 3
```

```
< 9245879883    /u01/
> 9292575701    /u01/
< 3820439441    /u01/app
> 3820754656    /u01/app
< 424912        /u01/app/cominder
> 738124        /u01/app/cominder
< 172714        /u01/app/cominder/LOG
> 172948        /u01/app/cominder/LOG
> 317074        /u01/app/cominder/temp
< 4096  /u01/app/cominder/temp
....
< 2170372282    /u01/inst.ora
> 2216752885    /u01/inst.ora
> 20010526      /u01/inst.ora/java.dropme
< 2064689619    /u01/inst.ora/rdbms
> 2091059696    /u01/inst.ora/rdbms
> 26370077      /u01/inst.ora/rdbms/OPatch.dropme
```

OS garbage collector

- OS garbage collector example:

```
TARGET=${1}
echo " Looking for large archives"
find ${TARGET} -name "*.tar" -size +102400000c -exec ls -l {} \;
find ${TARGET} -name "*.gz" -size +102400000c -exec ls -l {} \;
find ${TARGET} -name "*.Z" -size +102400000c -exec ls -l {} \;

echo " Looking for large logs"
find ${TARGET} -name "*log" -size +102400000c -exec ls -l {} \;

echo " Looking for dumps"
find ${TARGET} -name "*dmp"  -exec ls -l {} \;
find ${TARGET} -name "core"  -exec ls -l {} \;

echo " Looking for audit and lost traces"
find ${TARGET} -name "*.aud" -mtime +30 -exec ls -l {} \;
find ${TARGET} -name "*.trc" -mtime +30 -exec ls -l {} \;
```

TOOLS
for space
management

Tools for Design

- Know what your data is – document any changes to Data Model / Segments
- Estimate number of rows, bytes for tables for 1-3-5 years
- Extent and Block Space Calculation and Usage in Oracle Databases [Note ID 10640.1]
- DBMS_SPACE specification – in Oracle documentation
- Compare several options for Master data, Backup, Stand-by, TEST, etc.
- Compare several options for storage hardware
- Index Rebuild, the Need vs the Implications [Note ID 989093.1]
- FORECASTING DATABASE DISK SPACE REQUIREMENTS: A POOR MAN'S APPROACH by Edward L. Trettel

Tools for Maintenance / Support – expectations

- Requirements:

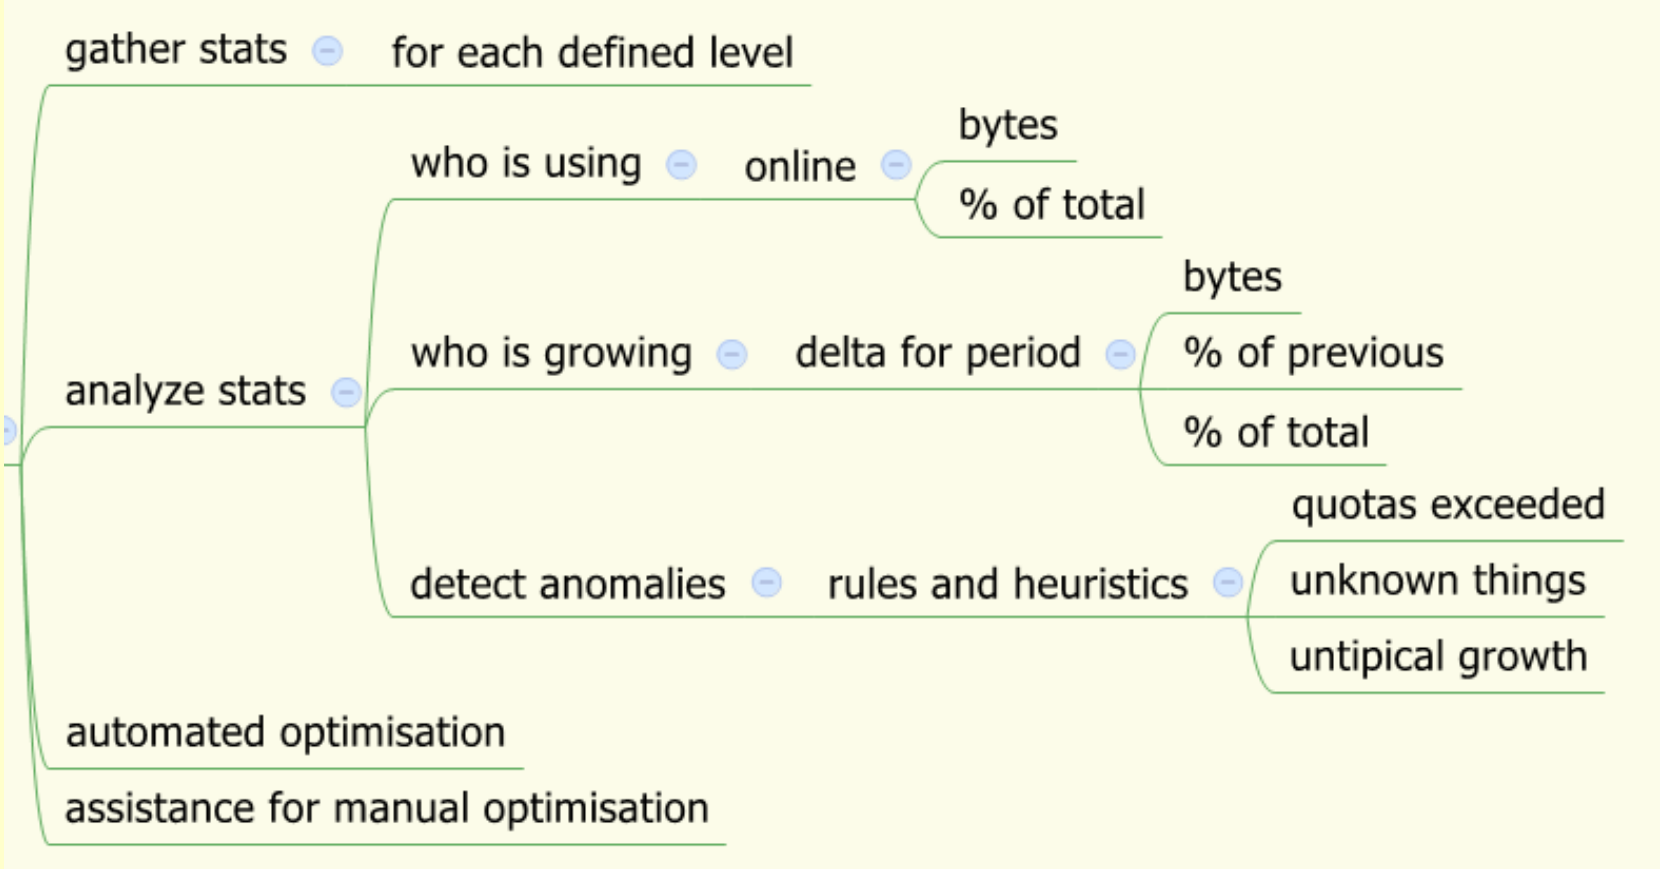

Re-Active SC

- Tools:
  - On-line monitoring
  - Statistics
  - “Who is using” tools
    - OS (win): SpaceSniffer, TreeSize
    - OS (Unix/Linux): du + find
    - DB: select … from dba_segments
  - Scripts
    - to find top usage by...
    - to find anomalies
    - garbage collectors
    - offload data/files
    - backup and remove, or compress

Tools - available

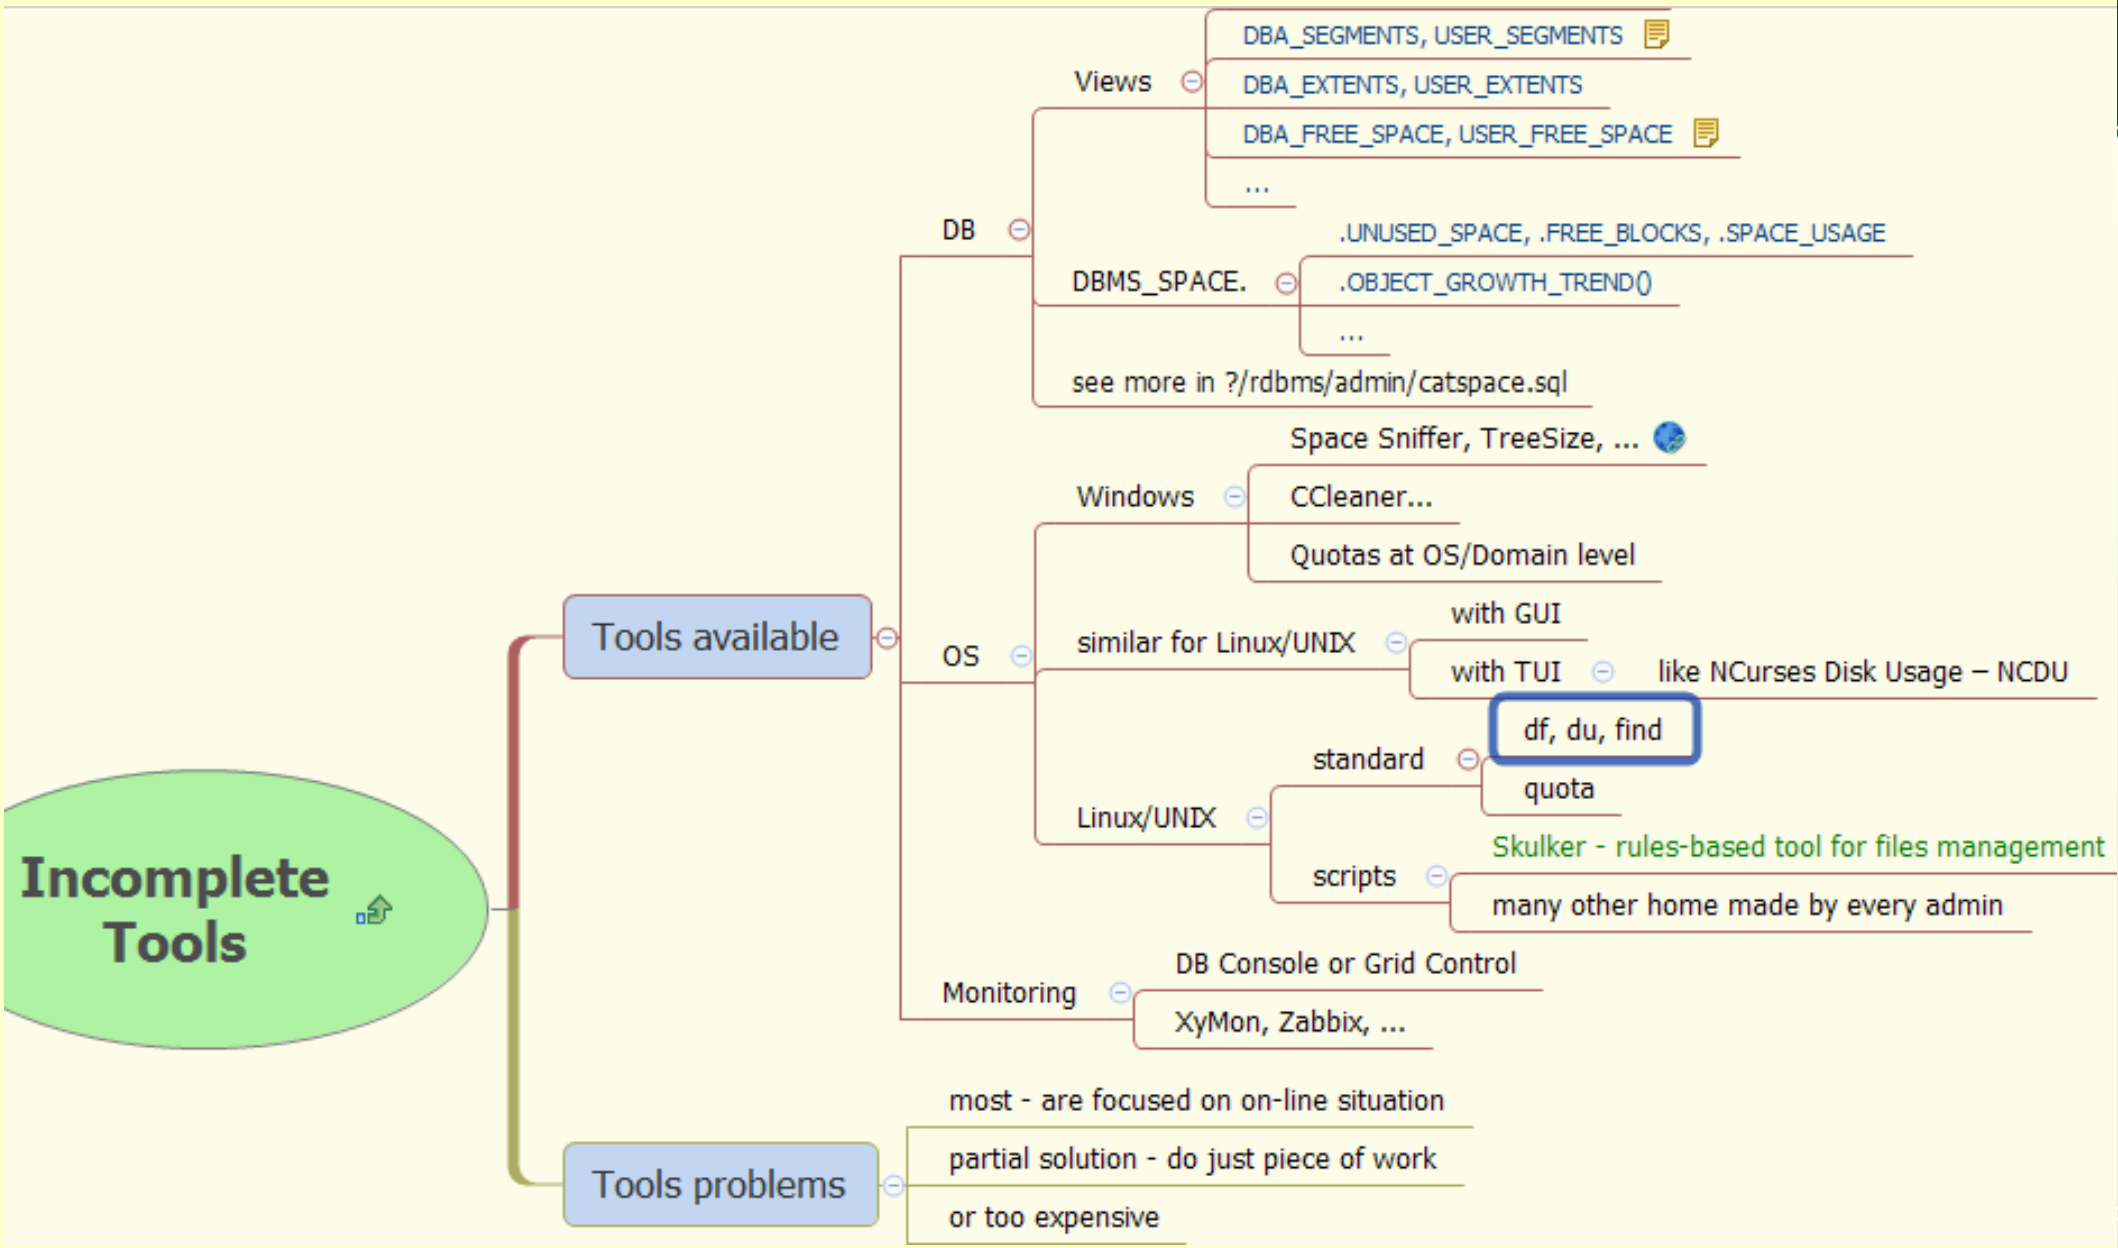

Tools - available

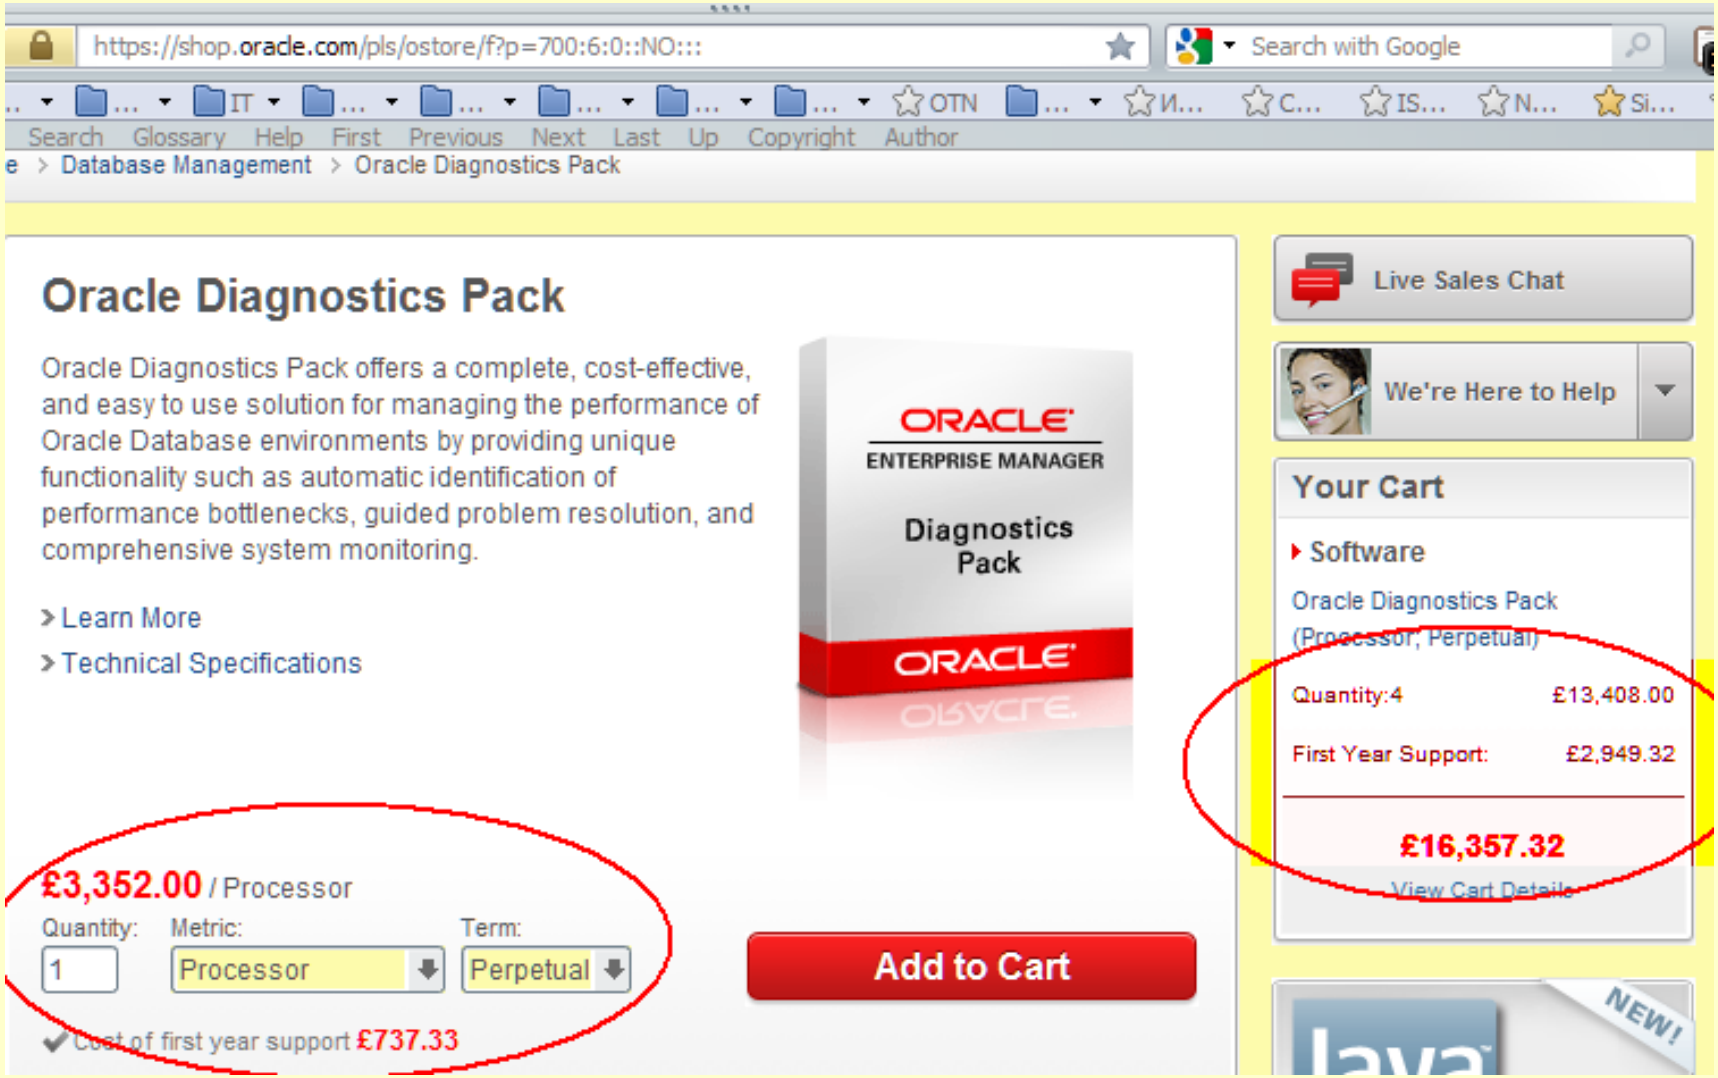

Tools for Support (fire-fighting)

NCurses Disk Usage – NCDU (for OS)

- Collects file-system directories/files size information
- Show this as tree
- Can export/import collected data in JSON format

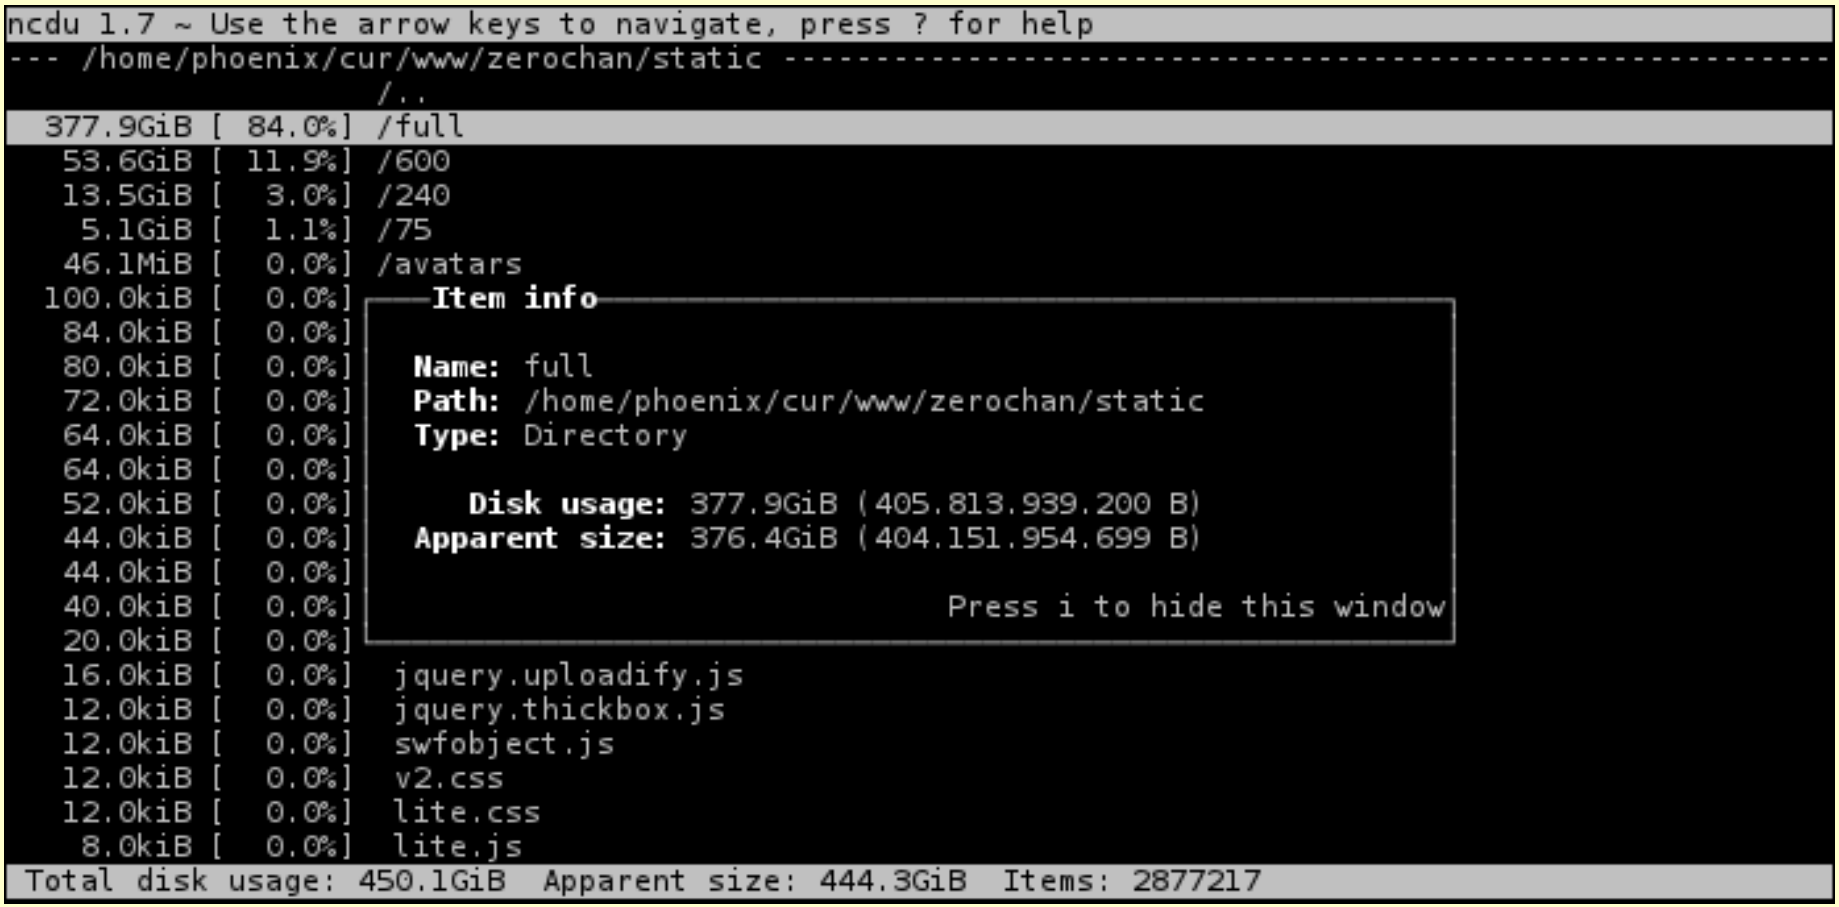

Tools for Support (fire-fighting)

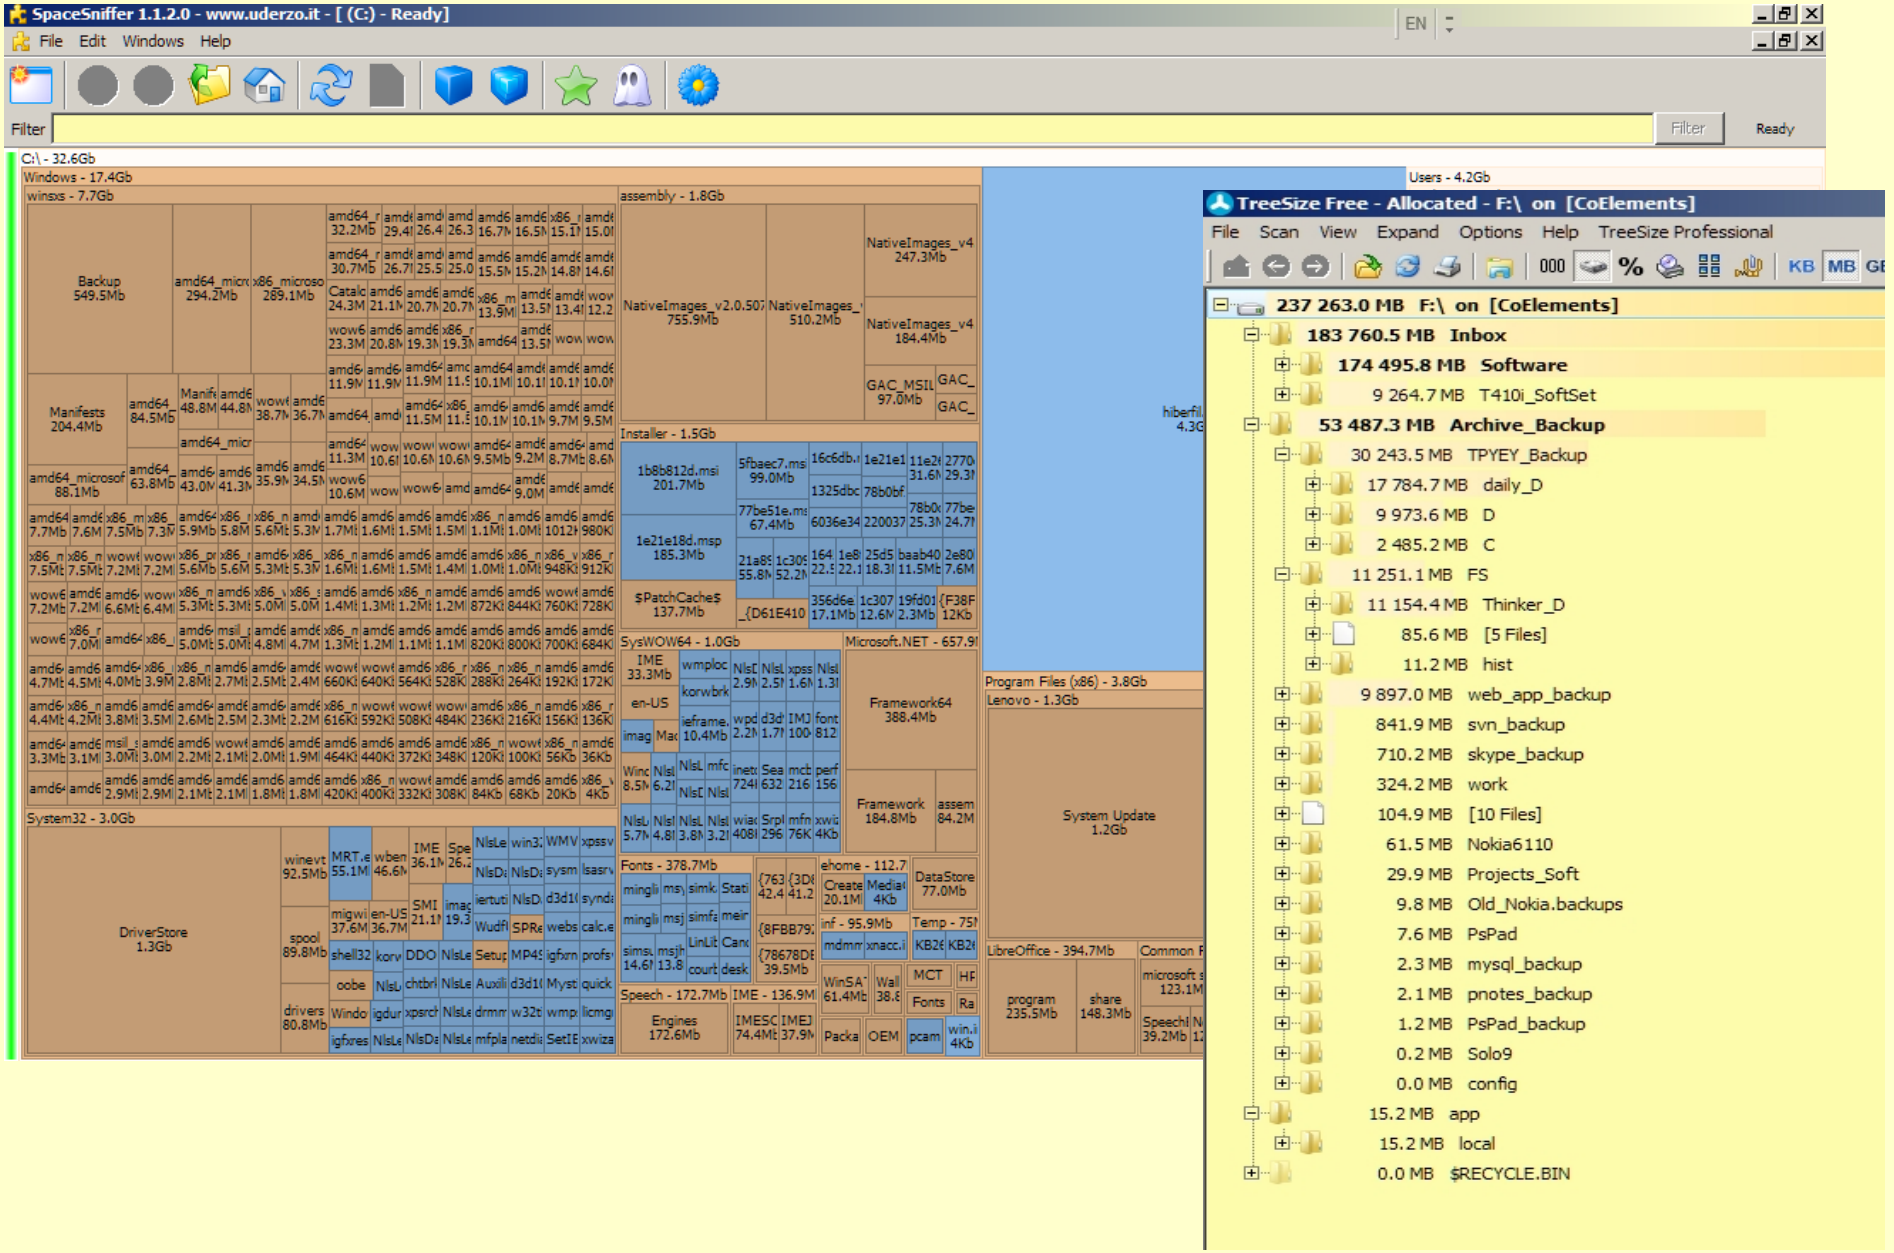

Tools needed for Maint./Support/Re-Design

- Requirements:

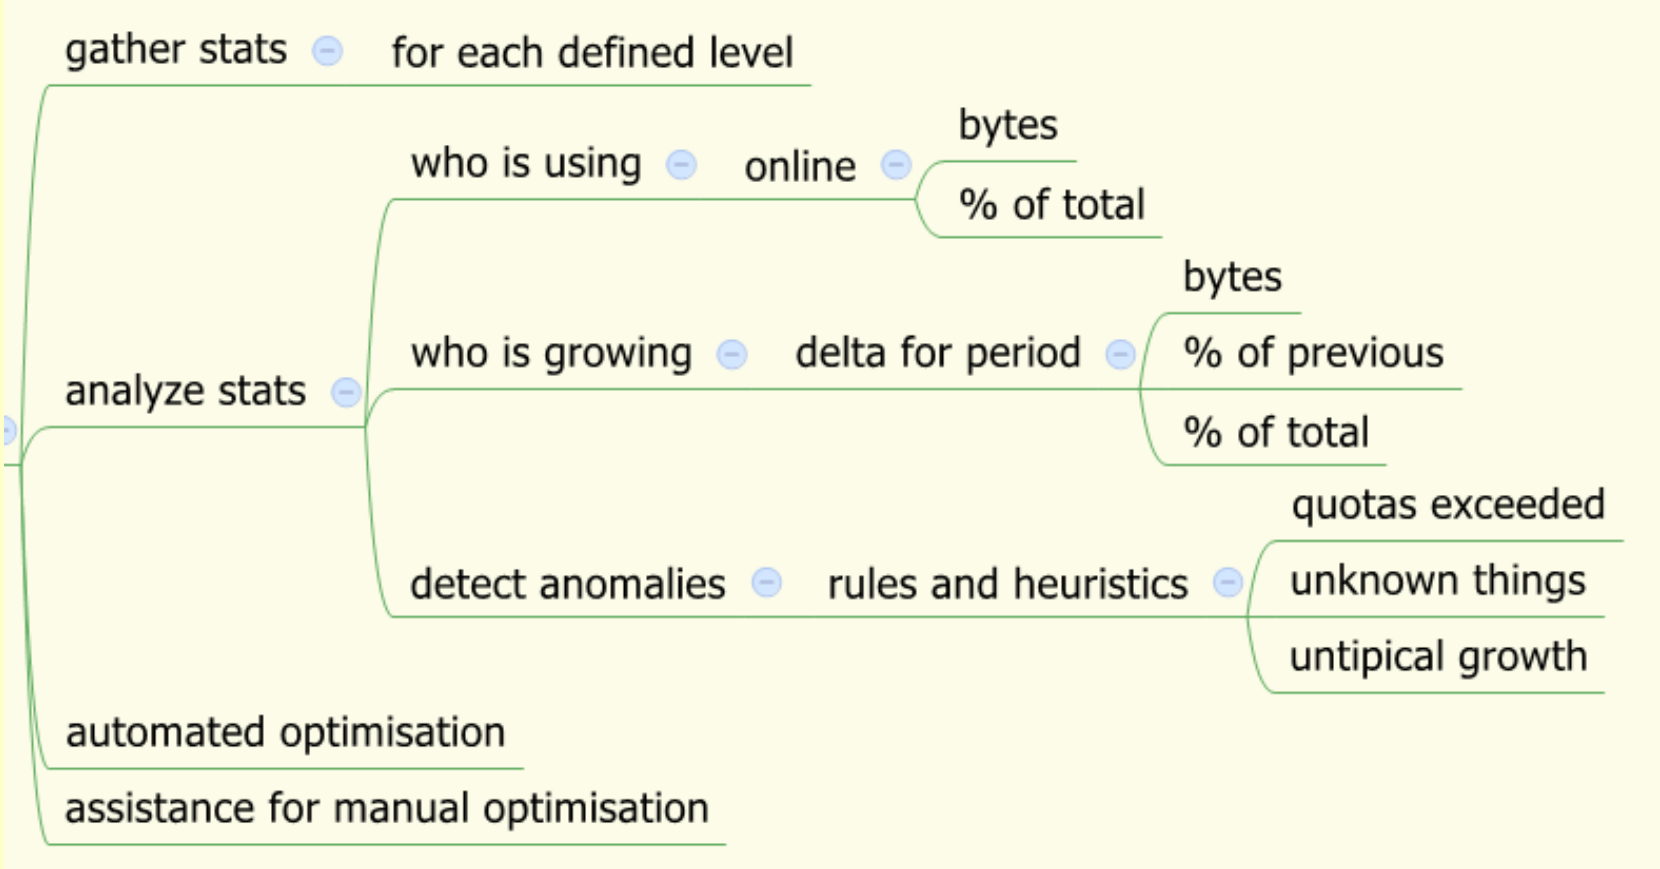

Tools needed for Maint./Support/Re-Design

- **Like what?!**

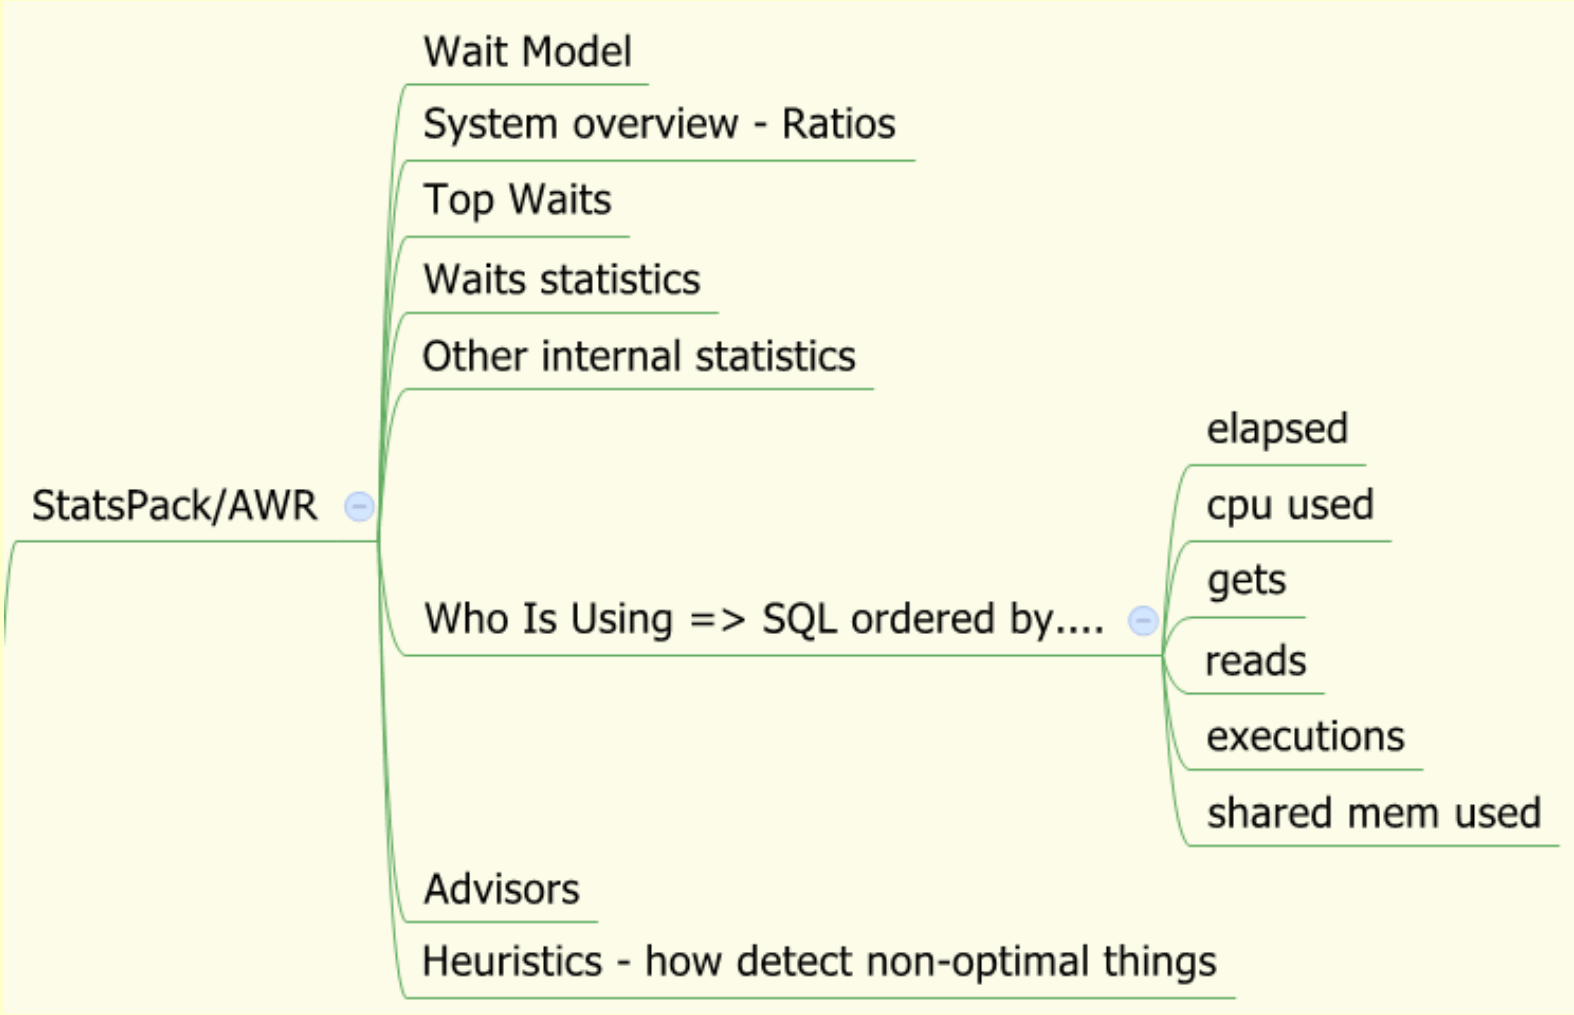

Tools needed for Maint./Support/Re-Design

- We need similar for pro-active space management:

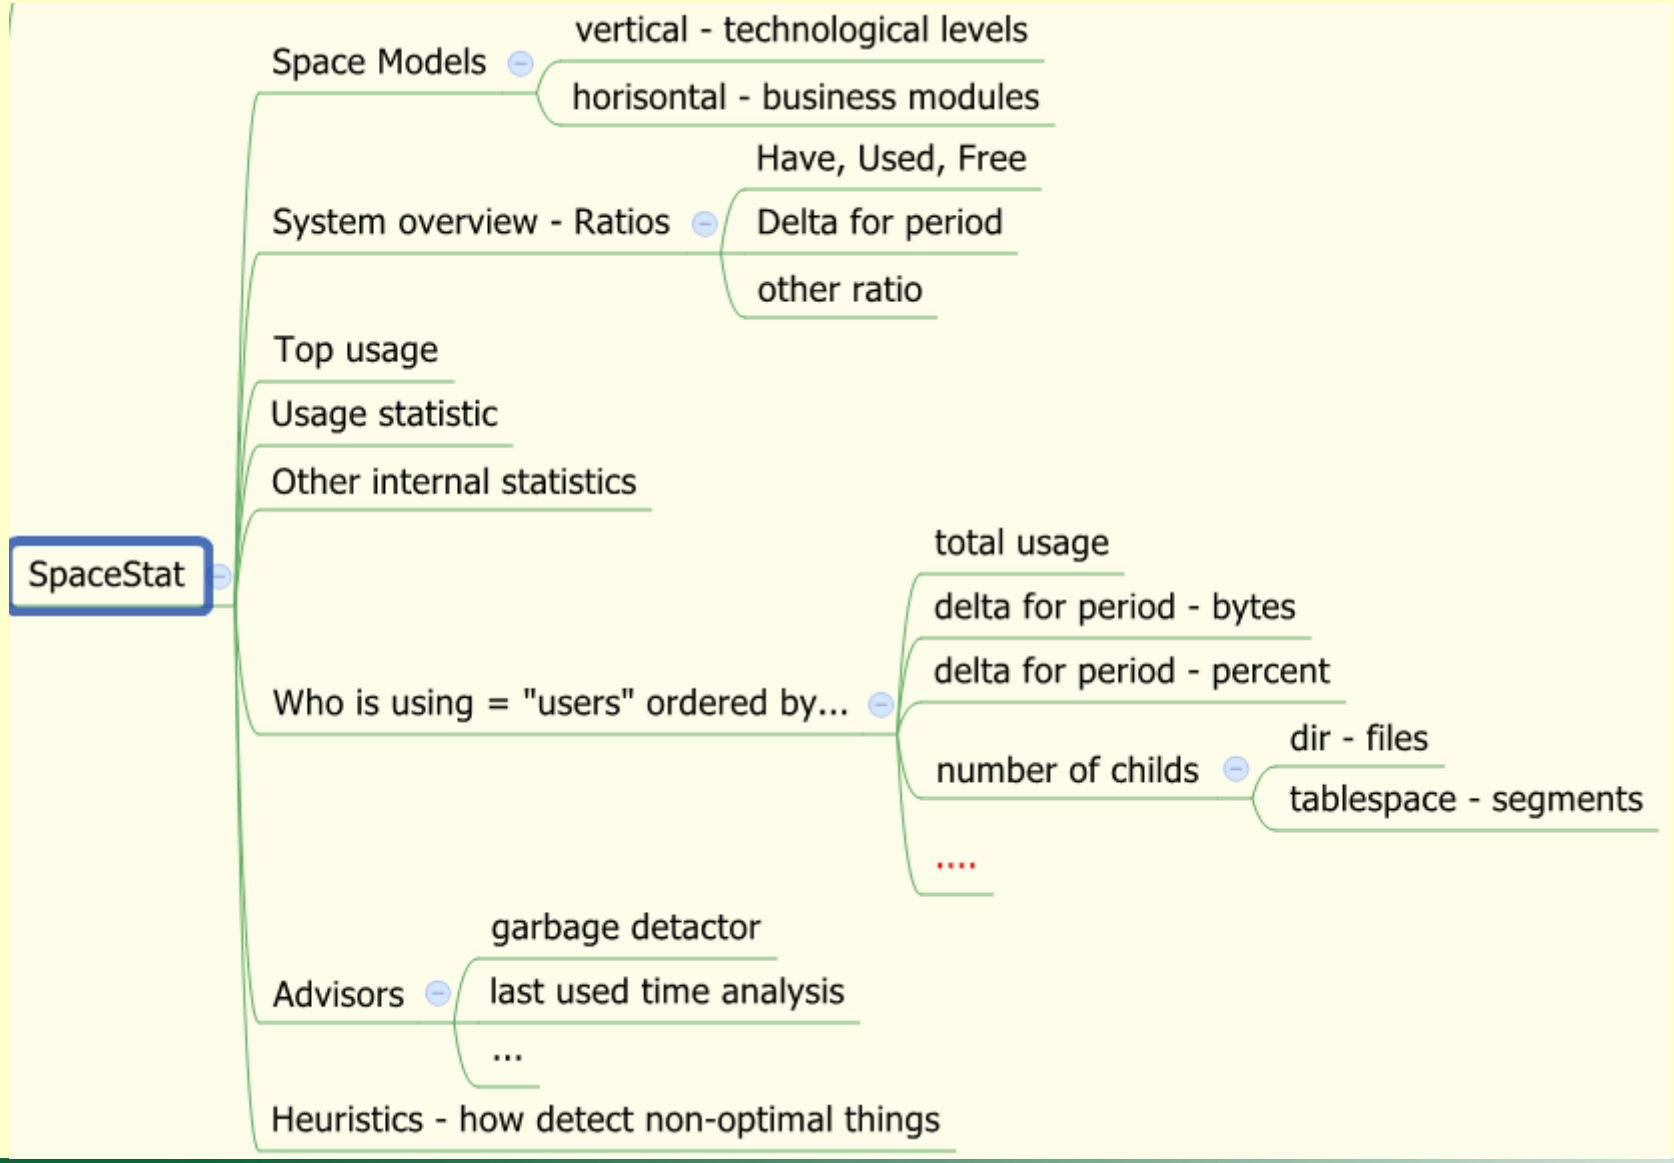

Tools

- Have You ever heard Oracle (IBM, Microsoft, ...) has SpaceStat?
- No?!
- And me too! :)

- Because Oracle (IBM, Microsoft, ...) does not have it yet!

Tools

- But, really nothing prevents us to invent it!

*And it was developed by **CoMinder SpaceStat:***

- *in 2011 for our projects use first;*
- *published 01 of June, 2012 for free use;*
- *still live development – new version planned in June, 2017*

CoMinder SpaceStat as it used to be

(and can be developed by any Oracle DBA)

SpaceStat

- CoMinder SpaceStat:
  - Application based on SQL, Shell scripts and some objects created in target DB
  - Collects DB space info (can be extended to process uploaded OS sizes info)
  - Snapshots like in StatsPack
  - Based on DB views: DBA_SEGMENTS, DBA_DATA_FILES, …..
  - Reports current state and/or changes for period
    - Usage by files, file types
    - Usage by segments
    - Usage by users
    - Other (can easy develop new reports)
  - **It is free**
    - Available for download since 1st June, 2012 at http://www.cominder.eu/

SpaceStat

- DB – Summary

```
SPACESTAT report for

DB_NAME          DBID CREATED              LOG_MODE
--------- ---------- ------------------- ------------
DEMODB       77043597 13.05.2011 21:32:19 ARCHIVELOG
------------------- ----------------------------------------
Hostname:             abc-oracle04.abc.com
Instance:             DEMODB
Started:              01.02.2012 23:06:00
Server IP:            10.10.10.10
OS:                   x86_64/Linux 2.4.xx

 Begin Snap:94  DEMODB 01-04-2012 01:01:03
   End Snap:158 DEMODB 01-05-2012 01:01:03
    Elapsed:30 days

=== Files Size ===
     OLD_GB      NEW_GB    DELTA_GB   D_PCT   AVG_DAY_GB
---------- ---------- ---------- ------- ------------
     285.08      299.97       14.89    5.22         .496
=== By Files type ===
FILE_TYPE            OLD_GB      NEW_GB    DELTA_GB   D_PCT
-------------- ---------- ---------- ---------- -------
DATA_FILE            277.08      291.97       14.89    5.37
TEMP_FILE              8.00        8.00         .00     .00
```

SpaceStat

- DB – Summary

```
=== Segments Size ===
     OLD_GB      NEW_GB    DELTA_GB   D_PCT   AVG_DAY_GB
---------- ---------- ---------- ------- ------------
     260.48      273.07       12.59    4.83         .420
=== By Segment type ===
SEGMENT_TYPE         OLD_GB      NEW_GB    DELTA_GB   D_PCT
-------------- ---------- ---------- ---------- -------
TABLE                148.80      157.07        8.27    5.56
INDEX                100.57      103.27        2.69    2.68
TYPE2 UNDO             5.27        6.57        1.30   24.69
LOBSEGMENT             5.72        6.05         .32    5.67
CLUSTER                 .10         .10         .00     .06
LOBINDEX                .03         .03        -.00    -.23
ROLLBACK                .00         .00         .00     .00
CACHE                   .00         .00         .00     .00
TABLE PARTITION         .00         .00         .00     .00
NESTED TABLE            .00         .00         .00     .00
```

SpaceStat

- DB – Segments changes

```
=== New SEGMENTS existing in 157 DEMODB 01-05-2012 01:01:01 only (or size=0 in 93 DEMODB 01-04-2012

OWNER           SEGMENT_TYPE SEGMENT_NAME                    TABLESPACE_NAME         NEW_MB
--------------- ------------ ------------------------------- -------------------- ------------
STEN            TABLE        JN_BAAK_DETAILS                 JN_TBL                  136.00
STEN            INDEX        JN_BAAK_DETAILS_ROW_ID_I        JN_IND                   52.00
RSJOBMAN        TABLE        RW_SERVER_QUEUE                 TOOLS                    26.00
STEN            INDEX        JN_BAAK_DETAILS_IDS_I           JN_IND                   24.00
STEN            TABLE        JN_BAAK_FAK_ALL_IN              JN_TBL                   22.00
```

```
=== Old SEGMENTS exist in 93 DEMODB 01-04-2012 01:01:01 only (or size=0 in 157 DEMODB 01-05-2012

OWNER           SEGMENT_TYPE SEGMENT_NAME                    TABLESPACE_NAME         OLD_MB
--------------- ------------ ------------------------------- -------------------- ------------
STEN            TABLE        TESTFACT_ERROR_2011_07_15        STEN_TBL                183.00
STEN            TABLE        TOGF_TRANS_20100608_BCK         STEN_TBL                 41.00
STEN            TABLE        YDELSER_20110831                STEN_TBL                 13.00
STEN            TABLE        YDELSER_20110704                STEN_TBL                 12.00
STEN            TABLE        TUNDE_31122010_16               STEN_TBL                  5.00
```

```
======= CHANGED SEGMENTS  =========

OWNER           SEGMENT_TYPE SEGMENT_NAME                    TABLESPACE_NAME         OLD_MB       NEW_MB      DELTA_MB
--------------- ------------ ------------------------------- -------------------- ------------ ------------ ------------
STEN            TABLE        JN_BAAK_GODS                    JN_TBL                 16633.00     18041.00      1408.00
STEN            TABLE        JN_BAAK_LIN                     JN_TBL                 12223.00     13567.00      1344.00
STEN            TABLE        JN_BAAK                         JN_TBL                 12295.00     13570.00      1275.00
STEN            TABLE        DWH_FACT_BAAK_LIN               STEN_AVG_TBL            4169.00      4864.00       695.00
STEN            INDEX        PK_TESTFACT_FLD_IOT              EDIF_BIG_TBL          11136.00     11584.00       448.00
STEN            TABLE        DWH_FACT_BAAK_LIN_BAK           STEN_AVG_TBL             496.00       896.00       400.00
PERFSTAT        TABLE        STATS$SQLTEXT                   PERFSTAT                 432.00       819.94       387.94
STEN            INDEX        JN_BAAK_LIN_ROW_ID_I            JN_IND                  2040.00      2360.00       320.00
```

DB space details

- DB – Top 100 segments

```
select TOT.* from (
select B.OWNER, B.SEGMENT_TYPE, B.SEGMENT_NAME, B.TABLESPACE_NAME,
       B.BYTES/1024/1024 MBYTES, B.BYTES/S.SBYTES*100 PCT
  from DBA_SEGMENTS B,
      (select sum(bytes) SBYTES from DBA_SEGMENTS) S
order by B.BYTES DESC, B.OWNER, B.SEGMENT_TYPE, B.SEGMENT_NAME
) TOT where rownum<101;
```

OWNER	SEGMENT_TYPE	SEGMENT_NAME	TABLESPACE	MBYTES	PCT
STEN	TABLE	JN_TAAK_GOODS	JN_TBL	17593	6.26
STEN	TABLE	TSS_SEGMENTS	STEN_TBL	14745	5.25
STEN	TABLE	TAAK_LIN	TAAK_BIG_TBL	13351	4.75
STEN	TABLE	JN_TAAK	JN_TBL	13126	4.67
STEN	TABLE	JN_TAAK_LIN	JN_TBL	13119	4.67
STEN	INDEX	PK_EDIFACT_FLD_IOT	EDIF_BIG_TBL	11456	4.08
STEN	TABLE	TAAK_GOODS	TAAK_BIG_TBL	7042	2.51
STEN	TABLE	BOGF_TRANS	BOGF_BIG_TBL	6912	2.46
STEN	TABLE	ADRESSE	ADR_BIG_TBL	5888	2.10
STEN	LOBSEGMENT	SYS_LOB0000027537C00005$$	STEN_SML_TBL	4676	1.66

DB space details

- DB – Top 20 users

```
select TOT.* from (
select B.OWNER, sum(B.BYTES/1024/1024) MBYTES, sum(B.BYTES)/S.SBYTES*100 PCT
  from DBA_SEGMENTS B,
       (select sum(bytes) SBYTES from DBA_SEGMENTS) S
group by B.OWNER, S.SBYTES
order by MBYTES DESC
) TOT where rownum<21;
```

OWNER	MBYTES	PCT
STEN	263224.19	94.11
SYS	6007.65	2.15
PERFSTAT	3569.56	1.28
TAYB	973.19	0.35
ANI	900.94	0.32
CPROTO	727.69	0.26
ABW	707.00	0.25
SYSKNL	524.31	0.19
HKISKO	357.63	0.13
GOTICH	313.19	0.11
ASRNAF	264.19	0.09
CPHNRA	121.25	0.04
SPACESTAT	86.69	0.03

DB space details

- DB – Tablespace info

TABLESPACE_NAME	FALLOCATED	USED_SIZE	FREE_SIZE	RESERVED	MAXTBS	MAXRESFILE	MAXSEG
BAAK_BIG_IND	12800.00	12741.375	58.625	3642.625	16384.000	1920.938	64.000
BAAK_BIG_TBL	31744.00	31685.000	59.000	1083.000	32768.000	399.938	64.000
TAYB_IND	128.00	44.375	83.625	467.625	512.000	467.625	1.000
TAYB_TBL	928.00	926.250	1.750	7265.750	8192.000	7265.750	8.000
DWH_AVG_IND	1345.00	1145.000	200.000	7048.000	8193.000	7048.000	4.000
DWH_BIG_TBL	5889.00	5761.000	128.000	2432.000	8193.000	2432.000	128.000
DWH_SML_IND	257.00	33.625	223.375	8159.375	8193.000	8159.375	.125
DWH_SML_TBL	97.00	50.500	46.500	8142.500	8193.000	8142.500	.125
KNL_AVG_IND	64.00	33.063	30.938	2014.938	2048.000	2014.938	1.000
KNL_AVG_TBL	128.00	120.125	7.875	1927.875	2048.000	1927.875	8.000
KNL_SML_IND	112.00	104.125	7.875	1943.875	2048.000	1943.875	1.000
KNL_SML_TBL	272.00	267.250	4.750	1780.750	2048.000	1780.750	8.000
PERFSTAT	5120.00	3569.625	1550.375	1550.375	5120.000	1550.375	8.000
SPACESTAT	97.00	86.750	10.250	1961.250	2048.000	1961.250	1.000
SYSTEM	810.00	801.625	8.375	1247.375	2049.000	1247.375	8.000
TOOLS	37.50	36.188	1.313	32731.797	32767.984	32731.797	1.000
UNDOTBS1	10240.00	5471.250	4768.750	4768.750	10240.000	4770.750	8.000
USERS	217.00	213.250	3.750	7978.750	8192.000	7978.750	1.000
USER_DAT	7424.00	5106.563	2317.438	3085.438	8192.000	3085.625	8.000
USER_DATA	48.00	.188	47.813	1023.813	1024.000	1023.813	.063
XDB	48.00	44.750	3.250	2003.250	2048.000	2003.250	.063

FALLOCATED – sum of tablespace files size

USED – by segments = FALLOCATED – FREE_SIZE

MAXTBS – sum(maxbytes) for tablespace files

RESERVED = MAXTBS – USED = logically available while files reach it's max.size

MAXRESFILE – max (MAXBYTES of file minus used space in that file) – **1 extent should reside within 1 file!**

MAXSEG – biggest segment extent in tablespace, next new extent will be at least that size!

DB garbage collector

- DB garbage collector example:

```
set lines 300
set pages 1000
-- candidates for removal:
select OWNER, SEGMENT_NAME, SEGMENT_TYPE, BYTES/1024/1024 MB
from dba_segments
where owner not in ('SYS','SYSTEM','PERFSTAT','WMSYS','OUTL')
and (segment_name like '%BCK%'
   or segment_name like '%BACKUP%'
   or segment_name like '%BACUP%'
   or segment_name like '%BAK%'
   or segment_name like '%TMP%'
   or segment_name like '%TEMP%'
   or segment_name like '%COPY%'
   or segment_name like '%200%'
   or segment_name like '%2010%'
   or segment_name like '%2011%'
   )
order by OWNER, SEGMENT_TYPE, SEGMENT_NAME;
```

CoMinder SpaceStat now

(ver. 0.56)

CoMinder SpaceStat

Subject: CoMinder SpaceStat Summary Report for databases: ALL.ABC, Time Period: 2017.04

Scope:	Files:				Segments:			
DB:	At start:	At end:	Delta:	Avg per day:	At start:	At end:	Delta:	Avg per day:
PROD1	1043.638	1068.394	24.756	0.8252	973.882	968.360	-5.522	-0.1841
PROD2	260.251	267.034	6.783	0.2261	216.317	219.175	2.857	0.0952
ITEST	821.613	824.270	2.656	0.0885	757.654	755.347	-2.307	-0.0769
RTEST	186.786	192.911	6.125	0.2042	167.310	167.656	0.346	0.0115
UDEV	784.752	788.596	3.844	0.1281	741.105	741.785	0.680	0.0227
ABCSCAN	179.841	185.278	5.438	0.1813	167.542	172.805	5.264	0.1755
IBOXTEST	47.315	47.565	0.250	0.0083	39.753	40.073	0.321	0.0107

You can always do Improvement Request for this report.

CoMinder SpaceStat

Subject: CoMinder SpaceStat Report for database: PROD1, Time Period: 2017.04

SPACESTAT report for Oracle Database
Database Info:
==========

DB_NAME	DBID	CREATED	LOG_MODE
PROD1	770435978	13-MAY-11	ARCHIVELOG

PROPERTY	VALUE
Hostname:	gdc-oracle01.abc-local.local
Instance:	PROD1
Started:	09.05.2017 21:05:51
Server IP:	10.45.101.55
OS:	x86_64/Linux 2.4.xx

=== Files ===
Begin Snap:3657 PROD1 01-04-2017 01:01:02
End Snap:3717 PROD1 01-05-2017 01:01:02

=== Segments ===
Begin Snap:3658 PROD1 01-04-2017 01:01:02
End Snap:3718 PROD1 01-05-2017 01:01:03
Elapsed:30

CoMinder SpaceStat

=== Segments Size ===

METRIC	OLD_GB	NEW_GB	DELTA_GB	D_PCT	AVG_DAY_GB
DB.DBS	973.882	968.360	-5.522	-0.57	-0.1841

=== By Segment type ===

SEGMENT_TYPE	OLD_GB	NEW_GB	DELTA_GB	D_PCT
TABLE	493.474	491.975	-1.499	-0.30
INDEX	432.212	437.311	5.099	1.18
LOBSEGMENT	30.735	31.431	0.696	2.27
TYPE2 UNDO	17.303	7.476	-9.828	-56.80
CLUSTER	0.110	0.110	0.000	0.00
LOBINDEX	0.047	0.057	0.010	20.99
ROLLBACK	0.000	0.000	0.000	0.00
NESTED TABLE	0.000	0.000	0.000	66.67
CACHE	0.000	0.000	0.000	0.00
TABLE PARTITION	0.000	0.000	0.000	0.00

CoMinder SpaceStat

=== Old SEGMENTS exist in 3658 PROD1 01-04-2017 01:01:02 only (or size=0 in 3718 PROD1 01-05-2017 01:01:03) ===

OWNER	SEGMENT_TYPE	SEGMENT_NAME	TABLESPACE_NAME	OLD_MB
SCAN	TABLE	PA_TMP_YDELSER46	USER_DAT	2.00
SCAN	TABLE	PA_TMP_KUNDE_TO_DEL	USER_DAT	0.81
SCAN	TABLE	PA_TMP_YDE_TABEL46	USER_DAT	0.25
ANALYZER	TABLE	BB_INVOBJ_EXPTIONS	SPACESTAT	0.06
DEVZVI	TABLE	SQLN_EXPLAIN_PLAN	USER_DAT	0.06
SCAN	TABLE	PA_TMP_BOOK_OVERVIEW_CFG_8906	USER_DAT	0.06

6 rows selected.

======= CHANGED SEGMENTS =========

OWNER	SEGMENT_TYPE	SEGMENT_NAME	TABLESPACE_NAME	OLD_MB	NEW_MB	DELTA_MB
SCAN	INDEX	PK_EDIFACT_FLD_IOT	EDI_BIG_IOT	72064.00	73856.00	1792.00
SCAN	TABLE	DWH_FACT_BOOK_LIN_AGNT	DWH_BIG_TBL	17536.00	18176.00	640.00
SCAN	TABLE	MSG_SEGMENTS	MSG_BIG_TBL	43648.00	44288.00	640.00
SCAN	TABLE	EML_EMAIL_OUTBOX	EDI_AVG_TBL	6524.00	7060.00	536.00

CoMinder SpaceStat

=== Top 100 SEGMENTS existing in 3718 PROD1 01-05-2017 01:01:03 ===

ROWNUM	OWNER	SEGMENT_TYPE	SEGMENT_NAME	TABLESPACE_NAME	MBYTES	PCT
1	SCAN	INDEX	PK_EDIFACT_FLD_IOT	EDI_BIG_IOT	73856.00	7.45
2	SCAN	TABLE	MSG_SEGMENTS	MSG_BIG_TBL	44288.00	4.47
3	SCAN	TABLE	BOOK_LIN	BKG_BIG_TBL	30976.00	3.12
4	SCAN	TABLE	JN_BOOK_LIN	JN_BIG_TBL	29312.00	2.96
5	SCAN	TABLE	JN_BOOK	JN_BIG_TBL	20608.00	2.08
6	SCAN	TABLE	JN_BOOK_GODS	JN_BIG_TBL	20480.00	2.07
7	SCAN	TABLE	BOGF_TRANS	BGF_BIG_TBL	20096.00	2.03
8	SCAN	TABLE	DWH_FACT_BOOK_LIN_AGNT	DWH_BIG_TBL	18176.00	1.83
9	SCAN	TABLE	JN_CODECO	JN_BIG_TBL	18048.00	1.82
10	SCAN	INDEX	PK_EDIFACT_SEQ	EDI_BIG_IND	17280.00	1.74
11	SCAN	TABLE	ABC_MRN_IMPORT_DATA	NS_BIG_TBL	16768.00	1.69
12	SCAN	TABLE	EDIFACT_SEG	EDI_BIG_TBL	16256.00	1.64
13	SCAN	LOBSEGMENT	SYS_LOB0000027537C00005$$	NS_LOB	15908.00	1.60

CoMinder SpaceStat roadmap

download for free
and more details at:

http://spacestat.lv/

What else can be done?!

- Ask Your system administrators, developers for help, hints and ideas
- Review references at this file end
- Contact me by e-mail: a.chervonets@cominder.eu :
    - we have large experience with space management, that can not fit into just one presentation :)
    - explain your “pain”
    - I will help to find solution
- Follow our web-site for tech notes and SpaceStat updates.
- Share experience and you may find event more ideas in your blog comments

Summary

Summary

- “There is no spoon!”
  - **Boy: Do not try and bend the spoon. That's impossible. Instead only try to realize the truth.**
  - **Neo: What truth?**
  - **Boy: There is no spoon.**
  - **Neo: There is no spoon?**
  - **Boy: Then you'll see that it is not the spoon that bends, it is only yourself.**
- There is ”No Silver Bullet” too:
  - Any more. And never was
    - Read for more details: "No Silver Bullet – Essence and Accident in Software Engineering" Frederick Phillips Brooks in 1986
- **“Everything You do not know - is not Your favour!”**
- To Reduce or not to reduce – You should choose it yourself!
- Because “You are – the One!” Responsible for the Results!
- **And You have enough information to make the right choice!**

Summary

- There may be benefits / drawbacks to keep DB size smaller or bigger!

  factors to balance: Time, Price (TCO), Performance, Space
- Anyway – used space management, capacity planning - is mandatory:
  - during Design
  - Integral part for maintenance
  - We should be prepared for incidents
- Any database can be made smaller – plenty of techniques
- It is possible to stay alive... and make life better
- Be prepared – keep your TOOLS ready to use!

Make the right choice….

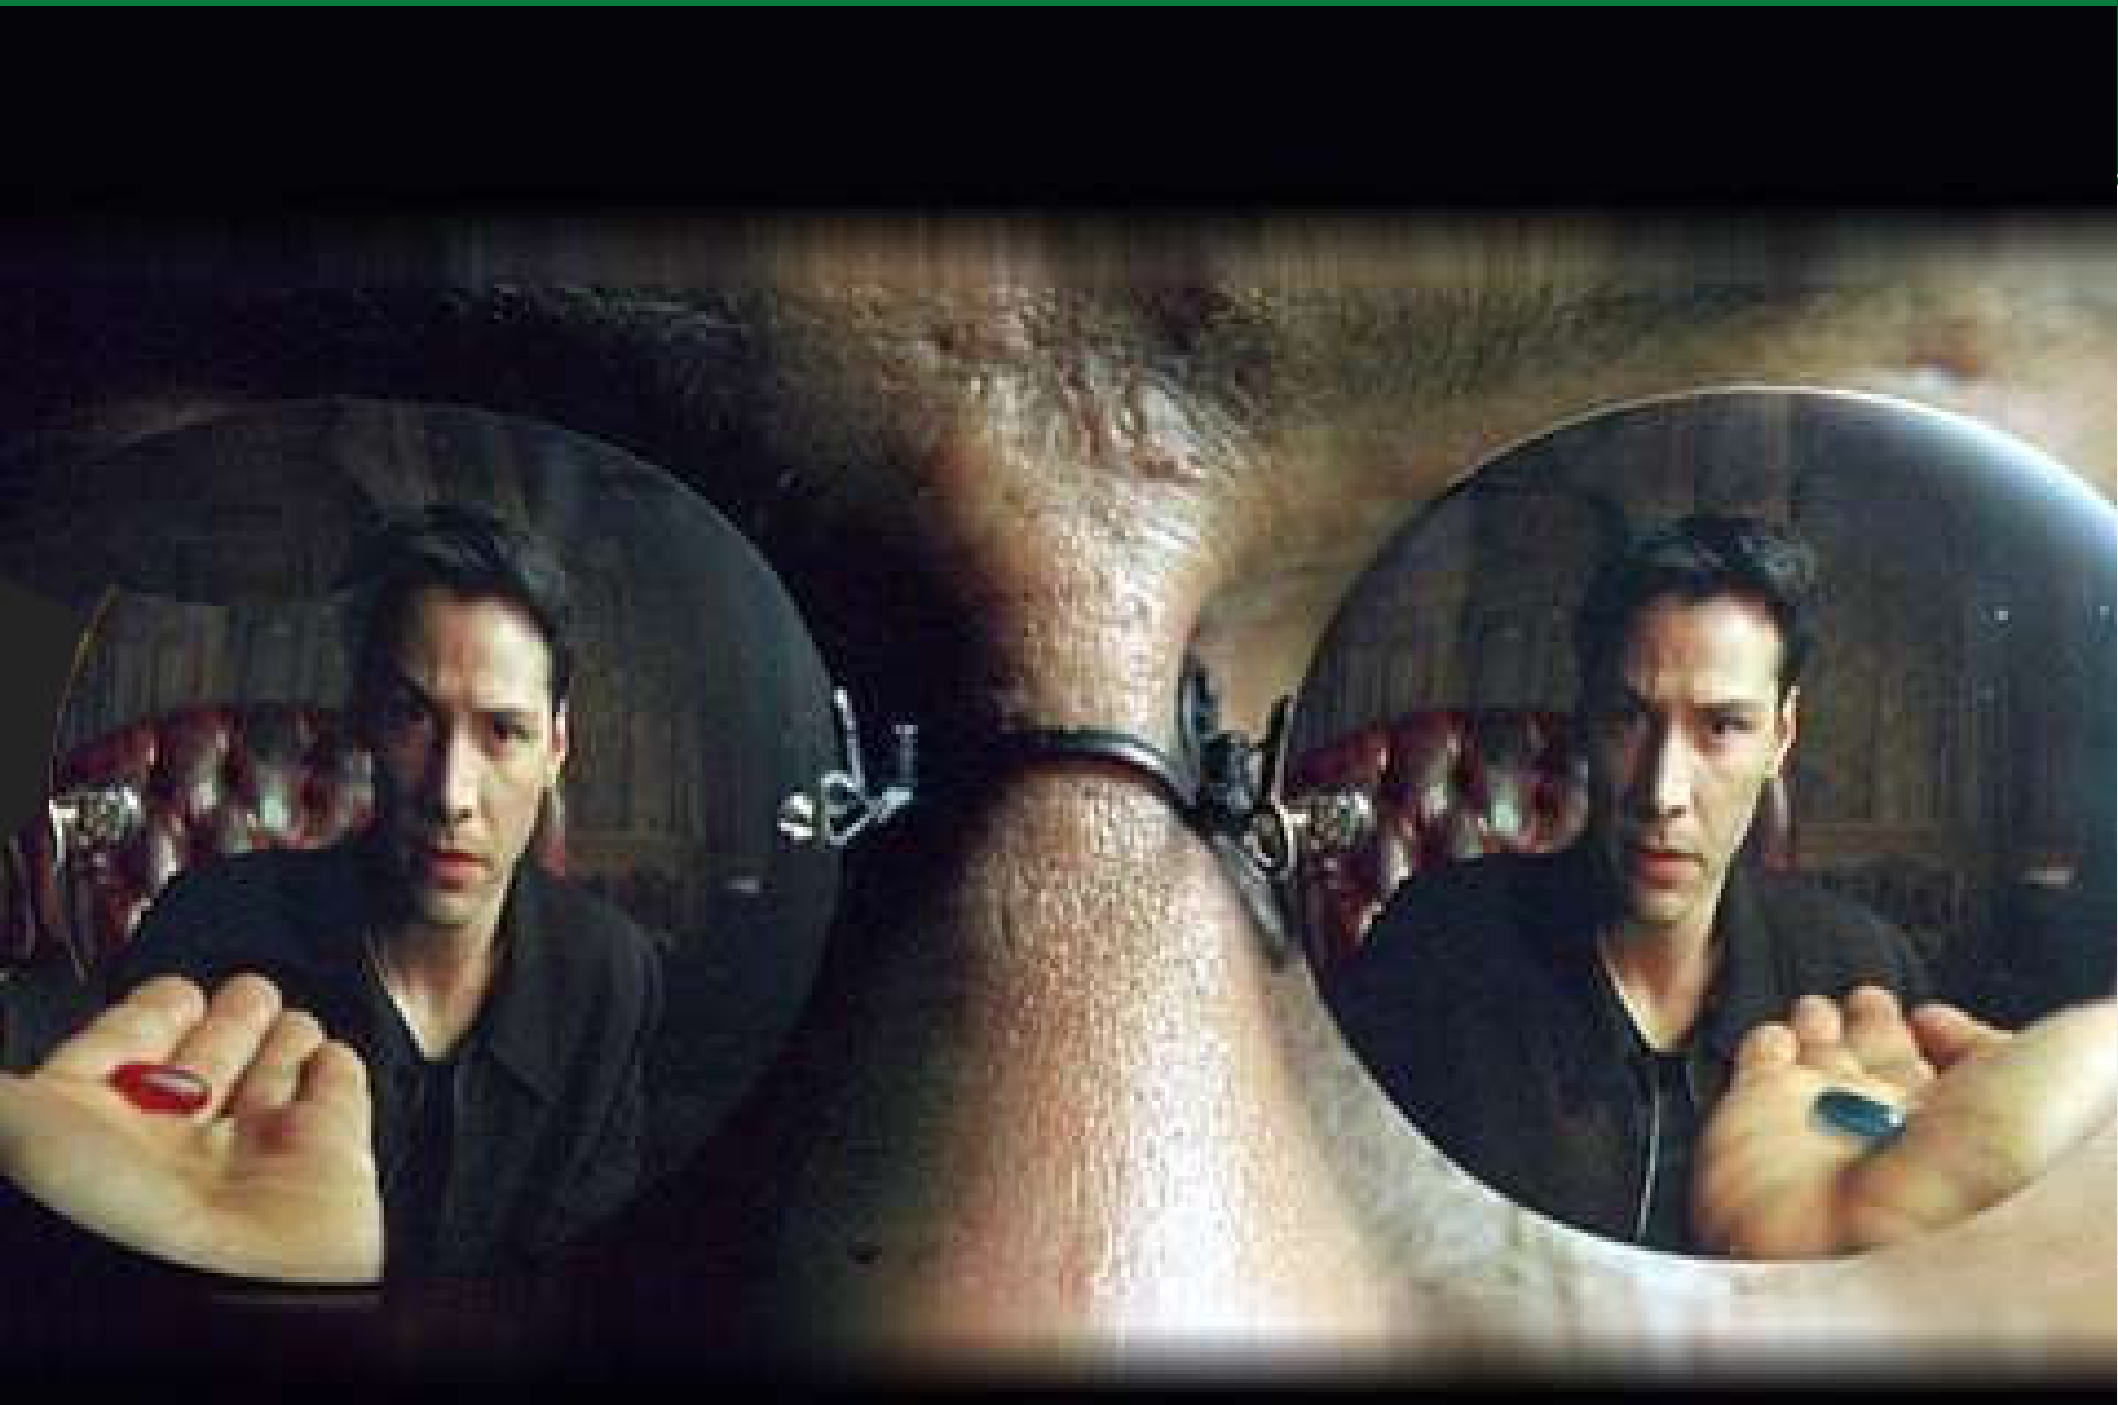

Links

- DBMS_SPACE specification: https://docs.oracle.com/database/121/ARPLS/d_space.htm#ARPLS056
- ?/rdbms/admin/catspace.sql
- Extent and Block Space Calculation and Usage in Oracle Databases [Note ID 10640.1]
- Index Rebuild, the Need vs the Implications [Note ID 989093.1]
- SAFE method: HOW TO STOP DEFRAGMENTING AND START LIVING
  - Attachment to Note ID 10640.1
  - http://www.indiana.edu/~dbateam/Documents/fragment.pdf
- SAME method: Stripe And Mirror Everything – http://docs.oracle.com/cd/B28359_01/server.111/b32024/vldb_storage.htm#BABHJEED
- AVOIDING A DATABASE REORGANIZATION - Understanding, detecting, and eliminating harmful database fragmentation by Craig A. Shallahamer http://www.allenhayden.com/cgi/getdoc.pl?file=reorg.pdf
- All About Oracle Database Fragmentation by Craig A. Shallahamer http://resources.orapub.com/product_p/dbfrag.htm

Links

- Filesystem Hierarchy Standard
  http://www.pathname.com/fhs/pub/fhs-2.3.html
- FORECASTING DATABASE DISK SPACE REQUIREMENTS:A POOR MAN'S APPROACH by Edward L. Trettel
  http://regions.cmg.org/regions/mspcmg/Presentations/Presentation03.doc
- Strategies for Solving the Datacenter Space, Power, and Cooling Crunch: Sun Server and Storage Optimization Techniques
  http://www.oracle.com/us/products/servers-storage/servers/sparc-enterprise/sun-datacenter-space-power-wp-075961.pdf
- System Administration Toolkit: Monitoring disk space and usage
  http://www.ibm.com/developerworks/aix/library/au-satdiskmon.html
- CentOS: Chapter 7. Implementing Disk Quotas
  http://www.centos.org/docs/5/html/Deployment_Guide-en-US/ch-disk-quotas.html
- Chapter 15. Optimizing Disk Space at "UNIX Power Tools"
  http://docstore.mik.ua/orelly/unix3/upt/ch15_01.htm
- Forty Ways To Free Up Disk Space - http://technet.microsoft.com/en-us/library/cc750370.aspx

Links

Software Projects:

- Skulker2 - https://www.openhub.net/p/skulker2 , https://code.google.com/archive/p/skulker2/
- CCleaner - http://www.piriform.com/CCLEANER
- TreeSize - http://www.jam-software.com/freeware/
- SpaceSniffer - http://www.uderzo.it/main_products/space_sniffer/
- NCurses Disk Usage (NCDU) - https://dev.yorhel.nl/ncdu

- CoMinder SpaceStat - http://spacestat.lv/

This presentation has been made with the following tools assistance:

- LibreOffice - https://www.libreoffice.org/ - free office suite
- XMind - http://www.xmind.net/ - one of the most popular Mind Mapping tool on the Planet

off-topic:

- Recommended Book to read:

“Time Management for System Administrators” *(By Thomas A. Limoncelli)*

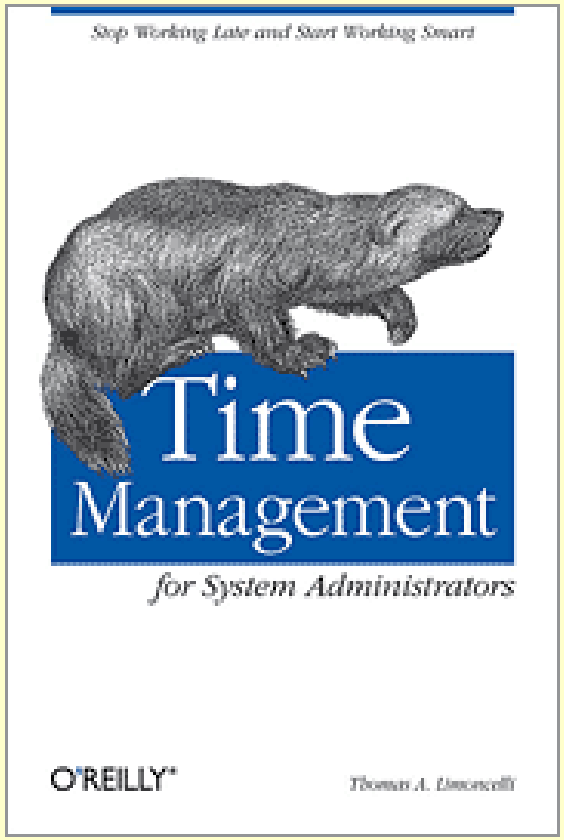

Q + A

CoMinder

Harmony 2017 Conference
17-18 of May, 2017
Helsinki, Finland

Thanks for attention!

Andrey Chervonets (Андрей Червонец)

e-mail: a.chervonets@cominder.eu
LinkedIn: http://www.linkedin.com/in/andreychervonets
web: http://www.cominder.eu