

Deconsolidate your databases

Ilmar Kerm

2021

WHOAMI: Ilmar Kerm

- Senior database administrator in Kindred Group
- Working in IT since 2000
- Working with Oracle database since 2003
- Working at Kindred Group since 2007/2015
- From LAMP developer to Oracle DBA
- Oracle ACE

Blog: <https://ilmarkerm.eu>

Email: ilmar.kerm@gmail.com

Twitter: [@ilmarkerm](https://twitter.com/ilmarkerm)



Outline

- This is a story about how we suddenly found ourselves with overconsolidated database environment.
 - Why did we reach this point?
 - What challenges did we face
 - Search for technical alternatives
 - Living with large number of deconsolidated databases

Safe Harbour statement

This presentation does not make any claims about Oracle licensing.

There are different options on how to pay Oracle licensing fees (counting CPU cores is not the only option), so it is up to each company to assess the technical solutions in the framework of their own contract with Oracle.

Recommendations in this presentation are mine. I'll try to justify them all. No absolute truths, I'll cover my thinking and you are welcome to say this is rubbish 😊

I don't mind coding things myself.

Environment overview

- One of the largest online gambling services in Europe. High transaction rates, high availability requirements.
- So far mostly serving customers from one datacenter in Europe, but globalisation pressure is growing.
- Many hundreds (and growing) of microservice applications, with ever changing requirements.
 - Hard to predict data volumes/load even for the application teams.
- Database is mostly Oracle RAC 19c.
 - Multitenant
 - Custom hardware, not Exadata/ODA

Consolidation benefits

- Simplified management
 - Monitoring
 - Backups
- Simplified disaster recovery
 - Only one database to switch/fail over to alternate site
 - Switchover/failover is logically consistent over all applications
- Less hardware wastage
 - Even with spiky workload less resources needed to just blow hot air
 - Less storage footprint
 - Shared storage is smart – silent compression/deduplication, also between volumes.

Consolidation problems

- Apps can explode in size
 - Managing databases with hundreds of TB in size starts to get problematic
 - Replicating very busy transaction rates over thousands of km starts to get problematic and less reliable
- Service glitches can affect everyone, more consolidation, more apps affected by the same glitch
 - Database software problems (RAC waits)
 - Storage/networking glitches
- With growing database instance size (both vertical and horizontal) more likely to hit Oracle bugs that they have not considered/tested
- RAC startup/shutdown brownouts get longer and more impactful with larger buffer cache

Cloud pressure

- Pushing into new global markets pressure from infrastructure teams to standardise into Cloud (on-prem) setup.
- Oracle RAC requires quite a bit of specialised hardware and networking, that do not really fit into Cloud era (both public and on-prem).
 - Interconnect networking requirements
 - Shared storage
- Even if you can get RAC running on Cloud, not happy with stability and performance.
 - Not talking about Oracle virtualisation / Cloud on customer here. IT world is bigger than Oracle.

Oracle single instance

- Compared to RAC single instance has immediate drawbacks
 - Only vertical scalability
 - What to do when need more hardware resources?
 - Unplanned maintenance is downtime
 - Hardware problems
 - OS crashes
 - Oracle instance crashes
 - Planned maintenance is downtime
 - Database patching
 - OS patching
- How to make it (almost) highly available?
 - Must keep the ability to have regular automated maintenance with minimal glitches to apps
 - (Before 21c) Even RAC planned maintenance causes brownouts – bigger SGA, bigger problem

Oracle Data Guard

- Oracle Enterprise Edition includes Data Guard
 - Warm standby database solution – **shipping redo to standby databases**
 - **Physical** so it is simple and reliable, no complicated DML extraction and replay
 - Always transactionally consistent
 - Allow changing database roles
- Additional Active Data Guard option
 - Enables opening standby for reading
 - Allows sessions to do limited writing (TEMP locally and redirect to primary)
 - Automatic block repair
- Target can also be far-sync instance
 - Database instance without datafiles, only receives and distributes redo

Data Guard Log shipping modes

Latency numbers are from OpenStack cloud with libvirt networking.
Your experience will vary.

- ARCH **+0ms**
 - Full archive log is shipped after logswitch.
 - Set ARCHIVE_LAG_TARGET to set log switch upper limit interval
- ASYNC **+0ms**
 - Redo is streamed to standby in real time. No effect on user session.
- FASTSYNC aka SYNC NOAFFIRM **+0.3ms**
 - User commit waits until standby has received the redo (in memory). **Data loss possible in case of multiple failures.**
- SYNC **+0.6ms**
 - User commit waits until one standby has written redo to standby redo log.

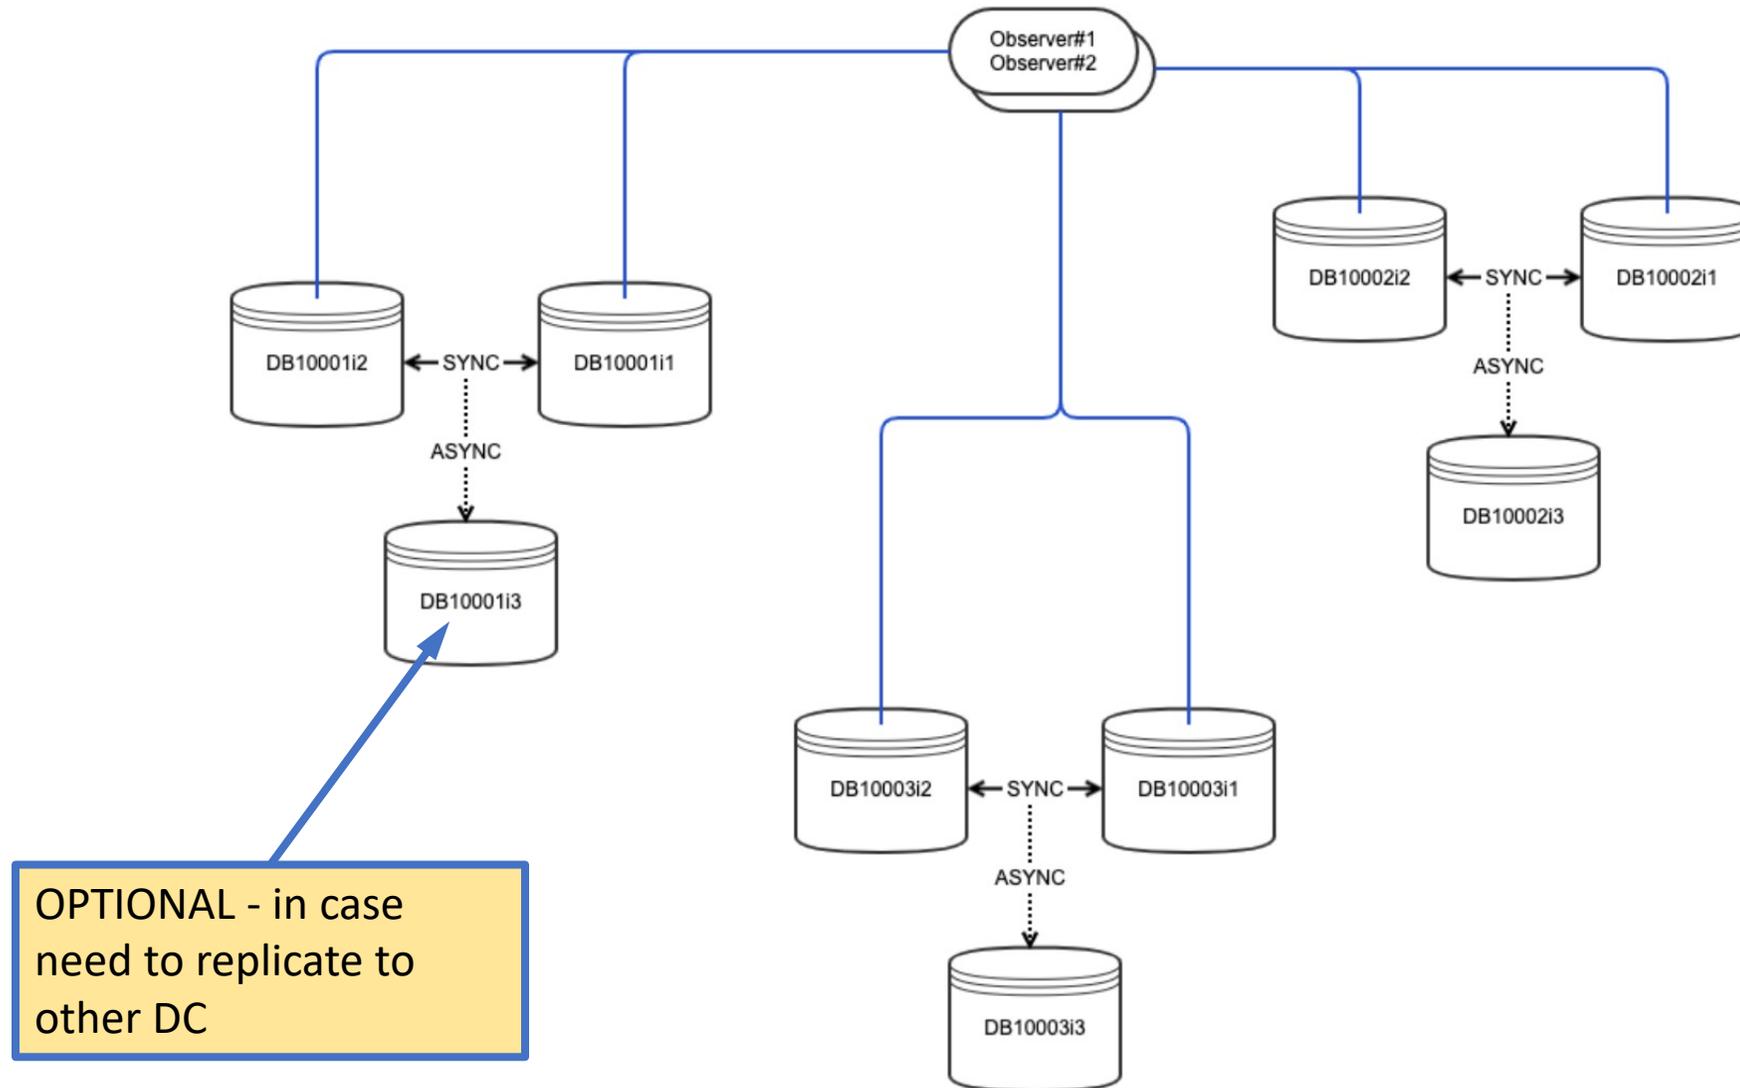
Data Guard protection modes

- MaxPerformance
 - Logs are shipped to standby asynchronously. Standby has no effect on primary.
 - ARCH/ASYNC
- **MaxAvailability**
 - User waits for standby database to acknowledge the commit.
 - **Primary keeps working if no standby is available.**
 - SYNC/FASTSYNC
- MaxProtection
 - User waits for standby database to acknowledge the commit.
 - **Shuts down primary if at least one standby does not acknowledge.**
 - At least two standbys recommended.
 - SYNC

Oracle Data Guard with Fast-Start Failover

- Database is in MaxAvailability or MaxProtection mode
- External Data Guard observers monitor databases and initiate failover
 - Multiple observers possible and strongly recommended
- Observers take action if primary is down or serious error happened
- Observers actually do a FAILOVER, not switchover
 - To bring previous primary back, it needs to be reinstated - Flashback database is required
- FSFO is performed only if databases were in sync
 - No data loss
 - Can allow data loss if FastStartFailoverLagLimit is set (don't)
- Really old stable non-sexy tech, that fits perfectly to cloud era

Overview of the setup



Clients connecting to single instance “cluster”

- Oracle drivers can know multiple addresses for a database
 - **SCAN style** – DNS name resolves to multiple IP addresses and client tries them one by one until it can connect. Recommended.
 - List them all in connection string
- Use long connection string and specify `TRANSPORT_CONNECT_TIMEOUT` to avoid waiting for TCP timeout.
- 19c JDBC driver has problems with (`LOAD_BALANCE=on`).
- How would client know if it connected to the correct listener?
 - Use services and make sure service only runs when database is PRIMARY
 - No problem, if standby is MOUNTED. With ADG some care is needed.

Example long connection string

```
(DESCRIPTION=
  (FAILOVER=on)
  (CONNECT_TIMEOUT=4) (TRANSPORT_CONNECT_TIMEOUT=2)
  (RETRY_DELAY=2) (RETRY_COUNT=10)
  (ADDRESS_LIST=
    (LOAD_BALANCE=on)
    (ADDRESS= (PROTOCOL=TCPS) (HOST=oracle.db.environment.net) (PORT=1523)) )
  (CONNECT_DATA= (SERVICE_NAME=app1.pt1.env) )
  (SECURITY=
    (SSL_SERVER_CERT_DN="cn=oracle.db.environment.net") )
)
```

oracle.db.environment.net resolves to:
10.10.10.10
10.10.10.11
10.10.10.12
10.10.10.13



Testing DG role transition and FSFO

- During switchover
 - Writes are blocked immediately with ORA-16456
 - Reads are allowed as long as possible
 - Downtime reading: 9s, downtime writing: 35s
- During observer initiated failover
 - Observer first waits for a timeout (default 30s), then initiates failover
 - Downtime 1min 3s
 - Old primary is reinstated automatically
- During user commanded failover
 - `dbms_dg.initiate_fs_failover`
 - No waiting, immediate failover, downtime 38s
 - **Observer does not reinstate old primary automatically**

How to perform planned maintenance

- Complete maintenance first on standby
- Start standby
- Wait for `V$DATABASE.FS_FAILOVER_STATUS` to become `SYNCRONIZED`
- Perform a `SWITCHOVER`
- Perform maintenance on former primary
- Make sure automation does not take down both at the same time, not even listeners

Single instance with FSFO problems

- More CPU/memory required
 - I'd recommend keeping FSFO target in MOUNTED mode, ADG has restrictions/less reliable. MOUNTED is solid.
 - If ADG instance crashes – it will not open again, needs some recovery in MOUNTED mode
 - Not all operations can be applied in ADG – like when primary as opened in startup upgrade mode
- More storage required
 - Secondary standby copy
 - Simple cloud storage cannot compress/deduplicate
- Longer COMMIT times for applications
 - Before COMMIT returns successfully it needs to commit changes on standby also
- Some overhead from mandatory flashback database logging

Living with large number of database instances

- Automation for every aspect of the lifecycle
 - Provisioning
 - Patching
 - Maintaining database settings
 - Monitoring

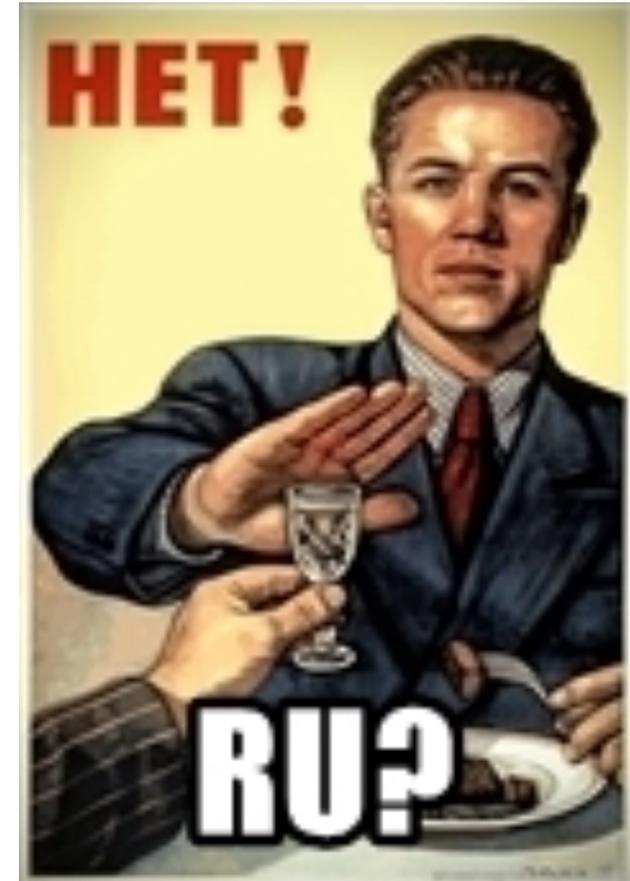
Short slide about provisioning

- AUTOMATE AUTOMATE AUTOMATE
- Terraform is great for provisioning virtual hardware
- Anti-affinity placement policy
 - To make sure standby and primary are not deployed to same hardware
 - Same goes for observers
- Ansible is great for configuring operating system
- Ansible is great for deploying Oracle software
 - Oracle has Fleet Patching and Provisioning with strange overly complex architecture, short Ansible playbook will do the same thing
- SystemD is great (and standard) for startup and shutdown
 - No need to add (and maintain) Oracle Restart if you just use it for automatic startup
- Liquibase is great for deploying common code to databases

Few more words about patching

CONTROVERSIAL

- Automation can be deadly
 - Be VERY careful that automation does not take down primary and standby at the same time. Not trivial to recover if everything shuts down during observer initiated failover.
- Linux OS components significantly easier to patch than Oracles
 - Avoid unnecessary Oracle software – like Oracle Restart
 - Talking about non-database software here 😊
- You do not have time to care and drill down for each beloved database pet
 - Use only solid and stable Oracle releases – always RUR, never RU in production.
 - RU is the release candidate branch – full of deadly undiscovered bugs.
 - RUR has the same content as RU, with **additional regression fixes** and security updates. **Doc ID 2337415.1**



Securing single instances

- **RAC and Oracle Clusterware Best Practices and Starter Kit (Linux) (Doc ID 811306.1)**
 - “Ensure that the Linux Firewall (iptables) is disabled. This is especially important for the private interconnect as this will cause instability within the cluster.”
 - Applies only for interconnect, but recommendation in the note disables iptables completely. SCARY!
- No such problem with Single Instance.
 - Just close everything except listener ports
 - FirewallD is good and easy to use.
- Not using TCPS yet? Please do. And close down that unencrypted 1521.

Maintaining database settings

- Liquibase is a great tool for delivering in common database changes
 - All your common DBA management packages, jobs, triggers, resource manager settings, common database settings...
 - Probably your application team is already using it for apps
 - By default declarative way of defining database objects (XML/JSON/YAML)
 - Liquibase support built in to SQLcl
 - But also works with plain SQL files
 - Probably not the correct way to deliver changes that require SYSDBA
 - But these changes should be very rare
- Regular job that will automatically connect to every CDB/PDB and execute Liquibase changelog

Automatic database startup on boot

- Oracle Restart
 - Out-of-box support for all Oracle database components and services
 - But an extra Oracle software to manage and patch
- SystemD
 - Standard powerful Linux core component
 - Most Linux distributions use it
 - Replaced old System V and other similar init systems
 - Can create units, units can have dependencies to other units
 - Before starting up database instance, mount this filesystem first
 - Can also monitor started processes and take action if they are down

/etc/systemd/system/oracle-listener.service

[Unit]

Description=The Oracle Database Listener

After=syslog.target network.target

[Service]

LimitNPROC=131072

LimitNOFILE=131072

LimitMEMLOCK=5525094

Type=forking

User=oracle

Group=oinstall

Restart=on-abnormal

RestartSec=5

Environment=ORACLE_HOME=/u01/app/oracle/product/19-clean-19081

Environment=LD_LIBRARY_PATH=/u01/app/oracle/product/19-clean-19081/lib

Environment=ORACLE_BASE=/u01/app/oracle

Environment=TNS_ADMIN=/u01/app/oracle/network

WorkingDirectory=/tmp

ExecStart=/bin/bash -c '/u01/app/oracle/product/19-clean-19081/bin/lsnrctl start'

ExecStop=/bin/bash -c '/u01/app/oracle/product/19-clean-19081/bin/lsnrctl stop'

[Install]

WantedBy=multi-user.target

Monitoring

- Treat your database farm truly as a cattle
- Create monitoring solutions primarily to report only on important events and give overall picture
- Keep the possibility to drill down, but not as a primary tool
- Grafana is a good, easy to use and popular tool for creating custom dashboards and alerts
 - Oracle also provides OEM repository data source for Grafana

What about Oracle Enterprise Manager?

- OEM is excellent at drilling down and to give overview of one database.
- Rather poor at giving overall view.
 - OEM has a lot of information for drilling down into details
 - OEM has poor and lack of any details overview of multiple targets

How to find, what is important to monitor

- Oracle has existing instrumentation and tools to use as data source
 - ADR
 - XML alert log / XML listener log
 - Automatic Database Diagnostic Monitor - ADDM
 - Autonomous Health Framework – AHF
 - Thousands of metrics in the database
 - Pick what is important for you, for example service level metrics from V\$SERVICEMETRIC
- Don't forget the basics
 - All databases will crash if disk space runs out

- Really important to monitor that standby is in sync
 - Fast-start failover will not occur if standby is not in sync
 - Columns in primary V\$DATABASE
 - FS_FAILOVER_MODE, FS_FAILOVER_STATUS, FS_FAILOVER_OBSERVER_PRESENT



Automatic Database Diagnostic Monitor - ADDM

- Built in database feature since 10g
 - NB! Built into database
- Analyses data stored in AWR snapshots to find performance problems and their root causes.
 - Executed automatically after every AWR snapshot.
 - Provides recommendations on how to correct the problem.
- Can report on issues that would otherwise be unnoticed
 - Broken network equipment -> ADDM reported on unusual GC lost blocks immediately.
- Sorted by impact (according to average active sessions)
 - Can be used for alerting, but be careful about the noise (yes, that SQL ID consuming a lot of database time is OK)

Autonomous Health Framework - AHF

- Collection of tools analysing (using applied ML) various database system diagnostic data sources to find issues and targeted resolutions.
 - Working independently from the database
- Helps to collect relevant data for support cases
- Has JSON output
- Newish automation tool, but suffers from Oracle americanisms
 - No time zone support
 - Why use ISO standard time format than you can invent your own or use american one??

```
[~]$ ps ax | grep pmon
2933 ? Ss 0:30 ora_pmon_pt1mt101i2
```

```
[~]$ kill -9 2933
```

```
[~]$ tfactl events -last 4h
```

```
Output from host : pt1mt101-2
```

```
-----
```

```
Event Summary:
```

```
INFO :0
```

```
ERROR :2
```

```
WARNING :1
```

```
Event Timeline:
```

```
[Oct/10/2021 07:59:31.000 UTC]: [db.pt1mt101i2.pt1mt101i2]: MMON (ospid: 3053):  
terminating the instance due to ORA error 472
```

```
[Oct/10/2021 07:59:31.000 UTC]: [db.pt1mt101i2.pt1mt101i2]: System State dumped to trace  
file /u01/app/oracle/diag/rdbms/pt1mt101i2/pt1mt101i2/trace/pt1mt101i2_diag_2992.trc
```

```
[Oct/10/2021 07:59:36.000 UTC]: [db.pt1mt101i2.pt1mt101i2]: Instance terminated by MMON,  
pid = 3053
```

```
[~]$ tfactl events -last 4h -json
```

```
[  
  { "hostname":"pt1mt101-2",  
    "Result":[  
      {  
        "timestamp":"20211010075931000",  
        "event_id":"64C2DB65B5648D8FE053D598EB0A0D86",  
        "name":"Terminating instance due to error",  
        "content":"MMON (ospid: 3053): terminating the instance due to ORA error 472",  
        "collectionid":"20211010075932pt1mt101-2", "source":"alert_pt1mt101i2.log",  
        "source_path":"../trace/alert_pt1mt101i2.log",  
        "component":"RDBMS", "database":"pt1mt101i2", "instance":"pt1mt101i2",  
        "alertlevel":"1"  
      }, {  
        "timestamp":"20211010075931000",  
        "event_id":"D302ECF5F6237B86E0431FC0E50A45A3",  
        "name":"System State dumped",  
        "content":"System State dumped to trace file ../trace/pt1mt101i2_diag_2992.trc",  
        "collectionid":"20211010075932pt1mt101-2", "source":"alert_pt1mt101i2.log",  
        "source_path":"../trace/alert_pt1mt101i2.log",  
        "component":"RDBMS", "database":"pt1mt101i2", "instance":"pt1mt101i2",  
        "alertlevel":"2"  
      }, ...  
    ]  
  } ] ]
```

XML alert log

- In addition to text alert.log, Oracle database related software is also writing logs in common XML based format, with extra information, for example severity.
 - V\$DIAG_ALERT_EXT – no GV\$ equivalent!
 - \$ADR_BASE/diag/rdbms/db/instance/alert/log.xml
- Severity in 19c RDBMS is well populated – **Alert if SEVERITY < 16**
 - 1: CRITICAL: critical errors
 - 2: SEVERE: severe errors
 - 8: IMPORTANT: important message
 - 16: NORMAL: normal message
- Problem – not all related messages are grouped together
 - For many components just seems to be wrapper around println to alert.log
- Recommend collecting entries from OS logfiles, not V\$DIAG_ALERT_EXT
 - RDBMS can crash – reason can be in alert.log
 - No GV\$ view

- Separate component VOS for instance startup/shutdown/crash events
 - What is VOS? Seems very important. Virtual Operating System?
 - VOS (ks) – from oradebug doc component.

DBS / alert.log - OTHER ☆ 🔗

Log alert_VOS Cluster name All Host All Instance All Message text regex .+

Log entries ▾

Time ▾	Log	Cluster	Host	Instance	Message	Severity
22:13:20	VOS	ocdev1se1	odev1cdb	odev1cdb1	Starting ORACLE instance (normal) (OS id: 26512)	NORMAL
22:13:11	VOS	ocdev1se1	odev1cdb	odev1cdb1	LMON (ospid: 11561): terminating the instance due to ORA error 472	SEVERE
18:13:23	VOS	oc2swe4	dwqa2	dwqa22	Starting ORACLE instance (normal) (OS id: 2859)	NORMAL
18:13:23	VOS	oc2swe4	dwqa2	dwqa21	Starting ORACLE instance (normal) (OS id: 8110)	NORMAL
18:13:19	VOS	oc2swe4	dwqa2	dwqa21	Shutting down ORACLE instance (abort) (OS id: 17017)	NORMAL
18:13:19	VOS	oc2swe4	dwqa2	dwqa22	Shutting down ORACLE instance (abort) (OS id: 29943)	NORMAL
16:33:05	VOS	- none -	ci10001i2	ci10001i2	PGA_AGGREGATE_LIMIT specified is high	IMPORTANT
16:32:58	VOS	- none -	ci10001i2	ci10001i2	PGA_AGGREGATE_TARGET specified is high	IMPORTANT
16:32:56	VOS	- none -	ci10001i2	ci10001i2	Starting ORACLE instance (normal) (OS id: 16779)	NORMAL
16:32:29	VOS	- none -	ci10001i2	ci10001i2	Shutting down ORACLE instance (immediate) (OS id: 16421)	NORMAL
11:40:53	VOS	- none -	spt1cdb	spt1cdb1	Starting ORACLE instance (normal) (OS id: 18956)	NORMAL

19c RDBMS XML alert.log sample messages

```
<msg time='2021-10-08T16:53:34.772+02:00' org_id='oracle'  
comp_id='rdbms' type='UNKNOWN' level='16' host_id='oc4b01n01'  
host_addr='10.10.10.10' pid='92697'>
```

```
  <txt>Thread 1 advanced to log sequence 132126 (LGWR switch),  
current SCN: 7073581543261
```

```
  </txt>
```

```
</msg>
```

19c RDBMS XML alert.log sample messages

```
<msg time='2021-05-25T10:06:54.591+02:00' org_id='oracle'  
comp_id='rdbms' msg_id='ksbsrvn_opt:7871:323864049' type='ERROR'  
group='background_proc' level='8' host_id='oc4b01n01'  
host_addr='10.10.10.10' module='MMON_SLAVE' pid='73401'>
```

```
  <txt>Process P0BZ submission failed with error = 20
```

```
  </txt>
```

```
</msg>
```

19c RDBMS XML alert.log sample messages

```
<msg time='2021-10-08T16:36:54.510+02:00' org_id='oracle'  
comp_id='rdbms' msg_id='975027453' type='INCIDENT_ERROR'  
group='Access Violation' level='1' host_id='oc4b01n01'  
host_addr='10.10.10.10' pid='54966' prob_key='ORA 7445 [kghalf]'  
downstream_comp='KGH' errid='187724'  
detail_path='...diag/.../b2prdwmtoc4bel11_ora_54966.trc'>  
  <txt>Errors in file ../diag/.../b2prdwmtoc4bel11_ora_54966.trc  
(incident=187724):  
  
ORA-07445: exception encountered: core dump [kghalf()+724]  
[SIGSEGV] [ADDR:0x401800] [PC:0x12D8B044] [Invalid permissions for  
mapped object] []  
  </txt>  
</msg>
```

19c RDBMS XML alert.log sample messages

```
<msg time='2021-07-12T17:33:19.853+02:00' org_id='oracle' comp_id='VOS'  
msg_id='ksuitem opt:1046:874893762' type='ERROR' group='VOS' level='2'  
host_id='oc4b0In01' host_addr='10.10.10.10' pid='13161' attn_msg_id='35782660'  
id_suffix='DB Termination' ora_id='2' db_id='b2prdwmtoc4bell1'  
sess_id='33.27701'
```

```
  callstack='dbklaWriteAttentionLogTextVaList()+556&lt;-  
dbklaWriteAttentionLogText()+173&lt;-ksuitem opt()+2030&lt;-  
ksucln check_fatal_death()+790&lt;-ksuxfd check_dead()+276&lt;-  
ksuxfd()+710&lt;-ksuxfl()+558&lt;-ksucln coord_main()+778&lt;-  
ksbrdp()+1123&lt;-opirip()+541&lt;-opidrv()+58I&lt;-sou2o()+165&lt;-  
opimai_real()+173&lt;-ssthrdmain()+417&lt;-main()+256&lt;-  
__libc_start_main()+245' flags='15'>
```

```
<txt>PMON (ospid: 13161): terminating the instance due to ORA error  
12752</txt>
```

```
<cause>The instance termination routine was called</cause>
```

```
<action>Check alert log for more information relating to instance termination  
rectify the error and restart the instance</action>
```

```
</msg>
```

Cluster name
 Database
 Instance
 Message severity
 Message text regex

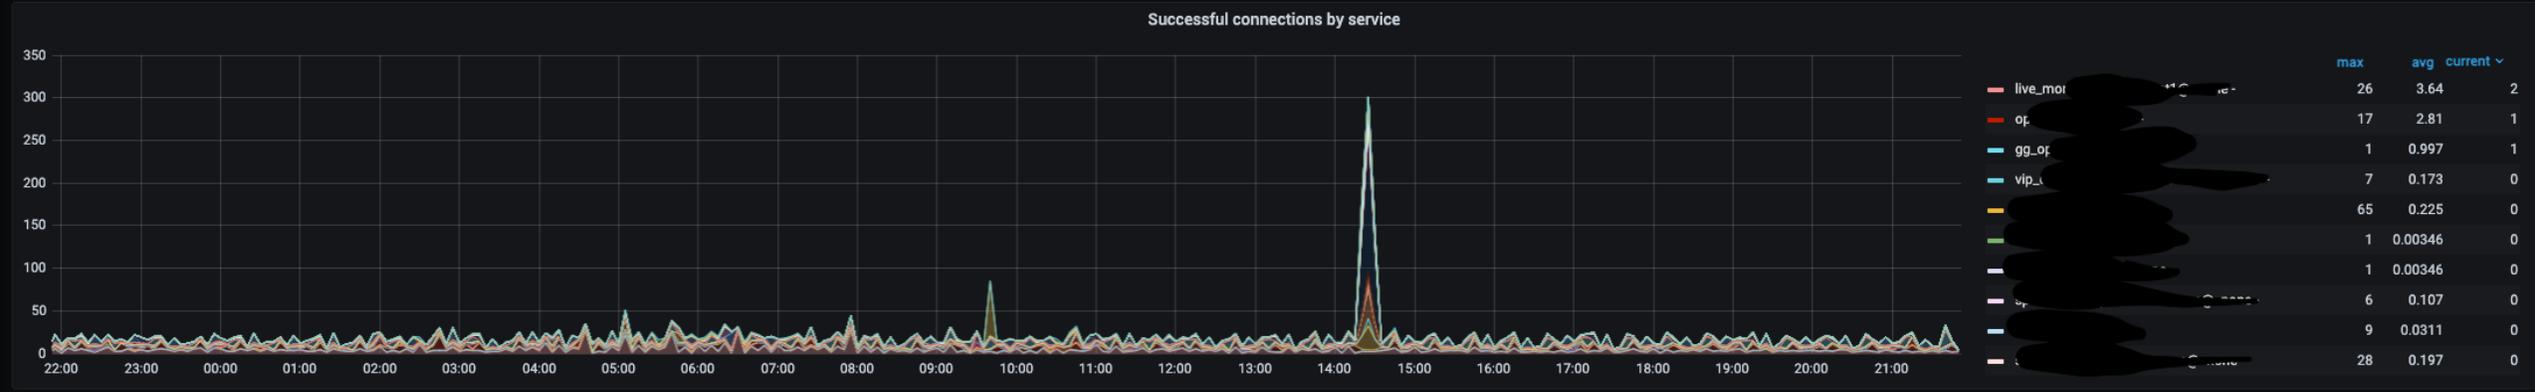
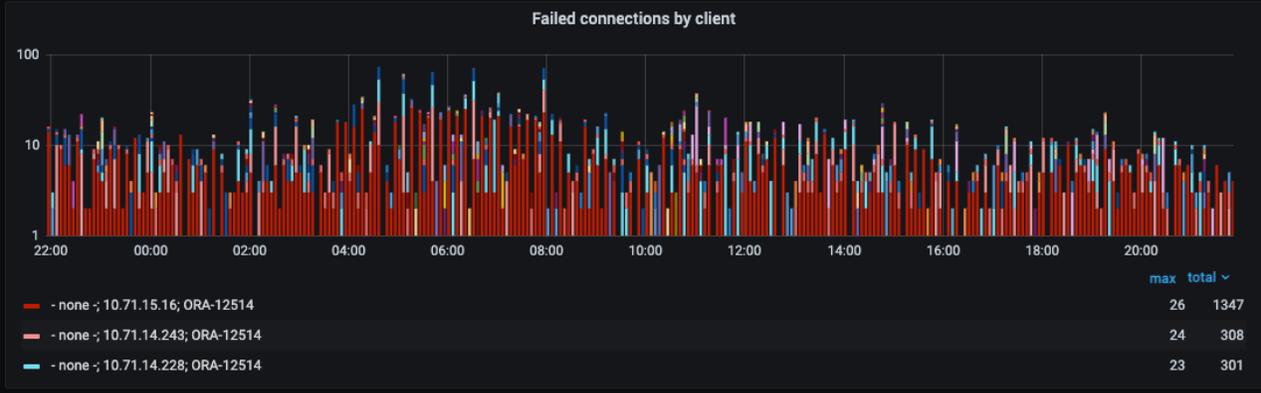
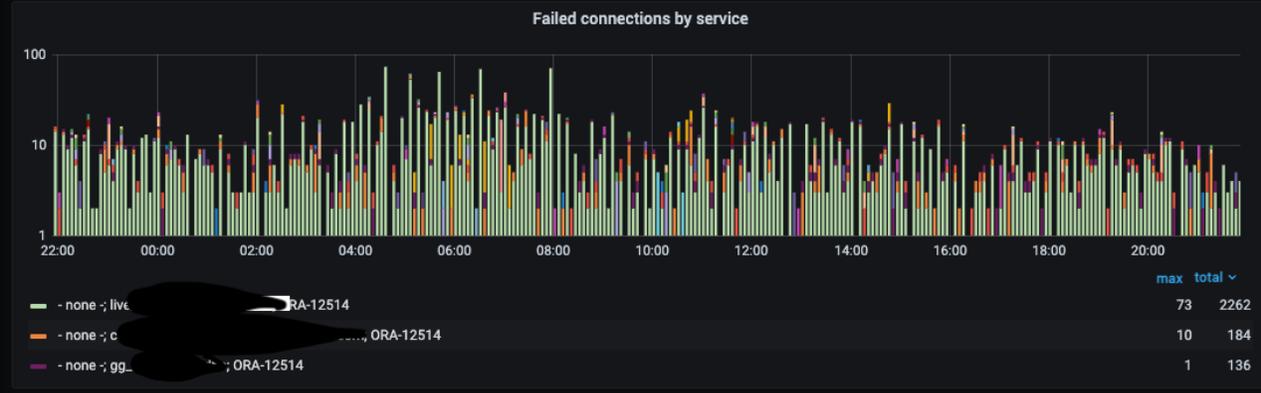
Log entries

Time	Database	Instance	Problem	Container	Message	Severity
16:36:54	b2prdwmtoc4bel1	b2prdwmtoc4bel11	ORA 7445 [kghalf]	-	Errors in file /u01/app/oracle/diag/rdbms/b2prdwmtoc4bel1/b2prdwmtoc4bel11/trace/b2prdwmtoc4bel11_ora_54966.trc (incident=187724): ORA-07445: exception encountered: core dump [kghalf()+724] [SIGSEGV] [ADDR:0x401800] [PC:0x12D8B044] [Invalid permissions for mapped object] []	CRITICAL 1
00:04:22	odwh2cdb	odwh2cdb4	ORA 603	CDB\$ROOT	Errors in file /u01/app/oracle/diag/rdbms/odwh2cdb/odwh2cdb4/trace/odwh2cdb4_mmmL10675.trc (incident=120458) (PDBNAME=CDB\$ROOT): ORA-00603: ORACLE server session terminated by fatal error ORA-27504: IPC error creating OSD context ORA-27300: OS system dependent operation:select (rcv) timed out failed with status: 0 ORA-27301: OS failure message: Error 0 ORA-27302: failure occurred at: skgxpvfynet ORA-27303: additional information: rval 5	CRITICAL 1
00:03:22	odwh2cdb	odwh2cdb4	ORA 603	CDB\$ROOT	Errors in file /u01/app/oracle/diag/rdbms/odwh2cdb/odwh2cdb4/trace/odwh2cdb4_mmmL10675.trc (incident=120457) (PDBNAME=CDB\$ROOT): ORA-00603: ORACLE server session terminated by fatal error ORA-27504: IPC error creating OSD context ORA-27300: OS system dependent operation:select (rcv) timed out failed with status: 0 ORA-27301: OS failure message: Error 0 ORA-27302: failure occurred at: skgxpvfynet ORA-27303: additional information: rval 5	CRITICAL 1
00:03:18	b2prdwmtoc4bel1	b2prdwmtoc4bel14	ORA 603	-	Errors in file /u01/app/oracle/diag/rdbms/b2prdwmtoc4bel1/b2prdwmtoc4bel14/trace/b2prdwmtoc4bel14_ora_74941.trc (incident=219990): ORA-00603: ORACLE server session terminated by fatal error ORA-27504: IPC error creating OSD context ORA-27300: OS system dependent operation:select (rcv) timed out failed with status: 0 ORA-27301: OS failure message: Error 0 ORA-27302: failure occurred at: skgxpvfynet ORA-27303: additional information: rval 5 ORA-01017: invalid username/password; logon denied	CRITICAL 1
00:03:13	b2prdwmtoc4bel1	b2prdwmtoc4bel14	ORA 603	-	Errors in file /u01/app/oracle/diag/rdbms/b2prdwmtoc4bel1/b2prdwmtoc4bel14/trace/b2prdwmtoc4bel14_ora_74941.trc (incident=219989): ORA-00603: ORACLE server session terminated by fatal error ORA-27504: IPC error creating OSD context ORA-27300: OS system dependent operation:select (rcv) timed out failed with status: 0 ORA-27301: OS failure message: Error 0 ORA-27302: failure occurred at: skgxpvfynet ORA-27303: additional information: rval 5	CRITICAL 1
00:02:51	b2prdwmtoc4bel1	b2prdwmtoc4bel14	ORA 603	-	Errors in file /u01/app/oracle/diag/rdbms/b2prdwmtoc4bel1/b2prdwmtoc4bel14/trace/b2prdwmtoc4bel14_ora_74300.trc (incident=219982): ORA-00603: ORACLE server session terminated by fatal error ORA-27504: IPC error creating OSD context ORA-27300: OS system dependent operation:select (rcv) timed out failed with status: 0 ORA-27301: OS failure message: Error 0 ORA-27302: failure occurred at: skgxpvfynet ORA-27303: additional information: rval 5 ORA-01017: invalid username/password; logon denied	CRITICAL 1
00:02:46	b2prdwmtoc4bel1	b2prdwmtoc4bel14	ORA 603	-	Errors in file /u01/app/oracle/diag/rdbms/b2prdwmtoc4bel1/b2prdwmtoc4bel14/trace/b2prdwmtoc4bel14_ora_74300.trc (incident=219981): ORA-00603: ORACLE server session terminated by fatal error ORA-27504: IPC error creating OSD context ORA-27300: OS system dependent operation:select (rcv) timed out failed with status: 0 ORA-27301: OS failure message: Error 0 ORA-27302: failure occurred at: skgxpvfynet ORA-27303: additional information: rval 5	CRITICAL 1
00:02:22	odwh2cdb	odwh2cdb4	ORA 603	CDB\$ROOT	Errors in file /u01/app/oracle/diag/rdbms/odwh2cdb/odwh2cdb4/trace/odwh2cdb4_mmmL10675.trc (incident=120456) (PDBNAME=CDB\$ROOT): ORA-00603: ORACLE server session terminated by fatal error ORA-27504: IPC error creating OSD context ORA-27300: OS system dependent operation:select (rcv) timed out failed with status: 0 ORA-27301: OS failure message: Error 0 ORA-27302: failure occurred at: skgxpvfynet ORA-27303: additional information: rval 5	CRITICAL 1
22:13:45	b2prdwmtoc4bel1	b2prdwmtoc4bel14	ORA 600 [ksfdids4]	-	Errors in file /u01/app/oracle/diag/rdbms/b2prdwmtoc4bel1/b2prdwmtoc4bel14/trace/b2prdwmtoc4bel14_ora_10018.trc (incident=219630): ORA-00600: internal error code, arguments: [ksfdids4], [Invalid argument], [], [], [], [], [], [], [], [], ORA-19624: operation failed, retry possible ORA-19624: operation failed, retry possible ORA-06512: at "SYS.DBMS_BACKUP_RESTORE", line 1714 ORA-00600: internal error code, arguments: [ksfdids4], [No such file or directory], [], [], [], [], [], [], [], [], ORA-19502: write error on file "/nfs/backup/b2prdwmtoc4bel1/b80b0pma_561512_1_1", block number 28169280 (block size=16384) ORA-17500: ODM err:KGNFS WRITE FAIL:Stale File handle ORA-06512: at "SYS.DBMS_BACKUP_RESTORE", line 1707	CRITICAL 1
22:13:44	b2prdwmtoc4bel1	b2prdwmtoc4bel14	ORA 600 [ksfdids4]	-	Errors in file /u01/app/oracle/diag/rdbms/b2prdwmtoc4bel1/b2prdwmtoc4bel14/trace/b2prdwmtoc4bel14_ora_10018.trc (incident=219629): ORA-00600: internal error code, arguments: [ksfdids4], [No such file or directory], [], [], [], [], [], [], [], [], ORA-19502: write error on file "/nfs/backup/b2prdwmtoc4bel1/b80b0pma_561512_1_1", block number 28169280 (block size=16384) ORA-17500: ODM err:KGNFS WRITE FAIL:Stale File handle	CRITICAL 1

Listener log

- Listener log is a very good indication if clients have trouble connecting to the database:
 - Are the services up?
 - What clients connect the most?
 - What clients fail to connect?
 - What protocol is used to connect? Who has not migrated to TCPS yet?

Cluster: none Service regex: .+ FQDN regex: .+ Client IP regex: .* Listener type: listener Group by: cluster_name



Log entries

Time	Instance	Problem	Message	Severity	Count
16:36:54	b2prdwmtoc4bel11	ORA-7445 [kghalf]	Errors in file /u01/app/oracle/diag/rdbms/b2prdwmtoc4bel11/b2prdwmtoc4bel11/trace/b2prdwmtoc4bel11_ora_54966.trc (incident=187724): ORA-07445: exception encountered: core dump [kghalf()+724] [SIGSEGV] [ADDR:0x401800] [PC:0x12D8B044] [Invalid permissions for mapped object] []	CRITICAL	1

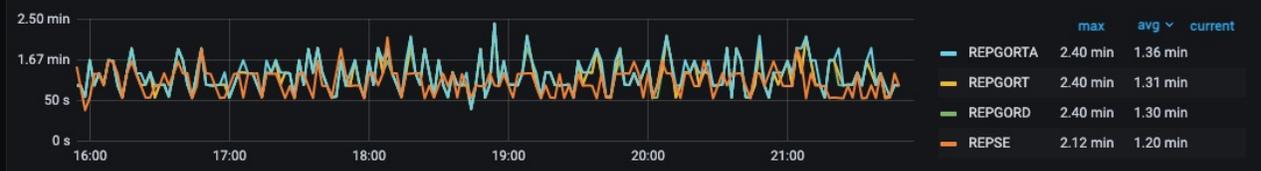
Failed connections by service



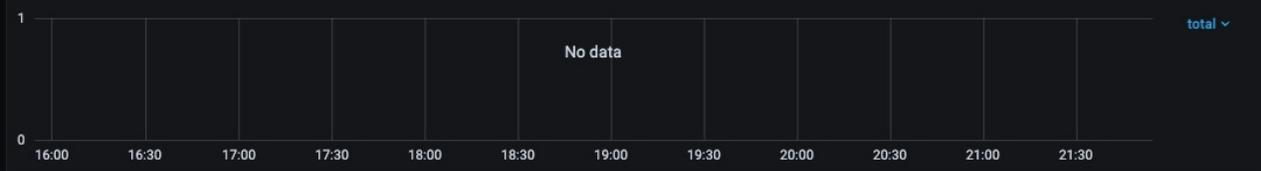
ORA-20 maximum number of processes exceeded



Goldengate end-to-end lag



Dictionary / binary version mismatch



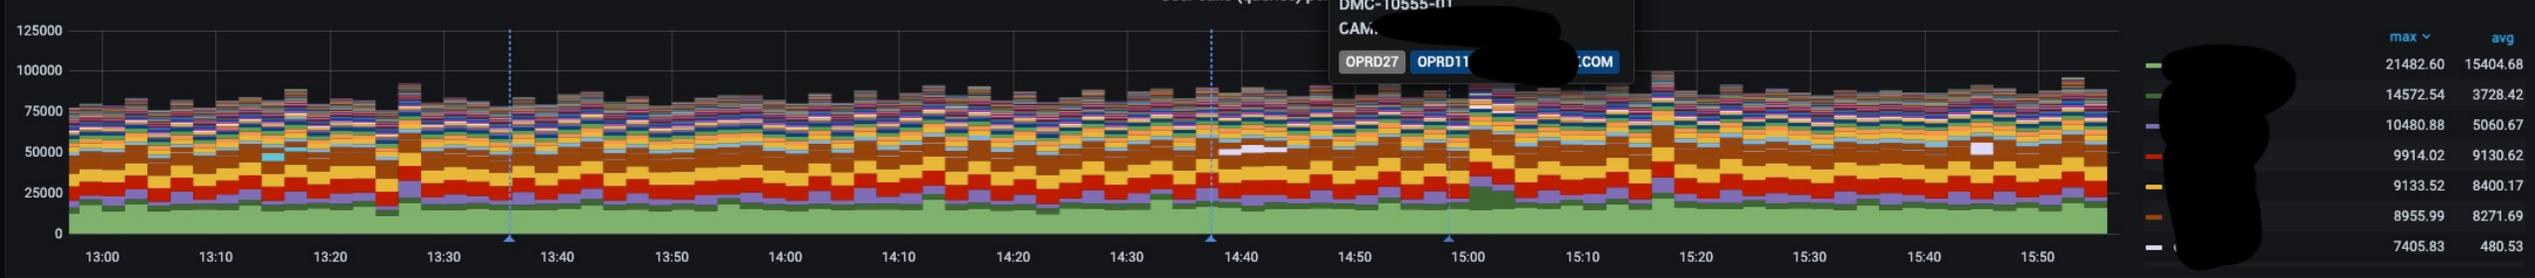
Environment PROD1 Market MT1 Service regex .+

Instance shutdown and startup Database errors Patching events Liquibase deployments

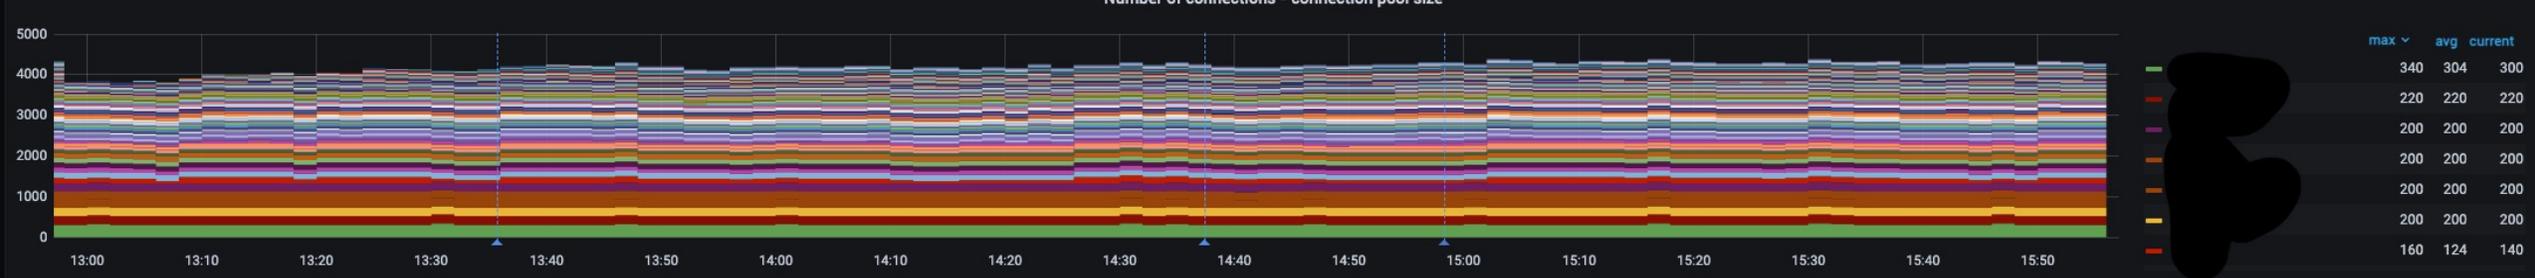
Number of database sessions that are doing some work in the database (not idle)



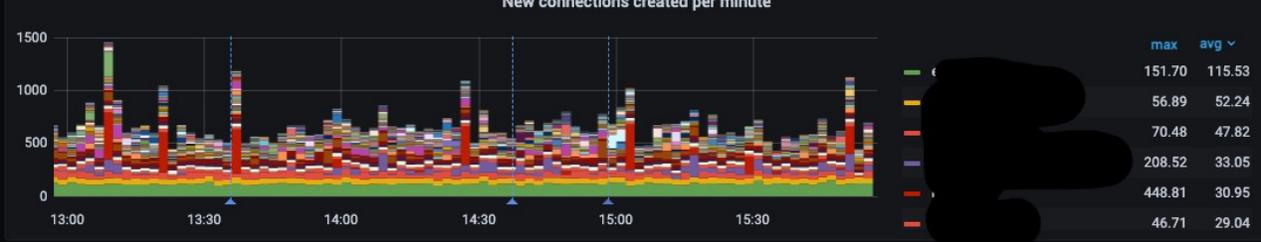
User calls (queries) per second



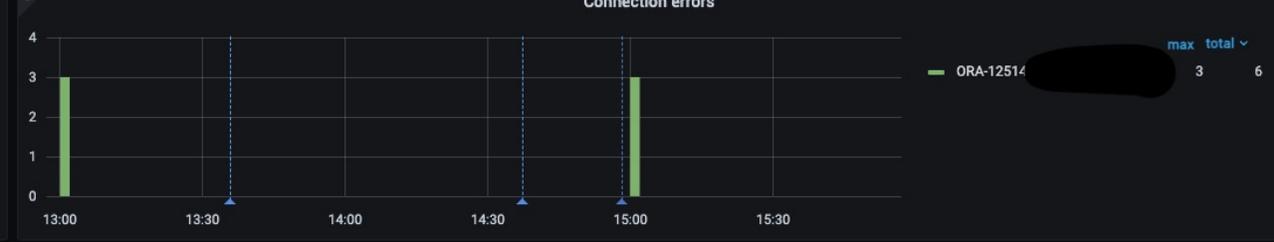
Number of connections - connection pool size



New connections created per minute



Connection errors



Redo generated - how much data changed



Waits due to application behaviour - row locks etc



Moving forward

- Oracle Connection Manager
 - As a first point of contact for clients
 - Would provide TLS encryption and (extra) connection pooling

Tools

- <https://ilmarkerm.eu/blog/2021/01/mining-log-files-under-adr-base/>
- <https://ilmarkerm.eu/blog/2021/01/monitoring-oracle-listener-from-grafana/>
- <https://ilmarkerm.eu/blog/2021/01/system-monitoring-agent/>
- <https://ilmarkerm.eu/blog/2021/01/managing-oracle-database-listener-with-systemd/>

Thank you!

@ilmarkerm

<https://ilmarkerm.eu/>