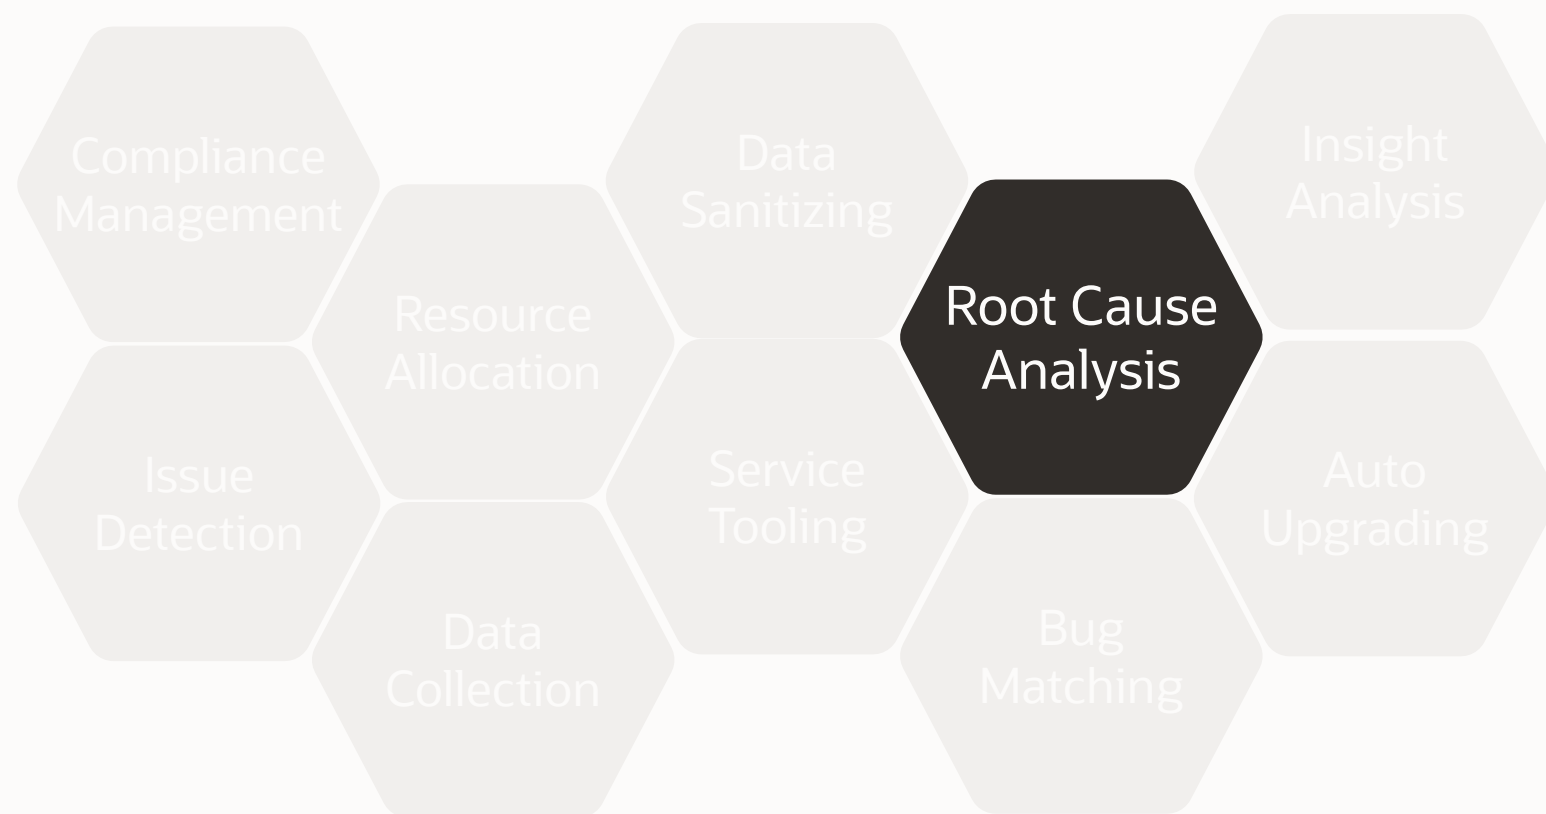




ORACLE

Autonomous Health Framework


Root Cause Analysis



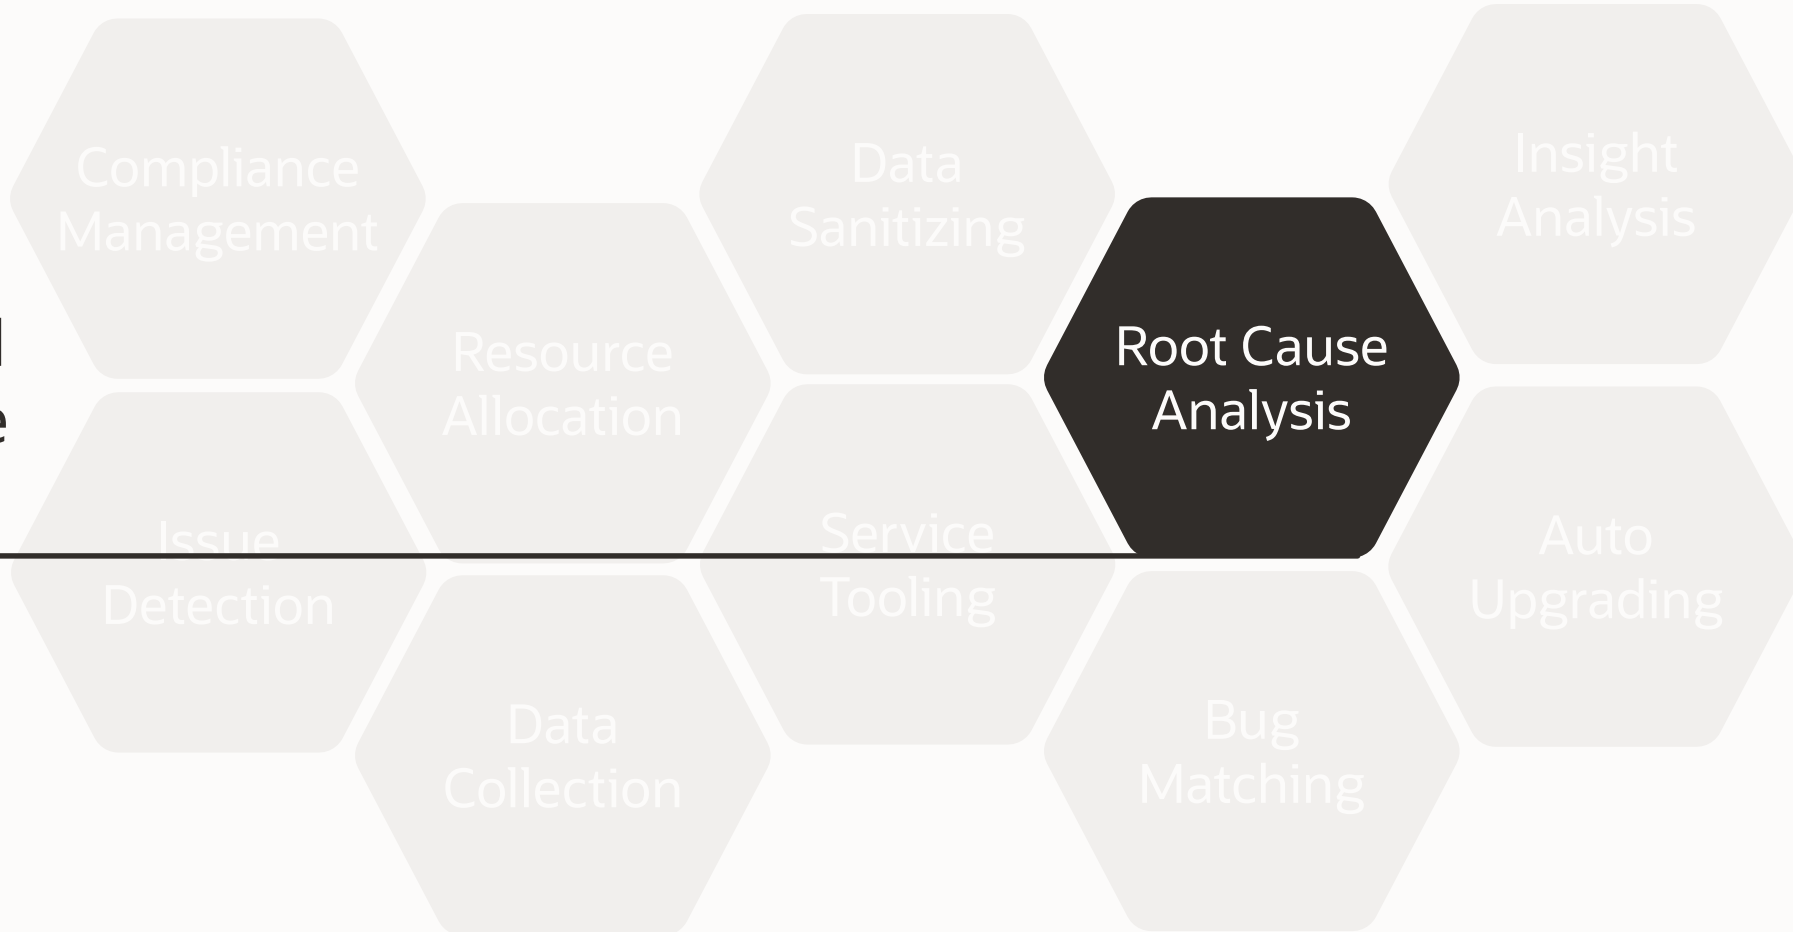
Oracle Autonomous Health Framework



Oracle Autonomous Health Framework



AI Powered
Root Cause
Analysis



Find if anything has changed



Has anything changed recently?

```
tfact1 changes
Output from host : myserver69
-----
[Apr/28/2023 04:54:15.397]: Parameter: fs.aio-nr: Value: 95488 => 97024
[Apr/28/2023 04:54:15.397]: Parameter: fs.inode-nr: Value: 764974 131561 => 740744 131259
[Apr/28/2023 04:54:15.397]: Parameter: kernel.pty.nr: Value: 2 => 1
[Apr/28/2023 04:54:15.397]: Parameter: kernel.random.entropy_avail: Value: 189 => 158
[Apr/28/2023 04:54:15.397]: Parameter: kernel.random.uuid: Value: 36269877-9bc9-40a3-82e0-
1619865096f2 => 7551c5e7-c59f-40fa-b55f-5bd170e8b1ab
[Apr/28/2023 05:46:15.397]: Parameter: fs.aio-nr: Value: 119680 => 122880
[Apr/28/2023 05:46:15.397]: Parameter: fs.inode-nr: Value: 1580316      810036 => 1562320
768555
[Apr/28/2023 05:46:15.397]: Parameter: kernel.pty.nr: Value: 19 => 18
[Apr/28/2023 05:46:15.397]: Parameter: kernel.random.uuid: Value: 37cc31aa-ee31-459e-8f2a-
0766b34b1b64 => f5176cdc-6390-415d-882e-02c4cff2ae4e
...
```

Has anything changed recently?

```
...
Output from host : myserver70
-----
[Apr/28/2023 04:54:15.397]: Parameter: fs.aio-nr: Value: 95488 => 97024
[Apr/28/2023 04:54:15.397]: Parameter: fs.inode-nr: Value: 764974 131561 => 740744 131259
[Apr/28/2023 04:54:15.397]: Parameter: kernel.ptype.nr: Value: 2 => 1
[Apr/28/2023 04:54:15.397]: Parameter: kernel.random.entropy_avail: Value: 189 => 158
[Apr/28/2023 04:54:15.397]: Parameter: kernel.random.uuid: Value: 36269877-9bc9-40a3-82e0-
1619865096f2 => 7551c5e7-c59f-40fa-b55f-5bd170e8b1ab
[Apr/28/2023 05:46:15.397]: Parameter: fs.aio-nr: Value: 119680 => 122880
[Apr/28/2023 05:46:15.397]: Parameter: fs.inode-nr: Value: 1580316      810036 => 1562320
768555
[Apr/28/2023 05:46:15.397]: Parameter: kernel.ptype.nr: Value: 19 => 18
[Apr/28/2023 05:46:15.397]: Parameter: kernel.random.uuid: Value: 37cc31aa-ee31-459e-8f2a-
0766b34b1b64 => f5176cdc-6390-415d-882e-02c4cff2ae4e
[Apr/28/2023 16:56:15.398]: Parameter: fs.aio-nr: Value: 97024 => 98560
```

Analyze Cluster Health

CHM data usage for diagnostic analysis

➤ CHM Analyser

- A python based utility that looks at CHM data and generates reports with relevant findings for CPU, memory, network, disk, I/O based on critical OS signatures.

➤ CHM Visualisation

- Built Using DASH application for viewing and analysing CHM data graphically.

➤ CHM Lrg Analyser

- A python based utility that analyses and maps diffs from watson.log and CHMOS data on a timeline. Critical system metrics like CPU utilization, Memory Utilization, CPU steal percent, load averages, etc are analyzed to detect if a diff has occurred due to depletion of critical system resources and also identify the processes which are causing system resource depletion.

CHM Metrics Viewer

“**oclumon**” - CHM metric viewer CLI to query the expanse of metric and metrics sets

- Python based utility (compatible with both Python 2.7 and Python 3.8)
- Can query metric data both in historical and continuous mode
- Directly queries compressed files (without uncompressing)
- Provides expansive set of CLI options to view metric data

Querying SYSTEM Metric Set

System metric-set detailing the summarizing state of the system.

ex : oclumon dumpnodeview local -system

```
-----  
Node : den03ceb Clock : 2023-04-28 07.31.34  
-----
```

SYSTEM:

[CPU] ←

```
pCpus[#]:4, cores[#]:4, vCpus[#]:4, cpuHT:Y, osName:Linux, chipName:Intel Core Processor (Haswell, no TSX, IBRS), usage[%]:2.7, system[%]:1.02, user[%]:1.68, nice[%]:0.0, ioWait[%]:0.15, steal[%]:0.0, cpuQ[#]:0, loadAvg1:0.69, loadAvg5:0.76, loadAvg15:0.6, intr[#/s]:3597, ctxSwitch[#/s]:8027
```

[MEMORY] ←

```
totalMem[KB]:14770492, freeMem[KB]:734904, avblMem[KB]:9257560, shMem[KB]:1830780, bufferAndCache[KB]:10574276, buffer[KB]:258528, cache[KB]:10315748, pgCache[KB]:7738936, slabReclaim[KB]:2576812, swapTotal[KB]:15359968, swapFree[KB]:14498448, hugePageTotal[#]:0, hugePageFree[#]:0, hugePageSize[KB]:2048, swpIn[KB/s]:0, swpOut[KB/s]:0, pgIn[#/s]:1, pgOut[#/s]:59
```

[DEVICE] ←

```
disks[#]:2, ioR[KB/s]:1, ioW[KB/s]:59, numIOs[#/s]:15
```

[NETWORK] ←

```
nics[#]:2, rxTotal[KB/s]:18.04, txTotal[KB/s]:9.38, nicErrs[#/s]:0
```

[PROCESS] ←

```
procs[#]:446, procsOnCpu[#]:1, procsBlocked[#]:0, rtProcs[#]:13, procsDState[#]:0, fds[#]:21056, sysFdLimit[#]:6815744
```

[NFS] ←

```
nfs[#]:87
```

Querying Process Aggregate

➤ Metrics aggregated by Process Groups (DB FG/DB BG/Other/Clusterware)

- ex. OTHER group is consuming ~96% (across 350 processes) and DB BG ~4%(across 75 processes) of total 29.56% CPU utilization.

- ex : `oclumon dumpnodeview local -procagg`

```
-----  
Node : den03ceb Clock : 2023-04-28 07.35.32  
-----  
  
PROCESS_AGGREGATE:  
category cpuWeight[%] cpu[%] rss[KB] shMem[KB] thrds[#] fds[#] processes[#] sid  
OTHER 95.93 28.36 5147760 86344 1055 6495 350 N/A  
MDBBG 3.36 0.99 19986580 1084556 83 4638 75 -MGMTDB  
CLUST 0.70 0.21 117908 102372 4 10 1 N/A  
MDBFG 0.00 0.00 2253296 994304 7 234 7 -MGMTDB
```



Querying Process Metric Set

Processes ordered by CPU, RSS, IO and Open FD's

ex : oclumon dumpnodeview local -process

Node : den03ceb Clock : 2023-04-28 06.42.04

PROCESS:
[CPU]

name	pid	pri	cpu[%]	vmem[KB]	rss[KB]	shMem[KB]	fds[#]	thrds[#]	ioT[KB/sec]	state
python	23673	20	8.34	195068	14412	6928	5	1	0.00	S
ora_vktm_-MGMTDB	18841	-2	1.66	2442188	28104	24184	54	1	0.00	S
java	19078	20	1.24	4813692	207476	5360	120	28	0.02	S
osysmond	29857	20	1.22	385116	117744	102372	10	4	0.15	S
orarootagent	4380	20	0.94	1433248	30308	5772	125	14	0.00	S

[RSS]

name	pid	pri	cpu[%]	vmem[KB]	rss[KB]	shMem[KB]	fds[#]	thrds[#]	ioT[KB/sec]	state
ora_m000_-MGMTDB	1958	20	0.00	2525688	1155016	1088284	76	1	0.00	S
ora_m002_-MGMTDB	19028	20	0.00	2558296	1125832	1056432	76	1	0.00	S
ora_m003_-MGMTDB	10088	20	0.18	2525592	1003764	951680	76	1	0.00	S
ora_m001_-MGMTDB	18654	20	0.00	2541848	950980	899164	75	1	0.00	S
oracle-MGMTDB	30646	20	0.17	2464824	925180	916904	16	1	0.00	S

[IO]


name	pid	pri	cpu[%]	vmem[KB]	rss[KB]	shMem[KB]	fds[#]	thrds[#]	ioT[KB/sec]	state
osysmond	29857	20	1.22	385116	117744	102372	10	4	0.15	S
rwhod	1447	20	0.19	6488	64	0	4	1	0.15	S
ora_lgwr_-MGMTDB	18910	-2	0.00	2461940	41720	34688	62	1	0.14	S
jbd2/vda1-8	533	20	0.00	0	0	0	2	1	0.07	S
ora_ckpt_-MGMTDB	18916	20	0.18	2461940	85928	78996	74	1	0.05	S

[FD]

name	pid	pri	cpu[%]	vmem[KB]	rss[KB]	shMem[KB]	fds[#]	thrds[#]	ioT[KB/sec]	state
ohasd.bin	2431	20	0.57	2398124	77148	18492	300	60	0.00	S
oraagent	4001	20	0.56	2323864	66408	19376	258	31	0.00	S
crsd	3116	20	0.37	2425116	80548	18508	218	62	0.00	S
osysmond	30461	-30	0.70	1756256	107308	34108	196	24	0.04	S
orarootagent	2479	20	0.19	1698328	29176	12532	194	22	0.01	S

Querying Device and NIC Metric Set

- Devices details ordered by **service time**.



```
-----  
Node : den03ceb Clock : 2023-04-28 07.31.34  
-----  
  
DISK:  
  name  ioR[KB/s]  ioW[KB/s]  numIOs[#/s]  qLen[#]  aWait[msec]  svcTm[msec]  util[%]  type  
  vdb    1.80      46.90     12.00        0.00      0.00         0.43        0.55     DISK  
  vda1   0.00      130.40    28.00        0.00      1.00         0.14        0.41     PARTITION  
  vda    0.00      130.40    28.00        0.00      1.00         0.14        0.41     DISK
```

- Network interfaces ordered by **net transmission rate** .

ex : oclumon dumpnodeview local -nic

```
-----  
Node : den03ceb Clock : 2023-04-28 07.37.53  
-----  
  
NIC:  
  name  rx[KB/s]  tx[KB/s]  total[KB/s]  rxErr[#/s]  txErr[#/s]  rxDscrd[#/s]  txDscrd[#/s]  
  eth0   2.51     37.02     39.53         0           0           0             0  
  lo     7.65     7.65     15.29         0           0           0             0
```



Querying CPU and File System Metric Set

- Individual CPU Core Details (ordered by usage)

```
-----  
Node : den03ceb Clock : 2023-04-28 07.31.34  
-----  
CPU:  
  cpuId  usage[%]  system[%]  user[%]  nice[%]  ioWait[%]  steal[%]  
    3    31.84    13.79    18.05    0.00    0.20    0.00  
    2    31.28    13.29    17.99    0.00    0.00    0.00  
    0    29.37    13.88    15.49    0.00    0.00    0.00  
    1    28.01    12.67    15.33    0.00    0.20    0.00
```

- File System Details

ex : oclumon dumpnodeview local -filesystem

```
-----  
Node : den03ceb Clock : 2023-04-28 08.15.12  
-----  
FILESYSTEM:  
  mount  type  total[KB]  used[KB]  avbl[KB]  used[%]  ifree[%]  tag  
    /    ext4  51473888  33168200  15667916  68.00    86.00    GRID_HOME
```



Metrics Repository

- Metric Repository is auto managed on the local filesystem (location and repository size are configurable)
- “Nodeview” samples are continuously written to repository (JSON record)
- Historical data is auto-archived into hourly zip files (for extended retention)
- Archived files are automatically purged once default retention limit is reached (default : 100MB)
- ex : Snapshot of metric repository directory

```
utsising@den02sxr /scratch$ ls -al /scratch/chmdata/json
total 10764
drwxr-xr-x 2 root      root      4096 Sep  6 09:00 .
drwxrwxrwx 3 utsising dba        4096 Sep  6 01:56 ..
-rw-r--r-- 1 root      root      77530 Sep  6 02:00 chmosdata_den02sxr_2020-09-06-0800.log.gz
-rw-r--r-- 1 root      root    1440450 Sep  6 03:00 chmosdata_den02sxr_2020-09-06-0900.log.gz
-rw-r--r-- 1 root      root    1425036 Sep  6 04:00 chmosdata_den02sxr_2020-09-06-1000.log.gz
-rw-r--r-- 1 root      root    1475615 Sep  6 05:00 chmosdata_den02sxr_2020-09-06-1100.log.gz
-rw-r--r-- 1 root      root    1415354 Sep  6 06:00 chmosdata_den02sxr_2020-09-06-1200.log.gz
-rw-r--r-- 1 root      root    1498497 Sep  6 07:00 chmosdata_den02sxr_2020-09-06-1300.log.gz
-rw-r--r-- 1 root      root    1477239 Sep  6 08:00 chmosdata_den02sxr_2020-09-06-1400.log.gz
-rw-r--r-- 1 root      root    1465193 Sep  6 09:00 chmosdata_den02sxr_2020-09-06-1500.log.gz
-rw-r--r-- 1 root      root     717271 Sep  6 09:04 chmosdata_den02sxr_2020-09-06-1600.log
-rw-r--r-- 1 root      root      6358 Sep  6 01:56 chmosmeta_211000.json
```

Command line operations

Check for Health Issues and Corrective Actions with CHACTL QUERY DIAGNOSIS

```
chactl query diagnosis -db oltpacdb -start "2023-04-28 01:42:50" -end "2023-04-28 03:19:15"
```

```
2023-04-28 01:47:10.0 Database oltpacdb DB Control File IO Performance (oltpacdb_1) [detected]
2023-04-28 01:47:10.0 Database oltpacdb DB Control File IO Performance (oltpacdb_2) [detected]
2023-04-28 02:59:35.0 Database oltpacdb DB Log File Switch (oltpacdb_1) [detected]
2023-04-28 02:59:45.0 Database oltpacdb DB Log File Switch (oltpacdb_2) [detected]
```

Problem: DB Control File IO Performance

Description: CHA has detected that reads or writes to the control files are slower than expected.

Cause: The Cluster Health Advisor (CHA) detected that reads or writes to the control files were slow because of an increase in disk IO.

The slow control file reads and writes may have an impact on checkpoint and Log Writer (LGWR) performance.

Action: Separate the control files from other database files and move them to faster disks or Solid

State Devices.

Problem: DB Log File Switch

Description: CHA detected that database sessions are waiting longer than expected for log switch completions.

Cause: The Cluster Health Advisor (CHA) detected high contention during log switches because the redo log files were small and the redo logs switched frequently.

Action: Increase the size of the redo logs.

Command line operations

HTML diagnostic health output available (**-html <file_name>**)

Timestamp	Target Information	Event Name	Detected/Cleared
2023-04-28 01:49:30.0	Host svr02	Host CPU Utilization	detected
2023-04-28 01:49:50.0	Host svr01	Host CPU Utilization	detected
2023-04-28 05:54:55.0	Host svr01	Host Memory Consumption	detected
2023-04-29 03:40:00.0	Host svr02	Host CPU Utilization	cleared
2023-04-29 03:40:05.0	Host svr01	Host CPU Utilization	cleared
2023-04-29 03:40:05.0	Host svr01	Host Memory Consumption	cleared

Problem	Description	Cause	Action
Host CPU Utilization	CHA detected larger than expected CPU utilization on this node. The available CPU resource may not be sufficient to support application failover or relocation of databases to this node.	The Cluster Health Advisor (CHA) detected an unexpected increase in CPU utilization by databases or applications on this node.	Identify CPU intensive processes and databases by reviewing Cluster Health Monitoring (CHM) data. Relocate databases to less busy machines, or limit the number of connections to databases on this node. Add nodes if more resources are required.
Host Memory Consumption	CHA detected that more memory than expected is consumed on this server. The memory is not allocated by sessions of this database.	The Cluster Health Advisor (CHA) detected an increase in memory consumption by other databases or by application not connected to a database on this node.	Identify the top memory consumers by using the Cluster Health Monitor (CHM).



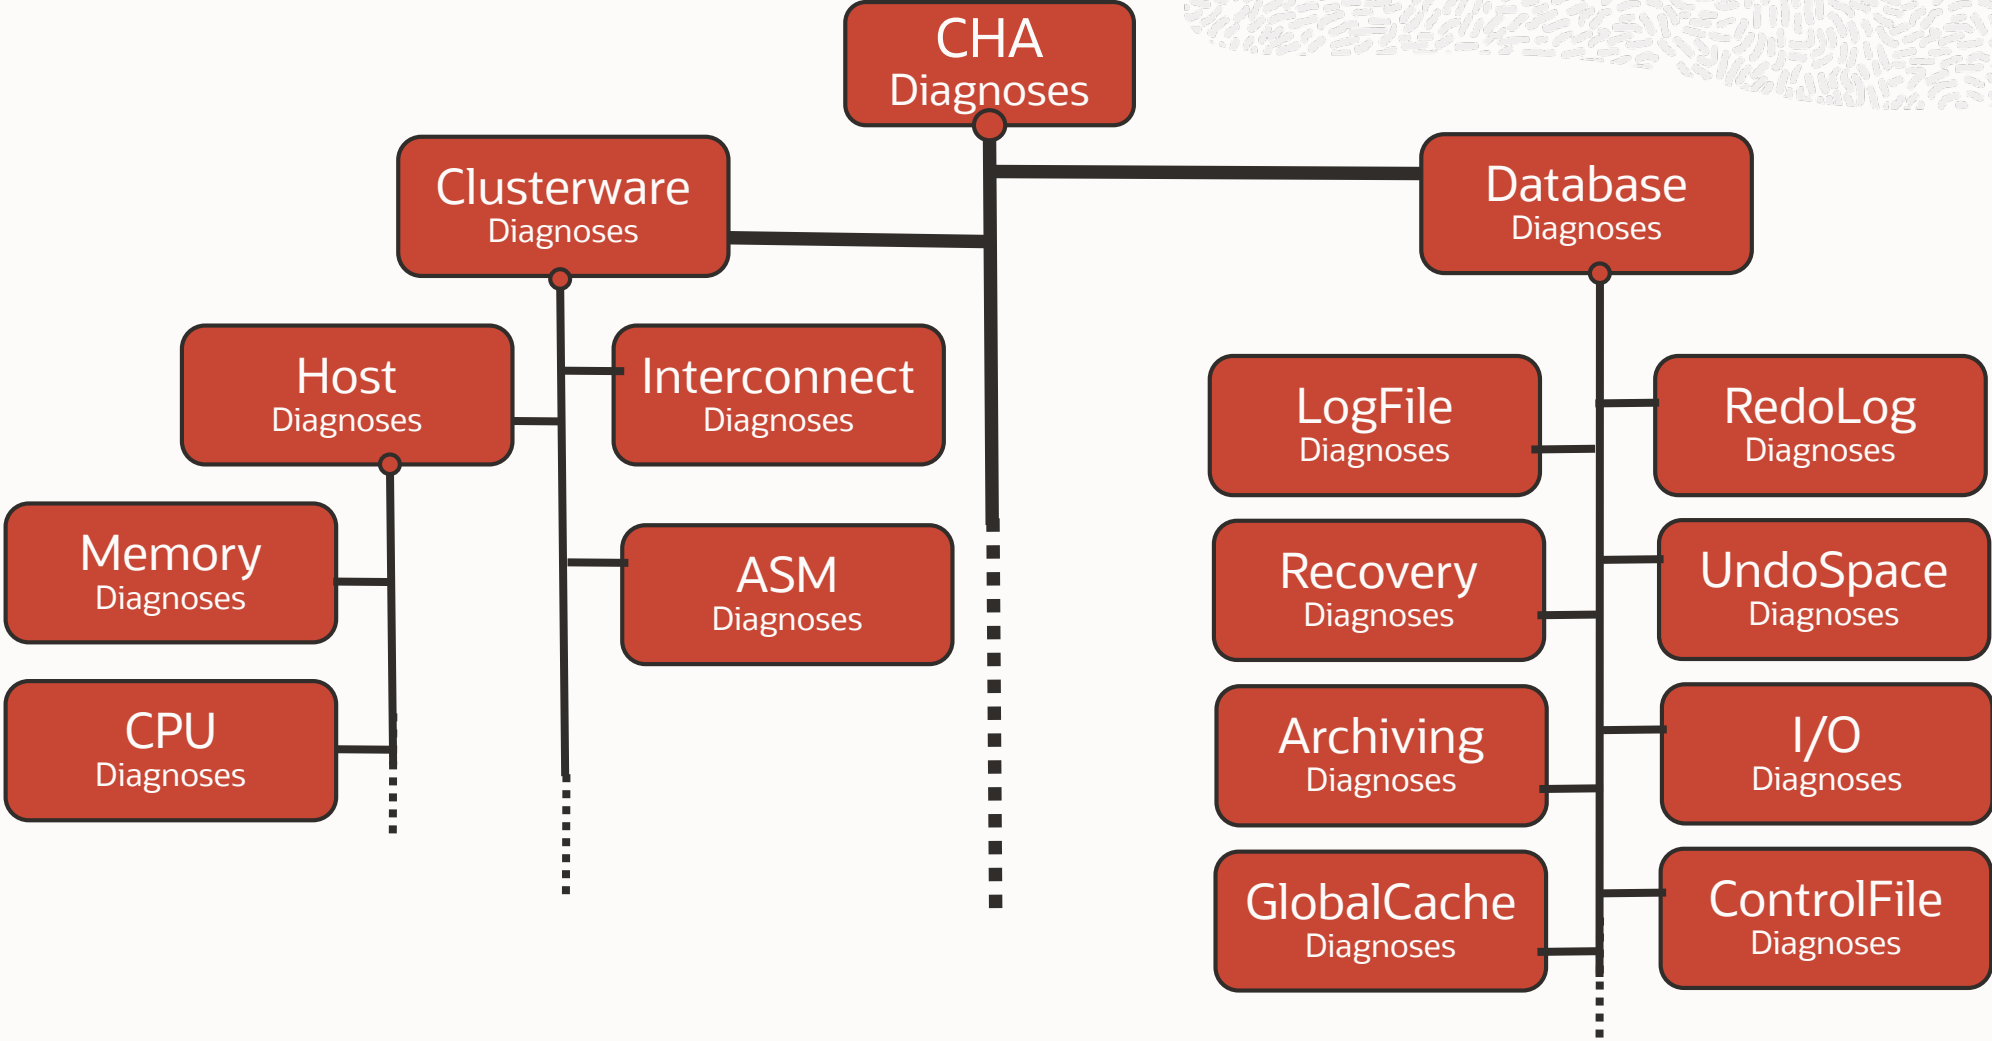
Diagnose cluster health



```
chactl query diagnosis -db oltpacdb -start "2023-04-28 01:42:50.0" -end "2023-04-28 03:19:15.0"
```

```
2023-04-28 01:47:10.0 Database oltpacdb DB Control File IO Performance (oltpacdb_1) [detected]
2023-04-28 01:47:10.0 Database oltpacdb DB Control File IO Performance (oltpacdb_2) [detected]
2023-04-28 02:52:15.0 Database oltpacdb DB CPU Utilization (oltpacdb_2) [detected]
2023-04-28 02:52:50.0 Database oltpacdb DB CPU Utilization (oltpacdb_1) [detected]
2023-04-28 02:59:35.0 Database oltpacdb DB Log File Switch (oltpacdb_1) [detected]
2023-04-28 02:59:45.0 Database oltpacdb DB Log File Switch (oltpacdb_2) [detected]
```

CHA Diagnoses Focus Areas



Critical CHA Diagnoses and Their Impacts



<u>Diagnosis ID</u>	<u>Description</u>	<u>Impact</u>
CHA_PRIV_NW_PATH	CHA detects abrupt, significant decrease in message traffic on the cluster Interconnect	Instance Eviction Node Eviction
CHA_PRIV_IC_LOSS	CHA detects slow response times for Global Cache messages	Hang Instance Eviction
CHA_GCS_BUSY	CHA detects a capacity issue in the Global Cache Services	Hang
CHA_PRIV_NETWORK_M SG	CHA detects Socket Buffer Overflows	Hang Instance Eviction
CHA_GC_NIC_CONFIG	CHA detects network packets are discarded by private network interface	Hang Instance Eviction Node Eviction
CHA_GC_IPC_CONGESTI ON	CHA detects global cache messages on the private interconnect are lost	Hang Instance Eviction

Critical CHA Diagnoses and Their Impacts



<u>Diagnosis ID</u>	<u>Description</u>	<u>Impact</u>
CHA_GC_CLUSTER_RECONFIG	CHA detects an instance is joining or leaving the cluster or a PDB open/close is taking long time	Application Performance Issue
CHA_GCS_LOG_SYNC	CHA detects slow log file writes are leading to increased contention on other instances	I/O issue
CHA_REDO_LOG_SIZE	CHA detects database sessions are waiting longer than expected for log switch completions	Hang
CHA_FRA_PERF	CHA detects database sessions are waiting longer than expected for log file switches	Hang
CHA_FRA_SPACE	CHA detects database sessions are waiting for a log switch complete	Hang
CHA_UNDO_SPACE_CONTENTION	CHA detects high contention for undo segment space	Hang

Critical CHA Diagnoses and Their Impacts



<u>Diagnosis ID</u>	<u>Description</u>	<u>Impact</u>
CHA_LOG_SYNC_HANG	CHA detects commits were blocked for several seconds	Hang
CHA_GC_HANG	CHA detects global cache messages are taking long time	Hang
CHA_CFENQ_HANG	CHA detects processes are waiting for the Control File enqueue for several seconds	Hang Instance Eviction
CHA_REDODG_HANG	CHA detects the Log Writer Process (LGWR) is waiting several seconds for a log write completion	Hang Instance Eviction
CHA_DATADG_HANG	CHA detects database read I/O's take several seconds to complete	Hang
CHA_MEM_RUL	CHA detects high rate of memory consumption,	Instance Eviction Node Eviction

CHA Diagnoses, Over 50 ... and Expanding



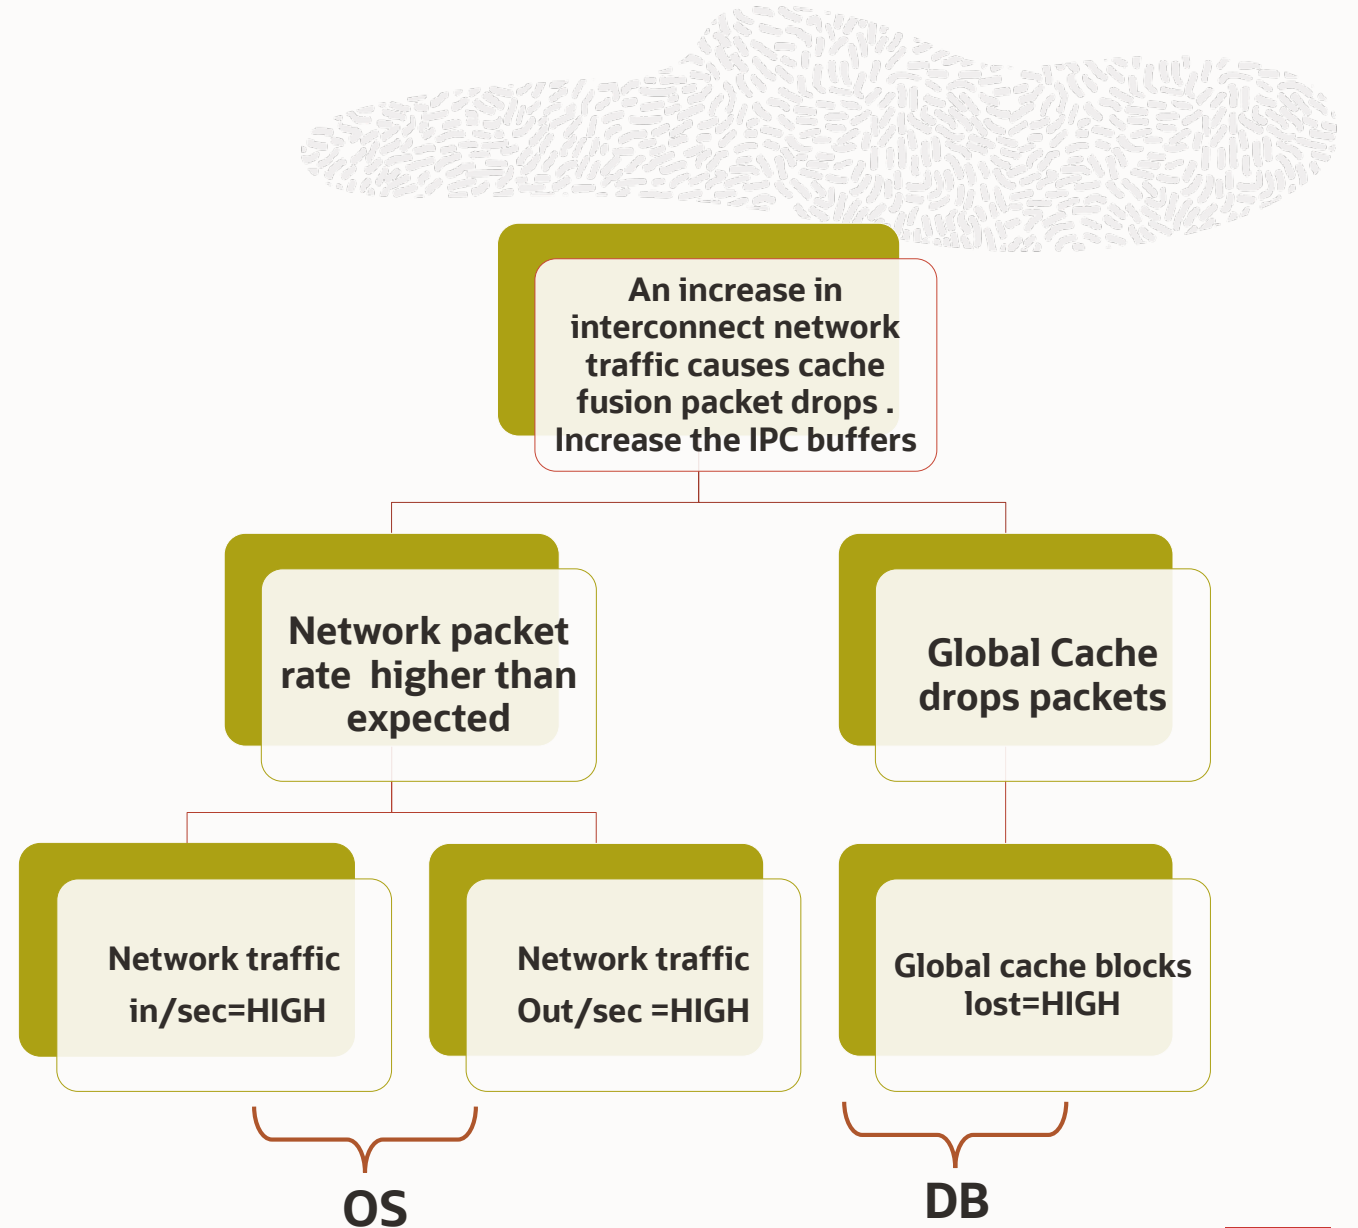
CHA_PRIV_NW_PATH	CHA_FRA_PERF	CHA_DATADG_HANG	CHA_SESSION_RATE_WARNING
CHA_PRIV_IC_LOSS	CHA_FRA_SPACE	CHA_MEM_RUL	CHA_SESS_PGAPROC_CUM
CHA_GCS_BUSY	CHA_UNDO_SPACE_CONTENTION	CHA_SESSION_RATE_WARNING	CHA_GCS_DRM
CHA_PRIV_NETWORK_MSG	CHA_LOG_SYNC_HANG	CHA_PHYSICAL_WRITES	CHA_GC_LOG_SYNC
CHA_GC_NIC_CONFIG	CHA_GC_HANG	CHA_REDO_WRITES	CHA_GC_CLUSTER_RECONFIG_TIME
CHA_GC_IPC_CONGESTION	CHA_CFENQ_HANG	CHA_RECO_CKPT	CHA_NODE_MEM_SWAP
CHA_REDO_LOG_SIZE	CHA_REDODG_HANG	CHA_SB_CPU	CHA_NODE_CPU

CHA Diagnostics Model

Why CHA Diagnoses are Correct?


They are not based on ad hoc rules!

At the core of CHA diagnostics model, an inference engine is utilized, it knows how to combine probability-measured uncertainty with our domain experience.



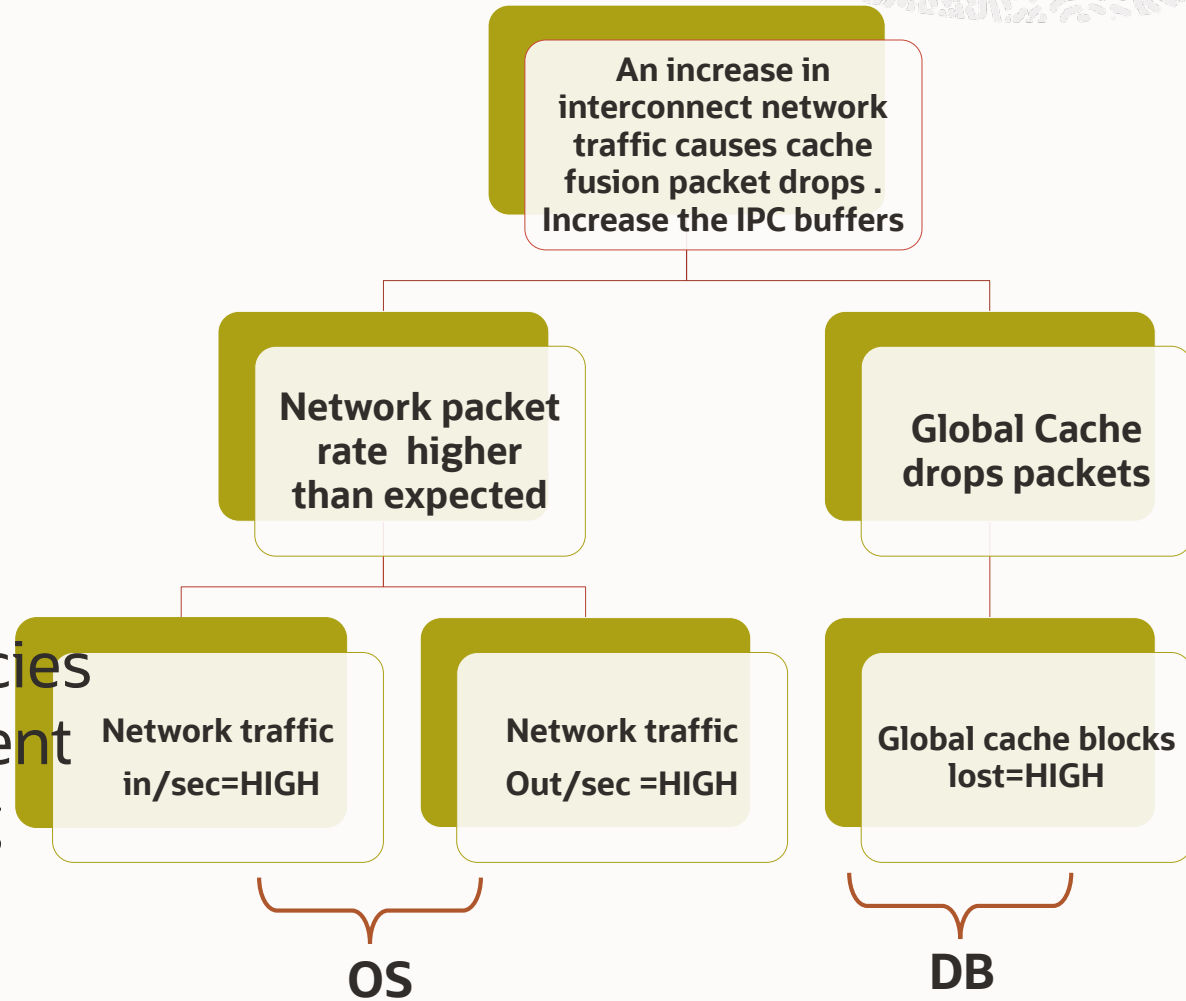
CHA Diagnostics Model

CHA Diagnostics Logic: What is in a “Model” ?




When multiple faults are detected, they are passed as evidence to a Probabilistic Bayesian Belief Network for Cause and Effect Analysis

Prior Probabilities and Dependencies are determined during development Based on historical cases and bug database



Cluster Health Advisor

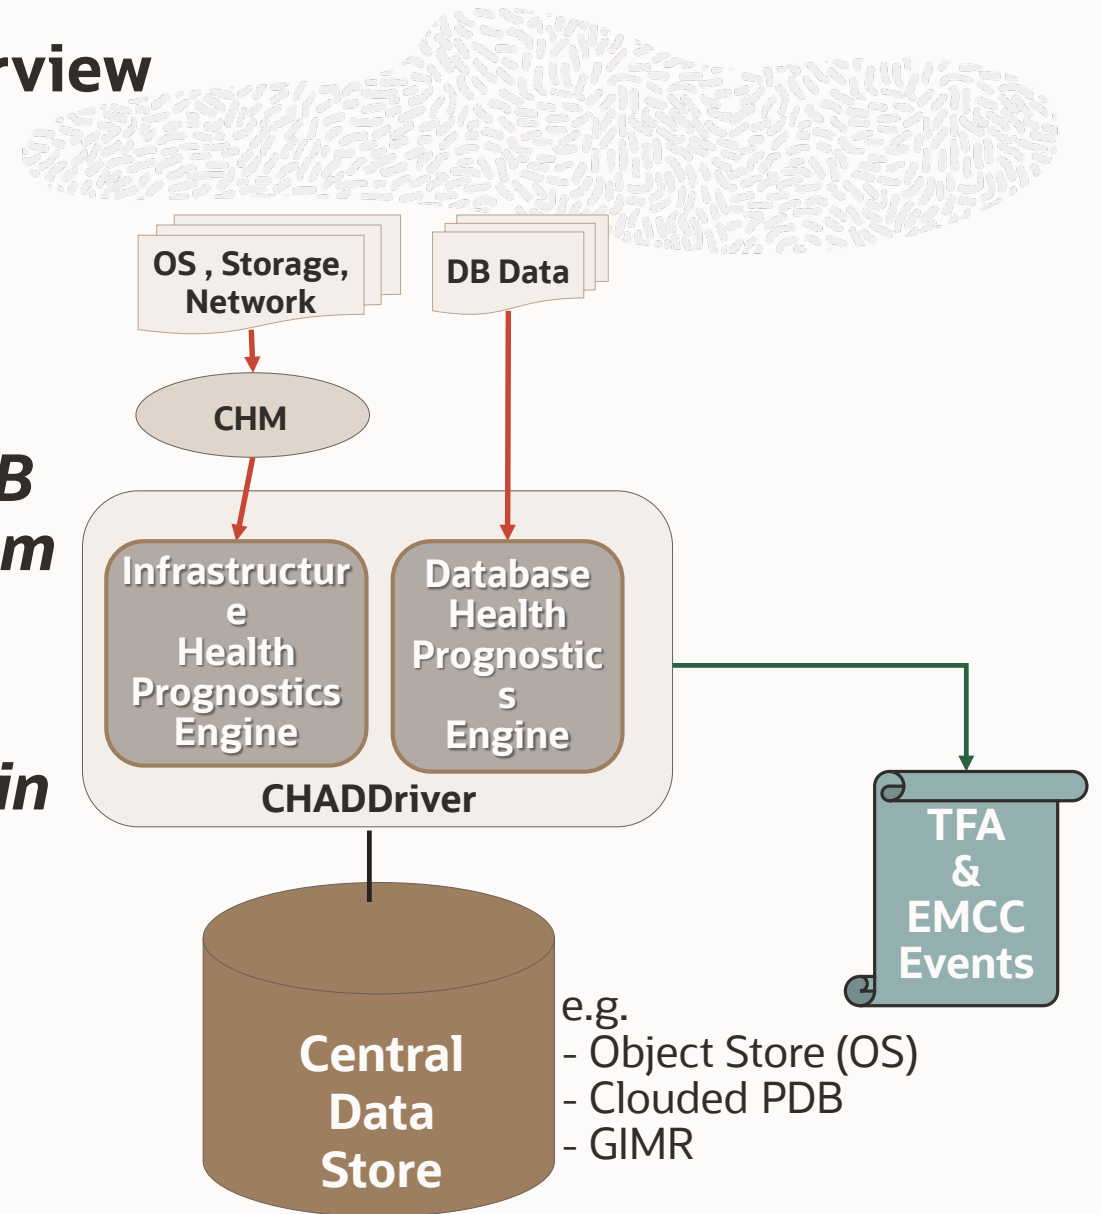
Proactive Health Prognostics System



1. Architectural Overview and GI Integration

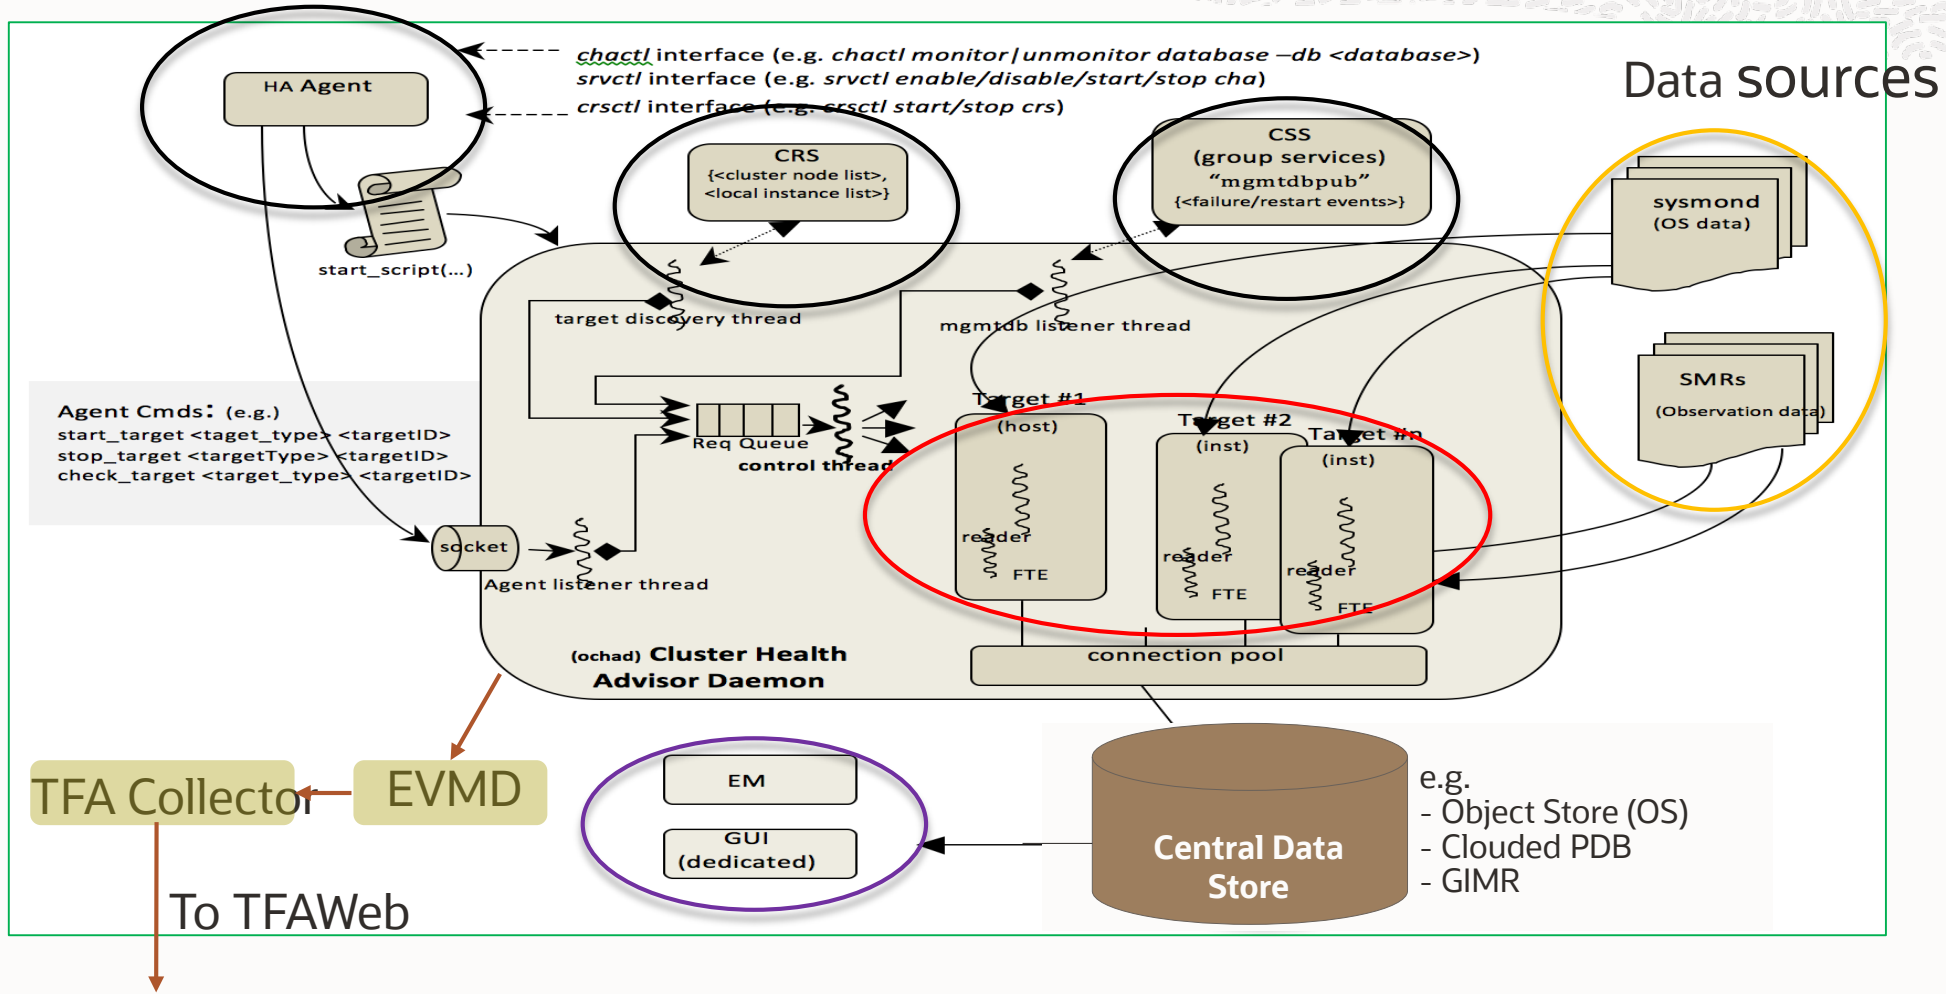
Cluster Health Advisor Architecture Overview

- **cha – Cluster node resource**
- Single Java oracle.cha.server.CHADDriver daemon per node
- Reads **Cluster Health Monitor data and DB System and session statistics directly from memory**
- Monitors GI, Storage, Network, Clusterware and DB **and performs anomaly detection in real-time**
- **Sync data, analysis, and evidence** in a central Data Store
- **Sends events** to TFA & EMCC Incident Manager



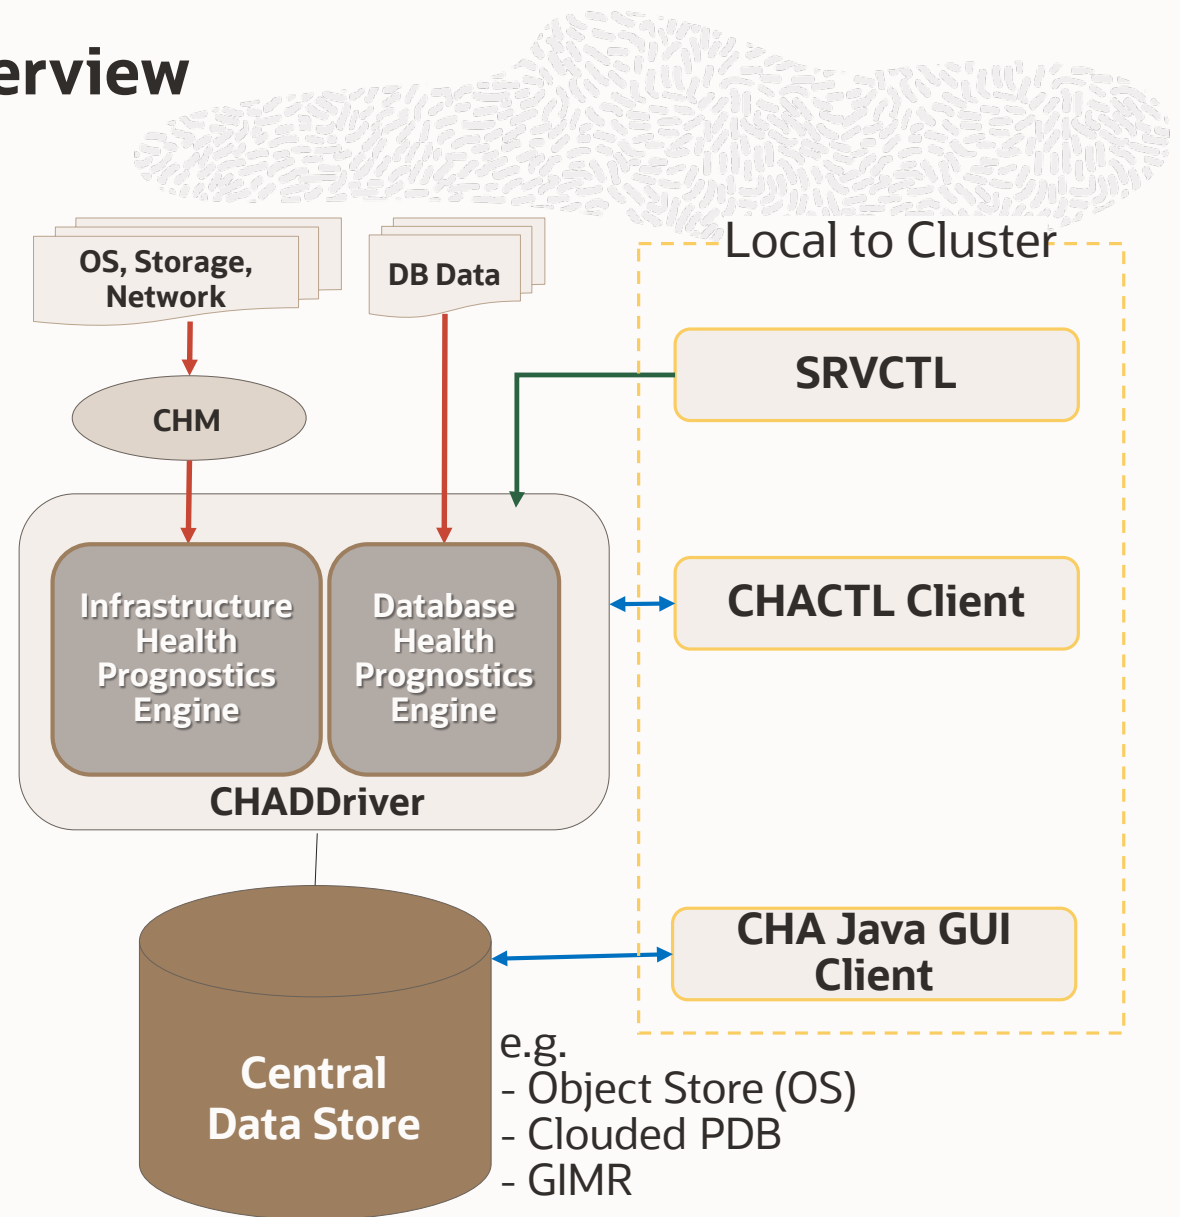
Cluster Health Advisor Architecture Overview

Fully Integrated into Clusterware: CHA Resource and Daemon



Cluster Health Advisor Operation Overview

- SRVCTL lifecycle daemon management
- Enabled by default - Activates when 1st RAC instance starts
- Comprehensive **CHACTL command line** tool for all local operations
- **TFAWeb** provides GUI interface to CHA diagnostic events
- Monitoring has no impact on DB performance or availability
- **Light footprint:**
 - ~ 2% of a single HT core for : 9 targets
 - ~ 5% of a single HT core for: 36 targets (512 MB -Xmx max heap memory)



Cluster Health Advisor

Proactive Health Prognostics System



2. In-Memory Real-Time Data Acquisition

Cluster Health Advisor

Data Sources



Operating System and Infrastructure data on each Cluster Node

DB System Statistics (e.g. physical reads , user calls) for each monitored instance and pdb

DB Session History (ASH enhanced by event history) for each monitored instance

Data from all sources is synchronized and sampled at the same frequency and forms and input vector

CHA also collects selected statistics at the PDB level.

Cluster Health Advisor

Data Sources and Data Points



A CHA *Data Point* contains > 150 signals (statistics, metrics , events) from *multiple sources*

OS, ASM , Clusterware, Network → | ← DB (ASH, AWR session, system and PDB statistics)

Time	CPU	ASM IOPS	Network % util	Network_Packets Dropped	Log file sync	Log file parallel write	Gc cr request	Gc current request	Gc current block 2-way	Gc current block busy	Enq: CF - contention	..
15:16:00	0.90	4100	13%	0	2 ms	600 us	0	0	300 us	1.5 ms	0	

Statistics are collected at a **1 second internal sampling** rate , synchronized, smoothed and aggregated to a Data Point **every 5 seconds**



rwsbj06.us.oracle.com:1 (rwsbj06:1 (oraha)) - VNC Viewer

OurCloud

- slcac454clu
 - Hosts
 - slcac454
 - slcac455
 - Databases
 - oltpacdb
 - oltpacdb_1**
 - oltpacdb_2

Nov-09 07:51:00

Instance oltpacdb_1, Expert mode.

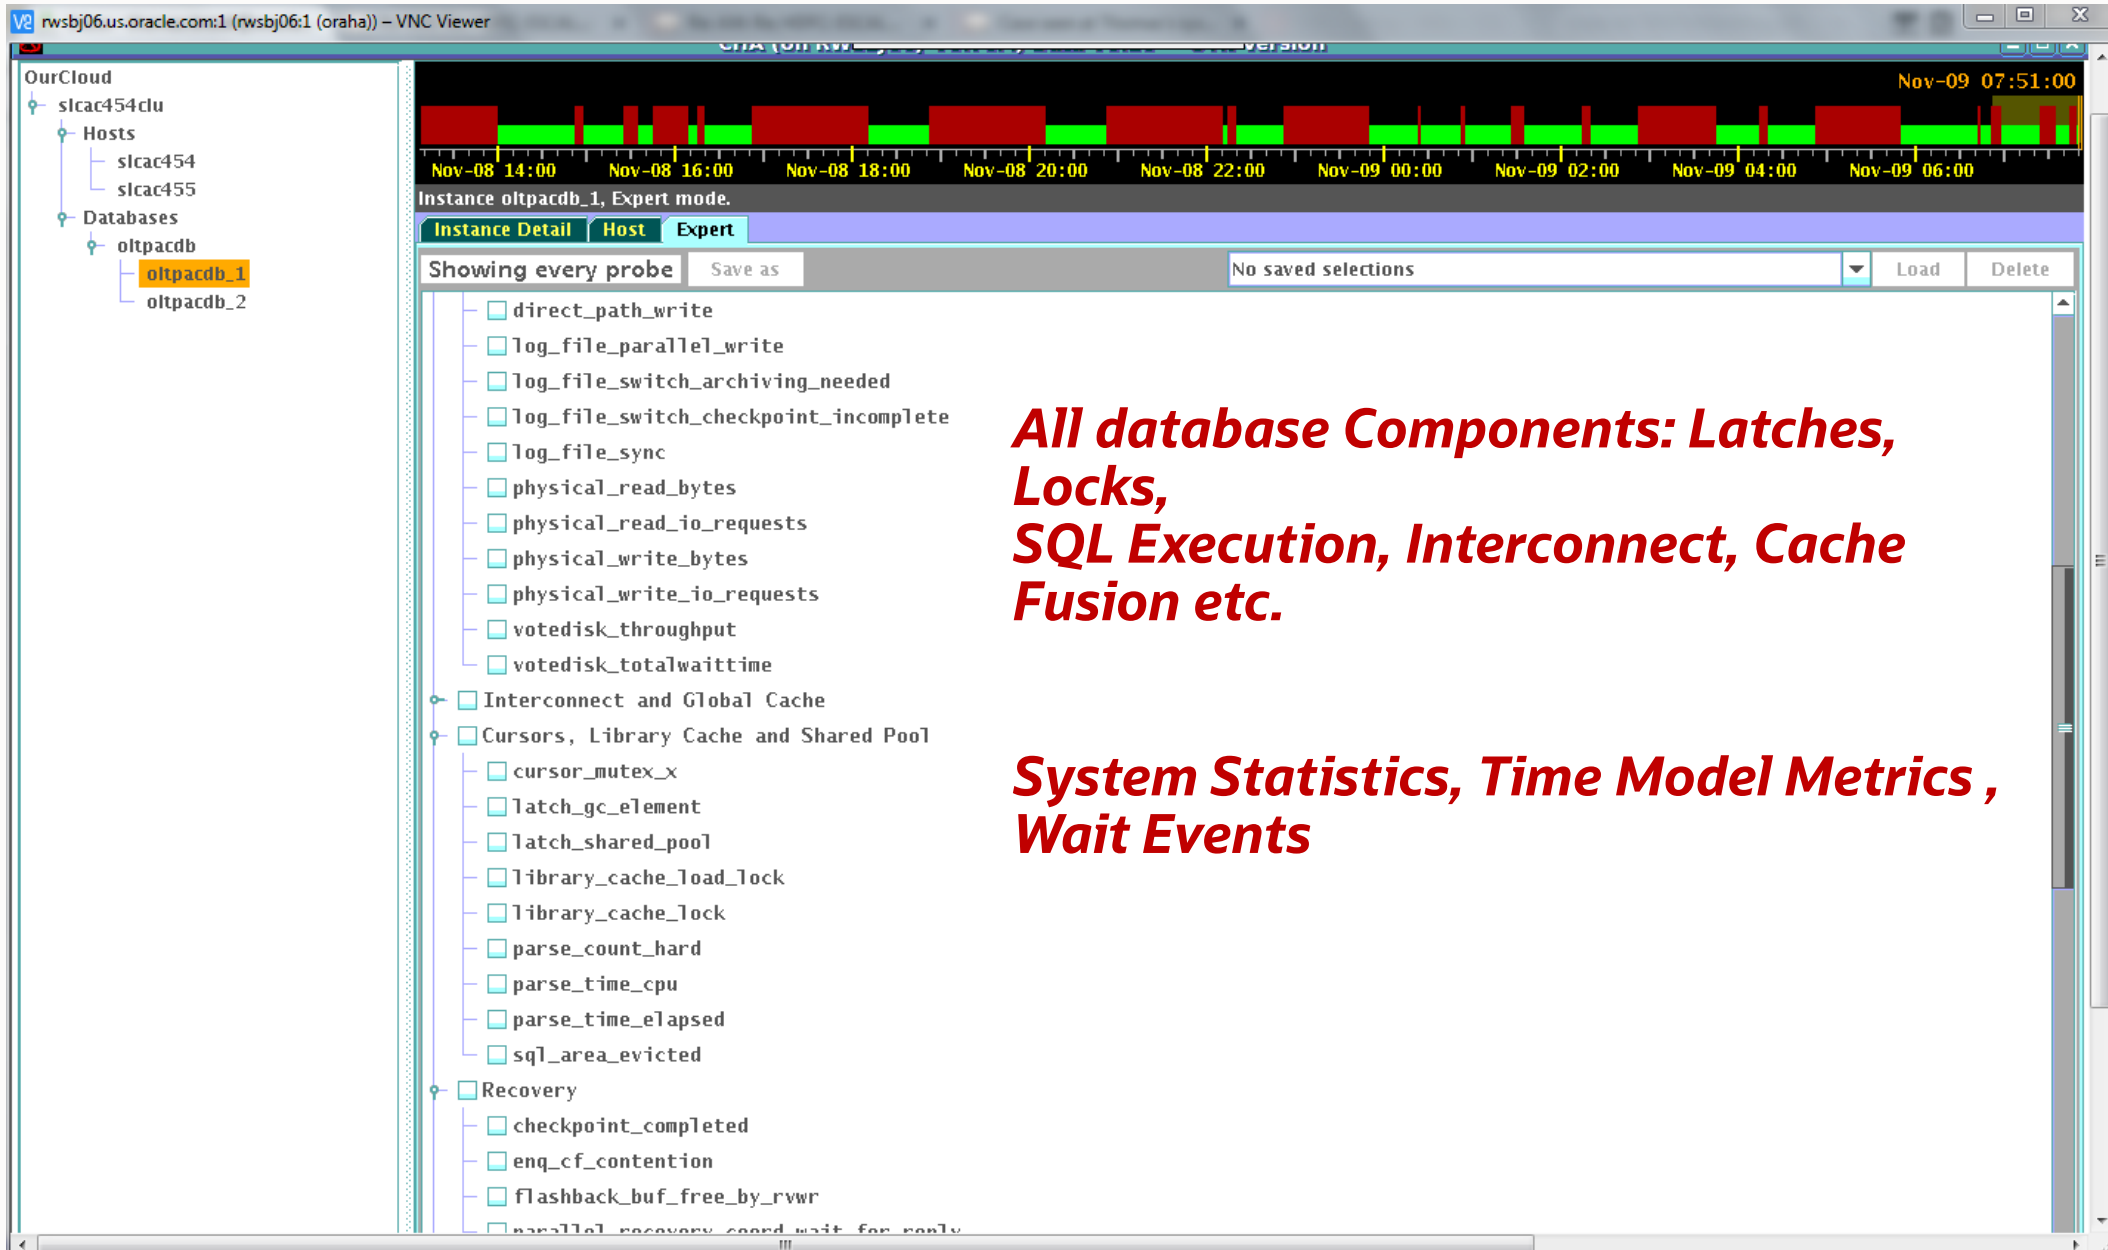
Instance Detail Host Expert

Showing every probe Save as No saved selections Load Delete

Reset

- Model Calibration Descriptors (MCD)
- CPU
- Memory
- Swap Space
- IO by Component
- Disk IO
- Interconnect and Global Cache
- Cursors, Library Cache and Shared Pool
- Recovery
- Reconfiguration
- Buffercache, Transactions and Datalayer
- Processes, Sessions and Connections

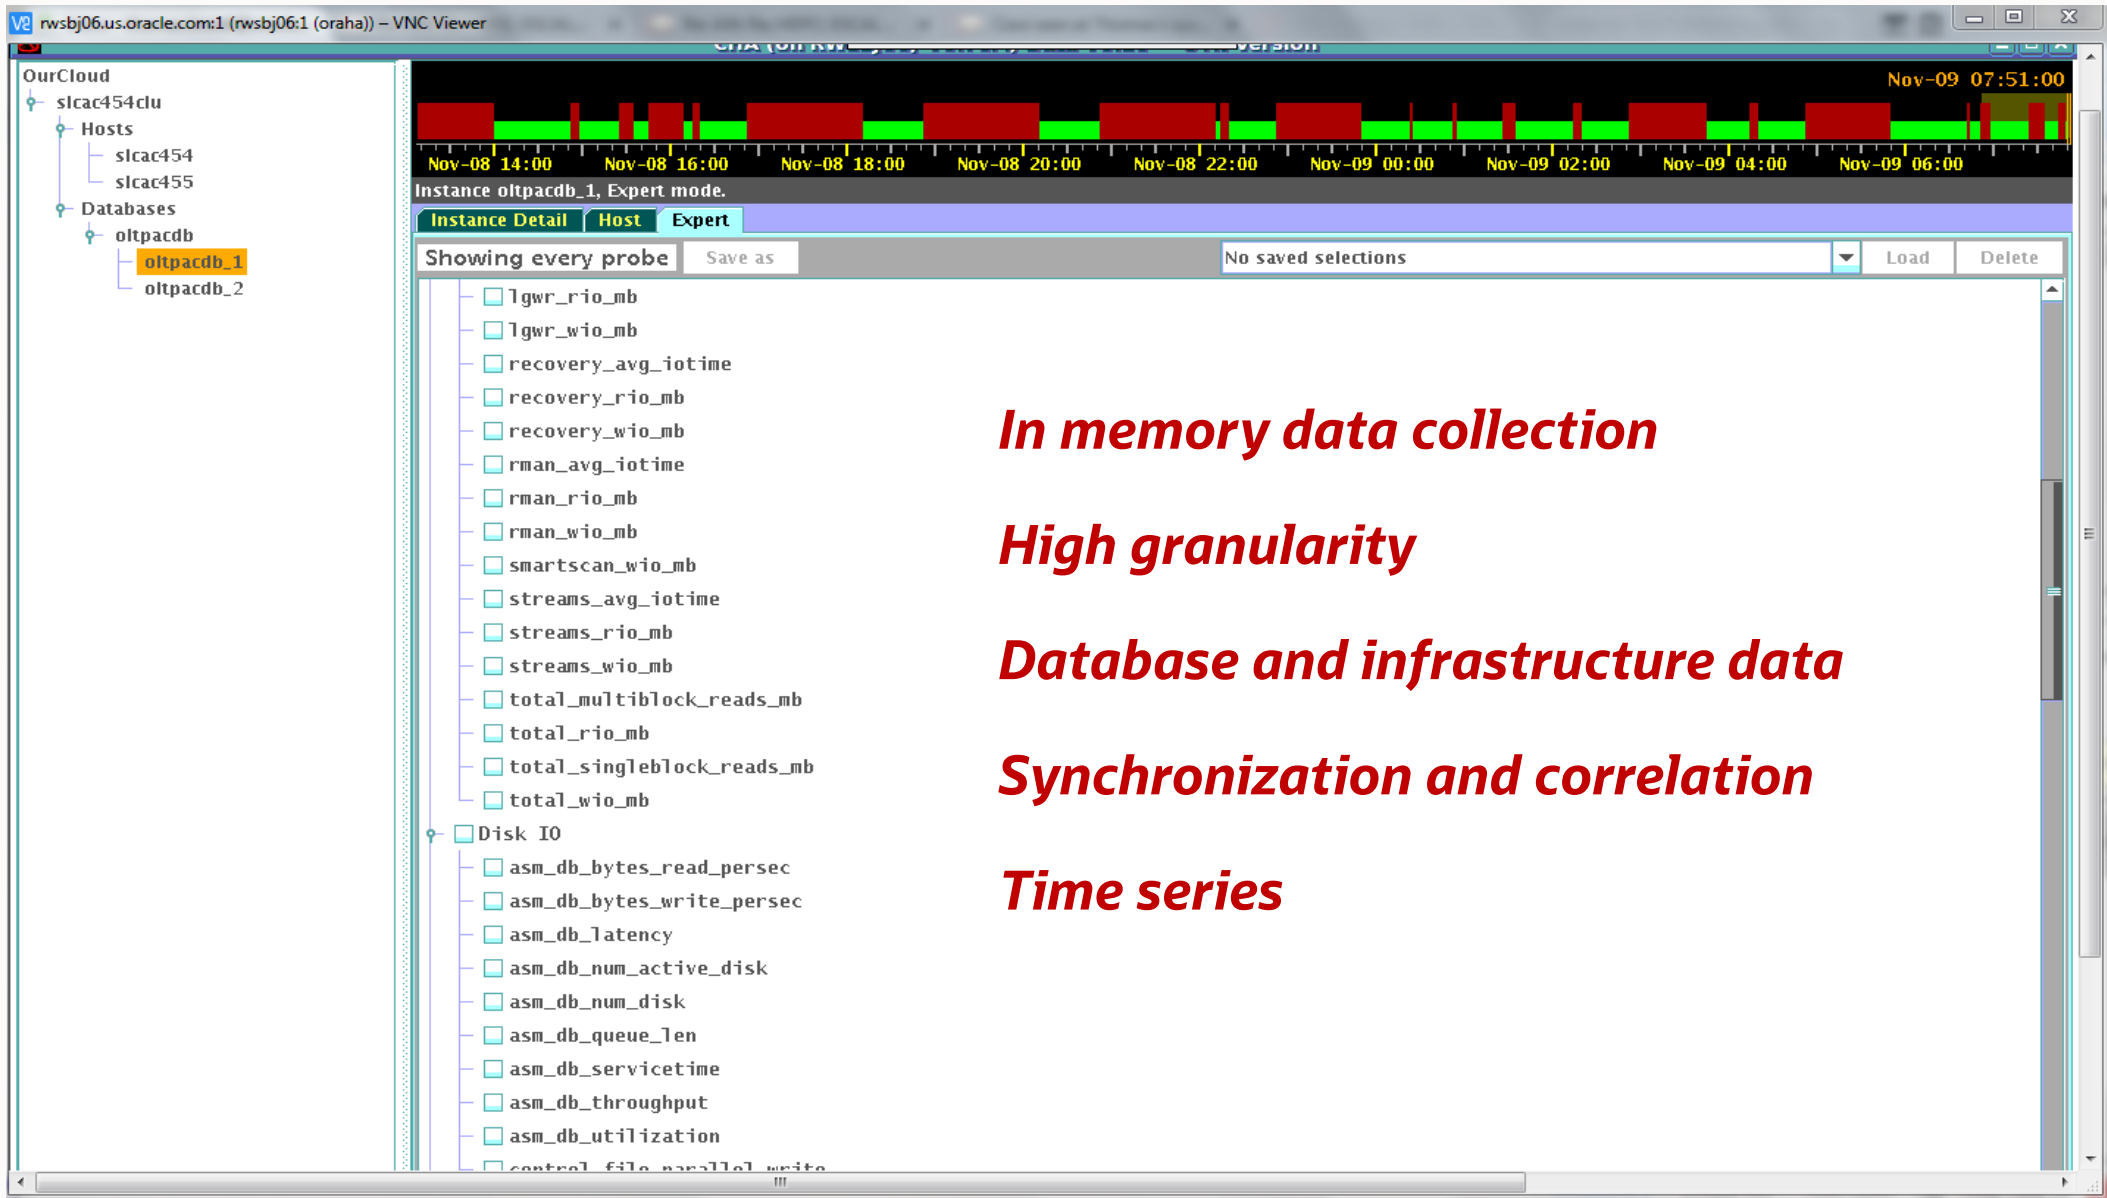
**Statistics from OS, Storage,DB, Network etc.
Grouped into component groups**



All database Components: Latches, Locks, SQL Execution, Interconnect, Cache Fusion etc.

System Statistics, Time Model Metrics, Wait Events





In memory data collection

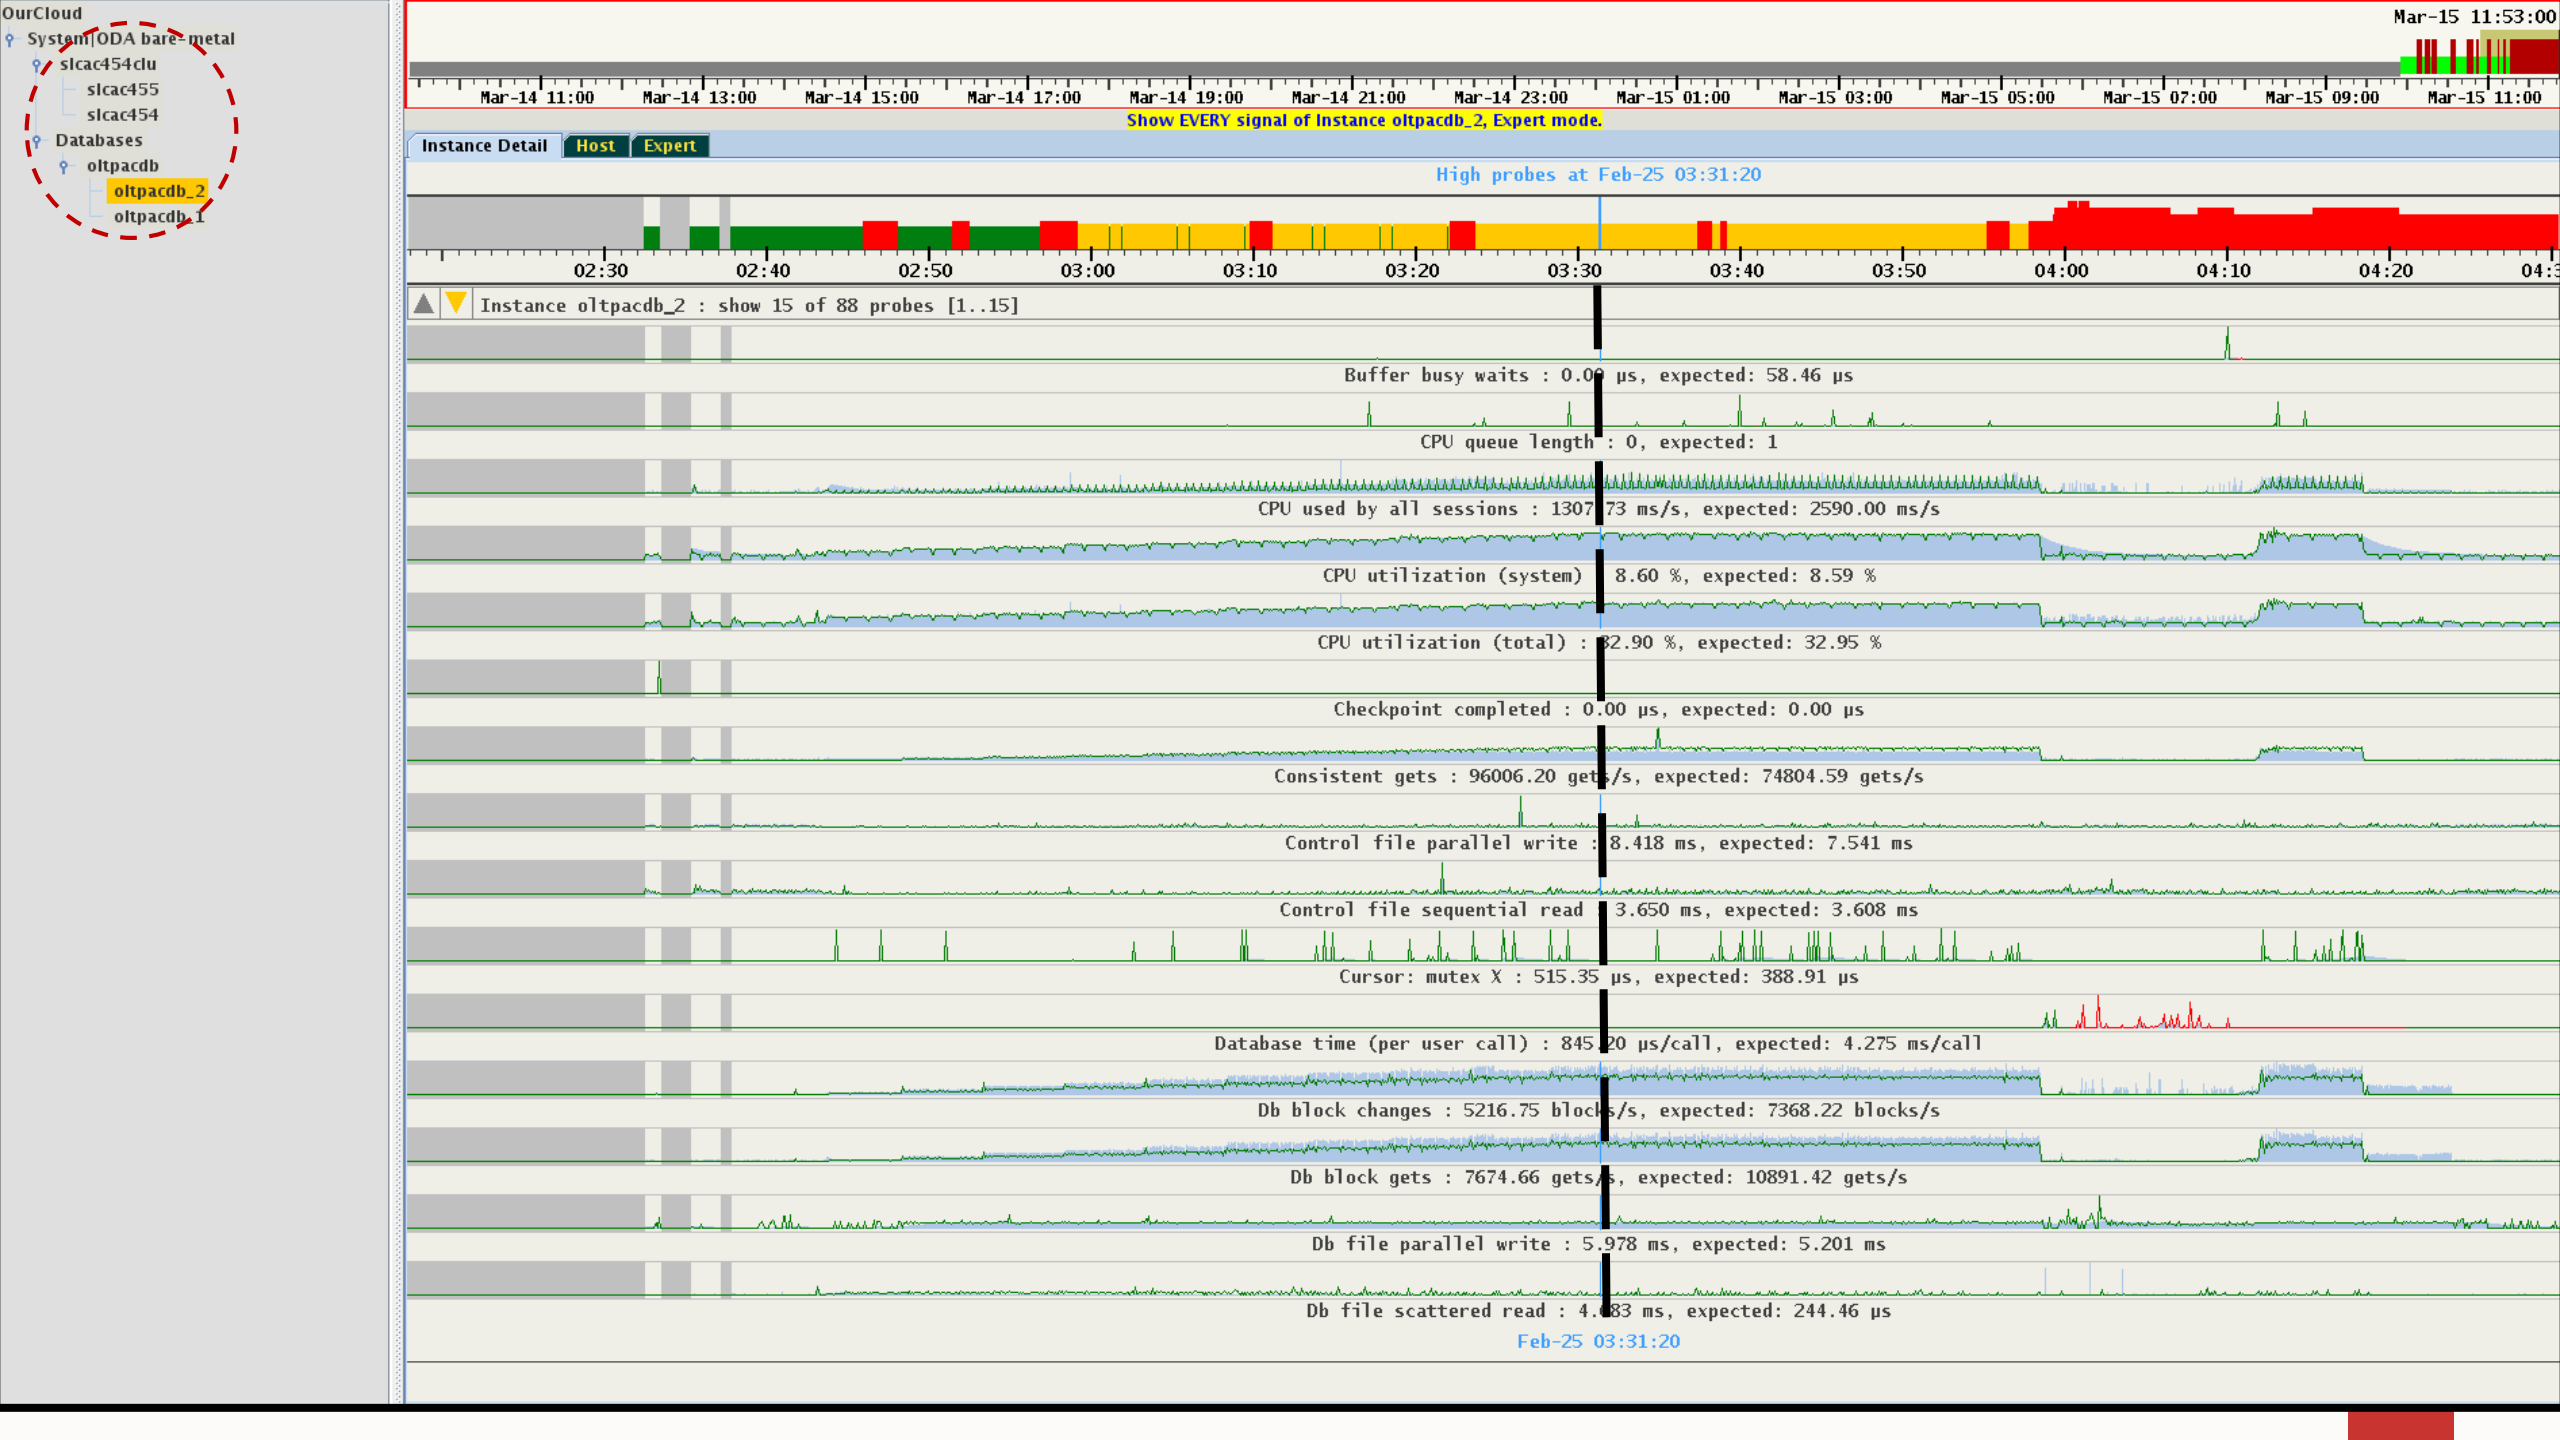
High granularity

Database and infrastructure data

Synchronization and correlation

Time series





Cluster Health Advisor

Proactive Health Prognostics System

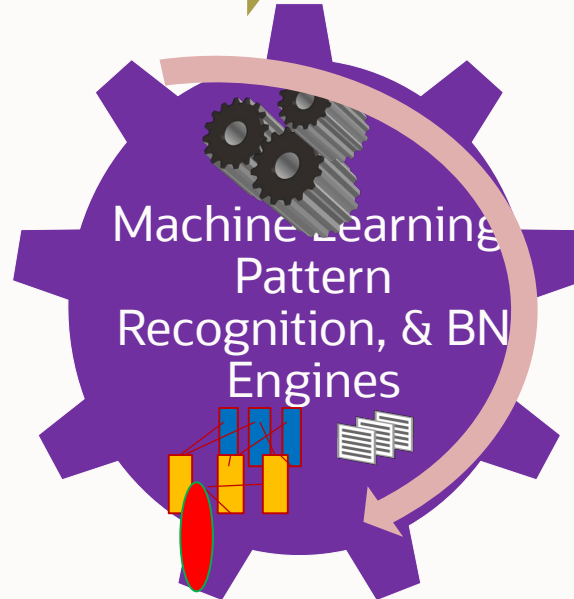
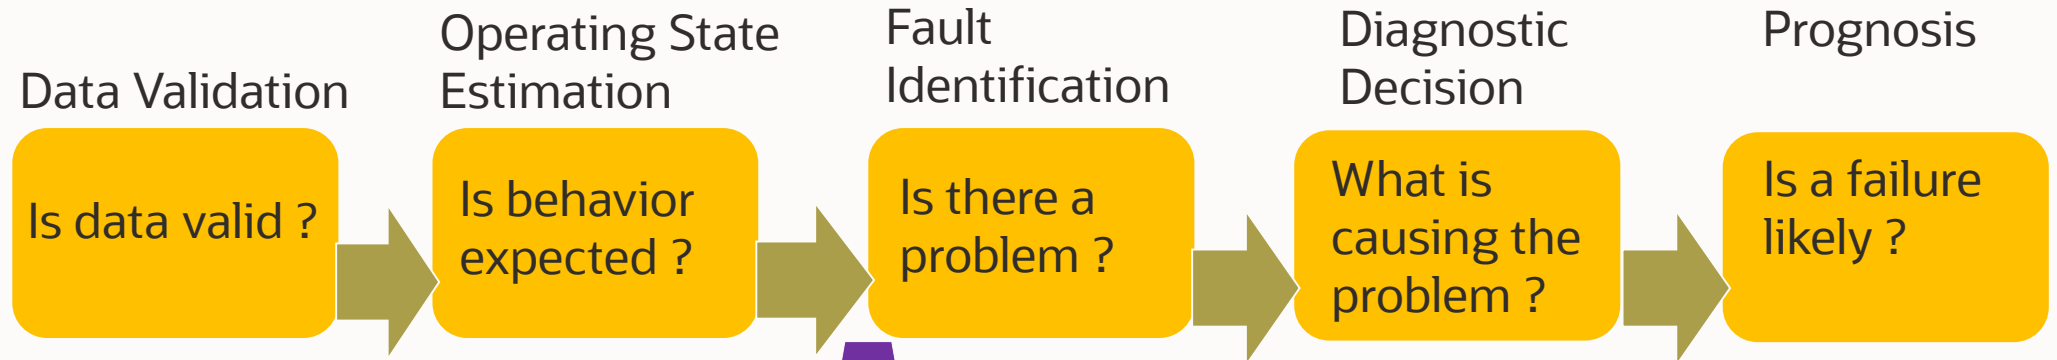


3. Prediction of Normal Operation and Anomaly Detection

Cluster Health Advisor

CHA Operational Flow : Anomaly Detection -> Diagnostics -> Prognosis

For each data point ...




Cluster Health Advisor

- ***Machine Learning and Statistical Inference*** address ***scale, dynamics and interdependency*** in Clusters and Clouds
- ***An ML Model*** is ***an in-memory representation*** of a ***normally behaving application*** over time, ***learned from historical operational data*** , in the form of a collection of ***vectors of operational data***
- The ***similarity or distance*** of a monitored data point to a vector in the in-memory model is the ***basis to for a comparison between the normal data and the actual data***

Models Capture all Normal Operating Modes

Normal load and response time variation and cyclical behavior of workloads



In-Memory Reference Matrix
(Part of "Normality" Model)


Log file sync	5.4	9.0	4.0	*
Log file parallel write	2.3	4.1	1.7	*
User commits	10000	21000	4400	*
IOPS	2500	4900	800	*

A model captures *the normal load phases* and their performance characteristics. One could say **that the model REMEMBERS the normal operational dynamics**



Cluster Health Advisor

CHA Model: Find Similarity with Normal State Vectors



In-Memory Reference Matrix
(Part of "Normality" Model)

Log file sync	*	*	4.0	*	*	*
Log file parallel write	*	*	1.7	*	*	*
User commits	*	*	4400	*	*	*
IOPS	*	*	800	*	*	*

observed values
(Part of a Data Point)

Predicted values

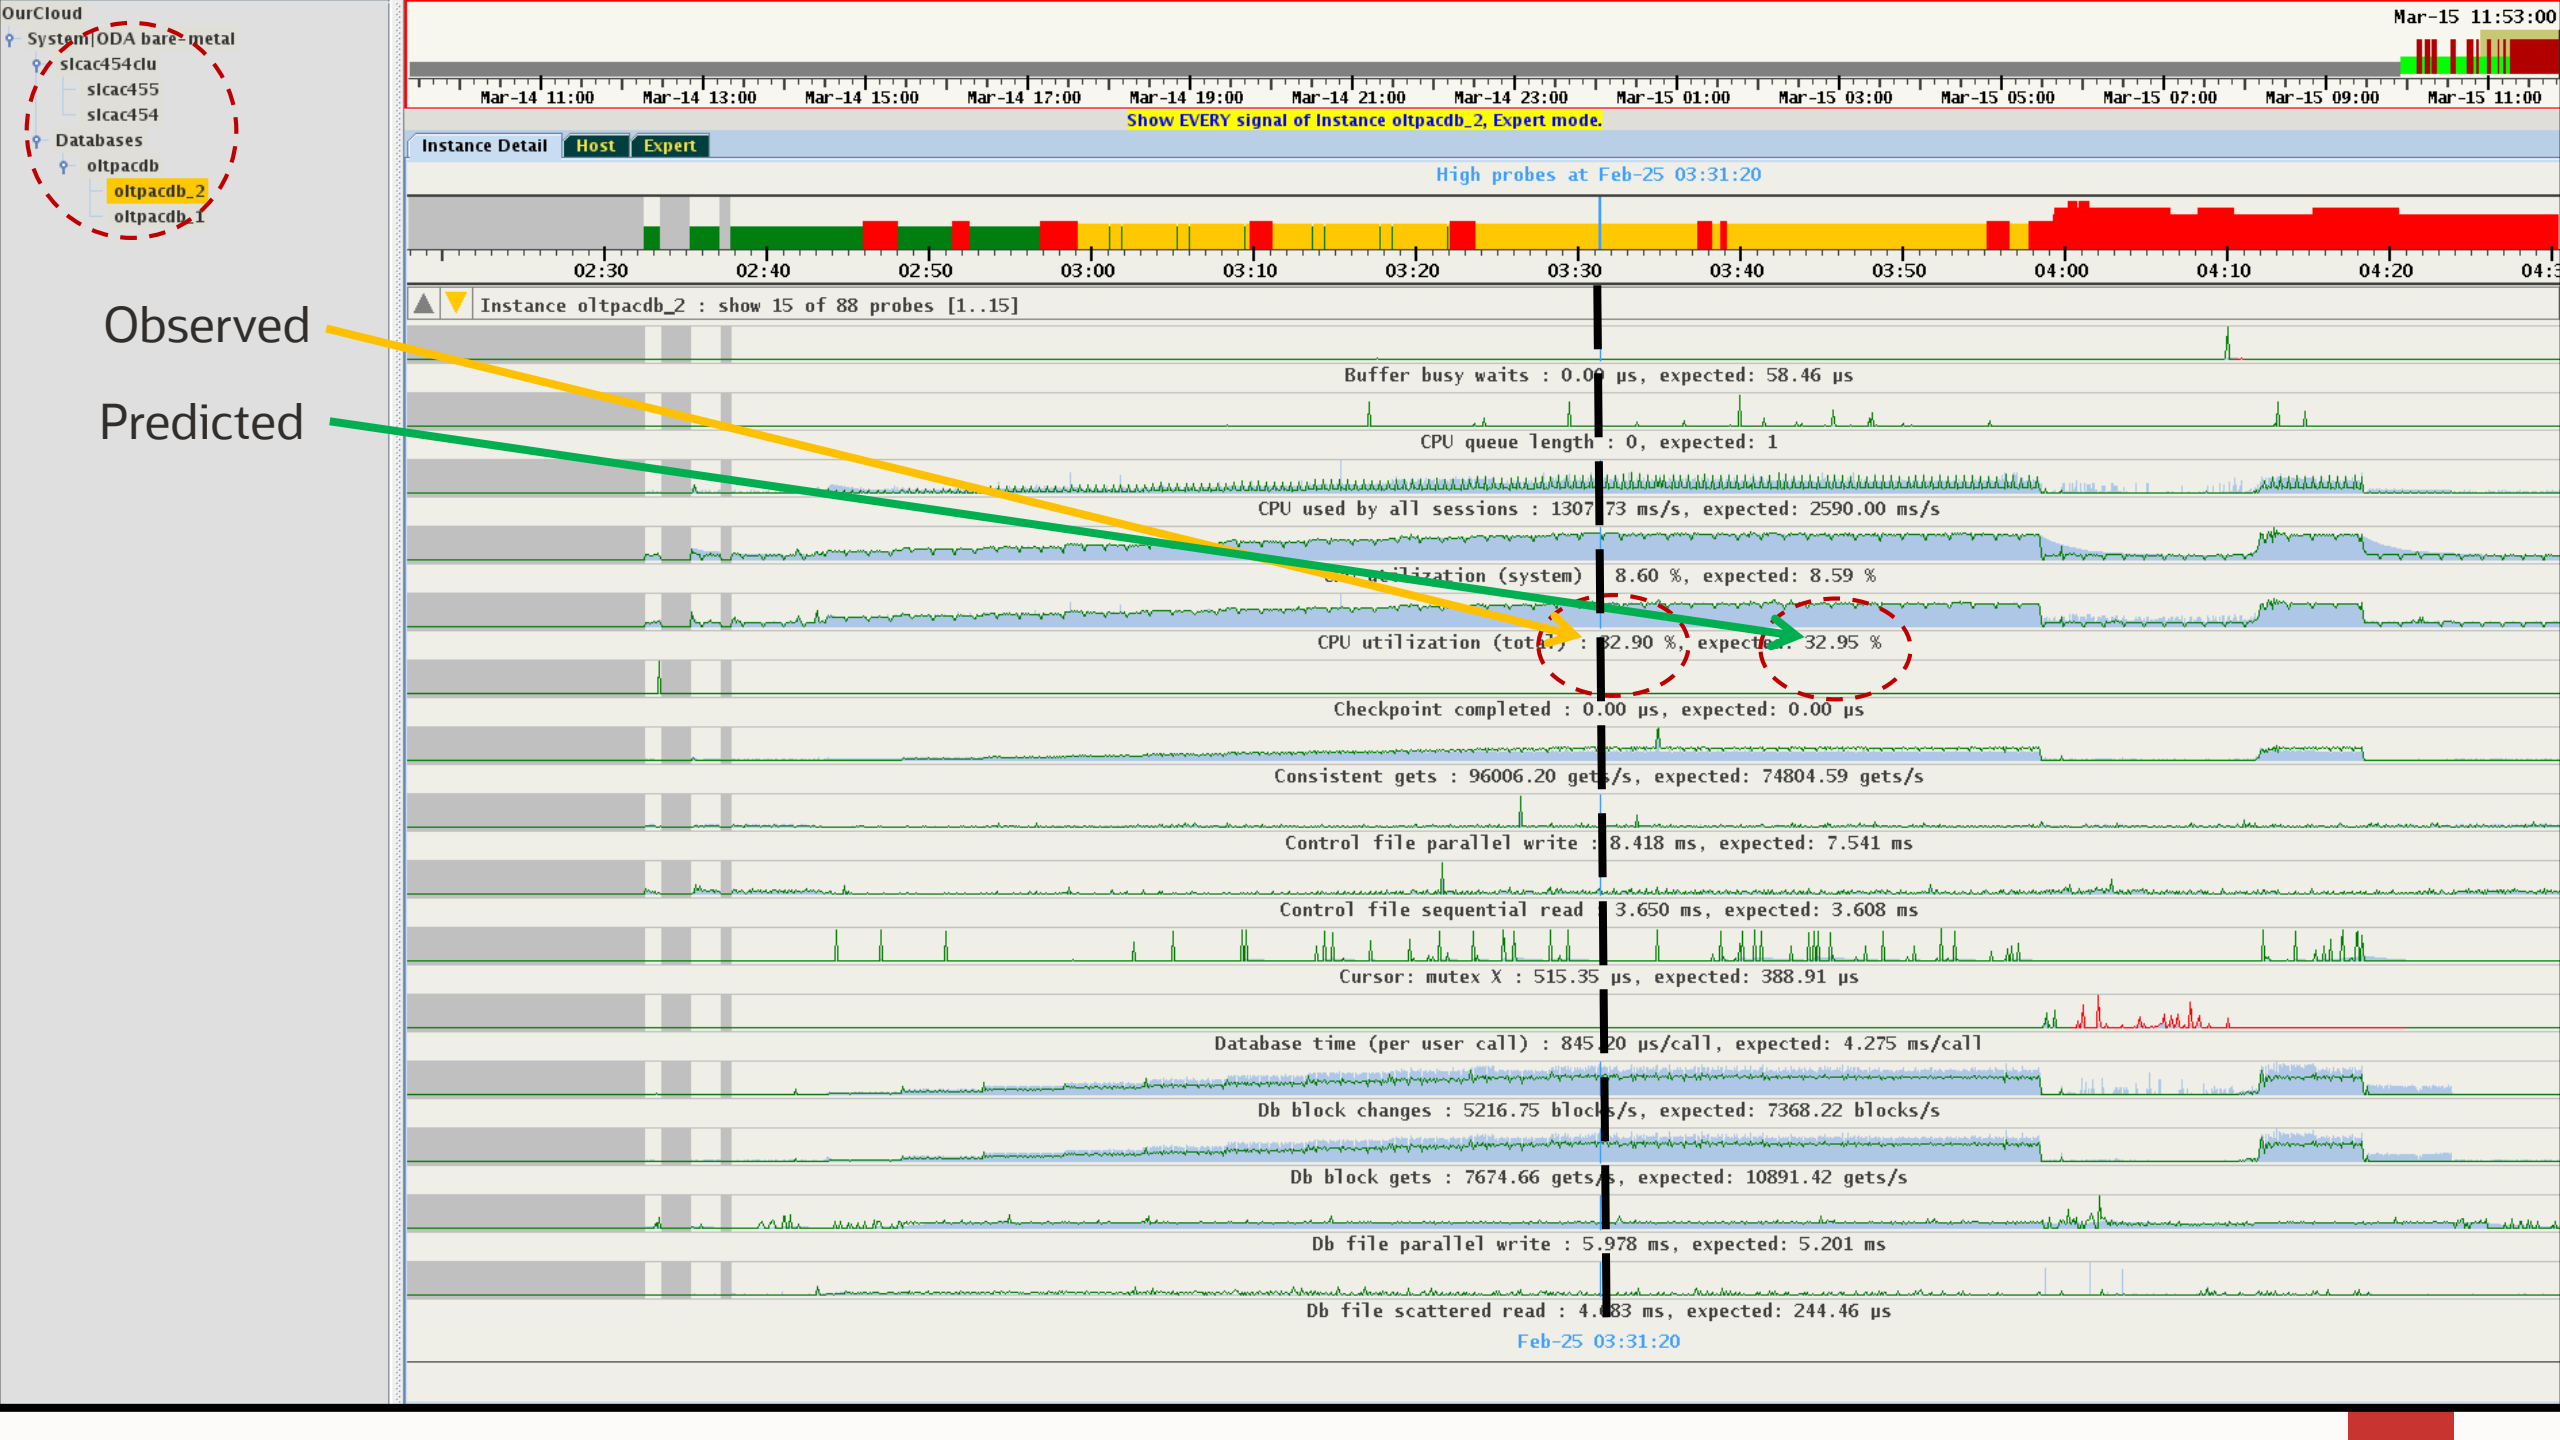
Residual

Log file Sync	9	4.3	4.7
Log file parallel write	5	1.8	3.2
User commits	600	4440	3840
IOPS	980	1100	280

CHA estimator/predictor (ESEE):
"based on my normality model, the value of log file sync should be in the vicinity of ~ 4ms, but it is reported as 9ms, this is causing a residual of ~ 5 ms in magnitude",

CHA fault detector: "such high magnitude of residuals should be tracked carefully! I'll keep an eye on the incoming sequence of this signal log file sync and if it remains deviant I'll generate a fault on it".`





Cluster Health Advisor

Fault Detection and Diagnostic Inference



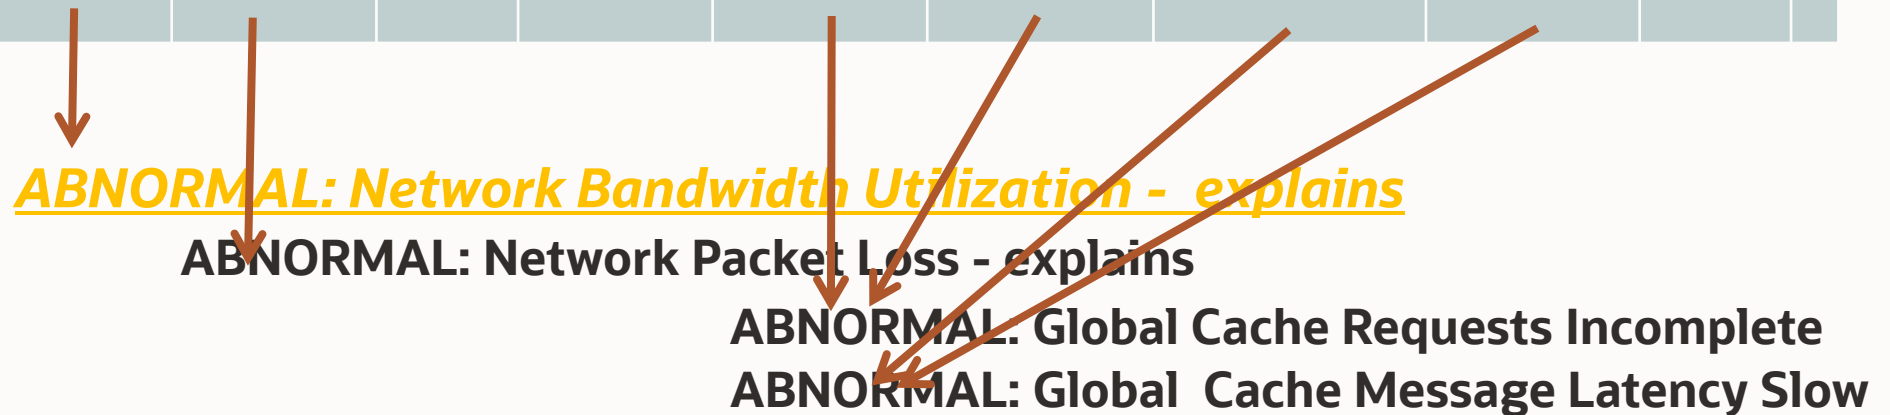
Input : Data Point at Time t

Time	CPU	ASM IOPS	Network % util	Network_Packets Dropped	Log file sync	Log file parallel write	Gc cr request	Gc current request	Gc current block 2-way	Gc current block busy	Enq: CF - contention	..
15:16:00	0.90	4100	88% 39%	105 0	2 ms	600 us	504 ms 200 us	513 ms 250us	2 ms 250 us	5.9 ms 1.5 ms	0	

Fault Detection and Classification

15:16:00	OK	OK	HIGH	HIGH	OK	OK	HIGH	HIGH	HIGH	HIGH	OK	
----------	----	----	------	------	----	----	------	------	------	------	----	--

Diagnostic Inference

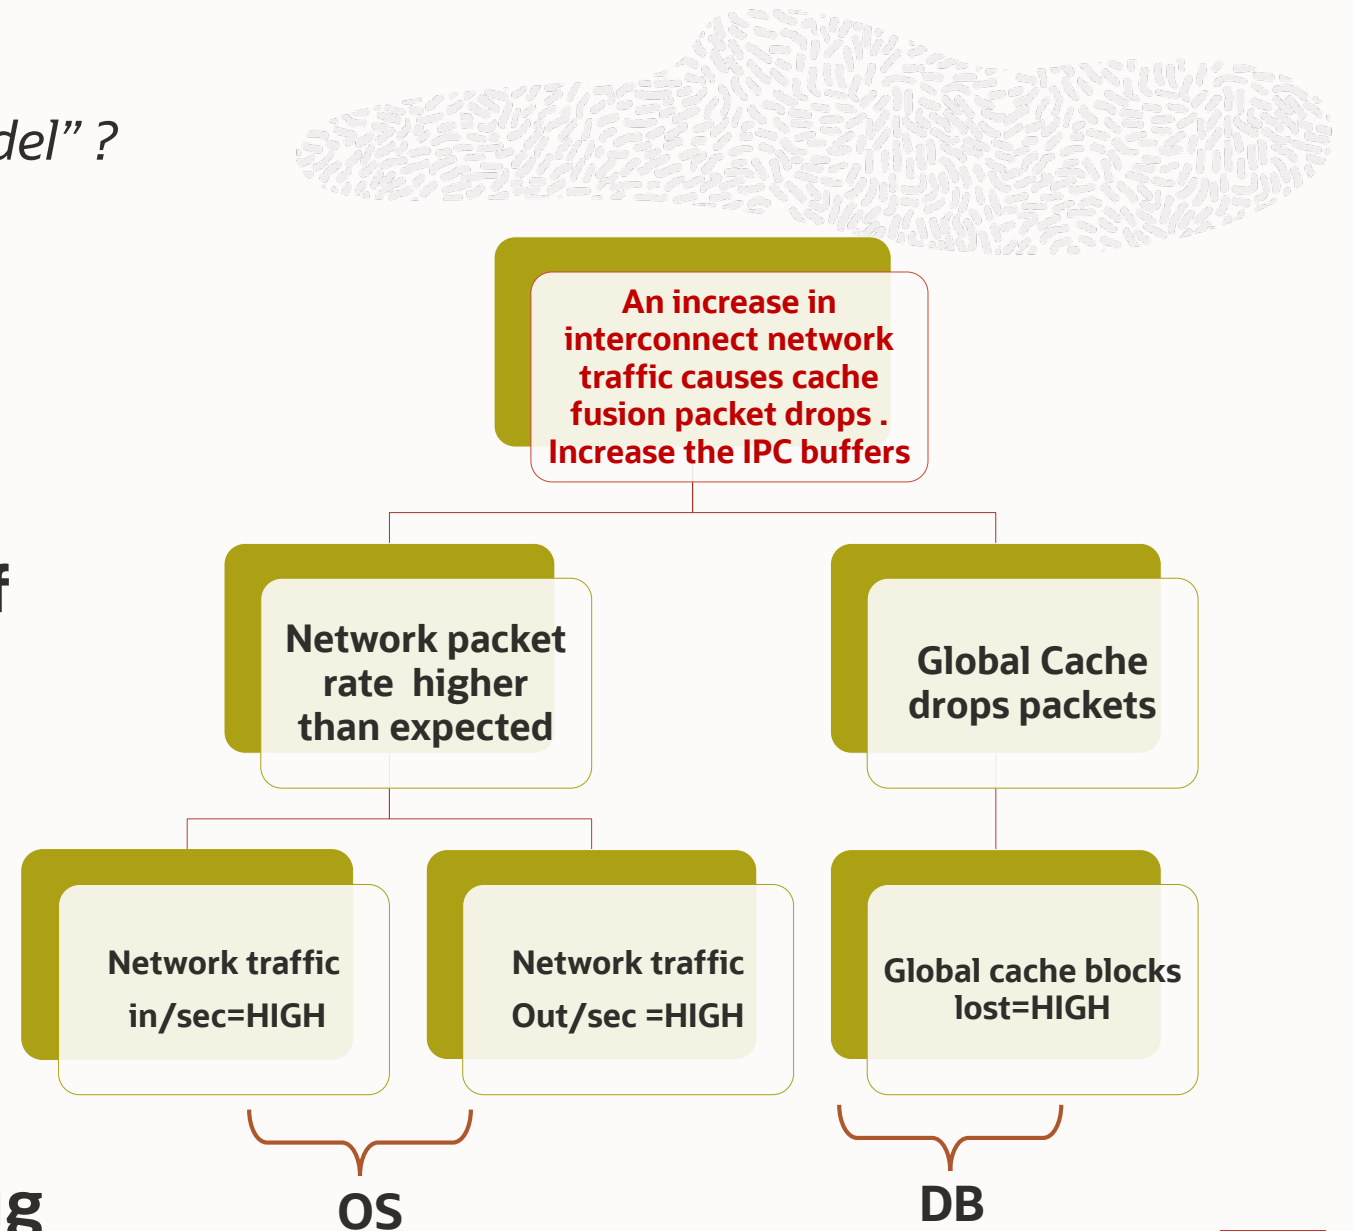


Cluster Health Advisor


CHA Diagnostics Logic: What is in a “Model” ?

When multiple faults are detected, they are passed as evidence to a Probabilistic Bayesian Belief Network for Cause and Effect Analysis

Prior Probabilities and Dependencies are determined during development Based on historical cases and bug database



- OurCloud
 - slcac454clu
 - Hosts
 - slcac454
 - slcac455
 - Databases
 - oltpacdb
 - oltpacdb_1
 - oltpacdb_2**



Instance Detail **Host** Expert

Alarm at Nov-15 06:55:45



Instance oltpacdb_2

91.5% The Cluster Health Advisor (CHA) detected that global cache messages are taking a long time.

Buffer busy waits*: 40.959 s

DB FG Wait Ratio*: 99.4 %

Database time (per user call)*: 1.369 s/call, expected: 22.537ms/call

Gc buffer busy acquire*: 908.390 ms

Gc cr request*: 11.684 s

Gc current request*: 8.597 s

Nov-15 06:55:45



OurCloud

slcac454clu

Hosts

slcac454

slcac455

Databases

oltpacdb

oltpacdb_1


oltpacdb_2

Showing abnormal probes of Instance oltpacdb_2

Probe Selection Active

Instance Detail Host

Alarm at Nov-15 06:54:20



▲ ▼ Instance oltpacdb_2 Shows 2 of 6 probes [1..2]

▼ 1: The Cluster Health Advisor (CHA) detected that global cache messages are taking a long time.

Cause:

CHA detected that global cache messages have not completed and are blocking database sessions, because a process on another instance of this database was pinning a block in the buffer cache and was waiting for a resource.

Activated at:

2017-11-15 06:52:15

Corrective Action:

Check whether incidents are detected and diagnosed on other nodes and instances in the cluster using the command 'chactl query diagnosis'.

Buffer busy waits*: 119.906 s

DB FG Wait Ratio*: 98.3 %

Nov-15 06:54:20


- OurCloud
 - slcac454clu
 - Hosts
 - slcac454
 - slcac455
 - Databases
 - oltpacdb
 - oltpacdb_1
 - oltpacdb_2



No probes found in selected set, shows EVERY probe. Host slcac454, Expert mode. Probe Selection Active

Host Detail Instances Expert

Alarm at Nov-15 06:54:20



Host slcac454

53.7% The Cluster Health Advisor (CHA) detected that more memory than expected is consumed on this server.



Physical memory utilization*: 98.0 %, expected: 88.8 %



Swap Utilization*: 8.4 %


Nov-15 06:54:20

oldashburn-preview-
iad101605exd-d0-05-06-cl-07-
09-clu01-atpd-exa-2itfs1-
hozel4qo.underscore.iad187-
hozel4qo1

Metrics


Graph

Clear





Performance Load & Run Queue - 1 Issues

Show Details




Health TFA - 6 Issues

Show Details




API Status Error - 17 Issues

Show Details



CHA Global Cache - 1 Issues

Show Details



Apply Filters ?

Time Filter ?

Applied : 2020-01-27 06:43:06 - 2020-02-10 12:00:00

Available : 2020-01-27 06:43:06 - 2020-02-10 12:00:00

- Timeline
- System Info
- Browse Files
 - List Files
 - By Directory
 - uc
 - By File Type
 - Clusterware (209)
 - Database (1959)
 - Exawatcher/OSW (21)
 - OS (2)
 - TFA (1122)
 - Generic (1703)
 - By Host
 - SELIICBL00433-RAC1 (74)
 - seliicbl00433-rac1 (3179)
 - Unknown Host Files (1763)
 - By Database
 - EOCECM (10)
 - BSCS (52)
 - Unknown (3109)
 - bscs (197)
 - eocecm (1648)
 - By Collection
 - TFA
- Analyzers
 - Block Dump Analyzer
 - Cluster Health Advisor
 - Instance Eviction
 - Node Eviction

Cluster Health Advisor

Found 3 ABNORMAL events

This Service Request contains Cluster Health Advisor (CHA) analysis dump files. CHA Diagnoses and identifies the most likely root causes for faults in OS/GI/ASM and Database and provides corrective actions for targeted triage.

Please download the CHA GUI application to your desktop to analyze the CHA analysis files. More information about CHA GUI and download options are available in MOS Note [2340062.1](#)

Timestamp	seliicbl00466-rac2	seliicbl00433-rac1
2020-02-10 02:25:50	67.9	ASM I/O Service Time [+]
2020-02-10 03:04:10	67.9	ASM I/O Service Time [+]
2020-02-10 06:02:10	87.7	Host Memory Swapping [+]
2020-02-10 06:29:10	53.8	Host Memory Consumption [+]
2020-02-10 06:29:55	87.7	Host Memory Swapping [+]
2020-02-10 07:21:10	53.8	Host Memory Consumption [+]
2020-02-10 07:21:32	87.7	Host Memory Swapping [+]
2020-02-10 07:44:20	53.8	Host Memory Consumption [+]
2020-02-10 07:44:20	53.8	Host Memory Consumption [+]
2020-02-10 11:02:20	53.8	Host Memory Consumption [+]
Description The Cluster Health Advisor (CHA) detected that more memory than expected is consumed on this server.		
Cause CHA detected an increase in host memory utilization because the number of databases processes on this server probably increased.		
Action Validate the CHA memory findings for all databases on this node.		
2020-02-10 11:05:20	53.8	Host Memory Consumption [+]



Analyze AHF Diagnostic Collections in MOS

ORA-00600-Troubleshooting Tool

Describe Problem
Upload Files
Review Recommendations

Back
Step 1 of 3
Next
Cancel

What is the Problem? ?

What would you like to do ?

Troubleshoot a new issue
 Review a troubleshooting report

Tip

This tool will provide recommendations to resolve ORA-600 issues based on details found in the uploaded IPS or Trace/Incident files. Click [Document 1521912.1](#) to see why you should use this tool!

To fully benefit from this tool all requested files should be uploaded to this tool. For details regarding the requested files and how to obtain them, see [Document 1521912.1](#) (ORA-600 Troubleshooting Tool).

If you don't have a trace file please use the [Document 153788.1](#) (ORA-600 / ORA-7445 Error Look-up Tool).

Press the NEXT button to continue.

Guided Resolution is always available from



ORA-00600-Troubleshooting Tool



Upload Files

Choose the radio button for one of the below sets of requested files to use for troubleshooting
 -Diagnostic files will be analyzed and a personalized solution will be provided if exists
 -SR fields will be automatically populated if you choose to create an SR

Click the UPLOAD button after choosing files from your local file system to use for troubleshooting.

Refer to [Document 1521912.1](#) to see why you should use this tool.

Note: Do not submit any personal information of European residents, protected health information subject to HIPAA, or any other sensitive personal information (such as payment card data) that requires protections greater than those specified in the [Oracle GCS Security Practices](#).

Select File Group	<input checked="" type="radio"/>	* TFA package (Recommended approach.)	<input type="button" value="Choose File"/>	myserver69....C_2018.zip
Select File Group	<input type="radio"/>	* IPS Package	<input type="button" value="Choose File"/>	No file chosen
Select File Group	<input type="radio"/>	* tracefile	<input type="button" value="Choose File"/>	No file chosen
Select File Group	<input type="radio"/>	Alertlog(Optional)	<input type="button" value="Choose File"/>	No file chosen

Tip

To obtain the most accurate diagnosis, upload the requested files.

Ensure you are uploading the correct files from the instance in which you are having issues. File upload combinations are:

- TFA is the recommended approach and will gather all relevant diagnostics for the problem using one command **"tfactl diagcollect -srcd ora600"**. For more information refer to TFA Document [Document 2165632.1](#)
- Incident Packaging Service (IPS) package or any archive file containing trace files with the error
- Alert log AND Incident dump file (Release 11g or higher)/Trace file (Prior to Release 11g)

For details regarding the requested files and how to obtain them, see [Document 1521912.1](#) (ORA-600 Troubleshooting Tool).



ORA-00600-Troubleshooting Tool

Describe Problem
 Upload Files
 Review Recommendations

Back
 Step 3 of 3
 Finish
 Restart

Troubleshoot a new issue: ORA-00600

Enter a report name and click 'Save' Save

Copyright (c) 2018, Oracle. All rights reserved. Oracle Confidential.

★ **ORA-600 [ktfbtgex-7] "Bitmap Managed Extent goes past end of file" (Doc ID 178643.1)** 📄 To Bottom

Modified: 04-Aug-2018 Type: REFERENCE +/- ✉ 🔗 🖨️

Note: For additional ORA-600 related information please read [Document:1092832.1](#)

This article represents a partially published OERI note.

It has been published because the ORA-600 error has been reported in at least one documented bug and/or confirmed Support Related Article.

Therefore, the SUGGESTIONS section of this article may help in terms of identifying the cause of the error.

Tip

Based on the data available, it was not possible to identify a known bug or solution. A knowledge document containing details for the error is provided when available. In most cases the knowledge document will contain a list of known bugs which have been resolved for the error.

For further assistance click on the 'Create SR' button below to log an SR with Oracle Support.

Help us improve this tool by providing feedback using the feedback form on this page.



ORACLE