

ORACLE

# Analysis of Database Issues using AHF and Machine Learning

OGBEMEA Oct 2022

**Sandesh Rao**

VP AIOps for the Autonomous Database

 [@sandeshr](https://twitter.com/sandeshr)

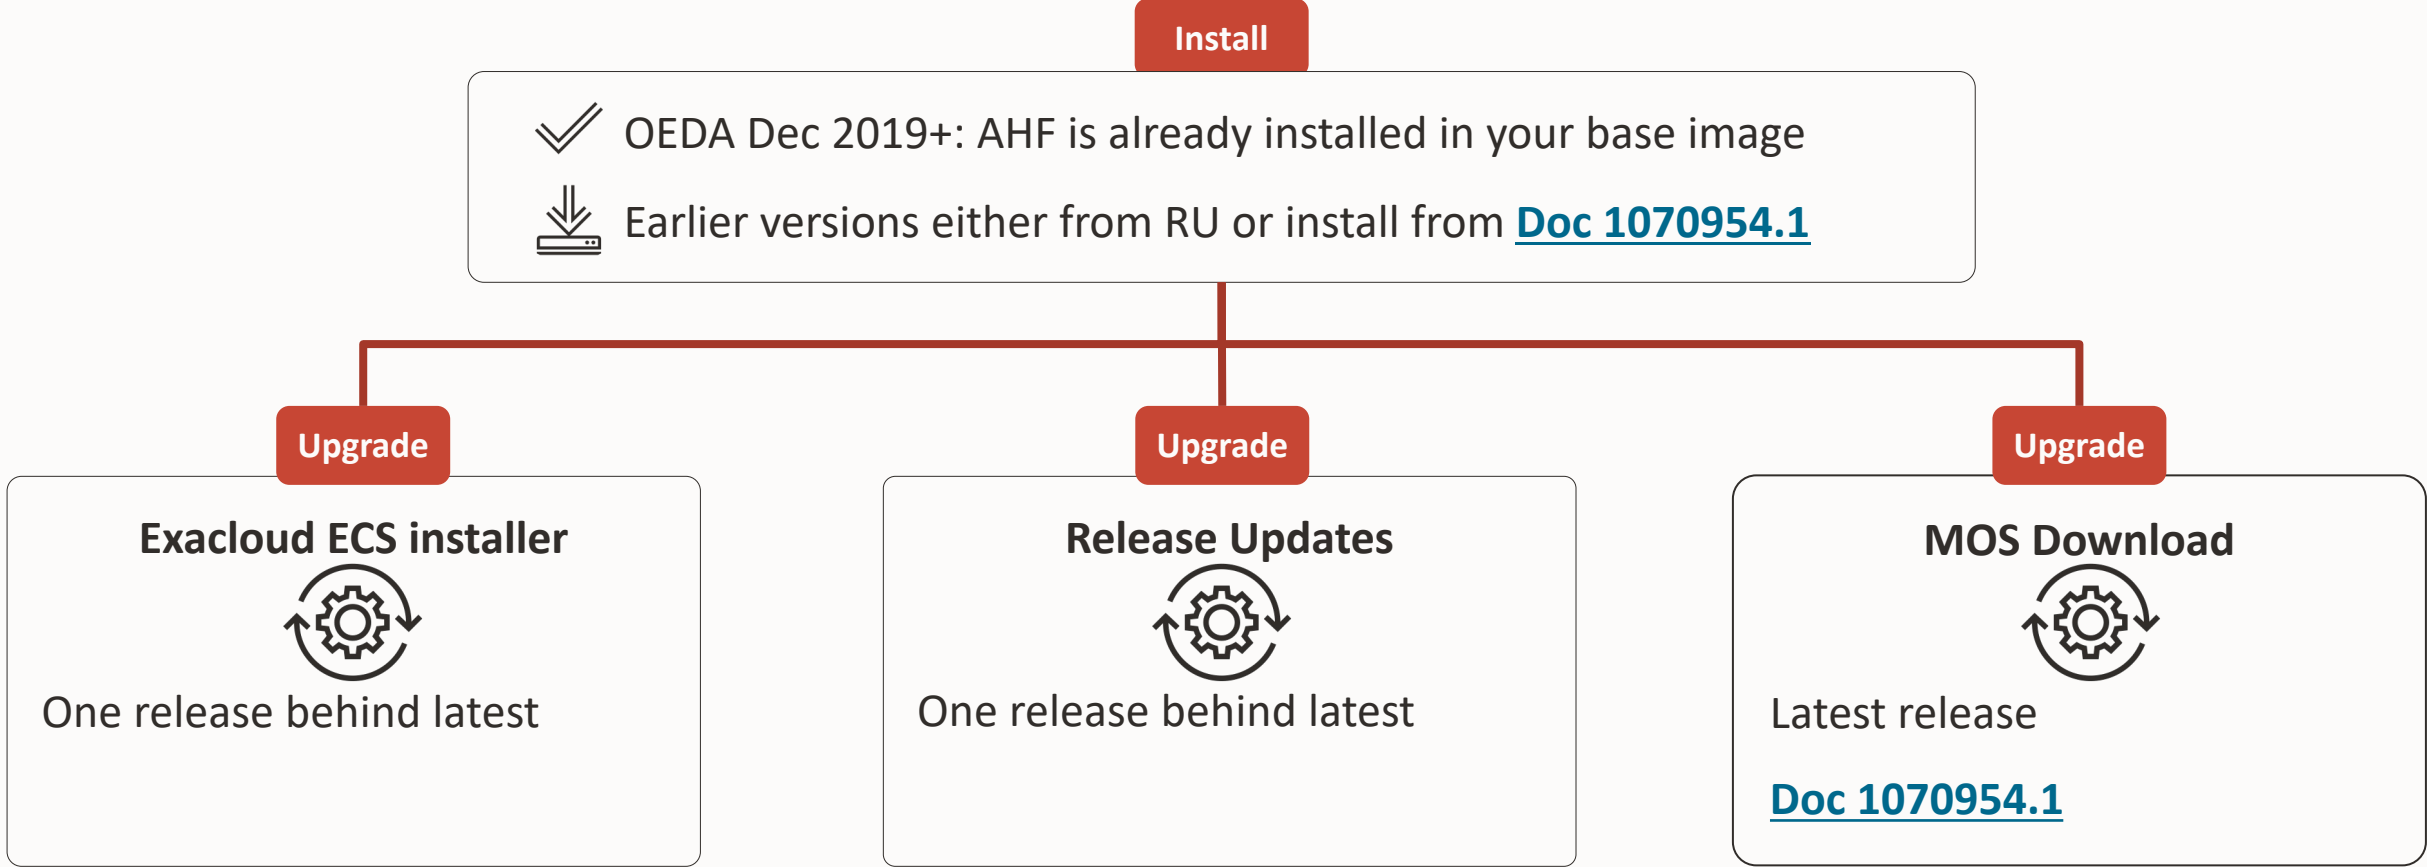
 <https://www.linkedin.com/in/raosandesh/>

 <https://www.slideshare.net/SandeshRao4>

# What is AHF



# Installation and staying up to date



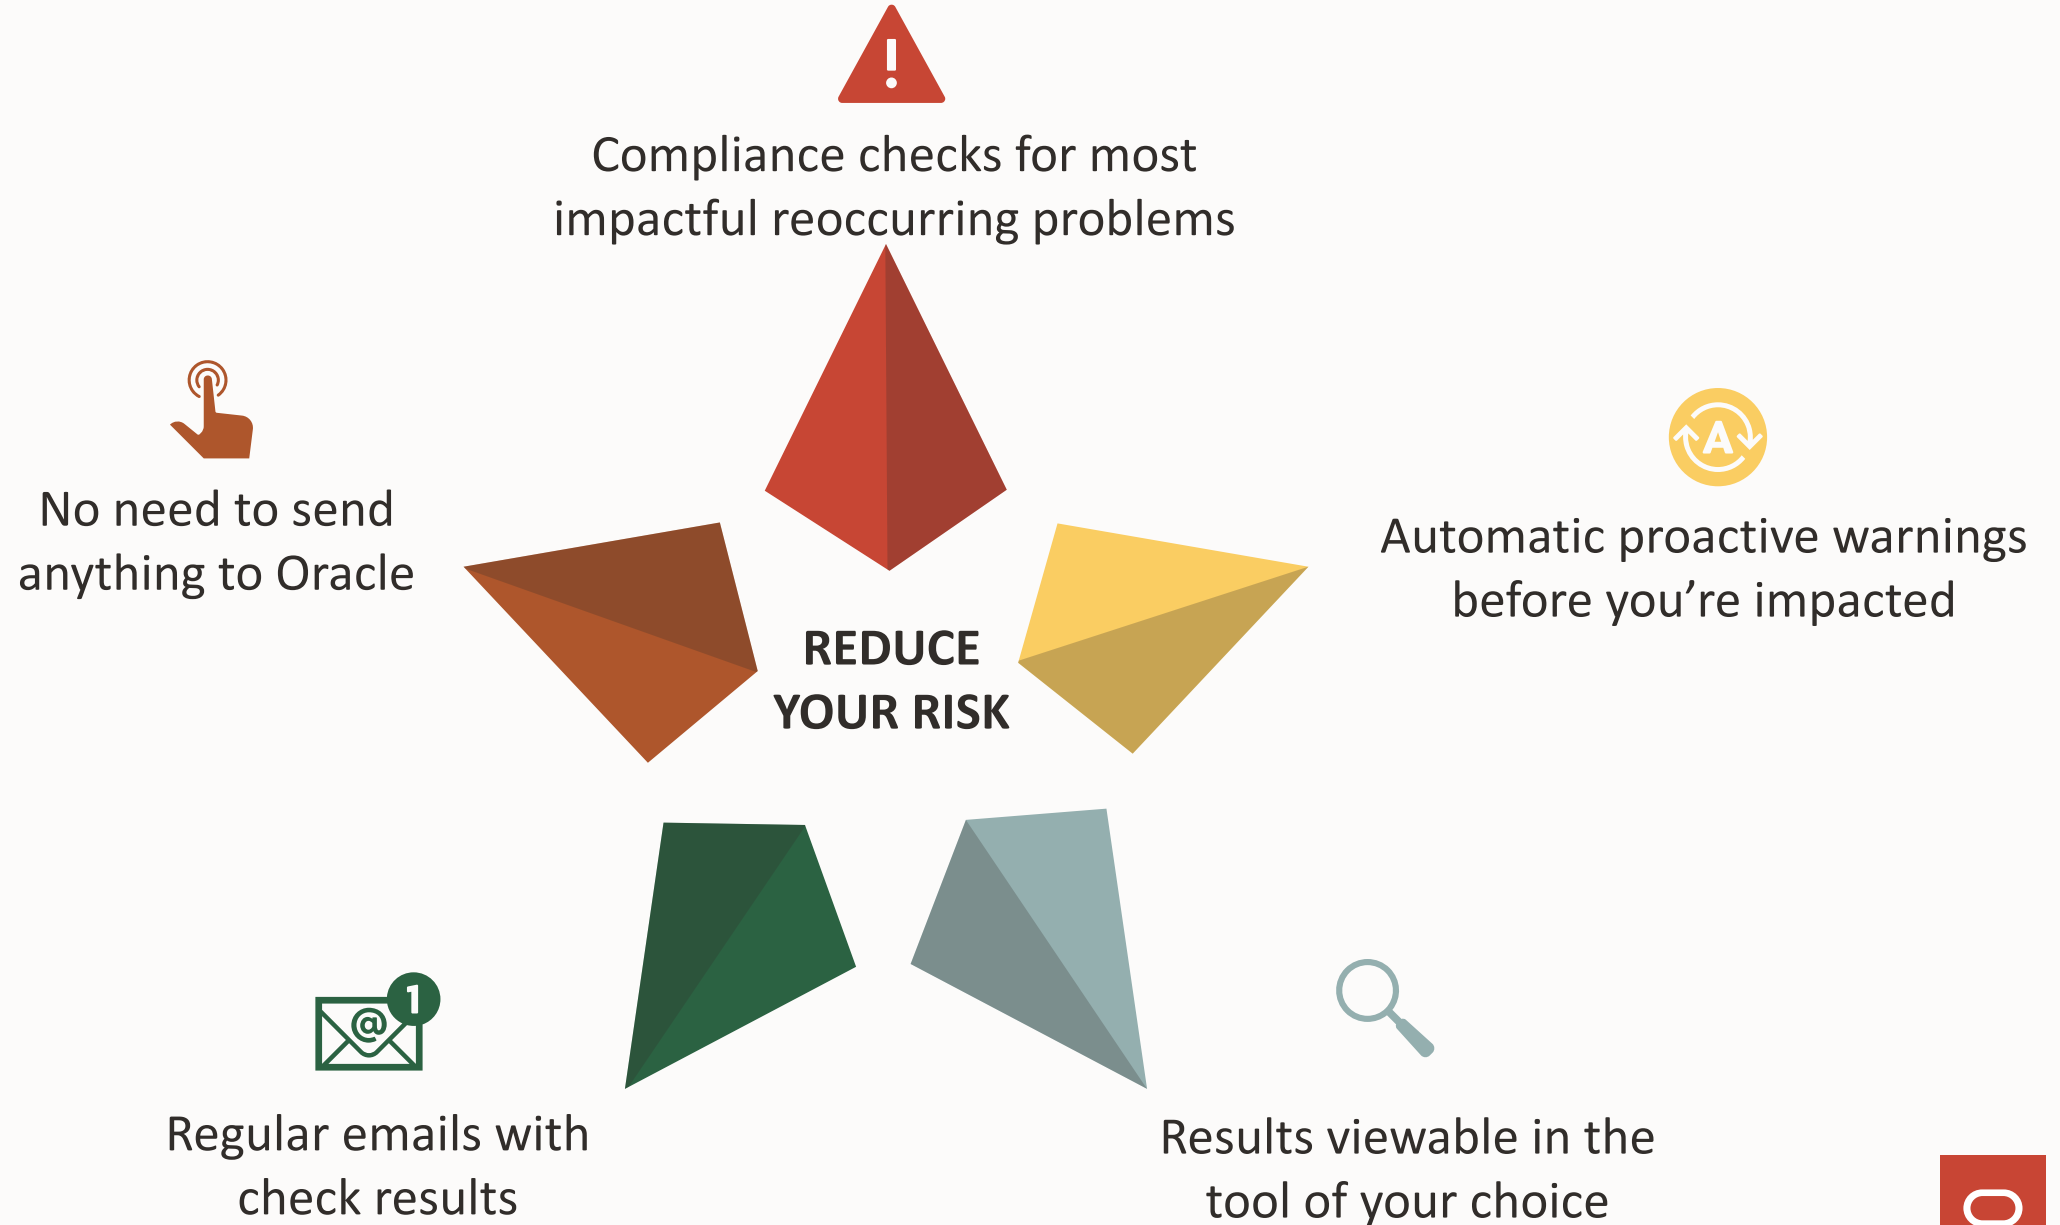
## Auto Upgrade

Automatically upgrade AHF if a new version is available at the software stage location

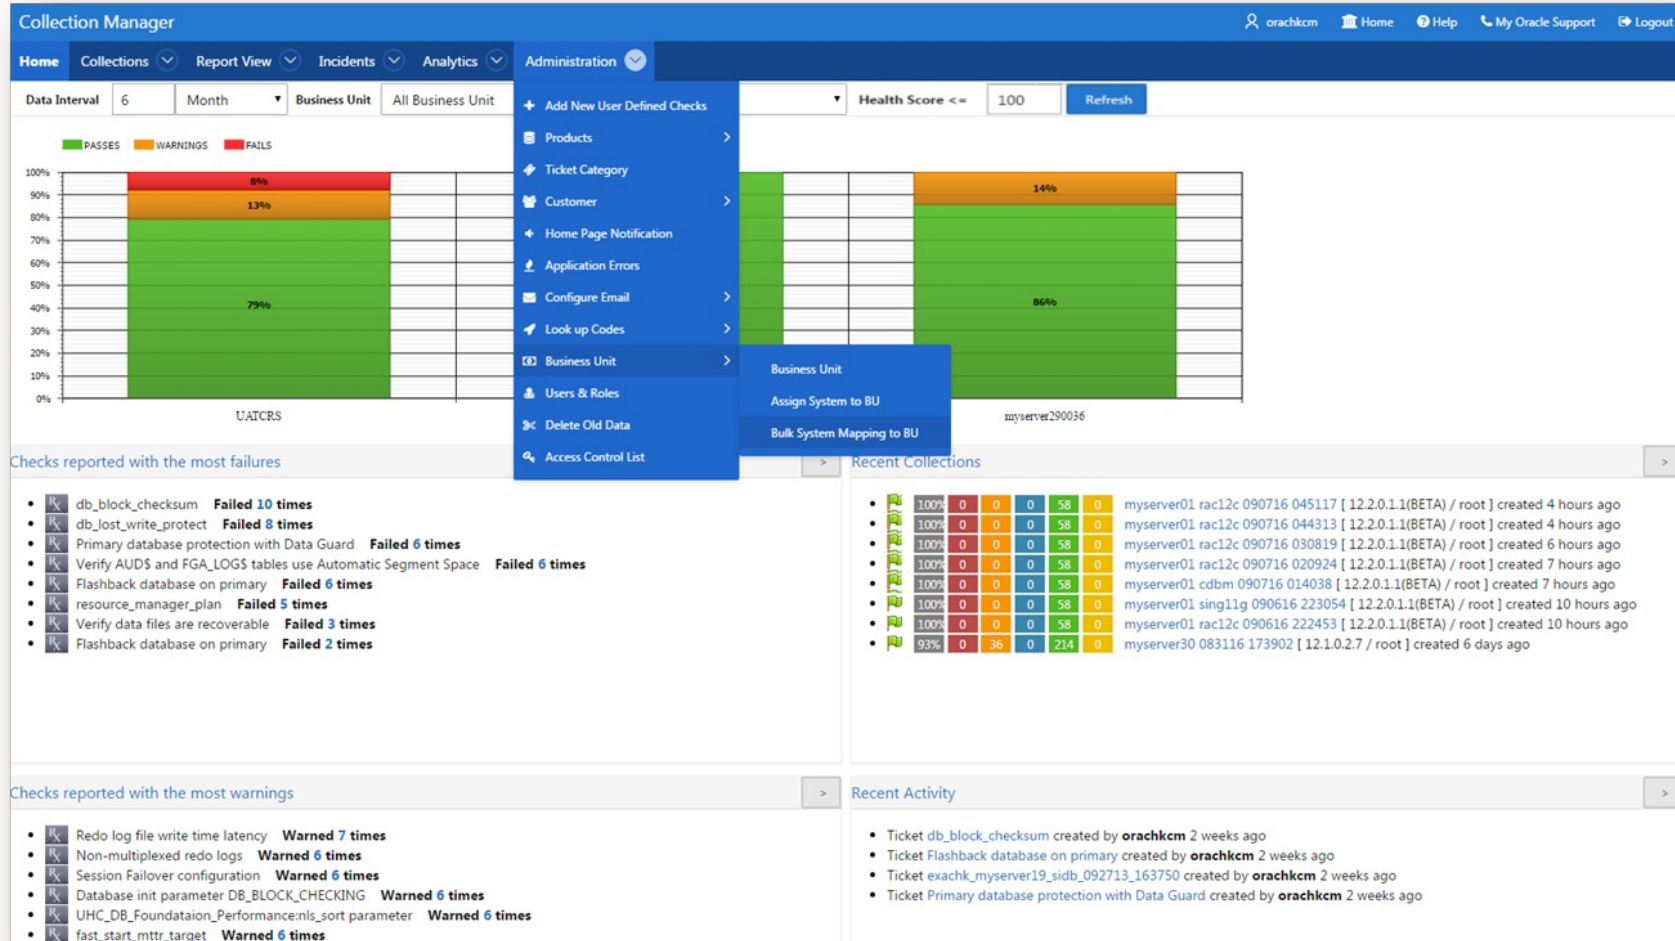
```
ahfctl setupgrade -autoupgrade <on/off> -swstage path
```

Example setting autoupgrade to check for a new version in stage location every 30 days

```
ahfctl setupgrade -all
Enter autoupgrade flag <on/off> : on
Enter software stage location : /scratch/ahf_stage
Enter auto upgrade frequency : 30
AHF autoupgrade parameters successfully updated
Successfully synced AHF configuration
```



# Oracle Health Check Collection Manager Dashboard




# The GIMR – Your Oracle Cluster Diagnostics Repository

## 1. PROS

- Stores Autonomous Health metrics for real-time and post-mortem analysis
  - Cluster Health Monitor (CHM)
  - Cluster Health Advisor (CHA)
  - DB QoS Management (QoSM)
- Default 72 hours of storage
- Minimized resource footprint
- Built-in Automatic Lifecycle management
- Automatic HA failover support
- No DBA management required

## 1. CONS

- Requires minimum 30GB of shared disk
- GI Patching and Upgrade integration requires significantly longer maintenance window



# Local Option

**Oracle Grid Infrastructure 21c Installer - Step 7 of 17**

**21<sup>c</sup> ORACLE**  
Grid Infrastructure

### Create GIMR Option

The Grid Infrastructure Management Repository(GIMR) is an essential component for complete operation of the Autonomous Health Framework, that offers enhanced real time diagnostics and performance management, and Fleet Patching and Provisioning. The components that depend on the repository in whole or in part are Cluster Health Advisor, Cluster Health Monitor, QoS Management, Fleet Patching and Provisioning and Cluster Activity Log. It is best practice to install this option and failure to do so could compromise timely resolution of issues as well as available functionality for patching.

Select one of the GIMR configuration options

- Use a Local GIMR database  
The GIMR database will have to be configured later in a separate RAC Database Oracle Home that is installed on all cluster nodes.
- Use an existing remote GIMR database  
Specify a credential file:
- Do Not use a GIMR database



# How To Install a Local 21c GIMR in 3 Steps

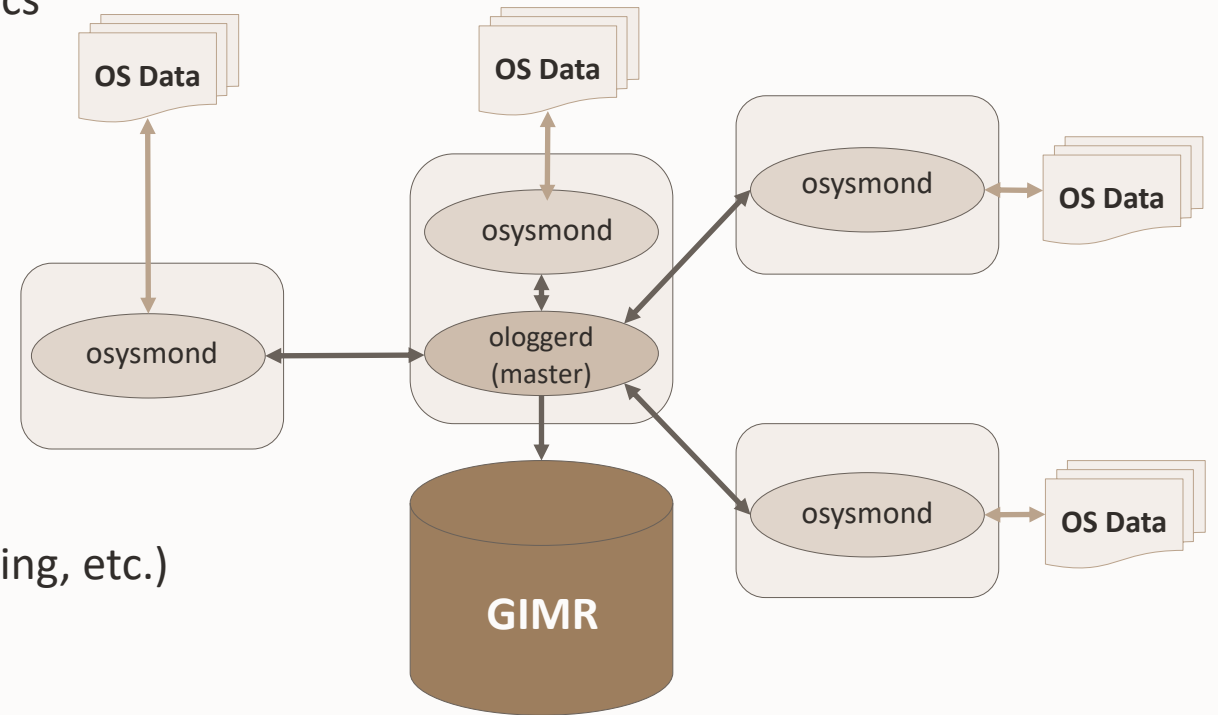
- 1. Install the Oracle 21c Grid Infrastructure with Default GIMR Option.**
  - If using ASM, create a disk group for the GIMR (ex: MGMT)
- 2. Install an Oracle 21c Database Home in a separate directory as the GI User.**
  - Install on all nodes as you would an Oracle RAC database.
- 3. Create the GIMR Database**
  - `OH/bin/mgmtca createGIMRContainer [-storageDiskLocation disk_location]`



# AHF Component - Cluster Health Monitor (CHM)

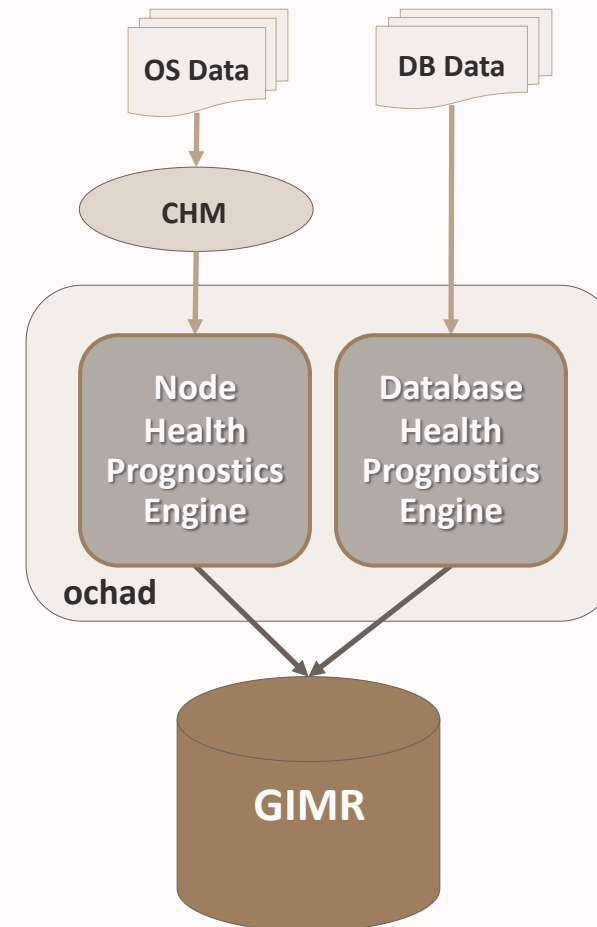
Generates view of Cluster and Database diagnostic metrics

- Always on - Enabled by default
- Provides Detailed OS Resource Metrics
- Assists Node eviction analysis
- Locally logs all process data
- User can define pinned processes
- Listens to CSS and GIPC events
- Categorizes processes by type
- Supports plug-in collectors (ex. traceroute, netstat, ping, etc.)
- New CSV output for ease of analysis



## AHF Component - Cluster Health Advisor (CHA)

- Always on - Enabled by default
- Detects node and database performance problems
- Provides early-warning alerts and corrective action
- Supports on-site calibration to improve sensitivity
- Integrated into EMCC Incident Manager and notifications



# Oracle Database Alert Log

---

## PRO

- Destination for Important DB Events
- Single file to monitor by DBAs
- Many tools available to parse
- Supported by TFA for generating alarms

## CONS

- Includes both critical and non-critical events
- Includes messages not intended for DBAs
- Inconsistently reports severity level
- Can report unintuitive cause and action
- New undocumented messages in every release

## The Curated Solution - New 21c Attention Log

---

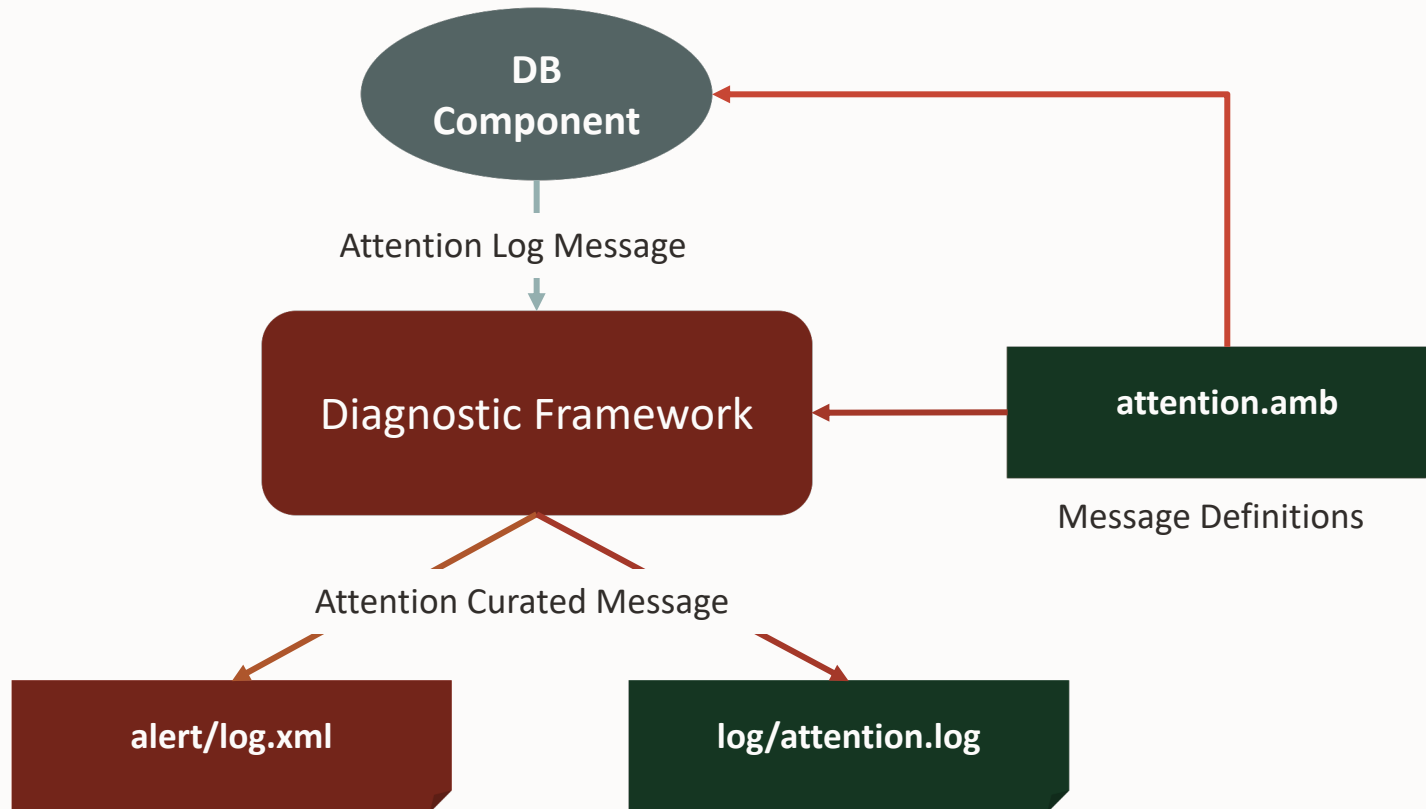
Contains only important events requiring customer attention

Includes documented set of messages and attributes

All Messages include these attributes:

- Type
- Urgency
- Scope
- Target User
- Cause and Action
- Additional debug information

# Oracle Database Attention Log Message Flow



# Attention Log Curation - Message Attributes

TYPE	SCOPE	TARGET USER
<ol style="list-style-type: none"><li>1. Error</li><li>2. Warning</li><li>3. Notification</li></ol>	<ol style="list-style-type: none"><li>1. Session</li><li>2. Process</li><li>3. PDB-Instance</li><li>4. CDB-Instance</li><li>5. CDB-Cluster</li><li>6. PDB-Persistent</li><li>7. CDB-Persistent</li></ol>	<ol style="list-style-type: none"><li>1. App-Dev</li><li>2. Sec-Admin</li><li>3. Net-Admin</li><li>4. Cluster-Admin</li><li>5. PDB-Admin</li><li>6. CDB-Admin</li><li>7. Server-Admin</li><li>8. Storage-Admin</li><li>9. DataOps-Admin</li></ol>
URGENCY		
<ol style="list-style-type: none"><li>1. Immediate</li><li>2. Soon</li><li>3. Deferable</li><li>4. Info</li></ol>		

# Database Health - Applied Machine Learning

Discovers Potential Cluster & DB Problems

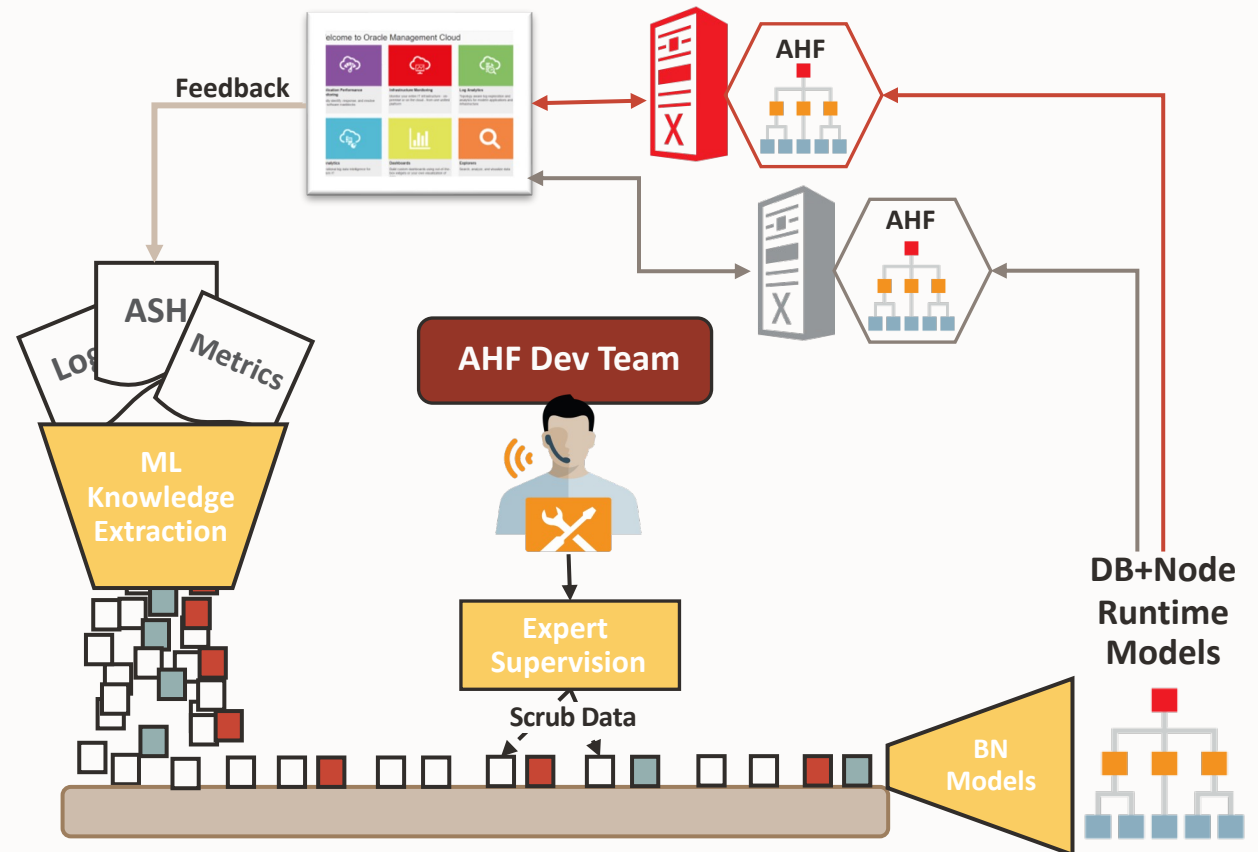
Actual Internal data drives model development

Applied purpose-built Applied ML for knowledge extraction

Expert Dev team scrubs data

Generates Bayesian Network-based diagnostic root-cause models

Uses BN-based run-time models to perform real-time prognostics

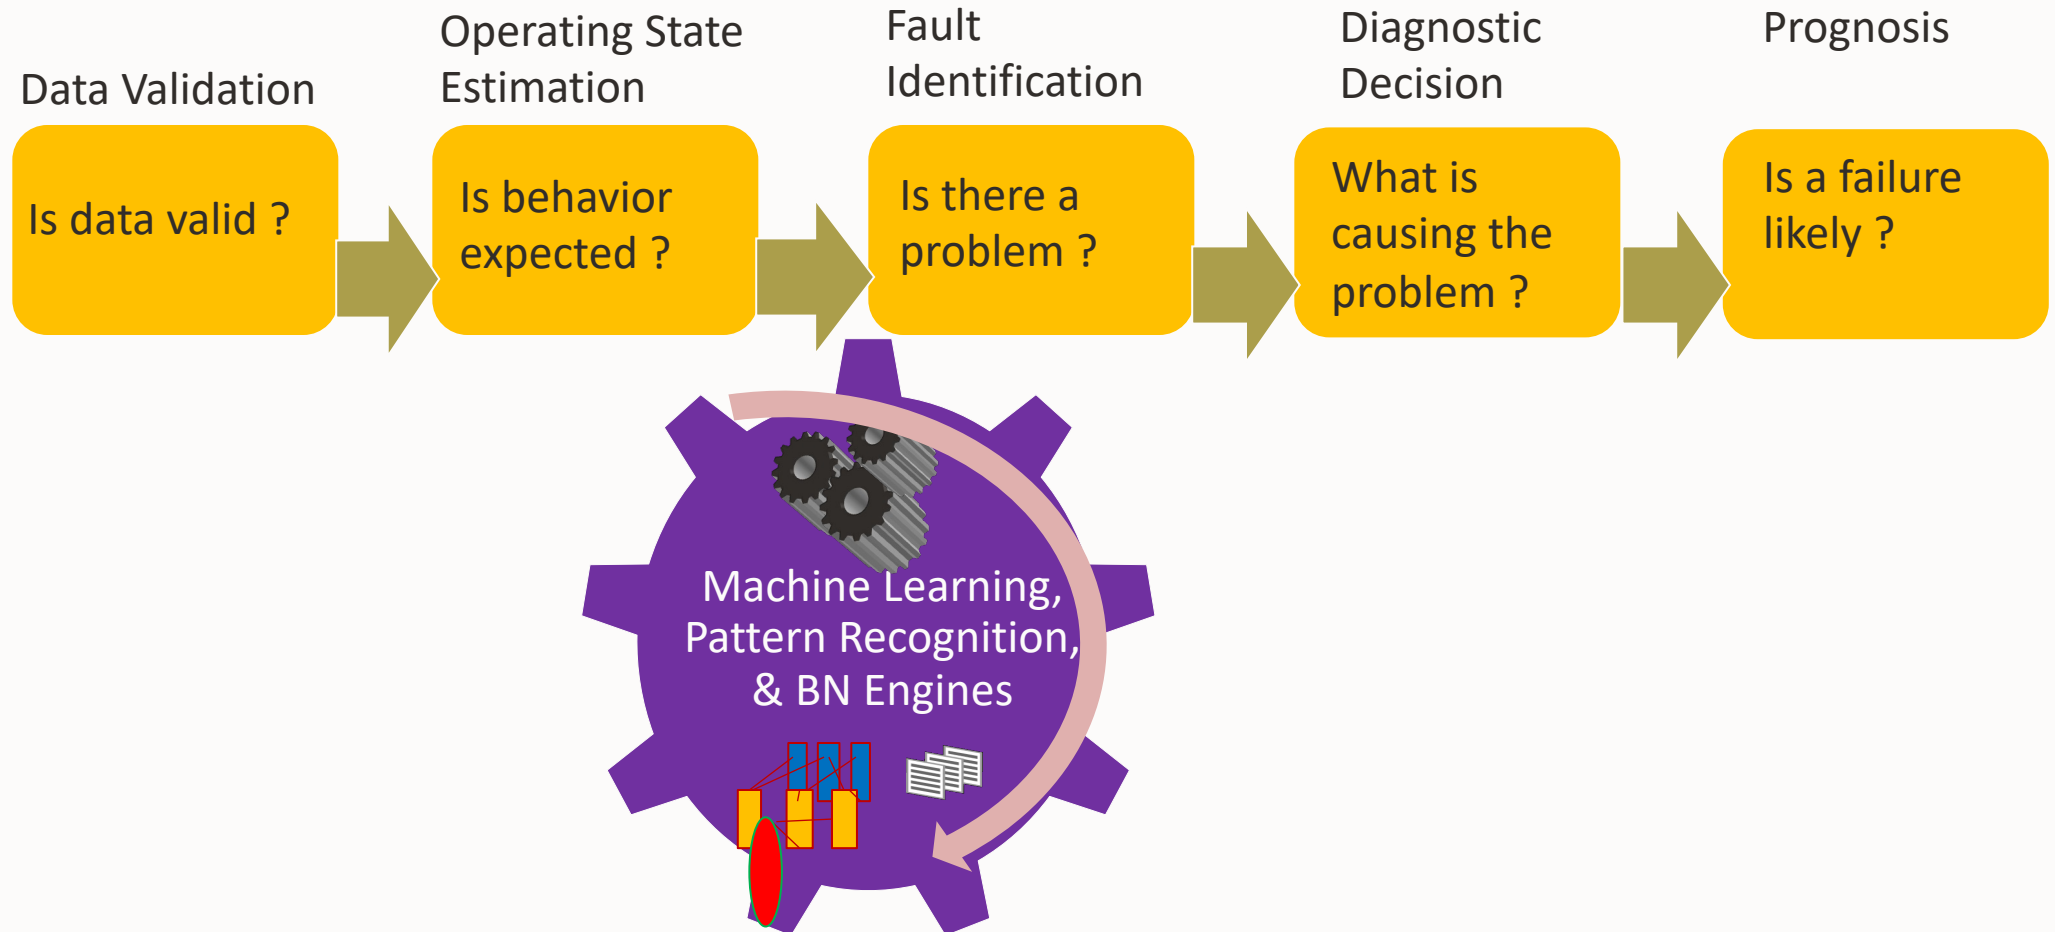




# Cluster Health Advisor

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

*For each data point ...*



## AHF and Machine Learning

- ***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***

# Step by Step on how to perform an autopsy of a problem

- ORAchk/EXAchk provides a **single source for all upgrade checks**
  - ORAchk , EXAchk checks , Database AutoUpgrade checks , Cluster Verification Utility (CVU) checks
- AHF proactive findings
- Tail Alert and Attention logs and build the timeline
- Dig deeper – AHF to either collect data , analyze
- Performance Issue
- OS or Network Issue
- Node Evictions ?
- AHF Insights and AHF Scope

## ORAchk/EXAchk provides a **single source for all upgrade checks**

- To check an environment before upgrading run:

```
orachk -preupgrade
```

- To check an environment after upgrade run:

```
orachk -postupgrade
```

- To check an environment for best practice violations

```
orachk
```

- Setup Collection Manager to aggregate the results for a birds eye view
  - **Autonomous Health Framework (AHF) - Including TFA and ORAchk/EXAchk (Doc ID 2550798.1)**

# Calibrating AHF to your RAC deployment

Choosing a Data Set for Calibration – Defining “normal”

```
chactl query calibration -cluster -timeranges 'start=2021-12-01 07:00:00,end=2021-12-01
13:00:00'
Cluster name : mycluster
Start time : 2021-12-01 07:00:00
End time : 2021-12-01 13:00:00
Total Samples : 11524
Percentage of filtered data : 100%

1) Disk read (ASM) (Mbyte/sec)
    MEAN      MEDIAN      STDDEV      MIN      MAX
    0.11      0.00      2.62      0.00      114.66
    <25      <50      <75      <100      >=100
    99.87%   0.08%   0.00%   0.02%   0.03%

...

```



# Calibrating AHF to your RAC deployment

Choosing a Data Set for Calibration – Defining “normal”

```
...
2) Disk write (ASM) (Mbyte/sec)
   MEAN      MEDIAN    STDDEV    MIN      MAX
   0.01      0.00      0.15     0.00     6.77
   <50       <100      <150     <200     >=200
   100.00%   0.00%    0.00%    0.00%    0.00%
...
```



# Calibrating AHF to your RAC deployment

Choosing a Data Set for Calibration – Defining “normal”

```
...
3) Disk throughput (ASM) (IO/sec)
    MEAN      MEDIAN    STDDEV    MIN      MAX
    2.20      0.00      31.17    0.00     1100.00
    <5000     <10000    <15000    <20000   >=20000
    100.00%   0.00%    0.00%    0.00%    0.00%

4) CPU utilization (total) (%)
    MEAN      MEDIAN    STDDEV    MIN      MAX
    9.62      9.30      7.95     1.80     77.90
    <20       <40       <60       <80      >=80
    92.67%   6.17%    1.11%    0.05%    0.00%
...

```



# Calibrating AHF to your RAC deployment

Create and store a new model

```
chactl query calibrate cluster -model daytime -timeranges 'start=2021-12-01 07:00:00, end= 2021-12-01 13:00:00'
```

Begin using the new model

```
chactl monitor cluster -model daytime
```

Confirm the new model is working

```
chactl status -verbose
```

```
monitoring nodes svr01, svr02 using model daytime  
monitoring database oltpacdb, instances oltpacdb_1, oltpacdb_2 using model DEFAULT_DB
```





# Command line operations

Check for Health Issues and Corrective Actions with CHACTL QUERY DIAGNOSIS

```
chactl query diagnosis -db oltpacdb -start "2021-12-01 01:42:50" -end "2021-12-01 03:19:15"
2021-12-01 01:47:10.0 Database oltpacdb DB Control File IO Performance (oltpacdb_1) [detected]
2021-12-01 01:47:10.0 Database oltpacdb DB Control File IO Performance (oltpacdb_2) [detected]
2021-12-01 02:59:35.0 Database oltpacdb DB Log File Switch (oltpacdb_1) [detected]
2021-12-01 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
2016-07-03 01:49:30.0	Host svr02	<a href="#">Host CPU Utilization</a>	detected
2016-07-03 01:49:50.0	Host svr01	<a href="#">Host CPU Utilization</a>	detected
2016-07-03 05:54:55.0	Host svr01	<a href="#">Host Memory Consumption</a>	detected
2016-07-04 03:40:00.0	Host svr02	<a href="#">Host CPU Utilization</a>	cleared
2016-07-04 03:40:05.0	Host svr01	<a href="#">Host CPU Utilization</a>	cleared
2016-07-04 03:40:05.0	Host svr01	<a href="#">Host Memory Consumption</a>	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 applications not connected to a database on this node.	Identify the top memory consumers by using the Cluster Health Monitor (CHM).



# Autonomous Health – Database Performance

## Data Sources and Data Points

A *Data Point* contains > 150 signals (statistics and events) from *multiple sources*

OS, ASM , Network → | ← DB ( SH, 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	

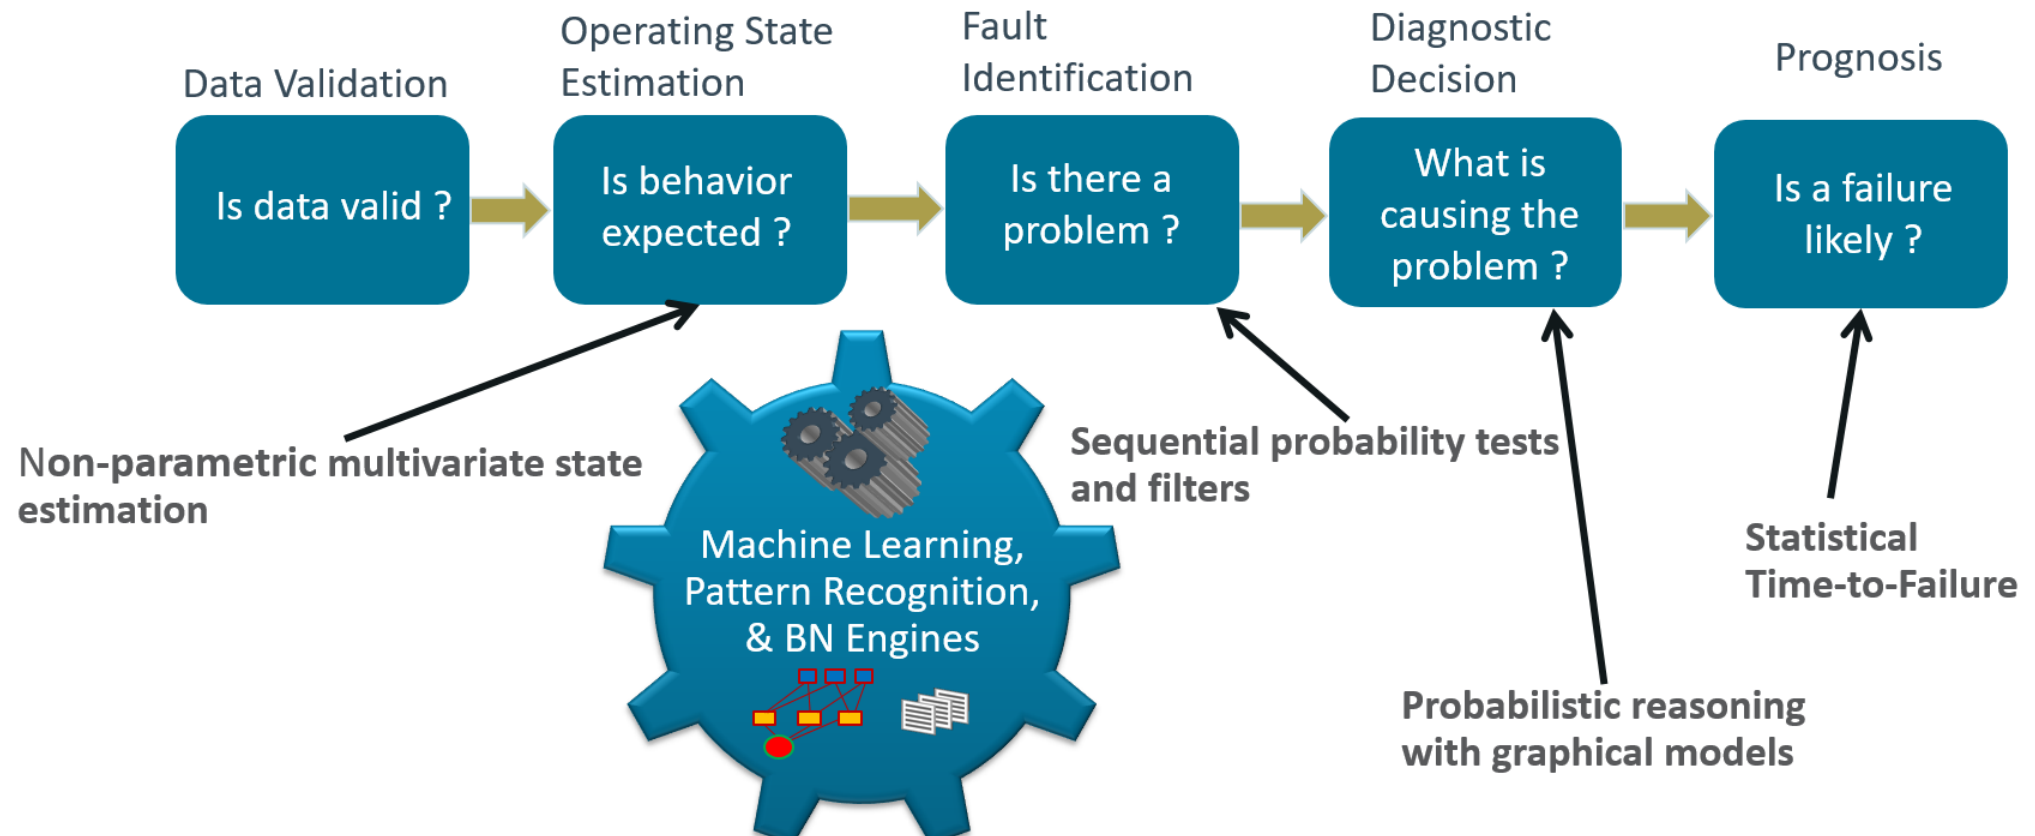
Statistics are collected at a **1 second internal sampling** rate , synchronized, smoothed and aggregated to a Data Point **every 5 seconds**



# Autonomous Health – Database Performance

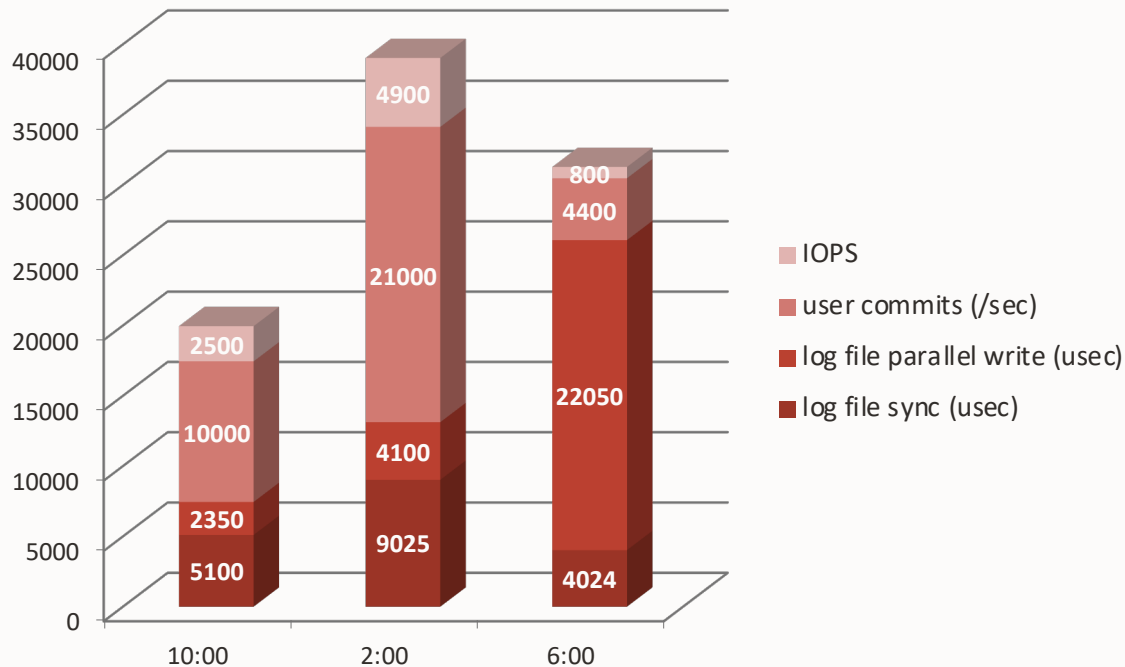
## Data Flow Overview

*For each data point ...*



# Models Capture all Normal Operating Modes

Models Capture the Dynamic Behavior of all Normal Operation



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

IOPS	####	2500	4900	800	####
User Commits	####	10000	21000	4400	####
Log File Parallel Write	####	2350	4100	22050	####
Log File Sync	####	5100	9025	4024	####
...	...	...	...	...	...

A model captures *the normal load phases* and their statistics over time , and thus the characteristics for all load intensities and profiles . During monitoring , *any data point similar* to one of the vectors is NORMAL. One could say that *the model REMEMBERS the normal operational dynamics over time*



# Autonomous Health – Database Performance

*Inline and Immediate Fault Detection and Diagnostic Inference*

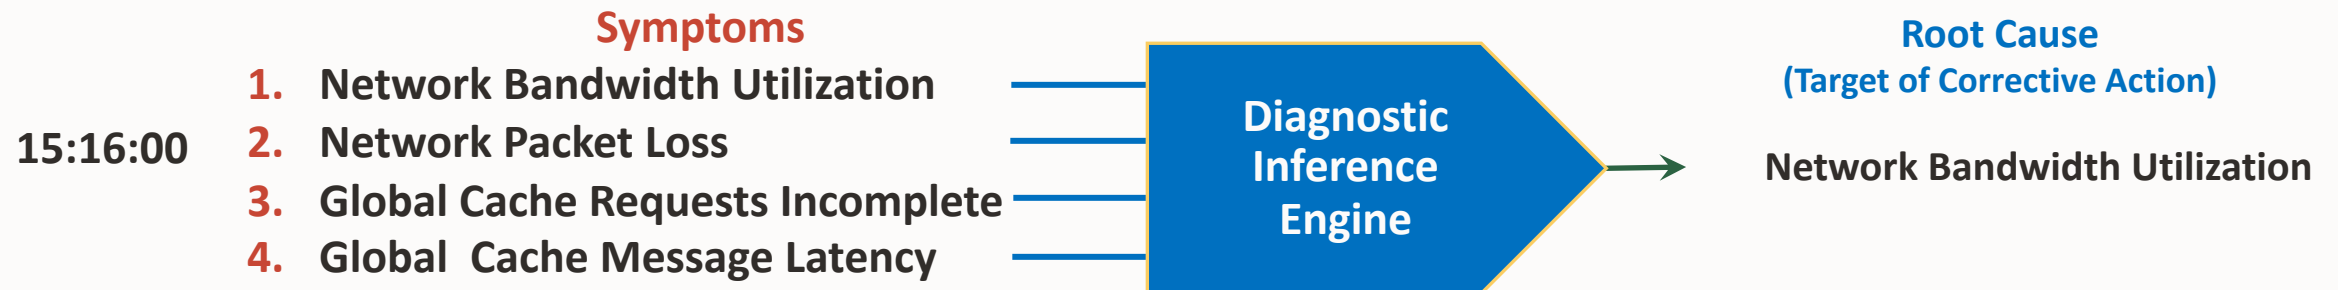
Input : Data Point at Time *t*

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%	105	2 ms	600 us	504 ms	513 ms	2 ms	5.9 ms	0	

Fault Detection and Classification

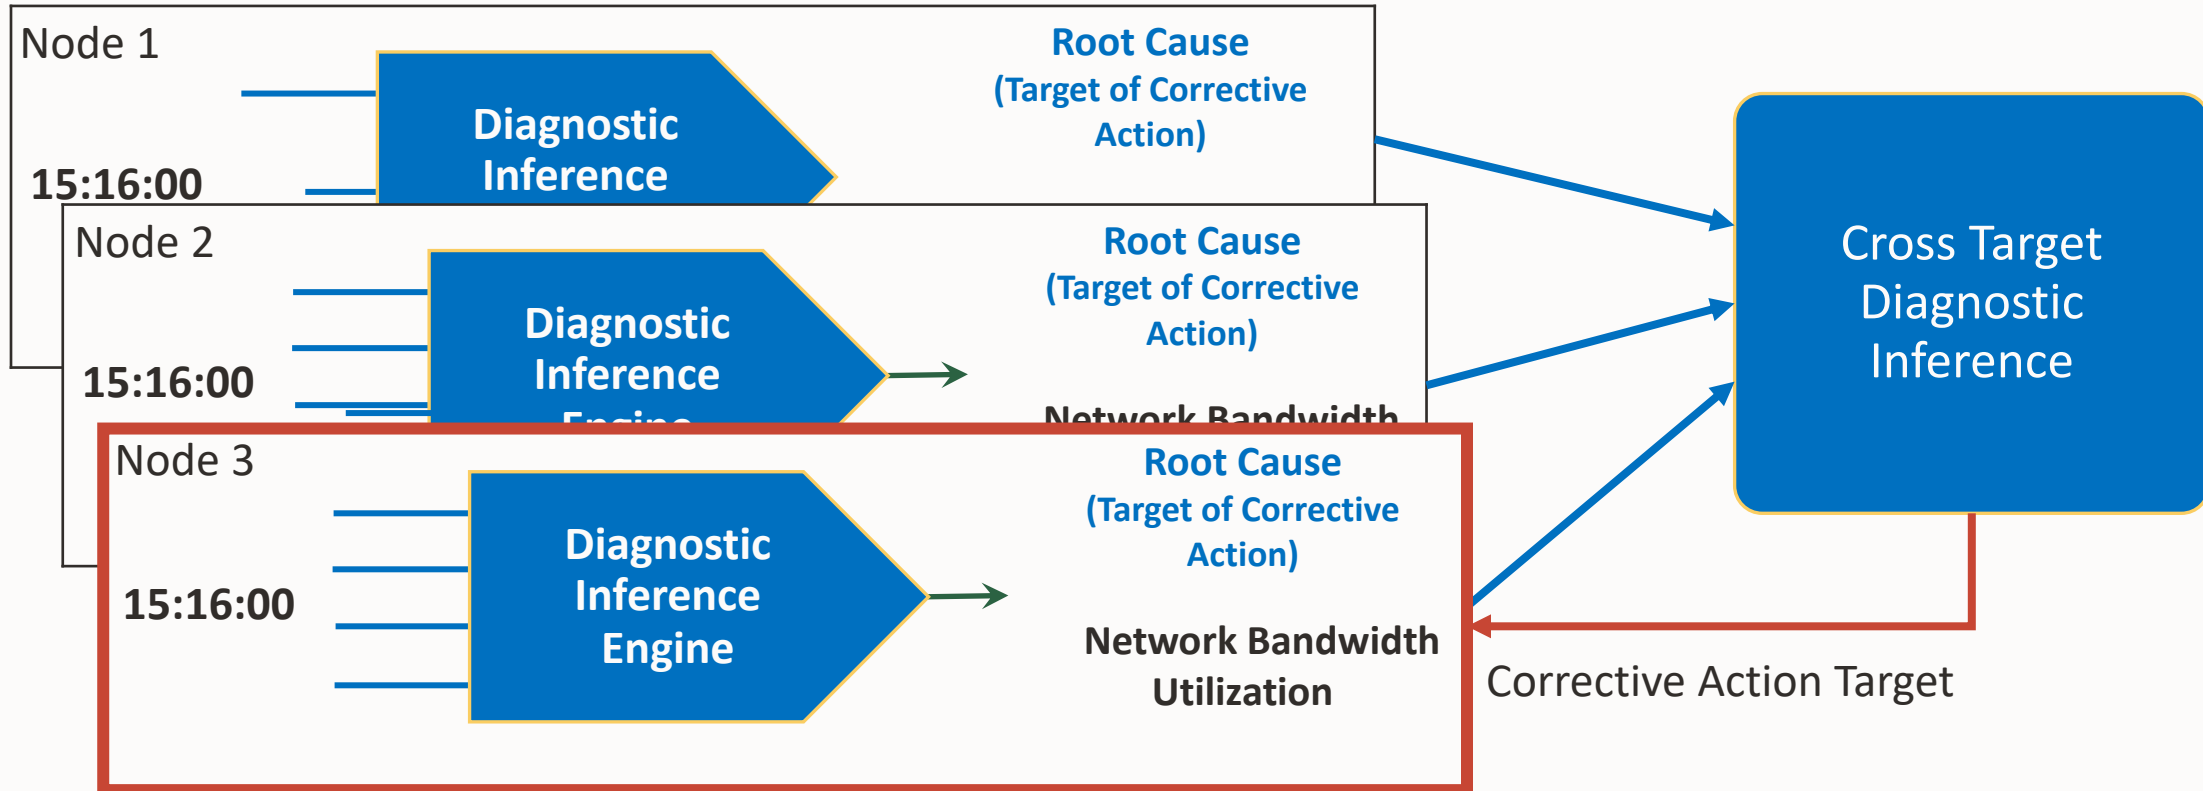
15:16:00	OK	OK	HIGH 1	HIGH 2	OK	OK	HIGH 3	HIGH 3	HIGH 4	HIGH 4	OK	
----------	----	----	--------	--------	----	----	--------	--------	--------	--------	----	--

Diagnostic Inference

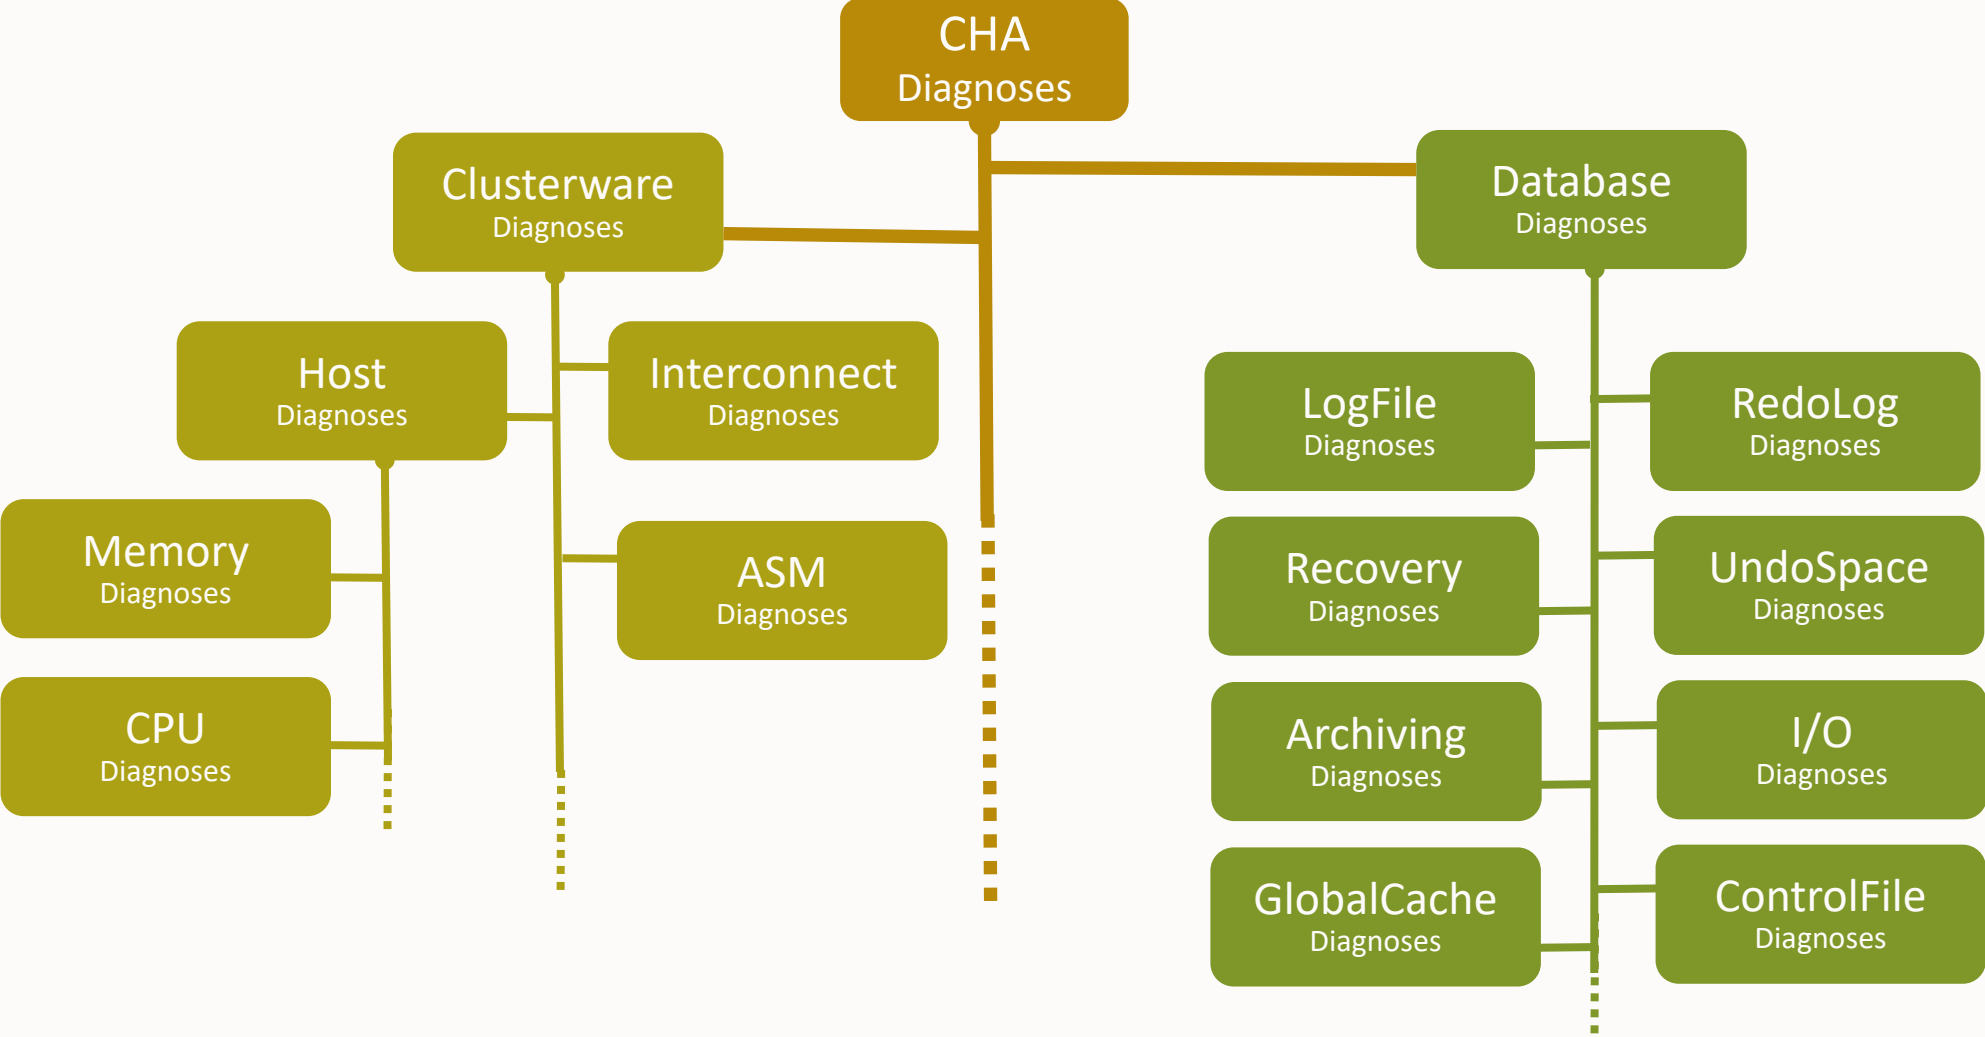


# Autonomous Health - Cluster Health Advisor

## Cross Node and Cross Instance Diagnostic Inference



# 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_MSG	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_CONGESTION	CHA detects global cache messages on the private interconnect are lost	Hang Instance Eviction

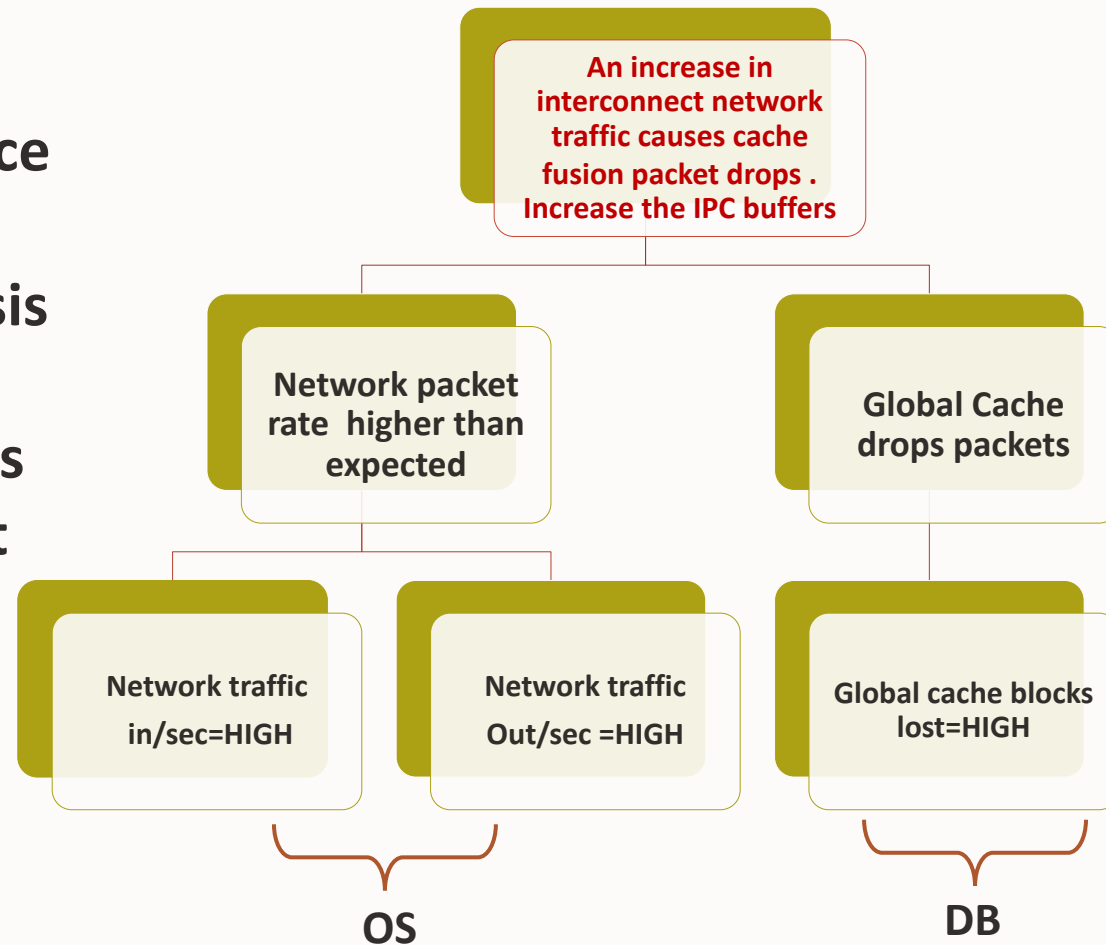


# AHF Diagnostics Model

AHF 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



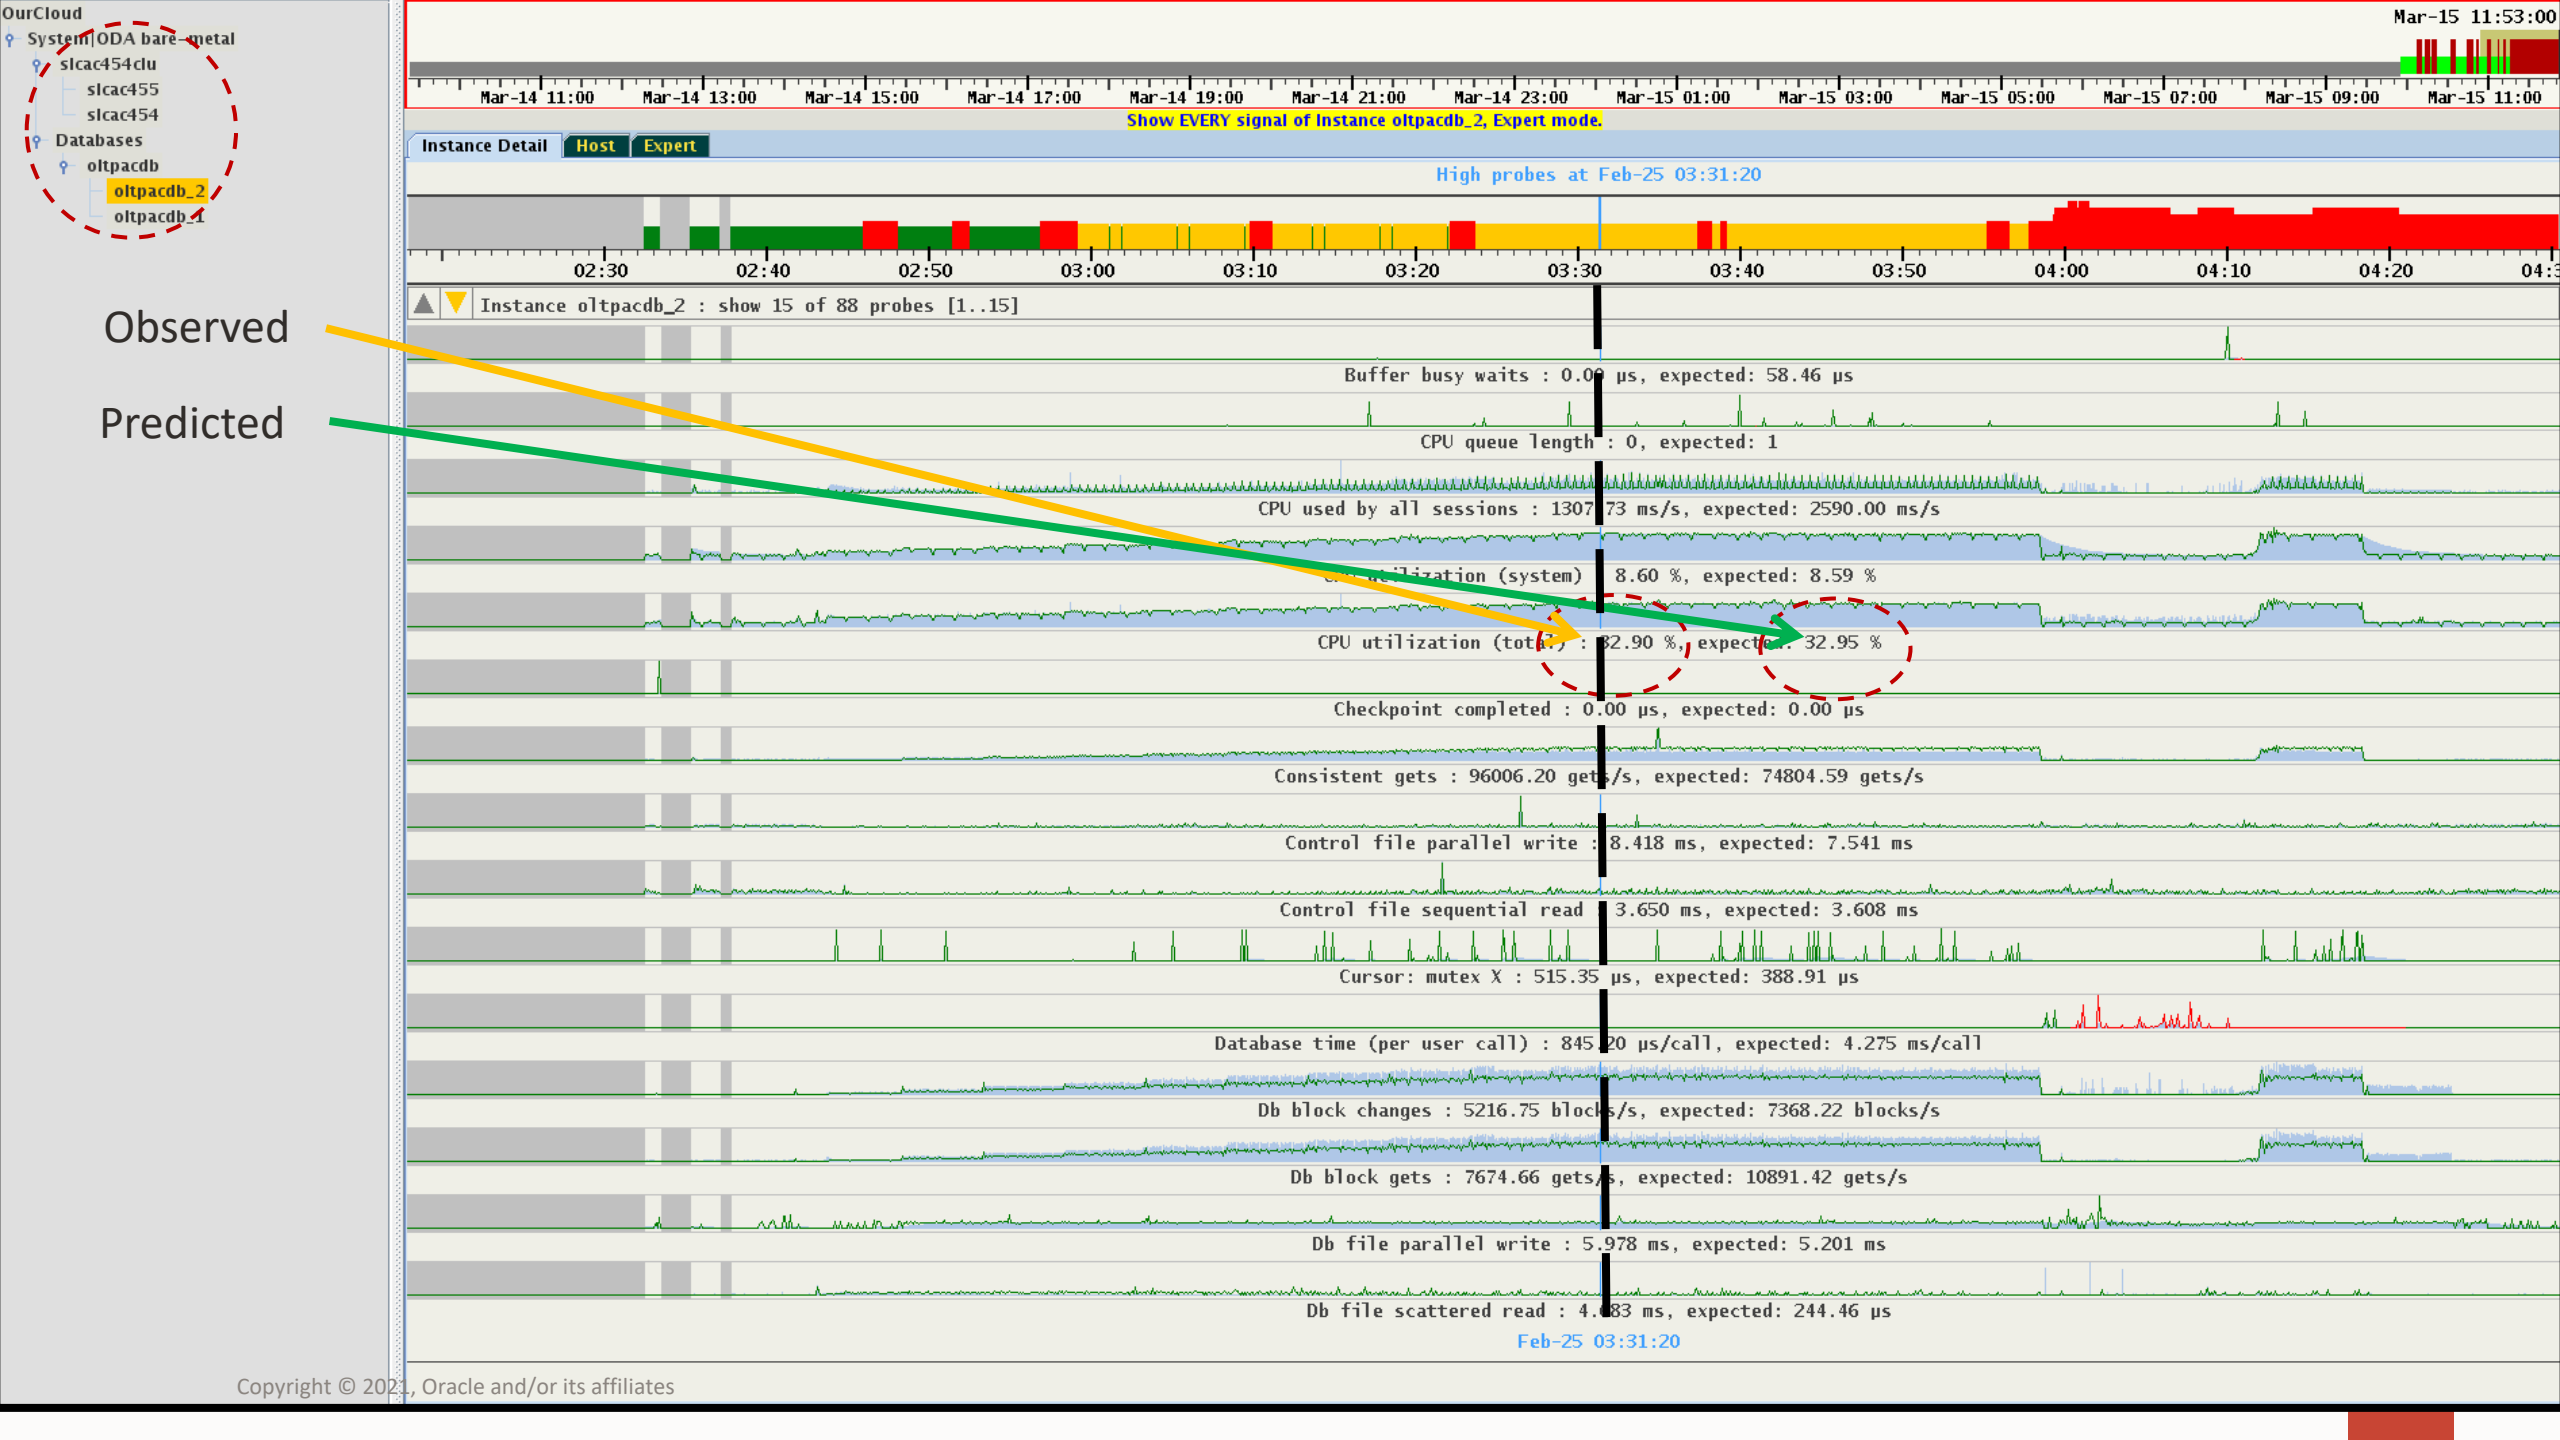
# What's AHF Scope (new in 22.3)

The screenshot shows the Oracle Cloud AHF (Autonomous Health Framework) interface. On the left, a tree view shows the hierarchy: OurCloud > slcac454clu > Hosts > slcac454 > Databases > oltpacdb > oltpacdb\_1. The main area displays a timeline from Nov-08 15:00 to Nov-09 06:00 with a red and green bar chart. Below the chart, it says "Instance oltpacdb\_1, Expert mode." and has tabs for "Instance Detail", "Host", and "Expert". A control bar includes "Showing every probe", "Save as", "No saved selections", "Load", and "Delete". A "Reset" button is also present. A list of component groups is shown with checkboxes:


- 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**






- 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



CHA (on SLCAC455.US.ORACLE.COM) V0.76.4, Data V0.20 -- OTN Version

Nov-15 08:34:00

Showing abnormal probes of Instance oltpacdb\_2 Probe Selection Active

Instance Detail Host

Alarm at Nov-15 06:54:20

06:50 07:00 07:10 07:20 07:30 07:40 07:50 08:00 08:10 08:20 08:30

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

## The Curated Solution - New 21c Attention Log

---

Contains only important events requiring customer attention

Includes documented set of messages and attributes

All Messages include these attributes:

- Type
- Urgency
- Scope
- Target User
- Cause and Action
- Additional debug information

## Example Attention Message Definition – CDB Warning

```
// TYPE          - 1 error, 2 warning, 3 notification
// URGENCY       - 1 immediate, 2 soon, 3 deferable, 4 info
// SCOPE         - 1 session, 2 process, 3 pdb-instance, 4 cdb-instance, 5 cdb-cluster, 6 pdb-
persistent, 7 cdb-persistent
// TARGETUSER    - 1 app-dev, 2 sec-admin, 3 net-admin, 4 cluster-admin, 5 pdb-admin, 6 cdb-
admin, 7 server-admin, 8 storage-admin, 9 dataops-admin

ID::2000
TYPE::2
URGENCY::1
SCOPE::4
TARGETUSER::6
TEXT::Parameter %s specified is high
CAUSE::Memory parameter specified for this instance is high
ACTION::Check alert log or trace file for more information relating to instance configuration,
reconfigure the parameter and restart the instance
STARTVERSION::21.1
```



## Example Attention Log Curated Message – CDB Warning

```
[
IMMEDIATE      Parameter SGA_MAX_SIZE specified is high
CAUSE:         Memory parameter specified for this instance is high
ACTION:        Check alert log or trace file for more information relating to instance
                configuration, reconfigure the parameter and restart the instance
CLASS:         CDB Instance / CDB ADMINISTRATOR / WARNING / AL-2000
TIME:         2020-05-01T11:09:02.223-07:00

ADDITIONAL INFO: -
WARNING: SGA_MAX_SIZE (6144 MB) is too high - it should be
         less than 5634 MB (80 percent of physical memory).
]
```

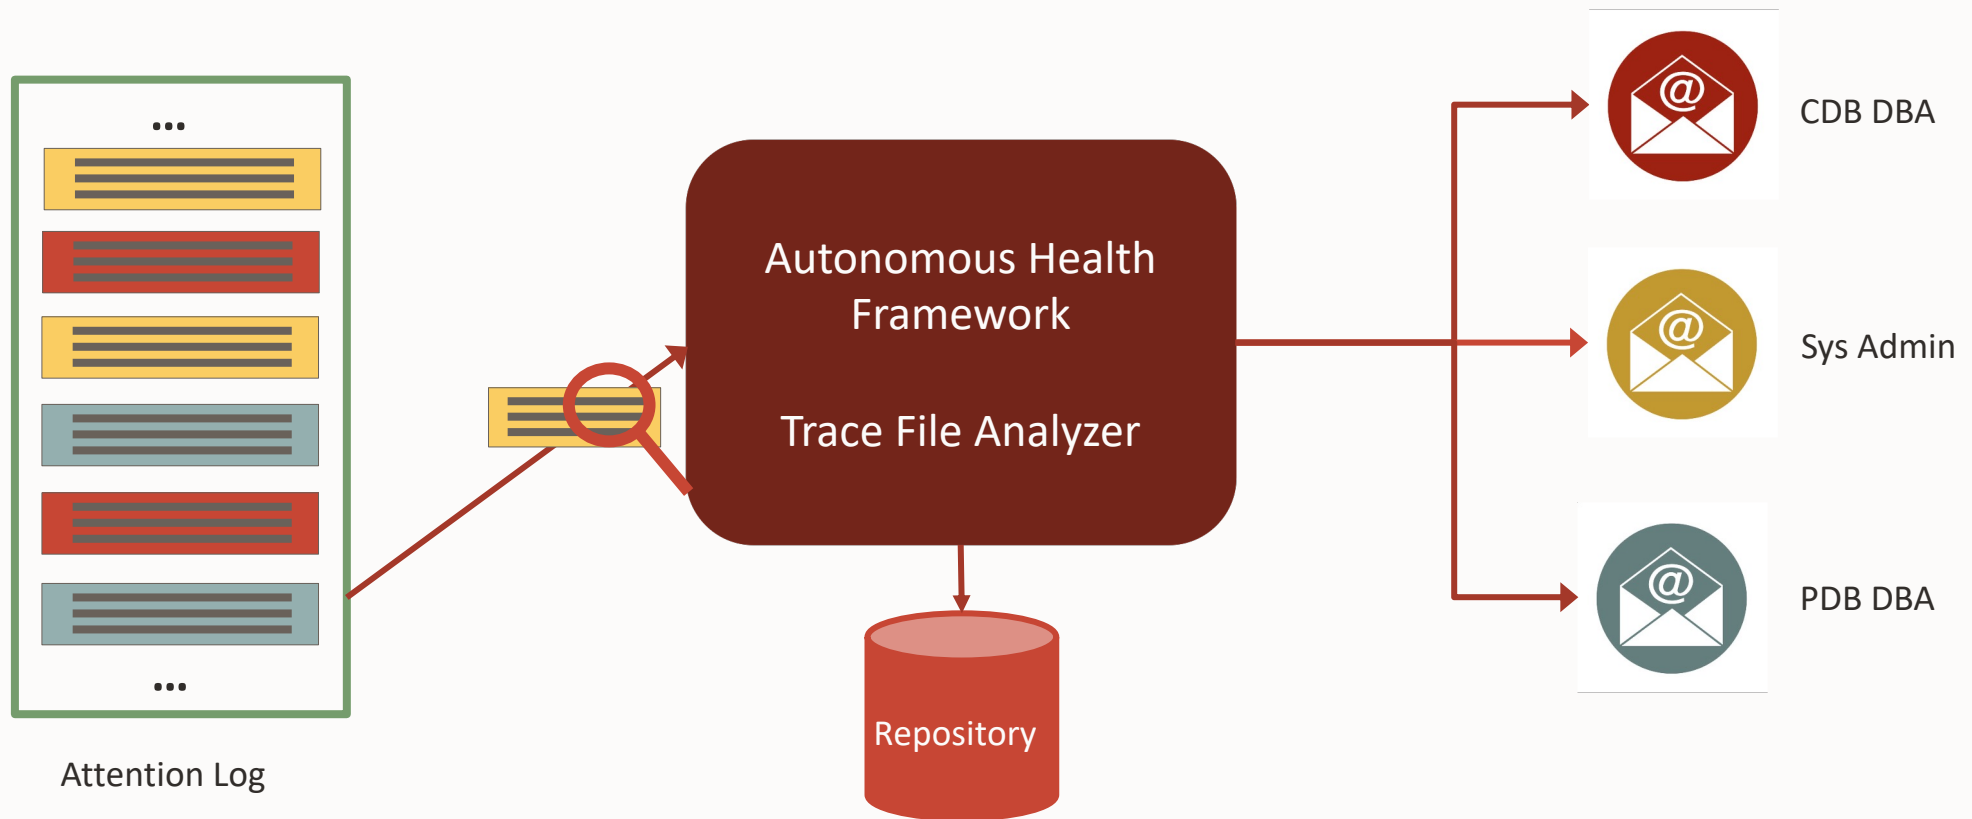
## Example Attention Log Curated Message – CDB Error

```
[  
  IMMEDIATE Shutting down ORACLE instance (abort) (OS id: 8394)  
  CAUSE: A command to shutdown the instance was executed  
  ACTION: Check alert log for progress and completion of command  
  CLASS: CDB Instance / CDB ADMINISTRATOR / ERROR / AL-1002  
  TIME: 2020-05-08T17:09:33.773-07:00  
  
  ADDITIONAL INFO: -  
  Shutdown is initiated by sqlplus@den02tlh (TNS V1-V3).  
]
```

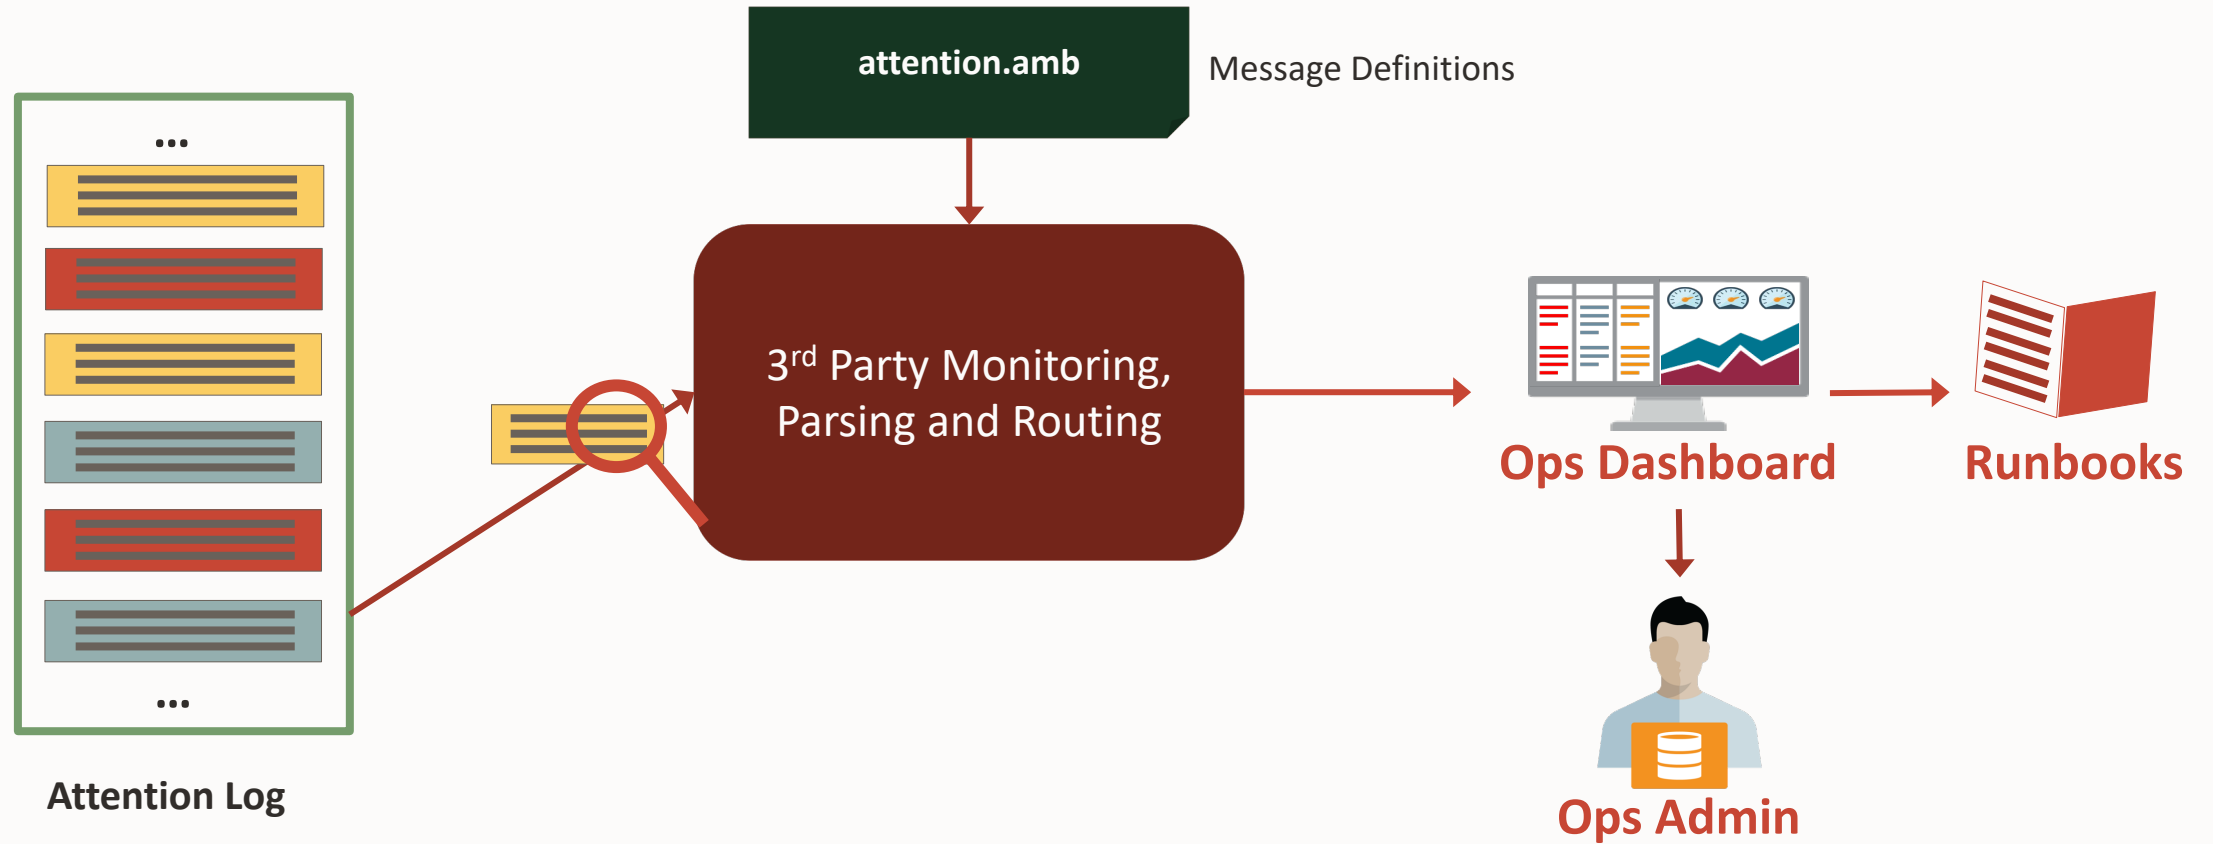
## Example Attention Log Curated Message – Server Warning

```
[  
  SOON           Heavy swapping observed on system  
  CAUSE:         Memory usage by one more application is leading to heavy swapping  
  ACTION:        Check alert log for more information, use tools to analyze memory  
                  usage and take action  
  CLASS:         CDB Instance / SERVER ADMINISTRATOR / WARNING / AL-2100  
  TIME:          2020-05-01T11:09:02.223-07:00  
  
  ADDITIONAL INFO: -  
  WARNING: Heavy swapping observed on system in last 15 mins.  
  Heavy swapping can lead to timeouts, poor performance, and instance eviction.  
]
```

# Attention Log Use Cases – Autonomous Health Framework



# Attention Log Use Cases – 3<sup>rd</sup> Party Monitoring



## Has anything changed recently?

```
tfact1 changes
Output from host : myserver69
-----
[Dec/01/2021 04:54:15.397]: Parameter: fs.aio-nr: Value: 95488 => 97024
[Dec/01/2021 04:54:15.397]: Parameter: fs.inode-nr: Value: 764974 131561 => 740744 131259
[Dec/01/2021 04:54:15.397]: Parameter: kernel.pty.nr: Value: 2 => 1
[Dec/01/2021 04:54:15.397]: Parameter: kernel.random.entropy_avail: Value: 189 => 158
[Dec/01/2021 04:54:15.397]: Parameter: kernel.random.uuid: Value: 36269877-9bc9-40a3-82e0-
1619865096f2 => 7551c5e7-c59f-40fa-b55f-5bd170e8b1ab
[Dec/01/2021 05:46:15.397]: Parameter: fs.aio-nr: Value: 119680 => 122880
[Dec/01/2021 05:46:15.397]: Parameter: fs.inode-nr: Value: 1580316      810036 => 1562320
768555
[Dec/01/2021 05:46:15.397]: Parameter: kernel.pty.nr: Value: 19 => 18
[Dec/01/2021 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
-----
[Dec/01/2021 04:54:15.397]: Parameter: fs.aio-nr: Value: 95488 => 97024
[Dec/01/2021 04:54:15.397]: Parameter: fs.inode-nr: Value: 764974 131561 => 740744 131259
[Dec/01/2021 04:54:15.397]: Parameter: kernel.ptty.nr: Value: 2 => 1
[Dec/01/2021 04:54:15.397]: Parameter: kernel.random.entropy_avail: Value: 189 => 158
[Dec/01/2021 04:54:15.397]: Parameter: kernel.random.uuid: Value: 36269877-9bc9-40a3-82e0-
1619865096f2 => 7551c5e7-c59f-40fa-b55f-5bd170e8b1ab
[Dec/01/2021 05:46:15.397]: Parameter: fs.aio-nr: Value: 119680 => 122880
[Dec/01/2021 05:46:15.397]: Parameter: fs.inode-nr: Value: 1580316      810036 => 1562320
768555
[Dec/01/2021 05:46:15.397]: Parameter: kernel.ptty.nr: Value: 19 => 18
[Dec/01/2021 05:46:15.397]: Parameter: kernel.random.uuid: Value: 37cc31aa-ee31-459e-8f2a-
0766b34b1b64 => f5176cdc-6390-415d-882e-02c4cff2ae4e
[Dec/01/2021 16:56:15.398]: Parameter: fs.aio-nr: Value: 97024 => 98560
```



## tail files

```
tfactl tail alert
```

```
Output from host : myserver69
```

```
-----
```

```
/scratch/app/11.2.0.4/grid/log/myserver69/alertmyserver69.log
```

```
2021-12-01 23:28:22.532:
```

```
[ctssd(5630)]CRS-2409:The clock on host myserver69 is not synchronous with the mean cluster time. No action has been taken as the Cluster Time Synchronization Service is running in observer mode.
```

```
2021-12-01 23:58:22.964:
```

```
[ctssd(5630)]CRS-2409:The clock on host myserver69 is not synchronous with the mean cluster time. No action has been taken as the Cluster Time Synchronization Service is running in observer mode.
```

```
...
```





## tail files

...

```
/scratch/app/oradb/diag/rdbms/apxcmupg/apxcmupg_2/trace/alert_apxcmupg_2.log
```

```
Wed Dec 01 06:00:00 2021 VKRM started with pid=82, OS id=4903
```

```
Wed Dec 01 06:00:02 2021 Begin automatic SQL Tuning Advisor run for special tuning task  
"SYS_AUTO_SQL_TUNING_TASK"
```

```
Wed Dec 01 06:00:37 2021 End automatic SQL Tuning Advisor run for special tuning task  
"SYS_AUTO_SQL_TUNING_TASK"
```

```
Wed Dec 01 23:00:28 2021 Thread 2 advanced to log sequence 759 (LGWR switch)
```

```
Current log# 3 seq# 759 mem# 0: +DATA/apxcmupg/onlinelog/group_3.289.917164707
```

```
Current log# 3 seq# 759 mem# 1: +FRA/apxcmupg/onlinelog/group_3.289.917164707
```

...

## Around 100 problem types covered

Database areas

Errors / Corruption

Performance

Install / patching / upgrade

RAC / Grid Infrastructure

Import / Export

RMAN

Transparent Data Encryption

Storage / partitioning

Undo / auditing

Listener / naming services

Spatial / XDB

## Other Server Technology

Enterprise Manager

Data Guard

GoldenGate

Exalogic

[Full list in documentation](#)

```
tfact1 diagcollect -srdc <srdc_type> [-sr <sr_number>]
```



# Manual collection vs TFA SRDC for database performance

## Manual method

1. Generate ADDM reviewing [Document 1680075.1](#) (multiple steps)
2. Identify “good” and “problem” periods and gather AWR reviewing [Document 1903158.1](#) (multiple steps)
3. Generate AWR compare report (awrddrpt.sql) using “good” and “problem” periods
4. Generate ASH report for “good” and “problem” periods reviewing [Document 1903145.1](#) (multiple steps)
5. Collect OSWatcher data reviewing [Document 301137.1](#) (multiple steps)
6. Collect Hang Analyze output at Level 4
7. Generate SQL Healthcheck for problem SQL id using [Document 1366133.1](#) (multiple steps)
8. Run support provided sql scripts – Log File sync diagnostic output using [Document 1064487.1](#) (multiple steps)
9. Check alert.log if there are any errors during the “problem” period
10. Find any trace files generated during the “problem” period
11. Collate and upload all the above files/outputs to SR

## TFA SRDC

1. Run

```
tfactl diagcollect -srdc dbperf [-sr <sr_number>]
```



## One command SRDC

### **tfactl diagcollect -srdc <srdc\_type>**

- Scans system to identify recent events
- Once the relevant event is chosen, proceeds with diagnostic collection

```
tfactl diagcollect -srdc ORA-00600
Enter the time of the ORA-00600 [YYYY-MM-DD HH24:MI:SS,<RETURN>=ALL] :
Enter the Database Name [<RETURN>=ALL] :

1. Dec/01/2021 05:29:58 : [orcl2] ORA-00600: internal error code, arguments: [600], [], [],
[], [], [], [], [], [], []
2. Dec/01/2021 06:55:08 : [orcl2] ORA-00600: internal error code, arguments: [600], [], [],
[], [], [], [], [], [], []

Please choose the event : 1-2 [1]
Selected value is : 1 (Dec/01/2021 05:29:58 )
```

## One command SRDC

All required files are identified

- Trimmed where applicable
- Package in a zip ready to provide to support

```
...
2021/12/01 06:14:24 EST : Getting List of Files to Collect
2021/12/01 06:14:27 EST : Trimming file :
myserver1/rdbms/orcl2/orcl2/trace/orcl2_lmhb_3542.trc with original file size : 163MB
...
2021/12/01 06:14:58 EST : Total time taken : 39s
2021/12/01 06:14:58 EST : Completed collection of zip files.
...
/opt/oracle.ahf/data/repository/srdc_ora600_collection_Wed_Dec_1_06_14_17_EST_2021_node_loca
l/myserver1.tfa_srdc_ora600_Wed_Dec_1_06_14_17_EST_2021.zip
```

## Understand Database **log disk space usage**

```
tfactl managelogs -show usage
```

```
...
```

```
|-----|
|                Grid Infrastructure Usage                |
|-----+-----|
| Location                                             | Size |
|-----+-----|
| /u01/app/crsusr/diag/afdbot/user_root/host_309243680_94/alert | 28.00 KB |
| /u01/app/crsusr/diag/afdbot/user_root/host_309243680_94/incident | 4.00 KB |
| /u01/app/crsusr/diag/afdbot/user_root/host_309243680_94/trace | 8.00 KB |
|-----+-----|
| Total                                             | 739.06 MB |
|-----+-----|
```

```
...
```

Use **-gi** to only show grid infrastructure

## Understand Database **log disk space usage**

```
...
+-----+
|                Database Homes Usage                |
+-----+-----+
| Location                                             | Size |
+-----+-----+
| /u01/app/crsusr/diag/rdbms/cdb674/CDB674/alert      | 1.06 MB |
| /u01/app/crsusr/diag/rdbms/cdb674/CDB674/incident  | 4.00 KB |
| /u01/app/crsusr/diag/rdbms/cdb674/CDB674/trace    | 146.19 MB |
| /u01/app/crsusr/diag/rdbms/cdb674/CDB674/cdump    | 4.00 KB |
| /u01/app/crsusr/diag/rdbms/cdb674/CDB674/hm       | 4.00 KB |
+-----+-----+
| Total                                               | 147.26 MB |
+-----+-----+
```

Use **-database** to only show database

Copyright © 2021, Oracle and/or its affiliates



# Understand Database log disk space usage variations

```
tfactl managelogs -show variation -older 30d
```

```
Output from host : myserver74
```

```
-----  
2021-12-01 12:30:42: INFO Checking space variation for 30 days
```

```
.....  
|                                     |                                     |                                     |  
|                                     |                                     |                                     |  
+-----+-----+-----+  
| Directory | Old Size | New Size |  
+-----+-----+-----+  
| /u01/app/crsusr/diag/asm/user_root/host_309243680_96/alert | 22.00 KB | 28.00 KB |  
+-----+-----+-----+  
| /u01/app/crsusr/diag/clients/user_crsusr/host_309243680_96/cdump | 4.00 KB | 4.00 KB |  
+-----+-----+-----+  
| /u01/app/crsusr/diag/tnslsnr/myserver74/listener/alert | 15.06 MB | 244.10 MB |  
+-----+-----+-----+  
| ... |
```

Directory	Old Size	New Size
/u01/app/crsusr/diag/asm/user_root/host_309243680_96/alert	22.00 KB	28.00 KB
/u01/app/crsusr/diag/clients/user_crsusr/host_309243680_96/cdump	4.00 KB	4.00 KB
/u01/app/crsusr/diag/tnslsnr/myserver74/listener/alert	15.06 MB	244.10 MB





# Run a database log purge

```
tfactl managelogs -purge -older 30d
```

```
Output from host : myserver74
```

```
-----
```

```
Purging files older than 30 days
```

```
Cleaning Grid Infrastructure destinations
```

```
Purging diagnostic destination "diag/afdbboot/user_root/host_309243680_94" for files - 0 files deleted , 0 bytes freed  
Purging diagnostic destination "diag/afdbboot/user_crsusr/host_309243680_94" for files - 1 files deleted , 10.16 KB freed  
Purging diagnostic destination "diag/asmtool/user_root/host_309243680_96" for files - 1 files deleted , 10.16 KB freed  
Purging diagnostic destination "diag/asmtool/user_crsusr/host_309243680_96" for files - 2 files deleted , 29.18 KB freed  
Purging diagnostic destination "diag/tnslsnr/myserver74/listener" for files - 2 files deleted , 29.18 KB freed  
Purging diagnostic destination "diag/diagtool/user_root/adrci_309243680_96" for files - 2 files deleted , 29.18 KB freed  
Purging diagnostic destination "diag/clients/user_crsusr/host_309243680_96" for files - 2 files deleted , 29.18 KB freed  
Purging diagnostic destination "diag/asm/+asm/+ASM" for files - 2 files deleted , 29.18 KB freed  
Purging diagnostic destination "diag/asm/user_root/host_309243680_96" for files - 2 files deleted , 29.18 KB freed  
Purging diagnostic destination "diag/asm/user_crsusr/host_309243680_96" for files - 2 files deleted , 29.18 KB freed  
Purging diagnostic destination "diag/crs/myserver74/crs" for files - 2 files deleted , 29.18 KB freed
```

```
...
```



# Monitor Database performance

```
tfact1 run oratop -database ogg19c
```

```
Oracle 19c - sh1 02:01:57 up: 4.3d, 2 ins, 128 sn, 1 us, 10G sga, 84.6%db
```

ID	%CPU	%DCP	LOAD	AAS	ASC	ASI	ASW	IDL	MBPS	%FR	PGA	UTPS	RT/X	DCTR	DWTR
2	19.2	13.7	10.2	39.2	11	4	22	26	100M	7	3.1G	656	56m	17	81
1	20.3	16.9	15.6	42.0	15	1	22	27	104M	14	2.7G	719	48m	16	78

EVENT (C)	TOT WAITS	TIME(s)	AVG MS	PCT	WAIT CLASS
log file sync	17841193	266406	14.1	39	Commit
DB CPU		154327		23	
gc current block busy	71129790	108578	1.5	16	Cluster
gc buffer busy acquire	8896661	73845	9.9	11	Cluster
enq: TX - row lock contention	3220952	71866	24.0	11	Application

ID	SID	SPID	USR	PROG	S	PGA	SQLID/BLOCKER	OPN	E/T	STA	STE	EVENT/*LA	W/T	
1	35	35023	TPC	tpcc	D	5.3M	4dxgf9wkjvzbz8	INS	0	ACT	I/O	db file s	29m	
1	932	35058	TPC	tpcc	D	5.3M	ahlvrygbhqdry	DEL	0	ACT	WAI	enq: TX -	12m	
2	6	7416	TPC	tpcc	D	5.2M	5zbjl8g6dz0gk	SEL	0	ACT	I/O	db file s	10m	
1	995	35051	TPC	tpcc	D	6.3M	462cz6g854c88	INS	0	ACT	I/O	db file s	10m	
2	1282	7328	TPC	tpcc	D	3.4M	5j4ntzvntxqxm	UPD	0	ACT	I/O	db file s	9.4m	
2	71	7343	TPC	tpcc	D	5.4M	5j4ntzvntxqxm	UPD	0	ACT	I/O	db file s	7.5m	
1	1251	35081	TPC	tpcc	D	4.3M	57kvfpfmx9801	UPD	0	ACT	WAI	enq: TX -	6.2m	
1	1447	35015	TPC	tpcc	D	5.1M	236ksg44h80ht	UPD	0	ACT	WAI	gc buffer	5.7m	
1	1220	35079	TPC	tpcc	D	4.3M	57kvfpfmx9801	UPD	0	ACT	WAI	enq: TX -	5.1m	
1	4	35021	TPC	tpcc	D	5.3M	4j8r6zsm6upcv	UPD	0	ACT	WAI	gc curren	5.0m	
2	902	7376	TPC	tpcc	D	5.1M			0	ACT	WAI	log file	4.4m	
1	34	35159	TPC	tpcc	D	5.3M	gz8100xp0fbc9		0	ACT	CPU	cpu runqu	3.9m	
1	1446	35097	TPC	tpcc	D	5.4M	4j8r6zsm6upcv	UPD	0	ACT	WAI	gc buffer	3.7m	
2	582	7356	TPC	tpcc	D	5.1M			0	ACT	WAI	log file	3.6m	
1	102	35164	TPC	tpcc	D	4.5M	236ksg44h80ht	UPD	0	ACT	WAI	gc curren	2.6m	
2	1473	49080	B/G	LGWR	D	2.7M			2:1	4.3d	ACT	WAI	LGWR any	2.2m
1	1	31199	B/G	LG00	D	2.4M			2.8d	ACT	CPU	cpu runqu	2.2m	



## Querying metrics in AHF

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

ex : oclumon dumpnodeview local -system

```
-----  
Node : den03ceb Clock : '2020-09-08 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 : '2020-09-10 05.22.39'  
-----  
  
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 : '2020-09-08 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 : '2020-09-08 08.36.29'  
-----  
  
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 : '2020-09-09 23.48.24'  
-----  
  
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 : '2020-09-08 08.47.59'  
-----  
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 : '2020-09-08 08.49.39'  
-----  
FILESYSTEM:  
  mount  type  total[KB]  used[KB]  avbl[KB]  used[%]  ifree[%]  tag  
    /    ext4  51473888  33168200  15667916  68.00    86.00    GRID_HOME
```








# AHF Insights (coming in 22.3)









AHF Insights - Summary AHF23.0.0.0

System State : **Stable** | System Type : Exadata | Time Range : 25-MAR-2022 08:50 to 25-MAR-2022 10:50 **(2 Hours)**

## System Topology

 <b>1</b> Cluster <b>GI Version : 21.3.0.0.0</b>	 <b>5</b> Databases <b>5 CDB(s) [ 5 PDB(s) / [ 5 open ] ]</b>	 <b>2</b> Database Servers <b>X4-8</b>	 <b>12</b> Storage Servers <b>X5-2, X5-3</b>	 <b>1</b> RDMA Network Fabric Switch <b>ROCE Switch</b>
---	--	---	---	--

## Insights

 <b>8</b> Timeline <b>All Log &amp; Metric events</b>	 <b>3</b> Operating System Issues <b>Memory, Network Issue</b>	 <b>17</b> Best Practice Issues <b>Health Score : 86</b>	 <b>12</b> System Change <b>12 changes in 30 days</b>	 <b>2</b> Recommended Software <b>All Components</b>
 <b>2</b> Management Server <b>2 Uncleared Alerts</b>	 <b>RPM List</b> <b>List of RPMs</b>	 <b>Database Parameters</b> <b>List of Database Parameters</b>		



ORACLE

# Thank you

OGBEMEA Sep 2022

**Sandesh Rao**

VP AIOps for the Autonomous Database

 [@sandeshr](https://twitter.com/sandeshr)

 <https://www.linkedin.com/in/raosandesh/>

 <https://www.slideshare.net/SandeshRao4>