ORACLE

# Autonomous Health Framework

## Resource Allocation

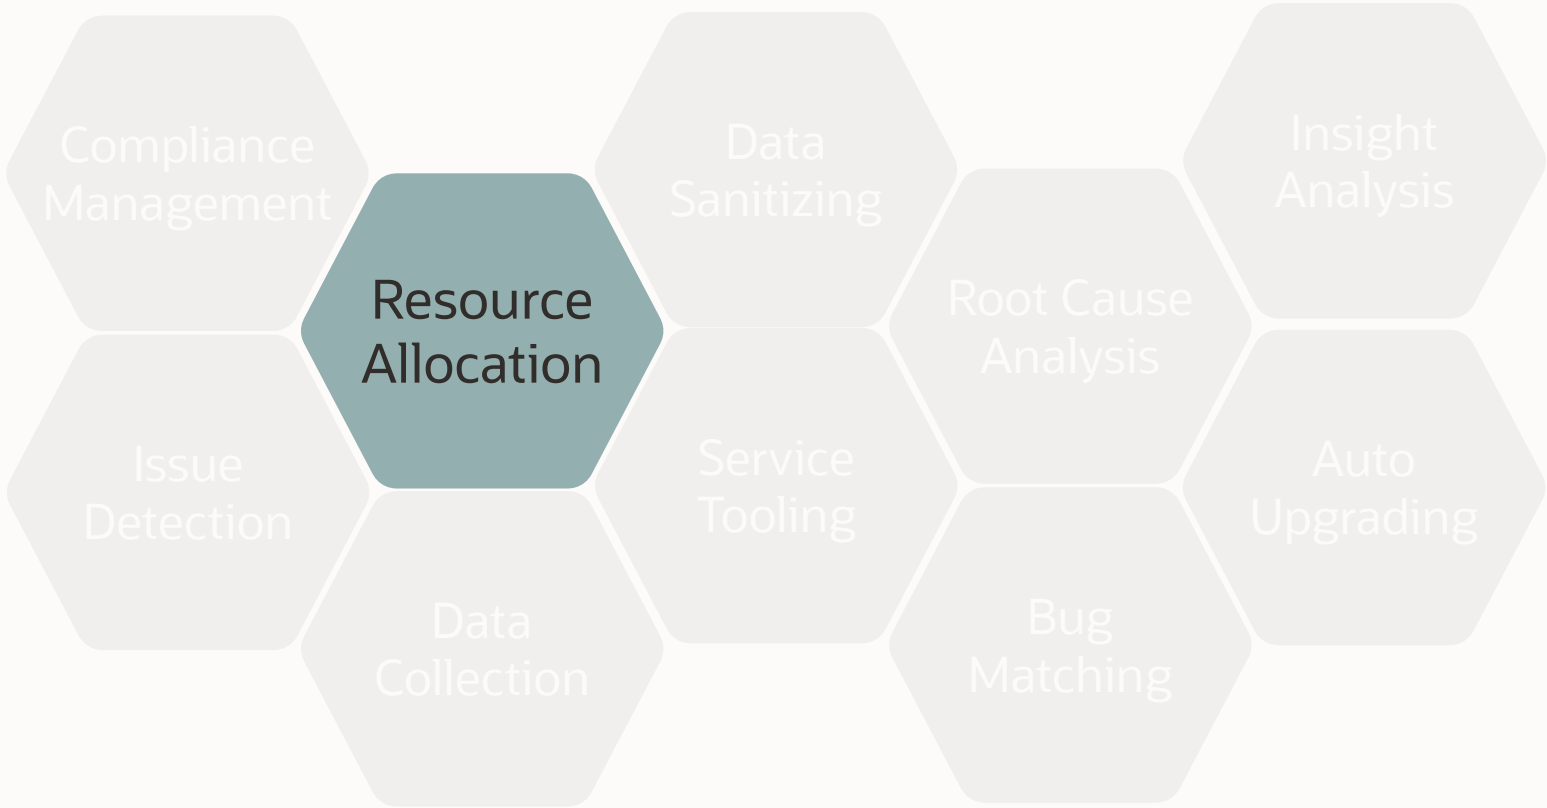

# Oracle Autonomous Health Framework

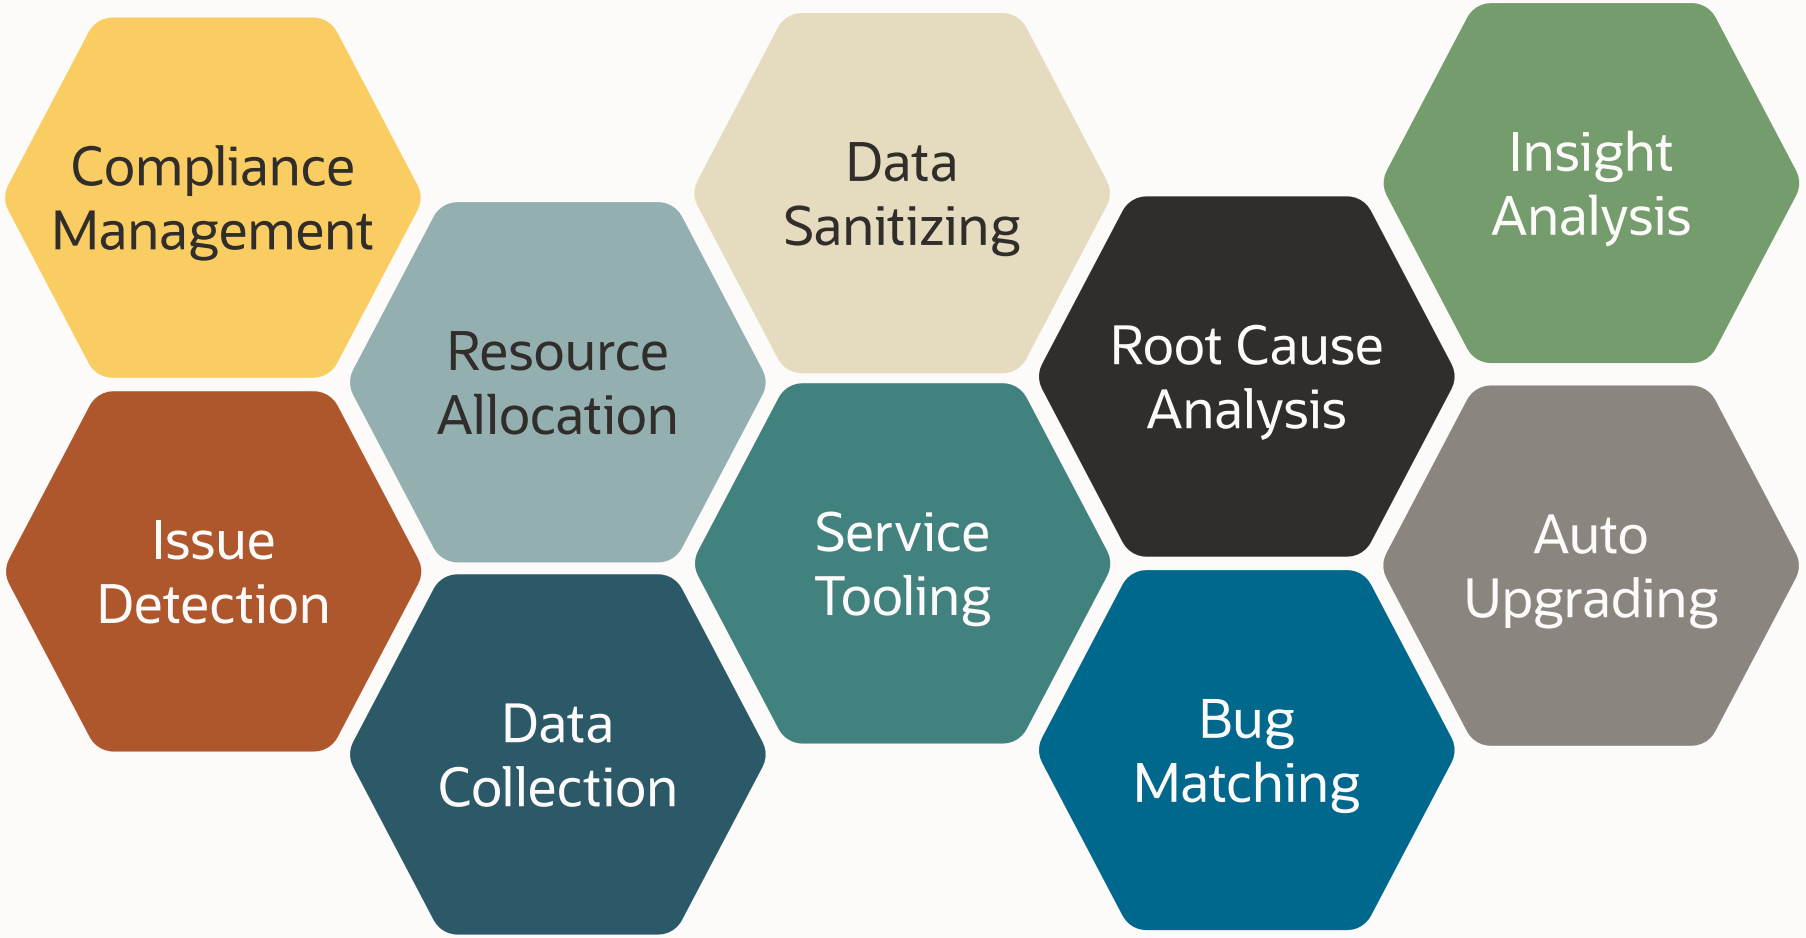

# Oracle Autonomous Health Framework

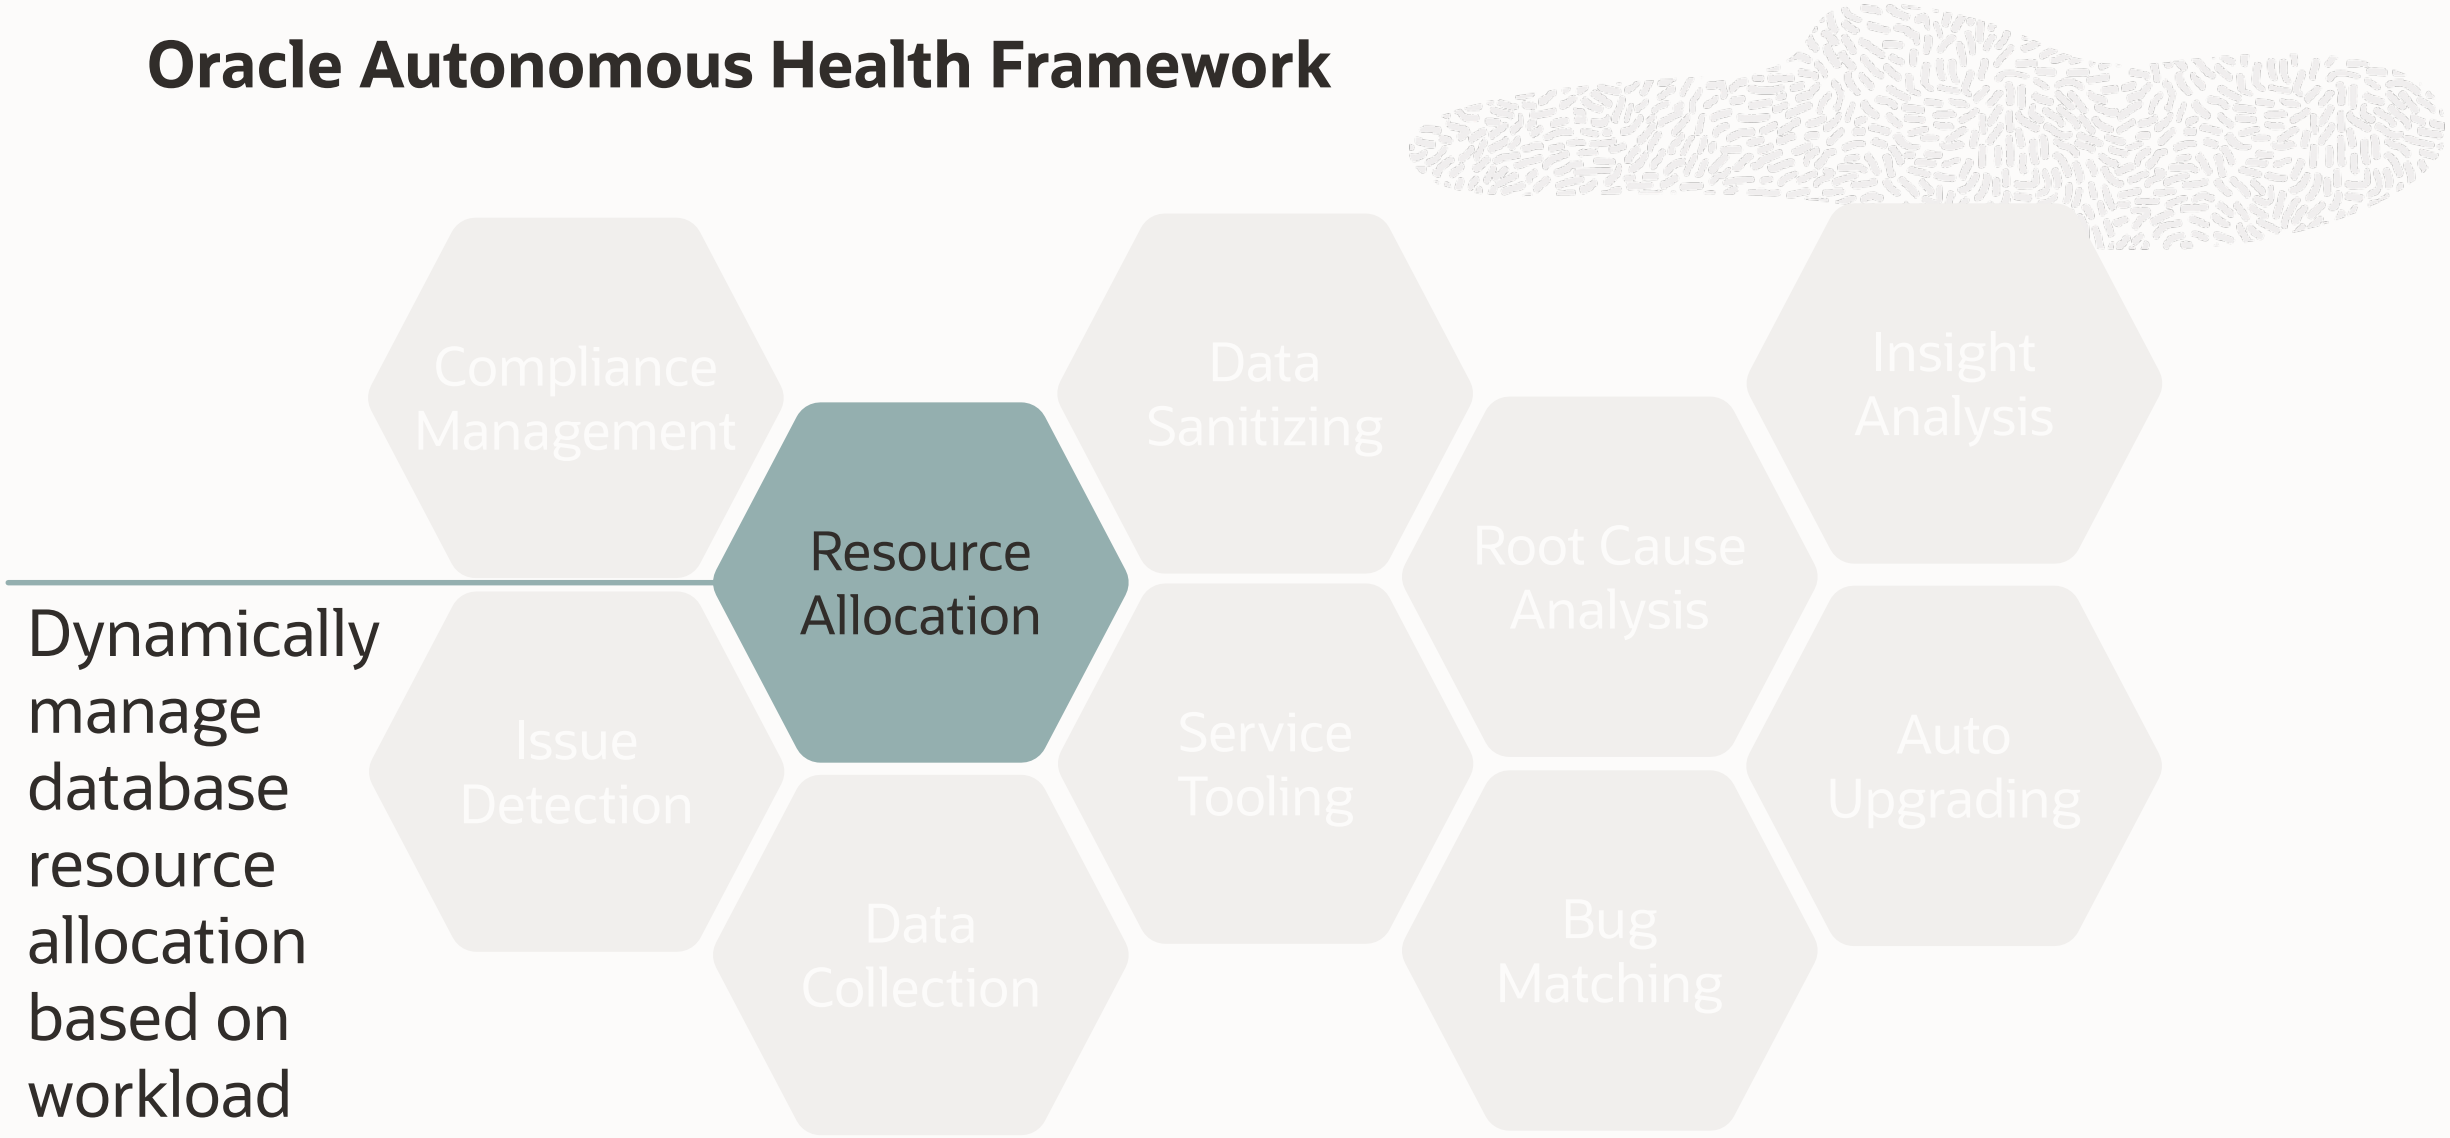


Dynamically manage database resource allocation based on workload

# Session Performance

Preventing a session from blocking performance of others by holding critical resources

# Database Hang Management – Applied Machine Learning

Discovers and Resolves Runtime Database Hangs and Deadlocks

- Actual Internal and External customer data drives model development
- Purpose-built diagnostic technology used for knowledge extraction
- Expert Dev team scrubs data
- Hang Heuristic  Engine created and deployed in database background
- HM uses run-time engine to perform real-time DB hang detection and resolution

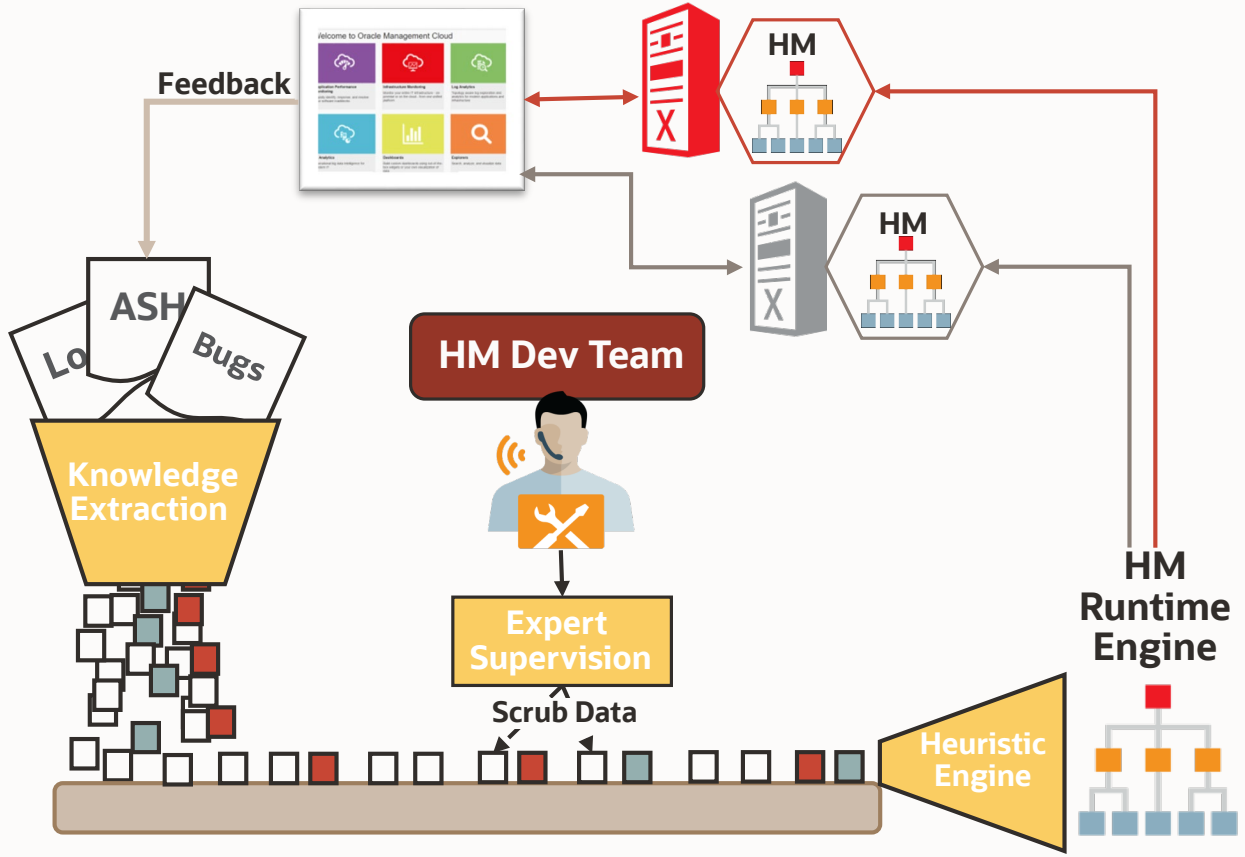

# Database Hang Management – Session

## Autonomously Preserves Database Availability and Performance

- Monitors Session snapshots for progress
- Evaluates potential hangs over time with based upon Wait Graphs
- Analyzes hang chain of sessions to identify blocker/victim
- Verifies blocker session is hung for time period to become victim
- Uses heuristic model to determine victim resolution
- Terminates session or process to resolve

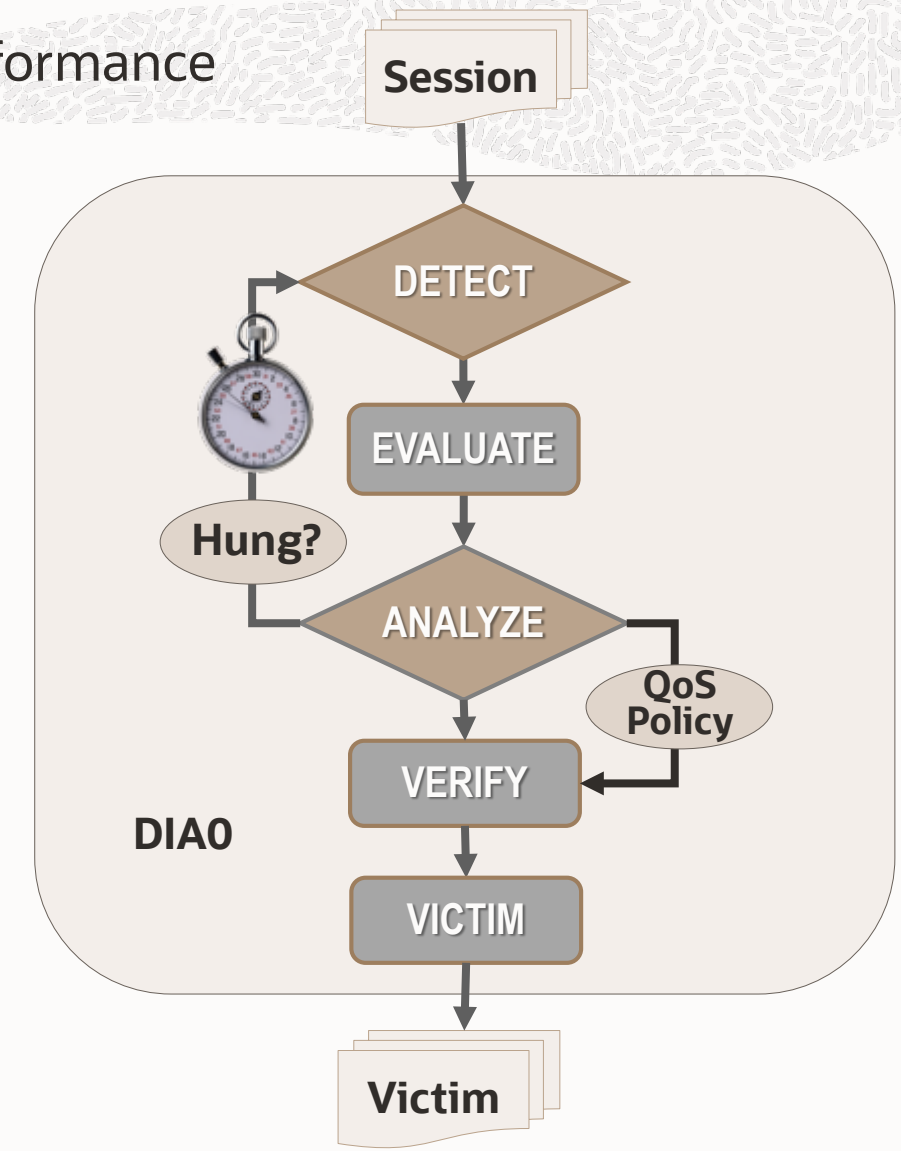

# Database Performance

Preserving instance performance  when database resources are constrained

# Database Health - Applied Machine Learning

## Discovers Potential Cluster & DB Problems

Actual Internal and External customer 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

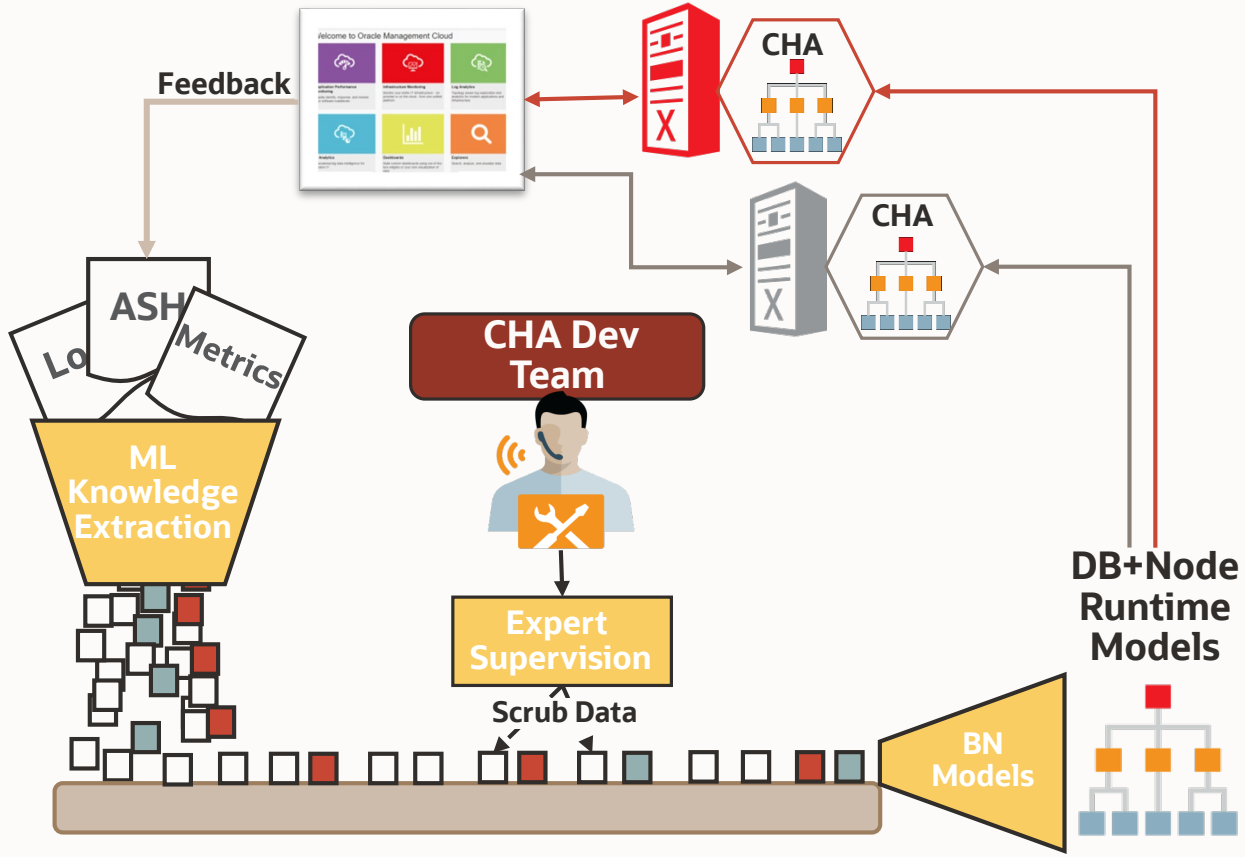

# Autonomous Health – Database Performance

## Database Data Flow Overview

Reads OS and DB Performance data directly from memory

Uses Machine Learning models and data to perform prognostics

Detects common RAC database problems

Performs root cause analysis

Sends alerts and preventative actions to Cloud Ops per target

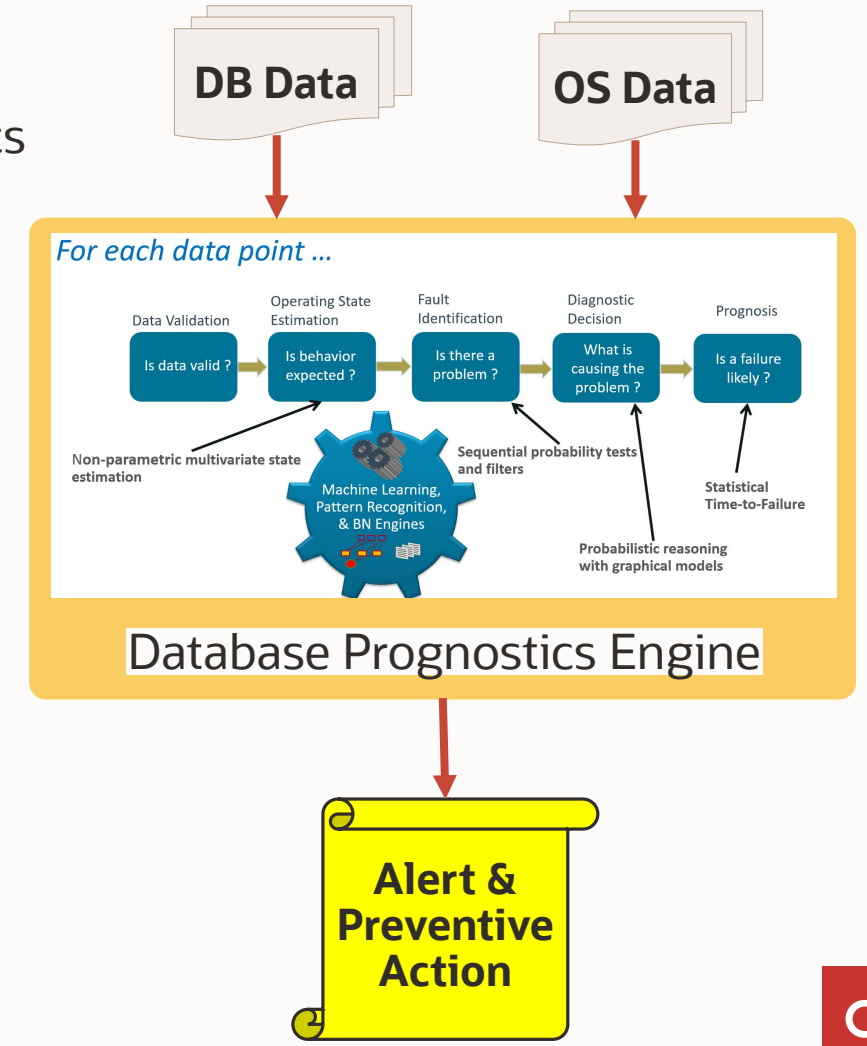

# Autonomous Health – Database Performance

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

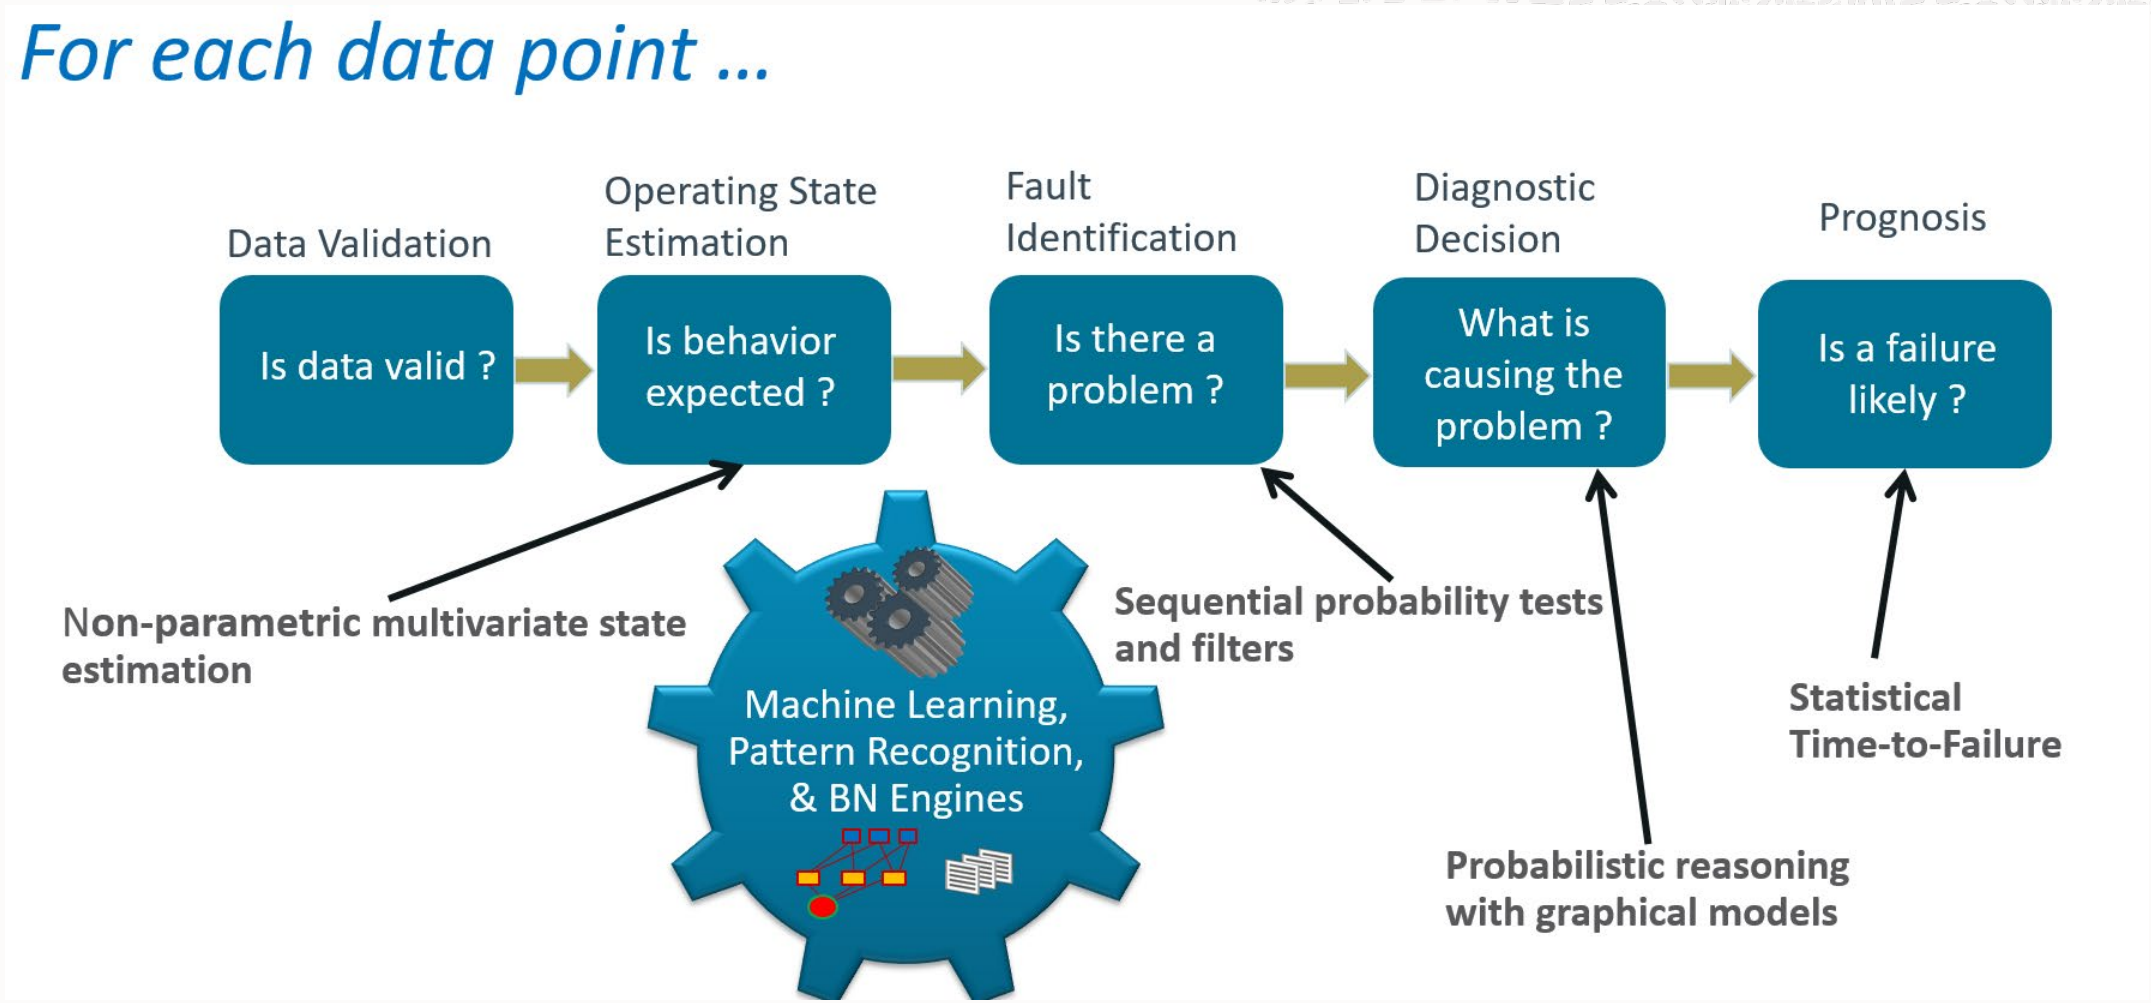

# Models Capture all Normal Operating Modes

Models Capture the Dynamic Behavior of all Normal Operation

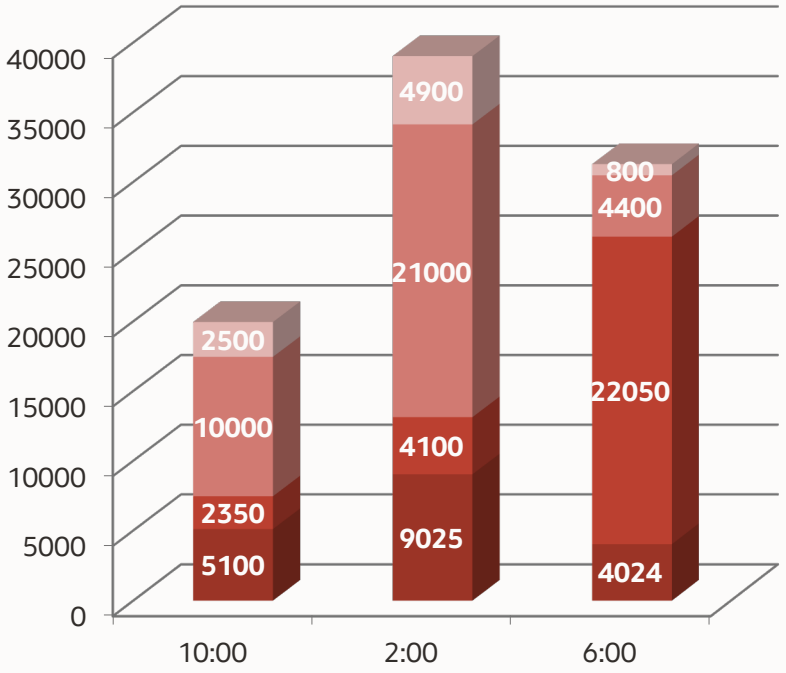


| In-Memory Reference Matrix (*Part of "Normality" Model*) | | | | | |
|---|---|---|---|---|---|
| **IOPS** | **####** | **2500** | **4900** | **800** | **####** |
| **User Commits** | **####** | **10000** | **21000** | **4400** | **####** |
| **Log File Parallel Write** | **####** | **2350** | **4100** | **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

## CHA Model: Find Similarity with Normal Values

| 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 | ##### |
| … | … | … | … | … | … |

**Observed values (*Part of a Data Point*)**

| |
|---|
| 10500 |
| 20000 |
| 4050 |
| 10250 |
| … |

Observed - Predicted =

**Residual Values (*Part of a Data Point*)**

| |
|---|
| 5600 |
| -1000 |
| -50 |
| 325 |
| … |

**CHA estimator/predictor (ESEE): "*based on my normality model, the value of IOPS should be in the vicinity of ~ 4900, but it is reported as 10500, this is causing a residual of ~ 5600 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 IOPS and if it remains deviant I'll generate a fault on it*".**

# 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**

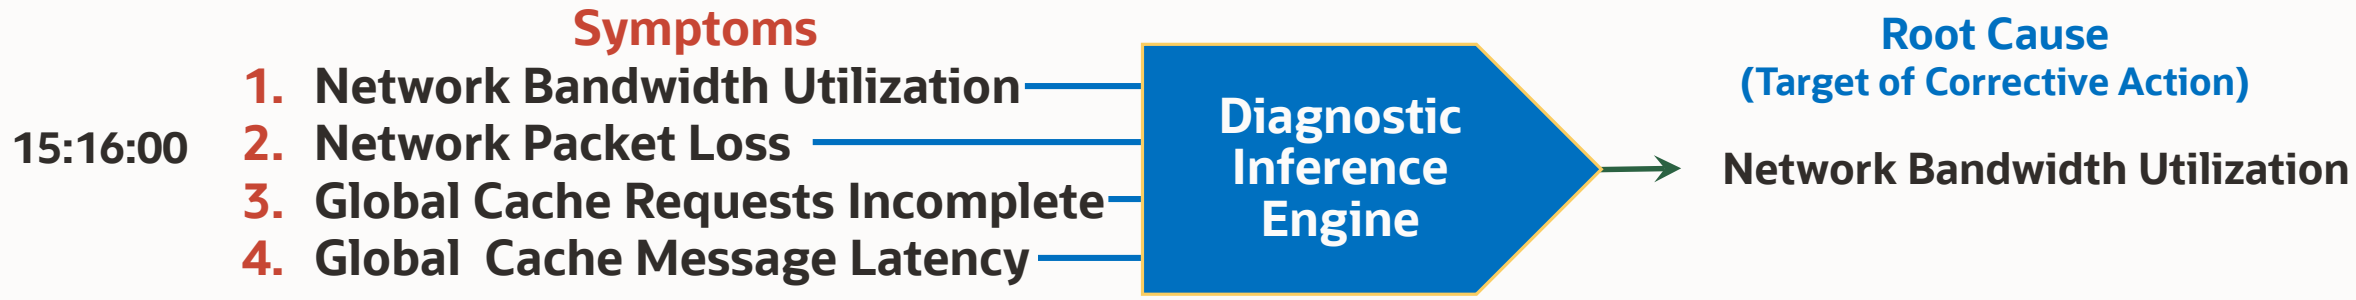

# Autonomous Health - Cluster Health Advisor

*Cross Node and Cross Instance Diagnostic Inference*

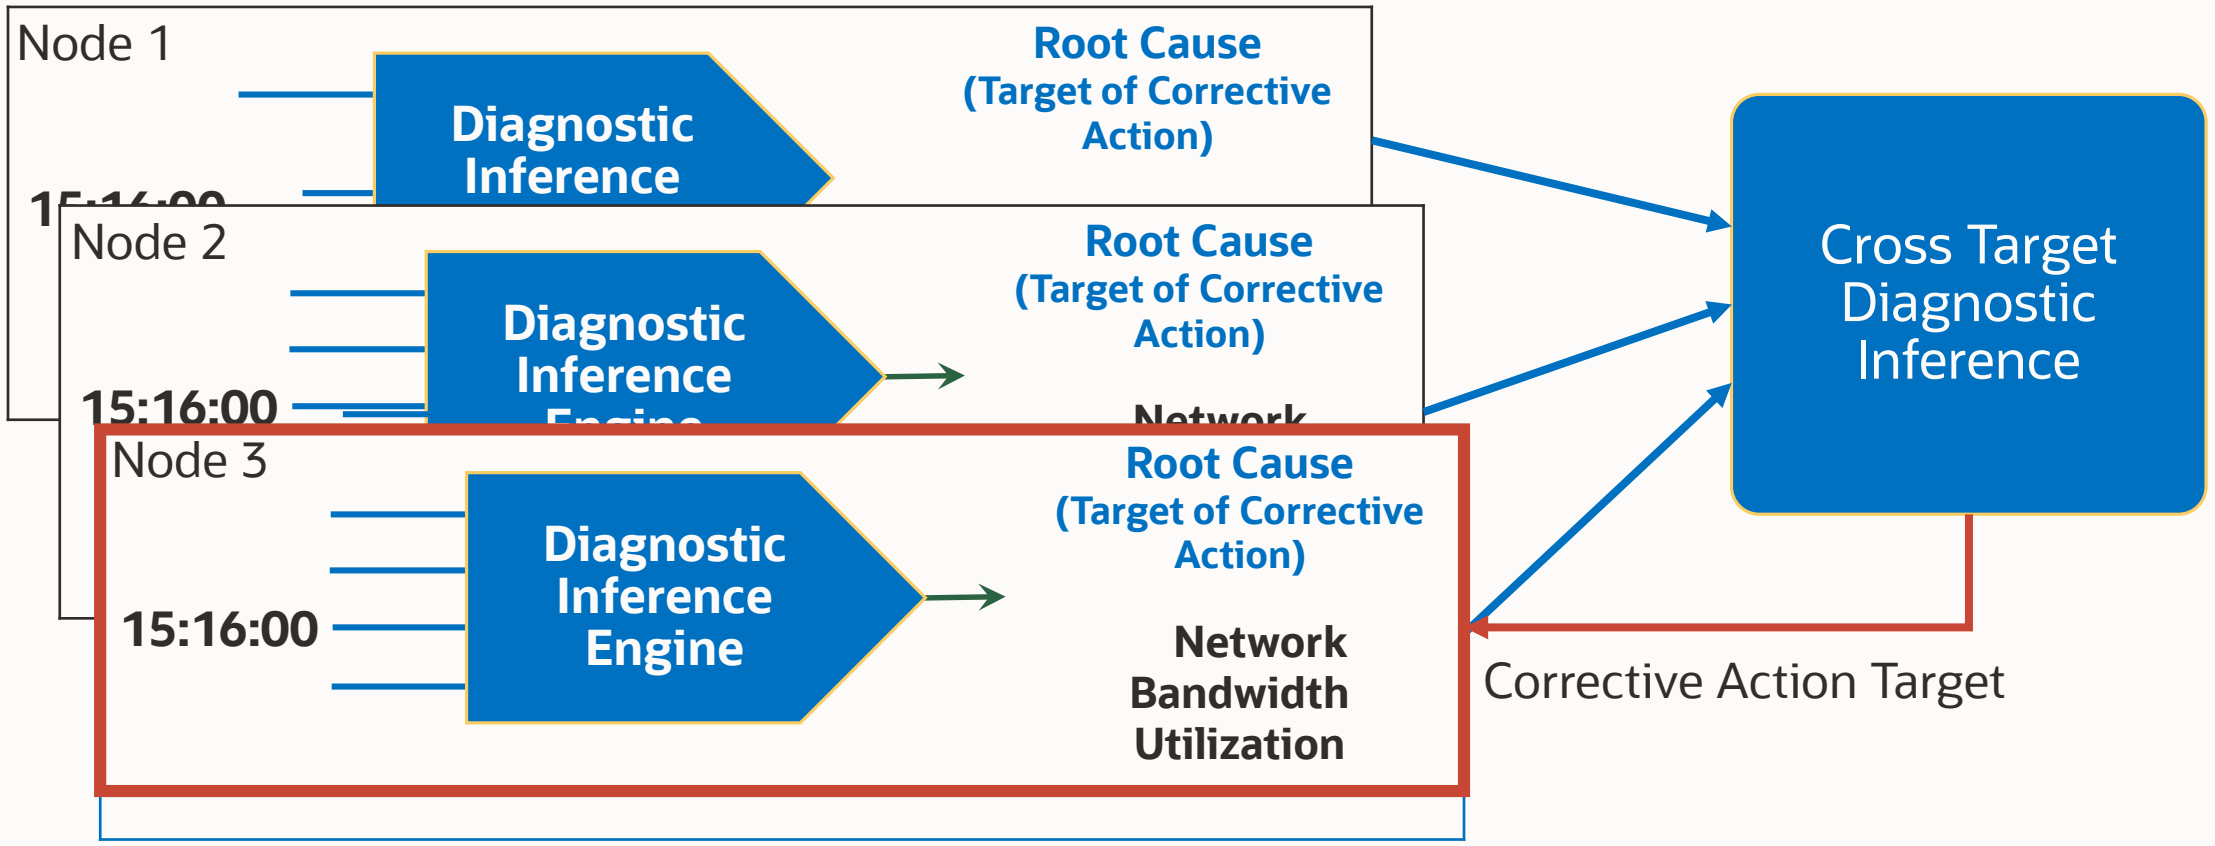

# Infrastructure Performance

Preserving database availability when an ASM instance is hung or blocked

# Database Hang Management - Infrastructure

## Autonomously Preserves Database Infrastructure Availability

Monitors Session snapshots for progress

Evaluates potential hangs over time with based upon Wait Graphs

Analyzes hang chain of sessions to identify blocker/victim

Discovers blocker is located in ASM instance

Requests ASM terminate session or instance relying on Flex ASM for recovery

Detection and resolution is bi-directional

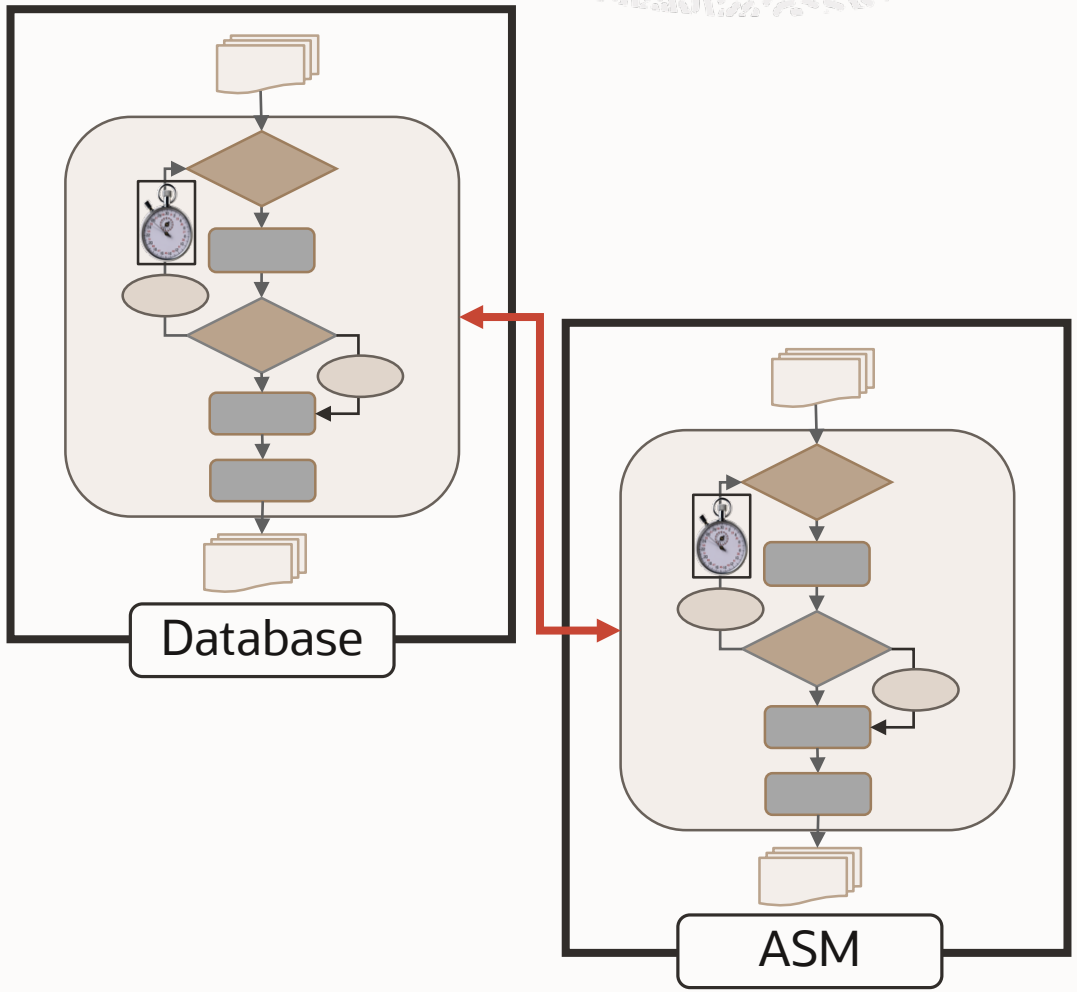


ORACLE