Beyond Diagnostics and Tuning:
Using Machine Learning and Data Analytics to Maximize DB Performance

Mughees A. Minhas
VP, Product Management
Oracle Corp.

ORACLE

Agenda

1. Database Management: Landscape and Challenges
2. Oracle Management Cloud
3. Real-Time Analytics
4. Anomaly Detection and Machine Learning
5. Q&A

DB Management and Monitoring Trends

- Heterogeneity and hyper-extensibility
  - Ability to monitor and manage open-source, relational, NoSQL databases running on traditional hardware, cloud, or software defined entities (e.g., VMs, Dockers)
- Hybrid/Multi-cloud management
  - Administer databases seamlessly across cloud boundaries offering single pane-of-glass for monitoring and management
- Demand for real-time analytic offerings
  - Problem avoidance (vs. fast resolution)
  - Need for more algorithmic decision-making

Management Challenges

END USER EXPERIENCE / ACTIVITY	**Real Users Synthetic Users**
APPLICATION	**App Metrics Transactions**
MIDDLE TIER	**Server Metrics Diagnostics**
DATA TIER	**Alert Log and Trace Files**
VIRTUALIZATION TIER	**Host Metrics VM Metrics Container Metrics**
INFRASTRUCTURE TIER	**CMDB Tickets Alerts**

Too many tools

Too much data

No insight

ORACLE

Powered by AI & Machine Learning

END USER EXPERIENCE / ACTIVITY

APPLICATION

MIDDLE TIER

DATA TIER

VIRTUALIZATION TIER

INFRASTRUCTURE TIER

Real Users
Synthetic Users

App Metrics
Transactions

Server Metrics
Diagnostics

Logs

Host Metrics
VM Metrics
Container Metrics

CMDB Tickets
Alerts

Unified Platform

ANOMALY DETECTION

CLUSTERING

PREDICTION

CORRELATION

Analytics and ML

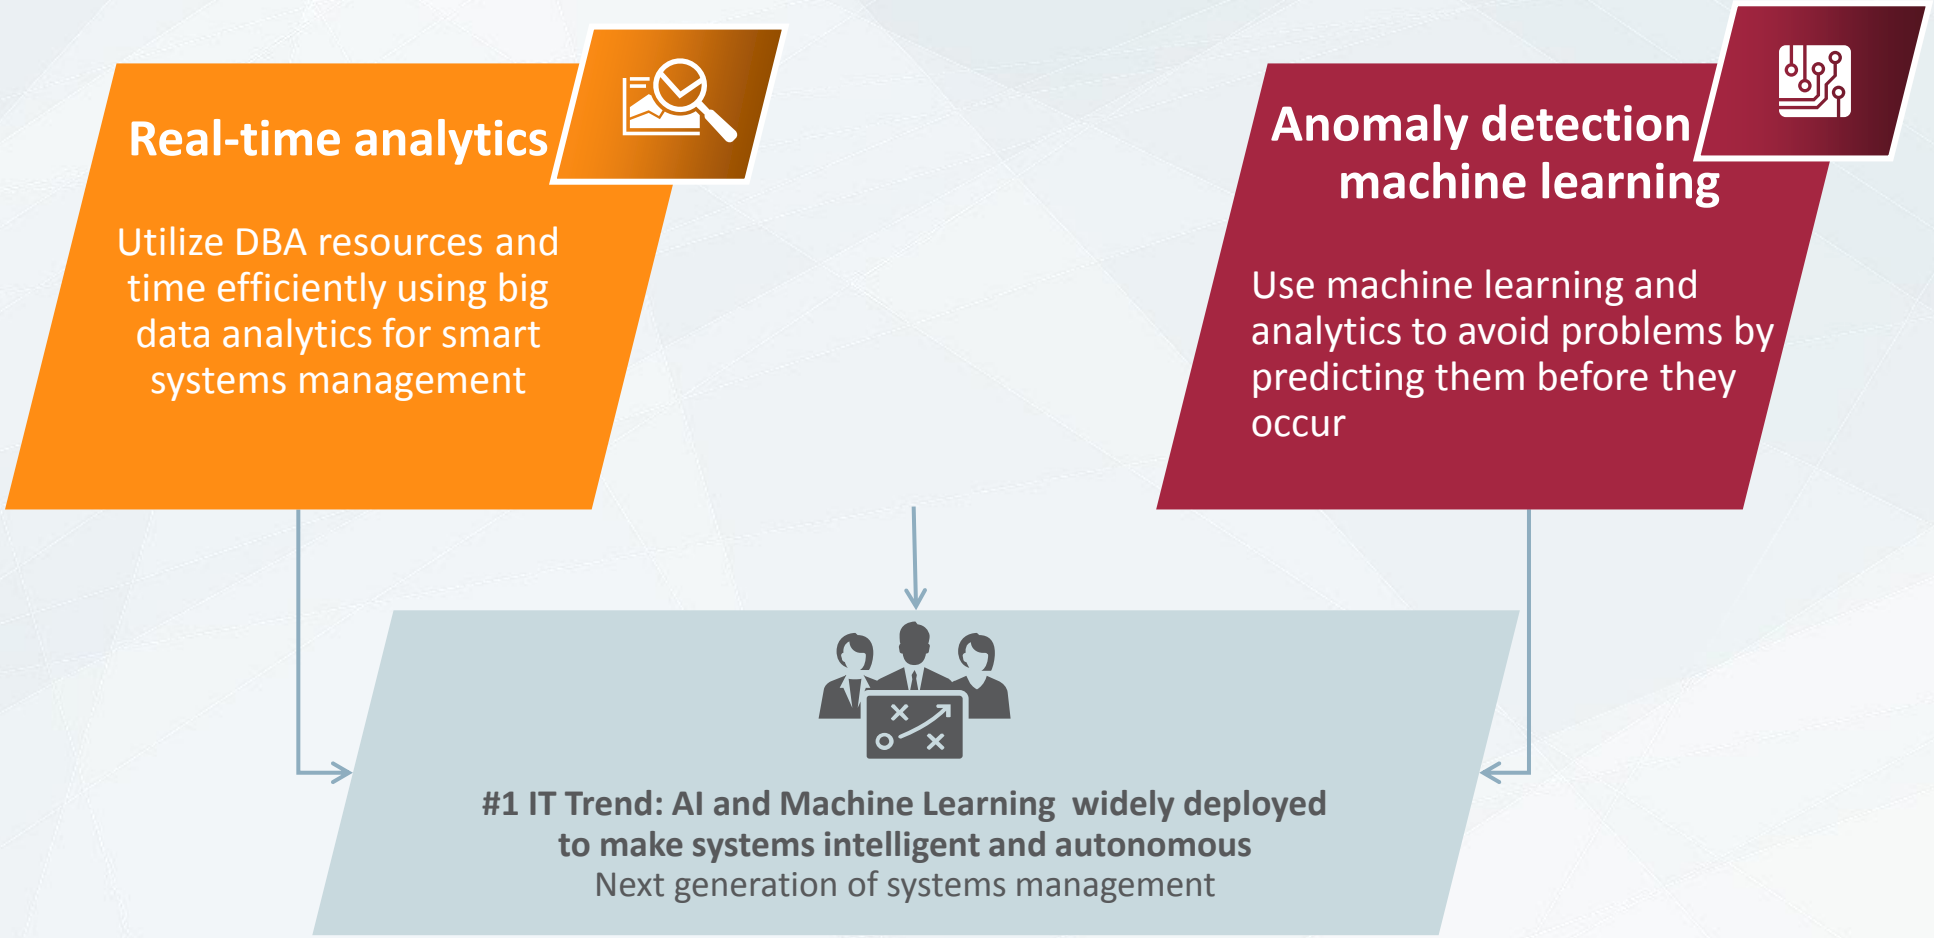

Real-Time Analytics

Real-Time Analytics: Find Needle in the Haystack

Performance Analytics

- Analyze database performance based on long term performance data
- Get insights into database performance
  - Analyze performance degradation by response time
  - Identify inefficient databases by wait time
  - Analyze performance of highly variant SQLs and databases
  - Databases which are increasingly inefficient
  - Top SQLs

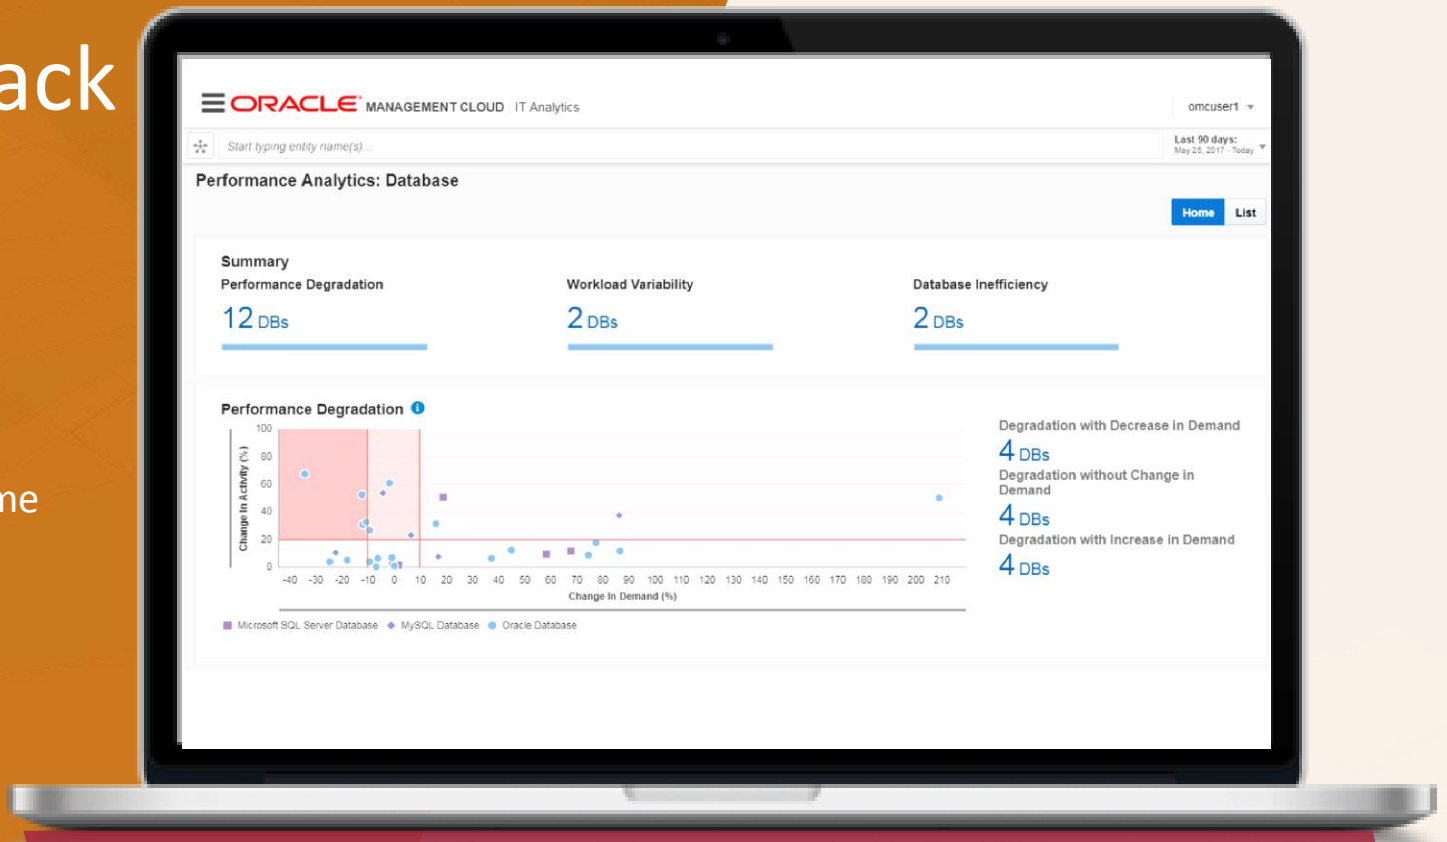

Maximize your cloud investments

- Achieve highest level of database performance
- Consolidate databases for better price to performance

Performance Analytics

Analyze performance degradation by response time

- Degradation with decrease in demand
  - Change in Response Time > +20% and Change in Demand < -10 %
- Degradation without change in demand
  - Change in Response Time > +20% and Change in Demand between -10% and +10%
- Degradation with increase in demand
  - Change in Response Time > +20% and Change in Demand < +10 %

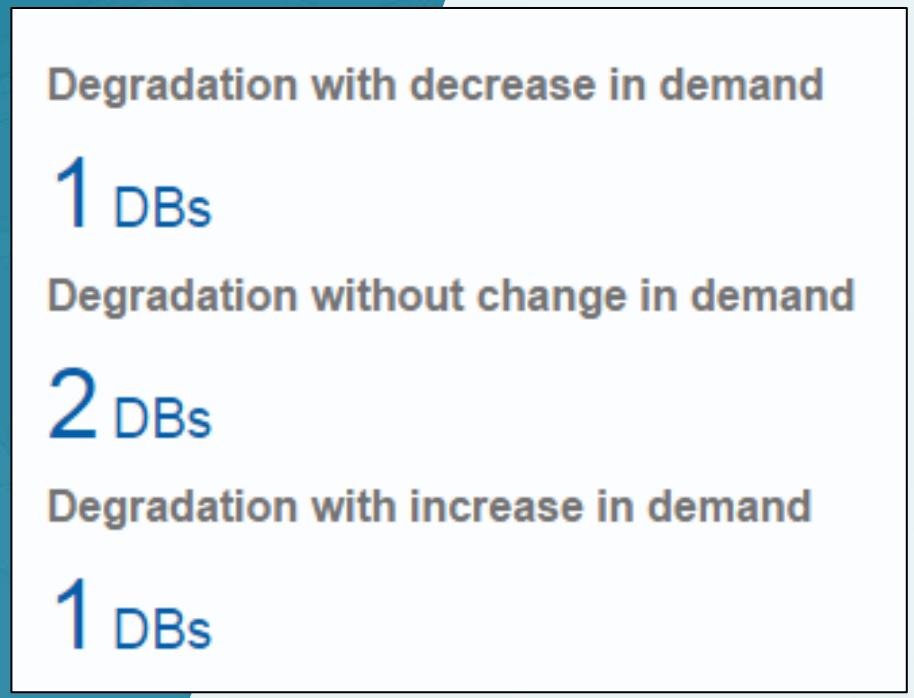

Configurable thresholds for “Change in Response Time” and “Change in Demand”

- **Default Values:**
  - **Change in Response Time - 20%**
  - **Change in Demand - 10%**

Performance Analytics

Analyze inefficient workloads

- Inefficient wait time is in Active Sessions which are not part of CPU Time, IO wait Time and idle wait events
- Inefficient (%) = Inefficient Wait Time / DB time
- Increasingly inefficient among the most inefficient databases = DBs with Inefficient (%) > 50 % and Inefficiency (%) Increasing over time

Inefficient databases to focus on

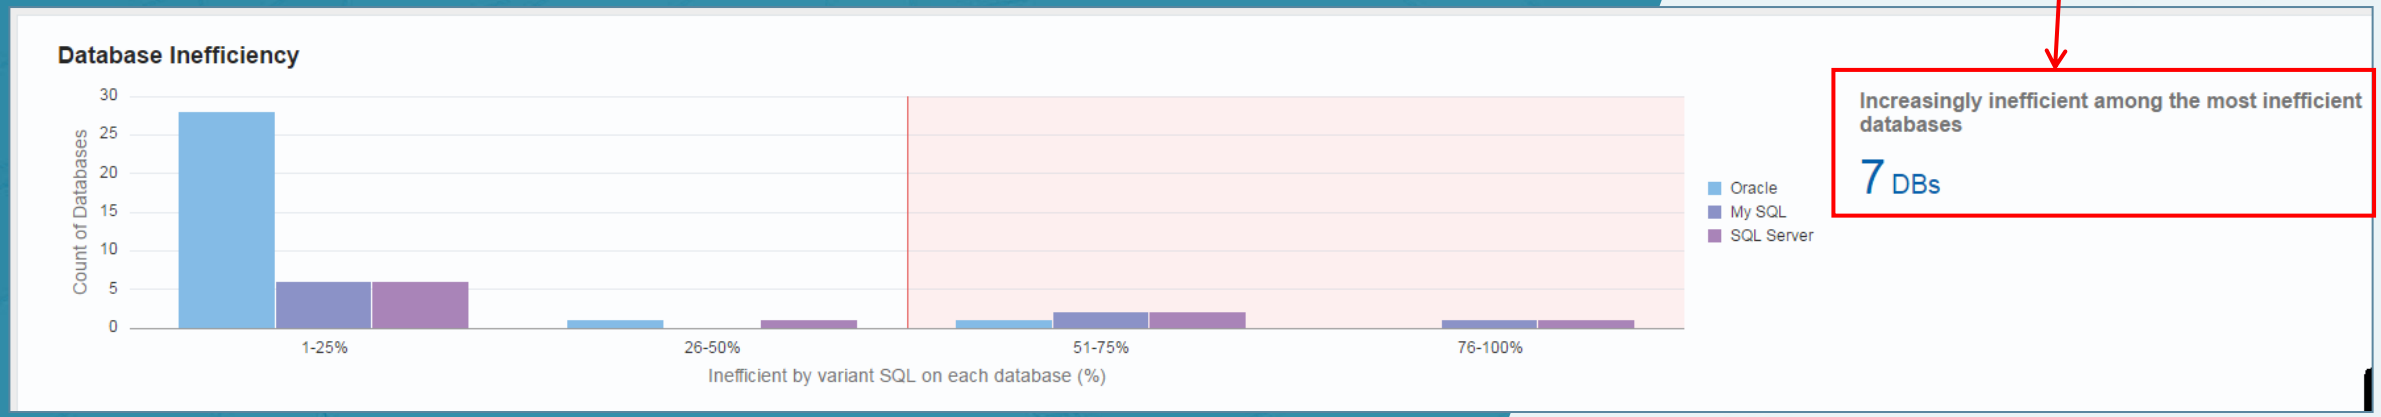

Performance Analytics

Analyze workload stability

- Predict the relative variability of a SQL over a period of time
- Focus on databases with variant workload performance

Databases with variant workload performance to focus on

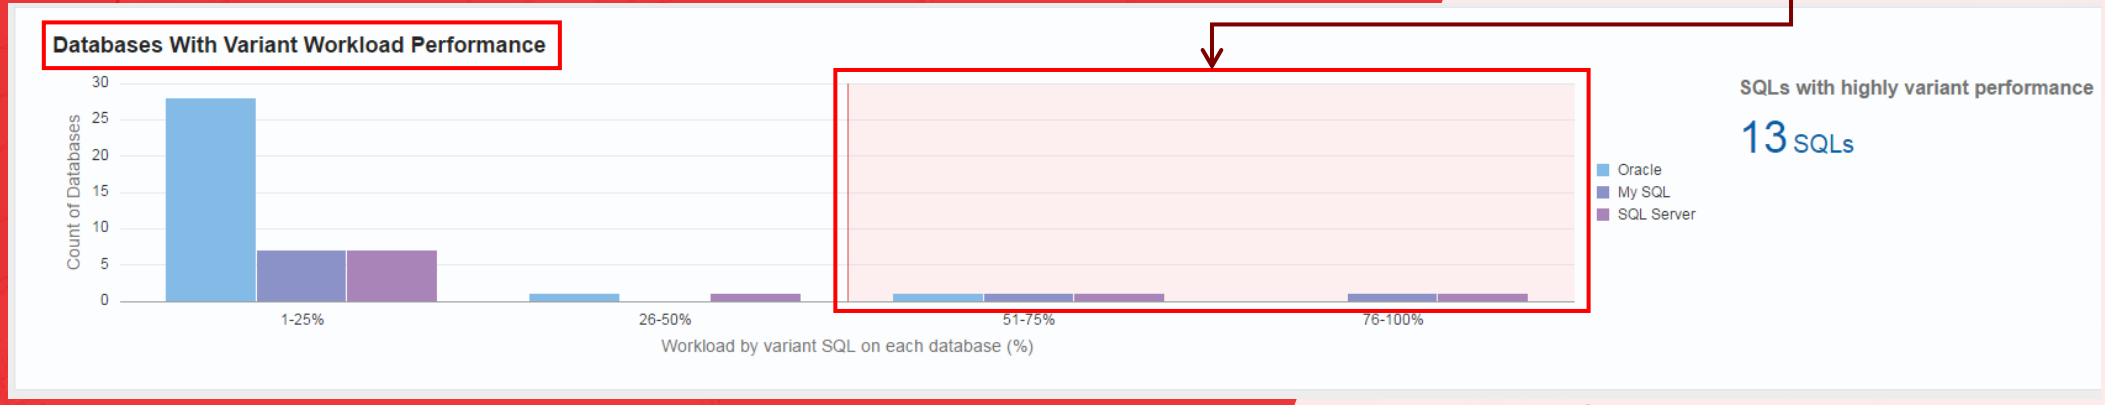

All databases with variant workload performance

Performance Analytics

Analyze SQL performance

- Analyze SQL performance based on long term ASH data and across entire fleet of databases
- Helps find systemic root cause of SQL performance issues
  - Compare performance trends across databases, drill-down to execution plans for analysis
  - Key insights into SQL performance issues
  - Trends and forecast SQL response time and other key SQL performance metrics
  - Was there a plan change for the SQL?

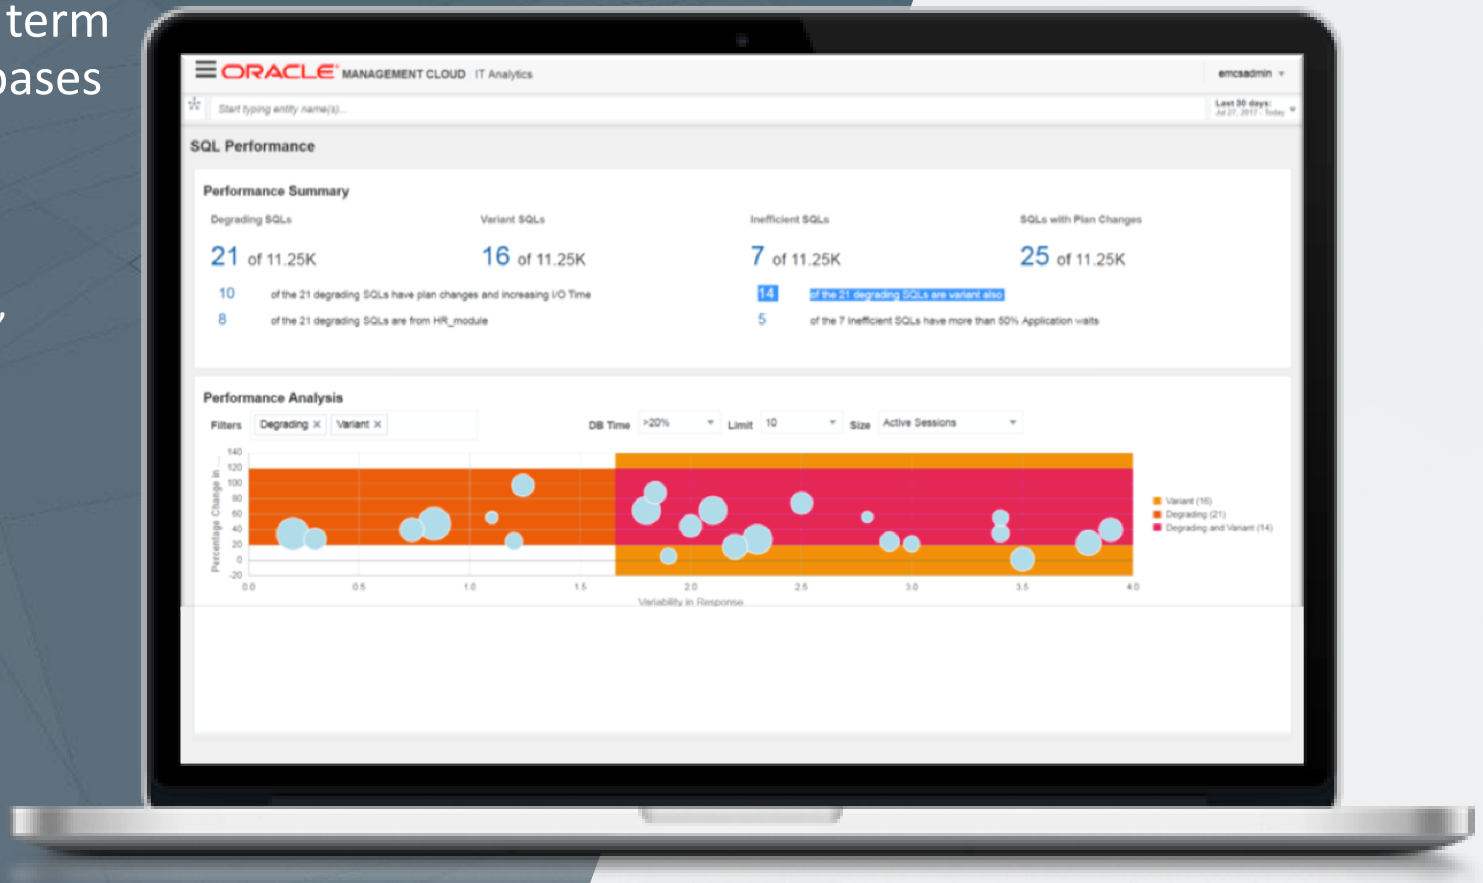

SQL Performance Analytics

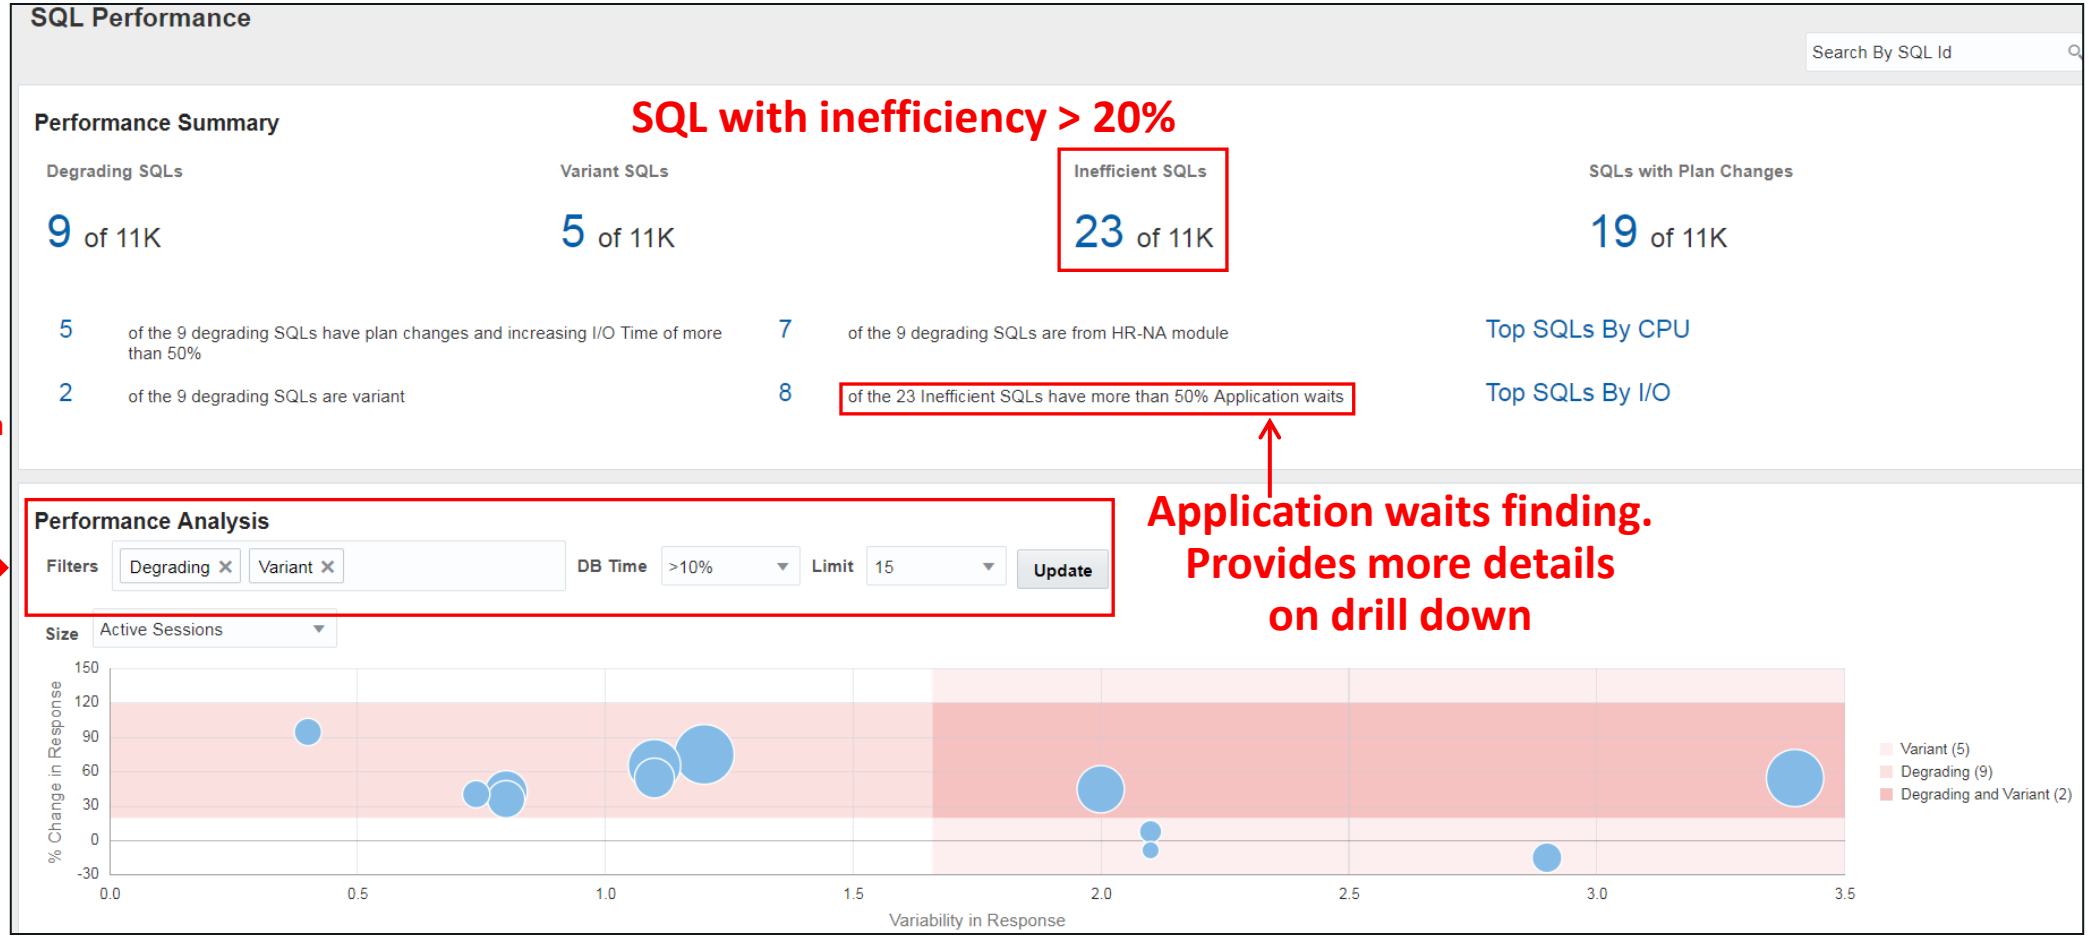

Use filters to explore and narrow down problem SQLs

Application waits finding. Provides more details on drill down

Deeper SQL performance analysis

SQL Performance Analytics

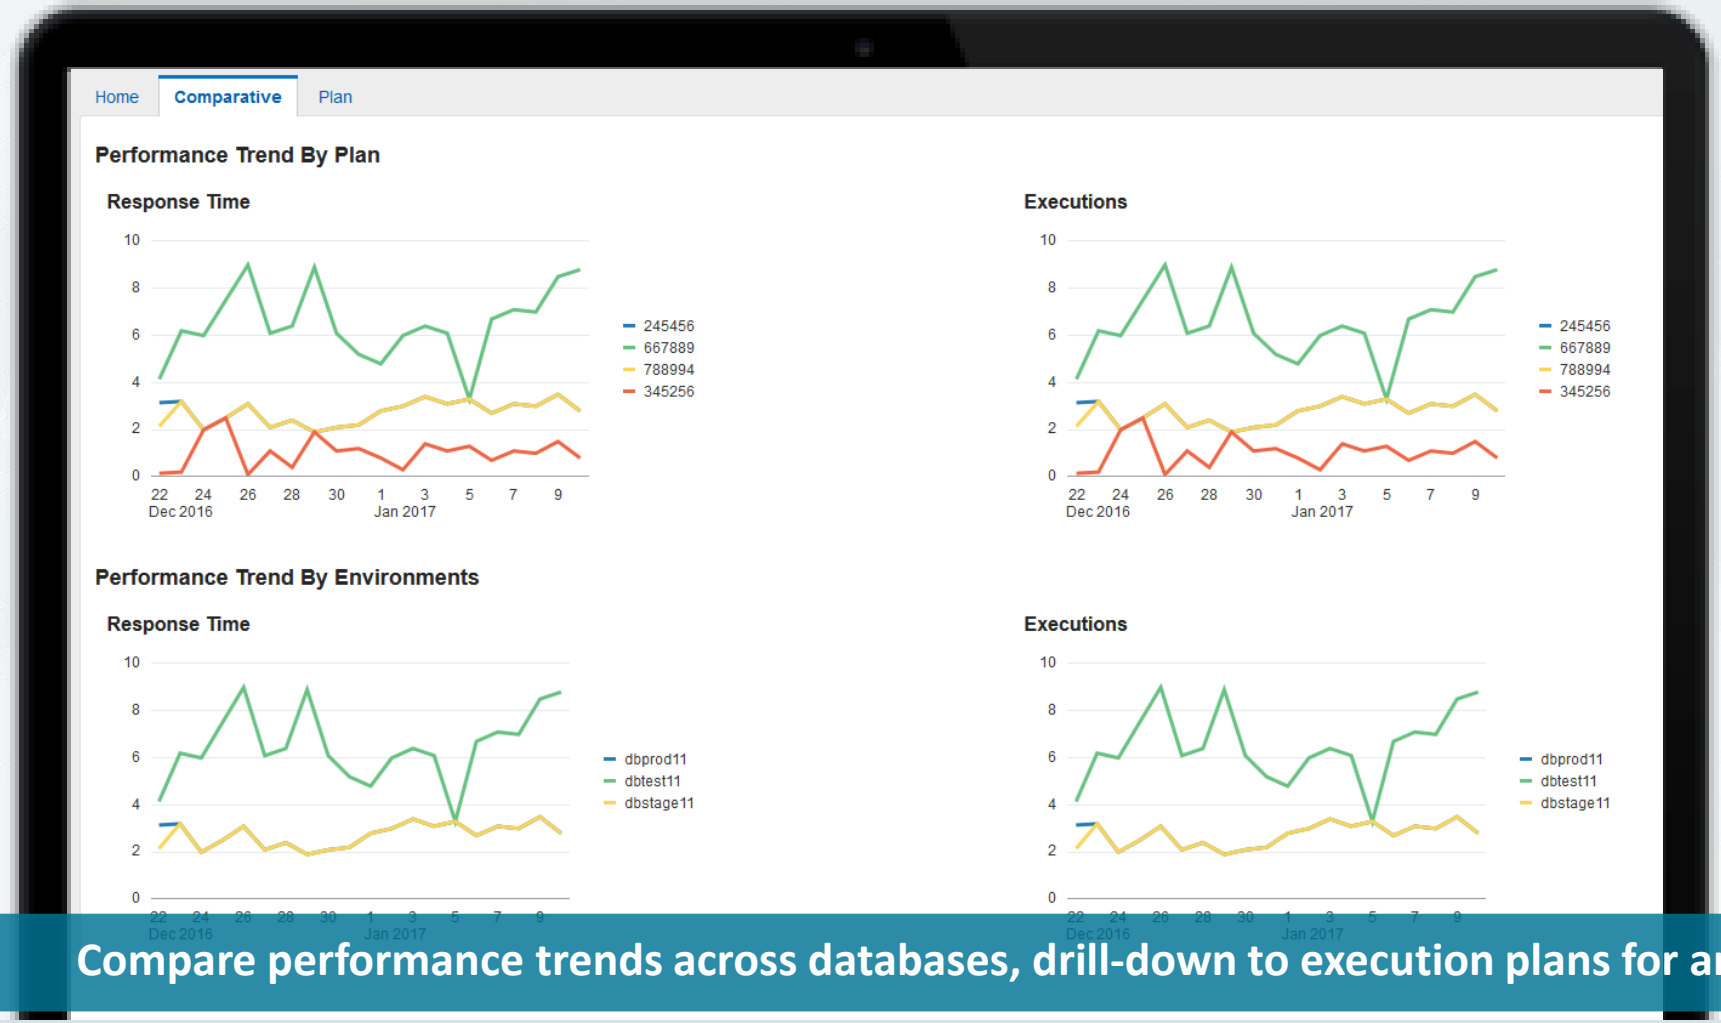

Compare performance trends across databases, drill-down to execution plans for analysis

SQL Performance Analytics

Identify Top SQLs by CPU and IO and GROUP BY various dimensions

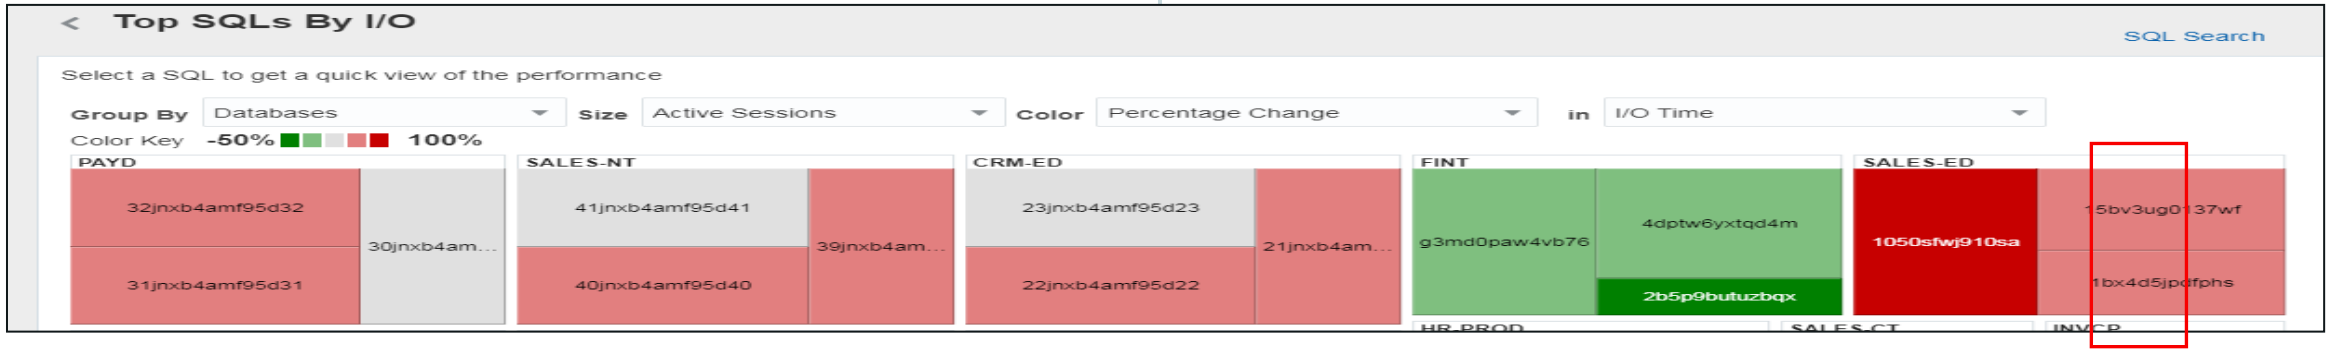

Utilization Analytics

- Analyze current utilization of the hardware & software
- Understand future resource requirements for host, VM based on current consumption
- Forecast future database resource requirements based on current consumption
- Consolidate additional databases or workloads on hosts

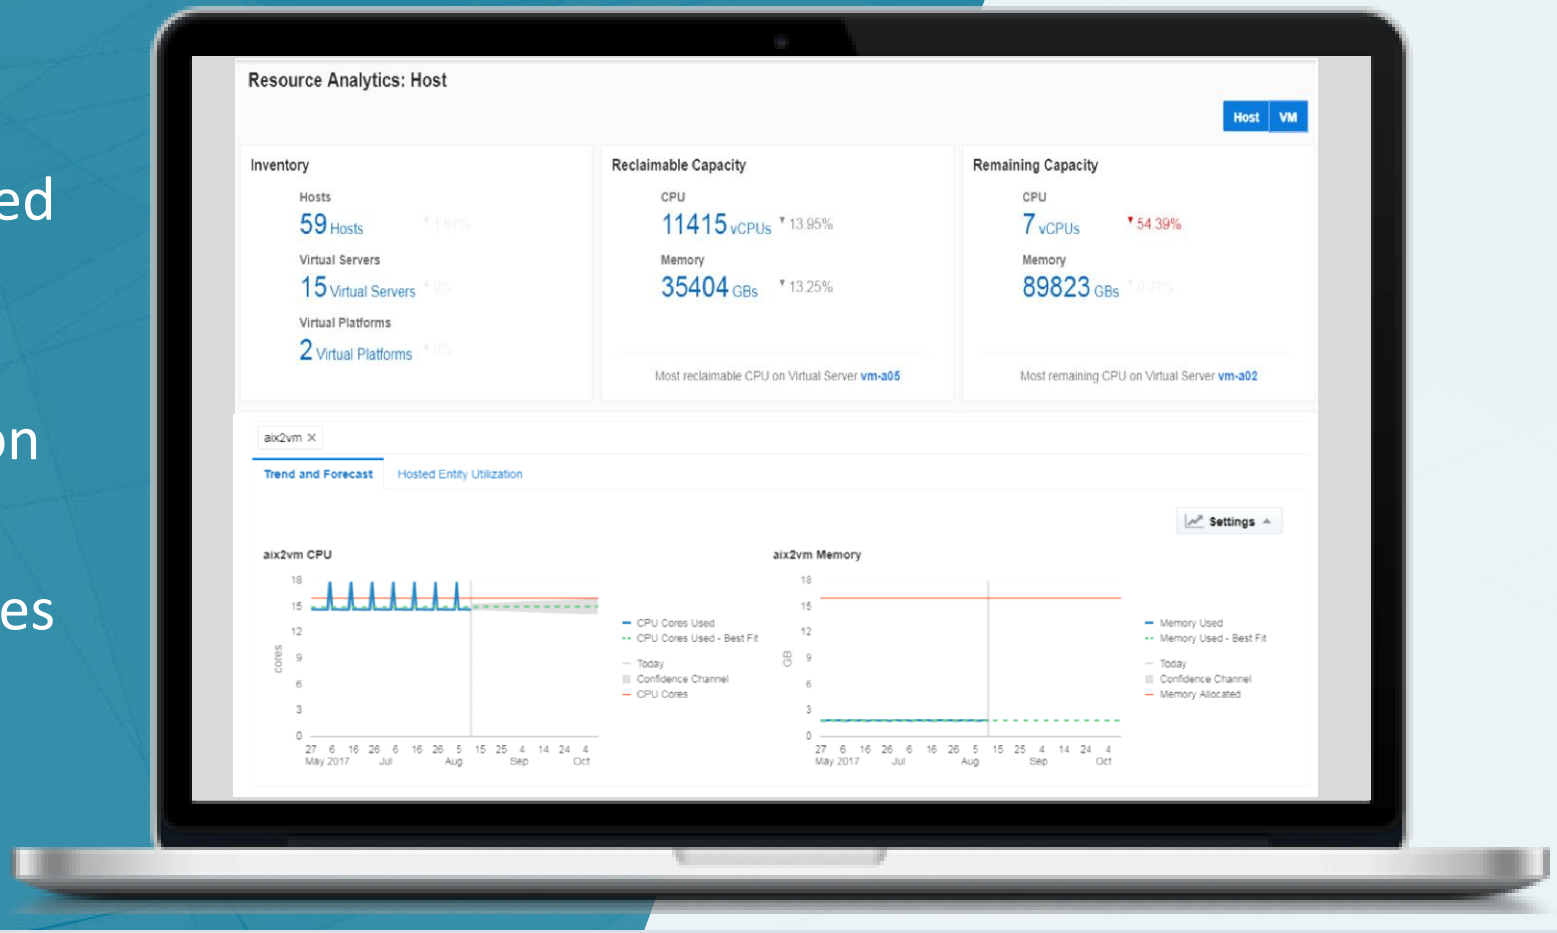

Ad-hoc Performance Analytics

ASH Analytics

- Performance activity data gathered sampling technique
  - Implemented in the database kernel for Oracle
  - Similar sampling technique for non-Oracle Databases
  - Efficiently handles volumes of data by down sampling appropriately
- Data sampled every 1 second
- DB time as a key measure and common currency cross databases for SQL statistics
- Active sessions for DB Activity based stats

Ad-hoc Analytics Demo

Ad-hoc Analytics

- I have a CDB with several PDBs in it
- DBA gets complaint from some users of poor performance
- Uses ASH Analytics to diagnose problem root cause

Reactive Performance Management

Analyzing transient performance problems using ASH Analytics

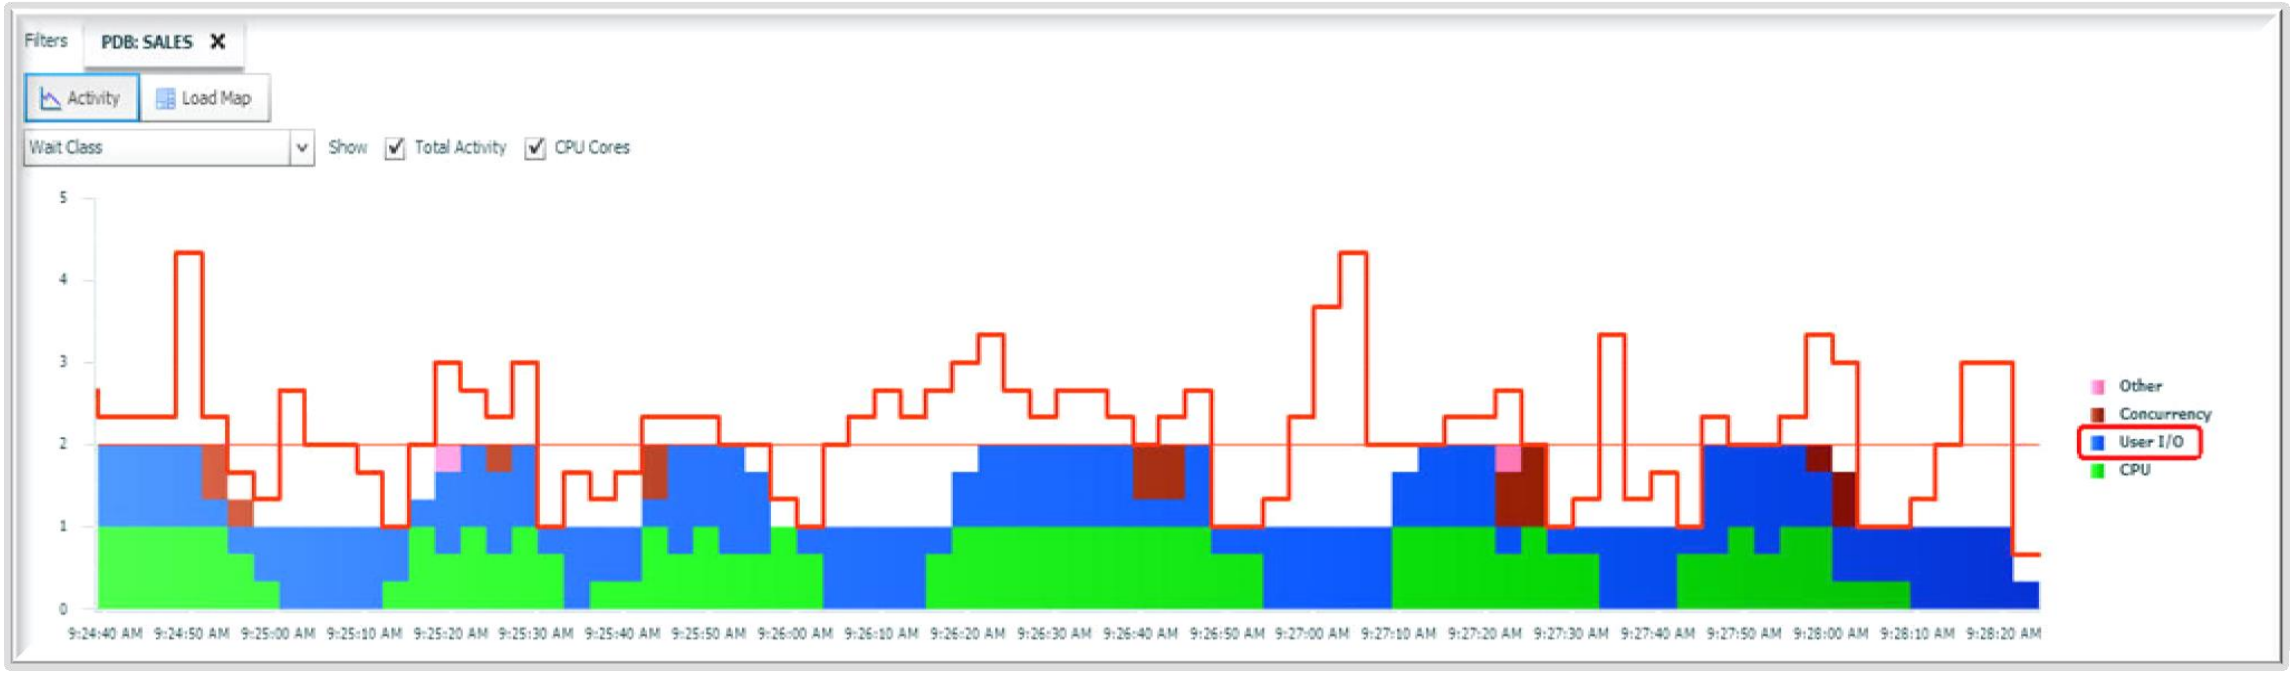

ASH Analytics identifies User I/O as the problem

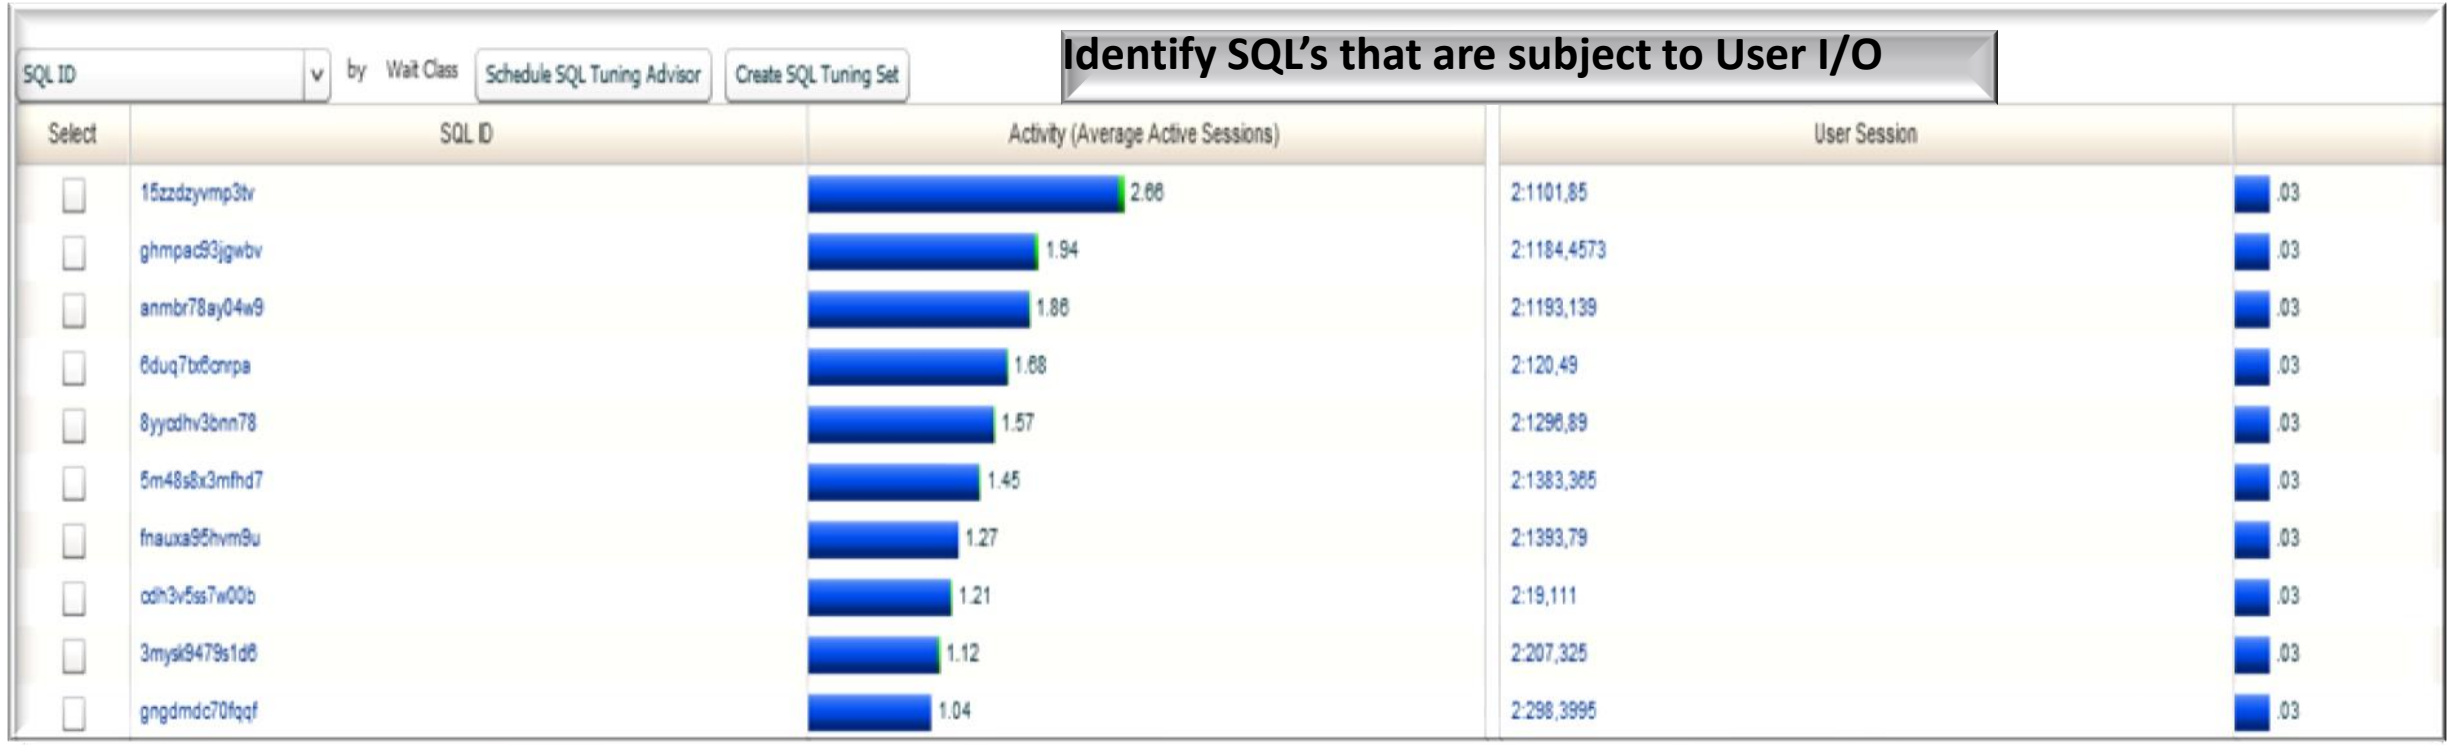

Anomaly Detection and Machine Learning

Smart Monitoring with Anomaly Detection

- Baselines automatically calculated to determine normal behavior
  - Tracks and adjusts for daily and weekly seasonality
    - For example, daily seasonality: Load from 9-10 am is expected to be higher than 6 pm – 7 pm
    - Weekly seasonality: Load on Mondays 9 am – 10 am is expected to be higher than Fridays 9 am – 10 am
- Auto-detection of performance anomalies
  - Anomaly = value outside expected baseline
  - Alert when performance is abnormal
  - Metrics from related entities, highly correlated with metric alert

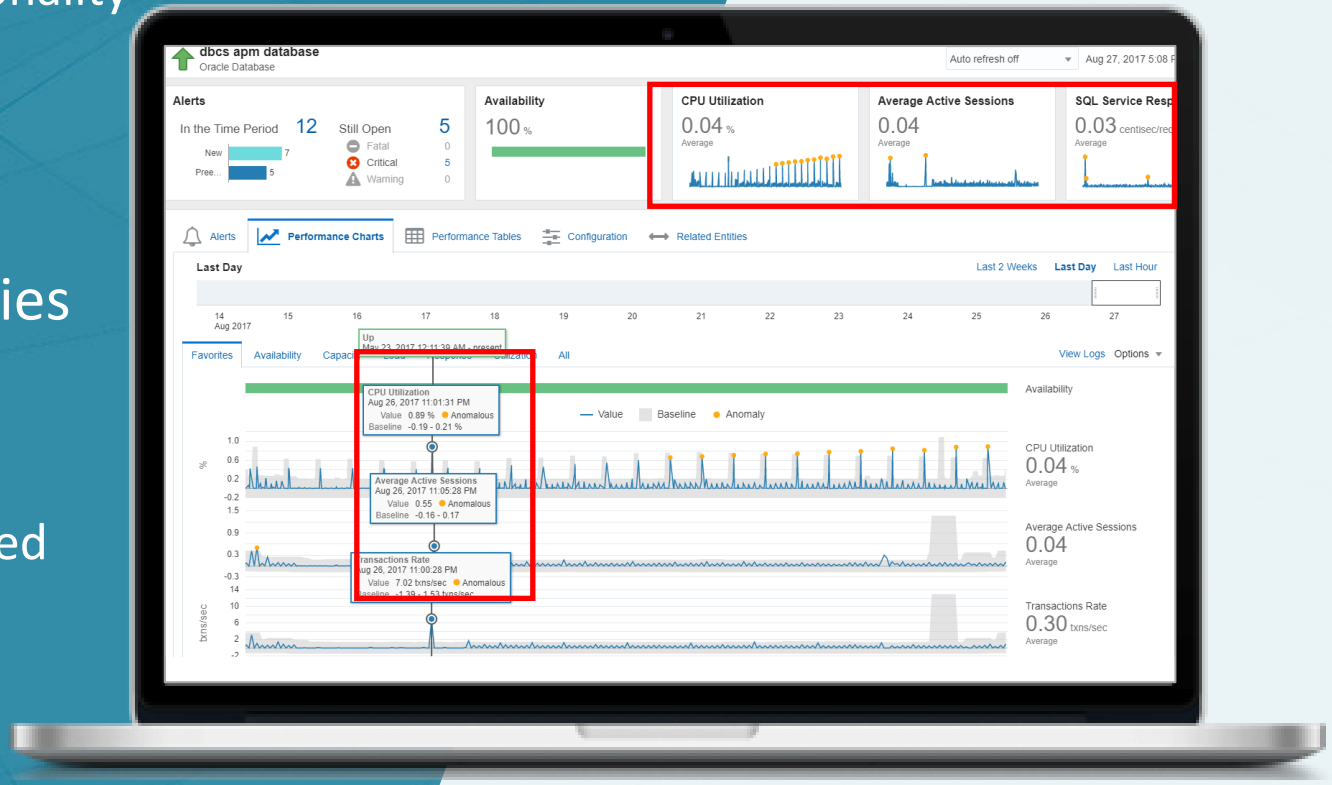

Diagnose Errors Using Machine Learning

Analyze and explore log data efficiently

- Specialized machine learning algorithms enable smart clustering of log events based on common signature
  - Reduce millions of log events into smaller set
  - Enables rapid troubleshooting by quickly identifying outliers and anomalies
  - Identify constants and variables in log records; bucketize variables for analysis
  - Topology-aware log exploration
    - Example: Cluster log records from application server and associated databases
  - Show trend of each clustered group; correlate clustered events that show similar trends
    - Example: Storage block corruption log records and DB write errors rising at the same time

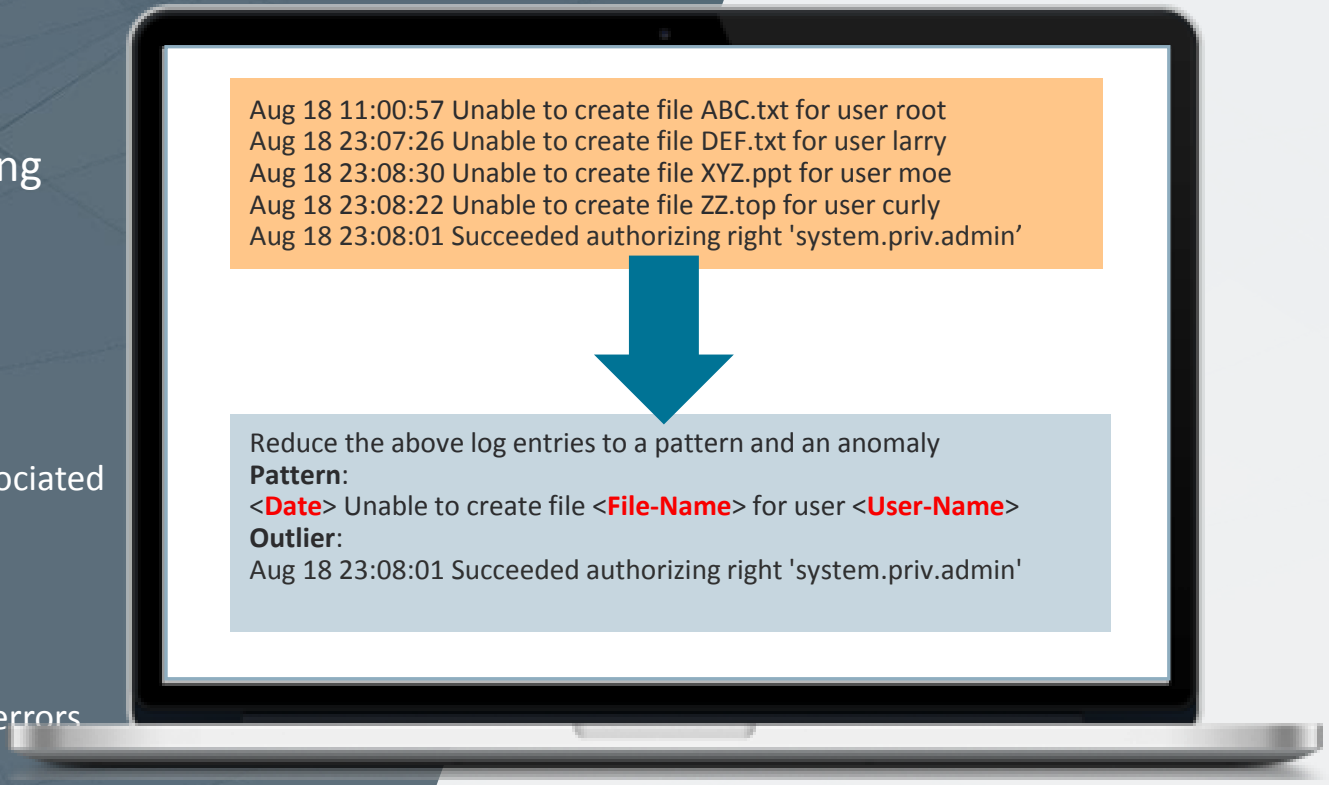


ORACLE®

Hybrid Cloud Management

Fleet Management

- Single pane of glass management view for all databases – on-premise and cloud
- Highlight active services, services with problems (alerts, failed backups), workload spikes, workflow status, etc.
  - Identify needle in the haystack - outliers based on CPU, I/O, memory and other important statistics
- Drill-down to respective single service pages for active management operations, troubleshooting and diagnostics
- Support for active operations at fleet level
  - E.g., set an init.ora parameter for all databases supporting an application

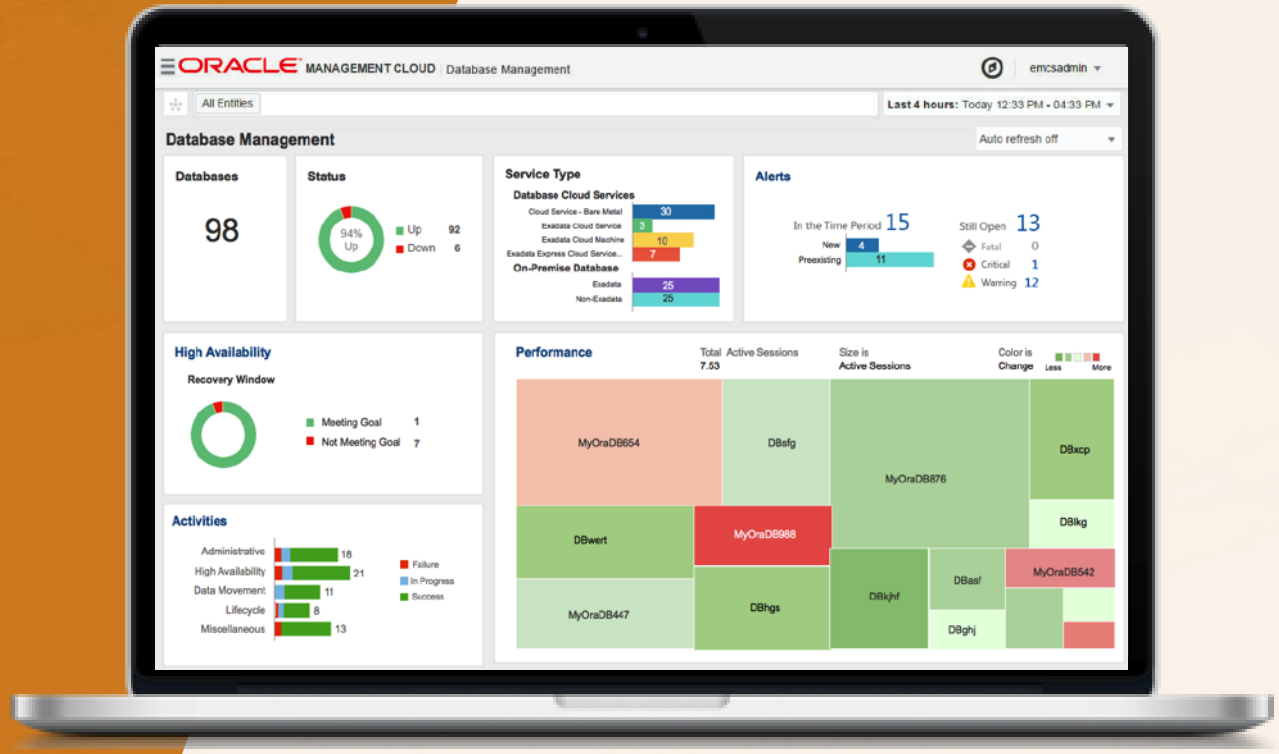

High Availability

- High availability fleet page
  - Protect your data by enabling business rules
  - Availability status of all managed databases
  - Backup/recovery related alerts for drill down
  - Identify databases that need attention – if meeting recovery window and unprotected data window goals
  - Supports Active operations at fleet level
    - Example: Backup all development databases as a single operation

Data Movement

Managed Workflows facilitate workload movement from on-premise to cloud

Deploy

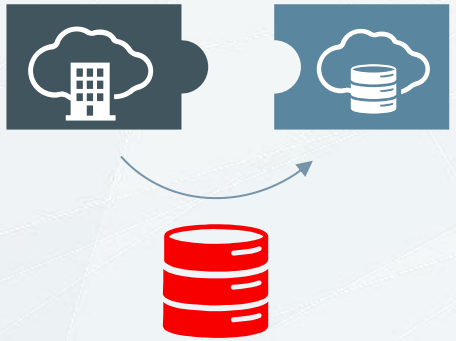

Move To Oracle Cloud

Lift and Shift workloads to Oracle Cloud with continuous and scheduled refresh of data from test master

Create

Test/Dev Creation

Clone production systems to cloud, track cloning lineage, data masking and TDE support for sensitive data

Lift and Shift from On-premise to Cloud

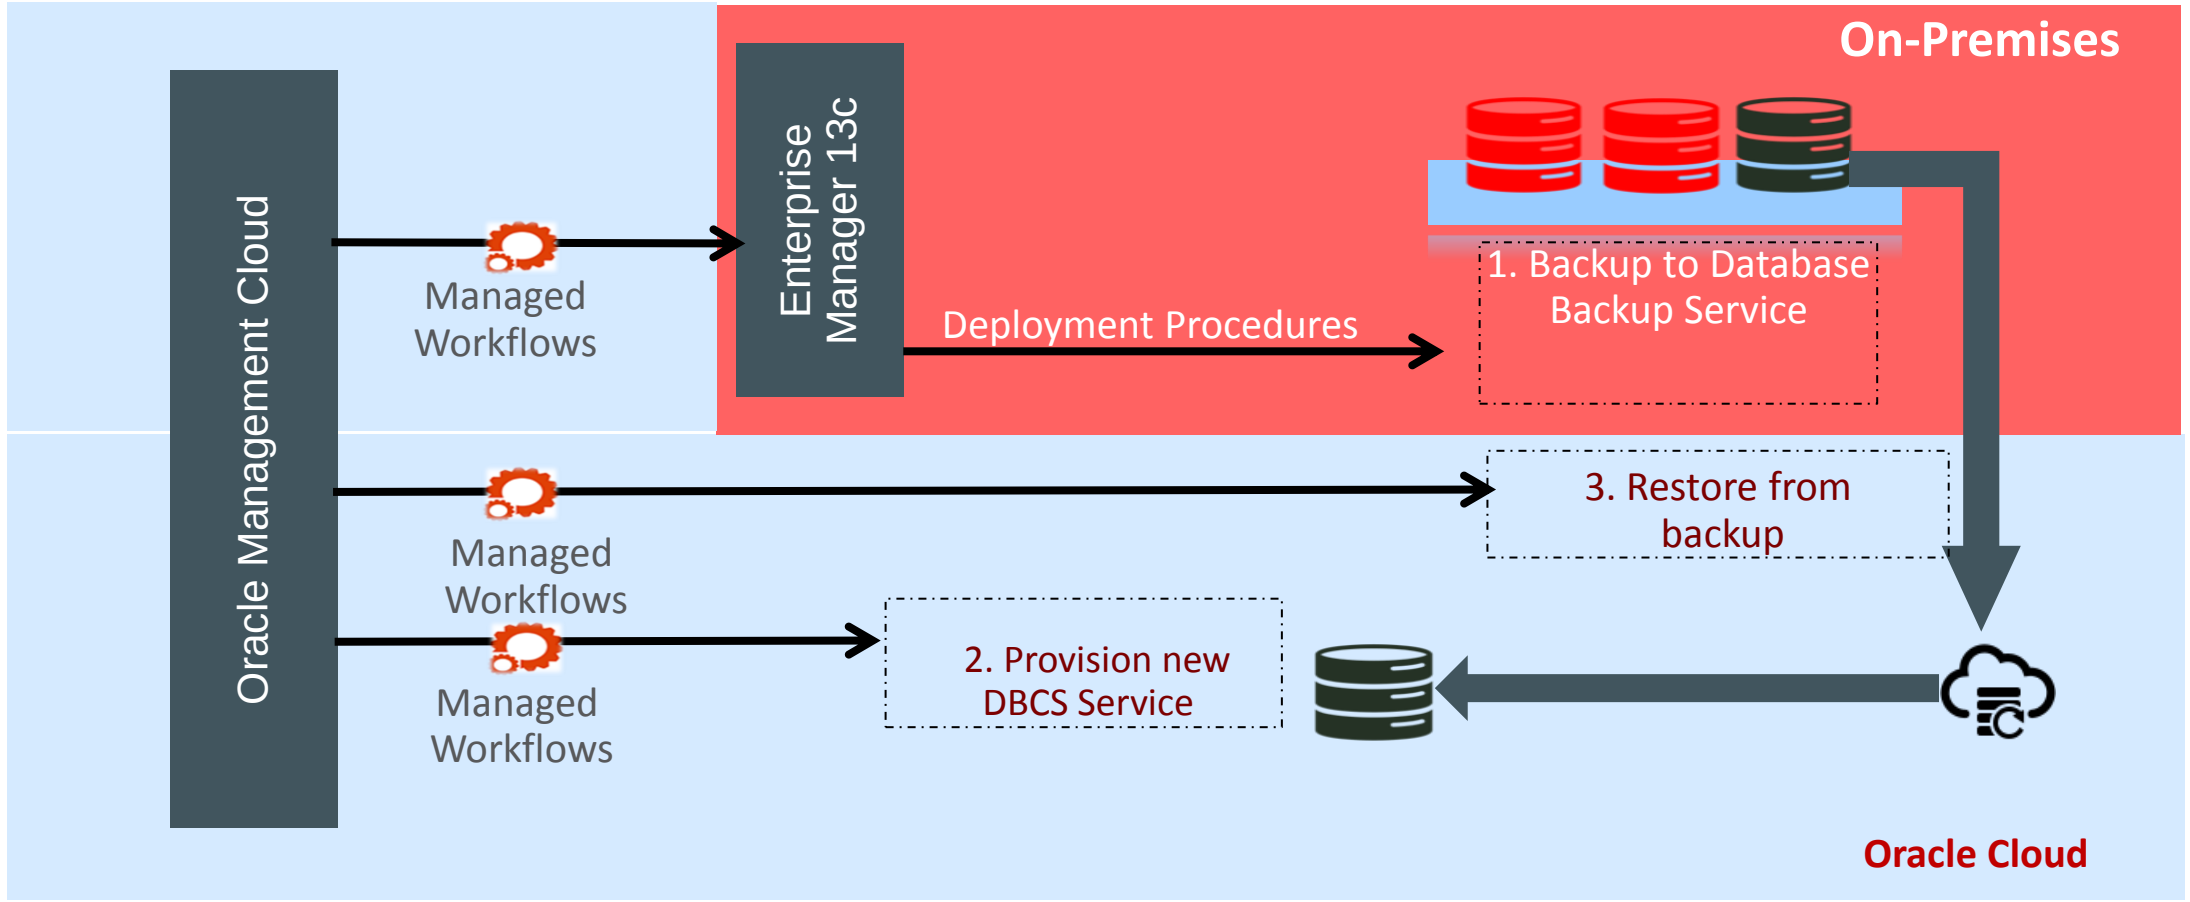

Test/Dev Creation and Data Refresh

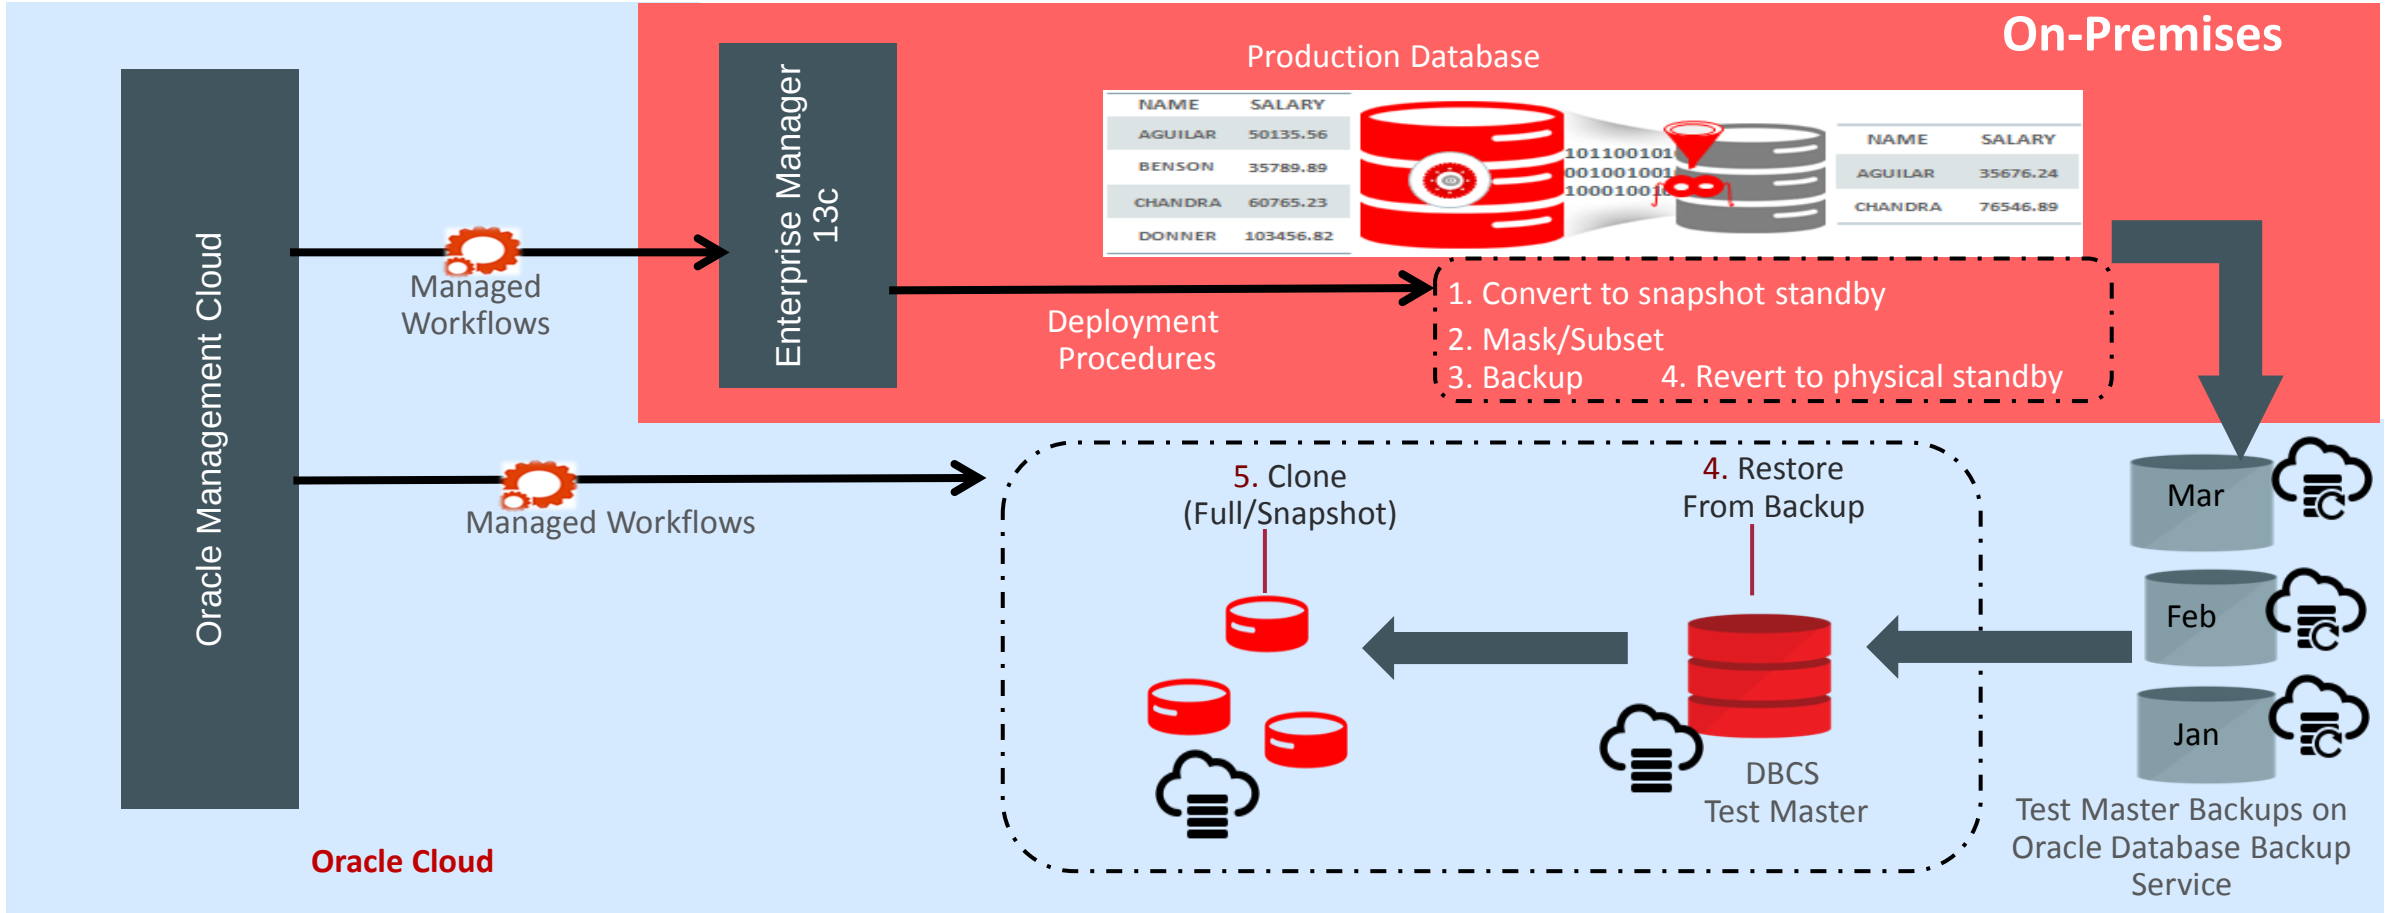

Modern DB Management Using ML and Analytics

Comprehensive Monitoring

Hybrid cloud management

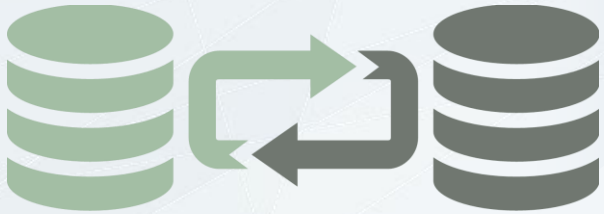

Gain 360-degree insight into the Oracle DBCS performance, availability & capacity planning

Designed to support IT operations, DevOps, LOB, and C-level executives

Real-time Analytics

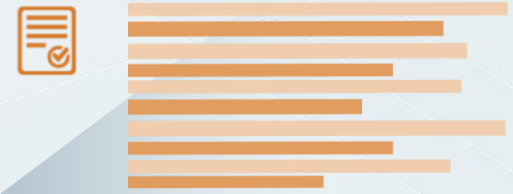

Anomaly Detection through ML

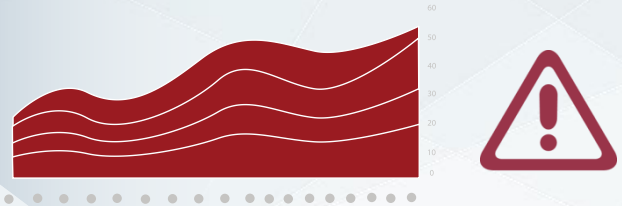

Hybrid Cloud Management

DB Management: Solution Overview

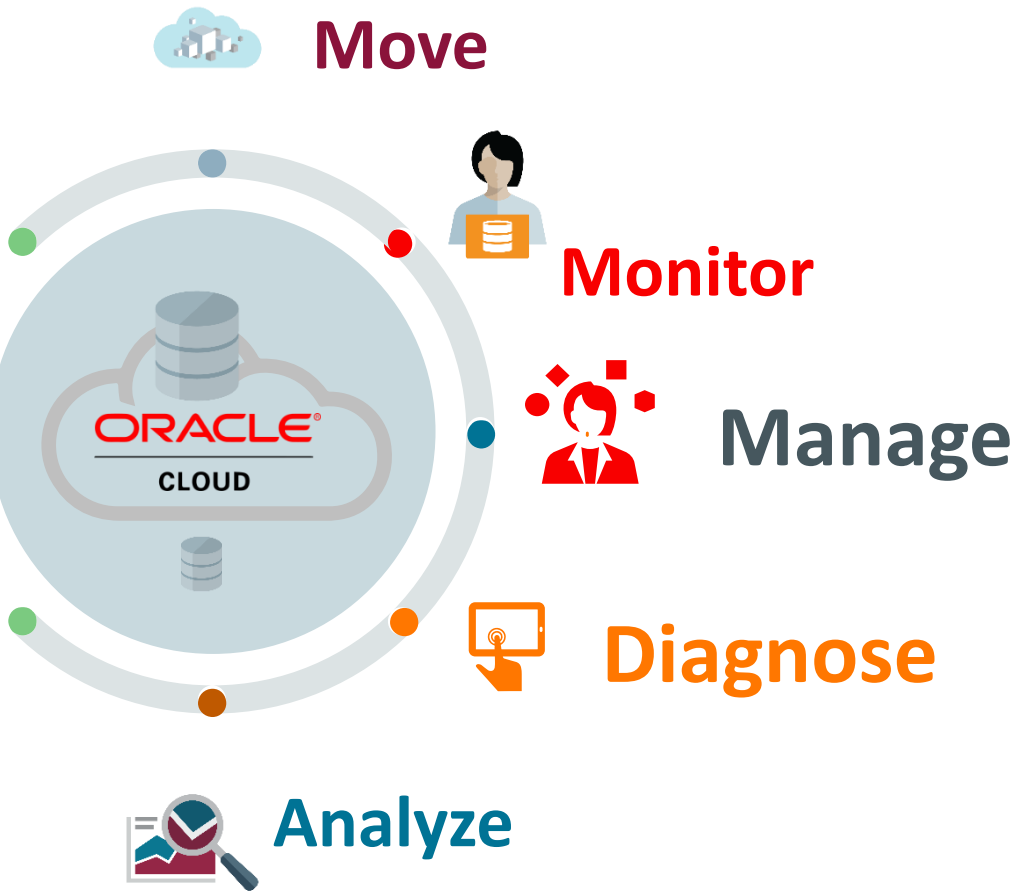

Analyze

- Performance Analytics for database
- Predict and plan for future demand
- Forecast future database resource requirements based on current consumption
- Reduce total cost of ownership

New Capabilities for Autonomous Database

- Move databases from on-premise to cloud using automated managed workflows
- Monitor key application metrics with alerting and anomaly detection
- Diagnose and improve application SQL logic
- Analyze logs, performance and utilization trends based on historical data