

AWR.ai: New LLM-Powered Investigations into Database Performance

—
Derik Harlow

Product Management – Observability and Management

April 30, 2026

Safe harbor statement

The following is intended to outline our general product direction. It is intended for information purposes only, and may not be incorporated into any contract. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. The development, release, timing, and pricing of any features or functionality described for Oracle's products may change and remains at the sole discretion of Oracle Corporation.

Agenda

1. AWR and LLM: Background and challenges
2. AWR.ai: GenAI powered solution
3. SQLPerf.ai: SQL performance AI analysis
4. Extensible ADDM
5. Conclusion

Critical Tool with Learning Curve

Requires domain expertise

AWR requires existing knowledge to utilize as a performance investigation tool:

- Complex system details
- Raw data, text only reading and interpretation
- Requires knowledge of detailed wait events
- Difficult to diagnose heavy workload
- Single issue can expand to other sessions making root cause difficult to identify
- Multiple AWR snapshots can spread across single problem
- Scoped investigation to single database and SQL load at the time
- AWR contains ADDM report but deep technical knowledge is still required

	Snap Id	Snap Time	Sessions	Cursors/Session
Begin Snap:	18314	01-Mar-25 14:00:02	19	2.7
End Snap:	18320	01-Mar-25 20:00:38	14	2.6
Elapsed:	360.60 (mins)			
DB Time:	2,944.21 (mins)			

Top 10 Foreground Events by Total Wait Time

Event	Waits	Total Wait Time (sec)	Avg Wait	% DB time	Wait Class
resmgr.cpu quantum	2,810,900	59.2K	21.08ms	33.5	Scheduler
DB CPU		23.3K		13.2	
db file sequential read	23,793,240	18.7K	784.64us	10.6	User I/O
log file sync	3,910,836	18.3K	4.68ms	10.4	Commit
db file parallel read	1,150,787	5576.6	4.85ms	3.2	User I/O

SQL ordered by Elapsed Time

- Resources reported for PL/SQL code includes the resources used by all SQL statements called by the code.
- % Total DB Time is the Elapsed Time of the SQL statement divided into the Total Database Time multiplied by 100
- %Total - Elapsed Time as a percentage of Total DB time
- %CPU - CPU Time as a percentage of Elapsed Time
- %IO - User I/O Time as a percentage of Elapsed Time
- Captured SQL account for 48.5% of Total DB Time (s): 176,653
- Captured PL/SQL account for 63.6% of Total DB Time (s): 176,653

Elapsed Time (s)	Executions	Elapsed Time per Exec (s)	%Total	%CPU	%IO	SQL Id	SQL Module	PDB Name	SQL Text
61,521.22	2,194,969	0.03	34.83	18.11	20.78	3n4tdaggd9b9r	JDBC Thin Client	ORCLPDB	BEGIN :1 := orderentry.neworde...
18,726.90	2,745,555	0.01	10.60	28.44	0.00	Zyhuvsmwggggn	JDBC Thin Client	ORCLPDB	BEGIN :1 := orderentry.browsep...
15,056.07	2,152,894	0.01	8.52	7.06	40.73	3fw75k1snsddx	New Order	ORCLPDB	INSERT INTO ORDERS (ORDER_ID,...
13,065.00	24,999,179	0.00	7.40	29.03	0.02	c13sma6rkr27c	New Order	ORCLPDB	SELECT PRODUCTS.PRODUCT_ID, PR...
9,995.03	274,970	0.04	5.66	7.45	49.71	ajjggrmacwv34		ORCLPDB	BEGIN :1 := orderentry.browsea...
9,018.07	31,801,884	0.00	5.10	32.98	0.00	0y1prvxgc2ra9	Browse Products	ORCLPDB	SELECT PRODUCTS.PRODUCT_ID, PR...
7,504.72	822,175	0.01	4.25	15.89	29.65	21ss57x3bykjp3		ORCLPDB	BEGIN :1 := orderentry.newcust...
6,706.63	547,397	0.01	3.80	10.39	41.55	06gapn9jpn8p1		ORCLPDB	BEGIN :1 := orderentry.updateC...
6,351.11	5,774,940	0.00	3.60	12.63	41.36	frxuxz264k87	New Order	ORCLPDB	INSERT INTO ORDER_ITEMS (ORDE...
5,890.61	2,217,478	0.00	3.33	7.04	47.56	g81cbrq5yamf5	New Order	ORCLPDB	SELECT ADDRESS_ID, CUSTOMER_ID...
5,481.22	255,452	0.02	3.10	4.64	59.16	7t0959msvy15g	Browse and Update Orders	ORCLPDB	SELECT ORDER_ID, ORDER_DATE, O...
4,683.59	479,595	0.01	2.65	5.97	49.66	8zz6y2yzdqjpp0	Update Customer Details	ORCLPDB	SELECT CUSTOMER_ID, CUST_FIRST...
3,481.55	46	75.69	1.97	9.26	6.99	cb5vgtz0hwc57	DBMS_SCHEDULER	ORCLPDB	SELECT NVL(SUM(TIME_WAITED/100...
3,369.09	274,728	0.01	1.91	9.47	40.19	8tt8far209h4n		ORCLPDB	BEGIN :1 := orderentry.process...
2,657.76	266,438	0.01	1.50	6.84	45.83	7hk2m2702ua0g	Process Orders	ORCLPDB	WITH NEED_TO_PROCESS AS (SELEC...
2,620.48	579,302	0.00	1.48	10.86	43.74	gh2g2tynpcpv1	Update Customer Details	ORCLPDB	INSERT INTO CUSTOMERS (CUSTOM...
2,089.98	194	10.77	1.18	0.01	0.00	0268x6996tvbv	DBMS_SCHEDULER	ORCLPDB	lock table wrs_optstat_histhe...



AWR.ai



AWR assistant

How Do We Make AWR More Approachable?

AWR.ai strives to turn AWR non-experts into experts.

AWR.ai offers several key features:

No More Complex Queries

Users no longer need to write intricate SQL statements to retrieve insights - just ask in plain language.

Enhanced User Experience

An intuitive and user-friendly interface allows interaction with AWR through conversational prompts or GUI interface.

Smarter, More Capable Answers

AWR.ai can handle complex, nuanced questions that traditional tools are unable to process or interpret, can map questions directly to underlying AWR views and even **correlate or compare multiple snapshot intervals.**

AWR.ai integrates with existing tools.

The screenshot displays the AWR.ai interface. On the left, a chat window shows a user asking for SQL with the highest regression for a specific time range. The AI response includes the SQL ID, previous and current plan hash values, and a bar chart showing a significant increase in average elapsed time.

AWR.ai

Powered by Oracle Generative AI. Check responses for mistakes.
Do not share any sensitive information with the chat agent.

Oct 9, 2025, 8:48:20 PM GMT-00:00

Show the SQL with the highest regression for time range Oct 9, 2025, 5:49:25 PM GMT-00:00 to Oct 9, 2025, 6:20:34 PM GMT-00:00

Oct 9, 2025, 8:48:20 PM GMT-00:00

SQL ID: [69w4xmrp0d1gy](#) has regressed.
Previous plan hash value: 1080033758
Current plan hash value: 4068247346

Oct 9, 2025, 8:48:36 PM GMT-00:00

Average Elapsed Time increased by 24,900.3% compared to the previous 7 days.

Plan Hash Value	Average Elapsed Time (s)
1080033758	~24
4068247346	~2400



What is AWR.ai?

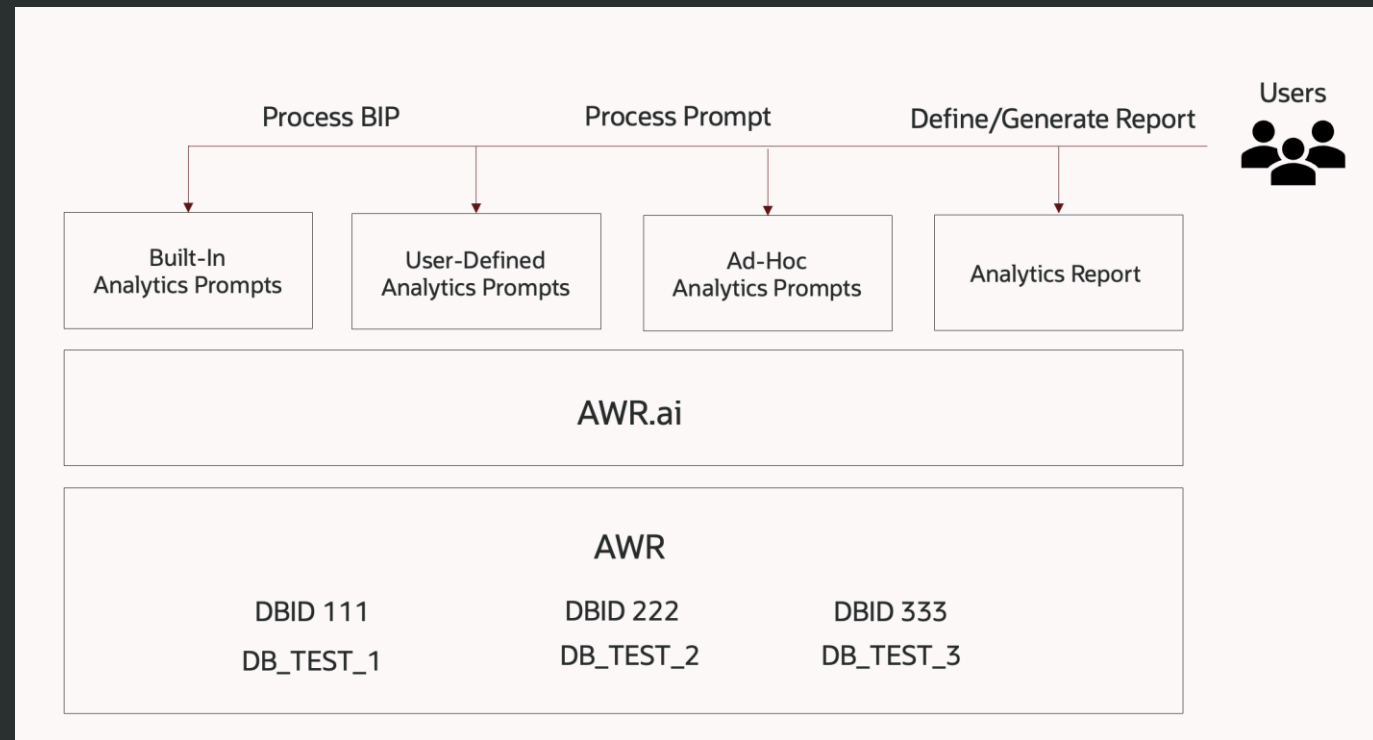
AWR.ai enables users to access powerful AI-driven AWR analysis through four core functionalities delivered in a database package:

Built-in Analytic Prompts (BIP): A collection of built-in analytic prompts to help users deep dive into AWR data. *Responses are backed by guaranteed accuracy.*

User-defined Analytic Prompts (UDP): Allow AWR experts to go beyond the Built-in Prompts by defining own prompts in natural language. Once a UDP is defined, it will be saved and reusable in the future.

Ad-hoc Analytic Prompts (AHP): Similar to User-defined Prompts, but it is one-time use and is not for future use.

Analytic Report: Allow users to create fully tailored reports by grouping multiple Prompts together. The prompts can be all types of Prompts mentioned above including BIP, UDP and AHP.



AWR.ai from Performance Hub

AWR.ai will be integrated into Performance Hub to enhance database monitoring capabilities.

- Interactive UI with seamless navigation to SQL Details
- AWR.ai renders output in a human-readable format including tabular data as well as charts
- Persists through navigation within Performance Hub for consistency
- Out-of-the-box prompts or natural language queries from the user

The screenshot shows the Oracle Enterprise Manager 24ai Performance Hub interface. The main content area displays the 'Performance Hub' section with a 'Quick Select' dropdown set to 'Custom' and a 'Time Range' of 'Oct 9, 2025, 2:13:00 PM - 6:35:00 PM'. Below this is a line chart titled 'Activity Summary (Average Active Sessions)' showing 'Maximum Threads' over time. The chart shows a peak in activity around 4:00 PM. Below the chart is a table titled 'Average Active Sessions' with columns for 'SQL ID', 'Activity (Average Active Sessions)', 'SQL Plan Hash', 'SQL Type', and 'Service'. The table lists several SQL IDs with their respective activity levels. On the right side of the interface, the 'AWR.ai' chat window is open, displaying a greeting: 'Hello, AWR.ai is an AI-based assistant that can answer your questions related to AWR data.' Below the greeting are several suggested prompts: 'Show the total number of snapshots', 'Show IO stats from load profile', 'Show time model for DB time per snapshot', and 'Show the SQL with the highest regression'. The chat window also includes a search bar and a 'Type a message' input field.



SQLPerf.ai



AI-assisted analysis of SQL performance

Why SQL Performance Tuning Is Hard

- **Complex queries** with multiple joins, subqueries, and nested views are hard to analyse.
- **Execution plans** require deep knowledge of data cardinality, join order, and optimizer behaviour.
- **Changing data volumes** and distributions can degrade previously good plans.
- **Indexing decisions and optimizer hints** are tricky and can backfire if misused.
- **Diagnosing** waits, latches, I/O patterns, and dynamic SQL adds complexity.
- Correlated subqueries and functions in **WHERE** clauses hinder efficient execution.
- **Limited visibility across** environments makes tuning repeatable only with difficulty.



Analyze and Improve SQL Performance with SQLPerf.ai

Obtain recommendations about SQL execution to improve its performance:

Integration with SQL Monitoring:

- SQL Monitoring Identifies poor performing or most resource intensive queries
- Quickly launch AI analysis of query performance
- Utilize with existing statistics in Performance Hub for guided resolution

The screenshot displays the Oracle Enterprise Manager 24ai interface for SQLPerf.ai. The main panel shows 'Real-time SQL Monitoring for SQL 8gux7xcyxu2aa'. The 'General' section indicates the query is 'Completed' and provides details like 'SQL Text: WITH /*+ monitor */ avg_by_prod AS (SELECT pr...', 'Execution Started: Oct 9, 2025, 8:48:18 PM GMT-00:00', and 'User Name: BADPLAN_DEMO@CDB1_PDB1'. The 'Time & Wait' chart shows a total duration of 3.42 m, with Database Time at 3.52 m and Wait Activity at 100%. The 'I/O' chart shows 12 K Buffer Gets, 33 K I/O Requests, and 6.7 GB I/O Bytes. The 'Plan Statistics' table lists operations such as 'SELECT STATEMENT', 'COUNT STOPKEY (SEL\$3)', 'VIEW (SEL\$CA0CB92C)', 'FILTER (SEL\$CA0CB92C)', 'SORT GROUP BY', 'HASH JOIN', and two instances of 'TABLE ACCESS FULL (SEL\$CA0CB92C)'. The 'SQLPerf.ai' panel on the right provides AI-generated insights, including a question: 'Q: Can a simple stats refresh fix the 1M row under-estimate?' and a response: '- Yes - very likely. When the optimizer sees "IS NOT NULL" it relies on the NUM_NULLS column statistic. If statistics are stale (NUM_ROWS & NUM_NULLS do not describe current d selectivity is mis-calculated.)'. It also lists two remedies: '1. Refresh table statistics with a decent sample size.' and '2. Add (or validate) a NOT NULL constraint on PRODUCT_ID.' Sample code is provided at the bottom.

Operation	Object	Information	Line ID	Timeline	Executions	Est. Rows
SELECT STATEMENT			0			1
COUNT STOPKEY (SEL\$3)			1			1
VIEW (SEL\$CA0CB92C)			2			1
FILTER (SEL\$CA0CB92C)			3			1
SORT GROUP BY			4			1
HASH JOIN			5			1
TABLE ACCESS FULL (SEL\$CA0CB92C)	FACT_SALES		6			1
TABLE ACCESS FULL (SEL\$CA0CB92C)	FACT_SALES		7			1



Powerful Insights to SQL Execution Details

Obtain recommendations for execution plan:

- Root cause analysis for poor performing queries
- Identifies bottlenecks in execution plan
- Provides feedback on performance improvement such as join orders, index utilization, statistics, etc.
- Recommendations provided for improved plan performance
 - Detailed resolution plan to improve query execution
 - Guided resolution with code

The screenshot displays the SQLPerf.ai interface. The main window shows a query execution plan analysis for 'Step 10 - HASH GROUP BY'. The analysis includes a question: 'Q: Can a simple stats refresh fix t', followed by an answer: 'Yes - very likely. When the optimizer uses stale statistics (NUM_ROWS & ...). Two most relevant remedies are listed: 1. Refresh table statistics with ... and 2. Add (or validate) a NOT NULL constraint. Sample code is provided for both remedies.

SQLPerf.ai

Sep 30, 2025, 10:14:07 PM GMT+05:30

Oct 9, 2025, 8:55:19 PM GMT-00:00

Plan Line 6 – TABLE ACC

Q: Can a simple stats refresh fix t

• Yes – very likely. When the optimizer uses stale statistics (NUM_ROWS & ...)

Two most relevant remedies

1. Refresh table statistics with ...
This updates NUM_ROWS and ...
2. Add (or validate) a NOT NULL constraint ...
Once the column is declared N ...

Sample code

```
-- 1. Refresh statistics
exec dbms_stats.gather_table_stat
  ownname => 'BADPLA
  tabname  => 'FACT_S
  estimate_percent => 100
  method_opt => 'FOR ALL

-- 2. Enforce NOT NULL if the dd
alter table BADPLAN_DEMO.FACT_SA
  modify (product_id not null
```

Step 10 – HASH GROUP BY

Fixing the wrong join order

1. Gather fresh and more detailed statistics (histograms and column-group/extended stats on the join and group-by columns) with DBMS_STATS so the optimizer stops under-estimating the left side of the join.
2. Create a SQL Plan Baseline after one correct execution (with proper join order) so future parses reuse the proven plan even if estimates drift.

Extensible ADDM



Custom extension rules for ADDM

Challenges With Custom Scripts for Performance Analysis

- DBAs often rely on **custom scripts** for performance analysis.
- **Inconsistent formats** make it hard to compare or consolidate results.
- **Manual scheduling and orchestration** increase maintenance overhead.



Extend ADDM for your own Database

Extensible ADDM will allow DBAs or developers to write their own analysis logic and register them as ADDM rules

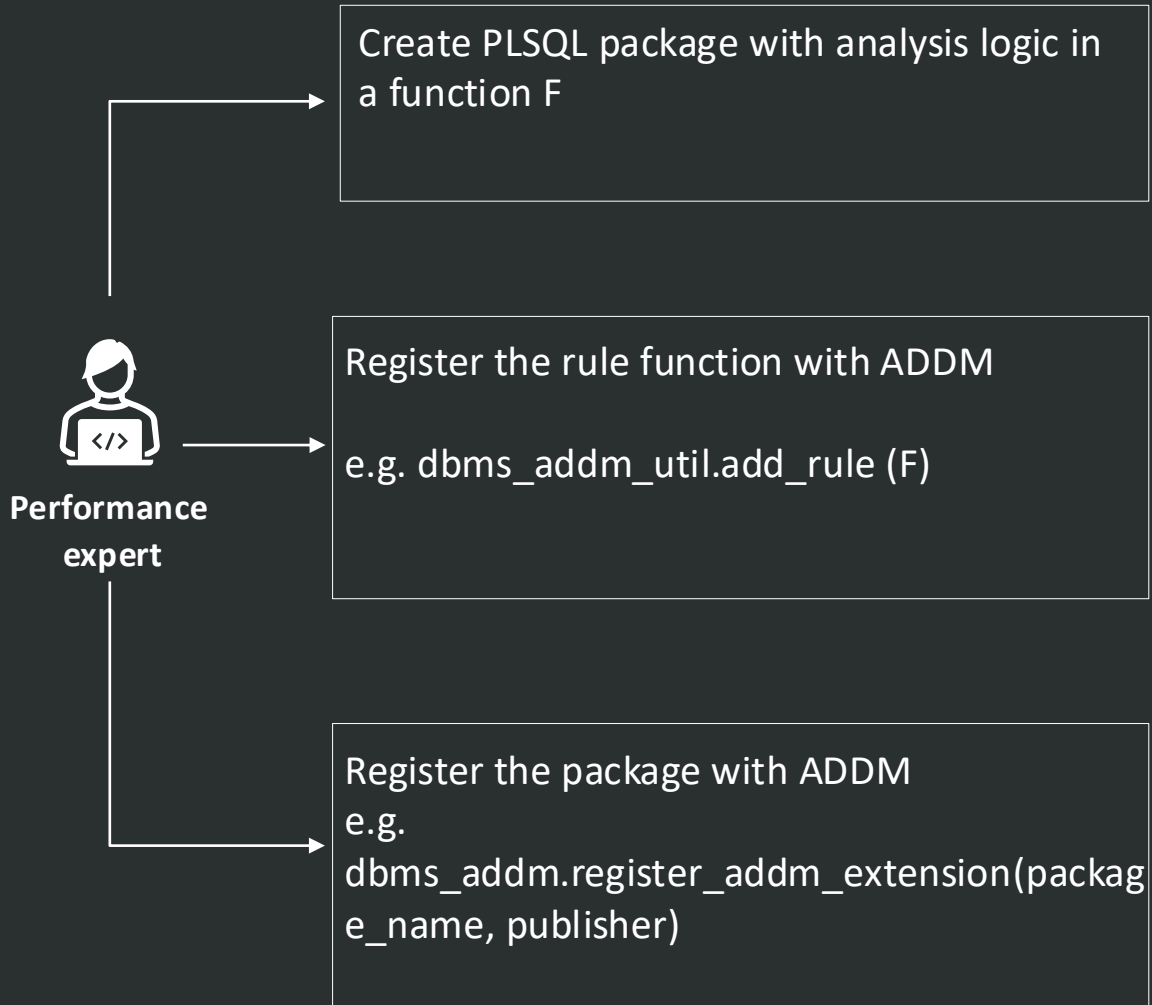
- Built-in library of PL/SQL APIs to help write extension rules
- Reduction of custom code deployments

Extension rules are enabled by DBA to be run automatically by ADDM with every AWR snapshot

- Extension rules will run as the DB user that registers them, not as SYS
- Strict time limits will be enforced for rules
- Findings and recommendations from extension rules will appear in existing ADDM report
- Identified with special “Publisher” attribute to distinguish from built-in rules provided by Oracle

Extensible ADDM

Custom extension rules for ADDM



Finding 1: Usage of DBMS_LOCK

Publisher: PerfDBA

Impact is 17.24 active sessions, 99.48% of total activity.

Waits from DBMS_LOCK package were consuming significant database time.

Recommendation 1: Application Analysis

Estimated benefit is 17.36 active sessions, 99.48% of total activity.

Action

Investigate application logic using DBMS_LOCK package. Consider using finer granularity of locks.

[Sample ADDM report](#)

AskEM



GenAI Assistant

Navigating Complex Observability Environments

- **Fleet-level complexity:** Analyzing performance across large-scale deployments is difficult.
- **Scope definition:** It's hard to narrow down which systems or components to investigate first.
- **Outlier detection:** Identifying performance anomalies across Exadata or shared infrastructure takes time and expertise.
- **Drill-down challenges:** Moving from fleet-level views to pinpointing specific database or instance issues is often cumbersome.
- **Telemetry sources:** Data comes from databases, hosts, Exadata, applications, and more.
- **Deployment flexibility:** Collects data from on-premises, multicloud, and hybrid environments.



Observability Telemetry and Challenges

Problem statement:

Heterogenous observability telemetry

- Databases, hosts, Exadata infrastructure, applications, etc.

Diverse telemetry origins

- On-premises, multicloud, hybrid environments

How do you leverage power of AI to provide rich user experience for self-diagnostics?

AskEM is the solution:

- Enable users to ask questions conversationally (Natural Language Processing)
- Leverage telemetry data to provide contextually relevant, reliable, and efficient responses
- Dynamically generate metric widgets in a dashboard framework
- Deliver capabilities for both on-premises and cloud EM deployments
- Streamline search of product documentation, knowledge bases, etc.

Solution: Observability GenAI Assistants

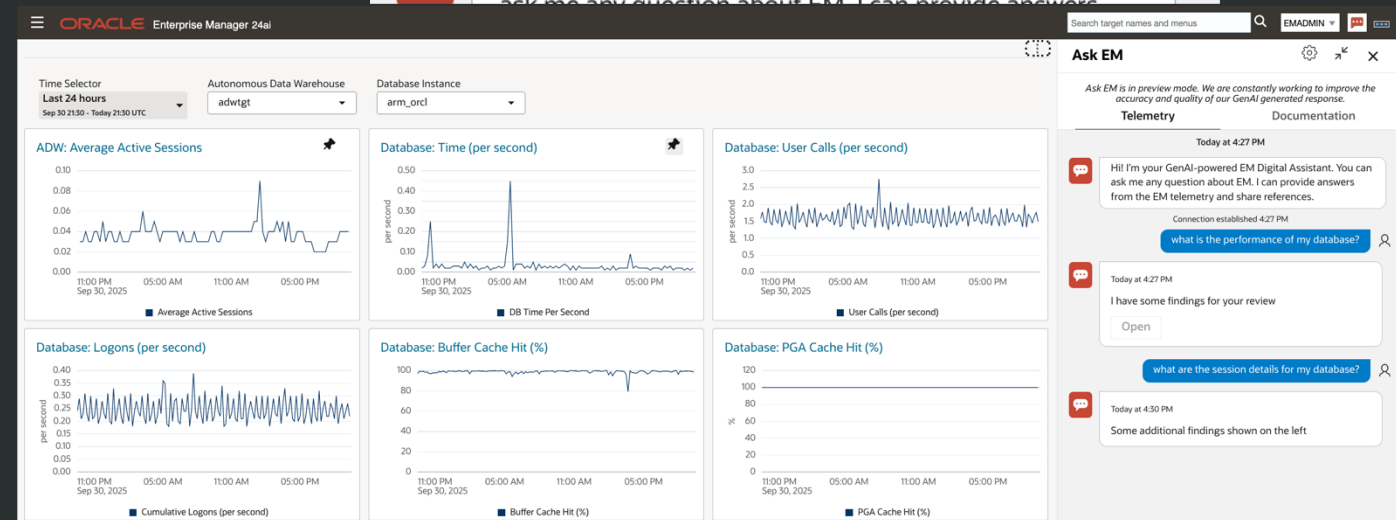
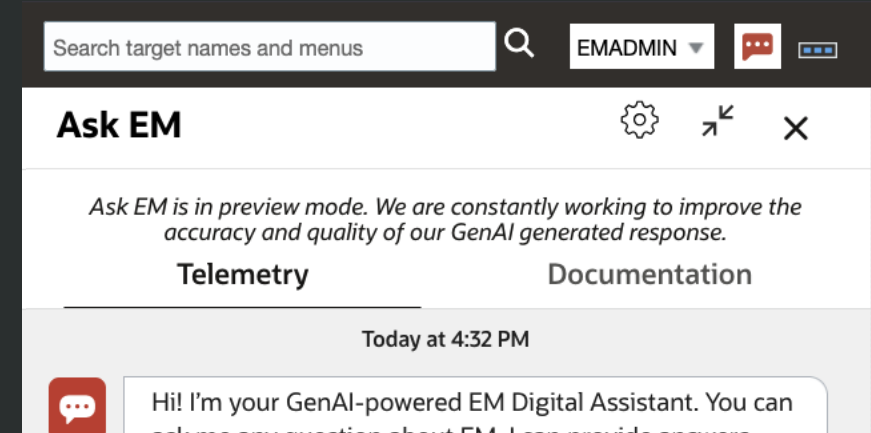
Deliver fast and reliable insights

State-of-the-art AI assistants for EM and OCI Observability & Management

- Advanced AI-driven chatbot technology
- Augmented with metrics metadata and Oracle documentation
- Trained on EM repository metadata

Enhanced user experience

- Native language interface to the observability telemetry
- Delivers quick and relevant responses
- Shorten learning curve for new users, reduce reliance on subject matter experts



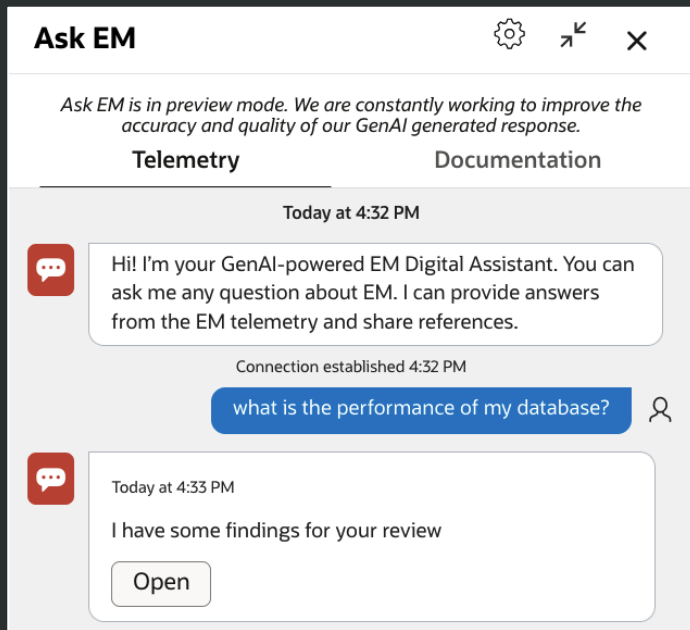
Some GenAI-powered observability use cases

Ask in natural language and get AI-generated responses



Troubleshooting of app & infrastructure performance, errors, availability, and more

Q: Why is the manufacturing database slow?




Summarization of results

Q: What kind of Exadata I/O patterns do I have?



Dashboard generation

Generates metric widgets for simplified data exploration



Product usage help

Covers frequent user actions like agent deployment, config changes, API lookup, etc. via Oracle documentation

Q: Can I get information on buffer busy waits?



Q&A

Learn More

Web: oracle.com/enterprisemanager

Videos: youtube.com/OracleEnterpriseMgr

Blogs: blogs.oracle.com/observability

Docs: <http://docs.oracle.com/en/enterprise-manager>

[Try it now](#)



Hands-on-labs

Oracle Cloud Free Tier

Always Free

Services you can use for unlimited time



30-Day Free Trial

Free credits you can use for more services

www.oracle.com/cloud/free