INTELIGENCIA ARTIFICIAL
MÁS ALLÁ DE LA PALABRA DE MODA

Miguel Balsa
SOLUTIONS ENGINEER
DYNATRACE

dynatrace
IMMERSE

AGENDA

- - Artificial Intelligence **-** The Big Picture
- - Causal AI
- - **Hypermodal AI** for Unified Observability and Security
- - **Dynatrace Davis AI -** Use cases

AI: THE BIG PICTURE

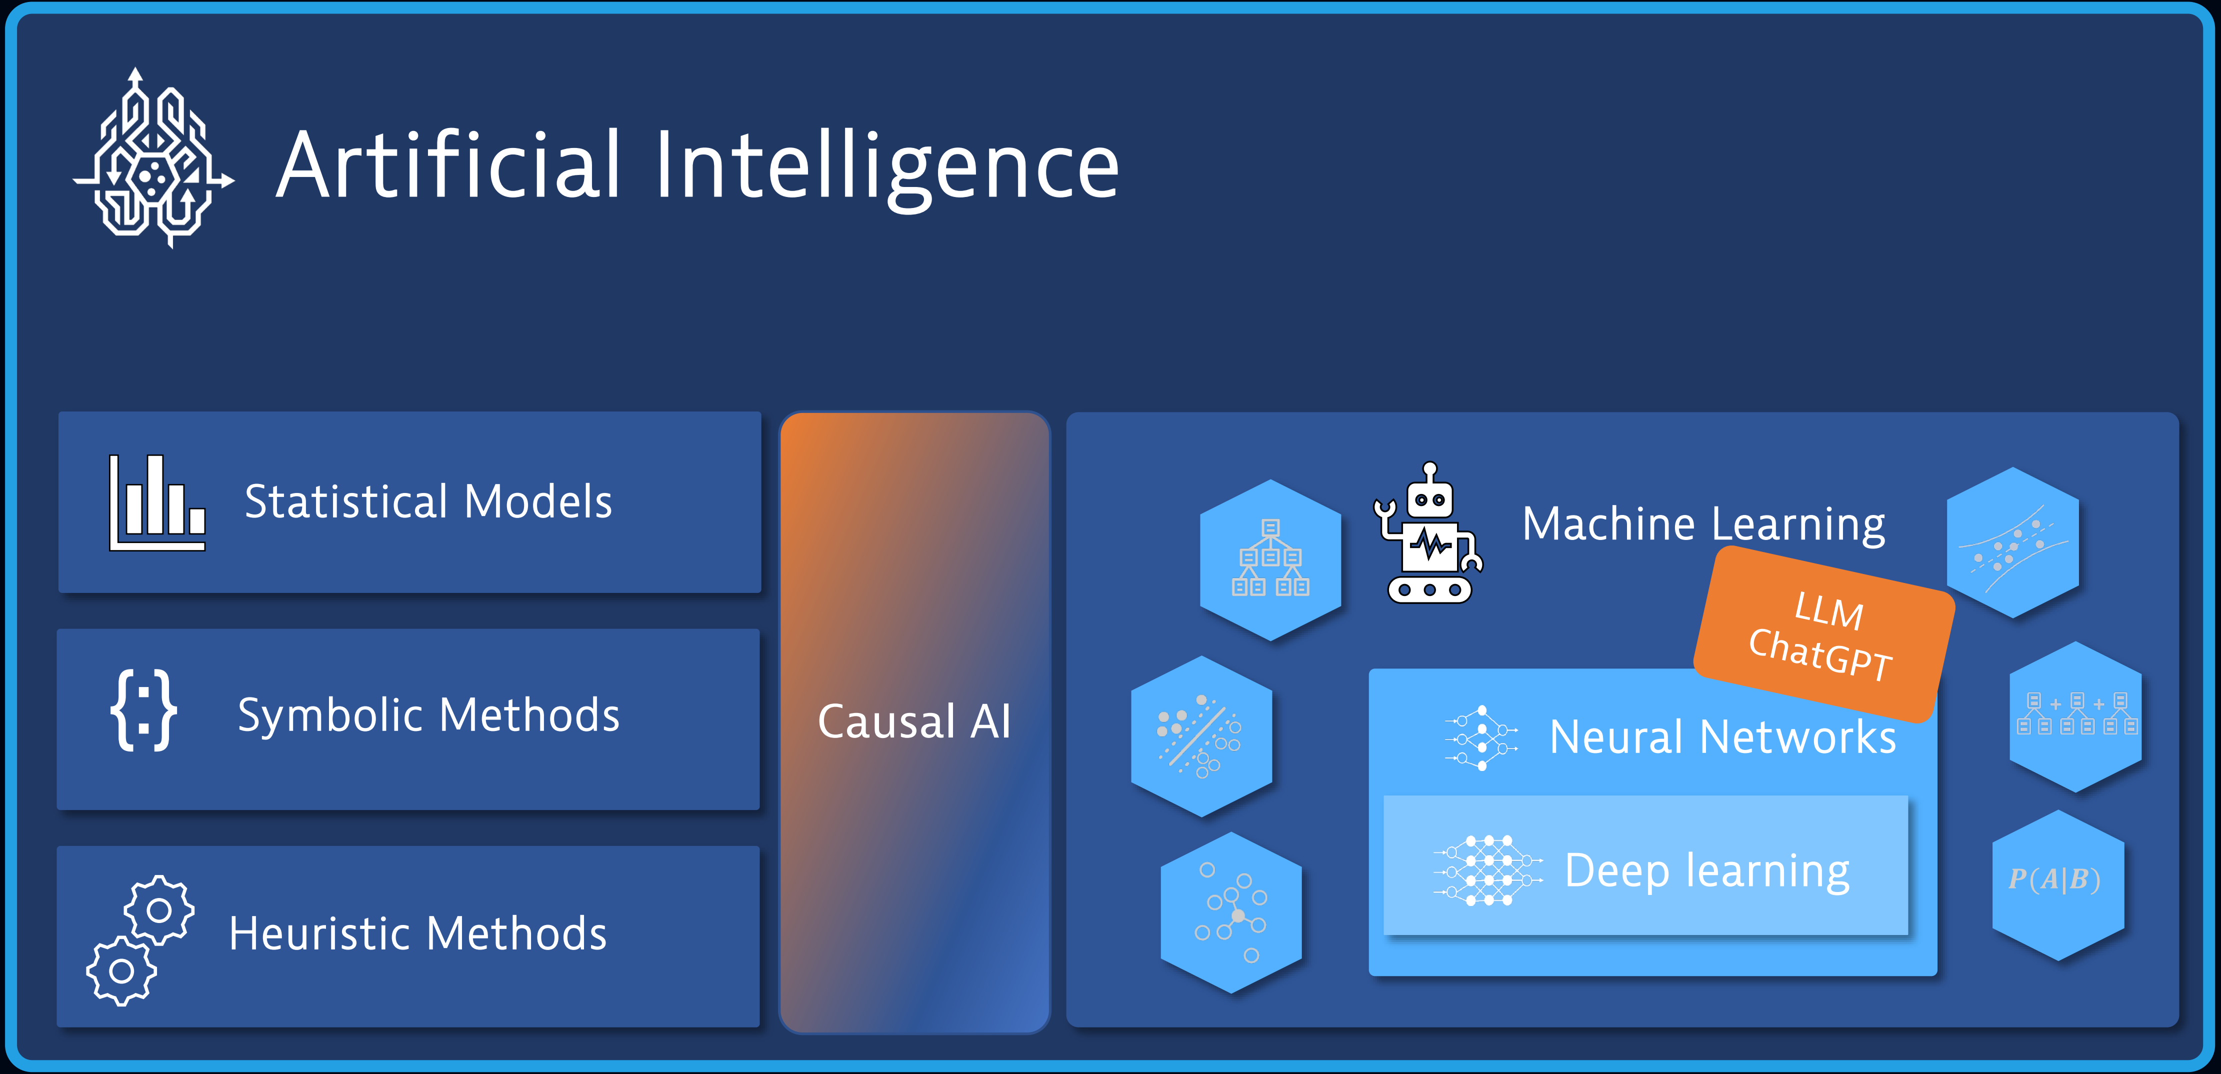

Artificial Intelligence

Statistical Models

Symbolic Methods

Heuristic Methods

Causal AI

Machine Learning

LLM
ChatGPT

Neural Networks

Deep learning

LLM: CHATGPT IS A PROBABILISTIC LANGUAGE MODEL

Probabilistic mapping
from the **recent history (context)**
to the next word!

The best thing about AI is its ability to

learn	4,5%
predict	3,5%
make	3,2%
understand	3,1%
do	2,9%

context

conditional probabilities (top 5)

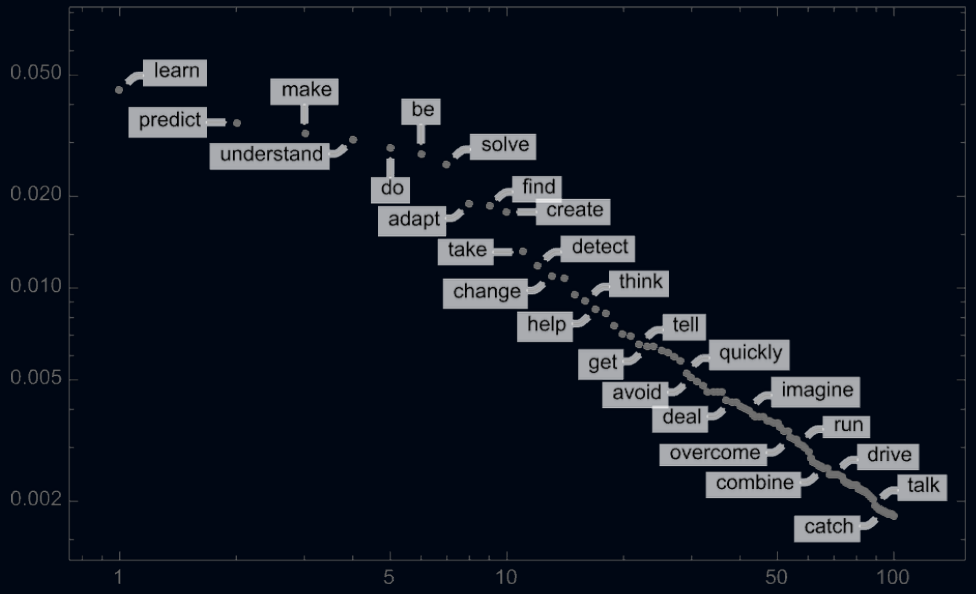

more conditional probabilities

CAUSAL AI

- AI technique used to determine the **exact underlying causes and effects of events** or behaviors.
- Uses **fault-tree analysis** to determine system-level failures based on component-level failures.
- Draws on supporting data, such as **relationships**, **dependencies**, and other **context among network entities** and **events**.

Uber
Causal**ML**

DoWhy

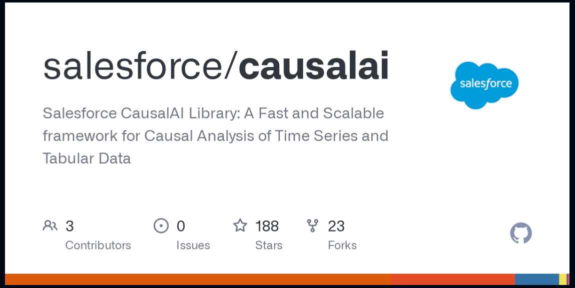

salesforce/**causalai**

Salesforce CausalAI Library: A Fast and Scalable framework for Causal Analysis of Time Series and Tabular Data

3	0	188	23
Contributors	Issues	Stars	Forks

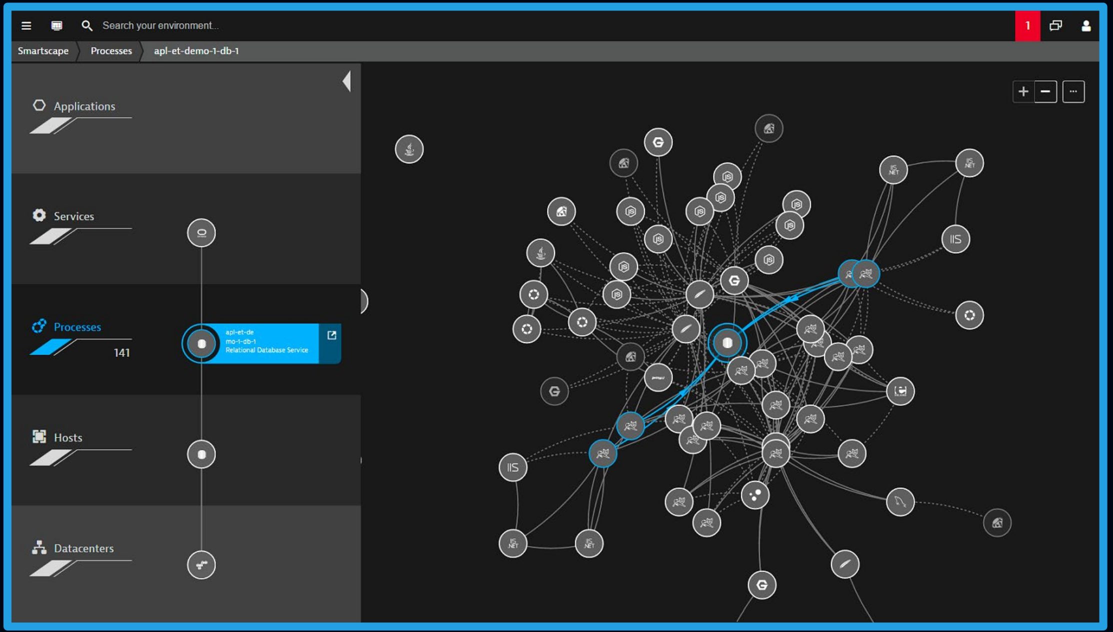

REDEFINING GENERATIVE AI THOUGHT LEADERSHIP

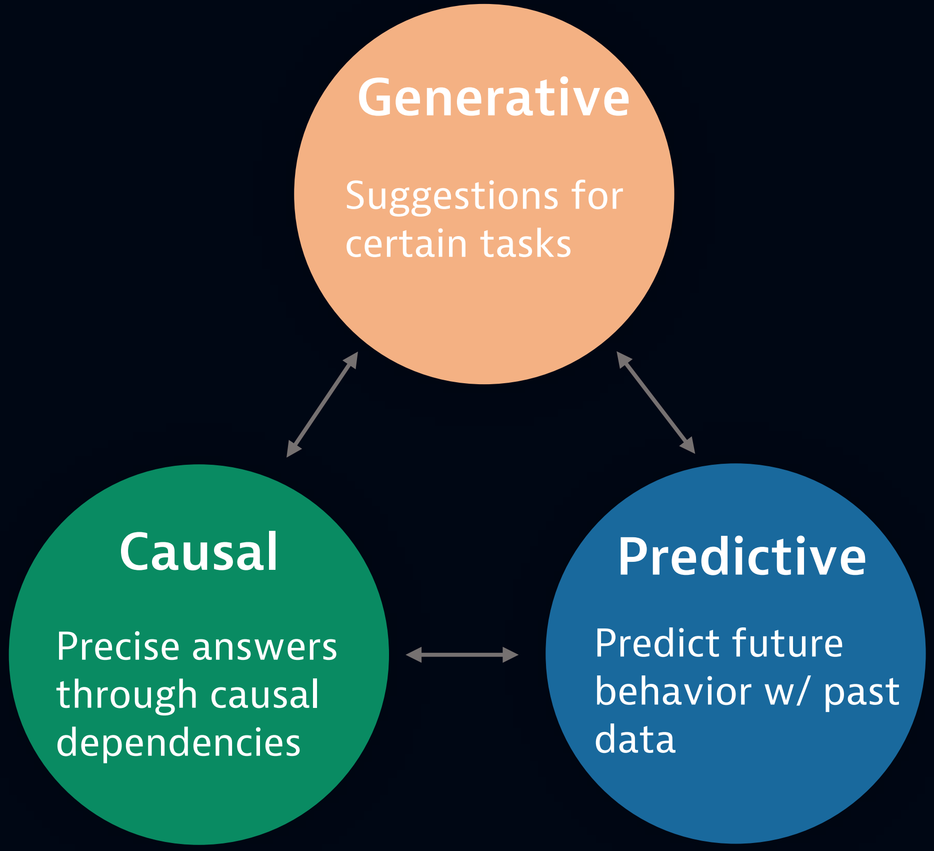

Generative

Suggestions for certain tasks

Causal

Precise answers through causal dependencies

Predictive

Predict future behavior w/ past data

Dynatrace Introduces the Industry’s first Hypermodal AI

- **Causal AI**: Assist in the identification of problems and root causes.
- **Generative AI**: Assist in the creation of DQL, dashboards, metrics… and much more.
- **Predictive AI**: Helping to understand the future behavior of systems/applications.

HYPERMODAL AI FOR UNIFIED OBSERVABILITY AND SECURITY

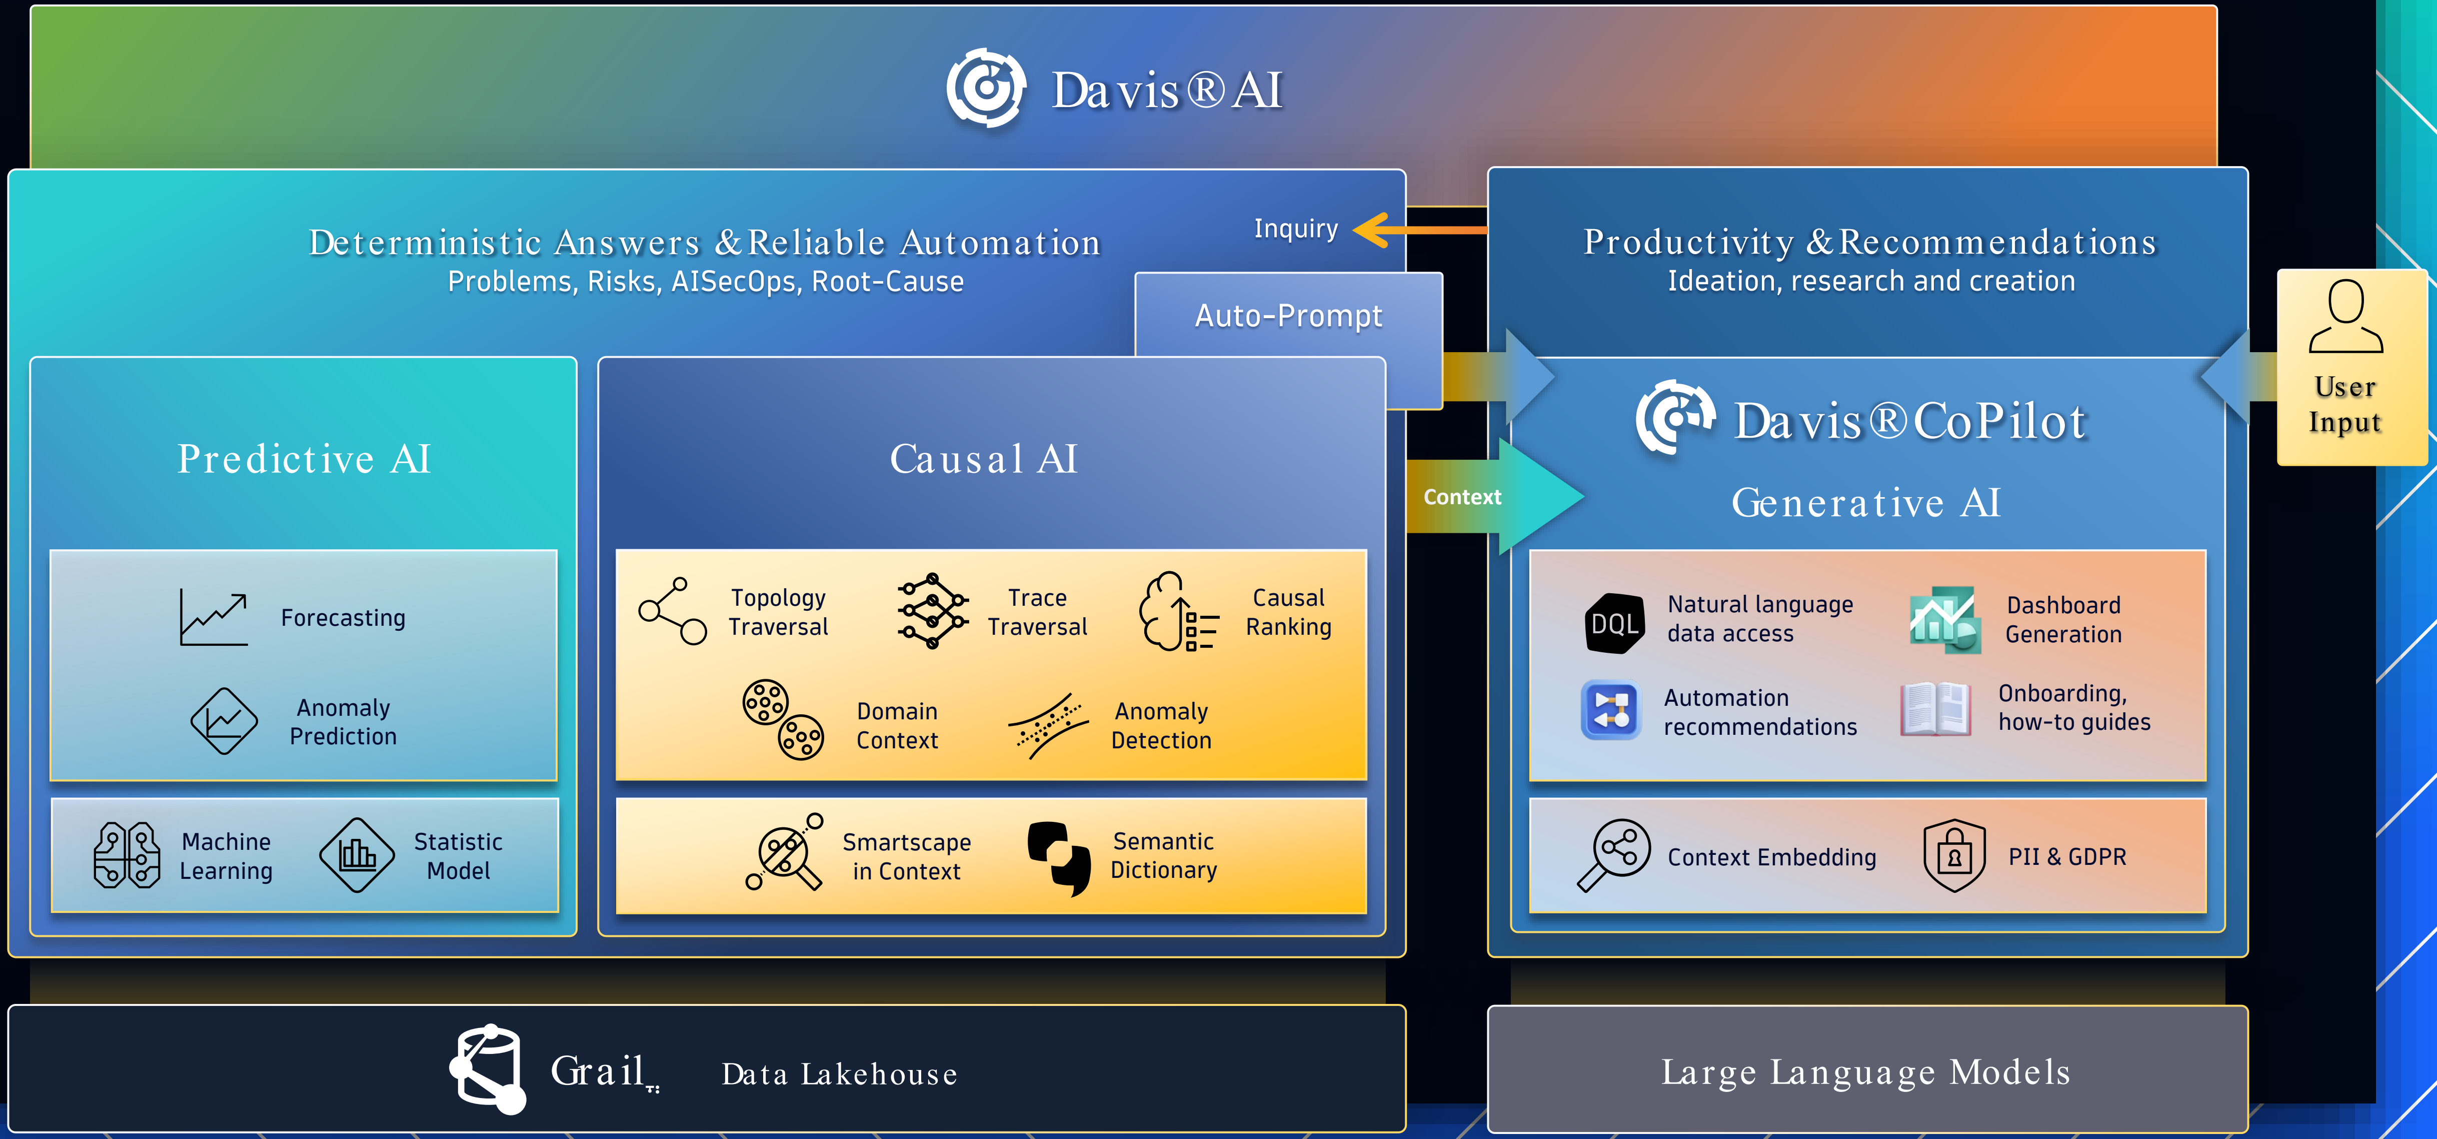

Davis® AI

Deterministic Answers & Reliable Automation
Problems, Risks, AISecOps, Root-Cause

Inquiry

Auto-Prompt

Predictive AI

- Forecasting
- Anomaly Prediction
- Machine Learning
- Statistic Model

Causal AI

- Topology Traversal
- Trace Traversal
- Causal Ranking
- Domain Context
- Anomaly Detection
- Smartscape in Context
- Semantic Dictionary

Context

Productivity & Recommendations
Ideation, research and creation

User Input

Davis® CoPilot

Generative AI

- DQL Natural language data access
- Dashboard Generation
- Automation recommendations
- Onboarding, how-to guides
- Context Embedding
- PII & GDPR

Grail Data Lakehouse

Large Language Models

DAVIS AI – USE CASES

DAVIS ENRICHES LLM PROMPT WITH CONTEXT

Davis has access to Smartscape and knows all dependencies and has full real-time context → used to automatically engineer a prompt

Regular Prompt

How many users were impacted by yesterday's outage?

Unenriched answer – no real help

This
This depends
This depends on
This depends on many
This depends on many different
This depends on many different factors
This depends on many different factors, for
This depends on many different factors, for example
This depends on many different factors, for example, how
This depends on many different factors, for example, how many
This depends on many different factors, for example, how many users
This depends on many different factors, for example, how many users were
This depends on many different factors, for example, how many users were online
This depends on many different factors, for example, how many users were online yesterday
…

Engineered Prompt

Generate a query based on the domain specific language Dynatrace Query Language (DQL). Start with problem "P-230771722" and find the related application based on the affected host. Then, correlate RUM error events based on application showing the distinct number of user sessions. Kubernetes node 77652 depends on loyalty service ….

Question: **How many users were impacted by yesterday's outage?**

Davis-enriched answer → **USEFUL QUERY**

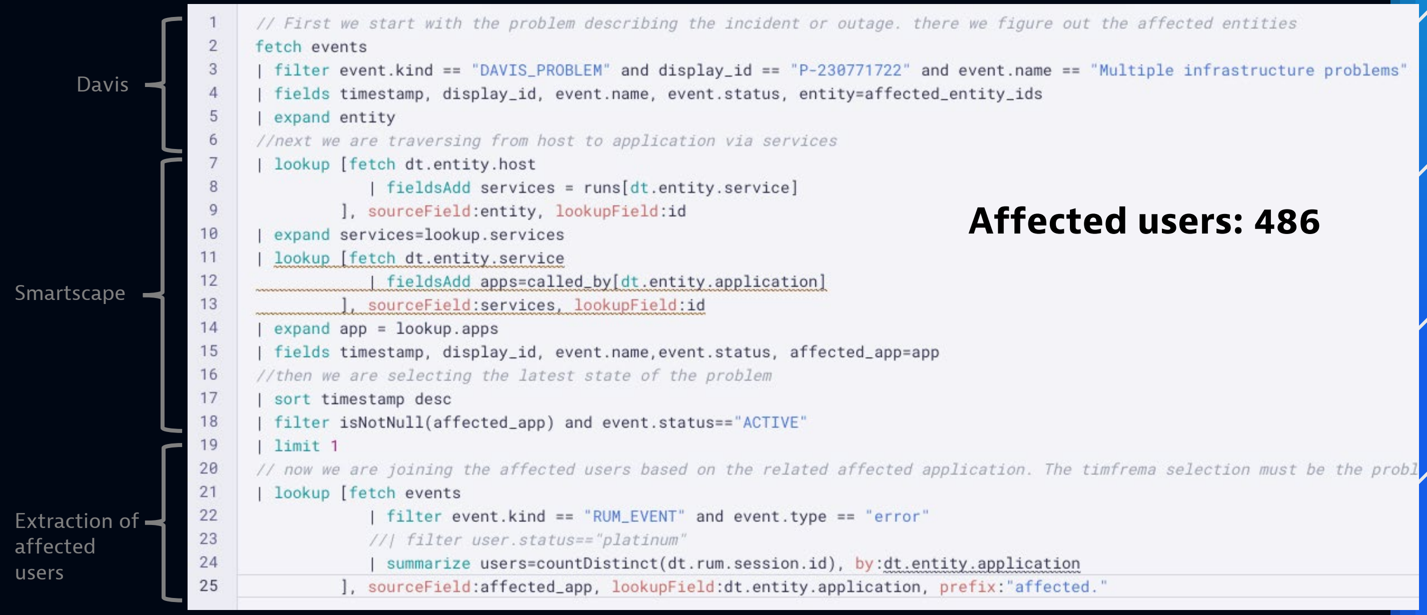

```
// First we start with the problem describing the incident or outage. there we figure out the affected entities
fetch events
| filter event.kind == "DAVIS_PROBLEM" and display_id == "P-230771722" and event.name == "Multiple infrastructure problems"
| fields timestamp, display_id, event.name, event.status, entity=affected_entity_ids
| expand entity
//next we are traversing from host to application via services
| lookup [fetch dt.entity.host
            | fieldsAdd services = runs[dt.entity.service]
         ], sourceField:entity, lookupField:id
| expand services=lookup.services
| lookup [fetch dt.entity.service
            | fieldsAdd apps=called_by[dt.entity.application]
         ], sourceField:services, lookupField:id
| expand app = lookup.apps
| fields timestamp, display_id, event.name,event.status, affected_app=app
//then we are selecting the latest state of the problem
| sort timestamp desc
| filter isNotNull(affected_app) and event.status=="ACTIVE"
| limit 1
// now we are joining the affected users based on the related affected application. The timfrema selection must be the probl
| lookup [fetch events
            | filter event.kind == "RUM_EVENT" and event.type == "error"
            //| filter user.status=="platinum"
            | summarize users=countDistinct(dt.rum.session.id), by:dt.entity.application
         ], sourceField:affected_app, lookupField:dt.entity.application, prefix:"affected."
```

Affected users: 486

USE CASE: ROOT CAUSE ANALYSIS

Automatically analyze the root cause of a problem

Davis® AI

Causal AI

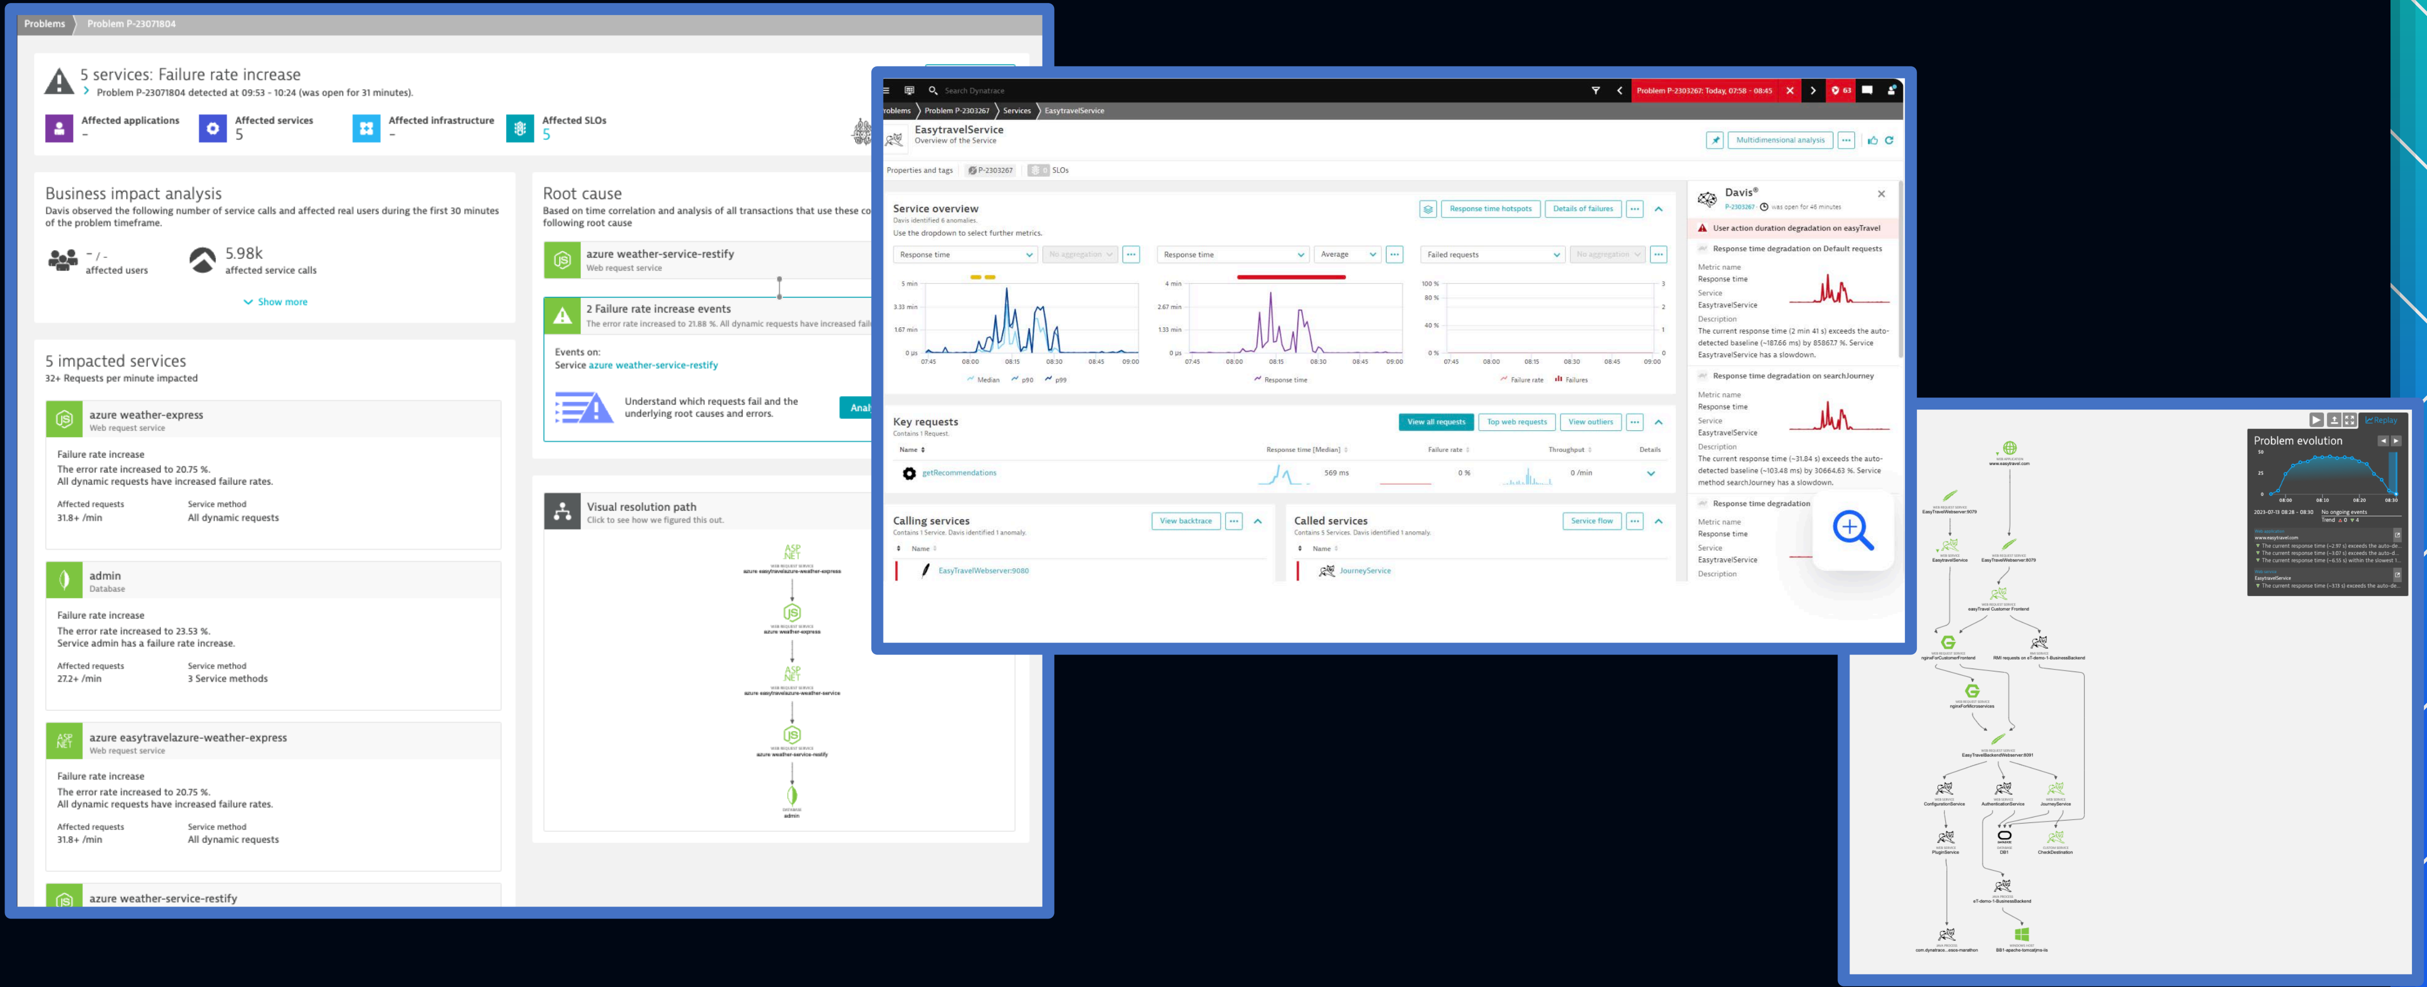

Problems › Problem P-23071804

5 services: Failure rate increase

› Problem P-23071804 detected at 09:53 - 10:24 (was open for 31 minutes).

Affected applications	Affected services	Affected infrastructure	Affected SLOs
-	5	-	5

Business impact analysis

Davis observed the following number of service calls and affected real users during the first 30 minutes of the problem timeframe.

- / - affected users

5.98k affected service calls

Show more

5 impacted services

32+ Requests per minute impacted

azure weather-express
Web request service

Failure rate increase
The error rate increased to 20.75 %.
All dynamic requests have increased failure rates.

Affected requests	Service method
31.8+ /min	All dynamic requests

admin
Database

Failure rate increase
The error rate increased to 23.53 %.
Service admin has a failure rate increase.

Affected requests	Service method
27.2+ /min	3 Service methods

azure easytravelazure-weather-express
Web request service

Failure rate increase
The error rate increased to 20.75 %.
All dynamic requests have increased failure rates.

Affected requests	Service method
31.8+ /min	All dynamic requests

azure weather-service-restify

Root cause

Based on time correlation and analysis of all transactions that use these co
following root cause

azure weather-service-restify
Web request service

2 Failure rate increase events
The error rate increased to 21.88 %. All dynamic requests have increased fail

Events on:
Service azure weather-service-restify

Understand which requests fail and the underlying root causes and errors.

Visual resolution path
Click to see how we figured this out.

azure easytravelazure-weather-express
azure weather-express
azure easytravelazure-weather-service
azure weather-service-restify
admin

Search Dynatrace

Problem P-2303267: Today, 07:58 - 08:45

roblems › Problem P-2303267 › Services › EasytravelService

EasytravelService
Overview of the Service

Multidimensional analysis

Properties and tags · P-2303267 · 0 SLOs

Service overview
Davis identified 6 anomalies.
Use the dropdown to select further metrics.

Response time · No aggregation
Response time · Average
Failed requests · No aggregation

Response time hotspots · Details of failures

Key requests
Contains 1 Request.

View all requests · Top web requests · View outliers

Name	Response time [Median]	Failure rate	Throughput	Details
getRecommendations	569 ms	0 %	0 /min	

Calling services
Contains 1 Service. Davis identified 1 anomaly.

View backtrace

Name
EasyTravelWebserver:9080

Called services
Contains 5 Services. Davis identified 1 anomaly.

Service flow

Name
JourneyService

Davis®
P-2303267 · was open for 46 minutes

User action duration degradation on easyTravel

Response time degradation on Default requests

Metric name
Response time
Service
EasytravelService
Description
The current response time (2 min 41 s) exceeds the auto-detected baseline (~187.66 ms) by 85867.7 %. Service EasytravelService has a slowdown.

Response time degradation on searchJourney

Metric name
Response time
Service
EasytravelService
Description
The current response time (~31.84 s) exceeds the auto-detected baseline (~103.48 ms) by 30664.63 %. Service method searchJourney has a slowdown.

Response time degradation

Metric name
Response time
Service
EasytravelService
Description

Problem evolution

USE CASE: PREDICTIVE OPERATIONS

Davis® AI

Predictive AI

Automate manual operations task and solve problems before they occur.

For example: Predictive disk resizing, predictive autoscaling

1. **Choose the trigger schedule, check once a day, hour, week**
2. **Forecast all disk free space**
3. **Generates a report of all disks that will run out**
4. **Choose a delivery action, e.g.: send by Slack, email, etc**

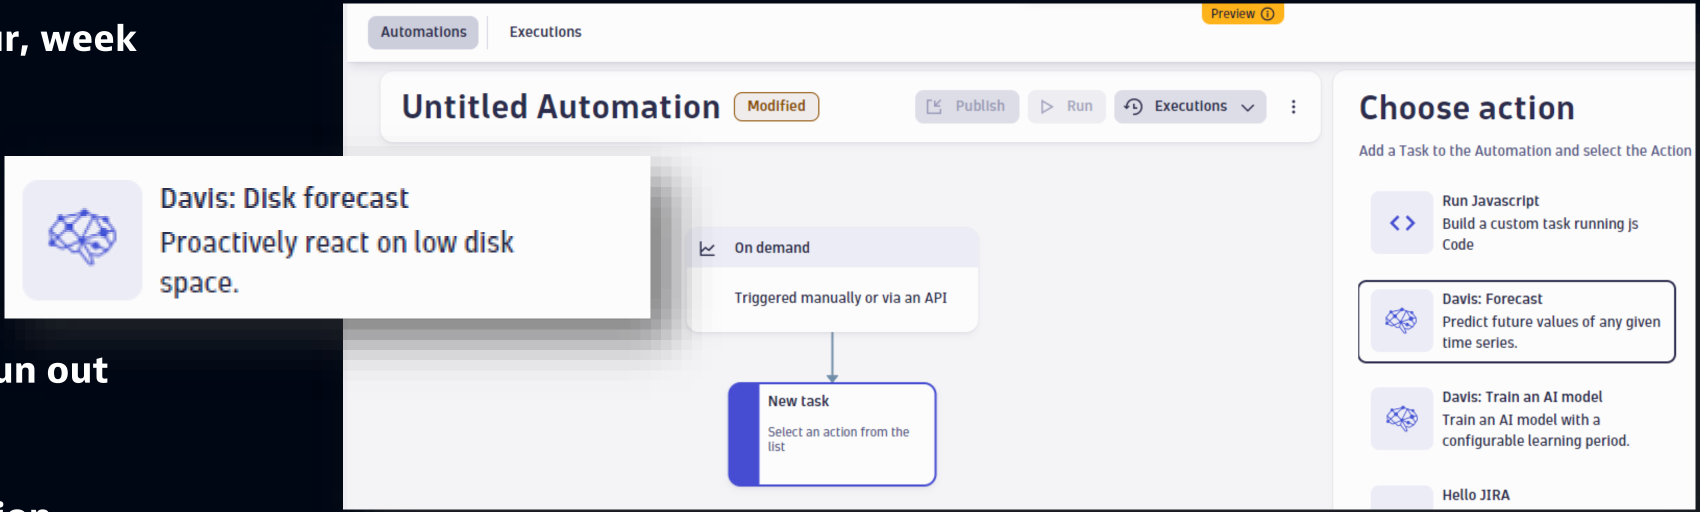

Davis: Disk forecast

Proactively react on low disk space.

Untitled Automation

On demand

Triggered manually or via an API

New task

Select an action from the list

Choose action

Add a Task to the Automation and select the Action

Run Javascript

Build a custom task running js Code

Davis: Forecast

Predict future values of any given time series.

Davis: Train an AI model

Train an AI model with a configurable learning period.

Hello JIRA

https://www.dynatrace.com/news/blog/stay-ahead-of-the-game-forecast-it-capacity-with-dynatrace-grail-and-davis-ai/

USE CASE: AI-BASED AUTOMATION

Davis® AI

Causal AI

Predictive AI

Use custom AI analysis steps to build advanced answer-driven automation

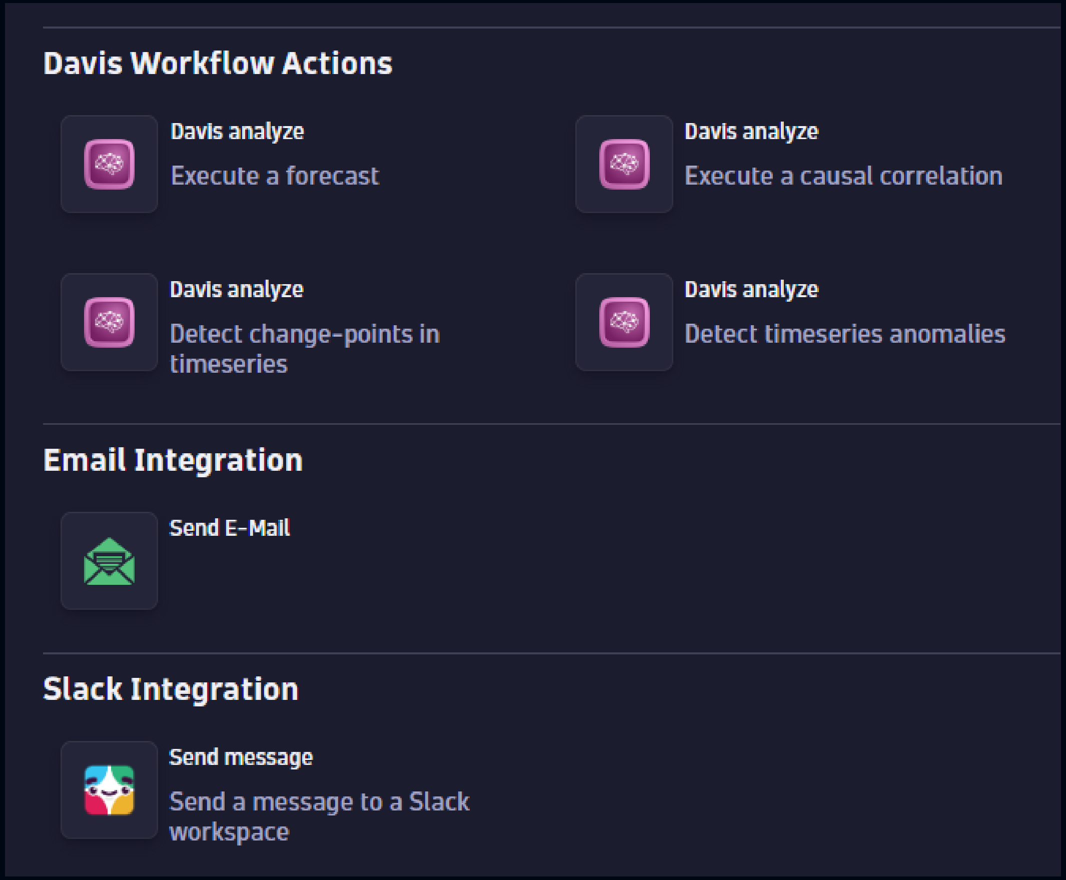

Davis Workflow Actions

Davis analyze
Execute a forecast

Davis analyze
Execute a causal correlation

Davis analyze
Detect change-points in timeseries

Davis analyze
Detect timeseries anomalies

Email Integration

Send E-Mail

Slack Integration

Send message
Send a message to a Slack workspace

Automate AI/ML

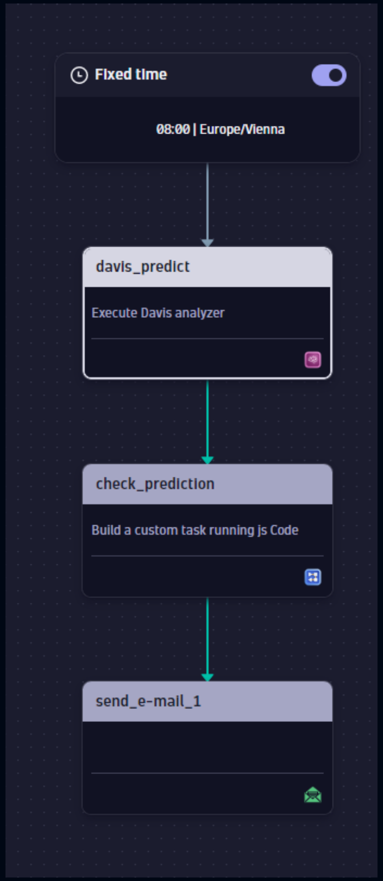

Fixed time

08:00 | Europe/Vienna

davis_predict

Execute Davis analyzer

check_prediction

Build a custom task running js Code

send_e-mail_1

USE CASE: NATURAL LANGUAGE-BASED PRODUCTIVITY & ONBOARDING

Davis® AI

Causal AI

Davis CoPilot™

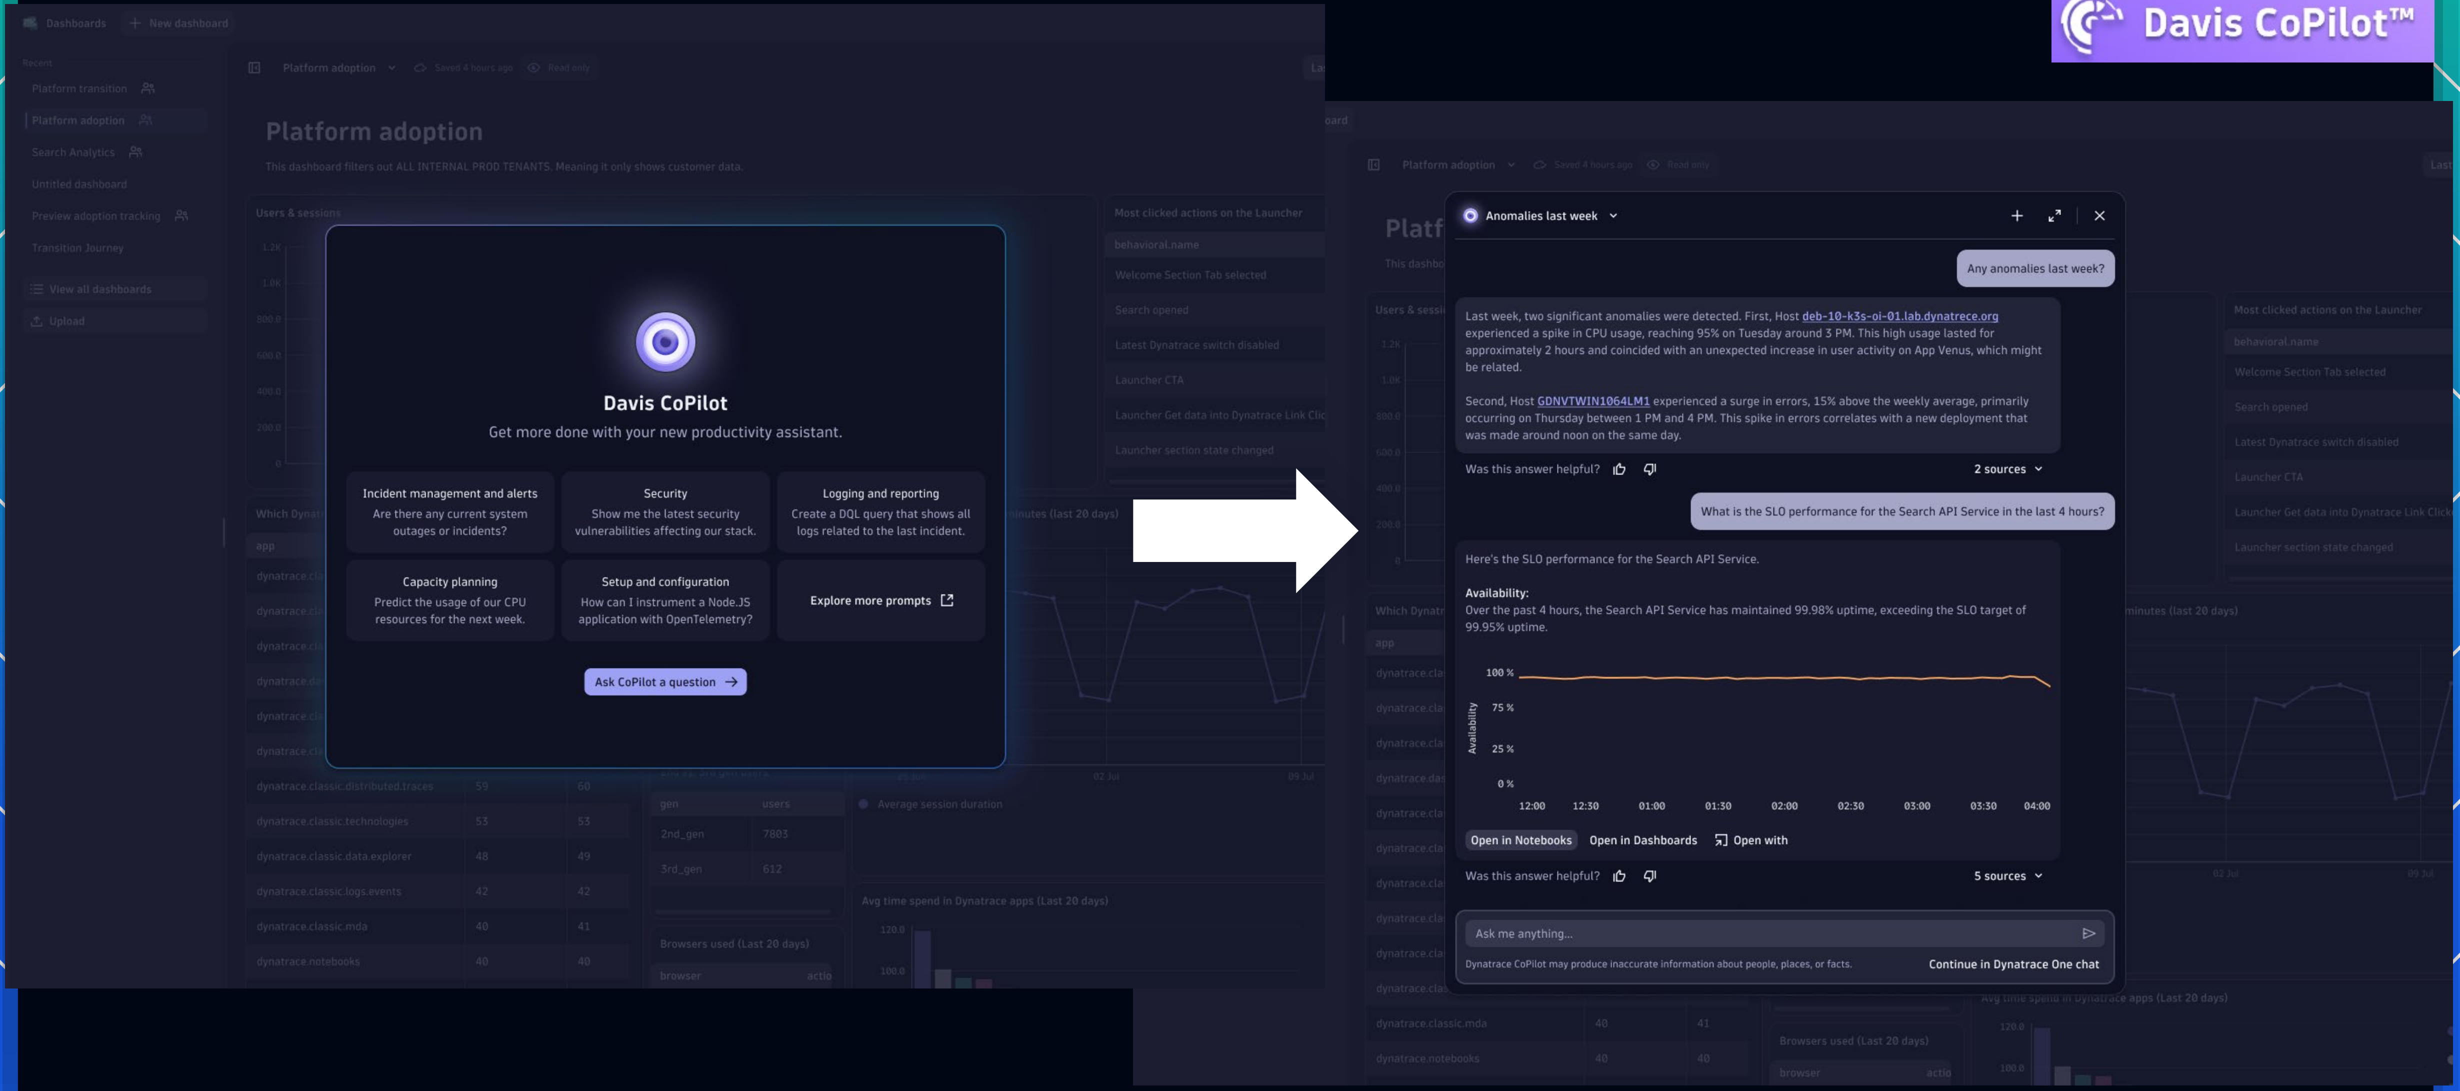

Platform adoption

This dashboard filters out ALL INTERNAL PROD TENANTS. Meaning it only shows customer data.

Davis CoPilot

Get more done with your new productivity assistant.

Incident management and alerts Are there any current system outages or incidents?	Security Show me the latest security vulnerabilities affecting our stack.	Logging and reporting Create a DQL query that shows all logs related to the last incident.
Capacity planning Predict the usage of our CPU resources for the next week.	Setup and configuration How can I instrument a Node.JS application with OpenTelemetry?	Explore more prompts

Ask CoPilot a question →

Any anomalies last week?

Last week, two significant anomalies were detected. First, Host deb-10-k3s-oi-01.lab.dynatrece.org experienced a spike in CPU usage, reaching 95% on Tuesday around 3 PM. This high usage lasted for approximately 2 hours and coincided with an unexpected increase in user activity on App Venus, which might be related.

Second, Host GDNVTWIN1064LM1 experienced a surge in errors, 15% above the weekly average, primarily occurring on Thursday between 1 PM and 4 PM. This spike in errors correlates with a new deployment that was made around noon on the same day.

Was this answer helpful?

What is the SLO performance for the Search API Service in the last 4 hours?

Here's the SLO performance for the Search API Service.

Availability:
Over the past 4 hours, the Search API Service has maintained 99.98% uptime, exceeding the SLO target of 99.95% uptime.

Was this answer helpful?

Dynatrace CoPilot may produce inaccurate information about people, places, or facts.

USE CASE: AI-ASSISTED NOTEBOOKS

Davis® AI

Causal AI

Davis CoPilot™

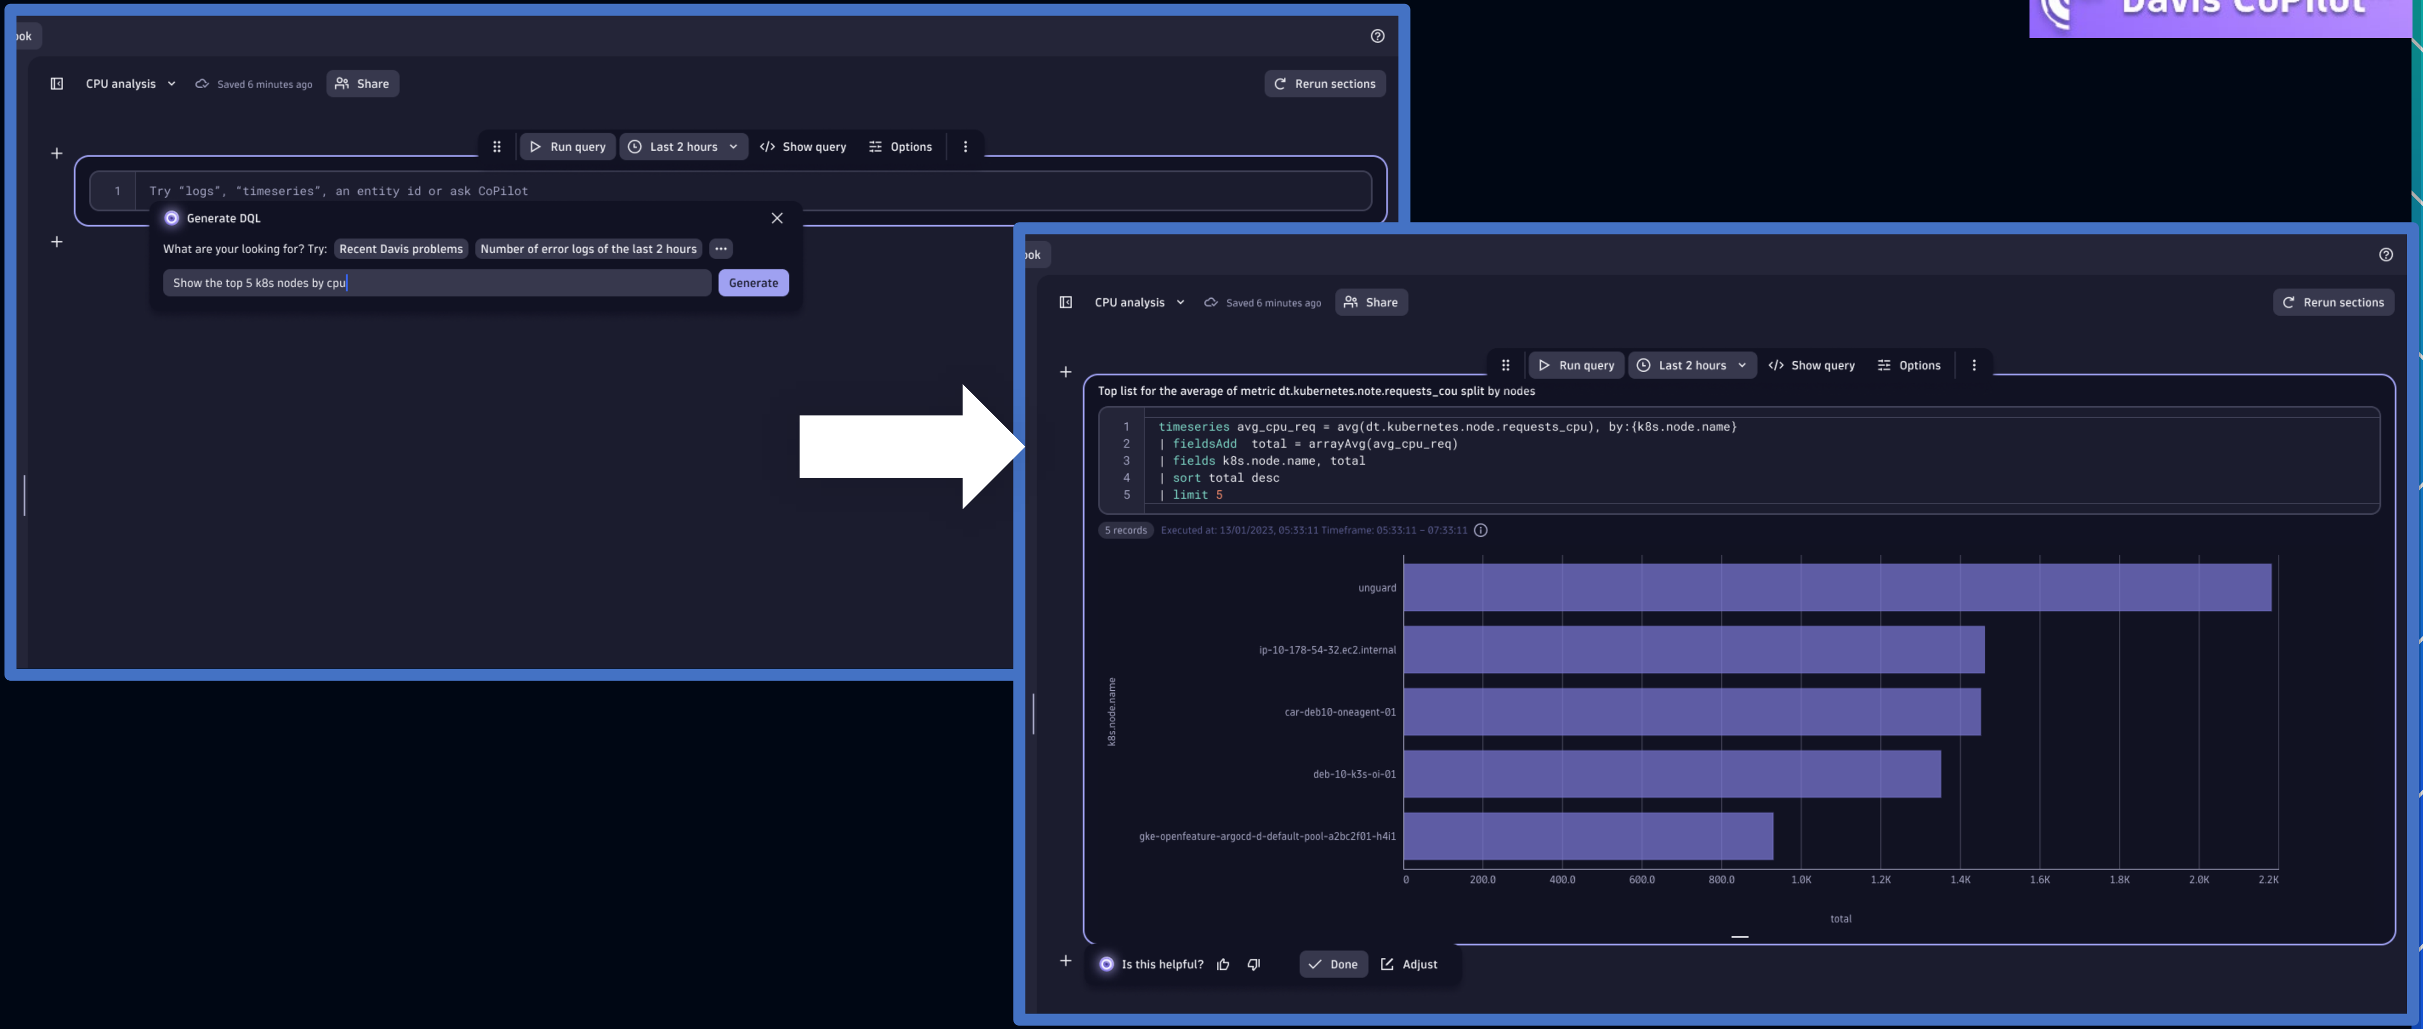

Try "logs", "timeseries", an entity id or ask CoPilot

What are your looking for? Try: Recent Davis problems · Number of error logs of the last 2 hours

Show the top 5 k8s nodes by cpu

Top list for the average of metric dt.kubernetes.note.requests_cou split by nodes

```
timeseries avg_cpu_req = avg(dt.kubernetes.node.requests_cpu), by:{k8s.node.name}
| fieldsAdd  total = arrayAvg(avg_cpu_req)
| fields k8s.node.name, total
| sort total desc
| limit 5
```

5 records Executed at: 13/01/2023, 05:33:11 Timeframe: 05:33:11 – 07:33:11

USE CASES: REMEDIATION PROPOSAL

Davis® AI

Causal AI

Davis CoPilot™

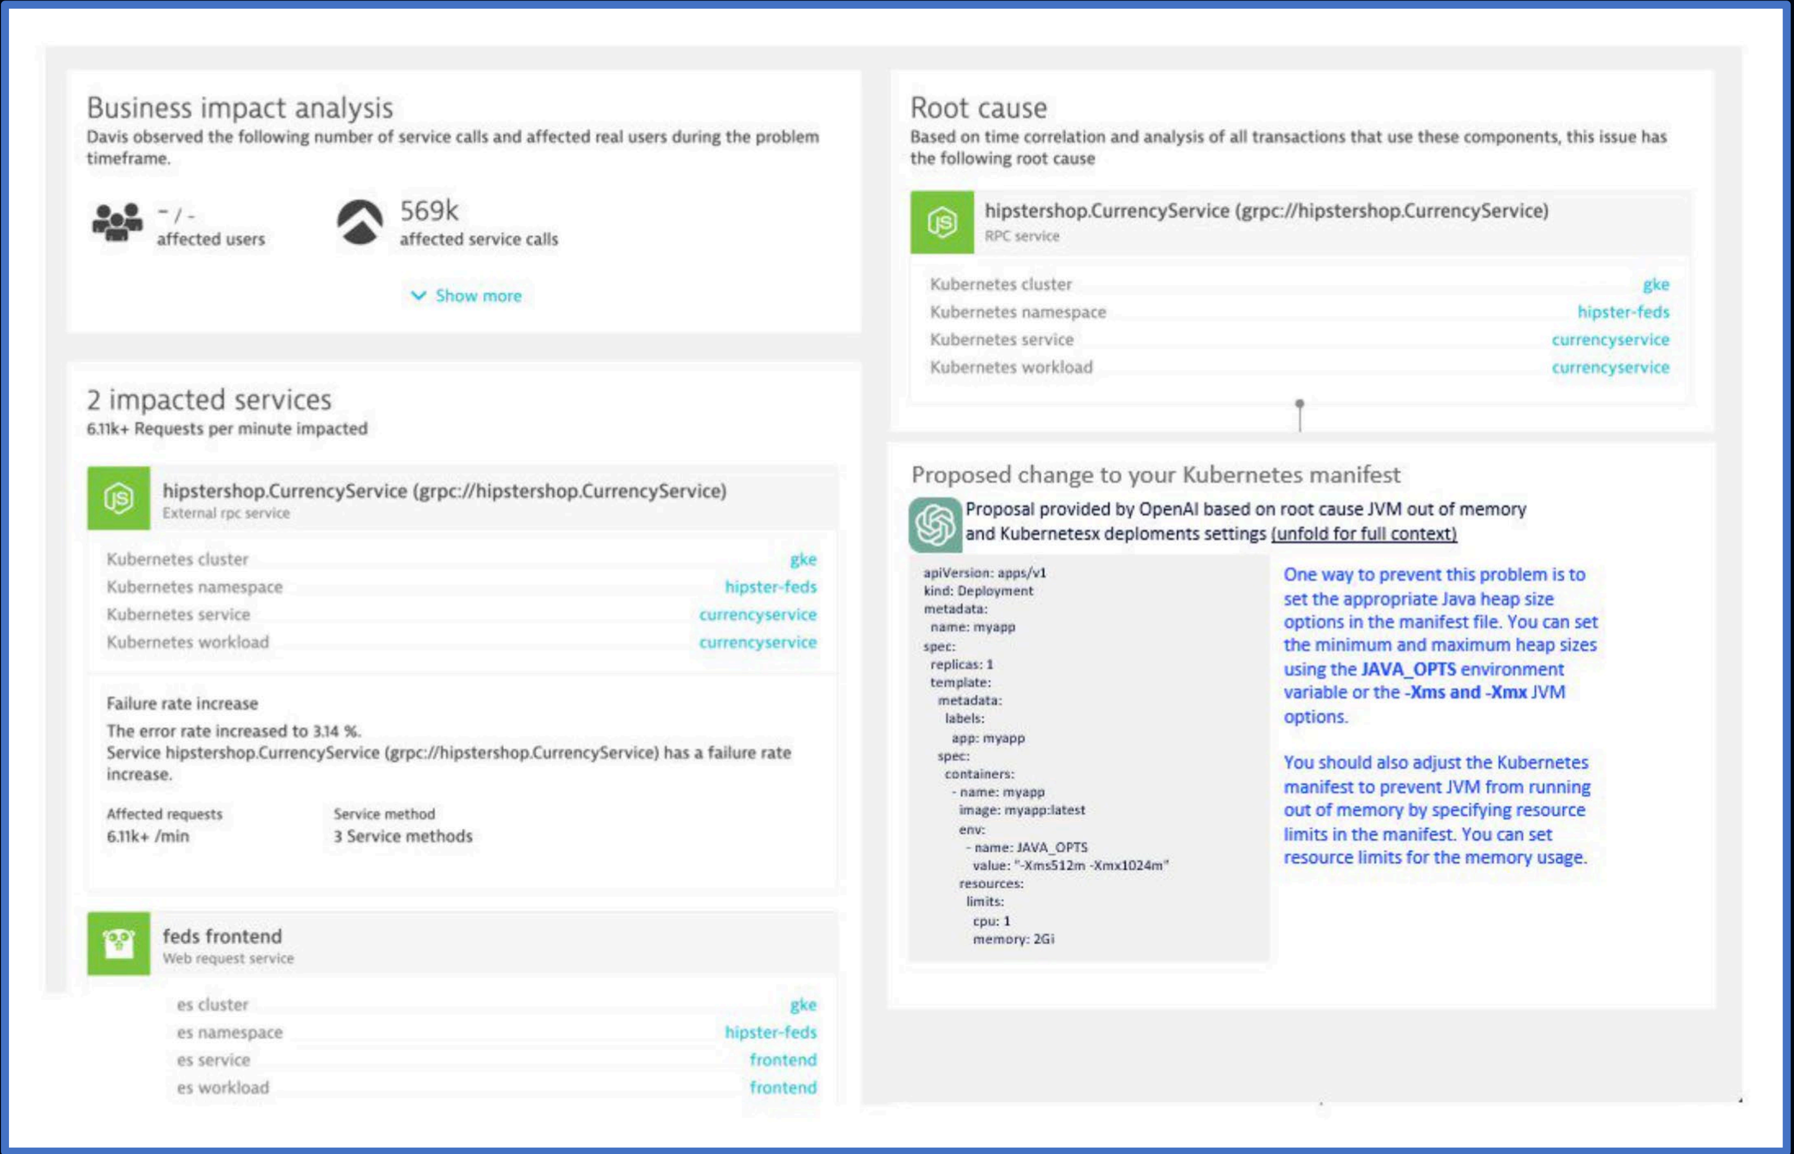

Business impact analysis

Davis observed the following number of service calls and affected real users during the problem timeframe.

- / -
affected users

569k
affected service calls

2 impacted services

6.11k+ Requests per minute impacted

hipstershop.CurrencyService (grpc://hipstershop.CurrencyService)
External rpc service

Kubernetes cluster	gke
Kubernetes namespace	hipster-feds
Kubernetes service	currencyservice
Kubernetes workload	currencyservice

Failure rate increase

The error rate increased to 3.14 %.
Service hipstershop.CurrencyService (grpc://hipstershop.CurrencyService) has a failure rate increase.

Affected requests	Service method
6.11k+ /min	3 Service methods

feds frontend
Web request service

es cluster	gke
es namespace	hipster-feds
es service	frontend
es workload	frontend

Root cause

Based on time correlation and analysis of all transactions that use these components, this issue has the following root cause

hipstershop.CurrencyService (grpc://hipstershop.CurrencyService)
RPC service

Kubernetes cluster	gke
Kubernetes namespace	hipster-feds
Kubernetes service	currencyservice
Kubernetes workload	currencyservice

Proposed change to your Kubernetes manifest

Proposal provided by OpenAI based on root cause JVM out of memory and Kubernetesx deploments settings (unfold for full context)

```
apiVersion: apps/v1
kind: Deployment
metadata:
 name: myapp
spec:
 replicas: 1
 template:
  metadata:
   labels:
    app: myapp
  spec:
   containers:
    - name: myapp
     image: myapp:latest
     env:
      - name: JAVA_OPTS
       value: "-Xms512m -Xmx1024m"
     resources:
      limits:
       cpu: 1
       memory: 2Gi
```

One way to prevent this problem is to set the appropriate Java heap size options in the manifest file. You can set the minimum and maximum heap sizes using the **JAVA_OPTS** environment variable or the **-Xms and -Xmx** JVM options.

You should also adjust the Kubernetes manifest to prevent JVM from running out of memory by specifying resource limits in the manifest. You can set resource limits for the memory usage.

dynatrace Hypermodal AI for Unified Observability and Security

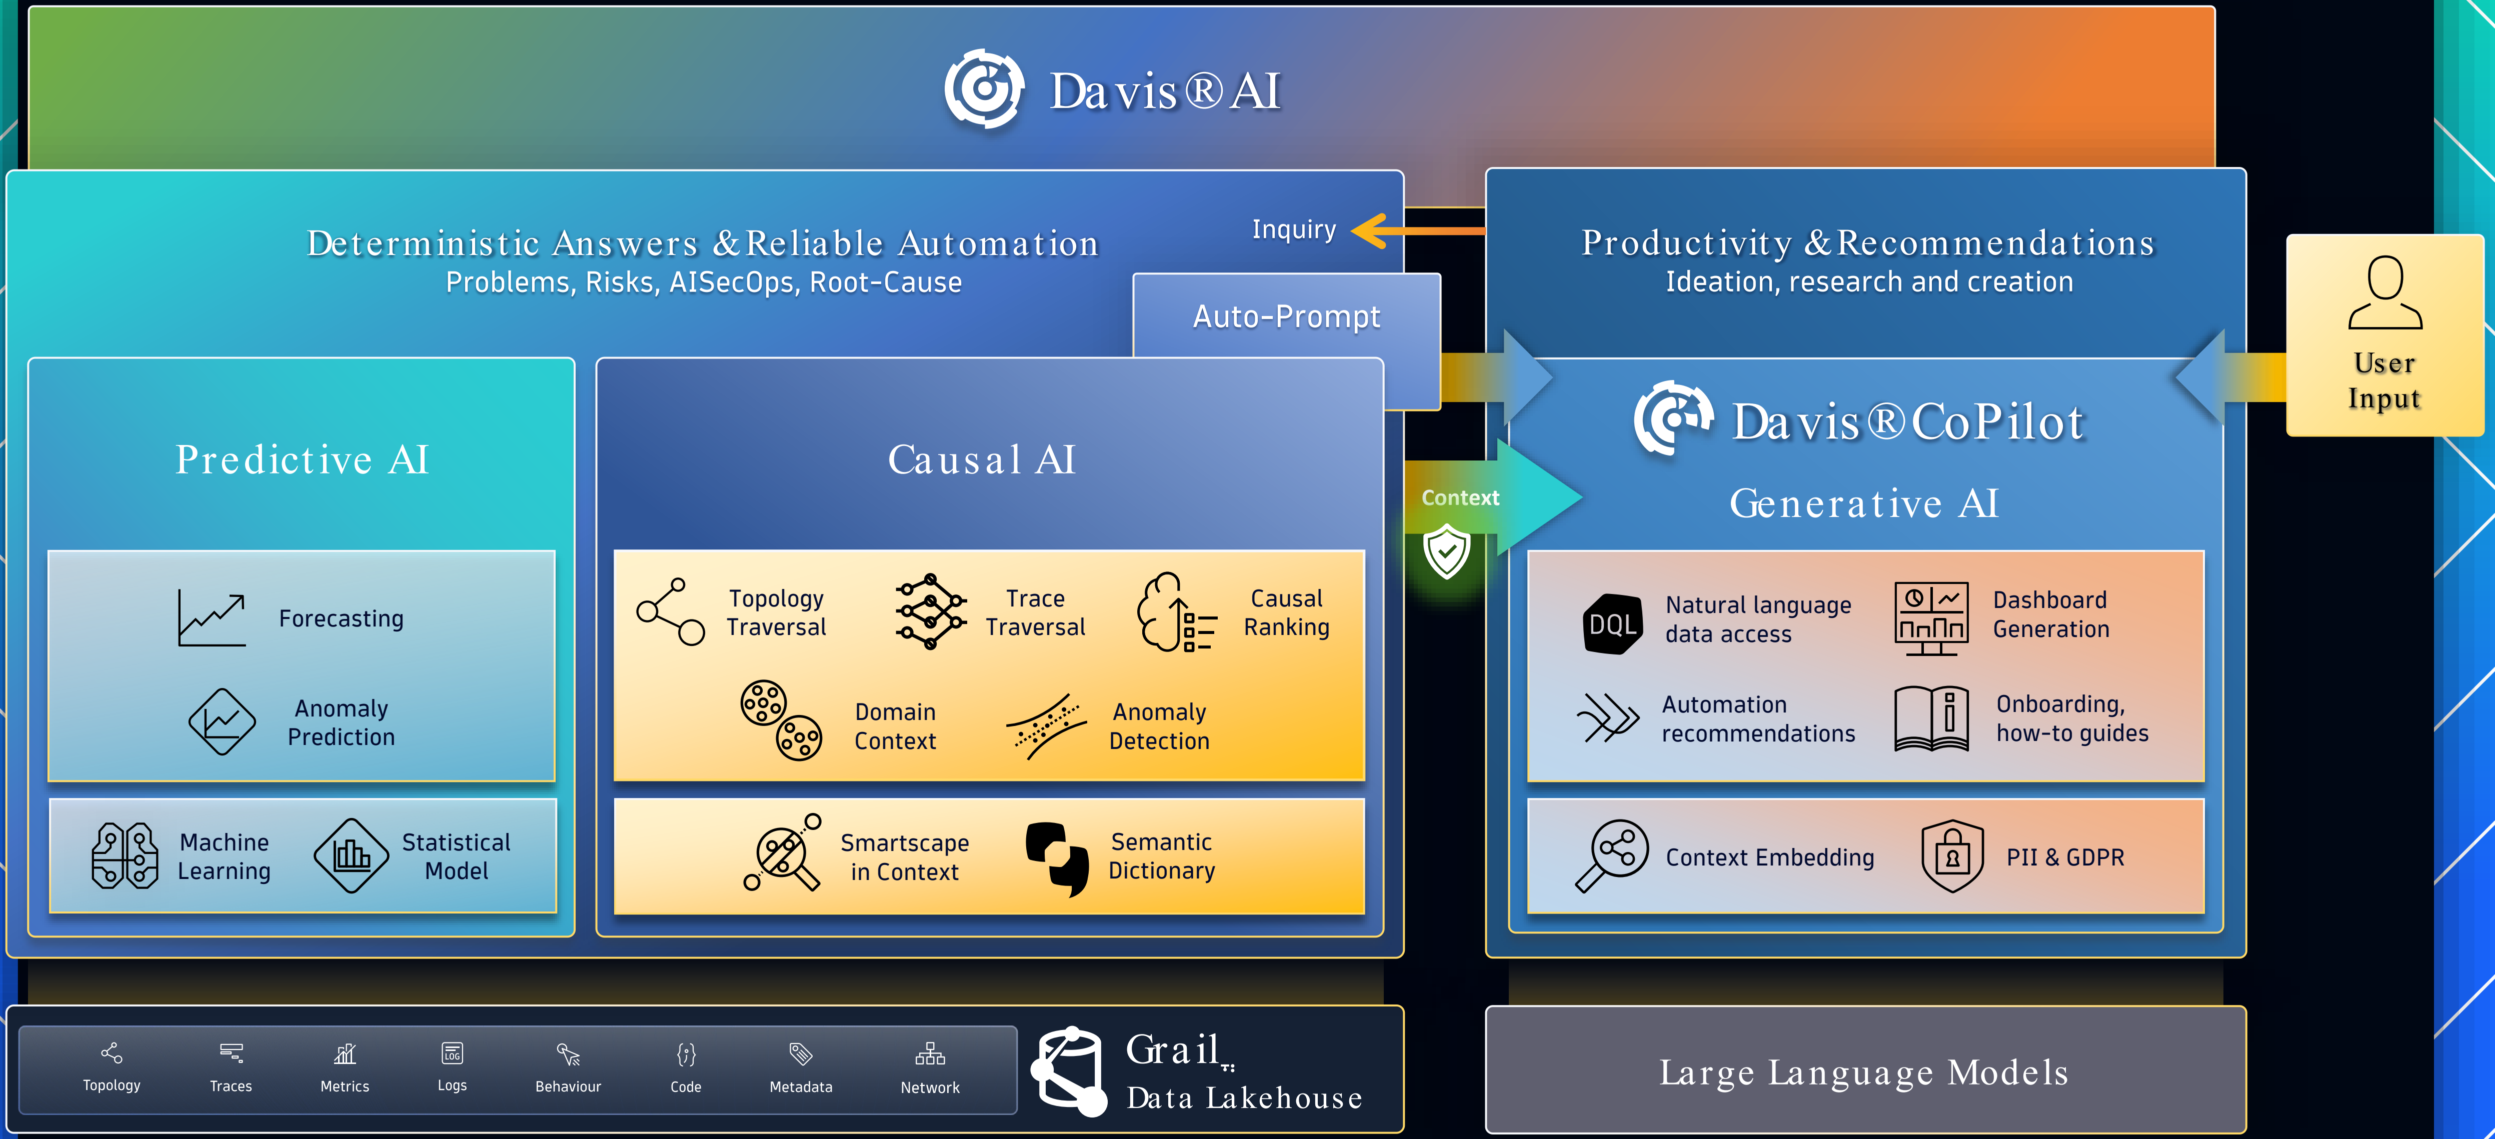

Davis® AI

Deterministic Answers & Reliable Automation

Problems, Risks, AISecOps, Root-Cause

Inquiry

Auto-Prompt

Predictive AI

- Forecasting
- Anomaly Prediction
- Machine Learning
- Statistical Model

Causal AI

- Topology Traversal
- Trace Traversal
- Causal Ranking
- Domain Context
- Anomaly Detection
- Smartscape in Context
- Semantic Dictionary

Context

Productivity & Recommendations

Ideation, research and creation

Davis® CoPilot

Generative AI

- DQL Natural language data access
- Dashboard Generation
- Automation recommendations
- Onboarding, how-to guides
- Context Embedding
- PII & GDPR

User Input

Topology, Traces, Metrics, Logs, Behaviour, Code, Metadata, Network

Grail™ Data Lakehouse

Large Language Models

dynatrace

IMMERSE

📍 MADRID

📅 21.11.23