

# Dynatrace Powered Innovation

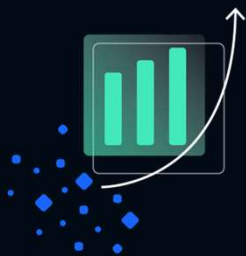
Cloud done right.



PRESENTER

Jon Griffiths  
Field CTO, EMEA

# What's driven our innovation



Data collection  
explodes



Making sense  
from data



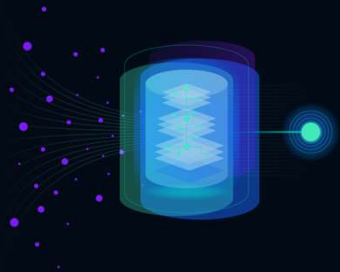
Observability and  
security converges  
with need for  
automation



BizDevSecOps &  
Everything-as-Code

# Key innovations that will drive us forward

---



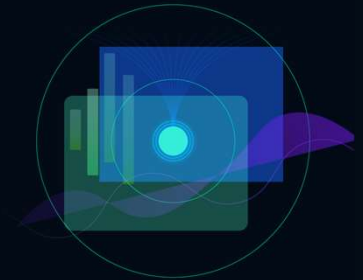
Grail



AutomationEngine



AppEngine



New User Experience



Analytics and Automation for Unified Observability and Security **CLOUD DONE RIGHT.**



Infrastructure  
Observability



Application  
Observability



Security  
Protection



Security  
Analytics



Digital  
Experience



Business  
Analytics



Automations



Custom  
Solutions

Platform



AutomationEngine



AppEngine



Smartscape®



Davis® AI



Grail™



Hub

Unified Ingest



PurePath®



OneAgent®



Topology



Traces



Metrics



Logs



Behavior



Code



Metadata



Network



# Next Gen Storage



## Purpose-built data lakehouse

Focus on observability, security and automation  
Optimized for Dynatrace AI to process billions of dependencies

## Unified storage for observability data

Single data store for metrics, traces, logs and more  
Data stored in context within real-time dependency model

## >100x more scalable

100 TB/day data ingest per tenant (Q2 CY23)  
Aiming for 1000 PB/day depending on future market needs

## Cost efficient storage

Retain data for more than 12 months  
No need to manage cold, warm or hot storage tiers

## Instant answers

No index, no schema definition, no rehydration  
Data analytics at any time with Dynatrace Query Language

## Query performance up to 100x faster

Massive parallel processing on 1000s of nodes  
~1TB in ~1s @ 1000 cores

# DQL:ONE Language for all data in dynatrace

```
fetch dt.entity.host, from:-3d
| filter contains(entityName, "dtp-dev")
| lookup [fetch logs | filter matchesPhrase(content,"deployment") | parse content, "ld 'version: ' ld:version ld"
      | sort timestamp desc | fields latest = version, timestamp, dt.entity.host
      ], sourceField:entityId, lookupField:dt.entity.host
| lookup [fetch events | filter event.kind == "DAVIS_EVENT" and event.name=="CPU saturation"
      | sort timestamp desc | limit 1 | fields timestamp, dt.entity.host
      ], sourceField:entityId, lookupField:dt.entity.host
| lookup [timeseries avg(dt.host.cpu.usage), by:{dt.entity.host}, filter:contains(dt.entity.host,"HOST")
      ], sourceField:entityId, lookupField:dt.entity.host
| fields entityName, version = lookup.latest, deployment = lookup.timestamp,
      recentHighCPU = lookup.timestamp >= now()-3h, cpu_last10m = lookup.avg
```

entityName	version	deployment	recentHighCPU	cpu_last10m
CI-dtp-dev3 Cas-3	1.257	1/3/2023, 1:22:39 PM	true	87.3

↓  
OBSERVE

↓  
AUTOMATE

↓  
PREDICT



# AutomationEngine™

Intelligent Automation  
with answers from data.

## SECURE by design

compliance built-in, as data stays  
within Dynatrace platform;  
EdgeConnect for secured remote  
actions

## INTELLIGENT

Answer driven automation  
with Davis Causal AI, precise  
analytics, context, custom  
logic, feedback-loops

## CLOUD-NATIVE

Use SNOW for human centric  
workflows,  
Dynatrace for cloud-native, real-  
time, data-centric

architected for enterprise  
scale with

**MILLIONS**  
of Automations/day

## EFFICIENT

with rapidly growing  
number of  
ready-made ecosystem  
integrations, consolidates  
many home-grown tools

## BizDevSecOps

Automation for Observability, Security  
and Business use-cases. From simple to  
complex, all made easy.

# Workflows



## Graphical workflows

Drag and drop tasks and draw dependencies in no/low code graphical editor



## Schedules

Time-based execution according to flexible rules



## Edge Connect

Connect Dynatrace SaaS securely to on premise systems



## Actions library

Integrate any system using predefined Actions or write your own



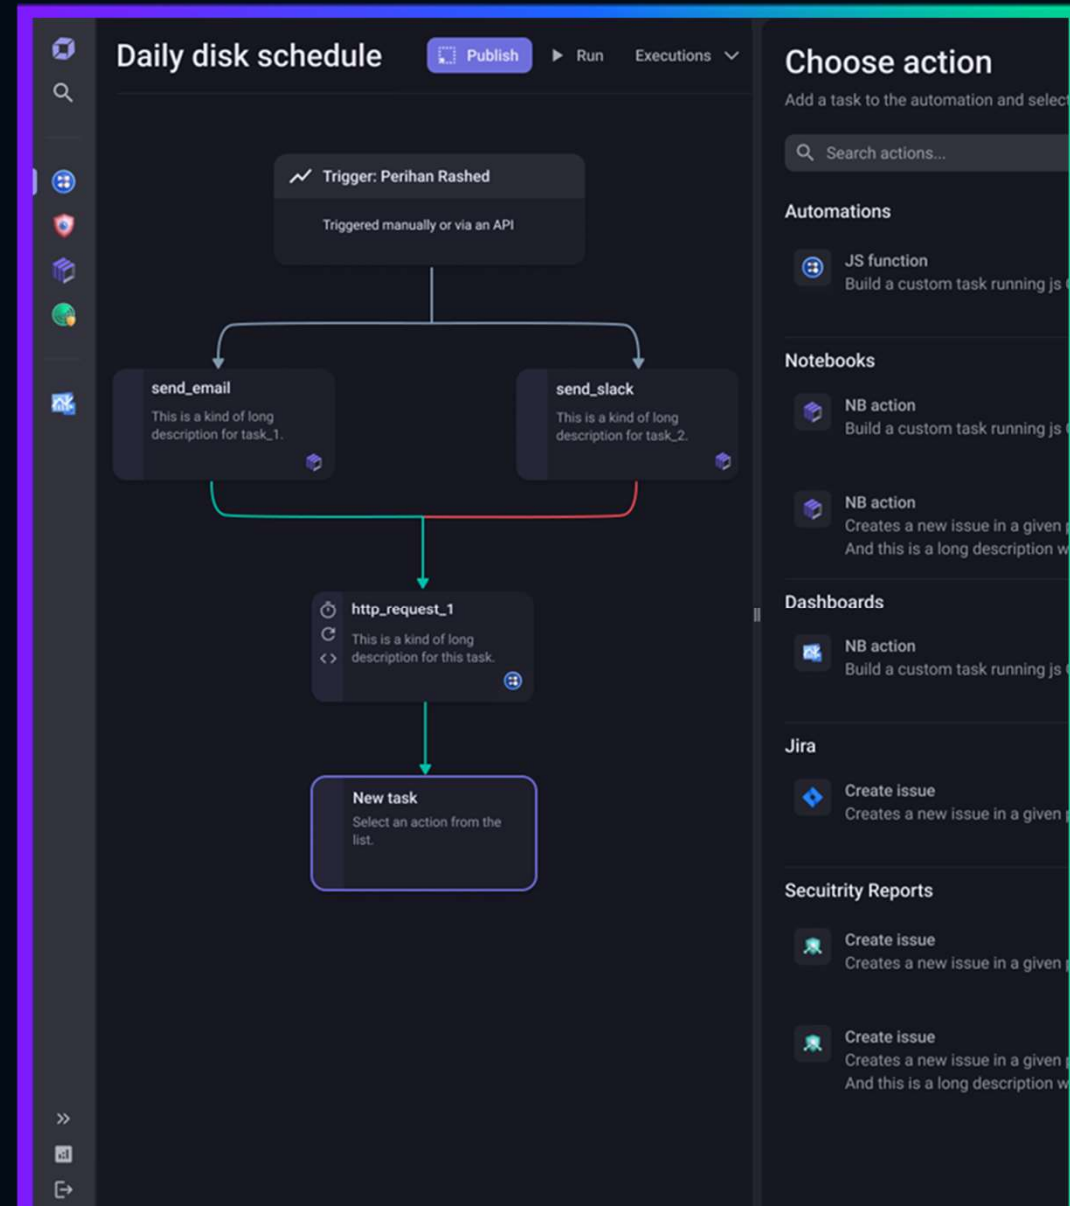
## Event triggers

Automatically react to any Dynatrace Events or Problems

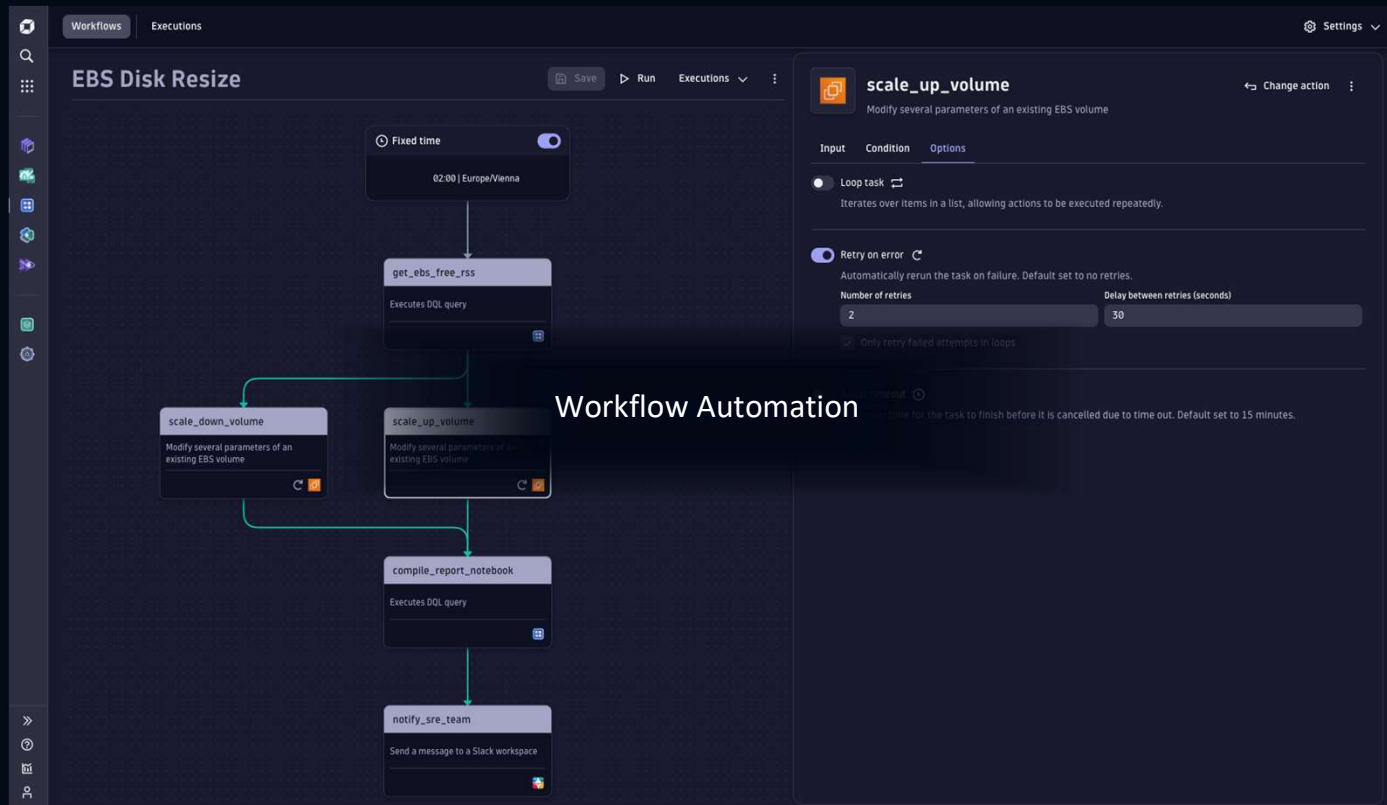


## AI-Trigger

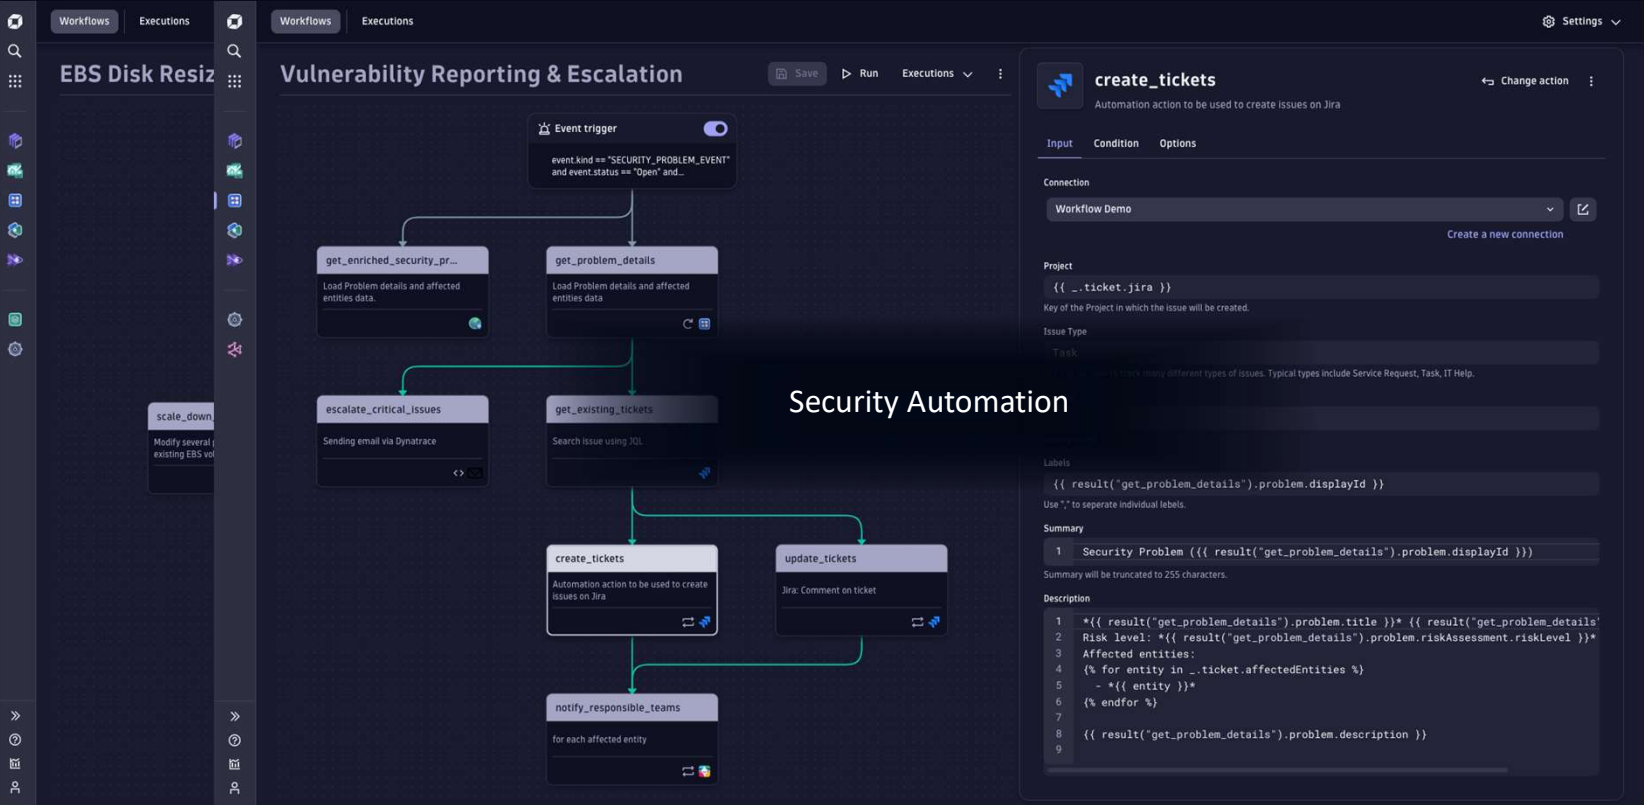
Build and run based on Davis Analyzers



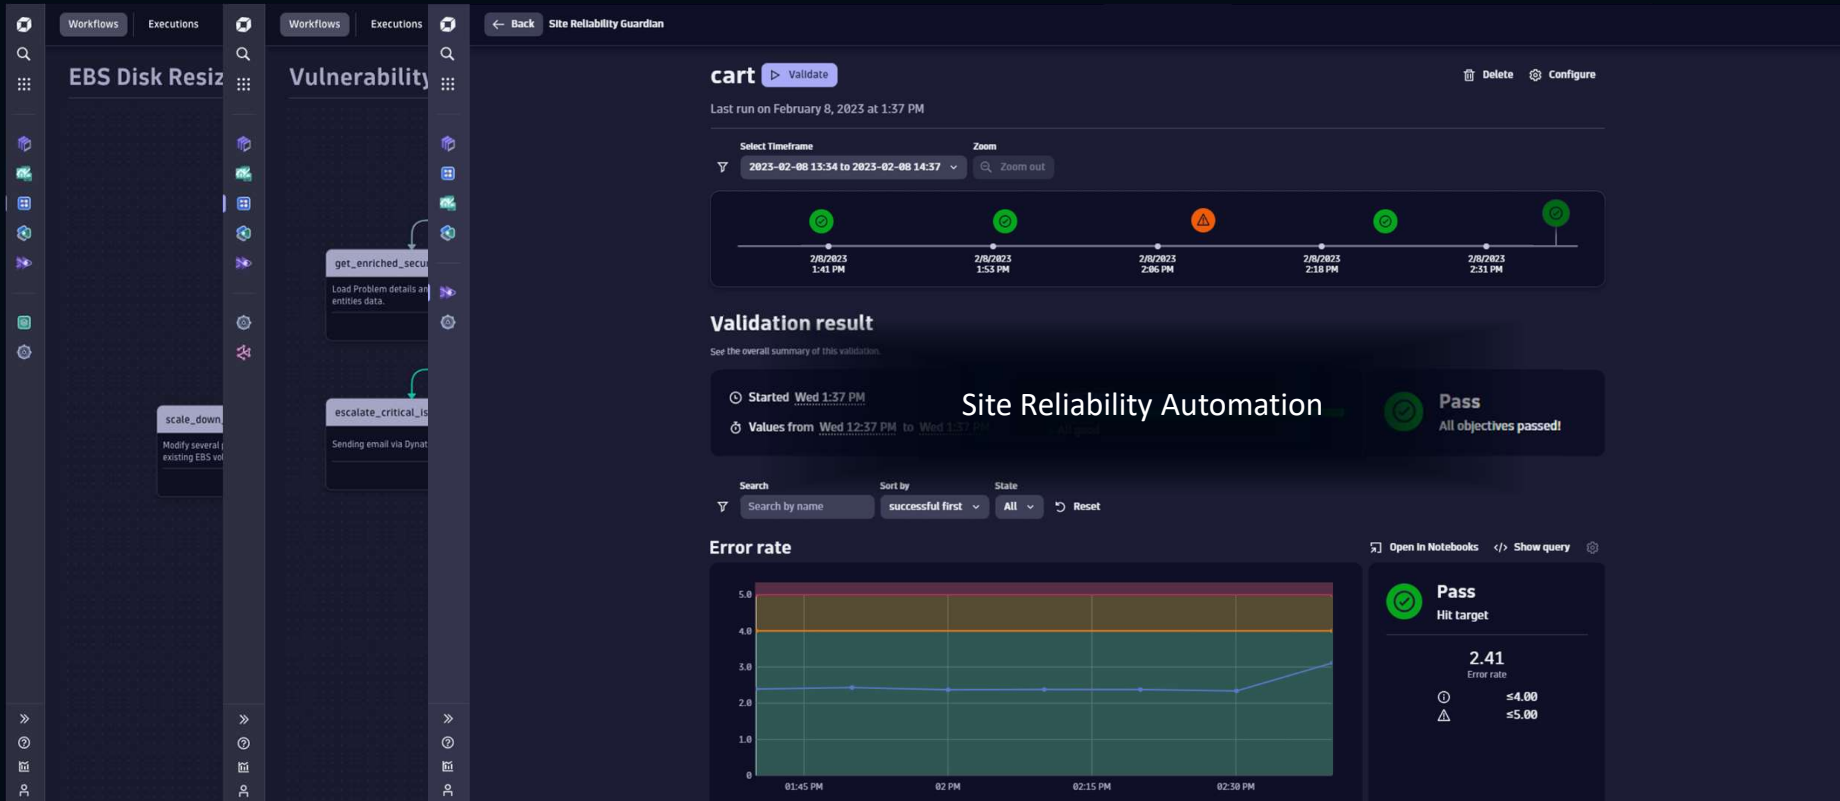
# Use Cases



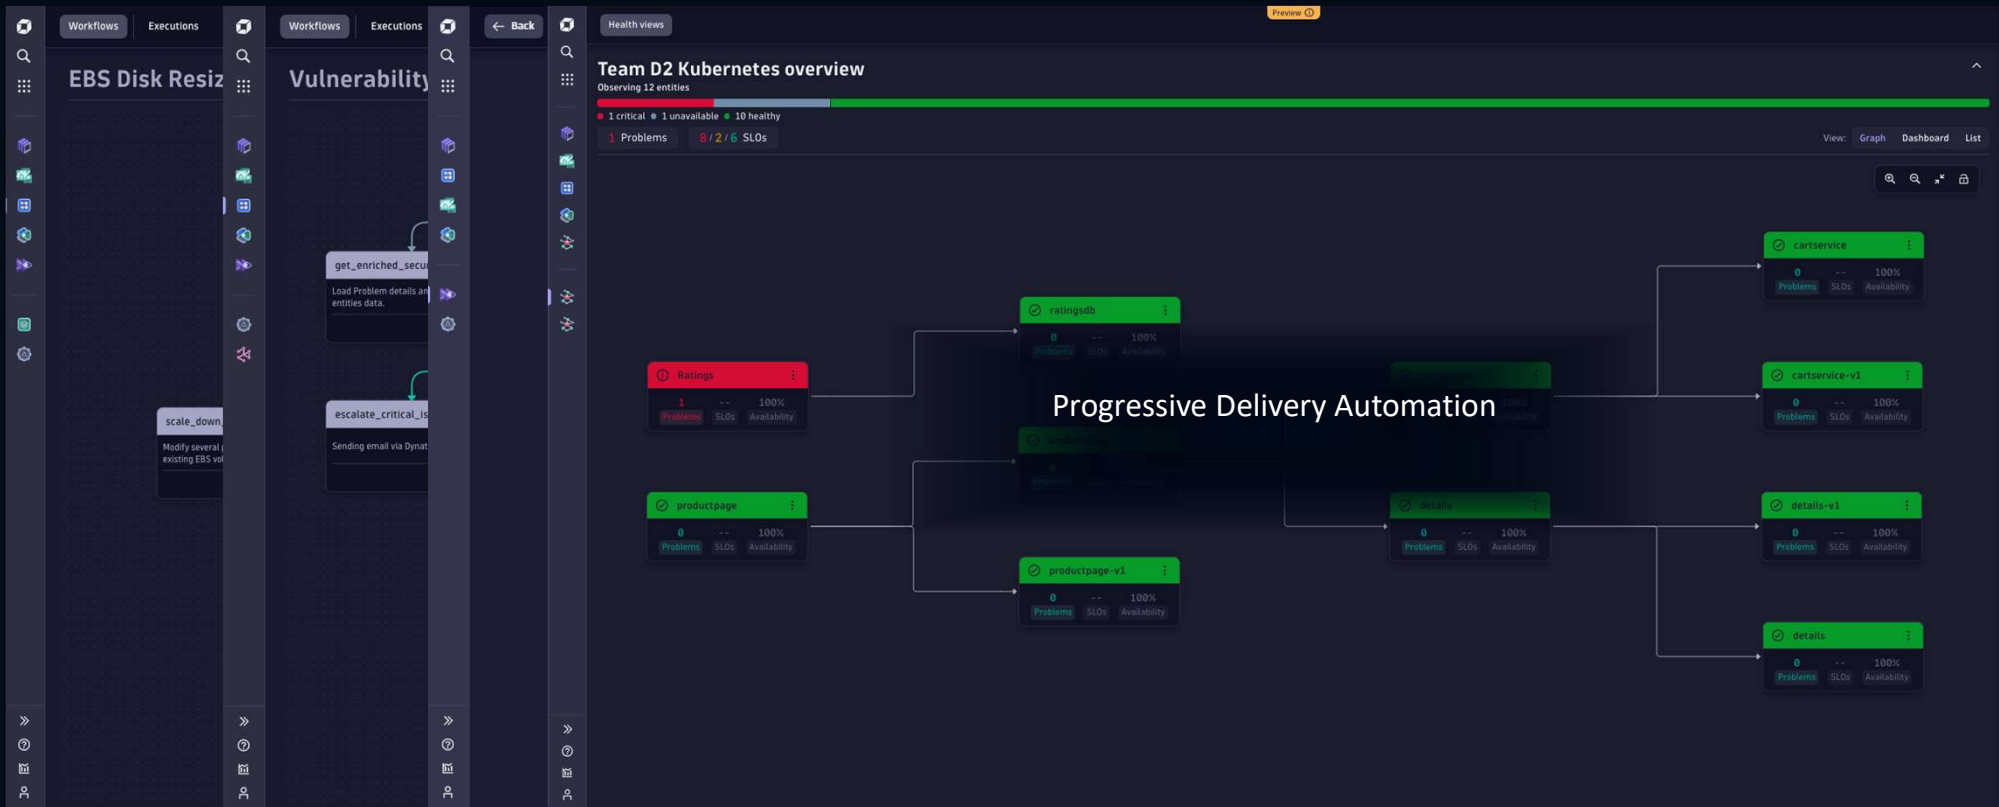
# Use Cases



# Use Cases



# Use Cases



# Use Cases

The screenshot displays a development environment with a dark theme. On the left, there are two panels showing workflow visualizations: 'EBS Disk Resiz' and 'Vulnerability'. The central part of the image shows a code editor with a file explorer on the left listing files like 'application-web', 'auto-tag', 'dashboard', 'management-zone', 'owners', 'ephemeral.yaml', 'owners.json', 'remediations', 'strategies.yaml', 'workflow.json', and 'synthetic-monitor'. The main code editor shows a 'workflow.json' file with a 'remediations' section. A red box highlights a specific configuration within the 'remediations' section, which includes a 'config' object with 'remediations' and 'escalation\_channel' properties. To the right of the code editor, a visual workflow diagram is shown, illustrating the sequence of steps in the automation process. The diagram includes steps like 'run\_validation', 'derive\_ownership\_info', 'message\_on\_warning', 'message\_on\_fail', 'ask\_for\_approval', 'promote', and 'rollback'. Arrows indicate the flow between these steps. A text overlay 'Automation-As-Code e.g. Kubernetes, GitOps' is positioned over the diagram, with arrows pointing to the workflow steps and the code editor.

Automation-As-Code e.g. Kubernetes, GitOps

# AppEngine™



brings custom logic to  
your data to extend use-  
cases and value



# AppEngine™

GDPR, governance, auditing,  
data quality, low-effort yet  
powerful analytics

## Logic to data

high performance, low cost  
eliminated data-flow and  
privacy challenges  
encapsulated data capture,  
analytics and automation

SSO, IAM, hiQgh-  
availability, security,  
secrets, low-code

## Secure app runtime

JavaScript/TypeScript  
functions  
automations  
code shielding technology

BizDevSecOps  
automation,  
shift-left

## Integrates your environment

fetch or import external  
data, execute actions  
secure remote invocations  
with EdgeConnect  
integrations hosted within  
Dynatrace

# Apps for everything and anything

Business Flow

Configurations > 0

Order fulfilment

Key performance

+11

Revenue: 3

Business Flow

Flow Funnel

US-MI

eu-0

us-6

US-1

AT-6

US-1

PL-6

PL-C

PL-F

Choose an option

DevSecOps Lifecycle Coverage

Overview

Percentage of container images are scanned. See which containers have remediation, and govern with a unit development to runtime.

Global

Breakdown of workloads scanned all environments to identify and

9,765 workloads

14% not covered

22% partially covered

Dev

Pre-deployment | Development

8,997 container

95% not covered

Salesforce Insights

Dashboard

Investigate User A

Settings

Verin

The GitHub / failed builds

This app

Using build

Notification

Fetch data

How to view

Try it live

Survey - P

Survey - C

Survey - P

What's next

Fork this

Carbon Impact

Dashboard

Instances

last year

Carbon footprint

3.25 t

CO<sub>2</sub>e in selected time period

2023/01/01 00:00 - 2024/12/01 23:59

3.57 t

CO<sub>2</sub>e in previous time period

2022/01/01 00:00 - 2022/12/01 23:59

-0.32 t

Change

Optimization recommendations

14 entities

Observed idling

View details

12 entities

Identified for scale down

View details

Data center emissions

Filter

Filter table

Data center	Instances idling	Instances scaling	Energy [kWh]	CO2 equivalent
us-east-1-virginia	8	3	2556.84	
us-west-1-california	2	4	1798.11	
eu-west-1-ireland	1	2	1084.63	
apac-sydney-global	0	1	790.53	
latam-brazil-local	1	0	606.95	

Accumulated carbon footprint over time

Carbon versus Business

Carbon

Business

500

400

300

200

100

0

Requests

Errors

Response time

538 ms

What's next?

Fork this app on GitHub and learn how to write apps for Dynatrace.

Fork on GitHub

CONFIDENTIAL

16

Launcher

SearchCTRL K

Apps

Dashboards

Notebooks

Workflows

Hub

Good morning, Grace

What's new

Notebooks

Dashboards

Workflows

Hub

Smartscape Health View

Site Reliability Guardian

Network Observability

DevSecOps Lifecycle Coverage With Snyk

Business Flow

Carbon Impact

Observe and explore

Notebooks

Dashboards

Dashboards Classic

Data Explorer

Metrics

Logs

Smartscape Topology

Containers

Docker

Tests

Technologies and Processes

Network

Infrastructure Monitoring

Kubernetes

Cloud Foundry

AWS

Azure

VMware

Extensions

Network Observability

Cloud Automation

Workflows

Releases

Service-Level Objectives

Site Reliability Guardian

Applications and Microservices

Smartscape Health View

Frontend

Application Security

Security Overview

Third-Party Vulnerabilities

Digital Experience

Mobile

Web

Business Analytics

Business Flow

Carbon Impact

My custom apps

My First Custom App

Cloud

17

New Dynatrace experience

Strato Design System

Dashboards

Notebooks

# Notebooks

Petabyte-scale data exploration & analytics for real-time insights.

- Notebooks are interactive apps
  1. Write a DQL (Dynatrace Query Language) query or code snippet
  2. Get the resulting table or chart immediately.
- Ask follow-on questions at the speed of your thoughts, get trend forecasts and assistance by Davis AI.
  - Query, analyze, and visualize all data in Dynatrace and even integrate external data.
- Collaborate with teams using interactive, data-driven notebooks with comments.

DQL Query

Resulting chart

Dynatrace function

Resulting chart

Comments

The screenshot displays the 'My Analytic Notebook' interface in Dynatrace. The notebook is titled 'My Analytic Notebook' and was saved 4 minutes ago. It contains the following sections:

- DQL Query:** A query to fetch logs from the current time minus 90 minutes, filtered by log level 'WARN' or 'ERROR', summarized by 1-minute bins, and sorted by count descending. The query is: 

```
1 fetch logs, from:now()-90m
2 | filter loglevel == "WARN" or loglevel == "ERROR"
3 | summarize count = count(), by:{bin(timestamp, 1m), loglevel}
4 | sort count desc
```
- Resulting chart:** A bar chart showing the count of logs over time, with the y-axis ranging from 0 to 15.0 and the x-axis showing time from 12:00 PM to 1:00 PM.
- Dynatrace function:** A code snippet for an async function that fetches weather data from the OpenWeatherMap API. The code is: 

```
1 export default async () => {
2   const json = await fetch('https://api.openweathermap.org/data
3   return json.main.temp;
4 }
```
- Resulting chart:** A large display showing the current temperature as 23.52°C.
- Comments:** A section titled '# Summary' containing a text block: 'This is what we've found using the following query' followed by a code block: 

```
1 # Summary
2
3 This is what we've found using the following query
4 ```dql
5 fetch logs
6 | filter dt.entity.host == "HOST-E73A576280286480"
7 | summarize count = count(), by: {'1m interval' = bin(timestamp, 1m)}
8 ```
```

## COLLECT

Observability, Security  
and Business data



Traces



Metrics



Logs

+



Topology



Behaviour



Code



Metadata



Network



Problems

+



Business  
Events



Threats



Vulnerabilities

Deep, context rich, full stack

## CONTEXTUALIZE

Scalable data  
capture and ingest



OneAgent



API



OpenTelemetry



Semantic Dictionary



Ecosystem Integrations



Synthetic

Automatically captured in  
context & pre-processed

## PROCESS

Graph powered Massive  
Parallel Processing  
Lakehouse



Grail

Stored contextually  
with massive parallel  
processing and graph  
query technology

## ANALYZE

Causal & Predictive AI



Davis AI



Smartscape

Accessed by our causal  
AI for analysis &  
answers

## AUTOMATE

Enterprise low-code  
programmability and  
data-science



AutomationEngine



AppEngine



Dynatrace  
Query Language

Post-process,  
combine, act

## SOLVE

Answer driven  
Automation and  
Security



Observability



Security



Automation



Custom Apps

Powering automation,  
orchestration,  
prevention and protection



Hypermodal AI for Unified Observability and Security



## Deterministic Answers & Reliable Automation

Problems, Risks, AI SecOps, Root-Cause

### 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

Inquiry

Auto-Prompt

Context

## Productivity & Recommendations

Ideation, research and creation



Davis® CoPilot

### Generative AI



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

CONFIDENTIAL



20



CLOUD DONE RIGHT