Red Hat Summit

Connect

dynatrace

Optimiza tus Operaciones de TI

Automatización inteligente con IA Davis y Ansible Automation para una mayor eficiencia

Miguel Balsa
Solutions Engineer
DYNATRACE

Red Hat

Red Hat & dynatrace

Pre-integrated full stack solutions

Red Hat OpenShift

Red Hat Enterprise Linux

Red Hat Ansible Automation Platform

dynatrace Software Intelligence Platform

OneAgent | PurePath | Smartscape | Grail | Davis AI incl. AIOps

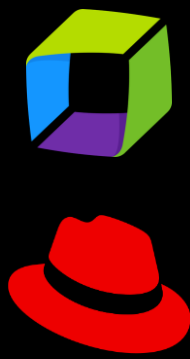

Combining to solve cloud-native challenges

Achieve better business outcomes

Manage microservice and cloud complexity at enterprise scale

Automate IT operations intelligently with causal AI and precise insights

Exceed DevOps and SRE service objectives with one tightly integrated solution

What is Auto-remediation?

Auto-remediation, or **self-healing**, is a workflow that triggers and responds to alerts or events by executing actions that can prevent or fix an issue.

Auto-remediation significantly **reduces MTTR**.

Types:

1. Automated remediation of a **know problem** or **frequent issue (proactive)**
2. Automated rollback of a **problematic change (reactive)**

The timeline of an incident

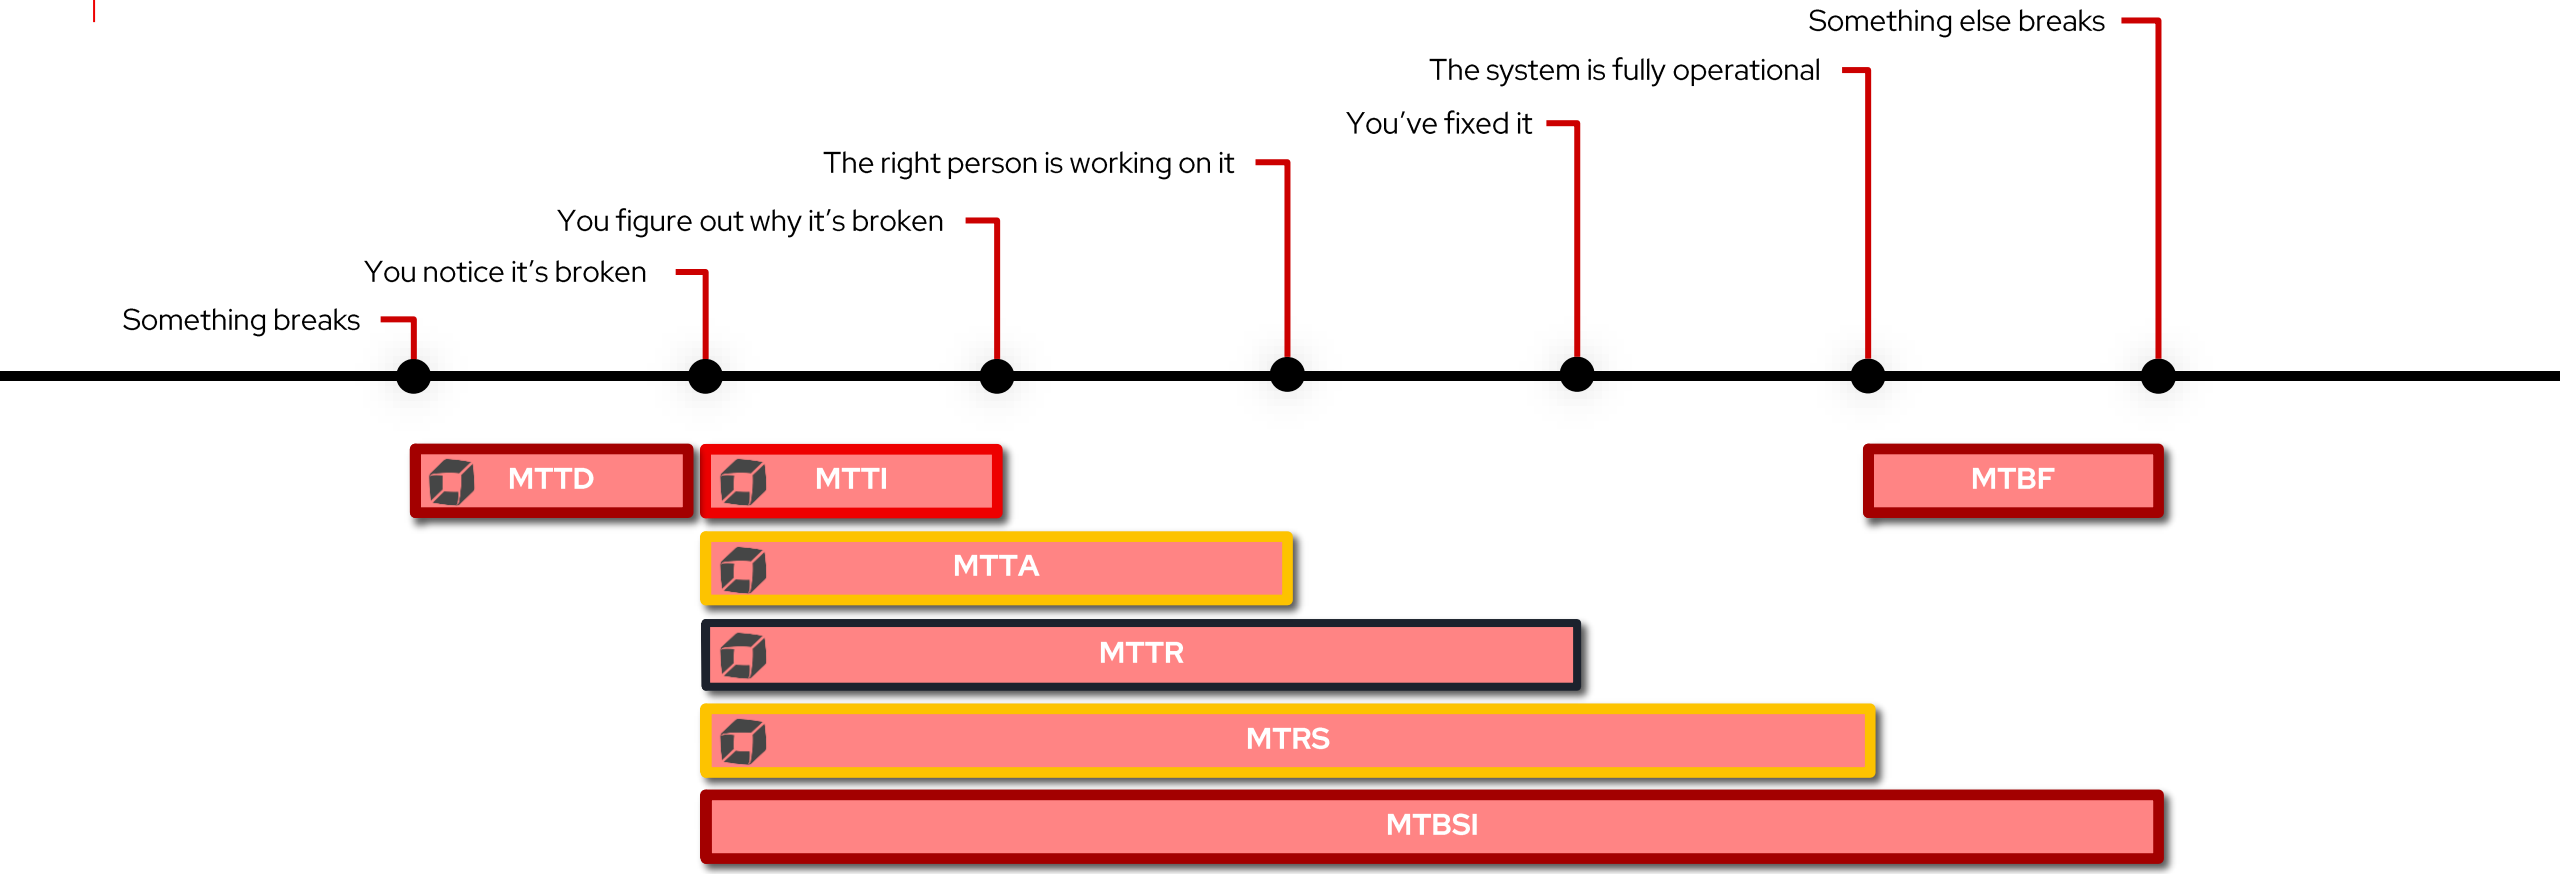

MMTD: Mean Time to Detect MTTI: Mean Time to Investigate MTTA: Mean Time to Action MTTR: Mean Time to Repair MTRS: Mean Time to Restore Service MTBF: Mean Time Between Failures MTBSI: Mean Time Between Service Incidents

The Diagnostics of the unknowns

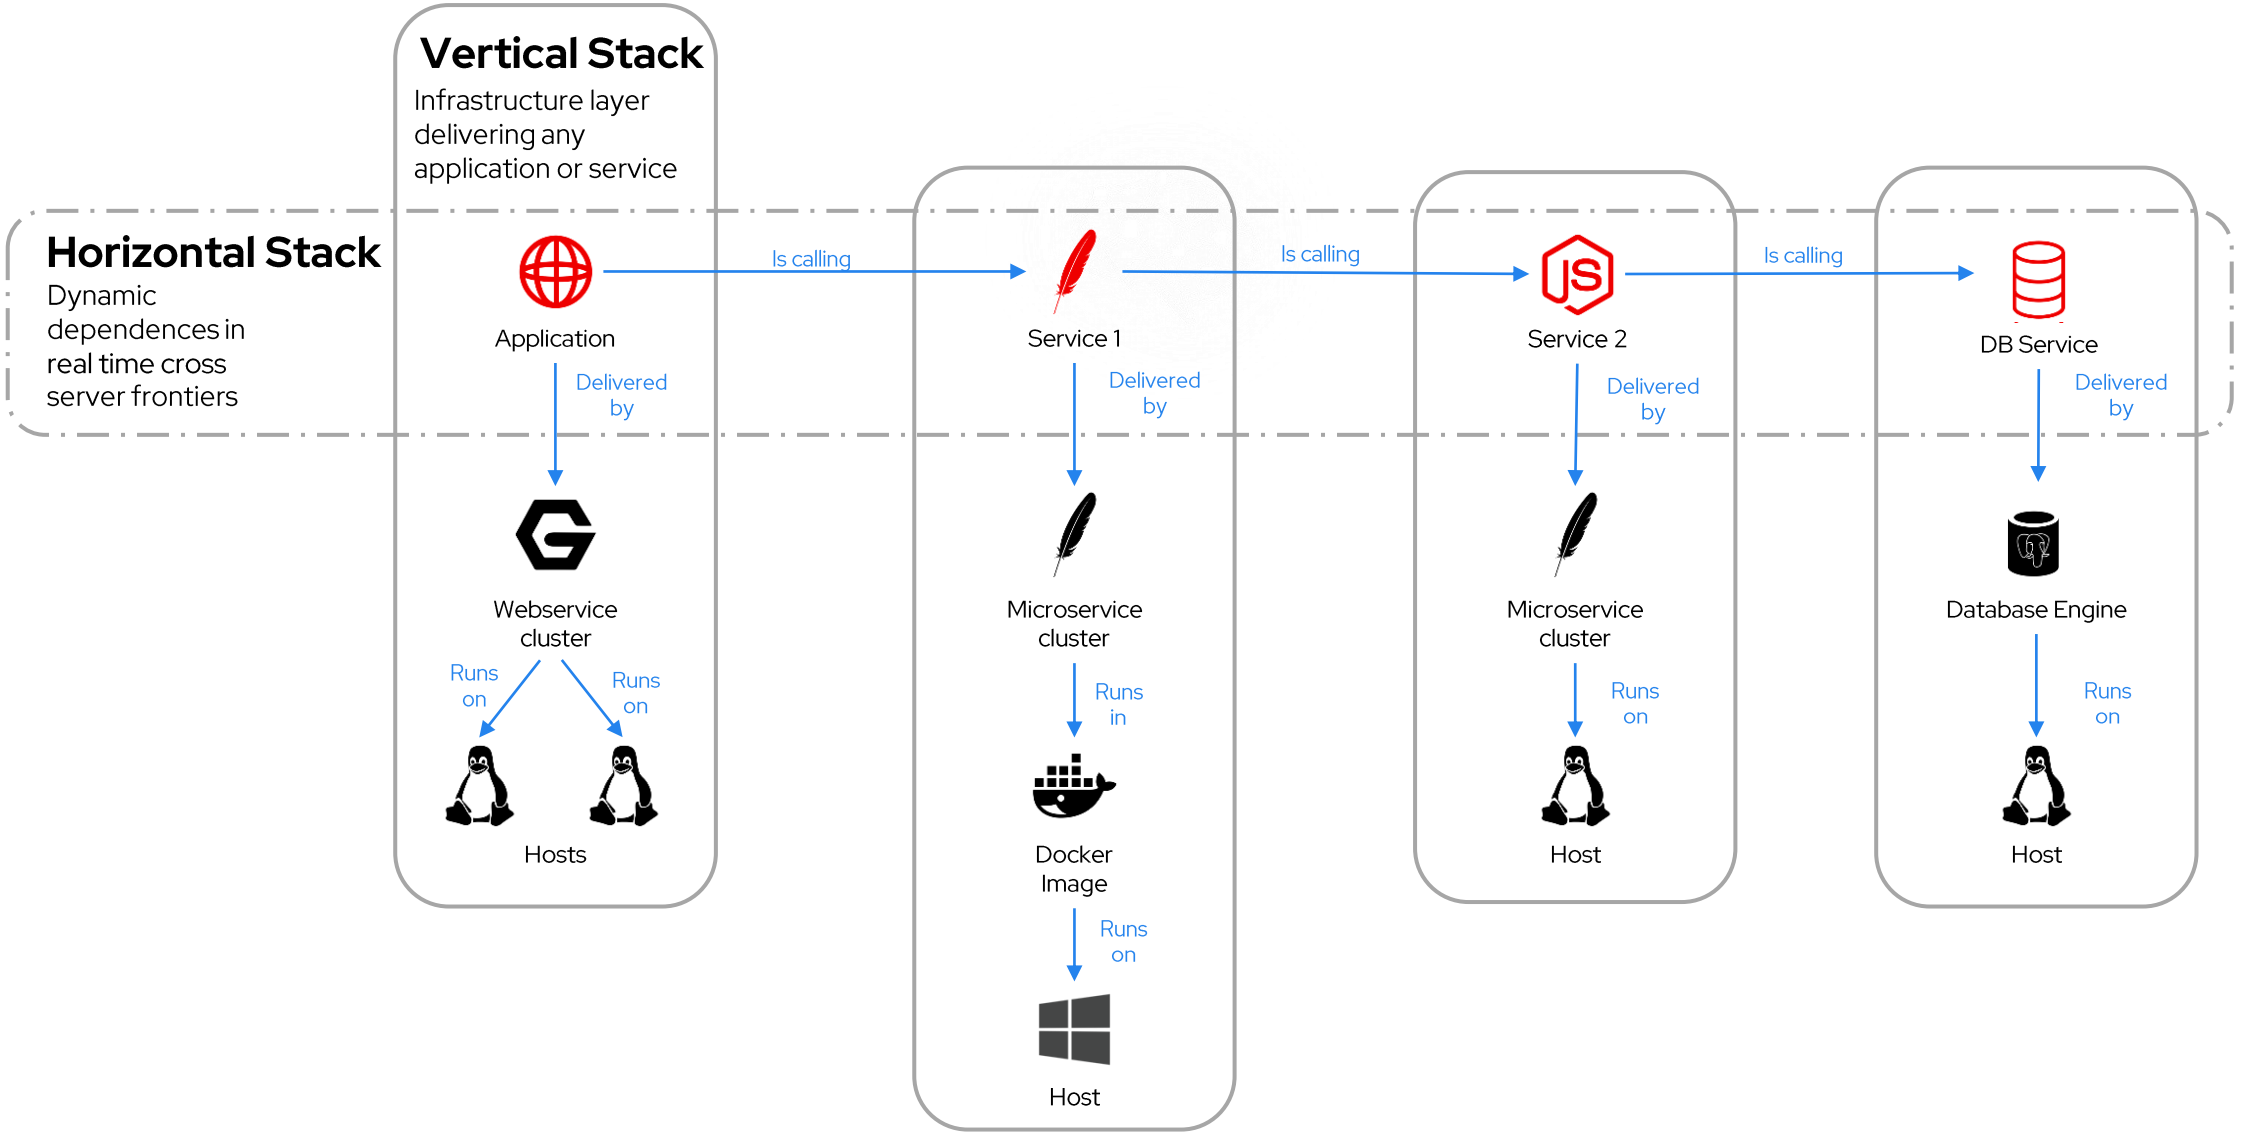

GO BEYOND DASHBOARDS AND GET PROACTIVE ANSWERS

Dynatrace continuously observes, learns and auto-adapts to changes in real-time to detect problems automatically (even the ones you never thought of.)

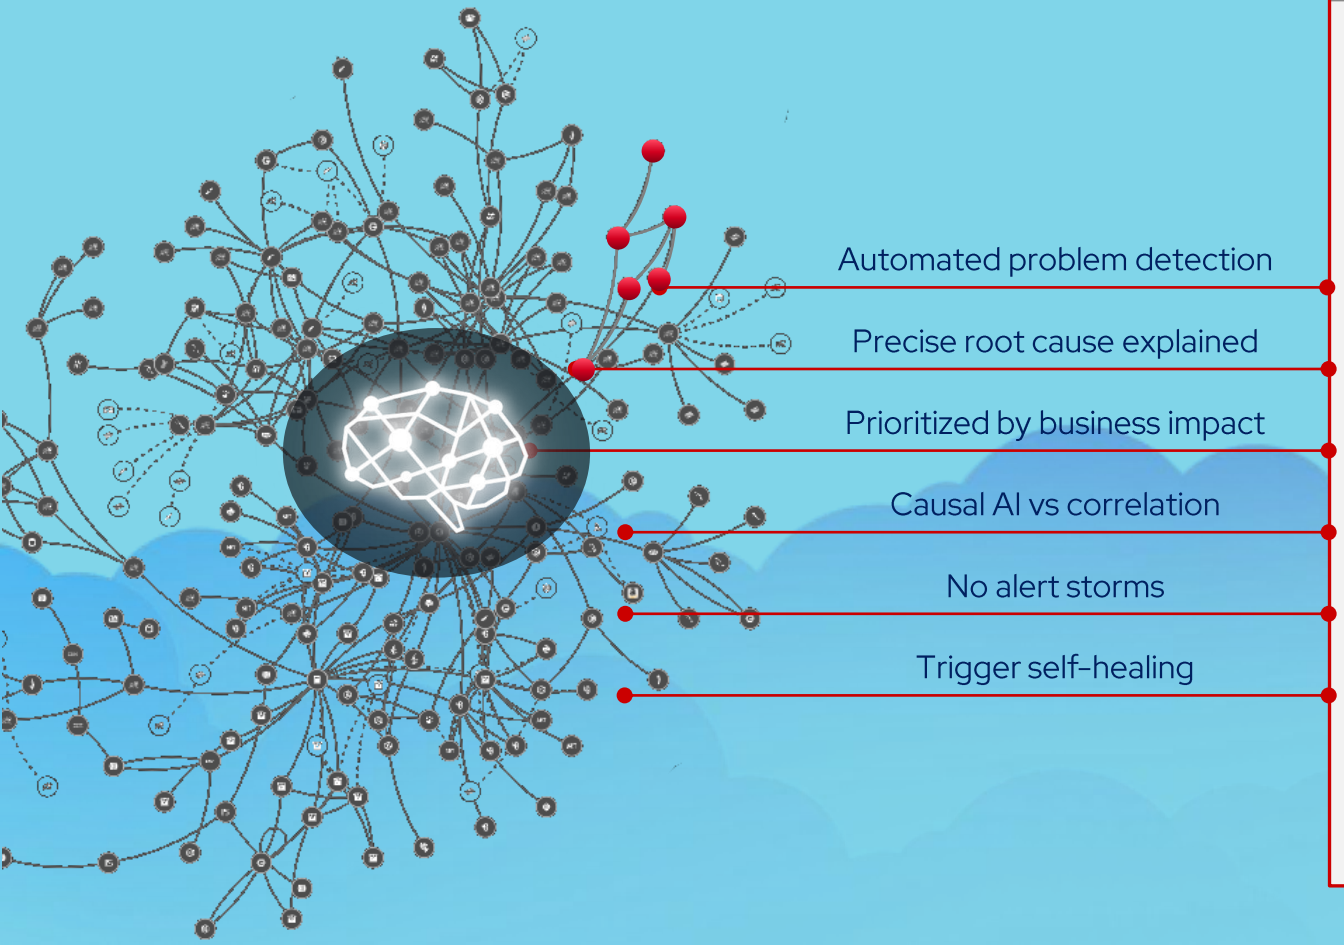

- Automated problem detection
- Precise root cause explained
- Prioritized by business impact
- Causal AI vs correlation
- No alert storms
- Trigger self-healing

> "Dynatrace's AI-driven answers are helping the business to remediate problems quicker, meaning we're able to spend more time innovating and less time solving problems."
>
> — Anish Patel, Principal Systems Engineer CENGAGE

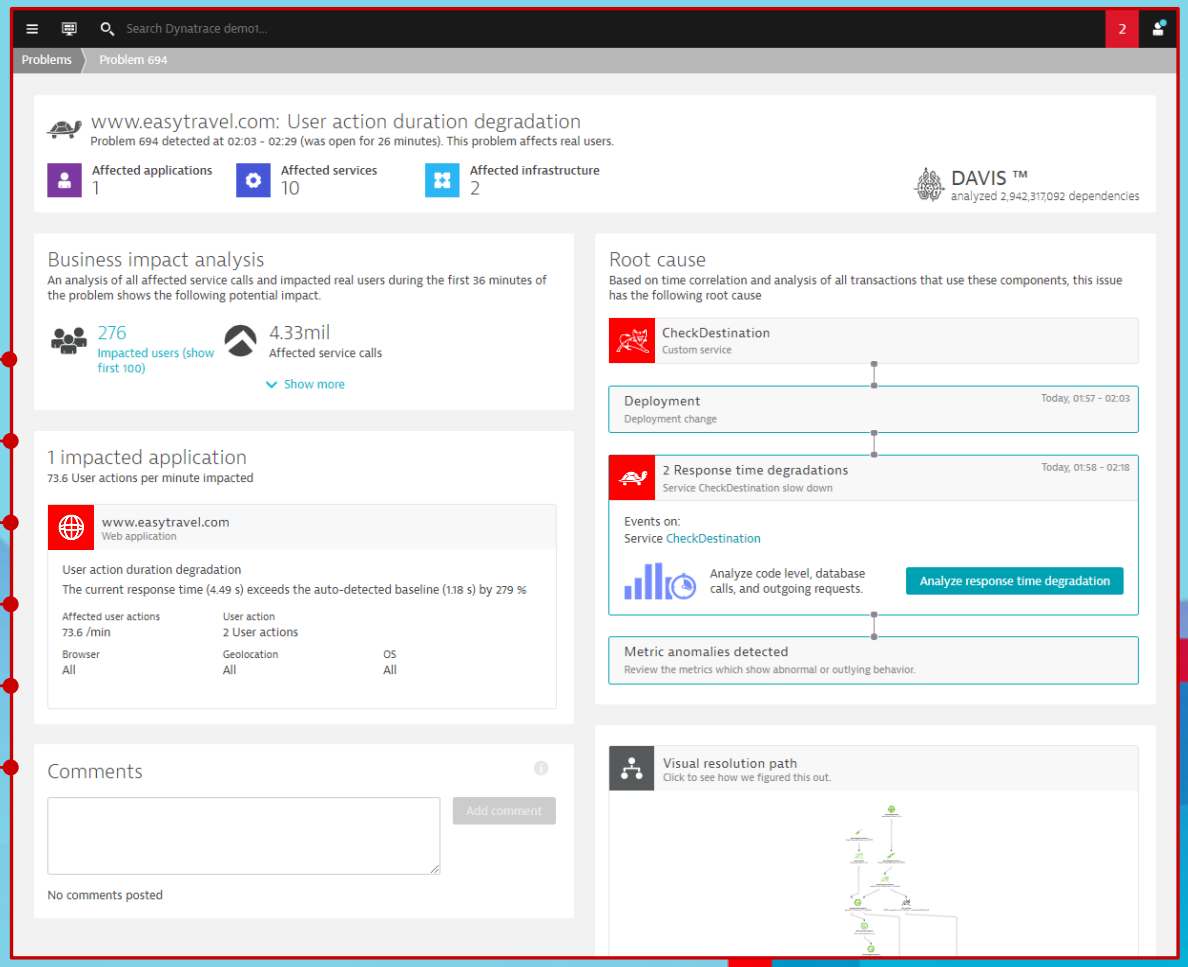

Same problems seem to repeat themselves

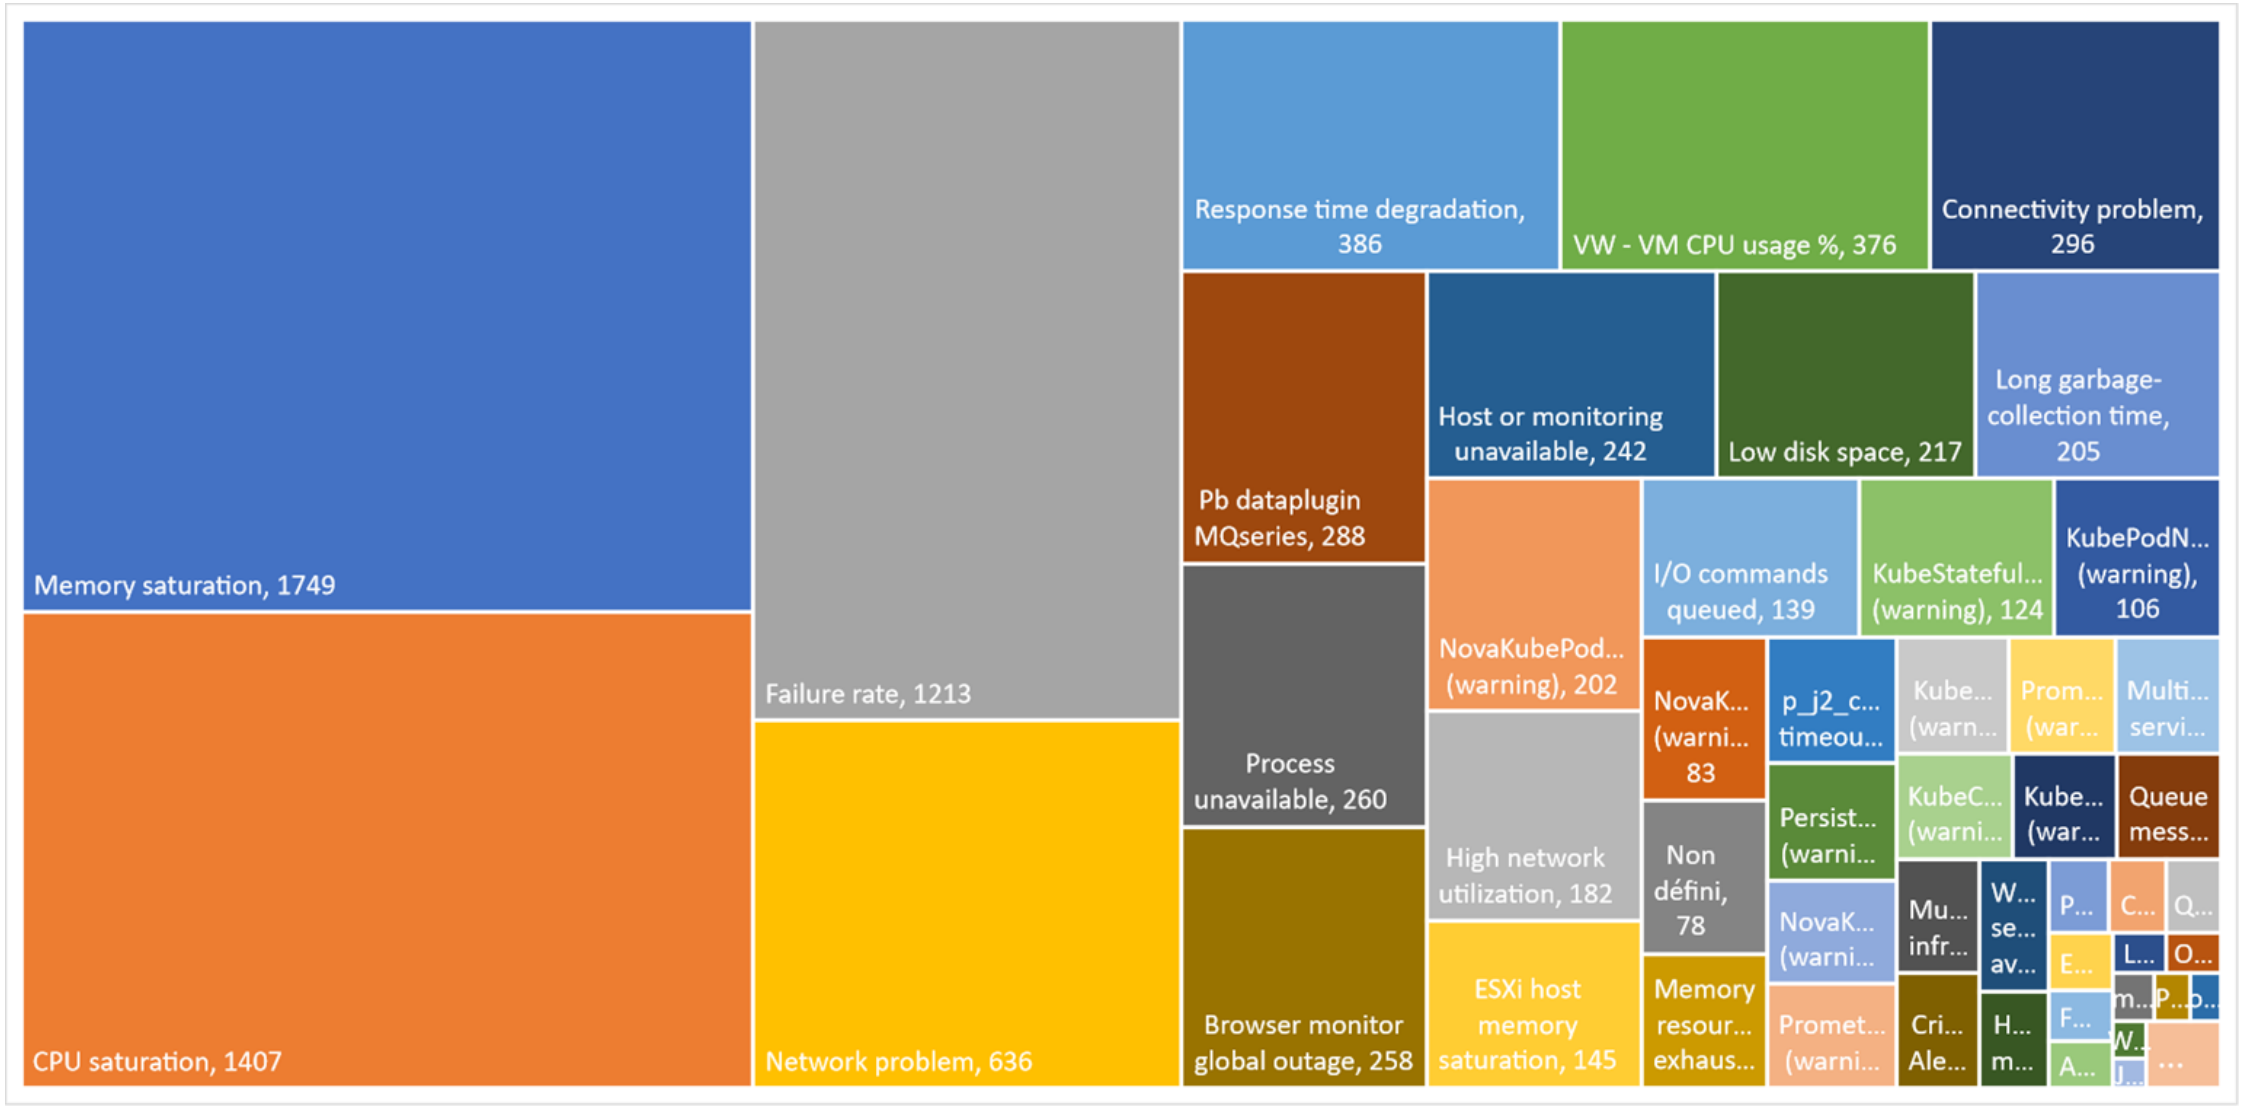

AUTOMATE YOUR PROBLEM REMEDIATION

Dynatrace automatically eliminates the noise and provides the precise, reliable root cause which is essential for self-healing, autonomous operations.

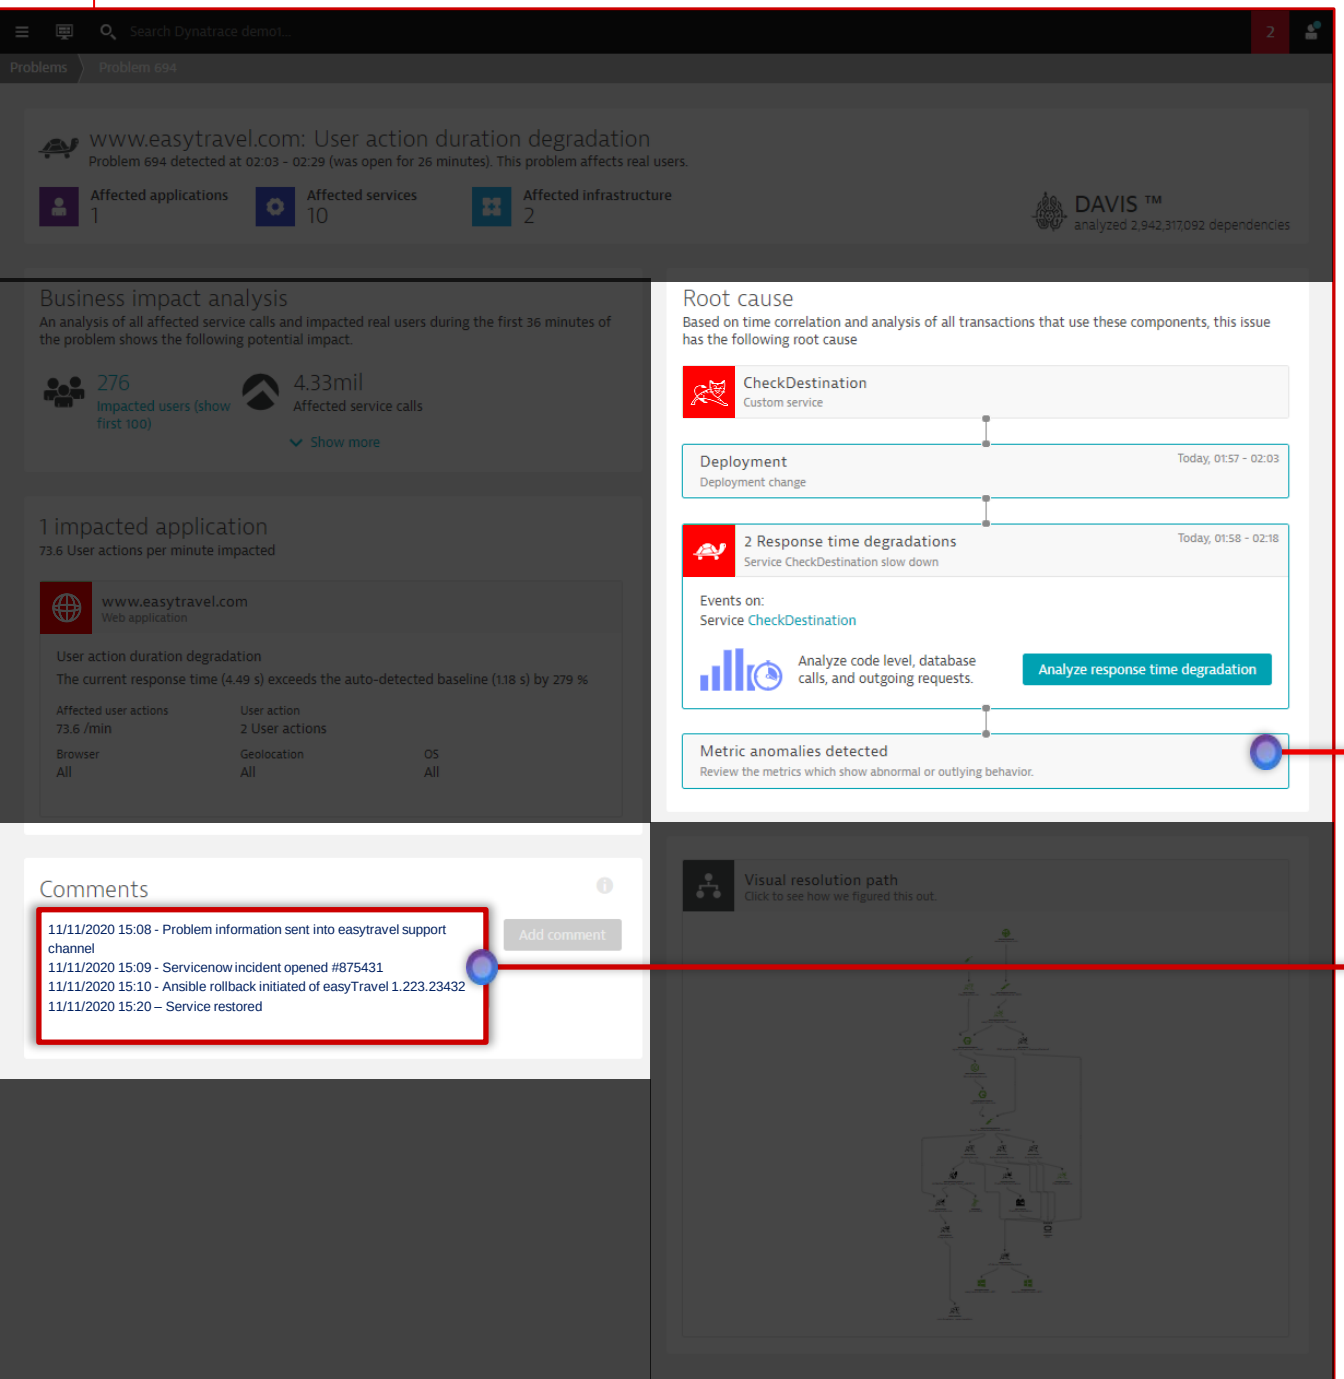

Integrations

Open problem ticket and assign it to the right team

Send details into support channel

Trigger Correct Remediation Action

Benefits of Auto-Remediation

dynatrace

1. Reduction of alerts, incident tickets and MTTR > **Cost reduction**
2. Ability to preempt outages, improved uptime
3. Reduction of risk and human error
4. **Improved user experiences**

Red Hat

New Innovation

for an even better-together value for our customers

dynatrace

Red Hat

Grail

Architected with game-changing attributes

1st	**1000 petabytes/ day ***	**BILLIONS**	**TRILLION**	**0 SECS**	**5-100x****	**90%*****
All-in-one platform that brings observability, security, and business together	all retained for more than 12 months	of dependencies for analytics in context and causal AI	cardinalities over a years period	no index, no rehydration, no schema creation, no extra parsing step	faster to query massive parallel processing (MPP) on 1000s of nodes	Reduced planning and configuration work in complex environments
PLATFORM PLATFORM PLATFORM PLATFORM PLATFORM	INGEST INGEST INGEST INGEST INGEST INGEST INGEST INGEST INGEST INGEST	CONTEXT CONTEXT CONTEXT CONTEXT CONTEXT	CARDINALITY CARDINAL CARDINALITY CARDINAL CARDINALITY	SCALE SCALE SCALE SCALE SCALE SCALE SCALE SCALE SCALE SCALE	SPEED SPEED SPEED SPEED SPEED SPEED SPEED SPEED SPEED SPEED	EASE EASE EASE EASE EASE EASE EASE EASE EASE EASE EASE EASE EASE EASE EASE

*The Grail™ architecture enables Dynatrace to scale an entire multi-tenant cluster to that dimension, however we reserve right to deliver based on actual market need

** potential combined performance gain on complex queries, leveraging massive-parallel processing, faster parsing compared to regular expressions, and large datasets that would require rehydration or similar

*** up to 90% reduced work due to unified Grail storage in contrast to common siloed data storage approaches with preparing ingest data-sources like metrics, traces and logs for pre-aggregation and elimination of high cardinalities

Dynatrace Query Language (DQL)

Purpose-built
for observability and security

Powerful
for even advanced use-cases, parsing built-in

Effortless
migration from e.g. Splunk

```
fetch       logs, from:now()-20m
| filter    endsWith(log.source,"/media/datastore/server-data/log/audit.config.change.
            and dt.host_group.id == "cluster_deve2e"
| parse     content, "timestamp('yyyy-MM-dd HH:mm:ss'):ts
                      ld json:settings
                      ipaddr:client_ip //IPv4/6"
| fields    ts,
            type = settings[eventType],
            tenant = settings[tenantId],
            user = settings[userId],
            change = settings[jsonPatch]
| filter    in(type,array("UPDATE","DELETE")) and user != "unknown"
| summarize creates = countIf(type=="CREATE"), upd = countIf(type=="UPDATE"), del = co
            by:{tenant, user}
| fieldsAdd changes_per_min = (upd + del)/20
| sort      changes_per_min desc
```

DQL

dynatrace

Red Hat

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

Unique Kubernetes Orchestration with Dynatrace Kubernetes Operator and Keptn Life-Cycle Toolkit

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.

AutomationEngine

dynatrace

RED HAT ANSIBLE & DYNATRACE 'INTELLIGENT OPERATIONS' ARCHITECTURE

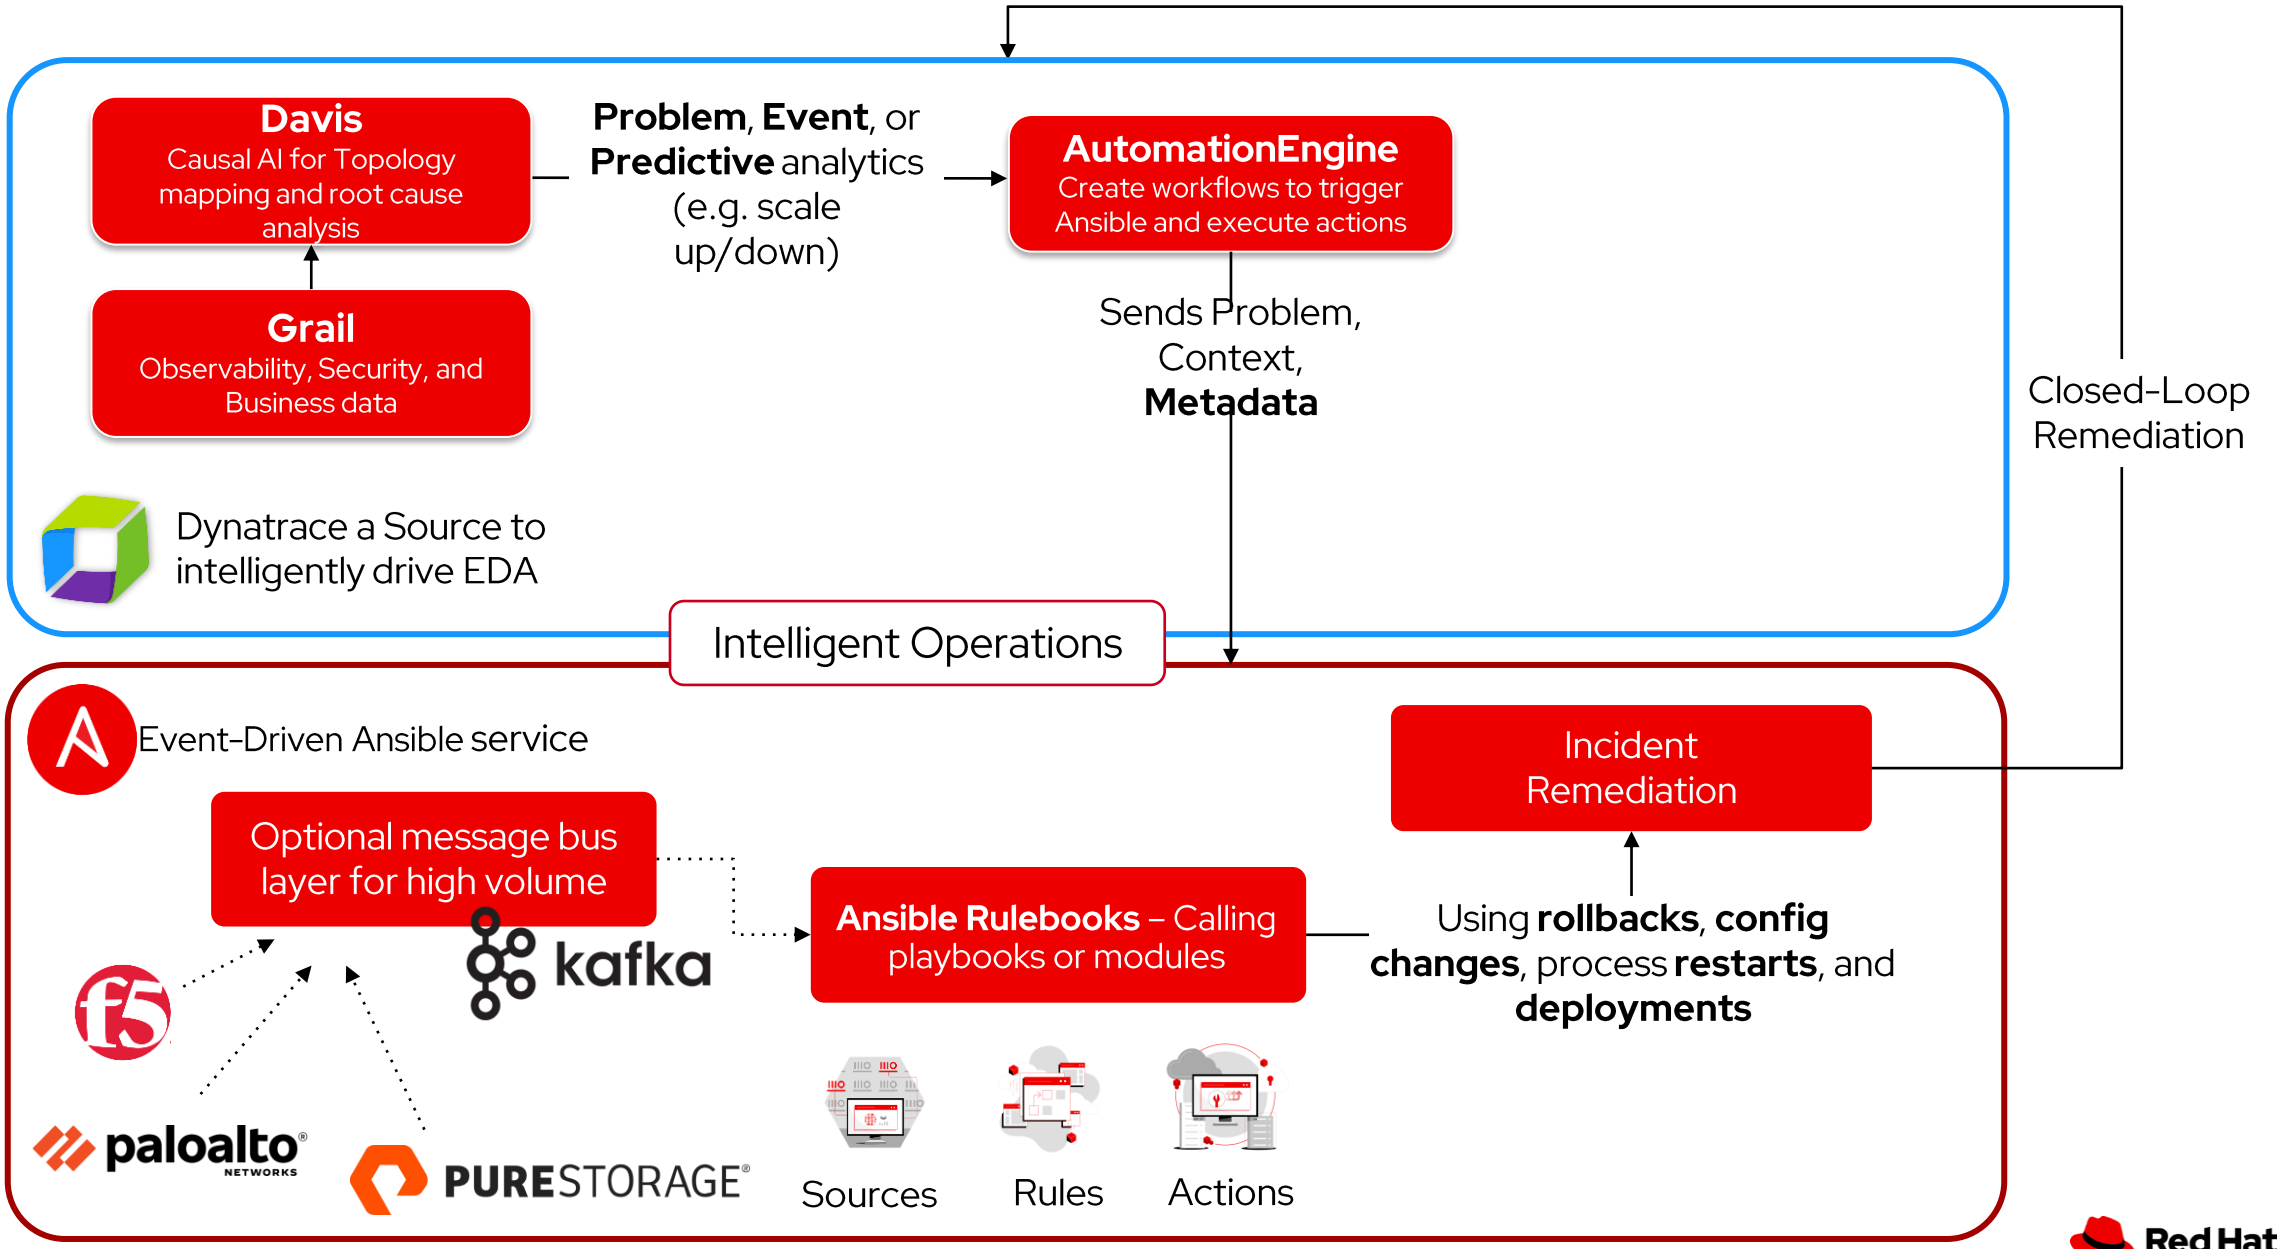

Automated Problem Remediation

Root Cause Identification

Automatically provide high fidelity root cause and configuration item (CI) data with **Dynatrace problem events**

Problem Remediation

Ansible automatically **triggers remediation scripts to resolve** common application problems

Problem Recovery

Automatically log remediation activities, validate recovery, and communicate status

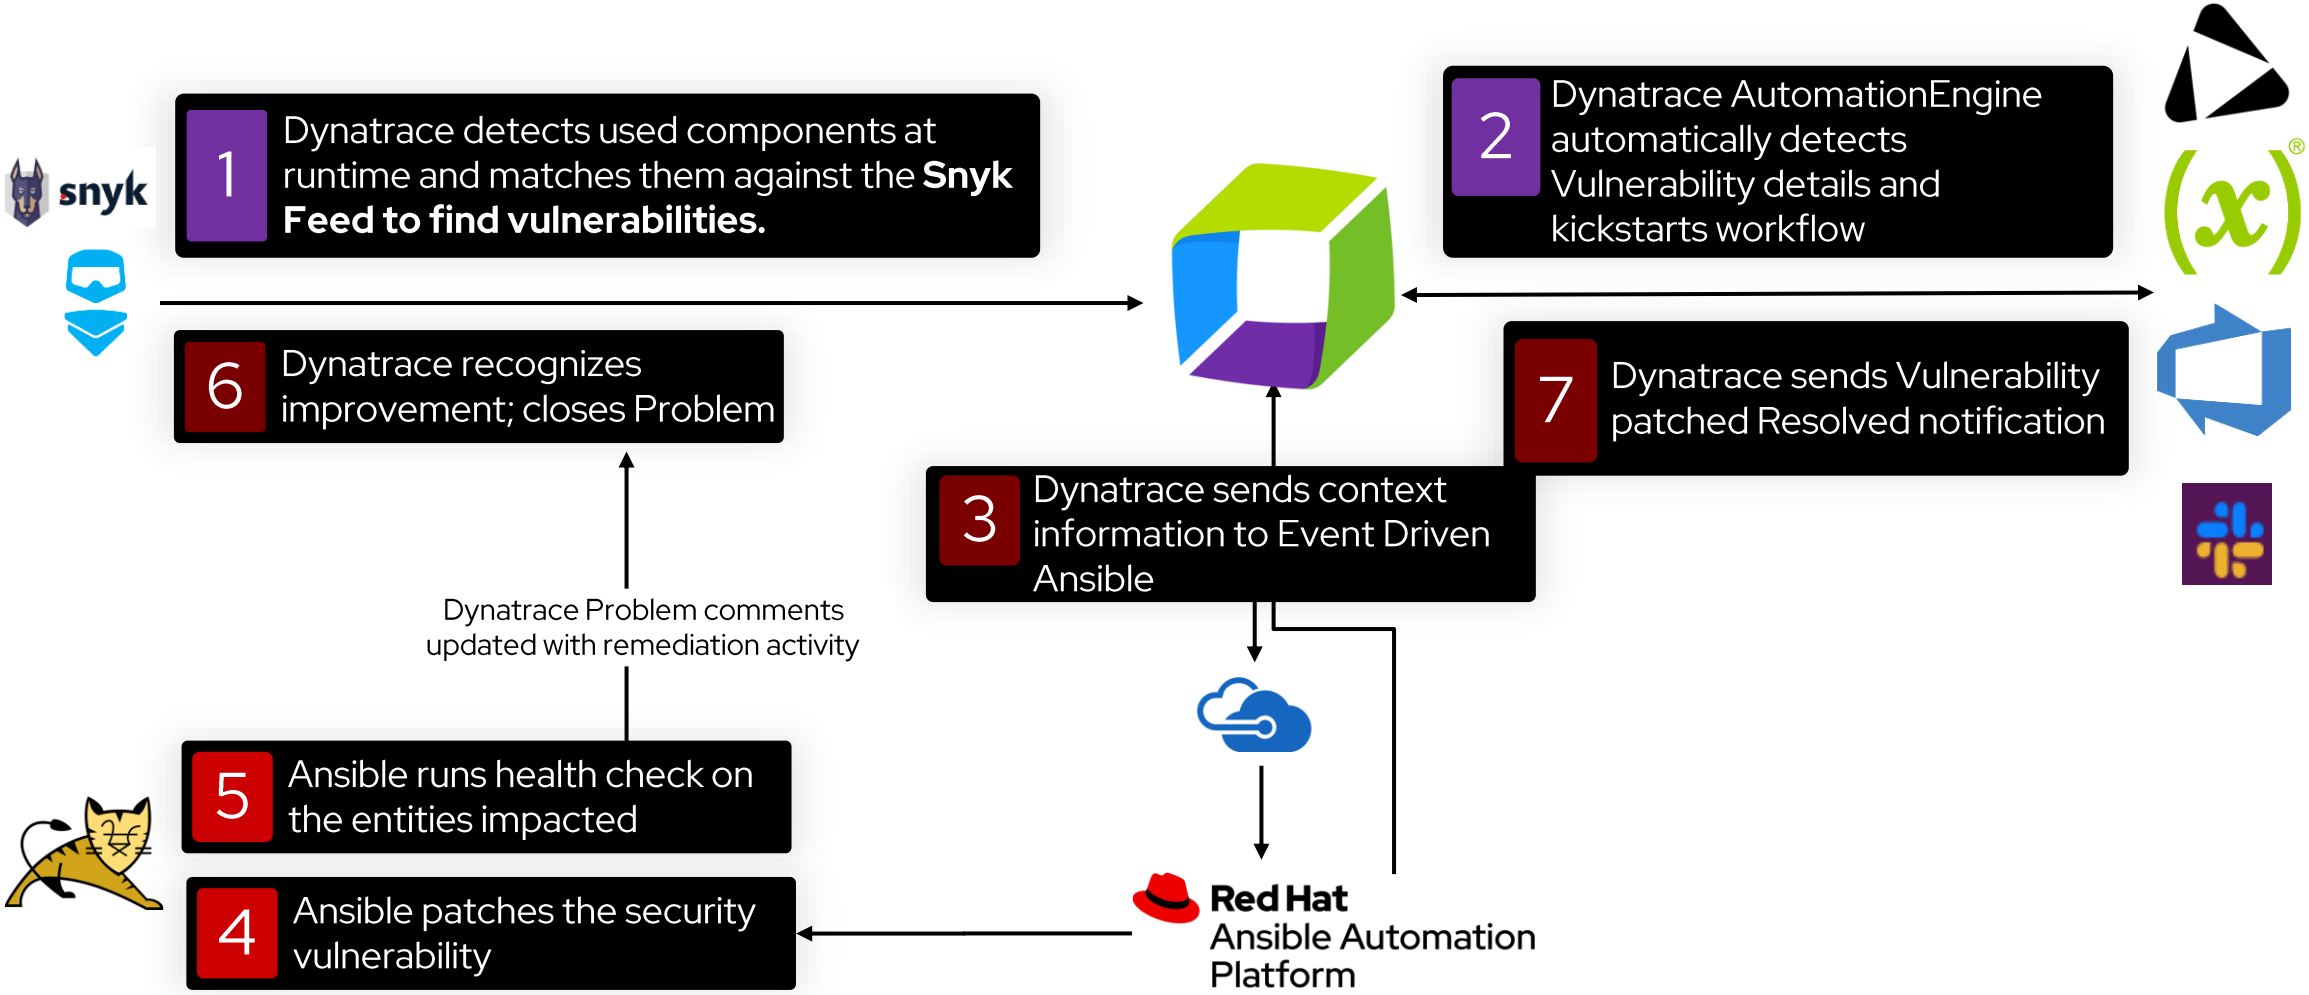
snyk
1 Dynatrace detects used components at runtime and matches them against the **Snyk Feed to find vulnerabilities.**
2 Dynatrace AutomationEngine automatically detects Vulnerability details and kickstarts workflow
6 Dynatrace recognizes improvement; closes Problem
7 Dynatrace sends Vulnerability patched Resolved notification
3 Dynatrace sends context information to Event Driven Ansible
Dynatrace Problem comments updated with remediation activity
5 Ansible runs health check on the entities impacted
4 Ansible patches the security vulnerability
Red Hat Ansible Automation Platform

Red Hat

Solution Architecture for APP HEALING FOR a CUSTOMER

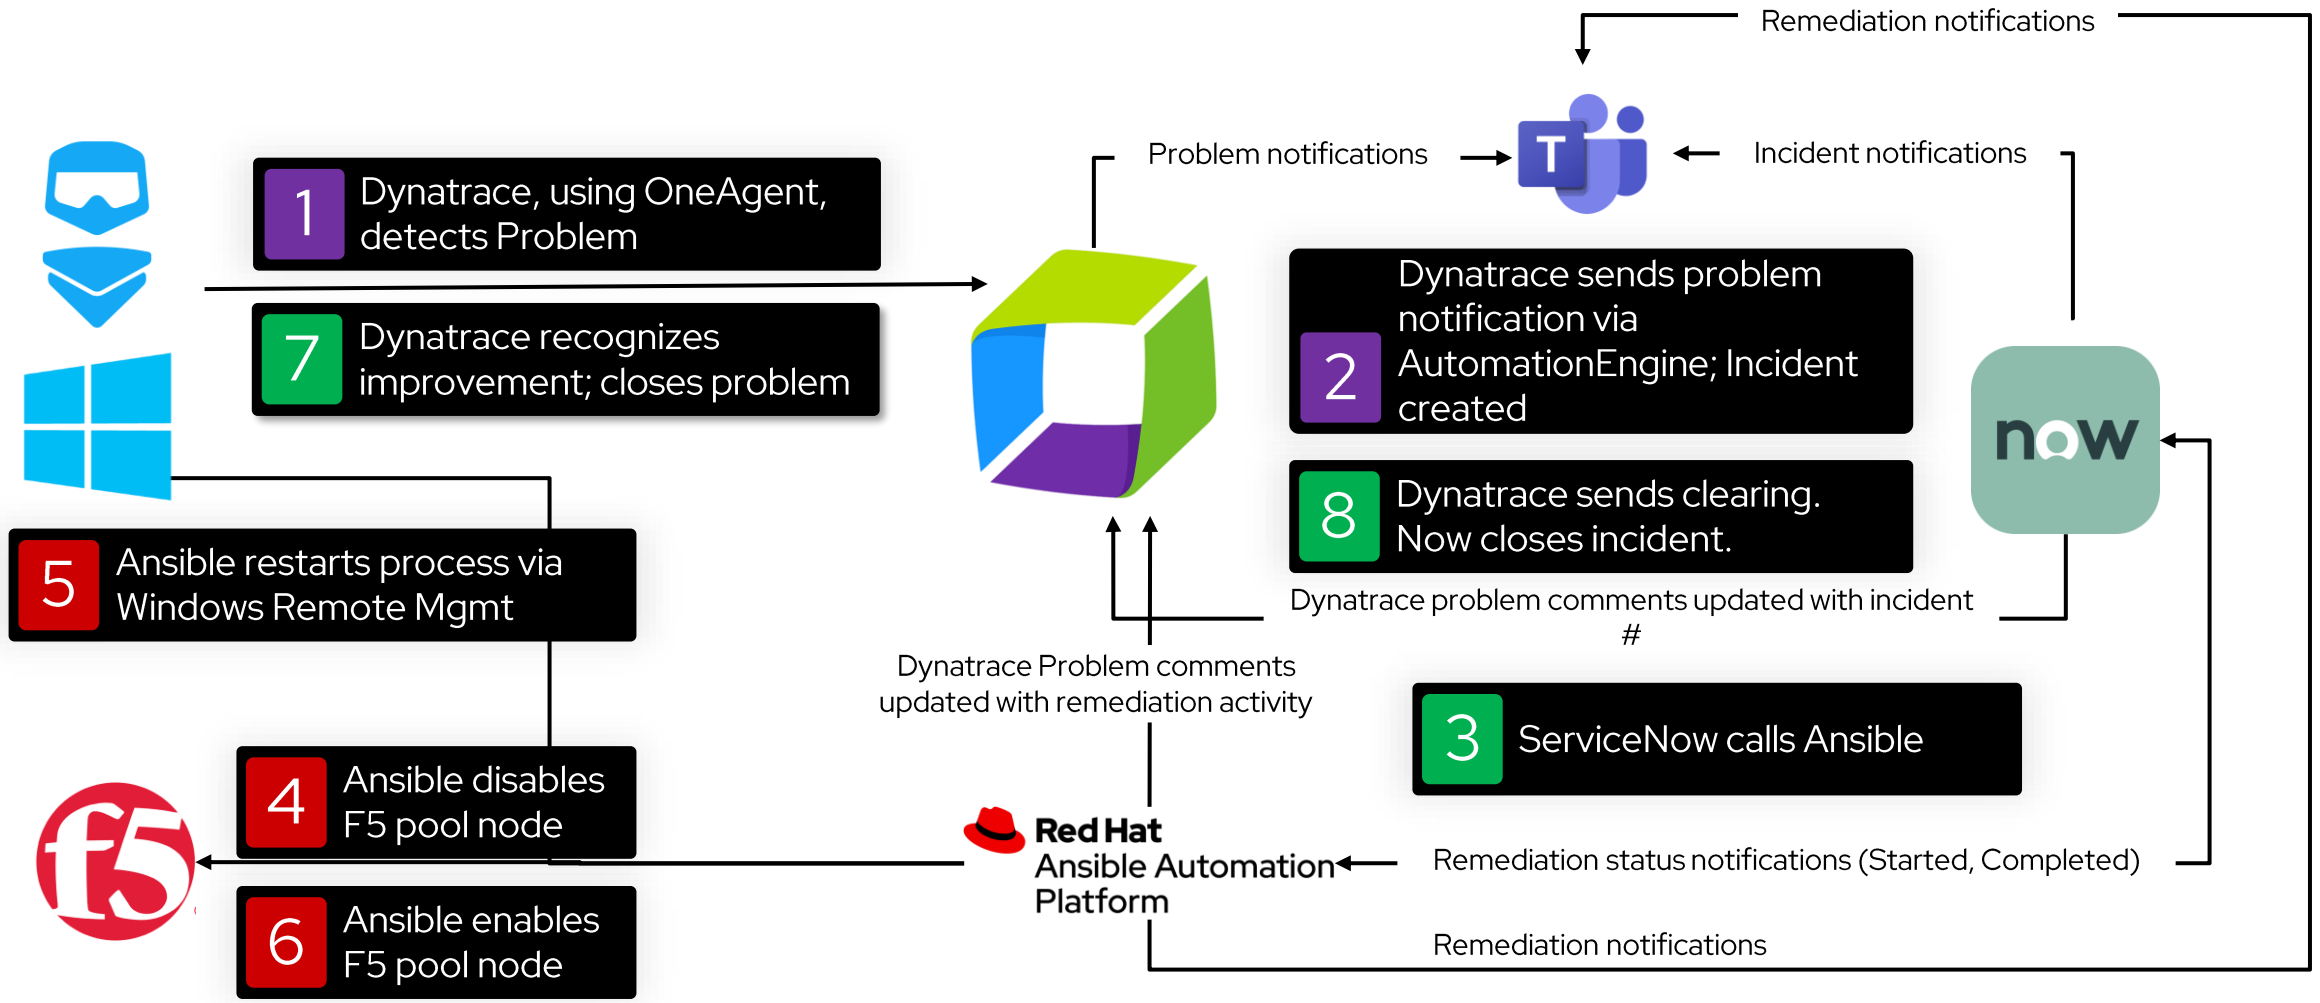

Red Hat

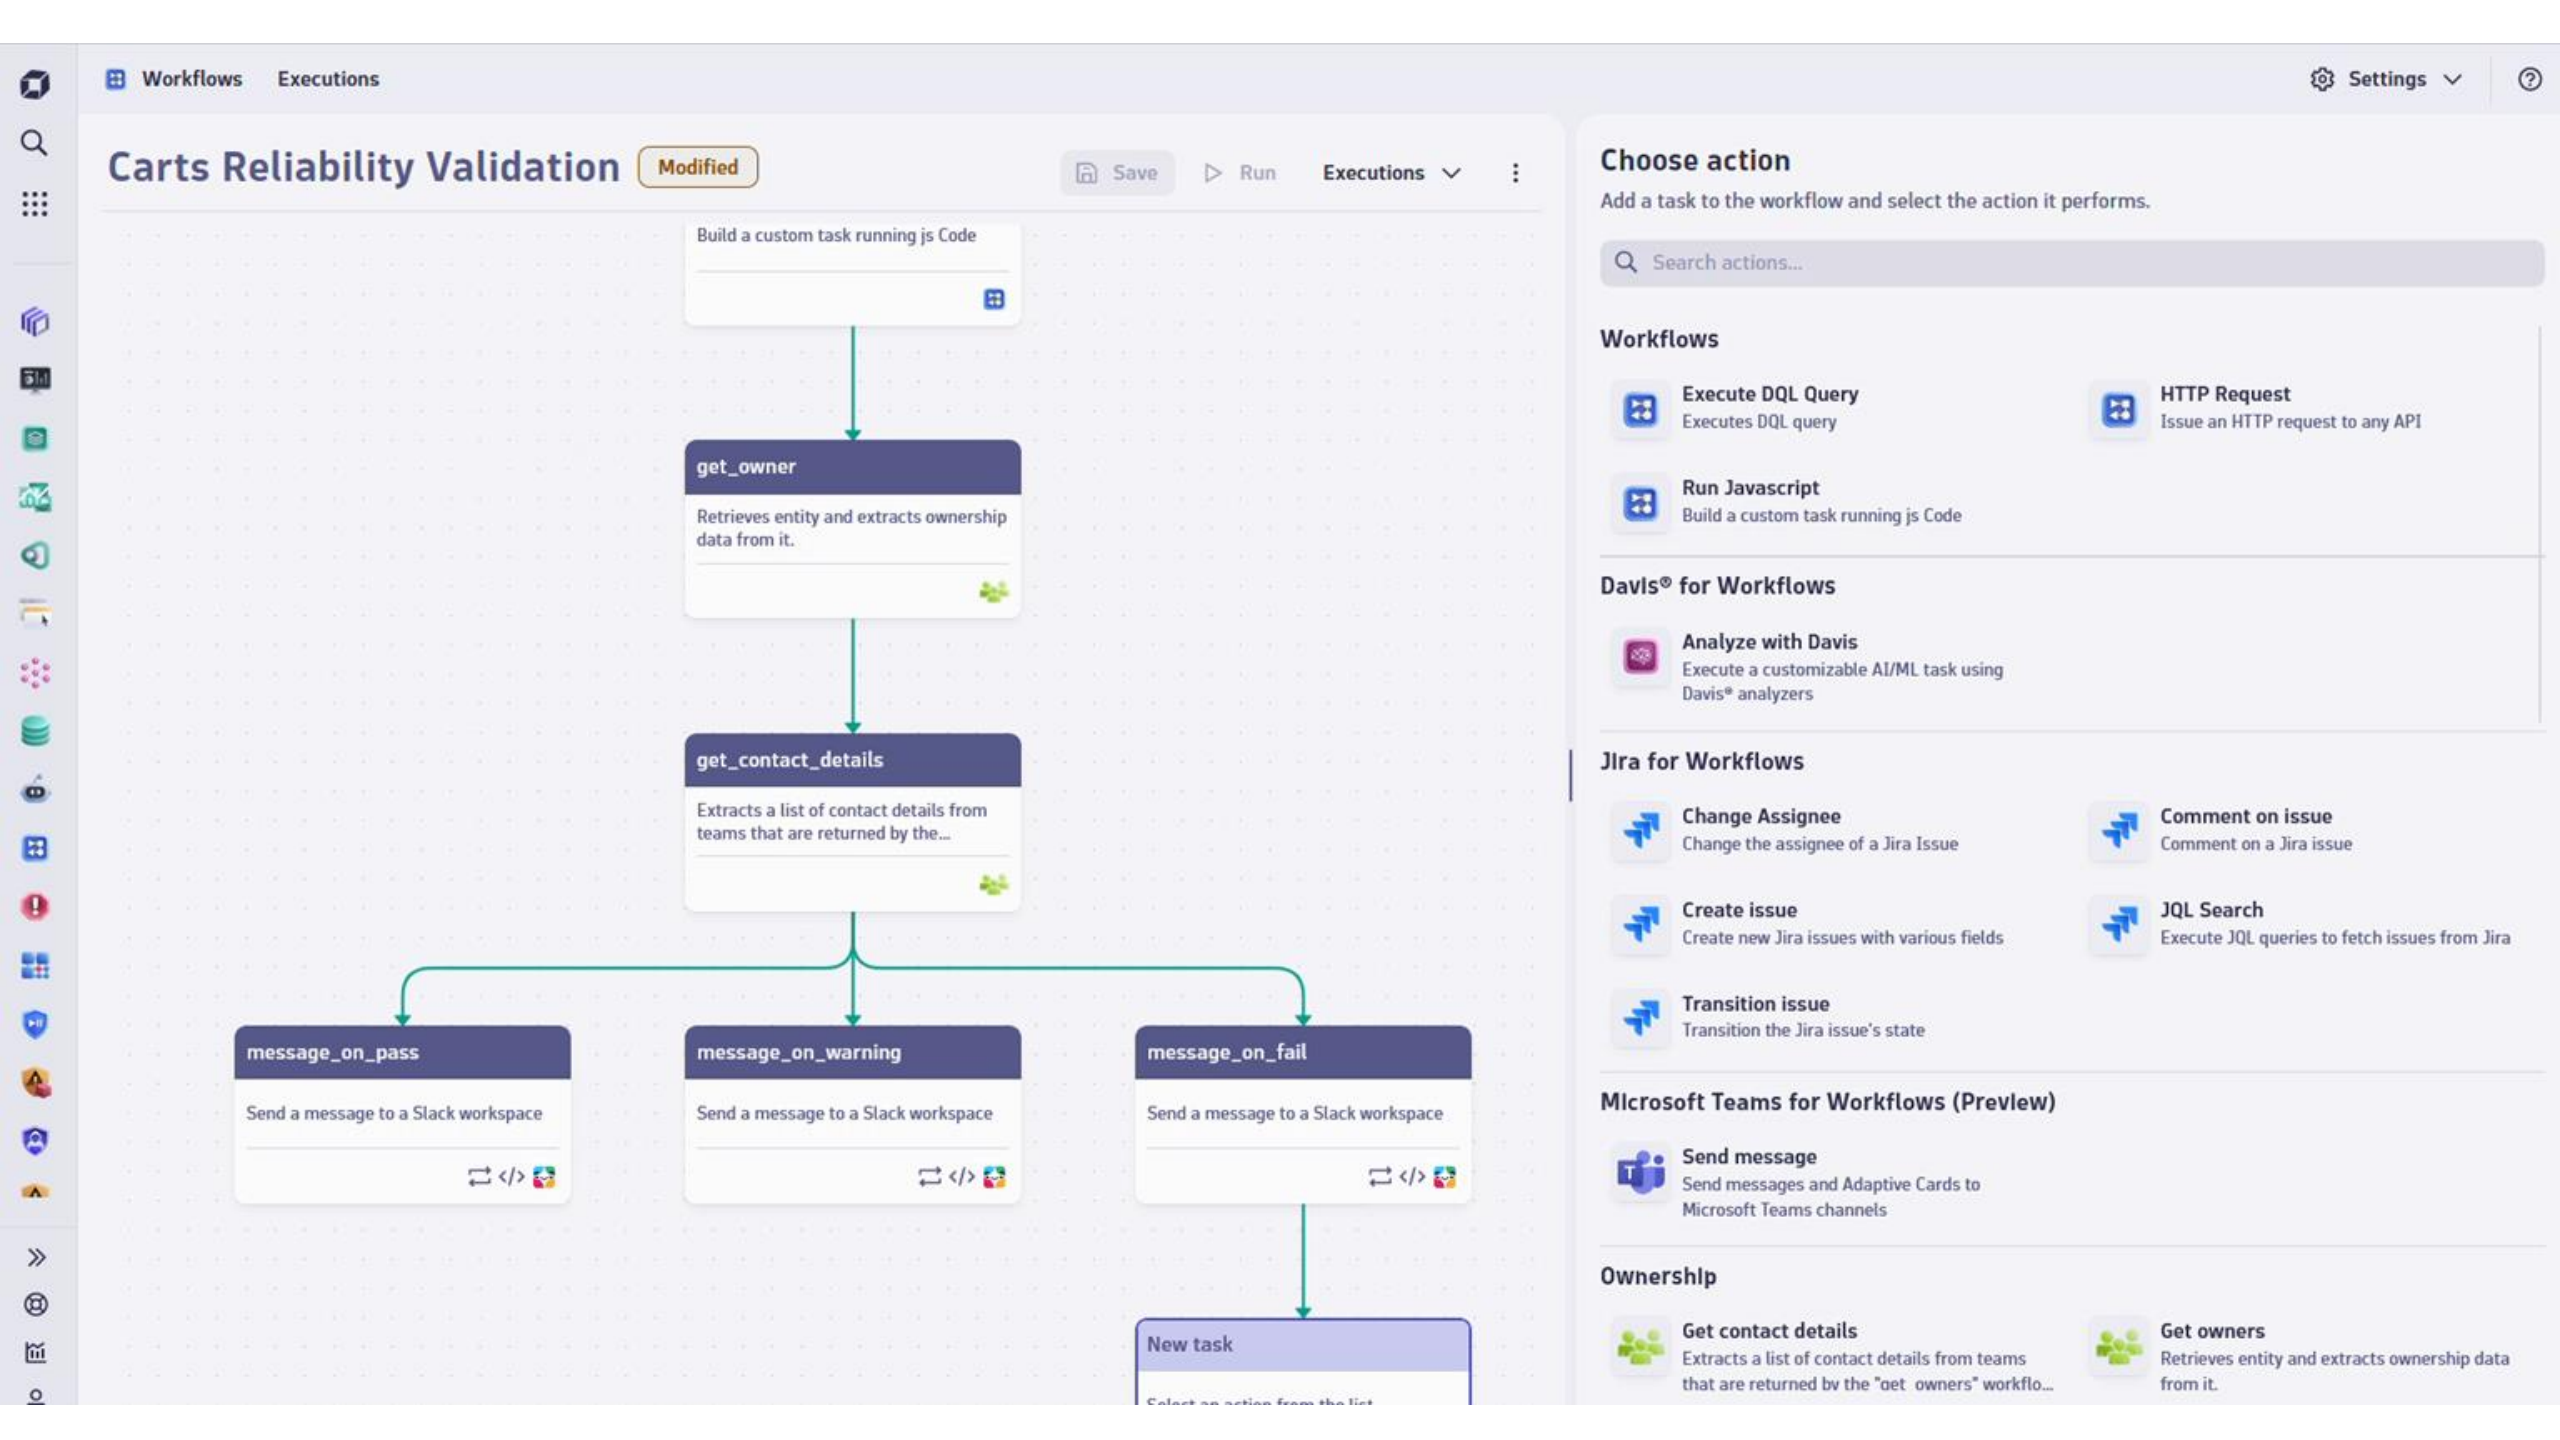
Workflows
Executions
Settings
Carts Reliability Validation
Modified
Save
Run
Executions
Build a custom task running js Code
get_owner
Retrieves entity and extracts ownership data from it.
get_contact_details
Extracts a list of contact details from teams that are returned by the...
message_on_pass
Send a message to a Slack workspace
message_on_warning
Send a message to a Slack workspace
message_on_fail
Send a message to a Slack workspace
New task
Choose action
Add a task to the workflow and select the action it performs.
Search actions...
Workflows
Execute DQL Query
Executes DQL query
HTTP Request
Issue an HTTP request to any API
Run Javascript
Build a custom task running js Code
Davis® for Workflows
Analyze with Davis
Execute a customizable AI/ML task using Davis® analyzers
Jira for Workflows
Change Assignee
Change the assignee of a Jira Issue
Comment on issue
Comment on a Jira issue
Create issue
Create new Jira issues with various fields
JQL Search
Execute JQL queries to fetch issues from Jira
Transition issue
Transition the Jira issue's state
Microsoft Teams for Workflows (Preview)
Send message
Send messages and Adaptive Cards to Microsoft Teams channels
Ownership
Get contact details
Extracts a list of contact details from teams that are returned by the "get_owners" workflo...
Get owners
Retrieves entity and extracts ownership data from it.

```yaml
- name: Listen for events on a webhook
  hosts: all
  sources:
    - ansible.eda.webhook:
        host: 0.0.0.0
        port: 5000

  rules:
    - name: Problem payload Dynatrace for CPU issue
      condition: event.payload.problemTitle contains "CPU saturation"
      action:
        run_job_template:
          name: "Remediate CPU saturation issue"
          organization: "Default"
    - name: Problem payload Dynatrace for App Failure rate increase issue
      condition: event.payload.problemTitle contains "Failure rate increase"
      action:
        run_job_template:
          name: "Remediate Application issue"
          organization: "Default"
```

```
hosts: all
  ## Define Dynatrace source for events
  sources:
    - dynatrace.eda.dt_esa:
        dt_api_host: "https://xxxx.live.dynatrace.com"
        dt_api_token: "xxxxx"
        dt_entity_tags: "entityTags(\"EDA Priority:High\",\"key1:value1\")"

  ## Define the conditions we are looking for
  rules:
    - name: Problem payload Dynatrace for App Healing
      condition: event.title == "Failure rate increase" and event.rootCauseEntity is
defined
      ## Define the action we should take should the condition be met
      actions:
        - run_playbook:
            name: playbooks/remediate-dynatrace-securitychange.yml
        - run_playbook:
            name: playbooks/dynatrace-update-problem-comments.yml
```

5 steps to start your auto-remediation success WITH ANSIBLE

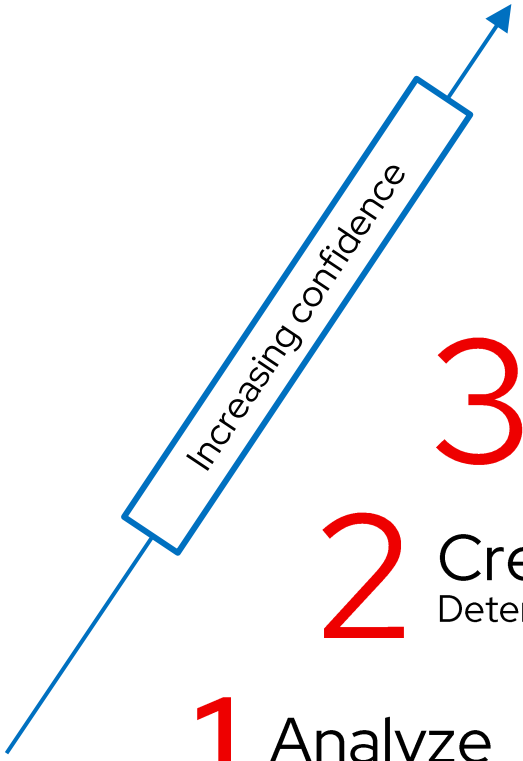

5 Automate
Fully automate the remediation actions with a proven track record end-to-end.

4 Integrate
Create approval-based triggering of remediation with Change Management.

3 Trigger
Trigger the playbooks manually when a Problem is detected by Dynatrace, learn and repeat.

2 Create
Determine current rulebooks and playbooks on Red Hat Ansible that can be manually triggered.

1 Analyze
Evaluate your problems for repetitive patterns to identify automation opportunities.

Red Hat Summit

Connect

dynatrace

Thank you

Red Hat