Red Hat Summit

Connect

Red Hat Observability

Challenges of modern observability

Radek Vokál <rvokal@redhat.com>
Lead Observability Product Management

Red Hat

With great help from Observability PMs

Vanessa Martini (Observability Analytics, UI)

Roger Floren (In cluster monitoring, Cluster Observability, Power monitoring)

Red Hat

Opening **Statements**

- Red Hat is not a **observability** company.
- Red Hat does not have stand-alone **observability** products…however….
  - OpenShift, OpenStack and RHEL offer **observability components and capabilities**.
  - Insights, OpenShift Virtualization, RHOAI contain features to **enhance observability**
  - Ansible can **automate** observability functionality.

This session will discuss

- **Challenges** in the Red Hat Observability stack is helping to overcome
- Overview of **interesting features** Red Hat Observability provides

redhat.com/observability

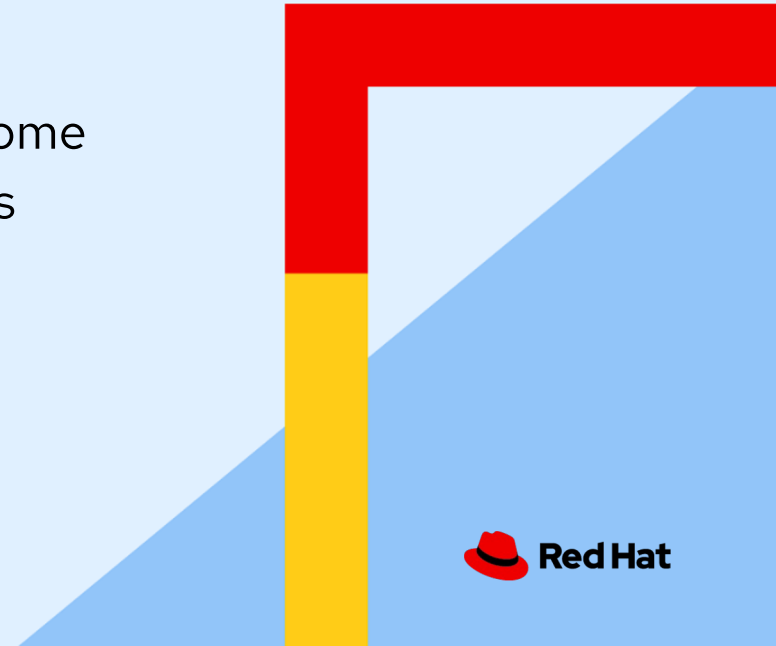

Observability Value:

Three pathways to platform & application excellence

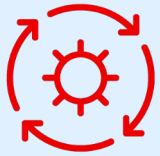

Monitor

- Platform, Services & Applications
- On-prem and managed OpenShift
- Complex environments, Edge, OpenShift AI, OpenShift Virtualization

Get Answers

- Analytics & Red Hat domain knowledge
- AI and Proactive analytics
- Data driven decisions

Maximize Effectiveness

- Improve experience with OpenShift platform by data driven features
- Efficient and secure platform and applications
- Value add features in Red Hat products (Support efficiency, Updates, Insights services)

redhat.com/observability

Red Hat

What if there's no Observability?
The rise of solar energy

1.3 million solar panel installations across the entire UK.

Solar panel installation rates tripled in last year

Growing need to connect/disconnect new power sources

Data-driven decisions need to be made within minutes

Wrong decisions can have **drastic implications**

High demand for a **reliable and scalable** observability solution

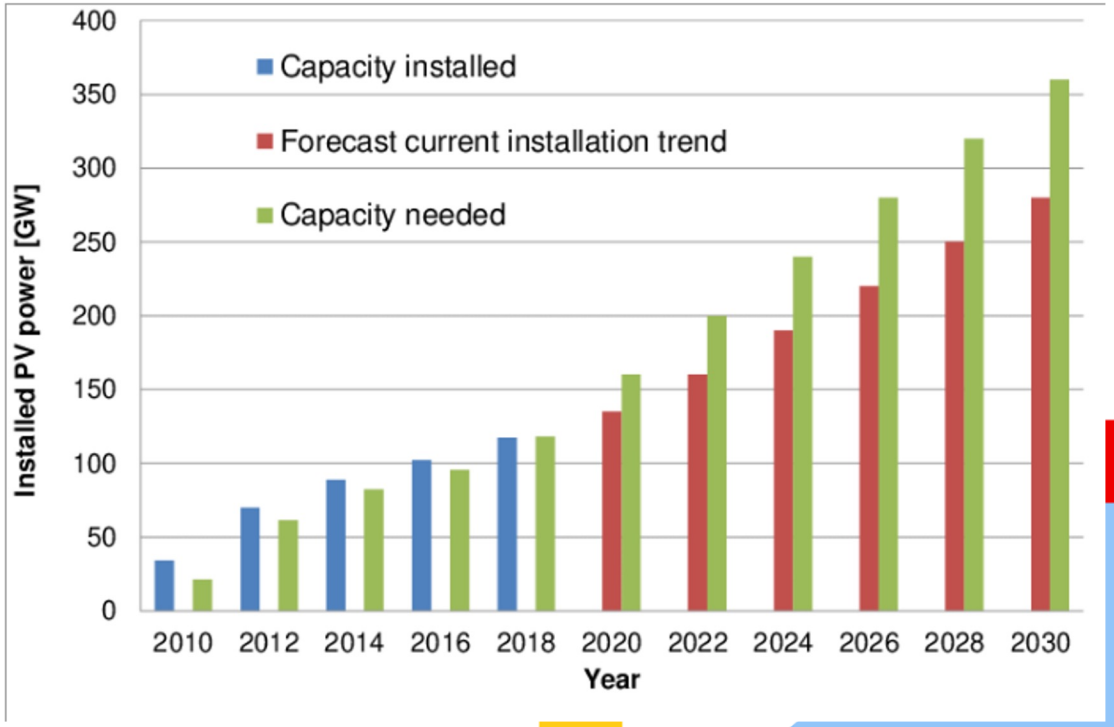

Red Hat

Challenge #1: (ever)Growing complexity

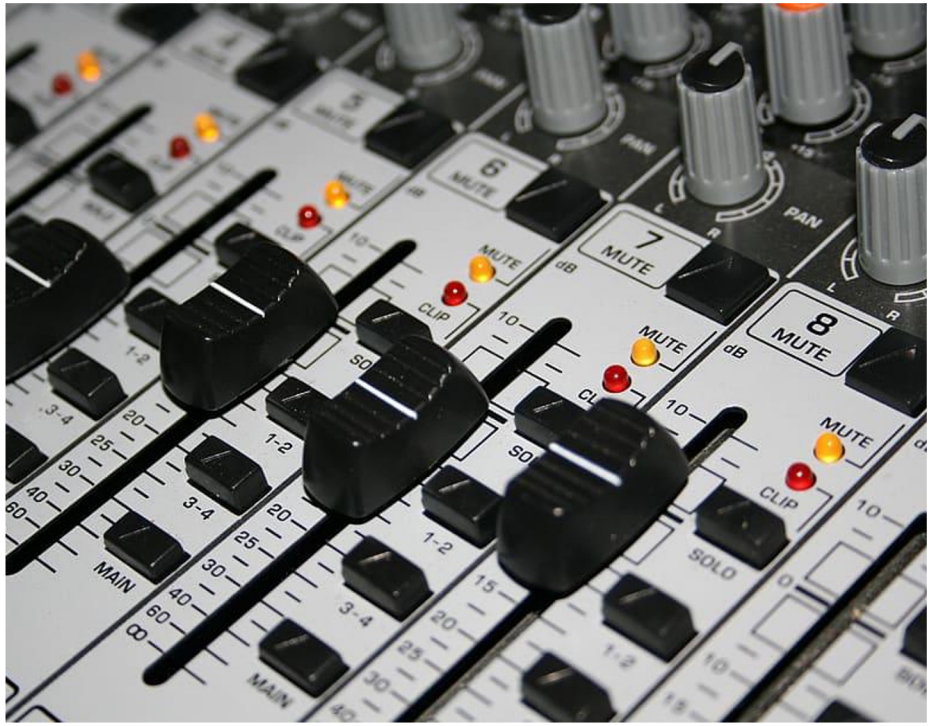

Kubernetes has become the standard application platform; complexity remains #1 challenge

Growing maturity of customers and applications, rise of microservices, new workloads (AI) and use cases

Lack of centralized visibility

Infrastructure complexity, hybrid cloud, bare-metal vs. public cloud, GPU/CPU, Intel/ARM, virtualization

Resource utilization and capacity planning

Red Hat

Challenge #2: Signal Noise & Troubleshooting Complexity

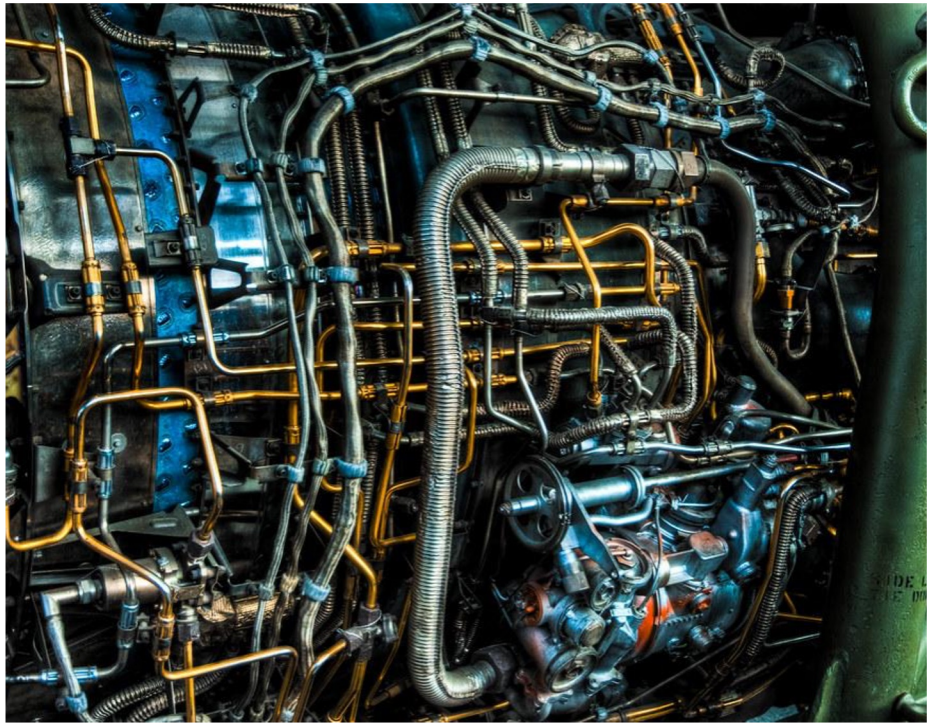

Platform built on signals, encouraging to leverage signals, producing too many signals

Growing data volume, low signal-to-noise ratio

Components are independent, hard to find correlations

What is the most important signal? Where do I start with triaging issues? What should I prioritize?

Different signal types, different data schemas, hard unification and data correlation

Red Hat

Challenge #3: Resource Control & Cost

Cloud scales, developers allocate resources, bills go up!

High overhead cost

Complex capacity planning, on-prem vs. public

Underutilized/over-provisioned resources

Insufficient data granularity and data transparency

Red Hat

Challenge #4: Scalability & Governance

Increasing volume of data generated and growing demand from users to leverage data

Different teams/stakeholders accessing different data - Admins, operatorions, finance, developers, ...

Data overload - leading to slowdowns, errors and oversight

Security and compliance risk - lack of data privacy/governance

Slower innovation, hard to adopt new technologies with a lack of scalable observability

Red Hat

The Red Hat difference:

Core elements of integrated observability excellence

We are

- A distribution of open source observability components
- Components integrated, secured and supported by Red Hat
- Components for collecting, storing, processing, analyzing and visualizing observability signals
- Capabilities to address operational needs of our platforms, improve efficiency and troubleshooting
- Guidelines for observability in specialized use cases (Edge, AI, Telco, ..)
- Partner in the ISV ecosystem for integrations

We are **NOT**

- A monitoring software company
- A logging solution company
- An ITSM software company
- Network solution company
- Dynatrace (or Datadog or name other 3rd party SaaS monitoring tool)

Red Hat

Red Hat Observability: Open Source to Enterprise Value

A Framework for Modern Observability

Data Collection

Metrics

System health and performance

Logs

System generated events for debugging and analysis

Traces

Sequence for events and transactions

Other

NetObserv Insights

Data Storage

Metrics

Multi-tenant, highly scalable

Logs

Distributed, scalable, HA log aggregator

Traces

Distributed, scalable, HA, multi-tenant Distributed tracing backend

Data Delivery

- Data standardization
- Data forwarding
- Data migration & transportation

observability operator

Metrics

Logs

Traces

Power monitoring

Data Visualization

- Single-/multi- cluster experience
- On-prem / SaaS

OPENSHIFT Cluster Observability

Red Hat Advanced Cluster Management for Kubernetes

Cluster Observability Operator

MultiCluster Observability Addon

Data Analytics

- Issue navigation
- Event correlations
- Predictive analytics
- Remote Health/OCP features
- Support efficiency

OPENSHIFT

korrel8r

Update risk

Red Hat Insights

Insights Advisor

Cost Management

Rightsizing

ProductSense

Foundation | **Integration** | **Value Add**

Data **Collection & Storage**

Prometheus

Monitoring & **alerting** for distributed environments

Strengths

Time-series data collection & **flexible** querying via PromQL

OpenTelemetry

APIs & **SDKs** for instrumenting, generating & exporting telemetry data

Strengths

Simplification of observability data collection & great **compatibility** with analysis/visualization tools

Grafana loki

Horizontally scalable **log aggregation** system

Strengths

Simplification of operations & reduction of storage **costs**

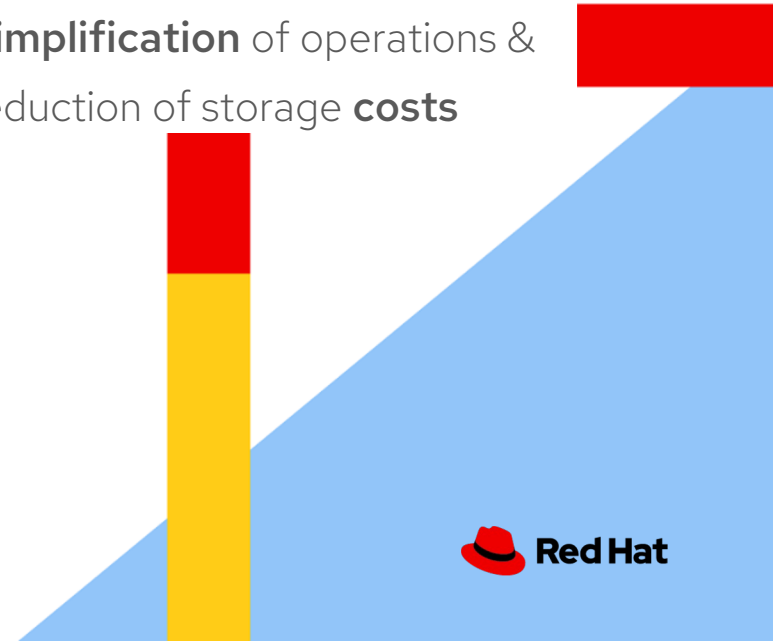

Red Hat

Data **Visualization: Perses**

- **Standard** Observability data visualization tool - **CNCF Sandbox**
- Enables ***dashboards-as-code***
- **GitOps** friendly
- Embeddable with **NPM packages**
- Supports **plugins**

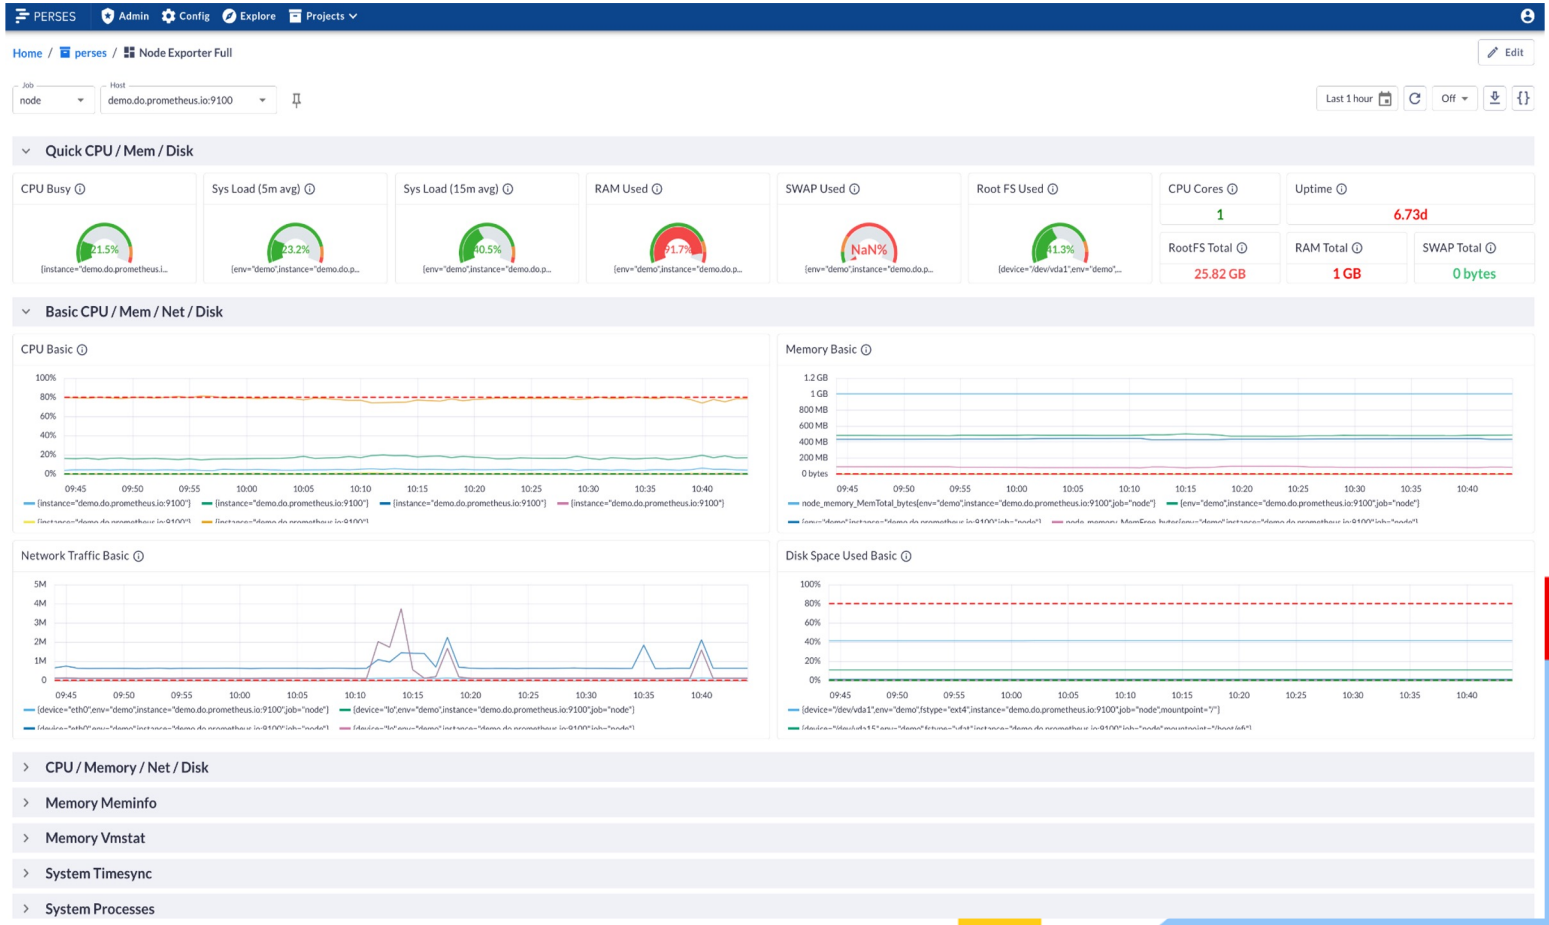

Red Hat

Data **Analytics: Troubleshooting journey**

Korrel8r

- **Correlation** of observability signals
- **Rules** define relationships between signals
- Reduces **troubleshooting time**

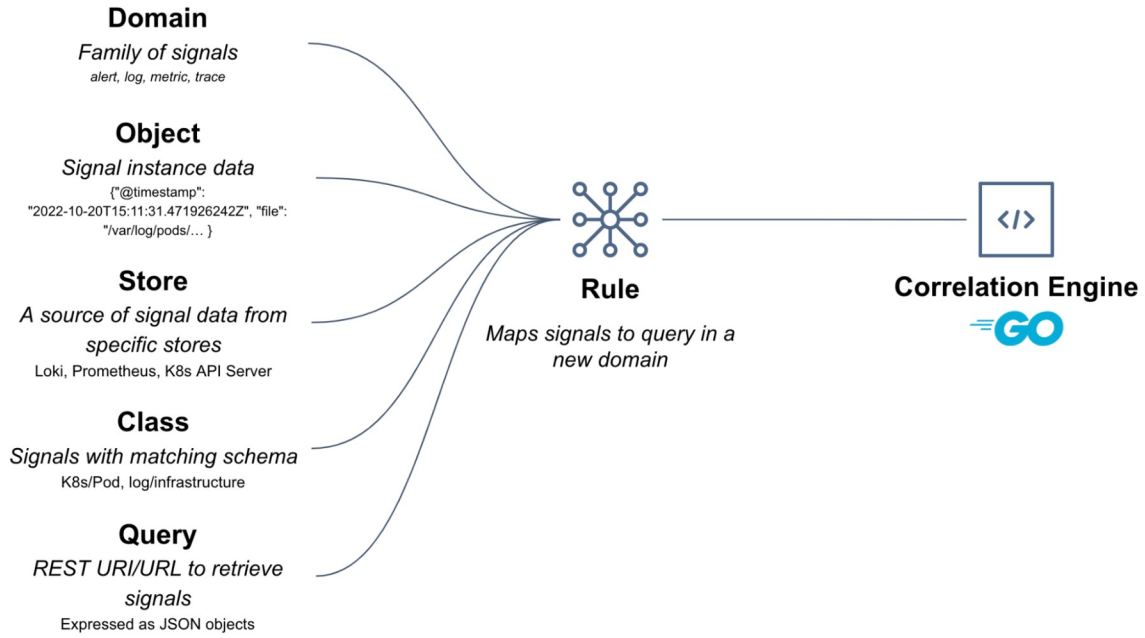

Incident Navigation

- **Grouping** of events and signals
- **Time and context** based relationships
- Ability to **deep dive and highlight potential root cause**

Red Hat

Single- and Multi-Cluster Observability

Cluster Observability Operator

- **Single pane of glass** for different signal types
- Integration with OpenShift WebConsole **UI and analytics components**
- Default (opinionated) **configurations**
- **Profiles**

Multi-Cluster Observability Addon

- Multi-cluster setup for **different signal types**
- Tight integration with **Advanced Cluster Management**
- Selecting only the most relevant signals
- Adapting to different infrastructure
- **UI and analytics integration**, fleet-level views

Red Hat

Power **Monitoring: Kepler**

- Kubernetes-based Efficient Power Level Exporter
- eBPF-based **CNCF Sandbox** project
- Probes **performance** counters
- Uses ML models to estimate **workload energy consumption**
- Exports stats as **Prometheus** metrics

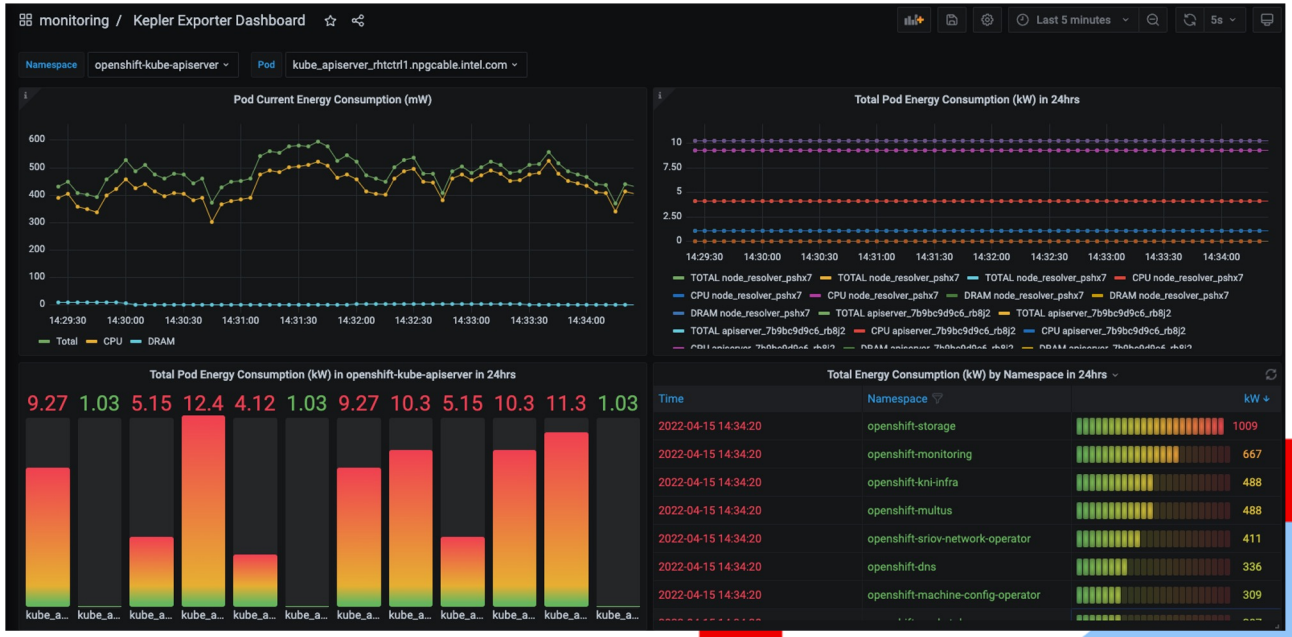

Red Hat

GenAI: **OpenShift Lightspeed**

OpenShift Lightspeed

- **Generative AI** based virtual assistant
- Chat interface built into the OpenShift UI console
- **Leverages** the latest, robust OpenShift **documentation**
- Question relevance validation/rejection
- Pluggable **LLM providers/models**

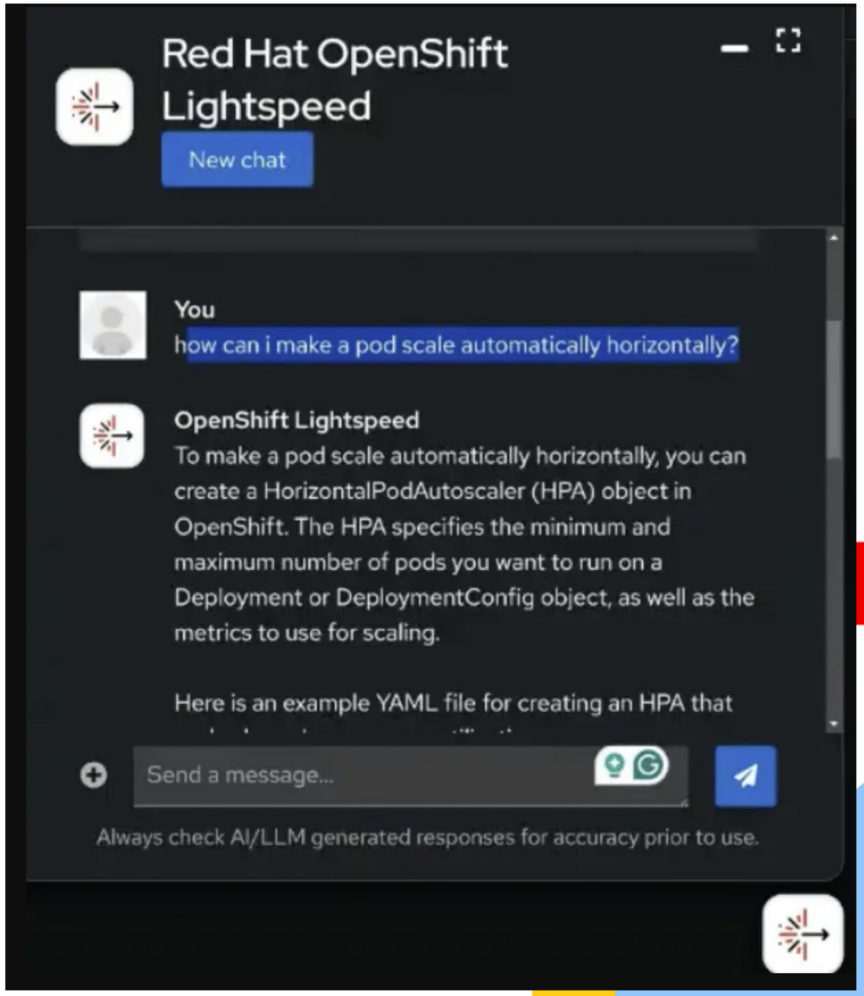

Red Hat

2025 focus

- ***Edge* clusters**: real-time processing & low-latency
- Defining the role of **traces** - rising technology
- AI-based insights - **predictive AI** on the rise!
- Need for industry-accepted standards - data **standardization**
- Defining observability requirements through **infrastructure-as-code**
- More focus on **cost management** & **sustainability** tooling
- Focus on **automated remediation and AIops**

Red Hat

Red Hat Summit

Connect

Thank you

linkedin.com/company/red-hat

facebook.com/redhatinc

youtube.com/user/RedHatVideos

twitter.com/RedHat

Red Hat


Red Hat

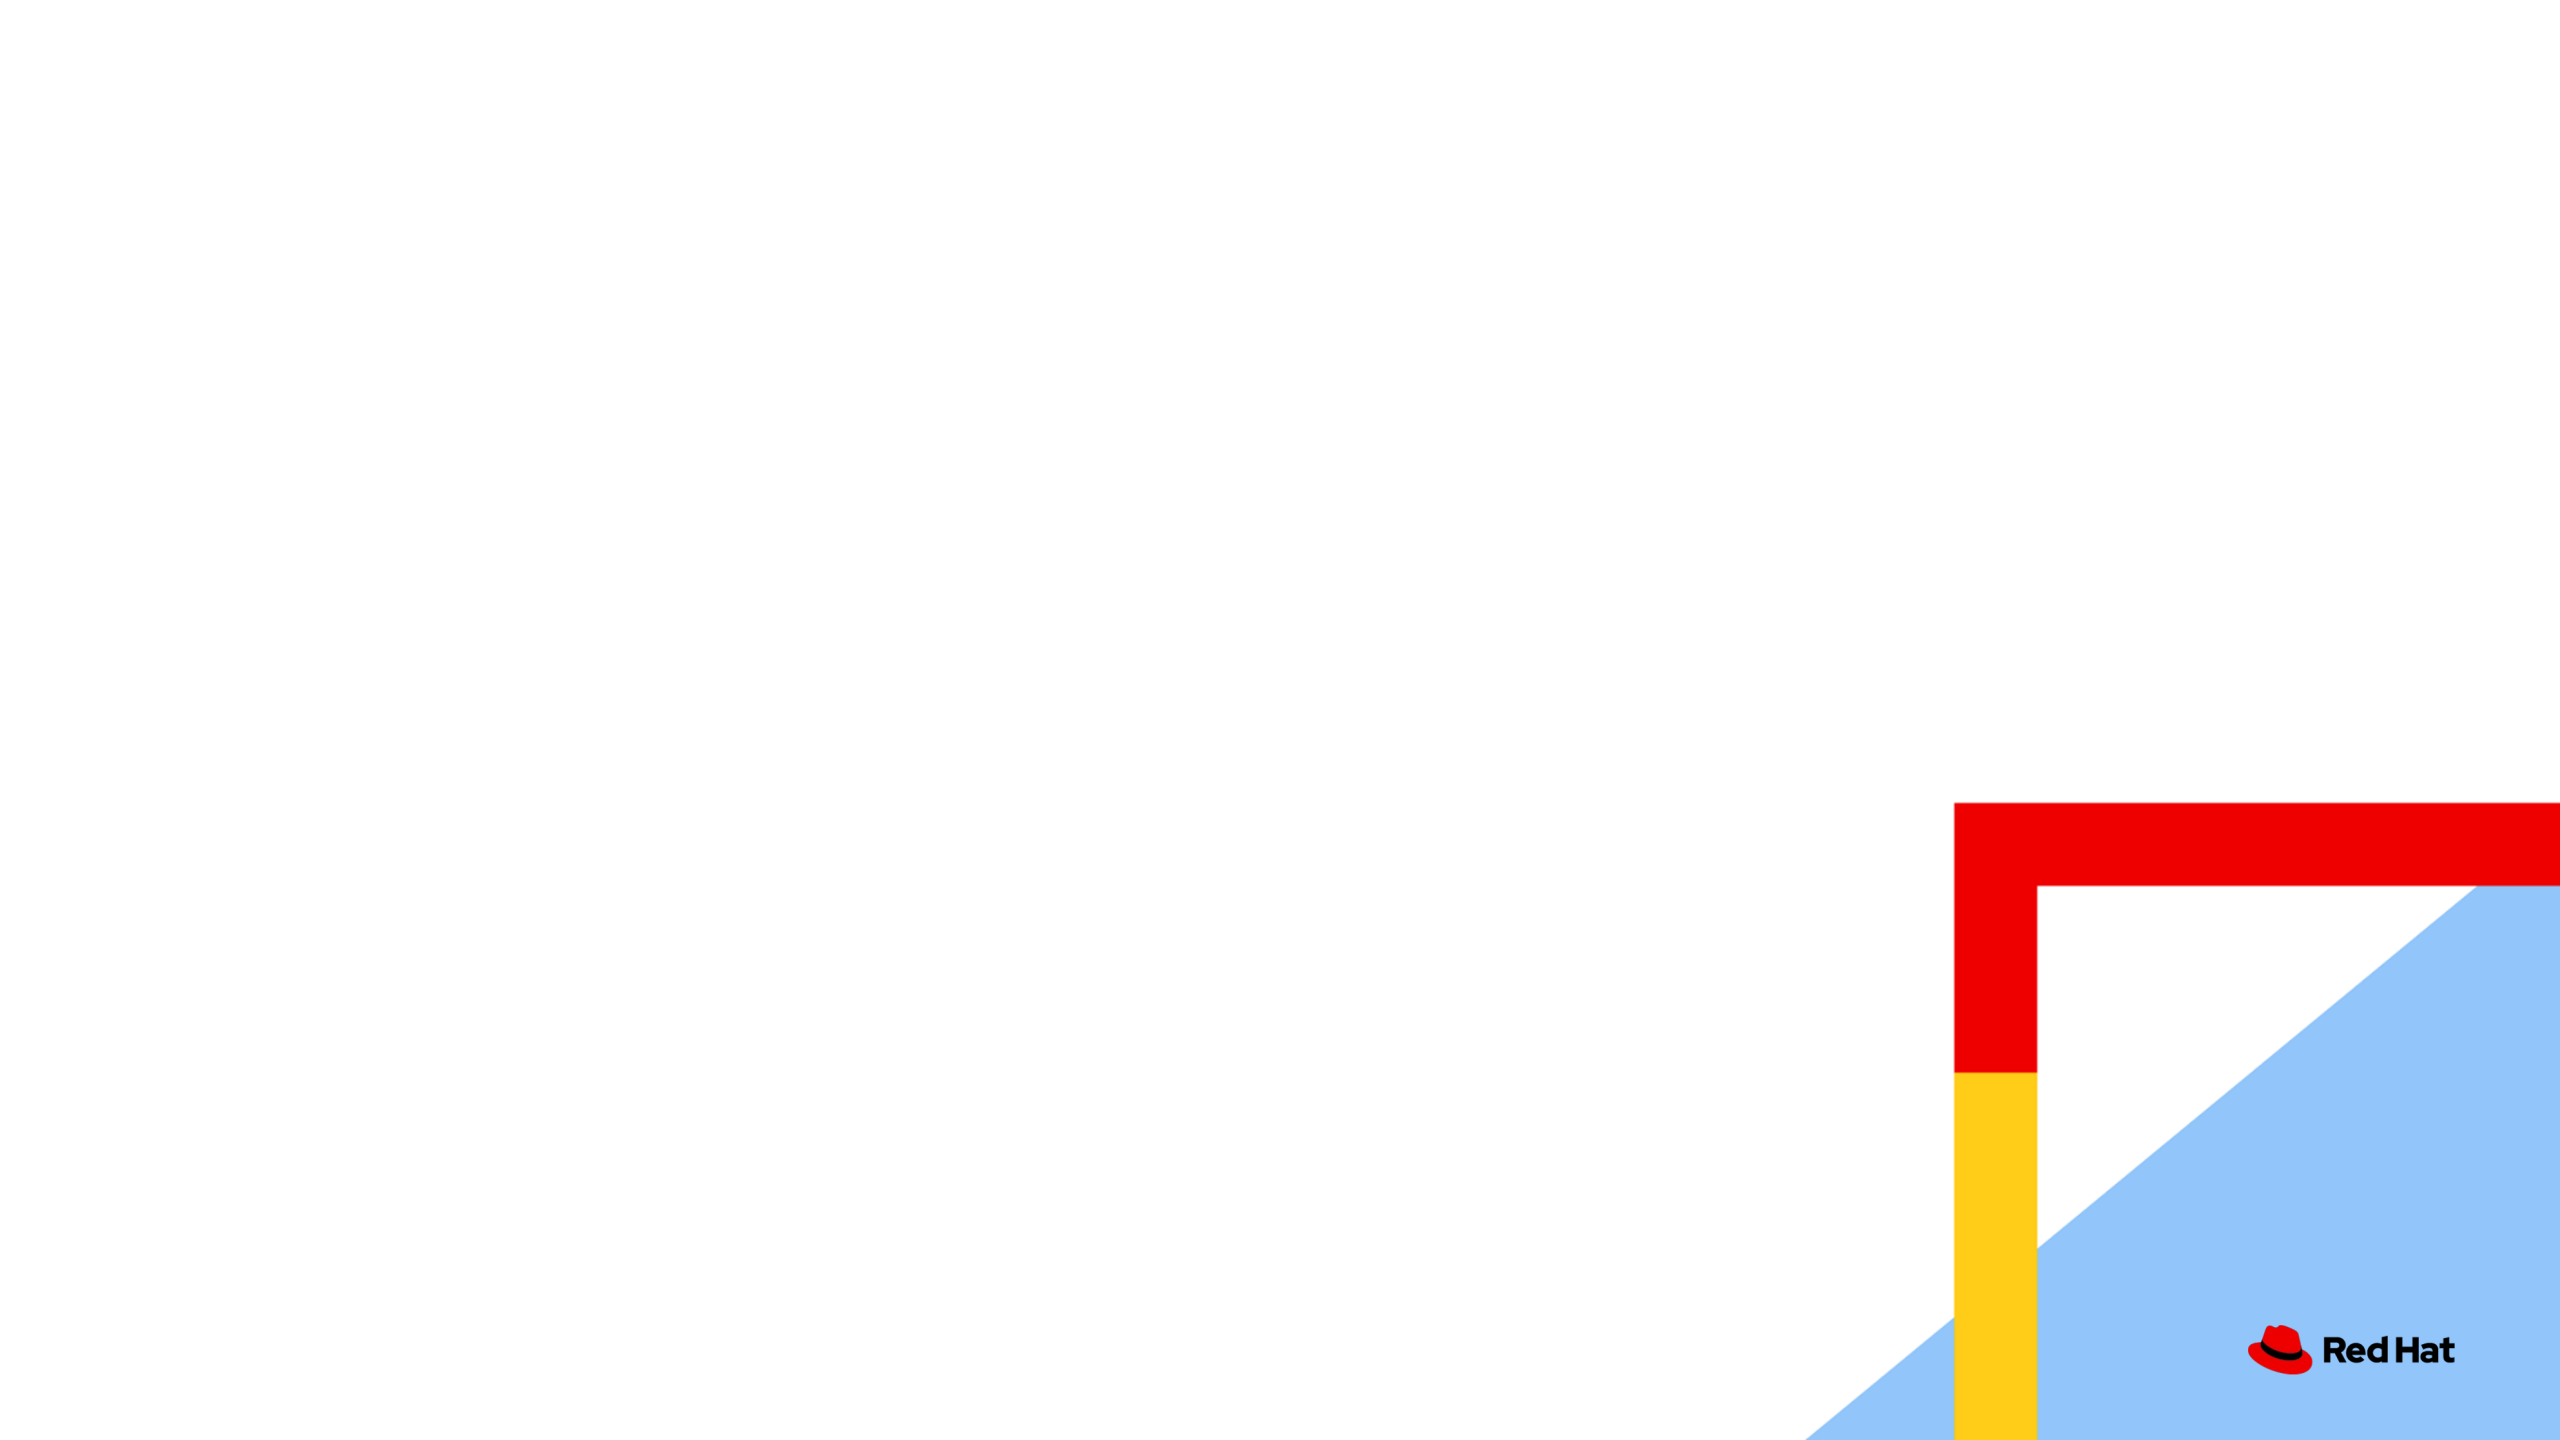
Red Hat