

Red Hat
Summit

Announcement Overview

Red Hat Summit 2024

Product & technology innovation news



Red Hat
Enterprise Linux AI
Developer Preview



podman
AI Lab project now available



KONVEYOR
GenAI project availability soon

New features **generally available**
for
 **Red Hat**
OpenShift AI

Extending **Red Hat Lightspeed** to
 **Red Hat** Enterprise Linux  **Red Hat** OpenShift
Technology preview in 2H CY24

Policy as Code for
 **Red Hat**
Ansible Automation Platform
Technology preview soon



Red Hat
Trusted Software Supply Chain



InstructLab

Image mode for
 **Red Hat**
Enterprise Linux
Technology preview



Red Hat
Connectivity Link
Developer preview

Ecosystem expansion news

Acceleration for OEMs and GPU providers:



Integrated ISV applications:



Software-defined vehicles:



Industry alliances:



Customer news



Adopted Red Hat OpenShift AI to expand AI usage across Uruguayan government agencies



Drove new innovations in maternal-fetal image using Red Hat OpenShift for AI



Used Azure Red Hat OpenShift to get to market faster and more collaboratively



Selected Red Hat Enterprise Linux as the operating system platform for the El Capitan exascale supercomputer



Worked with Red Hat to adopt InnerSource practices for greater collaboration and faster innovation.



Enhanced network uptime, consistency and efficiency with Red Hat Ansible Automation Platform

The cloud is hybrid. So is AI.



AI is not defined by a single footprint



AI strategies start with a hybrid cloud strategy



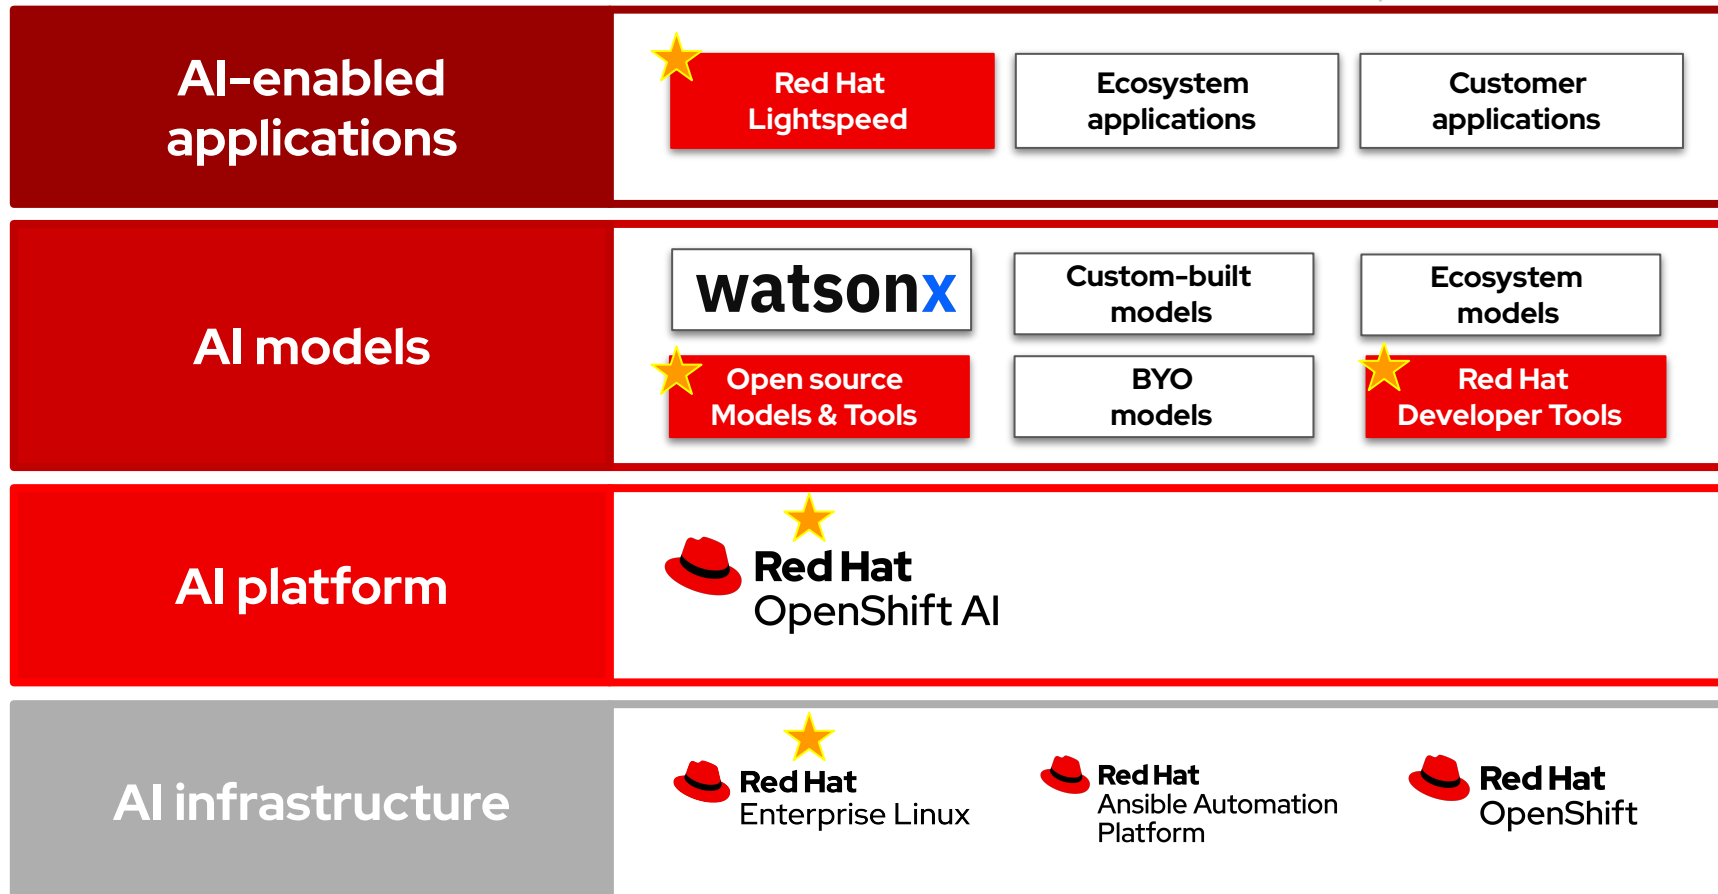
Flexibility and choice are paramount to AI success in the enterprise



Existing investments must work in support of, not against, AI workloads and models

Red Hat delivers the *consistent, trusted and comprehensive* AI platforms for the hybrid cloud.

Red Hat is accelerating AI across the open hybrid cloud



Red Hat's AI certified partner ecosystem



Integrated ISVs

Anaconda
Elastic
IBM
Intel
NVIDIA
Pachyderm
Run:ai
Stability AI
Starburst



AI & general ISVs

C3.ai
Cloudera
Couchbase
Crunchy
Enterprise DB
Guise AI
H2O.ai
Hugging Face
Lightbend
Palantir
Redis Labs
SAS




Delivery partners

Accenture
BCG
CapGemini
Deloitte
HCL Technologies
Infosys
Kyndryl
World Wide Technology



Cloud partners

AWS Google
Azure IBM Cloud



Hardware/OEMs

AMD
Dell
Cisco
HPE
Intel
Lenovo
NVIDIA
Qualcomm
Supermicro





Red Hat OpenShift AI

Hybrid MLOps platform

Collaborate within a common platform to bring IT, data science, and app dev teams together

Available as

- **managed cloud service**
- **traditional software product on-site or in the cloud**




Model tuning & alignment

Conduct exploratory experiments in JupyterLab with access to core AI / ML libraries and frameworks including PyTorch and InstructLab using our notebook images or your own.



Model serving & monitoring

Deploy and scale models across any cloud, fully managed, and self-managed OpenShift footprint and centrally monitor their performance.



Lifecycle management

Create model pipelines for tuning and alignment (manual and InstructLab), validation and delivery of models across your enterprise.



Resource optimization and management

Optimize resources across users for use in tuning and inference. Share resources, projects, and models across users

Red Hat's view of AI

Today

"Open" models
Closed evolution

- LLMs flourish in the open, **but are limited without being truly open source**
- Users for communities but **can't contribute**

Tomorrow

Open source
models
Open evolution

- **Bring open source concepts to AI models** and the tools for open evolution
- Enable users to **create and merge contributions to LLMs**, increasing efficiency and performance

Execution

Enterprise products
driven by upstream
communities

- Red Hat and IBM will release **Apache 2.0-licensed models and tools** as an open source project
- A new Red Hat offering, **RHEL AI**, brings these capabilities to the enterprise



Red Hat Enterprise Linux AI

Foundation Model Platform

Seamlessly develop, test and run best of breed, open source Granite generative AI models to power your enterprise applications.

The model is the new platform.



Open Granite models

Highly performant, fully open source, collaboratively developed Granite language and code models from the community, fully supported & indemnified by Red Hat and IBM.




InstructLab model alignment

Scalable, cost-effective solution for enhancing LLM capabilities efficiently for a wide range of applications, making knowledge & skills contributions accessible to a wide range of users



Optimized bootable model runtime instances

Granite models & InstructLab tooling packaged as a bootable RHEL image, including Pytorch/runtime libraries, hardware optimized inference for Nvidia, Intel and AMD that can run anywhere and provides onramp to OpenShift AI for scale and lifecycle & watsonx for agent integration and governance.

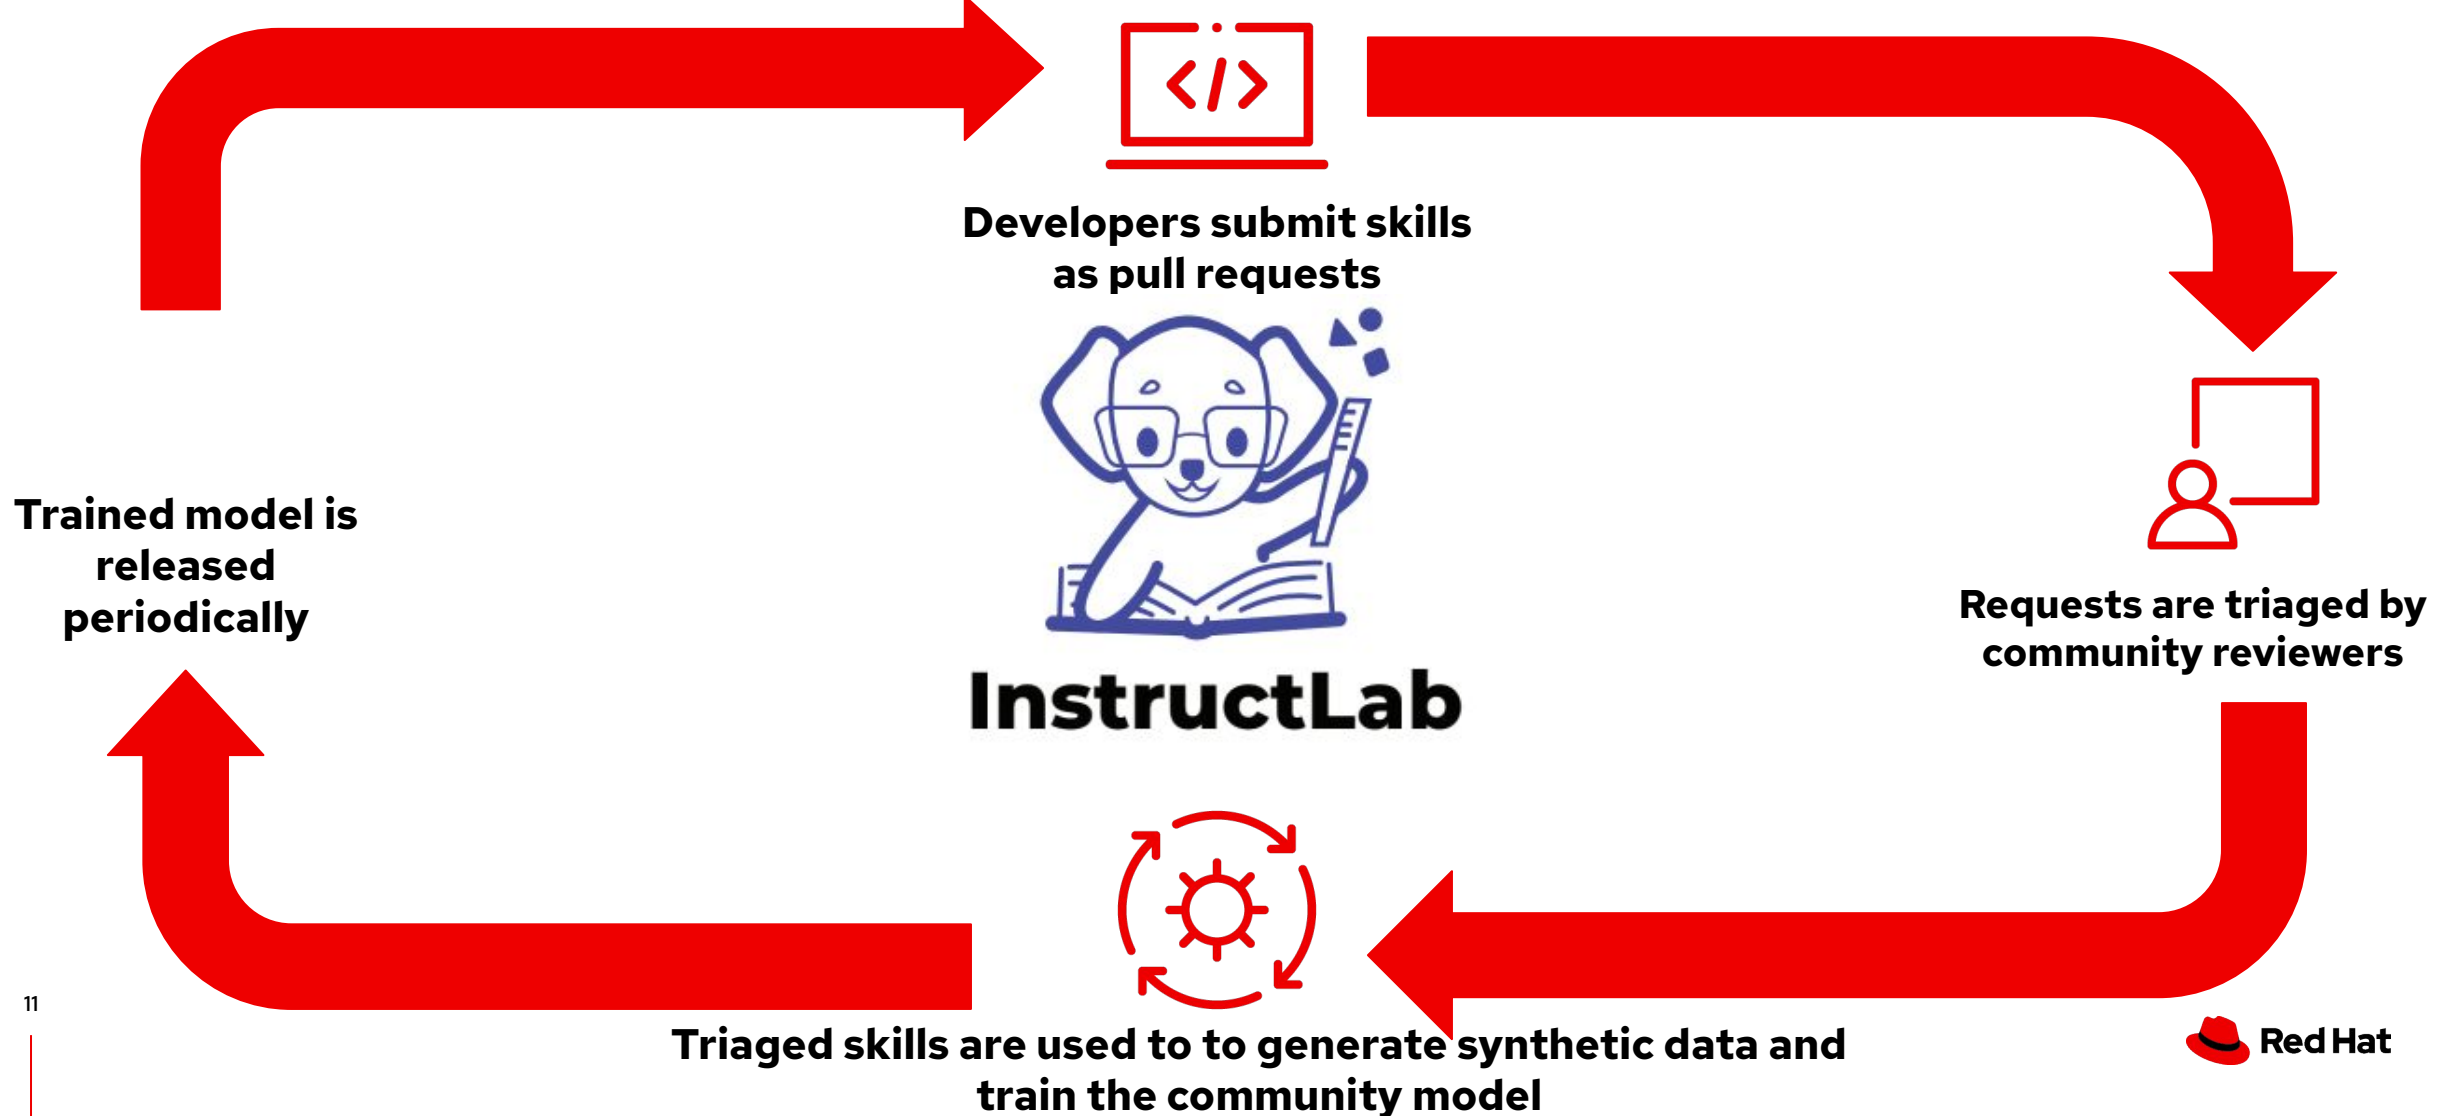


Enterprise support, lifecycle & indemnification

Trusted enterprise platform, 24x7 production support, extended model lifecycle and model IP indemnification

Introducing: **InstructLab**

Open source community project for GenAI model development



Introducing: **image mode** for Red Hat Enterprise Linux

Making operating systems as simple as containers

```
FROM quay.io/rhel9/rhel-9-bootc:latest

RUN dnf -y install [software]
[dependencies] && dnf clean all

ADD [application]
ADD [configuration files]

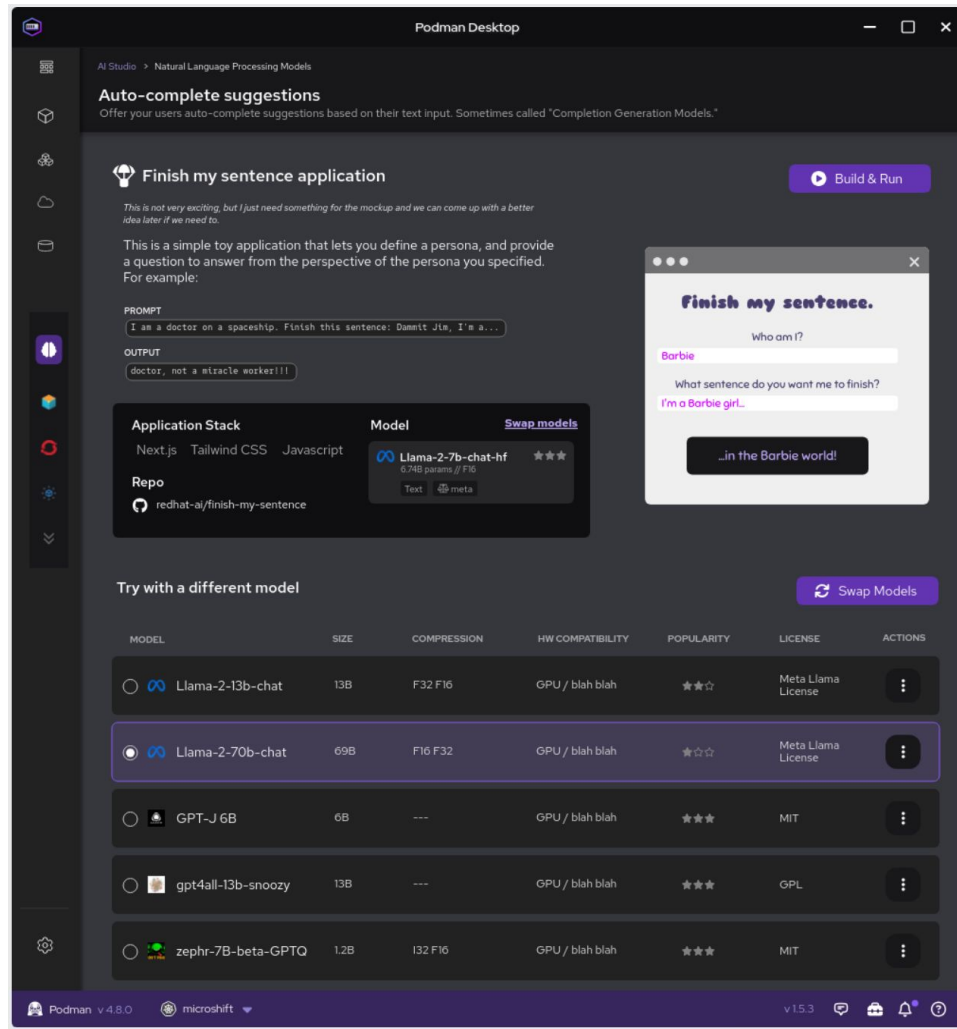
RUN [config scripts]
```

Allows teams, infrastructure, and ecosystems to converge on a single container-native workflow to manage everything from their applications to the underlying operating system. build, test and distribute the operating system as if it was any other container.

- ▶ **All customers** benefit from the simplicity and portability across their traditional and hybrid cloud environments
- ▶ **DevOps teams** finally have a native OS with their favorite CI/CD & GitOps workflows solving last-mile problems
- ▶ **Security teams** will be delighted by the cryptographic validation, transparent provenance, and tool consolidation
- ▶ **Partners** will love how easy it is to build and distribute on a trusted base operating system

Introducing: Podman AI Lab

Simple developer access to local containers and AI



- ▶ Run and code against local models quickly on your laptop (Mac, Windows, & Linux)
- ▶ Accelerate AI adoption, by easing concerns around data access, data privacy, & security
- ▶ Local developer workflow for model fine-tuning
- ▶ Path to production - Easy to package and deploy apps and models direct to OpenShift AI all the way down to bare metal
- ▶ Simple access to Red Hat developer subscriptions

Announcing: **Red Hat Lightspeed** across Red Hat platforms

Intelligent, natural language GenAI processing capabilities designed to extend existing IT skills



OpenShift Lightspeed will be available in Technology Preview later in 2024

RHEL Lightspeed is currently in the planning stages and availability will be announced



Improve the **productivity and efficiency** of ops and developers by integrating AI into cluster administration and the operating system



Simplify enterprise planning and administration, improve performance and enhance security



More easily navigate the complexities of enterprise IT in the hybrid cloud


Announcing: **Konveyor GenAI for the Konveyor Community**



- ▶ GenAI applied to application modernization efforts
- ▶ Workflow-integrated LLMs
- ▶ Generated code directly within IDEs
- ▶ Successful migrations build strong recommendations
- ▶ Roadmap for Red Hat migration toolkit for applications


Introducing: **Policy as Code for Red Hat Ansible Automation Platform**

Solution for Governance, Risk and Compliance (GRC) in complex, AI-infused IT environments



Check policy enforcement during the development cycle

Accelerate policy checking during automation content development



Integrated mechanism for policy enforcement based on unique needs

Discretionary or mandatory checks before or during automation runs



Reporting to aid in compliance status and audit tracking

Policy compliance and auditing reporting more time-efficient

Introducing: **Red Hat Connectivity Link**

A streamlined and composable approach to application connectivity



Simplify connecting disparate apps and services

Integrated, unified control plane for managing modern workloads across the hybrid cloud



Seamless application connectivity across Kubernetes clusters

Multi-cluster ingress management unifies ingress management




Automate critical connectivity needs

Load balancing, resource utilization, disaster recovery and more

Red Hat Summit

Build a **trusted**
software supply chain
with Red Hat



Software supply chain attacks: a matter of when, not if

Ransom paid but a mere fraction to the overall
downtime and recovery costs of a data breach



742%

average annual increase in
software supply chain
attacks over the past 3 years¹

45%

of organizations worldwide
will experience supply
chain attacks by 2025²

1 in 5

data breaches are due
to a software supply
chain compromise³

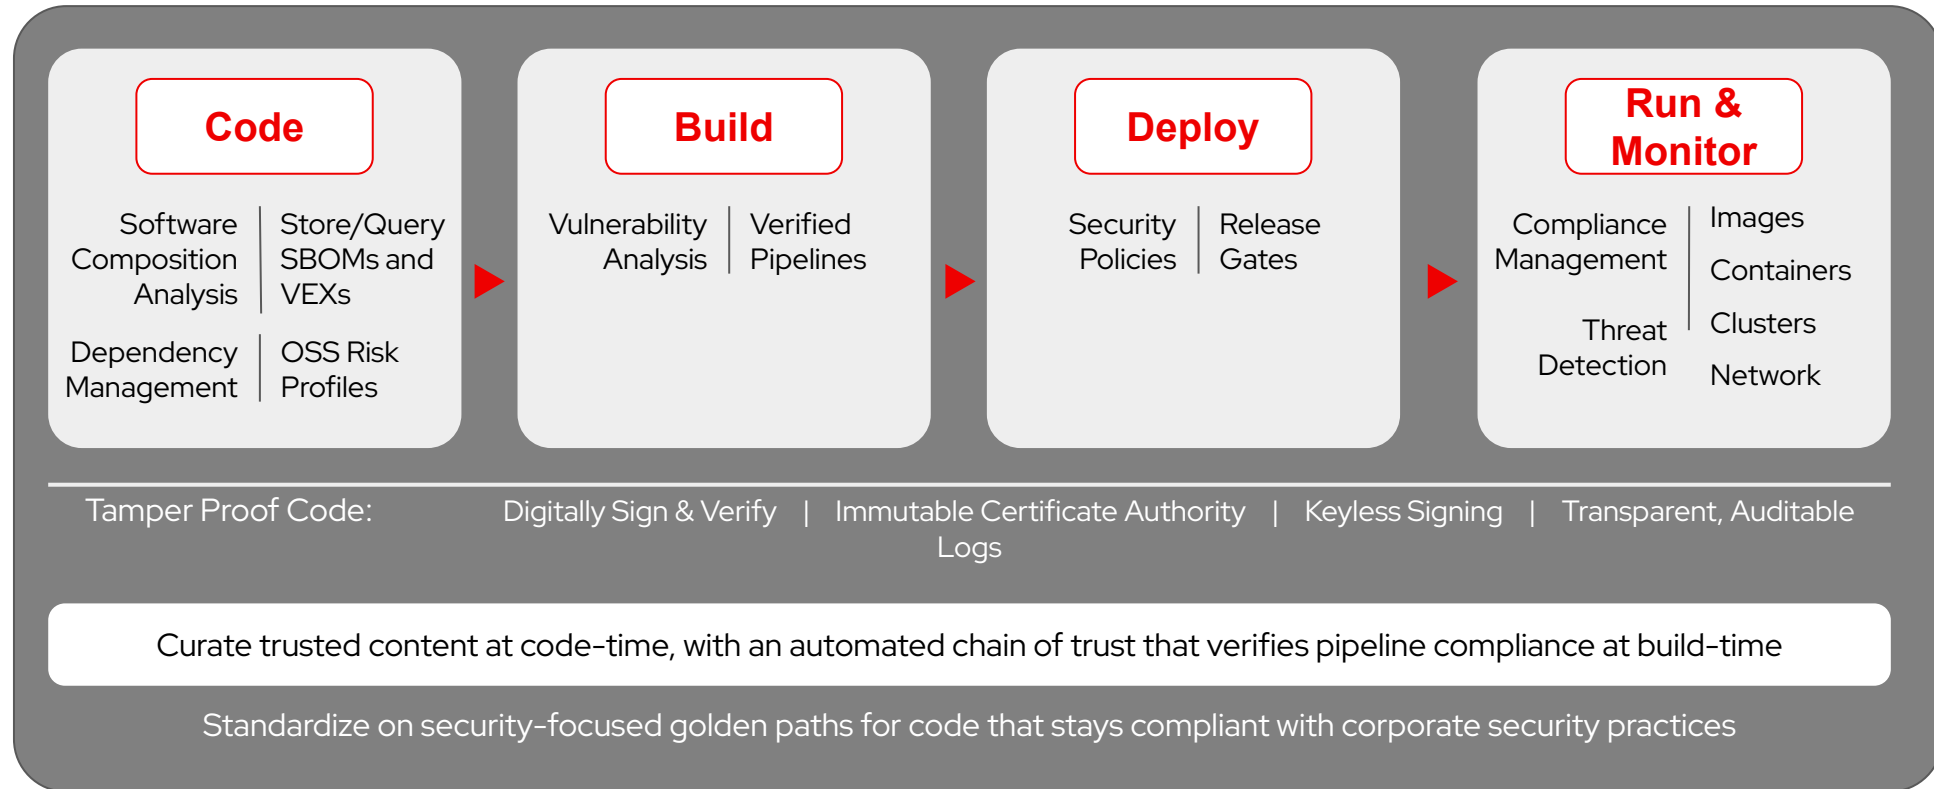
71%

YoY increase in cost
of average ransom
payment⁴



Accelerate Innovation that Safeguards User Trust

Delivered with integrated security guardrails at every phase of the software development lifecycle



Build and deploy platform, pipeline and applications as-code to an auditable, declarative state that's continuously monitored

Red Hat® Summit: Connect 2024

Join us as we explore the future of AI, hybrid cloud, open source technology, and IT.

WEST

- ▶ Bellevue, WA - Sept. 19
- ▶ Sacramento, CA - Sept. 26
- ▶ Scottsdale, AZ - Oct. 3
- ▶ Los Angeles, CA - Oct. 16
- ▶ San Jose, CA - Oct. 23

EAST

- ▶ Arlington, VA - Sept. TBD
- ▶ Waltham, MA - Sept. 10
- ▶ New York City, NY - Sept. 12
- ▶ Philadelphia, PA - Oct. 3
- ▶ Detroit, MI - Oct. 9
- ▶ Jacksonville, FL - Nov. 6
- ▶ Alpharetta, GA - Nov. 19

CENTRAL

- ▶ Chicago, IL - Oct. 1
- ▶ Austin, TX - Oct. 3
- ▶ Minneapolis, MN - Oct. 10
- ▶ Dallas, TX - Oct. 15
- ▶ St. Louis, MO - Oct. 30

CANADA

- ▶ Ottawa, ON - Sept. 17
- ▶ Montreal, QC - Oct. 9
- ▶ Vancouver, BC - Oct. 24
- ▶ Toronto - Nov. 5

Learn more:

red.ht/SummitConnect24



Red Hat
User Group Community



Red Hat

Red Hat
Summit

Thank you



[linkedin.com/company/red-hat](https://www.linkedin.com/company/red-hat)



[facebook.com/redhatinc](https://www.facebook.com/redhatinc)



[youtube.com/user/RedHatVideos](https://www.youtube.com/user/RedHatVideos)



twitter.com/RedHat