Red Hat Summit

Connect

Cloud Development Environment with Red Hat OpenShift Dev Spaces and Personal AI Assistants

Red Hat

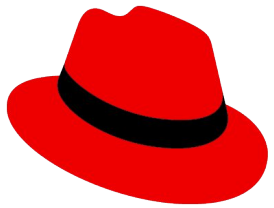

Red Hat

Ilya Buziuk

Principal Software Engineer

Dr.-Ing. Manuel Hahn

Associate Principal Specialist Solution Architect

Red Hat

Cloud Development Environment (CDE)

Created using Gemini

Red Hat

Why CDEs are gaining popularity?

Cloud Spreading

Low Latency Networking

Large Codebases / Monorepos

Laptop / Desktop Compute Power Plateauing

Developer Productivity

Code Leaks

Red Hat

Red Hat OpenShift Dev Spaces

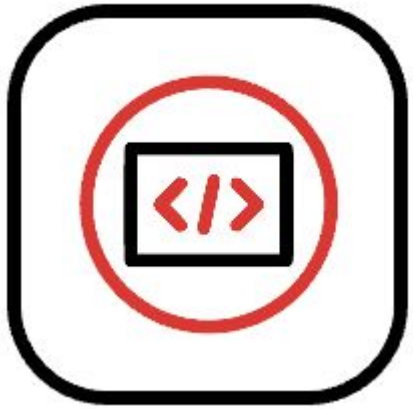

Kubernetes-based Cloud Development Environments for Enterprise Teams

Red Hat

Devfile

An open standard defining containerized development environments.

Red Hat

Devfile Dev Spaces CDE

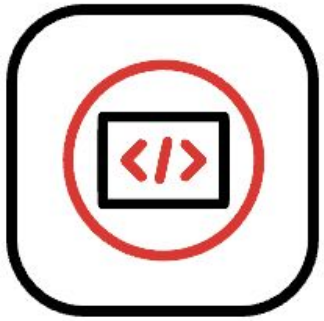

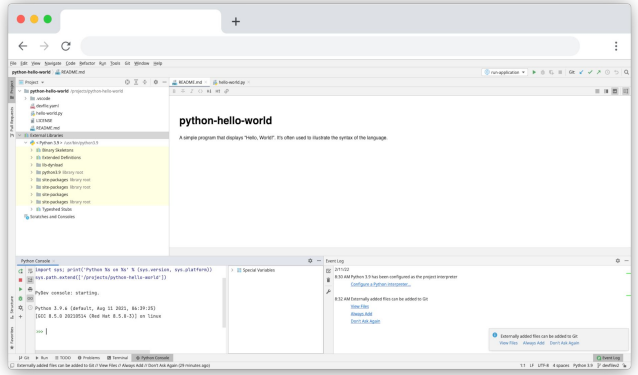

```
schemaVersion: 2.3.0
metadata:
  name: sample-java-devfile
components:
  - name: tools
    container:
      image:
       registry.redhat.io/devspaces/udi-rhel8:3.2-26
      memoryLimit: 3Gi
      cpuLimit: 2
      endpoints:
        - exposure: public
          name: my-spring-boot-api
          protocol: https
          targetPort: 8080
      volumeMounts:
        - name: m2
          path: /home/user/.m2
  - name: m2
    volume:
      size: 5G
commands:
  - id: 1-build
    exec:
      component: tools
      workingDir: ${PROJECT_SOURCE}
      commandLine: mvn -DskipTests clean package
```

Red Hat

Development Image with all the dependencies

```
schemaVersion: 2.3.0
metadata:
  name: sample-java-devfile
components:
  - name: tools
    container:
      image:
       registry.redhat.io/devspaces/udi-rhel8:3.2-26
      memoryLimit: 3Gi
      cpuLimit: 2
      endpoints:
        - exposure: public
          name: my-spring-boot-api
          protocol: https
          targetPort: 8080
      volumeMounts:
        - name: m2
          path: /home/user/.m2
  - name: m2
    volume:
      size: 5G
commands:
  - id: 1-build
    exec:
      component: tools
      workingDir: ${PROJECT_SOURCE}
      commandLine: mvn -DskipTests clean package
```

Red Hat

Full control over required resources

```
schemaVersion: 2.3.0
metadata:
  name: sample-java-devfile
components:
  - name: tools
    container:
      image:
       registry.redhat.io/devspaces/udi-rhel8:3.2-26
      memoryLimit: 3Gi
      cpuLimit: 2
      endpoints:
        - exposure: public
          name: my-spring-boot-api
          protocol: https
          targetPort: 8080
      volumeMounts:
        - name: m2
          path: /home/user/.m2
  - name: m2
    volume:
      size: 5G
commands:
  - id: 1-build
    exec:
      component: tools
      workingDir: ${PROJECT_SOURCE}
      commandLine: mvn -DskipTests clean package
```

Red Hat

```
schemaVersion: 2.3.0
metadata:
  name: sample-java-devfile
components:
  - name: tools
    container:
      image:
       registry.redhat.io/devspaces/udi-rhel8:3.2-26
      memoryLimit: 3Gi
      cpuLimit: 2
      endpoints:
        - exposure: public
          name: my-spring-boot-api
          protocol: https
          targetPort: 8080
      volumeMounts:
        - name: m2
          path: /home/user/.m2
  - name: m2
    volume:
      size: 5G
commands:
  - id: 1-build
    exec:
      component: tools
      workingDir: ${PROJECT_SOURCE}
      commandLine: mvn -DskipTests clean package
```

Expose application for testing and collaboration purposes

Red Hat

Persist maven/npm/python repositories across workspace restarts

```
schemaVersion: 2.3.0
metadata:
  name: sample-java-devfile
components:
  - name: tools
    container:
      image:
       registry.redhat.io/devspaces/udi-rhel8:3.2-26
      memoryLimit: 3Gi
      cpuLimit: 2
      endpoints:
        - exposure: public
          name: my-spring-boot-api
          protocol: https
          targetPort: 8080
      volumeMounts:
        - name: m2
          path: /home/user/.m2
  - name: m2
    volume:
      size: 5G
commands:
  - id: 1-build
    exec:
      component: tools
      workingDir: ${PROJECT_SOURCE}
      commandLine: mvn -DskipTests clean package
```

Red Hat

```
schemaVersion: 2.3.0
metadata:
  name: sample-java-devfile
components:
  - name: tools
    container:
      image:
       registry.redhat.io/devspaces/udi-rhel8:3.2-26
      memoryLimit: 3Gi
      cpuLimit: 2
      endpoints:
        - exposure: public
          name: my-spring-boot-api
          protocol: https
          targetPort: 8080
      volumeMounts:
        - name: m2
          path: /home/user/.m2
  - name: m2
    volume:
      size: 5G
commands:
  - id: 1-build
    exec:
      component: tools
      workingDir: ${PROJECT_SOURCE}
      commandLine: mvn -DskipTests clean package
```

Add shortcuts for favourites commands

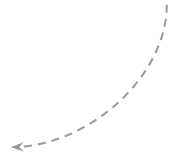

Red Hat

Demo time...

Red Hat

Personal AI Assistant

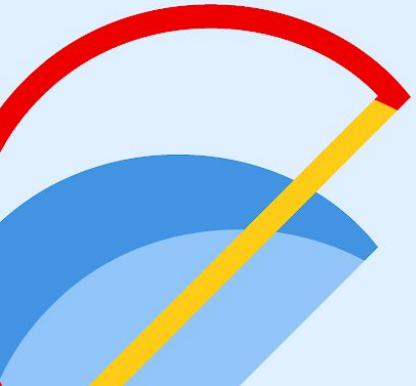

Red Hat

Why Personal AI Assistants?

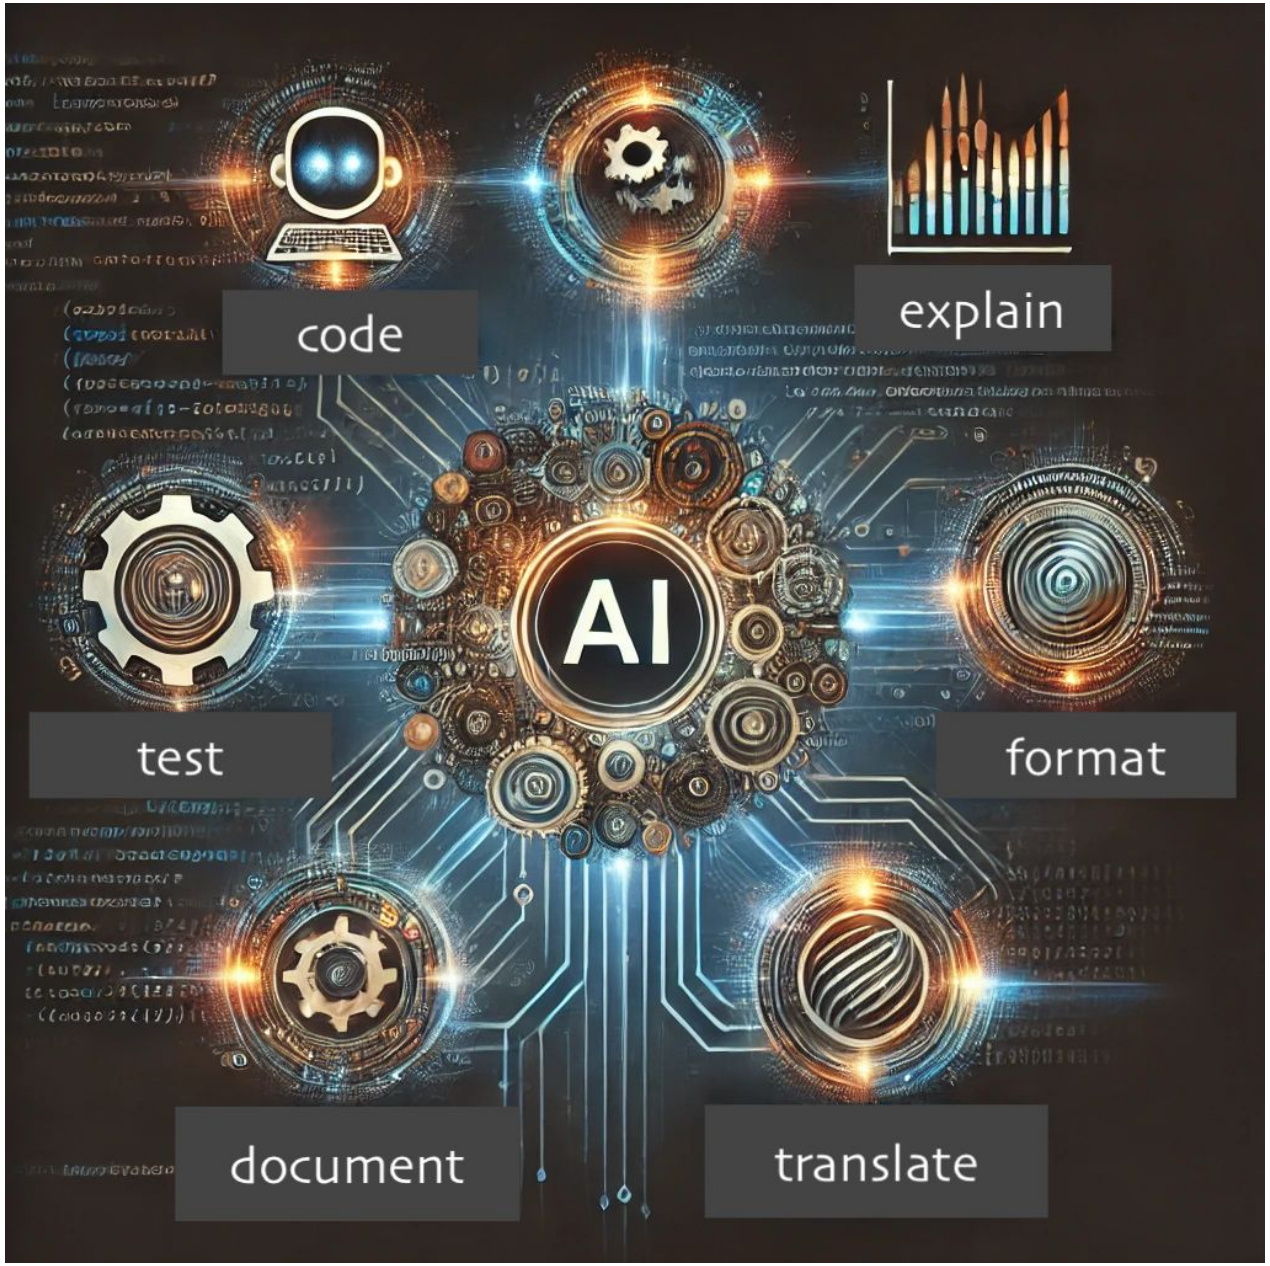

Created using OpenAI & DALL-E

Red Hat

The Open-Source Personal AI Assistant: Continue

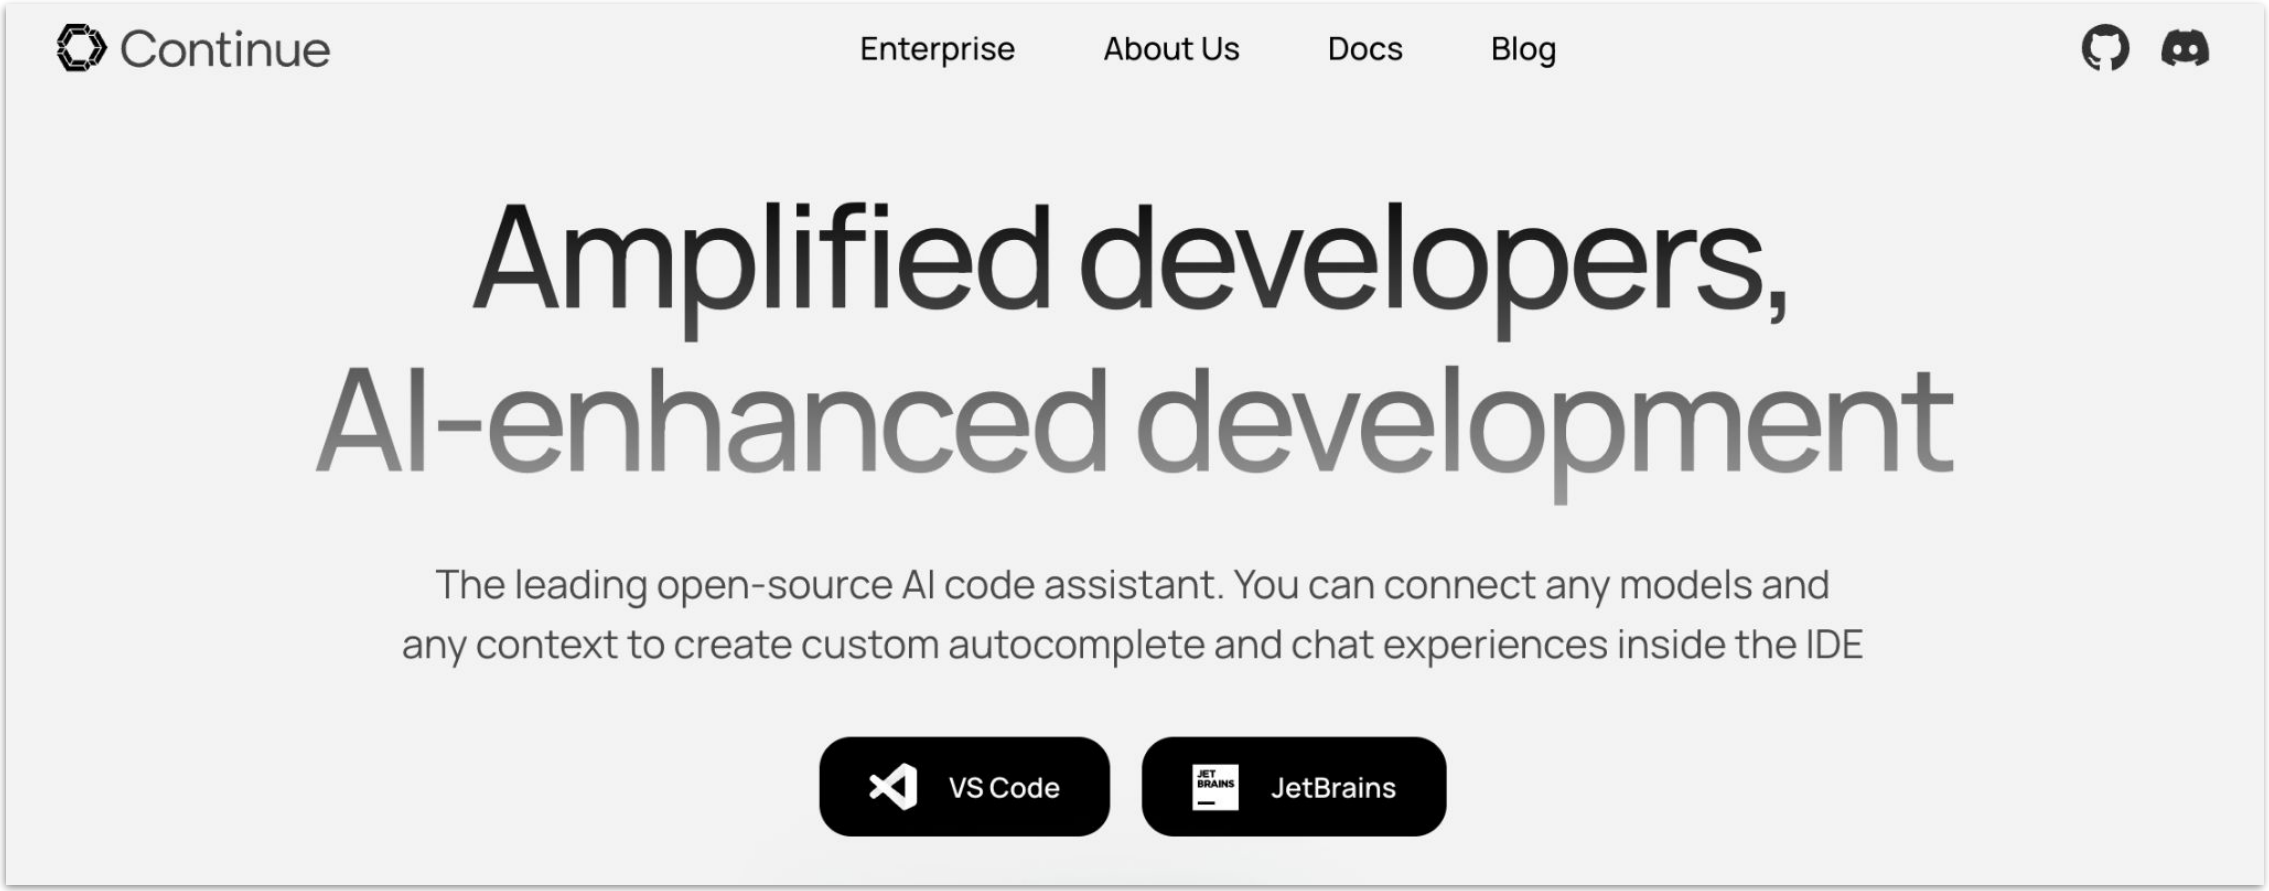

https://www.continue.dev | https://github.com/continuedev/continue

Red Hat

Demo time...

Red Hat

Large Language Model (LLM)

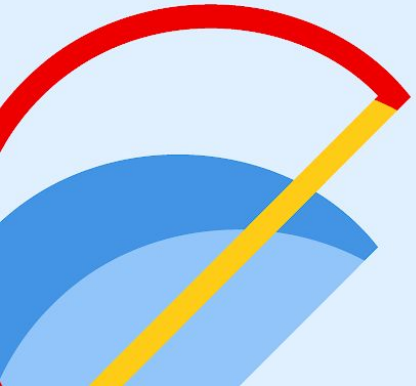

Red Hat

Where to locate LLMs in the AI jungle?

Predictive AI

- Predicts or forecasts outcomes based on historical data
- Training a model **from scratch**
- Needs **moderate** amounts of data

Generative AI

- Generates new, original content, based of trained data
- Using a **foundational model**
- Needs **massive** amounts of data
- **LLMs** are foundation models for natural language processing

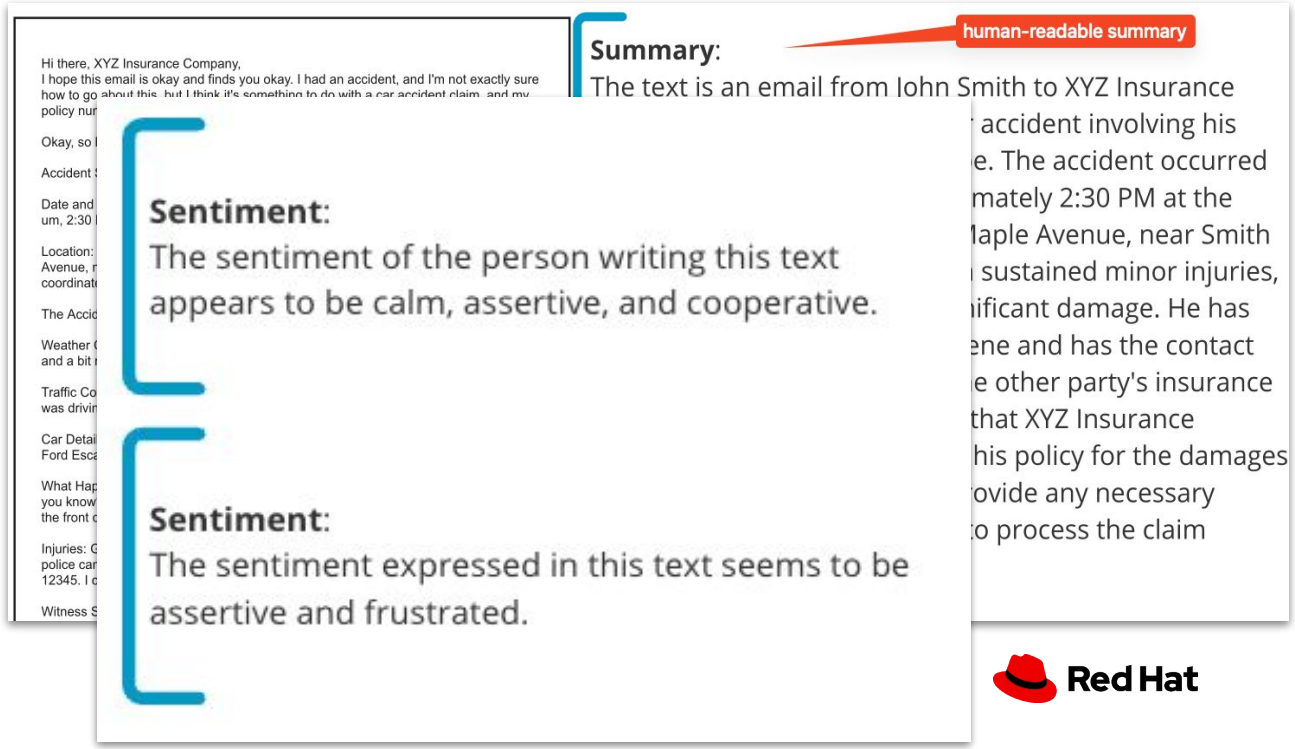

Red Hat

How to choose the right LLM?

What to consider?

- Dataset (e.g. quality?)
- Licensing (e.g. permissive, IP indemnification?)
- Deployment options (on-prem or online-only?)
- Resource needs (e.g. RAM, storage, GPU needed?)
- Customization options (e.g. fine-tuning, RAG, LAB?)
- …

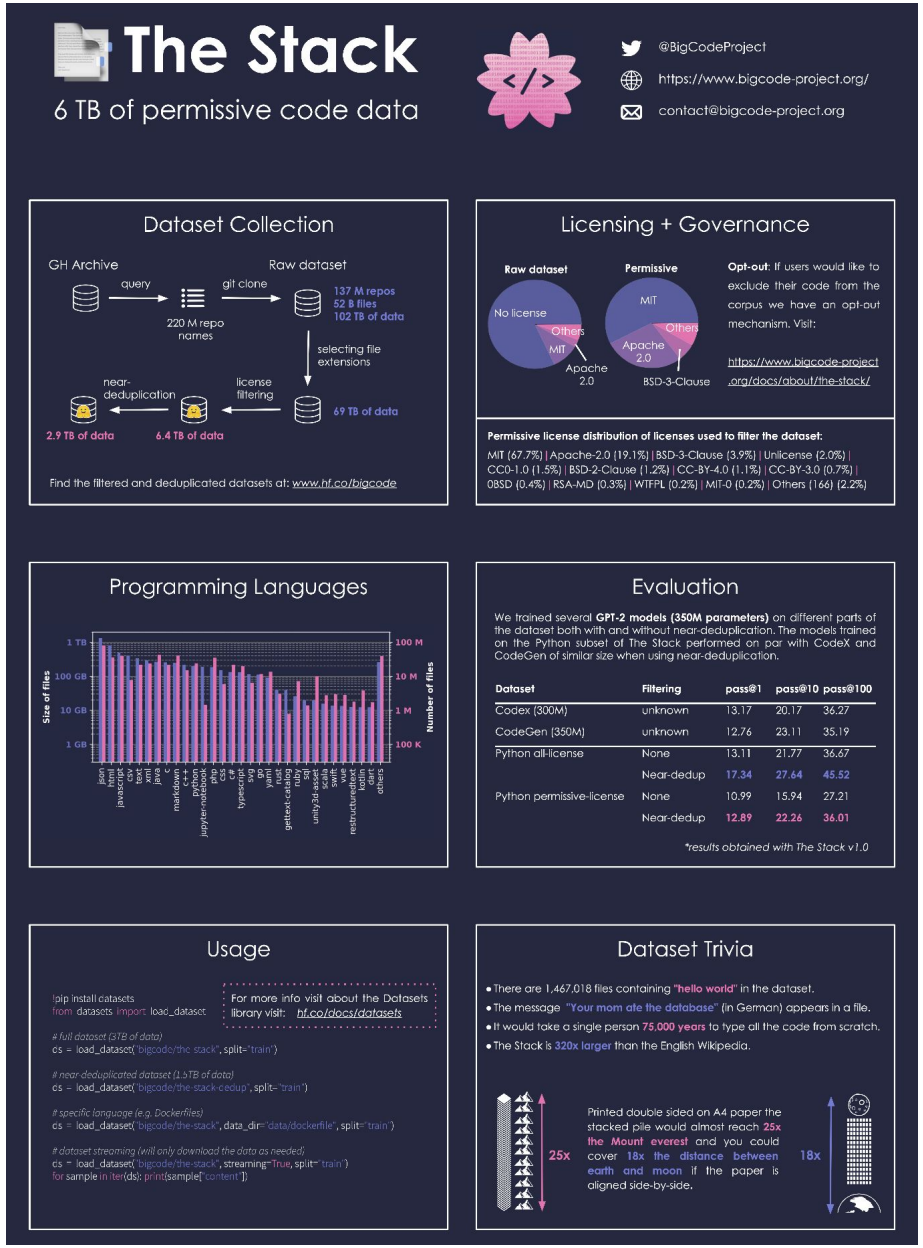

Red Hat

How to choose the right LLM?

What to consider?

- Dataset (e.g. quality?)
- Licensing (e.g. permissive, IP indemnification?)
- Deployment options (on-prem or online-only?)
- Resource needs (e.g. RAM, storage, GPU needed?)
- Customization options (e.g. fine-tuning, RAG, LAB?)
- …

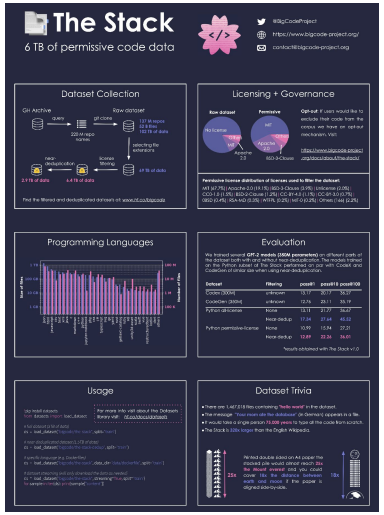

Why open source models?

- Transparency
  No black box, training data, community-driven development, …
- Control
  Stability & consistency through deliberate LLM changes, …
- Flexibility
  On-prem or cloud, customization like fine-tuning, not relying on a single provider, …
- …

Red Hat

Demo time...

Red Hat

How to get started?

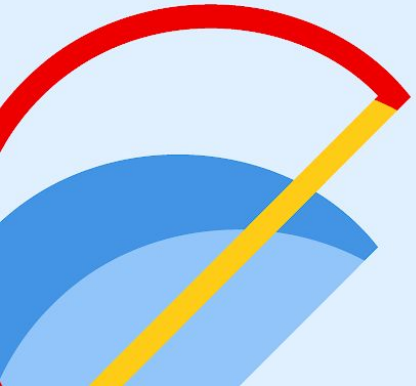

Red Hat

Developer Sandbox

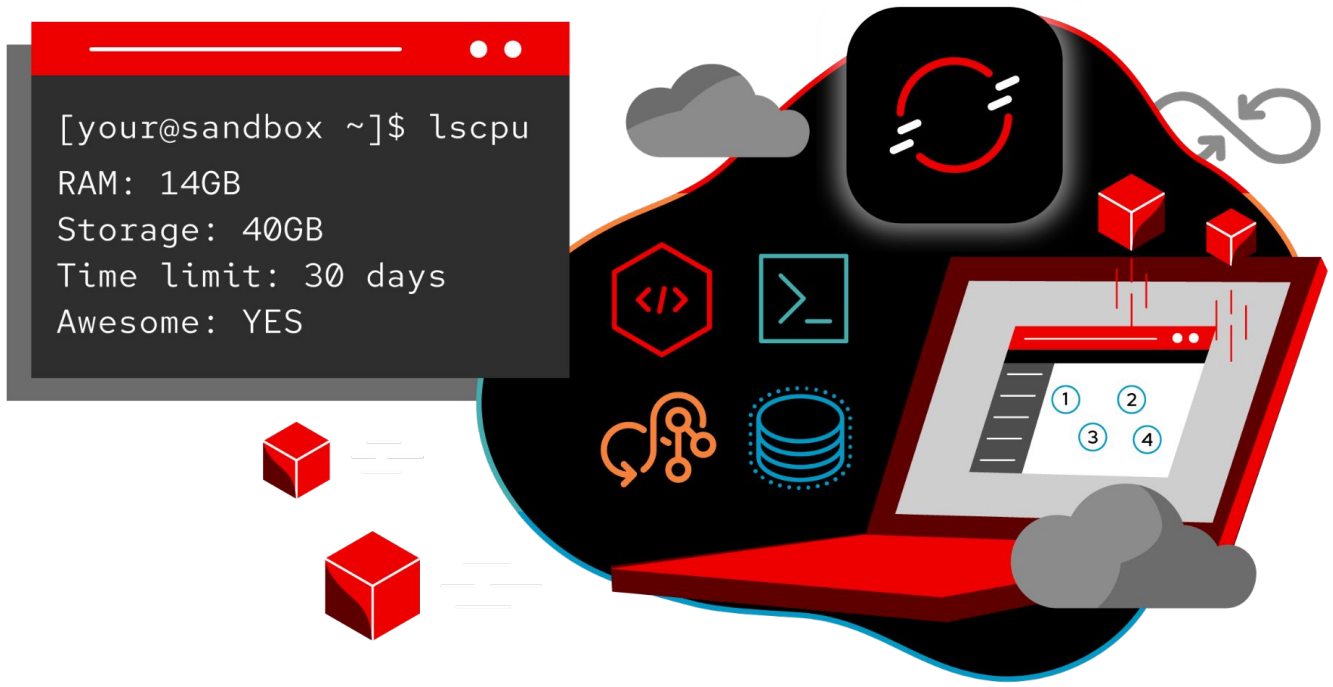

Red Hat

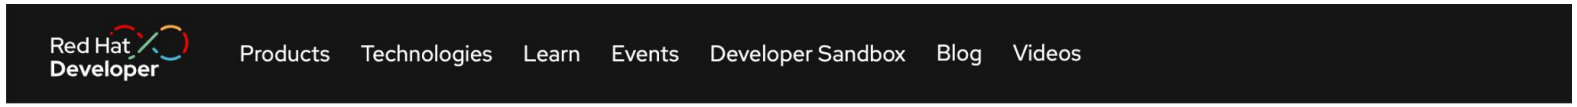

Integrate a private AI coding assistant into your CDE using Ollama, Continue, and OpenShift Dev Spaces

Level up your cloud development environment

August 12, 2024 | Ilya Buziuk, Manuel Hahn

Related topics: AI/ML, Containers, Developer Productivity, Developer Tools

Related products: Developer Sandbox, Developer Tools, Red Hat OpenShift Dev Spaces

Share:

Red Hat

Questions?

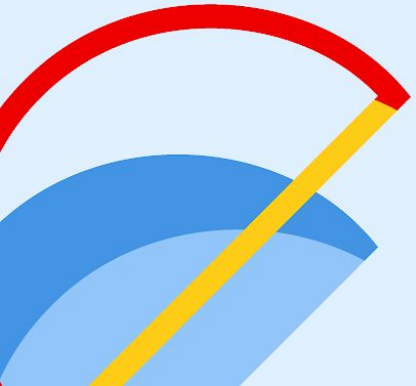

Red Hat

Session: 14:40 - 15:10

Jetzt Session bewerten!

Einfach QR-Code scannen, Session wählen und bewerten.
Vielen Dank!

red.ht/rhsc24-de-s4

Red Hat

Red Hat Summit

Connect

Thank you

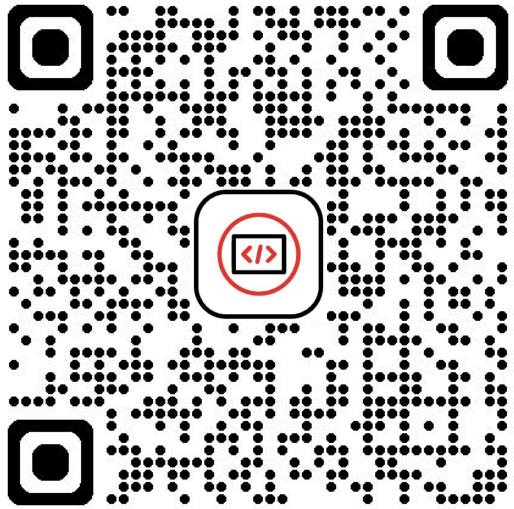

linkedin.com/company/red-hat

facebook.com/redhatinc

youtube.com/user/RedHatVideos

twitter.com/RedHat

Red Hat