Red Hat

RFI Response

JPL Open Sourced Science (OSS) for NASA Earth System Observatory (ESO) Mission Science Data Processing Study

National Aeronautics and Space Administration - Jet Propulsion Laboratory

February 1, 2022

Proposed to:	National Aeronautics and Space Administration - Jet Propulsion Laboratory Glenn.E.Campbell@jpl.nasa.gov
Red Hat point of contact:	Jackie Balenger, Business Development Manager jbalenge@redhat.com
Prepared by:	Dan McGuan, Specialist Solution Architecture Loren McMillan, Sr. Technical Writer

Treatment of Red Hat's proposal

This response document contains Red Hat trade secrets, including but not limited to sensitive pricing information, and may not be disclosed to a third party without Red Hat's prior written consent. The exchange of information under this response shall be protected under any applicable agreement between the parties related to the treatment of confidential information. In the absence of such agreement, the recipient of this response shall protect the information contained herein in the same manner as it protects its own confidential information and may not be shared with any third party.

This is a proposal for a potential business relationship, and is neither a contract nor an offer that could be accepted by the other party. If the parties decide to enter into a business relationship, they will memorialize the terms of that relationship in a separate agreement between them or their respective business partners.

Red Hat

Contents

Company Information 1

Areas of Inquiry 2

- Data processing system architecture 2
- Open science 3
- Component technologies 4
- Downstream interoperability 5
- Other Recommendations 6

Relevant Publications and Resources 7

Company Information

A. Business Name and description:

Red Hat, Inc. - We're the world's leading provider of enterprise open source solutions, using a community-powered approach to deliver high-performing Linux, cloud, container, and Kubernetes technologies. We help you standardize across environments, develop cloud-native applications, and integrate, automate, secure, and manage complex environments with award-winning support, training, and consulting services.

B. Point of Contact:

Jackie Balenger, Business Development Manager
jbalenge@redhat.com

C. Past Performance

Edge computing in action: Space

Solution: NASA is accelerating research with a containerized ML application running aboard the ISS. A proof of concept with a Hewlett Packard Enterprise (HPE) computer running Red Hat Enterprise Linux ran from 2017 to 2019. In February 2021, NASA sent the Spaceborne Computer-2 to the ISS. The compact edge device runs Red Hat CodeReady Containers, which are one-node OpenShift clusters. Developers build the code on the ground and push it to the ISS. Results of edge analysis are available immediately to ISS personnel and also relayed to scientists on the ground. Read more about the ISS edge computing solution.

NASA's Jet Propulsion Laboratory Powers Planetary Exploration with Red Hat OpenStack Platform

(Synopsis) Red Hat enables technology behind planetary exploration, helping JPL maximize server and storage capacity to process flight projects and research data through an OpenStack private cloud.

NASA's JPL has led many significant achievements in space exploration, from creating America's first satellite to sending a spacecraft to every planet in the solar system and launching all four of the Mars rovers. Today, these exploratory missions rely more significantly on cloud computing capabilities to process requests from flight projects and researchers working with mission data. Traditionally, most of JPL's infrastructure was housed in on-site server hardware.

JPL engineers built an on-site OpenStack cloud for the large and flexible cloud computing capacity they could offer mission scientists and engineers, planning to move critical compute activities that needed to be on-site into a more efficient private cloud architecture.

Areas of Inquiry

1. Data processing system architecture

Big data problems arise with the ever-increasing volume, velocity, and variety of data in integrated, event driven, tactical environments. Demand for an overabundance of data services to assist analytics and real-time decision-making has driven government and industry to deliver big data solutions. Red Hat software solutions deliver a flexible, scalable, modular and secure integration platform, based on open standards, open interfaces, and open source software to address JPL's challenges.

Red Hat Enterprise Linux (RHEL): The world's leading enterprise Linux Platform provides a stable, consistent and more secure operating system (OS) footprint. It is an open source operating system providing the foundation from which you can scale existing apps—and roll out emerging technologies—across bare-metal, virtual, container, and all types of cloud environments.

Red Hat Process Automation Manager: provides rule-based processing, workflow definition and execution, and complex event processing. Process Automation Manager contains a rules engine based on the JBoss Drools project and a workflow process engine based on the jBPM project. Red Hat Process Automation Manager is a Business Process Modeling Notation (BPMN) 2.0 compliant workflow engine.

Red Hat OpenShift Container Platform: Deployment flexibility enables component deployment across multiple hosts, virtual machines, or containers. Red Hat OpenShift Container Platform provides a container management and hosting solution using Linux containers based on Docker and Kubernetes upstream technologies.

Red Hat OpenStack Platform: Red Hat OpenStack is a highly scalable cloud platform capable of satisfying almost all work loads. In terms of compute, deployments of Red Hat OpenStack using just three dedicated controllers have supported up to 300 compute nodes, each supporting numerous VMs. Controller nodes can be split into their sub-components - these components can be scaled across many nodes to enable even a greater volume of compute nodes to be supported. Storage limitations are down to the selected storage for the VIM in question. If selecting Ceph, the maximum number of supported storage nodes is 72. The amount of storage at each storage node and the type will vary depending on use case - but maximum storage capacities running into **dozens of petabytes** are currently possible to support mission science processing objectives.

Processing capabilities include preparation of data for later interaction with analysts and analytics processes or persistence to the storage tier. In this tier we route the data collected from a variety of sources to different destinations and classify the data flows. Data Processing is the first location in the enterprise data pipeline, where analytic processes are performed against incoming data. Coupling Camel data and message routing with Red Hat Process Automation Manager provides a robust platform for routing data to the various analytics processes. The use of standards and specifications defines the tagging and normalization required by the system to support current and future intelligence requirements.

2. Open science

A mission science data processing system should have the following characteristics/capabilities:

- A modular, flexible, scalable and secure integration platform based on open standards, open interfaces, and open source software.
- Implement a modular architecture. This approach enables organizations to incrementally take advantage of innovations and adopt best-value services provided by industry - now and into the future.
- Portability. Red Hat OpenShift Container Platform and Red Hat Enterprise Linux provides a single common, and consistent application platform to facilitate application workloads to be portable, hybrid, and multi-cloud footprints.
- Enable continuous integration and continuous delivery (CI/CD) automation, and enable modern agile and DevOps/DevSecOps practices.
- A highly-secure server. Red Hat's military-grade, highly secure and stable Linux server OS serves as a foundation for every Red Hat solution, service and workload. Red Hat invests in security, providing an environment that meets Common Criteria and can be installed as a STIG compliant OS.
- Distributed integration: A few dozen high-level integration patterns reflect enterprise work and dataflows. When these integration patterns are deployed within containers, they can be deployed at the scale and location needed for specific applications and teams. This approach represents a distributed integration architecture, rather than the traditional centralized integration architecture, and allows individual teams to define and deploy the integration patterns that they need with agility.
- Stable, well-managed Application Programming Interfaces (APIs) have a huge effect on collaboration between teams, development, and operations. APIs wrap key assets in stable, reusable interfaces, which can be used and reused as building blocks across the OMB. APIs can be deployed together with containers to different environments, allowing different users to interact with different sets of APIs.
- For both API and distributed integration technologies, containers work as the underlying deployment platform. Containers allow the exact service to be deployed within a specific environment in a way that is easy and consistent to develop, test, and maintain. Because containers are the dominant platform for DevOps environments and microservices, using containers as the integration platform produces a much more transparent and collaborative relationship between development and infrastructure teams.
- With zero downtime deployments, the enterprise will ensure mission critical applications are always available, while continuously deploying patches and new capabilities.

An example of Red Hat's experience in open science, at Lawrence Livermore National Lab, can be found in this narrative: In science and in space: Red Hat leads the way for supercomputing.

3. Component technologies

Rules-based invocation of integration routes, via Red Hat Decision Manager and Red Hat Integration, as well as dynamic message routing, via Red Hat AMQ Interconnect, enhances behavioral code for sensing. Red Hat AMQ is a standards-based, reliable open source messaging platform that provides real-time communication between applications, services, and devices (both within and across clusters). The reactive asynchronous core is non-blocking and uses reactive listeners to guarantee JPL maximum response and resource efficiency and enables load processing with fewer hardware resources.

The behavioral code for decisioning can be maximized with system-based BPMN, leveraging Decision Model and Notation (DMN), with the ability to integrate AI/ML models using PMML. Further, the behavioral code for acting can be enhanced using strongly written event topics for actors, based on the service registry and AMQ Streams, coupled with compliance and custom resource definition (CRD) – based approach for enforcing the desired actor state (e.g., fully autonomous, human-on-the-loop, and human-in-the-loop).

This combination of "reactive control system" and "deliberative control system" provides the greatest flexibility in behavioral code programming as it permits isolated instances of self-governed discovery and manipulation of sensory data within a broader model-based reflex control (reactive) framework. Red Hat's graphical composition tools, expanded capabilities with projects like Kogito, and Integrated Development Environment support for Visual Studio Code, Eclipse, Intellij, and browser-based Red Hat CodeReady Workspaces all facilitate minimizing the burden and improving approachable software development for behavioral code and models.

Red Hat OpenStack Platform delivers an integrated and open foundation to create, scale, and manage a secure and reliable public or private OpenStack cloud. It delivers a production-ready cloud platform that combines the Red Hat-hardened OpenStack infrastructure, co-engineered and integrated with Red Hat Enterprise Linux. Ease of use through unified management. Red Hat OpenStack Platform can help the NASA IT operations teams automate the provisioning of bare-metal Red Hat Enterprise Linux resources for OpenShift Container Platform. This automation extends the Red Hat OpenStack Platform capability to provide NASA with a unified cloud solution for virtualized workloads, and can be useful for enterprises evolving toward cloud-native or containerized applications on bare-metal servers.

Red Hat OpenShift Container Platform is a container application platform that includes an enterprise-grade Linux operating system, container runtime, networking, monitoring, registry, and authentication and authorization solutions. It provides a consistent hybrid cloud foundation for building and scaling containerized applications and is based on a RHEL foundation. Red Hat OpenShift is supported on bare metal, virtual, hyper-converged, private cloud, public cloud and/or hybrid environment(s). Red Hat OpenShift is extensible to emerging Kubernetes-based frameworks and technologies, making it easier for teams to do more with containers. Benefit from greater productivity with Knative, Istio, and machine-learning / artificial intelligence.

4. Downstream interoperability

Interoperability, workload portability, and flexibility of open source software to enterprise environments should be characteristics of an open hybrid cloud strategy. Everything above and below the OS is abstracted—every environment, every app—thereby providing consistent interaction with any app, in nearly any environment, without retooling the app, retraining people, splitting management, or sacrificing security. And because it is all open source, your data and solution is **agnostic to the cloud provider(s)** hosting the infrastructure. The vision extends out to the edge by enabling deployments of applications, from core to the edge of the network, to use the same open hybrid cloud platform and all its benefits. Red Hat Cloud Suite includes an open source operating system, virtualization tools, a private cloud platform, public cloud interoperability, a containerized app-development platform, and unified management system.

By building **microservice architecture-based systems,** teams can integrate their components through APIs, which provides a loose coupling between interdependent components. This enables component development, in the appropriate languages and use of the best technical solutions for a given problem set, as well as minimization of the impact of changes to an underlying component's implementation.

API description formats act as a contract that end users can utilize to understand how to best work with your API. This contract is language-agnostic, and readable by both humans and machines, helping to streamline adoption and improve interoperability between applications.

Interfacing between various data exchanges and across an ecosystem that spans a multitude of legacy, current, and modern standards, protocols and mechanisms should be a key priority. Red Hat Integration, the collection of Red Hat integration capabilities, supports this breadth of interoperability, as it is built on open source technologies that rely on industry standards, protocol and patterns. Red Hat Integration is comprised of **Red Hat AMQ**, a distributed traditional messaging broker and real-time message streaming platform; **Red Hat Fuse,** an integration platform to provide connectors to link to various systems and data, data transformation, service composition and orchestration; and **Red Hat 3scale API Management**, to share, manage, and secure APIs. These three products can be deployed individually, but are tightly aligned to provide a modern, agile, API-first development approach to developing enterprise integration platforms.

Interoperability performance measurements and metrics. Red Hat OpenShift provides verifiable performance measurement and metrics, to evaluate interoperability. RHEL enables system administrators to monitor system performance using the Performance Co-Pilot (PCP). PCP is a suite of tools, services, and libraries for monitoring, visualizing, storing, and analyzing system-level performance measurements.

5. Other Recommendations

To minimize risk, we recommend partnering with an organization that has a history of collaboration with NASA - Red Hat is ideally suited to support NASA's goals and has extensive past performance, throughout the U.S. Government - in particular NASA and the Department of Energy. We are the market leader in open source software. We have been heavily engaged in emerging technology initiatives throughout the Government's scientific agencies. Our performance includes, but is not limited to, the following programs of record:

- NASA International Space Station - SpaceBorne Computer 2: a collaboration with HPE to fly custom hardware to the ISS with Red Hat OpenShift on board for astronauts to sequence DNA on orbit.
- NASA JPL Joint Community Cluster: an OpenStack environment designed for various mission simulation requirements, including the Perseverance Rover.
- Department of Energy CORAL - Summit and Sierra Supercomputers: Red Hat Enterprise Linux (RHEL) is the standard operating system on both machines, which were the #1 and #2 fastest machines when they were deployed.
- Oak Ridge CADES - Compute and Data Environment for Science: an on demand HPCaaS environment for research at the Spallation Neutron Source, hosted by ORNL OLCF.
- NASA Orion SOCRRATES and PLATO- Flight software test and evaluation system: a custom built OpenStack IaaS designed exclusively to run multiple versions of the Orion's flight software in parallel, both natively and on SIMICS, to simulate software readiness for vehicle flight.

Predominantly, Red Hat Enterprise Linux is called out by name in the NASA End User Services Program Office memo (NASA-STD-2804, 3.1.3 and 4.4), as the Linux operating system of record, NASA-wide. To quote the memo:

> (*3.1.3) "The selection of standards for a broad and cost-effective infrastructure using commercial off-the-shelf and well-supported open source products to the greatest extent practical:*
>
> - *Interoperability both within and when used remotely to NASA*
> - *Flexibility for future growth*
> - *Consistency with generally accepted consensus standards as much as feasible*
> - *Security for NASA systems and data"*
>
> ***(4.4) "All new and refreshed Linux systems must run a supported RHEL distribution. Vendor provided and -supported versions of applications shall be used. The version of application the vendor provides in their update stream shall supersede any listed in the CRC."***

Since NASA JSC's own in-house developed applications (Trick, GUNNS, etc.) are all compiled to run on Linux, we are the ideal partner for this pursuit. Red Hat has the capability and the experience, within a variety of environments, that align with NASA JPL's needs.

Relevant Publications and Resources

Red Hat Customer success stories

Red Hat Products and Solutions

Red Hat website

Why build a Red Hat cloud?