

TURN-KEY DEVOPS & DATA FABRIC FOR THE EXPANDING FRONTIER

A platform for enabling the modernization of development and integration practices at scale to power data-driven innovation

Challenge

Agencies require the ability to develop their needed advanced capabilities and integrations in an agile manner. Deploying to distributed hybrid clouds is now the standard, and as additional integrated systems create more disparate data there's the crucial need to utilize Machine Learning to gain actionable insight.

Solution

From accelerating DevOps and the System Development Life Cycles to streaming data from the Edge and IoT enabling actionable insight with Machine Learning and AI, a complimentary set of Enterprise Open-Source solutions can provide the platform required to enable workloads at any scale, in any environment.

Benefits

- A unified platform that operates and scales across hybrid and multi-cloud
- Enhanced security across the infrastructure, platforms, and applications developed and deployed
- Best practices, support, and training from industry-leading Enterprise Open-Source leaders

As agencies and their prime contractors and partners collaborate to advance their IT systems, there are often challenges presented. Traditional infrastructure solutions and methods of developing and deploying systems are shackling the progress of innovation. More contemporary practices and technologies must be employed in order to meet the demands of a modern agency - hybrid and multi-cloud environments, agile development across different platforms, and data computation from the Edge and IoT with Machine Learning to aid in data-driven development.

Agility can make or break a mission. Commercial organizations have been benefiting from Cloud-Native technologies and DevOps practices for years and have established best practices and thriving communities that foster innovation. Federal agencies can benefit from the groundwork put forth by these commercial entities to quickly bootstrap these same practices and systems in order to maintain the competitive edge that is demanded at home, on the battlefield, and along expanding frontiers.

Some challenges can be presented when architecting next-generation systems - leveraging Public and Private Clouds requires interoperability on an infrastructure layer, software development lifecycles can quickly become convoluted, and the management of scale-out and scale-up devices is difficult at best. When considering DevOps practices it's important to not focus solely on the technologies but rather the capabilities it provides and the cultural impact on your teams. As your organization uses these next-generation technologies and becomes more agile in their processes your library of delivered applications and systems will quickly grow, how can you identify the untapped value of data ingested by these systems?

Infrastructure Requirements

Utilizing both Public and Private Cloud services requires infrastructure and platform components that are interoperable and can be orchestrated across disparate infrastructure and hardware, be that Amazon Web Services (AWS), OpenStack, on-premise, or even embedded field deployments. **Red Hat Cloud Suite** is a set of technologies that provides the base infrastructure components to deploy a private cloud with **Red Hat OpenStack**. Red Hat OpenStack can be deployed onto bare-metal and virtualized systems thus providing a Private Cloud locally, in the cloud, and at the edge. The primary architecture components of a Red Hat OpenStack private cloud are the Undercloud, or the OpenStack Director, and the Overcloud which is the consumed private cloud. The Undercloud is either a single or a set of three clustered nodes and is used to manage and provision the Overcloud. The Overcloud is the target architected Red Hat OpenStack private cloud and production environments start with 3 Controller nodes, 3+n Compute nodes, and 3+n Ceph Storage nodes.

Integrated Automation

A homogenous cloud platform can only take your hybrid and multi-cloud strategies so far - there comes a point where automation is required to integrate with and orchestrate across differing systems. **Ansible Tower by Red Hat** can provide the glue needed to bring these environments together. It utilizes simple & easy to read/write YAML, is cross-platform, and agentless by utilizing standard transport agents such as SSH and WinRM. Ansible Tower can be deployed onto traditional server platforms or as a containerized solution and presents a centralized and managed automation platform with a CLI, Web UI, REST APIs, and more. Ansible Tower is able to orchestrate workloads across servers, networking devices, cloud infrastructure, databases, and more to perform tasks such as, but not limited to, application deployment, provisioning and configuration management, continuous delivery, and security & compliance.

Cloud-Native Platforms

With cloud platforms the consumption of compute, networking, and storage resources are abstracted into easy to use APIs and user-centric web portals.

This offers additional agility in provisioning environments and allows for a much greater level of manageability. While this provides a great deal of value, cloud-native solutions require more contemporary platforms such as containers and the orchestration of those containers at scale in a secure and orderly fashion is of extreme importance. **Red Hat OpenShift Container Platform**, a component of the Red Hat Cloud Suite, provides you the enterprise-ready Kubernetes container platform to run these cloud-native solutions that are required as our needs continue to evolve.

Red Hat OpenShift orchestrates container workloads across a cluster of nodes, traditionally in a 3 Master, 3 Infrastructure, and 3+n Compute node configuration. These nodes can be bare-metal, virtual machines on-premise, in the cloud, or integrated as a layer on top of a Red Hat OpenStack Private Cloud. Red Hat OpenShift also provides all the needed networking, routing, service layer, image repository, and other components that are vital to a container platform that are not offered with other orchestration platforms.

For IT Operators, Red Hat OpenShift provides an easy to manage environment with control over quotas, role-based authentication, and a standardized operating environment that features a security inheritance model.

Developers will benefit greatly from the self-service portal and catalog of quick-deploy offerings such as Nginx, Tomcat, JBoss EAP, and even more with the inclusion of additional service brokers such as the AWS Service Broker which allows the provisioning of resources in AWS right from the catalog. This and the native integration with your DevOps CI/CD toolset means your developers can focus on what they do best and not on the underlying platform.

Management & Day 2 Operations

We require the ability to deploy VMs, containers, and orchestrate our software development lifecycle. This entails the management of different systems. **Red Hat CloudForms** and **Red Hat Satellite**, both part of Red Hat Cloud Suite, allow you to manage your hybrid and multi-cloud environments with a set of loosely de-coupled but highly cohesive tools. Red Hat CloudForms provides insight into the current state of your cloud environments, such as AWS and Red Hat OpenStack, their resources, and aids in meeting compliance requirements.

The deployment process of Red Hat CloudForms is based on an OVA image that can be imported into modern hypervisors and cloud platforms such as AWS, and is integrated with these cloud platforms via APIs.

Red Hat Satellite allows you to centrally manage your Red Hat environment. It is traditionally deployed as a VM, can be made highly available, and multiple Satellite servers can be set into different environments and managed from a central Satellite server. Red Hat Satellite and Red Hat CloudForms integrate to provide you an aggregated view across all your managed environments and with the inclusion of Red Hat Ansible allows you to run playbooks across your infrastructure.

Defense In Depth

As DevOps enables incredible speed in developing and deploying software the double-edge to that agility starts to become more pronounced. **Atomic Secured Linux** provides an OSSEC-based Multi-Level Security solution for your infrastructure and cloud. Atomic Secure Linux provides a security blanket across your Linux systems with intrusion prevention, application firewalls, and role-based authentication control. Atomic Secure Linux installs onto a Linux system with a few commands without replacing core components and provides a CLI, Web UI, and APIs to interact with.

As we embed security in multiple layers across all of our environments, **SysDig** provides additional monitoring and security measures for our container platform. SysDig integrates with the Kubernetes orchestrator to provide enhanced metrics, based on Prometheus, showing the health and lifecycle of your containers and applications deployed in them. SysDig also provides security enhancements for the containers running and the applications inside with image scanning and runtime protection. SysDig is agent-based and is installed on each node running container workloads.

Secure Software Factory

The most capable and secure platform is still only the canvas on which to develop. Your CI/CD engine is what drives the software development life cycle. Traditional methods of deploying a CI/CD engine, such as Jenkins, can become overwhelming quickly. Different configurations, dependencies, and plug-ins

found on individual developer environments don't often match what is in production, and the large centralized monolithic build server suffers from the same issues of stacking incompatibilities and does not scale.

CloudBees is the industry-leading provider of Enterprise DevOps and **CloudBees Core** is their CI/CD engine. CloudBees Core is based off Jenkins and extends additional security and compliance measures, enhanced role-based authentication, and simplifies infrastructure to offer CD as a Service that can scale out new Masters and Team Masters on demand from a centrally managed CloudBees Operations Center. It can be deployed as a traditional Java application via a WAR file, or can quickly be integrated into Red Hat OpenShift as a container-based service. From the CloudBees Operation Center, you can set configuration and governance models for your different teams and pipelines, set their curated plugin selection, and build Distributed Pipeline Architectures to scale out efficiently. CloudBees Core can deploy to a container in Red Hat OpenShift, a VM in Red Hat OpenStack, or into an EC2 node in AWS, and has over 1,600 plugins to use throughout your software development lifecycle.

Aside from their best-of-breed technology, CloudBees was founded by industry veterans from Red Hat and JBoss, and their CTO is the original developer of Jenkins. They're the most authoritative source for best practices for your DevOps campaigns and their SMEs are able to guide your teams along every step of the way.

Data as a Platform

The new agile landscape has more applications deployed and more frequently. These applications, be it living as a service or deployed to an IoT device, generate streams of data that can be used to deliver upon better mission objectives. Before you are able to ingest, infer, and validate different models against this growing data plane, considerations for the platform must be made. **Hortonworks CloudBreak** allows you to easily deploy scalable clusters for different data workloads, such as analyzing data at rest or transforming streaming data in motion. Hortonworks CloudBreak is deployed into a VM, and uses PostgreSQL as a storage back-end. There is an API, CLI, and Web UI presented and connects to public and private clouds via APIs in order to deploy.

To deploy a cluster with Hortonworks CloudBreak one simply selects an Ambari Blueprint that fits their desired cluster type, target the cloud provider such as AWS or an OpenStack cluster, and in a few moments you'll have a brand new cluster for your data-driven workloads.

Data as a Platform

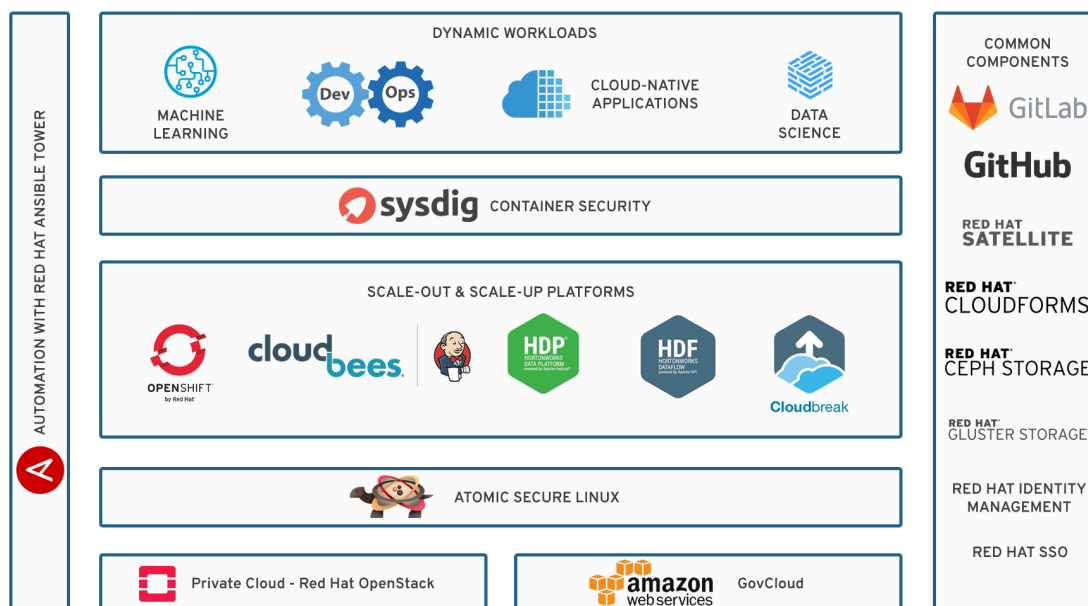
To enable your developers, engineers, and scientists identify, ingest, and innovate upon their data sets you need a platform that can scale and can get out of the way so focus can be set on generating actionable outcomes and not spending time operating a cluster. **Hortonworks Data Platform (HDP)** is a full-stack Hadoop powered platform that allows you to securely store, process, and analyze your data at rest with full governance across your entire data plane. HDP provides the ability to perform Extract, Transform Load (ETL) transactions distributed and at scale, in-memory processing, stream processors, and Machine Learning for better optics into advancing mission objectives. HDP is primarily based on Java applications integrated and orchestrated with YARN that provides APIs, CLIs, and Web UIs. HDP is

easily deployed via Hortonworks CloudBreak into cloud infrastructure providers. The architecture consists of Compute, Master, and Worker nodes that can scale to meet demand. The latest version of HDP brought more agile functions such as being able to compose your data science workloads into a container allowing for portability and segmentation of workloads, their configuration, dependencies, and plugins.

Streaming Analytics

As the number of systems grows, so does the amount of data generated at these endpoints and across their transit. Being able to analyze this data while in motion securely without taxing the limited compute resources on these nodes is possible with **Hortonworks DataFlow (HDF)**. HDF provides real-time acquisition, transformation, routing, and delivery of streams of data and can be event-driven. HDF is powered by Apache NiFi which was a project created and open-sourced by the NSA. It is easily deployed via Hortonworks CloudBreak, and there are components for MiniFi deployments at the edge and on IoT devices.

High-level Architecture



Source: <https://fiercesw.com/wp-content/uploads/2019/01/FAAStack.png>

Professional Services, Support, and Training

Deploying scale-out and scale-up infrastructure and platforms requires coordination across multiple components and teams. Navigating the cloud-native and DevOps landscape can be challenging, and the

expertise required to quickly execute effective Machine Learning workloads is sometimes difficult to source. All of these vendor technologies center around enterprise open-source software and include offerings from the 3 largest open-source companies in the world. These providers have an extensive history of Professional Service engagements with some of the largest agencies and corporations that extends past the traditional deployment with the guidance of your teams aided by custom generated playbooks. Ongoing training and certification options are available to enable your team to draw in the best practices and expertise of these vendors, bootstrapping your path to setting a center of excellence.

Price Lists

We can provide additional pricing details upon request

[Red Hat](#) • [Cloudera/Hortonworks](#) • [Cloudbees](#)

cloudera

Accelerating data management innovation from the Edge to AI.

Past Performance: [Public Sector Customers](#)

Management: Cloudera was founded in 2008 by some of the brightest minds at Silicon Valley's leading companies, including Google, Yahoo!, Oracle, and Facebook. And in 2011, 24 engineers from the original Hadoop team at Yahoo! spun out to form Hortonworks. Both companies were founded with the belief that open source, open standards, and open markets are best. [Cloudera Leadership Bios](#)



The world's leading provider of open source, enterprise IT solutions

Past Performance: [Customer Success Stories](#)

Management: Red Hat leaders apply open source principles—like open exchange, transparency, and meritocracy—to their management styles. It helped them become the first billion-dollar open source company, and it's how they will continue to thrive. [Red Hat Leadership Bios](#)



Building the World's First End-to-End Continuous Software Delivery System

Past Performance: [Cloudbees Customers](#)

Management: Backed by Matrix Partners, Lightspeed Venture Partners, Verizon Ventures, Delta-v Capital, Golub Capital and Unusual Ventures, CloudBees was founded in 2010 by former JBoss CTO Sacha Labourey and an elite team of continuous integration, continuous delivery and DevOps professionals. [Cloudbees Leadership Bios](#)