Managing complex workloads with confidence

Table of contents

Workloads for the digital business

For digital businesses, high-performance workloads provide the foundation for everything from maintaining operations to innovating new products and services. Hosted in the datacenter or in the cloud, workloads comprise a variety of applications and services that require continuous maintenance and performance so they run as expected.

As digital businesses grow in size, complexity, and service offerings, so does the diversity of the workloads they run. Whether used to provide email services, data backup and disaster recovery, or customer relationship management, many workloads fall into five categories:

1. IT infrastructure workloads
2. Security
3. Data workloads
4. Analytics workloads
5. Database workloads

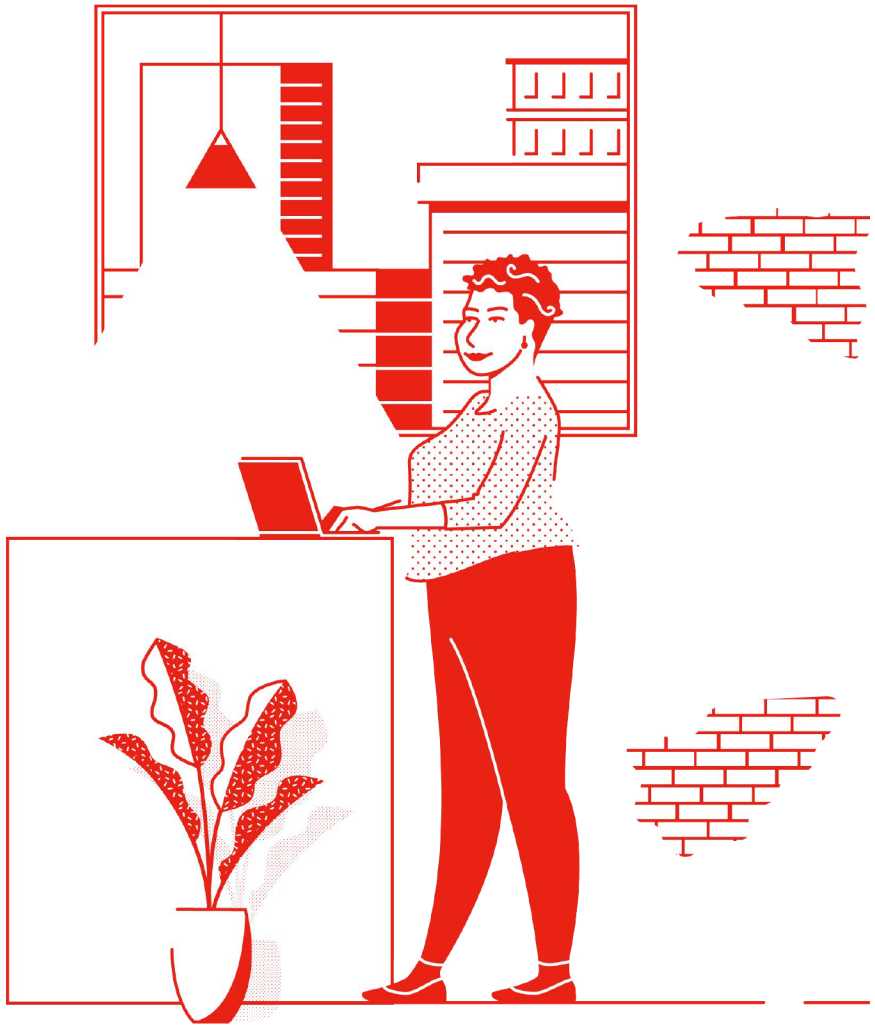

The role of open source in workload management

As organizations increasingly move to hybrid and multicloud environments, the complexity of managing workloads in these environments has grown. Enterprise IT organizations face challenges including having the flexibility and confidence to move these workloads from datacenters to the cloud; ensuring systems are up to date and easily scalable; proactively monitoring, analyzing, and remediating any security vulnerabilities; and staying compliant with industry standards.

Open source technologies play a key role in helping enterprise IT organizations address these challenges. Developed on open source software with built-in security, Red Hat® Enterprise Linux® provides a high-performance, scalable, and resilient infrastructure foundation for critical workloads.

Beyond the operating system, Red Hat provides technologies for workload management including automation, containers, business intelligence and analytics, and cloud infrastructure management. Red Hat's portfolio gives enterprise IT organizations the choice, control, and freedom they need to confidently manage their workloads.

Gain choice, control, and freedom with workloads running on Red Hat

IT infrastructure workloads

Security

Data workloads

Analytics workloads

Database workloads

Chapter 1

IT infrastructure workloads

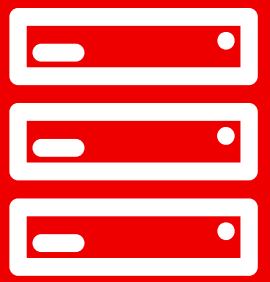

Businesses feel the pressure to improve IT operational efficiency, agility, and innovation. As technical environments become more complex, the IT organization often finds itself unable to meet the time-to-market speeds that business leaders require to satisfy customer demands. Moving to the cloud and developing within containers adds new layers of complexity.

The journey to reduce costs and complexity in IT involves the progression of five stages:

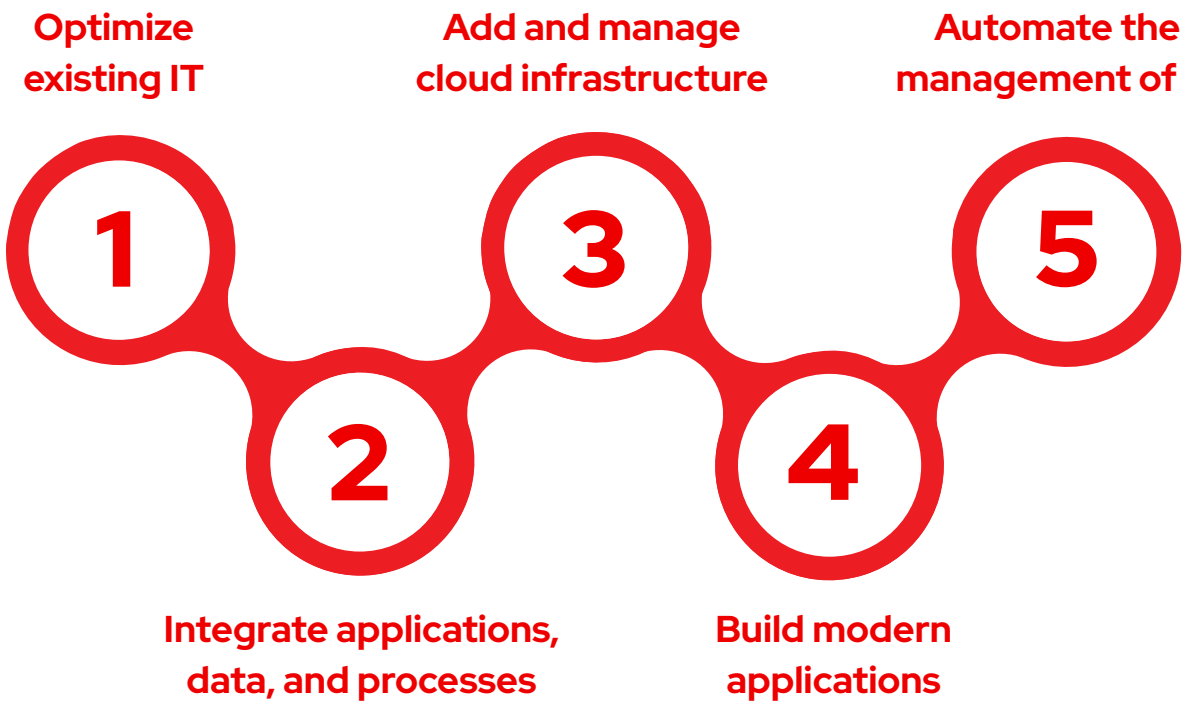

Management with automation

Automation is an essential and strategic component of IT management. For virtualized, cloud, and bare-metal environments, automation delivers a consistent, efficient approach. The technology reduces human errors, builds consistency across multiple users and environments, and speeds management processes to allow more time and resources for innovation.

Beyond automation for virtualized and cloud environments is the occasional need for the provisioning of bare metal. The right automation platform integrates with datacenter management tools to both invoke and enact the steps required for provisioning.

Virtualization management

Software tools for managing virtual environments simplify resource administration, enhance data analysis, streamline operations, monitor virtual environments, and automatically enforce rules. With large and diverse virtual environments, virtualization management tools are particularly useful to reduce manual effort in provisioning, ongoing operations, and compliance.

Public and private cloud management

In the cloud, management tools control operations in private and public cloud environments including data, applications, and services. Management tools ensure that cloud computing resources work efficiently and more securely so that workloads operating in every cloud environment operate as expected.

Many organizations build modern, cloud-native infrastructure using virtualization software technology. Today, more than **85%** of workloads are running on virtualized infrastructure.[1]

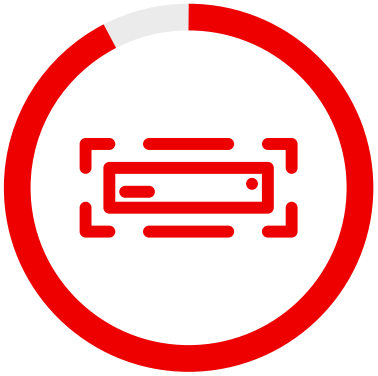

[1] IDC. "Preparing IT Infrastructure for Cloud-native Application Deployments," September 2019.

Chapter 2

Security

Security is a top concern for all IT organizations. Building security measures into the IT environment involves maintaining the confidentiality, integrity, and availability of sensitive information while blocking access to sophisticated hackers. IT security involves establishing a set of layered defense security strategies to protect these applications, infrastructures, networks, and data. Many times, security is done as a checkpoint at the end of a development cycle. Organizations now look for an integrated security process to adapt faster and more efficiently—with security that is built in at all layers of development and infrastructure life cycles, rather than bolted on.

Traditionally, IT security was focused on fortifying, maintaining, and policing the datacenter perimeter. However, that perimeter is dissolving. The way enterprise IT organizations develop, deploy, integrate, and manage IT has changed. With the growing use of public and hybrid clouds and containers, traditional methods of managing security are becoming increasingly insufficient.

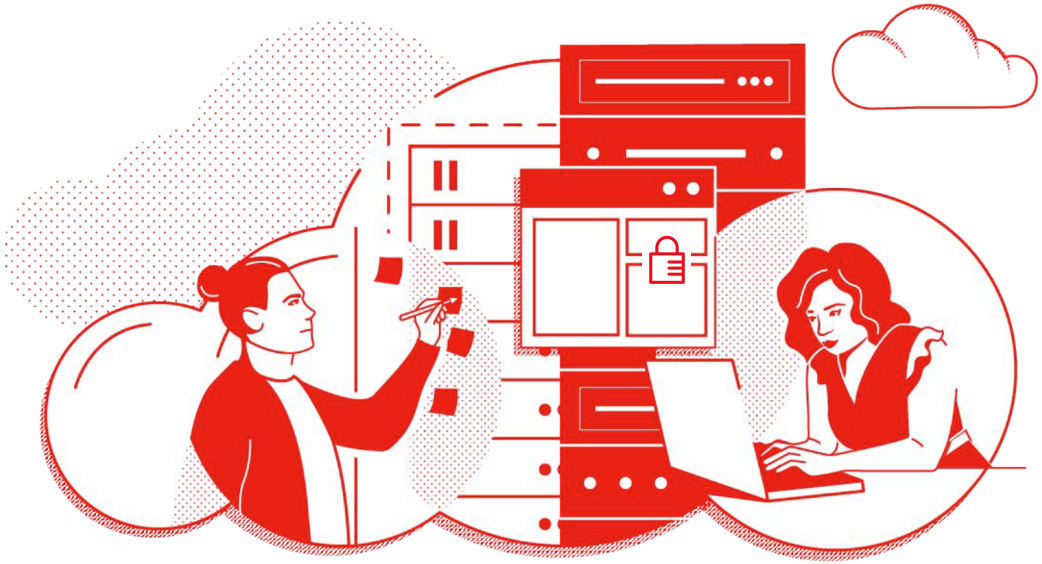

Container security

Containers make it easier to build, package, and promote an application or service across different environments and deployment targets. There are many security benefits of a containerized infrastructure, such as having a standard, hardened infrastructure with applications in line with defined security policies. Container security needs to be integrated and continuous. It must provide security for the container pipeline, application, environment, and infrastructure while integrating with enterprise security tools and meeting or enhancing existing security policies.

Cloud security

While many businesses understand the benefits of cloud computing, they are often deterred by the potential security threats. Cloud security protects data, applications, and infrastructures in cloud computing. For hybrid cloud computing specifically, the IT organization can keep sensitive data off the public cloud while still taking advantage of the cloud for data that involves less risk.

Security for cloud environments—whether on public, private, or hybrid clouds—is similar to security for any on-premise IT architecture. High-level security concerns, like data exposure and leaks, weak access controls, susceptibility to attacks, and availability disruptions, affect traditional IT and cloud systems alike.

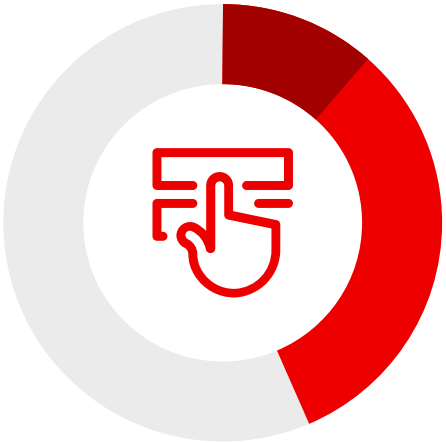

In 2019, **31%** of organizations experienced unauthorized access to cloud environments or cloud assets by outsiders, compared to **19%** in 2017.[2]

[2] SANS Technology Institute, "2019 SANS State of Cloud Security Survey," April 30, 2019.

Chapter 3

Data workloads

Data is growing quickly as the number of connected devices rises. Traditionally, data was seen as a liability, considering the resources it requires for storage and security. Today, the quality and value of data have improved. Businesses are using data as an asset, improving decision making with business intelligence and visualization.

Data is only valuable if it can be protected, processed, understood, and acted upon. With business intelligence and visualization, businesses can:

- Quickly learn and understand data, historical patterns, and trends.
- Make data-driven decisions with the help of aggregation, analysis, and visualization.
- Use data to manage business processes and policies.
- Analyze data from business operations across the organization.
- Gain user-friendly analysis with software that is easier to operate than traditional statistical analysis software.

Business intelligence

Business intelligence helps users make decisions based on current insights. Business intelligence tools and applications analyze data from business operations and transform raw data into meaningful, useful, and actionable information.

The timeliness of data is a concern for many organizations. The business intelligence tool must allow users to harness and analyze real-time data to bring speed and efficiency to decision making while having confidence in the relevance and value of the data.

Data visualization

Visualization is how data is viewed and used. It helps tell a story from data, enabling users to understand and interact with data patterns, trends, and insights by transforming it into a visual context. Data visuals can take many forms including tables, graphs, charts, and images.

Data visualization lets users from across the organization interact with data in a way that makes sense for their business function. In the past, business users had to rely on IT to access data and run reports for them. This process often took days or weeks. The data often became stale or the user's questions may have changed. Having self-service access to data in a user-friendly, visual format lets users ask questions of their data and follow the data where it leads them, on their own timeline.

To be effective, business intelligence and visualization technologies need to run on a flexible, comprehensive, and reliable foundation. By optimizing the infrastructure, the organization can continue to gather, analyze, and use data even as the technology stack changes.

Worldwide revenues for big data and business analytics (BDA) solutions are forecast to reach **US$189.1 billion** in 2019, an increase of **12%** over 2018.[3]

[3] IDC, "IDC Forecasts Revenues for Big Data and Business Analytics Solutions Will Reach $189.1 Billion This Year with Double-Digit Annual Growth Through 2022," April 4, 2019.

Chapter 4

Analytics workloads

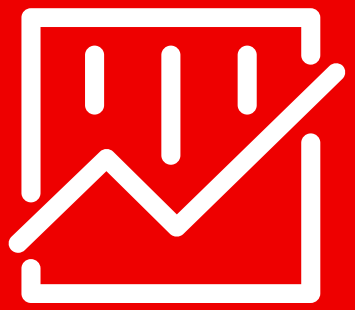

Business intelligence and business analytics both provide insights on past and present data to define future decisions. They differ in that business intelligence is focused on in-the-moment analysis whereas business analytics is a data science focused more on modifying raw data into a meaningful format. The analytics process is what brings business users to a place where they can perform business intelligence, including data modeling, data cleansing, predictive analytics, forecasting, and advanced statistics.

Business analytics is a complex but meaningful process. Data scientists analyze the data for insights and recommendations to take to the business. Data engineers identify, assemble, and manage the right tools into a data pipeline to best enable data scientists. Finally, on the infrastructure side, administrative users work with the infrastructure to provide the basic data services. Along the way are the challenges of integration, storage capacity, and shrinking or flat IT budgets.

Completing this work successfully and providing business users with valuable data for decision making requires several data management techniques:

Machine learning and artificial intelligence

Predictive analytics

Data mining

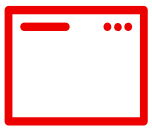

Apache Hadoop

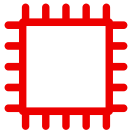

In-memory analytics

Machine learning and artificial intelligence

The advent of machine learning has created one of the most revolutionary technologies in the analysis and understanding of data. Machine learning is a subset of artificial intelligence that gives computers the ability to learn without being explicitly programmed. It automates the process of analytical model building and makes it possible to quickly and automatically produce models that can analyze bigger, more complex data and deliver faster, more accurate results.

Predictive analytics

Perhaps as groundbreaking as machine learning, predictive analytics helps businesses use their data to make predictions about the future. Predictive analytics uses historical data, statistical algorithms, and machine learning techniques to identify the likelihood of future outcomes. With a better understanding of future possibilities, the organization can feel more confident about business decisions and strategies.

Data mining

Data mining is the process of finding anomalies, patterns, and correlations within large data sets to predict outcomes. With the primary goal of uncovering and discovering knowledge, data mining extracts patterns and insights to identify hidden patterns, predict future trends, and analyze behaviors to allow businesses to make decisions from large volumes of data.

"Analysts resolve cases quicker, with much better results. For example, Red Hat Virtualization has eliminated latency for network traffic, and reusing code dramatically reduces our development time. Reports that used to take 48-72 hours can now be produced in a day."

Juan Carlos Biacchi

General Manager, Information Systems and Technology, Dirección Nacional de Migraciones de la República Argentina[4]

[4] Red Hat case study, "Argentina's migration department unifies national security data with Red Hat."

Apache Hadoop

An open source software framework, Hadoop stores large amounts of data and runs applications on hardware clusters. A complete industry of data analysis has evolved from the development of Hadoop. The technology sits at the center of an ecosystem of big data technologies that are used to support advanced analytics including predictive analytics, data mining, and machine learning.

In-memory analytics

In-memory analytics enables users to analyze data from system memory to uncover immediate insights and act on them quickly. Removing data preparation and analytical processing latencies, in-memory analytics enables organizations to stay agile and make better, faster business decisions.

Chapter 5

Database workloads

Database workloads process and store data while enabling data availability and data integrity. Through this process, IT organizes, stores, and shares data to ensure the availability, stability, and timeliness of data for better decision making.

Before data can be mined for insights and value, business users must be able to access this data quickly, accurately, and safely—whether the data lives on-premise, in the cloud, or at the edge.

In fact, the value of a company can be measured by the performance of its data. In this way, the performance of the database can have a significant effect on the organization's ability to use its data to make timely decisions, track important metrics, and even provide products and services to customers.

When creating a high-performance database, enterprise IT must focus on these data qualities:

Availability

Stability

Security

Timeliness

Modernization

Availability

Many businesses rely on data to deliver products and services to their customers. Data availability refers to the accessibility of data, particularly in the disruption of the network. If the business loses access to critical data, business and IT operations could come to a halt, resulting in financial losses and damage to the organization's reputation. Several factors like data redundancy, data loss prevention, and data backup and recovery play a key role in keeping data available and accessible.

Stability

The business-critical nature of database workloads requires a foundation that is stable and reliable, delivering data resiliency and consistency to the organization. A stable platform and operating system provides high performance and "always on" stability and availability in even the most demanding physical and virtual environments.

Security

Protecting the organization's valuable data is a primary consideration for any database deployment. Fines and penalties for data breaches can cause financial and reputational harm. With the right security measures in place, the business can protect its own data and that of its customers. A high-performance database built on a stable platform can help the organization apply a set of standards and practices to protect data from a wide variety of threats.

"Our clients can now access their data at any given point. In the BPO [business process outsourcing] industry, the more visibility you can give your clients, the happier and more productive they are. So this is a game changer for us."

Nicolaas Botes

Executive Director of Enterprise Architecture, TMG Health[5]

[5] Red Hat case study, "TMG Health accelerates to real-time data access for clients."

Timeliness

Data must be timely to deliver the insights and value the business expects. The timeliness of data affects the accuracy of insights used to make decisions. In any type of data analysis, timeliness and accuracy are fundamental prerequisites. Outdated, inaccurate data can compromise the usefulness of the insights it generates.

Modernization

Modern approaches to database management require a focus on emerging technologies that enable advanced analytics. Such technologies include predictive analytics, artificial intelligence, guided analysis, containers, Internet of Things (IoT), and machine learning—all of which create greater value from the organization's data. A high-performing database that allows for rapid querying, multiple types of sensor data, and advanced statistical modeling help generate insights with a significant degree of precision.

In addition to creating a high-performance database, data storage is an important consideration. Data storage has come a long way since the days of disk systems. Relational databases are the most widely used database types. However, with growing volumes and types of data being generated, the enterprise IT organization needs database technologies that store and manage structured and unstructured data. In addition to relational databases, database types include hierarchical databases, network databases, object-oriented databases, and graph databases. The complicated nature of database types requires a stable and secure platform as its foundation.

Gain the power of the Red Hat portfolio for workloads

Red Hat offers a broad portfolio of open source tools that help enterprise IT organizations reduce costs and complexity, use data in decision making, and improve database performance—all with embedded security protections. With Red Hat, your organization can modernize its existing IT while shifting investments to innovation.

Red Hat Enterprise Linux

Red Hat Enterprise Linux provides an intelligent, stable, and security-focused foundation for modern, agile business operations—reducing the friction and cost of cloud deployment while accelerating time to market for critical business workloads. Complete, built-in automation and insights let the organization deliver services faster and more consistently across on-premise and cloud infrastructure.

Red Hat Ansible Automation Platform

Red Hat Ansible® Automation Platform is a powerful tool for the orchestration of enterprise environments, allowing the organization to automate tasks such as configuration management, provisioning, workflow orchestration, application deployment, and life-cycle management. One of the solution's hallmarks is its low barrier to entry with simple, agentless IT automation technology.

Red Hat Satellite

Red Hat Satellite is an infrastructure management tool specifically designed to make it easier to deploy, scale, and manage Red Hat infrastructure across physical, virtual, and cloud environments. The tool gives users the ability to provision, configure, and update systems to keep them running efficiently, with security, and in compliance.

Red Hat Insights

Red Hat Insights provides highly scalable, prescriptive analytics across physical, virtual, container, and private and public cloud environments. It analyzes IT infrastructure against Red Hat's knowledgebase to provide a real-time assessment of risks related to performance, availability, stability, and security.

Red Hat Smart Management

Red Hat Smart Management helps organizations optimize and manage their Red Hat Enterprise Linux environments. It gives the organization the choice to manage the operating system in the cloud or on-premise, all with a single subscription to address key IT challenges around vulnerability, compliance, provisioning, and patching.

Red Hat OpenShift Container Platform

Red Hat OpenShift® Container Platform helps the organization develop, deploy, and manage existing and container-based applications seamlessly across physical, virtual, and public cloud infrastructures. Built on proven, open source technologies, it helps application development and IT operations teams to modernize applications, deliver new services, and accelerate development processes. Integrated with OpenShift Container Platform, Red Hat OpenShift Container Storage provides software-defined storage that runs anywhere OpenShift does, on-premise or in the public cloud.

Red Hat OpenStack Platform

Red Hat OpenStack® Platform is a cloud computing platform that virtualizes resources from industry-standard hardware, organizes those resources into clouds, and manages them so users can access what they need when they need it.

Red Hat Enterprise Linux for Microsoft SQL Server

With one platform from Microsoft SQL and Red Hat, your organization can consolidate database systems to provide developers with fast, reliable, and more secure access to data. It offers fast data performance and runs efficiently across multiple environments, whether on-premise or in the cloud, to meet critical demands.

Red Hat Enterprise Linux for SAP Solutions

Red Hat Enterprise Linux for SAP® Solutions offers a powerful, consistent platform for enterprises to run all of their applications and take advantage of both scale-up and scale-out configurations on a single server.

Red Hat Ceph Storage

Red Hat Ceph® Storage is an open, massively scalable storage solution for modern workloads like cloud infrastructure, data analytics, media repositories, and back up and restore systems.

Red Hat Gluster Storage

Red Hat Gluster® Storage is a software-defined storage platform designed to handle general purpose workloads like backup and archival as well as analytics. Ideal for hyperconvergence, it is cost-efficient and can be deployed on bare metal, virtual, container, and cloud environments.

Red Hat

Contact us

For more information, visit www.redhat.com.

About Red Hat, Inc.

Red Hat is the world's leading provider of enterprise open source software solutions, using a community-powered approach to deliver reliable and high-performing Linux, hybrid cloud, container, and Kubernetes technologies. Red Hat helps customers integrate new and existing IT applications, develop cloud-native applications, standardize on our industry-leading operating system, and automate, secure, and manage complex environments. Award-winning support, training, and consulting services make Red Hat a trusted adviser to the Fortune 500. As a strategic partner to cloud providers, system integrators, application vendors, customers, and open source communities, Red Hat can help organizations prepare for the digital future.

facebook.com/redhatinc
@RedHat
linkedin.com/company/red-hat

NORTH AMERICA
1-888-REDHAT1
www.redhat.com

EUROPE, MIDDLE EAST, AND AFRICA
00800 7334 2835
europe@redhat.com

ASIA PACIFIC
+65 6490 4200
apac@redhat.com

LATIN AMERICA
+54 11 4329 7300
info-latam@redhat.com