

Delivering on the Promise of Big Data



Solution
Partner

Enterprises gain powerful, competitive advantage with scalable and efficient platforms driven by Cisco, Hortonworks, and Red Hat innovation.

What if your company could collect, store, and analyze any amount of data, regardless of its source, format, or volume? What if you never had to throw away potentially valuable data because it is too expensive to store it? What if you could break down silos and derive insights at scale? Today, a number of organizations around the world are already able to use disruptive technologies to handle massive amounts of data and extract intelligence from it. In fact, many of these companies never discard data anymore. They have built scalable, efficient, and flexible systems with integrated and validated solutions from Cisco, Hortonworks, and Red Hat. With Hortonworks and Red Hat running on Cisco® Unified Computing System™ (Cisco UCS®) servers, big data success is now achievable for enterprises of all sizes.

Challenge: Data Platform That Delivers Speed, Performance, Cost Savings

The amount of data companies deal with is growing every day in volume, velocity, and variety. It's difficult enough to store it all, let alone turn it into a competitive advantage. And companies that don't use data strategically face a dubious future. Accenture found in a recent [survey](#) that 79 percent of respondents agree that "companies that do not embrace big data will lose their competitive position and may even face extinction."

Using massive amounts of data to derive new insights is more than ever the key to success. This is easier said than done. With billions of people and devices producing data, the challenge is to find a data platform that delivers outstanding performance, speeds deployment of new resources, and manages those resources cost-effectively.



Big data management and integration with Hortonworks and Red Hat delivers massive scalability, superior efficiency, and dramatically lower total cost of ownership.

This ideal data platform should be equipped to:

- Scale to handle ever-increasing amounts of data without breaking the corporate budget. By working seamlessly with hardware, the platform should expand at low cost with plug-and-play economy and efficiency.
- Work with existing technologies. Companies can't afford to rip and replace their enterprise applications, and they don't need to. A big data platform should be able to integrate with existing investments.
- Simplify management with automation capabilities so that IT personnel can quickly add servers and make updates across the entire application system. This would free staff from time-consuming tasks while increasing efficiency in the data center and reducing management costs.

Combined Solution: Cisco, Hortonworks, and Red Hat

By tightly integrating Hortonworks Data Platform, an open source Apache Hadoop data platform for the enterprise, with the Cisco UCS platform and Red Hat Enterprise Linux, companies can achieve a solution that enables big data success.

Working together, Cisco, Hortonworks, and Red Hat provide a comprehensive system that meets today's data imperatives: supporting an enterprise-grade Hadoop deployment, integrating data across the enterprise, and supporting the deployment of innovative applications.

Big data management and integration with Hortonworks and Red Hat can be combined easily with existing investments, enhancing them and extending ROI without rebuilding a top-to-bottom infrastructure. Going forward, it accommodates new innovations and scales easily, so enterprises can start small and grow their data deployment in a step-by-step fashion.

Enterprises can blend the Hortonworks platform with their computing, networking, and storage resources to run Hadoop clusters. And after the solution is in place, an enterprise can put all its data to optimal use. For example, unstructured data can be collected in Hadoop and integrated with information spread across many applications and systems to derive maximum value.

Powering this integration is Red Hat middleware, which sits in front of multiple data sources and allows them to be treated as a single source. As a result, data can be delivered in the required form, at the right time, to any application or user. With the Red Hat solution, an organization can combine historical data from Hadoop and other

Challenges

- Provision quickly and consistently.
- Scale easily and cost-effectively.
- Support innovative big data applications with flexibility.
- Reduce system management costs.

sources with real-time data stored in the data grid. This data mashup gives enterprises greater insight into operations and lets them rapidly react to changing market conditions.

In addition, Red Hat Gluster Storage offers open-software-defined storage that can be accessed through industry-standard interfaces. Engineers at Red Hat and Hortonworks have developed a plug-in so that the MapReduce programming model can run directly on top of Red Hat Storage. As a result, developers and architects can use it as a central repository without having to move data in and out of HDFS (Hadoop Distributed File System).

The Hortonworks and Red Hat solution also expands easily and at a low cost. Organizations can deploy Cisco UCS servers and build storage clusters on top of these servers instead of purchasing a large and expensive storage appliance. Rather than commit significant resources up front, a company can start with a few Cisco UCS servers running storage software from Red Hat and run Hortonworks Data Platform on top of them. As capacity grows, additional storage clusters can be added for easy scalability.

Since the Cisco UCS platform is a policy-based infrastructure that handles monitoring, management costs are kept to a minimum. The Cisco Unified Fabric network architecture manages everything from a single point, so customers cut expenses.

Cisco, Hortonworks, and Red Hat have conducted extensive testing and validation to provide a stable, robust, and fully tested platform architecture for big data.

Big data is part of the respective DNA at Cisco, Hortonworks, and Red Hat. Because all three companies build on open standards and work with open interfaces, their joint offering integrates easily with any data center and works with virtually any solution.

Business Results: Scalability, Efficiency, Lower Costs

Proving the ROI of your big data project is critical to its overall success. The Cisco approach to big data doesn't require extra resources or head count. Incorporating Hadoop to modernize your data architecture can lead to significant cost savings. Thanks to Hortonworks Data Platform, you can move high volumes of existing data to Hadoop, offload processing workloads, and enrich your data architecture with additional types of data to create new business value.

While data volumes continue to grow considerably, low-value workloads such as ETL (extract, transform, and load) consume ever more processing resources, and new types of data can't easily be captured and put to use. Organizations struggle with escalating costs, increasing complexity, and the challenge of expansion. Hadoop provides a way to address these challenges, moving high data volumes, offloading ETL processes, and enriching existing data architectures with new data for increased value.

The Hortonworks and Red Hat solution delivers massive scalability. Even as the volume and variety of their data explodes, enterprises have the capacity to store and manage it all.

Solution

- Ensures big data success
- Integrates with existing investments and accommodates new technologies
- Scales rapidly and at low cost, with minimal management demands

The fabric-based architecture of the Cisco UCS platform, which easily integrates storage, server, and network, delivers great efficiency. This infrastructure enables businesses to manage up to 10,000 Cisco UCS servers as if they were a single pool of resources to support very large data clusters. Cisco Unified Fabric technology dramatically lowers costs by reducing the number of switches and cables a company requires.

Enterprises today must be able to capture intelligence from both data at rest in the data center and real-time data at the edge of the network. The broad portfolio of Cisco UCS systems provides the flexibility to process data where it makes the most sense.

Cisco UCS Integrated Infrastructure for Big Data enables powerful management-automation capabilities. Cisco UCS Manager abstracts all configuration, identity, and I/O connectivity information into a Cisco UCS service profile that can be applied to other servers. This intelligent programmability enables IT staff to rapidly and consistently deploy new servers, restore servers, and update releases across an entire network of Cisco UCS servers. Additional Cisco UCS servers can be deployed 84 percent faster and with fewer steps than in traditional environments, while automation frees staff from tedious, time-consuming chores. The entire data center becomes more efficient, resulting in a 51 percent reduction in management costs.

Meanwhile, Cisco SingleConnect technology—an easy, intelligent, and efficient way to connect and manage computing in the data center—offers a single 10 Gigabit Ethernet unified fabric cable that supports LAN, SAN, and management traffic, while Cisco UCS Director Express for Big Data provides a single management pane to easily deploy, adapt, and scale Hadoop big data infrastructure. As a result, new nodes can be deployed in minutes instead of hours or days. Additionally, Cisco Security Management Suite supplies a framework of next-generation security management tools designed for management and policy administration of the Cisco Self-Defending Network.

Together Cisco, Hortonworks, and Red Hat deliver the enterprise-grade performance, availability, and security that are critical to any enterprise big data deployment.

Next Steps

To learn more about the many ways Cisco, Hortonworks, and Red Hat big data solutions can benefit your organization, visit http://www.cisco.com/c/dam/en/us/td/docs/unified_computing/ucs/UCS_CVDs/Cisco_UCS_Integrated_Infrastructure_for_Bigdata_with_Hortonworks.pdf.

Business Results

- Cisco UCS platform requires no extra resources, and Hadoop allows for dramatically lower overall storage costs.
- Cisco UCS servers deploy 84 percent faster, and automated management frees staff time.
- More efficient data center results in a 51 percent reduction in management costs.