

Gain insight using next-generation applications with IBM Spectrum Conductor

Efficiently analyze, access and protect data with an integrated application- and data-optimized platform



Highlights

- Accelerate time to insight up to 60 percent by abstracting IT complexity with policy-based, workload-aware resource management¹
- Ensure maximum availability and security of shared services with intelligent application and data lifecycle management
- Choose the optimal deployment model for your organization whether it's on premises, in the cloud or an integrated infrastructure solution

Organizations of all types are realizing data is a highly valuable business asset, and that extracting full value from that data provides a critical competitive advantage. However, these organizations face growing challenges. Businesses are accumulating not just data pools or data lakes, but oceans of data. The volume of data pouring in makes it difficult for businesses to take advantage of the insights it can provide.

The new generation of cloud-native applications required to analyze this data increasingly rely on open source frameworks such as Hadoop and Apache Spark. These frameworks have the potential to deliver significant benefits. However, to mitigate the risks and complexities associated with adopting this technology, organizations typically deploy them in isolated clusters, where each is dedicated to a specific application or department (Figure 1). This resulting cluster sprawl decreases infrastructure utilization while increasing time-to-results and costs.

Moreover, many new applications are composed of microservices that run in lightweight container environments such as Docker for greater efficiency. Containers provide integrated technology that lets development and IT operations teams build, ship and run distributed applications anywhere. But once container applications are running, resources must be effectively managed for optimum performance and asset utilization—especially as the container environment scales from a few small clusters to hundreds or even thousands.




Figure 1. Integrating new-generation applications and open source frameworks into existing infrastructure can lead to siloed, underutilized infrastructure and data replication.

Welcome to a new way of thinking about infrastructure

IBM Spectrum Conductor enables organizations to accelerate business insights from all their data by leveraging the most current scale-out applications, open source frameworks, in-memory analytics, NoSQL databases, cloud-native application architectures and container environments. Traditional IT server configurations, hypervisor environments and storage silos do not work well for these modern scale-out applications and frameworks because traditional systems are not optimized for distributed computing. They are optimized for traditional workloads such as Microsoft Exchange, SQL databases and data warehouses as well as virtual desktops.

Today, line-of-business managers and application owners are looking to adopt or improve the efficiency of distributed applications and microservices. They are asking their IT architects to find ways of adopting these technologies as quickly as possible—cost-efficiently and without impacting service levels. IT managers are building data oceans for analytics to distill insight from vast amounts of structured and unstructured data, and embracing software defined infrastructure that supports such an environment.

Whether an organization starts from an application- or data-centric viewpoint, IBM Spectrum Conductor provides dynamic sharing of resources, automated management of all resources across all applications, and an architecture for infrastructure management, storage of unstructured data and data lifecycle management.

In short, IBM Spectrum Conductor offers IT managers and architects a new way of thinking about their infrastructure—one that enables IT to unleash the potential of a new generation of applications.

Organizations need an environment where they can easily analyze, access and protect their data to accelerate business results while controlling costs. IBM® Spectrum Conductor meets those challenges by enabling organizations to transform their infrastructure into a tightly integrated application- and data-optimized platform.

Applying high-performance computing techniques in business environments

Technologies, tools and techniques once thought to be only for high-performance computing and analytics workloads are now being applied to business workloads to manage and analyze data. Organizations can use these capabilities to gain the improved performance and capacity needed for big data analytics, as well as to make data accessible to the right applications precisely when needed.

These advancements add value for businesses by integrating intelligence for both the compute and storage aspects of the infrastructure, and by delivering an intelligent software layer that optimizes the placement of data and application workloads across the infrastructure. IBM has delivered these proven capabilities to clients for years and has enhanced them in its latest software defined infrastructure solutions. For example, Citibank has used IBM software defined infrastructure to gain a 100x performance improvement.²

IBM Spectrum Conductor is based on this proven technology and offers superior performance at dramatically lower costs so organizations can create the innovations that take their business to a higher level.

Delivering an integrated platform for dealing with data growth

With data volumes continuing to grow, organizations want a complete solution that answers three questions:

- How do we accelerate results by rapidly and efficiently adopting next-generation analytics applications and frameworks?
- Are there ways to access and share big data at the speed of business while reducing infrastructure cost?
- Can we cost-efficiently protect and manage data and applications throughout their lifecycle?

IBM Spectrum Conductor enables organizations to answer these questions in the most cost-efficient way possible through the following solution capabilities:

- **Analyze:** Workload-, resource- and data-aware management software increases utilization of existing resources and speeds analysis. An organization can share resources and data across applications, users and lines of business, as well as easily capitalize on technologies such as Apache Spark and Docker.
- **Access:** Policy-driven workload placement, combined with global shared data access, lets organizations share data and resources across distributed teams or data centers (Figure 2). Features include data migration, automatic failover and seamless file system recovery. Advanced, latency-aware file management, routing and caching help optimize placement and performance.
- **Protect:** For data and application lifecycle management and protection, IBM Spectrum Conductor has enterprise-class features such as encryption to protect data integrity and end-to-end checksum to catch errors. A multi-tenant integrated application and data fabric enables organizations to eliminate resource and data silos while protecting applications and data throughout their lifecycle.




Figure 2. Optimized data placement and global access provided by IBM Spectrum Conductor enable collaboration and resource sharing.

Providing multidimensional scalability

IBM Spectrum Conductor offers multidimensional, independent scaling of storage capacity, compute resources, application versions and more (Figure 3). Organizations can future-proof their data center, avoiding rip-and-replace upgrades because they can easily and independently add resources as needed—for example, to accommodate new workloads, rising data volumes or more users.

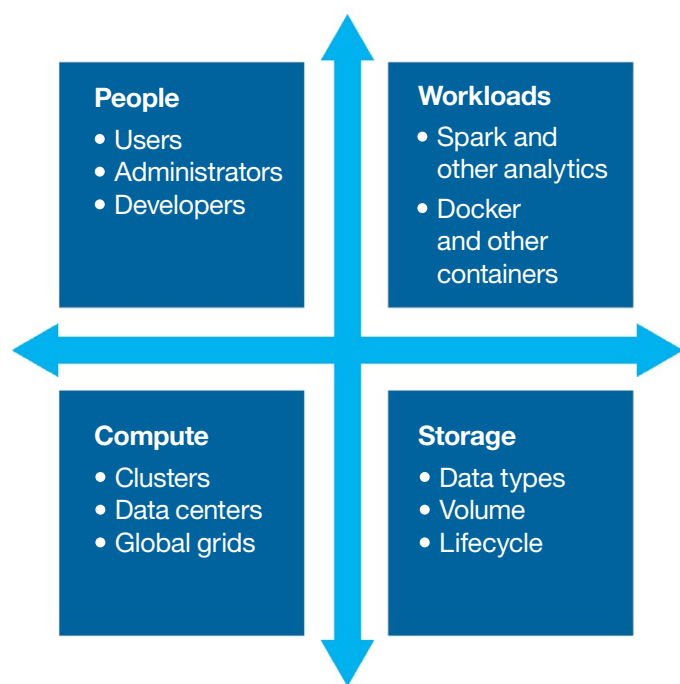


Figure 3. IBM Spectrum Conductor delivers multidimensional scalability including storage, compute, people and workloads.

Boosting performance, simplifying management and reducing costs

With IBM Spectrum Conductor, organizations benefit in multiple ways. They can:

- Eliminate silos with a multi-tenant, integrated application and data platform
- Simplify administration by deploying single-pane management and monitoring across a scale-out, distributed infrastructure
- Improve data availability with global, shared access to ensure data is available exactly when needed
- Obtain faster business insights from data while reducing both capital and operational expenses
- Future-proof the data center with multidimensional scaling—including independent scaling of compute and storage infrastructure
- Reduce storage costs and improve performance with automated multitier, hybrid cloud storage including flash, disk and archive

Workload and resource management capabilities in IBM Spectrum Conductor help minimize wait time so each application receives the resources and data it needs as efficiently as possible. This accelerates results and enables organizations to better utilize their existing infrastructure and defer additional capital investment.

IBM Spectrum Conductor is simple to deploy and reduces operational costs through management efficiencies, decreasing administrative overhead for monitoring, alerting and reporting. Organizations can view and monitor activity throughout the stack—from the data to the server and up to the application—for a single, correlated view of activity. Many routine administrative tasks are automated.

Scale with confidence across the organization

IBM Spectrum Conductor is based on software defined infrastructure technology proven in some of the world's most demanding environments ([learn more](#)). Here is a quick look at the numbers:

- Production customers with over 100,000 cores
- Thousands of nodes deployed
- Hundreds of applications on a single cluster
- Over 1 billion tasks per day
- Multiple sessions running over 1 million tasks per session

The solution is also fully resilient, with built-in high availability.

Leveraging new technologies including Spark and Docker

IBM Spectrum Conductor offers the flexibility of deployment that is essential in today's changing IT environment.

Organizations are witnessing fast-paced evolution in open source frameworks for big data analytics such as Hadoop MapReduce and Spark. The new generation of big data analytics and cognitive computing is often unable to tolerate the overhead of a hypervisor and must run on bare-metal operating systems. And IT is increasingly turning to lightweight container environments such as Docker to package application software.

IBM Spectrum Conductor is designed to meet these evolving needs. It speeds deployment and simplifies management across physical and containerized resources while adding enterprise-grade capabilities such as built-in high availability and optimized utilization with a global resource manager.

For example, IBM Spectrum Conductor provides automated deployment and management of distributed Docker-based application services and data dependencies, with end-to-end monitoring of the application, container and distributed infrastructure. By running multiple container-based services on shared infrastructure, IBM Spectrum Conductor enables individual applications to take full advantage of available resources to meet service-level objectives. The unified management interface allows administrators to see where application services are running and how resources are being allocated.

IBM Spectrum Conductor with Spark

Apache Spark offers compelling performance advantages as an open source, big data analytics framework. However, implementing Spark poses significant challenges, including investment in new expertise, tools and workflows. Setting up ad hoc Spark clusters can lead to inefficient resource use, as well as management and security challenges. [IBM Spectrum Conductor with Spark](#) helps address those issues. It integrates a Spark distribution with resource, infrastructure and data lifecycle management, streamlining the creation of enterprise-grade, multi-tenant Spark environments. To help manage fast-moving Spark lifecycles, IBM Spectrum Conductor with Spark supports simultaneous running of multiple instances and versions of Spark.

Accelerating your success

IBM Spectrum Conductor is the next evolution of IBM software defined infrastructure technology—a highly scalable fabric that enables organizations to analyze, access and protect their data with optimum efficiency. As more waves of data continue to roll in, IBM Spectrum Conductor helps reduce storage and management costs, accelerate time to insight and enhance your organization's competitive edge.

Why IBM?

IBM Spectrum Computing offers a comprehensive portfolio of software defined infrastructure solutions designed to help your organization deliver IT services in the most efficient way possible, optimizing resource utilization to speed time to results and reduce costs. These offerings help maximize the potential of your infrastructure to accelerate your analytics, high-performance computing (HPC), Hadoop, Apache Spark and cloud-native applications at any scale. The core value of the portfolio is simplifying simulations and analysis to help uncover insights and get higher-quality products to market faster.

For more information

To learn more about IBM Spectrum Conductor, contact your IBM representative or IBM Business Partner, or visit: ibm.com/systems/spectrum-computing/products/conductor/

Additionally, IBM Global Financing provides numerous payment options to help you acquire the technology you need to grow your business. We provide full lifecycle management of IT products and services, from acquisition to disposition. For more information, visit: ibm.com/financing



© Copyright IBM Corporation 2016

IBM Systems
Route 100
Somers, NY 10589

Produced in the United States of America
June 2016

IBM, the IBM logo, and ibm.com are trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at “Copyright and trademark information” at ibm.com/legal/copytrade.shtml

This document is current as of the initial date of publication and may be changed by IBM at any time. Not all offerings are available in every country in which IBM operates.

The performance data and client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions.

THE INFORMATION IN THIS DOCUMENT IS PROVIDED “AS IS” WITHOUT ANY WARRANTY, EXPRESS OR IMPLIED, INCLUDING WITHOUT ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND ANY WARRANTY OR CONDITION OF NON-INFRINGEMENT. IBM products are warranted according to the terms and conditions of the agreements under which they are provided.

¹ [STAC Report: Spark Resource Managers, Phase 1](#) (March 28, 2016)

² Greg Lavender, “Optimize Your Data Economics,” IBM Edge 2014, www.youtube.com/watch?v=CfO5zl3l6oc



Please Recycle
