White Paper

# Cisco and Greenplum Partner to Deliver High-Performance Hadoop Reference Configurations

September 2012

## Contents

# Cisco and Greenplum Partner to Deliver High-Performance Hadoop Reference Configurations

White Paper
September 2012

## Highlights

**Optimized for Performance**

- Greenplum and Cisco deliver an integrated Hadoop solution specifically engineered to handle the most demanding Hadoop workloads.

**Cisco UCS Creates a Flexible Appliance Platform**

- The Cisco Unified Computing System™ (Cisco UCS®) provides a flexible, high-performance platform that can be optimized and easily scaled for any size of Hadoop cluster.

**Ease of Management**

- Cisco UCS Manager provides unified, embedded management of all computing, networking, and storage-access resources.

**Choice of Configurations**

- The solution provides a choice of Cisco UCS configurations, letting organizations select performance and capacity as their needs dictate

**Enterprise-Class Support and Services**

- The Greenplum MR on Cisco UCS Reference Configurations combine the support and services of two of the world's largest technology companies.

Greenplum MR on Cisco UCS provides companies with an integrated Hadoop solution that delivers advanced performance, full data protection, no single point of failure, and improved data-access features that can expedite the implementation of big data analytics environments.

## Next-Generation Hadoop Solution

The worldwide leader in data center networking, and now a leading competitor in the server market, Cisco is partnering with Greenplum to provide a best-in-class big data solutions that meet a range of needs. The Greenplum MR on Cisco UCS® Reference Configurations deliver integrated, end-to-end software and hardware infrastructure that accelerates big data initiatives with a choice of performance and capacity. The combination of world-leading Cisco Unified Computing System™ (Cisco UCS) and Greenplum MR enables companies to significantly reduce time-to-value and the operating expenses associated with Apache Hadoop implementations.

### Greenplum MR: Hadoop Reengineered

Greenplum MR, based on the MapR M5 Distribution, is an implementation of the Apache Hadoop stack that enables near-real-time collection and organization of high volumes of structured, semistructured, and unstructured data distributed across a cluster of servers. Greenplum MR provides direct data input and output to the cluster with MapR Direct Access Network File System (NFS), offers real-time analytics, and is the first distribution to provide true high availability at all levels. Greenplum MR introduces the concept of logical volumes to Hadoop: a means of grouping data and applying policy consistently across an entire data set. Greenplum MR provides hardware status information and control with the MapR Control System, a comprehensive user interface that includes a heatmap that displays the health of the entire cluster at a glance.

Cisco UCS and Greenplum MR can help businesses manage many different big data scenarios. The examples in Table 1 show how the Greenplum MR on Cisco UCS Reference Configurations can accelerate big data initiatives.

**Table 1**. Sample Use Cases for Cisco UCS and Greenplum MR

| Scenario | Cisco Greenplum MR Reference Configuration Capabilities |
|---|---|
| Content management | Collect and store unstructured and semi-structured data in a fault-resilient, scalable data store that can be organized and sorted for indexing and analysis. |
| Batch processing unstructured data | Batch-process large quantities of unstructured and semi-structured data: for example, data warehousing extract, transform, and load (ETL) processing. |
| Medium-term data archive | Archive data (medium-term, 12 to 36 months) from an enterprise data warehouse (EDW) database management system (DBMS) to increase the length of time that data is retained or to meet data retention and compliance policies. |
| Integration with data warehouse | Transfer data stored in Hadoop to and from a separate DBMS for advanced analytics. |
| Customer risk analysis | Perform a comprehensive data assessment of customer-side risk, based on activity and behavior across products and accounts. |
| Personalization and asset management | Create and model investor strategy and goals based on market data, individual asset characteristics, and reports entered into an online recommendation system. |
| Trade analytics | Analyze historical volume and trading data for individual stock symbols, variable cost of trades, and allocation of expenses. |
| Credit scoring | Update credit-scoring models using cross-functional transaction data and recent outcomes, to respond to changes such as the collapse of bubble markets. Sweep recent credit history to build transactional and temporal models. |
| Retailer compromise | Prevent or catch fraud resulting from a breach of retailer cards or accounts by monitoring, modeling, and analyzing high volumes of transaction data and extracting features and patterns. |
| Miscategorized credit card fraud | Reduce false positives and prevent miscategorization of legitimate transactions as fraud, using high volumes of data to build good models. |
| Next-generation credit card fraud | Perform daily cross-sectional analysis of portfolio using transaction similarities to find accounts that are being cultivated for eventual fraud, using common application elements, temporal patterns, vendors, and transaction amounts to detect similar accounts before the fraud is perpetrated. |
| Customer retention | Combine transactional customer contact information and social network data to perform attrition modeling to learn social and transaction markers for attrition and retention. |
| Sentiment analysis (opinion mining) and bankruptcy | Find better indicators to predict bankruptcy among existing customers using sentiment analysis from social networking, responding quickly before the warning horizon. |

### Greenplum MR

- Complete enterprise-class solution
- Elimination of the common problems experienced with HDFS
- Direct access through NFS
- High availability through JobTracker enhancements and a no-NameNode architecture
- From two to five times faster performance than other Hadoop distributions
- Advanced management for clusters of all sizes
- Robust data protection features, including snapshots and intercluster mirroring
- Comprehensive network of enterprise business intelligence tools

As shown in Figure 1, the major components of Greenplum MR provide a Hadoop solution that is easy, dependable, and fast, with components that include:

- **Advanced storage services:** A replacement for the Hadoop Distributed File System (HDFS), MapR advanced storage services provide multidimensional scalability and accelerate MapReduce performance. The services allow random read and write operations while automatically compressing data in real time.
- **MapR Heatmap:** MapR Heatmap provides visibility, access, and tools that offer insight into the state of the cluster. Graphical and programmatic interfaces are designed to scale with the largest clusters.
- **MapR Control System:** MapR Control System provides real-time monitoring of the cluster health, including alarms to notify you of conditions that need to be corrected. Alarms can also be configured to trigger email notifications.
- **MapReduce:** Part of the Apache Hadoop framework, MapReduce simplifies the creation of applications that process large amounts of unstructured and structured data in parallel. Underlying hardware failures are handled transparently for user applications, providing a reliable and fault-tolerant capability.
- **Hive:** Hive is a data warehouse system for Hadoop that facilitates data summarization, impromptu queries, and analysis of large data sets. This SQL-like interface increases the compression of stored data for improved storage resource utilization without affecting access speed.
- **Pig:** Pig is a high-level procedural language for processing data sets in parallel using the Hadoop MapReduce platform. Its intuitive syntax simplifies the

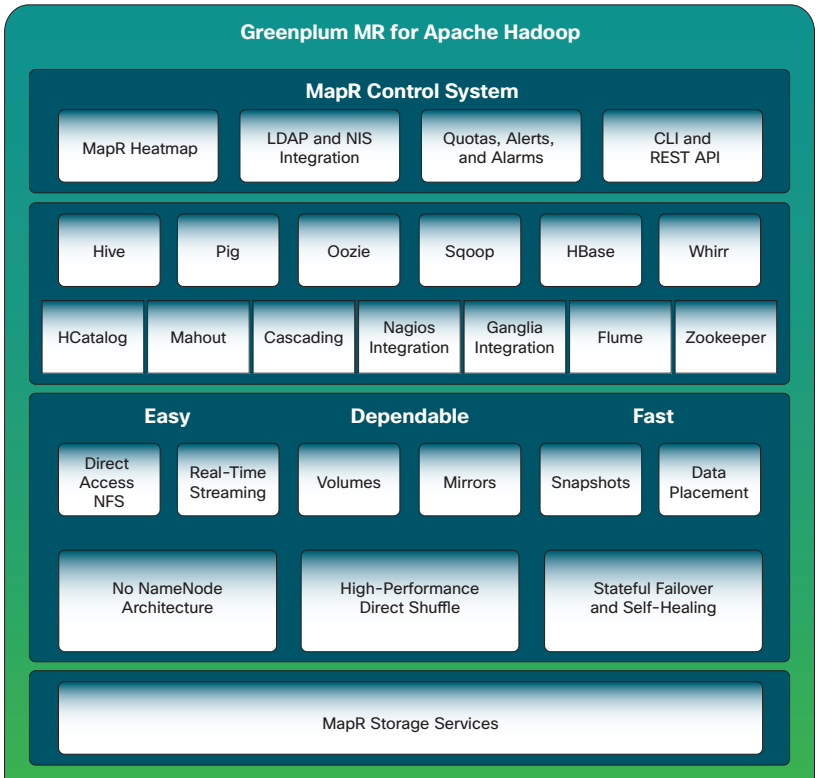


**Figure 1.** Greenplum MR Architecture and Components

development of MapReduce jobs, providing an alternative programming language to Java.

- **HBase:** HBase is the distributed, versioned, column-oriented database that delivers random, real-time, read-write access to big data.
- **ZooKeeper:** ZooKeeper is a highly available system for coordinating distributed processes. Applications use ZooKeeper to store and mediate updates to important configuration information.

## Cisco UCS: The Exclusive Platform for Greenplum MR

Validated through an extensive testing and development process at Greenplum and Cisco, Cisco UCS is the exclusive hardware platform for Greenplum MR. Cisco UCS innovations combine industry-standard, x86-architecture servers with networking and storage access into a single converged system (Figure 2). The system is entirely programmable using unified, model-based management to simplify and accelerate the deployment of enterprise-class applications and services running in bare-metal, virtualized, and cloud-computing environments.

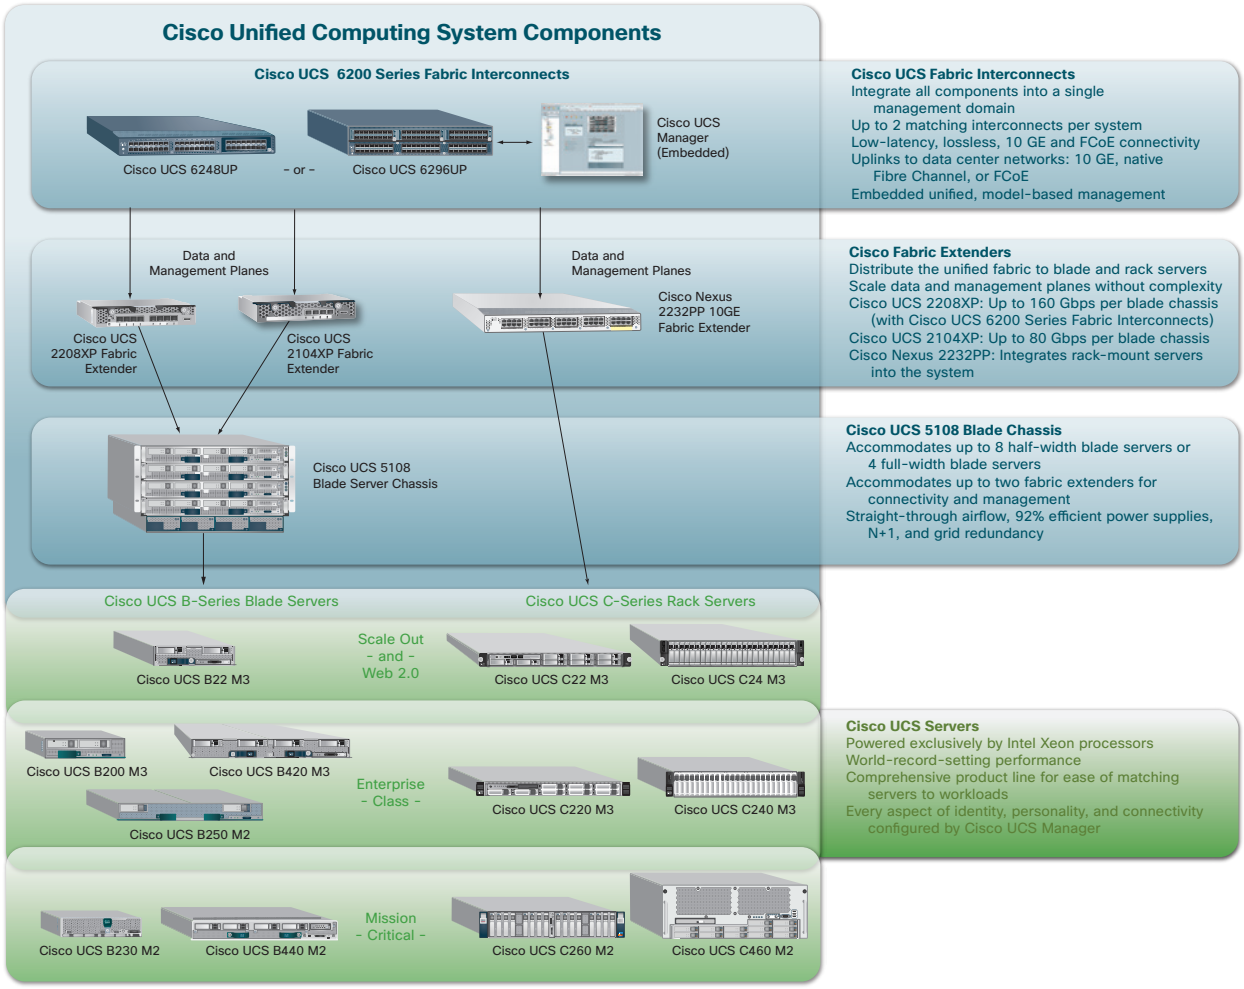


**Figure 2.** Cisco UCS Components

### Cisco UCS

The Cisco Unified Computing System delivers a radical simplification of traditional architecture with the first self-aware, self-integrating, converged system that automates system configuration in a reproducible, scalable manner

- More than 60 world records on critical benchmarks
- The benefits of centralized computing, through a single point of management, delivered to massive scale-out applications
- Self-aware and self-integrating system
- Automatic server provisioning through association of models with system resources
- Standards-based, high-bandwidth, low-latency, lossless Ethernet network

Cisco UCS helps organizations gain more than efficiency: it helps them become more effective through technologies that foster simplicity rather than complexity. The result is a flexible, agile, and high-performance platform that reduces operating costs with increased uptime through automation and enables more rapid return on investment (ROI).

## Reference Configurations

The Greenplum MR on Cisco UCS Reference Configurations are based on Cisco's big data common platform architecture (CPA), a highly scalable architecture designed to meet variety of scale-out application demands with transparent data integration and management integration capabilities using the following components:

- **Cisco UCS 6200 Series Fabric Interconnects:** The Cisco UCS 6200 Series Fabric Interconnects are a core part of Cisco UCS, providing both network connectivity and management capabilities across Cisco UCS 5100 Series Blade Server Chassis and Cisco UCS C-Series Rack Servers. Deployed in redundant pairs, the fabric Interconnects offer line-rate, low-latency, lossless 10 Gigabit Ethernet connectivity and unified management with Cisco UCS Manager in a highly available management domain.
- **Cisco UCS 2200 Series Fabric Extenders:** Cisco UCS 2200 Series Fabric Extenders behave like remote line cards for a parent switch and provide a highly scalable and extremely cost-effective unified server-access platform.
- **Cisco UCS rack servers:** Specific models are used to support the base, high-performance, and high-capacity configurations:
  - **Cisco UCS C210 M2 General-Purpose Rack Server:** Cisco UCS C210 M2 servers are general-purpose 2-socket platforms based on the Intel® Xeon® processor 5600 series. These servers support up to 192 GB of main memory and 16 internal front-accessible, hot-swappable, Small Form Factor (SFF) disk drives, with a choice of one or two RAID controllers for data performance and protection.
  - **Cisco UCS C240 M3 Rack Server:** Cisco UCS C240 M3 Servers are designed for both performance and expandability over a wide range of storage-intensive infrastructure workloads. Each server provides sockets for up to two processors from the Intel Xeon processor E5-2600 product family and up to 768 GB of main memory. Up to 24 SFF or 12 Large Form Factor (LFF) disk drives are supported, along with four Gigabit Ethernet LAN-on-motherboard (LOM) ports.
  - **Cisco UCS virtual interface cards (VICs):** Unique to Cisco, Cisco UCS VICs incorporate next-generation converged network adapter (CNA) technology from Cisco and offer dual10-Gbps ports designed for use with Cisco UCS C-Series Rack Servers. Optimized for virtualized networking, these cards deliver high performance and bandwidth utilization and support up to 256 virtual devices.

- **Cisco UCS Manager:** Cisco UCS Manager resides in the Cisco UCS 6200 Series Fabric Interconnects. It makes the system self-aware and self-integrating, managing all the system components as a single logical entity. Cisco UCS Manager can be accessed through an intuitive GUI, a command-line interface (CLI), or an XML API. Cisco UCS Manager uses service profiles to define the personality, configuration, and connectivity of all resources within Cisco UCS, radically simplifying provisioning of resources so that the process takes minutes instead of days. This simplification allows IT departments to shift their focus from constant maintenance to strategic business initiatives. It also provides the most streamlined, simplified approach commercially available today to firmware updating for all server components.

Organizations deploying Hadoop clusters have different needs depending on the nature of their computational applications and storage requirements. Cisco understands these important distinctions and has structured its Hadoop reference configurations to accommodate a range of diverse needs

- **Base reference configuration:** Cisco's base reference configuration is designed for organizations that require balanced computing and storage capacity. The reference configuration uses Cisco UCS C210 M2 servers, each with two Intel® Xeon® processors X5675, 96 GB of memory, and 16 1-TB 7,200-rpm SATA disk drives (Figure 3).

- 2 Intel Xeon Processors X5675
- 96 GB Memory
- 1 LSI 6G MegaRAID 9261-8i Card
- 16 1-TB 200 RPM SFF SATA Disk Drives
- Redundant Hot-Swappable Power Supplies and Fans
- 1 Cisco UCS P81E VIC (2x 10 Gbps)
- Embedded Cisco IMC (2x 1 Gbps)
- 2 Integrated Gigabit Ethernet Ports
- Red Hat Enterprise Linux Server Standard

**Figure 3.** Cisco UCS B210 M2 Cluster Node for Base Configuration

- **High-performance reference configuration:** Many organizations need to optimize computing and memory performance in their Hadoop clusters, so Cisco's high-performance reference configuration uses Cisco UCS C240 M3 servers, each with two Intel Xeon processors E5-2665, 256 GB of memory, and 24 1-TB 7,200-rpm SFF SATA disk drives (Figure 4).
- **High-capacity reference configuration:**. For organizations that require abundant storage capacity for the Hadoop cluster, the high-capacity reference configuration uses Cisco UCS C240 M3 servers, each with two Intel Xeon processors E5-2640, 128 GB of memory, and 12 3-TB 7,200-rpm SAS disk drives (Figure 4).

- 2 Intel Xeon Processors E5 Family
- Cisco UCS VIC 1225 (2x 10 Gbps)
- Embedded Cisco IMC (2x 1 Gbps)
- LSI MegaRAID SAS 9226CV-8i Card
- 4 Integrated Gigabit Ethernet Ports
- Red Hat Enterprise Linux Server Standard
- Redundant Hot-Swappable Power Supplies

High-Performance Configuration (shown):

- 256 GB Memory
- 24 1-TB SATA 7200 RPM SFF Disk Drive

High-Capacity Configuration (not shown):

- 128 GB Memory
- 12 3-TB SAS 7200 RPM LFF Disk Drives

**Figure 4.** Cisco UCS C240 M3 Cluster Node for High Performance and High Capacity

Table 2 summarizes the capabilities of the three reference configurations.

**Table 2**. Base, High-Performance, and High-Capacity Reference Configurations.

| | Component | Base Configuration | High-Performance Configuration | High Capacity Configuration |
|---|---|---|---|---|
| Server-level capacity | Cisco UCS servers | Cisco UCS C210 M2 | Cisco UCS C240 M3 | Cisco UCS C240 M3 |
| | Processor | Intel Xeon X5675 | Intel Xeon M5-2665 | Intel Xeon M5-2640 |
| | Memory | 96 GB | 256 GB | 128 GB |
| | Storage | 16 1-TB SFF 7.2K-RPM SATA | 24 1-TB SFF 7.2K-RPM SATA | 12 3-TB LFF 7.2K-RPM SAS |
| Rack-level capacity (16 servers) | Processor cores and threads | 192 cores and 384 threads | 256 cores and 512 threads | 192 cores and 384 threads |
| | Memory | 1536 GB | 4096 GB | 2048 GB |
| | Typical user storage capacity (3-way replication, compressed) | 320 terabytes (TB) | 480 TB | 720 TB |
| | I/O bandwidth | 21 GBps | 32 GBps | 17 GBps |

| | Component | Base Configuration | High-Performance Configuration | High Capacity Configuration |
|---|---|---|---|---|
| Cisco UCS Domain-level capacity (10 racks) | Processor cores and threads | 1920 cores and 3840 threads | 2560 cores and 5120 threads | 1920 cores and 3840 threads |
| | Typical user storage capacity (3-way replication, compressed) | 3.2 petabytes (PB) | 4.8 PB | 7.2 PB |
| | I/O bandwidth | 210 GBps | 320 GBps | 170 GBps |

The Greenplum MR on Cisco UCS Reference Configurations come in single- and multiple-rack form factors. The single-rack configuration consists of two fully redundant Cisco UCS 6248UP 48-Port Fabric Interconnects (for up to five racks) or Cisco UCS 6296UP 96-port Fabric Interconnects (for up to 10 racks) along with two Cisco Nexus® 2232PP 10GE Fabric Extenders (Figure 5). Each node in the configuration connects to the unified fabric through two active-active 10-Gbps links using a Cisco UCS VIC (for data traffic) and Cisco Integrated Management Controller (IMC; for management traffic). Multiple-rack configurations include the components for a single rack and two Cisco Nexus 2232PP fabric extenders for every additional rack.

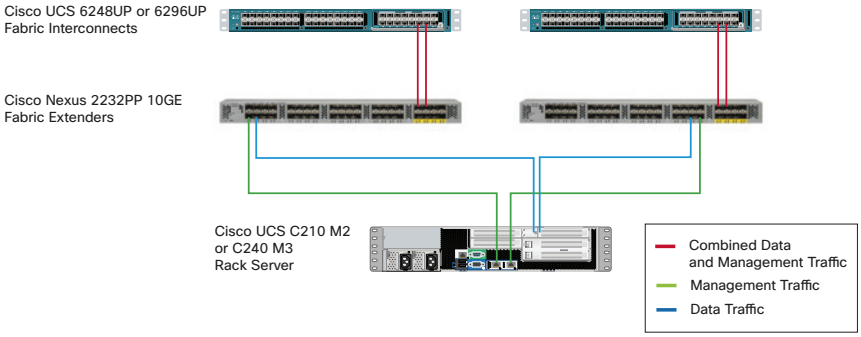


**Figure 5.** Cisco UCS Fabric Architecture

Figure 6 shows the components of the single-rack and multiple-rack configurations. Table 3 provides general guidelines for the number of instances of each service to run in a solution.

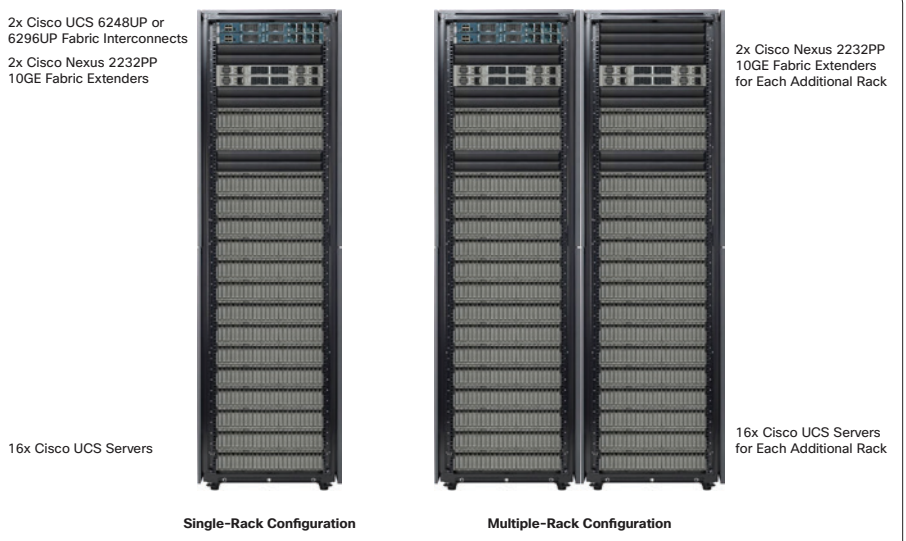


**Figure 6.** Base, High-Performance, and High-Capacity Configurations Can Be Built as Single- and Multiple-Rack Configurations

**Table 3**. Node Recommendations for Greenplum MR Services

| Service | Number of Nodes (per Rack) |
|---|---|
| Container location database | 1 to 3 |
| FileServer | Most or all nodes |
| HBase Master | 1 to 3 |
| HBase RegionServer | Varies |
| JobTracker | 1 to 3 |
| NFS | Varies |
| TaskTracker | Most or all nodes |
| WebServer | 1 or more |
| Zookeeper | 1, 3, 5, or a higher odd number |

## Excellence from Cisco and Greenplum

Hadoop implementations can present a number of challenges to enterprise environments. Many of these challenges arise from the dichotomy between the introduction of innovative new technology and the enterprise-class performance, reliability, and support demanded by mission-critical systems. The collaboration between Cisco and Greenplum is specifically designed to provide a solution to these challenges. The joint Cisco and Greenplum solution delivers all the characteristics expected of a fully integrated solution, including radically simplified deployment and management, high availability, excellent performance, exceptional scalability, and enterprise-class service and support.

### Complete Big Data Analysis Solution

The comprehensive solution from Cisco and Greenplum helps organizations deploy big data solutions quickly, with validated configurations that scale easily and predictably, as demand dictates. The Greenplum MR on Cisco UCS Reference Configurations provide an end-to-end solution that has been tested and validated and that enables enterprise customers to accelerate big data initiatives.

### Designed for High Availability and Reliability

Every component in the Cisco and Greenplum MR Hadoop solution is fully redundant. The Greenplum MR architecture provides JobTracker high availability (HA) and a no-NameNode architeture to prevent lost jobs from causing time-consuming restarts or failover incidents. Combining the core networking capabilities of Cisco, this solution can be extended to include remote mirroring to help ensure data reliability by synchronizing a copy of the cluster's data at a remote site so that data analysis can continue in the event of a disaster. Locally, snapshots protect data from application errors or accidental deletion. Snapshots also enable easy data recovery to a specific point in time by simply copying a file or directory from the snapshot location to the desired destination directory.

### High Performance and Exceptional Scalability

Cisco UCS unified fabric architecture provides fully redundant, highly scalable lossless 10-Gbps unified fabric connectivity for big data traffic. Powered by the latest Intel Xeon processor, the Cisco and Greenplum solution delivers best-in-class performance and internal storage capacity to gain at least two times faster performance than Apache Hadoop. The Greenplum MR on Cisco UCS Reference Configurations can easily scale to support a large number of nodes when required by business demands. The advanced management capabilities of Cisco UCS radically simplify this process with a single point of management that spans all nodes in the cluster.

### Simplified Management

Hadoop implementations tend to involve very large numbers of servers. In traditional environments, it can be challenging to manage these large numbers of servers effectively. Cisco UCS Manager delivers unified, model-based management that applies personality and configures server, network, and storage connectivity resources, making it as easy to deploy hundreds of servers as it is to deploy a single server. Additionally, Cisco UCS Manager can perform system maintenance activities such as firmware update operations across the entire cluster as a single operation. To ease the monitoring of these large clusters, the Greenplum MR control system raises alarms and sends notifications to alert IT personnel about cluster health and the status of services.

### Coexistence with Enterprise Applications

In addition to introducing a Hadoop deployment, organizations need ways to transfer data transparently between their enterprise applications and Hadoop. This solution can connect, across the same management plane, to other Cisco UCS deployments running enterprise applications, thereby radically simplifying data center management and connectivity.

### Rapid Deployment

Deployment of large numbers of servers can take time. Systems need to be racked, networked, configured, and provisioned before they can be put into use. Cisco UCS Manager uses a model-based approach to provision servers by applying a desired configuration to the physical infrastructure quickly, accurately, and automatically. The capability to create consistent configurations improves business agility and eliminates a major source of errors that can cause downtime.

### Enterprise Service and Support

Enterprises want know that the vendors providing a solution have the expertise to help them quickly proceed through the design, deployment, and testing of strategic big data initiatives. Businesses also need to have confidence that if a critical system fails, they will be able to get timely and competent support. The Greenplum MR on Cisco UCS Reference Configurations bring together world-class service and support from long-time collaborators Cisco and Greenplum.

## For More Information

For complete details about Cisco UCS, please visit http://www.cisco.com/go/ucs.

For more information about Greenplum MR on Cisco UCS, please visit: http://www.cisco.com/go/greenplum.

Cisco and Greenplum Partner to Deliver High-Performance Hadoop Reference Configurations
September 2012


**Americas Headquarters**
Cisco Systems, Inc.
San Jose, CA

**Asia Pacific Headquarters**
Cisco Systems (USA) Pte. Ltd.
Singapore

**Europe Headquarters**
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at **www.cisco.com/go/offices**.





LE-35101-04 09/12