

Powerful Duo: MapR Big Data Analytics with Cisco ACI Network Switches

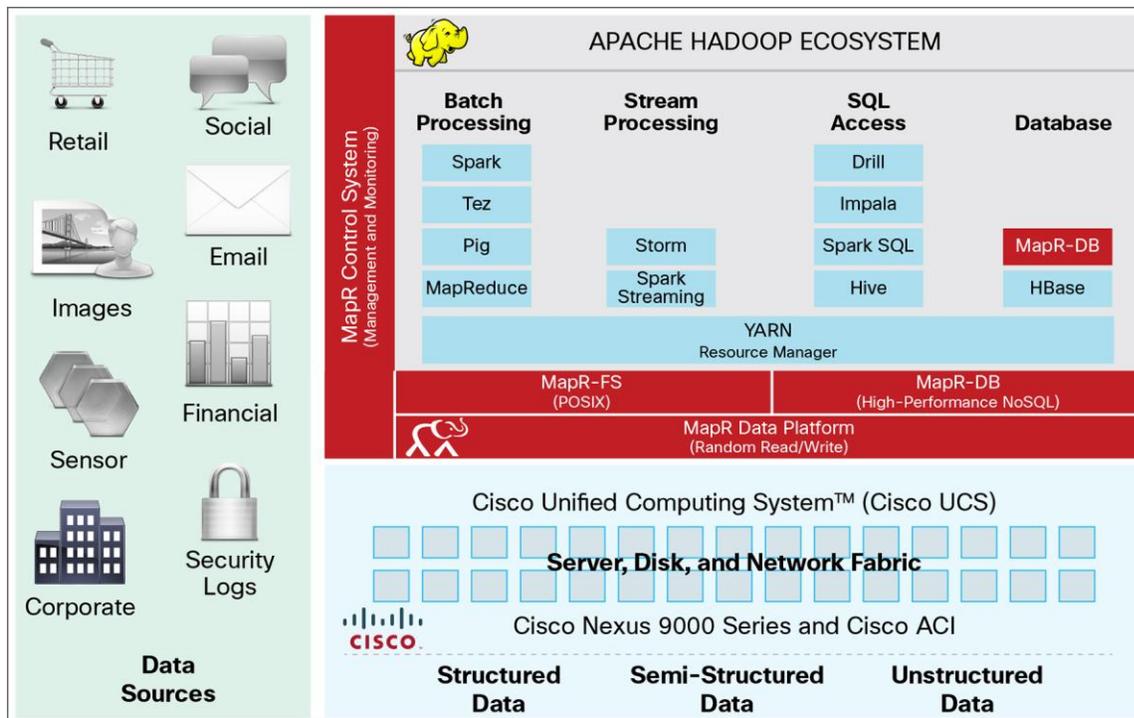
Introduction

For companies that want to quickly gain insights into or opportunities from big data - the dramatic volume growth in corporate and user-generated data, from nearly unlimited sources - the first step is choosing the best IT infrastructure.

MapR and Cisco: Enterprise-Class Production-Ready Hadoop Infrastructure

As a technology leader for Hadoop, MapR provides an enterprise-class, high-performance big data solution that can be quickly developed and is easy to administer. With significant investments in architectural innovation, the MapR distribution delivers more than a dozen tested and validated Hadoop software modules over a fortified data platform (Figure 1).

Figure 1: Apache Hadoop Ecosystem



The joint solution of MapR Hadoop paired with Cisco® Application Centric Infrastructure (ACI), Cisco Unified Computing System™ (Cisco UCS®) and Cisco Nexus 9000 Series switches working together, can instantly and cost-effectively scale capacity and deliver exceptional performance for the growing demands of big data processing, analytics, and storage workflows. For larger clusters and mixed workloads, ACI uses intelligent, policy-based flowlet switching and packet prioritization to deliver:

- Throughput on demand
- Leading-edge load balancing across the Hadoop cluster
- Agile, automated configuration of the cluster topology

Hadoop Has Changed the Data Center

Adoption of Apache Hadoop for big data workloads has tremendously increased the amount of data that enterprises store in data centers. Hadoop's promise of inexpensive and scalable storage and its highly scalable computational capabilities have changed the IT industry. Organizations can scale, with relative ease, from a few nodes to a few hundred nodes.

As the number of nodes increases, so does the workload burden on the network fabric that interconnects all the nodes. To avoid bottlenecks, fabric bandwidth and throughput are critical to helping ensure that all network pipes remain clear to facilitate data movement and data and analytic processing. As a result, fast growing scale-out requirements are pushing data centers toward 10- and 40-Gbps network access and aggregation layers.

At the same time, big data clusters are also evolving. The single-process batch-and-store jobs of the past have given way to multiprocessing, in-memory databases. Hadoop became popular so quickly because it allowed organizations to run a job in minutes instead of requiring days as with traditional approaches. Now organizations are asking whether these same jobs can be run in seconds. They are also asking whether growing workloads can still be completed quickly, and whether larger clusters can be set up easily.

Features: Technology Highlights

- New-generation data architecture with single infrastructure for data storage and processing
- High-performance and very large-scale potential
- Flexible data structure support: structured, semi structured, and unstructured
- Dynamic load balancing
- Innovative TCP flowlet-based switching
- Near-real-time database and supporting infrastructure
- Lower total cost of ownership (TCO) compared to traditional data warehouse approaches

Network Fabric Challenges for Big Data Implementations

Organizations need to look closely at traditional network approaches to massive big data workloads to determine whether they really deliver the value for their size. Traditional approaches do not really help with modern big data workloads because packet interference and bottlenecks can grow exponentially. What organizations need are innovative solutions to intelligently manage provisioning, data flow, visibility, and instrumentation. Modern data centers require networking solutions with these properties:

- Congestion awareness across the network fabric
- Maximum throughput utilization not limited by hashing algorithms
- Dynamic path determination that avoids congestion for higher-priority workflows
- Real-time awareness and distinction between small and large workflows
- Programmable capabilities
- Single point of management

Within a cluster interconnect fabric, multiple links are available to funnel traffic. Multiple links for traffic can operate either on a per-flow basis or a per-packet basis.

- Per-flow switching decisions are robust and do not need many changes at the server or application layer. However, because the path of a flow can't be changed after initialization, per-flow switching may not provide the best performance when the potential for path and link congestion is considered.
- Per-packet switching decisions can help achieve maximum throughput. However, this approach can lead to packet reordering, which may have a negative impact on application performance throughput.

Benefits

- Ingest data from a variety of data sources
- Deliver service-level agreements (SLAs) with confidence
- Maintain better control of costs as data growth increases dramatically
- Create flexibility to prioritize big data and Hadoop workloads
- Optimize performance to address fabric congestion
- Establish a data-based platform for the future

Cisco ACI Flowlet Switching for Optimal Performance Throughput

Unlike per-flow or per-packet switching, Cisco ACI fabrics use a novel approach defined as flowlet-based switching. Typical TCP flows often have gaps between packets. Cisco designed the ACI fabric to use these gaps and divide a single flow into a number of flowlets, which are smaller portions of the TCP flow. Flowlets then become bursts of packets (from a single flow) routed independently.

To achieve performance optimization, the intelligence of ACI determines whether the time required to split a flow and switch flowlets across separate paths is less than the time required to switch the original flow intact but with large gaps. If the time is less, then independent flowlets are switched onto different paths to travel from point A to point B. At the same time, while still making dynamic decisions about switching, the ACI fabric avoids packet reordering.

Thus, flowlet switching, with fabricwide congestion awareness, helps overcome network bandwidth utilization limits commonly seen with traditional (Equal-Cost Multipath (ECMP) based) multilink network designs, which typically use hashing algorithms to determine link paths (Figures 2 and 3).

Figure 2: Without Path-Awareness Link Use

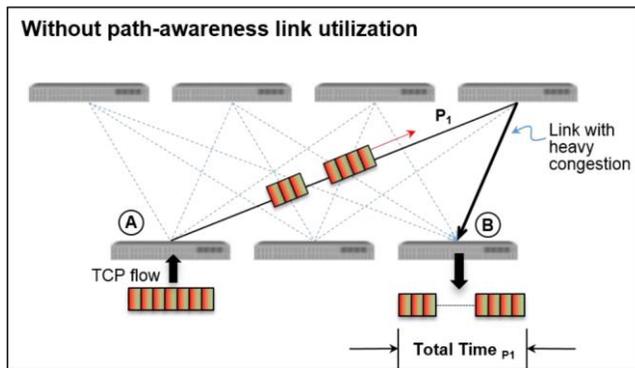
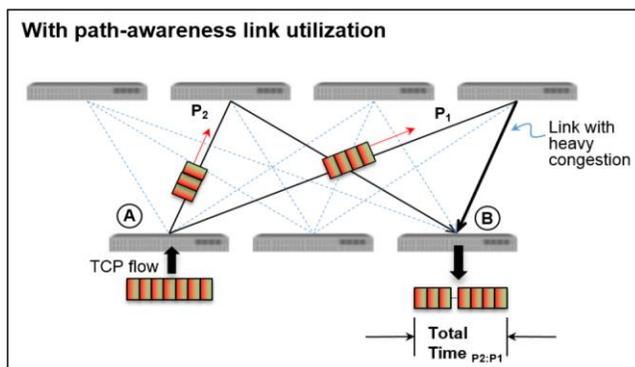


Figure 3: With Path Awareness Link Use



Dynamic Load Balancing

To achieve load balancing, Cisco ACI uses real-time path congestion metrics. Two dynamic load-balancing (DLB) modes are available, applied according to the amount of gap required to detect the start of a new flowlet:

- Aggressive DLB
 - Uses a relatively small inter-flowlet gaps
 - Provides very good load-balancing performance because a high number of rebalancing opportunities are available
 - Small chance of occasional packet reordering
 - Overall performance for normal TCP traffic increases
- Conservative DLB
 - Uses large inter-flowlet gaps so that packet reordering is avoided

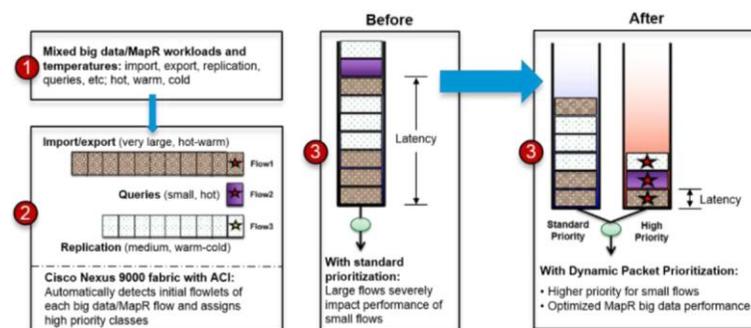
Packet Prioritization Helps Ensure Low-Latency Processing of Important Data Queries

Big data workflows typically are characterized by large to extremely large data sets. However, when you consider the entire data workload environment - from data ingestion, to data protection, to processing of MapReduce jobs, to data analysis - the data mix is a wider cross-section. This cross-section includes small and medium-sized data workloads. Workloads may also range across those with high, medium, and low database processing urgency.

With traditional fabric interconnects, small and urgent data workloads, such as database queries, may suffer processing latency delays because larger data sets are being sent across the fabric ahead of them. This approach presents a challenge for instances in which database queries require near-real-time results.

Cisco Nexus® 9000 Series Switches with Cisco ACI increase performance by prioritizing small workloads for processing, resulting in lower-latency performance throughput (Figure 4).

Figure 4: Improving Latency and Throughput



With ACI capabilities, the result is faster throughput for mixed MapR cluster data workloads, data sets, and data urgency levels. Latency-sensitive operations are prioritized over bulk transfers, such as file-system replication, or batch analytics.

Conclusion

The DLB and packet prioritization capabilities of Cisco ACI complement the big data analytics and storage of MapR-based infrastructure. You can optimize performance throughout all layers of the joint solution. From data ingestion to data analytics, with the MapR and Cisco combined solution you can deliver with confidence a range of possibilities to meet the needs of business executives, managers, and users and data scientists.

For More Information

<https://www.mapr.com/partners/partner/cisco-delivering-advanced-performance-hadoop-workloads>



© 2015 Cisco and/or its affiliates. All rights reserved. Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.com/go/trademarks. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)

[MapR](#) delivers on the promise of Hadoop with a proven, enterprise-grade platform that supports a broad set of mission-critical and real-time production uses. MapR brings unprecedented dependability, ease-of-use and world-record speed to Hadoop, NoSQL, database and streaming applications in one unified distribution for Hadoop.

C22-734763-00 05/15