Cisco and Big Data

What Is the Value of Big Data in Government?

Big Data refers to data sets that are too large to manage and analyze with traditional data management tools and techniques in a reasonable timeframe. These large data sets are often unstructured, containing document files, log files, events or message streams, and images.

The value of Big Data in government is yielding insights to support missions ranging from citizen services and regulatory compliance to defense and intelligence analysis.

What Problems Does It Help Solve?

Big Data use cases in government include:

- Cybersecurity and threat analysis
- Image processing
- Defense and anti-terror programs
- Budgetary control and management
- Asset control and assessment
- Environment monitoring
- Telemetry analysis

Cisco and Big Data

One of the most popular software frameworks for Big Data analytics is Apache Hadoop™, an open-source framework that enables applications to scale up to thousands of nodes and petabytes of data.

Government organizations can build a cost-effective, high-performance, and resilient Hadoop platform using Cisco® Unified Computing System™ (UCS) Servers, Cisco Nexus® 5000 Series Switches, and Cisco Nexus 2000 Series Fabric Extenders (Figure 1). This architecture supports the dynamic workloads and traffic patterns typical of Big Data analysis.

In particular, Cisco Nexus Series Switches deliver the scalability and low latency required for Hadoop's three data-processing phases:

- Writing data as it is streamed or bulk-delivered into the Hadoop Distributed File System (HDFS)
- Processing data in parallel on large clusters, using MapReduce
- Reading data from HDFS for user viewing or further analysis

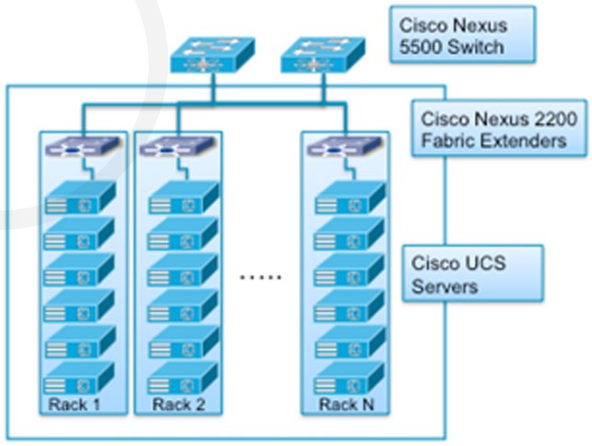

Figure 1- Hadoop Cluster Based on Cisco UCS Servers and Cisco Nexus Series Switches

What Are the Benefits of Cisco UCS and Cisco Nexus as a Hadoop Platform?

High Performance

Cisco UCS servers hold nine world records for performance on leading industry benchmarks, including for two-socket servers typically used for Hadoop*. Cisco UCS C210 Rack-Mount Servers topped all competitors on the SPEC®jbb2005, SPECint-rate_base2006, and SPECfp-rate_base2006 benchmarks.

* *www.cisco.com/en/US/prod/ps10265/at_work_promo.html#~industry_benchmarks*. Data from April 2011.

Cisco and Big Data

At-A-Glance

Factors contributing to record-setting performance of Cisco UCS servers include:

- **Low-latency switching:** Cisco Nexus Series Switches provide 10-Gbps, nonblocking, line rate Ethernet cut-through switching ports.
- **High memory capacity:** Hadoop performs best when it can process data in memory, and Cisco UCS servers provide industry-leading memory density. Cisco UCS C260 M2 Rack-Mount Servers provide up to 1 TB of memory.
- **High disk capacity:** In the Hadoop framework, every compute node stores data. Each Cisco UCS C-Series Server has its own internal small form-factor (SFF), SAS, or SATA disk drives, and can also connect to external Fibre Channel SAN storage, iSCSI-blocked storage, or network-attached storage (NAS). Cisco UCS C260 M2 Servers can be configured with up to 16 internal SAS, SATA, or SSD drives, providing up to 10 TB of capacity.

Investment Protection

The Cisco platform provides a migration path to 10-Gbps Ethernet, because agencies can start with their existing 1-Gbps Ethernet servers and later upgrade to Cisco UCS servers. The Cisco Nexus 5000 Series Switch already supports 10-Gbps Ethernet, so all you need are 10-Gbps Cisco Nexus 2000 Series Fabric Interconnects.

The Cisco platform for Big Data processing also scales easily as you add data. A single Cisco Nexus 5548 Switch can support up to 512 10-Gb Ethernet servers or 768-Gb Ethernet servers, through up to 16 Cisco Nexus 2000 Series Fabric Extenders.

Moreover, Cisco has pretested and validated Hadoop cluster designs that can be used for Big Data build-outs.

Why Cisco?

Government entities that want to analyze Big Data can come to Cisco for everything they need to build a Hadoop platform: low-latency switching fabric, industry-standard servers based on Intel x86 processors, and professional services. The proven performance of the Cisco UCS reduces risk. With 5400 public- and private-sector customers, the Cisco UCS has the second largest market share for Intel x86 processors in North America, and the third largest in the world*.

For More Information

To read more about Cisco UCS, visit:
www.cisco.com/go/unifiedcomputing

For SPEC benchmarks on Cisco UCS servers, visit:
www.cisco.com/en/US/prod/ps10265/at_work_promo.html#~industry_benchmarks

To read more about Cisco Nexus Series Switches, visit:
www.cisco.com/go/nexus

* *IDC Worldwide Quarterly Server Tracker*, Q1 2011, May 2011.