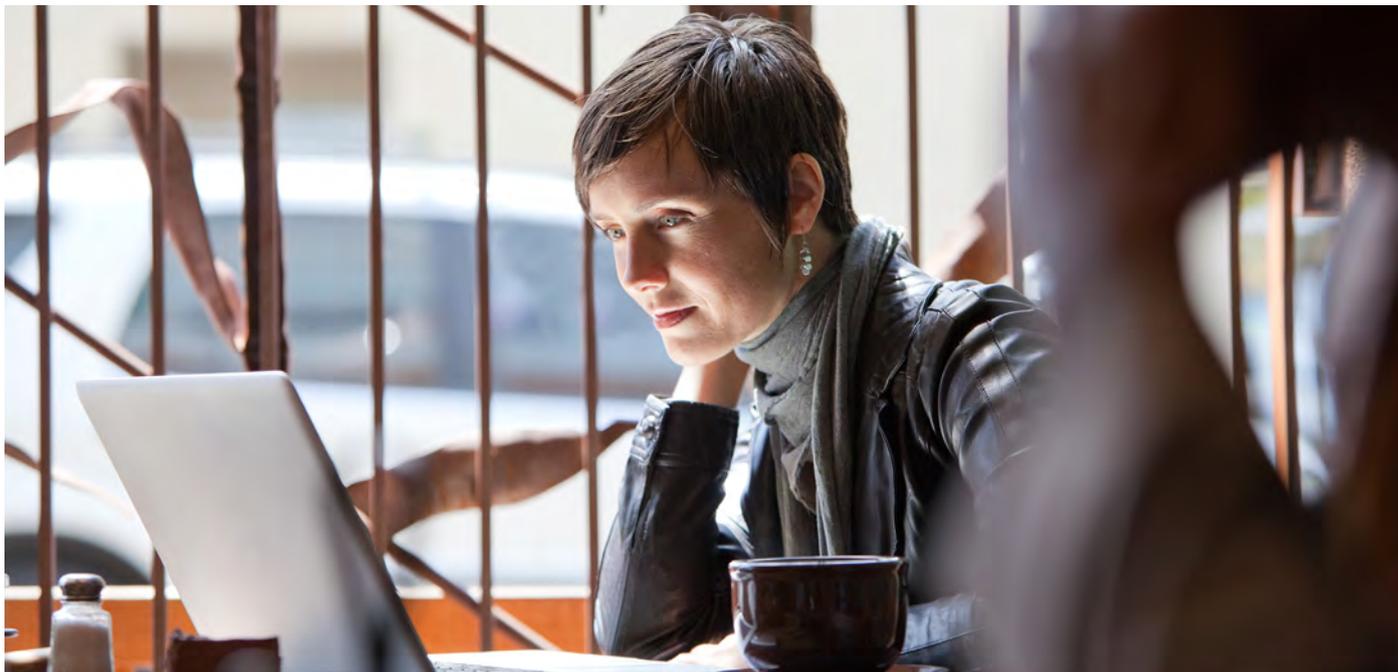


Online Ad Leader Builds Platform for Global Growth



Cisco Nexus helps Criteo deal with big data and optimizes Hadoop Analytics as foundation for worldwide growth

EXECUTIVE SUMMARY

Customer Name: Criteo

Industry: Digital Advertising

Location: France

Number of Employees: Over 700

Challenge

- Improve business processes and enhance growth capabilities
- Increase organizational resilience and ability to handle big data

Solution

- Cisco Nexus with FabricPath

Results

- Hadoop cluster throughput more than doubled from 40 to 100Gbps
- Data center scalability improved from 26 to virtually unlimited number of racks
- High availability assured through FabricPath

Challenge

Criteo has built a global business by predicting what ads consumers might like to see and click through. The ability to target individual ad impressions to individual users with personalised creative banners is now a reality and it's being done at scale. Dedicated exclusively to online media, the company uses complex analytics with a highly accurate product recommendation engine to display personalized advertisements.

By tagging visitors to its partners' websites with anonymous browser cookies, the company can monitor when an individual visits a website belonging to one of 5000 publishers in more than 30 countries around the world. At no point does Criteo collect personal data, such as name or address but, in order for its business model to work, the company has to know the browsing navigation to show the best and most relevant ads to consumers. And it has to be able to work very rapidly.

In fact, Criteo understands and predicts user intent to drive people to a website: from the moment when a visitor is arriving at a given site, it decides on whether it's worth showing an advertisement there, and creates and delivers a completely personalized ad on the page that the person is viewing. All this happens in the time it takes for a browser to load a page. From a technical perspective, this process requires massive computing power. Criteo has gained speed with a worldwide network of data centers to reduce web application latency.

The company has data centers in Europe, in the United States, and in Japan. Within each data center, computing power is provided by horizontal build-outs of commodity servers arranged in multiple clusters. Some of these servers are dedicated to handling incoming traffic and delivering highly-personalized ad banners to end users. They include web servers, caches, and real time database applications. Other clusters are devoted to the data analytics involved in creating these banners.



“We more than doubled the internal throughput of the cluster from 40Gbps to 100Gbps. Nexus completely unleashed the performance of the system.”

Julien Simon
Vice President of Engineering
Criteo

Criteo decided to upgrade its data processing platform from SQL to Apache Hadoop. With its parallel processing capabilities, Hadoop is a very chatty application. It splits processing tasks across multiple servers to improve performance and, in the process, creates a heavy switching load. Soon after implementation of Hadoop, Criteo realized that its data center switches could not cope.

“Our nodes all had 1Gbps interfaces, but we were not seeing that level of performance,” says Julien Simon, Vice President of Engineering at Criteo, “and we realized our switches were dropping packets. This was to be expected, as Hadoop is very heavy on the LAN, but at that point we understood we needed to upgrade the switches.”

Solution

Criteo compared switches from a number of vendors, but none offered the performance needed. However, Criteo was considering a move to Cisco Nexus® switches for its production network and, when the company tried them on the Hadoop cluster, the packet loss problem was solved immediately. “Cisco Nexus worked really well,” says Simon. “That meant we were able to deploy the same access switches for our web production and Hadoop clusters. It’s a good thing to have the same equipment across the board.”

The Criteo infrastructure team embarked on a nine-month project to upgrade all of the company’s 150 data center switches to Cisco® Nexus with FabricPath. The team chose Nexus 2248TP-E Series Fabric Extenders with extended buffers and Nexus 5548 Series Switches to make up the data center access layer. For Level 2 and Level 3 virtual data center aggregation and interconnection, Nexus 7000 Chassis were adopted. Cisco ASR 9006 Series Aggregation Services Routers were used for edge routing and transits.

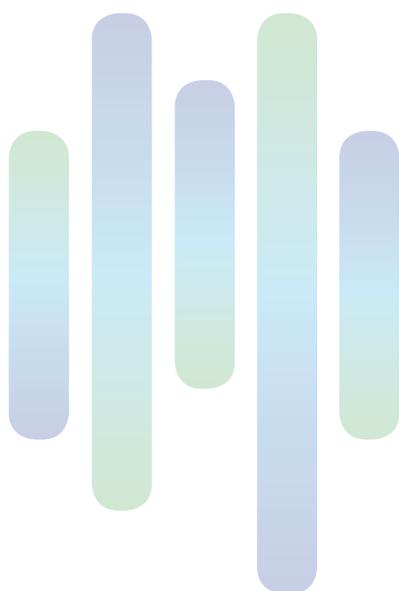
Results

Criteo currently has one of the largest Hadoop deployments in Europe. Moving to Nexus has made a significant impact on its performance. “We more than doubled the internal throughput of the cluster from 40Gbps to 100Gbps,” Simon says. “Nexus completely unleashed the performance of the system.”

This improvement extends to the rest of the data center, despite the fact that traffic loads have greatly increased since the infrastructure was installed. Incoming Criteo traffic grew significantly in the last year alone, while its real time applications were receiving over 30 billion HTTP requests per day, with the Hadoop system receiving 20 terabytes of new information every 24 hours.

Criteo has over two petabytes of analytics data and forecasts a significant increase by the end of the year. Yet Cisco Nexus is helping ensure the output from the analytics cluster is quickly fed back into real time Criteo operations. “That feedback loop needs to be ultra-fast,” says Simon. “We need to shorten it as much as possible. It’s about speed, speed, speed, end-to-end. When you’re delivering banners, every millisecond counts.”

At the same time, the Cisco Nexus deployment has enhanced data center scalability. “With the previous architecture, the number of servers we could deploy on a single aggregation switch was restricted,” says Simon. “Worse, we were limited to 26 racks per data center.” Using the Nexus FabricPath architecture, Criteo now has almost limitless scalability; new servers can be added-in anywhere and configured via the Nexus 5548 and 2248TP-E access architecture, which can exceed 1000 server ports per Nexus 5548. “One of the key things with Nexus is the ability to scale horizontally,” says Simon. “Now there’s no theoretical limit to the number of racks we can have. We can keep growing without any blocking points.”



“We are presently building a new Hadoop cluster, generating still more network traffic and more terabytes of data to move around. Cisco is a key partner for us. The quality of our infrastructure is what makes our performance so strong. It needs to be rock solid.”

Julien Simon
Vice President of Engineering
Criteo

Criteo regards Cisco FabricPath as a good match for its needs, because the company uses commodity servers to spread the load and cannot be sure all servers working on an application are co-located. With Nexus that arrangement is not a problem. At the same time, being able to deploy servers more easily means routine work can be delegated to system administration staff. “We deployed 1500 new servers in 2012, and the growth is still climbing, but the number of people in system administration has not grown in the same proportion,” says Simon. “So we’re able to do more with the same resources.”

Lastly, the Cisco Nexus platform means Criteo benefits from improved business resilience. “High availability is a priority for us,” Simon says. “One of the benefits of FabricPath is that there is no Spanning Tree, no blocking links, so any interruption is going to be circumvented very quickly.”

Next Steps

Criteo has standardized on Cisco Nexus for data center switching, which provides the company with a solid foundation for growth. “We are presently building a new Hadoop cluster, generating still more network traffic and more terabytes of data to move around,” says Simon. “Cisco is a key partner for us. The quality of our infrastructure is what makes our performance so strong. It needs to be rock solid.”

For More Information

To learn more about the Cisco architectures and solutions featured in this case study, please go to:

www.cisco.com/go/nexus

www.cisco.com/go/fabricpath

Product List

Routing and Switching

- Cisco Nexus 2248TP-E Series Fabric Extenders
- Cisco Nexus 5548 Series Switches
- Cisco Nexus 7000 Series Switches
- Cisco ASR 9006 Series Aggregation Services Routers



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.

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)