EMULEX

SOLUTIONS BRIEF

CONNECTIVITY

Boost Hadoop Performance with Emulex OneConnect™ 10GbE Network Adapters

Increased performance, streamlined management and industry-leading reliability

At a Glance

It is estimated that we collectively create 2.5 quintillion (25^{18}) bytes of data on a daily basis. To put that into perspective, approximately 90% of the data in the world today has been created within the last two years. The sources driving the creation of this data—consumer participation on the web, social media, financial transactions and medical records, to name a few—are growing at an equally impressive rate. Data generated by these disparate sources, and the rate at which it is created, is the definition of "Big Data".

Before these large datasets can be mined for valuable and actionable information, they must be imported into a Hadoop cluster at high speeds. The performance capabilities of today's commodity servers have quickly shifted the performance bottleneck to the legacy 1GbE network. Migrating from a 1GbE to a 10GbE network using an Emulex OneConnect OCe11102 adapter increases Hadoop's data tansfer rate by four times.

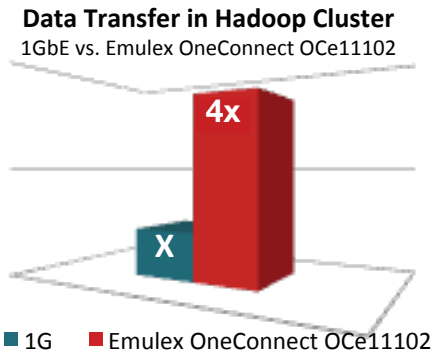

Products

- Emulex OneConnect OCe11102
- Emulex OneCommand™ Manager

Solution Benefits

- 4x data transfer rate
- Accelerated transaction response time
- Simplified management
- Broad Linux support
- Lower capital and operating expenses

Hadoop on 10GbE

As Big Data gets even bigger, the demand to increase processing power, storage capacity, storage performance and network bandwidth grows even greater. To meet this demand, more racks and data nodes may be added to Hadoop clusters. But this solution is neither economical nor efficient. It requires more space, increases power consumption and adds to management and maintenance overheads, while ignoring the fact that slow 1GbE interconnects continue to handicap Hadoop's node to node and overall input/output speed and cluster performance.

With its bullet-proof reliability and server manufacturer preference, Emulex has emerged as the leader in 10GbE -based port market share. Emulex delivers the benefits of 10GbE with its OneConnect OCe11102 10GbE Network Adapter, built to meet the challenges of Big Data and the need for more efficient and faster Hadoop network interconnects.

Higher Performance

Increased capacity, greater multi-core processor performance and faster storage access speed, available with today's commodity servers deployed in Hadoop clusters, cannot fully be utilized while 1GbE network pipes remain congested. Offering 300% Hadoop performance improvement over 1GbE networks, Emulex OneConnect OCe11102 is the necessary and obvious solution to meet the growing demands of Big Data.

The Emulex third-generation OneConnect OCe11102 network adapter provides superior performance, measured by throughput, CPU efficiency and transactions per second with competitive latency, and is clearly the best choice for Hadoop.

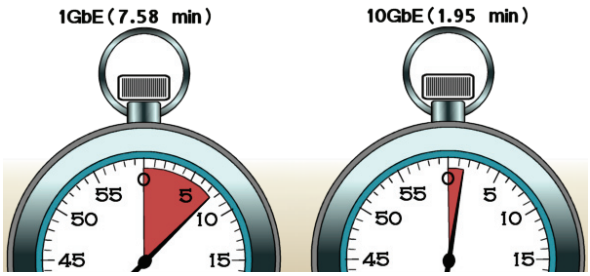

HDFS 'put' operations are four times faster using Emulex OneConnect adapters.

Boost Hadoop Performance with Emulex OneConnect 10GbE Network Adapters

Lower Capital and Operational Expenses

Providing 300% better throughput than 1GbE, each of the two OCe11102 ports eliminates the need for four 1GbE network adapters. Consequently, fewer adapters and switch ports are required to handle the same workload, and less money is spent on power and cooling.

Operational cost reduction efforts extend beyond reducing the number of physical adapters. With increased network throughput, data ingestion and data replication, operations take less time. As a result, more CPU cycles become available for data analysis, improving overall response time and reducing the required number of cluster nodes. Reducing cluster size and transaction response time results in significant cost reductions associated with implementing, operating and maintaining a Hadoop cluster. Emulex has a world-class engineering team that has been perfecting I/O performance at the CPU and server level, thanks to innovative design and management tools.

Extensive Linux Support

Linux is the operating system of choice for Hadoop. Emulex, a key contributor to the open development community, works closely with market leading Linux distribution suppliers to ensure the latest drivers are delivered and supported.

Offering the most advanced Linux driver architecture and management applications available as a CLI and GUI, Emulex provides support for all leading Linux distributions including Ubuntu, CentOS, Red Hat Enterprise Linux and FreeBSD.®

Management Leadership with OneCommand

The Emulex OneConnect OCe11102 10GbE Network Adapter continues Emulex's tradition of delivering reliable, high performance, scalable and manageable connectivity. Emulex OneCommand Manager provides administrators with an easy-to-use tool for network discovery, diagnostics and network-wide driver and firmware updates. Available in both a GUI and programmable CLI, OneCommand Manager is a very popular tool. In a recent survey of nearly 1000 Emulex end users, 75% of them used OneCommand Manager.

Successful deployment and management of hyperscale infrastructures such as Hadoop and High Performance Computing (HPC) clusters require maximizing efficiency while reducing cost. OneCommand Manager enables customers to efficiently automate and dynamically provision their Ethernet environments to reduce management overhead, maintenance overhead and operational cost.

Emulex is the Market Leader

Emulex leads the market in 10GbE port market share (source: Crehan Research and the Dell'Oro Group). With over 100 patents, as well as research projects with leading universities, such as Carnegie-Mellon and University of California at Irvine, and propelled by superior products, Emulex has established a strong market position in both storage and local area networks, delivering network adapters that offer enterprise-proven reliability complemented by exceptional performance.

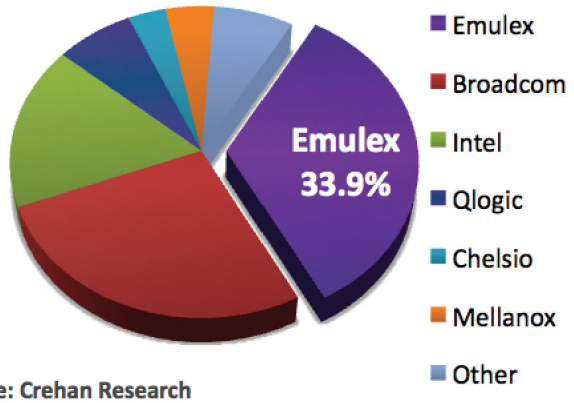

Source: Crehan Research

Emulex OneConnect adapters have been chosen and qualified by nine out of ten Tier One OEMs as their 10GbE platform. This includes standard stand-up adapters for rack-mount servers and custom form factor adapters for blade servers manufactured by HP, IBM, Dell, Cisco and Fujitsu Technology Solutions.

OneConnect adapters are also included as LOMs for high-performance servers from HP and IBM. As a result, Hadoop clusters can standardize on OneConnect adapters across a wide range of servers from multiple sources.

Boost Hadoop Performance with Emulex OneConnect 10GbE Network Adapters

Feature	Benefit
Speed	**Increased performance** with 300% data transfer rate increase over 1GbE **Lower CapEx** through increased performance reduces required cluster size **Increased ROI** through increased cluster node efficiency
Simplified management	**Reduced OpEx** through simplified management via a single console; OneCommand Manager allows version management and testing by supporting common NIC drivers, common firmware and boot code, consistent NIC teaming **Hyperscale infrastructure solution**
Broad Linux support	**Advanced driver architecture** for leading Linux distributions including Ubuntu, CentOS, Red Hat Enterprise Linux, and FreeBSD **Flexibility** in choice of operating systems

Summary

Tests have shown that the Emulex OneConnect OCe11102 10GbE Network Adapter boosts Hadoop's data transfer rate by 300% providing significant performance improvement and cost savings. A single intelligent framework, OneCommand Manager enables customers to efficiently automate and dynamically provision their Ethernet environments while reducing operational cost. Providing broad Linux support, Emulex offers customers flexibility in their choice of operating systems.

World Headquarters 3333 Susan Street, Costa Mesa, CA 92626 +1 714 662 5600
Wokingham, UK +44 (0) 118 977 2929 | **Munich, Germany** +49 (0) 89 97007 177
Paris, France +33 (0) 158 580 022 | **Beijing, China** +86 10 68499547
Tokyo, Japan +81 3 5325 3261 | **Bangalore, India** +91 80 40156789

Connect with Emulex

twitter.com/emulex friendfeed.com/emulex bit.ly/emulexlinks bit.ly/emulexfb

EMULEX®

www.emulex.com

12-0616 · 1/12