


Don't Lose Application Performance In an I/O Digital Traffic Jam

Intelligent Ethernet from Cavium improves
performance by over 55%



During heavy I/O congestion an application workload utilizing a Cavium FastLinQ NPAR partition maintains 100% of its performance

EXECUTIVE SUMMARY

Does it matter how many customer orders you can process? Of course it does, and it can matter most when network traffic is at its peak. You must ensure your customers can reach your business applications with a network that will deliver the performance your customers have come to expect. With a Cavium™ FastLinQ® Adapter, your applications are guaranteed the bandwidth to process orders, perform database queries, run simulations, and all the tasks that keep your business running.

Traditionally, to provide quality of service (QoS) for a specific application you would have to dedicate more physical resources such as more 10GbE ports or adapters. With Cavium Switch Independent NIC Partitioning (NPAR) you can allocate a specific portion of the available bandwidth to an application, virtual machine, or other key resource to assure that resources have the bandwidth to operate at peak efficiency even during the busiest times.

To illustrate how well applications respond to Cavium's NPAR, we exercised a data warehouse application using SQL server. During heavy network conditions, the application performed 300,000 more transactions per minute (TPM) with NPAR protection. 300,000 more TPM means a lot more customers are processing a lot more orders.

OVERVIEW

We first measured the baseline performance of our application by determining how many transactions per minute (TPM) could be completed when the application had exclusive access to the 10GbE port.

Next, we added more application workloads to increase the overall I/O traffic on the 10GbE port. The results we measured were stunning:

- During heavy I/O congestion, an application running on a standard 10GbE LOM lost 55% of its performance due to competing network traffic.
- During heavy I/O congestion, an application workload utilizing a Cavium FastLinQ NPAR partition maintained 100% of its performance even in the presence of competing heavy network traffic.
- Cavium Switch Independent NPAR maintains quality of service (QoS) and ensures service level agreements (SLAs) are met.




Figure 1. SQL 2014 Three Node AlwaysOn HADR Cluster

Servers

- Intel Xeon E5-2640, 2.5GHZ, Dual Sockets, 12 Core each, 24 total cores, 32GB memory
- 10GbE LOM
- Cavium FastLinQ 10GbE QLE3442-SR-CK Intelligent Ethernet Adapter
- QLE2562 8Gbps FC HBA

Switch

- 8Gb FC Switch
- 10GbE Ethernet switch

Storage

- DataCore All Flash Array with (4) 300 GB SSD, 1.2TB usable space

Database

- The databases were built with an OLTP DB performance tool using a TPC profile, at 25, 50 and 100GB.
- TPM measured at 300 Virtual Users

METHODOLOGY

To create the real-world architecture similar to a common data warehouse environment depicted in Figure 1, we virtualized the front-end server and connected via 10GbE to our back end SQL 2014 Three Node AlwaysOn HADR cluster. We installed and configured an OLTP DB workload simulator application on our front-end server and ran tests with 300 Virtual Users. The server running this client side application was virtualized using Windows Hyper-V.

Application performance protection was achieved by first creating a Cavium NIC Partition on the Cavium FastLinQ 10GbE QLE3442-SR-CK Intelligent Ethernet Adapter installed in the host machine of the virtualized application servers. The resulting partition, called a “Physical Function” appears to the VMs and applications like a fully featured physical NIC port with its own IP address. In reality, the Physical Function is a logical division of a 10GbE physical port and the bandwidth allocated to that Physical Function can be set to a guaranteed minimum.

We allocated 20% of the total network bandwidth to the physical function. The OLTP workload simulator application was then configured to access the SQL database using that physical function. Next, we ran a baseline test with no competing network traffic and measured about 835,000 TPM. After establishing this baseline, we ran a high traffic contention test with multiple VMs running multiple applications all sending I/O requests through the same Cavium FastLinQ 3400 Series Adapter, but on a separate physical partition from the virtual user's partition. As Figure 2 illustrates, the Cavium FastLinQ 3400 Series Adapter maintained 100% of the baseline performance even when this heavy I/O condition was occurring.

Finally, we ran tests with a popular 10GbE LOM that did not offer NPAR as an option. The LOM took a staggering 55% performance hit when network traffic was heavy. This equated to over 340,000 TPM lost—340,000 less Web server queries, processing 340,000 less real-world orders. The basic 10GbE LOM without Cavium NPAR technology could not protect this critical application's performance.

CONCLUSION

We showed that our application could perform over 300,000 more orders during a time of heavy network congestion. What are your orders worth to your business? If each order were for a \$1.00 widget for example, then your company just made \$300,000.00 more dollars per minute. In addition to increasing revenue from higher performing applications, you must consider the total cost of ownership (TCO).

Instead of purchasing too many 10 GbE ports, Cavium NPAR allows network administrators to configure a 10GbE port as if it were multiple physical ports. They can divide a single adapter into as many as four partitions per port. This means that a dual-port Cavium FastLinQ adapter can yield the functionality of eight 10GbE ports! And, an administrator can apply specific, dynamic QoS to each of these eight Cavium NPAR physical functions.

If you are considering a simple 10GbE LOM as your primary method of connecting to your network, do not risk a performance degradation of your money making applications on ports that get stuck in a digital traffic jam. To ensure your products can be ordered quickly all the time, you might want to reconsider that LOM purchase, or consider augmenting a simple LOM with a fast and smart Cavium FastLinQ 3400 or 8400 Series Adapter.

ABOUT CAVIUM

Cavium, Inc. (NASDAQ: CAVM), offers a broad portfolio of infrastructure solutions for compute, security, storage, switching, connectivity and baseband processing. Cavium's highly integrated multi-core SoC products deliver software compatible solutions across low to high performance points enabling secure and intelligent functionality in Enterprise, Data Center and Service Provider Equipment. Cavium processors and solutions are supported by an extensive ecosystem of operating systems, tools, application stacks, hardware reference designs and other products. Cavium is headquartered in San Jose, CA with design centers in California, Massachusetts, India, Israel, China and Taiwan.

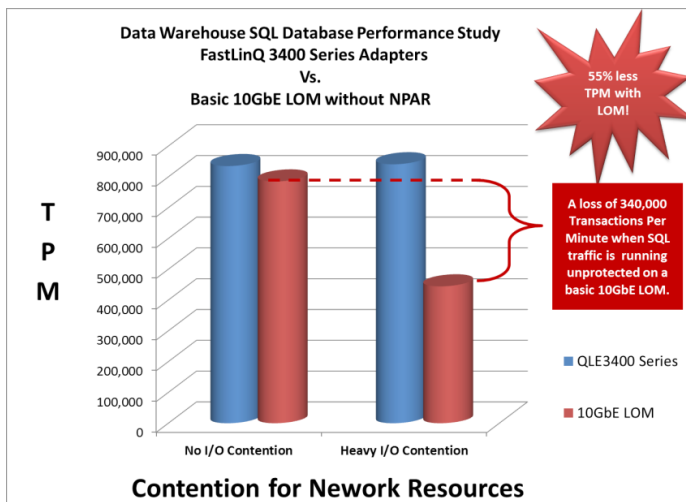


Figure 2. Application Performance Comparison results: Cavium FastLinQ NPAR vs. Basic 10GbE NIC or LOM



Follow us: [f](#) [g+](#) [in](#) [t](#) [You Tube](#) [RSS](#)

Corporate Headquarters Cavium, Inc. 2315 N. First Street San Jose, CA 95131 408-943-7100

International Offices UK | Ireland | Germany | France | India | Japan | China | Hong Kong | Singapore | Taiwan | Israel

Copyright © 2015 - 2017 Cavium, Inc. All rights reserved worldwide. Cavium and FastLinQ are registered trademarks or trademarks of Cavium Inc., registered in the United States and other countries. All other brand and product names are registered trademarks or trademarks of their respective owners.

This document is provided for informational purposes only and may contain errors. Cavium reserves the right, without notice, to make changes to this document or in product design or specifications. Cavium disclaims any warranty of any kind, expressed or implied, and does not guarantee that any results or performance described in the document will be achieved by you. All statements regarding Cavium's future direction and intent are subject to change or withdrawal without notice and represent goals and objectives only.