


Cavium FastLinQ 3400 Series Intelligent Ethernet Adapters for Flexible Networking

Practical Use Cases for Oracle RAC Database



The Cavium FastLinQ 3400 Series Intelligent Ethernet Adapters offer Switch Independent NIC Partitioning (NPAR) that provide the flexibility to effectively utilize the entire bandwidth available in these adapters.

BACKGROUND

The Oracle RAC (Real Applications Cluster) database is installed on servers (nodes) that are managed by Oracle's Cluster services. Each node must have at least two network adapters or network interface cards (NICs): one for the public network interface and one for the private network interface (the interconnect). This paper refers to them as the "public connection" and the "private connection."

An Oracle RAC database may be consolidated onto an Oracle Grid, sharing the servers with other RAC databases. This places even more demand on the Ethernet network to provide low latency and high bandwidth for the applications running on top of the databases.

Applications use the public connection to access, load, and back up the database.

The private connection includes a block buffer cache in each machine within the cluster. Oracle refers to this as the "Cache Fusion." This means that the block buffer cache from each machine is "fused" into one buffer cache that encompasses all machines in the RAC cluster.

An Oracle RAC database has multiple nodes that access and update data that are shared by all of the nodes in the cluster. The Database Management System (DBMS) uses the private connection for the global lock management to manage the consistent read that is enforced for queries when data being read is also being updated. If a node needs a database block that is held by another node in the cluster, the block is sent over the private connection.

Updating a table on one or more nodes places a lot of traffic on the private connection. This load increases when there are multiple Oracle RAC databases sharing the servers in the grid.

USE CASES

The following examples illustrate configurations for three different cases that use Oracle and the Cavium™ FastLinQ® 3400 Series Intelligent Ethernet Adapters.

Two Cavium FastLinQ QLE3442 Adapters without NPAR

The first use case illustrates two QLE3442 adapters used in public and private connections without enabling the NPAR function. This shows how the QLE3442 might be used as if it were any standard 10Gb dual-port NIC adapter. Port 0 on each card team together at the OS level to share the ports and increase bandwidth and redundancy. Port 1 on each card is managed by Oracle Clusterware and the load is balanced between the adapters. In this configuration, the public connection has 20Gb available bandwidth and the private connection has 20Gb available bandwidth (Figure 1).

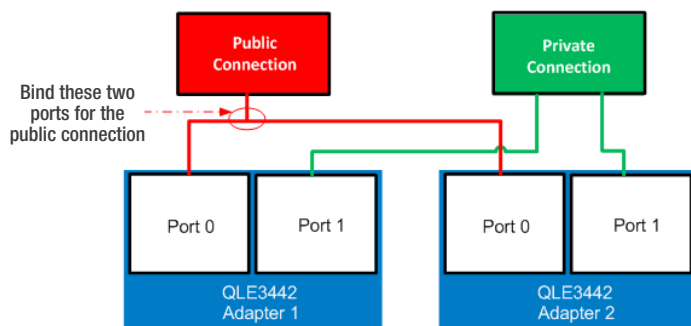


Figure 1. Public and Private Connections Using QLE3442 Adapters without NPAR Enabled

Two Cavium FastLinQ QLE3442 Adapters with NPAR Used to Create Two Partitions Per Port

The second use case illustrates two QLE3442 adapters used with NPAR configured on each port to divide them into two partitions. Each partition appears to the OS as a separate physical function (PF) or port. Partition 1 (PF1) on each port for both cards are teamed together at the OS level to share the ports and increase bandwidth and redundancy. Partition 2 (PF2) on each port for each adapter are managed by Oracle Clusterware and the load is balanced between all of the partitions.

Note: Oracle Clusterware will actively use the four Ethernet ports for the private connection, although more may be assigned.

The NPAR function has a quality of service (QoS) setting that guarantees a given bandwidth to each partition on the configured port. In this example, the QoS is set at 30% on partition 1 (public) and 70% on partition 2 (private connection) for each adapter. This means that the public connection, partition 1, will receive no less than 30% of the available 10Gb of bandwidth from each adapter port. And partition 2 (private connection) will receive no less than 70% of each port's bandwidth. In this example, with everything demanding the maximum bandwidth, partition 1 would receive 12Gb for the public connection and partition 2 would receive 28Gb.

The QoS settings here will permit the private connection to use up to the maximum bandwidth on the four partitions, up to 40Gb if the public connection does not need it. Or, if the database is being backed up across the public connection, the QoS will allow the backup to reach the bandwidth that the backup server can accept.

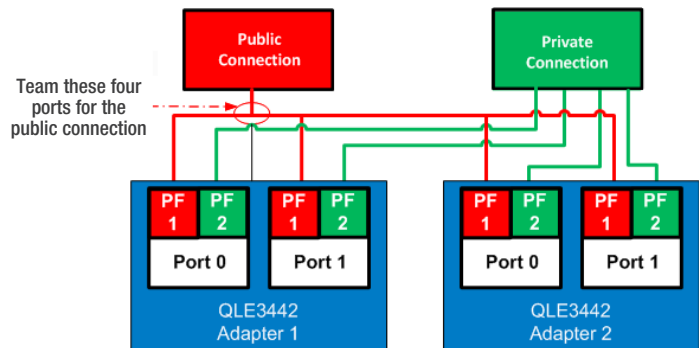


Figure 2. Public and Private Connections Using QLE3442 Adapters with NPAR Used to Create Two Partitions Per Port

Two Cavium FastLinQ QLE3442 Adapters with NPAR Used to Create Two Partitions on Port 0

The third use case illustrates QLE3442 adapters used with NPAR configured on port 0 to divide it into two partitions. Each partition appears to the OS as a separate physical function (PF) or port. Port 1 on each adapter is completely dedicated to the private connection with a maximum bandwidth of 20Gb. This example also connects to the backup VLAN. This ensures that LAN traffic for backing up does not interfere with the LAN connection for the network.

The QoS will permit the public connection and the backup VLAN to use the bandwidth that is needed. To achieve this configuration without NPAR would require a third 10GbEGB NIC, additional switch ports and cabling.

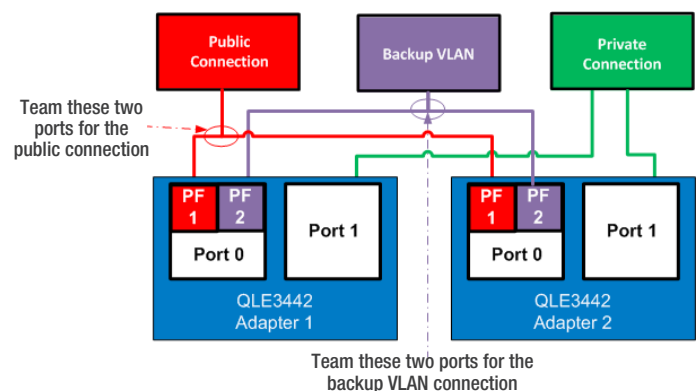


Figure 3. Public and Private Connections Using QLE3442 Adapters with NPAR Used to Create Two Partitions on Port 0

Note: This sample configuration without NPAR requires a third 10Gb adapter, additional switch ports and cabling.

RECOMMENDATIONS

The Cavium FastLinQ QLE3442 dual-port Intelligent Ethernet Adapter provides 10Gb per port, which adds up to 20Gb for each adapter. For high availability, administrators can connect two QLE3442 adapters to different 10Gb switches. For example, they can connect port 0 on each QLE3442 to Ethernet switch “A” and port 1 to Ethernet switch “B”.

LEARN MORE

- For details about Oracle Clusterware, read [Oracle Real Application Clusters \(RAC\) and Oracle Clusterware Interconnect Virtual Local Area Networks \(VLANs\) Deployment Considerations](#).
- For information on the Cavium Switch Independent NIC Partitioning, read [Enhancing Scalability Through NIC Switch Independent Partitioning](#).

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.



Follow us:      

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.