



Red Rock Canyon

Technology

Network Preprocessing and Load Balancing

July 2015

Preface

Network application, be it a firewall, an attack mitigation system, layer 7 load balancer, or other, all deal with traffic coming their way out of a tap device, or through a direct inline network connection. Application cannot predict or profile the traffic streaming in before it reached the system. Once traffic swarmed in, filtering and redirecting can be performed. For instance, a security implementation wishes to differentiate LAN from WAN traffic is able to do that once it has arrived into system, using CPU cycles and system memory. Moreover, if an application wishes to receive traffic in via an array of interfaces, in a balanced manner, for processing optimization purposes, then an interaction with an external device such as a switch, is required [1].

In the absence of a switching device at hand, application designer would opt for one of the two. Either to make some use of the bridging or filtering capabilities of the OS (e.g., bond device interface, or `netfilter` under Linux); or have the application deal with this burden. On both cases, the host CPU is the engine that should perform this network preprocessing.

Silicom's approach is to offer a new capability to the application, to preprocess network traffic before it reaches the application and its business logic; thus letting the application to focus on its business logic. All of this is done free of CPU power, and in a totally standard yet innovative manner.

Preprocessing with Smart NIC

Silicom offer a 4 x 10GbE network adapter that features an embedded fully capable network switch on it. This simple concept immensely leverages and expands the capabilities of a network application, by letting it powerfully control and manage elaborate schemes of network engineering, to its benefit. The

advantages of using such preprocessing are the following:

Traffic engineering – Application is now able to perform actions on the traffic, whether ingoing or outgoing, that were not easily achievable before, like NAT (and PAT), policing/QoS, five tuple **filtering**, **link aggregation**, network sampling, smart **mirroring** and many more capabilities that are offered by a modern switching fabric.

Close to application – The close physical location of the switching fabric as well as the application means very fast on-the-fly configuration. The access to the configuration plane of the switching fabric does not require SNMP, Netconf or other. The access is to an on board device, with **minimal latency**.

Industry standard – Intel® components are all over. The switching fabric as well as the MAC are all intel® components, with comprehensive software support already part of mainline open source projects.

More capable than CPU – Dedicated switching fabric is more capable then a standard general purpose CPU as it comes to network processing. Fully aware of the fact that more and more vendors of network application shift to standard COTS hardware, Silicom offers RRC based solution to adhere to this trend; let alone the fact that such switching fabric is expected to be seen as an integral part of COTS server within the next twi to three years.

True preprocess – Last but not least, the effect of network preprocessing on the overall system behavior and performance is visible not only due to the larger networking features set that a switching silicon offer compared to a general purpose CPU, but also due to the fact that all is done as a preliminary stage, completely offloaded off the CPU and system.

Load Balancing

Intel® Red Rock Canyon chip set offer a good deal of load balancing options for incoming traffic. Load balancing ports are grouped together to form load balancing group (LBG), over which a traffic can be engineered to be load balanced according to several simple yet powerful schemes.

Table 1 - Load Balancing Hashing

L3 hash		L3 / L4 hash
Source MAC address		Source IP address
Dest. MAC address		Dest. IP address
Ethertype		Source port
VLAN ID		Dest. port
VLAN Priority		DSCP
Symmetric MAC		ISL
		Symmetric IP/port
		ECMP
		IP protocol

Having load balanced ingress traffic can serve well for packet processing within the system. Ingress traffic become more predictable, and processing threads can be spawned as per the expected streams of traffic. Table 1 summarize the parameters that are hashed to become the basis for traffic distribution across ports.

Yet, it should be noted that switch ports on adapter are connected both towards host, as well as towards external I/O ports (see Figure 1). The switching fabric can be set to load balance traffic towards host, as well as towards external I/O ports.

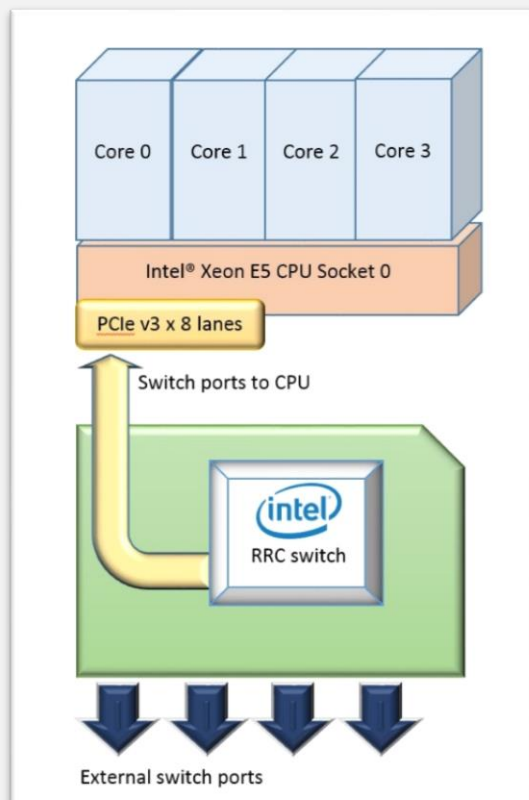


Figure1 - RRC Adapter Load Balancer

REFERENCES

- [1] Packet classification as a means of acceleration:
<http://ieeexplore.ieee.org/application/enterprise/entconfirmation.jsp?arnumber=6857541&icp=false>