

Network I/O Challenges of Data Center Virtualization

FORMS OF VIRTUALIZATION THROUGH THE YEARS

The term virtual machine (VM) was first used in the 1960s. Since then, the use of varied forms of virtualization as an efficiency mechanism has grown dramatically. A recent study found that in 2010 alone, nearly 60% of data centers are expanding their virtualization deployments.

Virtualized network infrastructure has been around for years. Notable technologies, like Ethernet VLANs, IPSec and SSL VPNs, and L3 VPNs via MPLS or virtual routing, are all examples of tried-and-true technologies for virtualizing networks. These techniques allow a single set of physical resources to be shared among a diverse group of users, providing isolation, performance guarantees and security. Each of these benefits can also apply to application hosting platforms. Mechanisms to virtualize servers are now becoming commonplace with server virtualization seen as the key to the convergence of networking and computing in the data center. According to research firm Gartner,[®] there will be about 58 million deployed VMs by 2012.*

WHAT IS SERVER VIRTUALIZATION?

To realize the numerous benefits of data center consolidation, many under-utilized physical servers with a single operating system (OS) and, typically, one hosted application are being replaced with a smaller number of systems supporting multiple guest OSs and application instances. This VM approach can dramatically improve the efficiency, power utilization and availability of costly hardware resources, such as CPU and memory. The “one server, one application” model is quickly becoming a thing of the past.

HOW SERVER VIRTUALIZATION CHANGES DATA CENTER DESIGNS

This profound change in data center architectures has important implications for parts of the network infrastructure, as well as the virtual server itself. Before this virtualized server model, when applications were uniquely tied to physical compute resources, all of the required networking functions happened outside the server. Operations like packet switching and routing, as well as network security functions like firewalls and intrusion prevention, occurred in discrete devices that were layers away from the hosting servers. In addition, the server infrastructure was typically front-ended by application load balancers used to evenly spread traffic across the hosting platforms. This network infrastructure sprawl led to very complex and power-hungry data center architectures.

Changing the ratio of applications to servers from 1:1 to N:1, as is the case when IT managers virtualize their hosting resources, has important implications for the servers themselves, thus changing the way we must think about and architect products for the data center. Current-generation multi-

NETRONOME WHITE PAPER

EXECUTIVE OVERVIEW

Fundamental changes are occurring in the data center; and these changes will have dramatic and lasting effects on how we think about designing computer networks and hosting platforms.

In previous years, servers simply acted as application hosting devices. Owing to the explosion of multicore CPUs and virtualization, the server now needs to look more like a hardened network element with the ability to perform extensive networking and security processing. A missing element in these designs is the server's I/O. IT architects need to consider how to link virtualized network infrastructure and virtualized server platforms without degrading performance.

NETRONOME

The Flow Processing Company

netronome.com

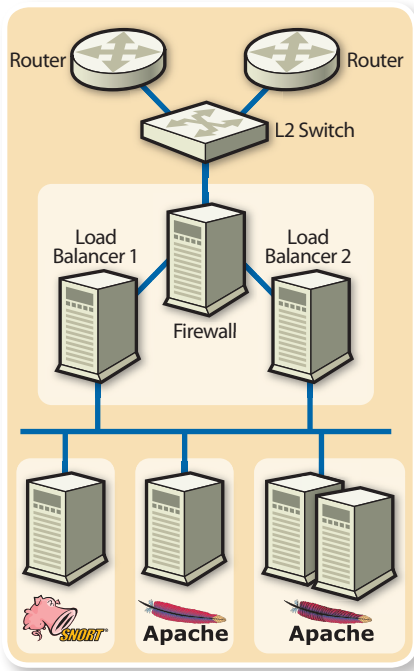


Figure 1. Network infrastructure and hosting platforms are physically separated in the data center.

core x86 servers can have up to six or more cores on a single CPU die, with core densities expected to climb to eight or 12 in the coming years. These servers can support many applications per physical device. When a single physical server supports many applications, IT managers and architects must consider how this radical change affects their networks. This change requires that, in addition to application and control-plane processing, networking and security functions previously housed outside the x86 server must now collapse inside the virtualized multicore platform to provide the same set of functions previously handled in separate and discrete pieces of hardware. Special processing, such as packet classification, flow-based load balancing, active flow state management and flow pinning, L2 switching, L3 forwarding and QoS treatment, must collapse out of discrete devices and into the virtualized x86 server for all of the same reasons they were separated in yesterday's discrete 1:1 host-to-application model.

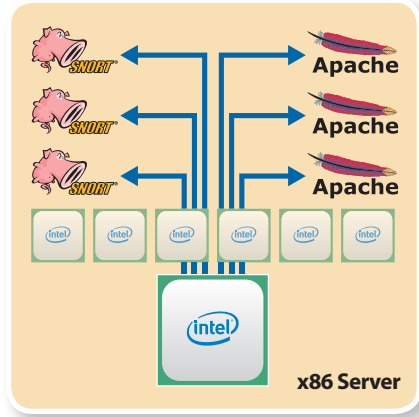


Figure 2. Multicore CPUs support multiple application instances.

Specifically, when traffic is received at the server, a variety of packet-processing and security workloads must be applied to the data. Packets must be classified to determine the destination core/application. Traffic must have security policy applied to it to avert threats. And it may also have to be encrypted/decrypted as with IPSec and SSL traffic. Finally, an L2 switching decision must be made to place the packets into the correct VM. Adding to these requirements, this processing must now also occur in a stateful fashion to assure that packets belonging to the same flow all reach the same destination application.




Figure 3. Virtualized servers need to support networking requirements.

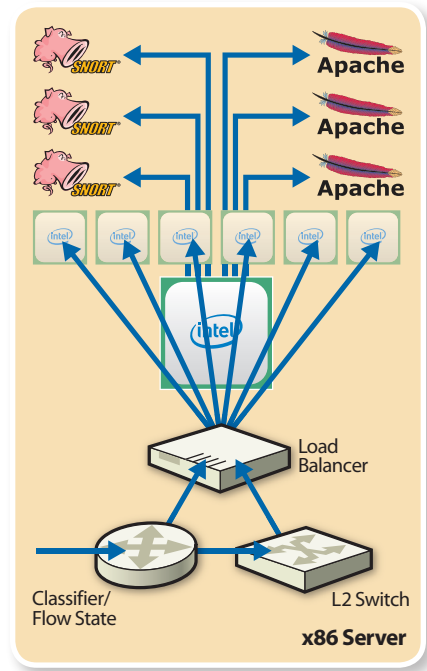


Figure 4. Networking functions occurring in software wastes x86 CPU cycles.

This packet processing is often implemented in software as a component of the VM hypervisor. However, multicore x86 is not optimized for network and security processing, thereby creating a huge performance problem in these devices. This x86-only architecture is unable to capture packets and process flows at the highest speeds because packet-processing interrupt handling, security processing, classification and forwarding are all tasks that waste valuable CPU cycles. Performing these varied tasks in software results in poor server performance and high power consumption per gigabit per second of throughput, ultimately reducing application performance. The conclusive result is that this new architecture requires more servers to support the compute load, eliminating the desired effect that virtualization was ultimately supposed to remedy.

ENTER I/O VIRTUALIZATION

The missing link between virtualized networks and virtualized servers is the I/O from the network interface into the respective VMs. An innovative architecture for virtualized servers combines workload-optimized network flow processors (NFPs) and x86 multicore processors across a high-performance virtualization-aware PCIe communications path. The NFPs provide hardware-based offload of network and security processing from x86, returning to those processors all of the CPU cycles and allowing them to focus on application and control-plane processing. This architecture supports efficient delivery of data to VMs at extremely high rates — 20Gbps and higher.

These special flow processors are optimized to offload burdensome workloads from general-purpose multicore CPUs. The flow processors handle lower-layer packet processing and accelerate higher-layer flow and application-level processing. This accelerated architecture uses the network-optimized flow processor cores for packet classification, stateful flow analysis, deep packet inspection, cryptography and dynamic flow-based load balancing. In addition, the critical function of demultiplexing traffic into specific VMs is implemented in the flow processor; and the L2 switching no longer needs to occur in the VM hypervisor. This I/O virtualization (IOV) provides guaranteed bandwidth, latency and isolation of traffic across multiple cores and VMs. Without IOV, many of the benefits provided by OS virtualization would be lost.

Conclusion

Building products in the with the considerations covered in this paper will enable more efficient use of scarce processing resources in the most efficient and “green” way possible to support the huge bandwidths that users demand.

Daniel Proch is the director of product management at Netronome Systems, and can be reached at daniel.proch@netronome.com.




Figure 5. Using flow processors for network and security processing, as well as VM demultiplexing, saves CPU cycles.

*Source: *Gartner Says 16 Percent of Workloads are Running in Virtual Machines Today*, October 21, 2009 press release (<http://www.gartner.com/it/page.jsp?id=1211813>).

© Netronome is a registered trademark and The Netronome Logo and “The Flow Processing Company” are trademarks of Netronome Systems, Inc. All other trademarks are the property of their respective owners.
© 2010 Netronome Systems, Inc. All rights reserved. Specifications are subject to change without notice. (10-10)



Netronome has operations in:
USA (Pittsburgh [HQ], Santa Clara & Boston),
UK (Cambridge), **Malaysia** (Penang),
South Africa (Centurion) and
China (Shenzhen, Hong Kong)
info@netronome.com +1 877 638 7629
netronome.com