CAVIUM
Partner

PL-81060

WIN Enterprises White Paper

A New, Economical Approach to Implementing Content Delivery Network (CDN) Systems

www.win-ent.com

A New, More Economical Approach to Implementing Content Delivery Network (CDN) Systems

PL-81060 High Performance Storage

Why were CDNs implemented?

CDNs aren't a new phenomenon; they have been around since the early nineties. As Internet activity ramped up the really active data centers, websites, and on-line applications were prone to crashing during high-traffic periods. Data was still centralized, creating bottlenecks as a multitude of simultaneous data requests poured in to the central server. This created what amounted to self-inflicted DDOS (Distributed Denial of Service) attack as the system was slowed.

Navigating the Cloud without CDN

We are all familiar with TCP/IP, the behind the scenes protocol and addressing method that enables any data requested on the Internet to navigate from its source through the cloud to be delivered to the correct machine and IP address.

While IP is dealing with the addressing portion of the problem, TCP enables the data to traverse the Internet by dealing with multiple servers before it reaches its final destination. Unfortunately, TPC is "chatty" and slow by today's standards as it handshakes its way from server to server navigating the cloud. Versus a CDN, regular un-optimized TCP has a goal of making the least number of hops in its journey, versus producing the lowest overall round-trip time, which would deliver a better user experience. A CDN optimizes TCP for the lowest round trip time.

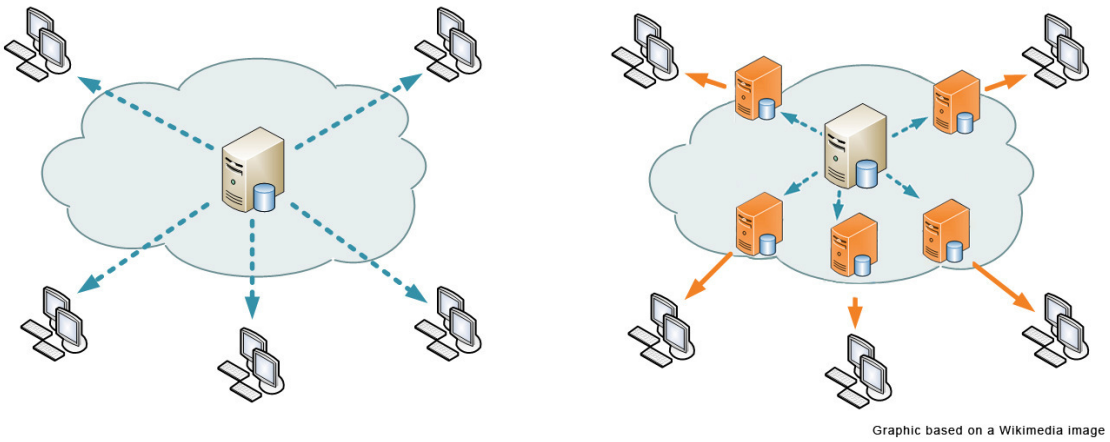

Graphic based on a Wikimedia image

The configuration on the left shows a typical TCP/IP Internet network. This is subject to traffic jams at the central server site. The one on the right represents a CDN and shows how the network is architected to push out replicated data to the edge of the network so that it's close to the end-user providing rapid data access. In addition, this offloads the central data center to keep it running smoothly.

Benefits of CDN

Even when the information system is worldwide in scope, the goal of CDN is to deliver local performance with a consistent and pleasant user experience.

A fast response time is provided through the tactic of data replication. The replicated data is generally static data (documents, images, videos, etc.) stored at a centralized data center. Different analogues are used, but typically new data is pushed out from a central server site to all the remote nodes. This makes data quickly accessible and more reliably available as non-working nodes can fail-over to healthy ones when needed to deliver requested data in a timely fashion.

The PL-81060 delivers LAN performance at 10 or 40-Gbps with full, end-to-end virtualization support. This translates to low latency and a better user experience. In addition, as the load on central data center is reduced due to receiving fewer direct requests, the overall system performance and reliability is improved.

Implementations of CDN

Adding new value-added services is the way organizations typically expand their product offerings, whether the organization is a service provider, SaaS, or video download provider like Netflix. Deployments of the remote servers can be accomplished any number of ways, for instance:

• By deploying dedicated PoPs around network's edge

• Partnering with a regional ISP and utilizing its existing servers

• Installing your own dedicated hardware at the ISP partner's PoPs

Does everyone need a CDN?

For smaller websites and sites with relatively light traffic, latency is not a big problem. For many websites virtual real-time delivery is not expected, so they don't really require a CDN. However, for busy sites where performance is linked to the lifeblood of the company, a CDN may be a necessity for company survival. In these cases the more robust the data delivery solution the better.

The CAVIUM® ThunderX™ processor, is a cost effective, high performance low thermal processor implemented in ARMv8 architecture for maximum economy and lowest Total Cost of Ownership. PL-81060, uses dual CAVIUM ThunderX™ processors to provide Data Centers, Service Providers, SaaS companies, and other cloud-centric businesses with a solution for deploying their remote servers more economically to support Content Delivery Networks (CDNs).

The dual processors of the PL-81060 feature up to 48 cores each (i.e., a maximum of 96 cores per board). At a single core performance rate of 2.0 GHz x 96 cores, the device delivers 192 GHz of compute power for data storage and dissemination activities. One processor of the device is dedicated to I/O handling, the second to the acceleration of any task it receives.

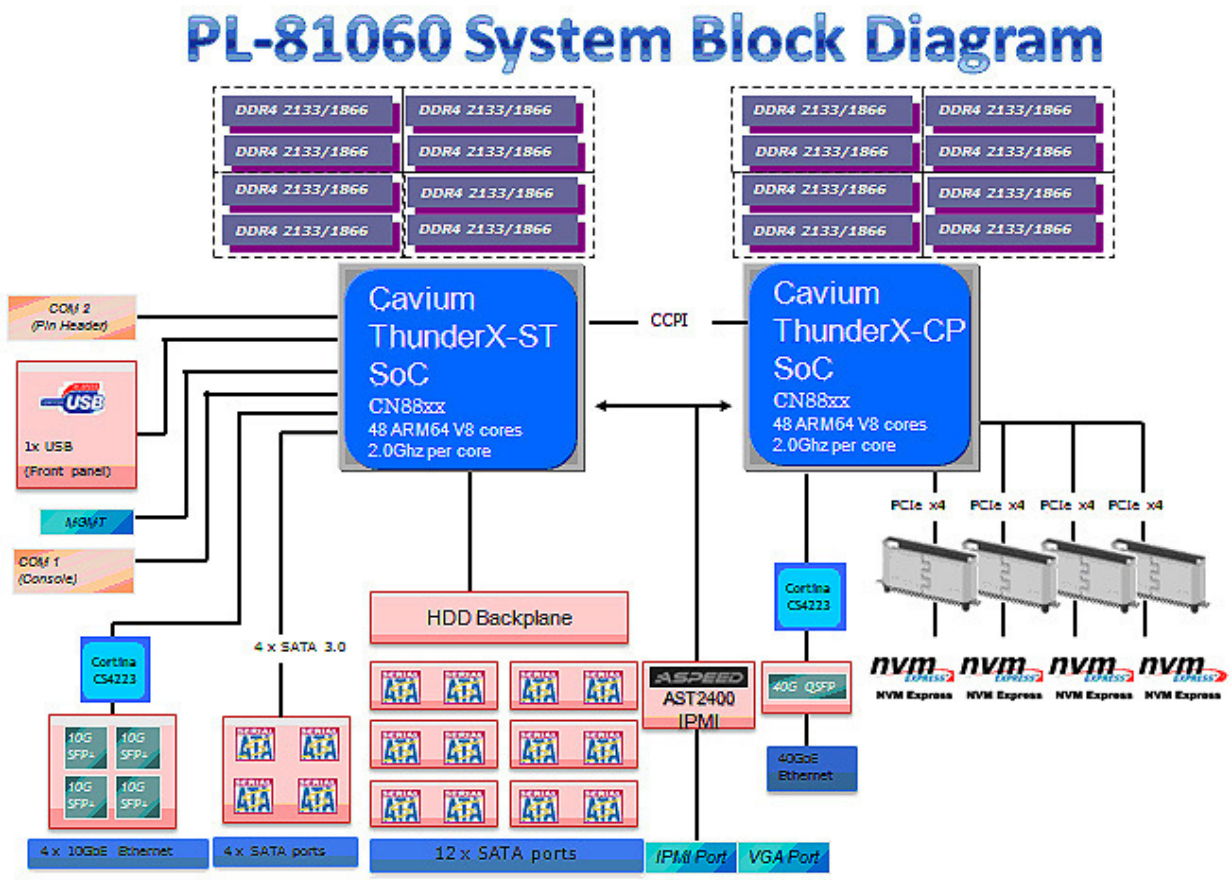

The PL-81060 features dual Cavium ThunderX™ processors. The first processor handles I/O functions; the second accelerates whatever functions directed to it. In general, the second processor enhances delivery of static data to/from the central server site and to the end-user.

Twenty-four 3.5-inch SATA HDDs provide up to 240TB of storage. Four 2.5-inch SATA or NVMe SSD provide up to an additional 8TB for system buffer cache, for maximum 248 TB of storage. The SSDs also provide fast upload and download capability as data is replicated, stored and distributed to requesting end users within the geographic area served by the PL-81060. This capability is key to the function of PL-81060, as its goal is to provide local accessibility to the end users that are relatively close to the edge of the network. It's easy to see that in a large and busy data center a centralized system would be slow and subject to strain on the central server.

It's ARM heritage means the PL-81060's offers cost effective pricing, low power consumption, and lower heat production. Lower heat and power means the unit delivers exceptionally low Total Cost of Ownership (TCO) over its lifetime. In addition, LAN performance of 4x 10 GbE and 1x 40GbE support a fast and consistent end-user experience.

Lower unit cost enables placement of more CDN servers at the network edge to provide a greater number of your end-users with fast, high-quality data and application access.

WHY IS PL-81060 THE PLATFORM TO MEET TODAY'S CDN MARKET NEEDS?

The more robust an organization's presence at the edge of the cloud, the better the quality of service it can provide to the end user.

• What attributes are needed at the edge?

Ubiquitous presence (meaning a healthy number of servers to provide rapid response). As an economical ARM-based solution, the PL-81060 enables you to deploy more instantiations in the field versus traditional and more expensive x86 approaches.

• Servers must be able to rapidly download / upload data originating from the central data center

The PL-81060 is configured with four SATA or NVMe SSDs that deliver a total of 8TB of rapid-response storage (caching).

• Servers require massive storage capacity to handle the replicated data coming from the central data center

PL-81060 offers 240TB of storage through a maximum of 24x 3.5 SATA HDDs, plus an additional 8TB through four SATA or NVMe SSDs.

• CDN server must have rapid communications capability

PL-81060 provides 4x 10/40GbE (processor #1) and 1x 40GbE (processor #2) fiber ports to speed upload/download and to enable rapid data dissemination to the end-user.

PL-81060 Total Cost Ownership

Architecture	Intel x86			CAVIUM/ARMv8 SoC		
Item	Part#	List Price	Power	Part#	List Price	Power
CPU(2)	E5-2690V4	$2090x2	~270W	ThunderX (2)	$1350x2	~240W
Chipset	C61X	$60	7W	Integrated	$0	0
LoM	X520	$50	15W	Integrated	$0	0
Storage HBA	9300-8i	$150	15W	Integrated	$0	0
		$4,440	307W		$2,700	240W

PL-81060 Virtualization Applications

CPU	Intel E5-2690V4	CAVIUM ThunderX
Configuration	2X	2X
Number of cores	28	96
VMs / core	2	1
Number of VMs	56	96
IO subsystem	External	Integrated
IO Virtualization	External of Software	Integrated

At the platform level we estimate a 40% savings in purchasing the PL-81060 versus other available units. In addition these units offer a lower Total Cost of Ownership (TCO) and more performance per dollar and per watt.

PL-81060 and the Cavium® ThunderX™ processor

The dual processor implementation is currently the most popular architecture for data centers. At a total of 96 cores at 2.0GHz per processor for the PL-81060, that's 192 GHz of power are brought to bear on the problem of data replication and dissemination from a central data center. In implementing a CDN with the PL-81060, a service provider can designate certain cores for certain tasks and/or specific customers enabling greater business efficiency.

The SoC design of the Cavium ThunderX™ enables greater operational efficiency versus implementing a processor solution that requires incremental chips on the board, as is the case with non-SoC x86 solutions. This SoC approach reduces processor footprint, the amount of heat generated, and adds to the efficiency of the motherboard by eliminating the distance between features (i.e., chipset, LoM, and Storage HBA).

The ThunderX™ represents a partnership between ARM and Cavium. In this partnership, the ARM architecture has been licensed and implemented by Cavium on a grander scale than heretofore. This is a departure from the past where the ARM architecture was usually implemented for lower-end industrial or hand-held consumer devices. Cavium chose to utilize the ARM architecture because it could deliver on the market's required lower cost per unit without sacrificing performance or reliability. This is consistent with the direction of the general server market, which is showing movement toward ARM technology.

Consistent with a trend toward open systems, Cavium ThunderX™ is designed to support the Linux operation system. The PL-81060 currently supports the following versions of Linux.

• ubuntu 14.04/16.04

• CentOS 7.1 64 bit

• Red Hat 7.2

• Fedora 21 64 bit 7.2

• Open SUSE 13.2 64 bit

x86 System Conversions

Because Cavium ThunderX™ resides in the Linux world, open-source cross compilers can be used to port applications that began life as x86-based applications. The caveat is that this can be straightforward as long as hardware has not been permanently altered by the software operations; for instance, changes that have been implemented within the original processor's registers. This would have to be reviewed as part of your system evaluation.

CONCLUSION

Organizations with highly active data centers, websites or service hubs need CDN to provide high quality of service (QoS) for their global end-users. PL-81060 is a CDN server with high-performance and low TCO that enables these organizations to deploy more points-of-presence (PoPs) at the edge of the cloud to provide end-users with fast, reliable service.

www.win-ent.com