

Military Surveillance and Communications Boosted by Intel® Xeon® Processor-Based VPX* Server Board

Concurrent Technologies more than triples the number of CPU cores in a rugged microserver form factor.

“This new 3U VPX* board is further proof of our commitment to improving performance within rigorous SWaP constraints, utilizing the latest processors from Intel.”

Glen Fawcett, CEO of Concurrent Technologies

More CPU Cores, Same SWaP

With an increased focus on situational awareness on the battlefield, there is a need for much more computing power in military vehicles such as aircraft, tanks, and Humvees. These vehicles provide warfighters with vital, real-time information generated by compute-heavy command, control, communications, computers, intelligence, surveillance, and reconnaissance (C4ISR) systems. Concurrent Technologies is making it easier for equipment manufacturers to pack more computing power into a small, convection- or conduction-cooled system by offering a 3U VPX* server board for clustering and storage applications.

The TR C4x/msd board features the new eight-core Intel® Xeon® processor D-1548 that has sufficient PCI Express* (PCIe*) lanes to construct a rugged microserver with five boards and up to 40 CPU cores. Compared to previous 3U VPX solutions, this board enables system designers to more than triple the number of cores while satisfying the same rigorous size, weight, and power (SWaP) constraints. The processor is suitable for both manned and unmanned vehicles (UAV/UGV/USV/UUV), where a relatively small weight decrease is enormously important. For instance, Concurrent Technologies is able to trim approximately 750 g (1.6 lbs) without sacrificing performance for systems that can now use one slot instead of two.

The number of CPU cores per TR C4x/msd board will further double (to 80 or more) in early 2016 when 16-core versions of the Intel® Xeon® processor D-1500 product family become available. This solution brief describes the TR C4x/msd board and the key capabilities Intel and Concurrent Technologies bring to military C4ISR system design.

Exceptional Performance-per-Watt for Military Systems

The Intel® Xeon® processor D-1548 was designed for dense and ruggedized systems requiring high performance, low power, and long-life support. In military systems where space is a premium, integrated Platform Controller Hub (PCH) technology and Intel® Ethernet in a ball grid array (BGA) package enable an exceptional level of design simplicity. The Intel® Xeon® processor D-1500 product family represents up to 2.3x greater performance per watt than the Intel® Core™ i7-4700EQ processor, which is a popular choice for VPX* board designers. The new processor family offers a scalable product lineup spanning two to eight cores, support for up to 128 GB of high-speed DDR4 memory, up to 12 MB of last-level cache, and two integrated ports of 10 Gb Intel Ethernet for ultra-fast connectivity.



Table of Contents

More CPU Cores, Same SWaP. 1
 Exceptional Performance-per-Watt for Military Systems 1
 Consolidating Complex Applications 2
 SWaP 3
 Non-transparent PCI Express* Bridging 4
 Virtualization Technology 4
 Compute-Intensive Military Applications 6

Consolidating Complex Applications

As the sophistication and complexity of military systems increase, equipment manufacturers need to deliver more computing performance in the same SWaP envelope for applications such as:

- C4ISR systems
- Synthetic aperture radar (SAR)
- Light detection and ranging (LiDAR)
- Multispectral sensor integration
- Image or signal processing
- Electronics warfare

For example, some of these applications are used by the U.S. Army Virtual Cockpit Optimization Programs (VCOP) to provide situational awareness, sensor imagery, flight data, and battlefield information in a clear and intuitive manner.¹

Processor Migration

Addressing this requirement, Concurrent Technologies greatly boosted the performance of its VPX board design by migrating from the Intel® Core™ i7-4700EQ processor to the Intel Xeon processor D-1548. This is demonstrated in Figure 1, which shows a 45 percent

performance increase using the PassMark* CPU benchmark.² Moreover, the Intel® Xeon® processor accomplishes this without increasing the board's size or thermal envelope.

Server Clusters

The transition to the Intel Xeon processor provides another performance improvement by way of enabling systems to cluster five boards (Figure 2) together, compared to three boards with the Intel® Core™ i7 processor. This represents 66 percent more computing resources and, when combined with a more powerful processor, yields a mathematical performance increase of approximately 2.4x. The Intel Xeon processor achieves higher cluster density by integrating a non-transparent PCIe bridge, which will be discussed in more detail.




Figure 2. The TR C4x/msd board supports clusters of five boards or more.

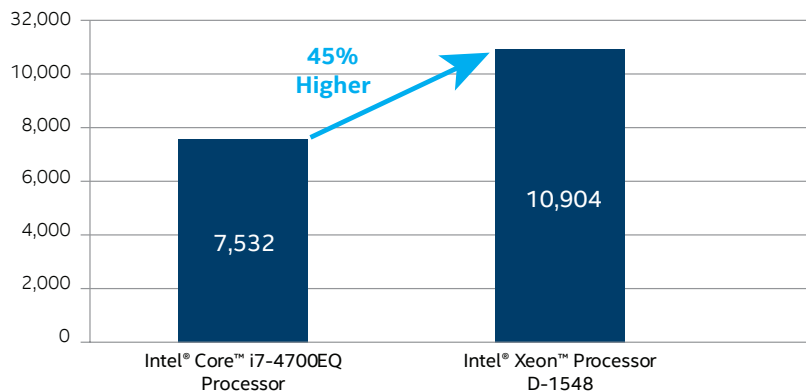


Figure 1. Intel® Xeon® processor D-1548 scores 45 percent higher than the Intel® Core™ i7 processor in PassMark* CPU benchmark tests.

Interconnect Fabric

For developers designing server clusters, Concurrent Technologies hides the complexity of the interconnect fabric and allows standard clustering applications to run "out of the box." This is achieved with Concurrent Technologies' Fabric Interconnect Networking Software (FIN-S)—a family of software packages that allows applications on multiple processor boards to communicate with each other. A high-performance, low-latency, message-based library enables direct, zero-copy, application-level communications between boards.

FIN-S provides an easy and high-level solution compared to a basic OS driver. It enables an application to see a cluster of nodes as though they were connected via an Ethernet socket-based structure. In other words, FIN-S makes the PCIe links between TR C4x/msd boards look like standard Ethernet sockets, greatly simplifying software development and enabling more deterministic communication with low latency, which can be vital for real-life applications in a military environment. This capability allows any standard, high-performance computer (HPC) application to see and make use of the total number of CPU cores in the cluster without modification.

To address different use cases, FIN-S is split into two functional packages: Embedded Clustering (EC) and Device Communication (DCOM). The EC package enables multiple Concurrent Technologies boards to communicate with each other over a socket interface to provide a clustering solution. The DCOM package enables point-to-point communication between third-party and Concurrent Technologies boards for solutions including FPGA/DSP/GPU content.




Figure 3. 3U VPX* conduction-cooled board.

SWaP

As a 3U VPX server board, the Concurrent Technologies TR C4x/msd is 3.9 inches by 6.3 inches (100mm x 160mm), which is about the size of a postcard. Figure 3 shows the aluminum metal plate that covers the computing system and conducts heat to the board's frame, allowing the module to be conduction cooled.

Even though the TR C4x/msd is small and light, it has a lot of memory, storage, and I/O capabilities, as shown in the block diagram in Figure 4. The board comes with up to 32 GB of DDR4 ECC DRAM and an optional on-board, 2.5-inch solid state drive (SSD) or a SATA-based Flash drive module. There are multiple SATA600 ports for storage and two built-in 10 Gb Ethernet channels for network connectivity, satisfying key attributes for high-performance, embedded, server grade applications.

The TR C4x/msd supports the VITA 46.11 specification for unified management within a VPX environment with two 10GBASE-KX4 data plane (VITA 46.7) and two 1000BASE-BX data plane (VITA 46.6) ports.

There is a mix of I/O interfaces, including up to 16 PCIe Gen 3 lanes, four USB ports, three RS232 ports, and general-purpose I/O, most of which are supported by the Intel Xeon processor D-1548 in system-on-chip (SoC) fashion. Intel offers embedded processors, like the Intel Xeon processor D-1548, for a seven-year extended supply life; in addition, Concurrent Technologies has standardized using long-life Intel® Ethernet controllers to help support long equipment lifecycles. At the board level, Concurrent Technologies offers up to 15 years lifecycle support, including extended manufacturing and repair elements. Furthermore, based on 30 years of experience designing embedded processor boards, Concurrent Technologies rates Intel as their most trustworthy and reliable supplier in terms of living up to lifecycle promises.

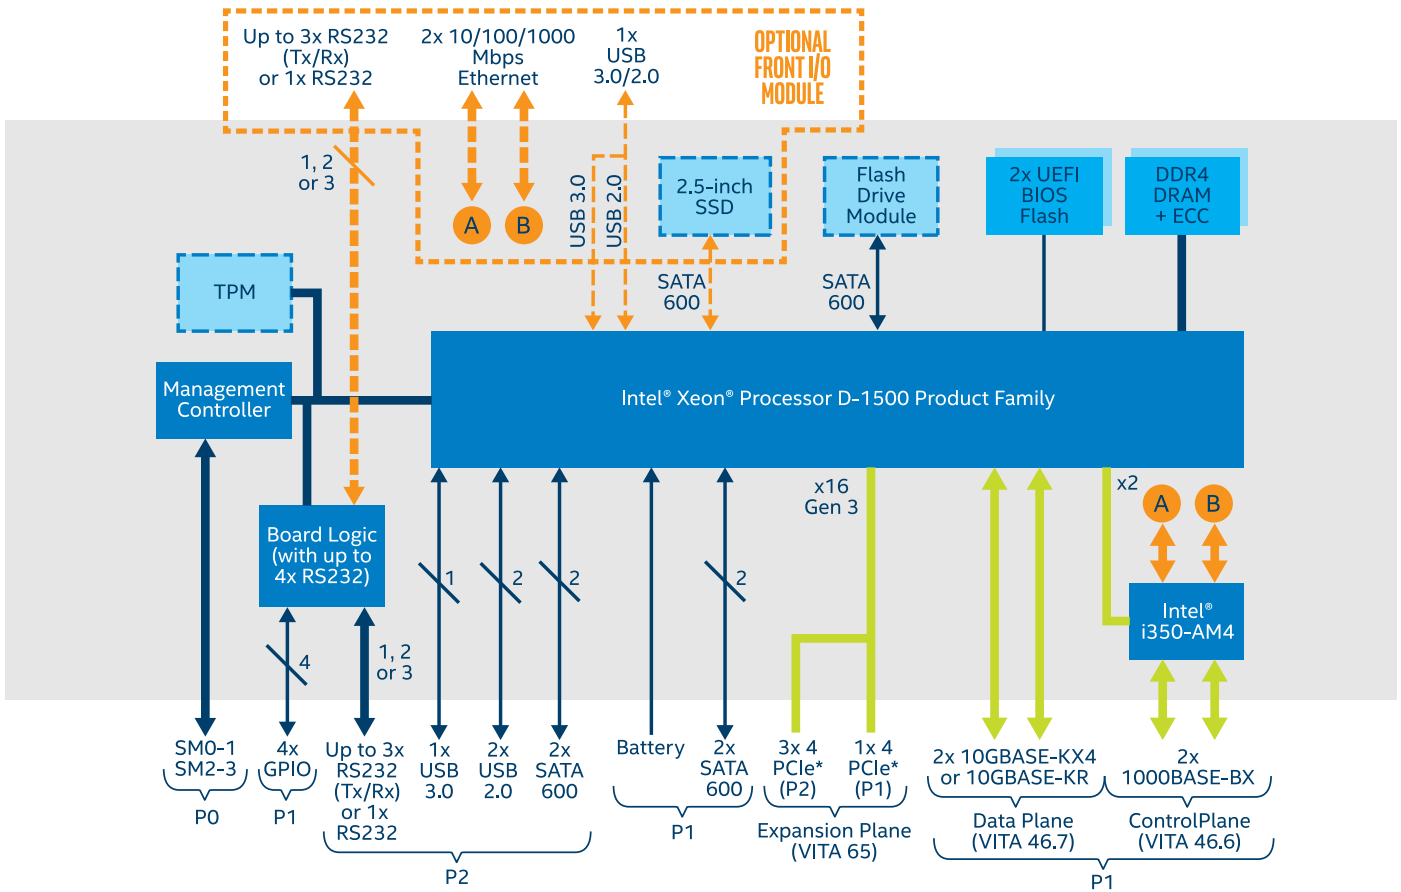


Figure 4. The TR C4x/msd board is rich in memory, storage, and I/O options.

Non-transparent PCI Express* Bridging

Critical for reducing the cost of a multi-processor embedded cluster is the non-transparent PCIe bridging in the Intel Xeon processor D-1500 product family. This capability allows small clusters to be constructed in a full mesh configuration with 3.9 GBps PCIe links between each TR C4x/msd board without using a switch board.

A non-transparent bridge is used in clusters for backplane networking and memory sharing between two systems. It allows one of the processors in the system to act as the root complex responsible for enumerating and configuring the system, and handling serious error conditions.

Each of the other processors has its own address-space layout and bus-address range, and the bridge performs PCIe address translation on bus messages as they come through the bridge.

The bridge adds address-domain isolation between the bus segments and translates addresses between the processor domains. Without this capability, the different processors would battle for control of the system.³

Virtualization Technology

Virtualization technology is a tool system that developers can use in a multitude of ways: partition workloads on a server, run different types of operating systems, simplify the migration of legacy applications, and so on. In a typical battlefield system, virtualization

allows real-time and general-purpose operating systems to run at the same time, which may be needed to support both time-critical (e.g., radar, communications) and non-time-critical applications (e.g., vehicle diagnostics), respectively.

From a performance perspective, virtualization introduces a small, but not insignificant, amount of latency; however, the Intel Xeon processor D-1500 product family helps minimize the performance impact by incorporating two technologies:

- Cache Monitoring Technology (CMT)
- Cache Allocation Technology (CAT)

CPU Basics

The performance of most CPUs is highly dependent upon the availability of data and instructions to the execution unit. To lower data latency, the execution unit is surrounded by small pieces of high-speed static RAM (SRAM) known as cache memory. This approach reduces the need for the CPU to fetch data from high-latency system memory (e.g., hard drives), thereby avoiding substantial delays.

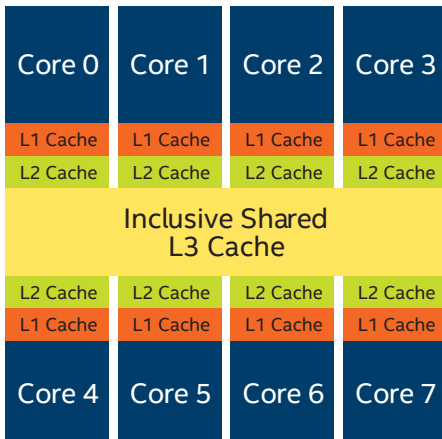


Figure 5. The L3 cache is shared by the CPU cores.

The Intel Xeon processor D-1500 product family implements a large shared last-level cache (Figure 5), which improves the performance of applications running in virtual machines (VMs) running on the processor. However, when VMs contend for L3 cache space, there could be a significant drop in performance and determinism. This can be avoided with these technologies:

Cache Monitoring Technology (CMT)

This technology allows an operating system or virtual machine monitor (VMM) to determine how much L3 cache each software thread (application) is using. This is valuable information because it identifies which applications are consuming large amounts of L3 cache and potentially degrading the performance of other applications. For example, Figure 6 shows a low-priority application (orange) that is using a lot of cache, making less cache available to speed up a higher-priority application (green).

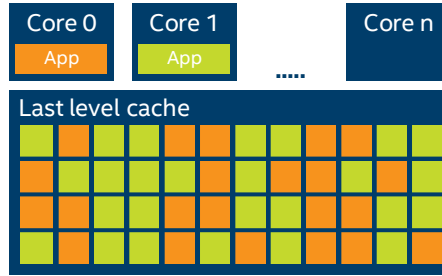


Figure 6. Cache Monitoring Technology (CMT) measures cache usage.

In the virtualized environment, where multiple VMs would be sharing this L3 cache, such “noisy neighbor” behavior can adversely affect the performance of other VMs, bringing down overall system performance.

There are several ways to take advantage of cache monitoring to optimize the overall system performance:

1. Move L3 cache-intensive applications to another socket (processor).
2. Schedule L3 cache-intensive applications when time-critical applications are not running, in order to optimize the amount of shared cache available to each application at any given time.
3. Generate performance histories to correlate available cache space and application performance. This information can be used to implement cache-aware scheduling to ensure applications have the necessary cache available to meet performance targets.

Cache Allocation Technology (CAT)

After measuring L3 cache usage at the application level with Cache Monitoring Technology (CMT), developers can use Cache Allocation Technology (CAT) to intelligently partition the L3 cache. This is shown in Figure 7, where a low priority-application (orange) is assigned a relatively small amount of L3 cache, giving a higher-priority application (green) access to more cache.

Developers can use cache allocation technology to increase determinism by prioritizing L3 cache access:

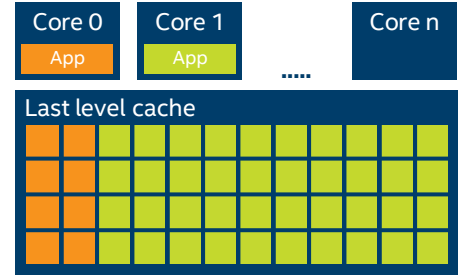


Figure 7. Cache Allocation Technology (CAT) assigns cache partitions to applications.

1. Assign high-priority applications enough dedicated L3 cache to avoid having their data and instructions evicted by other applications.
2. Isolate low-priority applications by limiting their access to L3 cache.
3. Avoid unnecessary cache evictions that reduce performance.

In a study by Wind River, CAT dramatically improved interrupt determinism, as seen in Figure 8. The left side shows interrupt latency without the technology ranging from 7 to 10 microseconds but, with the technology, the right side shows the interrupt latency for all samples was approximately 7 microseconds.⁴

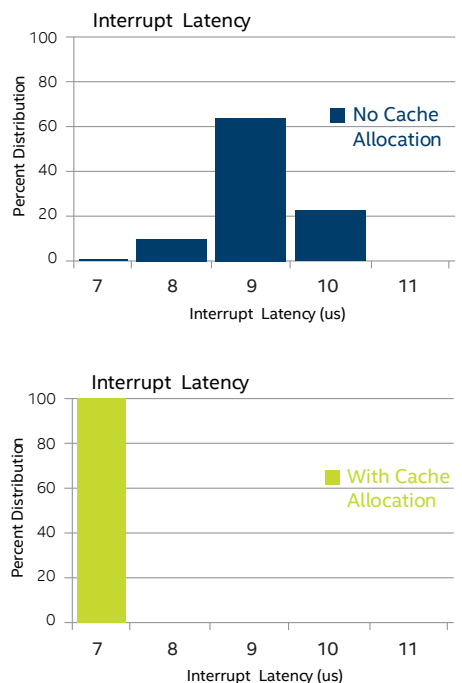


Figure 8. Cache Allocation Technology (CAT) improves interrupt latency.

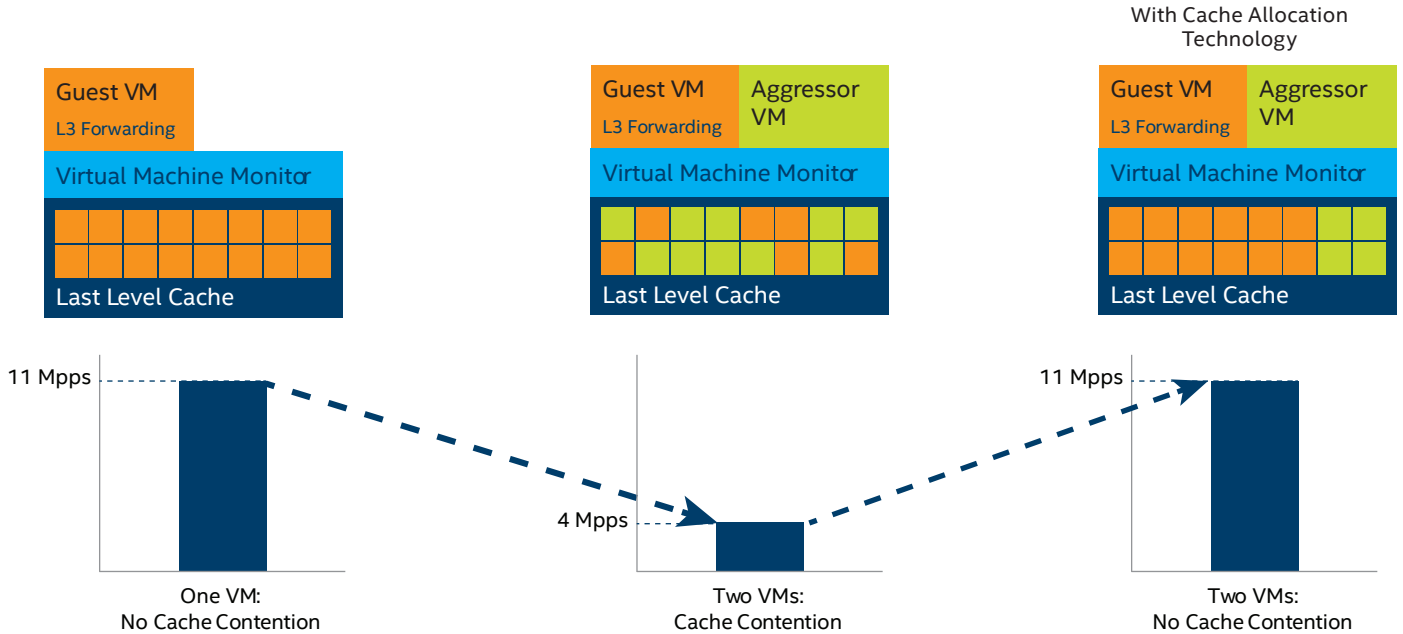


Figure 9. Packet processing determinism improves with Cache Allocation Technology.


The benefits of CAT can also be seen in a virtualized packet processing application. The QoS sample application implements a basic packet processing pipeline consisting of a packet classification and scheduling stage, which results in the selection of a high-priority or a low-priority queue. In this two 10 Gbps port configuration, the platform delivers 11 million packets per second (Mpps) of 64-byte packet throughput, as depicted in the leftmost pane of Figure 9.

When an aggressor VM is introduced (middle pane), which is a “noisy neighbor” because it takes over a substantial portion of L3 cache, packet-processing application performance drops to 4 Mpps. The rightmost pane depicts application of cache allocation, thereby limiting the aggressor VM’s access to L3 cache, after which the packet-processing application performance goes back to the original 11 Mpps.⁴

Compute-Intensive Military Applications

Manufacturers of battlefield military equipment are pressured to increase the computing power of their systems to improve situational awareness and a host of other tactical capabilities. Concurrent Technologies and Intel are helping the cause by enabling many more CPU cores within the same thermal envelope, compared to their previous 3U VPX processor boards. The TR C4x/msd board, based on the Intel Xeon processor D-1548, is an ideal choice for decision-making applications running in small and harsh environments.

For More Information
Learn more at intel.com/iot.



1. Edited by Michael J. Smith et al., “Usability Evaluation and Interface Design,” Lawrence Erlbaum Associates, 2001, www.crcpress.com/Usability-Evaluation-and-Interface-Design-Cognitive-Engineering-Intelligent/Smith-Koubek-Salvendy-Harris/9780805836073.
 2. Sources: www.cpubenchmark.net/cpu.php?cpu=Intel+Xeon+D-1540+%40+2.00GHz&id=2507 and www.cpubenchmark.net/cpu.php?cpu=Intel+Core+i7-4700EQ+%40+2.40GHz&id=1897.
 3. Larry Chisvin, “Non-transparent bridging allows multiprocessor design with PCI Express,” August 2, 2004, www.eetimes.com/document.asp?doc_id=1150824.
 4. Performance tests and ratings are measured using specific computer systems and/or components and reflect the approximate performance of Intel® products as measured by those tests. Any difference in system hardware or software design or configuration may affect actual performance. Buyers should consult other sources of information to evaluate the performance of systems or components they are considering purchasing. For more information on performance tests and on the performance of Intel products, visit Intel Performance Benchmark Limitations, intel.com/content/www/us/en/benchmarks/resources-benchmark-limitations.html.