



Intel Xeon E5-2600 v3 Codename Haswell-EP Launch



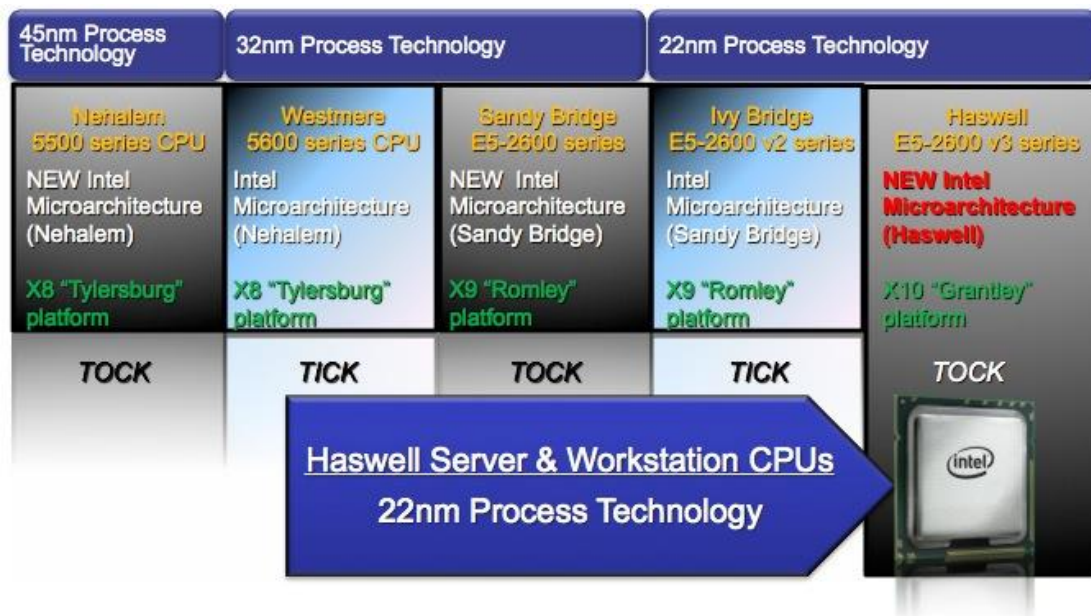
Introduction

It's that time of year again when Intel finally release their latest enterprise processor for the dual processor segment to the eagerly awaiting professional market. Following on in the traditional early September launch time frame, the [Xeon E5-2600 v3 processor series](#), codename Haswell-EP has been officially launched, finally allowing us at [Bostonlabs](#) to go through all the exciting details of the processors which we've been testing secretly in our labs for some time.



A Xeon E5-2600 v2 (left) pictured with the new Xeon E5-2600 v3 Series (right)

As a “Tock” in the Intel “Tick/Tock” release cycle, Haswell is an all new core microarchitecture, meaning that the core has been vastly redesigned and introduces many new features, yet stays on the same 22nm lithography as its predecessor – Ivy Bridge.



Intel®'s Tick / Tock roadmap to date from Nehalem to present

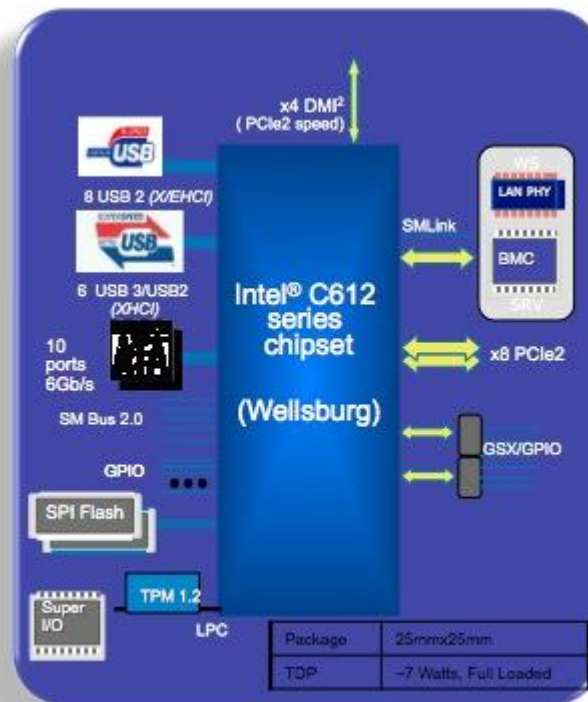
The Haswell microarchitecture itself, has actually been around for more than 12 months in the guise of the lower end Xeon E3-1200 v3 series processor, so the processor feature and instruction set is somewhat familiar, however Haswell-EP takes the changes that its entry level sibling delivered and expands on them significantly.

A new platform generation codename "Grantley" was also launched to compliment Haswell -EP, replacing "Romley". This brings in a new processor socket "Socket R3" which is similar to the outgoing "Socket R" and uses the same 2011 LGA pins, but these are laid out in a slightly different configuration to avoid confusion and accidental insertion of the older model. As a result, and due to other factors such as the newer memory controller found in Haswell-EP, the two platforms are not compatible and the E5-2600 v3 processor cannot be installed into an older "Romley" platform motherboard.



A Xeon E5-2600 v2 (left) pictured with the new Xeon E5-2600 v3 Series (right)

Additionally, “Grantley” brings in a new Platform Controller Hub (PCH); the C612 series, codenamed “Wellsburg”, which delivers well received new features over the previous C600 series, including up to 6 x USB 3.0 ports and up to a massive 10 x S-ATA 3 (6Gbit/s) ports



The Intel® C612 Chipset Block Diagram

Returning our focus back to the processor itself, there are numerous major enhancements to both the microarchitecture and the feature set. Below is a comparison of the E5-2600 v2 with the v3, detailing the major changes between these 2 generations.

	HASWELL-EP (Intel Xeon E5-2600 v3)	IVY BRIDGE-EP (Intel Xeon E5-2600 v2)
Number of Cores	18 cores TDP: Up to 145 W (Server); 160 W (WS)	12 cores TDP: Up to 130 W (Server); 150 W (WS)
Chipset	C612	C602
Socket	Socket - R3	Socket - R
Memory	4x DDR-4 channels Up to 2133MHz RDIMM, LRDIMM	4x DDR-3 channels Up to 1866MHz RDIMM, LRDIMM
QPI	2xQPI Up to 9.6 GT/s	2xQPI Up to 8 GT/s
PCIE	PCIe 3.0 40x PCIe	PCIe 3.0 40x PCIe
PCH Specifications	Up to 6x USB3 , 8x USB2 ports, 10x SATA3 ports	No USB3 Support 2x SATA3; 4x SATA2 USB 2.0

As you can see, one of the major enhancements is the increase in core count to 18, a huge leap of 50% over the previous generation. The addition of DDR-4 support with a frequency of up to 2133MHz – an increase of 14% in raw clock speed, alongside the 20% increase in QPI bandwidth, it's clear to see that the v3 series brings significant performance enhancements over the v2 generation.

Uncore Voltage / Frequency Scaling (UFS) enables the core and uncore (processor components not on the core but essential for processor performance) to be treated independently and run at different states. For example an LLC / Memory bound application no longer drives the main core frequency high for no reason, wasting energy.

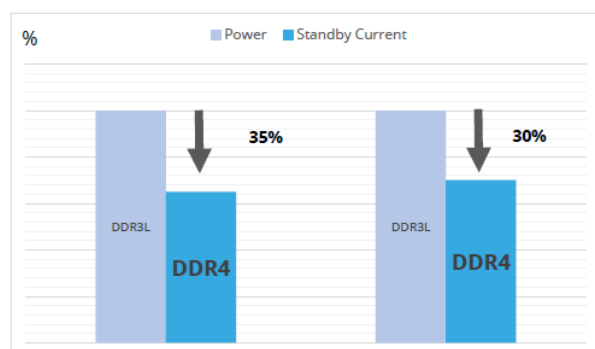


A DDR3 Module (top) compared to a DDR4 Module (below) Notice the elongated contacts in the centre

DDR4 not only brings an increase in performance from the increased frequency which it delivers to the platform, but also a decrease in power consumption through its lower operating voltage of just 1.2V.

DDR-4

- Up to 35% power savings compared to DDR-3L (2 DPC)
- Up to 50% at 3 DPC configurations
- 50% performance boost in bandwidth



This lower voltage can make as much as 4W difference per DIMM when measured at the wall, so this is not an insignificant change, given that almost all systems will ship with 8 DIMMS, making a 32W saving per system.

Additionally, there is a smaller decrease in speed when using multiple DIMMs per channel, resulting in higher performance for large memory deployments of up to 50% over DDR-3.

DIMMs Channel	/ DDR3 1.5v	DDR3 1.35v	DDR4 RDIMM	DDR4 LRDIMM
1	1866	1600	2133	2133
2	1600	1333	1866	2133
3	1066	800	1600	1600

Finally, but not least significantly, the inclusion of AVX 2.0 and the Haswell New Instructions (HNI) bring huge performance increases to the platform.

Intel sum these up in the following bullet points:

~10% higher IPC over IvyBridge Core (not counting new instructions)

Better branch prediction

Deeper buffers

Larger TLBs

More execution units

Improved front-end

Core improvements to feed FLOPs

2x L1 & L2 cache BW

Better misaligned memory operations (important for vectorization)

Increasing per core performance via power efficient features

Virtualization: Lower VMX round trip latency

Synchronization (multi-threaded/core scaling): Lower cache lock latency

Fused Multiply Add (FMA): 2x FLOPs/core vs. Sandy Bridge/Ivy Bridge

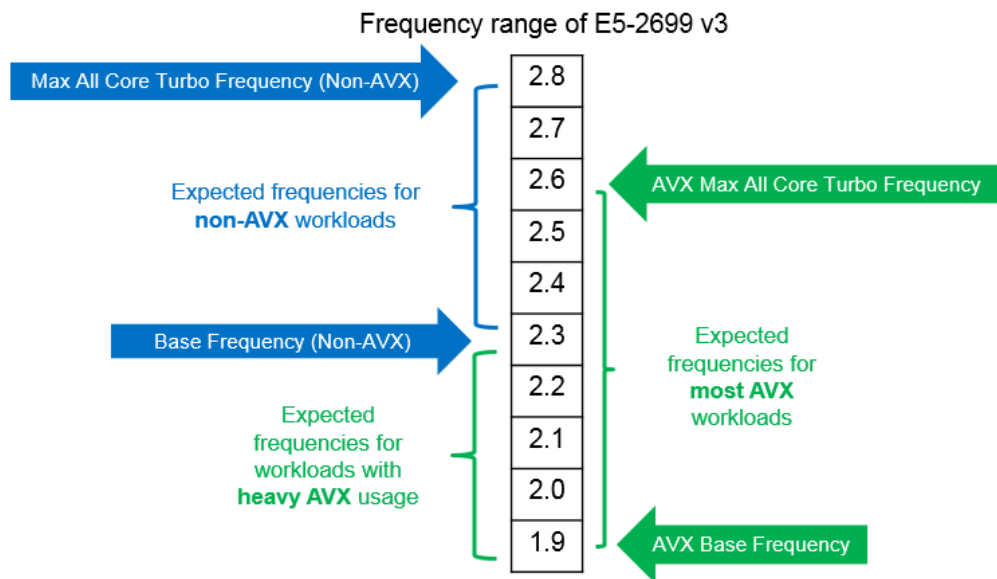
AVX-Integer: Extend 256-bit vector operations to include integer

Big Number Arithmetic (Crypto): Accelerating RSA and GMP (GNU Multiple Precision Arithmetic)

The key enhancement driving increase performance listed above is the introduction of the AVX 2.0 instruction set and its FMA feature to this processor architecture. This particular feature was actually already launched in the Xeon E3-1200 v3 series, and has been very successful in blazing a trail for new code in the ISV market.

As stated above, this delivers double the number of FLOPs per clock than in the previous generation processor architecture, increasing performance by as much as 100% where these new instructions apply. In testing for example, Intel's own optimised linpack benchmark was previously seeing high scores of sub 500 GFLOPs for the range topping processor in the v2 range; whereas with the v3 range, the E5-2699 v3 is able to deliver just over 1 TFLOP with current optimisations – a huge leap in performance.

There is one caveat to the delivery of AVX2 with this processor generation however, in that the AVX2 operational core frequency is actually lower than the standard core frequency due to increased thermal requirements of the relevant portions of the core. This is detailed in the below diagram, courtesy of Intel.



AVX Frequency Range Example – E5-2699 v3

Therefore calculating theoretical performance is a little more complicated than with previous generations as the frequency is less predictable.

Moving onto the SKU's which Intel launched with you can see that they've stayed with tradition and have a replacement model for most of the existing v2 SKU's, the majority of which have an increased core count or frequency over the outgoing models which enables them to typically perform better than the existing model.

There are still the same segments of processor to choose from, Basic, Standard, Advanced and Segment Optimised, with each layer adding more features like Turbo, Hyper Threading and an increasing clock speed or number of cores to enhance performance.

CPU Model	Cores		Frequency		Max Memory Frequency		L3 Cache		TDP	
	v2	v3	v2	v3	v2	v3	v2	v3	v2	v3
E5-2603	4	6	1.8GHz	1.6GHz	1333MHz	1600MHz	10MB	15MB	80w	85w
E5-2609	4	6	2.5GHz	1.9GHz	1333MHz	1600MHz	10MB	15MB	80w	85w
E5-2620	6	6	2.1GHz	2.4GHz	1600MHz	1866MHz	15MB	15MB	80w	85w
E5-2623	x	4	x	3.0GHz	x	2133MHz	x	10MB	x	105w
E5-2630	6	8	2.6GHz	2.4GHz	1600MHz	1866MHz	15MB	20MB	80w	85w
E5-2630L	6	8	2.4GHz	1.8GHz	1600MHz	1866MHz	15MB	20MB	60w	55w
E5-2637	4	4	3.5GHz	3.5GHz	1866MHz	2133MHz	15MB	15MB	130w	135w
E5-2640	8	8	2GHz	2.6GHz	1600MHz	1866MHz	20MB	20MB	95w	90w
E5-2643	6	6	3.5GHz	3.4GHz	1866MHz	2133MHz	25MB	20MB	130w	135w
E5-2650	8	10	2.6GHz	2.3GHz	1866MHz	2133MHz	20MB	20MB	95w	105w
E5-2650L	10	12	1.7GHz	1.8GHz	1600MHz	2133MHz	25MB	25MB	70w	65w
E5-2660	10	10	2.2GHz	2.6GHz	1866MHz	2133MHz	25MB	25MB	95w	105w
E5-2667	8	8	3.3GHz	3.2GHz	1866MHz	2133MHz	25MB	20MB	130w	135w
E5-2670	10	12	2.5GHz	2.3GHz	1866MHz	2133MHz	25MB	30MB	115w	120w
E5-2680	10	12	2.8GHz	2.5GHz	1866MHz	2133MHz	25MB	30MB	115w	120w
E5-2683	x	14	x	2GHz	x	2133MHz	x	35MB	x	120w
E5-2687W	8	10	3.4GHz	3.1GHz	1866MHz	2133MHz	20MB	25MB	150w	160w
E5-2690	10	12	3GHz	2.6GHz	1866MHz	2133MHz	25MB	30MB	130w	135w
E5-2695	12	14	2.4GHz	2.3GHz	1866MHz	2133MHz	30MB	35MB	115w	120w
E5-2697	12	14	2.7GHz	2.6GHz	1866MHz	2133MHz	30MB	35MB	130w	145w
E5-2698	x	16	x	2.3GHz	x	2133MHz	x	40MB	x	135w
E5-2699	x	18	x	2.3GHz	x	2133MHz	x	45MB	x	145w

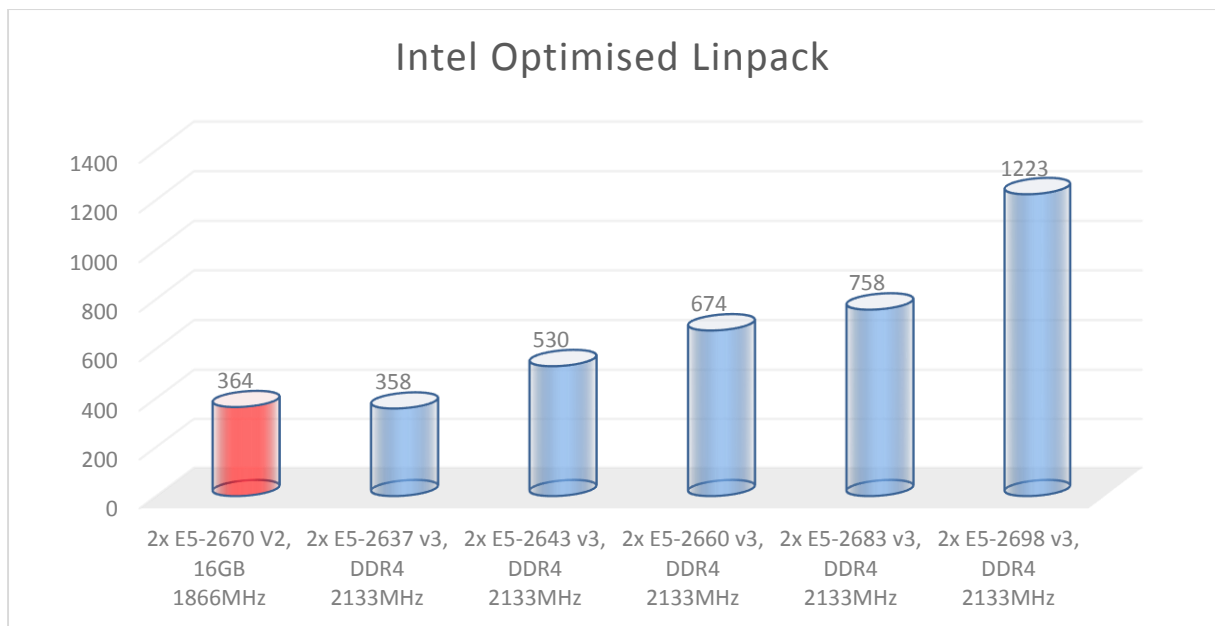


To give us an idea of how this new generation compares to the last generation in performance, we've run a few key benchmarks on a select set of processor models.

Intel Optimised Linpack – for Linux

Linpack is a well-known industry standard benchmark used to determine the floating point processing power of a system or a cluster of systems by asking it to solve complex mathematical problems.

Our test utilises a customised problem size to optimally use approx. 128GB of memory including the operating systems overhead to give optimum results.



Graph key for "Intel Optimised Linpack": Measured in GFLOPS – the higher the result, the better the performance. The results in blue are E5-2600 v3 series CPUs and red results are E5-2600 v2.

The trend from left to right shows a linear increase in GFLOPs across the range of processors sampled. The only acceptance to this is the 2670 v2 versus the 2637 v3. This at first might look surprising at first glance however there is a huge difference in core count here; the 2670 v2 has a mighty 10 cores and the 2637 v3 has only 4. As you can see the Haswell core actually increases Linpack performance significantly over the older Ivy Bridge. In fact, per thread (our Linpack is run with Hyperthreading enabled) we're seeing 18.2 GFLOPs for the E5-2670 v2 and a massive 44.75 GFLOPs for the E5-2637 v3.

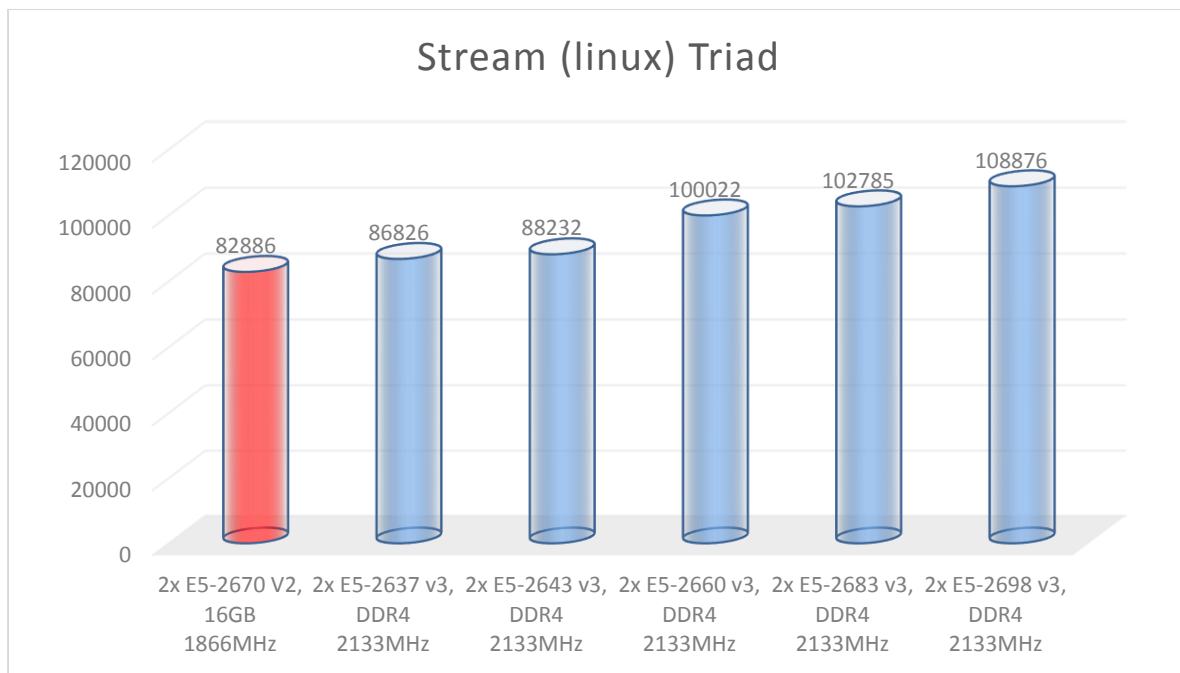
The gains shown in the graph and the breakdown above are brought about partly from the increases in core count, DDR generation 3 to 4, VRMs on die and register changes. AVX 2.0 however, plays a big part in the GFLOPs performance increase seen, and can result in a 100% improvement in some cases. This performance is not uniform across the board however, as applications have to be AVX 2.0 enabled to benefit. Clock for clock, the increase in performance between generations is somewhere between 15-20% if AVX 2.0 cannot be utilised.

Stream memory benchmark

STREAM benchmark is a simple synthetic benchmark program that measures sustainable memory bandwidth (in MB/s) and the corresponding computation rate for simple vector kernels. The benchmark carries out 4 operations, but we will concentrate on Triad.

Triad uses fused multiple-add operations (FMA) that are an important operation in many basic computations, such as dot products, matrix multiplication, polynomial evaluations, Newton's method for evaluation functions, and many DSP operations. This benchmark can be directly associated with application performance as a result. The FMA operation has its own instruction set now and is typically done in hardware and this is especially true with the Haswell New Instructions. Consequently, feeding such hardware operations with data can be extremely important – hence, the usefulness of the Triad memory bandwidth benchmark.

Our test utilises a problem size to optimally use approx. 128GB of memory, including the operating systems overhead and we disable hyper-threading to give the best performance output possible.



Graph key for "Stream memory benchmark Copy/Scale/Add/Triad": Measured in MB/s – the higher the result, the better the performance. The results in blue are E5-2600 v3 series CPUs and red results are E5-2600 v2.

Stream performance summary

The triad results show left to right what we expect with a generation change in DDR from three to four. The E5-2600 v2 generation is represented here in red by the E5-2670 v2 sporting DDR-3 at 1866MHz and shows a performance of 82886MB/sec. The V3 range with DDR-4 of 2133MHz shows performance between 86000 to 108000MB/sec depending on the model.

Looking at similar models of processors, the percentage difference between the E5-2660 v3 and the E5-2670 v2 is a respectable 20%. If memory performance is a major attributing factor towards your application performance then making the jump to the v3 platform with the new memory controller plus DDR4 DIMMS provide a non-insignificant increase in performance. Not to mention one that is delivered with lower power consumption, making a worthwhile investment. With DDR-4 prices somewhat higher than DDR-3 right now however, we might have to wait for up to 1 year to get this benefit at a level cost.

Prior to the launch of the Xeon E5-2600 v3 series processor, the team at Boston have been busy in the labs testing and validating a whole series of validated solutions, all available for immediate release.

We've even been working on a range of solutions designed to cool the top TDP processor from the range, the Xeon E5-2687W v3 at a staggering 160W each. Cooling such a processor is obviously a concern, so Intel deem it only suitable for workstation use, where there is plenty of room for cooling and larger heat sinks. At Boston, we already have a SKU with liquid cooling ready and validated for the E5-2687W v3, the Boston Venom 2401-12T.

This full tower system designed for creative professionals also sports a Quadro K6000 graphics card, up to 1TB of DDR-4 memory, a Blu-ray rewriter drive and up to 6 x S-ATA HDD's or SSD's.



The Boston Venom 2401-12T and its liquid cooled E5-2687W v3 processors

In traditional tier one designed, a 1U rack mount chassis will be restricted to 135W TDP processors, and up to 145W TDP processors in a more spacious 2U enclosure. This means that if you want to use 2 of the range topping Xeon E5-2699 v3 processor

with 18 cores each, you will need to consider a 2U or above chassis. However, here at Boston we have new range of server solutions available with data centre optimised cooling enabling us to use the range topping 160W TDP models in a 1U or 2U form factor. The clever cooling design features the processors situated side by side; instead of inline. This helps decrease the temperature of the 2nd CPU by up to 10 degrees Celsius, prolonging the CPU lifetime, lowering fan speed, lowering noise and increasing turbo performance.

Another example solution released today is the Boston Value Series VS360p. This has support for 2 x Xeon E5-2600 v3 processors of up to 145W TDP, 1TB of DDR-4 Registered memory, 8 x SAS 3 2.5" HDD's or SSD's and additionally has 2 x 2.5" NVME SSD bays, all within a space saving 1U package.



Right from launch, Boston have a range of solutions available from launch which are Xeon E5-2600 v3 series ready and have been thoroughly tested and validated with the latest operating systems.

For more details on the range, please download our exclusive X10 solutions catalogue from our website at the following [here](#).

If you would like some more information or a custom quotation please contact us on +44(0)1727 876 100 or email sales@boston.co.uk to speak to one of our technical account managers who will gladly make sure you get the best personalised solution for you.

Stay tuned to bostonlabs.co.uk for our next article in the coming days where we put the Xeon E5-2600 v3 series through its paces with our standard benchmark suite and see how it compares to the v2.



BOSTON

Servers • Storage • Solutions

Powered By



SanDisk



UK Headquarters

Tel: +44 (0) 1727 876 100

E-mail: sales@boston.co.uk

www.boston.co.uk

India

Tel: +91 22 6700 0897

E-mail: sales@bostonindia.in

www.bostonindia.in

Germany

Tel: +49 (0) 89 9090 199-3

E-mail: info@bostondeutschland.de

www.bostondeutschland.de