Business Brief

Enterprise & Government
Hybrid cloud

Solving today's critical business challenges with cloud infrastructure

Deliver digital transformation in uncertain times with distributed cloud computing architectures.

Scale today's and tomorrow's workloads with Intel's trusted technology portfolio, optimized for diverse cloud environments and supported by the industry's broadest ecosystem.

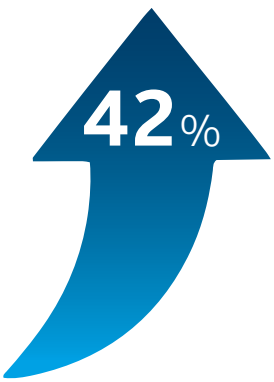

2nd generation Intel® Xeon® Scalable processors deliver an average of **42%** workload-optimized price/performance improvement compared to previous generation[1]

These are times of intense change for businesses, with new working models and economic uncertainty pushing digital transformation to the forefront of the enterprise agenda. Navigating this changing environment requires agile, trusted and scalable solutions to ensure business continuity and long-term success. So, it's not surprising that these challenges are driving the adoption and integration of hybrid and multi-cloud architectures – distributed cloud computing. In this paper, we will explore how Intel® technologies can help organizations meet their short-term, longer-term and persistent challenges.

Respond to evolving business needs today and tomorrow

Scaling your workloads from edge to cloud, securely and reliably while unlocking the full potential of the cloud, is a key route to accelerating transformation projects successfully. Intel's data-centric technology portfolio provides a trusted technology foundation for smooth transitions and migrations, with improved operational efficiency to help you meet these goals.

Delivering performance on your terms
Business needs are evolving fast, and enterprise infrastructure must be able to respond at equal speed. This might mean scaling quickly on-premises or optimizing workload performance in the cloud to meet changing demands. From on-site data centers, to the cloud and out to the edge, 2nd generation Intel® Xeon® Scalable processors are the most tested, performance-proven, and globally available solutions. This makes them ideal for providing the responsiveness that today's business environment requires. With workload-optimized performance, you can achieve on average 42 percent price/performance improvement generation over generation[1], ensuring optimal performance with improved operational efficiencies across your data center and cloud services investments.

Connecting people, fast
The ability to support thousands of remote workers is an immediate concern for many enterprises. Using Intel® Ethernet 700 or Intel® Ethernet® 800 series paired with 2nd gen Intel Xeon Scalable processors running Single Root-IOV (SR-IOV) optimized with Data Plane Development Kit (DPDK), you can deliver the network performance and scale necessary to support more than 100,000 remote workers using a VPN (see Box out 1). Intel Ethernet 700 series gives you the agility to optimize network performance, to respond to changing network requirements, and enable new services as needed.

Minimize costs while delivering on service-level agreements (SLAs)
Ensuring optimum performance of the systems your customers and employees use has never been so critical to business success – and with economic uncertainty, doing so cost efficiently is paramount. Improving operational efficiency is one route to minimizing costs while continuing to deliver on strict service-level agreements (SLAs). With Intel® Optane™ persistent memory and Intel® Xeon® Scalable processors, you can do more with less. Together they allow you to improve virtual

machine (VM) density and reduce latency per core, creating the highly efficient cloud and data center infrastructure services that today's environment demands. With the Intel® Xeon® Platinum 8260 processor, you can achieve up to 35 percent more VMs at up to 27 percent lower cost[2]. Your storage choices are another important factor governing the performance and efficiency and of business-critical apps in your data center and in the cloud. With Intel® Optane™ SSDs you can Increase VM density up to 60 percent and reduce cost up to 25 percent on VMware vSAN[3] delivering the performance these workloads demand while lowering overall costs.

Infrastructure challenges need agile solutions

With unprecedented pressure to adapt, teams are struggling to deliver smooth, easy, cost-efficient migrations. It's here that Intel's data-centric portfolio can help, as it is built on decades of optimizations from system to application level.

Hardware and software optimizations embedded into the technology enable you to adapt to changing business priorities quickly. Intel engineers work with thousands of independent software vendors (ISVs) and open source projects to verify, test, and optimize applications for Intel Xeon Scalable processors delivering over five generations

How Intel IT created 100,000 mobile workers over a weekend

Like many organizations, Intel has had a remote work capability in place for years, but the Covid-19 pandemic resulted in more than 100,000 employees working remotely.

Intel IT was already prepared for half of Intel's workforce to be remote at any one time – but suddenly, it had to double this capability in a single weekend.

Rather than the traditional approach of adding more capacity on-premises, taking up to three months, Intel IT moved the most bandwidth-intensive workloads, such as audio and video conferencing, to the cloud. With Intel's multi-cloud strategy, the company quickly scaled up its use of as-a-service offerings so employees could easily access tools and capabilities that did not require a VPN.

Intel deployed additional VPN appliances globally on premises and in public cloud environments to provide burst capacity to handle the new demand. To ensure a quality user experience for remote workers, the team focused on keeping latency consistently low while scaling up the number of VPN users. The team investigated several technology options with their public cloud provider and narrowed down their solution. Using Intel's broad technology portfolio, Intel® **Ethernet 700 series** with Single Root-IOV (SR-IOV), **Intel® Xeon® Scalable processors**, and the **Data Plane Development Kit** (DPDK), the team was able to optimize their solution to achieve the required performance both on premises and in the cloud. SR-IOV enables data to bypass the software switch in the virtual machine manager and go directly to and from the virtual machine (VM).

Intel IT was able to add 50 percent more users per VPN server without any increase in latency.

A key discovery for the team was that they could not scale the network/VPN capacity and performance needed without Intel® network interface cards (NICs). Intel IT replaced third-party NICs on-premises with Intel Ethernet 700 series with SR-IOV and requested Intel Ethernet 700 series public cloud instances from their cloud service provider (CSP).

The combination of Intel® technology on premises and in the cloud meant that by Monday morning, Intel employees around the globe were productively working from home.

About SR-IOV: SR-IOV provides a virtual function (VF), which is a lightweight virtual Peripheral Component Interconnect express (PCIe) device pointer unique to that VM. By accessing this unique handle, the VM will communicate directly with the PCIe NIC on its own channel. The NIC will then directly sort incoming packets to the memory address of the virtual NIC of the VM, without having to communicate with the hypervisor at all. **More information**.

of compatibility and scale. Key parts of the cloud stack are also uniquely optimized for Intel Xeon Scalable processors, creating smoother transitions for your workloads from edge to cloud.

- **Artificial intelligence (AI):** Intel enables future-ready AI frameworks to support a wide range of popular AI applications. Intel® Deep Learning Boost (Intel® DL Boost) brings new embedded performance acceleration for AI workloads in the 2nd gen Intel Xeon Scalable processor, with up to 30x performance improvement for inference workloads compared to the previous generation[4].
- **Workload performance acceleration**: Unique acceleration technologies from Intel can deliver performance gains. Intel® Advanced Vector Extensions 512 (Intel® AVX-512) on Intel Xeon Scalable processors, for example, delivers up to 68 percent performance gain in high-performance computing (HPC) financial services and insurance (FSI) applications[5].
- **Infrastructure virtualization:** Intel® Virtualization Technology (Intel® VT) built into Intel Xeon Scalable processors delivers accelerated performance for virtual environments, with live application migration on-premises or to the cloud. Seamlessly migrate your VMware vSAN environments into the public cloud or to another private cloud with no downtime, using VMware vMotion and Intel Xeon Scalable processors with built-in Intel VT.
- **Analytics & database:** There is a long history of database applications being optimized for Intel platforms. For example, with SAP applications hosted on Intel technology-based cloud infrastructure, SAP customers can improve run times and significantly boost memory performance. Only Intel Xeon Scalable processors are certified for SAP on infrastructure-as-a-service (IaaS) instances[6].

Scale your business with confidence

As your infrastructure scales, so do security threats and business risk. Securing your infrastructure begins with the silicon. Intel delivers unique acceleration technologies to help keep data protected at rest and in transit without disrupting performance. Intel® Software Guard Extensions (Intel® SGX) delivers hardware-based isolation and memory encryption that provide enhanced protection. Intel® Key Protection Technology (Intel® KPT) in Intel® QuickAssist Technology (Intel® QAT) helps secure keys in the hardware to improve performance. Intel AVX-512 delivers cryptographic performance for data at rest and in transit. Intel is a founding member of the Confidential Computing Consortium, a cross-industry effort to define and develop the enterprise-grade building blocks, foundational services and frameworks needed to better protect workloads with extremely sensitive data in the cloud.

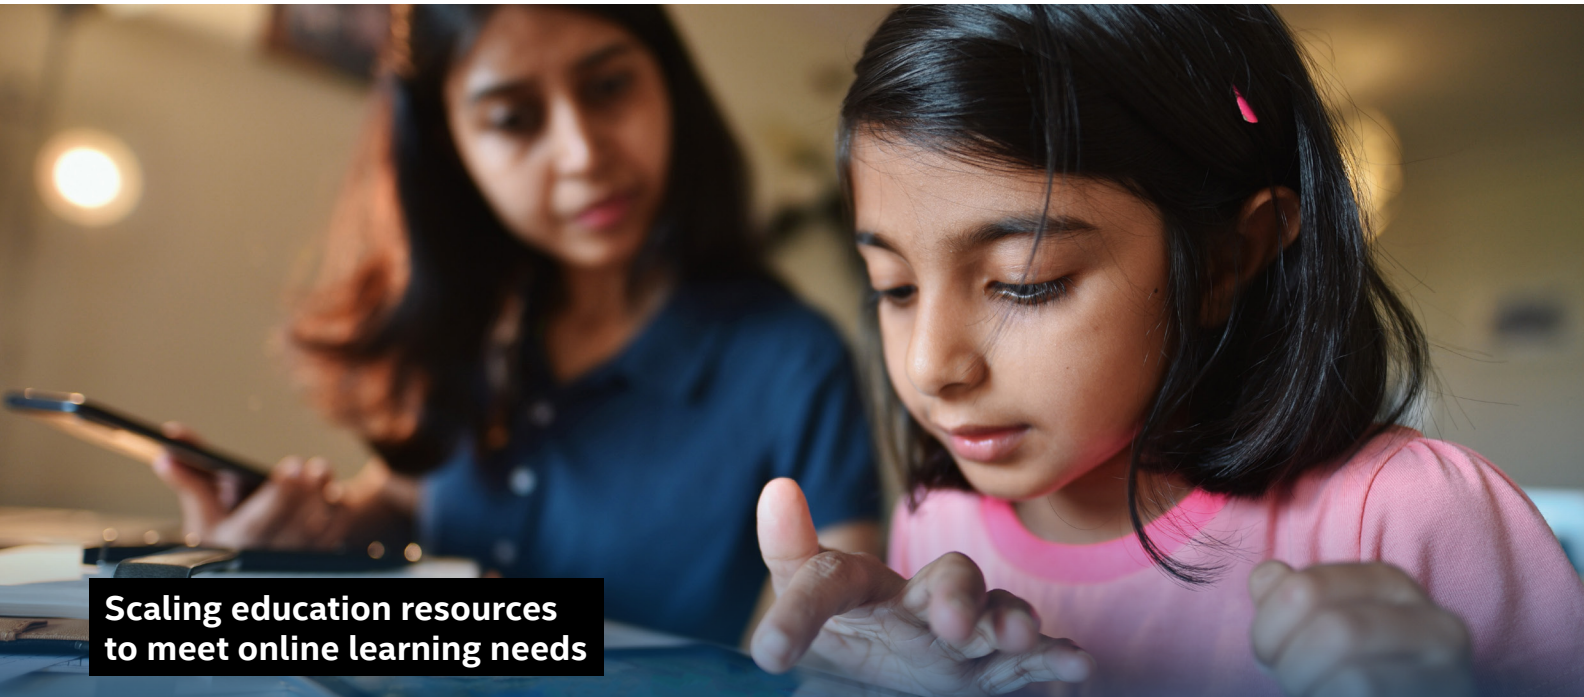

Scaling education resources to meet online learning needs

With schools closed throughout India due to Covid-19, and the clock ticking on the academic year, Intel teamed up with Amazon Web Services (AWS) and CareerLauncher to deliver a feature-rich cloud-based educational portal in just two weeks.

CareerLauncher works directly with educational institutions to provide distance learning through its portal, Aspiration.ai. It needed to create virtual classrooms for tens of thousands of new students, fast. Aspiration.ai is built on AWS cloud instances with artificial intelligence (AI) acceleration provided by Intel® technology. This meant that CareerLauncher could scale up quickly and seamlessly to support a much larger group of users with an impressive array of AI-powered services.

CareerLauncher enlisted the help of Intel and AWS, who created Project Aspiration 2020 to facilitate the implementation. Working together with the state government, the team established the cloud computing and storage resources required to create the educational experience that would support students and their teachers. With AWS and Intel at its side, CareerLauncher was able to offer live streaming of classes, archive and replay of lessons, learning games to motivate students, and advanced analytics to help teachers gauge student success.

In just two weeks, CareerLauncher trained hundreds of teachers and administrative staff to use its platform. It is now serving 300 teachers and 160,000 students across 100 schools – with plans to support over 1 million students in future.

Summary

Cloud computing and cloud transformation are a top priority for enterprise architects and technology innovation leaders. Whether you are looking for new ways to exploit the evolving models of cloud computing for business success or scaling infrastructure to meet the ongoing demands of your business, Intel is your cloud computing technology partner. Through a broad data-centric technology portfolio, long-term co-engineering relationships with ISVs, original equipment manufacturers (OEMs), and CSPs, Intel technologies will help you move faster, store more, and process everything. Get started on your cloud transformation journey today.

Learn More

- **Learn:** The Six M's of Cloud Transformation
- **See solution options from Intel's ecosystem:** VxRail Launch with Intel® Optane™ persistent memory
- **Learn more about cloud technologies at the:** Intel Cloud Hub

Accelerate time to market with Intel® Select Solutions

Bring all the advantages of cloud services to your on-premises data center in a simple and accelerated deployment option. Choose verified solutions that offer tuned and optimized infrastructure for private cloud strategies at premium performance with software-defined compute, storage, and networking capabilities.

- **Intel® Select Solution for Google Cloud's Anthos** allows you to quickly and easily deploy and manage apps and workloads using virtual machines (VMs) or Kubernetes containers across your on-premises environment, on your hybrid cloud, or with your preferred cloud provider.
- **Intel® Select Solution for Microsoft Azure Stack HCI** provides a simplified, low-cost, hyperconverged infrastructure tailored to the different compute, memory, and storage needs at the edge, in the data center, and on your Microsoft Azure public cloud.
- **Intel® Select Solution for VMware vSAN** combines optimized Intel® hardware components to enable organizations to quickly deploy reliable, comprehensive VMware vSAN built on a performance-optimized vSphere integrated infrastructure.
- **Intel® Select Solutions for Open Cloud** delivers a full-stack solution with OpenStack, Ceph, Kubernetes, and other open source infrastructure software as core components.

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors.

Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more complete information visit www.intel.com/benchmarks.

[1] For configuration details, visit www.intel.com/benchmarks (Intel® Xeon® Scalable Processors - claim #35).

[2] For configuration details, visit www.intel.com/benchmarks (Intel® Xeon® Scalable Processors - claim #20).

[3] Source – Intel. Estimated HW,SW,MEDIA,MAINT costs = $255,000 ; Estimated power & infrastructure costs = $15,000. Source: Intel. Power, Cooling, Consolidation savings configurations based on 1) HDD: 7.2K RPM 4TB HDD, AFR of 2.00% and 7.7W active power, 24 drives in 2U (1971W total power) https://www.seagate.com/files/www-content/datasheets/pdfs/exos-7-e8-data-sheet-DS1957-1-1709US-en_US.pdf. 2) Intel® SSD D5-P4326: 20W active power, 0.44% AFR, 64 drives in 2U (1280W total power) https://ark.intel.com/content/www/us/en/ark/products/186675/intel-ssd-d5-p4326-series-15-36tb-e1-l-pcie-3-1-x4-3d2-qlc.html; Cooling cost based on deployment term of 5 years with Kwh cost of $0.158 and number of watts to cool 1 watt 1.20. Based on 24 3.5" HDDs in 2U and 64 E1.L SSDs in 2U. Hybrid storage based off using Intel TLC SSD for cache.

[4] For configuration details, visit www.intel.com/benchmarks (Intel® Xeon® Scalable Processors - claim #3).

[5] FSI Kernels- Baseline: Intel® Xeon® Platinum 8268 processor configuration: Intel "Wolf Pass" platform with 2-socket Intel® Xeon® Platinum 8268 processors (2.9GHz, 24C), 12x16GB DDR4-2933, 1 SSD, BIOS: SE5C620.86B.02.01.0008.031920191559; Microcode: 0x500001c, Red Hat Enterprise Linux* 7.7, kernel 3.10.0-1062.1.1. FSI Kernels v2.0: Geomean (3 workloads: Binomial Options, Black Scholes, Monte Carlo), AVX2_256 build, Intel® Compiler 2019u5, Intel® Math Kernel Library (Intel® MKL) 2019u5, BIOS: Binomial (HT ON, Turbo ON, SNC OFF, 2 threads/core), Black Scholes (HT OFF, Turbo ON, SNC OFF, 1 threads/core), Monte Carlo (HT ON, Turbo ON, SNC OFF, 2 threads/core). Test by Intel as of 11/1/2019, w/AVX-512: Intel® Xeon® Platinum 8268 processor configuration: Intel "Wolf Pass" platform with 2-socket Intel® Xeon® Platinum 8268 processors (2.9GHz, 24C), 12x16GB DDR4-2933, 1 SSD, BIOS: SE5C620.86B.02.01.0008.031920191559; Microcode: 0x500001c, Red Hat Enterprise Linux* 7.7, kernel 3.10.0-1062.1.1. FSI Kernels v2.0: Geomean (3 workloads: Binomial Options, Black Scholes, Monte Carlo), AVX-512 build, Intel® Compiler 2019u5, Intel® Math Kernel Library (Intel® MKL) 2019u5, BIOS: BIOS: Binomial (HT ON, Turbo ON, SNC OFF, 2 threads/core), Black Scholes (HT OFF, Turbo ON, SNC OFF, 1 threads/core), Monte Carlo (HT ON, Turbo ON, SNC OFF, 2 threads/core). Test by Intel as of 11/1/2019

[6] https://www.sap.com/dmc/exp/2014-09-02-hana-hardware/enEN/iaas.html

Performance results are based on testing as of dates shown in configurations and may not reflect all publicly available updates. See backup for configuration details. No product or component can be absolutely secure.

Your costs and results may vary.

Intel technologies may require enabled hardware, software or service activation.

Intel does not control or audit third-party data. You should consult other sources to evaluate accuracy.