White Paper

Top Benefits of Intel® Technology in the Cloud

Consistent, predictable performance; a broad, optimized software ecosystem; best performance-per-dollar; and wide availability across multiple cloud service providers are compelling reasons to use Intel® technology for your cloud workloads.

Table of Contents

Reason #1: Experts in cloud computing optimizations for enterprise workloads 2

Reason #2: Greater choice and mobility to accelerate cloud deployments – public, private, hybrid and multi-cloud 2

Reason #3: Proven consistent performance from data center to cloud to edge 2

Reason #4: Half the cores for equal performance with lowest total costs 3

Reason #5: Most trusted hardware for cloud-based enterprise workloads 4

Reason #6: Easiest manageability for cloud environments 5

Waves of digital transformation continue to sweep across markets and enterprises. Cloud-based IT strategies promise gains in agility, security and efficiency. These gains come at the cost of growing complexity, as enterprise cloud architectures span multiple vendors and clouds, and integrate with a proliferation of edge devices. The challenges of workload placement and migration across an increasingly distributed multi-cloud infrastructure are getting more complex to solve. Workloads, too, are becoming more complex, ranging from new cloud-native apps to artificial intelligence (AI) to established services.

As cloud adoption grows, enterprises seeking to gain the benefits of cloud computing are asking important questions. How can we seamlessly migrate our workloads? Will we have to retest everything or refactor a lot of code? How can we be sure that infrastructure can meet the needs of our users or application demands? What about long-term flexibility, costs and scalability? The answers to these questions have everything to do with the underlying cloud infrastructure services—both hardware and software—and the workloads being considered.

Every application and workload has unique infrastructure requirements, and these requirements evolve over time. For example, the performance of workloads such as scientific modeling, analytics, security. and data compression can be accelerated with the Intel® Advanced Vector Extensions 512 (Intel® AVX-512) instruction set. AI workload performance can be significantly improved with Intel® Deep Learning Boost (Intel® DL Boost). Rooted in silicon, Intel® hardware-based security creates a trusted foundation for protecting data in all its phases: at rest, in flight and in use. Working with top independent software vendors (ISVs) and customers optimizing applications running in the cloud, Intel® technology-based solutions help reduce risk and complexity, contributes to optimal performance and improves operational efficiencies across the data center and cloud services.

Intel technology solutions are solving real pain points and business challenges on-premises, through communication networks, at the edge and in the public cloud. This pervasiveness enables enterprises to unlock new business opportunities, accelerate and modernize their business or migrate to the public cloud with ease. For example, you may be moving a workload from on-premises to the public cloud, or from one public cloud service provider (CSP) instance to a newer instance or to expand to more regions. In each scenario, the same optimized, tested and validated Intel® architecture eases migration headaches and delivers predictable, consistent application performance.

Decades of innovation and optimization with the world's leading technology and service providers means that Intel not only has a deep understanding of enterprise use cases and needs, but also a fine-tuned, globally deployed

architecture you can count on. Virtually all of the major CSPs offer a wide variety of workload-tuned instances featuring Intel® Xeon® Scalable Processors and other Intel technologies across their regions and availability zones. That's why each day the world's clouds runs on Intel. We'll be here for your needs today and deliver the solutions you need to build your business tomorrow.

Reason #1: Experts in cloud computing optimizations for enterprise workloads

Just as all workloads are different, so is each business. Innovation is accelerating at a rapid pace and business-critical enterprise workloads like VMware, SAP, Microsoft SQL Server, Salesforce, Oracle and data analytics are migrating to the cloud. Gartner predicts that worldwide public cloud spending by end-users will grow 18 percent from 2020 to 2021.[1] (In contrast, the global data center IT expenditure is expected to grow only 6 percent in the same time frame.[2]) Through long-standing relationships with top ISVs, original equipment manufacturers (OEMs) and CSPs, Intel collaborates daily with the cloud ecosystem, developing and optimizing Intel technology to run with the best performance-per-dollar. This includes collaborating with CSPs to develop purpose-built instances, such as Microsoft Azure DCsv2 Virtual Machines, designed to better support demanding tasks and workloads such as encryption and decryption of data, in-memory databases and data analytics.

Intel Xeon Scalable Processors provide leadership performance on latency-sensitive workloads such as database, edge and eCommerce applications to meet stringent customer requirements and service-level agreements. These processors are the most widely deployed computing engine in public cloud services, providing agility and flexibility across all major CSPs. Across the top global CSPs, more than 83 percent of all cloud instances are powered by Intel technology.[3] Plus, you can choose from hundreds of different Intel technology-based, workload-tuned instance types tailored to fit different use cases and data sizes. Many of these instances offer access to other innovations from Intel, such as Intel® field programmable gate arrays (Intel® FPGAs), Intel® SmartNICs, Intel® QuickAssist Technology and Intel® Optane™ persistent memory (PMem).

Reason #2: Greater choice and mobility to accelerate cloud deployments – public, private, hybrid and multi-cloud

Choose the ecosystem, and choose wisely. Considering the cost of hardware is important when choosing cloud infrastructure but is only part of the holistic performance and performance-per-dollar picture. If a particular platform technology is offered only by a single CSP, you run the risk that workloads migrated to that environment may not work without significant refactoring and revalidation. Such efforts can expose you to substantial costs and create risk of incompatibilities. 64 percent of surveyed enterprise IT decision makers discovered applications needed to be rearchitected after infrastructure changes were made.[4]

In addition, such a proprietary environment may make it difficult to move workloads to other CSPs if at some point it makes sense to do so. Also consider that most enterprise applications (such as SAP, VMware, Microsoft SQL, Oracle, Salesforce, IBM Db2 and others) do not have commercial support for non-x86 architectures. The same is true of many popular open-source software components, tools, orchestration and container images.

In contrast, the vast majority of enterprise application and open-source projects are developed first for Intel architecture. Consider that Intel has been the leading Linux Kernel contributor for the last decade. Intel also has two decades of open-source experience with global scope, and is active in more than 650 open-source projects ranging from KVM to Kubernetes to TensorFlow. By investing in Intel technology-based cloud infrastructure, you can count on seamless enterprise workload migration within and across CSPs as well as on-premises. Put simply, you can get the software you want, when you want it, how you want it and where you want it—worry-free.

Get the best performance from software optimizations. Intel spends a great deal of time and effort making sure its customers get the best performance from their enterprise applications. Optimizing applications, tools, libraries and frameworks lets them take best advantage of Intel technology. For example, Intel and SAP have worked together so that SAP HANA takes advantage of Intel Optane PMem. Most of the leading AI frameworks, like OpenVINO™, TensorFlow, PyTorch and MXNet have been optimized by Intel to take advantage of Intel AVX-512 and Intel DL Boost. These optimizations can accelerate AI inference by as much as 30X, driving faster insights.[5] Over many years, Intel has grown teams of developers with deep experience in tuning and validating compilers, languages and middleware to work best on Intel architecture.

Intel also provides high-touch support for 700+ industry-leading cloud and enterprise ISVs, helping to optimize application performance on next-generation Intel architecture spanning cloud, visual cloud, enterprise and high-performance computing (HPC).

Software Optimization Results

The following performance improvements compare the performance of software optimized for 2nd Generation Intel® Xeon® Scalable processors to non-optimized performance.

- 30X faster AI inference performance with Intel® Deep Learning Boost and software optimizations and 6X better throughput/$ on Intel® Xeon® Scalable processors compared to a graphics processing unit (GPU)[6]
- 5X improved automatic speech recognition throughput[7]
- 3X improved rendering throughput[8]

Reason #3: Proven consistent performance from data center to cloud to edge

Cloud-based workloads benefit from having scalable architecture across cloud to edge to on premises because that way, you can quickly adjust infrastructure sizing to meet business need and can easily migrate workloads from one environment to another with confidence. Intel offers

Workload Benefits of Intel® Technology-based Cloud Instances

The following examples of performance improvements compare newer cloud instances, powered by 2nd Generation Intel® Xeon® Scalable processors, to older instances using previous-generation Intel® processors.

- AWS results:
  - MySQL + Intel: Up to 1.69X better performance and up to 1.42X better performance-per-dollar[9]
  - PostGres + Intel: Up to 1.41X better performance and up to 1.18X better performance-per-dollar[10]
- Azure results:
  - Microsoft SQL + Intel: Up to 1.54X better performance and up to 1.34X better performance-per-dollar[11]
  - Wordpress + Intel: Up to 1.51X better performance and up to 1.28X better performance-per-dollar [12]
- Google Cloud results:
  - Microsoft SQL + Intel: Up to 1.68X better performance and up to 1.64X better performance-per-dollar[13]
  - MySQL + Intel: Up to 1.21X better performance and up to 1.09X better performance-per-dollar[14]

purpose-built architecture to deliver consistent performance on a wide range of workloads. This architecture is pervasive from edge to data center to cloud. Some platform architectures are only available at very specific CSPs, or in limited data centers. But Intel-based instances have the largest, most globally available footprint across all the leading CSPs.

What's more, Intel's commitment to providing consistent, predictable application performance improvement with each new generation of Intel® processor, leads to excellent performance per dollar spent on cloud services. Choosing the right generation infrastructure for your cloud workloads can make a significant difference in your performance and performance-per-dollar. For example, 2nd Generation Intel Xeon Scalable Processors provide up to 2X improvement compared to previous-generation processors.[15]

Also, some platform architectures may provide great performance some of the time, but the underlying architecture design can introduce unpredictable latency into workloads such as databases accessing large amounts of memory and storage. Intel® Mesh Architecture provides consistent and predictable, workload-tuned performance even as you scale your cloud instances or VMs up to the largest sizes. Intel Xeon Scalable processors offer the most consistent performance and performance-per-dollar for key workloads, whether those workloads are running on-premises, at the edge or in multiple clouds.

Reason #4: Half the cores for equal performance with lowest total costs

Often, cloud cost assessment tools are driven by instance cost alone; but this approach can be misleading. Many factors affect true cloud infrastructure cost. For example, many enterprise ISVs charge per core. Choosing a server node that has fewer cores but provides similar performance to a higher-core count node can lower software licensing costs. Intel® Server Platforms provide the highest virtual machine (VM) density in the industry, which means you can do more with less. Besides core count, other aspects of total costs to consider include migration ease, compatibility, and vendor lock-in risks.

Intel has collaborated with Densify to develop tool that helps optimize cloud, container and VMware infrastructure. This tool optimizes Intel technology based instances for performance and price. You can access the tool at densify.com/intel.

Continuous Innovation in Security

Data is one of an enterprise's most valuable assets. The exponential growth in data is fueling a paradigm shift where concerns for digital security, safety, privacy and an expanding attack surface require a new approach to protecting data. Intel is doubling-down on our Security First Pledge with 3rd Generation Intel® Xeon® Scalable processors, delivering new capabilities to better protect data at scale for cloud and enterprise. For example, these processors feature new crypto accelerators coupled with software innovations up and down the stack to deliver breakthrough performance for the industry's most widely deployed cryptographic algorithms. Here are some examples of the performance benefits of these crypto accelerators:[16]

- Up to 2x faster public key encryption and decryption
- Up to 4x faster and stronger encryption with Intel AVX-512
- Up to 3x faster secure hash performance

Mainstream Enterprise Workload Results

The following performance improvements compare the performance of enterprise workloads running on 2nd Generation Intel Xeon Scalable Processors to older systems.

- Oracle Exadata + Intel: Up to 10X reduction in I/O latency[17]
- SAP + Intel: Up to 12.5X faster data load at startup[18] and up to 20% lower system cost [19]
- VMware VDI + Intel: Up to 87% more virtual desktops and up to 16% reduced total costs[20] with up to 75% faster access[21]

Infrastructure Decision Factors

Research indicates that infrastructure choice is not only about performance-per-dollar parity between various architectures. There's a significant amount of risk associated with adopting a new processor architecture:[22]

- 64% of organizations discovered that applications needed to be changed or rearchitected because of a change in the underlying infrastructure.
- 49% of organizations have grappled with hardware incompatibilities after changing the application infrastructure.

98% of organizations revalidate application behavior after changing the underlying infrastructure.[23] And that revalidation can be expensive (and time-consuming). The same survey revealed that CAPEX and OPEX savings are of the lowest priority (see graphic below). Performance, scalability, and reliability—all of which Intel processors excel at—are all more than twice as important as mere cost. [24]

Measuring success: what matters more than cost? Everything

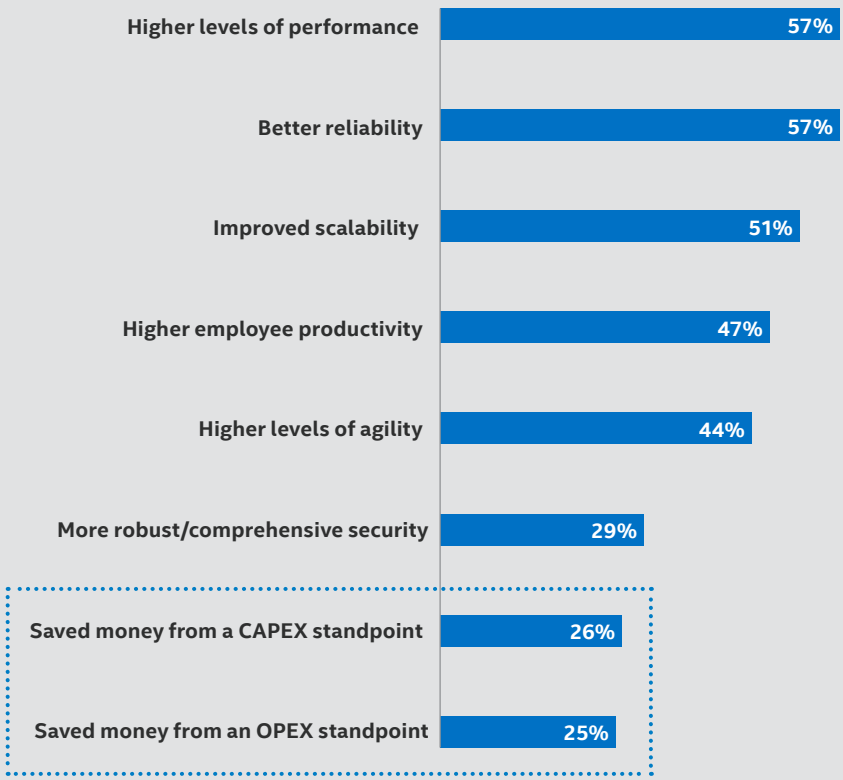

Question text:
How does your organization measure success when implementing an infrastructure change? (Percentage of respondents, N-301, five responses accepted)

Figure 1. The top five areas where organizations measure success of modernizing their IT infrastructure include improvements to performance, reliability, scalability, productivity and agility. Cost savings was cited the least.

Reason #5: Most trusted hardware for cloud-based enterprise workloads

Recent research shows that more than 75 percent of executives consider trust in the use of the cloud an important or leading priority.[25] Intel technology is most-trusted across all major CSPs, as evidenced by the fact that top ISVs like Oracle, SAP and VMware certify their cloud environments only or primarily on Intel[26]. For example, there are hundreds of certified SAP instances for the cloud, and the majority of which are available only on Intel architecture. Also, most smaller ISVs optimize their applications for Intel technology first.

Also, hybrid cloud offerings from the world's leading CSPs were offered first and run primarily on Intel architecture. For example:

- Amazon Web Services (AWS) Outposts supports multiple Intel-powered Nitro-based EC2 instance types.
- Intel and Microsoft have worked together on Azure Stack, integrating the Intel Xeon processor with Hyper-V and providing great performance and reliability.
- Intel and Google Cloud have co-engineered a Google Cloud's Anthos Intel® Select Solution, which is an Intel Reference Architecture that has been rigorously benchmark tested with Anthos in the Intel Lab.
- VMware Cloud is offered today by AWS, Azure and Google Cloud and is **only** certified on Intel architecture.

These endorsements mean that you can count on vendor support and stability for enterprise product offerings based on Intel technology.

It's been almost two decades since Intel created hardware based Virtualization Technology upon which most enterprise and public clouds are built today. Intel has been co-engineering and developing cloud technologies since the start of public cloud more than 15 years ago; that co-engineering collaboration with the world's leading CSPs continues today. With millions of Intel CPUs running in the cloud across all major CSPs, you can have greater confidence in seamless migrations as you move your workloads from on-premises to cloud, from one cloud instance to another or from cloud to cloud. In contrast, moving from Intel technology to a different platform technology or vice versa may require significant application revalidation or re-factoring, and performance may be substantially different.

Reason #6: Easiest manageability for cloud environments

No IT shop wants to have to learn a whole new set of management tools just for cloud workloads, or to maintain two sets of management tools (one for cloud and one for on-premises). Popular hybrid cloud stacks such as AWS Outposts, Azure Stack, Google Cloud's Anthos and VMware Cloud are optimized for Intel architecture. These stacks provide intuitive management capabilities that can lower cloud adoption barriers.

Also, Intel processors feature built-in telemetry that can be used to achieve closed-loop automation to orchestrate containers, optimize power consumption and streamline root-cause analysis. For example, node-level telemetry can identify workloads that do not have enough cache or memory, while cluster-level telemetry can help optimize placement decisions, scaling, and lifecycle management. Using the telemetry capabilities of Intel architecture, you can achieve:

- Operational efficiency
- Infrastructure efficiency
- High utilization
- Excellent performance
- High reliability

38 percent of network outages take more than a day to find and resolve, and improving network resilience is the number 1 priority for 49 percent of global IT leaders.[27]

For more information about Intel architecture and telemetry, read the white paper, "Building Intelligence into Your Cloud And Data Center."

Cloud Workload Management Using Intel Technology

Intel has incorporated metrics similar to those in Collectd into cAdvisor, Google's Container Advisor project. cAdvisor provides insight into hardware utilization at the container level (the telemetry in cAdvisor can also be applied to VMs). For more information, read the blog.

Ready to start—or extend—your cloud journey with Intel technology?

- Read the white paper, "Top Clouds Are Powered by Intel."
- Visit https://www.intel.com/content/www/us/en/cloud-computing/overview.html.
- Contact your Intel representative for more information.

[1] https://www.crn.com/news/cloud/10-future-cloud-computing-trends-to-watch-in-2021

[2] https://www.gartner.com/en/newsroom/press-releases/2020-10-07-gartner-says-worldwide-data-center-infrastructure-spending-to-grow-6-percent-in-2021

[3] Copyright Liftr Insights - Cloud Components Tracker – January 2021

[4] The Impact of Evolving On-premises Environments, Enterprise Strategy Group, June 2020

[5] https://www.intel.com/content/www/us/en/products/docs/processors/xeon/2nd-gen-xeon-scalable-processors-brief.html, endnote 4.

[6] How We Scaled Bert To Serve 1+ Billion Daily Requests on CPUs - Roblox Blog

[7] https://software.intel.com/content/dam/develop/external/us/en/documents/qihoo-cs-329387.pdf

[8] https://cdw-prod.adobecqms.net/content/dam/cdw/on-domain-ca/industries/media-and-entertainment/cdw-studiocloud-render-study-intel-optane.pdf

[9] Performance varies by use, configuration and other factors. Learn more at http://www.intel.com/PerformanceIndex. Testing done by Principled Technologies on 11/12/2020. For configuration and test details, visit http://facts.pt/H9kHsIV and http://facts.pt/cTvUFqi for explanations (commissioned by Intel). AWS EC2 on demand price as of 11/12/2020 for US East-1f region: M5n $0.476/hr, $0.952/hr, $3.808/hr vs M4 $0.40/hr, $0.80/hr, $3.20.hr for 8, 16, 64 vCPUs

[10] Performance varies by use, configuration and other factors. Learn more at http://www.intel.com/PerformanceIndex. Testing done by Principled Technologies on 11/12/2020. For configuration and test details, visit http://facts.pt/xHrpYj1 and http://facts.pt/3YlY74Z for explanations (commissioned by Intel). AWS EC2 on demand price as of 11/12/2020 for US East-1f region: M5n $0.476/hr, $0.952/hr, $3.808/hr vs M4 $0.40/hr, $0.80/hr, $3.20.hr for 8, 16, 64 vCPUs

[11] Performance varies by use, configuration and other factors. Learn more at http://www.intel.com/PerformanceIndex. Testing done by Principled Technologies on Sept 2020. For configuration and test details, visit http://facts.pt/ccqc4xe (commissioned by Intel). Azure on demand price as of 11/3/2020 for US East region: Eds_v4 $0.288/hr, $1.152/hr, $4.608/hr vs Es_v3 $0.252/hr, $1.008/hr, $3.629/hr 4, 16 and 64 vCPUs

[12] Performance varies by use, configuration and other factors. Learn more at http://www.intel.com/PerformanceIndex. Testing done by Principled Technologies on 11/3/2020. For configuration and test details, visit http://facts.pt/xif2usg (commissioned by Intel). Azure on demand price as of 11/3/2020 for US East region: Dds_v4 $0.226/hr, $0.904/hr, $3.616/hr vs Ds_v3 $0.192/hr, $0.768/hr, $3.072/hr 4, 16 and 64 vCPUs

[13] Performance varies by use, configuration and other factors. Learn more at http://www.intel.com/PerformanceIndex. Testing done by Principled Technologies on 11/3/2020. For configuration and test details, visit http://facts.pt/0u75KOz & http://facts.pt/NZJ1Hqd for explanations (commissioned by Intel). Azure on demand price as of 11/3/2020 for US East1 region: N2 $0.777/hr, $3.108/hr vs N1 $0.760/hr, $3.040/hr for 16 and 64 vCPUs.

[14] Performance varies by use, configuration and other factors. Learn more at http://www.intel.com/PerformanceIndex. Instance (GCP N1 Highmem & N2 Highmem us-central1) testing done by Intel in August 2020. OS: Ubuntu 20.04 LTS, Kernel: 5.4.0-1021-gcp, WL: MySQL 8.0.20, OLTP HammerDB v3.2, 64GB innodb_buffer_pool_size, 800 warehouses, 38 users (peak value), 32 vCPU, 208 GB Mem for N1, 256 GB Mem N2, 32 Gbps network, Persistent Disk SSD 500GB. GCP on demand price as of 11/12/2020 for US Central1 region: N2-highmem-32 $2.096/hr, N1-highmem-32 $1.893/hr

[15] https://www.intel.com/content/www/us/en/products/docs/processors/xeon/2nd-gen-xeon-scalable-processors-brief.html, endnote 3.

[16] Average, depending on workload. For more complete information about performance and benchmark results, visit www.intel.com/benchmarks.

[17] https://www.oracle.com/engineered-systems/exadata/database-machine/

[18] https://www.intel.com/content/www/us/en/big-data/partners/sap/fast-sap-hana-start-times-persistent-memory-infographic.html

[19] https://discover.sap.com/sap-hana-dc-persistent-memorynew/en-us/index.html

[20] https://blogs.vmware.com/virtualblocks/2020/06/18/intel-select-solution-vmware-horizon-vdi-vsan/

[21] https://www.dandh.com/media/pdf/pages/landing-bizcontinuity/Intel_VDI-Overview.pdf

[22] https://www.intel.com/content/www/us/en/cloud-computing/esg-on-premise-infrastructure-infographic.html?wapkw=Considerations%20and%20Drawbacks%20of%20Modernization

[23] Copyright Liftr Insights - Cloud Components Tracker – January 2021

[24] https://www.intel.com/content/www/us/en/cloud-computing/esg-on-premise-infrastructure-paper.html?wapkw=Considerations%20and%20Drawbacks%20of%20Modernization

[25] "Trust in cloud technology and business performance." The Economist, 2016. https://lp.google-mkto.com/rs/248-TPC-286/images/EIU_Trust_in_Cloud_Technology_FINAL.pdf

[26] https://www.sap.com/dmc/exp/2014-09-02-hana-hardware/enEN/#/solutions?filters=iaas

[27] ComputerWeekly. "Multiple network outages costing enterprise businesses millions of pounds a year." May 2020. computerweekly.com/news/252483308/Multiple-network-outages-costing-enterprise-businesses-millions-of-pounds-a-year.

Performance varies by use, configuration and other factors. Learn more at www.Intel.com/PerformanceIndex.

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.

0321/SMEY/CAT/PDF ♻ Please Recycle 346592-001EN