

The Intel logo is displayed in white lowercase letters on a white rectangular background. The background of the entire advertisement is a night-time aerial view of a city skyline, featuring numerous illuminated skyscrapers and buildings, with a prominent, very tall, slender skyscraper in the center-right. The sky is a deep blue with some light clouds.

intel.

The Cloud Continues to Evolve

More accessible,
more powerful with Intel

Powerful clouds lead to new opportunities

Digital transformation continues to sweep through markets, disrupting traditional IT solutions and services while generating new opportunities for growth and innovation. Faster, more-available, and more-secure cloud infrastructure means that businesses can design and deliver services at breakneck speeds and support advanced use cases like analytics, AI machine learning, and HPC in the cloud. The main barrier to entry is complexity, as workload placement and migration become more challenging across increasingly distributed hybrid and multicloud environments.



Pain points that businesses like yours face:

- Difficulty navigating distributed environments across private and public cloud infrastructure
- Workloads are becoming more complex with AI and established services
- Training for new solutions is costly and takes time

Intel-powered clouds simplify and solve problems:


- Private and public clouds run on the same validated Intel® architecture, easing migration headaches
- Intel offers key optimizations for fast AI and cloud microservices
- Many developers are already trained on x86 solutions

Experts in cloud computing optimization

Whether it is Intel® processors, Intel® storage or networking technology, or Intel® software and toolkits, the cloud is powered by Intel. Over 50,000 globally available instance types from leading cloud service providers (CSPs) operate on the Intel® Xeon® Scalable platform.¹ Intel at the heart of the cloud means that your business can easily access or scale up with familiar solutions—and integrate the latest cloud infrastructure into legacy investments—to help realize value fast.

Real value, realized

Here's what real-world businesses accomplished with Intel-powered cloud



Key takeaway

Pull ahead of the pack

In this eBook you'll learn about the key technologies that enable fast and resilient cloud infrastructure to help piece together a comprehensive strategy for success. Whether you're just starting your cloud journey or looking to augment your current cloud efforts, Intel leadership in the cloud is your ticket to discovering your advantage and pulling ahead of competitors.

Cloud 101:

What is the cloud?

Simply put, the cloud is computing resources that are abstracted and made available remotely through public or private infrastructure. Rather than running workloads of bare metal servers, businesses can run workloads within a collective, virtualized pool of computing resources that scale across any number of servers.

Why the cloud matters

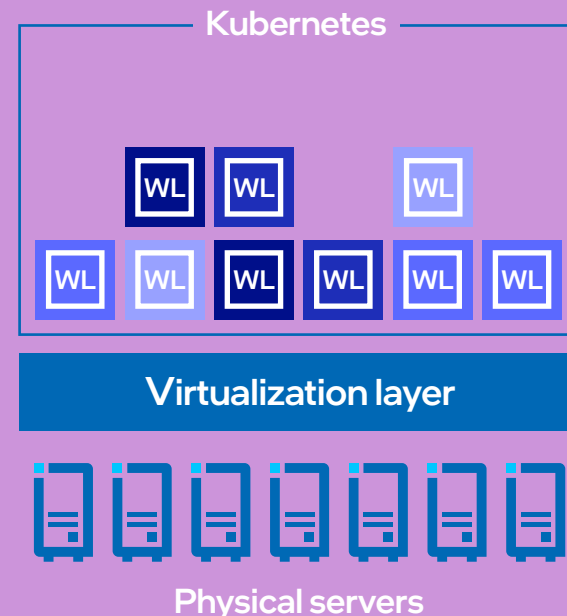
Because the cloud makes compute resources more available and cost efficient, many business-critical workloads are migrating to the cloud, including VMware, SAP, Microsoft SQL Server, Salesforce, Oracle, and data analytics. Gartner predicts that global cloud spending will increase by 18 percent from 2020 to 2021,⁵ compared to just 6 percent for data center IT spending over the same time frame.⁶



Cloud vocabulary

Cloud computing nomenclature is extensive, but this section will narrow the focus to three key concepts to fuel your understanding of how modern clouds operate.

- **Virtualization:** The methodology of abstracting computer resources from bare metal environments into virtual machines (VMs). VMs run their own operating system, and multiple VMs can run on a single bare metal server.
- **Container/containerization:** Smaller, virtualized instances that run a single app or workload. Containers do not need their own operating system, so they can spin up faster and more efficiently than VMs. Docker is the common framework for creating and deploying containers.
- **Orchestrator:** A platform that allows users to deploy and manage multiple containers in a cloud environment. Kubernetes is the standard platform used to orchestrate containers.





Cloud models

Cloud models explain the how of building a cloud, including who is doing the work of setting up the infrastructure (servers, switches, cabling) and where that infrastructure is physically located.

- **Private cloud** refers to a cloud infrastructure built and hosted on-premises at the business location. Benefit: maximum control over how resources are used.
- **Public cloud** refers to cloud infrastructure built and managed by a third party. Key examples of public CSPs include Amazon Web Services (AWS), Google Cloud, Microsoft Azure, and Oracle. Benefit: elastic resourcing purchased with OpEx budget and no capital investment required.
- **Hybrid cloud** is when a business combines their own private cloud infrastructure with public cloud resources to achieve best-of-both-worlds control for sensitive workloads and resource elasticity for seasonal spikes.
- **Multicloud** is the usage of cloud instances from multiple CSPs to prevent dependency on a single provider or to take advantage of multiple capabilities that no single provider offers.

Real value, realized

Here's what real-world businesses accomplished with Intel-powered cloud

Consolidation:

25,000 branch servers ► One unified cloud platform⁷

State Bank of India (SBI)

Details:

SBI converted their 25,000 individual branch servers to a unified cloud platform that all branches could access remotely. Previously, a downed branch server would take up to six days to restore. With a private cloud, SBI was able to provide consistent uptime while reducing their management overhead.

Service delivery models

Now that you grasp how clouds can be deployed, the next question is, what can you do with them?

Software as a Service (SaaS):

Software hosted in the cloud and accessed via the internet.

Example: **Salesforce, Microsoft Teams, OneDrive**

Infrastructure as a Service (IaaS):

Virtualized or physical compute resources that you rent and access remotely. Most public cloud offerings fall into this category.

Example: **AWS, Azure, Google Cloud, Oracle Cloud**

Platform as a Service (PaaS):

Cloud-hosted software frameworks used to create and manage your own apps.

Example: **Google App Engine**

Everything as a Service (XaaS):

A catchall category for any other independent service that a business or a user can access remotely.

Example: **AI as a Service (AIaaS), Video Safety as a Service (VSaaS)**

Key takeaway

A foundation for strategy

These are just the puzzle pieces that make up the larger picture of how cloud computing works to benefit your business. With Intel's guidance, you can start planning an effective cloud strategy that supports the IT functions you need right now and accommodates for future growth.



■ Cloud architecture in depth

The Intel® portfolio offers end-to-end solutions across processor, storage, networking, and software toolkits to help optimize cloud price performance and security. Choice in infrastructure matters, because switching or migrating to a new infrastructure can impose substantial costs as IT departments need to rearchitect applications for new environments. Also consider that most enterprise applications—such as SAP, VMware, Microsoft SQL, Oracle, Salesforce, and IBM Db2—do not offer commercial support for non-x86 architecture, and most open source projects are first developed for Intel-enabled platforms. By investing in Intel® technology, you benefit from ecosystem familiarity with Intel® hardware and software solutions, as well as seamless migration across Intel-based private and public cloud resources.

CSP collaboration for better clouds

Intel works closely with major CSPs like Alibaba, AWS, Baidu, Google, Microsoft Azure, and Oracle to ensure that the latest technologies and hardware solutions are integrated into their cloud platforms. Across top CSPs worldwide, a majority of all cloud instances are powered by Intel® technology.¹ This means that even if your business only uses public cloud resources, you can still benefit from the exclusive capabilities and performance levels made possible by Intel® architecture.

Intel® Xeon® Scalable processors

The heart of the cloud is powered by Intel Xeon Scalable processors, which deliver high core counts, high memory capacity, PCIe Gen 4 support, and exclusive features:

- **Intel® Virtualization Technology** enables cloud architects to host more VMs per server.
- **Intel® Software Guard Extensions (Intel® SGX)** isolates workloads in memory to help protect data.
- **Intel® AVX-512 and Intel® Deep Learning Boost (Intel® DL Boost)** reduce computational requirements for AI training and accelerate AI inference.⁸

Performance highlights

3rd Gen Intel® Xeon® Scalable processors:

- **25x better inference throughput** with INT8 data type vs. AMD Milan with FP32⁹
- **1.58x higher performance** for cloud microservices vs. previous generation¹⁰

Storage and memory

Intel® Optane™ DC SSDs

These high-capacity SSDs with options for PCIe interfaces that position data closer to the CPU offer a low-power, low-TCO option for data storage in the cloud.

Intel® Optane™ persistent memory

High-capacity memory modules support intensive I/O and help protect against data loss by enabling memory to persist through system shutdowns and reboots.

Networking

Intel® Ethernet and Intel® Silicon Photonics

Intel is the Ethernet market share leader. The Intel® Ethernet Network Adapter 800 Series delivers network speeds up to 100GbE to help eliminate data bottlenecks. Intel Silicon Photonics also offers a selection of optical transceivers that help cloud builders transition to fiber for even faster network speeds.

Intel® Tofino™ Ethernet Switch ASICs

Intel Tofino is a P4-programmable switch ASIC that offers more flexibility than traditional fixed-function ASICs. Programmability enables cloud builders to customize switch behavior, optimize data flow for specific workloads, and support advanced telemetry for easier network management and improvement.

Acceleration

Intel® field programmable gate arrays (FPGAs)

Programmable accelerators for SmartNIC integration and system-on-chip (SOC) solutions allow for greater parallelization of workloads, or faster data encryption at each node in the cloud. FPGA behavior can also be changed after deployment, making it easy to reuse or reconfigure hardware as you continue to add to your investments.

Key takeaway

Trust in your infrastructure

Make sure your first step into the cloud is smart, informed, and tactical. Intel in the cloud means more than just processors and covers a vast portfolio of solutions along with public cloud integration. This high degree of trust, ubiquity, and familiarity of Intel in the cloud helps businesses achieve low cost and efficient-workload migration.

Cloud security: Protection enables performance

Data security helps make the cloud possible by helping to protect businesses and limit their liability and risk. Without robust security, you are both limiting the benefits you can realize with the cloud and exposing yourself to greater harm. The best approach will include hardening your cloud platform from the ground up, with hardware-enabled security.



The role of data security in the cloud

When conceptualizing the key value of data security for cloud computing, experts today are focused on three areas: confidentiality, integrity, and availability (CIA).

Confidentiality

Only the correct entities, people, and apps have access to the right data.

Integrity

Data is not modified or tampered with in any way.

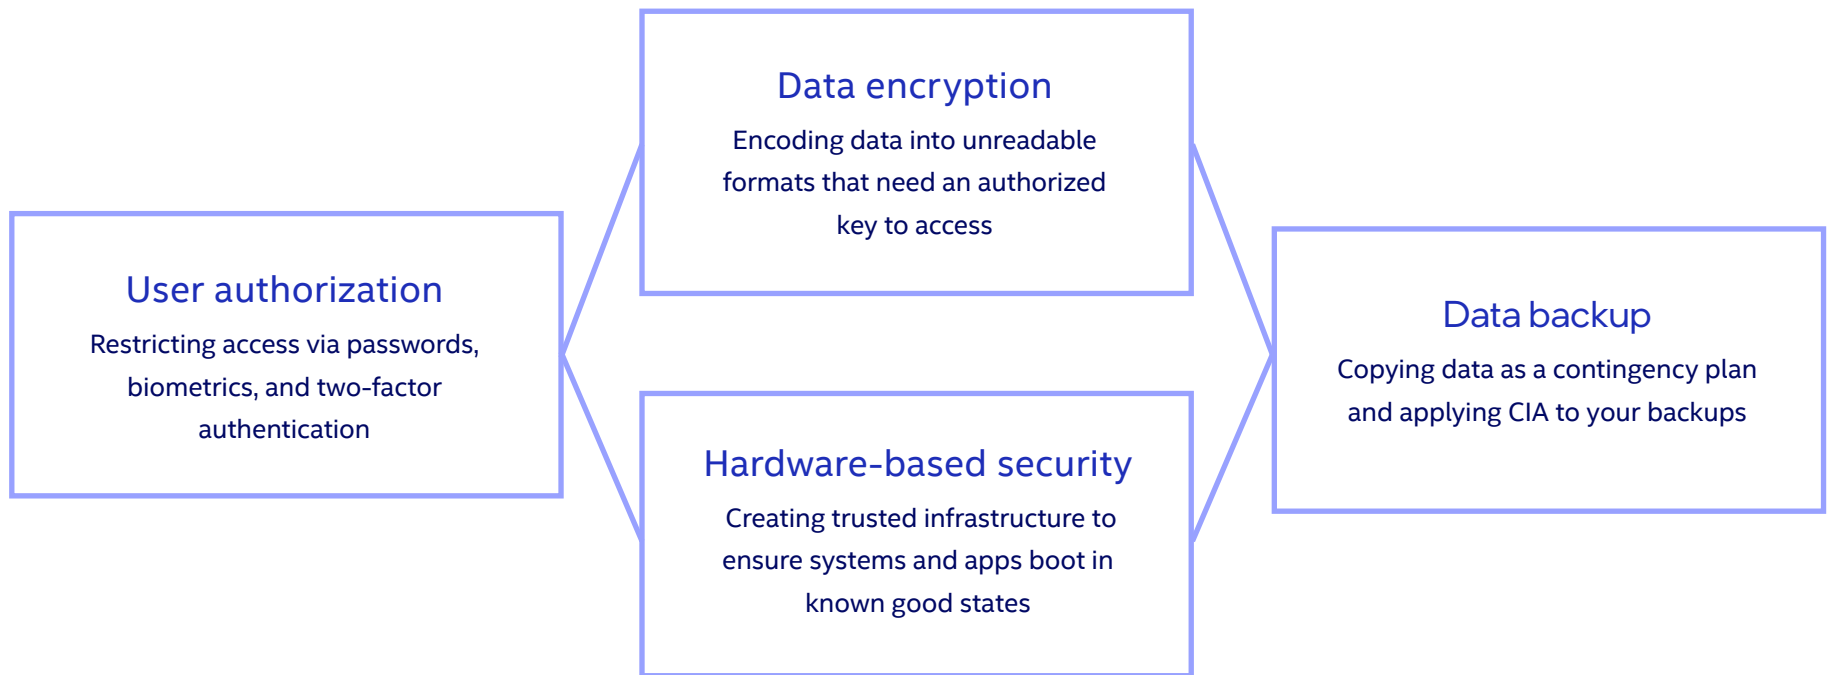
Availability

Data is always accessible when you need it.

Most digital threats, whether they are malware or more sophisticated attacks like privilege escalation, will attempt to compromise one of these three vectors. Attackers will try to either access confidential data, corrupt data by making it false or misleading, or block your access to the data you need.

A comprehensive strategy for protection

To help build toward CIA in your cloud infrastructure, your efforts should align with these four areas: encryption, authorization, backup, and hardware. Even if your business rents public cloud resources and you don't maintain on-premises infrastructure, these procedures still apply.





Real value, realized

Here's what real-world businesses accomplished with Intel-powered clouds.

Faster AI training:

UC San Francisco (UCSF) validates clinical AI models in days instead of years.¹¹

Details:

UCSF uses the Fortanix Confidential Computing Enclave Manager platform, enabled by Intel® SGX, to run applications in private regions of memory called enclaves. This allows organizations to validate clinical AI algorithms with shared data sets while protecting patient privacy and keeping IP confidential.

Data sovereignty

Data sovereignty refers to the jurisdiction that local governments may claim over your data, depending on your data's physical location. Sovereignty is important to cloud use cases specifically because the physical servers that underpin a cloud infrastructure can be located anywhere in the world. When renting cloud infrastructure from public CSPs, inquire about where your data is stored and processed and which laws may apply.

Key takeaway

The value of a strong defense

Venturing into cloud computing without a strong understanding of digital security needs can result in more harm than benefit. By investing in strong security as well as hardware-enabled protections, you can help safeguard the time, energy, and capital that your business invests in the cloud, as well as the business agility that comes out of it.

Public clouds: Extend your capability

Public clouds allow businesses to extend beyond their capabilities with scalable resources on demand. These resources enable separate business units to access high-performance clusters with operating expenditures (OpEx) rather than investments in capital expenditures (CapEx) like new hardware. Even if you invest heavily in a robust private cloud infrastructure, public clouds can sometimes offer specialized resources or support workloads like AI and HPC that you may not currently support on-premises.



Leading CSPs enabled by Intel

Intel collaborates closely with CSPs like AWS, Google Cloud, Microsoft Azure, and Oracle to engineer and optimize cloud infrastructure for the highest performance per cost. Customers and businesses like yours ultimately benefit with faster time to results, lower subscription fees, and access to exclusive Intel® technologies and toolkits.

AWS

Intel-based AWS instances use Intel® Xeon® Scalable processors and Intel® AVX-512 across a range of general purpose and compute/storage/memory-optimized cloud solutions.

Success metrics by workload

Intel-enabled AWS cloud instances vs. older instances with previous-generation processors:¹²

- MySQL + Intel: Up to 1.69x better performance, up to 1.42x better performance per dollar
- Apache Spark + Intel: Up to 1.72x better performance, up to 1.44x better performance per dollar

Google Cloud

Google Cloud Platform N2 and C2 machine types use Intel Xeon Scalable processors and Intel AVX-512. Google Cloud and Intel have also coengineered an Intel® Select Solution, Google Cloud Anthos, with rigorous testing in the Intel lab.

Success metrics by workload

Intel-enabled Google Cloud instances vs. older instances with previous-generation processors:¹³

- PostgreSQL + Intel: Up to 1.21x better performance, up to 1.18x better performance per dollar
- Streaming Analytics + Intel: Up to 1.51x better performance, up to 1.47x better performance per dollar

Microsoft Azure

Microsoft Azure virtual machines enabled by Intel Xeon Scalable processors support targeted workloads for everything from general purpose to confidential computing with Intel® SGX, to hyperconverged infrastructure with Intel® Optane™ technology.

Success metrics by workload

Intel-enabled Azure cloud instances vs. older instances with previous-generation processors:¹⁴

- Microsoft SQL + Intel: Up to 1.54x better performance, up to 1.34x better performance per dollar
- WordPress + Intel: Up to 1.51x better performance, up to 1.28x better performance per dollar

Oracle

Oracle integrates Intel Xeon Scalable processors, Intel® Ethernet products, and Intel® SSDs into their cloud service offerings.

Success metrics by workload

Intel-enabled Oracle X9 cloud instances vs. older X7 instances with previous-generation processors:¹⁵

- Cloud HPC + Intel: 3rd Gen Intel® Xeon® Scalable processors are expected to drive up to 30% higher performance

A man with glasses and a beard is sitting at a desk in a server room, looking at a computer monitor. The room is dimly lit with blue and green lights from the server racks. The man is wearing a grey t-shirt and has a lanyard around his neck. The computer monitor shows a dashboard with various charts and data. The overall atmosphere is professional and technical.

Onboarding made easy

Given the expansive offerings from CSPs, it can be difficult to know which cloud instances to choose and which add-on services to choose. You may even encounter challenges migrating your data to cloud platforms or setting up the complex workloads seen in HPC. Intel ecosystem partners like [Ronin](#), [Six Nines](#), and [OnScale](#) offer onboarding services that recommend specific instances, configurations, and proprietary UI or dashboard tools to simplify cloud management.

Pro tip: Understanding data migration fees

Public CSPs make it easy to onboard your data for little to no charge. However, most service contracts will set high fees for migrating your data to other platforms or even for pulling your data out of the cloud. Make sure you understand what's involved when transitioning to a new service provider.

Real value, realized

Here's what real-world businesses accomplished with Intel-powered clouds.

Faster results:

Runtime reduced from 11 minutes to 1.5 minutes on Google Cloud.¹⁶

onscale.

Details:

OnScale built a digital prototyping SaaS offering with Google Cloud machine types enabled by Intel® Xeon® Scalable processors. For an approximate 4x increase in core hours, a mechanical simulation ran 8x faster.¹⁶

Key takeaway

Embrace agility with public cloud offerings

Public and hybrid clouds offer a level of flexibility beyond on-premises infrastructure alone. Don't pass up the opportunity for a more agile platform while still benefiting from Intel-enabled cloud performance and security.



■ Optimizing for cost efficiency

The goal of optimization is to do more with less. In the case of public CSP resources, this means achieving the same or similar performance levels with fewer cores or core hours, which can also help reduce software licensing fees. For private clouds, optimization means fine-tuning infrastructure for better performance with less overhead while minimizing downtime.

Forecasting CSP capabilities

Telemetry tools can help you measure the performance of public cloud infrastructure. This is useful when shopping among different CSP providers or considering a migration to a new provider. Intel® Cloud Optimizer, jointly developed by Intel and Densify, can help provide recommendations for CSP offerings based on your specific workloads. Learn more at densify.com/intel.

Improving performance by visualizing performance

Deploying telemetry tools may require you to have access to your cloud infrastructure and network, which may or may not be available with CSP offerings. The good news is that software-level optimizations may already be in place to help speed time to results with fewer core hours on Intel-enabled CSP instances. For private clouds, you can extend your efforts even further with telemetry and visualization applications.

Cloud infrastructure telemetry

Intel® Telemetry Collector is a tool that gathers and visualizes cloud performance data using information provided by sophisticated on-chip hardware called Intel® Performance Monitoring Units (Intel® PMUs).¹⁸ This tool supports quick analysis of cloud workload distribution, hardware utilization, network performance, as well as system power and thermals. IT administrators can use this data to focus their efforts on achieving higher density, dynamic power capping, more-efficient thermal controls, and performance optimization throughout your cloud infrastructure.

- What** Total infrastructure readout of utilization, power, and thermals
- How** Intel Telemetry Collector, embedded Intel PMUs
- Why** Achieve higher density, more-efficient thermal control, better performance

Cloud network telemetry

In-Band Network Telemetry (INT) is a methodology that focuses on network performance by gathering metadata from data packets as they travel through nodes in your cloud infrastructure. INT can be enabled on nodes with programmable Ethernet switch ASICs like Intel® Tofino™ and feed the data to Intel® Deep Insight, a platform for understanding and visualizing telemetry data.

For enabled nodes, the switch reads packet metadata to ascertain information such as packet origin, destination, and departure/arrival time stamps. Network admins can then use this data to monitor for dropped packets and measure latency with granular-level detail in near-real time, making it easier to identify and troubleshoot issues faster.¹⁹

- What** Network monitoring of dropped packets and traffic latency
- How** Intel Tofino Ethernet switch ASICs, Intel Deep Insight
- Why** Identify network problems in near-real time and troubleshoot faster

Key takeaway

Taking cloud infrastructure to new heights

Without automation and optimization, you're leaving potential performance and TCO improvements on the table. Telemetry tools from Intel allow your IT department to refine your cloud platforms for higher levels of density and efficiency.

AI and HPC in the cloud

Cloud computing has evolved to offer near-similar performance levels as on-premises infrastructure for AI and HPC applications that are the most compute intensive, memory intensive, and storage intensive. This higher ceiling on performance opens new opportunities to drive high-value use cases for analytics, deep learning, and simulation that were previously off-limits.



Hardware AI acceleration, toolkit optimization

As mentioned in previous chapters, Intel® Xeon® Scalable processors offer embedded AI acceleration with exclusive features: Intel® AVX-512 and Intel® DL Boost. Intel also provides optimized distributions of leading AI frameworks like OpenVINO™, TensorFlow, PyTorch, and MXNet to take advantage of Intel AVX and Intel DL Boost.

For example, the Intel® Distribution of OpenVINO™ toolkit accelerates AI vision workloads, and the Intel® Distribution of TensorFlow accelerates AI inference for deep learning models. Combining hardware-enabled AI acceleration with optimized software can yield substantial improvements for AI in the cloud.

Confidentiality

Only the correct entities, people, and apps have access to the right data.

Real value, realized

Here's what real-world businesses accomplished with Intel-powered clouds.

Faster AI:

3.78x faster AI text recognition²¹

NAVER

Details:

NAVER, a South Korean CSP, deployed the Intel® Distribution of OpenVINO™ toolkit in its Character Region Awareness for Text (CRAFT) detection model, resulting in a massive 3.78x speedup for AI text recognition workloads.

Unified programming for HPC clouds

HPC cloud configurations integrate a multitude of heterogeneous accelerators and processors to achieve the highest performance levels for the most intensive workloads like genomic sequencing and computational fluid dynamics. Intel offers essential HPC optimizations through the Intel® oneAPI toolkit, which includes Intel flavors of familiar code libraries like Intel® MKL and Intel® MPI, along with Intel® compilers and the Intel® oneAPI rendering toolkit. HPC brings all new levels of complexity to cloud infrastructure, but a good place to start is with **Intel® Select Solutions for Hybrid Cloud**, which help simplify deployments with proven, prebaked configurations.

Real value, realized

Here's what real-world businesses accomplished with Intel-powered clouds.

Faster CAE:

15x faster prototyping²²

ROYAL ENFIELD

Details:

Royal Enfield used the Intel-enabled Altair HyperWorks HPC cloud appliance to accelerate computer-aided engineering (CAE) by a factor of 15x for its new motorcycle designs, with 25 percent fewer prototypes.²²

Key takeaway

AI and HPC are no longer niche

AI and HPC in the cloud are becoming more accessible for enterprises, opening a new dimension for competitive advantage along the lines of deep learning and simulation. Intel provides the AI and HPC tools and guidance you need to turn opportunity into real value.




intel.

The cloud is more powerful with Intel

Intel-enabled cloud computing provides high levels of performance, efficiency, and security while reducing the friction of new cloud deployments, so businesses can drive value and competitive advantage. Use Intel guidance to direct the conversation within your organization around building out cloud capacity to handle your business needs and accommodate for future growth.

Get started at intel.com/cloud



Footnotes and citations

1. "Liftr Cloud Components Tracker," Liftr Insights, April 2021. <https://liftrinsights.com/cloud-components-tracker>
2. "Kingsoft Cloud Improves Response Time for CDN," April 2021. <https://www.intel.com/content/www/us/en/customer-spotlight/stories/kingsoft-cloud-cdn-customer-story.html>
3. "HYHY: Full-Cycle AI-Based Medical Imaging Solution," April 2021. <https://www.intel.com/content/www/us/en/customer-spotlight/stories/hyhy-customer-story.html>
4. "Prog-It: Agile, All-in-One Private Cloud System," August 2020. <https://www.intel.com/content/www/us/en/customer-spotlight/stories/prog-it-customer-story.html>
5. "10 Future Cloud Computing Trends To Watch In 2021," CRN.com, November 2020. <https://www.crn.com/news/cloud/10-future-cloud-computing-trends-to-watch-in-2021>
6. "Gartner Says Worldwide Data Center Infrastructure Spending to Grow 6% in 2021," Gartner, October 2020. <https://www.gartner.com/en/newsroom/press-releases/2020-10-07-gartner-says-worldwide-data-center-infrastructure-spending-to-grow-6-percent-in-2021>
7. "State Bank of India: Future-Ready Infrastructure," December 2019. <https://www.intel.com/content/www/us/en/customer-spotlight/stories/state-bank-of-india-customer-story.html>
8. Various researchers have demonstrated that both deep learning training and inference can be performed with lower numerical precision, using 16-bit multipliers or fewer for inference with minimal to no loss in accuracy. <https://www.intel.com/content/www/us/en/artificial-intelligence/solutions/lower-numerical-precision-deep-learning-inference-and-training.html>
9. See [45] at www.intel.com/3gen-xeon-config. Results may vary.
10. See [98] at www.intel.com/3gen-xeon-config. Results may vary.
11. "Intel® SGX Helps UCSF Propel Medical Device Innovations," October 2020. <https://www.intel.com/content/www/us/en/newsroom/news/ucsf-propel-medical-device-innovations.html#gs.1aendw>
12. 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 USD 0.476/hr, USD 0.952/hr, USD 3.808/hr vs M4 USD 0.40/hr, USD 0.80/hr, USD 3.20/hr for 8, 16, 64 vCPUs.
13. Testing done by Principled Technologies on 11/12/2020. For configuration and test details, visit <http://facts.pt/OQ5h9xl> and <http://facts.pt/5MVUbdA> for explanations (commissioned by Intel). GCP on-demand price as of 11/12/2020 for US East1 region: N2 USD 0.389/hr, USD 0.777/hr, USD 3.108/hr vs N1 USD 0.380/hr, USD 0.760/hr, USD 3.040/hr for 8, 16, 64 vCPUs.
14. 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 USD 0.288/hr, USD 1.152/hr, USD 4.608/hr vs Es_v3 USD 0.252/hr, USD 1.008/hr, USD 3.629/hr for 4, 16 and 64 vCPUs.
15. "10nm Intel® Xeon® Scalable Processors to Power Oracle's Next-Gen HPC Cloud," September 2020. <https://newsroom.intel.com/articles/10nm-intel-xeon-scalable-oracle-hpc-cloud/>
16. "OnScale Makes Digital Prototyping Accessible," February 2021. <https://www.intel.com/content/www/us/en/customer-spotlight/stories/onscale-customer-story.html>
17. "Node Feature Discovery (NFD)," Kubernetes Documentation Site, 2020. <https://docs.01.org/kubernetes/nfd/index.html>
18. "Cloud Telemetry: Advancing Your IT Strategy." <https://www.intel.com/content/www/us/en/cloud-computing/telemetry.html>
19. "In-band Network Telemetry Detects Network Performance Issues," 2020. <https://builders.intel.com/docs/networkbuilders/in-band-network-telemetry-detects-network-performance-issues.pdf>
20. See [118] at www.intel.com/3gen-xeon-config. Results may vary.
21. "NAVER: Low-Latency Machine-Learning Inference," January 2021. <https://www.intel.com/content/www/us/en/customer-spotlight/stories/naver-ocr-customer-story.html>
22. "Royal Enfield: CAE Steps Up to an HPC Appliance," June 2020. <https://www.intel.com/content/www/us/en/customer-spotlight/stories/royal-enfield-customer-story.html>

Notices and disclaimers

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

Performance varies by use, configuration, and other factors. Learn more at 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.

Results that are based on systems and components as well as results that have been estimated or simulated using an Intel® Reference Platform (an internal example new system) or Intel® analysis or architecture simulation or modeling are provided to you for informational purposes only. Results may vary based on future changes to any systems, components, specifications, or configurations.

Intel may change availability of products and support at any time without notice. All product plans are subject to change without notice.

Your costs and results may vary.

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

Customer is responsible for safety of the overall system, including compliance with applicable safety-related requirements or standards.

© Intel Corporation. Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation or its subsidiaries. Other names and brands may be claimed as the property of others.

US/07/2021/PDF/JH/CMD