Innovate More. Spend Less.

Maximize value with better performance and efficiency. Take advantage of AI to enter new markets, grow revenue and innovate beyond your competition, while reducing TCO.

Every business must maximize the value it gets from its computing environment while striving to get the best performance. Increased performance and efficiency support growth, new opportunities and improved competitiveness, and it is critical to every aspect of every business, from day-to-day tactical operations to long-term strategic guidance. Likewise, improving cybersecurity is a constant requirement, to avoid the potentially crippling costs of a breach in time, money and reputation. The technology costs to reach these goals are a vital driver for growth, but left unchecked, they can interfere with the bottom line.

Intel provides two primary approaches to help companies reduce costs — modernizing and consolidating the environment with refreshed technology and optimizing existing solutions to reduce TCO. Select the Modernize or Optimize icon below to explore either option.

Modernize old tech

In many cases, systems from three or four years ago are not capable of meeting today's demands. It's also important to consider which technology providers can deliver the best results, including TCO. Upgrade to the latest Intel technology to:

- **Consolidate infrastructure**. Supporting the same workload capacity with fewer servers consumes less space, power, software licenses and other supporting resources, which helps bring down operating costs.
- **Take advantage of AI**. Enter new markets, grow your revenue and innovate beyond your competition.
- **Improve cybersecurity**. The costs of a breach in time, money and reputation can be crippling to a business, and investing in modern technology is a wise investment to avoid it.
- **Improve competitiveness**. Modernization positions the business to roll out new services and experiences more effectively, avoiding opportunity costs by being ready for new opportunities.
- **Reduce energy consumption**. Modern servers provide higher performance per watt to lower operating costs, and they are more reliable, which reduces the IT burden.

Top **5** benefits of modernization with Intel

Delivering innovative business models and services often increases the demands on your business's IT infrastructure, pushing it beyond the scale it was originally designed to deliver. Modernized infrastructure is needed, with accelerated AI throughput and more performance per core to support new deployment models, achieve novel goals, and meet evolving application and workload demands.

1. Save money.

Get unmatched TCO when upgrading from 1st Gen Intel® Xeon® to 5th Gen Intel Xeon CPUs.

UP TO **77%**[1] reduction in TCO

UP TO **94%**[1] reduction in server count

UP TO **90%**[1] reduction in power

2. Use fewer servers.

Save power and money on new server purchases, deploying fewer 5th Gen Intel Xeon processor-based servers to meet performance and TCO goals.

Comparisons to deploying 50 servers with 3rd Gen Intel® Xeon® processor

	Web (NGINX TLS)	Data Services (RocksDB)	Artificial Intelligence (NLP w/ BERT-Large)	Artificial Intelligence (Recommender w/ DLRM)
	50 SERVERS	50 SERVERS	50 SERVERS	50 SERVERS
Number of 5th Gen Intel Xeon processor-based servers	**29 SERVERS**	**14 SERVERS**	**13 SERVERS**	**10 SERVERS**
Lower fleet energy	870 MWh	1,482 MWh	1,644 MWh	1,705 MWh
TCO savings[2]	**$442K[3]**	**$1,198K[4]**	**$1,273K[5]**	**$1,425K[6]**

3. Be more power efficient.

Modernizing infrastructure using Intel Xeon processors delivers TCO advantages that are even more substantial when replacing older equipment.

5 YEAR REFRESH

5th Gen vs. 1st Gen Intel® Xeon® Processors

Web - NGINX TLS[7]

NGINX

UP TO **40%** TCO Savings[2]

REPLACE **50** 1st Gen Xeon WITH **7** 5th Gen Xeon

~840 MWh Fleet Energy Savings[2]

5th Gen vs. 1st Gen Intel Xeon Processors

BERT-Large - AI - NLP[1]

Google BERT

UP TO **77%** TCO Savings[2]

REPLACE **50** 1st Gen Xeon WITH **3** 5th Gen Xeon

~1700 MWh Fleet Energy Savings[2]

NEW SERVER DECISION

5th Gen vs. 3rd Gen Intel Xeon Processors

Data Services[4]

RocksDB

UP TO **61%** TCO Savings[2]

BUY **14** 5th Gen Xeon VS. **50** 3rd Gen Xeon

~1480 MWh Fleet Energy Savings[2]

5th Gen vs. 3rd Gen Intel Xeon Processors

AI - Recommender[6]

DLRM

UP TO **72%** TCO Savings[2]

BUY **10** 5th Gen Xeon VS. **50** 3rd Gen Xeon

~1700 MWh Fleet Energy Savings[2]

4. Get more performance than with AMD.

In mainstream deployments, 5th Gen Xeon bests competition in throughput and efficiency on workloads that matter most to customers.

5th Gen Intel® Xeon® 8592+ (64C) vs AMD EPYC 9554 (64C)[8]

Higher is Better

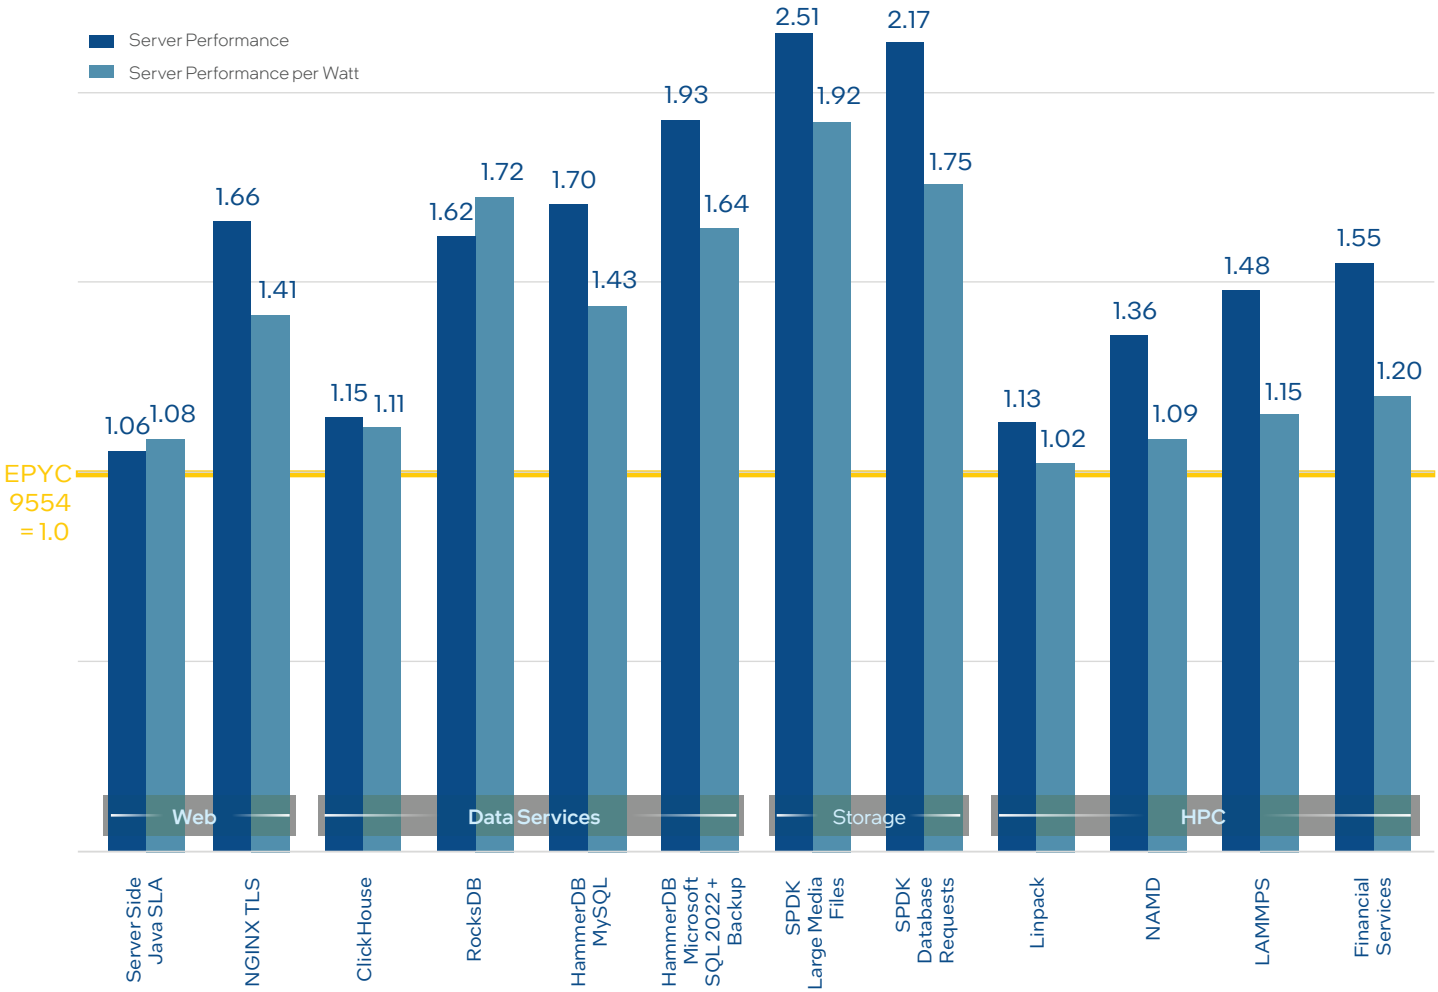

5. Better utilize infrastructure to support critical apps.

Achieve better cost savings and sustainability versus the competition without compromising on performance.

A comparison against 50 4th Gen AMD EPYC 9554 servers

	Web (NGINX TLS)[9]	Data Services (RocksDB)[10]	Data Services MySQL[11]	HPC Monte Carlo[12]	AI - NLP DistilBERT[13]
	50 SERVERS	50 SERVERS	50 SERVERS	50 SERVERS	50 SERVERS
5th Gen Intel® Xeon® Servers	**31 SERVERS**	**31 SERVERS**	**30 SERVERS**	**28 SERVERS**	**15 SERVERS**
Fleet Energy Saved[2]	489.7 MWh	1218.1 MWh	684.0 MWh	585.8 MWh	1496.5 MWh
TCO Savings[2]	$444K	$471K	$509K	$561K	$1,300K
TCO Delta	**21% savings**	**22% savings**	**24% savings**	**27% savings**	**62% savings**

Better utilize infrastructure to support critical apps

Intel invests heavily in co-engineering relationships across the industry with leading software vendors, equipment manufacturers and system integrators. Early and ongoing enablement helps ensure that popular enterprise software provides the greatest performance and cost efficiency possible, whether in the cloud or on-prem. In fact, 90% of developers are using software developed or optimized by Intel.[14]

The benefits of Intel enablement for the software ecosystem are compounded in the complex combinations of solutions that are the backbone of modern enterprises. The new Express Storage Architecture (ESA) introduced in VMware vSphere 8.0, together with the latest Intel technologies, enables generational performance and latency improvements for VMware vSAN implementations. ESA is a capability of vSAN that processes and stores data with improved efficiency, scalability and performance. For more information, read the solution design brief, "Boost Performance and Lower Latency with VMware vSAN 8 and 4th Gen Intel Xeon Scalable Processors."

Recent testing uses vSAN ESA on 4th Gen Intel Xeon processors with four nodes to compare the HCIBench throughput to vSAN OSA (Original Storage Architecture) on 1st Gen Xeon processors with four nodes. The results not only showed the potential for lower operational costs with lower hardware, space and energy requirements, but also more than 7.4x improvement in performance. This work also projects a 10.5:1 server-consolidation ratio from 1st Gen to 4th Gen. Learn more in the blog, "Beyond Savings: How Server Consolidation with VMware vSAN 8 Boosts Performance by more than 7.4x!"

Where infrastructure and software are outdated, the business might be challenged to securely run database and analytics workloads across hybrid, private/public cloud and on-prem resources. Modern solutions can draw on optimizations for a broad set of workloads, from databases and web serving to VDI and storage infrastructure. They support any type of cloud deployment and readily combine on-prem data with cloud analytics. IT can manage the entire data estate in a unified way, with greater operational efficiency and security for everyday tasks such as adding more data and users. Streamlined implementation and management reduce initial and ongoing costs.

Customer Callout

Netflix uses AI inference extensively for video delivery and recommendations, and it relies on Intel's AI software suite and Intel® Xeon® processors for the full end-to-end pipeline: engineering data, model creation-optimization-tuning and deployment. Ongoing collaboration between Intel and Netflix on profiling and architectural analysis helps break through performance bottlenecks. Learn more by reading the blog "Deploying AI Everywhere at Netflix."

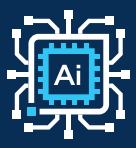

Cost effective considerations for deploying AI

Integrating AI into your environment unlocks advantages in agility, innovation and security. It helps transform data center and cloud environments, streamlining and automating operations for greater efficiency and speed. That can help reduce costs by optimizing resource usage and reducing downtime, while scaling dynamically to make infrastructure more adaptive.

- **Optimize your clouds with AI**: Dr Migrate, Densify and Intel® Granulate™ all use AI models for analysis that improve cost efficiency at every stage of the cloud migration journey. Learn more.
- **AI on Cisco:** Balance performance with cost using hardware you already operate for other workloads. Built-in accelerators instead of discrete devices reduce energy use, operating costs and environmental footprint. Learn more.
- **Deploy generative AI cost-effectively:** Extend existing infrastructure using Lenovo ThinkSystem servers, without investing in dedicated accelerators. Learn more.

Customer Callout

Legal firm Ropers Majeski collaborated with Intel, Activeloop and ZERO Systems on a generative AI solution to relieve knowledge workers of manual tasks such as documenting, filing, timekeeping, storing and information retrieval. The automated solution reduced costs by increasing worker productivity by 18.5% while also improving accuracy. Learn more by reading the customer story, "Ropers Majeski Boosts Productivity and Accuracy."

Optimize existing tech

Many companies seeking reduced costs by migrating to cloud infrastructure fall short of their goals. In fact, they find that public cloud adoption has actually caused their costs to rise. Optimizing performance and tuning cloud instance choices is a vital component of attaining the full TCO savings potential from cloud adoption.

Moving to the cloud can save you money or cost you money.

Why does the cloud seem more expensive?

- Developers overprovisioning
- Poor cloud density
- Paying for hardware with features that have not been turned on, optimized or tuned
- Purchasing more cores than you need
- Workloads might be on older hardware than you realize
- Not utilizing all the compute resources that you're paying for
- Deploying apps into the cloud without knowing what resources to assign to those applications

Optimize what you're already using

Cloud instance optimization

Cheap instances can actually cost you more

skip to subsection ›

Pre-migration assessment

AI-enabled planning to optimize your migration path

Dr Migrate

skip to subsection ›

Cost improvement

Advanced machine learning and analytics to recommend optimal instance choices

Densify

skip to subsection ›

Automated optimization

Using AI to tune apps while executing in real time to get more value from resources

intel Granulate

skip to subsection ›

Cheap instances can actually be expensive

Regardless of which public cloud provider you use, hundreds of instance types will be available to choose from. It's common for customers to rely on automated recommendations from the CSP as a way of cutting through that complexity. While those recommender systems typically make good, generalized suggestions, they may fall short on providing the most cost-optimized approach possible.

In fact, your choice of instance type is central to whether cloud technology delivers cost benefits or becomes a liability. With more performant instances, you may be able to deploy smaller or fewer instances, reducing your rental fees and license costs.

Time is money. Cloud pricing is calculated from how much time it takes to run a workload to completion. In general, the less time it takes your workload to run, the less you pay.

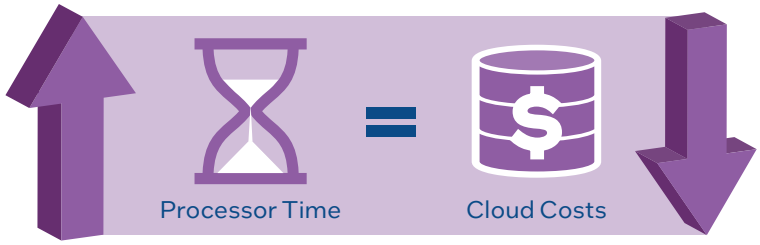

Reduce cloud-licensing costs. Selecting a higher-performing processor can contribute to lower cloud licensing costs. Many application subscriptions are based on the number of vCPUs in the user's cloud instance.

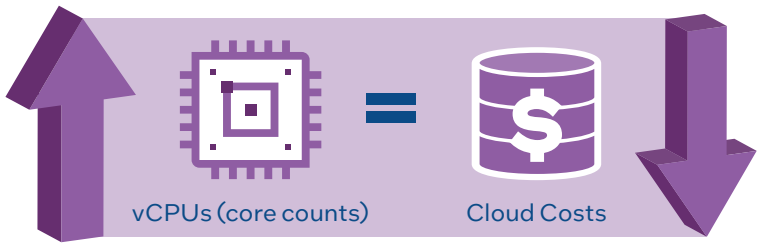

Boost performance per dollar. Your instance type choice can also enable you to run more transactions per second per dollar. You start lowering instance costs as soon as you open them.

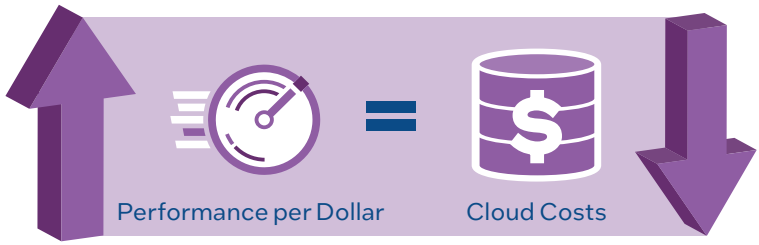

One important consideration with any instance recommender, whether automated or manual, is that most of them will not help you find the optimal settings for a selected instance. A misconfigured cloud environment can generate recurring extra charges until you can troubleshoot and resolve the problem(s). An analyzer tool, such as the Intel Granulate optimizer and migration tool for Intel-based instances, can help ensure that your cloud environment is properly configured. For more information, read the technical research study, "Cloud Computing: Why You Should Be Looking Under the Hood."

There are continuous cost and performance improvements as new public cloud instances are introduced from the major providers. One innovative example is the new AWS M7i-flex instances, which are designed to provide cost savings where workloads do not require full resource availability all of the time. The instances guarantee full performance 95% of the time and at least 40% performance the remaining 5% of the time, in exchange for a 5% discount to customers. According to AWS, M7i-flex instances provide up to 19% better price performance than previous M6i instances.[15] To learn more, see the blog, "Meet the Latest Amazon EC2 Family Members Featuring Intel Processors – M7i and M7i-Flex."

Customer Callout

GUNPOWDER®

The ability to reduce costs by optimizing cloud instances is on full display in Google Cloud-based rendering operations by Gunpowder, a provider of film visual effects. The company credits reduced compute instance time for helping it be more competitive in an industry where price wars can be fierce, landing new business and building a stronger brand. Learn more by reading the customer story, "Gunpowder Cuts Digital Rendering Time and Cost."

Guide your migration path: **Dr Migrate**

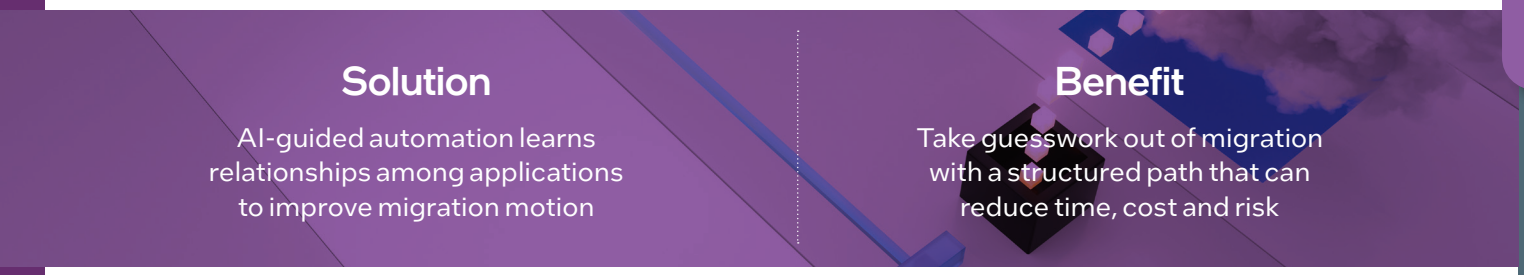

Solution

AI-guided automation learns relationships among applications to improve migration motion

Benefit

Take guesswork out of migration with a structured path that can reduce time, cost and risk

Dr Migrate by LAB3 is an important cloud tool for migration assessments. Dr Migrate provides an AI-guided framework that helps simplify and accelerate cloud migrations. The tool automatically analyzes applications, workloads, connections and resource needs to develop a comprehensive migration plan supporting business goals.

This automated approach to cloud migrations driven by machine learning learns how your applications interconnect and identifies which applications to migrate first and which out-of-date apps you should remove, tuning migration efforts to help lower TCO.

Drive efficiency: **Densify**

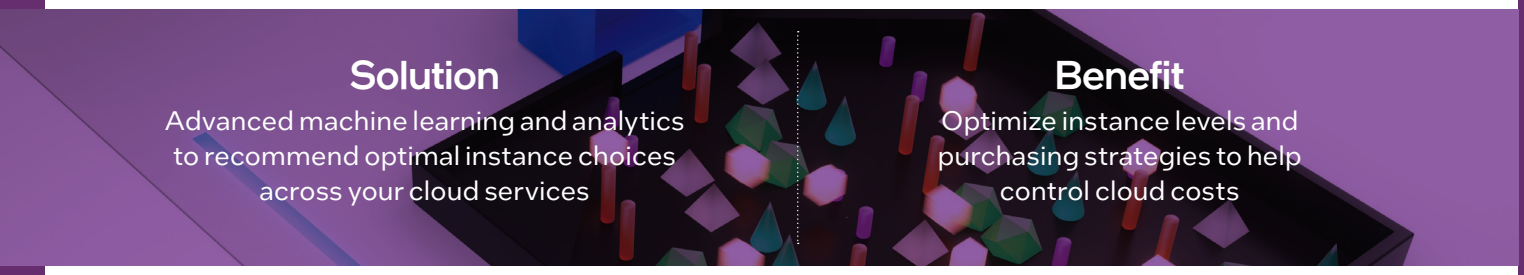

Solution

Advanced machine learning and analytics to recommend optimal instance choices across your cloud services

Benefit

Optimize instance levels and purchasing strategies to help control cloud costs

Optimize your existing cloud infrastructure with Intel Cloud Optimizer by Densify. It offers best-in-class modeling for right-sized and cost-efficient infrastructure, employing a scientific machine learning approach to understanding workload optimization. Densify offers instance-level optimization to help reduce costs at major CSPs, including AWS, Azure and GCP.

- **Measure the efficiency** of your cloud, container and server resource utilization.
- **Get precise recommendations** for cloud instance cost and performance improvements.
- **Optimize instance levels** and address purchasing strategies simultaneously.
- **Enable long-term, continuous optimization** with simplified integration into a cloud management stack.

Real-time optimization: **Intel® Granulate™**

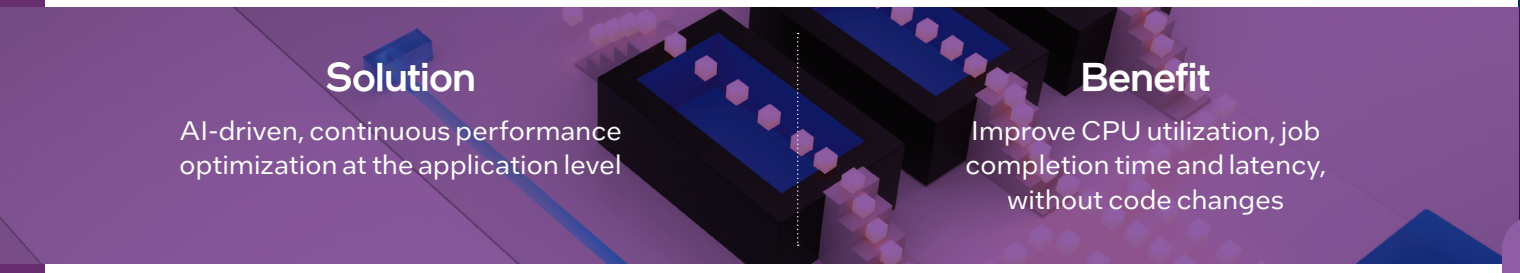

Solution

AI-driven, continuous performance optimization at the application level

Benefit

Improve CPU utilization, job completion time and latency, without code changes

Intel Granulate uses AI and machine learning to map your service's data flows and processing patterns, so it can automatically optimize runtime-level resource management. Its autonomous optimization service addresses inefficiencies in 80% of cloud workloads. Intel Granulate analyzes your application and deploys a customized set of continuous optimizations at runtime, which enables deployment on smaller compute clusters and instance types, potentially reducing costs.

- **Easy to implement**. Implement automated optimization without having to change your code. No developer intervention is required to set it up.
- **Helps even if you're already optimizing**. Help improve performance without re-architecting or recoding, even if you've already employed autoscaling or other optimization methods.
- **Find savings automatically**. Intel Granulate delivers automated continuous optimizations to maintain peak performance without intervention or maintenance.

Intel Telemetry Collector (ITC) can work alongside or independently of Intel Granulate, to monitor and analyze what apps are using the most memory, where resource contention is an issue and where you are using the most power. For more information, read "Cloud Telemetry: Advancing Your IT Strategy."

Customer Callout

Coralogix uses Intel® Granulate™ to reduce compute costs by 45% while cutting average rules-processing time by 30%, increasing throughput by 15% and reducing CPU utilization by 29%. Intel Granulate real-time continuous optimization enables Coralogix to deliver these benefits while continuing to provide the same QoS as before. Learn more by reading the case study, "Coralogix reduces EKS cluster costs by 45% in 2 weeks."

Coralogix

For more information on all the optimization tools:
"How to Get the Most Out of Your Cloud Without the Expense."

Getting **Started**

This section provides a list of resources to get you started.

Implement with these solution providers

- Work with Dell. Dell builds on Intel technologies to deliver performance and power efficiency advantages for advanced workloads.
- Engage with Lenovo. ThinkSystem servers and ThinkAgile hyperconverged infrastructure solutions provide a flexible, solid foundation for innovation.
- Modernize with HPE. Drive winning outcomes and set the stage for future growth with flexible, cloud-smart solutions engineered for the edge.
- Connect through the Intel Partner Directory. This ecosystem offers a broad spectrum of solutions to deliver advanced features and capabilities for the enterprise.

Optimize specific workloads

- Transformative cost benefits with Intel and Google Cloud. Scalable solutions provide compelling TCO for the broadest range of evolving business needs.
- NLP energy cost savings with Red Hat® Open®Shift®. Modernizing with 5th Gen Intel Xeon processors can increase both performance and performance per watt for NLP inference on Red Hat OpenShift.
- Server consolidation with VMware vSAN. Updating hardware along with vSAN software reduces resource requirements for your server fleet as it boosts performance.
- Intel and vSAN modernization. Modernizing infrastructure with vSAN can help lower TCO with performance improvements and increased operational efficiency.
- Intel and Cloudera Data Platform. Faster, easier data management and analytics reduce operational overhead, accelerate time to value and increase control over infrastructure costs.
- Apache Spark cost efficiency on AWS. Increasing performance of decision support systems while lowering costs delivers more value from data within a fixed budget.
- Microsoft Azure Arc on Azure HCI. Combined compute, storage and networking in a single system helps reduce costs with lower power consumption, space requirements and cooling costs.
- Microsoft SQL Server on Intel Xeon processors. Power savings, significantly easier administration and unified data governance and management reduce TCO for database deployments.

Get started with cloud optimization

- Pre-migration planning with Dr Migrate
- Intel Cloud Optimizer by Densify
  - Densify self-guided training. Separate training paths are available for cloud engineers and container users, along with access to Densify online help.
  - Densify resource library. This curated set of materials can help you get the most out of Densify in your environment.
- Intel Granulate

Fine-tune timing and scalability

- Intel Xeon Processor Advisor. Tailor product and solution recommendations for systems and instances, access up-to-date specifications and calculate TCO and ROI for data center solutions.
- Intel Optimization Hub. Select the best mix of technology building blocks such as hardware accelerators, software builds, open-source libraries and drivers, recipes and benchmarks. Optimizations as code are provided in a curated repository across use cases and workloads.
- Intel Developer Zone. Explore development topics, resources and subscriptions including programs, tools, documentation, training, technologies, events and more.

[1] Measurements on Natural language processing/BERT-Large; estimated over 4 years. See [T7] at intel.com/processorclaims: 5th Gen Intel Xeon Scalable processors. Results may vary.

[2] Estimated over four years.

[3] See [T9] at intel.com/processorclaims: 5th Gen Intel Xeon Scalable processors. Results may vary.

[4] See [T10] at intel.com/processorclaims: 5th Gen Intel Xeon Scalable processors. Results may vary.

[5] See [T11] at intel.com/processorclaims: 5th Gen Intel Xeon Scalable processors. Results may vary.

[6] See [T12] at intel.com/processorclaims: 5th Gen Intel Xeon Scalable processors. Results may vary.

[7] See [T6] at intel.com/processorclaims: 5th Gen Intel Xeon Scalable processors. Results may vary.

[8] 5th Gen Xeon Mainstream Workload Performance.

Server-Side Java SLA
Intel Xeon 8592+: 1-node, 2x INTEL(R) XEON(R) PLATINUM 8592+, 64 cores, HT on, Turbo on, Total memory 1024GB (16x64GB DDR5 5600 MT/s [5600 MT/s]), BIOS 3B05.TEL4P1, microcode 0x21000161, 2x Ethernet Controller X710 for 10GBASE-T, 1x 1.7T SAMSUNG MZQL21T9HCJR-00A07, Ubuntu 22.04.1 LTS, 5.15.0-78-generic, Server-side Java SLA throughput. Test by Intel as of 10/06/23.

AMD EPYC 9554: 1-node, 2x AMD EPYC 9554 64-Core Processor, 64 cores, HT on, Turbo on, Total memory 1536GB (24x64GB DDR5 4800 MT/s [4800 MT/s]), BIOS 1.5, microcode 0xa10113e, 2x Ethernet Controller 10G X550T, 1x 1.7T SAMSUNG MZ1L21T9HCLS-00A07, Ubuntu 22.04.3 LTS, 5.15.0-78-generic, Server-side Java SLA throughput. Test by Intel as of 10/24/23.

NGINX TLS
Intel Xeon 8592+: 1-node, 2x 5th Gen Intel Xeon Scalable processor (64 core) with integrated Intel Quick Assist Technology (Intel QAT), QAT device utilized=4(1 active socket), HT on, Turbo off, SNC on, with 1024GB DDR5 memory (16x64 GB 5600), microcode 0x21000161, Ubuntu 22.04.3 LTS, 5.15.0-78-generic, 1x 1.7T SAMSUNG MZWLJ1T9HBJR-00007, 1x Intel® Ethernet Network Adapter E810-2CQDA2, 1x100GbE, NGINX Async v0.5.1, OpenSSL 3.1.3, IPP Crypto 2021.8, IPsec MB v 1.4, QAT_Engine v 1.4.0, QAT Driver 20.l.1.1..20-00030, TLS 1.3 Webserver: ECDHE-X25519-RSA2K, tested by Intel October 2023.

AMD EPYC 9554: 1-node, AMD platform with 2x 4th Gen AMD EPYC processor (64 cores), SMT on, Core Performance Boost off, NPS1, Total memory 1536GB (24x64GB DDR5-4800), microcode 0xa10113e, Ubuntu 22.04.3 LTS, 5.15.0-78-generic, 1x 1.7T SAMSUNG MZWLJ1T9HBJR-00007, 1x Intel® Ethernet Network Adapter E810-2CQDA2, 1x100GbE, NGINX Async v0.5.1, OpenSSL 3.1.3, TLS 1.3 Webserver: ECDHE-X25519-RSA2K, tested by Intel October 2023.

ClickHouse
Intel Xeon 8592+: 1-node, 2x 5th Gen Intel Xeon Scalable processor 8592+ (64 cores) with integrated Intel In-Memory Analytics Accelerator (Intel IAA), Number of IAA device utilized=4 (1 sockets active), HT on, Turbo on, SNC off, Total memory 1024GB (16x64GB DDR5-5600), microcode 0x21000161, 2x Ethernet Controller 10-Gigabit X540-AT2, 1x 1.7T SAMSUNG MZQL21T9HCJR-00A07, Ubuntu 22.04.3 LTS, 6.5.0-060500-generic, ZSTD v1.5.0, QPL v1.3dev, accel-config-v4.1.1, clang13, Clickhouse 21dev, Star Schema Benchmark, Query 4.1, tested by Intel October 2023.

AMD EPYC 9554: 1-node, AMD platform with 2x 4th Gen AMD EPYC processor (64 cores), SMT on, Core Performance Boost on, NPS1, Total memory 1024GB (16x64GB DDR5-4800), microcode 0xa10113e, 2x Ethernet Controller 10G X550T, 1x 1.7T SAMSUNG MZQL21T9HCJR-00A07, Ubuntu 22.04.3 LTS, 6.5.0-060500-generic, ZSTD v1.5.0, clang13, Clickhouse 21dev, Star Schema Benchmark, Query 4.1, tested by Intel October 2023.

RocksDB
Intel Xeon 8592+: 1-node, 2x 5th Gen Intel Xeon Scalable processor 8592+ (64 cores) with integrated Intel In-Memory Analytics Accelerator (Intel IAA), Number of IAA device utilized=8(2 sockets active), HT on, Turbo on, SNC off, Total memory 1024GB (16x64GB DDR5-5600), microcode 0x21000161, 2x Ethernet Controller 10-Gigabit X540-AT2, 1x 1.7T SAMSUNG MZQL21T9HCJR-00A07, Ubuntu 22.04.3 LTS, 6.5.0-060500-generic, QPL v1.2.0, accel-config-v4.0, iaa_compressor plugin v0.3.0, ZSTD v1.5.5, gcc 10.4.0, RocksDB v8.3.0 trunk (commit 62fc15f) (db_bench), 4 threads per instance, 64 RocksDB instances, tested by Intel October 2023.

AMD EPYC 9554: 1-node, AMD platform with 2x 4th Gen AMD EPYC processor (64 cores), SMT on, Core Performance Boost on, NPS1, Total memory 1024GB (16x64GB DDR5-4800), microcode 0xa10113e, 2x Ethernet Controller 10G X550T, 1x 1.7T SAMSUNG MZQL21T9HCJR-00A07, Ubuntu 22.04.3 LTS, 6.5.0-060500-generic, ZSTD v1.5.5, gcc 10.4.0, RocksDB v8.3.0 trunk (commit 62fc15f) (db_bench), 4 threads per instance, 28 RocksDB instances, tested by Intel October 2023.

HammerDB MySQL
Intel Xeon 8592+: 1-node, 2x Intel Xeon Platinum 8592+, 64 cores, HT on, Turbo on, NUMA 2, Integrated Accelerators Available [used]: DLB 8 [0], DSA 8 [0], IAX 8 [0], QAT 8 [0], Total memory 1024GB (16x64GB DDR5 5600 MT/s [5600 MT/s]), BIOS 2.0, microcode 0x21000161, 2x Ethernet Controller X710 for 10GBASE-T, 1x 1.7T SAMSUNG MZQL21T9HCJR-00A07, 2x 1.7T SAMSUNG MZWLJ1T9HBJR-00007, Ubuntu 22.04.3 LTS, 5.15.0-84-generic, HammerDB Mv4.4, MySQL 8.0.33. Test by Intel as of 10/04/23.

AMD EPYC 9554: 1-node, 2x AMD EPYC 9554 64-Core Processor, 64 cores, HT on, Turbo on, NUMA 2, Integrated Accelerators Available [used]: DLB 0 [0], DSA 0 [0], IAX 0 [0], QAT 0 [0], Total memory 1536GB (24x64GB DDR5 4800 MT/s [4800 MT/s]), BIOS 1.5, microcode 0xa10113e, 2x Ethernet Controller X710 for 10GBASE-T, 1x 1.7T SAMSUNG MZQL21T9HCJR-00A07, 2x 1.7T SAMSUNG MZWLJ1T9HBJR-00007, Ubuntu 22.04.3 LTS, 5.15.125-0515125-generic, HammerDB v4.4, MySQL 8.0.33. Test by Intel as of 10/05/23.

HammerDB Microsoft SQL Server + Backup
Intel Xeon 8592+: 1-node, 2x 5th Gen Intel Xeon Scalable processor 8592+ (64 cores) with integrated Intel Quick Assist Technology (Intel QAT), Number of IAA device utilized=8(2 sockets active), HT on, Turbo on, SNC off, Total memory 1024GB (16x64GB DDR5-5600), microcode 0x21000161, 2x Ethernet Controller 10-Gigabit X540-AT2, 7x 3.5T INTEL SSDPE2KE032T807, QATZip 2.0.W.1.9.0-0008, Microsoft Windows Server Datacenter 2022, Microsoft SQL Server 2022, SQL Server Management Studio 19.0.1, HammerDB 4.5, tested by Intel October 2023.

AMD EPYC 9554: 1-node, AMD platform with 2x 4th Gen AMD EPYC processor (64 cores), SMT on, Core Performance Boost on, NPS1, Total memory 1536GB (24x64GB DDR5-4800), microcode 0xa10113e, 2x Ethernet Controller 10G X550T, 7x 3.5T INTEL SSDPE2KE032T807, Microsoft Windows Server Datacenter 2022, Microsoft SQL Server 2022, SQL Server Management Studio 19.0.1, HammerDB 4.5, tested by Intel October 2023.

SPDK 128K QD64 (large media files) / SPDK 16K QD256 (database requests)

Intel Xeon 8592+: 1-node, 2x 5th Gen Intel Xeon Scalable processor (64 core) with integrated Intel Data Streaming Accelerator (Intel DSA), DSA device utilized=1(1 active socket), HT on, Turbo on, SNC off, with 1024GB DDR5 memory (16x64 GB 5600), microcode 0x21000161, Ubuntu 22.04.3 LTS, 5.15.0-78-generic, 1x 894.3G Micron 7450, 4x 3.84TB Samsung PM1733, 1x Intel® Ethernet Network Adapter E810-2CQDA2, 2x100GbE, FIO v3.34, SPDK 22.05, tested by Intel October 2023.

AMD EPYC 9554: 1-node, AMD platform with 2x 4th Gen AMD EPYC processor (64 cores), SMT on, Core Performance Boost on, NPS2, Total memory 1536GB (24x64GB DDR5-4800), microcode 0xa10113e, Ubuntu 22.04.3 LTS, 5.15.0-78-generic, 1x 1.7T Samsung PM9A3, 4x 3.84TB Samsung PM1733, 1x Intel® Ethernet Network Adapter E810-2CQDA2, 2x100GbE, 1x Ethernet Connection X550T for 10GBASE-T, FIO v3.34, SPDK 22.05, tested by Intel October 2023.

LINPACK
Intel Xeon 8592+: 1-node 2x Intel Xeon 8592+, HT on, Turbo on, SNC2, 1024 GB DDR5-5600, ucode 0x21000161, Red Hat Enterprise Linux 8.7, 4.18.0-425.10.1.el8_7.x86_64, HPL from MKL_v2022.1.0, cmkl:2023.2.0, icc:2023.2.0, impi:2021.10.0. Test by Intel as of October 2023.

AMD EPYC 9554: 1-node, 2x AMD EPYC 9554, SMT on, Turbo on, CTDP=360W, NPS=4, 1536GB DDR5-4800, ucode=0xa101111, Red Hat Enterprise Linux 8.7, Kernel 4.18, AMD official binary. Test by Intel as of March 2023.

NAMD (Geomean of apoa1_npt_2fs, stmv_npt_2fs)
Intel Xeon 8592+: 1-node 2x Intel Xeon 8592+, HT on, Turbo on, SNC2, 1024 GB DDR5-5600, ucode 0x21000161, Red Hat Enterprise Linux 8.7, 4.18.0-425.10.1.el8_7.x86_64, NAMD v2.15alpha, cmkl:2023.2.0 icc:2023.2.0 tbb:2021.10.0. Test by Intel as of October 2023.

AMD EPYC 9554: 1-node, 2x AMD EPYC 9554, SMT on, Turbo on, CTDP=360W, NPS=4, 1536GB DDR5-4800, ucode=0xa101111, Red Hat Enterprise Linux 8.7, Kernel 4.18, NAMD v2.15alpha, cmkl:2023.2.0 icc:2023.2.0 tbb:2021.10.0.

LAMMPS (Geomean of Polyethylene,DPD,Copper, Liquid Crystal,Atomic Fluid,Protein, Stillinger-Weber,Tersoff, Water)
Intel Xeon 8592+: 1-node 2x Intel Xeon 8592+, HT on, Turbo on, SNC2, 1024 GB DDR5-5600, ucode 0x21000161, Red Hat Enterprise Linux 8.7, 4.18.0-425.10.1.el8_7.x86_64, LAMMPS v2021-09-29, cmkl:2023.2.0 icc:2023.2.0 tbb:2021.10.0, impi:2021.10.0. Test by Intel as of October 2023.

AMD EPYC 9554: 1-node, 2x AMD EPYC 9554, SMT on, Turbo on, CTDP=360W, NPS=4, 1536GB DDR5-4800, ucode= 0xa101111, Red Hat Enterprise Linux 8.7, Kernel 4.18, LAMMPS v2021-09-29, cmkl:2023.2.0 icc:2023.2.0 tbb:2021.10.0, impi:2021.10.0. Test by Intel as of March 2023.

FSI KERNELS (GEOMEAN OF BINOMIAL OPTIONS, MONTE CARLO, BLACKSCHOLES)
Binomial Options
Intel Xeon 8592+: 1-node 2x Intel Xeon 8592+, HT on, Turbo on, SNC2, 1024 GB DDR5-5600, ucode 0x21000161, Red Hat Enterprise Linux 8.7, 4.18.0-425.10.1.el8_7.x86_64, Binomial Options v1.1, icc:2023.2.0 tbb:2021.10.0. Test by Intel as of October 2023.

AMD EPYC 9554: 1-node, 2x AMD EPYC 9554, SMT on, Turbo on, CTDP=360W, NPS=4, 1536GB DDR5-4800, ucode=0xa101111, Red Hat Enterprise Linux 8.7, Kernel 4.18, Binomial Options v1.1, icc:2023.2.0 tbb:2021.10.0. Test by Intel as of March 2023.

Monte Carlo
Intel Xeon 8592+: 1-node 2x Intel Xeon 8592+, HT on, Turbo on, SNC2, 1024 GB DDR5-5600, ucode 0x21000161, Red Hat Enterprise Linux 8.7, 4.18.0-425.10.1.el8_7.x86_64, Monte Carlo v1.2, cmkl:2023.2.0 icc:2023.2.0 tbb:2021.10.0. Test by Intel as of October 2023.

AMD EPYC 9554: 1-node, 2x AMD EPYC 9554, SMT on, Turbo on, CTDP=360W, NPS=4, 1536GB DDR5-4800, ucode=0xa101111, Red Hat Enterprise Linux 8.7, Kernel 4.18, Monte Carlo v1.2, cmkl:2023.2.0 icc:2023.2.0 tbb:2021.10.0. Test by Intel as of March 2023.

Black-Scholes
Intel Xeon 8592+: 1-node 2x Intel Xeon 8592+, HT on, Turbo on, SNC2, 1024 GB DDR5-5600, ucode 0x21000161, Red Hat Enterprise Linux 8.7, 4.18.0-425.10.1.el8_7.x86_64, Black Scholes v1.4, cmkl:2023.2.0 icc:2023.2.0 tbb:2021.10.0. Test by Intel as of October 2023.

AMD EPYC 9554: 1-node, 2x AMD EPYC 9554, SMT on, Turbo on, CTDP=360W, NPS=4, 1536GB DDR5-4800, ucode=0xa101111, Red Hat Enterprise Linux 8.7, Kernel 4.18, Black Scholes v1.4, cmkl:2023.2.0 icc:2023.2.0 tbb:2021.10.0. Test by Intel as of March 2023.

[9] See [T203] at intel.com/processorclaims: 5th Gen Intel Xeon Scalable processors. Results may vary.

[10] See [T202] at intel.com/processorclaims: 5th Gen Intel Xeon Scalable processors. Results may vary.

[11] See [T201] at intel.com/processorclaims: 5th Gen Intel Xeon Scalable processors. Results may vary.

[12] See [T204] at intel.com/processorclaims: 5th Gen Intel Xeon Scalable processors. Results may vary.

[13] See [T206] at intel.com/processorclaims: 5th Gen Intel Xeon Scalable processors. Results may vary.

[14] Global Development Survey conducted by Evans Data Corp., 2021.

[15] https://www.intel.com/content/www/us/en/newsroom/news/4th-gen-intel-xeon-momentum-grows-in-cloud.html#gs.4hpul6.

Performance varies by use, configuration and other factors. Learn more on the Performance Index site.

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 does not control or audit third-party data. You should consult other sources to evaluate accuracy. Intel technologies may require enabled hardware, software or service activation.

0224/MH/MESH/PDF 353914-001US