

Solution Brief

Public Clouds
AI-Enabled Optimization



Low Cost, High Confidence with Intel Cloud Optimizer by Densify

The Intel Cloud Optimizer (ICO) by Densify offers a flexible, transparent recommendation engine with a multicloud view, helping organizations choose the best-fit public cloud instances for every workload, continuously—and Intel covers the cost when you get started.



“The Intel Cloud Optimizer by Densify recommends the best cloud instances running Intel® technology for your workloads, with meaningful recommendations based on pricing, stability, performance, or other parameters that are important to your business. This is a tremendous value, and Intel covers the cost.”

—Kevin Johnson, general manager of Intel's Cloud and Enterprise Solutions and Technology Sales Group

Public cloud service providers (CSPs) like AWS, Microsoft Azure, and Google Cloud Platform (GCP) are essential to any cloud strategy among enterprises globally, and businesses increasingly depend on them. The valuation of public cloud markets is forecasted to grow to USD 596B by 2027 at a compound annual growth rate (CAGR) of 14.6 percent, driven by improving scalability, efficiency, and data access.¹ The key to competitive advantage in an environment where every business is using the cloud depends on getting the most compute value per dollar spent.

Challenge: Optimization is complex, with many tangential paths

When leveraging cloud infrastructure, numerous factors make it difficult to balance performance and uptime with cost efficiency: what to buy (family), how much to buy (size), and how to ensure the optimal features that will allow specific workloads to perform and scale cost effectively in the cloud (features). These important factors are impossible for humans to determine optimally using manual or estimation methods.

Major CSPs, including AWS, Azure, and GCP, offer onboarding and optimization services free of charge. However, these services and analyses often fall short because they focus on purchase plans and billing optimizations, only offer one-size-fits-all tooling, lack granular controls, are too basic in their analytics, and do not explain how conclusions are reached. What's needed are analytics with the rigor required to derive precise answers and eliminate the guesswork, so decision-makers and cloud architects can be confident they are using the right cloud instances for the right workloads.

Solution: Intel Cloud Optimizer by Densify, with costs covered by Intel

The Intel Cloud Optimizer (ICO) is a Software as a Service (SaaS) technology that offers analytics capability and machine learning (ML)-informed recommendations, above and beyond what CSPs offer. This solution provides a scientific basis for cloud instance recommendations based on a large variety of critical inputs, including up to 90 days of workload history (optionally up to a year), business unit (BU) policies, and technical constraints, all while helping connect end users to the best Intel-enabled resources across a multicloud environment. If a business has workloads running in AWS, Azure, and GCP, the optimizer can recommend the ideal instance to choose across all three providers.

Intel covers the cost when using the optimizer, so users can achieve low billing without any additional software costs. Kevin Johnson, general manager of Intel's Cloud and Enterprise Solutions and Technology Sales Group, comments, “The Intel Cloud Optimizer by Densify recommends the best cloud instances running Intel® technology for your workloads, with meaningful recommendations based on pricing, stability, performance, or other parameters that are important to your business. This is a tremendous value, and Intel covers the cost.”




Figure 1. It can be impossible to manually determine the optimal balance of size, family, and features when choosing cloud instances for your workloads.

Intel Cloud Optimizer key benefits:			
<p>Fast results, fast value</p> <p>Get actionable recommendations within 24 hours, on average.</p>	<p>Granular control</p> <p>Input specific billing and workload requirements.</p>	<p>Exceptional transparency</p> <p>Understand how and why the optimizer makes recommendations.</p>	<p>Costs covered by Intel</p> <p>For up to one year, unlock high value without additional spending.</p>

How it works: Cost-saving recommendations within 24 hours

ICO is a dedicated SaaS offering, and one organization's data is not stored with another's data. "There are no agents to deploy, and this read-only technology is very rapid to enable. It takes less than one hour to connect the optimizer, and users will see their first results and recommendations in 24 hours," says Chuck Tatham, chief revenue officer at Densify. An organization simply inputs their CSP cloud credentials, at which point ICO will begin ingesting up to two months of workload history maintained by the CSP for all identified cloud instances. As the optimizer continues to run, it will store and maintain up to 90 days of workload history to further learn from and refine its recommendations.

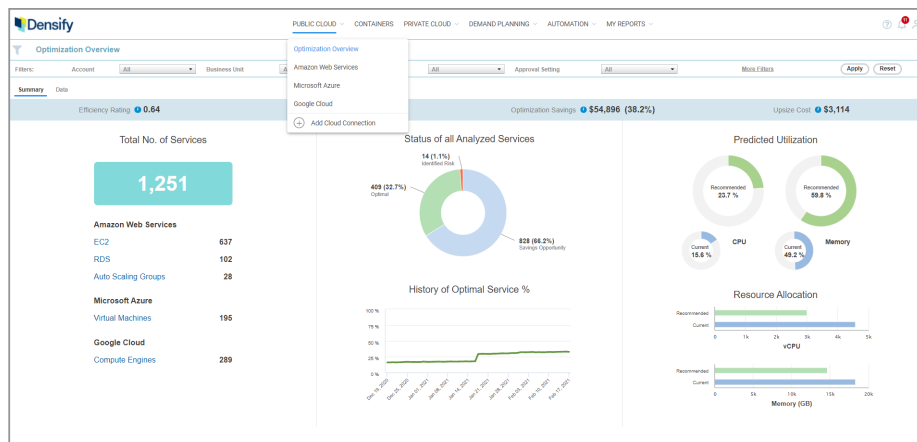


Figure 2. The ICO dashboard makes it easy to visualize savings opportunities across AWS, Azure, and GCP.

Intelligent recommendations that cover all the bases

ICO can be powered with more than just workload history when analyzing CSP instance usage. Users can leverage additional technical or business constraints and flags to ensure that answers are truly "fit for purpose" and actionable. The optimizer uses proprietary analytics models to learn from resource usage patterns and build a profile for each workload across compute and database² instances and then recommend an ideal instance type for each workload. These recommendations are actionable because they include all the required considerations as well as actual detailed workload behaviors, enabling application owners to confidently take action to switch to the recommended instance type to help realize their goals for each workload, such as cost efficiency, uptime, and/or performance thresholds.

Organizations generally begin by making these changes manually, but as the optimizer proves itself over time and confidence grows, platform owners can deploy integrations to Densify APIs to automate the continuous implementation of recommendations.

"For a production-critical workload, you might be more focused on uptime and performance and more conservative on optimization. Whereas with less-sensitive workloads, you can be more aggressive with optimization. The [Intel Cloud Optimizer's] analytics can be tuned to fit the sensitivities of the business."

—Chuck Tatham, chief revenue officer at Densify

Step 1:

Connect to CSP instances

Users entitle ICO to ingest data from their CSPs on a global or granular basis by account, project, or subscription.

Step 2:

Review recommendations

Actionable recommendations are available within 24 hours.

Step 3:

Choose to act

Users can manually switch to recommended CSP instances or use APIs to enable automated switching.

Optional:

Generate reports

Users also get full reports that can be extracted to tables and CSV files for bulk analysis.

Figure 3. Users can get started in minutes and see recommendations in 24 hours, on average.

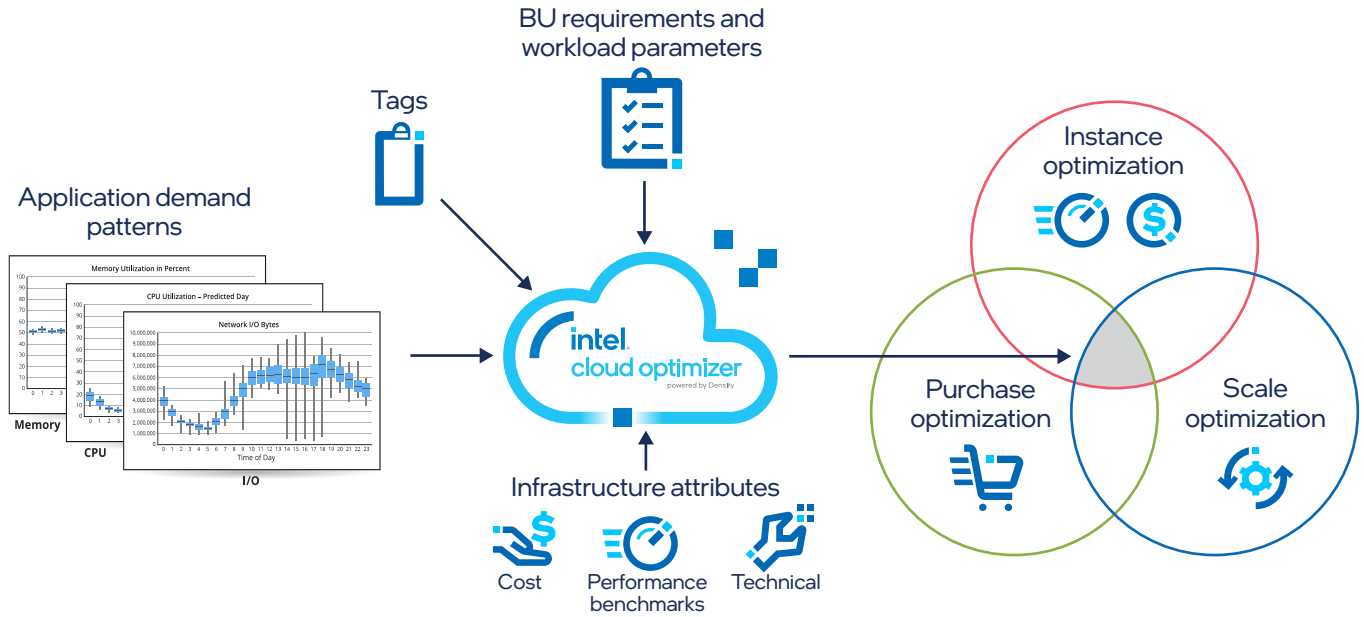


Figure 4. Sophisticated Densify IP powers the optimizer’s ML engine, which considers numerous factors to generate recommendations for the right balance of performance, stability, uptime, and cost efficiency.

Merging hardware expertise with cloud expertise

Intel sought out Densify to help develop a unique flavor of the optimization engine that could showcase the advantages of Intel-enabled nodes, with no obligation to transfer workloads. The goal with ICO was to help customers continuously optimize on the best Intel technology for their specific business needs. With Intel’s expertise in hardware capabilities and Densify’s expertise in the cloud, customers benefit from a best-of-both-worlds approach driven by a deep understanding of cloud use cases and the tools that help enable them. For example, key optimizations that accelerate AI workloads, such as Intel® AVX-512 Vector Neural Network Instructions (VNNI), often go unnoticed or overlooked by conventional algorithms.

Customer success with a major software company

One of the top 20 software vendors in the world by revenue deployed the optimizer for 13,000 AWS instances and was able to achieve an impactful cost savings in a matter of months. Beyond instance optimizations, the company’s development team also noted that the optimizer helped them visualize and demonstrate where their cloud budget was being spent and what it was being spent on. This level of visibility helped build confidence in the team’s efforts while paving the way for higher cost efficiency, without a huge time investment from an already-busy development team. One of the main challenges with any optimization effort is finding the time to optimize; in this case, Densify helped by enabling the company to automate comparing hundreds of different cloud instance options on AWS.

Recommendations evolve with ML over time

The first time a connection is made to ICO, the optimizer analyzes up to two months of workload history (the standard window of historical data stored by CSP metrics systems). The ML engine will continue to refine its recommendations with increasing amounts of data as the data aggregates. For example, over time the algorithms can learn to recognize and anticipate seasonal or time-of-day surges for a given workload. The longer the ML engine learns, the better it gets.

Continuous optimization is everything

Intel Cloud Optimizer helps businesses choose and update the right cloud instance for their workload, even as workloads change. Not only does this impact billing, but it also helps ensure smooth operation for key workloads and data access.

- **Robust stability:** Avoid shorting cloud resources to key workloads
- **High utilization:** Avoid overallocating capacity to instances
- **Few bottlenecks:** Avoid leaving CPU, memory, or other resources stranded

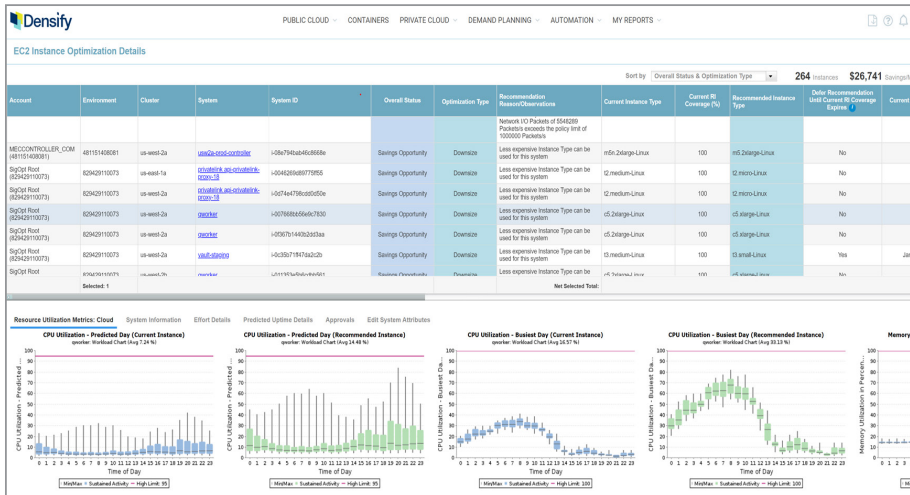


Figure 5. ICO can deliver recommendations within 24 hours on average, made accessible through the Densify web portal.

Building trust through transparency

“Recommendations are only as useful as they are trustworthy. Otherwise, application owners may not act on the recommendations they receive. A key driver in building trust is transparency into how the Intel Cloud Optimizer derived a given recommendation,” says Brenda Bollinger, director of business development at Densify. All recommendations made by the optimizer are accompanied by rich detail on what went into the analysis (policy, workload history, technical constraints) to help paint a full picture of how the optimizer generates conclusions. This proof enables the confidence to actually make the change and reap the benefits.

Differentiated capabilities with Densify

The Intel Cloud Optimizer goes above and beyond what other optimizers can currently do:

- Multicloud analysis across AWS, Azure, and GCP
- Granular control instead of one size fits all, with the ability to ingest technical or business policies
- Deep, industry-proven analytics that also provide supporting detail on why a recommendation is made
- Rich recommendation reports with graphs and exportable data tables
- Extended look-back period with up to 90 days of workload history

About Densify

Based in Toronto, Canada, Densify helps businesses worldwide optimize the cost and performance of business-critical applications through the continuous matching of workload demands and hybrid infrastructure supply.

densify.com

“Recommendations are only as useful as they are trustworthy. Otherwise, application owners may not act on the recommendations they receive. A key driver in building trust is transparency into how the Intel Cloud Optimizer derived a given recommendation.”

—Brenda Bollinger, director of business development at Densify

Learn more

Get started now at densify.com/intel.

Intel will cover the costs for the base cloud optimizer offering.



1. “Public Cloud Market Size Worth \$596.0 Billion by 2027 | CAGR: 14.6%: Grand View Research, Inc.,” pnewswire.com, May 2020. <https://www.pnewswire.com/news-releases/public-cloud-market-size-worth-596-0-billion-by-2027--cagr-14-6-grand-view-research-inc-301052701.html>
 2. Relational Database Service (RDS) support is currently only available with AWS.

Notices and disclaimers

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.

No product or component can be absolutely secure.

Your costs and results may vary.

© 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.

1221/SMR/CMD/PDF