

IT@Intel: Validating and Evolving Intel IT's Multicloud Strategy

A decade into our multicloud journey, our business-value-first approach is clearly driving business value, while we continue to adapt the strategy to changes in business needs and the ecosystem

Intel IT Authors

Sachin Ashtikar
Principal Engineer

Munir Ghamrawi
Principal Engineer

Ravi Hoskote
Principal Engineer

Sridhar Mahankali
Principal Engineer

Subbarao Mantha
Principal Engineer

Tony Valverde
Principal Engineer

Joshua Wise
Principal Engineer

Table of Contents

Executive Summary	1
Background.....	2
The Next Generation of Our Multicloud Strategy.....	3
Our Vision for a Unified Hosting Portal for Our Cloud Management Platform	3
Embracing Open Source for DBaaS...4	
Honing Our Application Hosting Decisions.....	4
Updating Our CaaS Environment... 6	
Supporting the Growth of AI and Automation.....	6
Next Steps.....	6
Conclusion.....	6
Related Content.....	7

Executive Summary

Intel IT adopted a “right workload, right place” cloud strategy nearly 10 years ago. This strategy has accelerated application development and deployment agility, achieved cost-responsible scalability, and enabled our business partners across Intel to achieve their business goals.

A decade of validation has underscored that emphasizing business value in cloud hosting decisions is advantageous. However, we continue to adapt our multicloud strategy as both business needs and the IT ecosystem evolve.

Some of those adaptations include the following:

- Developing a unified hosting portal for our cloud management platform.
- Embracing open-source offerings for our database as a service (DBaaS) solutions.
- Honing our application hosting decisions so our use of software as a service (SaaS), public cloud, and private cloud is optimized.
- Rapidly scaling our containers as a service (CaaS) environment.
- Supporting the growth of AI and automation.

But the cloud journey is never finished. New trends in the enterprise—such as the increasing use of AI, as well as innovations in the IT landscape—mean that we must continue to evaluate and evolve our multicloud strategy to consistently support the business in the most cost-efficient manner possible.

Intel IT Contributors

Munir Ghamrawi, Principal Engineer

Donald Pearson, Industry Engagement Manager

Acronyms

CaaS	containers as a service
DBaaS	database as a service
IaC	infrastructure as code
IaaS	infrastructure as a service
IAO	IT IDM 2.0 Acceleration Office
SaaS	software as a service
XaaS	everything as a service

Background

Nearly a decade ago, we formulated a cloud strategy that helped us decrease application time to market, enable agile application development, optimize cloud costs, and help our customers to achieve their business goals. A key element of our initial strategy was to put business value first, letting the business needs define the workload selection and hosting environments for a particular process or goal.

Unlike many enterprises that pursue a “cloud first” approach, our “right workload, right place” approach frees Intel from force-fitting a workload into a hosting environment that may not be suitable. Instead, our multicloud strategy entails a combination of private and public cloud resources (hybrid cloud and multicloud) that are cost- and agility-optimized. Business needs drive application needs, which in turn drive workload placement.

Right workload, right place decisions are based on the “three Cs”:

- **Capability.** Defines what functionalities the workload needs, then assess the best capabilities that are offered for functionality, performance, latency, and data gravity.
- **Controls.** Verify that the hosting environment has the necessary controls for the workload and data. Examples of controls include security, privacy, and regulatory compliance.
- **Costs.** Define peak and baseline needs for workload agility and cost optimization.

We have observed substantial business benefits from our business-value-focused approach to multicloud:

- Intel gains a high level of business velocity and agility, with built-in redundancy and resiliency as a result of everything-as-a-service (XaaS) capabilities.
- A cloud-native environment furnishes a consistent multicloud user experience across the enterprise.
- A simplified cloud stack provides outstanding application portability.
- We enable best-fit workload placement through proactive architecture guidance and governance, which is supported by automation.

The IT IDM 2.0 Acceleration Office (IAO) is an excellent example of where our multicloud strategy delivers significant business value. We formed IAO in 2023 with the purpose of working closely with business units and functional teams to accelerate Intel’s IDM 2.0 evolution¹ (see [sidebar](#) for details).

Industry trends lend further credence to our multicloud approach. For example, a Barclays survey revealed that 83% of leaders surveyed were planning to bring workloads from the public cloud back on-premises to private cloud infrastructure.² This evolution from a “cloud first” mentality to one that is more similar to our long-standing “right workload, right place” strategy is driven by several factors, including cost, data security, and the realization that an overly simplified lift-and-shift approach to moving services to the cloud—without refactoring—falls short of unlocking the value of the public cloud.

Our Multicloud Strategy in Action

Intel is undergoing a huge transformation that involves separating product research and development and manufacturing into separate entities, as well as ramping up manufacturing production capacity for third-party customers. We are meeting the increasing enterprise demand for AI and supporting Intel’s ongoing transformation through the IT IDM 2.0 Acceleration Office (IAO), working closely with business units and functional teams as they bring a new internal foundry model to life.

IAO has embarked on a massive upgrade to Intel’s enterprise resource planning (ERP) system to robustly support Intel’s products and the Intel foundry business units. This brings an incredible opportunity to reduce costs, with a goal to remove about 40% of our legacy applications over the next three years. Other IAO focus areas will include supply chain planning and data analytics. IAO will help the internal foundry model restore Intel’s healthy margins, establish a competitive cost structure, and enhance the culture and incentives required to deliver process and product leadership.

IAO uses different cloud environments for different purposes, which illustrates how hosting decisions based on cost, capabilities, and controls are made:

- The transformed ERP system, certain boundary applications, and integration tools are hosted on-premises.
- In contrast, IAO’s data lake, data warehouse, and analytics tools are hosted on public cloud infrastructure.
- IAO’s best-in-class SaaS solutions in the areas of finance, planning, customer relationship management (CRM), human capital management (HCM), and B2B integrations are also integrated to build a complete end-to-end solution.

¹ “IDM 2.0” is a major evolution of Intel’s integrated device manufacturing (IDM) model. For more information, read Intel’s [press release](#) from March 2021, “Intel CEO Pat Gelsinger Announces ‘IDM 2.0’ Strategy for Manufacturing, Innovation and Product Leadership.”

² RCR Wireless News, September 2024. <https://www.rcrwireless.com/20240909/network-infrastructure/survey-says-83-of-cios-are-embracing-private-cloud>

However, while we can acknowledge that we made a wise choice by starting out on the right multicloud path from the beginning, we never stop searching for ways to provide even more business value to Intel by evolving our multicloud strategy as we see changes in business needs and the ecosystem.

The Next Generation of Our Multicloud Strategy

Since the publication of our previous multicloud strategy white paper,³ we have added new capabilities and evolved others.

The following sections provide more details about some of these adaptations:

- Developing a unified hosting portal for our cloud management platform
- Embracing open-source offerings for our database as a service (DBaaS) solutions
- Honing our application hosting decisions
- Rapidly adopting and scaling our containers as a service (CaaS) environment
- Supporting the growth of AI and automation




Figure 1. We are developing a unified hosting portal that will simplify access to the cloud management platform.

³ See our IT@Intel paper, "Intel IT's Multi-Cloud Strategy: Focused on the Business."

Our Vision for a Unified Hosting Portal for Our Cloud Management Platform

Our multicloud strategy supports multiple private and public clouds as well as hybrid deployments; for this reason, we are developing a unified hosting portal for our cloud management platform (see Figure 1). This portal will simplify processes for our business partners at Intel and deliver business velocity and agility. We aim to use the portal to transform hosting services and capabilities by using automation and AI technologies to accelerate customer value and increase IT efficiency.

When complete, the portal will enable cloud consumers to access a common dashboard that includes available cloud hosting capabilities. The dashboard will provide quick access to a services-and-API catalog, automated orchestration, infrastructure as a service (IaaS), backup as a service, storage as a service, CaaS, DBaaS, public cloud integration, and more. Some unique features of the portal roadmap include the following:

- **Workload guidance.** When cloud consumers enter the hosting portal, they will interact with an automated workload guidance advisor that recommends workload placement. The automated advisor will be available 24/7 and will provide immediate access to the relevant cloud services. Cloud consumers won't have queue up in a weeks-long waiting list to talk with a cloud broker agent.
- **Application template and blueprint catalog.** This feature will provide templates for cloud consumers, which they can customize based on their needs and deploy automatically across multiple environments. This feature will accelerate application implementation while significantly reducing operational overhead for deploying applications.
- **Infrastructure as code (IaC).** Cloud consumers will be able to tap into a rich catalog of APIs to consume infrastructure services. IaC is a method of provisioning infrastructure using scripts instead of manual processes. This will enable us to manage infrastructure in the same way we manage application code today—versioned, reviewed, tested, and deployed.
- **Intelligent monitoring and notifications.** The portal will enable us to use AI to perform smart monitoring of consumed cloud services and send performance notifications to the consumer based on customized threshold settings and other data. The exact monitoring capabilities will vary based on the hosting environment. For example, we already use in-house tools for application and infrastructure monitoring in our private cloud. For workloads hosted in the public cloud, we take advantage of the cloud service provider's native monitoring and alerting capabilities.

Overall, the unified hosting portal will deliver ease of use, faster workload placement decisions, and cloud-native application portability. Because the portal will be integrated with our application portfolio management system, it will support the application lifecycle management process. The portal is part of our three-year rolling roadmap for multicloud strategy, and we will continue to improve it over time.

Embracing Open Source for DBaaS

As part of our application platform suite, we offer internally developed DBaaS capabilities, which speed development and increase redundancy and resiliency through database abstraction. With DBaaS, developers no longer have to request infrastructure, build the database, and manage it—today, they can quickly deploy a secure, scalable, and highly available new cost-optimized database with a few mouse clicks.

We embrace open-source DBaaS solutions for optimized total cost of ownership and scalability without requiring per-core licensing—something we call “cost-responsible scalability.” Our DBaaS framework enables us to quickly deploy and support new database technologies as the open-source application landscape and diversity of use cases continue to evolve rapidly. Currently, **nearly 2,500 applications at Intel—some owned by Intel IT and some by the business units—use DBaaS.**

Intel IT Multicloud Hosting Options




Figure 2. Where an application is hosted depends on the “three Cs” (capabilities, controls, and cost structure) as well as other factors.

Honing Our Application Hosting Decisions

Hosting choices span private cloud in our own data centers, various public cloud service providers, and SaaS (see Figure 2). For example:

- We typically host differentiating applications in our private cloud.
- We use SaaS extensively for portions of customer relations management, human capital management, logistics, and other applications where Intel does not have a differentiating advantage.

- We use public cloud to support our enterprise cloud analytics strategy and other applications that may require certain specialized cloud-native capabilities or benefit from large on-demand scalability requirements.⁴

Several factors help determine the optimal hosting choice for an application:

- **Workload type.** Applications with a predictable, 24/7 workload are often good choices for private cloud hosting. On the other hand, applications with an unpredictable workload or quickly increasing demand might be best suited for the public cloud’s rapid elasticity and scalability.
- **Lifespan and lifecycle.** Applications with a short-term lifecycle can be deployed quickly in the public cloud. Applications that will be used long term can generally be hosted most efficiently on-premises. Similarly, applications that are being tested or used in a proof of concept can take advantage of a public cloud sandbox that offers rapid deployment. Established production workloads are often better hosted on-premises.
- **Controls.** Public cloud service providers may or may not deliver the right security and privacy controls for certain applications.
- **Capabilities.** The public cloud often offers innovation technologies—like analytics as a service and machine learning as a service—more quickly than is possible in an on-premises data center.
- **Cost.** To decide whether public or private cloud is best, we determine a five-year total cost of ownership. All factors are considered, including hardware, software, integration, hosting, and resourcing costs. We also consider any available enterprise or incentive discounts that a cloud service provider may have available.
- **Talent pool.** We generally do not move non-cloud-native applications to the cloud—it doesn’t make good sense to simply move technical debt to the cloud. For example, if a business unit doesn’t have the talent pool to develop and deliver a cloud-native application, we usually recommend it be hosted privately.
- **Data and process affinity.** Data integration requirements between systems can affect the hosting decision because these systems may need to be placed close together if there’s high data affinity. For example, an application may need general ledger account data and sales transaction data—both of which are hosted in our private cloud. Another application may need human capital management information, which is hosted on external SaaS.

Application rationalization is integral to making informed, optimal cloud hosting decisions and reducing technical debt. As described in our earlier cloud strategy paper,⁵ the application rationalization process determines if an application is producing business value and if it is effectively using the resources assigned to it from a business-value perspective.

⁴ See our IT@Intel paper, “Modernizing Enterprise Data Analytics Using Databricks in the Cloud.”

⁵ See our IT@Intel paper, “Intel IT’s Multi-Cloud Strategy: Focused on the Business.”

As we continue our journey of reducing technical debt, we've modernized a substantial number of applications, modified business processes to better take advantage of investments in core platforms, and removed a significant number of applications from our IT landscape (see Figure 3). Recent adaptations to our application rationalization process include improved application categorization, better data quality, and increased in-house application support.

Application Rationalization Results




Figure 3. Application rationalization has enabled us to dramatically reduce the number of applications in use at Intel.

Application Categorization

We emphasize achieving better application categorization by recognizing that application categorization is not simply bi-modal (public cloud app or an internal app). The distinctions between applications are more nuanced than that. We now categorize applications as follows:

- Discrete—hosted exclusively in the public cloud.
- Internal hybrid—application environments that can be hosted across public and private cloud but with no external exposure.
- External hybrid—application environments with components that are reachable from the internet and appropriately secured.

Regardless of the type of application, we provide centrally managed guardrails, security, and agility to our business partners at Intel.

Better Data Quality

Another aspect of our application rationalization process is improving application data quality. We are committed to creating a single version of application truth that is available to all business units and employees. With trusted, centralized data about applications and platforms, employees can quickly access pre-approved tools, which facilitates rapid adoption. In addition, the application metadata includes key contacts, which enables employees to follow up with questions if needed.

In-house Application Support

In our own experience, along with conversations with other enterprises, we have found that application support cost structures vary significantly between private data centers, managed service providers, or hyperscalers' IaaS

hosting services. This is especially true with commercial off-the-shelf (COTS) applications, which often require specific configurations or significant refactoring of existing business processes to deliver business success.

By building in-house support capabilities for making the right application configuration and business process adjustments, we can help ensure the support staff understand both the business needs and the application's capabilities to address business goals.

We try to maximize aspects of our contractual relationship with cloud service providers, such as choosing the optimal service-level agreement. We have created an internal website that helps application developers evaluate private cloud data center offerings and advantages versus an enterprise-managed service provider or SaaS alternatives. Often, we have found that an on-premises or co-location environment provides a valuable alternative to IaaS cloud hosting offerings. As mentioned earlier, we see that many enterprises are also repatriating certain workloads to on-premises data centers and private clouds to better fit their business needs.

Intel® Tiber™ AI Cloud

In 2022, Intel's Software and Advanced Technology Group (SATG) worked with Intel IT to transform the formerly known Intel® Developer Cloud (DevCloud) to be more scalable and broader in scope. The underlying infrastructure that supports today's Intel® Tiber™ AI Cloud was developed by several Intel IT teams that determined which data centers could optimally host the services, set up the network, and manage load balancing.

The expanded cloud offering, first available in 2023, provides various Intel® toolkits and libraries; AI frameworks and tooling; and AI foundation models. Internal and external developers can use these resources to accelerate AI development using Intel® architecture-optimized software on the latest Intel® Xeon® processors and GPU compute. By the end of 2023, the Intel Tiber AI Cloud was being used by 10,000 unique active customers every day.

Software engineers can access the Intel Tiber AI Cloud for various resources:

- Get hands-on experience with the latest Intel® products and build AI skills.
- Access and evaluate pre-release Intel platforms and associated Intel architecture-optimized software stacks.
- Accelerate AI deployments with the latest machine-learning toolkits.

Intel IT development teams use the Intel Tiber AI Cloud to test and develop solutions using the latest Intel® hardware and optimized software without downloading or configurations. In other words, Intel IT helped build the Intel Tiber AI Cloud, which is now being used to "engineer our way into the future."

Updating Our CaaS Environment

When we introduced the CaaS platform as part of our cloud strategy, we wanted to provide a standard, secure, scalable platform for modernizing business workloads. CaaS delivers a vast array of business benefits, including resiliency, portability, flexibility, and scalability while keeping cost, capability, and controls in mind.

CaaS abstracts complexity away from the user and enables a standardized, secure, and efficient solution for the development and maintenance of container-native applications. Containerized applications are quickly deployable, offer easy scalability, and are inherently portable because they can execute consistently across diverse environments. As part of our application rationalization process, we often refactor applications to CaaS to increase portability. The CaaS environment also enables application developers to develop with a cloud-aware mindset, so applications are more cloud-native.

In the last six years, we have seen remarkable growth in adoption of our CaaS platform. While the growth did not come without challenges, it was not surprising due to the agility CaaS provides. Today, our CaaS platform is one of the prominent platforms for developers at Intel. From simple web applications to complex build pipelines, from data processing platforms to critical databases, we have seen a large spectrum of workloads migrating to our CaaS platform.

Like our DBaaS platform, our CaaS environment is increasingly driven by open-source innovation and cost savings. Open source allows us to cost-efficiently scale our application infrastructure, support new and innovative usage models, adapt to business requirements quickly, and lower the risk of vendor disruptions.

Supporting the Growth of AI and Automation

AI is a driving force behind Intel's ongoing transformation and AI solutions are integral to Intel's business-critical activities. For example, as Intel's manufacturing capacity ramps up for IDM 2.0, engineers are deploying AI across hundreds of manufacturing use cases—including closed-loop control of manufacturing equipment, computer vision-based quality control, and data-driven preventative maintenance. Also, we have built a new generative AI (GenAI) platform that will enable teams to safely use GenAI without requiring them to rebuild common solution components, thereby reducing security risks and improving time to value.

To enhance repeatability and scalability, we have integrated industry-standard, automated quality control principles and tools throughout the AI model lifecycle. Our vision is for AI solutions to understand the quality state of their functions and respond independently to that state without any human intervention.

We are also transforming hosting services and capabilities by using automation and AI technologies to accelerate customer value, expand capabilities, and increase IT efficiency—allowing us to scale AI as fast as the business needs it. For example, AI and GenAI help to enhance hosting services through indexing of enterprise content sources, support ticket generation, and workload placement guidance. We also offer a hosting chatbot that can help end users access information about Kubernetes and containers and obtain hosting environment telemetry. Our work-in-progress automated IaC capabilities will use AI workload guidance to accelerate delivery, further enabling the automation of continuous integration/continuous delivery pipelines.

Next Steps

We know that our multicloud strategy is providing business value today, but we also know that we must continue to evolve that strategy as business needs change. Areas that we intend to explore in the next few years include the following:

- Expanding storage as a service.
- Improving private and public cost efficiency by deploying higher-frequency CPUs with fewer cores to reduce licensing costs.
- Using the Intel® Tiber™ AI Cloud to explore how Intel® technologies like the Intel® Geti™ platform, Intel® Gaudi® AI accelerators, and Intel® Threat Detection Technology (Intel® TDT) can improve on-premises cloud offerings.
- Evolving our AI roadmap to keep pace with rapidly changing and maturing AI technology.

Measuring success is key to knowing if our multicloud strategy is meeting our business goals. In the past, we used surveys and focus groups to engage with Intel business units, and we are communicating with other companies through interviews, workshops, conferences, and an analysis of industry cloud strategy trends. We plan to assess the business value of our multicloud strategy every six months internally and every two years externally.

Conclusion

By focusing our cloud strategy on business and application needs rather than on infrastructure, we can provide optimal business value to all of Intel's lines of business through a consistent user experience and best-fit workload placement. Overall, we have developed a model for an enterprise-wide cloud strategy that applies to both IT and Intel's business units and can maximize the value of the cloud across the enterprise.

While our initial strategy of business-focused cloud decision-making appears to have worked well so far, we constantly monitor new trends that are taking place inside Intel as well as in the industry. We are committed to continually learning and will evolve our multicloud strategy so that it always delivers business agility at the right cost.

Related Content

If you liked this paper, you may also be interested in these related stories:

- Data Center Strategy Leading Intel's Business Transformation
- Cloud Containers as a Service – Wins, Trends and Strategies
- Building a Multi-Cloud-Ready Enterprise Network
- Driving Digital Acceleration with Application Portfolio Management
- Enterprise Technical Debt Strategy and Framework

For more information on Intel IT best practices, visit intel.com/IT

IT@Intel

We connect IT professionals with their IT peers inside Intel. Our IT department solves some of today's most demanding and complex technology issues, and we want to share these lessons directly with our fellow IT professionals in an open peer-to-peer forum.

Our goal is simple: improve efficiency throughout the organization and enhance the business value of IT investments.

Follow us and join the conversation on [X](#) or [LinkedIn](#).

Visit us today at intel.com/IT or contact your local Intel representative if you would like to learn more.

