WHITE PAPER

Data Center
Cloud

Optimal Workload Placement for Public, Hybrid, and Private Clouds

Strategies for evaluating your business needs, technical requirements, and ecosystem offerings before selecting a cloud service

Public and private cloud services provide on-demand access to computing resources for rapid provisioning and release with minimal management.

Table of Contents

Executive Summary

Cloud technologies have improved time to market (TTM), lowered operational and capital expenditures, and provided organizations with an ability to dynamically adjust provisioning to meet changing needs globally. The term "cloud" remains nebulous as cloud implementations include ever-changing technology offerings, making it difficult to discern how to optimize a cloud solution to meet your business needs. Additionally, the needs and goals of each organization and industry differ, making it impossible to adopt a one-size-fits-all cloud strategy—or even the same strategy for each workload within an organization.

Understanding your workload attributes (performance, security, integration, and data volume) is also critical in making a cloud hosting decision. It is important to consider the cumulative impact of these attributes on your workload placement decisions.

This document explains in detail the differences in workloads and how you can guide your organization to the optimal cloud solution given the many options available today. For example, email is well-suited to public clouds, whereas scientific research is typically not. While some of these placement decisions are more obvious, such as readily available commercial software-as-a-service (SaaS) solutions, most require consideration of additional factors.

This document is designed to assist with decision making for existing, traditional workloads that fall in the center of the cloud spectrum, where there is no straightforward answer. Most of these workloads require a comprehensive understanding of the business needs, gathering data to assess your current environment, analyzing options, enhancing selected workloads for the best cloud performance, moving the workloads to the most suitable environment, and continuously monitoring performance to ensure business needs are met.

Intel Data Center Group Authors

Julie Moore, Program Director

William Giard, Principal Engineer

Danielle Austin, Enterprise Cloud Marketing Specialist

Overview

The term cloud has become a catch-all, so it is important to build a common understanding of the terms used in this document. For additional terminology and definitions not covered in the following sections, refer to Appendix A – Additional Terminology and Definitions at the end of this document. The following are recommended for characterizing workloads during the decision process (also see Figure 1):

- **Business considerations.** These include the top business problems that your organization is working to solve, the main use cases you want to enable or enhance—time to market (TTM), agility, and legal and regulatory.
- **Technical considerations.** These include attributes such as performance, security, integration, data volume, and workload elasticity.
- **Ecosystem considerations.** These include factors such as software as a service (SaaS) maturity, cloud service provider (CSP) offerings, or the market accessibility of cloud expertise.
- **Other considerations.** These include a range of issues like existing applications and their cloud-readiness, licensing, and organizational practices, such as disaster recovery.

In general, cloud services provide on-demand access to a shared pool of computing resources designed for rapid provisioning and release with minimal management or provider interaction. As you examine your cloud infrastructure requirements, your first step is to identify which cloud services best fit those needs. You will also assess which deployment model is best for you.

Cloud Implementation Models

There are three cloud implementation models:

- **Private.** Private clouds are virtualized compute, network, and storage resources that support orchestration and automation, and are deployed within a customer's private network. They may reside on-premises, managed internally. They can also be located off-premises, managed by a third party and connected through virtual private networks (VPNs). Managed solutions may be provisioned as single or multi-tenant configurations.
- **Public.** Public clouds are virtualized compute, network, and storage resources that are offered and managed by a third party outside of the customer's private network. Resources are hosted in a multi-tenant configuration in an external data center which may be distributed geographically.
- **Hybrid.** Hybrid clouds provide an approach that combines workload components from both private and public cloud solutions.

Cloud Service Types

There are three primary types of cloud services:

- **Software as a service (SaaS).** SaaS includes network-based services offered through commercially available software that is running on private or public clouds. SaaS examples include applications like customer relationship management (CRM), productivity and collaboration tools, online file storage, and backup services. SaaS options generally offer excellent agility and service quality, as well as simplified internal management with little or no capital expense.
- **Platform as a service (PaaS).** Offered in both public and private clouds, PaaS provides a pre-provisioned environment with an operating system, abstracted middleware, and infrastructure that allows developers to rapidly deploy applications without having to procure or provision servers. PaaS enables rapid application deployment through self-service, on-demand tools, resources, automation, and a hosted platform runtime container in private or public clouds.

Workload Placement Decision Factors

Business Considerations
- Agility/time to market
- Legal/regulatory
- Business asset control
- Global reach
- TCO reduction
- Service level agreements

Technical Considerations
- Performance
- Security
- Back-end integration
- Data volume
- Workload elasticity

Ecosystem Considerations
- Mature SaaS offerings
- Cloud service provider offerings without lock-in
- Cloud expertise accessibility

Other Considerations
- Traditional/legacy applications
- Global data center operations
- Application licensing
- Business continuity and disaster recovery

Figure 1. The decision factors for workload placement can be divided into four main categories: business, technical, ecosystem, and other considerations.

- **Infrastructure as a service (IaaS).** IaaS is virtualized hardware (compute, storage, and network) delivered as code. IaaS gives developers more control over the entire application stack, or when the application requires isolation from other applications within the stack. In addition, an IaaS implementation exposes infrastructure services as APIs and provides a user interface that makes it easy for developers to consume compute, network, and storage.

Characterizing Workloads and Deciding on Placement

Your highest priority in deciding workload placement is understanding your organization's business needs and pain points. What are the issues that keep organizational leaders up at night—growth, data control, cost, IT scale? Use these priorities to guide your decision process.

Business Considerations

While technical considerations are important, organizations may already have firm positions about the business needs that will dictate the cloud strategy. Before going into detail on specific workloads, review the following general rules of thumb. These can often spare you from unnecessary workload analysis (see Table 1).

Further, different industries will prioritize the business considerations differently. For example, in the Academic segment, Intel has consistently found that the three major business considerations for workload placement include:

1. Legal requirements: Are there regulations that require educational institutions to keep data on site?
2. Tolerance: What is an acceptable level for trade-offs (e.g., slightly reduced performance or lack of customization for a reduced data center footprint)?
3. Risk: What risk is the organization willing to accept? Can the solution go offline? Who owns the data in the cloud?

Using the above business considerations, Intel has been able to successfully advise the right public and private/hybrid cloud mix for our K-12 and higher-education customers.

Example: Financial Institution – Sensitive Data Challenge

A Fortune 500 financial institution had an initiative to move applications and data to the public cloud. During the process of developing the strategy, it was discovered that their corporate policy prohibited placement of personally identifiable information (PII) and other sensitive data beyond their internal network/firewall. Although many security standards are supported by public cloud providers, because of the internal policy, the organization kept their data on-premises.

Table 1. Business Considerations For Cloud Implementations

Attribute	Public Cloud Indicators	Private/Hybrid Cloud Indicators
Organization Size and Maturity	• Small or startup organization • Little or no IT infrastructure • Limited in-house IT investment	• Large or mature organization • 500+ physical servers running at 50%+ capacity • Large in-house IT investment
IT Engineering Team	• Little or no in-house IT support • Limited cloud expertise available	• Large in-house IT support • Deep cloud technical bench
Financial Strategy	• OPEX or subscription/payments preferred • No funding for initial data center deployment • Lighter data volumes • Mainstream business processes	• CAPEX and depreciation preferred • Large IT/capital budget • Large data volumes • Many customized business processes
End-User Location	• Global customer base – requires global entry point for applications and operations • End-user latency concerns – customers in remote locations • Funding multiple data centers is not cost effective	• Country restrictions on Internet – private WAN connection to private data center required • End-user locations do not contribute to latency concerns • Large corporation with global but consolidated end-user locations – multiple private data centers are cost effective • Data sovereignty restrictions
Compliance and Control Regulations	• No or low regulations or compliance requirements	• Major regulations or compliance requirements • Data sovereignty restrictions (PII or controlled technology)
Service Level Agreement (SLA) Flexibility	• Flexible SLAs • Risk-accepting of Internet/service provider failures • Contract can be placed to penalize providers for latency/downtime or to ensure redundancy	• Restrictive SLAs or 100% availability required at all times • Risk-adverse to SLA failures – trust private infrastructure over the Internet/service providers
Business Asset Control (Risk) Tolerance	• Organization trusts third parties to manage data • Business policies permit data residing outside the firewall	• Requires absolute control of business data and intellectual property (IP) • Failing to maintain IP and data control may result in the loss of critical business assets

Technical Considerations

After reviewing critical business considerations, select the technology that supports your goals. Some applications are better suited for public clouds, while others are better for private clouds. The following overview provides the foundational elements to understand before you consider more complex workload placement.

In 2016, Intel conducted more than 125 customer and systems integrator (SI) focus groups worldwide across all major industries. This internal research identified the four most important technical characteristics that help determine cloud workload placement. Note that priorities differ based on workload type; however, Security was identified as "high priority" most frequently across all groups. The top four technical characteristics are as follows:

- **Performance.** The primary use cases which apply to workloads with very high performance scores (see Figure 2) include the following:
  - Performance and latency in relation to an end user's location, such as engineering solutions that reside physically near engineering departments.
  - Performance for resource-intensive transactions (compute, memory, and I/O) with guaranteed quality of service and response agreements.

  See "Performance Questions" for additional considerations.
- **Security.** Some applications process and house data—intellectual property (IP), personally identifiable information (PII), and personal health information (PHI)—that could cause harm to the organization if affected by malicious or accidental actions. This rating also incorporates whether or not security solutions are broadly available for a particular workload—security solutions are fairly mature for email workloads, for example.

  See "Security Questions" for additional considerations.
- **Integration.** Connections to other databases, frameworks, applications, workflows, and endpoints present challenges to both traditional and cloud migrations. The complexity and quantity of integrations impact the workload placement because of the increased cost to integrate into multiple clouds. Each integration must be assessed, modified, and refactored to meet the operational level agreement (OLA).

  See "Integration Questions" for additional considerations.
- **Data volume.** There are two major factors: data size and location (where the data is created and managed). Large datasets can be challenging to transfer across distances. For example, network log data (a large local dataset) for analytics would be cost prohibitive to transmit and store externally, and it would significantly extend the time to achieve insights.

  See "Data Volume Questions" for additional considerations.

By scoring the above technical considerations for a specific workload, then adding together the individual scores for performance, data volume, integration, and security, you can calculate a total attribute score. In comparing the multiple attribute scores across workloads, their suitability to the public or private cloud becomes clearer. For example, workloads with significant performance needs, security requirements, multiple backend integrations, and large data volumes are better on private clouds. Workloads requiring minimal performance, integration, or storage requirements tend to be better for public clouds and possibly SaaS solutions (see Figure 2).

This model helps explain why an application such as CRM and its business process workloads can be successfully implemented using SaaS in public clouds, whereas research and development workloads, such as engineering or industrial visualization, remain predominantly private.

The following sections delve into the above considerations and provide key questions you should answer to determine the best cloud implementation for your workloads.

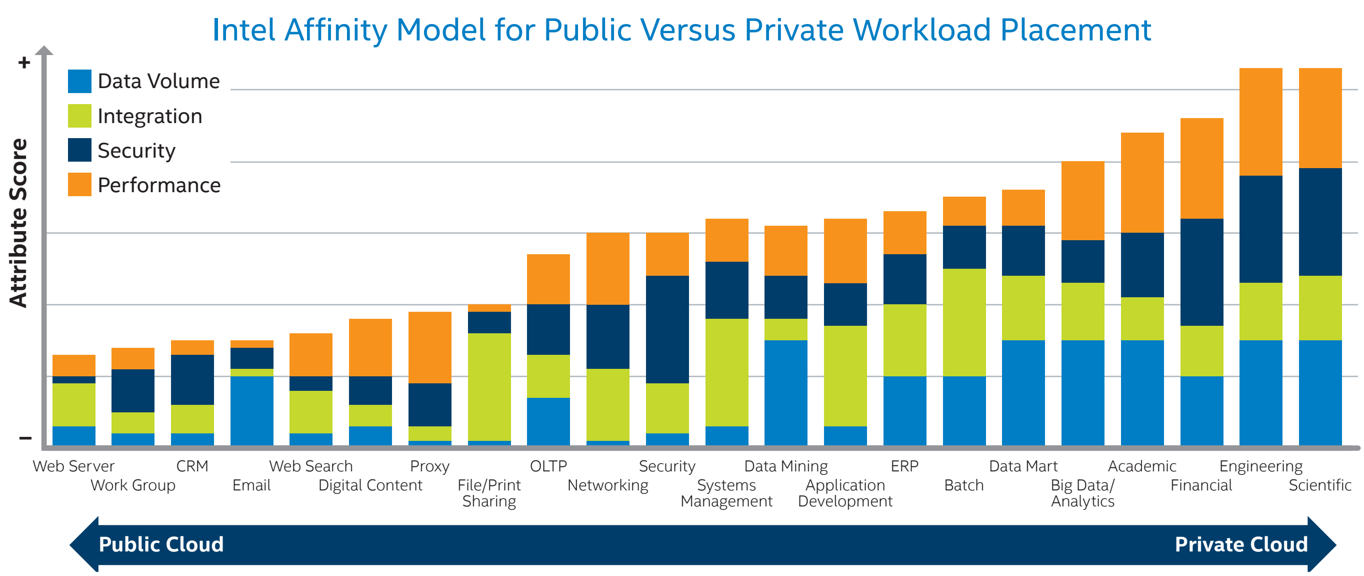

Figure 2. The technical workload characteristics determine whether an application is better suited to public or private clouds. Source: Intel internal analysis, December 2016. Note: Different businesses will have different workload deployments. These deployment differences may influence the attribute score.

PERFORMANCE QUESTIONS

How important is application performance to your competitiveness and financial results? Application performance areas to consider include latency-sensitivity, burst capacity, memory or I/O bottlenecks, and restrictive service level agreements (SLAs).

How sensitive to latency is your application? Some applications and databases must respond within a defined time frame to meet end-user expectations. Private or hybrid clouds may be a good option where organizations control the location of their users and data centers, or the connections between them.

Do any of your workloads suffer memory or I/O speed bottlenecks? Applications that perform specific computations or I/O-intensive tasks may be better suited for private or hybrid cloud environments where performance is more easily tuned.

Do your applications need to meet SLAs that require performance monitoring? When your organization requires a specific level of performance to consistently remain within thresholds, a private cloud or a managed services infrastructure can provide reliable, detailed, and visible performance monitoring and alerts.

Would your workloads benefit from burst capacity to cover irregular demand? For workloads that require extensive computing power at irregular intervals, such as online retail holiday sales or Internet voting, public clouds provide a distinct advantage as long as the application can scale across multiple instances and does not have latency dependencies.

How much elasticity do your workloads require? Elasticity is among the top reasons organizations pursue public clouds because it provides a way to adapt to a workload's changing needs—the degree to which demand grows, shrinks, pauses, is transient, or is unpredictable. Workloads that take advantage of elasticity include prototyping, development, test, batch processing, seasonal bursting, month-end processing, and testing new hardware capabilities. As a core capability of cloud architectures, elasticity can alleviate both capacity and short-term compute constraints. While public cloud is a solid choice for elastic scale/descale needs, there are also private cloud solutions that take advantage of the scale-out architecture to manage spikes in demand and provide a good customer experience.

SECURITY QUESTIONS

Do your business practices and culture allow you to host critical information only on private infrastructure? IP security and policy compliance is clearly a top priority for every CIO, and some businesses categorically require certain information to reside only on internal infrastructure. Enterprises must determine if the target hosting platform supports all their data security requirements and enables them to mitigate risks. Be sure to understand your company's IP handling policies for sensitive information and culture before migrating workloads to any new hosting environment.

Does your organization actively employ best practices for attack surface reduction? Data encryption? Many sophisticated IT organizations have policies and practices to reduce the number of attack points across configuration, network access points, management interfaces, administrative accounts, and shared multi-tenancy. Typically, they also require protection controls for data at rest, runtime, and in transit while certifying the underlying infrastructure from malware and misconfiguration. IT organizations with a deep cultural history of data security practices, may be more comfortable maintaining sensitive workloads on premises, although public clouds are also very secure.

Is your organization required to perform regular security audits? The ability to perform security audits is a critical element in meeting security and regulatory attestation requirements. To prove they are secure, most IT organizations require transparency and tracking in environment management, changes, and transaction execution. If your organization is subject to these expectations, your security auditors may prefer more direct control over, and accessibility to, the on-premises environment. Many public cloud providers have also enhanced audit and compliance support, be sure security audit terms and SLA meet your requirements.

INTEGRATION QUESTIONS

To what extent do your applications depend upon other applications? Legacy systems integration can unexpectedly add weeks or even months to the implementation schedule. Very few enterprise workloads are stand-alone, and many have multiple components. They often require resources from other teams within the enterprise, and they integrate with other applications to support critical organization functions, such as enterprise resource planning (ERP) or product lifecycle management (PLM). They also contain critical data, such as financial, customer, or employee information. It is important to consider enterprise security and authentication systems like a directory services application, single sign-on, as well as identity and access management applications.

Applications that require integration with other systems may be limited by proximity, hardware, or security requirements. Some legacy or third-party applications may not have open APIs, making it challenging or impractical to integrate with new systems. It may be cost-prohibitive to replace, migrate, or modify these applications without addressing the dependencies.

These factors add complexity that is often underestimated. The amount of rework required to integrate services from other systems with external solutions can add significant time. Additionally, migrating an application to a different cloud can increase WAN traffic and costs if the dependent applications remain on a different cloud.

Most organizations find more value in connecting traditional and modern cloud applications through services instead of rewriting an entire application. This enables them to manage both traditional and modern cloud-aware applications on a common infrastructure, delivering new capabilities using cloud techniques alongside mature solutions. Such an approach delivers both agility and the cost-effectiveness of cloud solutions. IT organizations also gain value through automation.

Example: Major Online Retailer – Integration Challenge

A Fortune 500 online retailer manually entered orders for a critical supplier due to the supplier's public cloud application's inability to communicate with the retailer's third-party order entry system. Manually entering orders was not scalable and developing custom integrations was too complex to meet near-term business goals. As a result of numerous integration challenges, the retailer chose not to migrate their order entry application to the public cloud and instead selected a third-party integration platform to enable automated order updates.

DATA VOLUME QUESTIONS

How much data do you currently store and what are your data growth projections? Large data volumes can be expensive to store, difficult to migrate, and may increase transaction time. Organizations can find opportunities to improve internal storage infrastructure, reduce cost, and improve operational efficiencies by taking advantage of new storage innovation and cloud integration. When working with large data stores, the physical location is often very important. Locating the data close to the business applications that use the information is an important design decision for both public and private cloud solutions.

Example: Financial Institution – Data Storage Challenge

A Fortune 500 financial institution's data growth was expected to exceed their current IT budget within two years. They evaluated several options to consolidate storage, including using the public cloud and new hardware platforms. Due to the size and extent of their data, the cost of migration to and from the public cloud presented serious concerns. As a result, the organization invested in a modern, software-defined storage (SDS) platform for their internal data center. The system significantly reduced costs and can accommodate growth into the foreseeable future.

Ecosystem Considerations

Ecosystem considerations include factors such as SaaS, CSP, or the availability of cloud expertise.

Is a mature SaaS offering available for this business application? SaaS offerings are increasingly popular because they help organizations fund business solutions through subscription with the bulk of operational expenses directed toward staffing costs and software licenses. For some workloads, the agility benefits of SaaS exceed costs. As a prerequisite to undertaking SaaS, the organization should clearly understand the extent of business process re-engineering and technical integrations required to meet the business needs. If the organization's existing business processes are close to mainstream, or the technical evaluation team determines few major integrations are required, SaaS may be the right fit for the workload.

Which CSP services do you need? Public CSPs are delivering capabilities to better enable cloud technologies, introducing code constructs and services at a rapid pace. Larger CSPs offer developer frameworks and rich integrated development environment (IDE) tool suites. With each capability there are unique services and offerings. Before you plan to use specific services, be sure to verify the extent to which the service will increase your subscription costs.

Even with simple models of authentication or deeper platform-specific runtimes, removing the code or configuration-specific syntax, logic, or service dependencies becomes difficult. Once data has been moved and application code is rewritten, those workloads are much more difficult to migrate to other CSPs or back to private data centers.

How accessible is cloud expertise to your organization? Cloud technologies are still relatively new compared to traditional models, and many organizations have only begun the multi-year journey of training their development and operational engineers. Developers, for example, need to understand how to develop cloud-optimized applications. For organizations with a deep technical bench, chances are some early adopters have already acquired cloud expertise.

Some industries and geographies have limited cloud expertise or narrow internal knowledge of cloud technologies. With increasing cloud demand, these organizations may prefer public, hybrid cloud, or SaaS models, which allow them to contract expertise from their cloud providers.

Other Considerations

Are your applications designed for traditional computing or cloud computing environments? Placing an application that was not designed to share infrastructure—such as runaway processes, heavy network traffic, and high memory utilization—without isolation or tuning may result in performance issues for other workloads using the same resources.

New application development efforts are recommended to take advantage of cloud technologies, regardless of infrastructure hosting location or cloud architecture.

To realize the full value of investing in cloud, some organizations will need to refactor a portion of their high-value applications to exploit cloud technologies, especially for the major benefits of automated orchestration, auto-scaling, and improved failover and fault resiliency.

Existing applications can also be refactored to increase agility, such as web and application servers, without touching databases and traditional application engines that rely on virtual machine (VM) deployments and connecting them through services. This approach maximizes flexibility and return on investment, while increasing choices as technology evolves for both public and private clouds.

Example: Rushed Public Cloud Migration Challenge

This case study is an example of why understanding the business, technical, and ecosystem considerations is important to workload placement decisions. An IT department at a large retailer was mandated to move everything to the public cloud quickly to reduce the total cost of ownership (TCO) of applications. But within a month of the relaunch, a critical customer-facing application experienced significant issues because the refactoring requirements were significantly underestimated. It resulted in significant troubleshooting and code refactoring hours, as well as testing cycles. The problem was also highly visible, causing senior leadership to lose confidence in the organization's ability to deliver on their commitments. Performance issues were eventually resolved, but at a high cost to the organization and lost sales.

Do you have global end users but lack global data center operations? It may not be cost effective to have data center operations in every market, but regulations, end-user demands, or user experience may require a local presence. In these cases, public clouds offer the ability to extend an entry point, an application, or operations into another region.

Example: Global End-User Challenge

The users of a Fortune 500 online European transportation services company experienced high latency due to a security and compliance project. Presenting applications based in Europe to a U.S.-based computer system resulted in a poor user experience. To solve this problem, the organization used the public cloud to provide virtual desktops for its European customer service agents. This improved the experience for European users without additional requirements or risk to the organization's sensitive data.

Do you have any commercial application license agreements that prevent off-premises hosting? Contractual obligations as part of software supplier license agreements can impact placement. Some licensing models only support deployment on-premises or in specific vendor clouds. Additionally, some licensing models and keys rely on specific infrastructure configurations. It is important to assess software licenses simultaneously with workload placement to understand the deployment implications.

How important are business continuity and disaster recovery (BC/DR) practices to your workloads? BC/DR is among the most common public cloud usages as it enables data to be stored in multiple locations worldwide without the cost of data center facilities. For organizations with multiple data centers, cloud technologies can be used to quickly restore functionality at different sites. For workloads designed to take advantage of cloud-aware architectures, business continuity can be greatly improved.

Conclusion

There is no single perfect cloud model or set of services for every organization, but optimizing will yield the best benefits of each approach. Consider your business needs and the technical characteristics of your workloads—performance, security, integration, and data requirements. Then consider the cloud ecosystem—suitable CSPs aligned to business needs, SaaS maturity for particular workloads, and commercially available solutions that can work for your business.

Small Organizations, Start-Ups, Acquisitions, and those with Limited In-House IT

Smaller companies or new organizations may prefer public clouds or SaaS models, which allow you to rent services so that operational costs only increase with the growth and success of the business. Consider whether you need base compatibilities or differentiating capabilities. If you need only minimal functionality to conduct business, such as accounting, financial, CRM, human resources, marketing, sourcing, facilities, and office productivity, look first to SaaS offerings. Some may lack the full on-premises features; however, the flexibility they offer your organization can scale as your business grows.

For capabilities that differentiate your organization in the marketplace and enable business strategy execution, evaluate public cloud service providers for subscription-based services until your costs or special requirements demand its own infrastructure.

Medium and Large Organizations

For larger organizations, develop a strategy based on new and existing applications. For new applications, introduce a cloud development strategy to write code for multiple platforms that take advantage of cloud technologies. For existing applications, code will need to be written to address resiliency, failure, and latency, and APIs will need to become bandwidth-aware.

If you have flexible time and a small budget, you can use Gartner Research's workload assessment framework: Five Ways to Migrate Applications to the Cloud and manually develop your cloud strategy. We further suggest that organizations augment this framework by first retiring applications with low usage or duplicate business processes.

If you need faster results, many solution providers can also assist you in quickly performing infrastructure discovery and developing workload placement scenario analyses with sensible justifications.

Contact Us

We have evaluated several leading-industry solution providers and their philosophical alignment to the workload placement guidelines presented in this document. Contact your Intel Account Executive or connect with Intel IT Advisors by completing our short form at advisors.intel.com. Intel IT Advisors are experienced account executives and technical experts ready to help with your analysis or a referral to a suitable solution provider to meet your needs, budget, and timeline.

Find the solution that is right for you. Contact your Intel representative or visit **advisors.intel.com**.

Appendix A – Additional Terminology and Definitions

Bare metal: An environment in which a virtual machine (VM) is installed directly on hardware rather than within the host operating system. The term "bare metal" refers to a hard disk, the usual medium on which a computer's operating system is installed.

Cloud service provider (CSP): A company that offers cloud computing services such as IaaS, PaaS or SaaS for purchase through the Internet.

Cloud-aware: Applications that have the ability to leverage cloud technologies, most notably those of orchestration, auto-scaling, and improved failover and fault resiliency.

Latency: The amount of time it takes a packet to travel from source to destination. Together, latency and bandwidth define the speed and capacity of a network.

Service: A self-contained, modular piece of software that is used for exchanging data among systems across a network.

Virtual machine (VM): Software that can run an operating system and applications to imitate a physical computer. Multiple VMs can run on a physical computer, delivering an efficient computing environment.

VM farm: A networking environment, enabled through virtualization software, which provides multiple application and/or infrastructure servers to run on two or more physical servers.

Workload: An application or group of applications that deliver a specific business function. Workload tasks can also be defined as: an application and/or dataset as run, with reproducible measurements; an application under real load; and a capability that can conceivably run in different physical or virtual environments; something running at a given point in time.

Table A1. Workload Categories, Types, and Definitions

Type	Definition
APPLICATION DEVELOPMENT	
Development	This is software that supports the process of computer programming, documenting, testing, and bug fixing in order to create and manage applications and frameworks that ultimately results in a software product. Application development also includes orthogonal tasks embedded into application development workload.
BUSINESS PROCESSING	
Batch	These applications execute a series of jobs in a program on a server without manual intervention (non-interactive), typically high volume. Batch involves traditional legacy mainframe-type processes that execute business process transactions in a batch process. Batch also includes orthogonal tasks embedded into batch workload.
Customer Relationship Management (CRM)	These are applications that are used to manage and analyze customer interactions and data throughout the customer lifecycle, with the goal of improving organization relationships with customers, assisting in customer retention, and driving sales growth. CRM automates the customer-facing business processes within an organization: sales, marketing, and customer service. CRM also includes orthogonal tasks embedded into CRM workloads such as customer analytics and CRM databases.
Enterprise Resource Planning (ERP)	These applications manage and integrate the important parts of its organization. An ERP management information system integrates areas such as planning, purchasing, inventory, sales, marketing, finance, and human resources. ERP includes the following transactional tasks: enterprise-wide line of business (LOB) applications (e.g., Oracle*, PeopleSoft*, and SAP*); other business commerce applications that facilitate business transactions or other task automation over networks; and, departmental transactional applications that run on servers and storage but do not tie directly to other applications. ERP also includes orthogonal tasks embedded into ERP workloads such as ERP data warehouses.
Online Transaction Processing (OLTP)	These are applications that facilitate and manage transaction-oriented information, typically for data entry and retrieval transaction processing. OLTP includes associated databases, but does not include ERP workloads. OLTP also includes orthogonal tasks embedded into OLTP workloads.
COLLABORATION	
Email	Email servers are applications whose sole purpose is to act as a virtual post office. This category includes workloads associated with traditional email applications, whether or not the end user pays to use them. The workload associated with an individual using Google Gmail* would be allocated to the email workload. Email also includes orthogonal tasks embedded into email workload.
Workgroup	This is designed to help people involved in a common task to achieve their goals, by enabling them to more easily collaborate and share information. Software products such as email, calendaring, text chat, wiki, and bookmarking belong to this category, as do social media workloads. Workgroup also includes orthogonal tasks embedded into workgroup workload.

Type	Definition
DECISION SUPPORT SYSTEMS (DSS)	
Big Data/ Analytics	These workloads include technologies and architectures to process and analyze very large volumes of data, for a wide variety of types of structured and unstructured data, while enabling high-velocity capture, discovery, high-performance, and high-end commercial analytics, and simulation to uncover hidden patterns, unknown correlations, and other useful information. Only analytic workloads that fit the big data definition are included in this category. Big data analytics includes embedded workloads associated with application data store, smart data, fast data, and massive data. The big data analytics workload includes both discrete tasks and those often embedded in other workloads such as economic/financial, industrial, scientific research, etc. Big data associated with government R&D, academic/university, and national/industrial security is not considered here.
Data Mining/Data Analysis	These workloads involve infrastructure and tools that are used to access data warehouses for online analytical processing (OLAP), data mining, data visualization, web query tools, and so on. Embedded portions of data analysis/data mining are captured within respective individual workloads (ERP, CRM, etc.). Economics/financial workloads are excluded, and are included in a separate workload. Data analysis/data mining workloads that fit the big data definition are excluded, and are included in the big data analytics workload. Data analysis/data mining also includes orthogonal tasks embedded into the data analysis/data mining workload.
Data Warehouse/ Data Mart	These are central repositories of integrated data from one or more disparate sources, used for reporting and data analysis, and are considered a core component of business intelligence. Data warehousing/data mart includes discrete workloads associated with application data store, smart data, fast data, and massive data. Embedded portions of data warehousing/data mart and big data are captured within respective individual workloads (ERP, CRM, big data analytics, scientific research, national/industrial security, etc.).
Financial/ Economics	These include econometric modeling, portfolio management, stock market and economic forecasting, and financial analysis. The segment includes both trader and computationally intensive non-trader tasks. Economics/financial workloads that fit the technical computing definition are included. Economics/financial also includes orthogonal tasks embedded into economics/financial workload.
IT INFRASTRUCTURE	
File/Print Sharing	This software enables users on a network to share files and printers across different computers. This category also includes cache servers. Embedded portions of file and print are captured within respective individual workloads (ERP, CRM, etc.).
Networking	These workloads include networking applications that support foundational networking processes such as: directory, security/authentication, network data/file transfer, communication, and system data/file transfer. Embedded networking functions are captured within respective individual workloads (ERP CRM, etc.).
Proxy	These applications improve data center performance by storing and serving content from the network edge. Proxy caching includes data caching workloads. Embedded portions of proxy caching are captured within respective individual workloads (ERP, CRM, etc.).
Security	These applications are specifically designed for authentication and identification and typically performing “firewall” services, as well as Unified Threat Management and Deep Packet Inspection workloads. Embedded portions of security are captured within respective individual workloads (ERP, CRM, etc.). Security includes workloads that fit the technical computing definition, but excludes R&D workloads included in the national/ industrial security workload.
Systems Management/ Monitoring	These discrete tasks monitor and account for systems performance, resource planning, and resource allocation. It also includes infrastructure virtualization and monitor and control workloads. Discrete storage management tasks such as backup and recovery, data de-dup, compress/decompress, encrypt/decrypt, thin provisioning, etc. are included in the systems management workload. Embedded portions of systems management are captured within respective individual workloads (ERP, CRM, data warehousing, etc.).
RESEARCH AND DEVELOPMENT	
Academic/ University	This includes applications for scientific research and engineering R&D efforts conducted at public or private institutes of higher education, including systems sold for both research and educational activities, conducted at privately funded and/or non-profit research institutes that have a strong academic mission (i.e., work to extend the bounds of public knowledge). Applications are typically compute or data intensive and often require high-performance graphics. These users are less bound by strict economic constraints than those performing applications in product development environments. This segment includes NSF sites that are located at universities and includes orthogonal tasks embedded into academic/university workloads.

Type	Definition
RESEARCH AND DEVELOPMENT (CONTINUED)	
Engineering/ Industrial	These workloads obtain new knowledge that is applicable to a company's business needs, which eventually will result in new or improved products, processes, systems, or services that can increase the company's sales and profits. Examples of these workloads include: chemical engineering; computer aided engineering and manufacturing; electronic design and analysis; and industrial automation.
Scientific Research	This includes a variety of workloads that center on investigative research associated with the relationships among natural phenomenon, including the biological sciences, geosciences and geo-engineering, weather forecasting, and climate modeling. Excluded categories are technical computing workloads included in government R&D and academic/university. Workloads that fit the technical computing definition are included, as are orthogonal tasks embedded into scientific research workloads.
WEB INFRASTRUCTURE	
Digital Content Creation and Delivery (DCC&D)	These workloads are related to applications for 2D/3D animation, film and video editing and production, and multimedia authoring for both CD and web pages that utilize sophisticated graphics content. DCC&D includes workloads associated with the streaming/delivery of video and audio multimedia applications, including an Internet component. This category also includes servers used for image rendering, content management, and distribution of finished products for areas such as film, TV, commercial animation, advertising, product styling, and industrial design, as well as servers used for large-scale massively multiplayer online games (MMOG). DCC&D also includes orthogonal tasks embedded into DCC&D workloads. Embedded portions of DCC&D are captured within respective individual workloads (CRM, workgroup, etc.).
Web Search	This workload involves discovering information on the Internet and generating natural (organic) and sponsored (paid) search results. Graph search is also included here as well as future search innovations. Search also includes orthogonal tasks embedded into search workloads and includes both discrete tasks and orthogonal search tasks often embedded in other workloads.
Web Server	These applications support HTTP (Hypertext Transfer Protocol) to serve files that form web pages to users, in response to their requests, which are forwarded by their computers' HTTP clients. Web serving excludes search, but includes front-end web. This category also includes orthogonal tasks embedded into the workload. Embedded portions of web serving are captured within respective individual workloads (ERP, CRM, etc.).

Learn More

- Listen to the **webinar replay** of Critical Considerations for Implementing a Hybrid Cloud
- Read about **Intel® Xeon® Scalable processors**
- Read how **Intel® Select Solutions** provides a faster, easier choice to deploy workload-optimized configurations verified for balanced performance across compute, storage and network in the data-fueled, hybrid cloud era.

 1117/JMOO/KC/PDF 335054-002