

Hybrid Cloud 101

What is Hybrid Cloud?

A Crash Course on Combining Private with Public Cloud Infrastructure

What is hybrid cloud? Is the latest buzz about it all hype, or does it offer real-world opportunities to take the next step in cloud computing? Hybrid clouds combine two cloud delivery models—typically private and public—to help businesses better manage cloud service delivery. This brief provides you with a crash course on the benefits of a hybrid cloud model and how Intel can help.

The Path to Hybrid Cloud

Cloud technology has matured, opening up new possibilities for more elastic private, public, and hybrid models. Organizations now see [hybrid cloud](#) models as an enhanced way to get value from cloud computing by leveraging the sometimes-complementary benefits offered by private and public clouds—agility, cost efficiencies, and high availability of services. In particular, hybrid clouds offer businesses flexibility and choice that can help IT balance capital and operational expenses, make optimal use of in-house resources, and improve responsiveness to changing business requirements.

How each company manages this flexibility depends on their own business needs and IT requirements. Hybrid clouds can be used to:

- Implement disaster recovery strategies.
- Dynamically move workloads between clouds based on business needs.
- Deliver business processes as complete applications through software as a service (SaaS), such as customer relationship management or human resource applications.
- Help manage unpredictable peaks in demand through cloudbursting with multiple vendors, such as consumer-facing web services that may respond to seasonal peaks and valleys.
- Make services available quickly for a specific period of time or to capitalize on new business opportunities.

Intel: Making the Cloud Work for You

Intel wants to help you simplify delivery of your cloud services so your business can realize the full benefits of cloud computing. As a first step, we recommend that you implement a private cloud. Starting with a highly virtualized foundation, you can build a fully automated and orchestrated self-service cloud environment that enables you to offer infrastructure as a service (SaaS) and that scales as demand grows. This provides your organization with a certain level of cloud maturity and enables you to expand your cloud deployment to models that are more flexible.

For more information on developing private cloud services, read [Planning Guide: Private Cloud Infrastructure as a Service](#).

Why Hybrid Cloud Matters

Many organizations are having considerable success delivering private cloud services and may have dabbled in public cloud services as well. As the broader organization recognizes the advantages of cloud computing and gets on board, greater demand is put on existing services. A hybrid cloud enables IT to gain additional flexibility between private and public cloud resources via a single automated and orchestrated operating environment.

A hybrid cloud also enables IT to optimize for disparate needs and diverse workload requirements. Because these requirements will be different for each organization, hybrid cloud is not a “one-size-fits-all” solution. For example, quick-scaling, fast-to-market systems may have an affinity with public cloud. Legacy systems may have an affinity with either private cloud or public cloud SaaS for standardized functions. Core business systems are typically run on an in-house private cloud.

A hybrid cloud also enables an organization to balance its need to invest in on-premises cloud technologies (capital expenditures) with utilization of off-premises public cloud services (operating expenditures) using a “build the base, rent the spike” deployment model to achieve an optimized total cost of ownership (TCO).

Check out how Intel IT is approaching hybrid cloud for the company in the “[Create Advanced Cloud Computing Platforms](#)” video and [Hybrid Cloud—Key Design Strategies](#) white paper.

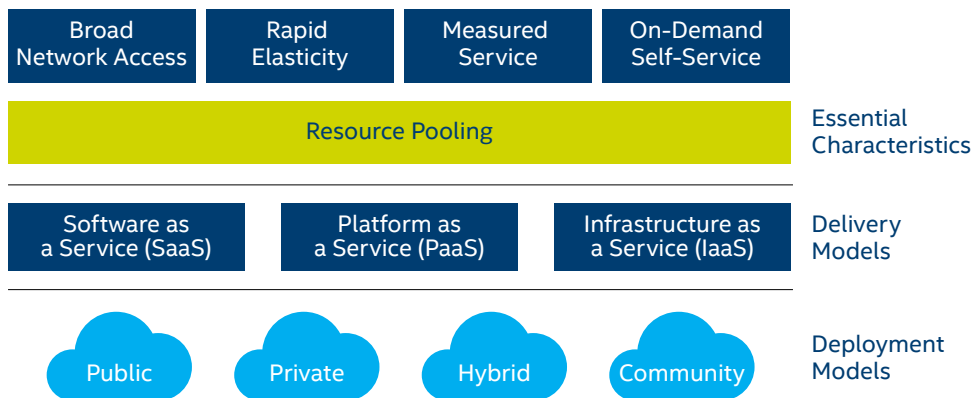
Hybrid Cloud Capabilities

Hybrid clouds combine two or more cloud delivery models—typically private and public—allowing both to remain as unique entities but binding the multiple models together via technology to enable data and application portability.

The National Institute of Standards and Technology (NIST) has described the essential characteristics of the cloud as broad network access, rapid elasticity, measured service, on-demand self-service, and resource pooling. These characteristics operate across three major cloud service layers: SaaS, platform as a service (PaaS), and infrastructure as a service (IaaS).

Service Layers: IaaS and PaaS

Learn more about these cloud services by reading the [planning guide for building private cloud infrastructure as a service \(IaaS\)](#) and the [platform as a service \(PaaS\) white paper](#). Plus, find out how Intel IT uses PaaS to extend the company’s private cloud.



Based on the U.S. National Institute of Standards and Technology (NIST) definition of cloud computing. (Source: *Cloud Computing Synopsis and Recommendations*. NIST Special Publication 800-146 (May 2012). <http://csrc.nist.gov/publications/nistpubs/800-146/sp800-146.pdf>)

Hybrid cloud technology is still maturing, and the combination of two different cloud environments presents IT managers with a number of technical challenges. In an ideal state, however, the design goals for hybrid clouds include these capabilities in addition to the NIST characteristics.

- **Integration of infrastructure and the application environment.** In a hybrid cloud, the ability to spin up virtual machines (VMs) for IaaS or combinations of IaaS and PaaS is ideally the same in both private and public cloud environments.
- **Portability of applications.** Using a cloud-aware development approach improves portability by building capabilities into applications so they work the same across cloud environments.
- **Interconnectivity.** Communication and interaction between two coexisting environments facilitate the easy movement of VMs, data, and applications between individual clouds.
- **Monitoring and management across cloud environments.** While monitoring is important for any cloud environment, visibility into system health across clouds is critical.

Technology Challenges

These additional characteristics of hybrid clouds present unique technology issues that must be addressed as you plan your hybrid cloud, including:

- **Integrated systems architecture across the environment.** Before a potential workload can become a hybrid deployment, you must have at least a rough architecture of where components, functions, and data will reside.
- **Infrastructure and application portability.** Infrastructure must be able to support the environment you maintain in your private cloud so that VMs can move back and forth and applications can work across a dynamic environment.
- **Security across cloud environments.** You must be able to maintain your security, compliance, and privacy requirements in the public cloud environment.

Visibility across cloud environments. Monitoring the environment is essential for cloud management. At the point where IT is offering externally hosted services, you must have a way to measure overall service availability to monitor third-party service level agreements.

For help planning your hybrid cloud, see [Hybrid Cloud Checklist: Operationalizing Your Hybrid Cloud](#).

Open, Extensible Cloud Ecosystems

The vision of open, extensible cloud ecosystems ties an organization's cloud services (public and private) together into an ecosystem with portability and interconnectivity. Portability and interconnectivity are issues at every layer of a hybrid cloud. The key to addressing these issues is open standards.

The [Open Data Center Alliance](#) has defined requirements for [usage models](#) in each service layer for interoperability and interconnectivity, including:

- **IaaS portability.** The ability to move physical or VM instances or images (complete with network connectivity and storage) between environments over short or long distances, with capabilities such as manageability, live and at-rest migration, performance, and distance.
- **PaaS interconnection and application portability.** The ability to move applications (and related logical data structures) between different PaaS environments—development and runtime—with cloud-aware applications that maintain attributes such as feature sets, configurability, and orchestration.
- **SaaS interconnection and portability.** The ability to connect or transfer business process functionality and information via SaaS and to create mash-ups from multiple SaaS and non-SaaS applications via compatible interfaces that exchange data smoothly.

Case Study: Intel IT and Hybrid Cloud

As part of a multiyear cloud strategy, Intel IT has moved toward developing a highly available, dynamic hybrid cloud environment, with three key design goals in mind:

- Design applications and the hosting environment for automated self-healing.
- Design the hybrid cloud to meet unpredictable demand automatically.
- Design cloud-aware applications that accommodate infrastructure outages and that can be concurrently active at multiple locations.

Intel IT has also developed a theoretical financial model showing that a hybrid approach can potentially result in growing savings over a purely outsourced hosting solution. For more information, read [Cloud Computing Cost: Saving with a Hybrid Model](#).

Cloud Management Platforms

A cloud management platform (CMP) is the integrated software that delivers service quality, security, and availability for workloads running in cloud environments. CMP offerings vary widely in terms of platform maturity, architecture complexity, and capabilities. Your choice of platforms can simplify extending manageability, automation, and orchestration into public clouds.

At minimum, a CMP should provide direct user access to the system, self-service capabilities and interfaces, a workflow engine, automated provisioning, and metering and chargeback functionality.

Hybrid clouds utilize more advanced capabilities, such as:

- Performance and capacity management
- Interoperability between private and public IaaS offerings
- Connectivity to and management of external clouds
- Application life-cycle support
- Back-end service catalogs
- Integration with external enterprise management system

Learn more about leading cloud management platforms in [Planning Guide: Private Cloud Infrastructure as a Service](#).

Cloud-Aware Application Design

Cloud-aware application development can take full advantage of underlying cloud infrastructure for improved scalability, performance, and resiliency. If a hybrid cloud is in your future, your applications can be designed now to include capabilities that minimize potential portability issues for the time when you do combine cloud environments, including:

- **Treat everything as a service.** Application capabilities should be partitioned into granular components that can be implemented, tested, and scaled separately.
- **Use representational state transfer (RESTful) APIs.** RESTful APIs enable easy reuse and scaling of application capabilities and shield applications from underlying technology implementations.
- **Separate compute and persistence.** Nothing is stored locally on the compute instance that is running the cloud application, providing deployment and scaling flexibility across environments.
- **Design for failure.** Although the goal is zero failure, in reality, components fail, services become unavailable, and latencies increase. Designing applications to gracefully survive failure enhances the user experience.

- **Architect for resilience.** An architecture designed with a focus on the mean time to recovery (MTTR) accepts imperfection and enables rapid identification and resolution of problems when they occur.
- **Operationalize everything.** All services should be easy to maintain and troubleshoot. Instrumenting, logging, and analyzing application behavior will lead to operational improvements.
- **Implement security at every layer.** A perimeter security approach is not sufficient in a public cloud. A more comprehensive approach is needed, such as encrypted transport into the cloud, secure coding and access control inside applications, and encryption at rest. The security of every API and all data should be tested and analyzed.

For more detail, read [Developing a Highly Available, Dynamic Hybrid Cloud Environment](#).

Security Across Cloud Environments

Combined cloud environments bring new security challenges. Even though private cloud environments are dedicated to a single organization and provide more control than public clouds, data and applications still extend beyond the perimeter of the data center. In addition:

- Virtualization aggregates the security risks of various application components and services into a single physical server platform.
- Shared technology such as CPU caches, graphics processing units (GPUs), disk partitions, memory, and other components were never designed for strong compartmentalization. Compromise of the hypervisor can, in turn, potentially compromise shared physical resources.
- Stealthy attacks on data center infrastructure are difficult to detect with traditional antivirus products. Cybercriminals use rootkit attacks to infect system components such as hypervisors, BIOSs, and operating systems and can hide malware that operates in the background and spreads throughout a cloud environment. Sophisticated platform attacks designed for stealth and control put sensitive data and intellectual property at risk.
- The proliferation of different types of client devices accessing cloud resources provides hackers with many potential access points and targets.
- Edge systems that interact inside and outside the organization, such as web servers, portal servers, e-mail servers, bridges, and routers, represent a growing attack target.

In public clouds, the boundaries between the data center and cloud providers are blurred, creating third-party dependencies for data protection. This can be particularly challenging for businesses trying to meet compliance demands. To comply, organizations need to be able to monitor and attest that security policies are being set and enforced across cloud environments.

The solution is to integrate security into every layer of the cloud. One way to enable the portability and scalability of security across cloud environments is to assign security policies for infrastructure and applications to specific VMs, based on their function. These policies are automatically assigned when that VM is provisioned. Another way is to build security into the hardware of servers and clients to protect infrastructure.

Learn more about Intel's hardware-based security technologies in [Real-World Guide: Intel Security Technology for the Cloud](#). For additional information about security in public clouds, view the "[Powered by Intel® Cloud Technology Program](#)" video.

How Intel Can Help

Intel makes the technologies that serve as the foundation for both private and public clouds. We're committed to helping simplify your cloud services delivery by:

- Delivering open-standards, optimized technology built to scale in the most demanding dynamic environments
- Providing resources to help you advance your cloud projects faster
- Addressing future challenges with ongoing innovations for cloud computing

Intel Resources to Learn More

The Intel® IT Center provides straightforward information that addresses each of the ways Intel can help IT pros implement strategic projects like hybrid cloud computing. For planning guides, peer research, real-world customer references, solution spotlights, and live events about hybrid cloud computing, visit intel.com/cloud.

Intel® Cloud Technology

Intel® Cloud Technology powers and protects the infrastructure stack from the ground up in both private and public cloud environments. A broad range of solutions developed by Intel® ecosystem partners are optimized for Intel® processors and take full advantage of our cloud technologies.

Powerful Intel® Xeon® processor E5 family-based servers come with built-in technologies that boost performance and provide advanced security capabilities.

- **Intel® Turbo Boost Technology** – Delivers more computing power when you need it and adapts to spikes in performance
- **Intel® Advanced Vector Extensions (Intel® AVX)** – Provides performance boost for floating-point-performance-intensive applications
- **Intel® Advanced Encryption Standard New Instructions (Intel® AES-NI)** – Increases encryption speed without associated performance penalties
- **Intel® Platform Protection Technology with Trusted Execution Technology (Intel® TXT)** – Provides enhanced security through hardware-based resistance to malicious software attacks that occur before the virtual machine boots, and provides the ability to verify hardware host integrity for compliance purposes

Find out more about the latest [Intel Xeon processor E5 family](#) and [Intel Cloud Technology](#).



Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Performance varies depending on system configuration. No computer system can be absolutely secure. Check with your system manufacturer or retailer or learn more at intel.com.

Any software source code reprinted in this document is furnished under a software license and may only be used or copied in accordance with the terms of that license.

INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF INTEL PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT.

A "MISSION CRITICAL APPLICATION" IS ANY APPLICATION IN WHICH FAILURE OF THE INTEL PRODUCT COULD RESULT, DIRECTLY OR INDIRECTLY, IN PERSONAL INJURY OR DEATH. SHOULD YOU PURCHASE OR USE INTEL'S PRODUCTS FOR ANY SUCH MISSION CRITICAL APPLICATION, YOU SHALL INDEMNIFY AND HOLD INTEL AND ITS SUBSIDIARIES, SUBCONTRACTORS AND AFFILIATES, AND THE DIRECTORS, OFFICERS, AND EMPLOYEES OF EACH, HARMLESS AGAINST ALL CLAIMS COSTS, DAMAGES, AND EXPENSES AND REASONABLE ATTORNEYS' FEES ARISING OUT OF, DIRECTLY OR INDIRECTLY, ANY CLAIM OF PRODUCT LIABILITY, PERSONAL INJURY, OR DEATH ARISING IN ANY WAY OUT OF SUCH MISSION CRITICAL APPLICATION, WHETHER OR NOT INTEL OR ITS SUBCONTRACTOR WAS NEGLIGENT IN THE DESIGN, MANUFACTURE, OR WARNING OF THE INTEL PRODUCT OR ANY OF ITS PARTS.

Intel may make changes to specifications and product descriptions at any time, without notice. Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined". Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them. The information here is subject to change without notice. Do not finalize a design with this information.

The products described in this document may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

Contact your local Intel sales office or your distributor to obtain the latest specifications and before placing your product order.

Copyright © 2015 Intel Corporation. All rights reserved. Intel, the Intel logo, and Intel Xeon are trademarks of Intel Corporation in the U.S. and/or other countries.

*Other names and brands may be claimed as the property of others

1215/RPC/OC/XX/PDF

330885-002US