



CRAFT AN EFFECTIVE HYBRID CLOUD STRATEGY

Tips for building a successful hybrid cloud deployment plan

TABLE OF CONTENTS

Introduction.....	3
CHAPTER 1	Determine where to put your workloads
CHAPTER 2	Deploy an SDI to support your environment
CHAPTER 3	Choose Intel to deliver high performance
Conclusion	14



INTRODUCTION

Cloud technology has become ubiquitous as a means to respond to disruptive change. To best take advantage of the efficiencies created by this technology, more enterprises are adopting a cloud infrastructure based on industry-standard platforms that allow them to future-proof for changing and transforming IT and evolving software stacks.

Specifically, they're turning to a hybrid cloud environment—which combines private cloud and third-party public cloud services—as the logical infrastructure for businesses and a key enabler for innovation, efficiency, and scalability.

Eighty-three percent of IT business leaders in a recent survey said they currently use

or plan to use a hybrid cloud.¹ And Gartner predicts that by 2020, 90% of IT enterprises will adopt hybrid cloud technology.²

But like any other IT environment, if deployment isn't first thoroughly considered and mapped out, the results could be suboptimal for your needs.

Fortunately, there are some best practices those who came before have discovered. Drawing upon those lessons, this eBook looks at how enterprise IT shops like yours can help ensure their hybrid cloud deployment is successful.

Chapter 1:

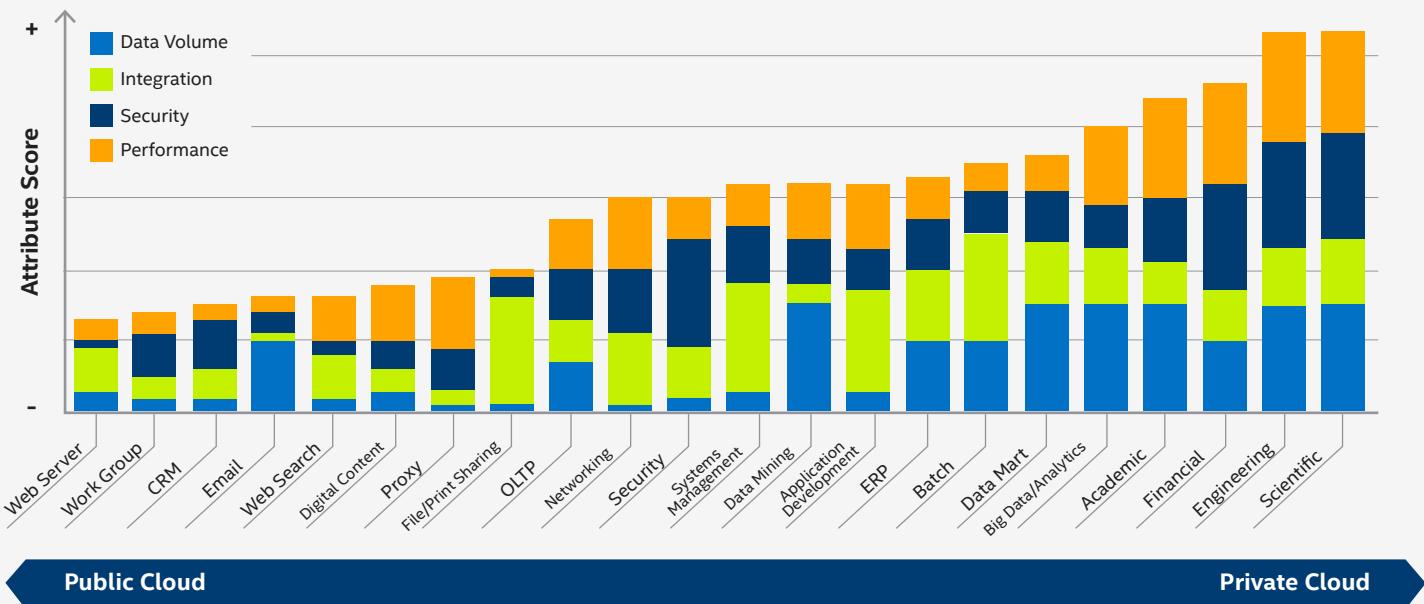
DETERMINE WHERE TO PUT YOUR WORKLOADS

So, you've decided to build a hybrid cloud environment. Now comes the hard part: How do you build one that will allow you to maximize its benefits and grow with you into the future? You'll need to understand your applications and the capabilities of your on-premises infrastructure and your clouds, both public and private, so you can determine where to place various workloads.

It's important to lead with your organization's specific business requirements as you determine workload placement. Consider the primary business needs and use cases your organization is working to address. Workloads should be positioned for the optimal architecture and distributed across public and private clouds.

In 2016, Intel conducted over 125 customer and systems integrator focus groups worldwide to determine optimal cloud workload placement. Intel's findings are key to helping their customers decide where they should place their workloads in the hybrid cloud for optimal total cost of ownership, workload performance and security. The Intel Affinity Model for Workload Placement (below) visually summarizes common enterprise workload's affinity to public or private cloud.

Intel® Affinity Model for Workload Placement



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.

As the model depicts, workloads that have substantial performance needs, security requirements, multiple backend integrations, and large data volumes are often well suited on-premises. Mission-critical workloads are also often ideal for on-premises because moving that data to the public cloud can sometimes delay time-to-insights or expose sensitive information, for example, research and development workloads.

Workloads with minimal performance, integration, or storage requirements tend to be better for public clouds and possibly software-as-a-service (SaaS) solutions. Customer relationship management systems, for example, are ideal for SaaS in public clouds. Public clouds also provide benefits for scaling compute intensive workloads, maintaining business continuity,

1



DATA SIZE AND LOCATION

The larger the data, the more problematic it is to move, possibly increasing expense and transaction time. Locating the data close to the business applications that use the information is an important design decision for both public and private cloud solutions. For example, the enterprise resource management system is likely an ideal application for the private cloud. A large local dataset such as network log data for analytics can be cost-prohibitive to transmit and store externally.

2



SECURITY

Organizations have often chosen to keep applications that involve confidential information on-premises because of perceived stronger security. But as services have matured, public clouds have proven to be very secure, as well. Understanding and asking about the underlying technology behind the CSPs' security is essential.

IT must determine if the target hosting platform supports all data security requirements and enables them to mitigate risks. To do so, get a firm understanding of your company's IP handling policies for sensitive information and culture before migrating workloads to any new hosting environment.

Many organizations want to extend security policies and tools from on-premises clouds to public clouds. Doing so requires an IT administrator to extend identity, IPAM (IP Address Management), and DNS solutions to the public instances. IT must be prepared to manage this extension as more workloads move to public and strategies become hybrid.

Most IT organizations require transparency in environment management, changes, and transaction execution as proof of security. If your organization is subject to these expectations, your security auditors may prefer the on-premises environment because it allows for more direct control over and accessibility to the information.

and providing a development environment to explore and test emerging workloads such as artificial intelligence.

Intel scored each workload based on four technical characteristics that can help guide an organization's hybrid strategy development. By scoring these technical considerations for each workload type, their suitability to the public or private cloud becomes clearer.

3

INTEGRATION

Very few enterprise workloads are standalone; they connect to other databases, frameworks, applications, workflows, and endpoints. They may also contain critical data. Applications that require integration may be limited by proximity, hardware, or security requirements. Some legacy or third-party applications may not have open APIs, making it a challenge (or impractical) to integrate with new systems.

Most organizations find more value in connecting traditional and modern cloud applications through services instead of rewriting entire applications. This enables them to manage traditional and modern cloud-aware applications on a common infrastructure. Such an approach delivers automation, agility, and the cost-effectiveness of cloud solutions.

4

PERFORMANCE

Some applications and databases must respond within a defined timeframe to meet end-user expectations. Applications that perform specific computations or I/O-intensive tasks may be better suited for private clouds where performance is more easily tuned.

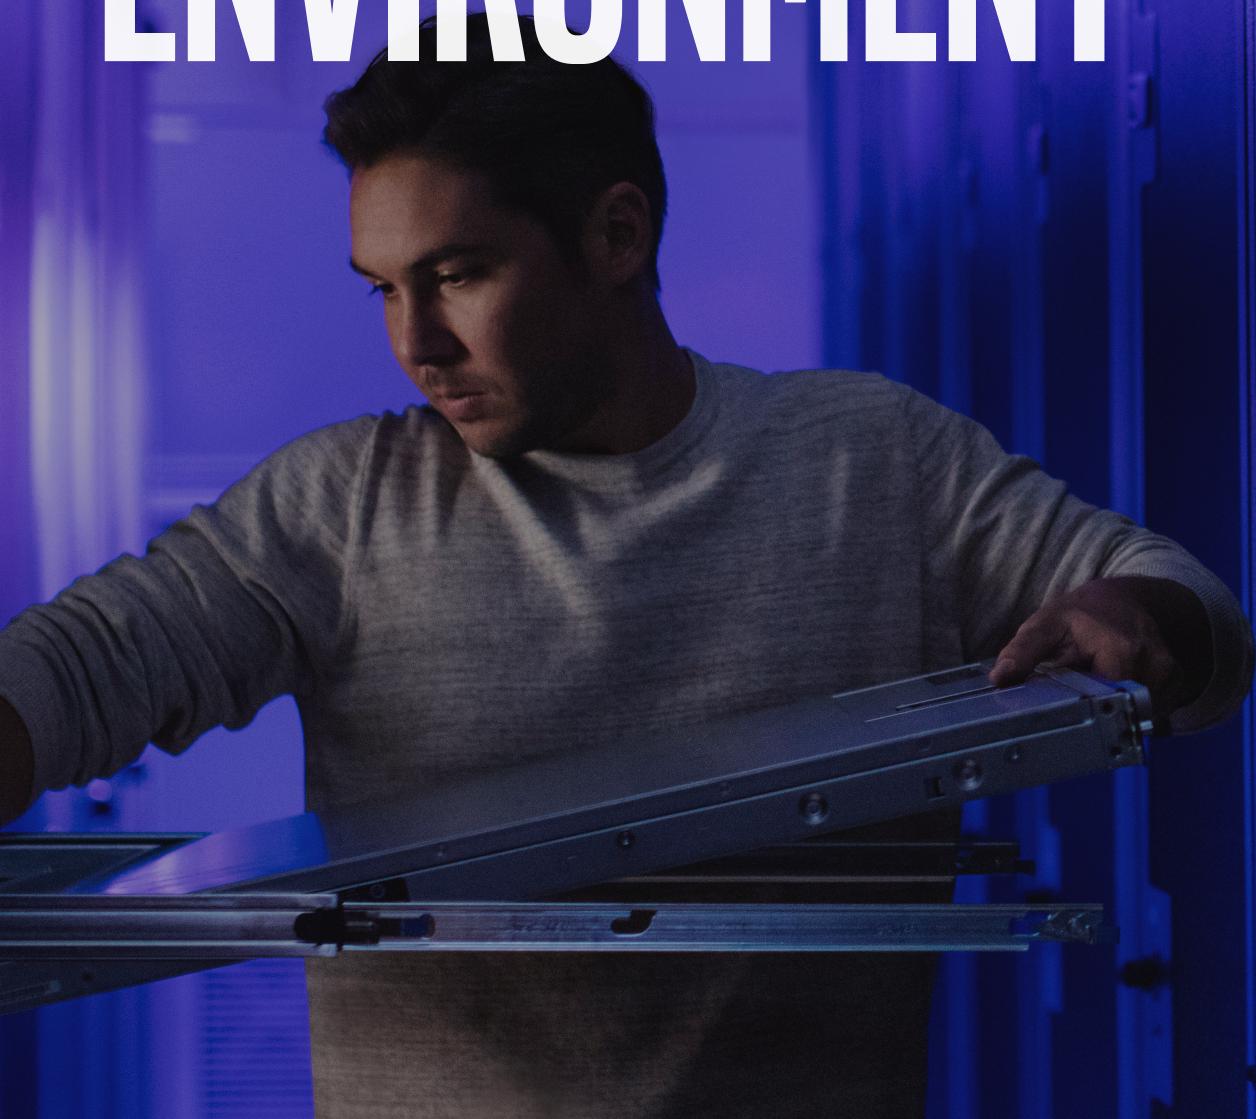
Workloads that require extensive computing power at irregular intervals, such as online retail holiday sales, get a distinct advantage from a public cloud because of the scalability, as long as the application can scale across multiple instances and does not have large latency dependencies.

Application performance areas to consider include latency-sensitivity, burst capacity, memory or I/O bottlenecks, and restrictive SLAs. Be sure to understand your workload requires before speaking to the CSP to be assured your requirements will be met.

To learn more about how Intel developed its evaluation technique for workload placement, reference the Optimal Workload Placement for [Public, Hybrid, and Private Clouds](#) white paper.

Chapter 2:

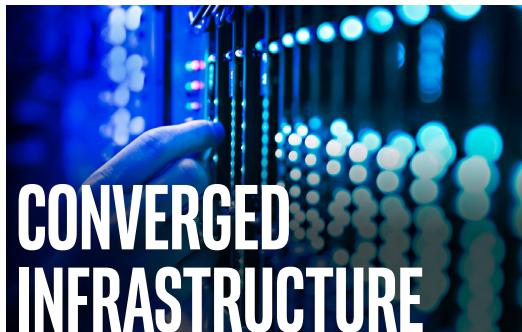
DEPLOY AN SDI TO SUPPORT YOUR ENVIRONMENT



Most enterprises are moving away from traditional IT models as their data centers mature. Many are instead choosing to backstop their networks with software-defined infrastructure (SDI). With SDI, servers based on a standard platform and running in VMs or containers are used for network, storage, and computing, which are presented as a pool of resources. That allows each application to scale across resources to meet workload needs.

SDI provides a comprehensive view of network infrastructure, and the pooled resources allow organizations to reduce the number of standard server models needed in the data center, reducing CapEx and OpEx. Custom APIs allow existing storage pools to interface with the SDI.

There are two preferred types of SDI setups in the market today:



CONVERGED INFRASTRUCTURE

Converged infrastructure (CI) has more integration complexity but allows for more easily managed individual hardware components for optimized performance and scale. Compute, storage, and networking hardware is tightly networked together with a management layer on top. Individual components are integrated by the end user.



HYPER-CONVERGED INFRASTRUCTURE

Hyper-converged infrastructure (HCI) abstracts compute, storage, and networking resources from the hardware and delivers them via hypervisor. HCI delivers performance, efficiencies, and unified management with a single pane of glass. Components are pre-integrated and optimized in the appliance by the vendor.

Either of these options, depending on your use case and workload demands, will ease the path to hybrid cloud by creating the efficiencies and control needed with cloud computing.

Chapter 3:

CHOOSE INTEL TO DELIVER HIGH PERFORMANCE



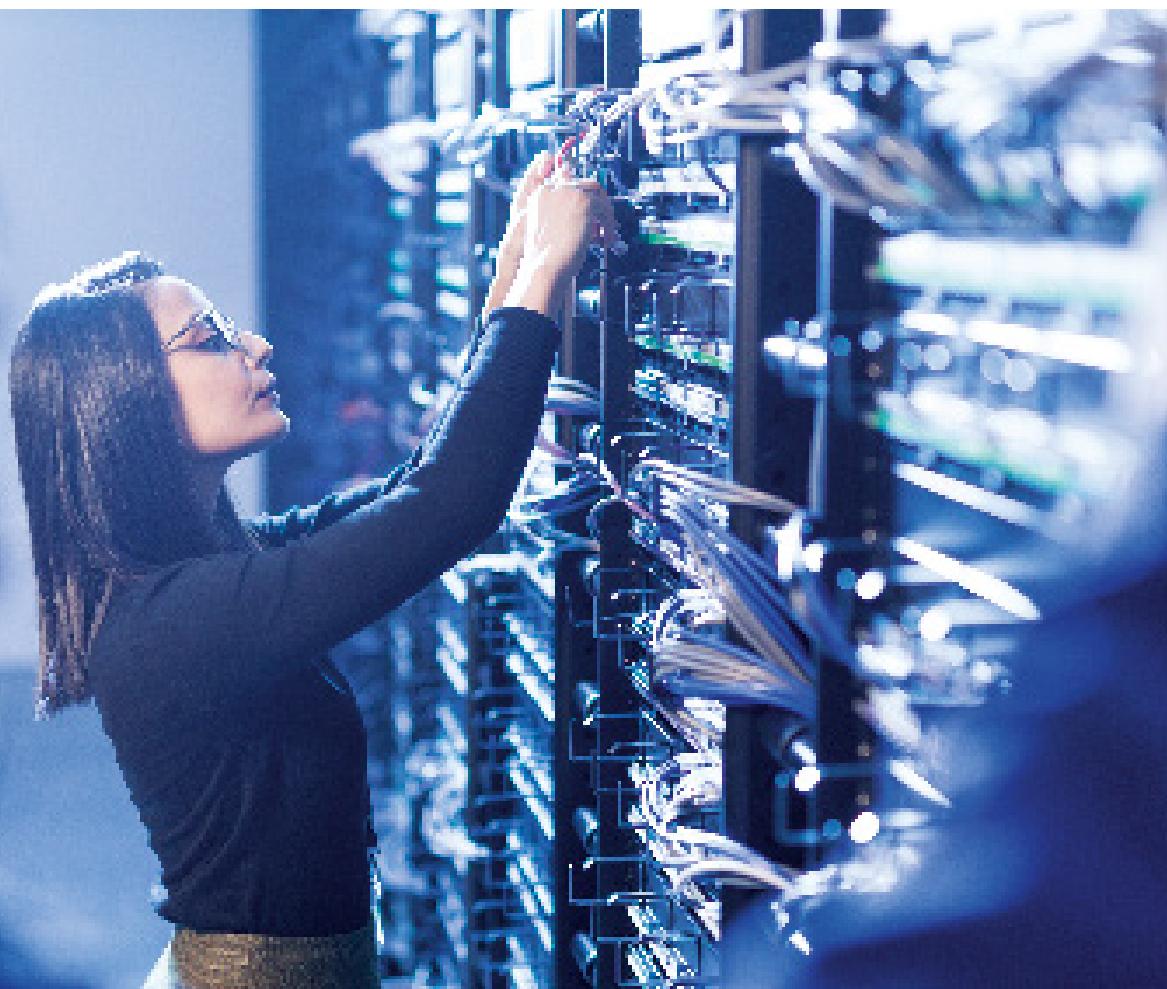
Transitioning to a hybrid cloud environment isn't a simple process. You'll want to identify trusted partners to help you build, which will ease the process and reduce time to value.

And no matter what hybrid cloud solution your organization deploys, Intel technologies deliver performance and efficiency across all clouds adopted by your organization.

The Intel® Xeon® Scalable processor family enables you to replace multiple old systems with fewer new ones, reducing data center footprint and allowing more of your limited IT budget to be allocated to growing your business.

New features such as Intel® Advanced Vector Extension 512 (Intel® AVX-512) improve performance and throughput for advanced analytics, high performance computing (HPC) applications, and data compression. Integrated Intel® QuickAssist Technology (Intel® QAT) speeds up data compression and cryptography, while the high-speed integrated Intel® Ethernet (up to 4x10GbE) reduces total system cost, lowers power consumption, and improves transfer latency of large storage blocks and virtual machine migration.

To accelerate deployment of workload-optimized data center infrastructure, look for an Intel® Select Solution from our valued partners. Intel® Select Solutions are a fast path to data center modernization with workload-optimized configurations. They help ensure proven performance and the ability to scale by running on the Intel® Xeon® Scalable platform.



Key benefits to investing in Intel® Select Solutions include:

- **Simplified evaluation.** IT spends a lot of time and money sorting through endless solution options for workload integration and the transition to software-defined infrastructure. Intel® Select Solutions are tightly specified in terms of hardware and software components to eliminate guesswork and speed decision-making.
- **Fast and easy deployment.** With pre-defined settings and rigorous system-wide tuning, Intel® Select Solutions aim to increase efficiency in IT's testing process, speed time-to-service delivery, and increase confidence in solution performance.
- **Workload-optimized performance.** Intel® Select Solution configurations are designed by Intel and our partners to deliver to a performance threshold for the workload.

SIMPLIFIED EVALUATION

ELIMINATES GUESSWORK
THROUGH TIGHTLY-SPECIFIED
HW AND SW COMPONENTS

FAST & EASY TO DEPLOY

SMOOTH DEPLOYMENT WITH
PRE-DEFINED SETTINGS
AND SYSTEM-WIDE TUNING

WORKLOAD OPTIMIZED

DESIGNED AND BENCHMARKED
FOR SPECIFIC WORKLOADS TO
DELIVER OPTIMIZED PERFORMANCE



CONCLUSION

In today's enterprise environment, having the IT flexibility, agility, and performance to respond instantly to changes is imperative. Increasingly, that means tapping into cloud computing. An intelligent cloud strategy will apply computing, advanced storage technologies, replication, and security to optimize workloads across the enterprise and cloud to maximize the potential of this crucial strategy.

Hybrid cloud offers the best of both worlds. But only if it has been thoughtfully, meticulously deployed.

With the Intel® Xeon® Scalable processor family, IT can be certain that its hybrid cloud infrastructure is powered by Intel's high performance and exceptionally scalable platform—one that is truly future ready to handle the complex requirements of modern data centers. In addition, the offerings in Intel® Select Solutions—both hardware- and software-based—can help you select and deploy modern data center solutions simply and quickly.

Partner with Intel on the journey to modernizing your data center and innovating your business.

LEARN MORE

¹ <https://www.rightscale.com/blog/cloud-industry-insights/cloud-computing-trends-2017-state-cloud-survey>

² Gartner, <https://www.gartner.com/newsroom/id/3666917>

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.

Intel, the Intel logo and Xeon are trademarks of Intel Corporation or its subsidiaries in the U.S. and/or other countries.

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

© Intel Corporation.