

## Cloud Computing Cost: Saving with a Hybrid Model

- A hybrid cloud model can help us achieve high levels of agility, scalability, and efficiency.
- When making hosting decisions, we consider not only cost but also business velocity requirements and capability maturity.
- Our theoretical financial model associates a hybrid cloud hosting solution with increasing savings.

Intel IT believes that a hybrid cloud model (a mix of public and private clouds) will enable us to dynamically adjust the amount of capacity we are using in the public or private hosting environment, thereby achieving high levels of agility, scalability, and efficiency. We plan to use cutting-edge public cloud providers while demand for a new service grows; once demand stabilizes and becomes predictable, we can move that service to our private cloud. This approach helps us avoid building excessive private cloud infrastructure for volatile demand peaks.

We weigh the advantages and disadvantages of several factors when considering cloud computing hosting options—we do not simply look for the least expensive solution. Business velocity requirements, capability maturity, and cost all contribute to our hosting decisions (see Figure 1).

We have developed a theoretical financial model for analyzing cloud computing cost. While this model is still in the proof-of-concept stage, we believe that it can provide valuable information about the future savings associated with a hybrid cloud hosting solution.

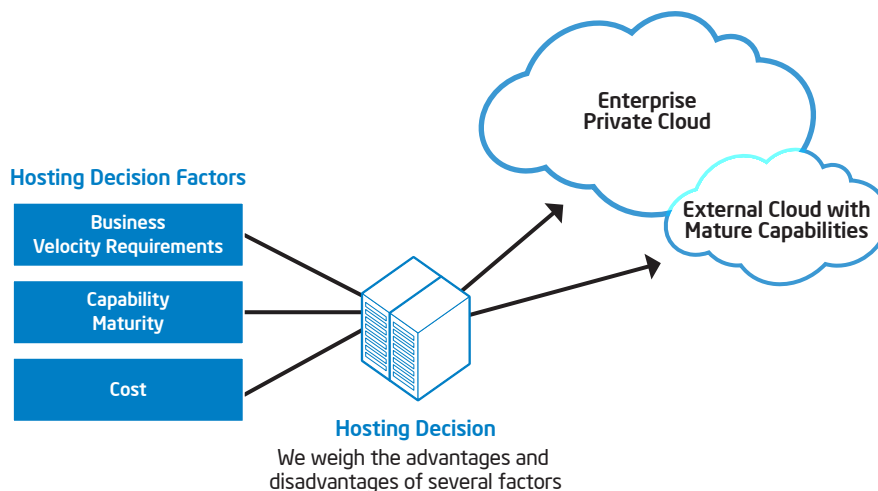


Figure 1. Our analysis of cloud computing needs includes business velocity requirements, capability maturity, and cost.

## Background

Intel IT needed to determine an investment strategy that could meet the growing capacity needs for our new businesses that required hosting solutions. We could have built out the capacity and capability internally. Or we could have outsourced everything to a public cloud provider. We chose a third option—a hybrid cloud model (a mix of public and private clouds). This flexible approach allows us to dynamically adjust the amount of capacity we are using in the public or private hosting environment, thereby achieving high levels of agility, scalability, and efficiency.

We are working to accelerate Intel's adoption of a hybrid cloud by establishing the following key design strategies:

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

Developing a hybrid cloud model required us to look at various trade-offs between the public and private cloud environments, choosing a mix that represents the best solution from several perspectives.

## Hybrid Cloud Computing Considerations

In our quest to accelerate transformation to a large-scale automated hybrid cloud infrastructure, we decided to use public cloud providers with the most up-to-date technology and service levels while demand for a new service grows. Once demand stabilizes and becomes predictable, we can move that service to our private cloud. This approach helps us avoid building excessive private cloud infrastructure for volatile demand peaks.

When analyzing cloud computing hosting needs, we look at more than cost; we also consider business velocity requirements and capability maturity. For example, one public cloud provider may be more expensive than another but may offer extra capabilities, such as enhanced performance, reliability, or security, that make

that provider the better choice for a particular service overall. Or outsourced hosting may be less expensive, but for a particular service that needs the highest level of security, we may need to host that service on our enterprise private cloud.

### BUSINESS VELOCITY

To support our goal of enabling developers to turn innovative ideas into production services in less than a day, we emphasize velocity. With the ever-increasing speed being demanded in the business environment, on-demand self-service has become the norm within our enterprise for compute infrastructure as a service. We continue to enable more advanced use cases for self-service, and the next phase of our cloud implementation will enable elastic capacity that can provision large numbers of virtual machines (VMs) within minutes and provide complete application environments with one click of a button or one API call.

In addition, we have moved forward with our platform-as-a-service implementation, which allows our software developers to focus on writing code instead of the complexities of the underlying infrastructure. To bolster our solution capabilities and capacity, we focus on hosting solutions that can automatically handle sourcing decisions at provision time and peaks in demand at runtime.

### CAPABILITY MATURITY

For us, a public cloud provider must offer the following core capabilities, which are the foundation of the Open Data Center Alliance Cloud Maturity Model:<sup>1</sup>

- **Federation.** Federation refers to the ability of identity and access management software to be able to securely share user identities and profiles. This allows users within a specific organization to use resources located in multiple clouds without having to generate separate credentials in each cloud. Federation allows IT to manage one set of identities,

<sup>1</sup> For more information on the Open Data Center Alliance Cloud Maturity Model, refer to the ODCA web page: [www.opendatacenteralliance.org](http://www.opendatacenteralliance.org); in particular, the documents listed under "Our Work." In addition to working with the ODCA, Intel is furthering its cloud computing vision through its participation in the [Intel® Cloud Builders](#) initiative and the [OpenStack\\*](#) project.

authorizations, and security review processes. From the user perspective, this ability enables seamless integration with systems and applications.

- **Interoperability.** There are two key concepts of interoperability. The first concept is the ability to connect two applications that are concurrently running in cloud environments. The second concept is the ability to easily port an application from one cloud to another. Both concepts involve the use of standard mechanisms for service orchestration and management, enabling elastic operation and flexibility for dynamic business models, while minimizing supplier lock-in.
- **Open standards.** The term "open" refers to both software and standards. Open source software operates at a fast rate of change supported by diverse, vibrant community updates. These frequent update cycles provide access to the latest features and capabilities, including performance and efficiency improvements. The use of common APIs or abstraction layers makes it easier for end users to rapidly consume cloud services from different providers to meet business requirements. Even if software is not open source, it should adhere to open standards, in order to maximize the benefits of cloud deployment.

When making hosting decisions we also consider security and health monitoring capabilities.

- **Security.** The public cloud providers we choose must feature a broad set of security capabilities that can protect data and intellectual property. Some technologies that can help increase security include private virtual LANs that isolate VMs and separate network and server administrative duties. Intel® architecture can help provide secure virtualization capabilities through hardware-assisted security, such as that provided by Intel® Trusted Execution Technology and Intel® Advanced Encryption Standard New Instructions. We are also exploring data anonymization, which can enhance cloud security.
- **End-to-end health monitoring.** Because of our high expectation for availability, health monitoring is one of the most critical elements of managing a hybrid cloud

environment. Our goal is to provide the experience of no downtime to both our Intel employees and our customers. This means all applications should attain 99.99-percent availability, which amounts to a little more than 52 minutes of annual downtime.

Monitoring must include facilities, network, storage, compute, and applications. Monitoring must also be automated and supported by advanced business intelligence tools. We examine service-level agreements (SLAs) carefully to verify whether the guaranteed levels of service match our needs. If a sufficient public SLA is not available, we may consider hosting a particular service internally.

We also consider the level of automation, because an automated cloud enables cloud-computing services and resources to be specified, located, and more securely provisioned with little or no human intervention. Yet another capability we consider is whether a cloud provider offers client-aware services. A client-aware, device-savvy cloud can intelligently take advantage of client capabilities to deliver the best user experience anytime, anywhere.

**COST**

In addition to analyzing the business velocity and capability maturity of hosting solutions, we have developed a theoretical financial model for analyzing cloud computing cost. We are currently testing this financial model through a multiphase proof of concept (PoC). During this PoC we are implementing a hybrid cloud hosting solution for both internally and externally facing offerings. We believe the theoretical financial model, illustrated in Figure 2, can provide information about the future savings associated with this hybrid cloud hosting solution.

- The orange line represents the cumulative investment that would be required for 100-percent outsourced cloud computing; the black line represents the growth in demand for VMs.
- The green (public) and blue (private) bars represent a hybrid cloud hosting solution.
- The gap between the 100-percent outsourced line and the top of the hybrid bars (orange shading) represents the potential cost savings associated with a hybrid model.

The theoretical financial model is a reasonable assessment of variable cloud computing cost. In reality, the hybrid cost savings may vary from those shown in the figure, depending on the rate of cost-efficiency improvements in both the public cloud and in Intel’s private cloud. Public and private hosting costs are both decreasing rapidly; however, we believe that due to the operational efficiencies we are putting in place, including new technologies, training efficiencies, and use of open source solutions, hybrid costs will decrease as fast as or faster than public cloud costs. It should be noted that the efficiencies public providers have introduced into large-scale computing are a catalyst to help direct and lead the way for all computing—including private clouds. We believe this optimization is healthy for the industry and forces a new approach, which in some scenarios means a disruptive change in solutions as well as the work force.

In the model, the exact ratio of public and private hosting varies over time. Our goal is to determine an efficient “private to public” ratio that could vary depending on the demand in terms of capacity and capability. Currently, we believe that ratio will be about two-thirds private to one-third public.

**Projected Costs for Cumulative Public and Private Hybrid Cloud Hosting (by quarter)**




Figure 2. Our theoretical financial model for a hybrid cloud hosting solution shows that the hybrid approach can potentially result in growing savings compared to using a purely outsourced hosting solution.

## Conclusion

In our PoC, a mix of public and private cloud hosting is enabling us to develop and launch products and services rapidly and efficiently, allocating resources to meet demand. Eventually, by transforming Intel's cloud usage model into a large-scale automated hybrid cloud infrastructure, we intend to achieve higher levels of agility, scalability, and efficiency.

Intel IT is executing our strategic focus on cloud computing, which emphasizes federation, interoperability, and openness. Therefore, when making cloud hosting decisions, we balance the business velocity and capability maturity of the hosting solution against an analysis of cloud computing cost. As the open industry grows for public clouds, we will place greater emphasis on providers that enable open standards and solutions. We will continue to validate and refine our theoretical hybrid cloud financial model as well as implement operational efficiencies that further improve our hybrid cloud deployments.

## Related Information

Also refer to the following related white papers:

- "Developing a Highly Available, Dynamic Hybrid Cloud Environment"
- "Overcoming Security Challenges to Virtualize Internet-facing Applications"
- "Best Practices for Building an Enterprise Private Cloud"
- "Enhancing Cloud Security Using Data Anonymization"
- "Planning for eDiscovery in the Cloud"

**For more straight talk  
on current topics from  
Intel's IT leaders, visit  
[www.intel.com/it](http://www.intel.com/it)**

## AUTHORS

**Bob Barnard**  
Finance Analyst,  
Intel IT

**Ruchi Bhargava**  
Hybrid Cloud Program Owner,  
Intel IT

**Das Kamhout**  
Principal Engineering,  
Intel IT Cloud Computing Lead

THE INFORMATION PROVIDED IN THIS PAPER IS INTENDED TO BE GENERAL IN NATURE AND IS NOT SPECIFIC GUIDANCE. RECOMMENDATIONS (INCLUDING POTENTIAL COST SAVINGS) ARE BASED UPON INTEL'S EXPERIENCE AND ARE ESTIMATES ONLY. INTEL DOES NOT GUARANTEE OR WARRANT OTHERS WILL OBTAIN SIMILAR RESULTS.

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.

Intel and the Intel logo are trademarks of Intel Corporation in the U.S. and other countries.

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

Copyright © 2013 Intel Corporation. All rights reserved.

Printed in USA

 Please Recycle

0913/WWES/KC/PDF

329131-001US

