IT@Intel White Paper

Intel IT
IT Best Practices
Enterprise Private Cloud Implementation
June 2014

Simplifying the Path for Building an Enterprise Private Cloud

Providing an opportunity to innovate from within, our private cloud plays a significant role in advancing our progress to a federated, interoperable, and open cloud.

Catherine Spence
Principal Engineer, Intel IT

Sridhar Mahankali
Cloud Architect, Intel IT

Munir Ghamrawi
Cloud Automation Engineer, Intel IT

Aaron Huber
PaaS Engineering Lead, Intel IT

Greg Bunce
Automation Engineering Lead, Intel IT

Executive Overview

Intel IT's enterprise private cloud enables us to deploy a virtual machine in under 10 minutes, provide on-demand self-service to Intel users, and offer a rich set of services for software development. Providing an opportunity to innovate from within, our private cloud plays a significant role in advancing our progress to a federated, interoperable, and open cloud.

Our recent focus on open source cloud technologies is simplifying and accelerating our users' consumption of private cloud services. By enabling a common orchestration layer and control plane, open cloud technologies help us increase automation of resource management and will provide greater business flexibility as we progress in our ability to consolidate public and private clouds much like a cloud service provider. At the same time, we continue to improve Intel's business agility and velocity while reducing operating costs. Intel IT now delivers more than 85 percent of new services in the cloud in our Office, Enterprise, and Services data center environments.

We see many additional benefits in our open source, open standards approach to our private cloud:

- Opportunity to take a leadership role in Intel's transformation to a large-scale automated hybrid cloud infrastructure and in the open cloud ecosystem
- Increased productivity through improvements in how we expose applications and data to Intel's application developers
- Greater redundancy, capacity management, and automation as we rapidly progress toward our goal of zero downtime for cloud-architected applications
- Improved service levels and compliance as we operate our hosting services similarly to a public cloud service

As Intel continues to innovate and grow, new lines of business—from our manufacturing and supply chain to sales and marketing—will create urgent, unique IT demands. We plan to meet these demands by continuing our cloud journey to a hybrid cloud model where we can dynamically move applications to the location—public and private, on and off premises—depending on what makes the most sense based on user experience, performance, cost, and the context of current cloud activity. The progress we are making in our private cloud today is leading the way to developing a common orchestration layer and control plane for providing Intel business units with public and private cloud services.

Contents

IT@INTEL

The IT@Intel program connects IT professionals around the world with their peers inside our organization—sharing lessons learned, methods, and strategies. Our goal is simple: share Intel IT best practices that create business value and make IT a competitive advantage. Visit us today at www.intel.com/IT or contact your local Intel representative if you'd like to learn more.

BUSINESS BACKGROUND

Intel IT operates a massive, worldwide computing environment that supports about 100,000 Intel employees and includes approximately 55,000 servers in 64 data centers. Employees connect to our services through more than 147,000 devices. While the majority of these devices are mobile business PCs, over 43,000 are handheld devices.

The need to support a large organization makes cloud operations vital to Intel's business. To provide the best service to our users, our cloud goal is to enable a highly available, shared, elastic infrastructure in which secure, standardized application platforms and data are delivered on demand through self-service portals. We work to achieve this goal through three cloud delivery models: public, private, and hybrid.

Through our public- and private-cloud services, we provide Intel business units with the following:

- Software as a service (SaaS). On-demand software applications or services
- Platform as a service (PaaS). On-demand software development and hosting
- Infrastructure as a service (IaaS). On-demand compute infrastructure
- Database as a service (DBaaS). On-demand database solutions

Along with other large enterprises, Intel is on a multiyear path to cloud maturation based on the Open Data Center Alliance cloud maturity model (see Figure 1). Our ultimate objective is to implement a federated, interoperable, and open standards-based cloud that enables hybrid applications spanning private and public clouds to become the norm. This mature cloud model will help us optimize costs, balance workloads, and enable seamless resource sharing between services for greater end-user productivity.

Our Private Cloud Journey

We began our cloud journey in 2009 to realize the cost benefits available in the more efficient use of server resources through virtualization and data center consolidation. We started with a private cloud and compute IaaS.

A private cloud is a single-tenant computing environment built on a highly efficient automated and virtualized infrastructure. On premises, it is an enterprise private cloud. When hosted by a public cloud provider, it is a virtual

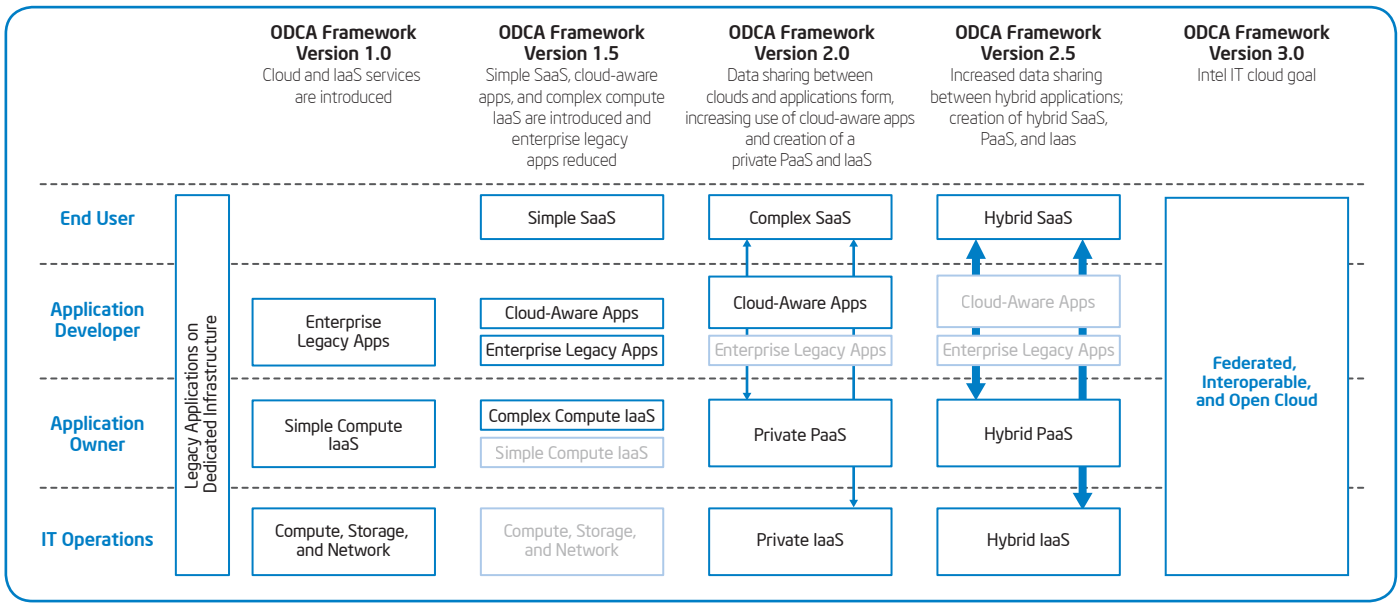

Figure 1. Intel IT is using the Open Data Center Alliance's Cloud Maturity Model to help guide our progress to a federated, interoperable, and open cloud, which industry organizations recommend.

private cloud. While our enterprise private cloud is a single-tenant environment in the sense that it is owned, used, and maintained by Intel IT, our private cloud can also be considered a multitenant environment because it serves multiple Intel business units. Groups of related applications from these business units are typically categorized as a tenant.

An enterprise private cloud offers many of the agility and efficiency benefits of public clouds without the potential risks associated with having highly sensitive applications and data outside the enterprise perimeter. Our initial strategy was to engage with cloud technology "from the inside out" by implementing a private cloud IaaS while utilizing SaaS in key areas through public cloud services.[1] The use of cloud technologies within the enterprise positions us to eventually take advantage of public clouds in a hybrid model, where components of applications run across multiple clouds.

We implemented a private cloud in 2010 initially for office, enterprise, and services applications. This first enterprise private cloud was primarily a virtualization of our data center server environment. We implemented it through existing enterprise manageability tools and solutions, coupled with integration software and databases. The result was a cohesive solution combining compute, storage, and network resources. We immediately saw cost-saving benefits from virtualization and consolidation, realizing better utilization of our server capacity. We later added self-service capabilities to increase end-user agility and productivity.

Why a Private Cloud?

Intel currently consumes services from public clouds to address specific business needs and locations, but we prefer to consume our own capacity before using an external provider. In operating our enterprise private cloud, we are not looking to directly compete with retail cloud providers. Our charter is to serve Intel business units and their users. However, as part of the cloud maturation process that we are working toward, the line between how an employee uses the enterprise environment and the public marketplace is blurring.

Considerable debate exists over whether public or private clouds are more cost effective. Public cloud providers are looking to make a profit; IT organizations are not for profit. As public cloud providers gain economies of scale, though, costs are shrinking and becoming more competitive.

We find that the most important reasons to continue developing our private cloud are flexibility and agility. It would be easy to get locked into a public cloud provider since there is little interoperability among clouds. Factors that cause high switching costs go beyond the workload itself to include everything around it, such as data, monitoring, service-level agreements, and security. As for agility, a private cloud enables our business units and their developers to rapidly acquire and manage their own cloud capacity, and—within predefined limits—to dynamically scale resources to meet their application needs.

To use public clouds effectively, we need to develop more cloud-aware applications that can intelligently use infrastructure and maximize cloud advantages. We continue to work with Intel's developers on expediting the production of such applications.

Addressing enterprise and business security concerns is also a priority for our cloud strategy. Intel's private cloud allows for ongoing evaluation of security controls and capabilities to address risks and concerns emerging from our continuing cloud adoption. Using this approach, we can easily adapt existing controls into our on-premises private cloud while exploring the transition to security capabilities delivered as cloud services. This approach is critical to addressing the evolution of enterprise security beyond the traditional perimeter defense model.

Our private cloud is already enabling us to extend higher levels of availability to all applications without the need for costly specialized hardware and software. This increase in reliability is due to new high-availability capabilities in virtualization software and the availability of mission-critical features, such as Machine Check Architecture Recovery in higher-end, industry-standard servers based on the latest Intel® Xeon® processors.

[1] For more information, see the IT@Intel white paper, "Developing an Enterprise Cloud Computing Strategy."

Among our measures of success, the following two stand out:

- **Establishment of a federated capacity,** which through the use of multitenancy and the pooling of assets increased resource sharing and delivered over USD 9 million in savings over three years.
- **Faster delivery of on-demand services**, allowing application developers to acquire a virtual machine (VM) with a service-level agreement in less than 3 hours. Previously, it took up to 90 days to provision a server in response to a request for new capacity.

Evolution to Open Cloud

The second and current instantiation of our private cloud is primarily based on open standards and open source technologies. It represents the next step in our journey to a federated, interoperable, and open cloud. Using open source software to create an open cloud provides a clear path to treating the public cloud environment as an extension of our data center services.

We consider open source solutions essential for accelerating customer application development and facilitating PaaS, DBaaS, and cloud-aware design principles. Open source cloud software offers us many advantages, including the following:

- Eliminates having to write, test, and implement our own cloud software
- Avoids the high switching costs and licensing costs associated with proprietary sources
- Lets us take advantage of the many open source options now available that did not exist when we first embarked on our private cloud initiative
- Enables us to benefit from and participate in the strong, vibrant open source community and enterprise-level support options available through third-party companies
- Allows us to take advantage of the frequent update cycles of the open source community, creating fast implementation of the latest features, capabilities, and performance improvements
- Enables us to write core code with other members of the community and share that code to move faster
- Allows us to choose to preserve aspects of our environment that we care about—such as infrastructure investments and virtualization expertise—while providing a constantly advancing software platform for delivery of our next-generation hosting capabilities

In addition to open source software, in many cases we use a traditional software model based on factors such as technology maturity, costs, or specific capabilities. To make those components work together effectively, we rely on open standards to address overall integration, optimization, and automation opportunities.

To make our cloud consumer friendly, we model our operations after cloud service providers. This helps our users understand our offerings and puts them in a position to compare and consume services. By making it easy to consume our IT services, users have more incentive to use our services first before paying for external services. By consolidating workloads we can get better cost and resource efficiency for Intel as a whole.

The provider model positions us for future hybrid scenarios where public cloud environments look and work like an extension of our infrastructure, providing applications and data as a seamless, secure user experience. Whether a cloud is internal or external, users need easy access to their applications and data in a secure fashion, anywhere and anytime from any device.

In the delivery of our cloud, our goal is to operate our private cloud as efficiently as a cloud provider does, running a cloud to scale using agile methodologies (as opposed to IT waterfall processes) and DevOps support.

Agile methodologies are a group of software development methods based on iterative and incremental development where requirements and solutions evolve through collaboration between cross-functional teams. We are using an agile approach in the development of our cloud. In addition, application development teams are using agile in their development of applications that land on our enterprise private cloud.

DevOps is a software development method that stresses communication, collaboration, and integration between IT engineering developers and IT operations professionals. By operating our cloud using the DevOps model, we are seeing the integration of duties between development and operations teams where developers largely take over support and management of their applications using highly automated, agile infrastructure. Increasingly, development teams use the cloud services that we offer to facilitate their use of DevOps in developing and supporting their applications.

The key to DevOps is the availability of web services at every layer to make it possible to automate common tasks. Automation promotes the ability to scale and reduces error-prone manual work, which makes it possible for a smaller team to support larger environments.

The above methodologies and technologies are important steps and capabilities for meeting our ultimate goal of a hybrid cloud model where we can dynamically move applications to the most appropriate location, on and off premises, based on user experience, performance, cost, security, and context of current cloud activity. In this age of cloud, we believe we need to be in the position to consolidate workloads and get the best value, in terms of cost and agility, whether using a public or private cloud. We also believe that to coordinate public and private clouds through a common orchestration layer, we first need to implement and operate such a layer seamlessly in our private cloud.

BEST PRACTICES

To operate our private cloud like a cloud service provider requires maximizing our resource efficiency while providing on-demand, self-service PaaS, as well as on-demand, self-service IaaS for compute, storage, and network options. To achieve these objectives, we are expanding our use of OpenStack* software, making it the single control plane for all our hosting virtualization environments.

Our platform also takes advantage of Cloud Foundry*, another open source project, to implement PaaS. To manage our cloud operations at each level of the cloud capability stack, we are following Information Technology Infrastructure Library* (ITIL) standards. In addition, we are continuing to advance our cloud solution delivery models to improve the productivity and experience of our customers.

In this section of the paper, we discuss our updated best practices and best-known methods (BKMs) in three areas:

- Cloud hosting strategy
- Cloud solution delivery models
- Cloud-aware applications

Our focus on these areas is helping us progress toward a federated, interoperable, and open cloud that enables hybrid applications spanning private and public clouds.

Implementing a Provider-Like Cloud Hosting Strategy

To provide excellent service, we believe that Intel IT must operate more like a cloud service provider in making the acquisition and use of services quick and easy. At the same time, we remain focused on cloud federation, maximizing agility and efficiency by enabling easy movement of workloads and services between IT infrastructures within the enterprise. Eventually, we plan to enable the same fluidity between private and public clouds. By acting as a cloud broker, we can optimize costs through consolidation and drive IT standard requirements, such as security, across our portfolio of cloud services.

A key measure of our cloud initiative success is our ability to provide on-demand self-service delivered through dynamic, flexibly configurable resource pools. Compute, storage, and network functions must all be abstracted and offered as services (see Figure 2). Services must be accessible over the Internet or corporate network, and resources must be shared among multiple users (multitenancy).

A final element of our cloud hosting strategy is regular server refreshes. To improve our ability to cost-effectively meet current and future compute needs within existing data centers, we are transitioning to servers based on the Intel® Xeon® processor E7 v2 family.

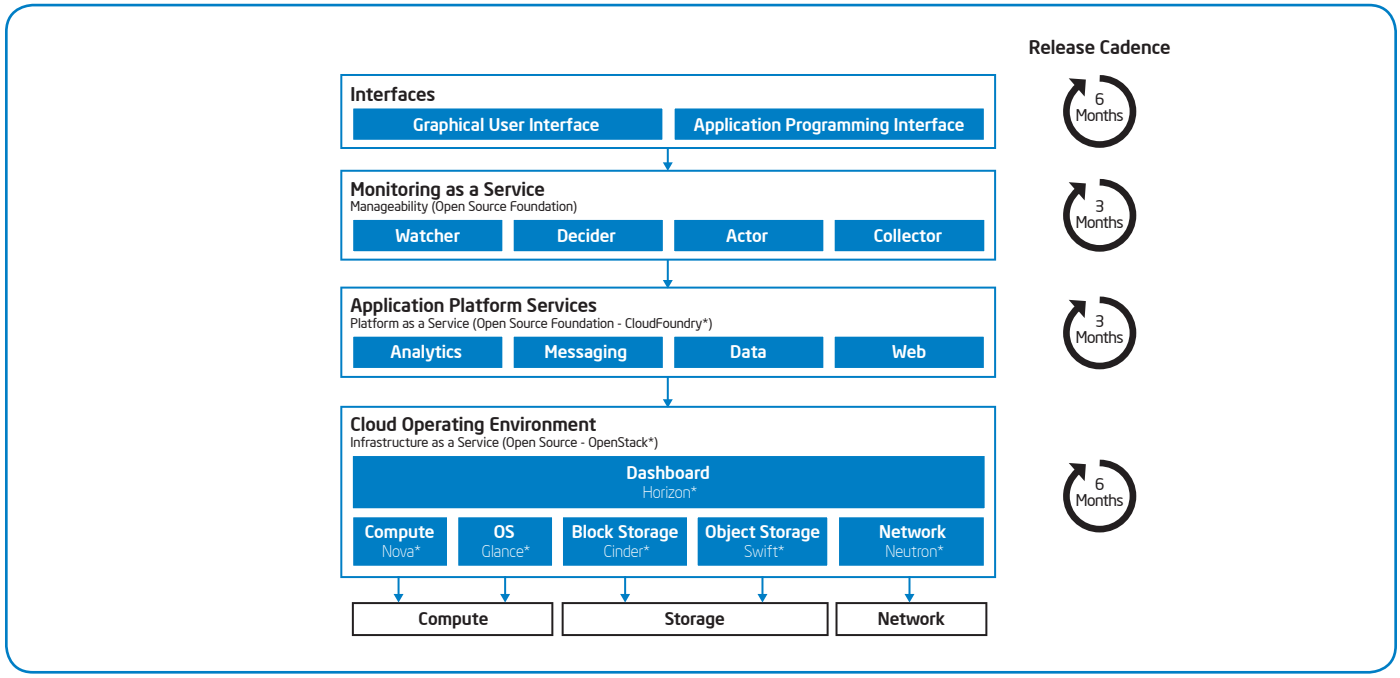

Figure 2. The Intel IT Cloud Platform Solution Stack showing open source and other components, and associated refresh cycles.

BKM 1: USE INFORMATION TECHNOLOGY INFRASTRUCTURE STANDARDS TO MANAGE RESOURCES

To manage our cloud operations efficiently at each level of the cloud capability stack, we implement ITIL standards (see Figure 3). These standards provide a holistic view for our entire cloud solution stack, including physical devices and containers, functions management, data integration and analytics, workflow automation, and service delivery.

Implementing ITIL standards enables us to enhance our use of automated monitoring and configuration tools and improve on-demand service by providing the following management layers:

- **Service delivery.** Provides IaaS and PaaS.
- **Workflow automation.** Establishes a modular, extensible framework that simplifies integration of many aspects of cloud computing and provides an on-demand, highly available, and scalable cloud computing infrastructure for rapid VM provisioning and deprovisioning.
- **Data integration and analytics.** Integrates data from all containers as well as from the functions management layer.
- **Functions management.** Watches and acts on configuration management, event management, change management, and capacity management.
- **Device or container.** Includes VMs, storage containers, and agents.
- **Physical devices.** Encompasses compute, storage, and network hardware platforms.
- **Facility.** Represents an entire data center.

We believe that using ITIL standards to institute a hosting automation framework that includes entitlement, quotas, transparent measured services, and data-driven business logic is an important best practice. These

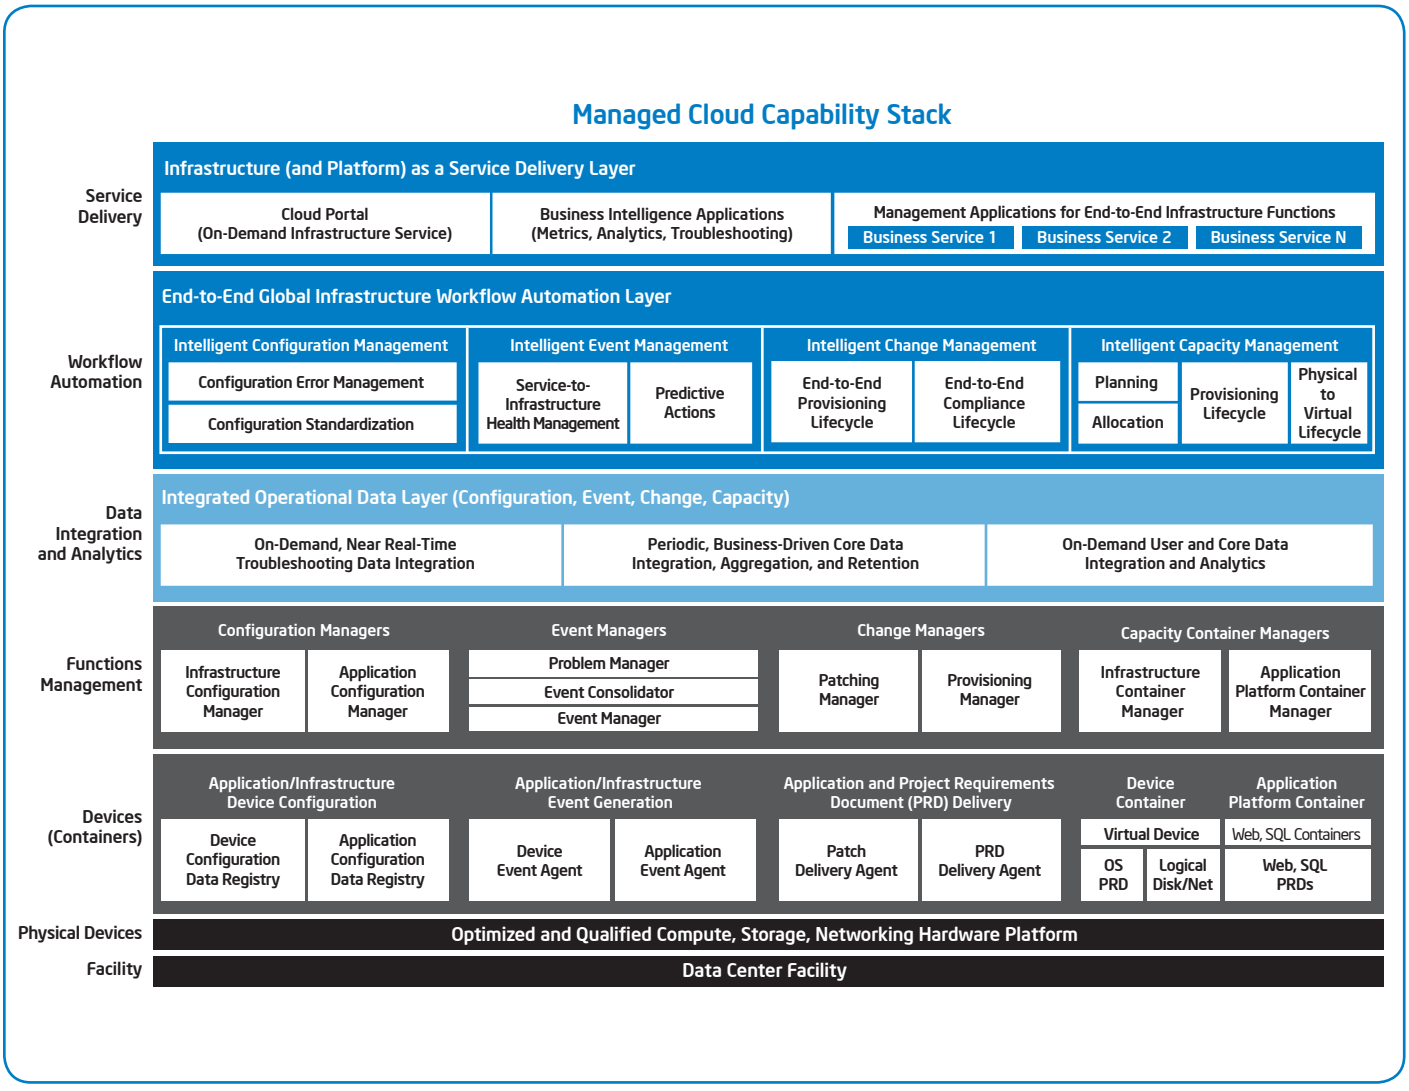

Figure 3. Information Technology Infrastructure Library* standards provide a holistic view of our entire cloud solution stack.

standards enable us to operate a true enterprise private cloud that provides a consumer-focused, self-service portal that improves our business agility and velocity.

BKM 2: IMPLEMENT AN OPENSTACK CLOUD

To expose core cloud infrastructure resources (compute, storage, and networking) in a programmable manner, we chose an open source cloud OS, OpenStack, and implemented it for all our hosting virtualization environments (see Figure 4). OpenStack enables us to advance from self-service for only compute (server) provisioning to self-service for storage and network elements as well.

As a convergence platform for IT hosting, OpenStack provides a dashboard that gives administrators management control. A web interface provides the self-service functionality, enabling users to select and provision resources readily. The combination of OpenStack with our automated monitoring and configuration tools allows us to more easily deliver rapid provisioning of new capacity for Intel's application developers.

BKM 3: MAXIMIZE AGILITY AND EFFICIENCY WITH A SINGLE CONTROL PLANE

Using OpenStack as our single control plane for our entire virtualization environment is a major part of our strategy in 2014.[2] OpenStack exposes our IaaS to the end user whether that user interacts with our cloud through the web portal, APIs, or the OpenStack command line interface (CLI). Instances of PaaS are provisioned on top of IaaS through the control plane.

Under the OpenStack layer is the hosting environment. This environment can be based on a commercial hypervisor or OpenStack hypervisor, such as Kernel-based Virtual Machine, and is abstracted from our users. Such abstraction makes sense because users do not care what the underlying infrastructure is; they just want a hosting service.

[2] More on the origins of our open source journey using OpenStack can be found in the IT@Intel white paper: "Accelerating Deployment of Cloud of Services Using Open Source Software."

Integrating OpenStack with our implementation of automated monitoring and configuration tools gives us the following advantages:

- Rapid provisioning of new capacity for application developers
- Support for an active/active application design that through greater redundancy and automation enables us to make faster progress toward our goal of zero downtime for cloud architected applications
- Shorter software upgrade cycles, enabling faster implementation of the latest improvements, features, and functionality, leading to more state-of-the-art, consumable cloud services
- Improved service levels and compliance

Current efforts include employing the OpenStack-based control plane to maximize the use of maturing open source capabilities while protecting existing infrastructure investments and providing a common interface for Intel's hosting environments. To achieve this balance, we are taking the following actions:

- Require suppliers to provide OpenStack APIs
- Deploy cost-optimized infrastructure through organic service growth, refresh, and capability maturity
- Transition to multi-hypervisors to manage our compute services
- Continue our workforce training to acquire open source skillsets
- Together, these steps will help improve the efficiency of our use of open software throughout our private cloud. They will also help ensure our smooth progression to a federated, interoperable, and open cloud.

MaaS

We are currently working to implement bare metal as a service (MaaS), a layer underneath IaaS. MaaS helps facilitate and automate the deployment and dynamic provisioning of physical hardware for complex scalable services, such as OpenStack cloud infrastructure.

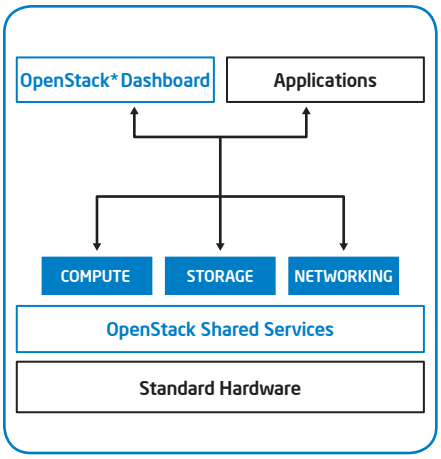

Figure 4. OpenStack* is an open source cloud OS that controls large pools of compute, storage, and networking resources throughout a data center. These resources are managed through a dashboard that gives administrators control while empowering users to provision resources through a web interface.

We have three primary infrastructure use cases for MaaS:

- Provision and manage IT infrastructure, particularly our own cloud infrastructure.
- Deliver on-demand, self-service, consumer provisioning and management of IT-hosted infrastructure. MaaS enables end users to provision physical devices just as they would provision VMs for workloads that demand it.
- Manage and provision non-IT managed infrastructure in hardware and software lab settings. In this use case, we provide the capability and the labs run their business.

We expect our first implementations of MaaS sometime in 2015.

BKM 4: OPTIMIZE DATA CENTER SPACE AND COSTS WITH REGULAR SERVER REFRESHES

A high rate of virtualization combined with a three-to-four-year cycle of server refreshes capitalizing on the latest Intel® Xeon® processors has enabled Intel IT to meet growing compute and storage demand, while significantly decreasing the number of required physical servers. From 2009 to 2012, we reduced physical servers by 60 percent and achieved an approximate 5.5x increase in the number of VMs.

Processor and server advances have allowed us to continue to meet our 35-percent annual increased capacity demands in the same amount of data center space. We find that optimizing the use of existing space is a better value proposition than building new data centers and delivers a faster return on investment. Building a data center has a long lead time; a server refresh delivers benefits in days. In addition, as long as server prices remain flat, the yearly performance gains from deployment of the latest Intel® processors come at no extra cost.

We are currently optimizing existing space instead of building new data centers by refreshing with servers powered by the latest Intel Xeon processor E7 v2 family. Based on Intel's 22nm process technology, these processors feature 50 percent more cores (up to 15), up to 30 threads per socket, and 25 percent more cache to deliver up to twice the average performance of previous versions.[3] This increased performance along with up to four times the bandwidth through integrated PCI Express* 3.08 has enabled us to realize 84 percent higher throughput compared to the previous generation Intel® Xeon® processor E7 4870.[4]

With three times the memory capacity of previous versions, the Intel Xeon processor E7 v2 family provides the flexibility to handle a variety of configurations and workloads, including data-demanding, transaction-intensive workloads. The processor's Machine Check Architecture recovery execution path allows the system to recover from certain errors that would otherwise be fatal. This feature and others help deliver an uptime equivalent to the best-in-class RISC-based platforms.

[3] For configuration details and other information regarding these performance claims, see the Intel product brief, "The Foundation for Better Business Intelligence."

[4] Ibid.

Advancing Open Source Solutions through Shared Communities

OpenStack* and Cloud Foundry* are developed and supported by a global collaboration of developers and cloud computing technologists. The goal is to make each project easy to implement, scalable, and feature-rich. The technologies in each case consist of a series of interrelated projects delivering various components for cloud software and hardware infrastructure solutions. All OpenStack and Cloud Foundry source code is available under the Apache 2.0 license.

As a large IT organization, we recognize the importance of working with open source communities to ensure that solutions meet enterprise needs. Intel is a Gold member of the OpenStack community, providing funding, strategic alignment, and expertise with the OpenStack mission. Intel IT staff is closely involved, contributing code, documentation, and reporting bugs and vulnerabilities—efforts all aimed at improving OpenStack's usefulness to Intel and other IT organizations. This helps ensure that IT operations that choose OpenStack for their cloud computing infrastructure derive maximum value from investments such as ours in Intel-based hardware.

Intel participates in the Cloud Foundry community as a Gold Member of the Cloud Foundry Foundation. In addition, as a member of the Cloud Foundry Community Advisory Board, Intel IT provides strategic technical feedback on the Cloud Foundry roadmap, advocates for enterprise requirements, and provides more tactical input on the day-to-day operation of the Cloud Foundry project. This includes management and governance of the community, and its outreach initiatives.

By working through collaborative open source projects such as OpenStack and Cloud Foundry, and through the Open Data Center Alliance, the Intel® Cloud Builders program, academia, and industry research centers, Intel is helping to further the vision and advance the development of open, standards-based solutions for cloud computing. Through these efforts, we continue to emphasize and support greater automation, federation, and cloud awareness.

Advancing Cloud Services Delivery Models

We continue to make significant improvements in the performance of our IaaS and PaaS (including DBaaS) delivery models and the self-service experience of our customers who use them. In most cases, the improvements increasingly use open source solutions—OpenStack for IaaS and Cloud Foundry for PaaS.

BKM 5: ENABLE SELF-SERVICE IAAS

IaaS is well suited for hosting commercial off-the-shelf software and custom applications requiring control over the entire application stack. Our business units using IaaS support the entire stack or make special arrangements for Intel IT to do it through a managed hosting service.

With our IaaS solution, we expose our compute, storage, and network infrastructure through APIs and a GUI to enable on-demand self-service. A number of improvements (described below) are underway in our existing compute and storage solutions. We are also implementing software-defined networking (SDN) to improve network flexibility and capability and to enable customers to acquire network services faster and more easily.

Compute Solution Improvements

On the compute side, we are transitioning to a multiple hypervisor cloud environment and increasing our use of open source technologies, including our monitoring solutions and database technologies, wherever we can. Equally important in IaaS is the infrastructure supporting compute. We are currently deploying new servers based on the latest Intel Xeon processor E7 v2 family. On the security side, Intel® architecture helps provide secure virtualization capabilities through hardware-assisted security. Specific solutions include Intel® Trusted Execution Technology[5] and Intel® Advanced Encryption Standard New Instructions.[6]

As an additional security measure, we are exploring data anonymization. Data anonymization is the process of changing text that will be used or published in a way that prevents the identification of key information. De-anonymization, the reverse process, cross references the anonymous data with other data sources to re-identify the anonymous data source.

Storage Solution Improvements

Block storage continues to be a central part of our cloud storage infrastructure. A boot volume gets created for each VM at provision time. In addition, consistent with our strategy to expose core IaaS components through APIs, we have enabled self-service for creating, deleting, and attaching additional volumes to VMs and using the OpenStack Cinder* API.

We are in the early stages of using web-addressable object storage in our private cloud environment. This type of storage typically is well suited for write once, read many (WORM) types of applications and can generally be implemented on relatively lower cost and lower performance disk drives.

Network Service Improvements

Having automated the provisioning of compute and storage, we are now implementing SDN to make the network component self-service. SDN separates the control plane (the element of the network used to configure the network) from the data plane (where the actual packet flow and traffic traverse the network). A key advantage of SDN is that it moves management of the data center network away from each individual switch, enabling configuration from centralized controllers. This centralization means users can now either go to a web portal or use APIs to request network services on their own and make them part of their application environment. This agility is the key benefit for the user, enabling provisioning a network in significantly less time.

We are integrating SDN as an overlay network on top of the virtual environment and providing core network connectivity (routing, switching, and network access control) for VMs in our cloud environment. Using SDN has allowed us to abstract the virtual network environment from the physical network (or the underlay) and has simplified the network configurations required on the physical network. The SDN solution, once deployed, integrates with OpenStack's Neutron* APIs and can be configured programmatically and on-demand using those APIs. As a best practice, we see SDN as a maturing technology that will continue to evolve over the next few years.

We are also transitioning to a newer distributed network access control model for our cloud environments. With this model, security groups and network access control lists implemented within the OpenStack Neutron and SDN framework are providing network access control for communications within the cloud environment. The traditional hardware firewalls are providing access control at the edge of the data center and for communications between cloud and non-cloud resources. We have also introduced self-service for the most commonly used load balancer and global load balancer configurations.

Foundation Infrastructure Automation

We are in the process of piloting a foundation infrastructure automation framework that automates launching our OpenStack-based cloud foundational infrastructure and controllers. This type of automation allows us to build and sustain multiple private cloud instances in a consistent manner. We have

[5] No computer system can provide absolute security under all conditions. Intel® Trusted Execution Technology (Intel® TXT) requires a computer with Intel® Virtualization Technology, an Intel TXT-enabled processor, chipset, BIOS, Authenticated Code Modules, and an Intel TXT-compatible measured launched environment (MLE). Intel TXT also requires the system to contain a TPM v1.s. For more information, see www.intel.com/content/www/us/en/data-security/security-overview-general-technology.html

[6] Intel® Advanced Encryption Standard New Instructions (Intel® AES-NI) requires a computer system with an Intel AES-NI enabled processor, as well as non-Intel software to execute the instructions in the correct sequence. Intel AES-NI is available on select Intel® processors. For availability, consult your reseller or system manufacturer. For more information, see Intel® Advanced Encryption Standard Instructions (AES-NI).

established a framework that positions us to pull the latest open source packages, patches, and releases from the trunk (main OpenStack development area), package it, and integrate it into our private cloud. This type of framework allows us to benefit from the pace of change and innovation that the vibrant open source community is introducing into the cloud orchestration framework.

BKM 6: PROVIDE COMPREHENSIVE PAAS OPTIONS

PaaS is a preprovisioned environment (OS, abstracted middleware, and infrastructure) for building and rapidly hosting custom applications in the cloud. With the use of PaaS, developers are in control from development to deployment—exponentially reducing time to production, optimizing the use of resources, and encouraging the development of cloud-aware applications. Developers manage their applications; Intel IT manages the underlying platform and infrastructure.

Through self-service, PaaS increases programmer productivity and enables us to extend the value of our private cloud to more groups and usages. By exposing cloud value and usage models to more users, PaaS supports our technology roadmap for eventually using hybrid (private-public) clouds to further increase scalability and cost efficiency.

Developers find that PaaS makes it easier to build cloud applications (see Figure 5). Applications built using PaaS get inherent cloud benefits, such as elasticity, high availability, on-demand access, and a metered, multitenant environment. Pilot work in 2012 demonstrated that PaaS can make it possible for developers to transition from innovative idea to production in a single day. Since then, we have focused on our production PaaS offering, which includes a major update to the underlying software platform and the addition of a custom-developed PaaS portal.

An Open Source PaaS Solution

In selecting a PaaS solution that could run on top of our IaaS solution, we wanted an open source solution appropriate for our enterprise private cloud that could deliver agility, simplicity, standardization, and efficiency. These requirements led us to Cloud Foundry. This open source solution met our requirements in terms of technical capabilities and a wide array of supported programming languages and frameworks, such as Java*, Ruby*, Python*, and PHP*. A different open source software project called Iron Foundry extends Cloud Foundry for .NET applications. The advantage here is that our developers can use the same toolset and platform for application deployment, which is a considerable benefit in terms of platform flexibility and addressing the current and emerging needs of our developer community.

The Cloud Foundry platform was developed by VMware and first released in 2011 under the Apache* License 2.0. Like OpenStack, a vibrant community supports Cloud Foundry, offering frequent contributions, knowledge sharing, and third-party support options.

Our Cloud Foundry PaaS stack delivers even more agility benefits than our IaaS because it requires no server provisioning step and significantly reduces the amount of support that developers perform as part of self-service hosting. Currently, we offer two instances of production PaaS: one for internal-facing applications and one for external-facing (Internet) applications.

Recent PaaS Innovations

Intel IT first delivered a PaaS pilot based on Cloud Foundry and Iron Foundry open source software, landed in our open cloud exclave environment built on top of an OpenStack foundation. Based on the results of the pilot we began to engineer production-level PaaS based on the second version of Cloud Foundry. This major release provided a new system architecture and new tools. Through this upgrade, we are introducing some new innovations:

- **Warden containers.** Running within a VM, a warden container is an isolated environment used to partition applications. The benefit and the goal of the warden is that it addresses the "noisy neighbor" problem where the behavior of one application may affect the other applications residing on a node. Each container has its own CPU, memory, disk, and network usage limits. Every time an application is pushed or scaled, it is hosted in its own container.
- **Build packs.** These self-contained packages contain the instructions for how to support a particular application stack. For example, there are build packs for languages such as PHP, Ruby, and Python. Build packs are a convenient way to package framework and runtime support for an application. Not having this support

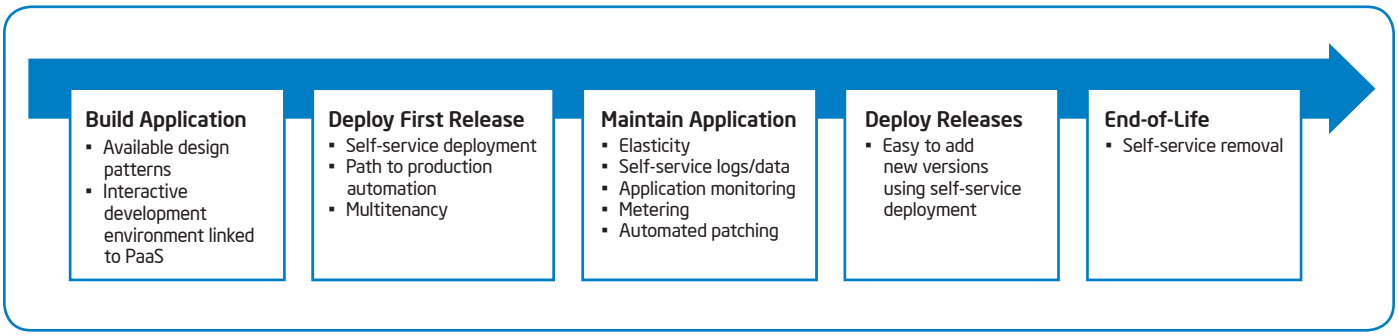

Figure 5. Our platform-as-a-service (PaaS) solution provides a quick, efficient way to build cloud applications.

tightly integrated into a platform makes upgrades, when required, less invasive. From a platform perspective, build packs also enable developers to use multiple versions of an application runtime for different applications simultaneously.

Build packs are attractive because they broaden the variety of programming languages and frameworks that PaaS can concurrently support. Developers who want to use an application framework without an available build pack have the option of creating their own build pack.

- **PaaS portal.** Cloud Foundry contains a CLI and APIs that enable developers to interact with PaaS instances. However, we are developing a PaaS portal that uses Cloud Foundry APIs to make it even easier for developers. The portal is an Intel IT custom-developed web-based interface. Over time, the portal will serve two main functions: (1) It creates a landing page on which all information for PaaS is collected and presented to current and prospective users, and (2) it provides a venue for users to sign up for accounts and deploy and manage applications.

  Users can still choose to utilize the CLI or the API. However, the PaaS portal has significant advantages:

  - Does not require a client download into the application development environment
  - Consolidates the capabilities together into an easy-to-use GUI
  - Provides a venue for future automation in areas such as high-availability scenarios including active/passive and active/active, security/governance automation, and linkage to other areas of application development

We anticipate over time that the portal will become the predominant way developers access resources and deploy and manage their applications.

Addition of DBaaS

DBaaS is an important subset in our PaaS offering. If an application requires a database, the customer can provision one through the DBaaS portal. This process involves specifying the database type (MySQL*, MongoDB*) and using the returned connection string. The connection string can be embedded in the application code and/or used with database management tools to create and manage application tables and data. We plan to increase our DBaaS offerings with additional choices and options.

The primary benefit of DBaaS for customers is increased agility through self-service. For Intel IT, benefits include reduced operational costs and better utilization of compute resources. As a best practice, we see DBaaS as an important step to running databases in a hybrid cloud in the future.

Each database in DBaaS inherits many standard features associated with DBaaS:

- **High availability.** Any database created through the DBaaS portal is created on multiple database instances for high availability of data. Any transaction that happens on any node is replicated to two additional nodes for data redundancy. Each database instance is on a separate VM to provide OS and machine redundancy that in turn is on a separate infrastructure host for hardware redundancy.
- **Security.** Each database instance is created with maximum security to ensure data safety. Each login is granted the proper permission to perform its task. Each application has a dedicated load balancer to isolate traffic. The infrastructures of all databases are protected with the right security permissions for access.
- **Disaster recovery.** All database instances are on a backup schedule that allows up-to-the-minute recovery in case of a disaster. We have a plan to offer automated disaster-recovery-site database failover when additional cloud sites are created.
- **Monitoring.** All databases are treated equally. Each database is added to our detail monitoring service for availability, failover, performance, and capacity. Alerts are sent to our system administration team for immediate resolution. Database specialists are alerted for immediate application assistance if needed. In addition, our DBaaS will offer application owners the ability to view their database service log to provide additional feedback.
- **Elasticity and right sizing.** DBaaS offers the capability for isolation from noisy neighbors. The capability enables us to port the database to any place in the cloud without impacting the application. Based on our monitoring alerts, we can add additional compute resources, such as memory, compute, or storage, to the instances when needed. We have monitoring in place to check the memory usage of the database and whether we need to increase the memory to the next size unit.
- **Index tuning and maintenance.** In addition to providing a database container, DBaaS monitors queries that are not tuned and that take a long time. We analyze any index optimization for running queries. We provide input to the developers when application code change is required. In addition, we provide standard maintenance procedures such as index fragmentation, log rotations, and security patching.

Enabling Cloud-Aware Applications

As Intel IT continues widespread adoption of cloud computing, Intel software development is undergoing a major shift. Intel architects and developers are learning to design cloud-aware applications that maximize cloud advantages, such as self-service provisioning, elasticity, run-anywhere design, multitenancy, and design for failure. Our developers are also increasingly using agile methodologies, taking advantage of the growing number of self-service tools and self-service data available in our PaaS offerings.

BKM 7: CLOUD-AWARE APPLICATIONS REQUIRE A DIFFERENT WAY OF THINKING.

To maximize cloud advantages, application developers need to incorporate new design principles. Four key areas we are addressing are the following:

- Be resilient to failure and latency.
- Don't be tied to specific infrastructure or location.
- Scale out/back based on demand.
- Provide self-service capabilities for the end user.

Our PaaS solution facilitates creation of such cloud-aware applications, providing templates, resource sharing, reusable web services, and large-scale multitenancy. These capabilities support our strategy of enabling developers to architect enterprise applications with a cloud back-end and multiplatform front-end, actively exposing and consuming web services offering built-in security.

BKM 8: NEW THINKING REQUIRES NEW TRAINING

To train developers in this new paradigm and validate our strategic agility target of innovative idea to production in one day, we are implementing one-day code-a-thons (hackathons) and other events. We believe that perfecting cloud-aware application development is a vital step to maximizing hybrid cloud advantages and bolstering the reliability, security, and agility of our enterprise applications.

RESULTS

Our open private cloud using OpenStack and Cloud Foundry has reduced the amount of time it takes to provision services and automatically resolve issues. We can now provision a VM in just 5 to 10 minutes, provide faster self-service to our customers, and offer a rich set of consumable services providing a foundation for leading-edge software development. We are making numerous advances to make our enterprise private cloud the hosting platform of choice for our cloud developers and our traditional enterprise IT teams.

Through increasing automation of resource management and other optimizations, the Intel IT enterprise private cloud is becoming the cornerstone in our goal of a federated, interoperable, and open cloud. It enables us to spend less time on engineering the core IaaS solution and more time on higher-level service areas, such as PaaS and DBaaS, which allow us to expose the more advanced services our application developers need to build applications that increase end-user productivity.

Major benefits we are seeing are increased customer self-service and control manifested in the following ways:

- Most new server requests in the Office and Enterprise environment use our self-service portal.
- On-demand self-service and automated sourcing decisions are now the norm, allowing users to secure infrastructure in less than 45 minutes.
- Application landing no longer requires multiweek lead times; innovative idea to production in less than a day is a reality for specific types of applications.
- Customer time spent managing application environment is significantly reduced.
- More advanced use cases for self-service are being continually added.
- Increased system resilience is being achieved in conjunction with cost reductions.
- Improved predictability and control of cloud resources provide valuable time savings.
- Reduced IT manual tasks and maintenance lessen IT service costs.
- Better capacity management improves use of resources.
- External cloud services are available for burst demand.

In the near term, we expect to achieve elastic capacity that can provision large numbers of VMs within minutes and provide complete application environments with one button click or API call. We are making it possible to land more application use cases in less than a day.

In our data centers, we are experiencing 80-percent effective asset utilization. We have reduced operating costs through larger pools in fewer data centers and established pervasive virtualization at our goal of 75 percent. We are also striving to achieve zero business impact through careful design of our automated, end-to-end service-managed cloud. And we are realizing increased availability through application design for failure and reduced mean time to recovery.

CONCLUSION

Having experienced early success with our enterprise private cloud, we continue to make strides in our journey to a federated, interoperable, and open cloud through implementing open solutions. Our current cloud infrastructure provides an excellent foundation for transforming our data center solutions into quickly obtainable, consumable services.

Our focus today is on automation and making our individual private clouds run like a single cloud through a common orchestration layer and control plane. While we will continue to use commercial software when it provides the best solution and fit for our cloud evolution, open source solutions such as OpenStack and Cloud Foundry are proving to be versatile tools for the building of an open, extensible framework for managing and exposing the various resources, such as compute, network, storage, and PaaS in our private cloud. Leveraging our OpenStack-based control plane, we will continue to maximize use of maturing open source capabilities, protect existing infrastructure

investments, and provide a common interface for Intel hosting environments.

Moving forward, we plan to take the following steps:

- Require suppliers to provide OpenStack APIs in their solutions.
- Deploy cost-optimized infrastructure through organic service growth, refresh, and capability maturity.
- Use PaaS to enable fast application deployment and host cloud application design patterns.
- Continue our evolution to hybrid solutions, increasing our ability to easily consume private and public cloud services.

As the open cloud industry continues to grow, we plan to place greater emphasis on selecting suppliers that enable open standards and solutions. We also plan to continue to validate and refine our cloud model, seeking to implement operational efficiencies that further improve our ability to cost-effectively advance to hybrid cloud deployments in the future.

RELATED INFORMATION

Visit www.intel.com/IT to find content on related topics:

- "Accelerating Deployment of Cloud Services Using Open Source Software"
- "Adopting Software-Defined Networking in the Enterprise"
- "An Enterprise Private Cloud Architecture and Implementation Roadmap"
- "Best Practice for Building an Enterprise Public Cloud"
- "Cloud Computing Cost: Saving with a Hybrid Model"
- "Enhancing Cloud Security Using Data Anonymization"
- "Extending Intel's Enterprise Private Cloud with Platform as a Service"
- "Maximizing Cloud Advantages Through Cloud-Aware Applications

ACRONYMS

BKM	best-known method
CLI	command line interface
DBaaS	database as a service
IaaS	infrastructure as a service
MaaS	metal as a service
PaaS	platform as a service
SaaS	software as a service
SDN	software-defined networking
VM	virtual machine

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