

Compliments of **IBM**

IBM Limited Edition

# Hybrid Cloud

FOR  
**DUMMIES**<sup>®</sup>  
A Wiley Brand

## *Learn:*

- Fundamentals of hybrid cloud
- How to manage visibility, control, and security
- How to handle deployment in hybrid cloud environments
- How to recognize business impact of hybrid cloud

**Judith Hurwitz  
Marcia Kaufman  
Daniel Kirsch**





***Hybrid Cloud***

FOR  
**DUMMIES<sup>®</sup>**  
A Wiley Brand

***IBM Limited Edition***

**by Judith Hurwitz, Marcia  
Kaufman, Daniel Kirsch**

FOR  
**DUMMIES<sup>®</sup>**  
A Wiley Brand

## Hybrid Cloud For Dummies®, IBM Limited Edition

Published by:  
**John Wiley & Sons, Inc.**  
111 River St. Hoboken,  
NJ 07030-5774  
[www.wiley.com](http://www.wiley.com)

Copyright © 2015 by John Wiley & Sons, Inc.

No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning or otherwise, except as permitted under Sections 107 or 108 of the 1976 United States Copyright Act, without the prior written permission of the Publisher. Requests to the Publisher for permission should be addressed to the Permissions Department, John Wiley & Sons, Inc., 111 River Street, Hoboken, NJ 07030, (201) 748-6011, fax (201) 748-6008, or online at <http://www.wiley.com/go/permissions>.

**Trademarks:** Wiley, For Dummies, the Dummies Man logo, The Dummies Way, Dummies.com, Making Everything Easier, and related trade dress are trademarks or registered trademarks of John Wiley & Sons, Inc., and/or its affiliates in the United States and other countries, and may not be used without written permission. IBM and the IBM logo are registered trademarks of International Business Machines Corporation. All other trademarks are the property of their respective owners. John Wiley & Sons, Inc., is not associated with any product or vendor mentioned in this book.

**LIMIT OF LIABILITY/DISCLAIMER OF WARRANTY: THE PUBLISHER AND THE AUTHOR MAKE NO REPRESENTATIONS OR WARRANTIES WITH RESPECT TO THE ACCURACY OR COMPLETENESS OF THE CONTENTS OF THIS WORK AND SPECIFICALLY DISCLAIM ALL WARRANTIES, INCLUDING WITHOUT LIMITATION WARRANTIES OF FITNESS FOR A PARTICULAR PURPOSE. NO WARRANTY MAY BE CREATED OR EXTENDED BY SALES OR PROMOTIONAL MATERIALS. THE ADVICE AND STRATEGIES CONTAINED HEREIN MAY NOT BE SUITABLE FOR EVERY SITUATION. THIS WORK IS SOLD WITH THE UNDERSTANDING THAT THE PUBLISHER IS NOT ENGAGED IN RENDERING LEGAL, ACCOUNTING, OR OTHER PROFESSIONAL SERVICES. IF PROFESSIONAL ASSISTANCE IS REQUIRED, THE SERVICES OF A COMPETENT PROFESSIONAL PERSON SHOULD BE SOUGHT. NEITHER THE PUBLISHER NOR THE AUTHOR SHALL BE LIABLE FOR DAMAGES ARISING HEREFROM. THE FACT THAT AN ORGANIZATION OR WEBSITE IS REFERRED TO IN THIS WORK AS A CITATION AND/OR A POTENTIAL SOURCE OF FURTHER INFORMATION DOES NOT MEAN THAT THE AUTHOR OR THE PUBLISHER ENDORSES THE INFORMATION THE ORGANIZATION OR WEBSITE MAY PROVIDE OR RECOMMENDATIONS IT MAY MAKE. FURTHER, READERS SHOULD BE AWARE THAT INTERNET WEBSITES LISTED IN THIS WORK MAY HAVE CHANGED OR DISAPPEARED BETWEEN WHEN THIS WORK WAS WRITTEN AND WHEN IT IS READ.**

For general information on our other products and services, or how to create a custom *For Dummies* book for your business or organization, please contact our Business Development Department in the U.S. at 877-409-4177, contact [info@dummies.biz](mailto:info@dummies.biz), or visit [www.wiley.com/go/custompub](http://www.wiley.com/go/custompub). For information about licensing the *For Dummies* brand for products or services, contact [BrandedRights&Licenses@Wiley.com](mailto:BrandedRights&Licenses@Wiley.com).

ISBN: 978-1-119-16886-7 (pbk); ISBN: 978-1-119-16887-4 (ebk)

Manufactured in the United States of America

10 9 8 7 6 5 4 3 2 1

## Publisher's Acknowledgments

Some of the people who helped bring this book to market include the following:

**Project Editor:** Carrie A. Johnson

**Editorial Manager:** Rev Mengle

**Acquisitions Editor:** Steve Hayes

**Business Development Representative:**  
Sue Blessing

# Table of Contents

<b>Introduction</b> .....	<b>1</b>
About This Book .....	1
Foolish Assumptions .....	2
Icons Used in This Book.....	2
<b>Chapter 1: Discovering the Fundamentals of Your Hybrid Cloud</b> .....	<b>3</b>
Defining Hybrid Cloud .....	4
Connection Points to Support a Hybrid Cloud.....	5
The Hybrid Cloud Imperative.....	10
Looking at Different Cloud Delivery Models.....	11
<b>Chapter 2: Managing Visibility, Control, and Security</b> .....	<b>17</b>
Gaining Visibility .....	17
Self-Service Provisioning.....	18
Managing Workloads to the Right Platform .....	19
Using Automation and Service Delivery to Operationalize the Hybrid Cloud.....	22
Securing and Governing Hybrid Clouds.....	23
Managing Workloads Dynamically.....	25
<b>Chapter 3: Application Development and Deployment in Hybrid Cloud Environments</b> .....	<b>27</b>
Improving Developer Productivity .....	28
Defining DevOps.....	29
Optimizing DevOps for the Hybrid Cloud.....	30
Enabling Delivery Automation .....	33
Platform as a Service (PaaS) .....	34
Application Release Automation (ARA) and testing services.....	34
Creating Continuous Delivery .....	35
DevOps frameworks .....	35
Reusable web services .....	36
Monitoring and Analytics.....	37
Looking at the Impact of Open Standards for the Hybrid Cloud.....	38

**Chapter 4: Data and Analytics . . . . . 41**

Finding the Value of Managing Data and Analytics in a Hybrid Cloud .....	42
Locating relevant data sources.....	43
Refining and enriching data.....	45
Loading, storing, and analyzing data .....	47
Data services to replicate, synchronize, and integrate data.....	49
Building, testing, and running advanced analytical models .....	49
Using Analytics to Deliver Business Results .....	50
Predictive maintenance .....	51
Fraud .....	52
Marketing and customer churn .....	53
Looking at Cognitive Computing.....	53

**Chapter 5: Data and Application  
Integration and Portability . . . . . 55**

The Imperative for Integration and Portability.....	55
Managing Data in the Hybrid Environment .....	56
Managing integration and data portability.....	57
The need for portability.....	57
Container technology supports portability .....	58
Developing standards to support workload portability.....	58
Managing Multiple Data Sources within a Hybrid Environment .....	59
Managing metadata .....	59
Challenges in integrating data.....	59
Managing security and governance.....	60
Replicating and synchronizing data.....	61
Integrating Applications across Environments and Gaining Portability.....	61

**Chapter 6: Recognizing the Business  
Impact of Hybrid Cloud . . . . . 63**

Understanding Why Hybrid Matters .....	64
Balancing Costs with Value .....	64
Getting to the Dial Tone.....	65
Looking at the Value of Flexibility in a Fast-Changing World.....	66
Setting the Stage for Action .....	66

# Introduction

---

**W**elcome to *Hybrid Cloud For Dummies*, IBM Limited Edition. Hybrid clouds are becoming the primary and pragmatic approach to leveraging a variety of cloud services that include public, private, and managed services. Increasingly businesses are thinking about the cloud as the way they can implement changes in both their technology platforms and their approach to business. Changing the way business services are delivered to customers requires that technology be available to facilitate that change.

The hybrid cloud means choice. You have the capability to use the right type of cloud service to support the right workloads at the right time. The benefit of the hybrid cloud is that it can be used to provide a set of well-managed services. These services include the capability to use self-service infrastructure services that can provide the scalability that you need for unanticipated opportunities. In addition, a hybrid cloud provides a continuous delivery platform for building, deploying, and managing applications.

Through the adoption of emerging technologies, such as OpenStack and containers, you can achieve portability of both data and applications. With the use of sophisticated cloud applications, you gain the right services to efficiently and effectively deploy the applications your users and customers need.

## *About This Book*

This book gives you insights into what it means to use a hybrid approach to cloud computing and the flexibility and choice that it provides. You will discover the type of technologies and approaches that are imperative for your organization to create the business value and flexibility the constituents you serve demand. This book helps business and technical leaders understand the hybrid cloud in context with the needs of the business.

## *Foolish Assumptions*

The information in this book is useful to many people, but we admit that we made a few assumptions about whom we think you are:

- ✔ You're already using various forms of cloud computing and are looking for a systematic way to help your company leverage the services you need across many different deployment models.
- ✔ You're planning a long-term strategy to move to a standards and architectural approach to the hybrid cloud that protects your investment and allows you to gain the capability to move workloads to the most cost-effective model.
- ✔ You understand the huge potential value of the data that exists throughout your organization.
- ✔ You understand the benefits of a flexible approach to DevOps based on continuous development and deployment of new applications that give your organization a competitive edge as business models change.
- ✔ You're a business leader who wants IT resources to be available like a utility that's optimized to leverage what you've already paid for so you can invest in areas that differentiate your business.

## *Icons Used in This Book*

The following icons are used to point out important information throughout the book:



Tips help identify information that needs special attention.



Pay attention to these common pitfalls of managing your foundational cloud.



This icon highlights important information that you should remember.



This icon contains tidbits for the more technically inclined.



# Chapter 1

---

# Discovering the Fundamentals of Your Hybrid Cloud

.....

## *In This Chapter*

- ▶ Defining hybrid cloud
  - ▶ Looking at the points to support a hybrid cloud
  - ▶ Understanding the business imperative for the hybrid cloud
  - ▶ Introducing cloud delivery models
- .....

**T**he world of computing has entered the next stage of a revolution. Only a few years ago, businesses were debating whether to remain with a data center or move to a public or private cloud. Today, companies are realizing that all these models of computing are becoming the foundation of the next generation platform. Therefore, the hybrid cloud is becoming the architectural framework that allows companies to use whichever deployment model best serves their business needs. Flexibility in the hybrid computing model comes from the capability to change deployment models whenever the business needs to change.

Business disruption is at the heart of the requirement for the hybrid cloud. Across all industries new competitors are emerging that are able to leverage technologies to move faster and be more nimble than larger, well established companies. Who could have ever imagined five years ago that the taxi industry would be upended by emerging vendors like Uber, Lyft, and SideCar? Who would have predicted that born-on-the-web companies like Airbnb and HomeAway would challenge

the long established hospitality industry? There are hundreds of examples of how traditional businesses are now required to transform their business models to compete with emerging born-on-the-web players that have no legacy systems and no on-premises environment.

In order to prepare for increased competition, organizations are rethinking many traditional ways of delivering services to customers and partners. These organizations need an IT environment that's optimized for speed, security, and flexibility. Predictability and consistency across workloads is also a priority for businesses. Hybrid cloud offers a unified environment across a combination of deployment models that helps companies become more nimble and proactive.

In this chapter, we define hybrid cloud and describe the key connection points for the hybrid cloud. The chapter explains the technology changes that are enabling this transformation, including the emergence standards that support modularity, including OpenStack services and container technologies.

## Defining Hybrid Cloud

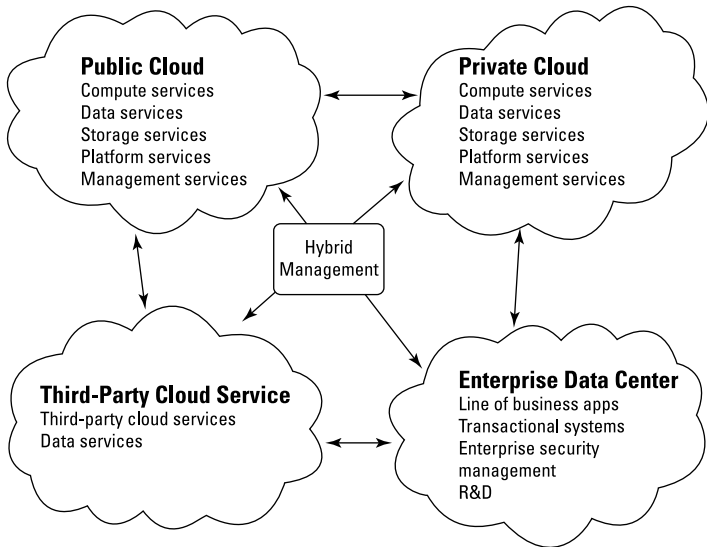
A *hybrid cloud* is an environment that integrates traditional IT with a combination of public, private, or managed cloud services. In essence, a hybrid cloud becomes a virtual computing environment that may combine services in a public cloud with services from a combination of environments to deliver the right service level and flexibility to meet emerging customer requirements. All these services need to be managed as though they were designed to behave as a single unified environment.



Users are more concerned with the quality and availability of services than the exact components of the supporting infrastructure. Increasingly, the actual delivery model of services is becoming invisible to the end customer. Users need to know that they can access the right service at the right time with the right service level. This hybrid cloud defines your computing environment and is used to run your business and satisfy the needs of your constituents.

Although each of these elements may be designed and managed by a variety of developers, they can be managed as a single system through standard APIs. A hybrid cloud can be

highly effective because as a distributed system, it can enable companies to leverage a series of services that are the best fit for the task at hand, as shown in Figure 1-1.



**Figure 1-1:** The hybrid cloud architecture.

## Connection Points to Support a Hybrid Cloud

In a hybrid cloud environment, everything needs to be connected so the end customer has a seamless experience. It's not enough simply to provide cloud services. Integrating best practices combined with new innovative technologies is a key to a well executed hybrid cloud. There are a number of enabling technologies that have to be put in place in order to enable a hybrid cloud to operate to support changing business requirements. The six key connection points include the following capabilities:

- ✓ Integration
- ✓ Data localization
- ✓ Operational visibility and management

- ✓ Security services
- ✓ DevOps
- ✓ Portability

## *Integration services*

The new generation of dynamic and customer focused applications requires integration on many levels. What services do you need to integrate data and process in your hybrid environment? You need to integrate data across your own company and with partner data and public data sources. For example, social media data (such as Twitter or Facebook) and third-party data (The Weather Company or U.S. Census Bureau) may need to be combined with internal data sources. In addition, incorporating managed data services like pricing services or credit checking services can provide consistency across applications and help companies to meet customer requirements for speed of delivery. Integration of data and processes needs to be seamless across mobile deployments, public and private clouds, and traditional data center environments. Having the ability to analyze this data in real or near-real time can improve insight about customers and grow the business.

## *Data localization services*

It isn't enough to simply provide techniques for integrating data across platforms. Some companies want data to be managed locally because of performance and manageability requirements. In some cases, there isn't a requirement to broadly share the data. In other situations, security and compliance requirements dictate the need for data localization. In addition, various European countries have strict governance rules defining where private data can be sent. For example, a company is required to indicate not only where data is stored but also which countries that data has passed through.

## *Operational visibility and management*

A customer that's leveraging services has no interest in what happens behind the scenes. Customers who want the hybrid

cloud environment really want the equivalent of a dial tone from their vendors. The customer wants to know that services will operate consistently and predictably. Therefore, a hybrid cloud can't work without the capability to manage all the elements of the environment as though they were a single unified environment.

Managing the variety of services that are part of a hybrid cloud requires you to have operational visibility to make sure that all processes, data, and services are working as designed. Accomplishing this task requires a platform that's able to discover the underlying services and infrastructure. It must be able to monitor those services so there's clear visibility so the entire environment can be controlled in a predictable manner. This means being able to monitor and control not just internal services but those services that are owned and controlled by third-party public cloud vendors or managed services providers.

## *Workload management*

Well-designed management tools need to be in place to ensure that the hybrid environment works well to support the business objectives. One of the fundamental differences between cloud computing and traditional computing is the way a cloud is designed to manage workload resources. Managing workloads is foundational to the cloud. Whereas a data center is designed to manage applications, a cloud is intended to manage a pool of resources, which is precisely what it sounds like — a set of shared, configured services that are independent of a physical location. In many situations, cloud service providers create a multi-tenant environment to support the deployment of these resources. Multi-tenancy enables the sharing of a service while keeping the data and configurations of individual customers separate.

For example, say that you're a cloud provider. You don't want your customers to have to make granular selections — for example, needing to select a specific server or storage system. Instead, you want to offer customers simplification and abstraction so they can simply say, "I need some more storage." Those storage resources are then pooled together from various physical systems to create a set of resources. In other words, customers don't need to know which storage system they're accessing.

## *Managing security*

In a hybrid world, security needs to be managed in conjunction with visibility and control. It also must be managed based on an open governance model. Therefore, a comprehensive approach to security is required to ensure that all the interaction points in a hybrid environment are tightly controlled.

Policies and business rules are needed to understand and monitor information about your workloads. For example, it will be important to determine where a workload needs to be located for either performance or compliance requirements. There's also a requirement to make sure that only authorized individuals are allowed to access and change data. Security is an essential component of the unified architecture for a hybrid cloud. If you have many different services, you need a consistent and predictable way to manage security even if your cloud environment changes.



The increase in connection points adds a multitude of vulnerabilities to the applications and overall IT environment. It is hard for many companies to keep abreast of changing security risks. Even the most experienced security expert is always playing catch up with the last security threats. One way to add security protection is by using a cloud-based development platform with built-in security services. When the hybrid cloud infrastructure is designed with a built-in security infrastructure, a less experienced professional will be able to safely manage the environment.

## *DevOps: Continuous Development and Operations in a Hybrid Cloud*

As IT focuses on creating a more efficient application life cycle, new approaches are required. Combining the development and operations of next generation applications is a top priority. Often called DevOps, this new dynamic development model needs to be optimized for the hybrid cloud. One of the benefits of the hybrid environment is to help companies quickly compose modern applications and access the data needed to support these applications. Consistency of outcomes becomes a priority to support the needs of the customer.

Increasingly, companies are adopting a Platform as a Service (PaaS) in order to develop and operate applications in a consistent and predictable way. Most importantly, organizations need a single unified way to manage and synchronize applications and data across public and private clouds. For example, you need a single federated catalog and operations console for public and private PaaS.

## *Portability and modularity*

As your organization demands the flexibility to change based on new business models, you'll need to move to an environment that supports portability and modularity. This approach allows application code to be easily moved from one environment to another. With emerging standards and open technologies, you become less dependent on one implementation of a cloud and gain interoperability no matter which vendor's services you use.

Two foundational services that are instrumental to establishing portability and modularity are

- ✔ **Containerization technology:** Containers are a technique to package an application or service, including its dependencies, so it can run in any Linux environment. In essence, a container creates an encapsulation for a set of logic and data. With this approach, it's possible to build microservices. A *microservice* is a well-defined service that can be used to execute a service in many different environments. A group of microservices can be combined to create a new application service. Because the microservice uses the same container encapsulation services, portability is ensured.
- ✔ **OpenStack environment:** OpenStack is emerging as another critical element in creating portability within a cloud environment. OpenStack is an open-source cloud operating system. More than 500 companies have joined the project, and a number of services have been built, tested, and deployed by many of the participants in the consortium. These services include networking, storage, identity management, orchestration, compute, and database. Companies that adhere to the OpenStack agreed on specifications can move workloads across a variety of clouds that support OpenStack services.

## The importance of open technologies

As a hybrid cloud model becomes ubiquitous, it's critical that customers have the confidence and assurance that the services they adopt are built on open technologies. These de facto or de jure standards are key to being able to provide interoperability and flexibility. Open technologies such as Linux, OpenStack, Hadoop,

Spark, and Docker are supported by many vendors. These open technologies also gain from the insights provided by sophisticated developers. The community approach means that innovation will continue to improve to support emerging customer requirements.

## *The Hybrid Cloud Imperative*

The hybrid cloud model provides for the combination of systems of engagement (like mobile) with systems of record. For example, one company may use its data center to manage customer transactions. Those transactions are then connected with a public cloud where the company has created a web-based, front-end, and a mobile interface to allow customers to buy products online. The same company may use a third-party managed service that checks credit for anyone paying on an installment plan. There may also be a series of public cloud-based applications that control customer service details. In addition, the company may need to use extra compute capabilities from a public cloud provider during peak holiday periods. In addition, many companies will use a public cloud to experiment with prototype business models.

The hybrid cloud is becoming the optimum approach to computing for companies that need to quickly change processes along with their requirements for compute, networking, and storage. For example, you may want to create a new mobile application that satisfies customer demands for access from any device. Customers want their suppliers to allow them to interact in their preferred deployment model. This point is especially important because emerging competitors typically provide a level of intimacy and accessibility to their solutions.





When hit with a competitive situation, suddenly you have to be able to demonstrate that your company can react quickly to retain market leadership. This means supporting the right deployment model while maintaining access to the right data and the right workloads. Ultimately, forward-looking companies think of their computing platform as a combination of services that can be brought together to support changing customer requirements. Companies see huge opportunity in mobile, big data, and social computing, but at the same time, they struggle with the data management challenges inherent in these new workloads. The data needed to support these workloads typically resides on a variety of different systems — ranging from internal IT systems to the public or private cloud.



Certainly some companies eager to get to market operate totally on a public cloud. However, long-term business and technical issues may emerge for 100 percent on the cloud companies. For example, as a company grows, the public cloud may not provide sufficient security and governance to protect customer data. In addition, most companies have systems that manage customer transactions and data about interactions. These systems are often best managed inside the data center to protect intellectual property and to manage performance. Costs are also a concern. The costs of public cloud services can add up quickly and will be a recurring expense that negatively impacts profits. You need to break down the barriers between your on-premises and cloud-based systems in order to gain the insights you need to service customers and grow your business.

## *Looking at Different Cloud Delivery Models*

Each of the cloud deployment models within a hybrid environment features different service models. Within cloud computing models, a number of different models and approaches exist that are optimized for executing specific tasks for specific workloads. These models are divided into the four areas covered in this section. They each provide a different capability that's important to implementing sophisticated services.

## *Infrastructure*

As its name implies, Infrastructure as a Service (IaaS) is the foundational cloud service. IaaS provisions compute, storage, and networking services through either a virtualized image or directly on the computer systems. In virtualization, a user accesses an image that contains the services needed to run a workload. When companies need maximum speed, they often use a native or bare metal implementation. A public IaaS is designed as a self-service environment so a customer can purchase a service such as compute or storage based on the instance of computing that's needed. Consumers can purchase an instance based on the amount of resources consumed over a specified period of time. When a consumer stops paying for the service, the resource disappears. In a private IaaS environment controlled directly by a company, those provisioned resources remain in place and are controlled by the IT organization.

The emergence of a Software-Defined Environment (SDE) provides a next-generation approach to IaaS and other cloud services. The goal of IaaS is to optimize the use of system resources so they can support workloads and applications with the maximum efficiency. An SDE is an abstraction layer that unifies the components of virtualization in IaaS so the components can be managed in a unified fashion. In effect, the SDE's intention is to provide an overall orchestration and management environment for the variety of resources used within an IaaS environment. Therefore, an SDE brings together compute, storage, and networking to create a more efficient hybrid cloud environment. It also enables developers to use a variety of types of virtualization within the same environment without the burden of hand-coding the linkages between these services.

## *Applications*

Today, most packaged software applications are available as Software as a Service (SaaS). SaaS is a defined application that's operated on a cloud service. Although these applications are delivered on a public cloud service, a trend exists in many enterprises to follow a similar approach with applications hosted on their private clouds. In other words, business users benefit from the ease of use and fast delivery

of applications delivered on the public cloud. IT can use its private cloud to host and deliver internal applications to meet the needs of its internal business users.

These applications are built to take advantage of IaaS. Therefore, like IaaS, SaaS is typically delivered in a multi-tenancy environment, offering load balancing and self-service provisioning. This means that multiple users share a physical computing environment with other users and companies. Each user's implementation is partitioned separately from other users.



One of the benefits of leveraging applications delivered via the cloud is that the consumer isn't responsible for software updates and maintenance of the application. However, unlike a traditional on-premises application, the user doesn't have a perpetual license for the application. Instead, the user pays on a per-user, per-month, or per-year basis.

Many SaaS applications are designed as packaged applications based on a business process such as customer relationship management or accounting. These applications are designed in a modular fashion so customers can select only what they require. For example, some accounting SaaS applications may have a foundation of a bookkeeping process and can expand into a complex online accounting system. Over the years, more and more areas of software have become available as a service, including the following:

- ✓ Collaboration
- ✓ Project management
- ✓ Marketing social media services
- ✓ Risk management
- ✓ Commerce solutions

SaaS implementations are expanding beyond the traditional packaged software. Increasingly, most emerging software platforms are implemented using cloud services as the preferred deployment model. Business users have recognized the benefits of leveraging applications as a service. For example, business users leveraging analytics as a service find that it's more affordable for solving complex problems. An analyst may need to build a predictive model to solve a specific problem in a

quick time frame. Instead of purchasing all the hardware and software, the analyst can leverage a sophisticated analytics application in the cloud. The analyst pays only for the capability used for that project. After the project is complete, no further financial obligation is required. The cloud offers the capability to solve a problem that leverages huge amounts of computing capability. There may also be the need to store the data and results from this analysis.

## *Software development and deployment*

Software development and deployment services are used to design, implement, and deploy applications and services in either a public or private cloud. Platform as a Service (PaaS) provides an underlying level of middleware services that abstract the complexity away from the developer. In addition, the PaaS environment provides a set of integrated software development tools. In some cases, it's possible to integrate third-party tools into the platform. A well-designed PaaS consists of an orchestrated platform to support the life cycle of both developing and deploying software within the cloud. A PaaS platform is designed to build, manage, and run applications in the cloud.

Unlike traditional software development and deployment environments, the software elements are designed to work together through Application Programming Interfaces (APIs) that support a variety of programming languages and tools. Within the PaaS environment are a set of prebuilt services such as source code management, deployment of workloads, security services, and various database services.

## *Business process*



One of the benefits of a hybrid cloud environment is the capability to link business processes from a variety of services and systems in order to satisfy a customer need. Flexibility is key for companies that want to change a process when the business changes. There are several different ways that business process is handled in a hybrid cloud environment. A number of business processes are mature and not likely to change quickly.

For example, some services include the detailed steps for paying employee salaries, including reporting, taxes, and so on. Likewise, well-honed processes track employee vacations based on company policy. However, other situations arise where business processes need to be able to be changed on the fly. For example, your company has discovered that an emerging competitor has come up with a novel approach to selling a product. If business processes can be modified quickly, your company can regain leadership.



Often customers demand that their existing suppliers change their business processes based on what they see from emerging vendors in the market.

To make the hybrid cloud concept work for your organization, you need a modular service architecture that provides for self-service, scalability, and flexible workload management. This level of flexibility demands that applications and business services be written and managed differently from traditional application architectures. Cloud services demand that each service be designed as a task without dependencies. Each service needs a standards-based API that allows for easier integration. Well-defined APIs enable services to be dynamically linked together to execute tasks. Through these APIs, modular services can link public cloud-based services such as vertical applications with traditional data center applications and services. This level of modularity helps companies build new value that was impossible with integrated code and applications.



## Chapter 2

---

# Managing Visibility, Control, and Security

.....

### *In This Chapter*

- ▶ Making sense of self-service provisioning
  - ▶ Understanding workload placement
  - ▶ Automating service delivery
  - ▶ Managing security and governance in a hybrid cloud environment
- .....

**A**s cloud computing gains more traction, organizations will begin understanding that their computing environment is a combination of a variety of resources that come together to create a unified computing environment. Customers and the users of these services won't think about the fact that some workloads live in the public cloud while others reside in the private cloud or the data center. In fact, the combination of all services becomes the definition of their hybrid computing environments. To make these disparate systems operate as one seamless environment, it's vital to have visibility, control, and security for each element and the combined elements. Achieving this vision requires work. However, many organizations are beginning to plan for a world where they can use all the elements in a cohesive environment. In this chapter, we discuss the requirements for hybrid cloud management.

## *Gaining Visibility*

When developing a hybrid cloud environment, one of the most important aspects is the ability to gain visibility and insight into the environment. In a hybrid cloud environment, an

organization leverages services from a number of different resources, including on-premises, SaaS vendors, IaaS, and possibly PaaS vendors. Therefore, administrators need a single dashboard that can give them insight into what resources are being used and how services are related to each other. Over time, these dashboards will gain more sophistication. For example, they're beginning to integrate advanced analytics, which help predict future capacity and budgeting needs.

## *Self-Service Provisioning*

One of the fundamental requirements for any type of cloud service is that customers can automatically provision (via a self-service portal) the type of compute, storage, and other services required. A simple example of self-provisioning that you may be familiar with is the banking ATM. Without the availability of the self-service ATM, banks would be required to use costly resources to manage activities of all their customers — even for the most common and repetitive tasks like a simple withdrawal or deposit. With an ATM, the customer makes a direct request to perform routine transactions that conform to predefined business rules.

For example, a customer must have an account to withdraw money. In addition, the customer can't take out more money than is in her account. There may be rules dictating how much money an individual can withdraw from the ATM at one time or in one day. This process is precisely how self-service works in a hybrid cloud environment. In a hybrid cloud environment, administrators can enable users to provision pre-defined resources based on a set of business rules. In this way, administrators and users are both satisfied. Users are given the freedom to accomplish their jobs with little red tape while administrators are able to control which users and workloads are using what resources.

When used in a controlled way, hybrid cloud environments allow administrators to place workloads and data in the right environment that's most appropriate for the required service level objectives, security, cost, and compliance. In addition, most organizations maintain a requirement to gain visibility into what cloud services employees and business units are using. This understanding of what resources an employee has access to is critical if an employee changes roles or



is terminated — you don't want an ex-employee having access to corporate intellectual property after he leaves the company.

## Managing Workloads to the Right Platform



At its core, a *workload* is an independent service or collection of code that can be executed. The application, operating system, and middleware can be considered workloads. From a management perspective, in order to meet the service level objective, workloads need to be executed based on

- ✔ Criticality
- ✔ Level of security
- ✔ The type of interaction with other environments

Of course, different workloads have different characteristics, and the platform for them to run in an optimized manner will differ. For example, an analytics workload may require compute-intensive statistical processing with complex queries. However, the analytics workload may only require resources once a week or once a quarter when analysis is conducted. Meanwhile, a transactional workload may necessitate round-the-clock availability and high speed but only require a small amount of compute power. Matching workloads with the right environment allows organizations to optimize workload efficiency while also only budgeting for the resources that are required.

Here's a list that explains some of the kinds of workloads you may find in a hybrid cloud environment and the compute requirements. Those workloads include the following:

- ✔ **Batch workloads:** These workloads typically operate in the background and conduct repetitive actions on large volumes of data. Examples of batch workloads include processing customer bills and account statements, archiving documents, or processing accounts payable and receivable. These workloads require considerable compute and storage resources; however, they aren't

time sensitive. Because the data in batch workloads is well documented and predictable, automating this type of workload is relatively easy. Generally, batch workloads are executed regularly when compute resources aren't needed for other business tasks. For example, batch workloads often occur overnight or on the weekend. Again, as in any cloud environment, the decision of where to execute batch workloads is determined by business rules, governance, and security regulations.

- ✔ **Analytic workloads:** Most organizations have grown their analytic workloads in terms of volume, complexity, and importance. In the past, many analytics workloads were treated like batch workloads — run in the back-office during off hours. Today, advanced analytics is driving critical business decisions and helping to deliver predictive recommendations at the point of decision. In addition, analytic workloads often require near real-time processing. Analytic workloads can be compute, memory, and storage intensive at unpredictable times. For example, heavy analysis may be performed for weeks or months when testing a new product but then ramped down after the product is released. Because of the unpredictable nature of analytics workloads, many organizations are looking to move these workloads to public clouds. By leveraging a public cloud for analytics workloads, an organization can provision the large resources that are required for short periods of time and then de-provision them when the project is completed.
- ✔ **Transactional workloads:** These workloads are the automation of business processes such as product lookups and order processing. Traditionally, transactional workloads were placed on a single system. However, with newer business models like e-commerce where transactions can reach across partners and suppliers, transactional workloads have become more complex. So, there's a need to focus on business processes with these transactional workloads. These workloads might consist of simple queries but require near real-time response rates. Depending on the cost-benefit analysis, it's likely that complex transactional workloads are best suited to a private cloud. In addition, transactional workloads often include customer Personally Identifiable Information (PII), credit card data, or other pieces of information that may need to be stored on-premises.

✓ **Database workloads:** A database workload must be tuned and managed to support the application that's using that data. Database workloads tend to require a high amount of Input/Out (I/O) cycles. In some situations, data workloads are small and self-contained; however, in other situations, the data workloads are huge, and the performance requires a sophisticated approach. For example, high-performance database workloads may be implemented on bare metal (directly on the hardware's operating system) to support the business requirement. Some IaaS vendors are offering customers access to bare metal implementations to support these workloads and have more control over the infrastructure.

## Increasing flexibility with cloud management and IaaS bare metal

A large multi-national conglomerate has business units in a variety of industries, including energy, health-care, and electronics. The company as a whole is focused on innovation and giving each business unit the tools it needs to advance. A major challenge for the company is that each user group has its own cloud infrastructure needs. If each group independently provisions resources, the company will have a series of fragmented environments with no ability to monitor costs, security, and governance. Therefore, the company needs a controlled way to provide the resources that all user groups require to achieve innovation.

The company worked with IBM to deploy IBM Cloud Orchestrator. The software gives the company an open framework for managing its heterogeneous cloud environment.

Hybrid cloud administrators are able to monitor where applications are running, who has access to what resources, and the costs of the environment while users are given the tools to accomplish their jobs with little administrative friction.

In addition, the company implemented SoftLayer for public cloud services. Two of the most important reasons why the company selected SoftLayer were because it is built on open-standard technology and the fact that SoftLayer offers bare metal servers.

Because the company runs a wide variety of workloads, access to bare metal is important. Although many workloads run very well in virtualized private or public clouds, there are still workloads that don't work well when virtualized. Speed and overall

*(continued)*

*(continued)*

performance were the overriding requirements for this company. For example, most workloads with high I/O requirements, such as databases, are poor candidates for virtualized environments. Applications that require high levels of hardware acceleration are also often difficult to run in a virtualized environment.

In addition, the company has some applications that require real-time responses, and delays of even a second or two could be very costly. Running these latency sensitive applications on bare metal allows them to eliminate the overhead of a hypervisor to improve overall performance.

## *Using Automation and Service Delivery to Operationalize the Hybrid Cloud*

What's the connection between workloads and workload management in the cloud? In a hybrid cloud environment, it's critical to have a technique for linking workloads together in a way that delivers on customer expectations. This is true whether you're a service provider offering either a public or private cloud to customers, or you're managing an internal private cloud to benefit internal customers and external customers and partners.

You may think that all you have to do is get some automation software (to automatically schedule resources and to perform some other functions associated with allocating resources) and you're set. However, when you look at workload management from an operational perspective, many other issues should be taken into account. Primarily, you have to determine how the collection of services will work together in terms of performance and quality of service. Depending on the use of the hybrid cloud, you have to address a variety of security and compliance requirements. Security is always a critical issue for companies. However, in a complex hybrid environment, it's important that all elements that make up the hybrid environment have the same consistency of security. In many instances, organizations must be able to demonstrate their compliance with industry rules. For example, in health-

care, when patient information is involved, an organization must be able to show that it's Health Insurance Portability and Accountability Act (HIPAA) compliant. In addition, there may be country-specific rules governing where customer and employee data must reside.

Other situations may have fewer restrictions because of the nature of the workloads and governing rules. In these cases, IT operations is free to move workloads to locations that make the most fiscal sense and have the bandwidth or capacity to meet the quality of service required by the business. In fact, the capability to change and move workloads based on business requirements is the heart of operational issues in the cloud.



Even when an organization looks holistically at managing a variety of workloads, change will be the norm. A hybrid cloud environment isn't static. From an operational perspective, organizations need to be prepared for the addition of new workloads that will change the way the hybrid environment operates. It is critical to be able to dynamically rebalance the workload and to ensure that security requirements continue to be met.

## Securing and Governing Hybrid Clouds

A company planning to secure its IT environment generally focuses on a broad range of vulnerabilities to its data center as well as ways to safeguard sensitive corporate, customer, and partner information wherever it's located. A hybrid cloud environment changes things because although ultimately it's your company's responsibility to protect and secure your applications and information, many challenges arise when you're working with an external provider.



Here are a few of those challenges:

- **Multi-tenancy:** In a multi-tenant architecture, a software application partitions its data and configuration so each customer has a customized virtual application instance. However, your applications and data exist on the same servers as other companies using the same service

provider, and these users are accessing their resources simultaneously. You may not know the names of the other companies that are sharing these servers. So, if one company's data or application is breached or fails, your application may be affected.

- ✔ **Attacks that affect you, even though you aren't the target:** If your company makes use of a public cloud, you may be the collateral damage in an attack — even if it wasn't meant for you. Consider a virus or denial of service (DoS) attack, for example. Because you're sharing an environment with others, even though you may not be a target, your resources may be affected, resulting in a service interruption, customer requests not being fulfilled, or a possible security breach.
- ✔ **Incident response:** In a cloud environment, you may not have control over how quickly incidents are handled. For example, some cloud providers may not tell you about a security incident until they've confirmed that an actual incident occurred. As a result, you won't know something has happened until it affects your business. Additionally, if you become aware of an incident, you may not have access to servers to perform an analysis of what went wrong.
- ✔ **Visibility:** The example in the preceding bullet illustrates that in many cloud environments you may not be able to see what your provider is doing. In other words, you may not have control over your visibility into your resources that run in the cloud. This situation is especially troublesome if you need to ensure that your provider is following compliance regulations or laws.
- ✔ **Non-vetted employees:** Although your company may perform extensive background checks on all your employees, you're now trusting that no malicious insiders work at your cloud provider. This concern is real because more than 50 percent of security breaches are caused by insiders (or by people getting help from insiders). If your company is going to use a cloud service, you need to have a plan to deal with inside as well as outside threats.
- ✔ **Data issues:** If you're putting your data in the cloud, you need to be concerned about a number of issues, including the following:

- Making sure no unauthorized person can access this data
- Understanding how this data will be segregated from other companies' data in a multi-tenant environment
- Understanding how your data will be destroyed if you terminate your contract
- Understanding where your data will be physically located
- Understanding how your data is treated as it moves from your location to your provider's servers

## *Managing Workloads Dynamically*

The most important challenge for companies is ensuring that it will be possible to create predictability when managing different workloads across a variety of environments. This means that people from across the company work in concert to protect the assets and provide seamless access to resources that users expect. Hybrid cloud management is intended to provide the capability to dynamically change as deployment models change and adjust the environment to new workloads and new services.





## Chapter 3

---

# Application Development and Deployment in Hybrid Cloud Environments

---

### *In This Chapter*

- ▶ Improving developer productivity
  - ▶ Defining DevOps
  - ▶ Optimizing DevOps for the hybrid cloud
  - ▶ Enabling delivery automation
  - ▶ Creating continuous delivery
  - ▶ Monitoring and controlling your organization's applications
  - ▶ Seeing the impact of open standards for the hybrid cloud
- 

**T**he market dynamics for most businesses today are leading to a faster development cycle. Customers have come to expect frequent updates to existing applications and a steady flow of innovative applications to address their needs. Businesses need to get new mobile applications up and running quickly while keeping costs under control. At the same time, because these new applications depend on data that's shared across internal and external systems, the need for security increases as well. Some of the key benefits of hybrid environments are that applications can be developed more quickly, enhanced more often, and deployed at a faster pace. In addition, hybrid cloud environments support organization requirements for increased security and decreased costs by improving developer productivity and the DevOps process.

## *Improving Developer Productivity*

Keeping pace with the needs of the business is hard. The business expects new applications to be exciting, practical, easy to use, and to deliver great value to its customers. End-users of applications delivered via the cloud have come to expect that any concerns related to the applications will be addressed immediately. As a result, the software release frequency has increased dramatically. Software development teams face a huge array of challenges that make it hard to reach the level of innovation and speed needed to support these business goals. Recognizing the issues that can lead to delays in development time and negatively impact software quality is the first step on the path to increasing developer productivity. Organizations need to address the following challenges:

- ✔ Developers on new projects may be forced to wait weeks or months for a new test lab to be provisioned.
- ✔ It takes a collaborative team of developers, IT management, and business users to build excellent applications. Everyone needs to be in sync and have a common understanding of goals and expectations.
- ✔ Collaboration is often made more difficult because teams are distributed across business units or different geographic regions.
- ✔ Software applications may be connected to multiple cloud services. Each cloud service needs to be monitored for changes and updates so these connections remain valid.
- ✔ Software needs to work across multiple platforms and devices.
- ✔ The software development process requires individual components that all need to work together in a service-oriented architecture.

The software development process needs to be improved to address these issues and improve developer productivity.

## Defining DevOps

If you want to support continuous innovation and speed to market in software delivery, you can't think about development and deployment as separate entities. Taking a DevOps approach means that all stakeholders are committed to the continuous delivery of great software. DevOps teams share a common culture focused on creating a smooth transition from design to development to production. In a DevOps environment, operations staff and developers typically use different elements of the same tools and techniques. Following consistent processes, standardizing tools, and increasing delivery automation helps to remove silos that once existed between development and IT operations. An important aspect of the DevOps process is to increase collaboration across all stakeholders to support rapid responses to customer feedback and continuous delivery of enhancements.

Remember, the end-user doesn't care how the application is developed or deployed. What the end-user is looking for are applications that deliver value at the right time.

DevOps has three main objectives:

- ✔ Speeding continuous innovation of ideas by enabling collaborative development and testing across the value chain
- ✔ Enabling continuous delivery of these innovations by automating software delivery processes and eliminating waste while still helping to meet regulatory concerns
- ✔ Providing a feedback loop for continuous learning from customers by monitoring and optimizing the software-driven innovation

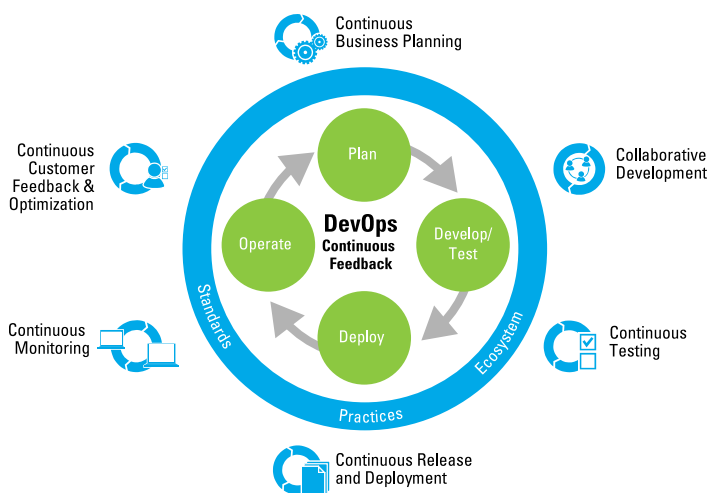
A DevOps approach needs to be complemented with additional business process changes that support innovation and rapid customer feedback cycles. You need to establish standardized, consistent, and repeatable processes for managing software quality from requirements definition to delivery. Companies are incorporating methodologies such as Agile, Scaled Agile, Lean Startup, and Design Thinking to build better applications and products while staying focused on solving problems that matter.

Agile and Scaled Agile methodologies support collaboration between teams and encourage increased involvement from end-users. This approach places higher value on the interaction between individuals than on the tools they use and the processes they follow. The emphasis is on rapid iterations of code development, feedback, and testing in order to increase responsiveness to change. The Lean Startup approach is consistent with Agile, given its focus on continuous testing of the vision for a new product. The Lean Startup methodology encourages teams to validate ideas and test possible solutions before committing any significant efforts. Design thinking emphasizes the user experience. What aspects of design will delight users and encourage them to feel a connection with the product or solution?

## *Optimizing DevOps for the Hybrid Cloud*

One of the most important benefits of the hybrid cloud is the flexibility you gain to select the right environment for your workload. DevOps can help to ensure that whatever platform you choose for your workload is a seamless experience for the end-user. When you use DevOps methodology to build, manage, and run applications, you're in a better position to ensure consistency across platforms in your hybrid environment. DevOps services provide a platform for managing the delivery life cycle from conception through testing and deployment by using iterative development processes. Taking a DevOps approach means that all parts of the organization are committed to the continuous delivery of software, including the development and deployment teams. DevOps teams share a common culture focused on creating a smooth transition from design to development to production. Operations staff may share the same tools and techniques as developers. When DevOps services are delivered in a cloud-based environment, it's easier to share project data across every team. As a result, these teams see improvements in collaboration for all stakeholders and higher levels of customer satisfaction.

Figure 3-1 illustrates a typical DevOps reference architecture.



**Figure 3-1:** The DevOps reference architecture.

The six key elements of this architecture are summarized here:

- ✔ **Continuous business planning:** In order to support your goal of continuous innovation, you need a flexible approach to business planning. The best plan is one that can be easily changed in response to customer feedback. You need to continuously measure your progress against changing customer requirements and ensure you are aligned with the needs of the organization.
- ✔ **Collaborative development:** Developers need a high level of collaboration to meet service levels in today's fast-paced environments. You need to remain focused on the customer experience throughout the software life cycle. Practices such as user interface (UI) prototyping, source code management, and code review can improve collaboration and keep all stakeholders moving toward a common goal.
- ✔ **Continuous testing:** You can't wait until the end of the development cycle to do your testing. If you delay testing for errors early, it becomes much harder to diagnose and correct later. By driving toward higher quality earlier in the life cycle, the test cycles are shorter, and there's less rework of previously developed software components. This shift-left testing approach helps developers and

testers collectively improve time to market and decrease software development costs.

- ✔ **Continuous release and deployment:** To achieve continuous delivery of changes in a shift-left testing approach, continuous deployment is a necessary practice and demands automation.
- ✔ **Continuous monitoring:** Developers need a way to monitor how their application is performing before it's deployed to production. Application performance monitoring can be used to provide early feedback in order to reduce the cost to fix errors and speed up the time to completion.
- ✔ **Continuous customer feedback:** Because you're continuously developing, testing, and releasing code at a fast pace, the capability to receive and incorporate feedback from customers is critical. You need to monitor the application in real time and make sure the end result is delivering on customer expectations.

## The benefit of using cloud-based DevOps services

A mobile software development organization needed to speed up the delivery of its projects to customers. Speed is a priority in all aspects of the development and deployment life cycle. The application needs to be developed, tested, and put into operation quickly. In addition, it needs to perform with the right level of speed and accuracy. The organization's customers are focused on delivering mobile apps that delight end-users with innovation, ease of use, exceptional performance, and speed. According to the company's chief creative officer, "When our customers are developing their mobile strategies, they need to get their apps to market as quickly as possible."

The company recognized that one way to speed up the application delivery process would be to improve the way stakeholders collaborate throughout the life cycle of the application. Prior business processes relied on spreadsheets and emails to share information on a project among key stakeholders. These processes led to unacceptable delays in the schedule and some serious software quality issues. In addition, the company's traditional methods of collaboration didn't maintain high enough levels of security for customers who were concerned about protecting their intellectual property using a DevOps approach.

This organization implemented cloud-based IBM Bluemix DevOps Services to build, manage, and run apps. This move immediately led to improvements in developer productivity. The bottom line benefit of this new approach has been to cut costs and speed up the application delivery process. Project startup costs associated with setting up the required infrastructure decreased substantially. One of the ways that the company used these DevOps services to increase development speed was improving both internal and external collaboration. For example, leveraging cloud-based collaboration tools has enabled the team to decrease the time required for daily team meetings. These team meetings are used to ensure that everyone on the team

is up to date on daily goals, accomplishments, and potential obstacles. The team still holds these important meetings, but it can be much more efficient because many concerns are addressed throughout the day by using the project management dashboard. In addition, other stakeholders are also able to access information about their projects in near-real time on a dashboard, which helps to ensure they're active participants in the development process.

Overall, customer satisfaction has improved. Customers receive their apps faster, have a better understanding of the development process, and can easily collaborate with the development team to ensure that their needs are met.

## *Enabling Delivery Automation*

In order to guarantee that your applications meet quality, security, and performance objectives throughout the life cycle, you need to remove the silos that have traditionally existed between development and IT operations. Following an automated DevOps approach helps to remove these silos and enables a continuous delivery of changes. However, your approach will vary depending on the types of applications you build. We characterize development teams according to two different categories based on the complexity of their applications and integration requirements:

- **Developers of cloud native applications:** Cloud native applications are often built for mobile platforms and social media. There are fewer dependencies to consider in the development process. Developers use web services for capabilities such as data management and analytics to speed up development time. In addition, developers may need to leverage cloud integration services to connect the mobile applications to their back-end services.

✓ **Enterprise developers of cloud enabled applications:** Developers working in enterprise hybrid cloud environments need to combine traditional applications (primarily developed to run in the data center and manage systems of record) with new applications developed in the cloud focused on systems of engagement. These applications have more dependencies, and require larger teams and more advanced processes and tools to automate the delivery pipeline.

Delivery automation can be achieved by implementing DevOps services either on-premises or in the cloud. Both categories of developers benefit from a Platform as a Service (PaaS) environment. The enterprise developers of cloud enabled applications benefit from incorporating a unified DevOps solution that supports more complex hybrid cloud applications. These services are introduced in more detail in the remainder of this chapter.

## *Platform as a Service (PaaS)*

PaaS provides hosted services used during development, such as middleware, operating systems, and security solutions. A well-designed PaaS environment can help to automate the software delivery process and improve collaboration between development and operations teams. Using a PaaS approach also ensures consistency because DevOps teams select from a shared set of tools and frameworks. Developing in a PaaS environment provides developers with a way to test, find, and fix errors faster, leading to improvements in software quality and faster development times. Developing in a PaaS environment reduces the complexity and increases consistency across the life cycle.

## *Application Release Automation (ARA) and testing services*

Applications Release Automation (ARA) also provides a platform of services to support DevOps automation. Capabilities go beyond typical PaaS services to support the more complex provisioning, integration, and testing requirements for complex hybrid cloud applications.



## Creating Continuous Delivery

DevOps services provide teams with a unified environment to ensure consistency across development and deployment. Different teams leverage the capabilities of shared tools and frameworks according to their requirements. These platforms provide the teams with a complete environment of application development and deployment services to support continuous delivery of applications. DevOps services can help to automate the application delivery pipeline — from provisioning to software build to testing and production. As a result, DevOps teams are able to start projects more quickly and stay focused on the task of delivering applications that meet customer expectations.

Examples of frameworks and tools that might be available in a catalog of DevOps services are highlighted in this section.

### DevOps frameworks

Developers typically use a development framework to help create a consistent structure for an application. While frameworks are also used in traditional software environments, these frameworks can be more easily shared across large distributed teams when used in a PaaS or DevOps services environment. Frameworks help to ensure quality across the application delivery life cycle. Examples of services available in DevOps frameworks include

- ✔ **Configuration management:** Developers use software configuration management services to keep track of all the changes, versions, and modules of code that are created during the software development process. The code is stored in an online repository. GitHub is an example of a popular hosted service used for this purpose. Software configuration management services help the developer to manage the sandbox environment used to create and test code.
- ✔ **Application build:** The build process contains several steps, including writing code, compiling code into an executable, running, and then testing the code. When creating an application, developers produce multiple modules of code each with its own dependencies. The build

services in a DevOps framework can help the developer keep track of the individual modules of code. Specific services are added to build a deployable application.

- ✓ **Testing:** In a DevOps environment, testing needs to be done continuously to allow for consistent feedback between the teams. You also need to have access to rapid feedback from test execution and customer usage to determine if your application meets customer requirements. An automated approach to testing increases speed to completion and quality of results. Testing needs to happen at the same pace as the coding. The overall goal is to shorten testing cycles and decrease costs by finding and fixing errors early in the code development process. DevOps testing services make it simple to access and isolate production systems, so developers can stop wasting time mocking up simulations and avoid the risk of discovering runtime errors in production or at late stages of delivery.

Testing frameworks include numerous services for testing, including user interface testing, load testing, regression testing, integration testing, and performance and scalability testing. If you're building applications with multiple connection points, you need an automated testing process that tests for all the hardware, software, and cloud-based dependencies.

Given the complexity of certain hybrid cloud applications, capabilities such as service virtualization can be applied to allow for integration testing at an earlier point in the development process. You can deploy virtual application environments so they can be tested prior to production, helping to reduce costly errors and improve time to market.

## *Reusable web services*

DevOps teams can leverage a catalog of reusable web services to speed development and deployment. Several examples include the following:

- ✓ **Security services:** There are many benefits to delivering software within an integrated PaaS environment that includes built-in security services. With so many different points of connection in a hybrid environment, a multitude of security vulnerabilities can be introduced during

the application development and deployment process. Many of these security risks can be avoided if your PaaS environment includes a best practices approach to security. Security services can be included in delivery automation to ensure that security testing is done regularly. When DevOps teams follow a consistent and automated approach to security, you can deliver more secure applications to the business.

- ✔ **Mobile application services:** Mobile applications are typically designed to analyze data from multiple sources in order to provide users with information they need on demand (for example, travel directions, restaurant suggestions, record of banking transactions). DevOps teams can leverage services designed to push information to all application users or to a set of users and devices.
- ✔ **Analytics services:** Businesses are increasingly applying data analytics to solve a wide range of problems. Having the capability to access and reuse pre-tested analytics services for application development improves quality and increases time to market. One example is an analytics service designed to leverage real-time geospatial analytics to track when devices move from a location.
- ✔ **Internet of Things (IoT) services:** The demand for applications that rely on streaming data from sensors is increasing at a fast rate. Your DevOps team can keep up with the pace of business requests by leveraging services to support these applications. For example, some Internet of Things (IoT) services focus on communication, enabling your application to communicate with and consume data collected by your connected devices, sensors, and gateways.

## *Monitoring and Analytics*

In hybrid environments, it becomes even more critical that the deployment team has automated processes to monitor and control the organization's applications. The IT organization may have direct responsibility for the deployment of applications developed internally, while also needing to maintain control over data and applications managed by SaaS providers. Ultimately, if there's a data breach or critical security patches that need to be applied to applications dispersed across many departments, an automated deployment solution with auditing will be indispensable.

You need to be able to monitor application and system performance. Application Performance, Usage Monitoring, and other similar tools provide insight into how your application is performing in production. Some monitoring tools provide insight into the health of your application as well as your environment. DevOps services environments typically have a dashboard for visualizing the metrics associated with these tools. For example, you can use the dashboard to monitor whether your services provider is up or down and is meeting its service level agreements (SLAs).



Ideally, in a hybrid environment, you should look for a dashboard that gives you uniform visibility across your own resources and those of your cloud services provider. Integrated diagnostics to monitor application performance will help provide the operations team with information to rapidly respond to production problems.

Application Analytics can be used to automatically identify patterns and anomalies in production metrics, helping to speed the identification of a problem and shorten the time to resolution. The analytics tool automates the collection, analysis, and correlation of events across federated data sources. DevOps teams can quickly identify problems in production and begin work on finding a solution. You can gain insight on how the application is being used and quickly assess how end-users are responding to a new application feature.

## *Looking at the Impact of Open Standards for the Hybrid Cloud*

Standards in a hybrid cloud are important because they help you improve quality, reduce cost, and improve choices. Without the broadly adopted open standards of Open Stack and Cloud Foundry for the models, formats, and conventions for interacting with the cloud, the hybrid cloud environment presents significant challenges. Simply put, without standards, or agreed upon approaches, moving your infrastructure or applications from one cloud provider to another, or from on-premises to a public or private cloud, is a difficult prospect that can slow an organization's development. Integrating your on-premises data center in a hybrid model

would be difficult without standards. Standards also help to ensure security and prevent you from becoming locked in to a proprietary software or cloud infrastructure. All these issues are key in a hybrid environment.

Open source DevOps environments abstract the underlying infrastructure you need to run a cloud. They also provide you with lots of choices regarding the frameworks and services you want to work with. For example, your open source PaaS should provide you with a variety of development frameworks and languages such as Java, Spring, Ruby, and Node.js. In addition, you will have a range of application services to choose from, including MySQL, MongoDB, PostgreSQL, Redis, and RabbitMQ. You should also be able to customize any of these options to meet the specific requirements of your organization.



## Chapter 4

# Data and Analytics

### *In This Chapter*

- ▶ Finding the value in managing data and analytics in a hybrid cloud
- ▶ Using analytics to deliver business results
- ▶ Looking at cognitive computing

**D**ata and analytics have become a top priority across industries. Traditional businesses are rapidly becoming disrupted by emerging companies that leverage data, analytics, and mobile technologies to service customers in new ways. When you observe how quickly new companies like Uber and Airbnb challenged the status quo for taxis and hotels, you realize that your industry may also be vulnerable. This isn't the time for a slow and easy transformation. Many companies are feeling the urgency to rethink their business models and build an IT platform that creates a competitive advantage in their industries.

A hybrid cloud infrastructure is imperative for a successful analytics strategy because of the need to support varying workloads and requirements for big data. Driving better business decisions with analytics typically demands bringing together multiple sources of data from different systems. For example, your organization may develop a predictive analytics model to improve your understanding of customer expectations and improve customer satisfaction. Your model needs timely and accurate data that may reside on different systems such as customer data stored on a Software as a Service (SaaS) application, product information stored in an on-premises database, and demographic data from a third-party data provider.

Hybrid cloud gives you the flexibility required to ensure that you have the compute capabilities you need to build and execute your predictive models in an efficient and cost-effective manner. For example, you can leverage high-powered compute resources in the cloud to meet service levels without consuming internal resources.

## *Finding the Value of Managing Data and Analytics in a Hybrid Cloud*

Organizations are gaining huge value by leveraging analytics in hybrid cloud environments. In order for analytics to help fuel innovation and drive new business models, your results often need to be available in real time. To meet these real-time requirements, your need for compute power and storage may be highly variable. You need the capability to access, move, refine, and analyze all types and volumes of data to support predictive models and emerging cognitive computing environments. In fact, managing data in preparation for model development and analysis consumes a major portion of the resources required for analytics projects. Cloud-based data services can reduce the time and minimize the investment required for data preparation and its inevitable growth. In addition, because of the cloud's capability to scale up and down, cloud-based data services can help reduce costs. Hybrid cloud provides the choice and flexibility to select the right IT environment, data sources, and services at the right time to ensure you meet your customer requirements.

Key considerations for managing data and analytics in a hybrid cloud are listed and described in the following sections:

- ✔ Understanding and locating relevant data sources from inside and outside your organization
- ✔ Refining and enriching your data to ensure reliability and quality
- ✔ Loading, storing, analyzing, and distributing application data



- ✓ Data services to replicate, synchronize, and integrate data
- ✓ Building, testing, and running advanced analytical models

## *Locating relevant data sources*

Business decisions need to be made based on constantly changing data from a variety of sources. Your data sources may include both traditional systems of record data (for example, customer, product, transactional, and financial data) as well as newer, diverse sources, such as external data feeds, mobile devices, sensors, and the web. These data sources can come in the form of both structured and unstructured data.

Structured data is typically stored in traditional relational databases and refers to data that has a defined length and format. Most organizations have a large amount of structured data in their on-premises data centers. Examples of structured data include numbers, dates, and groups of words and numbers called *strings* (for example, a customer's name, address, and so on). On the other hand, unstructured data is data that does not follow a specified format. Experts estimate that approximately 80 percent of an organization's data is unstructured; therefore, there's great opportunity in being able to analyze this type of data. Whether it's product reviews on a store's web page, interactions through a mobile application, or postings on social media, cloud, mobile, and social have contributed to a huge increase of unstructured data.

Some examples of both structured and unstructured data sources you may need for your analytics are detailed in this list:

### ✓ **Structured data sources**

- **Sensor data:** Examples include radio frequency ID (RFID) tags, smart meters, medical devices, and Global Positioning System (GPS) data.
- **Weblog data:** When servers, applications, networks, and so on operate, they capture all kinds of data about their activity.
- **Point-of-sale data:** When the cashier swipes the bar code of any product that you're purchasing, all that data associated with the product is generated.

- **Financial data:** A lot of financial systems are now programmatic; they're operated based on predefined rules that automate processes.
- **Weather data:** Sensors to collect weather data are being deployed across towns, cities, and regions to collect data on things like temperature, wind, barometric pressure, and precipitation. This data can help meteorologists create hyperlocal forecasts.
- **Click-stream data:** Data is generated every time you click a link on a website. This data can be analyzed to determine customer behavior and buying patterns.

#### ✔ Unstructured data sources

- **Text internal to your company:** Think of all the text within documents, logs, survey results, and emails. Enterprise information actually represents a large percent of the text information in the world today.
- **Social media data:** This data is generated from social media platforms such as YouTube, Facebook, Twitter, LinkedIn, and Flickr.
- **Mobile data:** This includes text messages, notes, calendar inputs, pictures, and videos as well as data entered into on third-party mobile applications.
- **Satellite images:** This includes weather data or the data that the government captures in its satellite surveillance imagery.
- **Photographs and video:** This includes security, surveillance, and traffic data.
- **Radar or sonar data:** This includes vehicular, meteorological, and oceanographic data.

Bringing all these varied data sources together for analysis can be a big challenge. This is one way that cloud services can help. Newer cloud-based data management offerings include data storage capabilities, data refinement, and embedded analytics technologies. The data you need for analytics may be in the public or private cloud, stored in your enterprise data center, or available from a third-party data service. Figure 4-1 illustrates examples of data sources and where they may be stored.

## Leveraging third-party weather data for your analytics model

Most businesses need to include both internal and external data to develop analytics models that provide actionable insights. For example, enterprises in industries ranging from retail to energy and utilities to insurance need to include weather forecasts in their business planning. Predictive analytics models that account for variations in weather do a better job of helping decision makers in many areas, including inventory planning, risk management, and setting prices. One third-party source for weather forecasts is The Weather Company. It collects data from weather sensors, aircraft, smartphones, buildings, and moving vehicles to forecast weather data that businesses can use in their own operational models. WSI, a division of The Weather Company, and IBM partnered together to make it easier

for businesses to include accurate and timely weather data into their operations and decision making. WSI's weather data services platform is hosted on the IBM Cloud and its weather data is integrated with IBM analytics and cloud services.

Having access to real-time weather insights can make a huge difference in the reliability of a company's analytical models. When this data is provided as a cloud service, companies can use the information to improve business outcomes. For example, insurance companies can potentially save millions of dollars per year in claims paid by using real-time weather data to alert policyholders of impending hailstorms. Sending a text message alert would give policyholders time to move their vehicles to safe spots before they are damaged.

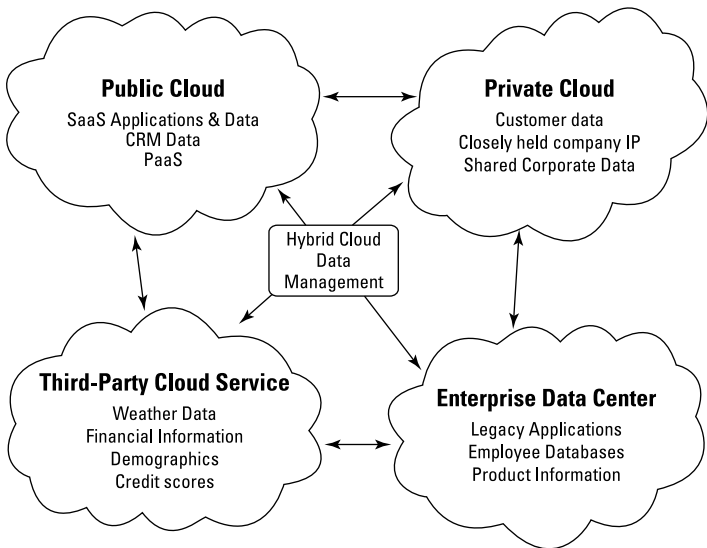
## *Refining and enriching data*

Application developers and business users need quick, reliable, and efficient access to data to build applications and actionable analytical models. While it can be very challenging to access the right data at the right time, this is only the beginning of many other time-consuming steps that need to be followed to ensure that your data can be trusted. Without a process for data refinement, the long lead times in accessing the right data can become a roadblock to further analysis.

A sample of the questions you need to ask to understand if your data is trustworthy includes

- ✔ Did you locate the most relevant data for your analysis?
- ✔ Do you understand the source for that data?
- ✔ Is your data up to date and accurate?
- ✔ Has your data been changed and, if so, who changed it?
- ✔ Is your data consistent across different departments (are definitions the same)?

Data refinement provides the foundation for building analytical models that deliver results you can trust. The process of data refinement helps ensure that your data is timely, clean, and well understood. In a hybrid cloud environment, companies can leverage cloud-based data access and refinement services.



**Figure 4-1:** Data sources in a hybrid environment.

## Defining big data

Big data is defined as any kind of data source that has at least three shared characteristics — extremely large *volumes* of data, an extremely high *velocity* of data, and an extremely wide *variety* of data. When dealing with big data, many organizations leverage a hybrid cloud environment because of its capability to scale. Scale and speed are critical capabilities in leveraging big data to help businesses improve business knowledge, anticipate change, and predict outcomes. Analytics depends on big data because of the need to look at very large volumes of data to uncover patterns and emerging trends. Without big data, you're trying to make decisions on small subsets of your data that may lead to misinterpreting a trend or missing a pattern that's just starting to emerge.

The hybrid cloud plays a major role in making big data actionable. The rate of growth in the volume of big data is so fast that it's becoming increasingly challenging to efficiently store, locate, and analyze

this data. Many organizations look to public cloud resources to store their big data. However, when the data contains proprietary information, many organizations are using tightly governed private cloud for big data storage.

Traditional business intelligence (BI) products weren't really designed to handle big data. They were designed to work with highly structured, well-understood data, often stored in a relational data repository. These traditional BI tools typically only analyze snapshots of data rather than the entire data set. Analytics on big data requires technology designed to gather, store, manage, and manipulate vast amounts data at the right speed, at the right time to gain the right insights. With the evolution of computing technology and the emergence of hybrid cloud architectures, it's now possible to manage immense volumes of data that previously could have only been handled by supercomputers at great expense.

## Loading, storing, and analyzing data

Various data stores are available for the structured and unstructured data sources you may need for your analytics, including relational databases, NoSQL databases, relational data marts, and Hadoop. While many companies

store traditional system of record data in their internal data centers, the movement to use cloud-based data stores for newer data sources is increasing. Cloud-based data stores provide the scalability and access needed to speed up the development and deployment of web and mobile applications. You can have access to sophisticated capabilities without the cost, complexity, and risk of managing the infrastructure. In addition to these various data stores, emerging data processing engines, like MapReduce and Spark, exist that are allowing organizations to process big data.



Hadoop is one of the most important technologies for managing large amounts of unstructured data. Organizations are implementing Hadoop both on-premises and in the cloud. Hadoop is an open-source community, codebase, and market for a big data environment that's designed, among other things, to parallel-execute code written to MapReduce. Text documents, ontologies, social media data, sensor data, and other forms of nontraditional data types can be efficiently managed in Hadoop. The benefit of using Hadoop is that you can quickly transform massive amounts of nontraditional data from raw data to structured data so you can search for patterns in that data.

Hadoop is particularly useful for managing big data in cognitive computing because it's easy to dynamically scale and make changes quickly. Hadoop provides a way to efficiently handle the problem of taking highly unstructured data, and you break it up into component parts to then solve the problem and produce results. Hadoop can be implemented on racks of commodity servers or included in a pre-optimized appliance on vendor-specific hardware. Three key capabilities associated with Hadoop are

- ✔ **Hadoop Distributed File System (HDFS):** A data storage cluster that's both highly reliable and low-cost used to make it easy to manage related files across different machines.
- ✔ **MapReduce engine:** Provides a way to distribute the processing of the analytics algorithms across a large number of systems. After the distributed computation is complete, all the elements are aggregated back together to provide a result.

✓ **Spark:** An in-memory compute engine based on a cluster-computing framework. This is in contrast to the disk-based MapReduce engine. Spark has made it possible to blend and analyze large scale, disparate data sources quickly and in an iterative way. Spark can run on top of other software, such as the Hadoop HDFS, or on its own. For developers, this flexibility means that they don't need to write different versions of Spark applications for different platforms. The high-speed of Spark makes it particularly well suited for emerging applications that are using more data sources, including IoT, machine learning, web, mobile, social, and business process.

## *Data services to replicate, synchronize, and integrate data*

The capability to replicate and integrate data across hybrid cloud environments is critical. Replicating data in essence gives developers a flexible data layer. Developers can write applications on top of this flexible data layer without worrying about where the data is located. Traditionally, developers and IT departments have had to make deployment decisions up front when launching an application. For example, if you wanted to launch an application with a database on-premises, you had to make that decision upfront. With the capability to replicate a database to the cloud, a database or application can be started either in the cloud or on-premises and then switch because of economic, performance, or security concerns. For more information on this topic, check out Chapter 5.

## *Building, testing, and running advanced analytical models*

Advanced analytics refers to a collection of techniques and algorithms for identifying patterns in large, complex, or high-velocity data sets with varying degrees of structure. It includes sophisticated statistical models, predictive analytics, machine learning, neural networks, text analytics, and other advanced data mining techniques. Some of the specific statistical techniques used in advanced analytics include decision tree analysis, linear and logistic regression analysis, social network analysis, and time series analysis.

Many organizations are using hybrid cloud architectures to apply these techniques in order to run analytics on shared data sources or take advantage of the economics of the cloud. These analytical processes help discover patterns and anomalies in large volumes of data that can anticipate and predict business outcomes. With the right level of advanced analytics, you can gain deeper insights and predict outcomes in a more sophisticated manner.

Below are two examples of advanced analytics:

- ✔ **Predictive modeling:** Predictive modeling is one of the most popular advanced analytics use cases. A predictive model is a statistical or data-mining solution consisting of algorithms and techniques that can be used on both structured and unstructured data (together or individually) to determine future outcomes. For example, a pharmacy might use predictive modeling to anticipate demand for flu medication.
- ✔ **Text analytics:** Unstructured data makes up a huge percentage of an organization's data. Therefore, text analytics — the process of analyzing unstructured text, extracting relevant information, and transforming it into structured information that can then be leveraged in various ways — has become an important component of advanced analytics. Text analytics is being used in all sorts of analysis, from predicting customer churn to reducing fraud, and understanding employee and customer sentiment via social media analytics.

## *Using Analytics to Deliver Business Results*

Companies are increasingly operationalizing analytics or making data analytics part of the business process. For example, statisticians at an insurance company may build a model that predicts the likelihood of a claim being fraudulent. The model, along with some decision rules, could be included in the company's claims-processing system to flag claims with a high probability of fraud. These claims would be sent to an investigation unit for further review. In other cases, the model itself may not be as apparent to the end-user. For example, a model could be built to predict customers who are



good targets for upselling when they call into a call center. The call center agent, while on the phone with the customer, would receive a message on specific additional products to sell to this customer. The agent may not even know that a predictive model was working behind the scenes to make this recommendation.

In addition to making data and analytics an integral component of the business decision-making process, companies are also finding new ways to monetize analytics. Analytics is also transforming roles because virtually everyone can be empowered to be a knowledge worker. Previously advanced analytics had been limited to highly skilled analysts, such as statisticians, data miners, and data scientists, who are now available to most people who need them. Now, no matter where you are in your organization, you can use analytics to discover insights, answer questions, and take action. Big data analytics can be used to derive revenue above and beyond the insights it provides just for your own department or company. You might be able to assemble a unique data set that is valuable to other companies, as well your own. For example, credit card providers take the data they assemble to offer value-added analytics products. Telecommunications companies are beginning to sell location-based insights to retailers. The idea is that various sources of data, such as billing data, location data, text-messaging data, or web-browsing data, can be used to make inferences about customer behavior patterns those retailers would find useful.

This section gives you some examples of how companies are using the cloud, big data, and analytics to deliver business results. In many of these cases, because of the complex nature of the challenge, a hybrid cloud becomes the preferred platform.

## ***Predictive maintenance***

Organizations in the manufacturing, transportation, mining, utility, and energy industries rely on large, expensive equipment to produce revenue. When equipment fails, or is taken out of commission for maintenance, the company may lose thousands of dollars an hour. In addition, unnecessary maintenance creates waste and takes equipment out of use. It has become a high priority for these organizations to use predictive analytics to anticipate and correct problems with equipment.

## Energy industry: Predictive maintenance

A large wind turbine company is using advanced analytics and a hybrid cloud environment to predict machine failure. The company has instrumented its turbines with hundreds of sensors that measure data points like vibration, temperature, and atmospheric pressure. This data is fed into the organization's analyt-

ics platform and is enriched with data from a third-party data source that provides weather data. After implementing the system, the organization was able to predict turbine failure up to two weeks before a breakdown was predicted to occur. Maintenance crews are now able to take proactive steps to prevent costly downtime.

## *Fraud*

Identifying and preventing fraud is complex — schemes are continually changing, unstructured data is often involved, and most importantly, a company doesn't want to cut off legitimate transactions. Preventing fraud is a top priority for organizations in a variety of industries, including financial services, insurance, and government.

## Insurance industry: Preventing fraud

A national health insurance company is using advanced analytics and the cloud to fight fraud. Before implementing the analytics solution, the insurer had a labor-intensive, manual document review procedure to spot potentially fraudulent claims. The company has moved claims data to a public cloud that is accessible to both employees and customers. It is using text analytics to understand the unstructured data contained in the claims data. It is then using social analytics to understand the

relationships between the involved parties and historical customers' records stored on-premises. All this data is fed into an analytics platform that scores the claim based on the insurer's business rules. The system prioritizes potentially fraudulent cases for further review. Customer satisfaction has increased because claims can be processed more quickly while the insurer has been able to prevent millions of dollars in fraudulent payments.

## *Marketing and customer churn*

With increased competition, capturing and retaining customers is a top priority for organizations across industries. By leveraging big data, analytics, and cloud services, organizations are able to spot high-value customers who are likely to churn. After a customer is spotted, proactive measures can be taken — for example, a customized incentive can be offered to retain the customer.

### **Telecommunications industry: Reducing customer churn**

A large telecommunications provider is using a cloud-based analytics platform to create more effective marketing campaigns and reduce customer churn rates. The company synchronizes its on-premises customer database to the cloud and feeds it into its analytics platform. Other data, such as demographic data, is also

added to the analytics process. The company created real-time scoring models to predict customer churn. Customer service representatives are given custom offers that they can give to customers as an incentive. The company has been able to reduce its customer churn rate while also improving customer satisfaction.

## *Looking at Cognitive Computing*

Companies have been successful at using analytics to understand both where they've been and how they can learn from the past to anticipate the future. They can describe how various actions and events impact outcomes. Although the knowledge from this analysis can be used to make predictions, typically these predictions are made through a lens of pre-conceived expectations. Data scientists and business analysts have been constrained to make predictions based on historical data. However, unknown factors always exist that can have a significant impact on future outcomes. The next generation of solutions combines some traditional technology techniques with innovations so organizations can solve vexing problems.

Cognitive computing is an evolution of technology that attempts to make sense of a very complex world that's

drowning in data. While cognitive computing is still in its early stages of maturation, a new era in computing will transform the way humans collaborate with machines to gain actionable insights. Because of the scale and scope of cognitive computing, a cloud-based approach has been the obvious choice for deployments. In addition, cognitive capabilities are being given to developers via Platform as a Service (PaaS) to embed cognitive computing into web and mobile applications.

Over the coming decade, you'll see cognitive capabilities built into many different applications and systems. Cognitive systems can be used to help with the transfer of knowledge and best practices in other industries as well. Early cognitive systems are designed to build a dialog between human and machine so best practices are learned by the system as opposed to being programmed as a set of rules. The following list includes some areas where cognitive computing is beginning to be applied. Use cases will continue to grow and this list is only a small sampling.

- ✔ **Healthcare:** Cognitive systems are being developed to ingest medical research, patient records, and best practices. The systems work with doctors to diagnose, patients and develop the most effective treatment plan.
- ✔ **Travel:** Systems are being developed to help travelers receive experience-based recommendations that coincide with the traveler's objective. These systems will provide the personal service of a travel agent while gaining the cost-saving of do-it-yourself travel booking.
- ✔ **Government:** Governmental organizations are using cognitive computing to run cities more efficiently. Data from each part of a city's infrastructure and services (for example, transportation, police, and air quality) can be fed into a cognitive system and decisions can be made that take into account all the ripple effects.
- ✔ **Legal:** With the help of law firms, technology vendors are creating cognitive systems that ingest laws and regulations, case law, and legal journals. These systems will help attorneys apply relevant arguments to their cases that they may have never discovered on their own.
- ✔ **Retail:** Systems are being developed to anticipate customer requirements and streamline the supply chain. These systems will help retailers anticipate trends and ordering needs many seasons in advance while also giving customers a customized experience.

## Chapter 5

---

# Data and Application Integration and Portability

.....

### *In This Chapter*

- ▶ Seeing the need for integration and portability
  - ▶ Working with data in the hybrid cloud
  - ▶ Juggling multiple data sources within a hybrid environment
  - ▶ Incorporating applications across environments and gaining portability
- .....

**C**ompanies are moving to the hybrid cloud to realize enormous business benefits in terms of flexibility, improved customer engagement, and cost savings. However, in order to realize these benefits, companies need to solve the challenges of integration and portability for data and applications. In this chapter, we discuss the need for integration and portability across a variety of cloud deployment models.

## *The Imperative for Integration and Portability*

Integration and portability are important in a hybrid cloud environment because of their flexibility. Technology needs of businesses are constantly changing. A decade ago, the data center was the primary deployment model for systems that operated the business, but things have changed significantly. Many companies choose to manage mission-critical systems, such as transaction management and customer-facing applications, on-premises in the data center. At the same time,

businesses are gaining benefit from Software as a Service (SaaS) applications that operate in a public cloud environment. When the data requires higher levels of security and compliance, many organizations choose to move those applications to private clouds.

While each deployment model supports needs of various departments and organizations, a requirement exists to be able to integrate data across this hybrid environment. Business leaders recognize that in order to create new business models to support customers they must be able to bring together data from different applications across traditional IT and public and private cloud services. Typically, the organization needs to integrate systems of engagement applications such as social media and customer management systems hosted in the public cloud with mission critical applications and databases stored on servers in its internal data center. Companies also have a need for efficient integration of applications across clouds.

The need for integration is always changing depending on the workloads and type of analysis. As companies begin to incorporate analytics into business processes, data integration requirements are becoming more complex. The volume and variety of data included in big data analytics workloads require systems that can scale to manage the required data integration and management. Increasingly, businesses are responding to customer and competitive pressures to deliver mobile-based applications in the cloud that leverage transactional systems in the data center. Many of these applications are emerging as new profit centers for the business.

## *Managing Data in the Hybrid Environment*

Most organizations have myriad data siloed across business units. Critical customer data is often stored in applications and business processes in business units, such as SaaS, public and private clouds, and data centers. Typically, business units make independent decisions to use a variety of SaaS applications to support their immediate requirements. To better service customers and develop innovative business models, organizations have begun the arduous process of

integrating data across these platforms. This task is complicated because data is often tightly woven into business processes. In addition, data isn't static; it's contently changing, so data integration isn't a one-time process. Data integration requirements change depending on the volume and complexity of the data and the speed at which it needs to be delivered across applications.

Many mobile applications are customer facing and require real-time responses to deliver results. If you want your customers to be able to select and buy products by using mobile applications, you need the ability to quickly access data from traditional IT systems, such as customer account records and product information. Therefore, there's an increasing need for rapid and secure integration between mobile environments and back-end systems for customer and transactional data.

## *Managing integration and data portability*

Companies need the flexibility to move between cloud-based Hadoop environments and traditional IT database environments depending on the economics. For example, when analyzing cloud-based social media data to gauge customer reaction to a new product, a cloud-based infrastructure may be the most cost-effective approach. It's easy and quick to create the Hadoop clusters required to search for patterns in large volumes of data. After you complete the first phase of your analysis, you may choose to bring a subset of the data into an internal NoSQL database to answer different types of questions. Each new analysis requires a different approach to data integration. You may need to correlate data across different sources that are managed in a Hadoop environment, the internal data center, or a private cloud.

## *The need for portability*

One of the primary benefits of the hybrid cloud is the capability to maintain agility and flexibility to support changing customer needs. Therefore, it is essential to be able to move workloads across deployment models. You need to be able to run the same application in different clouds and move between public and private clouds depending on the workload

and security and governance demands. You may want to move your workload if you encounter capacity constraints. With the diversity of current workloads, including big data analytics, capacity planning can't always be done. As a result, it helps to have the capability to easily burst to a public cloud and retain capacity in your private cloud for mission critical workloads.

## ***Container technology supports portability***

The emergence of enterprise container technology gives developers the ability to build applications that can be moved to different deployment models. Docker has emerged as the open-source approach for creating interoperable containers and as an alternative to virtualization. Many vendors have begun to standardize on Docker as a primary container technology. With Docker, it's possible to package an application and its dependencies into a lightweight container that can run on any Linux server. This means that containers are an efficient way to manage resource utilization. These containers help create an encapsulated environment for applications, making it simpler to deploy and move the application in either public or private clouds. In addition, Docker is useful in creating and managing microservices. *Microservices* is an architectural approach for building small business services that can be run inside a container. Combining several microservices can create a single deployable application. This level of modularity provides an organization with an efficient way of moving data and applications depending on changing business and technical requirements.

## ***Developing standards to support workload portability***

OpenStack is an open-source cloud computing environment. OpenStack began as a project in 2010 and has quickly grown and matured into a thriving organization with more than 500 members. The software allows IT organizations to control large pools of compute, storage, and networking resources throughout an environment, managed through a dashboard or via APIs. OpenStack has a modular architecture that's



made up of a number of component services. Some of the OpenStack services include compute, storage, networking, database, and identity.

Driven by customer demand, many cloud vendors have joined the OpenStack project and are contributing resources to further develop the environment. In addition, public cloud vendors have built their public clouds based on OpenStack standards. Organizations that build their private clouds based on OpenStack have the capability to more easily move workloads to OpenStack-based public clouds or burst into a public cloud during usage spikes.

## *Managing Multiple Data Sources within a Hybrid Environment*

The task of managing and integrating multiple data sources from a variety of applications requires you to understand the nature of the variety of data formats and capabilities. In this section, you examine the key requirements for data integration.

### *Managing metadata*

A critical component that enables the ability to integrate data is the metadata. *Metadata* is the definitions, mappings, and other characteristics used to describe how to find, access, and use a company's data (and software) components. In essence, metadata defines the meaning and source of your data. One example of metadata is data information about a customer account. Defined columns outline the account number, name, address, phone number, and account level.



Metadata is especially important in a hybrid cloud environment where data from multiple sources is blended. Metadata helps you to organize your data stores so you can have a well-designed way to consistently integrate blended data sources.

### *Challenges in integrating data*

Organizations often have large amounts of complex data that can be in a variety of formats. When managing and integrating

data, you need to be aware that there are challenges. For example, some data sources may not be trustworthy. In other situations, some of the data may include errors or extraneous data elements.



Some common data challenges that organizations need to plan for include the following:

- ✔ **Data can come from untrusted sources, and it can be dirty.** *Dirty data* refers to inaccurate, incomplete, or erroneous data. Dirty data may be as simple as a user who entered a misspelled word. However, it can be more complex; for example, in an Internet of Things (IoT) implementation, a sensor may be damaged or not properly calibrated, which can lead to faulty data and inaccurate alerts. Businesses must often integrate data from various sources. These may include both tightly managed internal sources as well as third-party cloud data providers. In addition, many organizations have cloud-based web and mobile applications where data is entered by customers and partners, or an analyst may select a data set that seems relevant only to discover that the information is out of date or inapplicable. These external sources aren't trustworthy.
- ✔ **Data can come from real-time feeds.** In many hybrid cloud scenarios, you need to analyze real-time data streams. Streaming data brings about an entirely different set of technical requirements, such as performance and computational capabilities that can be powerful when handled correctly. This real-time processing is especially important in IoT use cases.

## *Managing security and governance*

If your company is planning to use a variety of cloud services, you must have the assurance that there's a clearly defined set of security services, such as identity management, access control, and authorization. There needs to be a comprehensive security infrastructure provided at all levels for the hybrid cloud environment.

Governance is an increasingly complex issue in the hybrid cloud environment. There are often governmental or corporate issues for how data must be stored and when it can be

moved. For example, a governmental policy may require personal data to remain in the place where that data was created. In addition, an industry or corporate policy may state that all customer data must be stored within the data center.

## *Replicating and synchronizing data*

Replicating and synchronizing data is one of the key technologies that allow application developers to get data as close to the application as possible. Modern web and mobile applications bring a tremendous amount of information to customers and employees. For example, with a simple mobile application, a banking customer can gain access to her account information, transfer money, and schedule payments. The cloud-based, front-end of the application is very easy to use. Behind the scenes, the application is actually gathering data and writing data to a variety of data sources that are housed in back-end data center transaction management systems and SaaS-based CRM systems. Complicating the functionality of the application is that users expect high speed and will abandon applications that are slow to load.



To support user expectations of high speed, master-to-master replication and auto synchronization is an important capability. To provide fast access, many organizations deploy multiple databases in public clouds in a variety of geographies. Problems can quickly arise when databases are replicated. For example, if the same database is being read and written to, make sure that all other users are able to see the correct version of the data. Therefore, replicated databases must be automatically and continually synchronized.

## *Integrating Applications across Environments and Gaining Portability*



One of the greatest benefits of a well-executed hybrid cloud is being able to deliver applications on the most appropriate platform. A hybrid cloud environment allows you to align business requirements with the location of the application.

For instance, if an application is mission critical and must be available at all times, it should be deployed on the most dependable platform. On the other hand, if an application has a highly variable number of users it may be more appropriate to put it on the most scalable platform.

Take, for example, a fantasy football application — on the weekends the application may need to support millions of concurrent users while during the week very few users will be logged on. In addition, many organizations want the ability to move applications between public clouds, private clouds, and on-premises after the application has been deployed.

For example, a company in Germany may have very different deployment requirements based on internal business rules and government regulations than a company in Brazil or the United States. The ability to move applications has created a strong demand for application portability.

Many businesses don't want to be locked into dependence on a single cloud vendor. The vendor selected to manage a specific application workload may depend on performance priorities. If one provider can't meet the required SLAs or the cost becomes too high, the customer wants the ability to change vendors without expensive transition costs. The requirement has driven customer demand for standards to support portability. Therefore, there's a growing demand for cloud providers that adhere to standards.

Given the fact that users and the business expect an application to perform the same way no matter the underlying infrastructure, moving applications between environments isn't easy. Application portability is more complicated than just saying an application can be deployed in different environments. It's likely that the original environment and the target one have different underlying components. For example, the platforms can have different hypervisors, operating systems, security standards, storage and network models, database systems, and different management tools.

## Chapter 6

---

# Recognizing the Business Impact of Hybrid Cloud

---

### *In This Chapter*

- ▶ Knowing why hybrid matters
  - ▶ Looking at cost versus value
  - ▶ Getting consistent and predictable service
  - ▶ Being flexible in a constantly changing world
  - ▶ Setting the stage for action
- 

**M**ost of the discussions surrounding cloud computing tend to focus on the technology. However, the real value of cloud computing is the way it transforms the delivery of services to support business change and innovation. When you can select the deployment model that best supports the business, you gain tremendous value. Computing, indeed, has evolved from its original intent of storing and reporting on data. Today, hybrid computing environments are intended to help a company experiment with new business models without having to spend money on speculative endeavors. When business leaders have a variety of resources to experiment with different and unanticipated ways of doing business, they have a better opportunity to succeed. The hybrid cloud environment offers a streamlined and innovative foundation for creativity. In this chapter, we discuss the business impact of a hybrid cloud environment.

## Understanding Why Hybrid Matters

Hybrid cloud models require major changes within a company's IT organization. The first step is for management to understand where the valuable assets exist within the organization. Some departments or subsidiaries may use a set of applications that reside in the public cloud. Other business units may have on-premises applications to manage critical customer data. At the same time, IT has long managed complex transaction management systems. Management needs to bring together all stakeholders so there can be a common understanding of company goals. How are all these resources related to reach other? What's the benefit to customers to bring internal and external resources together in a seamless manner?

By focusing on linking resources together rather than isolating them, IT can become the engine of business change. Accomplishing this goal requires management to focus on three key elements:

- ✔ **Initiating a cloud strategy to change the dynamics of IT:** Because of its flexibility, the cloud can become a point of collaboration between business and IT to promote change.
- ✔ **Focusing on the right options based on the task at hand:** There isn't a single approach that works to support all situations. A hybrid approach is pragmatic.
- ✔ **Making resource manageability based on the customer experience a key driver for your strategy:** Customers increasingly expect services to be instantaneously available no matter which deployment model is used.

## Balancing Costs with Value



One of the most important considerations of the hybrid cloud is balancing the cost of services with the value to the business. Often a business unit will select a private cloud because the cost per CPU hour is incredibly low. However, you need to

know the full range of services you require. For example, are you experimenting with a prototype of something that may become a product offering in the future? Or are you simply creating a small set of services that will have localized use? Do you anticipate the need to move vast amounts of data across your enterprise? Will you need backup and disaster recovery services? Will you be including sensitive customer data within the application? After considering these questions, you can see that costs aren't always what they seem at first glance.

Another benefit of a hybrid cloud environment is to have the capability to use the best deployment model in terms of cost that offers the level of service demanded by your customers. Therefore, at times a public cloud service is ideal (testing, prototyping, managing huge volumes of unstructured data that are analyzed on a one-time basis). In other situations, a private cloud may be the best choice, such as when data is too sensitive to security or compliance to risk managing it in a public cloud. In addition, if the application becomes a profit center, it may make sense to keep the application in a private cloud to avoid cost overruns.

## *Getting to the Dial Tone*

There is no doubt that creating a predictable and well-managed hybrid cloud environment is complex. However, a well-architected hybrid cloud should be invisible to the business user. In fact, you may want to think about this hybrid cloud environment as providing the dial tone for the company. When you pick up a landline telephone anywhere in the world, you hear a tone that lets you know that everything is working as it should. Underneath that simple assurance of service is a complex network of services that has been established to support customer needs. While hybrid cloud is still in the early stages, the end goal is the same as with the telephone network. The customer wants to have a consistent and predictable level of service without ever having to see the underlying complexity that makes both the telephone network and the hybrid cloud operational.

## *Looking at the Value of Flexibility in a Fast-Changing World*

Increasingly, companies are beginning to understand that their competitive value in the markets they serve is directly related to the way they empower their ecosystem of customers, partners, and suppliers to be a strategic advantage. It's not enough to resort to the traditional techniques of managing these relationships. At the same time, companies are recognizing that their longevity depends on their ability to innovate and offer new products and services that leverage their intellectual property.

Without thinking about the use of IT in a new way, this transition can't happen. Organizations that are transforming themselves are moving away from thinking about IT as a backroom activity. These companies are putting IT in the forefront of empowering innovation and transformation. What does this mean? In brief, it means leveraging all the IT resources across the organization to create a flexible and scalable platform for innovation and change. This level of change requires a dramatic reinvention of IT. A hybrid cloud strategy that leverages the assets and computing models that exist with a continuum of emerging cloud models is at the heart of this transformation.

## *Setting the Stage for Action*

To change the concept of computing to a hybrid cloud model, management must first understand that the value of information technology is based on the capability to integrate and connect with resources both internally and externally. This method is significantly different than the traditional method of segregating elements of IT within different business units. By focusing on linking resources together rather than isolating them, IT can become the engine of business change. Accomplishing this goal requires that management focus on four key elements:



- ✔ Plan for change so your environment can adapt easily to new business requirements.
- ✔ Initiate a well-architected security and governance model.
- ✔ Create a plan that adheres to standards and open technology.
- ✔ Manage all your data across departments for integration and analytics.





## Leverage a variety of application services with hybrid cloud

Businesses want to implement change in both their technology platforms and their operations. Changing the way business services are delivered to customers requires that technology be available to facilitate that change. What's the conduit for this change? Hybrid cloud. Hybrid cloud maximizes flexibility. Use the right type of cloud service to support the right workloads at the right time. The benefit of hybrid cloud is that it can be used to provide a continuous delivery platform for building, deploying, and managing cloud applications.

- **Define hybrid cloud** — describe the key connection points for the hybrid cloud
- **Discover new technologies and approaches** — create business value and flexibility
- **Adopt emerging technologies** — use OpenStack and containers to achieve portability of both data and application
- **Use sophisticated cloud applications** — gain the right services to efficiently and effectively deploy the applications your users and customers need



Open the book and find:

- **Business imperative for hybrid cloud**
- **Introduction to cloud delivery models**
- **How to manage security and governance in a hybrid cloud environment**
- **How to manage data and analytics in a hybrid cloud**
- **Why hybrid matters**
- **Cost and value analysis**

Go to **Dummies.com**

for videos, step-by-step examples, how-to articles, or to shop!



# **WILEY END USER LICENSE AGREEMENT**

Go to [www.wiley.com/go/eula](http://www.wiley.com/go/eula) to access Wiley's ebook EULA.