

Brought to you by:



IBM Cloud Private

for
dummies[®]
A Wiley Brand

Understand the value
of private clouds

Discover the open-
source cloud platform

Learn the importance of
hybrid environments



Judith Hurwitz

Daniel Kirsch

UK Special Edition



IBM Cloud Private

UK Special Edition

**by Judith Hurwitz and
Daniel Kirsch**

**for
dummies[®]**
A Wiley Brand

IBM Cloud Private For Dummies®, UK Special Edition

Published by

John Wiley & Sons, Inc.

111 River St.

Hoboken, NJ 07030-5774

www.wiley.com

Copyright ©2018 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, or visit www.wiley.com/go/custompub. For information about licensing the *For Dummies* brand for products or services, contact BrandedRights&Licenses@Wiley.com.

ISBN: 978-1-119-51192-2 (pbk); ISBN: 978-1-119-51193-9 (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. Burchfield

Editorial Manager: Rev Mingle

Acquisitions Editor: Steve Hayes

Business Development

Representative: Sue Blessing

IBM Contributors: Mohammed

Abdula, Nancy Agosta, Roland

Barcia, Andy Gower, Sanjay

Joshi, David Lindquist, Eduardo

Protacinio, David Slovensky

Production Editor: Siddique Shaik

Table of Contents

INTRODUCTION	1
About This Book	1
Foolish Assumptions	2
Icons Used in This Book.....	2
CHAPTER 1: Explaining the Private Cloud	3
Setting the Context for the Private Cloud	4
The Rationale for the Private Cloud	5
Private Cloud Requirements	5
Dynamic provisioning.....	5
Workload portability.....	6
Security management	7
Integration and connectivity to access resources.....	7
Management of a hybrid environment.....	7
Defining IBM's Cloud Strategy	8
CHAPTER 2: How Customers Use the Private Cloud	11
The Business Imperative for the Private Cloud	11
Creating Cloud Native Applications.....	13
Integrating with Data and Applications	14
Moving Existing Workloads to the Cloud.....	15
Refactoring and Modernising	15
CHAPTER 3: Looking at the Technical Foundation of IBM Cloud Private	17
Features of IBM Cloud Private	18
The Core Components of IBM Cloud Private	18
Flexible choice of infrastructure	19
Kubernetes-based container platform.....	19
Cloud-optimised software and services.....	20
Integrated DevOps and management tools.....	21
CHAPTER 4: DevOps for the Private Cloud	23
The Changing Nature of DevOps.....	23
Continuous Integration and Delivery (CI/CD).....	25
The Role of IBM UrbanCode in CI/CD	26
Delivering a Better Customer Experience	27

CHAPTER 5:	Analysing Data in the Private Cloud	29
	The Value of Data Analytics.....	29
	The Role of Data Science in Data Analytics	30
	Understanding Machine Learning.....	31
	The role of the private cloud in advanced analytics	31
	Hybrid data management.....	32
	Advanced Analytics to Support the Hybrid Cloud	33
	Leveraging existing on-premises data sources.....	33
	Cloud native advanced analytics.....	34
	Hybrid cloud support for advanced analytics	34
CHAPTER 6:	Managing a Multi-Cloud Environment.....	35
	Setting the Stage of Managing Hybrid Computing.....	35
	Gaining Visibility	36
	Understanding the Rationale for Multi-Cloud Management	36
	Requirements for Managing the Multi-Cloud Environment	37
	Service catalog.....	37
	Configuration management	38
	Service automation and orchestration	38
	Creating a Balanced Platform with API Management.....	39
	Managing in a Microservices World	40
	IBM Cloud Automation Manager.....	40
	Terraform.....	41
	Service catalog within IBM CAM.....	41
	Service composer.....	42
	Monitoring a Multi-Cloud Environment.....	42
	Support for Multiple Monitoring Tools.....	42
	Setting the Stage for Action	43

Introduction

Hybrid cloud is redefining the way software is being built and executed within the organisation to support employees, customers and partners. This industrialisation of cloud computing has significant implications for businesses. Cloud services now must have the same resiliency, security, scalability and performance as the traditional data centre. Organisations need to take full advantage of both public and private cloud services to support a variety of workloads ranging from data and analytics to containers and microservices. This hybrid environment must be designed so the computing infrastructure is well managed.

As more businesses are focused on the cloud, they are looking to the private cloud as a way to ensure that they optimise their operational approach to leveraging cloud services. Combining the private cloud with public cloud services provides organisations with choice depending on the business requirements for security, compliance, elasticity and performance. The adoption of open standards such as Kubernetes, Cloud Foundry and Open Stack makes it possible to increase the portability of data and application services.

About This Book

As the hybrid cloud matures to support the enterprise, there is a requirement to provide a consistent and predictable foundation for enterprise computing. There is a need to support the development, deployment, management and control of the overall environment. While many business units have taken advantage of the public cloud because of its ease of provisioning services, many other organisations are looking pragmatically at the private cloud to support mission critical workloads that need to be protected behind the firewall. Enterprises require workload portability, security, integration and connectivity and overall management of workloads.

That's where this book comes in handy. *IBM Cloud Private For Dummies*, UK Special Edition, provides insights into the role of the private cloud and how it supports the changing requirements for computing. The book focuses on IBM's private cloud offering – IBM Cloud Private – that works in concert with IBM's public cloud services.

Foolish Assumptions

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

- » You're already familiar with cloud computing and need to understand the role of the private cloud and how it relates to your data centre and to the public cloud.
- » You're planning a long-term cloud strategy and want to understand the value of the private cloud and how it can be used to support your business goals.
- » You want to understand the IBM Cloud Private offering and how it can support your enterprise goals. You also need to understand how it works with both IBM Cloud and other public cloud services.
- » You want to understand how all the elements of cloud computing fit together and support software development, deployment, security, compliance and workload management.

Icons Used in This Book

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



REMEMBER

This icon highlights important information that you should remember.



TIP

Tips help identify information that is beneficial to you in some way. These tidbits might help you save time, money or a headache.



WARNING

This icon points out content that you should pay attention to in order to avoid problems.

- » Setting the context for private cloud
- » Understanding the rationale for the private cloud
- » Looking at private cloud requirements
- » Getting to know IBM cloud strategy

Chapter 1

Explaining the Private Cloud

Private cloud is emerging as an imperative for enterprise customers. Businesses require the flexibility and elasticity of the public cloud with the manageability, structure, architecture, control and security already supported within their own data centres.

The scope of cloud computing is changing with an emerging set of tools and approaches for managing workloads in a predictable manner. Depending on workloads and circumstances, companies are looking to a multi-cloud and hybrid set of services to meet their changing needs. There is a major transition happening with cloud computing as it matures: We can no longer simply think about an individual service; instead we need to think about the overall life cycle of computing based on cloud.

In this chapter, we explain the evolution of private cloud from a life cycle perspective and how public and private cloud services support changing workload requirements. This explanation is framed in terms of IBM's cloud strategy. We also introduce you to IBM Cloud Private.

Setting the Context for the Private Cloud

Just a few years ago, businesses discovered that cloud services were useful in providing a streamlined way for easily provisioning compute and storage services. The financial benefits were immediate and obvious. Many business leaders opted for the public cloud because of the ease of provisioning services. Other organisations, concerned about security, compliance and guaranteed service levels, insisted on leveraging private infrastructure services such as automated provisioning within their own data centre firewalls.

Over the past five years the market for cloud services has exploded as more and more businesses began to understand the possibilities of utilising cloud-based services as the foundation for computing.



TIP

Moving from simply using some compute or storage services to a full stack and infrastructure for cloud computing is a necessity. It's feasible to have a single way for supporting all workloads. In addition, there is an imperative to apply the same standards for accountability across workloads in a multi-computing environment.

While there have been debates about the value of a public versus a private cloud, these two deployment models are evolving so they can work collaboratively with each other. Think of the public cloud as the standardised Internet-based pay-for-service available on demand. In contrast, think of the private cloud as a set of software defined services with the same common application programming interfaces (APIs) provided by the public cloud. The emerging private cloud provides a set of modular life cycle services.



REMEMBER

The private cloud's common services reside in a company's data centre and have the APIs that allow those services to work with existing capabilities, including virtual machines, containers or full applications.

Increasingly, public cloud vendors are adding more enterprise services in order to link public cloud resources to line of business applications in the data centre. Likewise, private cloud services often take advantage of a variety of public or hosted services. What is the difference? Typically, public cloud services offer the ability to bring together a variety of services at time of need. In contrast, a private cloud environment is a more integrated ecosystem of services. Many organisations have discovered that both models are effective, depending on the nature of the workloads.

The Rationale for the Private Cloud

The architectural knowledge and programming ability is steep in a public cloud environment today. Over time, we expect that the public cloud will gain more abstraction to make public services act as though it were a data centre.

To be successful, organisations are discovering that they need a combination of public and private cloud services to meet the needs of their customers. For example, a born-on-the-cloud application without dependencies for complex services, or a service level agreement (SLA) can be developed in the public cloud. In other cases, an organisation wants to leverage existing core applications that can't be easily redesigned. At the same time, organisations often demand the performance of these integrated services, granular control for audit and performance, and indemnification. Organisations will typically have both of these requirements, resulting in the need for both public and private services. For example, you may run an application in the private cloud but use public cloud based video services or public datasets.

Private Cloud Requirements

A number of requirements are needed to effectively create and manage a private cloud. These services are listed in this section.

Dynamic provisioning

One of the benefits that developers like about the public cloud is the ability to provision compute and storage on demand. IBM Cloud Private, covered later in this chapter in the section 'Defining IBM's Cloud Strategy,' provides self-service provisioning. However, a difference does exist. The private cloud enables provisioning to be executed dynamically based on rules about the conditions that should trigger the addition of more compute or storage. These rules can also dictate where those resources should be placed. For example, an application requires additional storage.



REMEMBER

The rule in the private cloud may be established to select the most appropriate type of storage at a location that meets the compliance requirements for the enterprise. In addition, the rules can indicate how much money should be spent on that storage. For example, a pilot project built in the cloud may not need the

same speed of storage as that application will need once it's put into production. When you add machine learning algorithms to dynamic provisioning, you have the ability to trace the way workloads are managed. It can also provide the ability for self-healing of workloads that aren't executing according to requirements.

Workload portability

One of the most important imperatives for organisations is to be able to move workloads from one environment to another as their requirements change. For example, a workload might initially reside in the data centre. However, if the application begins to require more compute and storage resources, the organisation may decide to move it to a public cloud rather than acquiring more on-premises resources. Applications need to be designed for flexibility and change.



TIP

The value of a hybrid computing environment is that it provides the right mix of performance based on the use case. Complex workloads may benefit from on-premises API content that requires low latency and fast access to resources. You may have the requirement to use a private cloud facility and networking access to a set of on-premises resources. Workload portability also takes advantage of the modularity of containerisation because it supports polyglot development where developers use the language most appropriate to support the workload (issues related to DevOps are discussed in more detail in Chapter 4). One of the most important benefits of containerisation is that it helps developers abstract workloads into containers that have different underlying structure and performance characteristics. This approach means that the differences in language and resources don't impact performance and mitigate technical conflicts.

To make applications or application services portable requires that code be transformed so a variety of services can be linked together at runtime. This level of change is instrumental in fulfilling customer expectations to change services as new ways of conducting business emerge. After you modularise application services so they can be linked together, you have to be able to manage those workloads so they perform as expected without latency. Workloads need to be balanced. There need to be appropriate levels of security and compliance as the use of workloads change. These requirements such as load balancing, management and security are among the foundations of the modern private cloud. When

executed well, the private cloud has the predictability of the data centre with the agility of the public cloud.

Security management

Security is one of the top concerns for most enterprises. In many situations, security and compliance mandates are in place to protect sensitive customer data. In addition, protecting an organisation's intellectual property is critical. In the highly distributed environment of the cloud, security has to be implemented in a way that can change and evolve based on the use and demands of the workloads. These workloads may touch many different networks. In other cases, the workloads may integrate with data that is distributed across a multi-cloud environment.

Integration and connectivity to access resources

A private cloud environment must be able to connect to the needed data and application services across both the data centre and the public cloud. Services and data may need to be moved across networks or an algorithm may need to be executed at the source of the data. While many tools are available to provide connections between services, the private cloud typically requires the ability to manage latency, rules, access and authorisation, predictability, dependencies, and management of the connected services.

Management of a hybrid environment



REMEMBER

The private cloud is not a single unified environment typical of the data centre. The private cloud requires a software-defined layer that removes the dependencies connected with a single hardware or virtualised infrastructure. IT and business unit leaders need to have a consistent and predictable way to manage all these disparate services as if they were part of a single integrated system. A well-designed private cloud demands an intuitive method to determine how services are operating and monitor where the bottlenecks are in this distributed environment.

Defining IBM's Cloud Strategy

IBM's cloud strategy is to enable the hybrid enterprise. IBM offers two significant implementations of cloud: the IBM Cloud – a public cloud offering – and IBM Cloud Private.



REMEMBER

IBM Cloud Private is an application platform for developing and managing on-premises, containerised applications. This integrated environment includes the container orchestrator Kubernetes, a Cloud Foundry environment, a private image repository, a management console and monitoring frameworks.

The design of the IBM Cloud Private is fourfold:

- » A customer can leave an existing data centre application as is and provide a set of APIs that enables the flexibility of managing and monitoring workloads that interact with those existing applications in a more versatile and scalable manner.
- » The software-defined services enable workloads to be refactored so they are containerised without dependencies. These container-based services are then linked together with a set of consistent APIs. These are the same APIs that IBM has implemented in the public cloud.
- » These same container-based services are used to create cloud native applications that leverage common APIs and can easily connect with highly distributed data and services. For a deeper dive into the details of the IBM Cloud Private architecture, take a look at Chapter 3.
- » The environment provides the ability to apply advanced analytics, including machine learning, artificial intelligence and cognitive analysis to the cloud. This enhances the manageability and use of analytics in solving business problems.

A key design point of IBM's cloud strategy is to create a set of services that are designed to be linked together and managed through consistent and well-defined APIs.

IBM has exposed a consistent set of APIs across both its public and private cloud offerings. The IBM cloud strategy is based on a software-defined infrastructure. This means that the cloud can support whatever environment is already in place, including virtualisation, bare metal, containers and automation.



The IBM Cloud Private is logically a control plane on an existing data centre with connectivity to hosted and public cloud services. Therefore, IBM has designed a set of services that are layered on top of infrastructure services. This approach meets clients' many different needs based on the type of workloads that they need to support.

EMBRACING OPEN SOURCE

Open source pervades the IBM Cloud Private architecture. This is a pragmatic move since customers demand the ability to support multiple cloud architectures and want the ability to have portability and modularity. Open-source foundations are the basis of both IBM Cloud and IBM Cloud Private. For example, IBM supports an open container strategy. This means that IBM supports Kubernetes, Docker and Cloud Foundry, and popular programming languages, such as Java, Python, R, Node.js, and database implementations, such as MongoDB and Postgres.

The benefit to leveraging open source is clear; it enables IBM to provide a well-understood and portable set of tools to developers. In addition, given the breadth and depth of IBM service offerings, IBM can provide the support needed to use these open-source offerings.

IN THIS CHAPTER

- » Discovering the business need for private clouds
- » Creating cloud native applications
- » Securing and integrating data and application services
- » Extending the life of important existing applications
- » Refactoring and modernising for agility

Chapter 2

How Customers Use the Private Cloud

While the benefits of cloud computing are clear, businesses are increasingly finding the need to take different approaches to the cloud depending on the nature of each individual workload. For example, business leaders may want their workloads behind their firewall for security and compliance reasons. In addition, businesses want to leverage their data centre investment in their line of business (LoB) applications. In this chapter, we present key use cases and ways that companies are using the private cloud to meet their changing business needs.

The Business Imperative for the Private Cloud

As businesses begin to put cloud at the centre of their computing strategy and planning, the private cloud is becoming a core element. The bottom line is that not all workloads are the same. Various reasons and situations make a difference in terms of which platform is the most appropriate. The most obvious issue

for many businesses is security and compliance. Many companies want a guarantee that they have complete control over the security of their data centres. While a public cloud may be as secure as a data centre, IT leaders want to have direct control. Additionally, although many public clouds will enable you to store workloads based on compliance requirements, some organisations require a service level agreement (SLA) that guarantees adherence to policy.

Indeed, many existing workloads tend to leverage a variety of services. Some of these are tightly coupled within the data centre environment. Other services include dependencies on operating system and large code blocks that can't be easily transformed in the short term. For example, there might be a scheduler that determines when bills should be sent out to customers based on the products and services they purchase. This older code may include complex authorisation rules.



WARNING

Despite the desire to move away from these monolithic services, it is not cost effective to execute a transition all at once. Organisations need the ability to stage their transition in a way that protects the company's ability to deliver satisfaction to customers. Initially, an IT organisation might containerise applications or large components of a workload. This provides better agility in the interim. The next stage is to do an assessment in order to understand the costs and requirements for redevelopment. For more details on containers, head to Chapter 3.

The other major issue regards latency requirements. Many applications have dependencies for on-premises applications and data sources. If an organisation were to move these complex environments to the public cloud, there could be an impact on the speed of execution of a process or query. However, at the same time, it is important to be able to transform the data centre into a more agile computing model by transforming computing services. This transformation requires that underlying middleware, networking services, catalog services and rules engines be redesigned. Once these capabilities become modular services, it's possible to redefine computing as a web of linkable modules that can work together when needed. Because these services have been redefined, it is then possible to consume platform and infrastructure services in a more flexible as-needed manner. If more compute or storage is needed, it is not necessary to transform the data centre. Instead, to meet increased requirements, the IT infrastructure will be redefined so that additional compute or storage can be added on an on-demand basis from a public cloud.



TIP

Combining on-premises, private cloud services with public cloud services can foster a new level of pragmatism to creating customer value for your company. Many organisations have found that the traditional data centre model has been too rigid for changing business needs. The private cloud combined with public cloud services addresses many of the requirements that are mandatory for this new era of digital transformation.

Creating Cloud Native Applications

Many organisations want to the gain agility of cloud native applications but want to maintain control so applications can be deployed within the company's firewall. These organisations often have security and compliance reasons for a private cloud. In addition, many organisations want to have total control of their own infrastructure and don't trust their data to a multi-tenant environment. However, to be competitive, these companies need to create a new generation of applications.



REMEMBER

Cloud native applications are emerging as the most effective way of creating applications that are modular and scalable. The characteristics of cloud native applications needed are as follows:

- » **Applications are composed of multiple services.** These applications are typically built based on a microservices architecture and managed within a container framework.
- » **Cloud native applications have the benefit of enabling the enterprises to integrate services and data from third parties.** Cloud native applications have APIs that allow developers to easily access other resources. These APIs allow developers to more quickly create new capabilities.
- » **Each part of a cloud native application is elastic.** If some services need more resources, the cloud native application can acquire them without causing the overall application to slow down or fail.
- » **Native cloud applications are composable with a variety of microservices based on APIs.** This characteristic helps support a continuous delivery model so software can be frequently updated.

RESPONDING TO NEW COMPETITION

A wealth management company's business is being challenged by emerging born-on-the-cloud Fintech companies. As a result, its customers want the same intuitive and easy-to-use applications that competitors have created. The company decided to create cloud native applications in a private cloud to connect with on-premises applications as well as third-party services. The goal would be to continuously update the application as customers' expectations increased. The company wanted to be able to create a modular software approach that would help developers create software services based on standard APIs. The company used containers and a continuous integration DevOps approach to building new code. The resulting applications leveraged a new generation of modular middleware to easily connect with on-premises and public cloud services.

Integrating with Data and Applications

One of the benefits of a hybrid cloud environment is the ability to link to data and applications across internal and external systems. Even when companies adopt a cloud native application development and deployment approach, they will often need to reach out and efficiently leverage third-party data services, such as weather and demographic data. At the same time, there is a requirement to integrate with internal data sources such as transactional customer data and product inventory systems. Organisations can leverage core cloud services based on Watson APIs.

Data analytics is a core requirement for organisations moving to the hybrid cloud. In many situations, analytics is at the core of being able to provide predictive analytics into a business application. In addition, data scientists have determined that the private cloud is an ideal platform for creating and managing machine learning models. Often this data is so sensitive that organisations have a mandate to keep this data behind the firewall. For more details on data science and machine learning in the private cloud, check out Chapter 5.

MANAGING PRICING

A hospitality company wanted to come up with a more flexible pricing model that would take into account a variety of factors, such as weather and social media information about upcoming events in the area. The new application for the private cloud would leverage emerging machine learning techniques along with analysis of unstructured data. It would also integrate with cloud-based analytics services as well as with cloud-based data sources. The private cloud enables the business to keep its unique data behind the firewall. At the same time, developers can create new innovative services that can be integrated into the platform through APIs. The cloud environment offers more performance, scalability and modularity needed to compete as customer requirements change.

Moving Existing Workloads to the Cloud

What does it mean to move an existing application to the cloud? In a perfect world, it may be preferable to transform existing applications to a modern architecture. However, in some situations, you simply don't have the time to re-architect existing applications. Moving an existing application to the cloud allows those applications to gain access to use APIs to connect to a variety of cloud services and leverage the scalability and modularity of the cloud infrastructure. Some applications either are too brittle or are packaged applications that can't easily be changed. These applications can be shifted to the cloud as is. Other software such as middleware services can be moved to the cloud because licences (such as WebSphere) can be moved without additional cost to the cloud. For some companies, the migration of existing application to the cloud can be the first step toward a more flexible approach to computing.

Refactoring and Modernising

Ideally, organisations would be able to build all new applications that are streamlined and modern by design. The reality is that organisations have significant intellectual property embedded

inside applications that may have existed for decades. In many situations, these applications are the systems of record that are the engine for business continuity. Included in these applications are core services such as billing, transaction management, customer management and so on. For businesses with thousands of customers, there's no pragmatic way to start from scratch. Business rules and best practices are baked into these applications. Codified processes are based on compliance and governance rules.

However, many of these traditional applications are problematic. Many were written as integrated code where data, rules and processes were tightly integrated. Often it is difficult to update software as the enterprise adopts new business methods. Some organisations without the time or resources may decide to lift an existing application and move it to the cloud. It is possible to move middleware licences such as WebSphere from on premises to the private cloud so it can support the applications. This process is typically called *lift and shift*. While this process is pragmatic, it is only a respite from solving the problems associated with inflexible code.

The private cloud infrastructure can take advantage of microservices and container technology to begin transforming and modernising existing applications.

REDUCING CAPITAL EXPENSES

A manufacturing company was faced with emerging competitors that were able to customise products using a new set of applications. These competitors were able to anticipate customer preferences and optimise product customisation. The manufacturing company decided that moving to a more flexible services environment would help it compete. Developers spent time selecting core services from existing applications and creating new microservices. These new services were placed into containers with well-defined APIs. At the same time, other developers were creating new services needed to modernise the company's ability to offer more customised offerings to customers.

- » Looking at the features of IBM Cloud Private
- » Understanding the core components of IBM Cloud Private

Chapter 3

Looking at the Technical Foundation of IBM Cloud Private

IBM Cloud Private is a prepacked, enterprise-class solution that delivers a single platform located behind your firewall. You can leverage your on-premises software portfolio or easily integrate next-generation data and software optimised for cloud. IBM Cloud Private is built on open-source frameworks, including containers, Kubernetes and Cloud Foundry and offers flexibility, control, security and easy integration with public clouds. In addition, cloud management tools are included with IBM Cloud Private so you can govern and gain insight into multi-cloud infrastructures and applications.

IBM Cloud Private was developed based on IBM's expansive experience helping developers build cloud-native applications and refactoring monolithic applications into more flexible microservices-based applications.

In this chapter, we give you insight into what you have access to as an enterprise with IBM Cloud Private and we help you understand the core components of IBM Cloud Private.

Features of IBM Cloud Private

Organisations that adopt the IBM Cloud Private platform have access to

- » A unified installer to rapidly set up a Kubernetes-based cluster with master, worker and proxy nodes
- » The IBM Cloud Private cluster management console to manage, monitor and troubleshoot applications and clusters from a single, centralised and secure console
- » A private Docker image registry to provide a local registry service that functions in the same way as the cloud-based registry service in Docker Hub
- » Isolated tenant networks with Calico for improved performance and network isolation inside clusters
- » Robust monitoring and logging with ELK stack for collection, storage and querying of logs and metrics
- » The app centre as a centralised location to browse and install packages into clusters
- » Security management capabilities for a broad range of tools and services

To learn more, visit www.ibm.com/cloud-computing/products/ibm-cloud-private.

The Core Components of IBM Cloud Private



REMEMBER

IBM Cloud Private is not a monolithic platform; instead, it's a set of component services and infrastructure. The common denominator is that all the services are intended to leverage customers' existing infrastructure, software and services as well as providing open source at the core. All the services that comprise the private cloud are architected to operate in a collaborative manner as a unified system. IBM Cloud Private is comprised of the key components in this section.

Flexible choice of infrastructure

Most businesses don't own a single hardware or software platform in their data centres. Instead, they have a variety of systems and components that are core to day-to-day operations. Additionally, many different public clouds are used across hundreds of business units. The IBM Cloud Private environment can operate on any existing hardware environment, including Intel-based platforms and IBM Power Systems. IBM Cloud Private also supports OpenStack, VMWare, IBM Z, IBM Storage, IBM Hyperconverged Systems and third-party cloud providers. IBM Cloud Private is also designed to be compatible with leading systems manufacturers like Cisco, Dell/EMC, Lenovo, Intel and NetApp.

Kubernetes-based container platform



REMEMBER

Containerisation is at the foundation of how the private cloud is designed. The value of containerisation is clear. By placing code into containers, it is possible to create a service that includes all the dependencies needed to operate that service. A standardised application programming interface (API) supports each container. Therefore, a container may be used to manage newly built microservices or to refactor an existing application or service in new ways through encapsulation into a container.

IBM Cloud Private delivers the choice of open-source application runtimes, consistent with IBM public cloud offerings, Kubernetes and containers or Cloud Foundry technology. Customers can choose the prescriptive development approach of Cloud Foundry or the more customisable and portable approach of Kubernetes and Docker containers. Consistency with the IBM public cloud gives a choice of deployment models across the life cycle of workloads and a more consistent management experience as the workload matures and evolves.

IBM Cloud Private has also containerised middleware, data and analytics services. It also includes a unified installer to rapidly set up a Kubernetes-based cluster with master, worker and proxy nodes by using an Ansible-based installer. With containers as the core, IBM provides the same open-source application runtime services across both the IBM public and private cloud.



Containers manage the way both new and existing code can be combined to behave in a consistent and predictable manner. Using container architecture enables orchestration and management of applications and their services. Containerisation also supports the ability to port applications to different cloud platforms based on the need of the business.

Cloud-optimised software and services

One of the benefits of the container architecture featured in the private cloud is the ability to create a multi-tenant version, enabling you to set up different versions of a private cloud environment for each of your IT stakeholders. Each stakeholder gets a customised set of services within the private cloud based on a designed business process behind the firewall. Isolated tenant networks leverage Calico (a network policy service for Kubernetes clusters) to improve performance and network isolation inside clusters.



To ensure consistency, the private cloud provides a common catalog of services to accelerate developer productivity. Its common base services enable scalable management of microservices. In addition, the catalog enables applications to automatically scale both horizontally and vertically. Catalog services are core to managing and maintaining software services such as microservices and containers. The catalog structure makes it easier to govern, deploy and maintain software and services to support rapid development and test. Key services used to manage catalogs include Helm charts, Terraform templates and Cloud Foundry buildpacks.

IBM Cloud Private provides a broad range of easily accessible, consistent management and operational middleware, data and analytics services for both new cloud native applications as well as existing ones. New Kubernetes-based cloud native middleware and data platforms, such as Microservice Builder, Data Science Experience and API-Connect, work together with IBM Cloud Private to deliver the next generation of microservices. IBM Cloud Private enables the use of cloud-enabled, enterprise-grade middleware to leverage existing application development skills such as Java, Spring and Open Liberty. It also provides the API connectivity and management needed to integrate public, private and existing enterprise capabilities.



TIP

IBM Cloud Private includes capabilities to support development of new microservices-based applications, evolve and modernise existing applications and integrate new and existing applications. Along with the application runtime frameworks, IBM delivers a core set of management services for these frameworks and the applications being deployed on top. A common catalog of enterprise and open services is provided in order to accelerate developer productivity. This catalog is comprised of easily deployed and maintained software to enable developer, test and production use cases.

The application services such as middleware, data services and security services are key to making the private cloud operational and scalable. This set of common base services allows scalable management of microservices, including Istio (an open platform to connect, manage and secure Kubernetes), monitoring with Prometheus (open-source monitoring and alerting toolkit) and logging with the ELK Stack (logging with three popular tools: Elasticsearch, Logstash and Kibana).

IBM Cloud Private is designed to provide an end-to-end solution for applications, including popular open-source frameworks and languages, built-in DevOps, integrated monitoring and industry-leading enterprise solutions from the IBM middleware, data and analytics portfolios.

Integrated DevOps and management tools

Integration and management services are core to creating a cloud environment that behaves as though it were a single unified system. For example, all the services within the private cloud must have security at the core. These services include authentication, authorisation and identity support. They also include alerting, auditing, security tools for scanning Docker images and containers and other services.



REMEMBER

IBM Cloud Private includes a centralised, secure cluster management console to manage, monitor and troubleshoot applications and clusters. A new set of cloud management tools are integrated with existing enterprise management tools. Key capabilities included for cloud management are

- » API access to core platform to enable high-levels of automation needed for Continuous Integration & Continuous Delivery (CI/CD)
- » Integration of IBM and open-source DevOps tools and services to create code repositories, deployment, test automation and promotion to production
- » New capabilities to build, deploy and manage microservices-based applications, including Microprofile, Spring, Node, Microservice Builder and Istio
- » Integrated security assessment capabilities in the DevOps toolchain (DevSecOps with IBM Vulnerability Advisor for Containers)
- » Analytics-driven monitoring and logging capabilities built on open and extensible frameworks (Prometheus, Istio, ELK, Grafana)
- » Catalog of easily governable, deployable and maintainable software and services enabling rapid development and test (Helm charts, Terraform templates, Cloud Foundry buildpacks)
- » Best practice solution architectures and methods and IBM Cloud Garage services to enable enterprise transformations

THE PRAGMATISM OF A HYBRID ARCHITECTURE

Every organisation has a vast number of applications, data, platforms, networking, management tools and security requirements. Hybrid cloud architectures enable businesses to use their existing investments combined with modular, scalable, flexible services in a way that performs based on customer expectations. Combining the best of public and private cloud services is the most pragmatic approach to support constantly changing needs to innovate.

- » Looking at the changing nature of DevOps
- » Enabling continuous integration and continuous delivery
- » Understanding the role of IBM UrbanCode in CI/CD
- » Delivering a better customer experience

Chapter **4**

DevOps for the Private Cloud

The cloud is transforming the way organisations approach software development, deployment and operations. One of the key benefits of cloud architecture is that it abstracts application logic so applications can be dynamically designed using a new generation of cloud-based application development tools. In this chapter, we explain the shifts occurring around DevOps as they relate to cloud adoption and how companies can evolve their DevOps strategies to better align to private cloud implementations.

The Changing Nature of DevOps

The world of software design, development and deployment is rapidly changing with the evolution of cloud architectures and implementations. Modularity is at the core of the changing nature of the DevOps processes as it relates to code build,

deployment and release. Close alignment of DevOps processes with private cloud implementations is critical due to newly emerging changes:

THE IBM GARAGE METHOD

Having a strong DevOps culture is key for a successful IBM Cloud Private solution. The IBM Cloud Garage Method is IBM's approach to enable business, development and operations to continuously design, deliver and validate new functions. The practices, architectures and toolchains cover the entire product life cycle from inception through capturing and responding to customer feedback and market changes.

The IBM Garage Method consists of the following areas:

- **Culture:** Transform by combining business, technology and process innovations that help create teams that quickly learn from market experiences.
- **Think:** Incrementally deliver innovative solutions by using IBM Design Thinking and related design practices.
- **Code:** Adopt development practices to build cloud-native applications, release incremental function, gather feedback and measure results.
- **Deliver:** Accelerate time-to-market by using continuous integration, continuous deployment and automating repeatable and transparent processes.
- **Run:** Run highly available solutions on a cloud platform by using Cloud Foundry, containers or virtual machines. Run on a public cloud, a dedicated cloud, a private cloud or in a hybrid environment.
- **Manage:** Provide operational excellence with continuous application monitoring, high availability and fast recovery practices that expedite problem identification and resolution.
- **Learn:** Continuously experiment by testing hypotheses, using clear measurements to inform decisions and driving findings into the backlog so you can pivot.

For further information on IBM Cloud Garage method, visit www.ibm.com/cloud/garage.

- » Increasing adoption of building, deploying and managing microservices based applications across clouds requiring scalability and variable demand management
- » Availability of new, integrated DevOps tools and services for code repositories, deployment, test automation and promotion to production in cloud environments
- » Use of container technology for orchestrating microservices into applications

Given the profound changes in the approaches and architectures of cloud technology, DevOps organisations are now shifting focus to continuous integration and delivery in agile development environments. The need to drive down costs while speeding time to market and reducing risk is creating urgency to deploy solutions that combine DevOps processes with the private cloud model. The need for self-service, modularity, elasticity and performance is the foundation for the next generation of software design, development, deployment and operations.

Continuous Integration and Delivery (CI/CD)

Applications and microservices are being constantly updated as businesses create new services to support changing customer expectations. Therefore, the cloud has become the most effective and pragmatic vehicle for application development and deployment. Continuous integration and continuous delivery (CI/CD) techniques enable teams to deliver software faster with less risk. By emphasising the rapid delivery of shippable code through automated builds, tests and deployments, companies can realise the value of a continuous delivery model more rapidly.

With continuous delivery techniques, the business needs to quickly understand whether changes to applications have had the desired outcomes. The process involves frequent building and unit testing, frequent deployment and integration testing and ongoing release of code to ensure customers are satisfied with the changes.



Part of the process of CI/CD is the use of pipelines – a sequence of automated tasks that enable organisations to build, test and deploy services to the targeted system. The value of pipeline services, such as Jenkins, is the ability to provide developers with a view of the entire DevOps process.

Many organisations now have multi-cloud topologies that can make slow, error-prone manual code deployments untenable. Yet, when the IT organisation is required to support different data centres and multiple clouds, even scripted deployments can be risky, hard to track and expensive to maintain. As a result, multi-cloud topologies require deployment automation.

IBM uses open technologies like GitHub Enterprise and Jenkins to build CI/CD architectures. You can find these architectures at <https://goo.gl/1XU5qt>.

In one example, the pattern uses the Jenkins Helm Chart to install a Jenkins Master pod with the Kubernetes Plugin in a Kubernetes Cluster. Helm is Kubernetes's package manager designed to facilitate the deployment of prepackaged Kubernetes resources that are reusable. This set-up allows Jenkins to spin up ephemeral pods to run Jenkins jobs and pipelines without the need of always-on dedicated Jenkins slave/worker servers. This approach reduces Jenkins infrastructure costs. IBM Cloud Private also provides accelerators such as the Microservice Builder that help generate Jenkins Files for microservice applications.

Jenkins can be used in conjunction with IBM Cloud Private in a variety of situations. In addition, Jenkins can be used as a comprehensive deployment automation solution across public and private clouds. It can also be used to deploy decoupled containerised services that may be necessary to orchestrate containers with virtual machines.

The Role of IBM UrbanCode in CI/CD

The IBM UrbanCode suite of offerings is designed to support CI/CD. The UrbanCode Build, UrbanCode Deploy and UrbanCode Release products facilitate rapid feedback and continuous delivery in agile development environments. The suite also provides audit trails, versioning and approvals needed in production. These

offerings are intended to make releases more predictable, help reduce time between concept, rollout and feedback and accelerate time to market.

UrbanCode Deploy, for example, is an IBM release automation solution used by the DevOps team to put pipeline-automated tasks into production. It can scale up to enterprise-class deployments, handling thousands of servers. UrbanCode Deploy enables IT to deploy services to a heterogeneous deployment such as a public or private cloud. It provides the ability to audit and trace what is happening with application code. This is especially important when moving code from test into production. It is imperative in a hybrid cloud environment where application code is continuously changed to understand what is happening under the hood. For example, what changes have been made to code? Who has made those changes? Is the code secure or has someone opened a back-door into the code that could compromise security? UrbanCode Deploy provides auditability of code so the DevOps team can track the integrity of code.

Learn more about IBM UrbanCode at developer.ibm.com/urbancode.

Delivering a Better Customer Experience

Think about an airline and what it takes to provide seamless service to get passengers safely and quickly to their destinations. In a perfect world, customers purchase tickets, show up at the terminal, get on the plane and go. But, as we all know, things tend to be a lot more complicated. There are mechanical problems, scheduling conflicts, flight cancellations, delays on the runway and large natural disasters such as hurricanes and snowstorms. Airlines might have to reschedule hundreds of flights, rebook thousands of customers and reroute thousands of tons of luggage.

One airline decided that to be successful it needed to dramatically streamline the DevOps process to support the need for continuous integration and delivery of microservices. A key objective was to improve the experience for the airline's most valued customers. Monolithic applications were broken up into microservices through the use of the Microservices Builder. Some code that couldn't be rewritten was encapsulated. Microservices were

automated through pipeline services. In addition, microservices were combined into containers that were orchestrated to support a new, self-service model. Applications built on the private cloud were extended to leverage key services, like weather data, so that passengers could rebook flights while having visibility into weather conditions. The services were tested and put into production.

The results were significant. The airline was able to turn its core services such as seat booking, plane routing, record locator data and customer tracking into a series of microservices that are housed in containers for orchestration. The DevOps team was then able to quickly build and deploy the new services while minimizing risk.

This approach enabled both agents and customers to quickly rebook customers when situations changed. The DevOps processes used allowed for on-demand scaling and processing of data. The private cloud environment was able to handle the load and support the needs of the airline's customers. It was also able to seamlessly connect to the public cloud to access and leverage key services, like weather data. Because of the continuous integration and deployment approach now implemented across the private cloud, the airline will continue to add new services and new innovations that will help increase business and be more competitive.

- » Understanding the value of data analytics
- » Looking at the role of the data science in data analytics
- » Grasping machine learning
- » Valuing the hybrid cloud when performing analytics

Chapter 5

Analysing Data in the Private Cloud

The idea of leveraging a large amount of data to better understand the patterns and anomalies has been a decades-long goal for many businesses. A hybrid cloud environment is the ideal platform for many data intensive analytics projects because it's built on a distributed computing model. For example, an analytics-based application may rely on data from a variety of sources located on premises in legacy applications, on a private cloud and on public clouds. Analytics involving customer data and other sensitive data can be performed on premises in a private cloud. In addition, public cloud resources can be used to access high-powered compute resources on demand.

In this chapter, we put the innovation available in data analytics and machine learning models in context of the private cloud.

The Value of Data Analytics

Driving better business decisions with analytics typically demands bringing together multiple sets of data from different sources. For example, it may be necessary to develop a predictive analytics

model to improve your knowledge of customer expectations and trends. Your model needs timely and accurate data across your data system, a variety of Software-as-a-Service (SaaS) applications, social media data and demographic data. Being able to leverage the performance and elasticity of a private cloud can provide organisations with the ability to analyse complex data close to the source. In other words, many times the largest datasets are often managed within the data centre. Many organisations want to maintain data behind the firewall, for a variety of reasons. In addition, being able to implement analytics closer to the data is often more pragmatic due to an increase in speed and corresponding reduction in data movement.



TIP

One of the major benefits of data analytics in a cloud environment is the ability to provide a self-service interface. By abstracting heterogeneous data sources through a portal interface, more business professionals can analyse complex data. In this model, the data scientist is able to create models that can be translated into a set of self-service modules that are useful to other employees.

The Role of Data Science in Data Analytics

The potential value of data science is becoming clear to many businesses. However, finding the resources and talent to create and execute on models is complicated. Not enough data scientists exist to support the type of analytics that businesses want to do. To be successful, IT organisations have to find support services that implement predictability and repeatability in the data science process.

To support the challenges inherent in putting machine learning to use in the enterprise, IBM has created a framework called the Data Science Experience (DSX) to facilitate the collaboration with colleagues. Data scientists can create a repository of shared predictive and prescriptive algorithms, curated data and pretested models. The DSX environment has been architected with open-source components. By basing DSX on open source, data scientists can take advantage of machine learning libraries and other resources that are found within the open-source community. The platform supports the most popular open-source technologies, including R, Python and Spark. DSX provides a set of services that

can connect these enterprise data sources to a variety of machine learning tools through a user interface. Services provided within this framework include a variety of data science notebooks. In addition, DSX includes security mechanisms to prevent unauthorised access to data. Even in a shared environment, a user only has access to datasets that he or she is authorised to see and use.

Understanding Machine Learning

Machine learning is a powerful set of technologies that can help organisations transform their understanding of data. This technology approach is dramatically different from the ways companies have traditionally used data and created applications. Rather than beginning with business logic and then applying data, machine learning techniques enable the data to model the logic to support business goals. One of the greatest benefits of this approach is that you remove assumptions and biases.

Machine learning is a form of artificial intelligence (AI) that enables a system to learn from data rather than through explicit programming. However, machine learning is not a simple process. Machine learning uses a variety of algorithms to create models. The models iteratively learn from data to improve, describe data and predict outcomes.

You likely interact with machine learning techniques inside applications without realising. For example, when you visit an e-commerce site and start viewing products and reading reviews, you're likely to receive suggestions for other similar products that you may find interesting. Likewise, when you call a customer call centre, the operator may be given prompts and suggestions based on the data that you provided. These recommendations aren't hard coded by an army of developers. Instead, the suggestions are made via a machine learning model. The model analyses historical data along with the real-time data that you provide to make predictions on the next best action.

The role of the private cloud in advanced analytics

When approaching machine learning, many organisations have discovered that leveraging a private cloud environment is the most pragmatic way to ensure scalability, security and compliance.

The private cloud provides the platform behind the firewall that can give you the performance and elasticity to support the complexities of advanced analytics. Of course, many organisations take advantage of the public cloud to analyse large amounts of data. However, in some situations, this option is not viable. Businesses are beginning to apply machine learning models to strategic initiatives that have the potential to change the way the company does business. Many companies are therefore hesitant to manage that data in a public cloud because of privacy and security concerns. Even if a public cloud has a sophisticated security infrastructure, many businesses are more comfortable with internal control of these strategic assets. In addition, special security precautions must be taken around certain types of data, such as health care records, personally identifiable information, customer records and banking information.

In some situations, a company may decide to run some of its data analytics that involve the use of public data sources in the public cloud but move to the private cloud when applying the results to strategic business initiatives or customer data. When your data is stored in a private cloud, public cloud services such as IBM's Watson APIs can be used to apply cognitive analytics to a business problem, while ensuring the data remains safely protected in the private cloud.

Hybrid data management

In addition to leveraging data in a hybrid cloud environment, many organisations have data stored in heterogeneous environments. For example, you may have structured data in your transactional and CRM system and unstructured data in a data lake. In many cases, data is stored in databases from different vendors. The goal of hybrid data management is to be able to manage this data to support business requirements. Therefore, businesses need a way of managing and gaining insight from data no matter where it resides. IBM provides a Common SQL engine as part of its approach to hybrid data management. This allows organisations to leverage data where it sits across Db2, Oracle, SQL Server and other data sources. When combined with IBM's Db2 Warehouse, analytics can be performed where the data resides.

Advanced Analytics to Support the Hybrid Cloud

Organisations can't manage and analyse data in just one way. However, we have seen three common scenarios:

- » The majority of the organisation's data is stored behind the firewall and the company wants to perform analytics on premises.
- » The organisation or division within a company stores most of its data in the cloud and wants to analyse this distributed cloud data.
- » The organisation wants to use a hybrid approach where it brings together cloud and on-premises data in order to analyse them in context with each other.

Not surprisingly, most companies want to take the hybrid approach of analysing on-premises and public cloud data together.

These scenarios have one thing in common: They all want to take advantage of the emerging machine learning and AI techniques and apply them to the data. This is not a simple problem to solve. Organisations have to understand their data: the data structure, its origins and its veracity. It is too easy to make poor business decisions if the data isn't carefully curated.

We cover each scenario in this section.

Leveraging existing on-premises data sources

Data warehouses have been a staple of the data centre for decades. While these data systems are often complex, they include a rich set of resources to better understand customers and the business. However, organisations typically execute queries against the data rather than doing advanced analytics. Machine learning tools and frameworks are changing this. Many IT organisations want to be able to apply machine learning algorithms to this data. One of the most pragmatic approaches is to leverage machine learning

notebooks and frameworks to ingest data from the warehouse or the data lake in order to apply analytics. In addition, it is possible through connection software to ingest additional data sources into a framework to analyse data, gain insights and predict outcomes.

Cloud native advanced analytics

Both born-on-the-cloud and established companies have large amounts of data stored in private and public clouds. This cloud-based data is often stored in various SaaS applications, cloud databases and cloud storage repositories. While the data warehouse or data lake may be centralised in the data centre, cloud data is highly distributed. To perform advanced analytics and machine learning in such a distributed environment, companies need an analytics platform that can bring the data together.



REMEMBER

In addition, because you are bringing together different datasets from disparate applications and environments, you need to make sure you have a process in place to ensure version control.

Hybrid cloud support for advanced analytics

Many businesses are operating in a hybrid world where they have a data centre, a private cloud and use public cloud resources. However, in many cases, the business fails to look at all of its data in context. It will likely analyse customer data housed in the cloud CRM application separately from the on-premises sales and inventory data. However, the true value of machine learning and advanced analytics begins to emerge when these isolated datasets are analysed in context with each other. In this scenario, you want to keep the data where it resides. You, of course, wouldn't want to move certain data to the cloud and likewise, you don't want to transfer massive amounts of cloud data to your data centre. By analysing all your data in context with each other, you are able to begin spotting patterns that may challenge long-held beliefs about your company and identify new opportunities.

IN THIS CHAPTER

- » Setting the stage for managing a hybrid environment
- » Visibility as a core principle
- » The rationale for multi-cloud management
- » Multi-cloud management requirements
- » Tools and services to manage a multi-cloud environment

Chapter 6

Managing a Multi-Cloud Environment

To be able to respond to the changing needs of the business, IT requires a new approach to IT management. Business leaders demand that computing be managed in an organised and predictable way, no matter where the workloads reside or how they were designed. Deployment models have to be managed as a unified environment. Therefore, IT management requires the ability to manage a multi-cloud environment through standardisation and automation. This chapter examines the issues and technologies needed to manage a multi-cloud environment. In addition, we provide insights into how IBM Cloud Private has addressed multi-cloud management.

Setting the Stage of Managing Hybrid Computing

To make disparate systems operate as one seamless environment, it's vital to have visibility, control and security for each element and the combined elements. Let's say that you have 20 different

development teams in the organisation. Each group may select a different cloud platform and a variety of DevOps tools. This may work well in the early stages of creating and managing individual applications, but trouble often occurs when these applications and services become instrumental to the operation of the business, especially if these applications and services need to be shared across business units.

Gaining Visibility

One of the advantages of a private cloud is that all the elements of the environment are controlled by your organisation. As you transition to a hybrid environment, it is important to be able to gain visibility and insight into the environment. Increasingly, there are many techniques used in cloud environment, ranging from virtual machines (VMs) to microservices, API management platforms and serverless architectures (where the cloud dynamically manages the allocation of machine resources and containers).

In a hybrid cloud environment, an organisation needs to manage a variety of services, including data centre applications, Software as a Service (SaaS), Infrastructure as a Service (IaaS) and Platform as a Service (PaaS). Therefore, IT management requires a simplified abstracted technique or dashboard to manage the multi-cloud environment.

These management tools and dashboards aggregate data from different on-premises and cloud platforms and services. Some of these management tools are now using machine learning and advanced analytics that can understand the patterns of behaviour across a variety of cloud environments and can help ensure that the combined services act in the way they are intended.

Understanding the Rationale for Multi-Cloud Management

The multi-cloud environment requires a way to centrally manage workloads and their underlying resources across all the cloud services that you're using. What makes this complex? Each cloud environment includes its own set of services and ways of managing resources.

Requirements for Managing the Multi-Cloud Environment

To create an operational platform requires that all services be identified and managed in a unified manner. Therefore, you need a set of services to create a cohesive environment. We list those services in this section.

Service catalog

A *service catalog* is at the core of being able to manage services across a hybrid cloud environment. It is important to establish context across a variety of IT services no matter where they are physically located. A service catalog identifies and defines the services that are available for developers and users to meet business objectives. The catalog defines the parameters and characteristics of each service, for example, who is allowed to use the service, how the service can be used, what the security requirements of the service are and what the dependencies are.

The catalog allows IT to manage commonly deployed IT services to ensure compliance, consistency and security, all based on corporate policy. The IT services in a catalog can include provisioning services, storage services, VM images, authentication services and business processes.



WARNING

Each service that is identified and managed in the service catalog has to be carefully vetted. The service has to be identified both by IT and the business as an important resource to be used by multiple business units. After the service is identified, it must be tested for accuracy. These services also have to include the policy rules that govern its use. Who within IT or the business is allowed to use the service and for what purpose? What deployment models are permitted for operating this service? Can it run in a specific public or private cloud? How is the service integrated with other services? The service catalog provides guidance as to what services can be integrated via well-defined application programming interfaces (APIs).



TIP

One of the benefits of the service catalog is that it keeps the details of the service itself abstracted from the user of that service. The actual code is encapsulated into Docker images, containers or templates. The catalog also indicates how and when a service can be altered.

Configuration management

IT departments 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. Configuration management gathers data about the dependencies between services running in different deployment models. In a multi-cloud environment, different resources have to be managed, including bare metal services, VM images and container services. Each environment has different configuration management requirements that have to be carefully controlled. In a multi-cloud environment, there is a requirement to coordinate configuration management approaches.



REMEMBER

One of the important aspects of configuration management is to be able to implement governance and security rules within the multi-cloud environment. Considering that different cloud environments might have different configuration management, implementing governance and security becomes a challenge.

Service automation and orchestration

Service automation defines the automated process to deploy a service in a consistent manner. Service orchestration is focused on combining these services together to guarantee availability, resilience and performance. In order to bring services from a variety of clouds together, it is necessary to define the templates or patterns for how services must interact with each other so that it ensures compliance with governance requirements.

Orchestration is essentially when you want to build an application from existing services. This requires process management and orchestration through the use of API wrappers. If the service catalog is the way you identify and classify services, automation and orchestration are the techniques for linking together these services to create workflows. Lower level repetitive tasks can be managed using service automation. These are services that are not visible to the user. Instead, automation is used for tasks such as initiating a trouble ticket or provisioning a cloud instance.

In contrast, service orchestration is used when the task involves bringing together services defined in the catalog to create a new

business process. The value of service orchestration comes into play as we move to the As-a-Service model where we are taking a variety of predefined services and linking them together dynamically. This is in stark contrast to the traditional application that is written as an end-to-end process. Orchestration has to be implemented in context with application performance management. It is not enough to simply link services together. It is critical to ensure that when those services create a new business application, the performance reflects business requirements.

Creating a Balanced Platform with API Management

All the services mentioned in this chapter are necessary in order to create a well-functioning system made up of components from different on-premises, hosted and public cloud environments. Managing a multi-cloud environment requires standardisation and automation. Without standardisation, an organisation would have to understand all the details of every cloud environment and every platform. Without standardisation, a multi-cloud environment isn't practical. In fact, standardisation is required to be able to automate the steps required for the hybrid cloud. API management provides a consistent way to manage APIs as a life cycle. This is critical because APIs are also an essential means of sharing this intellectual property with customers and partners.

A typical API management platform or portal enables centralised administration to make deployment easier and safer. The API management platform provides a set of tools that helps to build, debug and deploy these APIs. The API management portal can also be used to discover what APIs exist and the rules that govern their use. The portal can also monitor the performance of APIs so they help manage the performance of the system.



TIP

For example, IBM offers IBM API Connect, designed to manage API throughout the life cycle, from creation to security and analytics on the use of the API. Check it out at developer.ibm.com/apiconnect.

Managing in a Microservices World

The benefits of microservices are becoming increasingly important for companies. (We discuss microservices in more detail in Chapters 3 and 4.) While microservices have clear benefits, the management of these services can be complex. Rather than having large applications that are managed as a single service, in a microservices and containerised world, you have applications made up of dozens or even thousands of microservices. Depending on your organisation's development approach, there is a strong likelihood that microservices will

- » Have different languages and different technologies
- » Be owned by different teams in the development organisation
- » Include their own version control and update underpinnings

Given the importance of microservices, how can you create a management environment that understands these services and enables them to behave well together? In essence, there needs to be a software-defined layer that sits on top of the variety of systems and enables these microservices to interact. Microservices include a set of standardised APIs that can interact with the software-defined layer.

IBM Cloud Automation Manager

IBM Cloud Automation Manager (CAM) is designed specifically to give the central IT team the ability to provide responsive management across many clouds supported within the business. IBM CAM defines a set of automation templates containing a set of resources and their related attributes to many environments, including VMware clouds, OpenStack, IBM Cloud, Amazon EC2 and Microsoft Azure.

In order to more easily manage a multi-cloud environment, IBM CAM supports templates designed in Terraform. These Terraform templates can be applied to a variety of environments. This approach provides you with the ability to manage multi-cloud

environments through templates that can be configured based on the supported infrastructure. The automation means that you can avoid manual coding. Therefore, you can combine microservices with VMs and containers. These services can be combined into a single object and can be placed into the service catalog.

Check out IBM CAM at www.ibm.com/us-en/marketplace/cognitive-automation.

Terraform

IBM CAM uses Terraform, an open-source tool that is supported by all the major cloud providers, as the configuration language. It is a software environment for building, changing and versioning IaaS. Terraform creates configurations based on the platform and then generates an execution plan. When the configuration changes, Terraform applies the appropriate changes. This tool can manage VM instances, storage, networks and applications, as well as build a graph of all computing resources and parallelise the creation and modification of any non-dependent resources. Terraform can provide automation delivery of application stacks stored in the IBM Cloud, Amazon's AWS, Microsoft Azure, Google Cloud Platform, VMware, OpenStack and PowerVC.

Because Terraform requires services to configure, patch and manage automation, Chef is often used. In addition, through the use of Terraform, it is possible to leverage IBM Cloud services such as the variety of Watson APIs, including text-to-speech operations.

Service catalog within IBM CAM

Through the use of a service catalog, the manager provides visibility at the core of the multi-cloud environment. The intention is to provide IT with a set of cloud services to help standardise deployment of the full cloud application stack. Built as a microservices application, IBM CAM provides a visual interface intended to compose data centre services and application services so they can be consumed in a self-service manner.

Service composer

The service composer is a graphical tool that uses a pallet canvas with a drag-and-drop interface to compose application services that can then be published in the self-service catalog. The service composer also supports Terraform configuration modules for VMs, load balancers and networks. The composer is intended to link together multiple activities, including automation templates, in the order that they should be executed.

Monitoring a Multi-Cloud Environment

IT must have the ability to monitor the variety of platforms and services within the hybrid cloud. Therefore, it is necessary to monitor the activity across microservices, transactional services, networking services and data services. All systems provide logs of their activity, including how they perform and what actions are being taken within and across services. Management software provides analysis, sometimes in real time regarding how well the systems are executing based on analysis of data.

For example, IBM Cloud offers IBM Cloud Monitoring. The offering gives DevOps teams insight into how applications are performing and using resources. In addition, the service allows teams to spot potential problems and address them before they cause a disruption. In addition, organisations typically acquire a variety of management and monitoring tools to support their workloads.

Support for Multiple Monitoring Tools

IBM recognises that most enterprises have implemented many different monitoring tools across the data centre and various cloud environments. Therefore, rather than build a separate set of monitoring tools, IBM provides a set of APIs that partners can use to integrate their monitoring offerings into the IBM Cloud Management platform.

Setting the Stage for Action

A private cloud provides the integration platform to rapidly develop new functions and APIs that integrate with existing applications and data. It allows enterprises to leverage existing infrastructure investments.

IBM Cloud Private offers an on-premises, integrated solution working across IaaS, PaaS and services. This solution is built on container and data services technologies with API consistency across public clouds. It provides a single platform that can run on existing on-premises infrastructure and contains the security, compliance, performance and cost efficiency that enterprises demand. IBM Cloud Private provides developers with choice of languages, frameworks, runtimes and services to build cloud-native applications and microservices.

The importance of standards in private clouds

As cloud computing matures into an enterprise imperative, organisations need to have a variety of deployment models based on the nature of the workloads being supported. To be successful, organisations need to be able to modularise software through the creation of microservices deployed in containers in a secure and scalable manner. In this book, you explore IBM Cloud Private and how it can be used in conjunction with IBM's public cloud as well as third-party cloud services.

Inside...

- What is a private cloud?
- Explaining core cloud services
- Understand microservices and containers
- Understanding the hybrid cloud
- Benefits of standards in portability
- Managing cloud services
- Explaining container cloud platforms



Judith Hurwitz, President, Hurwitz & Associates, is a consultant, thought leader, and coauthor of eight books, including Cognitive Computing and Big Data Analytics.

Daniel Kirsch, principle analyst, Hurwitz & Associates, is a researcher and consultant in cloud, machine learning, and security.

Go to **Dummies.com**®
for videos, step-by-step photos,
how-to articles, or to shop!

for
dummies®
A Wiley Brand



Also available
as an e-book

ISBN: 978-1-119-51192-2
Part #: KUM12391GBEN-00
Not for resale

WILEY END USER LICENSE AGREEMENT

Go to www.wiley.com/go/eula to access Wiley's ebook EULA.