

Making Everything Easier!TM

Tintri Special Edition

Building an Enterprise Cloud

FOR
DUMMIES[®]

A Wiley Brand

Brought to you by



Tintri

Lawrence C. Miller



About Tintri

Tintri puts the agility of public cloud inside your data center. Tintri's all-flash architecture uses building blocks similar to public cloud to simply manage enterprise and cloud native applications. Customers guarantee the performance of their applications, automate common IT tasks via Open APIs, troubleshoot across their infrastructure, and predict an organization's needs to scale — the underpinnings of a modern data center. That's why leading cloud service providers and enterprises, including Comcast, Chevron, NASA, Toyota, United Healthcare, and 20 percent of the Fortune 100, trust Tintri with enterprise cloud.

Visit www.tintri.com and follow @Tintri on Twitter.

Building an Enterprise Cloud

FOR
DUMMIES®

A Wiley Brand

Tintri Special Edition

by Lawrence C. Miller

FOR
DUMMIES®

A Wiley Brand

Building an Enterprise Cloud For Dummies®, Tintri Special Edition

Published by
John Wiley & Sons, Inc.
111 River St.
Hoboken, NJ 07030-5774
www.wiley.com

Copyright © 2017 by John Wiley & Sons, Inc., Hoboken, New Jersey

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. Tintri and the Tintri logo are trademarks or registered trademarks of Tintri, Inc. 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 978-1-119-41532-9 (pbk); ISBN 978-1-119-41533-6 (ebk)

Manufactured in the United States of America

10 9 8 7 6 5 4 3 2 1

Publisher's Acknowledgments

We're proud of this book and of the people who worked on it. Some of the people who helped bring this book to market include the following:

Project Editor: Martin V. Minner

Senior Acquisitions Editor: Amy Fandrei

Editorial Manager: Rev Mengle

Business Development Representative:

Karen Hattan

Production Editor: Siddique Shaik

Introduction

Cloud is changing how organizations deploy, manage, and support the applications that are critical to running their businesses. Organizations must support applications (such as enterprise resource planning, or ERP) and databases that tend to have predictable usage patterns with stability, reliability, and safety as key requirements, both on-premises and off-premises (hosted in a private or public cloud).

Additionally, organizations are now building applications to be *cloud-native* — designed to support new services that can be tested and deployed in days or even minutes (instead of weeks), updated daily, and scaled in real-time. The growth of these cloud-native applications has prompted more organizations to adopt Agile development methods.

An organization can support the needs of both enterprise and cloud-native applications by creating an environment that delivers the agility of public cloud in its data center.

About This Book

Building an Enterprise Cloud For Dummies, Tintri Special Edition, shows you a path to enterprise cloud — an architecture designed to meet the needs of both enterprise and cloud native applications, and to provide public cloud agility in your data center. This book consists of four short chapters that explore:

- ✓ Which cloud trends are driving enterprise cloud adoption (Chapter 1)
- ✓ What benefits organizations can reap from an enterprise cloud (Chapter 2)
- ✓ How to build an enterprise cloud infrastructure (Chapter 3)
- ✓ What features and capabilities you need in an enterprise cloud (Chapter 4)

Foolish Assumptions

It's been said that most assumptions have outlived their usefulness, but I assume a few things nonetheless.

Mainly, I assume that you're an IT executive, manager, cloud architect, or administrator with a good understanding of IT infrastructure and operations, but that your organization has recently started the journey to the cloud.

If these assumptions describe you, this book is for you!

Icons Used in This Book

Throughout this book, I occasionally use special icons to call attention to important information. Here's what to expect:



This icon points out information that you should commit to your non-volatile memory, your gray matter, or your noggin'.



This icon explains the jargon beneath the jargon and is the stuff legends — well, nerds — are made of!



Tips are always appreciated, never expected — and I sure hope you'll appreciate these tips! This icon points out helpful suggestions and useful nuggets of information.



This icon points out the stuff your mother warned you about. Okay, probably not, but take heed nonetheless — you might just save yourself some time and frustration.

Where to Go from Here

A journey of a thousand miles begins with a single step. This journey is only 24 pages, so you might just turn the page. But if you see a particular topic that piques your interest, feel free to jump ahead to that chapter. Each chapter is written to stand on its own, you're welcome to start reading anywhere and skip around to your heart's content!

There's only so much I can cover in 24 short pages, so if you find yourself at the end of this book thinking, "Where can I learn more?" just go to www.tintri.com.

Chapter 1

Establishing Your Cloud Options

In This Chapter

- ▶ Going beyond virtualization to the cloud
 - ▶ Demystifying cloud definitions
 - ▶ Introducing the enterprise cloud
-

The conventional IT model, which has been constrained by siloed, costly, and inflexible infrastructure, is giving way to cloud, which is designed to serve business applications with increased agility, productivity, and cost-efficiency. In this chapter, you learn about the rise of cloud computing and the emergence of the enterprise cloud.

When Did Virtualization Become a Cloud Imperative?

The cloud is everywhere today! In fact, you could say that the forecast for IT is mostly cloudy. Of course, “mostly cloudy” in IT doesn’t have the same gloomy implications as it does in meteorology. As of 2017, more than 75 percent of workloads have been virtualized, and IDC estimates that by the end of 2020, virtualized instances will represent more than 90 percent of workload instances deployed globally. Businesses have significantly increased their use of virtualization and containers in recent years, to achieve greater infrastructure cost efficiencies and scale. And, as it turns out, the adoption of virtualization is a prerequisite for — wait for it . . . the cloud!

Enterprises are increasingly deploying cloud technologies, including public and private cloud offerings, to deliver highly scalable and automated services delivered on demand. IDC's *CloudView Survey* revealed that respondents expect their budgets for traditional IT, which includes both in-house and outsourced deployments, to decline nearly 20 percent from 2016 to 2018, while their budgets for public and private clouds will grow nearly 30 percent and more than 50 percent, respectively, during the same period (IDC CloudView 2016 Survey, US40852416, January 2016).

Defining Cloud Types — NIST, Not NOAA

Although the cloud is everywhere, it is not every “thing.” Many technology vendors and service providers are re-branding their existing products and services as cloud offerings, eager to capitalize on the market opportunity in the cloud. This rush to position everything “as-a-service” or a cloud solution — a practice that Gartner refers to as *cloudwashing* — has obscured what the cloud is (and isn’t).

So, to “clear the air” of any confusion regarding the cloud, I’ll start with a few standard cloud definitions. And, who better to cite for standard definitions than the U.S. National Institute of Standards and Technology (NIST)?

While the U.S. National Oceanic and Atmospheric Administration (NOAA) defines 27 cloud categories based on nine cloud types, NIST defines just four deployment models and five essential characteristics. The four cloud deployment models are:

- ✓ **Public:** A cloud infrastructure that is used by multiple organizations (multi-tenant) and is owned, managed, and operated by a third party
- ✓ **Private:** A cloud infrastructure that is used exclusively by a single organization and may be owned, managed, and operated by the organization or a third party (or a combination of both) either on- or off-premises
- ✓ **Hybrid:** A cloud infrastructure that is composed of both public and private cloud models
- ✓ **Community:** A cloud infrastructure used exclusively by a particular group of organizations (not common)



NIST defines the following five essential characteristics of cloud computing:

- ✓ **On-demand self-service:** Cloud computing services, such as server time and network storage, can be automatically provisioned without requiring IT support.
- ✓ **Broad network access:** Services are available over the network and accessed through standard mechanisms (such as open application programming interfaces, or APIs).
- ✓ **Resource pooling:** Computing resources are pooled to serve many customers (multi-tenancy) and demand levels, and are dynamically assigned and reassigned, as needed.
- ✓ **Rapid elasticity:** Services can be provisioned and released, in some cases automatically, to scale (up/down and in/out) with demand.
- ✓ **Measured service:** Resource usage can be monitored, controlled, optimized, and reported.

Introducing the Enterprise Cloud — A Twist on NIST

Many organizations recognize that although public cloud delivers many benefits, it is not the right solution for all workloads. Moving applications to public cloud platforms can result in significant migration cost and effort, sometimes requiring applications to be recoded, reconfigured, refactored, and reintegrated. In addition, although public cloud infrastructure can scale applications with fluctuating demand, the unexpected cost from unpredictable data growth — or the cost of a large-scale cloud deployment — can quickly get out of control.

Private clouds provide many of the benefits of public clouds, such as resource pooling, rapid scaling, automation, and self-service, but also give organizations more control and flexibility over access to and usage of their applications. Private clouds are thus ideal for larger organizations or those organizations with strict data, governance, regulatory, and compliance mandates. Unlike public cloud offerings, which are designed primarily to support cloud-native applications, private clouds can satisfy the needs of both enterprise and cloud-native applications. Another unique benefit of private cloud is the ability to more easily customize the compute,

storage, and networking components to best suit the specific requirements of an organization.

Two factors — the compelling benefits of private clouds and the desire to have access to public clouds — have given rise to what is increasingly known as *enterprise cloud*, which is a cloud infrastructure deployed in an organization's own data center and connected to public clouds. An enterprise cloud provides many of the same benefits and capabilities as a public cloud, including autonomous services, automation, self-service, and analytics, with the added control, security, and support for enterprise applications of a private cloud.

An enterprise cloud should have the following key characteristics:

- ✓ **Web services architecture:** An enterprise cloud should be built using building blocks of interchangeable components that can be easily connected together to create a large number of useful web services. This design provides a common platform that allows multiple infrastructure components to communicate with each other. Infrastructure systems built this way can be broken down into multiple component web services, so that each of these services can be automated, deployed, modified, and then redeployed independently, without compromising the operation of the infrastructure.
- ✓ **Comprehensive suite of APIs:** An enterprise cloud should be based on a comprehensive set of modern, web-based APIs, including Representational State Transfer (REST), that provide programmatic access to a wide range of web services and third-party ecosystems. While non-web APIs are structured and rigid, requiring strict programming models, modern APIs designed for cloud are open and flexible. These modern APIs, which are easy to assemble, integrate, tear down, reconfigure, and connect to other services, underpin today's web services.
- ✓ **Right level of abstraction:** An enterprise cloud should support virtual machine (VM) and container level operations, which provides programmability at the level of individual applications. VM and container-level abstraction allows performance isolation at a granular level, making it possible to ensure performance of applications without manual intervention through automatic, policy-based quality of service (QoS) for performance tiers. The right level of abstraction is a prerequisite for automating many operational and technical processes, and self-service.

Chapter 2

Realizing the Benefits of an Enterprise Cloud

In This Chapter

- ▶ Getting the agility benefits of public cloud
 - ▶ Addressing performance with visibility and performance insights
 - ▶ Automating and orchestrating at cloud scale
 - ▶ Supporting DevOps
-

In this chapter, you explore four key benefits of an enterprise cloud architecture.

Enabling Autonomic Operations

Managing infrastructure can be a monumental challenge. Compute, storage, and networking must all work in concert or performance and reliability will suffer. Manually managing discrete units of computing, such as a virtual machine (VM), as well as the orchestration of its associated storage and network attributes, further compounds that challenge. When conventional storage systems are part of an enterprise cloud deployment, they bring tremendous management baggage because of legacy storage structure bindings — which are irrelevant constructs in the cloud.



Conventional storage architecture uses logical unit numbers (LUNs) and volumes to define a filesystem, and storage operations and performance characteristics are executed and assigned at that level. The lack of granularity creates numerous challenges including:

- ✓ Inflexible allocation of performance and capacity
- ✓ Wasteful over-provisioning to ensure peak and future performance needs
- ✓ Manual processes to reclaim space and continually tune storage as workloads change

Conventional architecture assigns VMs and containers to one or more LUNs, and each LUN typically contains many other VMs and/or containers. Performance resources are assigned at the LUN level and all the VMs and/or containers on a LUN must share the same resources. If you have a few resource-intensive VMs or containers (“noisy neighbors”) on a LUN, some workloads may not be able to access sufficient resources to perform satisfactorily — and you’ve got a conflict on your hands. To address these performance issues, an administrator has to randomly shuffle individual VMs and containers out of one LUN and into another, hoping for better results. Such a manual trial-and-error approach might be acceptable in a data center with a few hundred VMs and containers, but it doesn’t scale and is unsustainable in a cloud with hundreds of thousands, or even millions, of VMs and containers.

Enterprise clouds provide a platform to support scaling of storage as needed, but managing that storage as it scales can burn lots of time. That’s why autonomous operation is such an important underpinning of an enterprise cloud architecture.



Autonomous operation means that as you scale and spin up new workloads, no manual intervention is required to ensure performance across the environment.

How is autonomous operation in an enterprise cloud possible? An enterprise cloud architecture assigns every VM and container to its own performance lane. Each lane gets the exact resources required, so there is never any conflict or need for manual intervention. As you provision more

VMs, new lanes are added and the entire system functions autonomously.



Autonomous operation has three prerequisites:

- ✓ An architecture built around VMs and containers
- ✓ Clean APIs that allow integration with other infrastructure elements
- ✓ Automation tools that simplify scripting and eliminate the need for manual intervention

When these three pieces are in place, you've got a solution that just works.

Driving Scale and Efficiency with Automation and Orchestration

Automation and orchestration are critical to the efficient operation of cloud computing environments at scale. It simply isn't possible to maintain hundreds of thousands, or even millions, of VMs and containers, let alone optimize these workloads or provision new workloads at the speed of business, if human intervention is required at every step.



Automation refers to a task or function that is performed without requiring human intervention. *Orchestration* refers to the coordination or sequencing of automated tasks and/or functions to accomplish a defined process or workflow.

With a web services architecture, all tasks are performed at the VM or virtual disk (vDisk) level. As a result, they can be easily mapped to cloud services for end users. These value-added tasks cover advanced features for data synchronization, data protection, and quality of service (QoS).

Workflows are simple and reliable, so you can include value-added storage services as part of your self-service offerings. When creating VMs, for example, you can easily create an automation workflow to allow users to specify the synchronization, QoS, and replication settings they want.

Being Proactive with Real-time Analytics

Most virtual environments are tied to conventional storage architectures with LUNs and volumes, so they suffer significant shortcomings in identifying VM and container behaviors, creating barriers to real-time analytics. Conventional storage's resource pool structure has no intrinsic meaning for VMs and containers, making granular real-time monitoring for visibility on a per-workload basis and visibility into cloud latency issues all but impossible without massive additional effort. Without the monitoring visibility to support real-time cloud storage analytics, it becomes impossible to optimize an enterprise cloud.



Granular visibility drives real-time and predictive analytics across an enterprise cloud. For maximum value, these analytics must provide real-time actionable insights at a VM or container level, to rapidly identify factors — across the entire infrastructure — that could contribute to increased latency, degraded performance, or even downtime. Rapid identification and resolution of these factors in an enterprise cloud requires analytics that provide a real-time view of what is going on in the cloud's virtual infrastructure.

In addition, leveraging time and condition data to generate predictions of future behavior can help identify trends and requirements for additional capacity or performance before running up against roadblocks. Analytics tools can also enable what-if modeling; the next time the business asks you whether your enterprise cloud can accommodate another one hundred development servers, you won't have to do any calculus or rely on a SWAG (that's "Scientific, uh, Without Analytics Guess") — instead, you can provide a definitive answer about the impact of such a change in seconds.

Enjoying the Convenience of Self-Service

A central benefit of enterprise cloud is that non-experts, such as individual business units and DevOps teams, can control and manage their own cloud footprints — which removes dependencies on IT.

For example, traditional development processes, in which development and operations teams work in separate silos, can take days or weeks to deploy or modify infrastructure in support of new development initiatives (see Figure 2-1). A major driver of innovation and competitiveness within organizations today is DevOps — which empowers developers with self-service capabilities to spin up, modify, and shut down an infrastructure environment in minutes. To be successful, DevOps teams must have access to real-time copies of production data and the ability to ensure that changes in a master VM or container are instantly replicated across all the copies that have been generated. This requirement is particularly important for teams that rely on Agile development processes, which are built around a steady stream of smaller development efforts that drastically increase the number and rate of changes that an infrastructure needs to support.




Figure 2-1: DevOps accelerates the deployment of infrastructure in support of development projects.

Having widespread access to always current data in an enterprise cloud — without affecting the production environment — provides a major boost in productivity to an organization's DevOps teams. But providing real-time access and consistency among workload copies is a herculean effort without the proper tools, especially in the cloud. Trying to manage this effort with the limited services offered by conventional compute, storage, and network infrastructure is not feasible and it is becoming ever more difficult to meet the needs of DevOps teams.

Instead, DevOps teams need the ability to manage their own cloud footprint — within appropriate limits (enforced by

permissions), of course. If DevOps teams have to depend on others in the IT department for all of their infrastructure needs in the cloud, development cycles will inevitably be delayed and frustration will ensue.

At the same time, you don't want your developers burning cycles trying to figure out how to configure LUNs, volumes, and other infrastructure stuff — after all, that isn't their core competency (and a developer who knew how to write great code *and* configure infrastructure would be really dangerous — enough to possibly take over the world).

Instead, DevOps teams need simple and intuitive self-service tools with easy-to-use interfaces that empower them to provision their own VMs and containers, create policies, and tear down VMs and containers, as needed. With the right automation at work in the background, DevOps can engage a bot to complete tasks via chat-based, real-time collaboration and communications tools such as Slack. Going a step further, devices such as Alexa bring voicebots to automation and self-service, enabling developers to accomplish tasks using their voice. These interfaces are already familiar to DevOps teams and remove dependencies on IT.



Real-time storage management at the VM or container level gives DevOps the granularity and data consistency it needs to streamline development efforts and accelerate the release of new applications. Having access to current data and the ability to propagate it across workload copies allows DevOps to experiment with different modifications of applications before rolling them out to production environments. In addition, providing an option for rollover or temporary capacity expansion via the public cloud, as needed, further enhances enterprise cloud benefits to DevOps.

Chapter 3

Building an Enterprise Cloud Infrastructure

In This Chapter

- ▶ Choosing the best cloud for your applications
 - ▶ Getting started with the right architecture
 - ▶ Eliminating manual administration and processes
 - ▶ Growing your enterprise cloud
-

In this chapter, you learn how to build an enterprise cloud infrastructure for your organization.

Step 1: Assess Application Needs

Infrastructure is built to meet the needs of the applications that are deployed to run on it. While public cloud is well suited to cloud-native applications, that may not be the case for legacy enterprise applications. A central benefit of enterprise cloud is its ability to support *both* cloud-native applications *and* legacy enterprise applications.

When considering whether certain applications are candidates for public or enterprise cloud, it's important to account for these characteristics:

I✓ **Age:** If an application is more than ten years old, it may not be supported by public cloud. Legacy applications tend to have complex interface requirements and a long list of enhancements and support requests, making public cloud migration problematic — they will perform more predictably in enterprise cloud. Cloud-native applications can thrive in either public or enterprise cloud.

- ✓ **Data coupling:** If an application needs its data to be “local,” it may not translate well to public cloud, and should be considered a candidate for enterprise cloud.
- ✓ **Compliance:** There may be instances where an application is subject to regulation or compliance requirements that may not lend themselves to public cloud instances.
- ✓ **Team size:** For applications that are used only by an individual or small group in a single location, it may not make sense to migrate it to public cloud. For larger or more dispersed teams, you should consider both public and enterprise cloud.
- ✓ **Complexity:** If an application is tightly integrated with other applications or requires specialized infrastructure, it is less suited to public cloud. These types of applications are often resource intensive, which can run up public cloud bills, and require a degree of resiliency that can be offered only in enterprise cloud.
- ✓ **The cost of “undo”:** If an application is moved to the public cloud and the result is unacceptable, how easily can that application be moved back as a local instance? If failure cost is “high,” enterprise cloud reduces both risk and cost.
- ✓ **Other potential factors:** Agility, sensitivity to exposing an application programming interface (API) to the outside world, reliability, and performance are all important factors that need to be considered.

In the end, the primary drivers to migrate applications to public cloud are usually scalability/elasticity and agility/time to market. But rather than react to those requirements alone, it’s important to consider the characteristics in the preceding list — it may be that those applications can access the same agility as public cloud in an enterprise cloud, and get greater coupling, compliance, and recoverability.

Step 2: Establish the Right Architecture

Once your organization decides to move to the cloud, migrating certain infrastructures can be either a breeze or a migraine, both for setup and continuing operation, depending on initial

and ongoing design decisions. It is critical that configuration and management decisions work within the capabilities and limitations of the cloud. A properly architected enterprise cloud provides storage resources ideally matched for the requirements of enterprise applications at the individual virtual machine (VM) and container level.

Identify and allocate storage resources

Storage is fundamental to the success of any infrastructure, so taking the time to identify requirements and then operationalize them is a critical first step to building an enterprise cloud.

Unlike many legacy on-premises enterprise applications with dedicated physical resources, cloud applications and their resources are virtualized and shared. Conventional shared storage systems manage objects such as logical unit numbers (LUNs), volumes, or tiers, which have no intrinsic meaning for VMs and containers. Each new VM instance or container must be assigned to a specific storage LUN or volume. When input/output (I/O) requirements and VM behavior are poorly understood, a painful trial-and-error process ensues to make sure the storage needs of each VM and container are met.

Performance-sensitive applications can end up competing with other applications for resources, resulting in some of those applications getting starved by noisy neighbors. Without intelligent resources to automatically manage this problem, the only two options are to have applications that are slow and/or time out, or to heavily overprovision resources to deal with fluctuating demand and load levels. The former option is inherently untenable (and may generate a “resume updating event”); the latter is wasteful and can be prohibitively expensive.



Tintri Enterprise Cloud dramatically simplifies storage management in cloud environments. Tintri addresses noisy neighbor problems by eliminating LUNs and assigning every VM and container its own lane. Tintri acts as a single, federated (loosely coupled) pool. This provides broad automation of resource allocation and greatly simplifies your enterprise cloud environment. With Tintri, there's no conflict over resources or policies and therefore, no noisy neighbors.

Define granular policies

After you have identified storage resources and implemented an allocation schema, you need to define policies for data protection/replication, quality of service (QoS), and other operations, and assign them to service groups. How will I/O operations per second (IOPS) be allocated per VM or container? How do you manage mixed workloads with different requirements? What process will accommodate different replication requirements of individual VMs and containers?



With Tintri's advanced QoS capabilities, administrators can allocate exact minimum and maximum IOPS to each individual VM or container, not just to a LUN or volume. Unlike conventional QoS, which requires administrators to predict the right IOPS values, Tintri removes the guesswork by providing visual guidance on the QoS values to specify. Tintri's granular QoS policies help you manage mixed workloads with different service level requirements, while hosting multiple types of hypervisors with support for chargeback.

Step 3: Automate

The next step in building your enterprise cloud is to determine how storage resources will be accessed. To achieve the full benefits of the cloud, storage overhead must be minimized without sacrificing reliability and performance. Creating complex, event-specific scripting to push and pull data between applications and storage cannot scale and is prone to error. The way to avoid this problem is to build your storage and applications around a simple and comprehensive, yet clean, storage API.

Such an API should have granular command capability, coupled with the ability to bundle commands around a complete storage action. If a series of individual API puts and calls are required to meet application data needs, the associated network chatter and administrative oversight can bring operations to a near halt. Thus, it is critically important to avoid repetitive, manual efforts by leveraging automation scripts that can reduce entire storage operations to single (or just a few) API calls to execute the required actions.

Ideally, APIs can be organized according to functional buckets for easy management. Creating API categories around operations, such as provisioning and policy management, greatly simplifies storage administration, even in complex, large-scale cloud deployments.

APIs that are kludged together with redundant actions, inefficient organization and operation, and potential conflicts can reduce performance, complicate administration, and lead to application timeouts and downtime. APIs need to be cleanly and efficiently constructed to allow infrastructure to be assembled like building blocks that easily snap together and share information.



Open APIs make it possible for different building blocks in an enterprise cloud to snap together and communicate. They are needed to write automation scripts and connect to other elements of infrastructure. Tintri Enterprise Cloud is built with Representational State Transfer (REST) APIs, as well as a PowerShell toolkit and Python software development kit (SDK) so that customers can automate previously manual tasks.

Step 4: Analyze and Scale

Okay, you've got your storage needs identified, policies defined, and orchestration via APIs outlined. The next step is ensuring you have a strong foundation for understanding and adjusting your storage as needed. This requires powerful yet easy-to-interpret analytics and a robust platform for scale.

The resource pool structure in conventional storage (LUNs and volumes) has no intrinsic meaning in an enterprise cloud, making granular real-time monitoring on a per-VM or container basis — and visibility into enterprise cloud latency issues — all but impossible without massive additional effort. Granular visibility drives real-time and predictive analytics across cloud-based hosts, networks, and storage, thereby addressing a significant shortcoming of conventional storage. For maximum value, these analytics must provide real-time actionable insights, at a VM or container level, to rapidly identify factors that could contribute to increased cloud storage latency, degraded performance, or downtime.

Beyond the lack of visibility and granular real-time analytics, managing conventional scale-out storage for enterprise cloud environments can be exceedingly complex. Some of the storage management challenges associated with conventional scale-up storage in an enterprise cloud environment include:

- ✓ LUN and volume level data leads to bad guesses and poor decisions about optimal VM placement
- ✓ Poor migration recommendations based on incomplete data, and no visibility into the impact on performance or the time required to complete a migration
- ✓ Problem VMs and containers constantly getting bounced back and forth between arrays reacting to, but never resolving, performance issues



A major benefit of the cloud is rapid elasticity and, in case you missed (mist?) the NIST definition in Chapter 1, it's also one of the five essential characteristics of the cloud.



Tintri Enterprise Cloud provides cloud-based, real-time, and predictive analytics to improve enterprise cloud planning and operations. These analytics can crunch millions of data points from as many as 160,000 VMs over several years in less than one second. Tintri Enterprise Cloud delivers a flexible yet robust foundation for your enterprise cloud that:

- ✓ Optimizes VMs based on a complete picture of their storage capacity and performance needs
- ✓ Gives least-cost recommendations, saving time, bandwidth, and capacity, to maintain optimal VM distribution
- ✓ Allows recommendations to be reviewed and edited, and the outcome viewed before committing
- ✓ Learns every time recommendations are edited, and allows certain VMs and containers to be opted out of migration

Chapter 4

Ten Key Capabilities of an Enterprise Cloud

In This Chapter

- ▶ Providing scalability with open APIs and granular levels of abstraction
- ▶ Ensuring performance through optimization, analytics, and prioritization
- ▶ Simplifying management with automation and self-service

This chapter offers ten important capabilities to look for in an enterprise cloud platform. Make sure your solution has these characteristics:

- ✓ **Web services as building blocks:** Web services are the building blocks that enable organizations to design and build an enterprise cloud using a modular approach. This architecture enables organizations to easily snap together, build up, tear down, and reconfigure individual pieces.
- ✓ **Granular level of abstraction:** Simply put, what is the level of granularity at which individual compute, storage, and network performance, operations, and management actions can be taken? For example, in most conventional storage, it's logical unit numbers (LUNs) and volumes — irrelevant constructs in the cloud (and in most “non-storage admin” minds). For an enterprise cloud, virtual machines (VMs) and containers — the currency of cloud — is the required level of abstraction.
- ✓ **Open APIs:** Open APIs make it possible for different blocks of compute, storage, and networking to snap together and communicate. They are needed to write automation scripts and connect to other elements of an enterprise or public cloud infrastructure.

✓ **Massive scale-out:** Scale is a defining characteristic of the cloud. While many companies aspire to build cloud-scale infrastructures with agility and automation for diverse virtualized workloads, they are held back by a conventional architecture that's simply not designed to support those requirements. Organizations are forced to choose either limited scale-out that requires lots of hardware, or expensive and inefficient scale-out — without the benefits of VM and container level visibility, quality of service (QoS), and automation.

Instead, what if you could scale out your compute, storage, and network in an enterprise cloud the same way you do in a public cloud, by just provisioning additional capacity and letting the hypervisor optimize the pool of resources?

For example, conventional storage scale-out was not designed for today's virtualized and cloud applications. An enterprise cloud build on a modern scale-out storage platform has the following characteristics:

- **Federated storage pools:** Federated storage treats multiple devices — both all-flash and hybrid-flash nodes — as a single pool of storage.
- **Separate control and data planes:** Separation of the control and data planes ensures low latency and predictable performance for virtualized workloads.
- **Scale compute and storage independently:** Loose coupling of storage and compute provides maximum flexibility to scale compute and storage elements independently, as needed.
- **Intelligent storage management software:** Deliver consistent performance with optimized placement of VMs and containers in storage.
- **Cloud-scale storage:** An enterprise cloud storage platform must be designed to scale from small to extremely large (hundreds of thousands to millions of workloads) environments.

✓ **Public cloud integration:** An enterprise cloud without public cloud integration is just another private cloud. Integration shouldn't be limited to a single public cloud repository, but rather offer the flexibility to connect

to multiple public cloud instances, and moving applications between your data center and your public cloud instances should be “one-click” simple. That’s how you deliver true application mobility for your organization and ensure that each application runs in the environment that best serves its needs.

- ✓ **Workload optimization and predictive analytics:** Unlike conventional architectures that manage VMs and containers as groups of objects — without regard for the individual performance requirements and characteristics of each workload — a web services architecture manages each workload individually and optimizes the environment for each specific workload based on real-time, predictive analytics that drive decisions. In this way, you can ensure your workloads perform as needed in an enterprise cloud.
- ✓ **Automation and orchestration.** Building an enterprise cloud on a web services architecture dramatically simplifies the orchestration of VMs and containers at scale by automatically aligning compute, storage, and network resources to changing requirements without any IT intervention. For example, an organization can set up automation of policy management — replication, QoS, and more — when they have just a few hundred VMs, and those policies will continue to operate as their footprint scales to thousands upon thousands of VMs. And the automation of these policies can be managed through one central management console in only minutes per day.
- ✓ **On-demand self-service:** Who would’ve thunk that pumping your own gas would lead to self-service checkout lanes at the grocery store and in the cloud. Whether you call it self-service, shadow IT, or do-IT-yourself, end users have seemingly grown accustomed to fending for themselves in the cloud — and they like it!

So, to quote the O’Jays, The Kinks, and/or Sharon Jones and the Dap-Kings — depending on how old and hip you are: “Give the people what they want!” Any solution for the enterprise cloud should have a simple, intuitive interface that enables users to easily provision and manage their requirements in the cloud. For example, you can automatically spin up 100 VMs and apply granular load balancing policies in anticipation of a new software product launch that will generate extremely high demand, simply by instructing a voicebot using Amazon’s

Alexa or using a chatbot within Slack (a web-based real-time collaboration tool for teams).

- ✓ **Showback/chargeback:** Showback or chargeback capabilities are related to the NIST pillar of “measured service” as a fundamental requirement of cloud computing (don’t recall the five NIST pillars? Go to Chapter 1, go directly to Chapter 1. Do not pass “GO,” do not collect \$200).

You’ve got to be able to show different teams within your organization how much cloud resources they are consuming — and potentially bill their department for those costs. This practice can help ensure good behavior and appropriate use of the organization’s cloud footprint.

Showback/chargeback also addresses one of the key concerns of public cloud — that costs can sneak up on you. For example, you might have a workload that uses a lot of “free” network bandwidth in your data center, but your public cloud provider is going to run up your tab. With showback/chargeback, teams have a clearer sense of what resources they are actually using, so there are no surprises later if workloads are migrated to the public cloud — and it might help you identify workloads that should definitely *not* be migrated to the public cloud because of how they consume resources.

- ✓ **ChatOps:** A term popularized by GitHub, ChatOps is a collaboration model that connects people, processes, and technology (tools) in an automated and transparent workflow.

ChatOps provides a persistent chat-enabled workspace that directly integrates with the tools and technologies that teams use. For example, a service request to provision a container for a new application can be posted by an end user in ChatOps, and integrated bots and scripts can automatically provision the storage and other cloud resources, then notify the end user in real-time when it is completed. The end user can connect to the new application from within the ChatOps interface, eliminating the need to switch between different management tools and interfaces so that context is never lost.

With these ten capabilities in place, you’re set to deliver the agility of public cloud in your data center. Congratulations on building an enterprise cloud!

public cloud agility in your data center



tintri.com @tintri



Open the book and find:

- Choose the right cloud for your workloads
- Guarantee application performance
- See the future with predictive analytics
- Drive scale and efficiency with automation and orchestration
- Enjoy the convenience of self-service

Achieve the agility of public cloud in your data center

Organizations need the agility, automation, and self-service of the cloud to support both enterprise applications and cloud-native applications. This book explores the cloud options available today and provides advice to help you maximize the benefits of enterprise cloud.

- **Explore** cloud deployment models and the rise of the enterprise cloud
- **Design** an enterprise cloud that can snap together and scale like building blocks
- **Improve** agility and gain performance insights

Lawrence C. Miller has worked in information technology for more than 25 years. He has written more than 50 For Dummies books.

Go to **Dummies.com®**
for more!

FOR
DUMMIES®
A Wiley Brand



Also available
as an e-book

ISBN: 978-1-119-41532-9
Not for resale

WILEY END USER LICENSE AGREEMENT

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