

LEARNING MADE EASY

Ixia Special Edition

Cloud Visibility

for
dummies[®]
A Wiley Brand



Understand cloud
basics and strategies

—
Optimize cloud
for your organization

—
Achieve cloud visibility
and security

Compliments
of
ixia

Lawrence C. Miller

About Ixia

Ixia, recently acquired by Keysight Technologies, provides testing, visibility, and security solutions, strengthening applications across networks and cloud environments for enterprises, service providers, and network equipment manufacturers. Ixia offers companies trusted environments in which to develop, deploy, and operate. Customers worldwide rely on Ixia to verify their designs, optimize their performance, and ensure protection of their networks and cloud environments to make their applications stronger. Learn more at www.ixiacom.com.

About Keysight Technologies

Keysight Technologies is a leading technology company that helps its engineering, enterprise, and service provider customers optimize networks and bring electronic products to market faster and at a lower cost. Keysight's solutions go where the electronic signal goes, from design simulation, to prototype validation, to manufacturing test, to optimization in networks and cloud environments. Customers span the worldwide communications ecosystem, aerospace and defense, automotive, energy, semiconductor, and general electronics end markets. Keysight generated revenues of \$2.9 billion in fiscal year 2016. In April 2017, Keysight acquired Ixia, a leader in network test, visibility, and security. More information is available at www.keysight.com.



Cloud Visibility

Ixia Special Edition

by Lawrence C. Miller

for
dummies[®]
A Wiley Brand

Cloud Visibility For Dummies®, Ixia 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, Dummies.com, 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. Ixia and the Ixia logo are trademarks or registered trademarks of Ixia 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 978-1-119-42449-9 (pbk); ISBN 978-1-119-42447-5 (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: Magesh Elangovan

Table of Contents

INTRODUCTION	1
About This Book	2
Foolish Assumptions.....	3
Icons Used in This Book.....	3
Beyond the Book.....	4
Where to Go from Here.....	4
CHAPTER 1: Understanding Cloud Basics and Business Drivers	5
Defining Cloud Service and Delivery Models	5
Looking at Cloud Enabling Technologies.....	8
Seeing How Business Trends Are Precipitating Cloud Adoption.....	10
CHAPTER 2: Networking in the Cloud	11
The Network Perimeter Is Obscured by Clouds	11
Network Interconnects in On-Premises and Cloud Contexts	13
Considering Different Cloud Models	14
Ensuring Resilience in Cloud Environments.....	15
CHAPTER 3: Ensuring Security and Compliance in the Cloud	19
Recognizing Security Similarities and Differences in the Cloud	19
Gaining Visibility into the Cloud.....	23
CHAPTER 4: Monitoring Network and Application Performance in the Cloud	27
Looking at Metadata and Raw Packet Data	28
Fault Detection, Incident Analysis, and Reporting.....	31
Big Data for Analytics.....	32
CHAPTER 5: Exploring Cloud Use Cases	33
Private Cloud.....	33
Government.....	33
Health care.....	34

Public Cloud	35
Retail (e-commerce).....	35
Technology services.....	36
CHAPTER 6: Forecasting the Future of Cloud.....	37
The Cloud Is Here to Stay	37
The Cloud Will Evolve as Technologies Emerge.....	38
Contextual Cloud Networks	39
Fog and Mobile Edge Computing	40
CHAPTER 7: Ten Important Considerations for Your Journey to the Cloud	41
Migration Plan.....	41
Skills and Competency.....	42
Vendor Lock-In.....	42
Network Visibility.....	42
Application Monitoring.....	42
Security Best Practices.....	43
Compliance Requirements.....	43
Disaster Recovery.....	44
Performance	44
Pricing Models	44

Introduction

Business and technology are at a critical inflection point. Globalization, the Internet of Things (IoT), cloud, virtualization, and mobile devices are forcing companies to extend their network edge — often into places where they cannot easily gain visibility. Budgetary constraints, technical limitations, security concerns, and performance issues prevent moving entire IT infrastructures to the cloud, and that will continue to be the case for the foreseeable future. In the meantime, enterprises are implementing a hybrid model, sending a mix of critical and non-mission-critical workloads outside their on-premises environment.

Enterprise applications often comprise a mix of Software-as-a-Service (SaaS) applications and custom-developed applications running on Platform-as-a-Service (PaaS) and Infrastructure-as-a-Service (IaaS) in the public cloud. The business data exchanged among on-premises applications running in corporate data centers and branch offices, as well as public and private clouds, increases the complexity of the end-to-end, real-time visibility needed to identify and predict network outages, identify a security breach, and analyze mission-critical application performance issues. Moving forward, IoT will add additional visibility challenges. This prompts a need to rethink approaches to end-to-end visibility and security in an increasingly complex cloud-based world.

Looking at the cloud visibility challenge differently, a recent Cisco white paper predicted “annual global IP traffic [would] pass the zettabyte (or 1,000 exabytes) threshold by the end of 2016” and “nearly triple from 2015 to 2020.” To put that into context, Cisco estimates “it would take more than 5 million years to watch the amount of video that will cross global IP networks each month in 2020.” Looking through that volume of data for security risks will be extremely difficult — like looking for a needle in a haystack the size of Texas.

As IT decision makers are working to implement and manage viable hybrid networks, they operate in a business environment where application and network performance is essential to generating revenue and maintaining customer relationships. Monitoring tools access critical application data in these virtualized networks

and cloud environments to ensure the reliability, security, and performance of mission-critical services. As enterprises and service providers adopt new technologies like software-defined wide area networking (SD-WAN) and micro-segmentation to improve the security of their infrastructure and performance of their critical services, the need for pervasive end-to-end visibility is further amplified.

Ultimately, as enterprise IT teams move their critical workloads from on-premises data centers into virtualized, software-defined data centers (SDDC) and public clouds, they must address several important questions:

- » How can we ensure availability, reliability, and performance of our mission-critical applications?
- » How do we get relevant critical data to analytics and monitoring tools, regardless of where they are located?
- » How can we tell which applications are suitable for cloud and plan a successful migration?

This book helps you answer these questions for your organization's journey to the cloud.

About This Book

Cloud Visibility For Dummies, Ixia Special Edition, consists of seven short chapters that explore

- » Cloud basics and the technologies and trends driving cloud adoption (Chapter 1)
- » Network visibility challenges in the cloud (Chapter 2)
- » Security and compliance issues in the cloud (Chapter 3)
- » Network and application performance considerations for the cloud (Chapter 4)
- » Different cloud visibility use cases (Chapter 5)
- » Future developments and trends in the cloud (Chapter 6)
- » Important considerations to look at for your organization's journey to the cloud (Chapter 7)

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 are a technology professional (such as a network administrator or cloud architect) or decision maker (such as a CIO, IT director, or data privacy manager). I also assume you're working for a large enterprise, or possibly a mobile carrier or cloud provider implementing a cloud strategy or migrating to a cloud environment. This book is written primarily for technical readers, but just in case you aren't technical, I explain any terms and concepts that come up.

Beyond a general knowledge of data center and network technologies, I assume that you have only a basic awareness of the cloud and you're eager to learn more, including how to assess different cloud options and solutions, as well as best practices for executing a cloud strategy or migration. If you're instead looking for "worst practices," I recommend reading *Cloud Visibility For People Who Like to Change Jobs Frequently!*

If any of these assumptions describe you (aside from the one about "worst practices"), this book is for you! If none of these assumptions describe you, keep reading anyway. It's a great book and when you finish reading it, you'll have enough visibility into cloud visibility to be dangerous!

Icons Used in This Book

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



REMEMBER

This icon points out information you should commit to your non-volatile memory, your gray matter, or your noggin — along with anniversaries and birthdays!



TECHNICAL
STUFF

You won't find a map of the human genome here, but if you seek to attain the seventh level of NERD-vana, perk up! This icon explains the jargon beneath the jargon!



TIP

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



WARNING

These alerts point out the stuff your mother warned you about (well, probably not), but they do offer practical advice to help you avoid potentially costly or frustrating mistakes.

Beyond the Book

There's only so much I can cover in 48 short pages, so if you find yourself at the end of this book, thinking "Gosh, this is a great book; where can I learn more?" just go to www.ixiacom.com.

Where to Go from Here

With my apologies to Lewis Carroll, Alice, and the Cheshire cat:

"Would you tell me, please, which way I ought to go from here?"

"That depends a good deal on where you want to get to," said the Cat — er, the Dummies Man.

"I don't much care where . . . ," said Alice.

"Then it doesn't matter which way you go!"

That's certainly true of *Cloud Visibility For Dummies*, which, like *Alice in Wonderland*, is also destined to become a timeless classic.

If you don't know where you're going, any chapter will get you there — but Chapter 1 might be a good place to start. However, 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, so you can read this book in any order that suits you (though I don't recommend upside down or backward).

I promise you won't get lost falling down the rabbit hole.

IN THIS CHAPTER

- » Learning the basics of the cloud
- » Looking at key technologies inside the cloud
- » Recognizing opportunities and challenges in the cloud

Chapter 1

Understanding Cloud Basics and Business Drivers

In this chapter, you learn the basics about the cloud, including how the cloud “made it rain” for Amazon, key enabling technologies in the cloud, and the business benefits and technical challenges of the cloud.

Defining Cloud Service and Delivery Models

Remember when Amazon only sold books? Today, you can buy — or sell — just about anything on Amazon. With entire towns being built around Amazon’s distribution centers and innovative trials like drone deliveries, Amazon is disrupting the entire supply chain — both vertically and horizontally. You might get the sense that Amazon is everywhere, and if you consider the other half of Amazon’s, uh, amazon’ story — Amazon Web Services (AWS) — you wouldn’t be far off!

Originally conceived as a virtual platform for Amazon’s retail computing infrastructure, with the “possibility” of selling some excess server capacity as a service to customers, AWS officially launched in 2006. In the first quarter of 2016, this “possible” source of additional revenue contributed 56 percent of Amazon’s total profit for the quarter. Today, AWS is available in an ever-growing number of geographical regions on five continents and has more than a million customers in more than 190 countries.

Although Amazon had “first mover” advantage, Microsoft launched its Azure cloud service in 2010 and has established itself as the other major service provider in the public cloud market. Today, Azure is also available on five continents and has a data center footprint of similar size and scale as AWS. Still, according to Synergy Research Group, AWS had approximately 40 percent of worldwide market share in 2016, while Azure — along with the two other major cloud service providers, Google and IBM — had a combined total of 25 percent of worldwide market share.

You know the cloud is “everywhere,” but it’s important to understand that the cloud isn’t “everything.” To avoid confusion, I define a few standard cloud terms with a little help from the U.S. National Institute of Standards and Technologies (NIST). NIST defines three cloud deployment models, including:

- » **Public:** A cloud infrastructure that is used by multiple organizations (multi-tenant) and is owned, managed, and operated by a third party (or parties) on the cloud provider’s premises.
- » **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. However, you should be aware that in some contexts, *hybrid environment* or *hybrid network* refer to a combination of cloud and on-premises computing — a somewhat different meaning of *hybrid*.



TIP

Technically, a fourth cloud deployment model exists. However, *community* clouds (the “kumbaya” cloud) — private clouds for exclusive groups of tenants — aren’t too common, so you don’t hear much about them (and you have to know the secret handshake).

Three cloud models are available (see Figure 1-1), defined as follows:

- » **Software-as-a-Service (SaaS):** Customers are provided access to applications running on a cloud infrastructure but the customer has no visibility into, and does not manage or control, the underlying cloud infrastructure. Examples include email, customer relationship management (CRM), virtual desktops, and gaming.
- » **Platform-as-a-Service (PaaS):** Customers can deploy supported applications onto the provider's cloud infrastructure, but the customer has no visibility into, and does not manage or control, the underlying cloud infrastructure. Examples include databases, web servers, execution runtimes, and development tools.
- » **Infrastructure-as-a-Service (IaaS):** Customers can provision processing, storage, networks, and other computing resources and deploy and run operating systems and applications, but the customer has no visibility into, and does not manage or control, the underlying cloud infrastructure. Examples include virtual machines (VMs), storage, load balancers, and networking.




FIGURE 1-1: Examples of SaaS, PaaS, and IaaS.

NIST defines the following five essential characteristics of cloud computing:

- » **On-demand self-service:** Services can be unilaterally and automatically provisioned.
- » **Broad network access:** Services are available over the network through various platforms and devices.
- » **Resource pooling:** Compute, storage, and networking resources are pooled to serve various tenants 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 transparently monitored, controlled, optimized, and reported.

Looking at Cloud Enabling Technologies

Some key enabling technologies of the cloud include condensation, water vapor, and aerosols — oops, wrong cloud. Key enabling technologies for cloud computing include:

- » **Virtualization:** For many — if not most — organizations, the journey to the cloud began with virtualization, a key enabling technology for the cloud. Virtualization, of course, is not a new technology — it dates to the mainframe era — but widespread enterprise adoption over the past decade has fueled near-continuous innovation and an ever-growing array of use cases. What may have started as an IT initiative to consolidate physical servers, or to build a lab for development and testing, quickly morphed into a race to virtualize the entire data center with the answer to the question “What percentage of your infrastructure is virtualized?” becoming an all-important benchmark and a badge of honor for IT infrastructure teams.

Virtualization technology abstracts software (such as an operating system, application, database, or network function) from the underlying hardware (for example, servers and storage).



TECHNICAL
STUFF



TIP

» **Ubiquitous Internet access:** Like the sky in which clouds exist in nature, the Internet is the “Skynet” in which cloud computing exists — wait, terminate that analogy! But the Internet is, of course, what connects us all to the cloud. If not for high-speed, highly reliable, always-on access to the Internet, cloud computing would not be possible. The proliferation of mobile devices and Wi-Fi hotspots further drives cloud adoption and as 5G networks are introduced (beginning in 2020), Internet access will become truly ubiquitous.

Download *5G For Dummies* from www.ixiacom.com to learn more about 5G networks.

» **Software-defined everything (SDx):** SDx is an extension of virtualization that further abstracts an application or function from its underlying hardware, separating the control and data planes and adding programmability. Beginning with software-defined networking (SDN), SDx now encompasses software-defined storage (SDS), software-defined computing, software-defined security, software-defined visibility (SDV), and software-defined data centers (SDDC), among others.

» **Automation and orchestration:** *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. Both automation and orchestration are critical technologies in the cloud, enabling day-to-day tasks such as provisioning, patching, and resource management to be performed at massive scale — across hundreds of thousands (even millions) of servers and other cloud components.

» **Service-oriented architecture (SOA):** SOA is a software design in which modular web services are leveraged across a network to provide various application components. SOA enables businesses to improve agility and time-to-market (TTM), and is thus well suited for cloud computing applications — for example, through the use and reuse of Representational State Transfer (REST) application programming interfaces (APIs).

» **Microservices:** Like SOA, microservices are application building blocks comprised of small, independent processes and services. Microservices tend to be smaller than SOAs and an SOA can be comprised of multiple microservices — in that sense, a microservice could be a microcosm of an SOA!

- » **Containers:** Containers, such as Docker, are a type of operating system environment (OSE) virtualization that enables applications to be rapidly deployed (in seconds) and booted up (in fractions of a second). Containers allow a developer to package up an application with all the individual components it needs, such as libraries and other dependencies, and deploy it as a single package.

Seeing How Business Trends Are Precipitating Cloud Adoption

The cloud enables business agility through faster access to modern infrastructure, massive scalability and elasticity, higher availability, and faster time-to-market.

IT modernization and cost savings are among the top drivers for public cloud adoption. In RightScale's 2017 *State of the Cloud Report*, 94 percent of respondents (comprised of small-medium businesses and large enterprises) are in various stages of the journey to the cloud:

- » **Cloud watchers** (14 percent) are developing their cloud strategies and evaluating cloud options, but don't yet have any applications deployed to the cloud.
- » **Cloud beginners** (22 percent) have started working on initial cloud projects, but are still gaining comfort and experience in the cloud.
- » **Cloud explorers** (25 percent) have deployed multiple applications to the cloud and are exploring opportunities to improve and expand their cloud strategies.
- » **Cloud focused** (33 percent) organizations have, in many cases, adopted a "cloud first" or "cloud only" strategy, and are looking for opportunities to further optimize their cloud environments while reducing costs.

While the benefits of the cloud are many, accessing and monitoring traffic in the cloud is a challenge. Without granular access to traffic in the cloud, you may suffer from blind spots in your network that compromise application performance or security. Chapter 2 takes a closer look at these challenges.

- » Identifying network blind spots
- » Understanding visibility challenges in network interconnects
- » Looking at multiple cloud options
- » Addressing key visibility requirements in cloud environments

Chapter 2

Networking in the Cloud

In this chapter, you learn about network visibility challenges in cloud, traditional on-premises data centers, and hybrid cloud environments.

The Network Perimeter Is Obscured by Clouds

Moving applications and services to the cloud delivers increased agility at a lower cost — but there are many risks along the way and more complexity to manage when you get there. The challenge for IT is to ensure that the infrastructure delivering critical services and applications is reliable, fast, and secure. In addition, businesses need to control costs, which means exposing hidden problems or blind spots within the network, reducing mean time to resolution (MTTR) for problems, and creating a robust, resilient security architecture. Access to critical data in these virtualized networks and hybrid cloud environments, as well as dissemination to key analytics and monitoring tools, is more important than ever to ensure the reliability, security, and performance of mission-critical applications.

Unfortunately, most enterprises and service providers today struggle with these issues. Businesses experience several pain points, including:

- » Lack of intelligent visibility solutions for virtualized private or public cloud environments leads to elevated security threat exposure, as well as an inability to sufficiently monitor and troubleshoot critical events.
- » Blind spots have become such a serious security issue for enterprises and service providers that, according to the 2016 *Verizon Data Breach Investigations Report (DBIR)*, they prevent 75 percent or more of all businesses from knowing that they have suffered a security breach.
- » Most businesses do not get optimal results from their monitoring and security tools, even though they make significant tool investments.

Enterprises have invested in physical analytics and monitoring tools in their traditional data centers to achieve end-to-end network visibility across their physical networks with inline and out-of-band security and monitoring tools.

The physical monitoring infrastructure employs physical data access taps, aggregators and network packet brokers (NPBs), and various analytics tools connected to aggregators or NPBs for data collection and packet processing. Intelligent filtering (Layer 2 through Layer 7) after the NPBs aggregate raw data from the taps enables organizations to exclude certain traffic, such as streaming video, from inspection by your monitoring tools, thereby reducing bandwidth and load requirements. The combination of complete data access, intelligent visibility, and proactive monitoring creates a complete visibility architecture. Network taps can be installed on every network link that needs to be monitored. This includes taps at the top-of-rack switch, between the data center and the core network, and even inside virtualized servers to monitor traffic on virtual networks. Consequently, you can very easily see exactly what is going on with your applications in real time to address performance bottlenecks and possible indicators of compromise (IoC).

However, in virtualized data centers and cloud environments, this model breaks down because of the “virtual blind spot.” East-west

traffic between servers — particularly between virtual machines (inter-VM) and application or workload instances — in the data center or cloud typically never traverses a firewall and never hits a “wire” where it could be tapped for inspection, despite comprising the majority of network traffic today. This problem is further exacerbated by the distributed architecture and access requirements in the public cloud.



WARNING

In an effort to gain visibility of east-west traffic, many IT teams use techniques such as “hairpinning” to force network traffic through an inspection point. However, hairpinning increases network complexity and reduces efficiency by creating unnecessary choke points and potential points of failure in the network as well as adding congestion and latency on the network.

As more enterprises move their critical workloads from traditional on-premises data centers to virtualized, software-defined data centers (SDDC) and public clouds, they often face important questions associated with security, reliability, and performance of these services.

Network Interconnects in On-Premises and Cloud Contexts

Monitoring network interconnects between your on-premises data center and multiple cloud environments creates additional visibility challenges. Public cloud service providers offer different options for securely and reliably connecting to the cloud, but these options sometimes rely on specific hardware and configurations limited to certain vendors.

Additionally, virtual private networks (VPNs) that connect your on-premises data centers and users (and their mobile devices) to your various cloud environments encrypt traffic traversing the VPN, introducing further visibility challenges. The proliferation of VPNs and independent, “self-serve” network designs creates long-term support headaches for IT teams.

Considering Different Cloud Models

Organizations adopt different cloud models to support a diverse array of unique requirements. Options include:

- » **Public clouds:** Major public cloud service providers are rapidly expanding their data center footprints with hyper-scale deployments, characterized by continuous configuration changes based on demand. Although resource pooling and elastic scale are part of the cloud value proposition, the ability to monitor virtual traffic flows at the same scale has been limited.
- » **Private clouds:** Private clouds provide greater flexibility for organizations in their choice of deployments. However, having options can sometimes add complexity to the visibility architecture, particularly if it is based on proprietary technology. With multi-cloud prevalence, a visibility solution must be able to operate across platforms and offer a central way to configure across many platforms. For example, private clouds may use a variety of hypervisors and network virtualization platforms in their buildout, including:
 - Microsoft Hyper-V
 - OpenStack KVM
 - Oracle VM
 - VMware (vSphere, NSX, ESXi)
- » **Hybrid clouds:** Most modern enterprises do not live solely in public or private cloud environments but use a hybrid approach. You need to monitor your data in both the public and private cloud.

Many organizations run tens of thousands of instances (or virtual resources) in the public cloud. Often, these instances are logically separated into different virtual private clouds (VPCs) in Amazon Web Services (AWS) or network security groups (NSGs) in Microsoft Azure. Different departments or teams within an organization can build their own VPCs and NSGs, often leading to monitoring challenges.

For example, it's possible, even likely, that there will be overlapping IP addresses in different VPCs and NSGs. Managing

thousands of instances and the amount of traffic this creates, while dealing with potential IP address overlap, must be considered in a cloud visibility architecture.

One possible option is route data from the cloud to a physical network packet broker. However, this solution creates a choke point for traffic in the cloud and adds another layer of software to be managed.

An alternative approach uses peer-to-peer, containerized (such as Docker-based) agents in each segment rather than a virtual or physical packet broker. This peer-to-peer solution connects via a VPN with its own IP addressing scheme so traffic can be sent over the Internet instead of requiring a dedicated connection. Thus, this solution is more secure and less costly. The lightweight, containerized agents provide the ability to scale-out — as fast as the organization needs. Moreover, a containerized solution has the potential to be cloud agnostic, allowing seamless visibility across multi-cloud deployments that use the technology platforms of different vendors.



WARNING

According to RightScale's 2016 *State of the Cloud Survey*, 82 percent of enterprises have a multi-cloud strategy. On average, organizations are employing at least six clouds, evenly mixed between public and private. In the event of a breach, how can an enterprise assess who is at fault? What safeguards should you consider if your enterprise is deploying a multi-cloud approach?

Ensuring Resilience in Cloud Environments

Critical issues for enterprise IT teams to address when migrating mission-critical workloads to the cloud and implementing services across distributed, mixed on-premises and cloud infrastructures include:

- » **Infrastructure and tenant separation:** Private and public cloud service providers (CSPs) who own the virtualized infrastructure host workloads from multiple customers (tenants) on top of the same shared virtual fabric. Depending on how the CSP addresses confidentiality, integrity, and

availability of tenant workloads, this can increase the attack surface, risk compromise of sensitive customer data, and result in compliance and service-level agreement (SLA) issues. Since both the infrastructure owner and the tenant implement their own security, analytics, and application monitoring solutions, the design of intelligent visibility for data access and distribution must serve both the tenant and the infrastructure, separately.

Although the CSP does not and cannot generally store or access customer data, it still requires access to workload packet data to scan for botnets, enable distributed denial-of-service (DDoS) attack mitigation, and scan for vulnerability exploits, thus protecting customers from internal and external attacks. The tenant requires its own monitoring and visibility into workload packet data for big data analytics, access control, and higher level security intelligence based on internal security policies.

» **Getting the right data, to the right tool, at the right time in the right location:** Access to critical application data in these virtualized networks and cloud environments by monitoring tools is key to ensuring the reliability, security, and performance of mission-critical applications.

The enterprise branch office, for example, has limited or no local IT staff and relies on the fidelity of NetFlow for continuous network monitoring and maintaining application quality of experience. But, it is important to have granular access to application packet data in case of an event requiring further troubleshooting and fault analysis. Although some monitoring and analytics tools are deployed in a virtualized, private, or public cloud, most tools are deployed on-premises. However, as cloud adoption continues to increase, more tool providers are offering cloud options. Over time, a mass migration of tool vendors will occur.

In cases where your data access and monitoring occur within the same virtualized data center, copying raw packet data for continuous 24/7 monitoring is still not practical because around 80 percent of total data center traffic is east-west. However, getting the right information to the right tool in this environment is critical. Your solution must provide filtering and context-awareness of virtual traffic at the source, generating a continuous NetFlow feed for most applications, and packet data only where required

on demand. By ensuring that only the data that is needed is sent across the network, organizations can significantly reduce the cloud usage costs associated with their security and monitoring tools.

- » **Security:** Many organizations suffer a network crisis before IT departments realize the consequences of the loss of visibility. Security teams may not realize until the time of a malicious incident that they cannot see traffic between application or workload instances and other virtual resources. Without this visibility, they cannot detect and investigate the attack, identify compromised resources, take corrective action, and/or prevent future attacks. A virtualized data center or cloud environment is just like any other segment of your network — if it hasn't already been attacked, it's just a matter of time.

In the case of network performance monitoring, organizations typically do not appreciate the lack of visibility into virtualized resources until after they are well down the road of implementation. Virtualization and cloud migration projects, often driven quickly by valid business reasons, can introduce disconnects in the IT team, as network and security professionals often do not have access to the packet-level information that they need to do analysis.

- » **Elastic scale:** The most fundamental characteristics of an application or service designed to run in the cloud is elastic scale, characterized both by hyperscale deployments and by frequent changes to the current scale based on demand. On-premises or physical environments must rely on *vertical scaling* ("scaling up") — applying newer, faster, better, more expensive hardware to achieve scale. Throwing more hardware at a problem was a common approach. However, cloud environments are an innovation because they achieve performance through *horizontal scaling* ("scaling out") — adding virtual instances of the same tool or application to absorb the load. Once the enterprise moves application workloads into a cloud environment, IT needs to have a built-in strategy for achieving scale. This is true not only for the specific applications in question, but also for all the virtual network security and analytics tools such as intrusion detection systems (IDS) or forensic recorders. Your virtual tapping and virtual packet processing capabilities need to scale horizontally, on-demand, handling scale-out and



REMEMBER

scale-in events. They must be able to go from zero to thousands of instances and back to zero without requiring administrator interaction.

The ability to elastically scale out (and in) on-demand is a key characteristic and benefit of the cloud. Scaling up/down is also an option in the cloud, for example, in VM sizes, but scaling out is the native approach in a cloud architecture and is generally the preferred approach to address growth, capacity, and load requirements — particularly at massive scale.

With this constant scaling out (and in) in the cloud, monitoring and security tools are presented with ephemeral data, requiring tools in the public cloud to also scale dynamically, based on demand from the sources in the cloud.

- » **Performance:** A common misconception is that desired service performance in a cloud environment can be achieved by ensuring performance of a given workload, running as a VM. A cloud environment is much more likely to achieve performance through horizontal scaling. Thus, while maximizing the performance of an individual VM workload is important (scaling up), designing your cloud service to leverage the underlying virtualized infrastructure to ensure scale-out, elasticity, and multi-tenancy is also very important.
- » **Fault tolerance and reliability:** Virtualization and cloud technologies are complex. Many administrators underestimate the complexity involved in ensuring reliability, performance, and scalability of critical workloads in this mixed environment. This complexity drives the need for pervasive visibility that provides data access as well as intelligent packet processing and distribution, is also fault tolerant, highly recoverable from its own failures, and can scale as the service grows.

- » Comparing security issues in the cloud
- » Enabling end-to-end visibility in the cloud

Chapter 3

Ensuring Security and Compliance in the Cloud

Enterprise applications are increasingly being migrated from traditional on-premises data centers to the cloud. Providing visibility and security in such a diverse environment — consisting of public, private, and hybrid clouds, as well as traditional on-premises data centers — has become a huge challenge. Active threats, such as malware and data theft, must be identified and protected against in the cloud as well. Although these threats may operate similarly whether on-premises or in the cloud, securing your cloud environment is more nuanced. In this chapter, you learn about these challenges.

Recognizing Security Similarities and Differences in the Cloud

Security — and related issues of privacy, control, compliance, and transparency — have long been cited by organizations as a barrier to cloud adoption (see Figure 3-1). Organizations need to know that their data is secure, they can control access, and they can

maintain compliance with corporate and other relevant security mandates, such as the following:

- » **U.S. Health Insurance Portability and Accountability Act (HIPAA):** Sections 164.308 and 164.312 define logging and monitoring requirements for networks and systems that process and/or store protected health information (PHI).
- » **European Union General Data Protection Regulation (GDPR):** When it takes effect in 2018, the GDPR will establish protection requirements for organizations and cloud service providers regarding the personal data of EU citizens.
- » **Australian Privacy Principles (APP):** Contained in Schedule 1 of the Privacy Act, APP8 (cross-border disclosure of personal information) and APP11.1 (security of personal information) are both applicable to cloud environments.
- » **Canada Personal Information Protection and Electronic Documents Act (PIPEDA):** Organizations subject to PIPEDA must limit and monitor access to personal information.
- » **Payment Card Industry Data Security Standards (PCI-DSS):** Version 3.2 requirement 10 defines monitoring requirements for networks and systems that process sensitive data such as credit card numbers.



FIGURE 3-1: Security, privacy and control, compliance, and transparency concerns are the top barriers to cloud adoption.

The concern arises because, unlike a traditional on-premises data center, organizations working in the cloud have no physical access to the compute, storage, and network infrastructure — or even a defined space inside a data center. In the private cloud, physical access by customers is restricted and controlled by the data center operator, and operators offer varying levels of security. In the public cloud, customers have no physical access at all, because the environment is a distributed architecture designed for multi-tenancy.

Private cloud providers (such as VMware and OpenStack) and public cloud service providers (such as Amazon Web Services and Microsoft Azure) implement a shared responsibility model for security. The cloud service provider ensures that the cloud infrastructure itself is secure, while the enterprise is responsible for securing the actual services running on the cloud infrastructure.

The level of responsibility varies depending upon whether the service is Software-as-a-Service (SaaS), Platform-as-a-Service (PaaS), or Infrastructure-as-a-Service (IaaS). Understanding where your provider's performance and security responsibilities stop and your organization's start is vital (see Figure 3-2).




FIGURE 3-2: Shared security management responsibilities vary across SaaS, PaaS, and IaaS offerings.

Enterprises take on a larger share of responsibility for IaaS offerings than for PaaS and SaaS, including maintaining the integrity of their platforms, applications, and of course, data. In the case of PaaS, the enterprise is responsible for the applications but not the underlying operating system (OS) or middleware. Only under a SaaS model is the cloud service provider responsible for the application itself. Interestingly, SaaS is potentially the riskiest, because any employee can easily authorize use without the knowledge or oversight of IT, versus IaaS and PaaS solutions that typically require IT configuration and management. The role of a cloud access security broker (CASB) is to function as a gatekeeper, providing visibility into a SaaS environment. A CASB is typically a service or software that sits between the cloud and your on-premises environment.



TIP

The primary service offering from private cloud operators is typically IaaS, although PaaS and SaaS offerings may also be available. Public cloud operators typically provide a full range of IaaS, PaaS, and SaaS offerings.

Organizations must understand the security model of each cloud and plan accordingly. For example, Azure may have better capabilities in some areas, such as those relating to Microsoft Office applications, while Google Cloud may appeal more to developers.



REMEMBER

Security concerns should not deter organizations from using cloud services. The public cloud has the potential to be more secure if the right visibility and control solutions are used, including:

- » **Identity management:** Ensuring proper authentication of users and entities is a basic tenet of security — whether on-premises or in the cloud. Integration with existing directory services and support for multi-factor authentication are important capabilities.
- » **Access control:** Role-based access control (RBAC) helps organizations manage permissions and properly implement the principle of least privilege (a user or entity only has the rights necessary to perform a particular task or function).
- » **Data protection:** Encryption of data in transit and at rest is critical. Additionally, organizations must ensure that encryption keys are appropriately protected.

Gaining Visibility into the Cloud

Ensuring security in the cloud and hybrid cloud, as well as in on-premises environments, requires end-to-end visibility for proper monitoring and control. Simply accepting the cloud as a “blind spot” is not an option because the cloud is becoming the dominant mode of operation. A solution that can tap and filter data across these cloud environments is crucial to ensure complete visibility across the entire environment. Moreover, because the raw data can be overwhelming, you need to look beyond simple data access. Integrating a sophisticated security fabric ensures you can send the right data — without duplications, unnecessary protocols, and in the correct format — to the right security, analytics, and compliance tools.

The strong growth of public, private, and hybrid clouds, on top of traditional on-premises data centers and emerging enterprise private clouds, increases the size of the attack surface dramatically. Effective testing of new cloud-based applications requires simulation of cloud-scale applications and attacks. Likewise, effective monitoring in these multi-tenant environments, where usage is elastic and access is limited, requires its own type of visibility.

To secure cloud environments, a solution must have the ability to

- » Provide visibility to inter-VM and inter-instance traffic.
- » Work across multiple cloud platforms (or be provider agnostic).
- » Operate with or integrate with cloud-based security and analytics tools.
- » Retain the benefits of cloud, especially horizontal scale.

When evaluating visibility technologies, you must consider the advantages and disadvantages for how data is captured, processed, and managed in the solution to determine what is best for your environment.

Here’s an example. Using containerized software agents in public cloud has the following advantages:

- » Containerized agents reside within an instance or workload, so no security changes are required, meaning less likelihood of error or having an additional segment to manage.

- » No network changes are required for operation with networking services (such as load balancing).
- » No changes need to be made to the application being monitored.
- » Metadata about instances and workloads is readily available, simplifying analytics and other tasks.
- » Auto-scale is natively supported because the agent is embedded in the instance or workload.
- » Orchestration is easy because agents' deployment to instance and workload templates is automated.

A containerized software agent approach also eliminates the following potential disadvantages of traditional software agents:

- » Software agents compete for system resources. Look for containerized agents that have a small memory footprint and minimal processing requirements.
- » Agent-based models provided by vendors to support individual applications are often incompatible and have additive performance impact. It is preferable to use a single agent-based provider, which is integrated with security and analytics platforms, rather than multiple agents from different providers.
- » Smaller instances with already limited resources may be a challenge for installing traditional software agents.
- » Traditional software agents may use the same network interface as the instance; thus bandwidth sharing might be a problem. Look for containerized agents that can collect and leverage metadata instead of always relying on raw data (discussed in Chapter 4).
- » Misbehaving agents can be detrimental to production instances and workloads.

Early in a cloud migration, organizations may choose to backhaul data from the cloud to a legacy on-premises solution. This makes it easier to do a trial, cloud-based application deployment before decisions on native cloud monitoring have been made. For multiple reasons — elasticity, scale, and bandwidth cost — this is not an effective long-term solution. Consequently, it's a good idea to look for a containerized agent-based solution that makes it easy to achieve this kind of backhaul if that fits with your migration plan.

Finally, you should consider whether your solution presents any of the following issues:

- »» Creates a new network segment to manage.
- »» Introduces an additional layer of software where errors can occur.
- »» Doesn't sit within the existing security deployment, so the attack surface grows.
- »» Must be managed external to the cloud and is thus not capable of auto-scaling — human intervention and monitoring is a constant requirement.
- »» Relies on a single path or process so security becomes a single point of failure.

Though the approach of backhauling data may have a minimal immediate impact on how security is managed in an organization, it doesn't provide a sustainable solution. This method may be acceptable for pilot deployments or as a temporary quick fix, but it doesn't truly address your organization's visibility needs in cloud. Solutions that leverage the native capabilities of cloud technologies are more adaptable, scalable, and likely to evolve with the cloud technologies.

Finally, you should consider differences in how data is processed. A packet that is collected for visibility (for out-of-band monitoring) is duplicated. The original packet continues to be delivered to its recipient as defined in the packet header. The copy for visibility needs to be delivered to a destination other than the one defined in the packet envelope. Continuing with the earlier containerized software example, in this case, a packet can be encapsulated and sent through a secure overlay tunnel, where it transits the general network and is redirected to its new destination because of its new envelope. This all happens within the cloud environment, and captured packets can be delivered to another software agent potentially residing within a security or analytics tool instance or other destination. The agent running on the tool instance can terminate the overlay tunnel, decrypt and decapsulate the traffic, and present the traffic for further analysis and operation. Further advantages of this approach include:

- »» The overlay tunnel operates within the same networking constructs as the source instance, so all security constraints

are inherited from the source instance. That is, multi-tenant security boundaries put in place by the cloud service provider are adhered to without requiring any additional configuration.

- » The monitoring tools receive packets not originally addressed to the tool host, regardless of any security or architectural restrictions.
- » The visibility architecture fits within the same infrastructure and does not require separate management for the overlay.

Other approaches capture raw data or data that has been minimally processed and route it to an on-premises environment or, in some cases, a VM. This type of approach can be costly because even cloud-to-cloud connections are expensive and many cloud service providers charge per-gigabyte outbound data transfer rates. A solution that processes packets (the source data) in the cloud, whether public or private, will ultimately provide a better solution.



REMEMBER

The cloud by no means eliminates the need for strong, independent security and compliance practices at the application level. Quite the opposite is true. Organizations need to implement appropriate cloud visibility and control tools to mitigate security risks and threats.

- » Comparing metadata and raw packet data
- » Detecting faults in the cloud
- » Using big data analytics

Chapter 4

Monitoring Network and Application Performance in the Cloud

There are two important data types to consider for network visibility — data and metadata. *Data* is the actual content, for example, in a movie, whereas *metadata* is information about the movie, such as the title, cast, rating, and reviews.

Packet data is important for detailed analysis such as breach detection, data leak prevention, and troubleshooting; metadata is important for scale, macro analytics, and trend analysis. Having access to both is vital to achieving large-scale visibility that covers the entire spectrum of monitoring, analysis, and troubleshooting — as well as security scenarios.

In this chapter, I explain how you can use metadata to both manage and extract usable data from cloud application deployments. I also show how you can use metadata-based management to isolate critical packet-level data for specific scenarios.

Looking at Metadata and Raw Packet Data

Public cloud service providers offer a wide selection of instance types and sizes optimized to fit different use cases. Instance types comprise varying combinations of CPU, memory, storage, and networking capacity, and provide customers flexibility to choose the appropriate mix of resources for their applications and target workload. The network stack includes one or more instance sizes, allowing the customer to optimize resources to the requirements of the target workload. The network stack, which feeds the instance and, in turn, the application, is abstracted from the customer. This abstraction eases the provisioning and management of the physical infrastructure, but introduces challenges in providing visibility at different layers of the network stack — especially when the application misbehaves or there is a connectivity issue.

Within an on-premises or hosted private data center, your company's engineers have monitoring tools that can access the underlying infrastructure, easing problem diagnosis. In the public cloud, access is limited to log data, which does not provide any actionable information. Thus, packet data becomes increasingly valuable in troubleshooting performance issues. From the perspective of the OS and application, all services look like a physical network, although the actual network is abstracted away from the user by the cloud service provider, who may offer minimal networking metrics.

A major challenge in monitoring traffic between virtual workloads or applications (known as inter-VM or east-west traffic) in the cloud is the amount of data that must be sent to and processed by the monitoring tools. This traffic can create a significant burden on network and computing resources, as well as the monitoring tools, because it typically constitutes more than 80 percent of total traffic in the data center.

Although it may be tempting to deliver a full copy (raw data) of all inter-VM traffic to your monitoring tools, it may not always be required and, depending on your industry or applications, it may not even be practical — plus, it can be costly.

In most situations, you don't need to inspect all the traffic all the time, and some of the data may not be relevant. For example, in YouTube, Netflix, or other video streaming, the video payload is

heavy but is usually not of much interest to the monitoring tools, so filtering this traffic out often makes sense.

One viable option to lower the amount of traffic sent to your monitoring tools is to leverage metadata instead of raw data. Metadata, such as enhanced NetFlow, can significantly decrease the volume of network traffic and provide extremely valuable information.



TIP

Metadata can be used with NetFlow collectors to facilitate troubleshooting and threat detection.

Organizations increasingly rely on an expanded view of the characteristics that can uniquely identify personal and intellectual property stored in both public and private clouds for compliance, marketing, personalization, and more. High-level information, such as when and where something was changed or updated, the type and format of the data, the source, and the key characteristics (metadata) of the data, create new value streams and mitigate information risk. By leveraging metadata intelligently, organizations can extract new business value from information and potentially introduce new business models for value creation.

GOING BEYOND NetFlow WITH IxFlow

Ixia allows you to enrich NetFlow records with value-add extensions. With IxFlow, you can determine what additional information to send to your tools, such as:

- Geographical information (including region IP, latitude, and city name)
- Application ID or name
- Device types present on the network
- Browser types, including secure-sockets layer (SSL) cipher
- Subscriber-aware reporting (provides detail on application and handset or device type for mobile users)
- Hypertext transfer protocol (HTTP) uniform resource locator (URL) and hostname for web activity tracking
- HTTP and domain name system (DNS) metadata for rapid breach detection
- Transaction latency for application performance tracking

INTRODUCING IXIA'S CloudLens


Ixia's CloudLens is a platform for public, private, and hybrid cloud visibility that enables you to easily manage multiple segments effectively.

The platform's public arm, commonly referred to as CloudLens Public, is the first Software-as-a-Service (SaaS), network-level solution that provides Visibility-as-a-Service (VaaS). Designed from the ground up to retain the elastic scale, flexibility, and agility benefits of the cloud, CloudLens Public provides an intelligent and automated cloud VaaS that scales with public cloud infrastructures.

Ixia's CloudLens Public can auto-scale and its design is cloud-provider agnostic because it is cloud-native. At its core, it's an implicit microservices architecture that is orchestrated via application programming interfaces (APIs) — a serverless design. With this design, CloudLens can provide a highly available, scalable solution across platforms.

CloudLens Public has two core components (see the figure):

- A SaaS visibility management platform where users can configure visibility and define filtering.
- Sensors and connectors that are containerized, Docker-based software that sit within the source and tool instances respectively. Because the sensors and connector sit in the instances, they can access metadata, which they then share with the management platform.



Ixia's CloudLens Public is comprised of source and tool sensors and a centralized management platform.

Metadata information is actionable — it can help uniquely identify individual system characteristics like OS, kernel module, CPU, memory, cloud service provider, region, zone, and so on.

In the management platform, visibility is configured by drawing a connection (the secure visibility path) from source instances to relevant security and monitoring tool instances and defining filtering rules. This ensures packet data is processed and routed according to a company's policies. This visibility is scalable because the management platform has a smart search capability that allows users to create source groups and tool groups based on metadata as the criteria. The metadata can also be user-defined, allowing maximum flexibility. CloudLens Public uses metadata from cloud platform instances to classify them; because metadata inherently exists for each new instance that is created, the platform automatically knows how to treat it and which security and monitoring policies to apply. Consequently, packet data from instances is appropriately filtered and routed to security and monitoring tool instances, without requiring human intervention. This approach retains the scalability and elasticity of a cloud visibility solution.

CloudLens filters at the source sensor. At that time, packets that are collected for visibility are duplicated. The original packet is delivered as it was intended. The duplicate copy is encapsulated and placed on an encrypted secure visibility path, an "overlay tunnel," and is redirected to its new destination in monitoring or security tool instances based on its new envelope. The sensor on the tool instance terminates the overlay tunnel, decapsulates the traffic, and presents the traffic to the tool instance for further analysis.

Fault Detection, Incident Analysis, and Reporting

Fault detection is a challenge in virtualized cloud environments because of the transient nature of the cloud and because the customer does not own the infrastructure. Performance — such as a web application not letting users log in, or slow response times — can be used as a threshold to trigger an analysis of data to determine what is happening in the cloud environment.

A cloud visibility solution provides access to data for various monitoring tools, such as application performance monitoring

and network performance monitoring tools. These monitoring tools can either keep a running analysis of the data or metadata, or log the data for later analysis. When an incident occurs, that information is then accessed to track down what is happening and determine the root cause of the incident. Because it is not as straightforward as physical environments, you need constant access to see what is occurring in the cloud.

Proactive monitoring is another technique to enable cloud visibility. By deploying hardware or software endpoints in the network or in the cloud, data such as Internet Control Message Protocol (ICMP) pings or other synthetic traffic can be collected between cloud components to provide baseline measurements of normal activity in the cloud.

Big Data for Analytics

Data — and data about data — is being collected everywhere today. This rise of “big data” requires new forms of integration to uncover hidden value in diverse and complex datasets of massive scale. The public cloud is a powerful platform, not only to collect and store big data, but also to perform complex, large-scale computing tasks including database and application services.

The massive amounts of data collected, processed, and stored in the cloud presents a challenge for cloud visibility. Organizations must be able to extract and use big data to use it effectively. Additionally, they must be able to see what is happening in the cloud while they are processing these massive volumes of data — which requires visibility. Because this data is usually collected from a multitude of sources, integration is problematic. A cloud visibility solution acts as the glue that brings all this data together.

Although big data is a challenge for cloud visibility, it also plays a key role in enabling key insights for cloud monitoring, by essentially doing analytics on the volumes of big data being processed.

Basic metadata (discussed earlier in this chapter) can be used to automate classification of system data and can be used to improve performance metrics. The primary challenge in providing visibility with metadata is to capture, analyze, and process it in an expedient manner. Metadata-level information enables you to understand what is happening in the cloud. But analysis requires packet-level data and/or big data analytics.

- » Examining government and health care private cloud use cases
- » Delving into e-commerce and technology services public cloud use cases

Chapter 5

Exploring Cloud Use Cases

In this chapter, you learn about common use cases in the cloud for an end-to-end visibility architecture.

Private Cloud

In the following private cloud customer stories, you learn how Ixia visibility solutions helped a U.S. government agency and a health care company achieve complete visibility of east-west traffic in their private cloud environments.

Government

U.S. government agencies have policies and procedures for monitoring everything that occurs within their data centers. The IT team for a large federal agency was already using monitoring tools and taps to perform extensive traffic inspection and maintain a high level of visibility in the agency's physical networks. But many of its data centers had become highly virtualized with most servers being deployed as virtual machines (VMs). The team recognized that visibility gaps — created by “east-west” traffic flowing between VMs — existed in the data centers, and that a more reliable and holistic view of the network, traffic flows, and problem points was needed.

The agency's IT team implemented an end-to-end visibility architecture to expand beyond basic monitoring and debugging. The solution uses Ixia's CloudLens Private virtual tapping (vTap) capability and physical network packet brokers (NPBs) to bolster security, simplify management and configuration, and improve the efficiency of their monitoring tools. Benefits of the solution include:

- » East-west, inter-VM traffic is 100 percent visible.
- » Non-proprietary, tool-agnostic CloudLens Private vTap and NPBs send traffic to any existing security and performance monitoring tools.
- » Bandwidth and resources are saved by filtering traffic when it is tapped.
- » Deployment to an entire virtual data center is quick and easy with VMware's vCenter.
- » Hitless plug-in installation is possible and no VM maintenance mode is required.

The Ixia visibility architecture concept provides visibility across the agency's seven private clouds in a systematic way with a holistic solution instead of a product-oriented approach.



TIP

Ixia's CloudLens Private vTap supports vMotion — allowing running VMs to migrate from one physical server to another, without downtime and without losing monitoring efficiency.

Health care

A health care organization's data center was recently hacked, resulting in costly damage to both its reputation and bottom line. Its aging infrastructure was badly in need of upgrades to help meet regulatory compliance standards, including the Health Insurance Portability and Accountability Act (HIPAA) and Health Information Technology for Economic and Clinical Health (HITECH) Act.

The organization deployed an Ixia visibility architecture in its two new greenfield data centers for out-of-band (OOB) and inline security — including intrusion detection and prevention systems (IDS/IPS), secure-sockets layer (SSL) decryption, and data loss prevention (DLP) — as well as application performance monitoring (APM) and network monitoring (NetFlow).

Benefits of the newly deployed visibility architecture include:

- » Ability to capture all traffic seamlessly and integrate within any network configuration or security design.
- » Integration between all Ixia devices and technologies including iBypass, Vision xStream 40 network packet brokers, CloudLens Private vTap capability, Vision 7300, and Ixia's PacketStack and AppStack capabilities.
- » Repeatable architecture design for future data center buildouts.

Public Cloud

Public (and hybrid) cloud use cases can be somewhat more challenging than private cloud use cases. Because a public cloud is a multi-tenant environment, the visibility solutions available to a customer are typically more limited than in a private cloud. In the following customer stories, you learn how Ixia visibility solutions helped a retail company gain visibility in a hybrid cloud environment to automate the management of elastic load demand on its e-commerce site, and how a web teleconferencing company manages quality of service (QoS) with complete application visibility in the cloud.

Retail (e-commerce)

A large U.S.-based retail company with a traditional 3-tier (consisting of application, web, and database servers) architecture in its data center is moving to a hybrid cloud model (see Figure 5-1).




FIGURE 5-1: A retail company moving e-commerce infrastructure to a hybrid cloud model.

Customers access the company's e-commerce web servers from various devices, such as smartphones, tablets, and desktop PCs. Prior to deploying the Ixia CloudLens Public solution, the retailer had to create separate visibility policies for the web servers in its public and private clouds. To provide a concise and comprehensive solution, the company installed CloudLens Public containers on the web and application servers, thus creating a visibility path between the sources and various multi-vendor tools.

When demand on the e-commerce site increases, such as during the holiday season (including Black Friday and Cyber Monday), the web servers automatically expand into the public cloud to support the increased load — a practice known as *cloud-bursting*. In hybrid clouds, an application ordinarily runs in a private cloud or local computing environment. When the application requires additional computing power or storage, the app bursts into the public cloud to obtain the needed resources. The same visibility policies that are applied to the on-premises web servers are applied to the on-demand web server instances created in the public cloud.

Technology services

A large web conferencing company is planning to move its infrastructure, which provides online meeting, desktop sharing, and video conferencing capabilities, to the cloud to increase scaling capacity. During peak hours of operation — particularly during the early part of the workweek — the number of customers accessing these services increases exponentially. The company currently has no way to get cloud traffic to the company's quality of service (QoS) tools to monitor, detect, and diagnose voice quality issues. In addition, web traffic may spike at peak times, requiring a quick ramp-up in cloud-based application server and monitoring tool instances.

Although the company can run voice quality tools in the cloud, it has no way to get voice packets to those tools. At the same time, the company only wants to send voice traffic — not domain name service (DNS) or other application traffic. Finally, as traffic scales up and down, load must be distributed dynamically.

Ixia's CloudLens Public solution delivers the specified application traffic to monitoring tools in the cloud and allows the company to specify which traffic is sent to the selected tools and which traffic is excluded. In addition, CloudLens Public automatically and dynamically rebalances traffic from application instances to monitoring tools without requiring a packet broker or load balancer. All the intelligence is distributed to endpoints.

IN THIS CHAPTER

- » Calling for a “mostly cloud” future
- » Getting real about virtual and augmented reality and other innovations
- » Enhancing user experiences with contextual cloud networks
- » Moving the cloud closer to enterprise users and applications

Chapter 6

Forecasting the Future of Cloud

In this chapter, I dust off the crystal ball and lay out a few bold predictions about the future of the cloud.

The Cloud Is Here to Stay

Anyone who might have thought the Internet was just a passing fad might have similar misgivings about the cloud. However, the cloud — like the Internet — is here to stay.

In the near future, a corporate “no cloud” policy is likely to be as rare (and nonproductive) as a “no Internet” or “no cellphone” policy is today.

The cloud has also transformed the way we use the Internet. Rather than simply being a transport network, the cloud makes the Internet a destination. This change has major implications for Internet service providers (ISPs) that provide asynchronous

data rates (different download and upload speeds). Upload speeds will become as important as download speeds because users are increasingly uploading more content to the cloud.

Thus, the need for visibility in the cloud is likewise here to stay. This requirement will become even more important as the cloud continues to evolve and expand.

The Cloud Will Evolve as Technologies Emerge

The cloud will continue to evolve as new technologies emerge, including:

- » **5G:** The next-generation mobile networks, 5G, will provide aggregate speeds up to 20 gigabits per second (Gbps) and less than one millisecond (ms) latency. To put those numbers into perspective, today's fastest 4G LTE Advanced Pro networks advertise speeds up to 1Gbps with less than 10ms of latency. Fixed wireless access will be one of the earliest 5G use cases, providing 5G cellular service for residential Internet access.

Download *5G For Dummies* from www.ixiacom.com to learn more about 5G networks.



TIP

- » **Internet of Things (IoT):** Various industry estimates predict there will be anywhere from 20 billion to 50 billion smart, connected IoT devices by 2020 — everything from cars to washing machines to heavy industrial machinery. These devices will transmit and receive all sorts of data to various cloud environments. Ericsson predicts that monthly mobile data traffic alone could easily exceed 50 exabytes (EB) by 2021.

One terabyte (TB) equals 1,024 gigabytes (GB); one petabyte (PB) equals 1,024 TBs; and one exabyte (EB) equals 1,024 PB (or approximately one billion GB).



REMEMBER

- » **Virtual reality (VR) and augmented reality (AR):** VR technology creates a fully immersive, computer-generated experience that simulates or re-creates real-life situations and environments. In contrast to VR, AR layers computer-generated images and enhancements onto a real-world

situation or environment to provide a more meaningful context for user interaction. Next-generation VR and AR experiences will have “six degrees of freedom” (6DoF) — the next level of immersion — allowing users to move within and intuitively interact with the environment. 6DoF experiences, which are available in video games today, allow users to move spatially through the environment just by walking or leaning their heads forward. Many industries, such as tourism, education, and other forms of immersive video will flourish as 6DoF technologies evolve. These next-generation VR and AR solutions will leverage the cloud to upload and download immersive content for users.

As these technologies evolve, the need for pervasive visibility will, in many cases, become more critical. Certain IoT use cases (for example, self-driving vehicles and smart electrical grids) will require robust security to ensure safety.

Contextual Cloud Networks

Today, contextual computing provides users with a more personalized experience by delivering relevant content based on information including the user’s individual preferences, current location, time of day, and browsing history.

Contextual cloud networks will take this concept to a new level by interconnecting this data to deliver fully immersive experiences and services in the cloud. This, in turn, will lead to the emergence of specialty clouds that perform specific tasks or functions. For example, clouds built and optimized for transportation, health care, and other critical real-time functions; transactional clouds for retail and banking activities; social networking clouds and media/gaming entertainment clouds; and many others.

As the cloud becomes more purpose-built and personal, the need for visibility and security to protect privacy will become increasingly important.

Fog and Mobile Edge Computing

The cloud will also evolve as new innovations such as mobile edge computing (MEC) and fog computing develop. MEC is a network architecture that moves cloud and IT services to a mobile network operator's edge network to reduce network congestion and improve application performance. Autonomous (driverless) vehicles are an example of an application that may leverage MEC in the near future to provide extremely low latency and high bandwidth for critical functions.

Fog computing decentralizes computing resources, enabling processing, analysis, and applications to take place closer to the sensor, device, or system that creates the data. Fog computing thus extends and distributes cloud resources and services to

- » Increase efficiency.
- » Reduce network traffic and latency.
- » Improve security between the cloud and enterprise.

Today's hyperscale cloud providers will continue to deliver massive scale and flexibility to their customers for the foreseeable future. Edge data centers (largely comprised of hybrid clouds) will supplement these hyperscale cloud environments, providing faster responsiveness to customers that are geographically closer to these data centers, while leveraging the massive scale and elasticity of hyperscale clouds. Finally, microscale data centers (comprised of private and enterprise clouds) will permeate on-premises environments, providing real-time access to applications and data closest to the end user or consumer of information.

As the cloud edge becomes more dispersed, it will become less distinct, requiring tools that extend visibility into many diverse environments — microscale data centers at the bottom, pushing data up to edge cloud data centers that, in turn, push compute and data up to hyperscale cloud providers. To ensure visibility and security, you need a comprehensive cloud visibility architecture — an “all-seeing eye” — sitting atop this cloud pyramid.

IN THIS CHAPTER

- » Planning your trip to the cloud and matching your skills and experience
- » Avoiding proprietary cloud offerings
- » Ensuring complete visibility in your network and applications
- » Implementing security best practices and maintaining compliance
- » Developing a disaster recovery capability
- » Meeting application performance requirements
- » Aligning application workloads to cloud pricing models

Chapter 7

Ten Important Considerations for Your Journey to the Cloud

In this chapter, I describe ten important monitoring and visibility considerations for the cloud.

Migration Plan

Some of your application workloads may never be migrated to the cloud. It's important to develop a migration plan to ensure that only appropriate workloads are hosted in the cloud. Your plan can likewise ensure that a visibility architecture is designed to provide

end-to-end monitoring and control of applications wherever they are hosted — whether in public, private, or hybrid clouds, or in an on-premises data center.

Skills and Competency

When migrating to the cloud, look for cloud offerings, including your visibility platform, that align with the skills and experience of your existing staff. For example, if your IT staff has extensive knowledge of VMware vSphere or Microsoft Hyper-V, you should look for a cloud provider and visibility that supports those hypervisors. Similarly, if your developers use Docker containers extensively, look for a cloud offering that provides this option.

Vendor Lock-In

As with any technology solution, avoiding vendor lock-in is a key consideration. Look for cloud offerings that are standards-based and not proprietary. Application workload portability between on-premises data centers and different clouds, or even among clouds, should be a primary consideration. Design an end-to-end visibility solution that works across different clouds and on-premises data center environments, independent of a cloud service provider's offerings.

Network Visibility

A robust visibility architecture needs to monitor not only your network connections to the cloud, but also your network connections — including east-west traffic between VMs or instances — in the cloud. Also, not all workloads will necessarily move to the cloud, so look for a provider that can offer a complete visibility portfolio that interoperates with both physical and cloud environments.

Application Monitoring

Do your monitoring tools provide segmented or centralized application monitoring capabilities? For organizations that have tens of thousands of instances (or virtual resources) running

in the public cloud, you need both segmented and centralized capabilities.

For example, within Amazon Web Services (AWS), customers may segment their instances into different groups using virtual private clouds (VPCs). Similarly, Microsoft Azure customers might segment their instances into logical network security groups (NSGs). Your network and application monitoring tools need the ability to span these segments while maintaining their logical structures and providing centralized access to the data.

Ixia CloudLens Public provides this comprehensive, yet granular level of visibility for your network and application monitoring tools using peer-to-peer, Docker-based agents in each segment, such as a VPC. These agents communicate with each other and directly with a Software-as-a-Service (SaaS) based management component using metadata (discussed in Chapter 4), rather than backhauling all the traffic to a corporate data center.

Security Best Practices

Ensure that your cloud service provider(s) enforce security best practices, including the principle of *least privilege* (users only have the level of permissions necessary to perform an authorized task or function) and *defense in depth* (layered security mechanisms and controls to ensure robust security). Also look for security options that are integrated with your visibility solution, for a more seamless deployment, less hassle, and strong service-level agreements (SLAs) to ensure you can see everything in your on-premises, public cloud, and private cloud environments.

Compliance Requirements

Remember that cloud service options — such as Software-as-a-Service (SaaS), Platform-as-a-Service (PaaS), and Infrastructure-as-a-Service (IaaS) — are based on a shared security model, but the security of your data is ultimately your responsibility. This includes issues such as data locality (where the data is actually stored) and encryption.

While cloud service providers design their data centers and service offerings with security and compliance considerations top of mind, how you connect to and use the cloud will largely impact your security and compliance posture.



TIP

Ensure your visibility solution meets security and compliance requirements for complete visibility while providing the ability to keep personally identifiable information (PII), protected health information (PHI), and other sensitive data secure during packet inspection.

Disaster Recovery

A properly planned cloud migration can provide your organization with cost-effective disaster recovery capabilities. Ensure that your cloud provider's data centers are designed with a robust architecture to ensure resilience in the event of a disaster impacting the cloud provider directly, as well as failover capabilities in the event of a disaster striking your on-premises data centers.

Performance

Ensuring your cloud service provider can meet your application requirements in terms of application/network bandwidth and latency is another important consideration. Look for SLA performance requirements — not just availability — and implement a visibility architecture that helps you monitor and enforce SLA compliance.

Pricing Models

Many, but not all, cloud services use a subscription-based usage model. Matching your application workloads to the appropriate pricing model is critical to ensure you don't end up with unexpected charges every month. For example, if your applications upload significant amounts of data, it may be better to keep them on-premises rather than migrating them to the cloud if your cloud service provider charges for data egress. To achieve the benefits of the cloud across your environment, ensure your visibility, security, and monitoring solutions are designed and deployed to help you optimize your spend, whether on premises or in the cloud.



SEE INSIDE ANY CLOUD

Total Visibility
with CloudLens

ixia

www.ixiacom.com/cloud

Move critical workloads safely to the cloud

As enterprise IT teams move from on-premises data centers to virtualized, software-defined data centers and public clouds, they address these important questions: how can we ensure availability, reliability, and performance of our mission-critical applications? How do we get critical data to security, analytics, and monitoring tools? How can we tell which applications are suitable for cloud and plan a successful migration? This book helps you answer these questions for your organization's journey to the cloud.

Inside...

- Learn the reasons for cloud adoption
- Address network visibility challenges
- Ensure security in the cloud
- Monitor application performance
- Explore cloud visibility use cases
- Forecast the cloud's future

ixia

Lawrence C. Miller has worked in information technology for more than 25 years. He has written more than 60 *For Dummies* books.

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

for
dummies[®]
A Wiley Brand

ISBN: 978-1-119-42449-9
ixia part number:
915-8158-01
Not for resale

WILEY END USER LICENSE AGREEMENT

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