

the  
**GORILLA**  
**GUIDE**® to...



# Navigating Data Protection and Disaster Recovery

How Support for Container Data  
Protection Is Standardizing

**JOEP PISCAER**

# Navigating Data Protection and Disaster Recovery

---

By Joep Piscaer

## TABLE OF CONTENTS

---

Introduction.....	4
Understanding ‘as a Service’.....	5
SaaS.....	6
Public Cloud VMs and IaaS.....	12
PaaS and Serverless.....	15
Edge.....	17
On-Premises.....	19
End-User Devices.....	19
Things Will Never Be the Same.....	21

Copyright © 2022 by ActualTech Media

All rights reserved. This book or any portion thereof may not be reproduced or used in any manner whatsoever without the express written permission of the publisher except for the use of brief quotations in a book review. Printed in the United States of America.

**ActualTech Media**

6650 Rivers Ave Ste 105 #22489 | North Charleston, SC 29406-4829  
www.actualtechmedia.com

# Publisher's Acknowledgements

---

## EDITORIAL DIRECTOR

Keith Ward

## DIRECTOR OF CONTENT DELIVERY

Wendy Hernandez

## CREATIVE DIRECTOR

Olivia Thomson

## SENIOR DIRECTOR OF CONTENT

Katie Mohr

## PARTNER AND VP OF CONTENT

James Green

---

## ABOUT THE AUTHOR

**Joep Piscaer** is a seasoned IT professional, with 10-plus years experience as a CTO, head of IaaS and infrastructure, (enterprise) architect, and technical consultant. His specialization is in infrastructure, cloud, and way-of-work (DevOp, Infrastructure-as-Code). He has built Infrastructure-as-Code toolchains, IaaS platforms, transformed (infrastructure-focused) organizations to DevOps and Infrastructure-as-Code ways of work.

# Introduction

---

Welcome to The Gorilla Guide® To... Navigating Data Protection and Disaster Recovery, Foundation Edition! If you're confused about the ways to protect and recover your data in this new era of dispersed computing, you've come to the right place!

By “dispersed,” we mean that your data, applications, and even your infrastructure can live pretty much anywhere. That brings with it special challenges that weren't even conceived of in the days of self-contained data centers. One of the thorniest of those challenges is keeping the bad guys away from your data—then, if they *do* find a way in and wreak havoc, making sure that your business can recover from it.

---

**The typical slightly higher licensing or subscription cost of SaaS is offset by the savings in IT staffing costs.**

---

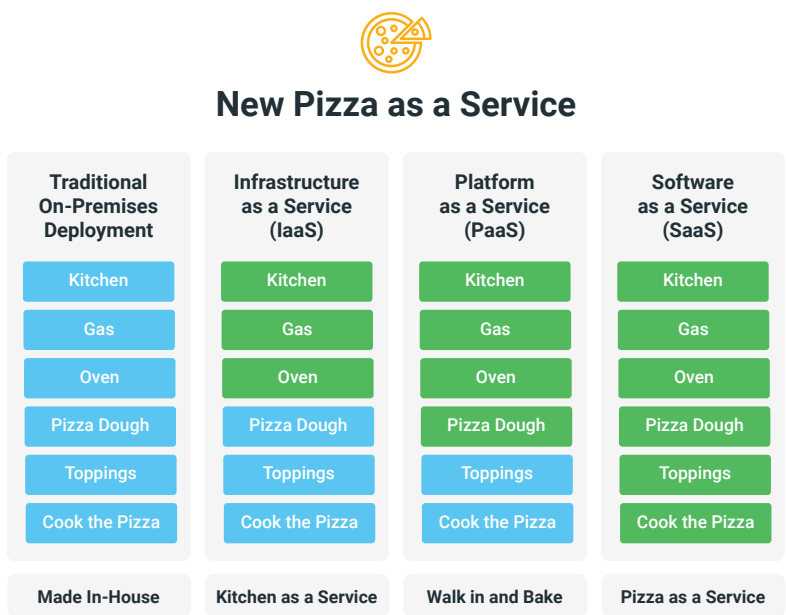
To that end, we'll look at data protection across the board—Software-as-a-Service (SaaS) applications, public cloud vendors, containers, Platform-as-a-Service (PaaS) and serverless, edge computing, and even on-premises and end-user devices.

With this knowledge, you'll be able to set yourself up for success for years to come. So let's enter the jungle together, and start hacking our way through!

## Understanding 'as a Service'

The Pizza-as-a-Service diagram in **Figure 1** helps clarify the differences between SaaS, PaaS, IaaS, and on-premises. You may find it useful throughout the article to refer back to it as you read.

Now we'll look at how each of these deployment and management models is impacted by the growth of the Internet.



**Figure 1: Pizza as a Service**

# SaaS

---

Let's start with one of the most-overlooked—and more difficult—ways to implement data protection: SaaS.

SaaS is great for commodity applications, like office productivity suites, email, accounting software, CRM and ERP, and it's no wonder that many enterprises have made the switch to SaaS for the bulk of their business application portfolio.



**SaaS is a category of web applications delivered as a service, over the Internet.** You, the customer, can just use the software, without worrying about operating, managing, or delivering the application. Its main benefits are easy onboarding and relative low cost due to economies of scale.

Microsoft, Google, and many others offer application suites that encompass the most common functionality businesses need, and they, and others, offer additional software for more specific applications for CRM, ERP, and marketing automation.

SaaS saves IT a lot of time and effort in operating and managing these applications, reduces the cognitive load on IT teams, and frees them up for more business-specific IT projects. The reduced operational complexity makes IT more agile

and more flexible, which translates to less organizational inertia. We all know that more nimble organizations are more likely to succeed because they can adjust to changing market circumstances more quickly.

The typical slightly higher licensing or subscription cost of SaaS is offset by the savings in IT staffing costs. The operational cost model, or OpEx, translates into lower long-term investments, allowing the cost to be adjusted monthly, quarterly, or yearly, which again offsets the slightly higher subscription cost. Flexibility can be of far greater value to an organization than the absolute lowest cost.

So far, so good: SaaS seems great. But it's not the end-all be-all, and IT still has to think about many nonfunctional aspects to guarantee availability, performance, security, and business continuity.

## **THEIR SERVICE, YOUR DATA**

What may come as a surprise, though, is that SaaS does not cover all of IT's responsibility.

The service delivers access to the application and includes keeping the service available according to the service-level agreement (SLA) and other service agreements regarding software and infrastructure scaling and upgrades, security, and more.

And while SaaS does have provisions relating to disaster recovery and data protection, those are put in place only to honor agreements to keep the service running.

But your data is still your data, whether it runs in SaaS or in a self-hosted application. The SaaS provider is not, and should not be, entirely responsible for your data. They merely provide disaster recovery and data protection for the service itself. And while they usually include some kind of bare-bones disaster recovery or data protection for your data, the scope of that protection is limited to whatever covers their liability. In other words, they're not in the business of providing data protection, and when push comes to shove, you're still responsible for protecting your own data.

## PROPERLY PROTECTING YOUR DATA

The standard data protection features offered probably won't be sufficient to meet your needs. That means you'll likely be on your own for more granular and more complete data protection.



**Adding third-party data protection to cover your SaaS applications is the first step toward properly protecting your data.** Factor in data protection, among other nonfunctional requirements, like security, as you select a SaaS vendor.

In practice, this means you'll have to resort to a third-party data protection solution that helps you granularly protect the data in that service, storing copies on different media and creating an off-site copy for DR purposes.

## NONSTANDARD DATA PROTECTION PROTOCOLS

Unfortunately, data protection solutions have to deal with nonstandard data protection protocols. Each SaaS service has its own methods of accessing the data inside the service for data protection purposes, and backup software vendors have to support each and every one. Because each integration is different, and all are complex to create and maintain, the coverage of any single data protection solution hasn't been great.

How different this is compared with the generally accepted standardization of virtualized environments, where VM-based snapshots, Microsoft VSS, and other application-aware mechanisms create wide coverage for granular, consistent backups across the board.

---

**Cloud and SaaS really have taken a few steps back in terms of data protection compatibility, and this is often unnoticed until it's too late.**

---

Cloud and SaaS really have taken a few steps back in terms of data protection compatibility, and this is often unnoticed until it's too late.

## SAAS PATCHWORK

Moreover, you'll likely use more than one piece of SaaS software, so the data protection puzzle becomes even more complex, because you'll have to find one, or more realistically two or three, backup solutions to cover your specific SaaS patchwork. Those specifics dictate which two or three backup vendors you should choose, increasing complexity and creating highly specific solutions that may be hard to maintain in the future.

In practice, most backup vendors will have support for the most popular SaaS services, including Microsoft<sup>365</sup> and Google Workspace, as well as applications like Salesforce, Oracle, and SAP, but as you venture into more specific and niche applications, backup support decreases rapidly.



**Determine which backup vendor most widely supports the current SaaS landscape to create the best coverage in a single backup solution.**

Repeat this with the remaining solutions if possible and resort to specific solutions where needed.

## MORE THAN JUST SAAS

Unfortunately, there's more to backup than just SaaS. Even if you're able to cover all of your SaaS applications, there are still traditional VMs, on-premises applications and databases, file servers, cloud file and object services, cloud VMs for

each of the public clouds, and containers. Reality is messy, and technologies often overlap for a good number of years: As one is sunsetting, the other is gaining traction.



**SaaS is only one category that requires data protection coverage.** Don't forget public cloud (VMs, object and file storage), on-premises, and container support. Look at the big picture and prioritize and optimize your data protection strategy globally.

As you investigate various solutions, you'll quickly see that the landscape of SaaS, on-premises VMs, cloud IaaS, file and object storage, containers, and on-premises applications is complex, and it's near-impossible to find the perfect mix of backup vendors to cover all your bases.

While the state of data protection in the public cloud and SaaS is improving, it's far from perfect. Consolidation will continue, with backup vendors increasing support for SaaS across the board and, hopefully, some standardization in data protection protocols in SaaS.

## Public Cloud VMs and IaaS

Data protection in IaaS is equally fragmented and diffuse. Many IaaS and public cloud VM instances offer granular backup, but almost always in a walled garden. This forces you to use a point solution with narrow support for just that one

public cloud service, further complicating the data protection landscape. But, luckily, most backup vendors include support for all public cloud VMs and IaaS in their products.

And as you move applications to either SaaS (for commodity applications) or containers (for self-developed applications), you'll need to adjust your data protection strategy to prioritize support for those.

## CONTAINERS

In container land, the separation of data across containers, cloud object and file storage, cloud-native databases, and many other discrete services causes data management issues. The proliferation of both stateful and stateless data across all of these services makes it harder to cover every service, and the risk of blind spots makes it easy to forget crucial parts of the application. And incomplete backups are useless backups, so complete coverage is mandatory.



**Your existing data protection software won't have support for all cloud-native storage services**, so be prepared to include one or more specific solutions for a couple of those cloud-native storage services.

Luckily, there's light at the end of the tunnel. While the patchwork of services around containers is still somewhat incomplete, support for container data protection is standardizing.

Kubernetes certainly plays a big role in this, and its increasing popularity means it'll become easier to standardize on Kubernetes as the next control plane for creating backups. This will help backup vendors increase application coverage more easily.

Native data protection at the container or Kubernetes level is, of course, important for technological reasons, as it enables optimal integration for storage and consistency. But those features are now table stakes, and not what sets modern data protection solutions apart.

What does set one solution apart from others is enabling consistent management across the patchwork of VMs, containers, and SaaS services used in your organization. The ability to define data protection policies across different technologies helps protect entire applications, but the immense fragmentation of data across container images, CI/CD pipelines, Git repositories, cloud file and object buckets, cloud-native databases, third-party APIs, and much more makes cloud-native data protection a challenge the industry hasn't conquered yet.

Wide coverage of all of these to protect cloud-native applications is increasing and will continue to expand. We see vendors searching for answers and taking baby steps in the right direction, with incumbent vendors acquiring cloud-native data protection startups and other vendors increasing their coverage and adding support for containers and cloud-native services.

## DEVELOPERS ARE TAKING THE DATA PROTECTION REINS

Containers create a shift in who defines, provisions, and operates infrastructure. Developers have more control than ever over the infrastructure that runs their applications, and are less and less dependent on infrastructure and Ops teams.

However, data protection risks falling into the void between ITops and Dev teams, and organizations need features to help development teams apply data protection policies in the application development lifecycle directly, instead of filing support tickets, doing handovers, and waiting for other teams.

This has led to a big shift in emphasis, with vendors decoupling the consumption of backup policies from the backup software itself, moving consumption to the developer's workflows instead.

By integrating into CI/CD pipelines and creating declarative backup policy-as-code solutions, vendors enable developers to apply data protection policies to their applications natively, without hopping between interfaces. This seamless, fully integrated experience is key for developers, for whom data protection may not always be top-of-mind.

---

**Containers create a shift in who defines, provisions, and operates infrastructure.**

---

This self-service policy-as-code approach works equally well for cloud and DevOps engineers alike, decentralizing the consumption of backup policies further while ITOps creates and retains control of the policy definitions, allowing for centralized policy definition to comply with regulations.

## PaaS and Serverless

Unfortunately, data protection for PaaS and serverless is much less mature than for the container landscape, and suffers from the same fragmentation issues as other cloud-native environments.

If you venture out into PaaS and serverless territories, you'll quickly see you're at the mercy of the vendor much more than with other platforms. That's especially the case in PaaS environments, where platform specifics make it harder for third parties to integrate, and you'll have to depend on the PaaS platform to provide you with the right data protection.



**With PaaS and serverless, your nonfunctional requirements should include due diligence on data protection.** You'll likely end up with multiple data protection solutions to safeguard the entirety of an application.

The lack of choice in data protection, security, and other non-functional requirements means PaaS is definitely less flexible than other approaches.

And as with cloud-native environments, be prepared to look at container images in container image repositories, code in the version-controlled code repository, stateful data in the object stores, and data in third-party (SaaS) services through their respective APIs. In order to protect the entirety of an application, you'll need to protect data in all of these, using more than one backup product.

You'll undoubtedly lose some overview and insights because of this; work on compliance and organizational ownership to battle the loss of insights across the patchwork. Decoupling the consumption of backup policies, putting the right policies into the hands of those who work with the containers, databases, and file and object services helps increase backup coverage across all of these.

PaaS and serverless are standardizing on Kubernetes as the underlying control and data plane. That means container images and the Kubernetes APIs are the building blocks for data protection. This is an area that both new and incumbent vendors are innovating in, integrating native support into their existing data protection products or building new solutions without being held back by old practices.

# Edge

---

Edge is a special use case in terms of data protection. For starters, it's not really one specific environment, nor does it have a very clear definition.

Generally speaking, edge means running applications, storing content, and performing data processing physically closer to the user for performance or cost reasons. The physical proximity makes delivering content (such as websites or streaming video) quicker and less expensive, and cheaper to analyze and process user-generated data before sending the much smaller resultant data to the cloud.

Once reserved for content delivery networks, edge computing has taken off in (mobile) gaming, AR/VR, CCTV video processing, industrial IoT, and more.

Until recently, though, most edge locations were pretty much black-box services with no access to or control over functions like data protection. You'd upload your application binary (or code) and the edge provider would run it.



**As with PaaS and serverless**, the edge is becoming yet another place to run your container workloads.

Modern edge computing opens up that black box by leveraging it via Kubernetes. As noted earlier, the standardization on Kubernetes vastly improves data protection capabilities.

---

**Remember that for container applications, the build pipeline that creates your application container images and artifacts is as important as the containers themselves.**

---

The standard disclaimer applies here, as well: The running container is only a small part of the overall data protection puzzle, and you need to take the stateful data storage on file, block, or object into account, as well as cover cloud-based databases, third-party services, and any other component that makes up your edge application. Remember that for container applications, the build pipeline that creates your application container images and artifacts is as important as the containers themselves.

## On-Premises

---

On-premises virtualized environments are still the primary compute platform for many IT departments, and they're using data protection solutions tailored to that situation. Those solutions are mature, work well, and are priced right.

So, don't change a winning team, right? Luckily, these solutions aren't standing still. In fact, they open up a way forward for those considering moving parts of their application landscape to the cloud.

By looking at the support these backup products have for cloud, containers, and SaaS, it becomes easier to execute cloud migration projects, as it doesn't require you to rip out the backup for that component and replace it with something new.

Backup vendors have been adding support for containers, public cloud, and SaaS for a while, and many of these features are stable and safe to use. That makes it the perfect time to take steps toward containers and the cloud, without compromising compliance and business continuity.

## End-User Devices



End-user devices are the final piece of our data protection puzzle. The global pandemic revolutionized work-from-home, and employees' devices rarely, if ever, make it back to the office. That means security and data protection are no longer confined by, or defined by, the physical office walls.

Instead, devices and data need to be protected regardless of physical location. Similarly, onboarding new employees and issuing new hardware is now location-independent. IT often ships hardware directly from the manufacturer or retailer

to the employee and the onboarding is a self-service matter with simple, easy-to-follow procedures.

This requires IT to support remote employees and devices by default, using the public Internet for onboarding. More often than not, the changes required to support employees remotely are substantial, including the choice of endpoint management systems, different approaches to security and networking, and vastly different end-user support.

---

## **Security and data protection are no longer confined by, or defined by, the physical office walls.**

---

In other words, these changes in remote working force IT to adopt more modern workplace support systems, including data protection, to be able to “get out of the way” and let employees do their jobs.

With a large share of employees being fully remote, on-premises starts to lose meaning, too. Whether their application runs on-premises, in the cloud, or somewhere else, it’s all remote from the employee’s perspective.

# Things Will Never Be the Same



Wow, what a ride! Navigating data protection and disaster recovery in the modern era is a complex undertaking, but we hope this Gorilla Guide established a foundation of knowledge.

The goal of this Guide is to empower you to control and direct the quickly shifting trends in data protection, as SaaS and Kubernetes become ever-larger pieces of the IT landscape.

One thing to be sure of is this: The data center has changed forever. If you're still doing things the old-fashioned way, you need to start looking at ways to change *now*. Otherwise, you risk getting left behind as your competition surges forward.

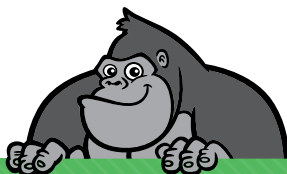
# About ActualTech Media



ActualTech Media is a B2B tech marketing company that connects enterprise IT vendors with IT buyers through innovative lead generation programs and compelling custom content services.

ActualTech Media's team speaks to the enterprise IT audience because we've been the enterprise IT audience.

Our leadership team is stacked with former CIOs, IT managers, architects, subject matter experts and marketing professionals that help our clients spend less time explaining what their technology does and more time creating strategies that drive results.



If you're an IT marketer and you'd like your own custom Gorilla Guide® title for your company, please visit <https://www.gorilla.guide/custom-solutions/>