

**THE
GORILLA
GUIDE TO...**®
EXPRESS EDITION



IT Resilience with Microsoft Azure

Trevor Pott

INSIDE THE GUIDE:

- Discover that IT resilience is more than just backup and disaster recovery
- Learn how cloud computing creates new IT resilience challenges and opportunities
- Tips for leveraging Zerto + Microsoft Azure to increase resilience

**TAKE A QUICK WALK THROUGH
THE IT JUNGLE!**

Compliments of

Zerto

THE GORILLA GUIDE TO...

IT Resilience with Microsoft Azure

Express Edition

AUTHOR

Trevor Pott

Copyright © 2018 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

Okatie Village Ste 103-157

Bluffton, SC 29909

www.actualtechmedia.com

TABLE OF CONTENTS

Chapter 1: Of Backups, Resilience, and Scotch.....	4
Backups: the 'Gateway Drug' to Cloud.....	6
Chapter 2: IT Resilience.....	9
Practical Differences.....	11
The Microsoft Azure Public Cloud.....	13
Why IT Resilience Is Important.....	16
Chapter 3: Challenges of Protecting Your Organization.....	20
Azure Limits.....	21
Connectivity Is a Challenge.....	23
Hybrid by Design.....	26
Chapter 4: Three Use-Cases for Using Zerto's IT Resilience Platform with Azure.....	29
Backup and Disaster Recovery.....	29
Application Migration.....	32
Solve End-of-Life Challenges for Windows Server 2008 and SQL Server 2008	33
Chapter 5: Do What You Do Best.....	35
Refocus on Digital Transformation.....	36
Zerto Enables IT Resilience.....	37

CHAPTER 1

Of Backups, Resilience, and Scotch

If you've been around the block in IT, you've probably got a war story or three. In the early days of the cloud, I was tasked with setting up a backup-to-the-cloud solution. A friend of the CTO convinced the CEO that cloud was the bee's knees, and that all our troubles would disappear if we used it.

I dutifully followed the various guides, how-tos, and documentation for getting my on-premises infrastructure to back up to the cloud, and was immediately consumed by the thousand other burning fires cluttering my desk. At the time, disaster recovery wasn't considered a priority, so no automation was set up, nor was any disaster recovery testing performed.

Eventually, disaster struck. The river overflowed its banks, and the data center flooded. The CTO, having bought into every single cloud marketing fluff piece in the airplane magazines, was convinced that getting everything back up and running would be as simple as pushing a button.

It was not.

All the critical VMs had static IP addresses, and there was no console access at the time. Not that this would have mattered anyway, because the GUI was installed on all Windows servers, many of which were still Windows Server 2003 at the time.

The end result was me, at 3 a.m. on a Sunday, tromping through a flooded, pitch-black data center, rummaging around in servers for any parts I could salvage, hoping I didn't get electrocuted. With salvaged parts in hand, I installed hypervisors on retired servers and some newly upgraded desktops until I had enough compute capacity to keep the company operating. The upcoming Monday was the busiest day of the year.

Shortly after noon on Sunday I went down to the local university with a bottle of expensive scotch and a NAS. The scotch was to bribe the local sysadmin to let me download all my backups from the public cloud using their delicious fiber optic connection. The NAS was to cart the data from A to B.

I got everything working just in the nick of time, but that particular backup to the cloud strategy was not exactly what you'd call "resilient."

Backups: the ‘Gateway Drug’ to Cloud

So, what is “resilience” in the IT world? Is it having backups in the cloud? If that’s your strategy, then it may be time for an update to your mental storage blocks. Making a backup of your data is simple: make a copy of your data. If your data doesn’t change, the backup remains valid for as long as you can read the data.

Data protection is a broader categorization that includes disaster recovery concepts. When nerds start talking about data protection, we’re not just talking about making a copy of data, we’re talking about having a copy of the entire apparatus necessary to make use of that data.

For many organizations, virtual machines (VMs) are the core of their IT. For some time now, data protection has been a fairly straightforward affair: not only do you need to copy your data, but also the VMs that operate on that data. Wrap it all up and send it somewhere that, should things go sideways, the VM copies can be spun up, and set about operating on the data once more.

For many organizations, this is the sum total of their interaction with cloud computing. Public cloud providers typically start out as nothing more

The Code Spaces Incident

One of the biggest mistakes an organization using a public cloud provider can make is choosing to believe that, because their IT is outsourced to the public cloud provider, they don't need to worry about data protection. The canonical lesson in this being a dumb idea is known as the "Code Spaces Incident."



The short version of the Code Spaces incident goes like this: Code Spaces built all of its IT with a public cloud provider. Both the production IT setup and the backups were accessible via a single administrative username and password. The bad guys obtained the relevant credentials and deleted the entire company. This is a bad approach to public cloud computing; it also emphatically makes the point that while backup to a public cloud can be a starting point, it falls hopelessly short of being sufficient to protect a modern business.

Much of what we see in our day-to-day regarding backup and data protection is out of date. It's time to stop talking merely about backups and data protection. It's time to have a conversation about *IT resilience*.

than a destination for backups. Often there's some handwaving about disaster recovery and/or data protection, but if we're all being perfectly honest with one another, we know that most organizations using the public cloud as a data protection destination haven't properly tested that it can be used to recover data and applications.

While data protection is often an organization's first foray into public cloud computing – sometimes known as the “gateway drug to cloud” – it's not long before the organization becomes comfortable running production workloads in the public cloud. Cost- and resource-strapped organizations (what IT organization is *not* resource strapped?) can realize savings and efficiency in the cloud.

Today, some companies put all their organization's IT into a public cloud. This Gorilla Guide Express will help those moving some or all of their infrastructure into the public cloud understand how to avoid the common rookie mistake of thinking that outsourcing IT absolves the company of the need for data protection. For this particular guide, the focus will be on Microsoft Azure as the public cloud target.

CHAPTER 2

IT Resilience

In this book, we'll use the term *IT resilience* to refer to a framework that prevents disruption to the business, whether it's avoiding, mitigating and remediating failures of information technology systems, or increasing availability during planned outages and migrations.

From a technology standpoint, IT resilience is made easier using Continuous Data Protection (CDP) technologies, which ensure that the absolute minimum



Continuous Data Protection (CDP) is a protection mechanism that allows organizations to continuously capture and track data modifications, automatically saving every version of the data that the user creates locally or at a target repository. Without snapshots or agents, writes are saved to a journal file. By utilizing changed block tracking, CDP allows users or administrators the ability to restore data to any point in time with remarkable granularity.

Pets vs. Cattle

Pets vs. cattle is perhaps best explained by Randy Bias: “In the old way of doing things, we treat our servers like pets, for example Bob the mail server. If Bob goes down, it’s all hands on deck. The CEO can’t get his email and it’s the end of the world. In the new way, servers are numbered, like cattle in a herd. For example, www001 to www100. When one server goes down, it’s taken out back, shot, and replaced on the line.”



Pets are workloads where, for whatever reason, rebuilding that workload from scratch would involve a significant time investment. Workload rebuilds for pets typically include application installs, Operating System Environment (OSE) customizations, and application configurations.

The closer a workload is to truly stateless, the closer it is to being “cattle.” Stateless applications store their data and configuration separately from the OSE and the application itself. The goal is to be able to throw away the OSE and the application(s) installed on it, whenever necessary. These applications can then be instantiated or scaled up and down as desired, with the separately stored data and configuration then being attached, and the workload moved into production.

Much of modern “cloud native” application design relies heavily on the idea of stateless applications. This type of application makes moving workloads between

infrastructures — both on-premises and public cloud infrastructures — much easier, though not all applications can be made stateless.

¹ <http://cloudscaling.com/blog/cloud-computing/the-history-of-pets-vs-cattle/>

of data is lost during a data protection event. Data protection events can be anything from ransomware attacks to DevOps teams writing the wrong things into their scripts and putting their organizations at risk.

Unlike backup or traditional data protection technologies, IT resilience is not only an insurance policy against unplanned outages. It's about handling planned outages and failovers, dealing with application scaling issues, finally making cloud bursting seamless, and slowly changing an organization's entire approach to IT so that there are fewer "pets" workloads, and more "cattle."

Practical Differences

There are some practical differences in the implementation of IT resilience strategies on different infrastructures. The traditional on-premises data protection approach involves an organization having

more than one data center. The organization doesn't have to own both data centers; it's quite popular for the second site to be in a co-location facility run by a service provider.

On-premises IT resilience is, hypothetically, a straightforward affair: duplicate everything in your production site precisely. The utopian version of data protection might presume that your organization has both the money and the will to buy two of everything, and implement everything identically on both sites.

In reality, this does not happen. On-premises data protection in the real world is usually about finding a way to get workloads that operate on this refresh cycle's equipment to "mostly sort of work" on the last refresh cycle's equipment. Pulling this off usually means making choices about what won't be spun up in the event of failover, because the second site's equipment doesn't have quite the same capacity as the production site. So many tough choices need to be made about what applications are "business critical": Is HR? Billing software? Payment and tracking systems? ERP? Think about the reaction of the team that IT determined was not "critical" enough to protect effectively. All applications are actually critical today, where no downtime is acceptable to either internal or external customers.

Other companies experience the polar opposite. They overprovision production IT resources and leave workloads running all the time, because it's easier to run those workloads all the time than it is to orchestrate their usage only when they're needed.

The problem with overprovisioning as the default is that it's expensive. Initially, IT managers thought public cloud would be a good solution for overprovisioning. However, public cloud providers charge customers based on what they provision, and leaving workloads provisioned when they aren't being used can cost a lot of money in the blink of an eye.

The Microsoft Azure Public Cloud

As awareness of when and how to efficiently make use of public cloud grows, it seems that for the foreseeable future, enterprise IT infrastructures will be a blend of on-premises and cloud known as “hybrid cloud.” Microsoft CEO Satya Nadella preached hybrid cloud repeatedly during the company's 2018 second-quarter earnings call, where he said:

“For me, it all comes down to really having an architectural advantage on what is a new secular trend. So when we think about the intelligent cloud and the intelligent edge and then bring that to the Azure

business, you can see it at each layer. When it comes to infrastructure, we're the only cloud provider that provides true hybrid cloud computing with Azure and Azure Stack. When it comes to the data tier, we have real uniqueness."

Microsoft's Azure bet is that a hybrid approach is going to win the cloud war, and they're setting up their cloud business to enable enterprises to consume cloud services with "intelligence at each layer."

One differentiator is that on Microsoft Azure, customers have the option to choose to use managed disks. Managed disks are more expensive than more basic storage types, but they come with more guarantees regarding reliability. Microsoft now offers zone-redundant snapshots for managed disks. Azure also offers snapshotting for SMB file shares and databases, including files, SQL data warehouses, blobs, disks, and more.

Cloud-native management has also improved in the last few years. Administrators no longer need to focus on duplicating all the management, automation, and orchestration tools. Public cloud providers expect those using their clouds to use the tools native to that cloud for management, automation, and orchestration.

Microsoft takes care of ensuring Azure’s management tools exist in all its data centers. It makes sure user accounts, logins, identity management, automation scripts, orchestration runbooks, and so forth are also replicated throughout its cloud fabric. Azure customers only have to worry about making copies of their workloads and data; Microsoft will replicate the rest.

Similar to on-premises approaches to IT explained previously, Azure’s management tools and features work best if your company is comfortable putting all IT into a single cloud provider’s basket. However, hybrid cloud and multi-cloud approaches to IT mean that multiple infrastructures are involved, and this greatly complicates matters. See **Figure 1** for an example.

Microsoft has been working on making interoperability, connectivity, and data protection between on-




Figure 1: Many organizations are leveraging multiple infrastructures both on-premises and in the cloud, and infrastructure sprawl adds significant complexity.

premises Microsoft environments and Azure seamless since Windows Server 2012. While its integration with Microsoft products is well advanced over competitors, there's still significant room for improvement in its hybrid cloud approach; and even that assumes that organizations run all-Microsoft environments entirely by the book.

Every major survey about IT decision-making intentions for at least the past two years has shown that hybrid cloud and multi-cloud are winning the day, including RightScale in its 2018 “State of the Cloud” survey.¹

In the real world, today's organizations don't get to choose either a simple on-premises identical-second-site unicorn or a single-vendor, cloud-native EASY button. In the real world, systems administrators have workloads all over the place, constant demands for change, and a need to make sure everything works – no matter what happens on which infrastructure provider.

Why IT Resilience Is Important

Today, IT is everywhere, and one result is that traditionally analog departments within organizations of all sizes are undergoing digital transformation. Internet

¹ <https://www.rightscale.com/lp/state-of-the-cloud>

of Things (IoT) devices are proliferating. Everyone has a smartphone. Seemingly everything relies on IT in some form or another, and for a lot more than filling out reports in Excel or answering the odd email.

One of the worst-kept secrets of IT is that in most organizations there's already more IT in use than administrators can properly keep running. Some organizations have all their various IT bits automated, orchestrated, and composable. Most don't.

Despite this, organizations are not just reliant on their data; they're reliant on IT automation. Most companies, for example, have a middleware solution of some sort. This middleware solution may be built on an industry-standard middleware package, but just as often it's a collection of scripts, runbooks, websites and databases cobbled together by in-house developers, contractors, and systems administrators over the course of several years.

Middleware packages are responsible for taking data from one application and feeding that data into another application. Data is often warehoused in the middleware application, with analytics being run on that data. The middleware application may even have a customer-facing component, especially for organiza-

tions whose customers are other organizations, rather than consumers.

In an example from my own career as a systems administrator, a middleware application I wrote was used to crack open .zip files submitted by customers using an ordering application. Metadata was extracted from an .xml file, parsed, cleaned, and injected into a database. Scripts would then run to inject that data into the invoicing and point-of-sale application, and to print order information onto a physical piece of paper.

The physical piece of paper followed the order through manufacturing, and ultimately ended up with the finished and packaged product in shipping. Shippers would use a barcode on the physical piece of paper to call up the invoicing information, print an invoice, and check the customer's shipping preferences. At this point, the shipper would see the options the customer had chosen for shipping, push a button to select the relevant courier, and the middleware would create a waybill with the relevant courier using the invoice data.

When this middleware was put into place, shipping errors dropped by more than 99 percent. The company was able to quintuple its volume using the same number of staff, and still cut its logistics costs in half, almost entirely due to error reduction. If any part of

the multitude of interconnected systems went down, however, the company was in a lot of trouble: There simply was not enough staff to handle the volume of orders going through any part of the manufacturing, sales, or shipping departments without the computers.

If any of the moving parts of middleware were lost, it could take months to rebuild, and the impact in terms of lost productivity, having to hire temporary workers to fill in, and increased error rate due to humans being more fallible than computers are all unknowns. Existing IT staff could barely cope with the constant changes being requested. Rebuilding all the various layers of automation would be next to impossible.

This is the discussion that needs to be had about IT resilience. Resilient IT is, at its core, about automation and orchestration of the movement and protection of workloads, so that there is no one script or workflow process that someone needs to code, manage, update or maintain. Unlike the example above, IT resilience is a framework for avoiding, mitigating, and remediating failures of information technology and systems, and increasing availability during planned outages and migrations. A resilient data center enables availability, regardless of the interconnections and relationships of IT components across multiple different infrastructures – including clouds.

CHAPTER 3

Challenges of Protecting Your Organization

We've discussed the some of the overall challenges of building resilient IT, including these realities:

- An increasing number of organizations use a combination of on-premises and public-cloud workloads, adding infrastructure complexity
- IT administrators often build custom one-offs for automation, configurations, and manifests that are baked into the management tools for the platforms on which those workloads execute. This poses a challenge for teams that must manage, update and maintain all of that custom configuration for running workloads.

In addition to these, there are protection challenges specific to Azure, including infrastructure limits, connectivity, and hybrid cloud complexity. This chapter addresses those issues and provides some tactics for avoiding pitfalls.

Azure Limits

While most people think of Azure as a monolithic entity, it is in fact a collection of multiple services, each created and maintained by different teams. There is no one-shot “back up and recover your Azure” configuration.

For those seeking to protect their Azure runbooks, for example, there are PowerShell scripts. Want to back up the API management service? Issue a POST command. The other Azure services all must be examined one at a time. And, like all public clouds, it's important to be aware of what you can and cannot do on the platform. **Table 1** shows some of Azure's basic parameters.

Infrastructure management tools contain a lot of data that's often taken for granted. The cold start-up sequence of various workloads is one example; it isn't abnormal to have some workloads that must be fully functional before others can be started. DNS servers and Microsoft Active Directory (AD) are examples of workloads that typically must be up before anything else. Similarly, many web services need to have the database up before the rest of the application can function.

Many applications consist of multiple workloads. Like any application, updates occur. One common data

Azure Limits

Virtual Machines	Default	Max
VM's per subscription, per region	10000	10000
VM total cores per subscription	20	varies
VM per series (Dv2, F, etc) cores per subscription	20	varies

Storage	Default	Max
Storage Accounts per region per subscription	200	250
Storage account capacity	500TB	varies
Maximum request rate (IOPs) per storage account	-	20k
Resources per deployment	800	800

Networking (ARM)	Default	Max
VNets	50	500
VNet peerings per VNet	10	50
Public IP address (dynamic)	5	varies

Table 1: Make sure Azure is sized properly for your environment before taking the plunge.

protection stumbling block is backups taken during an update sweep. If workload A is captured in an updated state, but workload B is captured in a pre-updated state, the application as a whole may not function if brought up at the disaster recovery location.

Simple data protection tools don't have the sophistication to deal properly with applications that consist of multiple workloads. They can't take into account problems like update skew, initiation order, or perform complex multi-workload testing to ensure that disaster recovery requirements will be met.

This awareness of workload orchestration complexity is part of what sets IT resilience apart from simple data protection. IT resilience solutions are themselves capable of workload automation and orchestration, making them a powerful tool for managing groups of interconnected workloads that span — and move between — multiple infrastructures.

Connectivity Is a Challenge

Once administrators have sorted out what needs to be backed up to where, and how it all needs to be automated on all sites, there's the small matter of shipping the bits from A to B. When shifting bits from one Azure region to the next, this is straightforward.

The overwhelming majority of organizations can treat Azure as having functionally unlimited connectivity between regions: It's highly unlikely that anyone reading this is going to be able to saturate the network links between Microsoft data centers.

Shipping bits from an on-premises data center to an offsite location, however, is more challenging. Most organizations have less internet throughput available than they'd like, and some may have regulatory compliance concerns about shipping certain types of data over the public internet, even if that data is encrypted.

One option is to seek dedicated connectivity to one's public cloud provider, such as Azure ExpressRoute (see **Figure 2**). ExpressRoute connections are dedicated, do not traverse the public internet, and can link an

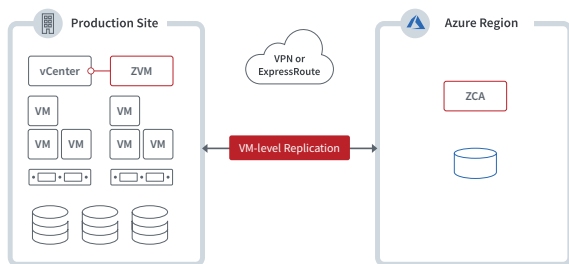


Figure 2: Using Azure ExpressRoute to link on-premises and cloud infrastructure.

organization's on-premises data center to the Azure cloud. This allows organizations to build truly hybrid applications, as well as to build resilient, multi-infrastructure IT that can survive outages both on-premises and in the public cloud.

Once the bits are shipped from A to B, workloads must function in their new environment. Networking can be a challenge. VMs taken from an on-premises data center may need some care and attention before they're ready for life in Azure.

On-premises admins are used to being able to open a console session to a VM that's having issues, and making changes. If a VM has a static IP address, for example, and is moved to a site with a different subnet, administrators of on-premises virtualization offerings have no problem changing the IP address. This is not necessarily true when talking about public clouds.

Azure recently added serial console access² to its list of features, but at the moment it's only valid for text-based interfaces. If your VM management approach relies on having access to a VM's console session GUI, you may be out of luck.

² <https://azure.microsoft.com/en-ca/blog/virtual-machine-serial-console-access/>

Azure normally relies on VMs having dynamic IPs, and uses guest agents to provision and manage VMs³. It's useful to install guest agents into VMs *before* sending them to Azure, as they enable much of the Azure guest management and orchestration functionality.

It's also important to consider the impact of moving VMs to Azure, or indeed to any new infrastructure. Windows VMs, for example, may need to be sysprepped, as the VM may undergo conversion during import, or may be the result of cloning.

Running into sysprep issues is common with classic data protection solutions, as they clone images for disaster recovery testing, but don't set up isolated networks. This results in VMs with identical names and IDs existing on the same network, which can cause all sorts of odd behavior. Modern IT resilience solutions are cloud-aware and avoid these common pitfalls.

Hybrid by Design

Despite the difficulties associated with trying to marry on-premises and public cloud infrastructures, Azure is a great choice to form part of an organization's IT resilience strategy because it's designed to integrate

³ <https://docs.microsoft.com/en-us/azure/virtual-machines/extensions/agent-windows>

easily with Microsoft products. Those products, like AD, form a core part of most IT deployments, and Microsoft has made great strides in integrating the on-premises and Azure-based offerings for many of its services and applications.

The availability of cloud-native versions of applications and services is another consideration when crafting an IT resilience strategy. It may not be necessary — or even advisable — to use Azure for data protection of certain applications. Where native versions of an application or service exist, it's usually possible to send

IT Resilience		
UNPLANNED	+	PLANNED
User Errors		Mergers & Aquisitions
Infrastructure Failures		Move to Cloud
Security & Ransomware		Datacenter Consolidation
Natural Disasters		Maintenance & Upgrades

Table 2: IT resilience covers more than just unplanned outages; in fact, more and more it's the planned outages that must be dealt with by companies.

only the data or configuration from the on-premises instance to a cloud-native one.

In each of the scenarios above – workload orchestration complexity, connectivity issues of solutions that aren't cloud-aware, and lack of native hybrid-cloud support – IT resilience provides several benefits beyond data protection. Making use of native hybrid cloud capabilities can make applications more flexible and responsive to change. IT resilience solutions provide a platform for managing, automating, and orchestrating applications across multiple infrastructures, and they protect against both unplanned and *planned* outages (examples of both kinds are shown in **Table 2**).

CHAPTER 4

Three Use-Cases for Using Zerto's IT Resilience Platform with Azure

Backup and Disaster Recovery

In the modern age, *any* loss of data is generally deemed unacceptable by the business. IT organizations strive to provide the lowest recovery times their budgets will allow. But there are many approaches to disaster recovery; what makes Zerto more efficient than other methods? Zerto is unique in leveraging Continuous Data Protection (CDP). Zerto CDP is hypervisor-based and doesn't utilize snapshots indefinitely to protect the VM; because of that, it suffers none of the usual performance penalty associated with snapshot-based data protection.

Zerto CDP enables you to deliver RPOs of seconds by replicating every change being generated real-time in concert with a journaling technology that keeps a log of all changes occurring during a customizable period — configurable down to the second. This technology significantly reduces the period of potential data loss

(the RPO), and thus the potential financial impact or reputational cost of failure.

Zerto's IT resilience capabilities include automatic and test-run failover to Azure, as well as automatic reverse protection of Azure workloads to on-premises data centers. Zerto can burst and expand compatible workloads to Azure as required, then either de-



PLANNED VS. UNPLANNED OUTAGES: ONE IS EXPECTED, BUT BOTH ARE EXPENSIVE

IT resilience takes the traditional notion of data protection—which is reactive—and adds to it the objective of accounting for planned outages, which is proactive. Even anticipated unavailability can cause a problem for the business; disruptive upgrades, workload relocation, and cloud migrations are all legitimate reasons for downtime, yet there is still a cost to the company for the outage. It would be better—and perhaps even mandatory in the future—for potentially disruptive activities like upgrades and migrations to be done without impact.

instantiate bursted instances, or migrate instances back on-premises when the time is right.

Additional advantages of using Zerto for disaster recovery in Azure:

- Reduces physical footprint, making it cost effective
- Compute capacity is only needed in case of DR event and DR test
- RPO in seconds, RTO in minutes
- Full application stack consistency
- Any application can be easily protected (SAP, Windows, Linux, Oracle)

Microsoft Azure Data Box, a gateway device that allows on-premises workload backups to easily consume Azure Storage and Zerto's IT Resilience Platform, delivers a direct route to utilize cost-efficient cloud storage to provide short-term and long-term retention of data and applications. Microsoft Azure's endless capacity, coupled with Zerto's incremental long-term retention capability, enables users to meet compliancy requirements for longer term retention without the traditional backup production impact.

Application Migration

This guide has primarily been focused on scenarios where it makes sense to move a company's DR site to Azure, but there are at least two other instances where Zerto helps companies automate and orchestrate the movement of applications for other purposes: application migrations, and solving challenges relating to end-of-life for Windows Server 2008 and SQL Server 2008.

Cloud vendors provide migration tools to aid companies migrating to cloud, as it's in their best interest to make it as easy as possible to move data to their platform. However, when dealing with hybrid and multiple clouds with diverse data sets, it's highly desirable to have a tool that unifies the experience, and that's where Zerto comes in.

Because migrating data takes time, and migrating datasets which have a high rate of change is even harder, Zerto is uniquely positioned to help you through this process. The near-zero RPO means that, given the proper bandwidth, even datasets with a huge daily change rate can have replicas synced to within minutes of the primary workload. When it comes time to swing over to the destination, you instruct Zerto to fail over with grace, and watch without fear as it

happens. Within moments, your high-change-rate workload is running in Azure, and you're headed home for the day with a pat on the back from your team.

Advantages of using Zerto for migrations to Azure include:

- Applications can be Windows or Linux based
- Supports SAP: M-series virtual machines are SAP HANA-certified with RAM sizes up to 4 TB
- Remove risk with the ability to test easily before commit
- Reduce physical footprint while increasing ability to scale

Solve End-of-Life Challenges for Windows Server 2008 and SQL Server 2008

End of support is quickly approaching for these Microsoft products:

- Extended Support for SQL Server 2008 and 2008 R2 will end on July 9, 2019
- Extended Support for Windows Server 2008 and 2008 R2 will end on Jan. 14, 2020

According to Microsoft, end of support means the end of regular security updates. They highly recommend upgrading to the most current versions for better performance, efficiency, and regular security updates. They are also providing unique incentives to take advantage of Azure to solve end-of-life issues. Zerto can make it easy to take advantage of these incentives and move to Azure.

CHAPTER 5

Do What You Do Best

While the line between data protection and IT resilience is somewhat subjective, there's a fairly simple way to tell which side of the divide your current IT strategy is closer to: Ask yourself whether the solution in use provides valuable functionality that goes beyond just duplicating your organization's production environment somewhere else. Data protection is the side of the scale that simply makes sure there's a copy of A residing in B, and that if A dies, B will be able to take over.

IT resilience, on the other hand, uses a variety of different tools to make workloads more flexible, creating an agile IT architecture that spans multiple infrastructures. While data and configurations are replicated between at least two different infrastructures, workloads can ideally execute on any of the infrastructures under management, depending on which infrastructure makes sense at the time. IT resilience allows you to get out of the business of running and maintaining disaster recovery data centers and lets you focus on what you do best.

Refocus on Digital Transformation

IT resilience doesn't spontaneously emerge from the ether. In fact, it's reasonable to say that, like any design philosophy, true IT resilience is more of an ongoing aspiration than an attainable end goal. There will always be new IT applications and services to consume, and they will all take time to integrate into the mix.

That said, quality IT resilience software can give administrators a place to start. Vendors like Zerto with a diverse customer base receive a diversity of feature requests and design their solutions to handle an ever-increasing number of unusual and edge-case scenarios. The more support that exists out of the box, the less systems administrators have to do.



Figure 3: The three pillars of Zerto's IT Resilience Platform

With the free time your sysadmins have, they can do fun new things like automate the coffee pot. Or even better: build neat, customer-facing software and IT services that increase value for customers and perhaps even generate new revenue streams.

Zerto Enables IT Resilience

Zerto provides an IT Resilience Platform™ that goes beyond data protection. Zerto aims to help organizations realize value from multiple infrastructures by providing a single platform for workload automation and orchestration across all infrastructures that an organization leverages. Zerto is focused on simplicity and ease of use, and nowhere is this more obvious than when using Zerto with Azure.

Zerto can be installed onto on-premises infrastructure in minutes. After the initial replication of a workload has been completed, Zerto's ongoing CDP workload protection has negligible production impact.

Zerto can be easily deployed in Azure directly from the Microsoft Azure marketplace, and can store both replica and journal data as cost-effective Blob storage. Zerto can create recovery VMs when needed, with pre-configured sizing and network settings.

Zerto's IT resilience philosophy is based on three pillars: continuous availability, workload mobility, and multi-cloud agility. IT resilience matters to organizations of all sizes, because agility and responsiveness to change have become competitive advantages. If you're ready to make your IT resilient, let the Zerto IT Resilience Platform take care of the hard stuff, so you can get back to more important things.