

PURE STORAGE PRESENTS

THE GORILLA GUIDE TO...[®]



Modern Data Protection

Ed Tittel

INSIDE THE GUIDE:

- How flash devices change the backup and recovery game
- The critical importance of cloud-based platforms and technologies for modern IT, including data protection and recovery
- Key industry trends driving modernization and adoption of flash-based storage technologies

**HELPING YOU NAVIGATE
THE TECHNOLOGY JUNGLE!**

 ActualTech Media
www.actualtechmedia.com

In Partnership With

 **PURESTORAGE[®]**

THE GORILLA GUIDE TO...®

Modern Data Protection

By Ed Tittel

Copyright © 2021 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

WITH SPECIAL CONTRIBUTIONS FROM PURE STORAGE

Esther Balestrieri, Director, Integrated Marketing

Roger Boss, Senior Solutions Marketing Manager, Data Protection

David Huskisson, Senior Solutions Manager, Portfolio Solutions

Biswajit Mishra, Director, FlashBlade Product Marketing

Marc Mombourquette, Senior Product Marketing Manager

ABOUT THE AUTHOR

Ed Tittel is a 30-plus year veteran of the IT industry who writes regularly about cloud computing, networking, security, and Windows topics. Perhaps best known as the creator of the *Exam Cram* series of certification prep books in the late 1990s, Ed writes and blogs regularly for GoCertify.com, TechTarget, ComputerWorld, and other sites. For more information about Ed, including a resume and list of publications, please visit EdTittel.com.

ENTERING THE JUNGLE

Introduction: The Importance of Business Resilience	8
Chapter 1: Setting the Stage for Protection and Recovery	10
Legacy Systems Establish a Base.....	12
The Cloud Changes Everything.....	14
Flash-Based Technologies Rule.....	16
Chapter 2: Beyond Backup and Recovery	19
The Old Backup and Recovery Mindset.....	20
It's Not About Speeds and Feeds.....	21
It's a Cloud World Now.....	22
Ransomware on the Rise.....	22
Today's Recovery SLAs Are More Stringent.....	23
Chapter 3: Key Industry Trends	24
Data Protection.....	24
Hybrid Data Lives Everywhere.....	26
Disaster Recovery.....	28
Storage in the Cloud.....	31
Pitfalls in Making Cloud Storage Work.....	33
Beware of Ransomware.....	33
Automation Is Key.....	34
Recovery Time and Recovery Point Objectives.....	35
Managing Changing Priorities.....	36
Protecting a Work-from-Anywhere Workforce.....	37

Chapter 4: Pure Modern Data Protection	39
Understanding Business-Centric IT.....	41
The Importance of Availability.....	42
Unified Fast File and Object.....	43
The Value of Rapid Recovery.....	45
Where SafeMode Snapshots Come into Play.....	46
Becoming Cloud-Ready.....	47
Maximizing Data Agility.....	47
Chapter 5: Real Customers, Real Cases	50
City of New Orleans.....	50
Domino’s Pizza.....	52
Make the Most of Flash Storage.....	53


CALLOUTS USED IN THIS BOOK



The Gorilla is the professorial sort that enjoys helping people learn. In the School House callout, you'll gain insight into topics that may be outside the main subject but are still important.



This is a special place where you can learn a bit more about ancillary topics presented in the book.



When we have a great thought, we express them through a series of grunts in the Bright Idea section.



Takes you into the deep, dark depths of a particular topic.



Discusses items of strategic interest to business leaders.

ICONS USED IN THIS BOOK



DEFINITION

Defines a word, phrase, or concept.



KNOWLEDGE CHECK

Tests your knowledge of what you've read.



PAY ATTENTION

We want to make sure you see this!



GPS

We'll help you navigate your knowledge to the right place.



WATCH OUT!

Make sure you read this so you don't make a critical error!



TIP

A helpful piece of advice based on what you've read.

INTRODUCTION

The Importance of Business Resilience

Welcome to The Gorilla Guide To...[®] Modern Data Protection.

In an age where people go online for everything from dog toys to paying utility bills, continuity of online presence and uninterrupted access to an organization's data are of vital importance. So is protecting that data from unwanted and unauthorized access and alteration.

This guide aims to provide executives, business and organization stakeholders, and IT managers and technical leads with the information they need to make sense of modern data protection. It also explores the key roles that storage solutions, software, and programmable APIs play in providing cloud-friendly, flexible, and adaptable solutions.

Because users are more inclined to conduct business in a variety of digital forms than ever before, the data they wish to access jumps to the forefront of what lets organizations and businesses function. That's because data is what drives applications and services. Likewise, data is where insights and value ultimately originate, and data is what organizations must manage, control, and protect to meet legal, regulatory, and compliance requirements.

When it comes to business continuity, the name of the game is to get back online as quickly as possible. That's because time is money when it comes to downtime. Use of proper, modern tools and technologies also means that automation serves to speed recovery times. In fact, automation also supports regular practice and testing for business continuity to ensure that recovery time objective (RTO) and recovery point objective (RPO) requirements are realistic and attainable. Legacy models for backup and recovery really can't cut it anymore, because they're not able to provide the flexibility and recovery speed (or failover to avoid recovery) that modern businesses need.

In Chapter 1 we examine the basics of backup and recovery so that readers can understand how these activities—and the platforms and solutions that support them—provide data protection and the ability to recover from interruption or disaster to get the organization back to “business as usual.”

CHAPTER 1

Setting the Stage for Protection and Recovery

In This Chapter:

- Legacy systems establish a base
- How the cloud changes everything
- Flash-based technologies rule

Although the concept of backup and recovery is neither novel nor new, organizational requirements for fast data restore in the wake of an unplanned interruption are more stringent than ever. As more organizations build their business and competitive edge atop a data foundation, rapid access to data for disaster recovery and data reuse becomes imperative.

In response, backup and recovery techniques, which have evolved and adapted many times since their introduction six decades ago, are poised to undergo a new and significant transformation, driven in particular by cloud and flash storage technologies.

But while data protection has always been vital to organizations, it doesn't always get enough attention. One reason is the IT technologies used just two or three years ago simply can't deliver the rapid recovery scenarios necessary in today's reality. Companies were disinclined to invest in the latest and greatest tape library or tools when they provided no real improvement to or benefits for recovery performance. Things have changed radically since then, and new technologies are forcing companies to rethink data protection priorities and investments.

The approach to backup and recovery has changed profoundly since 2015, as companies seek to do more with their backup data. It's not enough for data simply to get backed up—that same data can actually provide a competitive edge in the marketplace. The proliferation of technologies, in cloud compute and cloud storage, flash storage, and storage efficiency algorithms, has significantly altered the capabilities and intelligence that IT can provide to its internal customers by leveraging backup data.

Indeed, putting data to better use for internal customers is an enormous advantage. But another reason data protection enjoys elevated attention comes from the pressures that business leaders feel arising from growing risks of ransomware, breaches, hacks, and extended outages.

At the same time, virtualization and cloud technologies (especially the increasing adoption of and investment in hybrid multi-clouds) are changing IT's understanding of what data protection means and how it works. Assets may now be stored, backed up, and recovered across a broad range of locations. Key concerns that emerge from virtualization also carry over to data protection, backup, and recovery. These concerns include:

- Snapshotting, checkpointing, tracking, and synchronizing relevant applications and storage assets and resources
- Performance and resource consumption monitoring, to make sure that service-level agreements (SLAs) and other response time or user experience objectives are met, and cost controls maintained
- Enforcing security requirements for access controls, encryption for data in motion and at rest, and user and administrative privilege management, to guarantee the principle of least privilege and to audit access to sensitive data, use of administrative privileges, and more.



Organizations must understand that the best and final line of defense against ransomware is a known, good working backup (preferably one that's tamper-proof), that can be quickly and reliably restored to bring your business back online quickly.

And, finally, ransomware attacks have underscored the vast importance of access to tamper-proof, clean, and malware-free backups. Indeed, paying attackers does not guarantee that decryption will succeed, and explains why the FBI recommends against paying ransoms, period. Time considerations also argue against working through ransom payment processes, because mounting delays impose increasing opportunity costs and reputation damage.

Legacy Systems Establish a Base

Early on, the choice of magnetic tape as the exclusive medium for backup data arose from the high cost of disk media, as well as the intrinsic value of tape itself. That intrinsic value includes tape's portability (so that backup data could be transported to a safe offsite location) and the "air gap" it provides to backed-up data (tapes are offline, out of reach to hackers and malware). Recovering from tape has proved to be time consuming with shipping the tapes back to the site and loading each tape for restore; prone to failures as magnetic tape can be more fragile and break down over time; and may not meet the business SLAs on recovery times.

A disk-to-tape strategy for data protection persisted until the late 1990s, when distributed computing and lower-cost, higher-capacity disk drives appeared. In addition, distributed computing challenged tape-based backup in several ways. For one, sharing a tape system

required connectivity between servers, their storage, and the tape system, which ultimately became a key driver of storage area networks (SANs).

Adding to the complexity of physical connectivity was a need to schedule (1) use of the shared resource among multiple servers, and (2) processing workloads on servers and networks to accommodate backup processing and backup data traffic. As server farms moved to 24/7 processing schedules, time grew scarce for tape backups, and new solutions were sought.

Spinning disks have their limitations. They come with device identification issues, and repair and replacement challenges. Because hot-swap ability is a must, special-purpose enclosures and connections are required. When, as is often the case, RAID arrays are involved things get more complicated and time-consuming. Spinning disks are likewise prey to mechanical problems and failures, as even a cursory glance at the Backblaze Hard Drive Stats¹ will affirm. Trade-offs are inherent to RAID arrays, where increasing redundancy, performance, and availability all come with increased cost and complexity.

Because network-accessible hard disks seldom support retention locks or version histories, they often cannot act as “gold masters” for restoration. But if you look at recent snapshots as compared to previous ones, you can easily identify changes and revert to previous versions. In fact, sudden, significant changes in snapshot size, activity levels, and encryption behavior can signal a ransomware attack. Other technologies can offer improved protection, speed restore times, and work from immutable (tamper-proof) snapshots guaranteed to be clean and virus-free. Such technologies are covered in more detail in Chapter 4 of this Guide.

Legacy backup also typically occurs on a per-server or per-department basis. Alas, this creates collections of disparate, disjointed, and sometimes incompatible medias best understood as “backup silos.”

¹ <https://www.backblaze.com/b2/hard-drive-test-data.html>



A “gold master” is a recovery image or snapshot for a runtime environment that’s guaranteed to be clean and virus-free. It’s an essential point of departure for recovery efforts whenever malware attacks occur, but is especially useful in recovering from ransomware.

Gold masters are guaranteed to be free from malware, and should be used to recover from ransomware attacks. They will provide a clean replacement for infected or encrypted files, and undo ransomware’s impact.

Backup silos are suboptimal in today’s IT environments where customers are looking to improve efficiency and IT admin productivity.

The Cloud Changes Everything

In today’s modern world, the cloud has forever changed backup models. For one thing, how companies are using cloud services today differs from how they were used even just a few years ago. As organizations adopt multi-cloud environments, including both public and private cloud platforms, the number and kinds of uses keep proliferating. For example, the Flexera 2020 State of the Cloud² report says that enterprises use an average of 2.2 public and private clouds each. This includes both cloud platforms they use directly themselves, and cloud platforms built into Software-as-a-Service (SaaS) offerings such as Hybrid, Cloud Storage, Cloud Computing and Cloud Archive.

In some cases, cloud resources serve as adjuncts to on-premises infrastructure. In other cases, they provide specific services such as archival storage more cost-effectively than do-it-yourself alternatives. In yet other cases, companies instantiate entire workloads into the cloud,

² <https://info.flexera.com/CM-REPORT-State-of-the-Cloud>

and exploit connectivity between geographically dispersed cloud facilities to obtain failover and recovery services. There are a great many usage scenarios, and new ones pop up daily.

The latest analyses show some firms rehosting workloads originally placed in public clouds on cloud-based infrastructures created in private colocation or hosting facilities, or on-premises in their own data centers. This trend, primarily driven by cost considerations, security and compliance needs, and other factors, also determines where and how backup data is stored and used.

It could be appropriate, for example, for data to be mirrored between primary and secondary storage. In such a situation, both data and applications can simply fail over from one environment to another if a disruption occurs. An important consideration for such an approach, however, is to separate both storage environments (and their cloud hosts) geographically far enough to protect them from both going offline in a single disaster.

But with greater distance (more than about 50 miles) comes greater latency and “data deltas.” Those deltas represent differences between original data and a backup copy. Likewise, beyond that distance, synchronous replication gives way to asynchronous replication. In turn, this could lead to potential loss of data should a failover occur. The impact of data deltas must be considered whenever planning failover strategies. The strategy may call for logging tools and secure log transmission mechanisms to let replay techniques apply as much of the delta to the failover environment as possible after the fact. This is just one reason why ensuring robust replication is an important aspect of RPO planning and any business continuity or disaster recovery (BC/DR) implementation.

Without a doubt, the cloud’s allure is impossible to resist. It offers the unbeatable trinity of convenience (accessible from anywhere, anytime, to all parties willing to pay for access), scalability (highly virtualized cloud-based environments can scale up and down on demand), and

trading capital expenses against operational ones (the cloud puts the CAPEX burden on the provider, in exchange for pay-as-you-go charges for resource usage and consumption).

That said, the cloud's cost equations can be complex and imposing, particularly when it comes to backup and disaster recovery. For one thing, it's expensive to maintain disaster recovery instances for failover in the cloud. In general, it's expensive to host disaster recovery in the cloud while on-premises recovery is underway. And then, when that recovery is complete, it's expensive and slow to move back to on-premises from the cloud, because of egress charges for outbound data and limited bandwidth for moving it.

Flash-Based Technologies Rule

Flash-based storage uses non-volatile random access memory (NVRAM) chips to store information on silicon, rather than using magnetic remanence as with spinning hard disk platters or linear recording tape. Flash offers important advantages over hard disk or tape storage for backups. Those benefits include:

- **Performance characteristics:** Flash storage is significantly faster than hard disk storage, especially for random reads. Faster reads support faster restores of backup data whenever recovery is needed, either on a per-file basis, or for entire snapshots or images.
- **Physical attributes:** Flash storage is much more compact than hard disk drives, and thus can accommodate greater data density and capacity in the same physical space. When used along with a backup server or as a backup/restore appliance, flash storage delivers better economy and capacity for backup data writes.
- **Reliability:** Flash storage media is less prone to bit-level errors by at least a factor of 10, which makes it more reliable than hard disk media.

- **Solid state:** With neither motorized spindles or platters, and no read/write head assemblies to move around a spindle stack, the lack of moving parts in flash storage offers numerous benefits. Flash storage consumes less electrical power and generates less heat than hard disk arrays, delivering greater cost efficiencies.
- **Scalability:** Flash storage is inherently more scalable, because of its speed. Delays don't aggregate when multiple devices are used as with moving media.
- **Protocol usage:** Flash storage can employ networking protocols to extend its performance advantages to intersystem and intrasystem data copy, replication, and mirroring capabilities. It's faster, easier, and more efficient to use flash storage for real-time and near-real-time RPOs and RTOs, and for failover and failback scenarios.



Failback describes the process that should eventually follow failover. Failback occurs when operations resume on a primary machine or facility after they have been shifted to a secondary machine or facility because of a failover. During a site-wide failover, for example, all processing, network, and storage access gets shifted from a primary location to a disaster recovery (DR) location that may be on-premises at some other location, and/or into one or more cloud platforms. What failover does, failback returns to its original location (or its equivalent).

The superior speed and scalability of flash storage make it an increasingly popular and common choice for backup and recovery—especially for recovery or failover, where every second counts. It's also easy to mix and match flash-based devices, so organizations can avoid vendor lock-in and use the solutions that work best for their business needs. Although it remains more expensive than conventional hard

disks (although the differential between the two is closing quickly), flash storage is better suited for multi-purpose uses, and is no longer cost-prohibitive to deploy in volume. It's also the case that flash storage helps future-proof modern IT platforms. Incorporating such storage not only helps organizations meet today's increasingly aggressive demands, it also lets their IT platforms scale linearly and grow to meet future demands.

Figure 1 shows a simple speed comparison between a solid-state hybrid disk (SSHD) drive and a flash-based NVMe (TLC, QLC) drive. The flash device outpaces the hard disk by a factor of 25 to 965; most importantly, 4K random access is almost 57 times faster for reads, and 77 times faster for writes.

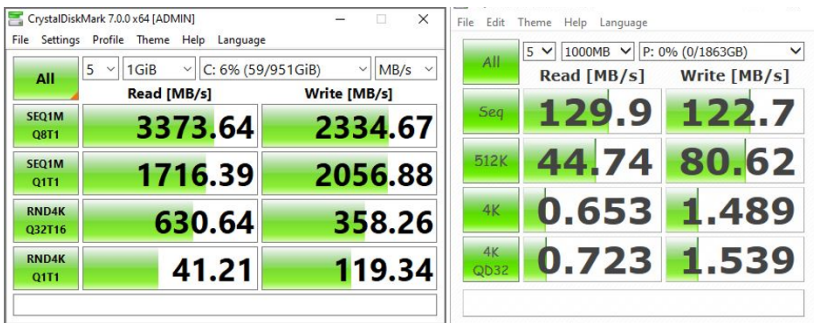


Figure 1: A flash device outperforms a spinning hard disk by 2,500% to 96,500%
 Source: [Output from CrystalDiskMark 8.0.2 Standard Edition](#)

CHAPTER 2

Beyond Backup and Recovery

In This Chapter:

- The old backup and recovery mindset
- It's not about speeds and feeds
- Ransomware is on the rise

There's more to bringing back business operations than the mere mechanics of backup and recovery. Organizations must always meet or exceed the requirements of the recovery SLAs to which they subscribe. In fact, there's a new, more modern approach emerging in the marketplace that goes beyond backup and recovery. Modern IT must be able to support a wide range of modern applications with diverse demands, including data protection applications, replication and snapshotting, transaction and activity logging, and more.

To those ends, performance becomes a key factor in numerous ways. Some hinge on reducing RTO and RPO intervals, others on improved throughput and productivity, and others on finding new innovative use for data already being acquired and backed up. Thus, a strong and wide-ranging data protection strategy supports the growing trend for organizations to become more data-driven, and to gather more insights and value from their data.

The Old Backup and Recovery Mindset

In the past, IT traditionally had been focused on RTOs and RPOs. For years RTOs/RPOs have been the standard by which IT success was measured. In many cases such measurements served to define SLAs between the business and IT. Business managers, on the other hand, focus on availability, business intelligence, expansion, and time to market.

Two types of events within an IT ecosystem can affect SLAs: business interruptions and declarations of disaster.

A business interruption is any event that might cause production or productivity to be hindered or come to a halt. Part of the definition is a “predetermined time,” or an agreed-upon time interval during which the business can tolerate downtime for some particular function or service. When that interval is exceeded—and applications and data remain unavailable to users and customers—the business interruption turns into a declaration of disaster.

A typical DR plan outlines the process or procedure for declaring a disaster. Once one is declared, you pull the cord and it’s all hands on deck to solve the problem. At that point, RPO becomes the focus, and the question becomes how long until “business as usual” can be resumed—however that’s defined. Addressing the needs of the business requires asking the questions most important to the business in their terms—not IT terms. This is what a modern enterprise needs, and what previous solutions did not provide.

The pressures that business exerts on IT pave the way toward a new framework to meet the business’s demands. It’s a modern approach, not only in the technology it uses, but in its practical emphasis on meeting the needs of the business, first and foremost.

Simply put, the chief reason data gets protected is to facilitate its recovery and restoration in case of emergency or disaster. In IT terms,

this is RTO. To the business, it may be a predetermined time, codified in an SLA. Even more narrowly, it might simply be called “availability.”

The point here is that the realities of a business interruption mean different things to different groups within an organization. That said, the outcome should always answer the same question: How rapidly can we return to business as usual? The cost of not doing business in the wake of a business interruption or a disaster can be staggering. A failure can even threaten the viability of the business itself. It’s this threat, and its accompanying urgency, that drives organizations to undertake data center modernization. Indeed, modernization inevitably begins with every organization’s crown jewels—namely, its data.

It’s Not About Speeds and Feeds

When you look at solutions to help modernize your company’s approach to data protection (you should also start replacing that phrase with “business protection”), it’s essential to consider desired outcomes primarily in terms of recovery or resumption of activities.

Companies naturally want the fastest solution on the market. A flash storage system can provide the necessary recovery speed, but the data protection software used is every bit as important as the hardware used for storage. Pure Storage offers the optimal solution for modern data protection. Think of Pure Storage as the high-octane fuel needed to squeeze maximum performance from a race car. Yes, of course there are alternative fuels available, some of them cheaper. But using lower-octane fuel slows down your car. The effect of slower storage on your data center is analogous, and leads to performance issues across the entire system.

As already mentioned, solid-state storage increases backup performance (a write operation). But it’s also significantly faster at delivering randomized read operations (such as for a restore) than hard disk or tape storage. Thus, the notion of “time to data” is a crucial component when it comes to achieving recovery and meeting RPOs.

To drive the point home: This isn't about speeds and feeds. Rather, it's about the business.

It's a Cloud World Now

As stated in the preceding chapter, there's no disputing that "the cloud changes everything." Modern backup, recovery, and data protection tools and platforms must be at home in this world. That means that capable and modern solutions must be able to back up to the cloud, recover from or to the cloud, and be ready to protect data wherever it lives on premises, at the edge, or in any and all of the clouds, private, public, and hybrid, that an organization uses. Because the cloud is everywhere and involved in at least some (or all) of an organization's applications, services, and data, that means backup, recovery, and data protection must also be able to accommodate, interact with, protect, and use the cloud efficiently and cost effectively.

Ransomware on the Rise

In early May 2021, the United States was rocked by a ransomware-induced fuel shortage along much of its eastern coast. Colonial Pipeline, a Georgia-based pipeline operator with lines through Texas, Louisiana, Mississippi, Tennessee, Georgia, the Carolinas, Virginia, Maryland, Pennsylvania, and New Jersey had to shut operations down for nearly a week because ransomware developed at an eastern European criminal gang named DarkSide was inflicted on Colonial Pipelines' billing systems and networks. The company ultimately paid a \$4.4M (75 BitCoin) ransom to help get its operations restarted on May 12, 2021, though its supplies of gasoline, diesel, and jet fuel did not return to normal for several days thereafter.

Ransomware is now regarded as not just a clear and present danger for businesses and organizations of all kinds and sizes, public and private. It's also a major source of expense and financial loss for companies and organizations that fall prey to such attacks. In November 2020,

[Cybersecurity Ventures](#) projected 2021 global losses from ransomware at over \$20B, “the fastest growing and one of the most damaging types of cybercrime.”

In fact, the only sure protection from ransomware is a known, good, tamper-proof backup that can restore systems to proper, uninfected, and unaffected operation quickly. This hinges on creating read-only backups that are carefully monitored, closely protected, and accessible only to highly trusted programs, users, and accounts. Paying the ransom is often insufficient to restore normal operations. Thus, Colonial Pipeline admitted to *The Wall Street Journal*³ that obtaining the decryption key did not provide everything it needed to restore its systems to working condition.

Today’s Recovery SLAs Are More Stringent

It’s also the case that many businesses and organizations bound to customers, clients, or end users by SLAs on recovery find themselves increasingly hemmed in by shorter recovery windows, higher uptime requirements, and, in general, less tolerance for service delays and interruptions. Ultimately, this means that the backup, recovery, and data protection solutions they deploy to meet such requirements must be faster, more resilient, and better able to withstand the rigors and demands of modern business. While flash memory-based solutions were already attractive, as discussed in the preceding chapter, the pressure that more stringent SLAs can bring to bear make them both irresistible and absolutely necessary.

³ <https://www.wsj.com/articles/colonial-pipeline-ceo-tells-why-he-paid-hackers-a-4-4-million-ransom-11621435636>

CHAPTER 3

Key Industry Trends

In This Chapter:

- Disaster protection and disaster recovery
- Hybrid data lives everywhere
- Protecting a work-from-anywhere workforce

The technology and industry landscape also defines how an organization should plan for and set recovery objectives. This landscape covers a variety of related tools and technologies, including data protection and disaster recovery. In the following sections, we'll review emerging considerations about which technology executives, stakeholders, and architects must be aware of in these areas.

Data Protection

Data protection technology is an area constantly changing, both on-premises (which includes edge locations at remote and branch offices, as well as core locations in data centers) and in the cloud. Numerous cloud service providers now offer Disaster Recovery as a Service (DRaaS) or Backup as a Service (BaaS). These services enable the replication of backup data between multiple cloud sites to provide four-nines or better availability for data and workloads.

Some organizations are drawn to the cloud to handle service and host backups for purely economic reasons. By leveraging cloud services, they can shed local backup infrastructure and software. This also

removes administration and supervision of backups from the duties of IT administrators who have better things to do with their time. In some cases, combined equipment and labor cost savings more than pay for outsourcing backup to the cloud. That said, the real value of cloud-enabled backup is its contribution to comprehensive recoverability for applications, especially cloud-native applications.

However, although data recovery is central to any successful restoration of operations, disaster recovery requires more than straight-up data recovery. Recovering from a facility disaster or an outage event with regional (or larger) impact requires rehosting the application as well. This means making applications available on-demand from different locations, and redirecting network traffic to access them directly and efficiently.

Two capabilities are key to cloud-based recovery scenarios. First, workloads must be re-instantiated rapidly on cloud-based hosts. Second, software-defined networks must be redirected to the new locations where those workloads run the necessary services and applications. If that happens, the entire business can “fail over” to the cloud in the event of a major disaster. If things work seamlessly, users may not even notice that any disruption has occurred.

Of course, failover strategies date back to the time of mainframe data centers. Their practices still define how on-premises failovers occur today. Failover has always required two separate (and sufficiently separated) data centers, each comparably equipped with processors and storage, and ongoing data mirroring between them.

The cost of maintaining duplicated data centers and keeping all hardware, software, and data synchronized was (and remains) huge. Thus, only those firms with the deepest pockets could afford such a strategy. As a less painful alternative, many firms simply made copies of their data—the only asset that’s absolutely irreplaceable—and hoped for the best when it came to replacing infrastructure. Failover required cobbling together network resources to put critical systems back into production. This is still true for many on-premises approaches today.

Synchronous and Asynchronous Replication with Pure Storage

Features like ActiveCluster and ActiveDR are included at no charge with the Purity operating environment.



Features like ActiveCluster and ActiveDR are free elements in the Purity operating system. Customers who could not typically afford this can explore deploying synchronous and asynchronous replication because added costs apply only to hardware.

This can make deployment costs easier to justify to executive management.

Cloud technology allows organizations to virtualize data center resources. Given an efficient backup and recovery program, data can reside in some cloud instance. From there, it can be used by locally re-hosted applications or accessed remotely via wide-area or metropolitan-area links from servers in business data centers or user facilities. The affordability of clouds has never been better.

Hybrid Data Lives Everywhere

The Hosting Tribunal regularly publishes cloud adoption and uptake statistics. In its January 19, 2021, blog post, “Cloud Adoption Statistics for 2021,”⁴ it reported that “Organizations leverage almost 5 different cloud platforms on average” including both private and public clouds. Then, too, organizations often seek to combine on-premises systems and assets with their cloud-based counterparts. The result is what’s called a *hybrid cloud* that brings all these systems together in

⁴ <https://hostingtribunal.com/blog/cloud-adoption-statistics/#gref>

as seamless and flexible an environment as an organization might be willing to design (see **Figure 2**). Tying all these environments together involves a great deal of effort and complexity. IT will often choose to rebuild or replace legacy tools and applications that cannot accommodate the APIs and connections that a hybrid cloud requires to operate. Where “remove and replace” with something (hybrid) cloud-native is not possible, an organization may have to accept diminished flexibility in positioning and running workloads.

In a hybrid cloud environment, managing data flows and workload placement becomes extremely important. Organizations may choose to run certain workloads on-premises and limit data outflows into the

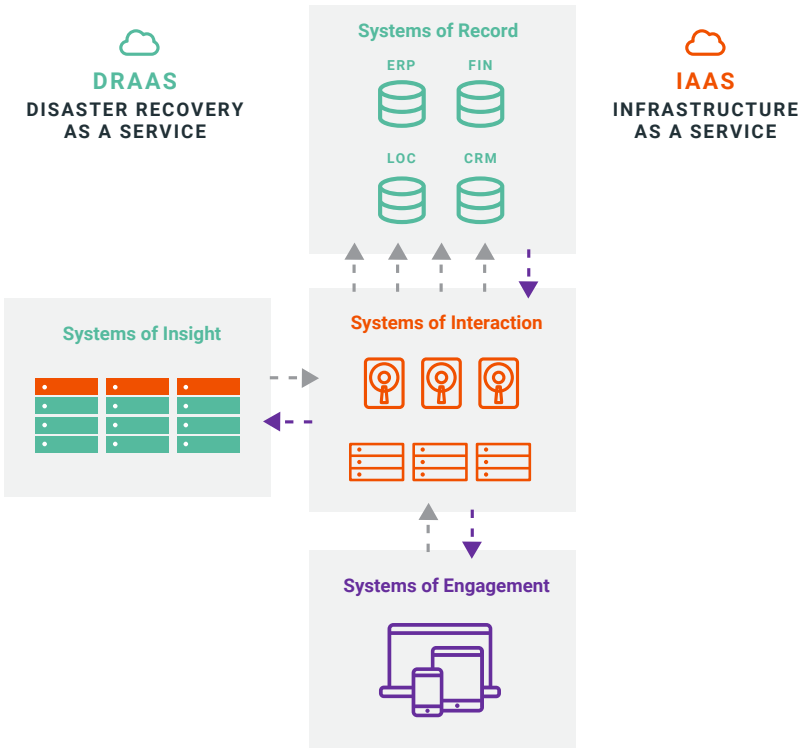


Figure 2: In a hybrid cloud model the corporate data center obtains select services from a cloud service provider (for example, Disaster Recovery as a Service or DRaaS) and additional compute, network, or storage resources on an as-needed basis from other clouds

cloud, because they wish to keep them confidential or to meet specific compliance and governance requirements. Other workloads may get positioned in the cloud, or placed at the edge to provide users with the best possible experience (low latency, high performance, better response times, and so forth).

In a hybrid cloud environment, organizations must juggle storage costs in a variety of forms. They must weigh the fully burdened costs of on-premises storage against the consumption and usage costs of cloud storage. The best economics are possible only when organizations can trade off unit costs for on-premises storage against comparable unit costs for cloud storage. Not surprisingly, in some cases this leads to “reverse migration” where data stored in the cloud moves back into a data center because it’s shown to be more cost efficient to house the data locally.

With the hybrid cloud in the picture, backup and recovery, along with BC/DR, become more fluid and complex. Organizations must understand the cost, security, and performance implications of creating and implementing hybrid cloud solutions. Only if solutions make sense should they be deployed. And if they do get deployed, they must be monitored and managed carefully to avoid unwanted or unexpected costs that sometimes loom large on the cloud side. Capacity planning and consumption planning turn out to be two sides of the same coin, where one side is in the data center and the other side is in the cloud.

Disaster Recovery

As with data protection, disaster recovery has on-premises and cloud-based options. On-premises, organizations must assess RTO and RPO intervals to choose among storage options that include tape, conventional hard disks, and flash elements. Tape makes sense only for long RTOs and RPOs, simply because of the time involved in reading from that linear medium (and, if necessary, transporting that data across a network with all the latency considerations thereby entailed). Hard

disk drives (HDDs) remain the most common choice in today's data centers, but can impose delays because of their data access speeds (recall the differences between HDDs and flash-based devices shown in **Figure 1** earlier in Chapter 1) and lack of scalability.

The advantages of architecting a modern approach to DR using flash are many. For one thing, all-flash performance can help firms struggling to meet data recovery standards as defined in IT SLAs. Flash storage can bring the speed of data and system recovery and restore up to the speed of data backup, which has been the focus of most improvements in data protection over the past 20 years.

By enabling multiple uses for backup data, flash extends its value beyond a traditional focus on risk reduction providing additional business value. If consumed properly, backup data confers both cost containment and improved productivity. Business management prefers IT initiatives or strategies that deliver value in multiple domains, so a modern data protection strategy based on flash storage is more likely to gain budget approval.

And, finally, all modern approaches must consider cloud-centered technologies. Solutions that make backup data portable to and from the cloud—as do flash storage technologies—offer two key benefits. First, they provide necessary separation between original and backup copies of data to ensure resilience in the face of facility and regional disaster events. Second, they provide a means to leverage cloud economics and flexibility needed to reduce cost and complexity in backup infrastructure.

By contrast, older and more traditional storage technologies—namely, tape and conventional hard disks—are typically too slow and cumbersome to meet today's RTO and RPO demands. Anything under four hours for either type of objective, in fact, is simply unattainable with either tape or spinning disk.

Likewise, the following issues all lean heavily in favor of flash-based storage:

- **Device size.** Because the volume provided inside data center racks determines how much storage capacity a rack can accommodate, the compactness of flash devices puts them ahead of both tape and disk drives.
- **Heat output.** Flash devices run much cooler than spinning drives.
- **Energy consumption.** Flash devices require less power than either tape or spinning drives.
- **Failure rates.** Flash devices fail less frequently (lower mean time before failure, or MTBF) and have far lower bit error rates than either tape or spinning drives.
- **Media management.** Flash devices can be accessed as soon as they're installed and connected. In contrast, tape devices must use libraries or human operators to swap media. Disk arrays impose performance penalties when drives fail, and require time-consuming operations to integrate replacement drives while rebuilding an array.

Thus, even though flash devices remain more expensive per storage unit, many organizations have found their other characteristics (especially support for shorter RPO and RTO intervals) compelling enough to make flash an essential data center storage technology. And with flash prices continuing to decline ([Gartner](#) Inc. says such costs have bottomed out, but that flash memory should remain at or near current levels for some time to come), the tilt toward flash storage continues to increase.

Storage in the Cloud

There are multiple storage scenarios to consider when the cloud comes into play. These may be characterized as follows:

- **Storage in the cloud.** This means consuming storage space in the cloud provider's environment, usually through some kind of vendor- and platform-neutral storage layer software and its available APIs. Costs are incurred for space consumed and for activity involved in accessing its contents.
- **Recovery in the cloud (failover).** This means using a cloud provider's facilities, including storage, to recreate a data center using cloud-based infrastructure. Storage is only a part of this kind of operation, albeit an essential one. Costs are incurred for instantiating entire IT infrastructures, including compute, memory, and network resources, as well as storage. Large recurring costs for storage consumption are typical for such scenarios.
- **Recovery from the cloud (failback).** When a data center comes back online, a cutover from the cloud environment used for failover is required. This takes time and involves substantial egress charges from the cloud provider, plus charges for resource usage and consumption while the failback process is underway. This activity can be costly, too.
- **Using Kubernetes-based containers in the cloud.** Because this approach best uses efficient, flash-based object storage, it typically consumes fewer resources (compute, networking, and memory) and less storage space (thanks to deduplication and compression) than file- or block-based alternatives. This translates further into lower overall costs and faster, more responsive containerized apps and services. Flash is the best technology currently available for cloud-native Kubernetes-based containers, apps, and services. As discussed in the section on hybrid clouds, calculating cloud costs can be messy and complex. Flash helps simplify those calculations and keep those costs down.”

There's considerable effort involved in translating between cloud-based storage costs and storage costs for data center-based systems and assets. The cloud's consumption-based, pay-as-you-go cost model depends on periodic billing based on storage space consumed, plus related usage costs for accessing and manipulating its contents and for sending and receiving related network traffic. Storage space costs are constant, while usage-based costs vary by activity and traffic levels.

The data center's fully burdened cost model is a whole different animal. It includes:

- Capital expenditures to acquire equipment, usually expressed through amortization and depreciation chargebacks to users. In addition to storage devices themselves, this equipment includes the racks in which they're mounted, the networking devices and cables that connect them to the network, servers and stations to configure and manage storage, heating and cooling systems, backup power generation facilities to keep data centers running when local power fails, and so on.
- Property costs to house the equipment, which may involve either capital expense (for company-owned property) or operational expense (for leased property).
- Overhead costs to cover planning, financing, purchasing and contract negotiation, and so forth.
- Lifecycle management costs to cover service and support contracts, and the necessary maintenance and upkeep for physical equipment and facilities.
- Operation costs, which include the power needed to run the equipment, the lights that are turned on, either more power or circulating water to keep the data center sufficiently cool to keep running, and so on. Also, usage and subscription costs for cloud-based services that the data center may consume—such as Storage as a Service—also count as operational expenses.

- Staffing costs, which include the fully burdened costs of employees who work in the data center, the IT staff that uses and manages local data stores and related systems, their managers, and so forth.

Putting all these costs together and working out unit costs for storage for both cloud and data center use takes time and effort. But it provides a helpful and rational basis for comparison, so that managers and stakeholders can define effective policies to guide storage placement and usage, and so that IT can monitor, manage, and tweak them as things change.

Pitfalls in Making Cloud Storage Work

Implementing data redundancy in the cloud can be a challenge, especially when it comes to balancing costs against capability, ensuring RTOs and RPOs can be met. Organizations will want to test and pilot prospective implementations before making any commitments, primarily because testing and practice are the only benchmarks that provide necessary insight into what happens when a real outage or disaster strikes.

It's also essential to establish the right priorities for recovery efforts. These priorities will drive the choice of a cloud positioning and recovery strategy that is sensible, affordable, and workable. Key items for consideration include line-of-business applications, email, CRM, e-commerce, and website presence. In short, anything that can materially affect business activity and capability had better be on that list, and its RTOs and RPOs must be tested against what on-premises versus in-the-cloud recovery can deliver.

Beware of Ransomware

According to [Cybersecurity Ventures](#), the estimated costs of ransomware attacks are growing furiously. From \$8B in 2018 to \$11.5B in 2019 and \$20B in 2020, that series indicates that something well north of

\$20B is likely for 2021. This omnipresent threat dictates that organizations do everything they can to protect and monitor their systems and data from ransomware attacks.



Be aware that ransomware attacks now increasingly seek to encrypt backup storage as well as primary storage holdings and filesystems. That dictates either an immutable, tamper-proof backup store, or physical isolation of backups from the network (and attack).

Automation Is Key

SLAs emphasize specific metrics that must be continuously monitored and compared to guaranteed levels or minima. Because time plays a vital role in certain SLA metrics, it is vital to respond to deviations below the minimum acceptable level in milliseconds. Thus, for example, consider uptime SLAs, which commonly fall in the four-nines to six-nines (99.99% to 99.9999%) range. (Uptime specifies the percentage of time a system or service stays up and running over some total time interval.) Five-nines (99.999%) uptime allows *five minutes* of downtime in any given month (except February, which gets only 4.5 minutes because it's so short).

Given requirements to meet such SLAs, speedy response is the name of that game. This makes automation essential, in order to read and respond to SLA metrics in tens to hundreds of milliseconds. Humans are lucky to respond to alarms or alerts in under a minute, which is unfortunately 600 to 6,000 times slower. Without a well-orchestrated and automated set of tools and responses, organizations cannot hope to honor SLAs. Well-tested automation also eliminates human errors or mistakes, which can slow responses even more.

RPO AND RTO




Figure 3: The differences between RTO and RPO

Recovery Time and Recovery Point Objectives

RTOs and RPOs assign hard and fast timing requirements for system availability. An RTO defines how long a computer, system, network, or application can be unavailable after a failure or disaster strikes (see **Figure 3**). Essentially, RTO states how long it should take to return something (computer, system, network, or application) to working condition when it fails or becomes unavailable. RTOs may be measured in seconds to days, but it's important to understand that choosing an RTO conveys how long the organization can go without the asset to which that objective applies before it causes intolerable or unaffordable damage.

RTOs should prioritize applications by their importance and potential losses that result from their absence. Resources must match the resulting intervals involved. RTOs under 10 minutes usually require failover services. Up to four hours, some kind of near-real-time failover is required. A four-hour RTO usually leaves enough time to perform bare-metal recovery and restore working applications and data access on-premises using a designated recovery team on staff. An RTO of eight or more hours usually permits IT to delegate recovery to a service provider.

RPO, on the other hand, refers to how much data can be lost before the organization experiences intolerable or unaffordable losses. An RPO is expressed as a time interval that stretches from the time of failure or disaster back in time to the most recent backup that precedes it. Thus, for example, an organization that backs up all or most of its data once a day (every 24 hours) should be prepared to sustain a loss of 24 hours' worth of data. For some applications this is tolerable; for others it could be catastrophic.

As with RTOs, shorter RPOs involve more expensive, time-sensitive technologies. RPOs under 10 minutes usually require real-time mirroring or replication solutions, so that another site maintains a live, available copy of all covered applications, data, and so forth, subject only to whatever data deltas (explained earlier in Chapter 1) apply to the link between the primary and secondary sites. Four-hour RPOs generally require snapshot replication (from primary to secondary), whereas eight-hour RPOs often work with existing backup solutions so long as they don't affect the performance of production systems significantly.

When RTO and RPO are both near zero, organizations must invest in technologies that combine continuous replication failover services. This gets your priority systems and services as close to 100% availability as the applicable SLAs from your service providers and your local infrastructure will allow. A five-nines SLA may sound like a good idea, but can come at a formidable cost. Only you and your organization can know if that cost is justified, so consider potential losses and impacts carefully.

Managing Changing Priorities

Establishing business continuity and supporting disaster recovery (BC/DR) is never a matter of "one and done." Regular practice is needed to make sure that your IT staff, stakeholders, and other players can execute their roles to meet applicable RTOs and RPOs. Over time, business conditions will change, as will available technologies. These

changes make BC/DR a perpetually moving target: priorities will shift, which makes recalculating RPOs and RTOs vital. New technologies can change what's possible—or at least what's affordable—and will require occasional reworking of the organization's technology investments and service purchases.

Best practice dictates checking disaster recovery regularly, and enacting full-blown recovery drills anywhere from annually to quarterly. Other tests should occur more frequently, including convening the DR team monthly to keep team members in sync. It's a good idea to bring backup personnel into the meeting once every two months, to convey current practices to those who must step in when primary team members might be unavailable. In general, intervals between tests should reflect ongoing changes in how the organization conducts itself, and how often network configurations, staff, technology tools and platforms, and compliance requirements (don't forget those!) change.

Protecting a Work-from-Anywhere Workforce

In view of current economic conditions, and the necessity for remote work scenarios, organizations must also think about protecting the assets that remote workers use and need to do their jobs from home or on the go. Organizations must revisit their technology choices and harden their security around the following assets to keep remote workers (and the apps, services, and data they use) safe and sound:

- **VPN.** The link from outside the network boundary into your systems and resources must be strongly encrypted and hardened against attack, yet still allow users to be productive.
- **Remote access infrastructure.** The servers, software (including protocols and services), and network connections that bring users to your boundary must also be robust and secure.

- **Endpoint protection.** Devices at the end of the remote link into your organization's networks should be protected against malware, sniffing, and vulnerabilities. Protection includes patching and updating remote systems, monitoring traffic for anomalies and attack signatures, and the like.
- **Edge infrastructures.** These are growing increasingly important as compute, network, and storage capabilities migrate closer to users out at the edge. These should get the same scrutiny and protection as your data centers, though at a smaller scale.
- **Edge data and applications.** Increasingly used to support IoT, data acquisition and analytics, and AI- or ML-based applications, these activities, too, must have strong protection and robust, secure access to storage, networking, and compute capabilities.

The migration of the security perimeter to the network edge and onto the proverbial last mile requires organizations to rethink and rebuild their remote access and networking capabilities. Secure, protected, and efficient storage has an important role to play in this effort.

CHAPTER 4

Pure Modern Data Protection

In This Chapter:

- Understanding business-centric IT
- The importance of availability
- The value of rapid recovery

The combination of requirements for storage and backup calls for Agile, secure, and integrated storage solutions, with cloud-native capabilities. Enter [Pure Storage](#)[®], a global technology company headquartered in Mountain View, California. Founded in 2009, Pure has been profitable since 2017. Through the years, Pure has acquired several companies including Kubernetes data service company Portworx.

At present, Pure Storage develops flash-based storage solutions for data center use. It also incorporates deduplication and compression software to reduce the amount of data written to and read from its drives, to boost capacity and speed throughput. Its primary hardware-based product lines, some of which runs its own custom-built OS named Purity, include:

- [FlashBlade](#)[®] for unstructured data is a Unified Fast File and Object storage platform that serves as the storage foundation for high-performance and high-throughput access to file and object workloads.

- [FlashArray™//C](#), based on QLC flash devices, provides an all-flash storage infrastructure that eliminates imaging delays and accelerates business-critical application performance and response.
- [FlashArray//X](#), based on high-end NVMe flash devices, also provides an all-flash storage infrastructure that eliminates imaging delays and accelerates business-critical application performance and response
- [FlashStack®](#), developed with Cisco, to provide a complete compute, network, and storage solution for modern IT infrastructures

AIRI®, developed with NVIDIA to extend NVIDIA's DCX A100 systems with FlashBlade devices, supporting a modern AI infrastructure that accelerates end-to-end GPU workflows

In addition, Pure Storage offers software and related services and solutions. These include:

- [Pure as-a-Service™](#), software-defined storage that works on-premises and in the cloud; it unifies your environment with a single subscription and one set of storage services.
- [Portworx®](#), a cloud-native, Kubernetes-based data services platform.
- [Pure1®](#), smart storage management.
- [Evergreen™](#), a storage subscription service that permits rapid upgrades and storage expansion without disruption.
- Pure [Cloud Block Store™](#), easy data mobility, data protection, and consistent, flexible block storage for AWS and Azure.
- [Purity](#), an operating system that runs on Pure Storage's own FlashBlade and FlashArray products to provide secure, highly scalable, and user-friendly storage and data management in hybrid cloud environments.

Thus, Pure Storage has the systems, software, skills, and knowledge to help enterprises implement the best and most cost-effective storage for cloud-native containerized apps and services, especially those within the Kubernetes ecosystem. In addition, the company's partnerships with Cisco and NVIDIA provide it with added synergies that organizations may find useful.

You needn't simply take our word for it, either. Pure Storage boasts a [Net Promoter Score](#) (NPS) in the top 1% of industry B2B scores. The company's latest audited score is more than double the B2B score of 40, and is placed among the highest such scores.

In 2021, Pure Storage [FlashBlade](#) and [FlashArray](#) received [recognition](#) as a Gartner Peer Insights Customers' Choice. This comes from [Gartner's Peer Insights](#) review program, which surveys customers to obtain feedback. Gartner provides unfiltered results that include numerous rave reviews. Pure achieved another distinction from Gartner: being positioned highest for their Ability to Execute and Furthest for Completeness of Vision in the 2020 [Gartner Magic Quadrant for Primary Storage Arrays](#).

Understanding Business-Centric IT

“Business-centric” is a concept that always seems to get people talking, but what does it mean in actual practice? Although it may seem obvious that the term focuses on the business, such a philosophy is often slower to permeate some parts of an organization than others. It's vital to recall that when you create a business-centric approach, you must focus on what's important to the business, not just what's important for IT success. That's why involving stakeholders from across the organization is important to bring focus to the right aspects and priorities and ultimately to making business-centric IT happen.

For example, if your company is in the hospitality industry, you probably depend heavily on online bookings, be they from your own or partner websites. You also depend critically on your point-of-sale

systems at various properties for check-in, guest billing, service orders, maintenance, and other day-to-day transactions and activities. When those systems are inaccessible, down, or not available for employees or customers to access, your business is in jeopardy.

The best way to think about IT doesn't turn on whether data is unstructured or structured, but in terms of the key mission of your business. This forcefully demands answers to the questions: What is your company's mission? In other words, what is it that makes your organization successful? Whatever the answer might be, that's where the primary focus should be for your version of business-centric IT.

An airline's chief mission, for instance, is safely moving passengers, crew, cargo, and planes from one destination to another. For such a mission to succeed, what platforms, systems, applications, data, and so on are required? Here, again, that's the primary focus for your version of business-centric IT. This list of key elements must then be accommodated within a technology solution that will help meet those expectations and thereby fulfill the mission.

This definition of purpose can all be accomplished without deep technical jargon or complexity. Rather, it works best when couched in business terms. In today's modern businesses, IT and the stakeholders need to collaborate more closely and effectively to ensure that the business meets its objectives, and doesn't waver from mission success.

The Importance of Availability

From a business perspective, availability is making applications, data, and processes accessible to users or to consumers. This means that however the organization defines such access, it allows "business as usual" to proceed unhampered and unhindered.

When key applications or data become inaccessible, the ripple effect can be horrendous across the board. Thus, the primary purpose of

modern data protection is to ensure data accessibility and availability. It's what lets organizations get back to business as usual as quickly as possible.

The best way to start this journey is to keep your data from being lost in the first place. Building data resiliency into your environment starts with a highly available infrastructure. One example comes from the Pure Storage FlashArray with [Purity ActiveCluster™](#). Its active-active synchronous replication provides a transparent, automatic, and non-disruptive failover between sites, as well.

Another example is Pure Storage FlashBlade, Unified Fast File and Object platform. FlashBlade's scale-out metadata architecture can handle tens of billions of files and objects with maximum performance and rich data services. [Purity//FB](#) supports cloud mobility with object replication and disaster recovery with file replication.

This kind of availability and performance is a good option to consider for the business continuity component of a data protection strategy.

Unified Fast File and Object

Many of Pure's Data Protection offerings take advantage of the company's Unified Fast File and Object (UFFO) platform. Even though the term may be new to some, UFFO storage is quickly becoming the only category of storage that can address modern digital transformation data requirements in many use cases, including data protection.

Organizations today need a storage solution that addresses modern data requirements, provides simplicity and multi-dimensional performance, and enables consolidation of key unstructured data workloads. These capabilities decisively eliminate storage silos, and provide profound returns on investment.

It's not just that the storage platform can store both file and object data, but it delivers outstanding performance for both. Its characteristics are best understood as follows:

- **Multi-Dimensional Performance** means very high throughput and IOPS with low latency to support multiple workloads simultaneously, including those with small or large files, sequential or random I/O access, batched or real-time jobs, and large numbers of files.
- **Intelligent Architecture** means that the storage system is built from the ground up to truly leverage the performance and efficiencies of flash storage. It is also simple to deploy, manage, and upgrade without requiring constant tuning. A modern storage solution must be simple enough so that its operation doesn't overwhelm storage admins. In fact, it relieves them of the mundane tasks involved in managing networking complexities when deploying the system, volumes, cluster pairs, aggregates, and flash caches or configuring replication.
- **Cloud-Ready** refers to its cloud-like agility, flexibility, and consumption choices with on-premises management and control.
- **Always Available** describes the capability of going beyond traditional platform resiliency. Maintenance operations, software upgrades, and capacity expansions are completed without disruption. The software design makes it possible for Pure Storage solutions to deliver high availability over multiple years and upgrade scenarios.
- **Dynamic Scalability** refers to the ability to seamlessly scale not only capacity but also performance, metadata, number of files and objects, and more.
- **Multi-Protocol Support** means that a single platform provides native file and native object protocol support without compromising performance or any functionality.

The Value of Rapid Recovery

Aside from a complete facility disaster or outright failure, one common issue that arises in most day-to-day operations is restoration of a corrupt, lost, or otherwise damaged file, directory, volume, or virtual machine.

When a business interruption hits your organization, the last thing business leaders want to hear is, “We’re working on it.” What they want to know is when business as usual can go on. That’s why recovery is such a critical component of your data protection strategy and is what success hinges on. Remember: Backup is critical, but the ability to quickly restore will define your success. Tech-speak no longer cuts it in the current age—talking in terms of availability and “business as usual” presents a business-centric approach that puts you in the good graces of the powers-that-be.

Keeping availability high, however, means keeping your infrastructure humming along. If you want to eliminate bottlenecks associated with traditional purpose-built backup appliances (PBBAs) using spinning disk, or if you’re still using tape as your primary backup destination, an all-flash solution is essential to bring your operations into the modern era.

And Pure is a leader in this space. The company first introduced FlashBlade in 2016, and since then has added significant features that provide compelling reasons to consider it when looking to upgrade your data protection environment. Note that there is no silver bullet when it comes to backup and recovery: It’s hard to architect a highly performant and resilient infrastructure. Going with FlashBlade, however, is a good place to start when modernizing a data protection strategy.

Why? FlashBlade requires no changes to your existing data protection software or the processes on which your IT organization has standardized. Its flexible nature allows your IT teams to offer a wide range of recovery options and multiple tiers of service.

Because Pure FlashBlade works with existing data protection software solutions, your IT teams can continue to service mission-critical recovery and compliance requirements. They can continue to protect their most essential data using preferred solution providers, including Commvault, Veeam, Cohesity, Rubrik and Veritas.

Where SafeMode Snapshots Come into Play

Pure Storage introduced SafeMode™ snapshots in 2019. This technology comes as a built-in feature on all Pure Storage systems, including FlashBlade and FlashArray. SafeMode snapshots create read-only snapshots of backup data and associated metadata catalogs after a full backup is performed. Organizations can recover data directly from these snapshots, to guard against ransomware or insider attacks (think: rogue administrator), and against backups and key data assets. SafeMode snapshots confer the following advantages:

- **Enhanced protection.** SafeMode snapshots cannot be deleted, modified, or encrypted. In fact, only authorized employees of your organization working with Pure Technical Support can configure the feature, modify governing policy, or delete such snapshots manually.
- **Backup integration.** SafeMode snapshots use the same snapshot process irrespective of the backup solution or native utility used to manage data protection processes.
- **Flexibility.** Admins can set their own snapshot cadence and deletion schedules for themselves, as part of flash system setup and maintenance.
- **Rapid Restore.** The underlying hardware delivers a massively parallel architecture. Its elastic performance scales with data to speed both backup and recovery.

Immutable snapshots such as SafeMode snapshots provide tremendous reassurance. Because they cannot be deleted, modified, or encrypted, they're essentially immune to ransomware attacks. Policy changes for making, storing, and deleting SafeMode snapshots occur only with vendor participation and permission, so an attacker's typical "next move" when stymied—namely, to use privileged access to reset permissions, then wreak havoc—is blocked in advance.

Becoming Cloud-Ready

If hybrid or multi-cloud isn't in your strategy today, it almost certainly will be in the future. Therefore, choosing a cloud-ready solution is just smart planning. Cloud storage is already in use today for a variety of purposes that include backup, DR, and long-term retention. If such capabilities aren't already in your portfolio, you should be considering them soon.

As you consider your storage needs, and the options available to help meet them, it's crucial to look for APIs and services that permit easy integration of storage as a standard service layer. These should be available in any and all of the clouds, both public and private, that you use, and should likewise be available in your on-premises applications and services. Containerized applications and services built within the Kubernetes ecosystem can deliver such capabilities both easily and readily.

Maximizing Data Agility

To begin with, it's interesting to consider what the term "agile data" really means. Let's start by considering the state of protected data in the past. Most of the time, it sat in a proprietary format on some cold device (tape or some other removable media). Thus, the most an organization could do with such protected data was report on files stored, creation and modification dates, and so on. But when that

protected data reached maturation, its life and usefulness expired. If the media couldn't be reclaimed and re-used, it was best destroyed for security's sake.

Today, that's no longer the case, because data is used for much more than recovery. Data agility is a way of extracting more value from backup data by using it in multiple ways, rather than letting it rot away on a hard drive under a mountain. For example, that protected data could be used to create a virtual lab for DevOps. That's a better use of that data, isn't it?

In the same vein, a flash-based storage strategy breaks down silos within backup data. Thus, it permits other workloads, ranging from analytics to application testing and development, to use that data when it isn't needed for business recovery. This flexibility is significantly different from purpose-built backup appliances (PBBAs), designed for and dedicated exclusively to backup and recovery.

Doing More with Your Data

The most interesting outcome from exercising data agility is that you will find more and better ways to use your data, as you and your staff become more familiar with what it can tell you about your business. Backup/recovery data becomes more accessible and usable when it can be accessed more easily and quickly. This in turn leads to more frequent and varied use of that data as analysts, researchers, and data scientists begin to understand what treasures are now at their disposal.

Organizations often start out with one or two pilot programs, but soon find themselves using the data for more complex and sophisticated purposes, particularly for training or improving AI- or ML-based systems and applications. Ultimately, this means organizations can take advantage of what they know, and what they learn, to do more for the business, and to be more efficient, productive, and innovative. It's a stealth benefit for modernization.



If this is appealing, consider that FlashBlade can become your data hub, serving DevOps needs and more, including as a target for analytics, AI, and ML, for training data, and for historical data analysis. And all this happens without an impact on production. For security, Pure provides data-at-rest encryption for all arrays to add an extra layer of protection against the bad guys. It provides more peace of mind for the business.

Data agility makes a world of difference to a business-centric approach. There's much more value to be gained by using assets than from allowing them to quietly expire and then get retired. It can also help IT by eliminating multiple silos, complexities, and additional cost burdens that would otherwise eat into annual IT budgets.

CHAPTER 5

Real Customers, Real Cases

In This Chapter:

- City of New Orleans
- Domino's Pizza
- Making the most of flash storage

Pure Storage prides itself on being a customer-first organization, and puts customers at the heart of every major decision it has made since the company was founded in 2009. The driving force is to deliver to its customers a better storage platform along with an enhanced, more modern experience for all those who work with and use its products, software, and services. In the sections that follow, you'll have a chance to read some customer stories and to learn more about best use cases where Pure Storage technology is particularly apt and helpful.

City of New Orleans

On December 13, 2019, the City of New Orleans fell prey to a [ransomware attack](#). After shutting down hundreds of its servers to forestall further damage, it needed a new IT infrastructure so it could test, cleanse, and restore its precious data. In under a week, Pure Storage helped the city's IT team migrate data to Pure Storage FlashArray NVMe and FlashBlade object storage. Its new storage architecture lowers risks with fast backup, restore, and immutable SafeMode snapshots that are immune to further ransomware attacks.

In this case, the Pure Storage FlashArray replicates data to a second data center for fast backup and restore. Pure Storage FlashBlade, with SafeMode snapshots, protects its backups from new ransomware attacks. New Orleans was able to transform its IT infrastructure quickly and efficiently and achieved the following benefits:

- New storage required only a minimal learning curve, and provided accelerated time to recovery. Ultimately, the city derived additional value from easier access to backup data for analysis and insights.
- Greatly reduced risk of repeat attack, with space-saving snapshots that were faster to backup and restore. Improvements in efficiency and resource consumption produced tangible cost savings.

Using SafeMode, recovering data in the wake of a ransomware attack follows a four-step process, as IT staff working with Pure Storage support team members learned:

1. Delete compromised data.
2. Reinstall backup software.
3. Point backup software at metadata catalog in a SafeMode snapshot.
4. Begin recovery.

Kimberly LaGrue, CIO for the City of New Orleans, is [quoted](#) as saying this about its disaster response to the attack: “For us, these were well-rehearsed plans that we just had to enact immediately.” Local staff responded to the attack by shutting down all 470 of the city government’s servers and thousands of endpoint computers. This brought services to a temporary halt, but limited the scope of infection, and allowed services to be brought back easily using proper recovery techniques already described.

With a new Pure Storage architecture in place, the City of New Orleans is better prepared to weather the next attack, and expects it will be able to respond more quickly to such things going forward.

Domino's Pizza

Domino's sells over 3 million pizzas each day from its 17,000-plus stores around the world. It collects huge volumes of customer data in that process. In fact, that data is vital to establishing customer preferences and order behavior, and to delivering positive customer experiences. To that end, Domino's uses Pure Storage to capture and analyze customer data rapidly, with the ability to scale as needed to support ongoing growth and expansion. Because Pure Storage captures customer data across all franchises, Domino's is able to learn and adapt to what customers want. It keeps innovating continuously to maintain its leadership position in the food delivery business.

To achieve business transformation, Pure Storage offered insights into customer preferences and behavior. In turn, this helped to fuel positive customer experiences, and increase repeat business frequency and order size. The flexible and adaptable cloud-based environment that Pure Storage provides also lets Domino's launch new applications rapidly, scaling storage as needed. Pure Storage technology has helped Domino's start to make inroads into autonomous delivery, artificial intelligence, and more.

On the IT side, Pure Storage has boosted performance and resilience to support Domino's rapid growth and innovation. The Pure Storage architecture lets the company analyze data at speed and scale, to feed it directly and quickly back into customer experiences. Furthermore, Pure Storage has helped Domino's simplify storage management, which makes it easier for the company to protect its valuable data while putting it to work for a wide variety of use cases.

Dan Djuric, Domino's VP of Global Infrastructure and Enterprise Information Managed, is [quoted](#) as saying "Our data works as hard as our delivery drivers. It's all about speed—the ability to ingest data, analyze it, and feed it back into actionable channels." In fact, [Domino's](#) considers itself to be "a technology company that sells pizza" and has long been a leader in online and mobile ordering with innovative

technology emerging from its Domino's Tracker online order tracking facility. Because more than half of its business comes from online orders, the company has plenty of data from which to work, analyze, and learn.

Make the Most of Flash Storage

This Gorilla Guide has covered a broad range of topics, including the state of the storage market, the challenges that face organizations seeking to modernize, the impact of cloud-based storage both public and private, and the insight, capability and solutions that Pure Storage can bring to bear for organizations seeking to build modern, flexible and adaptable IT infrastructures.

One message should be crystal clear from this screed—that old ways and tools for doing backup and recovery can't keep up with today's challenging environments and their associated RPOs and RTOs. As the cloud continues to change everything, there's more data to acquire, store, and protect than ever before, and that data lives everywhere from the edge to the core to the cloud.

As data remains the crown jewel of organizations everywhere, it's increasingly essential to make sure it's protected and resilient, and that it can be put to new, agile uses in all kinds of interesting and informative ways. By keeping data safe, secure, and easy to recover, organizations can respond effectively to business interruptions and even disasters, whether man-made (often by ransomware) or resulting from nature's wrath. Such organizations are the ones that will survive in the face of inevitable outages and interruptions. In fact, such organizations may not even notice or care when the power fails or an intruder gets in because they'll be ready for whatever happens next!

Is your organization ready? If you're unsure, please check out what Pure Storage can do to help. For more information, visit Pure's [Modern Data Protection Solutions](#) or send an email to info@purestorage.com.

ABOUT PURE STORAGE



Pure Storage gives technologists their time back. Pure delivers a modern data experience that empowers organizations to run their operations as a true, automated, storage-as-a-service model seamlessly across multiple clouds. Pure helps customers put data to use while reducing the complexity and expense of managing the infrastructure behind it. And our solutions help mitigate ransomware attacks with immutable snapshots and rapid, petabyte-scale recovery. Pure software and hardware solutions are easy to use and consistently evolving to better meet customer needs—without costly forklift migrations. For the past seven consecutive years, Gartner has named Pure a Leader in the Magic Quadrant for Primary Storage Arrays, and FlashArray and FlashBlade are 2021 Gartner Peer Insights Customers' Choice winners. With a certified customer satisfaction score in the top one percent of B2B companies, Pure's ever-expanding list of customers are among the happiest in the world.


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/>