

DATRIUM PRESENTS

THE GORILLA GUIDE TO...[®]



Failproof Cloud Disaster Recovery

Ed Tittel, James Green, Dan Keldsen, Alan R. Earls

INSIDE THE GUIDE:

- The Incredible Benefits of DRaaS
- Stop Ransomware Dead in Its Tracks
- Instant RTOs: Yes, You Can!

**HELPING YOU NAVIGATE
THE TECHNOLOGY JUNGLE!**

 ActualTech Media
www.actualtechmedia.com

In Partnership With

 Datrium[®]

THE GORILLA GUIDE TO...

Failproof Cloud Disaster Recovery

AUTHORS

Ed Tittel, James Green, Dan Keldsen, Alan R. Earls

EDITOR

Keith Ward, ActualTech Media

LAYOUT AND DESIGN

Olivia Thomson, ActualTech Media

Copyright © 2020 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

ENTERING THE JUNGLE

Introduction: No More Living in the Past	6
Chapter 1: No More Living in the Past	6
Chapter 2: Datrium DRaaS with VMware Cloud on AWS	7
Consumerization of IT—and DR.....	9
Simplifying Failover/Failback.....	12
Understanding DRaaS.....	14
DRaaS Connect.....	15
Chapter 3: Enable Effortless DR by Unifying Data Silos	17
The Datrium Automatrix Platform.....	19
Datrium DVX.....	22
Datrium DRaaS.....	26
Automatrix Platform Use Cases.....	28
Chapter 4: Livin' La Vida DR	30
Always-On Data Integrity.....	30
Compliance Audits? No Worries.....	35
Traditional DR—Recovering from Natural Disasters.....	39
Plays Well with Others.....	41
Chapter 5: Mastering the Art of Recovering from Disasters and Ransomware	44
The Threats Are Piling Up	44
Understanding the 3 Pillars of Backup and DR.....	46
How Modern Disaster Recovery Is Better.....	47
Datrium DRaaS with VMware Cloud on AWS	49
Be a Disaster Recovery Superhero	51


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!

INTRODUCTION

No More Living in the Past

Welcome to The Gorilla Guide To...[®] Failproof Cloud Disaster Recovery. If you're like most people on this side of the IT industry, you probably wonder how such a claim—failproof recovery—can even be made. After all, as anyone who's ever attempted to recover from a disaster knows, the process is filled with problems, from bad backups to messy refactoring of VMs to recovery time frames so slow that they can be measured with sun dials.

All that was true in the past. Legacy disaster recovery, or DR, was difficult enough with simpler infrastructures in which all your data, systems, and applications lived on-premises. In today's environments, with the focus on the cloud, Internet of Things, extreme mobility, and so on, it's simply impossible to use the old DR methods. What's required is to leverage the very technology that makes things complicated—the cloud—to smooth out the rocky road to fast DR that can have you up and running as fast you need to—even if that means near-instant recovery.

Yes, recovering that fast is not the impossible dream. It's a reality, and one you can take advantage of today. In this Gorilla Guide, we'll show you how to do just that. In these chapters, you'll learn how DR can be fast, easy—as easy as pushing a button—and, yes, failproof.

Too good to be true? Read on and see for yourself. It starts with the concept of disaster recovery as a service (DRaaS), and how one innovative company has started with that idea and built on top of it to take DR to an entirely new level. That company is Datrium.

CHAPTER 1

Datrium DRaaS with VMware Cloud on AWS

Datrium DRaaS takes a fresh, innovative approach to backup and DR, updating antiquated methods that are hopelessly incapable of meeting modern needs for data center uptime. It operates in the same way as most software as a service (SaaS) offerings—but unlike traditional DR platforms, which leverage on-premises orchestration software, Datrium DRaaS employs a SaaS model for orchestration and built-in backup. Datrium DRaaS combines cheap and deep storage, using Amazon S3 for backup snapshots of VMware virtual machines (VMs), with powerful, policy-driven DR orchestration and management capabilities.

In addition, you're paying for consumption-based economics for backups and on-demand compute when you need it. And since it's vSphere everywhere—both on-premises and in the cloud—admins enjoy operational consistency across environments—this results in greater efficiencies and saves even more money.

Datrium DRaaS simplifies traditional DR—a complex, human-driven, and multiplatform patchwork of systems—into a simple push-button environment for failover (and failback). The secret to Datrium DRaaS's instant RTO comes from its ability to start VMs directly from backup images in VMware Cloud on AWS. The VMs are powered on instantly via a live, cloud-native NFS datastore mounted by ESX hosts in the new software-defined data center (SDDC).

This means that the recovery process need not wait for rehydration of backup images to make them ready to run. Instead, the Datrium DRaaS

environment can simply spin up the set of snapshots needed to get recovery underway, as well as restore access to the IT infrastructure(s) that those snapshots capture (in the cloud or a different availability zone, rather than on premises in the data center or some other now-failed installation).

Datrium DRaaS also offers end-to-end encryption for all data (in motion and at rest) by encrypting all traffic (and files) under its control to boost security and block unauthorized access.

Through its orchestration services, Datrium DRaaS also provides essential DRaaS functions that include:

- Continuous compliance checks (to make sure DR is ready for invocation at any time)
- Periodic integrity checks to make sure that data and applications are sound and haven't been tampered with

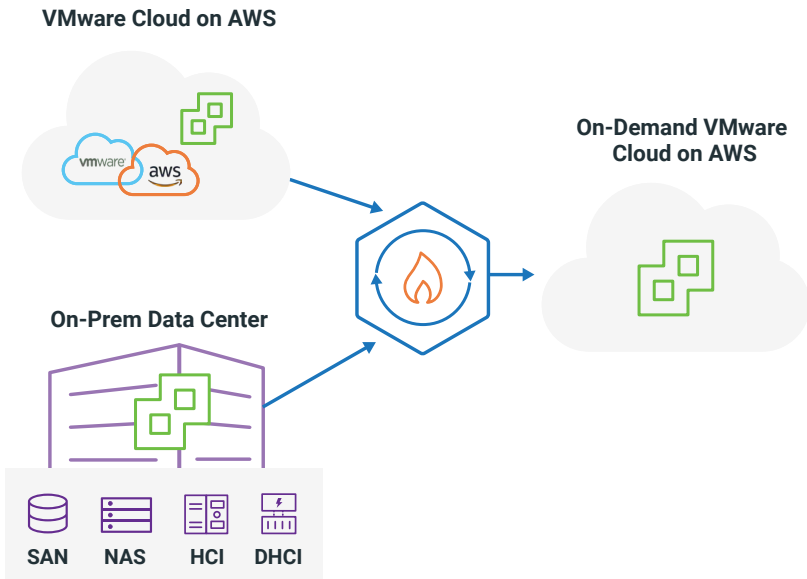


Figure 1: Simple, comprehensive backup and DR with Datrium DRaaS

- Reporting functions that make organizations audit-ready at all times
- DR test capability to show it's working and doesn't interfere with production IT environments

Figure 1 shows how Datrium DRaaS can back up VMware VMs either in an on-premises data center or in some VMware Cloud on AWS availability zone. At the push of a button, Datrium DRaaS can recover either kind of environment into an on-demand VMware Cloud on AWS.

That's because the VM snapshots stored in the recovery cloud can be started quickly, at any time, to take over for VMs from either the data center or another cloud environment.

Consumerization of IT—and DR

Datrium's goal in offering DRaaS is to create a simple, straightforward, cloud-based solution that supports both everyday backup/restore functionality and DR capability. This lets customers choose the objectives they want for Recovery Point Objectives and Recovery Time Objectives (RPOs/RTOs) to fit within their budgets and risk-management profiles.

At the same time, this approach suits the increasing consumerization of IT as it moves into the cloud, as a simple push-button alternative takes over for what used to be totally customized DR, tailored for each customer's individual circumstances.

Consumerization also yields a major reduction in the time, effort, and expense involved in setting up and maintaining DR regimes and services. More traditional methods of implementing DR are incredibly labor- and resource-intensive. Datrium DRaaS lets its users take advantage of a SaaS-based DR product that's easy to use and quick at failover/failback operations.

Undoing Complexity

The biggest value-add from Datrium DRaaS comes from its ability to undo or even eliminate the complexity normally associated with DR in any organizational setting, as well as deliver instant RTO. Datrium DRaaS offers a more reliable solution that reduces DR's overall complexity and costs.

Datrium DRaaS offers benefits beyond simplification, too:

- Switching from existing approaches and solutions usually means switching from multiple, loosely integrated point solutions (one or more each for backup, storage, cloud access, and so forth) to a single, coherent solution that combines backup and DR. Users can be confident that recovery works as it's supposed to, precisely when it's needed.
- Data moves offsite as part of instant failover (and failback) capabilities into (and out of) the public cloud. DR's inherent need to run "somewhere else" is baked into this solution.
- Data is always encrypted, whether in motion or at rest, so organizations are more secure. This not only offers protection against sniffing and snooping (penetration and breach attempts), it also fends off ransomware (unauthorized encryption makes data inaccessible to its intended users).
- Organizations can test DR at will, thanks to non-disruptive test facilities in the cloud—with no impact on production infrastructure, services, or data.
- Built-in checks ensure data integrity; built-in reports support audit and regulatory compliance, too.

Datrium DRaaS can reduce DR costs in various ways. First, those who operate secondary DR sites can shut them down. This can obviously be an enormous savings. Second, Datrium DRaaS recovery only needs to run when a disaster occurs, so costs of "hot" operation in the cloud

are vastly reduced. Third, because Datrium DRaaS combines DR and backup in a single solution, users need no longer acquire and manage such products separately.

Overall, these benefits not only reduce complexity from many sources and on many levels, they also improve the ROI in (and the value of) DR and backup.

Consolidation and Simplification

A major benefit of Datrium DRaaS is that there's no need for a secondary colocated or mirror site for failover when a disaster is declared. Customers can cut over from their primary data center running either Datrium's own DVX platform or, by leveraging the DRaaS Connect utility, any VMware-centric storage, to the cloud-based failover site in VMware Cloud on AWS.

In fact, prospective customers love the simplicity of Datrium DRaaS, as shown in the company's demo video¹. It presents seamless failover to the cloud as the result of a single press of a button. Once deployed, the ability to test DR at need and on demand without disrupting production computing is a huge incentive for many organizations to buy in.

Then, too, the solution's end-to-end encryption meets most companies' needs for data protection. Also, the ability to perform compliance checks—i.e., making sure DR works, and proving application and data preservation and integrity—is increasingly important for organizations.

Ongoing, built-in data integrity checks happen automatically four times a day, and built-in reporting meets monthly (or more frequent) audit requirements. In fact, the Datrium DRaaS environment supports continuous compliance checks which run every 30 minutes, as it constantly matches configurations against discovered items and

¹ Source: <https://youtu.be/ZkalbNIPgS4>

elements, and checks data consistency objectives against distributed business data.

Organizations that use DR runbooks soon realize they become obsolete almost immediately. The Datrium DRaaS continuous compliance checks overcome this uncertainty about whether the DR plan will work or not. In addition, Datrium DRaaS users are always audit-ready, because its continuous compliance checks generate auditable reports.

Simplifying Failover/Failback

For other solutions, failover to the cloud is complicated when re-instantiating images at the primary site into VMs in the cloud. The issue is that VMs from the primary site must be refactored into cloud-native formats upon failover.

Thus, failover involves a conversion process prior to start-up and access/operation. This takes more time to complete than a simple transfer of control takes in and of itself. It's not unusual for this process, often called rehydration, to take anywhere from several hours to as long as a many days, depending on the number and size of VMs involved.

Running the Runbook

A runbook—the set of operating instructions, procedures, and data sources and targets to enact DR—is a key element in commencing either failover or failback. Runbook configuration has historically been complex and highly technical.

For the 2020s, organizations need a runbook configuration process driven by business requirements, with smart technology behind the scenes handling the complex technical activities involved. Datrium DRaaS offers such capability, thanks to its powerful automation facilities and sophisticated configuration handling.



That transformation must also be run in reverse to restore normal operations at the primary site during failback. For such solutions, then, both failover and failback are more complex and time-consuming than they could or should be. Not so for Datrium DRaaS—both failover and failback occur in minutes. Neither requires an image transformation to run, whether control is passed to the data center or the cloud.

The real crux of DR involves ensuring that all critical elements are accessible, manageable, and under control during failover and failback. Those elements include:

- Primary storage (in the case of Datrium DRaaS, primary storage could be Datrium DVX; or, using the DRaaS Connect VM, any VMware-centric third-party storage)
- Backup
- DR Orchestration
- Security
- Mobility

Datrium DRaaS has all this covered, resulting in a significantly improved DR experience as well as improved RPO/RTO intervals. Along with a simplified recovery process (for both failover and failback), this means far less time, effort, and expense.

Datrium Failover

With Datrium DRaaS, the real effort goes into making the decision to invoke DR, which means moving operations from a primary data center to the cloud. Once that decision is made, failover is handled automatically with the simple push of a button. For example, organizations with a typical four-hour RTO (or equivalent SLA requirements) can take 3.75 hours to try to fix the primary site and avoid a disaster declaration altogether. The remaining 15 minutes will more than suffice to handle the failover to the cloud, should that prove necessary.

Datrium DRaaS

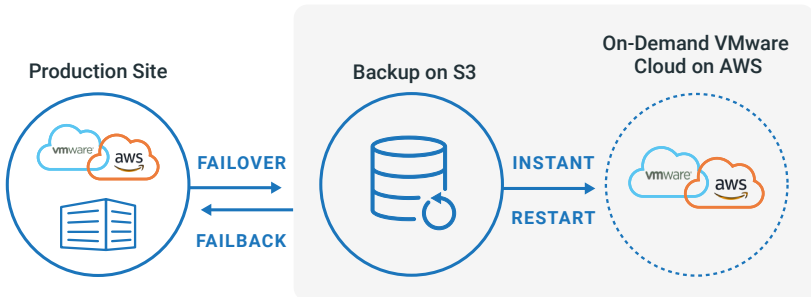


Figure 2: Failover/failback in Datrium DRaaS is automatic—just push a button

Datrium Failback

Failback reverses the failover process: VMs are moved from the cloud back to the on-premises data center environment. This, too, happens with a single UI selection in Datrium DRaaS. But because failback is entirely discretionary—at the command of the customer—this can occur whenever it makes sense to restore normal operations. Again, the cutover should occur within a 15-minute time window.

Figure 2 shows that failover/failback moves between a (primary) production site to Datrium DRaaS, with its run-ready backup VM snapshots ready to launch into VMware Cloud when recovery is needed.

Understanding DRaaS

Datrium DRaaS includes DR orchestration that executes DR plans and runbooks. It provisions and monitors software-defined data centers (SDDCs) that run in VMware Cloud on AWS. Thus, Datrium DRaaS offers full AWS integration with an organization’s primary data center. It includes the one-step “DR Button” to initiate DR, and requires no additional third-party hardware or software.

Datrium DRaaS represents an insurance policy against disaster, but includes the financial advantage that it doesn’t have to be consuming

expensive resources at all times (though an always-on pilot light option is available to those who need it).

Through a simple UI, teams set backup policies and DR runbooks. Tamper-proof backups can be created every few minutes, every hour, every day—whatever makes sense for the business. Backups are de-duplicated, compressed, and encrypted, then stored in their native format in S3. Compliance checks run every 30 minutes.

When disaster strikes, failover into Datrium DRaaS is started. Datrium DRaaS automatically provisions VMware resources and an SDDC in VMware cloud on AWS. Its stored backups are instantly powered on via a live cloud-native NFS datastore mounted by an ESX host in that SDDC, resulting in instant RTO. Unlike legacy backup-only solutions, there's no time wasted waiting for backup data to be copied into an SDDC before any VMs can be restarted.

DRaaS Connect

DRaaS Connect is downloadable, lightweight software that protects all VMware workloads running in the cloud and on-premises, providing that protection just minutes after downloading. As shown in **Figure 3**, DRaaS Connect for VMware Cloud works in the cloud to allow VMware Cloud on AWS in one availability zone to fail over to another availability zone. Likewise, it can also protect a data center running on a vSphere on-premises infrastructure, including SAN, NAS, HCI, and even DHCI.

Because of these factors, Datrium DRaaS is much more cost-effective than running a physical mirror or failover site of any temperature (“cold,” “warm,” or “hot,” to use standard but now outdated 20th century DR terminology). Organizations avoid excessive egress fees, and incur compute-related charges only when a DR infrastructure is “turned on” in the cloud following a disaster declaration via the DR button. Also, the combination of low-cost backup storage for snapshots, and on-demand compute capability that runs only when a disaster is declared, delivers exceptional economies of scale.

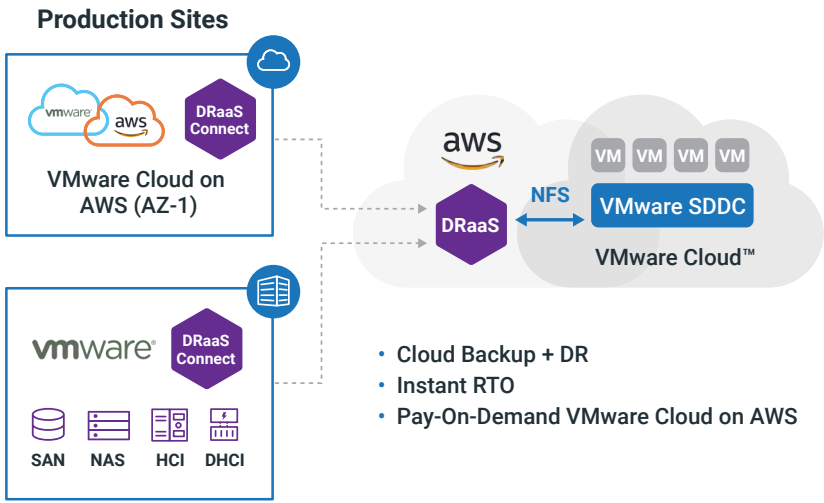


Figure 3: DRaaS Connect provides automatic failover from one availability zone to another

Now that you understand what makes Datrium DRaaS unique, let's turn to some of the revolutionary capabilities in the broader Datrium platform, and why it will make you re-think your entire DR strategy.

CHAPTER 2

Enable Effortless DR by Unifying Data Silos

It's no secret that IT infrastructure is becoming more fragmented. With a multicloud push underway by almost every organization with a significant IT footprint, data and workloads are being intentionally spread out and placed on isolated data islands. On each of those islands, a smorgasbord of single-purpose tools create their own data silos (see **Figure 4**).

There's a good reason for doing this, and there are valuable outcomes:

- Spreading data across multiple clouds helps organizations avoid cloud vendor lock-in
- Using on-premises clouds and any combination of public clouds, organizations choose the best cloud for a given workload type and leverage best-in-breed platform as a service offerings
- Avoiding putting all your eggs in one basket enhances availability and keeps a single cloud failure from bringing down your entire company
- Using data management tools that provide point solutions for a pressing pain point can resolve immediate tension with business leaders or auditors

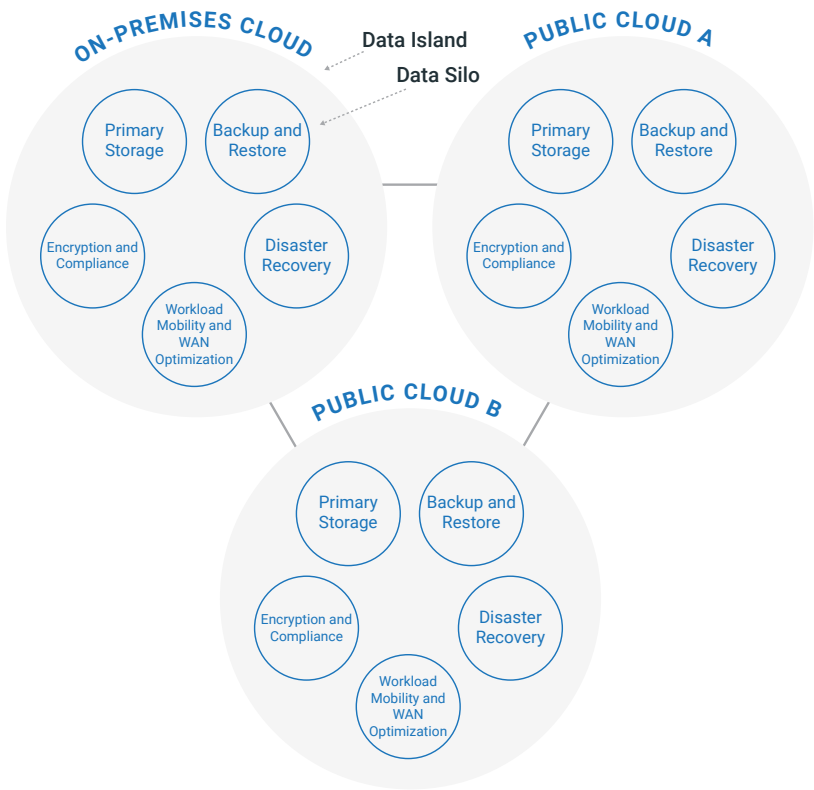


Figure 4: Multiple islands and silos of fragmented data threaten to undermine the multicloud vision that many organizations have

But there are also significant challenges with the ever-growing number of data islands and silos. For example:

- Each cloud has a unique user experience and a unique set of underlying assumptions that may or may not be familiar to IT administrators and developers
- Each point solution for data management has a unique view of the data and is managed from a separate console, creating a silo
- Storing data in multiple islands and silos can lead to duplication of data and storage inefficiency

- Data security becomes exponentially more complex to oversee as you add data silos

What organizations need today is a platform that unifies data silos and provides a consistent operational experience and a common data plane regardless of whether the data is on-premises or in a hyperscale cloud. The ultimate realization of the multicloud vision many organizations have depends on this unification.

In technology, there are often revolutions where a single, unifying technological breakthrough replaces a broad swath of discrete, stand-alone tools that have been commoditized with one powerful platform. Consider how the world was changed when Apple introduced the iPhone and replaced cellphones, MP3 players, PDAs, and cameras with a single device, which arguably did each thing better than the predecessors because of the common platform.

Data management and workload mobility is overdue for the same sort of breakthrough. The Datrium Automatrix platform is that revolution.

The Datrium Automatrix Platform

To adequately deliver an iPhone-like transformation to IT infrastructure, a solution needs to provide both a unified virtualization plane and a unified data plane.

Statistically, VMware vSphere and Kubernetes are far and away the industry mainstays when it comes to VM-based virtualization and container orchestration, respectively. Banking on these two platforms to be the unified virtualization plane is a relatively safe bet for the foreseeable future. VMware Cloud (VMC) on AWS gives organizations an opportunity to leverage the same virtualization plane in the cloud.

The missing piece, then, is the unified data plane, which can be found in the Automatrix platform. The Automatrix platform unifies five critical data functions into a single domain (see **Figure 5**).




Figure 5: The Automatrix platform unifies five critical data functions into a single domain

Primary Storage

Every IT infrastructure needs solid primary storage, and Datrium DVX provides enterprise-class storage and performance with a cloud-inspired model. Datrium DVX takes the best of trusted SAN storage and hyperconverged storage architectures and combines them into one high-performing, petabyte scale solution. A fully loaded DVX system is tested to deliver up to 18 million IOPS. The next section covers DVX in more depth.

Backup

Backup technology is due for a facelift, and there's no shortage of new entrants to this field. It seems that everyone and their brother knows that the world needs a new answer for backups. DVX uses its one-of-a-kind filesystem to perform and store backups in a totally new way and unlocks levels of recovery granularity and backup retention that were out of reach for most organizations just a few years ago.

Disaster Recovery

DR orchestration software has historically been incredibly complex. A few DR solutions on the market today make it much simpler than it used to be, but they also come with some important limitations that the Automatrix platform doesn't have. Thanks to Datrium DRaaS—a SaaS DR orchestration tool and part of the Automatrix platform—failproof push-button DR with VMware Cloud on AWS is within reach. Datrium DRaaS offers instant restore capability for near-zero Recovery Time Objective (RTO).

Mobility

In a multicloud world, the ability to move workloads and data from cloud to cloud at will is crucial. Having this capability allows you to avoid vendor lock-in. And make no mistake about it—hyperscale cloud providers *are trying* to lock you in. Hyperscalers provide plenty of tools to help you get data *into* their cloud. They provide very little help with getting out. The Automatrix platform neutralizes lock-in efforts regardless of the cloud.

Moving data between clouds is also challenging because of the laws of physics. Shuffling vast amounts of data around is tricky when you have to push a full copy of the data across the wire. Deduplication over WAN technology is required today if you want to migrate data between clouds in a reasonable amount of time.

Encryption

Security is a top concern for IT leaders today. Ensuring that your data platform is encrypting data in flight and at rest is an important step toward keeping your data safe. But without a unified platform like Automatrix, encryption is messy and comes with significant trade-offs.

The Automatrix platform is comprised of two main components: Datrium DHCI and Datrium DRaaS. Let's take a look at each a bit more in depth.

Datrium DVX

The engineering that went into the DVX platform is what allows Automatrix do the amazing things that it does. At a time when the industry hype was focused on a hyperconvergence architecture based on self-contained nodes with both compute and storage capacity, Datrium swam against the current and built a platform on a more loosely-coupled design where compute nodes remain stateless, but contain a local flash cache, and separate data nodes persist the data on cheap JBODs (**Figure 6**). Now that a handful of competitors have popped up, the industry has settled on a name for this architecture; it's referred to as disaggregated hyperconverged infrastructure, or DHCI. Datrium DHCI

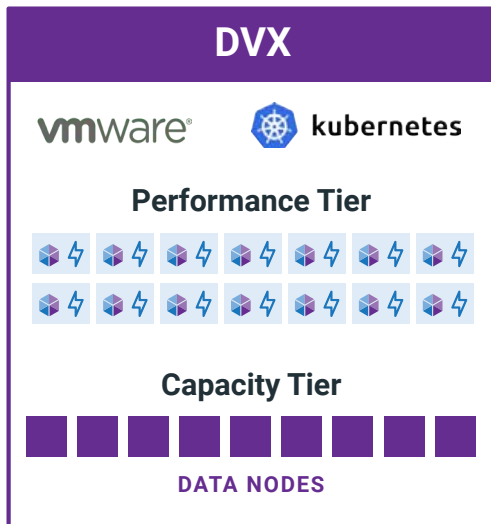


Figure 6: DVX has a two-layer design with a stateless performance tier and a durable, efficient capacity tier. This architecture is called disaggregated hyperconverged infrastructure.

can be configured with any compute node on the vSphere HCL and you can even use your existing equipment.

Efficiency and Durability to Boot

Deduplication, compression, and erasure coding are always on in DVX systems. And the system is constantly doing background integrity checks on the data. This ensures that your data gets maximum efficiency and durability without any knobs to turn.

Because the deduplication algorithm commonly delivers a 5x data reduction rate, a 1TB local flash cache in the performance tier often has a 5TB effective capacity. This means it's affordable to size the cache in a way that accommodates your entire working set. Thus, a very high cache hit ratio and workloads see sub-millisecond I/O performance.

Built-in Backup and Recovery

The shared data plane of the Automatrix platform allows very granular and efficient snapshots to be used for backup and restore, archive, and DR.

The DVX system takes advantage of its unique filesystem to snapshot VMs quickly and to store those snapshots without degrading performance. A single DVX system can snapshot thousands of VMs at a time and can store more than 1 million snapshots.

Snapshots are performed synchronously across the whole protection group versus sequentially—where one VM is snapshotted after the next and you have as many point-in-time snapshots as you have VMs. So, if you recover a whole protection group in the event of a disaster, all VMs are restored to exactly the same point in time thanks to the atomic snapshotting mechanism.

Rare is the environment where VMs don't depend on each other; so having the restore points be exactly the same across the whole protection group helps avoid application-level consistency issues during recovery.

The VM Write I/O Lifecycle

One key design feature that makes DVX interesting is the proprietary, log-structured filesystem. Here's the short version of the life of a VM write I/O in the DVX system (**Figure A**):



- DVX system receives incoming write I/O
- I/O is deduplicated, compressed, and checksummed in NVRAM; write is acknowledged
- In the background, the I/O is bin packed into an 8MB container object in memory
- Once the container is full, it's erasure coded and sequentially streamed to disk
- A copy is also placed in local flash on the performance tier

By taking a series of random I/O requests and sequentializing them in memory before interacting with slow disks, the filesystem overcomes the challenge of terrible random I/O capabilities with cheap spinning disks due to seek time. Effectively, storage capacity and storage performance have been decoupled.

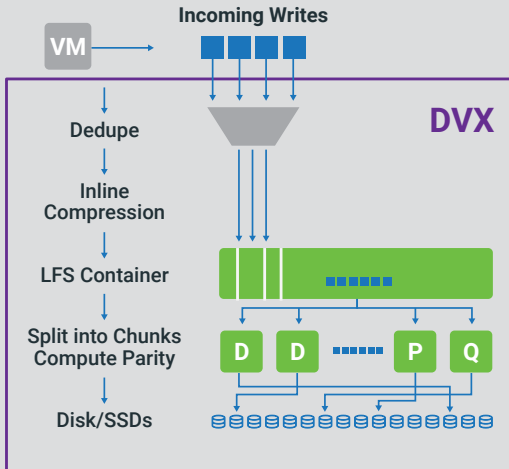


Figure A: The life of an I/O in Datrium's proprietary log-structured filesystem

Recover Instantly with DVX

Standalone backup systems have a primary goal of storing lots of backup data as efficiently as possible and improving backup and archive economics. That's great when you need to restore a file for an auditor. But that narrow objective can work against you in the case of sitewide disaster.

Recovery time is of the essence in a disaster, and when the focus of a backup system is purely on affordable long-term storage, restore performance can be lacking. Rehydration of deduplicated data can be extremely time-consuming at the exact moment when time is your most coveted asset. It can take many hours or days to restore from a legacy backup system, which is completely unacceptable in many cases today. Recovering files or VMs on DVX systems or with Datrium DRaaS is nearly instantaneous.

Make Your Data Invincible

Ransomware is maturing to the point where sophisticated strains will encrypt your backups, too. But that's a non-issue when your backup data is stored in a DVX system. Backup data is stored in a hidden, isolated namespace where backups are immutable. That means once a backup is written, it can never be changed—it's completely immune to ransomware and other types of data tampering.

Additionally, blanket encryption secures data at the host as soon as it's written to local flash. But because the DVX platform controls the encryption keys with an internal key manager, data can be encrypted in flight and at rest on the data nodes without impacting global deduplication and WAN optimization. Data on DVX systems is both safe *and* efficient.

Datrium DRaaS

DRaaS comes together with DVX and any other third-party vSphere storage to provide a true multicloud infrastructure with the operational simplicity of a single, unified system. Datrium DRaaS can take advantage of the high level of automation available with VMware Cloud on AWS to create an on-demand DR site where you only pay for what you use.

Recover in the Cloud

The underlying engineering that makes Datrium DHCI capable of the impressive things it can do on-premises was also instrumental in the design of Datrium DRaaS. Namely, the same sorts of architecture decisions that were required to decouple flash storage performance from spinning disk capacity and economics now make it possible to run a high-performing storage system in cloud compute instances with affordable and deep cloud object storage as the persistence tier (**Figure 7**).

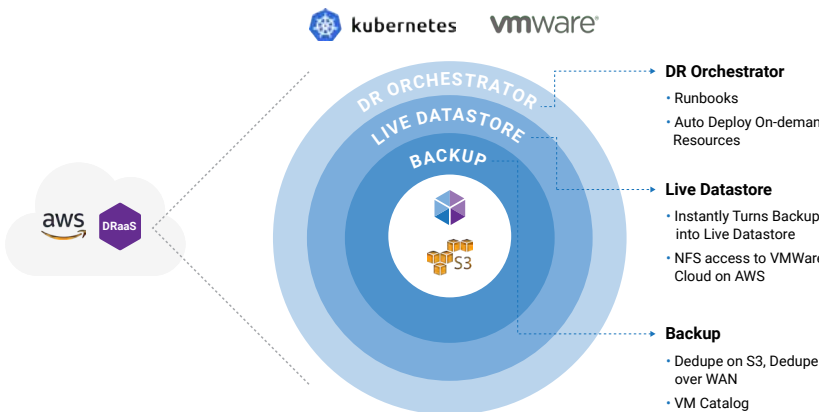


Figure 7: The architecture of of Datrium DRaaS

Datrium DRaaS allows you to use on-demand public cloud resources as a replication target for cost-effective backup to the cloud. Further, it enables fast VM and file recovery in the cloud to help you recover from disasters like a ransomware attack.

The Automatrix platform spreads the deduplication domain across all sites, which means that site-to-site replication—and, notably, failover and failback during DR—is very efficient because only the unique bits of data need to be copied. Transfers between sites are also forever incremental, which keeps the volume of intersite transfers as low as possible.

DR Runbook Orchestration

ControlShift is the multicloud DR orchestrator that completes the Automatrix vision. ControlShift is a SaaS application, so there's no infrastructure or software to manage on your end. It provides runbook automation and one-click failover.

With the Automatrix platform, your systems can run efficiently and securely whether it's on-premises or in the cloud. When you use Datrium DRaaS, you can easily orchestrate DR failovers and provide site-to-site mobility from data center to data center or from data center to public cloud (VMware Cloud on AWS).

It's not hyperbolic to call the on-demand DR capability a game changer for IT organizations. For testing or actual failover, Datrium DRaaS provisions a VMware software-defined data center (SDDC) in VMC. So, the only time infrastructure is provisioned and online is when you're using it. (Note, however, that a "pilot light" always-on option is available if the lowest possible RTO is necessary.)

Datrium DRaaS automatically destroys the VMC SDDC once testing or failback is complete, so you aren't paying for DR resources a moment longer than you need them. When a test or failover is not in progress, the only cost is for S3 storage for the backup data.

It's critical to business leaders and auditors to have confidence in your DR plan. Datrium DRaaS performs continuous compliance checks every 30 minutes to be sure you're ready to recover at a moment's notice.

DRaaS Connect

The power of Datrium DRaaS reaches outside the Automatrix platform as well. DRaaS Connect is downloadable, lightweight software for *any* vSphere infrastructure that enables customers to protect VMs just minutes after downloading.

- DRaaS Connect for VMware Cloud enables ControlShift to orchestrate failover from a VMware Cloud SDDC in one AWS availability zone (AZ) to another AZ
- DRaaS Connect for vSphere On Prem extends Datrium DRaaS capability to any vSphere on-premises infrastructure, including SANs, NAS, HCI, and DHCI

To have the option to replace a full DR site with an on-demand, consumption-based DR site orchestrated by Datrium DRaaS is a remarkable opportunity for IT organizations. It's hard to overstate just how transformational this shift could be.

Automatrix Platform Use Cases

The Automatrix platform is capable of taking on all kinds of enterprise IT challenges. Painting in broad strokes, there are three main buckets that Automatrix platform adoption fits into:

- DRaaS with VMware Cloud on AWS. Secondary DR sites are expensive insurance. By harnessing the power of the Automatrix platform, you can realize 10x lower disaster recovery costs by leveraging public cloud and only paying for what you use. And because VMC on AWS provides the exact same operational experience

as vSphere on-premises, IT administrators will be confident and effective in administering the DR site during a failover event.

- **DR Runbook Orchestration.** Since it's cloud-based and consumed as SaaS, ControlShift is easy to adopt. It provides failproof DR with continuous compliance checks and one-click failover and failback. ControlShift makes ransomware recovery to another site or to an on-demand DR site simple.
- **High-Performance Virtualized Applications.** As you learned, the underpinnings of the DVX platform allow it to deliver high performance and petabyte scale. DVX includes built-in backup and recovery, and always-on data efficiency.

All that great tech is in service of one primary goal: protecting your data and making sure it's available as soon as possible, even after a disaster event. In the next chapter, we look at the intertwined issues of data integrity, how the “bad guys” try to ruin it unless you pay up, and how Datrium keeps you safe and secure in the face of the bad guys' efforts.

CHAPTER 3

Livin' La Vida DR

Datrium's founders have deep experience in the issues of primary storage, backup and DR, VMs, and enterprise-scale solutions, with backgrounds from Data Domain, NetApp, VMware, and EMC.

Their combined experience provided insights into the root causes of why the best-laid DR plans go astray. What they found was that the technical approaches of the past fell short of a comprehensive solution. Addressing these problems more completely meant re-examining the core assumptions and approaches that made legacy methods vulnerable in the face of new realities.

What they determined is that today, DR can no longer just be a process (and a fragile, manual, and dangerous process at that). The reliance on people to plug the holes in the multi-tool tech stack required to do traditional DR doesn't scale for today's businesses.

That means it's time to bake DR into a product that simply works, throwing the dangerous aspects of traditional DR processes and piece-meal tech into the trash bin of past innovations.

Today's DR begins with data integrity, and includes thinking through the entire end-to-end lifecycle of data, from primary data, to backup, failover, failback, and the modern reality of VMs and containers vs. traditional filesystems.

Always-On Data Integrity

The core questions of DR are: "Do we have the right data to restore?" and "Can we restore quickly enough to meet our needs?"

Stepping away from thinking that DR is simply the raw ability to perform backup and recovery, Datrium recognized that data integrity is ultimately the foundation of any DR solution. If you can't rely on the data, the rest of the DR process doesn't really matter.

And having the right data doesn't help if it takes hours or days to restore.

It's the combination of data integrity and quick recovery that has been the elusive Holy Grail of DR. But like all technology innovations, it's

Always-On Data Integrity



Figure 8: Datrium Always-On Data Integrity

taken decades for multiple waves of technology innovations to finally come together to solve both problems.

So what is data integrity? And how far do you need to go to prove that you can trust the state of your data?

Plenty far, Datrium believes, which is why data integrity is obsessively addressed throughout the company's entire design and engineering process, as shown in **Figure 8**.

In many DR systems, there's just a single check on data integrity—confirming that the data written into the DR system is the same as the data coming from a source system. In most cases, that check is a checksum or hash of some kind, proprietary or built on publicly known algorithms. If the checksums of the source and backup don't match, the data can't be trusted, and it's time to try again.

In terms of risk management, however, that's a potential single point of failure, and could theoretically be hacked or corrupted to report success when the data had in fact been manipulated. If an attacker or malware can take over or bypass the checksum process, there's no way of knowing whether the system can be trusted.

Datrium, in contrast, treats every aspect of the data flowing through its system and back out to the source/recovery systems as though the data could be corrupted at any time. While many DR systems often throw in security and integrity at the last minute, Datrium has purposefully built its entire end-to-end design/engineering process to ensure that data integrity is the No. 1 concern.

Data Integrity vs. Ransomware

Traditional ransomware holds primary data hostage, which is dangerous enough. Today's ransomware not only targets your primary data, but backup copies, too.

The art of both security and resilience is to have layered defenses. It's a smart strategy, since everything breaks. Everything can be hacked.

Power fails. Fires blaze. Someone opens an email attachment they shouldn't have.

To protect against the newest waves of ransomware, for example, Datrium uses a layered encryption approach to secure all data in-flight and at-rest between the server host and the data nodes. Not only does encryption protect against drive thefts, but also against network sniffing in the data center. The encryption capability comes with an internal key manager, saving external key management costs and complexity. It's also FIPS 140-2 compliant.

In addition, backups are protected from ransomware by using immutable (i.e., unchangeable, read-only) snapshots that are stored in a separate namespace inaccessible to the network.

Further, backup data is isolated in a separate namespace so no outside user processes (such as a ransomware attacker) can access it. Groups of VMs are automatically protected based on organization-defined policies, and globally searchable from a single catalog.

Data and backups are automatically verified multiple times each day, so you know your data is always valid.

Backups can be mounted instantly without copying data, while in-use data is encrypted to reduce attack vectors. This approach is key to both speedy responses to ransomware attacks, and in keeping copies of infected data available in the backup for forensic analysis.

Walk-Through of Ransomware Detection and Recovery

When a customer finds that ransomware has started to encrypt their data, they can log into the Datrium interface and look at reports showing the total size and unique size of each group of VMs, as well as the individual snapshots for each VM.

Name	Taken timestamp	Origin DVX	Includes	Total size	Effective size	Destination DVX	Expiration	Uniquig size
SILVER Group - Half hourly - 2020-02-05T18:05 UTC	Feb-05 10:05 am (23m ago)	Prod-DVX	4 VMs (1 VSS)	72.9 GiB	415.3 GiB	DR-DVX	Feb-05 06:05 pm ...	3.2 GiB
SILVER Group - Half hourly - 2020-02-05T17:35 UTC	Feb-05 09:35 am (1h ago)	Prod-DVX	4 VMs (1 VSS)	72.9 GiB	415.3 GiB	DR-DVX	Feb-05 05:35 pm ...	2.5 GiB
SILVER Group - Half hourly - 2020-02-05T17:05 UTC	Feb-05 09:05 am (1h ago)	Prod-DVX	4 VMs (1 VSS)	72.7 GiB	415.3 GiB	DR-DVX	Feb-05 05:05 pm ...	2.9 GiB
SILVER Group - Half hourly - 2020-02-05T16:35 UTC	Feb-05 08:35 am (2h ago)	Prod-DVX	4 VMs (1 VSS)	72.7 GiB	415.3 GiB	DR-DVX	Feb-05 04:35 pm ...	3.0 GiB
SILVER Group - Half hourly - 2020-02-05T16:05 UTC	Feb-05 08:05 am (2h ago)	Prod-DVX	4 VMs (1 VSS)	72.6 GiB	415.3 GiB	DR-DVX	Feb-05 04:05 pm ...	5.3 GiB
SILVER Group - Half hourly - 2020-02-05T15:35 UTC	Feb-05 07:35 am (3h ago)	Prod-DVX	4 VMs (1 VSS)	72.5 GiB	415.3 GiB	DR-DVX	Feb-05 03:35 pm ...	2.5 GiB
SILVER Group - Half hourly - 2020-02-05T15:05 UTC	Feb-05 07:05 am (3h ago)	Prod-DVX	4 VMs (1 VSS)	72.5 GiB	415.3 GiB	DR-DVX	Feb-05 03:05 pm ...	2.9 GiB
SILVER Group - Half hourly - 2020-02-05T14:35 UTC	Feb-05 06:35 am (4h ago)	Prod-DVX	4 VMs (1 VSS)	72.9 GiB	415.3 GiB	DR-DVX	Feb-05 02:35 pm ...	3.0 GiB
SILVER Group - Half hourly - 2020-02-05T14:05 UTC	Feb-05 06:05 am (4h ago)	Prod-DVX	4 VMs (1 VSS)	73.0 GiB	415.3 GiB	DR-DVX	Feb-05 02:05 pm ...	2.9 GiB
SILVER Group - Half hourly - 2020-02-05T13:35 UTC	Feb-05 05:35 am (5h ago)	Prod-DVX	4 VMs (1 VSS)	72.9 GiB	415.3 GiB	DR-DVX	Feb-05 01:35 pm ...	1.1 GiB
SILVER Group - Half hourly - 2020-02-05T13:05 UTC	Feb-05 05:05 am (5h ago)	Prod-DVX	4 VMs (1 VSS)	72.9 GiB	415.3 GiB	DR-DVX	Feb-05 01:05 pm ...	1.1 GiB
SILVER Group - Half hourly - 2020-02-05T12:35 UTC	Feb-05 04:35 am (6h ago)	Prod-DVX	4 VMs (1 VSS)	72.8 GiB	415.3 GiB	DR-DVX	Feb-05 12:35 pm ...	5.3 GiB
SILVER Group - Half hourly - 2020-02-05T12:05 UTC	Feb-05 04:05 am (6h ago)	Prod-DVX	4 VMs (1 VSS)	72.8 GiB	415.3 GiB	DR-DVX	Feb-05 12:05 pm ...	6.1 GiB
SILVER Group - Half hourly - 2020-02-05T11:35 UTC	Feb-05 03:35 am (7h ago)	Prod-DVX	4 VMs (1 VSS)	72.6 GiB	415.3 GiB	DR-DVX	Feb-05 11:35 am ...	23.6 GiB
SILVER Group - Half hourly - 2020-02-05T11:05 UTC	Feb-05 03:05 am (7h ago)	Prod-DVX	4 VMs (1 VSS)	72.6 GiB	415.3 GiB	DR-DVX	Feb-05 11:05 am ...	80.2 GiB
SILVER Group - Half hourly - 2020-02-05T10:35 UTC	Feb-05 02:35 am (8h ago)	Prod-DVX	4 VMs (1 VSS)	72.6 GiB	415.3 GiB	DR-DVX	Feb-05 10:35 am ...	80.7 GiB
SILVER Group - Daily - 2020-02-05T08:01 UTC	Feb-05 12:01 am (10h ago)	Prod-DVX	4 VMs (1 VSS)	72.7 GiB	415.3 GiB	DR-DVX	Feb-11 12:01 am ...	3.6 GiB

Figure 9: When a virtual machine grows abnormally, it's often a sign that ransomware encryption is beginning

As the customer backtracks through the report, there's inevitably a time when a VM starts to grow abnormally, standing out from normal snapshots. This is a sign of ransomware encryption beginning (Figure 9).

Customers typically have new VM snapshots taken every 30 minutes (although they can take snapshots as often as every few minutes), so it's a matter of backtracking to at least the snapshot before the ransomware began its work, and restoring from there.

Datrium's approach to restoring a snapshot is different from most backup and recovery solutions. In this case, the VM snapshot is instantly made live, not by restoring a copy from a distant resource, but in simply mounting a clone of the snapshot and powering up the VM.

This approach is so radically different from traditional backup and recovery solutions that many customers are initially terrified of the idea of restoring hundreds of VMs at a time. Normally, it would take hours if not days to restore large snapshots with traditional backup and recovery solutions.

But Datrium’s approach is fundamentally different—it simply mounts the snapshot in place, leverages its two-layer storage architecture (a performance tier and capacity tier) to cache enough data to fire up the VM, and combines with streaming from disk to finish powering up the entire VM.

For customers who want to be even more cautious before restoring ransomware-impacted snapshots, Datrium supports a “Test Bubble”—an isolated sandbox—where the customer can verify that the snapshots in question are safe, before flipping them into production.

Given all this emphasis, how secure is Datrium’s platform? Well, Hinsdale Township High School District 86 found that the “all-in-one” DRaaS features of Datrium’s Automatrix platform reduced its security and availability risk profiles significantly, while minimizing the management overhead normally required in an extensively virtualized environment.

“Datrium is the ultimate solution we never thought was available in the marketplace. With the ControlShift application, it makes BaaS, DRaaS, and elastic replication a snap by making multiple copies of data available. Data is encrypted in transit and at rest, which makes it immutable against ransomware.

We feel uber confident in its ability to protect our data and recover quickly from an unfortunate situation. With other solutions, there would’ve been additional software, hardware, and complexity to accomplish what the Datrium solution provides. A single pane of glass to manage the hypervisor and Datrium is the icing on the cake.”

—Keith Bockwoldt, CIO, Hinsdale Township High School District 86

Compliance Audits? No Worries

Traditionally, the pain of fully vetting a DR plan with real-world testing is simply exhausting, not to mention potentially dangerous.

Datrium’s fundamental belief is that historically DR has been too complex, expensive, and unreliable. It’s too critical to leave purely to people caught up in the heat of the moment.

Just as automation and robots have enabled huge performance efficiencies and risk reduction for business workers, they can do the same for technical staff. The extra time and resources that can be reclaimed through the automation of core DR function makes it possible to redeploy technical staff to more proactive and value-add work. And with hundreds of tasks to coordinate for failover and failback, having a verified and automated DR runbook is critical.

The only way to be sure that the DR runbook works according to plan is to verify. The reality is that verifying a DR runbook and testing that the entire plan works is incredibly difficult with legacy approaches to DR. Automation is rare, and manual processes are the solution to cover gaps in automation, resulting in a fragile DR testing process.

The risk of infrequent testing is significant, and the dirty secret of DR. Infrequent testing creates significant risk, and the longer that testing is delayed, the longer it will take to fix any issues that may have accumulated since the last round of testing. Worst of all, infrequent testing all but guarantees that DR plans won’t work when actually needed.

Rather than verifying the plan as rarely as yearly, quarterly, or monthly, why not simply do it all the time? Modern DR like Datrium’s provides automated compliance testing, confirming a validated DR plan multiple times throughout the day. This ensures that any gaps in DR readiness will be minor and identified immediately.




Figure 10: Datrium ControlShift Plan Details dashboard

Datrium's DR processes clearly display the current overall state of managed systems in a single dashboard. When potential problems crop up, the interface allows you to dive directly into any impacted aspect of your DR plan that isn't in compliance.

Continuous compliance checks are run against every runbook every 30 minutes (**Figure 11**), to confirm that the fundamentals of the runbooks (sources, destinations, network paths, and more) are verified as still valid (**Figure 10**).

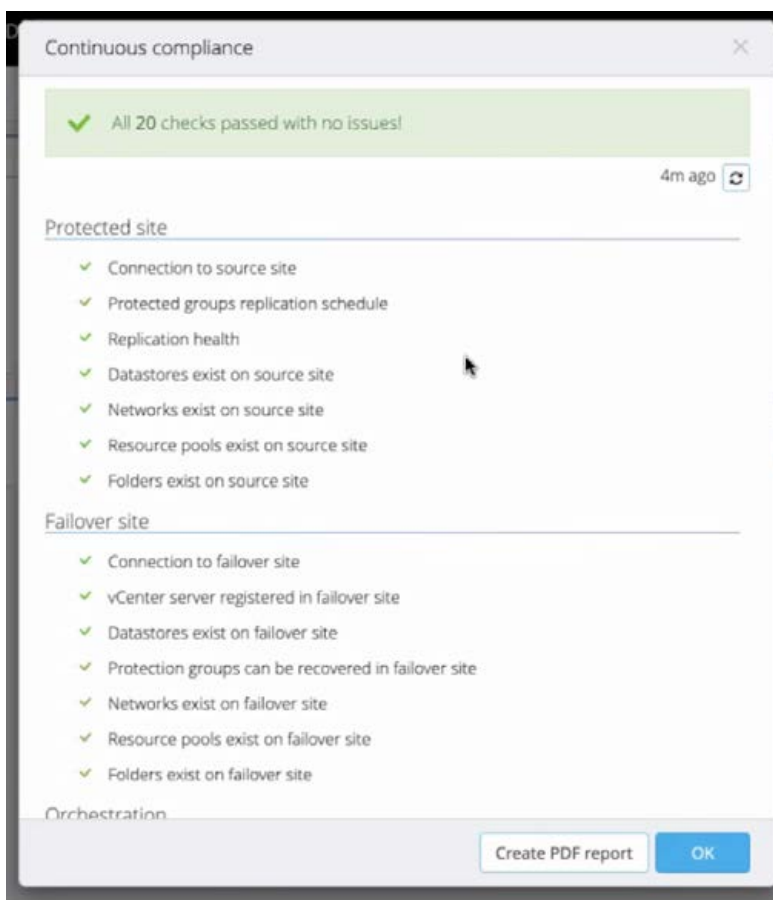


Figure 11: A Datrium ControlShift continuous compliance report

Additionally, backup snapshots are verified using fingerprinting of each snapshot. As an extra precaution, Datrium uses a blockchain-like approach to combine the collective fingerprints into a linked chain that becomes invalidated if any individual snapshot has been tampered with.

And backups are verified four times a day to make sure that the data itself can be trusted to match the known state at the backup’s creation. Finally, a test plan can be executed, exactly simulating a real DR event, automated, with confirmation that all parts of the plan have executed as expected.

Full visibility into the state of each stage of the plan is seen from within ControlShift (Datrium’s DRaaS Console), and can be reported out as a pre-formatted, human-readable report for auditors and management (**Figure 12**).

This may seem like a version of “trust but verify” taken to the extreme. But considering the stakes, it’s an approach that all but guarantees that your DR plans will be well executed, regardless of the cause of the disaster.

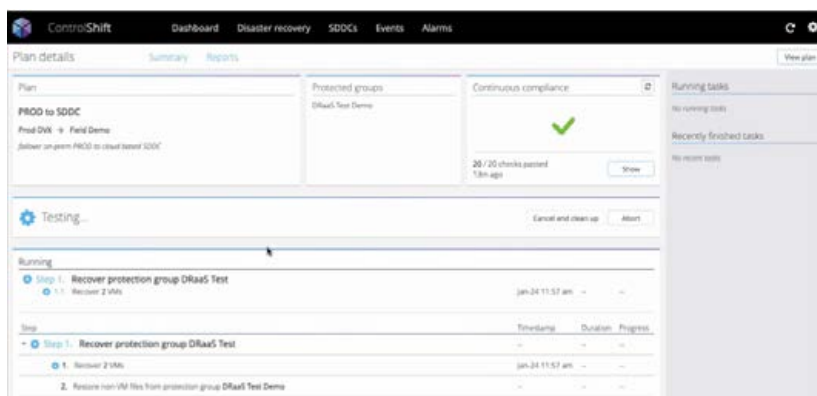


Figure 12: Executing a test plan in Datrium ControlShift

Traditional DR—Recovering from Natural Disasters

Speaking of disasters, natural disasters used to be the primary concern of DR—but today, natural disasters are simply one of many different threats a DR system must be ready for. They still happen, of course, but natural disasters are now less than half as likely to impact your organization as ransomware.

According to research from 2019, ransomware is the No. 1 cause of DR events—36.2% of known DR events, compared to 16.6% for natural disasters (see **Figure 13**).²

Those natural disasters still loom, though. Traditionally, the idea was to make sure that the secondary DR site was geographically distant so that the calamity wouldn't impact the recovery site.

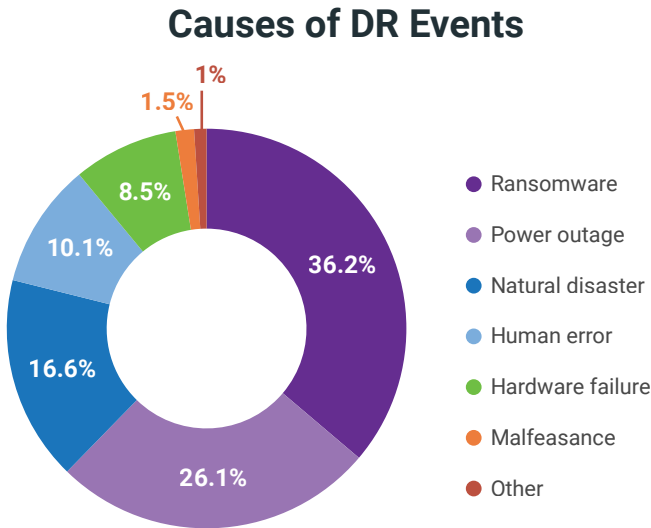


Figure 13: According to the report, “The State of Enterprise Data Resiliency and Disaster Recover 2019,” while ransomware is the No. 1 cause of DR events, power outages and natural disasters remain significant risks

² Source: “The State of Enterprise Data Resiliency and Disaster Recovery 2019”

But the age of on-demand everything has reduced the need for that paradigm. When you can instantly turn on thousands of compute cores to do large-scale analytics, or turn on render farms for the next superhero movie, why shouldn't you be able to orchestrate DR with the same ease?

Instead of “just-in-case” DR, it's now possible to use a “just-in-time” DR approach, using DRaaS. DRaaS takes advantage of everything that's been learned by cloud providers delivering on-demand capabilities, including the SaaS models popularized by Salesforce. Using a “pay as you go” model for DR by using AWS and VMware Cloud as your on-demand DR frees up significant time, money, and resources.

Needless to say, Datrium has you covered for “standard” DR events as well, with complete orchestration for on-prem to on-prem DR, or on-prem to cloud DR and back again. In the age of the cloud, it may be tempting to write off the need for on-prem to on-prem DR, but a small percentage of companies continue to use on-prem to on-prem DR for data that's deemed by corporate mandate to be too sensitive to store in the cloud.

For the great majority, though, the cloud's benefits are too great to ignore. One customer found that the built-in backup capabilities and ease of restore of Datrium's full virtualization-aware backup components removed the final barrier to eliminating tape backups, without the high cost of a traditional hot-site DR implementation:

“The bigger picture is that we now have a replicated data center in our DR site with Datrium DVX that is near hot—it is only 15 minutes behind our main data center. There's no need for us to go to backup tape to access last month's backup. We can recover data from 15 minutes ago and be up and running in a different data center.”

—**Source:** IDC interview

Plays Well with Others

A major reason for Datrium's customer successes is its integration with VMware. DRaaS Connect is downloadable, lightweight software for any vSphere infrastructure. DRaaS Connect enables customers to protect VMs just minutes after downloading, even without the use of Datrium's DHCI capabilities.

In addition to enabling cloud-based DR for VMs, DRaaS Connect for VMware Cloud enables DRaaS to orchestrate failover from a VMware Cloud SDDC in one AWS Availability Zone (AZ) to another AZ. DRaaS Connect for vSphere On Prem also extends Datrium DRaaS to any vSphere on-premises infrastructure, including SANs, NAS, HCI, and DHCI.

This flexible approach enables Datrium implementations to solve multiple problems at once with a single platform, including DR, storage management, and virtualization management.

Datrium's DR capabilities are built on top of a Disaggregated HCI (DHCI) foundation (however, it should be noted that DR doesn't rely on DHCI, although they're better together). DHCI is a next-generation hyperconverged infrastructure for VMware environments that runs apps with high performance and availability, using a two-layer design with separate performance and capacity storage tiers. This infrastructure allows customers to orchestrate DR, as well as manage and improve the performance of their virtualized infrastructure.

It may seem like Datrium's advantages sound too good to be true. But there are plenty of real-world examples of its power. One of those is Ultra Petroleum.

Ultra Petroleum

After trying multiple approaches to implementing DR, Ultra Petroleum found itself at a DR crossroads again. It had already tried the usual options for DR, including on-premises solutions, a colocation model (with

both backup and emergency DR sites in the same city), and a managed service model that physically separated the backup and DR sites.

There had to be a better way to reduce costs, give the company back more control over its own fate, and modernize the IT infrastructure to future-proof its DR capability.

Enter Datrium's DRaaS solution, built not just to solve the myriad challenges Ultra Petroleum was experiencing with DR, but to provide a future-proof foundation to build on:

"Datrium won on simplicity and savings. Built-in backup, built-in security, VMware compatibility, cost savings of 50%, future-ready ... What's not to like? We found a solution that finally fits our situation—today and tomorrow."

— Josh Rein, Network Manager, Ultra Petroleum

Perhaps most importantly, Ultra Petroleum could leverage the cost savings of both the technology and the personnel to operate it, and use those funds for future-focused business opportunities:

"This isn't just about making life easier for IT, it's about making IT more effective for our business. We were asked to cut our cost without cutting service levels, and we've come through. We manage dozens of systems, applications, and both IT-specific and business-specific technologies. With Datrium, we have more time, and we're saving money that can be re-invested in business priorities."

—Josh Rein, Network Manager, Ultra Petroleum

City Government

Public entities have also benefitted, like the IT director of a vendor that serviced a city government. The vendor needed to find ways to create an offsite DR site without the expense and complexity of a second physical data center, while also modernizing primary storage, and moving away from end-of-life implementations from a combination of three hardware and software suppliers. The core hurdle in this case

was that all of its data eggs (primary and secondary) were in several aging, expensive, and slow baskets.

The vendor went with a Datrium solution. The result was a radically simplified primary data center footprint, integrated management of primary storage, backup, security, mobility, and DR processes. The city also experienced significant improvements to backup Recovery Point Objectives (RPOs) and Recovery Time Objectives (RTOs).

Financially, the vendor was able to retire and reclaim the slated budget for yearly maintenance of three suppliers, reducing yearly costs by nearly \$150,000. It was also able to avoid the capital expenditures (CapEx) cost of a second physical data center location for DR purposes, and only incur operational expenditure (OpEx) costs when it wanted to test or put its DR plan into action during a DR event.

As you can see, Datrium's solution is out there working in the real world of DR. That's a good thing, because there are more ways to suffer a disaster than ever before. We'll take a look at some of them in the next chapter, as we survey the threat landscape that underscores the urgent need for valid, instant recovery, to keep your organization from experiencing a business-crippling meltdown.

CHAPTER 4

Mastering the Art of Recovering from Disasters and Ransomware

Remember the days when backup and DR were more afterthoughts than front-line concerns? Just do your backups, ship 'em to the mountain, and promptly forget about them until the next batch was shipped. In those days, disasters were relatively rare, and for most organizations, an extended outage wasn't crippling or, possibly, even fatal.

Those days are just a memory now. Today, the cost of downtime is huge, and the reputation harm can devastate even the largest businesses. Analyst firm IDC estimates that the typical cost of downtime (across small, medium, and large enterprises) is \$250,000 per hour, meaning it's something no company can afford to ignore.

In the past, DR has typically relied on a second DR site that mirrored the primary data center, making for a costly and complex solution. And even then, with all that cost and effort, recovery was fraught with problems—not the least of which was the too-often poor quality and unreliability of the backups themselves. Many an admin has tried to restore from disk or tape, only to find that the backups, since they hadn't been properly validated, were useless.

The Threats Are Piling Up

It's important to note that although the shift toward recovery was necessary, it doesn't mean that backups have become passé. They do have their place, and are a foundational component of the recovery process.

The shift, however, was a recognition of backups' big shortcoming in the modern era: they were an impediment to *fast* recovery. In this always-on world, you can't wait for hours or days for your data center to become functional again as you perform a lengthy rehydration of data from your backup systems.

Which brings us back to DR. Traditionally, the big worry was about natural disasters—fires, floods, hurricanes, tornadoes, and so on. But in recent years, disasters caused by humans have overtaken those fears, with ransomware leading the way. This scourge has grown to be Public Enemy No. 1 among enterprises. Striking without warning, ransomware can be more damaging than many traditional disaster scenarios. One significant reason is that it can be repeated over and over.

Humans can also bring on catastrophe by accident—opening that unknown email attachment or misconfiguring a server are just two of the most high-profile ways to bring down a data center. And let's not forget that hardware still fails, that software can still be buggy, and that zero-day attacks still happen regularly.

Yes, We WannaCry

Ransomware is growing, and the numbers are staggering. To take just one example, look at WannaCry's impact. According to PreciseSecurity.com, WannaCry was the most common crypto ransomware attack of 2019, accounting for 23.56% of all such assaults globally.

The report also estimates that about 230,000 computers were attacked during the year, at a total cost of \$4 billion.

Given that more companies than ever are deciding to pay the ransom, it's clear that ransomware will continue to wreak even more havoc in the future.

In other words, you can't wait anymore. You may be next.



It all adds up to a scary environment in which to operate. Smart organizations assume that a failure isn't a matter of "if," but "when," and plan accordingly.

That plan must now include speedy DR. To meet very aggressive Recovery Time Objectives (RTOs) and Recovery Point Objectives (RPOs), which are becoming more the norm than the exception, requires new thinking about DR, and applying modern, cutting-edge technology not available until recently. Yet, amazingly, most organizations continue to rely on the same approaches and technology as in the past.

Understanding the 3 Pillars of Backup and DR

The legacy approaches to backup are largely still with us, though some of the tools and methods have changed. But their reliance on manual methods rather than automation fails to address the biggest change in this area of IT—the necessity of minimal downtime.

Although a lot has changed, there are still three fundamental elements of backup and DR, also illustrated in **Figure 14**:

- **Data recovery** is where the value of backup is revealed. Putting backed-up data back into production and restarting systems and processes is especially difficult because of complexity. But it goes beyond just getting operations going again—the speed at which systems and applications are restored is key. If rehydration methods slow everything down, swift RTOs aren't possible.
- **Validation** involves not only testing plans for recovery, but also verifying that the right data comes back, and that nothing is lost or misplaced.
- **DR orchestration** is a means to automate the logistics of recovery and validate the results. When backup and DR functionality are converged, orchestration is the glue that holds it all together as a cohesive, efficient unit.

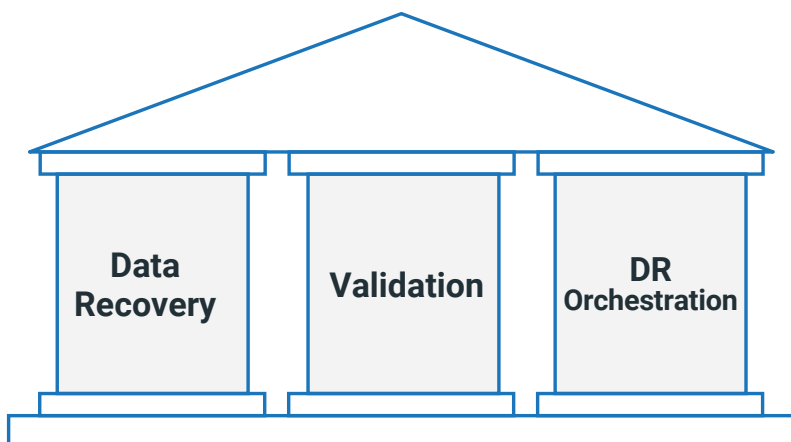


Figure 14: It can be helpful to think of data recovery, validation, and DR orchestration as the three essential pillars of backup

How Modern Disaster Recovery Is Better

The pillars may be the same as they've always been, but how backup and DR are best done today is very different. The limits of traditional DR are leading to reassessments and considerations of new, modern DR approaches.

Traditional DR was a largely manual, human-driven process that mostly ignored the potential of the cloud or looked at cloud capabilities through a legacy infrastructure lens. Simply throwing more “stuff” into the mix in an attempt to add functionality quickly created tool and infrastructure sprawl, which led some organizations to turn to expensive managed service providers (MSPs) to oversee their DR strategy. But considering the degree to which MSPs are struggling, and even going bankrupt, it's clear that the MSP approach to DR is proving unsustainable.

In contrast, modern DR sharply reduces the need for manual processes. When based on VMware Cloud on AWS, modern DR leverages the

scalability, availability, and most importantly, the on-demand nature of the cloud. Because of its inherent flexibility, it can provide multiple cost models, enabling organizations to strike the right balance between RTO, RPO, and cost.

Modern DR makes it far simpler to implement on-demand cloud DRaaS. And once in place, you can more easily manage and monitor costs and capabilities as part of your IT lifecycle.

With all the benefits and none of the high costs associated with duplicate data centers, new technology like Datrium DRaaS offers push-button DR that's easier to use while producing faster, more reliable results.

5 Ways To Get DR Satisfaction

Given that the requirements for DR have changed and a new approach is required, here are five ways to consider rethinking your approach to DR.

- **Don't use a DR site at all.** Instead, leverage the on-demand capacity of VMware Cloud on AWS to spin up temporary DR resources only in the event that you need them.
- **Maximize your budget's potential.** By recovering funds from unnecessarily expensive DR initiatives, you can tackle all sorts of new and exciting projects.
- **Say “No” to paying ransoms.** With the right converged backup and DR technology in place, there's no need to engage in a power struggle with ransomware developers. Instantly restore to just before the attack, plug up whatever hole they came in through, and move on about your business.
- **Focus on recovery.** The best backups in the world fall down when it comes to DR, unless the system is optimized for recovery speed. Many top-of-the-line backup products are fantastic at storing backups and archiving data; but they fail you at precisely the moment you need them most, which is during DR, by taking hours or days to rehydrate data and perform VM conversions in the cloud.

- **Believe in your DR plan.** With the right technology foundation, you can implement a DR plan free from the staleness and fragility characteristic of many legacy DR plans. With regular, verified recovery compliance, you can actually have faith in the integrity of your DR plan.

Datrium DRaaS with VMware Cloud on AWS

With those five considerations in mind, another one naturally springs to mind: Is there a single solution that addresses all of them?

There is, and it's Datrium DRaaS. A bold claim, no doubt, but Datrium DRaaS is a bold solution. We've already looked at Datrium DRaaS at length in this book, so in the interest of helping your retention of the information, let's walk through a quick snapshot (pun intended) of what it can do for your organization when the worst happens.

Through a simple UI, teams set backup policies and build automated DR runbooks. Tamperproof backups can be created every few minutes, every hour, every day—whatever makes sense for the business. Backups are deduplicated, compressed, encrypted, and stored in their native format in Amazon S3. Compliance checks run every 30 minutes to make sure the backup actually works.

You can be confident in the speed and validity of your DR plan with Datrium DRaaS. For testing or actual failover, an on-demand VMware SDDC is provisioned in VMware Cloud (VMC) on AWS. This means that the only time that infrastructure is provisioned and online is when you're actually using it. When a test or failover isn't in progress, the only cost is for S3 storage of all the snapshot data, which is globally deduplicated.

When disaster strikes, you simply initiate failover. The correct backup snapshots are instantly powered on via a live, cloud-native NFS data-store mounted by ESX hosts in that SDDC, resulting in instant RTO—a

capability all but impossible to achieve with any other backup system. Unlike legacy backup-only solutions, there's no time wasted waiting for backup data to be copied—also known as “rehydration”—into VMC before the VMs can be restarted.

Beyond rehydration, faster RTOs are possible because Datrium DRaaS doesn't involve VM conversion, also known as “refactoring.” Many

Defense in Depth

As you've seen, Datrium's design and engineering philosophy is to focus on delivering “always-on” integrity, rather than treating DR as a separate function. Although performance and many crucial capabilities are critical to any modern DR solution, data integrity is at the center, and it's the key to ransomware defense. Datrium employs a layered encryption approach so that data in transit and at rest is always protected. *Always.*



Encryption is done via an internal key manager that's FIPS 140-2 compliant, and immutable backup snapshots stored in an isolated, hidden namespace. Verification of data and backup occurs several times per day.

Ransomware can cause serious fallout in another way, as well: It can easily put you out of compliance, which can have its own devastating consequences.

Datrium understands this. With thousands of tasks to coordinate for failover and failback, Datrium DRaaS emphasizes the role of an automated DR runbook. Regular and automated compliance checking verifies the current state of DR components in a unified dashboard, so problems are easy to identify and address.

For instance, snapshots are verified using fingerprinting. A blockchain-like approach combines and links all the fingerprints to clarify whether any snapshot has been tampered with. A recovery plan is only validated if it contains known good backups.

backup solutions convert VMs to AWS' proprietary format during a failover, then convert back during failback. These conversions not only take a lot of time due to rehydration of the data, but things can go wrong during the conversion process, which can make for even longer RTOs.

Because Datrium has taken a VMware-centric approach, organizations can manage their on-prem and cloud environments the same way: New skills aren't required for administrators.

Take Control of Failback

Public cloud environments do come with a significant tax on data transfer, however, in the form of egress charges. Most public clouds won't charge you to put your data into their cloud, but instead charge to pull it out. This can really hurt a budget, especially when you're doing a failback operation.

This concern evaporates with Datrium DRaaS, in which only the unique, changed data is copied back to the primary site.

Be a Disaster Recovery Superhero

Modernization starts with understanding the entire end-to-end lifecycle of data, as well as up-to-date data protection needs. By now, you've seen how legacy DR is outmoded and insufficient. Emerging threats like ransomware require new thinking and improved technology. In addition, your infrastructure can no longer go down for hours or days at a time—you might as well put a "Closed for Business" sign on your front door.

Keep those doors wide open with newer, more comprehensive, and instantaneous DR that puts you back in control. Datrium DRaaS provides a modern experience that includes on-demand economics, failproof recovery, and the kinds of benefits only achievable with deep cloud integration. And it works with any VMware-centric primary storage for a powerful solution.

Throughout this Gorilla Guide, you've seen how much better your DR can be, by harnessing the awesome power of the cloud. You can, in fact, be a DR superhero in your organization, making sure that when a disaster strikes, full recovery is instantaneous. Your operations would be down for moments, completely unnoticed by the CEO and the rest of the world. The only one who would know you avoided a catastrophe would be you. Wouldn't that be great?

Just remember to use your powers for good!