The Key Benchmarks to Realizing Your **Cloud-Native Application Resiliency**

ABSTRACT

Cloud-Native Application Resiliency is Essential for Your Organization

If you're building your applications in containers, you're in good company. More organizations are building production-grade, net-new apps in Kubernetes than ever before — an estimated 70% by 2024 — up from just 10% in 2020, according to IDC[1].

Despite the many benefits of containers — speed, agility, and scalability, just to name a few — running containerized apps in production is a high stakes game. Your clients directly interact with them and are affected by outages, issues, and more. If your app goes down or gets hit by a ransomware attack, your reputation, customer base, and revenue are at risk.

In other words, the performance of your applications affects your business outcomes. So you don't want to take any chances.

That's why application resiliency matters so much. When your containerized apps are resilient, you can meet your compliance needs, hit service level agreements (SLAs), and quickly bounce back from outages.

So how can your Kubernetes apps — and the organization that depends on them — be more resilient? This eBook breaks down four benchmarks you need to be truly resilient and the right tools to help you boost resiliency.

[1] IDC FutureScape: Worldwide IT Industry 2022 Predictions.

Contents

Introduction .... Application Resiliency Benchmarks

Chapter 1 .... Recovery Time: The Key to Application Resiliency

Chapter 2 .... Migration Speed

Chapter 3 .... Disaster Recovery Time

Chapter 4 .... Recovery Cost

Chapter 5 .... How to Improve Your Metrics & Resiliency

Chapter 6 .... Conclusion

INTRODUCTION

In our unstable world, outages can happen at any time. And they often do — from digital disasters and ransomware attacks to unpredictable weather and human error. When the unexpected happens, your resiliency — or how quickly your organization can rebound when your data is compromised[2] — matters.

Being Truly Resilient Means:

Your applications run successfully in production without errors. ✓

You can easily change your infrastructure or move to a new environment if and when you need to. ✓

You can respin or recover your application into another cloud quickly and efficiently. ✓

So what's your application resiliency strategy?

You might rely on data protection methods like…

- Storage snapshots, provided by your storage vendor
- Traditional data protection technology
- Open source tools like Velero

But cloud-native environments are more complex with even more critical data to protect. That can be challenging if you're using free, built-in, or legacy tools like the above.

Limitations include:

Scalability: As your containerized environment grows, will your data protection scale with you? Or will you be limited by the number of clusters or containers? What about if you need to switch cloud providers?

Performance: How quickly can you reconstruct your application? Is all of your data and metadata accounted for?

Reliability: Can you trust that your solution is backing up all of your application data? Is there reliable support if you run into an issue?

Innovation: Will your vendor evolve with the technology so you're protected into the future?

Not sure if your tool is up to the task? We've got you covered.

Here's the key benchmark and supporting metrics that can help you determine how resilient your environment is and what you need to exponentially improve it.

[2] Spectra (2021). Data Resiliency: What You Need to Know. https://spectralogic.com/2021/06/10/data-resiliency-what-you-need-to-know-blog/

CHAPTER 1
RECOVERY TIME: THE KEY TO APPLICATION RESILIENCY

RECOVERY TIME: THE KEY TO APPLICATION RESILIENCY

It All Comes Down to Recovery Time

True application resiliency comes down to one big thing: time to recover. The faster you can recover, the less impact on your revenue, customers, and operations. And the more likely you are to hit your compliance requirements and SLAs.

That's why recovery time is so important and why all the resiliency benchmarks that follow affect it. So what recovery time metrics matter most and how is time the key to boosting your resilience?

Let's start with some definitions.

What is Recovery Time Objective (RTO)?

RTO is how much time it takes to restore your entire application — and all its resources — to the last known good state after an outage.

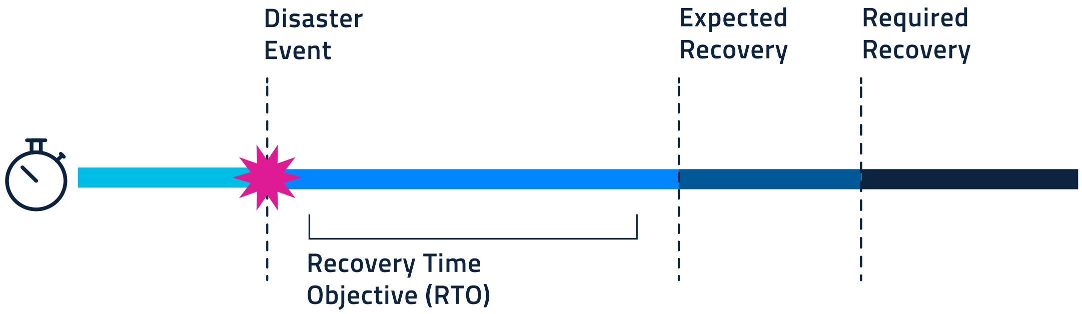

And it doesn't matter why that outage occurred, whether it's a disaster, ransomware attack, human error, or migrating to a new environment.

RTO factors in the amount of time that your app can be down before your business is significantly impacted, so it acts as a goal for your organization. It also takes into account things like:

People Resources: How long it takes to get the right people to execute the recovery process

Backup Access: How much time it takes to access backup data, on-prem or in the cloud

Transfer: How long it takes to transfer your data

System Restart: The time it takes to relaunch applications and load data into production

But RTO isn't the only time metric that matters. There are two components that contribute to your RTO — mean time to resolve (MTTR) and mean time between failure (MTBF). Here's why they're important.

Mean Time to Resolve (MTTR)

While RTO is your agreed-upon business objective for acceptable downtime, your mean time to resolve is about the actual time it takes you to recover after an outage. It consists of the average recovery time from each instance of downtime.

Your MTTR should help you set your RTO, and as it gets faster, you should adjust your RTO as well.

Mean Time Between Failure (MTBF)

RTO and MTTR focus on recovery from downtime; MTBF is about the opposite — how long an application stays available between outages.

In a Kubernetes or cloud-native system, application availability gets a major boost by enabling developers and site reliability engineers to create redundant services. MTBF is also handled at the application layer by enabling fault-tolerant topologies.

While resiliency takes into account RTO, MTTR, and MTBF, the main focus of this eBook is on recovery, so we'll dig into RTO and MTTR in more depth.

Benefits of Faster Time Metrics

You want to reduce time metrics as much as possible so you can get your applications up and running faster and with minimal impacts on your organization.

Faster recovery and restoration comes with big benefits for your business, including:

Your Customers

Limited downtime ensures that your customers have the best possible experience with your application.

Your Reputation

Increase trust with your brand by minimizing outages — and the media attention that comes with them.

Your People

Boost employee productivity through resilient infrastructure that keeps applications available and performing even during outages.

Your data protection solution should help you achieve faster recovery through various use cases. There's just one problem: Traditional methods have constraints that slow you down.

Legacy Backup Solutions Slow Down Your Recovery Time

If you're using a traditional method of data protection for your cloud-native apps, your recovery time will suffer. Here's why:

Manual Process

Many traditional methods rely only on snapshots as backup. But snapshots don't take into account the entirety of your cloud-native application, like your metadata. When disaster strikes, you're left stitching together snapshots manually, which adds time to your recovery process.

Increased Effort

Because your cloud-native apps are constantly changing, attempts to reconstruct them take not only time, but effort. It might not be possible to restore your entire application. And it'll take significant resources to try.

Data Loss

Most traditional data protection solutions focus solely on your data, without taking metadata — and the essential context it contains — into account. Without backing up all components of your app, you might lose critical data and metadata.

Even if you can access your metadata via the Git version of your control system, you'll be stuck manually correlating the data to the metadata.

So how do you improve your recovery time — and your resiliency?

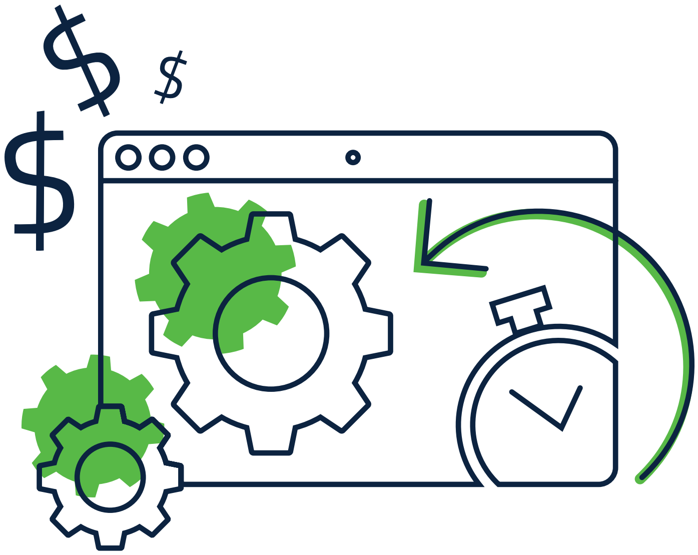

CHAPTER 1:
Recovery Time: The Key to Application Resiliency
CHAPTER 2
CHAPTER 3
CHAPTER 4
CHAPTER 5
CHAPTER 6

Cloud-Native Enterprise Solutions Give You Speed

Using an enterprise-grade, cloud-native platform can improve your recovery time by up to 80% — no small feat!

You can...

Restore your apps faster: When a disaster or outage hits, you can restore your application faster because your data is already backed up.

Create disaster recovery plans that fit your needs: Thanks to disaster recovery plans, you can automate your backups according to your needs. So you're covered whether you need to backup and restore across clouds or distributions, or in multiple namespace

Backup to any cloud, any storage, any time: Send point-in-time copies to the environment that works best for you, and avoid vendor lock in.

With a solution like this in your pocket, you can drastically boost your recovery time and your resiliency.

Now, let's take a look at a few specific situations, like disasters or migration, that cause outages and affect your recovery time. And how enterprise-grade solutions can increase your speed, no matter the situation.

CHAPTER 2
MIGRATION SPEED

MIGRATION SPEED

Next up on your journey to application resiliency is how quickly you can migrate your application to different environments.

Like we talked about last chapter, your recovery time, or the speed it takes you to rebound after an outage, is key to your resiliency — and happy customers, protected revenue, and a good reputation. And efficient migration is a part of that.

Why Migration Matters

In a multi-cloud world, you want to be able to move your application from one environment to another easily and with limited disruption. Here's why.

Stay Agile

The ability to migrate your application between distributions or across clouds gives you flexibility. You can work with a different vendor or backup to a secondary cloud without worrying about the downtime that comes with losing your primary cloud environment.

Save Money & Improve your Total Cost of Ownership (TCO)

Test and validate your application and associated components within a different cloud or distribution that fits your budget or performance requirements. As a result, you avoid vendor lock-in, save money, and improve your TCO.

The faster and more effectively you can migrate your applications, the better your recovery time. And the more resilient you are.

So what's your migration strategy?

Don’t Rely on Traditional Backup & Recovery

If you’re depending on traditional methods of backup and recovery to make migration easier, you might be in trouble.

Traditional Methods...

Lock you into a vendor: If your backup and recovery method is tied to your cloud or storage provider, you’ll spend more money and time making a switch or backing up a secondary copy elsewhere.

Take more time and resources to implement: You’ll need more people and more computing resources to migrate, adding time to the process.

Reduce your security: If your backup is vendor dependent, you’re at risk if that vendor goes down.

To be truly resilient, you need a solution that makes migration seamless — and fast — in any scenario.

Migrate Your Entire Application in 5 Min or Less

Using an enterprise, cloud-native solution can significantly increase your migration speed. We’re talking an entire app migration in 5 min or less. Reduce your downtime and satisfy your customers.

And the benefits don’t stop there.

1. Meet your SLAs and compliance requirements with minimal downtime.
2. Upgrade your environment with limited interruptions or errors.
3. Choose a storage and cloud agnostic platform to avoid vendor lock-in.
4. Increase security by backing up to a cloud or secondary cloud for guaranteed data protection.
5. Enjoy test/dev benefits by pushing real production data into your environments and trouble-shoot issues before they get to your customer.

Don’t let migration speed slow down your resiliency. Instead, find a solution that makes it easy to rebound after an unexpected outage.

CHAPTER 3
DISASTER RECOVERY SPEED

DISASTER RECOVERY SPEED

Another critical component of recovery time and resilience is disaster recovery (DR). Not just for your individual applications, but the entirety of applications that need to be available for your business to operate.

In other words, your minimum viable business (MVB), which includes a set of applications and sites that must be protected and recovered in case of a disaster.

To resume operations and avoid costly outages, you need to recover your minimum viable business quickly. So once again, speed matters. The faster you can recover your MVB, the more resilient you are.

Natural & Digital Disasters Damage Your Organization

A disaster recovery strategy isn't new — it's a must-have, especially as natural and digital disasters like ransomware happen more frequently than ever. And the consequences can be significant.

One of the Most Common Causes of Outages: Cybersecurity Events

The average downtime for a ransomware attack: **22 days**[3]

The average cost per minute of downtime: **$5,600**[4]

91% of organizations reported that just one hour of downtime cost more than $300,000[5]

To avoid critical impacts on your operations, revenue, productivity, reputation, and customers, you need a comprehensive DR strategy that gets your containers back up and running fast. And traditional methods of data protection aren't up to the challenge.

[3] Coveware. (2021, October 21). Ransomware attackers down shift to 'Mid-Game' Hunting in Q3 2021. https://www.coveware.com/blog/2021/10/20/ransomware-attacks-continue-as-pressure-mounts

[4] Gartner (2014, July 6). The Cost of Downtime. https://blogs.gartner.com/andrew-lerner/2014/07/16/the-cost-of-downtime/

[5] ITIC (2021). Hourly Cost of Downtime Survey.

Legacy Disaster Recovery Limitations for the Cloud

Legacy data protection solutions aren't built for the cloud, which can negatively impact your ability to recover from disasters efficiently — or at all.

Here are just some of the downsides:

Time to restore your site: Depending on the rate of change and how big your application is, it could take hours — or even days — to restore your site.

Service downtime: The longer your site is down, the longer your customers are without service.

Data loss: If you haven't backed up all of your applications, including their data and metadata, you might not be able to recover fully.

Complexity: Trying to force a legacy solution onto a containerized environment isn't easy. Your tool might not be multi-tenant or self-service, leaving you and your team to stitch the pieces together.

To make your cloud-native apps truly resilient, your disaster recovery solution needs to be cloud native, too.

How to Recover Your Entire Site From Disaster Faster

An enterprise-grade solution can exponentially speed up your recovery time. **In fact, you could recover your entire site from a disaster in less than 30 minutes.**

Here's what to look for:

Click-driven DR Plans: Set up disaster recovery policies for each individual application using built-in, click-driven workflows. Or create a custom policy based on your needs

Continuous backup: Look for a solution that works around the clock to limit your data loss.

Storage and cloud agnostic: Back up point-in-time copies to any storage or cloud, so you're not dependent on any vendor.

Multi-tenant, agentless, and self-service: Save time, money, and resources with the ability to restore your entire environment autonomously and efficiently.

Because your data is being backed up automatically and according to your personalized plan, restoration is easy and fast. That means limited data loss, improved recovery time, and a resilient business that can bounce back from whatever the world throws your way.

CHAPTER 4
RECOVERY COST

CHAPTER 1
CHAPTER 2
CHAPTER 3
CHAPTER 4: Recovery Cost
CHAPTER 5
CHAPTER 6

RECOVERY COST

Whether from ransomware attacks or human error, outages are expensive. The revenue you lose plus the cost it takes to get you back online adds up fast.

Like the previous metrics, cost and recovery time are connected. The longer your recovery time, the higher the cost of your outage.

An efficient recovery plan can reduce that cost (and time) and boost your data resiliency, too.

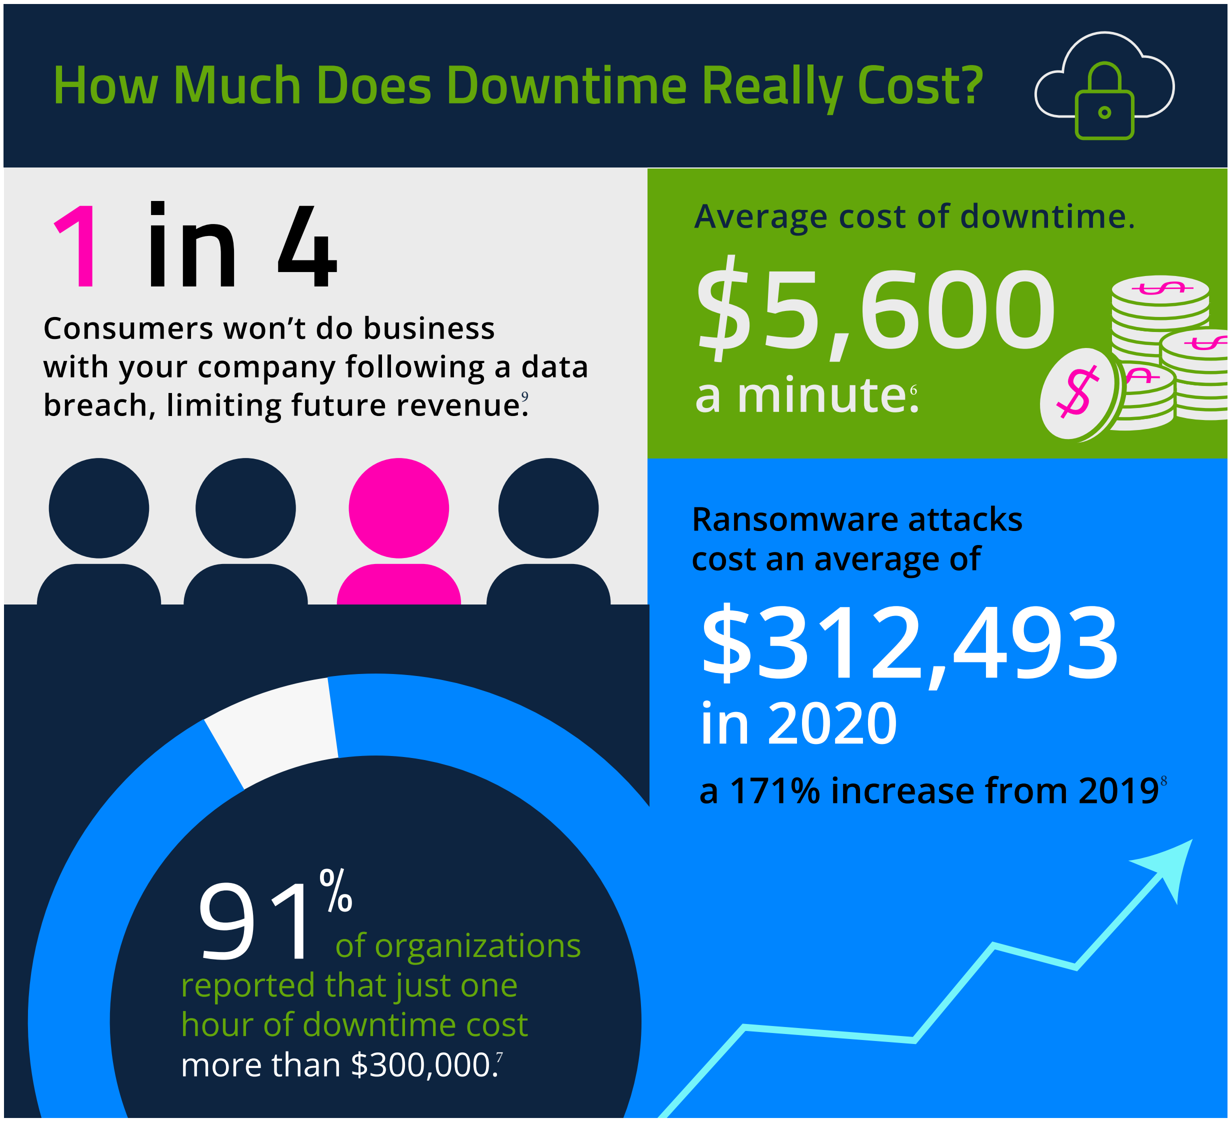

How Much Does Downtime Really Cost?

1 in 4

Consumers won't do business with your company following a data breach, limiting future revenue.[9]

Average cost of downtime.

$5,600 a minute.[6]

Ransomware attacks cost an average of

$312,493 in 2020

a 171% increase from 2019[8]

91% of organizations reported that just one hour of downtime cost more than $300,000.[7]

As you can see, unplanned outages cost your organization big time. Your data protection solution should help you offset these. But not all solutions are created equal, especially if you're running production-grade apps in containers.

[6] Gartner (2014, July 6). The Cost of Downtime. https://blogs.gartner.com/andrew-lerner/2014/07/16/the-cost-of-downtime/
[7] ITIC (2021). Hourly Cost of Downtime Survey.
[8] Palo Alto Networks. 2021 Unit 42 Ransomware Threat Report
[9] Security.org (2019, May 11). Public Awareness of Major Data Breaches. https://www.security.org/resources/data-breach-awareness/

CHAPTER 1
CHAPTER 2
CHAPTER 3
CHAPTER 4:
Recovery Cost
CHAPTER 5
CHAPTER 6

Hidden Costs of Legacy Data Protection

Here are some ways your data protection might be costing you money, instead of saving it.

More Downtime = More Expensive: Legacy providers can't keep up with the complexity of containerized environments, leaving you with the time-consuming task of manipulating raw data. And every minute costs you.

Manual Processes Take More Resources: You'll need more developers or staff to take on manual tasks, increasing costs.

Customer Loss: During outages, you lose the opportunity to bring in new customers and risk your current ones, too.

Ransomware Protection Not Included: Many traditional data protection methods don't include coverage for ransomware. So you're stuck paying for a separate solution on top of the one you already have.

Slash Your Costs & Increase ROI with Enterprise-Grade Data Protection

Using the right solution can 10x your ROI or more. But you have to know what to look for and where free tools fall short.

Avoid vendor lock-in: Tools that allow you to easily backup and recover from any storage or cloud provider give you flexibility. You can choose to write your backup copies to a cheaper secondary cloud.

Quick recovery based on policies you set: Because time and cost are so closely related, you want a solution that allows you to restore your most recent version as fast as possible. Automated backup policies make that easier.

Built-in ransomware protection: Look for a tool that helps you solve for all kinds of outages — ransomware attacks included — so you don't have to waste money on separate tools.

Easy user interface: A tool is only useful if you use it. Make sure your solution has a simple, user-friendly interface to reduce time and cost of training others to use it. Or building your own.

Full application coverage: Your solution should take all of the components of your application into account including metadata. That way, restoring to your last known good state is easy and doesn't involve manual — and costly —data manipulation.

If your enterprise-grade solution checks these boxes, the cost of downtime decreases, which means your application resiliency increases. Win!

CHAPTER 5
HOW TO INCREASE YOUR RESILIENCE

HOW TO INCREASE YOUR RESILIENCE

Cloud-native application resiliency depends on one big thing — recovery time. Reduce that, and your resiliency goes way up. But how?

Start with your tool.

Free tooling, traditional methods, and built-in protection via the cloud are simply not enough to make your data and organization resilient. You need an enterprise-grade solution that improves your recovery time and scales with you, no matter how big your cluster becomes or how many environments you need to protect.

These features and functionality will provide full coverage for your data:

Application Centric: Backup and restore all components of your applications including data, metadata, and all Kubernetes objects (namespaces, labels, Helms, and Operators).

Free cloud-native tools like Velero only backup by namespaces
and labels, which doesn't give you the granularity you need. You might backup more data than you need or miss something critical.

Scalable: Expand your applications easily with protection that scales with you. Add more clusters or an additional cloud? With an enterprise-grade solution, you're still covered. Need to recover one application — or thousands? Your recovery time should stay the same, no matter how many applications you have.

The performance of your free tool, however, decreases as you grow.

Simple management: Each containerized application creates a lot of data. You need an easy way to manage it all — multi-tenant and across clouds and clusters. Look for a solution that doesn't require you to learn a separate CLI. Instead, everything should be integrated into your Kubernetes API server.

Then, you can easily restore your data in your multi-cloud environment and track your backups, too.

Other tools don't offer this transparency or tracking.

Easy restoration: Restore your entire application, no matter where its components live, with easy-to-manage policies. Get granular control of your restores and make sure your metadata goes along for the ride.

	Enterprise-Grade Solutions	Free Cloud-Native Tools	Screenshot-Based Protection*
Application- centric	✓	●	●
Scalable	✓	●	●
Simple Management	✓	●	●
Easy Restoration	✓	●	●

*Includes snapshots from your storage or cloud provider

CHAPTER 6
CONCLUSION

CONCLUSION

For your cloud-native applications to be truly resilient, you need to be able to recover from an outage quickly and with minimal impact on your data, revenue, customers, and reputation.

And that comes down to improving your recovery time with a robust, cloud-native solution that scales with you and makes it easy to rebound, no matter the reason.

That's exactly what you'll find in TrilioVault for Kubernetes — application-centric backup and recovery that grows with you. Trust the protection of your cloud-native applications — and the organization they support — with Trilio.

Ready to Realize Your Application Resiliency Now?
Talk to a Cloud-Native Expert at Trilio Today.

Request Demo

trilio.io/request-demo

TRILIO.IO