Forget RTO, Redefine Recovery:

The impact of real-time data recovery on cyber resilience

brigantia + redstor

Forget RTO, Redefine Recovery:

The true cost of lost data and downtime can easily be thousands of pounds per minute. Recovery time will make or break your ability to mitigate the impact of lost or corrupted data.

That's why RTO (Recovery Time Objective) has been such a critical metric for data recovery and restore. This is the amount of acceptable downtime that can occur before business operations are significantly impacted.

But what if you could functionally forget about RTO altogether?

Many IT service providers and technology vendors focus on reducing downtime by minimising the time it takes to complete a full or partial system restore. For example, applying deduplication or compression to reduce file sizes, replicating files to a secondary location, or using high-performance storage arrays.

Unless you're willing to absorb the costs of fully duplicating IT infrastructure, none of these tactics entirely remove the lag between data loss and data access. No matter how small, any RTO objective still leaves end users struggling with downtime that damages their business.

At Redstor, we've pioneered a different approach. Rather than focusing on how long it takes to fully restore a system, we asked ourselves how quickly we could provide end users with access to their data while that restore takes place.

The solution we've landed on is simple, at least in principle — allow end users to instantly stream data from the cloud. This removes RTO as a practical concern for end users by providing functional access to all archived data on-demand.

Here, we'll look at how InstantData functions and what instant data access means for IT service providers and end users. Fundamentally, we'll explore how this change in perspective delivers far greater cyber resilience in an environment that's never presented more cyber risks. Let's get started.

Changing data recovery requirements

The business environment and technology landscape have undergone significant change in recent years. Companies are more reliant on data and cloud-based systems than ever before, resulting in trends that include:

Increased data volumes:

Businesses are generating and storing more data than ever. According to IDC, the global datasphere is expected to grow to 175 zettabytes by 2025. This exponential data growth necessitates more efficient and scalable data recovery solutions.

Cloud adoption

The migration to cloud services has changed the way data is stored and accessed. While cloud solutions offer flexibility and scalability, they also introduce new challenges for data recovery, such as data fragmentation across multiple platforms. It's critical for IT service providers to integrate SaaS-app, cloud and on-prem data restore capabilities.

Remote work

Dispersed workforces add to data management complexity. Decentralised data require robust data recovery solutions that can support a distributed environment and ensure seamless access to data regardless of location.

Managing these changes is made harder by the parallel expansion of cyber threats. A recent report by Cybersecurity Ventures estimates that cybercrime will cost the world $10.5 trillion annually by 2025. These threats have a direct impact on data recovery requirements:

Ransomware attacks

Ransomware has become one of the most prevalent and damaging cyber threats. Quick and reliable data recovery is essential to mitigate the impact of such attacks.

Data breaches

Unauthorised access to sensitive data can lead to significant financial and reputational damage. Effective data recovery solutions must ensure data integrity and protect against unauthorised modifications.

Malware and phishing

These common threats can corrupt or delete data, underscoring the need for reliable backup and recovery mechanisms to restore data to its original state.

Every cloud application and on-prem storage system presents one more avenue for attack. Larger volumes of data make it important to identify and restore only corrupted files or deleted data to limit the costs and time required for a full system restore.

It's also important to remember that critical data loss can occur accidentally, for example, because an employee deleted the wrong file.

Businesses want to navigate these risks with confidence. They turn to IT service providers to keep them safe and operating. IT service providers need simple ways to manage and protect this growing web of applications and volumes of data. They then need to deliver smooth, fast, and efficient recovery in every scenario to keep end users happy.

Data recovery as part of cyber resilience planning

“Cyber resilience” answers many of these threats and requirements. The National Institute of Standards and Technology (NIST) defines cyber resilience as “the ability to anticipate, withstand, recover from, and adapt to adverse conditions, stresses, attacks, or compromises on systems that use or are enabled by cyber resources.”

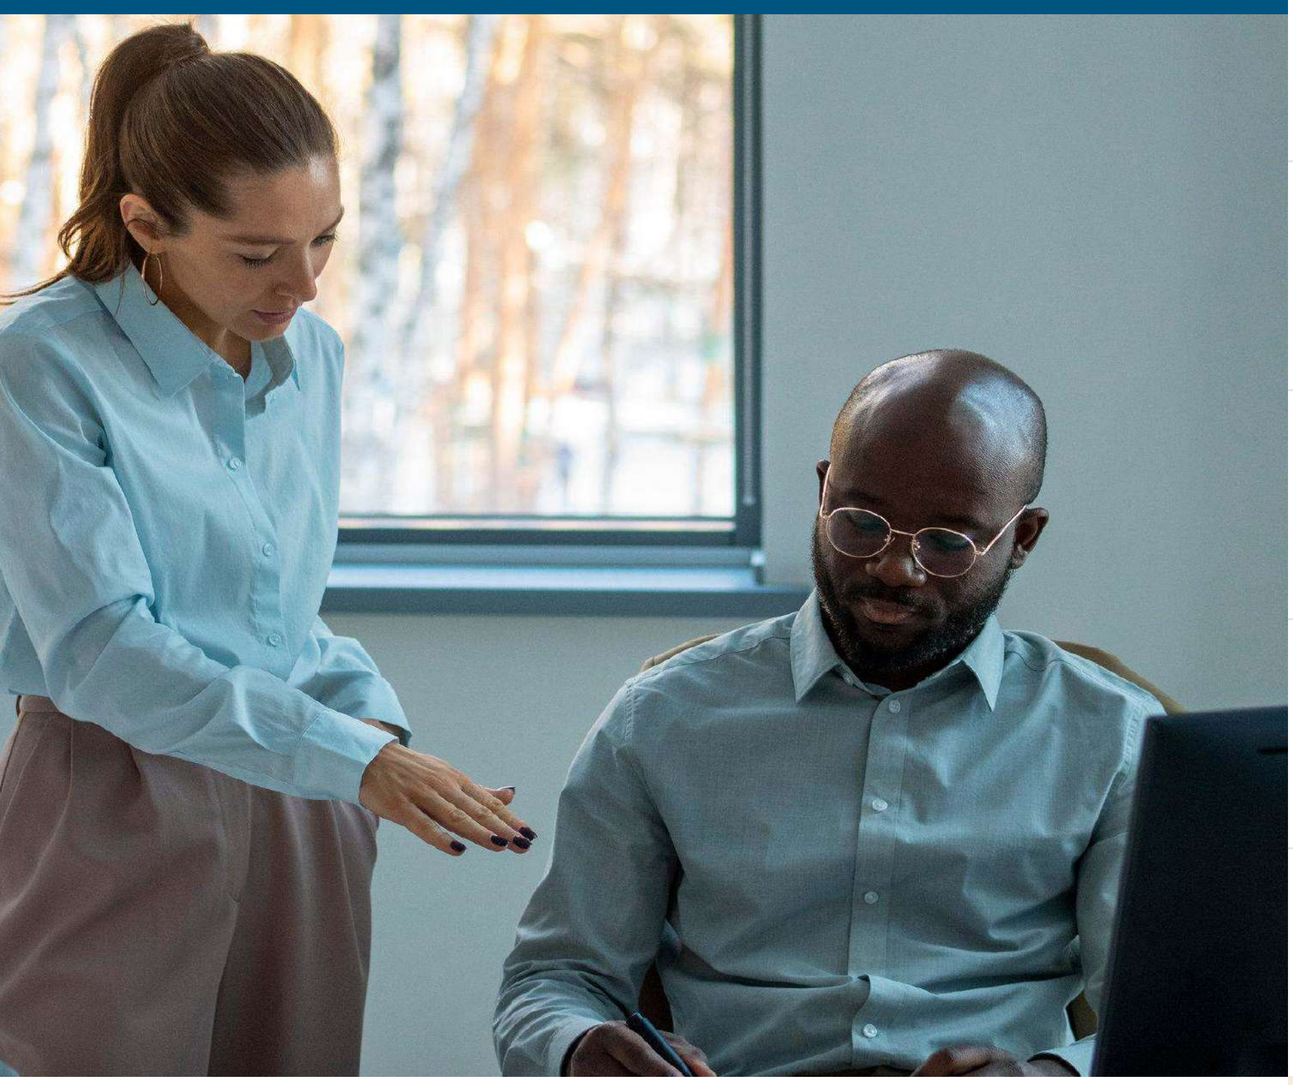

At least four of these seven elements are relevant to backup and recovery: know who your recovery service is delivered by; understand the procedure for data recovery; identify which assets are affected; and know which infrastructure and alternative sites to recover to.

However, cyber resilience fundamentally rests on effective data recovery at almost every level. That’s because a cyber resilience strategy aims to:

Minimise downtime

Your ability to continue operating following an attack is a central test of your resilience.

Ensure data integrity

Cyber resilience involves maintaining the integrity of data during and after a cyber incident. This includes ensuring that recovered data is accurate, complete, and uncorrupted.

Adapt to new threats

Cyber threats are constantly evolving. A resilient approach relies on flexible data recovery to shore up risk against unforeseen threats.

Ultimately, cyber resilience requires a data recovery system that can reliably provide access to clean data as quickly as possible. Your customers expect this. The question is whether or not your existing data recovery solutions can provide access to data quickly and reliably enough to stay competitive.

A comprehensive cyber resilience plan should contain seven key elements:

Accurate contact information:

How and where relevant participants can be reached.

Roles and responsibilities

Chain of command, who is doing what, where, and when.

Supply chain:

Vendors, insurers, legal.

Step-by-step guides:

Scripts/procedures that direct recovery actions.

Asset inventories

Equipment details, including network diagrams.

Infrastructure:

Alternative sites, telecommunications, and power requirements.

Messaging

What to tell customers and creditors and when and how to do so.

The traditional role of RTO

Recovery Time Objective (RTO) is the amount of downtime that can elapse after data loss before a business experiences intolerable consequences. This calculation is based on:

Productivity losses:

Affecting individuals or groups and their ability to complete their job tasks.

Business losses:

Affecting transactions and customer relationships.

In a traditional data recovery context, RTO is the critical metric that impacts these losses. This is because no work can be done using impacted data until a full recovery occurs. However, while RTO considers the cost of downtime, it simply seeks to minimise it, rather than avoid its impact altogether.

Delivering cyber resilience with InstantData™

Data recovery built for cyber resilience doesn't have time for RTO. You still need to set an RTO for the final restore point. However, end users cannot be held up by that process. This critical change is required to deliver real cyber resilience solutions.

InstantData is a feature of Redstor that operates on these principles. It delivers:

Immediate data access:

Users can begin accessing their data instantly, without waiting for the complete restoration process to finish. Critically, users can work on those files and save changes as they go. This data is then reconciled with the eventual full recovery, ensuring seamless business continuity.

Flexible data access:

Instead of retrieving entire datasets upfront, InstantData streams data directly from the cloud as needed. This approach reduces the burden on local infrastructure and accelerates the availability of critical information.

Together, these capabilities allow IT service providers to navigate the huge volumes of data generated by their customers, while still delivering instant access to data when needed.

How InstantData™ works

To operate, this technology leverages advanced cloud computing and data management techniques. Key components include:

Cloud integration:

Tight integration with cloud storage solutions allows seamless data streaming. This integration ensures that data is always available and can be accessed from anywhere, providing flexibility and reliability.

Smart indexing:

Intelligent indexing to manage and locate data quickly allows for rapid identification and retrieval of specific files or datasets, minimising latency and enhancing user experience

Stub files:

InstantData creates local stub files, which are initially empty placeholders for the actual data. When a stub file is accessed, the relevant data is retrieved and rehydrated, providing immediate access to the requested information. This ensures that critical data is available immediately, while less important data is restored in the background.

Efficient data compression:

Advanced compression algorithms that reduce the size of data packets without compromising integrity enable faster data transmission and access.

Using InstantData is then very simple. Depending on the system settings, end users and IT service providers can start recovery, which triggers an automated four-step process:

Step 1

Initiate recovery

When a recovery is initiated, InstantData maps the recovery data locally.

Step 2

Access data

Users can immediately access their data through stub files. These files are placeholders that retrieve the actual data from the cloud when accessed.

Step 3

Stream and rehydrate

As users access the stub files, InstantData streams and rehydrates the data on demand, ensuring immediate availability of critical information.

Step 4

Background recovery

While users work with their critical data, the full recovery process continues in the background, ensuring all data is eventually restored. While users work with their critical data, the full recovery process continues in the background, ensuring all data is eventually restored, saving the file over the top of the recovered files once it is complete.

This capability is integrated within the Redstor platform, which covers SaaS-apps, cloud storage, and on-prem servers within a single, multi-tenancy platform. This makes it simple for IT service providers to extend instant data protection across the expanding tech stacks of their customers.

The outcome is instant access to data when it's needed. InstantData gives your customers (and you) that recovery time back, because it only needs to deliver the changes that have happened since the last recovery point, and enables end users to work on files as the recovery takes place.

Recovery should be a competitive differentiation

Until it’s needed, data recovery can be an afterthought.
When it’s needed, little else matters.

Customers care about cybersecurity and cyber resilience, but are often more interested in preventing threats than recovering from them. Arguably, this is because traditional recovery methods are so time-consuming and costly.

A cyber resilience mindset requires a realistic perspective on risks. That doesn’t mean ignoring threat prevention. However, breaches will occur and not all threats can be identified. Important steps should be taken to mitigate, or even remove, the impact of these threats when they occur.

If demonstrated correctly, data restore capabilities can play a central part in selling cybersecurity and cyber resilience solutions. Redstor can help you do that with two additional capabilities:

Instant trials:

By operating in the cloud, it’s possible to get set up and trial Redstor in as little as 60 seconds. This means that your sales teams can show potential customers how the system works with live data, and demonstrate how quickly they will be able to access that data when needed.

AI-powered malware protection:

Instant data access only matters if you’re restoring clean files. Redstor applies an additional layer of cyber-protection that scans backups to stamp out malware and ensure instant recovery of usable data.

Regardless of whether your customers have on-prem infrastructure, cloud environments, SaaS applications or a combination of all of the above, Redstor offers protection of all estates. This is done within a single application, meaning you can protect your customers’ data no matter where they are on their data journey.

Not only does this mean Redstor can recover an entire estate in a lot less time than any competitors, it opens up expanding use cases by giving you and your customers instant access to any restore points in the system.

For MSPs helping customers recover from small mishaps, such as a deleted critical email, all the way through to restoration of a large file archive due to corruption or malware, the time invested is much the same. Traditionally the only option has been to roll back to the last full backup and wait out the recovery process causing business downtime.

By moving the data recovery process from hours or minutes to instant, IT service providers can differentiate an otherwise ‘me too’ offering and make better margins off recovery.

For customers, the recovery process becomes invisible. On its own, this is a differentiation. However, this shift will transform their cyber resilience planning, and allow you to deliver peace of mind.

If you want your customers to be able to recover data like it never happened, get in touch and experience it firsthand.

brigantia + redstor™