



Can You Tell if Your Data Recovery is Predictable?

Whether you've been hit with a ransomware attack or an internal mishap (malicious or human error), when disaster strikes, it's critical to quickly recover from data loss. The question is: if a disaster were to happen today, could you predictably recover your backup data - when and where you need it - without compromise? The ability to ensure predictable recovery gives you confidence in meeting your SLAs and in the resiliency of your organization. To assess your readiness for achieving predictable recovery, ask yourself these questions based on a framework of core recovery attributes.



1 Do you have 100% backup success rate?

Failed backup jobs and long backup windows that constantly bleed into production time can force compromise. This is because it becomes more convenient to skip protection of some mission-critical data. However, **data that is not protected cannot be recovered**. If you are getting anything less than 100% backup success rate, how can you meet your business SLAs? Achieving predictable recovery requires a backup solution that protects all mission-critical data without impacting normal production time or causing backup failures.



2 Can you do a global, Google-like search for any VM, file, or object?

Exponential data growth across workloads and geographical locations, combined with legacy siloed backup architecture, makes locating data challenging. In order to predictably meet recovery SLAs, you should be able to perform a simple, Google-like search for any VM, file, or object across all workloads and locations, all from a single UI. A modern backup solution will allow you to easily locate the data that you need across locations, and puts you in the desired workflow for ensuring predictable recovery.

Businesses Are Not Recovery Ready

- Ransomware attacks on enterprises are up 500% this year at a cost of \$11.5 billion
- Only 11% of IT decisions makers can restore data & apps within 3 days of an attack
- Only 25% of IT decisions makers are able to recover 75% to 100% of their data

"Organizations are reevaluating their current backup practices, as well as their recovery processes, with an eye to minimizing the impact an attack could have on their business."

- Forrester

**Forrester study of more than 300 IT infrastructure and operations decision-makers," October 2019.



3 Are your snapshots healthy enough to recover predictably?

IT organizations are used to verifying their backup copies, but the current process is manual, which is error-prone and non-scientific. Recovering from a bad snapshot can delay recovery but recovering from a compromised snapshot (with cyber vulnerabilities) can expose the organization to cyber threats.

So how can you ensure your snapshots are healthy? Automated policy-based backup verification gives backup admins much-needed visibility into the health of all their snapshots and helps them identify the good snapshot for recovery. Added visibility reduces the recovery process and also ensures no previously known/addressed cyber vulnerabilities are reinjected into the production environment by means of recovery.



4 Is there consistency between the original and backed up data and application?

Inconsistent data can lead to 'data not found error', or worse, permanent data loss. A backup solution with strict consistency ensures that the data is always first protected across the cluster before acknowledging the write back to the application. Achieving predictable recovery requires support for strict consistency in order to ensure backups are application and data consistent.



5 Is your RTO truly rapid?

Recovery time objective (RTO) is one of the key metrics used to measure the performance of the backup solution, and most modern solutions today support fast RTOs. As much as RTO is a technology discussion, it needs to be closely aligned with the business SLAs and agreed upon by the business and technology owners. In order to achieve predictable recovery, the backup solution needs to support rapid RTOs and IT owners need to clearly establish SLAs with their internal customers.



6 Can you ensure recovery at scale?

Trying to recover at scale without proper tools can be crippling to operations. Being able to recover only two or three VMs/objects at a time prolongs downtime, resulting in SLAs being driven by duration of recovery rather than business requirements. For predictable recovery, you should be able to recover any number of VMs/files/objects instantly. Rather than waiting for your backup solution to hydrate backup copies for recovery, deploy a backup solution that can maintain an unlimited number of fully-hydrated backup copies that can be instantly mounted. This makes data readily available, even while data is being restored in the background.



7 Is RPO allowing you to restore from any recovery point possible?

Restoring from the last backup seems logical but not always the best practice, especially after a ransomware attack when you might need to recover from an older, cleaner copy. No matter what the situation or business needs are, to achieve predictable recovery, the backup solution should allow frequent backups, the ability to store any number of snapshots without any performance impact, and the flexibility to restore to any previous point in time.



8 Can you recover anywhere?

In today's global environment, data and applications live in multiple locations, and recovery cannot be limited to the original location. Predictable recovery and meeting SLAs requires the flexibility to recover backup data and applications anywhere—on-premises or in the public cloud—and without any strings attached. Flexibility to recover to any target—original or alternative—gives confidence to users that their data and applications are not tied to a single location and can be moved to meet business SLAs.

Download this [Buyer's Guide](#) to learn more about modern backup and recovery.