



Creating and Testing Your IT Recovery Plan

Regular tests of your IT disaster recovery plan can mean the difference between a temporary inconvenience or going out of business.

WHITEPAPER

Quorum[®]
1-Click Instant Recovery



Testing your backups at least once per month is important to maintain engineering best practices, to comply with stringent standards for data protection and recovery, and to gain confidence and peace of mind. In the midst of disaster is not the time to determine the flaws in your backup and recovery system. Backup alone is useless without the ability to efficiently recover, and technologists know all too well that the only path from "ought to work" to "known to work" is through testing.

A recent study found that only 16 percent of companies test their disaster recovery plan each month, with over half testing just once or twice per year, if ever. Adding to the concern, almost one third of tests resulted in failure.

The reasons cited for infrequent testing include the usual litany of tight budgets, disruption to employees and customers, interruption of sales and revenue, and of course the scarcity of time. This survey covered mostly large enterprises, and the challenges are even greater for smaller firms.

According to the survey findings, half of the SMBs that have implemented disaster preparedness plans did so after experiencing an outage and/or data loss. Fifty two percent put together their plans within the last six months. However, only 28 percent have actually tested their recovery plans.


Yet new systems have arrived that allow daily automated testing of full recovery, putting such assurances in reach of every business. Backup without rapid recovery and testing will soon be as obsolete as buildings without sprinklers or cars without airbags.

Backup alone is useless without the ability to efficiently recover, and technologists know all too well that the only path from "ought to work" to "known to work" is through testing.


Define Your Objectives




Defines the length of time that an Entity can be unavailable before it impacts the company



The faster you want to recover, the more it costs



Does cost of recovery exceed the losses?



Many businesses have never tested the recovery process in the event of a server or site failure. With business continuity a core component of risk management, a well-rehearsed plan lays the foundation for confidence that your IT systems will work when needed most.

In discussions of IT disaster recovery, technical terms like recovery time objective (RTO) and recovery point objective (RPO) are often used, but what do they really mean? In practical terms, recovery time objective is the duration until a business can return to normal after the failure of a server or key data center, and recovery point objective is the place in the transaction flow where the business resumes.

- Recovery Time Objective (RTO)—How long can your business afford to be down?
- Recovery Point Objective (RPO)—How often do you backup? How much data can your business afford to lose in the event of a disaster?
- Level of Service (LOS)—What are your business' critical servers and essential units that cannot be disrupted?

Creating Your IT BC/DR Plan: Assess, Test, Repeat...

Often, business leaders might take for granted functional IT infrastructure, and fail to connect the dots from server downtime to business loss and failure. The RTO and RPO definition process helps highlight this connection.

When creating your IT Plan, consider the following key factors:

Assess Your Current Recovery plan and Your Company's Expectations

Implementing a disaster recovery plan includes documenting the process to bring a server or group of servers back online in the event of failure. An overlooked step in the process often flows from the assumption that an IT expert is always readily available. Due to the inherent unpredictability of a disaster, the IT staff that your company relies on may take time to find and start action. Considering this human latency when developing the recovery plan naturally highlights any undesirable complexity in the systems and processes, and the need to support recovery even with minimal IT expertise on hand.

Due to the inherent unpredictability of a disaster, the IT staff that your company relies on may take time to find and start action.



Questions to consider during assessment:

- Could a newly hired IT professional quickly handle the situation?
- Could a remote IT engineer talk a novice through the procedures?
- Could a smart phone web browser provide all needed access to bring your business back online?
- Could all this happen within the RTO and RPO requirements?


In addition to reviewing your IT Plan, survey your executive team to get a realistic picture of their expectations. You could spend too much time thinking of costly alternatives to cover aspects of daily operations that may not be critical. **When doing so, ask yourself and your executive team:**

- Specifically, what level of protection is necessary (RTO, RPO, LOS)?
- Which aspects of your company's business must stay operational in an emergency?
- Are your physical, as well as virtual servers, protected?

From Should-Do to Must-Do : Test Today!

The smartest approach to assessing where you currently stand is to test recovery for your servers and IT facilities as implemented today. If that out-of-warranty server running long past its expected lifespan is critical, take time to test how long it takes to recover to a functional state on an alternative platform. Most backup solutions provide a sense of false security. Even if your data is backed up, how long will it take to recover your systems? If you haven't tested it, assume it won't work. Most failures are first detected during recovery. By that time, even if the backup provider takes responsibility, rather than pointing out operational failures in your procedures, few practical alternatives may be available besides regretting avoidable omissions made painfully obvious in hindsight.

A well-developed IT disaster recovery plan will identify all key processes and expose any weaknesses, and the ideal way to uncover these is through testing.



To ensure you reach your objectives, perform a true recovery test on a critical server and capture these crucial observations:

- How long did recovery take?
- What data proved challenging to recover?
- Were all applications and related software returned to the exact state expected?
- Was the recovery process feasible for IT staff operating under stress with reduced tools?
- How would parallel recoveries amplify the challenges?

Learning from these questions on a single test will yield greater insight into your IT disaster recovery posture. Though obviously a sensible practice, human nature often postpones such disciplined testing, since historically it has been cumbersome, time-consuming, or simply impossible without unacceptable disruption.

New technology makes regular, even daily testing feasible. This automation provides a foundation for ongoing RTO and RPO reporting at a management level, allowing you to better estimate and mitigate risks to the business.

Without access to critical data in the first 24 hours after a crisis, forty percent of all businesses will fail. Such dire risk can be avoided by performing regular evaluations of your IT recovery process. Testing reveals not only whether the process can technically recover your servers, applications and data, but also the risk of any excess complexity.

Virtualization: Tool and Challenge for Rapid Recovery

Companies increasingly employ a mix of physical and virtual servers. Do virtual servers provide redundancy? Not automatically, and in practice, accidental deletion of virtual servers happens at least as often as failure of physical servers. Also, physical servers that are not easily virtualized often run the most important applications. Perhaps the domain server is virtualized, while the legacy database with crucial financial payload runs on a physical server too daunting to migrate. Often, site recovery proposals based on virtualization start with a theoretical assumption that all servers are virtualized, a practical impossibility in the hybrid physical/virtual environments most common in real business environments. Often, companies embark on a virtual solution but only implement it at their primary site because a replicated SAN, and an additional virtual host blade server for site redundancy is cost prohibitive.



With new business continuity solutions available today, companies can now easily and affordably implement remote site high availability without requiring complete hardware redundancy for the virtual environments.

Quorum's Recommendations

1. Identify all critical applications and servers. Include ancillary systems like domain servers.
2. In collaboration with business management and technical experts, set recovery objectives (RTO and RPO) that strike the right balance between risk mitigation and practicality.
3. Create a well-defined IT disaster recovery plan, and update it at least annually. Include allowances for locating and activating the right people.
4. Test your recovery process at least monthly. Choose the most critical servers, not just the most convenient.
5. Use test results to update your IT disaster recovery plan.
6. When reviewing potential solutions, include the recovery process a part of your evaluation. Test not only the technical backup capability, but also the complexity of the recovery.

Test your recovery process at least monthly. Choose the most critical servers, not just the most convenient.