IDERA® | SQL Server Whitepaper

PREPARE FOR THE WORST: HOW TO DEVELOP AND TEST A DISASTER RECOVERY PLAN

CONTENTS

PREFACE 3

WHAT IS DISASTER RECOVERY? 4

DEFINE A REALISTIC DISASTER RECOVERY PLAN 4

WHAT MAKES A SUCCESSFUL DISASTER RECOVERY PLAN? 5

ENSURE THE CORRECT BACKUP STRATEGY IS IN PLACE 7

PERFORM PERIODIC TESTING 8

FOCUS AREAS OF A DR TEST POST-MORTEM 9

WHAT’S YOUR DISASTER RECOVERY PLAN? 10

PREFACE

In the world of information technology, disasters come in many shapes and sizes. The loss of a single mission-critical database can be seen as a business-threatening disaster to some organizations. At the other end of the spectrum, unexpected events can force the complete shutdown of an enterprise's data center for an extended period of time. Falling between these two examples are a variety of circumstances that impact a portion of an IT environment and can cripple a company as they struggle to cope with the situation.

A disaster is usually thought of as being the result of extreme environmental or weather events such as a hurricane, flood, or fire. While this concept may be true when discussing the impact on municipal services or utilities, it is insufficient to describe the many ways in which an organization's computing landscape can be severely negatively affected. Additional concerns such as cyberattacks, faulty software, equipment failure, and human error can have catastrophic consequences for a company's IT resources.

WHAT IS DISASTER RECOVERY?

Disaster recovery can be defined as a set of policies and procedures designed to protect an organization from the negative impacts that accompany a disaster. Its goal is to recover the data, applications, and hardware required by the business so it can continue operations.

Disaster recovery (DR) is often used interchangeably with business continuity (BC) due to their similarity in purpose, but the two concepts are not identical. BC addresses the whole business and is a plan for returning it to full functionality after the impacts of an unforeseen event. DR is a subset of BC that specifically focuses on getting critical IT infrastructure, applications, and operations up and running. As such, it is a critical component of a BC plan and the one that directly affects an IT department.

DEFINE A REALISTIC DISASTER RECOVERY PLAN

Without a realistic disaster recovery plan, an organization is tempting fate regarding the ability of their IT department to restore their computing environment. There are many moving parts required to quickly and effectively address the sudden failure or shutdown of a company's computer systems. This paper will take an in-depth look at the components that need to be incorporated into the creation of a DR plan.

Before looking at the specific elements of a DR plan, there are two terms which need to be considered for each system and application that will be recovered. These are the recovery point objective (RPO) and recovery time objective (RTO). They are both important concepts when developing a DR plan but, as their names imply, they have different objectives.

The RTO is the target time defined for the recovery of IT services. This can vary based on the critical nature of systems and their tolerance for downtime. Setting RTOs may need to be done granularly based on the importance of specific systems. The time objective affects the types of measures taken to recover the systems and is a major part of the detailed recovery plan.

The RPO is focused on determining the ability of a business to recover without access to its data. It needs to be coordinated with the backup strategy to ensure that data is available to perform a point-in-time recovery. At best, disaster recovery will use data from the most recently created backups. The frequency of these backups is a critical factor when defining the RPO. If backups are being done once every 24 hours, then that is the minimum amount of time that should be built into the RPO. In a disaster the organization will need to go forward after potentially losing a whole day's worth of data, depending on the timing of backups and the catastrophe.

WHAT MAKES A SUCCESSFUL DISASTER RECOVERY PLAN?

A viable disaster recovery plan is a living document that should be updated regularly to reflect changes in the IT environment. Large corporations may have a designated team whose sole purpose is to keep the DR plan current with its evolving infrastructure and applications. Smaller organizations should assign a specific individual to oversee the DR plan and act as coordinator, even if it is a secondary function to their everyday activities. Someone needs to be responsible for keeping the plan in a functional state.

Here are some points that need to be considered when developing the DR plan. They should all become part of the disaster recovery planning document and be modified regularly to reflect changes to the computing environment or personnel.

- A current inventory of all hardware and software assets is an essential starting point for a DR plan. It should include information such as vendor technical support contract details, contact information, and contact methods. Resources should be prioritized in a way that makes sense to the business goals of the plan.

- Once a complete inventory is performed, RTOs and RPOs can be defined for individual systems or applications. Not all of an enterprise's computing resources carry the same level of significance to the business. It is good practice to use tiers to segregate systems in order of importance. Mission-critical systems should make up the top tier and be slated for quick recovery and preferential treatment by the team. Subsequent tiers contain other systems that are needed for a full recovery but are not required in the first wave of restores. Knowledge of the overall IT environment is crucial to making the right choices when organizing resources into prioritized tiers.

- Identifying the personnel responsible for the various aspects of the recovery process is the next step in plan development. Alternates should also be determined in the event the primary person is unavailable when needed. All individuals should be listed in the plan with their key responsibilities, their position, and emergency contact details.

- Communicating with the associated personnel and affected employees in the event of a disaster is key to an organization's ability to recover. This portion of the DR plan should address how to perform initial contact and notification of the DR team as well as the expectations for updates throughout the recovery process. Informing customers of the issue and providing updates on system recovery is also beneficial to maintaining confidence in the company.

- An alternate worksite needs to be defined in the DR plan for physical disasters that impact the workplace. The DR team may need to quickly travel to a site provided by a commercial disaster recovery vendor. Other personnel will need to have the capability to work remotely to get systems up and running again. The type of facility an organization employs for its recovery is based on its RTO. Hot sites enable the fastest recoveries but are more expensive than warm or cold sites that will require more time to perform the recovery at a reduced cost.

- Businesses that store and use sensitive information need to take precautions regarding how it is handled during disaster recovery. Very often a recovery operation uses some of the site vendor's human resources to augment their DR team. The plan needs to contain safeguards that prevent sensitive data from inadvertently falling into the wrong hands during the recovery.

- An adequate testing schedule needs to be built into the plan. As tests are conducted, improvements to the plan should be made based on results and gaps identified in the exercise's post-mortem.

There will be instances where only parts of the plan will be used, like in the failure of a single system or application. For this reason, the document should be segmented so that specific systems can be recovered without the need to consult the entire plan. Small disasters that impact a critical database require a different methodology than a data center that is flooded after a hurricane. An effective plan should be able to provide a course of action in either case.

ENSURE THE CORRECT BACKUP STRATEGY IS IN PLACE

The availability of backups is a crucial aspect of disaster recovery. Without access to backup data, systems cannot be recovered to the required state defined in the recovery point objective. When defining the recovery point and time objectives, it is vitally important to understand the way the systems are backed up, how frequently backups are made, and how easily those backups can be accessed.

Extensive infrastructures such as those in large computing environments pose complexities beyond simply verifying that backups are being successfully completed. Physical resource constraints often require backups to be performed in a staged manner with more important systems given priority. Differences in the dynamic nature of the data in question may necessitate that multiple or incremental backups be performed daily so recovery can come as close as possible to restoring systems to the state they were in before the disaster struck. This is why creating tiers in a DR plan that may have different RTOs and RPOs is often necessary.

Here's an illustrative example that drives home those points. Take a large data center that has numerous databases and servers in scope for disaster recovery. Daily backups will necessarily be completed at different times throughout the day, making it impossible to pinpoint an RPO time for the complete environment without neglecting the data changes that may have occurred in critical systems during the backup window. These systems may need to take additional partial backups to keep data as fresh as possible.

Data needs to be sent offsite for full protection in the case of a physical disaster impacting the data center. Traditionally this was done by backing up to tape and shipping the tapes to an offsite vault. Cloud backups have changed the dynamic somewhat, but recovering large amounts of data over the network can be too time-consuming to provide a realistic recovery strategy. Tapes are still used by many corporations as their disaster recovery media of choice.

Shipping tapes offsite implies that they are sent at a specific time and usually done daily. While they are all sent offsite at the same time, the data contained in the backups may have been generated throughout the day or at least during the overnight backup window. This means that the data on the tapes cannot recover all systems to the same RPO. Incorporating this fact into the DR plan and RPO expectations may impact the way the recovery is performed and the results it will provide when completed.

For this reason, a valid backup strategy requires multiple backups of critical and dynamically changing data during the course of the backup window. During recovery, systems can be brought up using full or incremental backups. After they are successfully brought online, further recovery can overlay database backups taken later in the process to achieve a recovery using the most recently backed up data. Therefore, the RPO of a database server and the actual database itself may differ and should be reflected in the DR plan.

PERFORM PERIODIC TESTING

It's safe to say that most if not all entities that create a disaster recovery plan hope they never have to implement it. Unfortunately many plans do indeed need to be executed to demonstrate the ability to recover from unexpected events. When a real disaster strikes is not the time to realize that the plan will not get the required systems up and running again in the time constraints defined in the recovery time objectives.

The only way to determine the probability of a successful recovery is to thoroughly test the plan. This should be done periodically, with a schedule built into the DR plan itself. Iterative changes to the plan resulting from the analysis of testing results will lead to improvements in the procedures used for future tests and potential real disaster recovery situations.

Testing can be performed in a variety of ways that provide varying degrees of confidence in the plan.

- A paper test is an abstract review of the procedures, timelines, and checklist that make up the plan. The DR team verifies the plan and makes any appropriate changes that need to be incorporated into the document.

- The next level of testing is a walk-through test. It is a more thorough review of the DR plan intended to further identify issues or inconsistencies in the plan and make the necessary changes before additional tests are performed or an actual recovery is required.

- Simulation tests conduct a recovery in real-time using physical machines and are used to gauge the level of accuracy in the plan and its ability to successfully recover systems. They can be done on a large scale where the complete plan is tested or to address specific systems or subsets of the environment.

- In a parallel test, the recovery systems are assessed on their ability to keep the business running in the event of a disaster. An organization's primary systems continue to process production work normally while a backup environment is used to test the validity of the DR plan in a more realistic setting. This type of test usually involves at least some members of the DR team traveling to the alternate site to perform the testing.

- The most exhaustive and complete testing is done via a cutover test. Here, the primary systems are disconnected and the recovery systems are used to supply the computing power for a company's production workload. This type of test should not be attempted until parallel tests have been run and any issues with the recovery plan have been ironed out.

Most companies that perform viable disaster recovery testing do so using simulation or parallel tests. Many outsourcing agreements include provisions to conduct these types of tests on an annual basis and DR teams meet regularly throughout the year to hone the plan and prepare for the test. Recovery exercises of this nature are watched closely by the corporate hierarchy who understand the importance of successfully recovering from a disaster.

FOCUS AREAS OF A DR TEST POST-MORTEM

Many specific areas of a computing environment can be the focus of a DR test. It is essential to document the test as it proceeds using checklists so actual results can be compared to abstract planning. Here are some things to look at to determine if the test was successful.

- Verify that all systems and applications that were in scope were recovered using the offsite media and DR documentation as the first priority. Comparisons to the expected and actual RTO can be used to tailor the plan going forward.

- Ensure that applications perform similarly on the hardware used at the recovery site, as this is another major factor that determines the success or failure of the DR plan. Specific application tests should be scheduled for verification after the systems have been recovered. Issues with incorrect RPO assumptions may come to light during these tests and allow for informed decisions to be made regarding the organization's backup policies.

- Establish that system recovery was prioritized correctly when developing the plan, and point out important changes that need to be made. Key applications that rely on other systems for support may have been inadvertently pushed to the top of the list. It may also become evident that some systems that were deemed expendable are required to fully recover the most important applications.

- Confirm that the proper personnel were assigned to the recovery team. In a testing situation, some resources need to be allocated to maintain the production systems while others work at the recovery site. Based on their performance during the test, modifications to the recovery roster may be necessary.

- Demonstrate that communication between the recovery personnel, management, and users is working according to plan. This is critical. While the team's focus is on the recovery, all associated stakeholders should be kept abreast of developments and issues that are impacting the exercise.

- Validate that the backup strategy and the procedures required are properly followed to deliver the necessary media to the recovery site on time. This is a critical factor whose failure can impact the ability of the team to complete the test in the designated time frame.

The company's disaster recovery coordinator should be prepared to lead the post-mortem meetings with an eye toward making improvements in the DR plan. After a detailed analysis of the test results, there are likely to be areas where changes can be made to streamline the process or to eliminate unnecessary systems from the plan. Best practices indicate that historical documentation should be kept so that tests can be compared as part of the process of developing the most comprehensive disaster recovery plan possible.

WHAT'S YOUR DISASTER RECOVERY PLAN?

The majority of organizations rely heavily on their databases and computer systems and would suffer irreparable harm if they were unavailable for any considerable length of time. Developing a robust disaster recovery plan with coordinated backup schedules and policies is an essential step in protecting against the loss of an enterprise's computing resources. Though under optimal circumstances it will never be used to deal with an actual disaster, it is essential, not optional, for your company to develop and test a disaster recovery plan. Your organization's survival may depend on it.

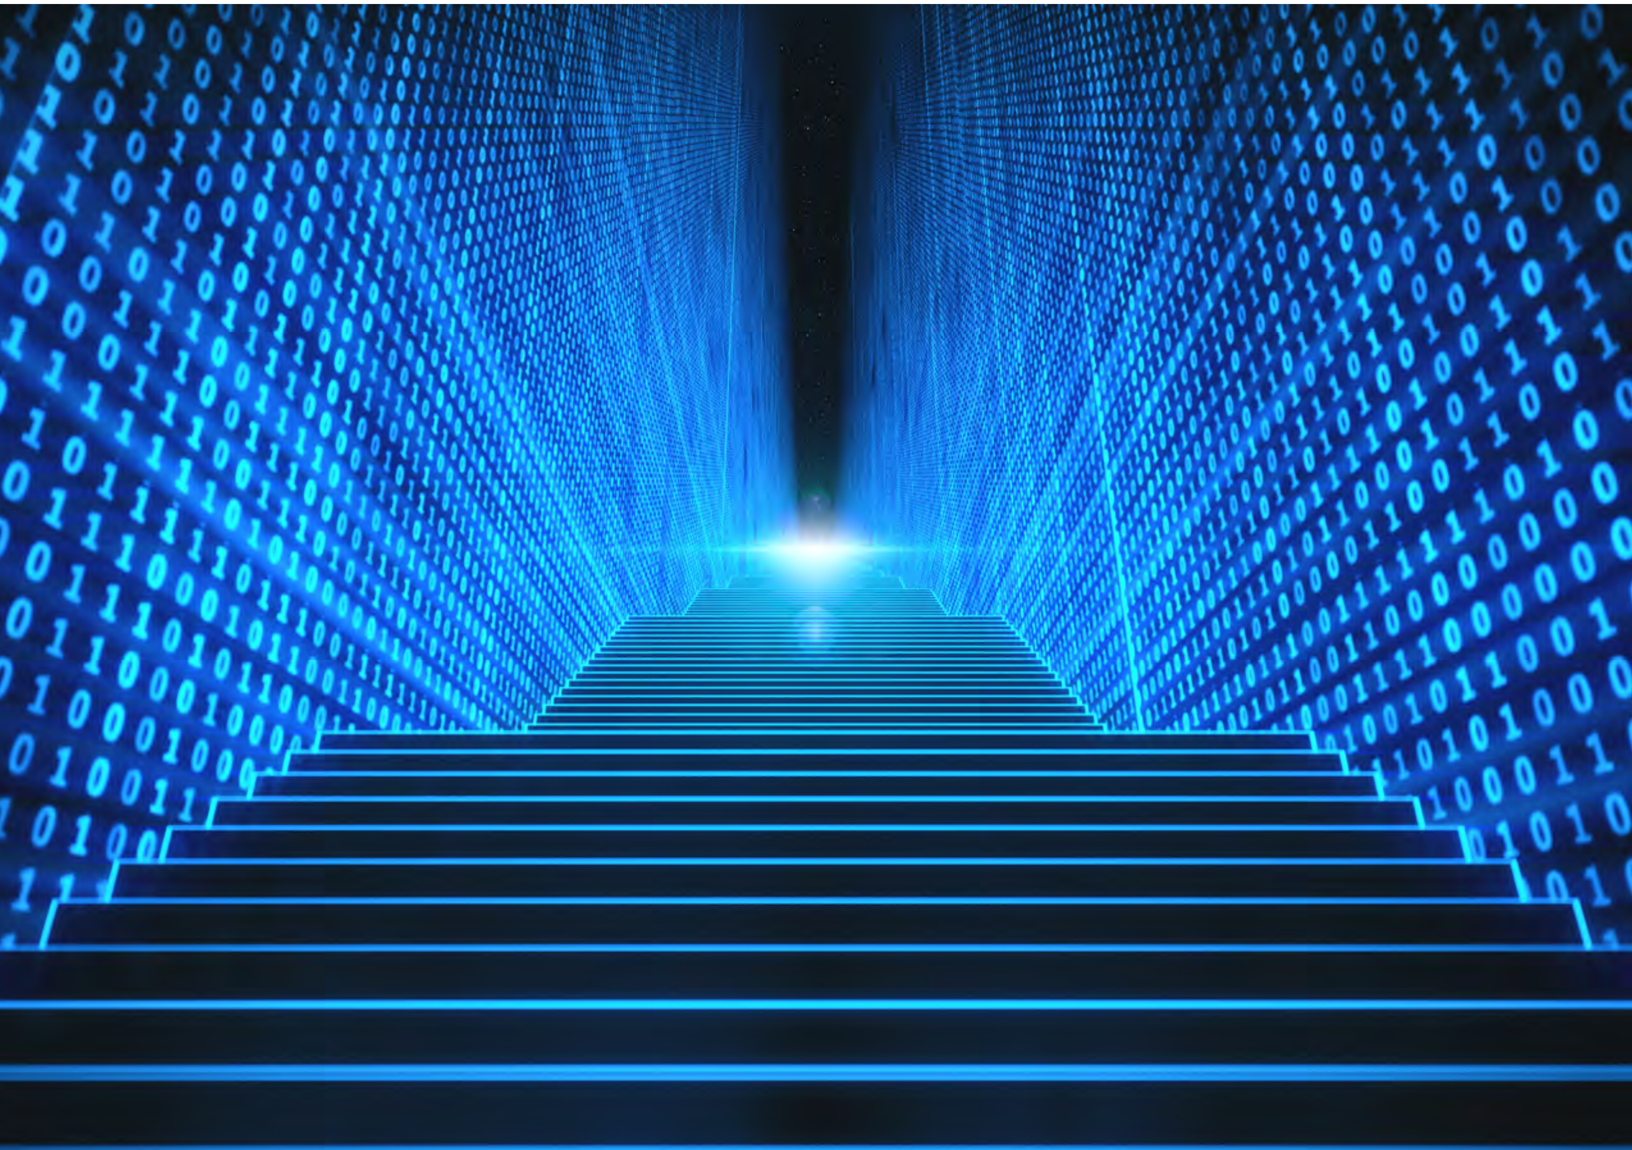

IDERA’S SOLUTION

IDERA’s SQL Safe Backup provides a high-performance backup and recovery solution for SQL Server. It reduces backup time by up to 50% over native backups and reduces backup disk space by up to 95%. It also enables automated backups of entire SQL Server infrastructures and ensures compliance with the backup and recovery policies of organizations. It turns backup files into virtual databases to query and modify data as if they were real databases without restoring. From tens of local SQL Servers to hundreds of global SQL Servers, it is the only backup and recovery solution that scales to meet the challenge.

Start a free, fully-functional, 14-day trial today!

Start for FREE

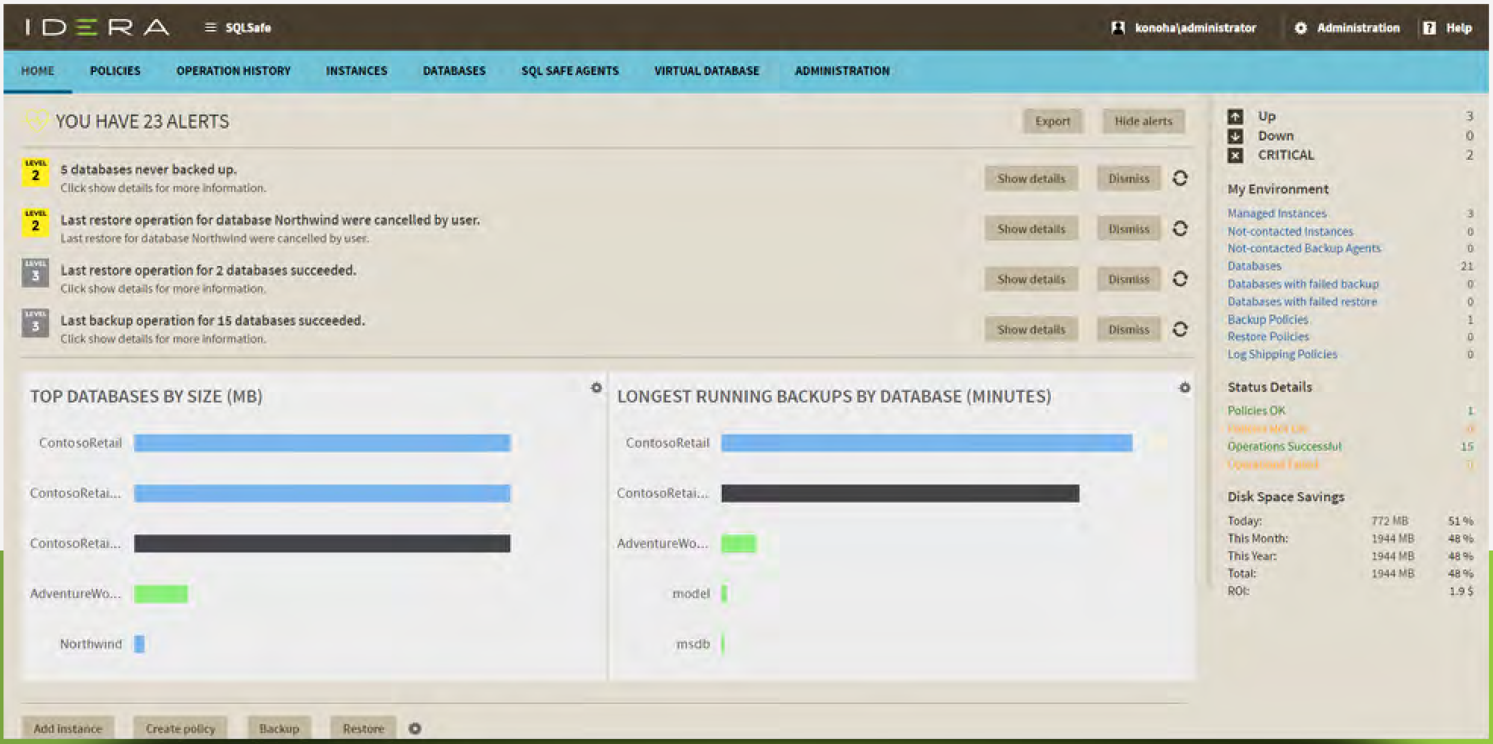

YOU HAVE 23 ALERTS

LEVEL 2 — 5 databases never backed up.
Click show details for more information.

LEVEL 2 — Last restore operation for database Northwind were cancelled by user.
Last restore for database Northwind were cancelled by user.

LEVEL 3 — Last restore operation for 2 databases succeeded.
Click show details for more information.

LEVEL 3 — Last backup operation for 15 databases succeeded.
Click show details for more information.

Up	3
Down	0
CRITICAL	2

My Environment

Managed Instances	3
Not-contacted Instances	0
Not-contacted Backup Agents	0
Databases	21
Databases with failed backup	0
Databases with failed restore	0
Backup Policies	1
Restore Policies	0
Log Shipping Policies	0

Status Details

Policies OK	1
Policies Not OK	0
Operations Successful	15
Operations Failed	0

Disk Space Savings

Today:	772 MB	51 %
This Month:	1944 MB	48 %
This Year:	1944 MB	48 %
Total:	1944 MB	48 %
ROI:		1.9 $