Real-time recovery

In information technology, **real-time recovery** (RTR) is the ability to recover a piece of IT infrastructure such as a server from an infrastructure failure or human-induced error in a time frame that has minimal impact on business operations. Real-time recovery focuses on the most appropriate technology for restores, thus reducing the Recovery Time Objective (RTO) to minutes, Recovery Point Objectives (RPO) to within 15 minutes ago, and minimizing Test Recovery Objectives (TRO), which is the ability to test and validate that backups have occurred correctly without impacting production systems.

Real-Time Recovery is a new market segment in the backup, recovery and disaster recovery market that addresses the challenges companies that have historically faced with regards to protecting, and more importantly, recovering their data.

Contents

Definition

A real-time recovery solution must contain (at a minimum) the following attributes: The ability to restore a server in minutes to the same, totally different or to a virtual environment to within 5 minutes ago and not require the use of any additional agents, options or modules to accomplish this. It must be able to restore files in seconds (after all, the only reason anyone backups is to be able to restore). It must perform sector level backups, every 5 minutes and have the ability to self-heal a broken incremental chain of backups should part of the image set get corrupted or deleted . It must deliver improved recoverability of data files and databases.

Classification of data loss

Data Loss can be classified in three broad categories:

1. Server Hardware Failure - Preventing a server failure is very difficult, but it is possible to take precautions to avoid total server failure through the user of Redundant Power Supplies, RAID disk sets.
2. Human Error - These disasters are major reasons for failure. Human error and intervention may be intentional or unintentional which can cause massive failures such as loss of entire systems or data files. This category of data loss includes accidental erasure, walkout, sabotage, burglary, virus, intrusion, etc.
3. Natural Disasters / Acts of terrorism – although infrequent, companies should weigh up their risk to natural disasters or acts of terrorism. How much data loss is the business willing or

able tolerate.

Platforms for data servers

Data servers can be either physical hosts or run as guest servers within a virtualization platform, or a combination of both. It is very common for a customer environment to have a mixture of Virtual and Physical Servers. This is where attention to detail must be given to the approach of protecting the data on these servers at regular intervals. There are distinct advantages in selecting a technology that is virtual or physical independent. This would limit the number of technologies that organizations will have to get trained on, skilled up on, purchase, deploy, manage and maintain. In an ideal world, if you can reduce the complexity of managing multiple products to protect your physical and virtual infrastructure you will reap the rewards. A technology that gets installed at the operating system level ensures consistency in an environment that is either physical or virtual and eliminates API compatibility or Disk Volume Structure limitations (e.g. Raw Mapped Devices, VMFS).

Strategies

Prior to selecting a real-time recovery strategy or solution, a disaster recovery planner will refer to their organization's business continuity plan for the key metrics of recovery point objective (RPO) and recovery time objective for various business processes (such as the process to run payroll, generate an order, e-mail, etc.). The metrics specified for the business processes must then be mapped to the underlying IT systems and infrastructure that support those processes.

Once the recovery time objective and recovery point objective metrics have been mapped to IT infrastructure, the DR planner can determine the most suitable recovery strategy for each system. The business ultimately sets the IT budget, and therefore the RTO and RPO metrics need to fit with the available budget. While the ideal is zero data loss and zero time loss, the cost associated with that level of protection historically have made high-availability solutions impractical and unaffordable. The costs of a Real-Time Recovery solution are far less than previous tape-based backup systems.

References

Retrieved from "https://en.wikipedia.org/w/index.php?title=Real-time_recovery&oldid=953038461"