





The Real Cost of System Protection, Disaster Recovery and Business Continuity

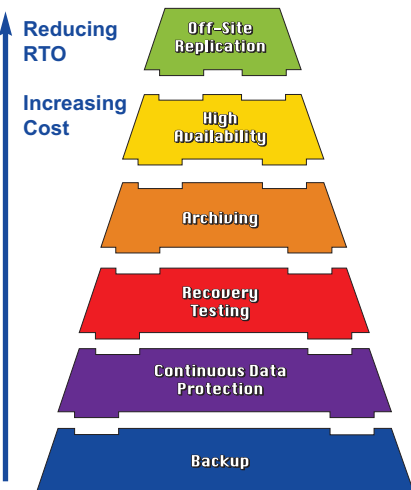
When most organisations think of disasters, the first thing that they consider are major disasters such as earthquakes, floods and terrorist attacks. Yet, natural disasters account for just 3% of all disasters. Hardware or systems failure accounts for 44% with Software Corruption, Human Error and Computer Viruses where historical data is required for recovery accounting for 53%.



In order to establish a successful Disaster Recovery/Business Continuity (DR/BC) Strategy, an organisation must first define their Recovery Time Objectives (RTO) and Recovery Point Objectives (RPO) for each of their systems for each of the disaster scenarios defined above. For example, an RTO of seconds may be suitable for an email system following a hardware failure, but following an earthquake an RTO of seconds could be considered futile if nobody has access to their desktops.

Recovery Time Objective (RTO) is defined as the maximum acceptable length of time to resume operations following an outage. Recovery Point Objective (RPO) is defined as the maximum acceptable age of data at the time of an outage, thus determining the age of the data when it is restored.

Most business people would want an RTO and RPO of zero seconds or, if they were in a good mood, a few minutes. They want all data accessible immediately after a disaster, and they can't afford to lose any data along the way. While this is actually possible with some technologies, the chief financial officer might not agree that the cost of these technologies makes it viable to use them on every computer in the data center.



BACKUP

The starting point for most organisations considering a DR/BC Strategy is their backups. Backups remain the most cost effective and least complex way to ensure protection for each of the disaster scenarios defined above. As most organisations perform backups every night, the RPO can be up to 24 hours. In the absence of regular system recovery testing, the RTO can be days or in many cases, never as the backup selections may not be sufficient to recover the system.

CDP

The implementation of Continuous Data Protection (CDP) will add some cost and complexity to the solution but has the advantage of reducing the Recovery Point Objective (RPO) to minutes as files are backed up as soon as they are saved.

Again, in the absence of regular system recovery testing, the RTO can be days or in many cases, never as the backup selections may not be sufficient to recover the system.

RECOVERY TESTING

In order to ensure that an organisations backups are sufficient for recovery in the event of a disaster, it is essential that system recovery is tested on a regular basis. Regular system recovery testing will considerably reduce the Recovery Time Objective (RTO) for most systems by eliminating any issues with the recovery process prior to a disaster actually occurring.

As part of the system recovery testing, provision should be made for the availability of suitable replacement hardware in the event of failure or the availability of an entire virtual infrastructure in the event of a major disaster. System recovery testing can also facilitate recovery to a warm recovery site in the event of a disaster destroying the primary site.

ARCHIVING

Many organisations do not archive older data from their servers. As a result, 70% of full backups consist of data which is over a year old. A successful archive implementation can reduce the time taken to backup and more importantly the time taken to restore, thereby dramatically reducing the RTO following a disaster.

HIGH AVAILABILITY

The implementation of high availability solutions such as clustering for critical systems can reduce the RTO to minutes in the event of hardware failure. A high availability solution may not protect against disasters which require recovery from historical data. As these high availability solutions are generally LAN based, they do not provide protection against natural disasters where recovery to a hot or warm disaster recovery site is required.


In the past clustering has proved an expensive and often problematic solution for organisations to implement. The costs associated with clustering such as dual path shared storage, software licences and redundant hardware all contribute towards pushing up the cost and complexity of a high availability solution. The introduction of virtualisation has made high availability a necessity for many organisations as multiple virtual servers are hosted on single physical servers thus creating a single point of failure.

OFF SITE REPLICATION

Off site replication to a hot or warm disaster recovery site can provide an RTO of minutes in the event of a disaster affecting the data centre. Depending on how it is implemented, off site replication can be an extremely expensive and complex solution.

Hot sites are fully equipped mirrors of existing data centres. This huge investment has all of the infrastructure needed to go live nearly instantaneously should a disaster occur. This includes cooling, network connectivity, servers, hubs, switches, and storage equipment and are typically synchronized continuously with the data centre for little to no loss of data during an event.

Warm sites have the cooling and electrical capacity, as well as pre-built communications equipment for immediate connectivity. There is usually some backup equipment ready for immediate use, such as servers and storage equipment. In some cases, a warm site will also house a storage appliance for periodic replication of data from the original data centre.



	Protection	Recovery Time Objective	Recovery Point Objective	Complexity	Cost
Backup		Days / Never	24 Hours		\$
CDP		Days / Never	Minutes		\$
Recovery Testing		Hours	Minutes		\$
Archiving		Under an Hour	Minutes		\$
High Availability		Minutes	Minutes		\$
DR Site Replication		Minutes	Minutes		\$

