



Empowering business to do more

The ROI of RPOs & RTOs

RPO (Recovery Point Objective) and RTO (Recovery Time Objective) measure the timescales that a particular solution can offer to recover your functions or data. This makes them one of the most important parameters to consider when choosing the correct DR solution for your business.

Disaster Recovery (DR) covers the technologies, services and procedures most businesses now realise they should employ to avoid the catastrophic loss of business critical computing services as a result of technology failure, cyber attack or environmental events.

RPO (Recovery Point Objective) and RTO (Recovery Time Objective) measure the timescales that a particular solution can offer to recover your functions or data. This makes them one of the most important parameters to consider when choosing the correct DR solution for your business.

Despite this being a measure of time and speed, it's important to note that the fastest or shortest doesn't always make the most sense. Some business functions would not necessarily require the lowest RPO or fastest RTO for one or all of their functions and thus there is no return on investment paying a premium for them.

Today, it is nearly impossible to highlight any modern business that would not require a stringent RTO/ RPO for all applications that support its primary line of business. There is, arguably, a distinction between businesses whose primary income stream is generated by On-Line Transaction Processing (OLTP) and those where online operations are important but can tolerate very short outages.

Recovery Point Objective (RPO)

Following a disaster, the RPO defines a moment in time that an application or service needs to return to in order to be acceptable to the business and customer. One example where an RPO needs to be extremely fast, perhaps zero, could be an online gaming site where it would be severely detrimental to the business to lose more than a second. At the other end of the spectrum, some businesses which don't primarily operate online could cope with a far longer RPO for some of their functions without any risk to brand or bottom-line.

Recovery Time Objective (RTO)

RTO predicts the time that an application or service will be unavailable while recovery is taking place. RTO can vary according to requirements from hours or days to zero. Costs of delivering quicker RTO are higher so again it's important to consider your requirements for different functions very strategically. Costs rise so dramatically with a zero RTO requirement that careful consideration should be given to the lost revenue versus the mitigation costs.

A longer RTO might be applied to a company bulletin board system or systems less critical to continued business operations such as staff holiday booking. In these cases, an RTO of 24 hours or no defined RTO would be fine. A business critical application, however, like an ecommerce platform, is very likely to require a very low RTO measured in minutes, seconds or even zero.

Node4 can address the range of recovery time and recovery point needs of a wide range of business types. Customers are able to choose the level of protection most appropriate to their business needs and separate business functions.

Technology options and their specific benefits:

Whether using cloud or a traditional on-premise data centres, the fundamental technologies remain the same. It's important to note, however, that hybrid cloud solutions bring a far greater range of options.

The broad range of solutions available from Node4 includes the following:

- Storage SnapShot
- Storage Replication
- Backup Software Application (Backup as a Service)
- Inter-Site Replication (Disaster Recover as a Service)
- Continuous Availability

Let's examine in more detail how each of these apply to RPO/RTO delivery to our customers.

Storage Snapshot – In line with most managed service and cloud providers, Node4 utilises enterprise class storage for our range of services. This allows us to implement an hourly snapshots schedule as standard. This means that the worst case RPO possible would be 59 minutes if a failure occurred. The RTO is typically rapid. The presentation of the 1 hour old data is near instantaneous.

Storage Replication – The Node4 storage infrastructure replicates the storage data and the logical containers where virtual machines are stored. Storage replication is the internal DR service within Node4 where customers choose to have cloud resources secured in a second data centre. This element of the service mitigates against the highly unlikely scenario of a systemic failure in the customer's set of virtual or colocated resources.

This typically employs the storage vendor's integral replication product. Within Node4, NetApp storage is heavily used for cloud and colocation solutions with the NetApp SnapMirror and SnapVault features utilised as highly mature and robust replication mechanisms.

The RPO for storage replication will typically be around 5 minutes but the RTO is longer than for snapshot restoration because making services available from the alternate location may involve administrative or configuration actions.

Backup Software Application (BaaS) – Node4 provides a cloud-based Backup as a Service (BaaS) facility supporting industry leading client software such as Veeam Cloud Connect for virtualised environments and nVault for end-point solutions such as laptops or single physical servers. These software solutions are typically implemented over Internet connections which has a direct bearing on the RPO/RTO measurements.

Restoring from backup data usually results in longer RTO time and less granular RPO. This is because a full or partial restoration of data is commonly from tape media which is a longer process than recovering from replicated copies on disk within local or secondary sites.

Backup software applications are not as well suited to low RTO as replication technologies. Nevertheless, your data is securely backed up for far less of a premium. It takes longer to reinstate but this will suit many businesses.

Intersite Replication (DRaaS)

This generic term differentiates replication of data that is operating at the hypervisor or server level rather than storage level. The main reason this type of solution might be applied is between data centres using dissimilar storage area networks (SANs). Equally, software (rather than storage based replication) can offer a more feature rich capability for both replication and backup capabilities.

Within Node4, the Zerto product suite is used for Disaster Recovery as a Service (DRaaS) between our VMware-based virtual data centres providing exceedingly robust and capable services.

In terms of RPO / RTO measurement, replication operating at the server or hypervisor level delivers essentially the same range as storage based replication with both RTO and RPO being as low as minutes/seconds. The limiting factors here tend to be bandwidth available between sites for synchronisation to take place.

The primary difference between storage and inter-site based replication is that the storage replication feature or software has no 'knowledge' of the data within Virtual Machines or volumes that are replicated. At the hypervisor level using tools such as Zerto, it is possible to isolate virtual machines, volumes or files for granular recovery.

Continuous Availability

This is a phrase used to describe technologies where no interruption to service is acceptable.

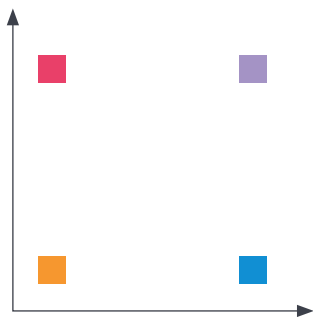
RPO / RTO terms are meaningless in this scenario where both are zero. The RPO is 'now' and the recovery time is instantaneous. How is this achieved?

In order to implement a continuous availability solution the environment requires a number of key components that can include some of all of the following:

- Application aware software, where the state of the application software is passed between physical
- or virtual servers (nodes) seamlessly on failure
- Clustering software such as Windows Server Failover Clustering (WSFC)
- Storage replication features where this isn't implemented at clustering layer
- Sufficient network bandwidth with low latency for full synchronisation between nodes

A commonly deployed continuous availability solution deployed within Node4 is Microsoft SQL Server High Availability (Always On) in conjunction with Windows Server Failover Clustering with high bandwidth / low latency networking over Node 4's private fibre based network).

This is how RPO and RTO applies to different types of organisations.



- **Very Low RTO/RPO**
Finance / stock trading / on-line booking / large on-line shopping businesses
This type of business will require continuous availability or replication that can deliver very low RPO/RTO. Solutions costs will be high.
- **High RTO/Low RPO**
Testing and some development environments
The loss of development code or data can be critical therefore the RPO is often very critical. However, the recovery time may not be so problematic. In most organisations, development or testing does not critically affect the income stream. The one exception to this is a software engineering business where the constant production of new code and testing activities are critical to the bottom-line. Such businesses may find themselves operating under low RPO/RTO requirements. Solutions costs will be at a medium level where requirements are typically met by backup solutions using frequent backup schedule.
- **High RPO/High RTO**
This typically applies to individual applications rather than business types. An extended recovery point metric implies acceptance of loss of data completely, which few businesses can tolerate for line-of-business applications. By comparison, lower-tier applications such as internal community bulletin boards, possibly even email in some circumstances, will not critically affect income stream.

A typical solution here might involve simple backup with a schedule of infrequent backups which in turn, minimise storage costs. This would be a cost-effective approach to meeting RPO/RTO requirements.
This is how RPO and RTO applies to different types of organisations.
- **High RPO/Low RTO**
Requirement for a high RPO implies a slow rate of data being updated therefore it is far less likely that important data is lost if the application is recovered to a much earlier time. However, a low RTO demands that the service is restored as soon as possible. A good example here is an on-line reference catalogues, a car-parts database for example.



Empowering business to do more