

Buyers' guide

VMware Disaster Recovery as a Service

At a glance

Using an on-premises site as a disaster recovery target is complex, expensive and unreliable. Disaster Recovery as a Service seamlessly delivered on top of a global public cloud has built-in advantages such as cloud economics, reliability, ease of use, flexibility to support infrequent but unpredictable characteristics of disaster scenarios and hence public cloud is an ideal DR target.

Key considerations

1. What are the RTOs required by different applications?
2. Does the service offer reliable infrastructure for the DR site along with failover automation and orchestration?
3. How complicated would it be to re-platform applications?
4. Can you run non-disruptive DR tests?
5. Does the service offer cost efficiency compared to existing alternate DR solutions?

Many organisations realise the importance of implementing a Disaster Recovery (DR) solution, or are required to do so to comply with government regulations. Maintaining a separate on-premises site to serve as a DR target is labour intensive and requires significant investment to deploy and maintain, especially when considering that it is not utilised on a day-to-day basis. Disaster Recovery as a Service (DRaaS) running on an elastic public cloud with built-in automation reduces the amount of under-utilised hardware and maintenance tasks, simplifies the deployment in case of a DR event, and increases the reliability of the DR solution with non-disruptive testing.

On-premises DR solutions are often expensive and require extensive expertise to deploy, maintain and operate. Deploying an on-premises DR site means paying in full for DR real-estate, hardware, and software, but using them only when the main data centre goes down. This makes it difficult for IT decision-makers to justify the high expenses of such initiatives. Even after a DR solution has been deployed, testing it is often labour intensive and disruptive. As a result, many organisations compromise on the level of protection their business-critical applications receive in the event of a disaster. To overcome these challenges, many organisations look to the public cloud. The goal of this guide is to help organisations understand the key factors they should take into account when considering public cloud as a solution for their DR needs.

According to Gartner report, DRaaS providers can be an excellent avenue for organisations wanting to "get out of the data centre business" and save money in the process. In fact, it is not uncommon for DRaaS prices to be 30% to 50% of what it would otherwise cost to build out similar capabilities.¹

Factor #1: Identifying the RTOs required by different applications

While organisations can certainly protect all their applications, doing so can be very expensive. Instead, organisations should categorise their applications based on their Recovery Time Objectives (RTOs), which is the acceptable amount of waiting time before an application comes back online. Some DRaaS solutions have RTOs ranging in minutes, others ranging in hours, and others in days.

Different business needs dictate different RTO levels. A revenue-generating application can rarely be down for long, while HR applications can typically take 8 hours or more to come back online without meaningful business impact. Naturally, the shorter the RTO requirements for an application, the more expensive it is to recover those applications in the required timeline. Organisations should therefore start by protecting business critical applications, and then, if enough budget remains, go on to protect lower tiers of applications. The desired level of RTO should then be the key factor when selecting a cloud-based DRaaS solution.

Apart from RTOs, the organisations should also consider whether the solution offers any capability to rapidly inspect various recovery points in a short period of time when recovering from a ransomware attack. This ensures rapid restoration of applications after ransomware attack.

Landscape of Disaster Recovery solutions

Data backup only:

These solutions replicate organisations' data to a second on-premises site or to the cloud. However, they leave organisations exposed to long down-times during disasters since applications do not have infrastructure to run on. They also do not include simple DR testing, and require extensive manual work once infrastructure is acquired.

Automated DR to an on-premise site or to a co-location:

Solutions of this kind reduce the amount of manual effort, but still require high capital investments in real-estate, hardware, and software that are not often used. Additionally, they are more difficult to scale.

Automated DR to data centres owned by DRaaS providers:

These solutions provide most of the benefits of automated DR to an on-premises site, and have a better cost structure that reflects the infrequent use of the DR target.

Customers of these solutions should assess the reliability of the DR infrastructure and the financial stability of the DRaaS provider.

Automated DR to a global mega-cloud:

These solutions provide most of the benefits of automated DR to an on-premises site, and have a better cost structure that reflects the infrequent use of the DR target.

Customers of these solutions benefit from reduced risk thanks to the reliable infrastructure, global availability, and financial stability of the mega-cloud provider. However, some of these solutions require re-platforming of customers' applications.

Factor #2: Determining if the service offers reliable infrastructure for a DR site along with failover automation and orchestration

It is relatively easy to backup data to the cloud, however, simply relying on backup alone exposes organisations to significant risks in the case of a disaster. If only data is copied to the cloud, organisations are left with the task of setting up a full environment, spinning up compute instances, moving data to the right cloud storage service, and setting up networking. Many of these tasks are highly manual and require significant time to execute. For applications with an RTO of two days or more, that is not an issue. However for revenue-generating applications, that is typically too long.

For more critical applications, organisations should choose cloud-based services that offer DR failover orchestration and automation. Such services deploy a DR environment based on a pre-defined runbook. They spin up the required nodes, power on VMs in the correct sequence according to the right dependencies, run scripts, and map IP networks automatically, all with very little human intervention. This ensures that critical applications can be powered up in time and minimises any business impact of a disaster.

Many vendors offer Disaster Recovery as-a-Service. The scale and sophistication of DRaaS solutions varies. Many vendors often lack the scale, reliability, financial stability, and global availability of the major cloud providers. Since organisations need to rely on DR solutions at critical times, when their main data centres are down, the reliability of the DR infrastructure is also a key factor to consider.

Factor #3: Evaluating the level of complication involved in VM format conversions

Many modern microservices-based applications are typically agnostic when it comes to which public cloud they run on. However, traditional applications, which are still very dominant in many organisations, are typically deployed as VMs. Different hypervisors have different VM formats, and many public clouds do not have the same VM formats as organisations' on-premises VMs. In order for applications written and deployed on one hypervisor to be used on another hypervisor, the VM disk format needs to be converted. VM format conversion is typically a long and complicated process, and organisations can spend many months in the process. More importantly, during this process, an organisation's applications are not protected in the case of a disaster.

Factor #4: The need to run non-disruptive DR tests

Creating a DR plan is not a one-time activity. Data centres are not static – existing applications get updated or replaced, and more applications are added over time. This results in a drift between an organisation's original DR plan and an effective DR plan that can keep up-to-date with the changing applications.

In order to make sure this situation does not occur, organisations need to test their DR plan often, with best practices suggesting at least once per quarter. Since these tests are not real disasters, they should not affect an organisation's current running applications. In other words, these tests need to be non-disruptive. Organisations need to easily define, maintain and test DR failover procedures in a constantly changing IT environment.

Furthermore, some organisations are required by law to perform DR tests and to present the results in an audit. A good DRaaS should offer customers extensive non-disruptive testing and provide detailed reports generated by these tests.

Disaster Recovery as a Service (DRaaS) Use Cases

Entirely new DR:

Meant for organisations that have only backups or do not have any DR plan in place.

Expand existing DR plans:

Some organisations already have an on-premises DR solution, but they use it only to protect a few workloads. With DRaaS, these customers can protect the rest of their workloads to the cloud, while keeping their existing DR plans unchanged.

Replace existing DR:

Some organisations have a mandate to reduce their on-premises footprint or "move to the cloud". DRaaS is a natural solution to move the on-premises DR site to the cloud.

DR between different cloud regions:

Even the largest public clouds have outages, making DR relevant to customers who are running apps in the cloud. Customers of DRaaS solutions can protect their applications between different cloud regions.

How can Xtravirt help?

Xtravirt is a cloud consulting and managed services business with a 15-year track record of delivering success across the customer lifecycle from strategy and design, through to deployment and lifecycle support.

As VMware's most qualified Principal Partner in the UK, Xtravirt offer a consultative and bespoke approach to protect and recover customers' mission critical data.

Whether it is a disaster recovery refresh or creating an entire cloud security roadmap, Xtravirt simplify every step of your DR and ransomware transformation, enabling you to achieve transformative cloud strategies at speed and scale.

xtravirt.com

Factor #5: Ensuring cost efficiency compared to existing alternate DR solutions

The DRaaS solutions need storage components in steady state to store the data that needs to be protected. Organisations need a highly efficient storage layer in the cloud to store this data in order to optimise their costs. In order to further reduce the costs, organisations should be able to spin up the infrastructure in the cloud only when needed during a DR testing or failover event.

Finally, organisations should also consider the factors such as whether the DRaaS solution always require bringing back all the data upon failback or does it allow optimised failback, what is the price and metric of the data being protected, what are the cloud provider's egress charges while transferring the data to/from the cloud infrastructure etc. These factors will significantly affect the DR costs.

Conclusion

As organisations turn to public clouds for Disaster Recovery as-a-Service, they should consider various factors of their DR strategy. A robust DRaaS offering should be able to provide the RTOs needed for business-critical applications. It should also offer orchestration and automation of the failover process, and non-disruptive testing. All this should ideally be done without the need to re-platform applications, and run on top of a reliable public cloud. And finally, the DRaaS should help customers optimise their DR costs.

VMware offers 2 in-house DRaaS solutions supported on VMware Cloud on AWS:

VMware Cloud Disaster Recovery offers on-demand disaster recovery, delivered as an easy-to-use SaaS solution, with cloud economics. It combines cost-efficient cloud storage with simple SaaS-based management for IT resiliency at scale. Customers benefit from consistent VMware operations across production and DR sites and a 'pay when you need' failover capacity model for disaster recovery (DR) resources. VMware Cloud Disaster Recovery can protect a very broad set of IT services in a cost-efficient manner, with fast recovery capabilities (On-demand DRaaS).

VMware Site Recovery™ for VMware Cloud™ on AWS offers customers a complete DR service. VMware Site Recovery can protect mission critical IT services that require very low RPO and RTO (Hot DRaaS). With Site Recovery, customers have access to global, reliable infrastructure, with the familiar interface of vSphere and vCenter, and without the need for re-platforming. Additionally, by leveraging widely proven and tested DR solutions such as VMware Site Recovery Manager™ (SRM), customers can orchestrate and automate failover, failback, and IP network remapping, and conduct non-disruptive testing that generate extensive reports.

-
1. [Gartner: "Reduce Costs and Piggyback DR Investments"](#), May 29 2020