


REDESIGN BACKUP STRATEGIES FOR NEXT-GEN DATA CENTERS



Mohamed Sohail

Senior Solutions Architect
Dell Technologies
Mohamed.sohail@dell.com

Amr Shaheen

Partner Technology Strategist
Microsoft
Amr.shaheen@outlook.com

George Crump

Founder, Lead Analyst
Storage Switzerland
Georgeacrump@storageswiss.com

Table of Contents

The Cyber-Attack Challenge.....	4
Responding to the New Threat Landscape	5
Protected Data is Now THE Target of Attacks.....	5
Protecting Protected Data	7
Modern Backup Strategies Need New Recovery Processes	8
Recovery Speed and Data Capture Intervals Are Still Critical.....	8
Introduction to Isolated Recovery	9
Modernizing Data Protection for Data Privacy & Ransomware	9
Air Gap solution design	10
Solution components.....	10
Compute vault	12
Management host	12
Backup application host	12
Recovery test host.....	12
Infrastructure services hosts	12
Low speed switch.....	12
High speed switch.....	12
Data Domain system.....	12
Considerations	13
A) Isolation.....	13
B) Backup.....	14
C) Recovery.....	14
Scenarios of implementation.....	14
Shared Switch	14
Dedicated Switch	16
Firewalled Vault.....	17
Appliance	18
Cloud Business Continuity and Disaster Recovery Strategy	19
1- DR for your Cloud VMs (IaaS VM Backup)	20
2- DR for your cloud-aware Applications	21
3- DR for your Hybrid Environment	23
4- Same Cloud DR between two regions.....	24
5- Multi-Cloud DR.....	25
Summary.....	26
References	27

Disclaimer: The views, processes or methodologies published in this article are those of the authors. They do not necessarily reflect Dell Technologies' views, processes or methodologies.

Forgive me, for I have sinned.

I have sacrificed backup consistency to get better benchmark numbers.

I have failed to write tests that simulate failures properly.

I have tested on too few nodes or threads to get meaningful results.

I have tweaked timeout values to make the tests pass.

I have failed to monitor my environment to find out where the real challenges.

I know I am not alone in doing these things, but I alone can repent, and I alone can try to do better. I pray for the guidance please give me the strength to sin. No more.

The threat landscape has changed, and backup strategies need to change to keep pace. In the past, an organization’s primary motivation for developing a backup strategy was recovery from a data center-wide disaster or to recover from data loss because of human error. While those concerns remain, the primary threat facing the data center today is cyber-attack. The organization needs to do everything reasonable to stop a cyber-breach with perimeter defenses, but they also need to make sure they have processes to limit the exposure of a successful breach and speed the recovery from a breach.

The Cyber-Attack Challenge

A data center disaster caused by natural circumstances (flood, hurricane, power outage) is challenging because of the logistics required to recover. In most cases, the organization needs to recover to an alternate facility, either to another one of their datacenters or in the cloud. The organization must also make sure that secondary locations are accessible to essential personnel to ensure that operations can be resumed seamlessly.

A cyber-attack is unique because, in most cases, personnel don’t need relocation, but hardware and software need a thorough inspection to make sure they are clean of attacking code. Like a natural disaster, a cyber-attack can create millions of dollars’ worth of damage but without causing any physical damage. Also, a natural disaster doesn’t typically continue to attack over and over again. Once the disaster has passed



organizations can focus on recovery efforts without much concern of a repeated attack. A cyber-attack, however, can continue to threaten the organization indefinitely. NotPetya was an attack on retail companies that are estimated in \$15 million per day was lost in forgone revenue. It spreads in seconds after the initial infection. Interconnected businesses help automate and streamline processes, however such attacks rendered useless hundreds of critical servers, desktops, and phones, impacting over 10,000 employees. The production of more than 15 factories was brought to a stop. This included real-time inventory management systems, where the downtime of these systems directly impacted the overall supply chain, impacting final assembly of goods.

The malware exploited known vulnerabilities in operating systems and found their way into third party software that is consumed by the organization. The malware was inserted into a patch of this software. Rather than pay the ransom, the organization decided to focus on recovery. Paying ransom is not recommended and would have not been effective in this case. Hackers promised a decryption key upon

ransom payment. Forensic analysis of the malware determined that a decryption key would not have allowed the data to be recovered.

Finally, while a natural disaster may attempt to destroy everything in its path including the data center and everything it contains including data, it does not typically target facilities and data that are off-site, thousands of miles away. However, some cyber-attacks specifically try to find and eliminate backed up data first before compromising the rest of the environment. These pre-strike efforts include trying to find all connected copies of data no matter how far away they are from the data center. Even attacks that don't specifically attack protected copies can accidentally find their way to that data if the organization does not take the right precautions.



Responding to the New Threat Landscape

The modern data center needs to respond differently to the new cyber-attack threat. In addition to off-site data copies for protection from a data center-wide threats like a natural disaster, the organization needs to protect off-site data to make sure that protected copies of data are isolated and minimize points of access to it. IT also needs to take extra time in recovery to make sure that it is not recovering compromised data or malware files that can re-start the attack.

Protected Data is Now THE Target of Attacks

When malware or some other form of cyber-attack breaches an organization, the backup process is supposed to be the last line of defense, enabling the organization to recover from the attack without much, if any, data loss. The problem is that cyber-attackers know that organizations count on the backup process and as a result, backup data is also a target of these attacks. Bad actors can compromise the backup process by either removing or corrupting backup data stores and configuration files, or by inserting nefarious code into the backup store.



Cyber-attacks can directly attack backup data by either encrypting the backup data store or completely removing it. If part of the backup policy is to replicate the backup store to an offsite location, the cyber-attack can follow the path to the replication site and also remove the DR copy of the backup data store. The success rate of the attack is dependent on what level of access the attacker gains, but all online copies of backup data are specifically susceptible. The attacker may also gain access to the backup software's configuration files and metadata history files.

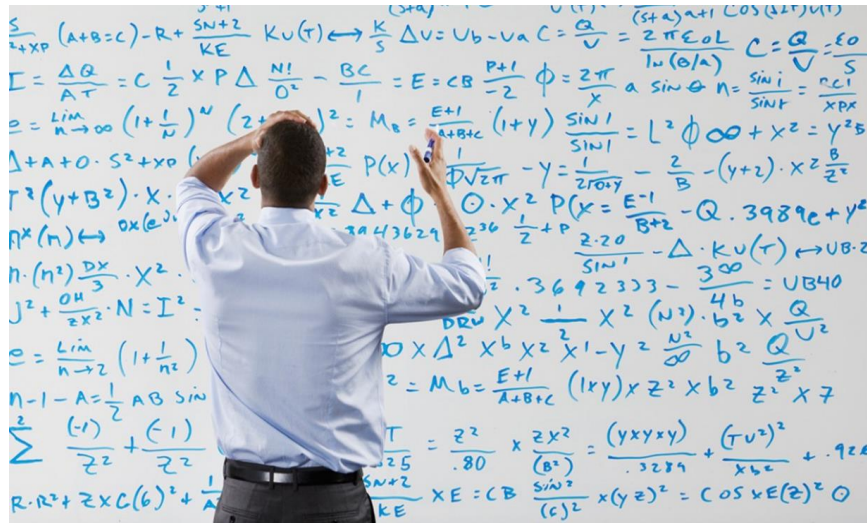
Once the attacker corrupts the backup data or backup configuration files, they can start attacking production data, knowing that the organization can't recover, and can hold the organization hostage. Once the organization identifies the attack on production storage, it turns to its backup data to recover, only to find that the attacker has compromised it as well. The only option is for the organization to deal with the public embarrassment of admitting to the attack. It may also need to pay the attacker to release their data.



A subtler attack, yet one that is simpler to accomplish is to insert nefarious code into the backup data store by placing attack trigger files on production storage but not triggering them until they have been backed up. Most backup solutions have no insight into what data they are backing up. They are just doing as ordered and backing up a specific file system, volume or mount point. If the attacker can insert their code into those file systems or mount points, then the backup system backs it up.

The bad actor only needs to be patient and not trigger the attack right away. Instead, they wait for their files to be backed up several times before executing. After the attack occurs, and IT eventually discovers it, IT then resorts to its normal process of recovery from the backup. One can assume that IT first removes the trigger files from primary storage, which then makes them believe they have stopped the attack. The problem is the restore process also restores the original trigger files along with all the other files. Once restored the trigger files activate again, compromising systems and data. The organization finds itself trapped in an attack loop.

Protecting Protected Data



Protected data is the data that the backup process creates and updates as part of the backup process. The rise in cyber-attack related disasters means that organizations need to look at new methods to make the protected data resilient to attack and it needs to change its recovery methods. Organizations need to make sure some of the protected copies of data have a gap in time from protected data. More importantly, some copies of the protected data need to be more difficult to access from the network. This data is typically considered off-line creating an air gap.

The problem with creating these gaps is data protection software continues to improve, and more of the protected data is online than in the past when tape was the only backup medium. Modern solutions can make copies of data more frequently than ever, which while improving recovery point objectives also exposes protected copies of data to corruption from cyber-attack. Software vendors need to deliver solutions that can quickly secure protected copies of data. IT professionals should also take steps to make sure that protected data is secure.



The most vulnerable data is data that is accessible via Windows SMB shares or on Windows volumes. A backup solution that is either Linux-based, or Windows-based but supports writing to an NFS or Object Storage mount point eliminates about 80% of the concern. Linux-based backup servers should similarly write data to an SMB mount or to a Cloud Storage location. To date, there is no known case of a cross-platform attack. Most cyber-attacks are written for a single platform and don't move across operating systems or mount types.

The next step is to create a gap in connectivity or an air-gap so that the cyber-attack can't follow the path to the protected data. Tape vendors claim superiority here since tape libraries aren't typically accessible by operating system mount points. Cloud vendors, however, can offer similar air-gapped protection by copying data via native cloud protocols and saving replicated backup copies to a write-once, read many (WORM) file system at a cloud provider. Object or Cloud Storage support provides the protocol switch described above and a WORM file system makes changes to the underlying protected data impossible. The organization can also go a step further and make sure the cloud account only provides minimal access credentials.



The advantage of a cloud copy is random access and no waiting for tape mounts. Data can also be quickly scanned for verification, a critical capability when dealing with attack loops.

Modern Backup Strategies Need New Recovery Processes

The default recovery protocol is to recover data as quickly as possible with minimal data loss, and there is increasing pressure on IT to meet these demands with greater regularity. Vendors continue to improve their ability to recover data rapidly and to reduce data loss. Malware and other cyber-threats, however, means that data centers can't focus only on rapid recovery; they must also have the option to perform an isolated recovery so that data is analyzed before being moved into production.

Recovery Speed and Data Capture Intervals Are Still Critical

While cyber-attack is a constant threat, it is unlikely that organizations are in the midst of an attack all the time. The majority of recoveries are a result of accidental user deletion or application corruption of data more so than the need to eliminate the damages caused by a malware attack, which means recovery speed and lack of data loss are the priorities the vast majority of the time.

Organizations can't afford to solely focus on recovery from malware or ransomware because the best practices of recovery require the organization to perform a staged recovery which adds time and loses data fidelity. IT needs to continue to invest in software and hardware solutions that improve the ability to recover rapidly and to capture backup copies more frequently.

Introduction to Isolated Recovery

When an organization needs to recover from a malware or cyber-attack, in most cases it knows that it is involved in that situation due to the scope of the restore request. In these situations, organizations need to perform a staged recovery. A staged recovery means restoring to an isolated section, sometimes called a **Sandbox**, of the data center or an isolated section of the cloud so that the backup data is verified before moving it to production storage. Once IT restores the data to the isolated area, it can use standard malware scanning solutions to verify that silent trigger files don't exist and that data within the restore set is free from corruption. If IT finds a malware trigger file or corrupted data, they need to move to the next previous backup set or try to manually extract malware and corrupted data from the currently restored set.



If the data set from a backup is found malware- and corruption-free, IT can execute the same recovery process to production stores. The process is of necessity a double restore, and IT needs to factor in the time involved in its recovery expectations.

If IT decides to use a recovered set with malware or corrupted data removed instead of utilizing a previous backup generation, backup administrators need to take great care when copying this data to production storage. A straight copy of data from the staged recovery area to production storage may not transfer all the file attributes correctly. The process is similar to migrating data from one storage system to another, and the organization may want to use a replication utility to perform the transfer.

Modernizing Data Protection for Data Privacy & Ransomware

To tackle such challenges, we propose two ways that are considered successful and can be easily combined to achieve the best-of-breed solutions for customers who have mission-critical applications, including those who can't be completely on cloud – i.e. Banking sector.

The rest of this paper details the above concepts, but it is critical that the organization place a high priority on backup modernization. There is a temptation, especially with backup, to take an “if it ain't broke, don't fix it” attitude. This attitude is dangerous. The threat landscape is changing, and organizations need to prepare now by constantly challenging their backup and recovery methodologies to make sure they are ready for the next potential threat.

Air Gap solution design

Solution components

Any device on the Internet with an open inbound port will be attacked. It's a matter of when, not if. This is the philosophy behind air gap: using PCs that are not connected to the internet, other devices nor the company's primary network. For high-assurance organizations, i.e. utilities, critical infrastructure, banks, government agencies and other heavily regulated companies, air-gapped devices can be a simple solution to today's complex data security challenges.

The idea behind air gap technology is simple; leave no doors or windows open, and criminals will have no way in and data no way out. There are very few ways to infiltrate air-gapped computers because data can only be shared to and from the machine via a FireWire connection, a USB flash drive or other external, removable media.

But as many IT teams have learned firsthand in recent years, air-gapped devices aren't immune to insider threats, zero-day attacks or the risk of coming in contact with malicious USBs. Stuxnet, a virus that wreaked havoc on centrifuges used at a uranium enrichment plant in Iran back in 2010, is one of the most notorious examples of compromising an air-gapped environment. The attackers first infected the PCs of external contractors programming the plant's systems in Iran. Unaware they had been breached, the contractors brought their infected laptops into the plant to transfer data to the air-gapped systems with a flash drive.

More recently, WikiLeaks released new Vault7 files revealing the details of malware aimed at infecting air-gapped PCs using USB drives. This leak, known as Brutal Kangaroo, included a user guide on "Drifting Deadline," malware designed to first infect a computer and then any thumb drive plugged into it. After infecting an air-gapped device, the malware would



perform an encore, employing a software called "Shadow" to create a custom covert network within the victim's closed network where the attacker could carry on freely with further attacks.



Figure 1: Zero-day threats

So how can high-assurance organizations protect their air-gapped devices? Many organizations in air-gapped environments turn to traditional, signature-based anti-virus solutions for additional protection, but they require ongoing, manual updates. These frequent signature updates are an enormous burden for IT teams, and sometimes IT falls behind on this time-consuming maintenance. Furthermore, signature-based anti-virus is inadequate protection against zero-day threats or newly-created targeted malware precisely because it hasn't been released yet – a requirement for these applications in order to have the malware signatures.

All this amounts to a defense that's less than airtight, which isn't good enough for high-assurance companies in heavily regulated industries.

Let's look at the basics of the solution design. We took Dell EMC as an example. Every customer however, can choose their customized version based on his or her operating conditions.

Compute vault

The vault includes compute and a hardened backup storage solution such as Dell EMC Data Domain. The compute can be virtual or physical on x86 systems.

Management host

The management host requires 1 vCPU, 2 GB of vMem, 4vDisks /root 60 GB, /isorec 100 GB, /var/mail 30 GB, /opt 10 GB. It must have dedicated NICs to Data Domain if it is physical.

Backup application host

The backup application host can be sized as needed. It communicates with the production backup application host.

Recovery test host

Can have as many as required. Must be sized to run the largest application.

Infrastructure services hosts

These can be consolidated and must be sized as needed. DNS is recommended, while AD/LDAP and Kerberos are optional. Also, an NSX with firewall and VPN are optional.

Low speed switch

Optional, used to connect to Data Domain management, interface and production NOC *(through VPN). A unique network subnet is recommended except when Data Domain replication interfaces are connected to the production network. This is used for remote management.

High speed switch

A high-speed switch connects production and inter-vault traffic. A Unique Network Subnet is recommended, except when Data Domain replication interfaces are connected to the production network.

Data Domain system

The Data Domain system must contain separate NICs to production Data Domain and for inter-vault traffic.

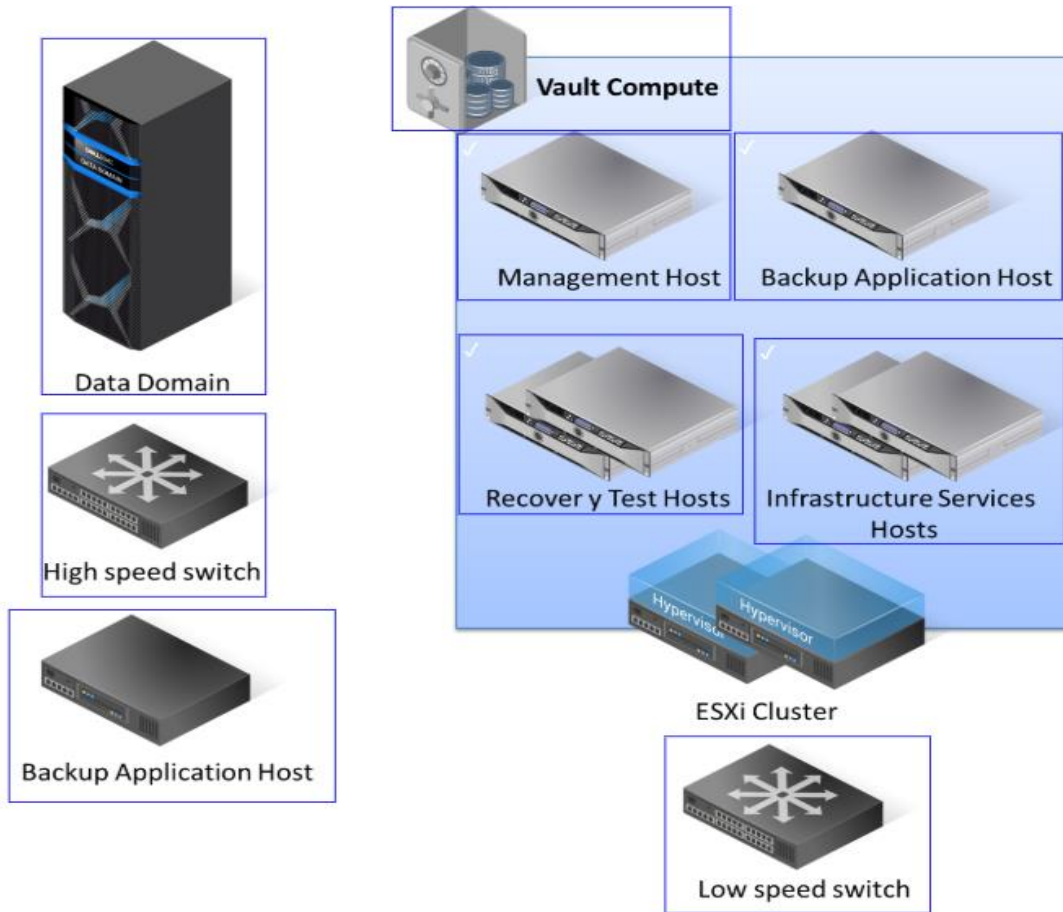


Figure 2 :Isolated recovery solution architecture

Considerations

A) Isolation

The solution must be isolated and off-network. It is protected in a limited access, secure location. Physical access to the isolated recovery vault environment must be protected with appropriate access controls.

Network connections outside of the vault are limited to only the necessary connections and protocols for data copy-in cycles and error reporting out.

None of the equipment in the vault, such as network switches, storage, and compute technologies should be shared with environments outside the vault.

When the data is not being copied to the vault, the network connection that is used is disabled and disconnected.

Wireless technologies are not allowed in the solution.

B) Backup

Backup is a point-in-time copy of data that may consist of a backup image or disk-based snapshot. The solution keeps multiple point-in-time, fail-safe recovery copies to meet business or regulatory recovery requirements.

Point-in-time copies are protected from deletion with an immutable technology, such as Retention Lock, WORM lock, etc. WORM technology has a configurable expiration. The retention lock can be compromised if the vault storage device is accessed and the storage device RAID groups are destroyed. Therefore, appropriate governance and access controls must be applied to the physical vault and vault storage device. Continuous replication-based technologies are not the same as isolated recovery solutions.

Be careful of the retention time and the number of copies because they can consume space that will not be cleared as a result of a cleaning due to the retention lock compliance. If compliance is used, the only way to expire the data is to wait for the retention time to expire. If capacity becomes an issue, capacity can be added to the Data Domain or other storage device.

C) Recovery

Isolated Recovery Solutions are designed to recover an organization's protected and dependent applications and data from a destructive cyber-event.

Recovery processes are designed to allow for partial recovery. If two non-dependent applications are protected and only one was attacked, the recovery process shall only require recovery of the affected application.

Recovery processes are designed to allow for the restoration of multiple point-in-time copies to enable the organization to find a usable copy. The Recovery Process is periodically tested to meet regulatory or business requirements.

Scenarios of implementation

Shared Switch

Figure 3 depicts a shared switch between the production infrastructure and the IR Vault. Isolation is achieved by configuring VLANs. A VPN is used to secure communications to NOC and Production Data Domain.

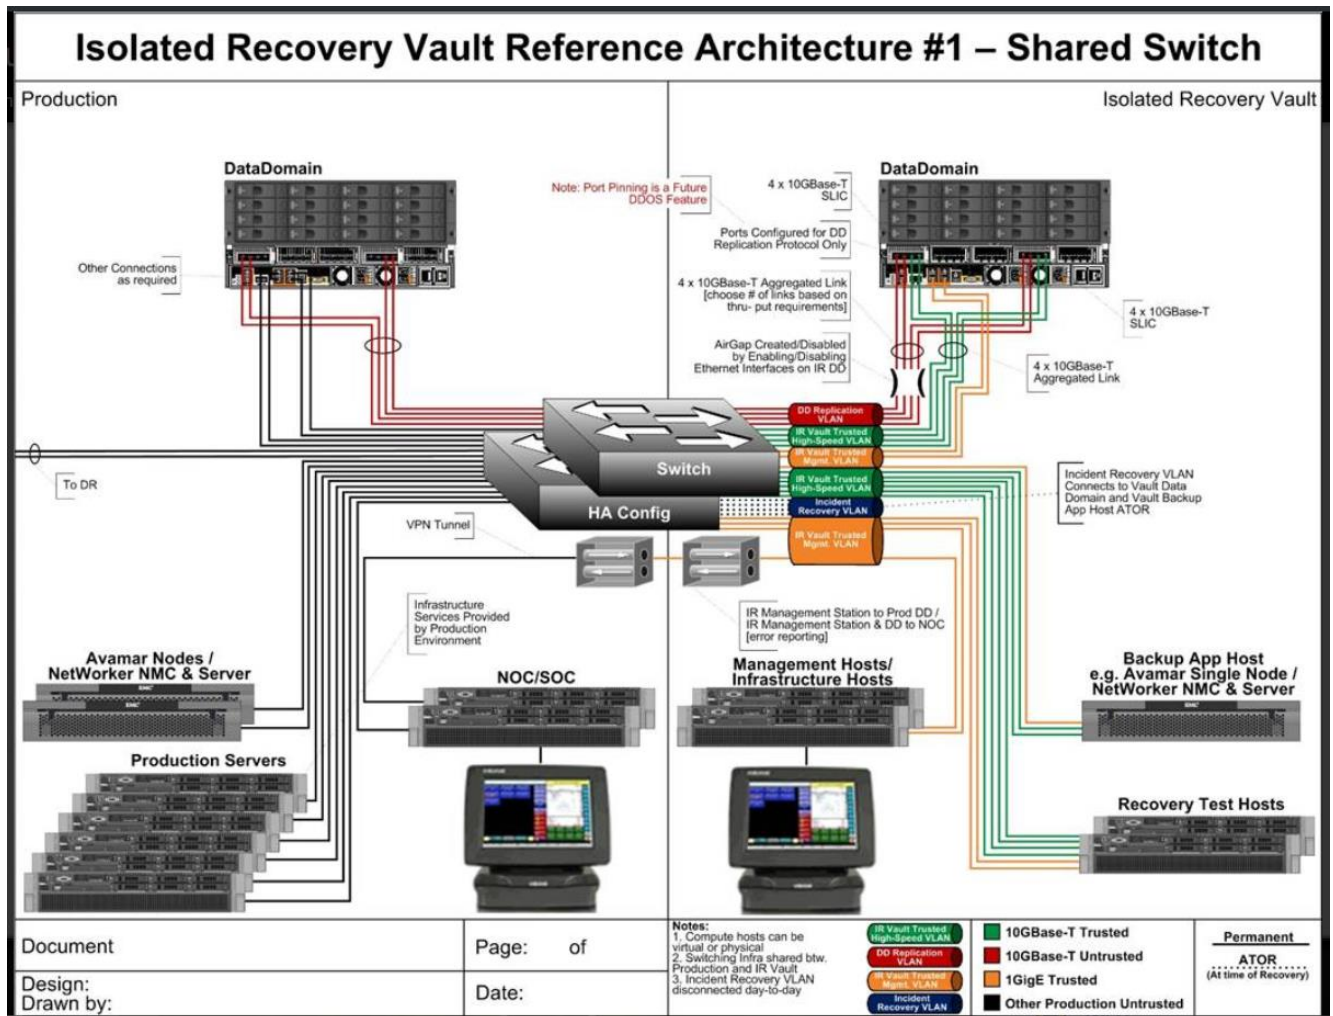


Figure 3: Shared switch option

The term enclave is defined as a section of an internal network that is subdivided from the rest of the network. The subdivision is done to limit internal access to a portion of a network.

PROS	CONS
Somewhat secure.	Not an enclave.
Inexpensive.	Switching infrastructure is shared with production.
	Data Domain has outward-facing interfaces that can accept any Data Domain-enabled protocol.
	If the switching infrastructure credentials are compromised, isolation is lost.

Dedicated Switch

Figure 4 shows a dedicated switch between the production infrastructure and the IR Vault. Dedicated switches in the IR Vault are separate from production switching infrastructure. Vault switches can be single switch or configured in high availability configuration.

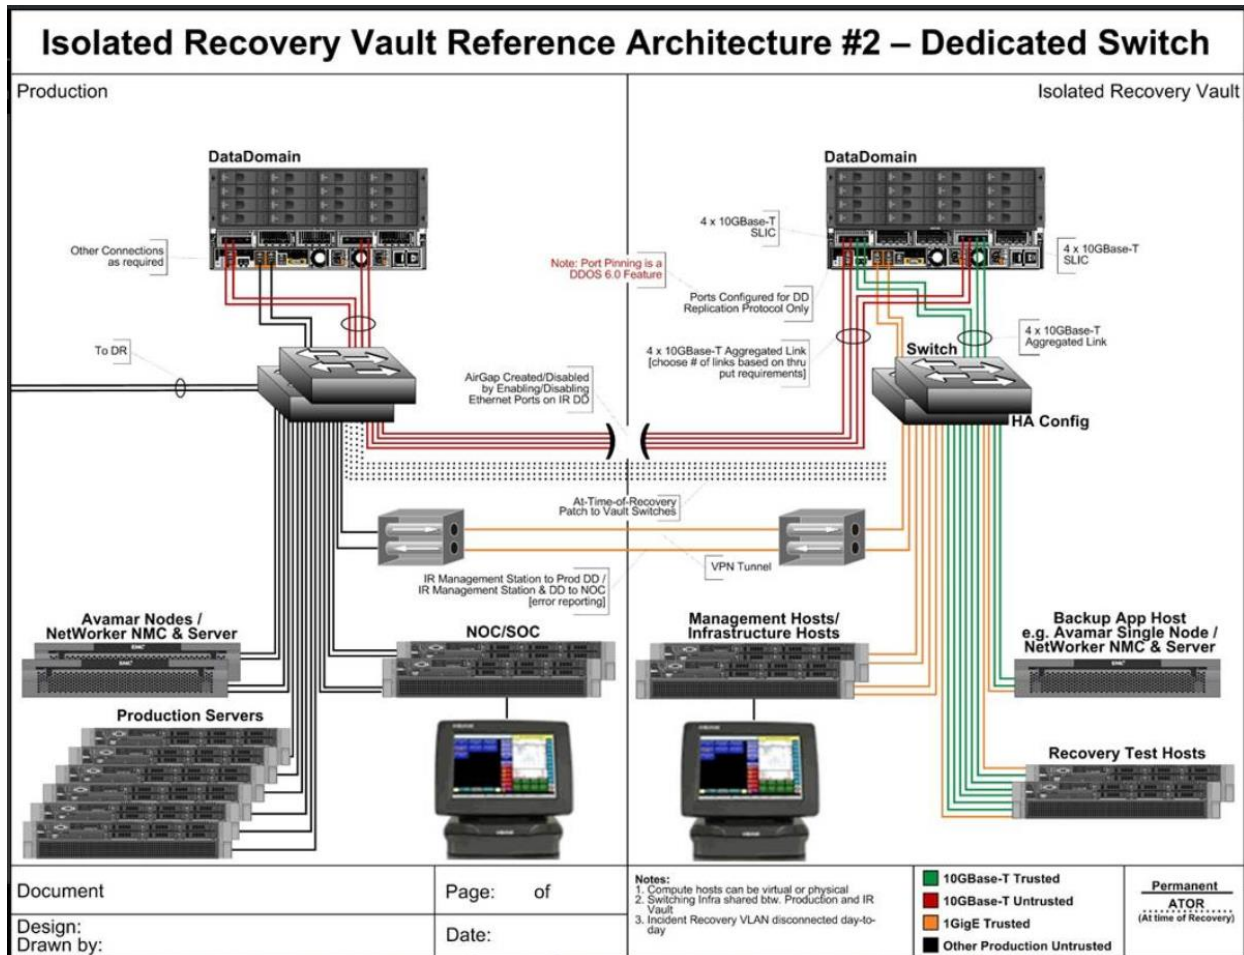


Figure 4: Dedicated Switch option

AirGap Control and isolation are provided by enabling and disabling the replication ports and using a VPN to secure communications to NOC and production Data Domain.

PROS	CONS
Nothing is shared with production.	More expensive solution since it requires more switches.
Somewhat secure.	Data Domain outward-facing interfaces can accept any Data Domain-enabled protocol.
Better isolation.	
Enclave enforced.	

Firewalled Vault

Figure 5 depicts a firewall in the Vault. Replication Circuits are run through a virtual or physical firewall. Packet inspection is not needed.

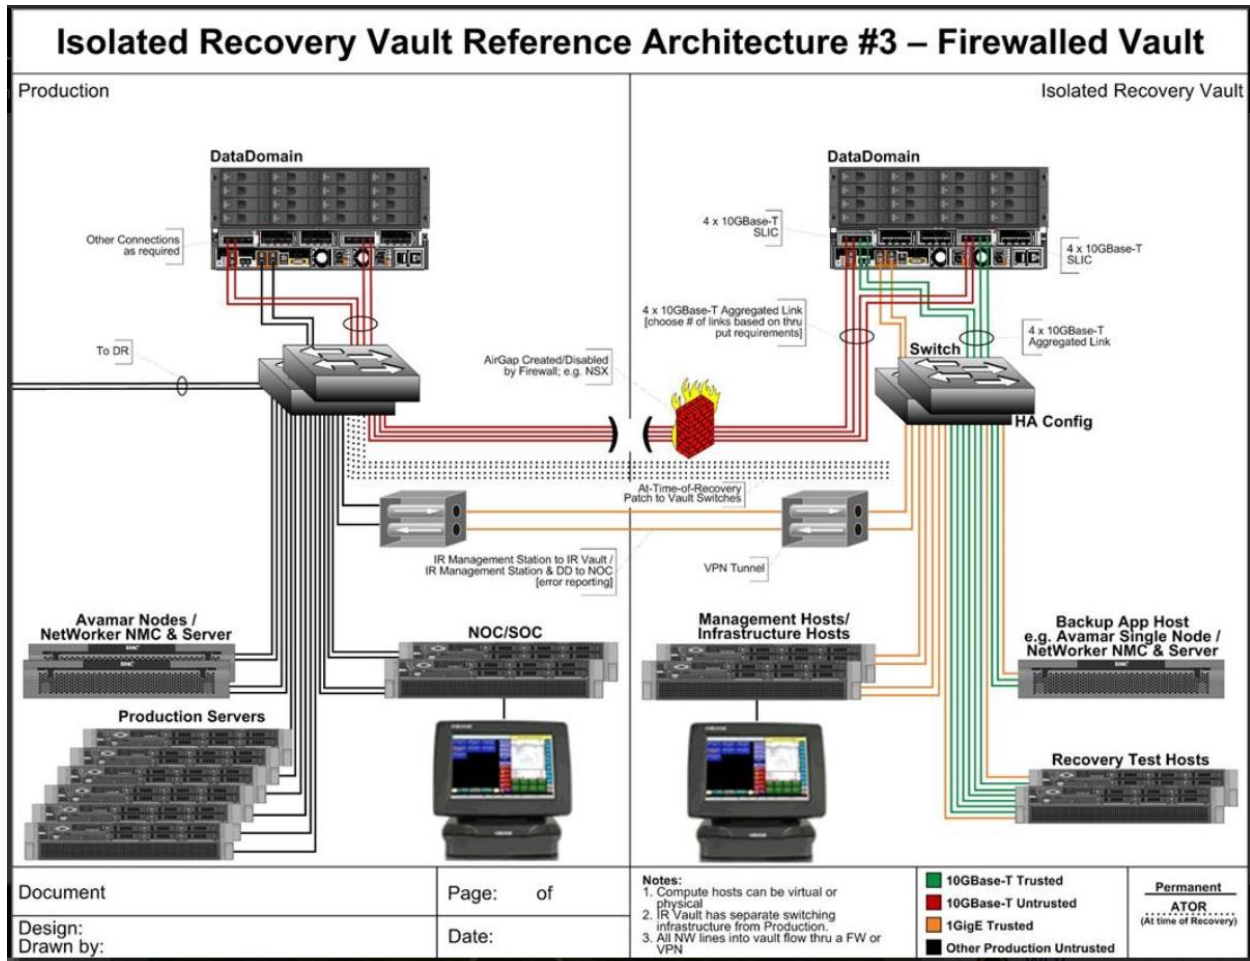


Figure 5: Firewalled vault option

The firewall provides isolation to secure the Isolated Recovery Vault. A VPN is used to secure communications to NOC and Production Data Domain.

PROS	CONS
Most secure.	More expensive solution due to the firewall costs.
Better isolation.	
Nothing shared with production.	
Enclave enforced.	
Allowed protocols from the untrusted side can be limited.	

Appliance

Figure 6 describes a hyper-converged appliance that provides compute, firewall, and test stations.

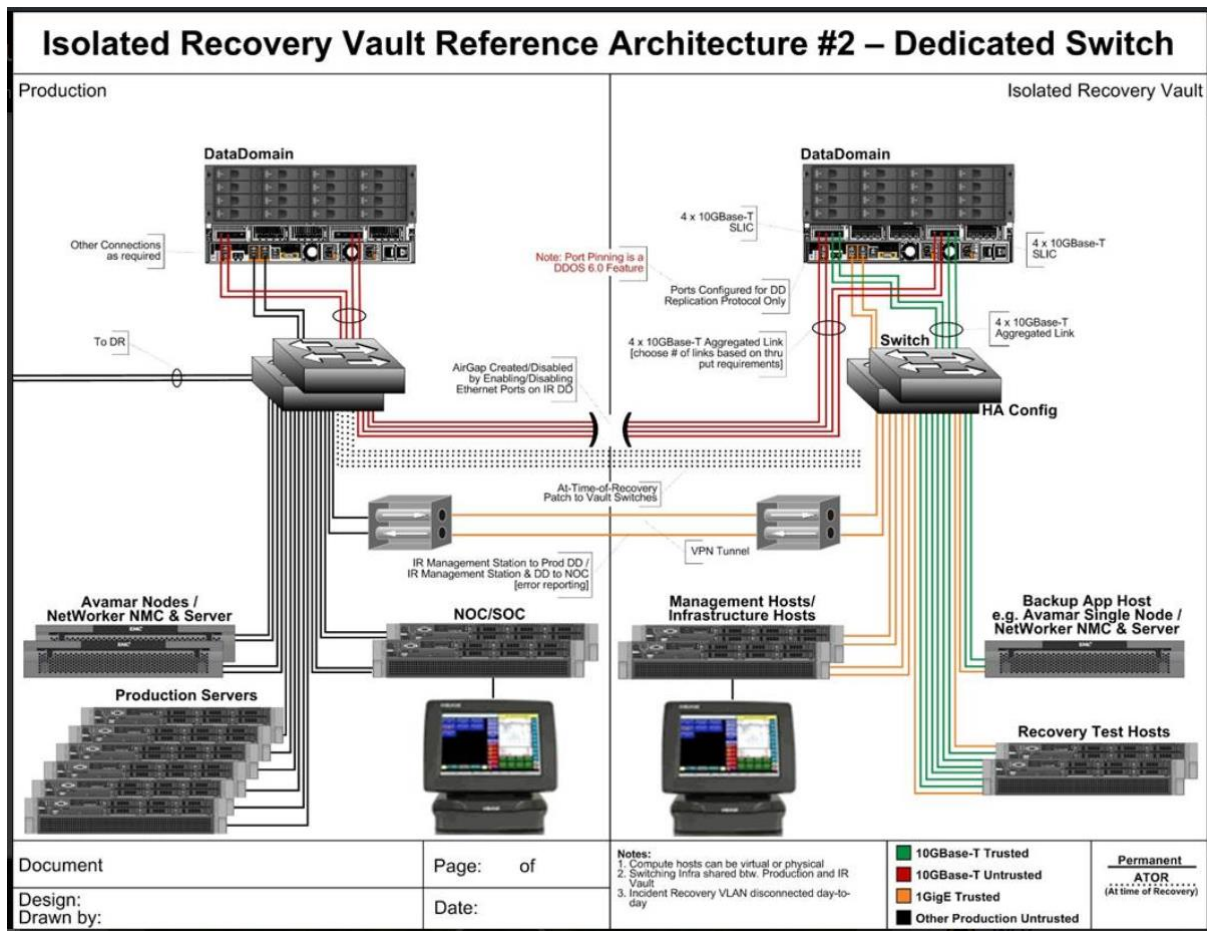


Figure 6: Appliance option

An NSX Firewall provides isolation. It secures the Isolated Recovery Vault. A VPN secures communications to NOC and the production Data Domain.

PROS	CONS
Most secure.	More expensive solution.
Better isolation.	
Enclave enforced.	
Allowed protocols from the untrusted side can be limited.	
User compute is secured with PCoIP desktops and Horizon View VDI.	
Dell EMC supports all Isolated Recovery Technology.	

Cloud Business Continuity and Disaster Recovery Strategy

Enterprises need to grow their data centers extensively, and the lack of hardware resources is a primary bottleneck to that growth. Companies are growing out of their hardware faster than their refresh cycles will support. Complicating that growth is the data these enterprises are creating. It is more critical (and confidential); which means, that they need to have a solid business continuity strategy in place.

When CFO's discuss business continuity, they cringe at the thought of doubling their current IT costs. This is why many companies are moving toward a strategy that includes cloud backup and disaster recovery. Most enterprises are confused about where to start or on how to plan their Cloud Business Continuity strategy. This section walks you through some of the main considerations to keep in mind when building a Cloud Business Continuity strategy.

A solid business continuity strategy must include 3 main pillars:

- 1- **High Availability**: meaning that when your applications or infrastructure have a failure, a second instance automatically replaces it in your primary site
- 2- **Disaster Recovery**: meaning that when your entire primary site is having a failure, you have a secondary site that can completely run all your workloads seamlessly
- 3- **Backup**: meaning that if any of your data bits get corrupted, deleted, or lost, you can restore a healthy version of your data at any point in time.



These three pillars are obvious to some people, but in execution there are some apparent challenges, such as reliability, costs, and complexity of setting up the above. This is where cloud comes to help, leveraging Cloud technologies for a business continuity plan can help you:

- 1- Reduce time to recovery (RPO & RTO)
- 2- Reduce operating costs
- 3- Reduce the complexity of setting up and managing the environment

One factor that is usually overlooked in Business Continuity and Disaster Recovery (BCDR) scenarios, is **Compliance**. Compliance towards customer data, country regulations, industry regulations, or even security regulations. In most cases, all cloud vendors compete to become more compliant in all these pillars, trying to offload this hassle from their customers and onto themselves. Simply stated, your **BCDR** plan should be **Simple, Secure, and Cost-Effective**.

Let's examine the most common scenarios for leveraging cloud in your BCDR strategy and drill down into the details for each scenario.

1- DR for your Cloud VMs (IaaS VM Backup)

This section discusses the considerations of backing up an enterprises cloud VMs. All cloud vendors offer backup capabilities that are built in to their VM management portals for easy and fast setup. These are the first line of defense in an enterprises BCDR strategy.

For example, Microsoft Azure uses Azure Backup Service to back up VMs. This feature uses the **VM Snapshot extension** in the case of Windows VMs, and the **VMSnapshotLinux** extension in the case of Linux VMs. As one can see in the architecture depicted in Figure 7, the Azure backup service can back up from the VM Backup extension, directly from the disks, or from the Vault (where snapshot copies are being stored).

One important consideration for enterprises planning to use the Azure Backup Service, is to keep in mind that this feature doesn't include data on the temporary storage (that is created by default with the creation of any VM). This temp storage is used to store the system paging files, and therefore can be erased in cases of your VM migrating from one host to another.




Figure 7: Azure VM Backup Architecture

The table below shows that the technology is similar at both Amazon and Microsoft; and probably other vendors as well. It is good to know the organization's options though, as some vendors are more advanced than others in some features. Selection is based on the exact specifics of the enterprise's environment, what features they may require and which vendor provides these features.

VENDOR	MICROSOFT AZURE	AMAZON WEB SERVICES
LOCALITY	Azure Only	AWS Only
LOCATION OF BACKUPS	Recovery Services Vault	Amazon S3
WHAT GETS PROTECTED	VMs + all Disks	VMs + all Disks
SNAPSHOT TECHNOLOGY	VSS	VSS
SNAPSHOT TYPE	Incremental	Incremental
REQUIRES AGENT	No	No

2- DR for your cloud-aware Applications

Backing up your applications and workloads are usually a challenging task and having the right tools for the job is key in your BCDR strategy. Applications are the essence of enterprises, and they are transforming to cloud-aware apps, which means they are designed to run on the Cloud and On-Premises alike. Therefore, regardless of what your application runs on from an OS, DB, or other applications, you should be able to back it up seamlessly and securely. Below are some key considerations to keep in mind for backing up your cloud-aware applications:

- a- Ensure that you have the lowest possible RPO & RTO for each of your applications and workloads and set up your strategy accordingly to achieve these numbers.
- b- For Multi-Tiered Applications, you need to ensure that all tiers are being backed up consistently to ensure that the application can be restored to any point in time in a consistent and non-corrupt manner.
- c- If you employ a lot of automation within your application environment, ensure that your backup plan is integrated in your automation scripts to ensure that your SLAs remain unimpacted.

Below are two examples of backing up your cloud applications in the Cloud.

Active/Passive (Full Replica)

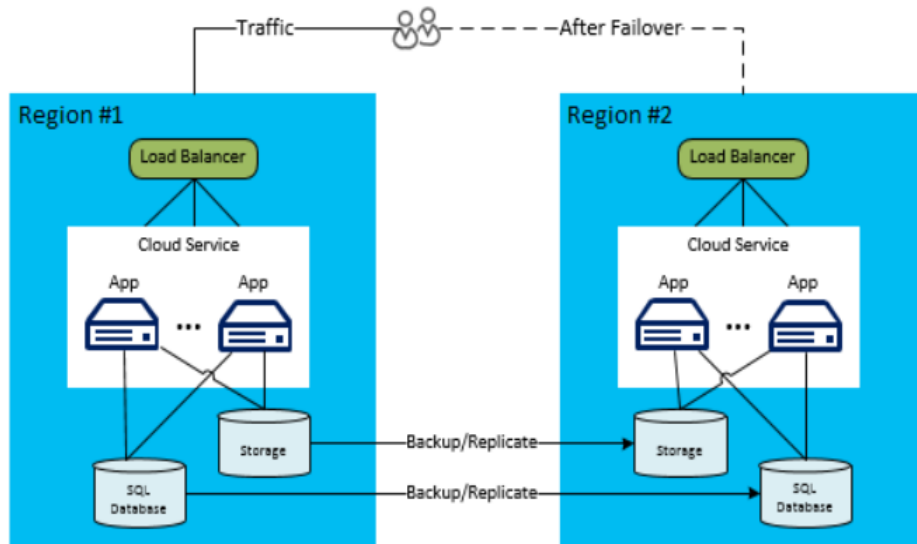


Figure 8: Active/Passive Cloud Application across two Cloud Regions

Active/Active




Figure 9: Active/Active Cloud Application across two Cloud Regions

3- DR for your Hybrid Environment

In present-day enterprises, cloud reliance is increasing by the day. However, with the restrictions of keeping some of the data on-premise, Hybrid Cloud scenarios play a major role in helping enterprises accommodate their scaling problems. One key scenario is leveraging Hybrid Cloud environments for Disaster Recovery, backing up the on-premise environment to a cloud data center.

When planning your Hybrid cloud DR, keep your focus on these key points:

- a- Is my backup tool hybrid-cloud ready?

You need to ensure that your primary backup application supports backing up your hybrid cloud management tool, as most backup tools are designed for either on-prem only, or cloud-only backups. Once you ensure this compatibility, your hybrid cloud management tool can leverage your backup tool for backing up your on-prem to Cloud.

- b- How do I ensure the security of my Hybrid cloud backups?

	ENABLED	DISABLED
Retention of deleted backup data	✔ Backup data retained for 14 days after delete operation	⚠ Instant deletion prevents recoverability from attacks
Minimum retention range checks	✔ Ensures more than one recovery point in case of attacks	⚠ Only one recovery point available for recovery
Alerts and notifications	✔ For critical operations like Stop backup with delete data	⚠ No security alerts or notifications for critical operations
Multiple layers of security	✔ Security PIN required for critical operations	⚠ Single layer of protection

Figure 10:Securing your Hybrid cloud Backups

When thinking about a Hybrid scenario for your DR, keep in mind that you will probably need to use additional tools to do the replication. You will not find any problems deciding on which tool to use, as each cloud vendor provide their own tools, or you may find vendor-agnostic tools like Veeam to help you with the replication. To give an example, refer to architecture in Figure 11 that explains how Azure Site Recovery is used for replicating your environment between your On-Premises data center and your Azure data center.

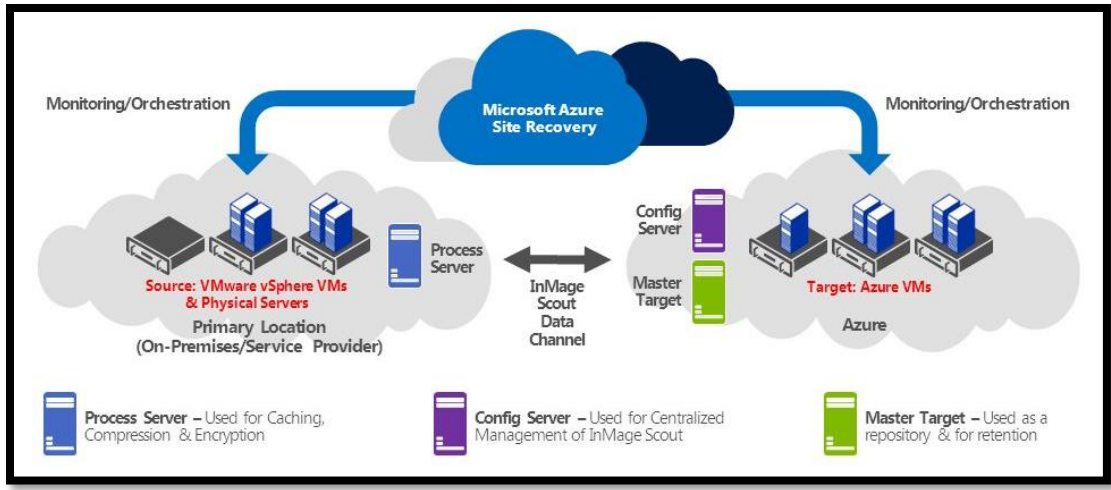


Figure 11: Using ASR from Hybrid Cloud DR Scenarios

4- Same Cloud DR between two regions

One of the easier options is to leverage your cloud vendors geo-distribution to set up your BCDR strategy. That way you can rely on having a single point of contact for all your support needs. This scenario may also be easier in terms on implementation and performance. This is due to the investment the cloud vendors have already made to ensure good latency between different sites and regions. There are many considerations to keep in mind when implementing this scenario, most of which are included in the table below provided by Microsoft for their Azure Platform.

Table 1: Considerations for Same-Cloud DR

	AVAILABILITY SET	AVAILABILITY ZONE	AZURE SITE RECOVERY/PAIRED REGION
SCOPE OF FAILURE	Rack	Datacenter	Region
REQUEST ROUTING	Load Balancer	Cross-zone Load Balancer	Traffic Manager
NETWORK LATENCY	Very low	Low	Mid to high
VIRTUAL NETWORK	VNet	VNet	Cross-region VNet peering

Once you have taken all the above into consideration and set your environment for same data replication, it is crucial to have a proper architecture of your environments. This will ensure you are leveraging all the right cloud services for the replication and that you have included all components of your application in your replicated version. Figure 12 provides an example architecture from AWS for a disaster recovery scenario between two cloud regions.

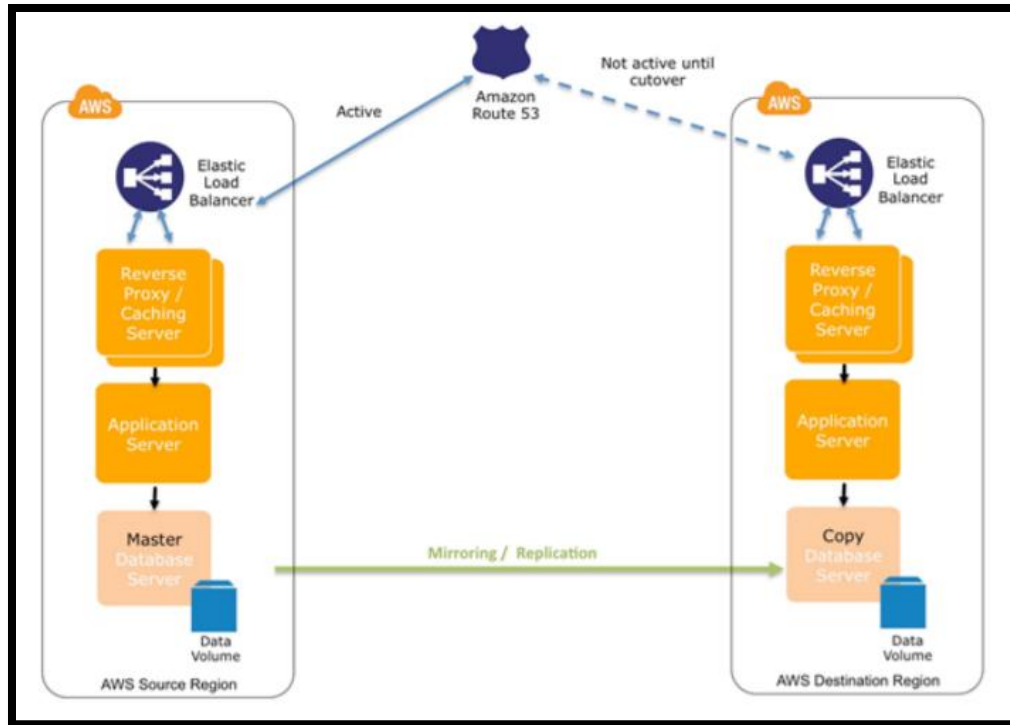


Figure 12:Replication between 2 AWS Regions

5- Multi-Cloud DR

As Vernal Vogel, Amazon CTO has said, “Everything Fails, all the time”. This scenario relieves you from relying on one cloud vendor for your BCDR strategy. So, you’re not just having a backup of your data on the cloud, but multiple clouds. Of course, such a scenario wouldn’t be used frequently due to its highly expensive nature, but just knowing that the option “Could” be available is re-assuring.

Though the scenario may be attractive to some companies with highly critical business applications, it may not be feasible all the time. Why, you ask? Simply because cloud vendors haven’t created a straight forward tool to integrate with one another. Therefore, this scenario would entail a lot of third-party tools and automation to achieve a true multi-cloud BCDR strategy. Also, keep in mind that it will be time consuming, and labor intensive to implement in an enterprise production environment.

Summary

When it comes to what you need to consider for your Cloud-based Backup Strategy, it is never a one size fits all scenario. The scenarios we have discussed above are merely a guide to some of the many paths that you can pursue for your cloud backup strategy. One thing for sure is that your Backup strategy cannot be complete without the cloud in today's world of Cyber-crime. As you may probably know by now, when it comes to securing your cloud environment, it is a shared responsibility to properly secure your environment. But opting for cloud means that you are offloading a big part of that responsibility to top cloud vendors that have heavily invested in security measures to help put you at ease knowing that your data is in good hands.

In the age of Cybercrime attacks, we notice that many cloud vendors have created cybercrime departments, with the sole purpose of identifying loopholes, and fortifying their data centers. Similarly, their investments are not just mitigating intentional attacks for secured data, but also planning for continuity beyond natural disasters, which means that by default most cloud vendors provide automatic replication of your data into multiple regionally located datacenters. This enables you to save on your capital investments of replicating your local data centers.

Of course, not all clients are for the idea of having their data hosted on the cloud, usually for reasons like data sensitivity, Geo-Political regulations and other sorts of legalities. Therefore, after going through the multiple scenarios above to help you design a solid BCDR strategy for your business, the recommendation is to leverage the best of both worlds, i.e. keep your sensitive data locally in your data centers and leverage the cloud features for all your non-sensitive data. With today's cloud networking features, you probably won't notice the difference of an extended data center in the cloud, but you will definitely notice the increased power and efficiency of your business.

References

<https://pl.atyp.us/tag/humor.html>
Backup consistency - a game of challenge.
<https://www.emc.com/collateral/whitepaper/recovering-business-destructive-cyber-attack.pdf>
https://www.pubnub.com/static/papers/loT_Security_Whitepaper_Final.pdf
<https://blog.dell EMC.com/en-us/dells-new-air-gap-security-solution-keeps-sensitive-data-airtight/>
<https://docs.microsoft.com/en-us/azure/backup/backup-azure-vms-introduction>
<https://docs.microsoft.com/en-us/azure/architecture/resiliency/disaster-recovery-azure-applications>
<https://docs.microsoft.com/en-us/azure/architecture/resiliency/disaster-recovery-azure-applications>
<https://docs.microsoft.com/en-us/azure/backup/backup-azure-security-feature>
<https://slideplayer.com/slide/8079156/>
<https://docs.microsoft.com/en-us/azure/architecture/resiliency/>
https://media.amazonwebservices.com/AWS_Migrate_Resources_To_New_Region.pdf

Figure 1: Zero-day threats.....	11
Figure 2 :Isolated recovery solution architecture.....	13
Figure 3: Shared switch option	15
Figure 4: Dedicated Switch option.....	16
Figure 5: Firewalled vault option	17
Figure 6: Appliance option.....	18
Figure 7: Azure VM Backup Architecture.....	20
Figure 8: Active/Passive Cloud Application across 2 Cloud Regions.....	22
Figure 9: Active/Active Cloud Application across 2 Cloud Regions.....	22
Figure 10:Securing your Hybrid cloud Backups	23
Figure 11: Using ASR from Hybrid Cloud DR Scenarios.....	24
Figure 12:Replication between 2 AWS Regions.....	25

Dell Technologies believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." DELL TECHNOLOGIES MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Use, copying and distribution of any Dell Technologies software described in this publication requires an applicable software license.

Copyright © 2019 Dell Inc. or its subsidiaries. All Rights Reserved. Dell Technologies, Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be trademarks of their respective owners.