COMMVAULT

WHITE PAPER

Safeguarding the Digital Enterprise

This document is intended to provide general guidance for a non-technical audience on how various digital asset protection methods work and the need for backup copies of the assets whether working in the cloud, on premises, or in a hybrid digital asset infrastructure.

Executive Summary

As organizations and individuals increasingly rely on cloud-based services to store, manage, and share digital assets, it is essential to recognize that these services, despite their numerous benefits, are not infallible. This white paper emphasizes the importance of backups for digital assets, even when replicated or hosted in the cloud, and offers guidance on creating a comprehensive data protection strategy to protect valuable data from loss, damage, or theft.

If Data is the new Oil, can I run my car on it?

In the digital enterprise, data is the lifeblood that powers communications, collaboration, and customer experiences and fuels modern business outcomes – from production lines to financial services to schools, hospitals, and government services. The increase in essential technology woven into our daily lives continues to increase our daily data dependency; if the availability, flow, and access to data is significantly interrupted or wholly compromised and lost – all technology-enabled activities stop cold. This creates a critical dependency on the data flow, which many associates with the "new oil" or lifeblood of the digital enterprise, workspace, and community. Leading organizations recognize these risks and build safeguards and redundancies to protect the data assets to minimize the risks of significant impacts and outages for that service.

Unlike other vital asset classes with high-demand/limited supply, such as precious metals or medical blood supplies, digital data assets are constantly growing in demand and supply as innovations tech-enable everyday life experiences. Unlike gold bars or coins with a well-accepted standard value, the value of your data can vary widely and fluctuate based on needs. The data-value equation is shifting daily, demonstrated in the headlines where criminal Cyber events can carry initial ransom incident costs exceeding $1.4M[1]. When the organization cannot quickly respond, remediate, and recover business services, the overall impact

While you may not be able to run your engine on Data, each Autonomous vehicle is estimated to generate over 4TB[2] of data per day from all the sensors and real-world context, comparable to the daily email content generated by a workforce of 250K users.

1. https://www.csoonline.com/article/3276584/what-does-a-ransomware-attack-cost-beware-the-hidden-expenses.html#:~:text=The%20average%20cost%20of%20a,%24733%2C000%20for%20the%20non%2Dpayers

2. https://premioinc.com/pages/autonomous-vehicle-data-storage#:~:text=Intel%20estimates%20that%20an%20autonomous,3%2C000%20in%20a%20single%20day.

can soon exceed $100M, resulting in major impacts on the external community as hospitals, schools, and government agencies go dark.

Figure 1 Key stats and Shifting Executive focus on Resiliency/Recovery

Given that data is an asset and a liability depending on how it is managed and utilized, in the cyber climate that we are in today, focus has shifted towards ensuring that the liability of data in terms of Privacy and security risks is mitigated as best as feasible. One of the key areas is by ensuring a cyber-resilient architecture is deployed and implemented with Zero-trust security principle, adhering to the philosophy of defense in layers, bolstered by a robust data protection solution that can protect and aid recovery of data at scale in the event of a crisis, Natural Disaster or Cyber. Undoubtedly, data protection serves as the essential “insurance policy” on which enterprises depend.

A comprehensive data strategy needs to balance risks, resources, and objectives.

While these risks can’t simply be dismissed as risk rejection is not an option in a dynamic and competitive environment where risks are inherent, they can however be mitigated with intelligent planning, operations, and a recovery-readiness strategy. This strategy acts as a set of guidelines to assess and classify the Applications and data assets that power the business services, which align with quantified service level objectives. This allows the operation plan to properly balance investments in people and resources to ensure the highest impact risks can be effectively mitigated.

The Cybersecurity and Infrastructure Security Agency (CISA) states “Maintain offline, encrypted backups of data and regularly test your backups.” Source: Stop Ransomware | CISA

The Zero-Loss philosophy recognizes the fact that data can’t be managed like the gold bars – considering the scale of data change rates, data volumes, scale and costs a universal plan would demand massive investments, infrastructure and create an unsustainable outcome.

Rather this strategy presents a guiding principle allowing organizations to establish a common understanding on a degree of potential loss – which may be realized as downtime, repeated actions (lost transactions or latest changes), or degraded services that follow a failure or data compromise – to align against Service Level Agreement (SLA) protection objectives that balance security, retention, scale, and recovery performance.

Today, the older principle that helped us reach the moon, "failure is not an option”, it is an accepted reality that must be planned into the strategy to ensure we can get back home.

Whether the damage is done by a natural disaster, hardware failure, user mistakes, data breach, or ransomware attack, you need to recover your data quickly with zero data loss and as minimal downtime to business as possible. While cyber and disaster recovery characteristics differ, the outcome remains the same: recover clean data quickly. Many perceive cyber recovery as another form of disaster recovery – the methodology of recoveries is similar, using data protection and recovery tools. Still, the salient difference remains in the remediation as the assurance of sanctity and business readiness of the recovered data are key facets especially in the cyber context. Cyber corruption can demand more surgical investigation and data curation actions to repair and sanitize your system before it can be returned to full-service, also eliminating any potential risk of reinfecting the environment.

Salient differences between Cyber Recovery and Disaster Recovery - ***https://www.commvault.com/blogs/the-path-to-recovery-readiness***

One needs to be mindful when narrowing the scope of a protection/recovery exercise exclusively to the “crown jewels”. Business Impact Analysis (BIA) is keystone effort, performed to help identify the mission criticality of data and applications to make the appropriate investments in its protection and recovery strategies. Not all data is created equal, hence understanding sensitivity of data and importance of the digital assets is core and foundational element towards designing a cyber resilient architecture. https://www.mckinsey.com/~/media/mckinsey/email/classics/2023/2023-03-18c.html

In simpler terms, Data Protection strategies which include both Backup and recovery of data, refer to the procedures involved in creating and storing data copies to safeguard organizations against data loss. The recovery process usually entails restoring the data to its original location, enabling employees and business to resume their work seamlessly. In more significant incidents, the data may be restored to an alternative location to serve as a substitute for the lost or damaged data. In cyber crisis context, this alternative location also serves as a location for data forensics to assess the usability of data by business before it can be presented back.

Shifting to the cloud does not shift the responsibility for data resiliency and recovery; it always remains the core responsibility of the user organization.

Cloud platforms can be credited with a paradigm shift allowing IT to provide agile services to the Application developers and owners based on preconfigured, on-demand services consumed from an elastic resource pool. Modern applications and data services offer a practical way to balance the needs of IT and DevOps to help drive faster, successful transformation jointly. These services allow designers, recommenders, and operations providers to deliver more dynamic infrastructures that can be aggregated to promote resource sharing and improve efficiencies while aligning different resource classes to match business SLA outcomes. One of the most potent advantages of operating from inside the Cloud is the openness and accessibility, allowing organizations to quickly readapt how they conduct and support business processes and user communities. Like our example referencing the explosion of data associated with adopting Autonomous Vehicles, cloud-enabled organizations also witness skyrocketing growth rates in data creation and consumption. The elastic cloud eliminates numerous scalability issues that previously constrained earlier generations relying on scale-up architectures and capital expenditures (Capex) for IT infrastructure. Instead, it delivers agility through cloud compute and storage elasticity in an Operating expenditure (Opex) model with ability to grow to unlimited mailboxes, and continuous data feeding and capturing for analytics applications and right-size based on fluctuating growth patterns.

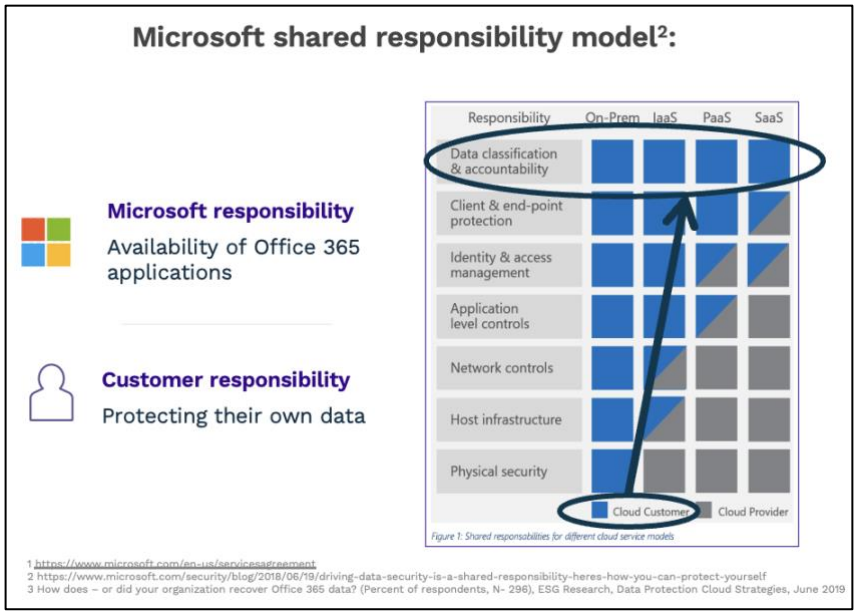

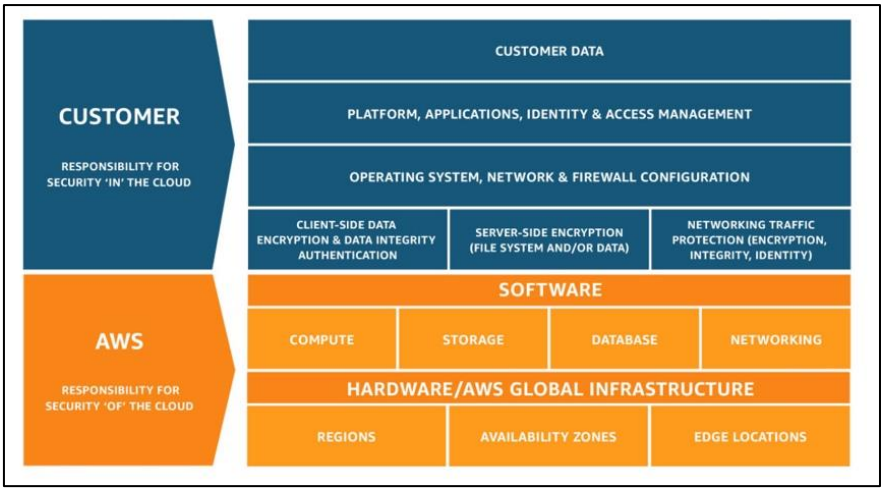

However, considering all the benefits of elasticity and agility that cloud platforms provide, a very common misconception is that the Cloud takes care of everything. Still, many realize that a crucial criterion for success in the cloud is driven by an understanding of the responsibility matrix, which continues to place the responsibility of managing and maintaining data entirely in the hands of the user. The cloud providers assume responsibility for the cloud resource services primarily focused on availability and durability to offset hardware or site loss. While the responsibility of data's resiliency, integrity, and recoverability sit firmly on the shoulders of the user's IT organization.

However, high availability and data resiliency often misunderstood as data protection.

1. **High availability** refers to the ability of a system or infrastructure to remain operational and accessible to users, even in the face of unexpected failures or disruptions. This is achieved through redundant components and failover mechanisms with replication to create data redundancy that can take over seamlessly in case of a failure. High availability ensures that users can access data and applications without interruption or minimal disruption.

- By that same token, Data Availability refers to the accessibility and readiness of data for use. Data availability ensures that the data is accessible to authorized users or systems whenever it is required.

2. **Resiliency**, on the other hand, refers to the ability of a system or infrastructure to recover quickly from a failure or disruption and resume normal operations. Resiliency is achieved through disaster recovery mechanisms, backups, and other strategies that allow the system to recover from a failure and resume normal operations with minimal downtime. This is not a standalone metric as it encompasses business continuity and recovery techniques to reduce the impact of disruptive events.
3. **Durability** is a critical component in today's cyber context as it refers to the continued persistence of data where the expectation is to ensure that the preserved data does not undergo data loss or corruption in any shape or form. Organizations take steps to enhance the durability of data by engaging in **immutability** and **erasure coding**.
4. **Data protection** refers to ensuring that data and changes/versions are secure, confidential, segmented, and available when needed. This includes controls and processes such as encryption, access controls, backup and recovery, and disaster recovery planning.

Recovery objectives must be balanced with access, response, retention, cost, and security.

Cloud providers use data replication methods and hardware redundancy to ensure data durability and availability by creating multiple copies of data and distributing them across different locations. This ensures that data remains available and accessible even in the event of hardware failures or other disruptions.

Cloud providers typically offer different levels of replication and redundancy depending on the needs and requirements of their customers. Some providers may offer multiple levels of replication and redundancy for different tiers of service, allowing customers to choose the level of data durability and availability that best meets their needs.

Most Cloud platforms provide built-in data protection techniques that rely on cloud-storage based replication with the option to choose snapshot-based replication, which are not without its own limitations and challenges.

Critical Challenges and Concerns with built-in Data Protection provided by Cloud platform vendors:

1. Cloud storage-based replication carries the inherent risk of replicating data corruption inadvertently across regions without validation of data which is further exacerbated by the fact that most cross geographical replications are asynchronous in nature. The asynchronous method replication of data does not guarantee that all the changes in the production region have been replicated in time to aid recovery for unplanned events.

    a. This puts the onus on the customer to ensure that there are smarter replication methods in place using 3rd party data protection solutions deployed in cloud, that ensure availability of data for business continuity purposes, providing the much-needed recovery assurance in case of any eventuality. This includes ability to restore at scale with capabilities around data sanitization and scrubbing before restoring into production.

2. While snapshots can be beneficial for recovering applications with aggressive Recovery Time Objectives (RTOs) and Recovery Point Objectives (RPOs) driven by business needs, relying on snapshots for long-term data protection requirements is not advisable. There are several reasons for this:
    a. Unpredictability of Cost of Operations in cloud - Snapshots can be costly and contribute to the accumulation of "static waste" when retained for extended periods. In cloud computing, the billing model operates on a pay-as-you-go basis, which can introduce the potential risk of encountering unexpected financial costs, commonly referred to

as "sticker shock." Though cloud platforms provide a wide array in choice of services, the cost associated with each of the services and its various options can result in reduced predictability of cloud costs over time. https://www.cio.com/article/403231/cios-contend-with-rising-cloud-costs.html

b. Operational Overheads - The process of recovering data from snapshots requires multiple steps to ensure transparency with the data stored within them. These steps often involve manual efforts and extensive scripting, which can lead to the accumulation of "technical debt" and a higher risk of human error.

What are critical criteria for selecting the right solution for protection of cloud assets?

1. **Simplification:** Solution that provides enhanced data management capabilities through automation and consolidation, supported by comprehensive reporting and a unified management interface, that can streamline operations and align them with business Service Level Agreements (SLAs).
2. **Risk reduction:** Mitigating risk by reducing downtime and data loss, improving recovery speed, and providing litigation support,
    - Cost predictability - keeping the risk of spiraling cost of operations under control.
    - Security – solution should provide ability to integrate with security ecosystem tools and have built-in security capabilities for enhancing overall security posture.
3. **Productivity gains:** Enabling gains in terms of both tactical operations and strategic planning, with intuitive approach to data protection and recovery at scale in cloud, that eliminates “guesswork” out of data protection.

Reducing Technical Debt by consolidating to a more unified and integrated approach

Today, many organizations still struggle to scale the infrastructure silos created to quickly address the initial wave of transformation as shared storage environments were introduced and tape-based backup settings were replaced with purpose-built backup devices. Increasingly many now recognize this as a form of increasing technical debt. They need to consider the increasing burden on the cost of operational time, imposed fragmentation (due to scale limits), and reliance on manual efforts to realign the environment to avoid failures. Technical debt measures the implied costs caused by choosing the easy (limited) short-cut solution rather than carrying a much higher burden set of constraints over time.

> Technical debt refers to the consequences of taking shortcuts or quick-fix solutions during software development, which can lead to suboptimal or inefficient code, resulting in maintenance issues, bugs, and other problems. It accumulates over time as teams prioritize speed over quality and can become increasingly difficult and expensive to address. Addressing technical debt is crucial for maintainable, scalable, and efficient systems.

This challenge can be characterized by the variety of data management tools and storage resources an organization manages, increasing complexity, and overall operating costs as more diverse expertise needs to be maintained. Storage resources that cannot aggregate horizontally into scale-out clusters increase the general waste or excess resources required to handle unexpected growth. Operation teams must constantly battle increasing consumption and shifting workloads to different infrastructure silos to avoid overrun conditions. This leads many to become extremely selective and restrictive on the scope of the data protection SLA to adapt to the limitations of the protection infrastructure; this may, in turn, reduce the frequency of the protection operations (Recovery Point Objectives) or more narrowly focus on which data set is protected.

Considering the heightened cyber risks of today, any operational reduction in the scope of data protection can have disastrous consequences for an organization. If critical applications cannot be recovered due to limitations imposed on the protection scope and challenges in scaling data protection operations, the impact can be severe.

When architecting data protection strategy, it is recommended to adopt and implement the following best practices:

1. **Perform Business Impact Analysis** - Identify critical business processes and the technologies that businesses rely on to ensure comprehensive protection of systems and data. Understanding the criticality of the applications helps answer key questions around operationalizing data protection – such as the following:
    a. what should be backed up?
    b. how often should the application be backed up?
    c. how quickly does the business want to recover the data?
2. **Identify Recovery SLAs** - Determine the Recovery Time Objective (RTO) by assessing how long your business can afford to be without critical processes. RTO (Recovery Time Objective) answers the question "How long can an application or system be down before it starts to impact business operations?" In other words, RTO represents the maximum acceptable downtime for a business, after which the system needs to be restored to minimize the impact on operations. It represents the target time for restoring the system after a disruption or disaster.

    Subsequently, business also needs to define the Recovery Point Objective (RPO) by establishing the acceptable level of data loss. RPO (Recovery Point Objective) answers the question "How much data loss can a business tolerate in the event of a disruption or disaster?" In other words, RPO represents the maximum acceptable amount of data that a business can afford to lose before it starts to impact business operations. It indicates the point in time to which data needs to be recovered to minimize the impact of the disaster on the business. Recovery Point Objective is directly correlated with the frequency of backups, as it is dependent on the specific requirements of the protected application.

    Now equipped with the above information, next logical step is to map the expected RTO and RPO to the applications based on their criticality.

Application	Criticality	RTO (Recovery Time Objective)	RPO (Recovery Point Objective)	Planned Data Recovery Drills for testing. (Cyber or Disaster scenarios)	Assessed Impact of Downtime on Business operations
Email system	High	1 hours	15 mins	6 times/ year (*Recommended)	$$$$
Legislative Management System	High	2 hours	1 hour	6 times/ year (*Recommended)	$$$$
Financial Management System	Medium	8 hours	4 hours	4 times / year	$$
Human Resources Management System	Medium	8 hours	4 hours	4 times / year	$$
Time and Attendance System	Low	24 hours	8 hours	2 times / year	$

Note: the above table is a sample of Applications mapped to their criticality and recovery SLAs of RTO and RPO. The actual values of RTO/RPO values may vary depending on the specific needs and requirements of the environment.

In the example in the table, the applications are categorized according to their criticality level, with High being the most critical and Low being the least critical. The RTO represents the maximum amount of time that the application can be down before it starts to impact business operations, while the RPO represents the maximum acceptable data loss in the event of a disaster. This table helps organizations prioritize their disaster recovery efforts by identifying which applications are most critical to business operations and how quickly they need to be restored in the event of a disaster.

Time and Attendance system which has an RPO of 24 hours will be backed up once in 24-hour cycle, while Email system which has an RPO if 15 mins would require more frequent backups, as often as every 15 minutes, to ensure all transactions are protected.

3. **Data Protection strategy** - Implement replicas of protected data by replicating them to another physical location or a secure cloud storage. This ensures you have multiple copies of backups for reliability. There is significant adoption of cloud storage as an option for secure and vaulted backup copies of data that aid in Disaster recovery events use cases.

   Cloud storage provides a reasonably affordable, resilient, and highly available storage target that can scale elastically without the need for purchasing and managing storage arrays. However, it is crucial to carefully consider the the use case fitment when leveraging cloud storage, to arrive at the most optimal way to address rapid recovery requirements from Cloud storage back to on-premises.

   Additionally, a rapidly evolving market is also seeing an adoption of cloud object storage for storing long term compliance requirements at more cost optimized cloud storage options. This is being viewed as data modernization with the replacement of Tape infrastructure with cloud archive tiers of storage, such as video recordings; Tape backups were popular in the early 2000s, using magnetic tape devices to store large amounts of data. However, management of Tapes can be difficult and are prone to data read errors due to degradation of tapes over time.

   In other words, a future-proof Data Protection solution should be capable of supporting various storages options including disk storage, tapes, and seamless integration with cloud storage providers such as AWS, Azure, OCI, GCP, and others.

   Improved Storage and Performance Efficiency - Backups frequently contain duplicate data, which can lead to inefficient storage utilization. A high-quality backup solution should include the ability to identify and remove duplicate data, retaining only unique information to minimize storage usage. The solution should be capable of effectively deduplicating all the digital assets and be compatible with various storage types. Deduplication helps optimize storage space consumption but also contribute to more efficient backup and recovery processes, ultimately reducing costs and enhancing overall performance. Additionally, adopting an incremental-forever backup model allows for optimal backup efficiency by avoiding the need to repeatedly back up unchanged data.

4. **Improved Security of Protected Data** - Enhance security posture by implementing Multi-Factor Authentication (MFA), Multi-Person Authentication, Privilege Access Management Solutions with tight implementation of Role Based Access Controls (RBAC) in conjunction with other controls to safeguard access to backup data. Backup data are often targeted by cybercriminals during ransomware attacks, as it has crippling effect on an organization by compromising backup data and their overall ability to recover from it.

   Data encryption is crucial for enhancing the security of data. When choosing a data protection solution, it's important to select a product that supports data encryption of data both in flight and for data at rest. The solution should possess the ability to support multiple encryption algorithms, ensuring the highest level of security for stored data. Additionally, it should include provisions for re-encrypting backups in case the organization adopts a new encryption standard in the future. This ensures that data remains protected and aligned with evolving encryption requirements and standards.

In addition to these, having the ability to **audit** actions and access with non-repudiable methods are core to securing the data protected copies of data.

5. **Testing and Updating Data Protection Plans and runbooks**: Finally, regularly test backups and recovery processes to ensure they work effectively. Conducting restore tests frequently based on criticality of the applications to verify the functionality, availability and business readiness of data and reliability of the backup system, gives the much-needed recovery assurance, with reliable estimate of the application's recovery time in the event of a disaster. Ensuring that these operational challenges are overcome with the power of a strong **automation** framework is key. Robust automation framework ensures minimizing human-prone errors and helps scale the operations as workloads increase.

In conclusion, implementing an effective data protection strategy is vital for safeguarding an organization's digital assets and ensuring swift recovery in case of data loss or disaster. This entails creating a comprehensive inventory of all digital assets, selecting a high-quality backup product that offers a range of features such as various storage targets, data encryption, and deduplication, and establishing a robust recovery testing plan. By incorporating these components, organizations can successfully protect their valuable data and remain well-prepared for any potential data-related challenges.

Appendix

The Importance of Backups for Digital Assets:

Protection Against Data Loss: Data loss is not just losing a file or an email, it can cause entire operational systems to fail. This could result in the inability to communicate using email, messaging systems such as MS Teams, or VoIP phone systems. Using high availability protection methods like replication provides some protection, but backup copies are essential in mitigating risk against data loss.

Human Error - Protecting digital assets is only as good as the weakest link.

Sometimes that can be the people that manage the assets. In a report by the World Economic Forum, it states "Businesses also operate in a world in which 95% of cybersecurity issues can be traced to human error, and where insider threats (intentional or accidental) represent 43% of all breaches" Source: WEF_The_Global_Risks_Report_2022.pdf (weforum.org). Having secure backup copies of digital assets greatly mitigates the risk of data loss due to human error.

Human Caused Disasters:

In the case of human caused disasters, we are not talking about human error but rather deliberate attacks. A disgruntled employee with intimate knowledge of digital assets, terrorist attacks, or state-initiated acts of war, should all be considered. Additionally, EMP, nuclear or dirty bombs could affect not only your digital assets but communication systems to access the assets.

Cyberattacks:

Ransomware attacks and other types of data breaches not only affect digital asset access but can cause critical operational systems to fail. This can result in the complete inability to access data such as files and email but also cause the inability to continue business operations. In May of 2021, Colonial Pipeline was hit with a ransomware attack. Operations were down for six days.

There are three major risks regarding cyberattacks:

1. Inability to access data.
2. Inability for operational assets to function.
3. Data exfiltration.

Backup copies provide a virtual and, in some cases, physical barrier making it difficult or impossible for a cyberattack to compromise the data. Backups also allow for longer retention beyond a ransomware's incubation period. Additionally, encrypting backup data eliminates the risk of data exfiltration from the backup copies.

Hardware and Software Failures:

Even in the cloud, hardware and software failures can happen. It is a shared responsibility between an organization and a cloud vendor to provide availability and disaster recovery. Backup copies protect data beyond normal replication. Since hardware and software failures can occur, even in the cloud, backups are critical. Along with shared responsibility is shared trust.

Natural Disasters:

Disasters such as hurricanes, tornados, tsunamis, or major floods are the common types of natural disasters. For a comprehensive disaster recovery plan, disasters including meteor strikes or solar flares should be included. These disasters range in size and scope but could affect digital assets, specifically communication. In these scenarios, the ability of the organizations to effectively communicate with citizens or customers will be critical.

Salient differences between Cyber Recovery and Disaster Recovery:

Whether the damage is caused by a natural disaster, hardware failure, data breach, or ransomware attack, you need to recover your data quickly with zero data loss and as minimal downtime to business as possible. While the characteristics of cyber and disaster recovery differ, the outcome remains the same: recover clean data quickly. Many perceive cyber recovery as another form of disaster recovery – the methodology of recoveries is similar, using backup and recovery tools, but the salient difference is that in cyber recovery use cases, the validity of data is in question so performing scrubbing of data before restoring to is critical to avoid any risk of reinfection.

Ensuring Business Continuity:

Using protection methods including replication, snapshots, and backups is a strong digital asset protection strategy. Business Continuity is more than just a protection strategy. It is a comprehensive plan for the continuation of business operations in the event of a disaster. Being able to recover a lost email is not the same as being able to continue to send and receive email when a disaster strikes. **Communication will be essential in these types of events.**

It is critical for your organization to have the expertise, resources, and knowledge on how to respond to a breach, cyber-attack, or disaster to maintain business continuity. You cannot afford to wait for an incident to occur to figure out how your organization will respond and determine what to do during a security incident. According to Rob Joyce, director of the US National Security Agency's cyber security arm, "the rest of the world should take the Russia-Ukraine war to heart and learn how Ukraine has learned to keep

its critical infrastructure online and running." This is from an article titled "Why a top US cyber spy urges: Get religious about backups." Source: Why a top US cyber spy urges: Get religious about backups • The Register

Having a sound multifaceted business continuity plan during a disaster is essential. This includes recovering digital assets from replicated copies and snapshots, if available, and backups when required. While recovery operations are ongoing, effective communication is key. This will include public communication through any means available as well as internal communications. Consider a ransomware attack that takes down all email communication. What alternative communication methods are available? Are they secure? Having a sound business continuity strategy will result in all of these issues being proactively resolved.