EVOLVING SOLUTIONS

POINT OF VIEW

Data Protection in a Cloud-Connected World

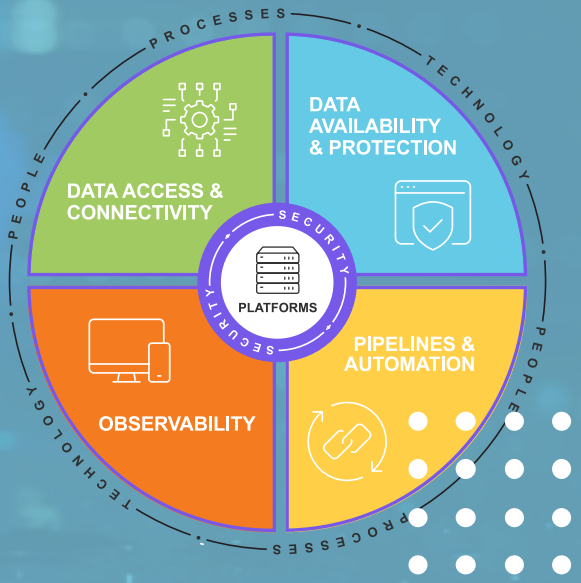

Environment and workflow are at risk for cyber attack

Once largely a “set and forget” proposition, data protection now requires a well thought out strategy that prioritizes the business value of various data sets. In the past, from a data protection perspective, all data was handled the same way. Servers were backed up to the same set of tape with the expectation that everything would be recovered at about the same time.
This is a risky situation. In the event of a cyber-attack, getting critical business functions up and running quickly is crucial to business continuity. But restoring the entire environment from tape may take days or weeks.

Today, data is being used for a wider variety of purposes. That means today’s data also has varying degrees of business value. For this reason, data recovery is becoming a more important discussion, particularly in terms of prioritizing data recovery and having a strategy for recovering high-value data more quickly than lower value data. For example, a revenue generating application needs to be back online faster than historical data being used for reporting.

In addition, recovering data in the cloud may be different than recovering it on-premise.

Understanding the specific requirements for different data sets can impact the kinds of data protection and recovery solutions you need to consider.

In an increasingly mobile, cloud- and internet-connected world, where loss of data integrity can take many forms -- such as corruption, unintentional deletion, and cyberattacks -- organizations are prioritizing how they protect and back up their crown jewels.

In this environment, it is critical to develop a data protection strategy designed to support business continuity. As you develop a strategy, keep these key questions in mind:

- What environments do your workloads reside on, and what are the resiliency needs for these environments?
- Do you have a strategy in place that protects not just the data, but the workloads dependent on that data; regardless of their design?
- Does your resiliency plan address availability, data protection, reliability and recovery for on-premise, off-premise and hybrid workloads?

Case Study

Property, Casualty, and Life Insurance Company

CHALLENGE:

Modernize an aging infrastructure with an outdated data protection solution to meet scalability, performance, and recovery requirements while maximizing investment value.

HOW WE HELPED:

- Conducted a comprehensive assessment of the existing data protection solution and identified infrastructure bottlenecks, including aging backend storage.
- Provided an in-depth analysis of the client's network design, uncovering root causes of performance issues and offering actionable recommendations.
- Proposed a cutting-edge solution featuring a greenfield deployment to allow for a parallel migration to the new, optimized environment.
- Explored the physical and logical network topology—becoming the only partner to deliver such a thorough evaluation.
- Delivered technical expertise in best practices, design, application, and support, showcasing the capability to partner effectively across vendors.
- Supported the client in vendor and partner evaluation by offering clear metrics and guidance tailored to their operational needs.

RESULTS:

- Delivered a robust framework for implementation and design of storage solutions and complementary technologies.
- Established a scalable repository for enterprise-grade data protection.
- Improved scalability, performance, and backup speed.

Environment Matters

The location of data, whether it be on-premise, in a public cloud, fully cloud-native, or part of a hybrid cloud environment, directly influences your data protection strategy.

For example, if you leverage microservices as part of a cloud-native application approach, your data and services may be spread across multiple clouds and multiple zones. As a result, you must use code (or scripts) to recreate infrastructure and reestablish data connectivity to get the same resilience you might have in a traditional environment. Additionally, the code used to provision your infrastructure must also be protected and available. In short, your environment directly impacts your data protection strategy.

In an increasingly hybrid cloud and cloud-native world, it is paramount to remember to protect not just the data in your environment, but also the data "of" and "about" your environment.

Setting Recovery Objectives

Setting recovery time objectives (RTO) and recovery point objectives (RPO) for data should be based on the business value and compliance requirements of the data you are protecting. RTO is the period of time it takes to recover a function. RPO is a measure of how much data you are willing to lose. For example, if you are backing something up every 24 hours, that indicates that losing a day's worth of data is acceptable.

To establish RTO and RPO for a particular data set, you can group data into two buckets. Data "of" is data of record, which, for a bank, is its most important asset and includes crucial information such as bank balances. This is high-value business data that needs to be accurate and current at all times, so the RTO needs to be as close to zero as possible.

Data "about" is information that is not supporting a critical business function, for example, information about accounts and where and how they do their banking. This data involves much less short-term business risk if lost or corrupted, so an RTO of hours or days may be acceptable.

Whether your data is on-premise, on a public cloud, fully cloud-native, or part of a hybrid cloud environment, these factors directly influence your data protection strategy.

Data “of” must be 100% accurate, so it is important to have that data’s business system back in place with the last known good copy of data. Data “about,” on the other hand, does not need to be recovered immediately to ensure business continuity.

Create an Agnostic Data Protection Workflow

The increasing prevalence of public and hybrid cloud changes the way on-premise infrastructure and workloads function. Public cloud workloads and cloud-native behavior rarely translate perfectly back to traditional monolithic, multi-tiered, on-premise infrastructure.

Pervasive cloud adoption is driving customers to treat their on-premise infrastructure as private clouds. This provides the agility of a public cloud, and because the cloud can be effectively leveraged in backup and recovery strategies, enables workloads to more gracefully transition between environments.

However, this trend -- the integration of on-premise infrastructure and public cloud -- calls for developing an agnostic data protection workflow that looks the same on-premise and in the cloud, whether private or public.

An API-enabled automation pipeline can help with this, reducing workflow complexities required to create an effective data protection environment. For example, if you are rolling out a new application, or your data protection needs change, you need a workflow that automatically makes these updates or changes to where your organization’s backup lives. Data protection strategies contain too many elements for a single human to track, effectively making automation a necessary component of your plan.

Security Measures to Prioritize in Hybrid Cloud Environments

When prioritizing security measures for data in a hybrid cloud environment, consider implementing a process for having a history of backups that can be evaluated for malicious software and always having a last known good copy of data to restore from. It is also important to evaluate immutable copies -- copies that cannot be changed, and indelible copies -- copies that cannot be deleted. It is not unusual for a malicious actor to delete a backup that cannot be changed, so you need protection against both change and deletion.

Another security measure to consider is having a policy where more than one administrator is required for certain delicate operations such as changes to the backup schedule or deletion of backups.

Backup and Recovery in Hybrid Environments

If your organization backs up its on-premise data to the cloud, it is important to consider where the workload will be recovered to. If it is coming back to the data center, you need to consider if you have enough bandwidth to handle the job. There also may be transformation considerations. For example, if you are backing up a VMware environment to AWS, the VMs may need to be converted to EC2 instances before they can be recovered. It is also important to ensure that your cloud backup is secure. The cloud provider does not do that for you.

Ideally, backup and recovery across clouds and on-premise is a cohesive, integrated process using the same software to handle both environments.

Unfortunately, there are very few vendors that support this.

Organizations need to focus on systems integration to create their own cohesive backup and recovery process. For example, if you are running Kubernetes in the cloud that is using data stored on-premise and you need to recover the entire application, you potentially need two pieces of recovery software to handle that. Organizations with this environment will need a way to integrate and orchestrate the process on both platforms. At this point, this is not being completely solved by AI, so it requires an integration strategy built on an understanding of the dependencies between the two environments.

Enforcing Security Policies Across Environments

To ensure security policies are enforced across both cloud and on-premise environments, organizations must have the proper tools to implement controls in these environments.

Additionally, using traceability is key to confirming that security policies are properly enforced in the software being used.

Security Automation for Data Protection

Automation plays a role in data protection and resilience, because once an intrusion is detected, the clock is ticking — you must respond as quickly as possible. If you have a more advanced cyber resilience environment that detects suspicious activity, automation can transmit the alerts to a centralized logging and notification system.

In addition, security orchestration, automation, and response (SOAR) is a useful framework for automating responses using pre-established automation playbooks, such as automatically starting an Ansible script to isolate certain a high-value environment or to take a snapshot of your transactional data, so in the event that a bad actor does get through, you have a copy of data to compare and contrast with known good copies.

Case Study

Client: Telecommunications Company

CHALLENGE:

Develop a comprehensive business continuity and disaster recovery plan tailored to the unique needs of each business unit to ensureoperational resilience and clarity across the organization.

HOW WE HELPED:

- Conducted thorough top-down audit to evaluate the impact of outages on various business units; engaged directly with end-users to gain application-specific insights.
- Interviewed stakeholders to gain new perspectives, refine requirements, and define effective disaster recovery strategies to ensure recovery practices align with business needs.
- Analyzed daily, weekly, and monthly workflows to define success metrics for disaster recovery and business restoration across the organization.
- Addressed critical business functionality restoration with processes that enable swift recovery of operations.
- Design scalable strategies to provide immediate value while allowing for future expansion of resilience efforts.

RESULTS:

- Enhanced resilience and preparedness so the organization can respond effectively to future disruptions.
- Delivered a reusable disaster recovery planning template for consistent implementation across business units.
- Developed comprehensive collateral to support business unit-specific recovery and continuity efforts.

4 Steps to Increase Data Protection and Guard Against Data Breaches

Although ransomware is a hot topic that often captures public attention, there is much more to data protection than guarding against ransomware alone. Data availability challenges can take many forms — from inadvertent data deletion to unexpected downtime.

Fortunately, there are numerous approaches your organization can use to mitigate data corruption. Here are four steps to help increase data protection and guard your organization against data breaches. These steps work best in combination, creating a comprehensive approach:

1. **Implement strong security policies:** The need for consistent, strong security policies has become more pressing when it comes to keeping track of different environments. As a result, an effective, end-to-end strategy starts with maintaining consistency across security policies.

2. **Consistently create and update data copies:** Backed up data copies allow you to restore data back to a point in time before data integrity loss.

3. **Establish complete traceability and visibility:** Seamless traceability from the controls to the policy to the tools that implement the policy creates an ideal environment in which you can quickly map issues or data integrity problems to the business impact. Data managers should also have complete visibility into the security patterns of users for logging, encrypting users' data, and stopping potential outbreaks. This will make it easier to know when a data integrity issue occurred.

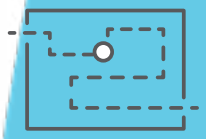

4. **Constantly monitor your data environment:** Everything must be monitored, from the firewall to your organization's internal components to data access patterns. Automating security helps enable continuous monitoring as well.

Leveraging AI to Enhance Data Protection Strategies

AI is increasingly becoming a useful tool for supporting data resilience and protection strategies. AI can play a significant role in analyzing metadata to understand what data contains PII or other regulated information. We are seeing a trend where organizations are using backup software to do automated data classification for this purpose, especially when using regulated data for analytics.

AI can do the heavy lifting to proactively detect abnormal patterns when writing to disk, something that is very difficult to do manually by sifting through logs after something has happened. These tools do not work right out of the box but need time to learn what normal is in your environment, typically 30-45 days, to capture normal behaviors for end-of-month processes and avoid false positives as you begin to trust the tool.

Addressing Data Protection Misconceptions

It is also important to identify and correct any data protection misconceptions before developing your organization's strategy. Here are a few common ones:

- **Public cloud is automatically safe and secure:** Just like any environment, you must attend to your organization's cloud security strategy. Cloud providers do not have this responsibility — you do. Make sure levels of security are appropriate for your business workloads and understand what your responsibilities are in a cloud provider's shared security model.
- **Data protection is an "all or nothing" undertaking:** In many cases, organizations believe they should either backup and restore absolutely everything or nothing at all. There is a middle ground here. Your organization should do what it can to protect its data and evaluate the cost and complexity against business risks and impacts.
- **Data should only be backed up to guard against disasters:** The fact is, data integrity can be damaged through small breaches and inadvertent data deletion -- not just full-on disasters or ransomware breaches.

Action Recommendations

Your data protection strategy will depend on your business needs, your industry, your data environments, the risks you might be exposed to, and more. There is not a "one-size-fits-all" approach that will meet all your needs, but you can follow these steps to develop a data protection strategy that is fit for your business:

Evaluate the data environments and the processes you need to keep data protected

Automate where possible and create an environment-agnostic data protection workflow that works on-premise, in the cloud, or in a hybrid environment

Ensure you have data protection best practices in place that you monitor and update regularly - from traceability to security policies to data

Create a plan to consistently maintain your data resilience strategy

Data Protection Is More Important than Ever

Years ago, data was an adjunct to the business. Today, data is the business and its value cannot be understated. While today's data protection needs are more complex than the past, the tools and technologies needed to protect data in a dynamic environment have evolved to support this vital need, including AI, which is quickly becoming a data center staple for efficient, error-free operations.

By leveraging the best practices listed above, your organization can better ensure business continuity when an adverse condition arises.

EVOLVING SOLUTIONS

We're here to help protect your organization's data.

Let's get to work.

 ESI-010325-X104