7 DATA PROTECTION CONSIDERATIONS TO REDEFINE YOUR BUSINESS

As the world confronts the challenges of cloud, big data and mobile computing, data protection is more critical to business success than ever before. Done the right way, data protection instills trust, ensuring our most critical business asset – our data – is protected wherever it is, against whatever happens.

Armed with this certainty, organizations are able to support new data center, Cloud, Mobile and Big Data environments with confidence, react faster to new business opportunities, and compete more effectively in the marketplace.

This paper looks at how organizations that redefine data protection can transform what's possible for their businesses, and what this transformation looks like from a business, technology and people/process standpoint.

A must-read for C-level audiences, business teams, and IT departments looking to elevate their businesses and their roles within them.

Introduction: it's all about mindset

One of the biggest challenges IT organizations face is building executive mindshare, and this is especially true for data protection. However, in today's always on, always connected world, it's critical. IT organizations have long been trained to think of their value in terms of CAPEX and return on investment; data protection teams in terms of backup speeds, recovery times and deduplication ratios; storage teams in terms of I/O and capacity, and so on. And while these metrics are important, they are not game-changers.

To be relevant, IT departments need to adopt a business-down, customer-centric view of the data center and their roles within it, and they need to help align what they do and the IT investments their companies make to business goals. This means thinking about how technology can be used to bring new applications online faster, reach new customers, deliver new value, drive innovation and outpace competitors.

Sub-Zero Group and Nielsen are good examples of companies that are continuously innovating around IT and are redefining the role of IT, and data protection in particular. They have adopted a new business-focused mindset when it comes to evaluating the technologies they are deploying, the IT processes they have in place to support their businesses, and the roles they have defined for both IT and business teams.

Although these companies are at very different stages of transformation and their environments are very different, they share a boardroom view of IT. Regardless of how much data they have, what applications they have used to create the data, or where their data resides, their number-one concern is making sure their data is always available and always protected wherever it is, whatever happens. In this way, data protection becomes an enabler of business, not an inhibitor.

Following are seven considerations every company should take into account; they will redefine your business.

EMC²

1 Information is king

Without question, enormous changes are at our virtual doorstep, and at the center of all this change is data, or information. In today's interconnected world, data has new business value. In fact, for many organizations, it is their number-one asset, worthy of being listed on the company balance sheet, capable of moving shareholder value, and responsible for driving innovation, revenue and efficiency, etc.

Like it or not, this is the new reality, and organizations that don't embrace it, whose business and IT teams ignore it, risk more than higher costs and lower productivity; they risk relevancy, or worse.

In fact, Gartner, which has been tracking closely the increasing economic relevance of data, has coined the term "infonomics" (information + economics) to help quantify data's new value and importance within and to organizations, and assess the implications for traditional IT operations, including data protection. Gartner VP of Research, Information Innovation and Strategy, Doug Laney explains in this webcast.

As such, the need to protect our information is more important than ever before, requiring a continuum of data protection solutions to ensure data is always available wherever it resides.

EMC²

2 IT must have a seat at the business table

The bridge between IT and business executives is closing as more and more organizations realize the business importance of their IT teams and are including them in business-strategy planning sessions.

As IT's center of gravity shifts from the back office to the board room, its value to the business is also changing. IT professionals are being seen less as cost centers and more as business drivers, and as this happens, IT compensation also trends upward.

This bodes well for IT professionals who get it (i.e., those who understand that their role is strategic first and tactical second).

EMC²

3 Trust binds IT to the business

When business teams don't get their needs met, the tendency is to bypass IT departments in search of the services they need. However, despite the disruptive shifts in technology and the way IT is consumed, IT departments can not only regain the confidence of business teams but also position themselves as trusted advisors to the business. The question is, how can IT re-instill trust where it has been lost?

Re-establishing trust requires IT teams and businesses to collaborate. This means IT teams need to communicate directly with the application, protection and primary storage, and VM teams, and, equally important, it requires IT teams to reach out to their business counterparts to establish relationships. Doing so positions IT teams as infrastructure brokers.

On the flip side, business teams must be open to engaging with IT teams in new ways. They must understand the value of the information these teams bring to the business and must be open to forging new relationships with them (see #3 below).

From a protection perspective, re-estabilishing trust boils down to three things: streamlining protection services, creating application-intelligent infrastructures, and enabling cross-infrastructure data mobility (see below).

Note: As an example, we've used the VM environment; however, the same principles would apply to application-direct (i.e., data dumps) and storage-level (i.e., snapshots) protection approaches.

Step 1—Streamline Protection Services

Poor service levels and lack of visibility into the protection copies have weakened the relationship between data protection teams and their customers (i.e., business and application owners). Most IT environments employ a bewildering assortment of protection products, or technologies, at different layers of the infrastructure (application, hypervisor, server, and storage), which result in unnecessary capital expense, operational complexity, poor oversight, and, yes, failed recoveries.

This creates the perfect environment for business teams to go rogue and deploy their own solutions. Frustration and dissatisfaction abounds on both sides, and the business ultimately suffers. However, by breaking down the protection silos and enabling their customers (i.e., business team leaders, application owners, etc.) to play a more active role in the protection strategy, IT departments can reestablish trust and enable the business.

Take the virtual machine (VM) team—one of IT's critical users. They ask for three things:

1. Access to data protection controls via standard VM interfaces.
2. The ability to manage protection at a VM-level of granularity, rather than larger storage containers.
3. High-performance unified protection management for VMs—rollback of a VM after a logical corruption, full failover after a catastrophic disaster, and the ability to test the recoveries.

EMC²

By delivering better performance and more visibility, the infrastructure team can deliver streamlined protection services. They can deliver trust.

Step 2—Create an Application-Intelligent Infrastructure

IT builds credibility with the business by delivering application service levels reliably and quickly. To better meet SLAs, IT must understand the relationship between application and infrastructure. While virtualization optimizes the IT environment in many ways, it has also made it much more difficult to know what is really happening. When diagnosing an application-level issue, it requires knowing what infrastructures depend on it. When assessing the potential impact of an infrastructure change, it requires predicting the upstream application impact. To cut through the layers of virtualization, IT needs two things:

- Holistic visibility into the environment: The VM admin needs control over both the protection policy and the protection infrastructure to get a full view.
- The ability to capture and validate application requirements: Applications have multiple components. To track the interdependencies, the VM admin needs a grouping mechanism (e.g., consistency groups). Then, to validate the grouping, the VM admin needs to be able to test that the application can be accessed.

By delivering a consolidated view of an application-intelligent infrastructure, the infrastructure team can deliver greater agility to the VM admin and to the business. This fosters trust.

Step 3—Enable Cross-Infrastructure Mobility

Nothing deteriorates trust and a sense of security in an IT organization more than being locked into one vendor. While unique functionality creates some vendor stickiness, the weight of the data creates the most powerful lock-in. The complexities of migrating large amounts of data, often stored in proprietary formats, make data mobility an enormous challenge.

With the proper architecture, however, the information infrastructure team can position themselves as trusted infrastructure brokers. By delivering cross-infrastructure mobility, the infrastructure team assures the business that they are putting the right data in the right place at the right time.

Successful teams focus on two areas:

1. Selecting the right layer for data management services (e.g., protection), considering both efficiency and flexibility: Many customers have chosen the VM as their unit of data management because it eliminates storage lock-in and links well to the cloud, via hybrid cloud.
2. Deploying efficient, reliable data movement technology: Customers must have complete confidence in the mechanism they choose to actually move the data. Since data may move over great distances and between clouds, network efficiency becomes even more critical. Moreover, since this is the company's data, proven reliability is the only option.

EMC²

4 BUSINESSES THAT CAN'T ADAPT RISK OBSOLESCENCE

Understanding the cloud and its potential value to the business has become critical. In fact, moving forward, a lack of understanding of the cloud won't be an option whether you're in the IT business or a different industry altogether.

In this cloud-centric world, applications are extremely agile; they're always on (or should be) and they're always connected. They are "born in the cloud" not in the data center; they're managed in the cloud; and they're protected across clouds.

The data (or information) these applications generate is different; it lives anywhere and everywhere… whether in the cloud under our control or hosted where it can be served up to anyone, anywhere, and at any time with the right permissions. In this world, metadata, or the information about our information, becomes important (see "The "Journey to Redefine" Hinges on Metadata"). It's essential for analytics, tracking, correlating, sharing, data movement (between private and public clouds), data management, compliance, security, and data protection.

Fear of the Cloud – And What It Means for Data Protection

Today, the majority of organizations still have an on-premise data center component. They're virtualized, but they're thinking of ways they can leverage a hybrid cloud to improve IT and (to a lesser degree) business processes.

For example, organizations are using the cloud in the following ways:

- As a tiered infrastructure to reduce CAPEX and OPEX expenses.
- As a smart cloud for application consistent disaster recovery spin-ups—in other words, even if you run an application in a virtualized environment, you can actually use the data in the cloud as a disaster recovery solution.

But for many IT organizations, there's also a strong element of fear of the cloud, specifically of business managers taking IT into their own hands and rolling their own applications in the cloud. While this concern is understandable, what's even more worrisome is the next wave of change. This is where the bulk of our applications are truly born in the cloud, we're generating lots of data away from the data center and we're looking to the cloud to help with "small data sprawl. This is the world in which the cloud becomes pivotal. However, as with most good things ("born in the cloud" being one of them), there are potential "gotchas" that could derail good business intentions if organizations aren't careful.

EMC²

5 CAPEX/OPEX IS NO LONGER THE ONLY FOCUS

To take advantage of the new business playing field, organizations need to 1) shift their center of gravity from product to customer and 2) bring IT into the strategic fold. This is a significant departure from how things have historically been done, particularly for IT, which has been focused largely on reducing capital and operational costs, managing the IT infrastructure and putting out fires.

However, when the conversation shifts to how data and technology can be leveraged to maximize business value, the cycle is broken. Business teams stop bypassing IT in search of better services, and IT teams become trusted advisors to business teams, enabling them to accelerate plans. When this happens, organizations are in a position to redefine what is "business" possible.

Consider Big Data. Whether you're a large enterprise with massive data volumes, a midsize organization with pockets of distributed data or a small business with data volumes that are large for your environment, your data has big value and requires improved levels of protection. Doing so will not only help reduce risk and costs, but also drive business growth and revenue.

6 AGILITY IS PARAMOUNT

What people love about the cloud is the "just-in-time" data center it delivers. They love the utility of it; they love the cost-efficiency of it. But what worries many business organizations is all the unanswered questions.

What happens when:

- You put all their data in the cloud and things change?
- Cloud prices go up?
- Business strategies (yours or your cloud provider's) change?
- Acquisitions occur?

When you own your own data center, it's simply easier to manage the migrations. However, if you don't own any infrastructure, things can get dicey if proper thought isn't given to planning, visibility and accountability.

What your business needs is an organization that looks after its data loads and is able to deploy and re-deploy from the cold storage or storage of their choice. If that choice is the cloud, then it allows you to move things around easily, transparently, and without obviating the business benefits of leveraging the cloud to full advantage.

Your business needs a direct path to safeguard its journey to the cloud so you can:

- Use multiple clouds.
- Redeploy applications and data in the cloud, and tier data as needed.
- Have visibility into your clouds to ensure your data is still going through it… just because you can't see the box or touch it, you need the ability to audit it as if were a physical box. Think "master librarian" or single source of truth.

EMC²

7 YES, METADATA!

Metadata—or the information about your information—is the smarts of your business; it is what enables IT to become a service provider to the business, and as such is the contract that binds IT to the business.

- There are three categories of metadata:
- Infrastructure metadata—Each element of the infrastructure generates vast amounts of data describing their status and behavior.
- Application metadata—Data generated by applications describing their state and relationship to other applications.
- Content metadata—User content is accompanied by information like the data owner, those who have permission to access data, and keyword tags about the data.

It enables IT to track, monitor, and analyze data across the business. Even more, it allows protection teams to become providers of data management services by unlocking the next generation of use cases tied to protection metadata, and this is huge. These services include:

- Security—Correlates infrastructure metadata (e.g., who is logging into what systems) with application metadata (e.g., what is running on that infrastructure) and content metadata (e.g., what data are they accessing/creating) to flag security and compliance issues.
- Availability—Correlates application metadata (e.g., what applications are being created or moved) to the infrastructure metadata (e.g., where is the load going to run and be protected) to predict availability issues.
- Cloud Compliance—This flow is similar to on-premise security, with one additional requirement. IT must work with the cloud provider to access the infrastructure metadata. Any cloud provider unwilling or unable to provide access to key log information is not mature enough to trust with critical application workloads.

When this happens, the protection team increases its value to business teams. By collecting, analyzing, and taking action on the metadata, IT can enable hybrid cloud mobility, analytics-driven automation, and public cloud data management. This is huge not only for your business but also for your data protection teams who are uniquely positioned to be the masters of the metadata, connecting them to the business in ways they've likely never imagined.

Hybrid Cloud Mobility

IT needs to put right data in the right place to run the business.

Organizations may need applications to run in different locations within their private or public cloud(s). That mobility may be for data availability, disaster recovery, application performance, expanded workloads, or a myriad of other reasons. Regardless, IT needs to dynamically leverage resources rather than compromise agility and capital efficiency by spending time and money deploying new infrastructure.

To create a viable hybrid cloud, you need infrastructure metadata. Moving workloads demands knowing where you can move, what mechanisms are available to enable the movement, and what secondary effects the move will have (e.g., on your protection copies, security policies, etc.). In other words, a functional hybrid cloud requires analyzing and acting upon infrastructure metadata.

Mobility isn't just about spinning up storage objects or a single database in an alternate location. Mobility demands that the entire business application continue to run. This means IT needs to understand, track, and ensure that each component of the business application fits together, even as modules migrate. Application metadata enables IT to track those relationships and dependencies.

Mobility is the lifeblood of the hybrid cloud, but it only works if you have expertise in infrastructure and application metadata.

EMC²

Analytics-Driven Automation

Big data and cloud intersect at metadata, where they can drive infrastructure automation. Environments are changing so quickly and becoming so complex that customers can no longer simply operate on instinct alone.

Infrastructure metadata can provide visibility and insight into a customer's environment. Most customers have invested in multiple analytics tools. They struggle, however, to realize the value from those investments. First, the amount of metadata is so vast that they end up with an ocean of information. Second, the customers that can analyze the data end up deluged with reports, dashboards, and charts. The problem? They have data and they have reports, but they cannot take action.

The Internet of Things is just that — a vast ecosystem of devices that are willing to share TBs of information with anyone who will listen. Like the early Internet, it is noisy, frustrating, and isolating. We need the Internet of Things to connect to one another and to build a Social Network of Machines. With these connections, we can drive action.

Let's bring it down to a specific use case. Customers want to know if they are getting their money's worth from their protection storage configuration (e.g., Is my dedupe rate, performance, or utilization high enough, etc.)? Should I rebalance my workload or upgrade my systems? Should I choose a different type of storage system?

Today, answering those questions feels like mixing experience with black magic, but it doesn't have to be. By using analytics-driven automation, IT organizations can provide value to the business. For example, systems can send telemetry about how they're doing. They leverage algorithms to detect which systems are in the real-world sweet spot and which one are outliers, and then recommend the actions to migrate the outliers into the sweet spot.

Cloud-Centric Data Management

Businesses love the agility of the infrastructure and platform service. However, as with any new technology consumption model, there's a lot of underlying chaos. Some applications will be born in the cloud, live in the cloud, and die in the cloud. Others will reside in the data center, and still others will bridge both worlds. With metadata, protection administrators can help preserve the business.

Here's how:

- First, the protection team must ensure that the business gets what it expects out of the cloud provider. That means pulling in the infrastructure metadata to validate performance, security, data protection and availability. Public cloud providers will differentiate themselves by offering the pertinent metadata necessary to prove that their customers get what they pay for. Trust, but verify.
- Second, as businesses deploy more critical applications to the cloud, they will need enhanced application protection options. This means that protection copies will need to move within a cloud, across clouds, and potentially back on-premise. In other words, the work for hybrid cloud mobility is a requirement for public cloud data management.
- Third, customers will need to search for data across their entire application space. Just as it means searching across active and protection copies, it also means searching across private and public clouds. They could try to run multiple searches across each different part of the infrastructure and stitch together the results, but that approach is both complex and error-prone. They would likely only do such a search in the direst of legal or compliance circumstances. They would never leverage the information for deeper business insight. To truly unlock the potential of the metadata, one needs a single source of truth – a metadata lake – to search across all of your information.

Applications will move to the public cloud. The protection team can extend their offerings to ensure that the business's investment is maximized, data is safe, and all metadata is available for search and analytics.

CONCLUSION

For more detailed information about these seven considerations, please check out the following sites and follow us on Twitter @guychurchward and @makitadremel. They are jam-packed with business, market and technology information that will help your organization redefine what's possible:

http://emc.com/protectionleader

http://emc.com/dataprotection

http://protectioncontinuumblog.emc.com

EMC²