

## **A Holistic Approach to Managing Data Protection Environments**

**MAY 2011**



Data protection solutions often get a bad rap under the premise that they are too cumbersome to manage. In fact, data protection solutions such as EMC NetWorker, Data Domain, Avamar etc. do an excellent job at doing what they have been designed to do – protect data. Furthermore, vendors have gone to great lengths to make such solutions simple to deploy and manage. So why is there this paradox?

The reasons for it have very little to do with data protection technology and more to do with the changing landscape in IT. First, the adoption of multiple solutions for data protection including disk based technologies to meet aggressive RPOs and RTOs leads to the creation of data protection silos that lack a unified management framework. Second, traditional approaches to data protection are becoming ineffective as companies adopt next generation cloud-based storage and compute technologies. These changes over-burden existing data protection resources in an unprecedented manner thus increasing the uncertainty that they will be able to effectively protect the business. This uncertainty drives companies to reevaluate how they approach data protection.

Fortunately, Data Protection Management (DPM) solutions provide relief to IT organizations that struggle to maintain control over their multi-tiered data protection infrastructure. They enable the use of existing backup, recovery, virtualization and replication technologies in an integrated manner and tie together the associated instrumentation into a single reporting, alerting and monitoring framework. DPM solutions thus enhance the visibility and control that IT organizations have over their data protection infrastructure in a cost effective manner.

The DPM solution from EMC is Data Protection Advisor. The Data Protection Advisor architecture is optimized to provide unified monitoring, correlation, analysis and reporting of a wide variety of industry leading backup and recovery software, replication solutions, virtualization infrastructure, and SAN and NAS solutions from EMC, VMware and other vendors. Simply put, Data Protection Advisor harnesses critical configuration and operational information from software and hardware components, and correlates and analyzes this information to provide multi-tiered reporting, monitoring and alerting capabilities via a customizable and an easy to manage interface. In this solution profile we will examine how DPM solutions such as Data Protection Advisor complement existing storage management instrumentation in providing capacity reporting, planning, and optimization as well as usage charge back capabilities.

## MANAGING DATA PROTECTION

Most IT departments today will readily admit that adequately protecting business data is one of the most challenging tasks they face today. In spite of the continually advancing capabilities of the individual components of a data protection infrastructure, data protection is no longer a singular entity, and managing it as a whole is a struggle for many. Data protection environments have become multi-tiered, multi-vendor and siloed in nature while the data they are supposed to protect only continues to grow. This problem is further exacerbated by the fact that resources are spread thin and often overlook gaps in data protection policies because of the need to sift through various user interfaces or tools. Let us examine these challenges in detail.

- **Managing data sprawl:** Data continues to grow at an unprecedented pace and application architectures become more complex day-by-day. This increases management complexity as operations teams constantly struggle to maintain protection. Moreover, such sprawl and complexity compels IT departments to deploy one-off solutions that create data protection islands and further increase management complexity. It is a vicious circle.
- **Dealing with resource constraints:** This increased data sprawl is unfortunately not accompanied by a proportionate increase in the IT budget or a proportionate increase in the number of resources managing the environment. IT resources are simply spread too thin, and there are too few drivers manning the wheel. It is no surprise that this increases the potential for operations lapses and companies are left exposed to irrecoverable data loss or corruption.
- **Tackling Multi-vendor or Multi-tier data protection:** Companies today are often choosing to go beyond simple tape based recovery solutions to hybrid array, disk and tape solutions in pursuit of more aggressive operational and DR RPOs and RTOs. It is also not uncommon to find that companies must support multi-vendor solutions after strategic shifts or M&A events. While there is a good amount of integration available with products from the same vendor, deploying solutions from different vendors means a nightmare for companies trying to obtain a unified view of their data protection domain. The end-to-end integration between many different data protection systems is often a challenging and manual process.
- **Dealing with infrastructure consolidation:** The adoption of server virtualization technologies, often from multiple vendors is morphing most traditional environments into ultra-dense consolidated entities. Meanwhile, next generation technologies like deduplication, virtualization and cloud computing are becoming commonplace in today's enterprise and seek to further raise the bar on consolidation. But traditional data protection policies and processes are often no longer effective in such an infrastructure and in turn, companies adopt specialized data protection technologies and practices to deal with this "reverse sprawl". Technologies such as those used to manage the virtual infrastructure, while solving specific problems for a specific component, only decrease the level of visibility in the data protection environments making it more challenging for IT administrators to detect and solve problems.
- **Challenges with unified reporting, alerting and monitoring:** Multi-vendor, multi-tiered environments also mean multiple points of reporting and alerting. While each product may be sufficient on its own, in growing environments, such instrumentation cannot scale. It may be impossible at scale to aggregate the output of such multi-product instrumentation to provide proactive correlation, forensics and analytics that assures that the business is meeting data protection goals. Multiple interfaces may make it impossible for multiple operations teams to concurrently work on multiple protection technologies while understanding the linked dependencies between them. For example, a change in replication for a critical system with an assumption that tape backup is in place for it can in fact leave it unprotected.

These issues result in companies losing confidence in the ability of their data protection infrastructure to meet the recovery objectives of the business. No one wants to know that in spite of

the enormous investments in the data protection infrastructure they remain exposed to the risk of data loss because of operations lapses or oversight.

## DPM SOLUTIONS RISE TO THE CHALLENGE

Fortunately, DPM solutions such as Data Protection Advisor take data protection infrastructure management to the next level. They seek to unleash the full potential of data protection environments by acting as “glue” that binds the various DP tiers into a single consolidated domain. By providing a single pane of glass, DPM solutions provide IT administrators with an additional set of hands and eyes on their data protection environment. IT environments can now leverage the intelligence and insight offered by such tools to solve many of the challenges mentioned above.

DPM solutions can be deployed alongside existing data protection environments and seamlessly integrate into the existing IT infrastructure. With an easy out-of-box setup, such solutions can be quickly configured to support a wide range of data protection applications, replication solutions, backup to disk platforms, tape libraries, SAN and NAS platforms and database applications. These solutions can be used to complement existing data protection environments and provide an additional layer of the following value-added functions:

- **Monitoring & alerting:** DPM solutions provide holistic, proactive and real-time alerting on the hardware and software components of data protection environments. Their analysis engines provide rules-based checking on incoming data in real-time to identify failure conditions, something that can't be done manually and at best is tedious with multiple point tools. The capabilities of DPM solutions therefore complement the native alerting capabilities of the installed products and provide an additional set of “always-on” eyes.
- **Troubleshooting & optimization:** The correlation algorithms found in DPM solutions can be used to provide key metrics related to backup, recovery and replication components of the data protection environment. In addition to real-time activity aggregated and presented in an easy to read manner, DPM solutions can also provide trending analysis. These features help IT administrators gain valuable insight into their environment – insight that aids in the troubleshooting and optimization the environment. The reporting capabilities of DPM solutions can assist IT departments identify misconfigured hardware and software, hot spots and other offending conditions that reduce the operational efficiency of the environment. DPM solutions thus assist IT organizations to improve service levels, reduce the burden on the operations staff and thus reduce budgetary strain by making asset purchases or refreshes less frequent.
- **Capacity planning & chargeback:** DPM solutions are an excellent tool to assist in deploying capacity planning and chargeback capabilities in a data protection environment. They can calculate historical tape and disk usage, estimate the amount of new capacity that needs to be purchased and the timeframe in which it needs to be installed in order to avoid any operational lapses. This makes budget planning more efficient and minimizes off-budget purchases for data protection environments. DPM solutions also have the functionality to provide departmental-level visibility for the services provided by IT for chargeback and planning purposes.
- **Reporting & policy enforcement:** Most DPM solutions feature a customizable real time reporting interface that makes the task of end-to-end reporting in multi-tiered and multi-vendor data protection environments easier and quicker, in many cases by as much as ten times. This information can be used to deduce whether policies and service levels are enforced consistently and if the company is exposed to any risks associated with data protection. The data gathered for these reports can also be exported to third party management applications for enterprise-level reporting and integration as needed.

## THE EMC DATA PROTECTION ADVISOR SOLUTION

Data Protection Advisor is EMC's data protection management solution. Data Protection Advisor consists of a powerful data mining database mated to an intelligent analysis and correlation engine. The central data mine drives the monitoring, reporting, alerting, troubleshooting, capacity planning and optimization capabilities. The analysis and correlation engine is the brain behind Data Protection Advisor and offers all of the features that make it stand out amongst the variety of DPM solutions in the market. Data Protection Advisor is available in the following variants – each designed to work in a standalone manner or as a complete solution.

- **Data Protection Advisor for backup** supports most industry leading backup technologies including those from EMC (Data Domain, NetWorker and Avamar), IBM, HP, CommVault and Symantec. It also supports a wide variety of tape and backup to disk hardware.
- **Data Protection Advisor for replication analysis** supports replication solutions including EMC RecoverPoint, VNX, CLARiiON MirrorView, Symmetrix SRDF and TimeFinder as well as industry leading databases such as Microsoft Exchange, SQL and Oracle. The replication analysis engine is an important oversight tool in environments as it can detect situations where the data may not be recoverable even if there are no warnings. The ability to map and catalog all items in the “replication chain” i.e. hosts, databases, storage configurations etc. and present them in a graphical view is unique to Data Protection Advisor.
- **Data Protection Advisor for virtual infrastructure** supports full integration with VMware vSphere and enables performance tracking, resource consumption, protection status and conflicts. It is an invaluable tool in densely consolidated virtual environments today. It helps IT administrators create a critical link between their virtual infrastructure and their data protection environment. DPA automatically discovers virtual machines as they are created and tracks their “protection” state even if they are moved to different hypervisors.
- **Data Protection Advisor for file servers** complements EMC Celerra and VNX environments and provides a single lens view for performance and utilization monitoring.

Data Protection Advisor's comprehensive monitoring, alerting and optimization capabilities extend the ones offered by native tools in any data protection environment. Additionally companies often deploy Data Protection Advisor to provide secondary reporting on their data protection environments specifically for backup success-failure reporting and to identify jobs that are likely to fail, slow down or exceed backup windows. Without Data Protection Advisor, these organizations often grapple with the number of backup recovery instances and replication environments, each having its own set of reporting tools that lack correlation. Data Protection Advisor allows them to proactively optimize their environment by eliminating the offending conditions – which many times are also detected by Data Protection Advisor.

## DATA PROTECTION MANAGEMENT USE CASES

The capabilities of Data Protection Advisor make it an invaluable tool in any data protection environment. Its benefits span many critical functions such as management of a *multi-tiered* backup ecosystem and provide a centralized view into replication deployments as an extension of the company's data protection environment. These use cases illustrate how deploying DPM solutions such as the Data Protection Advisor assist IT organizations in getting an integrated view of their data protection environments.

### ***Use Case #1: Single-lens view for multi-tiered data protection environments***

A typical data protection environment these days can consist of backup recovery software such as EMC NetWorker that manages a blended disk and tape environment including a backup to disk solution such as EMC Data Domain. Many IT departments also choose to deploy a solution such as

EMC Avamar for remote office consolidation or for backing up virtual infrastructures such as VMware vSphere. These applications are excellent at what they do and even better when used in concert.

A multinational service provider uses Data Protection Advisor to complement their data protection infrastructure that consists of NetWorker for their enterprise backups and Avamar for remote site backups. With over 4000 systems, 1.6PB of SAN and 900TB of NAS backed up into over 14,000 tapes, the challenges faced by the company were enormous. The scale of these challenges can be further appreciated if one factors the 600 database backup clients, over 1,300 database instances, 250,000 Microsoft Exchange mailboxes, 376 Lotus Notes servers and 278 Microsoft SQL Server and SharePoint databases. Data Protection Advisor brought the ability to view this entire data protection portfolio through a single console and provides a host of proactive monitoring, reporting and alerting capabilities that complemented what NetWorker and Avamar could provide. Data Protection Advisor in this environment generated significant quantifiable hard and soft dollar savings. The ability of Data Protection Advisor to optimize the environment resulted in a reduction in the amount of tape used by over 50% and the hardware budget by 60%. The ability to report and troubleshoot on backup failures before they occurred increased their average daily backup success rate to over 99.8% and their audit success rate to 100%. They have incurred zero service level agreement penalties since deploying Data Protection Advisor, indicating the level to which they have gained control of their backup operations.

### ***Use Case #2: Convergence between primary & secondary data protection***

The need for rigorous operational and disaster recovery RPOs and RTOs requires the deployment of a multi-tiered backup and recovery infrastructure. In such multi-tiered environments, backups to disk or tape provide a primary level of protection. For mission critical applications that require a more granular level of data protection an additional layer of array or appliance based technologies such as replication, snapshot and cloning are deployed. Such services are often integrated with the application recovery processes and tied to the backup application by using custom scripts or hooks provided by the array vendor. These types of environments often lack the orchestration layer that unifies them. This makes managing them a challenging task. DPM solutions provide this orchestration layer by tying together the backup and replication infrastructure. Data Protection Advisor supports EMC's flagship replication, snapshot and cloning solutions on VNX, CLARiiON, VMAX, Symmetrix and RecoverPoint.

An example of such a deployment is in EMC's own corporate IT environment. With a massive infrastructure to support the world leader in information infrastructure solutions, the task of data protection is no easy task. EMC's IT found it a challenging task to effectively capture and troubleshoot backup failures in their geographically distributed environment. Due to the immense growth in data fueled by M&A and organic growth, EMC IT found it challenging to manage and forecast capacity for their data protection environment. Additionally, the lack of visibility into the replication of its Oracle data warehousing and backup environments was also a constant challenge. To reign in these challenges, EMC IT deployed Data Protection Advisor to automatically monitor, analyze, and provide alerts and reports on its backup and replication environments.

Data Protection Advisor oversees EMC's entire global backup infrastructure, which today includes Avamar, Disk Library, NetWorker, Data Domain and tape libraries. Data Protection Advisor is used for managing backups. Data Protection Advisor is also used to provide host-level reporting on their server infrastructure that consists of standalone physical servers and virtual machines running on VMware vSphere. Applications and databases running on these servers include Microsoft Exchange, SQL and Oracle databases. EMC IT also recently deployed Data Protection Advisor to provide secondary reporting on their EMC Celerra NAS environment.


EMC uses Data Protection Advisor to report on and analyze their replication environments for gaps in meeting operational and disaster recovery RPOs and RTOs for their mission critical Oracle environments, one of which is their data warehouse instance that is protected by RMAN for recovery and SRDF for replication. The use of Data Protection Advisor also allows EMC IT to meet their service level objectives by deploying it for incident management – by proactively identifying and analyzing issues for faster resolution. Its holistic integration has helped in proactive license management and capacity management leading to over half a million dollars in savings. Data Protection Advisor's role based access control with a customizable web interface makes it easy to deploy in geographically distributed IT environments where several teams manage various parts of the data protection infrastructure.

EMC IT no longer has to manage separate (and often custom) interfaces for their primary, secondary and tertiary data protection infrastructure. The out-of-band architecture of Data Protection Advisor means that it can provide a truly global view of the EMC IT data protection infrastructure. Furthermore the analytics and correlation algorithms can be extended into the replication infrastructure and make the end-to-end enforcement of data protection policies and the compliance with service level objectives (SLOs) a seamless task.

## EASING THE JOURNEY TO THE CLOUD

Companies have to constantly adapt to a changing landscape. Whether it is the proliferation of mobile computing or something else, one thing is for sure – the appetite for data is only going to increase. The pressure on IT organizations to manage this data forces the adoption of newer technologies that are cost effective and sustainable in the long run. A cloud based utility model for storage and compute infrastructure is the latest solution to tackling the challenge of managing data growth.

A movement to the cloud and beyond means IT has to change the way in which they function. IT services can no longer be offered in disparate silos but rather as a consistent end-to-end service i.e. IT as-a-service. Most companies making this transition adopt a three phased approach that starts with consolidation and goes through improved service quality before offering IT as-a-service:



The three phases of journey to the cloud

- **Phase 1 – Consolidation:** In this phase, companies seek to consolidate their infrastructure in an effort to lower their costs. Most companies achieve this consolidation by heavily leveraging server virtualization technologies for their non-production environments. Such companies also seek out to maximize their asset utilization including that of their data protection infrastructure.
- **Phase 2 – Improved quality of service:** In this phase, companies focus on improving the service quality of their IT infrastructure. This focus is important as they continue to transition their production and mission critical environments into a virtualized infrastructure.

- **Phase 3 – IT as-a-service:** Having successfully driven down costs and improved the service quality of their infrastructure, companies can now make the final transition from being a reactive entity into one that thinks strategically and one that has completely decoupled its catalog of services from the physical assets that are used to provide them.

The “IT as-a-service” model focuses on delivering services in a technology-agnostic manner. It seeks to provide better predictability and control over the infrastructure and institute end-to-end and enforceable service level guarantees. Such end-to-end service level guarantees mean that all tiers of the data protection infrastructure need to be in lock-step with the storage and compute infrastructure - all or part of which could be in the cloud. It also means that data protection policies and processes are updated to factor in cloud as a service tier. This coupling – whether it is technology or process related – cannot be achieved without the use of “bridge tools” that enable IT organizations to obtain a unified view of the infrastructure as it morphs through this change.

DPM solutions such as Data Protection Advisor fall in this category. They make it easier for companies to transition into an IT to “as-a-service” entity and ultimately transition their infrastructure into the cloud. They do this by creating a valuable link between the data protection, compute and storage infrastructure – making the task of mobilizing virtual services into the cloud easier. Additionally DPM solutions provide continuity in the reporting, monitoring and analytics capabilities as components of the infrastructure are being moved into the cloud. And once the movement into the cloud is complete, DPM tools provide transparent pane of glass through which the entire infrastructure can be singularly viewed even when some or all of the underlying components may be cloud based.

## TANEJA GROUP OPINION

Keeping up with a dynamic ever-changing infrastructure is a constant challenge for most IT environments. On one side they have to constantly keep up with newer and more complex technologies in order to meet demands placed by the business. On the other side, they have to manage a constant increase in the amount of data being managed, ensuring that it is protected in the right manner and in a way that minimizes the risk to the business. And these tasks must be completed without any relief on the budget.

While IT will perhaps continue to struggle in many other areas, at least on the data protection side, DPM solutions offer much needed relief. In our view, DPM solutions are on their way to becoming a cornerstone of the data protection, recovery and management framework as they bring about a level of predictability and proactivity that IT departments have long struggled to enforce in their environments. There is a multi-domain and organizational IT culture impact from this fundamental capability change – armed with such capabilities, IT can easily make the transition from being a reactive entity to one that thinks strategically as an “IT-as-a-service” provider. In other words, DPM is an enabler for companies embarking on a journey to the cloud and beyond.

There’s a reason for the widespread industry conversation around “IT-as-a-service” and cloud-based IT architecture. That reason sits in the crosshairs of DPM solutions, and was its primary target long before the cloud phenomenon caught on. Fighting fires in data protection environments has long been an all too familiar story. With infrastructure changing at an ever more rapid pace, these fires will only get worse and eventually create a business disconnect that makes it impossible for IT to go to the next level. By making DPM solutions a part of their data protection portfolio, companies can choose to stop this burning and begin putting out fires even before they start. That’s what well managed IT services in this cloud vision of IT is all about.

---

*NOTICE: The information and product recommendations made by the TANEJA GROUP are based upon public information and sources and may also include personal opinions both of the TANEJA GROUP and others, all of which we believe to be accurate and reliable. However, as market conditions change and not within our control, the information and recommendations are made without warranty of any kind. All product names used and mentioned herein are the trademarks of their respective owners. The TANEJA GROUP, Inc. assumes no responsibility or liability for any damages whatsoever (including incidental, consequential or otherwise), caused by your use of, or reliance upon, the information and recommendations presented herein, nor for any inadvertent errors that may appear in this document.*