

Data Protection for EMC VSPEX Proven Infrastructure

Abstract

Data Protection is a critical element of supporting your applications in today's virtualized data center. Data Protection for VSPEX Proven Infrastructure offers the Backup, Recovery, Replication, Archive, Business Continuity, and Mobility that your applications require; leveraging proven, industry-leading components integrated into a seamless solution.

November 2013

Copyright © 2013 EMC Corporation. All Rights Reserved.

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

The information in this publication is provided “as is.” EMC Corporation makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

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

For the most up-to-date listing of EMC product names, see EMC Corporation Trademarks on EMC.com.

VMware, VMware Data Protection and VMware vCenter Site Recovery Manager are registered trademarks or trademarks of VMware, Inc. in the United States and/or other jurisdictions. All other trademarks used herein are the property of their respective owners.

Part Number H12017

Table of Contents

Executive summary	4
What is Data Protection?	4
Industry Trends and Objectives	5
Data Protection Challenges for Today's High-growth Virtualized Data Centers Challenges	6
Increasing back-up complexity	6
Accidental backup architectures.....	.6
Server optimization and traditional backup.....	7
Business Continuity and Application Mobility in Virtualized Environments	7
Data Protection Options	7
Backup and Recovery	7
Replication	8
Business Continuity and Application Mobility	9
Data Protection Recommendations	9
Data Protection for VSPEX Proven Infrastructure	11
Avamar	11
Data Domain	11
VMware Data Protection Advanced	11
RecoverPoint.....	13
VPLEX.....	15
Summary	16

Executive summary

Data Protection is a critical element in supporting your mission-critical applications in today's virtualized data center. Data Protection for VSPEX Proven Infrastructure offers the Backup, Recovery, Replication, Archive, Business Continuity, Continuous Availability, and Mobility that your applications require; leveraging proven, industry-leading components integrated into a seamless solution.

Data Protection for EMC VSPEX Proven Infrastructure is Simple, Efficient, and Flexible.

- Simple – Simplify Data Protection and application availability
- Efficient – Streamline Backup operations and deliver the protection for your most mission critical applications
- Flexible – Flexible levels and choices of protection designed to meet your business and application requirements

What is Data Protection?

In the broadest context the term "data protection" encompasses a wide range of different technologies, techniques and methodologies used to ensure the integrity, availability and currency of both the data and the underlying delivery systems in an IT environment. This solution brief reviews the available data protection technologies for VSPEX Proven Infrastructures providing the following capabilities:

- The ability to back up data and files for recovery in the event of loss or corruption
- The ability to guarantee both the availability and the integrity of data resident on a platform by creating backup versions
- The replication of data and files resident on one system onto another so that the loss of a system does not significantly disrupt service delivery
- Continuous data protection that allows recovery to any point in time (PIT)
- The ability to maintain business continuity in the event of the loss or failure (either temporarily or for a significant period) of an entire platform or parts thereof
- Maintaining identical data in two locations that are capable of simultaneous read/writes at both sites, if one site fails the other site continues to provide access to the data with no downtime.

Industry Trends and Objectives

Massive Data Growth, Virtualization and Backup Windows

Businesses often turn to virtualization in an effort to streamline operations and reduce data growth, however, virtual machine sprawl and large virtual machine images can compound the problem the solution sought to address. Massive data growth has become the biggest factor in driving change in the backup and recovery market. Virtualization also creates new challenges for backup.

The net result has been a migration from slower, less sophisticated tape-based approaches to faster disk-based ones. In particular, purpose-built backup appliances with deduplication are replacing tape. Disk-based appliances speed performance, reduce backup data volumes and efficiently send backup data off-site over existing WAN links for DR.

Reducing costs by selecting the appropriate level of Data Protection

While storage costs have declined over the years, it's certainly not free. Every customer wants zero data loss and zero downtime for 100% of their data. That is always true, until they see the price tag associated with 100% site redundancies with fully synchronous replication.

Responsibility falls to the IT department to offer fiscally responsible data protection options to the business units while ensuring Application SLAs are met. When formulating a data protection strategy it is imperative that one understand the unique availability requirements of one's applications and the financial and business costs associated with a loss or outage of said application. Two key classifications are important here, that are determined by the needs of the business:

- The Recovery Point Objective (RPO) - what is the maximum amount of data that may be lost in case of a system/site failure
- The Recovery Time Objective (RTO) - How fast must the business service be restored in the event of failure

As an example: If the data for a particular application is backed up every two hours and it takes an hour to recover operations from an alternate location, the RPO will be three hours and the RTO will be one hour. If a failure occurred and that particular problem could be resolved in less than an hour, it would not be necessary to declare a disaster and fail-over to another location. When a failure is confined to a single application as in the case of file corruption, the remedy should not disrupt other applications and services running on the same system/server. In such a scenario, data recovery should be implemented without the need for system fail-over. Determining the optimum data protection strategy

requires careful consideration of all these factors, then matching them up to the right backup and recovery option.

Data Protection for VSPEX Proven Infrastructure offers a range of options to meet the performance requirements and SLAs of your most mission critical applications while addressing the cost-efficiency requirements of your less critical ones.

Data Protection Challenges for Today's High-Data Growth, Virtualized Data Centers

Traditionally, data protection has primarily consisted of a combination of local and remote tape-based backup in addition to replication capabilities. Local backup/replication provides easier and faster recovery times in the event of data loss or system failure, due to higher local bandwidth connections. Replication or backup to a second location via the wide area network, although slower, provides a greater level of disaster recovery and business continuity.

Increasing backup cost and complexity

Over time, data protection techniques have developed to keep pace with the growing volumes of data and the increased criticality of IT services to the business. Originally, 100% of all data was backed up nightly. As the volume of data grew to the point that nightly "fulls" could not be completed in the given time window, weekday incremental copies combined with weekend full copies became the norm, further complicating the backup process. Although it met the backup window (at that time), it also added complexity and extended recovery times. More sophisticated organizations went a step further: they split backup workloads so that every night of the week, 1/7th of the data was fully backed up (rotating full backup schedule). While it was an efficient use of the backup window, it again made the administration and recovery all the more complex.

Accidental backup architectures

Given these backup and recovery challenges, data source vendors – that is, application, hypervisor, and storage vendors – have taken action in a way that's completely transformed the industry. These vendors are building tools (e.g., VMware vSphere Data Protection, array snapshots, Oracle RMAN) and optimization (eg, VMware's Changed Block Tracking, Oracle's Block Changed Tracking, and Microsoft's Block Level Change Detection) directly into their products, but with such rapid change and easily obtained options, many organizations have fallen into the chaos of

an ad-hoc backup approach. The application, virtualization and storage teams, dissatisfied with traditional backup, use the native point products now available to them. This has resulted in protection silos and an accidental backup architecture. It's accidental because no one would intentionally plan for a half-dozen unconnected protection tools with no central oversight and no cost controls. Accidental architectures increase cost, complexity and risk as your data protection environment scales.

Server optimization and traditional backup

A key system design goal is to drive utilization of the hardware as high as possible without negatively impacting application performance. The higher the number of virtual machines per server, the lower the overall infrastructure cost. Yet the most taxing workload on virtual machine resources is nightly backup. Backup uses large amounts of CPU disk I/O and network I/O and so the backup workload can easily become the limiting factor on how many virtual servers can run in a single physical server (blade).

Business Continuity and Application Mobility in Virtualized Environments

The widespread use of virtualization and the deployment of reference architecture-based solutions such as VSPEX Proven Infrastructure have created both new opportunities and new challenges. Virtualization supports the dynamic allocation of applications to virtual machines and can automatically distribute or redistribute workloads based on server availability. This provides a level of resiliency not previously possible. However, to provide for true business continuity and to reduce the need for "recovery" requires not only virtualization of the compute resources but also of the storage and networking capabilities so that applications can dynamically migrate between geographically dispersed servers.

Data Protection Options

Backup and Recovery

Backup with deduplication will have a significantly positive impact in addressing the challenges of massive data growth and the new demands of virtualization.

To overcome "backup bottlenecks" in a virtualized environment, a client-based data deduplication backup and recovery product is an excellent option to consider. Client based deduplication dramatically reduces the amount of data that must flow from each source server, through the network, to complete the nightly backup. It delivers an equally dramatic reduction in backup time. As a result, the bottleneck is eliminated and scaling up the virtual machine count can be achieved without adding more manpower, clock time or bandwidth.

VSPEX includes EMC backup and recovery solutions, featuring a combination of Avamar deduplication software and systems, and Data Domain deduplication storage systems.

VSPEX with EMC backup delivers a comprehensive solution to address virtualized application and virtual infrastructure backup requirements. Whether using Avamar only or a combination of Avamar and Data Domain, backups can be easily replicated off-site across a WAN without operator intervention. This is another compelling improvement over tape-based backup. It eliminates time and resources to clone a second set of tapes, and also the cost and security risk of physically moving tapes off-site.

One other consideration for protection of a virtualized environment is the need to have guest level backups for application consistent application recovery, and for full server image-level backups for rapid restore of a failed VM server. Image level backups are always crash consistent. Avamar provides both and depending on your unique IT requirements you may need to leverage both guest and image level backup to meet your requirements.

Replication

When a once-per day backup is not enough protection for a particular business application, continuous data replication provides a solution with lower RPO. A write splitter integrated within the array enables the replication of all writes to one or more remote locations. With synchronous replication, it is possible to achieve zero RPO. This means no data loss in the case of a site failure. Synchronous replication, however, is limited by distance, since longer distance may introduce some latency to the applications. Alternatively, asynchronous replication can be used without latency over any distance, but may be subject to a small data loss in the case of a production site disaster.

RecoverPoint is EMC's replication mechanism that allows not only remote, but also local replication to the same or different array. It creates a 'write journal' that enables recovery to any point in time (PIT) in the event of an operational failure. The journal also allows any PIT image access for testing and data repurposing. This can be done without stopping replication, since the journal acts as a buffer between the replication engine and the replica image being accessed. RecoverPoint further implements deduplication and compression to save bandwidth on the wide area network, and thereby reduce OPEX.

RecoverPoint is uniquely optimized for VMware environments by supporting any PIT with VMware vCenter Site Recovery Manager (SRM). Typically, SRM can only fail over or test the latest image; however, the EMC virtual storage integrator (VSI) allows selection of PIT from the RecoverPoint bookmarks, which can then be used for SRM operations.

RecoverPoint is also optimized to support Hyper-V environments. Its support of cluster shared volumes (CSV) and read-only image access at the remote DR site, enable smooth automated failover in a distributed Hyper-V environment. RecoverPoint further supports Offload Data Transfer (ODX) to enable performance optimization with supporting storage arrays such as the VNX in VSPEX solutions.

Business Continuity and Application Mobility

True Business Continuity requires more than point-to-point replication. It requires fully automated recovery from unplanned disasters as well as the capability to avoid downtime for planned maintenance activities.

Most of today's organizations operate 24X7, with most applications being mission-critical. Continuous availability of these applications to all users is a primary goal of IT. A secondary goal is to have all applications up and running as soon as possible if the applications stop processing. There are hundreds of possibilities for an infrastructure to be rendered inaccessible, from fires, flooding, natural disasters, application failures, or even simple mistakes in the computer room, and many of these are unforeseeable and outside of IT's control. Sometimes there are good reasons to take down applications, scheduled maintenance, tech refreshes, load balancing, or data center relocation. In all of these planned outage scenarios the outcome is the same, applications stop processing.

VPLEX is a storage virtualization platform designed for Business Continuity and Data Mobility within, across and between data centers and storage

arrays. The VPLEX Family is designed to remove the physical barriers of data centers and storage arrays, enabling hypervisors to utilize features such as VMware HA and vMotion beyond the boundaries of a single datacenter, across sites separated by synchronous distances, exactly the same as if the datastore were on the same array. VPLEX solutions are generally leveraged in situations that rely on transparent load sharing between multiple sites, planned workload relocation and minimized recovery time during the unplanned disruptions.

VPLEX provides unparalleled protection for the most critical applications in any data center while extending VMware features HA, DRS, and vMotion to make application mobility even better and easier, with the ability to stretch VMware capabilities over distance, and eliminate the need for host CPU cycles to do Storage vMotion . A single datastore can consist of storage from multiple arrays allowing applications to both read and write to these identical volumes simultaneously, providing true Active-Active storage. This functionality enables users to reduce Recovery Time Objectives to zero and eliminate data loss.

Data Protection Recommendations

Every company's data protection requirements differ depending on the size and scope of their IT operation and the importance of data and applications to the business. To be truly cost-effective your data protection strategy should incorporate the combination of backup, replication, and business continuity capabilities that best meet the needs of your applications SLAs.

Application data should be protected by the appropriate data protection solutions. For example, a loss of real-time financial transactions could have devastating negative impact on a business. Clearly this data easily justifies the highest levels of data protection possible - low RPO, low RTO. In this case the objective would be zero data loss. Similarly, it will also be important that the data be restored almost instantaneously. At the other end of the spectrum, data corruption on a desktop computer that is backed up nightly could be an inconvenience to a single user. It might mean someone needs to re-create a day's worth of word-processing or spreadsheets, but is not likely to alter the corporations' profitability.

When deciding which course of action to take to meet RPOs and RTOs it is important to consider different classes of failure and the circumstances surrounding them. Failure of a system component that can be easily replaced merits a different response than the loss of an entire data center or business location due to a disaster. Some failures can be more easily

corrected than others. The corrective course of action will depend on the RPOs and RTOs that have been established.

VSPEX solutions are fully compatible with a complete range of options to meet the performance needs of the most mission critical applications and the cost-efficiency requirements of less critical ones.

Protection Types	Recommended EMC Products and Data Protection Strategy	Notes
Backup and Recovery (All Data and Applications)	<ul style="list-style-type: none"> • Avamar • Avamar with Data Domain • VMware vSphere Data Protection Advanced (for smaller cloud environments) 	<p>Assumption is once-per-day backup, but variable according to methods, backup performance and more. RPO is 24 hours, maximum.</p> <p>Backups are often retained for 30 days or more.</p>
Replication (Business Critical Applications)	Backup and Recovery along with RecoverPoint for applications where daily backups leave business exposure	Replication does not eliminate need for daily point-in-time backups. Use both.
Mission Critical Application availability and Protection	<ul style="list-style-type: none"> • VPLEX for Mission Critical applications and data availability • RecoverPoint for Replication • Backup and Recovery 	Zero data loss. Near zero service interruption (Recover Time Objective) Use all three Data Protection Options.

Backup and Recovery: All corporate data and applications where once-per-day backup is an acceptable level of data protection. In event of disaster, data loss is limited to last backup completion, therefore, 24 hours max.

Replication: Data/applications where once-per-day backup leaves too much business exposure. Multiple copies-per-day reduces worst case data loss to completion of last replication.

Business Critical: - Business critical applications and data. Zero data loss in event of disaster is mandatory.

Data Protection for VSPEX Proven Infrastructure

Backup and Recovery

EMC backup and recovery is smarter backup. It consists of best of breed, integrated protection storage and software designed to meet your backup and recovery objectives – now and in the future.

EMC Avamar

EMC Avamar provides fast, efficient backup and recovery through a complete software and hardware solution. Equipped with integrated variable-length deduplication technology, Avamar facilitates fast, daily full backups for virtual environments, remote offices, enterprise applications, network-attached storage (NAS) servers, and desktops/laptops. Learn more: www.emc.com/avamar

EMC Data Domain

EMC Data Domain Deduplication Storage Systems continue to revolutionize disk backup, archiving, and disaster recovery with high-speed, inline deduplication for backup and archive workloads. Learn more: www.emc.com/datadomain

VMware Data Protection Advanced

VMware vSphere® Data Protection™ Advanced is backup and recovery software and a virtual appliance that provides proven, efficient and easy-to-use backup and recovery for virtual machines and Microsoft mission-critical applications.

With EMC backup solutions that include market-leading protection storage, deep data source integration and feature-rich data management services, you can deploy a flexible protection storage architecture that combats accidental architectures and allows you to scale while lowering cost and complexity.

Increase Scale

EMC is the backup industry's performance and scale leader. You'll reduce backup times by 90%, and our protection storage ensures fast, reliable disk based recoveries. Advanced integration with VMware speeds recoveries by 30 times – even delivers VM instant access. Smart replication moves only the data needed for efficient and fast DR protection – eliminating the need for tapes and trucks. And our protection storage systems are up to 27x more scalable than the nearest competitor, so you'll stay well ahead of the data growth avalanche.

Reduce Cost

Integration hardware and software makes it easy to deploy and manage EMC backup, while deduplication reduces the amount of storage and bandwidth required for onsite and offsite protection. In all, customers see an average 7 month payback, 10-30x reduction in backup storage required, up to 99% reduction in required bandwidth and an up to 81% reduction in time spent on backup administration.

Lower Complexity

EMC backup simplifies so you can go faster. End users get the visibility they need and the backup team gets the infrastructure control they need while eliminating routine tasks. This prevents end-users from creating protection silos, which creates complexity. It also provides backup teams with consolidated infrastructure and views that's essential for better management. With fewer moving parts and a central service for backup, you'll speed your organization's virtualization, application roll-outs, and business expansion – while risk decreases.

ADDITIONAL RESOURCES

- EMC Backup and Recovery Options for VSPEX Private Clouds
- EMC Backup and Recovery Options for VSPEX End User Computing with Horizon View
- EMC Backup and Recovery Options for VSPEX for Virtualized Microsoft SQL Server 2012
- EMC Backup and Recovery Options for VSPEX for Virtualized Microsoft SharePoint 2013
- EMC Backup and Recovery Options for VSPEX for Virtualized Microsoft Exchange 2013
- EMC Backup and Recovery Options for VSPEX for Virtualized Oracle 11gR2

Replication

RecoverPoint

Over the past 8 years, RecoverPoint has built a solid reputation in the market for data protection and disaster recovery with over 10,000 units shipped, close to an Exabyte of data protected and over 100 million run hours.

RecoverPoint provides integrated local continuous data protection (CDP) and continuous remote replication (CRR) as well as concurrent local and remote (CLR) data protection that allow for any-point-in-time data recovery. The RecoverPoint family protects data at local and/or remote sites with minimal impact to production application processing. It supports remote replication to multiple remote sites protecting against disasters and replication of branch offices data to a prime central office.

- RecoverPoint's local continuous data protection (CDP) provides synchronous local replication, enabling customers to roll back to any point in time for

effective recovery from events such as database corruption or loss of data due to viruses or human errors.

- RecoverPoint's continuous remote replication (CRR) provides dynamic synchronous and asynchronous remote replication for disaster recovery—with the option to switch between modes based on user-defined policies for throughput and latency—and unique bandwidth reduction that can save up to 90% on WAN expenses.
- RecoverPoint's concurrent local and remote replication (CLR) provides both operational recovery and disaster recovery. RecoverPoint supports all EMC block storage devices and heterogeneous storage environments using VPLEX storage virtualization.

RecoverPoint's operational-recovery capabilities provide storage administrators, IT managers, and CIOs, with their "Peace of mind" thanks to any point-in-time recovery based on journaling technology and application-consistent bookmarks. "Go back in time" restores storage to just a moment before a data corruption happened. RecoverPoint uses a write-splitter that is integrated within the EMC array to capture the changes made to the production volumes and replicate them at the local or remote copy. Connectivity to the remote site can be over either Fibre Channel (FC) or Internet Protocol (IP). IP is normally used for cost savings, or when link latency does not allow FC connectivity. RecoverPoint supports both synchronous and asynchronous replication over IP.

RecoverPoint offloads the processing required for replication from the storage array or application host onto external dedicated appliances. Up to eight appliances are used as a cluster, with at least one cluster at the production site and one at up to four remote sites.

With VNX storage, virtual appliances can be used instead of the dedicated hardware appliances, while supporting almost identical functionality, although lower performance and scale. The virtual appliances are easily deployed within a VMware environment using standard VMware tools, and can provide a cost-effective solution.

RecoverPoint Value Propositions

- Heterogeneous local and remote replication to protect any application, on any array, to any location and recover to any point in time
- Scales up and out - RecoverPoint's high capacity data availability solutions for Tier-1 can scale up to 8 RecoverPoint appliances (RPAs), as well as scaling out by using multiple clusters connected to the same array.
- Use RecoverPoint appliance (RPA) for maximum performance or, for small sites or less-demanding environments with a software-only virtual RPA (vRPA) that runs on a virtual machine

- Support multi-site configurations for customers requiring multi-site redundancy and remote office data protection
- Industry-first: Any Point In Time VMware Site Recovery Manager (SRM) Test/Failover.
- Low total cost of ownership, including savings through WAN dedupe, bandwidth reduction, sync replication capability over IP, journal compression, and snapshot consolidation technologies.
- Support demanding application needs (such as Oracle and VMware) requiring high throughput
- Enhanced manageability and flexibility with RP 4.0 REST API

ADDITIONAL RESOURCES

A full suite of RecoverPoint documents is available on the EMC Online Support site. For Additional information about RecoverPoint solutions for VSPEX please reference the following documentation:

- **Design and Implementation Guide: EMC VSPEX with RecoverPoint for VMware 5.5**
- **Design and Implementation Guide: EMC VSPEX with RecoverPoint for Microsoft Windows 2012 Hyper V**
- **Email:** [RecoverPoint Deal Support Desk](#)
- **ECN:** [RecoverPoint page](#)
- **Powerlink:** [Products → Software P-R → RecoverPoint Family](#)
- **Powerlink:** [Home > Solutions → VSPEX → Data Protection](#)
- **Sales Play:** [VSPEX with RecoverPoint play](#)

Business Continuity and Availability

VPLEX

EMC VPLEX delivers data mobility and availability across arrays and sites. VPLEX is a unique virtual storage technology that enables mission critical applications to remain up and running during any of a variety of planned and unplanned downtime scenarios. VPLEX permits painless, non-disruptive data movement, taking technologies like VMware and other clusters that were built assuming a single storage instance and enabling them to function across arrays and across distance.

Besides offering continuous availability of applications, VPLEX provides easy migrations, workload re-balancing, and dynamic upgrades of storage arrays. Users readily appreciate the ease of using vMotion and DRS to rebalance workloads between arrays without the complexity and time requirements of Storage vMotion and Storage DRS, and the ability to non-disruptively upgrade storage arrays during normal business hours.

VPLEX enables the ability to do workload balancing or application migrations across multiple VSPEX arrays and across regional data centers, with zero downtime during a workload transfer. When VPLEX distributed virtual volumes are configured, and using VMware HA, DRS, and vMotion, the following migrations are possible:

- Non-disruptive workload balancing across VSPEX arrays
- Non-disruptive migrations off of non-VSPEX infrastructure onto a VSPEX array
- Non-disruptive upgrades from older VSPEX to newer VSPEX arrays
- Non-disruptive datacenter moves

ADDITIONAL RESOURCES

- Design and Implementation Guide: EMC VSPEX with EMC VPLEX for VMware vSphere 5.1
 - Email: [VPLEX Support Desk](#)
 - Powerlink: [VPLEX Family](#)
 - Powerlink: Home > Solutions → VSPEX → Data Protection
 - Video: [VPLEX with VMware](#)
 - Video: [VPLEX for Migrations](#)
 - Sales Play: [VSPEX with VPLEX play](#)

Conclusion

VSPEX provides businesses the flexibility to choose hypervisors, servers and networking coupled with Data Protection options that are optimized for virtualization to deliver increased performance, reliability, and availability of your applications. VSPEX flexibly scales as your business needs change and through deep integration with our technology partners VSPEX creates an agile and flexible foundation enabling the delivery of new value-added services and protection to your rapidly changing IT environment and business needs.

VSPEX offers a robust set of solutions that support your data protection requirements offering Backup/Recovery, Replication, Archive, Business Continuity, and Application Mobility for your most mission critical environments. VSPEX Data Protection includes several data protection methodologies making it possible to select the optimal level based upon application requirements. Ultimately, the decision of what combination of data protection functionality to use depends on your data availability requirements.

