

HP STOREONCE RECOVERY MANAGER CENTRAL

Author: Chris M Evans

A TECHNICAL OVERVIEW AND COMPETITIVE COMPARISON

First Published: August 2015
Latest Update: August 2015
Document ID - LB1WP0068
Release 1

This White Paper was commissioned by HP Inc and written & produced by Langton Blue Ltd

Table of Contents

EXECUTIVE SUMMARY	3
INTRODUCTION	4
OBJECTIVE	4
AUDIENCE	4
CONTENTS OF THIS REPORT	4
THE CHALLENGE OF VM BACKUP	5
CONVERGED DATA PROTECTION	6
HP STOREONCE RMC: DEEP DIVE	8
DEPLOYMENT	8
SYNTHETIC FULL BACKUPS	9
MULTI-STREAMING SUPPORT.....	10
SYSTEM REQUIREMENTS & LICENSING	11
ISV ALIGNMENT.....	11
COMPETITIVE COMPARISON	12
VMAX.....	12
XTREMIO.....	12
CONCLUSION	13
CALL TO ACTION.....	13
REFERENCES	14
HP DOCUMENTATION	14
EMC DOCUMENTATION	14
MORE INFORMATION	15
THE AUTHOR.....	15

Executive Summary

The traditional approach to backup, based on agent technology, was developed in a time when application servers were physical machines, each with their own compute, memory, storage and networking resources. As server virtualisation became the dominant application deployment model, the derived benefits of hardware consolidation meant that traditional backup was no longer tenable and new approaches were required.

In the virtual world, a virtual machine, or VM is nothing more than a collection of files, representing what was previously a physical asset, such as a disk drive or LUN on shared external storage. Backups can be achieved through taking snapshots of the files themselves, either at the hypervisor or storage layer. Hypervisor snapshots provide data integrity, whereas array snapshots deliver the data mover performance required to deal with synchronising large amounts of changed data tracking.

One benefit offered by the original agent-based approach was in identifying and backing up only changed data since the last complete (or full) backup was taken. By comparison, simply backing up a snapshot provides no level of data efficiency, until the snapshot reaches the backup appliance, where some degree of de-duplication can be achieved.

With the release of HP StoreOnce Recovery Manager Central (RMC), HP has leveraged the benefits of array-based snapshots on HP 3PAR StoreServ with the de-duplication capabilities of StoreOnce to deliver a fast and efficient backup process that optimises both data transfer across the network and storage of backup images in the StoreOnce system. The StoreOnce RMC appliance running as a virtual machine is able to orchestrate the creation of HP 3PAR StoreServ snapshots and extract only the changed data within the snapshot for transmission to the StoreOnce backup system.

Snapshot data within the StoreOnce system is combined with an original full-image copy to provide what are known as synthetic backups; a full restore image made from overlaying the snapshot changes onto the full image – without having to ship the snapshot in its entirety.

HP StoreOnce RMC integrates into a VMware vSphere environment, providing the ability to snapshot, backup and restore virtual machines in a significantly more efficient manner than other solutions from competitors such as EMC. The ability to perform minimal impact backup/restore tasks allows RMC to be used by application owners to orchestrate backups based on business need. It also allows developers to use backup images to seed test and development environments far quicker than copying data or restoring from traditional backup.

In summary, StoreOnce RMC provides an integrated, optimised backup solution, exploiting the benefits of both the HP 3PAR StoreServ and StoreOnce architectures, delivering features for both data protection and application development.

Introduction

Objective

This report investigates the issues involved in ensuring that timely backups of data in virtual server environments are achieved, specifically where HP 3PAR storage is deployed. The report includes a detailed discussion of HP StoreOnce Recovery Manager Central (RMC), a software product from HP that integrates data protection between HP 3PAR storage and HP StoreOnce System. RMC currently supports virtual server environments, with the underlying technology designed to expand future capabilities into providing comprehensive data protection for applications.

Audience

Decision makers and technical influencers in organisations looking to streamline and improve the efficiency and reduce the cost of their backup solutions for virtual servers using HP 3PAR storage systems. The report is also applicable for those organisations looking to transform their data protection strategy to integrate rich application support.

Contents of This Report

- **Executive Summary** – a summary of the background and conclusions derived from Langton Blue research and analysis.
- **The challenge of VM backup** – a discussion on the challenges of taking efficient backups of virtual machines when using centralised storage platforms.
- **Converged data protection** – an outline discussion on bringing the benefits of external shared storage arrays together with backup solutions in an intelligent and integrated approach.
- **HP StoreOnce RMC: Deep Dive** – a detailed discussion on the implementation of StoreOnce RMC with HP 3PAR storage.
- **Competitive Comparison** – a brief comparison of the backup approach between 3PAR/StoreServ and EMC storage and BURA products.
- **Conclusion** – a summary of the benefits of HP StoreOnce RMC.

The Challenge of VM Backup

The widespread adoption of server virtualisation has proved to be a huge benefit for the IT industry, in terms of the efficient use of technology resources and improved time to delivery. Where server hardware used to take days or weeks to be approved for purchase and deployment, today a virtual machine (VM) can be spun up in minutes, fully configured to the needs of the end customer.

The ability to deliver on-demand VMs provides developers and other end users a significant productivity boost but also introduces issues around ensuring timely data protection of these resources. Physical servers were typically backed up across the network or through snapshots on shared storage infrastructure via proxy backup targets. In the virtual world, deploying backup agents on virtual machines makes little sense due to the potential effort and cost for what may be a very transient entity. Within the hypervisor there are opportunities to use changed data tracking as a method of capturing updates to virtual machines, negating the need to take repeated full image backups. However, these features can be slow and cause performance problems for the application, which is something that should be avoided on highly shared infrastructure.

Shared external storage continues to be a powerful way to implement a data protection strategy, specifically when using features such as integrated snapshots. Backups taken from snapshots are “point-in-time”, removing the backup window associated with traditional file-based backups over the network. In particular, on the HP 3PAR StoreServ platform, the benefits include:

- Offload of the snapshot process to the storage array, resulting in zero performance impact to the hypervisor and guests. HP 3PAR snapshots are managed through the manipulation of in-memory data structures (metadata) on the 3PAR appliance and so can operate at a much more granular and frequent basis than can be achieved using a hypervisor.
- Offload of the backup process and in particular the I/O load associated with taking backups from hypervisor-based snapshots. I/O load can be managed much more effectively at the storage array level.
- HP 3PAR StoreServ arrays are *thin aware*, including snapshots. This ensures that any snapshot will initially consume no additional storage space when first taken and after that point will only consume space as data is updated. In addition, snapshots benefit from all of the existing space efficiency technologies in HP 3PAR systems, including zero-page detect/reclaim, data deduplication and compression.

Although snapshots are an efficient method for creating point-in-time backup copies, there are a few caveats that have to be considered. Snapshots live within the array and are dependent on the source image from which they are derived, as they share common (unchanged) data. This means there is no protection with snapshots from hardware problems, such as data protection failure within a RAID group. In addition, if the source volume is deleted, the snapshots go too, making them unsuitable for the long-term retention of backup images. A snapshot is a point in time and typically “crash consistent” copy of a VM, however resilient data recovery requires application consistency that needs the use of additional software.

The goal for effective data protection is to be able to use snapshots to generate backups that can be stored and managed through a secondary storage device that is specifically designed to handle backup data. These systems (such as HP StoreOnce) are designed for high scale and high efficiency, implementing data de-duplication and federated data management that allow systems to grow organically over time without needing to perform complex data migrations.

Converged Data Protection

The challenge for the backup process is to use snapshots to generate backup images on the backup appliance in the most efficient manner and of course, to enable fast restore. The performance of a backup system directly impacts two metrics - RPO and RTO.

- **RPO – Recovery Point Objective** is a measure of the concurrency of data. An RPO of zero means data in the backup system is 100% up to date in comparison to the production environment. RPO values are typically measured in minutes.
- **RTO – Recovery Time Objective** is a measure of the time taken to restore data back to the recovery point. Depending on the recovery method, RTO values could be instant, or measured in minutes or hours.

A typical approach to data protection using snapshots is to implement a proxy server that acts as the data mover, taking the snapshot image and moving it to the backup platform. In this instance the data mover sees the snapshot in its entirety, or what looks to be a full copy of the data itself. This method of backup is highly inefficient in a number of ways:

- A full image of the VMs being backed up are both read from the storage and transported across the network, even if only a small fraction of the data on the VM has changed.
- The backup data has to be processed in full, to identify and make savings from deduplication and compression. Some backup applications deploy an agent onto a server that pre-processes the backup data before sending unique blocks to the backup system, however these consume system resources and have to be recreated in the event of a system loss.

The optimum solution is to transfer only the changed data in the snapshots across the network to the backup appliance. In order to achieve this, both the storage appliance and the backup system must have an awareness of the data structure and be able to track the snapshot changes against an initial backup image (typically achieved through metadata). When the data transferred is highly optimised, the backup appliance can be placed at a geographically distant location to the primary storage, without significant network costs.

Converged backup is a term used to describe the close integration of the primary storage and backup systems. Implementing convergence allows the backup process to be aware of exactly what constitutes changed data as updates are tracked and moved to the backup platform. Understanding the data at this level allows backup to be implemented in the most optimum way, ensuring that only changed data is transmitted across the network from the primary system to the backup device.

Convergence should be more than just providing data awareness to the backup application. When backup becomes a benefit rather than a burden to IT systems, then there is an opportunity to delegate responsibility for backup and restore away from the operational backup team and allow application and VM owners to manage their own backup and restore process. This also applies to the virtualisation administrator and can offer capabilities such as the ability to backup a VM at creation time, providing a recovery point before any applications are deployed. This delegation allows IT departments to improve their efficiency and deliver to stringent RPO/RTO requirements.

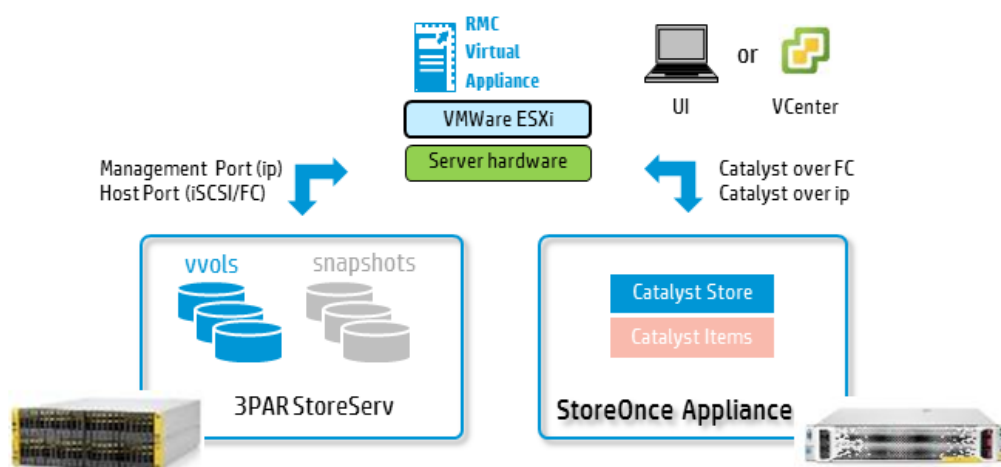
Delegation of backup tasks is achieved either through integrating the backup software into existing processes (in the case of virtual machines into VMware vCenter) or through the use of APIs

(application programming interfaces). Both of these solutions allow for either the automated or manual creation of backups without the intervention of the backup administrator.

An additional benefit is the ability to take a backup and use it as a method for generating VM images as part of a test/development cycle. Backups are an excellent source of production data to use for test, development and user acceptance testing, as directly copying production data would have an impact on the performance of production systems. However, restoring data from a backup image has a number of issues. Firstly, there is a problem with the time taken to restore all of the data from a VM or image backup across the network. Second, the large-scale movement of data means there is a risk of affecting production performance as the data is written to the array.

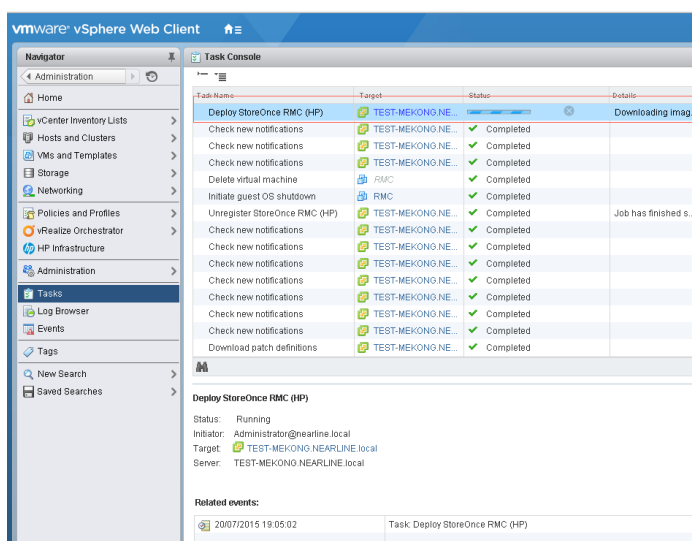
Snapshots provide a fast way to create VM image copies. They have the advantage of being lightweight (only changed/updated blocks of data consume physical space) and are efficient to take on the storage array. However, snapshots typically operate at the LUN level, which in VMware deployments translates to a datastore that can contain one or more VMs. As a result, bringing a duplicate datastore online, extracting only the required VM, moving it to a permanent location and deleting the datastore, becomes an arduous task.

RMC integration with VMware simplifies this entire process, providing a VM-centric view for administrators to be able to backup and restore entire VMs without having to manually manage the physical datastore/LUN issues. The process is simple and non-intrusive enough to allow VM owners to make their own replicas without risking production performance problems.



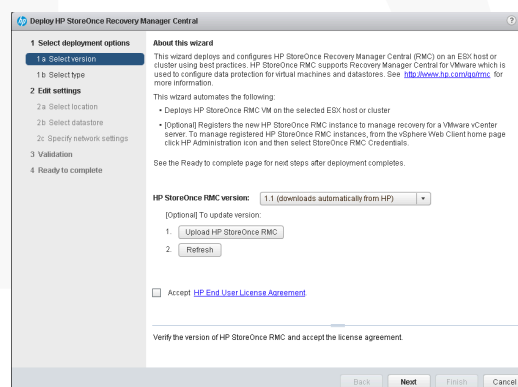
HP StoreOnce RMC: Deep Dive

HP StoreOnce Recovery Manager Central (RMC) is a software product from HP that delivers converged backup by integrating HP 3PAR StoreServ storage and HP StoreOnce Systems. HP StoreOnce RMC 1.1 for VMware (known as RMC-V) supports the ability for customers to snapshot and backup virtual machines from VMware vSphere using HP 3PAR snapshot technology, known as HP 3PAR Virtual Copy. RMC can be used on standard 3PAR Virtual Volumes, in which case the backup copy will appear as a “crash-copy” on restore. Alternatively, with VMware integration using HP OneView, virtual machines can be paused using the VMware VADP backup API, providing an application consistent backup image. RMC integration into HP OneView for vCenter allows the backup and restore processes to be managed at the granularity of a virtual machine or datastore and is therefore focused on data protection policies at the application level.



Deployment

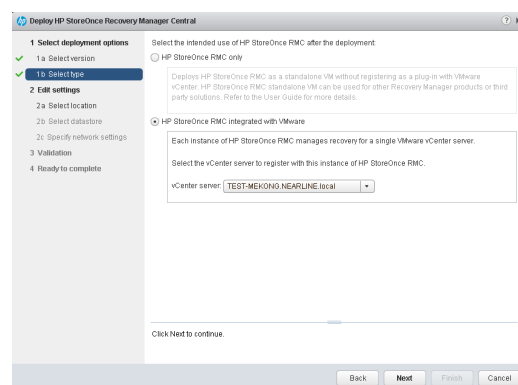
Recovery Manager Central is delivered as a vSphere virtual machine that can sit within or alongside the virtual infrastructure being backed up. RMC is downloaded from HP and deployed using HP OneView for vCenter (this is the preferred installation method). It is also possible to directly deploy the OVF package using either the vSphere Web Client or the native vSphere Windows Client. The installation and registration is performed using the HP OpenView for vCenter plugin, which also manages any installed RMC instance. Integrating the installation procedure within vCenter means VM administrators can easily perform the installation process.



RMC acts as both the initiator of a snapshot backup and the data mover, removing the need for a separate backup ISV (Independent Software Vendor) product. . The Express Protect feature allows RMC direct access to the 3PAR StoreServ metadata. With full backups the snapshot is created and then RMC moves the data blocks to the StoreOnce appliance. Data is part deduplicated in the RMC VM and only new data is actually copied to the StoreOnce catalyst store. If incremental backups are selected only the differences between the current and last snapshots are moved. Data is split into 4K blocks in the RMC VM and a hash generated. The hash is used to compare which blocks are new and need to be moved to the catalyst store. Unchanged blocks are not moved and a count is

increased in the StoreOnce appliance by embedded catalyst 'client' software embedded in the RMC code.

StoreOnce Catalyst is a proprietary HP protocol that enables backup clients to de-duplicate at source and transfer only unique data across the network for storage in a StoreOnce Catalyst store. The HP Express Protect feature of RMC utilises StoreOnce Catalyst within the RMC VM to ensure that data traversing the network is not only the changed data on the 3PAR Virtual Volume, but is unique within StoreOnce. The result is an extremely efficient, quick and ultimately scalable backup process. In comparison, third party software that doesn't have access to the 3PAR bitmap would need to read the entire snapshot image and transfer it across the network in order to back it up, even if the image was subsequently deduplicated.



The Catalyst software provides two additional benefits to the backup process; creation of synthetic full backups and improved throughput through multi-streaming of backup data.

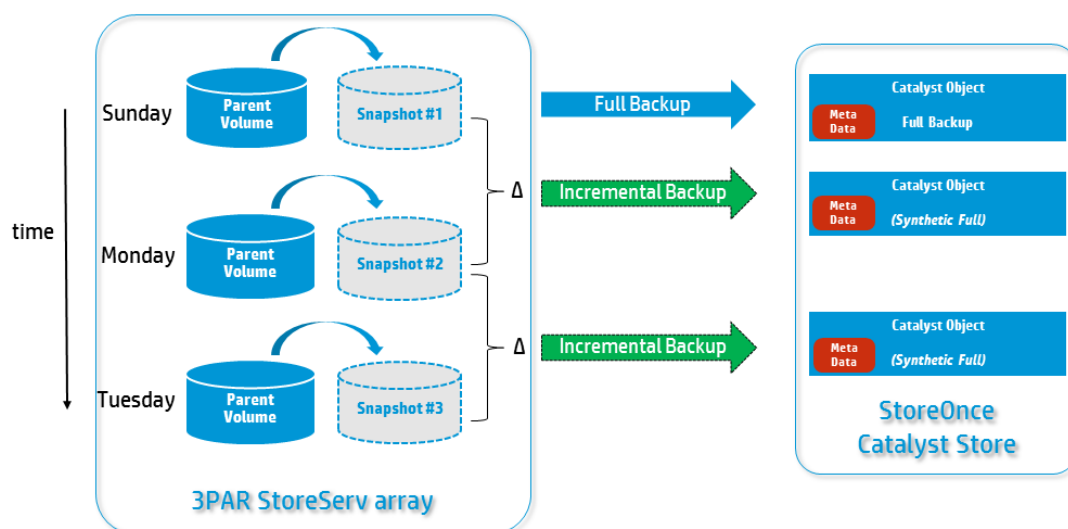
Synthetic Full Backups

Early backup applications used the concept of full backups followed by incremental backup as a method of minimising the overhead of network traffic on the backup process. A full-system restore then required accessing the initial full copy and applying all the subsequent incremental copies up to the required restore point. In this model, restores can be time consuming (especially when data isn't on disk) and error prone if any incremental is lost or damaged.

Disk-based backup systems allowed the ability to perform full backups and de-duplicate the data on ingest into the backup appliance. This provides an efficient use of backup storage (and the ability to safely perform a full restore) but is inefficient on network traffic and not scalable.

HP StoreOnce Catalyst provides the ability to de-duplicate the data at source (within the RMC VM) using Catalyst client software. Data transferred across the network consists of metadata, mapping any data already stored on the HP StoreOnce Catalyst store, plus new unique data. Within the StoreOnce system, the metadata and unique content is used to construct "synthetic full" backups that can be used to perform full restores.

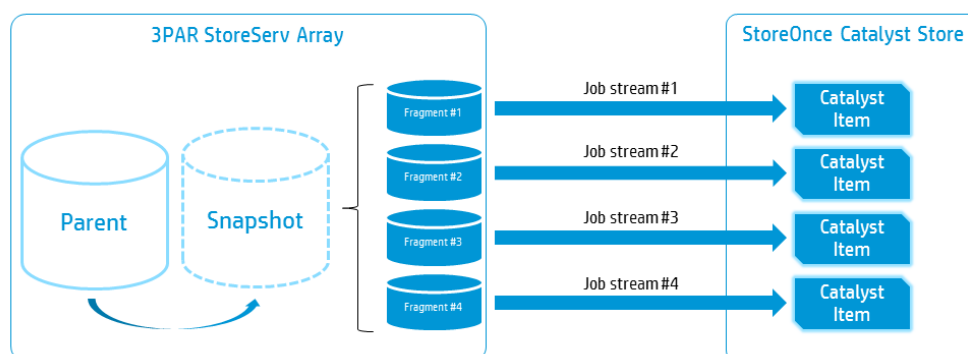
Synthetic backups have a great advantage over the traditional full and incremental model. There is no dependence between synthetic backup images (as each is independently created) and so any of them can be deleted with no impact. This flexibility means backup images can be taken more frequently and used to seed development environments with very little additional resource overhead. With 3PAR StoreServ and StoreOnce, backups do not have to be restored to the original 3PAR platform and have the ability to be restored to other storage systems, providing a degree of data mobility.



RMC full and incremental backups

Multi-streaming Support

Multi-streaming in StoreOnce Federated Catalyst provides the ability to improve performance when transferring data into a StoreOnce Catalyst store. This represents a significant improvement over existing proxy-based snapshot backups, which are typically single-streamed. HP StoreServ data is always stored in fixed block and because RMC has unique access to the internal array data it is possible to efficiently divide into multiple data streams.



In HP tests, StoreOnce backups were shown to provide a 17x improvement in backup times compared to backups using traditional ISV software. For restores, the recovery times were improved by a factor of 5x compared to traditional methods.

System Requirements & Licensing

Recovery Manager Central is supported on all HP 3PAR 7000, 8000, 10000 and 20000 series arrays with limited support for older F-Class and T-Class models (these do not support Express Protect). Snapshot management is available with HP 3PAR OS releases 3.1.2 MU3, 3.1.2 and 3.2.1. Express Protect requires HP 3PAR OS 3.2.1 MU1 as a minimum and licences for both 3PAR Virtual Copy and StoreOnce Catalyst. RMC for VMware is a free upgrade for existing Recovery Manager for VMware customers. RMC 1.1 now supports vSphere 6.0 and HP 3PAR Remote Copy snapshots. All physical and virtual StoreOnce systems are supported, running either 3.11.2 or 3.12.X software releases.

RMC can be accessed either through a dedicated web application (based on the “piano key” dashboard visualisation being introduced across HP management tools) or through integration within VMware vCenter through HP OneView for vCenter.

ISV Alignment

At first inspection it may seem that RMC represents a replacement for traditional ISV backup software, however the two will have the ability to work in conjunction. RMC provides the capability to access only the changed data within a snapshot and to deliver highly efficient data movement to HP StoreOnce, a feature that cannot be achieved by any ISV backup software. In contrast, ISV software provides a higher-level management position, delivering richer reporting capabilities than RMC. It is HP’s intention to enable RMC to be directed to take backups through an external API, allowing today’s ISV providers to support RMC as a data mover. This enables customers to use RMC and their existing ISV software as part of an integrated and heterogeneous backup strategy.

Competitive Comparison

StoreOnce Recovery Manager Central was compared to the data protection offerings from EMC, specifically covering the equivalent enterprise storage platforms VMAX and XtremIO. These platforms have been chosen for comparison as they are aimed at enterprise and all-flash workloads respectively.

VMAX

EMC ProtectPoint software provides the capability to deliver integrated backups between EMC VMAX storage and EMC Data Domain backup systems. The data within a VMAX storage array is protected using local snapshots. Snapshot images are transferred from the EMC VMAX by the ProtectPoint software to the Data Domain appliance using Fibre Channel connectivity.

The Data Domain appliance performs de-duplication on ingested snapshot data from VMAX, writing new unique data and identifying non-unique content that is then indexed and tracked through metadata. Full synthetic backups are then created from existing and new data.

ProtectPoint provides support for traditional workloads, such as physical servers running Oracle, SAP and SQL applications. The software operates through the use of an agent on each host server that is used to schedule the backup process. ProtectPoint is currently only supported on VMAX3 platforms due to the dependence on the integration between VMAX Hypermax OS and Data Domain. There is no support for virtual machine environments, for example VMware or Hyper-V. In comparison, RMC specifically supports VMware virtual server deployments.

ProtectPoint is much less efficient than Recovery Manager Central as all of the snapshot image data is transferred across the Fibre Channel network for the data ingestion process to be performed on the Data Domain appliance. In comparison, with RMC, only the changed snapshot blocks need to be read from the 3PAR array and these are then de-duplicated by the RMC VM before sending over the network the StoreOnce appliance.

ProtectPoint currently supports Data Domain models DD4500, DD7200 and DD9500. The lower end DD2200, DD2500 and DD4200 are not supported. By comparison, RMC is support across all StoreOnce systems, including those deployed as VSAs (Virtual Storage Appliances).

XtremIO

EMC currently provides no backup system-level support for the XtremIO all-flash platform. There is an integration point between XtremIO and RecoverPoint, however this is for remote replication rather than backup data protection.

Conclusion

Scalable data protection in virtual server environments requires a new approach to the process in which backup data is taken and secured. HP StoreOnce Recovery Manager Central provides HP 3PAR StoreServ and StoreOnce customers the ability to implement a highly efficient data backup regime, while ensuring that the impact on production services is at a minimum.

Compared to backup solutions from EMC, RMC is more efficient on data transfer across the network, provides better scalability through the use of StoreOnce Catalyst features and works with enterprise-class and all-flash 3PAR StoreServ solutions.

The speed and efficiency of backup with RMC means HP customers can use the RMC platform as a way of managing the creation of test and develop environments with very little additional storage resources. The VM backup and restore processes are achieved using a dedicated web application or integration with VMware vCenter, providing the ability to delegate data protection tasks to end-users.

HP continues to invest and develop RMC; release 1.1 introduces support for VMware vSphere 6.0 and HP 3PAR Remote Copy snapshots. Release 1.2 will introduce a new feature called Element Recovery Technology (ERT) that provides granularity of recovery to restore individual files from within a virtual machine.

In summary RMC provides the most flexible and efficient backup solution for customers using HP 3PAR StoreServ storage.

Call to Action

More information on Recovery Manager Central can be found on the HP RMC website at <http://www.hp.com/go/rmc>

In addition, the following links provide additional background information:

- [“1+1 = 3 when you integrate 3PAR and StoreOnce with RMC” Video](#)
- [HP StoreOnce Recovery Manager Central – Product Video](#)
- [HP StoreOnce RMC Chalk Talk](#)
- [HP StoreOnce Recovery Manager Central – Short Take \(Part 1\)](#)
- [HP StoreOnce Recovery Manager Central – Short Take \(Part 2\)](#)
- [HP StoreOnce RMC for VMware solution brief](#)
- [HP StoreOnce RMC Technical White Paper](#)
- [HP StoreOnce RMC Datasheet](#)
- [HP StoreOnce RMC Brochure](#)
- [HP StoreOnce RMC Business White Paper](#)

References

The following publicly available documents and references were used in the production of this report.

HP Documentation

- [Application-managed data protection for HP 3PAR StoreServ Storage](#) – (DocID: 4AA5-5865ENW)
- [HP StoreOnce Recovery Manager Central Data sheet](#) – (DocID: 4AA5-5172ENW)
- [Converged Data Protection HP StoreOnce Recovery Manager Central](#) – (DocID: 4AA5-5173ENW)

EMC Documentation

- [EMC ProtectPoint: A Detailed Review](#) (DocID: H13261)
- [EMC ProtectPoint: Data Sheet](#) (DocID: H13167)

More Information

For additional technical background or other advice on replication technologies, contact enquiries@langtonblue.com for more information.

Langton Blue Ltd is hardware and software independent, working for the business value to the end customer. Contact us to discuss how we can help you transform your business through effective use of technology.

Website: www.langtonblue.com

Email: enquiries@langtonblue.com

Twitter: [@langtonblue](https://twitter.com/langtonblue)

Phone: (0) 845 275 7085

Post:

Langton Blue Ltd
133 Houndsditch
London
EC3A 7BX
United Kingdom

The Author

Chris M Evans has worked in the technology industry since 1987, starting as a systems programmer on the IBM mainframe platform, while retaining an interest in storage. After working abroad, he co-founded an Internet-based music distribution company during the .com era, returning to consultancy in the new millennium. In 2009 he co-founded Langton Blue Ltd (www.langtonblue.com), a boutique consultancy firm focused on delivering business benefit through efficient technology deployments. Chris writes a popular blog at <http://blog.architecting.it>, attends many conferences and invitation-only events and can be found providing regular industry contributions through Twitter ([@chrismevans](https://twitter.com/chrismevans)) and other social media outlets.

No guarantees or warranties are provided regarding the accuracy, reliability or usability of any information contained within this document and readers are recommended to validate any statements or other representations made for validity.

Copyright© 2015 Langton Blue Ltd. All rights reserved. No portions of this document may be reproduced without the prior written consent of Langton Blue Ltd. Details are subject to change without notice. All brands and trademarks of the respective owners are recognised as such.