

Data Protection for Virtualized Environments

IBM Spectrum Protect Plus Delivers a Modern Approach

By Steve Scully, Sr. Analyst

February 2018



Evaluator Group



Evaluator Group

Executive Summary

Virtualization is continuing its march through the datacenter, fundamentally changing IT environments forever. No longer just used for server consolidation and VDI workloads, virtual machines (VMs) are the foundational technology for building a flexible and scalable data center for most IT organizations. In addition to VM growth, IT organizations have seen tremendous growth in the amount of data they have. Forecasts are for that growth to continue at a rapid pace and even accelerate as organizations address the deluge of data coming from next generation applications such as rich media or IoT networks. Virtualization is also changing datacenter responsibilities where many IT admins are generalists that approach tasks from a VM or application-centric perspective. The result is that data availability has become a responsibility shared by VM admins, application owners and backup admins.

Many legacy backup solutions were not designed to meet the needs of virtualized environments. The traditional structured approach of scheduled backup and recovery jobs does not meet the needs of more frequent and flexible backups by VM admins and application owners. These users often make frequent changes to VMs and applications and want the ability to run an ad hoc backup before implementing the changes. They also want the ability to quickly restore the environment to a known good point without needing to schedule a recovery job.

All these pressures are driving IT organizations to re-evaluate their data protection and data management strategies. The Evaluator Group has consulted with several IT organizations recently whose current data protection plans are no longer meeting the company's goals and need to be overhauled. Other organizations have inquired about new approaches to protect all the different types of data they have in a variety of places (at the edge, in the datacenter, in the cloud). Still others are looking to get more value from their data by reusing copies (or better managing the copies they already use) for additional activities such as DevOps, reporting and analytics.

In this Technical Insight, we look at IBM Spectrum Protect Plus, one of the latest solutions for protecting, retaining and reusing data in virtualized environments. It is the latest member of the IBM Spectrum Storage family of software-defined storage solutions. IBM Spectrum Protect Plus takes a modern approach to data protection by focusing on Operational Recovery of data rather than being a backup solution. It utilizes snapshot technologies to save immutable copies of the data which can be rapidly restored or repurposed for new uses. IBM recently launched the latest version of Spectrum Protect Plus improving the availability of the data repository, adding blueprints to simplify hardware selection and improving the performance of item level restores from SQL and Oracle databases.

IBM Spectrum Protect Plus

IBM has been focused for several years on delivering a broad set of software-defined storage capabilities to solve a variety of modern datacenter challenges. IBM recently launched IBM Spectrum Protect Plus, providing data protection, operational recovery, and data reuse for today's virtualized environments. It allows IT admins to better meet Recovery Time Objectives (RTO) and Recovery Point Objectives (RPO) with faster backup creation using policies, faster recovery through Internet like searches, right-click restores and restore in place to eliminate data movement. Additionally, it helps IT leaders increase the value of their data by using copies for new workloads such as DevOps, Test/Dev, Reporting and Analytics.

IBM Spectrum Protect Plus is designed from the ground up with a modern, virtualized architecture. It is designed to be easy to deploy and simple to operate by general IT admins and by end users. The user interface is tailored for non-backup administrators and uses policies to define data availability. It creates and maintains a global catalog of all the VMs and files it manages and allows for rapid recovery at very granular levels.

Designed for Data Reuse and Data Management

IBM Spectrum Protect Plus takes a modern approach to data protection by focusing on operational recovery of data (versus simply creating a backups) and is designed to be a data reuse and data management platform. Spectrum Protect Plus utilizes snapshot technologies to save copies of the data which can be rapidly restored or repurposed for various uses described above. IT admins determine the granularity of those snapshots to meet their RTO and RPO requirements.

Highlights

- Modern approach to operational recovery for data protection and data reuse
- Enables self-service access to data for new workloads such as DevOps, Reporting and Analytics
- Unlocks additional business value by reusing backup copies
- Centralized, role-based management allows authorized users access to the data based on their needs
- Integrates with IBM Spectrum Protect for enhanced retention and compliance

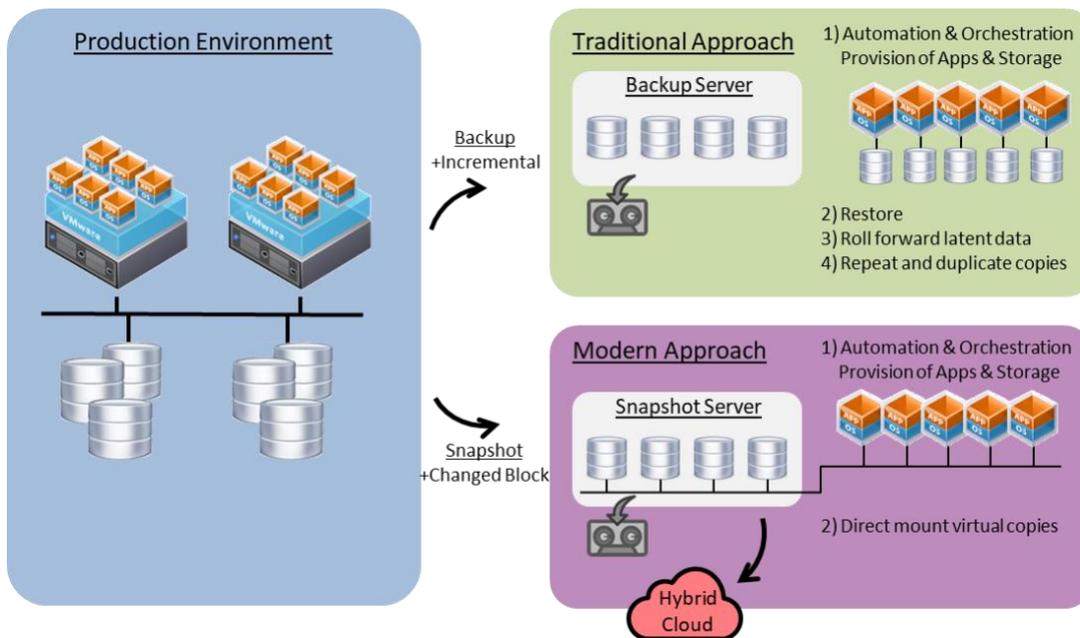


Figure 1. Modern Approach to Data Protection (Source: IBM)

Figure 1 shows the key differences between using a traditional approach with backup or media servers to create and store backup copies of data. Traditional backups can have their place in a data protection strategy, but they typically produce multiple copies of the data and can be slower to recover from. The modern approach used by Spectrum Protect Plus leverages data copy technologies via Application Programming Interfaces (APIs) in the hypervisors to efficiently snapshot the data. The snapshot server saves these recover points as immutable copies of the data in its native format.

If needed, data can be recovered instantly from any of the recovery points stored in the system. Or customers can spin up a new virtual machine with any recovery point as a read-write clone and use that VM for any application that needs a copy of the production data. These copies can also operate with various technologies such as data masking, which will protect sensitive company or personal information in the data while using the copies for DevOps and testing.

The flexible nature of these recover points allows them to be easily deployed into hybrid cloud environments as well. Optional API-based management capabilities enable IT admins to integrate Spectrum Protect Plus processes into workflow automation with scripts and third-party tools.

Designed In and For Virtualized Environments

As a more recent entry to the market, Spectrum Protect Plus was designed from the start in and for virtual environments. It currently supports VMware and Microsoft Hyper-V and is designed to enable hypervisor admins and application admins to manage their own backups and perform their own restores when needed. An important feature in the design is a centralized, role-based management of the backup and recovery processes. Administrators can define what actions different roles in the organization are allowed to perform and what access to the data they have.

Spectrum Protect Plus has updated the typical user experience with a Service Level Agreement (SLA)-centric approach that reducing the complexity found in other data protection solutions. It is designed to allow non-backup specialists to create and manage their own backups as well as service their own restores. Administrators assign a recovery service level VMs, using either predefined services levels (gold, silver, bronze) or their own customized policies to meet specific requirements. Service levels can be assigned individually, by the datastore, or globally to define all new VMs. Spectrum Protect Plus identifies any VMs that do not have an assigned recovery level and are therefore not protected.

Another differentiating feature of Spectrum Protect Plus is its global index and search capabilities. Once the availability policies are defined, Protect Plus creates a metadata catalog that records and indexes all VMs in the environment. Customers have the option to index all the files in the VMs as well. This allows IT admins to perform Google-like searches across the entire Spectrum Protect Plus environment including the ability to search across VMware and Hyper-V virtual machines in a single operation.

Evaluator Group Comments: IBM has included a number features that enhance the user experience with Spectrum Protect Plus. In addition to simplifying backup and recovery, the index and search features as well as data reuse capabilities make it more about data management than just backup.

Ease of use is an important characteristic in today's IT environments that have fewer backup specialists and more virtualization generalists. IBM Spectrum Protect Plus is simpler to use than many competitive solutions. The key for IBM will be to keep simplicity at the forefront of design as Protect Plus continues to evolve and grow over the years.

Deploys On-Premises or in the Cloud

IBM Spectrum Protect Plus is designed to protect virtual machines located on-premises to VMs running in public clouds. The Spectrum Protect Plus platform is delivered as an open virtual appliance (OVA) or self-extracting .exe for Hyper-V which provides platform administration and contains the policy catalog, role-based access controls, and the metadata information. It makes calls to the data protection APIs that exist in virtual environments and applications to create block-level, incremental forever backups based on the policy definitions. These backups are stored as immutable snapshot images in the Spectrum Protect Plus vSnap data repository.

The vSnap data repository of Spectrum Protect Plus is deployed using block storage from multiple storage vendors including all IBM block storage solutions. Customers can scale up the storage pool for more capacity (depending on the capabilities of the underlying storage hardware) or they can scale out the data repository with multiple vSnap Storage Managers.

IBM has defined several Blueprints as reference systems for deploying small, medium and large Spectrum Protect Plus solutions. These blueprints define all the hardware and infrastructure components needed to fully deploy the Spectrum Protect Plus solution. They can be used directly by customers and will likely be leveraged by IBM channel partners to develop Spectrum Protect Plus appliances.

IBM Spectrum Protect Plus has also been architected to operate with VMs in public clouds. The initial integration is with VMware operating in the IBM Cloud. For VMs and their associated data that were born in or migrated to the cloud, admins will see Spectrum Protect Plus listed as a service in the IBM Cloud Marketplace which they can add to their VMware instance with a single click. They define the protection policy at the time the VM is provisioned and it is protected right from the beginning as a service.

Evaluator Group Comments: The Spectrum Protect Plus architecture provides a scalable and flexible approach to data protection for virtual environments. The API-based management approach will allow admins to integrate with backup scripts, third party tools and application portals. It should also be a matter of time for future integrations as the virtualized environments and applications evolve.

Since many IT organizations are considering or already implementing a multi-cloud provider environment, IBM will need to integrate with other popular cloud platforms to meet customer expectations. With the architecture of Protect Plus, Evaluator Group believes this should only be a matter of time for IBM to implement.

Operates standalone or integrated with IBM Spectrum Protect

IBM Spectrum Protect Plus can be deployed standalone for use in virtual environments. Purchased standalone, Spectrum Protect Plus is simply priced by the number of VMs managed. Spectrum Protect Plus can also be integrated with the enterprise-grade IBM Spectrum Protect software for complete data protection of physical and virtual infrastructure on premise and in the cloud.

Backup copies can be offloaded to Spectrum Protect to create consolidated enterprise-wide backups for long term retention, data governance and compliance requirements. This allows Spectrum Protect Plus data to move to any of the various storage tiers (disk, tape, cloud data stores, and object storage) supported by IBM Spectrum Protect. Spectrum Protect Plus tracks the locations of its backup sets, so recovery from any tier is accomplished through the Spectrum Protect Plus user interface in a single step.

Spectrum Protect Plus is included in the IBM Spectrum Protect Suite as well as the more comprehensive IBM Spectrum Storage Suite.

Evaluator Group Comments: Ensuring that IBM Spectrum Protect Plus can be installed on an operated completely standalone is a good strategy by IBM. It expands the target market beyond the traditional enterprise to any customer (of any size) looking for a new approach to managing their virtualized environment.

IBM has a long history with enterprise data protection with their Spectrum Protect (formerly Tivoli Storage Manager) and many customers who have built their data protection processes with it. The ability to integrate Spectrum Protect Plus backups into their existing operations will be beneficial to those customers. Evaluator Group believes that customers of both will eventually want an integrated metadata catalog, enabling better discover and improved data management.

Conclusion

IBM Spectrum Protect Plus is a solution that fits a trend we are seeing in the industry with data protection moving towards various aspects of data management. No longer just a repository of backup copies, Spectrum Protect Plus offers data protection for VMs, applications and files with the ability to index and search all the data and to reuse copies of the data for other activities that derive more value from an organization's data. With its agentless architecture, SLA-approach and pre-defined policies, Spectrum Protect Plus can be deployed in very little time. It represents a good foundation for IT organizations to transition from a data protection strategy to a higher-level data management strategy.

About Evaluator Group

Evaluator Group Inc. is dedicated to helping IT professionals and vendors create and implement strategies that make the most of the value of their storage and digital information. Evaluator Group services deliver in-depth, unbiased analysis on storage architectures, infrastructures and management for IT professionals. Since 1997 Evaluator Group has provided services for thousands of end users and vendor professionals through product and market evaluations, competitive analysis and education. www.evaluatorgroup.com Follow us on Twitter @evaluator_group

Copyright 2018 Evaluator Group, Inc. All rights reserved.

No part of this publication may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording, or stored in a database or retrieval system for any purpose without the express written consent of Evaluator Group Inc. The information contained in this document is subject to change without notice. Evaluator Group assumes no responsibility for errors or omissions. Evaluator Group makes no expressed or implied warranties in this document relating to the use or operation of the products described herein. In no event shall Evaluator Group be liable for any indirect, special, inconsequential or incidental damages arising out of or associated with any aspect of this publication, even if advised of the possibility of such damages. The Evaluator Series is a trademark of Evaluator Group, Inc. All other trademarks are the property of their respective companies.