

DELL EMC DATA PROTECTION FOR VMWARE WINNING IN THE REAL WORLD

A Competitive Comparison Exposé

ABSTRACT

This white paper provides a deep dive analysis based on truly real world comparison of Dell EMC data protection vs. Veritas NetBackup for VMware backup and recovery.

November, 2016

Copyright © 2016 Dell Inc. or its subsidiaries. All Rights Reserved. Dell, EMC, and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be the property of their respective owners. Published in the USA, 11/16, White Paper, H14700.2

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

TABLE OF CONTENTS

EXECUTIVE SUMMARY	4
AUDIENCE	4
INNACCURACIES IN THE VERITAS 'BENCHMARK REPORT'	4
DELL EMC TESTING REPORT & ENVIRONMENTAL DETAILS VALIDATED BY ESG	4
Dell EMC testing methodology	4
Dell EMC data gathering methods	6
DELL EMC VS. VERITAS TESTING RESULTS	7
Not real world: Single stream	7
Real world: Multiple streams	7
Not real world: Initial backups only	8
Real world: Daily backup performance over 28 days	8
Not real world: No deduplication rates shown	9
Real world: Deduplication rates	9
Real world: Dedupe rates and replication impact	10
IN FOCUS: NETBACKUP ACCELERATOR IMPACT ON CATALOG SIZE	10
OTHER CONSIDERATIONS	12
Veritas' use of SAN transport vs. VMware API for data protection (VADP)	12
Veritas NetBackup 5330 maintenance and configuration	12
SUMMARY – WHY DELL EMC FOR VMWARE DATA PROTECTION?	12
Backup and recovery performance	12
Deduplication efficiency	12
Integration	13
Reliability	13
CONCLUSION	13
APPENDIX	14
Veritas test environment specifications	14
Dell EMC test environment specifications	14

EXECUTIVE SUMMARY

The goal of this technical benchmark testing white paper is to provide objective details surrounding the holistic view of data protection for VMware workloads in real world scenarios. The following is a comparison of these real world tests done and validated by ESG vs that same exact test harness applied to Veritas NetBackup 7.6.1 writing to a NetBackup 5330 appliance. In conclusion, it illustrates the superiority of Dell EMC's data protection solutions in realistic end user scenarios. To show the value of a real world setup, we also highlight gaps and inaccuracies with the Veritas funded benchmark test by Principled Technologies.

AUDIENCE

This white paper is intended for end users seeking the truth as well as Dell EMC sales representatives and engineers.

INNACCURACIES IN THE VERITAS 'BENCHMARK REPORT'

Veritas recently published a 'Benchmark Report' that touted results compared to their major competitors based on so called "real-world benchmark tests" by Principled Technologies. Below summarizes some of the major issues with this report:

Tested only initial backup performance, unknown size and no growth rates: In their benchmark comparison Veritas measured the initial first day back up and nothing more. They did not introduce any data change rates, nor did they stipulate how much data was being backed up. In the real world, a typical data protection policy is a 2 week or 4 week cycle, you know exactly how much data you are backing up and it changes and grows every day. Without the size of a backup, backups over time or a growth rate, one cannot gain a real understanding of the relative speed, efficiency and scalability of competing solutions.

Tested only with Avamar single node: Dell EMC's current best practice is to leverage Data Domain as the protection storage with Avamar software, so that customers can gain the efficiencies of multi-stream backup and recovery and leverage the additional benefits of Data Domain's protection storage. The benchmark tests used a single Avamar node and no Data Domain, which is not what most customers would deploy for VMware backup. In addition, although Avamar nodes can work as a standalone system, they are designed to work in a grid, which is the way customers would maximize performance without Data Domain.

Omitted critical information, not telling the whole story: There was no mention of deduplication, replication or catalog efficiency. Nor did they address that Veritas requires a net new full backup every 6 months and how you deal with this when you're also trying to simplify and meet backup windows.

Compared single stream Dell EMC recovery vs multi-stream Veritas recovery: This benchmark comparison tested a Veritas solution that leverages multi-stream recovery and compared it to a solution from Dell EMC that only leveraged single stream recovery capabilities by design. However, the results still shows that recovery from Avamar was still faster than NetBackup.

DELL EMC TESTING REPORT & ENVIRONMENTAL DETAILS VALIDATED BY ESG

Please download and review the complete ESG lab tested and validated data found here: <http://www.emc.com/collateral/analyst-reports/esg-lab-validation-report-emc-data-domain-avamar.pdf>

Dell EMC tested a Dell EMC data protection solution that included Data Domain and Avamar with the virtual environment hosted on a Dell EMC XtremIO all-flash storage array, which was validated by ESG, a leading industry analyst. Figure 1 provides a conceptual overview of the solution, which includes a VMware ESX hypervisor hosting multiple application VMs, the Avamar Virtual Edition VM, and proxy VMs. The Avamar Administrator dashboard is shown at the top. A proxy VM leverages VMware Change Block Tracking to identify changed data. Then, leveraging Data Domain Boost, parts of the deduplication process are distributed to Avamar, where client-side deduplication ensures only unique data is sent from the Avamar client to the Data Domain system. This enables significantly faster backups with reduced bandwidth requirements. The highly scalable Data Domain system can serve simultaneously as a target for other backup and archiving workloads, as shown on the right. For this testing, XtremIO was used as the primary storage platform to eliminate the possibility of any storage bottleneck.

DELL EMC TESTING METHODOLOGY

Dell EMC's testing methodology was focused on showing a true apples to apples comparison in a real world scenario. The tests performed simulated a 28 day (4 week) backup cycle using 200 virtual clients. The initial data generated reflects approximately 50

gigabytes of data per client. Approximately 4% modification and .125% net new data per client was generated for each backup cycle (per day). We then measured performance and deduplication efficiency through a full 4 week backup cycle using Avamar and an integrated Data Domain DD4500 system. We ran the following Veritas configuration though the exact same test harness (3 unique times); then compared the results using the data gathering methods described in the next section.

In order to show an apples to apples comparison, we tested Avamar (which always does 'forever full' backups by design) vs. NetBackup Accelerator (which is the closest NetBackup can come to a 'forever full' scenario).

Veritas documentation notes that you should not run an Accelerated full for more than 30 days in a row due to increased catalog size (aka "catalog bloat"). To ensure that speed and dedupe tests would be apples to apples comparison as possible; we tested Veritas Accelerated Full Backups every day (instead of the typical NetBackup protection cycle - a combination of full backups and incrementals). In addition, we tested the Dell EMC solution using Avamar integrated with a DD4500 because the DD4500 is most similar to the 5330 appliance (instead of using a higher end Data Domain system with industry leading performance and scale).

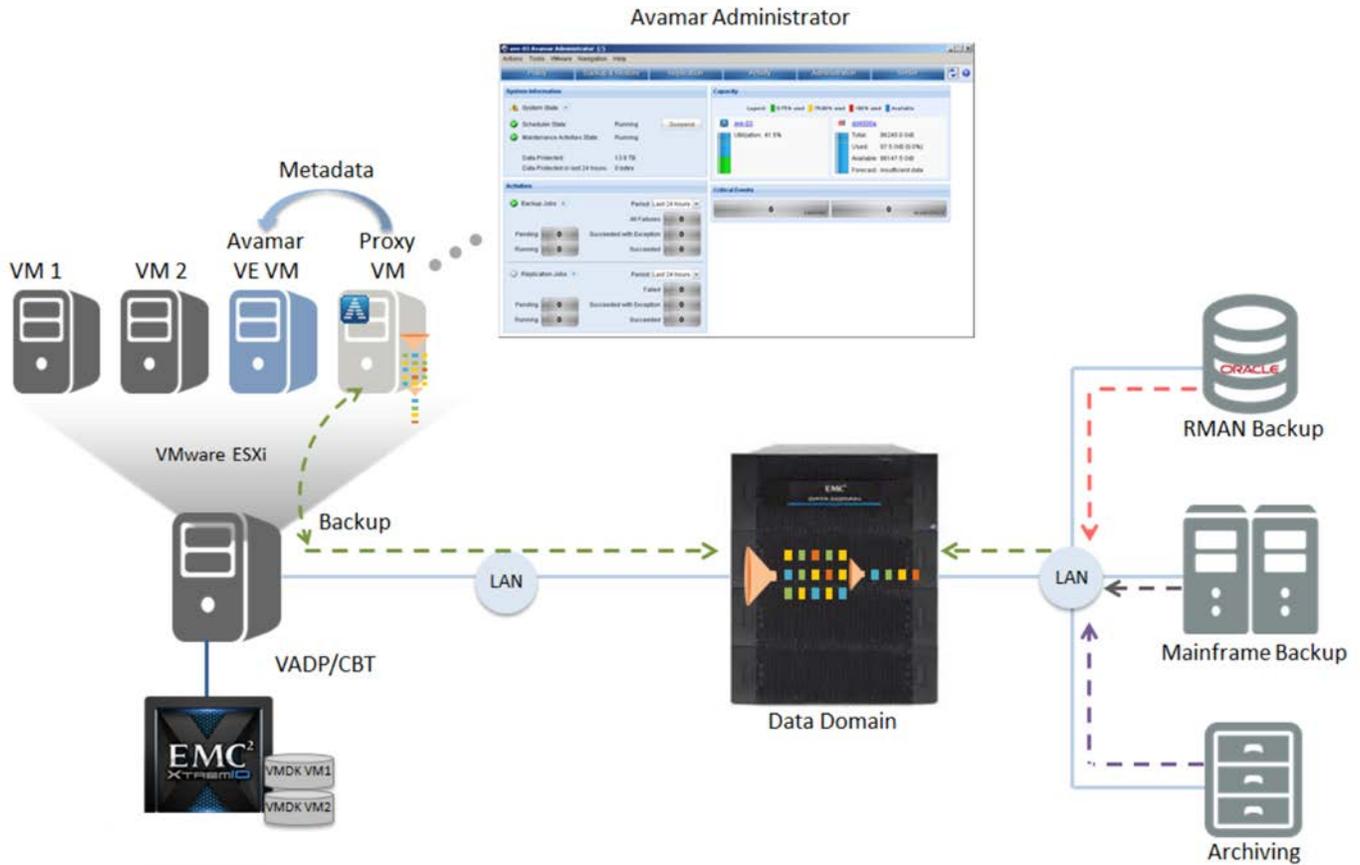


Figure 1: Testing Configuration Diagram

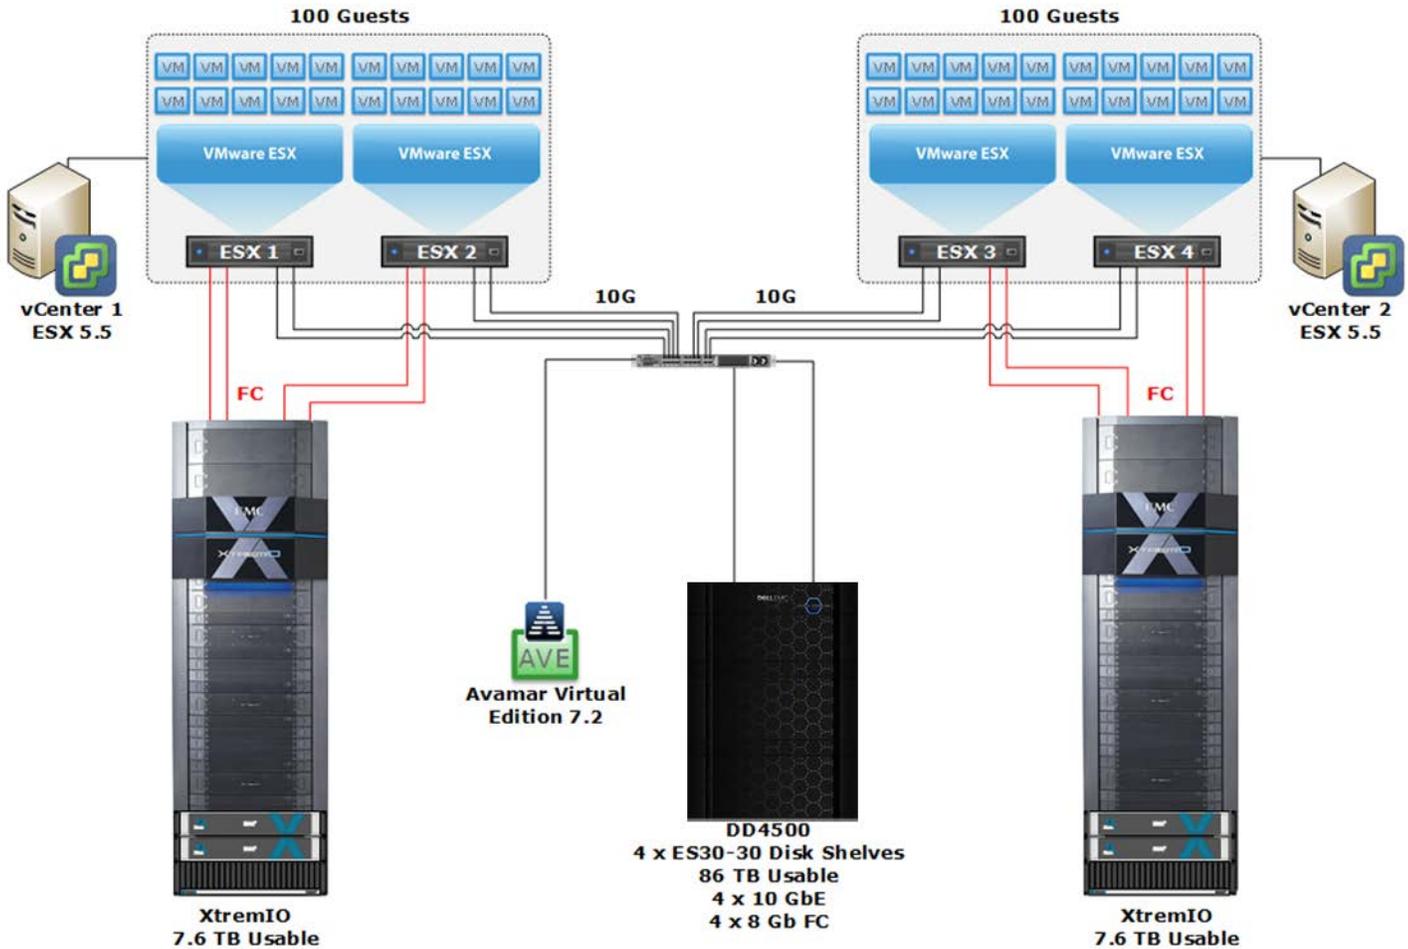


Figure 2: Testing Configuration Diagram

DELL EMC DATA GATHERING METHODS

Although the lack of reporting tools on the NetBackup 5330 hampers data gathering, for each run the following data was gathered:

- Screenshot of the Deduplication Report from the NetBackup 5330 Shell Menu. **Note:** This is a point-in-time report.
- Export of the NetBackup Activity Logs to text file.

All test runs were orchestrated using scripts to automate the backup and restore jobs as well as the gathering of critical information such as backup and restore runtimes, which were measured via the common Unix/Linux time utility.

Deduplication measurements were pulled directly from the Data Domain system using the `filesys show compression` command in the scripts controlling the backup jobs. The data was gathered immediately after each backup job was completed and written to a log file for later analysis.

CPU and other system resource information for various components in the test environment were gathered using Dell EMC Data Protection Advisor (DPA), which supports most backup applications, including Veritas NetBackup and Dell EMC Avamar. DPA was also utilized as a second method to verify deduplication and runtime results.

DELL EMC VS. VERITAS TESTING RESULTS

NOT REAL WORLD: SINGLE STREAM

Veritas testing compared their multi-stream recovery to Dell EMC's single stream recovery, which is not an apples to apples comparison. However, let's look at the chart provided in Veritas's Benchmark Report:

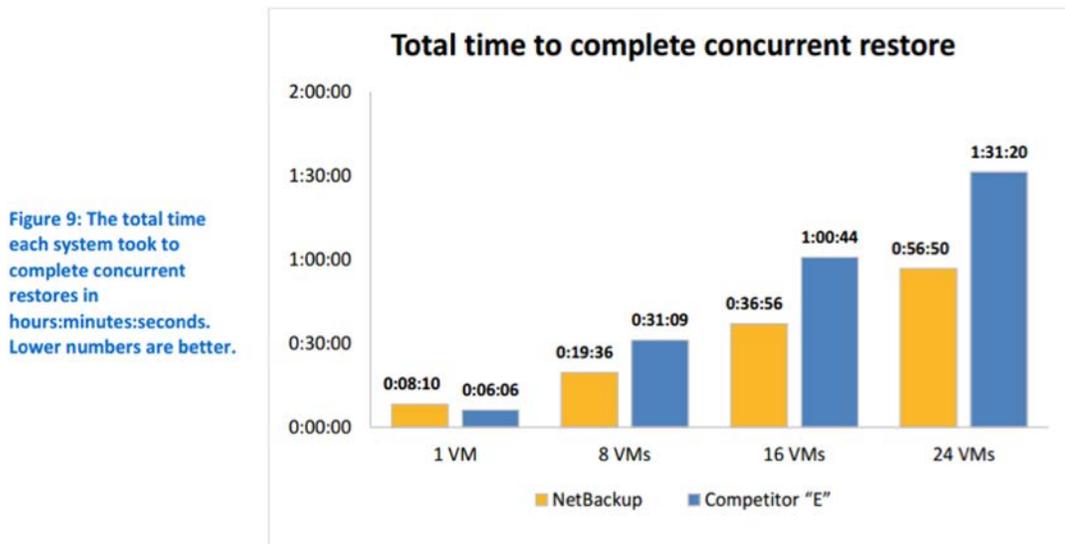


Figure 3: Veritas Single Stream Backup Results

The only apples to apples comparison on this chart is the first (1 VM restore), which shows Dell EMC outperforms Veritas. The other tests comparisons shown here are unrealistic as they show single stream recovery vs. multi-stream recovery.

REAL WORLD: MULTIPLE STREAMS

Dell EMC testing for recovery was performed with an apples to apples multi-streamed comparison. We tested recovery of 31 virtual clients, with an initial data generation of approximately 50 gigabytes per client. An initial backup was taken of all 31 clients after which one data modification cycle was performed to simulate data change to the clients between the last backup and the restore. These tests were then performed to determine recovery performance while simultaneously recovering 1, 10 and 20 clients (VMs). Figure 4 summarizes the average recovery performance provided by each solution with 10 clients (which is representative of the each recovery test):

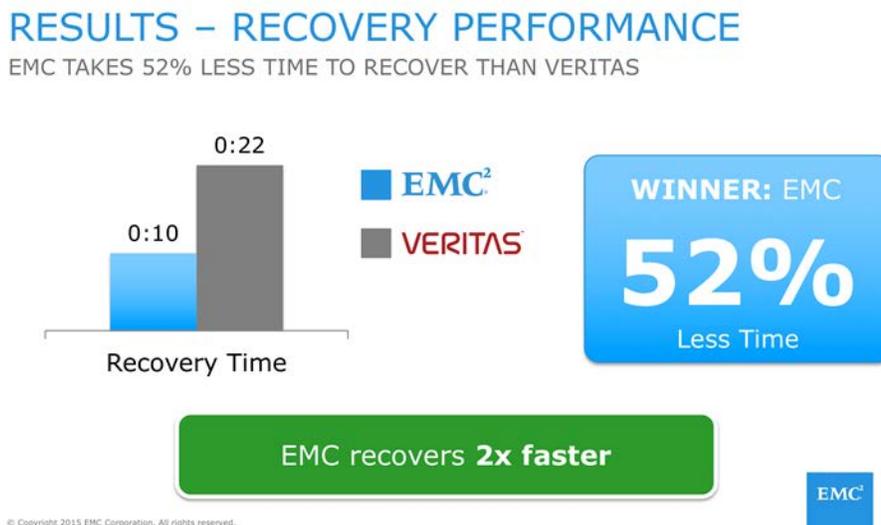


Figure 4: Dell EMC Multi-Stream Recovery Results

NOT REAL WORLD: INITIAL BACKUPS ONLY

Veritas performance testing showed a relatively meaningless comparison of 1 day of backups for each solution. Standard backup policies includes 2 to 4 weeks of retention – including a combination of full and incremental backups each day. Therefore, a real world apples to apples comparison requires at least 14 to 28 backups over time.

REAL WORLD: DAILY BACKUP PERFORMANCE OVER 28 DAYS

EMC testing for backup performance simulated a real world scenario with the following parameters for each solution:

- Testing used 200 virtual clients, initial data generation of approximately 50 Gigabytes per client
- Approximately 4% modification and .125% new data per client for each backup cycle
- Number of concurrent backup jobs increased with each run
- We ran test iterations of 50, 100, 150, and 200 concurrent clients

Figure 5 summarizes the average backup performance of each solution after the initial backup:

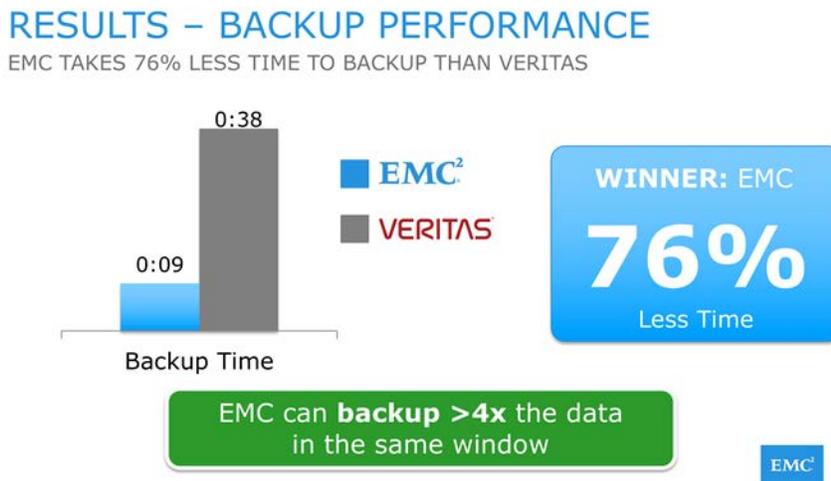


Figure 5: Dell EMC Multi-Stream Backup Results

In addition, Figure 6 shows the comprehensive results for backup performance with increasing workloads. Regardless of the % rate of change, NetBackup backup time stays consistent where Avamar backup time decreases. Veritas Benchmark testing only focused on the initial backup, so it failed to highlight these results.

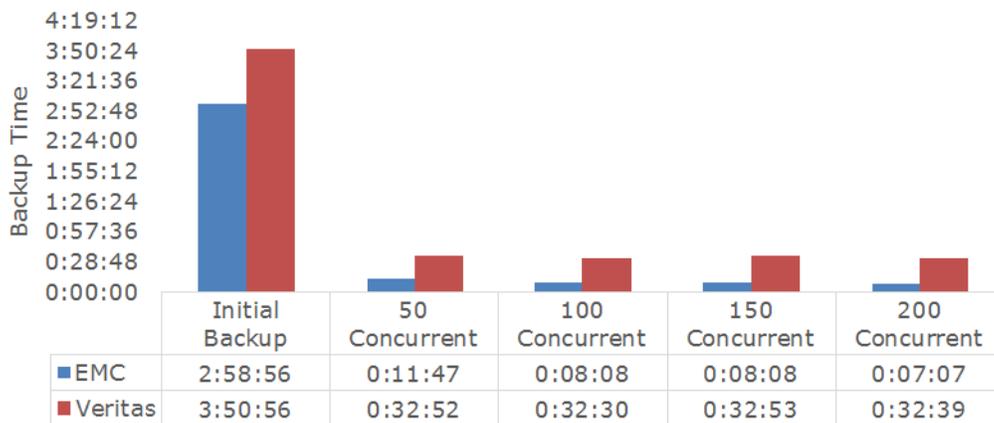


Figure 6: Dell EMC vs. Veritas Backup Performance

NOT REAL WORLD: NO DEDUPLICATION RATES SHOWN

Veritas testing neglected to show deduplication rates at all. Ensuring maximum dedupe rates for your protection storage solution is critical to minimizing footprint and bandwidth. This is a critical metric that all backup appliances should share freely.

REAL WORLD: DEDUPLICATION RATES

Dell EMC Testing compared deduplication rates of each solution, which highlighted the impact of Veritas' fixed-length vs. Dell EMC's variable-length deduplication. Figure 7 shows that after 28 days, the Dell EMC solution achieved a deduplication rate of 73x compared to Veritas 30x. This means that a NetBackup appliance would require 2.5x the capacity to protect the same size environment as Data Domain.

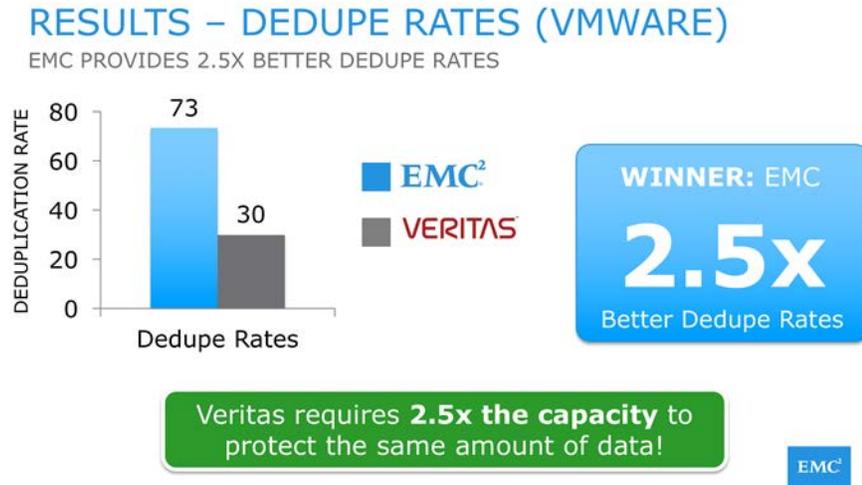


Figure 7: Dell EMC Deduplication Rates Test Results

In addition, Figure 8 shows that NetBackup deduplication efficiency plateaus after some time, while Data Domain variable-length deduplication continues to improve over the 28 days.

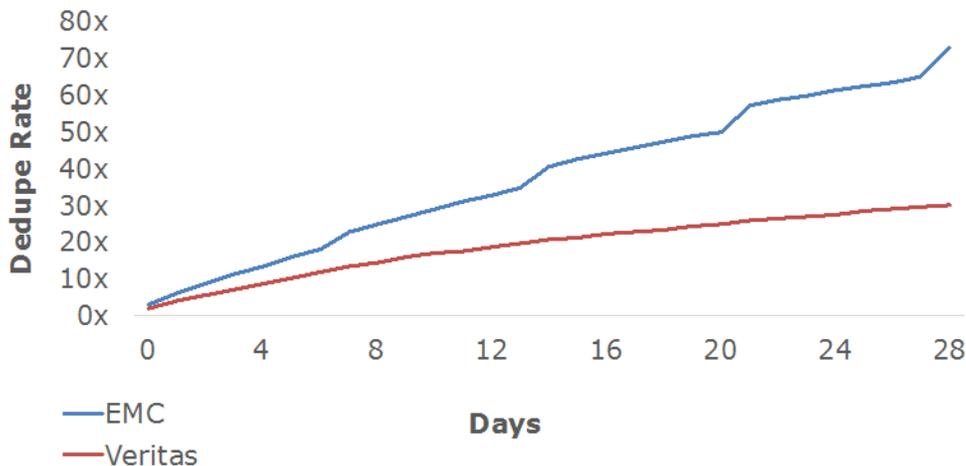


Figure 8: Deduplication Rate over Time

REAL WORLD: DEDUPE RATES AND REPLICATION IMPACT

In addition to the storage footprint, deduplication rates dramatically impact bandwidth requirements as well. Figure 8 illustrates that due to lower deduplication rates achieved, Dell EMC replicates 60% less data than NetBackup. For example, if the Dell EMC solution had to replicate 200 GB a day of changed data, it would take approximately 10 hours over a 45Mbps connection. By comparison, the Veritas solution would have to send 500 GB (2.5x more), which over a the same 45Mbps connection, would take 26 hrs. To achieve the same replication time with Veritas, a customer would have to pay for more replication bandwidth month after month. In addition, if a customer is using NetBackup, the bandwidth requirements would increase.

RESULTS – REPLICATION IMPACT

EMC REPLICATED 60% LESS DATA PER DAY THAN VERITAS

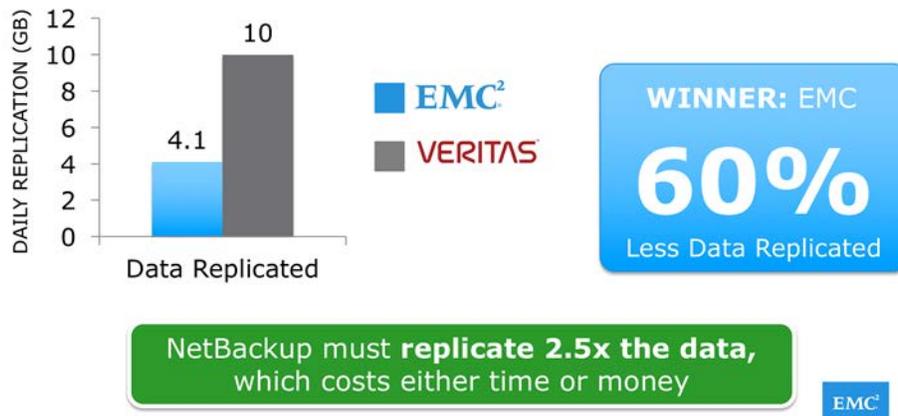


Figure 9: Dell EMC vs. Veritas Replication Impact Comparison Results

Warning: "Good Enough" Dedupe is *Never Good Enough* and it will cost you

You should question solutions where the terms "good enough" are used with respect to data deduplication. The significance of even a minor improvement in overall dedupe has a considerable downstream impact on overall solution cost. The impact can be felt by the amount of disk you need to acquire in order to meet retention requirements. In addition, this has a direct impact on the amount of rack, power, cooling, network ports, Fibre Channel ports and, most importantly, replication bandwidth required to replicate not only the daily change rate but also the catalog.

IN FOCUS: NETBACKUP ACCELERATOR IMPACT ON CATALOG SIZE

Avamar was purpose-built for full backups every day without compromise. To provide 'incremental forever' backups like Avamar's 'forever fulls', many customers use NetBackup Accelerator to perform daily full backups. However, NetBackup Accelerator has numerous drawbacks including dramatic catalog growth and the requirement for a full backup every 6 months. This is noted in Veritas best practices documentation, and our testing found the following:

Using NetBackup Accelerator for daily full backups causes 3 to 5x catalog growth. Although accelerating a full from an incremental is not a new concept, when implemented with a traditional backup application architecture, catalog bloat will occur. The more you use NetBackup Accelerator, the worse the catalog bloat gets and the longer the replication window if using AIR. The bigger the catalog get, the bigger risk you have for errors to occur within that catalog. NetBackup Accelerator's excessive catalog growth is noted in their documentation and also says that it needs to be regularly maintained.

To manage this bloat, a NetBackup Accelerated backup needs to be “rescanned”, which means a real full backup needs to be done, every 6 months. For VMware backups, CBT is reset a full backup to ensure nothing is missed, which means you need to plan for an even larger backup window each time. This requires complex scheduling to balance making your backup window and still ensuring every server is protected.

Figures 10 and 11 summarize the challenges that exist with NetBackup Accelerator:

- **Performance:** NetBackup Accelerator is still slower than Avamar with Data Domain (9 minutes vs. 38 minutes).
- **Complexity:** Real full backup required every 60 days, which requires extra time and planning.
- **Catalog Bloat:** NetBackup's catalog is 5x bigger than Avamar, increasing media server and replication costs.



Figure 10: Backup Performance Comparison

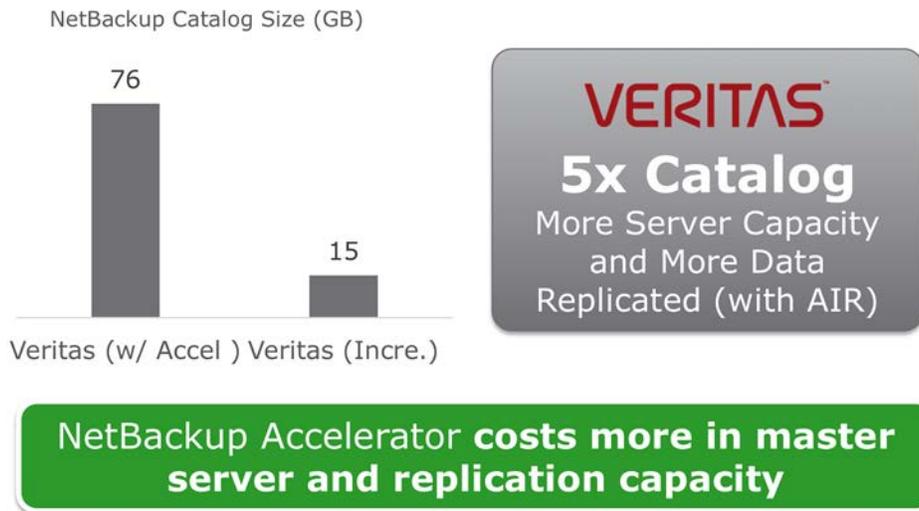


Figure 11: Backup Catalog Efficiency Comparison

OTHER CONSIDERATIONS

VERITAS' USE OF SAN TRANSPORT VS. VMWARE API FOR DATA PROTECTION (VADP)

The Veritas benchmark test used array-based snapshot technology to snap the NetBackup SAN, then leverage a protocol called "SAN Transport" to move data to the backup appliance for dedupe and storage. There are two things to keep in mind about this approach to VMware backup:

1. Leverage SAN transport at the array layer depends on each array vendor's unique API's and coordination with VMware. This means if a customer changes their storage layer, or moves to a cloud provider where they don't have access to the storage layer anymore, then they will have to re-architect their data protection strategy. In contrast, Dell EMC leverages VMware best practice for data protection including virtual proxies and VMware API for data protection (VADP), offering a more flexible and agile solution.
2. With the advent of vVols, VMware will no longer support the "SAN Transport" approach to backup. Therefore, any organization which has based their VMware data protection strategy on array based snaps and this old protocol will have to redesign to take full advantage of vVols.

VERITAS NETBACKUP 5330 MAINTENANCE AND CONFIGURATION

In Dell EMC benchmark testing, we observed that there was almost no performance improvement after the garbage collection process completed. In addition, cleaning is done every day during idle time periods, but active workloads take precedence over maintenance/cleaning processes. Depending on the daily workloads this will impact performance over time, results in even slower backup and recovery with NetBackup.

SUMMARY – WHY DELL EMC FOR VMWARE DATA PROTECTION?

Based on the ESG validated benchmark testing we complete, we have demonstrated that Dell EMC offers superior results in the real world where it counts: backup and recovery performance as well as deduplication efficiency. In addition, it is the only solution that was designed to ensure recovery of data with industry leading data integrity. Finally, the powering integration between Avamar and Data Domain provides the use of deployment of an integrated appliance, with the flexibility of an open protection storage system.

BACKUP AND RECOVERY PERFORMANCE

Dell EMC provides best in class backup by performing daily full backups while only processing the equivalent of an incremental and does so without compromising with side effects like catalog bloat. With Avamar you never have to schedule another real full backup job or plan for the time required to complete a full backup periodically. Dell EMC performance is maximized by combining Avamar 'forever fulls' with Data Domain Boost and Data Domain Stream-informed Segment Layout (SISL). Overall, this results in:

- Dell EMC takes 52% less time to recover than Veritas
- Dell EMC takes 76% less time to backup than Veritas

DEDUPLICATION EFFICIENCY

Data Domain variable-length deduplication ensures customers get the highest dedupe rates possible. In our testing, Dell EMC provided 2.5x better deduplication rates than Veritas for VMware backup. Superior dedupe matters because it means:

- Less protection storage needed
- Less bandwidth for replication required
- Lower operational complexity at scale with fewer systems to maintain
- Improved deduplication efficiency
- Reduction in overall costs over time

INTEGRATION

Along with superior performance and efficiency, Dell EMC offers choice and flexibility to protect data in the most efficient way possible. As stated by ESG, "Avamar and Data Domain are both industry-proven data protection solutions, and as an integrated solution they are even more powerful. Data Domain provides a single, highly reliable and scalable platform for consolidated backup and archiving that enables administrators in different roles to manage their own protection tasks without having to buy and manage separate infrastructure stacks. Avamar integrates with VADP to leverage CBT, speeding backup and restore, while delivering scalable VM backup using load-balanced proxy servers. Data Domain delivers data efficiency with variable-length global deduplication. Together, they provide an integrated, appliance-like solution that is easy to deploy and reduces capital and operational costs, while forming a protection foundation that matches the flexibility and agility of today's data center."

RELIABILITY

Dell EMC Data Domain systems is the only solution that validates 100% of your data to ensure it remains ready and available for restore throughout its life on the system. Dell EMC's patented Data Domain Data Invulnerability Architecture ensures that your data is always recoverable through end-to-end data verification, continuous fault detection & self-healing, fault avoidance & containment, and full file system recoverability. In comparison, NetBackup lacks any serious data integrity functionality. NetBackup, as do many competitors, offer some way to scan a small percentage of your data on a non-reoccurring basis, but warn you of huge performance hits if you do it or offer nothing beyond a onetime MD5 hash check and RAID 6.

CONCLUSION

In summary, the combination of Avamar with Data Domain performs faster, more efficient backups and recoveries, provides higher deduplication efficiency, offers better flexibility through broader integration, includes better data integrity protection and provides a simpler day to day operational efficiency than Veritas. All of these factors reduce end user costs and give Dell EMC customers confidence that they can quickly and reliably recover their mission critical data when needed.

APPENDIX

VERITAS TEST ENVIRONMENT SPECIFICATIONS

Backup Software	Disk Target
NetBackup 5330 v2.6.1 NetBackup 7.6.1 Media Server Role Installed 384 GB RAM 20 CPU Cores	NetBackup 5330 Appliance version v2.6.1 148TB Storage 5 Shelves – 12 drives per Shelf 4 x 1 GbE Interfaces 4 x 10 GbE Interfaces 5 x 8 Gb FC
Virtualization	Servers
VMware vCenter 5.5.2	4 x Cisco UCS 240 OS: SLES 12.3
Storage	Network
2 x Dell EMC XtremIO, 7.56 TB usable each	10GbE

Appendix 1: Veritas Testing Configuration Details

DELL EMC TEST ENVIRONMENT SPECIFICATIONS

Backup Software	Disk Target
Avamar Virtual Edition 7.2 2 x vCPU 6 x vRAM 900 GB Disk Space	Data Domain DD4500 DD OS 5.6.0.3 4 x ES30-30 Disk Shelves 4 x 10 GbE Interfaces (2 used in testing)
Virtualization	Servers
VMware vCenter 5.5.2	4 x Cisco UCS 240 OS: SLES 12.3
Storage	Network
2 x Dell EMC XtremIO, 7.56 TB usable each	10GbE

Appendix 2: Dell EMC Testing Configuration Details