

Harnessing the Performance of Dell EMC VxRail 7.0.100

A Lab-Based Performance Analysis

November 2020

H18552.1

White Paper

Abstract

This white paper details our testing of vSphere and vSAN 7.0 U1 with VxRail 7.0.100. It shows how the latest VMware releases affect key performance indicators and drive performance improvements for VxRail business workloads.

Dell Technologies Integrated Products

Copyright

The information in this publication is provided as is. Dell Inc. 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 software described in this publication requires an applicable software license.

Copyright © 2020 Dell Inc. or its subsidiaries. All Rights Reserved. Dell Technologies, Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Intel, the Intel logo, the Intel Inside logo and Xeon are trademarks of Intel Corporation in the U.S. and/or other countries. Other trademarks may be trademarks of their respective owners. Published in the USA 11/20 White Paper H18552.1.

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

Contents

- Introduction4
- Test results7
- Summary.....14
- References.....15
- Appendix: Test team and setup16

Introduction

VxRail overview

Dell EMC VxRail hyperconverged infrastructure (HCI) is the only fully integrated, preconfigured, and tested HCI system that is optimized for VMware vSAN technology for software-defined storage. Co-engineered with VMware, VxRail systems are integrated with existing VMware ecosystem and cloud management solutions. Those integrated solutions, which are optional, include VMware vRealize, VMware NSX, VMware Horizon, and all other solutions that are a part of the vast and robust VMware vSphere ecosystem. Our joint engineering with VMware and investments in VxRail HCI System Software provide customers with an optimized turnkey system at reduced risk.

VxRail HCI System Software

VxRail HCI System Software is a suite of integrated value-added software elements that extends VMware native capabilities and further reduces operational complexity. It encapsulates much of the key functionality that differentiates VxRail systems from vSAN Ready Nodes and other HCI solutions in the market. VxRail HCI System Software provides automation and orchestration for Day 1 deployment and Day 2 system-based operational tasks.

VxRail HCI System Software consists of multiple integrated software elements that extend VMware native capabilities to deliver a seamless and automated operational experience. Within VxRail HCI System Software, SaaS multicluster management provides global visualization, simplified health monitoring, and multicluster management through a cloud-based web portal. These features build upon life cycle management services to increase operational efficiency, especially for customers who have a large footprint of VxRail clusters and for whom managing at scale has been challenging.

Full-stack testing, validation, and upgrades

The VxRail engineering team fully tests and validates the entire stack. The team performs more than 25,000 test hours for every major release, minimizing testing and validation time and costs for our customers. The HCI System Software electronic compatibility matrix ensures that every iteration of possible upgrade paths is rigorously tested. This process offers multiple direct upgrade paths, which allows customers to bypass interim updates that they may have missed.

VxRail provides faster, more efficient upgrades in two unique ways:

- Leapfrogging capability—Customers can choose any VxRail version they want, with the assurance that it is in a validated state, and seamlessly bypass any in-between releases.
- Pre-staged upgrades—Each node is ready to go and provides estimated upgrade times, allowing customers to schedule upgrades at the optimal time for their organization.

VxRail customers can take advantage of the latest updates by applying the full-stack, single-click upgrades. Customers no longer have to verify hardware compatibility lists, run test and development scenarios, sequence updates, and so on. Intelligent life cycle management functionality automatically updates clusters with pre-validated, pre-tested

software and firmware components, ensuring that the infrastructure is in a continuously validated state.

VxRail HCI supports multiple generations of hardware in the same cluster, so that current and future customers can:

- Upgrade from many versions of VxRail and VMware software.
- Buy the technologies they need, when they need them.
- Future-proof their infrastructure.
- Tap into generational software enhancements without being forced to swap out infrastructure.

Latest updates and performance improvements

The most recent vSphere and vSAN updates bring several new features to the market. The features provide VxRail customers with significant performance improvements across nearly all workload types and up to 58 percent IOPS increase with database workloads. The performance increases are visible across all workload types, small block and large block, for both reads and writes.

To enable customers to take advantage of these performance enhancements, Dell offers single-click system upgrades across multiple generations of VxRail dating back to mid-2018. For a list of these software versions, see the [Dell EMC VxRail Release Notes](#) (login required). For additional VxRail documentation, see [References](#).

For customers who have already invested in the physical hardware technologies, no further hardware upgrades are needed. Customers can realize these performance improvements through the VxRail upgrades that are part of the VxRail maintenance contracts.

Key points:

- The VxRail platform is co-engineered with VMware and Dell for a fully integrated experience.
- Recent updates to vSAN have resulted in VxRail performance increases.
- Customers can easily take advantage of these increases with full-stack, single-click life cycle management.
- The ability to leapfrog releases gets customers from point A, B, or C to point Z so they can quickly realize significant performance increases.
- Dell Technologies provides software updates within 30 days of VMware software releases.
- Dell will continue to enhance LCM so that customers can easily take advantage of evolving technologies.

**Document
purpose**

This document describes our performance tests of VxRail 7.0.100 with vSAN 7.0 U1. It examines the most relevant performance test results and the impact that the performance improvements can have on various workload types and VxRail configurations.

Audience

This white paper is for organizations that want to consider moving their business application workloads to a VxRail system. It is also meant for Dell Technologies sales engineers and IT professionals who are assisting customers in their hybrid cloud journey.

**Authors and
contributors**

David Glynn, Senior Principal Engineer, VxRail Technical Marketing

Tony Foster, Technical Marketing Innovation Engineer

Bill Leslie, Senior Manager, VxRail Technical Marketing

Nelson Fonseca, Senior Principal Engineer, Product Technologist

Lakshmi Nagarajan, Principal Engineer, Product Technologist

Inigo Olcoz, Consultant Engineer, HCI & CI Technical Marketing

Robert Percy, Principal Engineer, HCI & CI Technical Marketing

Test results

Introduction

Our tests with vSAN 7.0 U1 on the VxRail platform showed significant performance gains, which, in turn, can result in simplified operations, faster host restarts, and the potential for cost savings.

Analysis of the test results led us to the following primary conclusions:

- vSAN is significantly faster.
- Moving to RAID 5 provides significant benefits.
- vSAN compression-only is nearly penalty-free.
- vSAN 7.0 U1 provides faster drive rebuild times.
- NVMe cache drives provide maximum storage performance.

Conclusion #1: vSAN is significantly faster

In VxRail 7.0.100, we have observed one of the most significant performance gains ever in the VxRail platform, leading to better performance and efficiency for mission-critical workloads, databases, and workload consolidation. We measured gains in various storage operations ranging from RAID 1 sequential reads to RAID 6 random writes. The improvements include up to a 69 percent increase in IOPS (for sequential writes in RAID 6) and 38 percent increase in throughput for random writes in RAID 6 configurations.

vSAN 7.0 U1 provides sweeping improvements. Significant highlights include:

- Considerable performance increases on RAID 5 and RAID 6 compared with previous vSAN versions
- Overall gains that significantly reduce the need to use RAID 1 configurations for many workloads

The following figure captures many of the improvements that we measured, representing IOPS and throughput gains in many types of storage operations on the VxRail platform, from vSAN 7.0 to vSAN 7.0 U1.

Significant across-the-board performance gains

Probably the most significant jump in performance gains for VxRail ever




Figure 1. Front-end I/O performance improvement of vSAN 7.0 U1 over 7.0

Conclusion #2: Moving to RAID 5 provides significant benefits

vSAN improvements to RAID 5 configurations bring it to near parity with RAID 1, providing 25 percent more usable capacity.

Many vSAN users have implemented RAID 1 configurations because the need for performance outweighed the need for efficiency. In VxRail 7.0.100, performance in erasure coding RAID 5 (“RAID 5”) configurations is greatly improved and almost entirely removes the trade-off between performance and efficiency.

vSAN 7.0 U1 enhances the read performance of RAID 5, bringing it in line with RAID 1 read performance. Improvements in the software for RAID 5 are responsible for a dramatic increase in performance between VxRail 7.0 and VxRail 7.0.100.

This increase is consistent across RAID 5 reads and writes. However, as our tests show, write-heavy operations might still benefit from a RAID 1 (mirrored) configuration as opposed to a RAID 5 configuration. For workloads that mainly consist of random or sequential reads, the performance enhancements offered by VxRail 7.0.100 could merit upgrading VxRail environments from previous versions and migrating to a RAID 5 configuration.

Additionally, by migrating to a RAID 5 configuration to capture performance benefits, the RAID set’s usable capacity increases from 50 percent with a mirrored configuration to 75 percent with a RAID 5 configuration. In other words, using a RAID 5 configuration provides 25 percent more usable capacity. Thus, environments in need of additional storage capacity can realize both capacity and performance increases by upgrading to VxRail 7.0.100.


You can update 36 different VxRail releases with a one-click upgrade to this latest VxRail release, going back to VxRail 4.5.211. This capability makes it easy to realize RAID 5

performance gains in most environments by simply scheduling a maintenance window to perform an upgrade.

The benefits that VxRail 7.0.100 provides to RAID 5 are significant, making VxRail with RAID 5 a powerful option in today’s modern data centers. The benefits include up to a 30 percent increase in random read performance compared to VxRail 7.0, as highlighted in the following figure. Combining this type of performance gain with the 25 percent gain in usable capacity over a mirrored configuration makes using RAID 5 much more advisable than before. We can make this assessment even without considering the enhancements to data reduction data services (described in [Conclusion #3: vSAN compression-only is nearly penalty-free](#)).

Percentage performance gains across all FTM

RAID 5/6 throughput performance gap vs. mirror is greatly reduced



64 KB block size, no dedupe, no compression

Figure 2. Performance gains for RAID 1, RAID 5, and RAID 6

Conclusion #3: vSAN compression-only is nearly penalty-free

The new compression-only feature has these significant benefits:

- Saves space while improving performance for compression-friendly workloads such as high-performance databases (that is, SAP and SQL).
- Stabilizes environments and reduces risk. By eliminating the deduplication requirement from the feature, users no longer have to worry that enabling better disk utilization will cause an entire disk group to fail if they lose just one drive.

Compression-only is a new and separate offering in vSAN 7.0 U1. Upgrading to vSAN 7.0 U1 provides database workloads with a performance boost of up to 58 percent while simultaneously unlocking the space-saving efficiencies of deduplication and compression.

The peak performance gains in vSAN 7.0 U1 are only part of the performance-improvement story. The following figure shows the performance curves with an RDBMS database workload using RAID 5. The best performance curves are long, flat lines. The further the curve is to the right, the higher the IOPS, and the flatter it is, the lower the latencies—both beneficial traits in storage performance. Overall, we can see that vSAN 7.0 U1 has performance gains, even when data services are used.

Massive improvements for database workloads (RDBMS)

Smoother and more predictable latency



Figure 3. Workload simulation—RDBMS, RAID 5

Before reviewing the gains with data services, consider the gains without them. Comparing the two green lines in Figure 3, we can see that peak performance gained a significant 58 percent additional IOPS with just the vSAN upgrade. That increase in peak IOPS also came with a 39 percent reduction in latencies.

Regardless of the platform and service, there is always a trade-off between the benefit that a data service provides and the impact that it has. The trade-off tipping point varies from customer to customer, but with deduplication and compression on vSAN 7.0 U1, that tipping point has shifted significantly.

Comparing the two blue lines in Figure 3, we can see that peak performance with deduplication and compression enabled gained a significant 36 percent additional IOPS and reduced peak latency by more than 16 percent. Further, compare vSAN 7.0 U1 with deduplication and compression enabled (the solid blue line) and vSAN 7.0 without deduplication and compression enabled (the dashed green line). Even with deduplication and compression enabled, the IOPS gain was more than 15 percent, and the drop in latency was more than 18 percent. Not only can performance be gained, but so can additional capacity.

Finally, consider the gray line in Figure 3, which represents a new vSAN data service. In vSAN 7.0 U1, VMware has made compression available as a data reduction data service by itself, without deduplication. This data service enables demanding workloads (such as

database workloads) to benefit from compression when they cannot take advantage of the space savings of deduplication. As compression-only did not exist in vSAN 7.0, we cannot show those gains. However, we can compare this latest release with and without compression-only enabled. Under vSAN 7.0 U1, enabling compression dropped peak IOPS a little more than 9 percent and increased latency a little more than 14 percent. For many, these levels of reduced peak IOPS and increased latency are a reasonable trade-off for the gains that compression can bring. The trade-off is especially justifiable for compression-friendly workloads such as large I/O block sizes in mixed workloads, RDBMS and OLTP databases, and VDI workloads using linked clones. If we compare vSAN 7.0 U1 with compression-only to vSAN 7.0 without data services, we still see a performance gain of 45 percent for peak IOPS with latencies dropping more than 30 percent.

If you are using deduplication on an early VxRail release, you might want to:

- Upgrade to the latest release.
- Disable deduplication.
- Enable compression-only to get between 50 and 70 percent more IOPS at lower latencies and a better failure scenario. For details, see the VMware blog [Space Efficiency Using the New “Compression only” Option in vSAN 7 U1](#).

Conclusion #4: vSAN 7.0 U1 provides faster drive rebuild times

vSAN 7.0 U1 provides up to 7x faster drive rebuild times, which:

- Reduces risk and exposure during drive failure
- Decreases the time required to get back to full performance
- Allows users to take advantage of increasing drive sizes

The improved performance of VxRail 7.0 U1 tips the scales for users, providing more options when balancing the trade-offs between cost, efficiency, and requirements.

As previously mentioned, regardless of the platform and service, there is always a trade-off between the benefit that a data service provides and the impact that it has.


In vSAN 7.0 U1, compression is available as a data reduction method on its own, apart from the deduplication and compression data service. With compression-only enabled, the trade-off for data reduction has been reduced. As discussed in [Conclusion #3: vSAN compression-only is nearly penalty-free](#), we saw that the performance impact was significantly reduced. Under compression-only, the failure of a capacity drive only affects that device and does not take the entire disk group offline as it does with deduplication and compression. This smaller failure domain shrinks the impact of a disk failure and reduces the potential amount of data that vSAN needs to rebuild when a disk failure occurs.

The following figure shows that vSAN 7.0 U1 provides important gains in drive rebuild transfer rates. It also shows that enabling deduplication and compression has a significant impact—up to more than 7x—on drive rebuild transfer rates. However, the impact on drive rebuild transfer rates is significantly reduced when compression-only is enabled. This faster transfer rate, coupled with the reduced failure domain and reduced performance impact, makes compression-only a viable option for even the most demanding

compression-friendly workloads. Such workloads include mixed workloads with large I/O block sizes, RDBMS and OLTP databases, and VDI workloads using linked clones.

Drive rebuild transfer rate improvements

Compression-only rebuild rates are 7x faster than dedupe + compression



With 15% OLTP 4 KB

Figure 4. Improvements in drive rebuild transfer rates

Conclusion #5: NVMe cache drives provide maximum storage performance

VxRail 7.0.100 unlocks the potential of NVMe functionality in cache drives, especially when supporting write-intensive applications such as in-memory databases. With the most recent improvements to vSAN, VxRail can now take full advantage of the speed and predictability delivered by high-specification drives, outperforming standard SSDs by up to 60 percent.

With previous vSAN versions, faster cache drives did not provide meaningful performance improvements, although they did provide more predictable latency. The following figure shows that at lower performance demands, the differences between cache drives is insignificant. As higher and higher performance is required, however, higher specification cache drives are required.

Maximum performance requires fastest cache drives

Higher IOPS, lower & more consistent latency



Figure 5. Maximum performance

We ran this test based on RAID 1, with no compression or deduplication set up with an RDBMS type workload. During the test, we saw the most significant performance improvement when testing large block operations on the most performant drives. This improvement can greatly benefit write-intensive operations such as those that take place in in-memory database workloads. To protect data from the risk of memory failure, in-memory databases persist in-memory data to storage media and periodically flush all changed data from memory to the underlying data volumes.

With mixed workloads and small block operations, performance was comparable regardless of the drive specifications. However, when we pushed the system, we began to see significant performance deviations depending on the drive’s specifications. In short, below a certain IOPS threshold, the cache drive is not a determining factor. Above that threshold, however, higher performance requires a high performing cache drive.

When we need to stress our VxRail system to achieve maximum storage performance, we have to include the highest-performing drives that are available. The investment will yield huge latency improvements and increased maximum obtainable throughput.

Summary

The vSphere and vSAN 7.0 U1 release provides numerous performance gains, which have been tested and documented by the VxRail performance engineering team at Dell Technologies. The benefits include lower cost, greater redundancy, and greater fault tolerance.

vSAN 7.0 U1 provides performance increases across all workload types, small block and large block, for both reads and writes, with greater increases in writes. Most significantly, performance is nearly equal to the performance of RAID 5 workloads. Because the storage policy can be changed without taking a workload offline, VxRail customers can now tap into the extra 25 percent capacity that RAID 5 provides compared to RAID 1.

VxRail customers can also now take advantage of compression-only, reducing the failure domain that is present with deduplication and compression, while still achieving most of the top-end performance improvements that 7.0 U1 provides. Customers who still want to use deduplication and compression will see better performance in comparison to vSAN 7.0 without those data services. Another benefit of compression-only is that it reduces throughput by only 16 percent, thus improving drive rebuild times.

VxRail customers who adopted NVMe cache drives and Intel Optane cache can see significant performance improvements compared to SAS cache drives. These drive technologies offer more predictable latency and higher top-end performance, especially for database workloads.

Given the recent findings and new capabilities, there is no better time to upgrade to the most recent version of VxRail and VMware software. Directly upgrading from previous VxRail software versions to the newest version is easy, so that customers can quickly take advantage of the latest improvements.

References

Dell Technologies documentation

The following Dell Technologies documentation provides additional information. Access to these documents may depend on your login credentials. If you do not have access to a document, contact your Dell representative.

- [Dell EMC VxRail Release Notes](#)
- [Dell EMC VxRail Administration Guide](#)
- [VxRail Documentation Quick Reference List](#)

VMware documentation

For additional information, see the following VMware documentation:

- [Space Efficiency Using the New “Compression only” Option in vSAN 7 U1](#)
- [What’s New in vSAN 7 Update 1](#)
- [VMware vSAN 7.0 Update 1 Release Notes](#)

Appendix: Test team and setup

Test team

Within Dell Technologies Corporate Systems Engineering (CSE), the VxRail Product Technology team is responsible for internal knowledge transfers, creation of technical how-to procedures, documentation in the form of planning and architecture guides, and evaluating and characterizing the performance of VxRail.

These product technologists maintain their focus on areas of expertise. The contributors to this paper are responsible exclusively for performance modeling and characterization. They are first to evaluate and characterize new technologies, such as Intel Optane Persistent Memory and new server platforms, as well as firmware, BIOS, and software code for VMware products, Dell PowerEdge servers, and VxRail HCI System Software.

The technologists use their findings to build the performance models of the VxRail Online Sizing tool. Dell sales teams use the sizing tool to ensure that the workloads that customers intend to run on the VxRail system will meet customers' performance expectations.

Test setup

VxRail configuration

We conducted our tests on the following VxRail configuration:

- Six VxRail P570F systems, each configured with:
 - Dual Intel 6226R processors (2.9 GHz)
 - 768 GB RAM
 - 25 GbE network connectivity
- Two disk groups:
 - Cache disk—Toshiba SAS PX05SMB080Y, 800 GB
 - Capacity disks—Three KIOXIA KPM5XRUG3T84, 3.84 TB
- VxRail HCI System Software 7.0.100, which includes vSphere and vSAN 7.0 U1

To evaluate platform behavior with even faster cache drives, as described in [Conclusion #5: NVMe cache drives provide maximum storage performance](#), we replaced the Intel processors and the Toshiba SAS cache disk with:

- One Intel Optane processor, 375 GB
- Two Samsung NVMe 1725b SSDs, 1,600 GB

VMs and working sets

In each VxRail node, we deployed two VMs. Each VM included:

- Eight VMDK disks
- 200 GB per VMDK

With this setup, we obtained the following working sets:

- **Total allocated space**—19,200 GB (1,600 GB per VM x 12 VMs)

- **Total working set size for OLTP and RDBMS**—7,200 GB (600 GB per VM x 12 VMs)
- **Total working set size for sequential workloads and random reads**—14,400 GB (1,200 GB per VM x 12 VMs)
- **Total working set size for random writes**—3,600 GB (300 GB per VM x 12 VMs)

Test methodology

We used a tool that uses vdbench and parameter files to run a suite of 17 tests. These tests were a mix of:

- IOPS compared with latency, emulating typical customer workloads:
 - OLTP—Random 70 percent reads and 30 percent writes
 - RDBMS—Mostly random 40 percent writes, mix of 8 KB and 128 KB; 60 percent reads
- Short duration reads and writes using different block sizes to assess front-end performance under optimal conditions
- Longer steady-state random and sequential writes to assess back-end performance

To prepare for the testing, at the start of each test suite, we used vdbench to fully write to the end of each VMDK with sequential 64 KB writes using random data. The test suite is run with all combinations of the following vSphere options, totaling nine sets of tests:

- **RAID options**—RAID 1 (FTT=1), RAID 5, RAID 6
- **Space efficiency options**—None, compression only, deduplication and compression