Hewlett Packard Enterprise

Technical white paper

Check if the document is available in the language of your choice.

VERY LARGE DATABASE DEPLOYMENTS ON VMWARE USING AN HPE SYNERGY AND 3PAR SOLUTION

Deployment white paper

Authors:
Robert Ladd, HPE Pointnext Services Master Consultant
Lee Humphrey, Presales Enterprise Architect

CONTENTS

DOCUMENT REVISION HISTORY

Document status: Final

Change request # (Optional)	Document version	Date	Prepared/Modified by	Reviewed by	Section and text revised
	V0.1	6/28/2019	Robert Ladd	Lee Humphrey	Template creation
	V0.2	7/10/2019	Lee Humphrey	Robert Ladd	Initial content creation
	V0.3	7/16/2019	Lee Humphrey	Robert Ladd	Added Oracle Backups section and various editing
	V2.1	4/6/2020	Robert Ladd	Robert Ladd	Accepted modification recommendations from peer review by Lee Humphrey
	V2.2	4/6/2020	Robert Ladd	Lee Humphrey	Added flow diagram of 3PAR backup and restore
	V2.3	4/7/2020	Robert Ladd	Lee Humphrey	Release candidate 1
	V2.4	4/8/2020	Robert Ladd	Lee Humphrey	Finial corrections
	V2.5	4/9/2020	Robert Ladd	Lee Humphrey	Finial release

PROPRIETARY NOTICE

SETTING THE STAGE

Many HPE customers have considered running large Oracle databases on VMware®. However, several of these customers have been discouraged from moving forward with this architecture primarily due to Oracle licensing considerations and concerns of the performance of a virtualized database. HPE Synergy compute infrastructure and 3PAR storage as well as using best practice design methodologies for databases in virtualized environments can deliver the performance customers require. This architecture allows HPE customers to meet their performance and availability requirements at a much lower cost. Additionally, customers are negotiating enterprise licensing with Oracle as it becomes a more cost-effective model for large Oracle deployments. HPE has been asked by customers to demonstrate running a **Very Large Database (VLDB)** in a VMware environment with some key requirements that represented the type of workloads they experience.

They provided four success factors:

1. Database transactional throughput provides equal to or better than the current environment
   a. Support 1 million transactions per hour with 1 ms or less disk I/O response times
   b. We estimate this to be about 3,000 8 KB IOPS (70% read, 30% write)
2. Perform traditional Veritas NetBackup backups utilizing HPE StoreOnce for Backups and Restores utilizing Oracle RMAN
3. No loss of functionality of maintenance and management of the Oracle Database
   a. DR off-site
   b. Client backup and restore execution
   c. HPE 3PAR Virtual Copy Backup and Restore execution (New concept)
4. 30–40 TB database size

SCOPE OF THIS DOCUMENT

The scope of this document is limited to a VLDB workload, backup/recovery utilizing standard VMFS datastores and an introduction to V-Vols using VASA containers. The goal of this design is to successfully back up the database without stunning the VM. This will be accomplished with utilizing a customized script to quiesce the database, then tell the 3PAR storage array to snap a consistent copy of the VMFS datastores and finally return the database to normal operation.

The next test will be with V-Vols and utilizing OneView for VMware vCenter® to use VMware base snapshots without stunning the database by using advanced features developed by HPE and VMware.

VLDB defined

VLDBs can fall into a few categories; from high transactional with heavy compute requirements to large storage and I/O requirements. Generally, these types of databases fall somewhere in between. This customer's database falls into a blended form of high transactional with large storage and minimal millisecond I/O response time from the storage subsystem. The reference database was Oracle 12c RDBMS with a SGA of 128 GB and disk storage capacity of 60 TB of data.

vSphere 6.5 and 6.7 capabilities

We reviewed the capabilities of VMware vSphere® 6.5 U2 and vSphere 6.7 U2 and based on the information, we do not believe that we are pushing any defined vSphere limits, nor do we believe that the performance requirements are beyond the capabilities of either version of vSphere. Below is a VMware provided chart of the options of interest for this white paper. More information can be gathered from VMware.

vSphere VM Maximums

TABLE 1. vSphere Virtual Machine Maximums

Category	6.5 U2 Limits	6.7. U3 Limits	Description
Compute (Virtual CPUs per VM)	128	256	Max. limit for cores per socket is 64 (128 is based on 2 sockets with 64 cores each)
Memory (RAM/SWAP)	6128/6128 GB	6128/6128 GB	Utilizing VMFS5 Datastore File formatting
Storage Virtual Adapters	4	4	Utilizing Paravirtualized Adapter type for best performance
Storage Virtual Devices per adapter	15	64	6.7 is with the use of Paravirtualized SCSI Adapters PVSCSI
Storage Device Total per VM	60	256	6.7 is with the use of Paravirtualized SCSI Adapters PVSCSI
Virtual Disk Size Per Storage Device	62 TB	62 TB	This is the maximum size of each virtual disk presented to Storage Adapter

Proof-of-concept parameters

This customer wanted as near a real-world lab test as possible. To accommodate this request, additional requirements were added to have multiple VMs defined in the environment running other workloads and for the lab to mirror their on-site environment (compute, networking, storage, and backups). Additionally, HPE wanted to prove the following proposed functionality:

- VMFS5, V-Vols, vSphere Virtual Distributed Switches.
- LACP/vPC with Synergy Link Aggregation Groups for improved bandwidth and availability.
- Utilize 3PAR and Virtual Copy Capabilities for providing fast backups and recoveries of large and/or mission-critical databases.
- Best virtualization practices for storage deployment of database virtual machine on vSphere.

POC test scenarios

- Phase 1: Deploy the methodology at HPE's lab (Networking, 3PAR Storage Snapshots, OneView for vCenter, VMFS5 Datastores and V-Vols on VASA).
- Phase 2: Execute the methodology demonstrate and document the process to the customers.
- Phase 3: Utilizing scripts backup the VLDB with 3PAR read-only and updatable snapshots and perform a recovery of the database to a DR database. This will show the power of the 3PAR Storage and the ability to recover quickly from a database outage or bring up a DR environment. We can also utilize this methodology to perform an initial refresh of a lower environments.
- Phase 4: On-site implementation of the methodology using the customer's physical data.

ENVIRONMENT OVERVIEW

Compute infrastructure

HPE Synergy (3) frame logical enclosure with 30 Synergy 480 modules and 3 Synergy 660 modules. The Synergy 480 modules are dual socket (Intel® x86 Gen10) with up to (24) DIMM memory modules. We utilized two of these Synergy 480 compute modules (blocked in green) in our testing.

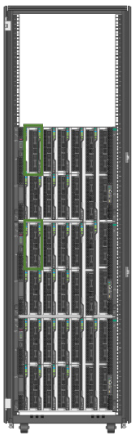

FIGURE 1. HPE Synergy 3 Frame Logical Enclosure example

Storage infrastructure

- HPE 3PAR 8450 all-flash storage array
- HPE StoreOnce 3640 backup appliance

Virtual infrastructure

- vSphere 6.5 VMware ESXi™ Cluster with vCenter 6.5. Later upgraded to vSphere 6.7 with vCenter 6.7.

Backup infrastructure

- HPE ProLiant DL380 for Veritas NetBackup media servers
- HPE StoreOnce 6600 backup appliance
- HPE 3PAR Virtual Copy (Snapshots)

Software infrastructure

- Red Hat® Enterprise Linux® 7.6
- Oracle 12.2 RDBMS Enterprise Edition
- Oracle 18c Grid Infrastructure
- HammerDB Benchmarking tool for the (TPC-H and TPC-C) Database workloads
- HPE OneView for vCenter

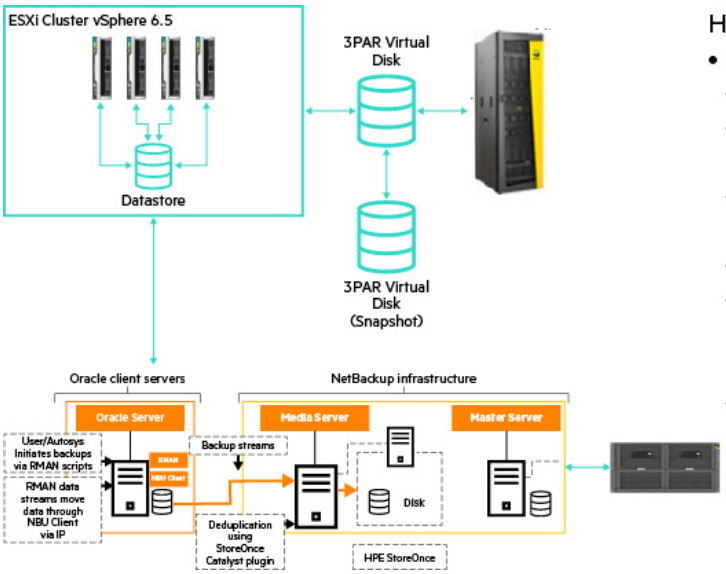

High-level overview of the off-premises test environment

- Environment
  - 3 Synergy Frame Logical Enclosure
  - (4) Synergy 480, 36 CPUs, 384 GB RAM, 24 GB/s Network, 16 GB/s FCoE
  - (1) 3PAR 8K All Flash Array; StoreServ, StoreOnce 3640 Gen10 Appliance
  - (4) VM 16-core 300 GB of RAM
  - Veritas NetBackup solution deployed that matches CNP deployment
    - DL380 utilized for NetBackup solution
  - 60 terabyte database was built with 40 terabyte of user data with 20 TB peripheral data as indexes, relationships, system data, and etc. making up the rest of the 60 TB

FIGURE 2. HPE VLDB off-premises test environment

ENVIRONMENT DETAILS

HPE Synergy

Optimize applications and service levels

The HPE Synergy 12000 Frame is architected as composable infrastructure (CI), which matches the flexible use of compute, storage, and network/fabric resources to an intelligent HPE management software architecture. These capabilities are presented as **infrastructure as code** to the user.

The combination of hardware flexibility with embedded intelligence enables auto-discovery of all available resources for quick deployment and use. Management of hardware by profiles defined in software allows fast repurposing of compute, storage, and fabric resources to meet workload demands.

Operations which use a single frame can scale to large environments of multiple units. Intelligent frame link technology provides easy scaling from a single HPE Synergy 12000 Frame to multiple frames by using HPE Frame Link Modules and HPE Synergy Composer. A total of 21 frames can be managed by two Synergy Composers, powered by HPE OneView. These Composers run in an HA configuration to provide redundancy to the frames (the frames can continue to run workloads without a Composer).

Move faster by accelerating application and service delivery

The HPE Synergy 12000 Frame provides auto-discovery for fast deployments. The flexible design also allows you to custom-fit multiple types of compute, storage, and fabric components in a variety of configurations to meet your workload requirements. Hardware installations are simplified by a simple step-by-step guided process and the intelligent HPE Synergy 12000 Frame design, which reduces the time to set up, debug, and check inventory during installation. A limited-access mode also guides technicians to correct installation errors. HPE OneView supports role-based access control and scopes to create fine grain access based on company employee's role in the infrastructure. High infrastructure bandwidth (over 16 Tb/s) helps you to perform well on today's data-intensive applications and on next-generation, mobile-first, cloud-native applications. The design also utilizes a true passive mid-plane to assure highly reliable operation. And it's photonics ready!

Work more efficiently by reducing operational effort and cost

The HPE Synergy 12000 Frame monitors status for operations by automatically reporting resources, installation information, and health status of the system or domain. Status is updated by the HPE Synergy Composer management appliance (which is powered by HPE OneView) via HPE Frame Link Modules. Updates, from a tested release set of component software, are securely handled using a dedicated data network (separate from the management network) with redundant path configurations. HPE Synergy out-of-band management provides a dedicated 10 Gbps bandwidth for additional security and control. The HPE Synergy 12000 Frame will fit in data centers with AC or DC power operations. The 10U form factor of the HPE Synergy Frame also fits in most standard racks and connects to industry standard power feeds for quick and easy installation.

Unlock value by increasing productivity and control

The HPE Synergy 12000 Frame provides fully air-gapped management (control) and data (production) networks for improved data efficiency and for management security. Separation of data and management traffic helps prevent denial-of-service attacks and increases security control. A single management domain can control up to 21 HPE Synergy Frames and 252 Synergy Compute Modules, and still have it presented as one. This can result in significantly fewer management uplink connections in larger environments. Every HPE Synergy 12000 Frame includes an automated cooling system with capacity to cool both today's components and tomorrow's hotter versions. The HPE Synergy design has cooling capacity with headroom to not limit component operation and to keep your workloads operating at peak utilization.

HPE Synergy 480 Gen10 Compute Module

Gain operational efficiency and control and deploy IT resources quickly for any workload through a single interface. HPE Synergy is a powerful software-defined solution. HPE Synergy Composable Compute resources create pools of flexible compute capacity that can be configured almost instantly to rapidly provision infrastructure for a broad range of applications. The HPE Synergy 480 Gen10 Compute Module delivers an efficient and flexible two-socket workhorse to support most demanding workloads. Powered by Intel® Xeon® Scalable Family of processors, up to 4.5 TB DDR4, more storage capacity and controllers and a variety of GPU options within a composable architecture. HPE Synergy 480 Gen10 Compute Module is the ideal platform for general-purpose enterprise workload performance now and in the future.

HPE Synergy 660 Gen10 Compute Module

Delivering a truly agile and efficient IT infrastructure requires the transformation of traditionally rigid physical systems into flexible physical or virtual resource pools. HPE Synergy Composable Compute resources create a pool of flexible compute capacity that can be configured almost instantly to rapidly provision infrastructure for a broad range of applications. The HPE Synergy 660 Gen10 Compute Module delivers higher performance and scalability for your demanding, enterprise data-intensive workloads. The powerful Intel Xeon Scalable Family of processors (2 or 4 sockets), broad memory footprint (up to 6 TB of HPE DDR4 memory in 48 DIMM slots) and flexible I/O configuration give your large workloads, such as structured databases and business processing applications, the compute power and memory support they demand.

HPE 3PAR StoreServ

Built to meet the extreme requirements of massively consolidated cloud service providers, HPE 3PAR StoreServ provides more than 3M IOPS and consistent sub-ms latency. Transform your midrange and enterprise deployments with solutions that scale from a few TB to more than 20 PB.

The HPE 3PAR StoreServ family of flash-optimized data storage systems modernizes your data center to handle unpredictable workloads effortlessly, and HPE has the only storage smart enough to guarantee 99.9999% data availability.[1] Get rapid and automated provisioning, multitenant design, hardware-accelerated deduplication and compression, and sub 1 ms latency—all in a tier-1 storage architecture that starts small and scales big. The following link will take you to the HPE 3PAR storage webpage, which will provide additional detailed information: HPE 3PAR StoreServ family and HPE Primera Storage family.

HPE StoreOnce

Next-generation HPE StoreOnce: The new, next-generation HPE StoreOnce system provides a built-for-cloud data protection platform that can scale from small remote offices to the largest enterprises and service providers. Next-generation HPE StoreOnce can help you reduce cost, risk, and complexity with flash-speed protection in your data center, as well as deliver low-cost archive and disaster recovery in the cloud.

Help optimize cloud investment

HPE StoreOnce with HPE Cloud Bank Storage delivers simple, efficient, and reliable backup to the cloud. It lets you seamlessly cloud-enable your storage, backup, and enterprise apps, as well as natively integrate your choice of cloud service provider without the need for a separate gateway or virtual appliance. You can reduce the time, cost, and network bandwidth for your cloud storage significantly with highly efficient deduplicated data transfer as well as store encrypted, self-describing backup data for simple cloud disaster recovery.

Take control of data growth: A scalable portfolio of physical and software-defined HPE StoreOnce backup appliances provides you with a wide range of pay-as-you-grow capacity points system. HPE StoreOnce makes every gigabyte of your storage go further with intelligent deduplication that can reduces your backup costs and footprint and in the cloud.

[1] Based on HPE analysis of public information for portfolios spanning IDC midrange and high-end price bands

Deliver on service-level agreements (SLAs): HPE StoreOnce helps you break the backup window with faster performance. Direct backup from your HPE 3PAR or HPE Nimble Storage array to HPE StoreOnce with HPE Recovery Manager Central (RMC) delivers faster backup and recovery for your enterprise apps. This backup comes with less cost and complexity, as well as having minimal impact on your production environment.

Get comprehensive protection against any threat

HPE StoreOnce gives you the confidence that your data is recoverable with a data protection platform that is reliable by design. Mitigate threats with 3-2-1 data protection[2] best practice, encrypted backups, data integrity checks, role-based access control, fast restores, and high availability storage of last resort. Make backups invisible to ransomware attacks with the highly reliable HPE StoreOnce Catalyst protocol.

Reduce complexity

HPE StoreOnce takes the time and stress out of data protection, reducing management time with multisystem and multisite visibility and management—from a single pane of glass. It empowers your application owners with self-service protection, enabling them to easily and efficiently control end-to-end protection for their applications—from the data center to the cloud—using familiar native user interfaces.

Experience agile software-defined protection for virtual and cloud environments

HPE StoreOnce VSA extends the deployment options for StoreOnce with the agility and flexibility of a virtual appliance, removing the need to install dedicated data protection hardware. All the features of the purpose-built HPE StoreOnce systems are available in a software-defined backup target of up to 500 TB usable capacity, which can be configured in 1 TB increments. An HPE StoreOnce VSA license server simplifies license management for large or dynamic VSA deployments.

HPE OneView for VMware vCenter

HPE OneView for VMware vCenter is a single, integrated plug-in application for VMware vCenter management, which enables the vSphere administrator to quickly obtain context-aware information about HPE Servers and HPE Storage in their VMware vSphere data center directly from within vCenter. This application enables the vSphere administrator to easily manage physical servers and storage, datastores, and virtual machines. By providing the ability to clearly view and directly manage the HPE Infrastructure from within the vCenter console, the productivity of VMware administrator increases, as does the ability to ensure quality of service. HPE OneView for VMware vCenter offers the following benefits:

- Simplified administration through integration of the physical and virtual infrastructure.
- Accurate problem indicators through the hardware events generated in the VMware vSphere Management Console.
- Single-click launch of trusted HPE management tools from the vSphere dashboard.
- Ability to proactively manage or view changes with detailed relationship dashboards of server, networking, and storage.
- Simplified on-demand server and storage provisioning.
- Visualization of complex configuration relationships:
  - Virtual machine mashup with storage
  - Peer Persistence volumes mashup
  - Virtual Connect end-to-end networking view

vSphere environment

The vSphere environment was composed of six (6) Synergy 480 compute modules with 32 cores and 384 GB of RAM each. In addition, the HPE Synergy 480 had a multifunction adapter with 12 Gb/s Network and an 8 Gb/s Fibre Channel. A separate vSphere environment housed the vCenter for this cluster as well as VMs for the Veritas NetBackup environment.

Virtual Machines were created for:

- 5 Oracle Databases Virtual Machines
  - 1 CPU socket 16 and 8 cores with
  - 128 GB of RAM and
  - 80 TB of VMFS5 storage with
  - 4 paravirtual storage adapters
  - 4 virtual NIC (Public, Backup, Interconnect_1 [ASM/Private], Interconnect_2 [Private redundant path])

[2] Defined by Carnegie Mellon recommended 3-2-1 method in publication titled Data Backup Options

- 1 HPE OneView for vCenter utilized 1 socket 2 cores with 16 GB of RAM
- 1 vSphere vCenter utilized 1 socket 4 cores and 16 GB of RAM

vSphere deployed configuration

The ESXi hosts are connected to the HPE 3PAR storage array. We followed the Oracle database and VMware vendor's best practices for datastore design. This practice is recommended for all database vendor types. Typically, customers assign datastores to hold multiple virtual machines. This works well for low transactional databases and for general applications. However, in a high transactional workload like a VLDB, this type of deployment provides less than desirable results, and performance suffers and then the customer feels it's a virtualization issue with the database, when in fact it is a design issue of the virtual environment in relationship to the database architecture characteristics.

Design considerations for databases

When deploying databases in a virtualized environment, it's important to be mindful of resource allocation and storage design and do not over allocated CPU and Memory. Over allocating CPU and Memory resources can do more harm than good for the database. The reason is that the VMware must unnecessarily manage more CPU and Memory schedulers when the system does not utilize all the allocated resources. The key point to make here is that you want to deploy a virtualized database just as you would a physical database. There is a bit of trial and error on allocation of these resources to find the optimal balance. Taking a database and running a few tests and reviewing performance reports (system, database, and VMware) under different workloads will help tune these parameters. Next is the storage layout; this is often the configuration point that is misconfigured.

NOTE

Configure the storage in a virtual environment as you would in a physical environment. What this means is that you wouldn't have different database servers accessing the same storage devices and expect it to perform well and so you should not do this in a virtual environment either.

Virtualized database storage design

When deploying databases in a virtualized environment you should apply the BUDR method. The BUDR method was developed by HPE Consulting based on our years of experience and analyzing what works and what does not work across many corporate environments.

What is BUDR? BUDR stands for Binaries, Utilities, Data, and Recovery. These are the 4 major categories of a database. From here, we can further subdivide these categories to meet the database needs.

- Binaries: This is a standalone 1 TB datastore, it contains the Operating System Binaries and Database Binaries and directory structures for each respectfully. This would be at least two separate VMDKs. One for the OS and one for the DB. You might need more for the database depending on the features you are implementing or multiple database versions, following vendor's documented process.
- Utilities: This is a shared with the binaries datastore, it contains third-party tools and utilities for managing and supporting the virtual database server. Tools like this are Backup, Scheduler/Automation and monitoring tools and utilities.
- Data: This is the **standalone** datastores. Data is broken down into separate datastores based on the category of the data. Data is broken down into subcategories of data. Data is random access in nature and contains the core information that users access.
- Recovery: Databases have a storage segments that stores a sequential data set for recovery and transitional transaction purposes. These transactions are logs, log backups, temporary segments, and in some cases online copies of the database for instant recovery and other key configuration segments.

Implementing BUDR

- Binaries and Utilities are typically a 1 TB thin provisioned datastore to the Operating System, database binaries and utilities.
- Data is sized in blocks (1, 2, 4, or 8 TB) sizes. To determine the block size for data, take the database data segment size and add + 50%. Compare that value to one of the container sizes previously mentioned until it fits within that container size. Once you hit the 8 TB size and you need more space then allocate another 8 TB container until the data segment requirement is met. As an example, a 60 TB database would require (9) 8 TB containers. The reason for 8 TB being an arbitrary maximum size, is through our testing of performance, replication to a secondary array and recovery techniques we found that this size fits quite well in all categories. Consider your storage area network infrastructure's capabilities in terms of bandwidth, capacity, and functionality. This will help determine an appropriate sizing recommendation for your datastore sizing.
  - Oracle, SQL Server, and PostgreSQL Databases have Data, Indexes, and System Tables

- Recovery is sized in blocks of (1, 2, 4, and sometimes 8 TB) sizes. Logs, log backups typically fall into 1 and 2 TB and you typically will have two of these acting as mirror for high availability. Temp Segments and Online Backups will follow in 2, 4, or 8 TB sizes with the backups typically following the same sizing recommendations as data.
  - Oracle databases have Flash Recovery Area, Redo Logs, Archive Logs, and Temp segments
  - SQL Server and PostgreSQL have Transactional Logs and Temp Segments
- When running Oracle on VMware it is strongly recommended to use ASM as the storage mechanism for the Oracle database. It provides better load balancing among the datastores. It also provides easier maintenance for storage backups, storage provisioning, and storage migrations. This document utilized ASM in the testing of Oracle VLDB and the results reflect both standalone ASM and Oracle Real Application Clusters (RAC).

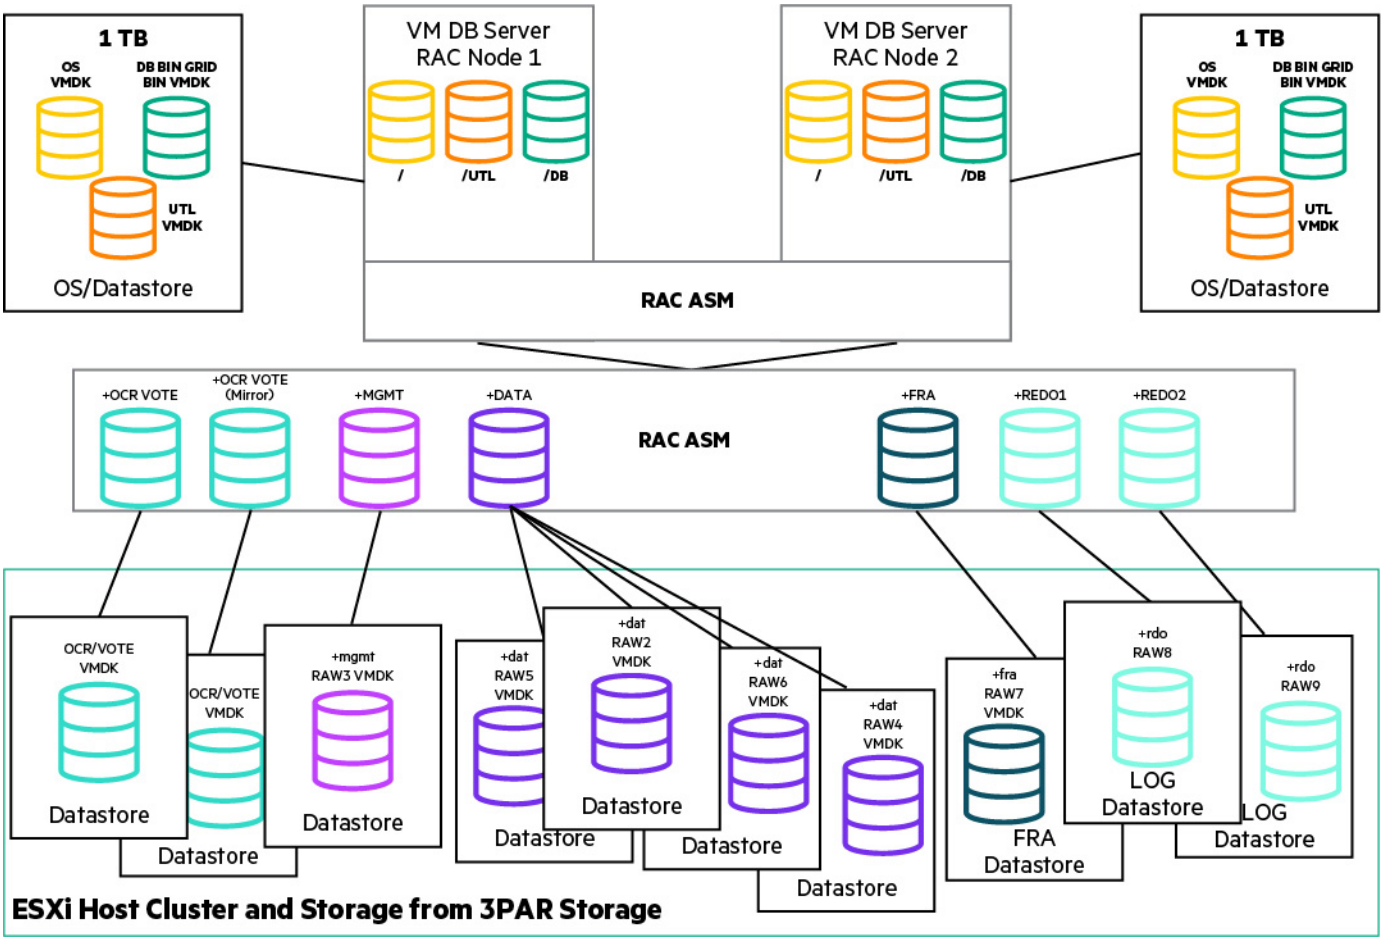

FIGURE 3. Oracle RAC datastore example

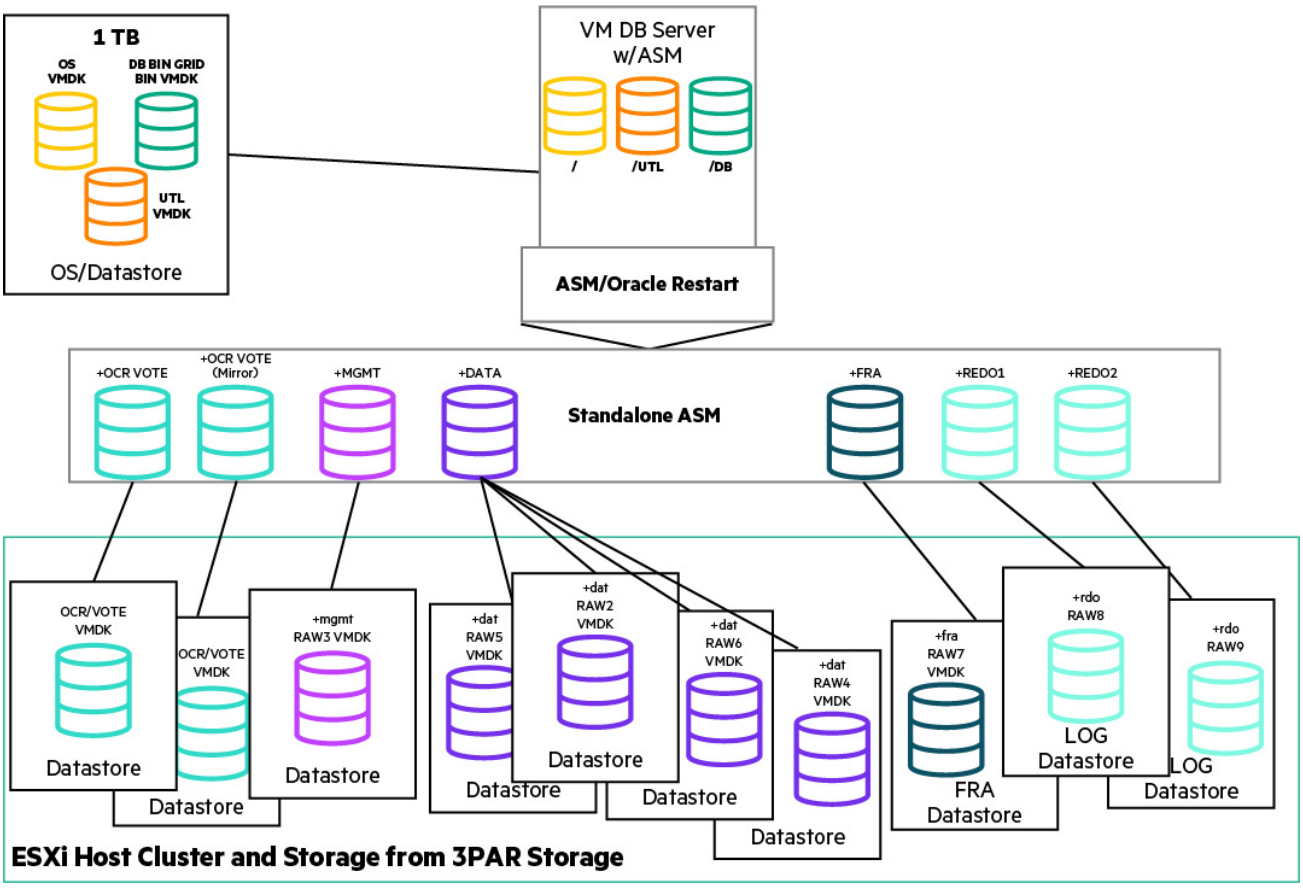

FIGURE 4. Oracle ASM standalone with restart example

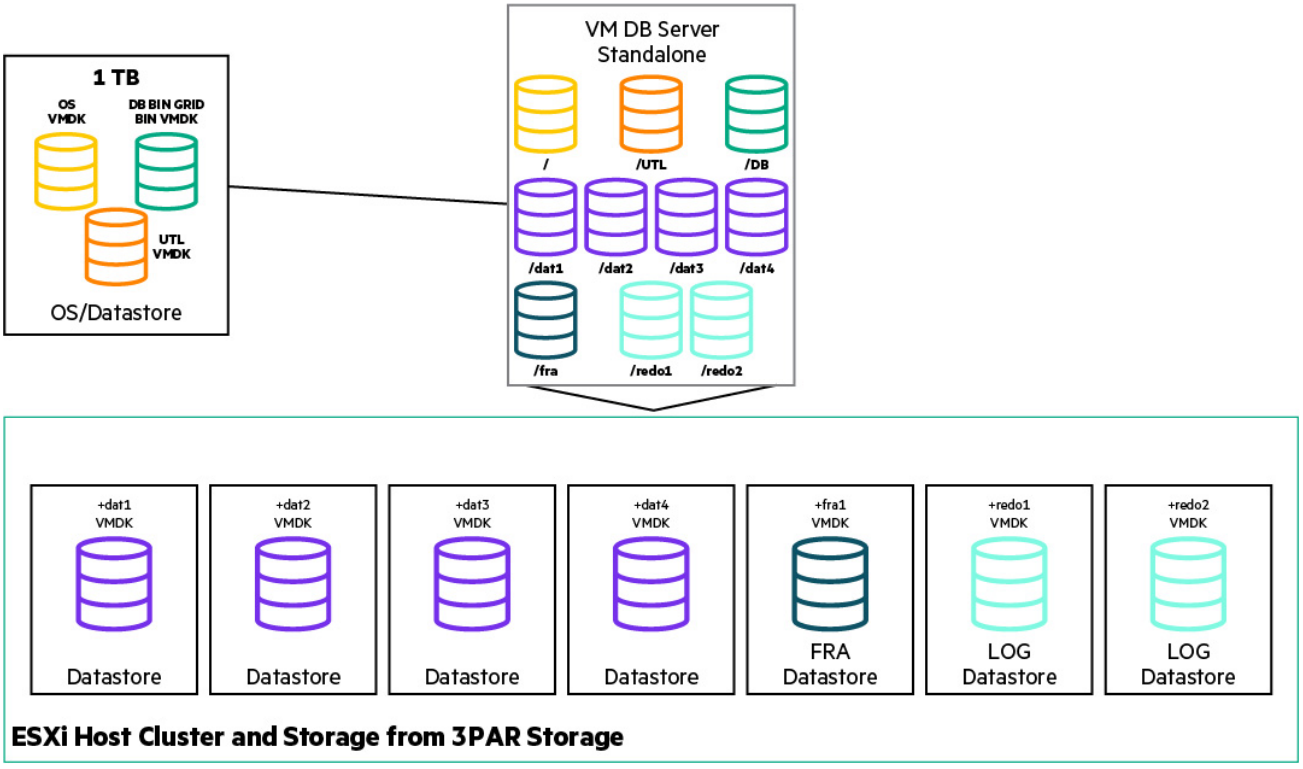

FIGURE 5. Oracle file system layout example

Veritas NetBackup environment overview

The Veritas NetBackup media server is running on a set of HPE ProLiant DL380 servers, which streams its backups over Fibre Channel to an HPE StoreOnce appliance. Data backup is performed by an Oracle RMAN NetBackup agent running on the Oracle Data Guard server. The backup is streamed from the Oracle Data Guard server on Ethernet/IP to the Veritas NetBackup media server, which streams the data over Fibre Channel to the HPE StoreOnce appliance. This phase also included two more mixed workload Oracle databases to provide other activity on the 480 Compute module to exhibit more of a real-world scenario and show that the 480 can run the Oracle VLDB with other concurrent Oracle workloads.

Oracle backups

Oracle backups utilize the RMAN integration with NetBackup and Oracle backup is performed using the Oracle RMAN interface. User or automated process initiates a RMAN backup on the Oracle database. RMAN channels are streamed to the NetBackup media server via IP by a NetBackup Oracle agent running on the Oracle Data Guard server. Users can perform hot, non-disruptive backups being either full or incremental backups of Oracle databases. The HPE catalyst plug-in on the media server performs deduplication and streams the data over Fibre Channel to the HPE StoreOnce appliance.

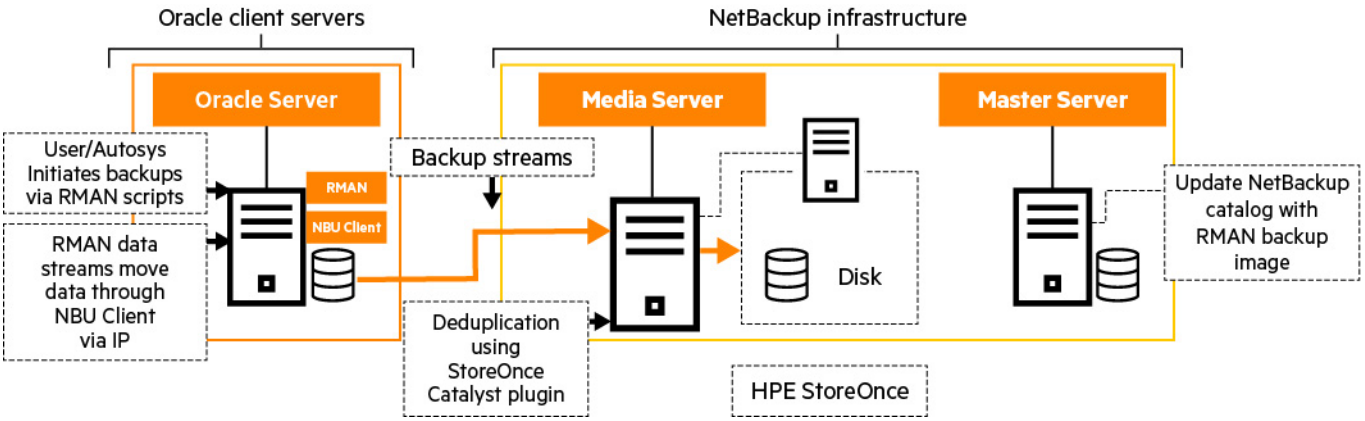

FIGURE 6. HPE StoreOnce integration with Oracle RMAN and Veritas NetBackup

HPE 3PAR Virtual Copy backup and restore

The benefits of the datastore design truly shines when it comes to backups and restoring a database. These backups happen in an instant and can be regulated as to how often and the number of copies that are kept. These backups can be mounted on backup servers and then backed up with the traditional backup vendor software. The 3PAR Storage Array Virtual Copy feature provides quick recovery times compared to traditional restore methods. HPE 3PAR Storage Array Virtual Copy provides a convenient mechanism for lifecycle management. Examples of these are staging, preproduction, and/or lower environments. Note: This does not replace your traditional backup and restore, it should be an enhancement to those backup methods.

HPE 3PAR Virtual Copy

The 3PAR Virtual Copy feature is an often-undervalued tool in a database toolkit that can be used to provide availability, reliability, and recoverability of mission-critical databases. This feature can be used whether the database is a few hundred gigabytes or a multiterabyte monster database. The key is a good storage design architecture that plays into the capabilities of 3PAR Virtual Copy feature. This feature creates what the industry understands as a point-in-time snapshot. This snapshot is thin at creation and grows very slowly over time. The growth is due to when data is changed and the 3PAR manages and monitors these changes. The Virtual Copy can be used as a separate copy of the data in that it is a snapshot, and it can be copied and promoted to a real volume. The Virtual Copy can be used to restore data back to the point in time of the creation of the virtual volume copy.

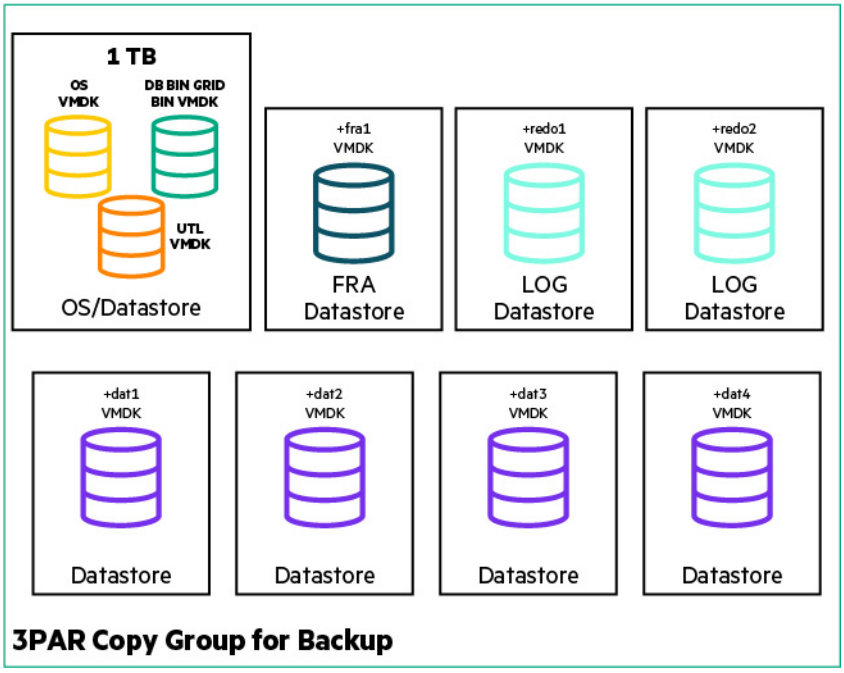

FIGURE 7. HPE 3PAR Virtual Copy using Consistent Snapshots for backup

When performing an HPE 3PAR Virtual Copy with databases, it is imperative that copy groups and consistent snaps are performed. A copy group is a function of the 3PAR Storage Array that allows you to take individual virtual volumes and group them together and perform commands on the group. The consistent option complements the group option and ensures that all the volumes in the group are snapped while providing a consistent time across all the volumes. This is important for databases as they are time sensitive with their volumes and a database will not come up if these volumes have different time values. When we perform the backup, we grab BUDR and this provides a complete point in time of backup. Again, we can take the BUDR and use it for the Disaster Recovery Image (DRI) and we can also use the DRI for refresh of another environment.

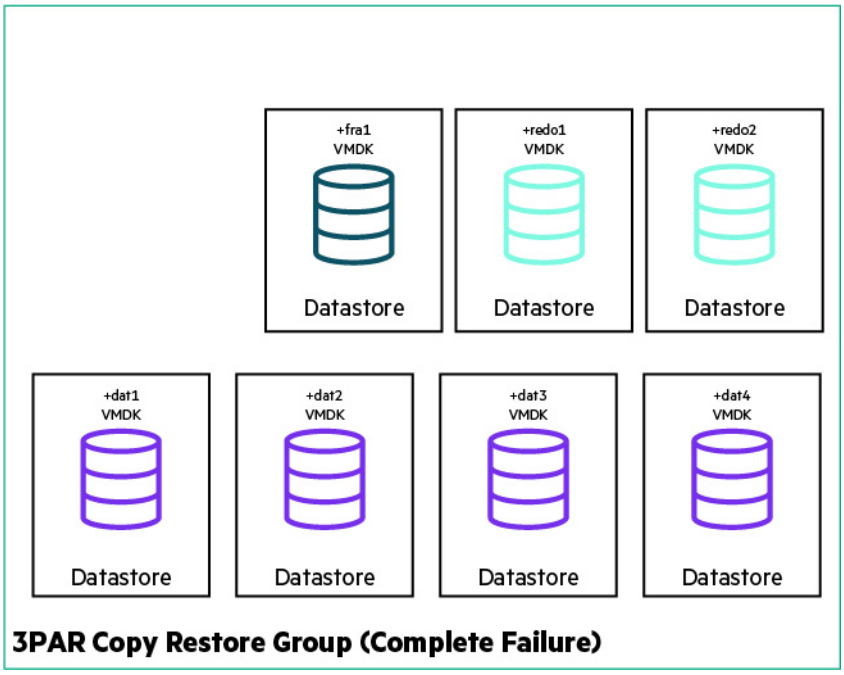

FIGURE 8. HPE 3PAR Virtual Copy promote for restore production database

When we have a complete database failure or corruption, we can use the 3PAR Virtual Copy to restore the database back to the point of the last Virtual Copy and then use traditional backups to roll the database forward.

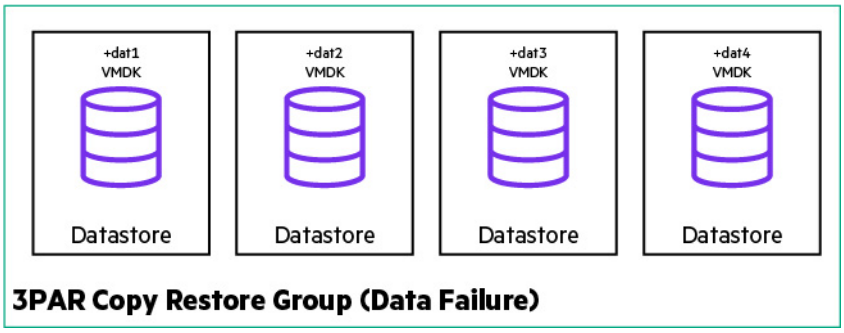

FIGURE 9. HPE 3PAR Data Failure Restore Group

When we have a data failure or corruption then we can simply use Virtual Copy to restore the data segments and not restore the recovery segments. When we do this the database will perform a roll-forward operation and the database will be brought to the point-of-failure time. This may need some actions from the traditional backup of restoring archive logs that had been backed up and purged due to space constraints. This restore process is usually very fast taking only minutes to hours versus hours to days to recover. This is based on the data rate of change between the time of snapshot and the time of failure along with the online archive logs that are available.

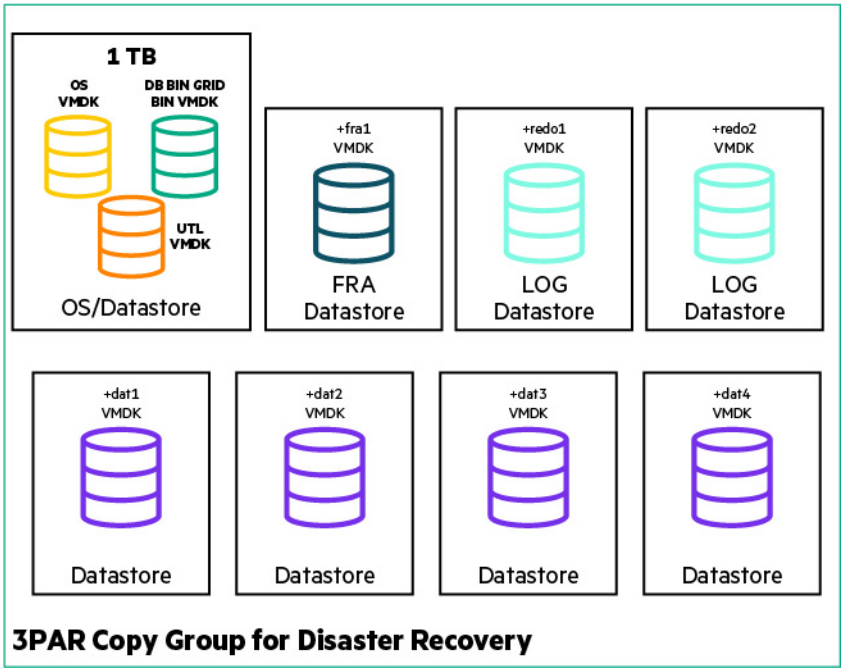

FIGURE 10. HPE 3PAR Disaster Recovery Group

We can take the BUDR image we created during our backup and mount these volumes to a server of similar characteristics and bring the database online and provide business continuity or offline backups of the database with traditional backup utilities. We could also use 3PAR replication and replicate these volumes to an off-site location and provide a remote disaster recovery solution.

As you can see these use cases are only the tip of the iceberg of available options that can be provided by utilizing a solid database storage architecture and the power of HPE 3PAR StoreServ arrays.

VLDB pilot goals

- Primary goals
  - Synergy can successfully run CNP's VLDB without loss of performance or functionality
  - Best practice of backup and recovery (3-2-1 methodology)
  - Best practice of datastore design for databases in a virtual environment
- HPE VLDB methodology
  - Create off-premises test environment
  - Run performance test with multiple workloads
  - Successfully backup and restore the primary database with 3PAR and NetBackup
- HPE VLDB off-premises testing results
  - Test results vSphere 6.5 U2 and 6.7 U3
  - Environment enhancements
- Next steps
  - Run on-premises validation
  - Document technology

VLDB support matrix

- Oracle RAC on VMware support
  - Oracle RAC in Virtual Machines
  - As per My Oracle Support document ID 249212.1, RAC 11.2.0.2 and later is supported by Oracle on VMware.
  - High Availability of Oracle Databases: VMware, Inc. vmware.com/content/dam/digitalmarketing/vmware/en/pdf/solutions/oracle/oracle-databases-on-vmware-high-availability-guidelines.pdf
- Source Database Version and Patch Level
  - OCW Patch Set Update: 11.2.0.4.190115
  - AIX
- Target Database Version and Patch Level
  - OCW Patch Set Update: 11.2.0.4.190115
  - Oracle Enterprise Linux 7 or RHEL 7
- 2-Node RAC Cluster with Data Guard

VLDB performance and functionality

- VLDB Defined
  - 8 TB or larger
- Performance
  - 1 million transaction per hour requirement
- Functionality
  - Database capabilities
    - Oracle Data Guard
    - Oracle Real Application Clusters
  - Database backup recovery
    - Veritas NetBackup
  - Database maintenance
  - Disaster recovery capabilities

HPE VLDB methodology—testing process phase 1 (vSphere 6.5)

- High-level overview of the off-premises testing
  - Performed performance tests on 3 databases
    - Utilized open source product called HammerDB to create the databases and perform user load on each database.

      35–47 users performing 1 million transaction each on each database
  - Performed HPE 3PAR StoreServ backups with/without load and restores
    - Backups—Virtual Copy Snap Volumes
    - Recoveries

      Point-of-backup time

      Point-of-failure time
  - Performed NetBackup RMAN integrated backups with/without load and restores
    - Backups—Standard RMAN backup with NetBackup integration
    - Recoveries

      Point-of-failure time

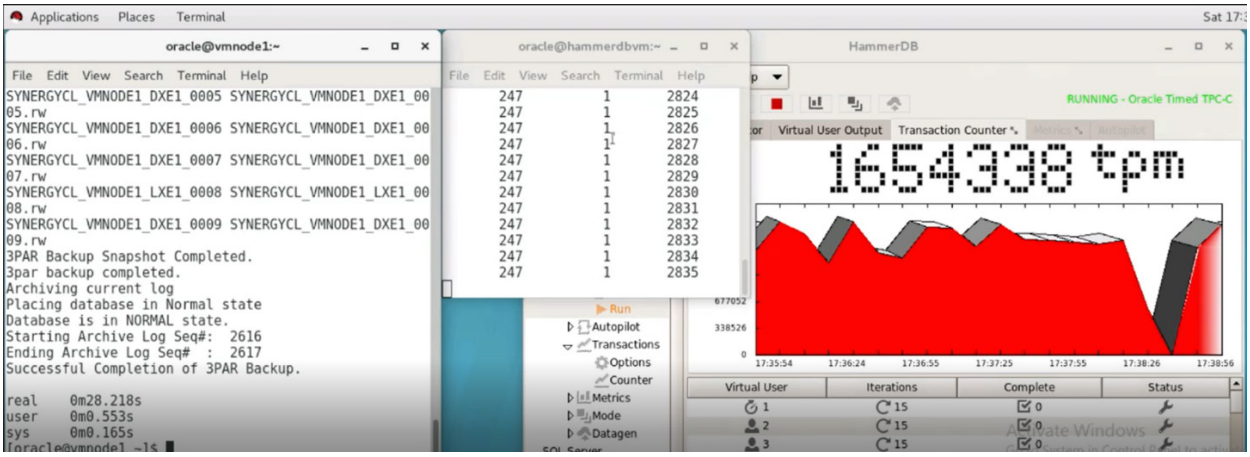

FIGURE 11. HPE VLDB Standalone Database Performance Load Testing results

- Synergy can successfully run VLDB Non-RAC w/Data Guard without loss of performance or functionality.
- 47 users running a million transaction each with zero think time.
- 1.65 million transactions per minute, with two additional databases with the same number of users running workload on same Synergy compute module 480.
- Notice the peak and valleys of the transaction graph.
- During this test we performed an HPE 3PAR Virtual Copy (Snapshot). You can see this is where the graph goes to a sharp drop. On average, this is about a 2–3 second drop in writes and then rapidly returns to full speed.
- What cannot be seen by this graphic, in the middle window is a report query running against the database. The query never stopped or hesitated during the 3PAR Virtual Copy.
- The far-left window is the 3PAR backup scripted process. This is the time to fully snap the BUDR environment. The fastest time that was recorded was 19 seconds and the worse time was 29 seconds. It is important to note that this is a 60-terabyte database plus the binary volume while running a 1.65 TPM load.

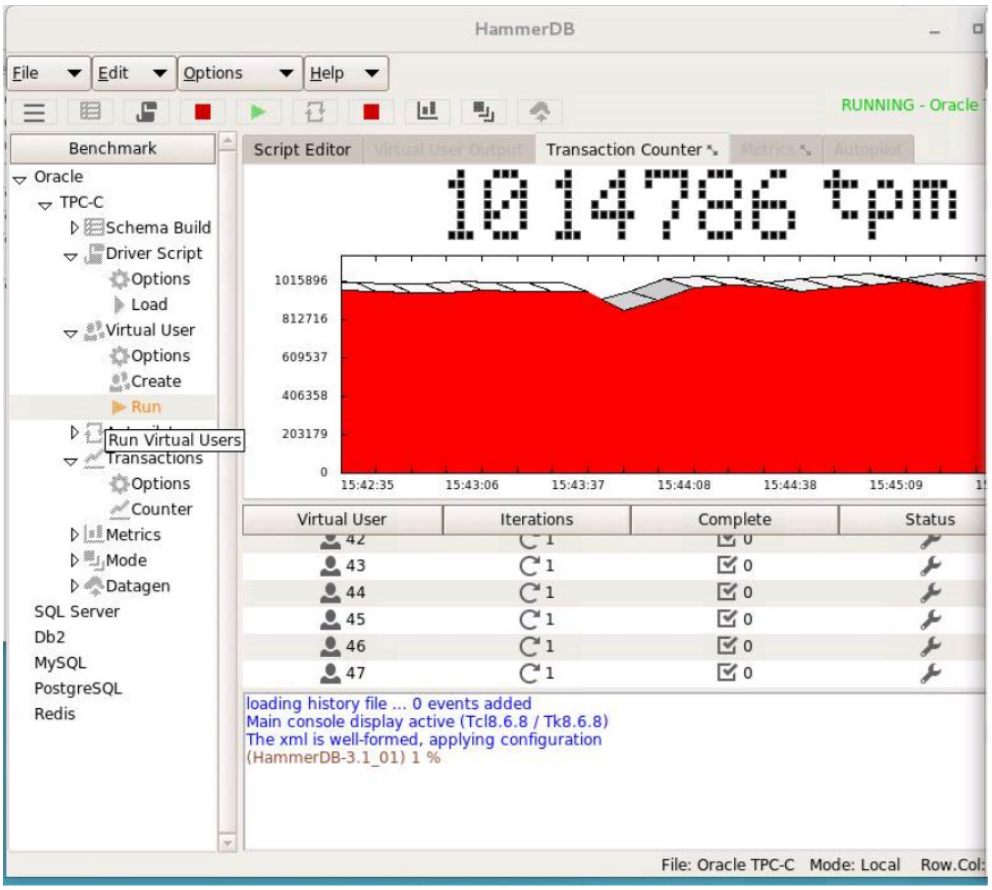

FIGURE 12. HPE VLDB Single-Node RAC Performance Load Testing results (vSphere 6.5)

- Synergy can successfully run VLDB RAC w/Data Guard with no other load on vSphere 6.5.
- 47 users running a million transactions each with zero think time.
- 1.01 million transactions per minute, while two additional databases with the same number of users running workloads on the same Synergy 480 compute module.
- Node 2 of the RAC was shut down for this initial test. This was to show what overhead, if any, presented any positive or negative effects.
- In this case the TPM rate went down with the RAC components.
- Notice the flatness of the transactions compared to the non-RAC database.

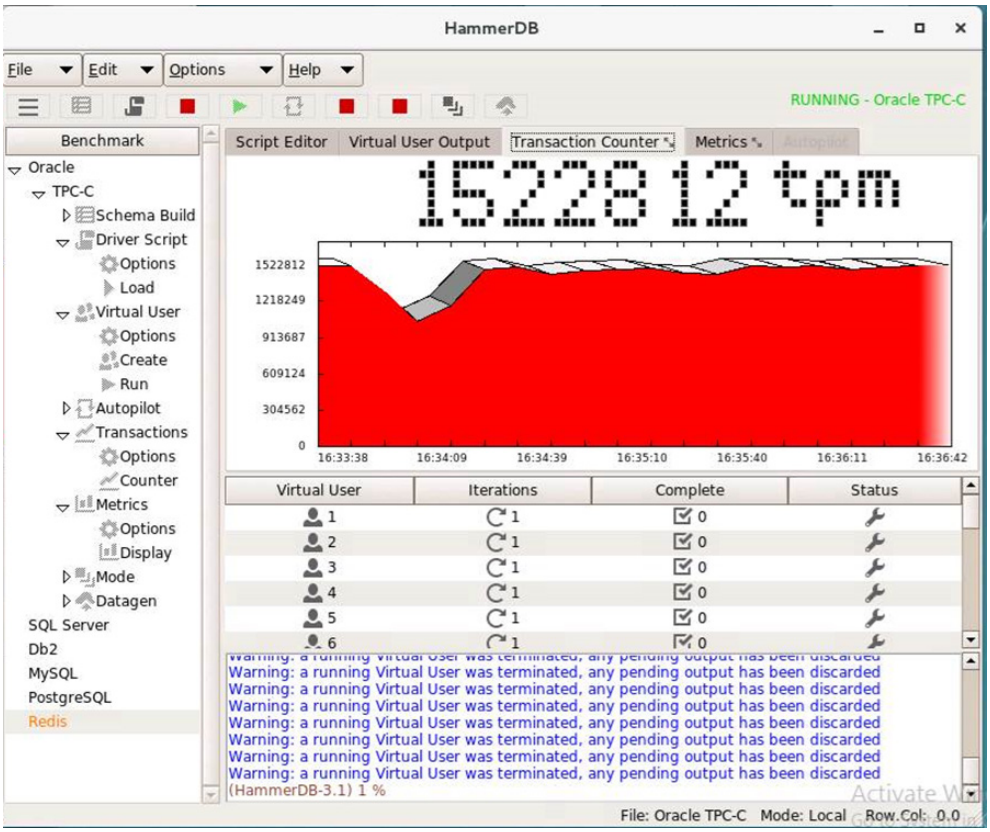

FIGURE 13. HPE VLDB Dual-Node RAC Performance Load Testing results (vSphere 6.5)

- Synergy can successfully run VLDB RAC w/Data Guard with vSphere 6.5.
- 47 users performing 1 million transactions each with no think time.
- Notice the increase of the TPM to 1.52 million transactions per minute, which provided a 50% boost in performance.
- Again, notice the relatively flat transactional graph.
- Oracle and Linux Performance Monitoring Statistics were reviewed. Upon review the HammerDB could not push anymore data. Nowhere did the performance reports show that there were waits for compute resources and that the queries could be tuned. However, wait until the next slide and we move to vSphere 6.7.

HPE VLDB methodology—testing process phase 2 (vSphere 6.7)

- **Upgraded from 6.5 update 2 to 6.7 update 3**
- High-level overview of the off-premises testing.
  - Performed performance tests on 3 databases.
  - **55–102 users performing** 1 million transaction per minute each on each database.
  - Room for several additional users.
  - The upgrade afforded more users and more throughput on both signal-node and 2-node RAC clusters.
- Figures 13 and 14 are running performance tests and monitoring the main database.
  - 1 and 2 RAC Cluster with Data Guard Databases are on separate ESXi Hosts.

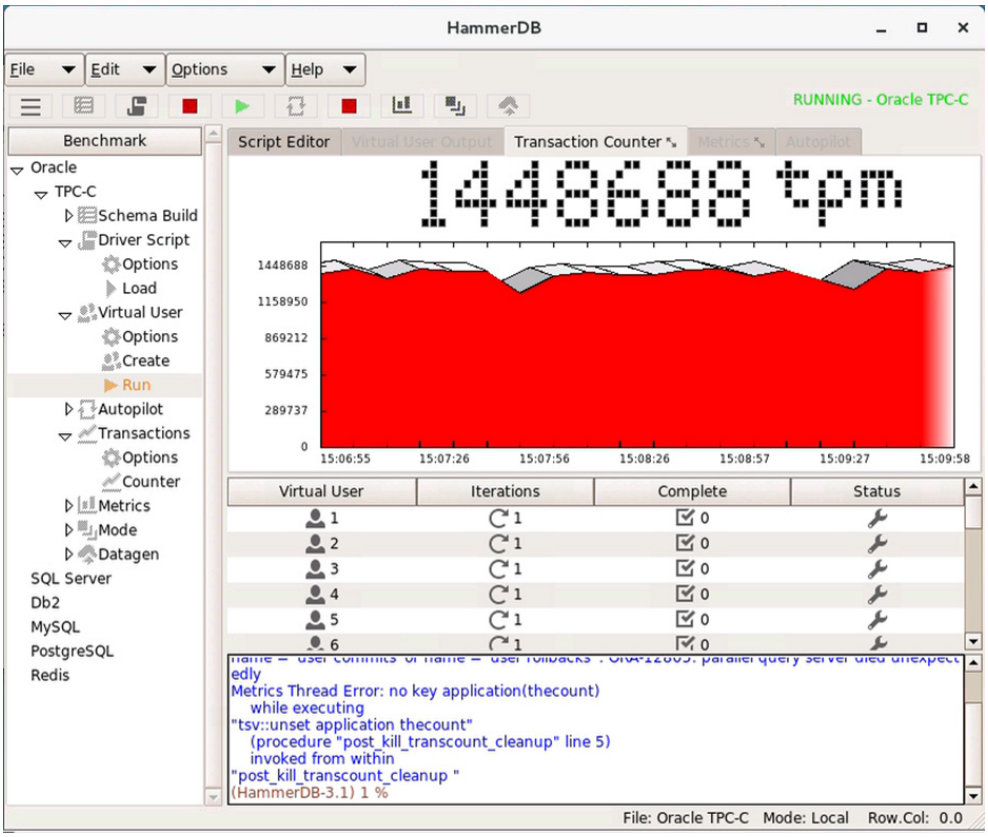

FIGURE 14. HPE VLDB Single-Node RAC Performance Load Testing results (vSphere 6.7)

- Synergy successfully ran a VLDB single node w/Data Guard and with vSphere 6.7 (55 users).
- Throughput increased nearly 50% for a single RAC node with only changing vSphere from 6.5 to 6.7.
- User base was increased nearly 20% from 47 to 55 users running heavy loads.
- Some additional peak and valleys in the throughput graph.

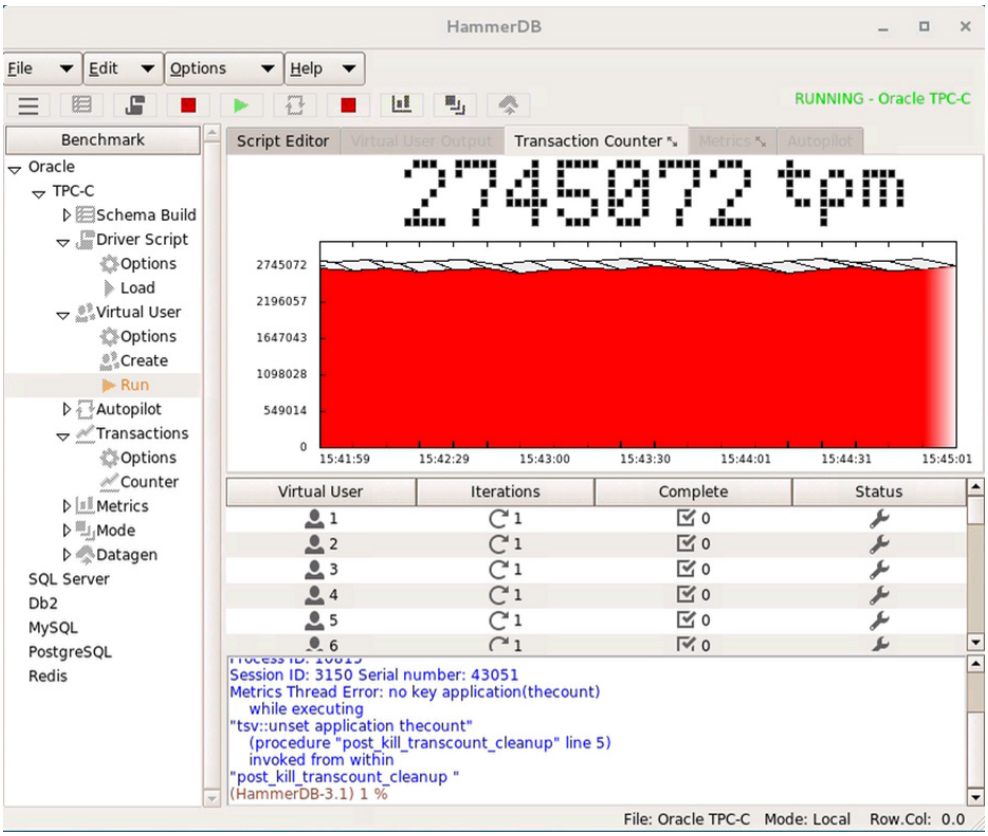

FIGURE 15. HPE VLDB Dual-Node RAC Performance Load Testing results (vSphere 6.7)

- Synergy can successfully run VLDB RAC w/Data Guard with vSphere 6.7 (102 users).
- We increased the user base approximately 85% from 55 to 102 users on vSphere 6.7 and 117% with 47 to 102 users on vSphere 6.5.
- The performance throughput of 80% increased from 1.52 million to 2.75 million TPM.
- As you can see there are some peaks and valleys and it maintained this run rate for the length of the one-hour test.

HPE VLDB Backup and Restore Testing

Perform HPE 3PAR backup of the VLDB under processing

Utilizing a few VLDB databases running on one Synergy module running mixed workloads. We performed the following tasks:

- Issued a backup of the VLDB database.
- Run a query against the database and perform two 3PAR snapshots.
  - Observe if the database becomes stunned and or if it crashes.

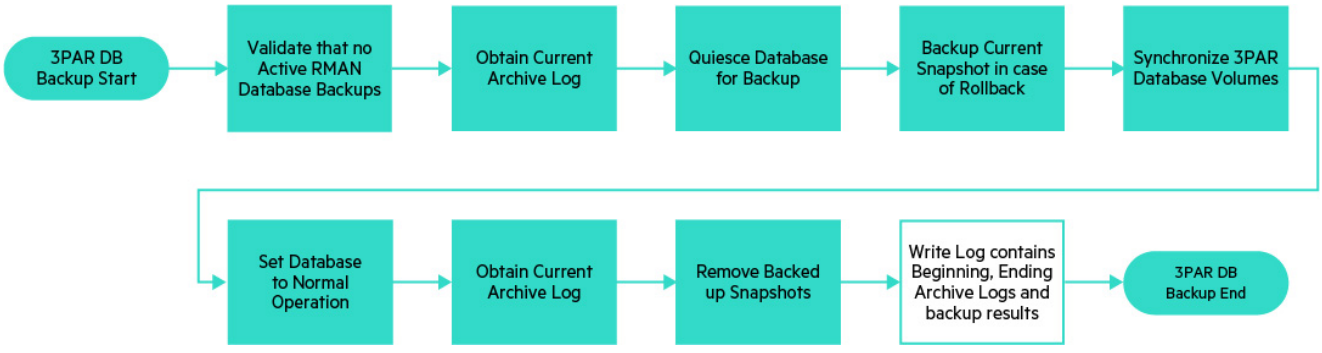

FIGURE 16. HPE 3PAR database backup process

Perform HPE 3PAR point-in-time restore of the VLDB

Utilizing 3PAR perform a complete restore of the VLDB at the time of backup. Observe timings of the two scenarios and identify any specific steps needed for the restore.

1. Utilize the snapshot as a quick recovery for a simulated disaster recovery scenario. Where you quickly need a read-writable copy of the database brought up for specific reason like a DR or specific table recovery.
2. Full restore of the VLDB from the snapshot. This would be an update of a snapshot to physical virtual volume.

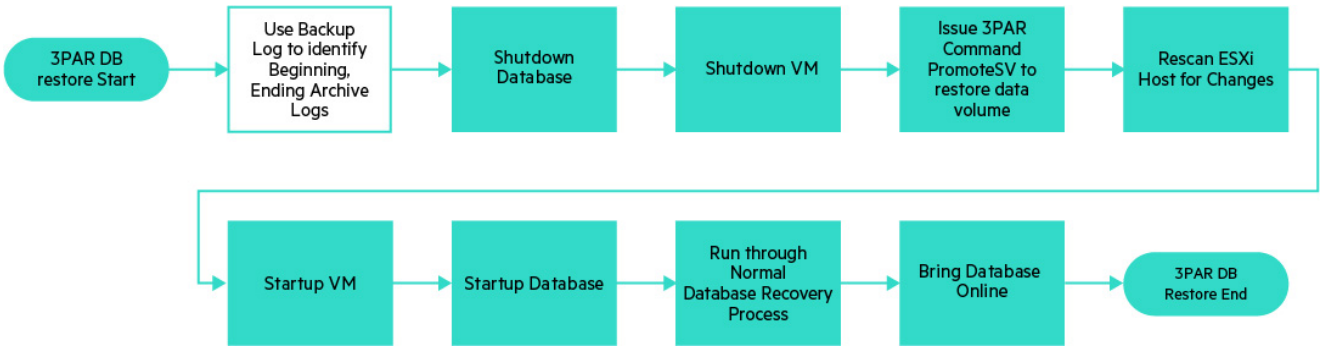

FIGURE 17. HPE 3PAR database restore point-in-time process

Perform HPE 3PAR point-in-time of failure restore of the VLDB

Utilizing 3PAR perform a complete restore of the VLDB at time of backup. Observe timings of the scenario and identify any specific steps needed for the restore.

1. Utilizing database vendor prescription for PTFR and 3PAR snapshots, recover to the database to the point of failure.

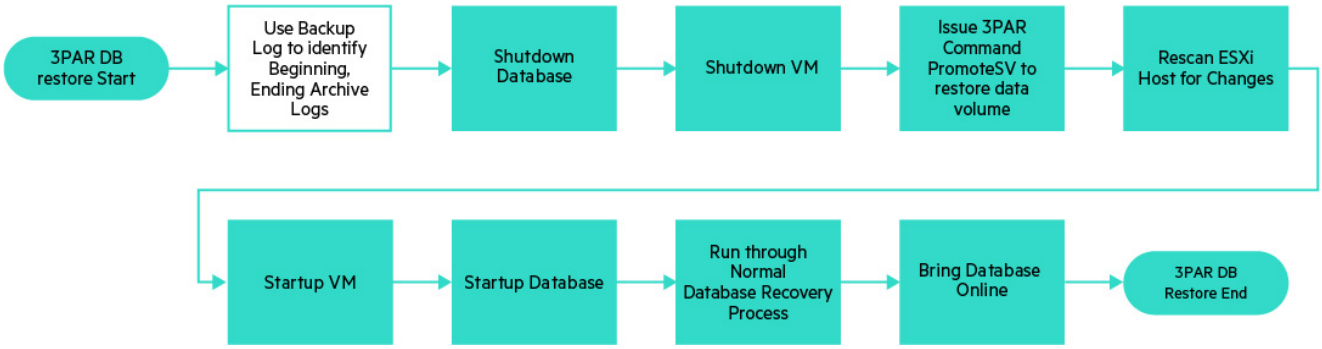

FIGURE 18. HPE 3PAR database restore point-of-failure process

Perform NetBackup point-in-time restore of the VLDB

Utilizing NetBackup perform a complete restore of the VLDB at time of backup. Observe timings of scenario and identify any specific steps needed for the restore.

1. Utilizing NetBackup prescription for PTR, recover to the database to the point of backup and record timings.

TABLE 2. Database Backup and Restore Testing Results

Database Backup/Restore Type	DB Run Status	Elapse Time
Veritas NetBackup and Oracle RMAN Backup	DB Running and processing transactions	13 hr and 50 min
Veritas NetBackup and Oracle RMAN Restore	DB Shutdown	13 hr and 11 min
HPE 3PAR Virtual Copy—Quiesce DB backup	DB Running and processing transactions	29 seconds
HPE 3PAR Virtual Copy—Restore	DB Shutdown	6 hr and 51 min

- HPE lab timings using out-of-the-box settings with no tuning. NetBackup Restores could be improved with tuning of restore threads.
- 60 TB database with 40 TB of raw data.
- 3PAR snapshot was taken at time of backup and in place when we performed the NetBackup Restore. When the restore ran, 3PAR recorded the changes as transactions occurred. This was to get a worst-case scenario.
- Typically, the size of the Virtual Copy will be much smaller and therefore run quicker.
- HPE 3PAR Restore was performed after a transaction performance run and this restore took about 15–30 minutes depending on the rate of change using the 3PAR.

HPE and VMware V-Vols, VASA, and OneView for vCenter architecture

Overview of HPE OneView for vCenter

- HPE OneView plug-in for VMware vCenter consist of several modules to manage servers, networks, or storage through a single pane of glass of both physical and virtual environments.
- It's an appliance that HPE provides via download from our secure site.
- Today we will focus OV4VC's roll in the use of VASA and V-Vols.
  - vSphere APIs for Storage Awareness (VASA) is a set of APIs allowing storage vendors to deliver storage related data into vCenter.
  - V-Vols are virtual volumes that are created by vCenter using the VASA protocol and instructs the storage array to present unit of storage to a VM.
  - OV4VC is a broker in this process that translates the vCenter commands to HPE Storage commands.

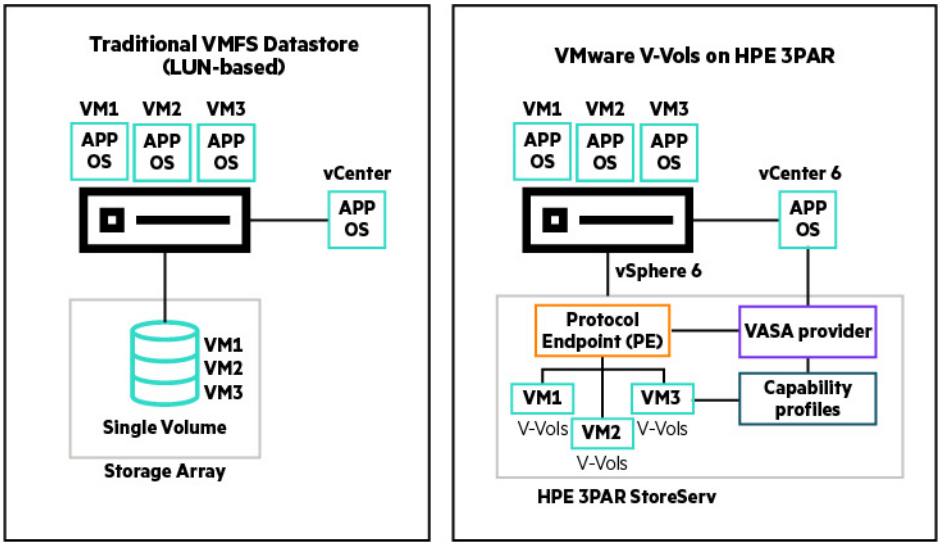

FIGURE 19. Overview of V-Vols/VASA for vCenter

- V-Vols introduces two big benefits: (1) VM-level granularity by introducing a one-to-one mapping of VMs to storage volumes, and (2) support for VMware's Storage Policy Based Management (SPBM) to simplify storage management and automate storage provisioning and space reclamation.
- Before the introduction of V-Vols, storage arrays primarily integrated with vSphere at the datastore level using VMware's VMFS. Moving forward, users can choose to use VMFS or V-Vols (or both), with V-Vols offering more advanced capabilities.

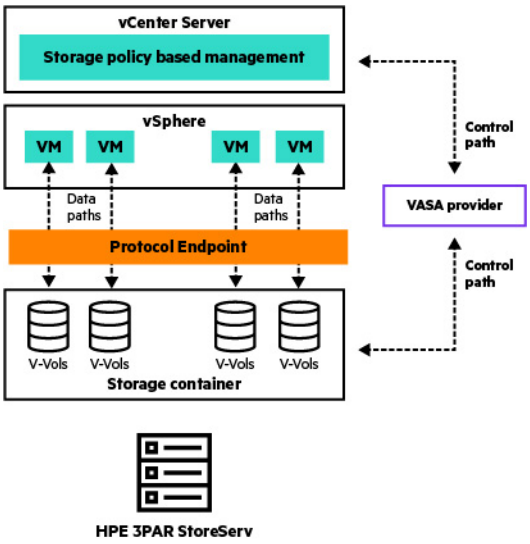

FIGURE 20. HPE 3PAR StoreServ Storage supported features with V-Vols

- The following HPE 3PAR features are supported with V-Vols:
- Array-based snapshots
- Space reclamation (VMware vSphere's UNMAP command)
- Array-based thin provisioning
- Thin deduplication
- Array-based replication
- Auto Zero Detect
- Adaptive flash cache

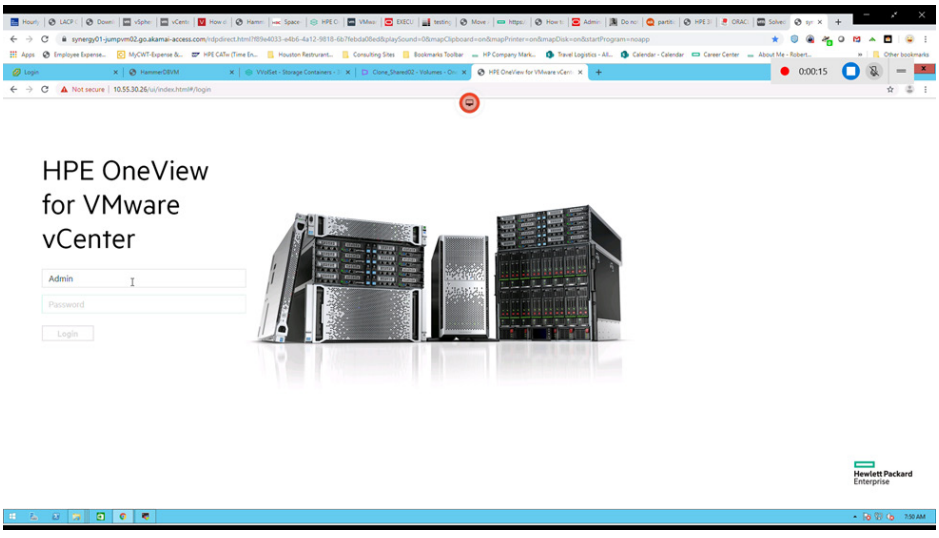

FIGURE 21. OV4VC login screen for appliance broker

- HPE OneView for VMware vCenter (OV4VC) seamlessly integrates the manageability features of HPE OneView, HPE ProLiant, HPE Synergy, HPE Virtual Connect, and HPE Storage into the VMware vCenter Server® console. By integrating HPE Converged Infrastructure management features directly into VMware vCenter Server, administrators gain insight and control of their HPE infrastructure supporting their virtualized infrastructure—reducing the time it takes to make important decisions, manage planned and unplanned downtime and do lifecycle management.

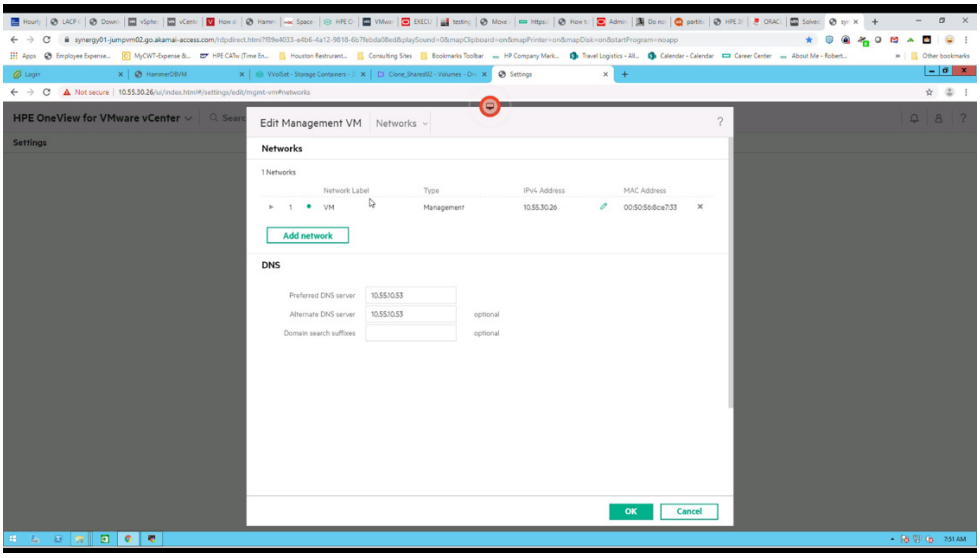

FIGURE 22. OV4VC add management VMware network

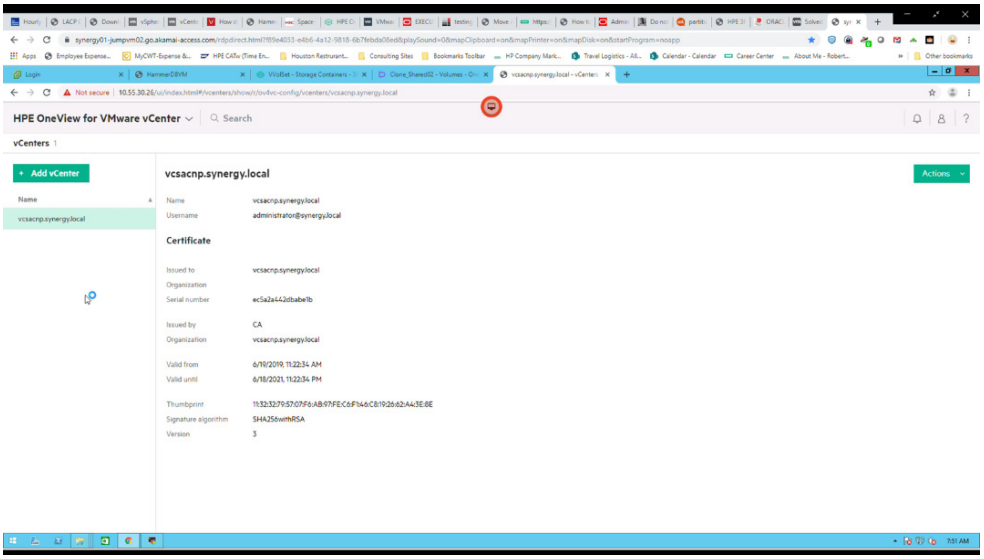

FIGURE 23. OV4VC add vCenter(s) to the broker

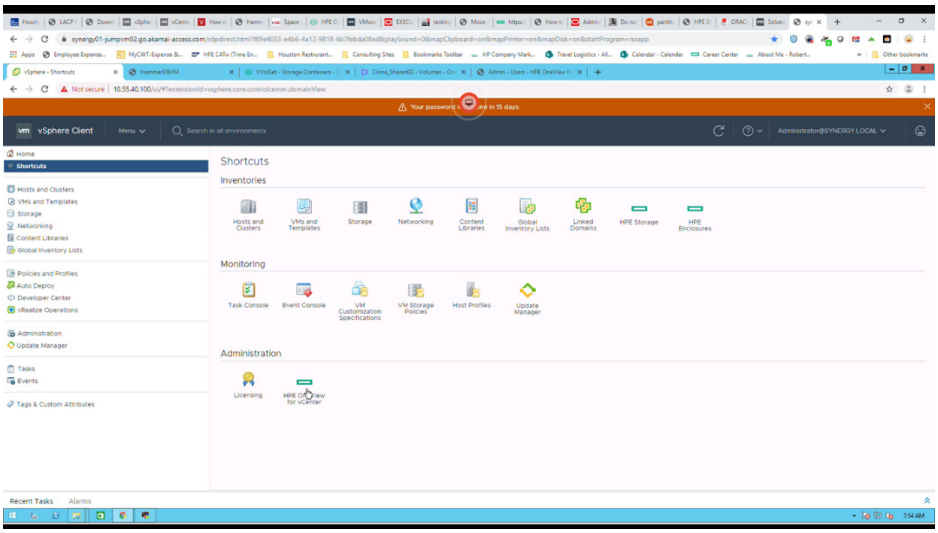

FIGURE 24. OV4VC post process

- Once vCenter and Storage Systems are defined within OneView for vCenter, additional icons become available.

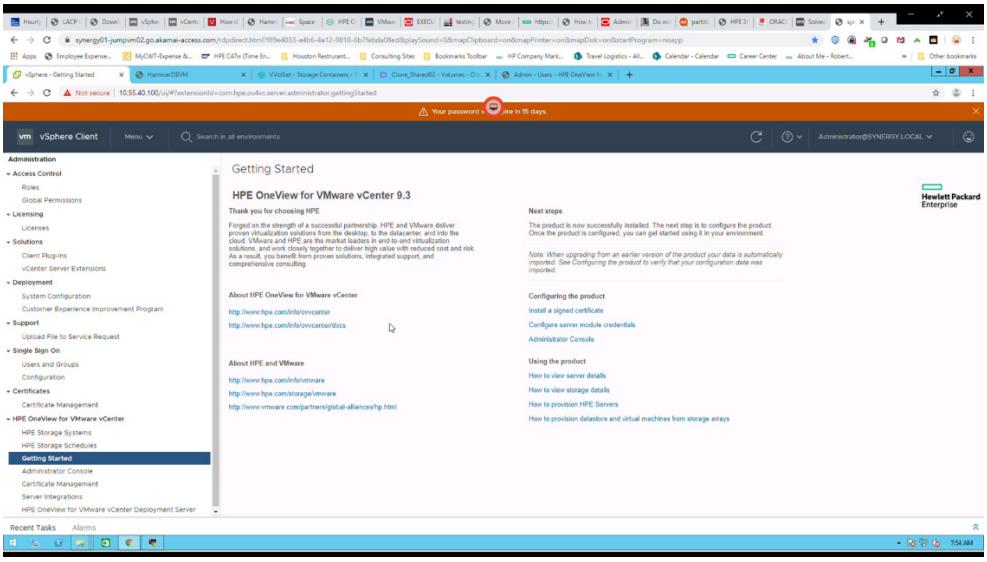

FIGURE 25. OV4VC post process—getting started

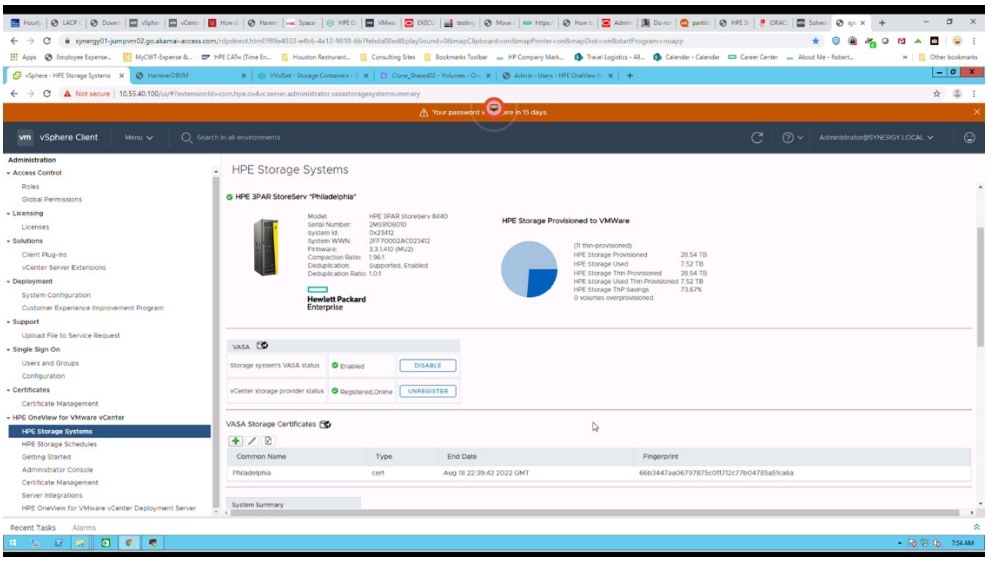

FIGURE 26. OV4VC Storage Systems and VASA status

This shows the storage devices and capacity available. You are able to validate that VASA is enabled and the vCenter storage provider is registered.

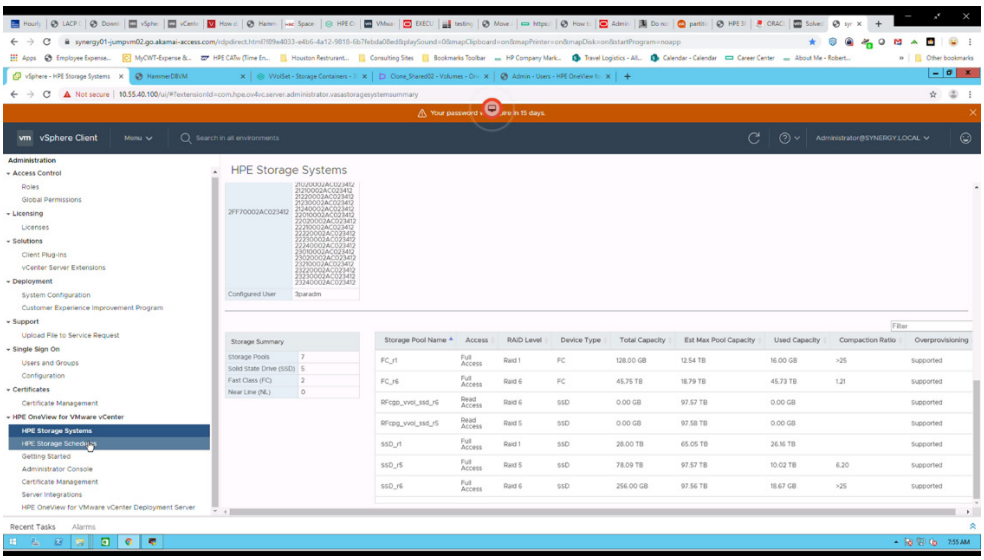

FIGURE 27. OV4VC HPE Storage schedules

This shows the storage pool and types that are defined on the storage.

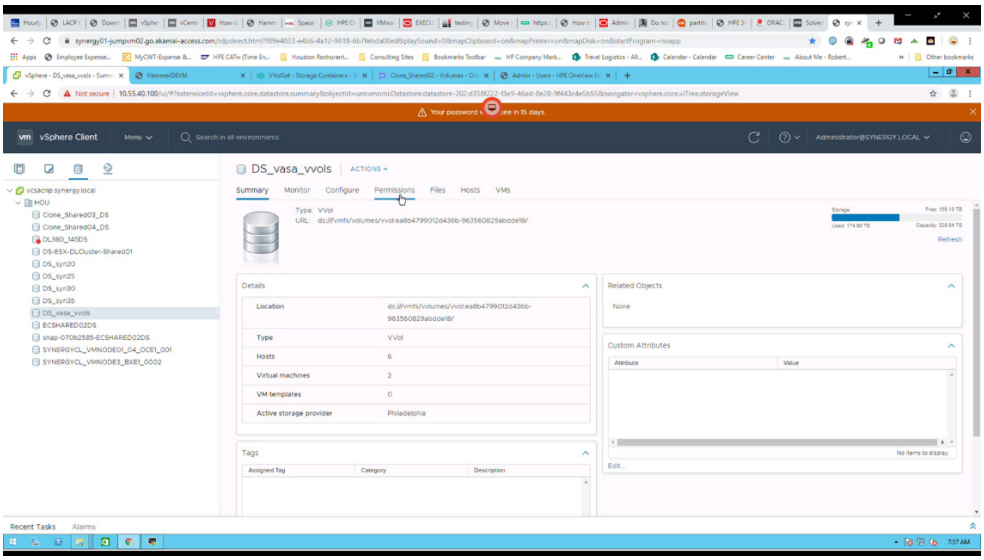

FIGURE 28. OV4VC VASA datastore container

After creating a VASA container and providing a name the container will be available and look like a datastore.

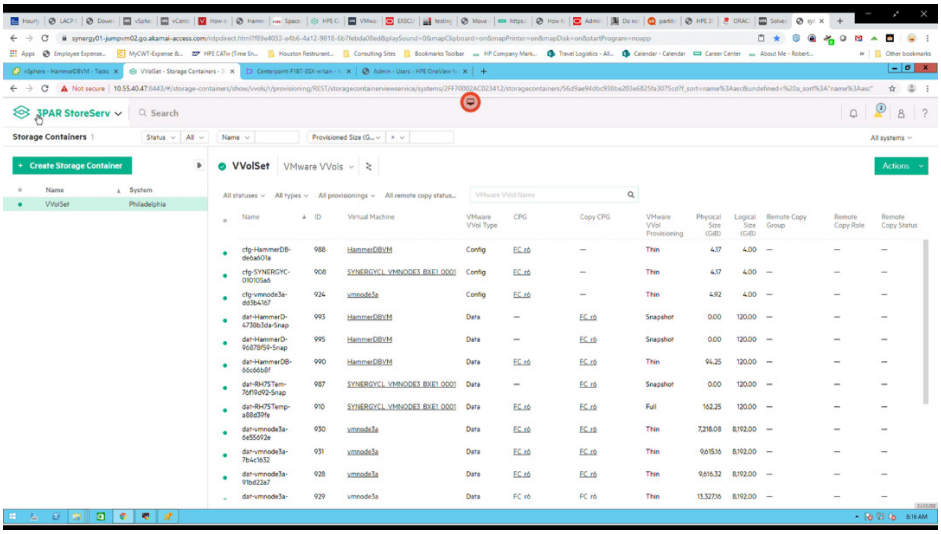

FIGURE 29. OV4VC contents of a VASA container

Exploring the details of the storage you will see LUNs much like you do with a normal HPE Storage Array LUNs. The major difference here is that each LUN on the array relates to a Virtual Machine as a file.

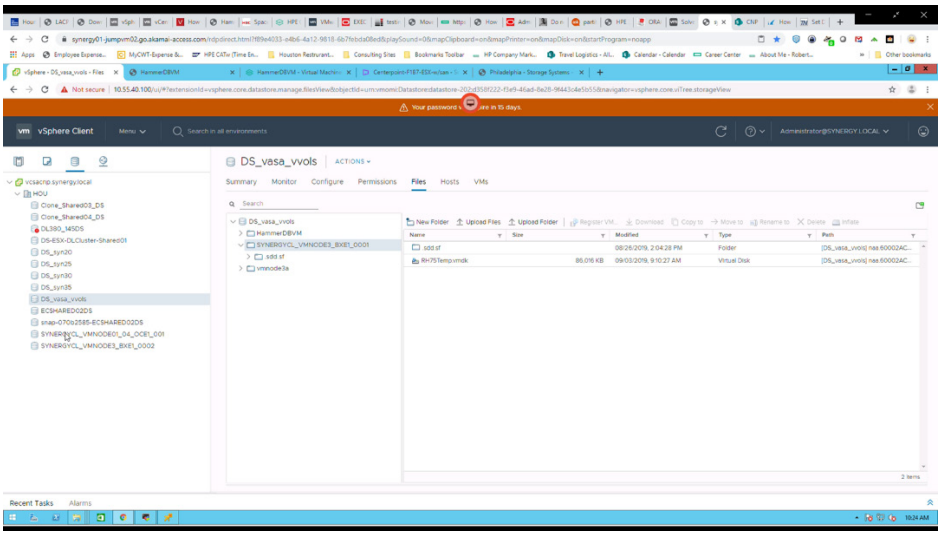

FIGURE 30. VMware file looks as a V-Vols

Exploring the details of the VASA datastore you will see files much like you do with a normal datastore. The major difference here is that each file here is a LUN on the array and not a VMDK file.

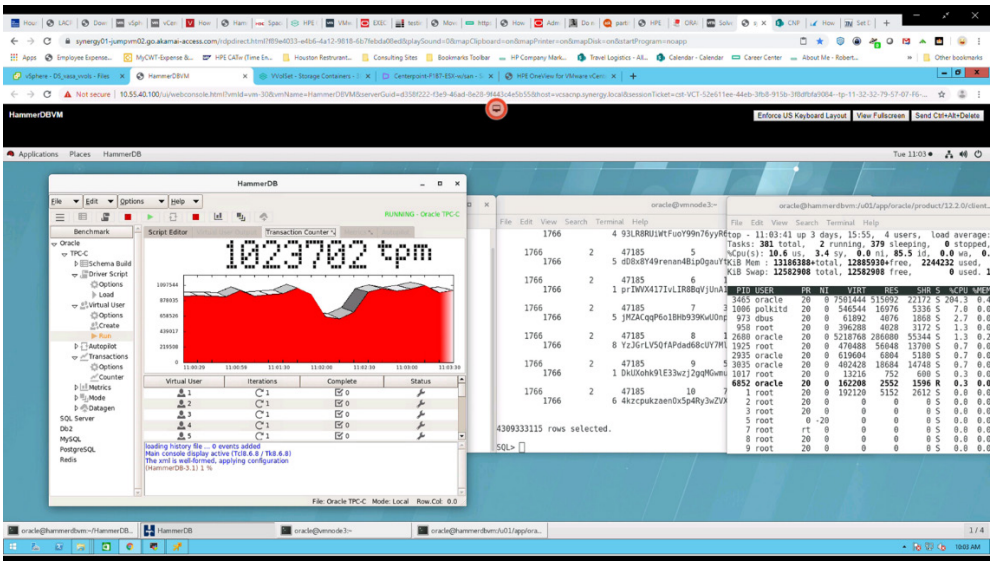

FIGURE 31. V-Vols performance snapshot

This was a standalone 60 TB database running 55 users. Each user is running 1 million transactions with no think time. This database was converted from VMDKs to V-Vols and there was no performance tuning performed. With a little additional analysis and better planning we could improve these values. However, utilizing applications and non-VLDB databases performance is well within an acceptable range and simplifies management of storage and virtual machines.

vSphere 6.7 hands-on review and performance results

- Utilizing vSphere 6.7 and V-Vols will allow for standard Microsoft Failover Clusters to be virtualized.
- Currently this is not an option unless you utilize RDM devices.
- This would help get higher levels of virtualization by virtualizing physical MSFC deployments.
- You can set tiers of storage and/or classes like Gold, Silver, and Bronze.
- VASA is large bucket.

REFERENCES

- VMware Configuration Maximums website: configmax.vmware.com/guest?vmwareproduct=vSphere&release=vSphere%207.0&categories=1-0
- Carnegie Mellon publication for US-CERT on 3-2-1 Data Backup Options: us-cert.gov/sites/default/files/publications/data_backup_options.pdf

LEARN MORE AT

HPE Education Services

HPE Pointnext Services

HPE Synergy

a00098833ENW, May 2020

Hewlett Packard Enterprise