

EDB Postgres™ Advanced Server Performance on EMC XtremIO

With VMware vSphere, VMware vCenter, and EMC PowerPath/VE

- Consistent performance
- On-demand scale-out capacity
- Agile XtremIO

EMC Solutions

Abstract

This white paper describes the advantages of deploying the EMC XtremIO all-flash array in a VMware vSphere environment to achieve optimum performance for EDB Postgres™ Advanced Server databases.

June 2016

Copyright © 2016 EMC Corporation. All rights reserved. Published in the USA.

Published June 2016

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

The information in this publication is provided as is. EMC Corporation 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 EMC software described in this publication requires an applicable software license.

EMC², EMC, PowerPath, XtremIO, X-Brick, and the EMC logo are registered trademarks or trademarks of EMC Corporation in the United States and other countries. All other trademarks used herein are the property of their respective owners.

For the most up-to-date listing of EMC product names, see [EMC Corporation Trademarks](#) on EMC.com.

**EDB Postgres™ Advanced Server Performance on EMC XtremIO
White Paper**

Part Number: H14581.2

Contents

Executive summary	4
Technology overview	6
Solution architecture	10
Solution design and configuration	12
Performance scalability testing and validation	14
Production and test/dev consolidation testing and validation	21
Conclusion	23
References	24

Executive summary

Business case Business needs are driving growth in the volume and velocity of data being collected. At the same time, the demand to quickly turn this data into information and gain insight into opportunities and risks is increasing.

Databases, such as the EDB Postgres™ Advanced Server (EDB Advanced Server) from EnterpriseDB®, are used to support business-critical applications. To deliver fast response times across the range of applications, these databases require storage designed for both low-latency transactional I/O and high-throughput analytic workloads.

Virtualization enables greater consolidation across multiple types of database workloads. Often, by consolidation, both online transaction processing (OLTP) production workloads and test/development (test/dev) workloads share the same servers and storage. Therefore, for the best performance, the underlying storage infrastructure must also be designed to handle the different workloads in a consolidated infrastructure.

The EMC® XtremIO® all-flash array effectively addresses the effects of virtualization on I/O-intensive database workloads with impressive random I/O performance and ultra-low latency. This applies equally to the consolidation of multiple workloads on a common storage platform.

XtremIO also provides new levels of speed and provisioning agility to virtualized environments with space-efficient snapshots, inline data reduction, and accelerated provisioning using VMware vStorage APIs for Array Integration (VAAI). The results include storage and database license savings, breakthrough simplicity for storage management and provisioning, and new capabilities for realtime analytics and test/dev cycles.

Solution overview Virtualizing database servers is a proven strategy for consolidating databases, but it presents unique challenges. When multiple applications are consolidated on fewer physical hosts, the I/O workload to the back-end storage can become highly random because the virtual machines share physical resources such as host bus adapters (HBAs).

XtremIO is built to perform in these demanding virtualized environments with high and consistent random I/O performance.

This solution shows the benefits of deploying an EDB Advanced Server database on VMware vSphere using XtremIO storage, including:

- Consolidated virtualized EDB Advanced Server databases on the same physical hosts and storage
- Consistent performance for application service-level agreements (SLAs) during workload scaling
- Quick creation of multiple databases for test/dev purposes with XtremIO snapshots

- Near-zero performance impact on the production workload when test/dev workloads run simultaneously

Document purpose This white paper describes and validates the solution, which deploys EDB Advanced Server on the XtremIO all-flash array in a consolidated vSphere environment, to achieve the best database performance.

Audience The primary audience for this white paper is database and system administrators, storage administrators, and system architects looking to improve database performance in a virtualized environment. Readers of this white paper should be familiar with EDB Advanced Server, EMC XtremIO, and VMware vSphere technology.

Key results This solution shows that XtremIO delivers:

- Consistent performance and high scalability on the EDB Advanced Server databases running on XtremIO. Table 1 shows the IOPS when database servers are added. For both types of workloads, the I/O latency is less than 1 millisecond (ms).
- Limited performance impact to the production workload when test/dev workloads are run on XtremIO snapshots. The total IOPS including read and write activities for the production workload was about 28,900, which we used as the baseline. When we ran the test/dev workload on snapshots together with the production workload, the IOPS on the production database decreased to about 28,000, which was about a 3 percent impact to the production workload.

Table 1. IOPS test results for workloads

Workload type	Performance statistics	One server	Two servers	Four servers
Read-only	IOPS	75,000	130,000	225,000
Mixed read/write	Aggregate IOPS (read and write)	~46,000	~75,000	~130,000

We value your feedback!

EMC and the authors of this document welcome your feedback on the solution and the solution documentation. Contact EMC.Solution.Feedback@emc.com with your comments.

Authors: Mark Wu, Aighne Kearney

Technology overview

Overview

The key technology components used in this solution are:

- EMC XtremIO
- VMware vSphere
- VMware vCenter
- EMC PowerPath®/VE
- EDB Postgres Advanced Server

EMC XtremIO 4.0

The XtremIO storage array is an all-flash system based on a scale-out architecture. The system uses EMC XtremIO X-Brick™ building blocks, which can be clustered together, as shown in Figure 1, to grow performance and capacity as required.




Figure 1. EMC XtremIO configuration for multiple X-Brick clusters

The XtremIO array architecture is specifically designed to deliver the full performance potential of flash while linearly scaling all resources such as CPU, RAM, solid-state drives (SSDs), and host ports in a balanced way. This enables the array to achieve any wanted performance level while maintaining performance consistency that is critical to predictable application behavior.

With clusters of two or more X-Brick building blocks, XtremIO uses a redundant 40 Gb/s QDR InfiniBand network for back-end connectivity between the storage controllers, ensuring a highly available, low-latency network. The InfiniBand network is a fully managed component of the XtremIO array, and administrators of XtremIO arrays do not need specialized skills in InfiniBand technology.

The cluster's performance level is not affected by its capacity utilization level, number of volumes, or aging effects. Performance is not based on a shared cache architecture and is not affected by the dataset size or data access pattern.

Due to its content-aware storage architecture, XtremIO provides:

- An even distribution of data blocks, leading to maximum performance and minimal flash wear
- An even distribution of metadata
- No data or metadata hotspots
- Easy setup and no tuning
- Advanced storage functionality, including inline data reduction and compression, thin provisioning, advanced data protection (XDP), and snapshots

XtremIO is the ideal storage platform for accelerating EDB Advanced Server databases, providing outstanding transactional speed with consistently low latency. XtremIO automatic balanced-data distribution enables consistent performance across the entire lifecycle of the volumes, regardless of I/O patterns, data placement, or cluster size.

With the benefits of flash-optimized algorithms and active/active controllers, XtremIO provides consistent performance even under mixed workloads. This performance stays completely balanced, and latency remains under a millisecond even at scale.

For more information, refer to the *EMC XtremIO Storage Array User Guide*.

XtremIO Management Server

XtremIO Management Server (XMS) is a stand-alone, dedicated Linux-based server that is preinstalled with the CLI and GUI and is used to control XtremIO system operations. XMS can be either a physical or a virtual server. The array continues operating if it is disconnected from XMS but cannot be configured or monitored.

XtremIO inline data reduction

XtremIO inline data reduction is achieved by using inline data deduplication and inline data compression. Data deduplication and data compression complement each other. Data deduplication removes redundancies from data before it is written to the flash media. Data compression compresses the deduplicated data before it is written to the flash media.

Because XtremIO is also a fully thin-provisioned storage system, more storage savings are realized for overall efficiency.

XtremIO snapshots

XtremIO snapshots are instantly copied images of volume data with the state of the data captured exactly as it appears at the specific point in time that the snapshot is created. The volume data state and the specific volume data can be accessed when required, even when the source volume has changed. Creating snapshots, which can be done at any time, does not affect system performance. A snapshot can be taken either directly from a source volume or from other snapshots within a source volume's group (volume snapshot group).

The original copy of the data remains available without interruption, while the snapshot can be used to perform other functions on the data. Changes made to the snapshot's source do not change or impact on the snapshot data.

XtremIO snapshots are read/write, although you can choose to mount the snapshots in read-only mode to maintain their immutability.

Snapshots and inline data reduction

XtremIO snapshot technology is implemented by using the content-aware capabilities of the system (inline data reduction), optimized for SSD media. It has a unique metadata tree structure that directs I/O to the right timestamp of the data. This enables efficient snapshotting that can sustain high performance, while maximizing the media endurance, both in terms of the ability to create multiple snapshots and the amount of I/O that a snapshot can support.

Snapshot creation and performance impact

When creating a snapshot, the system generates a pointer to the ancestor metadata (of the actual data in the system). Therefore, creating a snapshot is a very quick operation that does not have any impact on the system and does not consume any capacity. Snapshot capacity consumption occurs only if a change requires writing a new unique block.

XtremIO snapshots are space-efficient both in terms of additional metadata consumed and physical capacity. Snapshots are implemented using redirect-on-write method, where new writes to the source volume (or snapshot) are redirected to new locations, and only metadata is updated to point to the new data location. This method guarantees no performance degradation while snapshots are created.

Snapshot accessibility

Snapshots can be accessed like any other volume in the cluster in read/write access mode, and enable a wide range of uses, including:

- Recovery from logical corruption
- Backups
- Test/dev
- Clones
- Offline processing

VMware vSphere 6.0

VMware vSphere 6.0 is a virtualization platform that empowers users to virtualize any application, redefines availability, and also simplifies the virtual data center. An ideal foundation for any cloud environment, vSphere ensures a highly available, resilient, and on-demand infrastructure for computing, networking, and storage resources.

vSphere 6.0 features the following key scalability and performance enhancements, which enable a virtual machine to use more resources from the hypervisor:

- Support for vSphere clusters to scale to up to 64 hosts and 8,000 virtual machines in a single cluster
- Support for 480 logical CPUs, 12 TB of RAM, and 1,024 virtual machines in each hypervisor instance
- Virtual machine hardware version 11, which provides up to 128 vCPUs and 4 TB of RAM
- Expanded support for the latest x86 chip sets, devices, drivers, and guest operating systems
- Virtual Volumes feature that enables the external storage arrays to become virtual-machine-aware
- Support for nondisruptive live migration of workloads across virtual switches and vCenter Servers
- Expanded support for software-based fault tolerance for workloads with up to four vCPUs
- Support for the VMware Paravirtual SCSI (PVSCSI) adapter with virtual machines running Windows Server Failover Clustering (WSFC)

VMware vCenter 6.0

VMware vCenter is a centralized management system for a VMware virtual infrastructure. This system provides you with a single interface that you can access from multiple devices for all aspects of monitoring, managing, and maintaining a virtual infrastructure.

vCenter also manages some advanced features of a VMware virtual infrastructure, such as VMware vSphere High Availability (HA), VMware vSphere Distributed Resource Scheduler (DRS), VMware vMotion, and VMware vCenter Update Manager.

EMC PowerPath/VE 6.0

EMC PowerPath/VE provides intelligent, high-performance path management with path failover and load balancing optimized for EMC and selected third-party storage systems. PowerPath/VE 6.0 redistributes the I/O load across multiple paths between the vSphere host and an external storage device to perform load balancing.

PowerPath/VE fully supports the vSphere 6.0 virtualization platform. PowerPath/VE works with the VMware ESXi physical server as a multipath plug-in (MPP) that provides path management to hosts. The MPP is installed as a kernel module on the vSphere host. It plugs into the vSphere I/O stack framework to bring the advanced multipathing capabilities of PowerPath/VE, including dynamic load balancing and automatic failover, to the vSphere hosts.

EDB Postgres™ Advanced Server 9.5

EDB Postgres Advanced Server provides powerful performance and security enhancements for PostgreSQL, sophisticated management tools for global deployments, and database compatibility. It supports both mission and non-mission-critical enterprise applications.

Solution architecture

Overview

This section describes the solution architecture that provides the best cost-to-performance ratio for EDB Advanced Server mission-critical application environments.

In the solution, multiple EDB Advanced Server databases are deployed as virtualized databases on an XtremIO storage array consisting of two X-Brick cluster. Four ESXi servers are used in the solution to provide the virtualization layer.

To get the best performance of EDB Advanced Server and scalability, we deployed one virtual machine on each ESXi server, for a total of four virtual machines running the workload, including read-only and read/write mixed workloads. We created one 200 GB EDB Advanced Server database on each virtual machine, for a total of 800 GB in database storage running on XtremIO.

Features of EDB Advanced Server 9.5

Features in the 9.5 release that contribute to performance improvement include:

- Improved concurrency for shared buffer replacement
- Reduced number of page locks and pins during index scans, enabling index vacuums to be blocked less often
- More memory-efficient per-backend tracking of buffer pins
- Improved lock scalability to address problems on systems with multiple CPU sockets
- Optimizer ability to remove unnecessary references to left-joined subqueries
- Pushdown of query restrictions into subqueries with window functions
- Pushdown of a nonleakproof function into a security barrier view if the function does not receive any view output columns
- Improved bitmap index scan performance
- Increased number of buffer mapping partitions, boosting performance for highly concurrent workloads

Solution architecture

Figure 2 shows the reference architecture of this solution.




Figure 2. Solution reference architecture

The architecture is composed of the following:

- **Compute layer:** Includes four servers to enable a high-performing, consolidated, virtualized approach to database infrastructure, resulting in deployment flexibility without the need for application changes.
- **Network layer:** Includes two IP switches and two director-class SAN switches, which we configured to produce 108 GB/s active bandwidth. The SAN switches are designed for deployment in storage networks supporting virtualized data centers and enterprise clouds.
- **Storage layer:** Includes an XtremIO storage array with a two X-Brick cluster, with 30.49 TB of usable physical capacity.

The servers are installed with vSphere and configured as an ESXi cluster.

Hardware resources

Table 2 lists the hardware resources used in the solution.

Table 2. Hardware resources

Hardware	Quantity	Configuration
Storage array	1	XtremIO array including: <ul style="list-style-type: none"> • Single two X-Brick cluster (12U) • 800 GB SSD, 30.49 TB physical capacity in total
Servers	4	32 cores, 2.90 GHz processors, 512 GB RAM, including: <ul style="list-style-type: none"> • 1 x 1 Gb Ethernet (GbE) network interface card (NIC) • 1 x 10 GbE NIC
LAN switches	2	10 GbE
SAN switches	2	8 Gb/s Fibre Channel (FC)

Software resources Table 3 lists the software resources used in this solution.

Table 3. Software resources

Software	Version	Notes
EMC XtremIO Operating System (XIOS)	4.0.1-7	Operating system for XtremIO
EMC XtremIO	4.0	Storage array
EMC PowerPath/VE	6.0	Multipathing and load balancing with I/O path optimization
VMware vSphere	6.0	Hypervisor hosting all virtual machines
VMware vCenter	6.0	vSphere management server
Red Hat Enterprise Linux	6.5	Operating system for database servers
EDB Postgres Advanced Server	9.5	Enterprise database system
Pgbench	9.5	OLTP benchmark tool

Solution design and configuration

Overview

This section describes the XtremIO storage design and host SAN networking configuration used in this solution. You must understand the design considerations, including host connectivity for the XtremIO storage array, to ensure the best availability and performance for your hosts.

This section also describes virtual design and configuration for the vSphere platform, and application considerations for EDB Advanced Server.

EMC recommends that you check the latest host configuration guide for XtremIO and vSphere design considerations before building your environment.

Storage design

With the traditional storage design for the EDB Advanced Server database, multiple RAID groups of different drive types are created, each with different levels of protection and distributed across multiple controllers.

With XtremIO, all drives are under XDP protection, and data in the array is distributed across the X-Brick cluster to maintain consistent performance and equivalent flash wear levels. The workloads on the databases generate random I/O on XtremIO, and these are distributed evenly throughout the array.

Table 4 lists the storage layout for each database.

Table 4. Storage design for each EDB Advanced Server database on the XtremIO X-Brick cluster

Volume type	Size (GB)
Data	200
Log	10

Virtualization design

Virtual machine template configuration

The choice of a server platform for a virtualized infrastructure is based on both the supportability of the platform and the technical requirements of the environment. In production environments, the servers must have sufficient:

- Cores and memory to support the required number and workload of the virtual machines
- Connectivity, both Ethernet and FC, to enable redundant connectivity to the IP and storage network switches
- Capacity to withstand a server failure and support failover of the virtual machines

We used four physical servers configured as a vSphere HA cluster, with each server running an ESXi server. We configured the virtual machine template in the vSphere client according to the requirements and prerequisites for the database software, as shown in Table 5.

Table 5. Virtual machine template configuration

Component	Description
CPU	16 vCPUs
Memory	To generate stable IOPS, we reduced the memory allocated to the virtual machines as follows: <ul style="list-style-type: none"> • 512 MB for read-only workload • 1 GB for read/write mixed workload
Operating system	Red Hat Enterprise Linux Server Release 6.5

Component	Description
Virtual network interfaces	Eth0: Public/management IP network
Preinstalled software	EDB Postgres Advanced Server 9.5

After deploying the virtual machines, we added the disks as VMDK for the database storage.

For more information about best practices for VMware virtualization, refer to [References](#).

Enabling access to the storage devices

To implement storage design for each database, we did the following:

- Created two volumes in the XtremIO array for each database, one for datafiles and the other for log files.
- Created one data store on each LUN after presenting the data stores to the ESXi server, one for datafiles and the other for log files.
- Added two VMDK disks on each virtual machine on the two data stores. The two disks were spread across multiple VMware Paravirtual SCSI (PVSCSI) controllers to balance I/O as follows:
 - One disk (200 GB) used for datafiles was assigned to SCSI Controller 1
 - One disk (10 GB) used for log files was assigned to SCSI Controller 2

EDB Advanced Server design

We created four virtual machines to run four databases, with one database on each virtual machine, to simulate OLTP-like workloads.

Table 6 lists each database workload profile in the solution.

Table 6. Database workload profile for each OLTP database

Profile characteristic	Details
Database	EDB Postgres Advanced Server 9.5
Database type	OLTP-like
Database size	200 GB
Workload profile	OLTP workload simulated by pgbench
Data block size	8 KB

Performance scalability testing and validation

Overview

To measure the performance of EDB Advanced Server on XtremIO, we used two workloads: a read-only workload and a mixed (read and write) workload. For each type of workload, we began with one virtual machine, and then scaled out to two virtual machines, and then scaled out to four virtual machines, to test the scalability

of the performance. Each virtual machine was configured with one 200 GB EDB Advanced Server database.

We ran workloads generated by pgbench, and gathered the following performance metrics:

- I/O performance metrics (IOPS and latency) from storage
- Transactions per second (TPS) from pgbench

Notes on results

Benchmark results are highly dependent on workload, specific application requirements, and system design and implementation. Relative system performance varies as a result of these and other factors. Therefore, this workload should not be used as a substitute for a specific application benchmark when critical capacity planning or product evaluation decisions are contemplated.

All performance data contained in this report was obtained in a rigorously controlled environment. Results obtained in other operating environments could vary significantly.

EMC Corporation does not warrant or represent that a user can achieve similar performance expressed in transactions per minute.

Test objectives

The overall test objectives were to:

- Show the high performance achieved when virtualized EDB Advanced Server databases are running on XtremIO
- Show the high scalability of workloads running on XtremIO during the scale-up of EDB Advanced Server databases

Test scenarios

We conducted the following test scenarios on the solution:

- Read-only workload
- Mixed workload (read/write)

Test procedure

Data population

We ran the following pgbench command to populate four 200 GB EDB Advanced Server databases in four virtual machines, that is, one database in each virtual machine for a total of 800 GB data populated.

```
pgbench -i -s 7500 -index-tablespace=<tablespace name> -
tablespace=<tablespace name> <database name>
```

The options in the command are:

- **-i:** Run initialization mode
- **-s:** Multiply the number of rows generated by the scale factor. For example, -s 100 creates 10,000,000 rows in the pgbench_accounts table. Default is 1.
- **-index_tablespace:** Create indexes in the specified tablespace, rather than the default tablespace

- **-tablespace:** Create tables in the specified tablespace, rather than the default tablespace

Read-only workload test and scalability

We ran the following `pgbench` command to generate read-only workload against the EDB Advanced Server database on each virtual machine:

```
pgbench -h <VM ip> -p <port number> -c 50 -S -T 1800 -U EDB  
Postgres <database name>
```

We ran the command to generate 50 concurrent sessions with each session running the similar `SELECT` statements for 30 minutes.

The options in the command are:

- **-h:** The database server's hostname or IP address
- **-p:** The database server's port number
- **-c:** Number of clients simulated, which means the number of concurrent database sessions
- **-S:** Perform select-only transactions
- **-T:** Run the test for this many seconds
- **-U:** The username to connect to the database

We ran the test using the following procedure to validate the performance and workloads while scaling up both. For each step, we recorded the IOPS and latency from the XtremIO array and the TPS from the `pgbench` output:

1. Ran the command on one virtual machine to generate a read-only workload
2. Ran the command on two virtual machines simultaneously to double the workload
3. Ran the command on four virtual machines to quadruple the workload

Mixed workload (read/write) test and scalability

We ran the following `pgbench` command to generate a mixed workload with a 2:1 read/write ratio against EDB Advanced Server database on each virtual machine:

```
pgbench -h <VM ip> -p <port number> -c 100 -N -T 1800 -U EDB Postgres <database  
name>
```

The command generated 100 concurrent sessions, with each session running the similar `SELECT/UPDATE` statements for 30 minutes.

The options in the command are:

- **-h:** The database server's hostname or IP address
- **-p:** The database server's port number
- **-c:** Number of clients simulated, which means the number of concurrent database sessions

- **-N:** Do not update `pgbench_tellers` and `pgbench_branches`. This avoids update contention on these tables
- **-T:** Run the test for this many seconds
- **-U:** The username to connect to the database

We ran the test using the following procedure. For each step, we recorded the IOPS and latency from the XtremIO array and TPS from the database:

1. Ran the command on one virtual machine to generate a mixed read/write workload
2. Added one additional virtual machine into the workload and ran the command on the two virtual machines simultaneously, to double the workload
3. Add two additional virtual machines into the workload and ran the command on the four virtual machines, to quadruple the workload.

Read-only workload test results

Table 7 lists the physical read IOPS, latency, and TPS, which increased with the database host scaling. The IOPS increased along with the workload scaling while the latency was kept under 1 ms.

Table 7. Scaling of nodes and resulting increases in IOPS and TPS

Metric	1 virtual machine	2 virtual machines	4 virtual machines
IOPS	75,000	130,000	225,000
Read latency (ms)	0.24	0.35	0.44
TPS	22,632	37,234	63,006

Note: The IOPS and latency listed in Table 7 are the average numbers calculated from the results exported from XtremIO CLI using the `export-performance-history` command.

Table 8 shows the physical read IOPS captured from the XtremIO console.

Table 8. Physical read IOPS captured from XtremIO console

No. of virtual machines	IOPS screenshot captured from XtremIO console
One	
Two	
Four	

Figure 3 shows the results of the read-only scaling test. The IOPS increased as the workload scaled up.




Figure 3. IOPS and TPS scalability of read-only workload

Mixed workload (read/write) test results

Table 9 lists the physical read and write IOPS, latency, and TPS, which increased with the scaling of the database nodes. The IOPS increased as the workload scaled up, while the read and write I/O latency was kept under 1 ms. Scaling of nodes and resulting increases in IOPS and TPS.

Table 9. Scaling of database servers and resulting increases in IOPS and TPS

Metric	1 virtual machine	2 virtual machines	4 virtual machines
Read IOPS	~38,000	~60,000	~105,000
Write IOPS	~8,000	~15,000	~25,000
Total IOPS	~46,000	~75,000	~130,000
Read latency (ms)	~0.25	~0.35	~0.5
Write latency (ms)	~0.45	~0.61	~0.85
TPS	7,642	12,853	20,815

Note: The IOPS and latency listed in Table 9 are the performance numbers for the data volume. These are the approximate values captured from the XtremIO GUI rather than the exported IOPS and latency from XtremIO CLI. We did not use the exported IOPS and latency from the XtremIO CLI because the exported performance data based is on the array level rather than volume level and includes both production and snapshots volumes.

Table 10 shows the physical read and write IOPS captured from the XtremIO console.

Table 10. Physical read and write IOPS captured from XtremIO console

No. of virtual machines	IOPS screenshot captured from XtremIO console
One	
Two virtual machines	
Four virtual machines	

Figure 4 shows the results of the mixed workload (read/write) scaling test. The total IOPS (including read and write IOPS) increased as the workload scaled up.




Figure 4. IOPS and TPS scalability of mixed workload

Production and test/dev consolidation testing and validation

Overview

We validated the performance impact to the production workload when running OLTP-like workloads on multiple XtremIO snapshots. In this test, we provisioned five snapshots against the production volumes and mounted them to five additional virtual machines, which were used as five test/dev databases. We then ran five mixed workloads (read/write), simulated by pgbench, against the test/dev databases.

Test objectives

The overall test objectives were to:

- Show how easy it is to create snapshots in XtremIO
- Show that workloads on XtremIO snapshots have a limited impact on the production workload

Test procedure

We ran the test using the following procedure:

1. Ran the production workload with the following command and recorded the performance data for the baseline:


```
> pgbench -h <VM ip> -p <port number> -c 100 -N -T 1800 -U EDB Postgres <database name>
```
2. Created five consistent snapshots against the production volumes including data and log, while the production database was running, with no impact to the production workload.

Note: You can quickly create snapshots with a few mouse clicks using the XtremIO Storage Management Application console.

3. Mounted the five snapshots on five additional virtual machines respectively, which were on one physical server. Opened the databases in the virtual machines, which were used as test/dev environments.
4. Ran the same command as the baseline on the production database, while simultaneously running five workloads on the test/dev databases.
5. Ran the following command, which resulted in a total workload on the test/dev environments of about 30 percent of the production workload:


```
pgbench -h <VM ip> -p <port number> -c 2 -N -T 1800 -U EDB
Postgres <database name>
```

Test results

Table 11 shows the physical read and write IOPS, latency, and TPS, including baseline and production performances with test/dev workloads running simultaneously on snapshots.

Table 11. Baseline and performance with test/dev workloads running

Metric	Baseline	Production and test/dev running simultaneously on snapshots	
	Production	Production	Snapshots
Read IOPS	~38,000	~36,000	~8,500
Write IOPS	~8,000	~8,000	~2,500
Total IOPS	~46000	~44,000	~11,000
Read latency	~0.25	~0.3	~0.35
Write latency	~0.45	~0.53	~0.76
TPS	7,642	7,345	2,506

Note: The IOPS and latency listed in Table 11 are the approximate performance numbers captured from the XtremIO GUI rather than the exported numbers from the XtremIO CLI. We did not use the exported IOPS and latency from the XtremIO CLI because the exported performance data was based on the array level rather than volume level and includes both production and snapshots volumes.

Compared to the baseline, the workloads on the XtremIO snapshots have very limited impact on the production workload:

- The production IOPS decreased by about 4.3 percent, while the latency was kept under 1 ms for both baseline and production workloads when the test/dev workloads ran simultaneously.
- The production TPS decreased about 3.9 percent.

Conclusion

Summary

This validated solution uses EMC XtremIO as the back-end storage for the EDB Postgres Advanced Server database in a virtualized environment.

The XtremIO array can keep up with scaling on the host side. The performance scales when you add additional compute resources including CPUs, memory, HBA ports, and front-end ports, to provide higher IOPS for OLTP environments.

The XtremIO snapshot performance overhead is negligible when you take multiple crash-consistent snapshots of a production database and then run workloads against the production and snapshot databases.

Findings

The key findings of the solution are:

- You can achieve consistent performance and high scalability when running an EDB Postgres Advanced Server database on XtremIO, with low latency (less than 1 ms) and increased IOPS.
- Running test/dev workloads on XtremIO snapshots has limited performance impact (about 3 percent) on the production database workload.
- XtremIO configuration is fast and simple, requiring no storage tuning.
- For mixed workloads, XtremIO storage with EDB Postgres Advanced Server 9.5 provides significant improvement compared to EDB Advanced Server 9.4.

References

EMC documentation

The following documentation on [EMC.com](#) or [EMC Online Support](#) provides additional relevant information. Access to these documents depends on your login credentials. If you do not have access to a document, contact your EMC representative.

- [Introduction to XtremIO Virtual Copies White Paper](#)
- [XtremIO Data Protection \(XDP\) White Paper](#)
- [XtremIO Snapshots – Do you use film or digital photography? \(blog post\)](#)

VMware documentation

For more information about vSphere, refer to the [VMware vSphere 6 Documentation](#) information hub.

Other documentation

The following documentation on the EnterpriseDB website provides additional relevant information about EDB Postgres Advanced Server:

- [EDB Postgres Advanced Server technical documentation](#)
- [EDB Postgres white papers and data sheets](#)