

TECHNICAL WHITE PAPER

Rapid Restore of Oracle Using FlashBlade and NetBackup

Modernize data protection with all-flash storage and deduplication technology.

Contents

- Introduction..... 3
- A Modern Approach to Backup and Recovery..... 3
- Pure Storage FlashBlade 4
- Veritas NetBackup..... 5
- Solution Overview..... 5
- Performance Tests..... 6
 - Test Strategy..... 6
 - Performance Test Results..... 7
- Data Reduction Test..... 8
 - Test Strategy..... 8
- Best Practices..... 8
 - Provisioning for Deduplicated Data 8
 - Compression..... 9
 - Encryption..... 9
 - Disaster Recovery 9
 - NFS Options..... 9
 - Planning..... 9
 - Sizing..... 10
- Conclusion..... 10
- Appendix..... 11
 - Next Steps..... 11
 - Supporting Information..... 11
- About the Author..... 12

Introduction

This paper provides an overview of the integration of Veritas NetBackup™ and Pure Storage® FlashBlade™. It discusses best practices for using FlashBlade as an NFS target for NetBackup deduplicated data when backing up and restoring Oracle database(s), as well as performance and data reduction tests on the integrated solution. The target audience for this document includes, but is not limited to Oracle database administrators, storage administrators, IT managers, system architects, sales engineers, field consultants, professional services, and partners who are looking to design and deploy NetBackup for Oracle backup and recovery on FlashBlade as the secondary target. The content assumes a working knowledge of NetBackup, Oracle, Linux, server, storage, and networks.

A Modern Approach to Backup and Recovery

Around-the-clock global operations, along with the explosion of data growth, have impacted backup and recovery strategies at many organizations. Traditional backup methods can no longer meet your backup windows and recovery objectives. Meanwhile, growing demand for system resources and network bandwidth puts pressure on your overall application and database performance. This can expose your business to unacceptable risks. Modern approaches to backup and recovery include the introduction of high-speed, storage-based snapshot technologies and application integration. However, many organizations still face the need to manage snapshot copies separately from traditional backup processes, complicating data management and increasing cost and complexity.

In this complex environment, emerging technologies have changed the nature and perception of cold data. The most successful organizations no longer consider data a liability—instead, they view it as a value-creating asset. In the past, organizations approached backup data like an insurance policy, keeping it around in case of an unexpected event or for regulatory compliance. Today, data-first enterprises leverage their backup data to power test/dev environments, fuel datasets for AI and modern analytics workloads, and prepare for new data-intensive regulations like the General Data Protection Regulation (GDPR).

Data has evolved, creating new data protection requirements. Today, you're truly protected only when you can keep up with demanding SLAs and rapidly recover from any unexpected events. In addition, you need to modernize data protection to rapidly restore and reuse data to support GDPR and analytics initiatives.

Meet these challenges by combining Pure Storage solutions with modern backup and restore solutions like Veritas NetBackup that utilize deduplication technology. Pure Storage FlashArray™ enables you to accelerate the performance of your mission-critical workloads. And by integrating Veritas NetBackup with FlashBlade's scale-out, all-flash storage, you can bypass the limits of disk and tape to rapidly restore mission-critical workloads.



Pure Storage FlashBlade

Pure Storage developed the FlashBlade architecture to meet the storage needs of data-driven businesses. FlashBlade is an all-flash system, primarily optimized for storing and processing unstructured data. A FlashBlade system can simultaneously host multiple file systems and multi-tenant object stores for thousands of clients.

FlashBlade gives you a scale-out, all-flash storage system, powered by a distributed file system purpose-built for massive concurrency across all data types. It can scale up to multi-petabyte capacity with linear-scale performance, simply by adding a single blade at a time, up to 150 blades. Due to its native scale-out architecture and ability to drive performance for any type of workload, FlashBlade serves as a data hub that enables you to consolidate a range of workloads, from backup to analytics and AI, on a single platform.




Figure 1. FlashBlade components

Many organizations build their data-protection strategy with FlashBlade, enjoying rapid backup and restore performance while investing in a platform that enables the consolidation data lakes and other data silos. FlashBlade enables you to scale performance and capacity with five key innovations:

- **High-performance storage device:** FlashBlade maximizes the advantages of an all-flash architecture by storing data in storage units and ditching crippling, high-latency storage media such as traditional spinning disks and conventional solid-state drives. The integration of scalable NVRAM into each storage unit helps scale performance and capacity proportionally when you add new blades to a system.
- **Unified network:** A FlashBlade system consolidates high-communication traffic between clients and internal administrative hosts into a single, reliable high-performing network that supports both IPv4 and IPv6 client access over Ethernet links up to 160Gb/s.
- **Purity//FB storage operating system:** With its symmetrical operating system running on FlashBlade's fabric modules, Purity//FB minimizes workload balancing problems by distributing all client operation requests evenly among blades.
- **Common media architectural design for files and objects:** FlashBlade's single underlying media architecture supports concurrent access to files via a variety of protocols such as NFSv3, NFS over HTTP, and SMB (with Samba-level functionality) and objects via S3 across the entire FlashBlade configuration.



- **Simple usability:** Purity//FB on FlashBlade alleviates system management headaches. It simplifies storage operations by performing routine administrative tasks autonomously. With a robust operating system, FlashBlade is capable of self-tuning and providing system alerts when components fail.

A full FlashBlade system configuration consists of up to five self-contained rack-mounted chassis interconnected by high-speed links to two external fabric modules (XFM). At the rear of each chassis are two on-board fabric modules for interconnecting the blades, other chassis, and client systems using TCP/IP over high-speed Ethernet. Both interconnected fabric modules contain a control processor and Ethernet switch ASIC. For reliability, each chassis includes a redundant power supply and cooling fan.

The front of each chassis holds up to 15 blades for processing data operations and storage. Each blade assembly is a self-contained compute module equipped with processors, communication interfaces, and 17TB or 52TB of flash memory for persistent data storage.

The current FlashBlade system can support more than 1.5 million NFSv3 getattrs commands per second, >17 GiB/sec of 512KiB reads, or >8 GiB/sec of 512KiB overwrites on a 3:1 compressible dataset in a single 4U chassis with 15 blades. It can scale both compute and performance up to a 10-by-4U chassis with 150 blades.

Veritas NetBackup

Veritas offers industry-leading enterprise data-management solutions that integrate broadly across its product portfolio and build on the company's market-leading NetBackup data-protection solution. Long recognized as a market-share leader in enterprise backup and recovery software, Veritas NetBackup protects the largest and most demanding multicloud and data center environments. NetBackup provides breakthrough capabilities for virtualized and cloud-based deployments. Features like high performance, intelligent automation, and centralized management based on a flexible, multi-tier architecture enable it to adapt to the growing needs of a fast-paced, modern enterprise.

One of the hallmarks of enterprise IT is its heterogeneity. The wide variety of platforms, applications, and infrastructure often grows with the size of the enterprise. NetBackup supports a vast array of environments and integrates with every layer of the infrastructure stack to unify your entire data protection strategy. NetBackup software offers backup to storage targets including tape, storage array network (SAN), network-attached storage (NAS), public and private cloud.

A key NetBackup feature is its deduplication technology. Benefits include improved storage efficiency and reduction of bandwidth consumption and backup window, resulting in cost optimization. To implement a NetBackup Deduplication pool to NFS, Veritas has partnered with Pure Storage to enable its deduplication technology for FlashBlade NFS.

Solution Overview

The logical architecture of Veritas NetBackup with Pure Storage arrays is illustrated in Figure 2. You can host the source, an Oracle database, on any storage, but for fast backup and fast restore, it must be hosted on an all-flash array like Pure FlashArray. FlashBlade acts as an NFS storage target for backups conducted by NetBackup.

NetBackup can use FlashBlade as a target for deduplicated or non-deduplicated backups. This document covers using FlashBlade as storage for deduplicated backups. Utilizing FlashBlade as an NFS storage target for deduplicated data is



supported with NetBackup version 8.2 with EEB ET3981134 running on Red Hat Enterprise Linux (RHEL) 7.6 and onward. The maximum size of the media server deduplication pool (MSDP) on a FlashBlade is 100TB. This integration supports NetBackup features such as client-side and server deduplication, Accelerator, and Auto Image Replication (AIR). However, the integrated solution does not currently support instant access and universal shares.

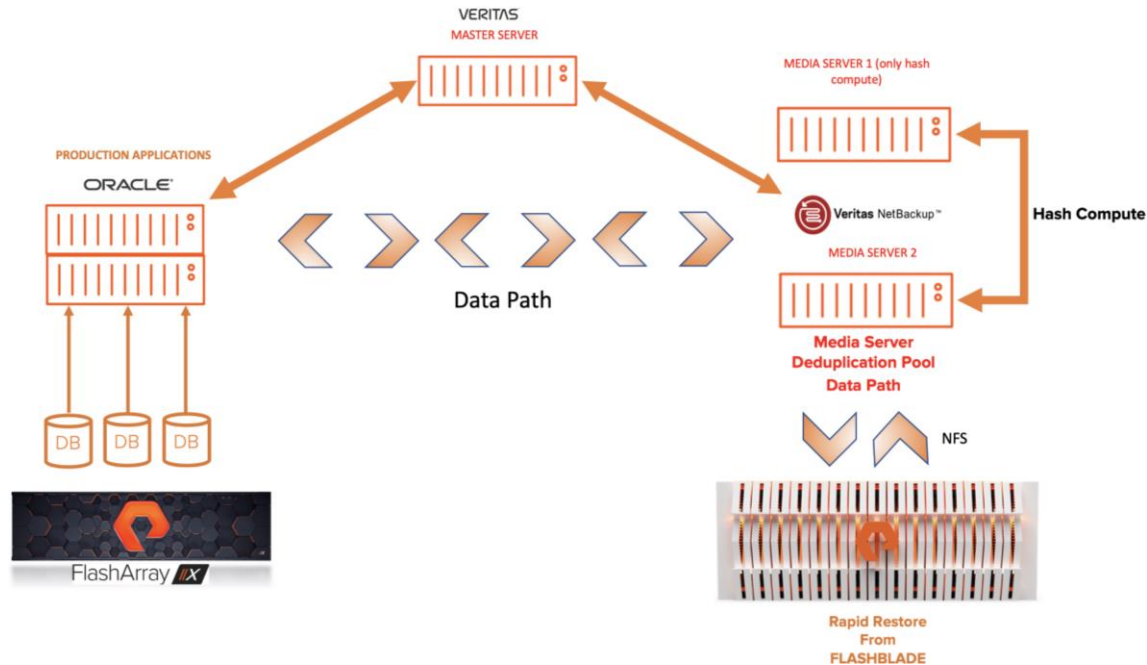


Figure 2. Pure Storage and Veritas architectural overview

Performance Tests

Pure conducted performance tests of backup and restore operations for Oracle database(s) to validate and showcase the performance of FlashBlade as an NFS target for NetBackup deduplicated backups. (Refer to the Appendix for more configuration and environment details.)

Test Strategy

For tests involving backup and recovery of deduplicated data to and from FlashBlade, the Pure team created eight volumes on FlashBlade and exported them to a single media server. The volumes were NFS-mounted on the single media server using eight different mount points and virtual IP addresses. These volumes were presented to NetBackup as a single media server deduplication pool (MSDP). In this test, the team configured another media server to load-balance the computation of the fingerprints. They configured NetBackup with a maximum of 256 MSDP streams and eight maximum parallel streams.

Pure next created several Oracle databases with different 1TB datasets to measure the throughput when backing up and restoring one, two, four, and eight databases. As for network connectivity, each database host had a 10 GbE interface. The media server consisted of four 25GbE network interface cards teamed to create a 100GbE pipe. The FlashBlade had a 40GbE pipe. The backup and restore speed were captured from Oracle's internal view command, `"v$rman_backup_job_details"`, showing the end-to-end performance times and not just the time spent at FlashBlade.



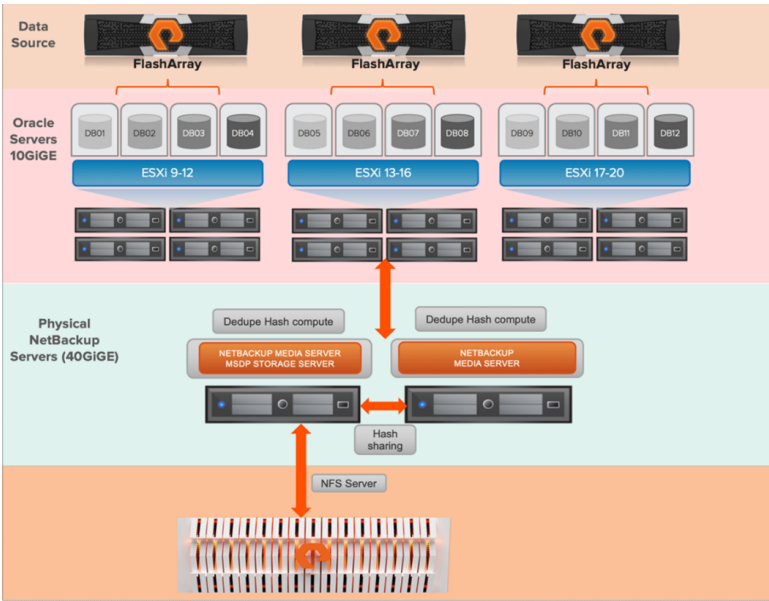


Figure 3. Test environment for NetBackup deduplicated data on FlashBlade

Performance Test Results

The results of these tests (Figure 4) indicate that FlashBlade as an NFS target for deduplicated data from NetBackup scales linearly as more databases are added. For full backups of eight databases, the tests achieved performance throughput of 11.27TB/hour. Throughput with a 5% data change was 15.06TB/hour and performance of restores for the same workload was 18.55TB/hour.

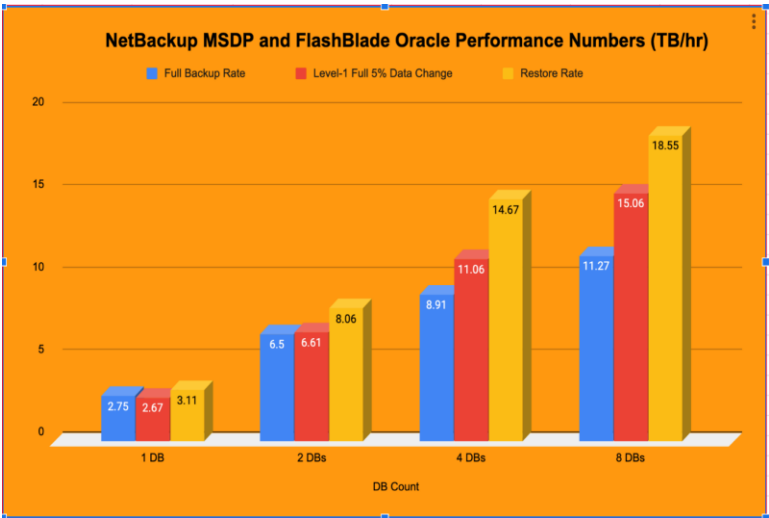


Figure 4. FlashBlade with NetBackup delivers scalable performance.



Data Reduction Test


Pure also conducted a study to illustrate the benefits of compression on deduplication offered by Veritas MSDP with FlashBlade compression while performing full backups of Oracle databases.

Test Strategy

The team performed a series of full backups of eight Oracle databases through Veritas NetBackup. As with the backup tests, Pure performed the level 0 full backup first, captured the data reduction numbers, and then introduced a 5% change into the database. Next, the team performed a level 1 full backup of the database and again captured the data-reduction numbers. We performed the series of level 1 backups five times to simulate five daily full backups with a 5% change introduced daily.

Data Reduction Results

The team observed data-reduction rates as high as 13x with compression enabled on both NetBackup and FlashBlade. The data reduction tests involved backing up eight Oracle databases with a total size of 6.4TB.



8 Oracle Databases total size 6.4TB

Figure 5. Data reduction test methodology and results

Best Practices

The following key recommendations and best practices can help simplify and accelerate planning, architecting, and deploying the integrated solution of FlashBlade with NetBackup.

Provisioning for Deduplicated Data

Pure has observed that performance increases when employing multiple volumes exported to multiple shares with different data virtual IP addresses, especially when combined into one MSDP on a single media server. Implementing multiple volumes allows for additional parallelism and paths to FlashBlade. As a best practice, the minimum volume size should be 10TB when utilizing multiple volumes in a single MSDP. Also, limit the number of filesystems per MSDP to no more than 12. As previously mentioned, 100TB is the maximum size of an MSDP supported for FlashBlade NFS. When deploying MSDP in this fashion, Pure recommends that you create a separate FlashBlade volume for the MSDP catalog (metadata).



If you aggregate multiple volumes to create an MSDP pool, we recommend that each volume grows at the same rate. For example, if you create eight FlashBlade volumes and aim to grow MSDP by 80TB, then you should expand each volume by 10TB. One of the criteria for determining where to send data (on which volume) is the available capacity on that volume. This ensures an even distribution of the data across volumes.

Compression

Both NetBackup and FlashBlade have compression enabled by default. When compression is enabled on both NetBackup and FlashBlade, data consumption on FlashBlade reduces by 30% and keeps TCO low. As a best practice, keep the default settings for compression on NetBackup for deduplicated data.

Encryption

FlashBlade encryption is enabled by default for data at rest. Keys are distributed over all the blades. However, if you require encryption “in-flight” to the storage target, then you need to implement NetBackup as well. NetBackup offers both policy-based encryption and encryption for deduplicated data.

Disaster Recovery

As a best practice, have a disaster recovery plan in case of catastrophic events. NetBackup Auto Image Replication (AIR) provides an effective way to do this. AIR allows you to replicate asynchronously NetBackup catalogs and data from one domain to another NetBackup domain in one or several geographical sites. You can use other ways to protect NetBackup from accidental deletion or corruption, including conducting regular backups of the NetBackup catalog.

NFS Options

The NetBackup deduplication technology does not allow the following nfs options:

```
Disabling client-side attributes caching with options: actimeo=0, acregmin=0, acdirmin=0, acdirmax=0  
Use of the soft option
```

NetBackup deduplication technology may cause issues if used with these options. We recommend that you keep these options at their default settings when mounting.

The following is the recommended nfs mount option to obtain the optimized performance:

```
hard  
rsize=524288,wsize=524288
```

Planning

Utilizing multiple media servers to load balance the fingerprint calculation improves overall throughput. However, add a load-balancing media server only if you observe high CPU utilization on the media server acting as a storage server.

There is a known limitation that the media server can be in a hung state if the NFS mount to FlashBlade is broken. It is highly recommended that you configure a dedicated media server solely for use as an NFS target since the media server may need to be rebooted in some cases.



Sizing

Your performance and capacity requirements will dictate the size of the media servers and FlashBlade. Parameters that may affect these requirements include:

- Size of source data
- Daily change rate of data
- Annual storage growth
- Desired retention
- Deduplication rate
- Conducting weekly, daily, or monthly incremental and/or full backups.

The Veritas and Pure account teams have tools to appropriately size the environment based your desired requirements.

In the context of storage consumption on FlashBlade with NetBackup, there is overhead associated with the backup images such as metadata directory, logs, and history information stored with the backup images. The size for the metadata files is negligible, however. Although with deduplication, additional metadata resides by default on the same storage target as the deduplicated data. As a best practice, it is best to specify a different FlashBlade NFS volume to hold this metadata if utilizing multiple FlashBlade volumes in one MSDP disk pool. There is an option to specify an alternate location for this metadata during the configuration of the MSDP.

Conclusion

FlashBlade as an NFS storage target for NetBackup deduplicated data offers an option for organizations seeking a storage-efficient solution that delivers “rapid restore” of their environments. Sending NetBackup-deduplicated data to FlashBlade over NFS is an exclusive integration. Overall, this integrated solution offers key advantages, including the ability to:

- Achieve superior rapid restore performance with near-linear performance growth from one to eight databases.
- Choose multiple tuning options with FlashBlade and support multiple, concurrent backup and restore operations that take advantage of FlashBlade as a storage platform for NetBackup.
- Gain support for key NetBackup features such as client-side and server deduplication, Accelerator, and AIR.
- Improve storage efficiency when utilizing both NetBackup and FlashBlade compression.



Appendix

The table describes the environment details related to the test studies of NetBackup-deduplicated data with FlashBlade as an NFS storage target when backing up and restoring Oracle databases.

Resource	Specifications	Software
Sources (e.g., Oracle database, VMs, clients)	Each ESXi host: 2 x Intel® Xeon® CPU E5-2697 v2 @ 2.70GHz – 12 cores per socket Memory: 512GB Network: Mellanox MT27500 family network adapter 8 Oracle Database Virtual Machines 16 vCPUs 64 GB RAM 1 x 1GbE/s network adapter for mgmt. 1 x 10GbE/s network adapter for data Each Oracle VM is capable of 1 GB/s to 1.2GB/s	ESXi Version: 6.7 Oracle Version: 12.2.0.1 Oracle Linux 7
BYOS NetBackup 1x NetBackup Master 2x Media Server	Media Servers: Dell R720: 2 x Intel Xeon CPU E5-2697 v2 @ 2.70GHz – 12 cores per socket Memory: 512GB NICS: 4 x 25GbE (teamed) Operating System: RHEL 8.0 NFS options: vers=3, hard, size=524288, wsize=524288 Master Server (VM): 4 vCPUs Memory: 32GB NICS: 1 x 10 GbE	NetBackup 8.2 + EEB Media Server – NBU 8.2 + EEB ET3981134 NetBackup Parameters: MSDP max streams: 256 Max Parallel Streams: 8
Pure Storage FlashBlade	15 blades x 17TB Blades 4 x 40GbE Ethernet (data) 2 x 1GbE redundant Ethernet (management)	Purity//FB 2.3.6
Pure Storage FlashArray	3 x //M70 Connectivity: 4 x 10GbE redundant iSCSI	Purity//FA

Table 1. Hardware and software specifications

Next Steps

- Learn more about the [Pure Storage and Veritas relationship](#) by visiting our web page.

Supporting Information

- [Pure Storage data-protection solutions](#)
- [Pure Storage FlashBlade](#)



About the Author



Mandeep Arora is the Pure Storage Data Protection Solutions Architect responsible for defining data protection solutions partnered with various backup applications. In this role, he defines solutions and reference architecture for Oracle, SQL, and VMware based workloads on the company's products and performance benchmarks. Mandeep has spent over 12 years in the data protection industry and has working knowledge of data protection products designed for businesses of all sizes. He started his career with IBM Tivoli Storage Manager in the core software development and test team, followed by Isilon Systems, where he was responsible for delivering the NAS backup solution to enterprise-class customers. He was also part of the Veritas storage solutions team and was responsible for technical relationships and advising partners on data protection for VMware.

©2020 Pure Storage, the Pure P Logo, and the marks on the Pure Trademark List at <https://www.purestorage.com/legal/productenduserinfo.html> are trademarks of Pure Storage, Inc. Other names are trademarks of their respective owners. Use of Pure Storage Products and Programs are covered by End User Agreements, IP, and other terms, available at: <https://www.purestorage.com/legal/productenduserinfo.html> and <https://www.purestorage.com/patents>

The Pure Storage products and programs described in this documentation are distributed under a license agreement restricting the use, copying, distribution, and decompilation/reverse engineering of the products. No part of this documentation may be reproduced in any form by any means without prior written authorization from Pure Storage, Inc. and its licensors, if any. Pure Storage may make improvements and/or changes in the Pure Storage products and/or the programs described in this documentation at any time without notice.

THIS DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID. PURE STORAGE SHALL NOT BE LIABLE FOR INCIDENTAL OR CONSEQUENTIAL DAMAGES IN CONNECTION WITH THE FURNISHING, PERFORMANCE, OR USE OF THIS DOCUMENTATION. THE INFORMATION CONTAINED IN THIS DOCUMENTATION IS SUBJECT TO CHANGE WITHOUT NOTICE.

Pure Storage, Inc.
650 Castro Street, #400
Mountain View, CA 94041

purestorage.com

800.379.PURE

