

TECHNICAL WHITE PAPER

# **Simplified Oracle Rapid Restore Using Pure Storage FlashRecover, Powered by Cohesity**

An integrated, all-flash data protection solution for rapid recovery at scale.

# Contents

**Introduction .....3**

**Business Environments .....3**

**Pure FlashRecover, Powered by Cohesity Overview.....4**

**Business Value .....4**

    Performance ..... 5

    Simple..... 5

    Scalable ..... 5

**Why Disaggregated Storage and Compute Is Important for Data Protection .....5**

    Unified Fast File and Object Platform ..... 6

**Use Case: Using Pure FlashRecover to Protect Oracle Databases .....7**

**Solutions Architecture Overview .....7**

**System Configuration .....8**

**FlashArray Configuration.....8**

    Pure FlashBlade Configuration ..... 10

**FlashRecover Cluster ..... 10**

    Oracle Database Configuration ..... 10

**Test Cases and Results ..... 12**

**Backup Tests ..... 12**

    Objectives .....12

    Test Procedure.....12

    Test Results .....12

**Restore Tests..... 13**

    Objectives ..... 13

    Test Procedure..... 13

    Test Results ..... 13

**Data Reduction Tests ..... 13**

    Test Procedure..... 13

    Data Reduction Results ..... 14

**Conclusion ..... 14**

**About the Author ..... 16**



## Introduction

Enterprises of every size are at an IT crossroads. Today's business world is more dependent on uptime of systems, applications, and data availability than ever before. When there's an outage, the entire business can stop. IT has to address these significant challenges around managing business-critical data and systems to ensure that the business can operate while balancing data protection, access, retention, recovery, and mobility. Database administrators and IT management personnel are faced with ensuring the availability of mission-critical business applications daily, for any number of stakeholders within the organization. Line-of-business owners, knowledge workers, developers, and even the CxO all need timely access to Oracle databases 24x7x365.

When an IT organization needs to protect Oracle databases, they need a solution that provides operational recovery with very low recovery time objectives (RTOs), and it must also address multiple recovery point objectives (RPOs). While there's no shortage of tools and solutions that offer protection for Oracle databases, many don't focus on the most important aspect: the recovery. When business-critical databases and data are unavailable, every second counts. So having a solution that can't meet the recovery goals of the business creates a liability.

This white paper focuses on protecting Oracle environments with Pure FlashRecover™, Powered by Cohesity®. You will be taken through the performance testing of FlashRecover in the backup and recovery of Oracle databases.

---

## Business Environments

Data continues to grow at an unprecedented rate, and the rise in the value and volume of data has coincided with the increased adoption of all-flash storage systems in enterprises worldwide. Yet many backup environments remain unchanged, unable to meet new demanding recovery requirements, cloud service demands, and the data reusability needs of modern use cases such as DevOps and analytics. The need to ensure flash-enabled, fast, reliable, and simple backup and recovery of data—at scale—is greater than ever before. Increased occurrences of ransomware attacks are a constant reminder that this is becoming an even more critical problem for organizations of all sizes.

Technology leaders are looking for better answers to the challenges associated with:

- The growing recovery demands of modern applications and data
- The move toward cloud services
- The heightened concerns over ransomware
- The need to address their infrastructure complexities

Pure Storage® and Cohesity forged a [strategic partnership](#) and brought to market a first-of-its-kind solution to address all of these challenges. [Pure FlashRecover, Powered by Cohesity](#) is an integrated, modern, all-flash data protection solution for rapid



recovery at scale. The jointly developed solution offers simple, fast, and reliable backup and recovery for enterprise data that also allows organizations to do more with their backup data, restoring petabytes of data in hours, not days or weeks.

## Pure FlashRecover, Powered by Cohesity Overview

All-flash storage dominates storage spend in enterprise environments, but the market wasn't always this way. A decade ago, Pure Storage disrupted the market by bringing wide-scale flash storage to the enterprise storage market, a disruption that forever changed customer expectations for their primary storage. Current business demands from our customers have called for another change, this time in the backup industry.

While backup vendors like Cohesity have disrupted the backup market, the demanding recovery requirements have meant that the purpose-built backup appliance (PBBA) market is impacted to an extent like ever before. Backup and storage markets are adjacent in many ways, and they share many similar customer pain points.

Pure FlashRecover, Powered by Cohesity is an integrated, modern, all-flash data protection solution for rapid recovery at scale. The jointly developed solution is simple, fast, and reliable, and it provides independent scaling of storage and compute for backup and recovery of enterprise data. With storage innovation from Pure Storage and modern backup from Cohesity, the integrated solution delivers:

- **High performance:** Up to three times faster backup and restore throughput than disk-based alternatives
- **Simplified management:** Ease of management, auto-discovery, and configuration of FlashBlade®, non-disruptive upgrades, native cloud integration for flash-to-flash-to-cloud backup, recovery, and archive
- **Scalable and efficient data reuse:** Disaggregated compute and storage for independent scaling of backup and recovery processes, reuse of backup data on FlashBlade for modern applications
- **Streamlined support and procurement for modern data protection:** Solution available from Pure Storage for single-point purchasing, deployment, and support, eliminating the need for organizations to engage two vendors
- **Recovery at scale:** Restore and recover petabytes of data in mere hours

The Pure FlashRecover, Powered by Cohesity solution has three components:

- **Cohesity DataProtect software:** Cohesity DataProtect is simple, comprehensive, enterprise-grade backup and recovery software for your traditional and modern data sources.
- **Cohesity-certified compute nodes:** Cohesity DataProtect runs on Cohesity-certified compute (PXG1) nodes for Pure FlashRecover with no local drives.
- **Pure FlashBlade:** The storage is provided by Pure FlashBlade, and the file systems are mounted to FlashRecover compute nodes via NFS v3.

## Business Value

Pure Storage's FlashRecover introduces tremendous value to enterprises in three key areas: Performance, simplicity, and scalability.



## Performance

There's value in finishing backups faster. Backups offer protection services to production applications and data. With data and application sets growing rapidly, creating backups in a timely fashion has become more difficult. By being able to back up more often and write quicker, a business can meet growing demand while meeting the most stringent recovery point objectives (RPOs).

While an "incremental forever" approach to backups can shrink backup windows, during recoveries there are no such optimizations. Restore performance is often dictated by how fast you can read from your backup target and how efficiently your backup target file system is designed.

Flash eliminates defragmentation-related concerns and inherently brings performance to restores. A high-throughput FlashBlade system offers massive recovery capabilities, allowing recovery of thousands of virtual machines or petabyte-scale environments.

## Simple

Backups have always been considered complex solutions. Cohesity has introduced a variety of innovations, including non-disruptive upgrades to backups. When combined with the non-disruptive upgrades for FlashBlade, the solution becomes simple and scalable. Pure FlashRecover will auto-discover and configure the storage required for backups in a simple wizard-driven approach.

You can now store data on FlashBlade for rapid recoveries, ransomware protection, and reuse of data, and seamlessly integrate with the cloud for longer term retention or for reuse of data in the cloud.

## Scalable

Backups are storage-intensive. By providing a disaggregated, hyperconverged model for backups, you can scale based on need. If additional storage is required, you can just add additional blades to the FlashBlade system and grow the environment. Independent scaling allows efficient use of resources.

You don't have to pre-plan their backup purchases, which the backup admins typically have no control of. They can grow their environment based on the changing business needs. By leveraging the unified fast file and object (UFFO) storage of FlashBlade, customers have the benefit of consolidating all their modern data needs.

## Why Disaggregated Storage and Compute Is Important for Data Protection

One of the unique design aspects of FlashRecover is its disaggregated storage and compute architecture. With this architectural approach, organizations can scale storage and compute separately. Why is this so important for data protection?

One of data protection's biggest challenges is the ability to scale out. Primary storage continues to grow at an unprecedented rate, and this inevitably impacts your data protection strategy, requiring it to keep pace with the needs of your primary storage. Traditional data protection solutions deliver this by adding a backup server and storage or another backup appliance. This will require not only deployment, but also calibrating backup and disaster recovery (DR) policies. The other factor to be considered is that backup deduplication only works at the backup server level, so careful policy management is needed to maintain levels of deduplication efficiency.



Contrast this with the approach that is allowed due to disaggregation simply add more storage. No backup policy maintenance is required; deduplication continues to be performed across all backed up data. A disaggregated solution delivers other advantages. If an organization wants to increase backup or recovery performance, they can simply add additional compute nodes. In a more traditional data protection model, this would most likely require a significant redesign. In the disaggregated model, it's simply a case of adding more compute nodes.

The disaggregated approach preserves a single name space, eliminates islands of deduplication, and simplifies the process of scaling a data protection solution—both in respect to adding more storage and more performance.

## Unified Fast File and Object Platform

Organizations today need a storage solution that addresses modern data requirements, provides simplicity and multi-dimensional performance, and enables consolidation of key unstructured data workloads. This will eliminate storage silos and provide investment protection for organizations at any point in their digital transformation journey.

Even though the term may be new to many, unified fast file and object (UFFO) storage is quickly becoming the only category of storage that can address these challenges. Specifically, UFFO storage meets the following requirements for modern digital transformation data:

- **Multi-dimensional performance:** Extremely high throughput and IOPS with low latency to support multiple workloads simultaneously, including those with small or large files, sequential or random I/O access, batched or real-time jobs, and large numbers of files.
- **Intelligent architecture:** The storage system is built for flash from the ground up to truly leverage the performance and efficiencies of flash. It is also simple to deploy, manage, and upgrade without requiring constant tuning. A modern storage solution must be simple enough so that the operations don't overwhelm storage admins with the mundane tasks of managing networking complexities when deploying the system, volumes, cluster pairs, aggregates, and flash caches or configuring replication.
- **Cloud-ready:** A cloud-like agility, flexibility, and consumption choices with on-premises control.
- **Always available:** The capability of extending past traditional platform resiliency. Maintenance operations, software upgrades, and capacity expansions are completed without disruption. The foundational software design makes it possible for the solution to deliver high availability over multiple years and upgrade scenarios.
- **Dynamic scalability:** The ability to seamlessly scale not only capacity but also performance, metadata, number of files and objects, and more.
- **Multi-protocol support:** A single platform provides native file and native object protocol support without compromising performance or any functionality.

FlashBlade, the storage backbone of FlashRecover, utilizes the industry-leading Unified Fast File and Object platform, accelerating unstructured data growth, evolving modern data and modern applications, the need for better, faster, real-time insights, the rise of fast objects and the transition to a cloud experience. Pure's UFFO technology helps shift inefficient, sprawling storage infrastructure with data silos to a single unified platform that delivers simplicity, performance, and consolidation.



Multi-dimensional performance delivers extremely high throughput and high performance to support multiple workloads simultaneously including workloads with any file size (small or large), sequential or random I/O, batched or real-time jobs, and a large number of files.

With Pure FlashRecover, Powered by Cohesity, IT leaders can effectively address their top-of-mind concerns around rapid recovery, ransomware protection, and the reuse of data.

## Use Case: Using Pure FlashRecover to Protect Oracle Databases

Pure FlashRecover achieves rapid restore at scale, defends all your backup data wherever it resides, and gets the most value from your data. This integrated solution delivers high-performance backup and restore, simplified management, and scalable and efficient data reuse.

Most environments today rely on virtual machines to run their applications and workloads. To illustrate the simplicity and performance of FlashRecover, we have outlined the setup of FlashRecover and tested the performance of FlashRecover in protecting and recovering virtual machines in subsequent sections.

Setup and creation of policies has been simplified as illustrated below:

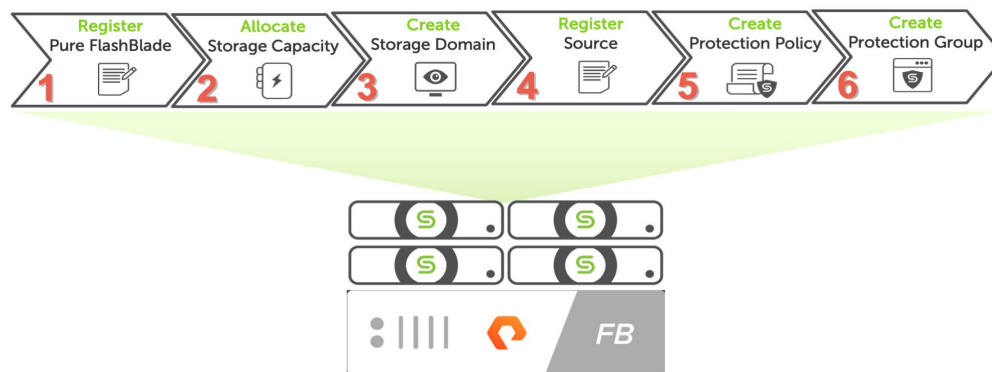


Figure 1. Six steps to protect Oracle databases with Pure FlashRecover, Powered by Cohesity.

## Solutions Architecture Overview

The source application Oracle databases can be hosted on any storage, but for fast backup and fast restore it is imperative that it is hosted on an all-flash array like Pure Storage FlashArray™.

For this test, four Oracle database servers are hosted across three FlashArray systems (with 1-1-2 databases per FlashArray system), that are connected to four ESX hosts over iSCSI. Datastores are created on the three FlashArray systems, and virtual machine disks (VMDKs) are created for every Oracle database server to host the database. In this white paper, the test will only illustrate the performance measurement using four Oracle databases. Datastores are created from the three FlashArray systems, and the VMDKs are created for every Oracle database server to host the database. Pure FlashRecover, Powered by Cohesity architecture is shown in the following figure.

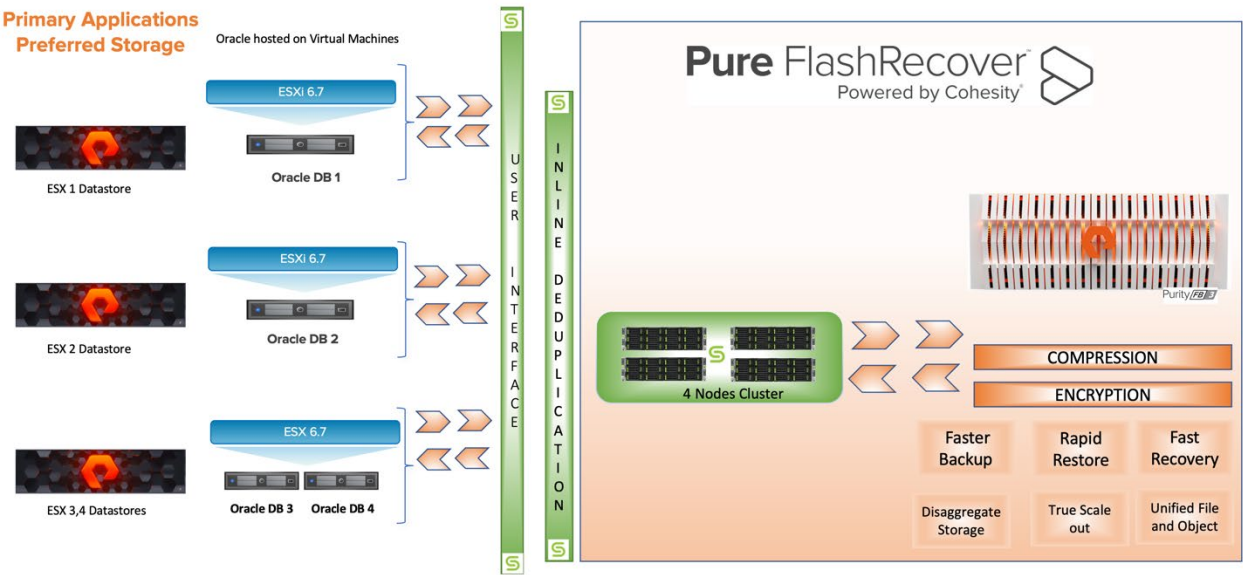


Figure 2. Every ESX server is equipped with a 10Gb/s network interface, which enables approximately 1Gb/d of bandwidth between the database servers, FlashRecover cluster, and FlashBlade. Hence, running a single database backup/restore on this environment can only yield a maximum of 1Gb/s of bandwidth. To demonstrate higher and scalable bandwidth, multiple database backups and restores across ESX servers are performed.

System Configuration

The system configuration details are presented below:

Compute	Network	Oracle VM Configuration
Each server includes: <ul style="list-style-type: none"><li>Two Intel-Xeon E5-2697 v2 @ 2.70 GHz (24 cores total)</li><li>512GB RAM</li></ul>	Mellanox MT27500 family network adapter	16 vCPUs 64GB RAM One 1Gb/s network adapter for management One 10Gb/s network adapter for data Five virtual disks (VMDK) from two datastores

Table 1. System configuration.

FlashArray Configuration

The following table shows the configuration of the three Pure FlashArray systems that are used for hosting the source Oracle databases as well as the Veritas NetBackup software.





COMPONENT	DESCRIPTION
FlashArray	3 X //M70
Capacity	34.90TB raw (base chassis) 20.94TB usable (before data reduction)
Connectivity	4 X 10Gb/s redundant iSCSI 1Gb/s redundant Ethernet (management port)
Physical	3U 5.12" X 18.94" X 29.72" FlashArray//M chassis
Software	Purity//FA 5.1.4

Table 2. FlashArray configuration.

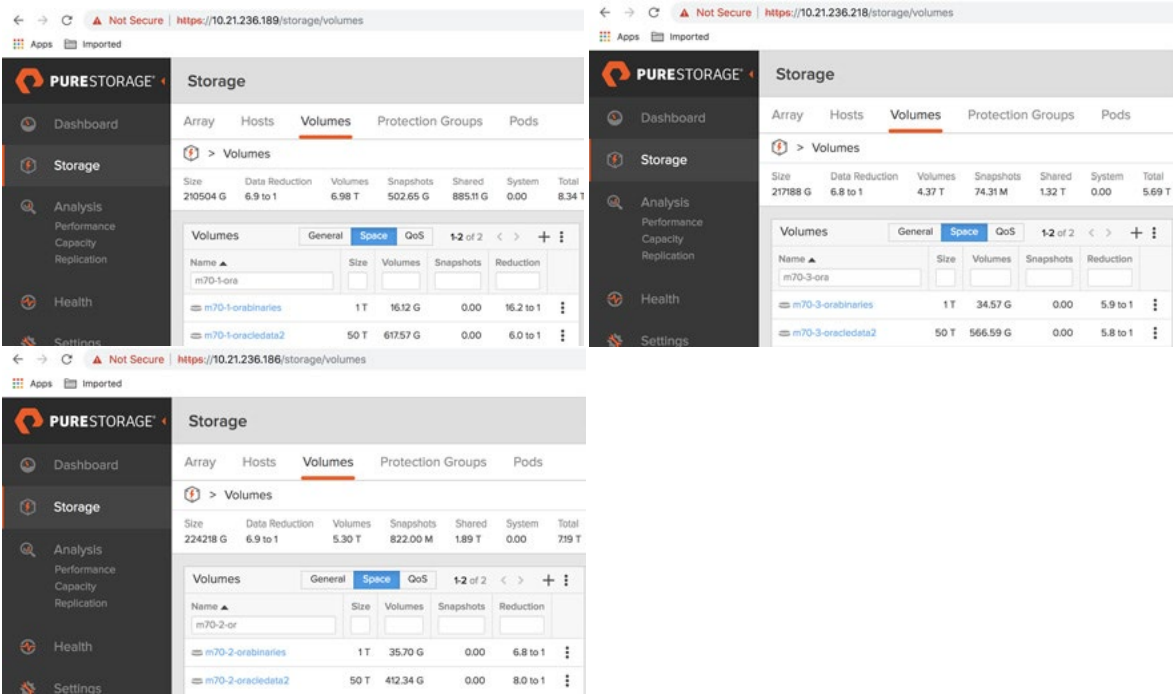


Figure 3. Two volumes per FlashArray are set up, one to host the root volume, Oracle binaries, and automatic storage management (ASM), and another to host the Oracle data.



ora-m701-db01 - View Hard Disks		
Hard disk 1	150 GB	O.S
Hard disk 2	250 GB	Oracle Binaries
Capacity		
VM storage policy		
Type		Thin provision
Location		m70-1-orabinaries (438 GB free)
Hard disk 3	50 GB	ASM
Hard disk 4	1 TB	Oracle data
Capacity		
VM storage policy		
Type		Thick provision eager zeroed
Location		m70-1-oracledata2 (44.8 TB free)
Hard disk 5	300 GB	Oracle FRA

Figure 4. Disk layout of the Oracle virtual machines.

Pure FlashBlade Configuration

Pure FlashBlade acts as a storage target in FlashRecover using the NFSv3 protocol to receive the backup. The FlashBlade system used in testing is configured as shown in Table 3.

COMPONENT	DESCRIPTION
FlashBlade	15 X 17TB blades
Capacity	240TB raw 162.46TB usable (with no data reduction)
Connectivity	4 X 40Gb/s Ethernet (data) 2 X 1Gb/s redundant Ethernet (management port)
Physical	4U
Software	Purity///FB 3.0+

Table 3. FlashBlade configuration.

FlashRecover Cluster

Oracle Database Configuration

Oracle 12c Release 2 (12.2) was installed on all database hosts, which were running Oracle Linux 7.4. A standalone database was created on ASM on all twelve servers. The ASM metadata and ASM instance were hosted on one of the virtual disks labeled with the ASM filter driver (AFD), while the Oracle data and FRA were hosted on two other virtual disks.



Component	Description
<b>RDBMS Version</b>	Oracle 12.2.0.1
<b>File system</b>	ASM
<b>Oracle Components</b>	Oracle Restart, ASM
<b>Oracle User/Group</b>	Oracle: oinstall (DBA, ASMDBA)
<b>Grid User/Group</b>	Grid:oinstall (ASMDBA, DBA, ASMADMIN)

Table 4 Oracle RDBMS configuration.

Disk group	Size	Redundancy	Allocation unit	Description
<b>ASM</b>	50GB	External	4MB	Small disk group for ASM metadata
<b>Data</b>	1TB	External	4MB	Data disk group for OLTP
<b>FRA</b>	300GB	External	4MB	Flash recovery area

Table 5. Oracle ASM configuration.

Grp Name	Path	Disk Name	Sector Size	Block Size	AU STATE	Group Total GB	Disk Total GB	Group Free GB	Disk Free GB
ASM	AFD:ASM1	ASM1	512	4096	4 MOUNTED	50	50	50	50
DATA	/dev/dg_oradata	DATA_0000	512	4096	4 MOUNTED	1,024	1,024	110	110
FRA	/dev/dg_orafra	FRA_0000	512	4096	4 MOUNTED	300	300	298	298

Figure 5. ASM details of an Oracle database.

ASM is set up under a **grid** user and the Oracle database is set up under an **oracle** user. The database was 905GB and populated with OLTP-type data using a modified SLOB benchmark to ingest more uniqueness to achieve 4:1 data reduction to mimic the real-world database data reduction most of our FlashArray customers experience. The database was populated with 816GB of data.

```
SQL> @showspace

Space Details
-----
Allocated Space :      905.01 GB
Used Space      :      816.46 GB
```

Figure 6. ASM details.



## Test Cases and Results

In the following sections you will find details on the three use cases that we tested:

- Backup tests
- Restore tests
- Data reduction tests

## Backup Tests

### Objectives

This test case illustrates full backups of Oracle database(s) through the Pure FlashRecover, Powered by Cohesity UI to demonstrate the backup rate.

### Test Procedure

A series of full backups of the Oracle databases were performed through the FlashRecover UI by varying the number of databases between 1, 2, and 4 with 16 channels. To simulate the real-world scenario, after the first level 0 backup, 5% of change was introduced in the database prior to the next level 1 full backup. The backup speed was captured from Oracle's internal view and ,hence, shows the end-to-end backup timing.

### Test Results

Based on our environment setup, where each database host is limited to 10Gb/s, single database backup achieved an average of 811Mb/s with 16 channels. The number of channels can be increased to get higher throughput on a single database backup. At present, there are some limitations on performance with a four nodes FlashRecover cluster when backing up all unique data, which will be addressed in a future release.




Figure 7. FlashRecover backup performance of an Oracle database.



## Restore Tests

### Objectives

This test case illustrates a full restore of Oracle database(s) through FlashRecover to demonstrate the rapid restore capability.

### Test Procedure

A series of full restores of the Oracle databases were performed through the FlashRecover UI by varying the number of databases between 1, 2, and 4. The restore speed was captured from Oracle's internal view and, hence, shows the end-to-end restore timing.

### Test Results

The restore tests show the biggest advantage of FlashRecover: rapid restore.




Figure 8. FlashRecover restore performance of an Oracle database.

## Data Reduction Tests

This test case illustrates the benefits of storage reduction, with inline deduplication offered by FlashRecover, with compression and encryption offered on the FlashBlade storage level.

### Test Procedure

In this test, a series of full backups of four Oracle databases was performed through the FlashRecover UI. We performed the level zero backup first, captured the data reduction numbers, and then introduced 5% change into the database. We performed level 1 full backup of the databases, and then again captured the data reduction numbers. We performed the series of full level 1 backups five times to simulate five days of daily backups with 5% change.



## Data Reduction Results

The test data-reduction rates were as high as fifteen times with inline deduplication enabled on the FlashRecover cluster and compression enabled by default on Pure FlashBlade storage.

### Data Reduction Test

Pure FlashRecover<sup>™</sup>  
Powered by Cohesity<sup>™</sup>




Figure 9. Data reduction test.

## Conclusion

Pure FlashRecover, Powered by Cohesity delivers a modern data protection experience, combining the best in flash-based unified fast file and object (UFFO) storage with an industry-leading data protection application from Cohesity, giving mid-sized and enterprise customers alike a highly performant solution designed for fast recoveries at scale.

FlashBlade as an NFS storage target for FlashRecover deduplicated data offers an option for organizations seeking a storage-efficient solution that delivers “rapid restore” at scale of virtual machines in their environment.

FlashRecover is an exclusive integration of the FlashBlade product from Pure Storage with Cohesity’s DataProtect products, combined to deliver the power and ease-of-use of Cohesity’s data protection software, with the highly performant, scalable flash storage backbone.

By leveraging the key strengths of the FlashRecover data platform to consolidate secondary storage workflows, this architecture offers the industry’s top compute platform and disaggregated unified flash storage with integrated backup and recovery software.



FlashRecover can deliver:

- Up to three times faster restore throughput than disk-based alternatives
- Improved storage efficiency when utilizing FlashRecover, Powered by Cohesity inline deduplication and FlashBlade compression
- Disaggregated compute and storage for independent scaling of backup and recovery processes
- Reuse of backup data on FlashBlade for modern apps
- Most simplified form of Oracle data protection



## About the Author



Mandeep Arora is the Pure Storage data protection solutions architect responsible for defining data protection solutions. He serves as the lead solutions architect for the FlashRecover data protection solution, creating reference architectures for primary workloads such as Oracle, SQL, and VMware.

Mandeep has spent more than 14 years of his career in the data protection industry, working with various data protection products meant for small, medium, and large enterprise businesses. He started his data protection career with IBM Tivoli Storage Manager in the core software development and test team, followed by Isilon Systems, where he was responsible for delivering NAS backup solutions for enterprise-class customers. He was also a part of the Veritas storage solutions team, where he was responsible for technical relationships and advised partners in data protection for VMware.

©2021 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](https://purestorage.com)

800.379.PURE

