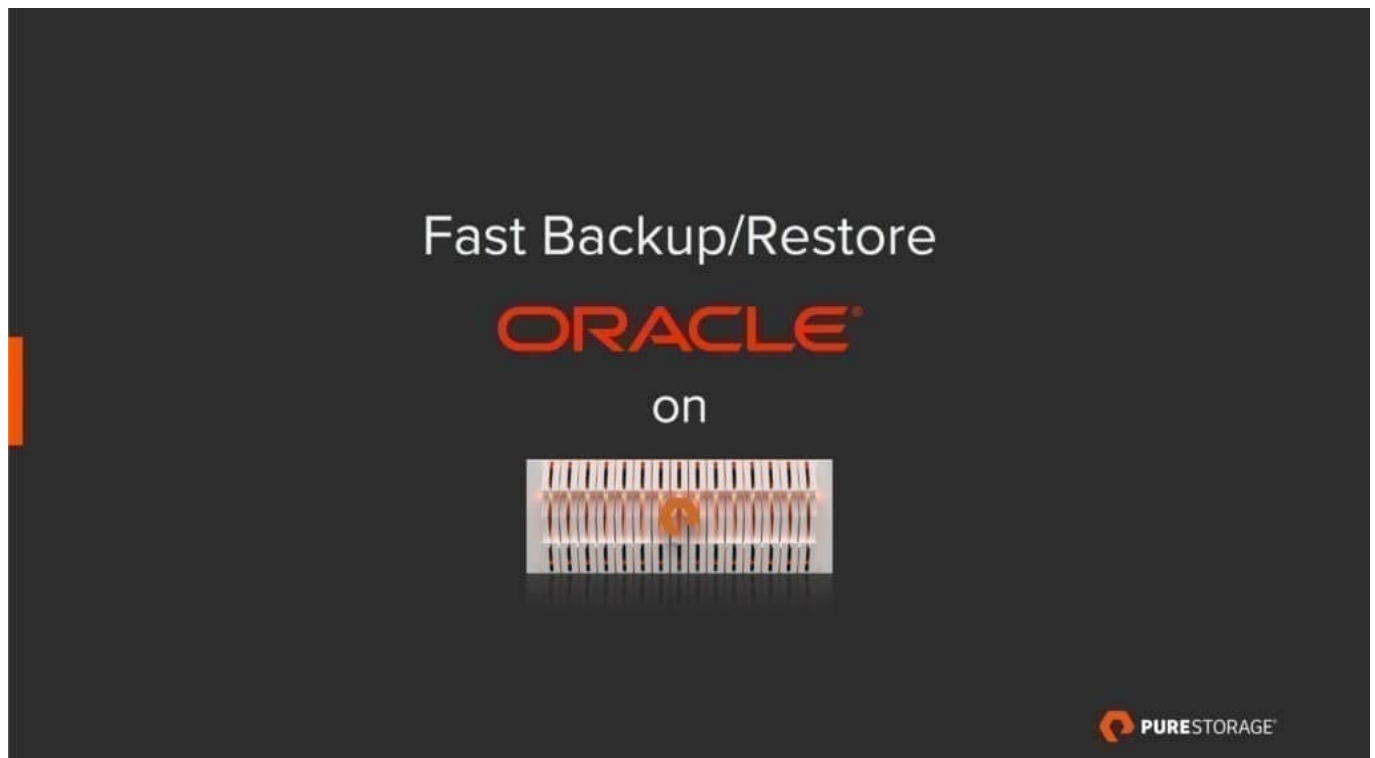


# Accelerated RMAN recipes for Oracle on FlashBlade



## Data Protection Challenges

The growth of data in recent years has been astounding. Modernization of database platforms and analytics along with big data has made data the primary asset of enterprises. While companies are modernizing their primary storage by adopting all flash systems like Pure FlashArray, their backup and restore systems are still using traditional disk based solutions. In conjunction with the data growth, customers are now challenged with daily backups taking over 24 hours breaching their service level agreements.

To address the speed of data protection, storage and backup vendors have come up with purpose-built appliances which may have accelerated the backup times but did not meet the recovery time objective (RTO) during restore as the data rehydration process during restore operations generates large random I/O access patterns on disk drives, resulting in poor performance.

In addition, the customers also face challenges keeping up with demands in their test and development environment, provisioning clones of production databases quickly for DBAs to iterate on. While Pure FlashArray's FlashRecover snapshot functionality allows customers to clone databases in seconds at no additional storage, many customers wanted to segregate their non-production environments from production and hence end up deploying additional infrastructure, including storage to support test and development environment. Ideally, it would make perfect sense to take advantage of the Oracle backups

to create the database clones to be used by testers and developers.

The speed at which an enterprise can restore data in the event of a failure, or quickly provision test/dev environment, defines its business edge and also defines a new class of solution area, called “Fast Backup/Restore”. This solution also puts the dormant backups to use by speeding up database clones.

## The Solution

With the introduction of FlashBlade™, a ground-breaking scale-out flash storage system from Pure Storage, the backup and restore challenges of enterprise customers can be addressed easily. Oracle customers can now direct their RMAN backups to FlashBlade and use dNFS (direct NFS) to accelerate backups significantly. The current generation of FlashBlade system can support a write rate of 4.5 GBps and a read rate of 16 GBps in a single 4U chassis.

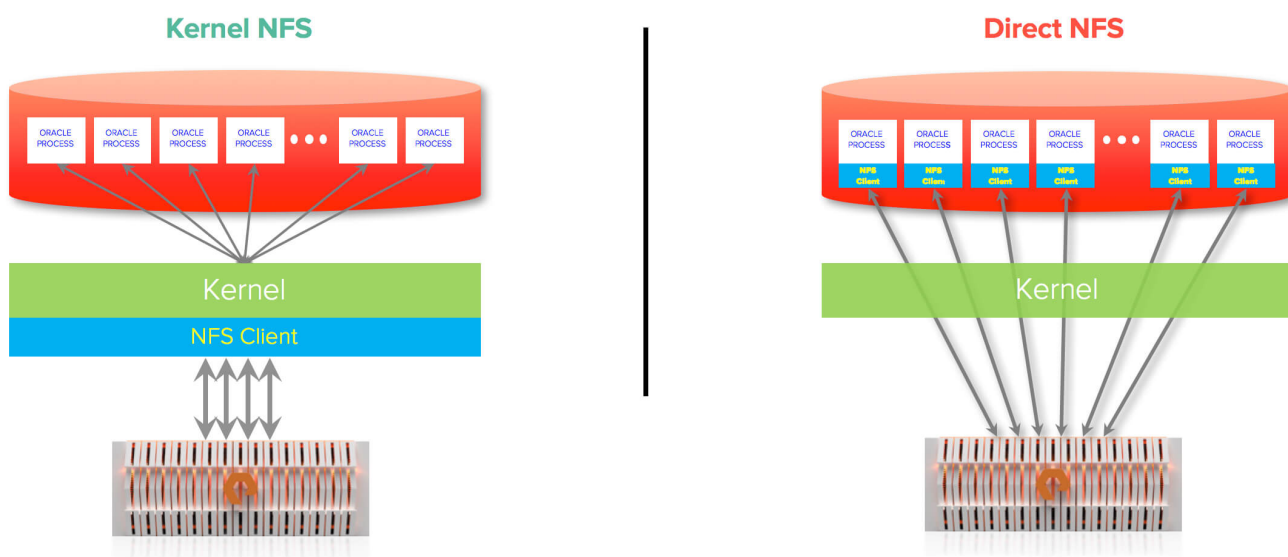
The sustained write rate of 4.5 GBps is equivalent to 15 TB/hour of backup rate. Meanwhile the restore rates from FlashBlade can be as high as 45 TB/hour, often limited by the storage system where the backup is restored onto.

The high bandwidth capabilities of FlashBlade along with RMAN feature of DUPLICATE can be used to clone databases very quickly from the periodic backups. This not only allows the customer to take advantage of the dormant backups but also helps validating their restore/recovery process.

## Why Direct NFS (dNFS)?

The scale-out FlashBlade system is file/object based and it supports applications using NFS, S3/object, SMB, and HTTP protocols. In our use case for Oracle RMAN backups and restore, the filesystems from FlashBlade would be mounted on the host using NFS protocol.

The challenge with the Kernel NFS is that it allows only a single connection to storage for every mount from the host. As FlashBlade’s massively distributed architecture performs best with parallel connections to the blades, dNFS is highly recommended as dNFS makes separate connections to the storage system for every server process.



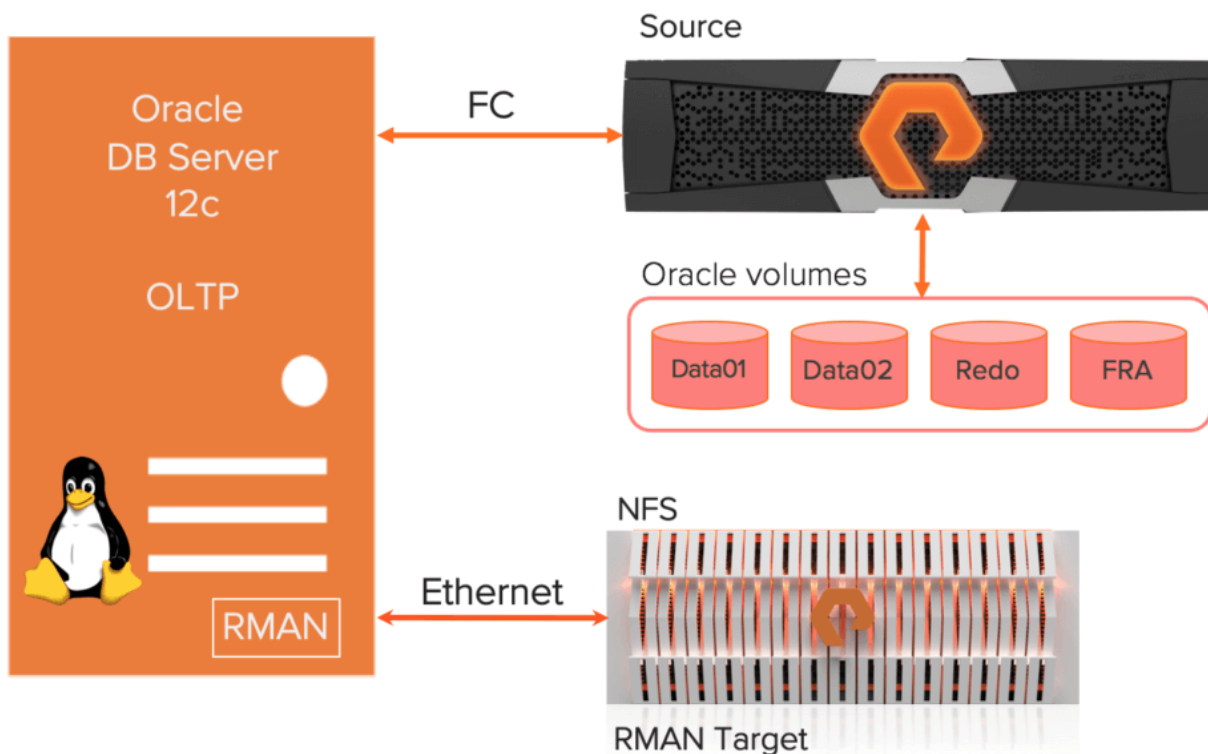
In addition, Direct NFS is capable of performing concurrent direct I/O by bypassing Operating System level

caches. It also performs asynchronous I/O, which allows processing to continue while the I/O request is submitted and processed. Other key feature of Direct NFS Client is the high availability. Direct NFS delivers optimized performance by automatically load balancing requests across all specified paths (up to 4 parallel network paths). If one network path fails, then Direct NFS Client will reissue I/O commands over any remaining paths, ensuring fault tolerance and high availability.

## Fast Backup/Restore - Setup

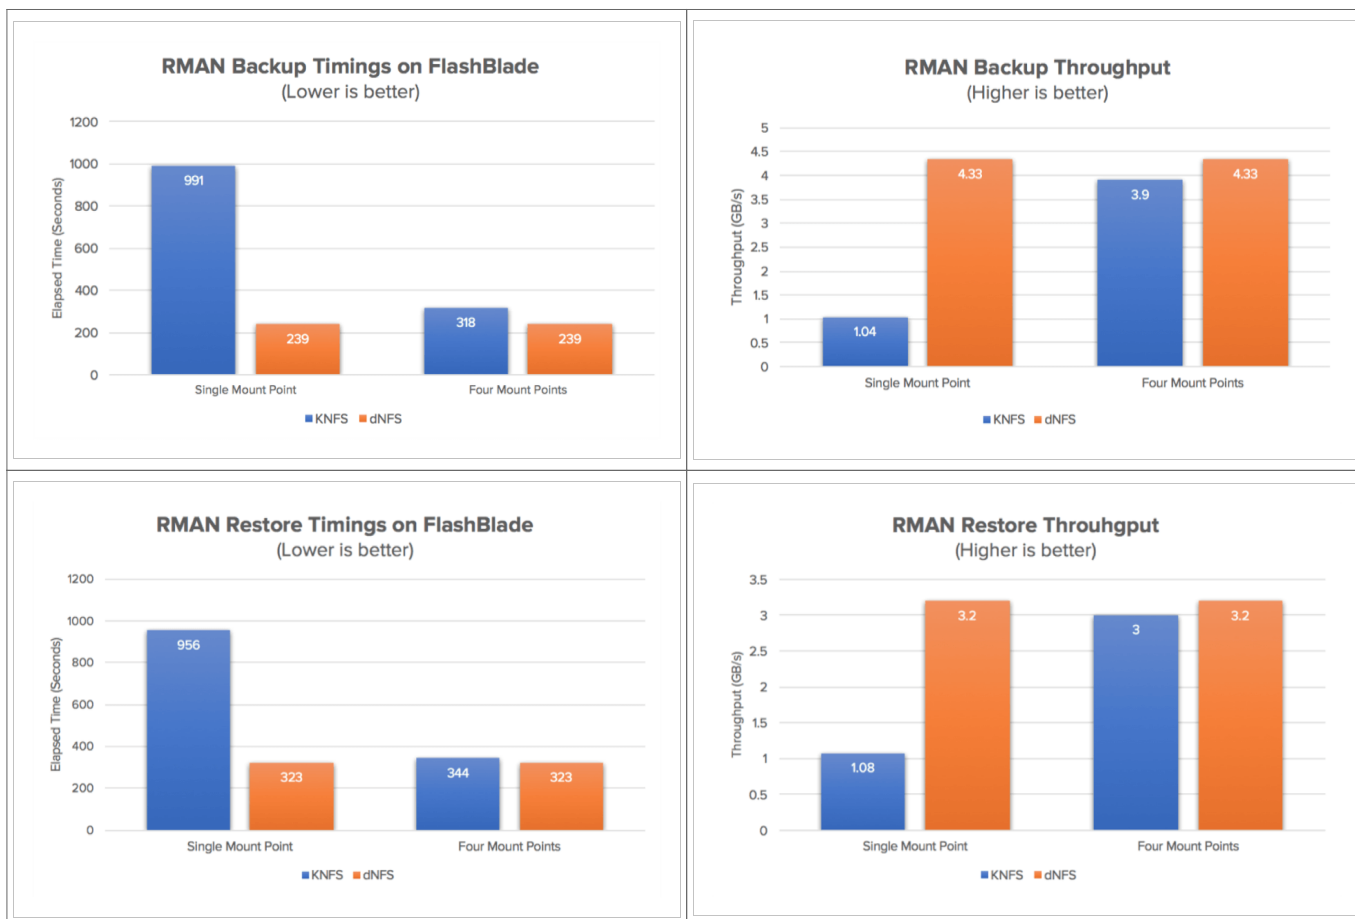
The setup is very simple.

1. Setup multiple subnets on FlashBlade (preferably four as Oracle dNFS can only support up to four network paths to the storage system).
2. Setup multiple network interfaces on the host and align it to the subnets on FlashBlade.
3. Setup filesystems (or volumes) on FlashBlade and mount them onto the database host over NFS protocol. This will be the target for the RMAN backups.
4. Enable dNFS on the database. Even though the source database is on a block based storage, when RMAN is invoked, any writes to FlashBlade will be managed through dNFS



## Fast Backup/Restore - Test Results

To illustrate the advantages of FlashBlade in supporting “Fast Backup/Restore”, we tested Oracle RMAN backup and recovery scenarios with and without dNFS. We performed four backup tests and four recovery tests. These tests include single vs multiple mount points across Kernel NFS and Direct NFS. The database was 1.01 TB in size.

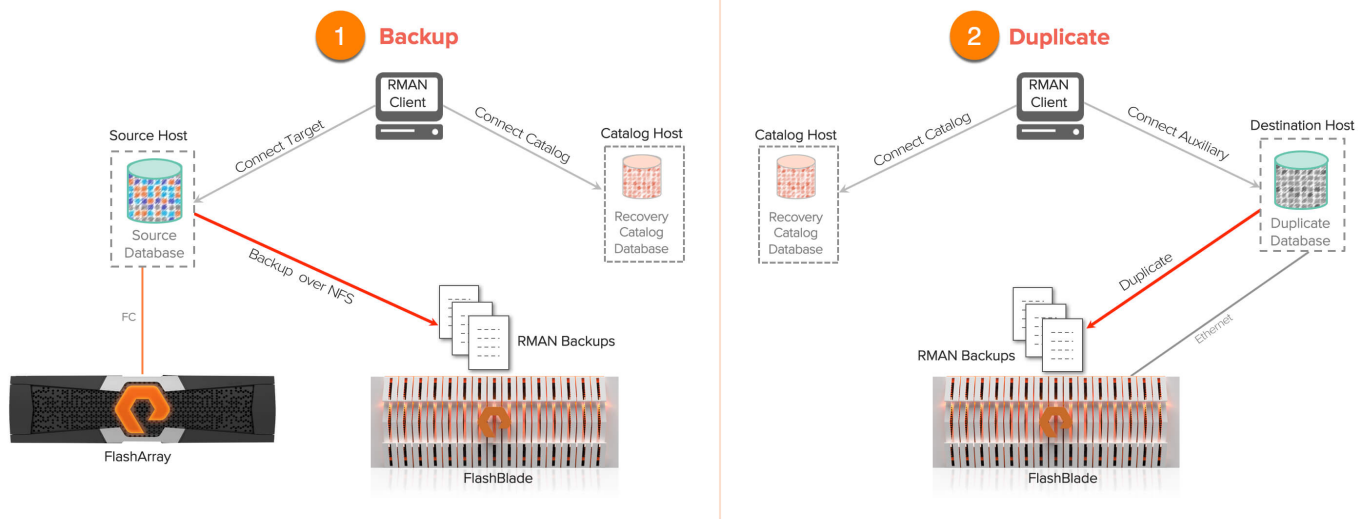


From these results, clearly dNFS with FlashBlade delivers higher throughput along with the benefits of high availability with multipathing like behavior. In the restore scenario, the read throughput of 3.2GBps is still impressive but it is limited by the target FlashArray's write performance. With more FlashArrays, FlashBlade can deliver higher performance. For more details on Fast Backup/Restore with RMAN, see the [White Paper](#).

## Accelerated RMAN Duplicate for Test/Dev

Even though RMAN allows active-database duplication, our requirement is to take advantage of the backup that is available on the FlashBlade and hence our recommendation is to use backup-based duplication, preferably with recovery catalog database which you might be already using.

The duplicate or clone process can be performed on the same server or onto a remote server. For the remote server, the key requirement is to make the backups available. This is very easy to accomplish with FlashBlade by mounting the backup filesystems on the remote host with the same mount point location as that of the source.

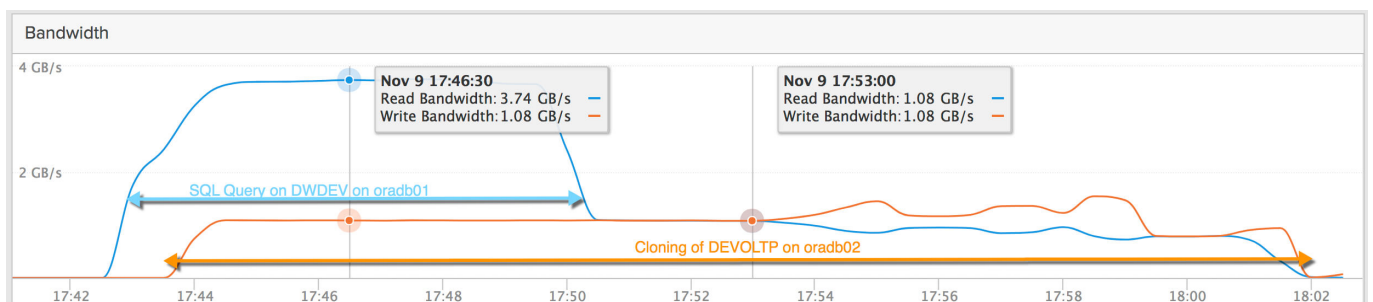


As you can see from the above diagram, the clone or duplicate process happens on the FlashBlade, meaning the backups are read from FlashBlade and new database is created on FlashBlade. The performance would be limited by the availability of the network interfaces and compute resources.

If your source database is very busy and running RMAN backups would impact the database performance, easy solution is to take FlashRecover snapshots of the source database, place it on a second host (Mount host) and perform backups from the mount host onto FlashBlade. This would certainly require RMAN recovery catalog database.

We tested the cloning process in our lab and documented the steps in the White Paper which would be available under the Data Protection section of our website. We had two databases OLTP (1.2TB) and DW (1.8TB) that were backed up on to FlashBlade and clones were created from these two backups. The clones were created under 20 minutes in both cases.

We also validated the cloned environment by running query against the cloned database while creating a clone of another database from the backup. The query on server oradb01 performed over 2.7 GBps of read throughput while the RMAN duplicate on the server oradb02 performed read and write bandwidth of 1.08GBps. The performance in these test cases were limited by the compute and network resources on hosts leaving lot of performance room in FlashBlade to accommodate further workloads.



Data is the new currency in the modern era. It must be analyzed, backed up, recovered, and iterated upon at the speed of modern businesses. For more detailed information on Fast Backup/Restore not just for Oracle but for other databases, please visit ["Next-Gen Data Protection"](#).