WHITEPAPER

GIVE YOUR ORACLE DBAS THE BACKUPS THEY REALLY WANT

The Business Value of Data Domain for Oracle

WHY YOU SHOULD TAKE THE TIME TO READ THIS PAPER

- **Speed up backups** (Up to 68TB/hr, Data Domain systems are 1.5 times faster than the closest competitor so you can complete full backups within backup windows.)
- **Make database administrators happy** (Give DBAs fast daily full database backups.)
- **Ensure database recovery** (Data Domain Data Invulnerability Architecture provides the industry's best data integrity protection for your storage of last resort.)
- **Simplify operations** (Stop managing tape cartridges and greatly simplify day-to-day and disaster recovery operations.)
- **Reduce bandwidth consumption** (Use up to 99% less bandwidth for your backups and replication.)
- **Maximize database backup storage efficiency** (Data Domain deduplication improves Oracle database backup efficiency.)
- **Increase backup success** (Maximize success through link aggregation & failover with less backup jobs to restart.)
- **Reduce licensing costs** (Reduce expensive Oracle backup application licensing costs by doing backups directly to a Data Domain system with Oracle RMAN.)
- **Let your DBAs have their cake and eat it too!** (Data Domain's Oracle optimized deduplication allows DBAs to maximize throughput with multiplexing and still achieve high deduplication efficiency.)

October, 2017

Table of Contents

EXECUTIVE SUMMARY

WHAT YOUR ORACLE DBAS REALLY WANT FOR THEIR BACKUPS

If you ask your Oracle® database administrators what they really want for database backup and recovery, they will likely tell you they would like to do daily full backups and control it themselves without any dependence on the backup operations staff. They also would like to know what backups are onsite and what backups are offsite for disaster recovery. Because full backups can take a long time, especially with large databases, DBAs also need backups that are fast enough to complete within available backup windows. For very large databases, it is more common to see a combination of periodic full backups with incremental backups. And finally, they want to retain multiple backups over multiple days so they can go back to them if needed. The perfect answer to all these requirements is performing backups directly to a Data Domain deduplicated protection storage system with the Oracle RMAN tools that DBAs already know and trust.

INTRODUCTION

This paper focuses on doing Oracle RMAN backups directly to a Data Domain® system using the Oracle RMAN Agent, part of the Dell EMC Data Protection Suite Family®. The purpose of this paper is to explain the technical and business reasons why Data Domain systems can give your Oracle DBAs exactly what they need for database backups.

AUDIENCE

This paper is intended for Dell EMC® customers, Dell EMC sales, Dell EMC systems engineers, Dell EMC partners, and anyone else who is interested in learning more about Data Domain technology and all the unique advantages that it can provide for Oracle database backups.

CHALLENGES WITH TRADITIONAL ORACLE BACKUPS

Before you can fully appreciate what Data Domain systems can do to improve Oracle database backups, we would like to review some of the factors that limit traditional backups. Then we will explain the ways that Data Domain systems gain efficiencies and all the benefits that you can take advantage of.

TRADITIONAL BACKUP LIMITATIONS

There are many factors that limit overall backup performance and reliability causing backups to run into production windows and reduced backup success. We will focus on the most common challenges:

- Limited number of physical tape drives
- Unreliability of physical tape drives and tape media
- Slow physical tape drives
- Tape drives shoe-shining if tape buffer empties
- Limited LAN bandwidth
- Inefficiency of CIFS and NFS for backups
- Limited number of Fibre Channel connections
- Speed of Fibre Channel connections
- Overworked backup servers

Many backups are limited to some degree by LAN bandwidth. Now add the backup inefficiency of the CIFS and NFS protocols over Ethernet and you have exacerbated the problem further. Other backups are limited by the number and speed of Fibre Channel connections. Backups over Fibre Channel can be limited by the speed of older infrastructure and the expense of upgrading to the latest port speeds every few years. Another common choke point in traditional backups is overworked backup servers especially during full backup periods. You also have to deal with slow or a limited number of unreliable physical tape devices and tape media errors. With physical tape, if you can't keep the data streaming fast enough, the tape buffer will empty resulting in dramatically reduced backup performance. And finally, without deduplication technology, too much storage capacity, data center footprint, and backup bandwidth is consumed. These problems result in database backups that frequently don't complete successfully within established backup windows.

DBAS AND BACKUP ADMINISTRATORS

In addition to the limitations mentioned above, there can be other organizational backup challenges associated with Oracle backups. In many cases, the Oracle DBAs report to one part of the IT organization and the backup or storage administrators to another. This can cause friction because the Oracle DBAs want to control their own backup and recovery and the backup or storage administrators don't like to be bothered by the DBAs making special requests all the time. Many times the DBAs want reports that they can't run themselves. In some cases, the Oracle DBAs will go out and find their own backup solution so that they can control everything themselves on not be dependent on backup operations. Backup and storage administrators are also nervous about giving DBAs control fearing they might consume too much shared backup capacity resources.

WHAT IF YOU COULD GIVE THE DBAS THE BACKUPS THEY REALLY WANT?

What if you could give your Oracle DBAs super-fast database backups over Ethernet or FC to an efficient deduplication storage platform, with total control and catalog awareness leveraging the RMAN interface they already know? Now imagine also reducing the workload on your current backup servers and achieve all of this while limiting the DBAs on the deduplication storage capacity that they can use in a shared duplication appliance? That's exactly what RMAN direct backups to Data Domain systems can do.

DATA DOMAIN SYSTEMS – LEADERSHIP FOR ORACLE BACKUPS

ORACLE RMAN DIRECT TO DATA DOMAIN - DEFINITION

Oracle RMAN direct to Data Domain refers to doing backups of your Oracle databases from the Oracle database server using the Oracle RMAN utility and the Oracle RMAN Agent directly to a Data Domain system without using a third party backup application. Your Oracle database server sends backups directly to a Data Domain deduplication storage target Data Domain BoostTM over Fibre Channel or Ethernet. Since doing backups with Data Domain Boost offers tremendous advantages over CIFS and NFS, we will focus on RMAN backups using DD Boost for the remainder of this paper.

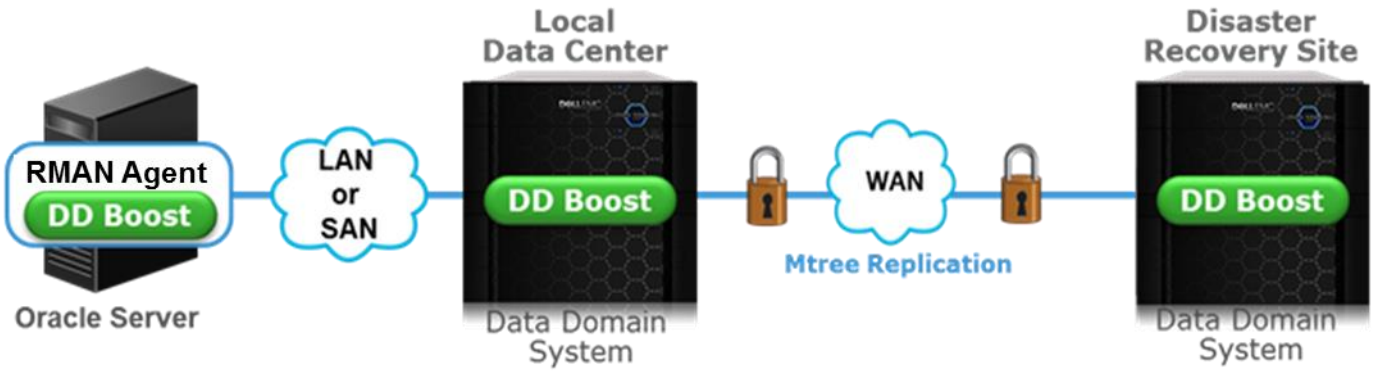

Figure 1: Oracle RMAN Agent with MTree Replication

INHERENT BENEFITS OF RMAN DIRECT BACKUPS TO DATA DOMAIN

RMAN direct backups to Data Domain bypass your backup servers completely. This means the Oracle database backup workload is removed from your backup servers giving back resources that can be used for other backups.

Many backup applications require expensive client and plug-in licensing software per Oracle database server being protected. Oracle RMAN direct to Data Domain may reduce or eliminate some of these expensive licensing costs because the databases are being protected by RMAN directly and not by a third party backup application.

DATA DOMAIN DATA INVULNERABILITY ARCHITECTURE

Ensuring data integrity should be priority one for your backup platform because it is typically the storage of last resort. When you try to recover data from backups, this is probably the last place that the data exists. When considering backup solutions, none of the other features and capabilities matter if the data cannot be recovered when it's really needed. No single protection mechanism can provide protection for all the different ways your data can be lost. The Data Domain Data Invulnerability Architecture includes 4 different protection mechanisms that together provide the industry's best protection for data integrity and recovery.

- End-to-end Verification
- Fault Avoidance and Containment
- Fault Detection and Healing
- File System Recoverability

BENEFITS OF DATA DOMAIN DATA INVULNERABILITY ARCHITECTURE

With the Data Invulnerability Architecture, you can reliably recover your critical data and trust that the data will be exactly as you expect it from your storage of last resort. No other vendor provides this same level of attention to data integrity. Data Domain systems check the data saved to the data sent, and ensures your data is stored correctly. In addition, the system takes precautions not to trash existing data by never appending new data to previous data, which ensures your data doesn't get overwritten. Data Domain systems also protect against data loss due to power failures or dual disk drive failures. Background data scrubbing also protects against bit flips with on-the-fly error correction, which automatically ensures your data stays recoverable and correct. And finally, unlike most vendors, Data Domain systems can be rebuilt from scratch in a reasonable timeframe to ensure you're up and running as quickly as possible. This commitment to ensuring data integrity should give you confidence that Data Domain systems will protect your data better than anyone else.

DATA DOMAIN SISL IS THE FOUNDATION FOR SPEED

The foundation for Data Domain system's industry leading performance is the Stream Informed Segment Layout (SISLTM) scaling architecture. Specifically, Data Domain systems perform 99% of the deduplication effort in CPU and RAM, which enables fantastic performance even with backup inefficient protocols like CIFS, and NFS. Data Domain systems can achieve backup speeds up to 68 TB/hour which is about 1.5 times faster than our nearest competitor. SISL means Data Domain systems are not spindle-bound for performance like other deduplication platforms. This is also the reason why Data Domain systems have dramatic increases in overall performance with each successive generation using the latest Intel processors.

BENEFITS OF DATA DOMAIN SISL

There are 2 important benefits from SISL – faster backups and investment protection. Most importantly, since Data Domain systems are the fastest in the industry, they will help you meet tight backup windows in the face of exploding data growth. Secondly, because Data Domain systems performance increases with Intel performance, it follows Moore's Law. This means that future Data Domain systems will continue to realize dramatic improvements in speed and scalability as future CPUs are used in new Data Domain systems. As new technology is introduced, many of our systems enable you to replace the controller with a next generation model while leaving all the backup data in-place. This investment protection ensures you can dramatically improve backup performance and scalability without disrupting operations.

DATA DOMAIN BOOST DISTRIBUTES THE DEDUPLICATION PROCESS

Data Domain Boost distributes parts of the deduplication effort down to the individual Oracle database servers leaving the Data Domain system to focus its energy on determining what is unique and writing the new data to disk. The larger the Oracle database shop the more significant this distribution is. For example, a backup shop with ten or more Oracle database servers would have ten server resources each doing some of the deduplication effort with DD Boost. Without DD Boost, the entire deduplication effort is being performed by the Data Domain system and all data must be sent from the source database server to the Data Domain system. Surprisingly, the effort to assist with some of the deduplication effort using DD Boost is less than the effort required to move full backups from the Oracle database servers to a target Data Domain system without DD Boost.

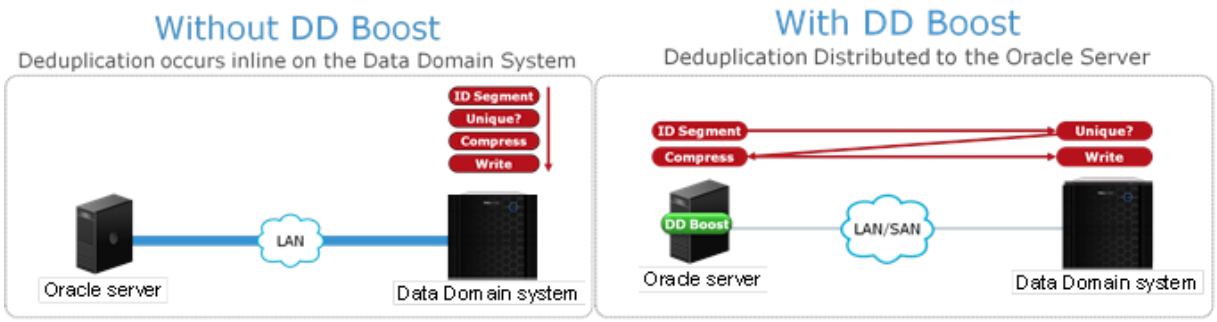

Figure 2: DD Boost distributed segment processing

BENEFITS OF DATA DOMAIN BOOST

Data Domain Boost improves the speed backups by 50% or more using the same network infrastructure that you have in place today. This is really a simple point. Would you like to perform your backups at the speed of X without DD Boost or up to 50% faster using the same exact hardware with DD Boost? This combination of Data Domain SISL and DD Boost is why a single controller Data Domain system can achieve performance of up to 68 TB/hr! This improved performance gives you the ability to finish daily full Oracle database backups within backup windows with breathing room for future data growth.

By distributing parts of the deduplication effort to the Oracle database server itself with DD Boost, the only data that actually has to be sent from the database server to the Data Domain system is new data that is unique. This means up to 99% less data is moved allowing you to more efficiently use the connectivity resources that you have today. Because DD Boost is being leveraged on the Oracle database server, this bandwidth advantage spans the entire backup path all the way from the Oracle database server to the Data Domain system.

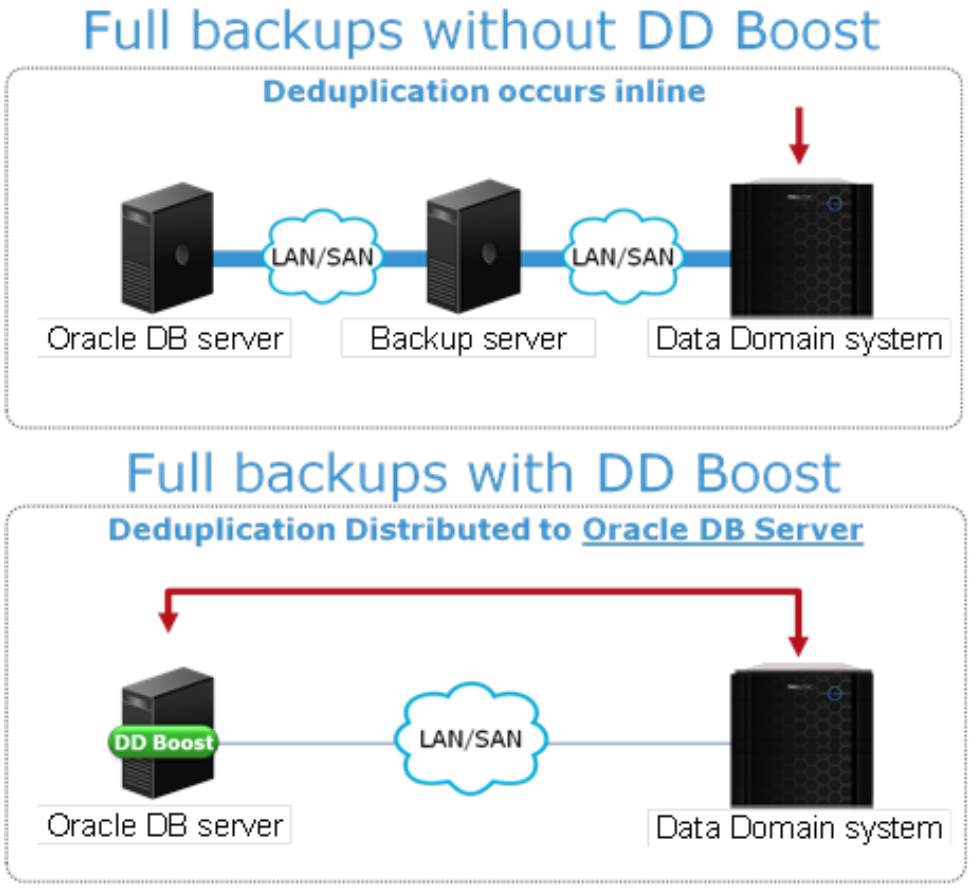

Figure 3: DD Boost minimizes bandwidth used

For Ethernet backup environments that may be experiencing LAN bandwidth saturation during full backups, DD Boost over Ethernet will provide serious performance improvements and may defer the costs of doing a LAN upgrade.

Figure 4 illustrates specific examples of just how much less data is moved using DD Boost for a 1TB backup at three different deduplication levels. At a 4:1 deduplication ratio only 256GB is sent to the Data Domain system, at 8:1 deduplication ratio only 128GB is sent to the Data Domain system and at 12:1 deduplication ratio only 84.4GB of data is sent to the Data Domain system for a 1,024GB backup!

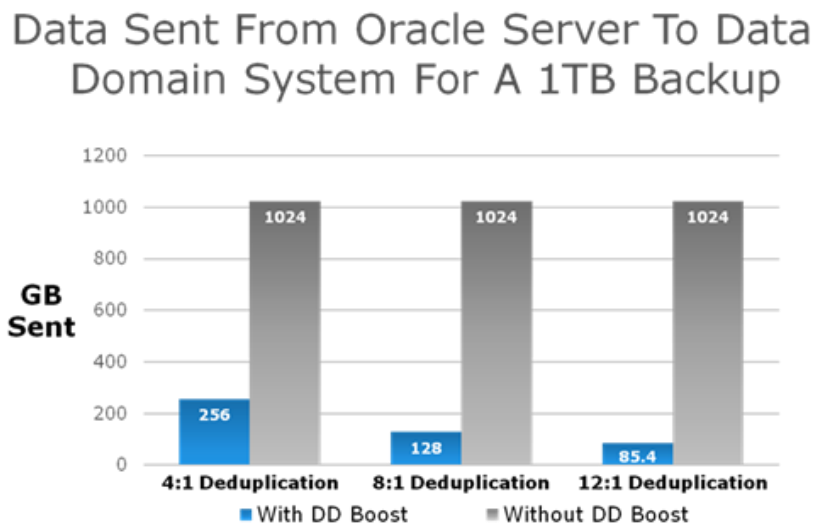

Figure 4: DD Boost means less data is sent to Data Domain

Data Domain implementations typically leverage multiple paths to the Data Domain system. One of the additional benefits of DD Boost is automatic load balancing of the backup workload across the available paths further maximizing performance and efficiency. Another benefit is automatic path failover in the case of one path having a problem improving the reliability of your backups. This means your backups continue to run even if you lose a backup path resulting in higher backup & recovery success rates and less work for your staff with fewer failed backup jobs to restart.

One advantage of distributing the deduplication process with DD Boost is not very intuitive. DD Boost actually reduces CPU utilization on the Oracle database servers even though they're executing parts of the deduplication process. As it turns out, sending data is a lot more resource intensive than executing parts of the deduplication process, so by sending less data, on average we reduce the backup impact on the Oracle server by 20 to 40%.

MANAGE SHARED DATA DOMAIN CAPACITY WITH LOGICAL QUOTAS & STREAM COUNTS

Data Domain systems can be shared for RMAN direct backups and many other backup or archive applications at the same time. Data Domain systems include a feature to set logical quotas and limit stream counts per MTree so that backup or archive applications can't consume all the capacity in a shared Data Domain system. You can set one threshold to send warnings and another threshold to stop any further backups.

BENEFITS OF DATA DOMAIN LOGICAL CAPACITY QUOTAS & STREAM COUNTS

Giving DBAs direct access to Data Domain is a good thing as long as they don't consume too much storage capacity or too many streams on a shared Data Domain system. Data Domain's logical quota & stream limit features provides the access that the DBAs want plus capacity & stream count protection for the backup or storage administrator. With agreed to in advance threshold settings, Data Domain soft quota capacity warnings, hard quota capacity limits, and stream limits give the backup administrators and DBAs the best of both worlds.

DATA DOMAIN ORACLE OPTIMIZED DEDUPLICATION

And if we haven't already blown you away with what Data Domain can do for Oracle data protection, we still have one more key Data Domain feature to brag about. Deduplication efficiency depends on backup streams looking very much the same from day-to-day. This is a result of typically starting backups at the same point and going in the same order each time. Data Domain uses the most efficient variable length segmentation approach to determine logical places to segment the incoming backup stream. The result of starting at the same place, going in the same order, and efficient variable length segmentation is achieving high deduplication ratios up to 30:1 with typical enterprise data and retention periods.

In the physical tape backup world, multiplexing is very common. Multiplexing means sending multiple backup streams to a single target device mixing the backup data in order to keep the tape buffer full for the physical tape drive so that it functions as fast as it can. Multiplexing is typically enabled by a setting in the backup application.

It is also very common to backup Oracle databases using multiple channels in order to improve overall backup performance. When the Oracle filesperset value is set greater than 1, multiple channels are used. This is, in fact, another way to multiplex backups.

Unfortunately the multiplexing methods mentioned above have a negative impact on typical deduplication efficiency because it varies the backup stream from day-to-day. When the backup data is mixed into a common stream, even though they may start in the same place, the data won't be in the same order every day. The result is significantly less deduplication efficiency and higher storage costs. This is true for any vendor's deduplication solution. For this reason, the standard deduplication best practice is to turn off multiplexing and set Oracle filesperset = 1.

Until now, customers have been forced to choose between high performance and high deduplication efficiency. With Data Domain Oracle optimized deduplication technology, you can use multiplexing and still achieve high deduplication efficiency for Data Domain systems dedicated to Oracle DB backups.

BENEFITS OF DATA DOMAIN ORACLE OPTIMIZED DEDUPLICATION

Data Domain's Oracle optimized deduplication technology means you no longer have to choose between maximizing speed or maximizing deduplication efficiency, you can have both at the same time for systems dedicated to Oracle DB backups. Multiple Oracle channels maximize database backup performance. Higher deduplication ratios mean less storage used, less bandwidth used, and in the end, less cost and complexity.

EFFICIENT REPLICATION FOR DISASTER RECOVERY

Data Domain Replicator provides a bandwidth efficient replication option enabling a cost effective alternative to physical tape based operational and disaster recovery. The Oracle RMAN Agent leverages bandwidth efficient Data Domain MTree replication. With RMAN direct to Data Domain with the Oracle RMAN Agent, the DBAs have a bandwidth efficient replication option to a remote site for disaster recovery data protection.

BENEFITS OF DATA DOMAIN REPLICATION FOR DISASTER RECOVERY

Oracle RMAN Agent backups direct to Data Domain with Data Domain Replicator means you can finally stop managing thousands of physical or virtual tape cartridges making your day-to-day operational backup environment simpler and less labor intensive. Thanks to Data Domain's efficient variable length segmentation, minimal bandwidth will be needed for replication keeping WAN costs low. This makes your disaster recovery environment simpler, more reliable, and more cost effective.

CONCLUSION

Oracle database backups include important challenges. After reading this paper, hopefully you have a better understanding how Dell EMC Data Domain systems for Oracle RMAN can dramatically improve the protection of your critical Oracle databases.

In summary Data Domain systems will help you:

- Ensure the recoverability of your critically important Oracle databases with Data Domain's Data Invulnerability Architecture on your storage of last resort
- Improve Oracle database backup performance by 50% or more, up to 68 TB per hour, which is about 1.5 times faster than our nearest competitor
- Use up to 99% less bandwidth, more efficiently using your existing network
- Remove the Oracle database backup workload from current backup servers freeing up backup server resources that can be used for other systems
- Maximize backup success, efficiency, performance, and reliability, with less jobs to restart because of automatic path load balancing and path failover
- Have happier Oracle DBAs by giving them control of their own operational backups & replication for Disaster Recovery
- Eliminate expensive Oracle database backup licensing costs
- Manage DBA consumption of your shared Data Domain storage capacity through the use of soft and hard logical capacity quotas
- Limit DBA consumption of shared Data Domain stream resources by leveraging stream limit controls
- Simplify day-to-day backup operations by eliminating tape cartridges with Oracle RMAN Agent backups directly to Data Domain
- Better leverage your existing infrastructure investments by supporting DD Boost over Ethernet or Fibre Channel
- Enjoy maximum speed and deduplication efficiency at the same time thanks to Data Domain's Oracle optimized deduplication

DELL EMC STORE: CONFIGURE AND COMPARE

Compare features, see options, and get pricing for Data Domain systems. Visit the Dell EMC Store now.

If you would like to know more about our Data Domain technology differentiation, please refer to our Data Domain Data Invulnerability Architecture, Data Domain SISL, Data Domain Replicator, Oracle RMAN Agent, Data Domain Oracle Optimized Deduplication, and Data Domain Physical Capacity Measurement white papers found on Dell EMC.COM. We encourage you to check out The Dell EMC Core Technologies blog on this and other Dell EMC Data Protection and Availability topics.

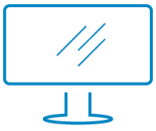

Learn more about Dell EMC Data Domain solutions

Contact a Dell EMC Expert