

# DELL EMC GDDR (GEOGRAPHICALLY DISPERSED DISASTER RESTART) for PowerMax 8000 & VMAX ALL FLASH 950F

ARM your data center for protection against disaster

## ABSTRACT

This paper presents an overview of GDDR software technology used to automate, react, and monitor (ARM) large scale mainframe and mixed mainframe-open systems environments providing continuous operations or automated failover during planned or unplanned events

September 2020

## Revisions

Date	Description
August 2019	Revision 0.4 which includes content related to GDDR 5.2 (Q3 2019) as well as reformatting and attributions of release 0.3
September 2020	Revision 0.5 which includes content related to GDDR 5.3 (Q3 2020) as well as reformatting and attributions of release 0.4

## Acknowledgements

This paper was produced by the following:  
Authors: Brett Quinn, Bruce Klenk, Paul Scheuer  
Support: Mainframe Corporate Systems Engineering

The information in this publication is provided as is. Dell Inc. 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 software described in this publication requires an applicable software license. Copyright © 2020 Dell Inc. or its subsidiaries. All Rights Reserved. Dell, EMC, and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be the property of their respective owners. September 2020. White paper h13805.5

# CONTENTS

<b>INTRODUCTION</b> .....	<b>4</b>
<b>KEY GDDR CONCEPTS</b> .....	<b>5</b>
Situational Awareness .....	5
Leadership Arbitration and Control .....	5
Survivor Recognition.....	5
<b>SUPPORTED TOPOLOGIES</b> .....	<b>5</b>
SRDF/S with ConGroup.....	6
SRDF/S with AutoSwap.....	7
SRDF/A .....	8
SRDF/STAR .....	9
SRDF/SQAR.....	10
GDDR Tape: Support with Dell EMC Disk Library for Mainframe (DLm) .....	12
<b>GDDR ENHANCEMENTS BY VERSION LEVEL</b> .....	<b>12</b>
Version 5.2 .....	12
Version 5.3 .....	13
<b>CONCLUSION</b> .....	<b>14</b>

## Introduction

GDDR is a mainframe software product that *automates* business recovery procedures by *reacting* to events that its *monitoring* capability detects in the data center. These three functions: automate, react and monitor (ARM) combine to enable continuous operations across both planned and unplanned outages. GDDR is designed to perform planned data center site switch operations as well as to restart operations following disasters ranging from the loss of compute capacity and/or disk array access, through to total loss of a single data center or a regional disaster, including the loss of dual data centers. GDDR achieves this goal by providing automation to complement the functionality of the Dell EMC hardware and software products required for business restart.

Because GDDR provides system restart following disasters, it does not reside in the same systems that it is seeking to protect. GDDR resides on separate logical partitions (LPARs) from those that run application workloads. For example, in a three data center SRDF/STAR configuration, GDDR is installed on a control LPAR at each site. Each GDDR node is aware of the other two GDDR nodes via network connections between each site. This awareness enables the monitoring that is required to react to disasters, identify survivors, nominate the leader and then automate the necessary actions to resume operations at one of the customer-chosen surviving sites.

To achieve the task of business restart, GDDR automation extends well beyond the disk layer where Dell EMC has traditionally focused and into the host operating system layer. It is at this layer that sufficient controls and access to third party software and hardware products exist to enable Dell EMC to provide automated recovery services.

GDDR is unique in that it uses an expert system—application of knowledge engineering—to dynamically create an automation script to handle a planned or unplanned event. As a result GDDR is a single product. It can handle the complexity and variability of 13 different customer configurations of sites and software products. Competitors use “versions” of their base software to deal with these configuration differences. This makes it difficult or impossible to move from one configuration type to another without completely re-engineering the solution. For GDDR environments, that’s simply a matter of re-describing the configuration via parameters and rerunning discovery utility software.

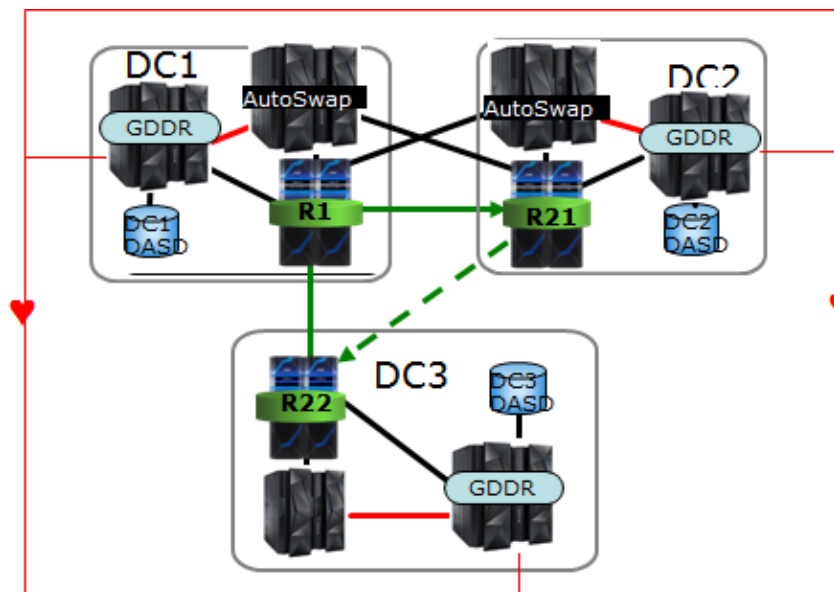
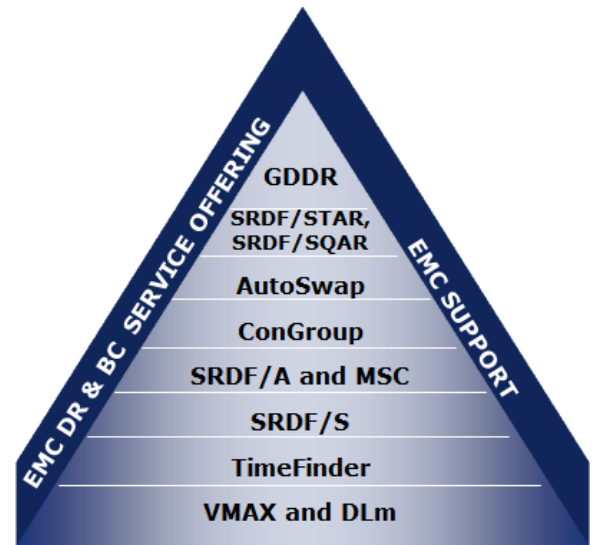


Figure 1: GDDR/Star with AutoSwap

## Key GDDR Concepts

GDDR brings some important concepts to the deployment and operation of Dell EMC business continuity technology:

- Situational Awareness
- Survivor Recognition
- Leadership Arbitration and Control

### Situational Awareness

GDDR brings **situational awareness** to Dell EMC business continuity technology. For example, GDDR is able to distinguish between network outages – SRDF link drops – versus real disasters. This awareness is achieved by periodic exchange of dual-direction heartbeats between the GDDR LPARs. It seems like a simple notion, but to a foundation technology such as SRDF/A there is no means to determine the difference between a link outage and a real disaster.

### Leadership Arbitration and Control

GDDR operates in a Master Owner/ Non-Owner in relation to other GDDR control LPARs. In a three site topology, the GDDR master C-System would normally reside at the DC2 location. However, if the DC2 location is destroyed – or the GDDR C-System itself fails, then one of the surviving GDDR C-Systems will assume the role of the GDDR Master.

Changes to GDDR configuration information can only be made on the GDDR Master C-System. GDDR propagates these changes to the subordinate GDDR systems using inter-system communications facilities built into GDDR.

Restart procedures following disasters are coordinated from the GDDR Master C-System.

### Survivor Recognition

Without automation software, replication technologies do not act on disaster situations to achieve recovery. GDDR has built in intelligence to “look out for” other GDDR systems; constantly checking for disaster situations and constantly ensuring that other GDDR systems are “healthy”. This constant checking allows GDDR to recognize and act on potential disaster situations, even if only one GDDR system survives.

## Supported Topologies

A Dell EMC GDDR complex consists of GDDR control systems (C-Systems), z/OS and open systems hosts, and Dell EMC PowerMax 8000 and or VMAX storage systems which support an organization's mission-critical workload. GDDR is unique in the ability of a single GDDR complex to support multiple z/OS parallel sysplexes and can manage both CKD and FBA disk, providing an enterprise-wide disaster restart solution.

Dell EMC GDDR is available in the following configurations:

- Two site:
  - SRDF/S with ConGroup — The 2-site SRDF/S with ConGroup configuration provides disaster restart capabilities at site DC2.
  - SRDF/S with AutoSwap — The 2-site SRDF/S with AutoSwap configuration provides for continuous availability through device failover between DC1 and DC2.
  - SRDF/A — The 2-site SRDF/A configuration provides disaster restart capabilities at site DC3.
- Three site:
  - SRDF/Star with ConGroup— The 3-site SRDF/Star configuration provides disaster restart capabilities at either DC2 or DC3. Concurrent and Cascaded SRDF support options further minimize the DC3 recovery time objective.

- 2-site SRDF/Star—A variant of 3-site SRDF/Star with ConGroup, this configuration supports a DC2 site with no host and the PowerMax 8000 or VMAX acting as a 'data bunker'.
  - SRDF/Star with AutoSwap— The 3-site SRDF/Star configuration provides both continuous availability between DC1 and DC2 as well as disaster restart capabilities at either DC2 or DC3. Concurrent and Cascaded SRDF support options further minimize the DC3 recovery time objective.
- Four Site:
  - SRDF/SQAR with AutoSwap — The 4-site SRDF/SQAR with AutoSwap configuration provides for continuous availability through device failover between DC1 and DC2 as well as continuous disaster recovery protection through redundant SRDF/A replication out of region to DC3 and DC4.

GDDR can be customized to operate in any of these configurations. GDDR functionality is controlled by a parameter library. During GDDR implementation, this parameter library is customized to reflect:

- The prerequisite Dell EMC software components
- The desired data center topology (two-site ,three-site, four-sire, synchronous or asynchronous, concurrent and/or cascaded).

### SRDF/S with ConGroup

The 2-site SRDF/S with ConGroup configuration provides disaster restart capabilities at site DC2. Figure 2 illustrates GDDR operation in the SRDF/S with Consistency Group environment.

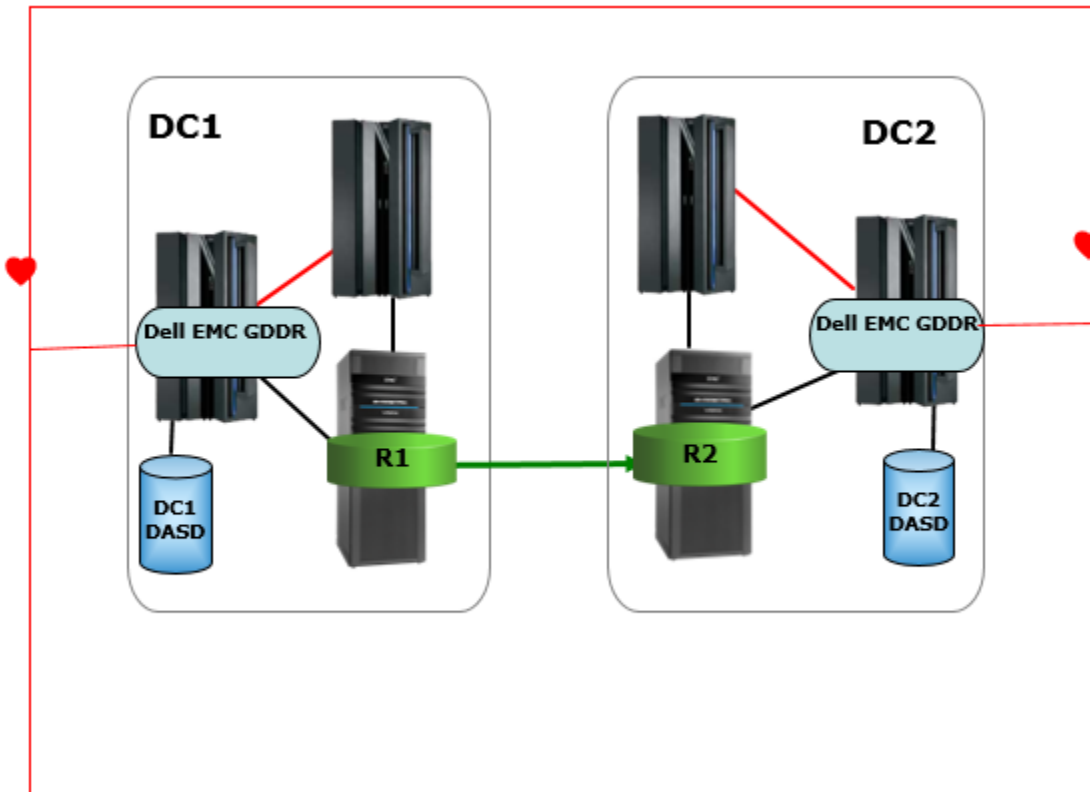


Figure 2 GDDR SRDF/S with ConGroup

Figure 2 shows the two GDDR C-Systems with their heartbeat communication paths, separate from the production disk and computer facilities. Each of the DC1 and DC2 production z/OS LPARs has Dell EMC Consistency Group (ConGroup) software installed. SRDF/S and ConGroup ensure that at the point that GDDR receives notification of an unplanned or failure event, a point of consistency is already achieved.

In this environment, GDDR can do the following:

- Manage planned site swaps
- Restart processing at the secondary site following unplanned primary site events
- Perform standard operational tasks:
  - IPL, system reset, activate, deactivate
  - Trigger stop/start of business workloads
- Actively monitor for unplanned/failure events
  - Sites
  - Systems
  - Loss of SRDF/S
  - ConGroup trip
  - Inter-site communication failure

### SRDF/S with AutoSwap

The 2-site SRDF/S with AutoSwap configuration provides for near-continuous availability through device failover between DC1 and DC2. Figure 3 illustrates GDDR operation in the SRDF/S with AutoSwap environment.

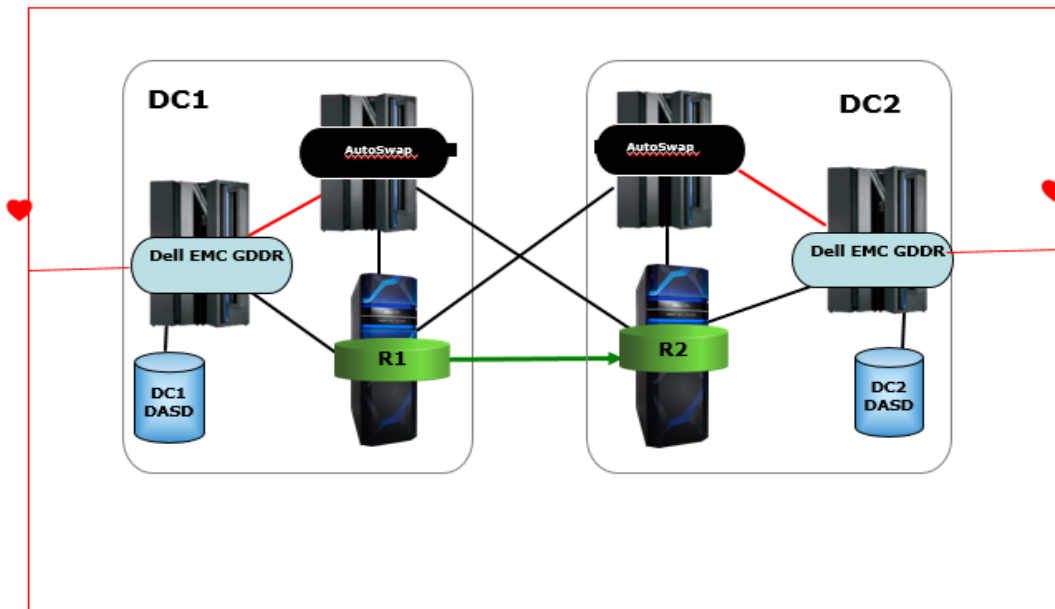


Figure 3 GDDR SRDF/S with AutoSwap

As Figure 3 shows, the relationship between the DC1 and DC2 sites is maintained through SRDF/S replication of primary disk images at DC1 to DC2. Both open systems (FBA) and mainframe (CKD) disk images can be replicated.

Figure 3 shows the two GDDR C-Systems with their heartbeat communication paths, separate from the production disk and computer facilities. Each of the DC1 and DC2 production z/OS LPARs has AutoSwap and Consistency Group (ConGroup) software installed. AutoSwap and ConGroup ensure that a point of consistency exists whenever Dell EMC GDDR receives notification of an unplanned or failure event.

In this environment, GDDR can do the following:

- Manage planned site swaps
- Manage recovery after unplanned site swaps
- Perform standard operational tasks:
  - IPL, system reset, activate, deactivate
  - Trigger stop/start of business workloads
- Actively monitor for unplanned/failure events
  - Sites
  - Systems
  - Loss of SRDF/S
  - ConGroup trip
  - Inter-site communication failure
  - AutoSwap events
  - Configure/reconfigure Couple datasets
- Manage coupling facilities policies

### SRDF/A

The 2-site SRDF/A configuration provide disaster restart capabilities at site DC3. Figure 4 illustrates GDDR operation in the SRDF/A environment.

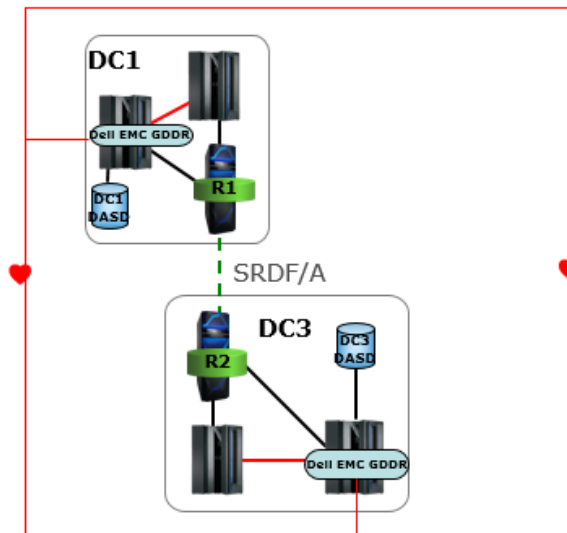


Figure 4: GDDR Two Site SRDF/A

As Figure 4 shows, the relationship between the DC1 and DC3 sites is maintained through SRDF/A replication of primary disk images at DC1 to DC3. Both open systems (FBA) and mainframe (CKD) disk images can be replicated.

Figure 4 shows the two GDDR C-Systems with their heartbeat communication paths, separate from the production disk and computer facilities.



GDDR does not have a requirement to “freeze” I/O to obtain a point of consistency. Multi-Session Consistency and SRDF/A provide the mechanism. At the point that GDDR receives notification of an unplanned or failure event, a point of consistency is already achieved through these foundation technologies.

In this environment, GDDR can do the following:

- Manage planned site swaps
- Restart processing at the secondary site following unplanned primary site events
- Perform standard operational tasks
  - IPL, system reset, activate, deactivate
  - Trigger stop/start of business workloads
- Actively monitor for unplanned/failure events
  - Sites
  - Systems
  - Loss of SRDF/A
  - Inter-site communication failure

### SRDF/STAR

The 3-site SRDF/Star configuration provides disaster restart capabilities at DC2 or DC3. Figure 5 illustrates GDDR operation in a concurrent SRDF/Star environment. GDDR can also be configured for operation in a cascaded SRDF/Star environment.




Figure 5: SRDF/Star with Autoswap

The relationship between the DC1 and DC2 sites is maintained through SRDF/Synchronous replication of primary disk images at DC1 to DC2. Both open systems (FBA) and mainframe (CKD) disk images can be replicated. In a concurrent configuration, the asynchronous relationship is between DC1 and DC3, while in a cascaded environment, the asynchronous relationship is between DC2 and DC3.

Figure 5 shows the three GDDR C-Systems with their independent heartbeat communication paths, separate from the production disk and computer facilities. Each of the DC1 and DC2 production z/OS LPARs has Consistency Group (ConGroup) installed.

In this environment, GDDR can perform the following tasks:

- Manage planned site swaps
- Manage recovery after unplanned site swaps
- Manage reconfiguration of the SRDF/Star environment between concurrent and cascaded topologies
- Manage reconfiguration of the SRDF/Star environment from cascaded to concurrent with a primary processing site move
- Perform standard operational tasks:
  - IPL, system reset, activate, deactivate
  - Trigger stop/start of business workloads
- Actively monitor for unplanned/failure events, including:
  - Sites
  - Systems
  - ConGroup trip
  - Loss of SRDF/S
  - Loss of SRDF/A
  - Inter-site communication failure

## **SRDF/SQAR**

The 4-site SRDF/SQAR configuration provides disaster restart capabilities at DC2, DC3 or DC4. Figure 6 illustrates GDDR operation in a SRDF/SQAR environment. This topology features redundant SRDF/A connections for continuous DR protection out of region and provides the ability to resume a two site SRDF/S operation in another region without having to perform a full resynchronization between the arrays.

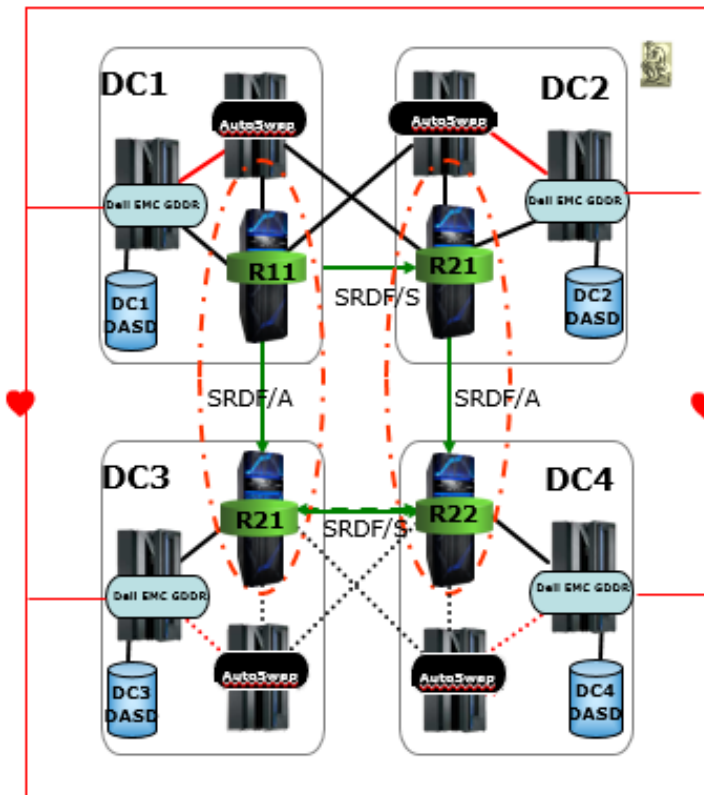


Figure 6: GDDR SRDF/SQAR

In this environment, GDDR can perform the following tasks:

- Manage planned site swaps
- Manage planned region swaps
- Continue remote SRDF/A replication following inter-site link failure
- Resume SRDF/S with AutoSwap protection in remote region following unplanned site/region outage
- Manage recovery after unplanned site swaps
- Manage recovery after unplanned region swaps
- Perform standard operational tasks:
  - IPL, system reset, activate, deactivate
  - Trigger stop/start of business workloads
- Actively monitor for unplanned/failure events, including:
  - Sites
  - Systems
  - ConGroup trip
  - Loss of SRDF/S
  - Loss of SRDF/A
  - Inter-site communication failure

## GDDR Tape: Support with Dell EMC Disk Library for Mainframe (DLm)

Since GDDR is able to produce consistency across open systems platforms and z/OS it was natural to apply this capability to the Dell EMC Disk Library for Mainframe (DLm), as the DLm is simply treated as another open systems host to GDDR. By including a DLm that uses PowerMax 8000 or VMAX storage as its back-end disk into a GDDR managed consistency group, GDDR is able to provide consistency across tape data and the tape file related meta-data stored on DASD, such as the tape catalog, ICF catalog, and DFSMSHsm control datasets. This concept is known as 'Universal Data Consistency' and is unique in the marketplace in its ability to ensure data integrity across tape and DASD in local and remote (synchronous and asynchronous) replication environments.




Figure 7: Universal Data Consistency

## GDDR Enhancements by Version Level

### Version 5.2

Added the following enhancements to release 5.1:

- 1. Support for PowerMax 800 and VMAX 950F All Flash array with Mainframe Enablers V8.x**
- 2. Exploitation of TimeFinder SnapVX:** SnapVX is a new local replication technology available in the PowerMax 8000 and VMAX 950F All Flash arrays that allows up to 256 point in time copies per volume in a very capacity efficient manner using new pointer based technology. While TimeFinder SnapVX is used by GDDR on PowerMax, VMAX3 and VMAX All Flash systems, TimeFinder Clone support has also been added to GDDR for use on older VMAX systems. Expanding support for both TimeFinder features allows GDDR to take consistent point in time copies in environments with multiple generations of VMAX arrays, protecting customers' investments in Dell EMC technology.
- 3. Support for Data Protector for z Systems (zDP):** zDP is a Dell EMC z/OS based application that utilizes SnapVX snapshots to enable rapid recovery from logical data corruption. zDP achieves this by providing multiple, frequent, consistent, point-in-time copies of data in an automated fashion across multiple volumes from which an application level recovery can be conducted. By providing easy access to multiple different point-in-time copies of

data (with a granularity of minutes), precise remediation of logical data corruption can be performed using storage or application-based recovery procedures. zDP provides the following benefits:

- a. Faster recovery times as less data must be processed due to the granularity of the available point in time data copies
- b. Cross application data consistency for recovery data
- c. Minimal data loss compared to the previous method of restoring data from daily or weekly backups. This is especially important for non-DBMS data, which does not have the granular recovery options provided by log files and image copies associated with database management systems.

GDDR 5.0 provides support for zDP, interfacing with zDP where required during GDDR automation tasks to ensure planned and unplanned actions execute successfully on zDP managed volumes.

With the release of PowerMax OS / HyperMax OS 5978 SR, (September 2019), zDP is supported in GDDR STAR and SQAR configurations.

## Version 5.3

Version 5.3 of GDDR adds the following enhancements to release 5.2:

1. **GDDR Tape support for 3-site long distance asynchronous replication for DLm with VNX and/or DD storage.**

This enables GDDR Tape to provide a full complement of failover, switchover, disaster recovery and disaster recovery test operations for DLm using VNX and DD backends in the following two new configurations:

  - a. GDDR for SRDF/Star-A with GDDR Tape for 3 site long distance asynchronous replication
  - b. GDDR Tape 3 site long distance asynchronous replication.
2. **GDDR Tape support for DLm's enhanced multiple independent tenants**

This GDDR feature adds support to control the configuration settings for up to 8 'tenants' within the DLm. A tenant is defined as an instance of a tape library and set of devices for a specific DLm user, such as a z/OS sysplex. This extends the current limit of 2. This capability requires DLm 5.3 as a co-requisite for customers using DD or VNX storage and a minimum of DLm 5.1 for PowerMax/VMAX backends.
3. **GDDR script protection from external device locking processes**

Leverages device expansion of 3390 devices from smaller to larger capacities online and while local and remote replication is active (subject to PowerMax OS supported replication features) in a GDDR-managed configuration by providing an interlock mechanism between GDDR scripts and DVE processes or other processes taking device locks.
4. **zDP Active Monitoring**

Enables support in GDDR of 1024 snapsets in zDP and monitors zDP operational status in order to ensure users are alerted for any zDP failure.
5. **Automated SRDF/A Recovery**

Provides automated recovery for SRDF/A failures (with gold copy protection using SnapVX) within GDDR SRDF/A Multi Cycle Mode environments
6. **TimeFinder/Clone precopy support**

Allows GDDR to utilize 'precopy' in TF/Clone commands so that background copy to the target volume occurs prior to ACTIVATE.
7. **Script execution Previewer**

Provide a list of the script steps, with meaningful descriptions, that will be executed during the current script, based on GDDR's response to planned or unplanned events.

## Conclusion

GDDR provides automation to Dell EMC's enterprise class business continuity solutions from the most complex and intricate four-site solutions through to the simpler two-site configurations. GDDR provides automation for both planned and unplanned outage management for both the z/OS layer and PowerMax 8000 and or VMAX storage and related software.

GDDR automation provides the correct steps, the correct commands, and the correct sequencing of orderly business continuity operations. Users deploying GDDR automation will realize these significant benefits from a GDDR implementation:

- Predictability of outcome
- Improved testability of business continuance plans
- Operational simplicity, allowing lower skilled personnel to perform business continuity operations

Consider shifting responsibility for your business continuity requirements to Dell EMC. GDDR's test hardened resiliency combined with global 7\*24 customer service will help to provide certainty when you need it most.