

# Dell Technologies GDDR (Geographically Dispersed Disaster Restart) for PowerMax

ARM your Data Center for Protection Against Disaster and Cyber Attacks

December 2022

H13805.6

## White Paper

### Abstract

This paper provides an overview of GDDR software technology used to automate, react, and monitor (ARM) large scale mainframe and mixed mainframe-open system environments for continuous operations or automated failover during planned or unplanned outages.

Dell Technologies

## Copyright

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 © 2019-2022 Dell Inc. or its subsidiaries. All Rights Reserved. Dell Technologies, Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Intel, the Intel logo, the Intel Inside logo and Xeon are trademarks of Intel Corporation in the U.S. and/or other countries. Other trademarks may be trademarks of their respective owners.

Published in the USA December 2022 H13805.6.

Dell Inc. believes the information in this document is accurate as of its publication date. The information is subject to change without notice.

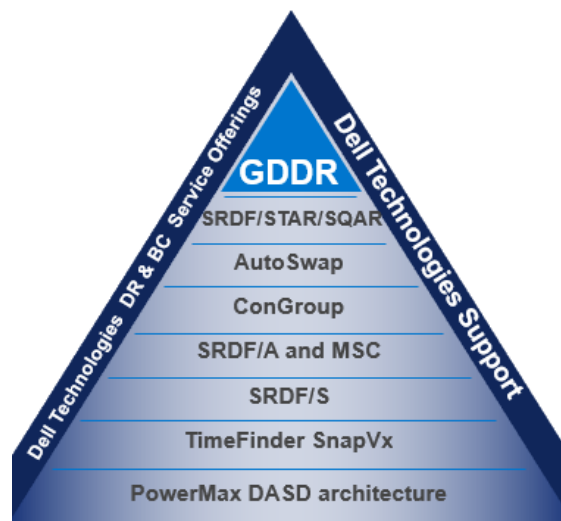
# Contents

- Executive summary ..... 4
- Key GDDR Concepts..... 6
- GDDR Topologies ..... 6
- GDDR Multi-Tenancy support ..... 13
- GDDR zCPA: Cyber Protection Automation for z Systems ..... 14
- GDDR Tape: Support with Dell Technologies Disk Library for Mainframe (DLm) ..... 16
- Conclusion..... 17

## Executive summary

### Overview

Dell GDDR is a mainframe software product that *automates* planned storage switchovers between datacenters as well as unplanned storage recovery by *reacting* to events that its *monitoring* detects in the data center. These three functions (*automate, react, and monitor* (ARM)) combine to enable continuous operations across planned and unplanned outages. GDDR is designed to perform planned data center site switch operations and to restart operations following disasters, ranging from the loss of compute capacity and/or disk and tape 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 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 protects. 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 through 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 business restart, GDDR automation extends 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 provide automated recovery services.

GDDR also provides automation for creating a physical cyber vault copy of data using its Cyber Protection Automation for mainframe (zCPA) feature, supported on all GDDR topologies. zCPA is designed to automate the periodic creation of space efficient snapshots in a separate cyber vault array, which can be optionally air-gapped from the production and disaster recovery arrays managed by GDDR.

In comparison to other mainframe storage switching and failover automation products, Dell GDDR is unique in that it is a single product, using artificial intelligence (AI) to dynamically create an automation script to handle a planned or unplanned event. The AI simplifies GDDR's design, implementation, and maintenance because only ONE product needs to be purchased, maintained, and reconfigured as storage configurations and datacenters change over time. It can handle the complexity and variability of 16 different customer configurations of sites and software products. Contrast this with competitors who rely on many unique "versions" of their base software to deal with configuration differences, making it nearly impossible to move from one configuration type to another without completely re-engineering the solution. For GDDR, it is simply a matter of using parameters to redescribe the configuration and rerunning discovery utility software.

## 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.
December 2022	Revision 0.6 which includes updated images and content on GDDR zCPA.

## We value your feedback

Dell Technologies and the authors of this document welcome your feedback on this document. Contact the Dell Technologies team by [email](#).

**Authors:** Justin Bastin, Brett Quinn, Bruce Klenk, Paul Scheuer

---

**Note:** For links to other documentation for this topic, see the [PowerMax Info Hub](#).

---

## Key GDDR Concepts

### SRDF and GDDR

GDDR uses industry-leading Dell SRDF replication at both synchronous (SRDF/S) and asynchronous (SRDF/A) distances between datacenters to accomplish its objectives.

### Situational awareness

GDDR brings **situational awareness** to Dell Technologies business continuity technology. For example, GDDR can 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 simple, but to the underlying 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 Primary Owner / Non-Owner state in relation to other GDDR control LPARs. In a three site topology, the GDDR primary C-System (Control 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 Primary.

Changes to GDDR configuration information can only be made on the GDDR primary 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 primary C-System.

Customer application systems are referred to as 'managed systems' or P-systems to distinguish them from GDDR C-Systems.

### 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.

## GDDR Topologies

A Dell GDDR complex consists of a minimum mix of two (local and remote) Dell PowerMax 8000, 2500, 8500, and or VMAX storage systems and one or more GDDR control systems (C-Systems), and z/OS (and even open systems) hosts.

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.

- Two Site:
  - SRDF/S with ConGroup — The two-site SRDF/S with ConGroup configuration provides disaster restart capabilities at site DC2.

- SRDF/S with AutoSwap — The two-site SRDF/S with AutoSwap configuration provides for continuous availability through device failover between DC1 and DC2.
- SRDF/A — The two-site SRDF/A configuration provides disaster restart capabilities at site DC3.
- Three Site:
  - SRDF/Star with ConGroup — The three-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.
    - Two-site SRDF/Star — A variant of three-site SRDF/Star with ConGroup, this configuration supports a DC2 site with no host and the PowerMax 8000, 2500, 8500, or VMAX acting as a ‘data bunker’.
  - SRDF/Star with AutoSwap — The three-site SRDF/Star configuration provides both continuous availability between DC1 and DC2, and 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 four-site SRDF/SQAR with AutoSwap configuration provides for continuous availability through device failover between DC1 and DC2, and continuous disaster recovery protection through redundant SRDF/A replication out of region to DC3 and DC4.


Any of these topologies can be augmented with one or more Cyber Recovery (CRx) sites when the Cyber Protection Automation (zCPA) feature is added to a GDDR configuration to populate a dedicated cyber recovery vault array.

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 software components
- The desired data center topology (two-site, three-site, four-site, synchronous or asynchronous, concurrent and/or cascaded, and any Cyber Recovery vault sites).

## SRDF/S with ConGroup

The two-site SRDF/S with ConGroup configuration provides disaster restart capabilities at site DC2. Figure 1 illustrates GDDR operation in the SRDF/S with Consistency Group environment.



**Figure 1. GDDR SRDF/S with ConGroup**

Figure 1 shows 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 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.

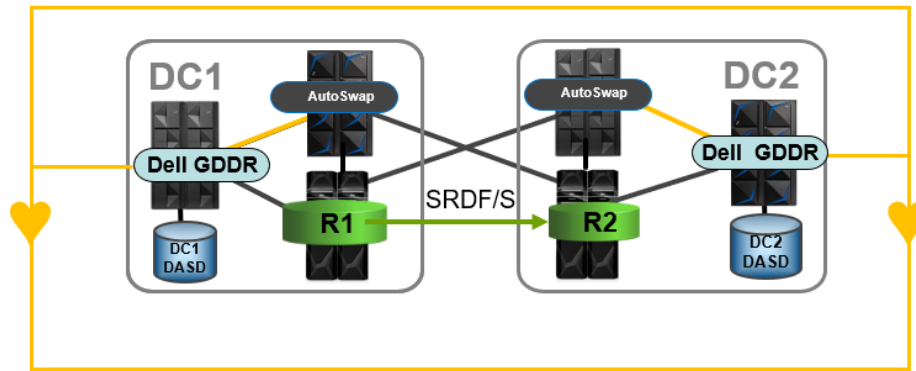
In this environment, GDDR can do the following:

- 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

In Figure 1, the two-site SRDF/S with AutoSwap configuration provides for near-continuous availability through device failover between DC1 and DC2. Figure 2 illustrates GDDR operation in the SRDF/S with AutoSwap environment.





**Figure 2. GDDR SRDF/S with AutoSwap**

In Figure 2, 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 2 also 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 provides for continuous IO operations in the event of an inability to access the primary (R1) array due to an array or site failure, while ConGroup ensures that the R2 arrays preserve a dependent write consistent image of the data in the event of an SRDF/S replication failure. In short, AutoSwap protects the FICON connections to the primary array (providing continuous data accessibility on the R2 array), while ConGroup protects the SRDF/S connections (preserving data integrity on the R2 arrays in the event of a 'rolling outage').

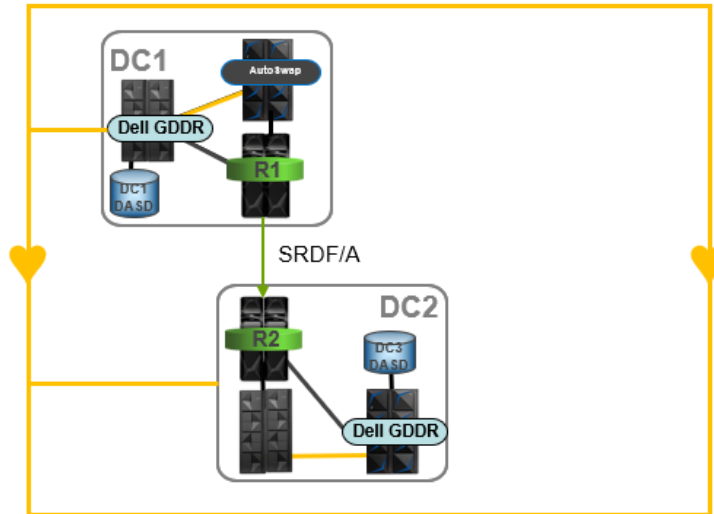
In this environment, GDDR can do the following:

- 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
  - Inter-site communication failure

- Manage coupling facilities policies

## SRDF/A

The two-site SRDF/A configuration provide disaster restart capabilities at site DC3. Figure 3 illustrates GDDR operation in the SRDF/A environment.



**Figure 3. GDDR Two Site SRDF/A**

As Figure 3 shows, the relationship between the DC1 and DC2 sites is maintained through SRDF/A 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.

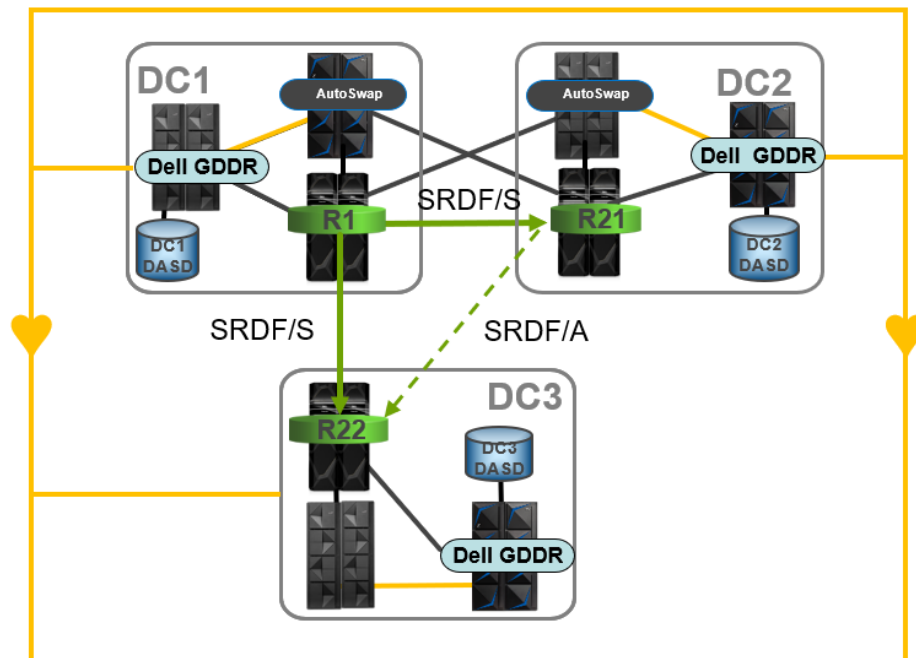
GDDR does not have a requirement to “freeze” I/O to obtain a point of consistency. Multi-Session Consistency and SRDF/A provide the mechanism. At the point that GDDR receives notification of an unplanned or failure event, a point of consistency has already been achieved through these foundation technologies.

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 three-site SRDF/Star configuration provides disaster restart capabilities at DC2 or DC3. Figure 4 illustrates GDDR operation in a concurrent SRDF/Star environment. GDDR can also be configured for operation in a cascaded SRDF/Star environment.



**Figure 4. 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 4 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.

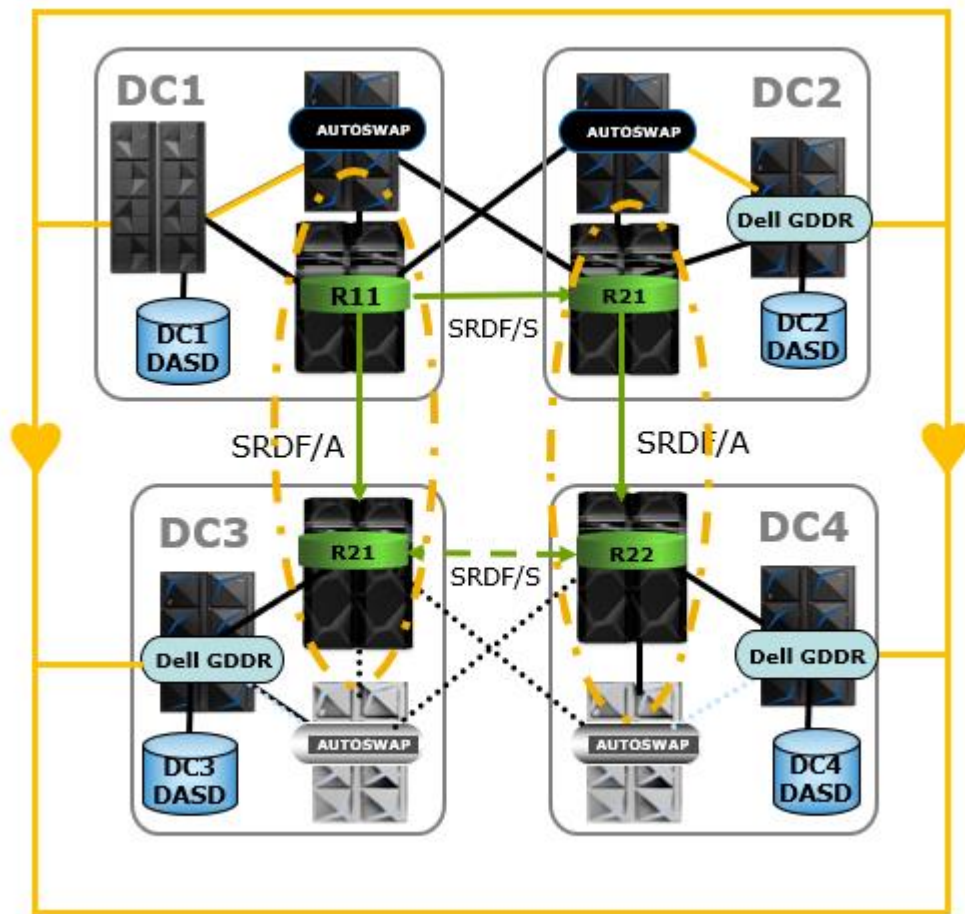
In this environment, GDDR can perform the following tasks:

- 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 four-site SRDF/SQAR configuration provides disaster restart capabilities at DC2, DC3 or DC4. Figure 5 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 5. 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 a 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 Multi-Tenancy support

GDDR is unique in its ability to support multiple instances of its control system within a single z/OS LPAR. While other business continuance automation solutions suffer from z/OS control system LPAR 'sprawl', requiring a dedicated z/OS LPAR per sysplex to act as a control system, GDDR can support up to eight C-Systems (called C-instances) within a single z/OS LPAR. This 'multi-tenancy' capability dramatically reduces the z/OS LPAR count when deploying multiple instances of GDDR in an environment.

In Figure 6, a three site Star with AutoSwap topology and a two site SRDF/A topology are deployed in a single C-system LPAR per site, with two of the GDDR C-system LPARs supporting two GDDR C-instances, one for each of the two deployed GDDR solutions (GDDR 1 and GDDR 2). This results in a 40% reduction in the number of z/OS LPARs required by GDDR (from 5 to 3).

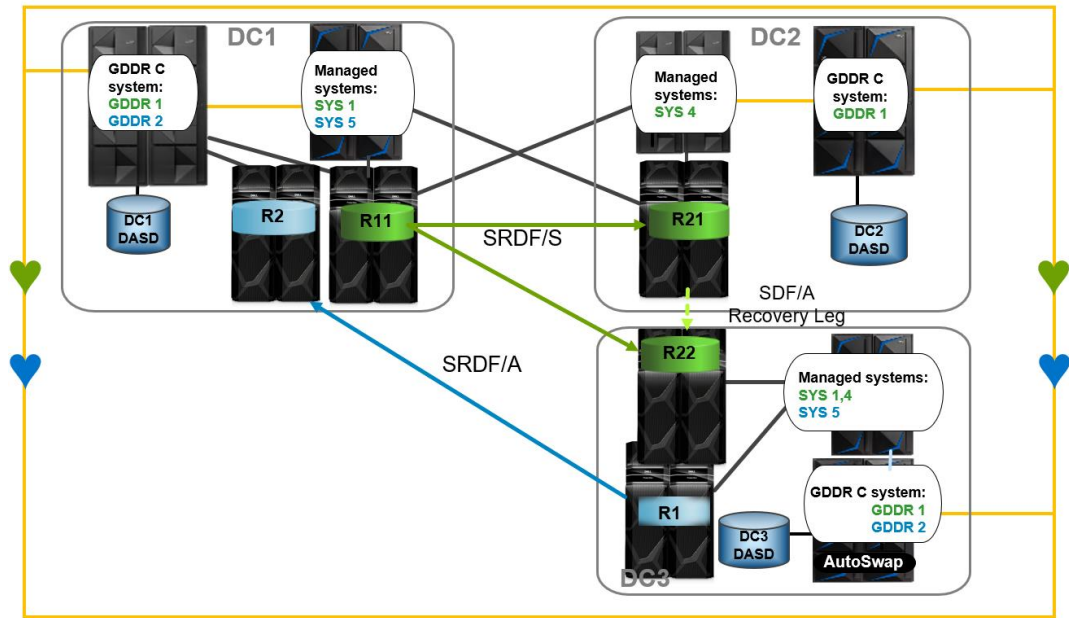


Figure 6. Three Site SRDF/STAR with Autoswap + 2 site SRDF/A

In Figure 7, a three site Star with AutoSwap topology, a two site SRDF/A topology, and a four site SQAR topology are deployed in a single C-system LPAR per site, with the GDDR C-system LPARs supporting multiple GDDR C-instances, one for each of the deployed GDDR solutions (GDDR 1, GDDR 2, and GDDR 3). This results in a 63% reduction in the number of z/OS LPARs required by GDDR (from 11 to 4).

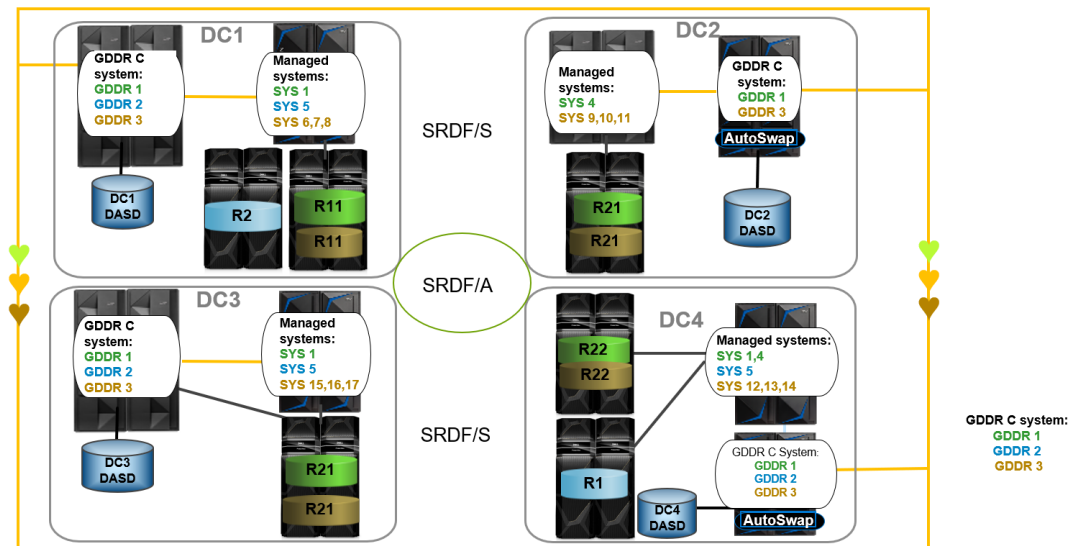


Figure 7. Three site Star with AutoSwap, 2 site SRDF/A, and 4 site SQAR

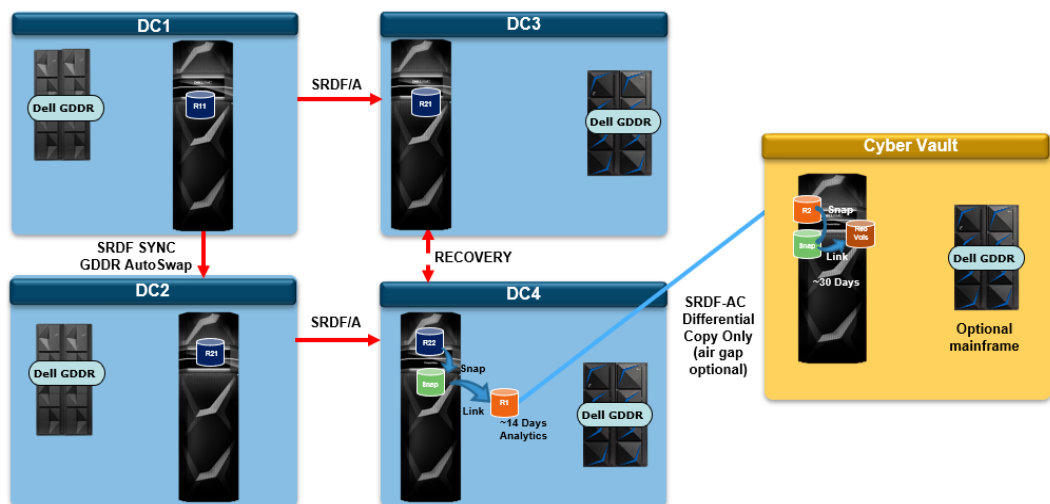
## GDDR zCPA: Cyber Protection Automation for z Systems

Over the last several years there has been an increasing requirement to protect data from logical corruption or malicious access from bad actors. Dell Technologies responded with Data Protector for z Systems (zDP), which provided automation for making up to 1024

space efficient snapshots per volume at frequencies as low as every five minutes in a VMAX or PowerMax array. zDP utilizes capacity in a production array, often at the disaster recovery site, to provide immutable snapshots that can be used for recovery from a cyber attack.

In addition to these 'virtual vaults', requirements have emerged for mainframes to utilize an isolated physical vault for DASD. This can be achieved by using an array dedicated to this purpose with either minimal connectivity (SRDF connections only) or no connectivity (physically air-gapped) in the production data center or at a separate site.

GDDR 6.0 has been enhanced with automation to provide for the ongoing population of a physical vault array on a continuous basis. This feature is called Cyber Protection Automation (zCPA) and can be deployed in any GDDR topology. An example of zCPA in a SQAR topology is shown below with one Cyber Recovery (CR) site. There can be as many Cyber recovery sites as there are sites in the GDDR topology.



**Figure 8. Four site SQAR with a physical cyber recovery vault**

GDDR zCPA works with zDP to provide three different methods for populating a physical cyber vault array, one that utilizes an air gap between the production site and the cyber vault site and two that do not utilize an airgap. In all methods, zCPA monitors and controls zDP snapshot automation to achieve the differential resynchronization of zDP snapshots in a production site array (DC4 in Figure 8) with the cyber vault array in the cyber recovery site CR1. The three methods are:

1. **zCPA Type 1** (not air gapped): Triggered by GDDR's detection of a new snapshot created by zDP in DC(x). GDDR automation drives differential resynchronization of the data to the cyber vault and resumes the zDP automation to create a snapshot in the cyber vault. There is no enabling/disabling of an air gap. This is the default type.
2. **zCPA Type 2** (air gapped): Triggered by the ENABLE (closure) of the DC(x)-CR(y) air gap by a cyber recovery server that resides in the vault site. zDP Snapshot creation is initiated and controlled by GDDR resuming and pausing zDP, controlling snapshot creation in the DC(x) and CR(y) sites. When data is propagated to the CR site, GDDR DISABLES (opens) the airgap by setting the SRDF links offline in the CR site PowerMax.

3. **zCPA Type 3** (not air gapped): Triggered by a SCHEDULE or INTERVAL parameter in GDDR. There is no enabling/disabling of an air gap. Data resynchronization is initiated and controlled by GDDR resuming and pausing zDP, controlling snapshot creation in the DC(x) and CR(y) sites.

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

GDDR Tape is a separately licensed product that monitors the Dell Disk Library for mainframe (DLm) operations for DLm systems with either a VNX, PowerProtect, Data Domain, PowerScale, or Isilon, and automates DLm Disaster Recovery (DR) testing, and failover and failback processes. GDDR Tape runs on the z/OS system ('managed LPAR') and requires no separate GDDR control system, although it can run in a GDDR Control system that supports GDDR for DASD.

GDDR Tape main activities include:

- Actively monitoring the managed environment and responding to exception conditions
- Testing disaster recovery from a snapshot at a remote site
- Running planned and unplanned failover and failback operations

When deployed with a supported backend disk system, GDDR Tape supports two topologies:

- Two-Site Long-Distance -- The two-site DLm Long Distance configuration provides disaster restart and disaster testing capabilities at a remote utilizing the asynchronous copy provided by the various supported backend disk systems.
- Three-Site Long-Distance -- The three-site DLm Long Distance configuration provides disaster restart and disaster testing capabilities at either of two remote sites supported by the asynchronous copy feature of various supported backend disk systems.

GDDR Tape also supports PowerMax as a backend disk subsystem. Because GDDR is able to produce consistency across open systems platforms and z/OS, it was natural to apply this capability to the Dell Technologies Disk Library for Mainframe (DLm), because the DLm is simply treated as another open systems host to GDDR. By including a DLm that uses PowerMax 8000, 2500, 8500, or VMAX storage as its backend disk into a GDDR managed consistency group, GDDR is able to provide consistency across tape data and the tape file related metadata stored on DASD, such as the tape catalog, ICF catalog, and DFSMSHsm control datasets. This concept is known as 'Universal Data Consistency' which is designed to ensure data integrity across tape and DASD in local and remote (synchronous and asynchronous) replication environments using all the topologies supported by GDDR for DASD.



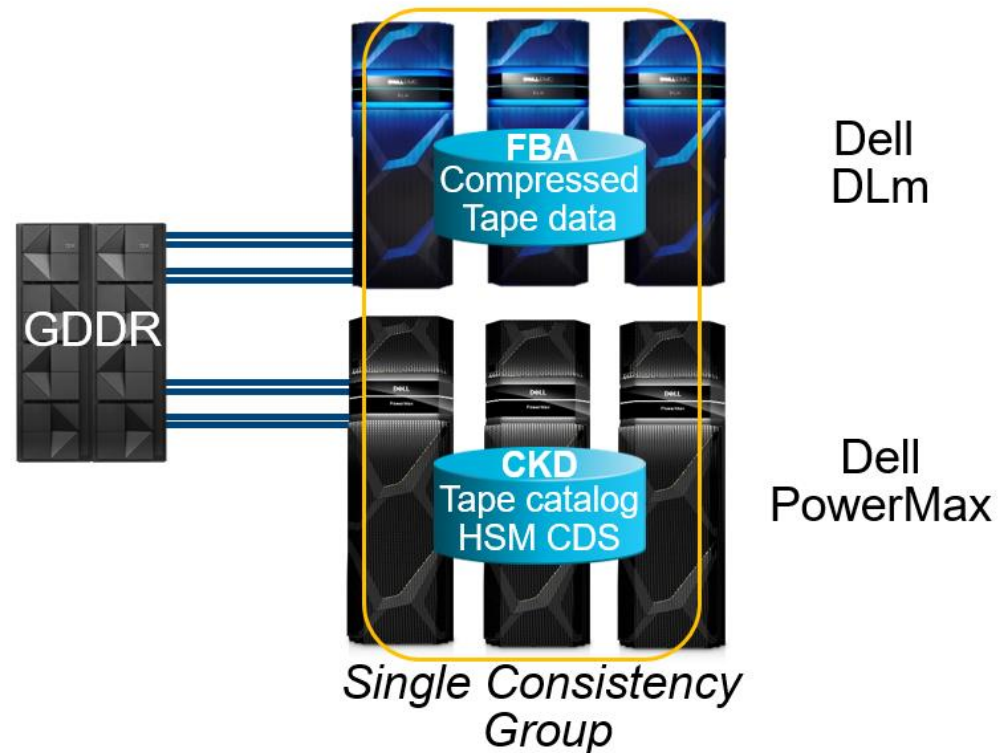


Figure 9. Universal Data Consistency

## Conclusion

GDDR provides automation to Dell's enterprise-class business continuity solutions by covering planned events, unplanned outages, and cyber protection 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, 2500, 8500, or VMAX storage and related software, and provides automation to populate physical cyber vault arrays.

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 covering both logical data corruption and physical disasters
- Operational simplicity, allowing lower skilled personnel to perform business continuity operations

Consider shifting responsibility for your business continuity requirements to Dell. GDDR's test hardened resiliency combined with global 7x24 customer service will help to provide certainty when you need it most.