

Enginuity 5876 New Features for the Symmetrix VMAX Family Storage Systems

A detailed review for Open Systems environments

Abstract

This white paper describes the new features that are available for the EMC® Symmetrix® VMAX™ Family storage systems. These features are available when running Enginuity™ 5876 and Solutions Enabler V7.4. Throughout this document, the term EMC Symmetrix VMAX Family is applicable to all VMAX 10K (VMAXe), VMAX 20K (VMAX), and VMAX 40K Series storage systems.

May 2012

Copyright © 2012 EMC Corporation. All Rights Reserved.

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

The information in this publication is provided “as is.” EMC Corporation 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 EMC software described in this publication requires an applicable software license.

For the most up-to-date listing of EMC product names, see EMC Corporation Trademarks on EMC.com.

VMware is a registered trademark of VMware, Inc. All other trademarks used herein are the property of their respective owners.

Part Number h10497

Table of Contents

Executive Summary	4
Introduction	4
Open Systems New Features	5
Federated Tiered Storage.....	5
TimeFinder enhancements	6
TimeFinder VP Snap.....	7
Clone Terminate with Target Not Ready.....	7
Auto-Terminate/Auto-Cancel logging.....	8
TimeFinder/Snap operations on VMAX 40K hardware	8
RecoverPoint enhancements	8
Array-based RecoverPoint Splitter	8
RecoverPoint settable device size and geometry	9
RecoverPoint managed dynamic Write-Protect-Bypass (WPBP)	9
Additional descriptive device tagging for RecoverPoint.....	10
RecoverPoint n-1 interoperability support.....	10
RecoverPoint and VMware integration	10
Federated Live Migration support	10
FAST VP enhancements	11
VP allocation by FAST Policy	12
SRDF coordination.....	12
Hardware Related Enhancements.....	13
Green power supply.....	13
Dynamic back-end improvements.....	13
Compression for Fibre RA.....	14
Miscellaneous VMAX 40K enhancements	14
Reporting	14
Dispersed system bay	14
RVA Direct Member Sparing.....	14
Conclusion.....	15
References	15

Executive Summary

Enginuity™ version 5876 is the latest release supporting the Symmetrix® VMAX™ Family storage systems. Enginuity 5876 expands EMC's leadership in virtual storage by federating and sharing cooperating pools of resources with automation and mobility to enable Infrastructure-as-a-Service (IaaS).

The capabilities of Enginuity 5876 to network, share, and tier storage resources allows data centers to consolidate their applications from dedicated arrays and to deliver new levels of efficiency with increased utilization rates, improved mobility, reduced power and footprint requirements, and simplified storage management.

Introduction

EMC introduces the Symmetrix VMAX 40K. The VMAX 40K is a new array that further extends EMC's lead in high-end storage. Symmetrix VMAX 40K and Enginuity 5876 build on the Symmetrix VMAX foundation of powerful, trusted, smart storage to provide even higher levels of performance, availability, and intelligence in the virtual data center.

The Enginuity 5876 operating system expands EMC's leadership in virtual storage by federating and sharing pools of storage resources with automation and mobility across the hybrid cloud. Enginuity 5876 is available across the Symmetrix VMAX family, including VMAX 10K (VMAXe), VMAX 20K (VMAX) and VMAX 40K Series storage systems. New Enginuity 5876 software capabilities described below dramatically simplify and automate the management and protection of information in both physical and virtual server environments.

Enginuity 5876 delivers the following new capabilities:

Automation and improved capacity utilization for IBM System z and IBM i environments:

- Fully Automated Storage Tiering for Virtual Pools (FAST VP) and Virtual Provisioning™ for IBM System z and IBM i environments
- VLUN VP migration support for IBM System z and IBM i environments

Increased efficiency

- SRDF® and SRDFe awareness for FAST VP environments

Increased consolidation and reduced TCO

- Federated Tiered Storage (FTS) to manage external storage arrays through a VMAX system*
- RecoverPoint™ Splitter
- TimeFinder® VP Snap

Ease-of-Integration and Management

- Unisphere™ for VMAX
- Dynamic back end
- Federated Live Migration cluster support

* Not available for VMAX 10K (VMAXe).

The features described in this paper apply only to Open Systems environments.

For information regarding the Symmetrix 5876 new features for IBM System z and IBM i platforms (mainframe systems), refer to the corresponding white papers that will be published in conjunction with this paper.

Open Systems New Features

The following features are available with Enginuity version 5876 with Solutions Enabler version 7.4 and support the Symmetrix VMAX Family storage systems.

Federated Tiered Storage

Federated Tiered Storage (FTS) allows the creation of external storage using EMC or third-party storage arrays. FTS enables data mobility between the VMAX 20K (VMAX), or VMAX 40K, and the external arrays, and data mobility between two external arrays, using existing VLUN migration technology. FTS does not require any new hardware. It leverages the available SAN ports (also usable as FAs and RFs) and is based on new Enginuity emulation functionality, known as DX (Disk Adaptor eXternal).

Table 1 shows the storage platforms that support FTS as a 5876 Enginuity feature. For more information on external array support, refer to the FTS support matrix published on www.emc.com/Powerlink.

VMAX 10K (VMAXe)	VMAX 20K (VMAX)	VMAX 40K
N/A	✓	✓

Table 1 - Symmetrix storage supporting FTS

Federated Tiered Storage (FTS) allows LUNs that exist on external arrays to be used to provide physical storage for Symmetrix VMAX. The external LUNs can be used as raw storage space for the creation of Symmetrix devices in the same way internal Symmetrix physical drives are used. These devices are referred to as eDisks. Data on the external LUNs can also be preserved and accessed through Symmetrix devices. This allows the use of Symmetrix Enginuity functionality such as local replication, remote replication, storage tiering, data management, and data migration with data that resides on external arrays.

Benefits of Federated Tiered Storage include:

- Simplifies management of virtualized multi-vendor or EMC storage by allowing heterogeneous arrays to be managed by SMAS, Solutions Enabler, and ProSphere.

- Allows data mobility and migration between heterogeneous storage arrays and between heterogeneous arrays and VMAX.
- Offers Virtual Provisioning benefits to external arrays.
- Allows VMAX enterprise replication technologies such as SRDF and TimeFinder to be used to replicate storage that exists on an external array.
- Extends the value of existing disk arrays by allowing them to be used as an additional storage tier.

FTS is only supported in Open Systems environments. It is not currently supported in System z and IBM i environments. Figure 1 shows the components of an FTS configuration using EMC external storage.

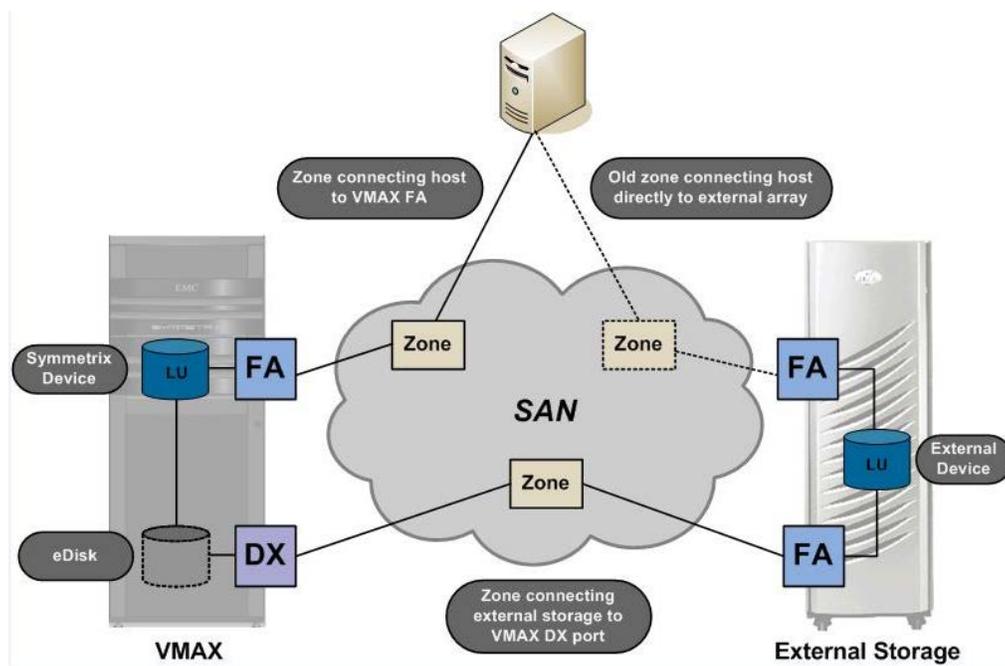


Figure 1. Federated Tiered Storage configuration

TimeFinder enhancements

The following new TimeFinder enhancements described in this section are available with Enginuity 5876 and Solutions Enabler V7.4. These TimeFinder enhancements are supported on the following VMAX Series storage systems as listed in Table 2.

VMAX 10K (VMAXe)	VMAX 20K (VMAX)	VMAX 40K
✓	✓	✓

Table 2 – Symmetrix storage supporting TimeFinder enhancements

TimeFinder VP Snap

TimeFinder VP Snap provides space-efficient snaps for virtual pool devices, while offering the efficiency of snap technology with improved cache utilization and simplified pool management.

TimeFinder VP Snap is available for all Symmetrix VMAX Family storage arrays running Enginuity level of 5876 and higher, and Solutions Enabler V7.4.

This new TimeFinder technology adds the ability to share capacity allocation for track groups originating from the same source volume. (This is similar to the way TimeFinder/Snap currently copies shared tracks.) TimeFinder VP Snap is easy to manage and combines the benefits of both full copy clone and space saving snap technologies.

TimeFinder VP Snap provides the ability for multiple clone VSE (Virtual Space Efficient) sessions to target thin devices and share extent allocations within the thin pool. This reduces the space needed for the storage of saved tracks. TimeFinder VP Snap further provides the ability to perform incremental restores from these VP Snap sessions, however, restrictions apply with restore. Figure 2 shows TimeFinder VP Snap with multiple targets, sharing track allocations within a single thin pool.

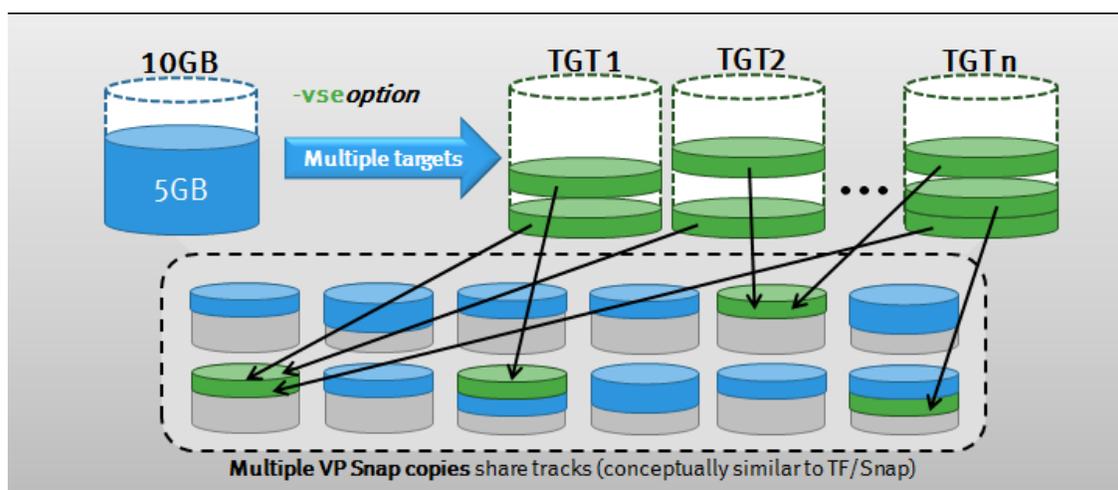


Figure 2. TimeFinder VP Snap sharing track allocations

Clone Terminate with Target Not Ready

Clone terminate allows users to specify the Target Not Ready option, which leaves the target device in either its prior ready or not-ready (NR) state at the completion of the termination operation. After completing termination, the default is set to Ready.

Auto-Terminate/Auto-Cancel logging

The Snap and Clone Auto-Terminate and TimeFinder Mirror Auto-Cancel features produce Audit Log and symapi log entries when sessions are removed.

TimeFinder/Snap operations on VMAX 40K hardware

Starting with Enginuity 5876, all snap sessions will automatically be created as Multi-Virtual Snap, regardless of how the SYMCLI_MULTI_VIRTUAL_SNAP environment variable is set. Only Multi Virtual Snap will be supported on VMAX 20K (VMAX) and VMAX 40K storage arrays.

Note: VMAX 10K (VMAXe) currently supports all TimeFinder/Snap operations.

RecoverPoint enhancements

RecoverPoint provides data-protection and replication technology that captures every write and saves it in a history PiT (point in time) journal. Data recovery operations can be performed locally and/or remotely by rewinding the target volumes back to a selected point in time by using earlier versions of data saved in the journal.

Table 3 shows the availability of RecoverPoint for the Symmetrix VMAX Family storage systems running a minimum Enginuity level of 5876 and Solutions Enabler V7.4.

VMAX 10K (VMAXe)	VMAX 20K (VMAX)	VMAX 40K
✓	✓	✓

Table 3. Symmetrix storage supporting RecoverPoint enhancements

Array-based RecoverPoint Splitter

Deploying RecoverPoint with Symmetrix arrays requires either a host-based splitter (KDriver) or a fabric-based splitter (Cisco SANTap or Brocade FAP). The Symmetrix splitter is supported on VMAX 10K (VMAXe) with Enginuity version 5875 and RecoverPoint version 3.4 SP1 and higher.

With the latest Enginuity release, the RecoverPoint splitter supports VMAX 20K (VMAX) and VMAX 40K running Enginuity version 5876 and RecoverPoint version 3.5.

RecoverPoint supports the following configurations:

- RecoverPoint CRR, CLP, and CDP
- Open Systems only

Figure 3 shows an active RecoverPoint CDP and CRR configuration.

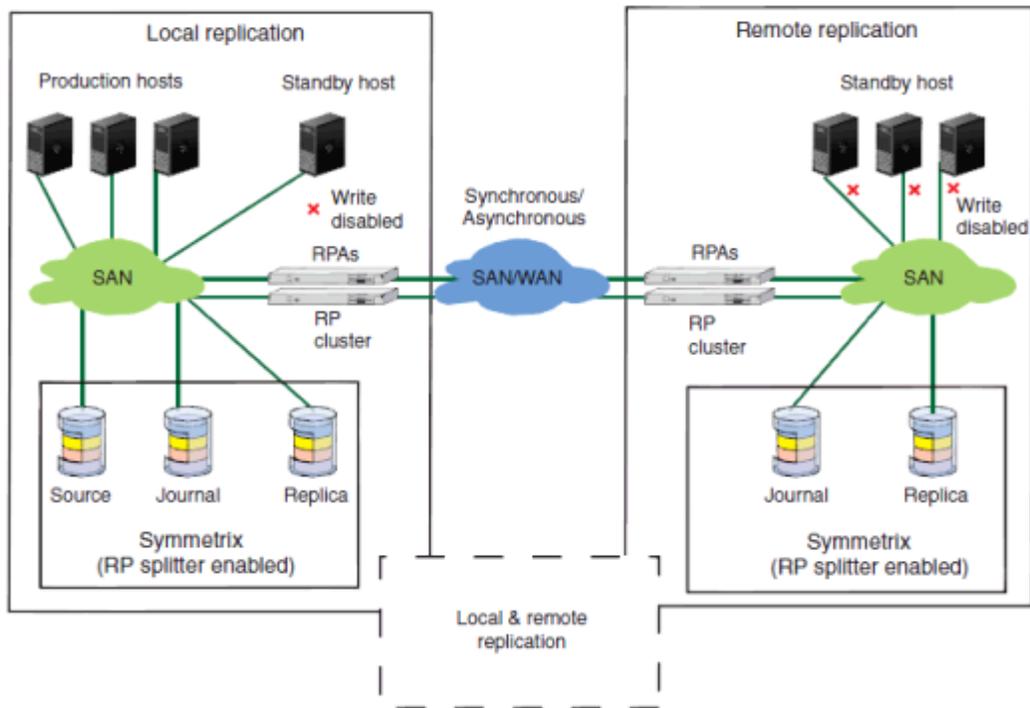


Figure 3. RecoverPoint configuration for local and remote replication

RecoverPoint settable device size and geometry

The RecoverPoint splitter supports replication from smaller devices to larger devices for the purpose of migration. Users can change the native Symmetrix device size, as well as the geometry of a device, in order to support smaller-to-larger device size RecoverPoint replications and migrations. You can specify different geometries in device sets for replication and migration. Typically, this is required in heterogeneous configurations when geometry and cylinder allocation units do not produce equal device sizes. RecoverPoint detects the size differences in the replica set and reduces the devices to a common size in blocks.

RecoverPoint managed dynamic Write-Protect-Bypass (WPBP)

RecoverPoint requires that users specifically enable the WriteProtectBypass feature of the array on the RecoverPoint initiator group. RecoverPoint manages this bypass feature, not the user.

However, by no longer requiring explicit removal of write protection for devices that were once masked to a RecoverPoint cluster, dynamic WPBP device protection eliminates a step in the configuration process.

This feature automates the WPBP flag setting after devices are masked to a RecoverPoint cluster without requiring any additional customer action.

Additional descriptive device tagging for RecoverPoint

With additional descriptive device tagging, RecoverPoint now allows devices to be explicitly tagged as repository, journal, source, or replica volumes. This feature enhances the current tagging mechanism, as well as the manageability of RecoverPoint.

Previously there was no easy distinction between repository, journal, source, or replica volumes. Now the additional descriptive device tagging removes the requirement to cross-reference with RecoverPoint management to determine how each volume is used.

RecoverPoint n-1 interoperability support

RecoverPoint n-1 interoperability supports RecoverPoint splitters between releases.

EMC VMAX 10K (VMAXe) first supported the RecoverPoint splitter with Enginuity version 5875. N-1 interoperability support establishes the backward compatibility for the 5875 splitter and 5876 RecoverPoint splitter.

RecoverPoint and VMware integration

VMware vSphere 4.1 introduces vStorage API for Array Integration (VAAI). The VAAI commands speed up certain operations when a VMware ESX server writes to a storage array.

RecoverPoint and VMware integration fully support hardware-assisted LUN locking, which provides a more efficient mechanism in clustered host environments.

Federated Live Migration support

Federated Live Migration (FLM) is EMC's technology for seamlessly migrating a host application's volumes from a source array to any EMC Symmetrix VMAX array without requiring application downtime. The new array volumes emulate the old array volumes' WWIDs and geometry, so that the I/O hosts perceive the new array's paths to be paths to the old array. This allows multi-pathing software to seamlessly switch to the new array without any stoppage of I/O. Because FLM is transparent to the application, storage administrators can initiate the migration tasks without downtime, which significantly speeds up migration times.

The following table shows the EMC storage platforms that are supported as FLM targets. At this time, only DMX storage can be a source/donor for FLM migrations.

VMAX 10K (VMAXe)	VMAX 20K (VMAX)	VMAX 40K
✓	✓	✓

Table 4. Symmetrix storage supporting FLM

Enhancements to the Open Replicator for Symmetrix (ORS) feature set enables Federated Live Migration support for SCSI-3 host clustering. Prior to implementing this new feature, a host shutdown was required during FLM pull operations from devices with SCSI reservations before the Open Replicator session could be created.

This cluster support provides for SCSI-2 and SCSI-3 persistent reservations for Windows 2003 and Windows 2008 clusters. Microsoft Windows 2003 cluster support requires SCSI-2 non-persistent reservations, while Windows 2008 cluster support requires SCSI-3 persistent reservations.

Cluster support requires Enginuity 5876 and Solutions Enabler 7.4 and higher. FLM migrations with lower Enginuity versions continue to require non-clustered host environments. Figure 4 depicts an FLM configuration with cluster support.

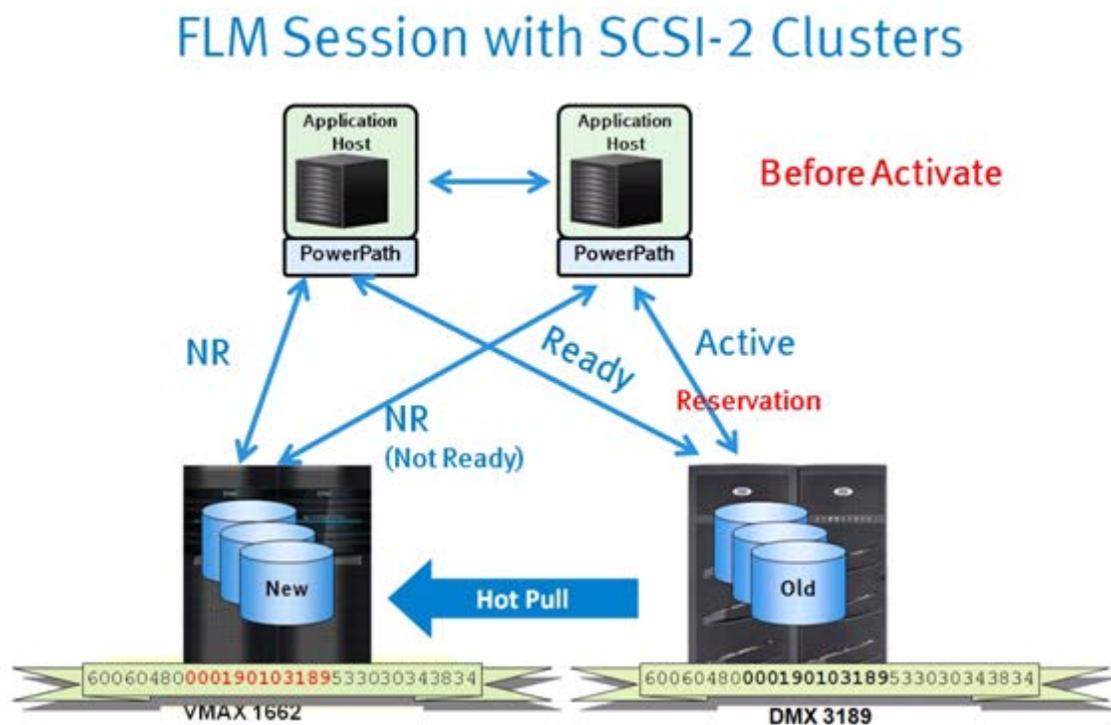


Figure 4. FLM session with SCSI-2 cluster

FAST VP enhancements

Fully Automated Storage Tiering for Virtual Provisioning (FAST VP) enhancements are available for the Symmetrix VMAX Family storage systems running Enginuity level of 5876 and higher and Solutions Enabler 7.4. Table 5 shows the supported Symmetrix VMAX Series storage systems.

VMAX 10K (VMAXe)	VMAX 20K (VMAX)	VMAX 40K
✓	✓	✓

Table 5. Symmetrix storage supporting FAST VP

FAST VP provides automated tiered storage for all EMC Symmetrix Family storage systems. It enables sub-LUN data movement for virtually provisioned storage devices, while providing dramatically increased capacity utilization and a reduction in time and complexity to manage storage.

VP allocation by FAST Policy

This feature allows new allocations to come from any of the thin pools included in the FAST VP policy that the thin device is associated with.

FAST VP attempts to allocate new writes in the most appropriate tier first. This avoids a double movement of initial allocation followed by storage tier promotion or demotion. If there is no information available to determine the most appropriate tier, then FAST VP allocates the extent from the pool the device is bound to.

If the pool chosen to allocate the data is full, FAST VP allocates to other pools contained within the FAST VP policy. As long as there is space available in at least one of the pools within the policy, all new extent allocations will be successful.

The allocation by policy feature is enabled at the Symmetrix array level and applies to all allocations for all devices managed by FAST VP. This feature is disabled by default. If the allocate by policy feature is disabled, new allocations will come from the pool the thin device is bound to.

SRDF coordination

Another core feature of Enginuity 5876 is the Symmetrix Remote Distance Facility (SRDF) enhancement to bring remote devices (R2) into consideration under a FAST VP policy. In Enginuity 5875, FAST VP operated independently on each side of the SRDF link. Promotion and demotion decisions made by FAST VP were based on the actual workload seen by each individual device.

While an R1 device would typically undergo a read and write workload mix, the corresponding R2 device would only see a write workload. (In 5875, reads against the R1 are not reflected across the link.). A consequence of this was that the R2 device data may not be located on the same tier as the related data on the R1 device.

Enginuity 5876 introduces SRDF awareness for FAST VP. The performance metrics collected for R1 devices are periodically transmitted across the link to the corresponding R2 devices.

On the R2 device, the R1 performance metrics are merged with the actual R2 metrics. FAST VP takes into account the workload on the R1 device and then makes promotion and demotion decisions for the R2 device data.

SRDF coordination is enabled, or disabled, per storage group. The default is disabled.

Hardware Related Enhancements

Green power supply

VMAX 10K (VMAXe)	VMAX 20K (VMAX)	VMAX 40K
N/A	N/A	✓

Table 6. Symmetrix storage supporting green power supply

EMC Symmetrix VMAX 40K uses a Climate Savers Gold energy efficient power supply for the D15 DAEs and engines. These power supplies improve the effective efficiency of the power delivery to meet 80 Plus gold standard for power supplies. Support for this green power supply also includes set up of adaptive cooling and power monitoring features.

Dynamic back-end improvements

VMAX 10K (VMAXe)	VMAX 20K (VMAX)	VMAX 40K
✓	✓	✓

Table 7. Symmetrix storage supporting dynamic back-end

Dynamic back-end is available for the Symmetrix storage systems specified in Table 10; however, these changes are not exposed to the user or other applications (Symmetrix Management Console, for example).

Ideally, configuration changes should cause absolutely no interruption to array processing, should be performed in parallel with other configuration changes, and should be finished in a short time. Previous Enginuity versions on EMC Symmetrix systems offered configuration changes that approached this methodology. Front-end mapping and masking, metavolume creation, and setting device attributes were defined as dynamic changes that used a new mechanism with no interruption to the storage array.

With Enginuity 5876, several back-end configuration changes are planned to take advantage of the dynamic methodology. Initial additions to dynamic capability are listed below, with more functionality scheduled for future releases.

As with other dynamic changes, the following back end operations cause no interruption, can participate in parallel operations, and are fast to complete. These improvements reduce the impact of configuration changes (including disk groups, physical drives, RAID groups, logical volumes, static RDF device associations, VLUN migrations, and sparing operations) without requiring an IML.

Compression for Fibre RA

VMAX 10K (VMAXe)	VMAX 20K (VMAX)	VMAX 40K
✓	N/A	✓

Table 8. Symmetrix storage supporting Compression for Fibre RA

This new 5876 feature allows the Fibre RA to use compression hardware to compress SRDF data for transfer over FC links. This capability is similar to the hardware compression available for the GigE RA. The same mechanisms are used to control the use of hardware compression, which is based on the SRDF dynamic parameter for compression for each RDF group.

SRDF compression is becoming more important for customers who increasingly deploy Symmetrix remote replication over longer distances. Currently, hardware compression is only supported on SRDF groups that are defined on GIGE directors.

With Engenuity 5876 and Solutions Enabler version 7.4, hardware compression for the Fiber RA is also supported on SRDF groups that are defined on Fibre RDF (RF) directors.

Miscellaneous VMAX 40K enhancements

Reporting

Solutions Enabler enhancements include the reporting of environmental data, such as fans and power supplies.

Dispersed system bay

The VMAX 40K storage system supports a dispersed system bay. The system bay may be split in two, separated by 25 meters. An even number of engines are in the first bay, while the remaining engines are in the second bay.

Using the Solutions Enabler CLI, the user can set the light to a blinking state on a dispersed system bay, as well as the drive bays. The Symmetrix Management Console (SMC) supports the ability to set these lights to a blinking state as necessary.

RVA Direct Member Sparing

Solutions Enabler enhancements provide reporting and events associated with the RVA direct member sparing capability. In drive-failure scenarios, the spared drive is added as another member of the Rebuild Group (RG), so in addition to rebuilding the data, users have the option to directly copy the data from the bad drive to the new drive.

Conclusion

EMC's latest hardware platform and software Enginuity features for the EMC Symmetrix VMAX Series storage systems support Cloud infrastructures, provide increased performance benefits, and more flexibility for users managing storage. EMC remains committed to green/eco-friendly products by continuing to provide features that support energy conservation.

References

Reference information and product documentation can be found at www.powerlink.emc.com including;

General

- *EMC Symmetrix VMAX 10K Series –Product Guide*
- *EMC Symmetrix VMAX 40K Series –Product Guide*
- *Implementing TimeFinder VP Snap for Local Replication*
- *Federated Live Migration Technical Overview –Technical Notes*
- *EMC RecoverPoint –Administration Guide*
- *FAST VP for EMC Symmetrix VMAX Theory and Best Practices for Planning and Performance –Technical Notes*
- *Symmetrix Remote Data Facility (SRDF) Product Guide*
- *Managing SAN-attached Disk Arrays with EMC Symmetrix Federated Tiered Storage*

Installation and Configuration

- *EMC RecoverPoint & the Symmetrix array-based Splitter*
- *EMC Solutions Enabler –Installation Guide*