



IDC Lab Validation Report

DELL EMC VMAX ALL FLASH

Essential Capabilities for Large Enterprise Mixed Workload Consolidation

By Eric Burgener, Research Director IDC Storage Team

Sponsored by Dell EMC | May 2017

Lab Validation Executive Summary

IDC validated six key features/functions of Dell EMC VMAX All Flash for Mission Critical Workloads:

- ✓ Introduction to VMAX All Flash Family
- ✓ Overview of Dell EMC Unisphere for VMAX
- ✓ Storage provisioning
- ✓ Hardware-assisted compression
- ✓ Point in time protection (SnapVX)
- ✓ Non-disruptive data migration (NDM)
- ✓ SRDF (replication, metro clusters)
- ✓ Transparent recovery from various failures



IDC Opinion

Across the six feature/function areas evaluated, IDC noted significant improvements in speed, ease of use, reliability, and cost-efficiency of operations that make a VMAX administrator's job notably easier. The storage management software released with the new VMAX All Flash portfolio make these systems the easiest VMAX systems to manage by far as well as expanding the performance and capacity capabilities of the line. IDC also validated other features (SRDF/Metro, Transparent Recovery) that further improve on the VMAX's "six nines availability" for which VMAX arrays are well known. These systems clearly have the features necessary to manage a mix of mission-critical enterprise workloads, now more cost-effectively than ever.

Dell EMC VMAX All Flash Family

The VMAX All Flash family is composed of 2 new arrays – the entry level enterprise 250F and the massively scalable 950F – with up to 4PB and 1M to 6M+ IOPS

- ✓ Much lower, more consistent storage latencies (90%+ lower)
- ✓ Higher throughput (IOPS) than VMAX 850F (>60%)
- ✓ Guaranteed data reduction ratios for mixed workloads of at least 4:1
- ✓ Much lower energy and floor space consumption (84%+ lower)

All comparisons relative to earlier generation VMAX hybrid arrays

Significant Flash Optimizations Improve VMAX Performance And Value



NEW

Array Family	250F	950F
# of V-Bricks	1-2	1-8
# Cores per System	96	576
Maximum # of Drives	100	1920
Maximum Capacity per Array (Effective)	1PBe	4PBe
Minimum capacity	11TBu	53TBu (Open/Mixed) 13TBu (MF)
Maximum IOPS	1 Million	6 Million +
Flash Drives Supported	960GB, 1.92TB, 3.84TB, 7.68TB, 15.36TB	

Technical data supplied by Dell EMC

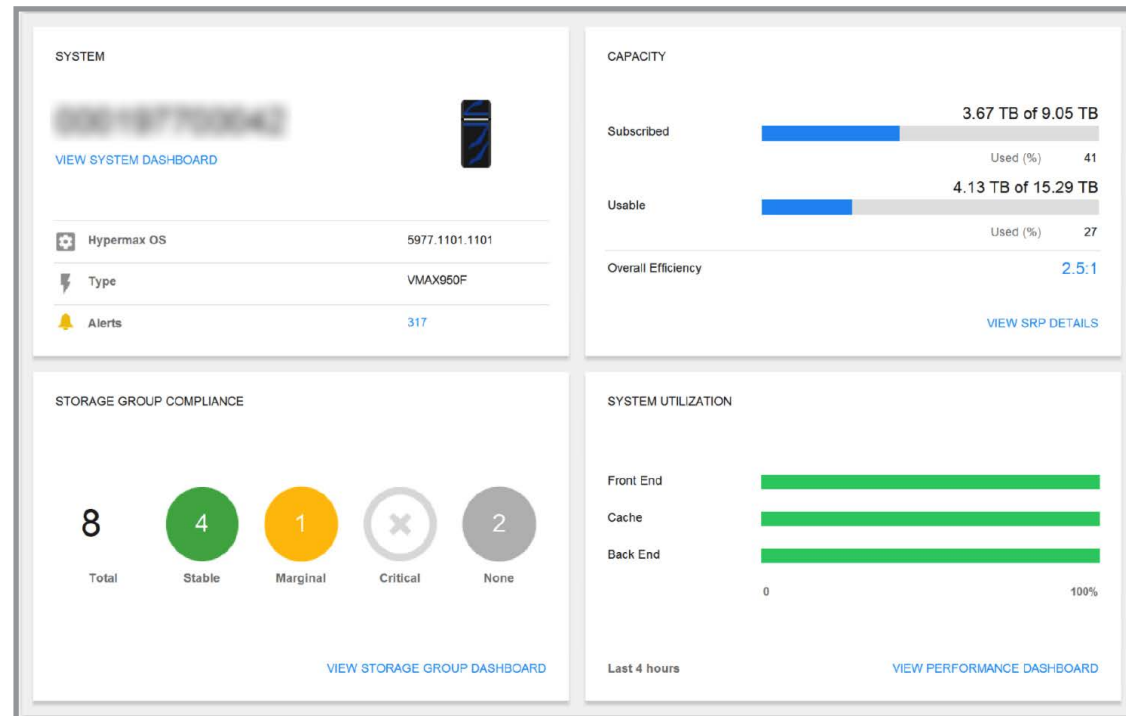
The VMAX All Flash family composed of 2 new arrays – the entry level enterprise 250F and the massively scalable 950F – with up to 4PB and 1M to 6M+ IOPS

Dell EMC Unisphere for VMAX

Unisphere provides a single, wizard-driven management domain for open and mainframe systems; the central dashboard shows “status at a glance” for the selected array with additional drill-down capability

The System panel shows the array ID number, the version of the HYPERMAX OS release, the type of array (here the VMAX All Flash 950F) and the top level compliance status with an overall array condition which can be either stable (green), marginal (yellow), critical (red) and not assigned (gray)

The Storage Group Compliance panel shows the total number of provisioned storage groups and what the service level objectives (SLOs) of the groups are (based on internal SLO latency expectations); here, 4 are meeting their requirements fully (green), 1 is not (yellow)



The Capacity panel shows the subscribed (allocated thin provisioned capacity) and the usable capacity with an Overall Efficiency (i.e. data reduction) ratio across all storage groups

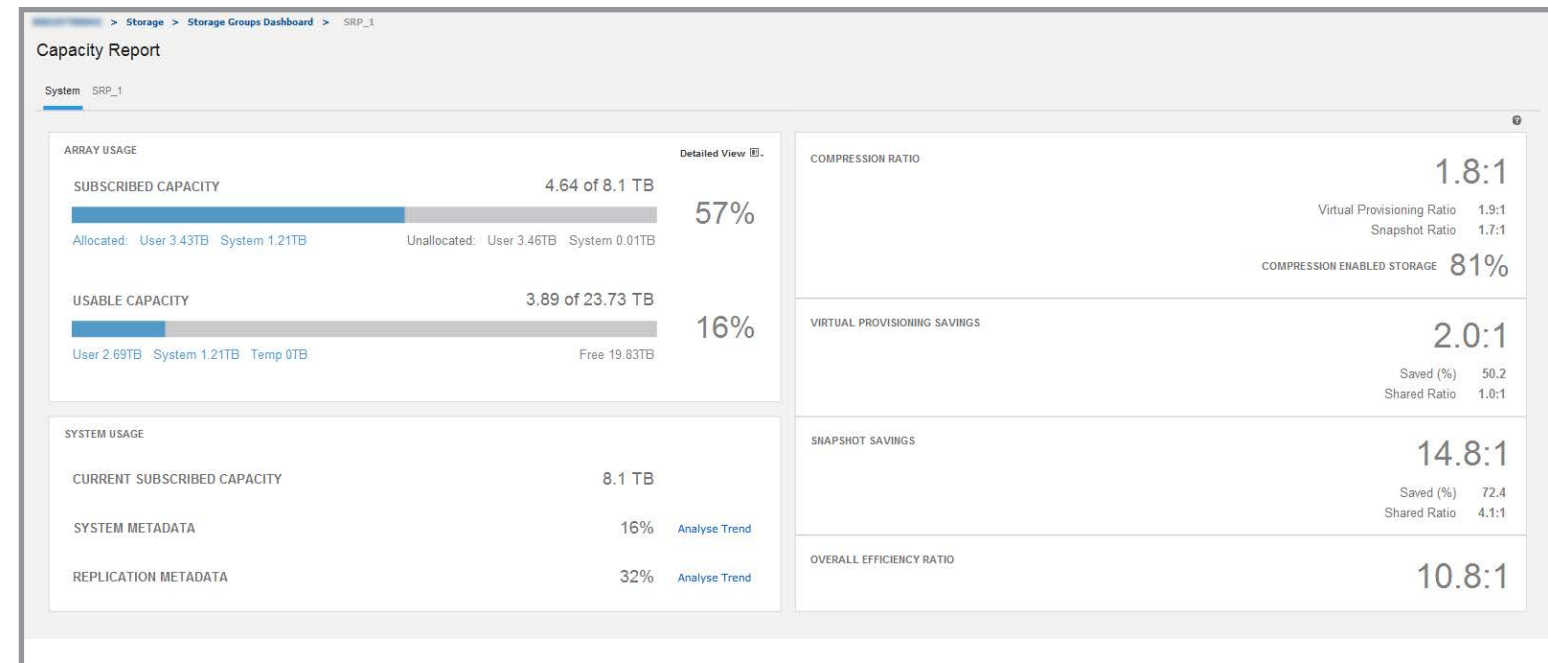
The System Utilization panel shows the percentage of front end (directors), cache, and back end (storage media) being used to process I/O; all green indicates that all resources are available and functioning normally

Dell EMC Unisphere for VMAX

The Capacity Report lists capacity consumption and usage as well as comprehensive space-efficiency metrics for the selected Storage Resource Pool

The Overall Efficiency panel provides add'l detail on contributors to the Overall Efficiency ratio, breaking out the contributions of adaptive compression, virtual (thin) provisioning, and snapshots, indicating what percentage of overall storage capacity in use is benefiting from adaptive compression (here, 97%), and what the percentage savings are due to both virtual provisioning and snapshots

The Array Usage panel shows the usable vs the actual subscribed capacity of the array for Storage Resource Pool SRP_1



The System Usage panel shows what percentage of the current subscribed capacity is used by system metadata and by replication metadata (if TimeFinder SnapVX is in use); it also provides an “analyze trend” option on metadata consumption

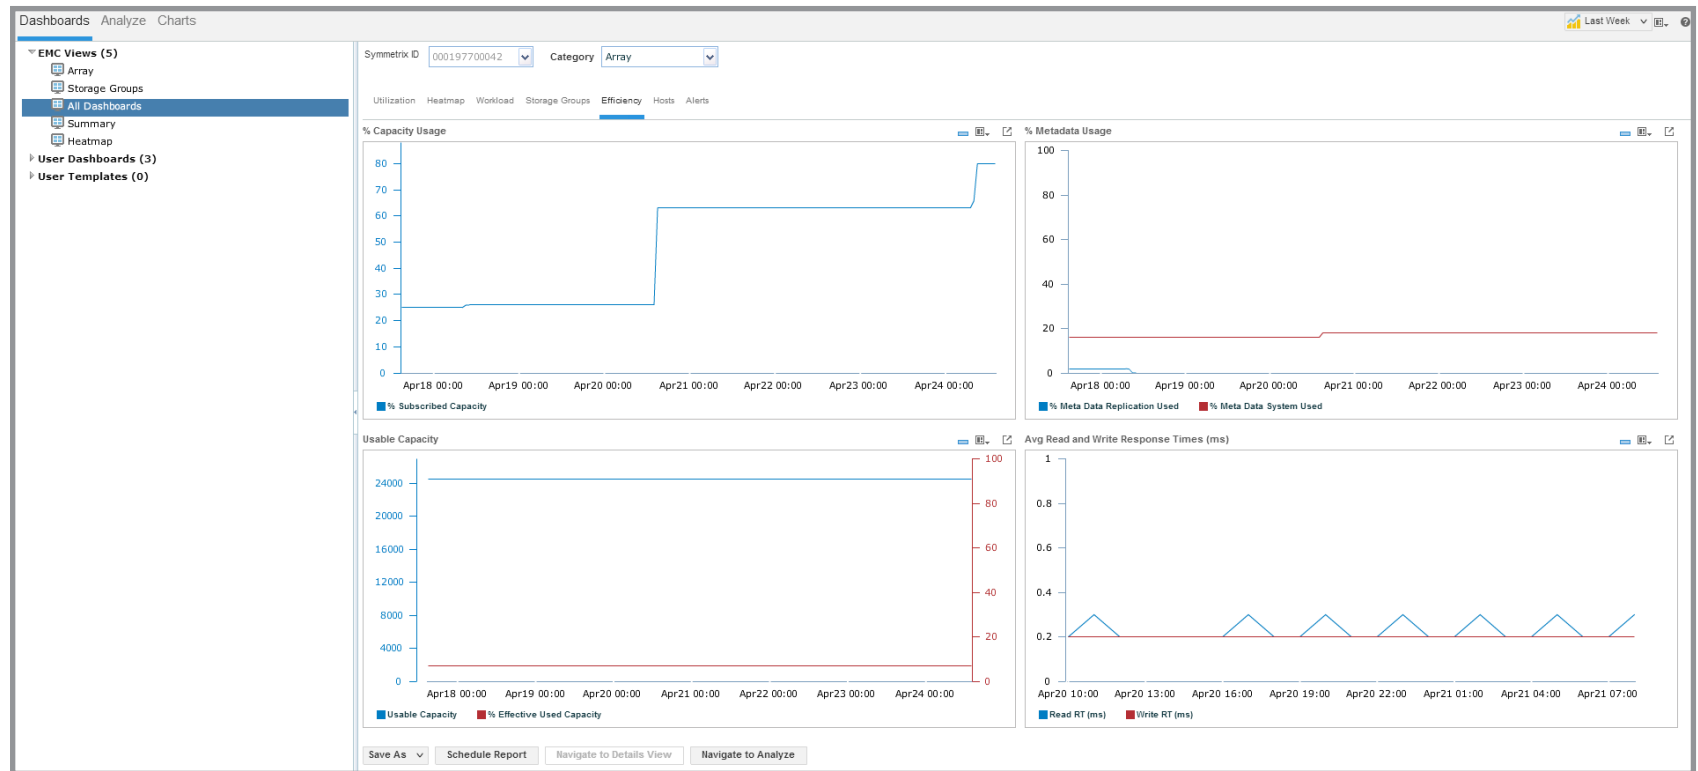
Dell EMC Unisphere for VMAX

Time-based trend analysis is available on a number of different metrics with a simple click

Add'l trending analysis at the Array level is available as a drill-down option

Here, we are showing 4 trend analyses over time – Capacity Usage, Metadata Usage, Usable Capacity, and Avg Read and Write Latencies.

These readouts are the primary statistics that let the administrator know the state of resource usage in the VMAX All Flash Array



Feature: Simplified Provisioning of Storage

Storage Manager Can:

- ✓ Provision and assign large amounts of storage in seconds
- ✓ Dispense with the use of spreadsheets to reduce storage provisioning times and improve reliability
- ✓ Use an automated Suitability Check to confirm that the configured system can meet the specified Service Level for the requested Storage Group

Improved ease of use



IDC Opinion

Storage provisioning times are reduced by over 90% versus the VMAX2. Provisioning reliability is improved through automation that is aware of host/port assignments. Unisphere can confirm the array's resources can successfully meet the service level of the storage as requested. These enhancements clearly make storage provisioning on the VMAX All Flash faster, easier, and more reliable.

Provisioning Storage

Selecting “Provision Storage to a Host” from among the top tabs on the home panel brings up the Provision Storage Wizard which can provision up to 200 volumes into a single storage group in about 30 seconds

Define the Storage Group name and assign it to an SRP (OS or MF in this system)

Define the Service Level (is always Diamond for All Flash, which means a service level objective of less than 1 millisecond latency), map it to a Workload Type (application), request the number of volumes and the capacity of each, and state whether compression should be used for the Storage Group

Select the host or host group which will use the Storage Group by selecting the relevant World Wide Name (WWN)

The wizard automatically checks the array's front end ports to determine which ports the host or host group is logged into and shows recent port utilization; the desired port(s) are then selected

Provision Storage

1 Create Storage

2 Select Host/Host Group

3 Select Port Group

4 Review

Storage Group Name: [] Storage Resource Pool: SRP_OS

Service Level: Diamond Workload Type: Not Specific Volumes: [] Volume Capacity: [] GB Avg. Resp. Time: 0.8 ms

Add Storage Group Total Capacity: 0.00 GB Total Service Lev: []

Set Host I/O Limits Compression

Next > Add to Job List Cancel Help

Provisioning Storage

The “Provision Storage to a Host” wizard automates much of the work admins used spreadsheets for in the past and includes a unique suitability check to confirm Service Level compliance for the provisioned storage

Review then brings up a summary screen which allows the admin to confirm their selections for the Storage Group creation

A unique feature allows the admin to run a “suitability check” to ensure that the specified Service Level can be met for the defined Storage Group; the suitability check shows the current load across the front end, back end, and cache and the anticipated new load with the best practice limit necessary to meet the defined Service Level

Once the admin determines that the Service Level can be met, they can either execute the Storage Group creation immediately or put it on the job list to be executed at some point later (for example, during a change window). Even if the workload analyzer indicates the workload may potentially exceed the performance envelope of the system the administrator can still allocate the capacity if they want to.

Provision Storage

Review & Finish

Masking View	Storage Group	Host	Port Group
Test_DB2_MV	Test_DB2	VHOL3049	Test_DB2_PG

Enable Compliance Alerts

System Suitability

✓ Storage creation on 2015/10/28/10:45 should meet the Service Levels

Best Practice Limit: 100%

Current Load: 22.30%
New Load: 23.75%

Front End, Back End

This System Suitability Check will check if the system can handle the new load being created by this Provisioning Request.

< Back Next > Add to Job List Cancel Help

Validated: VMAX All Flash Storage Provisioning Workflow

Feature Validation Summary: The new Provision Storage wizard reduces provisioning times by 90%+, automates time-consuming steps that in the past required separately maintained spreadsheets, and improves the speed, ease of use and reliability of storage provisioning.

Validation Process: IDC validated the new Dell EMC Provision Storage Wizard. We compared the time to create a single Storage Group with 200 volumes, assign it to a specific host, and confirmed the array's ability to deliver on the Diamond Service Level (latencies under 1 millisecond). The target array was a VMAX 950F, which is the new high end VMAX All Flash Family member. The new Suitability Check validated that the array could meet the specified Service Level automatically, another new feature that is not available on the flagship storage platforms of other established enterprise storage providers. Time to provision was 30 seconds, compared to a time of just over 12 minutes to provision this same storage under the older Dell EMC Unisphere for VMAX storage provisioning workflow.

Feature: Adaptive Compression

Highly efficient compression feature that works with ALL data services:

- ✓ The same trusted hardware compression used to compress SRDF data in-flight can now be applied to back-end data
- ✓ Compression engine resides in VMAX directors and compresses data as it is written to persistent storage
- ✓ Adaptive compression compresses less active data to keep access latencies to a minimum on active data
- ✓ Can be enabled at the Storage Group level at any time

Compression in VMAX All Flash lowers the effective storage costs

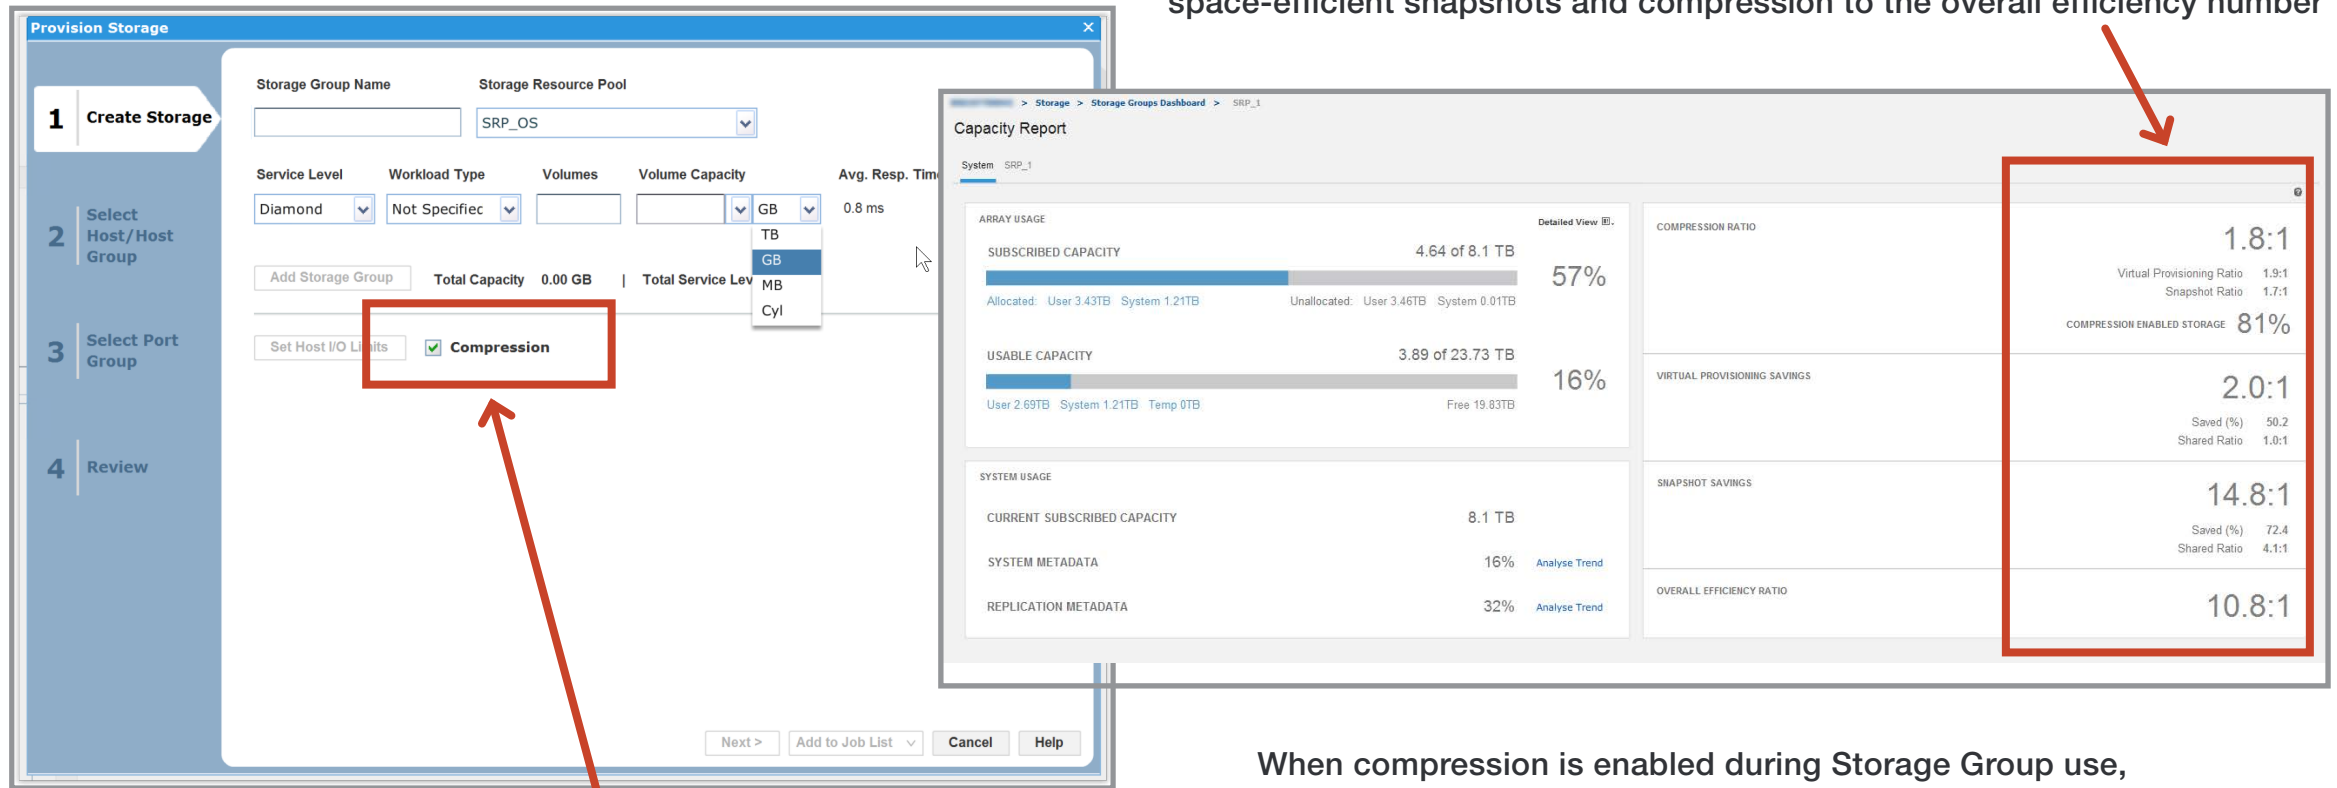


IDC Opinion

Compression in VMAX All Flash lowers the effective storage costs, while the adaptive implementation minimizes the latency impacts of these savings. In addition, granular enablement at the Storage Group level provides flexibility in use. With adaptive compression enabled, storage latencies increased by < 5% on average and were still consistently well under 1 millisecond.

Adaptive Compression

The Overall efficiency ratio is shown for the selected Storage Resource Pool (SRP_1), and broken down by the contribution of virtual (thin) provisioning, space-efficient snapshots and compression to the overall efficiency number



Compression can be enabled at Storage Group creation or at any time during a Storage Group's useful life

When compression is enabled during Storage Group use, it compresses the data as a background process so that real-time I/O latencies are not impacted

Adaptive Compression and Overall Efficiency

Dell EMC guarantees a 4:1 Overall Efficiency ratio for VMAX All Flash configurations in production use

- ✓ If that ratio is not achieved, Dell EMC will provide free add'l flash capacity to remedy the issue
- ✓ Unisphere provides data reduction detail by breaking space savings down by component (see slide 5) which includes compression (adaptive compression), virtual (thin) provisioning, and space-efficient snapshots
- ✓ The Overall Efficiency ratio (shown at the bottom of slide 5) accounts for the data reduction contribution of all storage efficiency technologies in use



This guarantee applies to the Overall Efficiency ratio and is subject to specific terms and conditions

Validated: VMAX All Flash Adaptive Compression Feature

Feature Validation Summary: We validated the granularity of the new compression for data-at-rest feature to be at the Storage Group / Application level, and turned the feature on and off for existing Storage Groups without any impact to Storage Group Service Levels over an observed 15 minute period with active I/O to the Storage Group. Adaptive compression is one of three storage efficiency technologies available on the VMAX All Flash (virtual (thin) provisioning and space-efficient snapshots are the two others). Achieved data reduction ratios just for adaptive compression are consistent with other Lempel-Ziv based lossless compression algorithms (in the range of 2:1).

Validation Process: IDC validated the ease of use and granularity of application of the new data-at-rest adaptive compression feature. We enabled compression during the creation of a new Storage Group (default behavior) and also selected a pre-existing Storage Group and enabled compression for that to view the impact. The target array was a VMAX 950F. Compression for the new Storage Group did not interfere with the ability of the Storage Group to meet its specified Diamond Service Level with an artificially generated workload over an observed 15 minute period. Compression for the existing Storage Group did not impact the Diamond Service Level either, although the entire Storage Group had not been compressed by the end of the 15 minute observation period. For this latter workflow, data is compressed as it is written to the Storage Group or compressed as a background process, so the time to compression varies with the capacity and write intensity against the selected Storage Group. During all adaptive compression testing, storage latencies stayed consistently well under 1 millisecond.

Feature: Point in Time Protection (SnapVX)

New SnapVX features include:

- ✓ Flash-optimized re-design uses redirect-on-write approach to enable higher performance and scalability with copy data management workflows
- ✓ Read/write snapshots reduce the number of steps required for common copy data management workflows

Improved performance, scalability,
and storage efficiency



IDC Opinion

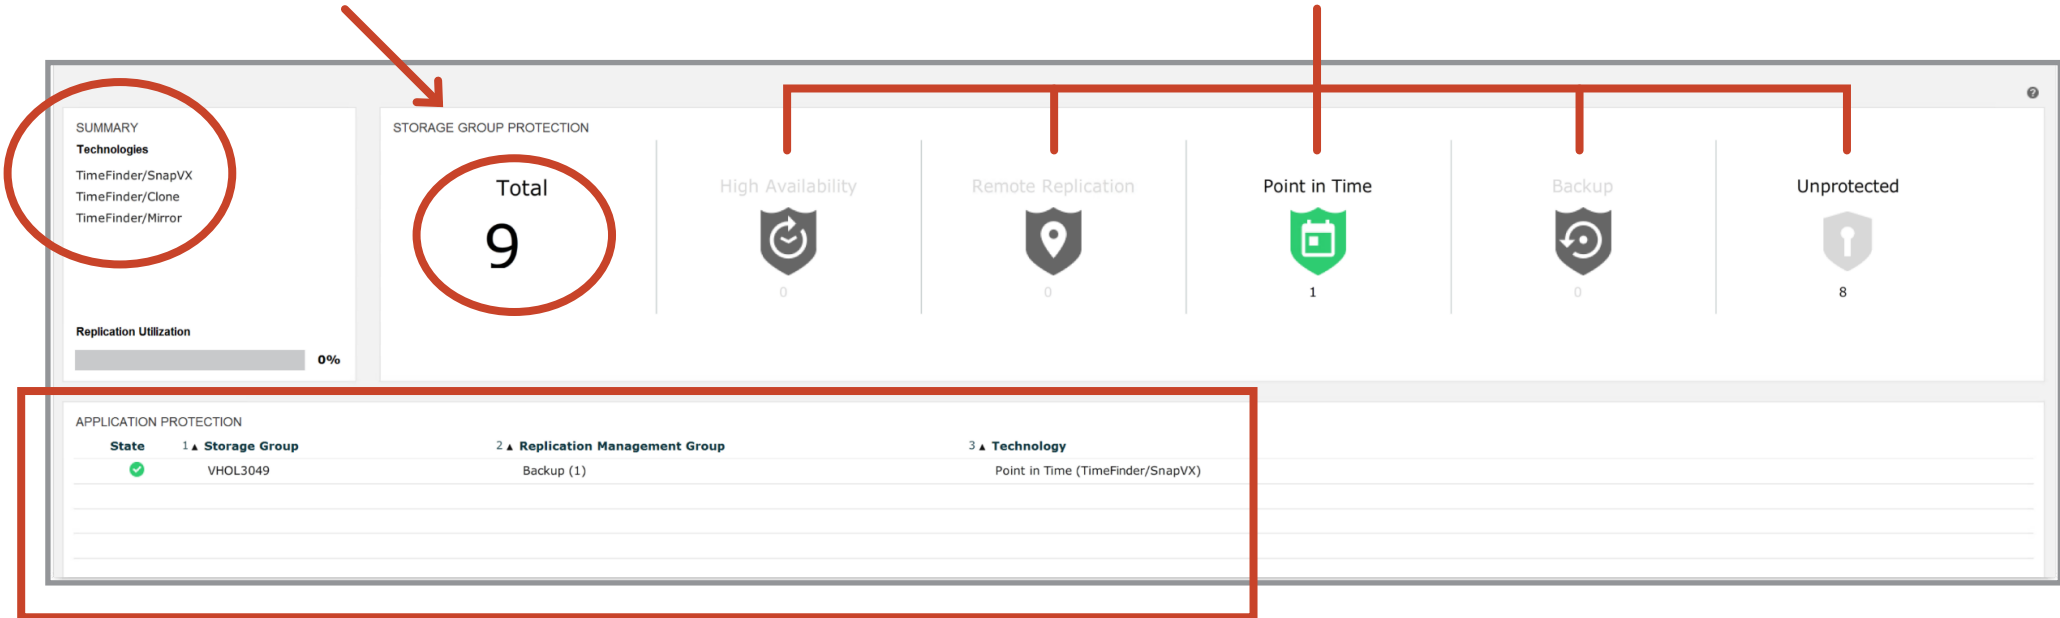
The Snapshot Creation Wizard is simple and intuitive, leverages snapshots that use a design that is better suited for common high activity copy data management workflows than copy on write approaches, and enables the immediate creation of space-efficient new snapshots or snapshots of existing snapshots. SnapVX includes the features administrators and application developers need in a snapshot implementation that will be used for copy data management tasks on a large enterprise array that is used for general purpose mixed workload consolidation.

Selecting Protection Levels

Snapshots are enabled as a type of “Protection Level”. There are four types of protection levels that can be selected against a Storage Group or a Storage Group can be specifically left Unprotected (no SRDF or SnapVX). The four types of Protection are High Availability (SRDF Metro), Remote Replication (SRDF, non Metro), Point in Time (SnapVX), and Backup (ProtectPoint). Please note that a single Storage Group can be protected by combinations of protection types.

Shows how many provisioned Storage Groups are being Protected

When a Protection Level is selected, the technologies in use are shown here; TimeFinder SnapVX is typically used for VMAX All Flash workflows



Storage Groups are volume groupings associated with applications. For each Storage Group, it shows the state of the Storage Group, the name, and any relevant Protection Levels applied, which in this example include a Storage Group (used for Backup) that is set up under the Point in Time Protection Level using SnapVX snapshots.

Protect Storage Group - Snapshot Creation

Point in Time Protection was selected from the previous screen, which brings up the Protect Storage Group Wizard

Admin defines whether this is a new snapshot or they will re-use an existing one, names the snapshot, and defines snapshot expiry (None means it remains until manually deleted). Customer can also select Time to Live which specifies a time when the snapshot will auto-delete

Advanced options include Both Sides (in SRDF sync any snapshot is created consistently on both sides) and, new with the Q2 2017 Software Release (SR), "Secure" (which prohibits snapshot deletion by users before its defined expiry time, useful in addressing hacking and ransomware attacks).

Details are then confirmed, and the snapshot is either created or added to the Job List for future creation

The screenshot shows the 'Protect Storage Group - vhol3050_Data' wizard. The interface is divided into three main sections on the left: '1 Select Protection Type', '2 Snapshot Details', and '3 Finish'. The '2 Snapshot Details' section is currently active. In the main content area, there are two radio button options: 'Create New Snapshot' (selected) and 'Reuse Existing Snapshot'. Below these is the 'Expiry Type' dropdown menu, currently set to 'None'. A link '<< Hide Advanced' is visible. A 'Secure' checkbox is checked. To the right, there is a 'Name' field with the value 'snap_1' and a 'Days' field with a dropdown arrow. Below the 'Days' field, there is a note: 'Secure snapshots may only be terminated after they expire or by customer-authorized EMC support.' and an 'Hours' field with a dropdown arrow.

Validated: VMAX All Flash SnapVX Feature

Feature Validation Summary: We validated the ease of snapshot creation along with the options available during creation, which included Type (new or re-use) and Expiry (Time to Live). We validated the SnapVX implementation (re-direct on write instead of the older copy-on-write) which enables increased performance and scalability for common high activity copy data management workflows. We did not confirm snapshot performance in usage, but note that two features of the snapshot implementation – re-direct on write architecture and the fact that mounted snapshots are mapped directly to all their data blocks (rather than through other snapshots) – will allow target volumes to deliver the same performance as source volumes. We also validated the new reporting capabilities which include status at a glance for all Protection Levels with drill-down capability for additional details within each Protection Level type.


Validation Process: IDC validated the speed and ease of use of snapshot creation. We validated that snapshot creation was space-efficient, observed the flexibility with which snapshots could be assigned to different hosts for various workflows (such as recovery or Dev/Ops), and confirmed that snapshots could be created off a single or multiple volumes (depending on how a Storage Group is defined) literally within just a few seconds without consuming any additional storage capacity. Storage capacity is used, however, when snapshots are written to (which we validated does not require separate clone creation when high performance is required as in the past). We also confirmed this implementation could take snapshots of snapshots and enable intermediate snapshot deletion (in a chain) without impacting the availability of any other snapshots in the chain.

Feature: Non-Disruptive Migration (NDM)

NDM features and enhancements:

- ✓ Enables easy, non-disruptive migration using reliable SRDF infrastructure
- ✓ Can migrate individual applications without taking downtime on the application
- ✓ Includes simple one-click “revert” option
- ✓ Can migrate from VMAX1/2 arrays to VMAX All Flash (or VMAX3) arrays and maintain SRDF/S if in use on source

Promotes high availability with flexibility across various data migration workflows



IDC Opinion

The new NDM workflow requires 65% fewer steps than the previous version, and removes the need to refer to external spreadsheets for port and host configuration data, making the process faster, easier, and more reliable. Automation preservation of the DR configuration allows regulated customers to seamlessly migrate while maintaining a synchronous remote replication at all times. SRDF/S can be maintained on the Source Array if needed. A new one-click “revert” option makes that process much easier when needed, and fulfills a common customer request to automate reversion in the middle of a non-disruptive data migration.

Data Migration Wizard

From the Storage Groups screen, select a Storage Group and launch the Data Migration Wizard

Note the capacity, number of volumes, and masking views of the Storage Group to be migrated; the FAST policy refers to the storage tiering policy for a Storage Group being migrated from a Hybrid Flash Array (which will no longer apply if the target is an All Flash Array)

In the Data Migration Wizard, select a target VMAX array for the migration and confirm the Storage Resource Pool to which it should be migrated

The screenshot displays the EMC Unisphere for VMAX v8.3.0.1 interface. The main window shows the 'Storage Groups' screen with a table of storage groups. The 'WM_T22' storage group is selected. A 'Data Migration Wizard - WM_T22' dialog box is open, showing the '1 Destination' step. The 'Symmetrix' dropdown is set to '000197000002' and the 'SRP' dropdown is set to 'SRP_1'. The wizard has three steps: 1 Destination, 2 Create or Prepare, and 3 Summary. Navigation buttons at the bottom include '< Back', 'Next >', 'Add to Job List', 'Cancel', and 'Help'.

Name	FAST Policy	Capacity (GB)	Volumes	Masking Views
B245	—	200.02	6	1
B246	—	12	4	1
▼ GA_DEMO	—	0.04	6	1
a215_gks	—	0.04	6	1
Mig1_sq	—	50	2	1
▼ RT_COMP	—	4	2	0
RT_COMP_1	—	2	1	0
RT_COMP_2	—	2	1	0
WM_T22	BETTER	101	1	1

Data Migration Wizard

In the Create or Prepare step, identify the Storage Resource Pool, whether the Storage Group should be compressed on the target, the masking views, port and host groups

Confirm selections in the summary step

Add to Job List (which includes a Run Now option)

The screenshot displays the EMC Unisphere for VMAX v8.3.0.1 interface. The main window shows a table of Storage Groups with columns for Name, FAST Policy, Capacity (GB), Volumes, and Masking Views. The 'WM_T22' storage group is highlighted. A 'Data Migration Wizard - WM_T22' dialog box is open, showing the 'Summary' step. The wizard summary includes the following details:

- Destination:** Summary page for creation. Running this wizard section will create the migration session for the selected source Storage Group.
- Storage Group:** WM_T22
- Symmetrix:** 000197000002
- SRP:** SRP_1
- Compression:** Yes
- Masking Views:**

Masking View	Port Group	Host/Host Group
LCSEE164_MV	LCSEE164_PG2	LCSEE164

At the bottom of the wizard dialog, there are navigation buttons: < Back, Next >, Add to Job List (with a dropdown arrow), Cancel, and Help.

Significant Under the Covers Automation Leveraging SRDF

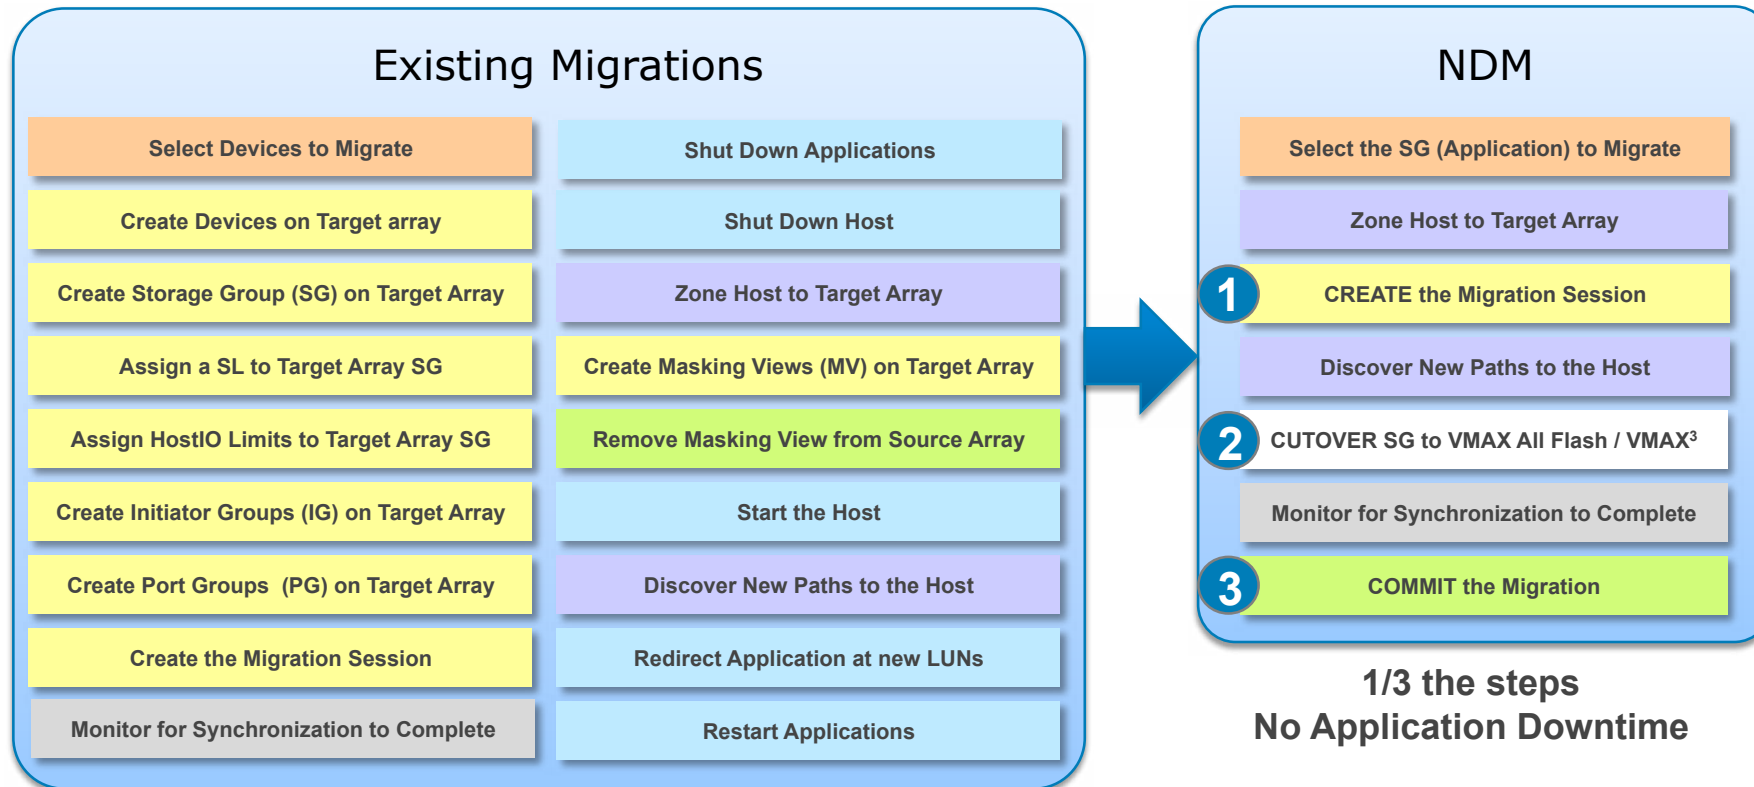
- ✓ Under Create, previously manual steps are automated (45 steps down to 17), including creation of the migration session, discovery of new paths on the target, mirroring of source zoning on the target, creation of target Storage Group, Service Level for the target is specified, and SCSI reservations and identity of source transferred to the target
- ✓ A non-disruptive re-scan is manually performed on the relevant host(s) to discover the target array
- ✓ Under Cutover, RDF connectivity is verified, arrays are run in a synchronized state but the paths to the Storage Group(s) being migrated from the source are inactivated, and invalid tracks on the source are automatically cleared as data is migrated
- ✓ Fully compliant with Powerpath and other selected multi pathing software, does not require separate SRDF licenses, and supports the same distance as SRDF Metro (up to 100 km)
- ✓ Under Commit, the RDF pair is terminated, identity of source is set to a new WWN (target now has original source identity in case it is re-deployed), and all Storage Group I/O is now being serviced solely by the target

Comparison

- With NDM, data migration requires 65% fewer steps than non-automated methods and reduces time and risk (there is no longer any need to reference separately maintained spreadsheets)
- With NDM, there are only three steps (create, cutover, commit) and this version includes a “one click” automated reversion workflow
- Under non-automated method, any host reboots or reversions were fully manual
- NDM is faster to setup, more reliable, and offers additional flexibility

Comparison of Data Migration Steps on VMAX (Before NDM and with NDM)

NDM offers simplified migration



**1/3 the steps
No Application Downtime**

LEGEND: Host operations | Array operations



Validated: VMAX Non-Disruptive Migration

Feature Validation Summary: We validated the new data migration workflow with 65% fewer steps and no references required to external spreadsheets, observed storage services being continuously available to an application running on a host during the migration of a Storage Group which represented 101GB of data on the source, and observed a successful reversion back to the original source VMAX, executed by a single click when the migration session was partially completed without incident.

Validation Process: IDC validated the operation of the Data Migration Wizard through the creation of a migration session between two local VMAX arrays connected over SRDF links and the several states (cutover-ready, cutover, and commit). The cutover-ready step performed a non-disruptive re-scan on the host to discover the target array and established synchronized I/O for the two arrays. The cutover step completed the SRDF synchronization of all data being migrated, invalidated the source array I/O paths to the Storage Groups being migrated, and established the target as the active array. The commit step stopped SRDF replication between the arrays, terminated the RDF device pairs, unmasked the source array, set the identity of the target array to be the same as the original source, and then assigned a new identity to the source array. **In the wake of migration, we confirmed all data had migrated and that the application with active I/O experienced no downtime during the migration. A second migration session was established for the express purpose of observing a reversion in the middle of a data migration, which successfully returned everything to the previous state with the source array servicing all active I/O without any disruption.**

Feature: SRDF/Metro

Simple SRDF/Metro creation workflow:

- ✓ Enables easy, non-disruptive creation of zero RTO HA configurations using reliable SRDF infrastructure
- ✓ Supports active/active Storage Groups and enables zero impact cluster host, application or virtual machine failover to an alternate VMAX
- ✓ Supports 3-way SRDF with far distant async remote site

Delivers highest levels of availability
with improved ease of use



IDC Opinion

The Protect Storage wizard makes SRDF/Metro creation fast and easy, requiring only six steps and automating much of the process which in the past required manual involvement for improved speed and better reliability. Different configuration options supported by the wizard provide the flexibility admins need in setting up these environments to meet various needs. SRDF/Metro uses proven SRDF infrastructure for synchronization, and supports failover options that were validated in VMware vSphere environments.

SRDF/Metro Configuration from the Protect Storage Group Wizard

Once the Protection Type is selected (High Availability Using SRDF/Metro), a target VMAX is selected, the SRDF is automatically selected, a witness type is selected (to resolve split brain scenarios), the Remote Storage Group is named, and a Service Level is selected (this will always be Diamond for an All Flash Array)

Selections are reviewed under Step 3, and an admin can then choose to either add it to a job list for future completion or execute creation immediately (Run Now)

Protect Storage Group - Metro_Cluster

1 Select Protection Type

2 Select SRDF/Metro Connectivity

3 Finish

Remote Symmetrix: [REDACTED]

SRDF Group: Automatic Selection

Establish Pairs: Yes

Managed By: Use Witness

Remote Storage Group Name: Metro_Cluster

Service Level Name: Diamond

< Back Next > Add to Job List v Cancel Help

Add to Job List

Run Now

Validated: Simplified SRDF/Metro Configuration

Feature Validation Summary: We validated the creation of an SRDF/Metro configuration between two local VMAX arrays, reviewing the speed with which it can be created and the ease of use offered by the Protect Storage Group wizard for establishing SRDF/Metro configurations. We reviewed the ability to provide an active/active data store across two local VMAX arrays in a VMware vSphere environment and the ability to failover a virtual machine in the event of an array failure. While we didn't test the 3 way SRDF/Metro solution we acknowledge this as a valued feature for customers with the highest levels of availability or compliance requirements.

Validation Process: IDC used the Protect Storage Group wizard to create an SRDF/Metro configuration for a single Storage Group. The configuration was established between a VMAX 250F and a VMAX 950F (two local arrays connected over 2 SRDF links running over FC). With this latest release, the creation workflow has been reduced to six clicks in a wizard which automates the required object creation on the target, making the operations not only faster but more reliable. We selected "Run Now" to create the configuration immediately, and creation began as a background process which did not interfere with the source array's ability to meet its Diamond Service Level during synchronization operations, which happened within minutes across the SRDF links for a 1TB Storage Group. From within vSphere, we performed a failover of a single VM from the original source array to the target array, which successfully re-established the VM in the new location in less than 2 minutes with no data loss. We also recognize that this functionality is applicable to clustered database use cases and that there are a number of VMAX customers using it with Oracle RAC and other clustered databases.

Feature: Transparent Recovery

Each VMAX Engine includes two directors (controllers) which operate in an active/active mode and ensure that the architecture suffers from no single point of failure

- ✓ Continuous operations w/o data loss from a total local RAID group failure with SRDF/Synchronous
- ✓ Transparent recovery w/o data loss from a director failure
- ✓ Failed resources can be replaced without downtime

Provides transparent recovery from various hardware failures



IDC Opinion

The active volume worked flawlessly during these fault injection tests, both upon initial failure and after re-establishing the failed resource. Fault injection included a dual local drive failure and a director failure. I/O was not interrupted, no data was lost, and in this case (with a relatively light workload) there was no long term impact to overall system performance (although that will vary based on how heavily loaded a system is when the failure occurs, and how many Engines are deployed in the system – the 950F supports up to 8 engines). Throughout the entire fault injection test, the array(s) continued to meet its specified Diamond Service Level.

Exploration of I/O Impact on Dual Local Drive Failure in a RAID 5 Configuration

- ✓ An SRDF/Synchronous configuration was set up using a VMAX 950F and a VMAX 250F where the data was mirrored to both locations; SSDs in each array were separately protected by RAID 5
- ✓ An artificial workload was generated to flow continuously against a volume that was mirrored to the two arrays with SRDF/Synchronous; this workload can be characterized as “light” as it was under 20% of each of the array’s rated performance capabilities
- ✓ Two SSDs in a single RAID 5 Raid Group in the VMAX 950F were simultaneously failed to create a dual drive failure scenario; the “failure” was validated running the Solutions Enabler command `symdisk -sid 0177 list -failed`. In addition, the failure can generate an Alert in Unisphere and will generate a dial home
- ✓ The impact on I/O was observed by continually watching the throughput and latency against the SRDF mirrored volume using iometer; upon failure, I/O was not disrupted as I/O to the failed devices continued to be handled by the remote mirror (through reads across the SRDF Link) with no impact to storage latency or throughput during this failure, and throughout the test the volume continued to meet its Diamond Service level
- ✓ This is significant because this type of dual failure in a stretched cluster configuration could lead to an outage whereas here the system just continues to access the data using the mirrored volume on the target array
- ✓ The “failed” SSDs were turned back on, and with no impact to storage latency or throughput the system performed a background resync (quite short since there was less than 5MB of writes to the mirrored volume during the outage) and returned the volume to a fully synchronized state across the VMAX 950F and the VMAX 250F

Exploration of I/O Impact on Director Failure

- ✓ An artificial workload was generated to flow continuously against the VMAX All Flash 950F, which was configured with two VMAX Engines (four directors); in normal operation an engine operates in an active/active mode with each director servicing roughly half the existing I/O workload
- ✓ A director failure was simulated by manually shutting down one of the directors
- ✓ The “failure” was validated by the `symcfg -sid 042 list -dir all` which indicated that the Director was offline. In addition, the failure can generate an Alert in Unisphere and will generate a dial home.
- ✓ At the time, the VMAX 950F had relatively low utilization (under 25% of its rated capability) based on the artificial workload it was handling; the impact on I/O was observed by continually watching the throughput and latency against the VMAX using iometer; upon failure, I/O was not disrupted but the throughput and latency showed a momentary latency impact in the range of 20-30% but still did not exceed the 1 millisecond response time required by the Diamond Service Level, and within several seconds returned to a steady state at roughly the same level as prior to the failure
- ✓ The failed director was turned back on, and within seconds the workload was re-distributed across both directors without any interruption in I/O
- ✓ The impact of a director failure on overall throughput and latency will vary depending on how heavily loaded a system is; because the VMAX was lightly loaded in this case, a single director could continue to service the I/O workload without a performance change but in a more heavily loaded VMAX the ultimate impact will vary based on the level of utilization; in either case, however, the I/O flow will not be disrupted

Validated: Transparent Recovery from Various Types of Hardware Failures

Feature Validation Summary: We tested the ability of an active volume running on a VMAX system to transparently recover from two failure types: a dual local drive failure in an SRDF/Synchronous configuration and a single director failure in a dual engine array. For the SRDF/Synchronous configuration we used a VMAX 950F and a VMAX 250F. During the failures, we observed the impact on I/O throughput and latency. We then brought the “failed” devices back on-line to observe the impact on I/O throughput and latency as the workload on the system returned to steady state (which for the dual drive failure required re-synchronization across the two VMAX arrays and for the director failure required the re-distribution of the workload across both directors). The VMAX 950F on which we performed the Director failure test had 2 Engines (a total of 4 Directors).

Validation Process: IDC validated that a VMAX active volume will transparently recover from both of these types of failures without interrupting storage services or any data loss, continuing to service I/O; for these fault injection tests, we established a light artificial workload for the array(s), manually failed the SSDs and the director (but not at the same time), observed the impact on throughput and latency, then turned the failed device(s) back on and observed the impact on throughput and latency and the time to return to “steady state” performance.

Validation Process

IDC performed the validation at Dell EMC's Labs in Hopkinton, MA. The storage test bed consistent of several VMAX systems, including a VMAX 250F and a VMAX 950F (both All Flash Arrays).

IDC Lab Validation Methodology

This Lab Validation Report provides a summary of an extensive validation process performed by IDC in collaboration with the supplier's teams. IDC relied on the supplier's equipment, facilities and their configuration to perform this validation. All of the tests were conducted during the presence of one or more IDC Analysts.

This report is meant to provide a quick set of inferences and insights for IT professionals and business decision makers seeking to perform further due diligence on the capabilities of the product and/or services that have been validated in this Report. However, the goal of this Report is not to supply detailed hands-on test plans and validation jobs. It is not meant to replace the evaluation process that most businesses will conduct before making any decision to purchase the product and/or services.

It is for this reason that this Report is not designed to be an all-inclusive document on all the capabilities of the product, but rather as a concise document that highlights select features/functions of products, their relative performance with respect to a traditional environment and the value these features bring to businesses looking to solve problems of scale and availability in large scale enterprise storage deployments.

Finally, even though this Report is a sponsored document, it is not meant to be an IDC endorsement of the product, service or the sponsoring supplier. IDC's opinions are its own and not influenced by the production of this document.

Validation Test Bed

This table provides a summary of test environment details for each feature validated.

Function	Components	Configuration Notes
Dell EMC Unisphere for VMAX	<ul style="list-style-type: none"> • VMAX 950F w/HYPERMAX OS 5977 Q2 SR (Beta) • Dell EMC Unisphere for VMAX 8.4 	VMAX 950F was configured with multiple volumes against which I/O workloads had been running using iometer and various VMAX features had been enabled; new reporting features which provided drill down detail on overall efficiency, subscribed vs usable capacity, and metadata consumption.
Storage Provisioning	<ul style="list-style-type: none"> • VMAX 950F w/HYPERMAX OS 5977 Q2 SR (Beta) • Dell EMC Unisphere for VMAX 8.4 	A Storage Group was created in the Open System Storage Resource Pool on the 950F using the Provision Storage wizard.
Adaptive Compression	<ul style="list-style-type: none"> • VMAX 950F w/HYPERMAX OS 5977 Q2 SR Beta • Dell EMC Unisphere for VMAX 8.4 • X86 Red Hat Enterprise Linux for load generation • vdbench was used for load generation 	A Storage Group was created with adaptive compression enabled, and adaptive compression was also enabled for a pre-existing storage group.
Point in Time Protection (SnapVX)	<ul style="list-style-type: none"> • VMAX 950F w/HYPERMAX OS 5977 Q2 SR Beta • Dell EMC Unisphere for VMAX 8.4 	VMAX 950F had multiple volumes in a Storage Resource Group, and SnapVX snapshots were created, stored, and used for single and multi volume snapshots, and Unisphere was used to review relevant new statistics on SnapVX usage.
SRDF/Metro	<ul style="list-style-type: none"> • VMAX 950F w/HYPERMAX OS 5977 Q2 SR Beta • VMAX 250F w/HYPERMAX OS 5977 Q2 SR Beta • 2 x CISCO UCS Server C240 w/ESX 6.5 	An SRDF/Metro configuration was set up for a single 1TB volume group between a VMAX 950F and a VMAX 250F using a local 8Gb/s Fiber Channel network. VMware Virtual Machines were failed over between arrays.
Non Disruptive Migration (NDM)	<ul style="list-style-type: none"> • VMAX 950F w/HYPERMAX OS 5977 Q2 SR Beta (Target) • VMAX 40K w/Enginuity 5876 + ePack (Source) • CISCO UCS Server U240 w/ESX 6.5 	NDM comes standard with UniShpere for VMAX. No additional software needs to be installed nor is additional Dell EMC licensing required. Select volume groups were migrated between a VMAX 950F and a VMAX 40K over a Fiber Channel network.
Transparent Recovery	<ul style="list-style-type: none"> • VMAX 950F w/HYPERMAX OS 5977 Q2 SR Beta • VMAX 250F w/HYPERMAX OS 5977 Q2 SR Beta • CISCO UCS Server U240 w/ESX 6.5 	A mirrored volume was set up across the two VMAX arrays using SRDF/Synchronous. Data was protected locally in each array using RAID 5. Each VMAX had two V-Brick engines (each of which had two directors).

Essential Guidance

Advice to Buyers: Customers looking to replace their primary storage platform which today is running a mix of workloads that include their most mission-critical applications should absolutely be looking at All Flash Arrays (AFAs). Performance, scalability, availability, enterprise-class data services, easy manageability, and overall maturity are all critical purchase criteria for datacenters looking for an enterprise storage platform that will support business growth over the next several years. Increasingly, customers expect those capabilities to be available with highly flash-optimized enterprise storage platforms.

Key areas for administrators to explore and understand when evaluating mission-critical storage platform alternatives include common workflows (like storage provisioning), snapshots, replication, and fault management. Other considerations which affect overall cost include storage efficiency technologies like compression and the ease with which data can be migrated, both for technology refresh and for easy integration into hybrid cloud environments.

As storage management tasks migrate away from dedicated storage administration groups more towards IT generalists, ease of use is critical. The use of wizard-driven workflows with built-in validation and other forms of automation are critical in maintaining the reliability of IT operations as businesses continue to grow. Certain tools are required for other critical workflows, such as snapshots for data protection and test/dev, and replication for disaster recovery. Vendor implementations should be evaluated for ease of use, flexibility to meet different workflow requirements, space-efficiency, and ability to support advanced configurations (like stretched clusters). For systems to achieve “six nines availability” they must offer “defense in depth”, and supporting technologies include hot plug components, RAID, space-efficient snapshots (to support stringent RPO/RTO requirements), flexible replication, and other features which support non-disruptive operations (for handling failures, upgrades, data movement, etc.). Data reduction features which help increase effective capacity while at the same time lowering overall storage costs without impacting storage latencies are a must-have. Easy-to-use tools which identify the contributions of each storage efficiency technology help IT staff to make better informed administrative decisions about how best to manage their storage resources.

IDC Opinion

IDC believes that these capabilities must be a core part of an enterprise storage infrastructure strategy, and urge prospective customers to make them important purchase criteria as they look to replace aging storage platforms that are used for the consolidation of mission-critical and other workloads. Enhancements evaluated in this lab validation make the current generation VMAX All Flash arrays significantly easier to manage, including features that improve the speed, reliability, and cost-efficiency of operations. The new VMAX All Flash portfolio provides an expanded range of performance and capacity scalability, is the easiest VMAX generation to manage by a wide margin, includes additional dashboard metrics that inform improved administration, and clearly meets the enterprise requirements for hosting a mix of mission-critical workloads.