EMC VMAX 10K Enhancements

Silverton Consulting, Inc.
StorInt™ Briefing

Introduction

EMC® has been an acknowledged leader in enterprise storage systems for almost two decades now. Over all that time EMC's Symmetrix® storage product line has exemplified rock solid availability, high input/output (I/O) performance and state of the art, advanced storage functionality. Not long ago, Symmetrix underwent a revolutionary re-incarnation that advanced well beyond the prior generation's Direct Matrix (DMX) design to the Virtual Matrix, or VMAX®, architecture of today. EMC Symmetrix VMAX continues their long tradition of enterprise storage excellence started with the very first Symmetrix system back in 1988.

Recently, EMC has enhanced their VMAX 10K storage system, the entry-point to the current Symmetrix enterprise class product family. The new VMAX 10K offers almost all the capabilities of the higher-end VMAX 20K and 40K storage systems but at a lower TCO with lower maximum capacity. The VMAX 10K permits mid-tier, open sytems enterprise customers to enjoy the same great availability and functionality of EMC's VMAX architecture at a more accessible cost.

EMC Symmetrix VMAX 5-9s+ availability

All VMAX storage systems are designed around a scale-out, clustered architecture with a high performance, dedicated interconnect that helps to

provide higher system availability, well beyond traditional dual controller designs. VMAX clustered system design supplies additional redundant hardware that complements the already existing high availability (HA) and data integrity characteristics available from Symmetrix Enginuity microcode. For example, HA capabilities such as VMAX's global storage cache, data integrity field check codes, and distributed architecture help to provide enhanced storage availability that exceeds 5-9s (99.999%) system availability, as measured by current field deployments.

A VMAX storage cluster contains one or more VMAX engines, each of which includes two storage directors. The VMAX architecture employs a unique and innovative global, storage system wide cache, whereby each VMAX engine contains a coherent copy of all the write data in other VMAX engines. This is done for a number of reasons but from an HA perspective, any VMAX engine director can take over the I/O activity of any other failing director, without data loss, even if the director is located in another engine.

Next, the VMAX system automatically appends a data integrity field (DIF) or check code to every block written to the storage. This DIF is present end-to-end (from front-end port to disk) and verified whenever data is moved throughout the system, such as when it is read into cache for host retrieval. Such capabilities enable VMAX to detect data corruption in real-time and potentially, correct it on the fly before the customer accesses their data. Also, similar capabilities are used to validate addressing and control information on all VMAX control paths. DIF checking on system addressing and control can often detect and recover from faults even before they impact customer data. While all enterprise class storage uses

some form of DIF, many other designs only verify data integrity at the back end (to the disk) instead of end-to-end.

Moreover, VMAX uses a distributed architecture to insure that no single port, director, or disk enclosure (DAE) outage can affect data availability. If a director does fail, the storage workload is distributed across the remaining directors to minimize any impact to performance.

Equally important, VMAX storage utilizes its redundant hardware and multiple instances of environmental monitoring/system management software to fence off sub-component failures from system activity. This allows VMAX storage to isolate failures to a single hardware (sub-) component and route subsequent activity around the fault providing continuous storage availability in the event of a component outage. In many other storage systems, the failure of one component shuts down the entire controller.

In addition, Symmetrix VMAX also offers online code updates that preserve all host/server IO activity while upgrading system microcode. This is akin to changing all four tires on a race car in the middle of a race, without a pit stop. It takes less than a minute to upgrade the Enginuity microcode across an entire VMAX array. All directors/all engines are initial microcode loaded (IMLed) with the new code, in one single, non-disruptive operation. Most other storage requires some down time to upgrade code, often calling for operators to take each controller offline to update its code before doing the same to its other controllers.

Finally Oracle and others have recently defined a new data integrity solution called Block Guard that insures data integrity all the way from the host to the storage system and back again. This new T10 committee Protection Information industry standard together with Oracle's Data Integrity Extension

(DIX) and VMAX storage provides an end-to-end data integrity check. Similar to DIF discussed previously, the DIX is created in the server O/S, checked at the host bus adapter (HBA) and then validated again at the VMAX storage front-end to verify that data has been transmitted properly. On the return path, the VMAX creates and appends the DIX to a block being read which is then checked at the HBA and at the O/S to verify data transmission integrity. Just as the VMAX uses end-to-end integrity checking within the array, Block Guard and DIX within Oracle extend data integrity assurances from the server all the way to the VMAX disks and back.

EMC Symmetrix VMAX core functionality

In addition to the excellent SAN block storage services utilizing Fibre Channel (FC), FC over Ethernet (FCoE) and/or iSCSI protocols, the VMAX product family also offers many advanced storage features. Some of these include:

- **Virtual Pools (VP) –** advanced storage thin provisioning that allows hosts to define more storage than is physically present by sharing empty, non-used capacity with other volumes. Storage is physically allocated only as it is actually consumed, reducing required raw capacity and costs substantially.
- **Fully Automated Storage Tiering for VP (FAST VP) –** automatic storage tiering that uses Enterprise Flash Drives (EFDs) together with additional tiers of spinning media to provide higher application performance at lower cost than achievable by rotating disks alone.
- **TimeFinder® -** local, on the subsystem, storage copy and clone capabilities that together with FAST VP provides a quick, space-efficient method to copy customer data for test/dev or data protection services.

- **Federated Live Migration (FLM) –** volume relocation capability that moves customer data volumes from one storage system to another for technology refresh and storage consolidation.
- **Symmetrix Replication Data Facility, SRDF®** - remote mirroring capability that helps data centers to recover from site failures by providing a consistent and coherent copy of data at remote sites. SRDF can operate both synchronously and asynchronously and has long been the "gold standard" of array-based replication.
- **Unisphere for VMAX** – host based administrative and management tool to provision, monitor and tune VMAX storage systems using a GUI.
- **PowerPath®** - host based multi-pathing facility that not only provides automatic alternate path failover but also can be used to supply concurrent, multi-path I/O activity to EMC VMAX storage.

There are many more advanced capabilities than what are listed above. All current VMAX storage systems can be purchased with the above functionality. However, VMAX 20K and 40K go well beyond these capabilities, e.g., FICON® mainframe attach, more storage connectivity, and higher levels of scalability and capacity.

EMC Symmetrix VMAX 10K

EMC introduced the original VMAX 10K storage system in 2011 as an entry-point to the VMAX product family. At that time, the VMAX 10K supported a subset of full Symmetrix functionality. Recently, EMC has updated the VMAX 10K with new hardware, enhanced functionality and significantly more configuration flexibility that more closely match VMAX 20K and VMAX 40K storage system capabilities.

New EMC Symmetrix VMAX 10K system hardware

The new VMAX 10K hardware continues the impressive evolution of Symmetrix storage that has been going on now for well over two decades. EMC has updated a number of system components to increase I/O performance and broaden configuration capabilities.

First, the new VMAX 10K can be configured with up to 1560 drives, 44% more than the previous maximum. Also, the new VMAX 10K now offers 50% more compute cores than the previous generation system, with up to 6 physical cores per director (12 cores per engine). Each compute core is now faster at 2.8GHz instead of 2.4GHz. Thus, the new VMAX 10K is capable of delivering more I/O and storage services than the older system.

In addition, the new VMAX 10K supplies faster internal and inter-cluster data transfers utilizing the latest Gen2 PCIe and Virtual Matrix RapidIO® interconnect fabric technologies. As such, the new data transfer hardware substantially increases data throughput for the VMAX 10K over the previous generation.

Furthermore, the new VMAX 10K now supports both large form factor (LFF) 3.5" and small form factor (SFF) 2.5" drives. The 2.5" drives can come packaged in new high-density, 2U, 25 drive enclosures such that a VMAX 10K cabinet can now hold twice as many usable drives as previously available. Also the 2.5" disk drives consume 1/3rd less power/cooling and weigh 1/3rd less than the 3.5" disk drives, providing more energy efficient storage with reduced floor loading. With the more dense drive enclosures, the new VMAX 10K can be configured with as few as 24 to as many as 1560 drives to support 1.5PB of usable storage. Furthermore, EFD, 2.5" and 3.5" disk drives can now be intermixed in the same cabinet for additional configuration flexibility.

Moreover, the new VMAX 10K can now be ordered in EMC supplied racks or in compliant customer supplied 19" cabinets. As such, customers now have the option of installing their VMAX 10K storage hardware in their own racks or cabinets, allowing them to standardize their datacenter infrastructure where required.

New VMAX 10K performance

With all the new and enhanced features, VMAX 10K storage performance improves substantially. The new VMAX 10K provides 100% more online transaction processing (OLTP) I/O operations per second (IOP/s) than the previous generation system. Further, the new system supports 30% more back-end bandwidth than previously available to improve data throughput intensive applications. As such, Oracle on VMware and virtualized OLTP-like applications now execute up to 90% faster on the new VMAX 10K storage systems.

New EMC VMAX 10K software features

In addition to all the new hardware, VMAX 10K software functionality has also been enhanced and now includes some capabilities previously only available on the larger VMAX systems. For instance, EMC Data at Rest Encryption (D@RE), Federated Tiered Storage (FTS) and full SRDF support to name just a few are now all available on VMAX 10K storage, making more true Tier-1 features available to mid-tier customers.

Data at Rest Encryption

EMC Data at Rest Encryption (D@RE) is a capability that secures customer data residing on drives in VMAX storage and is now available on the new

VMAX 10K. With D@RE functionality and encryption-enabled VMAX engines, customer data on all VMAX 10K disks is automatically encrypted using AES-256 so that data cannot be read without its cryptographic key. There is no effect on performance and no external encryption switches are required. EMC D@RE functionality fully meets and has been validated to FIPS 140-2 standards for data security.[1]

Moreover, D@RE key data can be supplied by an RSA Key Manager (RKM) external platform or alternatively, supplied internally by an embedded key manager within the VMAX 10K. In either case, all key data is secured and only available to authorized users, helping customers meet regulatory compliance and data privacy requirements.

Federated Tiered Storage

In addition, Federated Tiered Storage (FTS) provides storage virtualization and tiering for heterogeneous, external storage systems. The prior generation of Enginuity introduced FTS for VMAX 20K and 40K, and the latest version adds support for this feature for VMAX 10K. With FTS, the manageability, and functionality of VMAX 10K can be extended to external data storage. In fact, a VMAX checksum is appended to FTS and validated when read to include data integrity checking on external storage.

Also, FTS on any VMAX now supports using external storage as an additional tier in a FAST VP multi-tiered environment. Previously, FTS employed external storage as a third tier behind internal EFD, fast and slow disk storage, but with the new enhancements, external storage can be utilized as any tier. As

[1] The FIPS 140 certification (#1610, dated September 26, 2011) can be viewed here: http://csrc.nist.gov/groups/STM/cmvp/documents/140-1/140val-all.htm

such, if a customer has invested in an all-flash storage array, VMAX 10K can take advantage of its increased performance by using it as the highest performing tier in a FAST VP storage pool.

RecoverPoint and SRDF

RecoverPoint provides heterogeneous data replication and protection services for VMAX 10K storage. The VMAX 10K has the EMC RecoverPoint write-splitter integrated directly into the Enginuity microcode and hence, no longer requires a separate fabric splitting switch. This makes it much easier to deploy EMC RecoverPoint for VMAX 10K data protection.

Also, the new VMAX 10K now offers 2-site and 3-site SRDF configurations for open systems. With VMAX 10K using SRDF/STAR, three-way data replication, both within region and out-of-region disaster recovery capabilities are fully supported.

Other new VMAX 10K features

Furthermore, a pair of advanced storage capabilities has been added to all VMAX systems, including the VMAX 10K. Specifically,

- **Host I/O Limits** - this provides a quality of service capability for VMAX 10K on a storage or port group basis and can be ideal for cloud/multi-tenant application environments which need guaranteed service levels for client activity. With Host I/O Limits, the VMAX prevents any application or server from using more IOP/s or MB/sec performance than it was authorized to use. Using VMAX Host I/O Limits for multiple clients insures that no one application or user can starve others of their pre-allocated performance.

- **Data compaction –** this provides data compression to free up VMAX storage capacity and executes in the background on cold, idle data. This facility can optionally be used across any or all VMAX storage volumes. Data compaction makes the VMAX even more storage efficient, often storing twice as much idle data in the same storage capacity, thus substantially reducing the cost of storage.

VMAX 10K strong solution support

The VMAX 10K provides excellent host solution support. For example:

- **VMware –** VMAX 10K's full vSphere 5.1 management console integration uses vCenter plug-ins to help VMware administrators monitor and manage VMAX storage. Also, VMAX 10K fully supports the vSphere SAPI set including VAAI, VASA and SRM for increased VMware inter-operability using native VMAX functionality. In the near future, VMAX 10K will be supported by vCenter Operations Manager via an EMC Adapter, taking advantage of the new model for analytics by workload, capacity, and health.
- **Microsoft Hyper-V and Windows 2012 –** VMAX 10K's EMC Storage Integrator (ESI) provides application-aware provisioning using storage specific wizards to help administrators configure and manage VMAX storage under Microsoft Windows physical or virtual environments. Further, ESI can operate as a Microsoft Management Console snap-in or as a standalone Windows application. VMAX is also qualified with Windows 2012 (Windows 8) Server.
- **Microsoft Exchange Server, SQL Server and SharePoint – VMAX 10K's** FAST VP provides better I/O performance for Exchange and SQL Server storage while EMC SharePoint adapters offer specialized

services to ease the provisioning and management of SharePoint data for VMAX storage systems.

Unisphere for VMAX

EMC's Unisphere for VMAX storage management software has recently been updated to include full provisioning, monitoring and management of VMAX storage features.

Four hours to first host I/O

Installation, configuration and use of VMAX 10K storage has been fully optimized for ease of use and takes ~four hours to install and provision storage for host use. Most of that time is spent unpacking, installing and IMLing the hardware for the first time with only about four minutes devoted to configuring and provisioning host storage for I/O.

Summary

In short, all EMC Symmetrix VMAX storage systems supply leading high availability, data integrity checking and system fault tolerance that are state of the art in the IT industry today. These capabilities will help keep EMC VMAX storage operating long after other systems would have gone offline in the face of hardware outages. When system failures do occur, VMAX storage will detect and isolate them faster, before they can do any more harm, without impacting storage availability.

Moreover, the VMAX 10K's new hardware enhancements provide impressive I/O performance gains, roughly doubling the previous generation system's OLTP capabilities and 30% better system throughput. With its superior

performance and increased configuration flexibility, it's quite possible that VMAX 10K could function as your data center's enterprise storage for a long time to come.

When one adds all that to the VMAX 10K's equally impressive functionality enhancements such as D@RE, FTS and significant VMware/Microsoft solution support, the EMC VMAX 10K becomes an even more compelling enterprise storage solution. In the end, the new VMAX 10K system provides an excellent entry point to the VMAX storage family, that all by itself may be superior to other vendor's enterprise storage solutions available on the market today.

Silverton Consulting, Inc. is a Storage, Strategy & Systems consulting services company, based in the USA offering products and services to