WHITE PAPER

IBM DS8870 in the mainframe market- 2014

An Analyst Product and Market Review

By Josh Krischer,
Josh Krischer & Associates GmbH
May 2013, updated November 2014

Table of Contents

IBM DS8870 in the mainframe market

An Analyst Product Review

Josh Krischer is an expert IT advisor with more than 45 years of experience in high-end computing, storage, disaster recovery, and data center consolidation. Currently working as an independent analyst at Krischer & Associates GmbH, he was formerly a Research Vice President at Gartner, covering enterprise servers and storage from 1998 until 2007. During his career at Gartner responsible for high-end storage systems and spoke on this topic at a multitude of worldwide IT events, including Gartner conferences and symposia, industry and educational conferences, and major vendor events.

Executive Summary

Much is being published on various aspects of storage systems, their features, connectivity to servers, etc. Never the less very little is written on the storage platform on which the vast majority of world's crucial information is stored. Ninety-six of the world's top 100 financial institutions, most of the largest industry corporations, government organizations and many other entities, use IBM z Systems™ z as their major IT platform. This paper is about storage for IBM z Systems.

Three vendors manufacture high-end enterprise storage systems that connect to z Systems: EMC with VMAX2, Hitachi Data Systems with VSP G1000 (also OEM-ed by *Hewlett-Packard* Company), and IBM with the DS8000 series. All three systems are multi-platform products; however this paper will concentrate on features support, synergy, and compatibility with z Systems. The latest EMC high-end product, the VMAX^3, currently does not support FICON channels. According to IDC tracker, in 3Q 2014 the DS8000 had the highest market share for z/OS storage, 44%, EMC 29 and HDS 24%. HP's market share on z Systems is negligible. IBM is the architecture owner, and the two other companies provide z Systems compatible products. All the systems above are multi-platform products; however this paper will concentrate on features support, synergy, and compatibility with z Systems.

Storage Requirements and Challenges

Despite the weak world economy, the market for storage systems is not showing any indications of slowing down. It shrunk a little in 2009, but fully recovered in 2010 and 2011. In 2013 and 2014 storage growth was slowed down by the growing usage popularity of "storage economic" features such as thin provisioning, tiered storage, de-duplication and compression.

Availability and Business Continuity remain on top of the list of requirements. The nonstop global economy, fierce competition, and new levels of service requirements raise the requirements for business continuity. In the last decade, midrange storage users requested advanced functionality and flexibility, which allow them to increase utilization, decrease storage management efforts, and deploy disaster recovery schemes.

Scalability is a must. A storage system should be able to scale seamlessly in capacity, connectivity or performance, without decreasing service levels. It should support multi-tier storage media including Solid State Drives (SSDs), performance Hard Disk Drives (HDDs) in different capacities (Fibre Channel or SAS), and large capacity near-line Serial-Attached SCSI (SAS) or Serial Advanced Technology Attachment (SATA)] disks. In simple words, scalability should support tiered storage "in a box."

Performance has two aspects: the first is the throughput measured in number of I/Os per second (IOPS), and the second is the response time measured in milliseconds for hard drives. Performance should meet Service Level Agreement (SLA) requirements, regardless of the used capacity and the workload. Erratic performance levels irritate users more than a slightly slower but constant response times.

Advanced Functionality is required for better resource exploitation and to address storage management. The average organization's storage capacity grows by 20-40 percent per year while the size of the storage management staff generally remains the same. Data centers which supported terabytes at the end of the previous decade today support petabytes. The only way to cope with this capacity explosion is sophisticated functionality, advanced automation, and user-friendly management tools and interfaces.

Storage Efficiency in usage and energy consumption is equally important. It can be achieved with thin provisioning, tiered storage, automated data placement, deduplication, compression, small-form-factor HDDs, SSDs and virtualization. Efficient storage systems provide better storage utilization, which translates to lower capacity and lower capital and operational expenditure (CAPEX and OPEX), reduced floor space requirements, lower energy consumption, and improved administrators' efficiency.

The three-level processing structure and the powerful POWER7+ processors make the DS8870 the performance leader. The synergy between the DS8870 and z Systems brings additional performance improvements. The full redundancy design, embedded encryption, powerful and versatile data mirroring techniques and other high reliability design consideration ensure business continuity and data protection. Building the storage hierarchy with automated data tiering, small factor disk, near-line capacity disks SSDs, flash cards and a user-friendly GUI contribute to storage efficiency.

IBM's System Storage DS8000 series is a stable multiplatform high-end storage system able to support current and future user requirements. It uses state-of-the-art technology to achieve top performance levels and supports superior local and remote data mirroring techniques to ensure business continuity. The DS8870 offers unique features to createsynergy with z Systems which is explain further in the document. The DS8870 has all the required characteristics to be positioned on the top of enterprise storage systems for all platforms and in particular for z Systems.

IBM DS8870 Structure

Processors

IBM's System Storage DS8000 was announced in October 2004. This disk system uses a storage architecture based on standard components such as the IBM POWER+ processors also used in IBM Power Systems (formerly pSeries) (see Fig.1). The DS8000 was a follow-on to the Enterprise Storage Server (ESS) – codenamed "Shark", which was announced in July 1999. The DS8000 series uses three levels of processing: PowerPC and application-specific integrated circuits (ASICs) in the front and the backend adapters and POWER7+ as the main controller in a server-symmetrical multi-processing (SMP) design. The latest model of the IBM System Storage DS8000 series uses dual multi-core processors SMPs. The POWER7+ (4.228GHz) 2, 4, 8 and 16 core powerful processors supporting simultaneous multi-threading (SMT). One of the enhancements of the POWER7 and POWER7+ processors is the Simultaneous Multi-Threading (SMT4) mode. SMT4 enables four instruction threads to run simultaneously in each POWER+ processor core to maximize its throughput.

"The DS8000 series uses three levels of processing PowerPC and ASICs in the front and the backend adapters and POWER7+ as the main control in a server-symmetrical multi-processing SMP design."

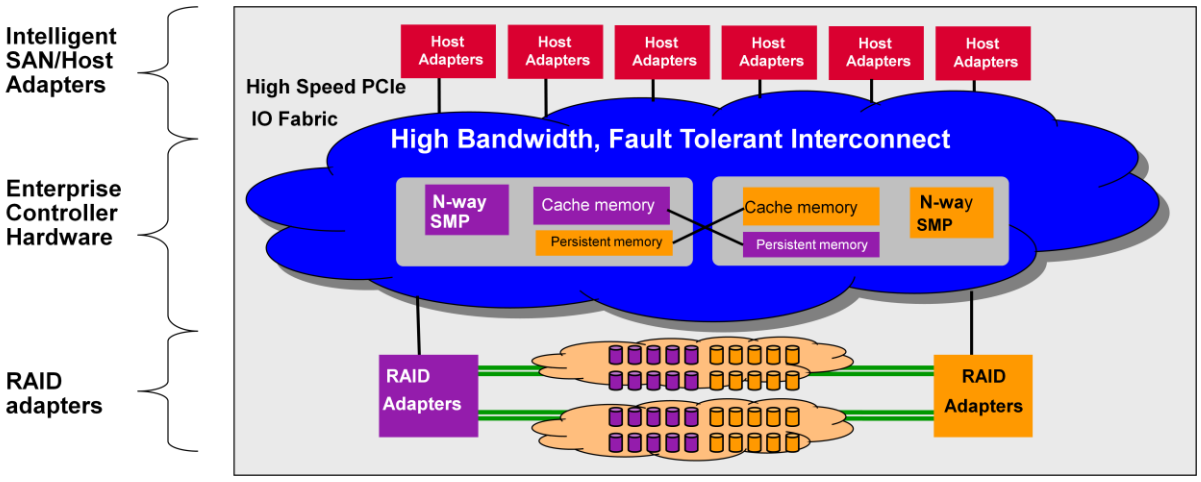

Figure 1. DS8870 High-end Storage System Block Diagram. Source IBM

These multithreading capabilities improve the I/O throughput of the DS8870 storage system. As opposed to common applications processing, where execution of an instruction may depend on a previous result, in the storage control unit the tasks are almost completely independent, which contributes to very high utilization of the multi-core, multi-thread processors.

Cache and Non Volatile Storage

Instead of using a dedicated cache, the DS8000's series cache is allocated as part of the IBM Power server memory. The Power server has a two-level cache (L1 and L2) in addition to its main memory, which creates three levels of memory hierarchy. The tightly

clustered SMP, the processor speed, the L1/L2 cache size and speed, and the memory bandwidth deliver better performance in comparison to a dedicated, single-level cache. Each side of the cluster has its own cache and the Non-Volatile Storage (NVS) on the other cluster side to provide better data integrity protection in case of a side failure. During normal operation, the DS8000 series preserves a copy of writes in cluster cache and a redundant copy of writes using the NVS copy in the alternate cluster. In case of extended loss of external power, batteries provide sufficient power to de-stage write data from cache to special disks reserved for that purpose. This “cross connection” protects write data against loss in case of power loss or other malfunction, even when one of the cluster sides suffers from a severe malfunction. The effective cache size approximately equals the installed capacity. The DS8000 series uses 4 Kbyte cache slots, which results in improved cache utilization compared to larger cache slots used in competitive high-end disk systems. It improves performance in particular for On-line Transaction Processing (OLTP) of databases, mail servers and most of the Enterprise Resource Programs (ERPs).

Cache Management Algorithms

In July 2007 IBM shipped a new microcode release with a significant improvement to the cache pre-fetching algorithm. The new pre-fetching sequential algorithm is called *Adaptive Multi-stream Pre-fetching* (AMP). AMP is applied on top of *Sequential Prefetching in Adaptive Replacement Cache* (SARC) [1] and significantly improves performance for common sequential and batch processing workloads. AMP optimizes cache efficiency by incorporating a self-optimizing pre-fetching algorithm, which is very responsive to changing workloads and delivers up to a two-fold increase in the sequential read capacity of RAID-5 arrays, including for smaller-than-maximum configurations. This technique also further reduces “pollution” of the cache with unnecessary data compared to usual pre-fetch algorithms.

Another feature is the Intelligent Write Caching (IWC), a.k.a. Wise Ordering for Writes. IWC optimizes management of write cache, helping to increase write I/O throughput to disk and reduce contention between read misses and write de-stages, thus enhancing overall I/O performance.

Front-end

The DS8870 supports from 2-to-32 host adapters (installed in pairs) and up to 128 Fiber Channel/IBM FICON host ports (four or eight host ports per adapter). Ports can be configured as 4 Gbps or 8 Gbps Fiber Channel host connections. This option allows connection flexibility and lower cost by saving ports. The ports can independently auto-negotiate to transfer data to match link speed depending on the host HBA. These ports are initially defined as switched point-to-point Fiber Channel protocol (FCP) however,

[1] Sequential Prefetching in Adaptive Replacement Cache (SARC) algorithm, which is a self-tuning, self-optimizing cache management technique for a wide range of sequential or random workloads

each port can also be set to either FCP, remote mirroring port or FICON, which provides investment protection for the users.
The adapter itself is driven by PowerPC processor and function-rich, high-performance, specialty ASICs installed in a PCIe Gen 2 slot.

Back-End

The system can support up to 16 4-port, 8 Gbps Fiber Channel adapters. The connection with the HDDs/SSDs is through 6 Gbps point-to-point switched SAS-2 connection to an 8 Gbps Fiber Channel backbone.

The host adapters include PowerPC processors and ASICs. IBM leveraged this strength by using function-rich/high performance ASICs. These ASICs are responsible for managing, monitoring, and rebuilding the RAID arrays. The symmetric dual path structure allows access to the disk enclosures from two independent networks, providing four access paths to each disk drive.

Scalability

The DS8870 [currently] supports from 8 to 1,536 (small-form factor 2.5 inch) or from 8 to 768 (large-form factor 3.5 inch) HDD or SSD disk drive modules plus 240 Flash cards (Hybrid). All drives in the DS8870 are encryption capable – the customer decides whether to activate encryption. Several different disk types (in mixed configurations) are supported using RAID 5, RAID 6 and RAID 10 techniques (same disk sizes have to be used in a RAID group):

- Flash / Solid State
  - 400 GB Flash FDE capable
  - 200/800/1600 GB SSD FDE capable
- Enterprise
  - 146 GB / 15,000 RPM FDE capable
  - 300 GB / 15,000 RPM FDE capable
  - 600 GB / 10,000 RPM FDE capable
  - 600 GB / 15,000 RPM FDE capable
  - 1.2TB / 10,000 RPM FDE capable
- Nearline
  - 4TB / 7,200 RPM FDE capable

Drive install groups; Flash cards available in groups of 16 or 14, SSD and Enterprise drives available in groups of 16and Nearline drives available in groups of 8.

Functionality

General Functionality

The DS8870 supports variety of advanced functions, including thin provisioning, advanced disaster recovery solutions, business continuity solutions, and advanced copy functionalities. As mentioned above, all disk drives in the DS8870 storage system support the *Full Disk Encryption (FDE)* feature. Because encryption is done by the disk drive, it is transparent to host systems, and can be used in for all platforms – z/OS, AIX, Linux, Windows, etc.Also, since each drive has its own encryption engine there is no performance degradation with added capacity. The *Storage Pool Striping* feature

distributes a volume's or LUN's data across many RAID arrays and across many disk drives, which helps maximize performance without special tuning and greatly reduce hot spots in arrays. The IBM *Easy Tier* feature can automatically move the data and optimize the use of each storage tier, particularly SSD drives. Data areas that are accessed frequently are moved automatically to higher tier disks; for example, to SSDs. Infrequently accessed data areas are moved to lower tiers such as Near-Line SAS drives without manual intervention.

z Systems functionality and the synergy between DS8870 and z Systems

IBM owns the z Systems mainframe architecture, which, from time to time, allows the company to offer exclusive features for at least a limited period. Other vendors purchase the specifications for such exclusive features to remain compatible. IBM provides them with technical specifications[2] of these features, but not the code itself, which the competing companies have to develop..
Typically, the other vendors offer their own version of the features after 12–24 months, but sometimes this takes much longer, and some features are not supported as of today..
The most important z Systems features are:

- **High Performance FICON (zHPF) –** zHPF offers high-performance data transfer, reduces the command overhead and thus better utilizes existing bandwidth. It has four functions, each of which requires cooperation between the z Systems server and the storage system. Selected zHPF functions are: multitrack (allow reading or writing more than two tracks worth of data by a single transport mode operation), extended distance, format writes, QSAM, BSAM, BPAM, and DB2 list prefetch.
- **Parallel Access Volumes (PAV)** and **HyperPAV** allow using multiple devices or aliases to address a single ECKD disk device.
- **I/O Priority Manager s**upports "importance" and "achievement" information provided by z/OS Workload Manager to manage execution priorities. The I/O Priority Manager Quality-of-Service (QoS) function integrates with zWLM and allows automated I/O priority management for mainframe applications. The zWLM monitors every I/O and manages I/O QoS through the Storage System based on zWLM Service Class. zWLM has the ability to provide information to the DS8870 I/O Priority Manager to allocate the required resources based on required goals. Thus, zWLM can now manage host and storage resources end–to-end to optimize the workload execution based on the specified Service Class.
- **Performance – DB2** Specialized cache algorithm can optimize DB2 list prefetch operations by multiple, parallel data fetches. Using simple FICON, the storage system would read a single page per protocol exchange. zHPF reduces the Host-to-Storage System I/O protocol to a single protocol exchange for the full I/O CCW chain. DB2

[2] What is licensed are specifications for interfaces, e.g., I/O commands to invoke a function. That does not include implementation information such as optimizations developed by IBM to support functions or design elements that are included to enable planned future enhancements.

List Prefetch Optimizer with zHPF, enables the storage system to read all 32 DB2 pages in parallel, transferring all 32 pages back to the host in a single exchange. At any point in time, DB2 has two List Prefetch I/Os outstanding; therefore the DS8870 is always reading 64 DB2 pages in parallel.

- **Performance – IMS** provides enhanced performance for IMS write-ahead data set (WADS).
- **Performance – zDAC** supports optimization to improve performance of z/OS Discovery and AutoConfiguration (zDAC).
- **Volume Management** supports dynamic volume expansion for standard (thick) 3390 volumes, Extended Address Volumes (EAV) – supports 3390 volumes up to 1 TB capacity.
- Many other features such as zHPF – format writes, zHPF – DB2 list prefetch cache optimization, Performance – DB2 (cache algorithm, Multiple Readers for z/GM, etc.

DS8000 series Local and Remote Copy Techniques

The DS8000 storage systems support wide range enterprise-level replication techniques for z Systems and other platforms to fulfill any requirements. All these techniques support data consistency. It supports:

Local replications:

- FlashCopy and FlashCopy SE
- Remote Pair FlashCopy (Preserve Mirror)

Remote Mirror and Copy:

- Metro Mirror
- Global Copy
- Global Mirror
- Global Copy Metro/Global Mirror
- z/OS Global Mirror
- z/OS Metro/Global Mirror

Local replications

Standard **FlashCopy** is a Point-in-Time (PiT) technique that uses a normal volume as target volume. This target volume must have equal or larger capacity. The space is fully allocated in the storage subsystem.

FlashCopy Space Efficient (SE) is a so-called “snapshot” technique that creates a virtual volume. At the creation, no space is allocated for this volume. Space is allocated just for updated tracks only when the source or target volumes are written.

Remote Pair FlashCopy a.k.a Preserve Mirror transmits the FlashCopy command to the remote site if the target volume is mirrored with Metro Mirror.

RAS and Disaster Recovery features

Reliability, availability, and serviceability (RAS) played important roles in the design of the IBM System Storage DS8870 structure and components. The DS8870 is based on a redundant cluster of POWER7+ servers. Also all other components are designed to ensure full redundancy. The technology sharing between the DS8000 storage and IBM

Power servers brings manufacturing costs and RAS advantages. The DS8000 series is actually the largest user of IBM POWER processors. By using the same components, IBM leverages economies of scale of larger production volumes. The collected field experience from thousands of servers helps to improve performance and reliability.

Geographically Dispersed Parallel Sysplex (GDPS)

z Systems supports Parallel Sysplex as a local or remote cluster. Up to 32 local or remote mainframes can participate in a single cluster. z Systems Parallel Sysplex also works in conjunction with IBM's disaster recovery software, called Geographically Dispersed Parallel Sysplex (GDPS). GDPS enables automated complete site fail-over with no or minimum loss of data. IBM's GDPS for z Systems is a multi-site application availability solution, with fast recovery time and highly-automated control. It manages application availability in and across sites for both planned maintenance and unplanned situations, such as a site failure or full-blown disaster.

In June 2011 IBM extended the zEnterprise's business resiliency via significant enhancements to GDPS, in particular an active/active configuration (in addition to the active/standby, which is commonly used). The GDPS active/active continuous availability is the next generation of GDPS and a fundamental shift from a failover model to a near-continuous availability model. IBM intends to deliver, over time, additional configurations that comprise GDPS active/active continuous availability – a solution for organizations using two sites separated by unlimited distances, running the same applications, and having the same data with cross-site workload monitoring, data replication, and balancing.

HyperSwap

Another availability function is HyperSwap. This function is probably the most important business-continuity and availability improvement for IBM mainframes and Power servers. While entire site outages are rare, hardware failures are more common.

"Another availability function is HyperSwap. This function is probably the most important business-continuity and availability improvement for IBM mainframes and Power Servers."

With the current integrated and complex application environments – assuming a highly-available, data-sharing Parallel Sysplex environment – storage system becomes a single-point-of-failure for the entire Sysplex. HyperSwap, which is used by multiple GDPS solutions, is controlled by GDPS automation and can eliminate an outage caused by planned maintenance or disk failure by reducing the time needed to switch disks between sites to a matter of seconds and allowing the primary site to use the secondary site's disk storage systems.

Basic HyperSwap between two remote or locally-installed storage systems provides automated fail-over for planned or un-planned outages and can be deployed with z/OS alone, without requiring multi-site GDPS. Like with GDPS, there is no equivalent functionality on any other platform besides z Systems and Power servers.

Case study of robust SAP on zEnterprise infrastructure

The Nationwide Building Society[3] is a large financial institution in the U. K. Its previous infrastructure was based on several platforms (some old and some new), which proved difficult to develop and manage. The Society wanted to modernize that architecture, improve operational efficiency, reduce costs, improve resilience, ensure future scalability, and guarantee customers' (internal and external) service satisfaction. After evaluating several options, the bank decided to select SAP on z Systems with GDPS/MzGM, and the application server on IBM Power servers with AIX. The major factors influencing this decision were the reliability, availability and scalability attributes of these platforms. An example of hardware configuration is shown in Fig. 2.

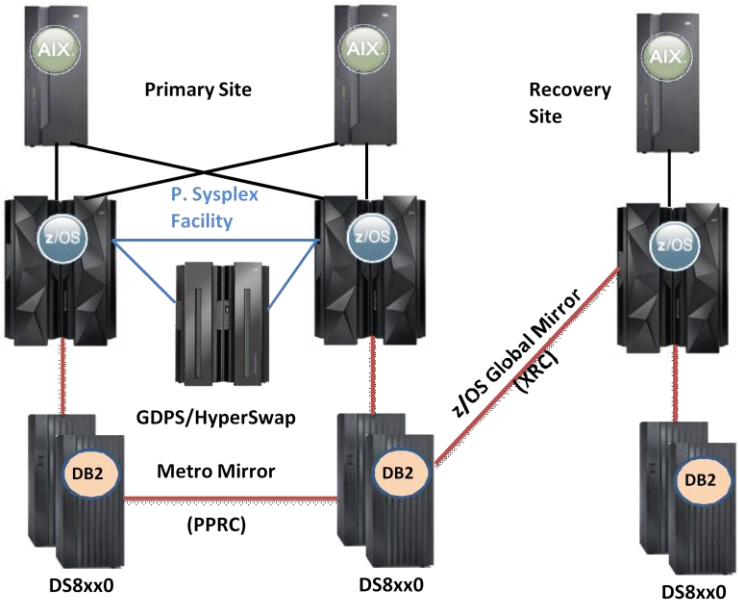

Figure 2: Example of GDPS/MzGM Hardware Infrastructure

Case Study – Advanced z Systems infrastructure for Kenya Power and Lighting Company (KPLC)

KPLC, East Africa's largest power distributor selected z Systems for an automated system that will provide a real time status of all business processes. This will enable Kenya Power to implement its strategic plans to expand its current 2.6 million-client base by another one million clients and to increase its current electricity generation capacity from 2025 megawatts to 5000 plus megawatts by 2015. This new infrastructure will use advanced IBM analytics to enable Kenya Power to study and compare real time and historical data to better monitor business operations and trends, and anticipate future

[3] Nationwide is the world's largest building society as well as the second largest savings provider and a top-three provider of mortgages in the UK. It is also a major provider of current accounts, credit cards, ISAs and personal loans. Nationwide has around 15 million members.

Customers can manage their finances in branch, on the telephone, Internet, and post. The Society has around 15,000 employees. Nationwide's head office is in Swindon with administration centers based in Northampton, Bournemouth and Dunfermline. The Society also has a number of call centers across the U. K.

electrical needs. The real-time analytics will be accessed through the cloud, allowing executives working remotely to use mobile devices to view data pooled from all company resources. The new system will consolidate data from ten key operational sources to provide a single view of enterprise data.
Working with its Business Partner the Symphony Group, IBM developed a complete solution for Kenya Power which includes:

- Hardware - System z BC-12 server, DS8870 storage, PureData for Analytics
- Software: zOS, DB2, Tivoli, Cognos, InfoSphere and SPSS
- IBM services and training

KPLC benefits from the new infrastructure include:

- More efficient resource management
- Improved customer services (call center activity, better identify customer needs and match new services to meet these needs)
- Lower operating costs (service, licenses, environmental) due to consolidation
- State-of-the-art technology hardware and software
- Robust infrastructure with better availability

DS8870 new features and enhancements

The IBM DS8000 pace of development is faster than the competition. In the last 10 years IBM launched seven models of the DS8xxx, more than a double in comparison to the competition. At least twice a year IBM introduces a new model, new features or both. Many of the new features are aimed to create synergy with IBM z Systems mainframes and servers based on POWER. Since the publishing of the original version of this paper in May2013, IBM enhanced the platform four times.

The following sections will describe the major enhancements in particular for z Systems.

DS8870 R7.1 Enhancements (June 2013)

Easy Tier V allows pinning volumes to a given tier and *Heat map* transfer, which enables the transfer of Easy Tier “learning metadata” to remote mirror targets. This unique feature (not available by the competition) saves time to built the “learning” optimization table on the recovery site and thus improved the performance after a planned switch or a disaster.

z/OS Soft Fence will allow a host system to put a volume into a "soft fenced" state and take it out of the “soft fenced” state. When a volume is in the soft fence state, the disk subsystem will prevent all reads and writes to the volume from any host system. It may be useful after a disaster preventing mistakenly corrupting data during the recovery process.

Timestamp synchronization for Sysplex and DS8000 - When a disaster occurs, Active/Active GDPS will need to understand the point of consistency for both software replication and GM in order to reconcile the two copies of the A/A workload data.

A new ICKDSF Verify Offline command that prevents accidental loss of data when initializing a volume with z/OS ICKDSF.

DS8870 R7.2 Enhancements (November 2013)

These were mainly hardware enhancements; the POWER7+ processors, new performance, and nearline HDDs and Microcode improvements enables up to 20% IOPS performance improvement with lower latency.

DS8870 R7.3 Enhancements (May 2014)

This revision improves the DS8870 Flash capabilities, delivering enhanced integration and optimization of flash technology in the DS8870:

Hybrid configuration: Up to 4 High Performance Flash enclosures per base (A) frame ,up to 36TB usable capacity.

All Flash configurations: Up to 8 High Performance Flash enclosures per base (A) frame with up to 73TB usable capacity.

The **High Performance Flash Enclosure ("HPFE")** is the most important part of the announcement. The 1U flash cards are directly plugged into the PCIe fabric of the two DS8870 POWER 7+processor complexes. All the RAID calculation is already hard-coded within the HPFE itself, which houses two RAID controllers and either 16 or 30 encryption capable eMLC flash cards, of 400 GB each and provides up to 12 TB raw / 9.2TB usable RAID5 capacity. The RAID groups are 6D+P with two spares per enclosure. Up to four High Performance Flash enclosures (HPFE) in both the base (A) frame and first expansion frame (B) provide more than 73TB usable capacity[4]. Each enclosure is independently and non-disruptively installed.

The RAID controller is specifically designed to utilize the performance capabilities of flash-based storage. This dense flash enclosure in the DS8870 provides up to four times the IOPS throughput in comparison to current 2.5″ SSDs, with the same amount of capacity.

One HPFE can execute 250,000 IOPS, accelerate database performance by up to factor of 3.2 times and deliver up to 3.4 GB/s data.

These enhancements also included **Easy Tier VI,** which improves the performance exploitations benefits of High Performance Flash.

[4] Increased capacity with support for 4 additional High Performance Flash Enclosures in expansion frame was announced in October 2014.

DS8870 R 7.4 Enhancements (October 2014)

In addition to an improved GUI to simplify storage management, the DS8870 got also new functionality such as:

- Improved disaster recovery protection with **Multi-Target Mirroring**: a single volume is the source for more than one synchronous or asynchronous remote copy simultaneous relationships.
- **Multiple Incremental FlashCopy**: incremental copy of up to 12 volume targets from a single source.
- Increased capacity with new drive options: **1.6TB flash drives 2.5" and 600GB 15K rpm**
- **Easy Tier 7th Generation**: advanced integration with z Systems enables applications to communicate performance requirements for optimal placement of data. Through the API, the application hints will set the intent, and EasyTier will move the data to the correct tier, which allows applications a direct way manage the Easy Tier "temperature" of application datasets. It allows DB2 to proactively instruct Easy Tier of the application's intended use of the data, reducing learning process time.
- **IBM zHyperWrite**: Improved DB2 Log Write Performance with DS8870 Metro Mirror. DB2 performs parallel log writes to primary and secondary volumes, which results of up to a 43% reduction in DB2 log write time and better performance.

The two last enhancements increase further the synergy between the z Systems, DB2 and the DS8870.

DS8870 Compared to Competition

IBM's System Storage DS8000 series is a stable multiplatform high-end storage system able to support current and future user requirements. It uses state-of-the-art technology to achieve top performance levels and supports superior local and remote data mirroring techniques to ensure business continuity. The DS8870 offers unique features to create synergy with z Systems, and IBM continues to introduce such features to increase its competitive advantages.

"The DS8870 offers unique features to create synergy with z Systems, and IBM continues to introduce such features to increase its competitive advantages..

Technology

As opposed other vendors, IBM is a huge technology company developing the required technologies for their products. IBM exploits synergies between the different divisions to produce state-of-the-art products with "tailored-in" technologies specially designed for these products. It takes several years to design a new storage system, which is why only a company with close relations with disk and flash technology can plan for using the latest technologies.

Remote Mirroring Techniques

IBM techniques seem to have a much more intelligent design in comparison to other vendors. Some major advantages:

IBM's Global Mirror function does not collect the modifications in cache as other vendors do, but uses a "track bit map" technique which consumes almost no cache capacity for temporary storing "delta sets". Global Mirror almost continuously transmits groups of consistent data to the secondary site as soon as possible after it is written in the primary. The actual "lag time" depends on the amount of data modified and the link bandwidth available. In practice, assuming that sufficient network bandwidth is available, the average lag time is on the order of five seconds.

Effective usage of the communications link bandwidth has significant impact on performance as well as the overall cost of a remote copy implementation. Global Mirror is designed to transfer data to the secondary as quickly as possible, which means immediately after receiving it if the link bandwidth is available.

Market

Only four vendors compete in the IBM mainframe storage market; EMC, HDS, HP (OEMing Hitachi as XP P950) and IBM. However currently only IBM offers state-of-the-art storage systems and tight system-z support. Since October 2007, IBM has accelerated its development rate, offering enhancements at the fastest pace in the industry. Examples of major enhancements include IBM's introduction of RAID-6 in August 2008, High-performance FICON for z Systems in October 2008, full-disk encryption and a solid-state drive (SSD) option in February 2009, and thin provisioning July '2009, etc. IBM sees the DS8870 as a strategic product and continues to announce enhancements at least twice a year.

In April 2010 IBM announced and delivered the IBM System Storage Easy Tier, which automates data placement within the DS8000 series system for performance and cost optimization. This includes the ability of the system to automatically and non-disruptively relocate data (at the extent level) across drive tiers (or within a tier, to less-loaded ranks), and the ability to manually relocate full volumes. This was the first sub-LUN automated data placement which was so required for optimal SSD and high-performance HDD utilization.

According to IDC figures the DS8870 is the market leader in storage for z Systems. In the 2Q14 EMEA Quarterly Disk Storage Systems Tracker from International Data Corporation, IDC wrote "*However, growth remains mixed across vendors and countries, with IBM's strong performance (mainly due to a revival in its DS8000 series) accounting for most of the growth in the subregion*."

Summary and Conclusions

IBM's System Storage DS8000 series is a stable multiplatform high-end storage system able to support current and future user requirements. It uses state-of-the-art technology to achieve top performance levels and supports superior local and remote data mirroring techniques to ensure business continuity. The DS8870 offers unique features to create synergy with z Systems and Systems i,p. The DS8870 has all the required characteristics to be positioned on the top of enterprise storage systems for all platforms and in particular for z Systems.

"IBM's System Storage DS8870 series is a stable multiplatform high-end storage system able to support current and future user requirements.

IBM, which 16 years ago lost some of its high-end disk enterprise storage market share, has managed an impressive come-back, delivering technology that should be put on a short list of any high-end storage procurement.
In addition to storage products, IBM delivers four server architectures, and thus the benefits of exploiting of the synergy between IBM servers and the DS8870 storage systems should be evaluated as well. z Systems mainframes and DS8870 systems have many synergies, in particular in Disaster Recovery and Business Continuity but also with DB2, etc.

The DS8870 pace of enhancements for hardware and features is industry leading. In addition to deploying the latest HDDs and SSDs technologies, the users of DS8870 may benefit from unique features such as *Heat Map Transfer* or the *High Performance Flash Enclosure.*

The operation GUI, ported from XIV, delivers the most user-friendly functionality in the industry and is also shared by the complete portfolio of IBM storage solutions.

Another IBM advantage is IBM Global Technology Services (GTS). GTS is the world-wide largest IT services organization with more than 50 years of experience in almost every vertical industry, with the ability to design anything from a basic data center infrastructure to the most complicated disaster recovery deployments.

Full redundancy, non-disruptive upgrades and maintenance, hot-swappable components, pre-emptive soft error detection and online microcode changes ensure high availability and data integrity. The advanced remote data replication techniques enable any scheme of disaster recovery deployments.

In summary, the DS8870 deserves to be put on a short list when considering a robust, high-performance, advanced storage system for z Systems or other platforms for the most demanding storage infrastructures.