



VNX: STORAGE TECHNOLOGY HIGH BANDWIDTH APPLICATIONS

Big data solutions for the mid-enterprise

EXECUTIVE SUMMARY

Businesses are constantly looking for new and innovative ways to succeed; to open up new business opportunities, or to perform their core competencies with increasing efficiency. One of the tools available to them to help with these strategies is the process of business intelligence (BI), where analytics are applied to disparate but related business information to uncover new business opportunities.

Here are some examples of how business intelligence is applied for competitive advantage:

- In financial forecasting, the ability to see profit and profit margin per customer per transaction for the top 10 percent of our customers allows us to identify the characteristics of the most profitable transactions.
- By monitoring the minute by minute power usage in a household, utility companies gain a far more comprehensive understanding of power usage patterns in a neighborhood, allowing them to optimize the power generation and supply process to ensure improved service levels with reduced costs.

Consider Facebook. The company is estimated to be worth around \$100 billion (May 2011) and much of this value is based upon the value it can potentially glean about its 600 million members. Information effectively turned into business opportunity can be worth BILLIONS.

Business Intelligence employs a number of data processing tools to a data set to turn data into decision criteria. Two of the key technologies fundamental in BI are decision support systems (DSS) and data warehousing (DW). Data marts may also be used, which are a subset of a data warehouse, typically oriented to a specific line of business.

As users recognize the differentiated value of data warehousing, the following characteristics need to be designed into the solution:

- **Currency:** the more up to date the data, the more relevant the decision making
- **Availability:** when critical business decisions and company revenue depend on the data warehouse, it must be always available
- **Scale:** the value of a data warehouse can be expanded by adding more data, allowing richer queries, so maintaining a highly performing solution as the implementation grows is important
- **Cost effectiveness:** flat IT budgets dictate the solution be as cost effective as possible to purchase and to operate, metrics such as \$ per MB/s and MB/s per rack U become highly relevant

It is critical that the infrastructure supporting the solution: server, network, storage, and application provide a robust, powerful, and flexible solution. The storage component is the data engine that stores and feeds the data warehouse while ensuring the integrity of the data feed, and is critical to maintain a valuable solution. EMC is the market leader in most major storage markets and offers the ideal storage solution for the mid-market, cost sensitive data warehouse customer.

The mid-range data warehouse market is ideally served by the EMC® VNX® series of storage platforms which provides:

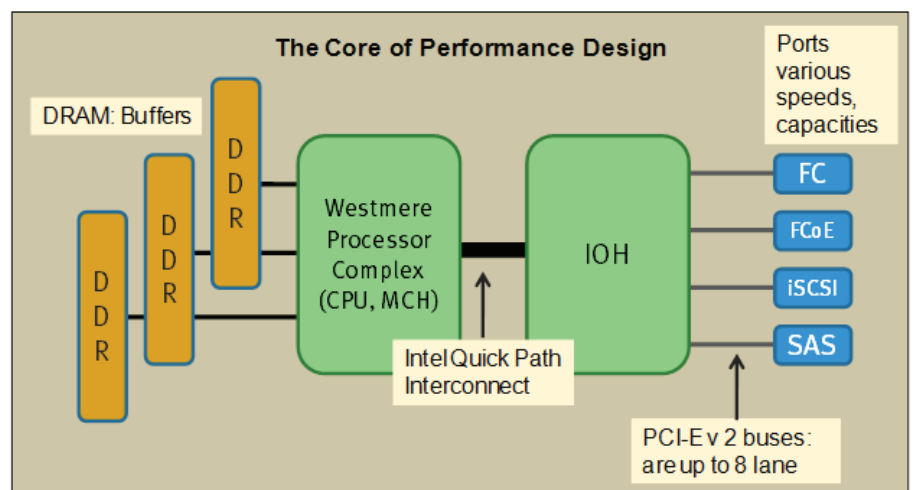
- **Capacity scale** from as little as a few TBs to multiple PB and 1000 drives
- Performance, using a high bandwidth storage controller design that delivers up to 10 GB per second for BI type workloads
- **Availability** with “designed for 5x9s” architecture (5x9s equals roughly five minutes of unplanned downtime per year)
- **Optimized TCO** via a simple management paradigm and extensive feature automation

SOLUTION OFFERING

The EMC VNX family delivers industry-leading innovation and enterprise capabilities for file, block, and object storage in a scalable, easy-to-use solution.



The VNX series storage platform is designed from the ground up to deliver unparalleled end to end performance to support the world's most demanding applications. There are at least four hardware components that factor into the bandwidth equation:



- The front-side connection ports (and protocol) used to connect servers to the storage system—low latency Fibre Channel or Fibre Channel over Ethernet is recommended for high bandwidth applications.
- The number of physical drives holding the needed data that participate in delivering data from the drives—the speed and type of drives affect what can be practically delivered by the design of the drives. VNX supports Flash (high-performance), SAS, and Near-Line SAS drives (high capacity).
- The storage system's CPU and memory capability for accepting incoming requests from the host, mapping and dispatching the I/O to the drives in the back end, and returning the data retrieved from the drives back to the servers through the front-side connections—VNX is powered by the latest Intel multi-core Xeon 5600 Series processors and DDR3 1333 MHz DRAM memory.
- The number of buses/channels over which the storage system processors can dispatch I/O to the back-end drives, and pull the data bits up from the drives—VNX implements a 4-lane 6 Gb/s SAS back-end and PCIe Gen 2 controller buses.

The table below outlines the configurations and performance capabilities of the VNX series platforms for data warehouse and similar workload types. Included in these configurations is a special VNX5500 offering specifically designed to provide a high-performance yet cost-effective solution for the mid tier (Note: the high-performance 5500 configuration requires VNX Block OE R31.5 code) the workloads in performance VNX5500 high bandwidth option is an ideal storage solution for big data analytics processing.

	VNX5100	VNX5300	VNX5500	VNX5500*	VNX5700	VNX7500
CPU	Intel Xeon 5600	Intel Xeon 5600	Intel Xeon 5600		Intel Xeon 5600	Intel Xeon 5600
Memory	8 GB	16 GB	24 GB	24 GB	36 GB	48 GB
Back-end SAS Buses	2	2	2	6*	4	4 or 8
Max Front-end FC Connections	8	16	16		24	32
DSS Bandwidth (MB/s)	2300	3600	4200	4200**	6400	7300
Data Warehouse Bandwidth (MB/s)	2000	3200	4200	6400**	6400	10000
Backup Bandwidth – Cache Bypass Mode (MB/s)	700	900	1700	1900**	3300	7500
Rich Media Bandwidth (MB/s)	3000	4100	5700	5700**	6200	9400

* The VNX 5500 Hi-Bw option consumes all the flex I/O modules in the system, and the bandwidth figures here are based on FC connectivity.

** To achieve the full data warehouse bandwidth, at least 130x15K SAS drives (or equivalent) are required.

COMPARED TO THE VNX5500 THE SPECIAL VNX5500 BANDWIDTH CONFIGURATION DELIVERS:

50 percent > bandwidth	25 percent better price/performance	43 percent more performance density
------------------------	-------------------------------------	-------------------------------------

SOLUTION APPLICATION DETAILS

Application	Outline	I/O Characteristics
Data warehouse—e.g., Oracle, MySQL, and Sybase	DW applications generate a very specific workload type that requires special consideration as to the storage configuration needed to most efficiently meet the requirements of the application. In most cases, data warehouses tend to emphasize the importance of sustained read bandwidth from the supporting infrastructure. While this is the basis for the solution discussed here, a flexible approach is required as the actual need is still often dependent on how the data model and the user applications are designed and implemented, and in many cases the actual I/O patterns may vary. The typical data warehouse workload accesses large chunks of disparate data that is then merged and aggregated to provide meaningful patterns.	512 KB Block Random Read
Decision Support Systems	DSS applications tend to generate slightly different workload types from data warehouse, in that they tend to involve highly sequential read workloads indicative of the need to read very large amounts of data requiring aggregation ahead of a complex data management query.	64 KB Block Sequential Read

Backup Processing— e.g., Symantec NetBackup, EMC NetWorker®, etc.	Backup has become increasingly prevalent in the disk storage space. As users recognize the limitations and reliability exposures of a tape based backup solution, they look to disk storage as the backup medium of choice. This is due to its increased (direct access) performance for both backup and restore, it's improved reliability and resiliency, as well as built in features such as replication, spin down, etcetera that ensure improved backup and restore service levels at a reasonable cost point. Consider that this also comes with the same simple management paradigm as the primary storage model, and it is obvious why storage subsystems such as VNX are used extensively in this space.	512 KB Block Sequential Write
Rich Media Applications— e.g., Adobe Flash, video editing and streaming, CCTV, etc.	Rich Media also requires that the underlying storage system is able to deliver many concurrent streams of (often) high definition video content to cable or web based users of rich video content.	256 KB Block Sequential Read

All these workloads have one thing in common, the need to process large volumes of data either into or out of the storage system, typically in a sequential manner, and hence bandwidth is the foremost performance vector to support.

SOLUTIONS BEST PRACTICES

Based on rigorous testing within the EMC solutions organization, a number of configuration recommendations are applicable to these high bandwidth applications. We will focus on the data warehouse use case although in many cases the recommendations are consistent across the other applications. The ultimate goal is to understand the requirements of the DW/DSS solution (capacity, performance, and resiliency) and design a solution that meets the end-to-end needs of the application. The solution must ensure that the drives deliver sufficient capacity and performance, and that the storage controllers and front end connectivity can handle the load. In addition the configuration should be optimized for cost and ease-of-use.

Drive types—In general, high bandwidth, primarily sequentially accessed applications, can very effectively leverage standard high-performance spinning disk media such as 15K rpm SAS drives. In many data warehousing and decision support system implementations, however, the workload is shared with other applications and certainly the system is going to be running a mix of simple and complex queries that may introduce different workload characteristics for the storage. The drive options and their performance capabilities (measured as large block bandwidth or small block 8 KB I/Os per second) within the VNX5500 are outlined below:

- Flash drives (Solid State Device): capable of around 250 MB/s and 3000 I/Os per second
- 15K SAS drives: capable of around 50 MB/s and 180 I/Os per second
- 10K SAS drives: capable of around 35 MB/s and 140 I/Os per second
- 7200 rpm Near-Line SAS drives: capable of around 25 MB/s and 90 I/Os per second

The general recommendation for data warehouse workloads is to primarily leverage 15K SAS technology, although if data is stored for long periods and fewer queries are run on older data then NL-SAS can be leveraged as the data ages. In addition, if the warehouse implementation makes extensive use of components such as data cubes, summary tables, indices, etcetera then using flash drives as part of a “pool based solution” provides significant value. Storage pools will be discussed later in this paper.

One further consideration would be physical drive size. With the availability in the VNX of 2.5” 10K rpm SAS drives, there is now a dense 10K spinning drive option that performs like a 15K 3.5” drive but supports 12.5 drives per cabinet U compared to five drives per U for 3.5” technology.

Connectivity—Historically, data warehouses have leveraged high bandwidth Fibre Channel connectivity, and while **state of the art 8 Gb/s FC remains the deployment option of choice**, EMC's use of plug-and-play flexible I/O Modules provides concurrent support of 10 Gb/s iSCSI and FCoE to ensure the solution aligns with the customer's infrastructure of choice.

FCoE usage is growing faster than any other protocol (>100 percent over the next five years according to IDC), and as a means of consolidating storage network and client network traffic over a single network infrastructure, provides compelling benefits to customers. FCoE is a very viable option for DW implementations, and even if a customer is migrating to FCoE and has a combination of FC and FCoE, the VNX series is able to support both connectivity models as needed.

iSCSI over 10 GbE can be considered, although the overhead of the iSCSI driver stack is a little more memory intensive than the FC or FCoE stacks, so where absolute performance is critical, choose one of the FC protocols.

If iSCSI is required due to the limitations of the server connectivity options, then where possible, use jumbo frames to improve IP network throughput. In addition, iSCSI can be direct connected to the VNX for a simpler configuration experience.

Note: The testing outlined in this document was based on FC connectivity and iSCSI. FCoE connectivity may not achieve the same bandwidth numbers.

Storage pools and Auto Tiering—EMC introduced the concept of storage pools in the EMC CLARiiON® CX4 product line, and supports the concept with the VNX. A key advantage with LUNs created out of storage pools is that the physical data layout of the LUN content is now under automatic management of the VNX storage system. Pools can contain a mixture of Flash, SAS, and NL-SAS drives, and introduce the concept of a virtualized storage environment. VNX does still support the traditional FLARE mechanism of provisioning RAID group-based LUNs.

When the workload is primarily sequential in nature, the deployment mechanism of choice is: **Traditional RAID Group LUNs**, which are preferred due to the ability to achieve the highest performance levels where the workload can be most cost effectively met by a single media type: 15K SAS drives. Generally, expect a slightly improved MB/s estimate per drive when traditional RAID Group LUNs are compared to Pool-based LUNs.

Considerations for using Pooled-based LUNs:

- Pools can be used where ease-of provisioning is important. The implementation should use a pool, and deploy fully allocated, or “thick” LUNs. “Thin” LUNs can be used in pools to reduce the cost of deployment, although performance considerations tend to limit their applicability to performance insensitive implementations.
- If the stored data clearly includes the “redundant data” created for the purpose of speeding up high impact queries, consider boosting performance with Flash drives with enough capacity to effectively house the expected “extra data.” See the FAST Cache section below.

Memory and FAST Cache—Typical DSS/BI implementations process very large volumes of data, and DRAM caches have little effect for anything other than write processing. **EMC's recommendation is to limit read cache to no more than 1 GB per system** to act as a read ahead buffer, although in many cases it may be advisable to not configure read cache at all. Write cache is recommended to be turned on to help with data ingest as well as for Data Sorting processes.

DRAM Cache can be expanded with the use of FAST Cache (a Flash based extension of classic cache) which may be more relevant than pure DRAM for DW implementations. Unlike the 24 GB of DRAM cache in the VNX5500, FAST cache can scale up to 1 TB, providing a major potential benefit to the data warehouse implementation, particularly for data cubes, summary tables, indices, and more.

SUMMARY

The VNX series offers unparalleled cost performance for many of the most critical high bandwidth data center application requirements. The new VNX5500 High-Bandwidth Option provides the availability, flexibility, and feature rich management options that ensure Business Intelligence projects are rolled out consistently and on budget—while also allowing the project to grow and deliver additional value as new needs are added.

GLOSSARY

Decision Support Systems—DSS include knowledge-based systems. A properly designed DSS is an interactive software-based system intended to help decision makers compile useful information from a combination of raw data, documents, personal knowledge, or business models to identify and solve problems and make decisions. In many cases DSS leverages a data warehouse engine to provide the specific data queries.

Data warehouse—A data warehouse is a database where data from production databases are copied to the data warehouse so that queries can be performed without disturbing the performance or the stability of the production systems. The warehouse contents are typically historical and static and may also contain numerous summaries to speed data access. The data warehouse is structured to support a variety of analyses, including elaborate queries on large amounts of data that can require extensive searching.

Data Mart—A data mart is a subset of a larger data warehouse which is often created for just one department or product line.

REFERENCES

VNX Family Data Sheet:

<http://www.emc.com/collateral/hardware/data-sheets/h8520-vnx-family-ds.pdf>

VNX Software Suites Data Sheet:

<http://www.emc.com/collateral/software/data-sheet/h8509-vnx-software-suites-ds.pdf>

Deploying VNX Unified Storage Systems for Data warehouse Applications:

<http://www.emc.com/collateral/hardware/white-papers/h8177-deploying-vnx-data-warehouse-wp.pdf>

VNX series Specification Sheet:

<http://www.emc.com/collateral/software/specification-sheet/h8514-vnx-series-ss.pdf>

Introduction to the VNX series White Paper:

<http://www.emc.com/collateral/hardware/white-papers/h8217-introduction-vnx-wp.pdf>

EMC FAST VP for Unified Storage Systems:

<http://www.emc.com/collateral/software/white-papers/h8058-fast-vp-unified-storage-wp.pdf>

EMC VNX Fast Cache White Paper:

<http://www.emc.com/collateral/software/white-papers/h8046-clariion-celerra-unified-fast-cache-wp.pdf>

EMC Unisphere—Unified Storage Management Solution White Paper:

<http://www.emc.com/collateral/software/white-papers/h8017-unisphere-element-manager.pdf>

CONTACT US

To learn more about how EMC products, services, and solutions can help solve your business and IT challenges, contact your local representative or authorized reseller—or visit us at www.EMC.com.

EMC², EMC, CLARiiON, CX, Unisphere, VNX, VNX5500, and the EMC logo are registered trademarks or trademarks of EMC Corporation in the United States and other countries. All other trademarks used herein are the property of their respective owners. © Copyright 2011 EMC Corporation. All rights reserved. Published in the USA. 10/11 EMC Perspective H8929