October 2011

MANAGEMENT BRIEF

Cost/Benefit Case for IBM XIV Storage

Comparing Cost Structures for IBM XIV and EMC VMAX/EMC VMAXe Systems

International Technology Group

609 Pacific Avenue, Suite 102
Santa Cruz, California 95060-4406
Telephone: + 831-427-9260
Email: Contact@ITGforInfo.com
Website: ITGforInfo.com

TABLE OF CONTENTS

List of Figures

EXECUTIVE SUMMARY

IBM XIV and EMC Symmetrix VMAX

IBM XIV Storage System and EMC Symmetrix VMAX are to a large extent contemporaries. XIV was launched in 2008. VMAX was introduced in April 2009 as the flagship EMC high-end disk system. It plays a central role in the company's Fully Automated Storage Tiering (FAST) strategy.

There has been significant demand for both systems. EMC reported that its high-end storage business – including the company's older DMX systems – grew by close to 25 percent during 2010, and by 15 percent and 25 percent respectively during the first and second quarters of 2011.

XIV growth, however, has clearly been higher. As of the end of 2009, approximately 1,000 systems had been installed. The current total appears to be close to 5,000. According to IBM, revenues have approximately doubled on a year-on-year basis.

One of the major drivers of XIV demand has been that overall costs are significantly lower than for VMAX systems and other high-end competitors. The XIV platform is built around comparatively inexpensive "commodity" components, software is offered in a single integrated package, and IBM pricing has been highly competitive. Storage administration and energy costs have proved to be less.

During 2011, both platforms have undergone major enhancements. EMC has delivered the second phase of FAST architecture, which automates allocation of data across flash memory (solid state), and Fibre Channel (FC) and SATA drives to improve overall performance. The company also introduced a new entry-level VMAXe.

The IBM XIV Storage System Gen3, introduced in July 2011, boosts capacity and performance. Systems employ latest-generation Intel Westmere processors, high-capacity serial attached SCSI (SAS) drives, increased cache, I/O upgrades and InfiniBand internal connectivity. The company continues to offer earlier models for environments with less exacting performance requirements.

XIV systems retain their cost edge. In six representative installations in large and midsize organizations, three-year costs for use of XIV systems average 69 percent less than for VMAX equivalents equipped with FC drives. Figure 1 illustrates this picture.

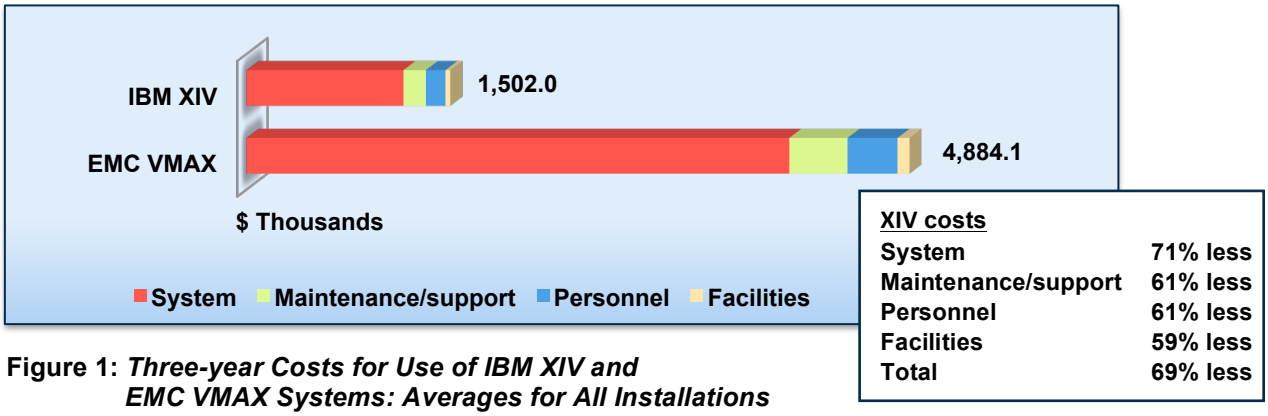

Figure 1: ***Three-year Costs for Use of IBM XIV and EMC VMAX Systems: Averages for All Installations***

In individual installations, three-year costs for use of XIV systems range from 65 to 73 percent less.

These comparisons are for financial services, government, health care, life sciences, manufacturing and telecommunications organizations with initial installed capacities of between 44 and 539 terabytes (TB). XIV configurations include Gen3 as well as earlier models.

In the same set of installations, three-year costs for use of current XIV systems average 67 percent less than for VMAX and VMAXe equivalents equipped with FC drives. Figure 2 illustrates this picture.

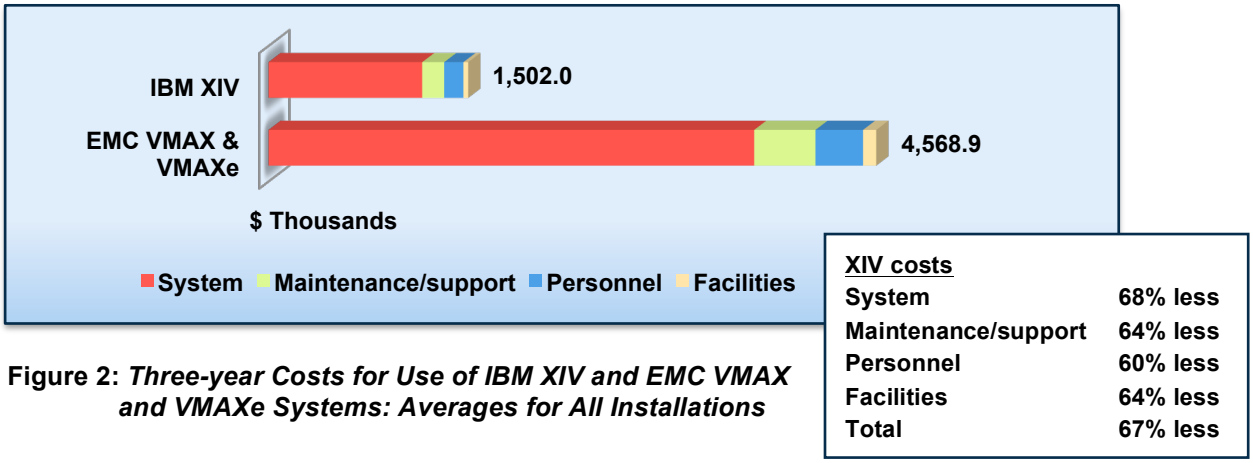

Figure 2: ***Three-year Costs for Use of IBM XIV and EMC VMAX and VMAXe Systems: Averages for All Installations***

Technology limitations restrict VMAXe use to smaller installations and less critical applications. Marginally lower VMAXe hardware and software costs are offset by maintenance outlays (VMAXe warranties are for next business day onsite coverage only), higher energy consumption (VMAXe systems support only single-phase power) and other factors.

IBM XIV and EMC FAST/FAST VP

According to EMC, during second quarter 2011, more than 90 percent of VMAX systems were shipped with FAST software, and flash and SATA drives. It can be expected that FAST-enabled configurations will become the norm for new VMAX deployments.

Use of FAST and FAST VP (Fully Automated Storage Tiering for Virtual Pools), according to EMC, improves VMAX performance and may reduce costs. Industry experience with storage tiering has, however, shown that such effects vary widely depending on application and workload characteristics. At the organizational level, the impact on cost structures appears to be more incremental than radical.

If FAST-enabled VMAX systems are employed in the same six installations as for other comparisons, three-year costs for XIV systems still average 61 percent less. Figure 3 illustrates this picture.

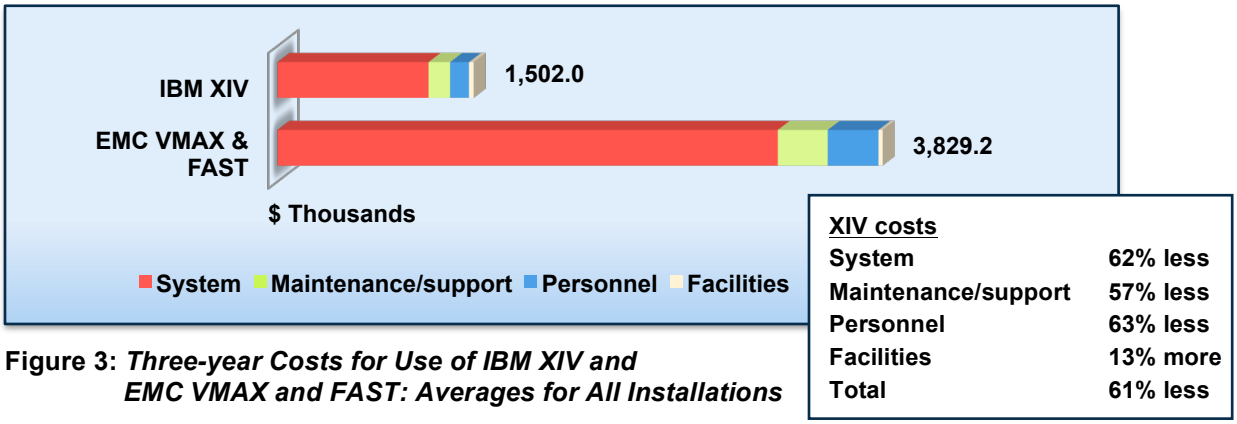

Figure 3: ***Three-year Costs for Use of IBM XIV and EMC VMAX and FAST: Averages for All Installations***

In this comparison, VMAX costs are for FAST-enabled systems providing the same usable capacity configured with solid state (SSD), FC and/or SATA drives. Configurations are based on performance requirements for applications supported by individual systems.

Lower VMAX hardware and energy costs are offset by costs for SSD drives, FAST software license and support fees, and higher personnel costs for performance tuning.

For all three sets of comparisons, costs include hardware acquisition and maintenance; software acquisition, licenses and support; storage administration personnel costs; and facilities, including energy and data center occupancy. Calculations allow for three-year capacity growth.

The "Systems" category in figures 1, 2 and 3 includes hardware and software licenses. Hardware and software costs for both platforms are calculated based on "street" (i.e., discounted) prices.

Further details of methodology, along with installations, configurations, staffing levels and cost structures may be found in the Basis of Calculations section of this report.

Performance and Availability

Critics of the XIV have often argued that performance and availability weaknesses meant that this platform was not viable as a "Tier 1" system. User experiences contradict such claims.

In a separate survey of 63 XIV users, more than 60 percent reported that XIV systems supported one or more Tier 1 applications. Organizations ranged from midsize businesses to Financial Times Global 50 companies. XIV systems also supported Tier 2 and Tier 3 applications, and 25 percent reported leveraging XIV single-tier architecture to consolidate storage platforms.

XIV systems were said to meet performance requirements for a wide range of applications. These included some of the largest SAP installations in the United States and Europe, along with other large-scale ERP systems, and business-critical systems in banking, health care, insurance, telecommunications, IT services and other industries.

Virtualized server farms using VMware; other x86 and UNIX hypervisors (more than 90 percent of users); and Microsoft Exchange and other e-mail systems (63 percent) were also commonly supported.

Users reported production XIV systems in the 30,000 to 50,000 I/Os per second (30-50K IOPS) range, and internal tests showing 50K to 100K IOPS. These results, it should be noted, were for earlier XIV models. It can be expected that the latest systems will at least double XIV performance, and that SSD caching scheduled for first half 2012 availability will boost it further.

XIV systems were said to be more effective than conventional high-end disk systems in handling mixed, fluctuating workloads. XIV parallel architecture enabled all system resources – including processors, disks and I/O – to be applied to all workloads. Embedded load balancing further improved mixed workload effectiveness.

This approach is well suited to heterogeneous environments such as those generated by storage consolidation, support for virtualized server farms (hundreds to thousands of diverse workloads may be running in individual partitions), distributed server clusters and cloud computing.

For all types of application, there was general agreement that XIV systems enabled extremely high levels of availability (terms such as "24/7" were commonly employed) and that recovery processes were both rapid and effective. Organizations that had replaced high-end disk systems reported that XIV systems delivered equivalent or superior capabilities to these.

Conclusions: Differences

The differences between VMAX and XIV systems are fundamental. They implement significantly different architectures. The VMAX platform is a conventional high-end disk system equipped with automated tiering capabilities that may improve performance and reduce costs for certain applications.

XIV systems represent the “Swiss Army Knife” of the storage world. They are designed to provide balanced performance and top-of-the-line quality of service for applications, and combinations of applications with widely varying characteristics, at a fraction of the cost of the VMAX platform.

There are also obvious differences between EMC and IBM product strategies. EMC strategy appears to revolve around maintaining the core VMAX platform “as is” while improving performance through tiering, and promoting various schemes for integration of servers, networks and storage with EMC products. IBM strategy focuses more on increasing the competitiveness of the core platform.

User experiences highlight a broader difference. XIV systems significantly reduced the complexities with which organizations had to deal in planning, provisioning and managing storage infrastructures. The XIV platform requires no tuning, and is easy to install, configure, expand and administer. The VMAX system is a great deal more complex, and FAST will make it even more so.

Complexity reduction was particularly valued by organizations facing rapid storage growth, increasingly heterogeneous workloads or both. Even in relatively small installations, however, XIV systems were seen as simplifying operations and minimizing overhead.

When the XIV system was first introduced, its combination of low costs, advanced functionality and Tier 1-class performance was widely seen as a “game changer” in disk systems. It may be, however, that its ability to minimize complexity in highly diverse, fast-changing conventional storage installations, as well as in emerging storage clouds, will represent an even more significant industry paradigm shift.

USER VIEW

Demographics

User Base

XIV clearly enjoys a strong appeal. As of this writing, close to 5,000 XIV systems have been installed worldwide. According to IBM, more than 1,100 of these were new storage customers for the company. Installations cover a wide range of organization sizes and industries. The smallest employ a single system, and the largest have more than two petabytes (PB) of installed XIV capacity.

In preparing this report, a survey of 63 XIV users worldwide was conducted. Demographics appear to be representative of the overall XIV base. The results of the survey, which covered such topics as numbers and capacity of systems, applications and user experiences, are summarized in this section. Technical issues are discussed in more detail in the following section.

The diversity of XIV deployments is striking. Organizations ranged from Global 50 corporations to businesses with fewer than 30 employees, while installations ranged from a single system with 27TB of usable capacity to more than 30 systems with a combined capacity of over 1PB.

Users included financial services (10), health care (9), manufacturing (9), IT services (6), insurance (5), media and telecommunications (3 each), distribution, retail, education and utilities (2 each) and construction and energy (1 each) companies, along with 7 government and 1 nonprofit organizations.

Applications

In all cases, organizations reported that XIV systems supported multiple applications. Most reported from 5 to 100 applications, although some reported 200 to more than 500. In many cases, overall totals were unclear. This was particularly the case where systems supported populations of virtualized servers.

Among 52 organizations that reported applications in detail, the principal types deployed were as shown in figure 4. Multiple applications in each category were often deployed.

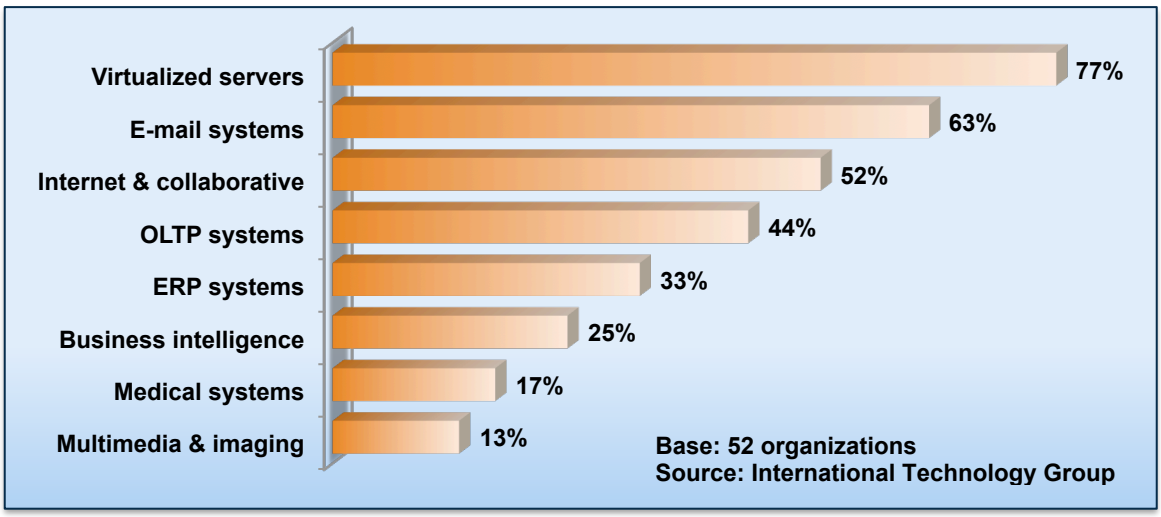

Figure 4: ***Principal Types of Application Deployed on IBM XIV Systems***

E-mail systems were predominantly based on Microsoft Exchange (27 cases) but included Lotus Notes and Domino, and open source mail applications. Internal mail systems ranged from fewer than 100 to 80,000 seats, although hosting companies reported even larger numbers – in one case, over 250,000.

ERP systems included SAP (nine cases), along with Oracle E-Business Suite, JD Edwards, industry-specific systems and others. Systems supported from 60 to over 5,000 users.

Other online transaction processing (OLTP) systems included core banking and insurance systems; billing and customer care systems in telecommunications, IT services and utilities; customer relationship management (CRM), finance and human resources systems; and a variety of industry-specific solutions. Medical systems included hospital and laboratory information systems, electronic medical records (EMR), picture archiving and communications systems (PACS) and others.

Business intelligence (BI) applications ranged from departmental data marts to multiple-terabyte Cognos, Oracle and SAS data warehouses.

Virtualization

Among the same 52 organizations, 40 (77 percent) reported that XIV systems supported virtualized server environments.

Of these, 33 (63 percent) reported use of VMware. In six organizations, XIV systems supported over 1,000 virtual machines (VMs), and one reported over 10,000. Hosting services companies employing XIV systems also reported heavy use of VMware for customer applications.

Other enablers were said to include IBM PowerVM and other UNIX-based hypervisors (8 cases), Microsoft Hyper-V and/or Virtual Server (7), Citrix XenServer (4) and IBM z/VM (1). A number of organizations also employed open source tools.

Benefits

Overview

Among the 63 XIV users surveyed, low overall costs were – by a wide margin – the most commonly cited benefit of employing XIV systems. Respondents cited lower hardware, software license and support, and facilities and administration costs compared to competitive systems.

The appeal of XIV systems was, however, not simply that they were inexpensive. Organizations typically conducted proof of concept (POC) or equivalent tests to compare performance and scalability of different systems, and evaluated comparative availability, recovery and other capabilities.

XIV systems were preferred because they offered equivalent or superior functionality to conventional disk systems – including high-end Tier 1 platforms – at significantly lower costs. Organizations reported a variety of other benefits of employing XIV systems, including those shown in figure 5.

A common experience among XIV users was that the system's key technology characteristics were mutually reinforcing. Grid architecture, for example, contributed to high levels of performance – particularly for mixed workloads – as well as availability.

The strengths of XIV snapshot technology (enabling extremely rapid creation of point-in-time copies of data) and thin provisioning (i.e., the ability to allocate storage capacity in real time in response to application requirements) also delivered multiple benefits.

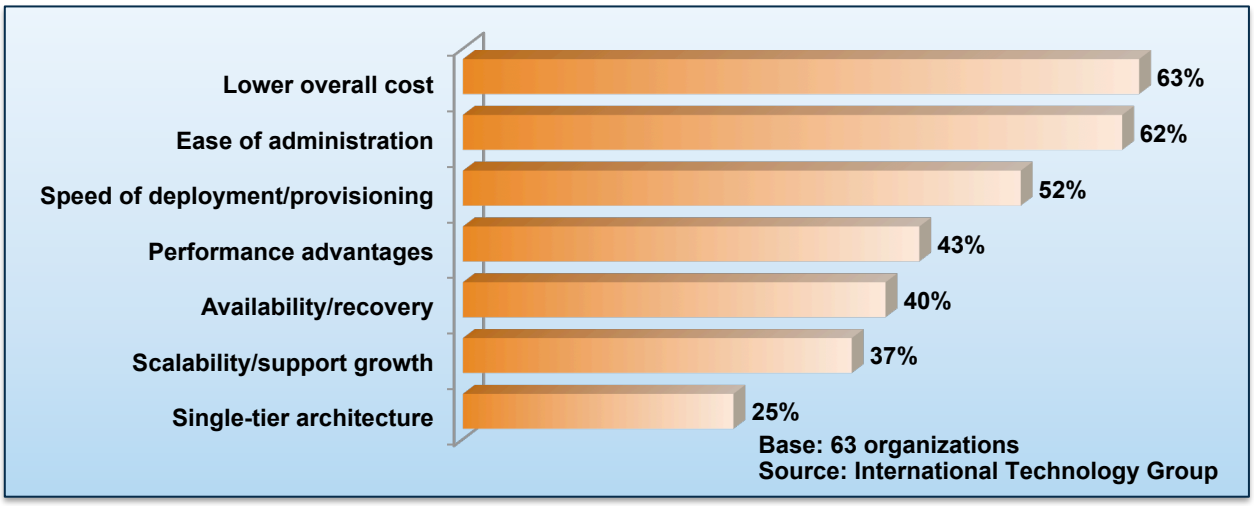

Figure 5: ***Principal Benefits of Employing IBM XIV Systems – User View***

Differential snapshot capabilities improved the speed and effectiveness of backup and replication processes. Both capabilities were said to materially improve capacity utilization and reduce storage administration overhead. Other examples were also cited.

Ease of Administration

Users repeatedly cited the comparative ease with which XIV systems could be installed, configured and managed. There was general agreement that administration tasks could be performed more rapidly, with less time and effort, than was the case for conventional disk systems with which they were familiar.

These reflect multiple XIV features, including the system's distinctive, high-productivity management graphical user interface (GUI); high levels of automation (including automated data placement, load-balancing and performance optimization functions); and the simplicity of the overall system architecture.

One key benefit was said to be that ***system setup and provisioning*** could be performed more easily and rapidly. XIV systems were in some cases brought into operation in "an hour…less than an hour…less than two hours…an afternoon…less than a day…four hours, from the loading bay to fully operational."

Provisioning was typically performed in a matter of minutes, or less than a minute, involving "simple point-and click…a few mouse clicks."

Other tasks said to benefit from distinctive XIV capabilities included hardware add-ons, creation and deletion of partitions, definition of volumes and snapshots, expanding logical units numbers (LUNs), management of cloning, backup and replications (XIV snapshot strengths simplified and accelerated these processes), along with data migration, system monitoring, troubleshooting and others. Performance monitoring and tuning tasks were largely eliminated.

It was noted that XIV manageability advantages were particularly valuable when new applications were deployed, or when workloads changed for other reasons. Gains in administrator productivity thus tended to be highest in organizations that experienced frequent changes in their storage environments.

Five companies providing managed hosting and/or ISP services reported that XIV manageability strengths translated into direct bottom-line benefits. Deployment times for new customer applications, as well as ongoing storage administration overhead were significantly reduced.

The XIV GUI was variously described as "very simple …very easy to use…outstanding…brilliant" and as requiring "very little training…virtually no training." Storage administrators could become proficient in its use "in an afternoon…in less than day…within a day…within days…within a few days…within a week." Others referred to "almost no learning curve…a near-zero learning curve."

Performance and Scalability

Most organizations had selected XIV systems only after tests had shown that this platform was capable of handling current as well as projected workloads. Typically, organizations required that systems be capable of handling three to five times current workloads, although some employed larger multiples.

XIV performance was, by wide margins, reported to be superior to that of existing disk systems. Overall, organizations reported that performance was superior to current-generation high-end and midrange disk arrays equipped with FC drives operating at 15,000 revolutions per minute (rpm).

This was the case for a wide range of applications. At the high end of the spectrum, XIV systems supported large-scale SAP ERP and other types of business-critical transaction processing systems, multiple-terabyte data warehouses and analytic systems, large Microsoft Windows Server and Exchange networks and VMware installations, high-volume streaming video and others.

Most comments about XIV performance did not deal with "raw" or "absolute" performance. Users repeatedly cited the ability of XIV systems to deliver consistent performance and meet service level agreement (SLA) targets for diverse, fluctuating workloads. This was particularly the case for workloads generated by virtualized server farms.

Several respondents attributed XIV strengths in this area to the system's architectural characteristics and load balancing capabilities. One noted, for example, that XIV's use of native wide striping allowed support for mixed workloads to be spread across all I/O facilities. This not only avoided bottlenecks, but significantly reduced system administration time and effort to maintain performance as workloads grew.

Other users noted that XIV systems were able to maintain performance even with high levels of data growth, capacity utilization or both. One organization, for example, reported no performance degradation at 90 percent utilization, and others reported similar experiences at 80 percent and higher. Most storage systems experience significant degradation at lower levels.

The ability of XIV systems to make extremely rapid copies of data was also cited as a major benefit by seven users. One organization was able to perform four differential snapshots per day over a seven-day period for a large (more than 7TB of raw data) Oracle database without performance degradation.

Availability and Recovery

Key XIV benefits were said to include hardware reliability (a function of massive disk redundancy) and extremely rapid data rebuilds in the event of a disk failure. Both are enabled by the XIV system's distinctive "RAID X" data distribution algorithm.

For many organizations, recovery times of 30 minutes for failed disks were considered acceptable, and were commonly specified in agreements. Users that had experienced failures reported, however, that disk rebuilds took between 3 and 15 minutes – the average was slightly over 7 minutes. In all cases, users experienced no performance degradation or interruption of service.

In some cases, larger-scale failures had been simulated. Results were similar. For example, two organizations simulated failure and recovery of entire 12-disk data modules, and another simulated near-simultaneous failure of two drives in different modules. In all cases, users again experienced no ill effects.

Also contributing to XIV availability were the ability to add data modules without taking systems offline, and snapshot capabilities that enabled organizations to minimize backup windows. These were variously reduced from 8 hours to around 15 minutes, and 6 to 10 hours to less than 5 minutes. Others reported similar experiences.

In addition, 27 out of 63 organizations (43 percent) reported use of XIV asynchronous or synchronous replication for disaster recovery failover. XIV systems were said to provide the same functionality as high-end tools at a significantly lower price.

TECHNOLOGY VIEW

IBM XIV Platform

Overview

IBM XIV systems implement an architecture, typically characterized as "grid storage" or "grid parallelism," which is entirely different from conventional high-end disk systems such as VMAX.

Its basic principle is that system resources (including processors, cache and disks) are packaged into self-contained nodes, and interconnected so that any node can communicate with any other without data passing through a centralized switch. Synchronization of resources across nodes is handled using massively parallel software.

This approach means that system-level performance is less dependent on such variables as controller throughput and disk drive speed. High levels of system-wide performance may be realized with comparatively low-speed – and low-cost – high-density SAS or SATA drives.

In the XIV design, workloads are spread across and executed by all system nodes, and all I/O resources are available to all applications. XIV algorithms also automate data placement, provide unified load balancing, and manage capacity utilization across the entire system.

Other capabilities include a distributed cache implementation that eliminates cache mirroring and locking latencies, thin provisioning, industry-leading snapshot technology, and a GUI-based administrator interface that is widely regarded as one of the most streamlined and productive in existence today.

The XIV operating system, along with tools for asynchronous and synchronous replication, snapshot copying and cloning, data migration, multipathing, host connectivity and management functions, are offered in an integrated single-price package. Providing comparable functionality for VMAX systems requires additional, separately charged tools.

Performance Characteristics

In performance terms, XIV architecture is significantly more efficient than conventional disk systems when handling diverse, mixed workloads. This is particularly the case when these are executing concurrently. In comparison, VMAX systems – as proponents never tire of pointing out – are better optimized to handle fewer, simpler I/O-intensive workloads.

An additional XIV benefit is that the combination of distributed data and high-performance snapshots enables data stores for VMware, Hyper-V and equivalent applications to be more easily copied and backed up. The efficiency of XIV snapshot technology means that disk-intensive snapshot processes do not disrupt production operations.

It can be expected that these strengths will be reinforced by use of SSD caching technology, which IBM has indicated will become available for XIV systems in the first half of 2012. This subject is discussed in more detail later in this section.

XIV systems software also provides performance optimization functions that, in conventional disk systems, are handled using separate tools and processes. This typically involves a great deal of work by system administrators. Tiering increases the amount of work required.

Availability Features

For availability purposes, the XIV system employs a RAID-like data distribution algorithm (sometimes referred to as RAID-X) that spreads data in small – one megabyte – blocks across all of the system's drives. While this approach reduces useable capacity, it provides a higher level of redundancy than conventional systems, and enables extremely rapid data rebuilds in the event of a disk failure.

In comparison, VMAX systems employ a combination of RAID 5 and spare drives to guard against the effects of drive failures. The RAID 5 technique spreads data across multiple drives, enabling it to be rebuilt if one drive fails. In a typical Tier 1 configuration, drives are configured in a 3:1 ratio, with one of every four drives acting as a failover device (i.e., useable capacity is 75 percent of physical capacity). Spare drives provide further protection.

Minimal requirements for system tuning, along with the ability to use XIV snapshots to handle data movement concurrently with production workloads, and other XIV capabilities also facilitate maintenance of high levels of uptime.

IBM XIV Gen3 Systems

The new XIV Gen3 hardware platform incorporates technology upgrades – summarized in figure 6 – that accelerate performance in a multiple areas.

Technology	XIV Second Generation Systems	XIV Gen3 Systems
Disk type	1TB, 2TB SATA	2TB SAS, 3TB SAS
Processor	Intel L5410 – 2.33 GHz quad core single-threaded	Intel E5620 – 2.4 GHz quad core hyperthreaded (2 threads/core)
Cache	8 or 16GB per module (max. 120 or 240GB)	24GB per module (max. 360GB) SSD Caching (future)
Cache-to-disk bandwidth	Max. 240 Gbps	Max. 480 Gbps
External connectivity	8-24 x 4 Gbps FC ports 0-6 x 1 Gbps iSCSI ports	8-24 x 8 Gbps FC ports 6-22 x 1 Gbps iSCSI ports
Internal connectivity	2 x 48-port Gigabit Ethernet (1 Gbps)	2 x 36-port InfiniBand (20 Gbps)

Figure 6: ***Principal IBM XIV Performance Features***

These upgrades are implemented in a mutually reinforcing manner; i.e., the effect on overall system-level performance is cumulative.

According to test results published by IBM, latest-generation systems deliver significantly higher performance than earlier equivalents. These results, which are summarized in figure 7, are consistent with the technology upgrades described above.

TEST/METRIC	XIV Second Generation Systems	XIV Gen3 Systems
Microsoft Exchange Solution Reviewed Program (ESRP) Storage Version 3.0 – 1GB mailbox quota	40,000 mailboxes	120,000 mailboxes
Representative Oracle Data Warehouse workload run with Oracle I/O Calibration (ORION) storage sizing suite	55,000 I/Os	115,000 I/Os
SAS Business Analytics Reports using Oracle Swingbench load generator	70 reports	207 reports

Figure 7: ***IBM XIV Comparative Performance Results***

Although actual results will vary by application and workload, an overall performance improvement of two times or more for XIV Gen3 systems can be reasonably expected. These will probably now deliver over 100K IOPS.

IBM has also indicated plans to offer an ***SSD read caching option*** in the first half of 2012. In addition, the company plans to employ 512GB SSD per module and up to 7.5TB per system to provide a secondary cache layer between DRAMs and disks.

IBM has targeted SSD caching at applications generating high proportions of random reads, such as messaging (e.g. Microsoft Exchange), file serving (e.g. Windows File Services) and transactional systems. XIV implementation of this technology is expected to reduce I/O latency by up to 90 percent.

EMC Solutions

EMC Symmetrix VMAX Platform

The EMC VMAX platform has replaced the EMC DMX generation of systems as the company's flagship high-end platform. In moving to VMAX, proprietary controllers were replaced by x86 servers, commodity components were drawn from the company's Clariion midrange platform, and RapidIO technology was adopted for internal interconnects.

RapidIO is an Ethernet and PCIe alternative originally developed by Motorola (now Freescale) Semiconductor. Currently, approximately 40 OEMs use RapidIO in communications, military, industrial and other electronics products.

VMAX systems are configured using System Bays, which contain engines and up 120 drives; and Storage Bays, which may contain up to 240 additional drives. Configurations may include up to eight engines and 2,400 disk drives, which may be SSD, FC and/or SATA devices.

Since their introduction, VMAX systems have employed Intel quad-core 2.33 GHz Xeon processors. The same processors also powered earlier XIV systems. XIV Gen3 models have been upgraded to latest-generation Intel Westmere technology.

EMC FAST/FAST VP

VMAX has been closely associated by EMC with the company's Fully Automated Storage Tiering (FAST) scheme. The original version of FAST, introduced in 2009, supported allocation and re-allocation of data only in LUN increments. This greatly limited its utility. FAST was upgraded in January 2011 to enable movement of data blocks with sub-LUN granularity.

FAST is implemented through FAST Suite software. This consists of the core FAST VP module and FAST Cache, which employs up to 2TB of SSD to provide caching for currently-accessed data located on FC and/or SATA drives.

The basic principles of tiering are well established. By placing the most frequently accessed data on SSDs, and the least frequently accessed on SATA drives, organizations may improve performance and reduce hardware and facilities costs. The extent of such gains, however, is dependent upon applications and workload mixes.

In cost calculations for this report, a variety of configurations were constructed for use of FAST with VMAX systems employing SSD, FC and/or SATA drives. Figure 8 shows examples.

System 1	System 2	System 3

4 x SSD (4%) 33 x FC 15K rpm (30%) 72 x SATA II 7.2K rpm (66%) **Total: 109 drives**	2 x SSD (3%) 25 x FC 15K rpm (36%) 42 x SATA II 7.2K rpm (61%) **Total: 69 drives**	2 x SSD (3%) 30 x FC 15K rpm (45%) 34 x SATA II 7.2K rpm (52%) **Total: 66 drives**

Figure 8: ***Examples of EMC VMAX FAST Configurations Employed for Cost Comparisons***

Although there is still comparatively little user experience with EMC FAST, certain drawbacks have become apparent.

One is that repeated movement of blocks of data between disks may cause bottlenecks that impair production performance. These typically become more serious as the size of blocks and the frequency with which they are moved increase.

This issue may be addressed by moving data in batch mode during off-peak periods. EMC supports this approach, which has been adopted by numerous early FAST adopters. Availability suffers.

A second drawback is that a great deal of manual intervention by storage administrators is still required. Although data movement may be automated, this is not the case for other processes. EMC supplies a tool, Tier Advisor, for modeling configurations against performance. Even using this, however, administrative overhead may be greater than for less complex single-tier environments.

In comparing FAST to XIV systems, a broader difference also emerges. FAST processes are essentially reactive – administrators respond to bottlenecks after these occur. In contrast, XIV load-balancing processes are dynamic, enabling genuine real-time identification and prevention of emerging problems.

Differences between FAST and XIV are particularly significant for organizations that must support workloads that experience frequent and/or major changes over time. In such environments, FAST's reactive approach not only increases system administration overhead, but also magnifies risks of inefficient capacity utilization, performance shortfalls and outages.

FAST, in short, addresses challenges that the XIV architecture avoids.

EMC Symmetrix VMAXe Systems

Introduced in July 2011, VMAXe is a downsized version of the VMAX platform. VMAXe cannot be upgraded to VMAX systems, and is technically constrained to limit potential competition with the latter. Major differences between VMAXe and VMAX systems are summarized in figure 9.

Unlike VMAX, VMAXe does not support Fibre Connection (FICON) and Count Key Data (CKD), which effectively preclude attachment to mainframe systems. FICON is the IBM implementation of the Fiber Channel protocol for mainframe channels attachment, while CKD architecture maps mainframe data to disk storage hardware. A large segment of the EMC VMAX business involves mainframe connectivity.

Support for single-phase power is apparently designed to restrict use of VMAXe systems to smaller data centers. It is widely used for smaller loads (e.g., for small business and residential sites), and in rural areas where distribution costs are comparatively low.

	VMAXe	VMAX
Maximum engines	4	8
Intel processor	Westmere quad 2.4 GHz	Harpertown quad 2.33 GHz
Maximum disks	960	2,400
Maximum cache per engine	96GB	128GB
Connectivity	64 x 8 Gbps FC ports 32 x Ethernet ports	128 x 8 Gbps FC ports 64 x Ethernet ports
CKD & FICON support	No	Yes
Power	Single-phase only	Single- & three-phase

Figure 9: ***EMC VMAXe and VMAX Systems – Summary***

Larger, more advanced commercial and industrial sites tend to employ three-phase power, which is more energy-efficient, or combinations of three-phase backbones and single-phase feeds. Most major platforms, including VMAX and XIV systems, offer single- and three-phase options.

VMAXe is supported for FAST, and offered by EMC in standard single-, two- and three-tier configurations that are summarized in figure 10.

Single-tier	Two-tier	Three-tier
100% x 450GB FC 15K rpm	3% x 200GB SSD 97% x 2TB SATA II	3% x SSD 32% x FC 15K rpm 65% x SATA II 7.2K rpm

Figure 10: ***EMC VMAXe Standard FAST Configurations***

VMAXe systems are said by EMC to be "100% thin-provisioned," which presumably enables faster system set-up.

VMAXe systems support most major EMC VMAX software tools, although the company has substituted RecoverPoint for Symmetrix Remote Data Facility (SRDF), its high-end solution for asynchronous and/or synchronous replication for disaster recover purposes.

SRDF is widely used on EMC VMAX and earlier DMX platforms to support high-end business-critical systems in large organizations. RecoverPoint is intended primarily for use with EMC midrange systems.

BASIS OF CALCULATIONS

Methodology

Cost comparisons presented in this report are based on profiles of disk system installations in six user organizations. These are summarized in figure 11.

	Government Organization	Manufacturing Company	Life Sciences Company
Business profile	Local government 20+ departments & agencies 5,000+ employees	Industrial equipment manufacturer $2 billion sales 8,000+ employees	Biomedical products manufacturer $700 million sales 1,500+ employees
Applications	ERP, HR, departmental, Microsoft Exchange & Hyper-V	ERP, CRM & SCM systems, BI, departmental, VMware, XenServer	ERP system, research applications, PowerVM, VMware
	Telecommunications Company	**Health Care Organization**	**Financial Services Company**
Business profile	Landline, mobile, Internet & cable services $3 billion sales 15,000+ employees	20 hospitals & clinics 200+ other facilities $5+ billion revenues 40,000+ employees	Diversified retail bank $10 billion sales $200 billion assets 30,000 employees
Applications	BI, CRM, ERP, operational, hosting services, Microsoft Exchange, VMware, Xen	HIS, EMR, HR, PACS, departmental, Microsoft Exchange, VMware	BI, CRM, ERM, financial, HR, Lotus Domino, PowerVM, Sun Dynamic Domains, Windows & Linux VMs

BI: Business intelligence
ERP: Enterprise resource planning
EMR: Electronic medical records
SCM: Supply Chain Management
PACS: Picture Archiving and Communication Systems
CRM: Customer relationship management
HIS: Hospital information system
ERM: Enterprise risk management

Figure 11: ***Installations Summary***

Installations were constructed using data on applications, workloads, hardware and software configurations, storage administration staffing levels and growth trends supplied by 27 users of IBM XIV and EMC VMAX systems.

Configurations and staffing levels were then determined for four scenarios for each installation:

1. ***XIV systems*** scenarios are for use of Gen3 and (in smaller installations) earlier systems.

2. ***VMAX systems*** scenarios are for use of VMAX systems configured with 600GB FC drives.

3. ***VMAX and VMAXe systems*** scenarios are for use of VMAX systems with 600GB FC drives, and VMAXe single-tier configurations with 450GB FC rpm drives.

4. ***VMAX with FAST*** scenarios are for FAST-enabled VMAX systems configured with SSD, FC and/or SATA drives.

Configurations and staffing levels were then determined for each installation and scenario. These are summarized in figure 12.

Government Organization	Manufacturing Company	Life Sciences Company	Telecom Company	Health Care Organization	Financial Services Company
IBM XIV SYSTEMS					
Beginning of period					
2 x 27TB **Total: 54TB** ***0.2 FTE***	2 x 55TB **Total: 110TB** ***0.25 FTE***	2 x 87TB **Total: 174TB** ***0.35 FTE***	2 x 125TB **Total: 250TB** ***0.35 FTE***	2 x 125TB 1 x 87TB **Total: 337TB** ***0.5 FTE***	3 x 125TB 1 x 82TB 1 x 55TB 1 x 27TB **Total: 539TB** ***0.85 FTE***
End of period					
2 x 43TB **Total: 86TB** ***0.2 FTE***	1 x 149TB 1 x 134TB **Total: 283TB** ***0.25 FTE***	1 x 161TB 1 x 149TB **Total: 310TB** ***0.4 FTE***	1 x 240TB 1 x 223TB **Total: 463TB** ***0.4 FTE***	2 x 223TB 1 x 134TB **Total: 580TB** ***0.6 FTE***	1 x 240TB 1 x 223TB 1 x 187TB 1 x 125TB 1 x 87TB 1 x 27TB **Total: 889TB** ***1.05 FTEs***
EMC SYMMETRIX VMAX SYSTEMS					
Beginning of period					
1 x 25TB 1 x 19TB **Total: 44TB** ***0.55 FTE***	2 x 57TB **Total: 114TB** ***0.55 FTE***	1 x 87TB 1 x 74TB **Total: 161TB** ***0.6 FTE***	1 x 121TB 1 x 101TB **Total: 222TB** ***0.85 FTE***	1 x 113TB 1 x 76TB 1 x 69TB 1 x 38TB **Total: 296TB** ***1.5 FTEs***	3 x 101TB 1 x 131TB 1 x 88TB **Total: 522TB** ***2.15 FTEs***
End of period					
1 x 50TB 1 x 38TB **Total: 88TB** ***0.6 FTE***	2 x 139TB **Total: 278TB** ***0.65 FTE***	1 x 164TB 1 x 145TB **Total: 309TB** ***0.7 FTE***	1 x 284TB 1 x 183TB **Total: 467TB** ***1.0 FTE***	1 x 214TB 1 x 151TB 1 x 132TB 1 x 63TB **Total: 560TB** ***1.85 FTEs***	2 x 239TB 1 x 183TB 1 x 107TB 1 x 101TB **Total: 869TB** ***2.7 FTEs***
EMC SYMMETRIX VMAX & VMAXe SYSTEMS					
Beginning of period					
2 x 23TB **Total: 46TB** ***0.5 FTE***	2 x 55TB **Total: 110TB** ***0.5 FTE***	1 x 87TB 1 x 74TB **Total: 161TB** ***0.55 FTE***	1 x 121TB 1 x 101TB **Total: 222TB** ***0.85 FTE***	1 x 113TB 1 x 76TB 1 x 69TB 1 x 38TB **Total: 296TB** ***1.5 FTEs***	1 x 132TB 3 x 101TB 1 x 88TB **Total: 523TB** ***2.15 FTEs***
End of period					
1 x 51TB 1 x 37TB **Total: 88TB** ***0.5 FTE***	2 x 139TB **Total: 278TB** ***0.55 FTE***	1 x 158TB 1 x 144TB **Total: 302TB** ***0.6 FTE***	1 x 284TB 1 x 183TB **Total: 467TB** ***1.0 FTE***	1 x 214TB 1 x 151TB 1 x 132TB 1 x 63TB **Total: 560TB** ***1.85 FTEs***	2 x 239TB 1 x 181TB 1 x 107TB 1 x 101TB **Total: 867TB** ***2.7 FTEs***
EMC SYMMETRIX VMAX SYSTEMS with FAST					
Beginning of period					
1 x 24TB 1 x 18TB **Total: 43TB** ***0.55 FTE***	2 x 56TB **Total: 112TB** ***0.55 FTE***	1 x 87TB 1 x 75TB **Total: 162TB** ***0.65 FTE***	1 x 121TB 1 x 101TB **Total: 222TB** ***0.9 FTE***	**1** x 125TB 1 x 73TB 1 x 67TB 1 x 36TB **Total: 301TB** ***1.5 FTEs***	1 x 133TB 2 x 101TB 1 x 100TB 1 x 87TB **Total: 522TB** ***2.35 FTEs***
End of period					
1 x 50TB 1 x 38TB **Total: 88TB** ***0.6 FTE***	2 x 138TB **Total: 276TB** ***0.65 FTE***	1 x 162TB 1 x 144TB **Total: 306TB** ***0.75 FTE***	1 x 283TB 1 x 183TB **Total: 466TB** ***1.1 FTE***	1 x 216TB 1 x 152TB 1 x 129TB 1 x 62TB **Total: 559TB** ***1.9 FTEs***	2 x 238TB 1 x 182TB 1 x 106TB 1 x 99TB **Total: 863TB** ***2.9 FTEs***

Figure 12: ***Configurations and Staffing Summary***

For all scenarios, configurations shown are for useable capacity. For EMC scenarios, usable capacities were calculated by ITG based on use of RAID 3 +1 configurations with eight spares per 100 FC and/or SATA drives, and a minimum of one spare for solid state drives. Allowance was also made for cache vaulting and Symmetrix File System space.

For calculation purposes, XIV models were equipped with the standard XIV systems software package, while VMAX and VMAXe systems were configured with EMC software providing equivalent functionality. These include the products shown in figure 13.

Function	VMAX	VMAXe
Operating system	Enginuity 5875	Enginuity 5875e
Replication	Symmetrix Remote Data Facility	RecoverPoint
Snapshot copies	TimeFinder	TimeFinder
Storage management	Symmetrix Manager	Symmetrix Management Console
Performance optimization	Symmetrix Optimizer	Symmetrix Performance Analyzer
Multipathing	PowerPath	PowerPath
Migration tools*	N/A	Symmetrix Migration Suite, Open Replicator for Symmetrix, Replication Manager

*Included in base software bundle

Figure 13: ***EMC Software Products Employed in Calculations***

The principal differences between VMAX and VMAXe software stacks are that RecoverPoint substitutes for SRDF for remote replication for failover and recovery, and management and optimization tooling is more limited. Cost calculations include RecoverPoint appliances installed at all sites where remote replication is employed. VMAX with FAST calculations also include FAST Suite software.

EMC offers a number of data migration tools as part of the base VMAXe software bundle at no additional cost. Not all VMAX and VMAXe systems were configured with all software products; e.g., SRDF and RecoverPoint software were employed only on systems whose contents were replicated to a second site for disaster recovery purposes.

Costs were calculated as follows.

System costs include initial hardware and software acquisition, as well as costs of subsequent upgrades over three-year periods.

As EMC offers a full three-year 24/7 onsite warranty for VMAX hardware, no maintenance costs were included in calculations for this scenario. The VMAXe warranty, however, covers only next-business-day onsite response. Calculations for this platform include upgrades to provide 24/7 onsite response.

EMC offers a 90-day software warranty for its software products. However, this applies only to defects in the media on which software is supplied. Since such defects are rare, no allowance is made for this warranty period in calculations.

Maintenance costs for XIV systems include 24/7 hardware maintenance and software support. Costs for all platforms were calculated using street prices; i.e., discounted prices paid by users.

Personnel costs for both platforms were calculated using an annual average FTE storage administrator salary of $71,295 increased by 51.2 percent to allow for benefits, bonuses, training and related items.

Facilities costs were calculated based on IBM and EMC specifications or, where sufficient data was not available from these vendors, were estimated by ITG. Calculations include costs of power consumed by data center infrastructure equipment such as power distribution systems, computer room air conditioning (CRAC) systems and chillers.

Power consumption was calculated based on specific utilization levels and hours of operation for each installation. A conservative assumption for average price per kilowatt/hour was employed to determine three-year power costs.

Occupancy costs were calculated using a conservative assumption for annual average cost per square foot for existing facilities (i.e., costs do not include new facilities construction). All calculations allow for capacity growth over the three-year measurement period.

All costs are for the United States.

Cost Breakdowns

Detailed cost breakdowns are presented in figure 14.

	Government Organization	Manufacturing Company	Life Sciences Company	Telecom Company	Health Care Company	Financial Services Company
IBM XIV SYSTEMS						
System	415.0	676.7	723.4	1,053.9	1,347.5	2,721.8
Maintenance/support	65.2	117.7	136.4	141.6	190.9	339.3
Personnel	64.7	80.8	117.5	118.6	177.9	307.2
Facilities	24.5	29.0	33.3	27.2	37.2	65.0
TOTAL ($000)	**569.4**	**904.2**	**1010.6**	**1,341.3**	**1,753.5**	**3,433.3**
EMC SYMMETRIX VMAX SYSTEMS						
System	1,256.9	2,764.8	2,890.2	3,703.1	5,134.8	8,243.8
Maintenance/support	163.2	372.0	355.4	255.9	485.8	939.2
Personnel	183.3	194.0	210.2	296.4	544.4	781.5
Facilities	25.3	52.0	74.8	69.6	118.0	189.9
TOTAL ($000)	**1,628.7**	**3,382.8**	**3,530.6**	**4,325.0**	**6,283.0**	**10,154.4**
EMC SYMMETRIX VMAX & VMAXe SYSTEMS						
System	883.0	1,954.6	2,202.2	3,703.1	5,134.8	8,031.5
Maintenance/support	218.9	429.1	485.8	255.9	485.8	886.6
Personnel	161.7	177.9	183.3	296.4	544.4	781.5
Facilities	33.0	68.8	102.6	69.6	118.0	205.0
TOTAL ($000)	**1296.6**	**2,630.4**	**2,973.9**	**4,325.0**	**6,283.0**	**9,904.6**
EMC SYMMETRIX VMAX SYSTEMS with FAST						
System	1,120.8	1,946.7	2,122.9	3,263.3	4,196.1	5,513.6
Maintenance/support	115.3	238.7	265.1	364.4	613.3	693.5
Personnel	183.3	199.4	226.4	323.4	555.2	846.2
Facilities	7.5	14.5	19.3	16.1	42.4	87.9
TOTAL ($000)	**1,426.9**	**2,399.3**	**2,633.7**	**3,967.2**	**5,407.0**	**7,141.2**

Figure 14: ***Three-year Cost Breakdowns***

ABOUT THE INTERNATIONAL TECHNOLOGY GROUP

ITG sharpens your awareness of what's happening and your competitive edge
. . . this could affect your future growth and profit prospects

International Technology Group (ITG), established in 1983, is an independent research and management consulting firm specializing in information technology (IT) investment strategy, cost/benefit metrics, infrastructure studies, deployment tactics, business alignment and financial analysis.

ITG was an early innovator and pioneer in developing total cost of ownership (TCO) and return on investment (ROI) processes and methodologies. In 2004, the firm received a Decade of Education Award from the Information Technology Financial Management Association (ITFMA), the leading professional association dedicated to education and advancement of financial management practices in end-user IT organizations.

The firm has undertaken more than 120 major consulting projects, released more than 250 management reports and white papers and more than 1,800 briefings and presentations to individual clients, user groups, industry conferences and seminars throughout the world.

Client services are designed to provide factual data and reliable documentation to assist in the decision-making process. Information provided establishes the basis for developing tactical and strategic plans. Important developments are analyzed and practical guidance is offered on the most effective ways to respond to changes that may impact complex IT deployment agendas.

A broad range of services is offered, furnishing clients with the information necessary to complement their internal capabilities and resources. Customized client programs involve various combinations of the following deliverables:

Status Reports	In-depth studies of important issues
Management Briefs	Detailed analysis of significant developments
Management Briefings	Periodic interactive meetings with management
Executive Presentations	Scheduled strategic presentations for decision-makers
Email Communications	Timely replies to informational requests
Telephone Consultation	Immediate response to informational needs

Clients include a cross section of IT end users in the private and public sectors representing multinational corporations, industrial companies, financial institutions, service organizations, educational institutions, federal and state government agencies as well as IT system suppliers, software vendors and service firms. Federal government clients have included agencies within the Department of Defense (e.g. DISA), Department of Transportation (e.g. FAA) and Department of Treasury (e.g. US Mint).

International Technology Group

609 Pacific Avenue, Suite 102
Santa Cruz, California 95060-4406
Telephone: + 831-427-9260
Email: Contact@ITGforInfo.com
Website: ITGforInfo.com