



# External Storage Systems:

*Market Analysis and Decision Criteria*

by John Cole  
Technical Architect  
Ahead

## Abstract

With a multitude of external storage systems confronting today's user, selecting the "right" one can be a difficult, time-consuming, and costly task. It's not uncommon for manufacturers to exacerbate this effort by creating high-level marketing descriptions that only focus on describing a product's target market, but lack feature and function definition. This paper analyzes five external storage market subspaces and provides six decision criteria to help you as a consumer determine the appropriate storage solution.

## Introduction

In an attempt to establish the terms of the competitive evaluation, many solutions claim to address multiple storage market segments. This paper will cover five market-agreed upon categories:

1. High-End (or Enterprise) systems
2. Enhanced Midrange (or Tier 1.5) systems
3. Midrange (or Mid-Tier) systems
4. Entry Level (or SMB) systems
5. Specialty systems

Don't be misled by category names, however; manufacturers often use product categories to denote the size of the customer being targeted—not the product's features. For example, a 50-person company might deploy a high end or "Enterprise" solution targeted to a large organization for critical data or system availability, while a Fortune 1000 company might purchase an entry level or "SMB" solution for a remote office or departmental need.

To help guide you to the system appropriate for your business, here are feature-based or functions-based descriptions of the five storage product categories, along with comments based on our experience and that of our clients.

---

---

*\$21 billion external storage market*

*<http://tinyurl.com/2010-growth>*

---

---

## External storage system categories and definitions

### High-End systems

The largest and most robust storage architectures, High-End systems often are referred to as Enterprise systems, due to their target profile customer. This segment is relatively small because a limited number of applications drive its requirements. The demands for uncompromising uptime and reliability are the key decision criteria. Typical customers in this segment tend to accept lagging behind the “latest and greatest” feature set. Instead, they favor higher data integrity and error correction over the latest functionality. Key target markets for Enterprise systems include financial services, highly regulated industries with downtime reporting penalties, healthcare services, e-commerce sites, logistics firms, and government agencies. In mixed environments these systems are deployed on the most mission-critical applications, usually only for Tier 1 and 2 needs.

High-End (or Enterprise) systems	
Defining characteristics	Leading vendors
N+2 to N+N redundancy of all internal components, including controllers	EMC (Symmetrix) Hitachi Data Systems (Virtual Storage Platform) IBM (DS8000 series*)
Mainframe connectivity (ESCON/FICON)	<i>*IBM's DS8000 series of storage is included in this market segment due to its support of mainframe and other high-end feature sets. However, it has a significant difference from the other vendors in this market space: The controller architecture of the DS8x00 is a dual-controller design, with each controller being an IBM P Series server. This means each controller has extremely high resiliency and redundancy, but at its core it is still a dual-controller architecture, with the inherent challenges of that design. IBM created this design intentionally, as it is the best architecture to support mainframe storage, providing a limited number of controllers to target for storage channels. In certain environments, this can be a strength in terms of performance or responsiveness. The dual-controller limitations have caused DS8x00 to fall behind offerings from the other manufacturers in recent years, resulting in a decline in interest from the general marketplace.</i>
Multi-site and/or multi-hop replication capabilities	
Non-disruptive upgrades	
Active/Active controller sets	
Multiple controllers (3+)	
Per terabyte licensing structures	
Focus on availability	

*Important points and common misconceptions*

- High-End systems are often assumed to be the highest performance systems available, but this is not necessarily valid. Designed for availability first and performance second, Enterprise systems frequently are outperformed by other architectures because of differences in error correction and event handling.
- High-End systems are often assumed to have the highest cost for an environment, but cost actually depends on the software options selected.
- High-End systems do not automatically improve the availability of an entire compute environment, because necessary steps for multi-pathing and connectivity resiliency still must be accounted for.

**Enhanced Midrange (Tier 1.5) systems**

The Enhanced Midrange, or Tier 1.5, system is a relatively new segment of the market. Typically, these solutions offer a scale-out design for storage controllers and go beyond the traditional dual-controller design of Midrange systems. However, they usually do not offer the true Active/Active or any-to-any model of the High-End systems. Instead, Enhanced Midrange systems purport to offer performance and availability near the levels of High-End systems, but at a price point closer to the Midrange offerings. As a result, they occupy a niche with service providers, which seek to provide IaaS solutions at an ever-more competitive price point to end users. Tier 1.5 systems, frequently deployed for application tiers 1, 2, and 3, have the least in common with each other—cost, not features and functions, drives the segment.

Enhanced Midrange (Tier 1.5) systems	
Defining characteristics	Leading vendors
N+1 to N+N redundancy of internal components, including controllers	EMC (VMAXe)* HP (3PAR) IBM (XiV)
Focus on scale-out with minimal management required	NetApp (FAS with cluster mode only) <i>*EMC's VMAXe offering is a version of its full Symmetrix platform, with only certain software capabilities removed (such as mainframe support, SRDF to a VMAX, and clones). All the hardware and availability components match the high-end version of the platform.</i>

*Important points and common misconceptions*

- Due to the unique nature of Enhanced Midrange controller designs, they often are compared to Enterprise or Midrange class systems, though neither is a real match.
- A system with more than two controllers does not automatically increase the availability of a particular volume in that system.
- Software implementation on these systems, with respect to controller cooperative behavior, is critical to understanding the strengths and limitations of these platforms.
- The emergence of these technologies has created significant overlap with existing storage architectures, and manufacturers often have multiple products to meet a specific need or requirement. This can result in confusion, with projects generating multiple proposed designs from the same manufacturer. To avoid this problem, provide clear and specific requirements so vendors can respond with their best alternative and demonstrate why their other products may fit, albeit imperfectly.

## Midrange systems

Midrange systems are the workhorse of the external storage market. Responsible for the lion’s share of the revenue, Midrange is the most contested and most populated section of the market. It’s also the breeding ground for much of the innovation in the storage market, thanks to the competitive forces driving the segment and the fact that they are usually built from commodity hardware and operating systems. Midrange systems, designed to be flexible and adaptable for a varying set of use cases across customer segments, are deployed as any mix of Tier 1, 2, 3, and 4 application support, depending on customer needs.

Midrange systems	
Defining characteristics	Leading vendors
N+1 to N+2 redundancy of all internal components	Dell (Compellent, EqualLogic)
Highly available upgrades	EMC (CLARiiON/VNX)
Active/Passive controller sets or LUN ownership model	Hitachi Data Systems (AMS)
Per-frame licensing model	HP (EVA)
Focus on ease of use and flexibility	IBM (DS5000 series–Engenio)
	NetApp (FAS, Engenio)

### Important points and common misconceptions

- Flexibility associated with connectivity is a key decision point when considering midrange systems. The support of native NAS (CIFS/ NFS) protocols and their use cases can be a key driver.
- While Midrange system maintenance or firmware upgrades often are described as non-disruptive, they usually involve controller rebalancing of volumes and/or operating for a period of time off a single controller. Understanding such impacts to your business is critical in choosing the right platform.
- These systems undergo the highest level of change and development in any given time period. When evaluating these systems, get an updated product roadmap if it has been more than 24 months since you received a full review of the product’s capabilities, functions, and development priorities.

## Departmental systems

Departmental, or SMB, solutions usually are deployed for a small environment and treated as a central storage pool for a limited number of hosts or as extended direct attached storage for a single host or cluster. In larger customers, departmental systems may start as a small island of storage to meet a specific need, but they often sprawl out, creating management problems. There is a place for SMB solutions in small and large customers, depending on circumstances. For small customers they can represent the entire storage footprint; for larger customers they're almost exclusively used as static applications, tier 4 applications, or in small remote facilities.

Departmental systems	
Defining characteristics	Leading vendors
Optional redundancy of internal components, including controllers	Dell (EqualLogic, PowerVault) EMC (AX/VNXe, Iomega)
Disruptive upgrades	HP (LeftHand, MSA)
Per-frame licensing model	NetApp (FAS 2000 series)
Focus on ease of use	
Limited feature sets	

### *Important points and common misconceptions*

- These systems often have low acquisition costs and are perceived as good values, but over time they can result in a high total cost of ownership due to increased management, system rebalancing, and high power/cooling costs per TB deployed, and limited scalability. If multiple systems are required to meet a storage need, a larger, consolidated system usually will be a better solution.
- Many times the low price point puts SMBs within a manager's purchasing authority, and it can be acquired without regard for corporate standards. This is particularly challenging now with 25-40 TB systems available out of the box.

## Specialty systems

The breadth and depth of specialty or single-purpose solutions offered in today’s storage market puts them outside the scope of this paper. Solutions in this segment usually have been developed to meet a specific use case or customer need. In some situations, this segment is served by a combination of software and commodity hardware; other times it is packaged as a complete offering.

Specialty systems	
Selected use cases	Leading vendors
Scale-out file systems	EMC Isilon, HP PolyServe, IBM SONAS, Gluster, others
Compliance-based archives with content application integration	EMC Centera, IBM GMAS, HDS HCAP, others
Data warehousing and decision-support appliances	EMC Greenplum, HP Vertica, IBM Netezza, others
VMware-aware storage	Tintri
SSD appliances for data caching or extreme low latency	Pure Storage, Violin Memory, Nimbus, Kove, others

---

### Why so many choices in the external storage systems market?

To understand the differences in today’s platform offerings, consider the storage industry’s history. You will see that the original use cases and design criteria continue to have an impact on the culture and design of the current-generation systems.

IBM and EMC were two of the first entrants in the external storage market, with offerings to complement mainframe systems in the late 1980s and early 1990s. Because these storage subsystems were attached to mainframes, they inherited the mainframe’s characteristics and design criteria: redundancy and uptime as the primary drivers, followed by performance.

Requirements of open systems teams in the mid-1990s to the early 2000s drove the next set of market entrants. Usability, flexibility, and performance were at a premium, followed by availability. Acquisition cost became a much more important factor in this space.

The next round of storage system entrants, entering the market from the early to mid-2000s, focused on newer connectivity protocols and addressed the gaps in customer needs ignored by the legacy mainframe-based or open-systems-based systems.

Today’s market is going through a burst of new systems geared exclusively towards hypervisor integration (VMware) or solid-state media. Manufacturers’ marketing

---

---

materials often attempt to explain their platforms in the context of current market drivers, which include:

- Virtualization integration (VMware affinity)
- Emergence of SSD as a primary tier for storage (Flash storage)
- Extreme storage density for SATA devices (3 TB+ per drive slot)
- Extreme data growth across enterprises, especially in unstructured content (scale-out file storage)
- Extreme demand for capacity growth in decision-support systems (big data)
- Connectivity convergence and the rise of IP-based storage connectivity
- Advanced data services, including snapshots, deduplication, thin provisioning, and object support

In each wave of this evolution, the existing platforms inherited new requirements to compete, but usually maintained their legacy heritage in the process. The new requirements may be at odds with the original design tenets of the frame, resulting in an end product that reflects design compromises or a dilution of core product strengths. This doesn't mean the latest product iteration will not meet current requirements, but it does force customers to understand the differences between current and prior product generations rather than rely on historical impressions. Assume that each higher-end segment is a superset of the functions of the lower tiers, with added features or functions resulting in a price premium.

---

## Selecting the appropriate architecture

It's important to recognize that although all the systems described here will store data, variations and differences will occur, based on each vendor's implementation. You will enhance your opportunity to make the best decision and to acquire a more efficient solution by clearly articulating your unique business requirements and objectives to your vendors and partners.

Our experience suggests that six factors merit your consideration as you select the appropriate storage architecture from all the alternatives in the market.

### 1. Availability

System availability and the tolerance for risk is one of the basic requirements determining the appropriate class of storage for an organization. As the tolerance for planned and unplanned downtime decreases, so do the options for storage architectures. If a storage architecture must provide continuously available operations to the business, only a High-End architecture will suffice.

### 2. Protection

Data protection, both local and remote, can determine the difference among storage classes and among vendors. The requirement for complex, multi-site, multi-hop, or long-distance replication without data loss also dictates a High-End architecture. The granularity of replication traffic and the amount of acceptable data loss can determine the underlying storage architecture.

### 3. Performance

Customers usually cite performance as their primary decision criteria, though it is often poorly understood as a purchasing metric. Most of today's systems can perform above and beyond what is needed

for an enterprise application in terms of latency, bandwidth, and total I/O operations. To accurately gauge your performance needs, be prepared to:

- Offer application or operating system statistics on storage counters, including I/O, bandwidth, queue depth, and wait values.
- Define acceptable response times in milliseconds by application tier.
- Describe the variances in your performance workload during a daily, weekly, and monthly cycle.
- Include your expectations for performance of backups, decision support load, batch process, and replication tasks.
- Articulate expected growth in performance needs, including new applications or versions.

#### 4. Scalability and growth

Scalability and growth are critical metrics in sizing the initial footprint of an array and planned growth requirements. Describe your expected capacity growth by application or by tier, at either monthly or yearly compounded rates. Industry-standard metrics for an enterprise application are 30-35%, although much of that growth is weighted to file systems rather than databases. When considering data growth, be mindful of drive slots available in the array, along with the size of each drive. Financial considerations are important in a scaling conversation, as a factor in planning how much capacity to purchase up front and when to add upgrades. Based on the architecture, each array has a scalability breaking point—typically operational complexity and cost—that must be considered.

#### 5. Operational complexity

Operational capabilities and avoiding complexity will help determine the number and type of storage architectures appropriate for a given environment. When possible, consolidating management tools or storage systems is preferable to limit administrative overhead.

Consider the different operational and replication run books required when deploying each additional architecture or platform to determine the cost/benefit of adding it to your environment.

## **6. Cost**

Cost is always a critical factor in determining appropriate storage architecture. There are multiple dimensions to costs, including acquisition, support, administration, physical footprint, and upgrades. Evaluating the business impact of system downtime and being able to quantify the importance of data availability will help you select the appropriate storage cost model.

## **Conclusion**

To select the appropriate storage architecture, first understand and articulate your specific requirements with respect to availability, performance, protection, scalability, complexity, and cost. Ahead has developed a methodology for scoring and weighting these metrics based on criteria critical to your organization and the success factors of your project. After quantifying your tolerance for each of these metrics, you can look at the most demanding set of requirements and determine which architecture satisfies them. This process will eliminate either the lower functionality or higher functionality segments, leaving you with a subset of the market to consider. From that point, you can look to the lowest cost remaining market segment that meets your minimum requirements.

Having reduced the number of architectures under consideration, you can then use your decision matrix to evaluate the offerings in that market segment and select the platform best suited to your organization's needs and objectives.

As a technical architect at Ahead, John Cole's responsibilities include emerging technologies, stack computing, product direction, and strategic vision for his clients. He has been involved in the VCE alliance as an EMC evangelist for over two years, and is recognized regionally as one of the thought leaders on private cloud and fully virtualized data center deployments.

John is a ten-year veteran of the IT industry, as both a customer and a vendor. During his six-year EMC tenure, he delivered projects for customers, worked as a pre-sales consultant, and managed a team of consultants. In his two-year role as a principal for EMC Corporation, he was the technical lead of the nine-state Midwest Commercial Division. He worked primarily with Mid-Market and Vertical Customers, including Healthcare and Higher Education.

#### **About Ahead**

*Ahead is a leading provider of next-generation data center solutions, with specific expertise in virtualization and cloud computing architectures. Ahead designs and implements agile, service-oriented architectures to help clients transform their data centers from a technology-focused environment to a service-delivery platform that brings agility to the business and drives out cost. Solutions are built using a unique, simple, scalable, and repeatable methodology that speeds delivery and lowers risk.*

*Ahead, purpose built to drive private cloud enablement, employs a unique methodology called **THINK | LOOK | PLAN | MOVE**.*

*In the **THINK** phase Ahead works with clients to identify challenges and to understand and define current business objectives and strategies.*

*In the **LOOK** phase Ahead identifies all service elements and dependencies, performs gap analysis, and evaluates appropriate technologies.*

*In the **PLAN** phase Ahead develops architectures and TCOs, along with tactical and strategic implementation plans.*

*Finally, during the **MOVE** phase Ahead initiates its plan and executes on delivery of its offerings.*

*For more details, visit [www.ThinkAheadIT.com](http://www.ThinkAheadIT.com).*

The material in this document is the proprietary information of Ahead, LLC.

All products, trademarks, and copyrights herein are the property of their respective owners.

©2011 Ahead, LLC. All rights reserved. revised 10/21/11