

COMPETITIVE ANALYSIS

IDC MarketScape: U.S. Open Systems Virtual Tape Library 2009 Vendor Analysis

Laura DuBois

Robert Amatruda

IDC OPINION

Disk is increasingly being placed in the data protection path to improve upon existing backup processes. Customers need faster backups to meet shrinking backup windows and shorter recovery times to satisfy service-level objectives. These firms have embraced virtual tape library (VTL) technology. VTL solutions enable an enterprise to protect its existing investments in tape infrastructure, people, and processes without introducing unnecessary risk into the data protection environment. Moreover, corporate information and institutional knowledge resident in existing tape infrastructures can be preserved according to established local, offsite, and/or long-term retention requirements. The VTL market is quickly shifting into a broader disk-based data protection market that includes not only tape emulation and the ability to export to physical tape but also support for deduplication, remote replication, standard network interfaces, and the ability to become a data protection platform for different backup, archive, and recovery workloads. This IDC study introduces a vendor assessment model called IDC MarketScape, which provides a quantitative and qualitative assessment of the characteristics that explain a vendor's success in the VTL marketplace. Evaluations are based on a comprehensive framework that assesses vendors relative to the criteria and one another and highlights factors most influential to success in the market, currently and in the next three to five years. IDC has identified the following criteria as factors that materially contribute to vendor success in the U.S. open systems VTL market:

- Brand, distribution channels, breadth of storage offerings, and access to customers
- Low TCO/ROI, including reduction of storage capacity and footprint requirements
- Feature functions, including scalability, performance, deduplication, replication, physical tape support, and a strategy that includes expansion beyond tape emulation

TABLE OF CONTENTS

	P
In This Study	5
Methodology	5
Situation Overview	6
Introduction	6
Market Strategies.....	8
Future Outlook	11
IDC MarketScape U.S. Open Systems VTL Market Vendor Assessment.....	11
Vendor Summary Analysis	12
Essential Guidance	21
Learn More	23
Related Research.....	23

LIST OF TABLES

P

- 1 Key Strategy Measures for Success: U.S. Open Systems VTL Market 8
- 2 Key Capability Measures for Success: U.S. Open Systems VTL Market 10

LIST OF FIGURES

P

1 IDC MarketScape U.S. Open Systems VTL Vendor Assessment 12

IN THIS STUDY

This IDC study evaluates the major vendors in the U.S. open systems VTL market. It utilizes a sophisticated scoring and ranking model based on both qualitative and quantitative criteria that results in a single, easy-to-understand graphical illustration of each vendor's position in the market. The study is supported by a discussion of the key criteria that contribute to a vendor's current and future success in the market and highlights the differentiators that play a role in each participant's placement on the graph.

Methodology

The IDC MarketScape is designed to provide an overview of the competitive fitness of major providers in the U.S. open systems VTL market. For the purposes of feasibility and timing, most material vendors in the open systems VTL market were included. A single chart displays each given company's market share and indicates whether it is over- or underperforming and how well it is suited to compete in the market today and in the future (three to five years from now). The accompanying text explains each contender's major strengths and weaknesses, threats and opportunities.

IDC employs the following method to arrive at each company's ranking:

- ☒ **Sources.** The study is based on a model that is populated with data provided to IDC from a vendor questionnaire, companies' quarterly and annual reports, earnings calls, industry analyst events, interviews with company representatives, IDC primary research, and news coverage.
- ☒ **Market share, growth rates, and revenue.** This IDC MarketScape covers U.S. open systems VTL revenue. For companies that do not publicly disclose this revenue, IDC estimates on revenue and growth rates are based on supplier discussions, public information, and industry knowledge.
- ☒ **Competitive fitness.** Each major competitor's preparedness for current and future market conditions is expressed as a set of two scores. One score expresses a given vendor's current capabilities, while the other expresses the appropriateness of its strategies for the future. IDC bases its assessment of future market conditions on what most likely will be the market's major trends and disruptors. Each of these two scores is broken down into three criteria (product offering, go to market, and business), each of which in turn is broken down into several subcriteria. Both criteria and subcriteria are weighted by importance. For each company, we score its qualities with regard to each of the subcriteria, assigning a numeric value between 1 (worst) and 5 (best). The IDC MarketScape model uses these values to calculate each company's score for each of the criteria and rolls these values up to arrive at the described set of two scores. These scores ultimately position each company in the IDC MarketScape chart.

SITUATION OVERVIEW

Introduction

The open systems VTL market was forecast to be a \$630.6 million market in 2008, with the market dominated by three or four suppliers. IDC defines the open systems VTL market as solutions that utilize virtualization software, hard disk drives, and one or more server engines to emulate either tape drives or tape libraries. Although these products are designed to write in tape format, they may or may not support the movement of data to physical tape cartridges or tape libraries. VTL capabilities can be packaged in a number of ways.

An *appliance* packaging is a single integrated solution or appliance (i.e., including VTL gateway/server head, VTL software, and disks). Conversely, a VTL *gateway* includes VTL software running on a server head that makes use of external disk (i.e., everything but the disk). VTL software can also run on a networked storage *controller*. Lastly, VTL software can be packaged as part of a *software suite* that allows for VTL functions and may include other data protection functions. Regardless of packaging, IDC research indicates that over the past several years, the customer drivers for investment in VTL solutions have resulted from the need to:

- ☒ Meet backup windows without compromise
- ☒ Provide faster restore and recovery times
- ☒ Enable rapid integration with existing backup applications
- ☒ Improve performance and utilization of backup resources
- ☒ Support physical tape for archive and regulatory requirements
- ☒ Lower the operational costs associated with physical tape handling

The open systems VTL market continues to evolve within the broader landscape of data protection and operational and disaster recovery objectives for firms of different sizes. Given these imperatives, key characteristics enable firms to better meet customer requirements and be successful in the VTL market.

Factors Impacting Success in the VTL Market

IDC has identified the following criteria as factors that are or will materially contribute to vendor success in the U.S. open systems VTL market:

- ☒ **Established storage brands.** Storage vendors with established routes to market, a large storage installed base to leverage, and products complementary to VTL will be best positioned to capitalize on the market opportunity, particularly due to their ability to offer a VTL solution as part of a larger customer initiative or investment. IDC expects this to remain true in the three- to five-year horizon.
- ☒ **Appliance packaging.** While the VTL capabilities are provided by software functions, the solutions with the most traction in the market to date have been delivered as integrated appliances. IDC expects to see these appliances begin to expand to offer a breadth of data protection functions including, but not limited to,

tape emulation. The need for a VTL pure-play appliance will diminish as the need for a data protection appliance escalates.

- ☒ **Working, proven deduplication.** Deduplication is a must-have for the VTL market. Some solutions have offered it for years, while others have not. Deduplication, including deduplication-aware replication, will continue to be a must-have capability across many different data protection workloads over the next three to five years. Suppliers that can offer deduplication across different data protection workloads will have an advantage over VTL pure-plays.
- ☒ **Solution minimizing trade-off between performance and optimization.** Large capacity backup environments may not tolerate a compromise in backup performance and threaten an available backup window. Smaller backup capacity environments may face less pressure in performance and completion. Solutions that can tune a deduplication process to fit into both environments without forcing a separate solution to be deployed or purchased will be at an advantage.
- ☒ **Support for VTL Fibre Channel interface alternatives.** Backup applications support disk as well as tape as a backup target. With this true, why are VTL functions necessary, and are there alternatives? VTL has served a purpose in introducing disk into an existing backup environment without disruption, supporting direct movement from disk to tape. However, some alternatives to a Fibre Channel VTL interface, including the NetBackup OST API and other application-specific interfaces, are emerging.
- ☒ **Expansion to adjacent data protection markets.** Over the next three to five years, data protection software vendors will continue to innovate on deduplication, support for disk, physical tape, and other integration of other complementary data protection workloads. Existing VTL appliance suppliers will need to continue to develop value-add functions in the appliance itself to warrant a dedicated system and support ingestion of many different types of data protection workloads.
- ☒ **Reduction of the VTL footprint.** In the next three to five years, a general-purpose data protection appliance that can serve as a target for many different types of data protection workloads will be well received by customers. However, suppliers that can leverage their existing storage portfolio and VTL IP to incorporate tape emulation capabilities of a VTL appliance into the storage controller logic, thus reducing the datacenter footprint, management, and cost requirements, will be successful in the long run.
- ☒ **Fit with cloud strategy.** A macrotrend within the IT industry is the use of private and/or public clouds as a means to provide IT functions as a service. These IT functions can include backup and disaster recovery. In the short run, VTL suppliers with a road map and strategy for the insertion of VTL functionality inside a cloud model will be well positioned. However, cloud architectures implemented in a grid architecture with geographic redundancy of data within the grid may obviate or reduce the need for VTL functions and physical tape in the long run.
- ☒ **Ongoing TCO reductions.** Suppliers with an aggressive road map for technology refresh, including higher-capacity drives, deduplication, increased processing power, and a smaller footprint at lower prices, will be well positioned.

These criteria will be imperative for continued displacement and/or augmentation of tape with disk. However, for some enterprise accounts, tape will remain.

- ☒ **Physical tape support.** Despite ongoing price-per-gigabyte declines in disk capacity in concert with storage optimization technologies, the use of physical tape will continue over the next three to five years, although use cases for the use of tape will change. The role of tape will continue to diminish but will exist in large, legacy datacenters with significant investment in tape automation. It will also continue to be used for long-term archive (deep archive), for regulatory purposes, or as a fourth or fifth copy of data, but perhaps on less frequent basis.

Also important are the performance capabilities of a VTL system and how much data can be sent to a VTL without compromising backup windows. Performance is difficult to measure due to variability in nonpublished and published performance numbers and absent third-party performance and scale test results.

Market Strategies

Tables 1 and 2 include market-specific weightings of the criteria used to measure vendor performance and their resulting placement in the IDC MarketScape.

Table 1 lists weightings of characteristics considered key indicators of a successful long-term strategy. Table 2 lists weightings of characteristics considered key indicators of current, short-term capabilities.

TABLE 1

Key Strategy Measures for Success: U.S. Open Systems VTL Market

Strategy Criteria	Criteria for Success	Subcriteria Weighting
Offering strategy	Development of offerings that will be relevant and attractive to customers over the next three to five years	
Road map	How well a VTL vendor's offering, architecture, and functional road map will meet future target customer needs; well-articulated expansion strategies into adjacent markets and how well the offering becomes an integrated platform for other data protection offerings	3.00
Delivery model	How well the VTL vendor is planning for delivery models that will match customers' shifting preferences for adoption/consumption and reduce the VTL footprint	3.00
Cost-management strategy	Extent of a VTL vendor's strategy for maintaining competitive costs and reducing internal costs associated with the offering	2.00
Portfolio strategy	How well the VTL offering is supported and enhanced by a portfolio of complementary offerings made available by the vendor or its ecosystem of partners	2.00
Offering strategy total		10.00

TABLE 1**Key Strategy Measures for Success: U.S. Open Systems VTL Market**

Strategy Criteria	Criteria for Success	Subcriteria Weighting
Go-to-market strategy	Strategies that maximize the connection between offering and customers, including choosing to target customer segments that offer the greatest opportunity over the next three to five years	
Pricing model	How well the VTL vendor's pricing model is aligned with customers' preference for acquisition of broader data protection offerings; VTL becomes a feature within a broader data protection appliance or moves into the storage controller	2.00
Sales/distribution strategy	How well the sales distribution strategy for the VTL solution is aligned with the way customers want to buy the offering (e.g., online, direct, indirect, etc.); movement of the solution down market increases the role of partners	3.00
Marketing strategy	VTL vendor has a robust game plan for all relevant facets of marketing that matches where the revenue is predicted to flow from over the next five years	3.00
Customer service strategy	How well the VTL vendor retains customers and continues to innovate in customer retention and service areas; high quality of service with low-cost execution offering service levels based on SLAs and variable requirements	2.00
Go-to-market strategy total		10.00
Business strategy	Strategies to grow the business are aligned with market trends and future opportunities over the next three to five years	
Growth strategy	Management within the VTL vendor has a strong formula for growth for the company and one that aligns well with the market trends anticipated over the next three to five years	3.00
Innovation/R&D pace and productivity	The company's innovation model for its VTL offering maximizes its potential to generate market value; the innovation must come in both new features and sustainable enhancements to ensure cost containment	3.00
Financial viability	How well the VTL vendor is generating, attracting, and managing capital to create market value; this is more important for emerging companies and/or VTL pure-play vendors	2.00
Employee strategy	The VTL vendor has a strategy for attracting, motivating, and retaining talent to maximize its opportunity for creating market value	2.00
Business strategy total		10.00

Source: IDC, 2009

TABLE 2

Key Capability Measures for Success: U.S. Open Systems VTL Market

Capabilities Criteria	Criteria for Success	Subcriteria Weighting
Offering capabilities	The offering's capabilities aligning well with current market needs and demands today	
Functionality	How well a vendor's current VTL offering and architecture matches the most important customer needs and requirements	3.00
Delivery model	How well a VTL vendor's offering meets with end-user preference for adoption and consumption	1.00
Cost competitiveness	How well a VTL vendor's cost structure is competitive in the market and the level of direct control it has over costs	2.00
Portfolio offerings	How well a VTL vendor is able to leverage adjacent portfolio offerings to drive attach of VTL and offer services pre- and postsale	1.00
Scale, availability, and performance	How well a VTL vendor's current offering matches the scalability, availability, and performance requirements for customers	3.00
Offering capabilities total		10.00
Go-to-market capabilities	Capabilities that maximize the connection between offerings and customers such as delivery, partnerships, pricing, distribution, marketing, sales, and service	
Pricing model	How well a VTL vendor's pricing model aligns with customers' current preferences for procurement	1.00
Sales/distribution	How well the VTL vendor is able to reach customers and distribute its product using direct and indirect selling	4.00
Marketing	How well the VTL vendor's brand (either product or company) is recognized in the storage industry, with sales support and marketing initiatives to ensure success	2.00
Customer service	How well the VTL vendor provides customer service and the control it has over high customer satisfaction	3.00
Go-to-market capabilities total		10.00
Business capabilities	Financial, employee, partner, and R&D management, among other capabilities, in agreement with current market opportunities	
Growth strategy and execution	How well the VTL vendor management team is executing on growth strategy and capitalizing on adjacent markets	4.00
Innovation/R&D pace	How well the VTL vendor's innovation is generating market value and looking to adjacent market opportunities	3.00

TABLE 2**Key Capability Measures for Success: U.S. Open Systems VTL Market**

Capabilities Criteria	Criteria for Success	Subcriteria Weighting
Financial viability	How well the VTL vendor is generating, attracting, and managing capital to create market value	2.00
Employee management	How well the VTL vendor is attracting, motivating, and retaining talent to create market value	1.00
Business capabilities total		10.00

Source: IDC, 2009

FUTURE OUTLOOK**IDC MarketScape U.S. Open Systems VTL Market Vendor Assessment**

The IDC vendor assessment for the U.S. open systems VTL market represents IDC's opinion on which vendors are well positioned today through current capabilities and which are best positioned to gain market share over the next few years. Positioning in the upper right of the grid indicates that vendors are well positioned to gain market share. For the purposes of discussion, IDC divided potential key strategy measures for success into two primary categories: capabilities and strategies.

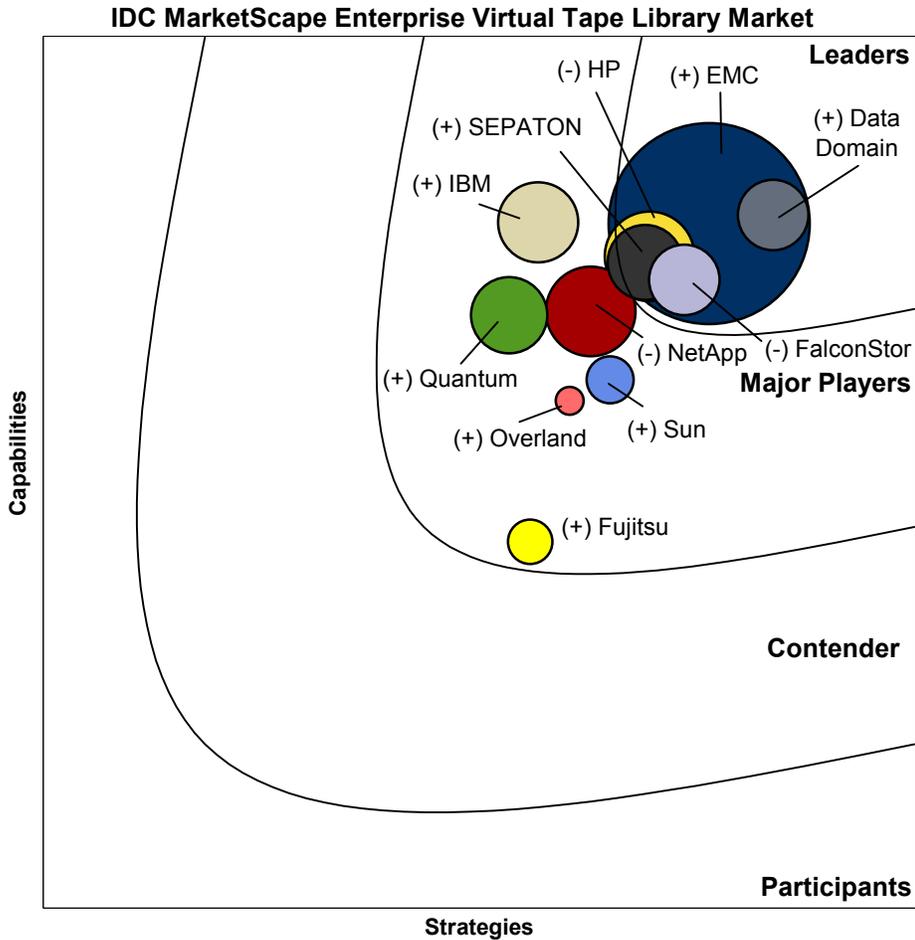
Positioning on the x-axis reflects the vendor's current capabilities and menu of services and how well aligned it is to customer needs. The capabilities category focuses on the capabilities of the company and product today, here, and now. Under this category, IDC analysts will look at how well a vendor is building/delivering capabilities that enable it to execute its chosen strategy in the market.

Positioning on the y-axis or strategy axis indicates how well the vendors' future strategy aligns with what customers will require in three to five years. The strategy category focuses on high-level strategic decisions and underlying assumptions about offerings, customer segments, business, and go-to-market plans for the future, in this case defined as the next three to five years. Under this category, analysts look at whether or not a supplier's strategies in various areas are aligned with customer requirements (and spending) over a defined future time period.

Figure 1 shows each vendor's position in the vendor assessment chart. Its market share is indicated by the size of the bubble and a (+) icon. The (+), (-), or () icon indicates whether the vendor is growing faster, slower, or even with overall market growth.

FIGURE 1

IDC MarketScape U.S. Open Systems VTL Vendor Assessment



Source: IDC, 2009

Vendor Summary Analysis

This section includes an overview of each vendor and qualitatively explains each vendor's position in the vendor assessment graph. The U.S. open systems VTL MarketScape includes the following major vendors (listed in alphabetical order): Data Domain, EMC, FalconStor, Fujitsu Siemens, HP, IBM, NetApp, Overland Storage, Quantum, SEPATON, and Sun. Of note, several of these suppliers carry several products and solutions with VTL functionality. In these cases, IDC selected one of a series of VTL products to compare with alternative suppliers.

EMC/Data Domain

EMC acquired Data Domain in July 2009 and announced that Data Domain will serve as the foundation for its backup and recovery product division. According to the company, the new product division will be led by Frank Slotman, who was the CEO of Data Domain. EMC has committed to investing considerably more resources into the new division, including more people and other products from its backup and recovery portfolio. Data Domain provides deduplication storage systems that dramatically reduce the amount of disk storage needed to retain and protect enterprise data. Combined with Data Domain Replicator software, the use of tape can be minimized or eliminated from both operational recovery and disaster recovery. Joe Tucci, EMC's CEO, said that EMC will accelerate the growth of Data Domain and EMC Avamar, and by the end of 2010, together they will represent more than a billion dollars in revenue.

Since January 2006, Data Domain has offered a VTL software option (licensed separately) that emulates multiple tape libraries over a Fibre Channel interface. Of the 3,500 Data Domain customers and 9,500 systems in production, approximately 10% in 2008 were using a VTL Fibre Channel interface to their storage systems.

Data Domain systems were sold predominantly via reseller partners. IDC believes that EMC will leverage its reseller channel and its direct sales force to extend the reach of Data Domain systems. Data Domain deduplication storage systems are packaged as appliances or as a gateway, and the VTL capability is a separately chargeable software option. Data Domain develops and owns its virtual tape library technology, which is part of the Data Domain operating system and architecture, as is its deduplication technology.

Data Domain offers virtual tape library support, but has expanded well beyond providing only tape emulation. It finds that in smaller enterprise accounts and the midmarket, firms take more of a tape replacement approach and often select a simple, easily deployed NAS interface. For large datacenter environments, firms have more of an investment in tape infrastructure built up, and as a result, a tape augmentation approach is often sought. In these environments, Data Domain more commonly sells its VTL Fibre Channel solution. For Data Domain, deduplication is part of the system architecture and is enabled with either interface. Data Domain continues to be a material provider of deduplication solutions with both its VTL and NFS/CIFS interfaces. Recently, the Data Domain OST interface has actually surpassed demand for VTL. As Symantec NetBackup users move to current versions of software, the OST interface has increased in attach. IDC estimates 15–20% of systems sold in the last several quarters have used the OST interface. EMC Data Domain differentiates itself in the marketplace by:

- Providing multiple types of workloads (backup, archive, replication) and interfaces (NFS, CIFS, VTL, OST) to a single deduplication system
- Offering well-proven inline deduplication technology and customer deployments in volume and a brand synonymous with deduplication
- Supplying a system design that scales performance with processor power as CPU architectures advance

- ☒ Providing advanced deduplication-enabled replication capabilities deployed in volume
- ☒ Capitalizing on EMC's brand, distribution channels, breadth of storage offerings, and access to customers

EMC Corp.

EMC acquired Data Domain in July 2009 and announced that Data Domain will serve as the foundation for its backup and recovery product division. According to the company, the new product division will be led by Frank Sootman, who was the CEO of Data Domain. EMC has committed to investing considerably more resources into the new division, including more people and other products from its backup and recovery portfolio. Data Domain provides deduplication storage systems that dramatically reduce the amount of disk storage needed to retain and protect enterprise data. Combined with Data Domain Replicator software, the use of tape can be minimized or eliminated from both operational recovery and disaster recovery. Joe Tucci, EMC's CEO, said that EMC will accelerate the growth of Data Domain and EMC Avamar, and by the end of 2010, together they will represent more than a billion dollars in revenue.

VTLs are important elements of this accelerated growth. In addition to Data Domain, the EMC VTL offerings include previous-generation Clariion Disk Libraries (CDL) as well as the current EMC Disk Library 4106, 4206, and 4406. EMC has indicated its plans to replace the EMC DL 1500 and 3000 models with Data Domain products. The 4000 series leverages Quantum software capabilities for deduplication. EMC intends to retain FalconStor Software for tape emulation while leveraging Data Domain technology for deduplication. EMC delivers VTL offerings to the midmarket as well as the enterprise. EMC VTLs are distributed through a combination of direct EMC, systems integrator, distributor, and reseller channels. Direct sales are conducted by core EMC sales and overlay backup and recovery sales specialists. IDC expects EMC to leverage its global distribution ecosystem to accelerate Data Domain growth.

The EMC Disk Library 4000 series of virtual tape libraries provides in excess of 1PB of capacity (in the 4406). The 4000 series allows management of all virtual and physical tape assets from one interface. Active engine failover provides for high availability, and disk spin-down and low-power 1TB SATA drives provide for excellent energy and capacity efficiencies.

With thousands of systems shipped from 2004 to 2009, the EMC Disk Library remains one of the most materially installed VTL products on the market today, with multiple thousands of customer deployments. However, only an estimated 10% of EMC's 2,000 VTL customers are making use of deduplication. EMC's dominance in the IDC VTL MarketScape analysis is a result of the significant installed base of customers using its EMC 4000 series Disk Library using FalconStor Software. IDC expects existing EMC Disk Library customers will migrate to VTL solutions with data deduplication capability using Data Domain technology. The company differentiates itself by:

- ☒ Leveraging its Clariion storage systems as the storage systems behind a VTL

- ☒ Leveraging Data Domain, with its well-proven inline deduplication technology and customer deployments in volume and a brand synonymous with deduplication
- ☒ Capitalizing on its brand, distribution channels, breadth of storage offerings, and access to customers
- ☒ Benefiting from a large installed base of VTL customers that are predominantly running the FalconStor-enabled libraries

Effective July 23, 2009 EMC acquired Data Domain for \$33.50 per share in cash.

FalconStor Software

Founded in 2000, FalconStor Software first developed a storage virtualization platform, called IPStor, for open storage virtualization and provisioning. FalconStor's intention was to use this virtualization platform as a base upon which to offer higher-level data protection services. In 2004, FalconStor released its FalconStor Virtual Tape Library (VTL) solution based on the IPStor architecture to streamline tape backup processes, remove tape management complexity, improve reliability, and accelerate backup and recovery. Following the release of the FalconStor VTL offering, in 2005, the company released its continuous data protection solution enabling local and remote recovery, using the same IPStor virtualization platform.

Today, FalconStor Software, a publicly traded company, develops a series of open storage and data protection solutions. FalconStor VTL is packaged in a number of ways to end-user customers as a virtual appliance, as a software suite, and as a storage appliance. FalconStor VTL Virtual Appliance for VMware infrastructure is a preinstalled, preconfigured, and ready-to-run software application packaged with an operating system inside a virtual machine. FalconStor VTL Enterprise Edition is a VTL software suite that can be run on FalconStor prequalified hardware. The VTL Storage Appliance is an integrated server, storage, and software system for SMB customers. FalconStor develops its own VTL software for its offerings, which are sold through a combination of resellers and OEM partners (e.g., Sun, IBM, EMC, COPAN, Spectra Logic). In 2008, FalconStor had two OEMs, EMC and Sun, which accounted for 20% and 13% of its total revenue, respectively. FalconStor has an estimated 3,100 customers using its FalconStor VTL branded product and over 100 customers using deduplication with its VTL products. FalconStor differentiates itself by:

- ☒ Offering a broader data protection platform business, including a strategy and product set that expands well beyond only tape emulation
- ☒ Leveraging a captive, large OEM customer base as a route to market for its solutions both internationally and domestically
- ☒ Intellectual property ownership of VTL technology that fuels a large percentage of the market

Fujitsu

Fujitsu's VTL solution is called ETERNUS CS. ETERNUS CS is the successor to the former CentricStor virtual tape solution from Fujitsu Siemens Computers, which became a 100% subsidiary of Fujitsu Ltd. on April 1, 2009.

Since 1999, CentricStor — now ETERNUS CS — has been tailored to address the data protection needs of large and medium-sized enterprise customers for open system and mainframe attachment. There are over 450 enterprises using ETERNUS CS in production worldwide. The global sales channel of Fujitsu for ETERNUS CS is mainly direct. ETERNUS CS has been successfully deployed in datacenters, primarily in telcos, banks, and insurance companies, and by the public sector for data protection and deep archiving. As an option to the attached tape storage capacity, the ETERNUS CS disk cache capacity can be scaled up to 1.3PB of raw capacity in a single configuration. For the attached host systems, ETERNUS CS supports up to 1,024 virtual tape drives and up to 1,500,000 virtual tape volumes. Furthermore, maximum sustained write throughput can achieve 4,400MBps. The current implementation of ETERNUS CS is focused on high-performance data protection, eliminating bottlenecks of backup and recovery data streams. All protected data can be stored policy-based on disk, tape or both, supporting multiple physical tape creation. Fujitsu's emphasis on high availability is based on automatic system failover as well as synchronous disk mirroring covering two sites. Fujitsu's road map will include some form of data deduplication, with details not available at the time of this analysis. Fujitsu differentiates itself by:

- ☒ Offering a data protection solution that provides both mainframe and open systems VTL support.
- ☒ Providing a scalable grid architecture that supports high-performance and capacity requirements for each class of datacenter environment
- ☒ Placing emphasis on high availability and disaster resiliency for protection against site outages

HP

HP has two VTL offerings, its D2D product family for the SMB and the VLS family for enterprise segments. The D2D family of VTL-enabled systems is based on technology developed and owned by HP and will shortly include a CIFS/NFS interface. The VLS family of systems is based upon an OEM relationship with SEPATON. Both product families are packaged and offered as appliances to their respective target markets. The product set under evaluation for this MarketScape is based on HP's relationship with SEPATON. Between the two product lines, HP has over 2,500 customers using its HP-branded VTL offerings. Deduplication is a capability in both families and has been shipping since June 2008. IDC estimates that approximately 30% of its 2,500 VTL customers are using deduplication.

The D2D product family starts at 1.5TB usable and scales to 18TB usable, supporting iSCSI and Fibre Channel interfaces. The D2D systems are all HP owned and developed technology. The D2D line includes an inline, hash-based deduplication that is optimized for SMB and midmarket price-sensitive environments. Developed out of HP Labs, according to HP, the D2D deduplication technology uses less memory for indexing and lowers the product cost. It also minimizes disk I/O requirements, and HP indicates its deduplication technology enables a match from memory approximately 98–99% of the time rather than going to disk.

The VLS family for entry-level enterprise to datacenter-class deployments has been optimized for scale, performance, and addressing growing backup windows. The VLS 12000 is a gateway product that supports connectivity with HP EVA storage. The VLS 9000 includes the disk storage capacity. The VLS family uses a postprocess, object-level-differencing deduplication approach. The object-level differencing leverages a real-time byte-level comparison process where the byte size can vary in length. With a VLS multinode system, the VLS can scale to thousands of megabytes per second and keep up with capacity requirements, scaling to 1.2PB (with 2:1 compression). According to HP, the current sweet spot in VLS deployments is a 60–70TB system with bandwidth requirements of 600MBps per node or 1,200MBps in a dual-node configuration. The VLS 9000 does support up to 8 nodes, resulting in 4,800MBps. HP negatives include being late to market with dedupe and smaller deduplication customer deployments and a lack of integration between its low-end and high-end offerings. HP differentiates itself by offering a range of VTL solutions serving low-end to datacenter environments, simplicity of use for SMB customers, and multinode scalability for enterprises.

IBM

IBM has participated in the mainframe VTL since 1997. IBM's mainframe attached Virtual Tape Server (VTS) is currently in its fourth generation. However, IBM has been providing open system VTLs since 2005 with its products based on FalconStor code. The IBM product set under evaluation for this MarketScape is based on technology acquired from Diligent Technologies. IBM made a strategic acquisition of Diligent Technologies in April 2008. The IBM/Diligent products are based on patented technologies. The ProtecTIER VTL platform features HyperFactor inline deduplication and provides scalable disk-based data protection solutions. The ProtecTIER platform uses inline data deduplication technology, enabling customers to protect their data without disrupting the existing policies, practices, and procedures in their datacenter. Currently, the number of IBM customers using ProtecTIER systems is in the hundreds across all industries in enterprise and midrange environments.

Currently, customers have the ability to replicate their ProtecTIER repository to remote locations using IBM's recently announced native IP replication. This provides tape cartridge-level replication from the primary site to a remote location. It is expected that IBM will provide many-to-one and eventually many-to-many replication in the near future. The ProtecTIER system does support direct attach physical tape support; however, both the VTL target and physical tape system maintain separate paths to the backup server. The backup application is used to move or copy data from the VTL to physical tape, which maintains catalog consistency. Security and encryption is supported by IBM's physical tape drives once the data is written to the cartridges.

A single ProtecTIER server can provide sustained throughput of up to 500MBps or more with inline deduplication. A dual-node ProtecTIER cluster can provide sustained inline deduplication performance up to 1,000MBps or more. IBM's emphasis for its ProtecTIER data deduplication solution is on cluster capabilities that deliver high availability, increased performance, and maximize the impact of deduplication. A ProtecTIER cluster provides a single deduplication repository with up to 256 virtual drives, 500,000 virtual tape cartridges, and 162 virtual libraries. It is expected that

IBM will provide 4-node, 8-node and, eventually, 16-node cluster configurations in the future.

IBM is differentiating its ProtecTIER systems by :

- ☒ Providing customers with a complete suite of disk, tape, software, and deduplication technologies to build an end-to-end data-protection framework
- ☒ Offering additional tools that would be deployed adjacent to a ProtecTIER system, including IBM Tivoli Storage Manager, IBM Tape Systems, IBM Fibre Channel Disk Systems, and IBM SATA Disk Systems as well as copy services such as IBM FlashCopy, Metro Mirror, and Global Mirror

NetApp

With 100% NetApp technology ownership, the VTL family provides a powerful combination of performance, scalability, simplicity, and cost savings. NetApp's VTL solutions are purpose built and tuned specifically for backup and recovery, providing the throughput needed to meet tightening windows. Patented NetApp Direct Tape Creation (DTC) technology reduces tape costs up to 50%, while driving continued value from existing tape assets with automated, high-speed tape creation. An embedded tape path provides up to 5.6TB/hour, so physical tapes can be created more efficiently and quickly and be easily sent offsite for disaster recovery. NetApp VTL Shadow Tape functionality offers a unique advantage for those customers using the embedded tape path. Shadow Tape copies enable the supported backup application to recall tapes that have been "ejected" from the system without operator intervention impacting the recovery time. The NetApp VTL systems have a single operating system that runs the tape emulation feature and deduplication and provides integrated disk management. IDC estimates NetApp currently has 800 VTL customers, with approximately 20% using its deduplication. NetApp differentiates its VTLs on:

- ☒ Performance and feature sets tuned for large enterprise datacenter accounts with significant investments in tape infrastructure and large-capacity data sets
- ☒ Support for physical tape features, including direct tape creation, tape smart sizing, and Shadow Tape copies for direct restore from disk
- ☒ Native, block-level deduplication in the core VTL operating system, which uses a rolling hash algorithm in combination with sequential byte-level comparison for maximum data integrity
- ☒ Single operating environment (The disk controller in the VTL appliance runs both the VTL tape emulation software and the integrated deduplication functionality. Also running on the NetApp disk controller are disk management capabilities for easy RAID group creation that can be done by a person unfamiliar with Fibre Channel storage.)

Overland Storage

Overland Storage offers a broad range of REO disk-based systems that are scalable to address the needs of small customers and small enterprises. Overland offers its

REO Protection OS, which provides a core set of volume and device virtualization, management, and connectivity capabilities.

The high-end REO 9100 VTL can be user defined as a tape library with up to 32 virtual tape drives, 512 virtual cartridges, 6 partitions, and/or any combination of up to 64 standalone virtual tape drives and/or disk volumes (LUNs). In addition, the 9100 can have a maximum of 4 FC ports and a usable capacity of up to 62TB, with optional expansion modules. However, a base REO 9100 system has a usable capacity of 20TBs, with 1TB drives, and is RAID5 protected. The REO 9100 offers sustained sequential write performance at 4TB/hour (assuming 2:1 compression) for a single-node system.

Overland's REO 4500 and REO 1550 are smaller VTL appliances that support usable capacities of 1–10TB, depending on the capacity of the drives and RAID set. However, the REO 4500 can be scaled up to 30TB with expansion modules. In addition, the REO 4500 has 2GbE (iSCSI) and two optional 4GB FC ports, whereas the REO 550 has one, 1GbE (iSCSI) port. Overland differentiates its REO VTLs by:

- ☒ Offering an extensive lineup of data protection solutions spanning disk and tape, including its ULTAMUS RAID FC storage arrays for online and bulk near-line storage requirements and Snap Server networked storage appliances for distributed offices and organization file sharing.
- ☒ Introducing its REO Business Continuity Appliance, which provides continuous data replication, enabling automated recovery of data and applications at both the local datacenter and remote disaster recovery levels

Quantum

Quantum offers an extensive line of disk-based backup appliances that can present as a VTL (Fibre Channel) or NAS (CIFS and NFS) called the DXi-Series. In addition, Quantum's DXi-Series supports policy-based data deduplication, remote replication, and direct tape creation. Quantum's DXi7500 Enterprise system is scalable from 18TB up to 220TB of usable capacity. In VTL mode, the DXi7500 can support up to 160 virtual tape drives and up to 30,000 virtual tape cartridges and can emulate nearly all of Quantum's tape libraries as well as support prevailing tape drive formats. In addition, the DXi7500 Enterprise system is replication compatible with Quantum's smaller DXi systems and capable of a 4TB/hour ingest rate, assuming a 2:1 compression, in native VTL mode with data deduplication enabled.

Quantum's DXi7500 Express model is deployed as a preconfigured solution with a usable capacity up to 11TB. The performance metrics are similar to those of the DXi Enterprise system. Both DXi7500 systems are designed to be interoperable with Quantum's DXi3500 and DXi2500-D disk systems. These systems have usable capacities of 1.2–4.2TB and are designed for distributed backup, data protection, and disaster recovery. The DXi2500-D systems include preconfigured data deduplication, NAS presentation, and replication. The DXi3500 can present as NAS or VTL. The DXi3500 and DXi2500-D are designed to be integrated with DXi7500 systems, using remote replication, to provide an edge-to-core data-protection strategy. Most recently, Quantum has enhanced its DXi-Series with a software update — DXi 1.2. The new software release has expanded support for direct tape creation as well as path-to-

tape support for Oracle Secure Backup and Atempo Time Navigator. In addition, the DXi7500 is now qualified to support Shadow Tape Creation with Symantec BackupExec, CA ArcServ, and BakBone NetVault. This is in addition to direct physical tape creation through Symantec's OpenStorage initiative. Quantum's key differentiators for its DXi-Series are its ability to offer customers:

- ☒ Integrated disk, tape, and management software solutions
- ☒ Storage systems that are optimized for network attached storage (NAS) and Symantec OpenStorage (OST) API architectures
- ☒ Backup software enabling backup, restoration, and protection for VMware virtual infrastructures and data using Quantum's esXpress software

SEPATON Inc.

SEPATON started shipping its VTL product line in 2003. Focused on enterprise-level scalability and performance, SEPATON's S2100 VTL appliances include VTL, data deduplication, and replication functions. SEPATON's VTL system is expandable up to 16 nodes and 2PB of capacity. SEPATON offers two different packages for its VTL solutions: the S2100-ES2 for large, enterprise datacenters and the S2100-DS2 for departmental or smaller environments. The ES2 published performance numbers quote speeds up to 9,600MBps for a 16-node configuration. Capacity starts at 10TB usable and scales to 1.6PB usable storage with built-in RAID 6. The SEPATON DS2 systems come with either 7TB or 14TB of usable capacity and offer 300MBps performance. The SEPATON DS2 can be deployed not only in smaller firms but also in smaller, regional datacenters (or remote branch) that serve as a satellite to a larger datacenter.

SEPATON is a software company that packages its solution as an appliance for ease of install. The first version of SEPATON's deduplication DeltaStor software came at the end of 2006 and was focused on NetBackup and single-node deployments. Since then, SEPATON has released DeltaStor Enterprise Edition 5.0 (GA May 2009), which is currently qualified with 6 nodes and is expanding to 8 nodes. SEPATON's DeltaStor deduplication uses a postprocess, content-aware approach. DeltaStor makes use of byte-level comparisons with delta differencing to get high deduplication ratios. *Content aware* refers to the system's ability to recognize particular backup application file formats and tune its deduplication algorithm for those formats to gain the greatest storage reduction. SEPATON's patented *forward referencing* uses the last data in as the reference data, and older data is made up of pointers to the newest backup. This allows for superior restore speed. SEPATON's deduplication starts once an individual backup job completes, and a nightly backup is made up of many jobs. It does not need to wait for all backup jobs to finish before deduplication processes start, thus making the term *postprocess* not entirely accurate. SEPATON calls this approach *concurrent processing*. The same technology enables its DeltaRemote technology, which replicates deduplicated data to a remote site.

SEPATON signed a major OEM deal with HP in 2005 for its VLS6105 and VLS6510 systems and renewed its agreement in 2008 to reflect new HP products that bring SEPATON's DeltaStor deduplication software to market as accelerated deduplication. In the VTL market. SEPATON is differentiated by:

- ☒ The DeltaScale architecture is a grid technology that offers performance and scalability, allowing for growth in capacity and/or performance independently. The grid architecture can support up to 16 nodes, with all nodes in the grid sharing responsibility for backup, restore, and deduplication processes. Each node is capable of up to 2TB/hour backup and recovery performance and can deduplicate 25TB of data in 24 hours.
- ☒ SEPATON also considers its content-aware, forward-referencing deduplication approach unique. DeltaStor restores data instantly. As a new backup is performed, the system replaces the previous backup as the baseline. Duplicates found in older stored data are replaced with pointers forward to the most recent backup. By keeping an unduplicated copy of the most recently backed up data, SEPATON claims DeltaStor software can restore data faster. These same benefits can be extended to a remote site with SEPATON's DeltaRemote replication software.

Sun

The Sun StorageTek Virtual Tape Library (VTL) products are based on its Solaris OS and utilize FalconStor's software for its VTL and deduplication capabilities.

The StorageTek VTL Plus is scalable from 8TB to 6,400TB of usable capacity, depending on model and configuration. In addition, VTL Plus configurations offer 800–1,200MBps per VTL engine throughput. Throughput and performance scales with size. The VTL Plus supports 16–256 virtual libraries, 128–2,048 virtual drives, and up to 131,000 virtual tape cartridges. The VTL Plus architecture supports IP replication with compression and encryption options, tape caching, and NDMP for policy management of physical tape libraries; the tape shredding compliance feature actively deletes data.

The StorageTek VTL Prime provides VTL and data deduplication capabilities in a single appliance. VTL Prime uses a postprocess data deduplication, delivering 700MBps performance and 65TB of useable deduplication repository. VTL Prime supports iSCSI and FC host interface. Sun's differentiators in the open market are:

- ☒ Support for Sun StorageTek ACSLS Manager software
- ☒ Strong heritage as a data protection vendor with its emphasis on integration of tape solutions for archive
- ☒ Breadth of server, storage, and IT infrastructure products as well as global support

ESSENTIAL GUIDANCE

Despite long-standing tape infrastructures, IDC observes that the VTL market is quickly shifting into a broader disk-based data protection market that includes not only tape emulation but also support for deduplication, remote replication, standard network interfaces, NAS, and the ability to become a data protection platform for different backup, archive, and recovery workloads. Therefore, it is essential for

suppliers to capitalize on the pure-play VTL market opportunity without losing sight of the broader disk-based data protection market and customer needs. As a result, IDC has the following observations and recommendations for suppliers:

- ☒ Large storage brands with a captive installed base, established routes to market, and a breadth of storage portfolio offerings are well positioned to capitalize on VTL deals that are part of broader datacenter initiatives such as disaster recovery, virtualization, and consolidation. IDC expects this trend to continue as the breadth of functionality needed by users expands well beyond tape emulation.
- ☒ In today's market, appliance packaging prevails. However, power, cooling, and space constraints continue to drive a need for reduced footprint, which will drive consideration of alternative packaging options. These include software appliances, features running on a storage controller, and/or running of multiple data protection services on a common platform. The need for a VTL pure-play appliance will continue to diminish over time.
- ☒ The disk-based data protection market will continue to consider capacity/backup windows and retention times as factors determining a deduplication approach. Firms with larger capacity backups and shorter retention times may gravitate to a postprocess deduplication. However, for firms with smaller capacity backups and longer retention times, an inline approach to deduplication will continue to prevail. Market factors illustrate that computing speeds and processor power continues to make inline deduplication more and more viable for larger-capacity backups, without compromising backup windows and performance.
- ☒ The need for a VTL Fibre Channel interface to deliver both performance and export to physical tape will diminish as other interfaces (NetBackup OST, etc.) that run over Fibre Channel continue to gain traction. These interfaces will offer the ability for backups to reside on local disk, remote disk, or tape while doing so in a backup application catalog-consistent fashion.
- ☒ The market will continue to move to satisfy a broader data protection appliance model, with support for different and multiple types of protection workloads and interfaces. A common software element will be provided to manage these appliances and interface with different applications and policy-driven data management features. As a result, VTLs must move from pure-play tape emulation to broader data features and services. Precedents include a shift from NAS to unified storage or from disk-based archive to archive services included within file-based storage platforms.
- ☒ IDC research shows there is a moderate opportunity for VTL in the midmarket, particularly for those using tape as an archive, thus requiring long-term retention. However, these firms typically do not have the resources, budget, or expertise to deploy a Fibre Channel SAN. Thus, suppliers with an iSCSI and/or NAS interface will be well positioned to capture this market opportunity.
- ☒ As VTL pure-play systems become part of large disk-based data protection strategies, tape emulation will either be embedded in the storage controller or be

part of a broader data protection system. In either scenario, tape emulation will increasingly tie not only to storage but also to solution sales. Suppliers must focus on developing a broader set of portfolio offerings and solution-selling skills for the long-term success of the technology.

- ☒ The IT industry is moving to a cloud-based model for the delivery of IT infrastructure services. VTL suppliers must consider their strategy and role in broader private or public cloud architectures and what the role of tape and tape emulation will be in these solutions. These architectures may very well warrant a tapeless approach, which eliminates the need for tape offsiteing, in which case the role of tape emulation will wane over time.

- ☒ In the long run, the life span of VTL as a technology is a function of the need to export to physical tape (for economic or regulatory purposes), the economics of disk versus tape (particularly for high volumes of data), the longevity of installed tape automation infrastructure, the retention times, and the standardization of other interfaces to get data to physical tape. Other factors will include power, cooling, and space requirements for disk versus tape.

LEARN MORE

Related Research

- ☒ *Worldwide Data Protection and Recovery Software 2009–2013 Forecast: The Year of the Cloud* (IDC #219111, July 2009)

- ☒ *Teck, a Canadian Mining Company, Implements Data Deduplication Technology as Part of a Global Standardization and Compliance Initiative* (IDC #216452, February 2009)

- ☒ *Data Protection in United States SMBs: Opportunity and Analysis* (IDC #216400, January 2009)

- ☒ *Worldwide Virtual Tape Library 2008–2012 Forecast and Analysis: Impact of Data Deduplication* (IDC #215644, December 2008)

- ☒ *Evaluating Tape's Evolving Role in the Data-Protection Market* (IDC #213051, June 2008)

Synopsis

This IDC study uses the IDC MarketScape model to provide an assessment of a number of vendors participating in the VTL market. The IDC MarketScape is an evaluation based on a comprehensive framework and a set of parameters that assess vendors relative to one another and to those factors expected to be most conducive to success in a given market during the short and long term.

"The VTL market is quickly shifting into a broader disk-based data protection market that includes not only tape emulation and the ability to export to physical tape but also

support for deduplication, remote replication, standard network interfaces, and the ability to become a data protection platform for different backup, archive, and recovery workloads," said Laura DuBois, program director for Storage Software at IDC.

Copyright Notice

This IDC research document was published as part of an IDC continuous intelligence service, providing written research, analyst interactions, telebriefings, and conferences. Visit www.idc.com to learn more about IDC subscription and consulting services. To view a list of IDC offices worldwide, visit www.idc.com/offices. Please contact the IDC Hotline at 800.343.4952, ext. 7988 (or +1.508.988.7988) or sales@idc.com for information on applying the price of this document toward the purchase of an IDC service or for information on additional copies or Web rights.

Copyright 2009 IDC. Reproduction is forbidden unless authorized. All rights reserved.