



I D C T E C H N O L O G Y S P O T L I G H T

Continuous Operations for Applications with Multiple Storage Array Access

December 2012

Adapted from *Worldwide Enterprise Storage Systems 2012–2016 Forecast Update* by Natalya Yezhkova, John Rydning, and Jeff Janukowicz, IDC #237886

Sponsored by EMC

IDC research forecasts that external storage capacity will grow by a factor of nearly six from 2011 to 2016, fueled by the expanding global economy, increased usage of mobile and social technologies, and big data. However, spending on services that manage the additional capacity is forecast to increase only by one-third. The exponential capacity growth versus incremental management spending has created a need for IT organizations to do more less. This trend, coupled with continued external cost pressures, has led IT organizations to investigate ways to leverage their current IT investments by increasing datacenter efficiencies, automating operations, and reducing complexity while achieving higher levels of service through continuous operations. Inhibitors to continuous operations have been current application and infrastructure capabilities, as well as the downtime and complexity associated with data migrations. Organizations have developed a few enterprise applications with workarounds in an effort to achieve continuous operations through implementing active-active (including dual storage) configurations, which can be costly and difficult to maintain. Organizations are seeking cost-efficient solutions to achieve continuous operations while reducing overall operational and management complexity.

This Technology Spotlight explores the trends and challenges related to continuous operations. The paper also discusses the role that EMC, with its VPLEX and VNX solutions, plays in this emerging and important market.

Introduction

Businesses are becoming more reliant on information to support revenue-driving applications, business decisions derived from data (i.e., analytics and big data), and mobile access to applications. Among the most significant issues for IT organizations are data availability and downtime, whether planned or unplanned. In many cases, business must drive toward increased availability, whether for regulatory or business reasons. Other issues many organizations struggle with include application complexity, datacenter usage, and data migrations.

Because the loss of processing capability and data unavailability can have disastrous consequences for both the organization and the IT manager, many organizations are now looking to achieving continuous operations. Recent weather issues such as hurricanes, which have severely disrupted operations in certain locations, have highlighted the need for geographically dispersed operations. In addition, not all downtime is unplanned, and a significant contributor to planned downtime is data migration. IT organizations are seeking solutions that can satisfy their availability needs for both planned and unplanned downtime. No longer can organizations truly achieve continuous operations in one datacenter. Therefore, many organizations are adopting a two-datacenter (or more) strategy. A majority of mature enterprise fault-tolerant applications were designed to have a single storage array as an active primary copy and a passive secondary copy for failover (within the datacenter or outside the datacenter). An example of this is VMware SRM. Therefore, two-datacenter strategies were deployed as active-passive to accommodate the application design. In an active-passive

configuration, the IT assets in the passive configuration were idle until the active assets transitioned the workload to the passive assets through a planned hand-off or unexpected failure. In the active-passive datacenter scenario, 50% of the IT assets were idle, essentially doubling the IT budget to achieve two-datacenter operations.

In an effort to increase datacenter utilization beyond 50% and continuous operations concurrently, IT organizations and vendors have developed interesting, but complex workarounds. For example, in some cases, IT organizations have deployed a complex host-based mirroring solution to achieve active-active between two datacenters. While this type of solution works and has vendor support, it is not the optimal solution because it is an add-on designed for single datacenter operations.

IT organizations' requirements for two-datacenter solutions include the following:

- Mission-critical availability
- Integration with existing applications and storage
- Ability to reduce or eliminate operational complexity
- Active-active processing between datacenters
- Increased utilization beyond 50% of IT assets
- Increased datacenter efficiency

Benefits of Highly Available Storage

High-availability clustered storage solutions are now being implemented to support an active-active two-datacenter strategy. This approach can overcome many of the limitations of host-based mirroring. Among the benefits of high-availability clustered storage are:

- Cost-effective storage capacity
- Reduction of downtime
- Improved performance
- Increased utilization rates
- Nondisruptive migrations and technology refreshes

Market Trends

Quite a few market forces are effecting change within the storage market. This section highlights the most common for IT organizations with regard to continuous operations requirements:

- **Minimized downtime with continuous operations.** Businesses have become more reliant on their data, and the trend to a globalized economy has extended businesses' data access requirements to 24 hours a day, 7 days a week, and 365 days a year. IDC research shows that the cost of downtime per hour varies per industry; however, examples in the following verticals show substantial financial (\$9.9 million), retail (\$397,000), and healthcare (\$157,000) impact. The cost impact of downtime has reduced or even eliminated the planned outage window for maintenance. Furthermore, the business expectation is 100% uptime, with many organizations demanding zero RTO and zero RPO. To deliver on expected business availability, applications and supporting infrastructure must be operational beyond a single failure, whether it is a single component, system, or datacenter.

- **Reduced application complexity.** Data migrations and active-active application demands result in an ever-growing level of complexity for IT organizations. Active-active application configurations can be very complex to implement (if even possible) and maintain. The scheduling, planning, and execution of data migrations require substantial amounts of additional labor, which drives up operational costs. Capacity and performance per storage array have grown substantially, which has led to the consolidation of an increasing number of applications per storage array. Data migrations for storage array refreshes have become increasingly painful due to the number of applications per storage array that need to be coordinated, scheduled, and moved. IT operations are seeking solutions to reduce operational complexity while delivering higher levels of availability.
- **Higher datacenter utilization with increased availability.** IT organizations want to have access to and process data simultaneously from a minimum of two datacenter locations. The primary drivers are to increase IT asset utilization beyond 50% (from active-passive configurations) and to reduce the number of idle assets. In addition, IT organizations aim to increase the performance of existing investments (or avoid the costs of purchasing new infrastructure) for continuous availability initiatives.
- **Automation.** Increased time-to-market pressures and the need to reduce the amount of labor required to deliver IT services have been the primary drivers for automation within IT organizations.

Considering EMC

Headquartered in Hopkinton, Massachusetts, EMC Corporation is a global developer and provider of information infrastructure technology and solutions. Working with EMC, IT organizations are able to leverage the company's more than 20 years of market experience in storing, protecting, and replicating mission-critical applications. EMC's VNX with VPLEX offers IT organizations continuous operations with distributed storage within a datacenter or for mission-critical and performance-demanding clustered applications. Use cases for VNX with VPLEX include the following:

- Minimize downtime with continuous operations
- Reduce application complexity
- Achieve higher datacenter utilization with increased availability
- Improve operational efficiency through automation

The EMC VNX Series is designed to provide advanced capabilities along with maximum performance while maintaining low cost per IOP and per capacity. VNX is a hybrid array offering optimal balance of performance, capacity, and cost. This balance is achieved through a blended pooled storage strategy where a small number of SSDs are utilized to increase the overall system performance and a larger number of HDDs are utilized to drive up capacity density while reducing the overall cost.

The VNX family has the following features/capabilities:

- Optimized for virtual environments
- Designed to take advantage of pooled storage (Flash, SAS, NL-SAS)
- Automated tiering of data with FAST Cache and FAST VP
- Easy-to-use centralized management with Unisphere

VPLEX is a continuous operations–enabling enterprise appliance that provides transparent data movement and high-availability functionality to applications over distance. Deployed between the compute and storage layers, VPLEX has the following features/capabilities:

- Distributed active-active data access from multiple datacenters
- Transparent data migrations between storage array within and between datacenters
- Scale-out capability by aggregating storage arrays into a single storage pool, enabling hosts with simultaneous read and write access to the same data, all while served from two or more storage arrays that are separated by distance

VNX is designed to provide capacity and performance at a midmarket price, while VPLEX provides sophisticated distributed storage access. Together, VNX and VPLEX can deliver active-active data access at midmarket storage cost. Typical use cases for the VNX with VPLEX solution include the following:

Continuous Availability Through Distributed Storage

- Designed for distance between datacenters
- Transparent failover between storage arrays (local and remote)
- Provide active-active processing and datacenter functionality to applications that were designed for single datacenter availability
- Remove single array requirement for VMware availability function
- Increased availability allowing for zero RTO and zero RPO over distance
- Full utilization of resources at both sites

Transparent Data Migrations

- Nondisruptive data movement between storage arrays
- Ease of scheduling data migration for applications (without downtime requirement)
- Heterogeneous storage array support
- Reduction in opex and professional services
- No additional host or storage software or licensing

VMware Availability Over Distance

- Increased availability through ability to implement VMware High Availability and VMware Fault Tolerance
- Balance resources between datacenters with VMware DRS and vMotion within metro distance
- No requirement for VMware SRM to provide availability
- Active-active datacenters providing continuous availability over distance
- Automatic, hands-free failover
- Dynamic cross-site application relocation for load balancing or maintenance
- Elimination of CPU cycles spent on data movement
- Certified by VMware

Stretched Oracle Real Application Clusters (RAC) 11g Across Datacenters

Oracle RAC is a continuous availability solution for Oracle databases. The typical Oracle RAC solution is deployed in a two-datacenter strategy as active-passive. In the VNX with VPLEX solution, Oracle RAC can be deployed in a two-datacenter solution as active-active with the following features/capabilities:

- Active-active datacenters providing continuous availability over distance
- Metro distance between datacenters (up to 5ms round trip time [RTT])
- Increased utilization through both datacenters running active-active
- No need for host-based mirroring or wasted host CPU cycles
- Simple management — no complex operational workarounds
- No customization — standard Oracle RAC implementation
- Certified by Oracle

Nondisruptive Migrations and Technology Refreshes

- Accelerated tech refreshes
- Removes the application downtime associated with storage array refreshes
- Enables storage optimization by enabling restructuring of data across storage arrays
- Enables datacenter migrations to be done as online events, faster and nondisruptively

Challenges

Many of EMC's customers are combining VNX and VPLEX to increase utilization rates and eliminate downtime for technology refreshes. In a crowded storage market, the company has to educate IT managers regarding the capabilities of its products and must continue to highlight its ability to deliver continuous operations due to the ability to move data from datacenter to datacenter irrespective of the underlying storage.

Conclusion

Business requirements have driven the demand for continuous data access while increasing pressure to reduce overall IT costs. To meet business demands, IT organizations are searching for solutions to increase datacenter utilization of IT assets beyond 50% and simplify operations while maintaining high levels of availability and data mobility. The majority of enterprise applications and infrastructure solutions were designed to be highly available in a single site, in single instance storage, or in active-passive configurations across two datacenters.

EMC offers an active-active infrastructure solution in the VNX with VPLEX solution, which was designed for distance, multiple datacenters, and data portability. IDC believes the market for continuous operations and data access for enterprise applications will continue to grow, and if EMC can address the challenges outlined in this paper, the company has a significant opportunity for success.

A B O U T T H I S P U B L I C A T I O N

This publication was produced by IDC Go-to-Market Services. The opinion, analysis, and research results presented herein are drawn from more detailed research and analysis independently conducted and published by IDC, unless specific vendor sponsorship is noted. IDC Go-to-Market Services makes IDC content available in a wide range of formats for distribution by various companies. A license to distribute IDC content does not imply endorsement of or opinion about the licensee.

C O P Y R I G H T A N D R E S T R I C T I O N S

Any IDC information or reference to IDC that is to be used in advertising, press releases, or promotional materials requires prior written approval from IDC. For permission requests, contact the GMS information line at 508-988-7610 or gms@idc.com. Translation and/or localization of this document requires an additional license from IDC.

For more information on IDC, visit www.idc.com. For more information on IDC GMS, visit www.idc.com/gms.

Global Headquarters: 5 Speen Street Framingham, MA 01701 USA P.508.872.8200 F.508.935.4015 www.idc.com