DELL Technologies

Solution Brief

Dell EMC PowerStore: Persistent Data Availability

Deploying to achieve business-critical data availability

Data availability

Today's business-critical application environments demand solutions which offer availability that is expressed in more than a simple number-of-nines of availability. While Dell EMC™ PowerStore™ is designed for 99.9999% availability[1], there is more to availability than only the probability that a system will be operational during a defined period. There are also required features and automated processes that need to be considered which enable non-disruptive operations so that the environment is persistent. These considerations in full provide the best opportunity for high availability and zero down time.

All Dell EMC storage solutions are architected for reliability, availability, and serviceability (RAS) as part of their design. Dell EMC storage products also include redundant hardware components and intelligent software architected to deliver extreme performance and availability—simultaneously. This combination results in unparalleled operational durability, while also leveraging components in new ways that decrease the total cost of ownership of each system.

All Dell EMC storage solutions can incorporate onboard and interoperable data protection functionality such as remote replication and snapshots of data, used to deliver nonstop business continuity. Solutions can also include intelligent software which automates the processes which are needed to seamlessly operate an always-on environment.

By combining storage solutions architected for RAS, comprehensive remote replication technologies, local snapshot technologies, and intelligent software which helps automate data center operations, Dell Technologies™ can provide customers with solutions which go beyond just a bunch of nines to an always-on persistent data availability environment.

Persistent solutions require that there be no single point of failure in the entire environment. This requirement spans the entire application infrastructure—from the hosts to the switches, to the storage, and to the data center itself. This means there must be complete redundancy end to end, and automation to detect and handle various fault conditions. For example, all the application components in a data center could be designed for the highest levels of availability. However, if a power-loss event occurs within the building or within the local region, the application could become unavailable unless redundancy is included at the data-center level.

High availability hardware resiliency

PowerStore appliances support clustering which allows for multi-appliance configurations, each appliance with its own independent fault tolerance. In addition, a single PowerStore appliance is a fully unified block and file environment packaged into a single 2U enclosure. Each PowerStore within the cluster consists of two nodes that make up a high availability (HA) pair. PowerStore appliances feature both fully redundant hardware and highly available software features that keep the system online in the event of component failures, environmental failures, and even failures of multiple components at the same time, such as an internal fan and a disk drive. Hardware can be replaced utilizing the redundant node architecture, keeping the system and data online and available.

PowerStore has a dual node architecture where each node is identical, giving the ability to serve I/O in an active/active manner. Active/active capabilities increase hardware efficiency since there is no requirement for idle-standby hardware. In this way, PowerStore efficiently makes full use of all available resources through a highly redundant architecture. If SAN or IP network port connectivity is lost, the system leverages redundant port connections, allowing hosts and clients to maintain access to the data. At the same time, dual-ported hard drives ensure that each node has seamless connectivity to the data.

PowerStore protects data writes using redundant shared write cache that is leveraged by both nodes simultaneously. For PowerStore 1000 to 9000 appliances, any incoming I/O that is determined to be a write is stored in DRAM memory and then copied to the NVMe NVRAM drives so that each node in the appliance can access the data. In PowerStore 500 appliances, since NVRAM drives are not supported, write I/Os are instead mirrored to DRAM memory in each

[1] Based on the Dell Technologies specification for Dell EMC PowerStore, April 2020. Actual system availability may vary.

© 2021 Dell Inc. or its subsidiaries. All Rights Reserved. Dell Technologies, Dell, EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be trademarks of their respective owners.

node. This enables the write cache to be preserved through hardware and software faults, as well as node reboots. Also, in the event of a power outage or temperature alarm on both nodes, the integrated battery backup unit (BBU) provides temporary power to the system, allowing the cache to be destaged.

High availability software features

Dynamic Resiliency Engine (DRE) ensures enterprise-class availability by achieving faster drive rebuild times with distributed sparing. DRE rebuilds smaller chunks of the drive simultaneously to multiple drives in the appliance. Intelligent allocation of unused user space is designed to replenish spare space for handling multiple drive failures automatically. Rebuild speeds are calculated intelligently when there is incoming I/O to ensure performance and availability during rebuild. DRE allows for flexible configurations that can lower TCO by expanding storage with a single drive at a time or adding different drive sizes based on storage needs. In PowerStoreOS 2.0, the added benefit of double drive failure tolerance further increases the number of concurrent failed drives the system can withstand.

Active/active architecture allows for configuring paths across nodes for both iSCSI and FC configurations. For multipathing, leverage software such as Dell EMC PowerPath™ on your hosts to utilize the full multipathing capabilities of PowerStore. Active/active architecture ensures host IO requests use the active optimized path for best performance. In the rare event that the host loses the active optimized path, the host will switch over to the active non-optimized path where the I/O request will be processed locally by the PowerStore node.

Snapshots are supported on block and file resources (volumes, volume groups, virtual machines, file systems, and thin clones). Snapshots ensure the availability of critical data with near-immediate reverts to a previous, known-good recovery point, in the event of an accidental deletion or malware outbreak. Snapshot rules can be configured to support a wide range of recovery point objectives.

Native replication is supported on block resources (volumes, volume groups, and thin clones). Native replication utilizes the iSCSI protocol to transfer data between two PowerStore clusters. Replication rules can be applied on a policy basis to provide various recovery point objectives for your storage resources.

Clustering provides the ability to scale up and scale out appliances independent of one another while being managed from a single management interface. This flexible deployment model enables PowerStore to start small with a single appliance configuration, then grow into a larger cluster by simply adding additional appliances online. As part of clustering, we also support non-disruptive migrations of volumes within the cluster.

VMware® integration is deep and rich with PowerStore. When administrators are leveraging virtual machines on the PowerStore system, native VMware features such as vSphere® High Availability (HA) are automatically enabled to ensure virtual machines remain available in the event of an outage. Other features, such as VMware Fault Tolerance (FT), can be enabled by administrators to guarantee the highest levels of availability.

Performance metrics in historical and real-time views enable monitoring for anomalies and troubleshooting performance issues. This enables administrators to monitor their environment and be proactive in preventing potential outages. Performance metrics are fully integrated into the PowerStore environment allowing for ease of use and customization.

Performance policies can be applied to block level resources (volumes, volume groups, VMware vSphere Virtual Volumes™ (vVols), and thin clones) based on a high, medium, or low setting. Applying these policies to the storage objects help prioritize I/O requests to different hosts and applications. For example, a storage administrator may apply a high setting for a volume that is utilized by a database application such as Microsoft® SQL Server. However, for other applications such as a test/dev vVol, the administrator may choose a low setting.

Non-disruptive upgrades to target PowerStoreOS software ensures the best possible resilience when it comes to PowerStore software and feature functionality. This is accomplished by upgrading one node at a time and ensuring all resources are running on the node not undergoing the upgrade. PowerStore also comes with **anytime upgrade** options which protect the customer's investment in the event they need to upgrade their hardware model.

CloudIQ and SupportAssist configurations enable proactive awareness of system health and performance as well as providing a holistic view of multiple storage systems regardless of if they are on the same network. SupportAssist allows for direct connection to support with automated service requests for hardware and software faults as well as remote connection to the array for faster troubleshooting and resolution.

© 2021 Dell Inc. or its subsidiaries. All Rights Reserved. Dell Technologies, Dell, EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be trademarks of their respective owners.

Configuring a high availability environment

A best practice for a highly available storage environment is to define, implement, and regularly test a business-continuity, data-availability, and disaster-recovery plan. Be sure to fully understand the implications of global settings and their impact on system operation and always record the original settings before making any configuration changes.

In addition to resilient LAN and WAN infrastructures and a redundant environment (power and cooling), having out-of-band access to and remote power control of all nodes and switches is invaluable when administering clustered systems. In a worst-case scenario, if a cluster does lose power, the write cache is battery protected and destaged to preserve any in-flight uncommitted writes to maintain write consistency.

Block guidance: Utilize redundant network switches at both edge and core connections to maximize path availability. Native asynchronous block replication should be used whenever there is a requirement for disaster recovery. Include snapshot rules in the protection policy for point in time recovery.

File guidance: For PowerStore T model appliances, file services run on the first two ports of the 4-port card known as the system bond. To ensure these ports have the most optimal performance and fault tolerance, it is highly recommended to configure multi-chassis link aggregation group (MC-LAG) on the network switches. For example, on Dell EMC switches, an administrator should configure Virtual Link Trunking (VLT).

Resiliency products: Leveraging products such as PowerStore metro node or Dell EMC VPLEX™ in front of PowerStore hardware enables the highest possible uptime for all block configurations. PowerStore metro node or Dell EMC VPLEX™ as a storage-virtualization layer, replicates data across two PowerStore appliances in the same data center or across metro distance data centers to increase data availability. This extra layer of separation provides extra fault tolerance.

With VPLEX, PowerStore volumes are mirrored, always mounted read/write and synchronized through VPLEX local/metro, providing continuous availability to hosts in the event one of the PowerStore systems becoming unavailable.

Measuring availability

In order to determine your storage system's lifetime data availability, we measure the percentage of time the storage system is operational—or servicing user read/write operations. The data we collect is modeled using product redundancy features, parts replacement rates, mean time to restore a hardware failure, and defined service support levels. Because of the redundancy and concurrent maintenance philosophy built into every Dell EMC storage product, it is important to note that a clear majority of the system issues and subsequent repairs and replacements will not affect your overall system availability.

In other words, we view data availability along with system durability from a total quality engineering perspective across all aspects of the array in combination with your data environment. In doing so, Dell Technologies continuously brings you closer to 100% data availability.

© 2021 Dell Inc. or its subsidiaries. All Rights Reserved. Dell Technologies, Dell, EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be trademarks of their respective owners.

By configuring your environment to maintain the highest availability, we will in turn measure our ability to deliver on that promise. Through ongoing monitoring by Dell Technologies using SupportAssist and CloudIQ features, we can determine field reliability and availability which is an integral part of the total customer experience. Throughout this monitoring process, all unplanned outage events are analyzed by engineering and service personnel in order to identify root-cause, and the findings are used towards continuous improvement of product, process, and personnel.

Dell Technologies can deliver on our promise of delivering you persistent data availability.

Other resources

- Dell EMC PowerStore: Introduction to the Platform
- Dell EMC PowerStore: Technical Primer
- Dell EMC PowerStore: Clustering and High Availability
- VPLEX Overview and General Best Practices
- Dell EMC PowerStore: Best Practices Guide
- Dell EMC PowerStore: Replication Technologies

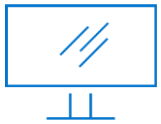

Learn more about Dell EMC PowerStore solutions

Contact a Dell Technologies expert

© 2021 Dell Inc. or its subsidiaries. All Rights Reserved. Dell Technologies, Dell, EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be trademarks of their respective owners. Reference Number: H18494.1