EMC PERSPECTIVE

Health Check: Helping You to Manage Your Midrange Environment

Uncompromised availability

Whether the system is a customer-facing online ordering application or a back-office communications platform, the business runs on information. The costs of data loss or system downtime are high. That is why managing business risk begins with managing data access, data availability, and application response times.

Historically, IT organizations have focused on achieving very high availability for their mission-critical applications, investing in both hardware and software solutions to keep the business running. Over time, with advances in technology, business-critical applications have migrated to midrange systems and storage arrays. More recently, IT organizations that need high performance, redundant data paths and power supplies, or "hot-plug expandability" can find it all in midrange arrays. These systems also promise flexibility in configuration, "just-in-time" storage allocation, and disaster recovery software. Internal messaging applications as well as customer-facing, OLTP applications can benefit from performance, scalability, and "five-nines" (99.999 percent) availability. With a well-planned deployment, IT organizations can deliver significant business benefit on a midrange platform.

Changing environments

But managing risk and meeting SLAs does not end at deployment; it begins there.

While initial implementations may meet SLAs, most IT environments are in an almost constant state of flux. The IT organization does its best to ensure adherence to both industry and storage vendor best practices to provide the highest quality of service—delivering high availability and continuously improving application performance. Yet, the dynamic nature of the business and its information needs often extend the storage infrastructure beyond its originally intended capacities and design.

The Information Explosion

International Data Corporation (IDC) has reported a greater than forty percent year-over-year growth in the disk storage market. As the information explosion continues, IT departments expand and reconfigure systems to provide the needed storage.

IT Growth and Consolidation

As companies grow, more desktops, servers, applications, and SAN storage are added. Mergers and acquisitions drive the integration of corporate infrastructures and test system compatibility and interoperability. At the same time, constant cost pressures drive consolidation. Eliminating capacity in one area may increase utilization in another. Changes in resource utilization affect system performance and necessitate configuration changes.

New Business Processes

In many cases, external forces bring pressure to bear. New industry regulations or heightened security awareness cause reconsideration of information lifecycle management (ILM) strategies and service-level agreements (SLAs). New applications (e.g., disaster recovery) introduce new issues and necessitate further configuration changes.

Vendor Updates

Finally, the technology itself changes. Vendors introduce new functionality or issue patches and, once again, configurations change.

In any of these cases, the storage configuration will change, but user expectations do not change. Long after installation, the storage system must continue to deliver rapid I/O and high availability for every application. Yet, with so many modifications to the original configuration, it is likely that the storage array is no longer optimally configured. And a sub-optimal configuration can lower performance and increase the risk of data loss.

Proactively managing risk

Most organizations have a plan in place for dealing with major catastrophes. Critical systems will fail over to a distant site, meeting requirements for recovery-point objectives (RPOs) and recovery-time objectives (RTOs).

However, many outages and performance problems occur, not as the result of a major disaster, but rather as a result of small missteps such as software faults, data migrations, upgrades, and operator error. In addition to the cost of system downtime, the loss of a server's data consumes valuable human resources to recover and restore. Issues such as performance bottlenecks, non-optimally configured volumes, or insufficient or misconfigured disk hot spares can slow or even halt productivity. Nearly half of data unavailability/data loss (DU/DL) events are associated with configuration errors.

Benchmarking the Configuration

Array performance and data availability keep business operations running smoothly. That is why, before making changes to any storage array, it is important to document the existing configuration and review the system's performance data. This benchmark gives the organization a starting point to:

- Determine whether increases in capacity will have unexpected impacts on disk response time and performance
- Determine whether the array has the capacity to support additional hosts
- Ensure compatibility of the configuration of components (e.g., hosts, switches, software)
- Compare to best practices to identify issues which may impact performance, data availability, or data loss

Utilization Creep

Capacity and performance utilization rates increase incrementally over time, placing increasing demands on the storage array. This may degrade performance and may result in inefficient configurations. Proactively checking the system can uncover these situations and ensure that SLAs continue to be met.

Availability Validation

In some cases, servers which were added to an array after the initial configuration may not be installed with the required high-availability software. While the servers may perform smoothly in this configuration, applications are at risk of an outage in the event of a path failure. A proactive check of the system should include a check of the high-availability configurations of all servers.

Maintaining Supported Configurations

IT organizations often create "gold copy" images for deployment on servers to make certain that each configuration is a uniform, supported configuration which includes all the add-on software for their typical server configuration. Often these gold copies are not updated when new revisions of HBA drivers or software revisions are released. This means the gold copies may be deployed with outdated software versions. Again, the system may perform smoothly and the problem may go unnoticed for months. However, if there were a need for vendor support, the outdated configuration may be the source of the problem, and the vendor may identify the configuration as non-compliant or unsupported from an interoperability perspective. That is why a proactive check should include an interoperability check of all hardware and all software versions and a check of compliance with vendor-recommended practices.

Balancing Performance and Capacity

Storage managers are quite familiar with the need to balance performance and capacity when configurations are designed. Over time, the addition of servers and/or storage capacity may upset this balance. LUNs may be balanced across I/O ports but, because some applications create more I/Os than others, cumulative traffic to the I/O ports may become significantly unbalanced. A proactive system check should consider this so that performance is optimized and SLAs are met.

Increasing IOPS

When additional hosts are added to an array, the I/O profile changes and Fibre Channel links may become over-utilized. This may degrade performance and eventually result in missed SLAs. To prevent this, a proactive system check should compare system I/O to the benchmark I/O profile established earlier or to vendor best practices and guidelines.

A proactive check enables the IT organization to uncover potential performance and availability issues before they cause significant problems. In many cases, a few minor adjustments or additions to existing configurations may extend the usable service life of an array, providing improved response times and additional headroom for growth. Most importantly, missed SLAs or situations resulting in data unavailability, in many cases, can be avoided.

The high rate of change in storage and server environments can cause these and other problems, which may impact uptime and time to recover in the event of a failure. Because high availability is a critical component of IT business value, IT organizations must anticipate and proactively manage these issues. In many cases, the best approach is to rely on the storage vendor for assistance.

Vendor-delivered health checks

Before making purchase decisions, IT organizations look carefully at vendor ability to support products after installation. There is good reason for this. When things go wrong, the vendor's support organization is usually in the best position to remedy the situation. It has the diagnostic tools, the skills to use them, and the body of expertise collected from many other customers. That is why having the storage vendor routinely check the health of storage arrays makes sense.

The storage provider benefits from the collective experience of many support professionals working on a range of storage configuration challenges. Their "best practices" have evolved from that experience. They have developed proprietary diagnostic, monitoring, and analysis tools which enable them to efficiently identify and resolve problems. In a relatively short time, the storage provider can review the status of the entire configuration (i.e., LUNs, metaLUNs, RAID groups, storage groups, storage processors, host path settings, connected switches and servers) to ensure optimal configuration and adherence to best practices.

When IT organizations use the storage provider to assess the system, they can rely on expert recommendations to reallocate capacity or modify the configuration in order to ensure optimum performance and availability.

EMC² where information lives®

EMC Corporation
Hopkinton
Massachusetts
01748-9103
1-508-435-1000
In North America 1-866-464-7381
www.EMC.com