Prescriptive Guide to Achieving Extreme Data Durability

An ENTERPRISE MANAGEMENT ASSOCIATES® (EMA™) White Paper
Prepared for Western Digital

April 2017

EMA™ IT & DATA MANAGEMENT RESEARCH, INDUSTRY ANALYSIS & CONSULTING

Table of Contents

Prescriptive Guide to Achieving Extreme Data Durability

Challenges to Long-Term Data Protection

Organizations are increasingly recognizing that data is their most valuable asset. Gone are the days when data was considered as principally a disposable resource—when the cost of storage systems was considered more valuable than the data it hosted. Today, businesses want to retain larger data sets for longer-term retention periods to support active archiving requirements, meet regulatory compliance commitments, and consolidate large record sets to perform analytical reporting and historical trending. As if this was not imposing enough pressure on data center managers, organizations are also demanding data be eminently available for immediate access. Modern business requirements simply no longer allow for data to be recovered from off-site archives or restored from tape libraries. Also, traditional backup and archive processes are not always reliable. Without continual, time-consuming testing of backup services (which is rarely performed), there is no way to ensure the data is even recoverable. Lacking the ability to easily and reliably access critical business data, organization face severe disruption to business performance, costly outages of production services, damaged reputations, lost customer revenues, and an inability to meet regulatory requirements.

Ensuring the integrity of data in long-term storage is the primary focus of solutions architected to support data durability. However, traditional RAID solutions that fall into this category have proven to be not very reliable at scale. Put simply, the larger the capacity of the drive in a RAID array, the greater the chance of bit errors during a rebuild, and with multi-terabyte drives now in common use, the risk of data corruption has increased exponentially. Additionally, RAID rebuilds take between 24 to 36 hours per terabyte to complete; during which time drives are stressed, increasing the chance of a second drive outage that would result in data loss. Even if the RAID device operates as advertised, it still will not protect the data in the event of a site failure, such as due to a flood, natural disaster, or other catastrophic event.

Ensuring the integrity of data in long-term storage is the primary focus of solutions architected to support data durability.

Achieving Extreme Data Durability

To achieve the extreme data durability necessary to meet today's more demanding storage requirements, organizations must adopt more robust technologies than found in antiquated RAID platforms. Object storage solutions are particularly suited to managing long-term storage of larger amounts of unstructured data. By managing data as individual, self-contained objects, rather than a file hierarchy or a collection of data blocks, unique descriptors (i.e. metadata) can be associated with data to ensure its optimal placement, enable more efficient indexing, and to provide intelligence on its use and storage requirements. Because object storage solutions are designed around a "shared nothing architecture," there is no single point of failure or scalability limits that would be imposed in more traditional storage approaches. This allows object storage environments to reliably grow from relatively small capacity solutions to supporting hundreds of terabytes or more of data while still achieving a high degree of reliability.

It is important to note here that "high-availability" is not the same as "data durability." Put simply, storage availability is ensured through hardware redundancy, while durability is achieved through data redundancy. This is an important distinction because data stored on a highly-available system can still be corrupt and irretrievable even though the physical storage platform devices are fully operational. Because of this, many storage vendors promoting high-availability solutions only commit to achieving 6 nines

(99.9999%) of data durability. This may seem on the surface to provide sufficient data reliability, but statistically it indicates one out of every million stored objects will not be accessible. In an environment that hosts just 10 terabytes of data with an average file size of 20kB, 500 files can be expected to be lost and unrecoverable with a traditional high-availability-focused environment. Financial institutions, for example, sometimes store data transactions representing tens of millions of dollars, which places the businesses at risk if that transaction happens to be one of the files that is lost. Storage solutions architected for data durability achieve a much higher rate of reliability because individual objects are recoverable from multiple physical locations.

Data durability with an object storage platform is typically achieved with the employment of a technology called, "erasure coding." Data objects are broken into fragments that are each expanded to include parity data which are then distributed across disparate storage capacity comprised of multiple drives, shelves, and even data center locations. This built-in data protection through distributed redundancy allows the storage solution to self-heal by simply reconstructing any data bits detected to be damaged from parity data stored elsewhere. In this way, there is never a need for rebuilding full disk drives to restore the data integrity of a single object, eliminating the downtime and hardware stresses commonly associated with traditional rebuild processes. Additionally, since data and fragments are distributed across multiple physical drives, shelves, racks, or even data centers, erasure coding ensures continuous data availability even if multiple drives should fail, up to the limits of the configuration.

While many object storage platforms employ erasure coding within a storage array or storage cluster, these are most commonly single-site solutions that are not designed to natively achieve full disaster recovery in the event of a site loss. Third-party backup, archiving, and replication solutions can certainly be employed to enable off-site data storage options, but these can be costly and often require failover downtime in the event of a site loss. A more pragmatic approach is to introduce a solution that geographically disperses erasure code fragments to multiple sites so that data is continuously accessible even if a full site is lost.

While many object storage platforms employ erasure coding within a storage array or storage cluster, these are most commonly single-site solutions that are not designed to natively achieve full disaster recovery in the event of a site loss.

As an example of a platform that offers extreme data durability at scale, the HGST ActiveScale™ storage solutions have been architected to deliver 17 nines (99.999999999999999%) of data durability. This exceeds even the most stringent service level agreements for availability meaning data will be accessible once the lights turn back on. Object storage systems supporting multi-site architectures, the ActiveScale platforms employ a patented BitSpread® technology to dynamically execute the erasure coding algorithm and perform hierarchical data spreading functions based on the existing state of the system. This enables higher levels of throughput, faster self-healing, and improved data integrity. The platforms also incorporate BitDynamics™ technology to perform continuous data integrity audits and automated repair functions to proactively guard against bit rot or other potential data corruption. The ActiveScale solution is eminently scalable with multiple editions designed to scale both up and out: the ActiveScale P100—supporting 720TB to 19PB raw capacity (or 508TB to 12.9PB usable capacity) and the ActiveScale X100—supporting 840TB to 52PB raw capacity (or 578TB to 33.8PB usable capacity). All solutions are managed from a centralized, single-pane-of-glass management environment (ActiveScale SM) per namespace, and provide information to cloud-base storage analytics (ActiveScale CM).

EMA Perspective

Today's IT-centric business models have little or no tolerance for data failures. The loss of key data records may result in significant impacts to business reputations and revenue attainment, and this challenge is only exasperated by the fact that data is continuing to grow exponentially. Traditional storage solutions simply were not architected to handle today's large data demands and replication and mirroring processes are slow, expensive, and resource intensive. Out of desperation, many organizations adopt costly solutions that are not aligned with their business objectives.

To responsibly achieve extreme data durability, organizations must adopt a strategic approach to architecting their storage ecosystem so that it meets current and expected requirements for scalability, high-availability, and performance. Additionally, solutions should be easy to manage to minimize the costs and efforts of operational support. For organizations with large data requirements, the key to achieving pragmatic and reliable data storage is the introduction of a platform, such as the HGST ActiveScale solutions, built on an object storage architecture with erasure coding resilience and the option for seamless distribution of data even across multiple data center locations.

About Western Digital Corporation

Western Digital Corporation (NASDAQ: WDC) is a provider of storage technologies and solutions that enable people to create, leverage, experience and preserve data. The company addresses ever-changing market needs by providing a full portfolio of compelling, high-quality storage solutions with customer-focused innovation, high efficiency, flexibility, and speed. Products are marketed to OEMs, distributors, resellers, cloud infrastructure providers, and consumers.

About Enterprise Management Associates, Inc.

Founded in 1996, Enterprise Management Associates (EMA) is a leading industry analyst firm that provides deep insight across the full spectrum of IT and data management technologies. EMA analysts leverage a unique combination of practical experience, insight into industry best practices, and in-depth knowledge of current and planned vendor solutions to help EMA's clients achieve their goals. Learn more about EMA research, analysis, and consulting services for enterprise line of business users, IT professionals, and IT vendors at www.enterprisemanagement.com or blogs.enterprisemanagement.com. You can also follow EMA on Twitter, Facebook, or LinkedIn.

Corporate Headquarters:
1995 North 57th Court, Suite 120
Boulder, CO 80301
Phone: +1 303.543.9500
Fax: +1 303.543.7687
www.enterprisemanagement.com

3556.032917