

HOW TO SELECT THE CORRECT RAID FOR AN EQUALLOGIC SAN

One of the most important decisions that you have to make when deploying a new storage solution is what RAID type(s) to use. Each RAID type has its advantages and disadvantages in terms of protection, efficiency (usable storage), and performance—both during normal operations and in terms of rebuild time. Regardless of who a storage solution is purchased from, RAID decisions will be based on the same set of criteria. First and foremost, the application's performance needs must be met, and this should be the primary decision factor when selecting a RAID policy. Once the application(s)'s needs are addressed, choosing a RAID policy comes down to two factors: cost—in terms capacity utilization efficiency—and protection level.

This white paper discusses how to select the right RAID policy for your Dell EqualLogic™ storage area network (SAN) arrays. First, each of the RAID types available with EqualLogic arrays is reviewed, in terms of performance, protection and efficiency. Next, we will discuss choosing a RAID policy within the EqualLogic SAN architecture to meet your performance and protection requirements.

OVERVIEW OF RAID TYPES AVAILABLE FOR EQUALLOGIC SANS

EqualLogic PS Series arrays support several different RAID types. Each type has been optimized to maximize its performance within the EqualLogic architecture. Within this architecture, RAID types are selected on an array by array basis, each array being configured with a single RAID policy. By default, this RAID policy focuses on two factors: the RAID type for the array and the number of disks reserved as hot spares. The RAID types available for use within each EqualLogic array are RAID 10, RAID 50, RAID 5, and RAID 6. Depending on the array model, each of these RAID types is implemented with differing numbers of drives dedicated to data, parity, mirrors, and hot spares based on the number of drives available in each model. Table 1 below provides a general definition of each RAID type.



EQUALLOGIC

Note: EqualLogic PS Series arrays implement each supported RAID type based on the number of disks in the array model. The diagrams in Table 1 are used only to illustrate the basic definition of each RAID type and do not represent the actual RAID implementation of a PS Series array.

Table 1: RAID Definitions

RAID	Description	Example
10	<p>RAID 10 (mirrored sets in a striped set) combines the best of two high performance RAID types — RAID 0 and RAID 1. A RAID 10 is created by first building a series of 2 disk RAID 1 mirrored disks, and then distributing data over those mirrors. Given 16x 1 TB disks, a RAID 10 implementation would be able to host 8 TBs of data.</p> <p>RAID 10 implementations can sustain multiple drive losses so long as no RAID 1 (mirrored) set loses both of its drives.</p>	
50	<p>RAID 50 (RAID 5 sets in a striped set) is created by first creating two or more RAID 5 sets, and then striping data over those RAID 5 sets. Given 16x 1 TB disks, a RAID 50 implementation would be able to host 14 TBs of data and reserving 2 TB of space for parity.</p> <p>RAID 50 implementations can tolerate a single drive failure per RAID 5 set.</p>	
5	<p>RAID 5 (striped disks with distributed parity) combines three or more disks in an arrangement where each stripe consists of N data blocks and 1 parity block (often referred to as N+1). Given 16x 1 TB disks, a RAID 5 implementation would be able to host 15 TBs of data and reserving 1 TB of space for parity.</p> <p>RAID 5 implementations can tolerate a single drive failure without data loss.</p>	
6	<p>RAID 6 (Striped set with dual distributed parity) extends the concept of RAID 5 to the use of 2 parity blocks per stripe (referred to as N+2). Given 16x 1 TB disks, a RAID 6 implementation would be able to host 14 TB of data and reserving 2 TB of space for parity.</p> <p>RAID 6 implementations can tolerate up to 2 drive failures at the same time without data loss.</p>	

Note: Dx = Data blocks, Mx = Duplicate Data, Px = “Vertical” Parity blocks, Qx = “Diagonal” Parity blocks

RAID SELECTION CRITERIA

When comparing RAID types, three specific characteristics should be considered: performance, storage efficiency, and data protection. These characteristics will determine which RAID type will meet your application, business and administrative needs. As a storage administrator, you must work with the application owner to determine which of these characteristics are more important. To help gain a better understanding, consider each in detail.

Performance

One of the most important requirements of any application—and thus the storage solution supporting it—is meeting the end-user’s expectations for speed and responsiveness. The performance characteristic of a storage solution’s RAID type has a direct impact on this application requirement. The faster the storage solution is in responding to requests, and to providing data, the faster and more responsive the application. RAID performance is made up of a mixture of factors including:

- **The number of disks dedicated to reading and writing of data**

RAID types that dedicate more disks to processing data will provide better performance. This is different from storage efficiency, which is focused on the number of disks that can be used to respond to a storage request. A RAID 10 array has a low storage efficiency requiring 50% of its disks to be dedicated to protecting data in the form of a mirror, but it can use both copies of the data in that mirror to respond to storage requests. This allows all of the disks (100%) to participate in processing data requests. In contrast, a RAID 5 array has a high storage efficiency characteristic since it only requires one disk’s worth of storage to protect the rest of the disks, but to have this high storage efficiency, it requires additional processing power that is taken away from processing storage requests.

- **The amount of processing required to provide data protection**

This factor is focused on RAID types that require parity calculation such as RAID 5 and RAID 6 arrays. Calculating parity is a very compute-intensive process and requires the storage array to focus some of its processing power to performing these calculations rather than simply sending or receiving data from an application. For every write operation, a RAID 5 array requires additional read operations, as well as processing to calculate the parity, before it can commit the data write request to disk. A RAID 6 array must do this twice as it requires the calculation of two different parity values. In addition, the RAID 6 array must do these parity calculations in sequence since the second parity must include the first parity calculation in its own process.

- **Rebuild time after recovering from a failure**

After a disk failure, a hot spare, if available, will automatically replace the failing disk and the array will enter a rebuild state where data from the failed disk is reconstructed using the data from the remaining disks in the RAID array. Depending on the RAID type, the amount of time and processing power required to complete this rebuild can have a dramatic affect on the array’s performance until the rebuild is complete. RAID types involving mirror sets, such as a RAID 10 array, can simply copy data from the surviving disk in the mirror pair. During rebuilds of RAID 10 arrays, there are no calculations to perform, and the other mirror sets within the RAID 10 set are not affected by the rebuild. RAID types involving parity calculations, such as RAID 5, RAID 50, and RAID 6 may see a drop in array performance. This is due to the need to read data from all the surviving disks in the failed RAID set and recalculating the parity information to determine the value of the lost data as part of the rebuild process

Storage Efficiency

In terms of storage efficiency, each RAID type has differing levels of storage efficiency. Storage efficiency indicates how much “usable storage” space is available after applying the requirements of the RAID type to the set of disks. With most RAID types, storage efficiency goes up as the number of disks increases. Others have a fixed efficiency in terms of percentage of raw disk space, regardless of the number of disks in the RAID set. For example, a RAID 10 array always requires 50% of its disks to be used to mirror data from the other 50% of disks, while RAID 50, RAID 5 and RAID 6 arrays’ storage efficiency will vary based on the number of disks in the RAID set as they require fewer disks to protect the rest of the disks in the RAID set, which are storing data.

Data Protection

Finally—and probably most obviously—each RAID type provides differing levels of protection. Some RAID types will protect from single disk failures, while others can protect the RAID set from multiple disk failures. RAID 10 can recover from multiple disk failures—one disk per mirror set. Therefore, an EqualLogic PS array with 16 disks—14 disk in the RAID set and 2 hot spares—consists of 7 mirror sets per array and could potentially continue to operate with 7 failed disks. On the other end of the spectrum, RAID 5 (and other single parity RAID type) arrays can only tolerate a single disk failure. RAID 6 and RAID 50 provide protection between those of RAID 10 and RAID 5, by allowing multiple disk failures—2 disks per array in the case of RAID 6 and 1 disk per RAID 5 set within the RAID 50 array—to occur before the array is disabled. Table 2 below provides relative ranking of each of the RAID types supported by EqualLogic in terms of these three characteristics.

Table 2: RAID Comparison

RAID	Reads (Random/Sequential)	Writes (Random/Sequential)	Relative Cost	Relative Protection	Rebuild Performance
10	★★★★★★★★	★★★★★★	\$\$\$\$	★★★★	★★★★
50	★★★★★★	★★★☆☆☆	\$\$\$	★★★	★★★
5	★★★★★★	★★☆☆☆	\$	★★	★★
6	★★★★★★★★	★☆☆	\$\$	★★★★	★

SELECTING THE RIGHT RAID POLICY

Selecting a set of RAID policies for your SAN will have a direct impact on the both the actual and perceived performance of each application that is using the SAN for its primary storage. As discussed earlier, each RAID type has differing levels of performance, cost, and protection. To be able to make informed decisions, understanding the reasons for deploying a new SAN will set some expectations. First, most SANs are deployed to consolidate storage into a single pool of storage to improve management, reduce the amount of isolated, unused storage capacity, and to better protect the data being stored—possibly due to regulatory requirements. Consolidation implies that multiple applications will be utilizing storage resources within the SAN. It is therefore important to understand each application’s data access profile. Using this information, the EqualLogic SAN can be designed to meet the requirements of all of the applications with minimal compromise by defining where each application’s data will be hosted within the SAN. This allows us to define RAID types that focus on your application’s requirements and grouping those application’s data into resource pools that require similar RAID policies.

UNDERSTANDING APPLICATION REQUIREMENTS

Understanding your application’s data requirements is the first step in selecting the right array model and RAID policy. Is the application accessing data in a random or sequential fashion? How responsive must the application be to the end user’s requests? How often does the data change (how many reads versus writes)? These are just a few of the questions that will affect SAN design and RAID policy selection along with how many RAID policies might need to be deployed.

The prevalent application environments are random access in nature, which are bound by IOPS and latency. Examples of these types of application are transaction oriented database applications, point-of-sale applications, and large email solutions. Since these types of applications require storage solutions that can generate high I/O with minimal latency, the storage solution should support RAID types that can be created with a large numbers of fast disk drives. Any RAID type supported by the EqualLogic PS Series array would be suitable to this type of access model, but RAID 10 and RAID 50 are probably the best here because these RAID types work better where requirements indicate a higher percentage of data writes while providing a large number of spindles to support higher read I/O and relatively better write I/O.

Less common are applications that access data in a sequential fashion such as data warehousing, video processing, and other large decision support applications. These applications require the storage solution to have the ability to transfer large amounts of data (usually reads and not writes) in a timely fashion and require very large network and storage configurations to meet overall throughput. A single storage array may fall short of the high throughput required by an application that accesses data sequentially, but this bottleneck can be removed by appropriately scaling the SAN, through the addition of other

EqualLogic arrays along with scaling the number of I/O paths or initiators from the host. RAID types that are most suitable for these types of applications typically are the most space efficient as well as read access friendly such as RAID 5, RAID 6, or RAID 50.

As mentioned at the beginning of this section, application requirements should play a primary role within the RAID policy selection process. While nothing can substitute for real-world testing, some general characteristics for various types of applications can provide some guidance. Table 3 lists some of the primary application types along with the suggested RAID policies that should be considered.

Table 3: Application Types and RAID Considerations

Application Type		Candidate RAID Policies
File/NAS Services		5, 6, 50
Messaging		10, 50
Database	Random, Transactional (OLTP)	10, 50
	Sequential (>75% Read), DSS/ Data Warehouse	6, 5, 50
	Sequential (> 25% Writes)	50, 6, 5

IMPACT OF THE VIRTUALIZED PEER STORAGE ARCHITECTURE ON RAID SELECTION

EqualLogic SANs consist of one or more PS Series arrays working cooperatively to provide optimal performance and protection to your applications via load balancing. Each array has a dedicated RAID policy allowing the array to automatically optimize its performance based on the characteristics of that RAID type, providing the best performance possible from any of the supported RAID policies. Each RAID policy consists of a RAID type and a defined number of disks designated as hot spares. By default, EqualLogic reserves 1 drive as a hot spare for RAID 5 and RAID 6 and reserves 2 drives as hot

spares for RAID 10 and RAID 50. The number of disks reserved as hot spares can be reduced, but only through the use of the command line interface (CLI). Dell does not recommend that the number of hot spares be reduced from the default number.

The power of the EqualLogic SAN design comes from combining multiple arrays. Each array in the SAN group can be configured with different RAID policy. Organizing arrays into different resource pools provides a tiered storage feature that may help to improve overall performance. Resource pools can host arrays with the same RAID policy as well as arrays with differing RAID policies. Each resource pool allows volumes to be load balanced across multiple arrays within the pool based on application performance needs.

When a resource pool contains arrays that all have the same RAID policy, additional arrays added to this pool not only increases the capacity of the pool, but also increases the performance through the addition of additional process power, additional network connections, and additional spindles to host each volume in the pool. When a resource pool contains multiple arrays with different RAID policies, the SAN can take advantage of the differing RAID types by automatically moving volumes (and the data they contain) onto a different RAID type depending on the usage patterns experienced as the volume is accessed over time. This process is automatic by default and happens without the awareness of the applications that are using the volumes.

Another feature of the virtualized peer storage architecture of the EqualLogic family is the ability to, with some limitations, change an array's RAID policy. So, even if you do not initially pick the right RAID policy for your application(s) or your application's requirements change, you might be able to change that RAID policy to a better RAID policy—all without the awareness of the application other than a temporary performance change as the

migration is performed. For small SANs consisting of only a few arrays, this RAID policy conversion ability is a great feature to have available. Table 4 lists the possible RAID policy migration options.

Table 4: RAID Migration Options

Current RAID Policy	Can Convert To
10	5, 50, 6
5	NONE
50	5, 6
6	5, 50

CONCLUSIONS

RAID selection is a key decision that must be made for any advanced, consolidated storage area network solution. Before any RAID policy choices can be made, you must understand what each RAID type is, what the pros and cons of each are in terms of performance, cost, and the amount of protection they provide and taking your applications storage performance requirements into account. Only then can you make informed choices during the planning phase of your storage design.

Within the EqualLogic architecture, RAID plays an important role in ensuring the storage solution that provides the best performance possible while also providing cost conscious data protection along with advanced features only available with a virtualized storage solution. This virtual peer storage model allows you to not only grow your SAN's capacity as your needs change, it also allows you to implement more advanced functionality and provides flexibility features such as automated volume migration, automatic distribution of volumes over multiple arrays with the same RAID policy and the ability to organize the storage into resource pools. All of these capabilities can be implemented without negative impacts on the operation of the host application and without bringing the storage down for maintenance.

SIMPLIFY YOUR STORAGE AT WWW.DELL.COM/Pseries



EQUALLOGIC

THIS WHITE PAPER IS FOR INFORMATIONAL PURPOSES ONLY, AND MAY CONTAIN TYPOGRAPHICAL ERRORS AND TECHNICAL INACCURACIES. THE CONTENT IS PROVIDED AS IS, WITHOUT EXPRESS OR IMPLIED WARRANTIES OF ANY KIND.