

RELDATA

FEATURE BRIEF: Overcoming RAID Group Limitations with R-Pools

With its revolutionary storage operating system featuring dynamic R-Pools, the RELDATA 9400 Storage System delivers true pay-as-you-grow scalability by overcoming the many limitations of a traditional RAID group architecture. R-Pools also provide increased performance and reliability, offering a radically improved management scenario for administrators.

The Current Limitations of Storage Systems Based on RAID Groups

Many of today's storage systems are architected around the concept of static RAID groups, introducing a wide range of management complications that plague administrators. All system expansion must be based around the existing RAID group construct, making scalability complex, costly, and rigid. For example, if the RAID groups are based on an 8-drive RAID set, drives must be added in 8-drive increments. All the drives in that drive set must be same size and speed. Once the drives are added to the system, the administrator must typically restripe the LUNs across the new drives in order to benefit from the performance of the recently added drives. In addition to the complexity and time-consuming nature of this process, the re-stripe is inherently risky.

Introducing Scalability with R-Pools

RELDATA's next generation dynamic storage pooling architecture completely removes the limitations associated with traditional architectures that rely on static RAID groups. R-Pools represent the missing link between the physical and virtual volumes, taking raw physical storage devices on one side and providing virtual LUNs with the required size, redundancy and storage class characteristics on the other.

An R-Pool is a collection of drives, with no restrictions on the size and speed of the drives comprising the pool. Data is striped across this set of drives, with RAID implemented as fixed size "chunks" of data – data, parity, or spare -- within the stripe. This architecture has three major benefits:

1. Different types of RAID layouts are supported across the same set of disks
2. Disks can be added to the pool, with no regard to RAID settings
3. Parity failures can be addressed at the stripe level, avoiding costly re-striping of an entire drive in most cases

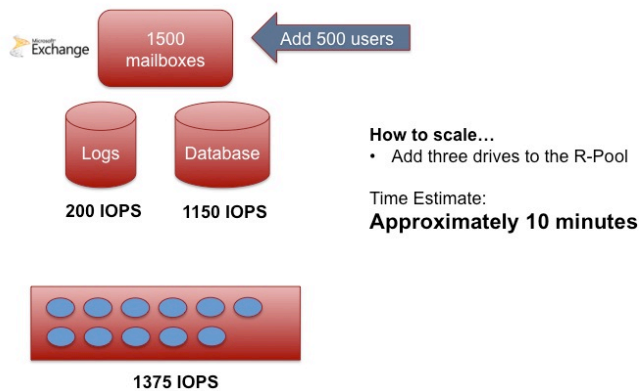
The R-Pool approach is much more granular than a RAID group, freeing administrators from dealing with drives in a static way and greatly simplifying the process of adding capacity.

Drive-by-Drive Scalability

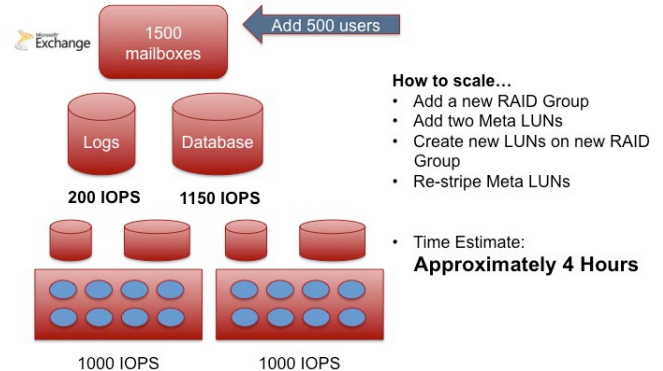
With dynamic R-Pools, you no longer need to add drives in fixed groups. Capacity can be added drive-by-drive, mixing drive types and sizes as desired, for true pay-as-you-grow scalability. As new, higher capacity drives become available, you can add them directly to the pool in combination with existing drives. New drives automatically add performance to the pool. Because striping for redundancy is done across all drives rather than across drives in a RAID group, the 9400 Storage System completely avoids a lengthy and risky re-build process upon the addition of new drives, instead automatically load balancing the data onto the new drives over time. All of the behind-the-scenes processing required to add capacity is completely automated by the 9400.

RELDATA

Adding Capacity to an R-Pool



Adding Capacity to a RAID Group



Increased Reliability and Performance

Within R-Pools, all the chunks of data are written to a stripe, and then the entire stripe is committed to disk together to a new physical location. This approach has two benefits:

- The “RAID write hole” which exists in most controller-based RAID implementations is eliminated
- The previous data is left in place

By not overwriting previous data, virtualization features such as snapshots are remarkably easy to implement. The storage operating system simply does not invalidate the older data, making snapshots, deduplication or block span copy very low cost activities. This approach results in substantially fewer I/Os than the so-called “copy on write” technology deployed in many legacy storage arrays.

As mentioned above, R-Pools handle parity failures at a granular level. This is important as today’s disk drives are 2TB in size, and 4TB and 8TB drives are coming soon. Restriping all that capacity will consume days and even weeks, leaving you exposed to another failure and probable data loss. Fixing parity problems at the stripe level delivers much higher overall reliability. First, in the case of the failure of a few sectors on a drive, only the small number of affected stripes are rebuilt. In the event of an entire drive failure, only those stripes that have been written need to be rebuilt. Both scenarios dramatically reduce the amount of data to be re-constructed and limit your exposure.

R-Pools also support the idea of “spare chunks” rather than spare drives. This allows re-builds for each stripe to start immediately, rather than waiting for new drives to be added. Lastly, you may elect much more aggressive RAID settings for key data volumes and not pay the capacity penalty for data that does not demand the additional protection.

Contact us today at: 888-734-7726 or sales@RELDATA.com to learn more about the RELDATA 9400 Storage System and our innovative architecture.