RELDATA

FEATURE BRIEF: Maximizing Storage Efficiency with Thin Provisioning

To maximize disk space usage and lower TCO, the RELDATA 9400 Storage System includes fully automated thin provisioning, freeing you from the hassles of storage management while delivering outstanding storage utilization rates.

Introducing Thin Provisioning

Under-utilization of storage capacity is a constant concern for IT administrators, driving up both capital and maintenance costs for storage. Thin provisioning allows administrators to optimize the utilization of available storage resources by providing on-demand allocation of blocks of data versus the traditional method of reserving capacity. With thin provisioning, storage utilization can be driven up from the typical 40-70 percent utilization rate towards 100 percent, with very little administrative overhead. You purchase less storage capacity up front, defer storage capacity upgrades, and save the operating costs such as power, cooling, and floorspace requirements associated with maintaining more storage capacity than necessary due to low utilization rates.

RELDATA's 9400 Storage System offers a very efficient thin provisioning capability featuring no reservations, no administrative overhead, and absolutely no performance impact. With the 9400, all block devices (LUNs) are automatically thin provisioned; the size of the LUN entered during setup acts as a quota, but the 9400 does not pre-allocate or reserve space based on the size of the LUN. As the administrator, you define a fixed size volume to a server, but storage blocks will be allocated to the volume only as they are written. All writes – even modifications to an existing block – are written as a new block. Blocks that contain data that is no longer needed are then de-allocated and added back to the free pool. The size of a LUN can be increased or decreased (assuming there is free space) at any time by any amount, without any effect on the thin provisioning mechanism.

Figure 1 shows an example of thin provisioning with the 9400. A 200GB iSCSI LUN is created and mounted to a Windows server as the G: drive. The Windows server "sees" a 200GB drive, but in fact no blocks are reserved on the 9400. If a 480KB file is then written to the G: drive, four 128KB blocks will be allocated to that volume from the free pool, and the data will be written to them. The volume only grows on the array based on writes, and any unused space is available for other volumes or NAS shares.

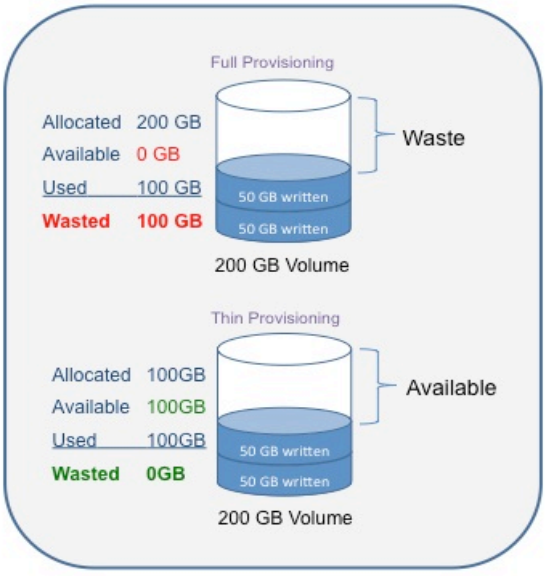

Figure 1. Comparing Full and Thin Provisioned Volume Allocation

Every iSCSI volume and fibre channel LUN on the 9400 is thin-provisioned, as well as all snapshots and clones. NAS shares are inherently thin provisioned based on their quota. No blocks are ever reserved or pre-allocated. If the 9400 reports that 10TB are used, all 10TB have been written to at least once. The 9400's thin provisioning feature is extremely efficient, and is in fact much more efficient than the thin provisioning implementations in most other systems. For example, there is no performance overhead to thin provisioning on the 9400. Each write to the array allocates a new block if needed. This allocation is done within the virtualization layer and effectively adds no

©2011 RELDATA. All Rights Reserved
1-888-734-7726 or 1-973-644-2770

latency to the I/O. Other systems may incur overhead for pre-allocation, which is one reason why thin provisioning is optional on many systems. There is also no administrative overhead to thin provisioning; no thresholds to set and manage, no snapshot reserves, no extents or stripe width settings. You simply create your LUNs and the system does the rest.

Cumulatively, you can define LUNs that present more space than what actually exists within the array. This is known as over provisioning and means that if all the capacity in all volumes and shares was written, the array would run out of physical storage. Over provisioning can be useful in situations where applications need very large volumes "just in case", but it is very unlikely that all of them will be filled up. An example would be a LUN containing database tables. You may want to define a 1TB volume to the database server so you don't need to worry about ever re-sizing it, even though the current database is only 300GB and growing slowly.

When the array is over provisioned, it is important to monitor the physical storage, especially in high-utilization scenarios. The 9400 Storage System does this monitoring for you, providing the necessary alerts at the macro and object level to ensure you are aware of any capacity problems with plenty of time to deal with them. If, as an administrator, you really don't ever want to worry about running out of physical space, simply do not over provision the array. In this scenario, you're completely safe -- no write to a volume can ever fail due to an "out of capacity" problem. The 9400 keeps you very informed on over provisioning, and will let you know where you stand every time you create or modify the size of a LUN or a share.

Thin Provisioning in a VMware Environment

VMware vSphere 4.0 has implemented virtual machine-level thin provisioning to address the earlier issues of artificially large virtual machine images that took up significant storage space. With past Hypervisor implementations, much of the virtual machine image was zero-filled, designed to accommodate crash files, diagnostics, etc. In this scenario, it was common for only a small proportion of the image to contain actual data. The VMware thin provisioning implementation addresses this issue, improving storage utilization dramatically for environments with a large number of virtual machines. However, in order to capitalize on this capability, the storage system must also thin provision the volumes storing the virtual machines. Without array-level thin provisioning, the storage system will reserve the volume capacity needed for the virtual machine images, causing you to forfeit the benefit of the VMware technology. The RELDATA 9400 Storage System's implementation of thin provisioning works in conjunction with VMware's feature to ensure optimal storage utilization in a VMware environment.

Deallocation of Blocks with the SCSI Unmap Command

One attribute of thin provisioning is counter-intuitive and can result in unused capacity within the array. The storage array currently has no way to know when a file is deleted from a file system, or what blocks may be freed up as a result. Therefore, once a block is written to, that space must remain as part of the volume, even if all the data has been deleted at the file system level. This is why the file system "% full" may not match the volume allocated size, and may in fact be much lower. This is an issue with all thin provisioning implementations. Capacity reclamation, or the ability to release blocks back to the pool upon file deletion, is an important new technology for maximizing capacity utilization. To work, this technology requires that the file system support the SCSI Unmap command; the file system must tell the storage array which SCSI addresses it no longer needs due to file deletions. A growing number of file systems support this command today, and support will be coming soon from additional vendors. RELDATA's thin provisioning technology supports the SCSI Unmap command, making the 9400 well-positioned to support this key technology as more and more file systems expand their support.

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

©2011 RELDATA. All Rights Reserved