

Virtual Provisioning of Drives for Windchill

2 Jun 2010 9:00 AM

0

Cost savings drives many IT departments to move databases to virtual provisioning. The cost savings come from better hardware utilization and from centralized administration. But, there are some situations in which such a move may result in a significant loss of Windchill performance. This post explains why virtual provisioning is not recommended for Windchill and what you can do to mitigate issues if you are required to use virtual provisioning.

In [earlier posts](#), I discussed logical drive allocation assuming that you use logical drives that are connected to physical drives. At this end of the enterprise spectrum is virtual provisioning of drives or storage pooling. This is one step or layer above traditional RAID and SANs. This is a layer because it sits over a RAID layer or over a SAN layer. But unlike SAN and RAID, it allows sharing of unused storage elsewhere. This is a basic assumption of virtual provisioning. This layer uses the term logical unit number (LUN), which is also used with SANs and RAID and can cause some confusion.

With the technologies of RAID and SANs, I will use the term traditional LUN. With virtualized drives, I will use the term thin LUN. Into this mix, I'll also use the term metaLUN, which is a LUN that sits on top of a hierarchy of other LUNs (of either type).

The specific guidance from virtual provisioning vendors is as follows:

- Performance critical: traditional LUN
- Tolerant of performance variation: thin LUN (It is more important to save space.)
- High throughput: metaLUN (i.e., striping across child LUNs)

In an MSDN blog post, SQL Server MVP Jonathan Kehayias commented on virtual provisioning in general:

[Virtual Provisioning] is not ideal for SQL Server. On any SAN you want dedicated disks and lots of little disks are better than less larger ones. We use an HP EVA at my company and you can measure the IO latency at times because of shared activity on the disks. I personally wouldn't recommend this configuration if you need/expect high performance from the IO subsystem for SQL Server. I am not saying that the technology is bad, but it's not ideal for SQL Server. You need to be able to dedicate your RAID groups and disk spindles to specific purposes for SQL. Data file Random IO should be on separate disks completely from Transaction Log Sequential IO, which should be separated from TempDB IO. You won't be able to ensure this occurs under the storage pool.
(<http://social.msdn.microsoft.com/Forums/en/sqldatabaseengine/thread/0d6b7ff0-3588-4140-8042-8f70bafc2186>)

Most thin LUN technologies have some form of a service manager (LUNSM) that identifies resource contention and performance and that provides information similar to that available from the `fn_virtualfilestats` function in SQL Server against a RAID, SAN, or physical drive. Logical drives are not tied to physical spindles but may float on an hourly or daily basis. Logical drives may be created from fragments across different spindles with sequential records being delivered from different drives.

The following table indicates some critical differences between a thin LUN and a traditional LUN.

Traditional LUN	Thin LUN
Absolute best performance.	Some performance enhancement if LUNSM is used to identify and manage resource contention within the LUN pool. Unfortunately, LUNSM does not automatically shift filegroups to better performing LUNs.
Most predictable performance.	
No concern with space efficiency.	Best space efficiency.
Most capital investment.	Lower capital costs and more energy savings.
Complex setup.	Easy setup.
SQLIO is recommended.	SQLIO is not applicable. (Physical media performance is not deterministic.)
New storage may require reconfiguration.	No host impact.
Often RAID 1 or RAID 5.	Typically RAID 5 or RAID 6.
	The thin LUN allows the easy shifting of a filegroup from one media (for example, cheap SATA SANs to striped RAID arrays).
Disk space is dedicated and deterministic.	Disk space is assigned when it is used from across arrays of disks.

To implement Windchill in a virtual provisioning environment, assign one thin LUN to each of the following:

- tempdb
- tempdb log
- wcAdmin (database) log
- INDX filegroup

- PRIMARY filegroup
- BLOBS filegroup
- WCAUDIT filegroup
- Backup drive

Then, use LUNSM to monitor resource contention and address that on a LUN-by-LUN basis. (See LUNSM for information about how to do this.) You may wish to preemptively move tempdb and tempdb log off the generic thin LUN to striped drives. These two files limit the maximum performance of the database.

Special Considerations If You Use a Thin LUN

If you use a thin LUN, keep these considerations in mind:

- Verify that instant file initialization is turned on. (It should always be turned on with Windchill and SQL Server databases.)
- Verify that write caching is turned off across each of the layers, down to the physical drive write cache.
- The SQL Server log file should use prewritten space for:
 - Optimum write performance. (You do not want the log to be slowed down by provisioning of more space.)
 - Guaranteeing that allocations are actually available.
- Do not use **sp_clean_db_free_space**.
- Thin pools should expand in increments of:
 - Five drives for RAID 5.
 - Eight drives for RAID 6.

Windchill is rarely database bottlenecked, especially with SQL Server 2005. If there is significant load on SQL Server or indications of performance problems, you should not use virtual provisioning. This is a high-end solution: with eight thin LUNs running RAID 6, you quickly count 64 hard drives as a starting point if you want complete thin LUN independence. Virtual provisioning reduces management complexities for a Windchill DBA. Complexities still exist, but they all are handled by the person who is managing the thin LUNs (i.e., layering management). Some key points to remember are as follows:

- There is potential for real and virtual cost savings with thin LUN, but it is not guaranteed.
- The literature gives many qualifications about when thin LUN should be used. Do not extrapolate beyond those qualifications. Performance is not a benefit.
- Tracking down performance problems due to hardware issues is very difficult. The physical drives that you used yesterday may not be the physical drives that you use today.

Additional Information

For more information, visit the following websites:

- Dell/EMC CX4 Virtual Provisioning Overview:
http://www.dell.com/downloads/global/products/pvaul/en/cx4_virtual_provisioning_wp.pdf
- EMC CLARiiON Virtual Provisioning:
<http://www.emc.com/collateral/hardware/white-papers/h5512-emc-clariion-virtual-provisioning-wp.pdf>
- EMC Symmetrix VMAX Virtual Provisioning Space Reclamation and Application Considerations:
<http://www.emc.com/collateral/software/white-papers/h6730-virtual-provisioning-space-reclamation-wp.pdf>
- HP StorageWorks SAN Virtualization Services Platform—Overview and Features:
<http://h18000.www1.hp.com/products/storage/software/sanvr/index.html>
- Webcast: Hyper-V & NetApp Storage, Demo Part III: Optimizing Storage
<http://virtualboytv.com/netapp4.aspx>

Ken Lassenen is part of the original team that created Dr. GUI of MSDN and specializes in new and resurrected commercial product architecture. He developed architecture for several Microsoft websites, including the original MSDN site and the current Microsoft Partner Network site. He's equally at home with SQL Server, XHTML, Section 508 accessibility standards, globalization, Security Content Automation Protocol (SCAP) security, C#, and ASP.NET server controls. When he is not having fun with technology, he enjoys taking lunch-break hikes in the North Cascades.

