



Storage Management First Steps - Improving Storage Performance

Prepared by: George Crump, Lead Analyst
Prepared on: 8/24/11

<http://www.storage-switzerland.com>

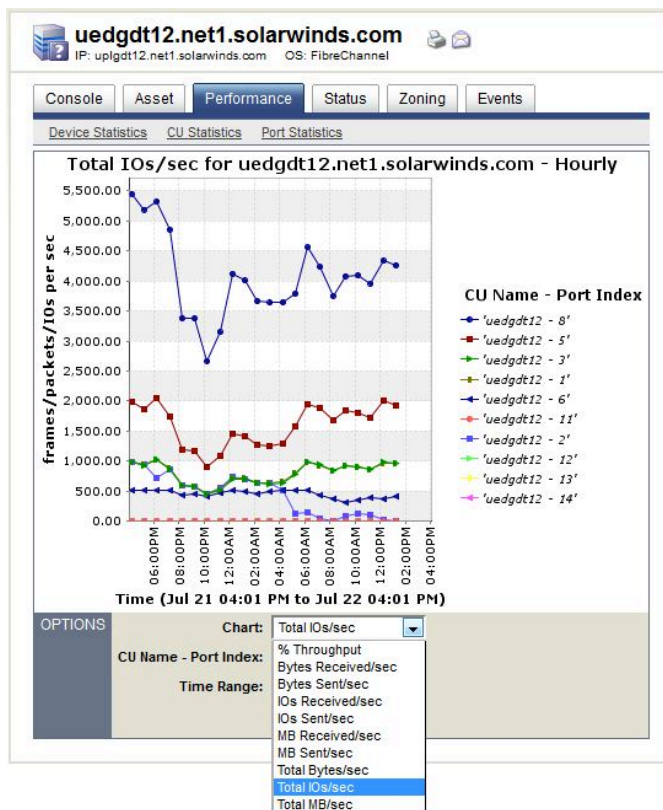
Copyright © 2011 Storage Switzerland, Inc. - All rights reserved

Storage Switzerland's article "[Storage Management First Steps - Inventory and Capacity Control](#)" covered how to quickly gain control of the storage environment through a better understanding of what is present (inventory) and how to reduce costs by increasing utilization (capacity control). The next step in storage management is to maximize performance through a better understanding of the storage infrastructure and how to maximize I/O utilization.

A Performance Inventory

While performance is not a static thing that can be inventoried like a storage array or a disk drive, it *is* something that can be quantified. The first step in performance tuning is to inventory what the potential performance capabilities are and the hosts or applications that are the consumers of that performance. This can be done by using a storage management tool like [SolarWinds Storage Manager, Powered by Profiler](#), to understand who those top consumers are. This will show the performance 'pressure points' in the storage infrastructure and help guide where to look for a performance problem.

Switch Performance



The first place to look is the storage switch. The speed of the storage system is irrelevant if the switch is performing at its maximum potential but still bottlenecking access to storage. The switch is also often an area where a little rebalancing of traffic can go a long way toward increasing performance. In fact, many extra switches and ports are purchased because storage administrators don't have visibility into the traffic going through their existing switch ports, making the actual utilization of total switch bandwidth remarkably low.

A storage management tool will be able to show the overall load on the switch as well as provide a drill down to see busy ports or inter-switch links (ISLs).

The first thing to look for, especially in a high switch-count infrastructure, is to examine the ISL traffic, the traffic flowing between each switch. Make sure that the speed of the ISL is set correctly. It's not uncommon to find switches where the ISL speed is slower than the speed of the switch itself, something that can happen as switches are gradually upgraded. For example, if a fabric consisted of just 2Gb FC switches and then one 4Gb FC switch was added, many administrators would hard set the ISL to 2Gb on each side. When the other 2Gb FC switch is finally upgraded the setting is often not changed to its full potential. Next look for ISLs that are being overused and see if an additional ISL or even an additional switch might better help traffic flow.

The next step is to drill down to the specific ports on the switch to see which are experiencing the most traffic. The value that a storage management tool brings is the ability to see this traffic across switches so that a comparison can be made. In many cases, two switches may be carrying the bulk of the traffic while two others are barely being used. Simply moving a server or two to another switch may result in better performance.

Once the fabric has been better balanced, then the amount of traffic going through certain ports can be examined to see how much of the available bandwidth is being consumed. If it's consistently hitting the threshold speed of the port/switch then an investment in a faster switch may be warranted. A storage management tool can ensure this step of purchasing more hardware is a last resort, after a complete optimization of the existing investment has occurred.

Storage System Performance

Once it is confirmed that the storage fabric is no longer a bottleneck, the next place to examine is the storage system itself. In these storage environments, multiple physical servers typically share the same physical disks. Understanding how many servers and how performance demanding those servers are is critical to maximizing performance.

Similar to the previous scenario with a busy switch, if several active servers are sharing the same physical drives, then performance is going to suffer. Creating separation and moving some physical systems to different drives can greatly improve performance. As was the case with switch performance, this

LUN Performance				
Show Description EMC CLARiON: LUN Performance Time Zone: US Central Standard Time (Chicago) Start: 07/15/2011 15:06 End: 07/22/2011 15:06 Data: Weekly Printer Friendly				
Row	Resource Name	LUN ID	▼ Total I/Os/sec	MB Read/sec
1	dhztbs12.net1.solarwinds.com	cgy-s102	453.38	0.00
2	dhztbs12.net1.solarwinds.com	cgy-s121	435.59	18.20
3	dhztbs12.net1.solarwinds.com	cgyexcp31-J-1TB	433.42	4.03
4	dhztbs12.net1.solarwinds.com	cgyexcp31-H-1TB	327.16	9.20
5	dhztbs12.net1.solarwinds.com	cgy-s104	266.39	0.00
6	dhztbs12.net1.solarwinds.com	cgy-s108	226.13	3.16
7	dhztbs12.net1.solarwinds.com	cgyexcp31-F-1TB	180.31	1.89
8	dhztbs12.net1.solarwinds.com	cgy-s106	161.91	0.00
9	dhztbs12.net1.solarwinds.com	cgy-s109	144.86	0.00
10	dhztbs12.net1.solarwinds.com	cgy-s103	125.77	1.81
11	dhztbs12.net1.solarwinds.com	cgy-s187-sata	93.09	0.55
12	dhztbs12.net1.solarwinds.com	cgy-s171-RAID10	92.25	0.00
13	dhztbs12.net1.solarwinds.com	cgy-s107	90.20	1.75
14	dhztbs12.net1.solarwinds.com	cgy-s186-sata	83.01	0.52
15	dhztbs12.net1.solarwinds.com	cgy-s151	81.22	0.00

can be done without buying additional storage hardware. Simply rebalancing workloads so that the performance demanding servers are spread out across the available disks will reduce head seek.

Finally, once the storage system has been better balanced, it's time to examine the performance of the RAID groups and LUNs to see if the amount of IOPS being generated is near the limit that the system can produce. Again, the storage management tool will report this type of data, sort for easy discovery as well as report on queue depths and response time/latency. With this data in hand, the decision can be made to investigate adding more or faster drives to bring queue depth down or to add solid state disk technology to reduce queue depth and latency at the same time.

This capability is critical, especially with virtualized storage systems like those offered by HP 3PAR, Dell Compellent and others, since these systems, as a default, spread data across all available drives. The assumption that adding more drives will increase performance may be inaccurate, since many drives are already utilized in the system. If the application is not experiencing long wait times due to its own queue depth, it's doubtful that adding drives will increase performance. Operating system tools will have difficulty providing this analysis.

Virtualized Infrastructure Performance

One of the biggest consumers of storage I/O is the virtual infrastructure. In a virtualized server environment, many relatively low-demanding hosts are all virtualized to create a few *very* demanding and *very* random I/O consumers. This is where products like SolarWinds Storage Manager bring significant value.

The first challenge in performance tuning a virtual infrastructure is being able to see "inside" of the host to understand which virtual machines are the large storage I/O consumers. Unlike the non-virtualized examination, each physical host can generate dozens of different I/O profiles. To further complicate the performance diagnosis, each disk area assigned to the physical host is most often shared with other physical hosts to enable functions like virtual machine migration. A poor performing virtual machine may be suffering because another VM is stealing those resources and starving the problem VM, not because of the quality of the resources available.

Products like SolarWinds Storage Manager can report on which virtual machines within each host are the most demanding. It may make sense to create multiple data stores on different physical disks so that busy VMs are not creating as much hard disk thrashing per physical spindle. To accomplish this requires that the software be able to perform an end-to-end map of the virtual infrastructure and lay that over the map of the storage infrastructure.

Sales Console

ConsoleStoragePerformanceAsset

Logical MappingVirtual DataVMOK FilesROM Mapping

Logical Mapping

Show Description

VMware VM Groups: Logical Mapping

Printer Friendly

Row	Virtual Machine Name	VMOK File Name	ESX Host Name	Cluster Name	Datastore Name	Extent	Total GB	% Used	Initiator	Protocol	Host LUN ID	Target Name	Target Size GB	Remote Device Name	Time
1	houcaesp01	iwddtfu12.vmdk	iwdfyq13.net1.solarwinds.com		iwds-001	vmhba1:0:1	1,207.00	41.97	32:11:11:2C:43:26:BA:69	FC	1	V2	1,207.17	hvcasr01.teckcominco.loc	07/22/2011 06:09:55
2	houcupp01	iwddvq12.vmdk	iwdfyq13.net1.solarwinds.com		iwds-001	vmhba1:0:1	1,207.00	41.97	32:11:11:2C:43:26:BA:69	FC	1	V2	1,207.17	hvcasr01.teckcominco.loc	07/22/2011 06:09:55
3	houcbq01	iwddcjq12_1.vmdk	iwdfyq14.net1.solarwinds.com		iwds-001	vmhba1:0:1	1,207.00	41.97	32:11:11:2C:43:26:9F:60	FC	1	V2	1,207.17	hvcasr01.teckcominco.loc	07/22/2011 06:14:55
4	houapp01	iwddbqq12_1.vmdk	iwdfyq14.net1.solarwinds.com		iwds-001	vmhba1:0:1	1,207.00	41.97	32:11:11:2C:43:26:9F:60	FC	1	V2	1,207.17	hvcasr01.teckcominco.loc	07/22/2011 06:14:55
5	houapp01	iwddbqq12.vmdk	iwdfyq14.net1.solarwinds.com		iwds-001	vmhba1:0:1	1,207.00	41.97	32:11:11:2C:43:26:9F:60	FC	1	V2	1,207.17	hvcasr01.teckcominco.loc	07/22/2011 06:14:55
6	houapp05	iwddbqq15.vmdk	iwdfyq13.net1.solarwinds.com		iwds-001	vmhba1:0:1	1,207.00	41.97	32:11:11:2C:43:26:BA:69	FC	1	V2	1,207.17	hvcasr01.teckcominco.loc	07/22/2011 06:09:55
7	houfp01	iwddgm12.vmdk	iwdfyq13.net1.solarwinds.com		iwds-0203	vmhba1:0:3	2,230.75	58.52	32:11:11:2C:43:26:BA:69	FC	3	V4	1,207.17	hvcasr01.teckcominco.loc	07/22/2011 06:09:55
8	houfp01	iwddgm12.vmdk	iwdfyq13.net1.solarwinds.com		iwds-0203	vmhba1:0:2	2,230.75	58.52	32:11:11:2C:43:26:BA:69	FC	2	V3	1,024.00	hvcasr01.teckcominco.loc	07/22/2011 06:09:55
9	houfp02	iwddgm13.vmdk	iwdfyq14.net1.solarwinds.com		iwds-0203	vmhba1:0:3	2,230.75	58.52	32:11:11:2C:43:26:9F:60	FC	3	V4	1,207.17	hvcasr01.teckcominco.loc	07/22/2011 06:14:56
10	houfp02	iwddgm13.vmdk	iwdfyq14.net1.solarwinds.com		iwds-0203	vmhba1:0:2	2,230.75	58.52	32:11:11:2C:43:26:9F:60	FC	2	V3	1,024.00	hvcasr01.teckcominco.loc	07/22/2011 06:14:56
11	houapp03	iwddbqq14.vmdk	iwdfyq15.net1.solarwinds.com		iwds-001_(1)	vmhba1:0:1	1,673.25	63.46	32:11:11:2C:43:11:10:D9	FC	1	N/A	1,673.38	Not Monitored	07/22/2011 05:39:55
12	chipsa27-a	dhztq138-a.vmdk	dhztq13.net1.solarwinds.com		dhz-s184-sata	naa.60060160b3d01a001250a8a0be15df11	1,406.75	3.46	DH2FTYQ13_vmhba3	FC	28	dhz-s184-sata	1,406.89	cgysar01.teckcominco.loc	07/22/2011 04:14:55
13	chipsa27-a	dhztq138-a_2.vmdk	dhztq13.net1.solarwinds.com		dhz-s184-sata	naa.60060160b3d01a001250a8a0be15df11	1,406.75	3.46	DH2FTYQ13_vmhba3	FC	28	dhz-s184-sata	1,406.89	cgysar01.teckcominco.loc	07/22/2011 04:14:55
14	chipsa27-a	dhztq138-a_1.vmdk	dhztq13.net1.solarwinds.com		dhz-s184-sata	naa.60060160b3d01a001250a8a0be15df11	1,406.75	3.46	DH2FTYQ13_vmhba3	FC	28	dhz-s184-sata	1,406.89	cgysar01.teckcominco.loc	07/22/2011 04:14:55
15	adm-sas-ds3	bs-snc-dc3.vmdk	tdofyq15.net1.solarwinds.com		tds-s103	naa.600221900080b9f000070c34bd1ceff	1,046.50	69.77	IQN.1998-01.COM.VMWARE:TODFTYQ15-34C54F19	ISCSI	3	N/A	1,046.71	Not Monitored	07/22/2011 02:59:54
16	adm-sas-ds3	bs-snc-dc3_1.vmdk	tdofyq15.net1.solarwinds.com		tds-s103	naa.600221900080b9f000070c34bd1ceff	1,046.50	69.77	IQN.1998-01.COM.VMWARE:TODFTYQ15-34C54F19	ISCSI	3	N/A	1,046.71	Not Monitored	07/22/2011 02:59:54
17	adm-sco-aps	bs-san-ap1.vmdk	tdofyq15.net1.solarwinds.com		tds-s103	naa.600221900080b9f000070c34bd1ceff	1,046.50	69.77	IQN.1998-01.COM.VMWARE:TODFTYQ15-34C54F19	ISCSI	3	N/A	1,046.71	Not Monitored	07/22/2011 02:59:54
18	adm-sco-aps	bs-san-ap1_1.vmdk	tdofyq15.net1.solarwinds.com		tds-s103	naa.600221900080b9f000070c34bd1ceff	1,046.50	69.77	IQN.1998-01.COM.VMWARE:TODFTYQ15-34C54F19	ISCSI	3	N/A	1,046.71	Not Monitored	07/22/2011 02:59:54
Sum		18 Values	5 Unique		5 Unique			47.69 (Avg)		2 Unique			3 Unique		6 Unique

Finally, it is important to know how storage performance is impacted as virtual machines are migrated to other hosts. This requires a historical tracking of virtual machine movement between hosts, something that SolarWinds calls “Time Travel”. Using SolarWinds Virtualization Manager, this feature allows the administrator to develop best practices on where and when to move a virtual machine.

Trending

Similar to how capacity utilization needs to be tracked over time, so does performance. A performance problem will not always happen when the administrator is staring at the console, nor do all spikes occur at the same time. Performance utilization needs to be captured so that it can be replayed to pinpoint when a performance spike occurred. This allows the balancing of the environment so that performance spikes can be distributed both across the infrastructure and throughout the day.

Smart Upgrades

The above performance examination is designed to expose opportunities for improvement in the current infrastructure without buying additional hardware. At some point capacity must be purchased however, as the environment can only be optimized so much. A storage management tool adds value during the upgrade process as well. Armed with an historical profile, the IT manager can recommend the capacity and performance upgrade that will deliver the greatest increase for the least amount of money. More importantly, the storage management tool delivers the confidence of knowing that the upgrade will actually improve performance.

Summary

Storage performance tuning is often looked at as a 'black art', something to be performed by storage 'wizards' armed with decades of experience. A storage management application can level the playing field and give the busy IT generalist the information they need to make significant improvements in their storage design. However, even the storage specialist will benefit from these tools, which can help them make quicker, more effective and less reactive decisions.

In today's virtualized environment where servers, networks and storage are all shared, manual tools simply can't keep pace and deliver the required performance diagnosis expertise. Tools are needed to extract maximum performance from the storage investment and to keep up with the dynamic data center.