SOLARWINDS WHITEPAPER

How Do I Know My SQL & Virtual Environments are ready for SSD?

solarwinds

Follow SolarWinds:

Most IT professionals will agree that Solid State Disks (SSD) are the new high performance alternative to existing physical disks. While SSD's offer significant improvements to performance, they come with a premium price tag. IT pros must ask themselves, is the performance gain worth the money spent if I added SSD to my environment? The problem is ascertaining the "SSD worthiness" of an environment because SSD's can be highly dependent on specific workloads. Another reason is the operating system tools aren't quite enough to manage storage performance because all the servers (virtual and physical) are accessing the same storage at the same time.

The key takeaway for this whitepaper is that operating system tools don't come anywhere near standalone performance monitoring systems. The reason is because servers, both virtual and physical, are trying to access the same storage at the same time. On that note, let's begin by looking at issues you may commonly face with your database and virtual server and the benefits of SSDs on your SQL Server® and virtual environments.

Importance of SQL Server

SQL Server is one of many critical applications that must be watched closely. While SQL Server databases no longer represent the largest capacity requirement in the datacenter, it should be noted that these databases can still represent applications with high performance requirements. Therefore, IT planners often place SQL data on high performance storage systems that are capable of supporting hundreds of high-speed hard drives. Not surprisingly, data centers are now looking at flash based storage to better address the performance demands of SQL Server based applications.

Since flash based storage comes with a hefty price tag, IT pros try to use it as sparingly as possible. For this very reason, organizations that use SQL Server can take full advantage of SSDs by determining which component would benefit most by being on SSD.

It's More Than Performance – Scale Reduces Cost

It's also important to understand that flash enables more than just great performance, it also enables scale. Thanks to SSD, a single database instance and/or server is able to achieve much higher user count and application scale per instance than it could with hard disk based technology. While storage I/O isn't the only SQL performance culprit, it's one of the most common. If storage I/O is the problem, leveraging SSD can potentially save the cost of having to buy additional server capabilities like RAM and processing power. It also eliminates the complexity of database "sharding" or clustering.

Sharding is the process of dividing up sections of a large database to execute on different servers with different storage I/O channels. SQL Server clusters introduce the complexity of cluster management and clustered file systems.

How Do You Know Your SQL Environment Is Ready For SSD?

The first item to look for when determining the SQL Server storage environment's SSD readiness is CPU utilization of the physical server. For SSD to provide any performance improvement, CPU utilization needs to be relatively low; typically less than 30%. Essentially, IT pros want the CPU to indicate that it's waiting on something, most likely the storage system. High CPU utilization is an indicator of poor code or simply underpowered processing abilities. In either case, these items need to be fixed prior to considering SSD.

It's important to capture CPU utilization statistics prior, during, and after a performance problem becomes apparent to users, in order to establish a meaningful baseline. Operating system tools are ill equipped to capture this information since they have to be manually executed. It’s important for the analysis of resource utilization to be continuous. Tools like those from SolarWinds® allow for CPU activity to be captured continuously and then “played back” to see what the utilization was during the performance problem. If that CPU activity was low while the performance problem was occurring, then it’s safe to assume that storage performance is the likely culprit.

Storage Statistics

Once it has been confirmed that I/O performance is the problem, IT pros can focus on storage. The primary statistic to look for is the queue depths of the storage system. Queue depth, as it relates to storage, is the number of pending I/O operations. Typically these are being queued up by the storage controller as the hard drives process all the I/O requests sitting in the queue.

Before SSD was available as a performance enhancing option, queue depth was typically reduced by adding more disk drives to the volume that the database was writing to. The more drives in the volume, the more I/Os that can be processed at any given time. In large database environments, this can lead to database volumes with hundreds of hard drives.

By comparison, a single SSD can often reduce to a zero queue depth that may have required dozens of hard drives. This makes SSD not only a better performance option, it also makes it less expensive. A single SSD can outperform an array of hard drives while using substantially less power and space.

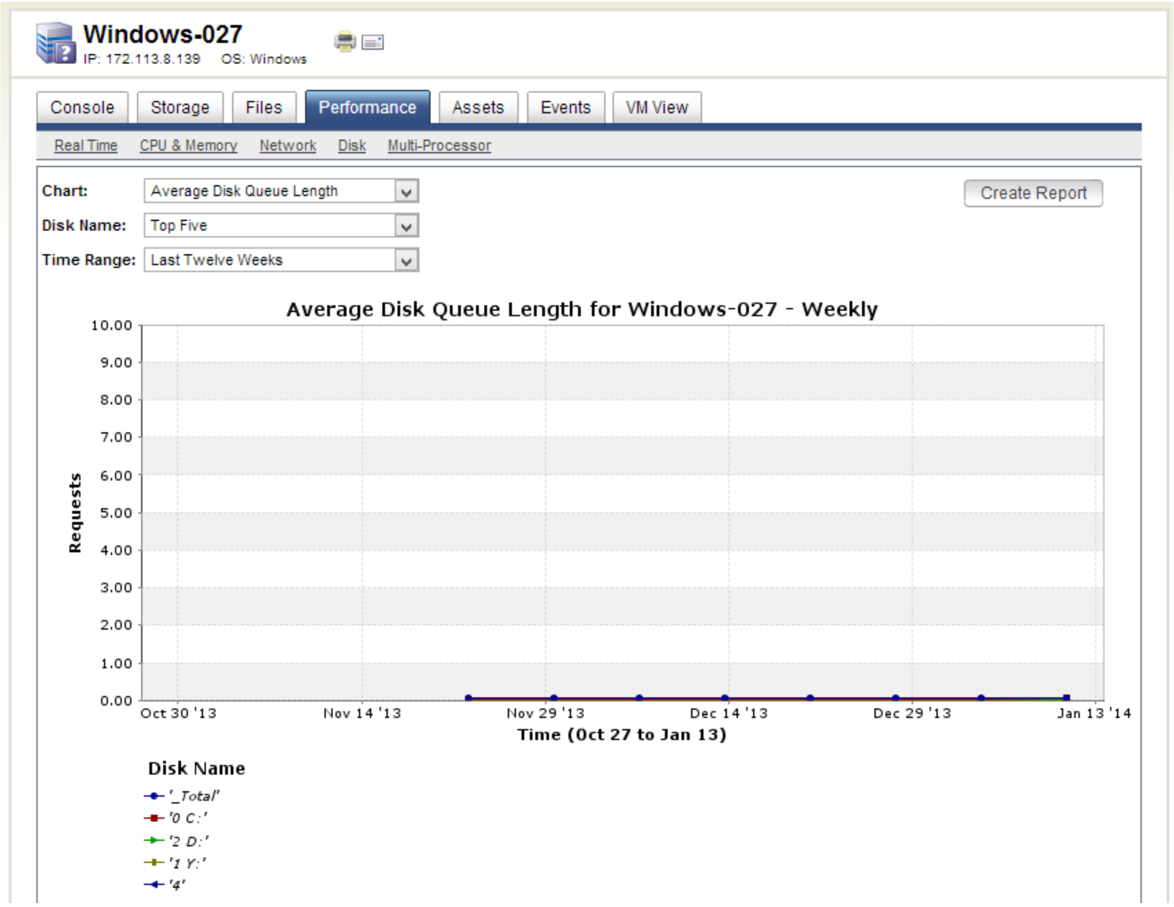

Windows-027

IP: 172.113.8.139 OS: Windows

Chart: Average Disk Queue Length

Disk Name: Top Five

Time Range: Last Twelve Weeks

Like CPU utilization, it's important to understand the queue depth of the database's volume prior to, during, and after the performance problem occurs. Once again, this requires that the volume be under constant monitoring with tools like those provided by SolarWinds. If during the performance problem, queue depth spikes very high, then adding drives or SSD will certainly help. In short, the higher the queue depth, the more appropriate it is to add SSD to the environment.

Getting Granular With Databases

The advantage of being able to focus specifically on SQL Server database performance, instead of the performance of the server as a whole, is that the performance variables can be more targeted. This "targeting" requires being able to "see inside" the database, so that specific files can be examined for performance problems. There are four main categories of files in a SQL environment: User Databases, System Databases, TempDB files, and log files. Each of these types should be examined for SSD worthiness.

The most obvious solution is to put the entire database environment on a SSD array. However, this can be impractical because the cost of an SSD that size might exceed the budget for the project. It also may be unnecessary because the bulk of the SQL environment's I/O activity is going to occur in the transaction logs. All new data and modified data first go to these logs. The logs are then written to the actual data file via a lazy write method, basically the write occurs when system activity is low. The result is that SQL doesn't typically try to write to both the logs and the data files at the same time.

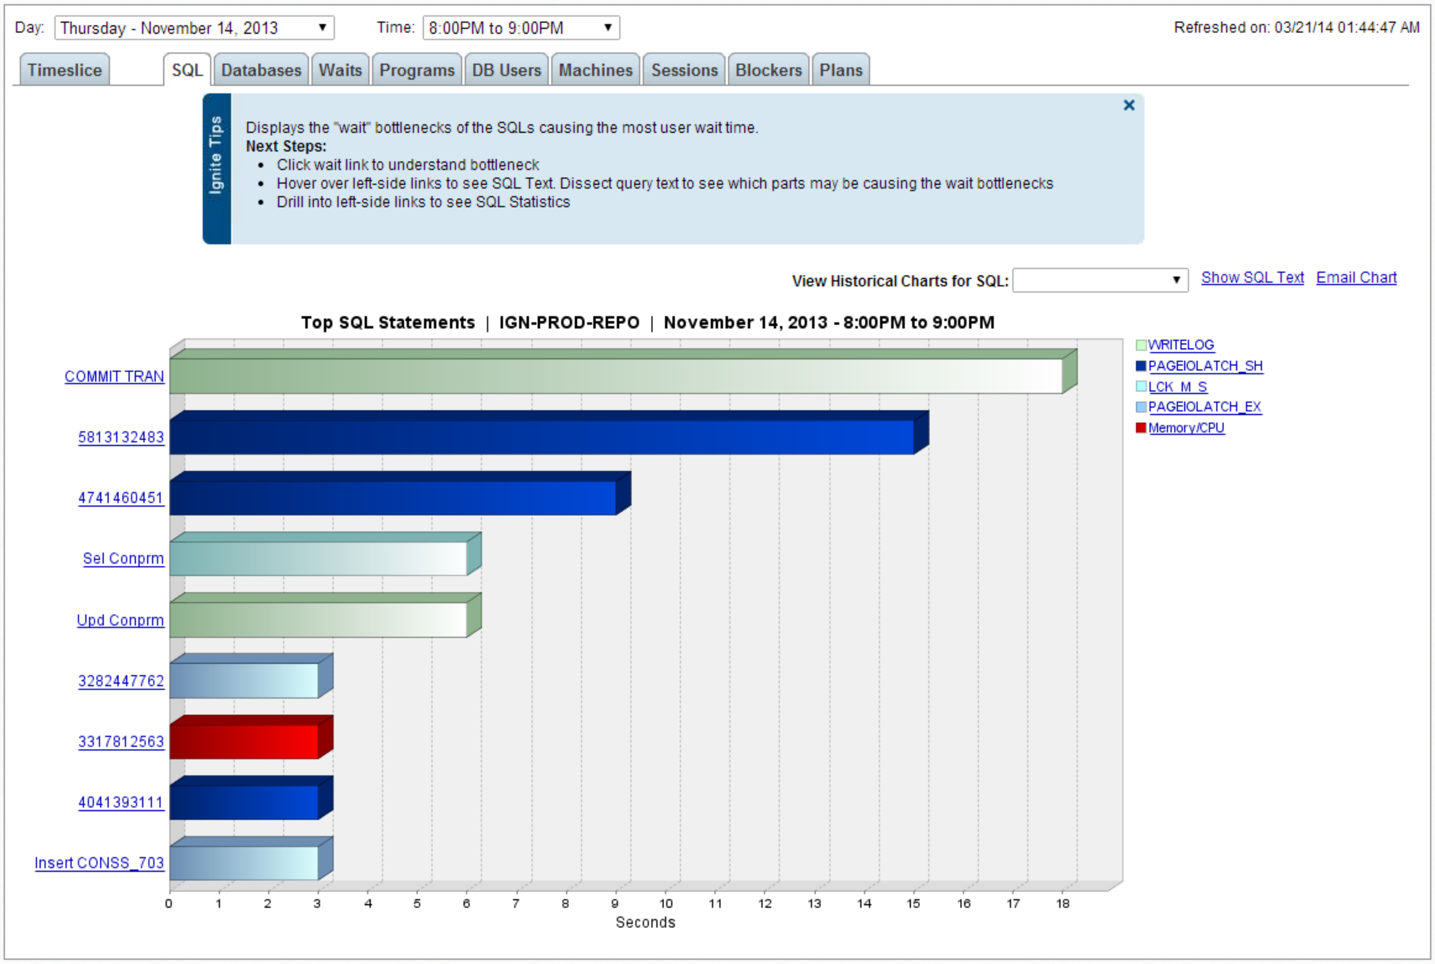

Day: Thursday - November 14, 2013

Time: 8:00PM to 9:00PM

Refreshed on: 03/21/14 01:44:47 AM

Displays the "wait" bottlenecks of the SQLs causing the most user wait time.
Next Steps:

- Click wait link to understand bottleneck
- Hover over left-side links to see SQL Text. Dissect query text to see which parts may be causing the wait bottlenecks
- Drill into left-side links to see SQL Statistics

View Historical Charts for SQL:

For many environments, SQL transaction performance can be substantially improved by just moving log files to SSD. The downside of flash technology, however, is that it has a limited number of times it can be written to. That number varies between Single-Level Cell (SLC) and Multi-Level Cell (MLC), with SLC allowing for the largest number of writes because of its increased durability. These logs should be monitored prior to SSD purchase to see how much data is written to them in a day. If it's high, SLC technology should be considered over the less expensive MLC, so it can provide uninterrupted performance acceleration to the environment.

Queries of database information typically count on index files. In the same way that transaction files are moved to SSD, index files can be moved to SSD to improve query results. But even though they are a subset of the actual data, indexes can be very large and it may be too expensive to place them on SLC based SSDs. Indexes however are much more read heavy and are more suitable for MLC based technology.

The key is knowing which of these files to place on SSD. Again with the appropriate tools, these files can be constantly monitored to see which ones are the most active during a performance problem. In most cases, the problem files are a fraction of the overall size of the environment. Some database environments can see significant performance improvement by placing a mirrored set of SSDs in the server and then storing all the user and system databases on a shared storage system. Moving these files to specific locations is also relatively easy to do within SQL configuration controls. A storage admin working with the DBA can place the larger, read-heavy index on a more cost effective per GB SSD tier using MLC flash and the smaller write heavy transaction logs on a smaller, but more write capable SLC tier.

A detailed wait type analysis (like the one provided in SolarWinds Database Performance Analyzer, formally Confio Ignite) provides the ability to see the impact drive configuration changes have on database response time. This gives IT professionals the ability to compare and contrast different storage configurations to strike a good balance between performance and cost.

Databases of almost any size should be a top consideration for SSD integration. Due to the granularity of the database structure, small components of the database can be moved to SSD, reducing the overall SSD investment. With proper monitoring tools, the correct components can be more easily identified. In the end, a surprisingly small amount of SSD can provide a cost effective, more scalable, higher performing database for the entire environment.

How Do You Know Your Virtual Environment Is Ready For SSD?

When storage I/O is in a virtual infrastructure, it's usually referred to as the I/O blender. The I/O blender is created when hundreds of virtual machines across dozens of virtual hosts all access the same storage volumes at the same time. This has a constraint on every aspect of the storage infrastructure, but especially on the hard disk system as it has to thrash back and forth to respond to all these requests.

The default answer for the I/O Blender problem has been to include flash based solid state disk (SSD) technology into your virtual infrastructure. But how do you determine if your virtual environment really needs SSD? And, if it does, where and how should you implement that investment? SSD technology is viewed by many as a performance machete that eliminates many storage I/O performance problems. However, it's an expensive technology and your budget may require that you take a more surgical approach; leveraging the technology more like a performance scalpel.

If your budget requires a more precise approach, how do you gather the information to make the right decision? In the past, pre-virtualized data center IT planners could possibly get by with standard operating system utilities like PerfMon to capture the data they needed to identify performance problems. But in the virtualized world, where almost every host component is shared with at least a dozen other virtual machines (VMs) and every storage component is shared across dozens of physical hosts, identifying performance bottlenecks is much more difficult.

What Do You Need To Track?

When deciding if an SSD makes sense for a non-virtualized application, the information to capture must follow a linear path of CPU utilization, network utilization, and storage performance metrics, like queue depth and latency. This information can be captured manually and is typically measured at specific points in time, not as a historical average. In the virtual context, all this information needs to be captured and then correlated against the same information from all the other VMs across all the physical hosts in the environment.

As a result, IT planners trying to determine the need for a SSD investment in a virtual environment will require the use of software tools to help capture and correlate that information. Manual spreadsheets are no longer adequate for conducting performance analysis.

CPU Utilization

The first step in any storage performance analysis is to determine if there truly is a storage performance issue. CPU utilization is generally a good indicator of storage performance issues. If CPU utilization is high, then in most cases this isn’t a storage performance issue, but either a lack of CPU resources or poor application programming.

In the virtual use case, this means capturing CPU utilization, both at the host and down to the VM level over a period of time. When a performance problem occurs, this performance history can be referenced to see if there was a spike in CPU utilization at the same moment the bottleneck took place. If the CPU utilization was under 50% during that moment, then more than likely there was a storage bottleneck causing the performance issue.

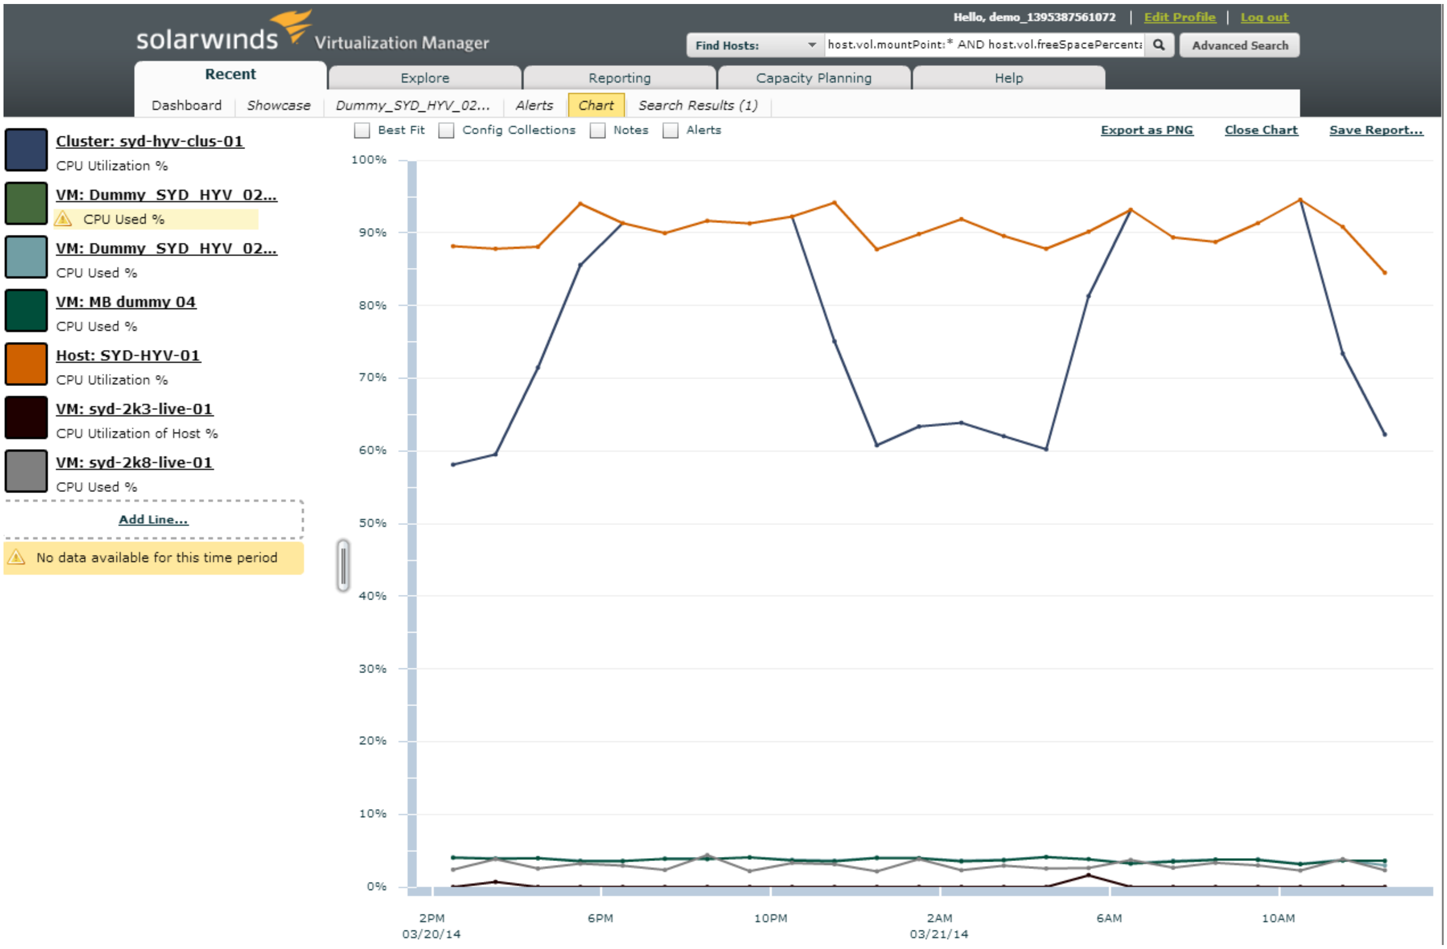

⚠ No data available for this time period

VMs CPU usage compared to the host they are running on

Storage Controller Utilization

The storage controller is the compute engine for the storage system. The utilization of its CPU can be easily measured, but like host CPU utilization, this parameter needs to be captured over a period of time so it can be compared to the moment when an application performance issue occurred. If the CPU inside the controller is running at greater than 60% utilization, then adding faster media like SSD likely will not help improve performance.

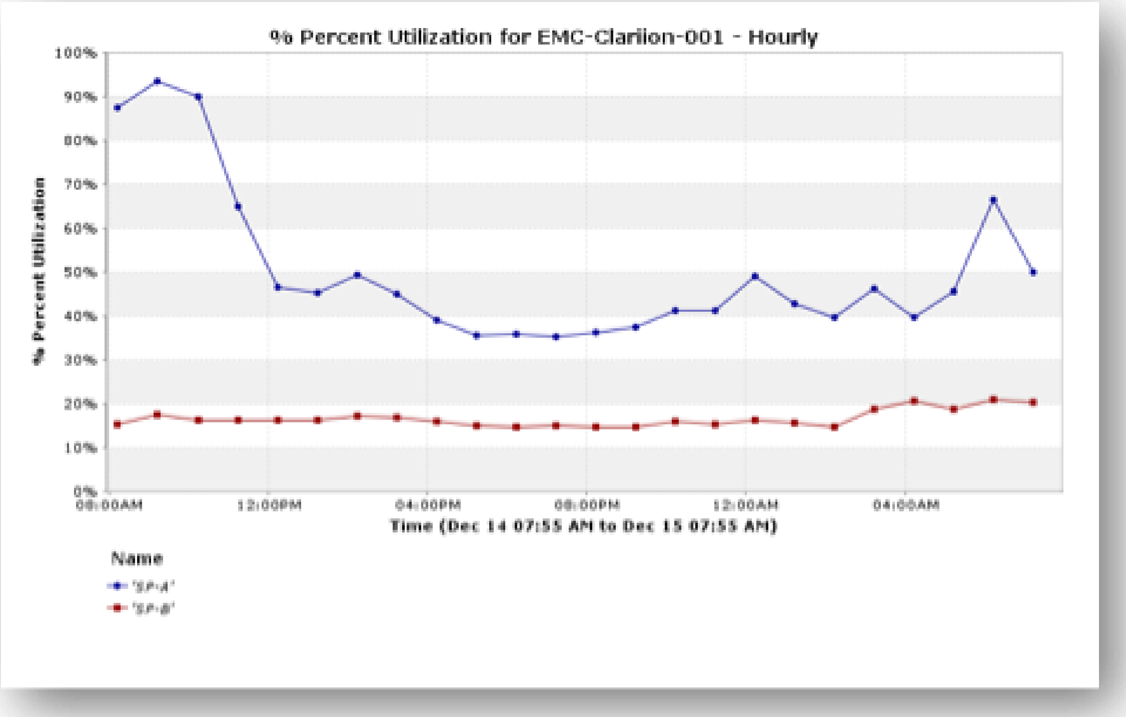

Controller Utilization over time

If adding SSD to the current system isn't effective, a decision must be made to either upgrade the existing system or add an additional storage system. Adding an additional system will require using something like Storage vMotion™ to migrate the performance-demanding VMs to that new system.

A cost effective alternative to a storage system upgrade could also be adding server based SSD to cache the performance demanding VMs locally to the physical server. Most server side caching software is read-only, so you'll need to determine if the performance demanding VMs in question have a heavier read I/O mix. If they do, a server side read cache with server side SSD can be a very cost effective resolution to performance problems, especially if there are only a handful of VMs experiencing problems.

If there are more than a handful of VMs that need the performance boost and if those VMs are scattered across several hosts, deploying an SSD appliance, as described above, specifically for these VMs may be a better approach.

Deciding between a server-based SSD or a network-based SSD requires understanding VM performance demands on an individualized VM basis. A tool that can provide a quick list of the top I/O consumers, per storage system regardless of the host it's on is essential. As well, the aforementioned need to determine read/write mix per a VM basis.

Storage Media Utilization

If the Host CPU and the Storage Controller CPU are not being fully utilized, then the next step is to determine if there's a device I/O problem. This is done by examining the IOPS and bandwidth utilization of the storage system; typically on a per LUN basis.

The first thing to look for is a few demanding VMs that are creating most of the I/O for a particular set of disks on the storage system. This requires a tool that can deliver an end-to-end view of the virtual environment, tracking a VM directly to the disks it's using. With this tool, certain VMs can be moved to new LUNs to obtain a substantial increase in response time. Moreover, this capability can save a substantial amount of money versus going to SSD.

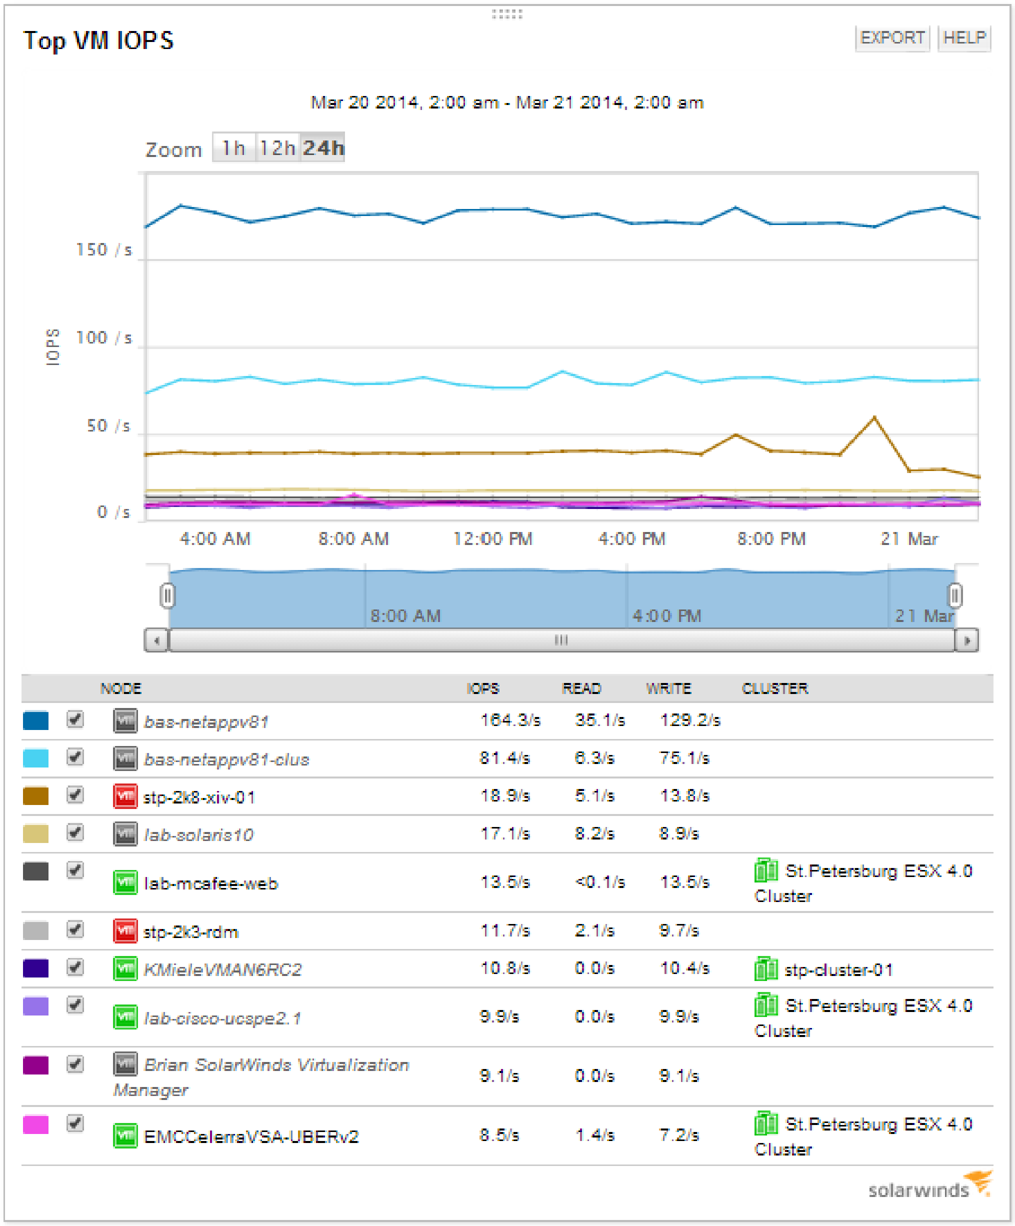

Top VM IOPS

		NODE	IOPS	READ	WRITE	CLUSTER
■	☑	VM *bas-netappv81*	164.3/s	35.1/s	129.2/s	
■	☑	VM *bas-netappv81-clus*	81.4/s	6.3/s	75.1/s	
■	☑	VM stp-2k8-xiv-01	18.9/s	5.1/s	13.8/s	
■	☑	VM *lab-solaris10*	17.1/s	8.2/s	8.9/s	
■	☑	VM lab-mcafee-web	13.5/s	<0.1/s	13.5/s	St.Petersburg ESX 4.0 Cluster
■	☑	VM stp-2k3-rdm	11.7/s	2.1/s	9.7/s	
■	☑	VM *KMieleVMAN6RC2*	10.8/s	0.0/s	10.4/s	stp-cluster-01
■	☑	VM *lab-cisco-ucspe2.1*	9.9/s	0.0/s	9.9/s	St.Petersburg ESX 4.0 Cluster
■	☑	VM *Brian SolarWinds Virtualization Manager*	9.1/s	0.0/s	9.1/s	
■	☑	VM EMCCelerraVSA-UBERv2	8.5/s	1.4/s	7.2/s	St.Petersburg ESX 4.0 Cluster

If the performance problem is more widespread than just a few VMs or if those few VMs need more performance than a dedicated disk can provide, then the performance capability of the media itself needs to be examined. Poor performance is often only evident for a few minutes of time, a narrow window to capture with manual tools.

Specifically, you want to find a period of time where either the IOPS, the bandwidth, or both essentially flat line from a performance perspective. This means that the storage system has hit a performance wall and assuming storage controller CPU utilization did not spike, adding flash SSD will help performance. The flash can be added with the confidence that it will justify the investment since it has been pre-determined that server and storage CPU resources are available to drive more data to a faster device.

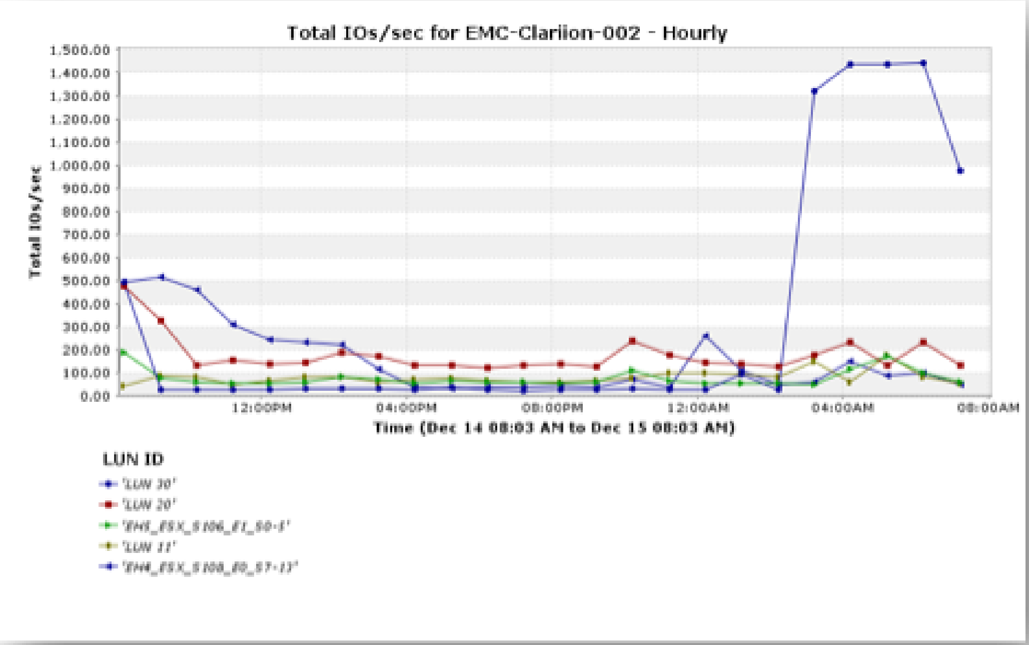

Flat Lining LUN

The ability to incrementally add flash to the existing infrastructure instead of resorting to purchasing a new system is an important justification for improving the monitoring process. This is because adding flash to a compatible storage system provides a performance boost to the entire environment. This process, assuming your vendor doesn't overcharge for a flash upgrade, should be less expensive than integrating a new SSD appliance and less complex than developing a multi-tiered server side SSD with shared storage infrastructure.

As previously mentioned, the fact that a performance bottleneck window may be as small as a few minutes or seconds explains why a solution that can provide an active analysis is important. Performance monitoring in a virtual environment essentially requires a TiVo like functionality where resource consumption can be "replayed" to find the point at which it occurred. The figure below shows how this can be achieved by looking at current and historical usage on a single LUN, including which VMs are driving the load and what other LUNs are sharing the disks.

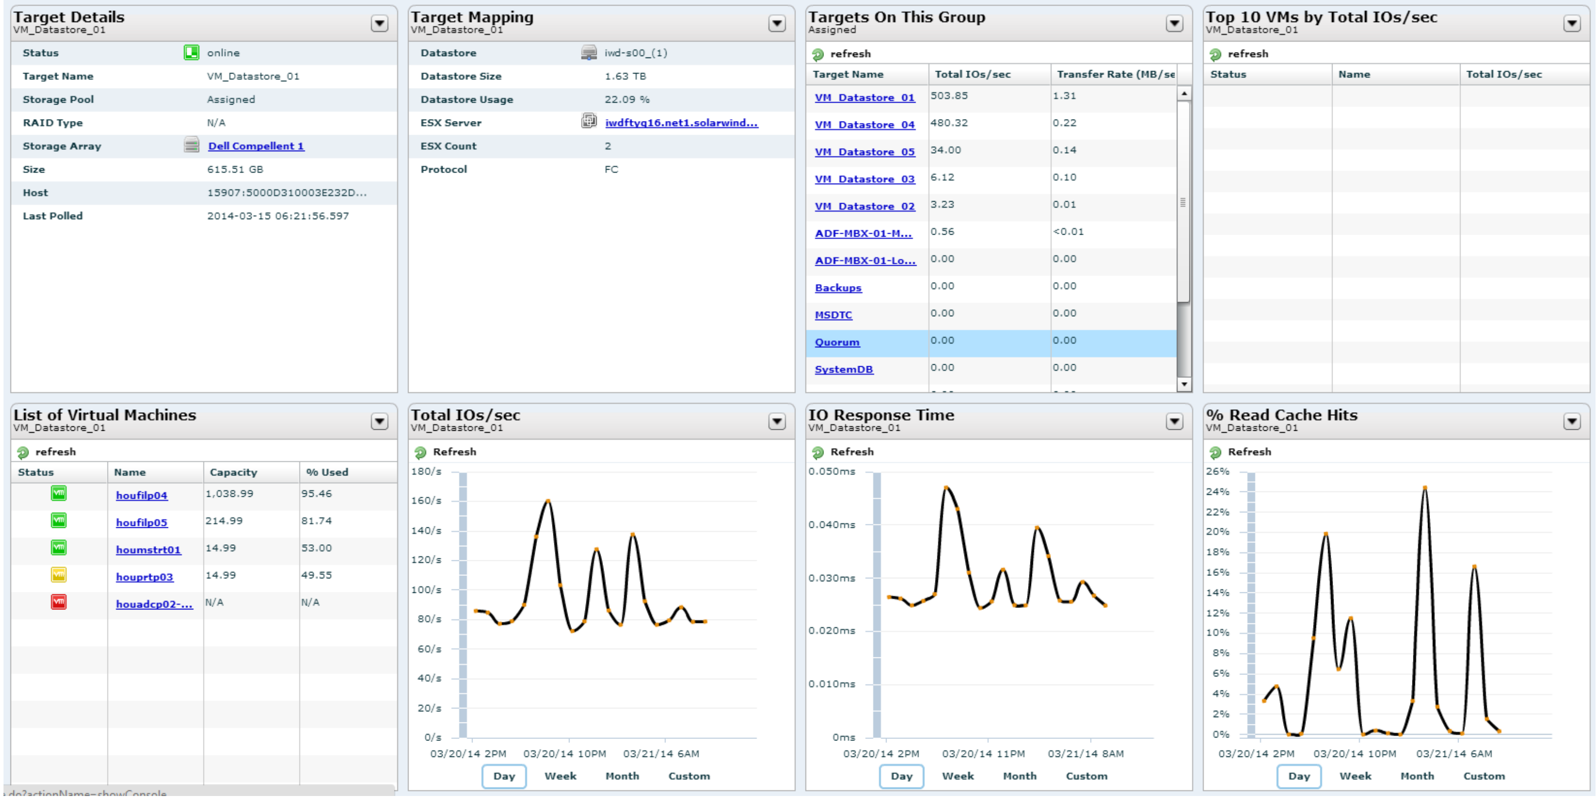

Target Details
VM_Datastore_01

Status	online
Target Name	VM_Datastore_01
Storage Pool	Assigned
RAID Type	N/A
Storage Array	Dell Compellent 1
Size	615.51 GB
Host	15907:5000D310003E232D...
Last Polled	2014-03-15 06:21:56.597

Target Mapping
VM_Datastore_01

Datastore	iwd-s00_(1)
Datastore Size	1.63 TB
Datastore Usage	22.09 %
ESX Server	iwdftyq16.net1.solarwind...
ESX Count	2
Protocol	FC

Targets On This Group
Assigned

refresh

Target Name	Total IOs/sec	Transfer Rate (MB/se
VM_Datastore_01	503.85	1.31
VM_Datastore_04	480.32	0.22
VM_Datastore_05	34.00	0.14
VM_Datastore_03	6.12	0.10
VM_Datastore_02	3.23	0.01
ADF-MBX-01-M...	0.56	<0.01
ADF-MBX-01-Lo...	0.00	0.00
Backups	0.00	0.00
MSDTC	0.00	0.00
Quorum	0.00	0.00
SystemDB	0.00	0.00

Top 10 VMs by Total IOs/sec
VM_Datastore_01

refresh

Status	Name	Total IOs/sec

List of Virtual Machines
VM_Datastore_01

refresh

Status	Name	Capacity	% Used
vm	houfilp04	1,038.99	95.46
vm	houfilp05	214.99	81.74
vm	houmstrt01	14.99	53.00
vm	houprtp03	14.99	49.55
vm	houadcp02-...	N/A	N/A

Total IOs/sec
VM_Datastore_01

Refresh

IO Response Time
VM_Datastore_01

Refresh

% Read Cache Hits
VM_Datastore_01

Refresh

Historical LUN Usage by VM

Storage Network Utilization

Once storage system performance attributes are determined and the host is ruled out as a performance culprit, a final area to explore is the storage network itself. This is important because it will better answer where SSD should be deployed, in the server or in the storage system. If the network is part of the performance problem, it may justify a network upgrade or the use of server-side flash as described above.

A storage network problem is likely if both host and storage controller CPU utilization rates are low. Further, if there's plenty of variation in the disk IOPS. This more than likely indicates a storage network performance problem. Essentially the host is waiting on its requests to be responded to and the storage system is waiting on requests from the host. The problem is then within the interconnected storage network. This can be solved by either a network upgrade or leveraging server-side flash. As with the other upgrades, it doesn't mean that the entire network needs to be upgraded, but a higher performance path needs to be put in place for the more demanding VMs data to travel.

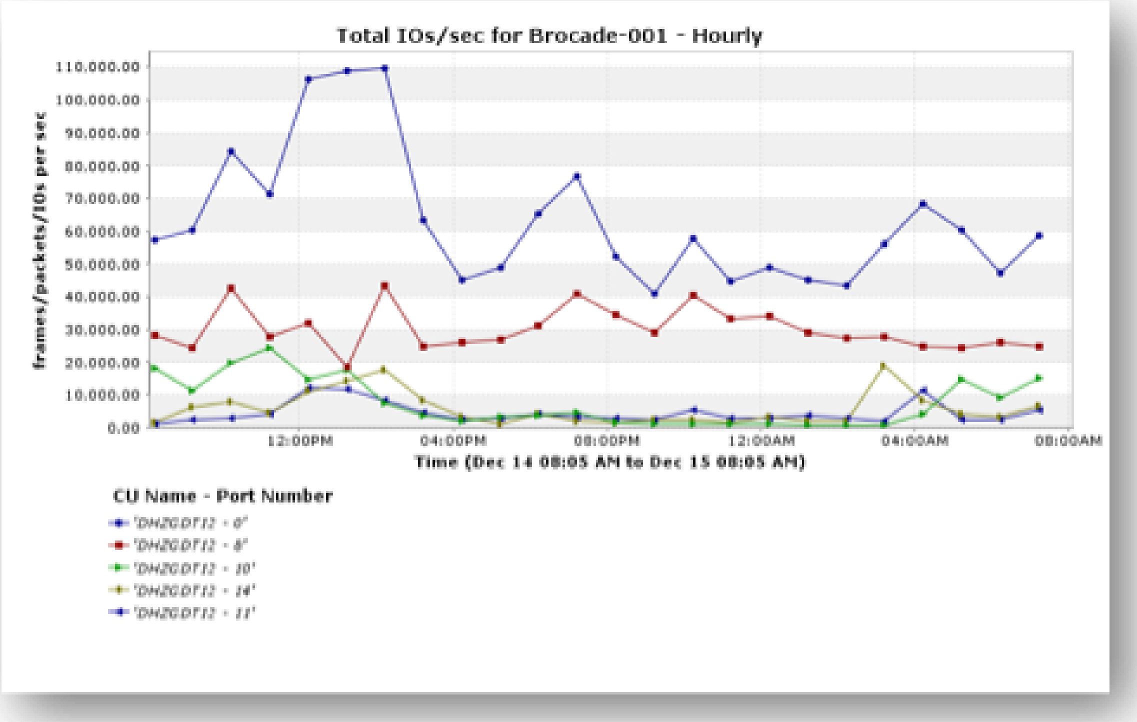

FC Switch port performance

Because IT is stretched so thin, the default answer to performance problems is to throw hardware at the problem. This is also a choice popular with vendors. In the modern data center, this means addressing the problem with flash, which is an expensive technology. Flash, because of its high performance and low latency, often fixes some of these problems. However, with proper analysis this “fix” can be more targeted, providing better results and lower cash expenditures.

SolarWinds Solutions for Monitoring Storage, Virtual & Database Environments

Storage Manager

- Monitoring storage performance & isolating hotspots in your multi-vendor SAN fabric
- Maps virtual machines to physical storage
- Automates storage capacity planning & reporting
- Simplifies analysis of storage usage & reclamation of storage space
- Monitor storage arrays from EMC®, NetApp®, HP®, IBM®, Dell®, Hitachi Data Systems® (HDS), and more
- Integration with Virtualization Manager and Server and Application Monitor for end-to-end visibility from application to virtualization to the storage

LEARN MORE » DOWNLOAD FREE TRIAL

Virtualization Manger

SolarWinds Virtualization Manager delivers integrated VMware and Microsoft Hyper-V capacity planning, performance monitoring, VM sprawl control, configuration management, and chargeback automation; all in one awesomely affordable product that's easy to download, deploy, and use.

- Real-time dashboards simplify Identification & troubleshooting of performance, capacity & configuration problems
- Identifies VM sprawl, helps you reclaim & optimize space & reduces licensing costs
- Integration with Server & Application Monitor provides application stack management from app to datastore
- Easy to download, deploy, and use – start monitoring your VMs in less than an hour

LEARN MORE » DOWNLOAD FREE TRIAL

Database Performance Analyzer for SQL Server

SolarWinds Database Performance Analyzer for SQL Server is a comprehensive database performance monitoring and analysis solution for DBAs, IT management, and application developers. Additionally, it eliminates performance bottlenecks, improves application service, and reduces overall cost of SQL Server operations.

- Resolve SQL Server performance problems by getting to the root of the issue
- Improve end-user application performance
- Get comprehensive visibility to both the physical and virtual resources impacting SQL Server performance
- Database Performance Analyzer for SQL Server is unique with no installed agent and less than 1% on monitored server

LEARN MORE »

DOWNLOAD FREE TRIAL

About SolarWinds

SolarWinds (NYSE: SWI) provides powerful and affordable IT management software to customers worldwide. Focused exclusively on IT Pros, we strive to eliminate the complexity in IT management software that many have been forced to accept from traditional enterprise software vendors. SolarWinds delivers on this commitment with unexpected simplicity through products that are easy to find, buy, use, and maintain, while providing the power to address any IT management problem on any scale. Our solutions are rooted in our deep connection to our user base, which interacts in our online community, thwack®, to solve problems, share technology and best practices, and directly participate in our product development process. Learn more at http://www.solarwinds.com.