

# Storage Options for VMware Environments

Sponsored by Dell and VMware



**Speaker: Greg Schulz, founder and analyst for server and StorageIO group**  
**Moderated by Collin Steele**

**Collin Steele:** Welcome to the SearchServerVirtualization.com webcast on ‘Storage Options for VMware Environments’. I am Senior Site Editor, Collin Steele, and our guest today is Greg Schulz, who is the founder and senior analyst for the server and StorageIO group. Thanks for joining us Greg.

**Greg Schulz:** Collin, my pleasure, nice to be with you today for this webcast. With that in mind, let us jump right into it. So, you see our first slide up, let us advance and what we are going to be talking about today is storage options for VMware environment. Collin, if we could jump to our first slide which is real quick, some background. I am not going to go through this slide in depth but you can come back and look at it. It is a little bit about who am I, who is server and storage I/O in the background as well as additional contact information with some followup information.

Let me jump to our first real slide which is storage for VMware environments. The theme here is the importance of storage for VMware environment. Here is the reality. Virtual machines or VMs need physical storage and they need that physical storage for several different things. Physical servers save their data, their persistent data which is on disk. Virtual machines or VMs are disk files that get loaded into server memory, applications and data require disk storage for consequently storage is also needed for backup/restore, BC and DR. So, there is a very, very strong affinity, dependency on both virtual machines, physical machines as well as storage, they go very much hand in hand. It is tough to have one without the other. With that in mind, let us jump to our next slide which is a graphic depiction of what we just mentioned on the previous slide. In other words visualization, seeing, is that important for the storage for the VM? This might be obvious, but it is great reinforcing the fact the virtual machine are effectively data structures that are mapped into memory but when they are not in executing they are a disk file. But in addition to the base virtual machine, there is also all the other files, the applications through the guest operating system, the actual data set things like that, so this picture here depicts that of storage and the dependencies of storage for both the virtual server as well as the applications that it is important. With that in mind, let us jump to our next slide.

On this slide, storage for VM environment, complimenting VMware and filling in the gaps. Hypervisors such as VMware provide a plethora of tools, plethora of plug-in. They have different capabilities, for example, with VMware you have facilities such as the site recovery manager which is a framework or collection of tools for gluing different pieces together. VMware while it has the site recovery manager, it also has a dependency on external third party technology such as replication, backups, snapshots, data protection tool, management, tiered storage and the like. The other aspect to this is there is the connectivity. There is I/O connectivity using serial attached SCSI SAS, iSCSI, fiber channel, FCoE, fiber channel over Ethernet or NAS, but down in the bottom there, where it comes back to we mentioned a few moments ago is that there are storage and I/O needs of aggregate performance, reducing latency, boosting IOPs and bandwidth in those aggregated environments. In other word, you have to support the availability, the reliability, and serviceability for data protection, backup, BC, restore, HA, DR, but it is also about optimizing performance, optimizing capacity and optimizing energy

efficiently, so often the focus around virtual machine and consequently storage for virtual machine is that around consolidation or writing up utilization, but there is another tenet which comes through in this chart is the importance of performing. In another words, as you aggregate multiple physical machines down onto a single machine running multiple virtual machines, you can introduce performance bottle neck, so multiple individual machines may not have had high-throughput, high bandwidth, high IOP, but as you aggregate them together, that aggregation can cause aggravation in the form of introducing bottlenecks, so in addition to looking at driving up CPU utilization, memory utilization, you also need to be that focused around storage, around being able to deliver the performance needs, performance criteria.

Okay, let us jump to our next slide. On our next slide, what is needed for VMware storage. Well, straight forward. Virtual machines need shared storage accessible by multiple physical servers. The caveat being there is you have got a standalone physical server that has no high availability, no business continuance, no failover type needs. All it is going to be just a single physical machine, no V motion, no BC, DR, just a single standalone machine, then you probably don't need to have shared storage. Everything else, shared storage is a necessity. Now, this is where things could become cloudy and that shared storage means a lots of different things. Usually when you think shared storage, you think, fiber channel FC, fiber channel over Ethernet FCoE, iSCSI, and NAS in the form of NFS, but shared storage can also be external shared, SAS serial attached SCSI. In another words, a radar ray that had serial attached SCSI that can be attached to two, four or more servers, similar to how you attach a fiber channel. The key notion here is shared storage using one of these mechanism, block or file. What gets interesting is depending on which vendor you talk to and if that vendor is an iSCSI only vendor, guess what they are going to say, "you need iSCSI for virtual machines. If they are fiber channel only vendors, you want to take a guess on what they are going to recommend, that's right, fiber channel only, likewise it is all the provide is NAP, you can guess that one, that is right. The way it all comes back to is a base requirement of having shared storage and that shared storage is there such that you can move virtual machines around that you can have failover, you can have load balancing, you could have high availability visits, continuous disaster recovery, that flexibility, that agility aspect of virtual machine.

Virtual machines also need basic reliability, availability and serviceability RAS features, in other words, redundant power supplies, redundant cooling hot-swap components, hot-swap disk drives. They need the ability to scale in terms of performance, scale in terms of availability, scale in terms of capacity and energy. There is an acronym there called PACE which is performance, availability, capacity and energy economic efficiency. Those are basic attributes of different applications. No particular environment, maybe you don't have a performance need, but you have the need for more space, but maybe you don't need lot of space to capture it but you need performance, or some combination, so that is where scale has different meanings and different dimensions that align to your particular need. Also important in what they need is integration with management tools, integration and plug-in with the VMware specific tools, whether it would be SRM for site recovery manager or the VMware console with that plug-in and also the ability to have a supported solution beyond basic interoperability. There can be some confusion here and

that you may have a product that is certified and you may have a product that is on the interoperability list, but that is not certified, so this is where some of that can become into marketing, some of it can become to what is actually a formally tested what is on the interoperability matrix? Here is another important one, it has to be easy to buy, easy to install, configure, and use it on a routine basis. So often I see products that are touted on their ease of use once you have got them in, their ease of installation, but I think equally important is how easy on the ongoing the day to day basis after you actually have it installed and are going to be using it over 3, 4, or 5 years, but also, how easy is that solution to acquire.

Let us jump to our next slide, that is what is needed for VMware storage, this is what is desirable okay. In other words, there is what I need to have in basic functionality, there is also what I want to have that would be nice to have, desirable, if I can afford it, if it is part of a solution, but they are not part of my base core requirement. Now you will get some push back on this one in terms of desirable feature. Some vendors will say that VMware VAAI API support is mandatory. Well it probably is, particularly if their product have that capability, if their messaging is very tightly aligned with it, yeah I would expect them to push in position that as an absolute requirement, but the reality is today because of lot of the VAAI functionality, API, the access to have VMware leverage, the underlying storage systems are still very much evolving. Today, I would put in desirable apps community and in a year from now I will probably move that into a required category. So that is one of the subjects to go depending on who you are talking to, but desirable features, management reporting or notification as well as basic usage metric, reliability and serviceability beyond what we talked about, in other words, additional hot swap capabilities, the ability to apply formal updates, yield apply code updates nondisruptively, but also while in use. Basic data protection, multiple RAID levels, snapshots, local and remote replication as well as tiered storage medium, solid state disk, or flash drive also known as EFD or enterprise flash drive, 10,000 or 15,000 (10K or 15K) SAS, and fiber channel hard disk drives as well as high capacity 1 and 2 terabytes SAS as well as SATA drives. So these are desirable features, they maybe part of your base and as I mentioned some of these are moving more and more into basic mandatory must have use the analogy jacks are better, in another words basic functionality.

Let us move onto our next slide. These are optional okay. What is optional. Intelligence power management, the ability to vary the performance and the power consumption based on the amount of work being done. Early generation power management or intelligent power management where storage is focused around pinning the dips down, in another words, if they are being used, power them down, but now we are into second and third generation intelligent power management, what we are seeing on the CPU or the core chips like those from Intel, those from AMD, is where you have different power modes. In another words, when there is a lot of work to be done, increase the frequency or the speed of the chip and when there is less work to be done, slow them down, so they consume less energy, so in other words, using energy more intelligently, using power intelligently that when there is work to be done, get the work done as fast as possible.

When there is not much work to be done, go into a slower state, a lower power usage state, but still be readily available to get the work done.

What is also becoming more and more popular, automated or system assisted data migration and tiering. Unified or multi-protocol support supporting both concurrent block and file as well as data footprint reduction techniques. In other words, support for archiving, compression, dedupe, thin provision, and space saving snapshots. Also application aware, application enabled data protection, leveraging snapshots, replication, continuous data protection as well as backup restore. Let us jump to our next slide. In other words what are your VMware external shared storage options? Well, there is quite a few. It could be as basic or as simple as storage shell attached to your physical server, PCIe, bus using a PCIe RAID card in the physical machine. It could be leveraging a virtualization appliance with storage attached to it. It could be a block, it could be a file, it could be a block and file, in other words, a unified or multi-protocol device. It could be a modular storage with expansion capabilities or it could be a large scalable enterprise class storage system depending what you need is. It could also be a clustered storage system.

Let us jump to our next slide here. Determining what VMware external shared storage you need. What it comes back to is this, determining your needs versus your wants. In other words, you must-have versus nice-to-have wish list. What this is really saying is that there is a lot of new things that are out there in the market place I would love to have, but there is a basic capability that I need to have to get my job done, to get my work done, to support the business activities that I need to accomplish. What if the nice-to-have or the one to have are affordable or they become part of that price where it becomes price effective, then they become more and more viable for what this is really separating is setting and establishing your basic needs, your wants, and your wish list and then in the course of going through your acquisition using that criteria. In other words, sticking to your guns, sticking to what you need to meet your basic objective, and then what else can you get per dollar spent? Having said that, what is your budget both capital and operating? What in other words, you must be able to get a deal on some storage especially right now, it is very cost effective, but what are the ongoing costs going to be billed for maintenance but also for ongoing software licensing or other fees. What are your performance needs? In term of read versus write, random versus sequential, big versus small IOPs, concurrent versus parallel. What kind of applications that you are going to be supporting. Also keep in mind that those individual servers that you are consolidating by themselves, they may not have much of a performance impact, but again keep in mind that aggregation can cause aggravation by introducing performance model and as you add them together.

The message around performance is we often hear about megabytes per second, gigabytes per second but keep IOPs, you are doing lots of small IOs, you are not going to see a large megabyte per second or gigabyte per second rate, and you are going to see lots of activity, so keep that in perspective. Couple other points here, what are your availabilities in data protection requirements? You need single or dual controller, you need high availability clustering, failover. What are your growth plans or requirements?

This is going to help determine whether you should buy something now that has plenty of growth potential, may even influence type of architecture solution you go for, so that you can pay as you go, add the technology as you grow, but then there is the other adage which is “If you know you are going to be growing and you know what your growth rates are say 18 months, 2 years, 3 years out, that large frame-based approach might actually become more attractive for you particularly if you can negotiate in pricing incentive or the add-on capabilities, so very, very important there is what are your growth plans or requirement and then this is the interesting one. What are your preferences? All those other things aside what is your preference for a vendor, architecture, protocol, block, file along with coexisting in your current environment.

Let us move on to our next slide, let us start to wrap up here. Closing comment, there are many different options for VMware storage. Identify what you need versus what you want as well as what you would like. List all of those out on a sheet, list all those out on your worksheet, on your comparison chart by also categorize that in terms of must have, need to have, would like to have, would want to have and then in your course of going through and evaluating the technology, use that as part of your criteria, in other words, what meets your base requirements, your basic needs and then what else can you get from those wish list items as a part of your solution without having to increase the acquisition costs outside of what your budget support. Also how will a technology or technique work for you? or are you going to have to work for it. In other words, there is technology that is easy to install it but how is it on an ongoing basis. Have a multi-year strategy, regular revisits, and update as needed and balance the old with the new, don't be afraid but look before you leap. Where you can learn more, there is information about my website [storageio.com](http://storageio.com), my blog [storageioblog.com](http://storageioblog.com) as well as on twitter at [storageio](https://twitter.com/storageio). If you have any questions, feel free to IM, tweet, e-mail. I will be happy to chat with you and with that I will turn it back to Collin.

**Collin Steele:** Greg Schulz, thank you very much for your time.

**Greg Schulz:** My pleasure.

**Collin Steele:** And thanks for viewing the [SearchServerVirtualization.com](http://SearchServerVirtualization.com) webcast.