



7 STEPS TO QUANTIFY A CONVERGED SOLUTION OPPORTUNITY

Sharon March

Manager—Field Technical Consultants/
Education Lead

Avnet Technology Solutions

sharon.march@avnet.com

EMC²

Table of Contents

Introduction	3
Layer 1 (Physical) Physical	7
Layer 2 (Data) Network	8
Layer 3 (Network) Compute	11
Layer 4 (Transport) Storage	13
Layer 5 (Session) SECURITY	16
Layer 6 (Presentation) Management	19
Layer 7 (Application) Applications	22
Conclusion	27

Disclaimer: The views, processes or methodologies published in this article are those of the author. They do not necessarily reflect EMC Corporation's views, processes or methodologies.

Introduction

We've all had that colleague. You know the one. The guy who has done hundreds of implementations and seems to be able to fly in and out of a customer site inside of a day, leaving behind a completely functional converged solution, installed and production ready. But how did he get to that point? If you ask him, he will probably say something along the lines of "I have a checklist that I go through in my head", or "Oh, I just know what to look for". Wouldn't it be nice to be able to take the "art" out of "architecture" and have some uniform structure around planning for implementation? In the age of doing more with less, repeatable procedures are becoming not only desirable, but required. However, in the world of converged infrastructure, design repetition is not always possible. The very nature of converged and the journey to the cloud mandates custom solutions, on some level, for nearly every business. Not everyone is at the same point in that journey to the cloud, and not everyone has the same resources, infrastructure or personnel to make that journey easy. Uncovering what a customer has and what a customer needs is always a challenge. Selling converged solutions adds another layer of complexity to that discovery process. It's not enough to look at just the compute or just the storage in a converged solution. You need to approach it from a much higher level, yet with a great amount of detail and with more focus on following the information through its entire path than ever before. Approaching the pre-sales engineering qualification stage of a converged sale with structure and forethought can minimize the on-site surprises that nearly inevitably happen when you have screwdriver in hand, implementing the new solution.

Too frequently, pre-sales engineers are brought into a sales cycle after the deal is closed, the purchase order is cut, and delivery is scheduled. At that point, a salesperson may think "let's just make sure we have all our ducks in a row for implementation". Unfortunately, sometimes the pre-sales engineer then becomes the bearer of bad news because cables are too short, cooling is inadequate, or applications are not ready for true convergence, at least how the customer envisions 'converged' in their environment.

This article will walk you through seven steps, loosely aligned with the 7-layer Open Systems Interconnection (OSI) model, to ensure that you ask the right questions, or start the right thought processes, before equipment arrives at a customer site. By following this methodology and proceeding through the steps, you can make sure that as many questions around the *entire* solution are addressed as possible, before you begin implementation. Questions like:

- Do you need to pull different power plugs?
- Are storage tiers moving or growing because of your cloud strategy?
- Where is your security policy starting and stopping or expanding to in this new converged world?
- Are your applications deployed and managed differently in a virtualized environment?
- How is your cloud strategy affecting your compute requirements, which may be affecting your virtual storage area network (VSAN) requirements?
- How is your mobile workforce using the converged infrastructure solution and how does that affect the design?

By asking questions as a structured flow, you may not only impact the final build and plan, as well as verify initial requirements, but you may also impact your scope of work and the service you provide. Answers to your questions may have the added benefit of leading you to new opportunities that the sales rep hadn't really considered, or allow the customer to reconsider a service that they previously may have minimized.

When selling a converged solution, you need to look at the solution holistically as it stands in the entire customer environment. The environment these days extends well beyond the physical confines of the data center. Additionally, as we abstract applications from hardware and sell different types of services to enable that path to cloud, we must also abstract our thinking and questioning from traditional, siloed sales. That is the only way a reseller can stay relevant now and for the future. To be effective in selling a truly effective converged solution, you must be prepared with a plan and a logic to qualify and quantify the solution as completely as possible. That is what this article will prepare you to do.

Data centers are still sold in a siloed manner, just wrapped up in virtual buzzwords, i.e. "hyperconverged", "converged infrastructure", and that all-encompassing nebulousness, "solution". But what does that all mean? More importantly, what does that all mean to the customer? A solution is only a solution when it solves a customer business problem to the customer's satisfaction. Anything less is a bunch of disparate and segmented hardware sitting in a data center that a customer has to manage and administer. The only way to get to a true solution is to make sure that you capture as many of the pain points as possible up front, during

the design phase. Why? To avoid triage after the fact when you are either doing the installation, conducting initial rollout, or worse yet, months after the fact when the customer is fully annoyed and convinced that you sold them a hunk of junk. So how do you capture as much information as possible before even a cable hits the ground?

Most pre-sales design engineers decide early in their careers that they want to focus on one solution stack or the other; FlexPod, vBlock[®], VSPEX[®], VersaStack, whatever. Once they select the solution stack, two things happen; one, they go about the business of getting all the appropriate certifications so the letters on their business card lines up with their chosen area of expertise and two, self-preservation. After a few not-so-smooth implementations, they decide “this will never happen again”, whatever “this” is, and they start a mental checklist. For example, if they sell a certain kind of rack only to discover upon implementation that the customer does not have a L6-20 power drop in their data center (let alone 4), they make a mental note to ask about power outlets. Likewise, if a design engineer sells a particularly rear-heavy chassis, only to find out that the customer has two post racks, they make a note to ask about rack posts and dimensions. The checklist grows as the engineer completes more installations, and unfortunately, for those living outside of his head, that list is typically a mental checklist based on an individual’s unique experiences. Usually, that list falls short in some form or another at some point of the implementation because very few people have installed every iteration of a solution stack, let alone published or posted their personal checklist. This article will help you start the process of formulating a checklist that you can customize and tweak to your chosen solution stack, vertical market, and/or area of expertise.

The checklist outline is loosely based on the 7-layer OSI model. As you move up the seven layers of the stack, you address each portion of a solution. The idea of this checklist is not to be all-inclusively suited for every situation, but to start the conversation; get the customer thinking holistically about the solution, so facets of the solution are not overlooked. Facets offering functions which could be vital to the success of the solution implementation and business structure. The checklist is also intended to help you expand your own horizons, and help you to intelligently inquire about aspects of the solution that may be outside your area of expertise. As you work through the process with your customer, your broadened knowledge will help you position yourself as a true “trusted advisor”, giving that term substance. Your area of expertise may be storage, but by asking a few focused, well-thought out questions about compute, you transition from “the SAN guy”, to someone who not only has the customer’s best interest at heart but who truly understands what the word “solution” means. It’s important to remember that

the definition of solution really only matters in how the customer defines it, not you and not your company. It's only a solution if it solves a problem. Using the checklist as an outline and then fine-tuning it to your business offerings enables you to work with the customer to develop a properly customized solution, and make yourself and your proposal matter to the customer.

The 7 layers of the OSI model are, in order:

1. Physical
2. Data
3. Network
4. Transport
5. Session
6. Presentation
7. Application

Correlating those layers to the components of a converged solution, you can logically work your way through all aspects of a converged solution, and start the customer thinking of not only how the solution is going to affect their data center and infrastructure, but how the solution is going to affect their business and business processes. In this article, each of the layers is associated with a part of the converged infrastructure solution about which you should be thinking and asking questions. They are associated as follows.

- Layer 1 (Physical) Physical – Cables, power, cooling, and other physical attributes of the data center
- Layer 2 (Data) Network – attributes associated with the network infrastructure
- Layer 3 (Network) Compute – compute requirements for the solution
- Layer 4 (Transport) Storage – storage requirements for the solution
- Layer 5 (Session) Security – security considerations for the solution
- Layer 6 (Presentation) Management – management options and considerations for the solution
- Layer 7 (Application) Application – software and business application requirements for the solution

Layer 1 (Physical) Physical

Converged infrastructure collapses a lot of components, and traffic from that collapsed infrastructure needs to travel properly to the users. Generally, we really don't care what the data center looks like, just as long as the cables connect to the right components, traffic gets where it needs to in a timely fashion, and the lights stay on. However, this base level requirement is something that cannot be taken for granted. You cannot assume that just because a customer wants unified traffic that their infrastructure is ready for unified traffic. You can't even make the leap of faith that the customer has Gigabit Ethernet (GigE) cabling, let alone cabling that can support Fiber Channel over Ethernet (FCoE). The good news is, this is typically the easiest and least intrusive place to start your questioning. This is where most IT folks are comfortable because this is where they started their careers. By simply asking the basic question "Do you have GigE capabilities, and if so, can you please describe the cabling to me?" you start in an area that many IT managers and personnel are extremely comfortable, and yet an area that is frequently overlooked. When you ask about the stuff that everyone else doesn't, you have differentiated yourself from the start and have made great gains of trust. You immediately show the customer that no detail is too small to consider, and everything in a converged environment matters.

On Layer 1, when discussing the physical attributes of the data center, what questions do you ask? Start literally from the ground up with power and cooling.

- Is there sufficient power for the draw of the converged infrastructure (CI) solution stack?
- Do you have the correct power plugs for the rack?
- Do you have enough of the correct power plugs available? Is it going to be hardwired into an uninterruptable power supply (UPS)?
- Is the UPS battery sufficient to sustain brownout or blackout conditions with the additional draw of the new CI solution stack?
- Is there sufficient cooling available in the data center for the British Thermal Units (BTUs) that the CI solution stack may add?
- Is the airflow coming from the top or the bottom or from side to side?
- Will that airflow affect the power supply airflow of the components in the cabinet?
- Do we need to move fans either in the room or within the solution itself?
- This leads you to the physical attributes of the data center itself.
- Where will the cabinets be located in the data center?

- Is the floor strong enough to hold the additional weight?
- If the customer is supplying or using their own cabinets, are they 2 post or 4 post?
- Is there a racking kit available to support the equipment in the customer supplied racks?
- Are the cabinets going to be right next to other cabinets, or are the racks going to be back to back?

Once you know exactly where the CI solution will be physically located, you can then ask about how it is being physically connected to the enterprise.

- Are the quoted cables long enough to thread through the floor or ceiling and then connect the equipment together the way the customer wants?
- Are the cables going through conduit and is there enough room in the conduit?
- Does the customer own all of the conduit?
- Is the cable insulated properly (plenum or riser) for how it will be run?

Since we are talking about physical attributes, here is an opportunity to talk about the logistics of getting the equipment to its final physical location.

- Is the door wide enough or tall enough to accept delivery of the solution?
- Is there a loading dock? (A surprising number of customers will answer “no” to this question.)
- Is there a freight elevator?

This particular layer of the conversation tends to raise more questions than answers, but that is a good thing. We’ve all been there; the \$50 cable holds up the \$5MM deal for a week, often at the end of a quarter or over a holiday weekend. Since this is the most basic and seemingly obvious line of questioning, frequently it is assumed that “someone” has already asked these questions or looked at these points. In a converged infrastructure sale, assumptions are a bad thing. If the answer isn’t written down in your project or pre-sales documentation, ask the question.

Layer 2 (Data) Network

In the OSI model, Layer 2 is the data layer. This is how your network sees various devices, so this is where we start talking about the customer’s network itself and how the network sees the CI stack. Assuming that the cabling has been qualified and the physical backbone is present to support a unified infrastructure, this is where you discover their network’s preparedness for

basic converged networking features. What questions do you ask to make sure that the network layout and design is ready to be planted on top of the physical cabling foundation you just quantified in step 1? Again, starting with a very basic question comes in very handy.

- Do they have a network diagram? A startling number of customers do not, so as in layer 1, assume nothing. If they do not have a diagram, then you have the opportunity to offer some of your own discovery and assessment services. This is where you also discover their preparedness for basic converged networking features such as Fibre Channel over Ethernet (FCoE) or voice over IP (VoIP), if those features are part of the proposal.
- Have they truly looked at their network layout and made adjustments where necessary to prepare the topology for converged?
- How is the enterprise organized and divided?
- What are the local area network (LAN) segmentation requirements, which may extrapolate into the virtual LAN (VLAN) requirements? Are there certain groups or LANs that need to be physically separated? Keep in mind that customers are at different points in their comfort level with virtualization, so they may not be comfortable converting their entire network to VLANs. There may be security, regulatory, or organizational reasons why a specific group must have their own physical network, rather than a virtual one. Plus, if your CI stack includes virtual network interface cards (VNICs), that may affect the cabling you just clarified and the upstream port count to either an aggregation or core layer from the converged infrastructure.

These answers may affect how you configure the networking portion of the solution and how you connect the CI solution stack to the rest of the enterprise.

Once you have the network diagram planted on top of the cabling properly, then you want to look at expected network performance. Experience has proven that while the network is not the only place where bottlenecks occur, it is usually the best place to start your throughput diagnosis. Make sure that you are looking not only at how the network is laid out, but how the network packets are travelling from point A to point B and if there any points of contention along the way. Then, make sure that the traffic pattern and flow is acceptable.

- Are there certain groups that require specific bandwidth?
- Are there quality of service (QoS) requirements?
- What is the network tiering structure?

- If network layers are collapsed, would the converged infrastructure stack perform better if those layers were separated out?
- Where does Layer 3 routing reside; within the CI stack or outside at the distribution or core layer?
- Ultimately, are you sure of your diagnosis? Do a simple trace route or packet trace, using after convergence implementation modeling tools, if available. The network slowdown that the customer identifies may or may not truly be the bottleneck. It may be a cabling or compute issue. Just be sure that your analysis doesn't change any downstream discoveries.

Of course, that's just talking about inside the enterprise; items that are within the customer's control.

- What about the items outside of the customer's control, like Internet Service Provider (ISP) connectivity? Where does that reside? Is it going to be embedded in the CI stack? Do you have just one ISP or multiple ISPs?
- Are there failover requirements that must be considered when implementing VSANs or port binding?
- Do you have a mobile version of your virtual desktop infrastructure (VDI) which is different than your desktop version and, if so, how many of your users are or will be remote?
- Is your ISP prepared for the growth that may occur when you have a more mobile workforce?
- Is your budget prepared to address the ISP?
- If you need more connections from the ISP as your mobile workforce grows, where are those connections coming into your CI stack?

The more time you spend getting a crystal clear picture of not only what the network looks like now, but what it is supposed to look like after implementation, the better. Keep in mind that the network is the one component that touches all of the others in the stack, so the quality and depth of your questions here will be important.

Layer 3 (Network) Compute

Layer 3 is the network layer in the OSI Model, but for purposes of this article, we correlate it to the compute portion of the CI stack. This is where you can start talking about compute requirements. Virtualized applications perform differently than in a physical implementation. Applications may be structured to have certain resources available to them, like central processing unit (CPU) threads or core, that would be decoupled and placed into a pool to be used by many applications in a virtualized environment. Most applications are not used 100% of the time, nor use 100% of the compute requirements. However, your customer probably has some key applications that must have compute resources available at all times, without risk of contention. This is again, where you want to go back to the basics.

- Have you gone through a virtualization sizing exercise like for example, your e-mail application, in a virtualized as opposed to a physical environment?
- Do the results of that exercise affect your blade layout?
- Are there threading or core requirements or limitations to the application, once it is virtualized?
- Can all the applications that you plan to use even be virtualized in the first place?
- Can the application be spread across multiple CPUs or multiple blades, and will that degrade or improve performance?
- Does the application have any performance spikes, like boot storms or running a batch job, that must be considered?

These limitations to virtualizing the application may affect how you actually configure each blade. Further, you may need blades that are dedicated to certain applications or groups of users simply so they have the compute power they need on demand. A classic example of this is your research, test, or development teams. Do the compute resources for testing and development need to be isolated so as to not risk bringing down production? Can research and development (R&D) itself even be virtualized? How many rollback versions of an application do you need not only access to, but run simultaneously, and does that affect the compute resources assigned to R&D?

This is also a good point to start considering memory requirements. We will categorize memory under compute, because it supports compute functions so directly. Similar questions about compute and CPUs can be asked about memory. You need to discuss not only the amount of memory that will be used in a virtual memory pool, but how it will be used.

- Do the applications need to have dedicated memory resources?
- Are you expecting the virtualized memory to be accessed at the application level in a cached file-system fashion, or at the operating system level in a page cache fashion?
- What is your expected and acceptable memory over-commitment ratio?
- How much swap space do you want to have in your secondary storage and what is your paging strategy (on-demand, pre-fetch, etc.)?

Similar to your storage discussion, as we will see, memory chunks simply thrown into a pool is not sufficient detail. How it is accessed and allocated plays a part in sizing enough memory so your bottleneck doesn't move from your network to your memory. Does the quantified memory now affect the number or type of CPUs and how they are laid out on the blades? Circle back through your process and questions to make sure nothing already determined has changed. The basis of a converged infrastructure is that all the components work together in a cohesive manner. Keep taking a step back after each quantification, so you can continue to look at the solution holistically.

High Availability (HA) also merits its own sub-category under compute because applications ultimately need CPUs to work. You have probably asked the question of which applications MUST be functional for the business, so that logically leads us to a HA discussion.

- What are the HA requirements?
- What is acceptable downtime? In other words, what is the tolerance for production failover?
- Where do the HA or disaster recovery (DR) or rollback compute resources reside?

Typically, we think of HA as internal to the data center but that may or may not be the case. This may affect how you want to lay the blades out in the chassis, depending on how or if HA works within the chassis or across chassis or possibly across sites. When you do fail over, is a performance drop acceptable? Does every application even need to fail over? This may help you size the backup converged stack a bit smaller or with less horsepower to help save the customer some money. Once you determine what the HA plan is and how it will be implemented, you may want to go back to Layers 1 and 2 to make sure that the failover will happen across the network as expected. You may want to revisit cabling and ISP or WAN connectivity, and then size the HA environment in a similar way as the production environment, starting at our Layer 1. Or if the failover is local, does placing high power-draw blades in one

chassis roll back into effecting your rack cooling and power? As you can start to see, your questions begin to build on each other. Think of the process as a 3-legged stool that must always be level. If you change one leg in any way, does it affect the level of the seat? Keep circling back to make sure that your already established quantifications don't change when a requirement is identified further up the stack.

Layer 4 (Transport) Storage

For the purposes of this article, Layer 4, the transport layer, is used as a memory jogger to talk about storage. As in, how is the data being transported once it is generated? As with memory, storage can no longer be considered as a giant vat attached to your all-powerful and critical compute portion. It has its own true value and set of considerations, especially if the customer's storage strategy is tiered. Once again, the simplest question is the best place to start. Let's assume that you have asked the basic storage questions of how much storage the customer has right now, how much they forecast needing in the next five years, and the RAID level they would like. These will at least get you into a sizing discussion about how many spindles your client will need.

However, in a converged infrastructure, the conversation needs to go much deeper than that. You not only need to discuss how much is being stored, but where it is being stored, and how the users are accessing the data. Is the data on premise, offsite, or a combination based on the structure? For example, archiving may be just fine off premise, but if a business lends itself to being exposed to lawsuits, you may need to consider ISP bandwidth and HA strategies to access those archived files and protect the company. How are users accessing the data and has implementation of the HA strategy changed with convergence? In other words, are there certain business-critical data that cannot tolerate any downtime or unavailability, regardless of the underlying infrastructure? If so, how is it protected now and how does that protection structure change in a converged environment? How is the storage structured? Does the structure change with a converged solution? Not just the size of each of your storage tiers, and where they are located for access purposes, but also the quality of service to deliver that data to end users in a virtual environment. How are the snapshots or backup copies being created and protected? How many snapshots do you want to have available, and what timeframe does "available" mean to you? What are the access requirements, and who can access the snapshots? Those answers may affect how the storage is tiered in a converged solution, which could alter not only your storage layout configuration, but your network configuration as well. Again, a diagram or discovery service can prove useful here, to not only uncover how much

data the customer requires and how much they are sizing for future growth, but how they have structured their SAN for availability and quality of access to business-critical data.

When discussing the structure of the SAN and the data, you cannot take any storage lingo for granted. Perhaps you are accustomed to thinking of storage as cache, access, and archive tiers, but is your customer? Do they only have cache and access? Do they only have cache and archive? It is possible that a customer may be moving from a hard tiered storage design to a more cache, pool-driven design, in an effort to take full advantage of the agility inherently offered in decoupling resources in a CI stack. If they are, you will also want to address not only storage tiering, but use and movement of that cache pool, once you have defined what they mean by “cache pool”, and how they plan to use it. By probing into the requirements of accessing business critical data, and how the CI stack will be expected to deliver that business critical data and to whom, you can make sure that your storage infrastructure design improves with the converged solution. Your questions here will probably be very similar to the memory questions.

- Is your rule of thumb sizing appropriate for the VSAN structure?
- Do they want to increase the size of available cache via the cache pool or inline storage pool in order to improve performance?
- Where do you want your swap space to reside?
- Do you now want different tiers to have different RAID levels?
- Is each tier and each logical unit (LUN) in the tier going to be virtualized?
- Is the storage going to need to be able to move with or tied to a specific compute node?
- Or are you going to tie just the metadata and protect the actual data?
- Have you discussed port binding in the network conversation and now need to bind tiers or slices of tiers to those ports?
- Once you bind the slices, is your network now adequate to support the data access demand?
- Is the compute now sufficient to refresh the data at an acceptable latency?
- Will tier sizes or the tiers themselves change to mitigate unacceptable latency?

You will also want to consider the traffic pattern of the storage packets, both upstream to the compute and downstream to the storage itself.

- How is the storage traffic flowing? When looking at the traffic pattern, you need to think of not only how the data is moving around the array and getting presented to the storage processor, but how is the data or request for data arriving at the array in the first place?
- Is the choice of storage switch correct for a CI stack or is it really just a SAN switch?
- Is the switch handling only storage traffic, or is it intended to be a unified switch performing some form of network routing as well?
- How does that answer affect switch connectivity to the array or into the SAN?
- Does the storage processor need to be changed to an FCoE storage processor?
- Does the storage controller even support FCoE in the environment's context or is a different storage controller needed?
- If the array is going to be virtualized for the VSAN portion of the CI stack, can it even be done in the first place and, if so, will it provide adequate access?
- Is the controller sized properly to not only support the tiering strategy but the virtualization strategy as well, virtualization meaning not only the VSAN but the virtual machines (VMs) and the VLANs?
- Where is the control of the LUNS in the VSAN? Is it at the storage processor or at the switch?
- Is there still a need for legacy fiber channel, and how does that affect the answers from either the compute or network or both conversations?

The point of this line of questioning is to remind you to keep circling back and verifying your quantification of the previous three layers. As you size and re-size the storage and follow the storage packet, make sure that a component or feature that you suggest to solve a storage performance problem, i.e. upgrading the storage switch or storage processor, does not introduce an upstream issue, like impacting network management traffic bandwidth or necessary fiber channel cables no longer being present in the infrastructure.

Typically, the storage conversation is where you need to put on your regulatory hat. You need to consider, as you size the VSAN and storage, how your new sizing and virtualized data protection strategy may be influenced by data retention regulations. Is or will the customer be subject to new regulations? How do those regulations affect the storage capacity and tiering strategy? Does the storage structure need to change because of new or upcoming compliance

regulations? This won't be just a conversation for healthcare or banking customers. Suppose that your customer is a small construction business who keeps all their data on premise. They may be subject to some of the same healthcare regulations around retention of employee health records that a larger, enterprise healthcare company may be.

You also need to look at required access to storage in the same manner that you looked at required access to applications when you discussed compute needs, particularly from a remote user perspective. Mobility and remote access to associated data needs to be part of the storage conversation if you are integrating it as part of the CI stack, and that access needs to be seamlessly integrated into your CI stack so that end users don't know or care how they are getting their data. Think of the games on your phone where your progress is saved to the cloud and you can access the game from multiple devices. Someone, or a group of someones, had to consider how users were going to access that data anywhere from any device, and tier the storage and access to the data such that online or frequent pay-for-features players had fast access, whereas the sporadic, only play for free players had perhaps a bit of latency in their games. The same type of thought process needs to happen when you are considering business data. With digital information and regulations around it becoming more and more pervasive, storage is one part of the solution that requires serious consideration. Since you are approaching the CI stack in a holistic manner, just as much thought and planning for storage needs to be applied as for compute or network. After all, your true value as a trusted advisor is to present the right solution, no matter where the path leads. You are taking a Sherlock Holmes approach, following the evidence to the only solution that makes sense for that environment, no matter how improbable it may seem for another.

Layer 5 (Session) SECURITY

This header is in all caps because this is a big one that people tend to overlook, especially if the solution you are offering is going to be a private cloud or internal stack. The perception is that because it is not exposed to the outside world, the standard enterprise firewall will be sufficient to protect the CI stack; nothing could be further from the truth. If the CI stack is going to be running the business or any significant portion of it, you need to keep in mind that most attacks come from inside the enterprise, not from the outside hacker.

While firewalls and intrusion prevention systems (IPS) do keep a significant amount of danger from entering the enterprise, it's really not a matter of if the data center gets breached; it's when. For all the time your customer has to spend protecting and guarding their data and business

applications, the hackers have more. Interjecting mobility into the mix introduces another level of exposure. Let's face it; when was the last time that the average user installed a firewall or data encryption on their smartphone? Unless security is front and center in your mind, chances are the user was more worried about getting a signal at the top of the hiking trail rather than making sure that signal was a secure channel. Now, of course nobody is taking about swapping your smartphone for a satellite phone you can use at the South Pole, but the fact of the matter is that the bell has been rung. The mobile workforce is here to stay, and everyone with a handheld device is now a security risk. Our mobile devices, especially with the open environment of application downloads, are the biggest threat to each of our enterprises. Thus, we need to ensure that anyone using that freshly minted converged solution is properly authenticated and authorized to access what they are trying to, no matter where they are, both internal and external to the network. If you have no experience around security, where do you start? What do you ask? Once again, the basics serve their purpose very well.

In speaking of security, you may need to come back to the customer with specific clearance or approval. Like storage, security can be its own, standalone conversation, and depending upon the industry or business, the participants in the conversation may need to be restricted. For now, let's assume that you have been invited and cleared to discuss your customer's security strategy. While you may not be able to have access to their rule base and security permissions, you can ask the right questions to make sure that your CI stack is prepared to deliver the proper security measures to the client and to their clients. Although this is far from a complete list, since every environment has their own set of unique security needs and circumstances, this should get you going in the right direction. Starting at the beginning, it's helpful again to assume absolutely nothing when it comes to security.

- The first question should be, do you have a security plan or practice in place?
- Is it compatible with the CI stack solution we are suggesting?
- Are there regulations that you must abide by?
- Are there measures planned or in place to separate functional areas?
- Are the measures constricting enough to secure critical or confidential data, but open and flexible enough to enable collaboration when appropriate?
- Are not only the business groups separated and isolated appropriately but do you know where each of the business groups are going outside of the enterprise?

- Can their clients access or not access data? In other words, if your marketing department frequently outsources pamphlet layout to a third party, has the marketing department granted that third party company access to corporate content behind the firewall, with the intent to make it easier to create literature?
- Has that access created a hole in your enterprise?
- Do you have a remediation strategy in place to not only discover breaches but prevent them from happening again? Can that remediation strategy discover all of the components you are proposing in your CI stack? In other words, is the CI stack introducing another layer of security complexity or removing it?

Again, this conversation alone can spin off into a whole other solution, which needs to be taken into consideration when implementing the converged solution. So keep in mind, even if the CI stack you are proposing is going to be “isolated” or “hidden from the rest of the enterprise”, you should be addressing security and its impact in nearly every scenario. The answers can play right back into the beginning of your discussion around cables and networks, all the way up through storage and how you may be securing parts of the SAN. Make sure that as you get the answers to questions as you move up the questioning stack, you keep circling back to the quantified answers you have already received to make sure nothing has changed.

Remote users and the security of those remote users introduce their own layer of security complexity. Threats are now focused on the behavior of the enterprise, not the devices themselves. Inherently, remote users are more exposed to threats than users at an office, physically behind a firewall. So when converging, especially in a VDI environment, you need to consider how that change in network access and user behavior is going to expose vulnerabilities and how you will catch the vulnerabilities before they impact your enterprise.

- Where exactly is the endpoint of your security policy in the enterprise? At the edge? At the mobile device?
- Are you going to issue phones that are in your control or are you going to use an access application such as Good?
- How does secure access from remote users impact your connectivity from your ISP? Or are they coming in through a virtual private network (VPN) tunnel that you establish and where does that VPN reside on the converged infrastructure?
- Are you going to implement single sign on (SSO)?
- Once a threat does happen, where does it go? How do you see that?

Careful attention should be paid to those remote user packets as far as accessing the CI stack and propagating threats throughout the enterprise. Follow the packet as it flows in and out of the enterprise and the converged solution. Make sure that as you plant security requirements and turn on security features to mitigate threats, you don't re-introduce bottlenecks to the underlying infrastructure or uproot improvements to the network, compute, or storage functions.

Layer 6 (Presentation) Management

This can be a big one, and also a polarizing one. You have the LAN guys who want to do it one way, the SAN guys who want to do it another way, and the security guys telling them they shouldn't be doing it at all. Almost job one, when you start talking about management, is to ask the IT administrators, "What's on your wish list?" Frequently, when talking about converged solutions, you want to steer the customer in a direction, so you ask closed-ended questions. Do you want this feature or that feature? However, when it comes to management and the management interface, that's an entirely different matter. This is what the administrators do every day. This is what they were hired to do. The way they see it, all of the hardware and software that you have talked about up to this point is simply adding to their task bucket of "things they need to tend to". So, asking the open-ended, "what do YOU want" here is very powerful. Not only are you no longer piling things onto their plate, but you are now positioning yourself in their corner to make their lives easier. Plus, by asking the question, you may open up opportunities you never considered.

Now just to be clear, ease of management and reducing total cost of ownership (TCO) is a huge and legitimate value proposition in a converged solution. Any sales rep worth their salt mentions that as part of getting in the door and closing the deal. The problem is the IT people actually performing the management tasks have probably heard it all before. Here is where your skill in asking open-ended questions and actively listening and documenting responses will be key in getting buy-in from this team. By digging down just a bit further, and asking questions of the folks who are supposed to be direct beneficiaries of that "ease of management", at their level and in their language, will advance your credible advisor status considerably. Do not minimize the impact that the line administrators can have on the success of your converged implementation. Tune in to their responses, paraphrase back, and get crystal clarity on how they expect the solution to actually work.

One of the biggest benefits in a converged solution is the holistic management aspect; being able to provision all of these formerly disparate components as a single entity. Collecting that list

of items that they simply must do and items that they want to do will be critically important. You need to be sure that you are drawing a solid line rather than dotted line to how they are going to benefit. Typically, when you ask what the benefit is, it's saving TIME, not money. Money to most IT folks is rather nebulous and abstract. Time, on the other hand, readily translates into weekends, nights, missed soccer games, birthday parties, date nights... the list goes on. If you can save time for these folks, you instantly become a hero. So here, when talking about management of the solution, make sure that you are circling back to the benefit of time.

- What do you do now on a daily basis, and what do you hope/think/understand that the new CI stack will do?
- What do you do to provision a user, and how long does it take? What about an application? A new company project? A new acquisition?
- How many different interfaces do you use and which are they?
- How do you collaborate amongst yourselves on tasks?
- What is the change management process once an IT project has been launched?
- How do you perform upgrades and refreshes?

By being tuned in and sensitive to their daily activities, you not only introduce the solution as a real solution to the people who are going to make it work, but you may also start the “you know, I always thought it would be a great idea if we could do XYZ. It could save the company a ton of money, because ABC process does nearly the same thing” conversation. Uncovering their pain points can lead you to management options either within the proposed converged solution that perhaps may have not been discussed, like being able to virtualize further up the stack so you can move VMs, VSANs, and users dynamically to perform upgrades during regular business hours, or complimentary management solutions, like offerings that put all of your management interfaces on a single pane of glass, or your own, custom- developed offering.

- For remote users, you have management requirements that are unique to those mobile users.
- What does your dashboard look like to monitor traffic, manage user access, and mitigate potential hazardous traffic?
- How is that remote access dashboard being integrated into the management interface?
- How do you determine how people are going to access what data from remote devices, and how is that policy enforced?

- What bandwidth, both inside the enterprise and across the WAN, are you going to permit?
- How do you authenticate remote identity and what impact does that have on performance?
- How much of your workforce will be remote vs. local? Will it be all the time?
- If this is a new offering to your user community, are your network bandwidth, wireless connections, and ISP all ready to deliver to remote users?
- Do you have different lightweight directory access protocol (LDAP) rules for local vs. remote users? What are the odds that the desktop and the mobile versions of a given application will be up and running at the same time, and how do you plan on managing data input conflicts such as occurs when a remote user who has a company issued laptop starts the day off in the home office, but then steps out to lunch and wants to check email. You can be sure they didn't shut their laptop off when they left.

Then, there is the discussion around high availability and disaster recovery situations. How does the management fail over? In all honesty, it's not the hardware that is the critical business element; it's the application and the data. You need to make sure that is what is functional. We go back to the cell phone example. Do you really care what server you download that application from, as long as the app store is available? How annoyed did you get when your web-based email wasn't accessible? Take that and extrapolate it to a business-critical application. Does a business truly care where the data resides, or do they care if they can get to it? In some instances, it does indeed matter, from a physical security perspective. But typically, availability of the applications and data is the key. You still need to manage that data and keep the business functioning at optimal capacity, even in the event of a failure. Frequently, the data and application fail over, but the management interface does not. Thus, in the event of a failure, legacy rules and policies are being enforced and implemented, which may or may not be acceptable. By stepping back and looking at the solution delivery from the management perspective, you continue to gain trust and buy-in for success.

It is a good idea at this point to circle back to all of the resulting quantifications derived from the previous five layers of discovery and get their input into your plans. You may discover someone in-house who is willing to pull new power drops, rather than needing to hire an electrician. You may discover someone who has legacy certifications in a particular network or storage switch that you had previously discarded because of perceived lack of customer experience. By engaging the people who will be ultimately responsible for the maintenance and upkeep of the

converged solution long after you walk out the door, you improve your chances of creating the “perfect” solution for the customer, and reduce the risk of post-order and post-implementation rework.

Layer 7 (Application) Applications

Applications, though most likely overlooked, especially need to be addressed before deployment. The applications are what keep the business up and running and are what the users see. They are the windows to your data center, so you want to make sure that your migration from siloed to converged is transparent to end users. Once again, the basics are the best place to start.

- Are the applications converged-ready?
- Can they be virtualized?
- Can they take advantage of multi-cores, multi-threads, or even multiple processors? If you introduce virtual machines, does your application licensing structure and cost change? (If the answer to this is “yes” usually it is for the worse, so you really want to make sure you have a grip on application licensing and rules)
- Do we need to set some resources aside for legacy applications that are critical and cannot take full advantage of the CI stack?
- Do we have to carve out parts of the converged solution for applications that must be physically separated from the rest of the resources, meaning that those resources are unavailable for the pool? This means you may have to go all the way back to cabling and networking.
- Are there bandwidths that you have to set aside for those legacy systems? Once identified, those applications to be virtualized and pushed out to end users typically fall under the catch-all umbrella of VDI, so as you launch that VDI environment, keep going back to the basics.
- What does the VDI look like? How are you going to implement it? Will it be a virtual environment that a user launches as a window on their own laptop, or will you issue laptops?
- Is the mobile version of your production applications different than the desktop versions, and how?
- If you have remote users, is there a different hardware pool you need to create to support the mobile version of the app?

Just from a generic, application delivery perspective, you will want the implementation of the converged solution to have as little impact as possible on the end users' daily work lives. Take some time to really dive into how the applications may act differently in a virtualized, converged environment.

You may also want to look at specific environments, such as test/development, as physically separate from production. A Development/Operations (DevOps) environment, for example, can really benefit from the reality of the converged infrastructure environment. By decoupling the hardware from the applications and the users, you are enabling a great amount of agility in your collaboration and test provisioning your environment no matter the physical location of your users. When something needs to be tested, you assign resources to it. When something needs to be rolled back and triaged, you reallocate resources or pull up the old profile to replicate the error. Converged solutions allow your resources to be available to you at will, so you can run your prototype on an environment that matches production and recreate environments for more accurate testing. However, you do need to have an accurate picture of how application development happens with your customer and the expectations of how similar a development environment needs to be to the actual operations and production environment to properly quantify for maximum customer benefit. Answers to how new applications are developed and how existing applications are tweaked and upgraded, and how applications, both new and upgrades/refreshes are rolled out to the enterprise, will effect nearly all of the CI stack.

- Do you need multiple versions of applications available to your testers?
- Once you have gone through testing to your satisfaction, how are you going to roll out the new applications or upgrades? Will they be staged, or will the rollouts be all or nothing?
- What are your plans in case of failure?
- How long will you allow access to older applications?

Once you determine that refresh and rollout strategy, you will want to circle back to the very beginning of our questioning stack model.

- Are you going to offer the same network bandwidth, or cut bandwidth down to legacy applications?
- How many compute resources will you dedicate to legacy vs. new applications and for how long they will be dedicated?

- How much storage and for how long you plan to store data in legacy application format?
- Will you continue to offer security to legacy apps at the same level as you do for the new?
- Will the data created in the new application have the same or different security levels?
- How will the data be accessed and by whom and for how long?
- Will there come a time when only admins or lawyers can access the old data?
- What time of day do you roll out applications, and are the new/improved tools going to impact performance at certain times of the workday?
- Is this performance impact something that needs to be addressed and mitigated by the converged solution?

This line of questioning highlights your interest in the continued success of how the converged solution will support the customer's business as it grows and adapts to market changes. By circling back to the beginning of the questioning stack, and keeping the flow of the logic in mind, you can organize your thoughts around quantified decisions and helping the customer succeed.

Line of business (LOB) applications also need to have questions answered to make sure that the CI stack is ready for those applications. The goal of a LOB discussion is not necessarily to upend the entire infrastructure, but to recommend improvements or changes to it so that the business elements are improved; implementing Quality of Service (QoS) differences to browsers vs. shoppers or granting customers credit faster, that sort of thing. One key question to keep asking throughout your analysis of LOB applications is "what if this doesn't work?" Not the whole solution necessarily, but what if a particular component isn't sized properly, doesn't function efficiently, or flat out isn't there? Help the customer think through the cascading effect of that failure or omission and you may expose a downstream opportunity. At the very least, you get the customer to think about their business application at a deeper level within the context of the converged infrastructure than they may have originally. You can then architect the previous quantified elements to mitigate around downtimes and drops in performance, so the result of the LOB initiative is positive. "Architect to mitigate", if you will.

- How are the different business goals dependent on each other?
- What happens if one business goal fails or is just not implemented because of cost or time?

- Are there one or two business goals that are critical above all others? Another way to ask this question is how much money will you lose if your data is lost for a day? Or an hour? Or if customer access is lost for a day? Or if customer access is slowed down by 50%?
- Who is involved directly and which functional areas are both directly and indirectly impacted? Finding out who is involved and who is going to be impacted is important.

Improving an LOB application sometimes has the unfortunate result of making something else more difficult. Not because the initiative failed, but because of the other things or functions that those directly impacted individuals may do. For example, if the LOB initiative is going to reduce bandwidth needs and free some up to run more critical batch jobs, is there enough horsepower on the local desktops to assume the offload? Again, get your customer to think in a cascading fashion. What is the result of the action going to do? Since Line of Business by definition, are the critical applications that are vital to running an enterprise, LOB opportunities force you to not only dive into the immediate details of the LOB application, but to follow the packets if you will, and make sure that you are continually looking at the solution holistically as it fits into the enterprise and how it impacts the enterprise. Lastly, what is the rollout plan for the improvements? The reason you are asking this very key question is so you can determine when your tweaking opportunity may arise. You may think you will be involved in the implementation and hear the feedback immediately, but then again, you may be turning your entire sizing and planning information over to the customer for their own implementation team.

Answers to these questions may cause you to go back and re-size some of the previous six layers of quantification. It's important to have these discussions now, not later. The more questions you ask, the better prepared you will be to act immediately on unintended and undesired results.

In summary, here are the seven steps to quantifying a converged solution.

- Layer 1 – Physical: address power, cooling, cabling and physical aspects of the environment
- Layer 2 – Network: address network topology as it pertains to the converged solution
- Layer 3 – Compute: address virtualized compute requirements
- Layer 4 – Storage: address virtualized storage requirements
- Layer 5 – Security: address local and remote security requirements
- Layer 6 – Management: address converged management requirements
- Layer 7 – Application: address virtualized application requirements

Which brings me to the final point. After you have gone through all seven layers of pre-sales analysis, you need to, wait for it...Write. It. Down. While that may seem obvious, it is not. Far too many pre-sales engineers have checklists floating around in their heads, rather than written down in a notebook. Thorough thought and planning through the entire pre-sales phase is critical to designing the proper solution. However, if it is not written down in a pattern that can be followed, it cannot be taught and cannot be replicated in a consistent manner. The more you can replicate, the easier it is to train others, and the more consistently positive your customer experience is with your go to market strategy.

Conclusion

This is not intended to be an exhaustive and unchangeable list of questions. There clearly are points that you specifically raise with every customer that are unique to your offering. However, you can use the structure of the seven lines of questioning correlated to the 7 layers of the OSI model to organize your own pre-sales process so you don't neglect any aspect of a converged solution offering. Add your own or your team's specific points of offering into the layer that you deem appropriate. Customize your questions to best suit your business model and service offering. Using this model to quantify and document your solution offering in an organized, consistent and repeatable manner will reduce your risk of overlooking critical features of the solution. Because you are looking at the solution from a higher, holistic level, while maintaining your technical credibility in asking pointed and focused questions around each facet of your converged infrastructure offering, you improve the validity of the offering. This is precisely why you were brought in to the customer in the first place; to provide quantified validation to the proposed solution. With an organized, repeatable, documented plan in place, you can walk a customer through the entire logic of their enterprise, making sure that the solution you are offering truly is a solution, solving that customer's particular business need, and genuinely becoming a trusted advisor.

Certifications held

- EMC VNX TA
- VMware VCP
- Cisco CCDP
- DC CCNP

EMC believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." EMC CORPORATION MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.