



Presented at the
ISPE – ISA Automation Forum - 2010

Living in a Virtual World with Tangible Results. Virtualization Overview and Real World Case Study.

Author Name	Andy Robinson (Avid Solutions) Bill Porter (APP Pharma)
Title	Living in a Virtual World with Tangible Results. Virtualization Overview and Real World Case Study.
Company	Avid Solutions/APP Pharma
Address	3471 Apex Peakway
City/State/Country/Postal Code	Apex, NC 27502
Telephone	919 468 4334
Fax Number	919 468 4335
E-mail	arobinson@avidsolutionsinc.com bporter@appharma.com

KEY WORDS

Virtualization, VSphere, VMWare, Case Study, Disaster Recovery

ABSTRACT

This presentation will deal with a real world pharmaceutical case study of a local operating company that has implemented a virtualized environment for their validated computer systems associated with aseptic filling and production environments. Bill Porter, Facilities Maintenance and Engineering Manager of APP Pharma in Raleigh will present the business side of case study and provide the audience the benefits and issues they have encountered using a virtualized approach for their information management infrastructure. Andy Robinson will present critical technical aspects of the project along with some detailed discussion on the decision making process in technology selections and sizing for the project.

Introduction

Earlier this year Avid solutions installed Phase 1 of a two part project to virtualize the majority of APP Pharmaceutical's manufacturing computer systems at their Capital Boulevard facility in Raleigh, NC. The intent of this paper is not so much to evangelize as it is to give a nuts and bolts view of why's and how's of virtualizing manufacturing systems in an operating facility. Just as with a project to purchase new hardware the requirements to establish a business justification is just as relevant with a virtualization project. Bill Porter from APP Pharma will discuss how he was able to make the business case for the project. Andy Robinson from Avid solutions will then go into great depth discussing the decision processes that led to software and hardware selections and design for the system.

Business Justification

The APP facility has a number of disparate SCADA and HMI systems distributed throughout the site. As is typical with most pharmaceutical installations many of the installations are aging and in need of some form of hardware upgrade to continue reliable operations into the future. The seed idea behind installing a virtual infrastructure started with the suggestion of replacing older operator workstations with new hardware. Unfortunately the time and expense to prequalify the new operator workstations was a challenge to any project financials, especially if the workstation itself was still functioning properly. An alternative approach to a complete replacement surfaced in the form of creating virtual images of the workstations and then running those virtual images on the newer hardware in the field. While discussing the options for operator workstation replacement the idea of going one step further and virtualizing the SCADA servers was considered. In addition to replacing individual control systems, virtualization of the systems on a server level would have the benefit of central backups and built-in redundancy. This would give the site disaster recovery solutions by prevention of possible failures with redundancy and virtual images of systems that could be restored in the event a failure occurred. This change also would relocate systems that were difficult to access in the production environment to a central accessible location in the facility. With a number of aging servers and an impending major upgrade to the Building Management SCADA system on the horizon it appeared that full blown virtualization might be a cost-effective means to replace the old hardware while minimizing risk to the existing systems as well as dramatically upgrade disaster recovery capabilities.

Choosing the Virtualization Software

To be fair there wasn't really a long detailed discussion of what virtualization solution to utilize. Our experience as integrators started with VMware Workstation somewhere around early 2007. During that year we experimented with both VMware Workstation and Microsoft's offering at the time, Virtual PC. It was our experience that VMWare Workstation was a far superior product in terms of stability and raw performance. As the team of engineers outgrew using VMWare Workstation on their individual machines Avid installed a VMWare ESX Server in late 2007. At the time Microsoft did not have an equivalent offering. In the years since Microsoft has introduced their competing product, Hyper-V. However, based purely on reading reviews and comparisons it appears that VMWare VSphere is still the clear leader in the market. Also, since the project was not installing a data center full of hosts, the actual unit cost of the software was a pretty minor player in the overall cost of the installed solution, less than 10%.

Sizing

No matter the project, whether it be installing a new PLC or putting in a new network, you should always start with trying to quantify how much of each item you'll need. Planning a virtual infrastructure is no different. For our infrastructure we needed to decide on items like host count, memory, storage space, networks, network interfaces, and power to name a few. Our first step was to catalog all of the machines the client thought they might virtualize over the next few years, including new machines they may want to install. This yielded an approximate makeup of machines below.

Category	Quantity
Process Control Servers (Historians, I/O Servers, etc.)	5
Operator Workstations	8
IT Type Servers (Domain Controllers, Backup Servers, etc.)	5

Table 1. Machine Categories

At first blush you may want to take a look at the categories and throw some general numbers at each category to come up with total requirements. The only problem with that is that you could have some major difference in requirements between systems in the same category. A prime example is that you are going to need much more hard drive space in a Historian than you will a typical I/O server. On the flip side depending on how much I/O you're pushing you may need more memory in your I/O server than your Historian that's historizing 500 tags. This led us to look at each machine individually and identify ballpark numbers for hard drive space and memory. A listing of those results is below (with names genericized for obvious reasons).

Machine	Service	Space Required (GB)	Memory Required (GB)
Server 1	Historian	200	2
Server 2	Historian	100	2
Server 3	Historian	100	2
Server 4	Historian	100	2
Server 5	I/O Server, Tag Server, Alarm Server	50	2
Workstation 1	Dedicated Operator Workstation	20	1
Workstation 2	Dedicated Operator Workstation	20	1
Workstation 3	Dedicated Operator Workstation	20	1
Workstation 6	Dedicated Operator Workstation	30	1
Workstation 7	Dedicated Operator Workstation	30	1
Server 6	Backup Exec Server	20	2
Workstation 8	Dedicated Operator Workstation	20	1
Server 7	Existing Print Server	10	1
Server 8	New Sharepoint Server	50	2
Server 9	New Domain Controller	10	1

Machine	Service	Space Required (GB)	Memory Required (GB)
Server 10	New SMTP Server	10	1
	Totals	830	23

Table 2. Listing of Machines with Resources

Something I see quite often is the casual thought that you need to size your XP workstation VM to 30 or 40 GB. I feel like this really is an artifact of the fact that these days you can't get hard drives much smaller than 250 or 500 GB so we forget how much room the OS really uses. Also, as long as we are only putting the required software on the machines, which typically isn't going to be much more than a couple gigabytes we can usually easily get away with a 10 GB hard drive for a typical workstation install. In a worst case scenario if you begin to run out of room the virtual infrastructure gives you multiple ways to expand a hard disk after initial install. We are also taking credit for VMWare's advanced memory algorithms and memory overcommitment technology. Because the machines don't actually use all of the memory allotted we can comfortably plan for more physical memory than is actually available on the machine.

Also, one thing you may pick up is that there are no workstations 4 and 5. During our initial planning we had included these in our configuration. However, we quickly realized that virtualizing these workstations wasn't practical as they utilized an ARCNet interface. This is a big red flag to watch out for when virtualizing machines, especially older ones. With the new VMDirectPath technology in VSphere we may have been able to make these machines work but for the time being we're going to leave these machines alone and revisit them in the future if the architecture of the target system is upgraded to utilize ethernet interfaces for the operator workstations.

These numbers are a good start but we would really be doing a disservice to the system if we only planned for today's requirements and didn't build in any capacity for future growth or more importantly underestimation of real needs. Thinking through this we went through each of the machines and assigned a future growth factor for the hard drive space. We didn't worry about memory too much because typically your memory requirements are somewhat stable. Usually if you are adding a large amount of load to a system you'll end up installing another server in parallel. I'll agree that this a pretty gross generalization but across the majority of servers this will typically hold true.

Machine	Service	Growth Factor	Planned Space (GB)
Server 1	Historian	100%	400
Server 2	Historian	100%	200
Server 3	Historian	100%	200
Server 4	Historian	100%	200
Server 5	I/O Server, Tag Server, Alarm Server	100%	100
Workstation 1	Dedicated Operator Workstation	25%	25
Workstation 2	Dedicated Operator Workstation	25%	25
Workstation 3	Dedicated Operator Workstation	25%	25

Machine	Service	Growth Factor	Planned Space (GB)
Workstation 6	Dedicated Operator Workstation	50%	45
Workstation 7	Dedicated Operator Workstation	50%	45
Server 6	Backup Exec Server	50%	30
Workstation 8	Dedicated Operator Workstation	25%	25
Server 7	Existing Print Server	25%	12.5
Server 8	New Sharepoint Server	50%	75
Server 9	New Domain Controller	25%	12.5
Server 10	New SMTP Server	25%	12.5
	Totals	830	23

Table 3. Future Growth Planning for Hard Drive Space

The next thing we need to consider is planning for machines that we don't even know about yet. To be on the safe side we applied a 50% factor to both the hard drive and memory plan. This yielded a total hard drive requirement of approximately 2.1 TB and a memory requirement of 35 GB. The two numbers will be used in the planning for our Storage Array and the amount of physical memory we will be ordering for our physical host.

The next aspect of the design involves planning out the networks. Because we will be using shared network storage our first decision was to utilize a completely dedicated storage network. Because backend storage is so critical to the proper operation of our system we also chose to install redundant switches in our storage network. In combination with network redundancy at the SAN and VSphere host level we have greatly minimized the possibility of a single network failure affecting the availability of our network storage device. For the communication network side we chose to install a single gigabit switch. There were a couple factors that led to this decision. First, the current network layout consisted of both PLC's and workstations communicating across the same physical network. While this is not an ideal long term design, our project wasn't going to attack this particular item. Also, the balance of the physical infrastructure was not in place to support a redundant backbone. If an upgrade takes place at a later date, upgrading to redundant physical switches is a relatively simple operation.

Once the networks were laid out the next task was to determine the network interface requirements for our host boxes. We started by deciding we would have redundant connections to our virtual machine network. From there we decided to utilize 4 network interfaces to connect to our storage network. The long term plan for this environment is to add more hosts to create a high availability and possibly fault tolerant cluster. Proper design of these clusters involves a dedicated network connection on each box for what is called VMotion traffic. Considering the relatively low cost of gigabit adapters we ended up installing 3 dual head NIC adapters. In combination with the provided 4 onboard gigabit NIC's this yielded a total of 10 gigabit network interfaces. This is one area where we should never run short of resources on our host.

One subtlety that we tend to think about in our process control designs is minimizing the impact of a single component failure. We also need to consider this in our network design. To that end when deciding which NIC's to use for a specific function we were careful to distribute the redundant pairs across the hardware. This means that we never use both ports on a single NIC for a redundant pair.

Instead we choose ports from different physical NICs for our redundant pairs. We also are careful to not utilize multiple ports from the built-in NICs for redundant pairs. It's a minor design element but it should provide for an even more robust system.

Storage Hardware

Without question the most critical element of any virtual infrastructure is your storage system. If you make a poor choice for host computer hardware or undersize your host hardware it's a relatively simple and comparatively inexpensive process to install new nodes in the system to distribute load. However, if you make a poor choice for your back end storage the implications are much more dramatic both in cost and potential downtime. More fundamental than migrating storage is simple availability during regular maintenance. Lower end storage arrays generally don't provide for online maintenance. Because these systems aren't designed with the intent of supporting 24x7 mission-critical operations many of their functions require reboots or taking the array offline. In general you can break network storage devices into 4 major classes. The first class of network storage devices typically uses single controllers and commodity class 7200 RPM SATA drives. These devices will have multiple ethernet interfaces for network redundancy along with redundant power supplies. However, the lack of redundant storage controller poses a risk not only if a controller was to fail but also effectively prevents any firmware upgrades while the system is running. The second class of storage devices typically upgrade to redundant storage controllers and faster enterprise class SAS hard drives. One of the more critical upgrades you'll get in this second tier of device is the ability to manage arrays on the fly. This was an especially critical feature in the design of our system. As the plan for this system developed it was decided that the installation would be performed in multiple phases. This required the installation of 4 300 GB SAS drives for primary storage and 3 1 TB drives for storage of backups during Phase 1. Phase 2 would install more 300 GB SAS drives to expand our primary storage. We would also be adding a second drive chassis and expanding our backup storage. Because we chose a slightly higher end SAN we would be able to not only expand our primary storage array on the fly without interfering with system operation but we would also be able to add a new chassis and migrate our existing backup drives between chassis. The third class of storage arrays begins to add more sophisticated software capabilities such as volume snapshotting and deduplication. Volume snapshots provides a mechanism to "save" the contents of a volume at a specific point in time. Then, at a later date the snapshot can be "mounted" to recover data that has been either lost or overwritten. Deduplication is a relatively new technology with some incredible implications for data storage. In its simplest form deduplication analyzes each block of data you are attempting to store and determines if it has already stored a matching block. If it has, the system will store a pointer to the matching block instead of storing the same block twice. Think for a moment how many blocks of data are probably duplicated across your windows environments. Even further if you are running 15 operator workstations all with the same software installed chances are they are about 90% the same. It is typical to achieve a 15-20x reduction in storage space via the use of deduplication technology. The fourth and final class of storage arrays dramatically increases the overall capacity to petabytes along with even more sophisticated software capabilities. As alluded to earlier in the description of these classes we chose a storage array in the second class. The ability to perform maintenance on the fly is critical for system that is expected to run without fail for years. We all know how difficult it can be to get system downtime for one computer. Once you put all your eggs in one basket with a single storage array the windows of downtime shrink exponentially. The cost factor between the second and third class of array for similar storage capacity

can easily be on the order of 3x to 4x. The additional software features did not support that kind of increased expense in a project such as this.

Storage Protocol Selection

Storage protocol selection can be a very complicated process with detailed analysis of workload, throughput, and management requirements. However, considering the scale and loading of our system the subtle differences between the two major protocols of choice, ISCSI and NFS, simply didn't warrant the protocol itself driving our SAN storage choices. The storage array we chose utilized the ISCSI protocol. Although not quite as easy to configure as NFS, the complexity of setting up ISCSI LUN's formatted with the VMFS file system were a pretty minor activity in the overall list of tasks the project required. A third protocol, Fibre Channel, was never seriously considered for the simple fact that Fibre Channel arrays are typically on the higher end of the spectrum and are also notoriously complicated to setup and maintain.

Backing up the system

Invariably the first or second question that is usually asked when discussing virtualization of physical computers is how you are going to back up the system. While unfamiliarity may make some people shy away from virtualizing their manufacturing environment I believe it is actually the most powerful argument for virtualizing your environment. The concept of agentless backups is a dramatic departure from having to purchase and maintain some kind of software agent on all of your physical servers. Almost all modern virtual infrastructure backup software utilizes an agentless scheme. There are a few key differentiators between some of the offerings. The first, and what I think is the most important differentiator, is the ability to utilize deduplication when constructing backups. As discussed earlier in this paper, deduplication can typically yield between a 15 and 20x compression factor for a relatively homogenous environment. This incredible compression allows for an unexpected benefit, the ability to keep many more backups than previously possible. A typical setup might be to keep the last 60 days of backups online and accessible at all times. Past that you may keep 1 week backups for the previous 6 months, then monthly backups past that. This is a granularity that you typically don't find using standard physical backup techniques, especially if you have limited storage capabilities. The second important differentiator between offerings is the use of "virtual appliances" to perform backups. Some software packages read and write data via the windows server on which the backup software is running. In a virtual environment this is very inefficient as it can effectively double the apparent traffic transmitted over the network. By spawning "virtual appliances" the more advanced packages offload all processing and network traffic to small specialized virtual machines running independently from the machine used for configuration of the backup jobs. This offloading mechanism allows the systems to scale much easier as your virtual infrastructure grows.

Now that we have daily backups running and our data is safely stored we need to backup our backups. Backups are no good if you end up with a disaster that takes out your entire rack or even worse your entire server room. There are two real options for guarding against this type of failure, both relying on the basic idea of backing up your data "offsite". One option would be to replicate all of your backup data to another storage device in another building. This method was not a practical option for this particular project. A second option is to utilize the old standby of tapes. While the technology may be dated the utility of being able to write large amounts of data to a media that's meant to be handled,

transported, and stored offsite is a perfect complement to our chosen backup scheme. The next choice involved backup frequency. Because tapes would only be taken offsite once a week there was no practical reason to perform tape backups more frequently. The combination of daily backups with replication to tape on a weekly basis provides an extremely resilient backup scheme that ensures even in a worst case scenario the virtual environment can be restored to a point in time less than 1 week old.

Conclusions

In the past when I would teach people about the concepts of basic process control I would emphasize the idea that all systems have variability and that variability has to be accounted for somewhere. This typically results in that variability being accounted for via the actuation of a valve or changing level in a tank. The complexity in computer system designs is no different. There is always a fixed amount of complexity that must be accounted for. In physical server environments this complexity is added and addressed slowly as you add new machines. There are typically break points where you need to add new network switches, new power sources, or new cabinets. Virtual infrastructures don't necessarily reduce the complexity, they simply demand much more planning up front to create a resilient system that can support 24x7 operations in a manufacturing facility. However, once that heavy lifting is complete the hurdle to adding new machines or maintaining existing systems is dramatically lowered.