



**BEST OF IT TRANSFORMATION
EMC WORLD 2015**

Ten ways to reduce cost while modernizing your IT

May 2015

V1.2

CHRISTOPHER KUSEK
CHIEF TECHNOLOGY OFFICER
XIOLOGIX, LLC.
CKUSEK@XIOLOGIX.COM
(888) 492-6843 EXTENSION 280

DISCLAIMER	3
WHERE DO WE GET STARTED.....	5
VIRTUALIZATION	5
CONSOLIDATION	6
BEYOND CONSOLIDATION.....	7
<i>Oversized virtual machines</i>	7
<i>Undersized virtual machines</i>	8
BUSINESS CRITICAL APPLICATIONS	9
<i>Reducing Cost with Business Critical Applications</i>	10
MODERNIZING YOUR VIRTUAL DATACENTER.....	12
STORAGE CONSOLIDATION.....	14
BACKUP.....	16
DISASTER RECOVERY AND BUSINESS CONTINUANCE.....	18
DISASTER RECOVERY.....	18
BUSINESS CONTINUANCE.....	20
CONVERGENCE	21
HYPER CONVERGENCE	22
CLOUD AND HYBRID CLOUD.....	23
SUMMARY	25
ABOUT THE AUTHOR.....	26
ABOUT XIOLOGIX	26

Disclaimer

These solutions are designed to help organizations address various datacenter modernization deficiencies. This document is intended to provide general guidance for organizations that are considering solutions to help them address such challenges. We encourage any organization that is considering solutions to engage appropriate legal, business, technical, and audit expertise within their specific organization for review of regulatory compliance requirements. It is the responsibility of each organization to determine what is required to meet any and all requirements. The information contained in this document is for educational and informational purposes only. This document is not intended to provide legal advice and is provided “AS IS”. We make no claims, promises or guarantees about the accuracy, completeness, or adequacy of the information contained herein. Nothing that you read in this document should be used as a substitute for the advice of competent legal counsel.

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.

Revision History

Date	Rev	Author	Comments	Reviewers
January 21, 2015	1.0	Christopher Kusek	Initial Release	Industry SMEs
April 20, 2015	1.1	Christopher Kusek	Editorial Updates	Industry SMEs
May 5, 2015	1.2	Christopher Kusek	Best of IT Transformation	Advisory Panel

EXECUTIVE SUMMARY

As it stands today there is a wealth of disparity between like and dissimilar organizations in how they run their IT organizations. Fortunately, where there has been disparity there has also been advancements made by vanguards of technology in the effort of data center modernization. Your organization may be adopting a few of the strategies discussed here today, you may be adopting some of these strategies, or you may be adopting all of them.

The intent of this paper is to provide guidance and direction on opportunities which exist in every data center to reduce operational and capital cost while also helping to extend the data center into the next generation of the modern data center. Although you may already have a firm grasp of some of the areas, I encourage you to review all of the items discussed as there will be breakouts which focus on elements of modernization in an effort to increase efficiency within those already existing investments.

The ten key focus areas which will be discussed throughout this paper are:

- Virtualization Consolidation
- Business Critical Applications
- Modernizing the Virtual Data Center
- Storage Consolidation
- Backup
- Disaster Recovery and Business Continuance
- Convergence
- Hyper Convergence
- Cloud
- Hybrid Cloud

While some of the solutions discussed may be the perfect match for your organization, no one solution or a combination of solutions will be a panacea for all organizations. Thus, it is important to align the data center benefit to the business need to avoid a solution mismatch. Last, these solutions do not have to be followed in a serial, one after each other fashion. Each item can and should be treated independently unless explicitly stated otherwise.

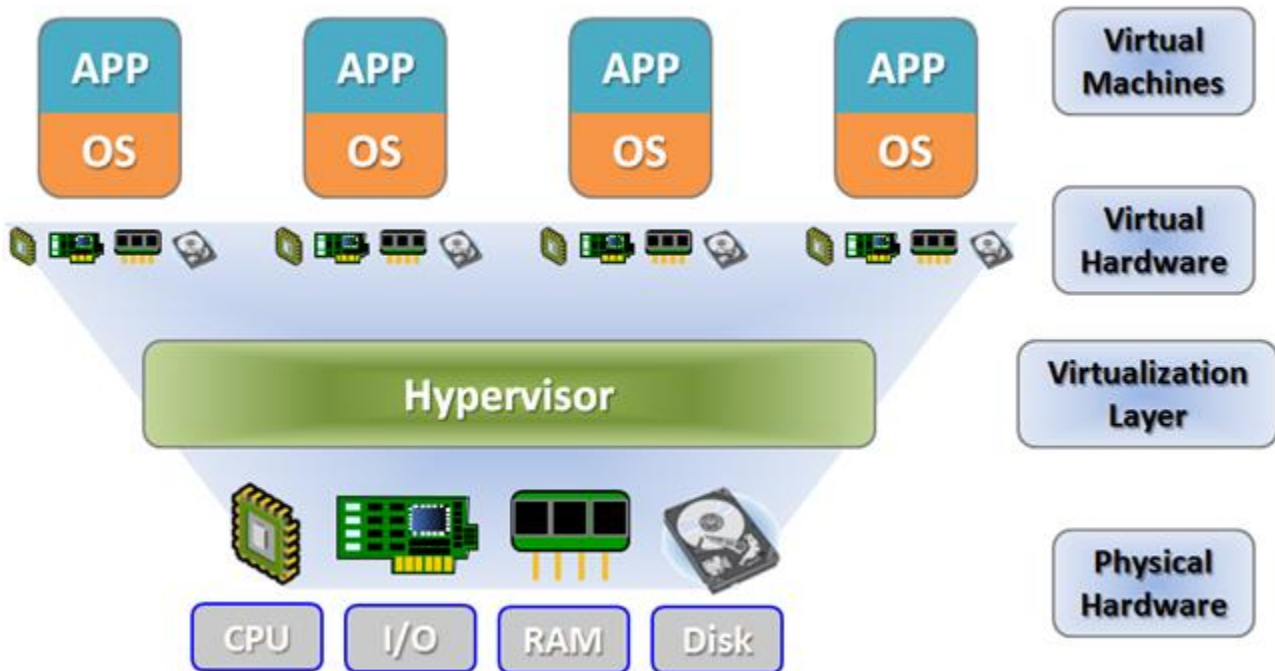
Where do we get started

While doing the research and analysis for this paper in ways to have a real impact both on the technology, advancement, and benefit to the data center while also providing value through reduced cost; the merits of Gartner, IDC, and other industry-leading analysts was considered, and then promptly thrown out the window. For what it is worth, you don't need a Magic Quadrant or a high-priced whitepaper from an analyst to tell you that paying less for modern infrastructure or doing things more efficiently than you are currently doing will result in cost savings. This leads us to kick off this section with something that industry analysts suggest that most organizations are already doing.

Virtualization

Virtualization is an ecosystem of software and hardware that come together to form an abstraction of the hardware and represents a departure from how computing has traditionally been done. It does this by providing a discrete abstraction of physical hardware components, CPU, memory, storage, and network. While virtualization isn't new – it's been around for nearly two decades in the x86 space – and adoption has progressively been on the rise, efficiency within the virtualization space has been below par for the most part.

Figure 1: A basic explanation of Virtualization from Hardware through to Virtual Machine



Consolidation

Traditionally, Consolidation is the first phase of virtualization adoption. During this phase, CapEx and OpEx savings are directly realized through the process of consolidating physical servers typically operating at 5-10% efficiency onto fewer servers as virtual machines (VMs) allowing these consolidated machines to increase server utilization upwards of 60-80%. Cost savings realized during the consolidation phase are typically experienced via:


- Reduced Power and Cooling requirements by decommissioning aging or legacy hardware
- Reduced Server maintenance and support agreements by fewer servers running higher workloads
- Reduced Rack Space utilization allowing for better allocation of the space allotted
- Reduced Network ports as more servers can communicate to fewer ports on the switches
- Reduced Storage requirements as more virtual machines are consolidated into data stores

Other benefits organizations will often realize include:

- Increased system utilization and uptime through high availability
- Easier testing and development with patching and snapshots
- Simplified system updates and migrations allowing for non-disruptive management of infrastructure
- Predictable high availability and disaster recovery
- Reduced deployment times for provisioning of new servers

While this isn't an exhaustive list, most organizations entering the first phase of consolidation when adopting virtualization tend to realize these benefits at excessive and exorbitant value.

Figure 2: Phase 1 of virtualization adoption – Consolidating and virtualizing non-critical applications



Beyond Consolidation

Some of you may be saying, “I’ve already virtualized and am well through the first phase of consolidation, so there’s no cost benefit to be had here”. However, this is where an important distinction of the first phase of consolidation must be understood, one which is often overlooked by no fault of business users and administrators.

“Physical servers do not virtualize well when they were oversized to begin with”

This is not to say that making a physical server into a VM does not work or does not work well at all. In fact, it is quite the opposite. When servers are converted from Physical to Virtual (A process referred to as P2V) the original server and all of its resources are typically taken at face value. Any physical server bought in the past 10 years will often be sized with two or more CPUs and will usually have a large amount of memory and disk space allocated to it. When that server, taken at face value, is virtualized it will operate wonderfully in the new virtual environment. However, the server is usually oversized for the requirements it calls for.

Note

Some virtual machines are not oversized and are undersized. These should be discovered and resolved.

Oversized virtual machines

By right-sizing virtual machines in your environment, you can start to realize real and actionable savings which can result in significant savings through a more efficient environment, reduced need to acquire new servers and storage, as well as providing better resiliency and availability to your existing virtual infrastructure. Through analysis of operational environments, the following virtual machines are typically oversized by default.

- Print Servers
- File Servers
- DHCP Servers
- Domain Controllers
- Syslog and Collector servers
- Security Scanner Appliances
- Web Servers

There is no perfect answer to what kind of resources should be allocated to servers in every environment as operationally, organizations will differ in their requirements and how they use applications, especially based upon the growth and size of an organization. Still, with that said, unless you have special justification and an aggressively high use of certain capabilities of applications, it’s difficult for me to imagine a Print server requiring more than 1 virtual CPU and 4GB of memory – and even that may be too high. This section isn’t presented to provide exact guidance of what is considered ‘oversized’ but more to recognize that there is a good chance that if un-tuned, you may have 90-99% of your virtual machines oversized by default and not even realize it.

The next item to discuss is not undersized VMs and Applications but instead Applications which are oversized but you may not think are oversized. This will vary based upon your applications usage, needs, and demands of users.

- SQL Servers
- SharePoint Servers
- Exchange Servers
- Oracle Servers

If left un-tuned and un-configured out of the box, applications such as these will consume all of the resources you provide them and more. It is important to not only monitor how your virtual machines are using their resources, but also to ensure that your applications are tuned to use them appropriately. SQL Servers by default will try to consume upwards of 2TB of memory, even if you only have 4GB allocated to them. This isn’t a problem unless your virtual machine tries to use the

resources allocated to it for other functions such as the operating system paging and so forth. Thus, when it comes to these higher critical applications you'll want to be cognizant of both the resources allocated as well as the configuration of the applications under the covers as they can mask the actual requirements of the system.

Undersized virtual machines

Undersized virtual machines tend to be a trickier beast to calculate for and understand. Interestingly, you may notice some similarity to the information above. Certain applications will absolutely be undersized due to the mere nature of how they function and how much in resources they're looking to consume. The applications below will traditionally jump to the top as reporting as undersized and impacting the ability of your applications to execute and perform their respective functions.

- Domain Controllers
- Security Scanner Appliances
- Web Servers
- SQL Servers
- SharePoint Servers
- Exchange Servers
- Oracle Servers
- Network and System monitoring servers and tools

You'll definitely notice a lot of overlap with the information above. The reason some applications made BOTH lists is because when certain applications are oversized they're considerably oversized, and when some applications are undersized they tend to be significantly undersized. One prescribed model is to supply applications with a uniform amount of resources, monitor for a set period of time, review performance and utilization of application resources and, if there is a need to allocate additional resources due to legitimate use, grant more. Some may argue that you should grant a large amount of resources and reclaim later if un-used. I welcome you to convince an application owner that they're not using their resources and shut it down to reclaim them. One of the benefits of virtualization has been the ability to add resources – even on the fly –when supported so as to increase capabilities within a server while reducing downtime.

With the servers above which are duplicated, it will really come down to getting a baseline understanding of what and where your applications needs are at and being able to apply those resources to them as appropriate. All servers and applications are not created equal; therefore you need not give every SQL server the same resources if their requirements vary significantly between them. It bears explicitly mentioning Security Scanner Applications here for a moment. It was mentioned above as oversized because certain applications tend to not operate too well in a multi-threaded fashion. When that is the case, additional vCPUs or Memory may result in no performance gain. Meanwhile, some security scanner applications, such as Nessus tend to operate far more efficiently with linear scale when additional vCPU resources are granted to them, so knowing the differences in your applications and where the benefits can be realized can result in a far better operational environment.

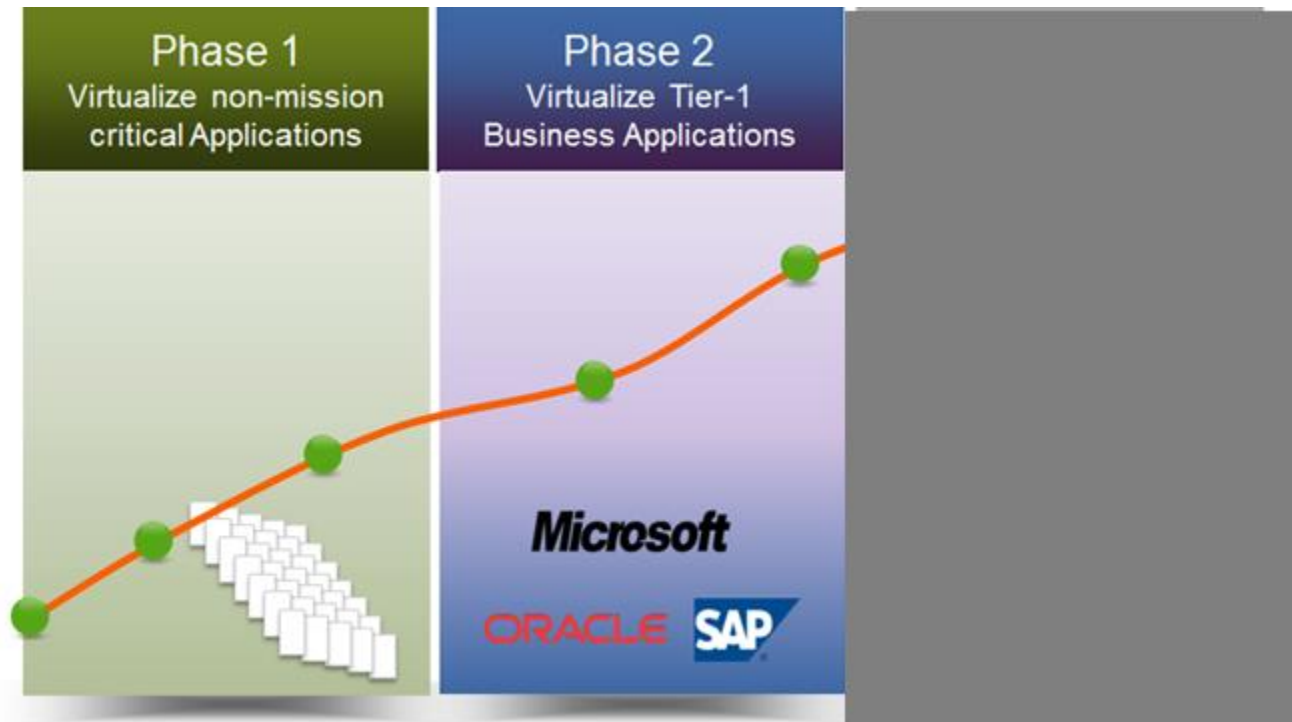
You may have noticed that I called out Network and System monitoring servers and tools. Of all of the most under-sized applications in the data center, the tools which can be used to collect information about what applications and servers are undersized or oversized tend to consume some of the most resources. So whether using a network monitoring tool which interrogates SNMP and drags Syslog data out of switches like SolarWinds Orion, NPM, etc.; or an application-centric monitoring tool which collects information directly out of the Hypervisor and their applications like VMware vCenter Operations (Renamed to vRealize Operations) or tools like VMTurbo or SolarWinds Virtualization Manager, these tools may require additional vCPU and Memory resources to keep up with the demands of data collection and constant analysis of that information.

While your organization may be starting the consolidation phase or be well beyond it, the need to regularly monitor your utilization and allocation and right-sizing where appropriate can be a major source of real cost savings and cost reduction for running your IT infrastructure.

Business Critical Applications

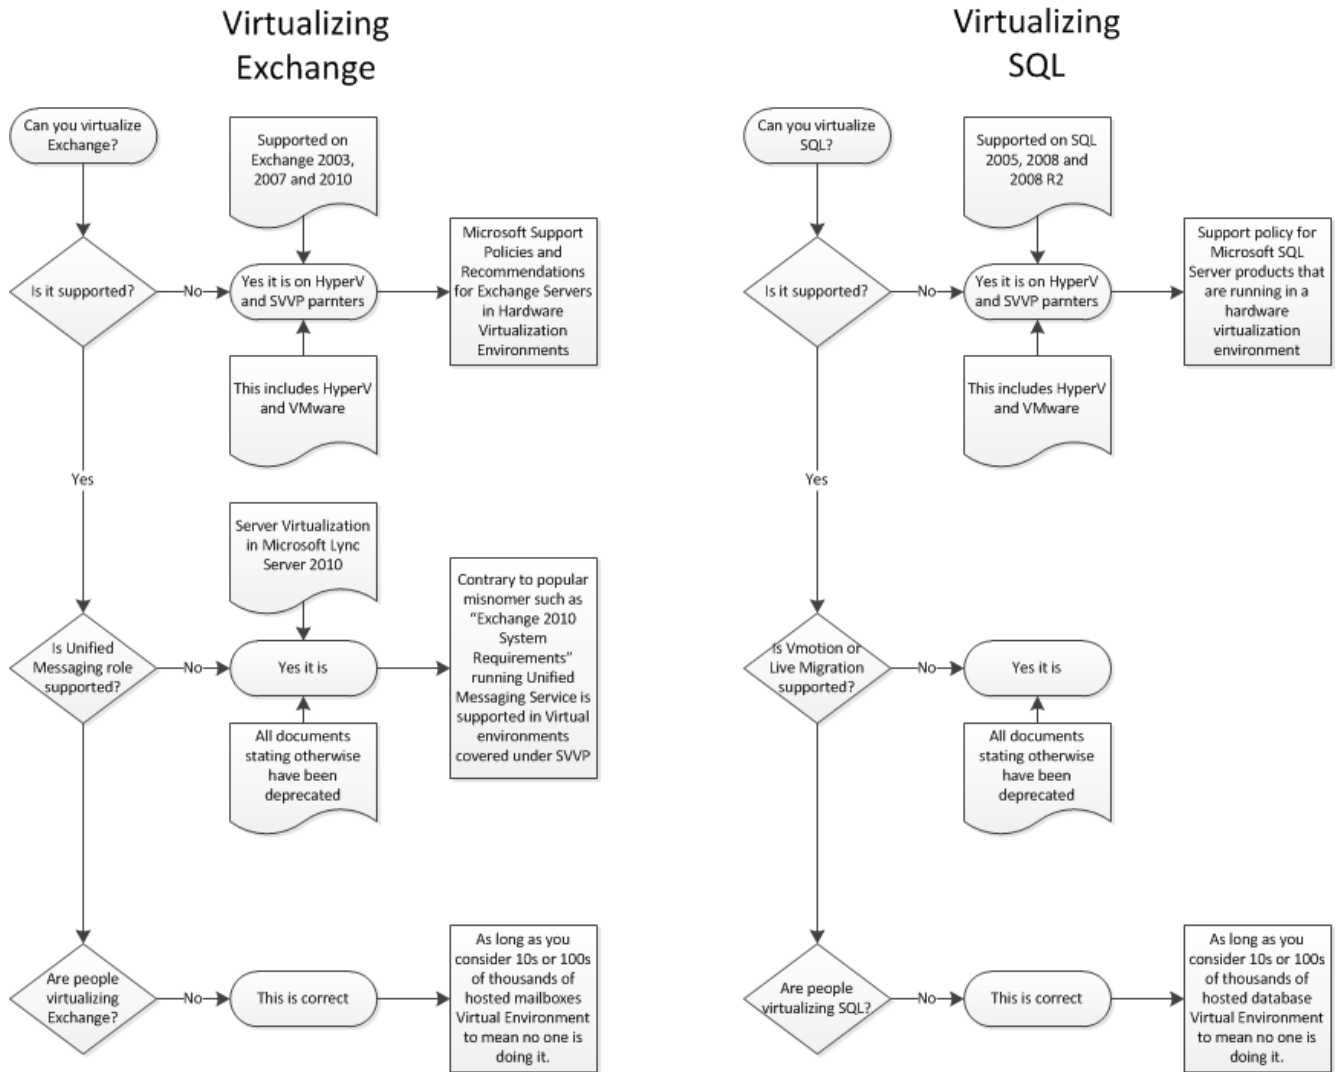
Moving up the stack of virtualization adoption as a means of reducing cost while modernizing the data center, we approach the second phase of virtualization adoption; Virtualizing Tier-1 Business-Critical and Mission-Critical applications. Any way you look at it, these are the important apps which your business runs on. Figure 3 shows some examples, but the applications most people associate with Phase 2 are COTS applications such as Microsoft Exchange, SQL and SharePoint, Oracle, SAP, and beyond. While custom-written applications can fall in this space, in the initial phase of getting comfortable while adopting Phase 2 it would be with those where companies have written agreements validating that their applications are supported in a virtual infrastructure.

Figure 3: Phase 2 of virtualization adoption – Virtualizing Tier-1 Business Critical Applications



It's important to reiterate the previous statement because the challenge has been brought up enough stating, "My application isn't supported to run in a virtual environment" an argument which was very true when brought up in the 2005. In 2015, that argument is no longer valid nor viable. Refer to the flowchart in Figure 4, published in 2011 showing what the support options for these applications at that time. Here in 2015, the support options only continue to improve, and the number of users adopting a virtualized platform for these applications grows at a steady rate.

Figure 4: Flowchart of support options for Virtualization of SQL and Exchange



Reducing Cost with Business Critical Applications

By now you're asking the question, "So you're saying that by virtualizing my business critical applications, I can reduce cost?" Yes and no. While there will certainly be some abilities to reduce cost which we will discuss, there are some capabilities which will not natively save you money. Examples of areas you'll be able to reduce cost include:

- Decreased downtime through utilization of vMotion for hardware maintenance
- No longer requiring complicated host clustering like MSCS or RAC by leveraging High Availability (HA)
- Ability to adopt a testable Disaster Recovery plan through using tools like SRM
- Ability to adopt a Business Continuance solution with Long Distance vMotion or Fault Tolerance (FT)
- Ability to increase allocated resources on demand with Hot-Add CPU and Hot-Add Memory
- Limited availability to reduce number of licenses by putting more users on fewer servers*

This is by no means a definitive list of every possible way you'd be able to reduce cost of your data center investment. However, if you're not utilizing these means today they'll often provide a value-add which otherwise was non-existent or extremely expensive to your organization, all the while helping to modernize your data center and applications.

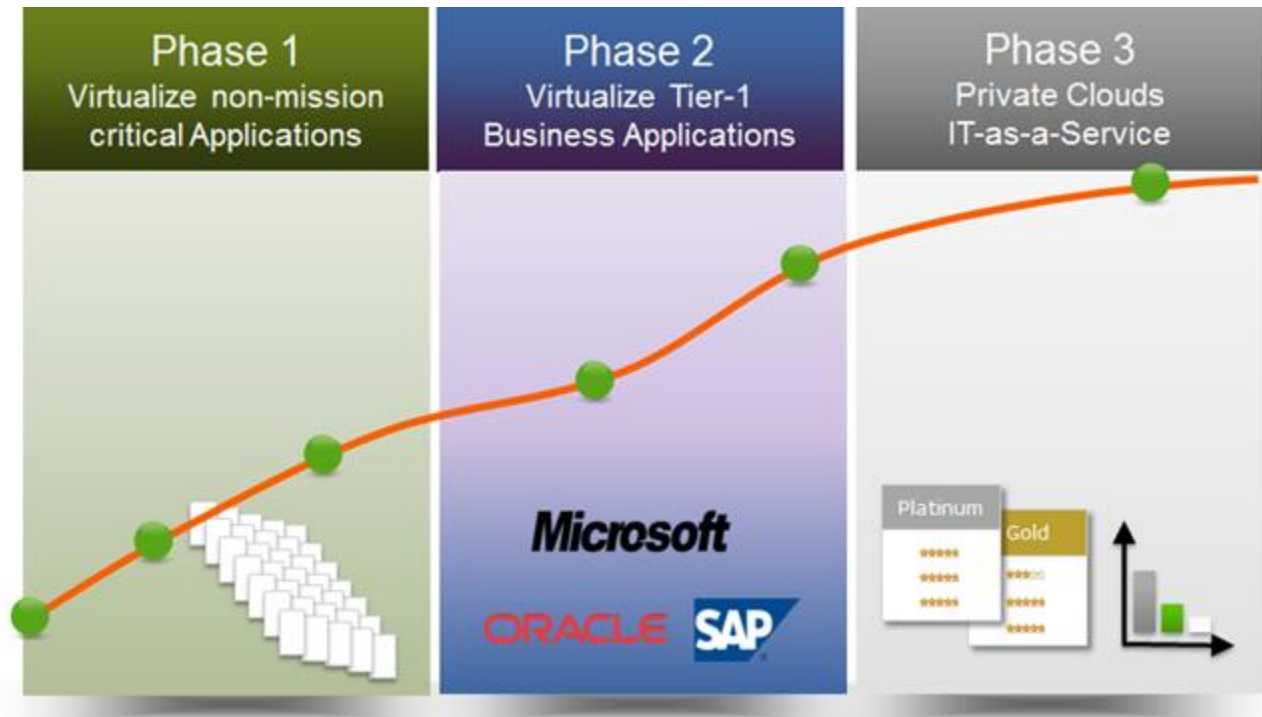
While I could go on about additional ways to reduce cost and gain benefits for your business-critical applications, there are numerous other sources which cover this area at length and in detail, including but not limited to:

- [VMware vSphere Performance: Designing CPU, Memory, Storage, and Networking for Performance-Intensive Workloads](#)
- [Virtualizing Microsoft Business Critical Applications on VMware vSphere](#)
- [Happy Virtual Holidays; Best Practices for Virtualizing Mission Critical Applications, Storage and Hyper-V!](#)

Modernizing your virtual datacenter

Phase 3 of virtualization carries with it many names; IT as a Service, Private Cloud, your personal Amazon, the data center portal, etc. While it can go by many names, they're all right in the right organization, and all wrong in the wrong organization. There are often many attributes which people will associate with this phase of the virtualization journey and there are entire business models driven around operating in this capacity. This is by no way a solution that every organization should strive for. While some of the capabilities this phase of adoption brings will be very valuable to some organizations and how they run their business and IT, other organizations may find them to have no function or purpose.

Figure 5: Phase 3 of virtualization adoption – IT as a Service and Private Cloud



Some of the common attributes associated with this phase of virtualization adoption happen to be:

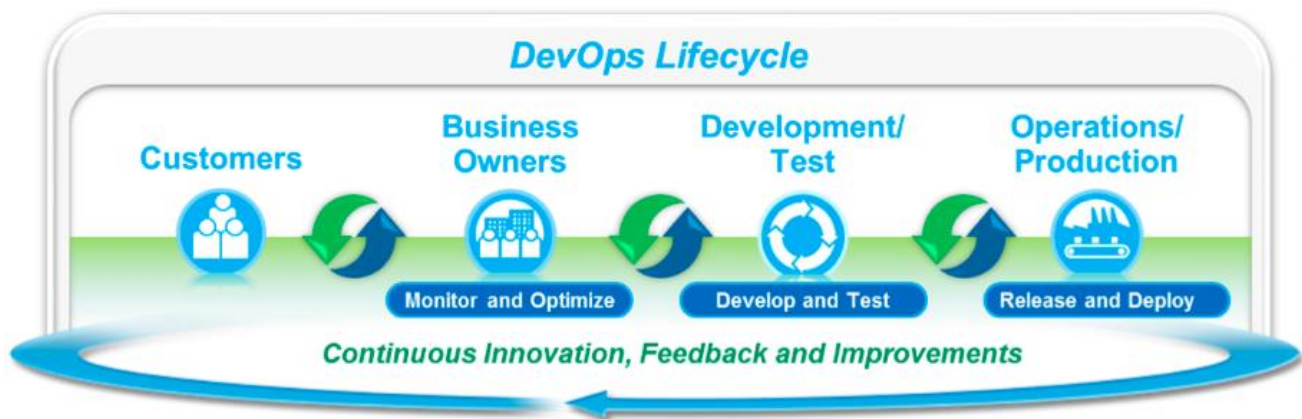
- Self Service Catalog
- Automation and DevOps
- Pay as you Go, Showback and Chargeback
- Automated Self-Service Catalog and AppStore
- Open API accessible infrastructure and adoption of OpenStack
- Defined SLAs with on-demand serviceability
- Tier based services (Platinum, Gold, Silver, Bronze)
- ITIL and IT Service Management

By looking at this list you may be saying to yourself, “We do all of that” or like I had listed above, “I don’t see the value in really doing any of that.” That’s why it’s important to understand what is important to your business, what is important to your IT organization, and especially – since this paper is discussing ways to reduce cost – where can it actually be used to reduce costs? If there are consistent, constant, repeatable tasks you are performing which are either time-consuming or error-prone, taking advantage of the capabilities of Automation immediately becomes a no-brainer which will help reduce errors and reduce cost in the bigger picture of time to deploy and time to market.

Unfortunately, other than a list of really pretty features and capabilities, most of these others are likely to incur additional cost before they actually result in cost reduction. For instance, being organizationally designed for things like Chargeback can provide a profit center to the business. However, situations where contracts and organization chargeback aren't already in place is wrought with burden and legal boundaries, many of which live outside of IT's grasp.

There has been an increased interest within the IT community and industry-wide in adopting OpenStack to help in deployment of new applications and especially so in the DevOps integrated space. If increasing the speed to market and adoption within the DevOps lifecycle is important to your business, this is an area – as seen in Figure 6 – where you may be able to realize cost reduction while providing a modernized capability to your development community.

Figure 6: DevOps Lifecycle



In this particular phase, there truly is no one right way to go about doing it. You could choose every attribute above and a hundred other and get nowhere if it is not the right fit for your organization. This is where sitting down and understanding what works best for your IT organization, business users, and ultimately, end-users becomes crucial. If there are ways to reduce time to delivery, time to availability, and minimize users from using capabilities in the form of "Shadow IT" and getting their services outside of the organization due to lack of agility, there may be absolute fits which will help your organization not only respond to the needs of today, but be able to react to the business requirements of tomorrow.

You may notice that higher up the stacks of virtualization, we begin to see fewer CapEx savings and even OpEx savings and started to see attributes such as reduced downtime, and higher resiliency, availability, and speed to delivery which are more difficult to place a dollar figure on. The further you dive into the complexities of technology the harder it is to squeeze blood out of that stone, but it doesn't make the benefits any less substantial. The more complete capabilities which are relatively easy and inexpensive to implement will likely provide the IT organization with attributes which were often extremely expensive in the past. In some cases, they were downright impossible to implement. The point is, do not be discouraged when hard cost reductions begin to wane. As we get more invested in the technology, the benefits and features significantly outweigh the associated costs.

Storage Consolidation


The discussion of Storage Consolidation is a far cry from virtualization consolidation yet all the more important and essential when it comes to reducing cost while bringing a sense of modernization to the data center. As our environments grow and shrink, they ebb and flow with the demands of the business. However, one thing tends to be constant; Users need to consume disk space and they'll never want to delete anything once it is created. In this era of Big Data and an infinitely growing user-set of information, consider regulatory compliance calls for saving 'x' data for 'y' number of years, with no deletion schedule in site. Clearly, storage is constantly growing with no end in sight.

Whether you have Direct Attached Storage (DAS) throughout your organization or SANs up to your ears, unless you made an investment within the past 18 to 24 months, it's likely there's a considerable value and benefit to be had by adopting some kind of storage consolidation strategy. Some of the key realized benefits you gain by consolidating are:

- Larger capacity drives at a much lower cost
- More reliable and resilient architectures with modern SAN hardware
- Better performance, Latency optimized, significantly higher IOPS
- Enhancements like Deduplication, Thin Provisioning, Compression
- Automatic Storage Tiering Policies moving data when hot or cold

This is just the tip of the iceberg when it comes to the capabilities which are available in modern data center SAN-type systems. We haven't even touched on all-Flash arrays, content addressable storage (CAS) or single file systems in the petabytes.

Figure 7: Fully Automated Storage Tiering (FAST) for automatic placement of data



Just from the few examples above and so many more, the abilities to reduce costs while also modernizing your data center can result in significant impact for your organization. However, if you still need some suggestions on how this can directly impact your bottom line, consider the following:

- SANs >5 years old cost significantly more in drive replacement and support than one purchased today
- SANs >5 years old typically have smaller capacity drives and take up significantly more rack space
- SANs >5 years old typically run slower with lower IOPS and higher latency times than a modern SAN
- SANs >5 years old both in hardware and software can't support the capacities or IOPS of today's SAN

We're just scratching the surface here but these are just a few examples of how you can upgrade your data center, advance your organization, and save a LOT in the process while addressing part of the bigger picture. What I've seen happen in the real world around this, is not merely replacing older SANs with newer ones, but instead replacing numerous older SANs with a few or only one newer one. Storage tends to be one of the more complicated areas to understand and also one of the easiest to work with if you have a background and understanding. I encourage you to engage an expert when selecting the right solution which fits your organization – especially when it comes to consolidation. This is one of the areas where leveraging the expertise of an objective Partner or VAR will come in handy in not only selecting the right solution for your business needs, but also to help ensure that you're addressing the full solution.


Backup

Backup is perhaps one of the least interesting areas when it comes to modernization. Yet this tends to be the most overlooked area in the data center, one which can determine whether you're paying more than you should or doing things efficiently and, let's not forget, legally. Does this story sound familiar to you? You buy and build an infrastructure and go to the well to ensure you have the appropriate funding needed to compute, analyze, and ultimately store the environment which runs your production, which runs your business. You find that your application size is 5TB so being conservative you support your backup solution to support backing up your application four full times, at 20TB in capacity. You start backing up your environment and find that 20TB isn't enough, and now that you have this infrastructure now everyone else wants to back up their data to it. You can support upwards of one to two days total of backup data.

Does that sound familiar? If it doesn't, I commend you but keep reading as there still may be reductions you can make while investing in your backup solution. On the other hand, if it does sound familiar, you are not alone. Unfortunately, this is a situation which plagues many organizations. When we discuss backup, many people will turn their thoughts to "Tape" or Grandfather-Father-Son backup rotation schemes. While backup itself hasn't changed much over the years, it's changed so dynamically in other respects that the types of savings you can realize can be a game changer.

Traditional backup solutions in the past would consist of a Weekly Full backup, taking a complete copy of all data and then an incremental or differential every day until the next Weekly Full is executed. While this would allow for a defined schedule of recovery points, in today's modern data center that solution is wasteful in both time and disk. Today there are solutions both hardware, i.e. EMC's Data Domain®, or software like EMC's NetWorker® which can provide features like deduplication, bandwidth saving through tape replication, Fast Backup and even Faster Recovery! Figure 8 shows an example of some of those savings realized without changing within the application itself and merely having software with some intelligence reading and writing that data out for backup.

Figure 8: EMC Networker Backup of Databases and Applications



500GB database, daily full backups, 10% change rate

operation	backup time	throughput	dedupe ratio
day#1 backup	30 mins	880 GB/hr	3x
day#2 backup	19 mins	1400 GB/hr	7x
day#3 backup	18 mins	1500 GB/hr	10x

Your mileage may vary

Speaking of Recovery, the whole reason backups are performed in the first place is the ability to successfully recover that data someday. Many Backup Administrators and IT departments may forget the main purpose of their backup is to be able to recover it, not to have successful backup jobs day in and day out. Where this tends to be most important is in litigious or regulatory-driven organizations whose Legal Dept. may not have communicated up the IT chain the importance of the recoverability, nor the cost associated with the inability to recover that data.

You may be wondering when I'm going to "show you the money" so to speak when it comes to backup, especially if you're performing a traditional type backup-to-tape or backup-to-disk model today. The need to retain information longer and longer – especially when it comes to storage data – tends to be a burden of the business. It then falls upon Backup to

ensure there is a copy of that information which is recoverable in the event of a catastrophic loss. Unfortunately, at the current scale of traditional backups, we'd run out of time in our backup windows as well as places to store the data sufficiently for any duration of time.

Figure 9: EMC Data Domain Backup over time

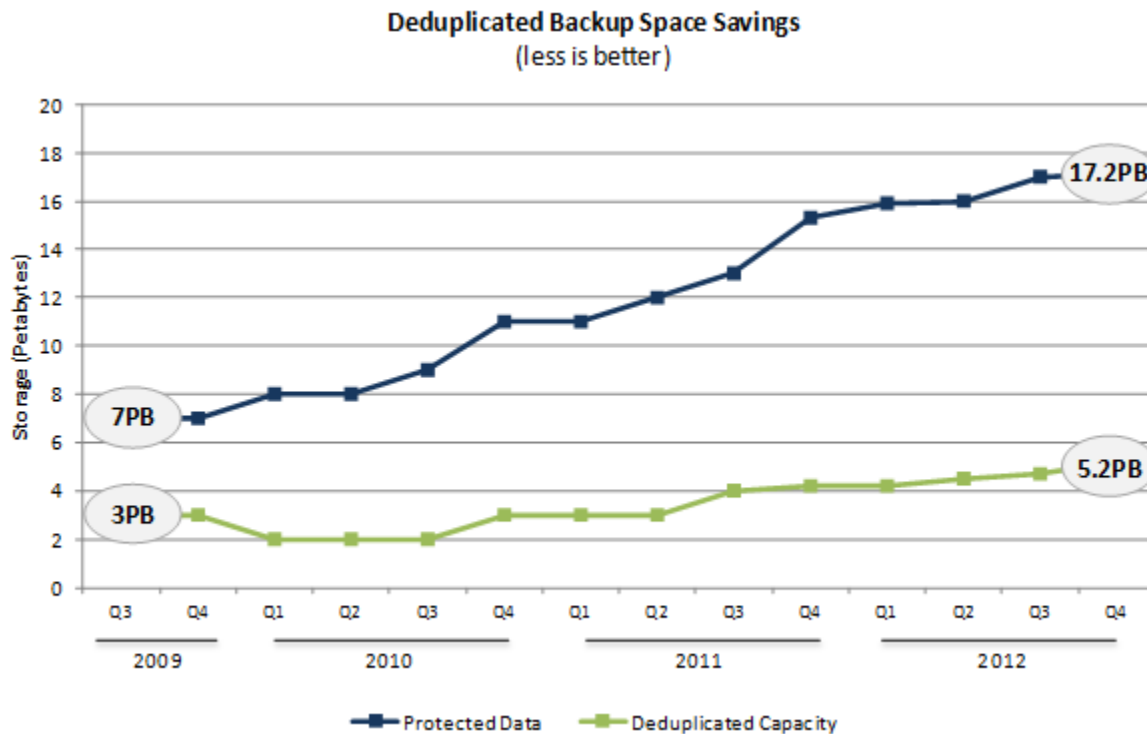


Figure 9 illustrates that, due to best-of-breed technologies like deduplication, while the amount of protected data has grown significantly year over year the amount of capacity used on disk seems to grow at a much more manageable rate. Suppose that every year you needed to increase your backup capacity to support your environment by 5 or more Petabytes just to keep the lights on and the backups flowing. Though results will vary depending on your backup environment, the type of data that you're protecting, and the schedule of backup, recovery, and retention, there is one thing on which we can all agree; the amount of data that we're being required to protect seems to grow every year, while our budgets to protect that data seem to shrink.

I acknowledge that backup may not be the most interesting solution to account for, to manage, and to back up. However, if you've not invested in a solution that saves the time it takes to back up and restore, and significantly decreases the amount of data you store vs protect, then this is a perfect opportunity to reduce cost for something you are doing today while producing value that is best realized when the environment needs it most.

Disaster Recovery and Business Continuance

I prefer to lump together Disaster Recovery (DR) and Business Continuance (BC), not because they're similar in the technologies or investments which are made to make them possible, but because a plan for one will often accommodate a plan for the other. As we discussed earlier around some technologies which can be adopted mutually exclusive from one another, this is one of those areas where we have some dependencies to move toward the best solution for our organization.

Disaster Recovery

DR is the fundamental approach used to ensure that services at another location are restored within a specified time (Recovery Time Objective [RTO]) and with an acceptable duration of time of data loss (Recovery Point Objective [RPO]), if your existing data center became a smoking hole in the ground, underwater by a tsunami, ripped apart by an earthquake, or somehow inaccessible to where your business is down for an undetermined period of time. It sounds so simple, so straight-forward, so awkwardly complicated!

Figure 10: Recovery Point Objective and Recovery Time Objective Chart

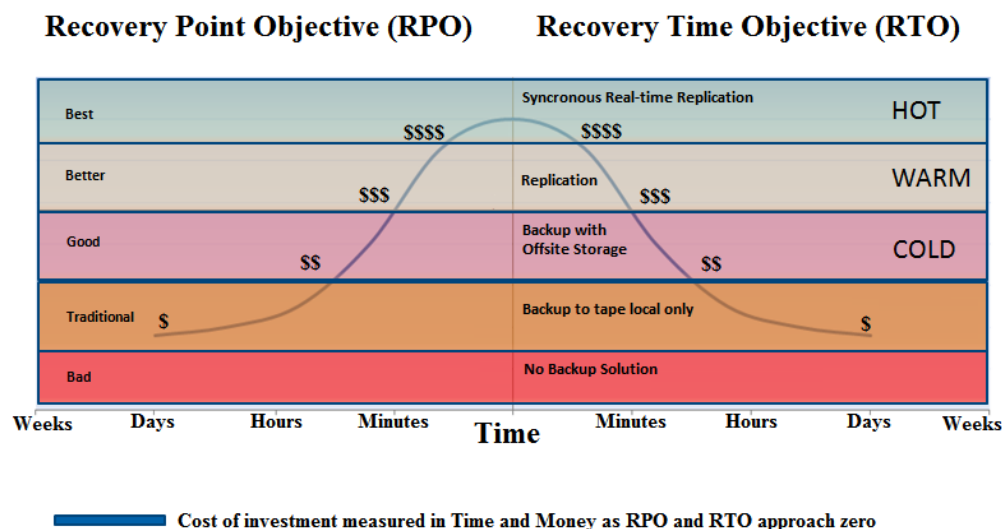


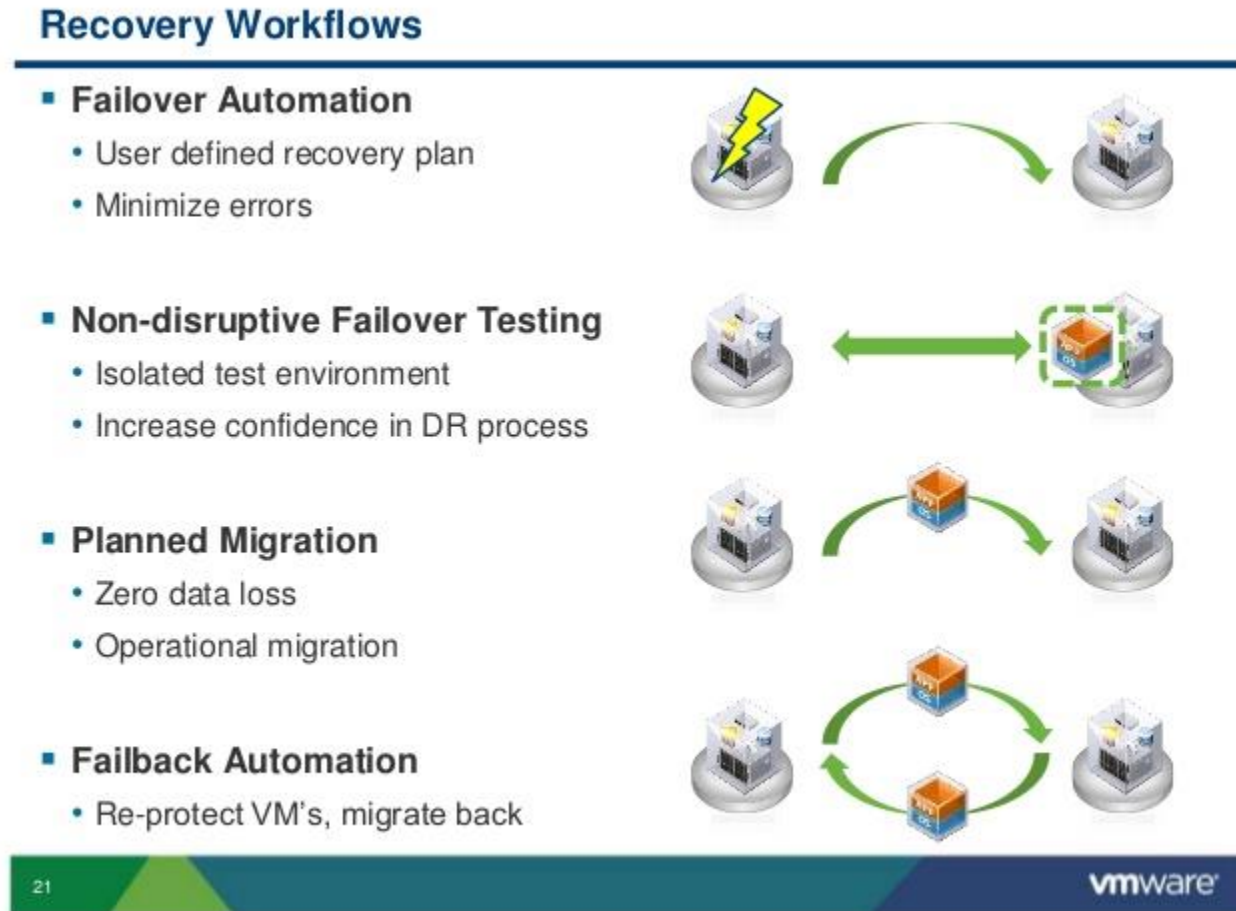
Figure 10 depicts how solutions become more expensive the closer businesses drive to meet an RPO and RTO of Zero time to recovery and Zero data loss. This is where the challenges of Disaster Recovery need to be established and accounted for. Common challenges to address when designing for DR are:

- Ensuring that application consistent data is replicated to a remote site
- Ensuring that networks are accessible at remote site
- Ensuring that the applications are started in the appropriate order when failing over for DR
- Ensuring that the DR plan is tested, testable and able to be successful

Those last few points are especially important because people have written DR plans for decades which when tested would not be successful, often through no fault of their own. Usually, failures could be attributed to limitations in technology or investment. Fortunately, solutions such as Site Recovery Manager (SRM) help mitigate most of the issues above while also automating most if not all of the actions performed. This will enable non-disruptive failover tests to ensure you'll be able to get services up and operational with minimal disruption in the event of an actual disaster.

Figure 11 gives an example of some of the workflow types that you can execute such as Failover testing, Automation of your failover, and especially the ability to automate the failback after the primary site is brought back online.

Figure 11: Recovery Workflows in SRM



Often not realized but important to keep in mind is that Disaster Recovery and even the Planned Failover of one site to another is a disruptive event, even when using an automated tool such as SRM. All services will cease to operate and function and will not become accessible until they've been restored at their destination. If performing this as a planned failover, all services at the source will be shut down. Only after they've all been terminated will they be restored to the destination.

Reviewing Figure 10, you'll see this is where we have to make the decision of what is important to the business to transcend beyond simply recoverable in a disaster event, to determining what must remain up and online without interruption despite a devastating failure or outage.


Business Continuance

Challenges when working with a BC solution tend to vary greatly from your standard DR solution. Instead of merely running automation or shutting things down, Business Continuance keeps running no matter what. To ensure services keep running without interruption as previously depicted in Figure 10, the cost goes up significantly. Not always in the cost to replicate or the cost of the technologies under the covers, but often the cost of the required low latency links between the sites to ensure you can have synchronous replication of data. What are some of the typical requirements for a Business Continuance solution that drive up the costs so high?

- Capacity required at both source and destination
- Maximum latency and Round Trip Time (RTT) needs to be below “x”
- Source and Destination networks must be spanned and routable
- Bandwidth required between source and destination can be significant

These requirements have changed over time and can vary based upon the solution. Hence, I’ve not specified the exact RTT or the amount of bandwidth required. Nonetheless, these figures are typically very specific and very controlled by the organizations with the Business Continuance solution to support. Figure 12 shows an example of the EMC VPLEX® solution providing business continuity between separate physical datacenters whereby you can move virtual machines logically between the locations without any impact to operations or function. While there are numerous solutions on the market which can provide these kinds of capabilities, adopting a Virtualization First model will enable you to do much more at the data center level rather than relying on numerous point solutions operating at the application level.

Figure 12: Example operations of VPLEX used for Data Center Business Continuance



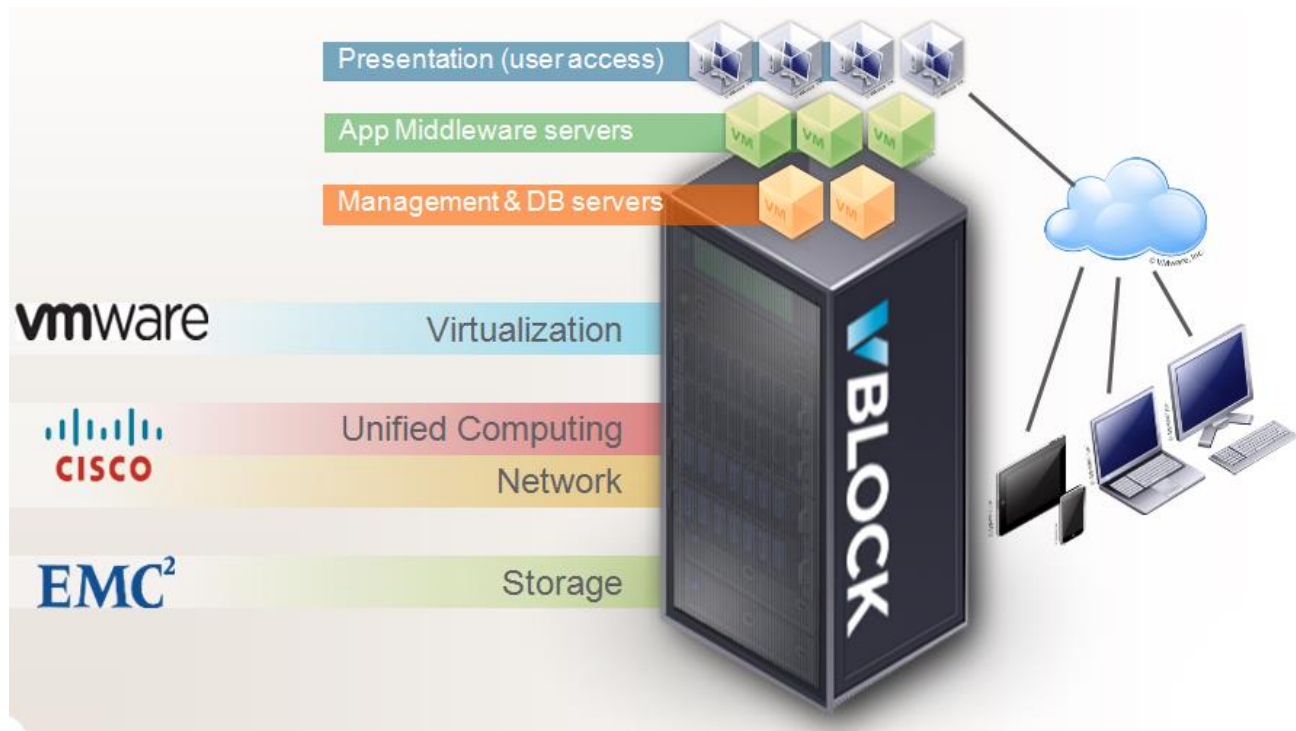
Here at the end of this very interesting conversation around Disaster Recover and Business Continuance you’re probably asking, “That was awesome, but I didn’t exactly see where I reduce cost, per se...” You’re absolutely right. The higher up the stack we go, the higher up the fruit tends to be. Consider the following, most organizations have a non-rationalized and non-functional DR plan, and no BC plan at all. They likely paid some company a lot of money to produce their run book for DR and never successfully were able to test it during a test event. Imagine if you can check that box, pass that test all the while actually having a successful run book which works and have it cost less. This is where we end up reducing cost while providing value which was otherwise non-existent. This like other solutions we discussed is not for everyone; though if you adopt a lot of the solutions discussed prior it’s possible that anyone and any organization can adopt this.

Convergence

The discussion of Convergence is not a new one; in fact it has been around for quite some time. But for those who are new to this idea, Convergence is the culmination of the ‘one throat to choke’ model whereby you source Storage, Network and Compute from a single ‘vendor’ which has undergone rigorous testing schedules to confirm and validate function and capability. I want to be clear that adopting Convergence in and of itself does not inherently reduce your cost if you were to acquire each of the parts piecemeal. This is not a CapEx reduction play; rather, it is treated as an OpEx reduction.

The rationale, should it work for your organization, is that you have a single supported solution from your vendor of choice. Thus, when it comes to set up, problems or support you call a single number instead of engaging your storage, network, and virtualization vendors separately. Figure 13 shows an example of how the architecture is laid out from the various layers.

Figure 13: VCE Vblock Architecture



For some organizations the OpEx savings can be significant as procurement and deployment time from order to production can be fairly significant for organizations that lack the agility or expertise to become operational in a timely fashion. These few examples have resulted in significant impacts for organizations whereby the Vblock turned into a saving grace and a win for the business. The converged solution is definitely a good fit for many organizations who want and need that single support number, that placed order to delivery, and that tested and proven software and hardware components in the box. While it may not be a fit for every organization, you should know it exists to determine whether it is indeed a fit for your environment.

Hyper Convergence


Hyper-Convergence, while sounding similar to what we've just discussed, is actually drastically different in numerous ways. When it comes to the points we'll be discussing, it's drastically different in many significant ways. There are numerous Hyper-Converged infrastructures available, some which are actually Hyper-Converged, and some which claim to be. For the purpose of this discussion the actual definition of a Hyper-Converged Infrastructure is compute, network, and storage resources combined and closely integrated with a hypervisor solution. The differentiator being that while some converged solutions will use a hypervisor, it comes down to the layers and levels of integration.

Figure 14 depicts VMware's EVO: RAIL solution which is constantly changing and evolving. A relatively new player to the market, it maximizes deep hypervisor integration by also being one of the first to truly adopt Software-Defined Storage (SDS) as part of the solution by leveraging vSAN in the infrastructure itself.

Figure 14: VMware EVO: RAIL at a glance

EVO:RAIL is an Ideal Building Block

- Simply add appliances for linear scaling and performance
 - Each HCIA adds compute, storage, networking, and management in minutes
 - New appliances are automatically added to the cluster
 - Version 1.0 supports 4 EVO:RAIL Appliances (16 ESXi hosts) in a cluster



The EVO: RAIL solution lets you pick your hardware vendor of choice and get an easy to configure, easy to deploy Hyper-Converged infrastructure that can scale and grow your virtualization environment today. How does this help reduce cost? Given current licensing costs for features like vSAN and others, I'm not suggesting it reduces costs through CapEx or licensing and support. But, it can possibly help reduce OpEx through the linear scalability of the system, not to mention other advancements being pursued in this direction.

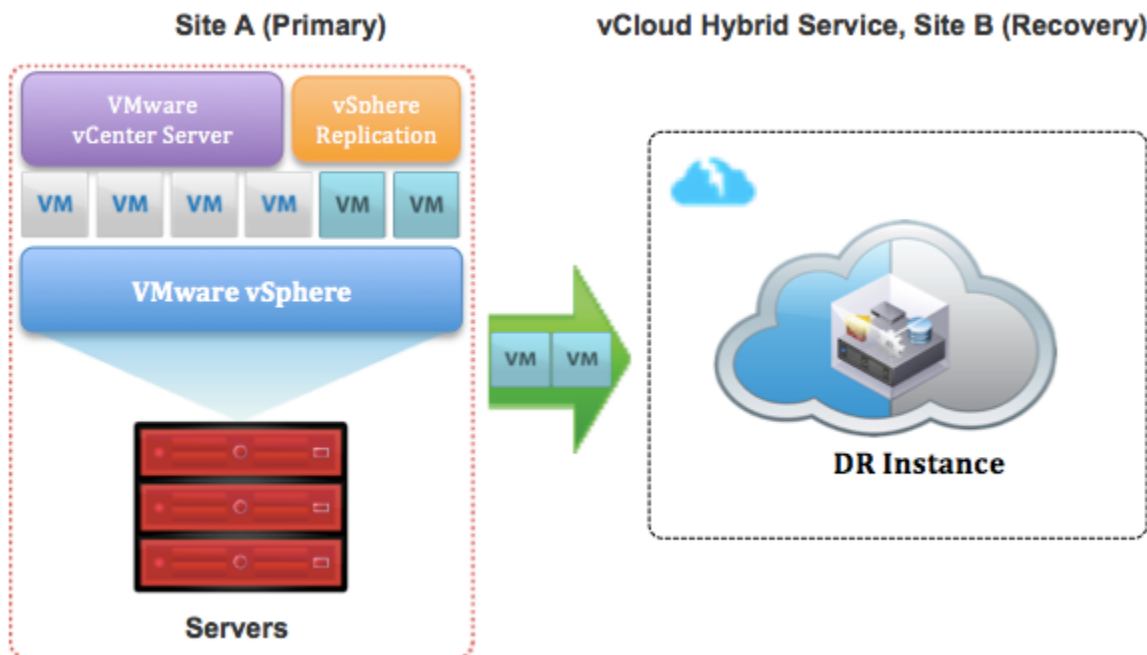
Cloud and Hybrid Cloud

Last, I want to address the 800lb gorilla in everyone’s organization. Every organization has one cloud initiative or another, whether it is “Cloud First”, or “Journey to the Cloud”, or “Cloudy with a chance of meatballs”. In any case, everyone has cloud on the mind. What does this mean for modernization of your data center, reducing costs, next generation, etc.? Depending on which investments in Cloud and Hybrid Cloud you’re looking to make, it can be rather significant. I’ll touch on a few of those areas below where a lot of people find their biggest bang for their buck.

- Disaster Recovery as a Service (DRaaS)
- Backup to the Cloud
- Platform as a Service (PaaS)
- Application / Development
- Cloud-Burst and Expanding On-Premises applications to Cloud spaces
- Remote Desktop in the Cloud
- Eliminate CapEx through Infrastructure as a Service (IaaS)

This merely touches the surface of how some people use cloud today. Figure 15 briefly touches on the vCloud Air solution supporting Disaster Recovery as a Service. In a solution like this, your virtual machines continue to run on-premises in your organization with virtual machines and data replicating off-site to vCloud Air. In the event of a catastrophic failure or a staged outage you failover services, networking, and connectivity to your DRaaS site and continue to run services there. Perhaps you’re now realizing that you don’t need to maintain a separate DR facility to support this. You would be absolutely correct. Your reduced cost is achieved through not having the CapEx of equipment sitting idle, facility, etc., and the adjusted OpEx to maintain the DRaaS site, cost, and so forth.

Figure 15: VMware vCloud Air (vCHS) leveraging Disaster Recovery as a Service (DRaaS)



This same approach of failing your applications or servers over to the cloud applies just as well if you're looking to Cloud Burst an application. This now crosses from simply a cloud service into the Hybrid-Cloud type solution where you have your data and network presented both in the Hybrid Cloud as well as part of your on-premises infrastructure. Whether this is a good fit for your organization comes down to contractual understanding, comfort and ease of use to accomplish things, and whether over the course of the operation of these services, you'll realize enough CapEx and potential OpEx savings to offset the cost of having this exist outside the walls of your data center.

Be aware that most of these Cloud and Hybrid-Cloud solutions do not have a good and viable backup solution today. They have 'solutions', but they're not something you'd want to rest your career on, so consider that also as part of the bigger picture cloud strategy you have to deal with.

The last of these items to discuss is the elimination of CapEx for an all-in IaaS solution in the Cloud. It's a strong possibility you can adopt this through Azure, AWS, Rackspace or vCloud Air. You can run all of your servers, services, applications and so forth entirely off-premises in the Cloud. Things to consider when it comes to making that decision include:

- Do we have any regulatory requirements?
 - Do we have security requirements?
 - Do we have data retention requirements?
 - Do we have to deal with Right to Audit requirements?
 - Do we have any contractual requirements written with customers?
- How much does it cost to run our applications in the Cloud?
- How much does it cost to backup and restore our applications in the Cloud?
- How long will our application live out here?
- What are our data storage and bandwidth requirements?

It is good to understand these questions at a high level as that can help make the determination if a particular application is 'Cloud ready' or even cloud possible. Using 'the cloud' doesn't mean it will be cheaper, sometimes not in the short-run, often not even in the long-run. You do save in up-front CapEx, but the longer an application runs and consumes bandwidth and capacity, the cheaper the physical alternative looks. There is no general purpose 'right' answer when it comes to determining what should or shouldn't live in the cloud. The biggest mistakes I see happen are moving applications off-site to the cloud and not backing them up, or expecting them to 'be backed up, because they're in the Cloud'. Not realizing you need to do something doesn't mean it reduces your cost, it just makes your cost burden that much more significant.

Certain applications are prime candidates for being moved to a cloud service provider and others are just a burden and at this time are still far too expensive, even with dropping rates. Understanding which is which, and what fits is why the Hybrid-Cloud strategy is such a powerful piece of your portfolio and infrastructure. Being able to run applications 'adjacent' where one is on-premises and the other is off-premises gives you a competitive differentiator both for your business and for your pocketbook.

Be aware of what your costs are because if you're a CapEx-heavy organization, adopting a cloud strategy may make you appear that you're wisely and aggressively saving, but in the end you may merely just be moving the buckets. And sometimes those buckets may be more expensive if not planned appropriately.

Summary

Every organization has a unique set of challenges. However, with minor exceptions we're mostly all using the same infrastructure, building the same type of data centers whether legacy, modern or next generation. A lot of challenges have been faced by people long before you and will continue to be faced long after you. Hopefully you have found useful some of these high level discussion points with practical examples of ways to reduce the cost within your IT organization while actively trying to modernize your data center. If you are able to find value in just one area to improve and benefit your organization, I consider this paper is an overwhelming success.

About the author

Christopher Kusek is Chief Technology Officer of Portland, OR based system integrator and VAR, Xiologix, LLC. With over 20 years of experience in IT as a Technology Evangelist and Industry leader, Christopher plays a key role in Xiologix's growth leading its storage and engineering practice to evaluate and architect solutions that meet client's tactical and strategic goals.

Prior to Xiologix, Christopher spent three years at EMC where he was the Global Lead for Cloud and Virtualization. In this position he was focused on the company's relationship and integration with VMware, VCE, Storage and Cloud Services. Christopher was also responsible for Product Management, Marketing and Sales. Before that, he worked at NetApp as a Technology Evangelist and Principal Consultant responsible for pre-sales and post-sales hybrid consulting engagements surrounding Virtualization, Storage, Microsoft and security solutions. Most recently Christopher has returned from spending the past two years responsible for theater wide infrastructure operations for the war effort in Afghanistan.

Christopher is connected to current and future trends; a strategist who comprehends complex business and technology problems uses his organization and leadership skills to solve them. An industry recognized expert Christopher is an EMC Elect and VMware vExpert, while also an accomplished speaker and author with five books published. With over 20 years of industry experience spend focused on creating innovative business solutions, Christopher's expertise is integral to Xiologix's ability to lead its clients to the right solution for their business.

About Xiologix

Xiologix is a leading provider of innovative information technology (IT) solutions serving corporate and public-sector customers. We provide technology solutions - including hardware, software, and services - to help customers resolve their most complicated IT needs. Our expertise includes enterprise architecture and high availability, infrastructure optimization, storage and resource management, identity management and business continuity. Xiologix is strategically partnered with Anitian, a leading information security intelligence firm, who delivers security assessment and compliance services for Xiologix's clients. Xiologix is headquartered in Portland, OR area with sales and engineering offices throughout the western region. For additional information call (888) 492-6843 or visit <http://www.xiologix.com> or <http://www.anitian.com>

Contacts

Xiologix

Sheryl Still

sstill@xiologix.com

888-492-6843 x254

503-691-4364 x254

Twitter: @Xiologix

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.