Data Center Consolidation

White Paper

September, 2016

World Wide Technology, Inc.

Preface

This white paper on data center consolidation was originally prepared as part of a proposal provided to a federal agency. The document content has been altered to remove the customer-specific information and to provide a more general document to address considerations for any organization contemplating a data center consolidation.

WWT brings an innovative and proven approach to how organizations evaluate, architect and implement technology. Taking an architectural, independent multi-vendor approach, WWT provides high-performing technology products, services and supply chain solutions to our customers around the globe.

With more than $7 billion in annual revenue, WWT is a financially strong, privately held systems integrator that ranks among the top tier of partners with Cisco, HPE, Dell Technologies, NetApp, VMware and Citrix. WWT offers uncommon financial strength and stability, allowing our customers to feel secure in long-term business partnerships.

WWT's Advanced Technology Center has the ability to provide proofs of concepts prior to implementation of a data center consolidation.

Data Center Consolidation: Recommendations and Considerations

WWT Recommendations

Efficiency

1. Assess power requirements to determine actual power required per rack rather than an arbitrary figure.
2. Include precision cooling and quality power techniques based on actual IT equipment power requirements.
3. For mobile environments, protect racks with field replaceable vibration isolators to protect the equipment based on their weight.
4. Provide a true multi-vendor approach. This requires support for equipment that is wider than the standard 19" rack width and airflow that is other than front-to-back (for example, some network equipment).
5. Maintenance should be performed inside of structure with access to both front and rear of racks with all exterior openings closed.
6. Provide a computational fluid dynamic (CFD) analysis of the proposed design with the proposal.

Things to Consider

Moving an application to a cloud model without a full diagnostic of any performance, extensibility, security, and governance aspects can result in outages, performance degradation, and data exposure, among other issues. Within the private cloud framework, there is more flexibility to act quickly; however, long-term planning will still need to realize the same information.

Figure 1: Applications: Things to Consider

World Wide Technology, Inc.

Table of Contents

World Wide Technology, Inc.

Data Center Consolidation Success Stories

.

Customer 1: US Federal Agency	
Scope	When a U.S. Federal Agency needed to consolidate multiple data centers into a cloud infrastructure in just 45 days, World Wide Technology Advanced Solutions (WWTAS) got engaged. In this unique program, in which the federal government chose to utilize a commercial enterprise for both integration services and delivery, WWTAS instituted the Factory Express Initiative to deliver 78 custom racks and six (6) complete cloud systems with FlexPod® Data Center Architecture in time to meet the client's aggressive deadline.
Solution	The WWTAS team is comprised of specialists in systems, virtualization, storage, testing, network engineering and database development with vast experience in high-availability clustering and enterprise systems management. Over 30 resources spent two months in the warehouse working 12-hour days, seven days a week to meet the client's aggressive timeline. Over the course of the project, WWT's OEM-certified professional services team used video conferencing to work directly with the remote customer to design custom assemblies, configurations, integrations, cabling and testing schematics and provide day-to-day project updates. As all of the FlexPods required configuration, custom cabling, testing, staging and shipping, WWTAS utilized WWT's Integration Technology Center (ITC) in St. Louis, Missouri. WWTAS immediately established 30,000 sq. ft. of dedicated space for the program comprised of 17,500 sq. ft. for inbound material staging and cabinet builds, 12,500 sq. ft. for staging, storage and crating, and 2,500 sq. ft. for testing. A transformer with 380 amps and 208 voltage was added to support increased power demands. Within this optimized setting, WWTAS simultaneously assembled, configured, tested and asset tagged all 78 racks so that they could be completed and delivered to the client at once.
Results	WWTAS packed, shipped and delivered optimized systems to the customer in time to meet their aggressive 45-calendar-day deadline. The project is ongoing and, to date, WWTAS has been awarded three increments totaling 12 systems, over 100 racks and over $40M in revenue.

Customer 2: Healthcare Organization	
Scope	This major healthcare organization's data center was out of power and recently had a Server Consolidation Assessment provided by VMware. Assessment results showed a total of 1200+ physical workloads that could be virtualized and consolidated to free both power and space.
Solution	WWT Virtualization Architects analyzed assessed data to assist in designing a virtual environment encompassing server compute, storage, and networking hardware. In addition, a detailed P2V plan and design was developed specific to the customer's environment which allowed an accelerated migration path of all 1200+ physical work loads into the newly deployed virtual environment. This accelerated migration path allowed the customer to consolidate all workloads in less than four months running at an average of 75 P2V's per week.
Results	Over 900 P2V migrations were completed in less than four months. The P2V run rate averaged over 75 migrations a week allowing the customer to quickly adopt almost complete virtual environment. In addition, over 200 servers were discovered to be eligible for decommission, and removed from the network. Power utilization within the data center has dropped significantly. Centralized compute infrastructure along with 10Gb network has increased utilization, performance, and allowed for future growth as demand needs excel.

Customer 3: Public Utility	
Scope	The customer launched a Data Center Modernization (DCM) project to upgrade its two existing Data Centers and build out a new Data Center. The two existing Data Centers are located in San Francisco (SF) and Fairfield (FF), California. The customer issued an RFQ for the network build-out at the three (3) Data Centers including the new Rancho Cordova location. In addition to procurement of network equipment, the RFQ requests for network design, configuration, implementation, and testing services at these Data Centers. The services cover Data Center LAN, MPLS core, Internet connections, firewalls, IDS/IPS, and global and local load balancers.
Solution	WWT worked with the customer resources and Cisco Advanced Service (AS) team to carry out the procurement, design, configuration, implementation, and testing of the network infrastructure portion of the DCM project. Local Area Network (LAN) Two physical LANs were installed each at the new Data Center and FF Data Center: production LAN and IBN. The production LAN has a pair of Nexus 7010s as core, with Nexus 5548, Nexus 2232, Nexus 2248, and Catalyst 6513 at distribution and access layers. The IBN has a pair of Nexus 5548 as cores; Nexus 5548 and Nexus 2248 at access layer. The San Francisco Data Center LAN will be upgraded with redundancy and additional capacity. Multiprotocol Label Switching (MPLS) Core The customer owns its ONS network. The customer is building its own MPLS network. A pair of ASR9006s at each of the three Data Centers will act as P Routers. The new Data Center and FF Data Center will each have an ASR 1002 as Router Reflector (RR) to managing the BGP Routing and a pair ASR1004s as PE Routers for WAN aggregation. Internet For the new Data Center and FF Data Center, each will have a pair of Catalyst 6513s as the Internet Gateways. The Catalyst 6513s will have ASA and ACE modules as Internet firewalls and load balancers. A pair of IPS4270s will connect to the Catalyst 6513s for Internet IDS/IPS functionalities. Firewalls and Internal IDS Apart from the ASA modules in Catalyst 6513 Internet Gateways, there will be two pairs of HA ASA5585s at the new Data Center and FF Data Center. The two pairs of ASA5585s will sandwich the internal load balancers on a pair of Catalyst 6506 with ACE modules. There will be another pair of ASA5585 with IDS modules for internal IDS/IPS needs. Load Balancers In addition to the ACE modules in Internet Gateway Catalyst 6513s and in Catalyst 6506s as the internal load balancers, each Data Center will have one GSS4492 for global load balancing functionalities. Security Area Network (SAN) Although Nexus platforms have built-in FCoE and SAN switching capabilities, efforts related to SAN and storage design, configuration, implementation, and testing are outside of the scope of this project.
Results	WWT performed the following services during this project Cisco Technologies, Security Technologies (Physical, Network, Endpoint, Application, Management, Etc.) and Router/Switch Configuration/BGP

Technical Considerations

1. Tier 1: System/Application Evaluation and Business Process Re-engineering:

- The key organizational construct to properly affect application evaluation and accelerate savings lies in the customer's ability to do proper testing evaluation and certification. To provide responsive, agile, and cost-effective interoperability, the tests, evaluations, and certifications construct is an essential element of the acquisition system. As you move forward with agile IT acquisition concepts, it must progress toward the use of distributed test methodologies that provide realism and promote effectiveness. The customer's early and continuous involvement throughout the IT acquisition life cycle will help ensure that both agility and cost saving are achieved.

  It is suggested that the customer include in the evaluation the ability to prototype, test and demonstrate applications and hardware solutions. In this manner, the customer will be able to more effectively evaluate applications and more accurately perform cost-benefit analysis, resulting in accelerated savings.

- Tools for Portfolio Management are evolving rapidly. There are over 100 tools for PPM and no one tool that is best for every organization and system implementation. Available tools differ greatly in what they do and how well they perform. No single tool does everything. Also, and most importantly, the right tool depends on the nature of the business, the kinds of projects to be conducted, the maturity of existing processes and associated tools, organizational culture, the degree of rigor desired, and what it can realistically bring to your decision-making processes.

- We have found that the dominant factors that must be considered when evaluating a system or application for either elimination or inclusion in an overall system strategy include:
  - Forward projection of technical requirements and technology trends that allow for a measurable assessment of how well systems and applications under consideration meet current and future needs.
  - Maintenance difficulty (MTBF, MTTR, downtime, workforce productivity impact, etc.) Is this system/application becoming a liability?
  - End of life/support considerations
  - Fit within architectural strategy
  - Evaluating all dependencies – systems, applications and organizational process
  - Financial considerations – Financial life of the assets, pending tech refresh and renewal schedule, cost benefits analysis (CBA), a new total cost of operations analysis and competitive impact
  - Suitability to task – current and future, operations design and delivery
  - Risk and potential impact to mission operations
  - Alternative solutions (as evaluated by CBA, requirement trends and risk analysis)

- Implementing technical and security standards to comply with specific application requirements is a continuous challenge government and commercial businesses face in today's environment. Industry offers many options for implementing and monitoring

security standards, however, maintaining processes and practices to remain compliant relies heavily on executing proper enforcement measures, and maintaining visibility to changes in the environment and threat/vulnerability landscape.

Common standards designed to protect privacy, provide security, promote technical compliance, and sustain environmental controls (FedRAMP, HIPPA, FISMA, SCAP, DIACAP, FIPS, 508, Energy Star, etc.), when combined, can be extremely cumbersome to maintain. Industry has answered the call by offering niche software tools that can provide positive, but inconsistent results. Unfortunately, there is no one product that spans the entire enterprise to provide a fail proof solution. However, by identifying critical processes and risk thresholds, combined with a trained team of experts to implement and execute a fully defined enterprise-wide continuous monitoring solution, the customer can expect to meet, as well as, maintain technical and security requirements.

2. Tier 2: System/Application Modernization:

- WWT uses a multi-phase approach to assessing and modernizing applications. Our methodology follows five phases, as shown below in Figure 2:

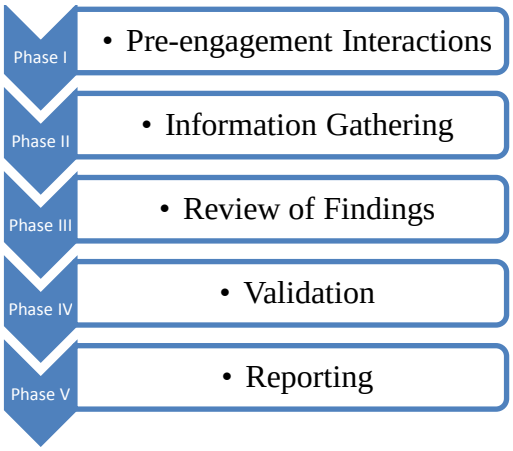

Figure 2: Five Phase Approach.

- **Phase I** – Pre-engagement Interactions: During the first phase, meetings are held with the government and other key stakeholders to review requirements, statement of work, methodologies (e.g., assessment, verification, testing, integration, validation, operations and maintenance), and appropriate protocols to be followed during the system modernization process.
- **Phase II** – Information Gathering: Gathering of information should consist of an automated high-level assessment. It should be all-inclusive and focused on the macro elements over a specific period of time in conjunction with face-to-face interviews and anonymous surveys to determine Technical Factors, Business Value, Cost, Staff / Skills Issues, and potential Vendor issues.
- **Phase III** – Review of Findings: These findings are then reviewed for false positives and evaluated to ensure each piece of the application has been through the automated analysis process. This also includes all third party libraries and components, as applicable.

- **Phase IV** – Validation: Results are then validated and described in a comprehensive report. This report details where in the application possible flaws and vulnerabilities exist, as well details on how they were verified.
- **Phase V** – Reporting: A comprehensive report detailing where in the application vulnerabilities exist, and providing detailed recommendations for modernizing applications

- The crucial and most dominate factors involved in any project of this scale to modernize and provide a pathway to application modernization and desktop optimization must include consideration of the following elements and are part of our assessment methodology discussed above:
    - **Cost savings**
        - FTEs
        - Licensing
        - Maintenance and Support
        - Less revenue or budget constraints
        - Moving CAPEX to OPEX
    - **Agility and elasticity**
    - **Quicker Time to Serve**
        - Automation
        - Service portfolio
        - ITIL, Governance
    - **Consolidation and Virtualization**
    - **Value creation**
        - Time to Value
        - Cost to Serve
        - Value of information
    - **End of Life**
    - **Rapid expansion or contraction**
        - Acquisitions
        - Down-sizing
        - Organic growth
        - Temporary need
    - **Easier innovation**
    - **Competitive advantage or necessity**
    - **Lower risk**

- Successful virtualization and consolidation of data centers depends greatly on following a disciplined system engineering process. Goals, objectives, and well thought out requirements are key components to success. It is absolutely critical to assess the environment in order evaluate how VDI fits into the overall enterprise desktop strategy. VDI can be effective for every enterprise, but it is not necessary suitable for every user. Some users will require traditional desktops due to their computing needs, while others would be the perfect candidates for virtualization. The goal of an assessment is to determine what are the user characteristics and the associated application workloads. The results of the assessment will guide the decision-making process throughout the

process. A typical assessment takes six weeks and includes a thorough data collection component, including getting information from all key stakeholders, and a detailed analysis.

Lakeside SysTrack is the leading and chosen solution for VDI assessment and monitoring. Other tools include Liquidware Labs Stratosphere, and Aternity Frontline Performance Intelligence Platform

Once the VDI assessment is performed, the data can be fed into the sizing model so that the solution can be sized accordingly based on the specific requirements. The critical performance parameters from the VDI assessment are average, max, and standard deviation. The five key metrics from a VDI assessment that will be used for hardware sizing are:

- CPU utilization
- Memory utilization
- Network bandwidth utilization
- Storage utilization
- Disk IOPS

It is important to note that sizing should not be done based on average, but rather, *peak average (average + standard deviation)*. Solutions that are sized based solely on average or max will typically be undersized or oversized.

Application virtualization lets us optimize the IT resources in the datacenter, minimizing wasted investment in servers, memory, storage or other components. By proactively allocating and adjusting capacity based on preset rules, capacity can be directed to where it's needed.

Further, intelligent capacity management allows for right-sizing the environment by both planning ahead for when additional capacity will be needed and by identifying and recovering unused resources.

Effective reporting, on both capacity and costs, helps us communicate IT's value and operating excellence to the business. However, reporting needs to be in the context of a dynamic infrastructure to be useful. It should:

- Leverage existing information so that applications can be quickly assessed and converted. Having the correct information available significantly speeds up the process of converting applications.
- Document complex tasks and make them repeatable. Tasks like provisioning or patching are often complex, multi-step processes. Because virtualization is a standard platform, tasks can more easily be made repeatable and successful.
- Use standardization to prepare, validate and deploy applications rapidly. Taking advantage of standardized platforms and templates enables efficient deployment.

- The major problem areas for industry in a large transformation project include but are not limited to:
    - Governance
    - Not understanding how to make architectures scale.
    - Not dealing with tenant management issues.
    - Security - not understanding that security is systemic.
    - Not understanding the proper use of services
    - Tossing technology at the problem.
    - Listening to the wrong people.
    - Managing by good guess.

  Many transformation efforts don't meet expectations because attention hasn't been paid to:
    - Efficient utilization of resources
    - Resource saturation
    - Lack of elasticity and scalability
    - Lack of security and governance
    - Long term performance and scalability
    - Bad or no tenant management
    - Lack of planning
    - API Integration
    - Workloads inappropriately split between locations and systems
    - Degree of automation and self service
    - Linking prior knowledge and data to the new environment
    - Cross border culture, compliance and regulatory issues

3. Tier 3: System/Application Virtualization:

- One of the main benefits of application virtualization is the ability to treat heterogeneous environments as if they were homogeneous; WWT believes the first step in any meaningful data gathering exercise is defining metrics of interest. For a true user-centric picture of application activity there's a core set of data that describe usage:
    - Platform dependency
    - User interaction
    - Resource consumption
    - Application complexity

Each one of these points is in turn composed of a wealth of important, illustrative information that provides direction for each of the distinct segments of users in every environment. A certainty of the current IT landscape is the demand of users to have access to applications on their terms. Agility is a paramount requirement for IT in all industries, but balancing security and stability with the right methods of delivery for all users can be extremely difficult. Knowledge of where users access applications and data drives later decisions for when applications can and should be delivered. Mobility and enablement of mobile devices is, rightfully so, a focus for the new "average" class of users that needs data and applications on demand, and starting the discussion with an analysis of current needs is a necessity. *Will highly mobile users require offline access? What backend data stores are in use? Are there concerns with methods of interaction available on the future platforms chosen?* Real user workflow information is pivotal in determining what options will meet the needs of an environment.

Insight into distinct work styles can also serve a broader purpose than enablement. By knowing where packages are installed, how many users really use them, and the type of usage for each application, significant cost savings through license reductions and productivity increases can be achieved. Key questions must center on differing modes of interaction: *How many users use this application simultaneously? Are users actively making content or viewing it? What user bases share software packages in common?* Use case planning allows for technical planning for infrastructure required to support mechanisms like XenApp/ThinApp/AppV application publishing that enable phenomenal flexibility for users.

Capacity planning for application virtualization is of paramount importance in exploring the feasibility of moving important applications to a new delivery mechanism. By first building the case for either developing net new delivery or expanding on current capacity planning becomes a straightforward architecture discussion. By understanding the resource demands of applications, their requirements for accelerated graphics, and the concurrency of usage a complete picture of provisioning needs can be developed to plan backend requirements to meet the needs of the user population. The concept of application complexity is directly related to this, and the amount of difficulty in delivery may play a large role in determining whether a package is installed directly into an image or delivered via virtualization.

Thorough environmental understanding is driven by familiarity with users, and key points that must be understood include:

- Platforms currently in use
- Installation record for all software
- Version information
- Where packages are used and who uses them
- Application resource consumption
- Complexity of packages
- Licensing requirements

- Application Virtualization solves a lot of issues in comparison with the traditional MSI packaging technologies. The advantages of application virtualization make it easier to create good functioning application packages. Note that the process for the creation of a standardized virtualized application can vary from application to application and each application virtualization platform has its own set of best practice guidance and design criteria. Of the many such platforms available, WWT recommends the use of any of, or a combination of XenApp from Citrix, App-V from Microsoft, ThinApp from VMware. Each solution while distinctly different generally adhere to the following general design patterns related to standardizing the creation of virtualized applications:
    - Use the operating system the application will be run on
    - Know the application installation process.
    - Make note of required reboots during the sequence process
    - Shutdown all other programs
    - Standardize Field Naming Conventions
    - Investigate the monitor result and remove unnecessary components
    - Document both the installation and virtualization steps. Although this sounds logical, many customer lack full documentation of the installation processes. This makes it very difficult to rebuild or upgrade an application. Therefore, it's essential that both the application installation steps as the virtualization steps are fully documented.
    - Applications virtualized are running in separated layers. Those layers cannot communicate with each other, therefore packaging an application that require other applications to function properly should be considered carefully. With the latest versions of application virtualization products, it is possible to specify that application in a separate layer can be called in another virtual layer. With this new feature there are three possibilities for middleware applications:
        - Middleware available in the operating system layer
        - Create a separate package for the middleware
        - Include the middleware in the same package
    - Install all components of the application
    - Revert the Snapshot to the default configuration
    - Disable "Auto Update" features. Some applications have the ability to check a web site or a server for the latest application updates. This feature should be turned off, as version control should be performed by upgrading the package in a controlled way by the IT department.
    - Package as much as possible in single pass of the installation phase. Package as much as possible in a single phase. Do not stop the record/monitoring process after each installer, unless the application requires a reboot. When dialog pops-up for reboot carry out that task, followed by stop record/monitoring step. The following application can be installed by selecting begin monitoring again.
    - Leave User Access Control enabled

- The use of virtualization can allow a company to move forward and accomplish more while using less computing hardware and financial resources. In order to accomplish this through the use of virtualization there are several things one must do to ensure a smooth transition to this technology. The path to virtualization is not easy and should not be taken lightly, but its benefits are quickly visible once the endeavor is complete.

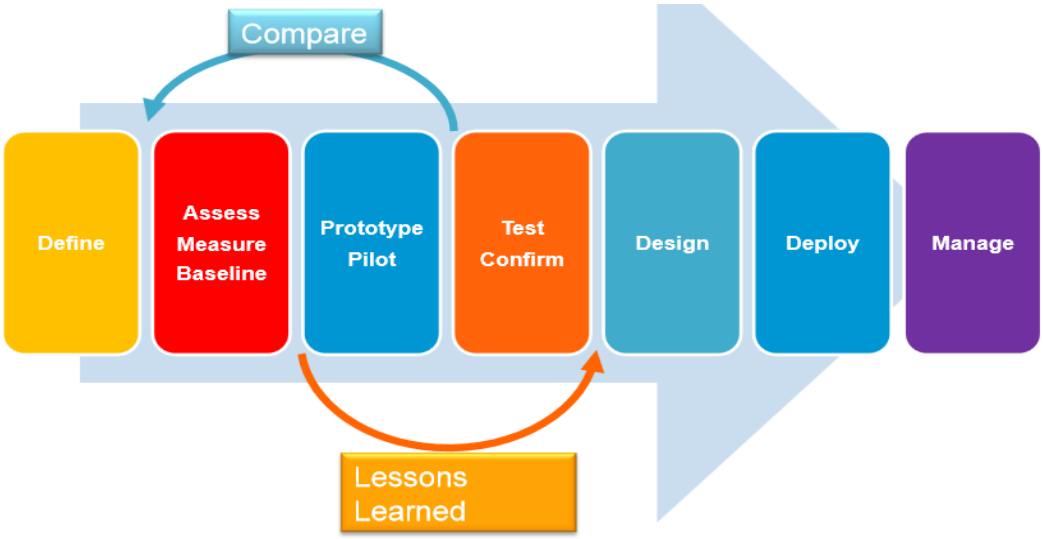

Figure 3: Path to Virtualization

As noted earlier the process for the creation of a standardized virtualized environment can vary from Vendor to Vendor. Each virtualization platform has its own set of best practice guidance and design criteria. Designing and implementing a virtualized environment requires time and a thorough understanding of your needs and available resources. Rushing through the planning phase can lead to more expensive measures later and reduce the efficiency of your virtualization project. If not done properly the move to virtualization can reduce current productivity and decrease availability of critical servers and resources.

It is reasonable to expect near total host server resource utilization when spread across multiple virtual machines. But keep in mind that all servers, virtual and physical, must have enough resources to perform their assigned tasks. An underpowered virtual machine or host server undercuts the benefits of virtualizing your environment. Pay attention to hardware compatibility with the software you intend to virtualize. `

Your software for virtualization, the virtual servers, and the host server all must have hardware compatible with the software you intend to use. Many virtualization packages will emulate a standard set of hardware regardless of that on which it resides. This will allow you to have a standard hardware setup for each virtual machine and reduce time and difficulty in managing driver versions across your virtual environment.

Virtualization increases the hardware requirements for your solution.
You must have enough processing power, storage, and RAM for your virtual machines, software, and virtualization overhead. This means your host servers must first have enough resources to support your virtualization software of choice, plus the operating system and software used within the virtual machines.

Know when not to virtualize. Part of ensuring your company's virtualization initiative is successful is knowing what to and what not to virtualize. Applications that make use of large amounts of a single server's available resources frequently are not good candidates for this technology. Their performance and reliability is typically better suited to their own hardware.

Understand why you want to virtualize your systems. The benefits of virtualizing are numerous and understanding which ones are best for your environment will ensure you choose the best software, hardware, and which systems to implement with this technology.

Once you have all of these items fully understood and ready for your initiative, it is time to move forward and decide which virtualization engine to use and prepare your plan. Failure to take all possible aspects of the plan and the risks into consideration can lead to more work and greater expense after the project is under way. The list of virtualization software is extensive and continuing to grow as companies create new products for the available operating systems. Each one has its own utilities, advantages, and disadvantages for use. Evaluate each one carefully before deciding which one to use.

- VMware (www.vmware.com)
- Microsoft (www.microsoft.com)
- Citrix Systems (www.citrix.com)
- Oracle (www.oracle.com)
- Virtual Iron (www.virtualiron.com)
- Parallels (www.parallels.com)
- InnoTek (www.virtualbox.org)
- Amazon EC2 (aws.amazon.com/ec2)

Clearly and uniquely identify all virtual machines.
In a virtual environment it is easy to have numerous machines doing similar tasks. It is also easy to move a virtual machine from one physical server to another, sometimes transparently. Due to this ease of movement and manipulation it is important to define and adhere to a set of identification standards that denote what the server does, what resources it supplies, and who has access to it. This identifier must be one that can move with the virtual machine regardless of hardware or network configuration as MAC addresses, IP addresses, and even server names can change during a server move. Keeping the identification of the virtual machine is important to prevent loss of access or location confusion.

Do not accept default file naming conventions as decided by the software as you install it.
In order to keep control of your data and know what types of files each virtual machine uses create a standard naming scheme that all virtual machines will use. This will allow you to know which virtual machines access which files, and who has access to those files.

Remember the command line interface.
Your virtualization software of choice should have a strong set of command line tools to allow scripting, scheduled tasks, and other maintenance utilities that will make activities such as rebooting and backing up simple and up to date.

Build a virtual machine library.
Rather than create, install, and configure each virtual machine from scratch, create a set of virtual machine templates for each server type you will need. For instance, if you have a standard setup for database servers, web servers, and file servers, create a base image for all these types of servers. Install your typical server-specific software in your image templates and store this securely. When you need to create a new virtual database server, you can retrieve the base server image and configure it to your current needs. This will decrease the time it will take to set up the machine properly, and ensure your new virtual machine meets your minimum standards of security and software configuration for its type of machine.

Maintenance
Once you have made your decision on which virtualization platform is best for your environment, you should begin researching the available automatic monitoring packages for virtual environments. Maintenance and monitoring are as important in the world of virtualization as the services they provide. Without the means of monitoring and reporting on your servers, you run the risk of losing control of your environment. Monitoring and reporting on them consistently and with a reliable system ensures they are available when you need them.

Monitor your entire environment, physical and virtual.
All of your servers are going to constantly need resources spread throughout your infrastructure. Put automatic monitoring systems in place to ensure none of your systems lose access or need resources that are suddenly unavailable. Pay close attention to your system's health. Network paths, storage locations, and data sources need to be constantly monitored to keep them online and free from errors that could interrupt your business flow. Also, policy compliance settings need to be monitored to ensure no users or virtual machines take up more resources than they are allowed. Using automatic monitoring methods will ensure your data center team is alerted to any anomalies when they occur and enable them to take action before such situations become critical.

Schedule and perform regular system audits.
Virtual machines are in a constant state of flux. They can be easily moved, renamed, or relocated to different host systems. Perform system audits to ensure each machine meets company security policies, but also ensure the work load is being spread evenly across the virtual servers. It is possible, through standard day-to-day operations as business needs grow and change, that one virtual server can be using all its resources while another is nearly unused. These audits will help you identify points in your infrastructure that need more resources and those that can use less, along with giving yourself a much needed opportunity to balance server usage and ensure resource availability.

Automate system audits.
It is possible that you may be managing more than a few virtual machines. In this case it is impractical to try to manually perform system audits for all of them. This is where automated tools become a necessity. These tools will monitor your virtual machines' performance and workload, recording statistics and reporting on situations where servers are performing less than ideally or being over-utilized. These tools can also make suggestions on what will be needed to alleviate strain or improve the performance of a given virtual machine.

Track each virtual machine's life in your data center.
Virtual machines can have nomadic lives as physical servers are upgraded or replaced. Virtual machines can be renamed, assigned new tasks, storage servers, and data locations with ease. This can cause loss of data access or loss of resources if care isn't taken. Each virtual machine needs to be tracked and all changes within it recorded including where it was created, when it was created, where it moved from, where it was moved to, and where it now resides. This information can help ensure each virtual machine is understood in its context and can ease the difficulty of understanding server needs when changes must be made.

Manage virtual servers in groups rather than all as individuals.
Not all servers will be able to follow the same rules for access and configuration. However, machines with similar job duties and user access can be managed as a group allowing for simpler management tasks that will be applied across multiple machines.

Manage ownership.
Each virtual machine should have a person or group responsible for its upkeep and maintenance once it is up and running. This will ensure the virtual machines are performing their required tasks correctly, and will also ensure that no machine is overlooked or neglected as business continues.

Know what's in your environment when it arrives.
Automatic discovery of new servers and new virtual machines is an expected ability of monitoring packages. IT managers cannot manage or control what they don't know about and these tools will ensure all new machines are detected, scanned, and audited for compliance with your standards. This will ensure that no virtual machine is overlooked during a maintenance cycle because no one knew it existed.

Know when to retire a virtual machine.
Like physical machines, virtual machines can outlive their usefulness or simply become obsolete. In many cases virtual machines were created to meet temporary needs for specific situations that have passed and the machines are no longer needed. In other cases, the business has moved on and the reason for a machine's existence has disappeared. Regular audits can ensure these unneeded machines are taken out of service and archived. Once this is done, resources can be freed up for other services that need them or new virtual machines can be put online to meet a current and immediate business need.

Real-time monitoring is the key to maintaining success in your virtual environment.
Many changes can happen very quickly and your support team must be aware as soon as problems arise to be able to respond and resolve without those changes affecting your data, business, or customers. These monitoring packages can help you identify under- and over-utilized systems, recognize virtual machines that have fallen out of compliance with your configuration requirements, identify and alert you of new virtual machines, prevent unauthorized virtual machines from coming online, tracking changes within virtual machines and within the servers on which they reside, and in general give you the ability to know and control your environment without the cost of a large IT team. Your business requires a high level of resource availability and will allow no excuses for extended outages. Real-time monitoring and reporting on all of your virtual servers allows you to meet this challenge on time and within budget.

- **What industry practices are used when application source code is not available?**
  When application source code is not available, the pathway forward is both limited and limitless. It is understood in the technology world that within most organizations there always seems to be at least one business critical application that resides on an old server in the previous datacenter and it is understood this server is never to be touched. *Why?* The Application vendor is out of business, the installed floppy disks have been lost long ago the reasons are many.

  As an IT organization, you have few choices what to do. When source code is no longer available we recommend the following industry practices:
  - Maintain obsolete hardware and O/S platforms
  - Reverse engineer of applications
  - Require the use of shims and or emulators
  - P2V of physical workloads
  - Hire IT staff with specific application knowledge
  - Evaluate and purchase modern software

4. Tier 4: System/Application and Data Migration and Hosting:

- When evaluating a hosting environment for a particular application, the decision for selecting a set of technologies are often narrowed down to these five factors:

Factors for Determining System/Application, Data Migration and Hosting	
Availability	The hosting environment should consist of a highly available solution that has multiple points of redundancy within both the hardware and software for any mission critical application to ensure it is available when needed. Enterprise class hardware and software with high levels of Quality of Service (QoS) and Quality Assurance (QA), which have a solid reputation within the data center, are recommended.
Accessibility	The hosting environment should always be accessible to its users. The environment should consist of industry standard Commercial Off the Shelf (COTS) hardware and software to maintain standard network compatibility and interoperability. Standardization of the hardware and software components of the hosting environment will help ensure accessibility.
Performance	Performance requirements may be the most difficult to quantify without a network evaluation prior to deploying the application in a hosting environment. Performance considerations must be made at the hosting environment, the network, and the user level. The hardware and software platforms must allow the application to perform sufficiently. The network must allow the application to interact with other network services and allow the user to access the application as needed without added latency.
Price	The hosting environment for applications must be sized appropriately to allow a solution to be built that is not over-scoped or under-scoped. Selecting intelligent hardware and software with high levels of efficiency, such as disk auto-tier technology, de-duplication, compression, integrated backup and recovery, and virtualization of hardware will allow the best utilization of resources across the platform.
Supportability	After the hosting environment is deployed it should be easily maintained by the support staff. The hosting environment must also have a clear and easy path for hardware and software in the future as technology and requirements change drastically every few years.

- To transfer large amounts of data across geographic regions, the most efficient methods include multiple components. Storage Area Network (SAN) or Network Attached Storage (NAS) appliances are best used to manage the actual replication of data. SAN and NAS technology that can de-duplicate and compress the data to ensure only unique blocks of data are transferred, improving the transfer capabilities significantly. To assist the data moving across the Wide Area Network (WAN), bandwidth accelerators are recommended to further reduce duplicate data, select optimal paths, perform Transmission Control Protocol (TCP) and User Datagram Protocol (UDP) optimization, prioritize network traffic, and application level protocol optimization.

- Data life cycle management tools exist for both hardware and software solutions. Intelligent auto-tiering of data can occur on NAS and SAN appliances, as well as through various specific software management solutions to include backup software. Archival of obsolete data should occur to less significant storage and relevant data should be available on-demand through primary data storage. Data retention rules can be set to customize data archival. Selection of the correct solution set requires understanding the data and the life cycle requirements.

- It is recommended that the transition occur through a phased approach, migrating non-critical applications first. COTS applications with acceptable outage windows should follow. Any custom applications or applications with minimal allowed outage periods should go last, after an evaluation of the network and performance requirements has taken place, and an approved migration strategy has been accepted. Applications can be migrated by creating a fresh install, backup and restore, virtualization, instant-virtualization, or physical transfer and integration to the new hosting environment.

- Network visibility products play a critical role in the visibility and security of any IT environment. Unlike traditional signature based tools we find these provide great insight on emerging threats. For example, most SIEM products won't recognize a small spike in network traffic by a few MBs. However, if the percentage of incoming versus outbound traffic changes significantly this could be an indicator or compromise. Network visibility products give an insight on network and virtual based traffic and are based on Netflow and NBAR stats.

  Google and other host providers deploy significant measures to maintain the highest level of security and integrity by employing tools that provide:

  - IP Policies
  - Public Key Infrastructure (PKI)
  - Cryptography
  - Digital Signatures
  - Secure Sockets Layer (SSL)
  - Encryption File System (EFS)
  - Authentication
  - Secure Network Connectivity
  - Firewalls
  - Secure Gateways
  - Detection
  - Compliance

  In order to ensure the integrity of customer assets, the host provider will need to conduct a thorough assessment of the customer's infrastructure, policies, and environment in order to design a plan that meets the government's needs.

- WWT recommends a solution that specializes in detecting and responding to risked based attacks. Including these tools in your overall security solution can typically evade traditional signature based approaches. This allows an organization to gain insight and visibility into the security posture of their enterprise. Lancope's Stealthwatch is one tool to consider and is DoD approved. Additionally, the customer should consider a solution that includes a management console, a flow collector, and flow sensors. Sizing the solution is directly dependent on the size of the network and how busy it is. For example, a relatively small deployment may consist of 1k flows while larger organizations will have between 10k to 25k flows.

- Commercial hosting environments are primarily priced by measure of processing, memory, data storage, and availability requirements. Operating Systems (OS) are virtualized and accessed through the commercial hosting platform where the application will run. Each OS and application will require a specific amount of memory, CPU, and networking to perform efficiently, as well as shared data storage space. Cost is based on the amount of resources required to meet the customer's demand. Capacity fluctuates with consumption and can be utilized by increasing (or decreasing) the

amount of resources required. The commercial facilities are extremely redundant in terms of power and connectivity, and infrastructure costs are included in the resource utilization charges so it eliminates added cost of real estate and power.

Specific commercial facilities are approved for high levels of data classification and must meet any and all requirements of the US Government to maintain classification levels. Network security boundaries can be owned and maintained by the customer within the hosting facility. Hardware maintenance can be limited to cleared personnel or owned completely by the customer, if required. Software maintenance can be performed remotely by customer personnel, also. Additional benefits include removing the necessity to purchase and maintain computing hardware, as well as, 24 x 7 support. Disaster Recovery and Continuity of Operations (COOP) options are also available through the use of one or more commercial hosting facilities.

- COTS applications can be easily integrated into industry standard hosting environments as they are designed to work with hardware and software platforms that are highly used and highly supported. GOTS applications may or may not fit within this model depending on the source-code, age, and hardware/software requirements. Migration to the hosting model might require new coding if the application requires unsupported technologies. An evaluation of GOTS applications is recommended prior to designing a hosting model.

- ISO 22301:2012, Societal Security – Preparedness and Continuity Management Systems – Requirements. Specifies requirements to plan, establish, implement, operate, monitor, review, maintain, and continually improve a documented management system to protect against, reduce the likelihood of occurrence, prepare for, respond to, and recover from disruptive incidents when they arise.

  ISO/IEC 24762:2008, Information technology – Security techniques – Guidelines for information and communications technology disaster recovery services. Provides guidelines on the provision of information and communications technology disaster recovery (ICT DR) services as part of business continuity management, applicable to both “in-house” and “outsourced” ICT DR Service providers of physical facilities and services.

ISO/IEC 27031:2011, Information technology – Security techniques – Guidelines for information and communication technology readiness for business continuity. Describes the concepts and principles of information and communication technology (ICT) readiness for business continuity, and provides a framework of methods and processes to identify and specify all aspects (such as performance criteria, design, and implementation) for improving an organization's ICT readiness to ensure business continuity. It applies to any organization (private, governmental, and non-governmental, irrespective of size) developing its ICT readiness for business continuity program (IRBC), and requiring its ICT services/infrastructures to be ready to support business operations in the event of emerging events and incidents, and related disruptions, that could affect continuity (including security) of critical business functions. It also enables an organization to measure performance parameters that correlate to its IRBC in a consistent and recognized manner.

Experience suggests the data center should provide the COOP as a service, the application owners should be responsible for ensuring the backup and recovery model is sufficient for their specific requirements and fit within the Enterprise strategy.

5. Tier 5: Integration of the entire DCC effort across Tiers 1-4:

- Engineering governance can be summarized as two questions: *Are we doing the right things? and Are we doing those things right?* Enterprise IT must always be on the lookout for the latest technologies that allows it to run with fewer resources while providing the infrastructure to meet current and future customer needs. Governance must address four aspects:
    - Meeting legal requirements for health and safety, probity, etc.
    - Ensuring the development, at acceptable risk, of competitive offerings for its customers
    - Ensuring the offerings are to specification
    - Delivering the offering to the customer to the business benefit of the enterprise.

From these, a given project will establish its own business objectives. For each of these objectives, it will be necessary to develop their own governance process which cover the business objective stakeholders with appropriate metrics and with a process owner responsible for good governance for that objective. Typically this individual turns out to be a Chief Engineer. Chief Engineers usually have responsibilities for all four of the aspects, above, so most of the governance processes will be owned by such an individual within the project, and streamlined processes can therefore be adopted.

Some governance mechanisms will already be in existence, design reviews, for example. Others may need to be extended, or developed; appraisals of individuals and their contributions, for example. Still others may be ignored; it is unlikely that the governance of all processes will be cost-effective or even feasible.

There are several key issues that must be borne in mind in developing governance:

- Essentially, governance involves people. Therefore, the processes and the metrics must be people-sensitive (i.e. as unobtrusive as possible).

  - Governance should measure only that which is necessary to achieve the business objective. It is a mistake to try to measure everything.
  - There must be clear responsibilities to act. This should include disciplinary processes as well.

- Interoperability is the primary concern between standardization and technical flexibility. There are more advantages to standardizing the architectural compliance of a hosting platform than disadvantages. Standardization equals interoperability, reduced training and supportability, simple upgrade paths, etc. Not all hardware and software can be considered a commodity. Storage vendors do not natively replicate data to each other. Third party solutions are inefficient and complicate the infrastructure. Standardization disadvantages exist due to the probability of favoritism and price gouging. An un-biased solution from a large Value Added Reseller (VAR) with a wide range of OEM partnerships would help provide a cost effective, and custom tailored solution by removing the element of favoritism. A large business would less likely be swayed by one OEM partner due to market share and economic stability.
- System Integrators (SIs), while well versed with government operations, tend to be manpower, process, and overhead heavy leading to less consolidation over time and higher costs. The current involvement of System Integrators in the customer's Data Center Consolidation efforts are part of the reason the customer is interested in hearing from industry to learn strategies to maximize cost savings and avoidance, increase efficiencies and standardization of the customer's Data Center consolidation and operations. In order for the government to attain a lowest cost solution to efficiently consolidate an increased number of sites with over 12,000 servers by FY19, or earlier, the government needs direct access to industry leading, OEM independent, technology firms with proven consolidation methodologies to develop a clear path to IT savings and efficiency.

  It would be in the government's best interest to have direct access to these firms or minimize System Integrator involvement to oversight during design and implementation. This project is best suited to a large technology VAR specializing in data center virtualization and consolidation, who can bring to bear hosting and small business partners with specialized experience to maximize results and reduce overall project costs.

- A detailed roadmap of people, processes, and technology is required and developed across the Tiers before the design and architecture can be determined and operational governance along with associated gates must be identified. As it relates to the organization's desire to consolidate, these are critical phases (at a very high level):
  - Datacenter Strategy Validation
  - Detailed DR, Cloud, and Systems Designs
  - Datacenter, System, or Application Migrations

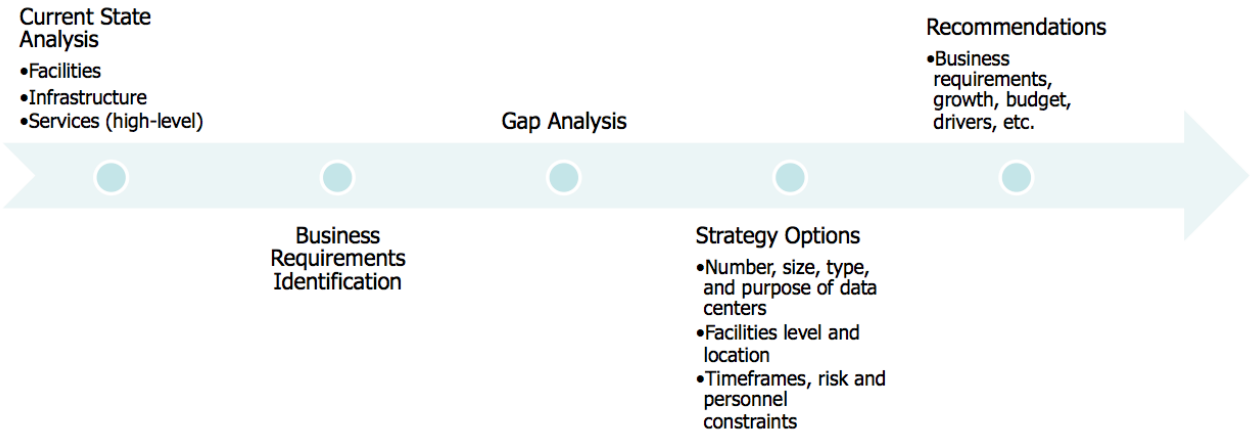

Figure 4: Critical Phases

Overall Approach:

- Collect current state facilities, architecture, operations, service level and financial data
- Review business requirements
- Develop strategy and tactical evaluation criteria and weighting
- Develop gap analysis of requirements versus current state
- Develop data center strategy and tactical options
- Rank options using criteria and weighting
- Facilitate executive review and decision making

The customer should consider that in order to meet its goals for data center consolidation and application optimization, a collaborative group of technology experts will be required to span across the four Tiers. The prime contractor should have the technical skills and experience to offer design and engineering alternatives while leveraging a select group of specialized partners to execute ideas to deliver and deploy a highly secure, yet automated hosted environment. Typical resellers and integrators are setup to cater to either technology or to applications, specifically. It is recommended to take a "general contractor" approach to consolidation and modernization of the datacenter. Areas such as business process re-engineering require specific skillsets. Application modernization is a different aspect from application virtualization, both of which require special skillsets over traditional server virtualization.

WWT refers the "layers of the data center" as: technology services, applications, infrastructure (including people and processes), and facilities. Each can be outsourced, separately, if in line with the mission's requirements. Often security requirements dictate that specific areas must be kept in house. For example, "moving to a hosted environment" might simply mean the facilities are owned externally, but the application and its infrastructure used to provide a service for the mission is still owned internally. Within these layers comes application evaluation, modernizing or virtualizing those applications, and how they are secured.

The process for application must be rolled up to the service layer for success. This is a direct tie into earlier mentions of how to assess an application. It is critical to place gates within the process that allows for key decisions to be made outside of organizational issues, "people and processes". Consider raw application to service mapping, costs (both OpEx and CapEx) that tie across to the mission's service and downstream into infrastructure and facilities. With those areas clear, decisions can be made if the application needs modernization, virtualization, decommissioning, etc.

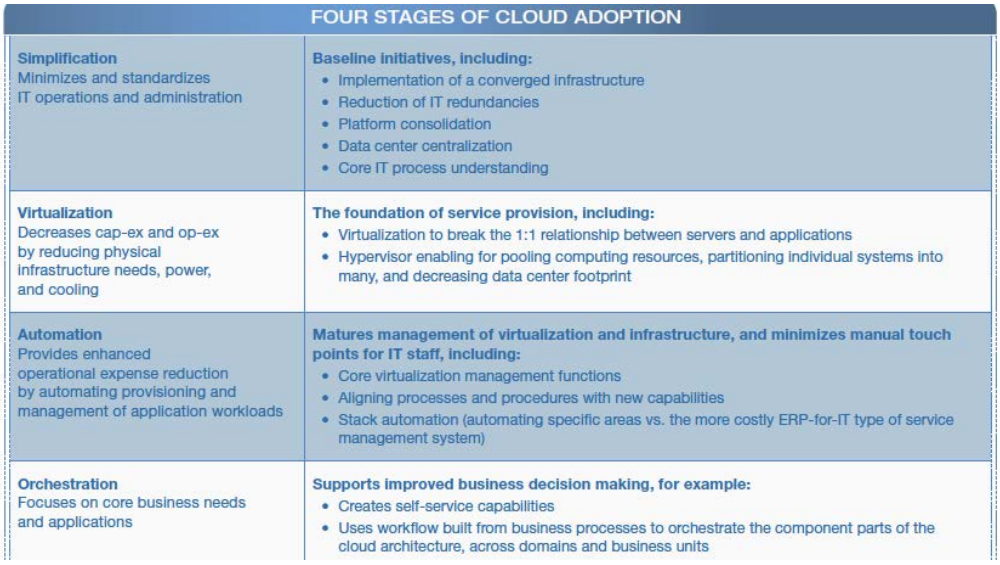

FOUR STAGES OF CLOUD ADOPTION	
Simplification Minimizes and standardizes IT operations and administration	**Baseline initiatives, including:** • Implementation of a converged infrastructure • Reduction of IT redundancies • Platform consolidation • Data center centralization • Core IT process understanding
Virtualization Decreases cap-ex and op-ex by reducing physical infrastructure needs, power, and cooling	**The foundation of service provision, including:** • Virtualization to break the 1:1 relationship between servers and applications • Hypervisor enabling for pooling computing resources, partitioning individual systems into many, and decreasing data center footprint
Automation Provides enhanced operational expense reduction by automating provisioning and management of application workloads	**Matures management of virtualization and infrastructure, and minimizes manual touch points for IT staff, including:** • Core virtualization management functions • Aligning processes and procedures with new capabilities • Stack automation (automating specific areas vs. the more costly ERP-for-IT type of service management system)
Orchestration Focuses on core business needs and applications	**Supports improved business decision making, for example:** • Creates self-service capabilities • Uses workflow built from business processes to orchestrate the component parts of the cloud architecture, across domains and business units

Figure 5: Cloud Adoption

While updating technology seems inconsequential, often integration of an agile, consolidated, next-generation data center significantly impacts the overall organization. For instance, re-structuring existing processes – currently policies no longer work in the new environment. Operational control of a next-generation, highly virtualized cloud (public or private) contains very different needs than the data center of days past. With this comes the need to focus on people, processes, and technology –equally- to be successful.

6. Metrics & Instrumentation

- Whether measuring the progress of a specific project or the security posture of an entire enterprise, some of the more dynamic tools and product suites available can be used independently or combined for continuous monitoring capabilities and may include:

Area of Measure	Available Instruments	
Compliance Automation & Monitoring	Security Associates MetricStream	Archer Xacta
Detection & Prevention	Trustifier Symantec McAfee Tenable	Gold Disk Retina ThreatGuard RSA
Project Management	Clarizen EasyProjects	LiquidPlanner

Area of Measure	Available Instruments	
Asset Management	SAP OpenView CA,	Altiris Maximo
Event Correlation	ArcSight Lexis Nexis	Splunk Loglogic
Data Collection	Security Associates,	BMC EMC
Work Flow	Integrify	
Security Information and Event Management (SIEM)	ArcSight Splunk Tripwire	NetIQ Symantec Quest Q1Labs
Incident Management	Remedy	Ad Hoc
Change Management	Security Associates	BMC EMC

While no one product offers true continuous monitoring of an IT enterprise, customized solutions are available using Commercial-off-the-Shelf (COTS) hardware and software combinations that can be scripted to accommodate any environment providing in-depth visibility, broad scope, and flexible control over security posture and compliance.

For instance, the customer will want to consider the following when implementing a fully integrated continuous monitoring and compliance solution to span across Tiers 1-5:

Step 1 **APPROACH**	**Step 2** **ANALYSIS**	**Step 3** **CONTROLS**
• Readiness Assessment • Governance (processes & policies) • Change Mgmt • Incident Mgmt • Risk Mgmt • C&A • Encryption • Disaster Recovery • Existing Technology & Gaps • Measurements • Asset Mgmt	• Asset Inventory Assessment • Discovery Assessment • Risk Mgmt. • Assessment Tools • Compliance Methodology	• Operational Controls • Management Controls • Technical Controls • Manual & Automated Processes • Customization
FRAMEWORK	**DELIVERABLES**	**REPORTS & MEASURES**
• Inventory Assessment • Existing Capabilities • Technology Map • Functional Compliance Map	• Gap Analysis • Remediation Plan • Implementation Plan • Maintenance Plan • Training	• Compliance • C&A • Critical POA&Ms • System Status & Changes • Security Posture & Status • Remediation • Incidents & Actions • Exposure & Risk • Notifications & Alerts • Cost & Efficiency • Executive Overview

7. Data Center Design

- Many organizations are moving towards either a containerized or modular data center in place of the traditional brick-and-mortar. There are five key reasons for this movement:
    - Time-to-Mission
    - Efficiency
    - Scalability
    - Color of Money for Purchase
    - Reduced O&M

Time-To-Mission

Time-to-Mission is the length of time between when a new data center is approved and when it can be occupied. While a traditional brick-and-mortar data center usually takes five years or more for funding, approvals, and construction, a containerized or modular data center can be onsite and operational in as little as 12 weeks, but more commonly 16-22 weeks. For ISO containerized solutions—we call them mobile data centers or MDC's—the equipment inside can be pre-racked, stacked, and configured so that when it arrives onsite, it can be operational even more quickly. We classify the MDC in to three categories depending upon how it will be used: *mobile*, *transportable*, and *fixed*. The mobile units are just that, highly mobile. They can be loaded in to a C-130, C-17, or carried by a Chinook.

Mobile use cases include:

- Tactical command centers
- Shipboard training environments that can be moved to land during a refitting
- Supporting infrastructure for an event such as Hurricane Sandy or the 9/11 attack
- Movement of large amounts of data during a data center consolidation project
- A temporary data center during a data center refresh
- Many others

We use the term *transportable* to mean that the MDC is primarily fixed, but can be moved several times during its life. This can be especially helpful for environments where space is being leased or in the event that a BRAC may occur. Using the transportable MDC, the organization is not investing in a brick-and-mortar building that will need to be vacated in a couple of years. It can also be used in office complexes to offset the cost of building out leased building space as a data center or a SCIF data center. The MDC can be located nearby or even in the parking lot or parking garage. Transportable MDCs also can make programmed IT refreshes faster and easier by replacing the entire container.

A *fixed* MDC is one that is likely to remain in the same location for all or most of its life. These units may be especially good candidates for green large adiabatic cooling systems.

Efficiency

One of the key advantages of a MDC or a modular data center environment is efficiency. Largely gone are the days of large CRAC or CRAH air units along the walls blowing cold air indiscriminately under raised floors to cool racks rated to support 25 kW. Those systems were inefficient and power consumption of equipment has dropped dramatically in recent years. Having performed many site assessments, we have found just about everything at one time or another gets stored under the raised floor blocking the airflow.

Today's best practices include the following to reduce the PUE to 1 or below.

- Isolation of the hot aisle or cold aisle or both
- Use of variable speed precision cooling systems
- Use of green techniques
- Location of the data center

While many of the following principles can apply equally to brick-and-mortar, MDC, and modular buildings, we find that the latter two are more likely to benefit the most since they are designed that way from "ground" up rather than being retrofitted with a variety of constraints that entails.

The primary intent of hot and cold aisle isolation is to eliminate the mixing of cold air and hot air in common spaces to maximize the cooling capacity. There are various arguments whether cold air isolation is better than hot air isolation and both can have valid points. Building a computational fluid dynamic analysis model during the design phase of a solution will ensure the design is valid and enable making design adjustments that would be difficult once the data center is built.

Unlike the CRAC of old that essentially floods the space with cold air, precision cooling provides the amount of cubic feet and temperature needed for the load of the associated racks. The cooling systems are more strategically selected and placed based on the equipment in the racks. In the early days, operators often wore parkas in the data center because it was kept cold. Today, the cold aisle can be much warmer. Many devices have operating temperatures of 95°F or higher. Raising the operating temperature to 85° or higher can have an enormous impact on reducing the cost of cooling as well as open alternative cooling solutions that may not have been possible before.

Many containerized and modular data center companies build one or two models with a fixed design and limited customization. Most convince the customer they need 25 kW of power to each rack. While that was once true, it is not common in today's market. We see an average of 8-12 kW per rack in roughly 85% of the cases, up to 15 kW in another 10%, and above 15 kW in only about 5% of cases. The cost to over engineer for power per rack that will not be used can easily reach in to multiple hundreds of thousands of dollars to purchase and easily double or triple the cost for continued service. In contrast, it is recommended that an analysis of the typical operating power the actual equipment uses is the best way to size your power, and therefore cooling,

requirements. *Then* design the containers and supporting elements to those requirements. This will also provide the best case PUE values.

Component	Custom Config	Model / ~List Price	Fixed Config	Model / ~List Price
Rack	20 x 9 kW = 180 kW		20 x 25 kW = 500 kW	
Chiller	180 kW = 51.18 tons → 60 ton	70 ton = $67,000	500 kW = 142.17 tons → 160 ton	150 ton = $107,000
UPS	190 kW x 1.2 = 228 kW → 250 kW UPS	$260,100	500 kW x 1.2 = 600 kW → 500 kW UPS	$361,000
Generator	300 kW generator	$80,500	800 kW generator	$318,000
Total		**$407,600**		**$786,000**
Customer saves approximately $378,400 at time of purchase alone. *Note: The figures shown are representative, but not final pricing.*				

Another design feature that we frequently see is that large double doors need to be opened to access the rear of the racks. The impact may not be so great when the container itself is inside a conditioned space, but can be disastrous when deployed outside, especially in austere environments. Opening the doors exposes the racks to rain, snow, dirt, sand, etc.; all things one does not want in a data center.

Scalability

An MDC approach offers scalability that is difficult to provide with a brick-and-mortar building. When designing a building, the first thing that usually happens is that the people take the size of the current data center and double or triple it to provide for growth over time because the building process takes so long, typically at least five years. The cost of the brick-and-mortar data center sky rockets and the efficiency plummets because the sizing of power (utility power, UPS, and generator) and cooling solutions now have to be based on the double or tripled design that may not be needed, at least for a while.

An MDC or modular building can be designed to scale as needed. Additional modules are added as necessary. The modules can be in various building blocks that have a number of racks with x, y, or z kW or a completely unique average per rack based on the actual load of the equipment to be installed. The supporting elements of power and cooling can be designed for these same building blocks, or they could be designed to support multiple building blocks for additional efficiency and reduced footprint. For example, instead of three modules each having a single 10-ton chiller, we could install a single 30 ton that would be more cost effective. We can design N+1 or 2N systems as well. A true ISO design can also be stacked when there is more vertical space than floor or ground space.

Reduced O&M

Everyone wants to reduce the O&M costs of the data center. Power and cooling are often the two largest costs for a data center. Using constructs that use green or more efficient designs and technologies will reduce the O&M costs of the data center.

- WWT employs these practices in all of our designs. We have achieved a 1.3 or better PUE in many designs. We look at the design from a holistic approach where everything has an impact on the final design in terms of both efficiency and function. WWT uses a single team to design, build, deploy, and support WWT's FAST product suite (micro-containers, MDC, and modular buildings). Members of the team are from two or three companies (depending on the project), but work as a single integrated team day-in and day-out. This approach promotes cross-discipline creativity and idea exchange to solve unique customer requirements as well as cross-discipline review of the entire design before getting to the fabrication phase.

  This is a different approach often taken by most other companies in this market where Company A subcontracts to Companies B, C, D, and E who each have different isolated parts of the project, who have very different areas of expertise, and who don't communicate with one another regarding the design. Only Company A knows the final design. This results design flaws, mistakes, and difficulty supporting the finished product. Two examples seen recently are the placement of an emergency power off (EPO) button being placed behind a rack where it could not be accessed and the inability of the customer being able to receive a full set of design floor plans and working drawings.

  WWT performs an analysis of the data center equipment that will be used to determine the actual power load and the most effective design and product selection. WWT also performs a computational fluid dynamics (CFD) analysis during the design phase to ensure the air flow will actually be as expected. It enables us to tweak the design for peak efficiency prior to the customer purchasing the solution.

  Because WWT is a system integrator with top tier reseller capabilities, we can provide a true multi-vendor best of breed solution regardless of whether the subject product is a server, storage, networking, precision cooling, ~~or~~ quality power and support for their physical footprint and airflow. We have a variety of products available depending upon the various use cases and intended locations. We have green adiabatic solutions for applicable locations or ruggedized solutions for austere environments. For example, a standard air conditioning system used in most parts of the United States will not perform as expected in the Afghan summer where it is 130°F. After about 90° the capacity of the standard unit begins to drop off sharply such that a 10-ton unit may produce only four tons at 135°.