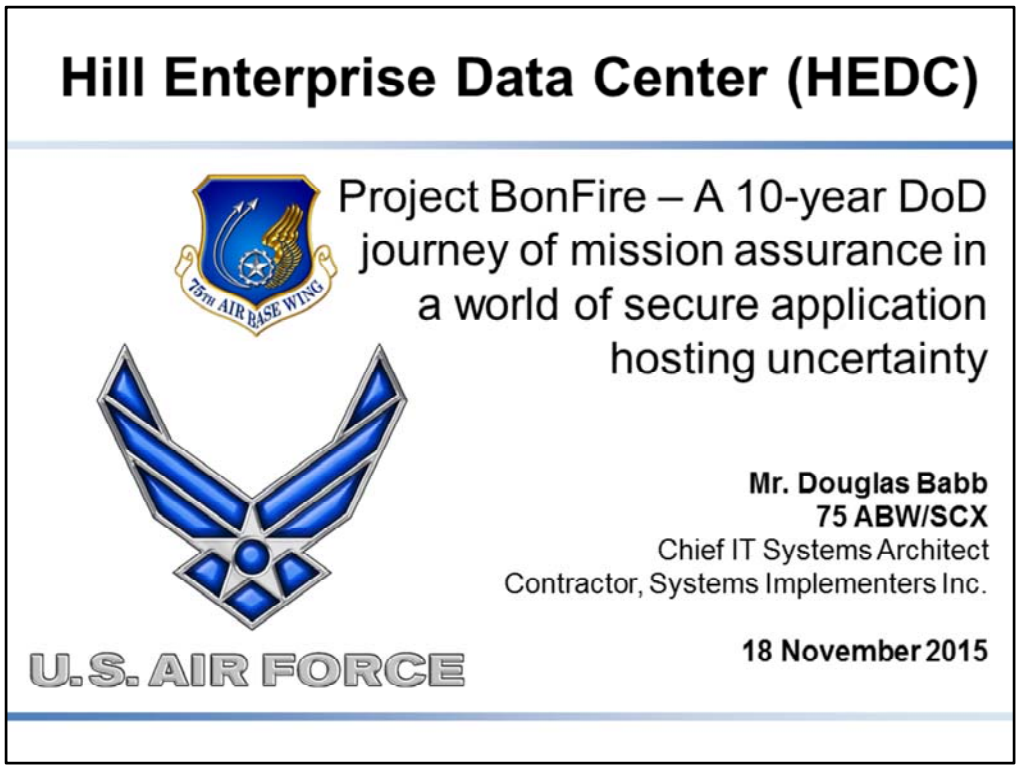

Igniting Change

"*Through Project BonFire, Hill Air Force Base updates data servers and storage systems to assure that 400+ apps for jet and missile maintenance are available 24 x 7.*

Design, reliability and screaming speed. When you think Air Force, these three things typically would be synonymous with jets. But they also drove the thinking of the Hill Air Force Base systems staff when they set out to upgrade the Air Force Materiel Command's data center operations.

The command at the northern Utah base repairs and maintains jet aircraft and intercontinental ballistic missiles. The mission: Keep these aircraft and missiles ever-ready for war. To do their jobs, military personnel rely on over 400 applications. Although the apps worked fine, in recent years, server sluggishness and downtime had become a problem, says Mike Jolley, chief of the Operational Policy Branch and program manager for the command's computer center. "

Source: FedTech » Magazine » "Igniting Change"
http://fedtechmagazine.com/article.asp?item_id=278

Other Articles about Project BonFire:
Government Computing News » "Order out of chaos"
http://www.gcn.com/print/26_16/44607-1.html

Washington Technology » "Something to celebrate"
http://www.washingtontechnology.com/print/22_14/31170-1.html

Gartner » "Linux Case Study: The U.S. Air Force Goes From Big Unix 'Iron' to x86 Linux in 290 Days", By George J. Weiss
Gartner ID Number G00156338, Publication Date: 26 June 2008

Business Wire » "Hill Air Force Base Named First Red Hat Innovator of the Year"
http://findarticles.com/p/articles/mi_m0EIN/is_2007_June_12/ai_n27269703

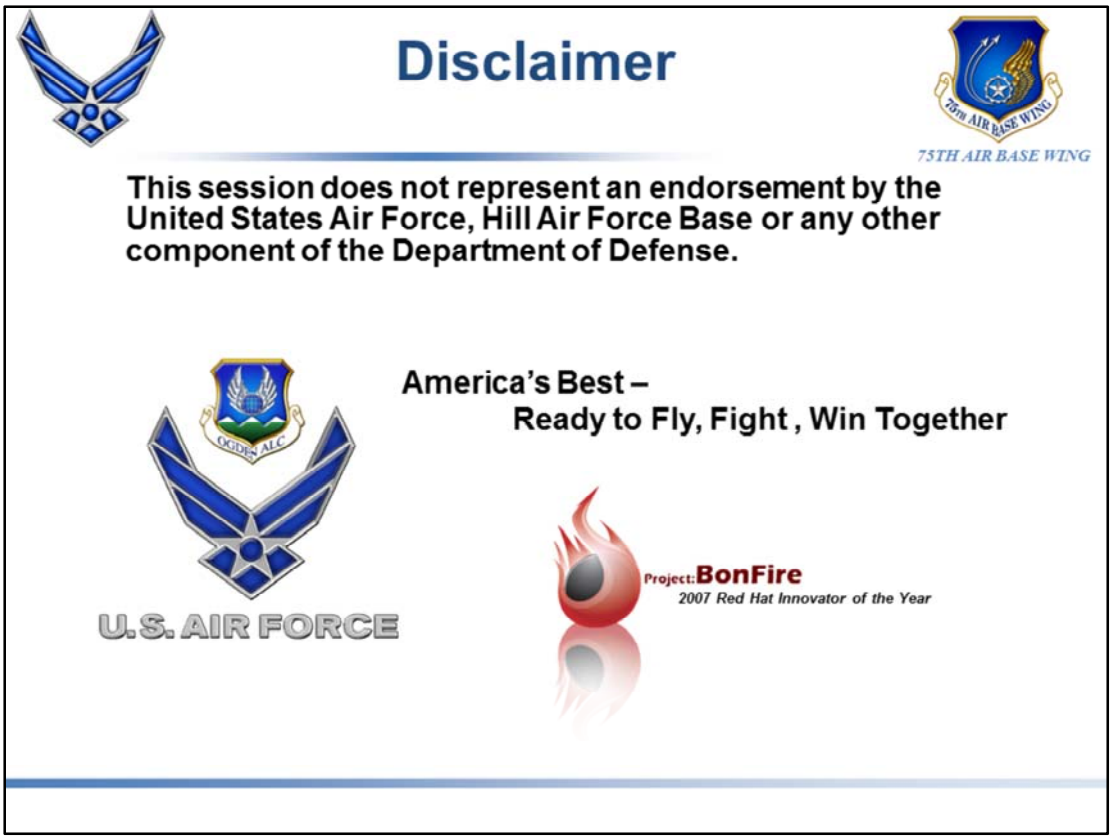

This session does not represent an endorsement by the United States Air Force, Hill Air Force Base or any other component of the Department of Defense.

Current Federal and DoD information technology policy is driving towards data center consolidation and towards leveraging the service and deployment models of cloud computing. As the Installation Processing Node (IPN) for Hill Air Force Base (HAFB), HEDC has the challenge of being both prepared to host a wide variety of applications (including legacy systems) and of being prepared to migrate and refactor applications towards a cloud deployment model. Many critical, vendor-supplied systems must be consolidated to meet Joint Information Environment (JIE) and Air Force Data Center Consolidation (AFDCC) timelines, but most are hardly "cloud native" applications. They are generally monolithic in nature and not typical candidates for a cloud deployment model. Nevertheless, HEDC has received hosting inquiries for large enterprise applications accompanied by the requirement that it be hosted in an IaaS-like cloud environment. The combination of massive data center consolidation and the push to move systems into a cloud model as the primary option will make the next five years challenging. However, HEDC has a successful track record of evolving infrastructure with abstraction and convergence. This evolution has allowed HEDC to modernize systems and host application workloads on commodity hardware and build an engineered Platform as a Service (PaaS) solution that can capably host large enterprise systems at a reasonable cost to system owners. With the challenges of the next 5 years in mind, HEDC has made recent strategic investments to continue

evolving the infrastructure and continue the success in this climate of consolidation and "cloud first."
Hyper-convergence, hybrid-cloud, software-defined networking, and next-generation abstraction via micro-services are all critical strategies for HEDC to fulfill its role as HAFB IPN as well as be a leader in promoting and fulfilling Federal, DoD, and AF information technology policy. HEDC is creating the ability to host a full spectrum of application workloads while constantly evolving the infrastructure to allow for meaningful application rationalization, i.e. application rationalization that allows current software investments to be modernized or "immigrated" to infrastructure that better conforms to the advantages of a cloud model without necessarily requiring a major re-factorization of the software platform. HEDC will be prepared to effectively deal with the next decade of application rationalization which will naturally focus more and more on cloud-native applications as both Federal/DoD policy and the larger course of information technology drive that reality.
Cloud-native application (NCA) rationalization can be broken down into (1) Cloud-Native: new "greenfield" development that is purpose built for the cloud, (2) Cloud-Immigrant: legacy systems that can be migrated more or less unchanged into the cloud and reasonably benefit from the cloud model, and (3) Cloud-Exile: legacy systems that require refactoring to appropriately leverage the benefits of the cloud model. The HEDC infrastructure will be prepared to host applications, provide services, and modernize systems within this moving and confusing spectrum. Legacy government off-the-shelf (GOTS) applications, which are common to an IPN, will still benefit from an enterprise-class infrastructure in which hyper-convergence features prominently. However, where applicable, other legacy GOTS will benefit from a hybrid-cloud capability that allows them to be deployed at HEDC or at an extension of HEDC running in infrastructure belonging to an authorized cloud service provider (CSP). Finally, all HEDC The Federal Government and DISA use the definition of "Cloud Computing" provided by the National Institute of Standards and Technology (NIST). NIST defines cloud computing as "a model for ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources (e.g. networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction." This definition further outlines cloud computing as "composed of five essential characteristics, three service models, and four deployment models."

The *essential characteristics* are outlined as
(1) **On-demand self-service**. "A consumer can unilaterally provision computing capabilities, such as server time and network storage, as needed automatically without requiring human interaction with each service provider."
(2) **Broad network access**. "Capabilities are available over the network and accessed through standard mechanisms that promote use by heterogeneous thin or thick client platforms (e.g., mobile phones, tablets, laptops, and workstations)."
(3) **Resource pooling**. "The provider's computing resources are pooled to serve multiple consumers using a multi-tenant model, with different physical and virtual resources dynamically assigned and reassigned according to consumer demand. There is a sense of location independence in that the customer generally has no control or knowledge over the

exact location of the provided resources but may be able to specify location at a higher level of abstraction (e.g., country, state, or datacenter). Examples of resources include storage, processing, memory, and network bandwidth."
(4) **Rapid elasticity**. "Capabilities can be elastically provisioned and released, in some cases automatically, to scale rapidly outward and inward commensurate with demand. To the consumer, the capabilities available for provisioning often appear to be unlimited and can be appropriated in any quantity at any time."
(5) **Measured service**. Cloud systems automatically control and optimize resource use by leveraging a metering capability at some level of abstraction appropriate to the type of service (e.g., storage, processing, bandwidth, and active user accounts). Resource usage can be monitored, controlled, and reported, providing transparency for both the provider and consumer of the utilized service.

The *service models* are outlined as (1) Software as a Service (SaaS), (2) Platform as a Service (PaaS), and (3) Infrastructure as a Service (IaaS).

The *deployment models* are outlined as (1) private cloud, (2) community cloud, (3) public cloud, and (3) hybrid cloud.

Special Publication (NIST SP) - 800-145

This project is within the context of the following law, instruction, and initiatives from Federal Government, DoD, and Air Force:
Federal Data Center Consolidation Initiative (FDCCI), 2010
National Defense Authorization Act (NDAA) FY2012, § 2867
OMB Federal Cloud Computing Strategy and Cloud-First Policy (Cloud-First), 2012
DoD Joint Information Environment Initiative (JIE), 2013
AFI 33-150, as amended December 2014
Air Force Data Center Consolidation Initiative (AFDCCI), 2015

BLUF: Current application rationalization process solely focuses on applications. Wholesale data center approach does not fit into the picture.

BACKGROUND: DoD / AF policy driving towards consolidation to CDC and / or an authorized cloud services provider (CSP). Landscape and path towards establishing hybrid cloud is unclear. Existing DISA CAP(s) not scalable and more than likely will increase costs to mission owners.

DISCUSSION: 75 ABW/SC is attempting develop a path forward to establish an

encrypted tunnel to a VPC via an authorized CSP for Level 5 applications and below can be hosted within the HEDC hybrid cloud. Intent is to establish a connection to the VPC equivalent to the level of current HEDC (e.g. Level 5) where existing MAC II & MAC III systems can be easily migrated to the commercial government private cloud. Many paths are potentially available to establish connectivity. Below are a few COAs to consider, diagrams and additional details can be located in Attachment 1.

- COA 1: Establish encrypted direct internet access link via commercial ISP to multiple authorized CSPs to reverse auction off data / applications to lowest cost providers
  -- PRO(S):
    --- Commercial ISP connections would provide low latency data center interconnects establishing viable COOP/DR functionality and paths to move applications to authorized locations
    --- CNDSP can be performed in the same manner as today by AFNET
  -- CON(S):
    --- Doesn't conform to existing main stream AF architectures
    --- Requires DoD waivers to establish commercial ISP
- COA 2: Leverage AF Gateways / JRSS stacks as approved connection points to Hybrid cloud routing around enclave boundary devices
  -- PRO(S):
    --- Encrypted Virtual Route Forwarding (VRF) to lower half of JRSS stack
    --- Data packet inspection by 26 NOS
  -- CON(S):
    --- Data migrations of hundreds of terabytes would take longer due to additional inspection points and packet sizing
    --- AF GWs / JRSS inspection adds latency for data packets to reach Hybrid cloud hindering user experience

COA 3: Leverage AF GWs / JRSS stacks routing traffic through local boundary devices into lower half of JRSS stacks
  -- PRO(S):
    --- Encrypted Virtual Route Forwarding (VRF) to lower half of JRSS stack
    --- Data packet inspection by 26 NOS
  -- CON(S):
    --- Data migrations of hundreds of terabytes would take longer due to additional inspection points and packet sizing
    --- Local boundary devices and AF GWs / JRSS inspection add latency for data packets to reach VPC hindering user experience

COA 4: Establish connectivity through DISA CAP
  -- PRO(S):
    --- Conforms to AF / DoD existing policies

--- Data packet inspection can be performed by AFNET and DISA

-- CON(S):

--- Data migrations of hundreds of terabytes would take longer due to additional inspection points and packet sizing

--- Local boundary devices, AF GWs / JRSS, and DISA CAP inspection points add latency for data packets to reach VPC hindering user experience

REQUESTED ACTION: get stakeholders of AF policy makers and CSP process owners together to formulate an acceptable way forward.

Many Issues In Current Environment:

- Application reliability problems
- Application performance problems
- Oracle consultants leaving after 6 years; lots of undocumented tribal knowledge
- Storage area network replacement requires reloading of operating systems
- Management complexity issues
- Server capacity utilization at maximum capacity
- Not enough software stacks acquired
- Server replacement required for old Oracle servers
- Inadequate documentation

The prior focus appeared to be on individual technologies and not the integrated solution it delivered or how it could be easily grown and sustained.

To correct this, it would be safer and faster to deploy a new environment utilizing commercial best practices leveraged in typical ERP implementations or major ERP upgrades utilizing much cheaper IT infrastructure technologies. Regression testing would be utilized to insure new system design was adequate based on captured baselines of the legacy environment. Regression testing would be a critical success factor.

Regression Testing: Identify, Analyze, Resolve, Document Then Repeat

1. Determine if the system design will meet the needs of the business *before* going live
2. Detect application bottlenecks resulting in slow time or down time in production
3. Know if system can scale to the desired level of usage in production, test and development
4. Disaster recovery is completely successful and repeatable meeting MTTR requirements
5. Quantifiable data on how the design performed and is improved

A Test Is Worth Many Expert Opinions…

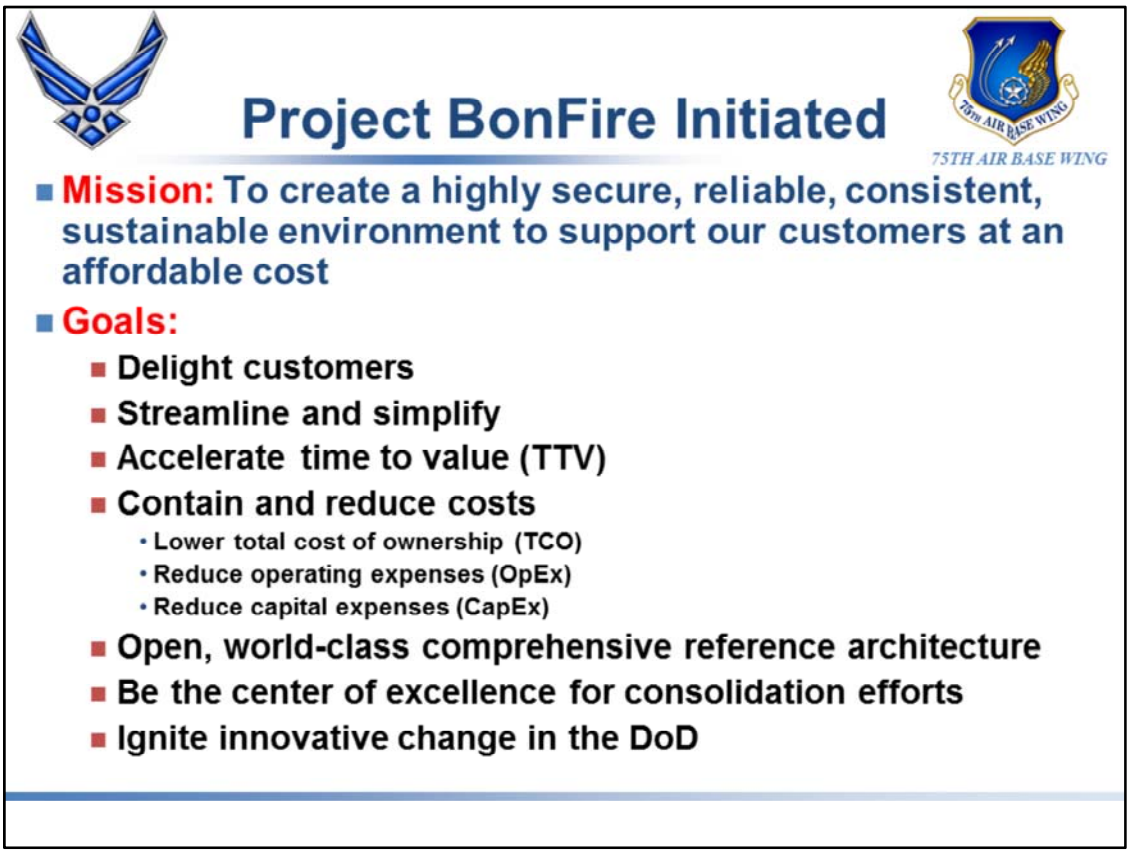

HEDC Mission Statement: We delight our customer by creating a modern, reliable, consistent, and elastic infrastructure at an affordable cost and by providing highly skilled professional services and administrative talent to help customers achieve rapid time-to-value.

HEDC is always evolving this infrastructure to manage ever more complex and expanding workloads creating greater value for customers by ensuring that their systems and workloads are never left behind or trapped on legacy and non-evolving infrastructures.

Streamline and Simplify: Reduce the number of different platforms and architectures that must be supported. Ensure that the technology stacks deployed work seamlessly, as if it were from a single vendor. Consolidate vendor relationships where possible and establish deep partnerships with a core set of technologies providers who work together now and in the future.

Contain and Reduce Costs: Embrace open platforms that cost less and scale more effectively. Embrace software-defined infrastructure. Automate the most labor intensive IT processes, like software maintenance and management. Ensure that the architecture works together – hardware, operating systems, services and key applications. Embrace vendors that promote open standards and interoperability. Ensure licensing terms cover the widest spectrum of customers.

Accelerate Time-to-Value: Pre-tested, validated best practices simplify and reduce deployment time for customers. They also ensure resources are put to work and not on the shelf. An example of this is dynamically provisioning a new or existing application from bare metal to a working environment utilizing a proven gold standard. This process, called “Bare Metal Provisioning”, ensures a consistent, reliable and scalable end result.

Secure Hyper-Convergence: Advance the ability of software-defined infrastructures and commodity-based storage hardware to host larger and more complex workloads in a meaningful way. In particular, increase the ability to host virtualized instances of Tier-1 and Enterprise application workloads providing the cost-savings of commodity hardware and operational benefits of abstraction to information systems that currently require dedicated hardware. As well, provide an infrastructure to support a greater services density achievable by fostering micro-services and a capability for supporting cloud-native application/services. The capabilities provided by hyper-convergence will be a key differentiator for HEDC

Implement a hyper-converged solution that will support encryption, dedupe, and CSP destaging. HEDC has already achieved a converged infrastructure with its large commitment to workload virtualization. The next step in this evolution is to achieve secure hyper-convergence. This will involve expanding the capabilities of software-defined networking to include greater security, performance, and better operational control over east-west traffic. Secure hyper-convergence will also involve expanding the capabilities of software-defined storage to include replication, encryption, dedupe, and destaging from higher performing to lower performing storage including the use of RAM as a high performance storage tier. VMware NSX will be the primary component of the software-defined network in this evolution, and Atlantis USX will be the primary component of software-only, hyper-converged storage. Atlantis USX hits all the objectives - Using RAM as a high-performance storage resource, encryption, dedupe, destage, and replication. HEDC will achieve a “Center of Excellence” arrangement with the hyper-convergence solution vendor Atlantis.

Secure Micro-Segmentation: Advance the ability to provide true isolation for multi-tenant environments. Currently most cloud or virtualization infrastructure still focus primarily on perimeter-centric security, which may isolate a group of tenants from threats originating from North-South traffic, but do not adequately isolate tenants from each other inside the perimeter (East-West traffic). If the perimeter security is compromised, unauthorized activities can ride along laterally with authorized activities. Unauthorized lateral traffic creates a major trust crisis in current converged infrastructure, and while it is possible to use traditional perimeter security methods to manage lateral traffic, it defeats the costs savings and economies of scale that software-defined networks provide. Secure Micro-Segmentation will be a key differentiator for HEDC allowing it to provide a "zero-trust" infrastructure between tenants securely managing all lateral traffic while still maximizing economies of scale in its hardware and software investments.

Implement VMware’s NSX with F5 VE’s. A software-defined network solution must provide isolation and feature-rich segmentation capabilities. By default, VMware NSX provides isolation of one virtual network from another virtual network without the need to configure firewalls, ACLs or VLANs to achieve isolation. This is true, even if a virtual network spans hypervisors. All traffic between hypervisors is encapsulated and virtual networks are isolated from the underlying physical infrastructure by default. NSX also manages segmentation programmatically ensuring that network services are created and distributed to the

hypervisor's virtual switch enforcing network services (L2, L3, ACLs, Firewalls, QoS) at the virtual interface. This eliminates the need to configure segmentation on physical networks and firewalls outside the virtual network and leaving enforcement open to provisioning errors. Furthermore, services can be inserted (service chaining) at the logical pipeline that flows through the virtual switch rather than forcing that traffic off the virtual network to the physical network to consume services provided by other vendors allowing tenants to take advantage of firewalls, application delivery controllers (ADCs), vulnerability scanners, etc. without exiting the virtual network and breaking the "zero trust" model. F5's Big-IP acquisition in FY15 is an important example of a product that will fully leverage service chaining and be inserted into the logical pipeline making an ADC service part of the East-West traffic rather than forcing tenants to make a lateral hop for this service.

Secure Micro-Services: Embrace secure micro-services as a means to create an HEDC common set of services expanding HEDC's software-as-a-service (SaaS) offerings and ensuring that HEDC has a solid ability to support cloud native applications and support appropriate candidates for refactoring into a micro-services model. Also, provide AF developers a modern set of tools and methods to create next-generation applications, allowing HEDC to create an attractive development-as-a-service (DEVaaS) infrastructure. For those services that are somewhat generic and often requested, package them as a micro-service and make them easily deployable, by themselves or in concert with other packaged services within an HEDC Service Catalog. For example, Apache HTTPD, MySQL, Tomcat-based applications are all good cloud-immigrants and candidates for micro-service deployment. Deploying common services this way will allow the HEDC to develop a catalog driven common computing environment (CCE) that can be updated and improved quickly. Furthermore, creating a robust development infrastructure that securely and responsibly brings developers back into the data center and allows them to continuously develop, test and integrate in a trustworthy manner will help to modernize AF development efforts and allow HEDC to better serve a customer segment that has not received a lot of attention but plays a crucial role in time-to-value. Micro-services with continuous integration provide an enhanced pipeline-to-production model that has been growing in relevance since 2013 and will soon become the standard alternative to both monolithic and typical SOA development models. Deploy software containers as the primary method to provide secure micro-services. While there are various ways to provide micro-services, methods based-on software containers reduce the complexity required by elaborate service frameworks and allow for a greater density of services at scale. However, the lightweight nature of software containers did not come without caveats. Early container solutions lacked an adequate framework to manage port abstraction, and prior to 2015, container solutions (for instance Dockers containers) lacked a comprehensive authentication framework to govern the interactions of containers. Furthermore, HEDC was reluctant to adopt a containers platform that did not leverage current investments. HEDC, along with other VMware customers, has put pressure on VMware to develop a micro-services strategy that leverages ESXi, vSphere, and NSX, and VMware announced in mid-April their container strategy. This strategy uses the Photon micro-OS on ESXi, leverages NSX for isolation and segmentation, and will also provide an authentication framework called Lightwave that can control access and authentication

between containers. VMware calls this their Cloud-Native Applications (CNA) initiative, and HEDC was involved with VMware in the early stages of this initiative. In August of 2015, VMware announced further integration of CNA with vSphere. The hyper-converged infrastructure envisioned in this 5-year proposal combined with micro-services will provide an especially dense services-at-scale solution that will drive data center energy and footprint efficiencies.

Leaf-Spine Network Architecture: Plan to migrate to a network architecture designed to scale HEDC's virtualization infrastructure. Traditional data center traffic, based predominantly on physical server instances, tends to produce more North-South traffic. As data centers achieve convergence and ultimately hyper-convergence, East-West network traffic begins to dominate as VM-to-VM traffic increases. This traffic pattern will increase even more with the addition of micro-services.

Shifting from a traditional three-tier or Core-Aggregate-Access model to a Leaf-Spine (Distributed Core) model will better support the traffic shift from a predominance of North-South traffic to the East-West traffic that accompanies growing software-defined abstraction of an organization's infrastructure. Typically, within the traditional three-tier model, VM-to-VM traffic is hindered by the spanning-tree protocol (STP) and by hair-pinning affecting the ability to scale network traffic within the virtualized network. The Leaf-Spine model accompanied by capabilities of VMware NSX will provide a scalable solution at a reasonable price as HEDC further embraces a secure, software-defined network infrastructure.

HEDC Service Catalog & Services Automation: Embrace a service catalog model that provides an automation framework for standing up services thereby standardizing and simplifying service deployments. This ability will allow HEDC to shift deployment tasks to lower-cost IT OPS personnel and make some deployment tasks available to customers directly via a self-help portal. While HEDC provides an engineered PaaS solution, it will increasingly be forced to deploy an IaaS to support IPN hosting requests that do not fall into the cost model of the PaaS. To embody the essential characteristics of the cloud model, the HEDC must develop and advance a service catalog supported with automation and self-help deployment capabilities to control IT operational costs and be prepared to compete in those instances where customers may demand self-help interaction with the infrastructure. Finally, the service modeling done as part of creating the service catalog and automation will provide the foundation for IT service management capabilities including showback/chargeback as well as the foundation of workload management for a hybrid-cloud infrastructure.

The best approach to building an HEDC service catalog is to build the catalog as an extension of automating IT services. The process of modeling a service and automating the deployment of that service creates a de-facto service catalog. Currently, HEDC has BMC's Cloud Lifecycle Management (CLM) to build the foundations of automation and a service catalog. However, no single automation tool may meet HEDC's needs. It will likely be necessary to incorporate both BMC's BladeLogic Server and PuppetLab's Puppet Enterprise

to enhance automation capabilities. At this point CLM will become the catalog framework, and then other automation services will be integrated into that framework as necessary.

IT Business Management Integration: Advance HEDC's ability to manage IT costs and customer commitments in order to align HEDC's infrastructure to customer business requirements and provide a "modern, reliable, consistent, and elastic infrastructure" at a reasonable cost.

Provide HEDC management with timely capacity utilization, growth prediction and service accounting/SLA information. Deploy ITBM framework that provides showback/chargeback billing as well as tracks IT costs (actual and predicted) for the organization. As well, provide a quality of service management component to track SLA commitment compliance and provide alerts when compliance is threatened. Where appropriate, present as much information as possible in a dashboard format to give HEDC management a simple and immediate way to consume ITBM information. Both VMware's vRealize Business (formerly IT Business Management Suite) and BMC's TrueSight Capacity Optimization product are candidates for the ITBM framework. vRealize Business has a built-in service quality management component where BMC's Remedy Service Level Management (SLM) can be used in conjunction with Capacity Optimization to complete the objective. The service modeling efforts performed as part of the HEDC Service Catalog and Services Automation effort provide the elements for creating chargeback/showback for services. As such, BMC's offerings may be best candidates for this objective.

Hybrid-Cloud Deployment: Plan for the ability to securely and quickly move workloads from HEDC's internal infrastructure to diverse cloud infrastructures. According to a recent study by Intuit and Emergent Research, 78% of small business in the US will be fully on the "cloud" by 2020. Federal and DoD policy indicates a push towards adoption of the cloud deployment model as well. Using a hybrid-cloud strategy will keep critical data and workloads in house but allow the movement of appropriate workloads to authorized service cloud providers, reducing costs, providing an elastic IaaS service offering, and giving the HEC the ability to perform application rationalization towards the a cloud deployment model as appropriate. The HEDC will extend their data center to the cloud with software and services that allow them to securely control the external compute resources and move workloads to and from the resources in a timely manner. Greater adoption of micro-services will likely accelerate the need for private cloud-to-public cloud federation. Authorized public CSP and DoD cloud infrastructure will become an elastic extension of HEDC's internal infrastructure and will provide better time-to-value as it will be possible to stand up appropriate services in minutes rather than days or months where additional internal infrastructure must be purchased and provisioned to support the service. It will also provide additional cost management options as HEDC will only pay for public cloud infrastructure that it uses. Services running in the CSP infrastructure can be deleted or moved back down to HEDC's internal infrastructure as it becomes available allowing HEDC to reduce costs as it shrinks after a prior expansion of services.

There are uncertainties about which CSPs will best conform to DoD standards in addition to FedRAMP. HEDC will begin the process of getting network access through DISA's Cloud Access Point (CAP) for at least three public authorized CSPs registered in the FedRAMP catalog, namely Amazon's AWS, Microsoft' Azure, and Quality Technology Services (QTS). AWS is the hands down leader in the market. Azure presents interesting opportunities for deploying Microsoft products and is building infrastructure within the DoDIN. QTS also has infrastructure within the DoDIN. HEDC make appropriate changes to HEDC's C&A or develop an entirely separate C&A to support the hybrid-cloud model. HEDC will begin with leveraging appropriate cloud service offerings (CSOs) such as data archiving on inexpensive cloud object-storage and will evolve an IaaS offering for those IPN customers that do not require an engineered PaaS solution. With cloud connectivity achieved, the HEDC will deploy a hybrid, scale out storage solution that uses local storage to cache data at the HEDC, but can securely and inexpensively archive that data into a CSP's infrastructure. This is a software defined solution that will further leverage the use of commodity-based storage hardware and will leverage the cloud to consolidate HEDC storage. Ideally, a version of this solution can also be applied to the next-generation WAN POP appliance, allowing the cloud to become an alternative site to archive data rather than bring it back to the HEDC.

Secure Micro-Service will provide a key mechanism to support cloud native applications within the hybrid-cloud deployment. For typical operating environment and cloud-immigrant support within the IaaS, VMWare's OpenStack Integration will provide a key support mechanism. The recent acquisition of VMware vSphere Enterprise Plus is a critical achievement to this end.

"Small Business Success in the Cloud," www.emergentresearch.com/future-of-smb-for-slidewharev2.html 07, August 2014

Geo-Distant COOP and Local "Sister" Data Center Regions and Azs: Create a geographically distant COOP to support HEDC operation continuity. This site must have adequate bandwidth and infrastructure to support disaster recovery of key HEDC services and for HEDC customers. Beyond better mission assurance and better service level agreement capabilities, an adequate geo-distant COOP will enhance HEDC's position that it actually functions as more than an IPN as defined by JIE. Rather it functions as a "super IPN."
Modern data centers use a sister data center model to support the realities of CAP theorem (consistency, availability and tolerance partition-ability) within the context of an application distributed across data centers. As the theorem goes, it is only possible to have two of the three desired attributes for an application at any one time in a transactional sense. So a sister data center within the distance limits of synchronous replication provides the best possible alternative to the hard limits imposed by the reality described by CAP theorem. Consistency at a geographically distant data center using asynchronous replication is generally also added to this scheme providing availability, partition-ability, and a "will be consistent." HEDC needs this local data center as a place to provide local failover capabilities with superior recovery point objectives and recovery time objectives, and to move workloads to accommodate maintenance in one of the two data centers.

HEDC will create a reasonable COOP capability at Nellis Air Force Base. This solution will primarily utilize the capabilities of EMC and HP storage arrays to provide replication services. However, hyper-converged infrastructure with native replication capabilities will compliment and perhaps replace these SAN arrays at some point.
HEDC will create a layer-3 connection between the data center in bldg. 1211 and the HEDC data center space in bldg. 891. Networking acquisitions for FY15 will make this possible. Furthermore, VMware NSX will enhance this solution allowing the HEDC to stretch subnets between the buildings. This will provide many options to move operating environments between the data centers and apply the same security model throughout.

Next Generation WAN POP Appliance: Enhance HEDC's "fog computing" strategy that establishes a small but powerful near-user compute and storage platform that is fully integrated with the HEDC infrastructure and that can be used in multiple scenarios including the next-generation WAN POP Appliance. HEDC was an early "fog computing" pioneer within the AF. Unlike cloud computing, fog computing is premised on putting small but necessary amounts of compute and storage infrastructure near the user. This infrastructure ties back into a primary data center(s). Many factors including performance and connection issues make fog computing a necessary strategy. HEDC will use hyper-convergence to radically shrink the next-generation of WAN POP appliance.

HEDC will use the same hyper-convergence building blocks being used in the data center to build the next-generation WAN POP appliance. Atlantis USX will feature prominently in this appliance and will allow it to shrink from one 42U rack to less than 10U. A version of the hybrid, scale-out storage solution will also be applied to the next-generation WAN POP appliance, allowing the cloud to become an alternative site to archive data rather than bringing it back to the HEDC. This "Fog-to-Cloud" capability will provide valuable options on how to best consume AFNet WAN bandwidth and HEDC SAN resources.

Next Generation of Performance for Analytics and BI for Tier1, Enterprise Applications: Ensure that HEDC has the best AF story for "jaw-dropping" performance of Tier1, enterprise application workloads, especially for the BI and Analytics requirements associated to that class of application. HEDC was an early adopter of flash-technology and used this technology to radically improve the performance of enterprise workloads. HEDC will provision a new-generation of BI/analytics grid that will take advantage of rack scale flash technology.

EMC recently acquired DSSD the makers of a unique all-flash, PCIe interconnected storage platform that brings the benefits of very fast flash to shared storage. Up to 96 servers can connect to and share this storage at PCIe gen3 interconnection speeds. Whereas the performance benefits of PCIe connected flash where isolated to a single server in the past and the performance benefits of shared storage on an all flash storage array were diminished by the slower connectivity of the array's bus, the DSSD promises the best of both worlds. HEDC has worked with DSSD/EMC and has secured one of the first DSSD's to ship in late 2015.

Data-As-A-Service (DaaS): Create a DaaS solution that will allow for quick and efficient data provisioning in environments with ongoing test and test configuration requirements and that will include an automated method to anonymize and mask data that will be used in minor environments. ERP systems require frequent environment refreshes for testing and upgrades. Shortening the window to stand up environments with the correct version of the data accelerates time-to-value on important projects. Also, a DaaS that can share data between platforms will decrease the time it takes to refactor or "immigrate" applications to alternative platforms such as the cloud. For, instance, a DaaS that can share data between an operating environment running on dedicated hardware and an operating environment running on as a VM or on a cloud machine image can significantly decrease the amount of time it takes to perform the migration. This is a key capability to deal with the complexities derived from the need to support a spectrum of deployment models across multiple infrastructures.

HEDC will develop a DaaS using Delphix. The Delphix solution manages virtual copies of data and can make these copies available to an operating environment as either a current copy or the current copy minus deltas so as to present a snapshot of the past. Virtual copies can be delivered in minutes rather than hours and can be pre-anonymized further saving an administrator from the effort of anonymizing before the data is presented to an operating environment.

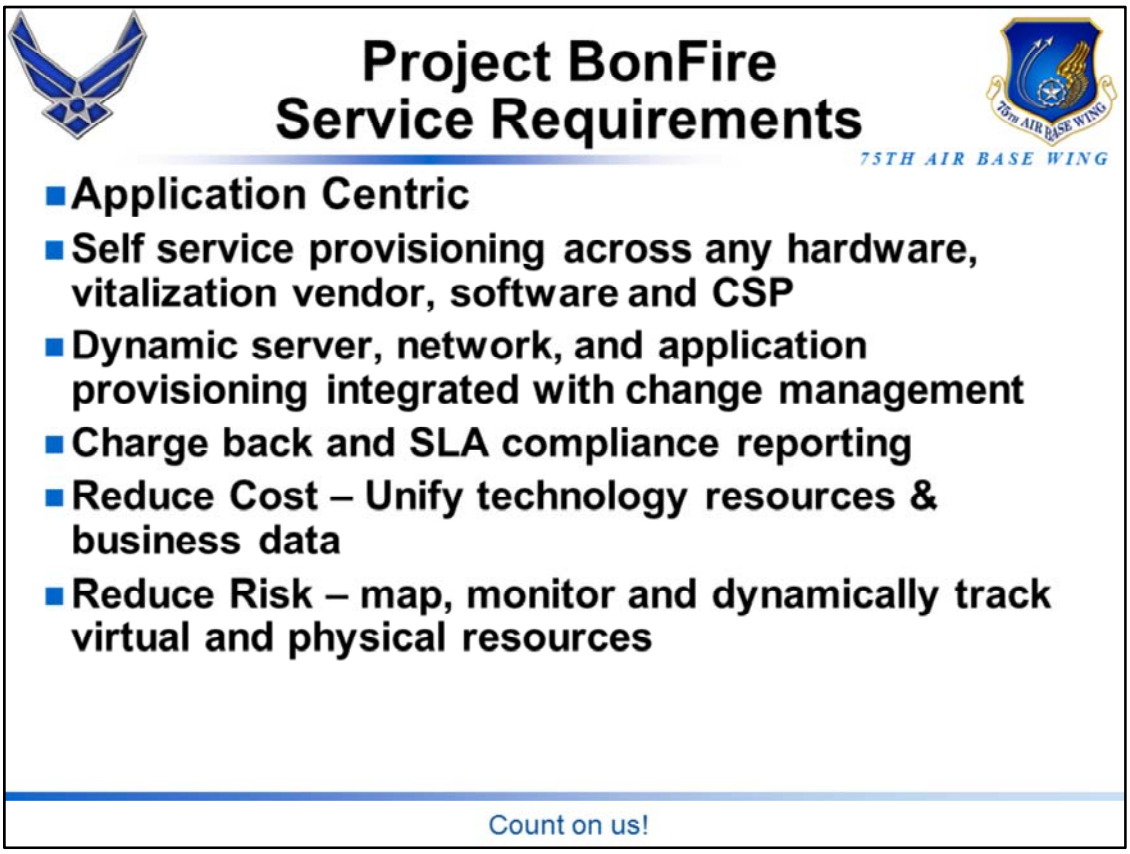

Streamline and simplify
Reduce the number of different platforms and architectures that must be supported. Ensure that the technology stacks deployed work seamlessly, as if it were from a single vendor. Consolidate vendor relationships where possible and establish deep partnerships with a core set of technology providers who work together now and in the future.

Contain and reduce costs
Embrace open platforms that cost less and scale more effectively. Automate the most labor intensive IT processes, like software maintenance and management. Ensure that the architecture works together - hardware, operating systems, services and key applications. Embrace vendors that promote open standards and interoperability. Ensure licensing terms cover the widest spectrum of customers.

Accelerate time to value
Pre-tested, validated best practices simplify and reduce deployment time for customers. They also ensure resources are put to work and not on the shelf. An example of this is dynamically provisioning a new or existing server from bare metal to a working environment utilizing a proven gold standard. This process, called “Bare Metal Provisioning”, ensures a consistent, reliable and scalable end result.

Often times technology surpasses slow TTV before a solution is even implemented!

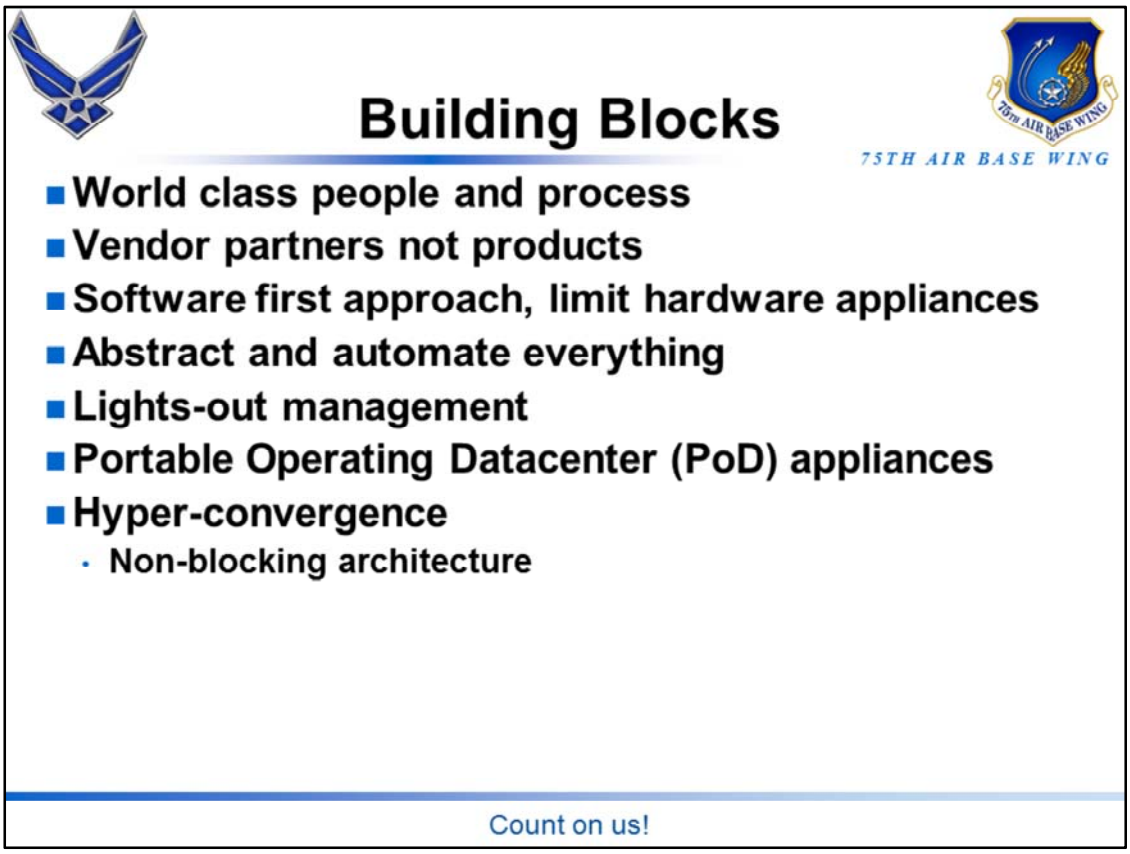

How was it done? – Process

More complex than technical issues

The BonFire project was made up of members from the Communications Squadron, IT, DISA, CMX, several contracting entities and airmen. Basically a regular project by today's standards.

The most important success factor is to have an executive sponsor for the project.

The contracting world will have disparate cultures or no culture at all as they come together to do a project. Developing a project culture is another tool for continuous improvement. One of the greatest challenges is to facilitate giving up separate agendas for one common goal. The doctrine of inclusion is the process of creating stakeholders of all the team members and their organizations. Those that would usually detract were given assignments and 'yoked up' to the project. Their success then becomes tied to the success of the project. Once the barriers are removed the technical skills and the personalities of the participants make the project an enjoyable endeavor.

The project was broken into several windows of effort. At each window the tasks were well defined with responsible parties and due dates. The daily focus meetings were attended by all team members to keep the effort on track and bring up potential issues including distractions from other efforts. After each window a couple of weeks were used to take care of other responsibilities that were of a lower priority and then the next window was started.

Celebrating the finish is as important as any other part of the project. There will always be another task ahead, but the gratitude shown for a job well done will be a lasting memory.

Caution: Icebergs Ahead

75TH AIR BASE WING

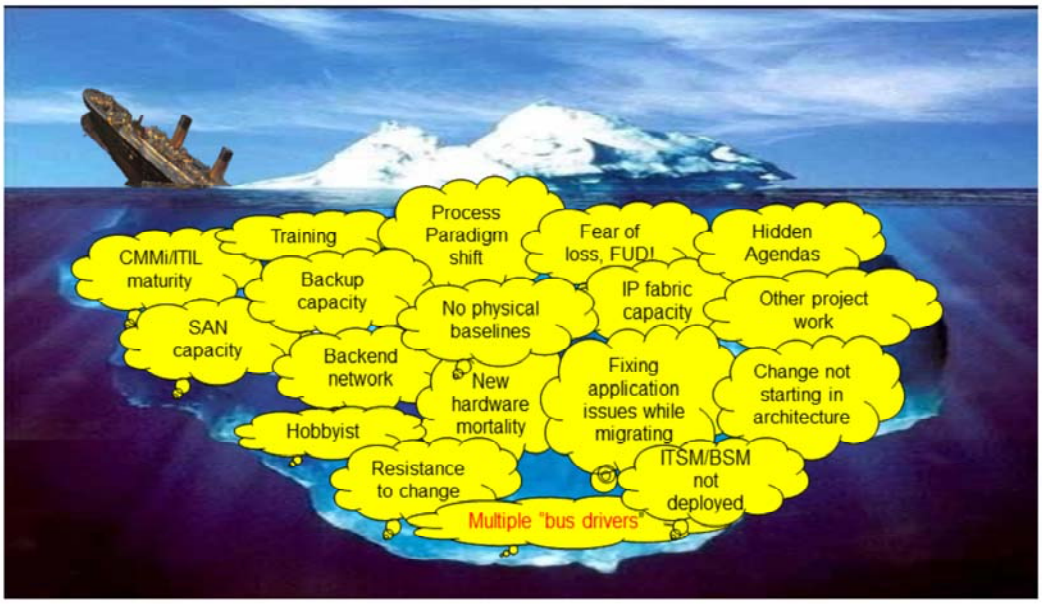

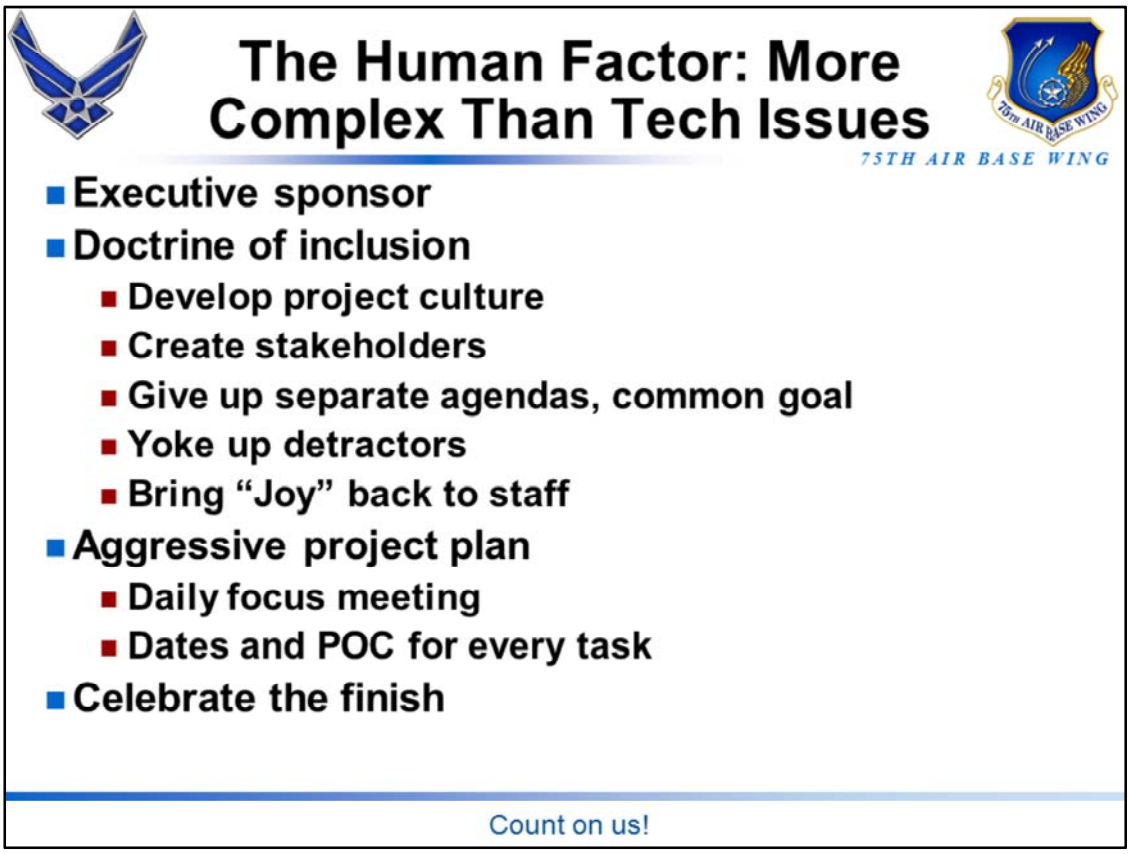

How was it done? – Process

More complex than technical issues

The BonFire project was made up of members from the Communications Squadron, IT, DISA, CMX, several contracting entities and airmen. Basically a regular project by today's standards.

The most important success factor is to have an executive sponsor for the project.

The contracting world will have disparate cultures or no culture at all as they come together to do a project. Developing a project culture is another tool for continuous improvement. One of the greatest challenges is to facilitate giving up separate agendas for one common goal. The doctrine of inclusion is the process of creating stakeholders of all the team members and their organizations. Those that would usually detract were given assignments and 'yoked up' to the project. Their success then becomes tied to the success of the project. Once the barriers are removed the technical skills and the personalities of the participants make the project an enjoyable endeavor.

The project was broken into several windows of effort. At each window the tasks were well defined with responsible parties and due dates. The daily focus meetings were attended by all team members to keep the effort on track and bring up potential issues including distractions from other efforts. After each window a couple of weeks were used to take care of other responsibilities that were of a lower priority and then the next window was started.

Celebrating the finish is as important as any other part of the project. There will always be another task ahead, but the gratitude shown for a job well done will be a lasting memory.

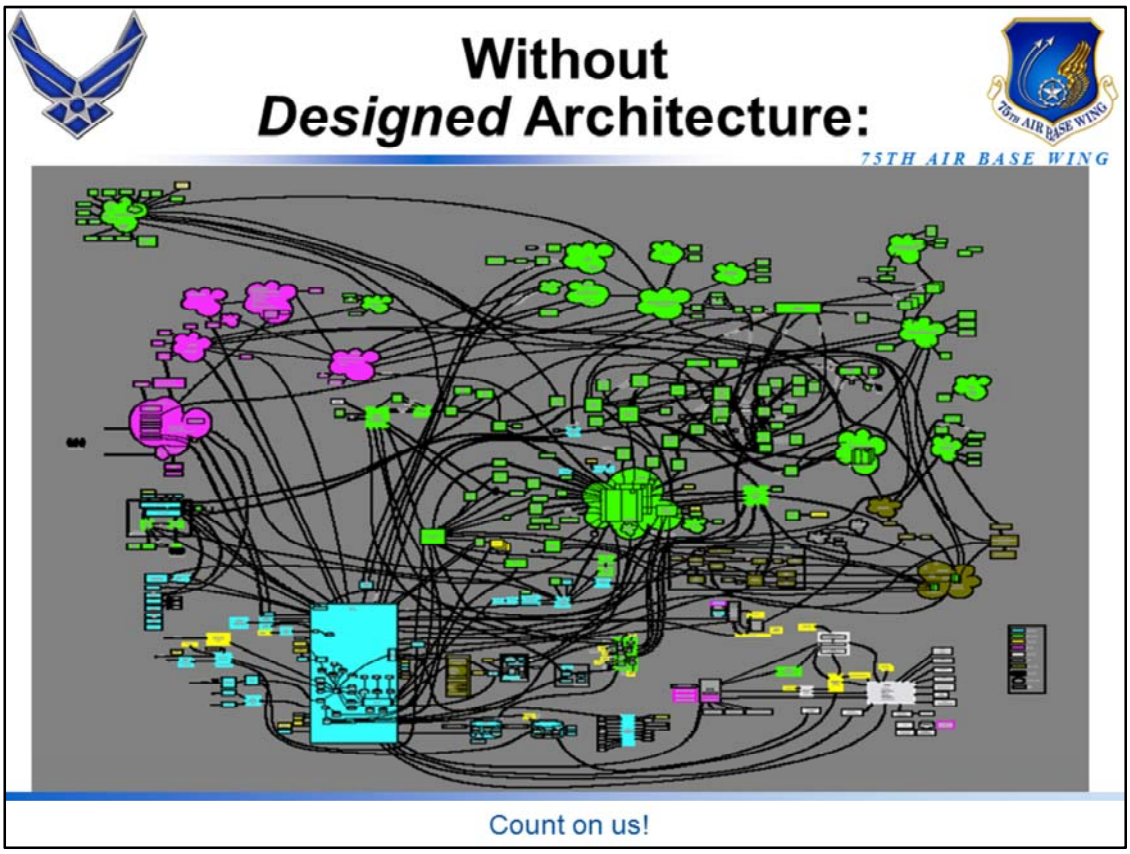

How was it done? – Design Thinking

Design thinkers speak the language of *viability* . "What is the real problem we are trying to solve?" They look to explore new opportunities independent of specific technologies.

More than just project doctrine, BonFire methodology includes design thinking, not critical thinking. Herbert Simon has defined "design" as the *"transformation of existing conditions into preferred ones"*. Design thinking is, then, always linked to an improved future. Unlike critical thinking, which is a process of analysis and is associated with the "breaking down" of ideas, design thinking is *a creative process based around the "building up" of ideas*. Design thinking process: Define, Research, Ideate, Prototype, Choose, Implement, Learn.

More technology than solutions: This is an architectural map of another actual IT department. Imagine trying to get any quality information out of this system! They have a very complex IT environment with many "best of breed" components that they integrated together themselves. They have a high proportion of customized software and interfaces. They have multiple, duplicate applications and versions of data. This creates the need for a very high amount of data movement between different systems. The environment is inflexible and difficult to change. They have not created their IT environment in accordance with a comprehensive architectural framework. This has compounded complexity and increased time to deliver new functionality. The environment is costly to maintain and operate and very costly to enhance or to integrate new functionality.

Key new architecture goals: Streamline and simplify. Reduce the number of different platforms and architectures that must be supported. Ensure that the technology stacks deployed work seamlessly, as if it were from a single vendor. Consolidate vendor relationships where possible and establish deep partnerships with a core set of technology providers who work together today and in the future. Manage infrastructure smarter and more effectively.

Contain and reduce costs: Embrace open platforms that cost less and scale more effectively without sacrificing security or stability. Automate the most labor intensive IT processes, like software maintenance and management. Ensure that the architecture works together - hardware, operating systems, services and key applications. Embrace vendors that promote open standards and interoperability.

Major IT trends used For project design:

- Virtualization - Abstraction of server, operating system, application and storage devices
- Grid storage and computing - Gaining IT power when you want it, where you want it
- Service-Oriented Architecture (SOA) - Simplify how information technology services are deployed, accessed and managed
- Information Lifecycle Management (ILM) - A policy-based approach to managing the flow of an information system's data throughout its life cycle: from creation and initial storage to the time when it becomes obsolete and is deleted.
- Maximum available architecture - Unified high availability (HA) solution for business services and the associated data storage

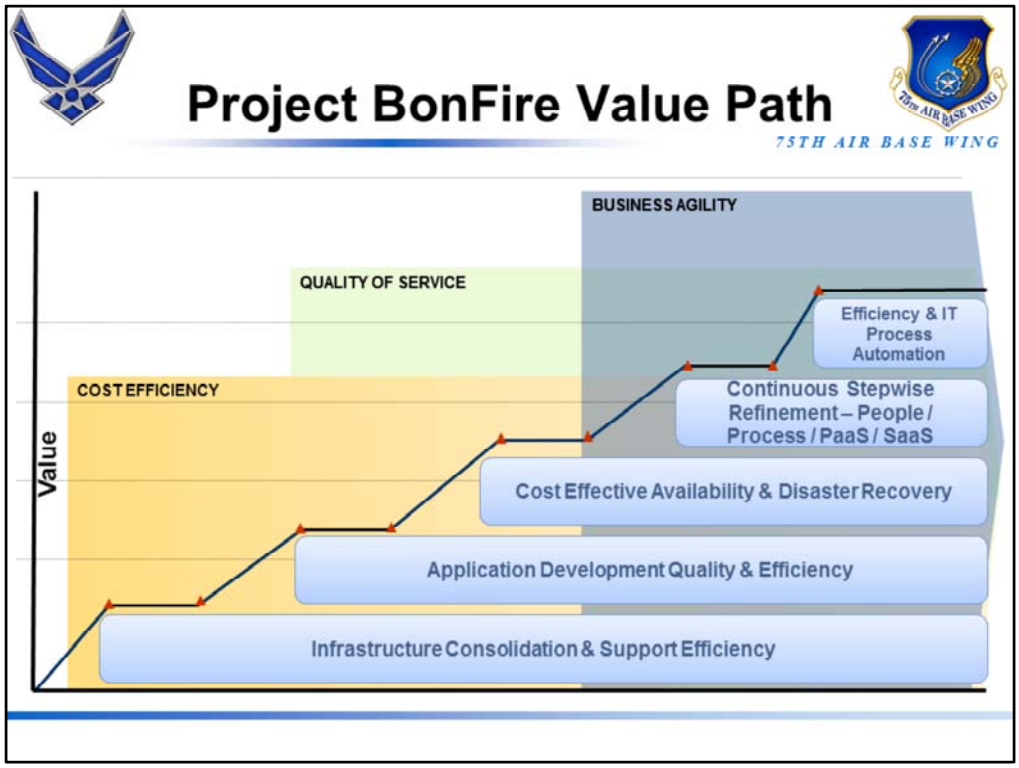

Project Value Path building from the ground up:

1) Phase 1 - Cost Efficiency
2) Phase 2 - Quality of Service
    1) Separation of application tiers, federate applications
3) Phase 3-5 - Business Agility

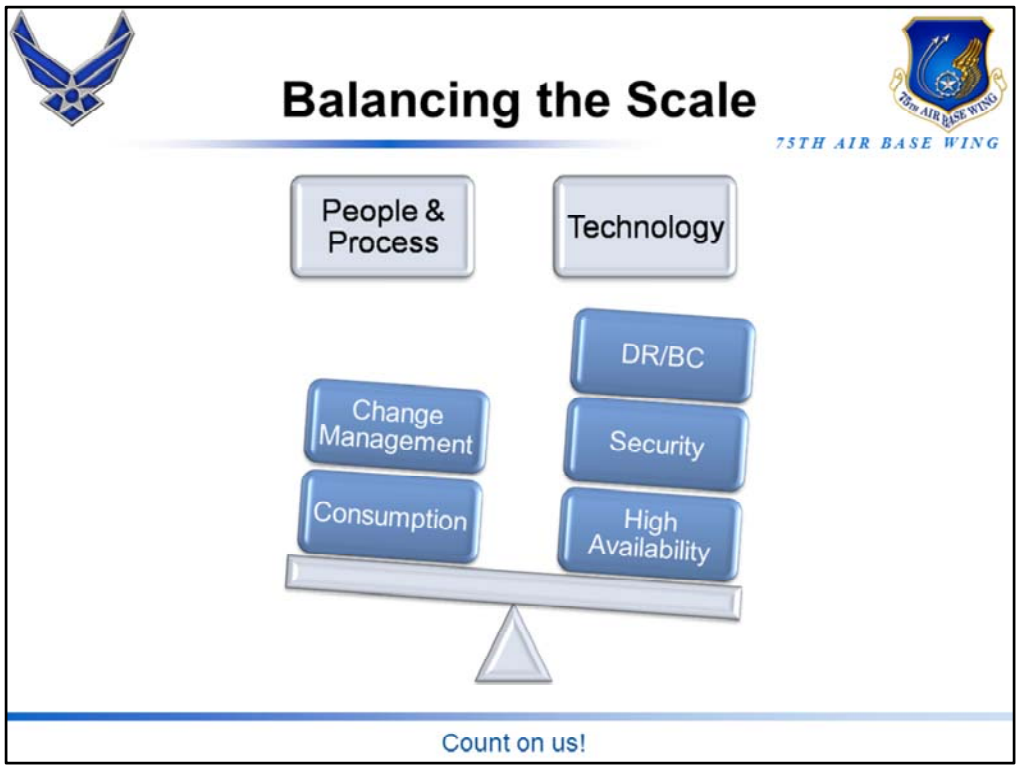

Successfully balancing the consolidation of applications while scaling service delivery in quality and capacity and reducing risk, costs and Time to Value (TTV) will not come from "Silver Bullet" technology or vendors:

Today's data center is increasingly more complex; there are too many disparate tools, too many manual processes, and adopting new technologies such as virtualization introduces its own set of challenges. The typical result is reduced visibility, control and performance of application service delivery and increased challenges in scaling operations without increasing risk of unplanned downtime.

Virtualized consolidation data centers and Clouds can and do go down, as evidenced by several major outages of public clouds operated by global, very tech-savvy corporations. The primary reason these crashes occur is when building these new environments, organizations are focused on "***Silver Bullet***" technologies like virtualization that enable data center consolidation and do not adequately plan for failure and mitigation of increased risk domains. The various resources needed to perform application hosting consolidation and to create clouds are adding more complexity to each application system, which puts mission-critical applications and data at much higher risk.

The Hill Enterprise Data Center (**HEDC**) has been actively consolidating applications since 2006 and is considered Operationally Mature (**OM**) and a Center of Excellence (**COE**) by our major vendors. Our focus has been to simplify, standardize, consolidate, automate, steadily refine then repeat our Enterprise Architecture (**EA**). We have iterated through five major EA versions during Project BonFire at the HEDC.

HEDC's EA supports our business service delivery with world-class heterogeneous geographically dispersed Clouds that are managed with the same set of tools across both our virtualized and physicalized environments. Our EA tools are server, storage, fabric, operating system, application system and virtualization ecosystem agnostic. Abstraction of these components has allowed us to obtain software and hardware components at some of the highest discount levels vendors have delivered to the DoD. The abstraction of these hardware and software components has further returned value by transforming the complexity to simplicity of migrations and expansions to refreshed and new technology in the HEDC cloud appliances.
Where possible, we have sought and obtained capability based Unlimited License Agreements (**ULA**) with fixed five years Total Cost of Ownership (**TCO**). Basing the ULA on capability and not specific products allows HEDC to not be locked out of new features which may require further unplanned CAPEX and OPEX. Security in Depth, Continuous Compliance, High Availability, Wide Area Backup/DR, Virtual Storage Management, Virtual Server Management, Archiving/eDiscovery, Major Databases, Major Application Servers, Business Intelligence and Business Service Management have all been considered and addressed.

Our EA has allowed our agility to increase while decreasing our costs and risks of data loss and unplanned service outage. We have not increased our staff since the Project BonFire kick off five years ago but have consolidated over 1000 servers and 1PB of data. Our next version 6 of HEDC will start being deployed in October 2012. With version 6, we will maintain and enhance end-to-end visibility and control of an increasingly heterogeneous environment, as well as optimize data center assets without affecting the availability of mission critical applications while reducing our TCO and TTV.

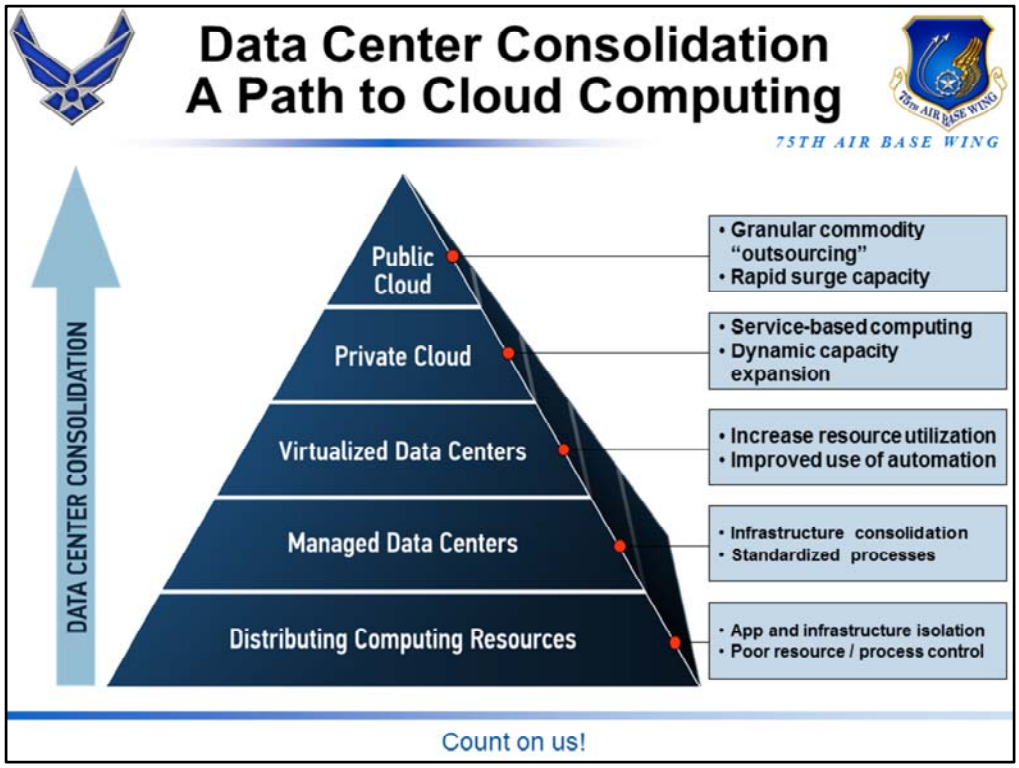
Data Center Consolidation
A Path to Cloud Computing
75TH AIR BASE WING
DATA CENTER CONSOLIDATION
Public Cloud
Private Cloud
Virtualized Data Centers
Managed Data Centers
Distributing Computing Resources
• Granular commodity "outsourcing"
• Rapid surge capacity
• Service-based computing
• Dynamic capacity expansion
• Increase resource utilization
• Improved use of automation
• Infrastructure consolidation
• Standardized processes
• App and infrastructure isolation
• Poor resource / process control
Count on us!

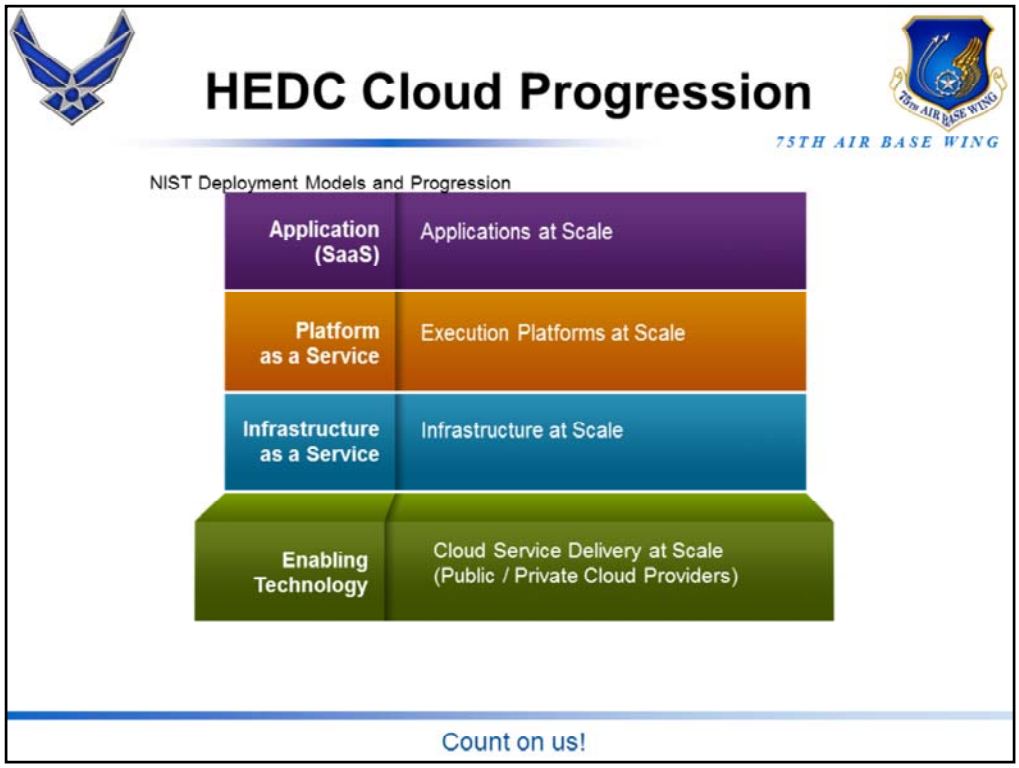

- There are 4 major categories in the Cloud Computing value chain. These are the target workloads and user base for each category
- The first category is Software as a Service: This is Applications services delivered over the network on a subscription basis.
- Then there is Platform as a Service which is Software development frameworks and components delivered over the network on a pay-as-you-go basis.
- The next category is Infrastructure as a Service where compute, network and storage delivered over the network on a pay-as-you-go basis. Amazon pioneered this with AWS (Amazon Web Service)
- And of course, there is an IT foundation that has to keep all this going—Cisco intends to be the leading provider of enabling technology to both the service provider and enterprise markets

HEDC Software Defined Progression
Lots Of *aaSes...

75TH AIR BASE WING

- Security as a Service
- Storage as a Service
- Database as a Service
- Environment as a Service
- Process as a Service
- Integration as a Service
- Management as a Service
- Testing as a Service
- Information as a Service

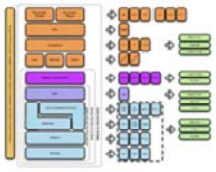

Count on us!

HEDC Focus Areas

75TH AIR BASE WING

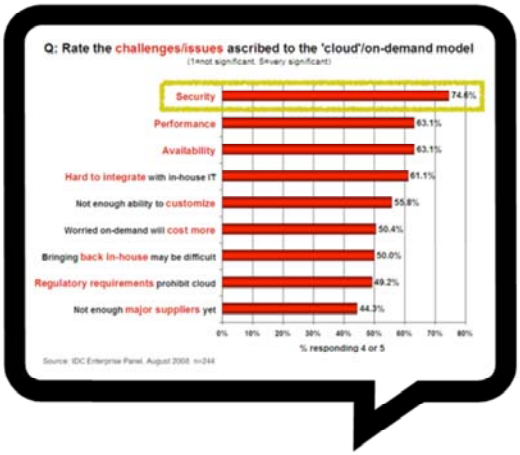

Count on us!

HEDC Current State

75TH AIR BASE WING

- Total # Applications – 450+
- Total # VMs – 1630+
- Total # Physical Servers – 375+
- Total Storage / Backup capacity – 96+ PB
- Green Revolution Cooling POD – SBIR project
  - An additional 256 hyper-converged nodes
  - Conservatively an additional 2500+ VMs
  - Deduplicated RAM capacity – 1+ PB
  - Deduplicated Storage capacity – 14+ PB

Count on us!

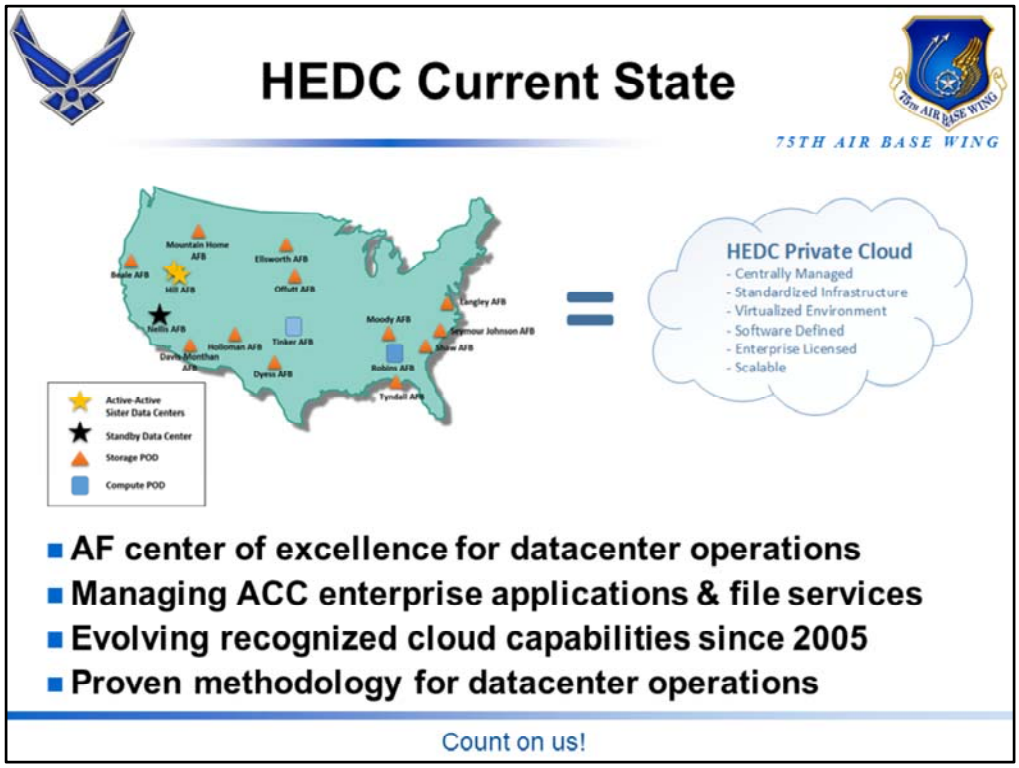

- Manage an environment spanning 15 physical locations consisting of 1630+ servers hosting 450+ applications
- Expanding to 16th location in support of Air Force Sustainment Center (AFSC) Continuity of Operations (COOP) Proof of Concept
- Data center architecture initially designed around hosting ALC ERPs and has grown to accommodate all manners of workload
- Capabilities delivered utilizing standardized infrastructure and software/virtual appliance components
- Environment is over 99% virtualized on shared or dedicated infrastructure
- Centrally managed by a team of 24 personnel
- Manpower requirements significantly reduced by employing automation and standardized processes

Stretched Metro Cluster AZs

75TH AIR BASE WING

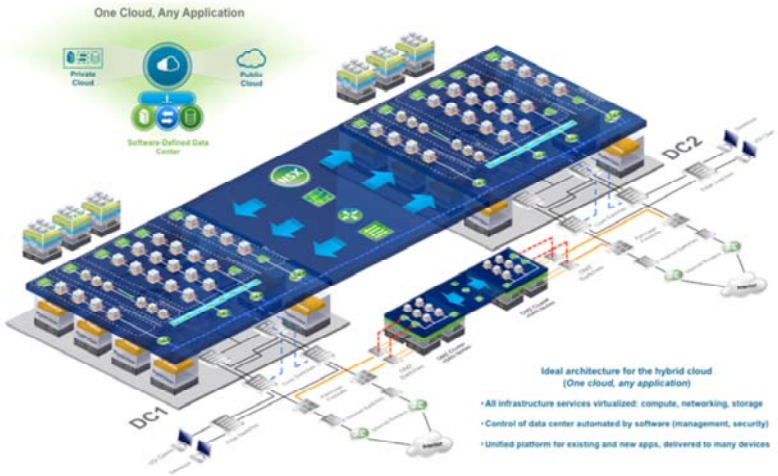

Count on us!

Three AZ, Two Region Disaster Avoidance

75TH AIR BASE WING

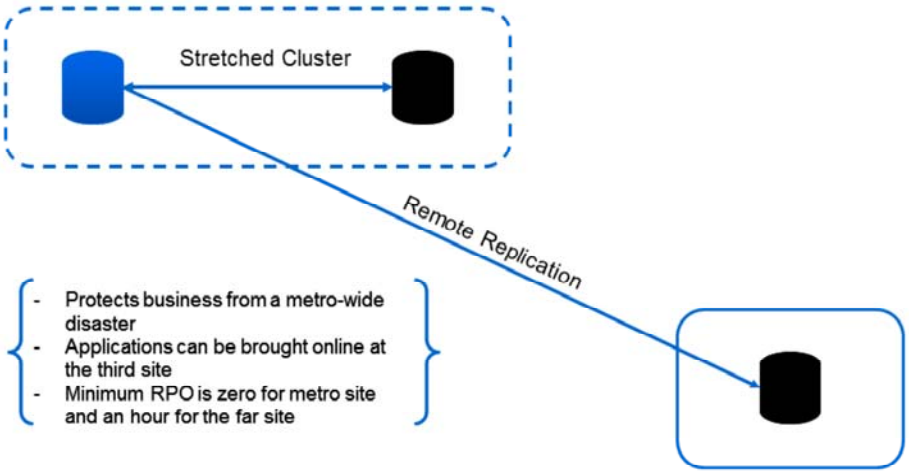

- Protects business from a metro-wide disaster
- Applications can be brought online at the third site
- Minimum RPO is zero for metro site and an hour for the far site

Count on us!

HEDC Earned Value Examples

75TH AIR BASE WING

- Stretched Metro Cluster AZs
  - Zero RTO/RPO
- Long distance Region AZs
- Performance and space improvement
  - 20/1 ratio in last 3 years
- Decreased application latency
  - In human terms from a month to a minute
- Software defined archive cloud
  - 15/1 less expensive
  - 300% faster

Count on us!

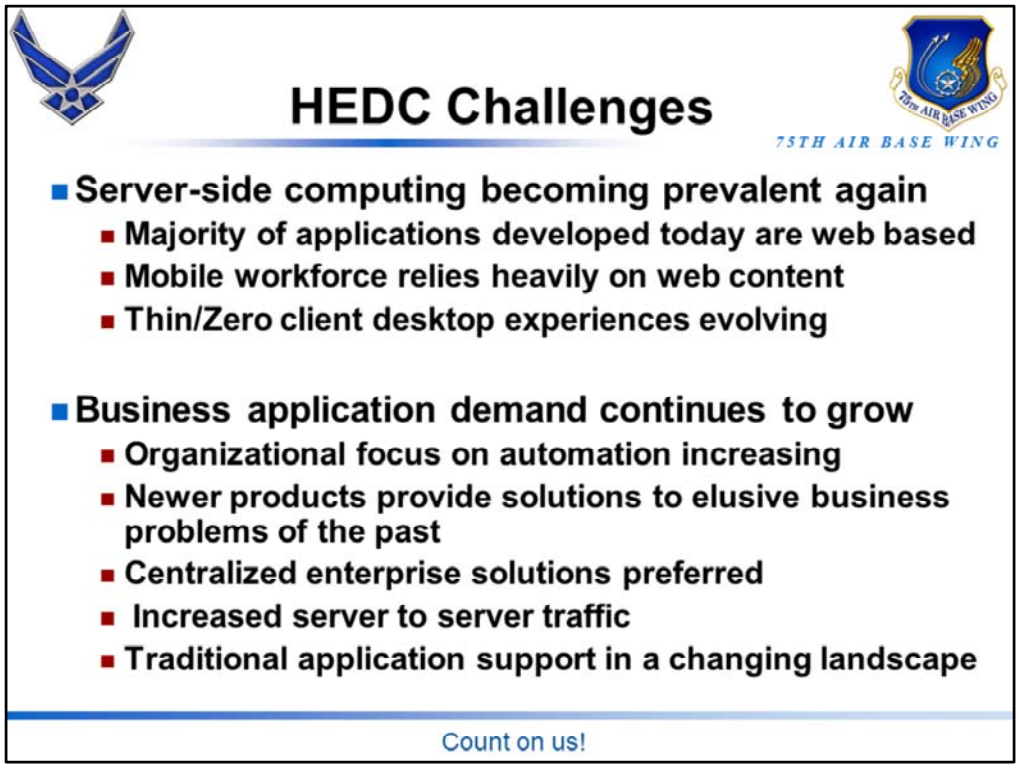

- Demands on server side compute increasing as more applications offload processing to backend processes
- More and more business apps are coming to light as data center consolidation takes hold
- Legacy applications or frameworks do not always translate to cloud environments; requires integration support from all parties

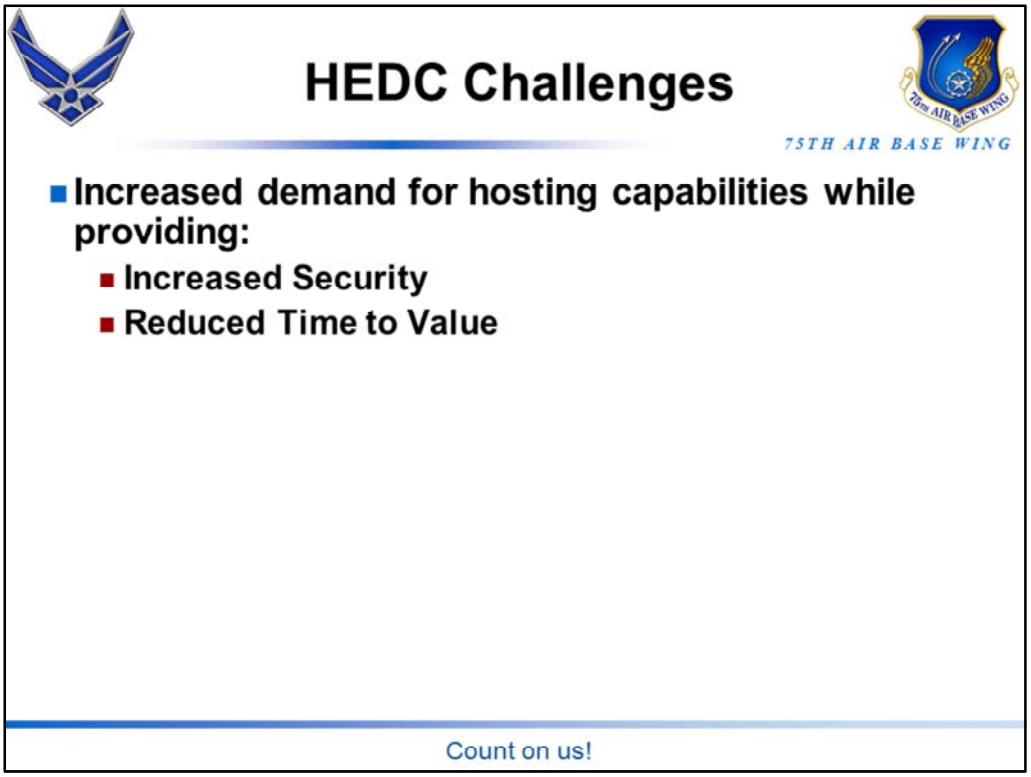

- Increased security and reduced TTV next area of focus in our architecture

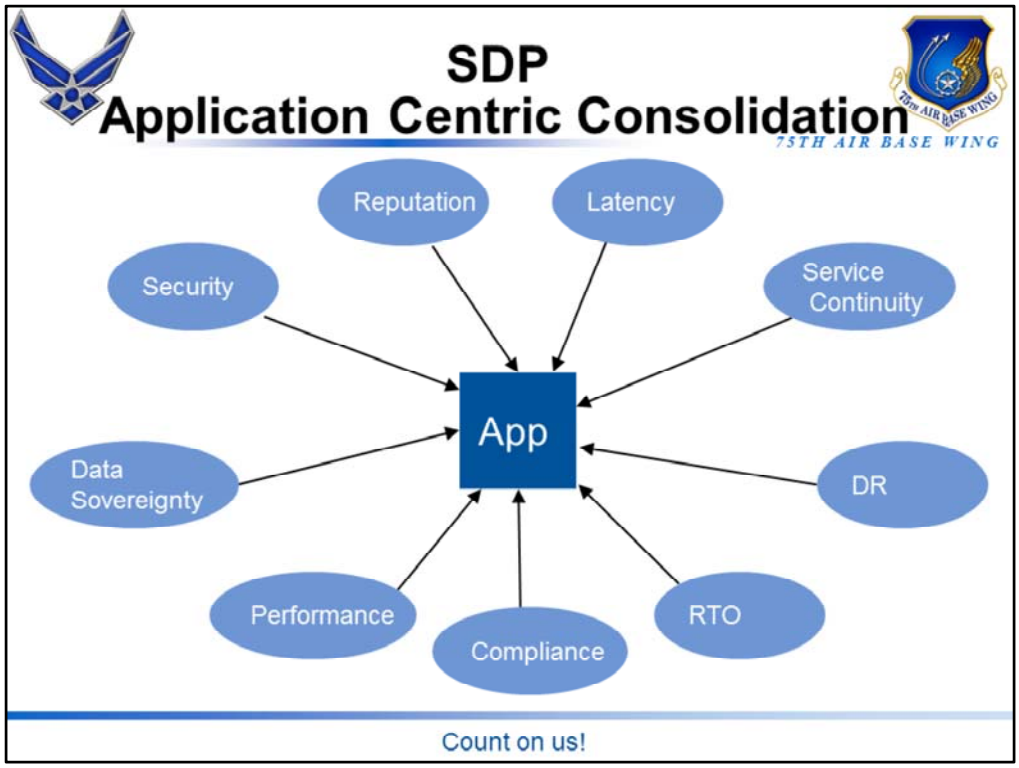
SDP
Application Centric Consolidation
75TH AIR BASE WING
Reputation
Latency
Security
Service Continuity
App
Data Sovereignty
DR
Performance
Compliance
RTO
Count on us!

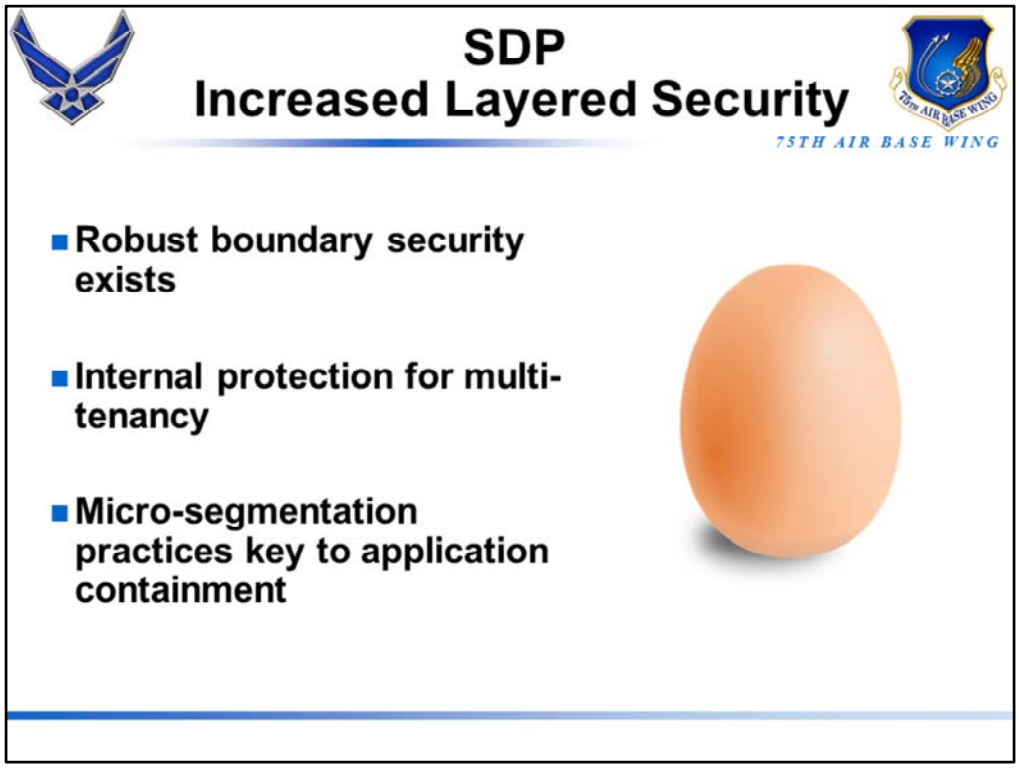

- We have to do multi-tenancy better than anyone else
- Recognizing that there needs to be a balance between micro-segmentation and consolidated services

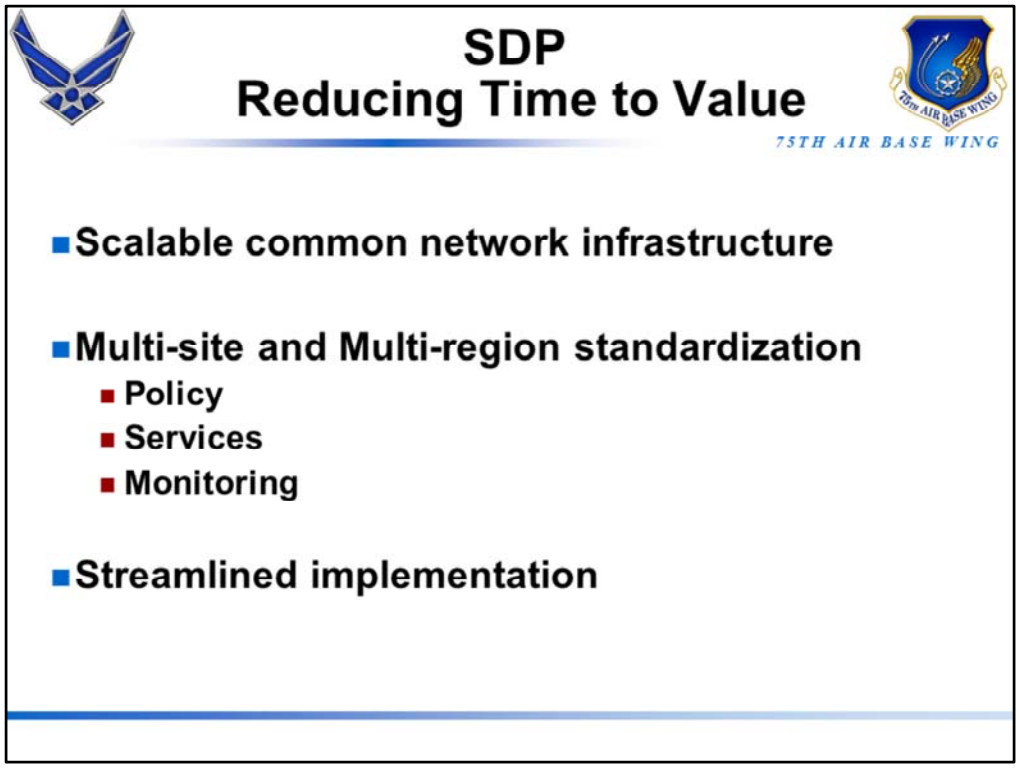

- Managing disparate networks across geographic locations proves challenging
- Standardization drives down sustainment and implementation costs

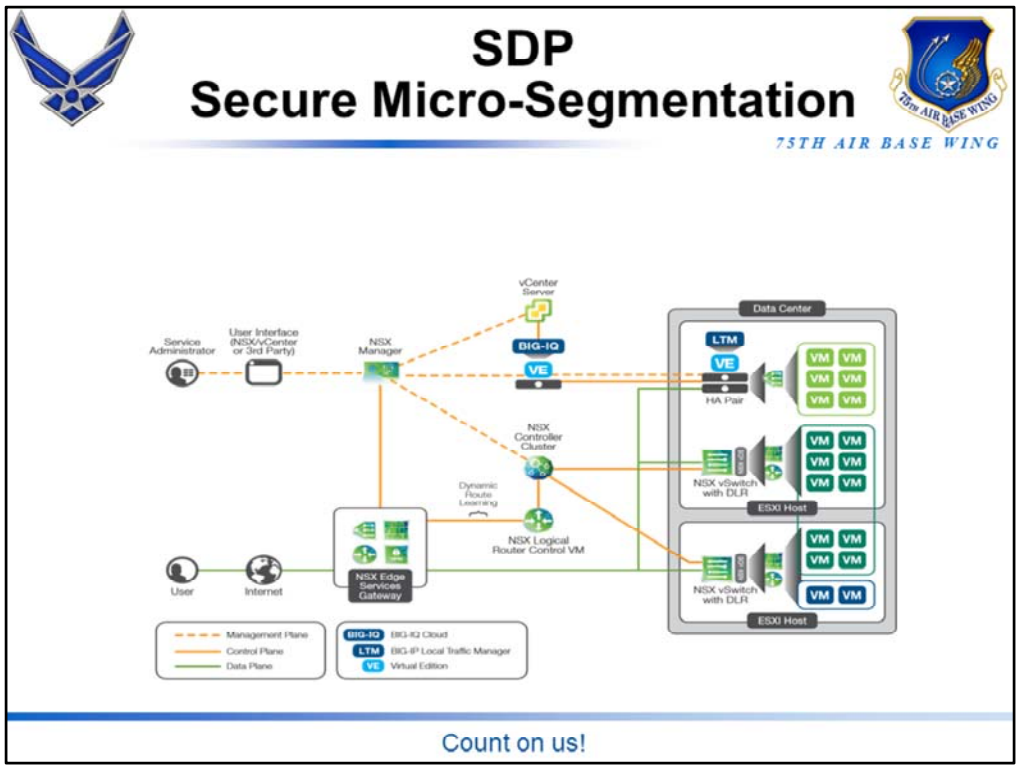

Without a flexible approach to risk management, which adapts to the onset of new technology paradigms, security silos using disparate approaches are created. These silos act as control islands, making it difficult to apply risk-focused predictability into your corporate security posture, causing unforeseen risks to be realized. These actualized risks cause an organization's attack surface to grow as the adoption of new compute technology increases, causing susceptibility to increasing advanced threat actors.

Micro-segmentation decreases the level of risk and increases the security posture of the modern data center. So what exactly defines micro-segmentation? For a solution to provide micro-segmentation requires a combination of the following capabilities, enabling the ability to achieve the below-noted outcomes.

Distributed stateful firewalling for topology agnostic segmentation – *Reducing the attack surface within the data center perimeter through distributed stateful firewalling and ALGs (Application Level Gateway) on a per-workload granularity regardless of the underlying L2 network topology (i.e. possible on either logical network overlays or underlying VLANs).*

Centralized ubiquitous policy control of distributed services – *Enabling the ability to programmatically create and provision security policy through a RESTful API or integrated cloud management platform (CMP).*

Granular unit-level controls implemented by high-level policy objects – *Enabling the ability to utilize security groups for object-based policy application, creating granular application level controls not dependent on network constructs (i.e. security groups can use dynamic constructs such as OS type, VM name or static constructs such active directory groups, logical switches, VMs, port groups IPsets, etc.). Each applcation can now have its*

own security perimeter without relying on VLANs . See the DFW Policy Rules Whitepaper for more information.

Network overlay based isolation and segmentation – *Logical Network overlay-based isolation and segmentation that can span across racks or data centers regardless of the underlying network hardware, enabling centrally managed multi-datacenter security policy with up to 16 million overlay-based segments per fabric.*

Policy-driven unit-level service insertion and traffic steering – *Enabling Integration with 3rd party solutions for advanced IDS/IPS and guest introspection capabilities.*

Alignment with emerging Cybersecurity Standards

National Institute of Standards and Technology (NIST) is the US federal technology agency that works with industry to develop and apply technology, measurements, and standards. NIST is working with standards bodies globally in driving forward the creation of international cybersecurity standards. NIST recently published NIST Special Publication 800-125B, "Secure Virtual Network Configuration for Virtual Machine (VM) Protection" to provide recommendations for securing virtualized workloads. The capabilities of micro-segmentation provided by NSX map directly to the recommendations made by NIST.

Section 4.4 of NIST 800-125b makes four recommendations for protecting virtual machine workloads within modern data center architecture. These recommendations are as follows

VM-FW-R1: In virtualized environments with VMs running delay-sensitive applications, virtual firewalls should be deployed for traffic flow control instead of physical firewalls, because in the latter case, there is latency involved in routing the virtual network traffic outside the virtualized host and back into the virtual network.

VM-FW-R2: In virtualized environments with VMs running I/O intensive applications, kernel-based virtual firewalls should be deployed instead of subnet-level virtual firewalls, since kernel-based virtual firewalls perform packet processing in the kernel of the hypervisor at native hardware speeds.

VM-FW-R3: For both subnet-level and kernel-based virtual firewalls, it is preferable if the firewall is integrated with a virtualization management platform rather than being accessible only through a standalone console. The former will enable easier provisioning of uniform firewall rules to multiple firewall instances, thus reducing the chances of configuration errors.

VM-FW-R4: For both subnet-level and kernel-based virtual firewalls, it is preferable that the firewall supports rules using higher-level components or abstractions (e.g., security group) in addition to the basic 5-tuple (source/destination IP address, source/destination ports, protocol).

NSX based micro-segmentation goes beyond the recommendations noted in the NIST publication and enables the ability for fine-grained application of service insertion (e.g. allowing IPS services to be applied to flows between assets that are part of a PCI zone).

* residing in separate trust zone outside the application's attack surface.

SDP
Increasing Context

75TH AIR BASE WING

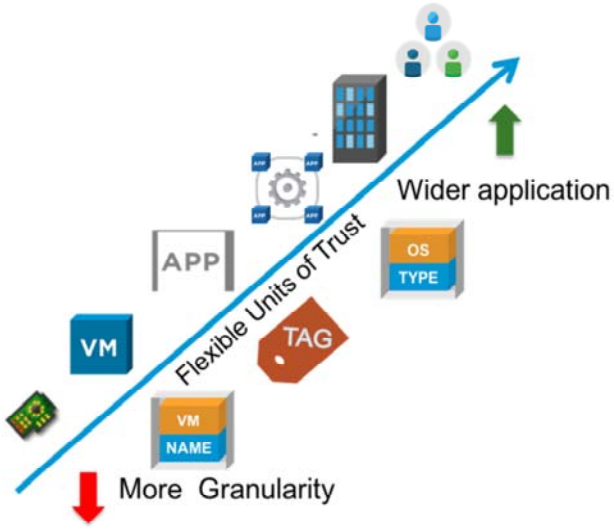

Count on us!

Better Together: F5 SDAS, NSX, ACI, CSPs

75TH AIR BASE WING

Count on us!

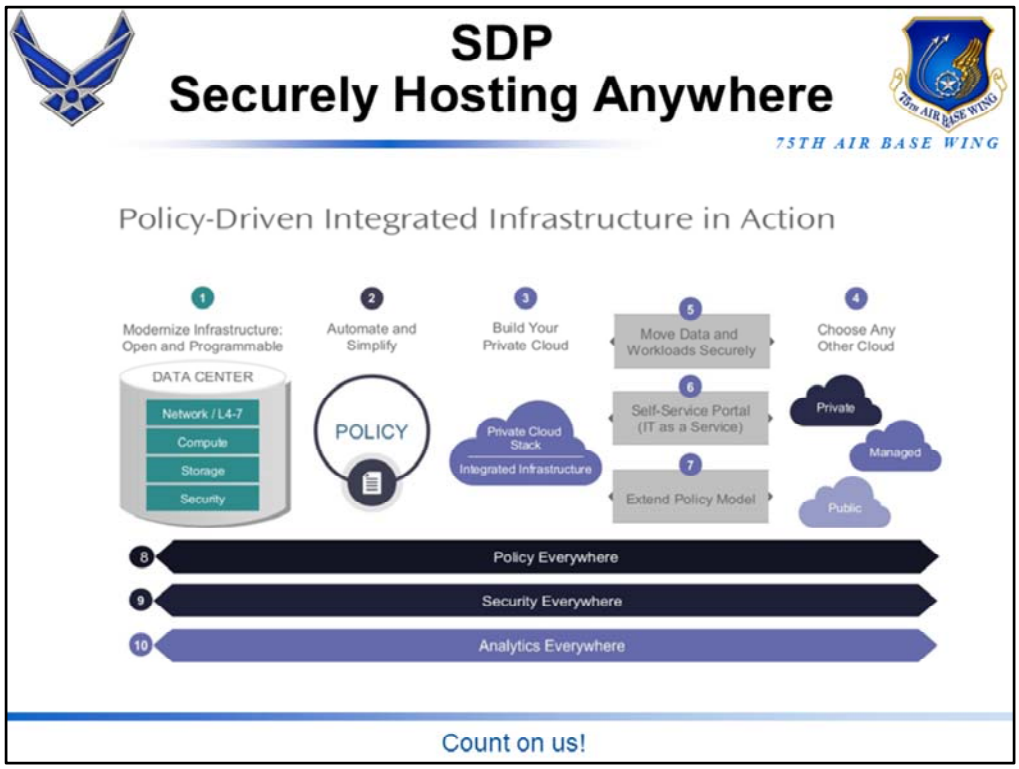

Software-Defined Application Services™ (SDAS)
These services include local and global load balancing as well as encryption, optimization, acceleration, access, and security services.
F5 to integrate their management and control planes to enable automated deployment and orchestration of application services.
SDAS is the next-generation model for delivering application services. It takes advantage of F5 innovations in scalability models, programmability, and an intrinsic decoupling of the data and control planes. SDAS creates a unique application service fabric capable of extending the benefits of F5 application delivery services to all applications, regardless of location. Examples of these SDAS services include load balancing, SSL encryption, DDoS protection, protocol management, connection management, TCP offloading, and content offloading. The BIG-IP platform combines all these services into F5 iApps® policies, thereby simplifying the service provisioning process via re-usable templates.

Policy-Driven Integrated Infrastructure in Action 4 Choose Any Other Cloud Managed Public Private Private Cloud Stack Integrated Infrastructure 3 Build Your Private Cloud POLICY Automate and Simplify 2 Move Data and Workloads Securely 5 Self-Service Portal (IT as a Service) 6 Extend Policy Model 7 9 1 Network / L4-7 Compute Storage Security Modernize Infrastructure: Open and Programmable DATA CENTER 10 8 Policy Everywhere

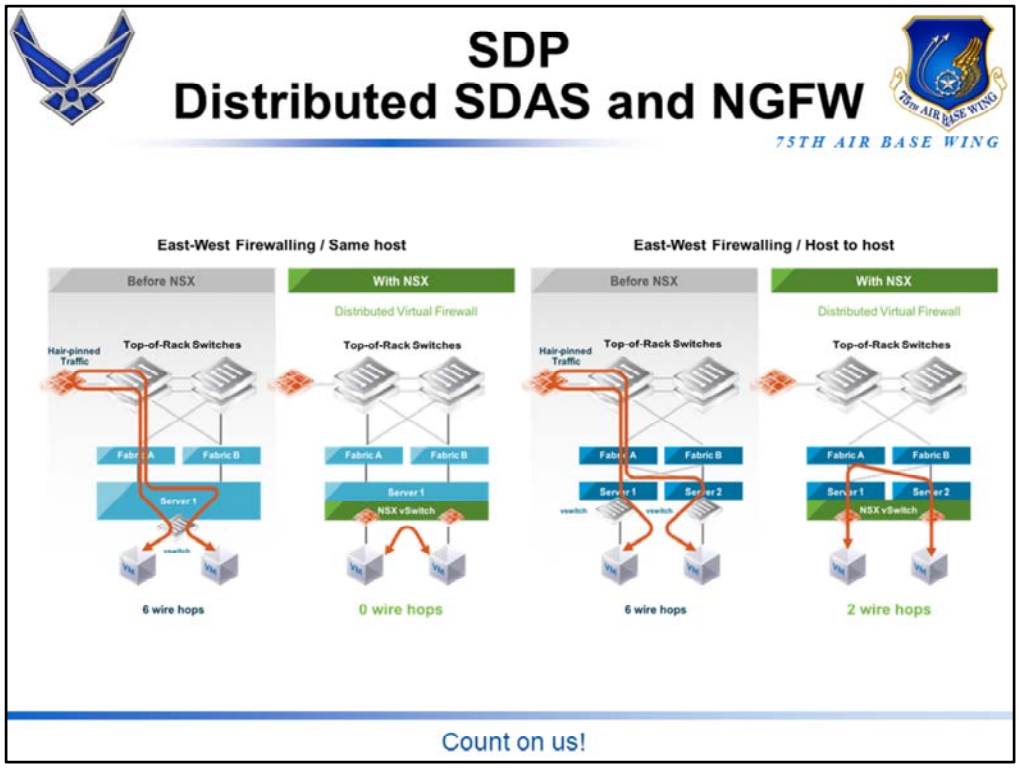

- Managing disparate networks across geographic locations proves challenging
- Standardization drives down sustainment and implementation costs

Cloud Deployment Models

75TH AIR BASE WING

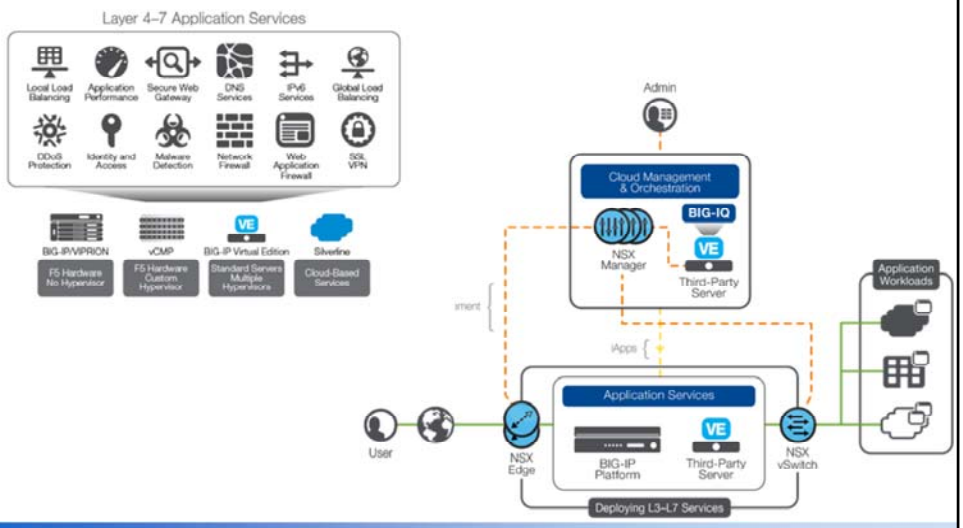

Count on us!

Cloud Native Apps

75TH AIR BASE WING

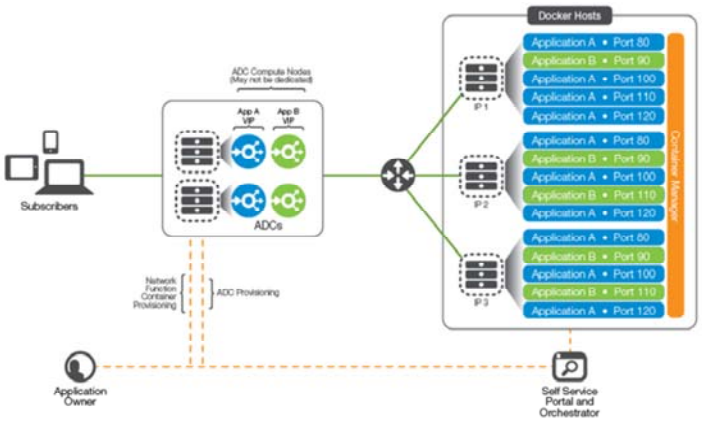

Count on us!

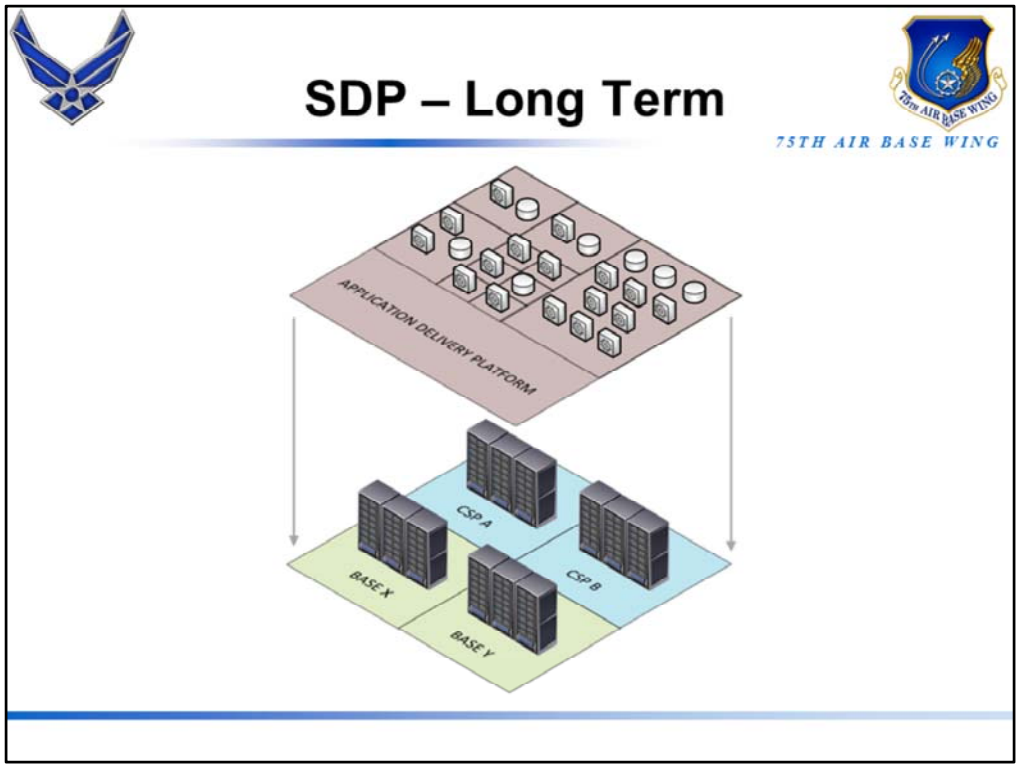

- We won't always control the physical infrastructure
- Platform needs to be agile enough to run on almost any infrastructure
- Integrate with cloud service providers to expand capabilities

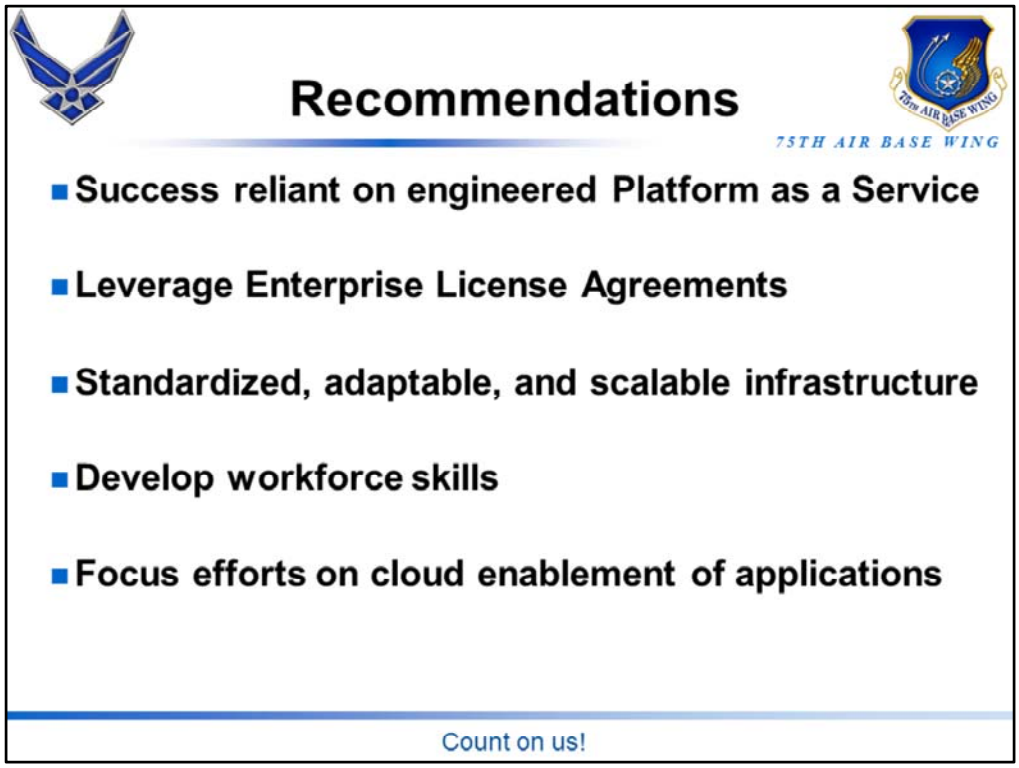

- PaaS is the core to any implementation. IaaS will no longer be an option as we transition to CSPs. We must look at how we can deliver business capabilities, regardless of the hosting environment
- ELAs and ULAs are the only cost effective way to provide PaaS; ensure cloud use rights are included
- Building block infrastructure for what you build in premise. Scalable up and then out
- Start training and getting your workforce ready now for these new concepts; be proactive, not reactive
- Expand integration capabilities and knowledge; more resources are required to integrate than ever before

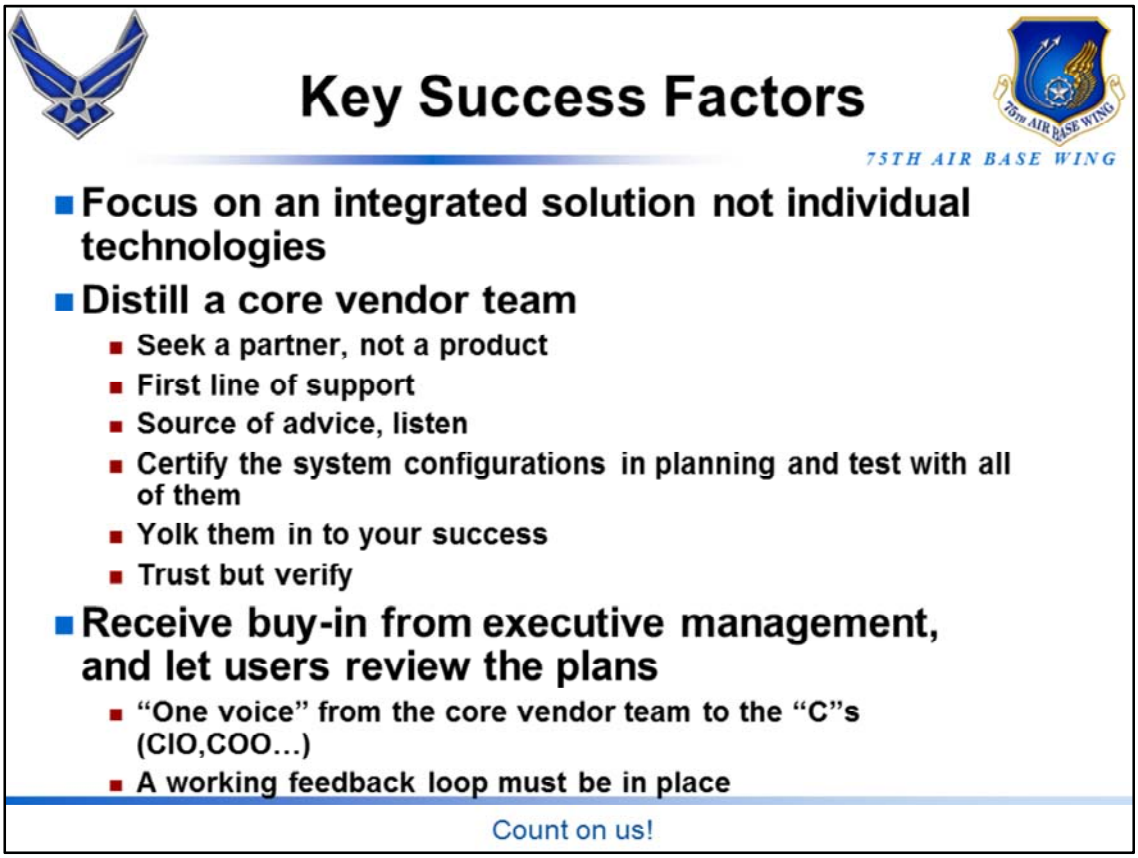

Project BonFire has proven that very large consolidation projects and data center refreshes of significant scale can be implemented with small teams quickly and concurrently

Lessons Learned:

- To embark on a project of this nature, you need a team of highly skilled and experienced engineers.

- You can't guess at performance and scalability. You need to run many proofs of concept as close to actual production as possible.

- It's necessary to distill many vendor relationships into fewer core members that have a vested interest in your success and whom you can trust, to avoid "finger pointing" if things don't go well.

- You need to follow best practices and life cycle management, and get users involved and committed in reviewing the new environment.

- Don't end with one success. Look to build on the foundation over time, instead of seeing it as an isolated, one-time project. For example, Hill Air Force Base is in the midst of building a grid of grids.

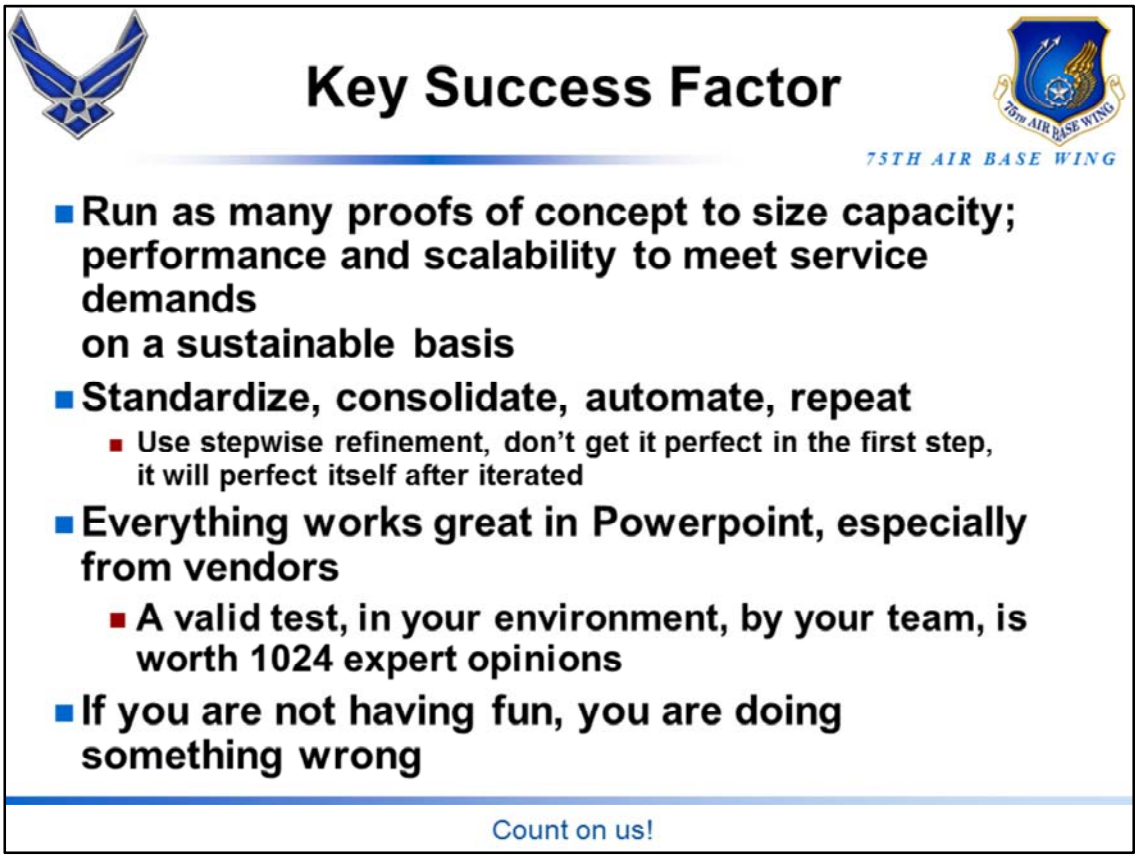

Project BonFire has proven that very large consolidation projects and data center refreshes of significant scale can be implemented with small teams quickly and concurrently

Lessons Learned:

- To embark on a project of this nature, you need a team of highly skilled and experienced engineers.

- You can't guess at performance and scalability. You need to run many proofs of concept as close to actual production as possible.

- It's necessary to distill many vendor relationships into fewer core members that have a vested interest in your success and whom you can trust, to avoid "finger pointing" if things don't go well.

- You need to follow best practices and life cycle management, and get users involved and committed in reviewing the new environment.

- Don't end with one success. Look to build on the foundation over time, instead of seeing it as an isolated, one-time project. For example, Hill Air Force Base is in the midst of building a grid of grids.

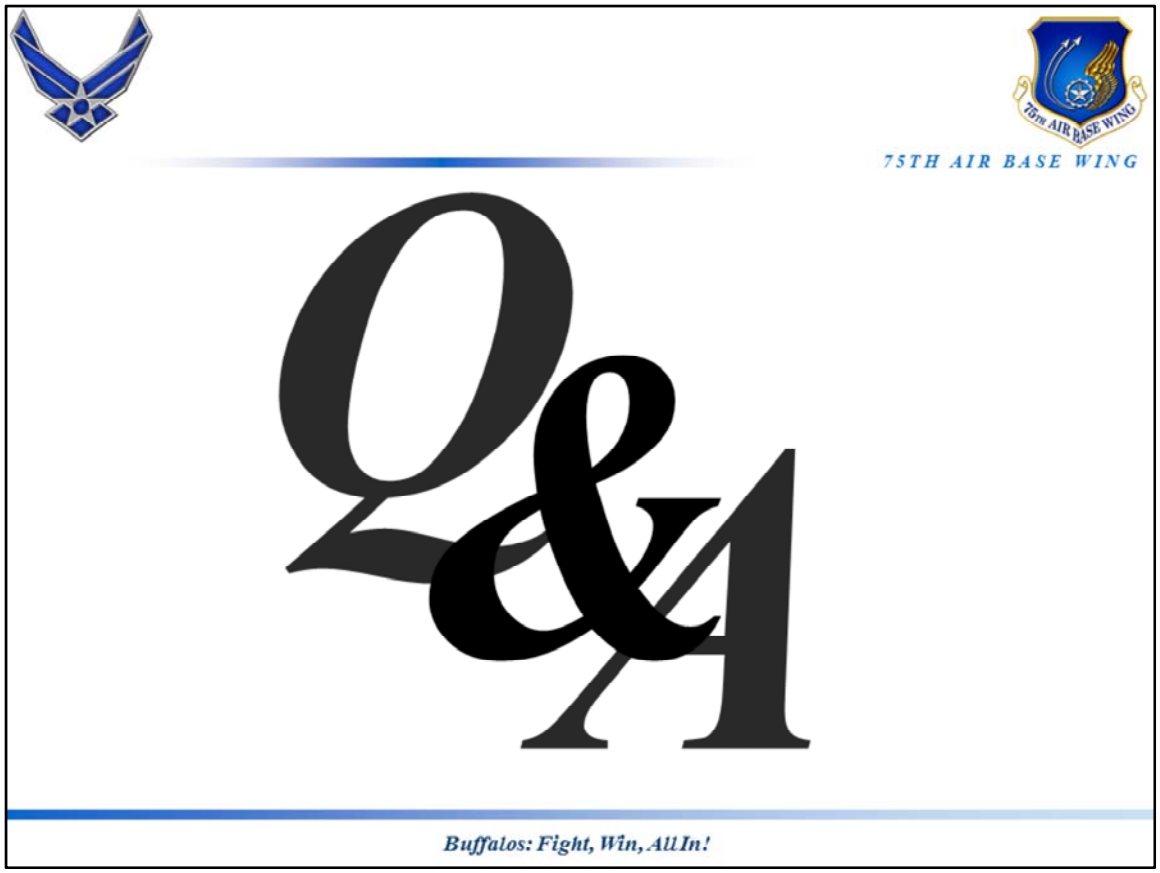
75TH AIR BASE WING
Q&A
Buffalos: Fight, Win, All In!

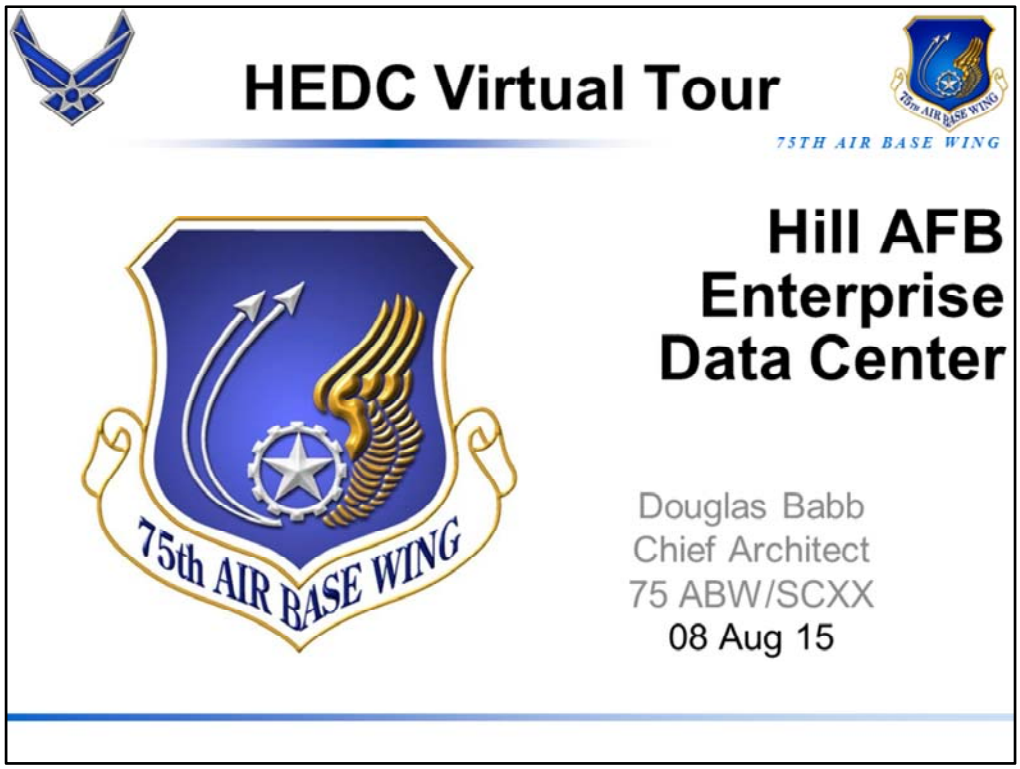

Igniting Change
"*Through Project BonFire, Hill Air Force Base updates data servers and storage systems to assure that 170 apps for jet and missile maintenance are available 24 x 7.*

Design, reliability and screaming speed. When you think Air Force, these three things typically would be synonymous with jets. But they also drove the thinking of the Hill Air Force Base systems staff when they set out to upgrade the Air Force Materiel Command's data center operations.

The command at the northern Utah base repairs and maintains F-16 and A-10 jet aircraft and intercontinental ballistic missiles. The mission: Keep these aircraft and missiles ever-ready for war. To do their jobs, military personnel rely on over 170 applications. Although the apps worked fine, in recent years, server sluggishness and downtime had become a problem, says Mike Jolley, chief of the Operational Policy Branch and program manager for the command's computer center. "

Source: FedTech » Magazine » "Igniting Change"
http://fedtechmagazine.com/article.asp?item_id=278
Other Articles about Project BonFire:
Government Computing News » "Order out of chaos"
http://www.gcn.com/print/26_16/44607-1.html
Washington Technology » "Something to celebrate"
http://www.washingtontechnology.com/print/22_14/31170-1.html
Gartner » "Linux Case Study: The U.S. Air Force Goes From Big Unix 'Iron' to x86 Linux in 290 Days", By George J. Weiss
Gartner ID Number G00156338, Publication Date: 26 June 2008
Business Wire » "Hill Air Force Base Named First Red Hat Innovator of the Year"
http://findarticles.com/p/articles/mi_m0EIN/is_2007_June_12/ai_n27269703

Overview

75TH AIR BASE WING

- Blocking Architecture
- Reducing Latency
- Non-Blocking Architecture
- NearLine Storage
- Hill Enterprise Data Center – Classified (HEDC-C)

Count on us!

Dielectric Fluid Submerged DC

75TH AIR BASE WING

Count on us!

Energy Reduction/Cost Savings

75TH AIR BASE WING

- Submerged DC cooling equivocates to $1 per Watt
- Reduces DC overall power use by 50%
- Submerged DC at 100°F as effective as 70°F air temp.
- Cost for a 10 Megawatt submerged DC by reduced cooling alone of 15% energy consumption translates into $1.5M savings annually

Count on us!

Blocking Architecture

75TH AIR BASE WING

- No Blades Servers Supporting Production Apps

Count on us!

Blade servers have become very popular in Data Centers due to their densities and reduced number of connections required

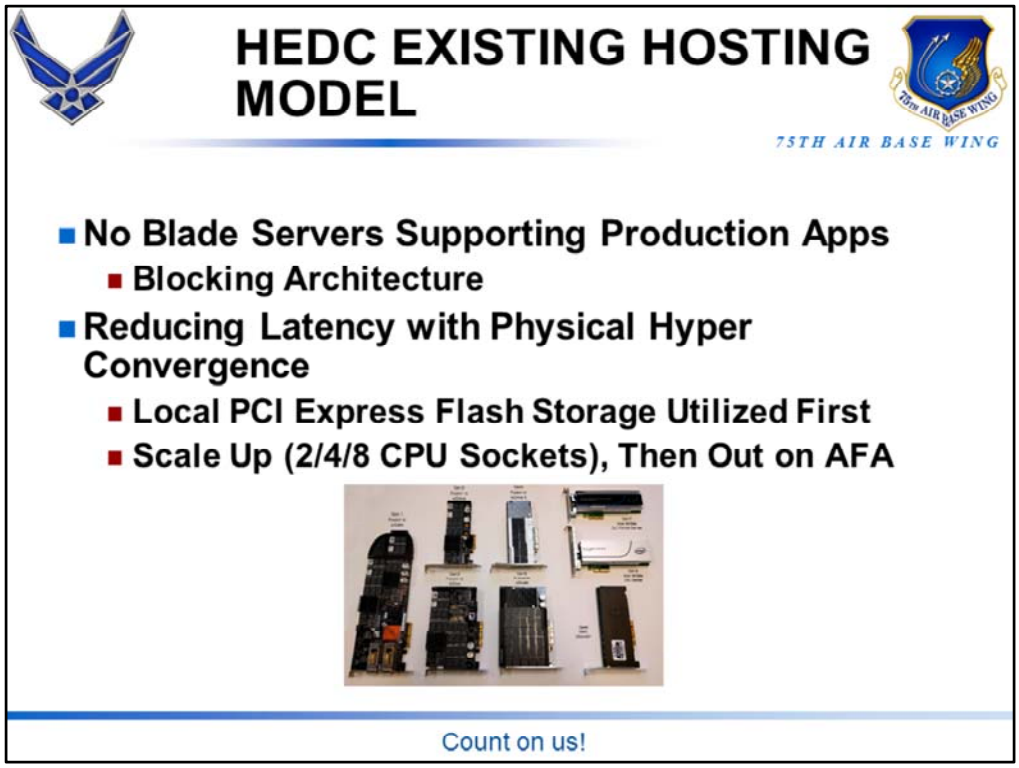

Riding the x86 Wave to reduce Latency Currently, Must Scale Up (2/4/8 CPU Sockets), Then Out on AFA

8 CPU sockets
- 160 vCPUs
- 4TB RAM

4 CPU sockets
- 80 vCPUs
- 1.5 TB RAM

2 CPU sockets
- 48 vCPUs
- 768GB RAM
- PCIe 3

Reducing Latency

Economically Maximize RAM

75TH AIR BASE WING

- Standard x86 servers
  - Fastest low power RAM possible
  - All slots utilized

Count on us!

Reducing Latency

Scale Up, Then Out

75TH AIR BASE WING

- Riding the x86 Wave
  - 8 CPU sockets
    - 1024 vCPUs
    - 4TB RAM
  - 4 CPU sockets
    - 288 vCPUs
    - 1.5 TB RAM
  - 2 CPU sockets
    - 72 vCPUs
    - 768GB RAM
    - PCIe 3

Count on us!

x86 is a family of backward compatible instruction set architecture based on the Intel 8086 CPU.

Reducing Latency
Non-Blocking Fabrics

75TH AIR BASE WING

- Stateless Multi-fiber Push On (MPO) structured glass fabric array
  - Breakout panels support 100 Gb/s per port
- 16 Gb/s storage conections, going to 32Gb
  - 2-6x per node
- 10/40 Gb/s Ethernet connections, going to 25Gb
  - 2-4x per node

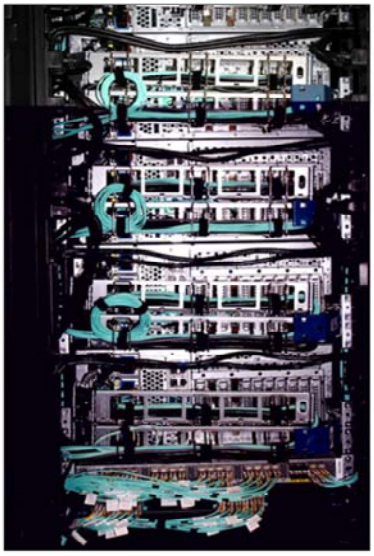

Count on us!

Reducing Latency
Fabric Interface Optimization

75TH AIR BASE WING

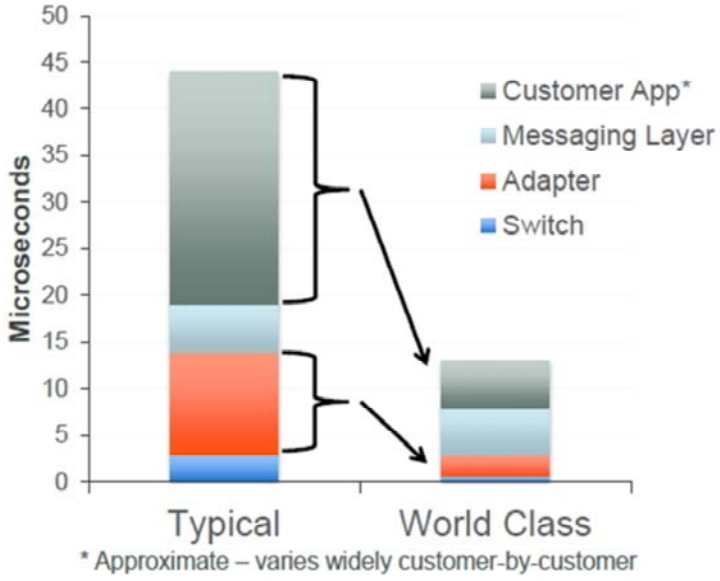

* Approximate – varies widely customer-by-customer

Count on us!

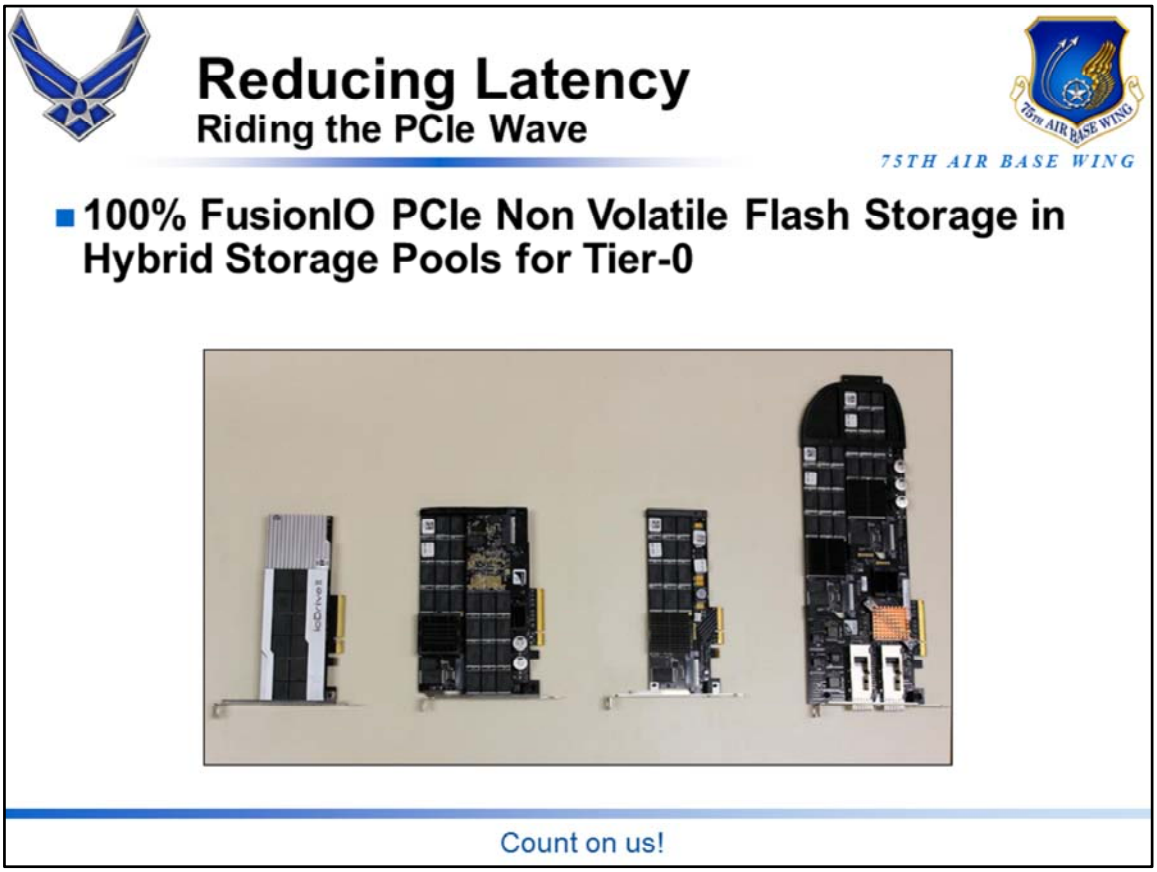

PCI Express (Peripheral Component Interconnect Express), officially abbreviated as PCIe, is a high-speed serial computer expansion bus standard designed to replace the older PCI, PCI-X, and AGP bus standards. PCIe has numerous improvements over the aforementioned bus standards, including higher maximum system bus throughput, lower I/O pin count and smaller physical footprint, better performance-scaling for bus devices, a more detailed error detection and reporting mechanism (Advanced Error Reporting (AER)), and native hot-plug functionality. More recent revisions of the PCIe standard support hardware I/O virtualization.

Reducing Latency
Riding the SSD Wave

75TH AIR BASE WING

- Solid State Drives Utilized in most nodes
  - Up to 7TB, Up to 24x per node possible
- Solid State NVMe 2x per node

Count on us!

Reducing Latency
Next DIMM NV RAM Wave

75TH AIR BASE WING

- Ultra Low Latency
  - 1-3 microsecond response times
- Up to 12TB RAM or NV Storage per node

Count on us!

Reducing Latency
Month to a Minute

75TH AIR BASE WING

The Access-Time Pyramid

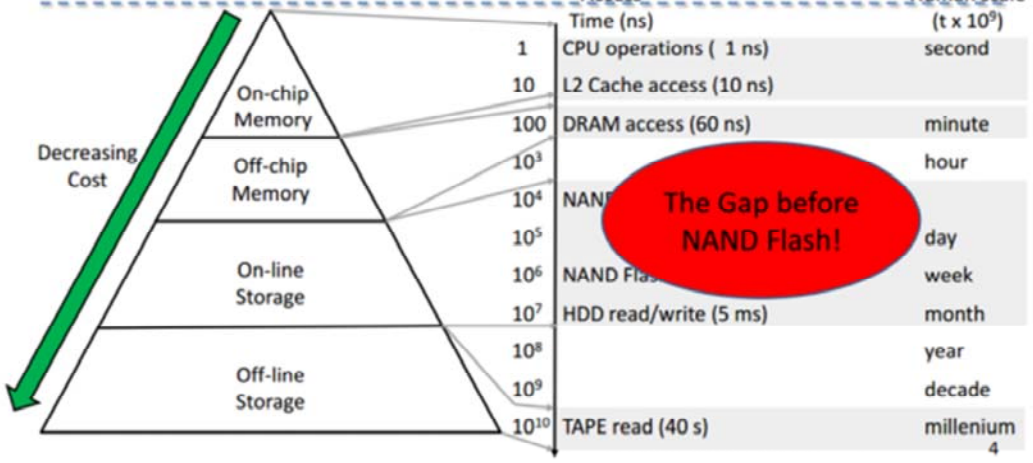

Count on us!

Reducing Latency
Flash Storage Area Networks

75TH AIR BASE WING

- 100% Solid State SANs for Tier-1 LEGACY
  - IBM v840'
  - HP 3PAR 7450's
- Atlantis USX
  - 100% RAM with de-stage to NVMe Flash

Count on us!

Non-Blocking Architecture

75TH AIR BASE WING

- Performance and space improvement
  - 20/1 ratio = 3 years

Count on us!

Nearline Storage Consolidation

75TH AIR BASE WING

- Hybrid software defined scale-out storage cloud
  - Archive cloud 15/1 less expensive
  - 90/1 global dedupe over 30 days
  - 200TB/hour hydration rate
  - Amazon S3 interface

Count on us!


75TH AIR BASE WING
Q&A