Hyper-convergence: Thin Client Solutions
ITES Contract No. W52P1J-16-D-0015

Wildflower International, Ltd.
1500 South St Francis Drive
Santa Fe, New Mexico 87505
505-466-9111

DUE: 22 August 2016

TABLE OF CONTENTS

1. THE WILDFLOWER SPEARS ARCHITECTURE AND USE CASES 1

1.1 SPEARS Architecture 1

1.2 Security Considerations 2
1.2.1 CONUS Security Considerations 3
1.2.2 OCONUS Security Considerations 3
1.2.3 TOC OCONUS Security Considerations 4

1.3 Power Considerations 4
1.3.1 CONUS Power Considerations 4
1.3.2 OCONUS Power Considerations 4
1.3.3 TOC OCONUS Power Considerations 5

1.4 Ease of Use, Management and Deployment 5
1.4.1 CONUS Location Ease of Use, Management and Deployment 5
1.4.2 OCONUS Ease of Use, Management and Deployment 6
1.4.3 TOC OCONUS Ease of Use, Management and Deployment 6

1.5 Availability 6
1.5.1 CONUS Availability 6
1.5.2 OCONUS Availability 7
1.5.3 TOC OCONUS Availability 7

1.6 Reliability 7
1.6.1 CONUS Reliability 7
1.6.2 OCONUS Reliability 8
1.6.3 TOC OCONUS Reliability 8
1.6.4 Commercial off the Shelf (COTS) 8

1.7 Scalability 9
1.7.1 CONUS 9
1.7.1.1 Design for Under 500 Users 9
1.7.1.2 CONUS Design for Over 500 Users 10
1.7.2 Bill of Materials for CONUS Architecture 10
1.7.3 OCONUS Design 11
1.7.3.1 Under 500 Users 11
1.7.3.2 OCONUS Design for Over 500 Users 11
1.7.3.3 Bill of Materials for OCONUS Architecture 12
1.7.4 Tactical Operations Center 12
1.7.4.1 TOC OCONUS Design for Under 500 Users 12
1.7.4.2 TOC OCONUS Design for Over 500 Users 13
1.7.4.3 Bill of Materials for TOC OCONUS Architecture 14

2. SUMMARY 14

PROPRIETARY STATEMENT
This document includes data that shall not be disclosed outside of the Government, used, or disclosed – in whole or in part – for any purpose other than to evaluate the Contractor. If, however, a contract is awarded to this offeror because of – or in connection with – the submission of this data, the Government shall have the right to duplicate, use, or disclose the data to the extent provided in the resulting contract. This restriction does not limit the Government's right to use information contained in this proposal if it is obtained from another source without restriction. The data subject to this restriction are contained in each sheet contained in this document.

The data in the pages of this document where so annotated contain trade secrets and commercial or financial information that are either specifically exempted from disclosure by statute or privileged or confidential within the meaning of the exemption set forth in Section 552(b)(3) and (4), respectively, of the Freedom of Information Act, 5 U.S.C. 552, the disclosure of which could invoke the criminal sanctions of 18 U.S.C. 1905.

EXECUTIVE SUMMARY

Wildflower has been a trusted partner to the United States Federal Government for over two decades. We have worked with our customers to deploy the latest IT solutions and address their demanding and complex business requirements.

Wildflower works with you to modernize critical information technology, so you can achieve mission goals, reduce operating costs, and improve performance. By delivering intelligent and cost-effective IT systems, we drive your business and mission forward. We incorporate existing capabilities to develop tailored, leading-edge solutions while ensuring you're on course to meet the demands of the future.

Wildflower considered the many challenges facing the United States Army's ability to deploy thin client solutions. We built our solution on a Hyper-converged model. Hyper-convergence offers the economic benefits of the cloud while delivering performance, high availability, and reliability.

The Army faces unique challenges when deploying solutions in CONUS and OCONUS environments. Stateside, facilities must conform to government-wide security policies, DoD policies, and site-specific requirements to protect information. Data is vulnerable to misuse and theft by end users via both intentional methods such as thumb drives, or external transfers. Unintentional misappropriation of data may occur when end users download malicious code. In OCONUS environments, information is even more vulnerable; for example, foreign nationals from rogue states as well as individuals working for hostile organizations are continually seeking access to the data in OCONUS facilities. This pursuit is via multiple attack vectors, which include stealing end user devices, hijacking end user devices, and sniffing end user devices. IT is faced with a constant barrage of these attacks especially when end users have multiple devices on their person in OCONUS fields of deployment

This document will outline Wildflower's SPEARS™ architecture as it applies to environments with greater or less than 500 users in both CONUS and OCONUS deployments. Additionally, we will articulate the value of Virtual Desktops/Thin Clients in the battlefield within a Tactical Operations Center (TOC) Design.

Wildflower's SPEARS architecture focuses on 6 key elements. They are:

- ✓ Security
- ✓ Power Consumption
- ✓ Ease of Deployment Given Conditions
- ✓ Availability
- ✓ Reliability
- ✓ Scalability

1. THE WILDFLOWER SPEARS ARCHITECTURE AND USE CASES

1.1 SPEARS ARCHITECTURE

Wildflower's virtual device architecture has been deployed in government and commercial installations from 50 to 90,000 devices. The device types vary from repurposed PC's to thin clients and simple mobile

tablets. This architecture is general purpose and can be deployed for either Citrix or Horizon (with minor modifications). Future iterations will include other alternatives.

Wildflower believes that OCONUS field deployments may have less technical team members present than CONUS deployments. Therefore, we recommend that OCONUS deployments use Horizon due to its easier initial deployment and lower time to maintain.

Ultimately any virtual device architecture will incorporate the following components: (see diagram below)

- Client Side Receiver – software installed on an end user device
- Storefront/App Store- authenticates users to the resource pool and provides access to apps
- Database- typically a SQL instance used to store configs and session information
- Delivery Controller- centrally manages the services that the hypervisor and delivery agents will deliver to end users; load balances user connections
- Management Console- access to and management of the deployment
- Director –monitor and troubleshoot an environment
- License Server- Manages licenses.

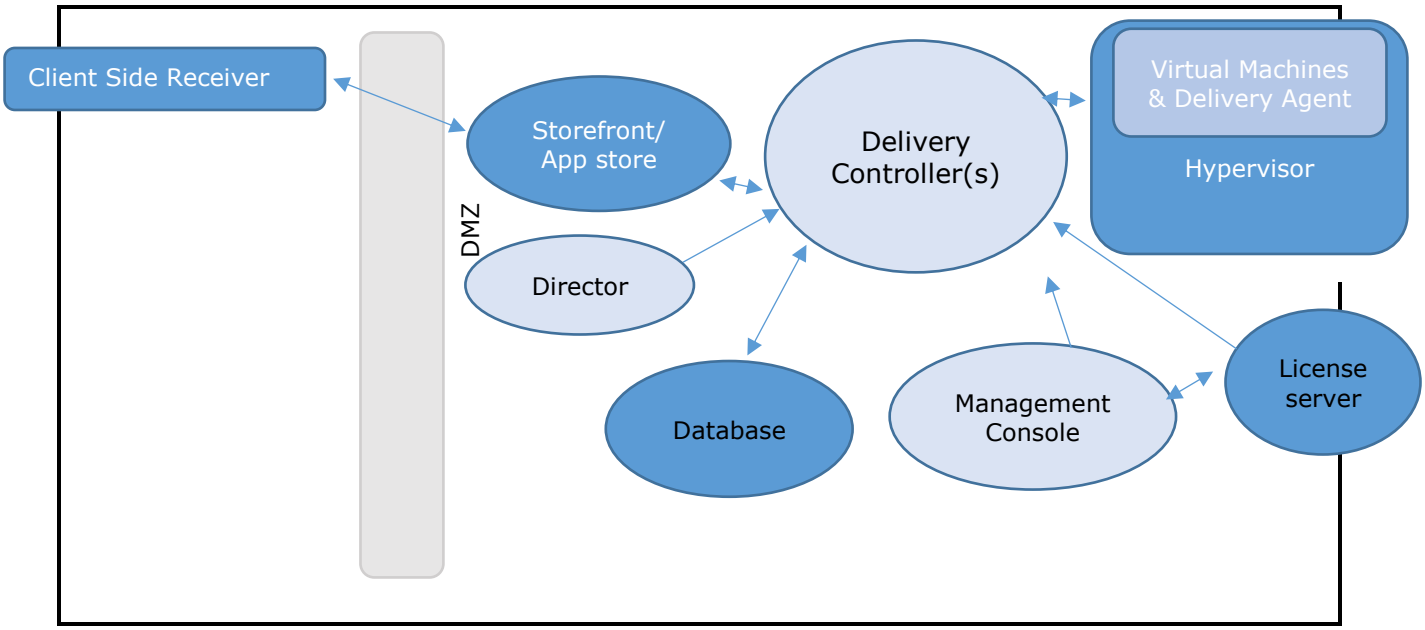

Virtual Device Software Architecture

Once the software architecture is defined, Wildflower implements SPEARS architecture to consider:

1. Security
2. Power Consumption
3. Ease of Use, Management and Deployment Given Conditions
4. Availability
5. Reliability
6. Scalability

1.2 SECURITY CONSIDERATIONS

The location's security policies should mandate a secure process for physical access to the facility, access to the location of the equipment, and include locked and keyed rack systems. The number one challenge a

controlled environment deployment faces is internal breach. Thus, locations must ensure all USB write capability is disabled on the equipment running the core of the architecture, the devices accessing the information, and ensure that all equipment is FIPS 140-2 compliant.
Wildflower advocates a strong security practice in the design of the physical security, paired with a strong digital security policy. Digital security policy must ensure that logging and other checks and balances are in place so no single end user can compromise the installation, intentionally or otherwise. All employees with access to modify profiles or administer Active Directory should be thoroughly screened. While end users access mission critical applications within the system, their sessions should be through secure tunnels and only use a thin client as a presentation layer to prevent removal of data. Policies should be in place for GEO fencing and the management of mobile devices within the facilities. These policies can be set in a profile management tool and incorporated into the design of the solution. Most secure OCONUS and CONUS sites are requiring mobile devices to be sequestered while on premises. If this is not the case where an installation is located, then a contextually aware security solution with MDM capabilities can be used to assist with security at the end points.

1.2.1 CONUS SECURITY CONSIDERATIONS

All installations and configurations of the systems will occur on site when installed at a CONUS site. This is unique to CONUS sites because of their locations on US soil. The benefit of an onsite installation is the shipped equipment will not have any critical information in transit. At a CONUS controlled site where the physical security is a level III or IV on the Interagency Security Committee Facility Security Level Determination Matrix, it is recommended to deploy the Citrix Workspace Suite with a high availability design as shown on page 2 of this document. While Citrix may be more difficult to deploy than Horizon, it provides greater end user granularity as well as better protection on multiple devices.

1.2.2 OCONUS SECURITY CONSIDERATIONS

OCONUS deployments can be in locations that may not be easily accessible, or may be in hostile locations that limit the ability to have equipment imported. The primary challenge may be physical access. Locations such as embassies have become vulnerable to attacks, and therefore OCONUS deployments must consider all previously mentioned security requirements, along with methodologies for data destruction and encryption key management.

One option for OCONUS locations is to deploy the solution in a secure CONUS facility and allow users to remotely access the system. There may be challenges with bandwidth and connectivity, which may render this option unacceptable. If it is necessary to deploy the solution OCONUS, Wildflower recommends preconfiguring the clusters prior to deployment. Additional security precautions should be implemented including secure, tamper proof shipping ideally via a US asset shipping program that is FIPS compliant. Upon arrival the product should be inventoried and a thorough virus scan should be administered. Additionally, OCONUS sites must make the key management server(s) for key encryption on the hard drives a remote server, or at a minimum one(s) that can be destroyed upon penetration of the facility.

Assuming physical security challenges can be managed, OCONUS sites are typically more vulnerable to digital attacks. These can be as simple as network taps and sniffing, or more complex using Trojans or viruses. In order to ensure no data is accessed on a physical device, all sessions should be non-persistent and no data should remain on the thin or zero clients. All USBs should have write access disabled and any mobile devices used with the deployment need to have MDM solutions and end point protection.

OCONUS sites should attempt to use application authentication remotely via SAML or Kerberos and thus not have applications onsite. Granted, this presents a trade off as it is better for security but it may increase latency. Lastly, authentication should include a contextually aware aspect. This could include a network handshake, an access point authentication, or another method that is location specific.

1.2.3 TOC OCONUS SECURITY CONSIDERATIONS

Tactical Operations in hostile locations brings a handful of challenges to the standard deployment, and render Tactical Operations Center activities distinctive. Common issues include increased security and the need for more redundancy in infrastructure.

TOC deployments may need to be mobile and may not be in a standard data center. A hyper-converged infrastructure may be deployed in a Tactical Wheeled Vehicle (TWV) at a Battalion level, or a field deployed data center at the Brigade level. Either way, the Company Level soldiers that are on the front line and communicate to the centralized resources via satellite or line of site (in the case of vehicle deployed solutions) have unique security concerns that must be addressed. Obviously, the physical challenges of a battlefield deployment take precedent when it comes to security. If a soldier's device is compromised all data and applications must be capable of immediately being wiped. The same situation must be considered for vehicle based deployments, as well as brigade level deployments in hostile areas.

Hostile arenas of combat may also have subversive tactics being deployed by enemy forces that affect the security of the deployment. From a digital perspective, local SPAN, Spoofing, Denial of Service and other tactics are being deployed by non-friendly forces to reduce front line communication capabilities. Often, soldiers that are in smaller Companies may not have access to line of site communication. Devices they employ may need to have application intelligence, a method to cache information, and provide 2-way communication when available. Authentication rules may need to be altered so that tokenization or other methods for authentication are used instead of standard AD acknowledgements. Using a process that has built in biometrics may be a way to overcome authentication issues that may be present in traditional architectures. For example, optical recognition in a multi factor authentication process is a method to incorporate additional security in a quick fashion. Other contextually aware methodologies should be considered as well. Common examples include using GPS to ensure a given device is within a set of geo defined parameters that an operation is occurring. If the device is outside of these boundaries, or does not authenticate within a given period of time, the device will be wiped.

Soldiers need access to applications that can manage devices such as drones that are used to gather field level intelligence. If this intelligence is stored on the end user device it is vulnerable; alternatively, it may be broadcast securely to a brigade level facility and then relayed to the appropriate device with no storage. Again, if communication is hampered and the device cannot receive this information, then a failback to a line of site communication protocol should be enabled. Another option may be a meshed network that is self-healing if a device is compromised.

1.3 POWER CONSIDERATIONS

1.3.1 CONUS POWER CONSIDERATIONS

Typically, secure CONUS sites have redundant power to the location. Wildflower's design approach recommends a minimum of dual UPS and dual PDUs with redundant power supplies in all components to ensure power is continuous and clean. Ideally, there is an N+2 design to the power grid and PDU deployment. All Wildflower CONUS designs will have a minimum of 2 power supplies in each configuration. As part of the final engineering design process, we will determine the type of power (110, 220, etc.), cable lengths, and locations.

1.3.2 OCONUS POWER CONSIDERATIONS

OCONUS locations may be sensitive to power fluctuations depending on their location globally. This being the case, all power should be conditioned to the equipment and proper UPS batteries or additional generators should be available in case of emergency. Wildflower recommends all OCONUS sites run a minimum of 3 PDU's and maintain an additional unit on site in case of failure. All hyper-converged OCONUS architectures will have a minimum of 3 power supplies. Prior to deployment a power assessment should

occur that documents the power type (i.e. 110, 208, 220, etc.) used in each facility. Once this is determined, it is essential to ensure the node manufacturer can support the given power infrastructure

1.3.3 TOC OCONUS POWER CONSIDERATIONS

Front facing Tactical Operations Centers tend to have limited power. This limitation tends to exponentially decrease the further removed and smaller a core operating unit is. For example, a company may have no method to recharge devices (minus solar power), so it is necessary for them to either rendezvous often to get new devices (thus non-persistent technology applies), or to return to a Battalion or Brigade and recharge. Most Battalions and Brigades are using either diesel generators or solar to power their facilities and servers.

Using a solid state drive in this environment lowers power consumption by about 30x versus traditional spinning disk. The 4 node cluster used for deployments under 500 users in Wildflower's solution uses just slightly more electricity than a common hair dryer. Field based deployments should consider limiting applications to save power, if resources are scarce. This is best done with a strategy that focuses on essential versus non-essential applications; only deploying essential applications in the field, and remoting into applications that are non-essential via thin apps/xen apps.

1.4 EASE OF USE, MANAGEMENT AND DEPLOYMENT

Wildflower believes all technical designs should be easy to install, operate and use. Given this philosophy we advocate the use of systems that are modular and highly integrated. The advent of hyper-converged infrastructures such as Nutanix, and Cisco, HP and Dell equivalents, are recommended, and may be interchangeably selected based on the customers' vendor management and relationships.

Hyper-converged architectures are designed to incorporate a true scale-out methodology. They are comprised of a number of nodes that are clustered and that each have the same basic resources available. A node is comprised of some amount of CPU, memory and storage, much in the same fashion a traditional server is designed. The major difference between simply using commodity servers to scale an architecture and using nodes to build a converged architecture is that the combination of nodes (known as a cluster) share resources with each other.

Clustering in the converged architecture provides 2 major benefits. First, an entire node can fail and a properly balanced cluster will remain functioning. Second, nodes can share resources with other nodes. This allows a node that may be running a resource intensive workload, which may normally run out of a resource like memory using a traditional server cluster, to share memory or other resources from another node. These two features allow converged architectures to have massive scalability and are the primary reason companies like Google, Yahoo, and Amazon, deploy them in their datacenters.

Clusters have the ability to provide greater up time, more availability, and easier methodologies to deploy. They are also capable of using replication tools to migrate data to secondary data centers in the case of emergency. Given the compact nature of a cluster (4 nodes can fit into a single 2U chassis), the entire installation can be picked up and moved to a new location. Keep in mind this benefit is also a threat since a cluster could be stolen if a perpetrator knew what they were looking for. Given this situation, Wildflower always recommends using self-encrypting hard drives on the system and a remote key manager.

1.4.1 CONUS LOCATION EASE OF USE, MANAGEMENT AND DEPLOYMENT

At a secure CONUS site, Wildflower would recommend a node based architecture that tightly integrates CPUs, storage and connectivity. This architecture should be connected via a minimum of a 10GbE network and capable of running in excess of 10 Gb/second based on multi-pathing and/or meshing of the switch architecture. Endpoints should reside on or have access to a minimum of a 1GbE network. Located on US

soil and accessible, installations can be done onsite versus remotely. Converged architectures can be brought into production quickly, and can repurpose existing network and security infrastructure.

Additionally, converged architectures that support thin client environments should have a profile management solution incorporated into their application design. When end users can access a self-service portal and request authorization for specific applications, it eases the help desk team's need to interact and provides the end user a faster response to their request.

1.4.2 OCONUS EASE OF USE, MANAGEMENT AND DEPLOYMENT

OCONUS sites may require hazard pay and can potentially be located in areas that are not easily accessible. Therefore, installations on these sites should be performed remotely if possible or pre-configured systems can be deployed. These systems must be pre-configured and tested prior to shipment and all encrypted drives must ship separately from each other and the key management server.

Self-service is essential for OCONUS deployments as staffing may be limited and communication may be challenging. Therefore, Active Directory integration is essential for authentication and profile management.

The thin clients used in OCONUS locations may be more difficult to replace and dispose of. It is recommended that thin clients used in these locations be standardized and updated frequently as a failure may not be capable of being managed as easily. An onsite depot of spare thin clients along with spare components for all of the hyper-converged systems is recommended.

1.4.3 TOC OCONUS EASE OF USE, MANAGEMENT AND DEPLOYMENT

TOC deployments need to be turnkey. All systems need to be tested and configured prior to entering the field of battle. The further in the field that the equipment is deployed, the lower the odds that someone with technical skills will be available to troubleshoot. All deployments need to be capable of repair and re-deployment via remotely managed processes. Brigade level deployments may have the traditional skill sets necessary, but company level deployments, especially those in TWVs, most likely will not. All equipment must be hardened and encrypted. Environmental factors such as dust, heat and vibration need to be considered and thus solid state drives versus spinning disks again are the logical choice to be deployed in the infrastructure. This minimizes heat dissipation and is not as affected by dust or vibration.

Operationally impactful applications such as PIRs must be preconfigured on devices and should allow commanders critical information requirements (CCIR) to be relayed both to the brigade and other platoon/company members via the TIC systems. Any Specific Intelligence Requirement (SIR) that needs to be relayed from the brigade to the deployed teams must be accessible via minimal connectivity. Often this information is relayed to the battalion via the TWVs and then sent to the soldiers via a meshed or line of site network. Additional considerations should be given for any Land Warrior-like systems that are currently in design by the DoD.

1.5 AVAILABILITY

1.5.1 CONUS AVAILABILITY

Secure CONUS sites typically can be accessed by badged vendor personnel for maintenance repairs and replacements which makes the need for onsite spares not as essential as OCONUS or TOC deployments. However, all systems should have components that are fully redundant and highly available in an active architecture that enables all resources to be used all the time by any transaction. Most architectures cannot provide this level of redundancy without relying on multiple local SANs for redundancy as the limiting factor is typically the Storage infrastructures' drive enclosures. RAID cannot account for a full Enclosure failure, and without a central repository for storage like a SAN or NAS, most hypervisors cannot perform

the necessary HA functions. Hyper-converged architectures eliminate this challenge by spreading the data over multiple nodes so at any point 2 or more nodes have the same data.

If the site experiences a full outage for any reason, a redundant configuration should be available with minimum downtime at a secondary site. Hyper-converged architectures lend themselves well to multi-site deployments given their ability to replicate not only data but the other two critical resources of CPU and connectivity as well. Hyper-converged architectures are also simple, scalable architectures which are easy to move and repurpose if a site will be out of commission for a longer period of time.

1.5.2 OCONUS AVAILABILITY

OCONUS locations must not only be fully redundant but should also incorporate replication and failover options to remote sites. This may be difficult given bandwidth limitations in some locations. If bandwidth is a factor, then WAN optimization and SDWAN options should be considered. All locations should implement secure tunnels from site to site for this replication. Options include VPN as well as high speed encryptors. Again, if the solution can reside CONUS and remote users can access the data in a presentation layer-only mode via a browser or chrome-like device, these options should be considered.

Given the potential for limited or interrupted bandwidth at OCONUS sites it is recommended that all OCONUS deployments have full spares kits on site and additional redundancy built into the designs.

1.5.3 TOC OCONUS AVAILABILITY

All TOC deployments need to have multiple redundant components available and should employ standardized FRU's. Hardware failures are unacceptable in the field of battle. Additional considerations are connectivity based. As a company or battalion deploys to a strategic target, their communications can be compromised if weather, topography, or technical issues occur. Therefore, the device the soldier uses to communicate with the centralized resource should contain some intelligence that allows them to still have critical application access regardless of connectivity.

This access challenge is known as persistence. Two major types of persistence exist: application and session. Applications may not have continual connectivity to the device, thus ideally the application resides on the device and the data is downloaded periodically. The same challenge arises when VDI-like systems are built out where session persistence is necessary. If a non-persistent desktop loses its session it will do the equivalent of a reboot which may force end users to re-authenticate, lose application data and ultimately jeopardize the individual. In order to minimize this risk, applications on the device should be levered with backend data flow occurring. However, when applications reside on a device it must be capable of being wiped using some type of MDM solution.

1.6 RELIABILITY

1.6.1 CONUS RELIABILITY

CONUS deployments occur where transportation and availability of equipment is standard. Less expensive equipment can be used as there is no need for rugged or other specialized equipment as might be the case in an OCONUS battlefield environment.

Wildflower does recommend using a brand name converged architecture that has 1000 or more units in service, a full 24 x 7 x 365 support infrastructure with proactive monitoring and a history of reliability. Each node should only be run up to a point where the total utilization of the system can still be distributed to the remaining nodes if one of the nodes fails. While 3 node architectures are supported, Wildflower recommends not implementing any systems under 4 nodes for additional reliability.

The reliability of the core architecture should be designed to overcome issues with end user devices. End users should be able to simply get a new device, use their credentials to log in, and have their applications available as if they were working on their original system. This is also an important construct that must occur between unlike devices, as many users choose to use a desktop at their office, a laptop while in transit, and a tablet during meetings or briefings.

1.6.2 OCONUS RELIABILITY

OCONUS implementation may require specialized hardware, as replacement parts with FIPS requirements may not be available. All drives should be solid state and encrypted with no return to manufacturer as the maintenance option. Two full spares kits should be onsite and replenished regularly for all parts.

Wildflower recommends using a brand name converged architecture that has 1000 or more units in service. The vendor must have a full 24 x 7 x 365 support infrastructure with proactive monitoring and a history of reliability. Each node should only be run up to a point where the total utilization of the system can still be distributed to the remaining nodes if one of the nodes fails. No implementation should occur with less than 4 nodes regardless of the number of end users.

1.6.3 TOC OCONUS RELIABILITY

OCONUS TOC deployments need a FIPS based system to have the highest reliability. All systems must be fully redundant. The Army has a saying, “one is none and two is one.” This same philosophy should be applied when designing a system. All parts should be field replaceable. Parts depots must have at least 2 spares of all parts due to the limited accessibility of some locations. End user devices must not only be ruggedized; they must also be highly reliable with no moving parts that can fail. At times, tradeoffs will be made between the size of the form factor and reliability. For example, ruggedized tablet devices that are now in the field tend to be no larger than 7” due to ease of management and weight. Minimal screen sizes mean applications need to be aware of the device they are being viewed on and the output must be adaptable to the device. Additionally, these devices should have some capability to ensure mission critical applications are available when connectivity may not be.

Applications such as drone management and mesh networking are critical to a company’s success. If applications are deployed at the battalion level and relayed to the company level, there is a potential that more of a thin application environment can be built out. However, if a single battalion level deployment is conducted there is potential for loss so it is recommended that multiple vehicles/sites are enabled for full redundancy even in hostile OCONUS locations. Potentially a battalion level deployment could be fully redundant with its brigade level counterpart. Typically, the brigade deployments are more of a true data center environment and should be capable of handling the necessary redundant communications in the event of a failure. Ultimately, though it is debatable if a true thin client deployment would work in the field of battle so instead a hybrid approach is advisable.

1.6.4 COMMERCIAL OFF THE SHELF (COTS)

The Wildflower design considers commercial off the shelf devices (COTS) to minimize cost and optimize the ability to scale a deployment in a non-proprietary fashion. Some of the benefits COTS provides thin client deployments are illustrated below.

Mobile Device Management

Most, if not all, COTS devices can be managed with a number of Mobile Device Management solutions. This allows an administrator to wipe a device or disable a device in case of a situation that compromises the device. Tools like Airwatch and XenMobile amongst others provide standardized and approved methodologies that incorporate DoD requirements for encryption and secure wiping. Additionally, MDM solutions can incorporate APIs that can be integrated into common DoD applications that can provide Geo-

tracking, device status and end user activity. Specifically, Wildflower recommends Airwatch for Mobile Device Management due to its ease of administration and deployment.

Profile Management Tools

COTS devices can be easily administered using standard profile management tools and group policies which can be integrated into an Active Directory or LDAP configuration. Tools like RES & Appsense, which both have been implemented for profile management throughout the US Government, allow for easy resource deployment, allocation, and repurposing. Pairing these technologies to manage COTS devices allows organizations to quickly assign assets if existing assets become antiquated or non-functioning. Specifically, Wildflower recommends RES Software for Profile Management due to its "Playstore" like distribution of end user applications and ability to integrate with AD and LDAP.

Bandwidth Optimization

COTS devices are continually evolving and new models are continually being developed which allows for a rapid evolution of the end user device and interface. COTS manufacturers are continually evaluating methods to minimize bandwidth requirements and optimize mobile traffic which will benefit TOC deployment in areas where bandwidth is constrained. Coupling new bandwidth optimized devices with traditional receivers from companies like Citrix and VMware allow tactical designs to minimize throughput for field deployments. This is especially useful in front line tactical deployments where satellite and line of site communications are common. Wildflower takes an agnostic approach to end user devices in the design of our virtual solutions. The benefit of this approach is that they can quickly be modified to support a given device type or a simple receiver can be incorporated of them and either Citrix or View can be deployed at the centralized location.

Application Persistence without Communication

Tactical operations in the battlefield may require applications that need to be available whether or not bandwidth is available. COTS devices, such as traditional mobile devices, can provide application persistence without communicating to a centralized computing resource if necessary.

Thus, allowing some intelligence on the device versus a simple zero client provides the added benefit of allowing the device to manage strategic assets like drones or other devices. Wildflower is agnostic as to the device the end user would use to access the centralized resources. However, if the end user's device can have some intelligence (i.e. Droid OS or IOS) additional applications are available. It is highly recommended that these devices do however have an MDM solution installed on them to ensure proper destruction can occur in the event a device is compromised.

1.7 SCALABILITY

1.7.1 CONUS

1.7.1.1 Design for Under 500 Users

CONUS Data Center deployments for under 500 users are standard and easily managed with a converged infrastructure approach. The hyper-converged architecture allows scaling using nodes. This methodology provides the system the ability to scale all resources simultaneously. Scaling using nodes eliminates bottle necks that may be created by traditional infrastructures that scale resources autonomous of each other. Additionally, traditional architectures may have had different vendors for their server, storage and connectivity and thus interoperability and scalability issues occurred.

Wildflower's converged approach recognizes that an average user in a standard environment (CONUS or OCONUS) uses roughly 2 vCPU's, 3GB vRAM, and 5GB of vDisk. This scales to 1000 vCPUS, 1500 GB of vRAM and 2500 GB of vDisk with up to 500 users. Using linked clones and non-persistent desktops,

deployments typically experience a 10/1 dedupe ratio and easily fit up to 500 users in a single 2U system with 4 nodes.

Having a 2U cluster that can support the entire location saves DC space and power in comparison to traditional architectures that need a minimum of 6U to incorporate (3) 1 U servers and a SAN. (Switches are necessary in both configurations and not considered for space savings.) The specific bill of materials for a CONUS deployment of under 500 users will be outlined in the following section.

From a software architecture perspective, CONUS deployments typically rely on Active Directory as the source of truth for end user roles and definitions. With fewer than 500 users, group policies can be simplified and instead individual profiles may be created for key users insuring availability and resource optimization for key users. The benefit of this in small implementations is the granularity of the definitions that can be created for security and accessibility.

Additionally, thin client deployments at CONUS locations with under 500 users typically minimize inventory requirements for their depots since spares are readily available.

1.7.1.2 CONUS Design for Over 500 Users

While CONUS deployments using a hyper-converged model are efficient for space savings as well as power utilization, the scalability of the infrastructure is the true benefit of the solution. For example, a typical 500 user deployment for a normal office environment with a mixed workload that was mentioned above can be run in a 4 node cluster. Adding additional users is as simple as adding additional nodes to the cluster. For example, if a node in the architecture contains 2 processors each with at least 18 cores (36 total cores), 512 GB of RAM and 2.4TB of self-encrypting Solid State Disk and each user in an environment gets 2vCPUS, 3GB of vRAM and 5GB of vDisk then the average node can support approximately 130 users per node with about a 70% utilization of the node. Thus the initial 4 nodes that supported upwards of 500 users can simply be upgraded to an additional xx number of nodes depending on requirements. Given the scalability of the hyper-converged architecture, to increase from under 500 users to over 500 users, is accomplished by simply adding additional nodes. Please note, it is best practice to always add nodes in a minimum of 2 at a time to ensure full redundancy.

Since the secure site CONUS deployments have traditional data center amenities and applications, they may be able to access critical systems like Active Directory. Assuming there is access to AD or an HR system that contains employee job classifications, Role based profiles can be developed and deployed. For example, an E6 may get access to only a handful of applications that fit their role and position as a staff Sargent, but an 02 First Lieutenant may need additional access and applications. Additionally, a self-service portal can be implemented that allows end users to quickly access applications and gain approvals without having to access help desk resources.

In locations with more than 500 users there may be a need to run terminals or shift work. If this is the case, a thin client architecture with only a presentation layer available to end users in a non-persistent session may allow the sites to minimize the number of thin clients necessary. In this architecture, no user can see another user's information even when using the same device.

1.7.2 BILL OF MATERIALS FOR CONUS ARCHITECTURE

This is a vendor agnostic example of a BOM for multiple user scenarios in a CONUS facility that has existing firewall, switching and overall networking considerations in place. Note the ease of scaling throughout the design. Each Node supports up to 130 users in the design below.

# of Users	# of Nodes Needed	CPU/Cores/Node	Memory/Node	Storage/Node
Up to 500 Users*	4 Nodes	2 x E5-2695v4-36 cores	512 GB RAM	2.4 TB Self encrypting Solid State Drives
1000 Users	8 Nodes	Same	Same	Same
5000 Users	40 Nodes	Same	Same	Same

*Wildflower recommends a minimum of 4 nodes in CONUS locations for redundancy in the case of failure.

1.7.3 OCONUS DESIGN

1.7.3.1 Under 500 Users

If the OCONUS deployments are in a secure facility that has a data center, then the ability to use a mixed disk configuration may be applicable. However, with the price points of spinning disk and flash converging and the reliability of flash over spinning disk, Wildflower recommends an all flash node configuration. This is for three reasons. First, OCONUS deployments may not have easy access to maintenance facilities and therefore need more reliable equipment. Second, OCONUS facilities may not be as physically accessible for repair technicians to enter and bring spare parts. Lastly, OCONUS locations may not have the necessary power to always run multiple spinning disks. Given these reasons, OCONUS deployments should not only be all SSD based but Wildflower also recommends that all FRU's have spares in a parts depot onsite and secure support tunnels set up for remote support.

In an OCONUS deployment of 500 users or less, standardization is essential. All OCONUS installations regardless of size should have a standardized non-persistent desktop. This would allow minimal data to reside on a users' device, which in the case of compromise or support is ideal. If possible, authentication and presentation of this desktop happens remotely via a connection to a secure CONUS facility. If that is not possible and assets must remain at the OCONUS site, then a standardized desktop and standardized backend hardware is an ideal situation from a support perspective. Hyper-converged architectures that can scale are perfect for this approach. Having a system where all the resources scale simultaneously eliminates bottlenecks created by traditional infrastructures, which scale resources autonomous of each other.

As mentioned in the CONUS section of this document, Wildflower's approach assumes that each user in a standard environment will not have an extraordinary workload and therefore will use approximately 2 vCPU's, 3GB vRAM, and 5GB of vDisk. This scales to 1000 vCPUS, 1500 GB of vRAM and 2500 GB of vDisk with up to 500 users. This specific bill of materials will be outlined in the following section.

1.7.3.2 OCONUS Design for Over 500 Users

OCONUS deployments may have limited space and power challenges. A hyper-converged architecture accomplishes this easily and also provides the added benefit of scalability. A typical 500+ user deployment for an OCONUS location with a standard mixed workload can be run in a cluster. Adding users is as simple as adding additional nodes to the cluster that was originally designed for the sub500 user facility. If a node contains 2 processors each with at least 18 cores (36 total cores), 512 GB of RAM and 2.4TB of self-encrypting Solid State Disk and each user in an environment gets 2vCPUS, 3GB of vRAM and 5GB of vDisk, then the average node can support approximately 130 users per node with about a 70% utilization using a standard mixed workload. The scalability of the hyper-converged architecture allows an OCONUS facility to scale from under 500 users to over 500 users by adding a simple set of additional nodes.

If the OCONUS deployments have traditional data center amenities and applications, they may be able to allow the Thin Client deployment to access critical systems like Active Directory. Assuming there is access to AD, role-based profiles can be developed and deployed. Additionally, a self-service portal can be

implemented that allows end users in larger deployments to quickly access applications and gain approvals without having to access help desk resources.

One major difference in OCONUS deployments versus CONUS deployments is the design may incorporate more Thin/Xen Apps and fewer desktops so that a browser-based approach can provide end users access. Ideally, thin apps are presented to the end user from CONUS locations via secure transmission that is encrypted. This is especially relevant if OCONUS deployments issue COTS based devices, as these devices are prone to hacks and thus a well-known MDM and mobile threat prevention tool should be used. If possible all endpoints in OCONUS locations should be encrypted if any data resides on them.

For hardware, it is recommended that OCONUS locations have (2) spare nodes with the same configuration as the 4 primary nodes. These nodes may not be used but should be tested and available.

1.7.3.3 Bill of Materials for OCONUS Architecture

This is a vendor agnostic example of a BOM for a few options starting at under 500 users and increasing to upwards of 5000 users in an OCONUS facility that has existing firewall, switching, and overall networking considerations in place where each node supports 130 users.

# of Users	# of Nodes Needed	CPU/Cores/Node	Memory/Node	Storage/Node
Up to 500 Users*	6 Nodes	2 x E5-2695v4-36 cores	512 GB RAM	2.4 TB Self encrypting Solid State Drives
1000 Users	10 Nodes	Same	Same	Same
5000 Users	42 Nodes	Same	Same	Same

1.7.4 TACTICAL OPERATIONS CENTER

1.7.4.1 TOC OCONUS Design for Under 500 Users

If a thin client deployment is in a field of operation with less than 500 users, it most likely would be integrated at the battalion level versus a brigade level facility. Since the field deployments in a TOC environment need to be ruggedized and may experience vibration, dust, and heat, solid state disks are recommended again over spinning medium. Additional considerations should be given to the hardware architecture such as motherboard replacement and methodologies to wipe a device. TOC OCONUS deployments may not have easy access to maintenance facilities and therefore need more reliable equipment. Spares should be field deployable; all configurations will therefore include 2 spare nodes for a total of 6 nodes per 500 users deployed.

A TOC OCONUS deployment of 500 users or less brings new challenges into the fold above and beyond those that exist for a standard OCONUS deployment. All OCONUS installations whether under or over 500 users should have a standardized non-persistent desktop. This situation is identical in the case of a TOC deployment with one major exception. In the case of a standard OCONUS deployment it was not advisable to have processing on the end user's device. For a TOC deployment, mission critical applications may need to run regardless of connectivity, and thus the end user's device may need more intelligence. In this instance, using a receiver based technology (i.e. Citrix receiver), a browser based login or a thin app will make more sense for the sake of usability and security. Any application that is on the device should be stateless. This means the applications that are on the device may serve a purpose but do not retain any long term information. A good example would be drone based remotes. These remotes can manage a drone and receive real time feedback from the drone but do not store this information. Instead that information is relayed directly from the drone to the Brigade. Using this approach, should an end user's device be compromised, sensitive information will not be jeopardized.

In a TOC deployment where space is limited, power is at a premium, and ruggedizing is a must, hyper-converged architectures that can scale are perfect. Having a system where all the resources scale simultaneously and are replaced easily eliminates both the bottlenecks that may be created by traditional infrastructures and churn associated with parts failure.

Wildflower's approach assumes that each user in a standard environment will have an average workload and therefore will use approximately 2 vCPU's, 3GB vRAM, and 5GB of vDisk. In a TOC deployment, this ratio should hold true minus any graphically intensive applications. Should graphically intensive applications be required, it is recommended to incorporate additional GPUs into the design. Holistically though the 1000 vCPUS, 1500 GB of vRAM and 2500 GB of vDisk should still support up to 500 users. Linked clones and non-persistent desktops provide approximately a 10/1 dedupe ratio, and therefore a cluster of only 4 nodes in a hyper-converged infrastructure can support up to 500 users. Field deployments will each receive 2 additional preconfigured nodes that will be configured and active. The specific bill of materials will be outlined in the following section.

1.7.4.2 TOC OCONUS Design for Over 500 Users

TOC OCONUS deployments may have limited space and power challenges. A hyper-converged architecture tackles this challenge easily and also provides the added benefit of scalability. A typical 500+ user deployment for a TOC OCONUS location with a standard mixed workload as previously mentioned can simply scale the 6 node cluster by adding additional nodes. Wildflower always recommends adding nodes in a minimum of a pair. Adding additional users is as simple as adding additional nodes to the cluster that was originally designed for the sub500 user group.

If the TOC OCONUS deployments are in a brigade facility with data center amenities and applications, they may be able to allow the Thin Client deployment to access critical systems like Active Directory. Assuming there is access to AD, role-based profiles can be developed and deployed. It is not recommended to have a self-service portal in TOC field deployments. Soldiers should not be requesting applications on the fly. The major difference in TOC OCONUS deployments versus standard OCONUS deployments is the design may incorporate more of a thick client approach. This is necessary in case there is limited connectivity in the field. Ideally, thin apps are presented to the end user from CONUS locations via a secure transmission that is encrypted. This is especially relevant if TOC OCONUS deployments issue COTS based devices are used.

TOC OCONUS deployments will be most successful if implemented at the brigade level. With more than 500 users, bi-directional communication becomes challenging if only satellite communication is available. This design would be similar to a WAN which would be deployed via GCS and could be augmented using the vehicle based FCBC2 systems if enough vehicles can be outfitted with proper equipment. If the communication is within the brigade's facility a more traditional terminal based approach could occur with non-persistent users logging in and authenticating via AD or LDAP. Field based applications on the other hand such as DSGS-A and PALANTIR could reside on the soldier's device but only receive information via satellite or FCBC2 communications. Ultimately, refresh rate, data transfer rate and field conditions will affect scale. One way to overcome some issues in the battlefield is to use meshed networks and peering. This methodology allows each device to be a node in a network that can continually redefine itself in the event of failure. If this system incorporates biometrics such as heart rate or movement, then a continuous feedback loop can be created allowing both the brigade to know the status of the individual soldier as well as the individual soldier having the capability to know the status of his team and auto-negotiate the connectivity of his device into a meshed network. Anytime this type of network is created though we must be careful to incorporate a methodology to wipe a device and disable it in the event of capture.

The portability of the specific mobile device must be taken into consideration. Weight and space are two major factors that affect all battlefield equipment. Minimizing both of these is essential. Many COTS based devices today are focusing on optimizing both of these and as these devices evolve and ruggedize the DoD should consider evaluating their use.

From the centralized resource perspective, a hyper-converged deployment fits well into the field deployable architecture models. The modularity and ability to quickly add nodes allows a tightly integrated approach to scalability. Specifically, when growing from under 500 users to over 500 users Wildflower recommends simply adding additional nodes. At a minimum, nodes should be added in pairs. Given the fact that the initial implementation already had 6 nodes and that each can support upwards of 130 people, Wildflower recommends simply adding an additional 1 node per 130 users for the sake of redundancy up to 1000 end users. Thus, a 1000 user deployment will have 10 nodes configured and active.

1.7.4.3 Bill of Materials for TOC OCONUS Architecture

This is a vendor agnostic example of a BOM for both an under 500 user implementation as well as up to 5000 users in a field deployed TOC OCONUS facility that has existing firewall, switching, and overall networking considerations in place:

# of Users	# of Nodes Needed	CPU/Cores/Node	Memory/Node	Storage/Node
Up to 500 Users*	6 Nodes	2 x E5-2695v4-36 cores	512 GB RAM	2.4 TB Self-Encrypting Solid State Drives
1000 Users	10 Nodes	Same	Same	Same
5000 Users	42 Nodes	Same	Same	Same

2. SUMMARY

Wildflower recommends a highly available architecture using a node-based, hyper-converged deployment that is highly scalable. This same architecture can be deployed CONUS, OCONUS and in TOC situations.

This system will start at a minimum of 4 nodes with: 2 x E5-2695v4- 36 cores, 512 GB RAM & 2.4 TB Self Encrypting Solid State Drives.

In OCONUS and OCONUS TOC deployments, an N + 2 design is recommended and thus 6 nodes are initially used for an under 500-person deployment. In CONUS deployments a simple 4 node architecture will support up to 500 users without the need for an onsite spares kit, whereas an OCONUS deployment will use a 6 node architecture for under 500 users and will require an onsite spare parts kit for all FRUs.

This solution will support 130 users per node and up to 4 nodes can be supported in a 2U chassis. Power and space savings using this architecture allow it to be deployed in both data center and field level situations. Additional design considerations will be reviewed on a case by case basis as use cases vary.

We modernize critical information technology so you can achieve mission goals, reduce operating costs, and improve performance. By delivering intelligent and cost-effective IT systems we drive your business and mission forward. We incorporate existing capabilities to develop tailored, leading-edge solutions while ensuring you're on course to meet the demands of the future.