Standardized Configurations

Requirement - Develop and maintain two thin-client user configurations for organizations consisting of less than 500 users and those with more than 500 users. (Note: Actual thin client configurations shall only be submitted post-award after CIO/G6 policy memorandum and NETCOM TECHCON have been finalized.) (Proposal shall address the following, in light of SOW paragraph 5.7.2)

Dell maintains active partnerships with all types of thin client computing solutions to include Citrix (ICA), Micrsoft (RDP), and several OS streaming technologies. As referenced in the recent Netcom Technical Authority Draft document addressing the Army enterprise thin client desktop architecture standardization, Dell has a proven track record of providing any of the solutions the Army requires. As a result of our flexibility, we believe we can offer the army the best solution to fit their very dynamic needs.

Reality dictates that all forts are not the same, and technology approaches will differ dramatically. Consistency is always the stated objective, although the criticality of the mission is foremost on our minds. As a result Dell would like to offer in response to this questioning, forefront technology that will meet the Army's needs and flexibility. In an effort to be compliant to your requirements, we will illustrate the following: Citrix and OS streaming. We will carefully define both of these environments, comparisons, and illustrations supporting environments with less than 500 users, as well as environments with more than 500 users.

Some key factors of success for the Armys' implementation of thin computing will be the following:

- Scalability: effective scalability and management of back-office costing to include servers, storage, labor rates
- Connectivity: access for off-line / disconnected / deployed users / networks in a failed state
- Flexibility: application flexibility.
- Transition: transitional phase from thick clients to thin clients. The ability to leverage current investments.

There is a direct relationship to the number of users that will be supported, to the servers and are required in the farm. Of course the number of servers, and the complexity of the servers, will dictate your labor costs in the reality of whether or not the solution will succeed or fail.

Scalability: One of the primary differences in technology is the location of the processing power. For example: any standard ICA/RDP thin client environment performs all the processing with in the server farm. Essentially this defines fewer users per processor, and to increase number of servers in the server farm to support your user base. In comparison, OS streaming technologies utilized processing power at the client, instead of the server farm, hence fewer servers are required with less complexity. If the business objective is to remove the complexity of image management, patch management, and security at the desktop, either one of the solutions/technologies are applicable. The Dell solution will illustrate the cost delta / total cost of ownership for both technologies. Dell is positioned to offer either of these technologies as they apply to the Army's environment.

Connectivity: Utilization of an embedded operating system at the thin client will provide users with limited functionality in the event the server farm is off-line. ICA/RDP environments are dependent on functionality of the server farm, and have many failsafe technologies in place to assure functionality. The challenge within the industry has not necessarily been the server farm, but the connectivity to the server farm. Reality: the Army will face similar challenges. OS streaming environments are also challenged by the inability to connect to the server farm. While maintaining the benefits of security, image management, and patch management in OS streaming thin client, if equipped with a hard drive can provide functionality of the local desktop while the network is in a failed state. The Dell solution can be molded to meet any connectivity challenge in your environment.

Flexibility: Application flexibility is always been the largest challenge to thin client environments. Point in case: as the department of the Navy began its consolidation efforts almost 5 years ago utilizing the NMCI contract, application integration was soon discovered to be their number one challenge. With well over 10,000 unique applications in the environment, they soon discovered that the applications were not supported in a server process environment. Not uncommon with industry standards, application integration into a standard ICA/RDP environment is less than 80%. The Dell solutions listed in this response allow the army to take advantage of a ICA/RDP technologies for the majority, while still utilizing similar technologies for the minority to accomplish business objectives. The Dell solution is flexible and will allow Dell and the Army to meet all the needs of local commands.

Transition: Transition from a thick client environment to a thin client environment will be imperative to the success of this project. Commanders will want to see a clear line between their current investment and the investment path of thin computing. Part of the marketing of this solution will be dependent upon the commanders ability to utilize current resources to move into a new technology. The resources reference will include current staff, training requirements, servers, desktops, and laptops. One of the stated objectives for the migration to thin computing is “TCO over the system life cycle is lower due to a longer client service life and centralized administration, even though implementation of thin client architectures requires an increased initial cost for servers and storage devices" (Army enterprise thin client desktop architecture standardization, Netcom technical Authority implementation memorandum, 2.0.b). The Dell solution offers a compelling transitional plan: the ability to replace expired workstations with an ICA/RDP thin client, while leveraging OS streaming technology to extend the life cycle of desktops and laptops that are currently in the environment.

ICA/RDP environment:

Servers: The standard calculation for processor utilization in this environment is 25 users per processor. Both manufacturers claim that new technology is enabling them to extend to 50 users per processor. If we assume there's truth in the 50 users per processor, utilizing dual processing servers equipped with dual core processors, and low usage for Microsoft office applications only, a standard dual processing server will handle approximately 100 users. Scalability of the solution has several factors that need to be considered:

- using the 50 users per processor model, and standardizing on a dual processing server, this calculation is very straightforward

- the type of users and the applications being served will define the cost of the server, and the number of users supported per processor. This calculation is often completed after the environment has been installed and metrics measured.

Below is an example of a 100 user environment.
Blades: The example below illustrates the use of blade servers as the thin client servers. The 100 user environment we will use two blades. This solution by default create single-point-of failure, which is the blade enclosure. Some environments choose to have two blade enclosures to support the servers, and eliminating the single-point-of failure. Either configuration is supported by Dell. The initial investment is greater for the blade servers when the quantity is less than 500 users (or five blades) in comparison to standard 1U servers, although as you continue to build out the blade infrastructure there is a cost savings in most cases.
Ethernet Redundancy: Note also on the drawing below the redundancy and ethernet switch fabric from the blade enclosure to the site ethernet backbone. Although illustrated here are ethernet switches, they could easily be passed through switches.
Storage: The storage illustrated here is fiber channel connected, although the storage could be iSCSI connected utilizing the ethernet switches and VLAN technology for performance.
Thin Client: Dell offers several thin client devices, from scaled down desktops to true "bricks”.

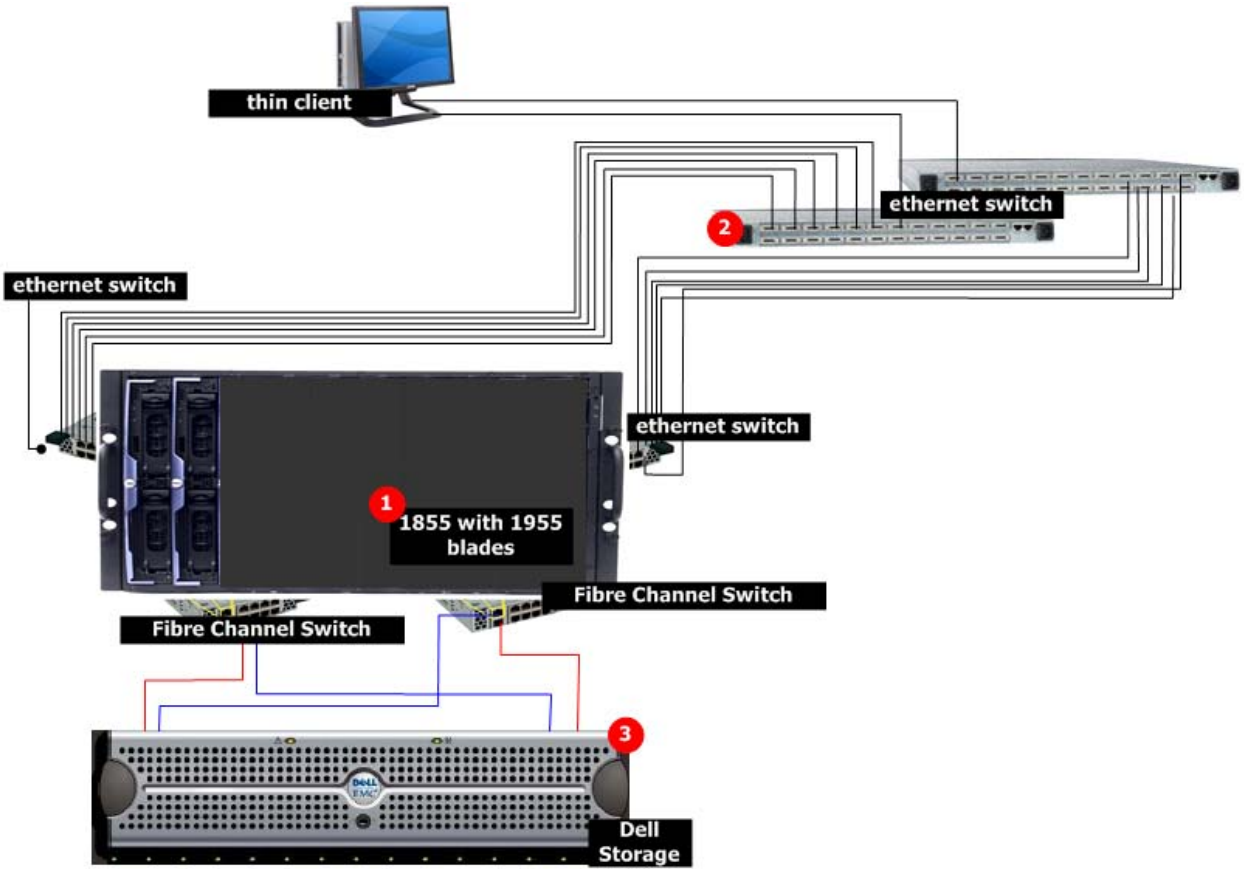

Below is an example of a 500 user environment.

Blades: The example below illustrates the use of blade servers as the thin client servers. The 500 user environment we will use five blades. This solution by default create single-point-of failure, which is the blade enclosure. Some environments choose to have two blade enclosures to support the servers, and eliminating the single-point-of failure. Either configuration is supported by Dell. The initial investment is greater for the blade servers when the quantity is less than 500 users (or five blades) in comparison to standard 1U servers, although as you continue to build out the blade infrastructure there is a cost savings in most cases.

Ethernet Redundancy: Note also on the drawing below the redundancy and ethernet switch fabric from the blade enclosure to the site ethernet backbone. Although illustrated here are ethernet switches, they could easily be passed through switches.

Storage: The storage illustrated here is fiber channel connected, although the storage could be iSCSI connected utilizing the ethernet switches and VLAN technology for performance.

Thin Client: Dell offers several thin client devices, from scaled down desktops to true "bricks".

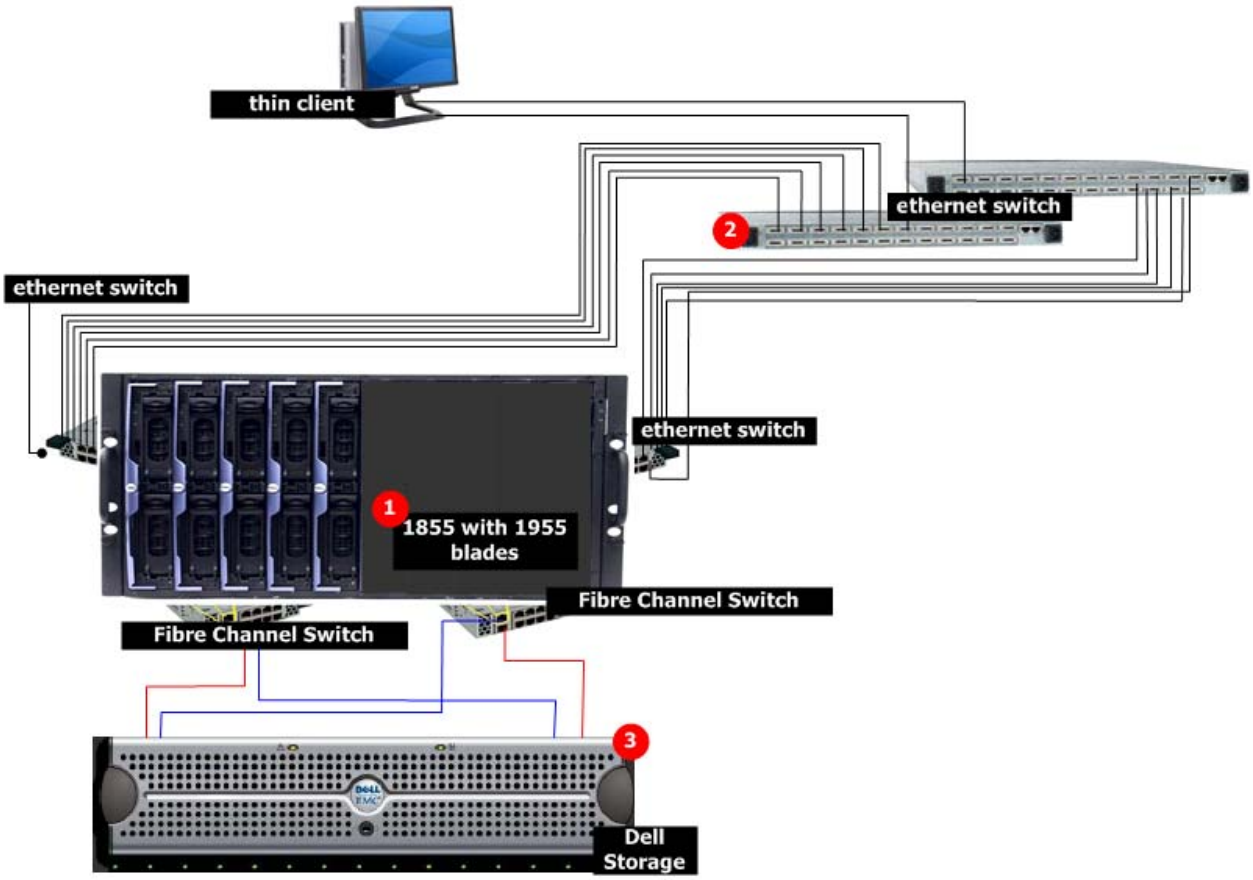

Below is an example of a 1500 user environment.

Blades: The example below illustrates the use of blade servers as the thin client servers. The 1500 user environment we will use fifteen blades. This solution illustrates the redundancy available in the blade architecture. Below you will discover that blade enclosure number one has a total of seven servers, while blade enclosure number two has a total of eight servers.

Ethernet Redundancy: Note also on the drawing below the redundancy and ethernet switch fabric from the blade enclosure to the site ethernet backbone. Although illustrated here are ethernet switches, they could easily be passed through switches.

Storage: The storage illustrated here is fiber channel connected, although the storage could be iSCSI connected utilizing the ethernet switches and VLAN technology for performance. In addition we notice an additional storage pod it has been added to our storage area network. The storage available for the blade servers can range from high performance fiber channel, low performance fiber channel, and serial attached ATA. Within the designated shelf for that storage you have the freedom to select which storage will serve your users most appropriately. All of the storage is fully redundant requiring only one storage array head to support up to 480 drives ranging from 73 GB to 500 GB per drive.

Thin Client: Dell offers several thin client devices, from scaled down desktops to true "bricks".

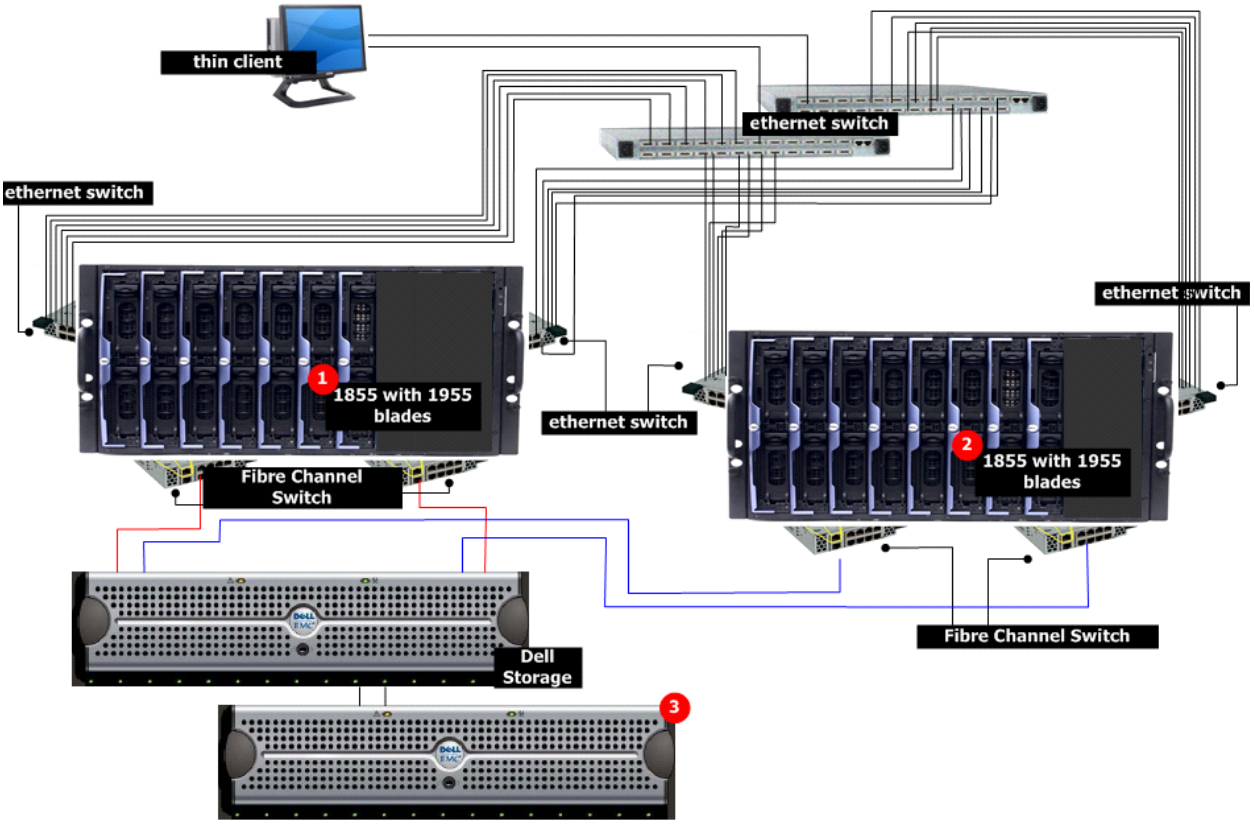

Storage: A storage foundation will be required from the initial implementation of your thin computing environment. Managing the initial investment will have a large impact on the success and adoption of thin client computing among stakeholders. Some steps need to be taken in establishing the type of storage that is required:

- how much space per user, based on quotas will be allocated

- what type of storage will be required? Maintaining the flexibility between high performance fiber channel, low performance fiber channel and serial ATA will be in the best interest of all parties.

- will the storage technology continue to perform optimally as more users are added? As bottlenecks are detected, in a non-disruptive manner, can performance be increased on-the-fly. Avoiding virtualized storage will allow you to adjust performance as required.

- Is the storage technology scalable? Can I start small and build?

Storage will be consistent regardless of the type of thin client computing environment utilizes.

Labor pool: The labor pool requirement to support and ICA/RDP environment is typically twice the cost of that for an OS streaming environment. Differences may be noteworthy:

- the location of processing: one of the primary factors in an accelerated cost for labor to support the ICA/RDP environment is the location of the computing power. Within an ICA/RDP environment all of the computing power for the thin client is performed in the server farm. Modifications and management of applications in the server farm can become complex. Many applications will require special attention in order to perform optimally in ICA/RDP client environment.

- the number of servers: as illustrated above, to support 100 users you will add a single dual processing server. The math is very basic: an increased number of servers calculates out to being increased number of servers that require maintenance, patch management, and server licensing.

Licensing: Server licensing and application licensing are considerations in the ICA/RDP environments. The client access license for an RDP environment is often included in the Army Microsoft licensing agreement. Although the licensing for the ICA environment is not a simple. Licensee will be required per desktop ranging from a price of $200-$500 dependent upon the feature set used. In addition, licensing has been required on the server is well for the ICA software. This does not take into consideration the server licensing costs that go to Microsoft (Windows 2003 for example).

OS Streaming Environment:

Servers: The standard calculation for processor utilization in this environment is 100 users per processor. If we assume there's truth in the 100 users per processor, utilizing dual processing servers equipped with dual core processors, and low usage for Microsoft office applications only, a standard dual processing server will handle approximately 200 users. Scalability of the solution has several factors that need to be considered:

- using the 100 users per processor model, and standardizing on a dual processing server, this calculation is very straightforward

- the type of users and the applications being served will define the cost of the server, and the number of users supported per processor. The type of users in the OS streaming environment range from shared to private. Dependent upon the adoption and type of client, numbers will vary.

Below is an example of a 100 user environment.
Blades: The example below illustrates the use of blade servers as the thin client servers. The 100 user environment we will use two blades for redundancy of the environment. This solution by default create single-point-of failure, which is the blade enclosure. Some environments choose to have two blade enclosures to support the servers, and eliminating the single-point-of failure. Either configuration is supported by Dell. The initial investment is greater for the blade servers when the quantity is less than 500 users (or five blades) in comparison to standard 1U servers, although as you continue to build out the blade infrastructure there is a cost savings in most cases.
Ethernet Redundancy: Note also on the drawing below the redundancy and ethernet switch fabric from the blade enclosure to the site ethernet backbone. Although illustrated here are ethernet switches, they could easily be passed through switches.
Storage: The storage illustrated here is fiber channel connected, although the storage could be iSCSI connected utilizing the ethernet switches and VLAN technology for performance.
Thin Client: Dell offers several thin client devices, from scaled down desktops to true "bricks". Application of standard desktops is also available with this model in a disk or diskless environment.

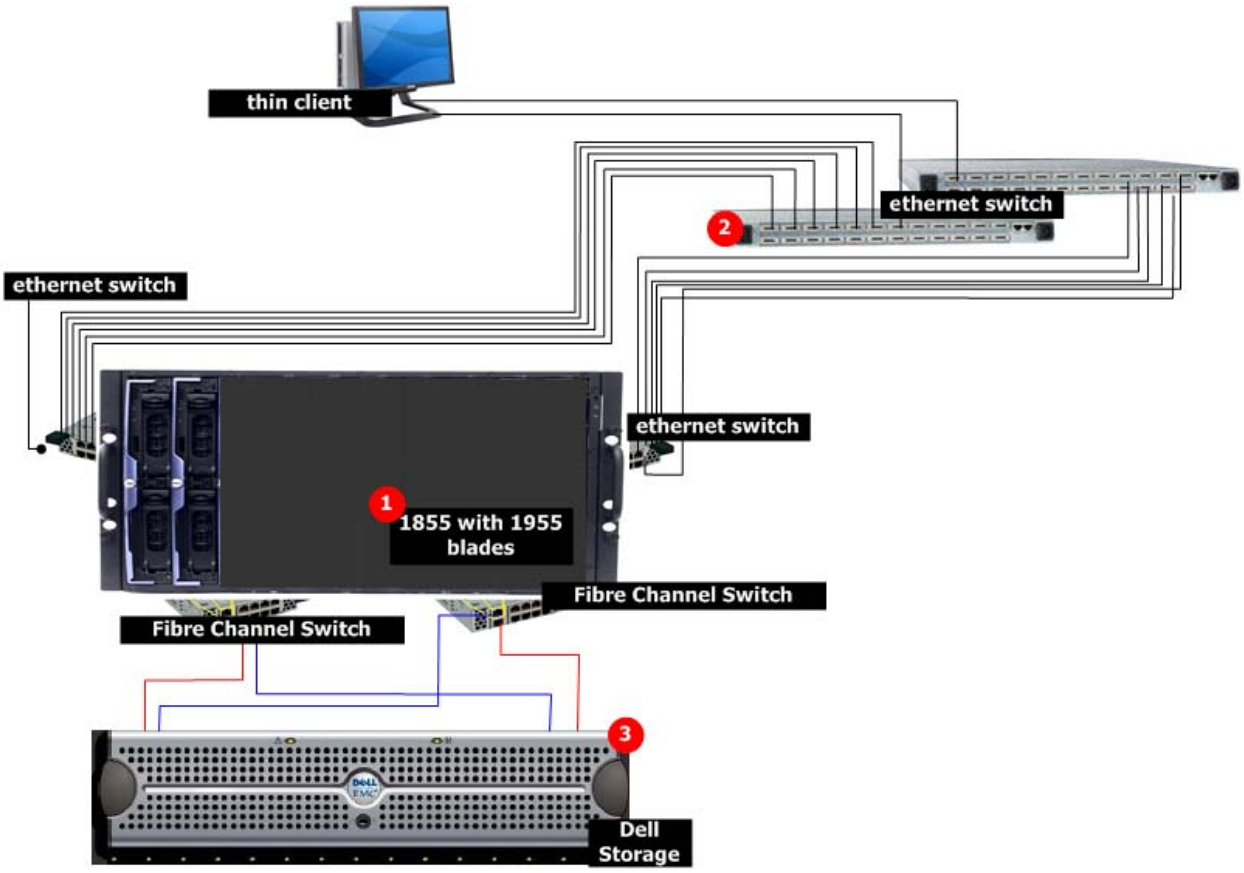

Below is an example of a 500 user environment.

Blades: The example below illustrates the use of blade servers as the thin client servers. The 500 user environment we will use three blades. This solution by default create single-point-of failure, which is the blade enclosure. Some environments choose to have two blade enclosures to support the servers, and eliminating the single-point-of failure. Either configuration is supported by Dell. The initial investment is greater for the blade servers when the quantity is less than 500 users (or five blades) in comparison to standard 1U servers, although as you continue to build out the blade infrastructure there is a cost savings in most cases.

Ethernet Redundancy: Note also on the drawing below the redundancy and ethernet switch fabric from the blade enclosure to the site ethernet backbone. Although illustrated here are ethernet switches, they could easily be passed through switches.

Storage: The storage illustrated here is fiber channel connected, although the storage could be iSCSI connected utilizing the ethernet switches and VLAN technology for performance.

Thin Client: Dell offers several thin client devices, from scaled down desktops to true "bricks".

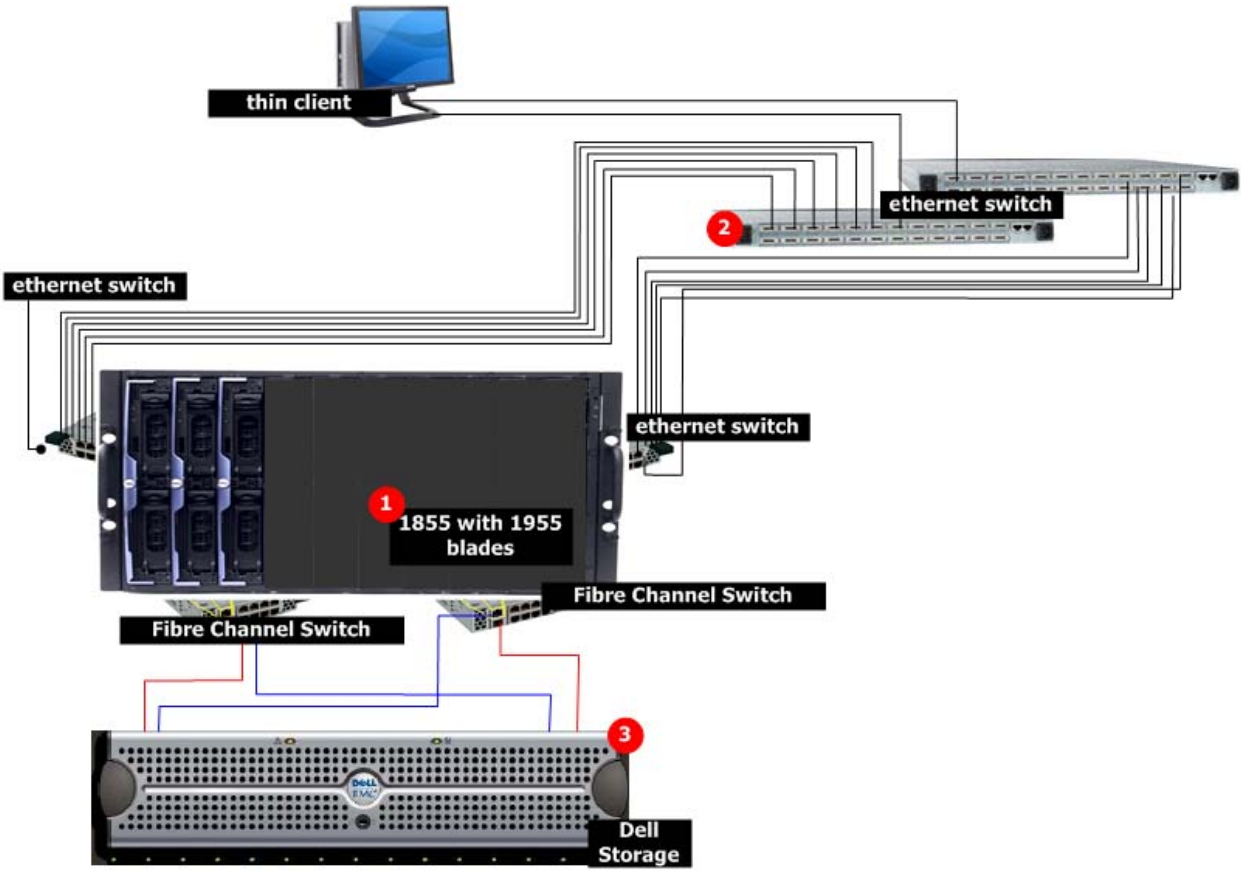

Below is an example of a 1500 user environment.
Blades: The example below illustrates the use of blade servers as the thin client servers. The 1500 user environment we will use eight blades.
Ethernet Redundancy: Note also on the drawing below the redundancy and ethernet switch fabric from the blade enclosure to the site ethernet backbone. Although illustrated here are ethernet switches, they could easily be passed through switches.
Storage: The storage illustrated here is fiber channel connected, although the storage could be iSCSI connected utilizing the ethernet switches and VLAN technology for performance. In addition we notice an additional storage pod it has been added to our storage area network. The storage available for the blade servers can range from high performance fiber channel, low performance fiber channel, and serial attached ATA. Within the designated shelf for that storage you have the freedom to select which storage will serve your users most appropriately. All of the storage is fully redundant requiring only one storage array head to support up to 480 drives ranging from 73 GB to 500 GB per drive.
Thin Client: Dell offers several thin client devices, from scaled down desktops to true "bricks".

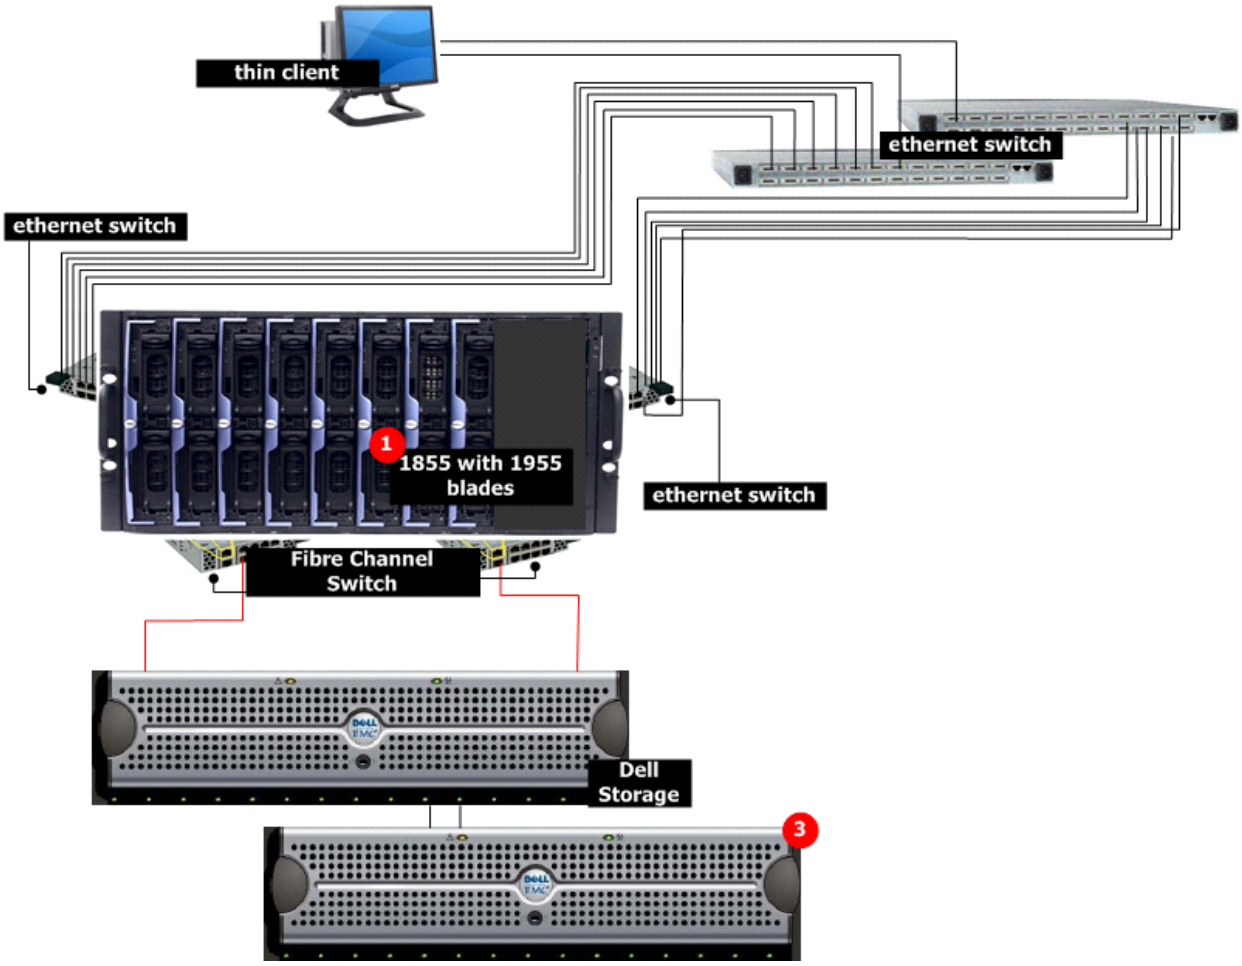

Storage: A storage foundation will be required from the initial implementation of your thin computing environment. Managing the initial investment will have a large impact on the success and adoption of thin client computing among stakeholders. Some steps need to be taken in establishing the type of storage that is required:

- how much space per user, based on quotas will be allocated

- what type of storage will be required? Maintaining the flexibility between high performance fiber channel, low performance fiber channel and serial ATA will be in the best interest of all parties.

- will the storage technology continue to perform optimally as more users are added? As bottlenecks are detected, in a non-disruptive manner, can performance be increased on-the-fly. Avoiding virtualized storage will allow you to adjust performance as required.

- Is the storage technology scalable? Can I start small and build?

Storage will be consistent regardless of the type of thin client computing environment utilizes.

Labor pool: The labor pool requirement to support and OS streaming environment is typically half the cost of that for an ICA/RDP environment. Differences may be noteworthy:

- the location of processing: Within an OS streaming environment all of the computing power for the thin client is performed at the client (with the local hard drive or without a local hard drive). Applications and integration with peripheral devices is dramatically simplified, as all of the processes are being driven at the local client, not in a server farm.

- the number of servers: as illustrated above, to support 100 users you will add a single dual

processing server. The math is very basic: an increased number of servers calculates out to being increased number of servers that require maintenance, patch management, and server licensing.

Licensing: Server licensing and application licensing are considerations in the OS streaming environments. Licensee will be required per desktop ranging from a price of $200-$500 dependent upon the feature set used. Always streaming technologies rarely charge of server charge for the use of their software, and because the number of users supported in an OS streaming environment per server is half that of an ICA such RDP environment, potential savings are notable.

Resources and Timeline

Identify resources and timeline to facilitate the development and maintenance of the proposed configurations;

Application integration is going to be a challenge for both technologies, as the facilities will have to understand exactly what is running on the user's desktop so that the services can be provided in a centralized thin computing environment. The number one challenges that most users do not comprehend completely the applications they use to perform their tasks. Some of the timelines associated with collecting the application integration piece are as follows:

-- **site survey/discovery** of all user desktops for applications in the environment. This timeline can be dramatically reduced if the utilization of SMS or similar technologies is applied, although the reality is that many sites are not fully functional with this level of technology. As a result collection in site surveys will have to be done using a network push to collect the information, or individually visiting each desktop. Estimated timeline for centralized collection of this information will range dependent upon the size of the environment, the complexity of the network infrastructure, and active directory policies and may restrict collection of certain data.

> **Thin Computing Choice:** Resources and timelines will be required for both ICA/RDP and OS streaming implementations.
>
> **Estimated Time:** to develop an automated solution to collect data from a local area network: estimated 120 man-hours. Complications that will arise as a result of a wide area network, network access, user rights, etc.: estimated 120 man-hours. Manual collection of data: 15 minutes per desktop (it is important to note that automated collection will most likely only result in 80 to 90% of all desktops and the manual collection may need to be used for zero to 20% of your desktops). Estimating timelines is very difficult without additional information, no guarantee of timeline is implied.

-- **creating the user profiles.** Once all of the data is been collected from the network, a user profile and network profile will need to be generated to define which applications will be available to the users in a thin computing environment. Often gathering of the information and utilization of technology is the easy part ..... soliciting conformity from all the end users and stakeholders will create the largest challenges. Resources and timelines will be required for both ICA/RDP and OS streaming implementations.

> **Thin Computing Choice:** Resources and timelines will be required for both ICA/RDP and OS streaming implementations.
>
> **Estimated Time:** development of the user and network profiles: estimated 60 man-hours. Solicitation and development of a “standard” desktop environment: the estimated timeline is impossible due to the number of variables. Estimating timelines is very difficult without additional information, no guarantee of timeline is implied.

-- **evaluating the correct thin computing technology to be applied**. Testing and validation are conducted during this stage. All applications that have been deemed as critical will have to be tested and validated to confirm the type of thin client computing technology that will best serve the end-user. A service level agreement will have to be established between the stakeholders and those developing the technology. It is always suggested this become a phased approach, although will be one of the quickest

ways to defeat this initiative if not executed properly and timely.

Thin Computing Choice: This primarily impacts the ICA/RDP environment. Remember that all of the applications in an ICA/RDP environment are run from the server farm, in comparison to an OS streaming environment that runs all the applications in local RAM.

Estimated Time: development time will have to be broken down by the number of applications the environment will be using. If the application does not require redevelopment, code modification: estimated 20-40 hours per application. Estimating timelines is very difficult without additional information, no guarantee of timeline is implied.

-- **migration and consolidation of user data**. Once the primary images been created, all the applications have been designed and optimized for the thin computing technology you have selected, data migration will need to begin. As with the site survey/discovery of the user's desktop for applications, discovering where all of the data is on a user's hard drive will be time-consuming. The use of automated tools will greatly enhance your ability to address this issue. One to define the data that will be migrated, access to appropriate storage will be the next challenge. Many of the clients it will be converted are at remote sites, and may not have the luxury of being close to an optimal ethernet environment so the migration of data has little impact on daily operations. Alternative methods of migration of user data from the thick client to the centralized storage will need to be thought of carefully, and will create a lot of work for the information technologists.

Thin Computing Choice: Resources and timelines will be required for both ICA/RDP and OS streaming implementations.

Estimated Time: the timeline variables will be dependent upon the network speed between the client and the location of the storage device. The type of storage selected will also have an impact on the migration of the user's data. For example: if I implement a storage device that only supports virtualized storage/shared storage where I can not adjust the read/write cache levels per raid group, and I am writing to slow ATA drives with the connectivity to a single 1 Gb port, my migration will be considerably slower in connecting to a server attached to high-speed fiber channel drives, optimized cache levels for read and write activities, traveling on a fully arbitrated 4 Gb storage area network. Estimating timelines is very difficult without additional information, no guarantee of timeline is implied.

-- **utilization of storage.** The type of storage purchased will have a direct impact on the scalability of the thin client environment. For example, SCSI or SAS drive technology is the most cost efficient technology for storage that can be purchased, although the solution is not very scalable. SCSI/SAS has a very attractive entry point when calculating the cost per megabyte/gigabyte, although as a storage requirement grows... and it will in the thin client environment, soon it will be realized that the cost per megabyte/gigabyte has increased dramatically. Discussed is only the hard cost associated with the hardware, but the real cost is soon discovered in the amount of time it takes to manage/and the lack of flexibility SCSI/SAS technologies offer.

This example applies also to network-attached storage, or any storage device that does not give the user (the Army) maximum utilization of disk space, as well as the choice in drive type to use. For example: if site anywhere purchase 10 TB of storage to support the thin client initiative, configure the storage using raid five, only to discover that the network attached storage required 30 to 40% overhead before raid was implemented.... in other words the 10 TB of storage that was purchased, once configured with no raid is actually 6 to 7 TB of usable space. Now add an additional 25% for raid operations, and you usable space becomes an estimated 4 to 5 TB. Hence, if your requirement was for 10 TB of usable space be centrally you will have to purchase 20 TB of raw storage. Many storage providers have less than 3% storage overhead to use their devices. Research and application of the proper storage device will impact timelines and cost.

Thin Computing Choice: Resources and timelines will be required for both ICA/RDP and OS streaming implementations.

Estimated Time: the timeline variables will be dependent upon the network speed between the

client and the location of the storage device. The type of storage selected will also have an impact on the migration of the user's data. For example: if I implement a storage device that only supports virtualized storage/shared storage where I can not adjust the read/write cache levels per raid group, and I am writing to slow ATA drives with the connectivity to a single 1 Gb port, my migration will be considerably slower in connecting to a server attached to high-speed fiber channel drives, optimized cache levels for read and write activities, traveling on a fully arbitrated 4 Gb storage area network. Estimating timelines is very difficult without additional information, no guarantee of timeline is implied.

-- **exception management:** as new applications are in reduced to the environment, and exception management process must be put in place. In an ICA/RDP environment this is absolutely critical to success, whereas in an OS streaming environment, because it processes are run at the local desktop, application integration is typically a non-issue.

Thin Computing Choice: This primarily impacts the ICA/RDP environment. Remember that all of the applications in an ICA/RDP environment are run from the server farm, in comparison to an OS streaming environment that runs all the applications in local RAM.

Estimated Time: development time will have to be broken down by the number of applications the environment will be using. If the application does not require redevelopment, code modification: estimated 20-40 hours per application. Estimating timelines is very difficult without additional information, no guarantee of timeline is implied.

Implementation and Promotion of Thin Client Solutions

Describe the extent of engineering and marketing efforts required to implement and promote thin client solutions over the life of the contract;

The extent of engineering required for the ICA/RDP and or OS streaming solutions is referenced above. The marketing effort required to promote the use of thin client technology is going to be dependent on some of the choices made for the following:

-- Exception management. What will be the timeline required for a user to have a new application added to the thin client environment? Will there be available resources in the future in order to satisfy these requests? The ICA/RDP environment will require additional marketing efforts, as the technology naturally creates challenges for application integration. The marketing effort will have to convince the end-user that a current application already approved and functional in the ICA/RDP environment will satisfy their request or the end-user becomes an exception.

-- Off-line or disconnected use. Thin computing technology will be best used when it is discovered that the user can utilize the same technology while connected to the network or traveling. The adoption of two different types of technology to satisfy these requests will create an increased workload for the information technologists.

The implementation of an ICA/RDP environment in comparison to an OS streaming environment will have a direct impact on the success of this solution. End users will react negatively if they feel they're losing functionality by going to a thin client environment. Because of the nature and influence of the stakeholders of those users, exceptions will need to be granted, and the popularity of the thin computing environment may potentially be minimal.

Thin Client Solution Benefits

Describe how your thin client solutions will be of benefit to the government over the life of the contract.

Business needs often are the driving force for many IT deployments. Some of the dynamics include:

- ✓ Aligning IT with business

- ✓ Standardizing systems management and desktop functionality while supporting flexibility for exceptions
- ✓ Maintaining stability via patch management and desktop management
- ✓ Ensuring a secure environment

Regulatory compliance continues to be initiated to provide protection from financial scandals, privacy protection, and information theft. Regulatory requirements that include Sarbanes-Oxley Act, Health Insurance Portability and Accountability Act (HIPAA), Gramm-Leach-Bliley Act (GLBA) and European Privacy directives persistently drive the need to streamline IT execution and management.

The infamous SQL Slammer worm that disrupted internet services in 2003 reiterated the requirement for proactive patch management at the server and desktop platforms, sky rocketing IT maintenance costs to capital impacting levels.

Many CIOs are faced with the reality… "Do more with less – Leverage IT to balance the bottom-line". As federal IT budgets soar out of control (Treasury Department 2003 – 2.6 Billion, Department of Homeland Security in 2002 receiving 45 Million) and many projects in failed status (according to CIO Magazine, May 15, 2006, Federal IT - federal agencies have lost over 1.6 Billion dollars from 2003 to 2007 and at risk an additional 11.4 Billion on projects that exceed 4 years behind schedule) the federal government is looking for alternative solutions to contain IT budgets.

Corporate America faces similar challenges with failed compliance to Federal Regulations and identity theft. According to CSO Perspectives Conference hosted in early 2006, David McIntyre, CEO of TriWest reported that 53 million identities have been stolen to date, costing companies between $40,000.00 and $92,000.00 per victim. Not to mention the 20% lost consumer base and 5% of consumers hiring lawyers. With record awards (recently ChoicePoint fined $15 million for a security breach) crippling corporations.

Containment of cost can be leveraged throughout the IT organization, to include computer resources. Differences exist between FAT and THIN clients. The FAT client is defined as a computer that hosts the operating system and applications locally, in comparison to THIN clients that receive their operating system and applications from a server in the environment. Cost containment is divided into the following (specifics on the calculations used for the statements below can be found in the appendix B: cost-containment calculations):

1. Hardware Refresh- Hardware is replaced, typically every three to five years.

- ✓ If your environment requires the flexibility to go from fat to thin-client, your hardware platform will be the same. In this model the cost delta per client is $1.29, favoring a thin client environment for cost savings.
- ✓ If your environment does not require the flexibility, and your new hardware acquisition can focus on thin client devices which are typically 50% less initial investment cost, the cost delta between fat client and thin client is $401.29 per client, favoring a thin client environment for cost savings. Using a 300 user environment as our standard model, the cost savings to replace 300 client machines as a hardware acquisition is $120,387.00 when favoring a thin client environment.

2. In-place migration/upgrade- Typically an operating system upgrade only. Existing hardware remains in place.

- ✓ The in-place migration/upgrade does not require a backup of users data, as the data exists on a centralized server – does not require a re-imaging of the client machine, as the image exists on a centralized server – and requires very little user productivity loss, as the task only requires a changing of the client machine

to point to the new image. Using a 300 user environment as our standard pricing model, the cost savings are in favor of the thin client environment, saving $82.50 per client or $24,750.

3. Maintenance- Daily/Annual maintenance associated with the computer. Computer failure rate of 15% annually is standard indicating that the failed computer will require new imaging, restoration of data from backup media, and user loss of productivity. Maintenance calculations are broken into two primary areas: failed computers that require a new image, and patch management.

- ✓ Utilizing a 15% average number for computer failures that require a reload of all applications and operating system, assuming a 300 user environment, dispersing the cost of the failed 15% as an individual cost per client across the 300 user environment. To restore the 15% failed fat client computers, it will cost $27.54 per client or $8,262 per year or $24,786 for three years. Continue to apply the variables as stated above, to restore the 15% failed thin client computers as a comparison, you will cost $1.07 per client or $321 per year or $963 for three years. When comparing fat client technology and thin client technology in an effort to restore the 15% of failed computers, in annual cost savings of $7,941 or $23,823 is available utilizing thin client technologies.
- ✓ Patch management is another cost required to run an information technology environment. Assuming we are working with a 300 client environment, to support a fat client environment with an average of 52 patches per year will create a cost burden of $111,067.96. When applying the same variables as stated above the thin client environment will create an annual cost burden of $17,778.87. The cost advantage for choosing a thin client computing environment, as stated utilizing the variables above, will be $93,289.09.

4. Security- Annual security costs associated with the security of the data on the network. Security includes the actual security of the files to include backup and restore, centralized user file accountability and security of data leaving the environment via USB device, etc.

- ✓ Security cost calculations are broken into the number of users that require support in the environment. The cost for 0 to 500 users in a fat client environment is 10% of the security administrators time requiring $7,700.00 and labor costs. The cost for 0 to 500 users in a thin client environment is 2% of the security administrators time requiring a labor cost of $1,540.00. Because the images are stored centrally, client security is simplified over a distributed model/non-centralized additional savings can be recognized with thin client computing. Annual savings can range between $6,160 and $50,050 depending upon the number of clients being supported in the environment. See Appendix B of this white paper for additional information and breakdowns.

5. Labor Force Skill set- The costing model for maintaining your environment. Variance is discovered from fat to thin client, to include cost differences between thin client technologies.