




Thin-Client Diet

MILITARY STEPS UP SWITCH TO STRIPPED-DOWN TERMINALS AND CENTRALIZED COMPUTING POWER.

By PETER A. BUXBAUM
MIT CORRESPONDENT



The military is stepping up efforts this year to deploy thin-client computing devices and systems, which use networks to provide processing power and storage to stripped-down user terminals. In doing so, analysts say, the Department of Defense is taking a vital step toward the culmination of network-centric warfare, bringing the power of the networked environment to end users at all levels.

All of the armed services have already invested in thin clients to one degree or another, and Lieutenant General Steven Boutelle, the Army G6/chief information officer, recently announced a major drive toward migrating Army desktops to thin clients.

The Navy has initiated a pilot program that has installed thin-client technology on three of its vessels, with an eye toward transitioning many more in the not-too-distant future.

Thin-client solutions represent a combination of hardware and software. On the software side, thin clients are best thought of as Internet browsers—programs that pull information stored elsewhere but offer little or no computing power of their own. On the hardware side, devices have been developed to accommodate this configuration—typi-

cally smart computer screens that have no disk drives or moving parts but may include other features, such as security-related card swiping devices.

Thin clients are entirely reliant on the network for their computing power. Thick clients, such as the traditional desktop computer, include varying levels of their own computing ability.

The march toward netcentricity is one motivating factor for the deployment of thin clients, but it is not the only one. Thin clients advance netcentricity by emphasizing the computing power of servers over end-user devices, thus maximizing the potential of the network.

Thin clients also offer cost advantages in the form of cheaper hardware, maintenance efficiency by concentrating information storage and applications centrally, and enhanced security by relieving the end user of the responsibility for potential breaches. In some cases, however, savings on end-user devices are offset by more expensive servers.

But the transition to thin clients is no panacea, and the vision of supplying thin clients to the lowest tactical level of warfighting is not currently achievable. At this point, thin clients are best suited for business and administrative applications. On the battlefield, they become less useful as one proceeds down to the lowest tactical levels, owing to network connectivity and bandwidth constraints.

These negatives are being worked on, however, and military planners foresee the day when warfighters in the field will be equipped with thin-client devices, thus completing the networking of the armed services.

TRANSITION BENEFITS

“Thin clients are small network devices that replace or are used in lieu of desktop and laptop computers,” said Greg Giaquinto, senior electronics analyst with Forecast International. “Thin-client devices are small, quiet and use one-sixth the energy of a regular desktop computer.”

All of which suggest that cost savings can be achieved by migrating to thin clients. A Hewlett Packard study indicates that replacing half of existing desktops in larger organizations with thin clients would reduce the total cost of ownership for desktops by 50 percent. “That would be a substantial cost saving to the armed services,” Giaquinto noted.

“Improved security over traditional desktop PCs is another benefit of thin client technology. This is a major reason why military and intelligence groups have been early adopters of thin-client architectures,” Giaquinto continued.

Thin clients contribute to enhanced security by concentrating security administration in the hands of trained professionals and taking it away from end users, who otherwise have to update their systems and applications against emerging threats. “You can’t ask users to be security experts,” said Alan Paller, director of research at the SANS Institute, a computer-security training organization. “That doesn’t work because attacks have moved from operating systems to users. You want to try to not allow users to hurt themselves. Thin clients are the only proven way of doing that, because you’re moving control out of the hands of users and putting it with someone who has better skill at maintaining security.”

On a strategic level, the move to thin clients means “getting with the power of the Internet,” according to Robin Quinlan, deputy director for joint forces integration in the Office of the Undersecretary of Defense for Acquisition, Technology and Logistics.

“It’s the next step in technology to move applications from thick clients and onto a distributed set of servers running enterprise applications. This reduces a lot of the systems-engineering issues we encounter,” said Quinlan, who prefers the terms “net-enabled” and “autonomous” for thin client and thick client, respectively.

“In a world with no defined enemies, it is important to have immediate access to relevant information from a variety of sources,” said Asim Khan, director of business development at Amir Technology Labs. “Thin clients make this more reasonable because the information only has to be uploaded to one server. It makes more sense in a tactical environment to deploy thin-client devices to the field.” AT Labs has worked with government clients in Chile, Dubai and Kuwait in deploying thin-client systems and is gearing up to do business with the U.S. military as well.

TACTICAL APPLICATIONS

But thin clients are not quite ready for prime time when it comes to tactical applications, according to Quinlan. “Data latency needs to approach zero, and quality of service needs to approach 100 percent at the lowest tactical levels,” she explained. “When you are faced with an incoming cruise missile and you have milliseconds to react, you can’t



Greg Giaquinto



David Bachman



Alan Paller

have an hourglass showing up on your computer screen.”

On the other hand, thin clients can now be effectively deployed at the operational level up through the strategic level, where near-real time and non-real time applications are running.

In addition, warfighters at the lowest tactical level will often not have access to the bandwidth necessary to fully take advantage of network connectivity. “Tactical level networks like radio systems have throughput in the kilobytes instead of the hundreds of megabytes,” she said. “We need to develop enterprise-level services and standards compatible with the tactical level and to field the actual pipes out to the troops.”

David Bachman, federal manager at Wyse Technology, sees other obstacles to the widespread deployment of thin client technology in the military at this time. Wyse, the largest manufacturer of thin clients, has installations in the Army, Air Force, Navy and Missile Defense Agency, as well as more than 30,000 within Veterans Administration hospitals.

“It takes time and thought to do a thin-client deployment,” Bachman said. “It’s easy to wait until the end of the year and buy the same PCs you did last year.”

Besides overcoming organization inertia, Bachman sees the need to promulgate specific standards for thin-client technology. “Right now within DoD there are standards for operating systems, applications and security on the desktop,” he explained. “That all needs to be examined for thin-client devices as well. These kinds of things don’t turn on a dime even if there is backing at the highest levels.”

To address this, Wyse has developed a hybrid to the traditional thin-client architecture. “With our Streaming Manager Software, we provide a no-excuses XP Professional user experience, with all the same OS, applications, security and peripheral support of a PC, but on a diskless thin-client appliance. We take the image you would use on a PC, and stream it on-demand over the network to the appliance. All the updates and patching can be done centrally. It is a new way of looking at thin clients, and our DoD customers are very interested because it falls into line with their existing desktop configuration guidelines.”

Although not strictly thin-client technology, ClearCube addresses the cost and efficiency considerations by centralizing the operations of multiple computers in a single data center, thus requiring no new standards for thin clients.

“The big move right now from a military standpoint is to deploy centralized computing solutions,” said Ken Knotts, a senior technologist at ClearCube. “Our architecture fits perfectly into that scenario.”

ClearCube’s configuration consolidates 112 PC “blades” onto a rack where all desktop data is stored. A “blade” refers to an individual computer processing unit, which in this case is located in a central data center rather than at an individual work station. Desktop devices known as i/ports and c/ports route connectivity to the blades in the data center.

“The good thing about the ports is that they have no moving parts and generate no heat,” Knotts said. “Since there is no local storage, the blades provide enhanced levels of security and manageability while providing the same end-user experience as do desktop PCs. From an IT perspective, everything is back at the data center, so that,

in case of a problem, a technician can swap blades and get a user up and running in three clicks and without leaving his desk.”

ULTRA-THIN TECHNOLOGY

Thin-client deployments in the Navy are designed to address requirements revolving around reducing space, weight, power; enhancing information security; and ensuring the ability to participate with coalition network requirements aboard ship, according to Mario Diaz, Navy sales manager for Sun Microsystems.

“We don’t look at thin-client technology by itself, but look instead at the network holistically to maximize flexibility, capability and value,” said Robert Wolborsky, program manager for network information assurance at the Navy Space and Naval Warfare Systems Command in San Diego.

The Navy uses the term “ultra-thin client” for its Sun deployment. This refers to the fact that the network appliances run no operating system, nor do they include any moving parts, Diaz said. The Sun system has thus far been installed on three vessels—the *USS Coronado*, *USS Mount Whitney* and *USS Blue Ridge*, with the prospect of further installations beginning in the next fiscal year, according to Diaz.

Sun’s ultra-thin client includes a keyboard, a mouse and a 17-inch monitor “equipped with additional smarts,” as Diaz put it, to provide the network connection. These additional features include a card reader used to authenticate and provide access to credentialed users, a graphics card and a Java card.

The other key Sun contribution to the system is its Solaris servers, which mediate sessions between the user appliances and the Windows servers that run the Microsoft Office applications.

“Those components, the ultra-thin client and the Windows server, are the two bookends,” Diaz explained. “In the middle, acting transparently and with minimal interference with the applications, is a Solaris session server which manages access of client devices on the desktop to the applications running on the Windows server. What the Navy is trying to do is to maintain the Microsoft applications environment as its applications of choice, while accessing those in a more secure manner and achieving reductions in space, weight and power.”

The Navy’s deployment of the Sun system also has the important goal of supporting network-centric warfare and collaboration with coalition partners, Wolborsky said. Sun’s Solaris servers have the capability to provide access to users based upon the level of classification of the data they are seeking. In other words, the systems allow users to access information at their specific classification level on a server that includes several security classifications, or “enclaves.”

At this point, this functionality is being deployed at the “secret-releasable” level, referring to secret data that can be shared with coalition partners, depending on existing agreements which have been negotiated between the United States and its coalition partners. This security-classification functionality is enabled by cross-domain technology.

“This is a high-assurance capability that allows you to operate in multiple security enclaves in the same network architecture,” Wolborsky explained. “We built the system to make sure there was no leakage or spillage of data.”



Asim Khan



Ken Knotts

Without that capability, he added, “each enclave would require a separate network. The thin-client solution has provided us tremendous benefits at the secret level.”

For all of these benefits, Worbolsky noted that costs savings provide the least dramatic benefit as far as the Navy’s business model is concerned. “The fact that our systems are afloat is not addressed in by commercial thin-client technologies,” he said. “One of our highest costs is in servers. So while we are saving at the front end in the form of less expensive work stations, we spend more on servers than do buyers of commercial off-the-shelf technology. Most organizations are spending \$5,000 for a standard rack where we have to spend \$50,000 to \$70,000 for a hardened server.”

Furthermore, Worbolsky acknowledged, thin-client systems at this point are best for running business applications. “Environments running real-time applications,” he added, “do not lend themselves to thin clients at this time.”

BATTLEFIELD GATEWAYS

That is why current battlefield and tactical systems must strive to achieve “an appropriate balance” between the deployment of thick-client and thin-client devices, Quinlan observed. Devices deployed to the battlefield for real-time fire control are more likely to be weighted in favor of thick clients, in order to obviate difficulties in the event of a disruption in network connectivity. Higher-level situational awareness systems at the theater and national level will more likely include devices based on network-centric thin clients.

The standards and enterprise services required to be adapted to the tactical level are “now just being thought through,” Quinlan said. “I would guess it is going to be a decade before we see enough progress in regards to standards, enterprise services and in terms of providing sufficient bandwidth to the battlefield.”

On the other hand, the difficulties associated with deploying battlefield thin-client technology are not an all-or-nothing situation. In the interim, until true thin clients are feasible for deployment to the lowest tactical levels, incremental improvements have been developed, according to Quinlan. The current fighting in Iraq has seen the Army sharing situational awareness information with warfighters through gateways.

“Using gateways is a fancy way of saying that they are connecting everyone up with commercial software on the network through which they can access information,” she said. “This represents a positive workaround that may not be optimal but that is helping the situation. This has been one of the positive lessons learned from the fighting in Iraq.”

The increased use of gateways could have an impact on warfighter hardware requirements as well, Quinlan noted. “If I have a gateway to data, I don’t need the computing power you have one echelon above me,” she said. “You can do all of the number crunching and send it to me. I can have a thinner client not loaded up with all the bells and whistles.” ★

Contact Editor Harrison Donnelly at harrisond@kerriganmedia.com. For more information related to this subject, search our archives at www.MIT-kmi.com.