



Missing Link 
Security Services
Mark Bouchard, Founder

SSL VPN: The Ideal Approach for Secure Access to Thin-Client Computing Applications

Table of contents

Introduction	2
The Thin-Client Computing Scenario	2
TCC Fundamentals	2
Why TCC Matters	2
TCC Challenges	3
Providing Secure Access for TCC Environments	3
The Typical Contenders	3
The Strengths of SSL VPN Technology	4
Selecting the Right SSL VPN Solution	4
Summary.....	6
About the Author	6



Introduction

The benefits of thin-client computing (TCC) are well established. Also known as server-based computing (SBC), this technology routinely helps IT organizations lower operating costs, reduce the risk of data loss, and improve user productivity. Discussed far less, however, is that maximizing these gains depends on establishing a means to secure remote access to applications that are provisioned in this manner. For that matter, both users and IT administrators will derive considerable value from having a single, consistent approach for enabling remote access to all of the organization's applications.

This paper reviews the benefits and challenges associated with thin-client computing, and it explains why SSL VPN technology is an ideal approach for securing access to TCC environments. It also provides guidance on how to select the right SSL VPN solution for the job.

The Thin-Client Computing Scenario

Although TCC has a high degree of popularity among IT shops, it is not without its challenges.

TCC Fundamentals

Thin-client computing is a server-centric processing model where the application software, data and primary source of CPU power reside on a network server, rather than on individual client devices. Keyboard and mouse inputs transmitted from the client direct the execution of the application on a central server. Display updates are then returned to the client, typically using a protocol designed specifically for this purpose. The most common examples of TCC technology include:

- Microsoft Windows Terminal Services, which uses Remote Desktop Protocol (RDP)
- Citrix XenApp (formerly known as Presentation Server, and before that MetaFrame), which uses Citrix Independent Computing Architecture (ICA)
- X Window System, which is common to many varieties of Unix and uses the X protocol
- A wide assortment of platform-independent Virtual Network Computing (VNC) products, which rely on the RFB (remote framebuffer) protocol.

Why TCC Matters

The popularity of TCC derives from its ability to help address three technical problems that organizations have historically encountered with their applications:

- The poor performance that results when client-server applications are not properly designed for operation over the WAN
- The deployment and maintenance headache that ensues from trying to keep up with new versions of client software, bug fixes and security patches
- The challenge of supporting a diverse set of client platforms when applications were originally designed to only run on a given operating system and/or hardware configuration

Alternately, from a business perspective, the benefits of TCC can be grouped into three categories.

Lower operating costs. Deployment and maintenance costs are substantially reduced with TCC, since practically all of the client software is run and controlled centrally.

Reduced risk. With TCC, sensitive data is kept in the data center and off of relatively vulnerable client devices. In addition, a smaller software footprint yields less surface area for attacks against client devices, and centralized control means that applicable patches and security countermeasures can be implemented more quickly and consistently.

Greater productivity. Users receive better and more timely technical support due to the centralization of both administrative personnel and the software/systems on which they operate. Other related benefits include improved stability and availability due to centralized control over client images, simplified compatibility testing, and the ease with which a centralized backup and recovery solution can be implemented.



TCC Challenges

Increasing user mobility is one of the most prolific trends with which today's organizations must contend. Traditional mobile employees, such as sales and field-services personnel, are being joined by a steadily growing population of telecommuters and "day extenders." Consequently, one of the key challenges pertaining to TCC is the need to establish a means to securely access associated applications remotely. At a minimum, users' identities must be verified, and the confidentiality and integrity of their sessions must be assured as they communicate over open/public networks.

Of course, the need to provide secure remote access is not the only challenge organizations face. Other factors that complicate TCC implementations include the inability to work offline, difficulty handling certain types of applications (such as those with intensive graphics or deep-rooted platform dependencies), and the effort and complexity required to ensure the scalability, availability and accessibility of associated server infrastructure. These issues reveal why TCC, despite its strengths, is actually not appropriate for all users, applications and organizations. Indeed, the vast majority of businesses deploy a wide array of applications in a wide variety of ways. It is very rare for an organization to rely solely on TCC.

To summarize matters from the perspective of the IT department:

- TCC is an attractive delivery model that needs to be supported
- This support needs to include having a means to securely access TCC applications remotely
- TCC applications are not the only ones that require such a capability
- Implementing several different secure remote access solutions would be highly inefficient

Then there's the user's perspective. What users really care about is having one consistent and reliable way to access all of the resources they need, whenever and wherever they need them.

Providing Secure Access for TCC Environments

The good news is that organizations have several options for enabling and securing remote access to their TCC implementations. The bad news is that the different approaches vary considerably in terms of efficiency, breadth of capabilities and overall degree of effectiveness.


The Typical Contenders

Conventional techniques and technologies include taking advantage of native security features, employing solution-specific security gateways or implementing an IPSec VPN.

Native security features. Most TCC products include a handful of native security capabilities. These incur no added cost, but are often very basic. Typical in this case is support for basic user authentication and having a means to encrypt session traffic (for example, by using SSL). Other limitations are that: these features only work for applications provisioned with the given TCC product; they may need to be configured separately for each user, application and/or TCC server; and they are inherently part of the application/server infrastructure and, therefore, contribute to its complexity and the related challenges of scaling and managing this infrastructure.

Solution-specific security gateways. Some TCC products include optional, standalone gateways or other add-on feature sets that provide core services to secure access to the TCC environment. In general, the associated security capabilities are still somewhat basic, but support for external credential stores and additional authentication mechanisms is fairly common. Even though this approach is decoupled from application/server infrastructure and enables a more unified configuration model, support is still limited to the applications hosted with the given TCC product. One example of an add-on feature set, as opposed to having a completely separate gateway, is the ability to employ Active Directory and Group Policies to centrally set and manage granular access control rules for Windows Terminal Services implementations.

IPSec VPNs. This approach also entails a standalone gateway and, as such, conveys similar benefits in terms of having a decoupled infrastructure and easier configuration. The primary security features it provides, once again, are user authentication and encryption. Depending on the selected product, it may also support host-integrity checking, network-layer firewalling and some measure of network intrusion prevention. Another significant advantage is that, like a private WAN connection, IPSec VPN enables access to virtually any electronic resource. In other words, it supports secure remote access for any and all types of TCC technologies an organization may be running, in addition to all other services and applications that a user might need. That said, there are definitely



some disadvantages to IPSec VPNs as well. First, they require special client software. This can be a major headache to deploy, only works for a subset of devices and platforms, and is impractical for non-employees and unmanaged devices. The second issue is that IPSec VPNs work by enabling a full network-layer connection. Therefore, additional containment infrastructure needs to be deployed to ensure that users, or worse, infected machines do not overstep their intended access privileges.

Although they are achievable and perhaps even a good fit in limited situations, all of the options discussed above have one or more significant drawbacks. Fortunately, one approach still remains to be explored—a solution that exhibits all of the positive aspects of the other options and none of the negative ones.

The Strengths of SSL VPN Technology

SSL VPN technology is particularly well suited as a secure remote access solution for TCC, based on the breadth and depth of security functionality it delivers, the ability to support multiple TCC methods and technologies at once, and the ability to address all of an organization's other remote access needs as well.

Robust security. In addition to flexible, multi-factor user authentication and encryption for all sessions, SSL VPN solutions typically support a wealth of advanced security capabilities, including: host integrity checking; granular definition and enforcement of access policies; control over specific end-user functions such as copy, print and save; and extremely detailed logging/auditing of user activities.

In-depth support for TCC. Unlike some of the other alternatives, SSL VPN is independent of the TCC technology being used. This means that there is no need to make any changes to an existing TCC implementation to get secure remote access to work. Yet, at the same time, there is no need to avoid changes, since the remote access infrastructure is capable of supporting multiple flavors of TCC technology simultaneously. For example, an organization that has historically used Citrix XenApp for its TCC solution does not have to worry about adding a VNC-based product to its environment—at least not in terms of the impact this will have on their secure remote access solution, because there isn't any. Furthermore, a decent SSL VPN solution will also provide a range of value-added capabilities, such as single sign-on (SSO), smooth roaming (in the event of intermittent connectivity), and intelligent delivery of client software.


Extensive coverage for other needs. SSL VPN technology provides a secure remote access solution not just for TCC applications, but for all other types of applications as well. This aligns favorably with the trend of widespread migration away from client-server and alternative architectures in favor of Web technologies. Enabling a wide range of access modes also ensures the ability to accommodate a broad array of client platforms, which is a key capability in this age of mobility and with the increasing prevalence and variety of user-owned devices such as iPhones, Windows Mobile Phones, PDAs, and laptops running various operating systems.

The net result is that a good SSL VPN solution has numerous advantageous. Chief among these is that it keeps users happy and productive by giving them a uniform way to access all of the resources they need to get their jobs done, at the same time that it reduces the cost and complexity of the organization's computing infrastructure and its operation.

Selecting the Right SSL VPN Solution

But not all SSL VPN products are created equal. Those with functional limitations or incomplete feature sets will not be able to fully deliver the advantages outlined in the previous section. To guard against making a choice that is less than ideal, organizations should evaluate candidate products against the following criteria.

Comprehensive Access. Ultimately, the goal is to be able to provide any user, operating in any location, with practically any type of device, access to just about any service or application. From a practical perspective such access will not always be allowed, but the point is to at least have the capability so that it can be utilized when the need arises. From a technical perspective, this entails supporting enough access modes to account for all types of applications, including Web, client server, legacy and multiple types of TCC. Furthermore, it is important to understand the dependencies and limitations for each of the access modes. What client operating systems are supported? What browsers are supported? What, if any, client software is required, can it be dynamically downloaded, and what technology (such as Active-X) and configuration dependencies (such as user must have administrative privileges) are applicable. An ideal solution is one that incurs the fewest dependencies while still supporting all of the organization's access needs.



Comprehensive Security. Not only must data be protected while it is in transit and for whatever time it resides on a client device, but it is also essential to protect the organization's overall computing environment from remote systems that have been compromised or otherwise infected. In other words, security capabilities must be thought of in terms of providing end-to-end protection, and should ideally include the following countermeasures:

- Strong encryption for all access and administrative sessions
- Multiple authentication mechanisms, both for flexibility as well as to account for varying degrees of trust and risk
- Granular authorization/access control that can be dynamically adjusted based on a wide variety of attributes (user role and location, strength of authentication, ownership and security posture of the client device, to name a few)
- Client-oriented features such as the ability to check the security posture of the remote device, the ability to clear the browser cache at the completion of an access session, and the ability to keep any downloaded data in an encrypted workspace, or else delete it once the session is terminated
- Gateway-oriented features such as a hardened operating system, embedded firewalling, and mechanisms to thwart denial of service (DoS) attacks
- Detailed activity logging for both user and administrator sessions to facilitate troubleshooting and help demonstrate compliance with applicable regulatory requirements

Another security capability to look for is single sign-on. Ideally, the implementation should be comprehensive in terms of addressing applications with a wide range of native authentication mechanisms (for example, forms/header/cookie-based, Basic Auth, NTLM). Support for Security Assertion Markup Language (SAML), which can enable both intra- and inter-organization SSO, is also an increasingly important feature.


Transparency and Compatibility. On the one hand, this involves minimizing the effort and investment required by the users that require access. There should be no need to acquire, operate and maintain any specific software or hardware at the remote end of the session. In addition, any dynamically downloaded software, such as agents or plug-ins used to support certain access modes or security features, should not disrupt or otherwise change the operation of any programs or the client system itself.

On the other hand, the same conditions should also apply for the party providing the access. The SSL VPN gateway should just "fit in." Little, if any, network re-configuration should be required. Furthermore, it should be able to operate completely independently or, optionally, it should be able to take advantage of any existing credential and attribute stores (for example, LDAP directories), access management software and portal software that the organization is already using. Most importantly, it should not require applications and other resources to be modified in any manner in order to be remotely accessible.

Ease of Use and Administration. This category of criteria is somewhat similar to the previous one. However, in this case it is more about the day-to-day experience of the users, as well as the folks in IT/security operations. For the users, the key to success is ease of use. The interface should be intuitive, if not familiar, and very easy to navigate. Users should not have to sign on more than once in a given session. Nor should they have to make any decisions (for example, in terms of the access mode to use), other than to select the resources they want to access. In the event that one or more groups of users will have access to multiple resources, then a customizable, portal-style look-and-feel will be appropriate.

For the administrators, it comes down to management functionality. A centralized management capability is essential, but it should also be possible to delegate policy creation to local administrators who may be more familiar with a specific group of users and the resources they are accessing. When it comes to the policy model, there should be flexible grouping of related items, as well as re-use and modularity of object definitions and policy fragments. Overall, there should be an ability to implement virtually any access rule an organization can articulate. In addition, real-time session monitoring is helpful for troubleshooting purposes, while extensive logging capabilities are needed to support capacity planning and compliance reporting activities.

Performance. This category of criteria is intended to cover more than just system capacity or throughput. Given today's highly collaborative applications, latency requirements should also be considered when evaluating performance-related features such as the count and type of processors being used, expandable memory, and enhanced techniques for handling and inspecting packets/sessions.



Scalability is another important factor, particularly when it comes to cost effectiveness. This will be determined in large part by a product's management capabilities, but can also be affected by support for advanced features such as clustering and virtual systems. The latter enables a single physical system to be used to provide access to multiple constituencies, but in a manner that keeps their policies, session processing and activity logs separate from each other.

Finally, there is reliability, which will continue to gain importance as providing secure remote access to computing resources is increasingly elevated to the status of "business-critical service." In this case, the primary feature to look for is support for multiple High Availability (HA) configurations and mechanisms (active-passive, active-active, stateful failover, session persistence). Secondary considerations include redundant components like fans and power supplies, and the minimization of moving parts (for example, by using Flash memory instead of a hard disk).

Summary

Thin-client computing technologies, such as Citrix XenApp and Microsoft Windows Terminal Services, are widely deployed in organizations today because they serve a valuable purpose. These solutions facilitate the smooth delivery of applications, many of which would otherwise exhibit poor performance and/or be costly to maintain. However, given the trends toward increasing levels of user mobility and inter-connectivity between businesses, simple application delivery is no longer sufficient.

Achieving the full potential of TCC now requires that IT organizations also establish a means to provide secure remote access to its TCC applications. In this regard, SSL VPN technology should be considered an ideal solution. By selecting a leading product, as defined by the criteria covered herein, organizations can efficiently provide secure access not just to their TCC environment(s), but also to all of the other applications and services that users need to get their jobs done.

About the Author

Mark Bouchard, CISSP, is the founder of Missing Link Security Services LLC, a consulting firm specializing in information security and risk management strategies. A former META Group analyst, Mark has assessed and projected the business and technology trends pertaining to a wide range of networking and information security topics for over 10 years. He is passionate about helping enterprises address their information security challenges. During his career, he has assisted hundreds of organizations worldwide with strategic and tactical initiatives alike, from the development of multi-year strategies and overall architectures to the justification, selection, acquisition, implementation and operation of security and networking solutions.

2000271-001 June 2008