



Building an Extranet Portal

Table of contents

Introduction	2
Why Build an Extranet Portal in the First Place?	2
Why SSL VPN Technology Makes the Most Sense for an Extranet Portal.....	3
The Benefits of SSL VPN Technology.....	4
Selecting the Right SSL VPN Solution.....	5
Summary.....	6
About the Author.....	7



Introduction

It is a rare business, or even public-sector institution, that can survive in a vacuum. For most organizations there are a set of symbiotic relationships that inherently depend, at least to some extent, on the sharing of information. For example, retailers and suppliers must communicate to express their desire to buy and sell products, to negotiate pricing, and to establish delivery schedules. Manufacturers have similar needs, but in both the up and down –stream directions. The economic benefits of specialization have also led to a steady growth in the number of business “partners” that the average organization maintains. Practically everyone relies on third-party suppliers to deliver components and services that are beyond their “core capabilities”. But even those organizations that have chosen the opposite path and which are highly vertically integrated have one external party they can’t ignore: their customers!

It may be hard to imagine now, but it was only two decades ago that the communications necessary to maintain and grow all of these relationships were predominately accomplished via phone, fax, and postal/courier service. The emergence of business computing alone didn’t really change this situation; however, the introduction of networking technologies certainly did. Even early solutions for electronic interactions (e.g., based on private WAN connections and electronic data interchange technologies) yielded significant gains in efficiency, cash management, and competitive positioning, typically as a result of enabling greater degrees of automation and more timely sharing of business critical information.

Of course, success over the years has led to demand from management for more of the same. Organizations are now looking to achieve even greater operational and strategic gains, not only by giving (as well as getting) deeper levels of access to electronic information and services, but also by doing so with a much broader set of external parties. Interestingly, the growing ubiquity of Internet services has both facilitated and exacerbated this situation. On one hand it has enabled relatively inexpensive any-to-any connectivity, but this in turn has opened the door to a flood of opportunities that are just begging to be enabled. Indeed, the challenge today is identifying an efficient and cost-effective solution that maximizes the ability to take advantage of these opportunities while also minimizing complexity and security related risks. To that end, this paper explores the virtues of an extranet portal and highlights SSL VPN technology as the best approach to achieve such a solution.

Why Build an Extranet Portal in the First Place?

By now the value proposition of providing external parties with access to various applications and information resources should be clear. But what makes an Extranet Portal a good approach to do so? For that matter, what is an Extranet Portal? The answer to these questions can be arrived at by examining the definitions of the two parts of this term.

Extranet. By way of lead in, an intranet is a private computer network that uses web technologies (e.g., browsers, http servers, and html content) and network connectivity to share part of an organization’s electronic information and resources with its own employees. Typically, an intranet is accessed locally over the corporate LAN, or remotely via private WAN links or secured public Internet connections. In contrast, an extranet is a private computer network for sharing part of an organization’s electronic information and resources with external parties (e.g., suppliers, service providers, partners, or even customers). Historically, extranets have not been limited to use of web technologies and applications, and access has been achieved either via private WAN links or appropriately secured public Internet connections. In any event, the relevant and somewhat anti-climatic point in this context is that “extranet” simply refers to who has access – not the general public, not the organization’s employees, but rather selected, pre-approved external parties.

Portal. A simple definition of a portal is that it is a site of sites – in other words a super-site enabled by a purpose-built portal application that serves as an organizational framework and launch pad to a collection of other, often inter-related sites. The term also typically connotes some degree of personalization and collaboration capabilities, as well as client platform agnosticism/flexibility. More generically the term portal has come to refer to any gateway that aggregates, organizes, and serves as the unified entry point to a collection of resources, not all of which are necessarily web based.

A classic use case for traditional portal software is to vastly enhance and simplify a corporate intranet. Instead of having to know, remember, and search for all of the internal services that have been “webified”, employees can simply turn to the organization’s intranet portal. From there they can see and access all of the information



and resources they are entitled to, such as: helpdesk services, IT services (e.g., software downloads, computer requisition), security policies, human resources policies and services, payroll services, corporate activities and information, project workspaces, knowledge management repositories, and so forth.

In a similar fashion, an extranet portal could be erected to meet the needs of an organization's customers, constituents, and/or partners. This sort of resource externalization will clearly have implications from a security perspective, which may in turn impact the technology used to build the portal capability. However, the benefits should remain the same as those which are derived from an intranet portal. In particular, the advantages of providing and controlling access via a portal include the following:

- **Portals are user friendly.** They help ensure uniformity of the user experience, facilitate identification and navigation of accessible resources, enable collaboration, and in many cases support single login capabilities. As such, they help ensure user satisfaction and recurring usage of the organization's electronic services.
- **Portals are IT friendly.** They utilize a familiar and simple paradigm which reduces the need for user training and help desk support. They provide a quick disconnect capability for terminating user access in the event of security or user/account specific issues. And in many cases they also help to unify, simplify, and reduce the amount of front-end infrastructure that is needed to make existing applications and services accessible.
- **Portals are business friendly.** They streamline routine operations and facilitate collaboration and timely information exchange, which are the keys to improved productivity, better business decisions, and sustainable strategic gains. In addition, they ensure a consistent "face" or "wrapper" (i.e., brand, style, and image) for the organization's electronic services.

Notably, even further advantages can be derived depending on the technology that is selected and implemented to build the extranet portal.

Why SSL VPN Technology Makes the Most Sense for an Extranet Portal

There are several approaches and related access and security technologies that can be used to implement an extranet, if not a full-blown extranet portal. However, they are by no means equal in their capabilities and cost effectiveness. This situation is made apparent by the following review of the most common techniques and technologies used to provide secure access to external parties.

Private 1:1 Connections. Having a dedicated leased line or frame relay circuit to an external party is certainly an option, and may even be warranted for particularly critical partnerships, or perhaps just ones that involve significant volumes of sensitive or bi-directional traffic. However, it is not a scalable approach, requiring a dedicated circuit for each partner. Furthermore, without significant investments in containment infrastructure it has the potential security risk of giving the partner access to a greater portion of the organization's environment than is really warranted, or wise.

Hosted Exchanges. The original version of this approach was the electronic data interchange (EDI) value added network (VAN). EDI VANs were essentially hosted messaging providers, accepting and delivering standardized (i.e., ANSI X.12) transaction data between subscribers. The primary benefit, beyond standardization, was the ability to conduct business with multiple partners over a single network connection. However, growing ubiquity of internet services steadily eroded that advantage. An inability to efficiently keep up with new transaction sets/types further contributed to their demise. The more recent incarnation of this approach is the extranet service provider (ESP). A classic example of an ESP and its corresponding extranet is ANXeBusiness Corporation and the ANX Network (aka the Automotive Network Exchange). Overall, ESPs have only been moderately successful. While they offload partner certification, connectivity, and interoperability issues, they typically involve relatively high operating costs, address only a subset of an organization's partners and interaction requirements, and still require significant back-end integration.

Conventional Web Servers. Returning to self-hosted solutions, a relatively basic approach to facilitate interaction with third parties is to "webify" resources that will be accessed. This involves creating/separating the presentation layer (i.e., user interface) for all web applications and building web front-ends for any non-web applications. Associated web servers are then hardened and deployed in a DMZ. Although suitable for organizations with fairly basic interaction needs, this approach has several potentially significant drawbacks. These include exposure of notoriously weak web servers to the Internet, substantial application development/transformation (and resulting limitations in terms of which applications can in fact be exposed), as well as a general lack of scalability.



Reverse Proxies. A common variation to the previous approach is to use a reverse proxy server as a somewhat artificial (or pre-canned and pre-hardened) front-end. This is particularly useful for applications where separation of the presentation and logic/data layers is problematic, or even impossible. It saves effort in terms of application development/transformation, often improves security, and can reduce costs if it is operated in a 1-to-many relationship (i.e., one proxy server acting as the front-end for multiple, internal applications). However, it is generally limited in that it can only support access to web applications. And even then, there may be a significant management burden due to the need to manually develop and maintain a map between external-facing URLs and internal content, along with an inability to support certain application techniques and technologies (e.g., client-side scripting, Flash). Furthermore, from a security perspective, it is important to acknowledge that most reverse proxy servers are really little more than slightly modified web servers.

Externalized Intranet Portals. Reverse proxies can also be used as a means to front-end traditional portal software, thereby making it accessible to external parties. Unlike just about every other approach, this unlocks and makes available all of the benefits associated with using a portal. But the issues with reverse proxies will still be applicable and, most notably, not all resources will be accessible using this technique.

Access Management (aka Web Single Sign-On) Technologies. Depending on their architecture, some of the representative access management products are essentially equivalent to a reverse proxy. The primary alternative architecture, involving the use of agent software, is really just a complementary add-on to the extranet approach that relies on conventional web servers. In this case, the access management product provides a front-end service that delivers single sign-on capabilities in combination with dynamic determination and enforcement of highly granular access rights. In other words, it enhances the usability, security, and cost effectiveness (e.g., via consolidation of rights management) of the solution, but does not significantly change the underlying architecture (i.e., putting web servers in the DMZ) or its limitations.

IPSec VPNs. The greatest virtue of IPSec VPN technology, is that like SSL-enabled web infrastructure (i.e., browsers and servers), it transforms relatively flexible and inexpensive public Internet services into a highly secure means of establishing any-to-any connectivity. Unlike SSL technology, however, it is not a transparent solution. Instead, it requires the deployment of IPSec gateway devices and IPSec client software. IPSec is also unlike SSL, which enables connections to specific applications, in that its sessions connect at the network layer. This introduces the need for additional containment infrastructure to adequately control the scope of access that is provided. Despite these drawbacks, the IPSec approach does have the advantage that, like private WAN connections, it can facilitate access to virtually any electronic resource, web and non-web alike.

Although achievable and perhaps even well-suited in limited situations, all of the options reviewed to this point have had at least a couple of significant drawbacks. Fortunately, one approach still remains to be explored – a solution that exhibits all of the positive aspects of the previously discussed options, and none of the negative ones.

The Benefits of SSL VPN Technology

SSL VPN technology is particularly well suited to building and securing extranet access solutions based on its ability to deliver the following, core capabilities:

- Transparency and cost savings result from not having to deploy and maintain gateways or client software at the external party's end of the connection. In addition, who owns and manages the client device and its type (e.g., PDA, laptop, desktop) does not matter, so long as it is running a compatible browser.
- Access to all IP-based resources results from an ability to support both application layer connections as well as full, network-level connections.
- All resources (including non-web applications) can be externalized via a portal interface, without the need to re-format, webify, or "publish" them. An automatic "re-writing" capability takes care of all internal-to-external mapping chores and enables the use of essentially all application techniques and technologies.
- A single system can typically support multiple, uniquely branded portals to enable a customized look and feel for different constituencies, including a mix of employees, partners, and customers.
- Both single sign-on and granular, dynamic access control is provided without the need to purchase separate access management software and deploy agents on DMZ-based web servers.
- Deployment costs, operational costs, and complexity are reduced and security is enhanced as a result of having a single-box solution in the form of a pre-hardened appliance. Having centralized policy management also simplifies matters.



- In addition to all of the security enhancing features that have already been mentioned, SSL VPNs also typically support a reverse proxy access mode (which enables brokering of user sessions and provides an opportunity for advanced inspection), client integrity checking (which makes user access conditional upon the security posture of their client platform), as well as highly granular activity logging.

The net result of these capabilities is that SSL VPN products are increasingly being used not only to provide access for remote employees, but also for an increasing number of external parties, and, in both cases, to a much greater variety of data and applications. Furthermore, SSL VPNs support having these resources be provisioned and accessed in a one-off fashion, or optionally in the manner of an extranet portal.

Selecting the Right SSL VPN Solution

Much like with the various options for implementing an extranet, not all SSL VPN products are created equal either. Those with functional limitations or incomplete feature sets will not be able to deliver the full set of advantages outlined in the prior section. To ensure against this situation, organizations should evaluate candidate products against the following selection 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 internal 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 inevitably arises. From a technical perspective, this entails supporting enough access modes to account for all types of applications (e.g., web, client server, legacy), including those with dynamically negotiated attributes (e.g., ports, addresses). Furthermore, it is important to understand the dependencies and limitations for each of these 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 (e.g., Active-X) and configuration dependencies (e.g., 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. In particular, it should be possible to access a variety of resources (e.g., virtually all web applications, file services, terminal access) without the need for any client add-ons.

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 calculated based on a wide variety of attributes (e.g., user role and location, strength of authentication, ownership and security posture of the client device);
- 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 attacks; and
- Detailed activity logging for both user and administrator sessions, to facilitate troubleshooting as well as to help demonstrate compliance with various regulatory requirements.

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



Transparency and Compatibility. On one hand this involves minimizing the effort and investment required by the parties gaining 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 (i.e., agents/ 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, essentially 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 (e.g., 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 externally 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 (e.g., 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 (i.e., 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 access to external parties 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 configurations and mechanisms (e.g., active-passive, active-active, stateful failover, session persistence). Secondary considerations include redundant components (e.g., fans, power supplies) and the minimization of moving parts (e.g., by using Flash memory instead of a hard disk).

Summary

Businesses and public institutions alike understand the operational and strategic benefits of enabling an increasing number of external parties with access to a growing variety of internal data, services, and applications. With a rapidly evolving technology landscape and ever increasing degrees of user mobility, they should also realize that maximizing these benefits depends on supporting access from anywhere and with any type of device (e.g., managed and unmanaged; PDA and laptop). The result is the need for a solution that efficiently, securely, and cost effectively provides such an extensive depth and breadth of access capabilities.

Faced with this set of requirements, organizations are advised to consider implementing an extranet portal, in particular by deploying SSL VPN technology. In terms of selecting a specific product, they are further advised to focus on those that are well aligned with the extranet use case. This can be established by evaluating candidates against key criteria in the areas of comprehensive access, security, transparency and compatibility, ease of use and administration, and overall system performance.



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 information security topics for nearly 10 years. He is passionate about helping enterprises address their information security challenges. During his career he has assisted hundreds of organizations world-wide with everything from strategic initiatives (e.g., creating 5-year security plans and over-arching security architectures) to tactical decisions involving the justification, selection, acquisition, implementation and operation of their security and privacy solutions.