



**Missing Link**   
**Security Services**  
Mark Bouchard, Founder

# Taking Secure Access to the Next Level – Achieving Granular Control that Really Works

## 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 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 individual security and privacy solutions.



## Introduction

---

These days IT is under pressure to ensure that applications and other information resources are broadly accessible. To drive top line growth and achieve bottom line efficiencies, it should at least be possible for virtually everything to be accessible to everyone regardless of their physical location or the type of client device they're using. Road warriors, telecommuters, day extenders, office-bound employees, service providers, strategic partners, temporary guests, constituents, customers, and even temporarily displaced users operating in disaster recovery mode must all be able to access the applications and services that are the lifeblood of the organization.

SSL VPN technology has proven to be a key enabler for meeting this need. In contrast to alternative technologies, such as IPSec, the ability of SSL VPNs to operate without pre-installed client software eliminates a significant administrative burden, reduces operational costs, and dramatically improves flexibility by extending applicability to unmanaged devices (e.g., public kiosks and computers owned by employees, customers, or partners).

At the same time, however, IT is also under pressure to scrupulously protect private information, provide robust security in general, and ensure compliance with prevailing regulatory requirements. As a result, adhering to the principle of least privileges is rapidly becoming a firm requirement, as opposed to just being a somewhat lofty, long-term objective. SSL VPN technology once again has the potential to solidly meet these needs (in large part due its NAC-like capabilities). However, this is only the case to the extent that a solution provides truly granular access control.

As this paper will explain, taking secure access to the next level by effectively implementing granular access control depends not only on the presence of related functionality but also on the efficiency and uniformity with which those features can be applied and the overall usability of the resulting configuration. In this regard, most SSL VPN products currently available in the market exhibit inconsistencies, dependencies, and other shortcomings that complicate matters considerably. Organizations, therefore, need to be mindful to avoid these all-too-prevalent pitfalls when selecting an SSL VPN solution.

## Core Capabilities

---

Notably, the goal of establishing granular control over access to computing resources is applicable not just for remote users, but for ones who locally connect to the enterprise network as well. This new reality is being driven by the profusion of itinerant and guest users, along with a host of regulations calling for organizations of all types to better secure their internal networks. And this is why it is now appropriate to think and act in terms of "secure access", as opposed to the narrower domain of "secure remote access".

In any event, to even begin to support granular control a secure access solution must incorporate certain foundational functionality. These core capabilities are aligned with the four principle classes of attributes used to translate high-level policies into detailed access control rules: user, client, network/location, and resource.

**User attributes** are fairly straightforward. Who a user is, how strongly his identity has been verified, and what role(s) he has are all typical factors used to govern access. Core capabilities relevant in this case include items such as: seamless support for a wide variety of authentication mechanisms (e.g., passwords, one-time passwords, tokens, certificates), extensive integration with common enterprise directories (which is where the bulk of role information already resides), and the ability to natively define user groups.

**Client attributes**, in contrast, can be relatively involved. What type of device is being used (e.g., PDA, smartphone, laptop, or desktop), who owns and/or manages it, specific configuration settings, and the presence and status of various pieces of security software should all be available as variables when creating access rules. The key enabling capability for capturing all of this information is commonly referred to as host integrity checking. And the potential complexity is due in large part to the vast scope of checks that can be performed and the subsequent challenge of mapping all of the permutations and combinations to a discrete outcome, such as allowing, denying, or restricting access in some manner.

**Network attributes** are intended to account, to the extent possible, for the type of access network and technology being used (e.g., public/private WiFi, public Internet). User location (e.g., local LAN, home office, partner facility) may also convey additional information regarding the inherent degree of privacy and security—or lack thereof—that can be expected. Host integrity checking is applicable once again, as are a range of passive techniques for detecting associated telltales.



**Resource attributes** are where the rubber meets the road. Until this point, all of the attributes have been used to granularly establish the level of trust that can be placed on the accessing entity and its associated conditions. In contrast, resource attributes establish the second half of the equation by accounting for the value and sensitivity of the resource being accessed. This enables access to be allowed or denied on the basis of whether the trust level for the first part of the equation exceeds the threshold corresponding to the value of the resource being accessed. Inherently wrapped up in this of course, is the resource itself and the granularity with which it can be defined and controlled. Relevant capabilities to look for in this regard include: the ability to allow access to email yet selectively block associated attachments; the ability to allow access to file services while blocking individual files; and, the ability to selectively allow/block access to individual functions within applications.

Technically, there is one other, somewhat secondary attribute class that deserves to be mentioned.

**Mitigating measures** are significant because they can provide an additional boost to the trust level of an access scenario, thereby paving the way to allow access to sensitive resources in what would otherwise be risky situations—for example, accessing company financials from a home PC. Often treated as a subset of the client attributes and typically available only with advanced SSL VPN solutions, the associated capabilities include:

- Recurring host integrity checks, which refers to the ability to periodically repeat the assessment of a client to detect any changes that occur as an access session proceeds;
- Cache cleaning, which refers to the ability to remove information remnants from browser and application-specific caches upon completion of an access session;
- Secure virtual workspace, which refers to the ability to create an encrypted workspace on the client device to help prevent data leakage;
- Information control, which refers to the ability to limit what a user can do with data that is accessed (e.g., copy/paste, save, print); and
- Client-based threat protection, which refers to the ability to dynamically download security modules at the start of an access session to help mitigate against any malware that might be present, such as keyloggers and trojans

Of course, the use of strong, multi-factor authentication would qualify as a mitigating measure as well.

## Equally Essential Elements

The thing to realize when it comes to actually achieving granular access control is that the core capabilities identified above, although necessary, are by no means sufficient. Indeed, there are several additional ingredients that are essential to getting the job done effectively and efficiently.

One of these items is policy and, although it is beyond the scope of a specific product or technology, it deserves to be acknowledged for two important reasons. First, having a detailed access policy is basically a prerequisite; without it, the degree to which access is controlled using the various capabilities will be haphazard, at best. The second reason is that policy provides a means, albeit imperfect, for extending beyond the technical capabilities of a given product. For example, an SSL VPN might not be able to prevent the use of USB memory sticks when access is made from an unmanaged device. But awareness of a corresponding policy could stop users from doing it anyway.

Returning to items that are directly pertinent to the selection of an SSL VPN solution, three other essential elements to look for are consistency, manageability, and usability.

**Consistency** can be an irksome issue. Having key features that work in some scenarios but not others only complicates matters. At a minimum, it means that administrators need to create multiple sets of rules or, worse, bend their policies to accommodate the product's idiosyncrasies. An example is when host integrity checking is not supported for the full range of client devices the organization is using, or across all of the different modes of access the product provides. Somewhat more subtle is when a feature is available but just works differently under different conditions. In either case, having greater entropy only serves to reduce administrative efficiency and increase the likelihood of configuration errors.



**Manageability** is another major factor in the quest for granular control. It should already be quite apparent given the number and variety of attributes that have been discussed, but the permutations and combinations that can be used to construct access rules is practically endless. This emphasizes the need for SSL VPN solutions to incorporate management features that foster ease of use and efficiency and help reduce complexity, such as:

- An intuitive, object-based, hierarchical policy architecture that enables flexible grouping of related items and reuse of rule “snippets”;
- A unified policy model where all attribute classes can be accounted for in a single, integrated rule;
- The ability to visualize object relationships;
- Pre-populated pull-down lists for establishing the specific host integrity checks to be performed;
- The ability to have user-oriented wildcards/variables to help reduce rule count (e.g., when providing granular access to file services);
- The flexibility to enable different access modes on whatever basis the organization deems applicable;
- The ability to test access rules before deploying them; and
- Ready access to detailed information about ongoing sessions to facilitate troubleshooting of mis-constructed rules.

Of course, even with such features it will still be possible to create an insanely complex configuration that is subsequently quite challenging to manage. From a best practices perspective, therefore, administrators should also help themselves by (a) not going overboard and (b) making an upfront investment in the form of access scenario profiling and reduction.

**Usability** is intended to account for the impact that greater granularity can have on end users. Establishing greater granularity is only a good pursuit up to the point that it confuses users or otherwise needlessly impedes them from getting their jobs done. Avoiding issues in this regard largely comes down to having the ability to make the user experience at once familiar, transparent, and well organized. Users should never have to select the mode of access or type of client software that will be employed in a given scenario. In addition, if users normally access resources clicking on their task bar, then that approach should be available to them. Alternately, if they are accustomed to a portal-style interface, then the arrangement of accessible resources should be intuitive and completely customizable—ideally by both the administrator and, at least to some extent, by the users themselves. Finally, it should be possible for users to receive clear notification of the specific reason when they are blocked from resources they are typically able to access (e.g., because they’ve failed a specific host-integrity check).

## Dirty Little Secrets

---

Unfortunately, but not surprisingly, downplaying features they don’t have, focusing disproportionately on the ones they do, and leaving little things like the degree of consistency for a given feature to the fine print is a relatively common practice for vendors of IT products—especially those that are trailing the leaders in their segment. For SSL VPN solutions, some of the more common shortcomings, in particular relative to achieving highly granular secure access control, include the following:

- Accessing some types of applications and services requires a pre-installed client.
- The extent of directory support and integration is relatively superficial, thereby limiting the depth of user attributes that are readily accessible.
- Host integrity checking is not available for all access modes.
- Host integrity checking is not available for a sufficiently broad range of client platforms.
- The variety and depth of inspections that are possible with the host integrity checking capability is fairly limited (e.g., can only look for antivirus and personal firewall software).



- Resources can only be specified at a relatively coarse level, such as IP address and top-level domain as opposed to individual applications, functions, and sub-domains or lower-tier URLs.
- An excessively rigid policy model restricts the manner in which attributes can be combined, reducing flexibility and limiting the scope of access rules that can be constructed.
- Cumbersome configuration capabilities require constant “flipping” between disparate screens, decreasing the efficiency of rule generation.
- Mitigating measures are not consistently applicable or, worse, are not even available in the first place.
- Some access modes and control features are not consistently applicable because they require the user to have administrative privileges on the client device.
- There is no ability to limit the number of access sessions a user has open at any given time (e.g., to help avoid cross-session leakage).

It is clearly in the best interest of organizations to be aware of these types of potential deficiencies. This way they can be used as an additional filter when selecting an enterprise SSL VPN solution.

## Conclusion

---

Businesses of all types and sizes are striving to provide broad access to their applications and related information resources. But they must do so in a way that is consistent with the somewhat competing objectives of protecting the privacy of sensitive information, maintaining the integrity and availability of their networked systems, and ensuring compliance with applicable regulations. In practice this means that access to resources will be limited, effectively, to the extent that it can be granularly controlled.

SSL VPN technology, at least in theory, addresses this situation by simultaneously delivering both broad access capabilities and robust control features—such as host integrity checking and fine-grained definition of resources. However, organizations must take care to not be misled by the mere presence of the latter. Unlocking the full potential of a secure access solution by achieving highly granular control also depends on the consistency, manageability, and usability of the related features. Gaps in feature coverage, an inflexible and inefficient policy architecture, and a cumbersome or cluttered interface for end users will all have the effect of impeding rather than promoting broad access to applications and services.