



Missing Link 
Security Services
Mark Bouchard, Founder

Internal Network Security


Go Inline for Greater Gains

Abstract

As enterprises consider internal or LAN security solutions, they face a confusing set of options. They can choose among solutions that reside at the client, the data center, or in the network. Network-based approaches are available as overlay or inline devices. This paper compares the relative strength of those approaches and concludes that network-based, inline solutions provide the greatest set of advantages.

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 spent the past decade assessing and projecting the business and technology trends pertaining to a wide range of information security topics. During this time he has assisted hundreds of organizations worldwide with everything from strategic initiatives (e.g., creating five-year security plans and over-arching security architectures) to tactical decisions involving the justification, selection, acquisition, implementation and ongoing operations of individual technologies/products. He also routinely works closely with the creators and sellers of information security solutions, helping them to better understand and meet the needs of the market at large.



Enterprises worldwide are waking up to the need to secure their internal networks. Indeed, recent surveys indicate that approximately two-thirds of organizations consider addressing internal threats to be their greatest current security challenge.

This situation derives in part from the need to comply with a considerable collection of imperative yet ambiguous privacy and security oriented legislation (e.g., S-Ox). More significant, however, is the fact that the inadequacies of a perimeter-only security strategy are repeatedly being revealed by the increasing frequency of malware gaining entry via alternate paths – in particular, the ever-growing population of local and remote connections needed to support mobile employees, guest users, and a vast array of business partners.

Not surprisingly, the vendor community has responded to this issue by launching a dizzying array of so-called internal network, or LAN, security solutions. Network admission control (NAC) is perhaps the most recognizable among these, but with many interesting variations and combinations of firewall, intrusion detection/prevention and identity management technology also being offered, NAC is far from alone. Consequently, the challenge now facing organizations is to navigate the various solution architectures and establish which approach yields the greatest effectiveness while incurring the least amount of disruption and expense.

To help organizations find their way, this paper will illuminate the relative strength of internal security solutions. The focus is on solutions that rely on inline devices deployed in the vicinity of end-user workstations. To be clear, an inline device is one that directly participates in the flow of communications traffic, typically processing the traffic in some fashion and selectively forwarding it on. This architecture stands in contrast to out-of-band devices, which can operate only on a passively captured copy of the traffic and then influence its flow indirectly by causing an inline device to take action according to its commands.

First Things First

To properly evaluate the various architectural approaches available, it is first necessary to acknowledge the objectives that internal security solutions are intended to address. From a technical perspective, the goals include (a) to prevent the spread of malware, and (b) to prevent the misuse of computing resources or corporate information. In addition, corresponding business-oriented goals are to:

- maximize availability of the network (which, in many instances, has become a mission-critical tool);
- control costs associated with owning and operating the network; and
- ensure the privacy and integrity of sensitive information.

Considering these goals and objectives in aggregate yields the following criteria for evaluating the approaches used to achieve internal security.

Scope of coverage refers to both the percentage of end-user connections that a solution can address and the amount of the network that is afforded protection.

Scalability involves the number of components needed relative to the degree of coverage provided. As a result, it loosely correlates with the cost of both acquisition and ongoing operations.

Degree of visibility and control refers to the granularity of a solution, primarily with regard to the networking/communications stack. In other words, does it have access to only network-level details (e.g., ports, protocols, IP addresses), or can it also decipher application-layer information? Most often, a high degree of visibility will also correspond to a high degree of control (e.g., the ability to prevent an end user from conducting web-email sessions, as opposed to having to shut down all HTTP-based traffic, or worse, all traffic from that user).

Degree of protection describes the scope and types of protective measures that a solution provides. Typical options include authenticating the end user, confirming certain configuration aspects of the user's computing station, and subsequently controlling the user's access to resources



contingent on these variables (i.e., identity and host state). Depending on the solution architecture, it might also be possible to conduct deep packet inspection and flow reassembly. These mechanisms facilitate more refined access control as well as intrusion detection/prevention capabilities.

Degree of complexity often, but not necessarily, correlates with the previously described scalability attribute. In general, it corresponds to the amount of “stuff” that must be implemented and the degree of effort required to make it all work together. In other words, how many existing or additional items (e.g., devices, software agents, protocols) must be deployed, configured, integrated, operated and maintained.

Not All Approaches are Created Equal

With the necessary groundwork in place, it is now possible to conduct a meaningful assessment of the four architectural approaches that dominate today’s internal security solutions. These approaches are classified according to the location where enforcement takes place and include: the client-based approach; the data center-based approach; and two network-based approaches, overlay and inline. To be clear, any given solution – or customer implementation – may actually combine elements from multiple approaches, but each is covered separately here to simplify the explanation and comparison.

The Client-Based Approach

This approach involves using native, operating-system capabilities or, more likely, one or more add-on software agents, to establish a security boundary at each client device. In general, because multiple types of security mechanisms can be accommodated – such as user log-on, anti-virus, personal firewall, intrusion protection – this approach has the potential to provide high degrees of visibility, control, and protection. Complexity can also be kept in check, assuming agent installs are kept to a minimum, upgrades are automatic, and a unified management console is available.

The need for a client-based component on each and every endpoint, however, will present a challenge, particularly in terms of scope of coverage. An ever-growing diversity of device types and client operating systems coupled with the inability to force the use of specific software on the increasing population of devices not owned by the organization will all but guarantee incomplete coverage. The result is a solution with a variable, if not outright unsatisfactory, degree of effectiveness.

The Data Center-Based Approach

Opposed to establishing a security boundary on each device, this approach is all about erecting DMZ-like defenses in the vicinity of the data center ... or wherever significant collections of servers needing protection reside – and therein lies the challenge of this approach. As with the client approach, the range of different security tools that can be implemented means the potential exists for high degrees of visibility, control, and protection. Even complexity can be addressed, especially as unified threat management devices become increasingly applicable even in high-capacity, high-criticality use cases.

However, also like the client approach, scope of coverage, and potentially scalability, are the Achilles’ heel(s) of the data center approach. Protection is afforded to only that portion of the environment that resides “behind the (data center) wall.” But what good is a pocket of protected resources if the rest of the networked environment is plagued with misuse (e.g., excessive peer-to-peer file sharing), malware, and, consequently, frequent disruptions to availability/access? Furthermore, scalability becomes an issue when multiple instances of the internal DMZ are needed because in reality, an organization’s servers (i.e., most sensitive resources) do not conveniently reside within a single, well-defined data center.

The Network-Based Overlay Approach

In general, solutions using an overlay are relatively complex because of the number of components involved. An overlay’s chief characteristic is that the policy decision point is separate from the policy enforcement point. For a typical setup – such as with Cisco’s version of NAC – one or more strategically deployed, out-of-band, policy-decision devices will coordinate the response/enforcement



activities of a larger, distributed cadre of inline devices (e.g., workgroup switches, routers, firewalls, other modified/purpose-built platforms).

Many overlay solutions also depend on a client-based agent to obtain and use information about client state (e.g., the presence of patches or anti-virus software) as part of the policy decision. Of course, those that do so are subject to the same limitation as the client-based approach in terms of scope of coverage. Alternately, some products attempt to overcome this shortcoming by utilizing “client-less” methods (e.g., network-based scanning) to obtain their information. However, the gain in coverage achieved with this variation will inevitably be offset by a dramatic reduction in the breadth and depth of information that can be obtained.

In any event, while the overlay approach ostensibly benefits from re-using existing infrastructure, it is also subject to the following potential drawbacks.

- Interoperability/integration must be achieved and subsequently maintained among the various components, which may involve significant upgrades to or wholesale replacement of existing gear.
- The degree of visibility, control, and protection is largely dependent on the capabilities of the policy enforcement devices. As a result, ordinary networking devices will yield less benefit than those which, for example, are capable of deep packet inspection.
- Both of the previous items/conditions can also lead to gaps and/or inconsistencies in the coverage provided, particularly when devices from different vendors are involved.
- Similar to the data center-based approach, scope of coverage will also be an issue for those overlay solutions that rely on policy enforcement points only or primarily at major network intersections (i.e., backbone locations), as opposed to being in closer proximity to the point of user/client connection.

The Network-Based Inline Approach


This approach relies on inline devices operating in close proximity to the point where client machines make their initial connection to the corporate network – in other words, coincident with a workgroup or distribution-layer switch, or at least in that general vicinity. Inline solutions will generally entail devices that have been purpose built to accommodate their top two distinguishing characteristics (i.e., beyond being directly in the flow of traffic). The first of these is the consolidation of policy decision and policy enforcement into a single platform. Eliminating the need for a chain of inter-device communications not only simplifies matters in general but also reduces the potential for encountering issues with the performance of end-user sessions.

The second characteristic is simply the presence of underlying capabilities that yield a high degree of visibility, control, and protection. For example, a good inline solution will incorporate and exhibit the following features and associated advantages, respectively:

- pre-admission protection mechanisms such as confirmation of user identity and, ideally, state of the client device (e.g., presence of anti-virus software), though the latter is less critical to this approach given the extent of other available capabilities;
- the ability to collect and correlate details of all end-user activity to support generation of appropriate, post-admission access control policies, and thereby reduce or even eliminate misuse of computing resources and corporate data;
- application-layer awareness to complement typical network-layer knowledge and thereby enable granular access control (as opposed to an all-or-nothing blocking scheme);
- direct control of the traffic stream to ensure reliable, consistent enforcement of policies, ideally via a range of mechanisms (e.g., blocking, segmentation); and
- advanced post-admission protection mechanisms, such as intrusion protection and content filtering, to granularly block malware while still allowing “good” traffic to proceed.

In addition, a good inline solution will also be characterized by:

- very good network coverage, as a result of the inline devices being positioned relatively close to the point where client machines are connecting;



- very good client coverage, as a result of being able to inherently see/control all client communications and, if supported, of being able to take advantage of client-based agents when they are present;
- reasonably good scalability, based on the ability of each inline device to address a large number of client connections; and
- support for one or more mechanisms for achieving high availability, to ensure that coverage which is otherwise very thorough can also be continuous.

Conclusions and Recommendations

It should be clear from this quick review that the inline approach, if executed properly, has significant advantages over the others. However, returning to an earlier point, the intent is not to imply that these four approaches are the only ways to achieve internal security or, for that matter, that they are mutually exclusive. After all, defense-in-depth will always be a best practice.

In addition, client devices are themselves resources that require protection, particularly given the increasing frequency with which they operate outside the boundaries of the corporate network. As a result, the best possible approach will most often be to complement an inline solution operating as the primary means of internal security with additional components that further address/bolster both endpoint and data center security.

Internal Security Approaches – Summary of Analysis

Approach	Client	Data Center	Network - Overlay	Network - In-Line
Scope of coverage	Low to medium	Low to medium	Low to medium	High
Scalability	Low	Medium	Medium to high	Medium to high
Degree of visibility and control	Medium to high	Medium to high	Low to medium	High
Degree of protection	Medium to high	Medium to high	Low to medium	High
Ease of implementation and operation	Medium	Medium	Low to medium	High