



White Paper

metagroup.com



800-945-META [6382]

January 2004

Securing Internal Networks: The Final Frontier

A META Group White Paper

“The proliferation of alternate paths into an organization, application-layer attacks, and devastating worms are all hammering home the conclusion that perimeter defenses must be complemented by a full range of internal security measures. Addressing this need will inevitably require implementing a combination of different types of security controls, though we expect products that are better tuned to the unique challenges of internal security will begin to emerge in 2004.”



METAGROUP

Introduction

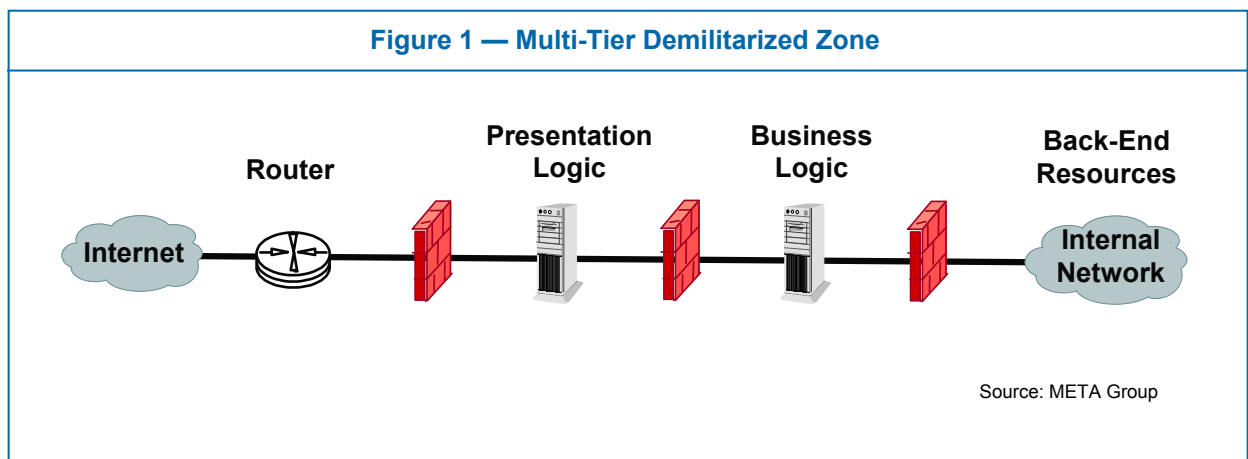
Despite significant investments in information security, organizations continue to be afflicted by cyberincidents. We believe a significant contributor to this condition is the relative lack of attention paid to securing internal networks and systems. This paper will explore numerous facets of that proposition, including: 1) observations on why addressing internal security is rapidly becoming a necessity; 2) enumeration of the various challenges that providing internal security entails; and 3) identification of the characteristics of an ideal solution.

Is Internal Security Really the Final Frontier?

Blasphemy! To security professionals, even remotely associating “final” and “information security” is a major transgression. After all, as they are fond of pointing out, information security is ideally an *ongoing* process, one that is executed by people and technology to fulfill and enforce a set of associated policies. However, given the minimal attention that internal security has been afforded to date, perhaps designating it as “final” is not entirely inaccurate.

Drawing a Distinction

Indeed, most organizations have focused their security strategies on establishing perimeter defenses. This can be characterized by the construction of multi-tiered demilitarized zones (DMZs — see Figure 1), typically incorporating controls such as firewalls, authentication services, intrusion detection systems, and antivirus scanners.



On the other hand, internal security involves implementing controls directly within the “back-end” environment. We will explore the details of what this entails in a later section of this paper. For now, it is important to realize that the goal in both

cases is actually the same (i.e., aside from the DMZ also being used to securely host externalized applications). Specifically, it is to protect the organization's most valuable computing resources, the majority of which, somewhat ironically, reside on the internal network.

Being in the Right Place at the Right Time

Given that they share a common goal, why then has internal security essentially been ignored, at least on a comparative basis? More importantly, is such treatment still appropriate?

Getting Started

It has been a long-standing perception among organizations that external threats are greater than internal threats. Whether this perception is accurate is the source of much debate, and even distraction. However, it really does not matter, because organizations also believe that reducing external threats is a more solvable problem. After all, with internal threats, the subject is one's own employees, and if they are intent on doing harm, how can they really be stopped?

The result has been to use whatever limited resources are available for information security to establish perimeter defenses. This is also consistent with our observation that only those 10%-20% of organizations with relatively mature security programs have managed to address internal security to a meaningful extent.

To be clear, the logic behind this historical approach makes perfectly good sense: given limited resources, get the most bang for a dollar. However, it should be equally clear that something is amiss.

Keeping Pace

Despite significant investments in perimeter security, successful attacks, as marked by the rate of acknowledged security incidents, continue to rise steadily in terms of both volume and impact. For insight on why this is the case, it is instructive to more closely examine the assumptions that underlie the effectiveness of a perimeter-oriented security strategy. Specifically, these are:

- That perimeter controls will consistently be completely effective
- That there are no alternate paths via which external threats can gain entrance to the environment
- That the threat from internal sources is indeed negligible

Rather than muddling through a heap of long-winded arguments, we will simply conclude on the basis of common sense that none of these has ever been a truly

safe assumption. However, more significant is that ongoing changes to the computing environment are further challenging their validity. The first two assumptions, in particular, are increasingly coming under fire.

Application-Layer Exploits

In recent years, application-layer attacks have become dramatically more common. Worms and blended threats such as Blaster, Bugbear, Slammer, and SoBig were among the most common and troublesome in 2003. In fact, more than half of the SANS/FBI top 20 threats to Internet security are categorized as application-layer attacks. This is unfortunate, because such attacks more easily evade commonly deployed perimeter controls, which historically have been focused on the network layer. Notably, such attacks also tend to evade those few controls that are applied within the internal network, such as antivirus scanners.

Proliferating Paths

Not so long ago the number of connections into an organization was relatively few, and the greatest challenge was identifying and eliminating rogue modems. However, falling prices for communications services and the evolution of various computing technologies have added dramatically to both the actual and potential paths into an organization's network. For example:

- Virtual private networking technology has facilitated a boom in terms of the number of connections that are maintained between business partners.
- Various pervasive computing technologies have been embraced to enhance productivity (e.g., wireless LANs, personal digital assistants with synchronization capabilities).
- Mobile and telecommuter solutions have seen widespread and increasing deployment on the basis of reducing operational costs and improving employee "quality of life."

Recent issues with mobile and guest users provide a good example that highlights the scope of the challenge at hand. Specifically, with Blaster, numerous META Group customers acknowledged that the worm did not penetrate their perimeter controls. Rather, their computing systems were thoroughly "taken down" after an otherwise mobile user visited a corporate office and connected an infected machine on the local-area network.

In any event, all these changes should help to crystallize two conclusions. First, drawing a distinction between an external threat and an internal threat is increasingly pointless. The source of a threat has simply become less relevant as network perimeters have become less well defined. The implication is that being effective from a security perspective will depend on establishing new/additional perimeters in closer proximity to the resources that are being protected. The

second conclusion is one that has always been true, but that should now be even more apparent. Specifically, perimeter-oriented security strategies, while at one time adequate, are not now and never will be sufficient.

In either case, the need for internal security is both obvious and growing. Yet, all too often, it is still being neglected. Our analysis indicates that a very legitimate and substantial reason for this is that implementing internal security is not a trivial matter. Unfortunately, it presents some unique challenges that tend to make achieving an effective “solution” both elusive and costly.

Security With Some Twists

One temptation when it comes to addressing internal security is to treat it identically to perimeter security. To some extent, this does indeed make sense. Many principles used to construct and operate perimeter security solutions apply equally well on the inside (e.g., defense in depth, hardening of critical systems). However, others do not. Having an out-of-band management network for the internal environment would simply not be practical from the perspective of cost and complexity.

The point is that the nature of the internal environment presents some unique challenges. A deceptively simple summary of these is that internal security basically involves more ... of everything:

- The scale of the environment requiring protection is significantly greater, involving numerous (sub)networks and potentially thousands of systems.
- The scope of the environment is significantly greater, involving a much wider variety of both business applications and underlying protocols — not just HTTP, FTP, SMTP, and the handful of others associated with the DMZ. Furthermore, many of these protocols are particularly weak. Their designers did not include many provisions for security under the assumption that the protocols would be sheltered as a result of operating only on internal networks.
- Along similar lines, there may also be the issue of having to manage more types of users or groups of users. With the external environment, different access control and authorization policies need to be established for public users and perhaps a half dozen other classifications of customer or partner. However, internally, the different roles can easily number in the hundreds, resulting in a much more complicated set of policies and controls.
- Finally, the internal network involves greater speeds or, more accurately, volumes of traffic. Internet connections and associated DMZ resources are rarely confronted with more than 45 Mbps, while internal networks and systems routinely operate at 2x-10x that throughput. As a result, any

controls that are implemented in the internal environment need to be capable of conducting the necessary inspections and dispositions at a much greater rate than we are typically accustomed to.

Clearly, any strategy, solution, or even product that is intended to provide internal security should ideally account for these items (see Figure 2).

Figure 2 — Comparing Perimeter Security and Internal Security

	Perimeter Security	Internal Security
Network Environment	<ul style="list-style-type: none"> ▲ 10's of systems to protect ▲ 10's of Mbps of traffic to inspect and mediate 	<ul style="list-style-type: none"> ▲ 1,000's of systems to protect ▲ 100's of Mbps of traffic to inspect and mediate
Application Environment	<ul style="list-style-type: none"> ▲ 10's of applications ▲ <10 protocols ▲ Standard/well-defined apps ▲ Stricter adherence to protocols ▲ Client-to-server applications 	<ul style="list-style-type: none"> ▲ 1,000's of applications ▲ >100 protocols (+ weaker) ▲ Homegrown applications ▲ Loose adherence to protocols ▲ Client-to-client applications
Management Environment	<ul style="list-style-type: none"> ▲ 10's of user groups to distinguish in terms of policy ▲ Block that which is unknown ▲ Typically centrally coordinated 	<ul style="list-style-type: none"> ▲ 100's of user roles/groups to distinguish ▲ Observe that which is unknown (do not disrupt) ▲ Typically locally coordinated

Source: META Group

Square Pegs and Round Holes

Considering all the relevant factors, it should be apparent that many conventional security products, while still applicable in most instances, are actually not well tuned to individually meeting the demands of internal security. For example:

- Antivirus products are limited by their dependency on foreknowledge of attack signatures. As a result, they are unable to assist, at least initially, with the worms that have so dramatically exposed in recent months the vulnerability of internal networks and systems.

- Intrusion detection systems have the same limitation, while emerging intrusion prevention systems, which attempt to overcome this condition by using various techniques, remain relatively immature.
- Most firewalls, as well as many other products, lack sufficient application coverage and performance capabilities.
- Switch and other network infrastructure-based products, while often excelling in terms of performance, typically lack the requisite visibility above the network layer.

Patching Not a Panacea

Another common reaction is to expect patch management to be a silver-bullet solution, essentially correcting all the vulnerabilities that internal systems have and thereby eliminating the need for any other controls. However, there are a significant number of fatal flaws to such an approach:

- Not all vulnerabilities involve design or coding errors; some are due to configuration errors. In those instances, patching is completely irrelevant, whereas other types of controls can still make a difference.
- Even when the vulnerabilities are due to coding/design errors, patches are not always available in time to make a difference. In the past two years, we have seen the period of time between the identification of a vulnerability and the development of an attack that exploits the vulnerability shrink dramatically. The result is often a hectic race, pitting software manufacturers against the hacking community.
- Even when there are patches for known vulnerabilities, implementing them can be far from a straightforward exercise. Patch quality has been poor with such great frequency in the past that time-consuming testing of patches is often deemed to be a critical prerequisite to deployment. In addition, patch management systems with sufficient degrees of automation as well as sufficient scope in terms of the systems and applications they can support are simply not available in the market today. Finally, the sheer volume of patches being issued complicates the entire process, also making it essentially impossible to keep up.

To be clear, our intent is not to dismiss patch management outright. On the contrary, we believe patch management will play a very significant role alongside numerous other components as part of a comprehensive information security solution, particularly as some of its current shortcomings are resolved.

Composing/Conducting an Internal Security Symphony

If patch management and the various conventional security products are not *the* answer when it comes to developing a solution to address internal security, then what is?

A Conventional Approach

As is often the case when it comes to information security, we believe that what is needed to properly address the challenges of securing internal networks is actually a blending, or combination, of many different controls. We break these into three distinct categories: fundamental, intermediate, and emerging.

Fundamental

At the core and essential to success of any solution are several fundamental controls. Without going into detail, these are physical security, user awareness, and antivirus scanners. We include antivirus measures in this category despite earlier declaring that they are not *the* answer because: 1) they are already quite ubiquitous; and 2) if properly maintained, they do validly contribute to the solution, by keeping a large collection of known exploits from being a nuisance, or perhaps even wreaking havoc on a recurring basis.

Intermediate

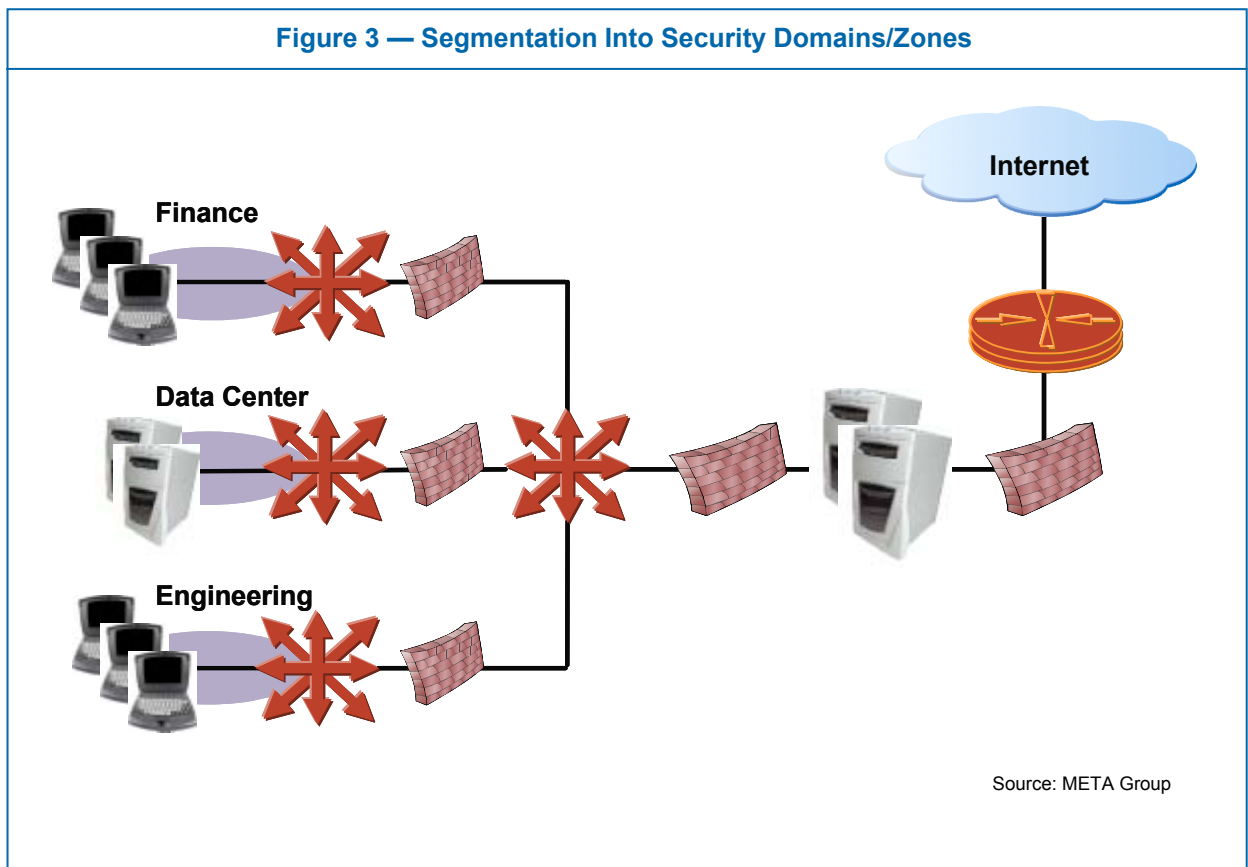
These measures are clearly a step above the fundamental ones, but there really is nothing special about them. They are all practically achievable and would not be considered aggressive or even particularly innovative. However, with the relative lack of investment in internal security, accomplishing the following measures would represent a better-than-average performance for an organization:

- Internal segmentation (e.g., using routers, switches, and virtual LAN technology) supports logically or physically separating resources that require different levels of security.
- Building off this, internal firewalling also facilitates segmentation, but does so with the added benefit of providing a more effective security barrier, where needed (see Figure 3).
- Operating system hardening is essential at least for the most critical application platforms.
- User administration involves explicitly provisioning which resources users have access to, versus just giving everyone access to everything.
- Finally, monitoring for suspicious activity within the internal environment is also important, and perhaps can be accomplished with judicious deployment of a few intrusion detection systems. However, we are compelled to offer a caution along with this suggestion, because the

operational cost of tuning and maintaining an intrusion detection system in an environment as highly diverse as most internal networks could be significant.

In general, we recommend that, at a minimum, regulated companies should be using all these controls, at least to some degree, by 2005. These include organizations that are subject to various privacy and security laws, such as the Health Insurance Portability and Accountability Act (HIPAA) and the Gramm-Leach-Bliley Act (GLBA) in the US, and perhaps the Privacy Directive for members of the European Union. A common component of many of these regulations is that impacted organizations must establish “comprehensive information security programs.” Thus, while we expect regulators to focus on well-known “basics” during their first rounds of audits against associated requirements, we also expect their investigations to focus soon thereafter on protective measures being taken for internal networks and systems.

Figure 3 — Segmentation Into Security Domains/Zones



Emerging

We characterize these additional measures for tackling internal security as ones that are likely to become best practices over time, though widespread adoption is still one to two years off. They include:

- *User life-cycle management:* This is a more comprehensive version of user administration that incorporates self-service capabilities and delegated administration. In general, it implies paying closer attention to the maintenance of user rights and privileges as their position and status within an organization change over time.
- *Personal and end-node firewalls:* The scope in this case is not just mobile users and their computing devices, but rather all users on the LAN.
- *Stronger user authentication:* Very likely, this will involve digital certificates in some manner, but the point is that reliance on simple username and password combinations is insufficient in the long run.
- *Fine-grain user/behavioral auditing and monitoring:* The ability to monitor object access and operations, along with enhanced capabilities for interpreting the data, will help detect and deter misuse.
- *Domain structuring and other trust management concepts:* In combination, these essentially yield a significantly enhanced version of segmentation. It is an entire methodology for identifying and establishing boundaries between collections of resources, while also identifying and establishing the appropriate controls that should be implemented within each collection. Perhaps most importantly, it is an approach that inherently incorporates and therefore ensures buy-in and participation by business-unit managers.

A Best-Fit Approach

The need to blend and combine multiple controls to achieve a meaningful degree of security is likely to never change. What does and should change, however, is the nature of these controls, or more accurately the products that represent them. In other words, stitching together a range of conventional products to achieve internal security is only one approach — albeit perhaps the most expedient and practical one available today. However, organizations should also be on the lookout for solutions that are better oriented toward the specific challenges of the internal environment. We expect such products, exhibiting at a minimum enhanced security, performance, and management capabilities, to emerge in the market beginning in 2004.

Enhanced Security

This involves four distinct capabilities, but consistent with the concept of “solution,” they should ideally all be available in combination as part of a single product.

Indeed, they should even be integrated, essentially acting as services to each other and thereby increasing the effectiveness of the overall solution:

- Attack protection is the ability to stop both known and unknown threats. The primary goal is to thwart worms and other malicious code, which have become so troublesome of late. In this regard, it essentially aligns with how intrusion prevention systems are being marketed today. However, our expectation is that it would not be a standalone capability, but rather a component of a broader solution that also incorporates the other services identified below. When considering this capability, we emphasize the need to evaluate two significant characteristics. The first is the presence of multiple, varying algorithms and mechanisms, which are essential for detecting suspicious activity in the first place, and then also for accurately classifying it as an attack (versus legitimate traffic). The second item is the existence and comprehensiveness of related research activities by the vendor of the product. This is absolutely essential to maintaining and even improving effectiveness of the capability.
- Segmentation and compartmentalization are actually a pair of closely related capabilities. The former entails being able to separate classes of resources into separate enclaves, and then mediating the flow of traffic between them. Compartmentalization is very similar in nature, but is focused more on stopping the spread of “infections.” Ideally, it enhances the permanent though semitransparent boundaries associated with segmentation by also enabling the dynamic establishment of conclusive boundaries (i.e., ones where no traffic is allowed to pass). In addition, these boundaries may be more targeted, pertaining to groups of machines or even individual machines within the predefined “segments” or enclaves. Furthermore, it would ideally be linked to the aforementioned attack protection capability, which would provide it with the requisite intelligence to operate in an automated fashion.
- Application awareness and control involve the ability to protect communications and computing resources on the basis of application-layer information, not simply network-layer details. As alluded to previously, this degree of granularity is increasingly essential because exploits are already tending to operate at this higher layer, rendering solutions with only network-layer visibility ineffective. However, it is also important when evaluating this capability to keep in mind that the scope of applications being used in the internal environment is much greater than the few that commonly transit perimeter defenses. So, this is no trivial capability, either to develop or to operate. It should ideally be extensible and even customizable, providing the opportunity for vendors and users alike to incorporate broader application coverage. Furthermore, any initial coverage

- should focus both on those applications that are most common and on those that are most vulnerable.
- The final item is also a matched pair of capabilities: admission control and quarantining. The idea is to have the ability to check the status of end-node stations with regard to security policies and settings before letting them connect to the network and communicate with any of the internal resources. Secondly, the quarantining part entails the ability to provide limited access by nodes that fail these checks to selected services, including the resources that can be used to bring them into compliance.

Enhanced Performance

The need for greater performance or, more accurately, greater throughput ratings has already been explained. In most instances, we expect this will be accomplished by better tuning/matching software and hardware (e.g., by developing custom drivers, by implementing reduced instruction sets) or by incorporating specialized hardware (e.g., ASICs, various add-on accelerator chips). However, it could also be accomplished, at least to some degree, by utilizing alternative or advanced techniques for implementing policies and/or inspecting communication sessions. For example, a dramatically simpler rule set will typically reduce the operations that a filtering device will have to execute.

Enhanced Management

This is arguably the most important area where change from conventional products should be sought and evaluated. First and foremost, support for scalable and flexible management of policies and configuration settings is required. After all, the diversity of the internal environment implies that a much greater quantity of “objects” will need to be accounted for. One potential way to facilitate this is to enable “default allow” policies, as an optional alternative to the more familiar “default deny” arrangements associated with perimeter defenses. This may seem odd at first, but is in fact an acceptable and even appropriate practice when: 1) the security filters are located closer to the users and further from the resources; 2) it is important to not prevent legitimate communication sessions from occurring; and 3) the associated security filters are not necessarily familiar with each of the hundreds or even thousands of applications being used.

For similar reasons, it is also important to have support for delegated administration. This enables those persons most familiar with a “local” set of resources (i.e., users, systems, and applications) to establish and tune the associated security rules. Of somewhat less importance, but still valuable, would be the ability to integrate with the same management applications used for the perimeter environment — at least for event management purposes.

Additional Criteria

Another desirable characteristic is that components of any solution should remain transparent to the internal network, at least from the perspective of not requiring significant changes in terms of addressing, routing, or physical arrangement.

Last, but certainly not least in terms of importance, ubiquitous coverage needs to be achievable at an affordable price. Solutions that require dozens of units/components at an aggregate cost of hundreds of thousands or even millions of dollars are simply not going to be adopted, at least not on a widespread basis.

Bottom Line

The assumptions that underlie the effectiveness of a perimeter-oriented security strategy are steadily being eroded. The proliferation of alternate paths into an organization, application-layer attacks, and devastating worms are all hammering home the conclusion that perimeter defenses must be complemented by a full range of internal security measures. Addressing this need will inevitably require implementing a combination of different types of security controls, though we expect products that are better tuned to the unique challenges of internal security will begin to emerge in 2004. Taking advantage of such products will be instrumental to economically achieving an effective internal security solution.

Mark Bouchard is a senior program director with Security & Risk Strategies, a META Group advisory service. For additional information on this topic or other META Group offerings, contact info@metagroup.com.



About META Group

Return On IntelligenceSM

META Group is a leading provider of information technology research, advisory services, and strategic consulting. Delivering objective and actionable guidance, META Group's experienced analysts and consultants are trusted advisors to IT and business executives around the world. Our unique collaborative models and dedicated customer service help clients be more efficient, effective, and timely in their use of IT to achieve their business goals. Visit metagroup.com for more details on our high-value approach.

