metagroup.com • 800-945-META [6382]

May 2004

Designing Security Domains to Manage Outbreak Risk

A META Group White Paper

"Securing only the perimeter of an organization is no longer sufficient in light of innumerable product vulnerabilities and the increasing complexity of enterprise networks and business inter-relationships. Security domains and trust levels help organizations provide more effective and targeted security by isolating resources according to the level of trust that is required by the business."

Contents

The Concepts of Security Domains and Trust Levels

Information security historically has been provided by creating a perimeter to isolate the organization from the outside world. This single-domain approach, developed over the past 30 years, has been successful in many organizations. However, this perimeter approach not only overlooks the possibility of internal threats, but also provides an attractive target for any intruder or malware code (e.g., viruses, worms) that can breach the external perimeter. Moreover, the complexity of large organizations and the volume of product vulnerabilities have resulted in numerous potential backdoors. Demand for an improved approach to designing secure external and internal access to information assets has resulted in rising awareness of "defense in depth" design principles. Two inter-related concepts that help coalesce a defense-in-depth strategy are trust levels and security domains.

Trust Levels

In its broadest sense, trust is simply the degree to which an organization can rely on an information system to be trustworthy in protecting confidentiality, integrity, and availability of its data and processes. Business organizations need to trust the information technology systems they use, but trust is not binary. Across the spectrum of no trust at all and complete trust, there are many possible intermediate points. Indeed, segmentation of this spectrum is facilitated by the concept of trust levels, which in turn represent relative levels of security.

META Group defines five trust levels:

- ***Trust level 1:*** For low security needs
- ***Trust level 2:*** For normal, but non-critical business uses
- ***Trust level 3:*** For critical or confidential business use
- ***Trust level 4:*** For tightly controlled or legally regulated use
- ***Trust level 0:*** For anything not meeting the tests required to be at trust level 1

Smaller organizations may need only one or two trust levels, while larger organizations may need many more. Trust levels are instantiated by applying generally accepted security practices and infrastructure to provide the requisite security for that level of trust. The amount of trust a business organization requires its IT systems to deliver must be defined by the business, with advice and guidance from the IT organization and the security organization. The security organization's responsibility is to develop policy and procedures, deploy appropriately configured security infrastructure to deliver identifiable levels of security, and explain to the business the risks and threats that generally are being considered.

Security Domains

Security domain structuring is an approach used to segment existing infrastructure into logical zones based on a common trust level. A security domain could be an isolated subset of the network, together with all the computing resources attached to that subset. Network isolation is provided through network configuration (e.g., virtual local-area networks [VLANs]) and/or internal firewalls, while the level of security results from implementation of the policies, processes, and security technology deployed within a domain as well as the isolation boundary that defines the domain edges. Most organizations are familiar with at least one security domain, the DMZ (demilitarized zone — see Figure 1).

The process of defining domains is an art, and this may require several iterations before the security manager can devise a practical domain structure. Many enterprise resources can occupy numerous, sometimes overlapping, domains. The security organization must decide which structure works best for its given organization.

Figure 1 — Multi-Tier Demilitarized Zone

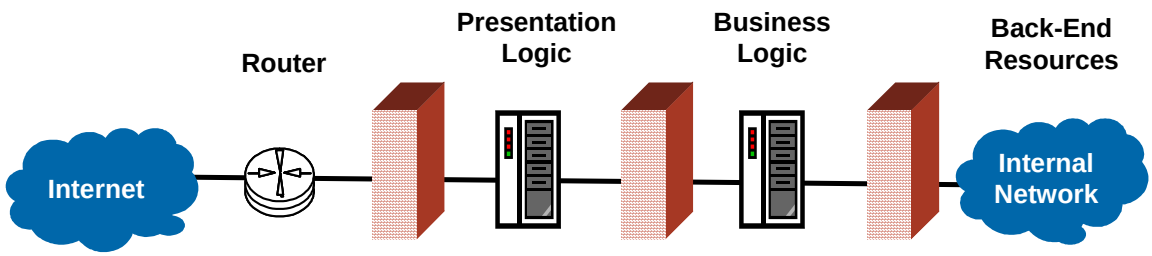

Source: META Group

Using Security Domains to Manage Outbreak Risk

The intelligence of attackers and malicious code developers as well as the increasing level of code vulnerabilities ensure that virtually any level of security eventually can be penetrated. Domain structures provide a series of gates, frequently with different detection and compliance mechanisms at each domain perimeter. The effort required to compromise each layer of security adds complexity and potential delay to the attack, which provides opportunity for alternating security infrastructure to trip up an attack, or provides time for the security team to detect and respond to the attack before damage is done. Moreover, leveraging common and adaptive technology services aids in reducing organizational complexity, which in turn improves security.

Segregating infrastructure into domains allows for increased resiliency in the face of an attack by confining an attack or infection to a section of the network, enabling business to continue in other domains. Worms such as SQL Slammer generate so much spurious traffic as they attempt to replicate that the network becomes unusable. Segmenting the network not only limits the spread of the attack, but also confines the denial-of-service (DOS) effect.

Figure 2 — Benefits of Security Domains and Trust Modeling

- Improved security
- More resiliency to attacks
- Improved scalability of security infrastructure
- Better communication and alignment of business, IT, and security
- Higher accuracy in estimating costs
- Reduced organizational complexity and increased skill reuse
- Potentially reduced security cost

Source: META Group

Security technology is among the least mature technology in use by information technology organizations and it does not scale well. Use of domains can result in localized implementations that do not suffer from scale-driven failures.

Domain structures allow targeted investment in security. Hardening all computer systems and components to the level of security required for the most valuable information asset in the organization would be prohibitively expensive. Security domains assist the security organization in matching the amount of security provided to the needs of the business through localized policy, process, and technology deployment. Moreover, by combining a finite set of security solutions with reusable technology components and skill sets, IT organizations are better able to accurately predict costs and the time needed to deliver secure solutions.

In order to determine the appropriate level of security to be implemented, it is essential to solicit input from various groups within the business and the IT organization. However, the complexity of the issue and excess technology terminology hamper effective communication. As a result, security often has been perceived as an inhibitor to business plans, rather than an enabler. Trust levels and security domains provide a common language that can be used by business and the IT organization to communicate more effectively. These concepts facilitate more effective discussion of security requirements and design considerations/tradeoffs with both the business and the various fractions within the IT organization, thereby

easing decisions about appropriate infrastructure, application security, and operational process.

Choosing an Effective Domain Structure

Determining appropriate security domain structure is as much an art as it is a science. In general, domain boundaries serve two purposes: they help determine defensible perimeters for the application of technology, and they determine the scope of security policies in manageable slices. Any grouping must keep these goals in mind.

There will likely be as many domain structures (see Figure 3) as there are IT organizations, yet most fall into one of the following six categories, or some combination or variation thereof:

- ***Geographic domains:*** A geographic domain consists of all the resources within one geographic entity, such as a building. Geographic domains are useful for single-purpose buildings, such as retail stores and small branch offices, or for federated entities such as business partners, government agencies, or customers. Wide-area network (WAN) connectivity provides straightforward domain boundaries between geographic domains.
- ***Organizational domains:*** Resources that can be easily assigned to organizational units (e.g., sales, finance, R&D, customer service, manufacturing) can be grouped into organizational domains.
- ***Administrative domains:*** These domains are defined by the administration of systems, technologies, and resources (e.g., database, server, desktop, mainframe). This type of domain represents tacit acknowledgement of existing siloed administration.
- ***Resource-based domains:*** Domain boundaries for resource-based domains are drawn by organizing the resources according to either *security class* or *application* (e.g., ERP, CRM, financial, office tools). Difficulty will arise if a computing system with a wide variety of resources belongs in multiple security classes.
- ***Technology-based domains:*** Similar to their use of administrative domains, enterprises commonly use a domain scheme based on technologies, with Unix, Windows, the data center, WiFi network, TCP/IP network, and the Internet each representing a domain.
- ***Life-cycle-based domains:*** These domains are based on product life cycles and are useful for production systems. However, they are somewhat more difficult to implement due to the domain boundaries being at such a high level.

Typically an application resides in one domain along with most or all of the end nodes and users accessing that application. Some applications (e.g., directory

services) reside in a “common services” domain, while the backbone network may be viewed as a separate domain for security purposes. One domain may be nested inside another, with the innermost domain running at a higher trust level than the outer domain (this is typically referred to as a “zoned” network), but the two must be isolated from one another.

A comprehensive domain scheme can and often does include different types of domains. The strengths of each domain scheme in doing an effective and useful job should be examined in terms of the natural boundaries of an organization, keeping in mind the original goals of identifying defensible perimeters and making the scope of security policies more manageable.

Figure 3 — Domain Structure

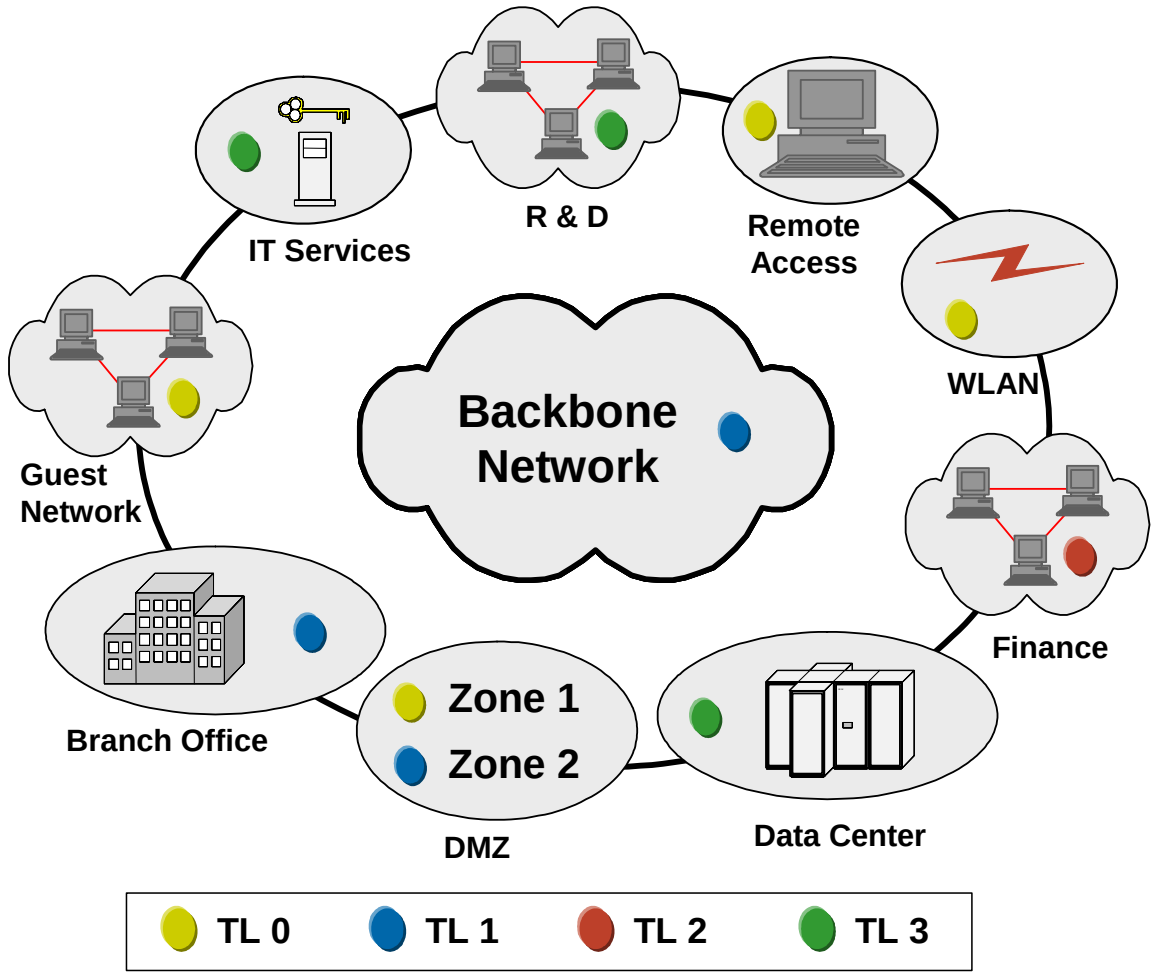

TL — Trust Level

Source: META Group

Implementing Security Domains

Organizations embarking on domain structuring should not attempt to map the entire infrastructure into domains initially. Smaller projects are required to understand what works best for the organization and to gain some early successes, which can demonstrate the value of the method to the business and the IT organization. Therefore, the first step is to identify the critical or significant resources within the organization. This list does not have to be complete; it can be refined later, but it should capture the obvious major components. The next step is to define a trust level for each domain. Of course, this implies that trust levels have been defined; if not, that must occur at this time.

Having identified the resources in an enterprise and organized them into domains, the security manager's next step is to determine the current state of security within the enterprise:

- How secure are resources now?
- Does the enterprise have policies governing their security?
- If so, are they being followed?

Next, a baseline of current security policy must be developed and the current state of security for resources within the enterprise assessed. For most organizations much of the work is already done for at least one nested domain — the DMZ. The DMZ is made up of two or more trust levels, to support an orderly transition from Level 0, the Internet, to Level X (i.e., the base level associated with the organization's internal network). A good first step in this area is to formalize the DMZ security domains and map the existing security procedures, policy, and technology into the definition of trust levels. After that, priority should be given to high-risk/high-value infrastructure to which it is relatively easy to apply domain principles, working down from there to low-risk/complex domains.

Figure 4 — 10 Steps to Creating a Security Domain Structure

1. Assess the current state of information security and network architecture
2. Identify critical/significant organizational resources
3. Group resources into security domains
4. Develop security policy to instantiate trust levels
5. Assign a trust level to each domain
6. Perform gap analysis and develop a project plan
7. Prioritize domain projects
8. Assign security administration responsibilities
9. Implement
10. Revise

Source: META Group

In most organizations the security group does much of the design work, creating policy and negotiating with the business the trust level of domains. Implementation will likely be conducted by a combination of infrastructure developers, security planners, and operations personnel. Clear roles and responsibilities must be delineated not only for the domain implementation project, but also for the ongoing administration.

Isolation Services

Fundamental to the nature of domains are the logical or physical choke points that serve as boundaries to other domains, a location that also enables enforcement of security policy to contain any malicious behavior. Therefore, a significant amount of effort will be involved in creating a set of isolation services that can be used at these locations. Isolation services are network- and systems-based tools and processes that enforce access control policies. Their primary function is segmentation and mediation of information flows between domains.

Mapping the domain structure on the physical network will likely be the biggest challenge for organizations, since different security domains with different trust-level requirements may exist on the same physical network. Using physical network separation should be reserved for the highest trust zones. Lower trust zones might use virtual LANs.

The type of cross-domain isolation services used will depend on the trust level of each domain and the primary risk and/or threat to that domain. The primary mechanism for cross-domain isolation will be firewalls. Enterprise firewalls base access-control decisions predominantly on packet header details (e.g., network-layer information). Blocking unwanted or ill-formatted communications in this manner is a useful service, but it may be insufficient in the face of application-layer attacks, which are detectable only through inspection of packet payloads and fully reassembled sessions. As a result, specific threats may require additional complementary security filters.

For example, the primary risk to the “research and development” (R&D) domain in our domain structure diagram (as shown in Figure 3) might be theft of proprietary information, which would indicate the predominant threat would be malicious hacking or insider theft. For the “sales” domain, the primary risk to the sales group might be the loss of productivity resulting from the ever-prevalent virus threat. Different isolation technologies (e.g., firewalls, VLANs) and security infrastructure (e.g., intrusion detection/protection, certificates, log analysis) offer protection for different types of risks or threats.

Network Worms: A Special Class of Virus

Network worms such as Slammer (see Figure 5) are a special class of virus threat that is not addressed adequately by existing host-based antivirus products. A "network worm" is a self-contained program (or set of programs) that is able to spread functional copies of itself or its segments to other computer systems by exploiting network connections. These viruses sometimes contain a destructive payload, but successful ones almost always create a denial-of-service effect that can impact business productivity or prevent business revenue. Domains where the risk of DOS or data loss is unacceptable must take special precautions against this threat.

Figure 5 — The Slammer Worm

The Slammer (Sapphire) Worm was the fastest computer worm yet. After its release on the Internet, it doubled in size every 8.5 seconds. Within only 10 minutes, it had already infected more than 90% of vulnerable hosts running Microsoft SQL. Although the worm did not have a malicious payload, network traffic resulting from its propagation routine rapidly saturated networks, bringing all work to a halt and blocking access to critical server resources, including airline reservation systems ATM and Internet sites.

Firewalls were effective in stopping Slammer by filtering UDP packets with a destination port of 1434. However, if the worm had exploited a vulnerability in a commonly used service (e.g., DNS at UDP port 53, HTTP at TCP port 80), such filtering could have caused significant disruption to legitimate traffic, with resulting denial-of-service more harmful than the worm itself. Moreover, once inside a network without isolation services, the worm was free to propagate until all SQL servers were taken offline and patched.

Source: META Group

Why Existing Infrastructure Is Insufficient

Network worm propagation techniques cause heavy traffic that typically saturates networks and prevents access to critical resources. Without internal isolation services or domain structure, it is impossible to stop the spread of the virus or contain its spurious DOS effect. Even with appropriate domain structuring, stopping the network worm at domain boundaries is difficult with existing tools and techniques.

Firewalls

Typical isolation firewalls enable routing policy to restrict users and applications to specific segments, ports, and protocols according to a relatively static policy. However, to a firewall, worm-infected machines may look like trusted computers/users using an allowed protocol/port combination. Some more advanced

firewalls offer application awareness and more granular control, enabling access rules based on a more thorough understanding of the protocols being processed. In addition, these firewalls can potentially check for a limited range of application-layer attacks. However, the application-layer capability of firewalls varies widely. The number of application protocols a firewall monitors may be limited, and the number of attack signatures is also not necessarily a complete set. Moreover, the ability of firewall vendors to rapidly assess new application vulnerabilities or exploits and distribute updates is limited.

Consequently, cross-domain firewalls have acted more like gates, or drawbridges, during recent worm outbreaks, providing a manual way to shut off traffic between domains to contain the spread of worms. But closing ports may choke off legitimate business traffic, and manually closing ports (assuming worm DOS traffic does not prevent management access to routing or firewalls) may be too slow to contain the spread. By the time policy changes can be made, the worm may have already propagated to other domains.

Antivirus Solutions

Network worms do not necessarily have a file component that can be identified by traditional host and perimeter antivirus software. Furthermore, traditional antivirus solutions rely heavily on post-outbreak analysis of virus characteristics to develop a unique virus signature that can be used by scan engines to identify viruses. The speed of worm propagation makes this type of reactive solution ineffective.

Intrusion Detection Systems

Intrusion detection systems (IDSs) are designed to monitor and log abnormal network behavior, but do not necessarily stop such traffic. Newer intrusion prevention systems (IPSs) are more adept at stopping positively identified threats and updating signatures more rapidly than are application-layer firewalls. Although it is increasingly common for IPSs to deploy new attack signatures based on evaluation of a new vulnerability or reactive signatures for specific exploits, reaction times of vendors vary significantly. Although IPSs can drop threat traffic, preventing worm proliferation, they cannot yet automatically repair infected systems.

Patching

Most worms attack known vulnerabilities for which a patch is available, indicating that patch deployment should be a priority. During the past two years, however, the period of time between identification of vulnerability and development of an exploit has diminished dramatically. The result often is a hectic race between the IT organizations getting, testing, and distributing patches to potentially hundreds of computers and the hacking community devising an exploit. 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, making it difficult for IT organizations to prioritize and keep current.

What Is Needed

Certainly, the aforementioned techniques are not useless, but a defense-in-depth strategy requires multiple layers. Indeed, we see these layers coalescing and overlapping in various product forms in the future. Yet to prevent a network worm infection, more protection is needed. Ideally, network worm protection would completely eliminate the risk of infection. However, realistic solutions will allow for the possibility of imperfect protection and include incident-response and clean-up tools (see Figure 6).

Figure 6 — Essential Elements of Network Worm Defense

- Perimeter isolation filtering based on: a) known exploits; and b) known vulnerabilities
- Rogue node and network connection detection and remediation
- Vulnerability analysis with practical mitigation advice correlated to asset information
- Identification, isolation, repair, and patching of infected nodes

Source: META Group

Protection

Effective network worm protection requires both automatic filtering at domain boundaries to prevent worms from penetrating the organization and implementation of proper domain structuring to contain worm outbreaks. Worm filters must identify and filter not only existing known exploit traffic, but also potential attack traffic — while allowing legitimate business traffic to continue to flow. Identification of all potential attacks is likely to be too slow of a process to implement on high-speed LAN networks.

A more practical approach would be monitoring for a set of potential attacks, targeting the existing universe of known current vulnerabilities. A combination of threat-specific attack signatures to prevent infection by known worms and a targeted set of potential-attack signatures to identify and filter zero-hour attacks is necessary. Correspondingly, the vendor of such a solution must have the resources to monitor global events, the technical expertise to analyze new vulnerabilities and exploits, and the distribution infrastructure to rapidly deploy advice and signatures to global clients.

In addition to automated policy filtering, a comprehensive solution would include vulnerability notification with a severity rating and some practical policy advice on how to minimize the risk of known vulnerabilities (e.g., filtering UDP packets with a destination port of 1434, turning off unnecessary SQL services). Linkage between system inventory and vulnerability reporting is also desirable. Instant visual indicators of the quantity and location/domain of vulnerable systems would enable a response that is correlated with the risk to the enterprise.

Closing the Back Door

Clearly, any worm defense solution must ensure that not only the domain perimeter is secured but also that each end node in the domain is not a potential backdoor for worms to traverse. In the past, this required auditing for rogue dial-up modems and, more recently, WiFi wireless access points. Today, with the increased deployment of laptop and other mobile computers, the node itself may become infected when it travels outside the domain and then carry that infection back when it rejoins the network. It is also common for a "shadow" IT organization to deploy rogue servers to fulfill departmental needs. These rogue servers should be automatically inspected to ensure conformity to security policy before being admitted to the network. Endpoint admission control (EAC) validates the security compliance (e.g., patch levels, antivirus update level, security software installed, security process running) on PCs and servers before granting network share rights, and ideally enables secure remediation rather than denying access outright.

Repair and Remediation

Finally, if a domain is infected with a network worm, incident response teams must rapidly identify the infected machines, take them offline, repair worm damage, and patch them before putting them back online. During worm outbreaks, the sheer volume of worm traffic can make remote identification of infected clients difficult. Incident response teams historically have assumed all clients are infected and have shut down network until remediation is complete, extending the DOS effect of the attack. Domain structures with appropriate isolation services enable networks to be brought back online, one domain at a time. Still, manually repairing and rebooting nodes in a network is time consuming and costly. Moreover, if repair is not 100% complete, re-infection is likely until systems have been patched. Effective worm protection would ideally include a means to remotely identify and automatically repair and patch infected machines.

True network worms are relatively rare, yet many of the outlined techniques and tools for dealing with network worms are applicable to a range of security threats. Therefore, preparing for the next Slammer Worm will also put the organization in good stead to defend against a number of less-tricky threats.

Bottom Line

Securing only the perimeter of an organization is no longer sufficient in light of innumerable product vulnerabilities and the increasing complexity of enterprise networks and business inter-relationships. Use of security domains and trust levels helps organizations provide more effective and targeted security by isolating resources according to the level of trust that is required by the business.

Network worms represent a special class of virus that is not adequately addressed by existing security technology and procedures. Domain structuring with appropriate cross-domain filtering helps isolate network segments, enabling business to continue even while sections of the network are under attack. An ideal domain-filtering solution for network worms will stop worm traffic based on both exploit signature and the targeted potential exploit. In addition, solutions should also protect mobile nodes and other potential backdoors, such as rogue network connections. Vulnerability analysis should be correlated with asset tracking information to enable rapid analysis of the severity of the risk as well as practical advice to minimize the threat. Finally, solutions should also include remote identification of infected nodes and provide repair and patching assistance to rapidly restore infected domains.

Authors: Peter Firstbrook is a program director with Security & Risk Strategies, a META Group advisory service. This white paper includes contributions from Christian Byrnes, vice president and director; Mark Bouchard, senior program director; and Michael Warrilow, research analyst.

For additional information on this topic or other META Group offerings, contact info@metagroup.com.

About META Group

Return On Intelligence[SM]

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.