

# Evaluating Unified Threat Management Products for Enterprise Networks

*Joel Snyder • Opus One*

## Table of Contents

Overview .....	1
Performance Requirements .....	1
Integration Requirements .....	3
Consolidation Support .....	5
Enterprise-class UTM features: Not a simple checklist .....	6
Platform Extensibility and Flexibility .....	8
Management Requirements .....	10
Conclusion .....	11

## Overview

The term Unified Threat Management (UTM) has as many meanings as there are products that carry the label. While UTM has primarily focused on the small- and medium-sized business (SMB) network, products are coming to market aimed at the enterprise. This white paper will help you understand the specific issues enterprises need to consider when assessing UTM products, and offer guidance on evaluation criteria for enterprise-class UTM.

**At its core, UTM brings together three main ideas: multiple security features, integrated on the basis of a mature firewall, deployed in an appliance form-factor.**

The intuitive appeal of UTM is obvious: why have two (or three or four) boxes perform separate functions, when a single box will do? As security threats to corporate networks have increased at an alarming rate, the number of devices combating these threats has grown at nearly the same speed. However, at some predictable point, it's not feasible to have every new threat addressed by its own dedicated device.

The reasoning behind UTM has resonated strongly with managers commanding SMB networks, as UTM firewalls — called such because the firewall is the undisputed lynchpin of the UTM product — have quickly become a standard offering from every vendor. In this market space, UTM firewalls, with combined features including anti-virus protection and intrusion prevention built into the same appliance, reduce costs and simplify configuration.

UTM products in larger enterprise networks are not an easy sell primarily because most UTM products are indeed aimed directly at the SMB environment and enterprise network and security managers haven't had reason to view them as appropriate parts of their security strategy. Fortunately for the higher end, this product deficit is quickly changing as enterprise-class firewall vendors are adding UTM features to their product lines.

Obviously, evaluation and design criteria for UTM in enterprise networks must be very different from those of SMB networks. When UTM concepts are brought to bear on large networks, in ways appropriate to those networks, they offer the network and/or security architects tremendous flexibility to control and mitigate the risks associated with security vulnerabilities.

Because UTM, in general, and especially UTM in enterprise networks is new, network managers need a framework to evaluate products and match them to enterprise requirements. We will now explore six unique issues for network and security architects to consider that should be addressed for any enterprise-sized deployment of UTM.

## Performance Requirements

The pivotal selection criterion for any security product is performance. As enterprise networks have become absolutely business critical. Poor performance, low throughput, high latency, or dropped packets caused by improperly sized security products, are completely unacceptable.

UTM architectures are especially vulnerable to the question of performance because measuring and reporting traditional metrics such as goodput (often called “throughput”) is a developing art rather than an agreed-upon science. In conventional performance measurement exercises conducted on traditional security devices such as firewalls, the metrics of connection rate, connection capacity, and goodput are easy to measure and report. For example, in measuring firewall performance, it's generally sufficient to hook well-known measurement tools such as Spirent® Communications' Avalanche™ or Reflector™ appliances on either side of a device, spin the dials to generate simulated traffic, and interpret the metrics for your own network. Vendors are happy to provide this information — and you can trust the results — because there are only a few scenarios to test and the tools are readily available.

With UTM, system performance is dependent on which features are enabled and how those features are configured. For example, turning on anti-virus scanning in a UTM device will slow down performance. Scanning both e-mail and HTTP traffic for viruses will slow down performance more than simply scanning Web traffic.

To further complicate the situation, performance will vary depending on what the actual data moving through the system comprise. If your e-mail includes numerous attachments, the UTM device will have to work hard. If most of those attachments are in compressed archives, such as ZIP files, the UTM device will have to work even harder. If many of your e-mail messages are in Japanese (or any language with double-byte Unicode representation), the virus scanner will have to work harder than if they were in English (or any language with single-byte characters).

**Enterprise network managers considering UTM devices should key in on products that have the ability to scale performance without requiring forklift upgrades.** A product design that accommodates scalability with the ability to drop in additional processor cards or accelerators, or to change out a processor card on the fly, is certainly an outstanding starting point.

However, the best products go further and optimize these modular hardware upgrades by using features such as internal load balancing. This is valuable because every UTM feature (such as anti-virus or intrusion prevention) has a different set of performance characteristics and simply turning on a feature will not cause a linear increase in load. For example, suppose enabling anti-virus features triples your CPU load — not an unreasonable assumption. But if your only option is to add a single anti-virus accelerator or dedicated CPU module, that limited expansion capability puts a ceiling on performance. However, if you have the option of adding two CPU modules, with load balancing between them, then the anti-virus load would minimally affect total system performance.

## Testing with Tools

*Test tools can be used to provide performance measures, as long as you are aware that actual performance may vary from the actual flows on a real network. When testing UTM devices, it helps to be as real as possible.*

*Based on our past history of testing UTM devices, keep in mind these three points when setting up the traffic mix to be pushed by these synthetic traffic generators:*

- *Every device will undergo a fairly constant (although often low-level) set of attacks from the Internet, and these need to be a part of any test. We have found that alerting systems and forensics databases may not scale with acceptable performance as they get filled up with data over a period of days or weeks.*
- *Traffic performance is based on the actual data, meaning you need to send as close to a 'real' data stream through the device as possible. You will need to send real threats, such as viruses (and not just the EICAR test virus, which is often artificially optimized for super-high performance in anti-virus products), in approximately the same proportion as your own data stream. For example, most enterprises have about a 1% rate of viruses in email once spam is taken out.*
- *Testing should be run to determine the performance effect of each independent variable, as well as the system as a whole. For example, if you are looking to add anti-virus and intrusion prevention, run three tests: anti-virus alone, intrusion prevention alone, and both combined.*

When evaluating performance, enterprise network managers will hit yet another complication because there is no agreed-upon testing methodology for UTM devices. The traditional metrics, such as connection rate, goodput, latency, and connection limits, are all valid, but vendors tend to play up their bigger numbers first, not really supplying enough detail to make an apples-to-apples comparison, and not offering enough scenarios to handle the variety of security configurations a UTM device might be expected to support. Without common methodologies and tools to test devices, **the only real way to test UTM devices is to put them into real networks and run them with real data passing through them.** (See “Testing with Tools” for some hints on how to test performance of UTM devices.)

#### KEY EVALUATION CRITERIA

- What is true device performance along the four common metrics (connection rate, simultaneous connection limit, goodput, and latency)? Performance must be measured with the actual features you want to use enabled, with the configuration you need, and across your real traffic flow.

## Integration Requirements

Enterprise-class UTM devices need to support the complicated network topologies present in larger corporations. Four key points of integration that require support include interface flexibility, dynamic routing, high availability, and scalability.

Based on their SMB roots, UTM devices have traditionally sat at the perimeter of the network, replacing an edge firewall. However, in enterprise networks, firewalls are scattered throughout the network to harden and protect it from both external and internal threats. Enterprise-class UTM devices need to offer flexibility to work both at the edge and deep

within the network. When a UTM device is properly designed, it becomes equally useful no matter where it is placed. For example, while an edge device may need only two or three interfaces (“inside” and “outside”), an internal firewall will need a much higher interface count (one for each server group) as well as VLAN capabilities, to support as many security zones as necessary.

Network managers of larger networks use dynamic routing protocols to simplify overall configurations and provide more robust service in the face of topology changes and service outages. Enterprise-class UTM devices must integrate with existing routing fabric and support common enterprise routing protocols, such as OSPF. When considering a UTM device for an enterprise network, support of the network’s native routing protocols, interface types (such as fiber or copper gigabit interfaces with VLAN capabilities), and scalability requirements (such as integration with an internal or external load balancer) is critical.

In addition to routing and interface flexibility, any critical network resource, such as a firewall acting as a choke point between network zones — a likely point of deployment for a UTM device — must be engineered for both availability in the face of component failure and scalability in the inevitable event of increasing loads.

High availability can mean many things, but the simplistic, core requirement here is that the failure of any part of a UTM device — whether hardware or software — should not interrupt the flow of traffic through the network. At the same time, availability brings the second element of scalability into the UTM picture. As large production networks become more and more critical to overall business operations, UTM devices must have the ability to scale in performance without interruption of service.

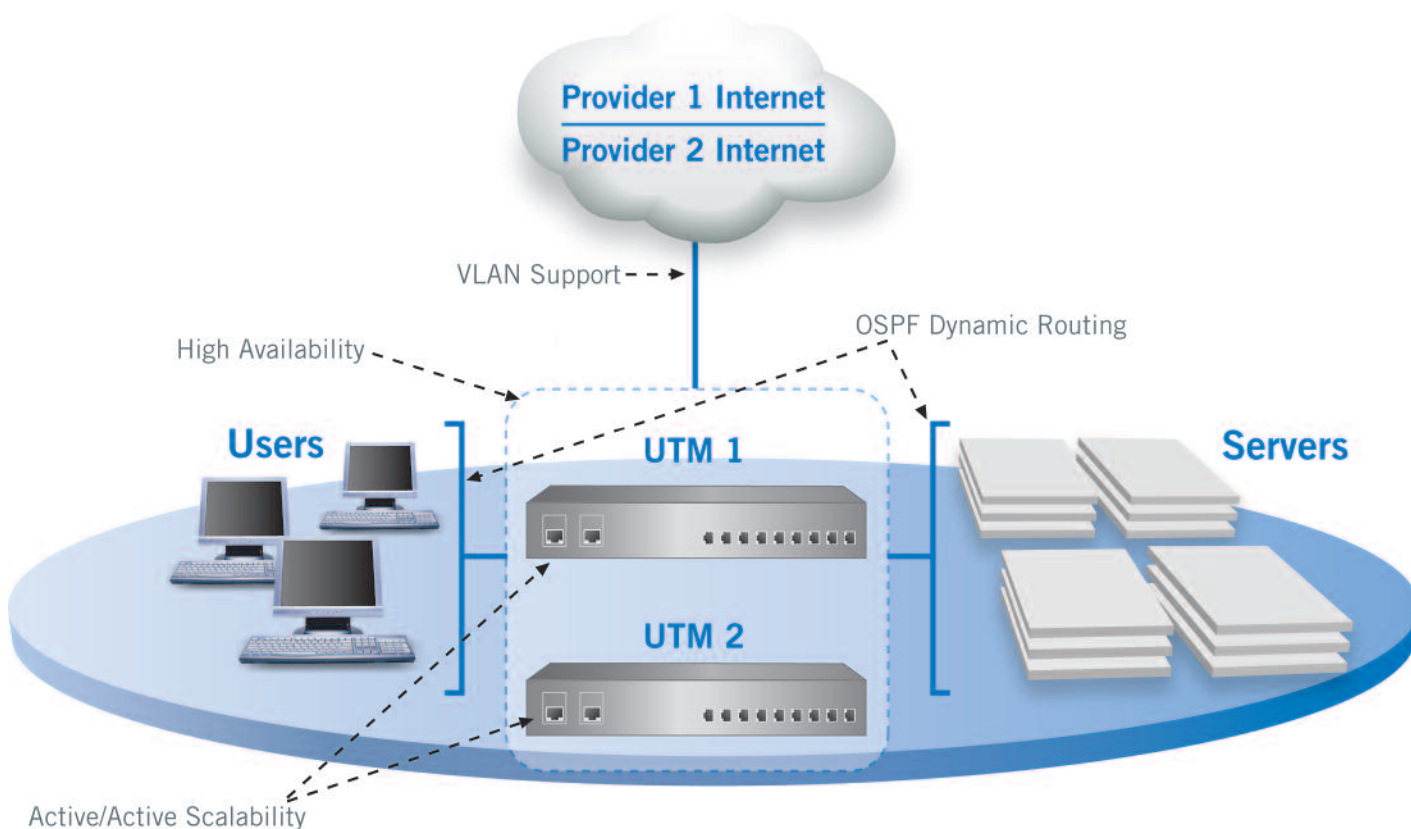
Measuring, testing, and evaluating scalability and high availability means understanding how the UTM device will interact with your network and then validating that it won’t take the network down either because of

component failure or system overload. Validation can be a difficult and expensive process, but it is critical in this environment. If you don't have in-house expertise to conduct testing, this is a definite place to consult outside help.

In many ways, these four requirements simply differentiate between basic SMB firewalls and traditional enterprise-class firewalls. However, "integration" for a UTM device will extend beyond these examples to include the necessity to mesh with existing desktop management systems, depending on the UTM services involved, or existing anti-virus systems or even existing intrusion prevention systems (IPS) and intrusion detection systems (IDS). Evaluating this aspect of UTM firewalls could be as simple as checking for support with your other security service vendor, but it's dangerous to assume that a check box on a web page represents true integration.

For example, the Network Access Control (NAC) juggernaut currently flowing down the riverbeds of security vendors is an obvious area where a UTM device should have full integration. But this area is in such flux that integration means more than finding a matching set of logos announcing partnership opportunities. Instead, you have to test the UTM firewall in place with any system it is going to touch.

An enterprise-class UTM device can accommodate all of these integration requirements. **The best devices minimize the number of components ranging from patch cords to external load balancers — by being designed explicitly to operate in non-stop environments.** In other words, the enterprise-class UTM device is designed specifically to integrate smoothly in any part of the network; it is both a security device and a network device.



*Enterprise-class UTM must smoothly integrate into the existing network, which means that features such as dynamic routing and VLAN support are required. Well-designed scalability and high availability strategies are also needed to support the needs of enterprise networks.*

### KEY EVALUATION CRITERIA

- Does the device support dynamic routing and can it integrate with your existing network routing fabric?
- Is the device designed for flexibility, with the ability to increase interface count as well as security zone count? Does it support different types of interfaces, and integrate with existing VLANs?
- Can the device be part of a high-availability configuration with the lowest possible number of external components? Does the device support high availability technologies such as clustering and active/active high availability to scale as needed?

## Consolidation Support

An older term in many enterprises is “consolidation”, more specifically “firewall consolidation.” The concept entails looking at the firewalls scattered around the network and considering whether one box might do where three, four or five sit now. Consolidation, or at least re-considering architectures to see if consolidation is appropriate, is a healthy activity.

Consolidation of other, non-firewall functions can also benefit the security manager. If network functions such as routing and load balancing can be consolidated into a single device, this reduces the tight interlock between network and security devices. A complex and less reliable topology can often be consolidated down to a smaller number of more functional systems. The result of device consolidation is greater freedom for security teams to make changes and updates without having to interact with the network team.

It’s important to understand that consolidation alone is not the sole purpose of UTM, but that UTM deployments benefit from consolidation. There are clear benefits to any enterprise in consolidating security functionality

into a smaller number of devices and management points such as cost savings, higher levels of reliability, and lower levels of complexity, to name a few. But simply sticking a bunch of firewalls into one large *über*-firewall doesn’t provide the same benefits as UTM.

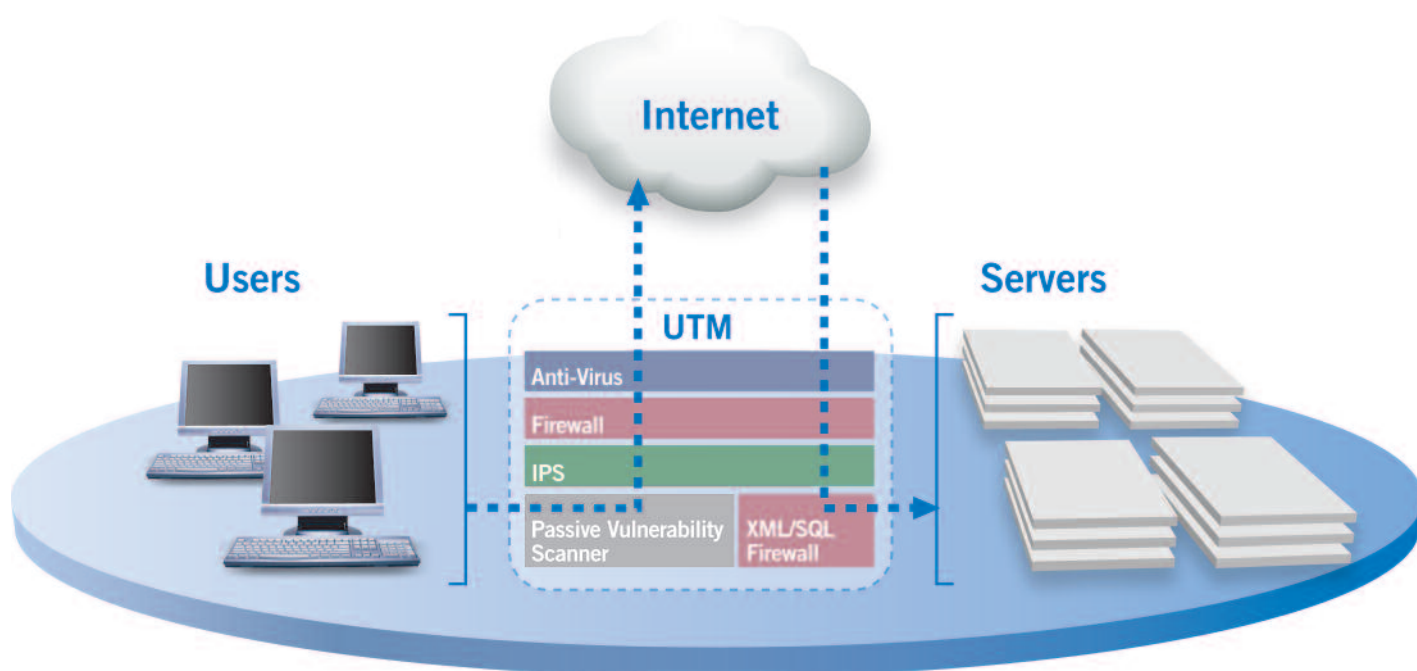
The goals of consolidation — higher reliability and simpler management — overlap with the goals of UTM. A firewall consolidation project may or may not have a UTM component and should be considered separately, using its own justification and evaluation criteria.

Adding UTM into a consolidated firewall brings even stronger benefits because it means that threat mitigation services can protect more than an Internet-facing barrier. Consolidated firewalls, by their nature, protect multiple security zones. The common argument that networks are becoming “de-perimeterized” is especially true in this environment, in which a single perimeter is replaced by a set of perimeters — all protected by the same device. By having the ability to watch most, if not all, of the traffic traveling between different security zones, the consolidated firewall becomes an obvious point to bring UTM services into play in a very cost-effective way.

**When firewalls have been consolidated in large networks that consolidated firewall becomes an ideal place to add in UTM functionality.** Therefore, a firewall consolidation discussion should always include the question “Will we be introducing UTM features here?” By pursuing a firewall consolidation project without first considering UTM, you risk revisiting the project sooner than you’d expect or to want to.

Consolidated and other internal firewalls add another twist to the UTM evaluation. A “traditional” UTM firewall generally includes anti-virus and IPS features as a minimum — features specifically designed to help in perimeter protection between corporate networks and the Internet. However, when the consolidated/internal firewall is used as the UTM base, internal threat management security functions are just as appropriate for addition to the UTM firewall. For example, Layer 7 application-specific firewalls, such as XML and database





*UTM supports consolidation by integrating multiple security functions into a single device. A critical requirement of a consolidated device is the capability to direct traffic flows through the UTM device, activating different functions, because not every security function applies to all flows.*

firewalls, all the way down to network visibility and vulnerability detection functions, facilitated by tools like passive network scanners, are all threat management functions that belong inside the network and fit well in a consolidated firewall offering UTM services.

When evaluating UTM for an enterprise network, **it's wise to not only focus on threats you want to counter at the edge, but also on threat management that is most appropriate deep in the network.** Enterprise-class UTM architectures can support these types of services.

#### KEY EVALUATION CRITERIA

- Does the device allow for 'internal' type UTM services, such as network visibility functions, application-specific firewalls, and passive vulnerability scanners?
- Does the device support firewall consolidation with multiple security zones, and allow different UTM functions to be applied against each zone?

#### Enterprise-class UTM features: Not a simple checklist

Because UTM has received a lot of attention from the press and analyst community — as well as firewall buyers — firewall vendors have worked to add UTM features onto their existing systems, either to keep up with the crowd or because the title “UTM” does indeed sell more boxes. Of course, UTM in the context of firewalls is not a particularly new idea. The actual thought of performing multiple threat mitigation processes on a single system is quite old. For example, the makers of proxy firewalls, the oldest firewall technology deployed in networks today, all argue their “deep inspection” provides additional threat mitigation.

However, IDC's Charles Kolodgy nailed a description onto the banner of UTM by saying that any UTM device must comprise firewall, IDS/IPS and anti-virus services. This edict had the unfortunate side effect of turning UTM into a simple checklist question: “do I have IDS/IPS and anti-virus? If so, the ensuing press release can read: It's a UTM!”

The critical issue is that UTM has to meet your network's requirements. It's very clear that a successful UTM deployment for your network might not include either IDS/IPS or anti-virus — but that doesn't mean it's not UTM.

The rush to create UTM products has created an additional problem: the match (or mismatch) between threat mitigation features and your requirements. IPS technology sets up a great example of this potential mismatch. Enterprise IPS can be a complex endeavor, with questions of management, forensics, signature tuning, as well as the base technology itself, which can range from rate-based to signature-based to anomaly-based detection and every combination in between. Most UTM products, especially those in the SMB space, simply take an open source tool, add a poorly designed GUI, and slap the IPS label on it. The result is not even a strong IPS service for an SMB environment, and it certainly isn't going to support enterprise requirements.

IPS isn't the only example of a touchy UTM feature. Anti-virus coverage, one of the most popular UTM features (because of IDC's definition), varies wildly from product to product. Some vendors support anti-virus scanning of only Web traffic. Others scan email. Some are configurable in terms of what they will scan, including protocol and port numbers; others are not.

**The key evaluation issue becomes whether or not the UTM feature set meets your needs**, not whether there is a check mark in the anti-virus box on the vendor's glossy list. The better UTM architectures are designed to meet the needs of enterprise networks, and can prove it by offering both the features you need and a full disclosure about what exactly each feature does.

Many UTM vendors try to derail this argument by pulling out a favorite codeword: "best of breed." "Best of breed" is supposed to explain how the particular product, selected by the UTM vendor, is the right one for you — because it's the "best!" The reality is that "best of breed" means different things to every

network. If the UTM vendor writes some component of its product, such as the virus scanner or the IPS, themselves — as is common in SMB-oriented devices — the term "best of breed" disappears from marketing literature to be replaced by another: "cost effective."

The result of this almost random decision about what products will go into a UTM firewall is that the quality of the pieces, whether "best of breed" or "cost effective," has received a well-deserved bad reputation. Firewall vendors have a relaxed attitude about the quality of the pieces they add to their UTM firewalls, driven more by marketing and partnership pressures than what is best for the customer. This attitude has created a bad taste in the enterprise security architect's mouth, where features are not there simply to fulfill a checklist.

In UTM products suitable for an enterprise deployment, you should expect that most vendors will integrate third-party products for almost all of their UTM capabilities, rather than try to build the technologies that will serve to be all things to all people. Good UTM architectures not only have well-respected security vendors represented, but also afford choice among the vendors included in the stack. This is particularly important for anti-virus services, where an enterprise needs to pick products that properly complement its server-side and desktop security strategies.

Just as important as breadth in vendor choice is depth in UTM features. Enterprise-class products can't stop at the anti-virus and IPS check boxes because enterprises need more. Good products allow the customer the flexibility to select UTM features based on their network's requirements, not on what the UTM vendor decided to put in place. Of course, no one expects full flexibility to run any application, but a good UTM product is one that gives the enterprise manager a broad choice of features making it easy for third-party security vendors to integrate technology into their platform.

What security tools can be expected in an enterprise-class UTM device? While it may seem like the firewall/VPN base is a given, that's not necessarily so — although



it's unlikely that anyone will try and introduce a device under the UTM banner that is not built on top of a solid firewall/VPN base. IPS, Web content filtering, and anti-virus/spyware/phishing/spam are likely to be the most popular and thus most readily available security features. IDS may be available, but is unlikely to be as good a UTM fit as IPS features are. Layer 7 application and protocol specific firewalls will be a strong differentiator between SMB and enterprise-class UTM devices, with XML, Web, and SQL firewall features likely candidates. Finally, security focused more on the end user and less on content, such as patch management, vulnerability assessment, and other NAC-oriented features will round out enterprise UTM.

### KEY EVALUATION CRITERIA

- Can each of the UTM features support your network's requirements for that feature, or are you being forced to compromise on any of the key services your network requires?
- Is the functionality of the UTM feature(s) you would deploy fully documented and described? Do you know what you're getting, especially with loaded terms such as 'IPS'?
- Do you have a choice of vendors when third-party products are integrated onto a UTM platform?
- Does the platform have sufficient breadth of add-on UTM features to meet your needs in edge, perimeter, and core environments? How much of our "wish list" of features is available today? How many will be available in 12 months?

## Platform Extensibility and Flexibility

The UTM marketplace is filled with fixed configuration devices, because they meet the needs of SMB networks at an affordable price point. Extensibility, as needed for larger deployments, drives up costs. That said, it is a clear requirement, not just for interface density, but also for interface flexibility.

An extensible platform can grow with hardware needs: more CPU, more memory, more interfaces and I/O. But additionally, an extensible platform should also grow "out of the box." Enterprise networks often need multiple devices, not just for high availability and scalability, but because they have multiple locations, complex network architectures, and rarely have just one of anything. Element management, where each device is managed separately, is not acceptable in this environment. Having multiple devices requires that management tools cleanly cross firewalls.

Performance needs (covered in greater detail above, see section on "Performance") also call for extensible platforms. Turning on UTM features always translates to a performance hit, meaning the ability to extend platform performance at the same time you extend the functionality is a requirement for enterprise-class UTM. The most appropriate UTM device architectures for enterprise networks are those that have the flexibility to turn features on while at the same time increasing performance.

When you choose to use UTM features in an enterprise network, you're putting more and more eggs in the same basket. An enterprise-class UTM architecture gives you the flexibility to scale in multiple directions at once and gives the security architect more options, rather than constraints.

**Extensibility and flexibility are issues that affect the internal architecture of any UTM device.** A naïve UTM implementation will assume that all packets flow through the device in the same way and that the order in which threat management elements, firewall, and VPN see each packet is fixed. Having a rigid order to

UTM functions and assuming every packet goes through every enabled function, is an inflexible design not appropriate for enterprise networks.

As a simple example, consider the relationship between IPS and firewall services. In some enterprises, security architects have elected to place IPS inside of the firewall to only see legitimate threats passed onto the network. In others, IPS goes outside the firewall to give a better ‘big picture’ of the threat landscape and the continuing attacks on the network. **Enterprise-class UTM devices should have the flexibility to order and reorder (and even duplicate) these functions as needed to meet the needs of the security architecture.** When a device doesn’t offer the flexibility to see and manage traffic flows, you’ll end up adding threat mitigation boxes on either side of the UTM device anyway — meaning that the UTM device is not really doing its intended job.


Solution vendors specializing in point products — Web proxies, IPS, Layer 7 firewalls — have tried to paint UTM as an alternative (and not a very good one) to their “best of breed” solutions. Sometimes this is a ludicrous

assertion — how much do Web content filters actually differ and are those differences really wide enough to claim a “best of breed” designation? Where a UTM vendor has restricted the breadth of applications as well as the depth of vendors, there is a bit of truth to this assertion. That’s why enterprise-class UTM products must have extensibility in these directions as well.

An important part of choosing a UTM solution is realizing that “threat management” is constantly being re-defined by the market. For every new threat, you have to ask: Is the right solution to scatter additional boxes around the network — the “new day, new threat, new box” style of security? Or, rather, is it to have an extensible platform that can be adjusted to support new types of threats without network re-engineering? In an enterprise network, the latter is obviously the more prudent course.

#### KEY EVALUATION CRITERIA

- Does the platform have expandability and extensibility, including everything from interface cards to additional CPUs?
- Does the platform support multiple device management, with true integration across devices?
- How much flexibility does the device give you for control within the system, to enable and disable and reorder UTM services?
- How much flexibility do you have in the breadth and depth of services offered?



*Flexibility in UTM's gives the network architect the option to have security services where they make sense, not in a fixed order. For example, an IPS might see traffic twice: once before it's decrypted to mitigate DoS attacks, and then again after it's passed through the VPN and firewall to look for different types of threats.*

## Management Requirements

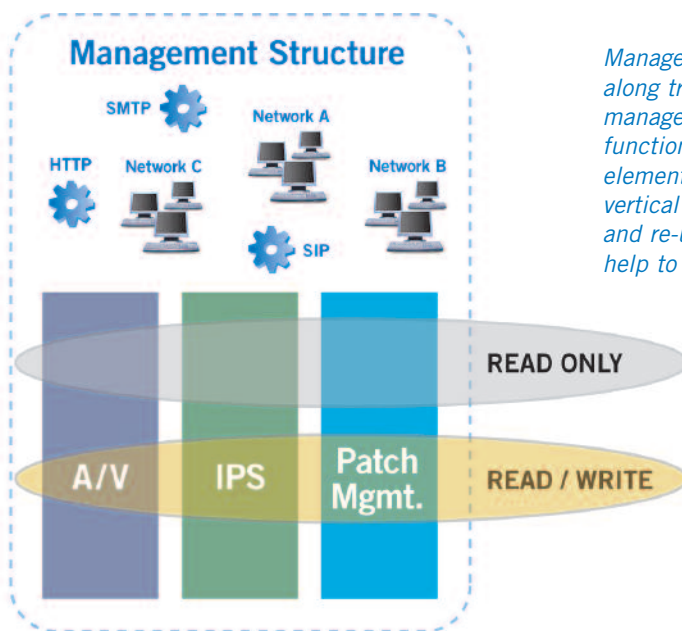
It's a pleasant thought to imagine that a unified GUI could offer the ability to configure and manage everything from setting up IP routing configurations to weeding through alerts on an IDS console. But the cold, hard reality is that different GUIs exist for a reason — the metaphor, layout, and work flow that you use in defining routing protocols and interface settings is very different from what you use in configuring a virus scanner.

UTM devices must tread this line very carefully and enterprise security architects need to evaluate capabilities with equal care. For example, most SMB UTM devices attempt to shoehorn an IPS console into an existing Web-based GUI. As any IPS manager knows, the requirements for these products are vastly different from those of a firewall, and may not even translate well to a Web-based interface. IPS GUIs are highly interactive, provide drill-down and forensics capabilities, and must be able to manage thousands of rules (as opposed to the

tens of rules that might be found on the accompanying firewall). Simply trying to shove the IPS into the firewall GUI inevitably leads to disaster.

Disparate GUI styles should not preclude management integration, though. In fact, a desirable enterprise UTM management framework doesn't attempt to integrate all aspects of all GUIs into one dizzying console on your screen. Instead, it keeps the important parts of each function intact, while sharing information and configuration capabilities as broadly as possible.

Separation of management tools is also important in enterprise networks because the engineers who manage different parts of the system are also diverse. In a small business, a single person might handle all security from desktop to firewall. In an enterprise, these are typically different people sitting on dissimilar teams. **An enterprise-class UTM device needs to meet the expectations of each of the security, networking, and desktop management teams which all play a role in managing it.**



*Management in enterprise-class UTM requires the flexibility to divide along traditional horizontal boundaries, such as operators, network managers, and auditors, as well as vertical boundaries, such as by function (e.g., anti-virus, firewall, patch management) or by network element (e.g., subnets or buildings or departments). Even when the vertical boundaries require multiple management systems, the use and re-use of common elements such as networks and services will help to ensure a tightly integrated and error-free UTM department.*

This situation drives a second management requirement: separation of duties and of powers. Enterprise-class UTM management needs both vertical (function-based) separation — to keep the security people from stomping on the toes of the network people — as well as horizontal (privilege-based) separation — to keep operators from changing things they shouldn't. As regulatory and compliance requirements stretch their evil fingers deeper into each organization, management separation can also be a vital part of this effort.

## Conclusion

UTM products originally were crafted based on the needs of smaller networks and smaller enterprises, and have seen broad acceptance in their large niche of potential installations. However, the concept of UTM has value in large networks and large enterprises as well. To support UTM in large networks, though, products must meet a very different set of requirements that set them apart from SMB-focused UTM firewalls. By going further in the areas of performance, network integration, support for consolidation, platform extensibility and flexibility, and management, UTM vendors can meet the needs of enterprise network managers. When UTM products reach to meet enterprise needs, the results are a powerful toolset that can displace traditional firewalls and give network managers greater flexibility and greater capability to solve their immediate security problems quickly.