Requirements for Enterprise UTM Products:

Evaluation Philosophy and Test Plan

Joel Snyder
jms@opus1.com
Opus One
8-June-2007

Table of Contents

non-Definition of Enterprise UTM

The term Unified Threat Management (UTM) has as many meanings as there are products that carry that label. While UTM has primarily focused on the small- and medium-sized network, products are coming to market that aim at the enterprise.

At its core, UTM brings together three main ideas: multiple threat mitigation features, integrated on top of a mature firewall, deployed in an appliance form-factor. However, deciding what UTM features are appropriate for an enterprise is a difficult proposition.

It is tempting to create a checklist to somehow distinguish between enterprise-class threat management and SMB threat management. We are going to avoid that particular hubris by simply stating what we are going to test (and what we are not going to test) based on our experience in designing and deploying security within enterprise-sized data networks.

Features Overview

Here are several lists that will help you determine if your product is a right fit for this test. We will test everything in the "Must" list, and you should not submit a product that does not meet these 7 basic requirements. If you have features on the "Probably" list, we will also test them, but these are not requirements to participate in the test.

We will evaluate your support for additional threat mitigation features (such as those on the "Don't Have to Have" list), but we will not consider the presence or absence of these features in evaluating your product. In other words, you don't have to have anti-malware to be an Enterprise UTM product, but if you do have anti-malware, it should be done correctly.

Things Your Product *Must Have* To Be Considered A Credible Enterprise UTM offering

1. Firewall with multiple zones
2. NAT
3. IPsec VPN for site-to-site
4. Global management capability
5. IPS or "Deep Inspection" capability
6. High Availability (Active/Passive HA)
7. Typical baseline firewall performance of 500Mbps to 1Gbps or more

Things Your Product *Probably Should Have* To Be A Credible Enterprise UTM Offering

1. Dynamic routing (OSPF or EIGRP or both)
2. Central Log Aggregation in Management Tool
3. Ability to scale 10/100/1000 port count up
4. VLAN capability
5. Clustering (Active/Active HA)

Things Your Product *Doesn't Have To Have* At The Enterprise Level (at least for this test)

1. Anti-Malware
2. Anti-Spam
3. Content Filtering

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 the measurement and reporting of 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.

With UTM, system performance will be 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. Additionally, scanning both e-mail and HTTP traffic for viruses will slow down performance more than just scanning the web traffic would. Implementation choices will also make a huge difference: some devices might scan all HTTP traffic, while others might only scan traffic with a particular MIME type.

Our testing will focus on how performance of the system degrades as various UTM features are put into play. We will start by setting a "baseline" of performance for each firewall using a fairly aggressive traffic mix of HTTP traffic. We've chosen HTTP because it tends to stress firewalls with a high connection establishment rate, and lets us get performance levels up to about 1Gbps fairly easily. We will not be testing at speeds in excess of 1 Gbps.

Our baseline will include a moderate firewall rule set of several hundred rules, but without NAT, and the same set of rules with NAT, and then again across an IPsec VPN.

From there, we will selectively enable UTM features (such as IPS, anti-malware, or content inspection)]] present in the firewall to see how performance changes as UTM features are put into place. We may also vary the UTM configuration (such as enabling different sets of IPS rules) to provide a number of reference points for each product.

We do not anticipate changing the traffic mix to test specific features (anti-spam is the obvious example here).

Although we will not test speeds higher than 1Gbps, we will evaluate your ability to scale up to higher speeds either in the product tested, or in other compatible products you offer.

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 configuration up 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 virus scanner.

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.

We will evaluate the management features by considering the philosophy of management in the product, and by looking at the following questions:

- Is the management system capable of reasonably handling hundreds of rules and tens of devices?
- Does the management system have logging and troubleshooting tools to help in day-to-day operations?
- Does the management system effectively integrate UTM features? How do different security and networking professionals within the enterprise make use of the management system(s) to do their jobs?
- How does the management system handle compliance requirements such as long-term log archiving and auditing of operations?

Additional Threat Mitigation Requirements

It seems that the minimum requirements for a UTM firewall that separate it from simply "just another firewall" have to include IPS. The reason IPS is a core requirement (and, for example, Anti-Malware is not) is that the fit between Layer 3 and Layer 4 defenses (typical firewall) and IPS is very symbiotic—whereas features such as anti-malware and anti-spam are often better handled via different data paths. With IPS, as well, there is little "post incident" management that is required. For example, when an IPS rule fires and a packet is dropped, you don't expect to have to put that packet into quarantine (as you might with anti-spam).

We will be focusing more specifically on IPS features, and do a more detailed evaluation of the IPS capabilities of each product.

We expect that most products will have a combination of signature-based IPS and heuristic IPS technologies. There is no specific requirement for one technology or the other, so long as whatever has been included is manageable and does its job.

We will be evaluating IPS primarily from a management and a functionality point of view. However, depending on the number of products submitted for this test, we may also do some limited additional performance testing with IPS at varying levels of attack.

Some specific questions we will try and answer about the IPS management system include:

- Can the security manager make sense of the rules (if signature-based IPS is included) and/or heuristics?
- Are there adequate alerting capabilities?
- How are logs and forensics handled?
- Do the threat mitigation features provided make sense in an enterprise setting?

Although this is not an IPS evaluation, per se, we will be taking a closer look at IPS functionality and management than other threat mitigation features which might be included in the product.

Of course, UTM systems will often have other threat mitigation features. Since we are not requiring any specific features (other than IPS and the firewall/VPN itself), we will evaluate features that are present. We will **not** be scoring products based on the number of UTM features present, and products with more features will **not** be treated preferentially to those with fewer. From our point of view, it is more important that you do a good job on the features that you have than that you have a pile of features.

Our goal in this section is to explain to network managers what the options are, and how well the threat mitigation is implemented in each product when present. This section will include some of the discussion of tradeoffs in implementation (such as scanning all traffic for malware, versus only specific protocols or ports). We believe that security managers will make their own decision about which features are requirements for their network and will be able to deduce the suitability of products for their deployment based on information presented in this test.

High Availability, Scalability, and Networking Requirements

Enterprise-class UTMs need to support the more exacting SLAs and complicated network topologies present in larger corporations.

Three areas we will specifically look at include high availability and scalability, VLAN and multiple interface support, and dynamic routing support.

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 being scattered throughout the network to harden and protect it from both external and internal threats. Enterprise-class UTM devices need to offer the flexibility to work both at the edge and deep within the network. For example, while an edge device may need only two or three interfaces (“inside” and “outside”, for example), an internal firewall will need a much higher interface count (one for each server group, for example) as well as VLAN capabilities thereby offering as many security zones as necessary.

Enterprise networks, both at the edge and in the core, have other characteristics that distinguish them from their smaller brethren and that effect the type of UTM necessary: dynamic routing, high availability (HA) needs, and scalability requirements. 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-oriented UTM devices must integrate with the existing routing fabric and support common enterprise routing protocols, such as OSPF.

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.

We will be specifically evaluating these features—when present—by answering the following questions:

- What high availability options are available? Does the high availability system work?
- What strategy does the product have for scalability? (We do anticipate testing the performance of "scaled" products, but only within the limits of our testbed)
- What dynamic routing options are present? (Please note that RIP doesn't count.) Do they work and will they integrate with an existing multi-vendor dynamic routing environment?
- What types of interface scalability are offered, not just in this product, but perhaps in other devices in the same product line?
- Is VLAN support available and how well is it integrated into the product as a whole?