

Making Everything Easier!™

Unified Threat Management

FOR
DUMMIES®

Learn to:

- Increase visibility and control of applications and threats
- Improve network performance and reduce complexity
- Consolidate management of essential security technologies

Brought to you by

FORTINET®

Ed Tittel



Fortinet (NASDAQ: FTNT) is a worldwide provider of network security solutions and the market leader in unified threat management (UTM). Our products and subscription services provide broad, integrated, and high-performance protection against dynamic security threats while simplifying the IT security infrastructure. Our customers include enterprises, service providers, and government entities worldwide, including the majority of the 2010 Fortune Global 100.

***Unified Threat
Management***
FOR
DUMMIES®

by Ed Tittel



WILEY

John Wiley & Sons, Inc.

These materials are the copyright of John Wiley & Sons, Inc. and any dissemination, distribution, or unauthorized use is strictly prohibited.

Unified Threat Management For Dummies®

Published by
John Wiley & Sons, Inc.
111 River St.
Hoboken, NJ 07030-5774
www.wiley.com

Copyright © 2012 by John Wiley & Sons, Inc., Hoboken, New Jersey

Published by John Wiley & Sons, Inc., Hoboken, New Jersey

No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning or otherwise, except as permitted under Sections 107 or 108 of the 1976 United States Copyright Act, without the prior written permission of the Publisher. Requests to the Publisher for permission should be addressed to the Permissions Department, John Wiley & Sons, Inc., 111 River Street, Hoboken, NJ 07030, (201) 748-6011, fax (201) 748-6008, or online at <http://www.wiley.com/go/permissions>.

Trademarks: Wiley, the Wiley logo, For Dummies, the Dummies Man logo, A Reference for the Rest of Us!, The Dummies Way, Dummies.com, Making Everything Easier, and related trade dress are trademarks or registered trademarks of John Wiley & Sons, Inc. and/or its affiliates in the United States and other countries, and may not be used without written permission. Fortinet is a registered trademark of Fortinet, Inc. All other trademarks are the property of their respective owners. John Wiley & Sons, Inc., is not associated with any product or vendor mentioned in this book.

LIMIT OF LIABILITY/DISCLAIMER OF WARRANTY: THE PUBLISHER AND THE AUTHOR MAKE NO REPRESENTATIONS OR WARRANTIES WITH RESPECT TO THE ACCURACY OR COMPLETENESS OF THE CONTENTS OF THIS WORK AND SPECIFICALLY DISCLAIM ALL WARRANTIES, INCLUDING WITHOUT LIMITATION WARRANTIES OF FITNESS FOR A PARTICULAR PURPOSE. NO WARRANTY MAY BE CREATED OR EXTENDED BY SALES OR PROMOTIONAL MATERIALS. THE ADVICE AND STRATEGIES CONTAINED HEREIN MAY NOT BE SUITABLE FOR EVERY SITUATION. THIS WORK IS SOLD WITH THE UNDERSTANDING THAT THE PUBLISHER IS NOT ENGAGED IN RENDERING LEGAL, ACCOUNTING, OR OTHER PROFESSIONAL SERVICES. IF PROFESSIONAL ASSISTANCE IS REQUIRED, THE SERVICES OF A COMPETENT PROFESSIONAL PERSON SHOULD BE SOUGHT. NEITHER THE PUBLISHER NOR THE AUTHOR SHALL BE LIABLE FOR DAMAGES ARISING HEREFROM. THE FACT THAT AN ORGANIZATION OR WEBSITE IS REFERRED TO IN THIS WORK AS A CITATION AND/OR A POTENTIAL SOURCE OF FURTHER INFORMATION DOES NOT MEAN THAT THE AUTHOR OR THE PUBLISHER ENDORSES THE INFORMATION THE ORGANIZATION OR WEBSITE MAY PROVIDE OR RECOMMENDATIONS IT MAY MAKE. FURTHER, READERS SHOULD BE AWARE THAT INTERNET WEBSITES LISTED IN THIS WORK MAY HAVE CHANGED OR DISAPPEARED BETWEEN WHEN THIS WORK WAS WRITTEN AND WHEN IT IS READ.

For general information on our other products and services, please contact our Business Development Department in the U.S. at 317-572-3205. For details on how to create a custom For Dummies book for your business or organization, contact info@dummies.biz. For information about licensing the For Dummies brand for products or services, contact BrandedRights&Licenses@Wiley.com.

ISBN 978-1-118-08701-5 (pbk); ISBN 978-1-118-08863-0 (ebk)

Manufactured in the United States of America

10 9 8 7 6 5 4 3 2 1



Publisher's Acknowledgments

Some of the people who helped bring this book to market include the following:

Acquisitions, Editorial, and Vertical Websites

Senior Project Editor: Zoë Wykes

Editorial Manager: Rev Mengle

Senior Business Development

Representative: Karen Hattan

Custom Publishing Project Specialist:

Michael Sullivan

Composition Services

Senior Project Coordinator: Kristie Rees

Layout and Graphics: Claudia Bell,
Samantha K. Cherolis

Proofreader: Jessica Kramer

Special Help from Fortinet:

Patrick Bedwell

These materials are the copyright of John Wiley & Sons, Inc. and any dissemination, distribution, or unauthorized use is strictly prohibited.

Table of Contents

Introduction	1
How This Book Is Organized	1
Icons Used in This Book.....	2
Chapter 1: The Need for Consolidated Network Security	3
The Good: More and Better Network Access	4
The Bad: Muddier Waters to Navigate and Secure.....	5
The Ugly: Criminal Mayhem and Mischief.....	8
Chapter 2: Why Traditional Firewalls Cannot Keep Up with Today's Threats	15
Inadequate Protection at Best.....	16
Stand-alone Products Mean Reduced Network Visibility	18
Multiple Devices Create a Performance Hit.....	21
Way-High TCO	22
No More Stand-alone Technologies?	23
Chapter 3: UTM Brings Order to the Chaos	25
Stepping Up to UTM.....	26
UTM: Flexible and Future-Ready	27
Integration beneath a“Single Pane of Glass”	29
Beware Frankenstein's Monster!.....	31
Accelerated Processing Speeds Business.....	32
Keeping Up with the Changing Threat Landscape	33
Chapter 4: A Closer Look at Unified Threat Management	35
Choosing the Right Inspection Technology for Your Network	36
General-purpose Processors versus ASICs.....	37
Chapter 5: Unified Threat Management in Action . . .	43
Parts that Make the Whole: Essential UTM Components	44
Securing the Network End to End	48
Use Cases	52

Chapter 6: Ten (Okay, Eleven) Key Questions to Ask When Evaluating UTM Solutions	55
How Does the UTM Solution Perform?	55
Which Security Technologies Are Included?	56
Which Network Features Are Supported?	58
Does UTM Support IPv6 Natively?	58
Does It Play Well with Virtual Environments?	58
Is the UTM Solution Scalable?	59
Does It Offer High Availability?	59
What Kind of Management and Reporting Does It Include?	60
How Does the UTM Solution Stay Ahead of Threats?	61
How Much Does It Cost?	62
What Support Options Are Available?	63
Glossary	65

Introduction

Very few security technologies leap out of nowhere to occupy a major segment of the market within a few years. Unified threat management, also known as UTM, is one of those few. UTM presents buyers with the opportunity to improve their network security and performance while reducing network complexity and costs.

Today's networking environment and threat landscape are changing constantly — new devices, applications, and threats appear almost daily. Organizations of all sizes struggle to enable secure access for mobile users on the latest smartphone or tablet while blocking the latest threats hidden inside application traffic from social media sites. It's a tough fight!

One traditional response to new threats has been to add new stand-alone security technologies to a network. Unfortunately, this adds complexity and cost, as each new technology means a new device to deploy, a new set of policies to configure, and a new management console to monitor. And, because each new stand-alone product isn't integrated with the existing product set, it also creates potential blind spots in an organization's security strategy, through which threats could pass undetected.

This book presents, explores, and explains the features and benefits that UTM can deliver to your network, within the context of today's rapidly evolving network and threat landscapes.

How This Book Is Organized

The chapters in this book demystify the components and technologies used in UTM, and then some.

- ✔ **Chapter 1: The Need for Consolidated Network Security** explains the challenges encountered when securing evolving networking environments, the need to control applications, the rise in crimeware, and more.

- ✔ **Chapter 2: Why Traditional Firewalls Cannot Keep Up with Today's Threats** explains the problems of legacy network security technology and much more.
- ✔ **Chapter 3: UTM Brings Order to the Chaos** describes how a consolidated approach to network security improves protection and performance while lowering cost and complexity.
- ✔ **Chapter 4: A Closer Look at Unified Threat Management** digs into more detail about the different inspection techniques and hardware processors that provide the muscle for UTM software.
- ✔ **Chapter 5: Unified Threat Management in Action** examines the components of UTM, from application control to web filtering, and how UTM has been implemented in real-world organizations.
- ✔ **Chapter 6: Ten (Okay, Eleven) Key Questions to Ask When Evaluating UTM Solutions** poses questions to ask before making any UTM purchase.
- ✔ **Glossary:** Here, we define acronyms and key terms.

Icons Used in This Book

Every *For Dummies* book sprinkles *icons* throughout its margins to call attention to text worth noting. Here are the icons in this book.



These are points to keep in mind as you immerse yourself in the world of unified threat management.



Sometimes we share technical information that goes beyond the basics. This icon tells you that the related text delivers more UTM details than absolutely necessary. You can skip it if you like!



This icon sets out on-target information to help you maximize the benefits of enhancing network security using unified threat management.

Chapter 1

The Need for Consolidated Network Security

In This Chapter

- ▶ Securing an ever-changing network
 - ▶ Reining in applications
 - ▶ Fending off crimeware
-

In today's increasingly mobile world, networks change constantly. New services, access methods, and even devices continue to show up in networks at a frenetic pace. Who could have foreseen the impact of Twitter and Facebook, or iPads and Android smartphones, only a few years ago? As a result, organizations of all sizes face challenges in maintaining control over network and security policies. Unfortunately, many organizations continue to take a traditional approach to network security in spite of changing software, devices, and user habits. Such a traditional approach cannot adapt to the latest trends and leaves organizations vulnerable to today's threats.

To keep up with the constant change, organizations of all sizes must adopt a different approach. But before we can discuss what's needed, in this chapter we need to take a look at how the network has evolved in the last few years, and the implications for network security: Some are good, some are bad, and some are downright ugly.

The Good: More and Better Network Access

The evolution and incredible proliferation of networking and security technology has produced many benefits across a wide spectrum of systems and users. These benefits include network access, anytime and anywhere. Users can utilize ever-improving network bandwidth across both wired and wireless networks for faster access to more data and services. At the same time, tremendous improvements in the range of applications (apps) available have boosted productivity across the board for all users.

Fast and secure remote access

Today, remote users access corporate data in many ways from around the block or around the world. They can use any device with a network connection, including smartphones and tablets, to access systems and networks inside a corporate firewall.

An amazing application explosion

Every day, the number of applications is exploding, with new apps of all kinds being launched, including tools, mobile programs, games, and widgets of every imaginable description.

Exploiting the Internet's speed and always-on access, today's applications make it trivial to deliver content to end-users. Versions crafted for mobile platforms, like iOS or Android, extend enterprise-focused applications to smartphones and tablets, beyond typical desktop or notebook PCs.

- ✔ Today's applications deliver better access to "back end" and internal data resources thanks to improved communications protocols and compatibility standards. This enables tremendous leaps in productivity among employees, partners, vendors, and customers.
- ✔ Delivery models for apps and content have also evolved significantly over the past decade. Applications may originate inside a network (also known as "on-premise," where they traditionally reside). Alternatively, applications may reside in the cloud, hosted by an application services

provider, or in some instances, may even be delivered in a combination of the two (known as a “hybrid” deployment).

- ✓ Finally, the growing use of social media such as YouTube, Twitter, Facebook, and Google+ in everyday business practices reflects a profound change in where and how potential customers, partners, and agencies gather data, influence buying decisions, and build brand awareness.

A lot of good things have happened on corporate networks, and a lot of those things involve new technologies and new uses for existing systems and information assets. But these changes also add to the burden of maintaining security and require forward thinking to ensure that levels of protection and security keep up with new uses, new forms of access, and new forms of communication.

The Bad: Muddier Waters to Navigate and Secure

Alas, the ongoing application explosion means that new apps show up and run on systems and devices on organizational networks every day. Too often, such applications fly under the radar of network administrators and appear without planning, proper licensing, or even informed consent. Worse, such applications can potentially inundate those networks on which they appear with unknown, questionable, or even outright malicious content. This can be very bad indeed!

New apps evade detection

Here’s why the application explosion currently underway on today’s networks is a real issue for organizations of all sizes: Traditional firewalls cannot detect these new applications. Traditional or “first-generation” firewalls rely on *port numbers* or *protocol* identifiers to recognize and categorize network traffic and to enforce policies related to such traffic.

This approach only works when applications use specific ports to request or deliver content, or require the use of unique and easily recognizable network protocols. For many

years, this was the standard approach for service delivery and application access on the Internet, and helped make it easy to ensure interoperability between different applications and networks.



Apps that use specific port numbers or protocols also made it easy for network administrators to block unwanted traffic: an administrator simply created a firewall policy to block traffic based on a port number or protocol type. For example, to prevent the transfer of e-mail across a network, a network administrator could block port 25, which is associated with the Simple Mail Transfer Protocol (SMTP). If an administrator wanted to stop File Transfer Protocol (FTP) traffic, an inherently unsecure way to exchange files, the administrator would block ports 20 and 21. Browser-based applications often use only two port numbers, each associated with a protocol vital to user productivity and responsible for the bulk of Internet traffic today. This means Port 80 for HTTP (HyperText Transfer Protocol), or Port 443 for HTTPS (the secure, encrypted version of HTTP routinely used for commercial transactions or any Web interaction involving sensitive data like online banking).

Think about what this means: All traffic from browser-based apps looks exactly the same to traditional firewalls. They can't differentiate among applications so there is no ready way to use them to block bad, unwanted, or inappropriate programs but to permit desirable or necessary apps to proceed unhindered. Blocking traffic on these ports or for those protocols essentially blocks all web-based traffic, including content and services for legitimate business use.

Explosions are messy

Today's continuing application explosion means that every day, more apps appear that traditional firewalls can neither detect nor control. Yet, establishing and maintaining control is not only desirable but is also downright essential because unchecked applications can wreak havoc. Web-based applications, because they pass unchecked through legacy firewalls, may not only import malicious or unwanted behavior and content onto an organization's networks, but they also may export proprietary, regulated, or confidential data out of its networks as well. These imports and exports of data and

content pose risks to property and competitive advantage and raise the specter of possible legal liability or compliance failure. Not pretty!

Consider these examples of the kinds of unwanted or undesirable side effects that unchecked application explosions can cause on an organization's networks:

- ✓ **Exposure to malicious content:** User-created content encompasses a wide range of threats, such as malware or links to malicious sites. Whether user content is a Twitter post, a Flickr photo upload, or a restaurant review, it may contain malicious code that could compromise user systems or even entire networks. Common risks include links to malicious sites where visitors are subject to *drive-by downloads* — downloading of malicious content without the user's knowledge.
- ✓ **Unwanted bandwidth consumption:** Bandwidth-intensive web-based applications such as YouTube can clog networks and impede delivery of business-critical content. File-sharing applications can bog down networks because of large file sizes and the sheer number of files being swapped.
- ✓ **Exposure to data leakage:** Apps that can accommodate outbound file attachments can permit employees to export sensitive, confidential, or protected information outside organizational boundaries and controls. This incurs potential civil and criminal liabilities, as well as loss of customer trust and brand equity.



This is a wide range of untoward, unwanted, or dangerous consequences that occur when new applications appear on an organization's networks — which pose problems that cannot go unsolved without incurring massive risks.

How apps avoid detection and control

New web-based applications often pose security concerns. Those that evade detection are particularly dangerous because they can introduce risk, impact performance, or compromise systems and networks without attracting notice from

system administrators. How? It happens daily on corporate networks because some web-based applications evade detection by utilizing a range of stealthy or obscure techniques:

- ✔ One application may tunnel inside another application to make it appear legitimate and to cloak itself with the mantle of a “trusted application.” Think of such application tunneling as a wolf in sheep’s clothing, and you’ll be right on target. It only takes one quasi-sheep in a whole herd of real sheep, and you start to appreciate the problems that application tunneling can pose.
- ✔ Some applications employ encryption algorithms to prevent inspection and identification of their content. Unless steps are taken to inspect and deal with that content, rogue applications can move traffic without notice or control.
- ✔ Other applications may connect to a system using dynamic port numbers, or masquerade as different applications. These techniques let them evade port-based firewall rules because they deliberately misrepresent themselves and their traffic. Malware, especially Trojan horse programs, is infamous for using such techniques to evade detection and for exporting sensitive or confidential data beyond organizational boundaries.

These disguises and evasive techniques make it easy for data theft or leakage to afflict organizations. This risk is acute in the new generation of application-borne threats designed to harvest confidential, propriety, or regulated data.

The Ugly: Criminal Mayhem and Mischief

Today’s threat landscape is chock-full of bad actors with lots of bad intent. Beyond the threats of damage or data loss from malware, professional criminals have gone online to ply their unsavory trades. Online crime is an ever-increasing part of the threat landscape, but it’s not the only kind of attack that can target corporations, organizations, and individuals.

Crimeware defined and deplored

Crimeware is a class of software that includes source code and developer kits, designed to help criminals compromise systems or build their own networks of compromised computers called *bots*, to do their bidding. Hackers who create these tools act like legitimate software development companies and offer various licensing options or annual support contracts to their criminal clientele.

Crime-as-a-service takes this a step further, and is delivered by malcontents who offer to steal data or mount attacks for a fee. These individuals have no obvious affiliations or causes (except making money). Such service providers work as contractors to deliver criminal actions online to the highest bidder.

Other common online crimes include *malicious mischief* (damage or theft for the thrill of it, instead of for a profit) that springs from political or personal malice. As we write this book, recent news is rife with reports of exploits from “hactivist” groups named “Anonymous” or “Lulzsec.” These shadowy groups target organizations’ user data and internal content for theft and disclosure as retribution against policies or political and social affiliations.

At the same time, there’s a growing cybercriminal network emerging from the Internet shadows, usually housed in countries where online law enforcement is lax or non-existent. Often, criminal enterprises are tolerated in exchange for a share of their proceeds. Tools of choice for these online black markets are crimeware and crime-as-a-service offerings.

In the sections that follow, we explore some of the ugliest elements this side of the online universe.

Botnets and cyberwarfare

A *bot* is an individual computer that has been compromised with malware that takes instructions from a criminal. Each such computer is called a bot (short for “robot”) and collections of such computers are called *botnets*. Criminal organizations may create and run their own botnets, but the biggest botnets (some of which include hundreds of thousands of

bots around the world) are usually available for hire. The crooks that operate botnets are called *botmasters* or *bot-herders*, and they take instruction from their criminal clientele to program their botnets to undertake various criminal projects. See Figure 1-1.




Figure 1-1: An example of a botnet.

Corporations, institutions, government agencies, and other large organizations can fall prey to various kinds of bot attacks. In every case, one computer inside the organization's boundaries will be compromised and added to a botnet. The compromised computer can then serve as a beachhead to attack other systems (internal or external), steal information, transfer funds, or perform other illicit activities.

Here are some common forms of attack that often produce more bots inside corporate boundaries than many IT professionals may be prepared to accept or believe:

- ✓ **Phishing-based attack:** An employee or other person operating a computer inside the corporate network clicks a link in an e-mail or other social media message and accesses a phishing website. Credentials are then phished and stolen from the victim. Oftentimes, this information is then used in a second-stage attack through some flavor of social engineering to claim additional victims.

- ✔ **Search Engine Optimization (SEO) attack:** SEO attacks succeed by hackers gaming popular search engines to include malicious websites among the top-ranked sites. This is achieved by using a botnet, creating hundreds of thousands of web servers and pages that use popular search terms. A potential victim then queries some hot topic, clicks on a malicious search result, and then lands on a dummy site that redirects them to a malicious site. Popular legitimate search sites are commonly targeted because they can produce the largest number of hits, and the largest number of potential victims. As with phishing, creation of a bot can open the door to network penetration and compromise.
- ✔ **Compromised legitimate websites:** Legitimate websites can become compromised in various ways without detection by the site operator. For example, ads being served by a third party that redirect visitors to malicious sites, injecting malicious code into the site itself that redirects users to malicious sites, or replacing legitimate HTML code that contains malware.
- ✔ **Internal infections:** Mass storage devices, from thumb drives to picture frames, all have the capability to carry viral payloads. This is often achieved by using an auto-run function, once the device is introduced to a network. However, some botnets can infect legitimate files on mobile devices that are executed on-demand by victims.

Ransomware: Pay up, or else...!

Ransomware is a special kind of malware to which individuals and organizations sometimes fall prey. Ransomware targets mission-critical systems or applications, and once installed, is very difficult to get rid of. By design, ransomware is nearly impossible to remove without access to keys or codes that only its developers possess. The original ransomware threat can be disinfected from a system; however the damage often cannot be reversed without engaging the attackers themselves.

Essentially, what happens with ransomware is the following: As ransomware infects a system, it makes the system unusable by encrypting various essential files (often using rigorous asymmetric cryptography) and disabling recovery/repair features and capabilities. It then announces itself to its victims

after taking control over those files or capabilities. Unless and until the victim makes payment to unlock what's being held for ransom, the files remain locked and the system unusable.

This gives the bad guys an unbeatable point of leverage in extracting payment from their targets. Only after paying a ransom does the victim receive the keys necessary to unlock the system or regain access to critical files. Other ransomware locks out systems at the boot level, so you may not even get into your operating system to attempt system recovery. Ransoms typically float around \$100 USD but can jump much higher if there's a high-profile target.



Once a user finds him- or herself reading a ransom note, data has already been encrypted or essential functions disabled, and the damage is essentially done. Traditional firewalls are oblivious to and impotent against such threats.

Phishing and spear-phishing attacks

In addition to botnet compromises described earlier in the section “Botnets and cyberwarfare,” outright phishing attacks against organizations also occur all the time. While these attacks may not always involve incorporation of the machine under attack into a botnet, they can and often do involve inadvertent and unwanted disclosure of sensitive and/or regulated information. This data generally includes information with financial value such as credit card numbers and credentials, user account names and passwords, and financial account information for banks, brokerages, or other financial service providers.

In the most benign form of phishing attack, an employee or customer of an organization clicks on a link in a fraudulent e-mail, tweet, Facebook post, and so on. This takes the person to a realistic-looking website that requests the individual to log in or otherwise divulge credential information. Cybercrooks use this information to impersonate their victims to steal money, services, or other items of value. Normal phishing attacks don't do much to target specific victims. Most phishing attacks create messages or lures that are designed to appeal to the broadest possible audiences. These messages, postings,

tweets, and so forth are then broadcast far and wide, and victims more or less select themselves for attack.

However, businesses and organizations are increasingly falling prey to a more targeted form of phishing. Known as *spear phishing*, such attacks are directed at employees or members of a particular organization. These attacks are often crafted to look like they are from an internal user or come from a trusted partner or service provider (in-house credit union, health insurance provider, and so forth).



Spear phishing uses social engineering to persuade users to open e-mail, tweets, or Facebook posts with links to malicious sites or malware attachments where applicable. Sometimes a link is not required, rather a direct communication request via e-mail or even the phone (known as “Vishing”). This information is carefully worded to be relevant and interesting to a specific individual or small target audience, to raise the odds of getting prospective victims to swallow their lures. Criminals take advantage of the wealth of online data individuals post about themselves on social media sites to create the content of the spear-phishing message. Spear-phishing attacks can also be used to infect a system with malware to gain desired information, often using infected documents (PDF, doc or docx, and xls orxlsx formats are popular for malicious payloads).

For example, some spear-phishing attacks have been reported to target accounting or financial services professionals within large corporations, because such individuals are likely to handle electronic funds transfers for their organizations. Targeted messages request these individuals to open a file to examine the details of a supposedly fraudulent or questionable transaction. The attachment instead contains malware that infects the target’s system.

From there, any systems these individuals interact with are monitored closely by the attackers. Credentials are stolen on login, and accounts are compromised. Cybercriminals have also developed sophisticated malware to use in conjunction with these attacks. The malware uses a technique known as *form injection*, which allows the attacker to inject messages and/or questions and fields into live browsing sessions, such as a secure banking site. This way, questions look like they are coming from legitimate websites, when in fact they are

being injected on the fly by attackers. It's a very effective way for criminals to phish further information from infected targets. See Figure 1-2 for more details.




Figure 1-2: How a spear-phishing attack may infect a victim's computer.

Traditional firewalls are unable to protect against these threats because they lack the ability to inspect the content or prevent access to malicious sites. They also offer no protection against most other forms of cybercrime. A different approach and a different kind of tool are clearly called for.

In Chapter 2, we outline what's needed to stop these threats.

Chapter 2

Why Traditional Firewalls Cannot Keep Up with Today's Threats

.....

In This Chapter

- ▶ Grappling with the problem of inadequate network security
 - ▶ Understanding how stand-alone security devices decrease network visibility
 - ▶ Evaluating firewall performance with stand-alone applications
 - ▶ Examining the costs of traditional security protection on your network
 - ▶ Getting a birds-eye view of the characteristics of a better solution
-

Because network traffic is so heavily web- and application-based, the ability of legacy security technologies to protect internal networks is increasingly challenged, to say the least. Attacks have become more frequent and more serious. Every day we read about cybercriminals whose attacks have been designed to avoid detection and who keep coming up with more and better ways to steal our information. Are they getting the upper hand?

As attacks and their consequences become more dangerous and costly, organizations are learning to their dismay that their existing security infrastructures just can't keep up. Traditional firewalls and security appliances have lost their ability to block such traffic. That means that prevailing tools and technologies can't prevent the worst and most nefarious attacks from affecting or damaging networks.

This chapter looks at traditional security technologies and how the makeup of today's network traffic affects their ability to protect an organization's networks. At the same time, we also look at how a mix of multiple legacy security components can decrease network performance and increase the total cost of ownership (TCO).

Inadequate Protection at Best

Those network devices that sit at the boundary between internal networks and the wild and often wooly Internet have a lot to do to maintain security nowadays. It's become commonplace to inspect and manage traffic crossing the network boundary to enforce policies, to block viruses and malware, to prevent potential intrusions or attacks, and to prevent unauthorized access. At the same time, it's necessary to screen outbound web traffic, to filter inbound traffic of all kinds, and to watch for and block potential intrusion and react quickly to signs of denial of service attacks.

That's a lot of diverse security and network technologies to operate and manage. It's no wonder that traditional methods just aren't keeping up with the onslaught of traffic and the many different kinds of filters that need to be applied to the mix in real time.

The problem with legacy firewalls

One of the first and most basic lines of network perimeter defense is a *firewall*, a device that inspects inbound and outbound traffic on a network. Firewalls were the first widely deployed network security technology when Internet use really began in earnest. The firewall's job is to inspect that traffic and to decide what traffic is allowed to go from outside to inside, and from inside to outside.

But network traffic has changed quite a bit in the past decade (read Chapter 1 for the big picture). A great proportion of traffic on the Internet, and on organizational networks, is now web-based.



Because of the commingling of modern application traffic with everyday web access, traditional port-based firewalls have essentially become blind to the most common type of network traffic. This means that they cannot distinguish different types of traffic that use the same port, they cannot detect applications that tunnel inside other applications, and they can't see inside encrypted packets either. They can't even block sneaky rogue applications that use nonstandard port numbers.

Evolving firewall technology still can't keep up

As simple rules for filtering traffic in a firewall proved to be too permissive, firewall vendors deployed new technology to stay abreast of the changing threat landscape. Thus, firewall technology evolved to incorporate more sophisticated methods to protect against new threats.

One important advance in firewall technology was the development of a *proxy server* (or “proxy” for short), which inserts itself between a client on one side of a network connection and a server on the other side. A proxy server sits transparently on the network and checks all packets for policy compliance before forwarding the packets. It reassembles all of the packets in a stream, examining a complete file attachment rather than looking at one packet at a time.

Stateful packet inspection firewalls were created to keep track of the state of network connections. These stateful firewalls distinguish legitimate communications from suspicious or malicious types of connections. They allow only packets that corresponded to a valid, recognized, active connection between a client and server. They reject other attempts to connect or transmit.

To provide basic inspection of contents, *deep packet inspection* (DPI) firewalls inspect the payload or data portion of a network packet. DPI provides some protection against attacks from viruses and worms, as it identifies and classifies network traffic based on a collection of characteristic patterns

called *signatures* found in the packet payload. Signatures support more comprehensive and effective control and filtering than methods that examine only packet header information. (There's more about DPI later in this chapter.)

In spite of these advances, firewalls could not stop the evolving threats. With each new firewall technology, the hackers devised new evasion techniques. New threats simply slipped through legacy firewalls and headed right onto the trusted network, attacking clients and servers with gusto and abandon. Network security vendors sought to fight back by creating new technologies to fight the new threats.

Stand-alone Products Mean Reduced Network Visibility

Over time, as the threat landscape evolved, organizations began adding specialized or *stand-alone security devices* (appliances or software) to their infrastructures. The intention was that these devices would provide workarounds to the limitations of traditional port-based firewalls.

Each device addresses a specific threat: One appliance provides malware screening, and another provides content filtering of websites and traffic. Still others detect and block intrusions, or add spam filtering and e-mail message handling. Put them all together on a network, and you could have half a dozen or more appliances, each inspecting the stream of traffic moving across the network. See Figure 2-1.



All of these devices were designed to improve a network's security by adding functionality missing from the firewall. However, a patchwork of stand-alone technologies can have the opposite effect on network visibility as well as performance. These threat-specific technologies don't talk to each other easily (if at all). They lack central management and monitoring because each product operates on its own. Plus, data from individual devices isn't aggregated to create a complete or holistic view. How can you manage the security of a network if you can't really see it end to end?




Figure 2-1: Multiple devices addressing incoming threats.

An administrator can buy yet another stand-alone tool to gain insight across this kind of patchwork environment. For example, a security information and event management (SIEM) tool collects information from network devices, evaluates it to identify security events, and provides some analysis. The drawback with this approach is that although it adds visibility by aggregating data from a range of products, it does not add centralized control. An administrator still needs to configure the various stand-alone devices individually.

You can also acquire a security management tool that collects, analyzes, and assesses security and risk information. But not every management tool is 100 percent compatible with every device or appliance, especially in a heterogeneous environment. Such a tool may provide adequate, consolidated reporting but probably won't offer configuration capability, which is essential to be able to react quickly to changes in the threat landscape.



By an unfortunate twist of fortune, stand-alone devices that are meant to shore up inadequate firewalls may themselves be susceptible to some of the same limitations as the firewalls. Many IPS or DLP systems, for example, cannot see into encrypted packets.

Examples of stand-alone security appliances

Here are some examples of typical stand-alone security appliances that are deployed in a network:

- ✔ **Virtual private network (VPN):** A VPN uses special protocols to move packets of information across the Internet securely. In general, VPN protocols encrypt traffic going from sender to receiver. This makes such traffic appear completely garbled to anyone who might intercept and examine those packets while they're on the Internet. VPNs use encryption to protect the traffic they carry from unauthorized access. Because the VPN packets wrap the encrypted data inside a new protocol envelope — a technique known as *encapsulation* — a VPN creates a private, encrypted “tunnel” through the Internet.
- ✔ **Data loss prevention (DLP):** DLP looks for confidential, proprietary, or regulated data leaving the network. It can prevent the accidental “leakage” of data, such as a customer service representative including account information in a response to an e-mail to a customer. It can also stop the wholesale loss of a proprietary data, such as a disgruntled employee forwarding an entire customer contact list to a prospective employer.
- ✔ **Intrusion prevention system (IPS):** An IPS acts as a network's watchdog, looking for patterns of network traffic and activity, and records events that may affect security. An IPS issues alarms or alerts for administrators, and is able to block unwanted traffic. IPS also routinely log information as events occur, so they can provide information to better handle threats in the future, or provide evidence for possible legal action.
- ✔ **Content filtering technologies:** These devices block traffic to and/or from a network by IP address, domain name/URL, type of content (for example, “adult content” or “file sharing”), or payload. They maintain a *whitelist* of trusted sites and a *blacklist* of forbidden sites to prevent users from violating acceptable use policies or being exposed to malicious content.

Multiple Devices Create a Performance Hit

In this section, we look at the performance impact of a traditional firewall environment when adding more stand-alone appliances. Each device performs its own inspection of network data, which means that data is inspected multiple times against multiple rule sets.

Going further, a device might be capable of deep packet inspection (DPI), which dives into a packet's payload for a detailed look at the contents. Such inspections look for the usual threats — viruses, spyware, spam, and so on — but may also detect unauthorized or questionable network protocols like those often used for music or video downloads. These are generally banned from a business network, not only because they can steal precious network bandwidth but also because they can leave an organization open to legal liability for copyright infringement.

DPI works by identifying and classifying network traffic based on a collection of characteristic patterns called signatures found in the packet payload. This supports more comprehensive and effective control and filtering than methods that examine only packet header information.

In particular, DPI is effective against attacks that can fit into a single packet, such as buffer-overflow attacks, denial-of-service (DoS) and distributed-denial-of-service (DDoS) attacks, as well as some intrusion attempts.

So, multiple inspections and especially DPI are all good, right? From a security standpoint, yes; but from a performance standpoint, not so great. All this inspection takes time and adds to the amount of *latency* that accumulates from multiple inspection processes and lowers the speed of data transmissions. For some industries that rely on real-time market data, such as financial services or brokerages, these delays can cause a competitive disadvantage.



Anytime you add a stand-alone device or “bolt on” additional capability to boost security rather than switch to an integrated solution, you will run into these kinds of limitations. These limitations can run the gamut but mainly involve higher latency, more complex management, or, ugh, incomplete or unreliable security (the very thing you were trying to improve in the first place).

Way-High TCO

The patchwork approach to security will not only slow your network — it can cost a lot of money! When you seek to compensate for an inadequate firewall, you must invest in multiple security appliances, thereby adding to your capital expenses (CapEx).

These appliances must be installed, configured, tested, and maintained. Multiple products mean many consoles that must be used to view data, configure policies, and run reports. These tasks require significant technician and/or administrator time, and thereby drive up your operating expenses (OpEx). Maintaining multiple products also requires multiple service and support contracts from different vendors, making the renewal process very costly and time-consuming to handle. Performing different software updates for different products, testing and ensuring the compatibility across all the platforms with different software releases, is another significant management overhead that impacts the cost of ownership.

You can invest in a SIEM (see the section “Stand-alone Products Mean Reduced Network Visibility” earlier in this chapter) to alleviate some administrative OpEx, but you cannot cut the majority of the OpEx cost, and you’ve increased your CapEx even more. Ouch!

Plus, anytime you install additional security equipment on a network, you’re increasing the number of potential points of failure. Any such failure can incur both OpEx and CapEx to effect repairs, depending on the type of failure and its severity, and the time and effort needed to restore operations to normal.



Here's the bottom line: As a network's complexity increases, so does the expense required to manage that network. In the long term, it's usually more cost-effective to fix the core problem than to limp along with a system that is weak at its seams. It's also wise to shy away from bandage on top of bandage.

Every administrator should carefully consider the TCO of add-on security appliances versus an integrated solution before making significant infrastructure commitments. The hidden costs to acquire and sustain a handful of stand-alone devices can be much higher than anyone might realize. Once committed to the stand-alone strategy, however, those costs become unavoidable!

No More Stand-alone Technologies?

If the approach of adding appliance after appliance can't really address your network's evolving security needs, what kind of approach might be worth considering in its place? The most desirable characteristics in such a solution would have to be

- ✓ **An integrated and holistic view of network security:** Just as the combination of stateful packet inspection and IPS technology enables technologies like DPI, a holistic and integrated view of network security permits all types of threat or traffic information to drive a comprehensive and well-informed information security strategy.
- ✓ **Optimal TCO:** Reducing the number of overall devices and simplifying management lowers both CapEx and OpEx. Furthermore, it also reduces annual maintenance and subscription costs as well as the number of vendor contracts.
- ✓ **Minimal latency impact:** Reducing the number of overall devices also means that data need only be screened, inspected, and filtered once as it passes through a security device. Moreover, such a device can be built around custom processors specifically designed to handle all those tasks as close to wire speeds as technology will allow.

- ✓ **Comprehensive security coverage:** Integration of security functions means that the same holistic view of security that supports maximum security also delivers the best and most thorough protection to networks where such technology is deployed.



Establishing and managing security controls across an entire network creates an approach called *defense in depth*. An attack that gets past any particular countermeasure must get through additional layers of security to access services or resources on the network.

Real defense in depth means erecting multiple layers of protection throughout entire IT systems and infrastructures. Defense in depth is designed to address security vulnerabilities that can come from people, technology, and operations or processes. It provides protection against today's sophisticated attacks like Advanced Persistent Threats (APT) that utilize multiple techniques to compromise a network. It also ensures that even if a threat is able to evade some technologies (such as anti-virus filtering), it will be stopped by other technologies (such as IPS).

We take a closer look in Chapter 3 at unified threat management, the consolidated solution that addresses the shortcomings of legacy approaches to securing the network and delivers integrated, effective security.

Chapter 3

UTM Brings Order to the Chaos

In This Chapter

- ▶ Understanding Unified Threat Management
 - ▶ Being flexible and ready for the future
 - ▶ Seeing and maintaining security through “a single pane of glass”
 - ▶ Avoiding a hodgepodge of devices
 - ▶ Realizing that accelerated hardware also accelerates business
 - ▶ Staying abreast of the changing landscape
-

Many organizations want to find some way out of the expensive, labor-intensive, unsecure, and chaotic mess of traditional firewalls supplemented with multiple stand-alone security technologies. Unified Threat Management (UTM) is the approach that many organizations have adopted to improve visibility and control of their network security while lowering complexity of their networks.

UTM creates an environment in which all network security falls beneath a single, consistent technology umbrella. UTM enables the consolidation of all traditional as well as next-generation firewall functions into a single device.

In this chapter, we begin by defining UTM. Then, we explore how UTM addresses the various issues and deficiencies of legacy security systems described in earlier chapters, to bring some order to this chaos — as the chapter title says — as well as improved security and manageability!

Stepping Up to UTM

UTM represents a significant shift in the evolution of network security technology and equipment. UTM generally refers to a security appliance that consolidates a wide range of essential network security functions into a single device, including next-generation firewall technologies like application control.

As network threats evolve and new threats emerge, network security must change and adapt to protect against such threats. This adaptability can make UTM difficult to define because the technologies included can vary from vendor to vendor. Over time, however, the collection of capabilities associated with UTM has consistently expanded, and this trend shows no sign of tapering off.

As we write this, the best UTM solutions for networks include the following core security functions in their arsenals:

- ✔ Network firewalls perform stateful packet inspection.
- ✔ IPS detects and blocks intrusions and certain attacks.
- ✔ Application control provides visibility and control of application behavior and content.
- ✔ VPN enables secure remote access to networks.
- ✔ Content filtering halts access to malicious, inappropriate, or questionable websites and online content.
- ✔ IPv6 support in all network security functions protects networks as they migrate from IPv4 to IPv6.
- ✔ Support for virtualized environments, both virtual domains and virtual appliances.

They also include additional security technologies that organizations can choose to deploy, including

- ✔ Data loss prevention that blocks accidental or intentional loss of confidential, proprietary, or regulated data.
- ✔ Anti-virus/anti-spam protection that prevents malicious payloads or unwanted messages from entering networks.

- ✓ Endpoint control that enforces corporate security policies with remote users and devices.
- ✓ Integrated wireless LAN (WLAN) controller to consolidate all wired and wireless traffic through a single device, to simplify policy creation and enforcement while reducing network complexity.

There's much more to modern UTM appliances than a laundry list of functions, however, as you'll see throughout the rest of this chapter (and the remainder of this book).

UTM: Flexible and Future-Ready

UTM delivers a flexible, future-ready solution to meet the challenges of today's networking environments. Not every organization is going to deploy every technology included in a UTM device; in fact, most do not. A UTM gives an organization the ability to deploy as many or as few technologies as it needs, when and where it needs them. The best UTM solutions employ a simple licensing model that includes all technologies, eliminating the need to purchase additional modules or to count users as an organization's security requirements change over time.

UTM delivers a versatile, powerful solution to meet the challenges of today's most complex networking environment. It overcomes shortcomings of non-integrated, legacy systems built around traditional firewalls and individual, stand-alone appliances and software.

The most advanced UTM devices may be described as possessing the following characteristics. These devices are

- ✓ **Flexible.** UTM devices are able to deploy multiple technologies to meet the unique and changing needs of modern organizations and an evolving threat landscape. Real-time updates ensure the technologies, policies, and other security measures are always current. And, with all the technology shipping with the device, adding additional features is just a few mouse clicks away.

- ✔ **Future-ready.** UTM devices are designed around versatile technologies that will remain valuable years into the future. These devices are ready to embrace changes to functionality and the networking environment while maintaining their performance. For example, support for 10 Gigabit Ethernet (GbE) and IPv6 today means that customers won't have to perform a "forklift upgrade" of their security infrastructure in the future.
- ✔ **Powerful.** UTM devices must keep pace with ever-increasing network performance requirements, to ensure that they don't become network bottlenecks.

UTM devices also provide better awareness and control into security issues and potential problems related to network traffic, user behavior, and application content. This increased ability to see security threats lets UTM devices support much more sophisticated detection and prevention technologies than traditional firewalls can support.

One of the technologies included in UTM devices mentioned earlier in this chapter — application control — makes a good case in point for more sophisticated detection and prevention.

Advanced application control permits administrators to establish more granular policies than simply Allow or Deny rules on a per-application basis. Application control provides the granular control over content and behavior that organizations need in order to manage today's application explosion. They can control what individual users can do, and when they can do it, by:

- ✔ Allowing certain functions but denying others within individual applications (such as granting access to Facebook chat but blocking the ability to click on links to other sites or download files)
- ✔ Allowing access to applications based on time of day (such as blocking access during normal work hours) or limiting the amount of bandwidth allocated to certain applications
- ✔ Establishing quotas or limits on time spent on non-essential sites (such as limiting access to auction or social media sites to only one hour per day for each employee)

Integration beneath a “Single Pane of Glass”

One significant advantage UTM devices have over legacy firewalls combined with multiple appliances is integrated management. With legacy systems, each component in the mix (firewall, VPN, application control, IPS, content filtering, and so forth) has its own management console. Operators need to keep multiple screens open and active so they must jump from one console to the next to see all aspects of their network security. This makes the job of trying to “connect the dots” and understanding the meaning of diverse events even more difficult.

The “single pane of glass” management console, by contrast, refers to one coherent and consolidated management interface that provides ready access to all configurations, management, and monitoring functions. Instead of separate consoles, however, this approach presents one consolidated high-level dashboard for all security functions. It also gives operators the ability to drill down into whatever level of detail makes sense for individual functions (or collections of functions) under their control.



Other benefits result from the holistic view of security that a consolidated interface can deliver. This is especially obvious when it comes to seeing and reacting to combinations of events that represent a sophisticated attack. Single-function views can't even see this kind of thing, let alone deal with it!

When security management is segregated, three types of failures may occur as a result:

- ✓ Failure to notice an incident or event
- ✓ Failure to act or respond in a timely fashion
- ✓ Failure to maintain compliance and audit controls

These three failures are far less likely when consolidating management into a single console. In addition to avoiding such critical failures, organizations receive additional benefits, which are covered in the sections that follow next.

Cost effectiveness

When a company consolidates network and security management, one of the earliest, most tangible differences is a smaller number of devices to manage, maintain, and replace. At the same time, staff previously responsible for now-obsolete equipment become available to perform other tasks and activities.

Improved awareness

Legacy equipment catches only what it can identify and measure. Taking this shortcoming one step further — legacy equipment is also unaware of today's sophisticated advanced persistent threats (APT) and polymorphic attacks:

- ✓ *APT* (also known as *multi-vector attacks*) uses multiple techniques to entice and attack victims. These techniques can include spam e-mails, spoofed e-mails targeting specific users, fake pop-up alerts, social media blog and forum spamming, botnet infection and propagation, poisoned search engine results, compromised banner ads and site content, and so forth.
- ✓ *Polymorphic attacks* are those in which malicious code changes regularly (such as TotalSecurity ransomware masquerading as anti-malware) to evade detection by traditional signature-based techniques.

For example, in some organizations the perimeter security mail systems, IPS, and web content filtering systems are managed independently. As a consequence of an APT attack, each of the managers in those respective areas may witness and act upon suspicious content or behavior (such as links to phishing sites, incoming malware files, spoofed e-mail messages, and an internal system trying to communicate with an external site). It could be a long time, if ever, before the IT staff combines the alerts from the individual systems and then draws the obvious conclusion, which is — a successful attack has occurred!

However, with constantly updated unified threat management across these areas, such strange behavior isn't witnessed separately but gets correlated within a single console. Only when this kind of information is correlated does the IT staff

understand the threat — for example, from an APT attack with interrelated behaviors.

Reduces the false sense of security

UTM devices, with their increased visibility and control, also reduce the potential for organizations to believe they're protected when they are not. An organization relying on its legacy approach to network security will be unaware of any unknown threats that pass by unseen. This causes the IT team to believe that the network is safe and secure, while undetected threats could in fact be both present and active.



No solution can work completely independently of human interaction. But a single dashboard delivers better security through consolidated, smarter management, thanks to integrated and correlated data.

Beware Frankenstein's Monster!

You know the story of Frankenstein's monster, built by Dr. Frankenstein and made of parts from several different bodies. One hand came from one body, the torso and legs came from another, and they were then all stitched together.

Many organizations take a similar approach to their network security. One appliance comes from one vendor, other devices from another, and so on, and like Frankenstein's monster, they all get stitched together. By the time all security needs are covered, it's a safe bet that there are more devices than necessary. Besides consuming more space, power, and cooling, extra devices also lead to performance bottlenecks and additional service and support costs.

To make things worse, unnecessary devices often duplicate processing because they are oblivious to each other's capabilities or tasks. As a result, one device may perform some processing and then send it downstream to the next device that then repeats the same processing. This may even occur repeatedly in device after device after device. . . .



These different devices aren't aware of each other because they aren't integrated. Given what could be months or years between installations, the IT staff may not see redundancy either. When considering the potential for excess equipment, compounded by redundant processing and added costs, the argument for a single-vendor solution grows stronger.

Accelerated Processing Speeds Business

When an organization creates its security solution using equipment from multiple vendors, there is a lot of inefficiency and redundancy in the inspection of data. Each device performs any content filtering that's needed, irrespective of whether such processing may be repeated from one box to the next. For example, a firewall does deep packet inspection (refer to Chapter 2), then sends a packet downstream to an IPS for content inspection. This means that a single packet will be completely dissected and analyzed twice, simply because the firewall can't share its work with the content filter. Such repeated behavior can slow traffic flows because these various devices must duplicate their inspection efforts. As a result, the network slows and application performance degrades. This is particularly troublesome when low latency is required, such as in voice, video, or financial trading applications.

Conversely, when functionality is integrated with a UTM device, individual components operate as parts of a single system. Those components are aware of each other and benefit from each other mutually. For example, processing required by many layers of protection occurs once, instead of repeated for each one. The result of this integrated filtering is to reduce the amount of processing time and resources required to fully inspect the data, which accelerates the data as it travels through the UTM device.

In Chapter 4, we explore in more detail the performance advantages that UTM devices offer over other network security solutions.



A single network security platform that performs many related functions is an excellent candidate for hardware- and software-performance optimization. Without a doubt, use of purpose-built, high-speed custom processors called ASICs enables integrated UTM devices to outperform traditional firewalls and stand-alone appliances at the hardware level (see Chapter 4 for a more detailed description of ASICs).

But what may be less obvious is that consolidating multiple functions on UTM devices permits its software to be optimized for overall performance as well. That's because designers know what kinds of data is being accessed — and how that data flows — and take advantage of shared processing, storage, and input/output activity. A single vantage point provides more opportunities for optimization than the fragmented, single-function-at-a-time viewpoint that applies to legacy firewalls with add-on appliances.



The ability of optimized software to take advantage of custom hardware in an integrated UTM device enables traffic to flow more quickly through the network. This accelerates the delivery of the data, improves the performance of the systems that utilize that data, and enables more work to get done. Ultimately, this flows directly to the bottom line, so that what is good for the network is also good for the organization as a whole!

Keeping Up with the Changing Threat Landscape

Another critical component of any network security strategy is the ability of the security technology to stay abreast of changes to the threat landscape. These changes can be in the form of a new variation to an existing threat, or the arrival of an exploit that targets a previously unpublished or unknown vulnerability.

Regardless of what type of new threat is involved, the rules, definitions, signatures, and configuration of the various devices that make up an organization's security infrastructure

must change quickly before damage can occur. The source for the knowledge used to make these updates comes from extensive threat research. Today, state-of-the-art threat research requires sophisticated knowledge of a wide range of vulnerabilities, exploits, and threats.

To keep current with the latest threat research tools and techniques, an organization would have to invest substantial resources. This would mean creating a sizeable in-house team and systems to gather and synthesize large volumes of data on a global basis, and then using that information to determine the necessary configuration changes. This is the only way to detect emerging threats and then quickly update the rules, policies, and signatures for any affected systems or devices. That's a very tall order, just like Everest is a very tall mountain!



For most organizations, creating an in-house security research team is too cost prohibitive. Fortunately, the alternative to a do-it-yourself approach is simple: Purchase a UTM product backed by an extensive global threat research team, deployed around the world to keep up with the evolving threat landscape. A large dedicated team will ensure that products they support are up to date, deploying the latest protection. Paying for security service is better and cheaper than building and maintaining your own in-house capability.

In Chapter 4, we take a look at the underlying technology that enables UTM to deliver the benefits described in this chapter. You also get to see it in action.

Chapter 4

A Closer Look at Unified Threat Management

.....

In This Chapter

- ▶ Comparing flow-based and proxy inspection
 - ▶ Examining general purpose processors and ASICs
-

UTM is a broad-based network security platform that represents the next stage of evolution for traditional firewalls. A variety of hardware and software elements protect networks, including firewalls, IPS, application control, content filtering, anti-virus and anti-spam software, and more. UTM delivers all these forms of protection on a single, integrated platform.

UTM is administered from a single interface as a collective system. This saves network administrators and engineers from installing, configuring, and maintaining multiple security devices, and from trying to make dissimilar platforms all work in unison to protect their network infrastructure. Because all protective functionality resides within one device and management interface, all of its protective methods are integrated and better able to provide “unified” coverage. Also, a single UTM device saves room in the server rack, utilizes less power, and generates less heat because it can be deployed with a single device that occupies fewer slots.

In this chapter, we explore some of the technologies in more detail and discuss the profound benefits of UTM.

Choosing the Right Inspection Technology for Your Network

UTM devices examine traffic to determine whether that traffic violates any policies. They can employ multiple inspection methods to address many different types of threats. Each inspection method has its own characteristics, advantages, and disadvantages.

Flow-based inspection

Flow-based inspection, also called stream-based inspection, samples the data entering a UTM device and uses pattern matching to determine whether there is malicious content. It does not fully reassemble files like proxy-based inspection, which we discuss in just a bit.



Flow-based inspection analyzes the data in chunks rather than fully reconstructing all elements of the communication, including files. This dramatically reduces the amount of data a UTM device needs to analyze.

Benefits of flow-based inspection

Speed is the chief advantage for flow-based inspection. Because there is less inspection of the data and less processing of the data, flow-based inspection is faster than proxy-based. Unfortunately, flow-based inspection methods are less complete than the proxy-based approach, meaning that some malicious content may be missed.



The advantages of flow-based inspection are also its disadvantages. Although ignoring specific content increases speed, any data carried by packets within a flow remains unexamined. Depending on the type of intrusion or attack, flow detection methods can miss certain exploits.

Proxy-based inspection

Proxy-based inspection reconstructs content entering the UTM device and performs a full inspection of the content, looking for possible security threats. It does not sample the data like flow-based inspection; instead, it inspects the full contents of a communication session, such as its files.

The UTM device acts as a proxy in the security inspection chain. The device downloads the payload of a content protocol session, re-constructs it, and then inspects the content. If the content appears clean, then the device sends the content to the client. If a virus or other security problem is detected, the device removes the problematic content before sending the file (or web page, or what have you) to the client.

Benefits of proxy-based inspection

The major advantage of proxy-based inspection over flow-based inspection is the level of security provided. Proxy-based inspection takes a more detailed look at content, making the proxy approach more thorough. It catches more malicious content and threats than the flow-based method.

On the other hand, the deep inspection of the proxy approach requires more processing power, which can lower network throughput and increase latency.



Proxy-based inspection provides a high level of security and offers enhanced protocol awareness, but at a cost of latency and lower throughput in high-speed environments.

General-purpose Processors versus ASICs

All tasks performed by firewall and UTM devices involve software running on hardware. In the preceding section, we examine the differences between different methods of network traffic inspection, including their pros and cons. Software

manages all those inspection tasks, but the tasks have to run on hardware platforms. Thus, the processors used to power the software make a difference.

General-purpose processors are designed to work in a wide variety of computer hardware environments and to process a wide variety of instructions. Think of a general processor as a one-size-fits-all, not-purpose-built chip intended to deliver a wide range of services.

By contrast, application-specific integrated circuit (ASIC) processors are designed to run instructions for a very specific application or narrow range of functions. ASIC chips are created to deliver specific functions, and in UTM devices, deliver accelerated processing of security functions.

Benefits of general-purpose processors

The main advantage of a general-purpose processor is that it's generic. General-purpose processing chips are designed to work in the widest possible range of hardware platforms under varying conditions and to manage many different sorts of software instructions. This type of chip design isn't locked into a specific environment, platform, or instruction set, so it can adapt to different computing and business requirements easily.

General-purpose processors do caching very well, which is important in platforms such as firewalls that require fast parsing of application data and speedy lookups in the cache. The flexibility of general-purpose-processor architecture is well suited for the diverse requirements of application caching for firewalls.



What works for general-purpose processors also works against them. This processor architecture operates okay with everything but doesn't excel at application-specific tasks. General-purpose processors are equally good and equally poor at what they do. If you need a processor suited to a very specific set of tasks for a particular hardware design, general-purpose processors seldom provide "the best fit."

Benefits of ASIC processors

The strength of an ASIC chip is in its specialization. An ASIC chip is designed to perform a single set of functions with optimal efficiency, performance, and scalability. For network security, that means creating ASICs designed solely to perform processing and inspection of data for security purposes. ASICs designed to meet those requirements will outperform generic chips in the same devices. In this case, an ASIC chip is one component in an overall, high-performance hardware platform designed to run network and security specific software.



The fastest-performing UTM devices on the market utilize a combination of ASIC processors and general-purpose CPUs to deliver very high throughput combined with very low latency.

Tuning ASIC processors for work and packet flows

Although ASIC chips can be engineered for many tasks, a few are specific to specialized UTM security operations, such as content and network processors. There are also other specialized processors that work with the custom ASICs, offering additional acceleration options.

Content processors

A *content processor* is designed to make high-speed comparisons of objects contained in network traffic to the characteristics of already known threats. The design of content-specific ASIC processors allows them to examine objects, which can be network packets, compressed files, or other objects in the traffic stream. Content processors can quickly organize and examine these objects for any unusual patterns or variants suggesting a threat. Content processor chips are engineered for protocol recognition and parsing, required to perform these highly specialized examinations.

A UTM by any other name

The traditional compilation of individual network security measures is becoming increasingly inadequate to protect infrastructures from newer threats such as botnets and APTs. Defense efforts must be unified to provide comprehensive protection against continually changing cyber threats. Because network security infrastructures vary depending on customer size and security requirements, UTM platforms need to be scalable. For example, network security requirements for small and medium-sized businesses (SMBs) will be very different than those for global businesses and organizations.

When considering the term “Unified Threat Management,” no one individual collection of technologies accurately describes the range of options. Which technology is considered a part of a UTM solution? It often depends on the individual vendor, as the technologies can vary significantly. There is a caveat, however.

Many UTM platforms are engineered to meet the minimum requirements of smaller networks and cannot scale to meet evolving security needs. These devices match the general description of UTM since they are single appliances that host multiple security functions, but they do not provide application awareness or

other critical functions, nor can they scale to meet performance requirements. Some of these basic UTM devices use first-generation firewalls and IPS, possibly with deep packet inspection, but are not fully featured enough to protect more complex networks from emerging threats.

Advanced UTM devices, on the other hand, are designed specifically for today’s high-performance networking environments. They offer enterprise-grade features such as 10 GbE support, 40 Gbps or higher firewall performance, delivered with the combination of custom processors and general purpose CPUs described earlier.

There is also some confusion in the market over the difference between UTM and next-generation firewalls (NGFWs). *Next-generation firewalls* are similar to UTM devices in that they are consolidated network security devices and operate as an inline security barrier with network security policy capabilities in real time. The most significant difference is that they provide a subset of the technologies included in most UTM solutions.

Because all UTM devices are not created equal, do your research, shop around, and select the best platform for your organization’s specific network security requirements.

Network processors

Network processors operate at the interface level to deliver an extremely low latency data path. These ASICs are engineered for high-speed processing, examining the characteristics of the traffic itself rather than objects contained within the traffic, looking for patterns indicative of threats.



One of the key tasks for this kind of processor is to accelerate firewall, intrusion prevention, and application control performance. Network processors, because they operate at the interface level, enable wire-speed performance, independent of packet size, with switch-like low latency for multimedia services, unicast, and multicast traffic. These processors also provide accelerated encryption for IPsec, TCP offload, Virtual Domain support, QoS, traffic shaping, and management logging.

Security processors

Another class of custom processors exists as well, known as *security processors*. These are a multi-core, multi-threaded processor combined with content processor logic. They operate at the interface or system level and add additional packet processing and acceleration options.

General, content, network, and security processor integration

Up to now, we've looked at the differences between general-purpose and ASIC processors. Within ASIC processors, we've examined differences between content and network processors. In UTM systems, different chip types may be integrated into an overall architecture for optimal handling of different threats. In Figure 4-1, you see a high-level architecture for a UTM device that employs all of these processor designs.

Network traffic enters the UTM through Ethernet interfaces and is immediately sent to the network and security processors for analysis. Traffic is then passed or blocked based on the results of that inspection, and passed traffic moves to the general processor unit. The general processor sends any traffic that contains objects requiring content examination to a content processor, which compares the object to prior threats.




Figure 4-1: High-level diagram of a UTM platform using onboard content, general, security, and network processors.

As you can see in Figure 4-1, network and security processors filter all network traffic entering the UTM. In addition to performing threat analysis, this design also takes on such tasks as general TCP processing, network address translation (NAT), and encryption/decryption. By comparison, a content processor is not placed in line with traffic flow and doesn't come into play unless the general processing unit notifies the content unit that some object requires examination, maximizing the efficiency of the inspection process.

In Chapter 5, we look at how the various technologies in UTM work together to provide comprehensive network security. We also provide several customer examples.

Chapter 5

Unified Threat Management in Action

In This Chapter

- ▶ Getting an overview of essential technologies
 - ▶ Protecting the network from end to end
 - ▶ Exploring how real-world organizations put UTM to work
-

Today's organizations can't be too careful when it comes to network security, especially when customer data and regulatory compliance are at risk. Many layers of security are required to keep attackers out, and to keep sensitive data from falling into the wrong hands. At the same time, achieving strong security shouldn't come at the expense of overloading busy network staff or affecting application or network performance. Unified threat management (UTM) systems consolidate security technologies into a single, dedicated device with a single management interface.

In this chapter, you get a close-up view of UTM in action, starting with an examination of some essential UTM technologies: intrusion prevention, application control, web filtering, data loss prevention, anti-spam, and anti-virus.

You also find out how UTM truly protects a network from its perimeter to the core, and you learn some key implementation tips. Finally, you're presented with brief case studies of three organizations that chose a UTM solution, and how such a transition improved business and customer service.

Parts that Make the Whole: Essential UTM Components

A UTM system is more than just a security box sitting on your network. It includes multiple functions, from application control to web filtering and more, all centrally managed and working in concert.

In the following sections, we look at core UTM technologies you can deploy to protect your network from today's threats.

Application control

Application control, an important next-generation firewall feature, gives you visibility into applications that generate traffic on your network, along with the ability to control those applications.

Application control can identify and control applications, software programs, network services, and protocols. In order to protect networks against the latest web-based threats, application control should be able to detect and control Web 2.0 apps like YouTube, Facebook, and Twitter. Enterprise-class app control provides granular policy control, letting you allow or block apps based on vendor, app behavior, and type of technology. For example, you can block specific sites, block only your users' ability to follow links or download files from sites, or block games but allow chat.

Another feature of application control is the ability to enforce identity-based policies on users. The UTM system tracks user names, IP addresses, and Active Directory user groups. When a user logs on and tries to access network resources, UTM applies a firewall policy based on the requested application or destination. Access is allowed only if the user belongs to one of the permitted user groups.

Traffic shaping, a method of guaranteeing network bandwidth to certain types of traffic, gives priority to some apps while limiting bandwidth for others. Application control provides traffic shaping, which is handy for keeping bandwidth-hungry apps, such as Skype, iTunes, or video-sharing sites, in check.

Application control also extends to network endpoints, from workstations to smartphones. Using an application control list on a security gateway, applications are allowed, monitored, or blocked at the gateway and at the endpoint.

Intrusion prevention systems (IPS)

An *IPS* protects your internal network from attacks that originate outside the network perimeter, as well as threats from within. IPS is also considered an important component of a next-generation firewall. The IPS component in a UTM solution offers a wide range of tools to detect and block malicious activity, such as:

- ✓ **Predefined signatures:** A database of attack signatures, such as those used against unpatched operating systems, is updated regularly.
- ✓ **Custom signature entries:** These entries pick up where standard IPS signatures may not provide complete protection, mainly against new or unknown attacks.
- ✓ **Out-of-band mode:** Also called one-arm IPS mode, an IPS can operate in this mode as an intrusion detection system, detecting but not acting on threats and attacks. Suspect traffic is analyzed on a separate switch port.
- ✓ **Packet logging:** This type of logging enables you to save network packets that match certain IPS signatures and analyze the log files with analysis tools.

Web content filtering

Web filtering lets you control what kinds of web content a user may view. By using web filtering, you can greatly reduce your exposure to spyware, phishing, pharming, inappropriate sites, website redirection, and other threats lurking on the Internet.

The web content filter feature scans the content for every web page that's accepted by a firewall policy. Content filters allow you to create a *blacklist* of banned words and phrases, and URL blocking lets you block unauthorized website addresses.

Category blocking is a third method of filtering web content, which relies on URL ratings to allow access to “good” sites and block access to “bad” sites.



Vendors understand the challenge of today’s security managers and administrators, so many vendors provide URL blacklists with their products to expedite this technology.

Anti-spam

Anti-spam filtering can block many Web 2.0 threats like bots, many of which arrive in your users’ e-mail boxes. Multiple anti-spam technologies incorporated into UTM can detect threats through a variety of techniques, including

- ✓ Blocking known spammer IP addresses to prevent receipt of e-mail from those addresses.
- ✓ Blocking messages where any URL in the message body can be associated with known spam.
- ✓ Creating a message hash, then comparing that value to hashes for known spam messages. Those that match may be blocked without knowing anything about their content.
- ✓ Comparing the IP address of the client and e-mail address of the sender to those on a blacklist/whitelist specified in a protection profile. Those on the whitelist get through; those on the blacklist are blocked.
- ✓ Conducting DNS lookup for the domain name when starting an SMTP session to see whether the domain exists or is blacklisted.
- ✓ Blocking e-mail messages based on matching message content with the words or patterns in a selected spam filter banned-word list.

Data loss/leakage prevention

Data loss prevention (DLP), also referred to as data leakage protection, helps prevent intentional or unintentional transfer of information to someone outside an organization. It can also apply to data that should remain within a department of an organization, such as personnel files or accounting data.

UTM DLP enables you to control data via inbound and outbound filtering, fingerprinting, and more. DLP filtering scans inbound and outbound files, looking for text strings and patterns, and then allow, block, or archive content based on matches with the DLP database. With fingerprinting, each document file is assigned a unique fingerprint. Based on the fingerprint, DLP prevents sensitive documents from being shared beyond the network.

Anti-virus

Anti-virus technology provides multi-layered protection against viruses, spyware, and other types of malware attacks. It enables you to scan e-mail for viruses, but it doesn't stop there. You can also apply anti-virus protection to File Transfer Protocol (FTP) traffic, instant messaging (IM), and web content at the network perimeter. Some solutions support Secure Sockets Layer (SSL) content scanning, which means that you can protect the secure counterparts to those types of traffic as well, such as HTTPS, SFTP, POP3S, and so on.



Need help sorting out the alphabet soup? HTTPS is short for “Hypertext Transfer Protocol Secure,” and POP3S stands for “Post Office Protocol 3 over Secure Sockets Layer.” SFTP means “Secure FTP.”

Essentially, a UTM virus filter examines all files against a database of known virus signatures and file patterns for infection. If no infection is detected, the file is sent to the recipient. If an infection is detected, the UTM solution deletes or quarantines the infected file and notifies the user.

Some configurable options in anti-virus profiles include

- ✔ **Anti-virus signature database:** Some vendors offer a choice of signature databases, so that you can determine the right tradeoff between performance and protection. A larger database increases accuracy but lowers system performance because it scans more records for matches.
- ✔ **File pattern:** Checks the filename against file pattern settings configured for the system.
- ✔ **File size:** Checks whether messages or attachments exceed a user-configurable threshold.

- ✔ **File type:** Applies a file type recognition filter that checks files against user-configured file type settings.
- ✔ **Grayware:** Checks any files that survive file pattern matching and virus scans for spyware.
- ✔ **Heuristics:** Checks files for virus-like behaviors or other known virus indicators.
- ✔ **Virus scan:** Checks any file that passes the file pattern scan for viruses.

Securing the Network End to End

Organizations need to secure their networks from the perimeter to the core. We discuss UTM features such as application control and web filtering earlier in this chapter. In this section, we address technologies that add more muscle to a UTM strategy all the way to the end-user level: integrated wireless LANs, endpoint control, and secure remote access.

Deploying integrated wireless LANs

Using separate LAN and wireless LAN (WLAN) security systems makes consistent policy enforcement tricky and causes potential blind spots. To control all traffic, both wired and wireless, you need a single set of policies.

A WiFi controller lets you integrate wireless networks into your current network infrastructure. Each WiFi network or service set identifier (SSID) is represented by a virtual network interface. You can impose the same UTM-based security policies and controls on the virtual interface as you do on a wired environment, eliminating potential blind spots. This approach also detects and eliminates wireless access points, helping you to prevent rogue or unauthorized access points from attaching to your network.



An integrated WLAN meets regulatory requirements like those included in the Payment Card Industry Data Security Standards (PCI DSS).

Protecting information in transit: IPsec and SSL VPNs

VPN technology allows you to encrypt traffic as it moves across the Internet from remote devices to corporate access points, preventing anyone who might intercept it from viewing its contents. There are two types of VPNs: IPsec and Secure Sockets Layer (SSL).

SSL-based VPNs allow connectivity from almost any Internet-enabled location using a web browser and its native SSL encryption. There are two different types of SSL VPNs:

- ✔ Clientless SSL VPNs require no specialized software and are ideal for deployments where you cannot easily provide a “client” to the user, such as employee-owned PCs, contractors, and business partners. Only web-enabled and some client-server applications, such as intranets, applications with web interfaces, and e-mail, can be accessed using a clientless connection.
- ✔ SSL VPN clients provide more complete access to network resources. A client resides on the remote device and allows access to the same applications and network resources that employees use in the office.

IPsec-based VPNs differs from SSL VPNs in how an encrypted tunnel is formed. However, like client-based SSL VPNs, IPsec VPN connections are established using pre-installed VPN client software on the user desktop. IPsec VPNs are also used for site-to-site connectivity where remote branch offices are connected over the Internet.

If you support a variety of remote users, consider deploying both SSL and IPsec technologies. For example, you could place an IPsec client on those systems under your control and use clientless SSL-based connections for those systems you don't control.



Encryption alone doesn't offer complete protection. Just because an outsider cannot decipher encrypted data as it passes over the Internet doesn't mean such information is “safe.” Without other protection, malware and legitimate traffic can traverse a “secure” VPN tunnel.

Are you really who you say you are?

Authenticating users before they access your network is vital. Unfortunately, in many situations, simple usernames and passwords don't cut it. Organizations often require stronger authentication for groups of users who have access to high value systems (network administrators or the Treasury group in Finance).

Two-factor authentication relies on the concept of "something you know" and "something you have" to provide stronger authentication than simple username and password combinations. The "something you know" is the user's own

individual name and password. The "something you have" is typically a token the user possesses; the token is either a small device or an application running on a mobile phone. Upon entering a password, the token provides a one-time password synchronized with an access server at headquarters. That way, any unauthorized individual must possess both the user's individual password and the token itself to gain access.

Advanced UTM devices support two-factor authentication, which greatly reduces the risk that unauthorized individuals can access network resources.

Just like being in the office: WAN optimization

Users accessing a network from a variety of locations expect their "office" environment to exist wherever they happen to be. This can be a problem because Internet speeds are often inadequate for data-intensive applications, or because of the cost of high-speed links. In such situations, WAN optimization plays a critical role using various techniques to improve WAN performance. These include protocol optimization, byte caching, web caching, SSL offloading, and secure tunneling to deliver improved application and network performance to make remote workers more efficient.

Here's more information on these techniques:

- ✔ **Protocol optimization:** Improves the efficiency of traffic that uses FTP, HTTP, TCP, and other protocols, accelerating network performance.

- ✔ **Byte caching:** Caches files and other data to reduce the amount of data transmitted across the WAN.
- ✔ **Web caching:** Stores or caches web pages and then serves them upon request to reduce latency and delays between the WAN and web servers.
- ✔ **SSL offloading:** Offloads SSL decryption and encryption from web servers onto SSL acceleration hardware, accelerating the performance of the web servers.
- ✔ **Secure tunneling:** Secures traffic as it crosses the WAN.

To reduce costs and make management easier, look for a security appliance that also offers WAN optimization features. Then, install the appliance between two private networks and the WAN. For best performance, install two appliances, one on either side of the WAN link, as shown in Figure 5-1. This configuration optimizes all traffic that traverses the WAN.




Figure 5-1: WAN optimization units located on both sides of a WAN and in front of both private networks.

Controlling endpoints

The “office” today no longer applies only to a building that houses employees and cubicles. Employees are increasingly on the move, using laptops, tablets, and smartphones to

access corporate data. They need corporate network access, but the mobile nature of the devices they use also presents added security challenges.

To control endpoints, use a remote access connectivity solution that:

- ✓ Incorporates IPsec and SSL VPN technology for secure remote communications
- ✓ Enforces policies on endpoints, such as disabling file sharing application or ensuring that an anti-virus application is running and up to date
- ✓ Uses web filtering and other UTM technologies to reduce potential for exposure to malicious web content and sites
- ✓ Uses a security appliance with built-in WAN optimization

Enforcing corporate data policies on all remote user connections helps make these endpoints as secure as possible.

Use Cases

Now that you understand how UTM works and its benefits, here's a look at how some organizations — from finance to retail to government agencies — have improved their security and productivity by putting UTM to work.

Midland Federal Savings & Loan

Part of Chicago's neighborhood fabric, Midland Federal Savings & Loan offers checking and savings accounts, loans, insurance, and more to local businesses and consumers. More than 100 employees work in its head office and three of its branches in southwest Chicago.

Protecting critical assets like customer banking and insurance information is a core requirement for this bank. A few years ago, Midland took a hard look at its IT infrastructure and security. The weakest point was that the three branches shared an Internet connection with the head office. If that connection went down or was compromised, the entire organization would be affected.

Working with a third-party solution provider, Midland decided to invest in a comprehensive UTM solution. It included VPN access, IPS perimeter anti-virus, spyware and malware protection, spam and e-mail content filtering, and website filtering.

This solution provided a higher level of protection for customer data. Employee productivity increased, too. Spam filtering, for example, blocked more than 500,000 spam messages per month. Likewise, the web filtering system blocked thousands of unauthorized websites monthly. Not only did this keep employees focused on their work, but it prevented malicious website redirection and reduced legal liability as well.

Because the UTM solution streamlined administration tasks and increased network performance, Midland Federal saw a return on its investment in the first year after implementation.

Valvoline Instant Oil Change

Many folks know about Valvoline Instant Oil Change (VIOC), a quick-lube chain with more than 800 shops in the USA. VIOC is based in Lexington, Kentucky, and is owned by Ashland, Inc.

VIOC's legacy dial-up network wasn't cutting it so the company decided to upgrade to a wireless network. Because VIOC accepts credit cards as payment for services, it must follow the Payment Card Industry Data Security Standard (PCI DSS). This standard includes stringent IT and security requirements.

Each shop required a separate, secure wireless network for its point of sale system, for wireless handheld scanners used by employees to service customers in the parking lot, and for wireless laptops for manager use. VIOC also wanted to offer free WiFi to customers waiting for oil changes. To release VIOC from liability, management needed a splash page that listed usage rules and required customer consent.

The company deployed an integrated solution at its Kentucky headquarters that included centralized network management, along with VPN and web content filtering gateways, and a backup appliance at an Ohio location for redundancy. Each VIOC shop uses a WiFi appliance to provide wireless infrastructure and secure SSIDs. Web content filtering prevents customers from accessing "inappropriate" sites.

This new solution lets Ashland meet regulatory compliance while its shops provide better customer service and have also increased employee efficiency.

PSC Info Group

Headquartered in historic Valley Forge, Pennsylvania, PSC Info Group provides statement printing and mailing, data management, and information management services. It has more than 1,000 clients that include healthcare providers, financial firms, government agencies, and utilities.

PSC handles upwards of one million client documents daily, receiving most of them over the web or via e-mail. Because many of its clients are in regulated industries, PSC requires tight security to protect client data.

The company had begun feeling some pain regarding security and expansion. For example, because of the web-based nature of file transfers, networks were sometimes overwhelmed with viruses. Also, the company wanted to switch from a site-to-site VPN network to a multi-protocol label switching (MPLS) to simplify management.

PSC chose an integrated UTM appliance with anti-virus, firewall, and IPS functions. The company installed a few of the appliances at entry points on the network. Because these new appliances did a better job of managing the network, PSC was able to eliminate many of its other security devices from different manufacturers. Consolidation reduced administration time and costs and the rigmarole of dealing with multiple vendors, contracts, and licenses.

Its new UTM appliances were automatically and continually updated and stopped web and e-mail threats from reaching their internal networks. PSC saw a dramatic decrease in the amount of malware they had been receiving: down from more than 5,000 viruses daily to less than 10, a 99.8 percent reduction!

Chapter 6

Ten (Okay, Eleven) Key Questions to Ask When Evaluating UTM Solutions

In This Chapter

- ▶ Checking performance levels to separate facts from marketing hype
 - ▶ Getting the security features you need today and tomorrow
 - ▶ Seeing how integrated and centralized management features expand IT's reach
 - ▶ Understanding the value of UTM
-

UTM is a fast-growing market. Research firm IDC expects UTM products to occupy 33.6 percent of the total network security market by the end of 2012. That huge percentage means a lot of your competitors are moving to UTM or have already installed it on their networks. Is it time to adopt UTM?

Because anything related to network security is an important decision, you should research and compare vendor offerings carefully. It's not easy. This chapter covers a number of key questions to ask before making any UTM purchasing decision.

How Does the UTM Solution Perform?

UTM products come in different sizes and “strengths.” Each product is designed to meet certain performance levels — throughput, latency, and so on — based on connection speed.

As networks migrate from technologies of 1 GbE to 10 GbE and faster, their security infrastructure has to keep pace. Otherwise, it becomes a bottleneck, slowing down traffic or missing threats because it can't keep up.

The key to high performance lies in a UTM's design. Components should be integrated, modular, and carefully orchestrated. Look for a vendor that offers a purpose-built solution, with high-performance hardware matched to specific security software and networking services. A solution should include network-level custom processes, such as firewalling and bandwidth control, as well as content-level custom processes, like web filtering and anti-virus protection, for maximum versatility and performance.



Check whether each vendor's offerings are verified by rigorous third-party certifications. As you narrow your list, read product assessments from third parties like ICSA Labs or NSS Labs for access to objective testing results. You also need to "kick the tires" to find out if a product performs as expected in your environment. Select three or four solutions to test in your shop. There's no better way to find out what works and what doesn't before you make a final selection.

Which Security Technologies Are Included?

One of the main benefits of a UTM solution is integration of multiple security features like firewall, application control, intrusion prevention, VPN, and other elements into a single product. But not every vendor offers the same feature set or the same functionality across all models in its product family.

Assuming that your shop needs well-rounded protection from daily threats bombarding your network, look for a UTM solution that includes the following features:

- ✓ Firewall
- ✓ Application control
- ✓ IPsec and SSL VPN
- ✓ IPS
- ✓ Web content filtering
- ✓ Anti-spam
- ✓ Data loss/leakage protection
- ✓ Anti-virus and anti-spyware protection
- ✓ IPv6 native support
- ✓ Traffic shaping/bandwidth control

Choosing the right mix of security technologies means knowing your needs today and into the future. What core features should you consolidate into a single device to replace a patchwork of other devices from multiple vendors? What new features are you likely to need in the next two to three years? Do you have remote users or branch offices to support over VPNs? Do you want to use a single management console to enforce policies for remote offices and branch offices, as well as the central office? All this makes a difference!



Newer technologies, such as application control, are also important considerations. Owing to the explosion in the number of web-based applications, organizations of all sizes want to be able to enforce granular policies — what applications users have access to, what they can do within those applications, and so forth. Find out how your short list of UTM solutions handles application traffic. Is the feature fully integrated into the management console, or provided as an add-on service?

A mature UTM solution should address all elements of IT infrastructure security, including networks, systems, services, applications, and data. Coverage should include everything from the network perimeter to the core, from DMZ web servers to choke points on internal networks, and to data on mobile devices.

Which Network Features Are Supported?

A UTM solution must be able to support a wide variety of network topologies and features — whatever you have in place now or expect to implement over the next few years.

Find out whether the UTM solutions you're researching allow you to monitor network latency and throughput, and automatically correlate it to event and network logging.

Finally, look for a UTM solution that lets you segment your networks for greater policy control and event isolation.

Does UTM Support IPv6 Natively?

With the exhaustion of IPv4 address space, IPv6 adoption is accelerating. It is essential for your network security solution to deliver the same functionality for IPv6 traffic as for IPv4 traffic to detect and handle threats delivered via IPv6 protocols.

Does your prospective UTM solution handle IPv6 seamlessly? Can it inspect IPv6 traffic or does it simply pass it along untouched? A UTM solution should natively support IPv6 in its firewalls, Domain Name System (DNS) in transparent mode, and within Session Initiation Protocol (SIP). A good solution should also be able to handle dynamic routing for IPv4 (such as RIP, OSPF, IS-IS, BGP, and multicast protocols) and IPv6 (such as RIPng, OSPFv3, and BGP4+), as well as policy-based routing.

Does It Play Well with Virtual Environments?

Some UTM solutions offer a flexible way to deal with virtual local area networks (VLANs), also known as virtual domains (VDOMs), and virtual appliances. VDOMs let you create separate security

domains for multiple departments or even separate organizations. Within each domain, you create separate zones, firewall policies, user authentication, and VPN configurations.

If you run a *multi-tenant* network or multiple networks, ensure that the UTM solution you're considering offers virtual network segmentation. This could be a big time-saver for your IT staff and boost the efficiency of your infrastructure. Also, if you have deployed virtual machines as a way to improve scalability and flexibility of your IT environment, make sure that your solution supports virtual as well as physical appliances.

Is the UTM Solution Scalable?

Are you planning to expand within the next year or two? Organizations change all the time, mainly to meet customer demands or differentiate themselves from competitors. Expansion is often part of that change.

Whatever UTM solution you choose should be flexible. You should be able to use the features you need now and “flip the switch” or add other features or more users down the road. A modular approach doesn't ding performance the way that some all-or-nothing solutions do.

Does It Offer High Availability?

With today's reliance on continuous communications, every organization, of any size, needs to be concerned about availability. A device failure can be devastating to any organization. High availability ensures that your network security is constantly available, even in case of a device failure.

When it comes to ensuring high availability, you can choose between active-active or active-passive architectures. An *active-active architecture* consists of one primary device that receives all communication sessions and load balances them among the primary unit and all of the subordinate units. An *active-passive architecture* consists of a primary device that processes data and one or more subordinate devices. The subordinate devices run in standby mode until needed. Also, look for bypass ports on the device itself to provide “fail open” network availability in the event of a device failure.

Clustering UTM solutions is a great way to achieve high availability. You connect cluster units to each other and to your network through switch. This redundant setup ensures that network traffic flows freely even if one unit fails. Without clustering, an organization with a single UTM unit that goes down can be cut off from the Internet, its branch offices, and remote users. If the branch offices route their traffic through the main office, they'd also be isolated.



For even higher availability, consider adding another Internet connection from a different Internet service provider (ISP) for redundancy.

What Kind of Management and Reporting Does It Include?

Integrated, centralized management is a must-have for UTM, but it's been hard to find in many solutions. Administrators must be able to use one interface to “see” the entire network and make configuration changes to policies and profiles. And within that system, policies for one application should be able to manage or affect other policies automatically. For example, web filtering, application control, and IPS profiles should be covered in firewall policies. That way, it's much easier to push updates from a central console to multiple devices.

A UTM solution with integrated management should also be able to share knowledge of events between modules. For example, if an anti-virus feature blocks or quarantines files, other modules should be made aware of such action. Updates made to one module should carry over to other modules.

On the reporting front, look for a UTM solution that offers a library of templates to generate commonly used reports with just a few mouse clicks. You also want the ability to generate custom reports quickly and easily. To cut down on administration effort and get the most from your purchase, look for a UTM solution that includes the following management features:

- ✔ **Centralized management:** Lets you remotely manage multiple devices at once and tweak policies and profiles across different modules.
- ✔ **Consolidated management:** “Single pane of glass” lets you remotely manage multiple devices at once and tweak policies and profiles across different modules.
- ✔ **Advanced management:** Lets you perform higher level management tasks, like role-based administration, detailed logging and reporting, and so on.

Of course, the UTM solution’s user interface must be web-based for ease of use and universal accessibility.

If you have a distributed network, check to see whether the UTM analysis solution offers both collector and analyzer modes. This way, analysis devices in remote locations collect the data and send it to the main analyzer in a store-and-forward manner. Such a solution makes it much easier to deal with the large amounts of data that can be generated by a large organization.

How Does the UTM Solution Stay Ahead of Threats?

With thousands of variations of malware, malicious sites, and exploits in the wild, today’s UTM solution must use a combination of cutting-edge technologies to beat back attacks. Vendors must also continuously send updates to their customers to ensure that customers’ devices and networks are fully protected.

UTM vendors enable customers to subscribe to security services and updates to continuously update their systems. However, not all research teams are created equal. For example, you’ll want to pursue the following topics:

- ✔ Which vendors have their own research teams that track emerging threats from around the world and determine how malware works? Relying on third-party updates can slow things down.

- ✓ Do the vendors have research teams in different countries, and do they deal with software and hardware vendors directly?
- ✓ What is each vendor's response time for taking action on newly discovered threats or their variants?

How Much Does It Cost?

The capital expenditure for any network security solution is a difficult deciding factor for buyers — but it's always a key element in the selection process. Piecing together an ad hoc security system that includes firewall, VPN, application control, intrusion detection, content filtering, anti-virus, and traffic shaping can quickly add up. There are also maintenance fees and subscriptions to consider for anti-virus, IPS, and content-filtering updates.

The more demanding your traffic needs, the more you'll have to pay, regardless of the technology. A UTM solution allows you to lower your costs by consolidating multiple systems, thereby lowering your total cost of ownership (TCO). Don't forget to take into account the cost of your IT staff. Because a new UTM solution is more efficient to operate and administer, your staff is free to perform other activities.



Not all organizations need every security feature available, which can drive up the capital cost of the solution and the costs involved with installation and maintenance. Pick the ones you need, and turn on others as and when you need them.

Simple licensing or not?

Check whether the UTM vendors require per-user licensing. Not all vendors do. Per-user licensing, or the need to license additional modules, can add more costly line items to your UTM budget. Look for a vendor that offers an unlimited number of users and features without additional licensing.

Weighing costs versus benefits

Be sure that you consider all of the benefits of a UTM solution instead of only the sticker price. Performance, scalability, and management are three of the most important factors that can tip the decision from one solution to another.

How much is a solution worth that can virtually eliminate throughput and latency problems, even with a high number of concurrent users? What if you need to expand the system fairly quickly? What is the cost benefit of being able to configure changes to handle growth rather than starting a new round of technology research? How important is centralized and integrated management that frees IT staff to focus on more pressing issues?

Streamlined performance, scaling to meet growth, and effective centralized management can be critical to organizations with lots of users to support.

What Support Options Are Available?

Your organization's IT staff is busy handling a wide variety of equipment and services. They can't know all the details of every piece of hardware or software in their infrastructure. Plus, you have to rely on the vendor for periodic updates to the firmware and software.

For the best shot at a successful UTM implementation, you need a vendor who gives you the support you need, whether it is around-the-clock global technical support or regional workday/work week support. The vendor should have a superior response rate, qualified personnel, and a reputation for strong customer service.

In addition, your organization may need onsite consulting for migration planning and implementation. Such specialized support often comes with a big price tag. But a vendor's engineers have deep skills in the technologies and software you're preparing to install, which can greatly speed up the entire process. Onsite consultants usually offer insights and tips to customers that extend well beyond the tasks at hand.

What about training?

Find out what training options are available, either directly through a vendor or a third party that's certified on the vendor's products.

Online training is essential for both controlling costs and giving busy professionals the ability to fit training into their schedules, rather than the other way around. This type of training should include high-quality videos and directed materials that come close to a classroom experience.

Can the vendor support global operations?

The global nature of some organizations makes even routine implementations and maintenance a challenge. If you operate in multiple countries, find out which of the vendors you're researching have local resellers and support teams in those countries. When an office's network security is at risk, you need solutions and support available as quickly as possible.

What are current customers saying?

Always check vendor references. If time allows, visit several companies that are using different UTM solutions. Find out about the level of administration needed to operate the products, how difficult they are to use, and other concerns that come up while researching products.

Organizations with high-volume environments need experienced vendors with superior networking products. When narrowing down your UTM choices, consider vendors with happy customers across a wide range of customers, industries, and government agencies.

Glossary



active-active architecture: Consists of one primary device that receives all communication sessions and load balances them among the primary unit and all of the subordinate units.

active-passive architecture: Consists of a primary device that processes data and one or more subordinate devices. The subordinate devices run in standby mode until needed.

advanced persistent threat (APT): An attack that uses multiple techniques to accomplish its goal of penetrating a network and compromising systems, including spam e-mails, fake pop-up alerts, and malicious social media content. Also known as a multi-vector attack.

blacklist: A list of banned words and phrases, or forbidden network and Internet sites.

bot: Short for robot, an individual computer compromised with malware that takes instructions from a command and control server in a botnet.

botnets: A collection of bot computers under the control of a botmaster. A botnet can include hundreds of thousands of bots around the world.

botmaster: Also called a bot herder, the person who operates a botnet. A botmaster takes instruction from its criminal clientele to program a botnet to undertake various criminal projects.

content processor: A processor designed to make high-speed comparisons of objects contained in network traffic to the characteristics of already known threats. The objects may be network packets, compressed files, or other items in the traffic stream.

data loss prevention (DLP): Also referred to as data leakage protection, a technology that helps prevent intentional or unintentional transfer of information to someone outside an organization.

deep packet inspection (DPI): Examining the payload or data portion of a network packet as it passes through a firewall or other security device. DPI identifies and classifies network traffic based on signatures in the payload.

defense in depth: An approach to security in which controls and technologies are established and managed across an entire network to create layers of security. An attack that gets past any particular countermeasure must get through multiple layers of security to access services or resources on the network.

drive-by downloads: File downloads infected with bots, Trojans, or viruses. Web users are exposed to drive-by downloads by clicking a link that is associated with a malicious website.

encapsulation: Wrapping encrypted data inside a protocol envelope for safe transfer over a network or the Internet.

firewall: A device that inspects inbound and outbound traffic on a network, allowing safe traffic and denying unsafe traffic.

flow: A set of packets that share common properties, such as the same source and the same destination. A flow can be a sequence of packets in a single session, or span multiple concurrent sessions.

flow-based inspection: Also called stream-based inspection, analyzes flows entering a UTM device before packets make their way to an internal interface.

latency: Delays in the speed of data through a device or network, caused by slow devices or multiple inspection processes.

malicious mischief: A common online crime that may include damage or theft for the thrill of it, instead of a profit motive. Malicious mischief stems from political or personal malice.

multi-tenant network: A large network, such as a data center network, that is logically divided into smaller, isolated networks.

multi-vector attack: See *advanced persistent threat*.

network processor: A processor that operates at the interface level to deliver an extremely low latency data path.

packet defragmentation: Reassembling fragmented packets. Network processors can quickly reassemble fragmented packets and examine them for their threat potential.

polymorphic attack: An attack in which malicious code changes regularly to evade detection.

port number: A numerical identifier, associated with a network address, to categorize network traffic and enforce policies.

proxy-based inspection: A UTM technique that reassembles files before performing security inspection.

proxy server: Hardware or software that sits transparently on a network and checks all packets for policy compliance before forwarding any of them.

ransomware: A type of malware that targets mission-critical systems or applications, locking essential files and disabling recovery/repair features and capabilities. Victims must pay the ransomware developer to “unlock” these files to use the infected system or application.

security processor: A multi-core, multi-threaded processor combined with content processor logic. A security processor operates at the interface level and adds packet processing and acceleration options.

signature: In network security terms, a collection of characteristic patterns found in a packet’s payload.

spear phishing: A targeted attack directed at specific employees, customers, or members of a particular organization.


stand-alone security device: A specialized security appliance or software. May act as a workaround to the limitations of a traditional firewall.

stateful packet inspection firewall: A firewall that keeps track of the state of network connections, allowing packets that correspond only to a valid, active connection between a client and server.

traffic shaping: A method of guaranteeing network bandwidth to certain types of traffic.

whitelist: A list of trusted network or Internet sites.

The threats stop here.



FORTINET®

Consolidated Network Security

Every day, Fortinet protects the networks of many of the largest and most successful organizations in the world. We secure their applications, users, and data from hidden threats. Our consolidated security technologies deliver unmatched performance and visibility. Learn how you can increase security, improve performance, and reduce complexity of your network.

Visit us at **www.fortinet.com** or call **1-866-868-3678** to find out how you can protect your network today.

Protect your network against emerging security threats!

No network security plan is complete without unified threat management, or UTM. UTM provides organizations with a very real way of addressing today's sophisticated security threats that continue to evolve. By incorporating UTM into your network security, you have unmatched protection against both known and emerging threats.

- **Why legacy security systems don't protect you** — what they can and can't provide
- **How network security evolved** — the evolution of security from firewalls to hackers to crimeware, and more
- **How security consolidation improves performance** — the benefit of purpose-built hardware and software
- **What deployment strategies are being used** — examine case studies to see how UTM works in the real world
- **How to evaluate new UTM systems** — what key questions should you ask before implementing a new system

Ed Tittel is a freelance writer, consultant, and expert legal witness. A veteran of the technology trade press, Ed has written or contributed to more than 140 books, including *Windows Server 2008 For Dummies*, *HTML, XHTML, and CSS For Dummies, 7th Edition*, and many more. When not writing, you can find him playing pool or in the kitchen with his lovely wife, Dina, cooking like crazy.



Open the book and find:

- How the threats continue to evolve and their impact on the user landscape
- Why traditional approaches to network security cannot keep up with the evolving threat landscape
- How you can get the control and visibility you want over your network, applications, and users
- What features and capabilities you need in your UTM system

Go to [Dummies.com](https://www.dummies.com)
for videos, step-by-step examples,
how-to articles, or to shop!

For Dummies®
A Branded Imprint of



ISBN: 978-1-118-08701-5
Book not for resale