The Laws of Vulnerabilities: Six Axioms for Understanding Risk

Global Data from 40 Million Security Scans over 40 Months Define Behavior of Vulnerabilities for Insight on Protecting Networks

CONTENTS

EXECUTIVE SUMMARY

A few years ago security professionals lived in constant reaction to sudden vulnerability exploits such as LoveLetter, SoBig, Slapper, Slammer and Blaster. Devising strategies to prevent exploits was difficult due to limited insight about the behavior of vulnerabilities over time. A rising volume of attacks posed other challenges. The use of automated attack tools eventually placed all Internet-connected systems with vulnerabilities under continuous attack, so CERT even stopped tracking the rising number of reported incidents after 2003.

The key difference today is that security professionals can have deeper insight and more technical options to proactively stop vulnerability exploits. Understanding the "enemy" is vital to winning a conflict. Understanding the behavior of vulnerabilities is essential to set effective security strategy and proactively implement security solutions.

This paper describes The Laws of Vulnerabilities, which are six axioms about the behavior of vulnerabilities gleaned from a continuous long-term research project launched by Qualys in 2002. We analyzed a global data pool of more than 40 million IP scans with QualysGuard, which is Qualys' on demand vulnerability management and policy compliance service. Data analysis revealed The Laws of Vulnerabilities, described below. Insight from The Laws helps security professionals to prevent exploits of IP-related vulnerabilities.

The Laws of Vulnerabilities

Half-life – Vulnerability half-life is 19 days on external systems and 48 days on internal systems; it doubles with lowering degrees of severity.

Prevalence – Half of the most prevalent critical vulnerabilities are replaced by new vulnerabilities each year.

Persistence – The life spans of some vulnerabilities are unlimited.

Focus – Ten percent of critical vulnerabilities cause nearly all exposure.

Exposure – The time-to-exploit cycle is shrinking faster than the remediation cycle.

Exploitation – Nearly all damage from automated attacks is during the first 15 days of outbreak.

Source: Qualys

DATA AND METHODOLOGY

The goal of Qualys' research was to understand how critical vulnerabilities behave over time in the real world.

Data was automatically and anonymously drawn from the largest collection of vulnerabilities in the world – the Qualys KnowledgeBase for QualysGuard. The KnowledgeBase contains signatures and identification statistics for more than 4,800 network vulnerabilities of varying severity within these categories:

Vulnerability Categories in Qualys KnowledgeBase	
Back Doors and Trojan Horses	Appliances
Brute Force Attacks	Information / Directory Services
CGI	SMB / Netbios Windows
Databases	File Sharing
DNS and Bind	SMTP and Mail
eCommerce Applications	Applications
File Transfer Protocol	SNMP
Firewalls	TCP/IP
General Remote Services	Web Servers
Hardware and Network	X-Windows

40,691,913

Total IP scans

45,378,619

Total critical vulnerabilities identified

1,595

Unique critical vulnerabilities identified

Data for this analysis was derived from 40,631,913 IP scans with QualysGuard conducted globally during the period of 8 September 2002 and 31 January 2006. About 70 percent of the data was from global enterprise scans and 30 percent from random trials of QualysGuard. All scan data was anonymously gathered without correlation to any specific user or system.

There were 45,378,619 critical vulnerabilities identified by these scans. A critical vulnerability provides an attacker with the ability to gain full control of the system, and/or leakage of highly sensitive information. For example, critical vulnerabilities may enable full read and/or write access to files, remote execution of commands, and the presence of backdoors. QualysGuard assigns vulnerabilities like these a rating of Level 4 or 5 – the most severe threats to network security. Vulnerabilities can stem from bad code, a variety of malware, or from errors in system or network configuration.

The scans identified 1,595 unique critical vulnerabilities out of 1,972 in the KnowledgeBase. This means 80 percent of known critical vulnerabilities showed up in real world scans.

Data during the last year and a half of the testing period was enhanced by rising scan statistics for devices with internal-facing IPs. Initial scans with QualysGuard were restricted to devices with external-facing IPs. Qualys later added capability to scan IPs on the intranet using a distributed scanner appliance. Currently, one-third of the devices scanned by QualysGuard customers are inside the network perimeter on an intranet.

The figure below shows the study's quarter-by-quarter volume of IP scans with QualysGuard segmented by external and internal targets.

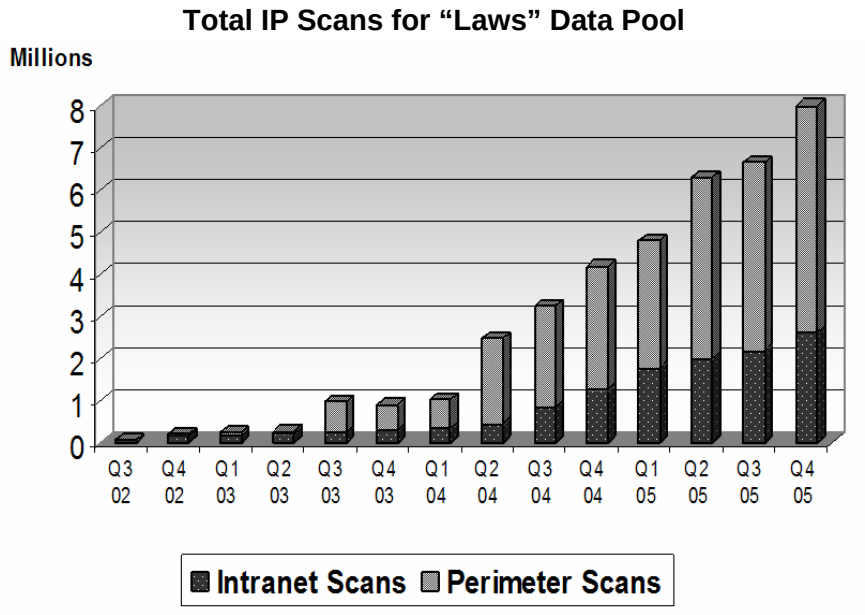

Qualys analyzed the vulnerability data with standard statistical techniques to identify:

- Window of exposure
- Lifespan of critical vulnerabilities
- Resolution response
- Trends over time
- Vulnerability prevalence

New Trend

60%+ of new critical vulnerabilities are in client applications.

Almost none are in a wireless device.

A major new trend is the shift in critical vulnerabilities from server to client applications. Earlier data in this analysis showed most vulnerabilities were in server applications such as web server, mail server, and operating system services. Data now show more than 60 percent of new critical vulnerabilities are in client applications such as web browser, backup software, media players, antivirus, Flash and in other tools. Vulnerabilities for client applications are subject to the same Laws of Vulnerabilities as those for server applications.

The data analysis also debunked a popular myth that wireless networks are a significant security vulnerability for enterprise networks. According to the data in this study, just one in nearly 20,000 critical vulnerabilities is caused by a wireless device. Future analysis will monitor this issue, especially as wireless becomes more widely adopted by enterprises throughout the world.

THE LAWS OF VULNERABILITIES

1 Half-Life

Vulnerability half-life is 19 Days on external systems and 48 days on internal systems; it doubles with lowering degrees of severity.

Half patched

Half of IPs with a critical vulnerability are still exposed after the respective 19 or 48 day half-life.

Half-life is the duration of half a process. The term often connotes danger. Half-life plays a critical role in protecting people, such as with radioactivity, or calculating the impact of improperly using an old drug. Half-life is equally important in understanding and preparing network defenses for malware and other vulnerabilities.

The data show that the half-life of critical vulnerabilities is shrinking. Our analysis for The Laws in 2003 found that half-life was 30 days, applicable mostly to external systems. Now the half-life for external systems has shrunk to 19 days. Half-life for internal systems is 48 days.

The meaning of these statistics is that for even the most dangerous vulnerabilities, it still takes organizations 19 days to patch half of vulnerable external systems. Patching half of internal systems takes 48 days – more than 150 percent longer than for patching external IPs! Exposure of unpatched systems continues during the significantly long period of half-life dissipation.

As an example of the Law of Half-Life, the illustration below shows a half-life time plot curve for the Microsoft Windows Color Management Module Remote Code Execution vulnerability following its appearance in July 2005. Scan data show incidents falling about every two weeks. Continued existence of the unpatched vulnerability triggered short bursts of new incident activity as the half-life curve fell over time.

Half Life of Microsoft Windows Color Management Module Remote Code Execution Vulnerability

CVE-2005-1219 – Released July 2005

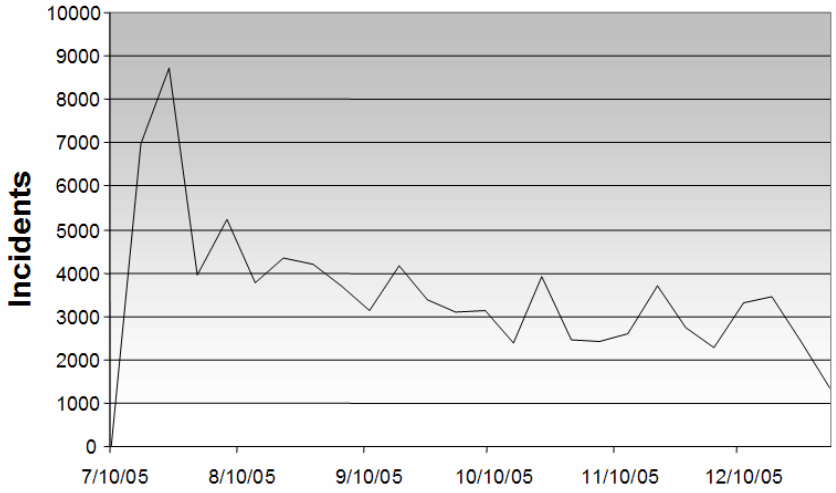

2 Prevalence

Half of the most prevalent critical vulnerabilities are replaced by new vulnerabilities each year.

Prevalence is the degree to which the vulnerability poses a significant threat. The ongoing global threats to people by dangerous viruses such as SARS or Avian Flu have significant prevalence until implementation of precautions reduces threats to a negligible level. With digital viruses and worms, we discovered similar trends.

Look Out

Half of critical vulnerabilities change every year.

We measured the prevalence of the most critical digital vulnerabilities for an 18-month period and learned that half are replaced by new vulnerabilities each year. This means there is ongoing change to the most important threats to our networks and systems.

The table below presents data to illustrate the Law of Prevalence. The most critical vulnerabilities found in customer security scans are listed by CVE number and ranked at three points in time: January 2005, June 2005 and January 2006. The gray highlight bars show a consistent shift in prevalence. After one year, eight of the top 13 vulnerabilities are still on that list. Five were replaced by new vulnerabilities.

Prevalence Table of the Most Critical Vulnerabilities

Vulnerability	CVE	Jan 2005	June 2005	Jan 2006
Microsoft Office XP Vulnerability Could Allow Remote Code Execution (MS05-005)	CVE-2004-0848	X		
Windows Media Player and Windows Messenger Remote Code Execution (MS05-009)	CVE-2004-1244	X		
Microsoft Server Message Block Remote Code Execution (MS05-011)	CVE-2005-0045	X	X	
Microsoft Windows OLE and COM Remote Code Execution Vulnerabilities (MS05-012)	CVE-2005-0047	X	X	
DHTML Editing Component ActiveX Control Remote Code Execution (MS05-013)	CVE-2004-1319	X	X	
Microsoft Hyperlink Object Library Buffer Overflow (MS05-015)	CVE-2005-0057	X	X	X
Microsoft Message Queuing Buffer Overflow (MS05-017)	CVE-2005-0059	X	X	X
Windows Multiple Denial of Service and Privilege Elevation Vulnerabilities (MS05-018)	CVE-2005-0061	X	X	X
Microsoft Exchange Server Remote Code Execution (MS05-021)	CVE-2005-0560	X	X	X
Microsoft SMB Remote Code Execution Vulnerability (MS05-027)	CVE-2005-1206	X	X	X
Microsoft Windows Web Client Service Remote Code Execution Vulnerability (MS05-028)	CVE-2005-1207	X	X	X
Windows Color Management Module Remote Code Execution (MS05-036)	CVE-2005-1219	X	X	X
Windows Plug and Play Remote Code Execution (MS05-039)	CVE-2005-1983	X	X	X
Windows Print Spooler Service Remote Code Execution (MS05-043)	CVE-2005-1984		X	X
Microsoft Windows Client Service For Netware Buffer Overflow Vulnerability (MS05-046)	CVE-2005-1985		X	X
Microsoft Plug and Play Remote Code Execution and Local Privilege Elevation Vulnerability (MS05-047)	CVE-2005-2120			X
Microsoft DirectShow Remote Code Execution Vulnerability (MS05-050)	CVE-2005-2128			X
Microsoft MSDTC and COM+ Remote Code Execution Vulnerability (MS05-051)	CVE-2005-1980			X
Microsoft Windows Graphics Rendering Engine WMF Format Code Execution (MS06-001)	CVE-2005-4560			X

3 Persistence

The life spans of some vulnerabilities are unlimited.

Many have experienced the frustration of having just patched a critical vulnerability, only to find that a variant exploit appears – and forces an immediate restarting of the patching process. The risk of re-infection also can happen when we deploy new PCs and servers with images of faulty unpatched operating system and/or application software.

Analysis of the data reveals that the life spans of some vulnerabilities are unlimited. One example is the SQL Slammer vulnerability, which demonstrated a nasty and persistent recurrence. Exploitation enabled a denial of service attack. Microsoft announced the existence of this vulnerability in July 2002 and published a patch at the same time. The chart below shows the first and biggest attack by a worm exploiting this vulnerability was in February 2003. The number of vulnerable systems dropped through March, then suddenly jumped to two-thirds of the original attack level and remained there for a few more months. Unpatched systems are still vulnerable to this threat today.

Persistence of MS-SQL 8.0 UDP Slammer Worm Buffer Overflow Vulnerability

CAN-2002-0649 – Released July 2002

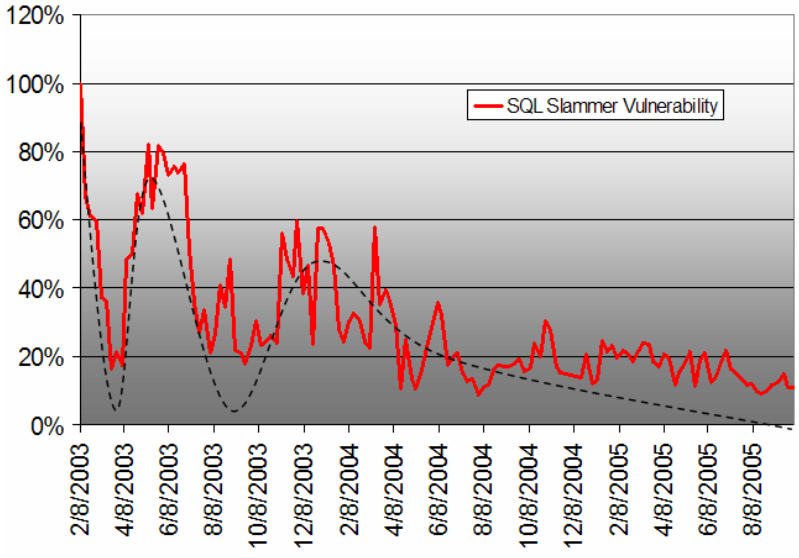

A dramatic demonstration of The Law of Persistence was the SNMP Writable vulnerability. Exploits of this vulnerability appeared in late 2002 and recurred with aggressive regularity for two years before subsiding in summer 2005.

Persistence of SNMP Writable Vulnerability

Multiple CVEs[1] – Released Feb. 2000

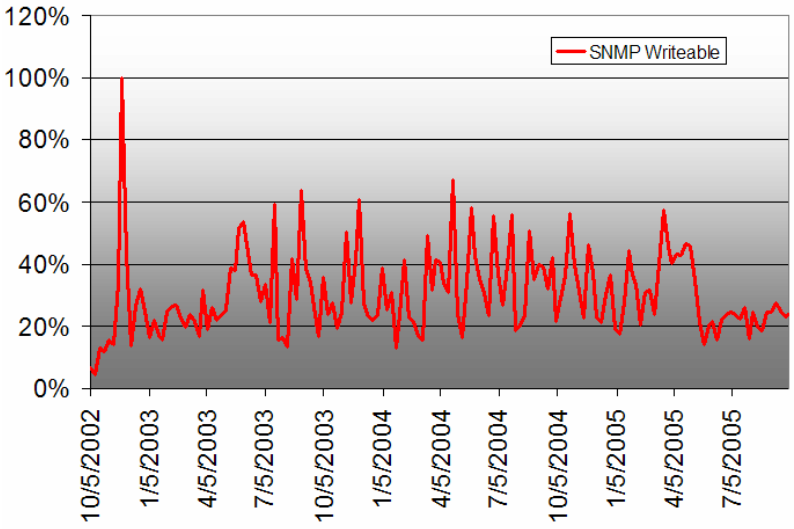

4 Focus

Ten percent of critical vulnerabilities cause nearly all exposure.

Work Smart

Patching the most critical vulnerabilities first eliminates most exposure.

The old 90 / 10 rule also applies to occurrence of critical vulnerabilities. The data in this study revealed that 90 percent of vulnerability exposure is caused by 10 percent of critical vulnerabilities. The figure below graphically depicts the exposure caused by 500 of the most critical vulnerabilities discovered during this study. The few distinct or small clusters of spikes correspond to specific vulnerabilities. Two had significant spikes of exposure: the Windows Null Session vulnerability and the NETBIOS / SMB Share Password vulnerability. Overall, only 10 percent caused significant exposure.

Security professionals can leverage the Law of Focus by targeting initial remediation efforts on critical vulnerabilities with the highest degree of exposure. Simply eliminating those vulnerabilities first will reduce 90 percent of the sources of risk.

[1] The SNMP Writable vulnerability had multiple CVE numbers, including CVE-1999-0792, CVE-2000-0147, CV-2000-0515, CVE-2001-0380, CVE-2001-1210, and CVE-2002-0478.

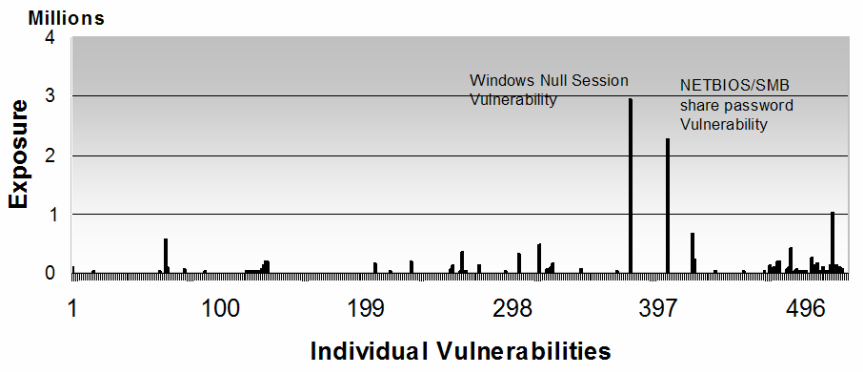

5 Exposure

The time-to-exploit cycle is shrinking faster than the remediation cycle.

Patch Faster

Accelerated exploits must be patched faster to eliminate system exposure.

Early data in this research project noted that that 80% of critical vulnerability exploits were available within 60 days of their public announcements. The updated Law of Half-Life shows this period is shrinking. Half-life is now 19 days for external systems and 48 days for internal systems. Since the duration of vulnerability announcement-to-exploit-availability is dramatically shrinking, organizations must eliminate vulnerabilities faster. The updated axiom restates the idea behind the Law of Exposure as 80 percent of critical vulnerability exploits are available within the first half-life after their appearance.

Exposure Curve of Critical Vulnerabilities

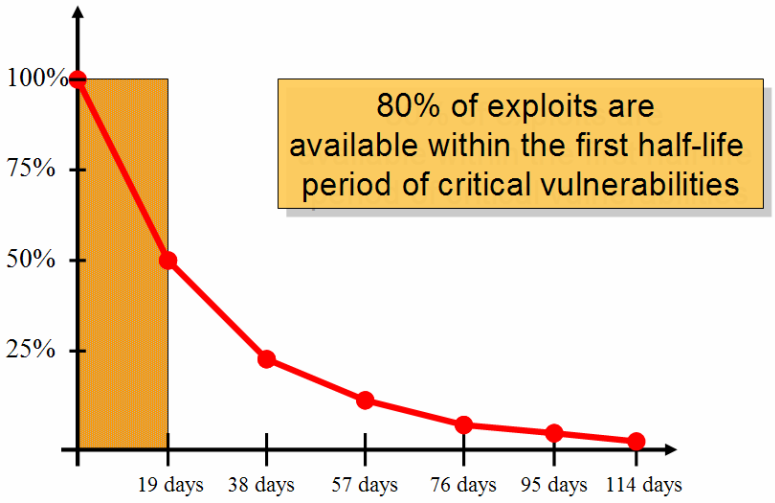

Some exploits are achieving the status of "zero-day" or "near zero day," meaning that the exploit is available on the same day of the vulnerability announcement. A recent example was the WMF vulnerability, also known as Microsoft Windows Graphics Rendering Engine WMF Format Code Execution (CVE-2005-4560). Exploitation of this vulnerability enabled execution of remote code and user account access. Exploitation was first observed in the wild on 26 Dec. 2005. Global scan data showed more than 50 websites were

infected two days later as exploitation took hold and then quickly expanded. Microsoft did not release a patch until 5 Jan. 2006.

The Zotob worm (CVE-2005-1983) is another recent example of quick exploitation. The worm is enabled by a stack-based buffer overflow in the Plug and Play (PnP) service for Microsoft Windows 2000 and Windows XP Service Pack 1. It allows remote attackers to execute arbitrary code and local users to gain unauthorized administrator privileges. Microsoft announced the vulnerability on August 11, 2005. Microsoft said exploit code became available the next day.

The rapid availability of exploits like these creates significant exposure for organizations until they patch all their vulnerable systems.

6 Exploitation

Nearly all damage from automated attacks is during the first 15 days of outbreak.

Automated attacks pose a special hazard to network security because they inflict damage swiftly with little time for reaction. The Law of Exploitation shows that severe damage from a vulnerability exploit is most likely to happen right after it appears. The most recent data show that initial period of severe damage is during the first 15 days of outbreak.

The graph below superimposes available outbreak data for six major vulnerabilities: Blaster, Code Red, Nachi, Sasser, Slapper and Zobot. For each critical vulnerability, the peak number of incidents occurs early after its respective appearance and swiftly drops off.

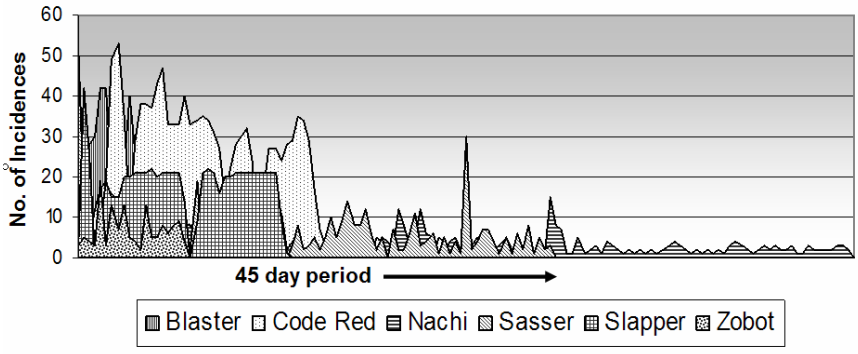

RECOMMENDATIONS

The Laws of Vulnerabilities demonstrate that known critical risks are far more prevalent than anyone has imagined. Data for our study document the persistent ability of attackers to gain full control of systems – including access to highly sensitive information such as financial data and intellectual property. The most effective thing organizations can do to mitigate fallout from vulnerabilities is to accelerate efforts to identify and remediate critical weaknesses. Continue use of

an automated vulnerability management system like QualysGuard will shorten half-lives of vulnerabilities and reduce risks for all organizations.

Qualys recommends that organizations regularly scan networks and systems for critical vulnerabilities and set a remediation goal of shortening the half-life by 20 percent by the end of 2006. Accomplishment of this goal will reduce the current half-life of external systems from 19 to 15 days, and of internal systems from 48 to 38 days.

ABOUT QUALYS

With more than 2,000 subscribers ranging from small businesses to multinational corporations, Qualys, Inc. has become the leader in on demand vulnerability management and policy compliance. The company allows security managers to strengthen the security of their networks effectively, conduct automated security audits and ensure compliance with internal policies and external regulations. Qualys' on demand technology offers customers significant economic advantages, requiring no capital outlay or infrastructure to deploy and manage. Its distributed scanning capabilities and unprecedented scalability make it ideal for large, distributed organizations. Hundreds of large companies have deployed Qualys on a global scale, including AXA, DuPont, Hershey Foods, ICI Ltd, Novartis, Sodexho, Standard Chartered Bank and many others. Qualys is headquartered in Redwood Shores, California, with European offices in France, Germany and the U.K., and Asian representatives in Japan, Singapore, Australia, Korea and the Republic of China.

Qualys, Inc.
1600 Bridge Parkway
Redwood Shores, Calif.
94065 — USA

800.745.4355
www.qualys.com