

When Malware Attacks (Anything but Windows)

Where is all the Macintosh malware? In a time when it would be conservative to say that one-quarter of all Internet-connected client PCs are compromised, it's curious that we've seen so little malware on Macs.

Dai Zovi and Shane Macaulay a MacBook Pro, Charlie Miller a MacBook Air, the individual teams US\$10,000, and the sponsor TippingPoint a pair of proof-of-concept attacks. Pwn2Own added a Windows machine to the competition this year, and it was compromised as well—many hours after the Mac fell.

Qualitative evidence shows that attackers no longer depend on vulnerable systems. Instead, modern attackers largely use social engineering to entice users to unwittingly relinquish control of their systems. Email-borne viruses, for example, have long depended on user interaction for their success. Over the past year, the Storm Worm crew has taken email-pushed malware to a new art, rotating across multiple pitches and attachment styles to build one of the largest botnets currently in existence. Their operation alone shows that the operating system's security posture has little to do with how many of the machines will be compromised.

We can accept the economic-motivation argument strictly due to Occam's razor, but it would be far more compelling if the current absence of large-scale malware could be explained through an economic model. Here, I describe a game theory-based model we at Cloudmark have developed to analyze the migration of attackers to emergent platforms.

Game theory: A brief introduction

Game theory is a mathematical tool that lets us reason about rational players' strategic interac-

ADAM J.
O'DONNELL
Cloudmark

Apple has highlighted its supposed invulnerability to viruses in various marketing campaigns, and for good reason. According to F-Secure, more than 250,000 new pieces of Windows malware were identified in 2007, and, if current trends continue, we are set for another 500,000 to appear by the end of 2008 (www.f-secure.com/2007/2/index.html). Meanwhile, the number of total pieces of Mac malware is rumored to be less than .1% of that total, with the majority of that total appearing only last year.¹

Here, I introduce a model based on game theory for predicting if, and when, Mac malware will arise based on a reasonable number of measurable parameters. But first, let's review a few theories on how Macs have been able to avoid malware.

Qualitative arguments for avoiding malware

Many Apple users would like to believe the *unique-population* argument—Apple users are collectively more intelligent, computer savvy, and attractive than the average PC user, which somehow makes their machines immune to the masses' malware problems. (In the interest of full disclosure, I switched from

a SPARCStation to an Apple in 2003.) However, this rationale is somewhat absurd when you consider that the recent increase in Mac market share largely results from previous PC users leaving their supposedly less-savvy user population to switch to the new platform.

Another, more plausible (and testable) argument is the *secure-design* argument. It states that the OS X's software architecture is inherently more secure than the PC's, so mass exploitation of Macs is more difficult.

Yet another possibility is the *economic-motivation* argument, which views the lack of Mac malware as resulting from a lack of economic incentive. The argument states that as Macs gain market share, attackers will pay more attention to the population and produce more Mac malware.

Recent ad hoc experiments have shown the secure-design argument to be far less likely than the economic-motivation argument. Consider the Pwn2Own contest, which pays participants a sizable reward for compromising different client systems during the annual CanSecWest conference. Mac computers fell to exploits in QuickTime in 2007 and in Safari in 2008, rewarding Dino

tions. Initially developed in the late 19th century and heavily researched throughout the 1950s, game theory formed the intellectual underpinnings for many major economic and political decisions made in the latter half of the 20th century, including the nuclear deterrent doctrine known as Mutually Assured Destruction. It has also proven useful for examining far less grave issues, such as the evolution of communication in animals, the behavior of prisoners during interrogation, and the interaction between attackers and defenders in information security. Herbert Gintis provides an excellent introduction to the topic in *Game Theory Evolving* (Princeton University Press, 2000).

A game consists of three components:

- *players*—actors that can make decisions in the game;
- *strategies*—choices a player can make; and
- *payoffs*—the economic gain or loss experienced when the players commit to strategies.

Our game consists of two groups of players known as *attackers* and *users*. Users can choose between two strategies, *defend A systems* or *defend B systems*. Attackers can choose to either *attack A systems* or *attack B systems*.

Defining the payoffs requires a few additional terms. Let f be the market share of *A systems*. Without loss of generality, assume that *A systems* is the market leader, and f ranges from .5 to 1. Let p represent the probability that a class of systems can be successfully defended. Assume that the success rate of defending *A systems* and *B systems* is the same. Finally, let v represent the value of systems to the attacker. Because this article is concerned with predicting the emergence of client-side malware issues, I restrict our definition of our systems to be client systems of

relatively similar capabilities and performance. This restriction lets us assume that the same monetization model, namely the client, will be used to send spam and keylog passwords rather than host malware or a phishing site, and the same model will be applied to both classes of systems.

If attackers attack an undefended pool of *A systems*, the attackers are given a payoff of fv , or the value of all the systems in the pool. Users playing the *defend A* strategy can limit the payoff to $(1 - p)fv$. The same reasoning provides the payoffs for attackers targeting *B systems*, where the payoffs for attacking undefended and defended systems are $(1 - f)v$ and $(1 - p)(1 - f)v$, respectively. The strategies and payoffs are shown in normal form presentation in Table 1.

Each player enters into either a single strategy or a combination of strategies to maximize their gain or minimize their losses. For certain market shares and filter accuracies, it's possible for the *attack A* strategy to strictly dominate the *attack B* strategy,

in which an attacker will always gain more value from the worst-case payoff of the *attack A systems* strategy than he or she will gain from the best-case payoff of the *attack B systems* strategy. For our payoff structure, if the payoff of *attack A systems/defend A systems* is greater than *attack B systems/defend A systems*, then a rational attacker will always play *attack A systems*. Rearranging terms, we see that if the ratio of the dominant systems to minority systems $f/(1 - f)$ is greater than $1/(1 - p)$, or the accuracy of the protection methods in terms of number of attacks caught for every attack missed, then it doesn't pay to play the *attack B systems* strategy.

From the attacker's standpoint, it's not worth attacking the minority system if the majority system's market share is greater than the accuracy of its protection methods. From the defender's standpoint, the protection mechanisms' effectiveness at defending majority systems determines when minority systems will be attacked.

Predicting the minimum necessary criteria that defines when



Table 1. The normal form of an attack-defense game.*

ATTACK	DEFEND	
	A	B
A	$(1 - p)fv$	fv
B	$(1 - f)v$	$(1 - p)(1 - f)v$

* p is the probability that a class of systems can be successfully defended; f is the market share of A systems; v is client systems' value to the attacker.

attackers will move to a new platform becomes a matter of measuring market share and the effectiveness of security mechanisms for the majority platform.

The malware tipping point

It's challenging to collect accurate data on both the relative market share of competing technologies and the effectiveness of our myriad security solutions in the field. For the Mac malware issue, we can gather market share figures for PCs and Macs directly via Web statistics; however, statistics on antivirus accuracy are notoriously unreliable. Net Applications (<http://market.share.hitslink.com/report.aspx?qprid=8>) and W3Schools (www.w3schools.com/browsers/browsers_os.asp) put the quantity of Macs in the field as percentage of total install base between four and seven percent. Quoted usage of Safari, the default Mac browser, is even lower (www.w3schools.com/browsers/browsers_stats.asp). Testing groups, such as AV Comparatives, put the accuracy of antivirus engines across all known viruses in the 98 percent range.² AV engines fare far worse on new malware samples, with the accuracy rates at 80 percent for the best engines and a median accuracy rate below 50 percent across all products tested.³

If we assume that the accuracy on new malware stays at an optimistic 80 percent and that the malware writer's economy remains constant, then the Mac platform won't become appeal-

ing to attackers until it makes up 1/6th of the market for client systems. Additionally, if we assume 4 percent market share and stagnant growth of the client market, Apple would have to convince three times as many PC users as there are Mac users to switch over to their platform before we reach the tipping point.

Much of this analysis depends on the assumption that large portions of the malware space will remain stable. There are several factors outside of our model that could hasten or postpone the arrival of Mac malware. An economic recession could cause companies to hold onto PCs and older software longer, leaving the infected systems online for longer than previously intended. A more efficient botnet aftermarket, where bot herders and spammers can better price each other's services, would delay the need for capacity increases afforded by botting more systems. On the other hand, large-scale cleanup services focused at infected PCs would cut off the botter's current supplies and push malware authors into new frontiers. Competitive malware, in which one bot tries to push the other one out, would have a similar effect.

It doesn't appear that we're in any danger of large-scale malware hitting the Mac community anytime soon. Our analysis is restricted to determining when we will see the appearance of large-scale monetized malware. It doesn't mean that Macs are immune to attack, and we should expect news of Mac malware to pop up regu-

larly. Malware authors will continually test the market conditions and look for the right time to begin exploiting the new platform. We must also be mindful of targeted attacks, as the value of the data contained on an individual system to an attacker might far exceed the value of the machine as a platform for sending spam.

Although the masses might be able to forgo protection schemes, users who either are or have reason to be paranoid should still protect their system as if it were under attack. Antivirus software, firewalls, good backups, and intelligent data hygiene are all necessary, even if your platform of choice isn't under constant assault. □

References

1. A. Kingsley-Hughes, "F-Secure: More than 100-150 Malware Variants Targeting Macs," <http://blogs.zdnet.com/hardware/?p=1021>.
2. A. Clementi, "Anti-Virus Comparative No. 17: On-Demand Detection of Malicious Software," AV-Comparatives.org, Feb. 2008; www.av-comparatives.org/seiten/ergebnisse/report16.pdf.
3. A. Clementi, "Anti-Virus Comparative No. 16: Proactive/retrospective test," AV-Comparatives.org, Nov. 2007; www.av-comparatives.org/seiten/ergebnisse/report16.pdf.

Adam J. O'Donnell is the director of emerging technologies at Cloudmark, an anti-messaging abuse company located in San Francisco. His research interests include distributed system security, network measurement, and writing random magazine articles. Contact him at adam@cloudmark.com.

Interested in writing for this department? Please contact editor Dave Ahmad (drma@mac.com).



\$29

New Lower Subscription Price!

IEEE
SECURITY & PRIVACY

Subscribe to our
magazine today
for only \$29—
our lowest price ever!

You'll receive 6 issues of today's
leading-edge, peer-reviewed
software development information.

Ask us how
you can get this great deal on
IEEE Security & Privacy magazine!

S&P is the premier magazine
for security professionals.
Every issue is packed with
tutorials, best practices, and
expert commentary on:

- attack trends
- cybercrime
- security policies
- mobile and wireless issues
- digital rights management
- and much more.

Subscribe at www.computer.org/services/nonmem/spbnr