

Reengineering the Internet for Better Security

Manoj Parameswaran, Santa Clara University

Xia Zhao and Andrew B. Whinston, University of Texas at Austin

Fang Fang, California State University, San Marcos

The growing proliferation of malware is raising doubts about the Internet's future. Current security measures primarily target inbound traffic, but service providers have no incentive to stop attacks and spam at the source. A proposed certification scheme motivates providers to control outgoing traffic, efficiently increasing overall security while preserving the Internet's open, decentralized structure.

Following a decade of dramatic growth, the Internet has come to redefine nearly all human endeavors, opening new horizons in commerce, government, science, medicine, education, and leisure. By shrinking time and distance, it has also accelerated globalization, connecting people and businesses worldwide. The Web is emerging as the dominant interface for information exchange and service delivery, and e-mail is becoming the communication tool of choice. At the same time, however, there is a growing perception of the Internet as an insecure environment, and these concerns may prevent the Internet from realizing its seemingly limitless potential.¹

The recent proliferation of malware—including viruses, worms, Trojan horses, spam, phishing schemes, distributed denial-of-service (DDoS) attacks, spyware, and adware—has made the Internet a harrowing experience for many individuals and a severe headache for organizations. In a 2005 survey by the Pew Internet and American Life Project, 22 percent of respondents reported reducing their use of e-mail because of spam, while 67 percent labeled the act of being online “unpleasant and annoying.”² Business leaders, concerned about the rising cost of managing Web-based security risks as well as productivity losses due to employee surfing, have contemplated giving up on the Internet altogether.³

Ironically, what has made the Internet so successful—its open and decentralized structure—is also what sustains malicious online activity. Information on newly discovered vulnerabilities propagates quickly, and tools to launch ever more sophisticated attacks are readily accessible. The growing availability of inexpensive personal computers and broadband connectivity, coupled with average users' poor efforts to secure their operating systems, has further facilitated large-scale intrusions, including the remote hijacking of such systems to launch zombie attacks.⁴

Security researchers have made tremendous progress in keeping pace with Internet threats, but there are limits to what technology alone can accomplish. While it is possible to proactively prevent some attacks, most solutions are responses to exploits of unforeseen security flaws, and the current framework encourages neither the dynamic assessment of security risks nor the optimal deployment of prevention and response measures.

Controls implemented by Internet service providers (ISPs), which are interested in protecting their own network—and customer base—from external attacks, predominantly target inbound traffic. However, there is no similar economic incentive to control outbound traffic, as the potential damage is to other networks. This lack of clear lines of accountability derives from both the

decentralized nature of datagram routing in the Internet as well as its decentralized organizational structure.

What the Internet needs is an institutional structure that strongly motivates ISPs, network service providers, equipment vendors, and users themselves to control attacks at their origin as well as to maintain security on a dynamic basis. One way to accomplish this is to introduce a certification mechanism that induces service providers to voluntarily accept some degree of accountability, without interfering with the underlying decentralized protocols. Such a mechanism could propagate incentives through the network, ensuring that distributed participants coordinate their efforts to increase security as well as reduce congestion.

OUTBOUND TRAFFIC CONTROL

Security controls that focus only on inbound traffic tend to be limited in their effectiveness. Such traffic has already traversed multiple domains and wastefully consumed network resources. Many of the attacks that originate from a single domain rapidly branch out toward many targets, making it much more difficult to control them at destinations rather than sources.

To improve Internet security, it is essential that service providers control outbound as well as inbound traffic. Outbound traffic control stamps out attacks at the source and thus stops them from spreading, without subjecting the network to congestion. Outbound control is especially effective when done by ISPs, which can leverage the direct relationship with their customers to hold them accountable and take punitive action against violators.

The importance of outbound traffic control is underlined by the unanticipated consequences of improved inbound controls. For example, the deployment of more effective e-mail filters has caused spammers to increase the volume of junk mail they generate to increase their chances of sneaking in messages, leading to significant backbone congestion. Thus, while inbound controls might stop specific attacks that arrive at a network, they are of limited benefit to the Internet as a whole and might even cause some harm.

The only way to effectively and efficiently secure the Internet is to block malware as it leaves a network. Some ISPs and e-mail providers do curtail outgoing malicious content, but the practice must be universal or nearly universal to work. Given the Internet's decentralized structure, economic incentives to carry out outbound as well as inbound traffic control must be designed in such a way that individual service providers' self-interested decisions collectively benefit the general Internet community.

Role of the Certifying Authority

There is an intense international debate on whether to introduce regulation to the Internet and on what future governance models for the Internet should be. Given the degree to which malicious traffic is undermining productive use of the Internet, many feel that ICANN or other agencies must make organizational changes to provide a more secure environment for both users and businesses.

In our proposed scheme, the certifying authority's main function is to align service provider incentives in various ways—for example, by supporting verification of certified traffic by issuing public/private keys to certified providers, by enforcing compensation payments and suspending defaulting providers from membership, and by arbitrating settlement disputes.

The CA thus introduces a limited degree of regulation to the Internet, tempered by the fact that participants choose their actions voluntarily based on self-interest. The presence of the CA does not alter the Internet's decentralized design—no change in domain or routing configurations is implied; rather, it transforms the organizational structure by mediating economic transactions among service providers.

SERVICE PROVIDER CERTIFICATION

Toward that end, we propose a security mechanism for service providers based on the notion of a *certifying authority*. Membership in the scheme is voluntary: Providers that choose to join pay a subscription fee to the CA and are called *certified providers*, while those who opt out are known as *noncertified providers* (traffic originating from each type is labeled similarly). The CA requires certified providers to compensate

- remote providers that receive malicious traffic from the certified providers' users, and
- their own customers who receive malicious traffic, regardless of the source.

The CA holds any certified source provider accountable for an attack originating from its domain, regardless of whether a human customer or a zombie node initiated the attack. To minimize compensation payments, certified providers are motivated to filter all outgoing traffic.

Certification likewise helps guide inbound traffic control policies. Certified providers need not filter incoming certified traffic because they are assured of compensation; they must decide whether to reject all incoming noncertified traffic or filter it before entry. The choice would depend on the provider's optimization strategy—for example, the costs of incoming malware from a particular noncertified provider could exceed the benefits of communicating with the users that provider services.

Game Theory

Game theory is a mathematical tool used to analyze situations in which interacting agents seek to maximize benefits, which are determined by players' interdependent actions. Game theory is applied in numerous fields including economics, sociology, evolutionary biology, political science, artificial intelligence, and military theory to determine how agents will behave and what the ultimate outcome will be.

Game-theoretic analysis uses the concept of the Nash equilibrium to characterize the likely choice of strategies by agents. In a Nash equilibrium, each agent chooses the best response to strategies that other agents employ, implying that agents' expectations are mutually correct and that they act rationally based on these expectations. No agent can gain by unilaterally deviating from a Nash equilibrium of strategies.

The CA can be a nonprofit agency, like the Internet Corporation for Assigned Names and Numbers (www.icann.org), or a private for-profit organization; the Internet's current governance structure contains successful examples of both models. The "Role of the Certifying Authority" sidebar describes the regulatory implications of the CA's functions against the backdrop of ICANN's impending expansion.

Our proposed scheme encourages service providers to take up an offensive, rather than a defensive, posture against intrusions and spurious traffic that exploit the Internet's distributed structure. Tying penalties to the source of malicious activity shifts the responsibility for security to the originating provider. This is particularly important for attacks that rapidly escalate by replicating and targeting multiple destinations.

The certification mechanism currently applies only to service providers supplying Internet connectivity, whether to residential or enterprise customers. However, certification could be extended to include, for example, corporate networks. Businesses are particularly concerned about network security and thus have a strong incentive, as well as ample resources, to participate in such a scheme. Coupled with insurance protection against potential damage, certification would help guide corporate investments in cybersecurity as well as reduce anxiety about the Internet's viability as a commercial platform. A more general approach would be to extend certification to ISPs, e-mail service providers, and institutions that administer a particular domain; in a multitiered approach, ISPs can delegate accountability to clients with large networks.

INCENTIVE PROPAGATION

The effectiveness of inbound traffic control depends on the ability to correctly identify the source of incoming packets. For example, spoofing, which involves manipulating the source addresses of IP packets to conceal the

originating node, constitutes a significant percentage of malicious activity on the Internet today and is very hard to combat—even when a spoofed packet is detected, determining its actual origin is difficult.

Our proposed certification mechanism will reduce such attacks, which rely on backbone networks for transportation, by motivating access networks to implement their own outbound traffic controls. Because a local ISP that receives potentially spoofed traffic might not be able to collect compensation from the source provider, it has an incentive to persuade its access providers to suppress spoofed traffic—for example, by denying source routing and Internet Control Message Protocol redirects. The interior providers in turn will demand router and switch manufacturers to embed provisions for such controls. Current initiatives by

Cisco Systems and Juniper Networks indicate that router design is already incorporating security measures.

Although our framework does not explicitly prescribe penalties for end users, it could transform both the customer-provider relationship and the technology used to secure edge devices. Because service providers are accountable for malware in outgoing traffic, they might seek to better protect their customers' PCs as a more cost-effective alternative to taking action against individual violators. Given that average users are not sophisticated administrators, service providers could require security-certified devices as a condition for connectivity and might even sell or license the sale of such devices. Another option would be to create a tiered service structure that allocates different levels of security assurance or user responsibility based on customer preferences. A casual user may choose a hardened device with limited access and low risk; a tech-savvy user may accept higher risk for flexible access.

Starting with the ISPs, incentives can thus spread to the users, interior network providers, and hardware and software vendors. Eventually, they can propagate throughout the Internet, or at least most of it. The certification mechanism's scalability makes it more sustainable than the current framework, which is based on simply adopting better filtering technologies. Further, the combination of outbound traffic control in edge networks and deployment of more secure equipment configurations will significantly reduce network congestion in the backbone due to malicious traffic.

A GAME-THEORETIC EVALUATION

We used game-theoretic analysis, described in the "Game Theory" sidebar, to evaluate the viability of our proposed incentive mechanism and to determine its implications in terms of provider actions and collective security.

Following standard economic theory,⁵ we divided service providers into two types, those with a low-risk security profile (A) and those with a high-risk security profile (B), with the distribution of types but not the classification of individual providers being common knowledge. A high-risk profile indicates that the provider's customer base is more prone to sending out malicious traffic, either intentionally or by having less securely configured machines.

This classification is not meant to represent the real world of network providers—a continuum of types would more accurately reflect reality than two discrete types—but to simply determine whether, given reasonable assumptions, service providers would choose certification, whether it would benefit the Internet in terms of security, and whether the incentive structure is sustainable.

Assuming that customers seek to communicate with as many other Internet users as possible, service providers face various choices:

- whether to subscribe to certification,
- how to price their services to account for both the subscription and penalty dues,
- how to control inbound traffic from certified and noncertified providers, and
- how to control outbound traffic destined for certified and noncertified providers.

Each provider chooses strategies to maximize its payoff based on the distribution of types in the Internet. All providers of the same type adopt similar strategies.

Our analysis indicates that with a nonprofit certifying authority, all service providers will choose certification, leading to a net increase in system surplus. With a profit-maximizing CA, different equilibria may exist depending on the proportion of A-type providers in the network. When this proportion exceeds a certain threshold, only A-type providers get certified, leading to a *separating equilibrium*. When the proportion of A-type providers is below the threshold, all providers subscribe to the certification scheme, leading to a *pooling equilibrium*. These are Nash equilibria, wherein each participant chooses the best response to others' actions.

Separating equilibrium

In the case of a separating equilibrium, the CA has enough A-type subscribers that it can afford to exclude the B-type providers by setting high subscription fees. For B-type providers, these fees, combined with the expectation of costly penalties due to the number of A-type providers that potentially receive traffic from them, makes certification prohibitively expensive.

Certified providers screen all outbound traffic destined for other certified providers to minimize compensation payments, but they have no incentive to control outgo-

ing traffic to noncertified providers. Certified providers also need not filter inbound traffic from other certified providers, as certification insures them against any potential loss.

When dealing with inbound traffic from noncertified providers, certified providers can choose to either block such traffic altogether or filter it for malicious traffic. Game-theoretic analysis reveals that the blocking strategy is dominant, as the noncertified providers' propensity to send malicious traffic offsets the value derived from any legitimate incoming traffic from them. In this scenario, the certified Internet is effectively closed off to the noncertified Internet with respect to inbound traffic.

Noncertified providers have no incentive to control outbound traffic, but they do invest in inbound control to retain customers.

Pooling equilibrium

In the case of a pooling equilibrium, the CA sets the subscription fee low enough to induce everyone to join, the high number of subscriptions compensating for the low fee. B-type providers choose to be certified along with A-type providers because they can draw on a potentially larger pool of providers for compensation, their need for insurance is lower, and their customers stand to benefit from being able to communicate with customers of all types.

Every service provider controls outbound traffic, while none implements inbound controls as all sources are required to pay compensation. The primary focus of investment thus shifts from ingress to egress, rendering control measures far more effective. Because customers can potentially communicate with more users than in the separating outcome, they are more willing to pay for the service.

B-type providers typically generate more attacks than A-type providers and thus are more likely to invest in outbound controls. However, even A-type providers will expend more effort filtering outgoing traffic than in the case of a separating equilibrium to minimize compensation payments to other certified providers.

INTERNET SEGREGATION

Critics might argue that the benefits of service provider certification carry a high cost: Because certified networks accept inbound traffic only from other certified providers, the certified Internet may reject much legitimate traffic from the noncertified Internet. However, the disutility of such *false positives* lies in the uncertainty they entail, not in the rejection of a legitimate service. For example, if your friend has a noncertified service provider, you know that you will not receive any message she attempts to send to you. In contrast, with currently deployed antispam measures, you might never learn that your filter terminated a legitimate message sent to you.

The segregation of service providers resulting from certification would naturally extend to users. While some users would clearly value the ability to send traffic to anyone, provided it is legitimate, others would be willing to forsake greater reach for more control over outbound content. The mere fact that a user subscribes to a noncertified provider does not imply intent to commit malicious activity, but it does suggest a higher tolerance for such activity.

However, segregation is not absolute. Those with a certified provider account can easily establish an alternate, noncertified Internet account. Any individual can thus maintain multiple digital identities to derive the value of membership in each community. The imposition of incentives does not cut spammers and hackers off from Internet access altogether; it merely limits their malicious activities.

Interestingly, some researchers have concluded that physically segregating the Internet would increase security—according to one study, for example, separating the IP address space into servers and clients would effectively curb DDoS attacks.⁶ However, such a strategy would require rebuilding the Internet from scratch, which is clearly infeasible. In contrast, our voluntary participation-based incentive mechanism achieves the benefits of segregation through economic means, robustly supporting communications without interfering with the Internet's basic structure.

Failure to combat the growing scourge of malware could lead to real fragmentation of the Internet: Academic communities could spin off to form their own private networks—as some are already doing—and enterprises could rely on private, value-added IP networks, while underground and edge networks proliferate on the side.

Service provider certification improves overall security without undermining the fundamental design philosophy of the Internet as an open, decentralized network. Choices by both users and service providers are voluntary, and digital identities remain connected. By imposing a virtual rather than a physical partition within the Internet, our proposed incentive mechanism encourages the formation of communities of interest, which is critical to both information sharing and productive activity. ■

References

1. D. Ropeik and G. Gray, *Risk: A Practical Guide for Deciding What Is Really Safe and What Is Really Dangerous in the World Around You*, Houghton Mifflin, 2002.
2. D. Fallows, "CAN-SPAM a Year Later," data memo, Pew Internet & American Life Project, Apr. 2005; www.pewinternet.org/pdfs/PIP_Spam_Ap05.pdf.

3. D. Talbot, "The Internet Is Broken," *Technology Rev.*, Dec. 2005/Jan. 2006, pp. 62-69.
4. Y. Huang, X. Geng, and A.B. Whinston, "Defeating DDoS Attacks by Fixing the Incentive Chain," *ACM Trans. Internet Technology*, Feb. 2005; <http://crec.mcombs.utexas.edu/works/articles/DDoS2005.pdf>.
5. M. Mas-Colell, M.D. Whinston, and J.R. Green, *Microeconomics Theory*, Oxford Univ. Press, 1995.
6. A. Greenhalgh, M. Handley, and F. Huici, "Using Routing and Tunneling to Combat DoS Attacks," *Proc. Usenix Workshop Steps to Reducing Unwanted Traffic on the Internet*, Usenix Assoc., 2005, pp. 1-7; www.cs.ucl.ac.uk/staff/M.Handley/papers/sruti.pdf.

Manoj Parameswaran is an assistant professor of operations and management information systems in the Leavey School of Business at Santa Clara University. His research interests include Internet governance and security, social computing, content distribution, and peer-to-peer (P2P) networks. Parameswaran received a PhD in information systems from the University of Texas at Austin. He is a member of the Association for Information Systems (AIS). Contact him at mparameswaran@scu.edu.

Xia Zhao is a PhD candidate in information systems at the University of Texas at Austin. Her research interests include information security, IT compliance, and electronic commerce. Zhao received an MS in control theory and control engineering from Tsinghua University, China. She is a member of the AIS and the Decision Sciences Institute. Contact her at xia.zhao@phd.mcombs.utexas.edu.

Fang Fang is an assistant professor in the College of Business Administration at California State University, San Marcos. Her research interests include prediction markets, network finance, and P2P network pricing. Fang received a PhD in information systems from the University of Texas at Austin. She is a member of the AIS. Contact her at fangfang@csusm.edu.

Andrew B. Whinston is a professor of information systems, computer science, and economics at the University of Texas at Austin, where he also directs the Center for Research in Electronic Commerce at the Graduate School of Business. His research interests include electronic commerce, knowledge management, online auctions, and financial markets. Whinston received a PhD in management from Carnegie Mellon University. He is a member of the ACM and the AIS. Contact him at abw@uts.cc.utexas.edu.