Communications of the ACM

Opinion

Computing Applications

Technical opinion

The Right Place at the Right Time

Examining the use of network location in authentication and abuse prevention.

By Geoffrey Goodell and Paul Syverson

Posted May 1 2007

The use of network location to draw conclusions about users has become quite commonplace on today's Internet. Numerous institutional subscription services use the IP address as the exclusive means of identifying users. An IP address is sometimes also used as a criterion for fraud detection and abuse prevention, for example, flagging discrepancies between geographic location associated with an IP address and that associated with a mailing address for credit card billing. For abuse prevention, many Web sites that allow public contributions (wikis, blogs, chat rooms, and so forth) simply block the IP addresses from which chronic abuse emerges. Using IP address to categorize or identify users seems like a reasonable approach in general, but there are some important caveats as well. In this column, we briefly examine the long-term architectural dangers as well as short-term policy risks as we assess the costs and benefits.

Back to Top

Practical Justification

There is little doubt that some ISPs are more vigilant than others in curtailing spam and abuse. Many network administrators have accepted the idea that it is easier or more effective to fight spam by restricting the flow of mail instead of addressing the security vulnerabilities that make spam feasible (compromised hosts send a large proportion of unsolicited mail) or the market forces that make spam desirable (financial incentives foster the development of spam distribution software and methods). At least for the present, knowing the ISP from which a connection originates certainly provides some statistically meaningful information about whether the user responsible for the connection is likely to engage in antisocial behavior. Similarly, teaching Bayesian spam filters about network location may enhance their effectiveness in reducing spam.

Providers of online subscription services can be reasonably assured that most connections from the IP address of institutions that provide limited access to their systems are from authorized users of those systems. Also, the overhead of setting up a mechanism that uses this information is far simpler than most alternatives.

In recent years, a market has emerged for so-called "geolocation" services, which provide a mapping from network location to physical, geographic locations. Service providers collect data from Internet service providers and resell it to geolocation customers in the form of datasets or permission to execute queries on their databases. Geolocation products have been developed for several uses, including fraud resolution, spam mitigation, targeted advertising, and digital rights management.

An important premise of geolocation is that there is much to learn about individual Internet users from how they are connected to the network, and such information can be used as a basis for implicitly categorizing them by risk level or market segment. Geolocation affords advertisers the possibility of offering products and services specific to particular localities. Content providers wishing to disallow people from certain countries or regions from accessing certain content may, to a significant extent though not completely, use network location information to achieve this result.

Similarly, it is reasonable that credit charges from locations that are far from the home of a credit card holder and that are known to be hotbeds of credit fraud merit close scrutiny. But this can have problems too: a large percentage of credit card use is associated with travel. Also, mail-order catalogs and, still easier to distribute, Web pages mean that even small local shops with only telephonic credit transaction clearing can have many remote customers. The credit industry typically errs on the side of giving users easy access to its services rather than denying undesired access when these goals conflict. The content industry typically adopts the opposite approach. For this reason, IP address is only one of many factors in authorization and identification for credit approval and fraud detection rather than an absolute or sole discriminator.

> **To the extent that it is useful, the ability to use network location as an indicator of identity is a technical shortcoming of the current Internet that can be overcome.**

So far it seems that the primary incentive for those who use network-layer information for application-layer decisions is to provide an expedient means of authorization, striking an acceptable balance between easy access for desired usage and adequate deterrence against undesired usage. In short, it works. To date, IP addresses have been a resource difficult enough to obtain or spoof that they have fulfilled this role in authorization and fraud detection. More proper authentication, however, would require user certificates, a public key infrastructure, or other mechanisms that have proved difficult to set up for large, relatively open systems and have not seen widespread user adoption where they do exist. And, as long as end-user systems not under institutional control remain as vulnerable as they now are to root-level intrusion, end-to-end authentication could also be an illusory approach to security. However, that is a much broader and separate problem than the one we are considering.

Back to Top

Immediate Side Effects

The ability to block abuser IP addresses is a powerful but ill-suited tool for some of the problems to which it has been applied: in a few cases, individual sites have blocked access from IP addresses in a broad geographic area. Two well-known examples are the blocking by a major-party presidential campaign of its Web site from IP addresses outside the U.S. just prior to the 2004 U.S. presidential election and the blocking of 2004 Olympics coverage from IP addresses inside the U.S. In 2002, Pennsylvania ISPs blocked access to 1.6 million innocuous Web sites in an attempt to satisfy a state mandate intended to curtail child pornography [1]. More recently, the major U.S. ISP Verizon was the subject of lawsuits when it began blocking all email from Europe and other continents by default as a spam deterrent [5]; since that time, Verizon has agreed to a settlement.

In general, binding the identity of users to how they are connected to the Internet is not only impossible, but also undesirable. The vast majority of Internet users have dynamically assigned addresses, and while they may use the same address from lease to lease, they also may not. Proxies, workarounds, dynamic addresses, mobility, system vulnerabilities, and other complications make network location useful as a heuristic at best. To the extent that it is useful, the ability to use network location as an indicator of identity is a technical shortcoming of the current Internet that can be overcome.

One of the main drawbacks to using network location for authorization is that legitimate users cannot access a service when away from their base institutions. It is possible for users to set up tunneling such that their access appears to be from an IP address within the permitted range, but this may be onerous, technically difficult, or simply not possible as a matter of policy. A major expense associated with deploying and maintaining a virtual private network (VPN) infrastructure derives from the need for individuals and businesses to access Internet resources that rely upon network location information to differentiate between valid and invalid users.

Furthermore, abusers can deploy and use proxies in unblocked locations fairly easily. Individual proxies themselves can be blocked, but with multitudes of newly compromised hosts emerging daily, the effectiveness of that approach is limited. Networks designed to protect honest users from traffic analysis, such as Tor [3], can be blocked because they explicitly provide a means of doing so, but abusers can take advantage of million-node botnets with no easily discernible pattern of IP address source. The result would seem to block the honest users from protecting themselves while leaving the abusers unblocked and more difficult to find.

Back to Top

Long-Term Security Risks

We have noted some immediate practical problems, but solutions that avoid these problems raise additional security concerns of their own. IP tunneling simply to allow use of institutional subscriptions when it is not otherwise needed is an extreme solution that may open the possibility of other intrusions to the institution. University librarians have long recognized the problem that authorization by IP address poses for remote users attempting to access institutional services. For example, consider what happens when an outsider uses university-based research projects, such as research on proxies or filtering, to access resources that are supposed to be limited to university access. This has been one of the prime motivators for development and increasing adoption of systems such as Shibboleth (see http://shibboleth.Internet2.edu/), which provides single sign-on and user-controlled credential management independent of IP address. This can help permit remote access without resorting to IP tunneling.

If legitimate users of credit systems have incentive via easier authorization of their transactions to route from an IP address associated with their home location, they reveal via routing information their interactions with merchants and financial institutions, not just to those principals but to observers as well. This may leave them more vulnerable to identity theft, spear phishing, and the like. And, as actual large-scale systematic fraudsters become aware of the use of authentication by IP address, they are provided with specific incentive to spoof authorized or trusted (low fraud) locations, or worse, to break into systems in or near those locations. This approach thus has the potential of greater vulnerability and risk for the legitimate user with a false sense of protection against the actual adversaries.

Back to Top

The Role of Network Access Providers

The fact that network location provides an effective way to assign blame for malicious activity raises questions about the extent to which network access providers ought to be responsible for the comportment of the systems for which their networks serve as attachment points to the Internet.

Regulators have substantial interest in supporting the principle of enforcement within the network, since the network could potentially provide convenient points of control for execution of policy. Furthermore, incentives for major telecommunications carriers generally encourage a movement away from uniform, open access and toward vertically-integrated "silos" in which carriers determine the set of resources that customers may access [2, 9]. So, there exist strong industry and regulatory forces to empower network operators at the expense of network neutrality and end-to-end connectivity. A recent ITU proposal advocates expanding the role of ISPs to require that they ensure traffic traversing their wires adheres to certain normative requirements [6].

Back to Top

Function Creep and Expedience

Making use of the routing infrastructure itself to protect participants from each other is an arms race that sacrifices as collateral damage the neutrality characteristics of the Internet that provide its principal advantages over alternative interconnection paradigms.

The consistent response that proposers of such methods offer to proponents of end-to-end services is that it may be too late to salvage such principles, and that compromise is in order. Indeed, network-layer techniques have shown promise as expedient short-term remedies to exigent security threats, and as a result, governments and regulators have called upon ISPs to implement technical solutions within the network [10]. Vint Cerf recently argued that it does not make sense to use the network to compensate for operating systems that protect themselves inadequately. In particular, having the network perform application-level security checks on traffic to attempt to authenticate end users or to look for viruses and malware makes it less efficient at moving data around. "It’s really hard to have a network-level thing do this stuff, which means you have to assemble the packets into something bigger and thus violate all the protocols," Cerf says. "That takes a heck of a lot of resources" [8]. In other words, if we start requiring the network to perform tasks other than routing, then we undermine the ability of the network to do its most essential job.

Another problem with such function creep is that it can become entrenched. Once a technique such as authentication by IP address is widely established, if legitimate technical reasons to substantially change how addressing and routing is done should arise, they may be more difficult to implement and establish because they are hamstrung with the use of IP address as authenticator or antifraud mechanism, even if that was originally only introduced as an expedient.

Back to Top

Separating Identification from Routing

IP addresses were introduced to allow the routing of IP packets. As we have already seen, if IP addresses are used for other purposes, and if identification and authorization become tied up in routing, then the purpose for which they were designed is undermined: both legitimate users and attackers end up using IP addresses not because of routing, but to appear as authorized users. Onion routing was introduced 10 years ago as an infrastructure that "separates identification from routing" [7]. "Parties are free to (and usually should) identify themselves within a message. But the use of a public network should not automatically give away the identities and location of the communicating parties" [4]. Anonymity from one’s communication partner is not the primary motivation for onion routing; users may simply need to protect their points of attachment from attackers, whether personal (such as stalkers or identity thieves) or enterprise (such as corporate competitors gathering intelligence). In each case, the security benefits of separating routing from identification are substantial, even if the challenges it poses to the security models of some services are similarly great.

How will increased user mobility, increased use of anonymity networks for security by honest users, and other factors interplay with use of IP address for authentication and authorization? Intuitively, it seems that these two technologies are headed for a clash.

Back to Top

Discussion

The extent to which network location information is used in authentication and fraud and abuse detection will have a substantial impact on the Internet environment for the next several years. Adversaries may or may not adapt to these techniques before the techniques become entrenched in the architecture of critical services. Although, if history is a guide, they will adapt at some point, and all the faster if IP address location technology increases their incentives to do so. If we are to avoid arriving, therefore, at an entrenched burden with no ultimate benefit, we must understand the technology that we are using to do the job.

In this column, we have set out some of the unanticipated ways in which relevant technologies interact. One response to understanding this is through governance and policy, but our focus herein is the technology itself. The complexity, brittleness, and overhead involved in the deployment of solutions that use network-layer address for authorization may stifle innovation in the future. But it's not all bad news. Systems like Shibboleth provide a technological path that can continue to lead institutions away from authentication by IP address.

Similarly, since network location is only one of many factors going into fraud detection scores used by the credit industry, the technical framework already exists to allow an abandonment of network location as a factor as it diminishes in significance. This need not mean the end of geolocation services either; for example, people will still want to know about nearby restaurants and services. Geolocation services will just need to be based on information other than IP address if they are to continue serving a useful purpose. If the existing security and functionality problems arising from IP address-based abuse deterrence do not lead to its abandonment, then the incentives it provides to network attackers ultimately should.

Back to Top

Back to Top

References

1. Bradner, S. Simple solutions are often wrong. *Network World Weekly 20* (Sept. 2004).

2. Bradner, S. Misunderstanding the fundamentals of telecom reform. *Network World Weekly, 26* (Sept. 2005).

3. Dingledine, R., Mathewson, N., and Syverson, P. Tor: The second-generation onion router. In *Proceedings of the Seventh USENIX Security Symposium*, August 2004.

4. Goldschlag, D.M., Reed, M.G., and Syverson, P.F. Hiding routing information. In R. Anderson, Ed., *Proceedings of Information Hiding: First International Workshop*, Springer-Verlag, LNCS 1174, May 1996.

5. Leyden, J. Verizon faces lawsuit over email blocking. *The Register* (Jan. 21, 2005).

6. Palfrey, J.G. Stemming the international tide of spam. *Trends in Telecommunication Reform 2006* (Mar. 2006), 111125.

7. Reed, M.G., Syverson, P.F., and Goldschlag, D.M. Proxies for anonymous routing. In *Proceedings of the 12th Annual Computer Security Applications Conference*, (Dec. 1996), 95104.

8. Talbot, D. The Internet is broken. *Technology Review 108*, 11 (Dec. 2005).

9. Van Schewick, B. Towards an economic framework for network neutrality regulation. In *Proceedings of the Telecommunications Policy Research Conference* (Sept. 2005).

10. Zittrain, J. The generative Internet. *Harvard Law Review 119* (May 2006).

About the Authors

Geoffrey Goodell (goodell@eecs.harvard.edu) is a fixed-income strategist for a securities firm in New York, NY.

Paul Syverson (www.syverson.org) is a mathematician in the Center for High Assurance Computer Systems at the Naval Research Laboratory in Washington, DC.

Join the Discussion (0)