A Whitepaper

Fixing the Internet:
A Security Solution

by Roger A. Grimes
InfoWorld Magazine security columnist
roger_grimes@infoworld.com

June 28, 2011

Version 2.1

Table of Contents

Abstract

The Problem

The Internet has a significant amount of malicious activities and security risk. Cybercrime is high reward coupled with low risk, and continuing to increase in scope and sophistication. Hackavists are breaking into more high profile company web sites and databases than ever before. And at the same time, society is ever increasing its reliance on the Internet for everyday activities, including mission-critical applications that should not be utilized over untrusted networks like the Internet. The increasing (or even sustained) level of maliciousness is colliding with society's increasing dependence on the Internet for legitimate needs, creating an unacceptable risk. If left unaddressed, it could lead to a tipping point event; or at the very least is already causing a significant opportunity loss (e.g. money spent on inadequate defenses, slower computing performance, and people unwilling to conduct meaningful transactions over the Internet, etc.). Few currently available or advertised solutions (known to me) appear ready to change that fact over the next 5 to 10 years.

The Solution

Although many in society think we must learn to live with the current level of maliciousness, there are open standard solutions that can significantly minimize Internet maliciousness in the mid- and long-term. It will take a global, coordinated, community-based effort, along with accepting an increased managed (both centralized and private) control presence. This sort of maturation, in other infrastructure platforms (e.g. water, fuel, electricity, land rights, etc.) has occurred throughout history, turning nomadic peoples and uncivilized societies into collaborative, productive centers where all citizens (participating or not) benefit.

Fixing the Internet's security problems will require a two-fold approach: a world-wide, global "dream" team of security experts working in concert to solve the systemic problems; and a global Internet security infrastructure solution that addresses and provides security protections holistically.

In order to be broadly accepted and used, the solution(s) to fix the Internet must:

- Use Open Standards
- Be Vendor and Platform neutral
- Use Open and Transparent Decision Processes
- Be Voluntary, Opt-In
- Be Performance Neutral, or Nearly So
- Integrate with Legacy Systems
- Not Be Disruptive to Users and Services

As difficult and complex as this seems at first, it can be accomplished. Contrary to established, knowledgeable critics, this goal is readily achievable, today, using already existing open standards.

This paper will present the underlying security problems with the Internet, provide a global framework upon which to build stronger, longer lasting Internet security solutions, and ends by presenting two possible solutions that could fit within that framework. Readers are invited to critique, support, or reject.

[Note: The ideas and recommendations contained in this paper are solely the responsibility of Roger A. Grimes (e:roger_grimes@infoworld.com). No vendor or sponsor has been involved in the creation, editing, or approval of this whitepaper.]

Revision History

Date	Author	Version	Change reference
5/1/08	Roger A. Grimes	0.9	Initial draft, reviewed by key people
5/8/08	Roger A. Grimes	1.0	First draft released publically to the Internet and introduced in InfoWorld security column
5/12/08	Roger A. Grimes	1.1	Added FAQ section to answer most commonly asked questions
2/19/09	Roger A. Grimes	2.0	1. Updated various text portions to be more inclusive 2. Integrated WS-* open standards into text 3. Added additional authentication scenario
6/28/2011	Roger A. Grimes	2.1	Minor updates and clarifications, re-released.

Author Bio

Roger A. Grimes
e: roger_grimes@infoworld.com

- 24-year computer security researcher (started fighting malware on Apple IIs, Commodores, Amigas, and DOS in 1987)
- Author of 4 books, co-author of 1 book, and contributing author of 3 books, and over 300 national magazine articles on computer security
- InfoWorld magazine security columnist\blogger since August 2005
- Principal Security Architect since February 2007 for Microsoft Information Security & Risk Management ACE Team
- Expertise in securing Microsoft Windows, but also runs OpenBSD and various Linux distros at home
- Runs 8 honeypots tracking malware and hacker behavior
- Former Ultimate Hacking Expert instructor for Foundstone
- 9 years of experience as a penetration tester
    - Have compromised every company, site, and device I've been hired to pen test, in every instance in 1-hour or less (with one exception that took 3 hours)
- Frequently invited guest speaker or keynote at national security conferences
- Has participated in designing large scale and national identity metasystems

Introduction

Fact #1 – The Internet is full of malicious behavior which is not expected to decrease significantly over the next 5 to 10 years

The Internet is over two decades old[1], and unfortunately, rife with malware and malicious activities. Spam currently compromises over 70% of all email traffic[2]. Some experts estimate that malicious activities compromise 2-6% of all Internet traffic today[3]. Phishing attacks are becoming more targeted, and successfully compromising both casual home computer users and Fortune 100 executives[4] alike. Hundreds of thousands to millions of malicious bots control vulnerable computers[5] - conducting identity theft, adware redirection, distributed denial of service (DDoS) attacks, privacy invasions, corporate and government espionage, extortion, child pornography, and other malicious objectives. Malware is getting increasingly sophisticated (e.g. fast-fluxing[6], server-side polymorphism[7], generic one-offs that will never exist again) and propagated through legitimate (compromised) web sites[8]. People's bank accounts and stock portfolios have been emptied. Over a quarter of all U.S. adults had their financial identity information compromised in one year alone[9]. Hackivist groups seemingly attack high-profile targets with impunity, bringing down defenders and public media sites they judge as adversaries (e.g. LulzSec, Anonymous, etc.).

It is highly likely that millions of dollars are being stolen on the Internet every day[10], not to mention credit histories ruined, and legitimate operations and people's lives disrupted. We almost never catch the criminals. Of the major crimeware gangs (e.g. Russian Business Network[11], Rockphish[12], etc.) we have never identified, much less caught and prosecuted a single member. Internet crime is high yield and low risk. Current anti-malware defenses are being challenged like never before to accurately respond, and it is highly unlikely that most of the traditional solutions will significantly reduce malware over the mid- and long-term.

Fact #2 – Society Is becoming increasingly reliant on the Internet for basic and mission-critical services

At the same time, more and more of society's activities are moving online. Internet-connected mobile devices and phones and social media sites (like Facebook and Twitter) are becoming a way of life for a significant part of the world's population. The movement of proprietary software systems into the Internet cloud is hastening the world's dependency on the Internet.

Even older, traditional computer users who don't willing embrace the latest Internet fad are being forced to do more online. What starts out as a public service convenience, turns into the primary way business is conducted, and leads to the only way business can be conducted. These include traditional commercial transactions (e.g. airplane tickets, concert tickets, paying bills, requesting services, etc.), as well newly evolving mission-critical applications that were never intended for an unsafe transport mediums. These include online healthcare records, software-as-a-service applications, university emergency alert systems, remote workers, online banking, television, and Voice-over-IP telephony.

Many mission-critical applications that the general public would never imagine were hosted on the Internet are. For example, the SQL Slammer[13] worm in 2003 compromised tens of thousands of unpatched SQL servers in under 10 minutes[14]. Hundreds of banks, including many of the world's largest banks were compromised and shutdown during that outbreak.

Since then, the incredibly appealing low price point of using the Internet as a VPN transportation pipe versus other alternatives has attracted more mission-critical applications to the Internet, not less. Many large-scale, city and regional supporting infrastructures are dependent on the Internet, and are being compromised over the Internet. Even the highest-risk, mission-critical applications (e.g. 911 response systems, public utilities, police systems, nuclear management facilities, etc.) that we are told aren't connected to the Internet, can easily be affected by Internet performance issues because they share strategic "choke points" along transmission lines and within telecom facilities.

Malware outbreaks affecting non-online public and private services is nothing new. Several past malware outbreaks (e.g. Iloveyou worm[15], Blaster worm[16], etc.) in the early 2000's affected integrated resources, causing telephone, pager, and cellphone disruptions, network news delays, and even the late delivery of basic goods and services. We are so inter-connected now with the Internet, that a single, widespread online attack will always impact the physical world. This fact isn't new. We have been lulled into a sense of complacency because no big Internet attacks have happened over the last few years. What is disturbing is the increased reliance on the Internet and what a new widespread disturbance would mean today, or in 5 years, or 10 years?

Fact #3 – Current Computer Defenses Are Inadequate

Current anti-malware defenses (e.g. antivirus, anti-spam, anti-spyware, firewalls, etc.) have proven mostly inadequate over the last twenty years and are ever decreasing in effectiveness. Whatever computer defenses vendors have come up with have been easily circumvented by faster reacting malicious hackers. Unfortunately, even though the current, traditional defenses are imperfect, end-users and business entities are forced to accept them (and their expense) because there is nothing else out there currently is better. Many vendors are trying to develop stronger, longer-lasting, harder-to-defeat defenses, but they are many years away from production release or require global adoption to work.

This brings up many important questions, including:

Why do we keep creating the same types of traditional point defenses against malicious computer activities when they so obviously don't work (with over two decades of largely imperfect anti-malware history as proof)?

How many people will not conduct legitimate business over the Internet (i.e. opportunity cost) today because of realistic, appropriate fear?

How much legitimate business does not happen over the Internet (i.e. opportunity cost) today because of realistic, appropriate fear?

How can we possibly be looking to put our personal medical records online[17] with the Internet's stability and security so much in question?

We are looking to improve the overall conditions of the world using the power of the Internet and yet we are inviting these technologically new users into an inherently unsafe place.

It is these colliding realities, rampant maliciousness and increasing reliance on the Internet, which makes the improved security of the Internet of vital importance.

A Solution Framework

Solving the Internet's security problems will require a global, community effort.

No Single Vendor Solution Is the Answer

No single vendor solution can make the Internet more secure, for the following reasons:

- The substantial security problems of the Internet are not a "product" or "protocol" problem. The underlying problems are systemic and affect every vendor, every product, every protocol; and which if fixed, would make the other point solutions more successful (or unneeded).
- Most vendor security responses are acute, point-specific, in nature, not focusing on the underlying strategic problem; resulting in inefficient "whack-a-mole" defense solutions. When one hole is closed, the hacker attacks another weak link.
- Security defenses evolve slower than malicious attacker techniques.
- Every network packet is exposed to the same levels of scrutiny (or lack of scrutiny) and given the same speed of delivery regardless of the demonstrated historical trust of the originating gateway (e.g. a packet from a trusted partner is treated identically to a packet from an untrusted source). For examples, traffic from the Russian Business Network IP space travels around the Internet and to your Internet egress point at the same speed as a long-time, trusted business partner or loved one.
- Most global security events (e.g. large bot DDoS attack, phishing and spam floods) are only noticed by a small set of selected vendors. If the data was shared faster, globally with everyone, the benefit would be greater.
- No globally accepted security initiative addresses the systemic problems.
- No global Internet servers or services address security broadly.
- No Internet global body has a charter focused solely on malicious prevention, although we have dozens covering response.
- Few currently proposed solutions (that I am aware of, with one exception covered below) will make a significant decrease in malicious attacks over the next 5 to 10 years.
- End-user education is highly overvalued (many end-users are not technologically sophisticated enough to recognize malicious events). We need to develop solutions that minimize asking the end-user to make trust decisions.
- There is no accountability for poor security or poor coding (e.g. some of the poorest security performers are gaining market share, and origination sites with consistently poor trust records can access destination resources at the same quality of service as proven trustworthy providers).

Real Solutions

Fixing the Internet's security problems will require a two-fold approach:

- A world-wide, global "dream" team of tactical security experts working in concert to design systems and protocols to solve the systemic problems
- A global Internet security infrastructure service (likened to DNS) that addresses and provides holistic security protections

Global Security Dream Team

The Internet is full of very bright, sometimes popular and accomplished, sometimes relatively known, security experts with good solutions to the Internet security problems. Unfortunately, their good ideas languish inside of their respective employers (due to competing self-interests), on Internet discussion lists only known to the lists participants, or in research papers left largely unread.

We have many times in the past, when faced with a seemingly insurmountable problem, gathered together the world's best minds in their respective disciplines and solved the unsolvable. Examples abound: clean water, vaccines, computers, nuclear energy, outer space programs, and ending wars. This point in time requires that we build another team dedicated to significantly improving Internet security.

Selected top vendors (open source and commercial) and independent security experts should be brought together for a period of 6 months to 2 years to debate the problems of the Internet and recommend strategic and tactical solutions. An open and transparent consortium should be created to facilitate these expert meetings, and participants should agree to work toward common, agreed upon objectives.

Note: Many existing national commissions already exist and have had the ear of high-level politicians, even the United States President. Unfortunately, all of the prior committees have been heavy on executive and strategic thinkers, but missing senior security tactical thinkers and technicians. While strategic and political committees are absolutely necessary, there has yet to be one designated to design and deliver actual solutions.

Team Makeup and Responsibilities

There should be a different, independent team created for each critical core component, which naturally seems to lie somewhat along the OSI model's definitions (e.g. Physical, Logical, Network, Session, Application, etc.). To that idea we should add other shared necessary components, such as Cryptography, Identity, ISP, IANA, Legal, Global Considerations, Privacy, Open source, End-User, etc.

There should be a larger, more strategic Executive committee team that helps coordinate and integrate the various lower component teams and provides strategic direction to each component committee. There should be a team leader (with only 1 vote and chosen by a majority of participants (each also with 1 vote)) of each component committee. Each component committee will be responsible for developing the tactical ideas to be passed along to the technical participants under each component. The technical participants will be responsible for coming up with technical solutions and standards to meet the tactical assignments. The technical (and end-user, public, and shared committees) will also be tasked with providing technical guidance to the higher committees (i.e. can the tactic be realistically implemented).

The figure below shows the basic consortium design along with the component committees.

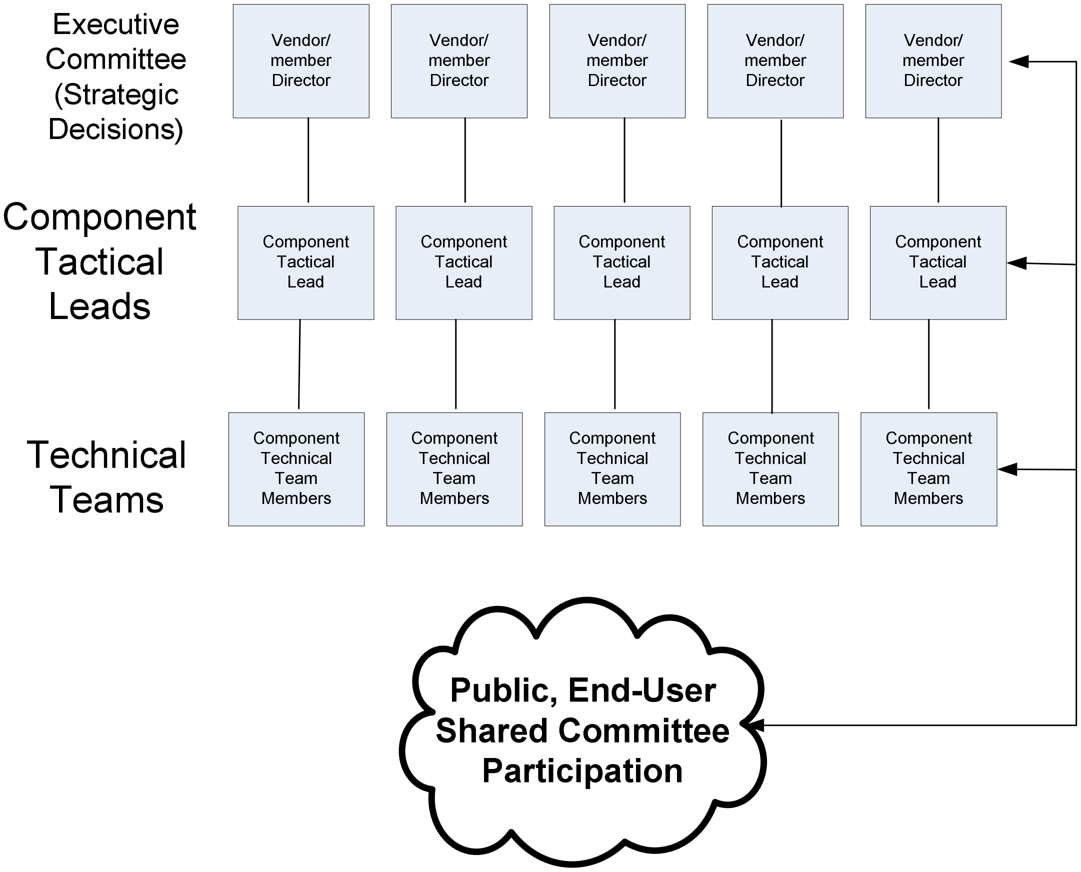

How Big of a Team?

Although any number I pick now is arbitrary, 10-20 participating members on the Executive committee and 10-20 members on each component committee seems a realistic starting number. Invite 5-10 vendor leaders into each component committee, and another 5-10 independent field experts. Initial component members (no more than half of the total members) could be chosen by the Executive committee, and additional members voted on by the original members by majority rule.

Example Vendor Participant Members

Participating vendors would have to dedicate and fund multiple original committee members, including:

- Senior Management (responsible for selecting Execute Director representative, non-voting)
- Execute Director (voting member, responsible for coordinating member's response and vote)
- Assistant to Director (logistics, minutes recording, etc.)
- Technical Lead for each tactical component the vendor is involved with
- Senior Technical Staff under each technical lead (although who participates here can vary according to need)

Thereafter, the community-based consortium would require ongoing, permanent (but revolving) members to address standard updates, either to address improvements, additional coverage, or to respond to vulnerabilities.

The Hardest Part

The hardest part of solving the Internet's security problems is not generating the technical security solutions. If you can solve the hardest part, the technical solutions will come easily. Getting vendors and independent experts to dedicate 6 months to 2 years of their life to a single, societal goal is among, if not the hardest part, of solving the Internet's security problems. Natural sustainability (usually revenue or earnings) dictates that members work on their own self interests to maximize revenue. How do we get vendors and individuals to give up potential, immediately recognizable revenue gains to concentrate on the greater good, which ultimately benefits themselves and the commons? I'm not sure. I'm hoping that if enough end-user interest is generated by this idea, governments will call upon citizens to do their civic duty and vendors will volunteer to participate as much is possible for a reasonable period of time. Or perhaps, a grant of some type could be awarded to offset the revenue reduction to benefit the greater good. This is a tough issue to solve.

Transparent and Open Submissions

It is important that every single consortium word, decision, and result would be posted on the Internet to be as transparent as possible. In order to get a world-wide community solution we need the community's trust. It must be supported by open source and commercial concerns.

One Member, One Vote, Public Participation

Each participating member would be given one, equal vote on all proposals, and additional members could only be added by majority vote. In order for this idea to be successful we must guarantee to participants (many of whom will be hesitant otherwise) that this is not vendor-specific dominated initiative. Multiple public and private participants can be present and engage in debate, but only one vote is allowed in a particular component committee.

The public will be invited to participate at multiple points and their comments and submissions reviewed (by a sub- or full-committee as each component committee deems appropriate); although in order for any idea or issue to be voted upon it must be brought into full committee by a voting member. All proposals must receive an up or down vote, and majority rules.

Why the 6 months to 2 Years Timeline?

I believe that time is of the essence, not because we don't have time necessarily, but we need to use time as a tool to minimize members debating details to death and getting lost in the weeds, and forgetting the overall goal. I propose the following time schedule:

- 2-3 months to organize the effort
- 6 months for the teams to meet and discuss possible solutions
- 6 months for public review and discussion
- 6 months for technical review and decisions, and the final vote on document 1.0
- 6 months to document decisions and release new security proposals to all vendors

Additionally, one of the primary problems why we have not solved the Internet's security problems is the relative speed at which malicious hackers move as compared to the security problem solvers. By proving that we can move quickly within a naturally bureaucratic system, it will provide some measurable disincentive to future malware writers. Plus we can use the lessons learned to move even more rapidly in the future, when responding to new challenges.

If we created a global consortium to concentrate on resolving the Internet's security problems, two years from now we would have new global, community supported Internet security standards, which could be implemented by participating vendors and individuals. At the end of two years, vendors and individuals could then take the time they need to implement the standards in their own way (or reject them and not participate directly). Legacy devices and software must be able co-exist and function with the newer devices. If done appropriately, no one is deprived of legitimate service, except the malicious hackers.

Other Solution Ideas:

- What the committees can't agree on will be tabled or split (for just that issue) so we can get the overall, strategic and tactical goals met. Let's vote on what we agree on.
- Solutions must be opt-in, with more "carrot" and little "stick". People choosing not to opt-in are only disadvantaged by not directly participating in better security.
- Solution must address all computer platforms (PCs, PDAs, cell phones, media players, TVs, etc.)
- Any response to hacker vulnerabilities against the new standards must be rapid. We want to demoralize the current and potential hackers, and show that the defenders can respond as quickly as the bad guys.
- There are human, process, and education elements to consider.
- We need strong global participation for global acceptance.
- Optional idea: Funding for the long-term community consortium members can be collected through some minor (voluntary) monthly or manual minimal fee collected at the Internet's egress points.

Challenges/Questions

- The normal issues associated with global, strategic direction without explicit authority.
- How to create enough self-interest to motivate major vendors and other needed participants to meet?
- How to be quickly responsive to changing malware tactics…must be built into process.
- Balkanization of committees, objectives, or protocols (that's why we will table and split when needed), but majority rules; let good ideas emerge, even if differing.
- Should this be an entirely new committee or rolled into some existing body (e.g. IANA, IETF, CERT, Trusted Computing Group, etc.)?

Global Internet Security Infrastructure Service

The Internet's major security problems cannot be solved by a single vendor or a vendor-specific solution. Whatever the solutions are coming from the above mentioned Internet security consortium, the outcomes will be global and require global, coordinate participation (in most cases). The Internet lacks any service or infrastructure dedicated to coordinating/advertising/publishing security services (again, think DNS). Accordingly, I propose building a global Internet infrastructure service to provide coordination, advertising, and publication of the various global security initiatives.

This idea is similar to an imagined cross between the global DNS infrastructure , a web services' Universal Description Discovery and Integration, UDDI[19] service, and the Trusted Computing Group's new IF-MAP standard, applied globally. The diagram below re-summarizes the concept.

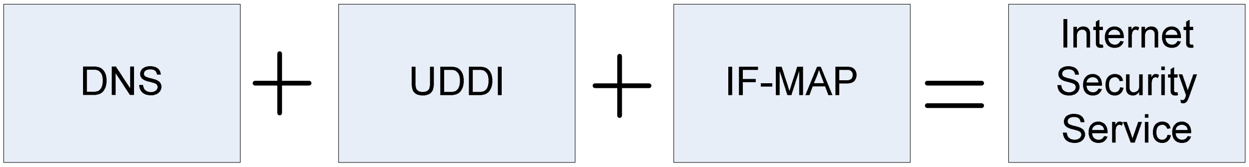

The new global Internet security infrastructure service should DNS-like in that there would be fault-tolerant, distributed "root" servers dedicated to directing querying clients to the appropriate security service server(s). It would be UDDI-like in that each participating global, sub-root server would to serve up IP addresses to the corresponding needed security services (and to advertise and publish such services). It would be IF-MAP-like in that the existing sub-root servers would allow participating members to report and respond in a global, holistic, multi-service manner.

If you are not familiar with IF-MAP, in a nutshell, the new Trusted Computing Group's (www.trustedcomputingroup.org) IF-MAP standard (https://www.trustedcomputinggroup.org/specs/TNC/IFMAP_FAQ_april_28.pdf) allows participating devices to report security events and receive notifications from other security devices to be able to respond in a coordinated fashion.

For example, if a firewall notes an unauthorized outbound stream that it recognizes as a bot spam stream, the firewall can contact the IF-MAP service, which can then contact a policy server that contacts another service that shunts the offending device off the network. The Internet security service would be similar to IF-MAP in that it would allow the coordination (i.e. reporting, advertising, direction, and response) of multiple disparate services, but be global in scale. Currently, the IF-MAP standard focuses on coordination within a single control domain. The Internet security service would be available for global coordination and direction, and should be integrated with private IF-MAP devices. The global Internet security service would have to be resilient, fault-tolerant, and cryptographically sound.

The following diagram gives an example of what the infrastructure might look like:

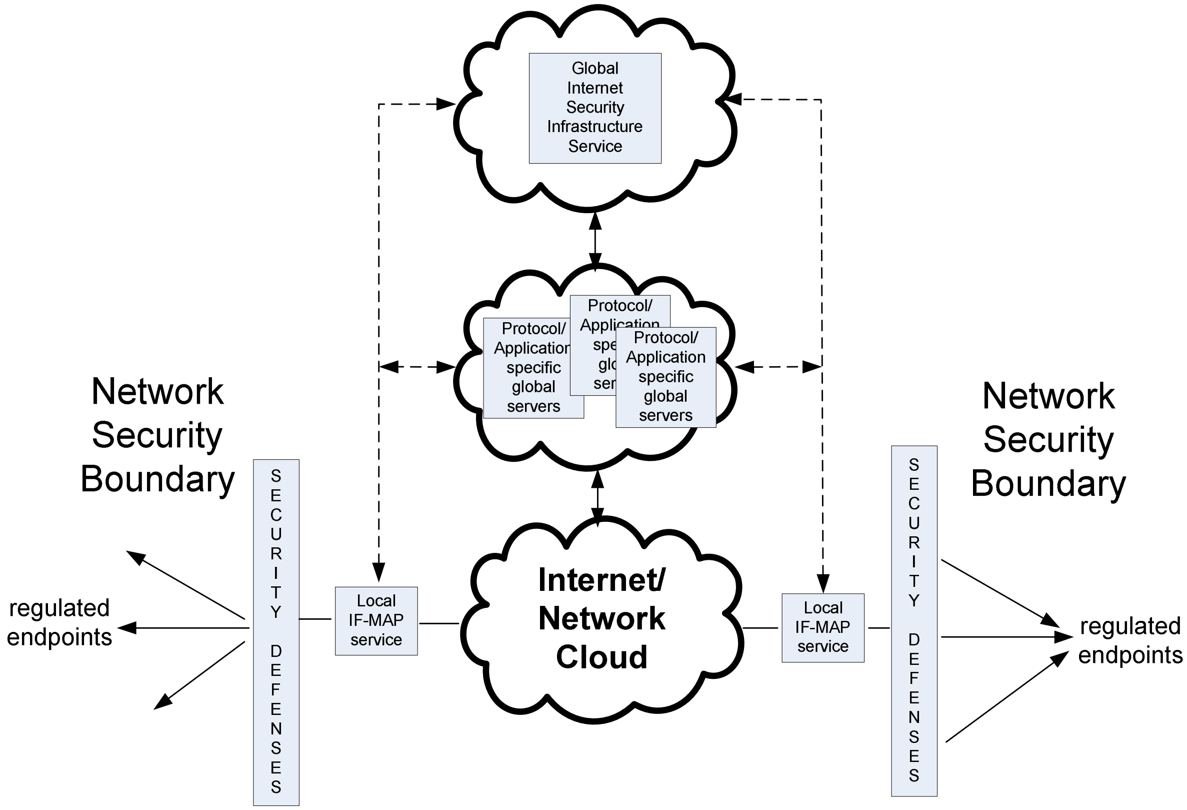
Global Internet Security Infrastructure Service
Protocol/ Application specific global servers
Protocol/ Application spe glo ser
Protocol/ Application specific global servers
Network Security Boundary
SECURITY DEFENSES
regulated endpoints
Local IF-MAP service
Internet/ Network Cloud
Local IF-MAP service
SECURITY DEFENSES
Network Security Boundary
regulated endpoints

The local IF-MAP services could take advantage of the global Internet security service, and be better able to respond (and report) threats. This would allow local security domains to respond quicker to threats noted by other partners, and be able to report local threats to other partners for their benefit. This sort of cooperative coordination has so far only realized in commercial, private, and more narrowly-focused public projects.

For example, several large anti-malware vendors (e.g. Symantec, Microsoft, McAfee, etc.) are able to capture and respond to large global threats because they have millions of participating nodes collecting and reporting statistics. Several open source and commercial anti-malware black lists have been around and used publicly for over a decade, albeit limited to a few uses (e.g. anti-spam, anti-phishing, etc.). There are several private groups, often led by anti-malware researchers, which collect and disseminate information to its members. Other groups, like SANS (www.dshield.org) collect limited information from participating members, and share the collected information publicly. These are all laudable goals, but suffer from limited membership or focus. A global Internet security service could collect information on a broader scope and its wider information used by more people. If global threat information was publicly communicated instantly, each participating entity, and the Internet, in general, would greatly benefit. Malicious hackers depend on the lack of global coordination to be successful. Let's take that strategic advantage way.

Example Scenarios Benefitting from a Global Internet Security Infrastructure Service

- Your network or web server comes under attack by a DDoS attack. Your local IF-MAP security device could connect to a root Internet security server and get directed to one or more services to allow an efficient response and defense to the attack. Your network could get subscribed on-the-fly to an anti-DDoS service, fire up additional availability resources on new IP spaces, or lead all the other participating networks into shunting off the offending bot-infected computers.

- Your company participates in a global whitelist/blacklist of IP addresses. Your company's whitelist/blacklist servers/service could contact the global root servers to get instantaneous updates of the Russian Business Networks' changing IP address space.
- Your anti-spam device or anti-phishing filter can learn instantly when a massive new spam or phishing attack occurs instead of waiting for a vendor update or allowing only the already existing global email servicers to learn about the attack.
- Supposed a MySQL-based Slammer type, zero-day, worm gets launched that can be successful against all existing, contactable MySQL servers on the Internet. Your firewall could be notified of the zero day attack and shut down the port until a better remedy is provided.

Regarding the last example. The original MS-SQL Slammer worm went off in the early morning weekend hours (in the United States). The majority of compromised servers occurred in under 10 minutes. Not only did the attack start and essentially end in under 10 minutes, but it was six to eight hours before the vast majority of the waking up Internet users (in the U.S.) learned of the attack, and began to respond. It seems unusually risky that we do not have devices ready to automatically respond to instantaneous global threats and are still relying on humans (which on average are asleep one-third of the time) to implement reactive solutions. It would be better if we had widespread, global early warning systems with rule-triggered IF-MAP devices to handle the initial response.

The following diagram shows some example coordinated services and proposed connection points.

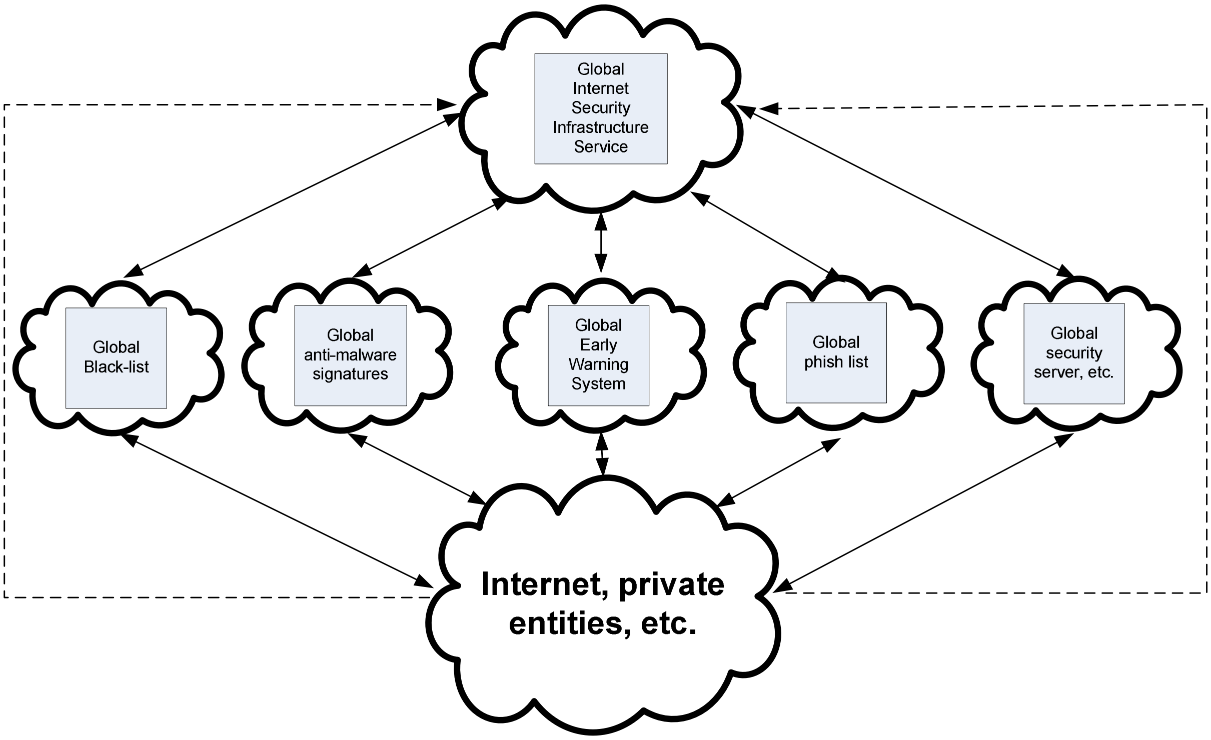
Global Internet Security Infrastructure Service
Global Black-list
Global anti-malware signatures
Global Early Warning System
Global phish list
Global security server, etc.
Internet, private entities, etc.

We know we need global, coordinated security early warning and responses, but we do not have a global security infrastructure to support this need. This type of solution would be in the end-users self-interests because it provides better, holistic solutions, and provide lower cost and better performance as Internet maliciousness decreases. It would be in the vendors self-interests because they get to develop a new stream of products and defense responses they haven’t even considered, yielding new customers and better solutions in an otherwise staid space.

Possible Solution #1– Replace Default Anonymity with Pervasive Identity and Integrity

by Roger A. Grimes, roger_grimes@infoworld.com

Abstract

The major underlying Internet security issue that is preventing a significant reduction in malicious behavior is the pervasiveness of default anonymity on the Internet. Because we can't identify malicious hackers with a high degree of confidence we cannot identify or hold them accountable. Internet crime is high-yield and low risk. If the Internet's model of default anonymity was replaced with default identity and integrity, the amount of maliciousness would significantly decrease.

I propose that every participating Internet component, hardware and software, be modified to provide increased identity and integrity assurance. Participating devices and users would provide improved levels of trust and be treated appropriately. All participating network traffic would be cryptographically tagged with a "trust level", which could be evaluated and acted upon accordingly. Each participating security domain would be responsible for assuring the trust and labeling of its egress traffic and responsible for acting upon tagged ingress traffic (and be held accountable for its attestations).

A security domain gateway device (called a " trust gateway") would perform the necessary trust labeling and evaluation. Every component (e.g. hardware, OS, network devices and pathway, identity, etc.) would end up being evaluated and assigned a numerical trust rating. Levels of trust, and how to obtain them, would be determined by a consortium of computer security experts, and published in an open, transparent manner.

Increased assurance levels would result in higher trust level ratings. For example, a user logging on with non-complex, short password would result in a lower trust rating than a user using two-factor authentication. Identity of participating nodes and users must be assured, but does not necessarily mean that each unique identity translates to a specific entity or user (i.e. user's real name).

All participating traffic would be encrypted and authenticated from origination to destination trust gateway end-points. Participating nodes and network traffic, demonstrating increased reliance and assurance, would undergo less inspection and given an increased quality of service. Nodes wishing not to participate would still be accepted and evaluated exactly as they are today, albeit with a lesser quality of service as compared to participating nodes. The solution would be vendor independent, transparent, open, voluntary opt-in, performance neutral, with least service and end-user interruption as possible, and driven by user and vendor self-interests.

Note1: The ideas and recommendations contained in this paper are solely the responsibility of Roger A. Grimes (e:roger_grimes@infoworld.com). No vendor or sponsor has been involved in the creation, editing, or approval of this whitepaper.

Note2: I accept that this particular solution will not make everyone happy. I'm bound to have critics that strongly disagree with it. However, it is my hope that this part of the whitepaper (*Possible Solution#1 – Replace Default Anonymity with Pervasive Identity and Integrity*) stands alone and is evaluated separately from the solution framework provided in the first part of the paper above.

What's Wrong With The Internet?

To understand how to improve Internet security you have to ask why things are as bad as they are. Most people when asked this question respond with problems (and solutions) that are pain point-specific (e.g. anti-virus technologies aren't accurate enough, we have to patch too often, software is always insecure, the end-user is the problem, etc.), but don't always focus on the strategic, underlying issues.

Security issues and solutions can be broken down into the CIA triad components: Confidentiality, Integrity, and Availability. All are important. But if you ask which one, if solved, would significantly decrease Internet maliciousness? It is without a doubt, Integrity. If we could confirm that the email is from who it says it is, we would end all spam and phishing. If we could confirm that the offered security patch is really from the vendor who says it is, we would not install malware. If we could identify the origination of released malware, we could track the hackers. If we could identify malicious hackers, we could arrest them. In fact, I can't think of a single significant, remaining Internet problem that isn't an identity or integrity issue.

Most of the Internet's infrastructure and its components run with default anonymity making it difficult to hold the majority of malicious participants accountable. Why do malicious hackers hack? Because they can do it with near impunity. Without greatly improved identity, integrity, and accountability, there can be no significant reduction in malicious Internet activity.

Solution

Build into the Internet pervasive, reliable, trustworthy identification and integrity into participating components and transactions, from source to destination. This will require a world-wide, community-based approach and the strengthening of every core component (called "trust components") along the OSI model, including:

- Hardware
- OS Boot Process and Loading
- Device and User Identity
- Network Stack and Protocols
- Applications
- Network Transmission Devices and Packets
- Communication Sessions

And it must be accomplished vendor independent, voluntary, opt-in, performance neutral, and with least service and end-user interruption as possible. An accepted solution must integrate legacy components while providing (voluntary) compelling reasons for consumers, vendors, and service providers to adopt solution-compatible components. I propose doing this by making each Internet egress network responsible and accountable for the security and trust of the endpoints in their network.

This applies to corporate environments, as well as, ISPs being responsible for the security of their end-user clients (to a variable degree). Each egress network access point would be known as a "trust network", and the management and technical teams responsible and accountable for implementing improved security trust mechanisms (e.g. egress filtering, two-factor authentication, anti-malware, secure coding practices, etc.).

A world-wide community consortium of computer security experts would transparently decide what levels of trust are assigned to the various trust components and how various trust networks earn increasing levels of trust. Egress points with poorly demonstrated levels of security will be given a low trust rating, and that rating known to all participants (e.g. world-wide trust rating list). This should encourage trust networks to improve their security to be rated higher, and at the same time hold accountable questionable networks (e.g. Russian Business Network's malicious IP space).

These global trust ratings would be sharable and available to each communicating trust network. Each receiving trust network can decide how to treat incoming traffic based on the originator's trust rating; and even provide custom trust ratings to trusted private trading partners (regardless of the packet's tagged trust). Traffic with higher ratings of trust should be inspected less and be delivered faster to end-points.

Trust Gateways

Each trust gateway should implement a trust gateway device (which can be a separate component or integrated into other egress/ingress point devices and software (e.g. ISA server). The trust gateway device is responsible for tagging egress traffic with a community decided upon trust rating, and appropriately handling (and handing off) incoming traffic based upon the trust rating with which it is marked.

Community-Based Trust Rating Server

A participating Trust Network's trust will be registered on a community-based Trust Rating Server. Trust gateways can periodically query the Trust Rating Server and download the trust ratings for various trust networks. This way we can update trust ratings and track when the bad guy networks move, and communicate that move to all participants. All network ratings, good and bad, will be readily available for inspection. We will have to build a process for rating and updating, efficiently. If a trust rating cannot be updated quickly and with integrity, the whole system breaks down. At first this may seem like an alien idea, but we have many such community-based servers, but none focusing on holistic trust.

How Trust Is Determined?

Every defined trust component (e.g. hardware, boot, OS, identity, software, network, etc.) contributes to the overall trust rating of the packet leaving or traversing a trust gateway device. Each trust component receives its own trust rating, and culmination of all trust component ratings leads to an overall packet trust rating. Each participating network transmission device is also assigned a trust rating, and the transmission path of each network packet from source to destination adds an additional network pathway trust rating. Thus packets sent along trusted network pathways are given higher levels of trust than those traversing lesser secured routers and devices.

For example, one-factor identity gets a lower rating as compared to two-factor, and so on. There will be a network device rating. Network routers without source routing enabled, fully patched, with strong passwords, without known vulnerable scripts, etc. will be given a strong rating.

The diagram below shows a logical representation of two packets with trust ratings, showing their individual component trust rankings and the overall packet trust ranking.

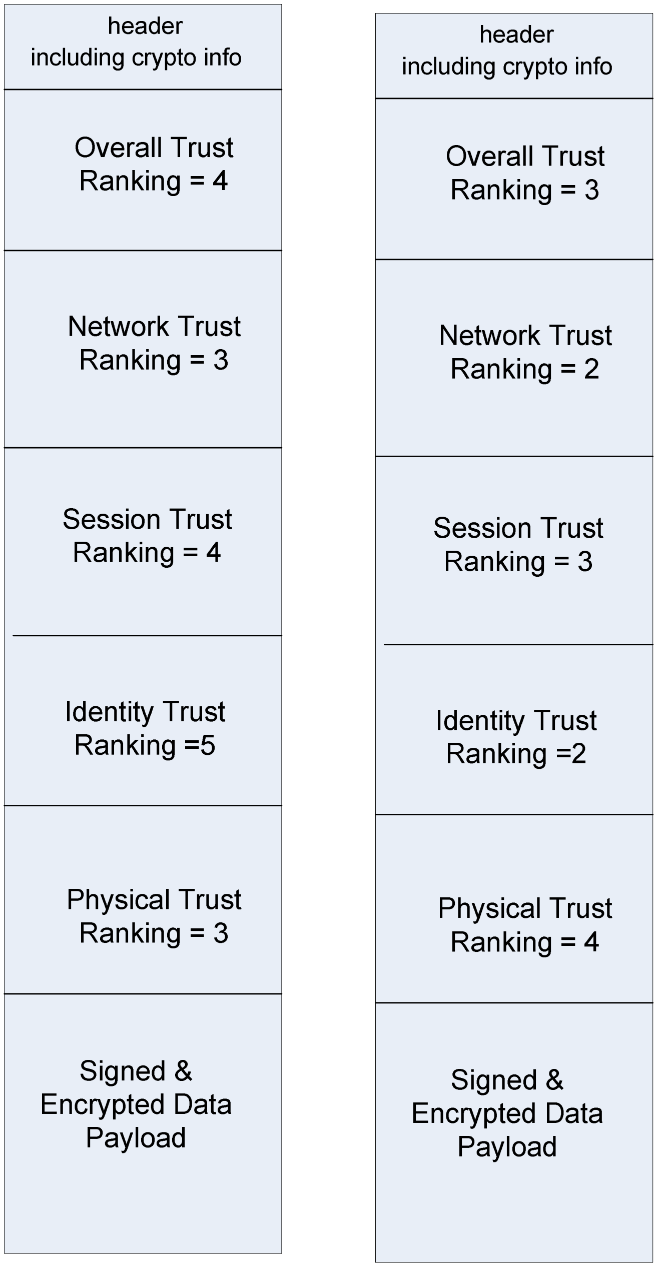

How a component is ranked will be determined by community-based decisions, and documented in a transparent, public-accessible document. It will be a common-criteria sort of document, but based on real, implemented security best practices. Most other common-criteria sort of guides are flawed because they end up being paper exercises and don't translate to real improvements in security. This document will be immediately usable and help all users and networks to improve security.

All component ratings end up generating the packet's overall trust rating, and both component and overall trust ratings are built into the network protocol for inspection by intermediate and final destination trust devices.

Ingress trust devices can treat network traffic differently depending on individual component trust rankings or rely solely on the packet's cumulative rating. This gives flexibility to ingress points that require different security policies (e.g. an online bank requires higher identity ranking while a network peering partner requires higher levels of network trust). Legacy devices will ignore the trust component, but pass along the trust components unmodified.

Trust ratings will be tagged into the traffic, and securely protected against unauthorized modification. Ingress trust gateways can rely upon the packet's attestation level and/or query a global community trust rating server to confirm the incoming security domain's historical trust ranking. If a particular security domain ends up being recognized as a poor trust decision, then the global trust rating servers can deliver that message to the ingress gateway device.

My idea is summarized in the diagram below.

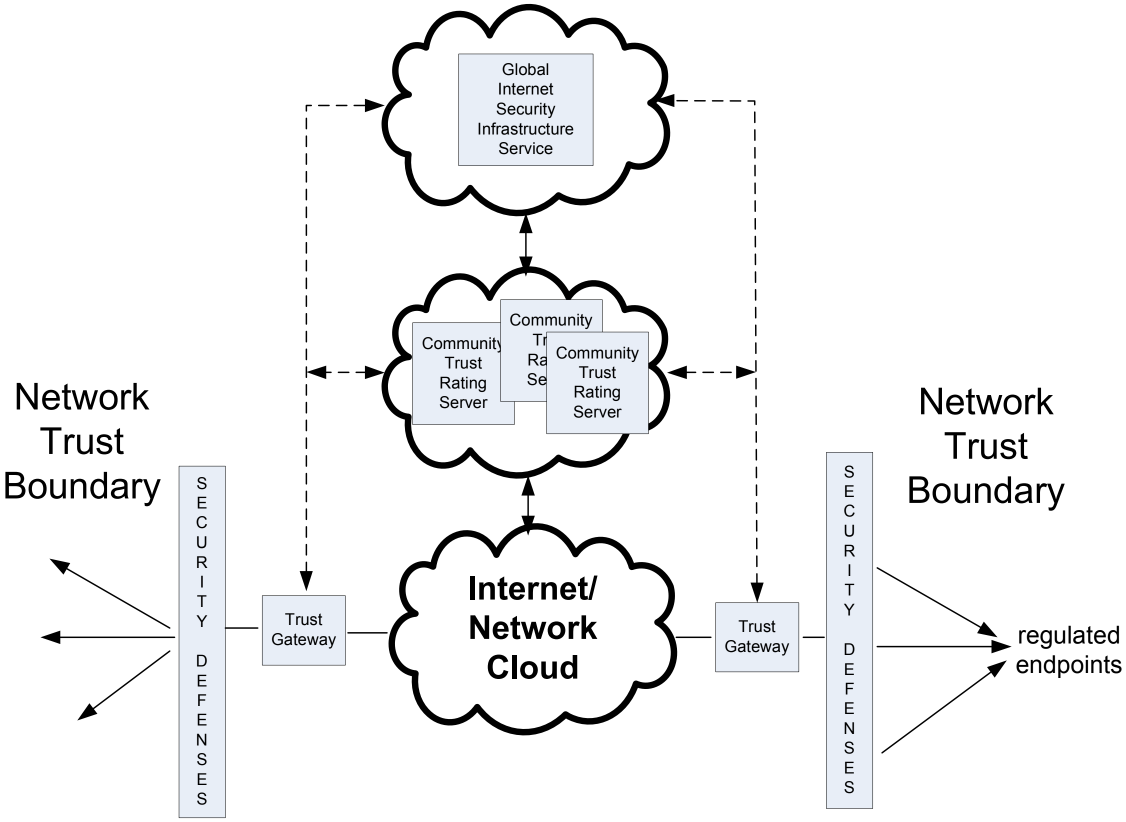

Thus, a roving malware network, constantly changing IP addresses could be tracked and identified by the global trust servers. No longer could malware writers hide behind fast-fluxing IP and DNS domain name changes. Another example, could be a previously highly trusted network or web site becomes infiltrated by malware. During the active attack, the compromised network or host could be assigned a lower trust rating, and that lower trust rating communicated to all participating parties. Once the malware was cleaned up and the network or host running clean again, its trust rating could be improved, maybe slowly at first. But certainly after a set period of time, it could regain its original trust rating, or actually improve it beyond the original if newer, more secure practices were used. Currently, there is no way for the Internet community to be aware that a particular, popular host or network is compromised. With more and more legitimate sites being used to host malware, we need some sort of warning system.

Integrity and Identity Without Personal Identification

Privacy proponents, of which I am one, might decline this solution on its face because of the forced identification to participate. It is true that in order for this solution to work, that the destination network must be able to rely on the identity of the originator. But this doesn't necessarily mean that the destination network knows the originator's true identity. There are mechanisms and companies dedicated to the idea of identity without personal identification. The idea is that I can prove my real identity to a trusted third party, who then gives me a global token that I can use on behalf of myself…or perhaps multiple tokens, unique to each use, so I can't be tracked or identified by anyone. This is known as *pseudo-anonymity*. Thus, Internet participants can choose to be truly anonymous, pseudo-anonymous, or authenticated along various levels of increasing trust assurance. True privacy advocates can choose not to be identified (i.e. remain anonymous), but it doesn't have to be a binary decision.

The destination network/host can choose whether to require the originator's real identity, or just a reliable proxy identity, or to accept truly anonymous connections, and treat received traffic accordingly. Originators may choose whether or not to participate with a destination network depending on the destination network's identity requirement. For example, my destination network may choose to drop traffic without a real person's identity attached to it, or just treat it differently than personally identified traffic. The idea is that right now all networks must accepted poorly authenticated traffic as the same level as more trusted traffic. This new solution would give both origination and destination networks a choice to handle trusted and untrusted traffic differently.

Cryptographically Sound

This solution requires that open cryptographic standards be employed to ensure that all participating transactions are secure, confidential, and have integrity. The participating, chained components in the trust pathway must cryptographically verify the next participating component (much like is done in the Trusted Platform Module chip today). Device and user identity must be cryptographically verified and attested. Each trust component and its trust ranking must be cryptographically verified and attested. Network traffic must be tagged in a cryptographically sound manner that detects unauthorized modification. Lastly, information sent is cryptographically protected (encrypted and signed) by default, and can only be read or verified by the destination network. Default encryption and signing of data is not required for this solution to work, but is encouraged to prevent unauthorized viewing and manipulation.

How To Satisfy the Remaining Critics and Non-Participants

This solution takes into account that initially some large portion of critics and end-point nodes will choose not to participate. This solution is an opt-in solution. If it provides a compelling reason to join, we can expect some of the critics to join as success is demonstrated. End-nodes not participating are not harmed beyond their current service levels and expectations, other than being given a lower quality of service rating as compared to more trusted traffic. If this solution significantly decreases malicious traffic on the Internet (rated at 2-6% of overall Internet traffic), even non-participants should benefit from increased performance, or at worst be performance neutral.

Possible Solution #2 – Global Identity Metasystem

by Roger A. Grimes, roger_grimes@infoworld.com

This solution proposes creating global infrastructure layers to provide one or more identity/authentication pairs to end-users from one or more Authentication Providers (APs) for use by content and service providers (let's call them Content Providers to simplify). Essentially, an End-User would request one or more identity/authentication pair from one or more Authentication Providers. Authentication Providers could provide password services, biometric identities, two-factor authentication tokens, smart cards, or whatever identity/authentication pair they want to offer- each with a defined trust assurance level.

Trust Assurance Levels (TALs) would be defined globally, published, and available for anyone to see. All participating Authentication Providers would have to build their identity/authentication pairs to meet a certain level of assurance as predefined in the TAL table. Example TAL table might look something like this:

TAL Value	Assurance Level	Authentication Type
0	None	Unknown connection
1	None	True Anonymous Connection
100	Low Assurance	Simple password, made up identity
500	Medium Assurance	Pseudo-anonymous identity using InfoCard, complex password and registered, verified identity
1000	Medium Assurance	Smart card, two-factor, identity verified by local proxy
65000	High Assurance	Three factor biometric identity, verified in person by certified representative, background investigation, etc.

End-Users would be free to obtain identity/authentication pairs from any participating Authentication Provider, and could have multiple identity/authentication pairs, and submit different ones to different Content Providers.

Authentication Providers would be audited by a central authority and given their own trust assurance level. Authentication Providers could not assign identity/authentication pairs above their own trust assurance level. Abuses by an Authentication Provider might result in censure or decrease in their trust assurance level.

Each participating Content Provider would re-code their site or applications to work with participating Authentication Providers. A Content Provider would designate what minimum level of assurance is needed for an end-user to connect to their content or service. It could be done at a domain or site level, down to as granular as a specific object. For example, a payroll processing company would allow anyone, anonymous or not, to connect and download public documents. However, to see individual paycheck results might require medium assurance. To withdraw payroll money might require high assurance.

When an End-User connects to the Content Provider's site, the Content Provider prompts the user for their identity/authentication pair, along with the minimum level of assurance needed. The End-User's computer would then securely supply the appropriate identity/authentication pair to the Content Provider's web site/application to begin authentication. In most cases, the Content Provider would not ever see the End-User's authentication token, just enough to identify the End-User's identity, the type of authentication token used, and the originating Authentication Provider's identity.

The Content Provider could pass along the submitted identity/authentication pair to an Authentication Provider for authentication. The Authentication Provider would approve or deny the identity/authentication pair, which the Content Provider could handle accordingly.

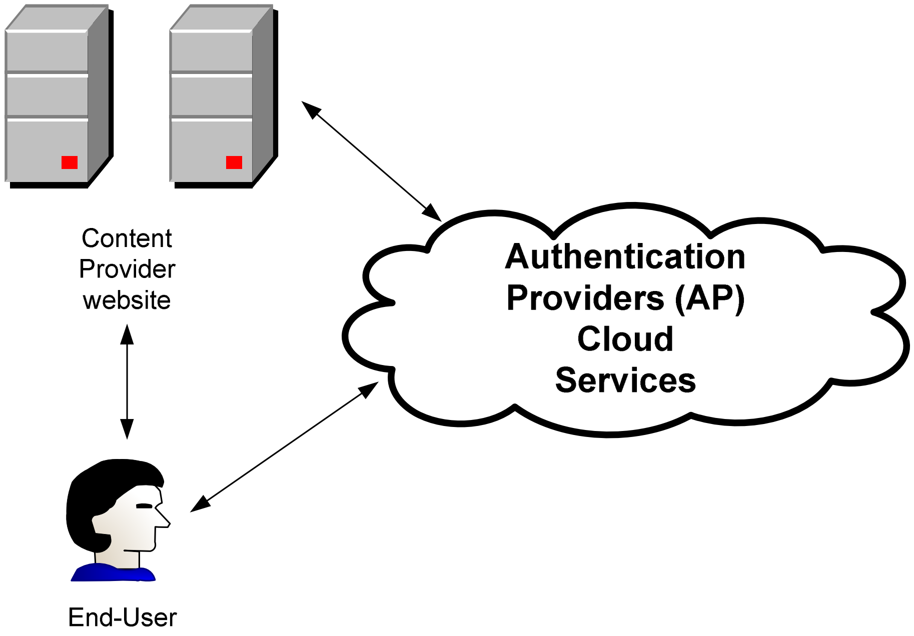

Essentially, you would have three, independent, but inter-connected layers, as shown below.

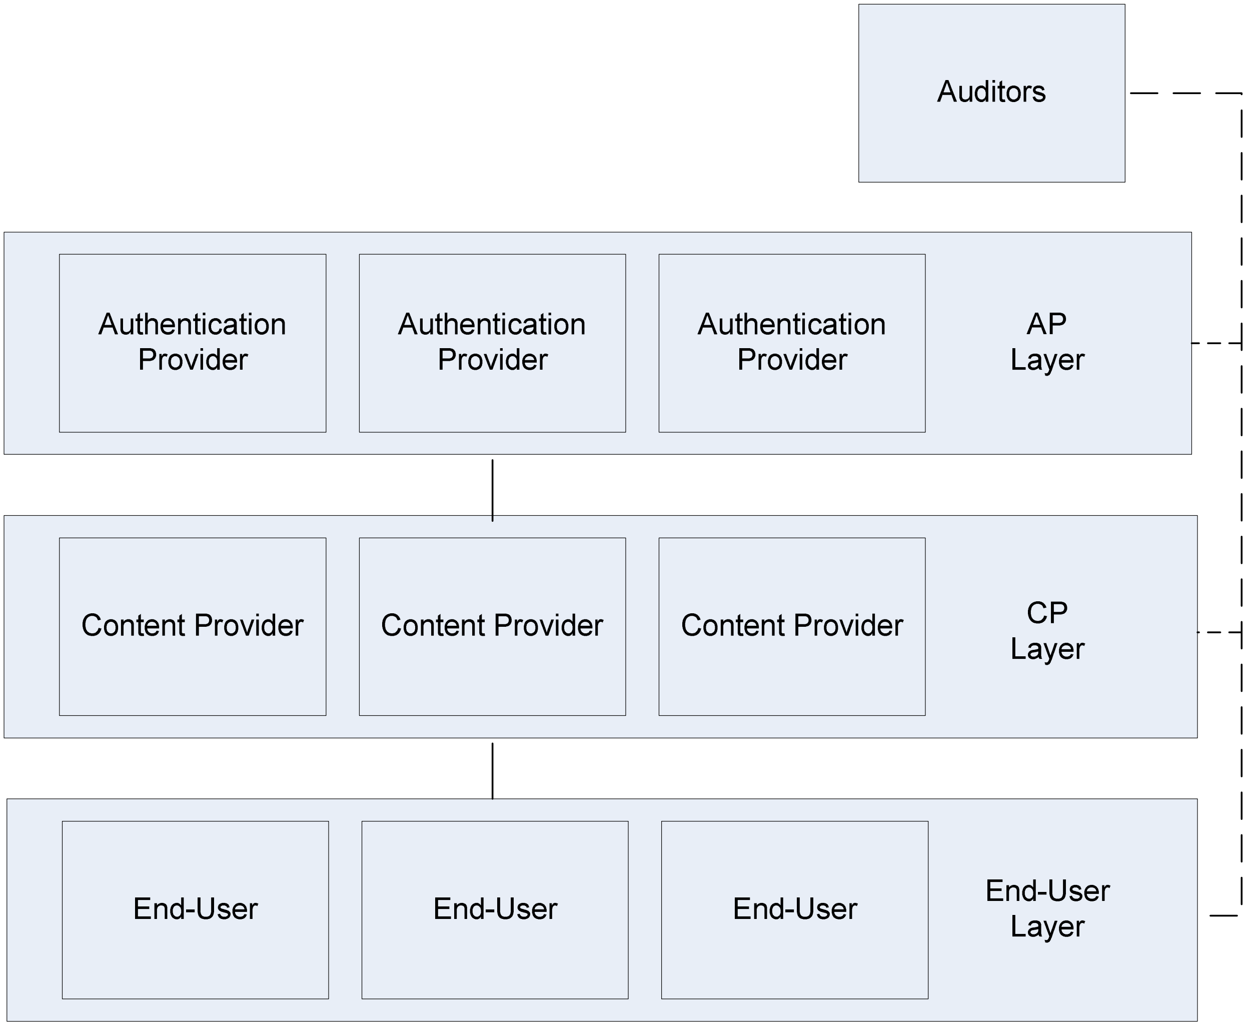

Connecting Existing Identity Systems

Each home computer user, business, enterprise, Internet Service Provider, and in some cases, entire countries have their own identity systems. This solution allows each individual identity system to be connected to the large identity metasystem using the appropriate protocols and coding, gateway, or service. With a gateway server or services, the Content Provider doesn't have to modify all their applications to take advantage of the global identity metasystem. See the diagram below.

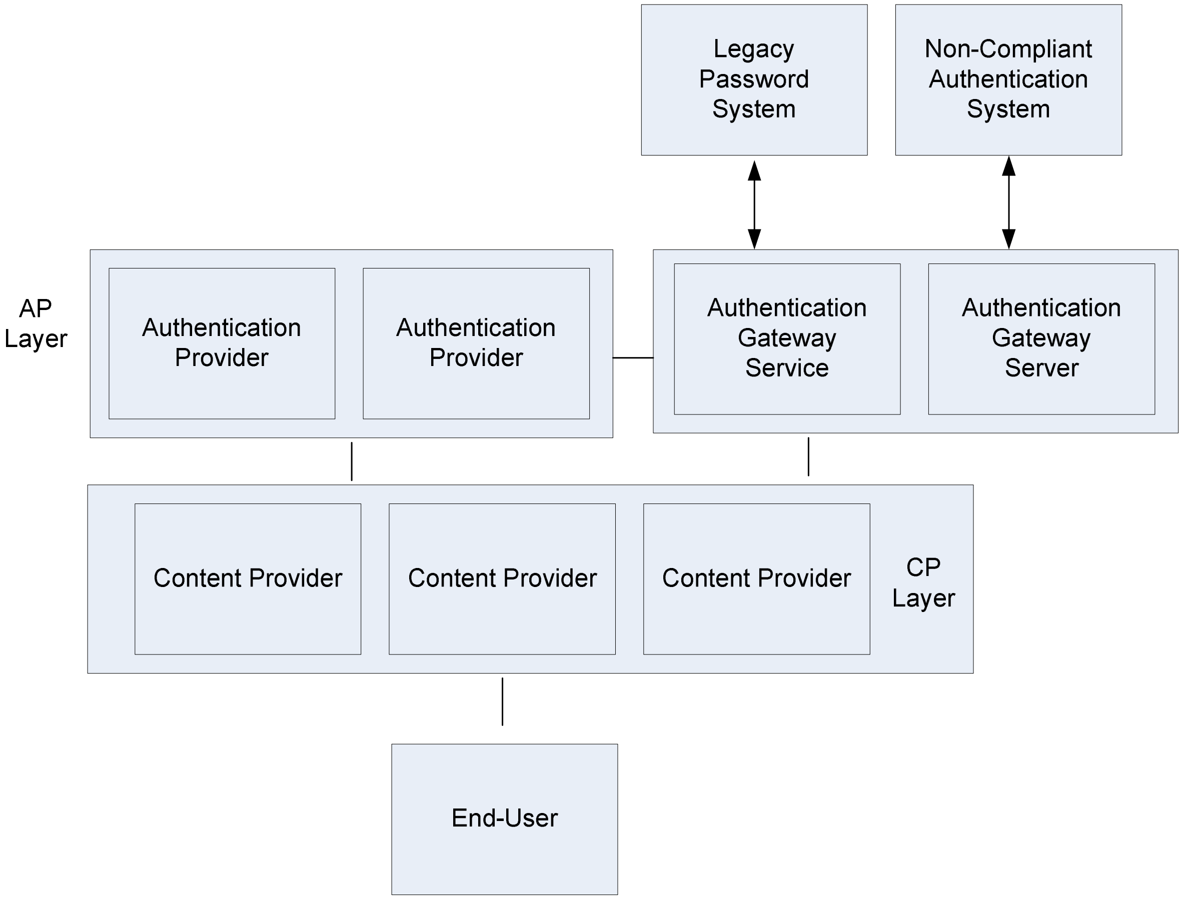

This open standards solution model has already been developed by vendors and supporting products are already available.

Open Standards Exist Today To Support Better Solutions

Today, there are enough existing open standards to support better solutions, including the possible solutions proposed above. These standards already have major industry support and products which implement them already exist. Those standards include:

- TCP/IP, especially IPv6
- Web Services (WS)
- Web Service Extensions (WS-*)
- WS – Trust
- WS – Federation
- Security Assertion Markup Language 2.0 (SAML 2.0)
- InfoCard
- DNSSec
- x.509 Digital Certificate Formats
- x.500 LDAP Directories
- Trusted Network Connect
- Network Access Control
- Trusted Platform Module chip

These standards and protocols can be used to make a more secure Internet.

FAQs

1. **Your solution decreases individual privacy. I'm completely against what you propose.**
   A: I, too, am a big privacy advocate, but I don't know of a solution that can significantly secure the Internet that doesn't involve improved, default, authentication and integrity. Let me know if you can think of one. Plus, my solution doesn't require that someone give up their anonymity, if they want to maintain it. Individuals can choose what level of identity or anonymity to give to a particular destination network. And the destination network can choose how to treat inbound traffic based upon the level of identity contained. Some hosts might choose to drop traffic, while others (I suspect the vast majority) will simply inspect the traffic more; while strongly authenticated traffic is given less inspection and faster transmission.
2. **Do you expect for your solution(s) to be adopted anytime soon?**
   A: It's highly unlikely in the near future, but dare to dream. I'm fairly confident that something along the lines of an Internet security service infrastructure will develop, because it is the only reasonable solution I can see for fighting larger, polymorphic threats. But overall, no, I don't think the world's vendors and security experts will come together to solve the Internet's big security problems until a tipping point event happens or the world's biggest governments get involved. Society, in general, is great at being reactive, and not so good at being proactive.
3. **Do any companies or entities currently support any part of your solution(s)?**
   A: Yes and no. No single company supports my exact solutions, but several already support similar ideas (or sometimes the exact concepts). Part of the reason I wrote this paper is that many of the ideas that I've been promoting for years, publicly and privately, are starting to become mainstream recommendations (e.g. Microsoft's End-to-End Trust initiative, Trusted Computing Group's IF-MAP standard, etc.), and I've been more right than wrong about the evolving threats. So, I thought by sharing more of my ideas in larger forums that people and companies with similar visions can come together and try to make a difference before the tipping point event happens.
4. **Would you be open to a public-private partnership, like what created the Internet?**
   A: Absolutely. If this solution is able to be accomplished, it will likely involve participation from both sides, and could likely involve national governments, as well. It is this necessary inclusion that makes it so hard to accomplish.
5. **Do any of your solutions offer enough significant advantages that vendors will be forced to adopt them?**
   A: No. That is a very large problem. How do you induce individuals and individual companies to act against their natural self-interests to do something for the great good? I'm hoping that significantly improved security, improved performance, and custom demand is enough to entice the initial players into the solution. After the big players and names are on board, the rest of the world should follow.
6. **Other Internet protocols, like DNSSec, SenderID, and IPSec, offer significant security improvements to the Internet, but haven't taken hold. How do you expect your idea to be any different?**

A: The problem with these other laudable protocols are that they are too limited in scope. Everyone knows that if you fix DNS, fix email, etc. that you are just fixing a point issue, and malicious Internet behavior will continue nearly unabated. My solution "fixes" all protocols. Fix the plumbing pipes and you don't have to fix nearly as much of the traffic in the pipes.

7. **You mention that there are few (i.e. meaning you know of some) defenses being developed that you think can significantly improve computer security. What are they?**
8. A: First, DNSSec, SenderID, and IPSEC are current protocols, that if adopted more completely would significantly improve security. Plus there are many new emerging defenses and protocols (End-to-End Trust, IF-MAP, any Trusted Computing Group standard, application signing, application and content whitelisting, Extended Validation SSL, Dshield-like data collection points, etc.) that appear to be very advantageous.
9. **Doesn't any solution of this type naturally discriminate against smaller companies and individuals who can't afford the newer stuff required to support the decision?**
   A: Yes, at least to some limited extent, whether intentional or not. It's like requiring a photo ID to vote. There's a valid reason to require a photo ID (i.e. voter fraud), but people who do not have easy access to a photo ID are discriminated against. There are many such decisions in the world (e.g. driver's license, social security card, passport, etc.). But with that said, it is my hope that the opt-in nature will allow, and very little discrimination (after all people not joining in will only be subject to the same scrutiny that they are today), will prove to be more like people moving from analog phone lines to broadband for Internet access (i.e. something people want to do). Many of the solution components (e.g. InfoCard, etc.) are zero cost.
10. **How can you realistically expect to increase security and not impact performance?**
    A: This is a difficult challenge, but with 2-6% of the Internet and 70-90% of all email being malicious in nature, if we can reduce those levels to near zero, it gives us a lot of room to play with before it actually slows down overall computing.
11. **Doesn't your solutions erode people's privacy?**
    A: Yes and no. Yes, at least a little, if you want improved security. No, if you choose not to participate. It's not a binary decision. We give up privacy all the time for more security (e.g. driver's license, city and community laws, etc.). And if you want to participate in better security without giving away your real identity, go pseudo-anonymous. Security and privacy are not completely exclusive of each other. Privacy isn't a binary choice anymore than security is.
12. **You propose that network traffic be encrypted and signed end-to-end. Won't many governments oppose this on the grounds that they need to inspect the traffic?**
    A: Probably, but so far there is no law that says I have to let the government read my information. Most governments have all sorts of rights and laws to try and read our traffic, but I don't know of any government law (I'm sure there are some) that requires people to let the government read it. For example, the U.S. government may sometimes have the legal right to capture your network traffic or listen on your phone calls, but there is no law saying that I cannot encrypt my phone call or network traffic further so they can't read it. Personally, I would strongly fight any law that says I have to show the government my information for basic services, and without a court order.

13. **You propose creating a new “dream team” consortium to solve the Internet’s major security issues. Doesn’t IETF, IANA, CERT, TCG, (or whoever), already exist to protect the Internet?**
    A: Nearly so. The Trusted Computing Group (TCG) is the closest model to what is needed. I’m open to imagining the security dream team as part of one of the aforementioned groups, as long as the team can act quickly (not something these former groups are always known for).
14. **How would the community trust rating server get populated with security domain trust rankings (i.e. would it be possible for a malicious person to maliciously malign my host or network in order to lower its trust rating)?**
    A: I’m not sure exactly how it would work, but yes, there would have to be protections in place to prevent malicious manipulation. This sort of thing is done for all sorts of services already with varying degrees of success.
15. **Why do you mention DNS in your solution as an example technology when DNS is so insecure?**
    A: For two reasons. First, it’s mainly mentioned as an example of a global, redundant, distributed infrastructure service. Attackers will try to take down any global security service, so we need to mention that it is possible, as demonstrated by DNS, to do it globally and secure. Second, DNSSec is one of three technologies (the other two being IF-MAP and Sender ID) that are true security solutions. I consider most other things security theater.
16. **How would you do X and XX in your solution?**
    A: I don’t have all the answers. That’s why I propose bringing together a dream team of experts under each component discipline to solve the tough technical challenges.

If you’ve reached this part of the paper, I thank you for your time and participation. Feel free to send comments to me at roger_grimes@infoworld.com.

Bibliography

[1] History of the Internet, Wikipedia, http://en.wikipedia.org/wiki/History_of_the_Internet

[2] MessageLabs Intelligence Report April 2008, MessageLabs, http://www.messagelabs.com/mlireport/MLI_Report_April_2008.pdf

[3] Internet has a trash problem, InfoWorld Magazine, April 1, 2008, http://www.infoworld.com/article/08/04/01/Internet-has-a-trash-problem_1.html; and Up to 3 percent of traffic is malicious says researcher, CSO Online, April 1, 2008, http://www.csoonline.com/article/326013/Up_to_Three_Percent_of_Internet_Traffic_is_Malicious_Researcher_Says; and A Peek at ISP DDOS, Spam Traffic Trends, DarkReading, April 1, 2008, http://www.darkreading.com/document.asp?doc_id=149866&WT.svl=news1_1; and 2% of Internet Traffic Raw Sewage, Arbor Networks, March 31, 2008, http://asert.arbornetworks.com/2008/03/2-of-internet-traffic-raw-sewage.

[4] Hackers harpoon executives, The Sydney Morning Herald, May 6, 2008, http://www.smh.com.au/news/security/going-after-the-big-phish/2008/05/06/1209839606696.html; and History and current status of phishing, Wikipedia, http://en.wikipedia.org/wiki/Phishing; and US Federal Court Subpoena Phish, Junkfax.org, http://www.junkfax.org/fax/phish/uscourtsPhish.htm; and U.S. District Court Subpoena, Snopes.com, http://www.snopes.com/fraud/phishing/subpoena.asp.

[5] Top Spam Bots Exposed, Secure Works, April 8, 2008, http://www.secureworks.com/research/threats/topbotnets; and Top botnets control 1M hijacked computers, ComputerWorld magazine, April 9, 2008, http://www.computerworld.com/action/article.do?command=viewArticleBasic&articleId=9076278; and Storm botnet, Wikipedia, http://en.wikipedia.org/wiki/Storm_botnet.

[6] From Fast Fluxing to Rockphish – Part 1, McAfee AVERT Labs, November 30, 2007, http://www.avertlabs.com/research/blog/index.php/2007/11/30/from-fast-flux-to-rockphish-part-1.

[7] What is server-side polymorphism?, Anti-Virus Rants, August 10, 2007, http://anti-virus-rants.blogspot.com/2007/08/what-is-server-side-polymorphism.html

[8] Compromised web sites serve more malware than malicious ones, Ars Technica, January 22, 2008, http://arstechnica.com/news.ars/post/20080122-compromised-websites-serve-more-malware-than-malicious-ones.html; and Malware spikes in 1Q as hackers increasingly infect web sites, InformationWeek, April 24, 2007, http://www.informationweek.com/news/internet/showArticle.jhtml?articleID=199201032; and Massive Site Compromise: The Siege Continues, TrendMicro Malware Blog, April 3, 2008.

[9] Privacy Protection: The Government is No Help, InfoWorld magazine, June 16, 2006, http://www.infoworld.com/article/06/06/16/79260_25OPsecadvise_1.html

[10] 2007 Internet Crime Report, FBI, http://www.ic3.gov/media/annualreport/2007_IC3Report.pdf and CSI 2007 Survey, FBI, http://i.cmpnet.com/v2.gocsi.com/pdf/CSISurvey2007.pdf

[11] Russian Business Network, Wikipedia, http://en.wikipedia.org/wiki/Russian_Business_Network.

[12] Rockphish, EU Spam Trackers, http://spamtrackers.eu/wiki/index.php?title=Rockphish.

[13] CERT Advisory CA-2003-04 MS-SQL Server worm, CERT, January 25, 2003, http://www.cert.org/advisories/CA-2003-04.html.

[14] SQL Slammer (computer worm), Wikipedia, http://en.wikipedia.org/wiki/SQL_slammer_worm.

[15] Iloveyou worm, Wikipedia, http://en.wikipedia.org/wiki/ILOVEYOU.

[16] Blaster (computer worm), Wikipedia, http://en.wikipedia.org/wiki/Blaster_worm.

[17] Google Health: A First Look, Official Google Blog, 2/28/2008, http://googleblog.blogspot.com/2008/02/google-health-first-look.html.

[18] One Laptop Per Child, One Laptop Per Child website, http://laptop.org.

[19] Universal Description Discovery and Integration service, Wikipedia, http://en.wikipedia.org/wiki/Universal_Description_Discovery_and_Integration.

20 Trusted Network Connect IF-MAP Announcement FAQ, Trusted Computing Group, https://www.trustedcomputinggroup.org/specs/TNC/IFMAP_FAQ_april_28.pdf.