



Collaborating to Secure Consumer Devices

Promoting Device Health for a Safer, More Trusted Internet

Kevin Sullivan
Senior Security Strategist
Trustworthy Computing
Microsoft Corporation

Microsoft[®]

Collaborating to Secure Consumer Devices

© 2011 Microsoft Corp. All rights reserved.

This document is provided "as-is." Information and views expressed in this document, including URL and other Internet Web site references, may change without notice. You bear the risk of using it. This document does not provide you with any legal rights to any intellectual property in any Microsoft product. You may copy and use this document for your internal, reference purposes.

Contents

- Considering Device Health 1
- Building on Existing Efforts 2
- A Vision of Healthy Devices 3
- An Opt-In Approach..... 5
- User Control and Privacy Considerations 6
- Collective Action 8
- Pursuing Progress 9

Considering Device Health

There are currently more than 2 billion¹ Internet users worldwide. Five hundred million of those connect with broadband. By 2012, more than 4.5 million people worldwide will be using mobile phones—many of which will be used to connect to the Internet. The number of devices connected to the Internet shows no sign of slowing. To the contrary, advances in technology innovation are expected to deliver even more Internet-enabled consumer devices, including televisions and kitchen appliances. At the same time, the increasing number and complexity of cyber-attacks targeted at these devices threaten our increasingly connected society.

Current data suggests that cybercrime is having a very real impact on consumers and businesses alike. Microsoft cleaned more than two million computers around the world in the second quarter of 2010 that were infected with malware². A recent Gallup³ crime survey reports more than ten percent of adults in the United States have been victims of online crime. Unprotected devices can increase the risks to consumers, businesses, and governments who rely on the Internet by providing criminals a platform they can use for attacks. Unhealthy or insecure consumer devices could be used to attack other systems including critical infrastructures. Therefore, in the quest for a healthier Internet ecosystem we must consider how to improve the health of devices.

Individuals and organizations have a wide variety of choices to help defend their devices against cyber threats, including antimalware solutions, firewalls and security updates. However, despite the broad availability of these technologies, they are not always fully deployed or kept up to date. Enterprises spend large amounts of time and money to manage risk with well-trained staff and sophisticated systems. Consumers, on the other hand, generally lack the expertise and ability to protect their devices themselves.

In addition, while individual defense technologies are required, they are no longer sufficient by themselves. Microsoft believes that society must explore a combination of collaborative actions by individuals, industry, and where necessary, policy makers to improve the health of consumer devices and consequently to reduce the overall risk for consumers, businesses, governments and critical infrastructures. Microsoft's earlier paper, [Collective Defense: Applying Public Health Models to the Internet](#)⁴ focuses on catalyzing the Internet ecosystem to address Internet Health. A device health model aimed at preventing the infection of consumer devices is a key element to achieving a safer, more trusted Internet.

¹ International Telecommunications Union, The World in 2010 ICT Facts and Figures

² Microsoft Security Intelligence Report v9, Microsoft.com/sir

³ Gallup Crime Survey, 12/2010 gallup.com

⁴ Microsoft.com/security/internethealth

Building on Existing Efforts

Industry and government have begun efforts to help protect consumers against online threats. For example, the Internet Industry Association of Australia launched its voluntary Internet Service Provider (ISP) Code of Conduct⁵, which creates a notification system for consumers with compromised computers and provides standardized resources to help clean them. In a similar example, in Germany the Anti-Botnet-Advisory Centre⁶ works with local ISPs to notify consumers infected with malware and provides them with tools and guidance to remove the infection from their computers. These efforts are important and must continue in order to help consumers stay safe against the wide range of online threats. However, these efforts also have well-defined, limited goals. These efforts are primarily reactive as they are initiated only after a consumer's device is compromised. These programs are intended to help a user remove an infection but preventing future infection is a far more complex challenge. A continuous cycle of notification and remediation can frustrate users and impede effectiveness. Current efforts to make consumers aware of online threats and the appropriate protections are important but lack the ability to deliver tailored notifications to consumers whose devices are at risk. Actively helping consumers secure their devices is a key first step in transforming the current posture from reactive to preventative.

⁵ <http://iia.net.au/images/resources/pdf/iiaCyberSecurityCodeImplementationDec2010.pdf>

⁶ <https://www.botfrei.de/en/index.html>

A Vision of Healthy Devices

The “Collective Defense: Applying Public Health Models to the Internet” paper calls for applying concepts from public health models to address cyber threats. One specific proposal in the paper was to promote efforts to better demonstrate the health of devices using a health certificate. The idea is that a consumer whose device lacks sufficient security protections to protect against online threats would receive timely and prominent notifications in the context of their online interactions. The consumer would then be connected with resources and guidance to install protections on their device before it can be compromised. At the RSA Conference 2011, Microsoft demonstrated a proof of concept⁷ with the intent of furthering the dialogue on applying public health model concepts to address cyber threats.

The [device health demonstration](#) involved a user who participates in a pilot program with her bank. The goal of the bank’s pilot program was to help the bank’s customers improve the health of their devices thus reducing the chance that their online banking credentials and/or identity would be stolen. As part of the pilot, the user installed software that can communicate that status of defensive technologies, such as antimalware software and security updates, to her bank as part of the online banking experience⁸. Before she could complete a high value transaction, she was notified that her device did not have one of the key security protections, an up-to-date version of anti-malware software. In this case, she was notified of the problem and provided with an informative message and link to update her antimalware software. After updating the software, she was successfully able to complete the online banking transaction with the added confidence that she is better protected by using a device that had up to date protections against a variety of threats.

⁷ This demonstration is not necessarily a final architecture or solution; rather it is intended as a proof of concept to stimulate necessary debate on the approach including full examination of security and privacy ramifications.

⁸ Online banking is used as a notional example throughout this paper; however, any provider of high value transactions could employ the device health model. Throughout this paper we will use the term “online service providers” to reference these providers.

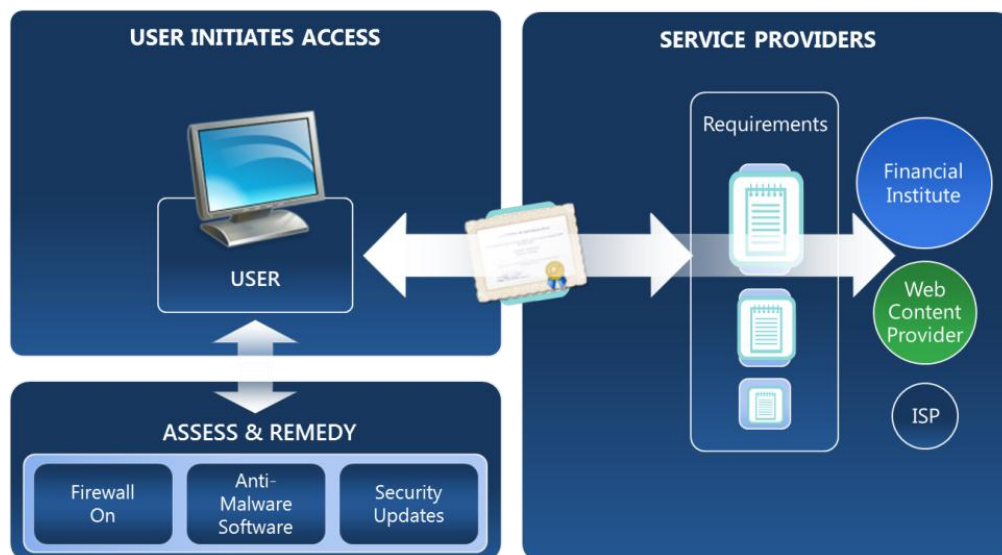


Figure 1 - Notional Device Health Scenario

Consumer security decisions can be aligned with the value of the asset or transaction that they are trying to protect and online service providers can specify controls commensurate with the risk. Furthermore, the increased security posture of each consumer device will help address the threats to other consumers, their devices and even critical infrastructures. With the proof of concept demonstration and this paper, Microsoft hopes to begin working across the ecosystem to turn this proposal into a roadmap for future solutions to improve device health and increase consumer safety online.

An Opt-In Approach

One of the first questions in response to this proposal is “should this experience be mandatory for all internet users?” Microsoft believes that enabling consumers and online service providers to take an opt-in approach is the best way to begin to implement this model. Consumers who participate in the device health model will help to protect against fraudulent use of their computer and online accounts. Participating online service providers may be able to realize reduced occurrence of fraud and associated costs. For example, banks or other financial service providers and their customers might be early adopters of the model since their interactions with consumers often involve high value transactions. More specifically, Microsoft expects the health level, attributes of defensive technologies that a consumer device demonstrates in order to participate in a transaction, will differ based on the nature of transaction. Other online service providers such as ecommerce sites and social networks could also offer device health checks to their users, again depending on the risk that each transaction presents to them and their customers. Over time, Microsoft believes that this approach will lead a majority of internet users to choose to participate in the device health model, which will result in increased security of the entire consumer device population.

While pursuing the device health model with an opt-in approach, consideration must be given to user control and privacy.

User Control and Privacy Considerations

To better protect internet users from online threats, the entire community of internet stakeholders should be involved. Determining the best path forward involves answering several key questions balancing user control, privacy, and security ranging from policy to business and technology.

Who decides the health requirements for devices?

Among the most pressing questions are “what are the health requirements for devices,” “which transactions do they apply to,” and “which organizations determine them?” These are not easy questions to answer based on the current threat landscape, and the future evolution of threats will make it even harder. Basic security measures such as antimalware software, automatic security updates and an active firewall help users prevent many computer infections and should be considered as a minimum baseline. However, attackers are always looking for ways to circumvent existing controls. Evolving threats require an evolving baseline of security controls. While this model cannot protect every user from every threat, it can help to improve the overall security posture of consumer devices, allowing for focus on the most serious threats.

It would be a good practice for all participating online service providers to notify users if their device does not have one of these basic security controls as demonstrated by the device health model. The model provides extensibility such that an online service provider could request their customers’ devices to exhibit a greater number of security controls. Some service providers – for example banks – may require a consumer device to exhibit more security controls than an internet service provider. While there may be legitimate reasons for this, such as banks looking to prevent specific financial malware, caution must be used to ensure that this ability is not abused.

Online service providers should not be able to require that their customers install arbitrary software on their devices. It is important to ensure that any required protection programs are used only for ensuring the health of the device and not for other purposes such as enforcing intellectual property rights. This might be accomplished by an independent certification process for device health requirements and protections. At a minimum, there should be a separation of duties between those creating the protections and those setting the requirements.

What happens if devices are not healthy?

If a consumer participates in the device health model and their device does not meet the health requirements of a particular service, what should the response be? An online service provider has a range of different responses. These include doing nothing, notifying the consumer that their device failed the health check, providing a differentiated or partially restricted experience, or blocking access⁹. Different responses could be appropriate in different scenarios as determined by the level of risk to the user or the service provider. For example, the presence of out-of-date antimalware signatures on a consumer's device might be communicated by the consumer device to the online service provider and could then trigger the online service provider to deliver a simple notification back to the consumer's device. Alternatively, in a situation where a consumer's device had not installed a critical security update necessary to combat a current internet worm, the user who had opted in to the program might have a restricted experience for that service until they install the necessary update.

What personal information is involved in determining device health?

The purpose of the device health model is to ensure the necessary protections exist on a device to help protect the consumer from harm. Determining a device's health should not require any personal information about users. The information passed from a consumer device to a service provider would be limited to the status of the protections (e.g. antimalware signatures are current or automatic security updates are on.) Any final architectural approach to delivering a device health solution to consumers should be implemented in a manner that protects privacy. In scenarios where a unique ID would be used to identify a particular device, the rules around any use of that ID must be clearly described and strictly enforced.

What about other types of devices?

Next generation computing devices such as mobile phones, televisions and other consumer electronics have ever-increasing computing ability and play a central role in our connected society. As more of these devices come online, today's nascent attacks against them may become commonplace and thus it will be important to help protect them against these threats. The opt-in approach enables consumers to decide if the health of their device would be checked before accessing certain services, so devices that do not support requests to demonstrate their health would not be impacted. Over time, Microsoft anticipates that the device health model will evolve to enable scenarios where devices can check their health and remediate issues automatically, creating a better fit for these next-generation computing devices.

⁹ Microsoft is committed to respecting and protecting the right to freedom of expression and recognizes that blocking access is an extreme response that is technically possible but socially unacceptable. Where possible, notifying customers and helping them remedy the problem is a preferred approach.

Collective Action

As no single entity can defeat global cybercrime by itself, members of the internet ecosystem must take collective action. The collective defense proposal calls for all members of the internet ecosystem to work collaboratively to help protect consumers who may be unaware that their device is vulnerable or compromised. The internet ecosystem ranges from software vendors to online service providers, consumers and governments. Each of these entities has a unique and valuable role to play and the device health model helps enhance the ability for a variety of ecosystem members to help their customers maintain the health of their devices in an extensible and user-friendly manner.

The goal of a device health model intends to keep consumers' devices from being infected with malware and reduce the number of infected devices going forward. However, the ecosystem must still collaborate to find ways of helping consumers with devices that have already been compromised. These consumers are the focus of many existing efforts including those in Australia and Germany mentioned earlier in this paper. Most efforts to notify and assist consumers with infected devices rely upon the unique ability of ISPs to match a report of an infected device to the customer whose device is infected and thus they often bear the potentially costly and labor intensive burden of notifying their customers and helping them resolve the issue. The device health model, implemented with an opt-in approach, could distribute the responsibility of helping consumers protect their devices across the internet ecosystem with the goal of reducing the number of infected devices. Participating consumers' devices could be kept secure with help from the service providers they interact with. An ISP could choose to participate in the device health model to reduce bandwidth and support costs or as a value added service to their customers. Alternatively, they may choose to limit their involvement to situations where a customer is actively under attack and use its capabilities to limit further damage.

Pursuing Progress

Microsoft and other members of the IT industry endeavor to help consumers have the most secure online experiences possible and to reduce the threat of botnets to critical infrastructures and other users' devices. The online threats to consumers are very real and growing and they need help to maintain the security of their computing devices. Microsoft believes that by working collaboratively the industry can make progress towards these goals by pursuing a device health model that provides consumers with simple and actionable guidance aided by prompts and notification as part of high value online transactions.

The principles of voluntary behavior and respect for user controls and privacy outlined in the "Collective Defense: Applying Public Health Models to the Internet" paper hold true here. The device health model outlined in this paper is based on both consumers and service providers voluntarily deciding to help manage the health of consumer devices. Efficient and effective device health implementations will lead large parts of the ecosystem to participate and improve internet security considerably. As the model and technology implementations evolve, user control and privacy must be considered at every step of the way.

There is no single product or service that can help to secure a global population of internet users. Much work is needed in both the technical, business and social domains before achieving this goal. There are protocols to be designed and implemented along with an evolving suite of security software. Additionally there are large outstanding questions about what the health requirements are, who sets them, and how are they updated. When will device health be implemented for the next class of consumer devices such as phones, televisions, or kitchen appliances? These decisions need to be worked out in an open and transparent manner that is inclusive of the entire internet community. It is Microsoft's hope that this technology demonstration and dialog encourage the future developments that are necessary in order to enable a collective defense of consumers against online threats.