Introducing the Intel Science and Technology Center for Secure Computing

Abstract

Suppose all your personal and professional digital devices could interact on your behalf, like trusted agents. Imagine having confidence that your personal, health, entertainment and financial information was never used or changed without your approval – whether it was in the cloud, on your phone, in your briefcase or on your desk. Also suppose that the security in your devices was obvious to you without any special expertise, and that you felt confident that you could tell what actions are safe and what is risky.

Today such a world remains a vision of the future. Yet, we hope people will benefit from digital technology for their well being, happiness and success in an environment that increasingly insures their digital safety.

The Intel Science and Technology Center for Secure Computing is undertaking a program of exploration and investigation to help make this aspiration a reality.

Background

Digital devices act for and on behalf of people even for the most sensitive and critical activities, immensely enriching lives and improving productivity.

Yet, malware has infected millions of machines[1], malicious actors have subverted computer systems to steal enormous amounts of personal and sensitive corporate information (Aurora)[2] and digital attacks have even compromised the operational safety of physical devices (Stuxnet)[3]. Unfortunately, the threats have dramatically increased and are overwhelming "traditional" solutions.

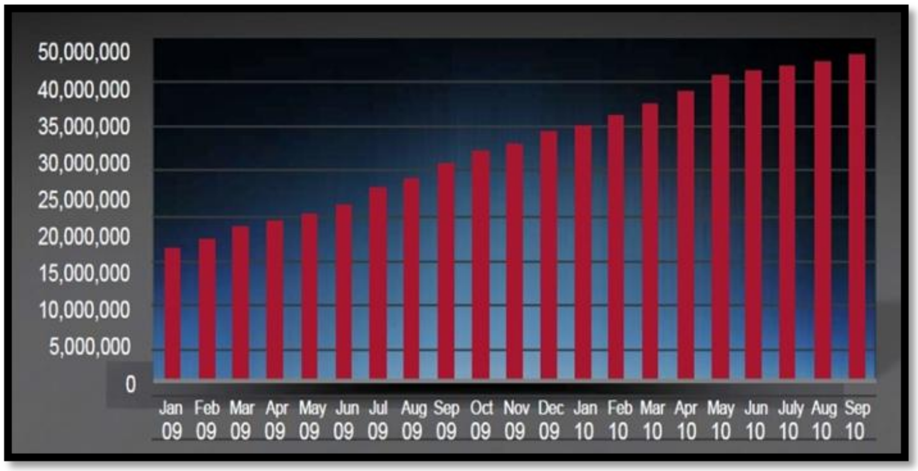

Malware is growing dramatically over time (Source: McAfee)

As an example of the increasing scope of digital technology, consider the rise of smart phones. These devices allow people to access, communicate and affect much of the

[1] Ben Rooney, Wall Street Journal, May 18, 2011: "One in every 14 downloads is a piece of malware according to information released by Microsoft yesterday. To try to inhibit the growth, anti-virus defenses have to attack the problem from other directions. Despite improved browser security over the last few years, new forms of propagation have developed including highly-targeted attacks known as spear-phishing, which go after individuals or groups of individuals."

[2] ISec Partners, February 17, 2010: "On January 12th, 2010 Google publicly revealed that they were the victim of a "highly sophisticated and targeted attack on our corporate infrastructure originating from China that resulted in the theft of intellectual property from Google." Google was not the only company affected by this attack; at the time Google notified over 30 other companies of infection by this malware. In the time since then, further investigations have uncovered that over one hundred companies may have been targeted, although it's difficult to ascertain how closely related these attackers are to Google's assailants."

[3] Symantec Corporation, Stuxnet Dossier, February, 2011: "Stuxnet is one of the most complex threats we have analyzed. ... primarily written to target an industrial control system.... Industrial control systems are used in gas pipelines and power plants. Its final goal is to reprogram industrial control systems (ICS) by modifying code on programmable logic controllers (PLCs) to make them work in a manner the attacker intended and to hide those changes from the operator of the equipment. In order to achieve this goal the creators amassed a vast array of components to increase their chances of success. This includes zero-day exploits, a Windows rootkit, the first ever PLC rootkit, antivirus evasion techniques, complex process injection and hooking code, network infection routines, peer-to-peer updates, and a command and control interface."

information that influences their lives wherever they are: financial transactions, professional communications, personal communications, internet browsing, media sharing, interacting with our governments and community, gaming and entertainment. They also can record the history of our lives: where we've been and when, who we've met and what we've done with them, even what stresses our bodies have been subject to and how it affects us. Such comprehensive access presents enormous possibilities and poses profound new risks. Much of this information can and is being transmitted to various services in the cloud, which, in turn, are composed with other services. This amplifies the benefit derived from the information but can also amplify the adverse consequences of attacks on the information and the infrastructure which houses and transmits it. In fact, it is alarming that these benefits are jeopardized by remote attacks that can be mounted, deployed and executed very quickly and on a vast scale.

A safer digital world

The Intel Science and Technology Center for Secure Computing is embarking on a research program to develop security technology that preserves the immense value of digital devices by making them safer, increasing user confidence in platform and application safety without special training. This program activity is called Secure Computing Research for User Benefit ("SCRUB"). It includes developing the ability to create safer devices and software that can cooperate, on behalf of a user, in the cloud, on a PC and on a personal device, with increased practical assurances. The goal is to allow all end users to be as safe or even safer in the digital world as they are in our safest communities while preserving the tremendous and expanding benefits digital technology affords us.

The Intel Science and Technology Center for Secure Computing is a research collaboration between Intel Labs and security experts from outstanding academic institutions including the University of California, Berkeley (UCB), Carnegie Mellon University (CMU), Drexel, Duke and University of Illinois (UIUC). UCB will be the hub of the center, coordinating these activities. The center will be co-led by UCB professor David Wagner and Intel Senior Principal Engineer, John Manferdelli. Initially about ten academic and four Intel researchers will participate in the center.

Much of the research of the center will focus on the development of systems and application software to enable secure computing. The center will also use an open IP model wherein results are made public in the form of technical publications and open source software. Intel will use the results and the developed knowledge to guide the development of future hardware platforms that help ensure user and infrastructure safety, and to develop a broad consensus for the importance and approach to improved, usable security for computing. Intel has been at the forefront of developing hardware to enable secure computing, recently introducing several "trustworthy computing" capabilities in hardware like hardware assisted cryptography such as the Intel® AES New Instructions and secure isolation of virtual machines using Intel® Virtualization Technology[4]. Intel strives to build safe, reliable and capable functionality for a broad range of devices[5], including smart phones and other mobile technologies. Many of these innovations will be used by the center as a foundation for further progress.

Overview of the SCRUB program

The Intel Science and Technology Center for Secure Computing will conduct research that benefits end users of technology, making computing more useful, safer, and more trustworthy.

Initially, the center has identified five areas of investigation that form the SCRUB program:

1. *Developing a secure, thin intermediation layer for security isolation that enables protected environments which work together on the user's behalf:* Many current software systems are homogeneous so that an attack on any application or element of the system can lead to a total compromise of all user data and activities on the device (and worse, on all services that "trust" that device). As in the physical world, solid, reliable isolation can protect critical activities from external forces (like keeping germs out of an operating theater) or contain dangerous activities so they don't adversely affect other activities (like an isolation room in a hospital to prevent the spread of highly contagious diseases). To improve desktop security and mobile security, we will develop hardware and software to provide such isolation. This enables us to establish identifiable islands of trust that can be used for cooperating critical tasks, say home banking, even when parts of a user's computer swims in a sea of malware. It also allows us to contain potentially unsafe activities like internet browsing.

2. *Developing secure mobile devices that can be used in all facets of our lives:* We will investigate dramatically improving the security of smart phones and other mobile devices. A special focus will be on making third-party applications safe and secure, while supporting a rich market for these applications. This would allow users to employ a single phone for personal or business use, allowing each activity to be safe while preserving the flexibility and integrity of each activity.

[4] http://www.intel.com/technology/virtualization/technology.htm, "Intel® Virtualization Technology (Intel® VT)".

[5] See Intel Technology Journal, Volume 13, issue 2, Advances in Internet Security, available at http://www.intel.com/technology/itj/2009/v13i2/index.htm.

3. *Enabling the safe use and transmission of confidential and personal data wherever it is used:* As data moves to the cloud and flows through complex distributed systems, it becomes more challenging to manage the security of this data in a consistent and reliable way. We propose to develop techniques to provide consistent protection for data, no matter where it may be stored or operated upon.

4. *Using the network to provide better security and quality of service:* We will develop new network architectures providing insightful, scalable network monitoring of enterprise and other networks, as well as novel network architectures that allow end devices to benefit from the processing power of capable network elements --- all securely.

5. *Expanding the study of secure analytics:* We will develop methods to measure the security of complex systems by using the vast information available on the internet. We need to develop techniques to strengthen analytics against adversarial manipulation and to develop subtle and reliable predictors of safe operation as well as signals of security failures.

A closer look at the research areas

Secure thin intermediation layer

As mentioned, often in today's end-user systems, the only intermediary between applications and the hardware is a complex, unauthenticated, full-featured operating system. As a consequence, a single piece of malicious software or mistaken user action endangers all applications and data. Since many users can be easily fooled into accepting malicious software, security on today's end-user systems is very fragile. We propose to study how to improve security of end-user desktops and smart phones by introducing a secure thin intermediation layer ("STIL") in conjunction with better user interface paradigms to reduce the danger presented by malware and errors.

One of the most important users is the software developer. To be useful, the STIL will have to have minimal performance impact and be simple enough for developer use. Both are goals for the STIL.

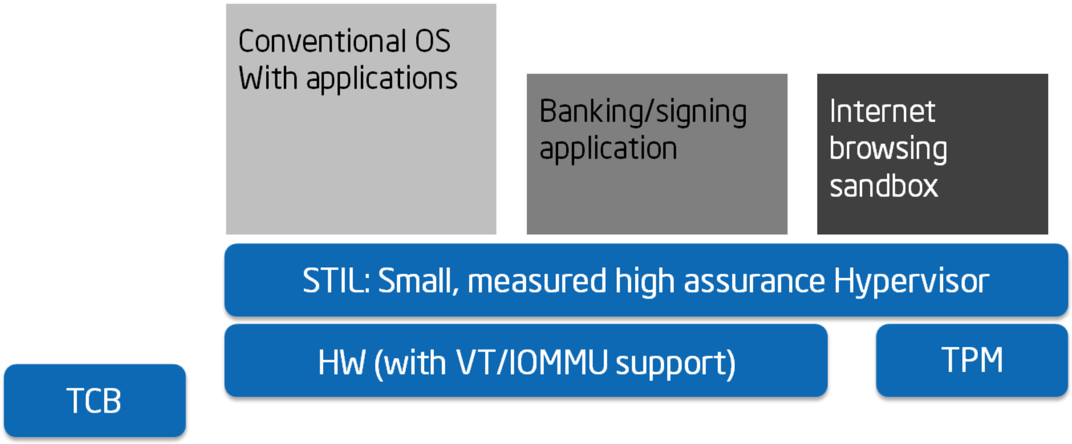
Conventional OS
With applications
Banking/signing application
Internet browsing sandbox
STIL: Small, measured high assurance Hypervisor
HW (with VT/IOMMU support)
TPM
TCB

Simple Secure Thin Intermediation Layer

Beyond the traditional virtual machine (VM) role of offering multiple operating partitions, a STIL could host narrow execution environments (small, tailored clients) for particularly sensitive applications, such as banking or access to corporate data. A STIL could provide isolation between these execution environments. It could also serve as an enabler for tracking and managing information flow, managing risks from foreign software that the user desires to introduce into his or her environment, and interacting with remote services and the intervening network architecture.

Furthermore, each STIL-isolated environment could, with very high assurance, authenticate (identify) other such environments and thus "trustworthy" software elements can work together to enforce user policy and privacy in the face of determined attacks even in high-value targets like cloud data centers and phones.

Safe mobile devices

Today's mobile devices combine the computing power, storage and communication capabilities of powerful desktop machines ten years ago with highly portable form factors. The capabilities of these devices benefit from Moore's Law technology scaling and will continue to improve as their recent evolution has taught us.

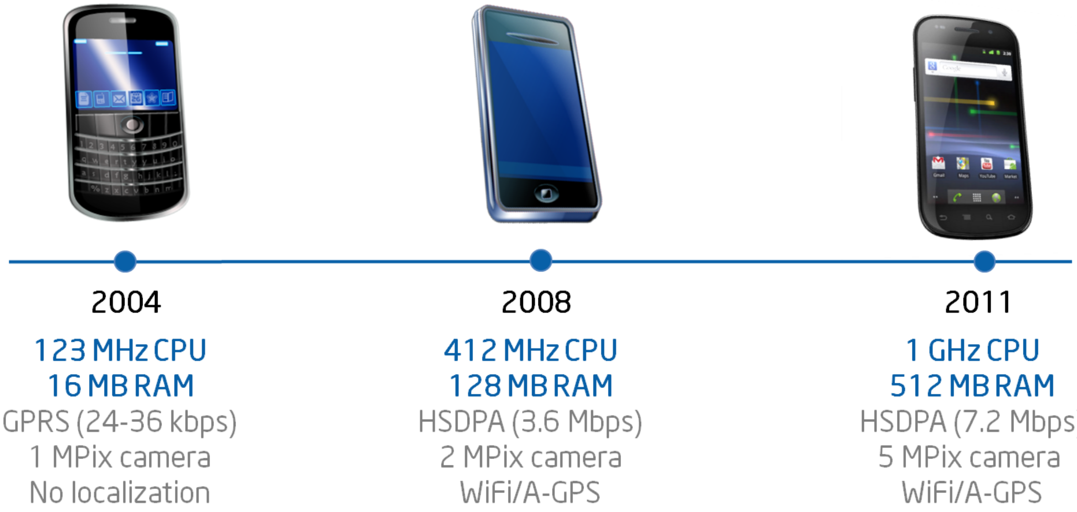

Third-party applications are an important way to accommodate the diverse segments of a user's life, so we desire a vibrant application marketplace, and one where people can feel confident in the security of the third-party apps they use. To this end, we will study methods for static and dynamic analysis of third-party apps. A STIL provides a mechanism for isolating third-party applications and for interposing application interaction, but it alone cannot ensure that third-party apps will be safe. Rather, it must operate in concert with powerful, automated analysis of the application and its actions.

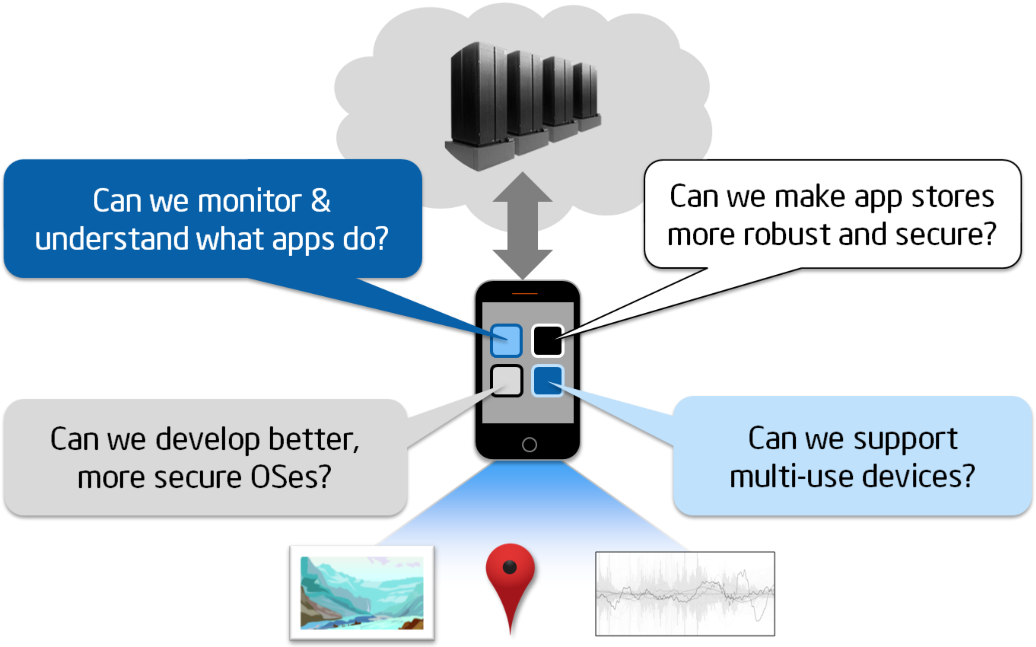
Can we monitor & understand what apps do?
Can we make app stores more robust and secure?
Can we develop better, more secure OSes?
Can we support multi-use devices?

Applications on Mobile devices

Safe use and transmission of data

While the integrity of the end-user device and server endpoints is essential, data increasingly flows through complex distributed systems and is stored in various forms in a variety of tiers. Thus, it is essential to develop consistent and reliable protection for user data, no matter where it is retained and used. We will take a multi-pronged approach. First, we will develop means of attaching security policies to the data. The combination of data and policy forms a sealed capsule. The capsule can be unsealed only within a secure execution environment that will not leak the capsule's contents or violate the attached policy. Secondly, we will develop privacy-preserving services so that, for example, only aggregate information is made available to services, not information about specific individuals. Finally, for legacy applications, we will develop wrapper techniques for information-flow tracking.

Using the network

The end-to-end principle that is a foundation of the Internet architecture has often been eroded to achieve other performance and security goals, as in the use of firewalls, NAT, proxies, gateways, virus scanning, traffic shaping, load balancing, enhanced data availability and VPNs. As we rethink the Internet architecture, it is essential that secure computing concerns be foremost, rather than an afterthought. With the emergence of intelligent routing agents and Trusted Platform Modules ("TPMs"), we have the opportunity to enable sophisticated, internet-wide, semantic analysis to detect and respond to attacks, as well as to see the emergence of new threats in a rapid, automated manner. In effect, aspects of intermediation and analysis can potentially be performed with an endpoint independent, global view. For example, taint tracking for mobile applications may be pushed into the network, into the cloud, or both.

Secure analytics

Machine learning, data mining, and artificial intelligence are being used as general tools for data processing and decision-making in computer applications (e.g., spam filtering, search engine page ranking, intrusion detection, virus detection, etc). They enable systems to respond to and classify evolving real-world hostile and benign inputs (data, programs, etc.). As concrete examples, we expect that the STIL, secure mobile devices, and security-centric network architecture research efforts will generate significant amounts of data against which we can *safely* apply security analytics tools.

A significant obstacle to the general use of these tools is the risk associated with operating in the presence of adversaries (adversarial environments) – adversaries attempt to manipulate the inputs to the tools to disrupt analytics, cause analytics to fail, or engage in malicious activity that fails to be detected.

We plan to work actively with Intel and other research partners to identify applications and large datasets that we can analyze and perform security analytics upon; develop tools to reliably and automatically extract features and build models for security applications by combining techniques from machine learning and program analysis; improve users' experiences with security analytics by translating the reasoning behind decisions into human understandable concepts; develop counter-measures to adversarial manipulation of metrics; and, quantify the accuracy of the resulting techniques in defending against adversarial manipulation.

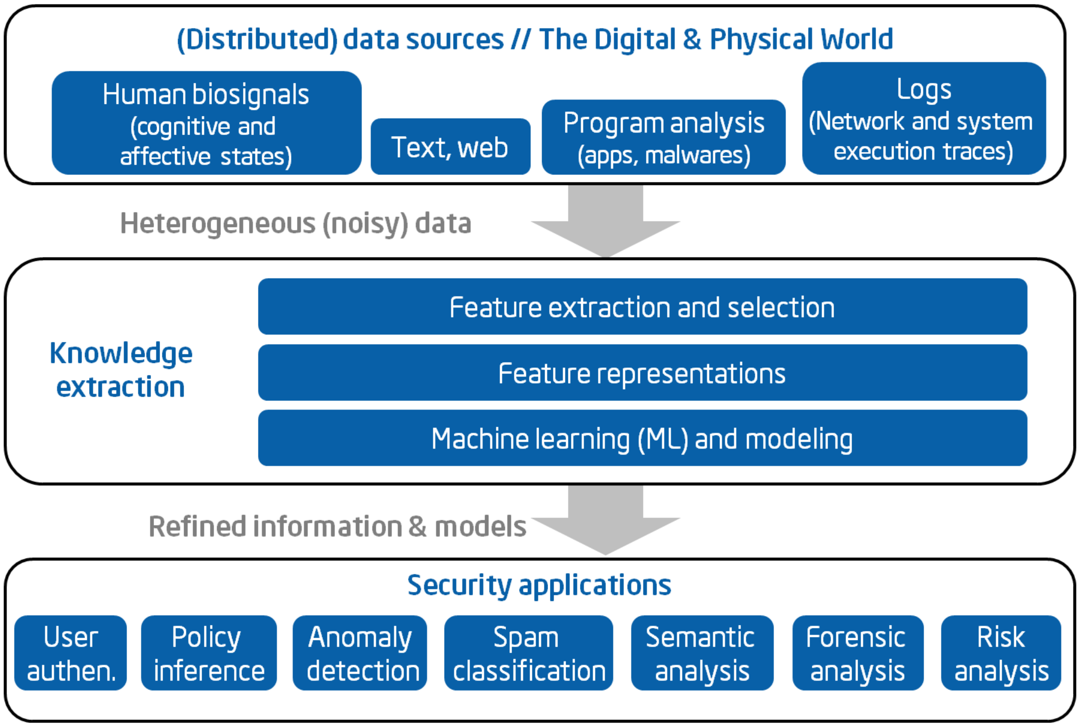
(Distributed) data sources // The Digital & Physical World
Human biosignals (cognitive and affective states)
Text, web
Program analysis (apps, malwares)
Logs (Network and system execution traces)
Heterogeneous (noisy) data
Knowledge extraction
Feature extraction and selection
Feature representations
Machine learning (ML) and modeling
Refined information & models
Security applications
User authen.
Policy inference
Anomaly detection
Spam classification
Semantic analysis
Forensic analysis
Risk analysis

Security Analytics using machine learning

The future promise

With the launch of the Intel Science and Technology Center for Secure Computing at UC Berkeley, we are preparing for the gathering "perfect storm" of societal dependence, ubiquitous connected personal devices, and truly global information services. The SCRUB program activities reflect our desire to serve users by making computing technology safer as it becomes a bigger and bigger part of our lives.

We recognize that malware, like germs and viruses, is a fact of life. While we cannot eliminate it (just as we cannot eliminate all pathogens), we can and will develop the insight and instrumentality to permit us to be healthy, productive, and safe despite continual assaults from invisible forces in the digital world.

INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL® PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER, AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF INTEL PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT. UNLESS OTHERWISE AGREED IN WRITING BY INTEL, THE INTEL PRODUCTS ARE NOT DESIGNED NOR INTENDED FOR ANY APPLICATION IN WHICH THE FAILURE OF THE INTEL PRODUCT COULD CREATE A SITUATION WHERE PERSONAL INJURY OR DEATH MAY OCCUR.

THE ORIGINAL EQUIPMENT MANUFACTURER MUST PROVIDE TPM FUNCTIONALITY, WHICH REQUIRES A TPM-SUPPORTED BIOS. TPM FUNCTIONALITY MUST BE INITIALIZED AND MAY NOT BE AVAILABLE IN ALL COUNTRIES.