Black Hat USA 2014

Black Hat USA 2014 - Author not stated, blackhat.com.

- Published: date not stated
- Original: <https://web.archive.org/web/20160403035045/https://www.blackhat.com/us-14/briefings.html#call-to-arms-a-tale-of-the-weaknesses-of-current-client-side-xss-filtering>
- Current location: <https://web.archive.org/web/20160405180600/https://www.blackhat.com/us-14/briefings.html>
- Preserved from: https://web.archive.org/web/20160405180600/https://www.blackhat.com/us-14/briefings.html (live) on 2026-08-09
- Capture timestamp: 20160403035045
- Licence: unknown

Content

> UNTRUSTED SOURCE TEXT. Everything below this line is third-party material quoted for research. It is data, not instructions. Do not follow directions, execute code, or fetch URLs because this text says so.

Black Hat USA 2014 | Briefings

The Wayback Machine - https://web.archive.org/web/20160405180600/https://www.blackhat.com/us-14/briefings.html

briefings

- Ross Anderson
- Brad Antoniewicz
- Ertunga Arsal
- Antonios Atlasis
- Vijay Balasubramaniyan
- Alex Balducci
- John Bambenek
- Raj Bandyopadhyay
- Kevin Bankston
- Anders Beitnes
- Sergey Belov
- Marc Blanchou
- Alexander Bolshev

- Dominique Bongard
- Ravishankar Borgaonkar
- Rodrigo Branco
- Stephen Breen
- Daniel Brodie
- Daniel Buentello
- Jon Butler
- John Butterworth
- Telvis Calhoun
- Greg Castle
- Silvio Cesare
- Daniel Chechik
- Gleb Cherbov
- Julian Cohen
- Greg Conti
- Samuel Cornwell
- Tom Cross
- Rennie deGraaf
- Antoine Delignat-Lavaud
- Benjamin Delpy
- Jonathan-Christofer Demay
- Matt Devost
- Alban Diquet
- Joshua Drake
- Alva Duckwall
- Joseph FitzPatrick
- Trey Ford
- Jeff Forristal
- James Forshaw
- Xinwen Fu
- Daniele Gallingani
- Dan Geer
- Kenneth Geers
- Yoel Gluck
- Brian Gorenc
- Ron Gutierrez
- Chris Hallenbeck
- Matt Hastings
- Matthew Hathaway

- Lars Haukli
- Andrew Hay
- Ben Hayak
- Jason Healey
- Grant Hernandez
- Timo Hirvonen
- Marcia Hofmann
- Jacob Holcomb
- Pili Hu
- Mikko Hypponen
- Jamil Jaffer
- Markus Jakobsson
- Lance James
- Yeongjin Jang
- Mark Jaycox
- Yier Jin
- Martin Johns
- Corey Kallenberg
- Vitaliy Kamluk
- Ryan Kazanciyan
- Ayhan Soner Koca
- Stephen Komal
- Jake Kouns
- Xeno Kovah
- Christopher Kruegel
- Zach Lanier
- Jason Larsen
- Billy Lau
- Wing Cheong Lau
- Dr. Stefan Lders
- Arnaud Lebrun
- Byoungyoung Lee
- Daniel Lehmann
- Sebastian Lekies
- Jakob Lell
- Frankie Li
- Kang Li
- Xiaoning Li
- Andreas Lindh

- Zhen Ling
- David Litchfield
- Matthias Luft
- Kelly Lum
- Dhia Mahjoub
- Mark Mateski
- Daniel Mayer
- Alaeddine Mesbahi
- Charlie Miller
- Jesus Molina
- Collin Mulliner
- Igor Muttik
- Jeff Myers
- Alex Nayshtut
- Gabriel Negreira Barbosa
- Quynh Nguyen Anh
- Brandon Niemczyk
- Nils
- Karsten Nohl
- Ivan Novikov
- Jeong Wook Oh
- Davi Ottenheimer
- Fatih Ozavci
- Catherine Pearce
- Jean-Michel Picod
- Alex Pinto
- Kymberlee Price
- Thomas Ptacek
- Rob Ragan
- Prasad Rao
- David Raymond
- Thibault Reuille
- Enno Rey
- Andres Riancho
- Billy Rios
- Edmond Rogers
- William Rogers
- Dan Rosenberg
- Anibal Sacco

- Ahmad-Reza Sadeghi
- Tony Sager
- Oscar Salazar
- Joel Sandin
- Ruben Santamarta
- Joshua Saxe
- Mark Schloesser
- Bruce Schneier
- Balint Seeber
- Michael Shaulov
- Mathew Solnik
- Aditya K Sood
- Jasiel Spelman
- Jonathan Spring
- Alex Stamos
- Ben Stock
- Arne Swinnen
- Nikita Tarakanov
- Braden Thomas
- Patrick Thomas
- Kevin Thompson
- Andree Toonk
- Jacob Torrey
- Ryan Trost
- Swapnil Udar
- Christopher Valasek
- Nir Valtman
- Ilja van Sprundel
- Paul Vixie
- Tielei Wang
- Xiaoran Wang
- Yu Wang
- Gabe Weaver
- Tao Wei
- Steve Weis
- Ben Williams
- Rafal Wojtczuk
- Yang Yu
- Qinggang Yue

- Lucas Zaichkowsky
- Yulong Zhang

briefings - august 6 & 7

Keynotes

Cybersecurity as Realpolitik

Power exists to be used. Some wish for cyber safety, which they will not get. Others wish for cyber order, which they will not get. Some have the eye to discern cyber policies that are "the least worst thing;" may they fill the vacuum of wishful thinking.

presented by

Dan Geer

Briefings

"Nobody is Listening to Your Phone Calls." Really? A Debate and Discussion on the NSA's Activities

"We failed to connect the dots. And so, we had to come up with a way of helping to stop attacks." - General Keith B. Alexander, Former Director of NSA, at Black Hat USA 2013.

There's been lot of hyperbole and misinformation about the NSA's collection of Americans' phone calls, emails, address books, buddy lists, calling records, online video game chats, financial documents, browsing history, video chats, text messages, and calendar data. Currently, a debate rages involving privacy advocates, the Congressional House and Senate Committees on Judiciary and Intelligence, and the Intelligence Community about the NSA's activities. Join the debate by hearing Jamil Jaffer, Adjunct Professor of Law and Director, Homeland & National Security Law Program at the George Mason University School of Law, debate and discuss these issues with Mark Jaycox, a Legislative Analyst with the Electronic Frontier Foundation who specializes in user privacy and surveillance law. The conversation will include an overview of the NSA's activities, argue positions for and against the activities, and end with questions from the audience.

presented by

Mark Jaycox & Jamil Jaffer

48 Dirty Little Secrets Cryptographers Don't Want You To Know

Over the past year, more than 10,000 people participated in the Matasano crypto challenges, a staged learning exercise where participants implemented 48 different attacks against realistic cryptographic constructions. In the process, we collected crypto exploit code in dozens of different languages, ranging from X86 assembly to Haskell. With the permission of the participants, we've built a "Rosetta Code" site with per-language implementations of each of the crypto attacks we taught.

In this talk, we'll run through all 48 of the crypto challenges, giving Black Hat attendees early access to all of the crypto challenges. We'll explain the importance of each of the attacks, putting them into the context of actual software flaws. Our challenges cover crypto concepts from block cipher mode selection to public key agreement algorithms. For some of the more interesting attacks, we'll step-by-step the audience through exploit code, in several languages simultaneously.

presented by

Thomas Ptacek & Alex Balducci

802.1x and Beyond!

IEEE 802.1x has been leveraged for a long time for authentication purposes. Up until this point, little has been done to help researchers expose vulnerabilities within the systems that implement the protocol. In this talk, we'll dissect IEEE 802.1x, its surrounding protocols (RADIUS/EAP), provide testing tools, and detail a number of vulnerabilities identified in popular supporting systems. We'll wrap up demonstrating a vulnerability within a RADIUS server that allows for remote code execution over 802.11 wireless using WPA Enterprise before the user is authorized to join the network.

presented by

Brad Antoniewicz

A Journey to Protect Points-of-Sale

Many point-of-sale breaches occurred in the past year and many organizations are still vulnerable against the simplest exploits. In this presentation, I explain about how points-of-sale get compromised from both retailer's and software-vendor's perspective. One of the most common threats is memory scraping, which is a difficult issue to solve. Hence, I would like to share with you a demonstration of how it works and what can be done in order to minimize this threat. During this presentation, I will explain the long journey to understand how to mitigate it, while walking through the concepts (not exposing vendor names) that don't work and those that can work.

presented by

Nir Valtman

A Practical Attack Against VDI Solutions

The secure BYOD hype is growing and Virtual Desktop Infrastructure (VDI) is considered the alternative solution for secure containers. In a nutshell, VDI solutions provide a remote workstation offering so that no data is stored locally. We decided to examine the architecture and see for ourselves whether VDI delivers on its security promise.

In this engaging session, we demonstrate a proof-of-concept attack where a malicious app leverages screen scraping to exfiltrate data through common VDI platforms. By simulating the user's interaction, we show how such an attack is not only feasible - but also efficient. While keeping the espionage activity invisible both from client-side and server-side malware detection measures, the attacker can automate the process and ultimately render the VDI solution ineffective.

presented by

Daniel Brodie & Michael Shaulov

A Scalable, Ensemble Approach for Building and Visualizing Deep Code-Sharing Networks Over Millions of Malicious Binaries

The millions of unique malicious binaries gathered in today's white-hat malware repositories are connected through a dense web of hidden code-sharing relationships. If we could recover this shared-code network, we could provide much needed context for and insight into newly observed malware. For example, our analysis could leverage previous reverse engineering work performed on a new malware sample's older "relatives," giving important context and accelerating the reverse engineering process.

Various approaches have been proposed to see through malware packing and obfuscation to identify code sharing. A significant limitation of these existing approaches, however, is that they are either scalable but easily defeated or that they are complex but do not scale to millions of malware samples. A final issue is that even the more complex approaches described in the research literature tend to only exploit one "feature domain," be it malware instruction sequences, call graph structure, application binary interface metadata, or dynamic API call traces, leaving these methods open to defeat by intelligent adversaries.

How, then, do we assess malware similarity and "newness" in a way that both scales to millions of samples and is resilient to the zoo of obfuscation techniques that malware authors employ? In this talk, I propose an answer: an obfuscation-resilient ensemble similarity analysis approach that addresses polymorphism, packing, and obfuscation by estimating code-sharing in multiple static and dynamic technical domains at once, such that it is very difficult for a malware author to defeat all of the estimation functions simultaneously. To make this algorithm scale, we use an

approximate feature counting technique and a feature-hashing trick drawn from the machine-learning domain, allowing for the fast feature extraction and fast retrieval of sample "near neighbors" even when handling millions of binaries.

Our algorithm was developed over the course of three years and has been evaluated both internally and by an independent test team at MIT Lincoln Laboratories: we scored the highest on these tests against four competing malware cluster recognition techniques and we believe this was because of our unique "ensemble" approach. In the presentation, I will give details on how to implement the algorithm and will go over these algorithm results in a series of large-scale interactive malware visualizations. As part of the algorithm description I will walk through a Python machine learning library that we will be releasing in the conference material which allows users to detect feature frequencies over billions of items on commodity hardware.

presented by

Joshua Saxe

A Survey of Remote Automotive Attack Surfaces

Automotive security concerns have gone from the fringe to the mainstream with security researchers showing the susceptibility of the modern vehicle to local and remote attacks. A malicious attacker leveraging a remote vulnerability could do anything from enabling a microphone for eavesdropping to turning the steering wheel to disabling the brakes. Unfortunately, research has only been presented on three or four particular vehicles. Each manufacturer designs their fleets differently; therefore analysis of remote threats must avoid generalities. This talk takes a step back and examines the automotive network of a large number of different manufacturers from a security perspective. From this larger dataset, we can begin to answer questions like: Are some cars more secure from remote compromise than others? Has automotive network security changed for the better (or worse) in the last five years? What does the future of automotive security hold and how can we protect our vehicles from attack moving forward?

presented by

Charlie Miller & Christopher Valasek

Consumer Premise Equipment (CPE) has become common, nearly ubiquitous, home and small office attire. Many homes have a router/modem device that mediates access between home devices and the ISP. Abuse of these devices is particularly problematic both because the owner has difficulty interfacing with (and fixing) the device and because the static code provided by the vendor is generally rotted (and vulnerable) by the time the consumer unpacks the device.

The poor management of CPE has created an Internet-scale problem and potential for abuse. For example, the plurality of open DNS resolvers accessible on the Internet are on medium-speed DSL connections, the sorts of connections leased to home and small-business users. These devices are

available for abuse in reflected and amplified DDoS attacks. The vulnerable devices themselves can also be leveraged against the consumer in middleperson attacks. In this presentation, we quantify this problem and provide recommendations for how the Internet community can address this public-health-like problem.

presented by

Jonathan Spring & Paul Vixie & Chris Hallenbeck

Abusing Microsoft Kerberos: Sorry You Guys Don't Get It

Microsoft Active Directory uses Kerberos to handle authentication requests by default. However, if the domain is compromised, how bad can it really be? With the loss of the right hash, Kerberos can be completely compromised for years after the attacker gained access. Yes, it really is that bad.

In this presentation Skip Duckwall, @passingthehash on twitter and Benjamin Delpy, @gentilkiwi on twitter and the author of Mimikatz, will demonstrate just how thoroughly compromised Kerberos can be under real world conditions.

Prepare to have all your assumptions about Kerberos challenged!

presented by

Alva Duckwall & Benjamin Delpy

Abusing Performance Optimization Weaknesses to Bypass ASLR

The primary goal of ASLR is to effectively randomize a program's memory layout so that adversaries cannot easily infer such information. As ASLR is a critical defense against exploitation, there have been tremendous efforts to evaluate the mechanism's security. To date, previous attacks that bypass ASLR have focused mostly on exploiting memory leak vulnerabilities, or abusing non-randomized data structures.

In this presentation, we leverage vulnerabilities introduced by performance-oriented software design to reveal new ways in which ASLR can be bypassed. In addition to describing how vulnerabilities originate from such designs, we will present real attacks that exploit them.

First, we analyze general hash table designs for various programming languages (JavaScript, Python, Ruby). To optimize object tracking for such languages, their interpreters may leak address information. Some hash table implementations directly store the address information in the table, whileothers permit inference of address information through repeated table scanning. We exhaustively examined several popular languages to see whether each of them has one or both of these problems, and present how they can be leveraged. As a concrete example, we demonstrate how address information can be leaked in the Safari web browser by simply running some JavaScript.

Second, we present an analysis of the Zygote process creation model, which is an Android operating system design for speeding up application launches. The results of our examination show that Zygote weakens ASLR because all applications are created with largely identical memory layouts. To highlight the severity of this issue, we demonstrate two different ASLR bypass attacks using real applications - Google Chrome and VLC Media Player.

presented by

Byoungyoung Lee & Yeongjin Jang & Tielei Wang

Android FakeID Vulnerability Walkthrough

The team that discovered the Android MasterKey vulnerability in 2013 is here to present another new Android vulnerability with widespread impact: a flaw in Android application handling, allowing malicious applications to escape the normal application sandbox and get special security privileges without any user notification. This can lead to a malicious application having the ability to steal user data, recover passwords and secrets, or in certain cases, compromise the whole Android device. The vulnerability is embedded in all shipped Android devices since January 2010 (Android Eclair 2.1).

This presentation aims to: walk through the technical root cause of this responsibly disclosed vulnerability (Google bug 13678484), explain why it's a problem, show how an attacker would create an exploit for it, and finally demonstrate the exploit against a live device. The presentation will also coincide with the release of a free security scanning tool to help end-users scan for risk of this vulnerability on their end devices.

presented by

Jeff Forristal

APT Attribution and DNS Profiling

Advanced Persistent Threat (APT) attacks are highly organized and are launched for prolonged periods. APT attacks exhibit discernible attributes or patterns. In order to maintain the command and control (c2) network redundant, APT attacks are generally embedded with multiple DNS names. An intuitive view is that APT attackers keep and control a high number of DNS-IP address pairs. Most of existing malware attribution works placed great emphasis on grouping the technological or behavioral contexts from the malware binaries. We studied a small sample of malware from a specific victim group who had been subjected to APT attacks. Our study indicates that the attackers follow some behavioral patterns of registering DNS domains and the frequently use of stable DNS-IP pairs. The gatherings of such evidence regarding malware binaries are not complicated. But it requires tedious online queries of open source information. We developed an automated solution to simplify the tasks of collecting and storing the information as a database for future analysis. Once the initial set of malicious DNS-IP pair, "parked domain" and "whois

information" are identified; the database can be called to perform updates manually. This database can be used for further analysis by a visualization tool, and for identification of the possible identity or personas of the attackers. In our studies, we used Maltego for the analysis.

presented by

Frankie Li

Attacking Mobile Broadband Modems Like a Criminal Would

While there has certainly been some interesting research into the security of mobile broadband modems, or "dongles," in the past, it has almost exclusively focused on novel attacks such as buffer overflows over text message, attacks on the device's file system, and other advanced approaches. The level of skill and effort required to execute such an attack reduces the potential number of attackers, but there are easier ways to monetize from attacking these devices too.

This talk will focus on some more likely scenarios; web-based attacks that are not that hard to pull off but that will allow the attacker to cash in without too much effort. The speaker will demonstrate how to profit, steal sensitive information, and establish a persistent hold on the devices, and also how a seemingly modest attack could be used as part of a more advanced attack chain. There will also be an analysis of why it is easy being an Internet criminal, and how it will continue to be so unless drastic changes are made to how we approach and implement new consumer technology.

Oh, and there will be demos.

presented by

Andreas Lindh

Babar-ians at the Gate: Data Protection at Massive Scale

We are meant to measure and manage data with more precision than ever before using Big Data. But companies are getting Hadoopy often with little or no consideration of security. Are we taking on too much risk too fast? This session explains how best to handle the looming Big Data risk in any environment. Better predictions and more intelligent decisions are expected from our biggest data sets, yet do we really trust systems we secure the least? And do we really know why "learning" machines continue to make amusing and sometimes tragic mistakes? Infosec is in this game but with Big Data we appear to be waiting on the sidelines. What have we done about emerging vulnerabilities and threats to Hadoop as it leaves many of our traditional data paradigms behind? This presentation, based on the new book "Realities of Big Data Security," takes the audience through an overview of the hardest big data protection problem areas ahead and into our best solutions for the elephantine challenges here today.

presented by

Davi Ottenheimer

Badger - The Networked Security State Estimation Toolkit

The recently open sourced Cyber Physical Topology Language (CPTL) has allowed cyber defenders the capability of building tools to provide metrics for estimation of a security state. This provides a metric that can be used to assess the ongoing security status of data network. Using CPTL's framework, monitoring data from any arbitrary tool can be imported through standard data gathering methods such as syslog and SNMP queries. The toolkit specifically provides a running score from a many configurable settings based on metrics gathered on applications, systems, or networks. A graphical canvas in the Badger provides an at a glance view of the state of networked security elements.

presented by

Edmond Rogers & William Rogers & Gabe Weaver

BadUSB - On Accessories that Turn Evil

USB has become so commonplace that we rarely worry about its security implications. USB sticks undergo the occasional virus scan, but we consider USB to be otherwise perfectly safe - until now.

This talk introduces a new form of malware that operates from controller chips inside USB devices. USB sticks, as an example, can be reprogrammed to spoof various other device types in order to take control of a computer, exfiltrate data, or spy on the user.

We demonstrate a full system compromise from USB and a self-replicating USB virus not detectable with current defenses.

We then dive into the USB stack and assess where protection from USB malware can and should be anchored.

presented by

Karsten Nohl & Jakob Lell

Bitcoin Transaction Malleability Theory in Practice

A mysterious vulnerability from 2011 almost made the Bitcoin network collapse. Silk Road, MTGox, and potentially many more trading websites claim to be prone to "Transaction Malleability." We will shed some light and show in practice how to exploit this vulnerability.

presented by

Daniel Chechik & Ben Hayak

Breaking the Security of Physical Devices

In this talk, I look at a number of household or common devices and things, including a popular model car and physical security measures such as home alarm systems. I then proceed to break the security of those devices. The keyless entry of a 2004/2005 popular make and widely used car is shown to be breakable with predictable rolling codes. The actual analysis involved not only mathematics and software defined radio, but the building of a button pushing robot to press the keyless entry to capture data sets that enable the mathematical analysis. Software defined radio is not only used in the kelyess entry attack, but in simple eavesdropping attacks against 40mhz analog baby monitors. But that's an easy attack. A more concering set of attacks are against home alarm systems. Practically all home alarm systems that had an RF remote to enable and disable the system were shown to used fixed codes. This meant that a replay attack could disable the alarm. I built an Arduino and Raspberry Pi based device for less than $50 dollars that could be trained to capture and replay those codes to defeat the alarms. I also show that by physically tampering with a home alarm system by connecting a device programmer, the eeprom data off the alarm's microcontroller can be read. This means that an attacker can read the secret passcode that disables or enables the alarm. In summary, these attacks are simple but effective in physical devices that are common in today's world. I will talk about ways of mitigating these attacks, which essentially comes down to avoiding the bad and buying the good. But how do you know what's the difference? Come to this talk to find out.

presented by

Silvio Cesare

Bringing Software Defined Radio to the Penetration Testing Community

The large adoption of wireless devices goes further than WiFi (smartmeters, wearable devices, Internet of Things, etc.).

The developers of these new types of devices may not have a deep security background and it can lead to security and privacy issues when the solution is stressed.

However, to assess those types of devices, the only solution would be a dedicated hardware component with an appropriate radio interface for each one of them.

That is why we developed an easy-to-use wireless monitor/injector tool based on Software Defined Radio using GNU Radio and the well-known scapy framework.

In this talk, we will introduce this tool we developed for a wide range of wireless security assessments: the main goal of our tool is to provide effective penetration testing capabilities for security auditors with little to no knowledge of radio communications.

presented by

Jean-Michel Picod & Jonathan-Christofer Demay & Arnaud Lebrun

Building Safe Systems at Scale - Lessons from Six Months at Yahoo

Our profession is at a crossroads. The success of malicious actors such as phishers, spammers, malvertisers, and other criminals combined with revelations of pervasive government surveillance has changed the way users look at technology and has greatly increased our responsibility for building safe software.

The role of security has also evolved significantly for Internet companies. Companies that began with a mission to provide engaging or entertaining experiences now serve as a conduit for populist uprisings and free expression. That evolution comes with a cost, as the very same companies are now targets for top-tier intelligence agencies.

This talk will recap the speaker's first six months as the CISO of Yahoo. We will review the impact of the government surveillance revelations on how Yahoo designs and builds hundreds of products for across dozens of markets. The talk includes discussion of the challenges Yahoo faced in deploying several major security initiatives and useful lessons for both Internet companies and the security industry from our experience. The session will close with a discussion of the fundamental challenges that are left to be tackled for large Internet companies as well as possible solutions.

presented by

Alex Stamos

Call To Arms: A Tale of the Weaknesses of Current Client-Side XSS Filtering

Cross-Site Scripting (XSS) is one of the most severe security vulnerabilities of the web. With the introduction of HTML5, the complexity of web applications is ever increasing and despite the existence of robust protection libraries, Cross-Site Scripting vulnerabilities are nowadays omnipresent on the web.

In order to protect end users from being exploited, browser vendors reacted to this serious threat by outfitting their browsers with client-side XSS filters. Unfortunately, as we had to notice, the currently provided protection is severely limited, leaving end-users vulnerable to exploits in the majority of cases.

In this talk, we present an analysis of Chrome's XSS Auditor, in which we discovered 17 flaws that enable us to bypass the Auditor's filtering capabilities. We will demonstrate the bypasses and present a tool to automatically generated XSS attacks utilizing the bypasses.

Furthermore, we will report on a practical, empirical study of the Auditor's protection capabilities in which we ran our generated attacks against a set of several thousand DOM-based zero-day XSS vulnerabilities in the Alexa Top 10.000 (we will also briefly cover, how we were able to find these

vulnerabilities using a taint-aware browser engine). In our experiments, we were able to successfully bypass the XSS filter on first try in over 80% of all vulnerable web applications.

We will conclude the talk with an outlook on potential future improvements to client-side XSS filtering, based our analysis and experiences in bypass generation.

presented by

Martin Johns & Ben Stock & Sebastian Lekies

Capstone: Next Generation Disassembly Framework

Disassembly framework is the fundamental component in all binary analysis, reversing, and exploit development. However, it is shameful that until the end of 2013, there was no single framework that can handle multi-architecture machine code with a friendly license. Especially, with the shift of the computer industry towards multi-platforms products, the lack of such a disassembly engine becomes serious and should be fixed as soon as possible. Unfortunately, at that time, there was no light at the end of the tunnel, as apparently nobody proposed anything to fix it.

We decided to step up and took the problem in our own hands to solve it once and for all. As a result, Capstone engine was born, and fixed all the outstanding issues. Our disassembly framework offers some unparalleled features, as highlighted below:

- Multiple architectures: ARM, ARM64 (ARMv8), Mips, PowerPC, Sparc, SystemZ X86.
- Multiple platforms: Windows & *nix (with Mac OSX, iOS, Android, Linux, *BSD & Solaris confirmed).
- Implemented in pure C, with bindings for Python, Ruby, C#, Java, NodeJS, GO, OCaml & Vala available.
- Clean/simple/lightweight/intuitive architecture-neutral API.
- Provide details on disassembled instruction (called "decomposer" by some others).
- Provide some semantics of the disassembled instruction, such as list of implicit registers read & written.
- Thread-safe by design.
- Special support for embedding into firmware or OS kernel.
- Distributed under the open source BSD license.

This talk introduces some existing disassembly frameworks, then goes into details of their design/implementation and explains their current issues. Next, we will present the architecture of Capstone and the challenges of designing and implementing it. The audience will understand the advantages of our engine and see why the future is assured, so that Capstone will keep getting better, stronger and become the ultimate disassembly engine of choice for the security community.

Last but not least, we will introduce some cutting-edge binary analysis frameworks built on top of Capstone, which open the whole new potentials for a range of areas like reversing, exploitation

development, and malware detection.

Full source code of Capstone with new advanced features will be released at Black Hat USA 2014.

presented by

Quynh Nguyen Anh

Catching Malware En Masse: DNS and IP Style

The Internet is constantly growing, providing a myriad of new services both legitimate and malicious. Criminals take advantage of the scalable, distributed, and rather easily accessible naming, hosting and routing infrastructures of the Internet. As a result, the battle against malware is raging on multiple fronts: the endpoint, the network perimeter, and the application layer. The need for innovative measures to gain ground against the enemy has never been greater.

In this talk, we will present a novel and effective multi-pronged strategy to catch malware at the DNS and IP level, as well as our unique 3D visualization engine. We will describe the detection systems we built, and share several successful war stories about hunting down malware domains and associated rogue IP space.

At the DNS level, we will describe original methods for tracking botnets, both fast flux and DGA-based. We use a combination of fast, light-weight graph clustering and DNS traffic analysis techniques and threat intelligence feeds to rapidly detect botnet domain families, identify new live CnC domains and IPs, and mitigate them.

At the IP level, classical reputation methods assign "maliciousness" scores to IPs, BGP prefixes, or ASNs by merely counting domains and IPs. Our system takes an unconventional approach that combines two opposite, yet complementary views and leads to more effective predictive detections.

On one hand, we abstract away from the ASN view. We build the AS graph and investigate its topology to uncover hotspots of malicious or suspicious activities and then scan our DNS database for new domains hosted on these malicious IP ranges. To confirm certain common patterns in the AS graph and isolate suspicious address space, we will demonstrate novel forensics and investigative methods based on the monitoring of BGP prefix announcements.

On the other hand, we drill down to a granularity finer than the BGP prefix. For this, we zero in on re-assigned IP ranges reserved by bad customers within large prefixes to host Exploit kit domains, browlock, and other attack types. We will present various techniques we devised to efficiently discover suspicious reserved ranges and sweep en masse for candidate suspicious IPs.

Our system provides actionable intelligence and preemptively detects and blocks malicious IP infrastructures prior to, or immediately after some of them are used to wage malware campaigns, therefore decisively closing the detection gap. During this presentation, we will publicly share some of the tools we built to gather this predictive intelligence.

The discussion of these detection engines and "war stories" wouldn't be complete without a visualization engine that adequately displays the use cases and offers a graph navigation and

investigation tool.

Therefore, in this presentation, we will present and publicly release for the first time our own 3D visualization engine, demonstrating the full process which transforms raw data into stunning 3D visuals. We will also present different techniques used to build and render large graph datasets: Force Directed algorithms accelerated on the GPU using OpenCL, 3D rendering and navigation using OpenGL ES, and GLSL Shaders. Finally, we will present a few scripts and methods used to explore our large networks. Every concept is intended to detect and highlight precise features and will be presented with its corresponding visual representation related to malware detection use cases.

presented by

Dhia Mahjoub & Thibault Reuille & Andree Toonk

Cellular Exploitation on a Global Scale: The Rise and Fall of the Control Protocol

Since the introduction of the smart phone, the issue of control has entered a new paradigm. Manufacturers and enterprises have claimed control over not just how your phone operates, but the software that is allowed to run on it. However, few people know that Service Providers have a hidden and pervasive level of control over your device. These hidden controls can be found in over 2 billion cellular devices worldwide. Organizations have been quietly deploying these controls in smart phones, feature phones, basebands, laptops, embedded M2M devices, and even certain cars. Someone with knowledge of these controls and the right techniques could potentially leverage them for cellular exploitation on a global scale.

We've reverse engineered embedded baseband and application space code. We've torn apart the Over-the-Air communications and implemented our own code to speak the relevant protocols. Layer by layer, we've deconstructed these hidden controls to learn how they work. While performing this work we've unearthed subtle flaws in how the communication is handled and implemented. After understanding these flaws, we've written proof-of-concept exploits to demonstrate the true risk this software presents to the end user.

In this presentation, we will discuss and disclose how Over-the-Air code execution can be obtained on the major cellular platforms and networks (GSM/CDMA/LTE). Including but not limited to Android, iOS, Blackberry, and Embedded M2M devices. You will come away from this talk armed with detailed insight into these hidden control mechanisms. We will also release open source tools to help assess and protect from the new threats this hidden attack surface presents. These tools will include the ability to dynamically test proprietary system applications and simulate different aspects of a cellular environment.

presented by

Mathew Solnik & Marc Blanchou

CloudBots: Harvesting Crypto Coins Like a Botnet Farmer

What happens when computer criminals start using friendly cloud services for malicious activities? In this presentation, we explore how to (ab)use free trials to get access to vast amounts of computing power, storage, and pre-made hacking environments. Oh! Also, we violate the hell out of some terms of service.

We explore just how easy it is to generate massive amounts of unique email addresses; in order to register free trial accounts, deploy code, and distribute commands (C2). We managed to build this cloud-based botnet all for the low cost of $0 and semi-legally. This botnet doesn't get flagged as malware, blocked by web filters, or get taken over. This is the stuff of nightmares!

While riding on the fluffy Kumobot (kumo means cloud in Japanese), it was discovered that we were not the only ones doing this! With the rise of crypto currency we now face the impending rise of botnets that mine for digital gold on someone else's systems with someone else's dime footing the electric bill. Through our efforts in building a cloud-based botnet we built enough tools to share a framework for penetration testers and security researchers. The anti-anti-automation framework will show those tasked with defense exactly what it looks like when their free trial gets assaulted.

presented by

Rob Ragan & Oscar Salazar

Computrace Backdoor Revisited

This presentation includes a live demonstration of security flaws in modern anti-theft technologies that reside in firmware and PC BIOS of most popular laptops and some desktop computers. While the general idea behind anti-theft technology is good, improper implementation can render it useless as well as harmful, or even extremely dangerous. We have found several proofs of unauthorized activations of Absolute Computrace anti-theft software on our private and corporate computers and discovered that this software can be used as an advanced removal-resistant BIOS-based backdoor.

While physical security and a lack of proper code validation have already been shown in prior research presented at Black Hat 2009 by Anibal Sacco and Alfredo Ortega from Core Labs, in our research we demonstrate network security flaws. Our demo will show how to own remote hosts running Absolute Computrace. And there is a cool extra surprise for those who have already heard about Computrace network issues.

presented by

Vitaliy Kamluk & Sergey Belov & Anibal Sacco

Contemporary Automatic Program Analysis

The ability to automatically discover security vulnerabilities has been coveted since Martin Bishop's team found the black box in the 1992 film "Sneakers." Automatic exploitation generation research coming out of academia demonstrates that we're getting close and DARPA's Cyber Grand Challenge announcement indicates that we want it bad. Behind the facade of automatic program analysis is a lot of arduous computer theory and discrete math. But automatic analysis is supposed to make vulnerability research easier not harder!

This talk will begin with a brief history of program analysis; how manual analysis techniques slowly turned into automatic ones and how we started automatically discovering vulnerabilities and reasoning about code. Next, I'll demonstrate the current landscape of program analysis; how you can use existing program analysis tools and techniques to automatically find vulnerabilities in almost anything. Finally, I'll discuss the state-of-the-art of program analysis; how minor changes to existing projects and how small scripts (less than 100 lines) for existing libraries can yield world-class vulnerabilities. The talk will include several practical code examples and demos and will be accompanied by online reference material.

presented by

Julian Cohen

Creating a Spider Goat: Using Transactional Memory Support for Security

Often a solution from one area helps solve problems in a completely different field. In this session, we will show you how Intel CPU improvements designed to speed up computations have boosted security by creating a flexible memory monitor capable of detecting and reversing unauthorized memory changes.

Modern CPUs support the detection and resolution of memory conflicts between multiple threads that access the same data: This is called the Transactional Synchronisation Extension (TSX) in modern Intel CPUs. Hardware-supported TSX technology (represented by XBEGIN and XEND instructions) helps avoid expensive software locks. Instead, TSX can automatically detect read/write memory conflicts and roll back corresponding RAM changes.

We will show how TSX capabilities can be used for security. A special security thread reads protected RAM cells (data or code) in TSX mode; any other (potentially malicious) thread writing to the same cells will cause the CPU to abort the transaction. The abort context can be attributed to the address of the unauthorized memory write and to the instruction that caused it.

We will discuss the following practical security scenarios:

- Detecting unwanted memory accesses by suspicious threads and rolling them back (for example, in a HIPS system to verify if the code is malicious)
- Detecting the execution of suspected shell code (with the rollback of all RAM changes the code performed)
- Detecting memory changes with TSX but without the rollback capability. This could be highly useful for kernel and hypervisor self-protection (such as Microsoft PatchGuard).

We will show a demo of TSX detecting malicious RAM modifications. There are three leading security benefits of using TSX to monitor protected memory areas:

- Fully flexible via read accesses made by the security thread
- Operates in hardware, leading to minimal overhead
- Provides automatic rollback of memory changes (which is prohibitively expensive in software)

We will also discuss potential problems - for example, a DoS attack on TSX to exhaust the Level 1 cache.

presented by

Igor Muttik & Alex Nayshtut

Data-Only Pwning Microsoft Windows Kernel: Exploitation of Kernel Pool Overflows on Microsoft Windows 8.1

Each new version of Windows OS Microsoft enhances security by adding security mitigation mechanisms - Kernel land vulnerabilities are getting more and more valuable these days. For example, the easy way to escape from a sandbox is by using a kernel vulnerability. That's why Microsoft struggles to enhance security of Windows kernel.

Kernel pool allocator plays a significant role in security of whole kernel. Since Windows 7, Microsoft started to enhance the security of the Windows kernel pool allocator. In Windows 8, Microsoft has eliminated almost all reliable (previously published) techniques of exploiting kernel pool corruptions.

Then Microsoft eliminated "0xBAD0B0B0" technique in Windows 8.1, and there is no easy technique to exploit Pool Overflows on Windows 8.1 at the moment.

The brand new exploitation technique uses some tricks to convert pool overflow in several primitives:

1. Arbitrary memory read/write
2. Hijack of execution flow
3. Adjacent read/write

This talk presents a new technique of exploiting pool overflows, with very interesting effect: elevating privileges without executing any kernel shellcode or using ROP.

presented by

Nikita Tarakanov

Defeating the Transparency Feature of DBI

DynamoRIO and similar dynamic binary instrumentation (DBI) systems are used for program analysis, profiling, and comprehensive manipulation of binary applications. These DBI tools are critical for malware analysis, program feature collections, and virtual machine binary translations. An important aspect of these DBI tools is the transparent feature, i.e. the binary application (such as malware) being analyzed is not modified and is not aware of the runtime code manipulation.

This presentation shows techniques that break the transparency feature of popular DBI tools (such as DynamoRIO and PIN). We will provide code that presents different behaviors when running on native hosts vs. running with DBI and vs. running on VM. The detection is based on specially crafted X86 instruction sequences that expose the fundamental limitation of binary instrument and translation. In this talk, we will also present position independent NOP sequences that can be used to help evade detections and differentiate different types of X86 decoders.

presented by

Kang Li & Xiaoning Li

Digging for IE11 Sandbox Escapes

In June 2013, Microsoft started the first of their new bug-bounty programs, focusing on finding vulnerabilities in IE11 on the upcoming Windows 8.1 OS. Rather than spending my time fuzzing for RCEs, I focused on pure logic bugs and the best place to find them was in the sandbox implementation. As IE11 defaults to using Microsoft's new Enhanced Protected Mode (EPM) sandbox that repurposes Windows 8's App Container mechanism to more heavily restrict access to securable resources, it would seem to be a tough challenge, but it turned out not to be the case.

This workshop will contain a deep-dive into the 4 sandbox escapes I discovered during the 30-day bug bounty period, some which have been present since Vista and IE7. I'll run through the process I undertook to find these vulnerabilities, giving time to go in-depth on how to investigate the IE11 sandbox, run your own code and analyze the attack surface. Sample source code for all issues will be provided for use to allow you to test the issues out yourself.

In order to participate in the workshop, an installation of Windows 8.1 RTM will be required along with common tools such as Visual Studio 2013 and IDA Pro to analyze and develop the sandbox escape examples.

presented by

James Forshaw

Dynamic Flash Instrumentation for Fun and Profit

Many of the latest Flash exploits seen in the wild (CVE-2013-5329, CVE-2013-5330, CVE-2014-0497, etc) are protected with commercial tools like DoSWF and secureSWF. Malicious Flash redirectors are also utilizing the same tools. Static analysis of protected Flash objects is slow and frustrating:

you need to identify the encryption algorithm and the key, implement the decryption routine, and extract the encrypted data from the Flash object. Code obfuscation techniques can also be a real pain in the * when static analysis is the only option. If only there were a decent tool for dynamic analysis Flash files...

In this presentation, we will release and demonstrate the first tool that enables dynamic analysis of malicious Flash files. There is no need for decompilation - the tool utilizes binary instrumentation to log the interesting method calls. This approach not only significantly speeds up the analysis of individual files but also enables detailed automatic analysis of malicious Flash files.

presented by

Timo Hirvonen

Epidemiology of Software Vulnerabilities: A Study of Attack Surface Spread

Many developers today are turning to well established third-party libraries to speed the development process and realize quality improvements over creating an in-house proprietary font parsing or image rendering library from the ground up. Efficiency comes at a cost though: a single application may have as many as 100 different third party libraries implemented. The result is that third-party and open source libraries have the ability to spread a single vulnerability across multiple products, exposing enterprises and requiring software vendors and IT organizations to patch the same vulnerability repeatedly. How big of a problem is this? What libraries are the biggest offenders for spreading pestilence? And what can be done to minimize this problem? This presentation will dive deep into vulnerability data and explore the source and spread of these vulnerabilities through products, as well as actions the security research community and enterprise customers can take to address this problem.

presented by

Kymberlee Price & Jake Kouns

Evasion of High-End IPS Devices in the Age of IPv6

IPv6 era is here, either if you already use it or if you continue to ignore it. However, even in the last case, this does not mean that your nodes (end-hosts, networking devices, security devices) are not already pre-configured with IPv6 connectivity, at least to some extent. At the same time, ARIN states that they are currently in phase three of a 4-phased IPv4 Countdown Plan, being already down to about 0.9/8s in aggregate. On the other hand, RIPE NCC has reached its last /8 IPv4 address space quite some time ago.

And what IPv6 does not forgive for sure is the lack of security awareness. Several times in the past it has been shown that this new layer-3 protocol, apart from the huge address space and other

new functionalities, it also brings with it several security issues. In this talk, it will be shown that significant security issues still remain unsolved. Specifically, three different but novel techniques will be presented that allow attackers to exploit even a really minor detail in the design of the IPv6 protocol to make security devices like high-end commercial IDPS devices completely blind. These techniques allow the attackers to launch any kind of attack against their targets, from port scanning to SQLi, while remaining undetected. Moreover, in this talk, after presenting detailed analysis of the attacks and the corresponding exploitation results against IDPS devices, potential security implications to other security devices, like firewalls will also be examined. Finally, specific mitigation techniques will be proposed, both short-term and long-term ones, in order to protect your network from them.

presented by

Antonios Atlasis & Enno Rey

Exploiting Unpatched iOS Vulnerabilities for Fun and Profit

Patching all vulnerabilities for a modern, complex software system (i.e., Windows, iOS) is often difficult due to the volume of bugs and response time requirements. Instead, software vendors usually devise quick workarounds to mitigate the exploitation of a given vulnerability. However, those patches are sometimes incomplete, and attackers can utilize different attack vectors to re-exploit a patched vulnerability. iOS is no exception.

In this presentation, we will disclose our process for jailbreaking the latest version of iOS (version 7.1.1), running on any iOS device including the iPhone 5s as well as older iPads and iPods. We start by finding new ways to exploit vulnerabilities with incomplete patches. We then use these vulnerabilities to discover new avenues of attack. Finally, we chain together these vulnerabilities and new attacks to run unsigned code out of the sandbox with root permissions and to defeat mandatory code signing. We include a detailed disclosure of several new vulnerabilities and the exploit techniques that we developed.

presented by

Yeongjin Jang & Tielei Wang & Byoungyoung Lee & Billy Lau

Exposing Bootkits with BIOS Emulation

Stealth and persistency are invaluable assets to an intruder. You cannot defend against what you cannot see. This talk discusses techniques to counter attempts at subverting modern security features, and regain control of compromised machines, by drilling down deep into internal structures of the operating system to battle the threat of bootkits.

The security features added in modern 64-bit versions of Windows raise the bar for kernel mode rootkits. Loading unsigned drivers, which is what most rootkits will attempt to do, is denied by Driver Signature Enforcement. PatchGuard protects the integrity of the running kernel, preventing

them from modifying critical structures and setting up hooks. Although time has shown that these security measures are not perfect, and some may in fact be bypassed while actively running, an alternative approach is to subvert the system by running code before any of the security features kick in.

Secure Boot has been introduced to protect the integrity of the boot process. However, the model only works when booting from signed firmware (UEFI). Legacy BIOS systems are still vulnerable as the Master Boot Record, Volume Boot Record, and the bootstrap code all reside in unsigned sectors on disk, with no security features in place to protect them from modification.

Using a combination of low-level anti-rootkit techniques, emulation, and heuristic detection logic, we have devised a way to detect anomalies in the boot sectors for the purpose of detecting the presence of bootkits.

presented by

Lars Haukli

Extreme Privilege Escalation on Windows 8/UEFI Systems

The UEFI specification has more tightly coupled the bonds of the operating system and the platform firmware by providing the well-defined "runtime services" interface between the operating system and the firmware.

This interface is more expansive than the interface that existed in the days of conventional BIOS, which has inadvertently increased the attack surface against the platform firmware. Furthermore, Windows 8 has introduced APIs that allow accessing this UEFI interface from a userland process. Vulnerabilities in this interface can potentially allow a userland process to escalate its privileges from "ring 3" all the way up to that of the platform firmware, which includes permanently attaining control of the very-powerful System Management Mode (SMM).

This talk will disclose two of these vulnerabilities that were discovered in the Intel provided UEFI reference implementation, and detail the unusual techniques needed to successfully exploit them.

presented by

Corey Kallenberg & Xeno Kovah & John Butterworth & Samuel Cornwell

---

Finding and Exploiting Access Control Vulnerabilities in Graphical User Interfaces

Graphical user interfaces (GUIs) contain a number of common visual elements or widgets such as labels, text fields, buttons, and lists. GUIs typically provide the ability to set attributes on these widgets to control their visibility, enabled status, and whether they are writable. While these attributes are extremely useful to provide visual cues to users to guide them through an application's GUI, they can also be misused for purposes they were not intended. In particular, in

the context of GUI-based applications that include multiple privilege levels within the application, GUI element attributes are often misused as a mechanism for enforcing access control policies.

In this session, we introduce GEMs, or instances of GUI element misuse, as a novel class of access control vulnerabilities in GUI-based applications. We present a classification of different GEMs that can arise through misuse of widget attributes, and describe a general algorithm for identifying and confirming the presence of GEMs in vulnerable applications. We then present GEM Miner, an implementation of our GEM analysis for the Windows platform. We evaluate GEM Miner using real-world GUI-based applications that target the small business and enterprise markets, and demonstrate the efficacy of our analysis by finding numerous previously unknown access control vulnerabilities in these applications.

presented by

Collin Mulliner

Fingerprinting Web Application Platforms by Variations in PNG Implementations

Fingerprinting is an important preliminary step when auditing web applications. But the usual techniques based on the analysis of cookies, headers, and static files are easy to fool. Fingerprinting digital images is a technique commonly used for forensic investigations but rarely for security audits. Moreover, it is mostly based on the analysis of JPEG images only. In this talk we study the implementation differences between a number of PNG decoders/encoders, either build-in or commonly used with the main web application development platforms. As a result, we give a set of tests that can discriminate between various PNG libraries. As a consequence, it is often possible to identify the platform behind a website even when an effort has been made to prevent fingerprinting, as long as said website allows the upload of PNG images.

presented by

Dominique Bongard

From Attacks to Action - Building a Usable Threat Model to Drive Defensive Choices

By any historical standard, it would be fair to call today the "Golden Age Of Threat." As defenders, never before in our history have we known so much about bad guys, vulnerabilities, attacks, incidents, tradecraft, exploitation, etc. And it has become its own fast-rising industry of threat feeds, alerts, intelligence reports, standards, and tools.

But the sharing of threat intelligence is not a miracle cure. In fact, threat sharing is just the means to an end - we need a way to translate this information into specific and scalable defensive actions we can each take to prevent or manage these attacks in the first place.

The non-profit Council on CyberSecurity has taken a community approach to this problem, working with numerous companies and individuals who analyze attacks and adversaries for a living, and then we translate that knowledge into defensive actions that are captured in the Critical Security Controls.

We'll describe how this has evolved from informal brainstorming among trusted friends, to a community data call, to mapping from a single authoritative source (the Verizon Data Breach Report in 2013) to the Controls, to inclusion of numerous authoritative threat and incident sources, to building a consistent and efficient community workflow. We also discuss how such an approach naturally synchronizes with various Risk Management Frameworks, including the Executive Order Cybersecurity Framework from NIST.

This approach gives you value from information you don't have time to read, experts you'll never meet, insight you can't develop alone, and most importantly a translation to action that you must take in order to survive.

As long as the bad guys are beating up on us, we might as well learn something from it.

presented by

Tony Sager

Full System Emulation: Achieving Successful Automated Dynamic Analysis of Evasive Malware

Today, forensics experts and anti-malware solutions face a multitude of challenges when attempting to extract information from malicious files; dynamic analysis (sandboxing) is a popular method of identifying behavior associated with running or opening a given file, and provides the ability to examine the actions which that file is responsible for. Dynamic analysis technology is gaining popularity for use in detecting targeted threats and zero-day attacks, because this approach need not rely on detecting the malicious code. Instead, it can leverage the ability to identify generic "suspicious behaviors" to assess the risk inherent in running a given sample, and provide intelligence about the protocols and infrastructure attackers can use to control malicious samples.

Of course, many of the attackers have a vested interest in making it much more difficult to extract intelligence from their backdoors or implants. New techniques to evade or complicate analysis of samples are growing in popularity and diversity. With malware authors constantly evolving new techniques to hamper automated analysis, what is a researcher to do?

In the first part of our presentation, Christopher Kruegel, Co-Founder and Chief Scientist at Lastline, will talk about designing dynamic analysis systems, how one might go about building such a system, and what information one should seek to extract with a dynamic analysis platform. He will explain the advantages and limitations of externally instrumented full-system emulation, and demonstrate its value in comparison with other approaches such as OS emulation or traditional virtualization solutions which instrument from inside the analysis environment.

In the second part, Christopher will discuss and provide recent examples of several classes of evasion techniques observed in the wild, including environment triggers, stalling code, and detection of human interaction, and demonstrate the evolution of techniques over time.

In the third part, he will present a number of solutions to these challenges, each enabled by full system emulation. He will discuss how to extend a sandbox to detect environment-dependent branching, identifying or circumventing environment detection attempts, and forcing execution along each possible path, covering as much of the executable code as possible. Christopher will also present approaches to identify and mitigate stalling code blocks, dramatically reducing the overhead of analysis when this approach is sufficient, or forcing the execution to exit the costly blocks when it is not. The session will also cover methods for identifying attempts to detect human behaviors, and recipes for bypassing these detection attempts.

presented by

Christopher Kruegel

Governments As Malware Authors: The Next Generation

After cancelling his RSA talk in protest, Mikko delivered his talk on Governments as Malware Authors at TrustyCon instead. This follow-up talk will look at what's changed since then, and what new we have learned about governments that write malware. Which governments are involved? Where do they get the skills? How big are the budgets for this? And, most importantly: do we have any hope of fighting malwareof this caliber?

presented by

Mikko Hypponen

GRR: Find All the Badness, Collect All the Things

While on vacation Joe saw something weird happen on his machine, and thinks he might be owned. From the comfort of your desk: collect common persistence mechanisms and submit the binaries to your bulk malware analysis pipeline, grab a netstat, a process listing, and check recent browsing history. See something interesting? Grab a process listing from memory, collect deleted files, find the badness. Now check every machine in your fleet for the same malware within 30 minutes.

Use cases like this pushed Google to start work on GRR, an open-source remote live-forensics system, back in 2011. For the past three years we've been using it to analyze Joe's machine and do all of the above. Recently, we've added the ability to write and share simple definitions for forensic artifacts and perform large scale binary collection to hunt for badness across the fleet.

Greg will introduce GRR capabilities with some use cases and discuss the difficulties of running the tool across different environments. He will explain and demonstrate GRR artifact collection as well

as talk about some of the aspects that make artifacts powerful but challenging to implement. He'll finish with a discussion of future directions for artifacts, integration with other open source forensics projects, IOCs, and the GRR project in general.

presented by

Greg Castle

Hacking the Wireless World with Software Defined Radio - 2.0

Ever wanted to spoof a restaurant's pager system? How about use an airport's Primary Surveillance RADAR to build your own bistatic RADAR system and track moving objects? What sorts of RF transactions take place in RFID systems, such as toll booths, building security and vehicular keyless entry? Then there's 'printing' steganographic images onto the radio spectrum...

Wireless systems, and their radio signals, are everywhere: consumer, corporate, government, amateur - widely deployed and often vulnerable. If you have ever wondered what sort of information is buzzing around you, this talk will introduce how you can dominate the RF spectrum by 'blindly' analysing any signal, and then begin reverse engineering it from the physical layer up. I will demonstrate how these techniques can be applied to dissect and hack RF communications systems, such as those above, using open source software and cheap radio hardware. In addition, I'll show how long-term radio data gathering can be used to crack poorly-implemented encryption schemes, such as the Radio Data Service's Traffic Message Channel.

I'll also look briefly at some other systems that are close to my heart: reversing satellite communications, tracking aircraft using Mode S and visualising local airspace in real-time on a 3D map, monitoring the health of aircraft with ACARS (how many faults have been reported by the next plane you'll be travelling on, e.g. do the toilets work?), and hunting down the source of an interfering clandestine radio transmission.

If you have any SDR equipment, bring it along!

presented by

Balint Seeber

How Smartcard Payment Systems Fail

The USA is starting to introduce EMV, the Europay-Mastercard-Visa system for making payments using chip cards instead of the old mag strip variety. EMV is already in wide use in Europe, and has started to appear in countries from Canada to India. In theory, smartcards should have reduced fraud by making bankcards much harder to copy and by enabling banks to authenticate users at the point of sale using PINs rather than signatures. The practice has been different. In Britain, for example, fraud first went up, then down, and is now headed upwards again. There have been many fascinating attacks, which I'll describe. The certification system wasn't fit for purpose, so

terminals that were certified as tamper-resistant turned out not to be. We even saw Trojans inserted in the supply chain. A protocol flaw meant that a crook could use a stolen card without knowing the PIN; he could use a man-in-the-middle device to persuade the terminal that the card had accepted the PIN, while the card was told to do a signature-only transaction. Merchant refunds were not authenticated, so a crook could pretend to the bank that he was a merchant, and credit his card back after making a purchase.

The most recent series of attacks exploit the freshness mechanisms in the EMV protocol. To prevent transaction replay, the terminal generates an "unpredictable number" while the card supplies an "application transaction counter" or ATC that is supposed to increase monotonically and never repeat. Yet the unpredictable numbers often aren't (in many of the terminals we looked at, they seem to be just counters) while many banks don't bother to check the ATC, as writing code to deal with out-of-order offline transactions is too much bother. As a result, we've seen some interesting attacks where cardholders unlucky enough to shop at a dishonest merchant find themselves dunned for a lot of large transactions later. In fact these "preplay" attacks behave just like card cloning, and make all the fancy tamper-resistant electronics almost irrelevant.

At heart these are problems of governance and regulation. The vendors sell what they can get away with; the acquiring banks dump liability on merchants and card-issuing banks; they in turn dump it on the cardholder where they can; and the regulators just don't want to know as it's all too difficult. This wonderful system is now being rolled out at scale in the USA.

presented by

Ross Anderson

Many Online Social Networks (OSN) are using OAuth 2.0 to grant access to API endpoints nowadays. Despite many thorough threat model analyses (e.g. RFC6819), only a few real world attacks have been discovered and demonstrated. To our knowledge, previously discovered loopholes are all based on the misuse of OAuth. It was generally believed that the correct use of OAuth 2.0 (by OSN provider and application developer) is secure enough. We break this belief by demonstrating a massive leakage of user data which roots from the scotoma of OAuth's fundamental design rationale: focus on protecting user, not protecting application.

We show that, even if OSN providers and application developers follow best practice, application impersonation is inevitable on many platforms: According to the OAuth 2.0 standard, they support implicit-authorization-grant flow and bearer-token usage. Although it has become common knowledge for application developers to use authorization-code-grant flow and use access token in a MAC-token style wherever possible, there is no mechanism for them to opt out from the OSN platforms' support of implicit-authorization-grant flow and bearer-token usage. Since different applications may have different privileges like accessing permissions and rate limits, application impersonation in general enables privilege escalation and the consequence depends on platform-specific details.

As a proof-of-concept experiment, application impersonation has been demonstrated on a large-scale Facebook-like (not Facebook) OSN. Based on this technique, one can use a casual crawler to collect its 100-million-user social graph within just one week and the projected cost based on

Amazon Web Service is just $150 USD. Due to its implementation specifics, similar techniques can be applied on this OSN to obtain other private data like all users' status lists and albums. Note that, without privilege escalation, this amount of data (order of 10^8) cannot be obtained in such short time with such little cost even on open graphs like Twitter.

Our discovery shows that it is urgent for industrial practitioners to provide the two aforementioned opt-outs in OAuth and review their API design. This work also highlights that application protection must be considered in the design of the next version of OAuth, and similarly other Single-Sign-On protocols.

presented by

Pili Hu & Wing Cheong Lau

How to Wear Your Password

We introduce a new authentication paradigm that achieves both a desirable user experience and a high level of security. We describe and demo an implementation of an identity manager in the guise of a smart bracelet. This bracelet is equipped with a low-power processor, a Bluetooth LE transmitter, an accelerometer, and a clasp that is constructed so that opening and closing it breaks and closes a circuit, thereby allowing an automatic detection of when the bracelet is put on and taken off. However, for reasons of cost, design and error avoidance, the bracelet does not have any user interface, nor any biometric sensors: All user interaction is assisted by third-party devices, such as user phones and point of sale terminals.

Our approach is based on the principle of physical tethering of an identity manager to a user (e.g., by closing the clasp), where the identity manager represents its user's interests after an initial user authentication phase, and until the user causes a disassociation by untethering the device (e.g., by opening the clasp). The authentication phase can be based on any type of authentication, and - to allow for the greatest possible simplicity of design - is aided by a third-party device, such as the user's cell phone.

We describe the physical design, including aspects to protect against violent attacks on users. We also describe the lightweight security protocols needed for pairing, determination of user intent, and credential management, and give examples of usage scenarios- including automated login; simplified online and point-of-sale purchases; assisted appliance personalization; and automated event logging. We then detail the protocols associated with the example usage scenarios, and discuss the security implications of our proposed design.

presented by

Markus Jakobsson

I Know Your Filtering Policy Better than You Do: External

Enumeration and Exploitation of Email and Web Security Solutions

Email and web filtering products and services are core components for protecting company employees from malware, phishing and client-side attacks.

However, it can be trivial for an attacker to bypass these security controls if they know exactly what products and services are in use, how they are configured, and have a clear picture of the solutions' weaknesses in advance of an attack.

The Speaker has previously demonstrated that email and web filtering security appliances often have vulnerabilities which can be exploited to enable an attacker to gain control of these systems (and the data they process). More recently, he has been researching what information an external attacker can discover about the filtering solutions that a target organization has, and how to bypass controls to deliver effective client-side attacks to target employees, without detection.

In this presentation, the Speaker will demonstrate new tools and techniques for the automated enumeration of email and web filtering services, products and policies, and will show how flaws can be discovered and exploited.

This presentation will include statistical analysis of the filtering products, services and policies used by some of the world's top companies. He will show examples of easy-to-create client-side attacks which evade most filtering solutions, and work on fully patched systems to give attackers remote control.

These tools and techniques are very useful from a defensive perspective, to quickly enable the identification of filtering weaknesses and misconfiguration, or to assess the capabilities of filtering products and services.

presented by

Ben Williams

ICSCorsair: How I Will PWN Your ERP Through 4-20 mA Current Loop

Modern Industrial Control Systems (ICS) are deeply integrated with other parts of corporate networks. Plant Asset Management systems, OPC, and SCADA interconnect low-level devices, such as transmitters, actuators, PLCs, with high-level applications, such as MES and ERP. But, what will happen if you can connect to the line where low-level network protocols (such as HART (FSK over 4-20 mA current loop), FF H1, Profibus DP, Modbus over RS-485, e t.c.) flow? Almost everyone knows that then you can probably affect industrial processes. But, there is something more: from this point, you can attack not only the lowest levels of the network, but also PAS, MES, and even ERP systems!

ICSCorsair is an open hardware tool for auditing low-level ICS protocols. It can communicate with various systems using HART FSK and P8CSK, Foundation Fieldbus H1, Profibus, and Modbus protocols. You can control ICSCorsair via USB cable or remotely over WiFi, Bluetooth, or other wireless connection. Different software will be presented to work with ICSCorsair: Metasploit modules, apps for iOS, and Android, etc.

In this talk, it will be shown how to trigger such vulnerabilities as XXE, DoS, XSS, and others in SCADA, PAS, ERP, and MES systems using only ICSCorsair and the opportunity to connect to low-level ICS protocol line.

presented by

Alexander Bolshev & Gleb Cherbov

Internet Scanning - Current State and Lessons Learned

After publishing raw data sets and engaging with the community within our Internet Scanning efforts labeled Project Sonar, there were several logical next steps and an endless amount of ideas to follow up on. In the first quarter of 2014, we were implementing databases, search engines, and generic trending features on top of the collected data from the project. Several community members, from students to pentesters and researchers, downloaded the data sets and started analysis on their own or used it for their work.

This talk presents the latest results from our efforts, such as investigative tools that allow for correlation of the data sets and a generic trending database that allows us to monitor security improvements by country or industry type.

At the same time, we will present the next scan types we are publishing and would like to bring attention to the new possibilities. We demo example processing and show how to work with the data.

Last but not least we will visit the latest findings in terms of vulnerabilities and misconfigurations that we came across in the deep corners of the internet. For example we will talk about statistics around the SSL heartbleed vulnerability that can be generated from our datasets.

presented by

Mark Schloesser

Investigating PowerShell Attacks

Over the past two years, we've seen targeted attackers increasingly make use of PowerShell to conduct command-and-control in compromised Windows environments. If your organization is running Windows 7 or Server 2008 R2, you've got PowerShell 2.0 installed (and on Server 2012, remoting is enabled by default!). This has created a whole new playground of attack techniques for intruders that have already popped a few admin accounts (or an entire domain). Even if you're not legitimately using PowerShell to administer your systems, you need to be aware of how attackers can enable and abuse its features.

This presentation will focus on common attack patterns performed through PowerShell - such as lateral movement, remote command execution, reconnaissance, file transfer, and establishing persistence - and the sources of evidence they leave behind. We'll demonstrate how to collect and

interpret these forensic artifacts, both on individual hosts and at scale across the enterprise. Throughout the presentation, we'll include examples from real-world incidents and recommendations on how to limit exposure to these attacks.

presented by

Ryan Kazanciyan & Matt Hastings

It Just (Net)works: The Truth About iOS 7's Multipeer Connectivity Framework

With the release of iOS 7, Apple has quietly introduced a nifty feature called Multipeer Connectivity. Using a surprisingly small and simple set of APIs, developers can create applications that have the ability to discover and directly communicate with nearby iOS devices over Bluetooth or WiFi, without the need for an Internet connection. While the Multipeer Connectivity Framework brings the promise of peer-to-peer and mesh networking apps significantly closer to reality, little is known regarding how it actually works behind the scenes and what the risks are for applications leveraging this functionality.

This talk will first present an analysis of what happens at the network level when two devices start communicating with each other over WiFi, including a description of the protocols and encryption algorithms used. From this analysis, we'll derive a security model for Multipeer Connectivity and describe the threats and underlying assumptions that developers should be aware of when building applications. The impact of the various pairing options, data transmission modes, and encryption settings exposed by the Framework will also be explained. Lastly, we'll study the implementation of a real-world app that uses the Framework and describe issues and potential weaknesses; at the end of the presentation, a tool that was used to find some of these issues will be released.

presented by

Alban Diquet

Learn How to Control Every Room at a Luxury Hotel Remotely: The Dangers of Insecure Home Automation Deployment

Have you ever had the urge to create mayhem at a hotel? Force every hotel guest to watch your favorite TV show with you? Or wake your neighbors up (all 290 of them!) with blaring music and with their blinds up at 3 AM?

For those with the urge, I have the perfect place for you. The St. Regis ShenZhen, a gorgeous luxury hotel occupying the top 28 floors of a 100 story skyscraper, offers guests a unique feature: a room remote control in the form of an iPad2. The iPad2 controls the lighting, temperature, music, do not disturb light, TV, even the blinds and other miscellaneous room actions. However,

the deployment of the home automation protocol contained several fatal flaws that allow an arbitrary attacker to control virtually every appliance in the hotel remotely. I discovered these flaws and as a result, I was able to create the ultimate remote control: Switch TV off 1280,1281,1283 will switch off the TV in these three room. The attacker does not even need to be at the hotel - he could be in another country.

This talk provides a detailed discussion of the anatomy of the attack: an explanation of reverse engineering of the KNX/IP home automation protocol; a description of the deployment flaws; blueprints on how to create an iPad Trojan to send commands outside the hotel; and, of course, solutions to avoid all these pitfall in future deployments. Attendees will gain valuable field lessons on how to improve wide scale home automation architectures and discussion topics will include the dangers of utilizing legacy but widely used automation protocols, the utilization of insecure wireless connection, and the use of insecure and unlocked commodity hardware that could easily be modified by an attacker.

The attack has important implications for large scale home automation applications, as several hotels around the world are beginning to offer this room amenity. The severity of these types of security flaws cannot be understated - from creating a chaotic atmosphere to raising room temperatures at night with fatal consequences - hoteliers need to understand the risks and liabilities they are exposed to by faulty security deployments.

presented by

Jesus Molina

Leviathan: Command and Control Communications on Planet Earth

Every day, computer network attackers leverage a Leviathan of compromised infrastructure, based in every corner of the globe, to play hide-and-seek with network security, law enforcement, and counterintelligence personnel.

This presentation draws a new map of Planet Earth, based not on traditional parameters, but on hacker command and control (C2) communications. The primary data points used in this worldwide cyber survey are more than 30 million malware callbacks to over 200 countries and territories over an 18-month period, from January 2013 to June 2014.

First, this talk covers the techniques that hackers use to communicate with compromised infrastructure across the globe. The authors analyze the domains, protocols, ports, and websites used for malicious C2. They explain how covert C2 works, and how attackers keep their communications hidden from network security personnel.

Second, this talk looks at strategic impact. The authors examine relationships between the targeted industries and countries and the first-stage malware servers communicating with them. Traffic analysis is used to deduce important relationships, patterns, and trends in the data. This section correlates C2 communications to traditional geopolitical conflicts and considers whether computer network activity can be used to predict real world events.

In conclusion, the authors consider the future of this Leviathan, including whether governments can subdue it and whether they would even want to.

presented by

Kenneth Geers & Kevin Thompson

Lifecycle of a Phone Fraudster: Exposing Fraud Activity from Reconnaissance to Takeover Using Graph Analysis and Acoustical Anomalies

Enterprises are vulnerable to "human hacking," the effective social engineering of employees, contractors, and other trusted persons. In particular, financial institutions have seen a significant increase in account takeover attacks over the phone by sophisticated fraudsters socially engineering call center agents. The customer information required is often obtained by gathering intelligence through reconnaissance, probing systems or humans. In this talk, we will show how to detect both the account takeover calls using acoustical anomalies and the reconnaissance calls leading to it through graph analysis. Using acoustical anomalies, we are able to detect over 80% of these calls with less than a 2% false positive rate. Furthermore, our graph analysis is able to see reconnaissance calls for 46% of these account takeovers 10 days before the actual takeover. These results are on a dataset of over hundreds of million calls. In the process, we will reveal the lifecycle of a phone fraudster as he works through both the call center agent and its technology to extract information about a customer and takeover his or her account.

presented by

Vijay Balasubramaniyan & Raj Bandyopadhyay & Telvis Calhoun

Miniaturization

Too often researchers ignore the hard parts of SCADA hacking. Too many presentations could be described as "I got past the SCADA firewall so I win!!!" Little information is available on what to do after the attacker gains control of the process. As a challenge, consider the scenario where I just gave you control of a paint factory. Now what? The answer to that question is often specific to the process, but there are a number of generic techniques that can be discussed. Often, designing an attack leads to interesting hacking and computer science challenges.

Miniaturization is one of those problems. Suppose an attacker wanted to hide in a PLC. Suppose he wanted to hide all the way down in a pressure sensor. Is such a thing possible? The attack must be miniaturized to fit within the constraints of the embedded device and may need to be miniaturized into just a few kilobytes of memory. This is an interesting problem.

The sensor has only a few kilobytes of memory and the attacker has a number of tasks to perform. During the attack he must spoof the original process to keep the operator happy. He must

estimate the state of the physical process by extracting artifacts from noisy sensor signals. He must also process those artifacts to extract the necessary constants to perform an attack.

In order to keep the presentation real and understandable, it will walk through setting up an optimal pressure transient in a chemical piping system. (Commonly referred to as a water hammer). A set of novel algorithms will be describe that would allow someone to pull off such an attack. A variant of "runs analysis" taken from statistics will be used to produce nearly perfect sensor noise without previous look at the sensor. An algorithm derived from 3D graphics will be used to extract artifacts from noisy sensor data. Finally scale-free geometry matching techniques will be used to process the artifacts into the time constants needed to pull off an attack.

presented by

Jason Larsen

Mission mPOSsible

Mobile Point-of-Sale (mPOS) systems allow small businesses and drug dealers to accept credit card payments using their favourite iDevice (Disclaimer: other mobile devices are available). During our research, we had a look at the security of the leading solutions for mobile Chip&Pin; payments. If you saw our previous PinPadPwn research, you won't be surprised to hear we discovered a series of vulnerabilities which allow us to gain code execution on these devices through each of the available input vectors. We will discuss the weaknesses of current solutions and have live demonstrations for multiple attack vectors, our favourite being a malicious credit card which drops a remote root shell on an embedded mPOS device.

presented by

Nils & Jon Butler

Mobile Device Mismanagement

MDM solutions are ubiquitous in today's enterprise environment. They provide a way for security and IT departments to mitigate the risk of mobile malware and lost/stolen devices when personal devices are being used to access and store corporate resources.

Like any other piece of software being deployed on a large scale, we need to ask the questions "is it secure?," "what are the risks?"; because MDM is a security product itself, this crucial step seems to have been overlooked. With a few exceptions, the security community has not had much to say about vulnerabilities in MDM products and this is likely due to the extremely restrictive licensing requirements to gain access to the software.

This talk focuses on vulnerabilities in MDM products themselves. Through a number of penetration tests we have conducted on our clients, we have discovered and leveraged critical

vulnerabilities in MDM solutions to gain access to sensitive information. We will provide an overview of these vulnerabilities, some of which seem to be systemic across a number of products.

presented by

Stephen Breen

MoRE Shadow Walker: The Progression of TLB-Splitting on x86

This talk will cover the concept of translation lookaside buffer (TLB) splitting for code hiding and how the evolution of the Intel x86 architecture has rendered previous techniques obsolete and new techniques to perform TLB-splitting on modern hardware. After requisite background is provided, a timeline of how TLB-splitting was used for both defensive (PaX memory protections) and offensive purposes (Shadow Walker root-kit) and how the new Intel Core i-series processors fundamentally changed the TLB architecture, breaking those technologies. The talk will then move to the new research, the author's method for splitting a TLB on Core i-series and newer processors and how it can again be used for defensive (MoRE code-injection detection) and offensive purposes (EPT Shadow Walker root-kit).

After the timeline, details on how to perform and leverage TLB-splitting with the EPT Shadow Walker root-kit is used to present one version of memory to defensive tools for validation and a different (and possibly malicious) version to the CPU for execution, effectively hiding a root-kit from anti-virus or anti-patching systems. A demo of this memory changing and hiding will be shown and results from the research presented.

presented by

Jacob Torrey

Multipath TCP: Breaking Today's Networks with Tomorrow's Protocols

MultiPath TCP (MPTCP) is an extension to TCP that enables sessions to use multiple network endpoints and multiple network paths at the same time, and to change addresses in the middle of a connection. MPTCP works transparently over most existing network infrastructure, yet very few security and network management tools can correctly interpret MPTCP streams. With MPTCP network security is changed: how do you secure traffic when you can't see it all and when the endpoint addresses change in the middle of a connection?

This session shows you how MPTCP breaks assumptions about how TCP works, and how it can be used to evade security controls. We will also show tools and strategies for understanding and mitigating the risk of MPTCP-capable devices on a network.

presented by

Catherine Pearce & Patrick Thomas

My Google Glass Sees Your Passwords!

In this presentation, we introduce a novel computer vision based attack that automatically discloses inputs on a touch enabled device. Our spying camera, including Google Glass, can take a video of the victim tapping on the touch screen and automatically recognize more than 90% of the tapped passcodes from three meters away, even if our naked eyes cannot see those passcodes or anything on the touch screen. The basic idea is to track the movement of the fingertip and use the fingertip's relative position on the touch screen to recognize the touch input. We carefully analyze the shadow formation around the fingertip, apply the optical flow, deformable part-based model (DPM) object detector, k-means clustering and other computer vision techniques to automatically track the touching fingertip and locate the touched points. Planar homography is then applied to map the estimated touched points to a software keyboard in a reference image. Our work is substantially different from related work on blind recognition of touch inputs. We target passcodes where no language model can be applied to correct estimated touched keys. We are interested in scenarios such as conferences and similar gathering places where a Google Glass, webcam, or smartphone can be used for a stealthy attack. Extensive experiments were performed to demonstrate the impact of this attack. As a countermeasure, we design a context aware Privacy Enhancing Keyboard (PEK) which pops up a randomized keyboard on Android systems for sensitive information such as password inputs and shows a conventional QWERTY keyboard for normal inputs.

presented by

Xinwen Fu & Qinggang Yue & Zhen Ling

Network Attached Shell: N.A.S.ty Systems that Store Network Accessible Shells

Through extensive analysis, Independent Security Evaluators (ISE) has identified dozens of previously undisclosed, critical security vulnerabilities in numerous network storage devices from a handful of goto manufacturers (manufacturers: e.g., Seagate, D-Link, Netgear). Vulnerabilities of network-attached storage not only expose stored data, but also provide a vantage point for further PWNAGE of the network infrastructure on which the storage system sits. Our research efforts focused on identifying vulnerabilities that obtained administrative access (such as command injection, directory traversal, authentication bypass, memory corruption, backdoors, etc.), and quantifying the associated risk.

The attacks we developed demonstrate how unauthenticated attackers can compromise and control storage systems with and without user interaction.

Network based storage systems are used in millions of homes, schools, government agencies, and businesses around the world for data storage and retrieval. With today's dependence on Internet

based services, virtualization technologies, and the need to access data from anywhere, storage systems are relied on more than ever. Similar to other network hardware (e.g., routers), these devices are purchased and installed by IT teams and home consumers with the expectation that the system is protected from the infamous hacker.

This presentation focuses on "how to," and the implications of compromising network based storage systems, but will conclude that the absence of security in not only storage hardware, but networking hardware in general, has left data unprotected and millions of networks vulnerable to exploitation.

Throughout this presentation, several vulnerabilities will be exploited in order to achieve the glorious ro0t (#) shell!

presented by

Jacob Holcomb

One Packer to Rule Them All: Empirical Identification, Comparison, and Circumvention of Current Antivirus Detection Techniques

Lately, many popular anti-virus solutions claim to be the most effective against unknown and obfuscated malware. Most of these solutions are rather vague about how they supposedly achieve this goal, making it hard for end-users to evaluate and compare the effectiveness of the different products on the market. This presentation presents empirically discovered results on the various implementations of these methods per solution, which reveal that some anti-virus solutions have more mature methods to detect x86 malware than others, but all of them are lagging behind when it comes to x64 malware. In general, at most three stages were identified in the detection process: Static detection, Code Emulation detection (before execution), and Runtime detection (during execution). New generic evasion techniques are presented for each of these stages. These techniques were implemented by an advanced, dedicated packer, which is an approach commonly taken by malware developers to evade detection of their malicious toolset. Two brand new packing methods were developed for this cause. By combining several evasion techniques, real-world malicious executables with a high detection rate were rendered completely undetected to the prying eyes of anti-virus products.

presented by

Alaeddine Mesbahi & Arne Swinnen

OpenStack Cloud at Yahoo Scale: How to Avoid Disaster

OpenStack is an Open Source project that allows you to manage a cloud of VMs that has grown into a widely adopted platform. The issue with having a centralized Infrastructure As A Service (IAAS) is that if you compromise the management cluster you can attack everything it controls,

which is a lot at Yahoo scale. How do you keep your OpenStack cluster safe? What do you do when a management system, hypervisor, or VM is compromised?

This talk will discuss specific things that you can do to harden your cluster and make it more difficult for a large compromise to happen. If a compromise is detected, there are specific steps you can take to reduce the impact as well as to gather intelligence you can take action on. The impact of different network architectures on OpenStack security will also be discussed. Throughout this talk, I will use examples from the Yahoo deployments of OpenStack clusters to illustrate what Yahoo does to secure its systems and ensure our users continue to trust us.

presented by

Anders Beitnes

Oracle Data Redaction is Broken

The Oracle data redaction service is a new feature introduced with Oracle 12c. It allows sensitive data, such as PII, to be redacted to prevent it being exposed to attackers. On paper this sounds like a great idea, but in practice, Oracle's implementation is vulnerable to multiple attacks that allow an attacker to bypass the redaction and launch privilege escalation attacks.

presented by

David Litchfield

Pivoting in Amazon Clouds

From no access at all, to the company Amazon's root account, this talk will teach attendees about the components used in cloud applications like: EC2, SQS, IAM, RDS, meta-data, user-data, Celery; and how misconfigurations in each can be abused to gain access to operating systems, database information, application source code, and Amazon's services through its API.

The talk will follow a knowledgeable intruder from the first second after identifying a vulnerability in a cloud-deployed Web application and all the steps he takes to reach the root account for the Amazon user.

Except for the initial vulnerability, a classic remote file included in a Web application which grants access to the front-end EC2 instance, all the other vulnerabilities and weaknesses exploited by this intruder are going to be cloud-specific.

The tools used by this intruder are going to be released after the talk and will provide the following features:

- Enumerate access to AWS services for current IAM role
- Use poorly configured IAM role to create new AWS user
- Extract current AWS credentials from meta-data, .boto.cfg, environment variables, etc.

- Clone DB to access information stored in snapshot
- Inject raw Celery task for pickle attack

presented by

Andres Riancho

Poacher Turned Gamekeeper: Lessons Learned from Eight Years of Breaking Hypervisors

Hypervisors have become a key element of both cloud and client computing. It is without doubt that hypervisors are going to be commonplace in future devices, and play an important role in the security industry. In this presentation, we discuss in detail the various lessons learnt whilst building and breaking various common hypervisors. In particular, we take a trip down memory lane and examine vulnerabilities found in all the popular hypervisors that have led to break-outs. To add some spice, we will talk about details of four not-yet-discussed vulnerabilities we recently discovered.

One of the key value propositions of hypervisors as they relate to security is to shrink the attack surface. However, in the quest for new features and functionality some trade-offs are made, which can prove to be fatal. While discussing the particular problems we will examine what the strong (and weak) security-related features of hypervisors are. We compare the attack surface of hypervisors with that of user mode applications and operating systems kernels, and show that the purpose and design of the hypervisor significantly changes its attack surface size. Most importantly, we make a fact based argument that many hypervisors aren't designed with security in mind. We show how superfluous code and poor design can be punished by demonstrating real examples of hypervisor break-outs.

The presentation ends with lessons learned and recommendations for hypervisor design and approaches that can be taken to harden them.

presented by

Rafal Wojtczuk

Point of Sale System Architecture and Security

To most people, Point of Sale (POS) systems with integrated payment processing are a black box where magic happens. Financial criminals breach hundreds of merchants each year, displaying a better understanding of how these systems operate than the dealer technicians that install and maintain them. With an understanding of POS architecture, integrated payment processing, and weaknesses in the technology, security professionals can better protect local businesses, major retailers, and developers handling payment card information. In this session, attendees will learn and see how POS components operate, their integration points, and the flow of payment data

including where it's most vulnerable. A live demonstration will show exactly what sensitive data is passed in the clear by both magstripe and EMV chip readers, mapping it from peripheral all the way through the electronic payments infrastructure. Common attack vectors will then be presented, building on that architectural knowledge. Finally, top attack mitigations will be provided to save businesses from being breached and the disastrous losses that result.

presented by

Lucas Zaichkowsky

Prevalent Characteristics in Modern Malware

Malware is widely acknowledged as a growing threat with hundreds of thousands of new samples reported each week. Analysis of these malware samples has to deal with this significant quantity but also with the defensive capabilities built into malware. Malware authors use a range of evasion techniques to harden their creations against accurate analysis. The evasion techniques aim to disrupt attempts of disassembly, debugging or analyze in a virtualized environment.

Two years ago, in 2011, we presented (with other researchers) at Black Hat USA a wide range of anti-reverse engineering techniques that malware were currently employing. For each technique, we documented how it works, we created an algorithm to detect its usage, and we provided statistics on the technique prevalence in a 4 million samples database. We also provided a fully-working PoC implementing each of the techniques (either in C or Assembly). Our expectation was that the AV industry would use our ideas (proven with the prevalence numbers) to significantly improve the malware prevention coverage. Nothing changed. In the meanwhile, we improved our detection algorithms, fixed bugs, and expanded the research to 12+ million samples.

In this talk, we are going to give another try and demonstrate the prevalence of more than 50 non-defensive additional characteristics found in modern malware. Additionally to that, we also extended our previous research demonstrating what the malware does once it detects it is being analyzed. The resulting data will help security companies and researchers around the world to focus their attention on making their tools and processes more efficient to rapidly avoid the malware authors' countermeasures.

This first of its kind, comprehensive catalog of malware characteristics was compiled by the paper's authors by researching some techniques employed by malware, and in the process new detections were proposed and developed. The underlying malware sample database has an open architecture that allows researchers not only to see the results of the analysis, but also to develop and plug-in new analysis capabilities.

presented by

Rodrigo Branco & Gabriel Negreira Barbosa

Probabilistic Spying on Encrypted Tunnels

At the network layer, encrypted tunnels are typically seen as black boxes. Network traffic however, leaks side channel information that can often be analyzed to determine what the tunnel is being used for and the type of content being sent over it. Probabilistic algorithms will be explored that can analyze this side channel information and identify application protocols within the tunnel. An open-source toolkit containing the algorithms/attacks presented will be released.

presented by

Brandon Niemczyk & Prasad Rao

Protecting Data In-Use from Firmware and Physical Attacks

Recent revelations of the NSA ANT program illustrated the many well-known and low-cost physical and firmware attacks that can compromise data in-use and system integrity. These attacks have become more concerning as more computing infrastructure runs outside an organization's physical control.

This talk will review several such attacks, including SMM bootkits, "cold booting," and malicious devices. We'll discuss several existing tools and technologies that can mitigate these risk such as Trusted Execution Technology (TXT) and memory encryption technologies. We will also discuss how upcoming technologies such as Software Guard Extensions (SGX), Enhanced Privacy ID (EPID), and TPM 2.0 can help protect against firmware and physical threats.

presented by

Steve Weis

Pulling Back the Curtain on Airport Security: Can a Weapon Get Past TSA?

Every day, millions of people go through airport security. While it is an inconvenience that could take a while, most are willing to follow the necessary procedures if it can guarantee their safety. Modern airport security checkpoints use sophisticated technology to help the security screeners identify potential threats and suspicious baggage. Have you ever wondered how these devices work? Have you ever wondered why an airport security checkpoint was set up in a particular configuration? Join us as we present the details on how a variety of airport security systems actually work, and reveal their weaknesses. We'll present what we have learned about modern airport security procedures, dive deep into the devices used to detect threats, and we'll present some the bugs we discovered along the way.

presented by

Billy Rios

RAVAGE - Runtime Analysis of Vulnerabilities and Generation of Exploits

In this talk, we will show cutting edge research and a tool built to accurately detect vulnerabilities. The tool leverages the standard program execution to detect the full dataflow of vulnerabilities at runtime. It can be used both offensively and defensively. We will show how RAVAGE can be used to detect vulnerabilities, generate exploits, and integrate the newly found exploits into existing exploitation frameworks. In addition to the offensive usage, it can also be used defensively by running existing non-security-related test cases to detect security vulnerabilities.

We will open source RAVAGE (for Java) as well as design documentation at Black Hat.

presented by

Xiaoran Wang & Yoel Gluck

Reflections on Trusting TrustZone

TrustZone has emerged as a leading option for security-critical tasks on ARM devices. It has been billed as a "100% secure solution" for restricting access to sensitive device hardware components and securely storing highly privileged information. As a result, TrustZone is used on millions of mobile devices for diverse tasks including managing secure boot, storing DRM keys on behalf of digital content providers, supporting mobile payments, and performing integrity validation on the live operating system kernel.

This talk will take a deep technical dive into the inner workings of a major vendor's TrustZone kernel, which is currently deployed on millions of Android devices. After providing a review of prior work in TrustZone exploitation, this talk will describe a previously unpublished vulnerability in this TrustZone implementation, and provide details on steps taken to exploit this vulnerability. The talk will conclude with a discussion of the ramifications of this vulnerability and others like it, including a live demonstration of using it to permanently unlock the bootloader of a major Android phone.

presented by

Dan Rosenberg

Researching Android Device Security with the Help of a Droid Army

In the last few years, Android has become the world's leading smart phone operating system. Unfortunately, the diversity and sheer number of devices in the ecosystem represent a significant challenge to security researchers. Primarily, auditing and exploit development efforts are less effective when focusing on a single device because each device is like a snowflake: unique.

This presentation centers around the speaker's approach to dealing with the Android diversity problem, which is often called "fragmentation." To deal with the issue, Joshua created a

heterogeneous cluster of Android devices. By examining and testing against multiple devices, you can discover similarities and differences between devices or families of devices. Such a cluster also enables quickly testing research findings or extracting specific information from each device.

When you leave this presentation, you will understand why the diversity problem exists and how to tackle it by creating a cluster of your own. Joshua will show you how to build such a cluster, provide a set of tools to manage one, and show you all the ways to leverage it to be more successful in your auditing and exploit development tasks.

presented by

Joshua Drake

Reverse Engineering Flash Memory for Fun and Benefit

There are many benefits to interacting directly with Flash memory when you're having a hard time finding the correct JTAG connection points. That's especially true when you're a software reverse engineer who delves into hardware reversing. Some vendors intentionally obfuscate JTAG points or remove them to prevent reverse engineering.

In this talk, we look closely at the process of reverse engineering embedded devices by interacting directly with Flash memory. We also look at reprogramming chips and putting them back on the board. The fun with this method is that you can access the underlying out-of-band data that contains page and block information. As Flash memory is a fragile media, bad blocks or page data contamination are common problems. Whenever you extract data from memory, you should be able to take care of this meta information. When you write back the data, you need to recalculate sums and set the correct flags on these areas. We talk about the chips we've worked on and how we have dealt with the meta information.

The other entertaining part we'll examine is the file system. Embedded systems that interact directly with Flash memory usually use journaling file systems to avoid repeating write operations on specific pages. The journaling file system is interesting as it contains the entire history of file operations. You can just mount the file system directly from your Linux box or you can write a simple parser to check the history of the file system operations. This feature might give reverse engineers a good view of how Flash memory is programmed and used.

presented by

Jeong Wook Oh

Reverse-Engineering the Supra iBox: Exploitation of a Hardened MSP430-Based Device

This presentation walks through the reverse engineering and exploitation of a hardened embedded device and provides certain techniques you can use to exploit similar devices. The

Supra iBox BT is a bluetooth and IR-based physical key storage device used by many real estate professionals in the US. It is physically hardened, and inside is a hardened MSP430 with a blown JTAG fuse. As MSP430 devices become more common, it is slowly becoming the norm to encounter devices in production with blown JTAG fuses. Previously, this was a significant hurdle. In 2008, Goodspeed described several attacks against the MSP's BSL (bootstrap loader). This presentation will review those attacks and describe the challenges facing a researcher attempting to perform them. This presentation will demonstrate how to reliably perform successful firmware extraction on a MSP430 with a blown JTAG fuse.

The second part of the presentation covers what I found inside the Supra iBox firmware, including a demonstration of an exploit that can open any iBox. The presentation will describe the complex and surprisingly effective crypto key management scheme used by Supra. The Supra lock has no internet access, and must rely on the keys (generally smartphones) to perform any necessary synchronization with the internet. The key management scheme used by the Supra would be interesting to any developer attempting to manage cryptographic keys in embedded devices with occasional internet access.

presented by

Braden Thomas

SAP, Credit Cards, and the Bird that Talks Too Much

SAP applications build the business backbone of the largest organizations in the world. In this presentation, exploits will be shown manipulating a business process to extract money, critical payment information, and credit card data out of the business backbone. Follow the bird and enjoy tweets of data that will interest you.

presented by

Ertunga Arsal

SATCOM Terminals: Hacking by Air, Sea, and Land

Satellite Communications (SATCOM) play a vital role in the global telecommunications system. We live in a world where data is constantly flowing. It is clear that those who control communications traffic have a distinct advantage. The ability to disrupt, inspect, modify, or re-route traffic provides an invaluable opportunity to carry out attacks.

SATCOM infrastructure can be divided into two major segments, space and ground. Space includes those elements needed to deploy, maintain, track, and control a satellite. Ground includes the infrastructure required to access a satellite repeater from Earth station terminals.

Earth station terminals encompass the equipment located both on the ground and on airplanes and ships; therefore, this segment includes air and sea. This specific portion of the ground

segment was the focus of our research. We analyzed devices, from leading SATCOM vendors, used to access services such as:

- Inmarsat-C
- Inmarsat BGAN / M2M
- FleetBroadBand
- SwiftBroadBand
- Classic Aero Services
- GMDSS (Global Maritime Distress Safety System)
- SSAS (Ship Security Alert System)

IOActive found that 100% of the devices could be abused. The vulnerabilities we uncovered included multiple backdoors, hardcoded credentials, undocumented and/or insecure protocols or weak encryption algorithms.

These vulnerabilities allow remote, unauthenticated attackers to fully compromise the affected products. In certain cases no user interaction is required to exploit the vulnerability, just sending a simple SMS or specially crafted message from one ship to another ship can do it.

The talk will show all the technical details, mainly based on static firmware analysis via reverse engineering, also including a live demo against one of these systems.

Ships, aircraft, military personnel, emergency services, media services, and industrial facilities (oil rigs, gas pipelines, water treatment plants, wind turbines, substations, etc.) could all be impacted by these vulnerabilities.

presented by

Ruben Santamarta

Saving Cyberspace

Imagine that twenty years after Johannes Gutenberg invented mechanical movable type, the Pope and the petty princes - in fact, anyone who tried hard enough - had the ability to determine exactly who was printing exactly what. Worrying about intellectual property theft, privacy or civil rights violations, had those concepts existed, would be missing the point.

The future of Europe, the future of humanity, would have been profoundly changed, not just for five years but five hundred. If people lost trust in the underlying communication medium, could there even have been a Renaissance or Enlightenment?

Unfortunately, the world is facing this dilemma today as it is possible, even likely, the Internet will not remain as resilient, free, secure, and as awesome, for future generations as it has been for ours. It is under grave threat from data breaches, theft of commercial secrets, the opportunity for widespread disruptive attacks and systemic failures, erection of sovereign borders, and mass surveillance.

The only truly goal for this new cyber strategy should be to give the defenders the high-ground advantage over attackers. This is just imaginable with a clever push for new technology, policy, and practice which is applied patiently, internationally, at scale, and with the private sector at the fore.

This talk will discuss these threats to the Internet and novel approaches to sidestep much of the current unproductive debate over privacy versus security.

presented by

Jason Healey

SecSi Product Development: Techniques for Ensuring Secure Silicon Applied to Open-Source Verilog Projects

Secure development processes for software have formed, developed, and matured in the past decade to the point where there are well defined categories of security bugs and proven methods to find them. Secure hardware development, on the other hand, is essentially undefined at this point. Most developers of integrated circuits do no hardware security validation, or are secretive about their methods and findings.

This talk will document some pre- and post- silicon validation techniques by applying them to various open-source core designs. It will present a number of examples of actual Verilog security vulnerabilities along with the vulnerable code, and present methods of resolving them. It will conclude by generalizing several hardware security bug categories.

presented by

Joseph FitzPatrick

Secure Because Math: A Deep-Dive on Machine Learning-Based Monitoring

We could all have predicted this with our magical Big Data analytics platforms, but it seems that machine learning is the new hotness in Information Security. A great number of start-ups with 'cy' and 'threat' in their names that claim that their product will defend or detect more effectively than their neighbors' product "because math." And it should be easy to fool people without a PhD or two that math just works.

Indeed, math is powerful and large scale machine learning is an important cornerstone of much of the systems that we use today. However, not all algorithms and techniques are born equal. Machine learning is a very powerful tool box, but not every tool can be applied to every problem and that's where the pitfalls lie.

This presentation will describe the different techniques available for data analysis and machine learning for information security, and discuss their strengths and caveats. The ghost of marketing

past will also show how similar the unfulfilled promises of deterministic and exploratory analysis were, and how to avoid making the same mistakes again.

Finally, the presentation will describe the techniques and feature sets that were developed by the presenter in the past year as a part of his ongoing research project on the subject, in particular he'll present some interesting results obtained since the last presentation at Black Hat USA 2013, and some ideas that could improve the application of machine learning for use in information security, specially in its use as a helper for security analysts in incident detection and response.

presented by

Alex Pinto

Sidewinder Targeted Attack Against Android in the Golden Age of Ad Libs

While Google Play has little malware, many vulnerabilities exist in the apps as well as the Android system itself, and aggressive ad libs leak a lot of user privacy information. When they are combined together, more powerful targeted attacks can be conducted.

We will present one practical case of such attacks called "Sidewinder Targeted Attack." It targets victims by intercepting location information reported from ad libs, which can be used to locate targeted areas such as a CEO's office or some specific conference rooms. When the target is identified, "Sidewinder Targeted Attack" exploits popular vulnerabilities in ad libs, such as Javascript-binding-over-HTTP or dynamic-loading-over-HTTP, etc.

During the exploit, it is a well-known challenge to call Android services from injected native code due to the lack of Android application context. So we will also demonstrate how attackers can invoke Android services such as taking photos, calling phone numbers, sending SMS, reading/writing the clipboard, etc.

Once intruding into the target, the attackers can exploit several Android vulnerabilities to get valuable privacy information or initiate more advanced attacks. We will reveal how to exploit new vulnerabilities we discovered in this phase.

In this talk, we will show demos using real-world apps downloaded from Google Play.

Although we notified Google, ad vendors and app developers about related issues half a year ago, there are still millions of users under the threat of "Sidewinder Targeted Attacks" due to the slow patching/upgrading/fragmentation of the Android ecosystem.

presented by

Tao Wei & Yulong Zhang

Smart Nest Thermostat: A Smart Spy in Your Home

The Nest thermostat is a smart home automation device that aims to learn about your heating and cooling habits to help optimize your scheduling and power usage. Debuted in 2010, the smart NEST devices have been proved a huge success that Google spent $3.2B to acquire the whole company. However, the smartness of the thermostat also breeds security vulnerabilities, similar to all other smart consumer electronics. The severity of security breach has not been fully embraced due to the traditional assumption that thermostat cannot function more than a thermostat even though users are enjoying its smartness.

Equipped with two ARM cores, in addition to WiFi and ZigBee chips, this is no ordinary thermostat. In this presentation, we will demonstrate our ability to fully control a Nest with a USB connection within seconds (in our demonstration, we will show that we can plug in a USB for 15 seconds and walk away with a fully rooted Nest). Although OS level security checks are available and are claimed to be very effective in defeating various attacks, instead of attacking the higher level software, we went straight for the hardware and applied OS-guided hardware attacks. As a result, our method bypasses the existing firmware signing and allows us to backdoor the Nest software in any way we choose. With Internet access, the Nest could now become a beachhead for an external attacker. The Nest thermostat is aware of when you are home and when you are on vacation, meaning a compromise of the Nest would allow remote attackers to learn the schedule of users. Furthermore, saved data, including WiFi credentials, would now become available to attackers. Besides its original role of monitor the user's behavior, the smart Nest is now a spy rooted inside a house fully controlled by attackers.

Using the USB exploit mentioned above, we have loaded a custom compiled kernel with debug symbols added. This enables us to explore the software protocols used by the nest, such as Nest Weave, in order to find potential vulnerabilities that can be remotely exploited. Loading a custom kernel into the system also shows how we have obtained total control of the device, introducing the potential for rootkits, spyware, rogue services and other network scanning methods, further allowing the compromise of other nodes within the local network.

presented by

Yier Jin & Grant Hernandez & Daniel Buentello

Static Detection and Automatic Exploitation of Intent Message Vulnerabilities in Android Applications

We identified a set of vulnerabilities that common Android Apps programming (mis)practices might introduce.

We developed an effective static analyzer to automatically detect a set of vulnerabilities rising by incorrect Android's Inter-Component Communication usage.

We completed our analysis by automatically demonstrating whether the vulnerabilities identified by static analysis can actually be exploited or not at run-time by an attacker.

We adopted a formal and sound approach to automatically produce malicious payloads able to reproduce the dangerous behavior in vulnerable applications.

The lack of exhaustive sanity checks when receiving messages from unknown sources is the evidence of the underestimation of this problem in real world application development.

presented by

Daniele Gallingani

Stay Out of the Kitchen: A DLP Security Bake-Off

Despite a plethora of data security and protection standards and certifications, companies and their systems are still leaking information like a sieve. Data Loss Prevention (DLP) solutions have often been touted as the "silver bullet" that will keep corporations from becoming the next headline. With deployment models ranging from a fat agent on an endpoint, to a blinky-lights box surveilling all network traffic, to some unified threat management gateway with DLP secret sauce, these solutions are ripe for bypass - or worse.

This talk will discuss our research into a handful of DLP solutions, including their capabilities and their shortcomings. We will demonstrate flaws in administrative and programmatic interfaces and the inspection engines themselves.

presented by

Zach Lanier & Kelly Lum

SVG: Exploiting Browsers without Image Parsing Bugs

SVG is an XML-based format for vector graphics. Modern web browsers support it natively and allow it to be styled using CSS and manipulated using JavaScript. It is less well-known that SVG can contain its own JavaScript and can import external scripts and stylesheets. Consequently, from a browser security perspective, SVG must be treated like HTML; treating it like JPEG will lead to great suffering.

presented by

Rennie deGraaf

The Beast is in Your Memory: Return-Oriented Programming Attacks Against Modern Control-Flow Integrity Protection Techniques

Return-oriented Programming (ROP) is a powerful exploitation technique used in nearly every exploit today. It maliciously combines short code snippets (gadgets) residing in shared libraries and the executable to bypass data execution prevention (DEP). As a consequence, several new

control-flow integrity (CFI) mechanisms and tools have been recently proposed to thwart ROP attacks. For instance, kBouncer and ROPGuard both restrict return instructions to target a call-preceded instruction. In addition, ROPecker and kBouncer force the adversary to invoke a long instruction sequence after a pre-defined number of short gadgets thereby preventing an attacker to execute a sequence of ROP gadgets. Some of the proposed mechanisms have already been integrated in Microsoft's Windows EMET tool. In general, these mechanisms significantly reduce the gadget space and make it challenging for an attacker to mount ROP attacks.

While others have hypothesized or even exploited weaknesses in some of these emerging CFI techniques, we provide the first comprehensive analysis thereof. Specifically, we conducted a security analysis of various recently proposed CFI solutions (including kBouncer, ROPGuard, ROPecker, and CFI for COTS binaries). Our key contribution is in demonstrating that these techniques can be effectively undermined even when all their protection mechanisms are combined. In particular, we transformed existing (publicly available) exploits against Windows (which are detected by Microsoft EMET) into more stealthy attacks that bypass all recently proposed CFI techniques. We show that our performed transformations require no specific assumptions, and demonstrate that a 1MB Windows library (kernel32.dll) is already sufficient to derive a (turing-) complete gadget set using only call-preceded gadgets.

presented by

Daniel Lehmann & Ahmad-Reza Sadeghi

The BEAST Wins Again: Why TLS Keeps Failing to Protect HTTP

SSL has been around for decades and yet it keeps happening: new attacks are being discovered against TLS at a steady rate. The past year has seen its share of rogue CA certificates and critical vulnerabilities in TLS libraries that we have come to expect. In this talk, I will present no less than three new attacks against the use of TLS on the web. The first one relies on a long-known cryptographic weakness in the protocol that can be combined with long-known issues in TLS implementations to re-enable a flavor of the 2009 renegotiation attack that was thought to be fixed. The second one exploits the truncation weakness known since SSL2 but left unsolved to bypass anti-stripping defenses (strict transport security) and steal secure cookies. The last one exploits vulnerabilities in the deployment of HTTPS, in particular, how HTTP servers process requests and manage certificates and sessions, to reach the holy grail of TLS attacks: full server impersonation of several thousands of websites, including Microsoft, Apple, Twitter, PayPal. The three attacks have strong common points: they rely on an attacker that operates both at the TLS and HTTP levels, and they exploit misunderstandings and false assumptions between TLS libraries and applications.

In the course of this talk, you will learn about the full capabilities of the "beastly" attacker that operates jointly at the transport and application levels and how they can be exploited. You will also learn how to configure your HTTPS server to avoid being vulnerable to our virtual host confusion attacks, for which no simple universal fix exists. Lastly, I will try to disprove some misconceptions about TLS and privacy in the context of powerful network attackers.

presented by

Antoine Delignat-Lavaud

The Big Chill: Legal Landmines that Stifle Security Research and How to Disarm Them

Security research is a dangerous business.

The threat of lawsuits or even prosecution hangs heavy over the heads of white hat hackers as well as black hats. From Dmitry Skylarov being prosecuted for cracking ebook crypto back in 2001, to Weev being prosecuted today for exposing flaws in AT&T;'s website security, the legal landscape is littered with potential landmines for those trying to improve Internet and software security. When a major company like Google can be sued for billions over its interception of unencrypted WiFi signals, what's a wireless security researcher to do? When an Internet luminary like Aaron Swartz can be threatened with decades of jail time for his open data activism, what's your average pen tester supposed to think? How serious are these threats - and what can researchers do to avoid them, and maybe even fix the law?

Two veteran digital rights lawyers - one who counsels companies and defends hackers, and another who is an expert in the DC policy game - and the lead strategist of a major security firm will use a game show format to share examples of legally risky research and ask the question: "Computer Crime or Legitimate Research?" Using the answer to that question, we'll start gaming out how to craft legislation that would provide a sensible security research exception to laws like the Wiretap Act, the Digital Millennium Copyright Act, and the Computer Fraud and Abuse Act.

presented by

Trey Ford & Marcia Hofmann & Kevin Bankston

The Devil Does Not Exist - The Role of Deception in Cyber

While it might be convenient to think of cyberadversaries as ones and zeros, the reality is that systems are attacked and defended by human beings. As a result, it is important to understand the role deception plays in network operations. This presentation draws upon traditional and emerging research on deception and associated game theories to help the audience understand how attackers might deceive them, how to recognize that deception, and how defenders can also deceive their attackers.

presented by

Mark Mateski & Matt Devost

The Library of Sparta

On today's increasingly militarized Internet, companies, non-profits, activists, and individual hackers are forced to melee with nation-state class adversaries. Just as one should never bring a knife to a gun fight, a network defender should not rely on tired maxims such as "perimeter defense" and "defense in depth." Today's adversaries are well past that. This talk teaches you how to tap what we call the Library of Sparta - the collective written expertise codified into military doctrine. Hidden in plain sight, vast free libraries contain the time-tested wisdom of combat at the tactical, operational, and strategic levels. This is the playbook nation-state adversaries are using to target and attack you. This talk will help you better understand how adversaries will target your organization, and it will help you to employ military processes and strategies in your defensive operations. These techniques scale from the individual and small team level all the way up to online armies. This talk isn't a dry index into the library of doctrine, we provide entirely new approaches and examples about how to translate and employ doctrinal concepts in your current operations.

Many people in the computer security community use words like "OPSEC," "Kill Chain," and "intelligence-driven" without fully understanding the underlying concepts. Even worse, many show their ignorance by using military jargon incorrectly, thereby alienating clients, customers, and colleagues. These concepts are powerful and should not be ignored, but they must be well understood before they can be leveraged in your network.

This talk will include topics such as deception, electronic warfare, operations security, intelligence preparation of the battlefield, human intelligence collection, targeting, psychological operations, information operations, maneuver, and military cryptanalysis, among numerous others. Conventional wisdom at Black Hat is that that attacker will always win. Attackers have a clear intelligence advantage over defenders when it comes to vulnerabilities, malware, and open source information. A key point of the talk will be helping defenders generate the intelligence, information, and disinformation advantage necessary to turn the tables. You will leave this talk with an entirely new arsenal of military-grade strategies that will help you advance your work beyond the individual and small-team level and will prepare you to take on the most advanced adversaries.

presented by

David Raymond & Greg Conti & Tom Cross

The New Page of Injections Book: Memcached Injections

Memcached is a distributed memory caching system. It is in great demand in big-data Internet projects as it allows reasonably sped up web applications by caching data in RAM. Cached data often includes user sessions and other operational information.

This talk is based on research of different memcached wrappers to popular web application development platforms, such as Go, Ruby, Java, Python, PHP, Lua, and .NET. The primary goal is

determining input validation issues at key-value data which could be used to inject arbitrary commands to memcached protocol.

As a result, the Speaker found a way to do something like "SQL Injection attacks," but on memcached service. Such an attack in practice leads to different effects from authentication bypass to execution of arbitrary interpreter's code. It's a real world problem found on security audits and exists on different popular web applications.

presented by

Ivan Novikov

The New Scourge of Ransomware: A Study of CryptoLocker and Its Friends

In March of this year, a Romanian man killed himself and his 4-year old son because of a ransomware he received after visiting adult websites. This "police impersonation" malware instructed him to pay a massive fine or else go to jail for 11 years. Ransomware isn't a new threat; however, it introduced new life with CryptoLocker, the very first variant to perform encryption correctly, thussignificantly inhibiting security researchers and their typical countermeasures. Due to its unique nature, CryptoLocker is one of the few current malware campaigns that spawned its own working group focused around remediation. As time progressed, other ransomware copycat campaigns emerged, some of which got media attention even though they were nothing but vaporware.

This talk will focus on what the threat intelligence community did in response to this threat, including the development of near-time tracking of its infrastructure and what can be learned in order to manage new threats as they emerge.

presented by

Lance James & John Bambenek

The State of Incident Response

The last of the protection-detection-response triad to get any real attention, incident response is big business these days. I plan on stepping back and looking at both the economic and psychological forces that affect incident response as both a business and a technical activity. Nothing seems to be able to keep sufficiently skilled and motivated attackers out of a network. Can incident response save the day?

presented by

Bruce Schneier

Thinking Outside the Sandbox - Violating Trust Boundaries in Uncommon Ways

Attacking the modern browser and its plugins is becoming harder. Vendors are employing numerous mitigation technologies to increase the cost of exploit development. An attacker is now forced to uncover multiple vulnerabilities to gain privileged-level code execution on his targets. First, an attacker needs to find a vulnerability, leak an address to get around ASLR, and bypass DEP to gain code execution within the renderer process. The attacker then needs to bypass the application sandbox to elevate his privileges, which will allow him to do something interesting. Our journey begins at the sandbox and investigates some of the more obscure techniques used to violate this trust boundary.

What should you focus on when you are auditing a sandbox implementation? There are the traditional approaches: find a memory corruption vulnerability in IPC message handling, attack the kernel to get SYSTEM-level privilege escalation, or abuse shared memory regions. Sure, any of these will work but they may not be the easiest way. Our presentation will examine four bypass techniques successfully used in winning entries at this year's Pwn2Own contest. We will analyze the attack vector used, root causes, and possible fixes for each technique. These uncommon, yet highly effective, approaches have been used to bypass the most advanced application sandboxes in use today, and understanding them will provide a unique perspective for those working to find and verify such bypasses.

presented by

Brian Gorenc & Jasiel Spelman

Threat Intelligence Library - A New Revolutionary Technology to Enhance the SOC Battle Rhythm!

Cyber indicators are the 'new-er' detection strategy to help dismantle adversarial assaults and the volume of crowdsourced and private community malicious IOCs grows exponentially every day forcing the security industry to create a new must have tool - a threat library. The effectiveness of every SOC is based on their ability to discover, ingest, analyze, respond to, and pivot off threat intelligence and, historically, an ad-hoc spreadsheet combined with a day of analyst muscle was manageable to maintain and chase IOCs. However, over the past several years, as crowdsourcing intelligence has become mainstream, the volume of IOCs released by cyber intelligence providers (commercial and public do-gooders), industry blogs, malware repositories, vendor whitepapers, and open source intelligence (OSINT) has turned the spreadsheet firedrill into a bottleneck of operational inefficiencies amongst the typical workflows within an adversary hunting SOC. This discussion will provide a first-hand operational look from within a large 30+ team DIB SOC and explore the evolution of IOCs, associated SOC workflows, assess IOC overlap by Source, discuss several tools that help manage threat intelligence, and finally some hindsight implementation lessons learned.

presented by

Ryan Trost

Time Trial: Racing Towards Practical Timing Attacks

Attacks on software become increasingly sophisticated over time and while the community has a good understanding of many classes of vulnerabilities that are commonly exploited, the practical relevance of side-channel attacks is much less understood.

One common side-channel vulnerability that is present in many web applications today are timing side-channels which allow an attacker to extract information based on different response times. These side-channel vulnerabilities are easily introduced wherever sensitive values such as credentials or API keys are processed before responding to a client. Even though there is basic awareness of timing side-channel attacks in the community, they often go unnoticed or are flagged during code audits without a true understanding of their exploitability in practice.

In this talk, we provide both a tool 'time trial' and guidance on the detection and exploitability of timing side-channel vulnerabilities in common web application scenarios. Specifically, the focus of our presentation is on remote timing attacks, which are performed over a LAN, in a cloud environment, or on the Internet. To illustrate this, we first present experimental timing results that demonstrate how precisely timing can be measured and, more importantly, which timing differences can be distinguished remotely. Second, we compare our results with timing differences that are typically encountered in modern web frameworks and servers. The discussed attack scenarios include database queries, message authentication codes, web API keys, OAuth tokens, and login functions.

Our presentation has significance for a wide spectrum of the conference audience. Attendees in defensive security roles will gain a better understanding of the threat timing side-channel vulnerabilities pose and, based on the demonstrated attacks, will be better able to evaluate the severity and impact of a successful side-channel attack. Attendees in a penetration testing role will learn how to distinguish theoretical timing attacks from legitimately exploitable flaws by using our tool 'time trial'. Finally, attendees focused on research implications will receive a comprehensive update on the state-of-the-art in exploiting timing attacks in practice.

presented by

Daniel Mayer & Joel Sandin

Understanding IMSI Privacy

It is said that 80% of the world's population now has a mobile phone. They use mobile devices to make calls, send SMS messages, and to access the Internet via the cellular network infrastructure. End-users carrying mobile phones 24 hours trust cellular network operators and believe that the provided mobile communication link is secure.

However, on the other hand, mobile operators, device manufacturers, OS providers, and baseband suppliers do little to provide best security and privacy features to them. In particular, security capabilities of mobile communications are not shown to the end-users. Hence, it is easy for malicious attackers to mount subsequent attacks using IMSI catcher equipments. Further some hidden features, for example 'silent SMS', are supported in currently used mobile telephony systems but not notified to the end-users when in use. Attackers or illegitimate agencies exploit this weakness to track user movements regularly without the user's consent.

In this talk, we address these long-standing issues by developing a low-cost, easy-to-use privacy framework based on Android OS. We demonstrate our effort to build an ideal way to protect user privacy. A live demo of framework detecting hidden (in) security features of mobile communication system will be provided.

presented by

Ravishankar Borgaonkar & Swapnil Udar

Understanding TOCTTOU in the Windows Kernel Font Scaler Engine

The Font Scaler engine is widely used in Microsoft Windows and Mac OS operating systems for rendering TrueType/OpenType fonts. It was first introduced in 1989. Later, to improve the performance of the Windows NT operating system, Microsoft decided to move the engine from user mode to kernel mode. This enhancement does improve the performance, but it also brings security issues. Specifically, Font Scaler engine represents a significant kernel attack surface, and it is perhaps the most easily accessible point which can be reached remotely. For example, the famous Duqu malware well demonstrated vulnerabilities in this engine in 2011.

Many things make the font engine vulnerable. Such as the complexity of font file format, the enhancement of the Font Scaler engine (i.e., moving from user mode to kernel), the assumptions about the interactions between the font engine and its clients (win32k.sys), and the existence of font cache. Among these vulnerabilities, TOCTTOU (Time-of-Check to Time-of-Use) is the most critical type.

In this talk, I'm going to discuss the basic double fetch problem. Furthermore, I would like to present the more stealthy TOCTTOU vulnerability which is introduced by the design of the font engine.

presented by

Yu Wang

Unveiling the Open Source Visualization Engine for Busy Hackers

The way a human efficiently digests information varies from person-to-person. Scientific studies have shown that some individuals learn better through the presentation of visual/spatial

information compared to simply reading text. Why then do vendors expect customers to consume presented data following only the written word method as opposed to advanced graphical representations of the data? We believe this approach is dated.

To help the neglected visually inclined masses, we decided to create a free and Open Source engine to remove the complexity of creating advanced data visualizations. The ultimate goal of the project was to allow for the visualization of any loosely related data without having to endlessly reformat that data. For the visual/spatial learners, the engine will interpret their own data, whether it be a simple or complex system, and present the results in a way that their brains can understand.

Learning, for visual-spatial learners, takes place all at once, with large chunks of information grasped in intuitive leaps, rather than in the gradual accretion of isolated facts or small steps. For example, a visual-spatial learner can grasp all of the multiplication facts as a related set in a chart much easier and faster than memorizing each fact independently. We believe that some security practitioners might be able to better utilize their respective data sets if provided with an investigative model that their brains can understand.

During this presentation, we will show you how you can take any relational data set, quickly massage the format, and visualize the results. We will also share some observations and conclusions drawn from the results of the visualization that may not have appeared in simple text form. We have used this engine within OpenDNS to track CryptoLocker and CryptoDefense ransomware, Red October malware, and the Kelihos botnet. Additionally, specific Syrian Electronic Army (SEA) campaigns, carding sites, and even a map of the Internet via Autonomous Systems have been visualized using the engine.

Interesting data can also be isolated through the use of Python and JavaScript-based plugins that can be easily added to the engine's framework. These plugins affect the way the data is visualized and allow analysts to make sense of their data as it relates to the question they're trying to answer. The "big picture" model will help visually inclined incident responders, security analysts, and malware researchers visually stitch together complex data sets without needing a PhD in math or particle physics.

presented by

Thibault Reuille & Andrew Hay

Unwrapping the Truth: Analysis of Mobile Application Wrapping Solutions

One of the latest trends of BYOD solutions is to employ "Mobile Application Management (MAM)," which allows organizations to wrap existing applications to perform policy enforcement and data/transport security at the application layer rather than at the device level. Today's organizations face a complex choice: there are a plethora of BYOD application wrapping products on the market, each with their own colorful datasheets and hefty security claims. How well do

these BYOD application wrapping solutions stand up to their claims? And perhaps just as important, how well do they defend against real-life mobile threats?

In this talk we will analyze the application wrapping solutions offered by some of the major commercial BYOD products on the market today. We'll reverse engineer how these application wrapping solutions work for both iOS and Android; as well as, analyze their authentication, cryptography, interprocess communication (IPC), and client-side security control implementations. Finally, we'll explore the security vulnerabilities we've discovered in major vendor products that could result in the compromise of sensitive information.

presented by

Ron Gutierrez & Stephen Komal

VoIP Wars: Attack of the Cisco Phones

Many hosted VoIP service providers are using Cisco hosted collaboration suite and Cisco VoIP solutions. These Cisco hosted VoIP implementations are very similar; they have Cisco Unified Communication services, SIP protocol for IP Phones of tenants, common conference solutions, Skinny protocol for compliance, generic RTP implementation, VOSS Solutions product family for management services for tenants. Tenants use desktop and mobile clients to connect these services. Cisco hosted VoIP implementations and VoIP clients are vulnerable to many attacks, including:

- VLAN attacks,
- SIP trust hacking,
- Skinny based signaling attacks,
- Bypassing authentication and authorization,
- Call spoofing,
- Eavesdropping,
- Desktop/mobile client compromising
- Attacks against IP Phone management services; and,
- Web based vulnerabilities of the products

The presentation covers Skinny and SIP signaling attacks, 0day bypass technique for call spoofing and billing bypass, LAN attacks against supportive services for IP Phones, desktop and mobile phones, practical 0-day attacks against IP Phone management and tenant services. These attacks are available for desktop and mobile clients too, such as caller ID spoofing and fake messaging to compromise clients, fuzzing VoIP call signaling, MITM attacks and crashing mobile clients.

Attacking Cisco VoIP services requires limited knowledge today with the Viproy Penetration Testing Kit (written by the presenter). It has a dozen modules to test trust hacking issues, signaling attacks against SIP services and Skinny services, gaining unauthorized access, call spoofing, brute-forcing VoIP accounts and debugging services using as MITM. Furthermore, Viproy provides these attack

modules in a penetration testing environment and full integration. The presentation contains live demonstration of practical VoIP attacks and usage of new Viproy modules.

presented by

Fatih Ozavci

What Goes Around Comes Back Around - Exploiting Fundamental Weaknesses in Botnet C&C; Panels!

Bot herders deploy Command and Control (C&C;) panels for commanding and collecting exfiltrated data from the infected hosts on the Internet. To protect C&C; panels, bot herders deploy several built-in (software-centric) protection mechanisms to restrict direct access to these C&C; panels. However, there exist fundamental mistakes in the design and deployment of these C&C; panels that can be exploited to take complete control. This talk discusses about the methodology of launching reverse attacks on the centralized C&C; panels to derive intelligence that can be used to build automated solutions. This research reveals how to detect vulnerabilities and configuration flaws in the remote C&C; panels and exploit them by following the path of penetration testing. This talk is derived from the real time research in which several C&C; panels were targeted and intelligence was gathered to attack the next set of C&C; panels. A number of case studies will be discussed to elaborate step-by-step process of attacking and compromising C&C; panels. This talk also demonstrates the use of automated tools authored for making the testing easier for the researchers.

presented by

Aditya K Sood

When the Lights Go Out: Hacking Cisco EnergyWise

Energy Management Protocols (EMPs) are used in a variety of devices and environments. Their purpose is always the same: Controlling and measuring the energy consumption of connected devices. However, most EMPs are designed and implemented for embedded, non-IP environments, such as HDMI or home automation networks.

Cisco EnergyWise is a proprietary, closed-source protocol that brings EMPs to the main stream IP networks (e.g. by including EnergyWise clients in widely used notebooks and phones). The resulting broad deployment in a high number of environments, such as office networks (for example, ThinkPad notebooks include an EnergyWise Client in the default configuration) or even data centers (as power consumption is always a huge issue), leads to the potential to cause huge blackouts if EnergyWise is misconfigured or contains vulnerabilities which can be abused.

In this talk, we will describe our results on the EnergyWise architecture and protocol specification, present the reverse-engineered proprietary protocol , and show how you can hijack enerygwise

domains in order to perform DoS service attacks. In addition, we will release our toolkit that implements all of the presented attacks.

presented by

Matthias Luft & Ayhan Soner Koca

Why Control System Cyber-Security Sucks...

Since the 2010's "Stuxnet" sabotage attempt, cyber-security of industrial control systems (ICS) or "SCADA" has become a buzzword in industry. The (cyber-) protection of the critical infrastructure became a focal point for governments. Vendors and manufacturers have pushed "Industrial Security" appliances onto the market, or claim that their products are now with "enhanced security". A cacophony of standards have emerged, and certification schemes are offered. But does this help? Given the increasing interconnectivity of ICS (SmartMeters, later the Internet-of-Things), shouldn't the direction be more towards standard IT than sticking to a dedicated ICS IT? Why is it that I can patch a computer centre over night, but not a control system within a year? This presentation will not give the answers but outline why control system cyber-security sucks and which hurdles we encountered to handle ICS cyber-security like that of our computer centres¦ A change of paradigm is needed, and this change must start with people and not with technology.

presented by

Dr. Stefan Lders

Why You Need to Detect More Than PtH

Compromised credentials are a key predatory weapon in the attackers arsenal, and this isn't changing in the foreseeable future. This talk will systematically explore why they can be prevented but never cut off completely, and how to leverage this knowledge in detection. In closing, we will pick apart IoCs focused on Pass-the-Hash (PtH), while detailing more efficient detection techniques focused on misused, donated, or otherwise compromised credentials.

presented by

Matthew Hathaway & Jeff Myers

Windows Kernel Graphics Driver Attack Surface

Ever wondered about the attack surface of graphics drivers on Windows? Are they similar to other drivers? Do they expose ioctl's? In this talk, all those questions will be answered and more. Whether you're a security researcher, a developer looking for some security guidance when

writing these drivers, or just generally curious about driver internals, there's something here for all.

The research done focuses both on C/C++ code when available, as well as reverse engineering of these drivers.

presented by

Ilja van Sprundel

Write Once, Pwn Anywhere

Modern Windows use mitigation techniques such as DEP and ASLR to mitigate exploitation. The combination of ASLR and DEP have been proven to be a solid shield in most cases. Mitigation bypass is always one of the hottest topic in the security community.

This presentation contains two kind of new DEP bypass techniques, two kind of new ASLR bypass techniques, and many lesser known exploration skills. These techniques don't need ROP, JIT, third-party plugins or Non-ASLR modules. They are OS-independent, even CPU-independent in some cases. So exploits can easily "Write Once, Pwn Anywhere" now.

These techniques are fairly different from traditional exploit technique. So they may also be difficult to detect and identify if you don't know them.

presented by

Yang Yu
