

The attached DRAFT document (provided here for historical purposes) has been superseded by the following publication:

Publication Number: **NIST Interagency Report 8151**

Title: *Dramatically Reducing Software Vulnerabilities: Report to the White House Office of Science and Technology Policy*

Publication Date: **November 2016**

- Final Publication: <https://doi.org/10.6028/NIST.IR.8151> (which links to <http://nvlpubs.nist.gov/nistpubs/ir/2016/NIST.IR.8151.pdf>).
- Information on other NIST cybersecurity publications and programs can be found at: <http://csrc.nist.gov/>

The following information was posted with the attached DRAFT document:

Oct 04, 2016

NISTIR 8151

DRAFT Dramatically Reducing Software Vulnerabilities: Report to the White House Office of Science and Technology Policy

NIST invites comments on Draft NIST Internal Report (NISTIR) 8151, *Dramatically Reducing Software Vulnerabilities -- Report to the White House Office of Science and Technology Policy*. The call for a dramatic reduction in software vulnerability is heard from numerous sources, recently from the February 2016 Federal Cybersecurity Research and Development Strategic Plan. The plan defines goals for reducing vulnerabilities in the near, mid and long term. This report addresses the first mid-term goal.

Email comments to: paul.black <at> nist.gov

Comments due by: October 18, 2016

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16

Dramatically Reducing Software Vulnerabilities

Report to the White House Office of Science and Technology Policy

Paul E. Black
Lee Badger
Barbara Guttman
Elizabeth Fong

Dramatically Reducing Software Vulnerabilities

Report to the White House Office of Science and Technology Policy

Paul E. Black
Lee Badger
Barbara Guttman
Elizabeth Fong
Information Technology Laboratory

October 2016



U.S. Department of Commerce
Penny Pritzker, Secretary

National Institute of Standards and Technology
Willie May, Under Secretary of Commerce for Standards and Technology and Director

46
47
48
49National Institute of Standards and Technology Interagency Report 8151
50 pages (October 2016)

50 Certain commercial entities, equipment, or materials may be identified in this document in order to describe an
51 experimental procedure or concept adequately. Such identification is not intended to imply recommendation or
52 endorsement by NIST, nor is it intended to imply that the entities, materials, or equipment are necessarily the best
53 available for the purpose.

54 There may be references in this publication to other publications currently under development by NIST in accordance
55 with its assigned statutory responsibilities. The information in this publication, including concepts and methodologies,
56 may be used by federal agencies even before the completion of such companion publications. Thus, until each
57 publication is completed, current requirements, guidelines, and procedures, where they exist, remain operative. For
58 planning and transition purposes, federal agencies may wish to closely follow the development of these new
59 publications by NIST.

60 Organizations are encouraged to review all draft publications during public comment periods and provide feedback to
61 NIST. Many NIST cybersecurity publications, other than the ones noted above, are available at
62 <http://csrc.nist.gov/publications>.

63

64

65 **Public comment period: *October 4, 2016 through October 18, 2016***

66 National Institute of Standards and Technology
67 Attn: Computer Security Division, Information Technology Laboratory
68 100 Bureau Drive (Mail Stop 8930) Gaithersburg, MD 20899-8930
69 Email: paul.black@nist.gov (SUBJECT="DRAFT NISTIR 8151 Comments")

70 All comments are subject to release under the Freedom of Information Act (FOIA).

71

72 Abstract

73 The call for a dramatic reduction in software vulnerability is heard from multiple sources,
74 recently from the February 2016 Federal Cybersecurity Research and Development Strategic
75 Plan. This plan starts by describing well known risks: current systems perform increasingly vital
76 tasks and are widely known to possess vulnerabilities. These vulnerabilities are often easy to
77 discover and difficult to correct. Cybersecurity has not kept pace and the pace that is needed is
78 rapidly accelerating. The goal of this report is to present a list of specific approaches that have
79 the potential to make a dramatic difference in reducing vulnerabilities – by stopping them before
80 they occur, by finding them before they are exploited or by reducing their impact.

81

82 Keywords:

83 Measurement; metrics; software assurance; security vulnerabilities; reduce software
84 vulnerability.

85 Acknowledgements:

86 Extravagant thanks to Joshi Rajeev rajeev.joshi@jpl.nasa.gov for contributions to Sect. 2.3
87 Additive Software Analysis Techniques.

88 Effusive thanks to W. Konrad Vesey (william.k.vesey.ctr@mail.mil), Contractor, MIT Lincoln
89 Laboratory, Office of the Assistant Secretary of Defense, Research and Engineering, for much of
90 the material in Sect. 2.5 Moving Target Defense (MTD) and Artificial Diversity. Much of the
91 wording is directly from a private communication from him.

92 **Table of Contents**

| | | | |
|-----|-------|---|----|
| 93 | 1 | Introduction | 1 |
| 94 | 1.1 | SCOPE of REPORT | 2 |
| 95 | 1.2 | METRICS | 3 |
| 96 | 1.3 | METHODOLOGY | 3 |
| 97 | 1.4 | REPORT ORGANIZATION..... | 4 |
| 98 | 2 | Technical Approaches..... | 4 |
| 99 | 2.1 | Formal Methods | 5 |
| 100 | 2.1.1 | Rigorous Static Program Analysis | 5 |
| 101 | 2.1.2 | Model Checkers, SAT Solvers and Other “Light Weight” Decision Algorithms .. | 6 |
| 102 | 2.1.3 | Directory of Verified Tools and Verified Code | 7 |
| 103 | 2.1.4 | Pragmas, Assertions, Pre- and Postconditions, Invariants, Properties, Contracts | |
| 104 | | and Proof Carrying Code | 7 |
| 105 | 2.1.5 | Correct-by-Construction and Model-Based Development | 8 |
| 106 | 2.2 | System Level Security | 9 |
| 107 | 2.2.1 | Operating System Containers | 11 |
| 108 | 2.2.2 | Microservices..... | 11 |
| 109 | 2.3 | Additive Software Analysis Techniques..... | 13 |
| 110 | 2.3.1 | Software Information Expression and Exchange Standards..... | 14 |
| 111 | 2.3.2 | Tool Development Framework or Architecture..... | 15 |
| 112 | 2.3.3 | Combining Analysis Results..... | 16 |
| 113 | 2.4 | More Mature Domain-Specific Software Development Frameworks | 18 |
| 114 | 2.4.1 | Rapid Framework Adoption | 20 |
| 115 | 2.4.2 | Compositional Testing | 21 |
| 116 | 2.4.3 | Conflict Resolution in Multi-Framework Composition..... | 21 |
| 117 | 2.5 | Moving Target Defenses (MTD) and Artificial Diversity | 22 |
| 118 | 2.5.1 | Compile-Time Techniques..... | 22 |
| 119 | 2.5.2 | System or Network Techniques | 23 |
| 120 | 2.5.3 | Operating System Techniques | 23 |
| 121 | 3 | Measures and Metrics | 25 |
| 122 | 3.1 | A Taxonomy of Software Metrics | 26 |
| 123 | 3.2 | Software Assurance: The Object of Software Metrics | 28 |

124 3.3 Software Metrology 28

125 3.4 Product Metrics 29

126 3.4.1 Existing Metrics 30

127 3.4.2 Better Code 31

128 3.4.3 Metrics and Measures of Binaries and Executables 31

129 3.4.4 More Useful Tool Outputs 31

130 4 Summary and Community Engagement33

131 4.1 Engaging the Research Community..... 33

132 4.1.1 Grand Challenges, Prizes and Awards..... 33

133 4.1.2 Research Infrastructure 33

134 4.2 Education and Training..... 34

135 4.3 Consumer-Enabling Technology Transfer..... 35

136 4.3.1 Government Contracting and Procurement 35

137 4.3.2 Liability..... 35

138 4.3.3 Insurance 35

139 4.3.4 Vendor-Customer Relations..... 35

140 4.3.5 Standards..... 36

141 4.3.6 Code Repositories 36

142 4.4 Conclusion 36

143 5 References38

144

145 **1 Introduction**

146 The call for a dramatic reduction in software vulnerability is being heard from multiple sources,
147 including the February 2016 Federal Cybersecurity Research and Development Strategic Plan
148 [FCRDSP16]. This plan starts by describing a well-known risk: current systems perform
149 increasingly vital tasks and are widely known to possess vulnerabilities. These vulnerabilities are
150 often easy to discover and difficult to correct. Cybersecurity has not kept pace and the pace that
151 is needed is rapidly accelerating. The plan defines goals for the near, mid and long term. This
152 report addresses the first mid-term goal:

153
154 *Achieve S&T advances to reverse adversaries' asymmetrical advantages, through*
155 *sustainably secure systems development and operation. This goal is two-pronged: first,*
156 *the design and implementation of software, firmware, and hardware that are highly*
157 *resistant to malicious cyber activities (e.g., software defects, which are common, give rise*
158 *to many vulnerabilities)*
159

160 Since it is central to the purpose of this report, we define what we mean by “vulnerability.” A
161 vulnerability is a property of system security requirements, design, implementation or operation
162 that could be accidentally triggered or intentionally exploited and result in a violation of desired
163 system properties. A vulnerability is the result of one or more weaknesses in requirements,
164 design, implementation or operation [Black11a]. This definition excludes

- 165 • operational problems, such as installing a program as world-readable or setting a trivial
166 password for administrator access.
- 167 • insider malfeasance, such as exfiltration ala Snowden.
- 168 • functional bugs, such as the mixture of SI and Imperial units, which led to the loss of the
169 Mars Climate Orbiter in 1999 [Ober99].
- 170 • purposely introduced malware or corrupting “mis-features” in regular code, such as
171 allowing root access by user names like “JoshuaCaleb.” We exclude this vulnerability,
172 because it is intentionally inserted. One assumes that a bad actor will fashion it to pass
173 review/quality control processes.
- 174 • software weaknesses that cannot be exploited (by “outsiders”) as a result of input
175 filtering or other mitigations.

176
177 Great strides have been made in defining software vulnerabilities, cataloging them and
178 understanding them. Additionally, great strides have been made in educating the software
179 community about the vulnerabilities, attendant patches and underlying weaknesses. This work,
180 however, is insufficient. Significant vulnerabilities are found routinely, many vulnerabilities lie
181 undiscovered for years and patches are often not applied. Clearly a different approach – one that
182 relies on improving software – is needed.

183

184 *Strengthening protection requires increasing assurance that the products people develop*
185 *and deploy are highly resistant to malicious cyber activities, because they include very*
186 *few vulnerabilities....* [FCRDSP16, p 17]
187

188 1.1 **SCOPE of REPORT**

189 The goal of this report is to present a list of specific approaches that have the potential to make a
190 dramatic difference reducing vulnerabilities – by stopping them before they occur, by finding
191 them before they are exploited or by reducing their impact.
192

- 193 • Stopping vulnerabilities before they occur generally includes improved methods for
194 specifying and building software.
- 195 • Finding vulnerability includes better testing techniques and more efficient use of multiple
196 testing methods.
- 197 • Reducing the impact of vulnerabilities refers to techniques to build architectures that are
198 more resilient, so that vulnerabilities cannot be meaningfully exploited.

199
200 The report does not segregate the approaches into these three bins, since some approaches may
201 include pieces from multiple bins.
202

203 The list of approaches for reducing vulnerabilities focuses on approaches that meet three criteria:

- 204 1. Dramatic impact
- 205 2. 3 to 7-year timeframe
- 206 3. Technical activities

207
208 *Dramatic.* This means reducing exploitable vulnerabilities by two orders of magnitude.
209 Estimates of software vulnerabilities are up to 25 errors per 1 000 lines of code [McConnell04,
210 page 521]. These approaches have been selected for the possibility of getting to 2.5 errors per 10
211 000 lines of code. The ability to measure whether an approach has a dramatic impact requires the
212 ability to measure it. Measuring software quality is a difficult task. A parallel effort on
213 improvements for measuring software vulnerabilities was pursued.
214

215 *3 to 7-year timeframe.* This timeframe was selected, because it is far enough out to make
216 dramatic changes, based on existing techniques, but not having reached their full potential for
217 impact. It is a timeframe that it is reasonable to speculate about. Beyond this timeframe, it is too
218 difficult to predict what new technologies and techniques will be developed, potentially making
219 their own set of dramatic changes on how IT is used. In the near future, the emphasis will be on
220 implementing techniques that are already being deployed.
221

222 *Technical.* There are many different types of approaches to reducing software vulnerabilities,
223 many of which are not primarily technical – from helping users meaningfully request security to

224 funding research and operational activities and training all parties, who design, build, test and
225 use software. During the development of this report, many ideas were put forward across this
226 broad span. The report only addresses technical approaches in order to have a manageable scope,
227 which builds on expertise available during the development of the report. These other areas are
228 critical, too.

229
230 During the drafting of this report, many excellent ideas were brought forth that are outside the
231 scope of this report and are summarized in Section 4 under Community Engagement. Examples
232 of these activities include:

- 233 • Improved funding
- 234 • Improving education
- 235 • More research for various aspects of software understanding
- 236 • Increased use of grand challenges and competitions
- 237 • Providing better methods for consumers of software to ask for and evaluate lower-
238 vulnerability software

239
240 This report excludes a discussion of vulnerabilities in firmware and hardware. This is not to say
241 that these are not critical. These can be addressed in another report. This report targets a broad
242 range of software, including government-contracted software, commercial and open source
243 software. It covers software used for general use, mobile devices and embedded in appliances
244 and devices. The goal is to prevent vulnerabilities in new code, in addition to identifying and
245 fixing vulnerabilities in existing code.

246

247 **1.2 METRICS**

248 There are multiple efforts to define software vulnerabilities, their prevalence, their detectability
249 and the efficacy of detection and mitigation techniques. The ability to measure software can play
250 an important role in dramatically reducing software vulnerabilities. Industry requires evidence of
251 the extent of such vulnerabilities, in addition to knowledge in determining which techniques are
252 most effective in developing software with far few vulnerabilities. Additionally, and more
253 critically, industry requires guidance in identifying the best places in code to deploy mitigations
254 or other actions. This evidence comes from measuring, in the broadest sense, or assessing the
255 properties of software.

256

257 **1.3 METHODOLOGY**

258 In order to produce the list of approaches, the Office of Science and Technology Policy asked
259 NIST to lead a community-based effort. NIST consulted with multiple experts in the software
260 assurance community including:

- 261 • Two Office of Science and Technology Policy (OSTP)-hosted inter-agency roundtables
- 262 • Half day session at the Software and Supply Chain Assurance Summer Forum

- 263 • Full day workshop on Software Measures and Metrics to Reduce Security Vulnerabilities
264 • Public comment 4-18 October 2016
265

266 **1.4 REPORT ORGANIZATION**

267 The report is organized into two major sections. The first enumerates technical approaches and
268 the second addresses metrics.

269 Section 2 covers technical approaches in dealing with vulnerabilities in software. These include
270 formal methods, such as rigorous static program analyses, model checkers and SAT solvers. It
271 also suggests having a directory of verified tools and verified code. This section addresses
272 system level security, including operating system containers and microservices. Additive
273 software analysis techniques are addressed. Finally, it discusses moving target defenses (MTD)
274 and artificial diversity. These include compile-time time techniques, system or network
275 techniques and operating system techniques.

276 Each subsection follows the same format:

- 277 • Definition and Background: Definition of the area and background
- 278 • Maturity Level: How mature the area is, including a discussion of whether the approach
279 has been used in the “real world” or just in a laboratory and issues related to scalability
280 and usability.
- 281 • Basis for Confidence: Rationale for why this could work
- 282 • Rational for potential impact
- 283 • Further Reading, papers, other materials

284 Section 3 covers measures and metrics. It is designed to encourage the adoption of metrics and
285 other tools to address vulnerabilities in software. It addresses product metrics and how to
286 develop better code. It also addresses the criticality of software security and quality metrics.

287

288 **2 Technical Approaches**

289

290 There are many approaches at varying levels of maturity that show great promise for reducing
291 the number of vulnerabilities in software. This report highlights five of them that are sufficiently
292 mature and have shown success so that it is possible to extrapolate into a 3 to 7 year horizon.
293 This list is not an exhaustive list, but rather to show that it is possible to make significant
294 progress in reducing vulnerabilities and to lay out paths to achieve this ambitious goal.

295

296 **2.1 Formal Methods**

297 Formal methods include all software analysis approaches based on mathematics and logic,
298 including parsing, type checking, correctness proofs, model-based development and correct-by-
299 construction. Formal methods can help software developers achieve greater assurance that entire
300 classes of vulnerabilities are absent and can also help reduce unpredictable cycles of expensive
301 testing and bug fixing.

302 In the early days of programming, some practitioners proved the correctness of their programs.
303 As the use of software exploded and programs grew so large that purely manual proofs were
304 infeasible, formalized correctness arguments lost favor. In recent decades, developments such
305 Moore's law, multi-core processors and cloud computing make orders of magnitude more
306 compute power readily available. Advances in algorithms for solving Boolean Satisfiability
307 (SAT) problems, decision procedures (e.g., ordered binary decision diagrams OBDD) and
308 reasoning models (e.g., abstract interpretation and separation logic) dramatically slashed
309 resources required to answer questions about software.

310 By the 1990s, formal methods had developed a bad reputation as taking far too long, in machine
311 time, person years and project time, and requiring a PhD in computer science and mathematics to
312 use. It is not that way anymore. Formal methods are widely used today. For instance, compilers
313 use SAT solvers to allocate registers and optimize code. Operating systems use algorithms
314 formally guaranteed to avoid deadlock. These are what Kiniry and Zimmerman call [Kiniry08]
315 Secret Ninja Formal Methods: they are invisible to the user, except to report that something is
316 not right. In contrast to such "invisible" use of formal methods, overt use often requires recasting
317 problems into a form compatible with formal methods tools. Most proposed cryptographic
318 protocols are now examined with model checkers for possible exploits. Practitioners also use
319 model checkers to look for attack paths in networks.

320 Despite their strengths, formal methods are less effective if there is no clear statement of
321 software requirements or if what constitutes proper software behavior can only be determined by
322 human judgment or through balancing many conflicting factors. Thus we would not expect
323 formal methods to contribute much to the evaluation of the usability of a user interface,
324 development of exploratory software or unstructured problems.

325 Formal methods include many, many techniques at all stages of software development and in
326 many different application areas. We do not list every possibly helpful formal method. Instead,
327 we concentrate on a few that may contribute significantly in the medium term.

328 **2.1.1 Rigorous Static Program Analysis**

329 Static analysis is the examination of software for specific properties without executing it. For our
330 purposes, we only consider automated analysis. Heuristic analysis is faster than rigorous
331 analysis, but lacks assurance that comes from a chain of logical reasoning. Some questions can
332 only be answered by running the software under analysis, i.e., through dynamic analysis.

333 Combining static and dynamic analysis yields a hybrid technique. In particular, executions may
334 produce existence proofs of properties that cannot be confirmed using static techniques only.

335 Many representations of software (e.g., source code, executables, requirements) may be statically
336 analyzed. Source code analysis, however, is the most mature. Many tools have been developed to
337 analyze software written in specific programming languages. One advantage of source code
338 analysis is that the context of problems identified in source code can be communicated to
339 software developers using a representation (the code itself) that is comprehensible to people.
340 When other representations are analyzed, an additional step is required to render a problem into a
341 form that people can first understand and then relate to a program under analysis.

342 According to Doyle's assessment [Doyle16], rigorous static analysis is superior in terms of
343 coverage, scalability and benefit for effort. A limitation is that it is difficult to specify some
344 properties in available terms.

345 Formal methods have shown significant applicability in recent years. For example, the Tokeneer
346 project shows [Barnes06, Woodcock10] that software can in some cases be developed with
347 formal methods faster and cheaper and with fewer bugs than with traditional software
348 development techniques. TrustInSoft used Frama-C to prove [Bakker14, Regehr15] the absence
349 of a set of Common Weakness Enumeration (CWE) classes in PolarSSL, now known as mbed
350 TLS. This approach is commonly used, and even mandated, in Europe for software in
351 transportation and nuclear plant control.

352 These developments illustrate a few among the many uses of static analysis. Going forward,
353 static analysis has the potential to efficiently preclude several classes of errors in newly-
354 developed software and to reduce the uncertainty regarding resources needed to reach higher
355 levels of assurance through testing.

356 **2.1.2 Model Checkers, SAT Solvers and Other "Light Weight" Decision Algorithms**

357 These algorithms can answer questions about desirable higher level properties, such as that a
358 protocol only allows sensitive text to be read if one has a key, that security properties are
359 preserved by the system, that an assignment of values satisfies multiple constraints or that there
360 are no paths to breaches via (known) attacks. These algorithms can also be applied to analyze
361 detailed design artifacts, such as finite (and infinite) state machines.

362 Doyle's assessment [Doyle16] is that model checkers can have excellent coverage and many
363 properties can be represented. Since the effort required increases exponentially with problem
364 size, there is always an effectual size limit, however. Problems smaller than the limit can be
365 solved quickly. Very large problems may require excessive resources or intensive human work to
366 break the problem into reasonable pieces.

367 Such techniques can be applied in essentially two ways. First, they can be used as part of
368 software in production. For instance, instead of an ad-hoc routine to find an efficient route for a

369 delivery truck, an application can use a well-studied Traveling Salesman or spanning tree
370 algorithm. Second, and perhaps more pertinent to the theme of this report, is to use the
371 algorithms to design or verify software.

372 **2.1.3 Directory of Verified Tools and Verified Code**

373 Software developers often must expend significant effort to qualify tools or develop program
374 libraries with proven properties. Even when a later developer wishes to use the results of such
375 work, there are no central clearing houses to consult. A list of verified tools, carefully
376 constructed libraries and even reusable specifications and requirements can speed the adoption of
377 formal methods. Such a tool library could facilitate wider use, with accompanying assurance, of
378 software with dramatically reduced numbers of vulnerabilities.

379 Many companies and government agencies evaluate the same tools or the same software for
380 similar uses. Since there is no way to find out who may have done related evaluations, each
381 entity must duplicate the work, sometimes with less knowledge and care than another has already
382 applied. It is especially challenging since many contracts discourage sharing results. [Klass16] A
383 repository or list would be of great benefit. Knowing about related efforts, developers could
384 contribute to one effort, instead of working on their own.

385 For instance, the Open Web Application Security (OWASP) foundation coordinated a project to
386 develop a shared application program interface (API), called Enterprise Security API (ESAPI).
387 The ESAPI toolkit “encapsulate[s] the key security operations most applications need.”

388 See Section 2.4 for a discussion of re-use of well-tested and well-analyzed code.

389 **2.1.4 Pragmas, Assertions, Pre- and Postconditions, Invariants, Properties, Contracts and** 390 **Proof Carrying Code**

391 Programmers generally have a body of information that gives them confidence that software will
392 perform as expected. A neglected part of formal methods is to unambiguously record such
393 insights. Variations go by different terms, such as contracts, assertions, preconditions,
394 postconditions and invariants. It cost programmers some thought to state exactly what is going
395 on using a language similar to code expressions, but such statements help. These are activated
396 (“compiled in”) during development and testing, then may be deactivated before release.

397 The benefit is that these formal statements of properties carried in the code may be used to cross
398 check the code. For example, tests may be generated directly from assertions. They may be
399 activated to perform internal consistency checks during testing or production. Faults can
400 therefore be detected much earlier and closer to erroneous code, instead of having to track back
401 from externally visible system failures. Such statements also supply additional information to
402 perform semi-automated proofs of program correctness. Unlike comments, which may not be
403 updated when the code changes, these can be substantiated or enforced by a computer and
404 therefore must continue to be precise statements of program features and attributes.

405 A striking example of how such formal statements could help is the 1996 failure of the first
406 Ariane 5 rocket launched. The Ariane 5 used software from the successful Ariane 4. Analysis
407 showed that a 16-bit integer could handle Ariane 4 speeds. However, higher Ariane 5 speed
408 values overflowed the variable leading to computer shut down and the loss of the vehicle. If the
409 code had a precondition that the speed must fit in a 16-bit integer, “Any team worth its salt
410 would have checked ... [preconditions, which] would have immediately revealed that the Ariane
411 5 calling software did not meet the expectation of the Ariane 4 routines that it called.”
412 [Jézéquel97]

413 **2.1.5 Correct-by-Construction and Model-Based Development**

414 In model-based development, a software developer creates and modifies a model of a system.
415 Behavior may be specified in a higher-level or domain-specific language or model, and then
416 code is automatically generated. Much or all of the code is generated from the model. This is one
417 correct-by-construction technique. This and others, such as design by refinement, aim to entirely
418 avoid whole classes of vulnerabilities, since the developer rarely touches the code. Code
419 synthesis like this is useful in fewer situations than other formal methods. Such models or
420 specifications may also generate test suites or oracles. They may also be used to validate or
421 monitor system operation.

422 According to Doyle’s assessment [Doyle16], program synthesis has an “A+” in coverage, “B” in
423 effort and properties, but “D” in scalability. When we can specify complete high-level models
424 for entire systems, or even subsystems, we call them languages and cease to consider them
425 unusual, but they represent a very substantial use of formal methods.

426 **2.1.6 Maturity Level**

427 Formal methods are today used (relatively invisibly) throughout the world. One of the most
428 pervasive applications is the use of strong type checking within modern programming languages.
429 Other, admittedly limited, uses are the algorithms of various software checking tools, some of
430 them built into widely used development environments (e.g., that tag inconsistent use of
431 variables, missing values or use of unsafe interfaces). In 2010, researchers at NICTA
432 demonstrated [Klein14] the formal verification of the L4 microkernel comprising about 10000
433 lines of C code.

434 **2.1.7 Basis for Confidence**

435 Assertions, and to a lesser extent, contracts, have been significantly adopted in high-quality
436 software. Their gradual improvement to encompass more advanced condition and API checking
437 is likely because they have already proven themselves in some developer communities. Many
438 tools now perform static analysis. A natural progression is to promote more and more advanced
439 forms of static analysis. Software proving based on techniques such as pre- and post-condition
440 satisfaction and proof carrying code have seen initial adoption in critical software; they require
441 more effort and cost, however, in some use cases they have been shown cost effective in the long
442 run: fewer or no fixes to deployed systems.

443 **2.1.8 Rationale for Potential Impact**

444 The greatest potential impact is likely in costs avoided for components that, over time, become
445 heavily relied upon. The heartbleed debacle is an example of a modest code base with outsized
446 importance: a judicious use of formal methods might have avoided the problem in the first place.
447 Generally, higher quality software, such as can be produced using formal methods, can be used
448 to lower long-term maintenance and replacement costs of software components. As noted in
449 [Woody14], unlike physical systems that wear out and eventually fail with greater frequently,
450 software systems generate failures when they are incorrect and the flaws are triggered by
451 environmental factors.

452 **2.1.9 Further Reading**

453 [Armstrong14] Robert C. Armstrong, Ratish J. Punnoose, Matthew H. Wong and Jackson R.
454 Mayo, “Survey of Existing Tools for Formal Verification,” Sandia National Laboratories report
455 SAND2014-20533, December 2014. Available at [http://prod.sandia.gov/techlib/access-](http://prod.sandia.gov/techlib/access-control.cgi/2014/1420533.pdf)
456 [control.cgi/2014/1420533.pdf](http://prod.sandia.gov/techlib/access-control.cgi/2014/1420533.pdf)

457 [Bjørner16] Nikolaj Bjørner, “SMT Solvers: Foundations and Applications”, Dependable
458 Software Systems Engineering, J. Esparza et. al. eds., pp 24-32, IOS Press, 2016. DOI:
459 10.3233/978-1-61499-627-9-24

460 [Boulanger12] “Industrial Use of Formal Methods: Formal Verification”, Jean-Louis Boulanger
461 (Ed), July 2012, Wiley-ISTE.

462 [Voas16a] Jeffrey Voas and Kim Schaffer, “Insights on Formal Methods in Cybersecurity”,
463 IEEE Computer 49(5):102 – 105, May 2016, DOI: 10.1109/MC.2016.131
464 A roundtable about formal methods with seven experts on formal methods.

465 [Voas16b] Jeffrey Voas and Kim Schaffer, (Insights, part 2), IEEE Computer August 2016

466 [Woodcock09] Jim Woodcock, Peter Gorm Larsen, Juan Bicarregui and John Fitzgerald,
467 “Formal Methods: Practice and Experience”, ACM Computing Surveys, 41(4), October 2009,
468 Article No. 19, DOI: 10.1145/1592434.1592436. Available at
469 <http://homepage.cs.uiowa.edu/~tinelli/classes/181/Fall14/Papers/Wood09.pdf>

470 **2.2 System Level Security**

471 When software is executed, the system context for the running software defines the resources
472 available to the software, the APIs needed to access those resources and how the software may
473 access (and be accessed by) outside entities. These aspects of a system context may strongly
474 affect the likelihood that software contains vulnerabilities (e.g., complex or buggy APIs increase
475 the likelihood), the feasibility of an attacker exploiting vulnerabilities (e.g., more feasible if
476 system services are reachable from outside) and the impact an attack could have (e.g., both
477 damage to system resources and mission-specific costs).

478 A long-standing goal of system designers is to build systems that are resistant to attack and that
479 enforce desirable security policies on both programs and users. Started in 1965, the Multics
480 system [Corbato65] combined a number of ideas (e.g., virtual memory, multi-processing,
481 memory segments) to implement a computing utility that could protect information from
482 unauthorized access by programs and users. Starting in the 1970s, a number of security policy
483 models were introduced to formalize the security responsibilities of the system layer. In 1976,
484 the Bell-Lapadula (BLP) model [Bell76] provided a formal expression of mandatory security for
485 protecting classified information: the BLP model allowed “high” (e.g., SECRET) processes
486 access to “low” (e.g., UNCLASSIFIED) information for usability but prevented “low” processes
487 from accessing “high” information. The noninterference model of [Goguen84] accounted for
488 indirect information flows, also known as covert channels. Biba’s integrity model expressed
489 [Biba77] mandatory security for integrity: it prevented possibly-malicious (low-integrity) data
490 from being observed by high-integrity processes, thus reducing the risk that high-integrity
491 processing and data might become corrupted. The type enforcement model of [Boebert85]
492 provided a table-based access control mechanism to allow data to be transformed only by pre-
493 approved programs. These security policy models provided necessary clarity regarding desirable
494 security properties, but using the models in real-scale systems posed usability problems for
495 system administrators, and software implementations of the models still contained exploitable
496 flaws.

497 In 1999, DARPA started the Intrusion Tolerant Systems (ITS) program predicated on the notion
498 that systems can be built to operate through, or “tolerate,” even successful attacks. A number of
499 other research programs followed that built on this idea. [Tolerant07] Essential concepts
500 explored by these programs included the structuring of systems with redundant and diverse
501 components unlikely to all be subverted by a single vulnerability, the introduction of new policy-
502 enforcing software layers and the use of diagnostic reasoning components for automated
503 recovery. The DARPA research thrust in tolerant systems recognized that the elimination of all
504 vulnerabilities from real-world systems is an unlikely achievement for the foreseeable future.
505 The research demonstrated substantial tolerance in red team testing (e.g., see [Pal05]), but the
506 approaches also imposed significant configuration complexity, reduced execution speed and
507 significantly increased resource (cpu, memory, etc.) requirements.

508 Recent advances, both in hardware and software, raise the possibility of developing security-
509 enforcing and intrusion tolerant systems that are both performance and cost effective. Such
510 systems have the potential to suppress the harms that software vulnerabilities can cause. On the
511 hardware side, the low cost multicore and system-on-a-chip processors are lowering the costs of
512 redundancy. On the complementary software side, emerging architectural patterns are offering a
513 new opportunity to build security and tolerance into the next generation of systems. Among
514 numerous possible patterns, two that appear promising are operating system containers and
515 microservices.

516 **2.2.1 Operating System Containers**

517 “A container is an object isolating some resources of the host, for the application or system
518 running in it.” [LXC] A container is, in essence, a very light weight virtual machine whose
519 resources (memory, disk, network) can be very flexibly shared with a host computer or other
520 containers. A container provides some of the isolation properties of an independent computer,
521 but a container can be launched in a fraction of a second on commodity hardware.

522 Container-based isolation can clearly reduce the impact of software vulnerabilities if the isolation
523 is strong enough. Container configurations, however, are complex: they determine numerous
524 critical elements of a container, such as how it shares its resources, how its network stack is
525 configured, its initial process, the system calls it can use and more. Although the market has
526 already embraced management systems, such as Docker [Docker16], that support the sharing of
527 container configurations, there is a need for tools and techniques that can analyze container
528 configurations and determine the extent to which they reduce security risk, including, e.g., the
529 extent to which they can mitigate the effects of software vulnerabilities.

530 Additionally, containers offer an opportunity to apply some of the traditional security models and
531 intrusion tolerance techniques using building blocks that favor efficiency and ease of
532 deployment. There is now a new opportunity to reevaluate which advanced security models and
533 intrusion tolerance techniques can become mainstream technologies.

534 Furthermore, because a container can be efficiently wrapped around a single run of a program, a
535 container might be configured to grant a program only the minimum level of access to resources,
536 thus following the principle of least privilege [Saltzer75]. Least privilege is a fundamental
537 principle for limiting the effects of software vulnerabilities and attacks. It is notoriously difficult,
538 however, to specify the minimal resources that a program requires. Rather than trying to solve
539 the problem in its full generality, one strategy is to develop analysis techniques/tools to generate
540 custom container configurations that approximate least-privilege for important classes of
541 programs. Due to the relative ease of deploying containers, such tool-assisted containers could
542 bring much more effective access control and safety to mainstream systems.

543 **2.2.2 Microservices**

544 Microservices describe “An approach to designing software as a suite of small services, each
545 running in its own process and communicating with lightweight mechanisms.” [Fowler14] The
546 essential microservices idea is not new: it has been explored using web services and in operating
547 systems based on microkernels such as the Mach microkernel [Rashid86], the GNU Hurd
548 [Hurd16] and the Web Services Architecture [WSA04]. The microservices approach, however,
549 structures services according to different criteria. As explained in [Fowler14], microservices
550 should implement individual business (or mission) capabilities, have independent refresh cycles,
551 be relatively easy to replace and be programming-language agnostic. In short, each microservice
552 should make economic and management sense on its own. At the same time, microservices may
553 rely on one another, which can support well-defined modularity.

554 This approach to system structure can result in a number of components whose interfaces are
555 explicitly defined and whose dependencies are similarly explicitly defined.

556 As a system operates and the flow of control passes between microservices, there is a natural
557 incentive to “batch up” inter-service communications to amortize boundary-crossing overheads.
558 While this kind of batching can increase latencies in some cases, it can also simplify inter-
559 component dependencies and possibly reduce the likelihood of software flaws and hence
560 vulnerabilities.

561 The deployment of software as collections of microservices raises a fundamental question: does
562 it make sense to build a “trusted microservice”? Even more ambitiously, would it be feasible to
563 develop microservices that are themselves reference monitors? The reference monitor concept
564 dates from the 1972 Anderson Report [Anderson72] and refers to a system component that
565 mediates all accesses to resources that it provides. A reference monitor is: 1) always invoked, 2)
566 tamperproof and 3) verified (i.e., small enough to be built with high assurance). As microservices
567 are becoming increasingly popular, the time may be right to research criteria for formulating
568 microservices that are trustworthy, or that are reference monitors, and to understand the security
569 limitations of the microservices architectural pattern.

570 By making component dependencies and interactions more explicit, microservices appear to
571 offer a new opportunity for interposition-based security enhancements. Wrapping layers inserted
572 between microservice interactions would have the power to augment, transform, deny and
573 monitor those interactions. Those powers could be used to restrict potential damage from
574 software vulnerabilities, but interposition can also destabilize systems and impose slowdowns. A
575 possible research thrust is to investigate interposition strategies that are compatible with
576 microservice based systems.

577 **2.2.3 Maturity Level**

578 Virtualization systems date from the 1960s. The LXC container form of virtualization began in
579 2008 and has been under active development since. A number of alternate lightweight
580 virtualization systems exist, for example BSD Jails, OpenVZ and Oracle Solaris Zones.
581 Containers are substantially deployed in clouds and on servers.

582 The current microservices terminology and design goals emerged by 2014. Earlier formulations,
583 such as tasks running on microkernels, predate the CMU Mach project’s initiation in 1985. Since
584 then, microkernel technology has been a subject of ongoing research and has been integrated into
585 significant commercial products, notably Apple’s OS X.

586 **2.2.4 Basis for Confidence**

587 The base technologies are widely used, and there is a recognized need for more automation in the
588 configuration of containers. So there could be demand pull. Because containers can be very
589 quickly created, tested and deleted, there is a good case that extensive testing could be done on
590 container configurations in a semi-automated manner. With respect to microservices, a growing

591 number of microservice frameworks indicates that the technology is growing in its popularity
592 and also that there is still room for enriching new microservices frameworks and for having the
593 enrichments adopted. Also, the modular nature of microservices may offer a pathway for
594 deploying more secure versions of microservices without significantly disrupting service to
595 clients.

596 **2.2.5 Rational for Potential Impact**

597 Operating system containers and microservices are already a significant part of the national
598 information infrastructure. Given the clear manageability, cost and performance advantages of
599 using them, it is reasonable to expect their use to continue to expand. Security-enhanced versions
600 of these technologies, if adopted, can therefore have a wide-spread effect on the exploitation of
601 software vulnerabilities.

602 **2.2.6 Further Reading**

603 [Fowler14] Martin Fowler, “Microservices: a definition of this new architectural term”,
604 <http://martinfowler.com/articles/microservices.html>, March 2014

605 [What] “What’s LXC?”, <https://linuxcontainers.org/lxc/introduction/>

606 [Lemon13] Lemon, “Getting Started with LXC on an Ubuntu 13.04 VPS”,
607 [https://www.digitalocean.com/community/tutorials/getting-started-with-lxc-on-an-ubuntu-13-04-](https://www.digitalocean.com/community/tutorials/getting-started-with-lxc-on-an-ubuntu-13-04-vps)
608 [vps](https://www.digitalocean.com/community/tutorials/getting-started-with-lxc-on-an-ubuntu-13-04-vps), August 2013.

609 **2.3 Additive Software Analysis Techniques**

610 Currently there are many different tools and techniques, both as open source and in commercial
611 products, to analyze software, and they check for myriad problems. Many of them can be
612 executed through a general Integrated Development Environment (IDE), such as Eclipse. But
613 current tools face a number of impediments. IDEs sometimes do not offer an “information bus”
614 for tools to share software properties. Each tool must do its own parsing, build its own abstract
615 syntax tree (AST), list variables with their scopes and attributes and “decorate” an AST with
616 proven facts or invariants. Some tools are built on a common infrastructure, like LLVM or
617 ROSE [Rose16], so they share code, but they must still do much of the analysis over again. In
618 addition, there are few standards that allow, say, one parser to be swapped out for a new parser
619 that runs faster.

620 Additive software analysis refers to a comprehensive approach for addressing impediments to the
621 use of multiple advanced software checking tools. The goal of additive software analysis is to
622 foster a continuing accumulation of highly-usable analysis modules that add together over time
623 to continually improve the state of the art in deployed software analysis. Additive Software
624 Analysis has three parts. First, it is documentary standards to allow algorithms and tools to
625 exchange information about software. Second, it is a framework or architecture to enable
626 modular and distributed development of software assurance and assessment tools. This

627 framework has a function similar to the Knowledge Discovery Metamodel (KDM) [KDM15] or
628 what is termed a black board in Artificial Intelligence (AI). Third, it is conceptual approaches to
629 aggregate, correlate or synthesize the results and capabilities of tools and algorithms. A key
630 output of additive software analysis will be a new generation of user-facing tools to readily
631 combine the outputs from different tools and techniques into unified, more comprehensive
632 assessments of a piece of software.

633 A comprehensive additive software analysis capability must facilitate tools working together
634 (hence, it must include standards), must provide building blocks to jumpstart new tool
635 development (hence, it must include a framework) and must facilitate integration and
636 interoperability among tools (hence, it must include techniques to combine analysis results).

637 **2.3.1 Software Information Expression and Exchange Standards**

638 Software assurance tools derive and store an enormous variety of information about programs.
639 Unfortunately, there is no widely-accepted standard for exact definitions of the information or
640 how it might be stored. Because of the lack of standards, developers must perform heroic feats to
641 exchange information with fidelity between different analysis tools and algorithms.

642 Merely passing bits back and forth between tools is of little benefit unless those bits convey
643 information that is understood the same way by tools. For example, “error,” “fault,” “failure,”
644 “weakness,” “bug” and “vulnerability” are related, but different, concepts. Without a standard, if
645 one tool reports a bug, another tool may understand “bug” to indicate a higher (or lower!)
646 potential for successful attack than the first tool’s assessment.

647 For example, a variety of kinds of formally defined information may be relevant for analyzing a
648 program:

- 649 • location in code.
- 650 • the variables that are visible at a certain location, with the variable types.
- 651 • possible values of variables at a certain location. This may include relations between the
652 values of variables, such as $x < y$.
- 653 • call traces and paths, that is, all possible ways to reach this point.
- 654 • attribution to source code locations for chunks of binaries and executables.
- 655 • possible weaknesses, e.g., possible BOF [Bojanova16], or the input that will be used in
656 an SQL query not filtered and therefore tainted.
- 657 • assertions, weakest preconditions, invariants and so forth.
- 658 • function signatures, including parameter types.

659 Program analysis can be applied at various stages of software development and to
660 representations of a program at different levels of abstraction. For instance, tools may operate on
661 the static structure of a program, such as its abstract syntax tree (AST), on representations that
662 represent data or control flow and even on semantic representations that encode functional
663 behaviors, such as weakest preconditions. We look at each of these categories in turn below.

664 **Abstract Representation:** Early static checkers usually had to include their own parsers for
665 building an AST to analyze. However, compiler writers realized the importance of developing
666 common intermediate representations (IRs) that are well-documented and easily accessible. For
667 instance, in version 4.0, the development team of the GNU compiler, gcc, [GCC16] introduced
668 the intermediate language GENERIC, which is a language-independent format for representing
669 source programs in any of several languages. As another example, the Clang compiler [Clang]
670 provides a well-documented AST that may be either directly accessed by third-party plugins or
671 saved in a common format, such as JSON, to be processed by third-party analysis tools. Other
672 compilers that provide well-documented interchange formats include Frama-C [FramaC] and the
673 ROSE compiler infrastructure [Rose16].

674 **Compiler Intermediate Representation:** Tools may perform in-depth analyses on intermediate
675 representations (IRs) that are closer to the final executable code generated by compilers. For
676 instance, the GNU compiler defines the GIMPLE format in which the original source program is
677 broken down into a simple three-address language. Similarly, the Clang compiler provides the
678 LLVM bitcode representation, a kind of typed assembly language format that is not tied to a
679 specific processor.

680 **Semantic Representations:** Tools that check functional correctness properties typically need a
681 representation that is more suited to expressing logical program properties than the
682 representations discussed above. While such representations are not as mature as ASTs and
683 compiler IRs, a few have gained popularity in recent years. For instance, the intermediate
684 verification language Boogie [Barnett05], which provides features such as parametric
685 polymorphism, universal and existential quantification, nondeterministic choice and partial
686 orderings, has become a popular backend for sophisticated checkers of both low-level languages,
687 like C and C++, and higher-level object-oriented languages, like Eiffel and C#. Boogie programs
688 can be translated into the SMT-LIB format [SMTLIB15], which allows them to be checked with
689 any theorem prover that accepts the SMT-LIB format. Another example of a common language
690 for semantic representations is Datalog [Whaley05], which has been used to build a variety of
691 tools for checking array bound overflows, finding race conditions in multithreaded programs and
692 checking web application security.

693 **2.3.2 Tool Development Framework or Architecture**

694 To foster new tool development, additive software analysis requires initial building blocks. The
695 key initial building block is a framework that can tie the capabilities of tools or techniques
696 together. Just like Eclipse greatly facilitates the improvement of IDE technology for developing
697 code, a framework for additive software analysis will aim to enable synergistic development of
698 software assurance and testing tools. This “framework” may be a separate tool, or it may be a
699 plugin or update to an existing IDE.

700 Broadly speaking, there are two common methods for frameworks to transmit information
701 between program analysis tools. The first is to integrate a checker as a plugin into an existing

702 compiler toolchain. Modern compiler frameworks like gcc, Clang and Frama-C make it easy to
703 write new plugins. Furthermore, plugins are often allowed to update an AST or intermediate
704 form, thus allowing plugins to make the results of their analysis available for use by other
705 plugins. For instance, the Frama-C compiler framework provides a library of plugins that
706 includes use-def and pointer-alias analyses that are often necessary for writing semantic
707 analyzers. The second method relies on a common format that is written to disk or sent via
708 network to pass information. An example of this is the Evidential Tool Bus [Rushby05] that
709 allows multiple analysis engines produced by different vendors to exchange logical conclusions
710 in order to perform sophisticated program analyses. An additive framework would support both
711 information transmission approaches in order to reuse existing efforts as much as possible.

712 The framework capabilities referred to in this section focus on information exchange among
713 tools, rather than development capabilities of frameworks discussed in Section 2.4.

714 **2.3.3 Combining Analysis Results**

715 With standards in place and a framework, we can get increased benefit by adding together or
716 combining different software analyses. There are three general ways that results of software
717 analysis can be added together. The first case is simply more information. Suppose the
718 programmer already has a tool to check for injection class (INJ) bugs [Bojanova16]. Adding a
719 tool to check for deadlocks could give the programmer more information.

720 The second case is confirmatory. The programmer may have two different heuristics to find
721 faulty operation (FOP) bugs [Bojanova16] that have independent chances of reporting true FOP
722 bugs and false positives. The framework could be used to correlate the outputs of the two
723 heuristics to produce a single result with fewer false positives.

724 The third case of additive software analysis is synergy. A research group with expertise in formal
725 reasoning about memory use and data structures can build upon a component developed by a
726 group that specializes in “parsing” binary code, thus creating a tool that reasons about the
727 memory use of binaries. Developers can experiment with hybrid and concolic assurance tools
728 more quickly. For instance, a tool may use a static analyzer to get the code locations that may
729 have problems then, using constraint satisfiers and symbolic execution, create inputs that trigger
730 a failure at each location.

731 **2.3.4 Maturity Level**

732 Many commonly used compilers, such as gcc, Clang and Frama-C, provide built-in support for
733 adding plugins that process and update AST and IR representations. Additionally, large
734 communities have developed extensive libraries of plugins and created wiki sites with tutorials
735 and reference manuals that lower the bar for new users to become involved. In the case of
736 semantic representations, the communities are smaller and the bar to entry is higher, though
737 languages like Boogie have been successfully used as the engine by several research groups for

738 building checkers for diverse languages like C [VCC13], Eiffel [Tschannen11] and even an
739 operating system [Yang10].

740 There are many current software information exchange systems, such as LLVM, ROSE, gcc's
741 GENERIC or GIMPLE and the Knowledge Discovery Metamodel (KDM). Efforts to consolidate
742 the output of tools, such as Tool Output Integration Framework (TOIF), Software Assurance
743 Findings Expression Schema (SAFES) [Barnum12] and Code Dx [CodeDx15], already
744 implicitly indicate classes of kinds of useful knowledge about software.

745 **2.3.5 Basis for Confidence**

746 The leading static analysis tools today have low false positive rates, which has led to increasing
747 adoption throughout industry and government organizations. This in turn has motivated compiler
748 teams to add support for plugins that can operate on internal program representations. There are
749 large and active user communities that are documenting interfaces and creating libraries of
750 plugins that can be combined to build complex analyzers. Indeed, the challenge is not whether an
751 additive software analysis approach might work, but in which to invest and how to tie them
752 together.

753 **2.3.6 Rationale for Potential Impact**

754 Early static analysis tools checked mostly syntactic properties of programs, enforcing coding
755 guidelines and looking for patterns that corresponded to simple runtime errors such as
756 dereferencing a null pointer or using a variable before assignment. As analyzers became more
757 sophisticated, they increasingly relied on more complex analyses of program structure and data
758 flow. Common frameworks that allow users to build small analysis engines that can share and
759 combine results will make it possible to build sophisticated analyzers that can find subtle errors
760 that are hard to find using traditional testing and simulation techniques.

761 Such frameworks and standards should allow modular and distributed development and permit
762 existing modules to be replaced by superior ones. They should also facilitate synergy between
763 groups of researchers. They should accelerate the growth of an "ecosystem" for tools and the
764 development of next generation "hybrid" tools. A hybrid tool might use a static analyzer module
765 to find problematic code locations, then use a constraint satisfier module and a symbolic
766 execution engine to create inputs that trigger failures. A growing, shared set of problematic and
767 virtuous programming patterns and idioms may ultimately be checked by tools [Kastrinis14].

768 **2.3.7 Further Reading**

769 [Bojanova16] Irena Bojanova, Paul E. Black, Yaacov Yesha and Yan Wu, "The Bugs
770 Framework (BF): A Structured Approach to Express Bugs," 2016 IEEE Int'l Conf. on Software
771 Quality, Reliability, and Security (QRS 2016), Vienna, Austria, August 1-3, 2016. Available at
772 <https://samate.nist.gov/BF>, accessed 12 September 2016.

773 [Kastrinis14] G. Kastrinis and Y. Smaragdakis, “Hybrid Context-Sensitivity for Points-To
774 Analysis,” *Proc. Conference on Programming Language Design and Implementation (PLDI)*,
775 2014

776 [Rushby05] John Rushby, “An Evidential Tool Bus,” *Proc. International Conference on Formal*
777 *Engineering Methods*, 2005.

778 **2.4 More Mature Domain-Specific Software Development Frameworks**

779 Briefly stated, the goal of this approach is to promote the use (and reuse) of well-tested, well-
780 analyzed code, and thus to reduce the incidence of exploitable vulnerabilities.

781 The idea of reusable software components, organized into component libraries or repositories as
782 mentioned in Sect. 4.3.6, dates from at least 1968 [Mcilroy68]. To make software reusable,
783 sharable software components can be packaged in a variety of building blocks, for example:
784 standalone programs, services, micro-services, modules, plugins, libraries of functions,
785 frameworks, classes and macro definitions. A set of such (legacy) building blocks typically
786 forms the starting point for new software development efforts. Or, more colloquially expressed:
787 hardly anything is created from scratch. The vulnerability of new software systems, therefore,
788 depends crucially on the selection and application of the most appropriate existing components
789 and on the interaction of new code with legacy components.

790 Although the unit of code sharing can be small, e.g., a single function or macro, there are
791 substantial benefits to using mature, high-value, components where significant investments have
792 already been made in design cleanliness, domain knowledge and code quality.

793 A software framework contains code and, importantly, also defines a software architecture
794 (including default behavior and flow of control) for programs built using it. A domain-specific
795 framework furthermore includes domain knowledge, e.g., GUI building, parsing, Web
796 applications, multimedia, scheduling. A mature domain-specific framework, once learned by
797 software developers, can enable quick production of programs that are well tested both from a
798 software perspective and from a domain knowledge perspective. In the best case, where a mature
799 framework is wielded properly by experts, there is a substantial opportunity to avoid software
800 mistakes that can result in exploitable vulnerabilities.

801 Unfortunately, the best case is difficult to achieve. Specifically, in order to realize the benefits of
802 mature frameworks, software developers must overcome several significant challenges.

803 **Finding Suitable Frameworks.** A plethora of frameworks exist. For example, a simple search
804 of github.com in September 2016 showed over 171 000 repositories having the word
805 “framework” either in their name or in their description string. The frameworks are implemented
806 in a wide variety of programming languages (PHP, JavaScript, Java, Python, C#, C++, etc.), and
807 many frameworks use multiple languages. Additional complexity results from a diversity of
808 package management and build systems that must be learned by potential framework clients.

809 Software development teams confront a significant challenge merely to survey the possible
810 frameworks that might support a project's requirements; the challenge is acute enough that there
811 is one project [TodoMVC16] that exists solely to help developers choose among available
812 (model-view) frameworks by showing a sample application implemented in multiple
813 frameworks, for comparison purposes. Assessing suitability in surveyed frameworks is a further
814 challenge. Many frameworks include some form of testing in their build processes, often unit
815 testing [Beck94]; such existing tests need to be assessed for sufficiency relative to a project's
816 goals.

817 **Learning new Frameworks.** Brooks said [Brooks95] that software embodies both "essential"
818 and "accidental" information. The essential information is about algorithms and fundamental
819 operations that software must perform. The accidental information is about interface details,
820 programming language selection, the names given to elements in a system, etc. Each framework
821 embodies both kinds of information, which must be understood at an expert level to safely
822 employ a framework for nontrivial applications. While an expert might already know much
823 essential information for a problem domain, the accidental information cannot be anticipated.

824 A quick perusal of a common data structure, the list, illustrates the fundamental difficulty. The
825 meaning of a list is well understood by most software developers, but the information required to
826 actually create and use a list data structure is quite different between competing environments.
827 For example, the Unix queue.h macros, Java collections, JavaScript arrays, Python's built-in list
828 and the C++ Standard Template Library list template, all implement the same basic idea, but
829 using quite different details. A software developer may be an expert in the concept of a list and
830 in some list implementations, but an absolute novice in the usage of the concrete list
831 implementation in a new framework. The developer must therefore expend time for the
832 unedifying learning of (often extensive amounts of) accidental information. If developers give in
833 to schedule pressure to minimize this preparatory work, novice-level framework-based software
834 may be produced, which is more likely to contain flaws and vulnerabilities.

835 **Understanding and Controlling Dependencies.** One framework may depend on others. The
836 resulting transitive graph of dependencies can be large, and framework users may easily find the
837 vulnerabilities in their projects dependent on possibly voluminous framework code included
838 automatically and indirectly by legacy package managers and build systems. The left-pad
839 incident of 2016 illustrates the danger. The heavily-used Node Package Manager maintains
840 numerous packages that JavaScript programs can easily refer to and use. When an ownership
841 controversy erupted in 2016, an Open Source author unpublished over 250 of his modules from
842 the Node Package Manager. One was the tiny function "leftpad," which adds padding of spaces
843 or zeros to strings. Thousands of programs, some very important, relied on leftpad and suddenly
844 failed until the unpublished package was "un-unpublished." [Williams16]

845 **Resolving Framework Composition Incompatibility.** Multiple frameworks may not be usable
846 simultaneously in the same program. Or, if they are, the order of their inclusion or the version

847 may be important, resulting in brittle code. In other cases, like the lex/yacc code generation tools,
848 explicit actions are needed to avoid name space conflicts in order to allow multiple instances of a
849 framework to coexist in a program. Such conflicts may be subtle. As Lampson points out
850 [Lampson04], each component may have a distinct “world view” and the composition of n
851 components can result in n^2 interactions.

852 These are long-standing challenges. Moreover, due to the large and growing number of
853 frameworks (of varying provenance and quality) currently available in Open Source via public
854 repositories hosted by repository-management entities such as GitHub, JIRA, Bitbucket,
855 CollabNet, etc., the difficulty of choosing a suitable framework may be more acute. This scale,
856 however, also represents an important opportunity: if even small improvements can be achieved
857 to how frameworks are found, learned, dependency-managed and composed, many software
858 vulnerabilities may be avoided.

859 A second significant development is the mainstreaming of software development (including
860 framework use) through copy/paste operations using software question/answer sites such as
861 stackoverflow or stackexchange. Although question/answer-based code reuse can be fast, it also
862 can result in poorly-understood and poorly-integrated solutions. The ability to get answers and
863 sample code for questions posed clearly can benefit developer comprehension, however
864 techniques are needed to avoid generating vulnerabilities when adapting others’ solutions.

865 Although these are significant challenges, the current state of the art provides opportunities to
866 leverage existing code and skills resources while augmenting them with new techniques and
867 tools.

868 **2.4.1 Rapid Framework Adoption**

869 Framework adoption is clearly impeded by the need to learn great quantities of accidental
870 information. Gabriel defines “habitability” as “the characteristic of source code that enables
871 programmers, coders, bug-fixers, and people coming to the code later in its life to understand its
872 construction and intentions and to change it comfortably and confidently.” [Gabriel96]
873 Recognizing the challenge of achieving habitability, Gabriel suggests the use of software
874 patterns to help developers quickly understand existing code, as well as to flag the use of
875 negative practices. Although not a panacea, patterns (e.g. [Gamma95]) can help bridge the
876 conceptual gap between framework providers and framework consumers. One approach to
877 facilitating this is to develop a set of patterns that encompass popular domains. An informal
878 survey in September 2016 of the top 10 most popular (“star’d”) and most “forked” repositories
879 on GitHub shows significant framework activity around Web application development, Front-
880 end Web development, operating system kernels, cross platform application frameworks, virtual
881 machine management, programming languages and asynchronous http servers. One approach to
882 speeding adoption is to formulate software patterns for some of these domains, with a focus on
883 harmonizing the accidental information between frameworks (so it need not be learned multiple
884 times) and to produce documentation for common use cases. Experiments can then measure the

885 effectiveness by comparing framework uptake both with and without the new pattern
886 information.

887 **2.4.2 Compositional Testing**

888 Advanced testing approaches hold promise to substantially increase framework robustness, and
889 furthermore, to build assurance for compositions of frameworks under various assumptions
890 regarding dependencies. Many frameworks currently employ only ad hoc testing. Others employ
891 standard unit testing [Beck94], practiced at varying levels of completeness. Recent advances in
892 the measurement of traditional test suite coverage provide an opportunity to compare
893 frameworks. Combinatorial testing [Kuhn10] has been used to improve on black box fuzz testing
894 as well as to test alternate software configurations. The many ways in which frameworks may be
895 customized or configured suggest a possible approach for gaining new confidence in the use of
896 software frameworks. By demonstrating high quality compositions, such testing also has
897 potential to highlight framework similarities, reduce learning curves and enable broader adoption
898 of well-tested, well-analyzed code.

899 **2.4.3 Conflict Resolution in Multi-Framework Composition**

900 In some cases, multiple frameworks can be used together concurrently without conflict. In
901 others, the composition details that allow concurrent use may be fragile. Dominant framework
902 patterns such as inversion of control (IoC) [Busoli07], also known as the Hollywood principle:
903 “don’t call us; we’ll call you,” may exacerbate this because each framework may assume that it
904 is defining the flow of control in an entire application. One approach for mitigating this is to
905 virtualize framework operations using, for example, lightweight operating system containers
906 [LXC] and then establish communication links between concurrently executing frameworks.
907 Another approach to conflict resolution is to employ software translation to rewrite frameworks
908 so that their overlapping elements become distinct. Pilot efforts can demonstrate the feasibility of
909 these and other deconfliction strategies and compare their costs and effects on application
910 vulnerability.

911 **2.4.4 Maturity Level**

912 The literature of software patterns is quite extensive and software testing is a relatively mature
913 subfield of computer science, practiced now for over 40 years. Frameworks themselves are now
914 a dominant unit of software sharing. The three supporting techniques listed in this section are
915 under continuous use and refinement.

916 **2.4.5 Basis for Confidence**

917 There is little doubt that patterns can be documented for several significant frameworks; rapid
918 uptake may be a more incremental than revolutionary improvement, but incremental
919 improvements should flow from investments in pattern documentation. The advanced testing
920 techniques that would be brought to bear on framework compositions, are relatively mature,
921 increasing confidence that framework integrations can be effectively tested.

922 **2.4.6 Rational for Potential Impact**

923 Code reuse is pervasive and seemingly accelerating; by investing in very popular frameworks,
924 any improvements will be widely relevant.

925 **2.4.7 Further Reading**

926 [Software16] “Software framework”, https://en.wikipedia.org/wiki/Software_framework

927 [TodoMVC16] “TodoMVC: Helping you select an MV* framework”, <http://todomvc.com/>

928 [Wayner15] Peter Wayner, “7 reasons why frameworks are the new programming languages”,

929 <http://www.infoworld.com/article/2902242/application-development/7-reasons-why->

930 [frameworks-are-the-new-programming-languages.html](http://www.infoworld.com/article/2902242/application-development/7-reasons-why-frameworks-are-the-new-programming-languages.html), March 2015.

931 **2.5 Moving Target Defenses (MTD) and Artificial Diversity**

932 This approach is a collection of techniques to vary software’s detailed structures and properties
933 such that an attacker has much greater difficulty exploiting any vulnerability. To illustrate,
934 consider one early, widely-used technique in this family: Address Space Layout Randomization
935 (ASLR), invented in 2001 by the PaX Team [PaX01]. When a program requests a buffer, the
936 easiest thing is to return the next available chunk of memory. This puts buffers in the same
937 relative location. Knowing this, an attacker can exploit a buffer overflow weakness (BOF)
938 [Bojanova16] in one buffer to, say, read the password that is in another buffer that is always 384
939 bytes beyond it. ASLR puts buffers in different (unpredictable) relative locations, so that the
940 above exploit is much harder.

941 The goal of artificial diversity and moving target defense (MTD) is to reduce an attacker's ability
942 to exploit vulnerabilities in software, not to reduce the number of weaknesses in software.

943 Diversification must, of course, be safe. That is, changes have no effect on normal behavior,
944 other than perhaps higher use of resources. Even with this constraint we can trade compute
945 power for increased granularity or thoroughness of diversification. The increased granularity is
946 presumed to offer better protection against exploitation of unknown vulnerabilities because of
947 the higher probability of affecting the location or value of some piece of information essential to
948 an attack. This tradeoff is similar to that for static analysis, referred to in Sect. 2.1.1 and 2.1.2:
949 the more resource invested, the higher the amount of assurance. The difference is that static
950 analysis provides assurance that the software does not contain vulnerabilities of specific types,
951 while MTD provides assurance that weaknesses of any type are expensive to exploit.

952 **2.5.1 Compile-Time Techniques**

953 Compile-time techniques are those applied automatically by a compiler. They may result in the
954 same executable for each compilation, such that the executable then chooses random behaviors
955 or memory layouts at run time, or they may result in a different executable at each compilation.

956 Some specific techniques are data structure layout randomization, different orders of parameters
957 in function calls, ASLR, instruction set randomization, data value randomization, application
958 keyword tagging and varied instruction ordering with operation obfuscation and refactoring.

959 The program information that is useful for proving that these diversifications are safe is also
960 useful for program analysis to find or remove vulnerabilities. The additive software analysis
961 approach, detailed in Sect. 2.3, is to use the same compute power to simultaneously detect or
962 remove weaknesses and to also randomize remaining weaknesses. These diversification
963 techniques could be tied into a static analysis tool through the additive analysis framework,
964 potentially with very modest resource expenditures.

965 Unfortunately, no tools do this today. Analysis software is usually run by the programmer, at
966 development time. Diversification typically only displays its benefit in the system test phase or
967 in the operation phase when it demonstrates resilience. At worst, diversification adds ambiguity
968 to test results and makes it more difficult to track down root causes of failures. To counteract this
969 disconnect between effort and benefit, programs that use diversification should be specifically
970 acknowledged, so customers know that they employ an extra layer of resilience.

971 **2.5.2 System or Network Techniques**

972 Some techniques at the system or network level are network address space randomization and
973 protocol diversity. These are likely to be dynamic in that they change on a regular basis. In many
974 cases, these are built on the assumption of a shared secret map from services to address or a
975 shared secret key, so an application can authenticate and get current information.

976 **2.5.3 Operating System Techniques**

977 An operating system (OS) may present different interfaces to different processes. These could be
978 dynamic, such as a random interrupt number assigned for each system service, or static, in which
979 the OS has several choices for each set of services. In the dynamic case, the linker/loader can
980 adjust each new executable to the assignments made for the process. As an example of the static
981 case, an OS presents a new process with a set C of memory management APIs, a set B of process
982 services, a set D of networking functions and a set A of I/O calls. Invasive code trying to execute
983 through that process would have to deal with $j \times k \times m \times n$ different OS interfaces in order to
984 succeed.

985 **2.5.4 Maturity Level**

986 Some moving target defenses are the default in many operating systems and compilers today.
987 There is intense research and entire conferences to understand limitations, costs and benefits of
988 current techniques and develop new and better techniques.

989 **2.5.5 Basis for Confidence**

990 The benefit in terms of number of attacks foiled, attackers discouraged or additional attacker
991 resources required is not known. However, many MTD techniques can be applied automatically,
992 e.g. by the compiler, at little cost of resources or run time.

993 **2.5.6 Rationale for Potential Impact**

994 MTD techniques can be applied to most programs and systems today, even static embedded
995 systems. Thus the scope of benefits is extremely large. The impact is not clear since most
996 techniques increase attacker's costs, not strictly eliminate vulnerabilities.

997 **2.5.7 Further Reading**

998 [Okhravi13] H. Okhravi, M.A. Rabe, T.J. Mayberry, W.G. Leonard, T.R. Hobson, D. Bigelow
999 and W.W. Streilein, "Survey of Cyber Moving Targets", Massachusetts Institute of Technology
1000 Lincoln Laboratory, Technical Report 1166, September 2013. Available at
1001 [https://www.ll.mit.edu/mission/cybersec/publications/publication-](https://www.ll.mit.edu/mission/cybersec/publications/publication-files/full_papers/2013_09_23_OkhraviH_TR_FP.pdf)
1002 [files/full_papers/2013_09_23_OkhraviH_TR_FP.pdf](https://www.ll.mit.edu/mission/cybersec/publications/publication-files/full_papers/2013_09_23_OkhraviH_TR_FP.pdf) Accessed 15 September 2015.

1003

1004 **3 Measures and Metrics**

1005 This section deals with metrics, measures, assessments, appraisals, judgements, evaluations, etc.
1006 in the broadest sense. Hence, code reviews and software testing have a place in this section. We
1007 have three areas of concern. First, encouraging the use of metrics. All the extraordinary metrics
1008 in the world do not help if nobody uses them. Also, nobody *can* act on metrics if the metrics are
1009 not produced and available. The Federal Government might motivate and encourage the use of
1010 software product metrics. Vehicles include procurement, contracting, liability, insurance and also
1011 standards as explained in Sect. 4.3. Software can also benefit from the programs and criteria of
1012 third-party, non-governmental organizations. Some possibilities are Underwriter's Laboratory
1013 Cybersecurity Assurance Program (CAP), Consortium for IT Software Quality (CISQ) Code
1014 Quality Standards, Coverity Scan, Core Infrastructure Initiative (CII) Best Practices badge and
1015 the Building Security In Maturity Model (BSIMM). Many of these include process metrics,
1016 which is the second area.

1017 The second area, process metrics includes hours of effort, number of changes with no acceptance
1018 test defects or acceptance test defect density in delivered code [Perini16]. These do not have a
1019 direct effect on the number of vulnerabilities, but the indirect effects are significant. For
1020 example, if developers are forced to frequently work overtime to meet a deadline or the schedule
1021 doesn't allow for training, the number of vulnerabilities is likely to be much higher. Other
1022 examples are software measures that indicate how much a new process step helps compared to
1023 the former practice or metrics that indicate parts of the process that are allowing vulnerabilities
1024 to escape. This approach of continuously improving the process is found in the highest levels of
1025 maturity models. It also allows groups to adopt or adapt methods and metrics that are most
1026 applicable to their circumstance. We do not discuss process metrics further.

1027 The final area of concern is metrics of software as a product, for instance, proof of absence of
1028 buffer overflows, number of defects per thousand lines of code, assurance that specifications are
1029 met or path coverage achieved by a test suite. The Software Quality Group at the U.S. National
1030 Institute of Standards and Technology (NIST) organized a workshop on Software Measures and
1031 Metrics to Reduce Security Vulnerabilities (SwMM-RSV) to gather ideas on how the Federal
1032 Government can best identify, improve, package, deliver or boost the use of software measures
1033 and metrics to significantly reduce vulnerabilities. The web site is
1034 <https://samate.nist.gov/SwMM-RSV2016.html>. They called for short position statements, then
1035 invited workshop presentations based on 10 of the 20 statements submitted. The workshop was
1036 held on 12 July 2016. The full workshop report is available as NIST SP-XXXX. Much of this
1037 section is informed by the results of the workshop. Ideas were often brought up by one person,
1038 discussed and elaborated by others, then written or reported by yet others. Hence it is difficult to
1039 attribute ideas to particular people in most cases. We thank all those who participated in the
1040 workshop and made contributions, large and small, to the ideas noted in the report.

1041 We distinguish between metrics and measures. A metric is a simple, basic assessment or count
1042 with a clear value. A measure, on the other hand, is derived from other metrics and measures.
1043 Measures are often surrogates for properties that we would like to be able to determine. For
1044 instance, number of buffer overflow weaknesses is a metric with a reasonably clear definition. In
1045 contrast, code security is a measure that is only loosely related to the number of buffer
1046 overflows. The absence of flaws does not indicate the presence of excellence.

1047 **3.1 A Taxonomy of Software Metrics**

1048 Software metrics may be classified along four dimensions. The first dimension is how “high-
1049 level” the metric. Low-level metrics are below semantics, such size of a program, number of
1050 paths, and function fan in/fan out. High-level metrics deal more with what the program is meant
1051 to accomplish. The second dimension is static or dynamic. Static metrics are those apply to the
1052 source code or “binary” itself. Dynamic metrics apply to the execution of the program. The third
1053 dimension is the point of view. It may be either an external view, sometimes called black box or
1054 functional, or an internal, “transparent” view, referred to as white box or structural. The fourth
1055 dimension is the object of the metric: bugs, code quality, and conformance.

1056 Software metrics may be divided into two broad categories as to whether they are low-level or
1057 high-level. Low-level metrics are generally widely applicable. High-level metrics, in contrast,
1058 deal with the relation between the program, as an object, and the developer or user, as a sentient
1059 subject. It is in this interaction between object and subject that quality arises, as Pirsig said.
1060 [Pirsig74] Analogously to low- and high-level metrics, there are low-level vulnerabilities and
1061 there are high-level vulnerabilities. Some low-level vulnerabilities are buffer overflow, integer
1062 overflow and failure to supply default switch cases. These low-level vulnerabilities can be
1063 discerned directly from the code. That is, one can inspect the code or have a program inspect the
1064 code and decide whether there’s a possibility of a buffer overflow (BOF) [Bojanova16] given
1065 particular inputs. There is no need to refer to a specification, requirement or security policy to
1066 determine whether a buffer overflow is possible.

1067 On the other hand, high-level vulnerabilities cannot be discerned solely by reference to the code.
1068 A human reviewer or a static analyzer must refer to requirements, specifications or a policy to
1069 determine high-level problems. For instance, failure to encrypt sensitive information generally
1070 cannot be discerned solely by code inspection. Of course, heuristics are possible. For example, if
1071 there is a variable named “password,” it is reasonable for a static analyzer to guess that variable
1072 is a password and should not be transmitted without protection or be available to unauthorized
1073 users. But neither tool nor human can determine whether or not the information in a variable
1074 named “ID” should be encrypted or not without examining an external definition.

1075 Having access to a requirements document for a security policy does not allow the quality of
1076 software to be assessed in all cases. Requirements documents typically deal with the behavior of
1077 the program and what the program uniquely needs to do. It is difficult, and perhaps impossible,
1078 to specify formally that code should be high quality. Software architecture is an attempt to define

1079 the structural components that distinguish good and useful software from software that is error-
1080 prone, difficult to debug, brittle or inflexible.

1081 The second dimension of classifying metrics is most apparent in testing. Test metrics
1082 conceptually have two parts: test generation or selection and test result evaluation. Test metrics
1083 generally answer the question, how much of the program (internal) or the input space (external)
1084 has been exercised? Test case generation is necessarily static, while evaluation is usually
1085 Θ dynamic, that is, based on the result of executions. In many test metrics, the two parts are tied
1086 to each other. They include a step like, choose additional test cases to increase the coverage, thus
1087 the dynamic part influences the static part. Testing is usually referred to as a dynamic technology
1088 since program execution is an essential part of testing. That is, if one comes up with test cases
1089 *but never runs them*, then no assurance is gained, strictly speaking. Of course, in most cases the
1090 thought and scrutiny that goes into selecting test cases is a static analysis that yields some
1091 assurance about the program.

1092 The third dimension is the point of view, either external or internal. External metrics are
1093 typically behavioral conformance to specifications, requirements or constraints. They are often
1094 referred to as “black box” or behavioral. These metrics are particularly useful for acceptance
1095 testing and estimating user or mission satisfaction. It matters little how well the program
1096 functions or is structured internally if it does not fulfil its purpose. In contrast, internal or
1097 structural metrics primarily deal with, or are informed by, the code’s architecture,
1098 implementation and fine-grained operation. Metrics in this class are related to qualities such as
1099 maintainability, portability, elegance and potential. For instance, external timing tests may be
1100 insufficient to determine the order of complexity of an algorithm whereas code examination may
1101 clearly show that the algorithm is order $\Theta(n^2)$ and will have performance issues for large inputs.

1102 Determining how much testing is enough also shows the difference between internal and external
1103 metrics. External metrics, such as boundary value analysis [Beizer90] and combinatorial testing
1104 [Kuhn10], consider the behavioral or specification in computing how much has been tested or
1105 what has not been tested. On the other hand, internal metrics include counts of the number of
1106 blocks, mutation adequacy [Okun04], and path coverage metrics [Zhu97]. The two approaches
1107 are complementary. External testing can find missing features. Internal testing can bring up cases
1108 that are not evident from the requirements, for example, switching from an insertion sort to a
1109 quick sort when there are many items.

1110 The fourth dimension to classify metrics conceptually divides them into three types. The first is
1111 presence (or absence) of particular weaknesses such as buffer overflow (BOF) or injection (INJ)
1112 [Bojanova16]. Note that the absence of flaws does not indicate, say, resilient architecture. The
1113 second type is quality metrics that directly measure that code, or parts of it, is excellent.
1114 However, we only have proxies for “quality,” like maintainability, portability or the presence of
1115 assertions. The third type is conformance to specification or correctness. This third type of metric
1116 must be specific to each task. General requirement languages and checking approaches are

1117 available. Because of the profound differences between these three types, there is no one security
1118 or vulnerability metric or measure that guarantees excellent code.

1119 **3.2 Software Assurance: The Object of Software Metrics**

1120 Software assurance, that is, our assurance that software will behave as it should, comes from
1121 three broad sources. The first is the development process. If software is developed by a team
1122 with clear requirements, are well trained and who have demonstrated an ability to build good
1123 software with low vulnerability rates, then we have confidence or assurance that software that
1124 they produce is likely to be have few vulnerabilities. The second source of assurance is our
1125 analysis of the software. For instance, code reviews, acceptance tests and static analysis can
1126 assure us that vulnerabilities are likely to be rare in the software. We can trade off these two
1127 sources of assurance. If we have little information about the development process or the
1128 development process has not yielded good software in the past, we must do much more analysis
1129 and testing to achieve confidence in the quality of the software. In contrast, if we have
1130 confidence in the development team and the development process, we only need to do minimal
1131 analysis in order to be sure that the software follows past experience.

1132 The third source of software assurance is a resilient execution environment. If we do not have
1133 confidence in the quality of the software, then we can run it in a container, give it few system
1134 privileges, then have other programs monitor the execution. Then if any vulnerabilities are
1135 triggered, the damage to the system is controlled.

1136 With research we may be able to give detail to the mathematical formula that expresses our
1137 assurance: $A = f(p, s, e)$ where A is the amount of assurance we have, p is the assurance that
1138 comes from our knowledge of the process, s is assurance from static and dynamic analysis and e
1139 is the assurance that we gain from strict execution environments.

1140 **3.3 Software Metrology**

1141 To have a coherent, broadly useful system of metrics, one must have a solid theoretical
1142 foundation. That is, a philosophy of software measurement. This section addresses questions
1143 such as, what is software metrology? What is its purpose? What are the challenges unique to
1144 measuring software, in contrast to physical measurement? What are possible solutions or
1145 potential approaches?

1146 Software metrics have well known theoretical limitations, too. Analogous to Heisenberg's
1147 Uncertainty Principle in Physics, Computer Science has the Halting Problem, Rice's Theorem
1148 and related results that show that it is impossible to correctly determine interesting metrics for *all*
1149 possible programs. Although this is a caution, it does not mean that all useful, precise, accurate
1150 measurement is impossible. There are several ways to avoid these theoretical road blocks. First,
1151 we may be satisfied with relative properties. It may be satisfactory to be able to determine that
1152 the new version of a program is more secure (or less!) than the previous version. We need not
1153 have an absolute measure of the security of a program. Second, a metric might apply only to

1154 program that do not have perverse structures. A metric may still be useful even if it doesn't apply
1155 to programs consisting solely of millions of conditional go-to statements with seemingly
1156 arbitrary computations interspersed. Nobody (should) write programs like that. Finally, society
1157 may decide that for certain applications, we will only build measurable software. Architects are
1158 not allowed to design building with arbitrary structures. They must run analyses showing that the
1159 design withstands expected loads and forces. Instead of writing some software and trying to
1160 show that it works, the expectation might change to only writing software that definitely satisfies
1161 its constraints and requirements.

1162 Computer programmers use the phrase “it’s not a bug: it’s a feature” half-seriously. Its sue
1163 highlights that bugs and features are entities that are related somehow. Let us assume that a
1164 program can be characterized as a set of features. (The notion that a program is a set of features
1165 is the basis of some size metrics. For example, Function Points attempts to capture the notion of
1166 a basic operation or function.) Saying that a program “has a bug” means it is a buggy version of a
1167 “good” program. Both the good program and the buggy version are programs. According to the
1168 assumption, both programs are a set of features. Therefore, the difference between the good
1169 program and the buggy program is some set of features—features added, removed, or changed.
1170 Hence, a precise definition is that a bug is the difference between the features you want and the
1171 features you have. In many cases, a bug may merely be an additional feature or one feature
1172 replacing another.

1173 We might contrast software metrology with physical metrology. In physical metrology the
1174 challenge is to precisely and reproducibly determine the properties of physical objects, events or
1175 systems. For software, on the other hand, most of the so-called measurement is merely counting.
1176 A case in point is that ASCMM-MNT-7: Inter-Module Dependency Cycles has a precise
1177 definition. [ASCMM16] It is not terribly difficult to write a program that precisely measures the
1178 number of instances where a module has references that cycle back a piece of software. The
1179 difference then is that physical metrology has clearly identified the properties that they want to
1180 determine, for instance, mass, length, duration and temperature. On the other hand, software
1181 metrology has a distinct gap. We want to determine measure high-level properties such as
1182 quality, maintainability and security, but we do not have precise definitions of those, and
1183 therefore cannot measure those directly. We can, however, measure many properties which are
1184 correlated with those high-level properties.

1185 Currently metrology relegates counting the number of entities to a second-class method of
1186 determining properties. Such counted *quantities* are all considered to be the same dimension one,
1187 sometimes called dimensionless quantities, although they may be different *kinds*.

1188 **3.4 Product Metrics**

1189 As much as good process is essential to the production of code with few vulnerabilities, the
1190 ultimate is to measure the code itself. As pointed out in the introduction to this section, measures
1191 of the software itself inform process improvement.

1192 Security or vulnerability measurement, in the broadest sense, which includes testing and
1193 checking, must be include in *all* phases of software development. Except for ambitious
1194 approaches like Clean Room, this kind of measurement cannot be left as a gate near the end of
1195 the production cycle.

1196 It is possible that software quality and security metrics may be the wrong emphasis to reduce
1197 software vulnerabilities. Such metrics may fade in emphasis as other software metrics have, for
1198 example cohesion and McCabe Cyclomatic Complexity. Perhaps the best approach is a “Clean
1199 Room” approach, in which metrics inform a decision to accept or reject and do not purport to
1200 establish an absolute certification of freedom from errors.

1201 **3.4.1 Existing Metrics**

1202 There are hundreds of proposed software metrics and measures, such as, lines of code, class
1203 coupling, number of closed classes, function points, change density and cohesion. Most of these
1204 are not precisely defined and are not rigorously validated. Worse yet, most of these only have
1205 moderate correlation with the high-level properties that we wish to determine in software. For
1206 instance, lines of code (LoC) capture only some of the variance in program capability. LoC for
1207 the same specification in the same language varies by as much as a factor of four, even when all
1208 programmers have similar expertise. On the other hand, LoC has a remarkably robust correlation
1209 with the number of bugs in a program. (This suggests that higher level languages, which allow a
1210 programmer to express functionality more succinctly, will lead to fewer bugs in general.)

1211 Even something as seemingly simple as counting the number of bugs in a program is surprisingly
1212 complicated [Black11b]. It is difficult to even subjectively define what is a bug. For example,
1213 one can write a binary search that is never subject to integer overflow, but the code is hard to
1214 understand. Dividing by zero may have a well-defined behavior, resulting in the special value
1215 “NaN”, but that is generally not a useful result. Bugs are often a cascade of several difficulties.
1216 Suppose (1) an unchecked user input leads to (2) an integer overflow that leads to (3) a buffer
1217 being allocated that is too small that causes (4) a buffer overflow that finally leads to (5)
1218 information exposure. Do we count this as one bug or five? If a programmer makes a systematic
1219 mistake in several places, say not releasing a resource after use, is that one problem or several?
1220 Rather than being the exception, these kinds of complication are the rule in software [Okun08].

1221 For any realistic program, it is infeasible to try every single possible input. Instead, one must
1222 choose a metric that spans the entire space. Some of these metrics are combinatorial input
1223 metrics [Kuhn10], mutation adequacy [Okun04], path coverage metrics [Zhu97] and boundary
1224 value analysis [Beizer90].

1225 There are far too many proposed measures to evaluate or even list here. We can state that, as
1226 alluded to above, metrics and measures should be firmly based on well-established science and
1227 have a rational foundation in metrology to have the greatest utility. [Flater16]

1228 3.4.2 Better Code

1229 Two workshop presentations, Andrew Walenstein’s “Measuring Software Analyzability” and
1230 James Kupsch’s “Dealing with Code that is Opaque to Static Analysis,” point the direction to
1231 new software measures. Both stressed that code should be amenable to automatic analysis. Both
1232 presented approaches to define what it means that code is readily analyzed, why analyzability
1233 contributes to reduced vulnerabilities and how analyzability could be measured and increased.

1234 There are subsets of programming languages that are designed to be analyzable, such as SPARK,
1235 or to be less error-prone, like Less Hatton’s SaferC. Participants generally favored using better
1236 languages, for example, functional languages such as F# or ML. However, there was no
1237 particular suggestion of *the* language, or languages, of the future.

1238 While code-based metrics are important, we can expect complementary results from metrics for
1239 other aspects of software. Some aspects are the software architecture and design erosion metrics,
1240 linguistic aspects of the code, developers’ backgrounds and metrics related to the software
1241 requirements.

1242 3.4.3 Metrics and Measures of Binaries and Executables

1243 Some workshop participants were of the opinion that there is a significant need for metrics and
1244 measures of binaries or executables. With today’s optimizing compilers and with the dependence
1245 on many libraries delivered in binary, solely examining source code leaves many avenues for
1246 appearance of subtle vulnerabilities.

1247 3.4.4 More Useful Tool Outputs

1248 There are many powerful and useful software assurance tools available today. No single tool
1249 meets all needs. Accordingly, users should use several tools. This is difficult because tools have
1250 different output formats and use different terms and classes. Tool outputs should be standardized.
1251 That is, the more there is common nomenclature, presentation and detail, the more feasible it is
1252 for users to combine tool results with other software assurance information and to choose a
1253 combination of tools that is most beneficial for them.

1254 Participants felt the need for scientifically valid research about tool strengths and limitations,
1255 mechanisms to allow publication of third party evaluation of tools, a common forum to share
1256 insights about tools and perhaps even a list of verified or certified tools.

1257 3.5 Further Reading

1258 [Barritt16] Keith Barritt, “3 Lessons: FDA/FTC Enforcement Against Mobile Medical Apps,”
1259 January 2016. Available at <http://www.meddeviceonline.com/doc/lessons-fda-ftc-enforcement-against-mobile-medical-apps-0001>

1261 [FTC16] “Mobile Health App Developers: FTC Best Practices,” April 2016. Available at
1262 <http://www.ftc.gov/tips-advice/business-center/guidance/mobile-health-app-developers-ftc-best-practices>
1263

1264 [Perini16] Barti Perini, Stephen Shook and Girish Seshagiri, “Reducing Software Vulnerabilities
1265 – The Number One Goal for Every Software Development Organization, Team, and Individual,”
1266 ISHPI Information Technologies Technical Report, 22 July 2016.

1267

1268 **4 Summary and Community Engagement**

1269

1270 In response to the February 2016 Federal Cybersecurity Research and Development Strategic
1271 Plan, NIST was asked to identify ways to dramatically reduce software vulnerabilities. NIST
1272 worked with the software assurance community to identify five promising approaches. This
1273 report presents some background for each of the approaches along a summary statement of the
1274 maturity of the approach and the rationale for why it might make a dramatic difference. Further
1275 reading was provided for each approach. Hopefully other approaches will be identified in the
1276 future.

1277

1278 These approaches are focused on technical activities with a three to seven-year horizon. Many
1279 critical aspects of improving software, such as creating better specifications, using the testing
1280 tools available today, understanding and controlling dependencies and creating and following
1281 project guidelines, were not addressed. While these areas fall outside the scope of the report, they
1282 are critical both now and in the future. Similarly, the report does not address research and
1283 development that is needed as part of a broader understanding of software and vulnerabilities.
1284 Topics such as identifying sources of vulnerabilities, how vulnerabilities manifest as bugs,
1285 improved scanning during development and use are also critical, but, again, outside the scope of
1286 this report.

1287 This section of the report outlines some of the needed steps for moving forward by engaging the
1288 broader community, including researchers, funders, developers, managers and customers/users.
1289 The section addresses: 1) engaging and supporting the research community, 2) education and
1290 training and 3) empowering customers and users of software to meaningfully participate by not
1291 only asking for quality, but pushing it.

1292

1293 **4.1 Engaging the Research Community**

1294 There are many approaches to engaging the research community beyond simply funding secure
1295 software research.

1296 **4.1.1 Grand Challenges, Prizes and Awards**

1297 Many organizations have announced grand challenges, some of which are general research goals
1298 and some are competitions. More secure software can be the focus of challenges or a side
1299 benefit, that is, the competition could be focused on a non-security goal, but require the winner
1300 to produce secure software. Many organizations use bug bounty programs to incentive the
1301 research community to find and notify organizations about bugs.

1302 **4.1.2 Research Infrastructure**

1303 There is a need for repositories of data related to secure software. Several very successful
1304 repositories exist, such as the National Vulnerability Database. However, many more are needed.

1305 There could be repositories to share related research as well as open repositories of source code,
1306 as mentioned in Sect. 4.3.6. There is also a need for a better understanding of weaknesses and
1307 bugs. For example, what proportion of vulnerabilities result from implementation errors and
1308 what proportion from design errors? Researchers need to be able to replicate results and test
1309 across different types of code. All of these activities require a large and public research
1310 infrastructure.

1311

1312 **4.2 Education and Training**

1313 The role of education and training cannot be overstated. This is the primary mechanism how new
1314 approaches are transitioned from the research community to both the development community
1315 and to the user/customer community.

1316 Education and training for the developer community needs to address both up and coming
1317 developers currently in the educational system as well as current developers who need to update
1318 their skills.

1319 Over the past couple of years, there has been a shift in focus in higher education to include a
1320 greater emphasis on designing software with security built in from the beginning rather than
1321 added afterwards. K-12 education has also seen growth in cybersecurity efforts – both from the
1322 user and producer perspectives. It is clear that computer science and cybersecurity come together
1323 in the issue of secure programming. Understanding the principles of cybersecurity are essential
1324 to making sure that software is secure, more and more academic programs are educating their
1325 students to program with security in mind.

1326 Current developers need to be exposed to new approaches and techniques. In order for
1327 developers to make changes, they need to see evidence that the new approaches and techniques
1328 will be effective, as well as training material. To complement the training of front-line software
1329 developers, managers and executives must also be educated in the risk management implications
1330 of software vulnerabilities and the importance of investing in cybersecurity and low vulnerability
1331 software. In order for this training to be successful, it, too, will require evidence that investment
1332 in secure software will be cost effective.

1333 It is currently unknown which pedagogical techniques are most effective. Early research has
1334 shown that providing developers with a better understanding of weaknesses creates better
1335 programs. [Wu11] Additional research, as well as training material ranging from use cases to
1336 how to guides will be needed for successful transition. The Federal government can lead by
1337 example by training its developer community.

1338

1339 **4.3 Consumer-Enabling Technology Transfer**

1340 One of the drivers for better software is if users, consumers and purchasers of software demand
1341 it. While the user community clearly wants higher quality software, it is difficult for them to
1342 meaningfully ask for it and know if it has been received. Improved metrics that are customer-
1343 focused are needed as are other policy and economic approaches. Policy and economic
1344 approaches are outside the scope of this report, but are critical to successful technology transfer
1345 for improved software. This section outlines some of these approaches that were discussed
1346 during the various workshops.

1347 **4.3.1 Government Contracting and Procurement**

1348 The Federal Government could lead a significant improvement in software quality by requiring
1349 software quality during contracting and procurement and by changing general expectations.
1350 Model contract language can include incentives for software to adhere to higher coding and
1351 assurance standards or punitive measures for egregious violations of those standards. Sample
1352 procurement language for cybersecurity and secure software has been published by the defense
1353 community [Marien16], the financial sector, the automotive sector and the medical sector. The
1354 focus on low bidder must include provisions for “fitness for purpose” that factor in
1355 considerations for secure software.

1356 **4.3.2 Liability**

1357 There is much discussion in the software community about liability including during the
1358 Software Measures and Metrics to Reduce Security Vulnerabilities (SwMM-RSV) workshop.
1359 Many felt that companies developing software should be contractually liable for vulnerabilities
1360 discovered after delivery. Many participants did not believe that there should be legal liability at
1361 this time. On the other hand, the language of such liability clauses needs to be strict enough to, as
1362 one participant wrote, “hold companies accountable for sloppy and easily-avoidable errors, flaws
1363 and mistakes.”

1364 Defining “sloppy and easily avoidable” is not a trivial matter. An additional complicating factor
1365 is that liability includes a concept of who is responsible. Responsibility may be hard to determine
1366 in the case of “open source” or freely available software.

1367 **4.3.3 Insurance**

1368 Cyber insurance is a growing area as cyber continues to grow in importance. The Financial
1369 Services Sector Coordinating Council (FSSCC) for Critical Infrastructure Protection and
1370 Homeland Security produced a 26-page document entitled Purchasers’ Guide to Cyber Insurance
1371 Products defining what this kind of insurance is, explaining why organizations need it,
1372 describing how it can be procured and giving other helpful information.

1373 **4.3.4 Vendor-Customer Relations**

1374 It would help end users if software has a “bill of materials” such that those using it could respond
1375 to a new threat in which some part of the software became a vector of attack. Users are

1376 sometimes prohibited by software licenses from publishing evaluations or comparisons with
1377 other tools. Georgetown University recently published a study of this issue. [Klass16] The study
1378 was sponsored by the Department of Homeland Security (DHS) Science & Technology
1379 Directorate (S&T), Cyber Security Division through the Security and Software Engineering
1380 Research Center (S²ERC).

1381 **4.3.5 Standards**

1382 The development and adoption of standards and guidelines, as well as conformity assessment
1383 programs, are used across multiple industries to address quality. The US system of voluntary
1384 industry consensus standards allows for great flexibility to address needs. In some cases, the
1385 Government (federal or state/local) set regulatory standards and communities often self-regulate.

1386 **4.3.6 Code Repositories**

1387 We explained the need for additional repositories of well-tested code in both Sections 2.1 and
1388 2.4. Code repositories promote code re-use and encourage organizations to test code by
1389 providing a location where the results can be published. Repositories can also contain examples
1390 of low bug densities projects such as Tokeneer. [Barnes06]

1391

1392 **4.4 Conclusion**

1393 The call for a dramatic reduction in software vulnerability is heard from multiple sources,
1394 including the 2015 Cybersecurity Action Plan. This report has identified five approaches for
1395 achieving this goal. Each approach meets three criteria: 1) have a potential for dramatic
1396 improvement in software quality, 2) could make a difference in a three to seven-year timeframe
1397 and 3) are technical activities. The identified approaches use multiple strategies:

- 1398 • Stopping vulnerabilities before they occur generally including improved methods for
1399 specifying and building software.
- 1400 • Finding vulnerabilities including better testing techniques and more efficient use of
1401 multiple testing methods.
- 1402 • Reducing the impact of vulnerabilities by building architectures that are more resilient, so
1403 that vulnerabilities can't be meaningfully exploited.

1404 **Formal Methods.** Formal methods include multiple techniques based on mathematics and logic,
1405 ranging from parsing to type checking to correctness proofs to model-based development to
1406 correct-by-construction. While previously deemed too time-consuming, formal methods have
1407 become mainstream in many behind-the-scenes applications and show significant promise for
1408 both building better software and for supporting better testing.

1409 **System Level Security.** System Level Security reduces the impact that vulnerabilities have.
1410 Operating system containers and microservices are already a significant part of the national
1411 information infrastructure. Given the clear manageability, cost and performance advantages of

1412 using them, it is reasonable to expect their use to continue to expand. Security-enhanced versions
1413 of these technologies, if adopted, can therefore have a wide-spread effect on the exploitation of
1414 software vulnerabilities throughout the National Information Infrastructure.

1415 **Additive Software Analysis.** There are many types of software analysis – some are general and
1416 some target very specific vulnerabilities. The goal of additive software analysis is to be able to
1417 use multiple tools as part of an ecosystem. This will allow for increased growth and use of
1418 specialized software analysis tools and ability to gain a synergy between tools and techniques.

1419 **More Mature Domain-Specific Software Development Frameworks.** The goal of this
1420 approach is to promote the use (and reuse) of well-tested, well-analyzed code, and thus to reduce
1421 the incidence of exploitable vulnerabilities.

1422 **Moving Target Defenses (MTD) and Artificial Diversity.** This approach is a collection of
1423 techniques to vary the software's detailed structures and properties such that an attacker has
1424 much greater difficulty exploiting any vulnerability. The goal of artificial diversity and moving
1425 target defense (MTD) is to reduce an attacker's ability to exploit any vulnerabilities in the
1426 software, not to reduce the number of weaknesses in software.

1427 A critical need for improving security is to have software with fewer and less exploitable
1428 vulnerabilities. The measures, techniques and approaches we have described will be able to do
1429 this. Higher quality software, though, does not get created in a vacuum. There must be a robust
1430 research infrastructure, education and training, and customer pull. Higher quality software is a
1431 necessary step, but it is insufficient. A robust operation and maintenance agenda that spans a
1432 system's lifecycle is still needed.

1433

1434 **5 References**

- 1435 [Anderson72] James P. Anderson, “Computer Security Technology Planning Study,” Air force
1436 ESD-TR-73-51, Vol II, Oct. 1972.
- 1437 [Armstrong14] Robert C. Armstrong, Ratish J. Punnoose, Matthew H. Wong and Jackson R.
1438 Mayo, “Survey of Existing Tools for Formal Verification,” Sandia National Laboratories report
1439 SAND2014-20533, December 2014. [http://prod.sandia.gov/techlib/access-](http://prod.sandia.gov/techlib/access-control.cgi/2014/1420533.pdf)
1440 [control.cgi/2014/1420533.pdf](http://prod.sandia.gov/techlib/access-control.cgi/2014/1420533.pdf)
- 1441 [ASCMM16] “Automated Source Code Maintainability Measure™ (ASCMM™) V1.0”,
1442 <http://www.omg.org/spec/ASCMM/1.0>, January 2016
- 1443 [Bakker14] Paul Bakker, “Providing assurance and trust in PolarSSL”, 2014
1444 <https://tls.mbed.org/tech-updates/blog/providing-assurance-and-trust-in-polarssl>, accessed 21
1445 June 2016.
- 1446 [Barnes06] Janet Barnes, Rod Chapman, Randy Johnson, James Widmaier, David Cooper and
1447 Bill Everett, “Engineering the Tokeneer Enclave Protection Software”, Proc. 1st IEEE
1448 International Symposium on Secure Software Engineering (ISSSE), March 2006. Available at
1449 http://www.adacore.com/uploads/technical-papers/issse2006tokeneer_altran.pdf
- 1450 [Barnett05] M. Barnett, B.E. Chang, R. DeLine, B. Jacobs and K. R. M. Leino, “Boogie: A
1451 Modular Reusable Verifier for Object-Oriented Programs,” *Proc. International Symposium on*
1452 *Formal Methods for Components and Objects (FMCO)*, 2005
- 1453 [Barnum12] Sean Barnum, “Software Assurance Findings Expression Schema (SAFES)
1454 Overview,” January 2012. Available at [https://www.mitre.org/publications/technical-](https://www.mitre.org/publications/technical-papers/software-assurance-findings-expression-schema-safes-overview)
1455 [papers/software-assurance-findings-expression-schema-safes-overview](https://www.mitre.org/publications/technical-papers/software-assurance-findings-expression-schema-safes-overview), accessed 8 September
1456 2016.
- 1457 [Barritt16] Keith Barritt, “3 Lessons: FDA/FTC Enforcement Against Mobile Medical Apps,”
1458 January 2016. Available at [http://www.meddeviceonline.com/doc/lessons-fda-ftc-enforcement-](http://www.meddeviceonline.com/doc/lessons-fda-ftc-enforcement-against-mobile-medical-apps-0001)
1459 [against-mobile-medical-apps-0001](http://www.meddeviceonline.com/doc/lessons-fda-ftc-enforcement-against-mobile-medical-apps-0001)
- 1460 [Beck94] K. Beck, “Simple Smalltalk testing: with Patterns,” The Smalltalk Report, 1994.
- 1461 [Beizer90] Boris Beizer, “Software Testing Techniques,” 2nd ed., Van Nostrand Reinhold Co.
1462 New York, NY, ISBN: 0-442-20672-0.
- 1463 [Bell76] D.E. Bell and L. Lapadula, “Secure Computer System: Unified Exposition and Multics
1464 Interpretation,” Technical Report No. ESD-TR-75-306, Electronics Systems Division, AFSC,
1465 Hanscom AF Base, Bedford MA, 1976.

- 1466 [Biba77] K.J. Biba, “Integrity Considerations for Secure Computer systems,” USAF Electronic
1467 Systems Division, Bedford, MA, ESD-TR-76-372, 1977.
- 1468 [Bjørner16] Nikolaj Bjørner, “SMT Solvers: Foundations and Applications”, in *Dependable*
1469 *Software Systems Engineering*, J. Esparza et. al. eds., pp 24-32, IOS Press, 2016. DOI:
1470 10.3233/978-1-61499-627-9-24
- 1471 [Black11a] Paul E. Black, Michael Kass, Michael Koo and Elizabeth Fong, “Source Code
1472 Security Analysis Tool Functional Specification Version 1.1,” NIST Special Publication (SP)
1473 500-268 v1.1, February 2011, DOI: 10.6028/NIST.SP.500-268v1.1
- 1474 [Black11b] Paul E. Black, “Counting Bugs is Harder Than You Think,” 11th IEEE Int’l Working
1475 Conference on Source Code Analysis and Manipulation (SCAM 2011), September 2011,
1476 Williamsburg, VA.
- 1477 [Black16] Paul E. Black and Elizabeth Fong, “Report of the Workshop on Software Measures
1478 and Metrics to Reduce Security Vulnerabilities (SwMM-RSV),” NIST Special Publication (SP)
1479 500-xxx, October 2016.
- 1480 [Boebert85] W.E. Boebert and R.Y. Kain, “A Practical Alternative to Hierarchical Integrity
1481 Policies,” Proc. 8th National Computer Security Conference, Gaithersburg, MD, p.18, 1985.
- 1482 [Bojanova16] Irena Bojanova, Paul E. Black, Yaacov Yesha and Yan Wu, “The Bugs
1483 Framework (BF): A Structured Approach to Express Bugs,” 2016 IEEE Int’l Conf. on Software
1484 Quality, Reliability, and Security (QRS 2016), Vienna, Austria, August 1-3, 2016. Available at
1485 <https://samate.nist.gov/BF>, accessed 12 September 2016.
- 1486 [Boulanger12] Jean-Louis Boulanger (Ed.) Industrial Use of Formal Methods: Formal
1487 Verification, July 2012, Wiley-ISTE.
- 1488 [Busoli07] Simone Busoli, “Inversion of Control and Dependency Injection with Castle Windsor
1489 Container,”
1490 <http://dotnetlackers.com/articles/designpatterns/InversionOfControlAndDependencyInjectionWithCastleWindsorContainerPart1.aspx>, July 2007. Accessed 29 September 2016.
- 1492 [Brooks95] F. Brooks, The Mythical Man-Month, Anniversary edition with 4 new chapters,
1493 Addison-Wesley, 1995.
- 1494 [Clang] “Clang: A C language family frontend for LLVM,” <http://clang.llvm.org/>
- 1495 [CodeDx15] “Finding Software Vulnerabilities Before Hackers Do,” available at
1496 <https://codedx.com/wp-content/uploads/2015/10/AppSec101-FromCodeDx.pdf>, Accessed 8
1497 September 2016.

- 1498 [Corbato65] F.J. Corbato and V.A. Vyssotsky, “Introduction and Overview of the Multics
1499 System,” 1965 Fall Joint Computer Conference. Available at <http://multicians.org/fjcc1.html>
- 1500 [Docker16] “Docker,” <https://www.docker.com/>
- 1501 [Doyle16] Richard Doyle, “Formal Methods, including Model-Based Verification and Correct-
1502 By-Construction,” Dramatically Reducing Security Vulnerabilities sessions, Software and
1503 Supply Chain Assurance (SSCA) Working Group Summer 2016, McLean, Virginia, July 2016,
1504 Available at https://samate.nist.gov/docs/DRSV2016/SSCA_07_JPL_FormalMethods_Doyle.pdf
- 1505 [FCRDSP16] Federal Cybersecurity Research and Development Strategic Plan. Available at
1506 [http://www.whitehouse.gov/sites/whitehouse.gov/files/documents/2016_Federal_Cybersecurity_](http://www.whitehouse.gov/sites/whitehouse.gov/files/documents/2016_Federal_Cybersecurity_Research_and_Development_Strategic_Plan.pdf)
1507 [Research_and_Development_Strategic_Plan.pdf](http://www.whitehouse.gov/sites/whitehouse.gov/files/documents/2016_Federal_Cybersecurity_Research_and_Development_Strategic_Plan.pdf)
- 1508 [Flater16] David Flater, Paul E. Black, Elizabeth Fong, Raghu Kacker, Vadim Okun, Stephen
1509 Wood and D. Richard Kuhn, “A Rational Foundation for Software Metrology,” NIST Internal
1510 Report (IR) 8101, January 2016. Available at <https://doi.org/10.6028/NIST.IR.8101>
- 1511 [Fowler14] Martin Fowler, “Microservices: a definition of this new architectural term”,
1512 <http://martinfowler.com/articles/microservices.html>, March 2014
- 1513 [FramaC] “What is Frama-C”, http://frama-c.com/what_is.html
- 1514 [FTC16] Federal Trade Commission, “Mobile Health App Developers: FTC Best Practices,”
1515 April 2016. Available at [http://www.ftc.gov/tips-advice/business-center/guidance/mobile-health-](http://www.ftc.gov/tips-advice/business-center/guidance/mobile-health-app-developers-ftc-best-practices)
1516 [app-developers-ftc-best-practices](http://www.ftc.gov/tips-advice/business-center/guidance/mobile-health-app-developers-ftc-best-practices)
- 1517 [Gabriel96] Richard P. Gabriel, “Patterns of Software: Tales from the Software Community,”
1518 Oxford Press, 1996.
- 1519 [Gamma95] Erich Gamma, Richard Helm, Ralph Johnson and John Vlissides, “Design Patterns:
1520 Elements of Reusable Object-Oriented Software”, Addison-Wesley, 1995.
- 1521 [GCC16] The GNU Compiler Collection, available at <https://gcc.gnu.org/>
- 1522 [Goguen84] J.A. Goguen and J. Meseguer, “Unwinding and Inference Control,” Proc. 1984
1523 IEEE Symposium on Security and Privacy, 1984.
- 1524 [Hurd16] “Hurd”, Web Services Architecture Workshop Group,
1525 <https://www.gnu.org/software/hurd/hurd.html>, 2016.
- 1526 [Jézéquel97] Jean-Marc Jézéquel and Bertrand Meyer, “Design by Contract: the Lesson of
1527 Ariane,” IEEE Computer, 30(1):129-130, January 1997. DOI: 10.1109/2.562936

- 1528 [Kastrinis14] G. Kastrinis and Y. Smaragdakis, “Hybrid Context-Sensitivity for Points-To
1529 Analysis,” *Proc. Conference on Programming Language Design and Implementation (PLDI)*,
1530 2014.
- 1531 [KDM15] “Knowledge Discovery Metamodel (KDM),” available at
1532 <http://www.omg.org/technology/kdm/>, accessed 8 September 2016.
- 1533 [Kiniry08] Joseph Kiniry and Daniel Zimmerman, “Secret Ninja Formal Methods,” 15th Int’l
1534 Symp. on Formal Methods (FM’08). Turku, Finland. May, 2008.
- 1535 [Klass16] Gregory Klass and Eric Burger, “Vendor Truth Serum”, Georgetown University,
1536 <https://s2erc.georgetown.edu/projects/vendortruthserum> Accessed 19 September 2016.
- 1537 [Klein14] Gerwin Klein, June Andronick, Kevin Elphinstone, Toby Murray, Thomas Sewell,
1538 Rafal Kolanski and Gernot Heiser, “Comprehensive Formal Verification of an OS Microkernel,”
1539 NICTA and UNSW, Sydney, Australia, *ACM Transactions on Computer Systems*, Vol. 32, No.
1540 1, Article 2, Publication date: February 2014.
- 1541 [Kuhn10] Richard Kuhn, Raghu Kacker and Y. Lei, “Practical Combinatorial Testing”, NIST
1542 Special Publication 800-142, Oct. 2010. Available at
1543 <http://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-142.pdf>
- 1544 [Lampson04] B.W. Lampson, “Software components: Only the Giants survive,” *Computer*
1545 *Systems: Theory, Technology, and Application*, Springer, 2004.
- 1546 [Lemon13] Lemon, “Getting Started with LXC on an Ubuntu 13.04 VPS”,
1547 [https://www.digitalocean.com/community/tutorials/getting-started-with-lxc-on-an-ubuntu-13-04-](https://www.digitalocean.com/community/tutorials/getting-started-with-lxc-on-an-ubuntu-13-04-vps)
1548 [vps](https://www.digitalocean.com/community/tutorials/getting-started-with-lxc-on-an-ubuntu-13-04-vps), August 2013.
- 1549 [LXC] “LXC”, Ubuntu 16.04 Server Guide, <https://help.ubuntu.com/lts/serverguide/lxc.html>
1550 Accessed 27 September 2016.
- 1551 [Marien16] John R. Marien, Chair, and Robert A. Martin, Co-chair, “Suggested Language to
1552 Incorporate Software Assurance Department of Defense Contracts,” Department of Defense
1553 (DoD) Software Assurance (SwA) Community of Practice (CoP) Contract Language Working
1554 Group, February 2016. Available at [http://www.acq.osd.mil/se/docs/2016-02-26-SwA-](http://www.acq.osd.mil/se/docs/2016-02-26-SwA-WorkingPapers.pdf)
1555 [WorkingPapers.pdf](http://www.acq.osd.mil/se/docs/2016-02-26-SwA-WorkingPapers.pdf) Accessed 6 September 2016.
- 1556 [McConnell04] Steve McConnell, “Code Complete,” 2nd Ed., Microsoft Press, 2004. ISBN:
1557 0735619670
- 1558 [McIlroy68] D. McIlroy, “Mass Produced Software Components,” 1968 NATO Conference on
1559 Software Engineering.

- 1560 [Oberg99] James Oberg, “Why the Mars probe went off course,” IEEE Spectrum, 36(12):34-39,
1561 December 1999. DOI: 10.1109/6.809121
- 1562 [Okhravi13] H. Okhravi, M. A. Rabe, T. J. Mayberry, W. G. Leonard, T. R. Hobson, D. Bigelow
1563 and W. W. Streilein, “Survey of Cyber Moving Targets,” Massachusetts Institute of Technology
1564 Lincoln Laboratory. Technical Report 1166, September 2013, Available
1565 [https://www.ll.mit.edu/mission/cybersec/publications/publications:files/full_papers/2013_09_23](https://www.ll.mit.edu/mission/cybersec/publications/publications:files/full_papers/2013_09_23_OkhraviH_TR_FP.pdf)
1566 [_OkhraviH_TR_FP.pdf](https://www.ll.mit.edu/mission/cybersec/publications/publications:files/full_papers/2013_09_23_OkhraviH_TR_FP.pdf) Accessed 15 September 2015.
- 1567 [Okun04] Vadim Okun, Paul E. Black and Yaacov Yesha, “Comparison of Fault Classes in
1568 Specification-Based Testing,” Information and Software Technology, Elsevier, 46(8):525-533,
1569 June 2004.
- 1570 [Okun08] Vadim Okun, Romain Gaucher and Paul E. Black, eds., “Static Analysis Tool
1571 Exposition (SATE) 2008,” NIST Special Publication (SP) 500-279, June 2009, DOI:
1572 10.6028/NIST.sp.500-279
- 1573 [Pal05] Partha Pal, Michael Atigetchi, Jennifer Chong, Franklin Webber and Paul Rubel,
1574 “Survivability Architecture of a Mission Critical System: The DPASA Example”, 21st Annual
1575 Computer Security Applications Conference (ACSAC 2005), pages 495-504, ISSN: 1063-9527,
1576 ISBN: 0-7695-2461-3, DOI: 10.1109/CSAC.2005.54, December 2005.
- 1577 [PaX01] “Design and Implementation of Address Space Layout randomization,”
1578 <https://pax.grsecurity.net/docs/aslr.txt>, cited in “Address space layout randomization,”
1579 Wikipedia. Available at https://en.wikipedia.org/wiki/Address_space_layout_randomization
1580 Access 15, September 2016.
- 1581 [Perini16] Barti Perini, Stephen Shook and Girish Seshagiri, “Reducing Software Vulnerabilities
1582 – The Number One Goal for Every Software Development Organization, Team, and Individual,”
1583 ISHIPI technical Report, 22 July 2016.
- 1584 [Pirsig74] Robert M. Pirsig, “Zen and the Art of Motorcycle Maintenance: An Inquiry into
1585 Values,” William Morrow & Company, 1974.
- 1586 [Rashid86] R. Rashid, “Threads of a New System,” Unix Review, Vol 4, No. 8, August 1986.
- 1587 [Regehr15] John Regehr, “Comments on a Formal Verification of PolarSSL”, 2015
1588 <http://blog.regehr.org/archives/1261>, accessed 21 June 2016.
- 1589 [Rose16] “ROSE compiler infrastructure,” <http://rosecompiler.org/>, accessed 8 September 2016.
- 1590 [Rushby05] John Rushby, “An Evidential Tool Bus,” *Proc. International Conference on Formal*
1591 *Engineering Methods*, 2005.

- 1592 [Saltzer75] J. Slatzer and M. Shroeder, “The Protection of Information in Computer systems,”
1593 Proc. IEEE vol. 63, issue 9, p. 1278-1308, September 1975.
- 1594 [SMTLIB15] “SMT-LIB: The Satisfiability Modulo Theories Library”,
1595 <http://smtlib.cs.uiowa.edu/>, June 2015.
- 1596 [Software16] “Software framework”, https://en.wikipedia.org/wiki/Software_framework
- 1597 [Tschannen11] J. Tschannen, C. A. Furia, M. Nordio and B. Meyer, “Verifying Eiffel Programs
1598 with Boogie,” *Boogie: First International Workshop on Intermediate Verification Languages*,
1599 2011.
- 1600 [TodoMVC16] “TodoMVC: Helping you select an MV* framework”, <http://todomvc.com/>
- 1601 [Tolerant07] “Tolerant Systems,” <http://www.tolerantsystems.org/>
- 1602 [VCC13] VCC verifier, available at <https://vcc.codeplex.com/>
- 1603 [Voas16a] Jeffrey Voas and Kim Schaffer, “Insights on Formal Methods in Cybersecurity”,
1604 IEEE Computer 49(5):102–105, May 2016, DOI: 10.1109/MC.2016.131
- 1605 [Voas16b] Jeffrey Voas and Kim Schaffer, (Insights, part 2), IEEE Computer, August 2016.
- 1606 [Wayner15] Peter Wayner, “7 reasons why frameworks are the new programming languages”,
1607 [http://www.infoworld.com/article/2902242/application-development/7-reasons-why-](http://www.infoworld.com/article/2902242/application-development/7-reasons-why-frameworks-are-the-new-programming-languages.html)
1608 [frameworks-are-the-new-programming-languages.html](http://www.infoworld.com/article/2902242/application-development/7-reasons-why-frameworks-are-the-new-programming-languages.html), March 2015.
- 1609 [What] “What’s LXC?”, <https://linuxcontainers.org/lxc/introduction/>
- 1610 [Whaley05] J. Whaley, D. Avots, M. Carbin and M. S. Lam, “Using Datalog with Binary
1611 Decision Diagrams for Program Analysis,” *Proc. Programming Languages and Systems*
1612 *(ASPLAS)*, 2005.
- 1613 [Williams16] Chris Williams, “How one developer just broke Node, Babel and thousands of
1614 projects in 11 lines of JavaScript,” http://www.theregister.co.uk/2016/03/23/npm_left_pad_chaos
- 1615 [Woodcock09] Jim Woodcock, Peter Gorm Larsen, Juan Bicarregui and John Fitzgerald,
1616 “Formal Methods: Practice and Experience,” *ACM Computing Surveys*, 41(4), October 2009,
1617 Article No. 19, DOI: 10.1145/1592434.1592436. Available at
1618 <http://homepage.cs.uiowa.edu/~tinelli/classes/181/Fall14/Papers/Wood09.pdf>
- 1619 [Woodcock10] Jim Woodcock, Emine Gökçe Aydal and Rod Chapman, “The Tokeneer
1620 Experiments”, in *Reflections on the Work of C.A.R. Hoare, History of Computing*, Chapter 17,
1621 pp 405-430, July 2010, DOI: 10.1007/978-1-84882-912-1_17.

- 1622 [Woody14] Carol Woody, R. Ellison and W. Nichols, “Predicting Software Assurance Using
1623 Quality and Reliability Measures,” CMU/SEI-2014-TN-026, Dec. 2014.
- 1624 [WSA04] “Web Services Architecture”, <https://www.w3.org/2002/ws/arch/>
- 1625 [Wu11] Yan Wu, H. Siy and Robin A. Gandhi, “New Ideas and Emerging Results Track:
1626 Empirical Results on the Study of Software Vulnerabilities,” New Ideas and Emerging Results
1627 Track at the 33rd International Conference on Software Engineering (ICSE 2011), Honolulu,
1628 Hawaii, May 21-28, 2011.
- 1629 [Yang10] J. Yang and C. Hawblitzel, “Safe to the last instruction: automated verification of a
1630 type-safe operating system,” in *Proc. 31st ACM SIGPLAN Conference on Programming
1631 Language Design and Implementation (PLDI)*, 2010.
- 1632 [Zhu97] Hong Zhu, Patrick A. V. Hall and John H. R. May, “Software Unit Test Coverage and
1633 Adequacy,” *ACM Computing Surveys (CSUR)*, 29(4): 366-427, December 1997, DOI:
1634 10.1145/267580.267590
- 1635