

BSIMM15

REPORT 2025



TABLE OF CONTENTS

PART 1: EXECUTIVE SUMMARY5

| | |
|---|---|
| WELCOME TO BSIMM15..... | 5 |
| BSIMM Terminology..... | 5 |
| BSIMM15 DATA HIGHLIGHTS | 6 |
| TRENDS AND INSIGHTS SUMMARY | 7 |
| Meeting Government Regulations..... | 7 |
| Secure Innovation..... | 7 |
| CALL TO ACTION | 7 |
| Plan Your Journey | 7 |
| Get a Handle on What You Have | 7 |
| Make the Right Investments..... | 7 |
| Some Reading Suggestions for This Report..... | 8 |
| THE BSIMM SKELETON | 9 |

PART 2: TRENDS AND INSIGHTS12

| | |
|--|----|
| THE EVOLUTION OF SHIFT EVERYWHERE | 12 |
| Governance and Automation..... | 12 |
| WHAT IS SHIFT EVERYWHERE, REALLY? | 12 |
| Enabling People..... | 13 |
| SOFTWARE SUPPLY CHAIN RISK MANAGEMENT | 13 |
| Software Bills of Materials..... | 13 |
| Vendor Management and Bespoke Software..... | 13 |
| REGULATORY PRESSURES | 14 |
| WRESTLING WITH NEW TECHNOLOGIES | 14 |
| FINDING PROBLEMS IN NEW TECHNOLOGIES..... | 14 |
| SECURITY SOLUTIONS FOR NEW SERVICES..... | 14 |
| LONG-TERM TRENDS | 15 |
| Expanding BSIMM's Reach | 15 |
| Software Security Awareness Training's Long Decline | 15 |
| TOPICS WE'RE WATCHING..... | 15 |

PART 3: FOR PRACTITIONERS17

| | |
|---|----|
| INTRODUCTION | 17 |
| REGULATIONS FOR PRACTITIONERS..... | 17 |
| THE REGULATORY LANDSCAPE | 17 |
| SSDF-BSIMM MAPPING UPDATED FOR BSIMM15..... | 18 |
| AI/ML..... | 19 |
| AI/ML EXPERT INSIGHTS: THE UNCERTAINTY PROBLEM | 19 |
| AI/ML AND BSIMM15 | 20 |
| Five Key Activities: Using the BSIMM to Deal with AI/ML..... | 21 |

| | |
|--|----|
| HOW TO BUILD OR UPGRADE AN SSI | 22 |
| CONSTRUCTION LESSONS FROM OUR PARTICIPANTS | 23 |
| Cultures..... | 23 |
| A New Wave in Engineering Culture..... | 23 |
| Understanding More About DevOps..... | 24 |
| Convergence as a Goal | 24 |
| FOR AN EMERGING SSI: SDLC TO SSDL | 25 |
| CHECKLIST FOR EMERGING SSIs | 26 |
| FOR A MATURING SSI: HARMONIZING OBJECTIVES..... | 28 |
| CHECKLIST FOR MATURING SSIs..... | 29 |
| FOR AN ENABLING SSI: DATA-DRIVEN IMPROVEMENTS | 31 |
| PARTICIPANTS | 34 |

PART 4: BSIMM PARTICIPANTS34

| | |
|--|----|
| ACKNOWLEDGEMENTS..... | 35 |
| A Baseline for SSI Leaders | 37 |
| Using A BSIMM Scorecard to Make Progress | 37 |

PART 5: QUICK GUIDE TO SSI MATURITY.....37

| | |
|---|----|
| Understand Your Organizational Mandate..... | 38 |
| Build the Scorecard | 38 |
| Make a Strategic Plan and Execute..... | 38 |

PART 6: SOFTWARE SECURITY INITIATIVE BY THE NUMBERS.....41

| | |
|--|----|
| ROLES IN A SOFTWARE SECURITY INITIATIVE | 41 |
| EXECUTIVE LEADERSHIP | 41 |
| SOFTWARE SECURITY GROUP LEADERS | 42 |
| SOFTWARE SECURITY GROUP (SSG) | 43 |
| SECURITY CHAMPIONS (SATELLITE)..... | 43 |
| OTHER KEY STAKEHOLDERS | 44 |


PART 7: THE BSIMM FRAMEWORK46

| | |
|--|----|
| CORE KNOWLEDGE..... | 46 |
| UNDERSTANDING THE MODEL | 47 |
| DETAILED VIEW OF THE BSIMM FRAMEWORK | 48 |
| The Detailed BSIMM Skeleton | 48 |
| CREATING BSIMM15 FROM BSIMM14..... | 51 |
| MODEL CHANGES OVER TIME | 52 |
| Where Do Old Activities Go? | 52 |

| | |
|---|---------------|
| DATA: BSIMM15 | 56 |
| Reading a Boxplot | 56 |
| Age-Based Program Changes | 57 |
| ACTIVITY CHANGES OVER TIME | 59 |
| PART 8: THE BSIMM ACTIVITIES | 62 |
| ACTIVITIES IN THE BSIMM | 62 |
| GOVERNANCE | 62 |
| Governance: Strategy & Metrics (SM) | 62 |
| Governance: Compliance & Policy (CP) | 64 |
| Governance: Training (T) | 66 |
| INTELLIGENCE | 68 |
| Intelligence: Attack Models (AM) | 68 |
| Intelligence: Security Features & Design (SFD) | 70 |
| Intelligence: Standards & Requirements (SR) | 71 |
| SDLC TOUCHPOINTS | 73 |
| SDLC Touchpoints: Architecture Analysis (AA) | 73 |
| SDLC Touchpoints: Code Review (CR) | 75 |
| SDLC Touchpoints: Security Testing (ST) | 76 |
| DEPLOYMENT | 78 |
| Deployment: Penetration Testing (PT) | 78 |
| Deployment: Software Environment (SE)..... | 79 |
| Deployment: Configuration Management & Vulnerability Management (CMVM) | 81 |

| | |
|---|------------|
| PART 9: WHAT CAN THE DATA TELL US..... | 85 |
| DATA ANALYSIS: SSG | 85 |
| SSG Characteristics | 85 |
| SSG Changes Based on Age..... | 87 |
| DATA ANALYSIS: SECURITY CHAMPIONS | 90 |
| DATA ANALYSIS: VERTICALS AND PRACTICES | 93 |
| VERTICAL SCORECARDS | 94 |
| PENETRATION TESTING | 99 |
| COMPLIANCE & POLICY..... | 99 |
| ARCHITECTURE ANALYSIS | 100 |
| ATTACK MODELS | 100 |
| TRAINING | 101 |
| IOT, TECH, ISV, AND FINTECH | 102 |
| FINTECH AND FINANCIAL..... | 105 |
| DATA ANALYSIS: LONGITUDINAL | 107 |
| Building a Model for Software Security | 107 |
| Changes Between First and Second Assessments | 108 |
| Changes Between First and Third Assessments | 110 |

PART I: EXECUTIVE SUMMARY



PART I: EXECUTIVE SUMMARY

In 2008, application security, research, and analysis experts set out to gather data on the different paths that organizations take to address the challenges of securing software. Their goal was to conduct in-person interviews with organizations that were known to be highly effective in software security initiatives (SSIs), gather details about their efforts, analyze the data, and publish their findings to help others.

The result was the Building Security In Maturity Model (BSIMM), a descriptive model—published as BSIMM1—that provides a baseline of observed activities (i.e., controls) for SSIs to build security into software and software development. Because these initiatives often use different methodologies and different terminology, the BSIMM also creates a common vocabulary everyone can use. In addition, the BSIMM provides a common methodology for starting and improving SSIs of any size and in any vertical market.

Since BSIMM1 in 2009, we've been early reporters on security program changes across people, process, technology, culture, compliance, digital transformation, and much more. Welcome to the BSIMM15 report, and thank you for reading.

WELCOME TO BSIMM15

If you're in charge of an SSI, understanding the BSIMM and its use by participants will help you plan strategic improvements. If you're running the technical aspects of an initiative, you can use the how-to guide (in Part 5) and activity descriptions (in Part 8) to help define tactical improvements to people, process, technology, and culture.

Each BSIMM annual report is the result of studying real-world SSIs, which many organizations refer to as their application or product security program or as their DevSecOps effort. Each year, firms in different industry verticals use the BSIMM to create a software security scorecard for their programs that they then use to inform their SSI improvements. Here, we present BSIMM15 as built directly out of the data we observed in 121 firms.

In the rapidly changing software security field, it's important to understand what other organizations are doing in their own SSIs. Comparing the efforts of more than 100 companies to yours will directly inform your strategy for improvement and growth.

BSIMM core knowledge is the activities we have directly observed in our participants—the group of firms that use the BSIMM as part of their SSI management. Each participant has their own unique SSI with an emphasis on the building security in activities important to their business objectives, but they collectively use the activities captured here. We organize that core knowledge into a software security framework (SSF), represented in Part 8. The SSF comprises four domains—Governance, Intelligence, SSDL Touchpoints, and Deployment—with those domains currently composed of 128 activities. The Governance domain, for example, includes activities that fall under the organization, management, and measurement efforts of an SSI.

From an executive perspective, you can view BSIMM activities as preventive, detective, corrective, or compensating controls

implemented in a software security risk management framework. Positioning the activities as controls allows for easier understanding of the BSIMM's value by governance, risk, compliance, legal, audit, and other executive management groups.

As with any research work, some terms have specific meanings in the BSIMM. The box below shows the most common BSIMM terminology.

BSIMM TERMINOLOGY

Nomenclature has always been a problem in computer security, and software security is no exception. Several terms used in the BSIMM have particular meaning for us, and the following list highlights some of the most important ones used in this document:

- **Activity.** Actions or efforts carried out or facilitated by the SSG as part of a practice. Activities are divided into three levels in the BSIMM based on observation rates.
- **Capability.** A set of BSIMM activities spanning one or more practices working together to serve a cohesive security function.
- **Champions.** A group of interested and engaged developers, cloud security engineers, deployment engineers, architects, software managers, testers, or people in similar roles who have an active interest in software security and contribute to the security posture of the organization and its software.
- **Data pool.** The collection of assessment data from the current participants.
- **Domain.** One of the four categories the framework is divided into, i.e., Governance, Intelligence, SSDL Touchpoints, and Deployment.
- **Participants.** The group of firms in the current data pool.
- **Practice.** A grouping of BSIMM activities. The SSF is organized into 12 practices, three in each of four domains.
- **Satellite.** A group of individuals, often called security champions, that is organized and leveraged by an SSG. This has largely been replaced by the term security champions.
- **Secure SDL (SSDL).** Any software lifecycle with integrated software security checkpoints and activities.
- **Software security framework (SSF).** The basic structure underlying the BSIMM, comprising 12 practices divided into four domains.
- **Software security group (SSG).** The internal group charged with carrying out and facilitating software security. The group's name might also have an appropriate organizational focus, such as application security group or product security group.
- **Software security initiative (SSI).** An organization-wide program to instill, measure, manage, and evolve software security activities in a coordinated fashion. Also referred to in some organizations as an application security program, product security program, or perhaps as a DevSecOps program.

BSIMM15 DATA HIGHLIGHTS

Use the information in this section to answer common questions about BSIMM data, such as, “What are some data pool statistics?,” “Which activities are most firms doing?,” and “How are software security efforts changing over time?”

Note: Items in italic purple refer to specific BSIMM activities in Part 8.

Activities are the building blocks of the BSIMM, the smallest units of granularity implemented across organizations to build SSIs. Rather than dictating a set of prescriptive activities, the purpose of the BSIMM is to descriptively observe and quantify the actual activities carried out by various kinds of SSIs across many organizations.

The BSIMM is an observational model that reflects current software security efforts, so we adjust it annually to keep it current. For BSIMM15, we’ve made the following changes to the model based on what we see in the BSIMM data pool:

- We moved the activities *enforce security checkpoints and track exceptions, define secure deployment parameters and configurations, have emergency response, and streamline incoming responsible vulnerability disclosure* because we now see them more frequently.
- We added the activities *create standards controlling and guiding the adoption of new technologies* and *protect the integrity of development endpoints* because we are beginning to see them more.

Unique in the software security industry, the BSIMM project has grown from nine participating companies in 2008 to 121 in 2024, now with approximately 3,400 software security group (SSG) members and 7,700 security champions. The average age of participants’ SSIs is 5.4 years, but the BSIMM project shows consistent growth even as participants enter and leave over time—we added nine firms for BSIMM15 and dropped 18 others whose data hadn’t been refreshed (see Figure 21). Increasingly, we see firms not just assessing their overall program but how well they reach into their business units and teams as well. While these business unit assessments provide valuable data to the firms that request them, the data from these assessments is not added to the data pool recorded in the annual report.

This 2024 edition of the BSIMM report—BSIMM15—examines anonymized data from the software security activities of 121 organizations across different verticals, including cloud, financial services, financial technology (FinTech), healthcare, independent software vendors (ISVs), insurance, Internet of Things (IoT), and technology organizations.

The *7 Habits of Highly Effective People* explores the theory that successful individuals share common qualities in achieving their goals and that these qualities can be identified and applied by others. The same premise can be applied to SSIs. Listed in Table 1 are the 10 most observed activities in the BSIMM15 data pool. The data suggests that if your organization is working on its own SSI, you should consider implementing these activities.

Table 2 shows some activities that have experienced exceptionally high growth over the past 12 months. Not surprisingly, some of these activities, such as *integrate*

software-defined lifecycle governance and *provide expertise via open collaboration channels*, are mentioned in the Trends and Insights section. In addition, the *streamline incoming responsible vulnerability disclosure* activity introduced in BSIMM12 continues its strong growth in observation count as we reported last year. Note that for some of the activities in Table 2, the growth in observation is a relatively new change. For example, the activity *have a research group that develops new attack methods* saw virtually no growth from BSIMM9 to BSIMM12 but had a significant jump in observation rates in BSIMM14, with BSIMM15 continuing that climb. Two years of growth suggests the change is meaningful and the activities are worth considering for your program.

In BSIMM14, we reported new growth after little change over time in the *enforce security checkpoints and track exceptions* activity. This activity has continued to grow in BSIMM15 as firms are able to take advantage of modern automation options in the development pipeline.

| BSIMM15 TOP 10 ACTIVITIES | |
|---------------------------|---|
| PERCENT | DESCRIPTION |
| 91.7 | [CMVM1.1] Create or interface with incident response. |
| 90.1 | [SM1.4] Implement security checkpoints and associated governance. |
| 87.6 | [CR1.4] Use automated code review tools. |
| 86.8 | [CP1.2] Identify privacy obligations. |
| 86.0 | [PT1.1] Use external penetration testers to find problems. |
| 84.3 | [ST1.1] Perform edge/boundary value condition testing during QA. |
| 84.3 | [SE1.2] Ensure host and network security basics are in place. |
| 81.8 | [AA1.1] Perform security feature review. |
| 81.0 | [CP1.1] Unify regulatory pressures. |
| 79.3 | [SR1.2] Create a security portal. |

TABLE 1. TOP ACTIVITIES BY OBSERVATION PERCENTAGE. The most frequently observed activities in BSIMM15 are likely important to all SSIs.

| BSIMM15 TOP 10 ACTIVITIES GROWTH BY COUNT | |
|---|---|
| INCREASE | DESCRIPTION |
| 10 | [T2.12] Provide expertise via open collaboration channels. |
| 6 | [SE1.4] Define secure deployment parameters and configurations. |
| 6 | [SE2.4] Protect code integrity. |
| 6 | [SE3.2] Use code protection. |
| 6 | [CMVM2.4] Streamline incoming responsible vulnerability disclosure. |
| 5 | [CP3.2] Ensure compatible vendor policies. |
| 4 | [AM2.8] Have a research group that develops new attack methods. |
| 4 | [AA2.2] Standardize architectural descriptions. |
| 4 | [CMVM3.3] Simulate software crises. |
| 3 | [SM3.4] Integrate software-defined lifecycle governance. |

TABLE 2. TOP ACTIVITIES BY RECENT GROWTH IN OBSERVATION COUNT. These activities had the largest growth in BSIMM15, which means they are likely important to your program now or will be soon.

TRENDS AND INSIGHTS SUMMARY

The BSIMM trends and insights we've identified are a distillation of software security lessons learned across 121 organizations that collectively have 11,100 security professionals helping about 270,000 developers do good security work on about 96,000 applications. Use this information to inform your own strategy for improvement.

Trends describe shifts in SSI behavior that affect activity implementation across multiple areas. Larger in scope than an activity, or even a capability that combines multiple activities within a workflow, we believe these trends show the way organizations are executing groups of activities within their evolving culture. In addition to the in-depth sections for practitioners on regulatory impact on application security in Part 3, which also covers how application security and the BSIMM are evolving in response to artificial intelligence (AI), we report on trends in more depth in Part 2.

MEETING GOVERNMENT REGULATIONS

In response to self-attestation requirements for selling software to the US government, many firms have put special emphasis on activities that can aid in compliance. One such mandate is the need to build software bills of materials (SBOMs), and we observed a 22% increase in [SE3.6] **create bills of materials for deployed software** and a 67% increase in observations of [SE3.8] **perform application composition analysis on code repositories**. Other areas include protecting source code, with observations of [SE2.4] **protect code integrity** increasing by ~20% while [SE3.2] **use code protection** increased by ~45%.

SECURE INNOVATION

In BSIMM15, we introduced [SR3.5] **create standards controlling and guiding the adoption of new technologies**, to measure how firms build security before the industry comes to a consensus on best practices. However, other activities can aid firms in preparing for these new challenges and opportunities. Building and testing potential attacks, for example, are on the rise, with observations of [AM2.8] **have a research group that develops new attack methods** increasing by ~30% and [ST3.5] **begin to build and apply adversarial tests (abuse cases)** observations doubling from BSIMM14 to BSIMM15. As firms capture innovative solutions in the form of secure design patterns, observations of [SFD3.3] **find and publishing secure design patterns from the organization** increased by ~45% and [SE1.4] **define secure deployment parameters and configurations** saw a ~15% increase over BSIMM14.

Take stock of your SSI. It's important to periodically look at your program through a different lens.

CALL TO ACTION

Use the information in this section to prioritize improvements in your SSI and perhaps also in the SSIs of your most important software suppliers and partners.

Every SSI has room for improvement, whether it's improving scale, effectiveness, depth, risk management, the framework of deployed activities, resourcing, or anything similar. The following suggestions represent the broad efforts we see happening in BSIMM participants, with various parts likely right for your program as well.

PLAN YOUR JOURNEY

- Take stock of your SSI. It's important to periodically look at your program through a different lens, and the BSIMM enables that. Use the guidance in Part 5 to create your own SSI scorecard and compare it to your expectations.
- Create a vision and a strategic plan. Use the activity descriptions in Part 8 when creating a prioritized action plan for business areas where your current SSI efforts fall short. Typical investment areas include risk management, digital transformation, technical debt removal, technology insertion, and process improvement.

GET A HANDLE ON WHAT YOU HAVE

- Inventory all your code. It's likely that you'll need specialized automation to keep track of all the code you write and all the code you bring in from outside the organization. A simple application inventory will be useful for some things, such as naming risk managers, but you'll quickly need specialized inventories, such as SBOMs, API and microservices lists, various as-code artifacts, code that is subject to specific compliance needs, and much more.
- Automate, automate, automate. Search for ways to eliminate error-prone manual processes and reduce friction between governance and engineering groups, including automating security decisions. This will require some policy-as-code effort and tools integration and might require bringing development skills into the SSG.
- Gather all the data. As more processes become code, and more policies and standards become machine-readable, day-to-day development and operations will generate significantly more telemetry about what's happening and why. Use this data to ensure that everything's working as expected.

MAKE THE RIGHT INVESTMENTS

- Innovate in digital transformation. Encourage your SSG and other security stakeholders to experiment with ways to deliver security value directly into engineering processes, especially where current security testing tools don't always keep up with engineering changes, such as with serverless architectures, single-page applications, AI, and zero trust.
- Secure the software supply chain. Nearly every organization today uses third-party code and provides code as a third

party to other organizations. While producing SBOMs is easy, the management of software, SBOMs themselves, vendors, and vulnerability information is much more complicated.

- Expand software security into adjacencies. Even perfect software can have its security undermined by mistakes elsewhere in the organization. Make explicit ties between the SSI and other security stakeholders working in areas such as container security, orchestration security, cloud security, infrastructure security, and site reliability.

In summary, the data shows that new SSIs—from just started to 18 months old—are typically doing about 33 BSIMM activities. These organizations are also beginning to scale these activities across their software portfolio, deal with all the change going on around them, and evolve their risk management strategy.

SOME READING SUGGESTIONS FOR THIS REPORT

- If you're experienced with the BSIMM, or if you need some content to help make your case with executive management, then Part 2 (Trends and Insights) is probably what you're looking for.
- If this is your first time with the BSIMM, we recommend first reading Part 5 for context and then returning here to decide what to read next.
- If you're starting an SSI or an SSG, or looking to mature an existing program, start with Part 5 (Quick Guide to SSI Maturity), then move to Part 3's How to Build or Upgrade an SSI, then read through the activities in Part 8.
- If you want to get right into the types of software security controls organizations are using in their SSIs, or if you are working on building out capabilities, then read Part 8 (The BSIMM Activities).
- If you want to see a summary of the BSIMM15 data, review Part 9.
- If you want to look at our analysis of the BSIMM data, review Part 9.



THE BSIMM SKELETON

The BSIMM skeleton provides a way to view activities at a glance, which is useful when thinking about your own SSI. The skeleton is shown in Figure 1, organized by domains and practices. A detailed version that includes activity references and level information is also available in Part 7. More complete descriptions of the activities and examples are available in Part 8 of this document.

Use this skeleton to understand the software security activities included in BSIMM15. A list of software security controls can be a very helpful guide here; the BSIMM project has worked since 2008 to ensure that its content matches real-world efforts.

| GOVERNANCE | | |
|--|--|---|
| STRATEGY & METRICS | COMPLIANCE & POLICY | TRAINING |
| <ul style="list-style-type: none"> Publish process and evolve as necessary. Educate executives on software security. Implement security checkpoints and associated governance. Publish data about software security internally and use it to drive change. Enforce security checkpoints and track exceptions. Create or grow a satellite (security champions). Require security sign-off prior to software release. Create evangelism role and perform internal marketing. Use a software asset tracking application with portfolio view. Make SSI efforts part of external marketing. Identify metrics and use them to drive resourcing. Integrate software-defined lifecycle governance. Integrate software supply chain risk management. | <ul style="list-style-type: none"> Unify regulatory pressures. Identify privacy obligations. Create policy. Build a PII inventory. Require security sign-off for compliance-related risk. Implement and track controls for compliance. Include software security SLAs in all vendor contracts. Ensure executive awareness of compliance and privacy obligations. Document a software compliance story. Ensure compatible vendor policies. Drive feedback from software lifecycle data back to policy. | <ul style="list-style-type: none"> Conduct software security awareness training. Deliver on-demand individual training. Include security resources in onboarding. Enhance satellite (security champions) through training and events. Create and use material specific to company history. Deliver role-specific advanced curriculum. Host software security events. Require an annual refresher. Provide expertise via open collaboration channels. Reward progression through curriculum. Provide training for vendors and outsourced workers. Identify new satellite members (security champions) through observation. |


| INTELLIGENCE | | |
|---|--|---|
| ATTACK MODELS | SECURITY FEATURES & DESIGN | STANDARDS & REQUIREMENTS |
| <ul style="list-style-type: none"> Use a data classification scheme for software inventory. Identify potential attackers. Gather and use attack intelligence. Build attack patterns and abuse cases tied to potential attackers. Collect and publish attack stories. Build an internal forum to discuss attacks. Have a research group that develops new attack methods. Monitor automated asset creation. Create and use automation to mimic attackers. Create technology-specific attack patterns. Maintain and use a top N possible attacks list. | <ul style="list-style-type: none"> Integrate and deliver security features. Application architecture teams engage with the SSG. Leverage secure-by-design components and services. Create capability to solve difficult design problems. Form a review board to approve and maintain secure design patterns. Require use of approved security features and frameworks. Find and publish secure design patterns from the organization. | <ul style="list-style-type: none"> Create security standards. Create a security portal. Translate compliance constraints to requirements. Identify open source. Create a standards review process. Create SLA boilerplate. Control open source risk. Communicate standards to vendors. Use secure coding standards. Create standards for technology stacks. |

| SSDL TOUCHPOINTS | | |
|---|---|--|
| ARCHITECTURE ANALYSIS | CODE REVIEW | SECURITY TESTING |
| <ul style="list-style-type: none"> • Perform security feature review. • Perform design review for high-risk applications. • Use a risk methodology to rank applications. • Perform architecture analysis using a defined process. • Standardize architectural descriptions. • Have SSG lead design review efforts. • Have engineering teams lead AA process. • Drive analysis results into standard design patterns. • Make the SSG available as an AA resource or mentor. | <ul style="list-style-type: none"> • Perform opportunistic code review. • Use automated code review tools. • Make code review mandatory for all projects. • Assign code review tool mentors. • Use custom rules with automated code review tools. • Use a top N bugs list (real data preferred). • Use centralized defect reporting to close the knowledge loop. • Build a capability to combine AST results. • Create capability to eradicate bugs. • Automate malicious code detection. • Enforce secure coding standards. | <ul style="list-style-type: none"> • Perform edge/boundary value condition testing during QA. • Drive tests with security requirements and security features. • Integrate opaque-box security tools into the QA process. • Drive QA tests with AST results. • Include security tests in QA automation. • Perform fuzz testing customized to application APIs. • Drive tests with design review results. • Leverage code coverage analysis. • Begin to build and apply adversarial security tests (abuse cases). • Implement event-driven security testing in automation. |

| DEPLOYMENT | | |
|---|---|---|
| PENETRATION TESTING | SOFTWARE ENVIRONMENT | CONFIGURATION MANAGEMENT & VULNERABILITY MANAGEMENT |
| <ul style="list-style-type: none"> • Use external penetration testers to find problems. • Feed results to the defect management and mitigation system. • Use penetration testing tools internally. • Penetration testers use all available information. • Schedule periodic penetration tests for application coverage. • Use external penetration testers to perform deep-dive analysis. • Customize penetration testing tools. | <ul style="list-style-type: none"> • Use application input monitoring. • Ensure host and network security basics are in place. • Implement cloud security controls. • Define secure deployment parameters and configurations. • Protect code integrity. • Use application containers to support security goals. • Use orchestration for containers and virtualized environments. • Use code protection. • Use application behavior monitoring and diagnostics. • Create bills of materials for deployed software. • Perform application composition analysis on code repositories. • Protect integrity of development toolchains. | <ul style="list-style-type: none"> • Create or interface with incident response. • Identify software defects found in operations monitoring and feed them back to engineering. • Track software defects found in operations through the fix process. • Have emergency response. • Develop an operations software inventory. • Fix all occurrences of software defects found in operations. • Enhance the SSDL to prevent software defects found in operations. • Simulate software crises. • Operate a bug bounty program. • Automate verification of operational infrastructure security. • Publish risk data for deployable artifacts. • Streamline incoming responsible vulnerability disclosure. • Do attack surface management for deployed applications. |

FIGURE 1. THE BSIMM SKELETON. Within the SSF, the 128 activities are organized into 12 BSIMM practices, which are within four domains.

PART 2: TRENDS AND INSIGHTS



PART 2: TRENDS AND INSIGHTS

BSIMM data originates in interviews conducted with member firms during a BSIMM assessment. Through these in-depth conversations, assessors look for the existence of BSIMM activities and assign credit for those that are performed with sufficient coverage across the organization, have the formality to be repeatable and consistent, and show the depth to be effective at managing associated risks. After each assessment, the observation data is added to the BSIMM data pool, where statistical analysis is performed to highlight trends in how firms secure their software.

You can use this information to understand what other organizations in your vertical are doing to then inform your own strategy.

This year has been defined by firms adjusting to the industry-wide shake-ups of the previous year as they continue to adjust to the challenges and opportunities presented by AI and large language models (LLMs) while building application security programs that are in compliance with government mandates. Companies behind DevSecOps platforms, cloud solutions, and security tooling are rising to the challenges from both the marketplace and attackers to make it easier for developers to automate security tooling and processes. This year marks the 15th iteration of the BSIMM15 framework, and we are reporting on long-term trends from the data pool.

THE EVOLUTION OF SHIFT EVERYWHERE

As automation and modern pipelines continue to handle more and more software development tasks, firms are integrating security governance and tests into their development platforms. While interest in shift left has only continued to grow over time, the BSIMM data shows firms adopting the shift everywhere philosophy. Organizations start to implement shift left to control the excessive costs, efforts, and risks associated with testing for security defects only just before promotion to production. This large friction point helped drive the development and adoption of SAST tooling, which drove the software industry to shift testing to the left in the SDLC, a place where vulnerabilities could be found and fixed faster and for less money. More recently, shift left was expanded to a broader testing philosophy where firms would also test designs and other development artifacts (e.g., golden masters, configurations, anything done as-code) as soon as they were ready, which marked the beginning of shift everywhere.

GOVERNANCE AND AUTOMATION

The core tenets of shift everywhere lie in taking advantage of automation to put data collection and decisions as close to the software development process as required. The activity [SM3.4] **integrate software-defined lifecycle governance** was introduced in BSIMM10 and has seen slow but steady progress, growing nearly 48% in the past year to reach 9.1% of the 121 firms. Additionally, because it is bad practice to mandate

WHAT IS SHIFT EVERYWHERE, REALLY?

To define shift everywhere, let's start by stating what it's not: it's not trying to do all the security things everywhere in all parts of the software lifecycle (SLC) all the time. Instead, shift everywhere is a philosophy; it's an approach to SLC governance that acknowledges the reality that consistently achieving acceptably secure software is a shared responsibility and that this responsibility traverses legal, audit, risk, governance, IT, cloud, technology, vendor management, and resilience, among others. Each stakeholder has their own business processes to execute, but each also needs to do their version of security sign-off, which requires understandable and usable telemetry from the SLC toolchain.

Not so long ago, the only view into the SLC afforded to stakeholders was, "We shipped it yesterday!" That was horrible then and is much worse now, mainly because automation generates telemetry that is easy to route to stakeholders. A shift everywhere approach starts by asking how these roles get the information they need, when they need it, in the processes they normally use, with little or no additional friction, then it bridges that gap, giving each role access to appropriate sensors, whenever they need it, from anywhere in the SLC. Shift everywhere is a philosophy about the security testing and sensors that generate information for all stakeholders in the company, it's not rooted in increasing the security spend or effort. Accordingly, shift everywhere is not adding more security for security's sake, instead, it's ensuring that every security stakeholder can knowledgeably make risk management decisions.

unsustainable security measures, the ease of automation has removed that barrier and allowed firms to [CR1.5] **make code review mandatory for all projects**, with an observation increase of about 72% since BSIMM10. Additionally, deciding when a test is required in the development is essential to ensuring that software is evaluated for risk at the most appropriate time, and observations of [ST3.6] **implement event-driven security testing in automation** grew by 43% from BSIMM14 to BSIMM15. Automating decisions and governance allows firms to manage risk in real time.

ENABLING PEOPLE

As firms further integrate security testing into the SDLC, developers are faced with an increasing duty to understand and respond to security defects as they're discovered and reported. One way that firms have found to reduce cognitive friction, or the mental effort spent, associated with this increased responsibility is to provide easy access to expertise via open collaboration channels. Observations of the activity [T2.12] **provide expertise via open collaboration channels** has shown strong growth over the last few years. Observation rates for this activity remained in the single digits from BSIMM-V through BSIMM12, where only ~5% of the firms did this activity, but starting with BSIMM13, observations began to rise fast, climbing to ~17% of firms in BSIMM13, to ~22% of firms in BSIMM14, to 31.4% of firms in BSIMM15. This represents more than 340% growth since BSIMM12. Since most firms are already taking advantage of collaboration tools such as Slack or Microsoft Teams, all it takes is to create a channel and then ask dedicated software security subject matter experts to answer questions as they arise. The long-term trend for this activity is one where firms found it was important 15 years ago as ~55% of firms did this in BSIMM1 and the activity remained popular until falling off in BSIMM-V. The availability of collaboration platforms combined with the need to provide access to security experts means that firms continue to shift everywhere through collaboration.

SOFTWARE SUPPLY CHAIN RISK MANAGEMENT

In recent years, the software supply chain has been put under the spotlight by regulatory requirements from the US government and European Union, and firms are continuing to expand their software supply chain risk management programs

to better keep up with regulatory challenges. Additionally, firms are continuing to mature their relationships with vendors as the entire industry is getting better at security.

SOFTWARE BILLS OF MATERIALS

When addressing risk in software, source code, or libraries sourced from the software supply chain, the first step is understanding which software components are present in the software and environment. Organizations are continuing to grow the capability to build SBOMs, with a 22% increase in [SE3.6] **create bills of materials for deployed software** and a 67% increase in observations of [SE3.8] **perform application composition analysis on code repositories**. Organizations continue to report issues with SBOMs, both in terms of the quality of the SBOMs that are created and how to make decisions with the SBOMs they receive from vendors.

VENDOR MANAGEMENT AND BESPOKE SOFTWARE

As we reported in BSIMM14, firms continue to change their relationships with vendors and expect vendors to be more mature in how they build secure software. Over the last several years, [CP2.4] **include software security SLAs in all vendor contracts**, [CP3.2] **ensure compatible vendor policies**, [SR 2.5] **create SLA boilerplate**, and [SR3.2] **communicate standards to vendors** have all seen growth since BSIMM9 and rapid growth since BSIMM11. [CP2.4] has grown from 36.9% of firms in BSIMM11 to 49.6% of firms in BSIMM15, a growth of over 34%. [CP3.2] has grown from 12.3% in BSIMM11 to 32.2% in BSIMM15, a 162% growth. [SR2.5] has gone from 33.8% in BSIMM11 to 51.2% in BSIMM15, a growth of over 51%. While [SR3.2] **communicate standards to vendors** saw a slight decrease between BSIMM13 and BSIMM14 as we reported last year, its growth has continued in BSIMM15, rising from 6.9% in BSIMM11 to 15.7% in BSIMM15, almost 127% growth. [T3.2] **provide training for vendors and outsourced workers**, while down since BSIMM11, managed to grow slightly from BSIMM14, where it was seen in 10.8% of firms to 11.6% of firms in BSIMM15. Firms are clearly expecting more from their vendors when it comes to the software these vendors are providing them.




FIGURE 2. VENDOR MANAGEMENT & BESPOKE SOFTWARE. Firms continue to increase their investment in managing their vendors and bespoke software development. Even [T3.2] increased in observation rates for BSIMM15 after a long period of decline.

REGULATORY PRESSURES

In BSIMM14, we added [SE3.9] **protect integrity of development toolchains** with the intention of tracking how many firms were meeting this particular regulatory requirement. In the past year, nearly half of the observations conducted under the BSIMM14 framework with this activity added have shown that firms are protecting the integrity of their toolchains by performing identity and access management and other security processes in and around the development environment. As a companion activity, BSIMM15 is introducing [SE3.10] **protect the integrity of development endpoints** to measure how firms are securing the workstations that access the various servers and services of the toolchain. Additionally, firms are protecting the code that they publish to better meet regulatory requirements, and observations of [SE2.4] **protect code integrity** have increased by ~20%, and observations of [SE3.2] **use code protection** have increased by ~45%. Finally, firms are feeling pressure to have a PSIRT functionality that can handle vulnerability reports and security bulletins, and there was a ~25% increase in observations of [CMVM2.4] **streamline incoming responsible vulnerability disclosure**. With the Cyber Resiliency Act moving through the EU regulatory process, we'll watch to see if mandated design review and security requirement-based activities increase in response.

WRESTLING WITH NEW TECHNOLOGIES

When we wrote about AI/ML in BSIMM14, we were watching for potential trends to measure and were waiting to see where this new technology would push the industry. Over the past

year, we have performed assessments on firms wrestling with the problems and opportunities that this new technology presented and have found that most firms are still struggling with defining and securing the new attack surface. To measure these efforts, we introduced a new activity in BSIMM15, [SR3.5] **create standards controlling and guiding the adoption of new technologies**. However, that new activity won't encompass the entirety of the AI/ML response, and we've found that firms have changed their approaches in how they find security issues and build solutions for this new technology.

FINDING PROBLEMS IN NEW TECHNOLOGIES

Before firms can build security solutions, they have to understand what problems they need to solve. Firms can better exercise new technologies by [AM2.8] **have a research group that develops new attack methods**, which saw a ~30% increase in BSIMM15. When firms seek to put their new attack methods into the SDLC, they can integrate those into test cases by [ST3.5] **begin to build and apply adversarial tests (abuse cases)**, observations of which doubled from BSIMM14 to BSIMM15. By finding new attack methods and putting those to the test, firms can better ensure that new technologies are resilient to novel attacks before the industry best practices catch up.

SECURITY SOLUTIONS FOR NEW SERVICES

When the industry hasn't caught up to new technologies and use cases, it's up to the SSGs to innovate and codify

internal practices. In addition to the new activity [SR3.5] **create standards controlling and guiding the adoption of new technologies**, firms can also enable developers with [SFD3.3] **find and publish secure design patterns from the organization** to keep from having to repeatedly reinvent the wheel, and observations of that activity grew by ~45% in BSIMM15. Additionally, as new architectures are understood, it's beneficial to codify security into the environment through [SE1.4] **define secure deployment parameters and configurations**, which saw a ~15% increase over BSIMM14. It's not enough to simply figure out solutions to secure new technologies—that innovation needs to be captured and codified by firms until the industry catches up.

LONG-TERM TRENDS

This BSIMM report marks the 15th iteration of the framework, and we have gone over the data for longer-term trends that are worth mentioning.

EXPANDING THE BSIMM'S REACH

One trend that has developed over the last few years is firms are increasingly moving beyond simple reassessments and are beginning to measure their software security program's penetration into their organizations by conducting BSIMMs for their business units. Since 2021, we have conducted 35 business unit BSIMM assessments. These assessments allow a firm to see how well its program has reached into these organizations and compare that to what the firm's BSIMM shows. When multiple business units are being assessed, a firm can compare the different organizations to each other, and many find that while some business units lag behind the main program, some business units go above and beyond. These organizations that go beyond can provide a foundation the firm can build on to expand its enterprise efforts. Business unit BSIMMs do not currently feed the overall BSIMM data pool that gets reported here annually, but that may change in the future as more business unit BSIMMs are conducted.

SOFTWARE SECURITY AWARENESS TRAINING'S LONG DECLINE

[T1.1] **Conduct software security awareness training** started off strong, with 100% of BSIMM1's nine firms doing this activity, but it dropped to 80% of the firms in BSIMM2 as we added

non-software security pioneers to the data pool—and that decline has continued ever since. In BSIMM15, its observation rate has fallen to an activity low of 51.2% of firms still doing basic software security awareness training. Part of this will always be driven by budgets as training takes a firm's staff away from the work that generates revenue, and it will always impact delivery schedules if engineers must step away from writing and testing code. Another factor is likely the "been there, done that" view of software security training. It is easy to take the view that "we did training, now we are good, we have other priorities," but this is a dangerous trap. Staffing changes all the time, and if there are no consistent efforts made, there is no guarantee engineers have ever received the needed training. On the positive side, anecdotal evidence suggests that firms that identified this gap in recent BSIMMs have taken steps to revamp their training programs and are scoring better upon their most recent reassessment.


TOPICS WE'RE WATCHING

This year continued to see changes in how firms automated their SDLCs, integrated new technologies into their software, and met regulatory challenges. The demands of cloud, toolchains, tools, application security adjacencies, AI, and government scrutiny are leading to a vastly increased program scope, which is in turn necessitating a new era of shared responsibility between SSGs and engineering.

Participant feedback indicates that the following might influence their future efforts:

- The continuing impact of regulations on the software industry as governments mandate security as an essential element of software products that run in critical infrastructure.
- How things change as the industry continues to mature and develop security solutions for and using AI, enabling software to move into new use cases.

PART 3: FOR PRACTITIONERS



PART 3: FOR PRACTITIONERS

INTRODUCTION

In Part 3, we take a look at industry trends that are important for software security programs as well as review how to build or upgrade a software security program.

REGULATIONS FOR PRACTITIONERS

THE REGULATORY LANDSCAPE

Guest Author: Tim Mackey

The regulatory landscape that modern companies face is challenging. Independent of industry, geography, and market size, all businesses face customer expectations around the protection of customer data, and recent legislative and regulatory actions, independent of jurisdiction, expand those expectations to include not only data protection and product safety but to now encompass how software powered products and services are developed.

By way of example, the following regulatory events have direct impact on SSGs in their respective markets:

- Section 3305 in the US Consolidated Appropriations Act, 2023, which added Section 524B to the Food, Drug and Cosmetics Act governing the US Food and Drug Administration, saw Section 524B become enforceable in October 2023. Section 524B obligates SBOMs and vulnerability disclosure programs details when a medical device manufacturer files for FDA authorization or approval.
- The US Securities and Exchange Commission (SEC) adopted regulations in December 2023 covering Cybersecurity Risk Management, Strategy, Governance, and Incident Disclosure by Public Companies. Those rules extend to include cybersecurity events impacting products or services sold by the covered company.
- The CISA Secure Software Development Attestation Form (OMB 1670-0052) was published in March 2024 and applies to software suppliers to the US government from September 2024. This form is one compliance element of President Biden's Executive Order on Cybersecurity (EO 14028) and directly aligns with practices in the NIST Secure Software Development Framework (SSDF).
- The European Union Cybersecurity Resilience Act, enacted in October 2024, stipulates that software subject to this EU regulation can be created anywhere but must be marketed in the EU.
- The EU Network and Information Security (NIS2) Directive primarily focuses on preparedness at the member state level but also seeks to promote security cultures within critical sectors. It is nominally effective as of October 2024, but there have been delays in its implementation at the member state level.
- The European Union Digital Operational Resilience Act (DORA) applies to financial institutions, brokerages, and securities exchanges and will be applicable in January 2026.

While each of these efforts has its own goals and technical requirements, one common theme is a need for meaningful and consistent processes and success metrics within development and SSG teams. Threat-informed risk management and the need for software producers to implement processes such as coordinated vulnerability disclosures form part of that benchmark and are arguably the first step in a regulatory-focused cybersecurity maturity model. We see such efforts reflected in BSIMM15 with the growth in [CMVM2.4] **streamline incoming responsible vulnerability disclosure** and [CMVM3.3] **simulate software crises**, and while the observed occurrence of [CMVM3.6] **publish risk data for deployable artifacts** is low, as regulatory disclosure requirements increase, we expect to see [CMVM3.6] become commonplace.

Of course, meeting regulatory requirements isn't something that happens in isolation. If your SSG security champions aren't aligned with your legal or governance team, then your organization runs the risk of investing in proposed requirements that might change before the regulation is finalized. Worse still, given that the penalties for non-compliance extend from fines to prison time, it's crucial that any required process changes for compliance align with the scope of each regulation and are prioritized based on business risk.

SSDF-BSIMM MAPPING UPDATED FOR BSIMMI5

The NIST SP 800-218 v1.1 Secure Software Development Framework (SSDF) standard was released in February 2022 and specifies 42 different tasks that organizations should do as part of a secure software development lifecycle (SSDL), as well as things to do to secure the development environment. The SSDF draws from a wide variety of existing frameworks as both a basis for itself and as references organizations can use. One of the most referenced frameworks is the BSIMM, and accordingly, the SSDF refers to it and Executive Order 14028, more than any other framework (see Figure 3).




FIGURE 3. SSDF REFERENCES TO OTHER FRAMEWORKS.

The SSDF’s references to BSIMM activities can help readers understand the intent behind the tasks in the standard and provide examples of activities that help with those tasks. Unfortunately, the NIST standard refers to BSIMM12, and between BSIMM12 and BSIMM15, many BSIMM activity numbers have changed as activities became more or less common. For example, the SSDF references SE2.6 from BSIMM12, but after the NIST standard was published, SE2.6 became more common and is now known as SE1.3. Looking for SE2.6 will not work unless you know activities have changed.

Since the SSDF standard was released, 13 activities have moved between levels, but because some are referenced by more than one SSDF task, a total of 16 tasks have been affected. Table 3 shows the updated BSIMM mappings for these 16 tasks.

| TASK | BSIMMI2 (ORIGINAL) | CHANGE | BSIMMI5 |
|--------|--------------------|---------|---------|
| PO.1.1 | SE2.6 | Becomes | SE1.3 |
| PO.1.2 | SM2.2 | Becomes | SM1.7 |
| PO.2.2 | T3.4 | Becomes | T2.11 |
| PO.4.1 | SM2.2 | Becomes | SM1.7 |
| PO.4.2 | SM2.2 | Becomes | SM1.7 |
| PW.1.1 | AM2.2 | Becomes | AM3.4 |
| | AM2.5 | Becomes | AM3.5 |
| | AA1.3 | Becomes | AA2.4 |
| PW.2.1 | AA1.3 | Becomes | AA2.4 |
| PW.4.1 | SR2.4 | Becomes | SR1.5 |
| | SR3.1 | Becomes | SR2.7 |

| TASK | BSIMMI2 (ORIGINAL) | CHANGE | BSIMMI5 |
|--------|--------------------|---------|---------|
| PW.4.4 | SR2.4 | Becomes | SR1.5 |
| | SR3.1 | Becomes | SR2.7 |
| PW.7.2 | CR1.6 | Becomes | CR2.8 |
| PW.9.1 | SE2.2 | Becomes | SE1.4 |
| PW.9.2 | SE2.2 | Becomes | SE1.4 |
| RV.1.1 | CMVM2.1 | Becomes | CMVM1.4 |
| | CMVM3.7 | Becomes | CMVM2.4 |
| RV.1.3 | CMVM2.1 | Becomes | CMVM1.4 |
| | CMVM3.7 | Becomes | CMVM2.4 |
| RV.2.1 | CMVM2.2 | Becomes | CMVM1.3 |
| RV.2.2 | CMVM2.1 | Becomes | CMVM1.4 |

TABLE 3. SSDF TASKS AFFECTED BY BSIMM MOVES.

BSIMM14 and BSIMM15 added activities that could be mapped to NIST standards but have not officially been mapped by NIST. Table 4 shows these unofficial mappings.

| TASK | BSIMM12 (ORIGINAL) | CHANGE | BSIMM15 |
|--------|--------------------|-----------|---------|
| PO.1.1 | None | Could Add | SE3.9 |
| PO.3.1 | None | Could Add | SE3.9 |
| PO.3.2 | None | Could Add | SE3.9 |
| PO.3.3 | None | Could Add | SE3.9 |
| PO.5.1 | None | Could Add | SE3.10 |
| PO5.2 | None | Could Add | SE3.10 |

TABLE 4. NEW BSIMM ACTIVITIES THAT COULD BE MAPPED TO THE SSDF STANDARD.

This update, which shows how the SSDF’s reference to BSIMM12 activities can be mapped to BSIMM15 activity numbers, should help people continue to use the BSIMM to understand what NIST intends with SSDF tasks.

AI/ML

Application security, like any area of technology, has to keep pace with rapid developments that seemingly upend everything once or twice a decade. The AI/machine learning (AI/ML) and LLM releases of the past few years are both exciting and terrifying for security professionals: exciting because new LLMs can perform many security tasks that were exclusively human, and terrifying because they can also perform many hacking tasks that were exclusively human. On a technical

The industry is changing, and the BSIMM is changing with it. To keep pace, we have added a new activity [SM3.5] that will measure an SSI’s formal activities to proactively evaluate new architectures, technologies, and processes and provide security practices when industry best practices haven’t been canonized yet. Additionally, we’ve highlighted the top five most impactful BSIMM activities that can help firms lean into the problems that AI/ML and LLM adoption and integration can cause. See Part 8 for more.

level, LLMs represent a new attack surface with a new class of vulnerabilities and security requirements that the industry is still figuring out in real time. As if that weren’t enough, LLMs seem to completely reinvent themselves every month by not only adding features and improving functionality but also by moving into completely new areas: first text and chat, then image generation, and at the time of writing, desktop productivity tasks.

To stay ahead of the innovation curve, developers need AppSec programs to provide solutions, guidance, and best practices that help them safely adopt new technology. In return, AppSec needs to rely on subject matter experts to both help them understand how these new technologies will be used as well as to get buy-in from development to stay flexible as the threat landscape evolves. This collaborative effort will often involve members from legal for support around regulatory updates, operations and IT for the software environment, and communications teams to keep everyone up to date.

AI/ML EXPERT INSIGHTS: THE UNCERTAINTY PROBLEM

Guest Author: David Benas

When dealing with an area as large as LLMs, building an understanding of the problem is important. When we talk to clients about what they’re trying to do and the problems they have doing it, we see a wide variety of pain points, but in general, the problem that everybody is struggling with is uncertainty. There isn’t a lot of well-understood guidance out there, so they’re having to find the answers themselves. Here are five areas we talk about a lot:

- What does security even mean in the context of AI/ML? When software blurs the lines between human behavior and rule-driven algorithms, what it means to be secure is different from securing a traditional web server or operating system. Writing policy is difficult because LLMs make security about behavior and secrets rather than business logic and vulnerabilities. Don’t forget about the traditional vulnerabilities, though, because it is still software, and old vulnerabilities can still happen.
- How can tools be used to effectively automate their way to meet security needs? The promise of LLMs is automating

While recent LLM developments are reshaping the software industry, firms have been using AI/ML models for years, leading to new terms around them. For example, let’s say an investment firm is building an in-house AI/ML model that consumes headlines, releases, and other business data to make stock investment predictions. These decision-helper models will be known as predictive models because they’re trained to predict results and outcomes based on inputs. LLMs consume prompts and then generate text based on those prompts, so they are called generative models.

- tasks, and security needs to be automated to keep up. Many of the security people we talk to struggle with trusting this new technology in a security role when it’s unknown how secure these tools are to begin with. They’re also concerned about the value of running automated scans against LLMs because of the risk of false negative results and existing rules.
- How can LLMs be used to improve the security posture? With the exception of writing some automated test case and pen testing scripts, most AppSec teams rely on off-the-shelf tooling for security, but adapting LLMs to security is actually a development effort, so AppSec professionals

are being forced to put on their developer hats or to enlist engineering's help in putting LLMs to work for security. Some tools incorporate AI/ML features, but as of today, most of these additions aren't yet revolutionary.

- Because AI/ML and LLMs can be used to perform tasks that may result in discriminatory biases, how can teams make sure that these implementations don't run afoul of government-mandated anti-discrimination laws? You can't sit an AI/ML model down in a sensitivity training classroom and expect it to correct its behavior like a human would, so teams struggle with making sure that their software doesn't propagate biases that it may have learned from its training data or learned on its own during the training process.
- The biggest challenge facing security teams and researchers at the moment is that teams can't reach into a model's engine and tinker with it to solve problems or improve performance, so researchers are actively exploring how to build an explainable AI/ML model that can have its rationale for decisions explained and understood. Until that happens, faulty decision-making rationales can result in unfixable security issues.

After talking about the problems they're facing, companies want to know if the effort they're putting in to update their security programs makes sense. They struggle with how quickly and radically things can change, and after having the rug pulled out from them a few times, teams are wary of investing heavily in a course of action that will be outdated after the next press release. The following list highlights just some of the common challenges faced by adopters of AI/ML:

- Models change from month-to-month, and industry-provided capabilities change from year-to-year, which means next year's must-have use case isn't even hinted at today. When building best practices, it's vital to know what those best practices will be used for, and that's not possible with how fast things move.

- A current limitation of AI/ML models is that functionality is constrained to using pre-trained models because going out and getting training data is costly, time-consuming, and often not legal. There's the ML 10x rule, for example, which estimates that for any given model's complexity, developers need 10x the training samples to train it properly. This also means that developers and security personnel are limited to what the established players have on offer.
- Another limitation lies in the disconnect between traditional security practices and new use cases. AI/ML software is often just bundles of Python software and pre-built training data that developers can source and use. Unfortunately, scanning these open source packages with software composition analysis, scanning Python with static application security testing (SAST) tooling, and running API pen tests fail to provide adequate coverage of new security issues like prompt injection.
- How do you even discover problems, then? Security testers are struggling with knowing even what to look at because the internal workings of models are incomprehensible to humans. The industry hasn't fully grasped how to use SBOMs to share risk data and make decisions, so what would an AI bill of materials (AIBOM) look like or do?
- Even if a vulnerability or security issue is discovered, developers can't go into the model and tune its responses manually, so security defects just get added to the backlog because they're unfixable. Until there are meaningful solutions to security problems, the defect backlog just grows and grows.

The majority of people we talk to are optimistic, though. At the rate that AI/ML and LLMs are maturing, they expect many of these questions will have good answers before too long.

AI/ML AND BSIMM15

The most common question we get asked is, "How is BSIMM15 going to change for AI/ML?" Well, here's the answer: we're updating and highlighting existing activities that have an impact on AI/ML security, and we've added, in our opinion, a long overdue activity around proactively planning to mitigate the impacts of new technologies on security.

NEW ACTIVITY: SR3.5 – CREATE STANDARDS CONTROLLING AND GUIDING THE ADOPTION OF NEW TECHNOLOGIES

While AI/ML may seem new and intimidating, SSGs have been faced with integrating security best practices into anything and everything new that their developers could think up for years. Since the first version of the BSIMM was published, SSGs have had to secure mobile applications, cloud environments, containers, new languages, and new frameworks, all while developers have shifted from waterfall and spiral development lifecycles to Agile lifecycles and DevSecOps cultures. Some firms have proactively created working groups that examine these new technologies and processes to understand potential security needs while the industry is still reeling from being turned upside down once again.

This activity is beneficial to firms looking to take advantage of innovations that are on the cutting edge of technology. It's not enough to have one smart individual who researches interesting new topics and then tells people about it, this is an activity for firms that rigorously examine and workshop the implications of adopting a new technology and then create actionable guidance for developers, IT, and operations to follow. We'll be looking for new security requirements and controls, updates to policy, tooling customizations, and other proactive methods that firms use to fill the gap left by the absence of industry best practices for brand-new technologies.

FIVE KEY ACTIVITIES: USING THE BSIMM TO DEAL WITH AI/ML

In addition to the debut of SR3.5, plenty of existing BSIMM activities can help teams address the AI/ML problem space. Whether the priority is keeping on the right side of emerging regulations, defending against emerging attacks, or enabling developers who want to take advantage of coding copilots, BSIMM15's five key activities can help:

- [AM1.5] **Gather and use attack intelligence.** While most firms may rely on automated tooling and skilled pen testers to thoroughly exercise an application's potential security vulnerabilities, that's not an option for technologies as new as LLMs. Instead, researchers are hard at work finding new ways to attack and defeat AI/ML safeguards or discover new vulnerabilities and exploits that don't exist in today's software types. It's up to individual organizations to have a designated threat intelligence function that collects publicly available research and helps everyone learn about these emerging vulnerability exploits.
- [AM3.4] **Create technology-specific attack patterns.** Once new exploits are collected, organizations should build a catalog of attacks specific to the use cases and types of AI/ML or LLM that developers are integrating or taking advantage of. These attack patterns should cover everything from novel attacks like prompt injection to traditional vulnerabilities that could apply to the API interface, cloud host, or other technologies that make up the complete implementation to uncover potential interactions, e.g., potentially using prompt injection to get an LLM to output malicious payloads that could compromise integrated databases or parsers.
- [SFD3.1] **Form a board to approve and maintain secure design patterns.** Once attack patterns are built for the organization's intended AI/ML or LLM uses, firms should seek to combine solutions and mitigations to the attacks in approved design patterns. These patterns could cover all parts of a scratch-built AI/ML-enabled application or provide security guidance when integrating commercially available AI/ML components.
- [SM3.5] **Integrate software supply chain risk management.** It's vital to remember that despite how new LLMs feel, they are still hosted in cloud environments or released as open source bundles of Python and training data. Firms should expand software supply chain risk management to account for attacks that could compromise third-party AI/ML services by fully documenting and understanding the shared responsibility model for cloud-hosted services, mandating privacy and security service-level agreements (SLAs) in vendor contracts, and expanding the data classification guide to cover types of data that can and can't be shared with third-party AI/ML solutions. Additionally, companies will want to build requirements around training data to prevent potential legal liability due to the inclusion of copyrighted data or that prevent the "Garbage In, Garbage Out" problem that can compromise the AI/ML model's integrity by sourcing problematic or false training data.
- [CR1.5] **Make code review mandatory for all projects.** For firms that aren't looking to build or maintain AI/ML applications or LLMs but still want to take advantage of the efficiency gains of this new technology, the first avenue for many of them is enabling developers to take advantage of AI/ML copilots that can speed up source code creation and updates. These copilots aren't infallible, and the source code they produce is likely to not be perfectly secure, so to address the potential for vulnerabilities included in source code generated by AI/ML copilots, we would advise checking the source code for those vulnerabilities. Making automated SAST scans mandatory will go a long way to making AI/ML-enabled development more secure.

WHAT'S NEXT?

As we conduct BSIMM assessments under the BSIMM15 framework, we'll continue to measure how firms are securing this exciting new technology. Look for an update in the next annual BSIMM report as firms continue to solve AI/ML security problems and industry best practices catch up.

HOW TO BUILD OR UPGRADE AN SSI

Putting someone in charge is just a first step in building an SSI, there will be iterations of planning, growth, measurement, and bridge-building. You can use the processes below to guide your SSI's growth from newly emerging through dependable maturity.

The BSIMM is not just a long-term software security study or a single-purpose SSI benchmarking tool—it also eases management and evolution for anyone in charge of software security, whether that person is currently in a central governance-focused position or in a more local engineering-focused team. Firms of all maturity levels, sizes, and verticals use the BSIMM as a reference guide when building new SSIs or evolving their initiatives through various maturity and stakeholder ownership phases over time.

We often refer to SSIs we've seen as being in one of three broad states—emerging, maturing, and enabling—which we describe as follows:

- **Emerging.** An emerging SSI has defined its initial strategy, chosen foundational activities (e.g., those observed most frequently in the data pool), acquired some resources, and created a general roadmap for the next 18 months. SSI leaders are likely resource constrained on both people and budget, so the SSG is usually small and uses compliance requirements or other executive mandates to drive participation and to continue adding activities. These leaders require strong, visible, and ongoing executive support to manage frictions with key stakeholders who are resistant to adopting foundational process discipline.
- **Maturing.** A maturing SSI has an in-place team, defined processes for interacting with software security stakeholders, and a documented software security approach that is clearly connected to executive expectations for both managing software security risk and progressing along a roadmap to scale security capabilities. A maturing SSI is learning from its existing efforts, likely making consistent, incremental improvements in the SSDL and key security integrations. Example improvements include:
 - Reducing friction across business and development stakeholders
 - Protecting people's productivity gains through automation investments
 - Building bridges to other parts of the firm through evangelism, defect discovery, software supply chain protection, and incident response
 - Undergoing a shift everywhere transformation to efficiently test software artifacts as soon as appropriate
 - Adjusting the security strategy to keep pace with changes in risk and risk management processes
 - Finding solutions to systemic problems and making them broadly available as reusable, pre-approved IP
 - Responding quickly when attacks or other circumstances uncover a lack of resiliency

- **Enabling.** An enabling SSI ensures that all stakeholders can meet their objectives without putting the organization at unacceptable risk. The following are important principles for an enabling SSI:
 - There is continuous evangelizing about the best way for all stakeholders to meet security expectations, ensuring that the path of least resistance for development and deployment is also the most secure path, along with investing to proactively overcome various people, process, technology, and cultural growing pains.
 - The evolutionary needs of the SSI are harmonized with the goals of business initiatives, such as digital transformation, open source use, and cloud adoption.
 - A mature and integrated response to process and technical risk invokes an innovation engine to make reasonably future-proof solutions.
 - The use of culturally engrained approaches to automation, blameless review of failures, and protection of critical resources—people, for example—allow more time to tackle security innovation.
 - A platform engineering perspective removes security activity silos and ensures that all telemetry and benefits are available to all stakeholders everywhere.

It's compelling to imagine that organizations could reach a state of emerging, maturing, or enabling simply by applying a certain number or mix of activities to specific percentages of the staff and software portfolio, but that doesn't happen. Experience shows that SSIs usually reach an emerging stage by organizing all the ad hoc software security efforts they're already doing into one program. SSIs usually proceed to the maturing stage by focusing on the activities that are right for them without regard for the total activity count. This is especially true when considering the complexity of scaling some activities across 100, 1,000, or 10,000+ applications or people.

Organizations rarely move their entire SSI from emerging to enabling all at once. We have seen SSIs form, break up, and re-form over time, so an SSI might shift between emerging, maturing, and enabling a few times over the years. In addition, capabilities within an SSI (e.g., supply chain security or training) likely won't progress through the same states at the same rate. We've noted cases where one capability—vendor management, for example—might be emerging, while the defect management capability is maturing, and the defect discovery capability is in the enabling stage. There is also constant change in tools, skill levels, external expectations, attackers, attacks, resources, culture, and everything else. You can use the BSIMM15 participants scorecard (see Figure 18 in Part 7) to see the frequency with which BSIMM activities are observed across all participants, but use your own metrics to determine if you're making the progress that's right for you.

CONSTRUCTION LESSONS FROM OUR PARTICIPANTS

The purpose of the BSIMM is to measure SSIs. While the BSIMM doesn't directly measure SSI architecture, evolution, or motivations, our experience with more than 283 organizations since 2008 has highlighted cultural differences in SSI implementations.

No SSI is built in a vacuum. Whether your SSI is just emerging or has some capabilities in the maturing stage, knowledge from both the struggles and successes of other organizations can save you time and disruption. As software security becomes an important goal for any organization, multiple internal groups might each be taking their own approach to their goals. Understanding and harmonizing these cultural and technological views into a single SSI is important to long-term success.

CULTURES

Whether implicitly or explicitly, organizations choose the path for their software security journey by tailoring goals, methods, tools, resources, and approaches according to their individual cultures. There have always been two distinct cultures in the BSIMM participants:

- Organizations where the SSG was started by executives in a central corporate group (e.g., under a CISO) as a full-time role and chartered with software security governance, including compliance, testing, remediation monitoring, and risk management. This SSG stayed in the corporate organization chart, had the power to enact organization-wide policy, and expanded its efforts outward through, for example, tooling and security champions. This path was seen most often in regulated industries such as banking, insurance, FinTech, and healthcare but was also seen in ISV and technology firms.
- Organizations where the SSG was started by engineering technical leadership (e.g., senior application architects) as a part-time role and focused on technical software security efforts, such as configuration hardening, technology stack standards, secure coding standards, and security tool integration, which was often done for a single toolchain or project. As evangelism efforts convinced other development projects to use the same technical controls, the technical leadership usually worked with a CTO, VP Engineering, or other technology executive to establish a centralized security function within the engineering domain. The centralized function—often still part time—then used its influence to establish its own type of governance, which was often peer pressure to set some development process, create and manage security standards, and ensure that the silos of engineering, testing, and operations were aware of and adhered to general security expectations. This path was most often seen in technology, cloud, and ISV firms but was also seen in other verticals.

Regardless of its origin point, each culture usually arrived at an SSI driven by a centralized, dedicated SSG whose function is to ensure that appropriate software security activities are happening across the portfolio, so nearly all SSIs that are more

Whether your SSI is just emerging or has some capabilities in the maturing stage, knowledge from both the struggles and successes of other organizations can save you time and disruption.

than a couple of years old are driven top-down by governance objectives, even those started by engineering for engineering. Evangelism, peer pressure, and local implementations go only so far in formally implementing software security risk management as a culture.

Today, as you start or plan a major revamp of your SSI, just get started. You can start in corporate, or you can start in engineering. You can start with governance as a top priority, or you can focus on some technical controls. In any case, history seems to show that SSIs gravitate toward a focus on policy along with process that ensures adherence. Yours likely will as well.

A NEW WAVE IN ENGINEERING CULTURE

Over the past few years, we're seeing a new wave of software security efforts emerging from engineering teams. These teams are usually responsible for delivering a product or value stream (such as is common within ISVs) or maintaining a technology domain (such as the "cloud security group" or a part of some digital transformation group). Some organizations refer to these collective security efforts as site reliability engineering, DevSecOps, or GitOps security, but some have no specific name for it at all.

At least three factors drive these new efforts:

- The confluence of process friction, unpredictable impacts on delivery schedules, adversarial internal relationships, and a growing number of human-intensive processes from existing SSIs; top-down governance doesn't fit culturally or technologically with new engineering processes.
- The demands and pressures from modern software delivery practices, be they cultural such as Agile and DevOps, or technology-based such as cloud- and orchestration-based; gates and checkpoints built for maximum assurance often cause unacceptable disruption in processes built for speed.
- The shift to engineer self-service, typically seen as self-service IT (cloud), configuration and deployment (DevOps), and development (open source use and continuous integration); the ability to instantiate infrastructure and pipelines is also the ability to integrate your own security tools and configurations.

This new software security effort is frequently happening independently from the lessons learned that an experienced SSG might provide. In addition, this effort is driving many application lifecycle processes ever faster, regardless of whether the organization is ready to do software security risk management at that speed.

The governance-oriented approaches we've seen for years, along with this new wave of engineering-oriented efforts, are increasingly coexisting within the same organization. In addition, they often have competing objectives, which is pulling traditional governance-driven programs into modern and evolving hybrids. Figure 4 shows this ongoing SSG evolution.

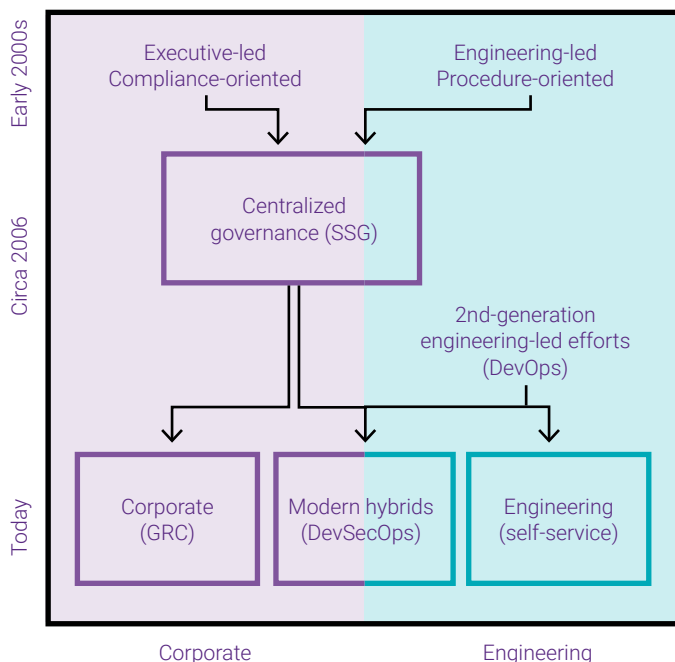


FIGURE 4. SSG EVOLUTION. These groups might have started in corporate or in engineering, but in general, they settled on enforcing compliance with tools. The new wave of engineering efforts is shifting where SSGs live, what they focus on, who is accountable for what, and how stakeholders work together.

The important lesson here is that this is likely happening in your organization as well—perhaps narrowly in a few development teams or perhaps broadly as a cultural shift across all of engineering. Taking an SSI to the maturing stage—and possibly to enabling as well—requires acknowledging this engineering effort and building bridges between all stakeholders who have ownership of the different aspects of software security. It also requires acknowledging that these different stakeholders have different business objectives and different views of risk, risk management, and risk tolerance relative to those objectives. Ensuring that everyone can meet their objective while also keeping the organization safe is a major goal for every SSI.

UNDERSTANDING MORE ABOUT DEVOPS

The DevOps movement has highlighted the tensions between established SSIs and engineering efforts that address software security their own way with their own processes. Given different objectives, we find that the outcomes desired by

these two approaches usually differ—rather than the top-down, compliance-driven style of governance-minded teams, these newer engineering-minded teams are more likely to prototype good ideas for securing software, which results in the creation of even more code and infrastructure on the critical path to delivery (e.g., security features, homespun vulnerability discovery, security guardrails).

Here, security is just another aspect of quality, and availability is just another aspect of resilience. To keep pace with both software development process changes (e.g., CI/CD adoption) and technology architecture changes (e.g., cloud, container, and orchestration adoption), engineering efforts are independently evolving both how they apply software security activities and, in some cases, what activities they apply. The changes these engineering teams are making include downloading and integrating their own security tools, spinning up self-service cloud infrastructure and virtual assets as they need them, following policy on the use of open source software (OSS) in applications but routinely downloading many other open source packages to build and manage software and processes, etc. Engineering efforts and their associated fast-paced evolutionary changes are putting governance-driven SSIs in a race to retroactively document, communicate, and even automate the knowledge they hold so that it can be useful to everyone.

Cloud service providers, software pipeline and orchestration platforms, and even QA tools have also begun adding their view of software security in their feature sets. For example, organizations are seeing platforms like GitHub, Azure DevOps, and GitLab competing by using security as a differentiator. Evolving vendor-provided features might be signaling to both the marketplace and adopting organizations that vendors believe security must be included in developer tools and that engineering security initiatives should feel comfortable relying on these external platforms as the basis of their security telemetry and even their governance workflows.

Again, the important lesson is that this is likely happening in your organization as well. Your path to an emerging or mature SSI must account for this federation of software security responsibilities and use of external providers, yet also enable every stakeholder to meet their business and security objectives.

CONVERGENCE AS A GOAL

We frequently observe governance-oriented SSIs planning centrally, seeking to proactively define an ideal risk posture during their emerging or early maturity phases. Initial uptake of the provided controls (e.g., security testing) is usually by those teams that have experienced real security issues and are looking for help, while other teams might take a wait-and-see approach.

We also observe that engineering efforts prototype controls incrementally, building on existing tools and techniques that already drive software delivery. Gains happen quickly in these emerging efforts, perhaps given the steady influx of new tools and techniques introduced by engineering but also helped along by the fact that each team is usually working in a homogenous culture on a single application and technology stack. Even so, these groups sometimes struggle to institutionalize durable

gains, usually because the engineers have not yet been able to turn capability into either secure-by-default functionality or automation-friendly assurance—at least not beyond the most frequently encountered security issues and beyond their own spheres of influence.

Engineering groups tend to view security as an enabler of software features and code quality. These groups recognize the need for having security standards but tend to prefer incremental steps toward governance-as-code as opposed to a large-manual-steps-with-human-review approach to enforcement. This tends to result in engineers building security features and frameworks into architectures, automating defect discovery techniques within a software delivery pipeline, and treating security defects like any other defect. Traditional human-driven security decisions are modeled into a software-defined workflow as opposed to being written into a document and implemented in a separate risk workflow handled outside of engineering. In this type of culture, it's not that the traditional SDLC gates and risk decisions go away, it's that they get implemented differently and usually have different goals compared to those of the governance groups. SSGs, and likely security champions groups as well, that begin to support this approach will speed up both convergence of various efforts and alignment with corporate risk management goals.


To summarize the lessons from our participants, scaling an emerging SSI across a software portfolio is hard for everyone, and stakeholders need to understand the lessons above before investing heavily in the journey from emerging to maturing.

Today's evolving cultural and technological environments require a concerted effort at converging governance and engineering objectives to create a cohesive SSI that ensures the software portfolio is appropriately protected.

FOR AN EMERGING SSI: SDLC TO SSDL

It's unlikely that any organization is doing nothing about software security. Even an organization without a formal initiative or a defined owner likely has some software security policy, AST, and/or processes for working with stakeholders. Provided below are actionable steps for consolidating an ad hoc effort into an emerging SSI. Keep in mind that most SSIs are multiyear efforts with real budgets, mandates, and ownership behind them, though. In addition, while all initiatives look different and are tailored to fit a particular organization, most initiatives share common core activities (see Table 10 in Part 7).

Figure 5 organizes the steps and suggested timeline to establish an emerging SSI, along with the associated BSIMM activities. It also includes a notional level of effort anticipated across people and budget, as well as estimated duration, all on a 1 to 3 scale. The effort and cost to reach each of these goals will vary across companies, of course, as they are primarily affected by risk objectives, organizational structure, and portfolio size. For example, deploying onsite static analysis across 10 applications using a common pipeline in one business unit will likely have a lower level of effort than deploying that same static analysis across 10 applications built in 10 toolchains in 10 business units.



The arrow of time (x-axis) is a notional order of efforts. Although this diagram appears to depict a waterfall process, many of these efforts will be happening at the same time and some will be repeated multiple times.

FIGURE 5. GETTING STARTED ROADMAP WITH NOTIONAL EFFORTS. This roadmap is supplemented with relative effort levels so that organizations can plan the resources needed for their emerging SSI.

Note that the getting started roadmap shown in Figure 5 includes some activities that have a high impact for emerging SSIs, even though they appear to be rarely observed in the BSIMM data pool. This happens because newly added BSIMM activities start with an observation rate of zero (e.g., [ST3.6] added for BSIMM11). These are foundational activities, even if organizations are just starting to add them to their journeys. Importantly, the steps described here are not specific to where in the organization the SSG is created—the SSG can be centralized in a governance group or an engineering group, or it can be federated across both. Regardless, governance and engineering functions will have to cooperate to ensure the achievement of organizational software security goals.

Note that an SSG leader with a young initiative (e.g., less than 18 months) working on foundations should not expect or set out to quickly implement too many BSIMM activities. Firms can absorb only a limited amount of technology, hiring, cultural, and process change at any given time. The BSIMM15 data shows that SSIs having an age of 18 months or less at the time of assessment (23 of 121 firms) have an average score of about 33, showing they are investing in software security but are not rushing ahead too rapidly.

Following are some details on the steps shown in Figure 5. The included activity references are meant to help the reader understand the associations between the topic being discussed and one or more BSIMM activities. Note that the references don't mean the topic being discussed is fully equivalent to the activity, for example, when we say, "...initial inventory [AM1.2]" (i.e., **use a data classification scheme for software inventory**), we don't mean that having the initial inventory encompasses the totality of [AM1.2], just that having it will likely be something you'll do on your way to implementing [AM1.2]. To continue using [AM1.2] as an example, most organizations will not set about implementing this activity and get it all done all at once. Instead, an organization will likely create an initial classification scheme and inventory, implement a process to keep the inventory up to date, and then decide how to create a view that's meaningful for stakeholders. Every activity has its own nuances and components, and every organizational evolution path for its emerging SSI will be unique.

CREATE A SOFTWARE SECURITY GROUP

The most important first step for all SSIs is to have a dedicated SSG that can get resources and drive organizational change, even if it's a group of one person coordinating organizational efforts. The SSG must understand which software security goals are important to the business and establish policy and process to drive everyone in that direction. At a minimum, the SSG should identify the risk management, compliance, and contractual requirements that the organization must adhere to [CP1.1]. Using awareness training [T1.1] to then help ensure that everyone understands their security responsibility is a common approach.

The SSG must work with engineering teams to establish a common understanding of the approach to software security, which might be, for example, to set up automated defect discovery, address security questions from developers with reusable security features [SFD1.1], and act as an advisor for design decisions [SFD1.2].

DOCUMENT AND SOCIALIZE THE SSDL

Next, you'll want to publish security policies and standards through established governance, risk, and compliance (GRC) channels to complement existing IT security standards or create those channels as necessary to secure the SDLC. The SSG can also create a security portal (e.g., website or wiki) that houses SSDL information centrally [SR1.2]. Similar to the approach for prioritizing defect discovery efforts by categorizing attacks and bugs [AM3.5, CR2.7], we observe these emerging SSIs driving initial standards creation from industry top risks, leveraging general sources such as MITRE, ISO, and NIST to form baseline requirements.

Getting the word out about the organization's top risks and what can be done about them is a key part of the SSG's job. We observe these leaders using every channel possible (e.g., town halls, brown bags, communities of practice forums, messaging channels) to socialize the software security message and raise awareness of the SSDL [SM2.7].

CHECKLIST FOR EMERGING SSIs

1. **Create an SSG.** Put a dedicated group in charge and give them resources.
2. **Document and socialize the SSDL.** Tell all stakeholders the expectations for software security.
3. **Inventory applications.** Decide on what you're going to focus on first, then apply good risk management.
4. **Apply infrastructure security.** Don't put good software on bad systems or in poorly constructed networks (cloud or otherwise).
5. **Deploy defect discovery.** Determine the issues in today's in-progress and production applications, then plan for tomorrow.
6. **Manage discovered defects.** Resolving issues reduces risk.
7. **Publish and promote.** Roll out the secure SDLC and promote it both bottom-up and top-down.

INVENTORY APPLICATIONS

One of the first activities for any SSG is to create an initial inventory of the application portfolio under its purview [AM1.2, CMVM2.3]. As a starting point, the inventory should include each application's important characteristics (e.g., programming language, architecture type, open source used [SR1.5]). Particularly useful for monitoring and incident response activities [CMVM1.1], many organizations will include relevant operational data in the inventory (e.g., where the application is deployed, owners, emergency contacts).

Inventory efforts tend to favor a top-down approach in the beginning, usually starting with a questionnaire to elicit data from business managers who serve as application owners, then using tools to find OSS. The SSG also tends to focus on understanding where sensitive data resides and flows (e.g., PII inventory) [CP2.1] and the resulting business risk level associated with the application (e.g., critical, high, medium, low).

When working with engineering teams, these efforts commonly attempt to extract software inventory data from the tools used to manage IT assets. By scraping these software and infrastructure configuration management databases or code repositories, the SSG crafts an inventory brick by brick rather than top-down.

Maintaining an application inventory is a capability to be built over time rather than a one-time effort. To remain accurate and current, the inventory must be regularly monitored and updated. As with all data currency efforts, it's important to make sure the data isn't overly burdensome to collect and is periodically spot-checked for validity. Organizations should favor automation for application discovery and management whenever possible.

APPLY INFRASTRUCTURE SECURITY

Bad infrastructure security can undermine good software security, which means the SSG must ensure that host and network security basics are in place [SE1.2, SE3.10] as well as cloud security controls [SE1.3]. Security engineers might begin by conducting this work manually, then baking these settings and changes into their software-defined infrastructure scripts [SE1.4] to ensure both consistent use within a development team and scalable sharing across the organization.

Forward-looking organizations that have adopted software and network orchestration technologies [SE2.7] (e.g., Kubernetes, Envoy, Istio) get maximum impact from them with the efforts of even an individual contributor, such as a security-minded DevOps engineer. Though many of the technologies in which security engineers specify hardening and security settings are human-readable, engineering groups don't typically take the time to extract and distill a document-based security policy from these codebases.

DEPLOY DEFECT DISCOVERY

Regardless of business drivers, one of the quickest ways of transitioning unknown risk to managed risk is through defect discovery—automated tools, both static and dynamic, provide fast, regular insight into the portfolio security posture, with experts doing the detailed testing for important applications [AA1.1, CMVM3.4]. While not necessarily done for the entire application portfolio, conducting some targeted vulnerability discovery to get a feel for the current risk posture allows firms to motivate the necessary conversations with stakeholders to gain buy-in and prioritize remediation. Organizations tend to determine their high-priority applications via risk ranking [AA1.4], then phase in a combination of manual testing techniques against these high-priority applications, relying on automated testing techniques for portfolio coverage.

Static and dynamic software testing techniques each provide unique views into an application's security posture. Static analysis can look for issues inside the code the organization develops [CR1.4] and inside third-party components [SR1.5]. Dynamic application security testing (DAST) [ST1.4] can uncover immediately exploitable issues and help provide steps to reproduce attacks. In addition, QA groups can help ensure that development streams are adhering to security expectations. All these testing results assist with prioritization and displaying impact to executive leadership.

Manual testing efforts generally start by bringing in third-party assessors [PT1.1] on a regular cadence, either upon major milestones or, more commonly, as a periodic out-of-band exercise to assess the most critical applications. Even where an internal penetration testing function exists, a third party periodically bringing in a unique perspective will be beneficial.

Note that engineering groups will tend to favor empowering pipelines and testers with automation and allow engineering leadership or individual engineering teams to define some aspects of mandatory testing and remediation timelines. It's important to ensure static, dynamic, and manual testing creates minimal unnecessary friction in engineering processes.

MANAGE DISCOVERED DEFECTS

Unaddressed security defects are unmanaged risks. At first, there will be a large backlog of discovered security defects that will have to be bundled and passed through the risk exception process and prioritized into the development backlog. After resolving the technical debt, the ongoing defect management process should be designed to deal with security defects as they are introduced to prevent their release into production systems.

When security defects are discovered, it is the responsibility of the SSI to make sure they are logged and tracked through to completion [CMVM1.3]. Security defects can come from diverse sources, including penetration testers [PT1.2], security tooling [CR1.4], and operations [CMVM1.2], but ideally, they are logged in a single source of truth for tracking purposes.

PUBLISH AND PROMOTE THE PROCESS

With a strategy in hand, an understanding of the portfolio, and security expectations set with engineering teams, the SSG documents the SSDL [SM1.1] and begins collecting telemetry [SM1.4]. The SSDL should include clearly documented goals, roles, responsibilities, and activities, but the most usable SSDLs also include process diagrams and provide contextual details for each stakeholder. Many organizations seeking to consolidate ad hoc efforts into an emerging SSI will find a variety of SSDLs in use across engineering teams. In these cases, the new SSDL might be a replacement for all such approaches, but it might also have some parts that are abstract enough to account for processes until they can be rolled into the new approach. Publication is a good time for the SSG to start a software security hub where the SSG can disseminate knowledge about processes and about software security as a whole [SR1.2].

In a top-down approach, organizations favor creating policy [CP1.3] and standards [SR1.1] that can be followed and audited like any other business process. Rather than documents, however, engineering teams might favor implementing their part of an SSDL inside of pipelines [SM3.4] and scripts [ST3.6] or by prescribing reusable security blocks that meet expectations. Over time, the SSG will also have to deliver some policy in the form of governance-as-code in engineering pipelines [SM1.4].

While executives have likely been engaged to get the SSI to this point, this is a good time to ensure that they're being regularly kept up to date with software security. Remember, executive teams need to understand not only how the SSI is performing

but also how other firms are solving software security problems and the ramifications of not investing in software security [SM1.3].

PROGRESS TO THE NEXT STEP IN YOUR JOURNEY

Usually done as part of moving to the mature stage, the SSG then proceeds to scale the SSI. For example, this scaling might be done by creating a security champions program, improving the inventory capability based on lessons learned, automating the basics, doing more prevention, and then repeating. As the initiative matures and the business grows, there will be new challenges for the SSG to address, so it will be crucial to ensure that feedback loops are in place for the program to consistently measure its progress and maturity.

FOR A MATURING SSI: HARMONIZING OBJECTIVES

With the foundations established, SSG leaders shift their attention to scaling risk-based controls across the entire software portfolio and enabling development to find and fix issues early in the software lifecycle. The SSI has likely reached the emerging stage across multiple capabilities (see Figure 5) and is maturing specific aspects of its initiative. This maturing includes both adding new activities and scaling existing ones (see Figure 6). It specifically includes building bridges between various software security efforts in corporate and engineering groups.




FIGURE 6. MOVING FROM EMERGING TO MATURING. Building an emerging SSI usually focuses on collecting activities into a single program. Moving from emerging to maturing requires ongoing iterative improvements and expansions. Piloting new capabilities (e.g., security champions or software supply chain risk management) likely requires reapplying the emerging approach for a specific set of activities.

This section on maturing an SSI repeats some of the foundational BSIMM activities from the “FOR AN EMERGING SSI” section. We do this because most organizations won’t treat SSI creation as a waterfall process. Instead, they will, for example, establish policy, set up a champions program, deploy defect discovery tools, etc., in overlapping, incremental improvement cycles. In addition, many organizations will determine in the emerging phase that some activities can wait a bit while engaging in other, more necessary, software security efforts. In either case, this is a good place for a reminder to keep working on foundational activities.

UNIFY STRUCTURE AND CONSOLIDATE EFFORTS

The first step is to ensure that there is a single SSI and to provide the proper resources for the owner tasked with shepherding the organization so the group can meet risk management objectives. At this point, the SSI might include multiple SSGs and owners (e.g., across major products or business units), so working to harmonize these efforts must be a key goal. The organization will want to ensure that the SSI is supported by a full-time team—an SSG—that can scale across the firm. Establishing this structure might not involve hiring staff immediately, but it will likely entail assembling a full-time team to implement key foundational activities central to supporting the assurance objectives further defined and institutionalized in policy [CP1.3], standards [SR1.1], and processes [SM1.1].

The SSG will require a mix of skills, including technical security knowledge, scripting and coding experience, and architectural skill. As organizations migrate toward their view of DevSecOps, the SSG might build its own software in the form of security automation, defect discovery in CI/CD pipelines, and infrastructure- and governance-as-code. SSGs often need to mentor, train, and work directly with developers, so communication skills, teaching ability, and practical knowledge are must-haves for at least some SSG staff. Essentially, the SSG is a group of people—whether one person, 10, or 100—who must improve the security posture of the software portfolio and all the processes that generate it, so management skills, risk management perspectives, an ability to contribute to engineering value streams, and an ability to break silos are critical success factors.

Within engineering teams, we see individuals taking on leadership roles such as product security engineer or security architect, while possessing functional titles such as Site Reliability Engineer, DevOps Engineer, or similar. Their responsibilities often include comparison and selection of security tools, definition of secure design guidelines and acceptable remediation actions, and implementation of infrastructure-as-code for secure packaging, delivery, and operations. Harmonizing leadership views across the SSG and engineering is a critical step to success.

CHECKLIST FOR MATURING SSIS

1. **Unify structure and consolidate efforts.** Formalize organization, staffing, objectives, budgets, and approach, then tell everybody about it.
2. **Expand security controls.** Increase program impact through policy, testing, training, and other quick wins.
3. **Engage development.** Use security champions to build bridges and harmonize software security objectives.
4. **Inventory and select in-scope software.** Expand the application inventory to include all software, not just applications.
5. **Enforce security basics everywhere.** Use automation to ensure that you run software only on good systems (cloud or otherwise).
6. **Integrate defect discovery and prevention.** Use automation and integration to scale and shift defect discovery and prevention everywhere.
7. **Upgrade incident response.** Ensure that software security experts are involved in all software security events and improve the program from lessons learned.
8. **Repeat and improve.** Growth does not happen in a straight line. You will have to revisit, remeasure, and replan multiple times.

EXPAND SECURITY CONTROLS

Next, it's time to use existing knowledge to choose the important software security activities to initiate, scale, or improve. This knowledge includes SSI scope, compliance, technology stacks, and deployment models, as well as the issues uncovered in defect discovery efforts. Common activity choices are policy [CP1.3], SDLC checkpoint conditions [SM1.4], testing [AA1.2, CR1.4, ST1.4, PT1.3, SR1.5], and training [T1.7], which are typically built out in a quick-win approach. When choosing and implementing new controls, it's often easier to get buy-in by showing adherence to well-known guidance (e.g., BSIMM, NIST SSDF, regulators) or choosing security controls that align with general industry guidance (e.g., OWASP, CWE, analysts). Ensure that activity selection includes an appropriate mix of preventive [SR1.1, SFD2.1] and detective (e.g., testing) controls to maximize positive impacts on the organization's risk posture.

ENGAGE DEVELOPMENT

As noted throughout this section, engineering teams are likely already thinking about various aspects of security related to design, configuration, infrastructure, and deployment. Engaging development begins by creating mutual awareness of how the SSG and development teams see the next steps in maturing security efforts—and successfully engaging early on relies on bridge-building and credentialing the SSG as competent in development culture, toolchains, and technologies. It also includes building awareness around which security capabilities constitute an SSDL and beginning to determine how those capabilities are expected to be conducted. Building consensus on what role each department will play in improving capabilities over the next evolutionary cycle greatly facilitates success.

To facilitate tool adoption, the SSG might dedicate some portion of their efforts or build a team of security champions [SM2.3] to serve as tool mentors to help development teams not only integrate the tools but also triage and interpret results [CR1.7]. Although the primary objective is to embed security leadership inside development, these individuals also serve as both key points of contact and interface points for the SSG to interact with engineering teams and monitor progress. Because they are local to teams, champions can facilitate defect management goals, such as tracking recurring issues to drive remediation [PT1.2]. The SSG can also roll out software security training [T2.9] tailored to the most common security defects identified through AST, often cataloged by technology stack and coding language.

INVENTORY AND SELECT IN-SCOPE SOFTWARE

It's important to take an enterprise-wide perspective when building a view into the software portfolio. Engaging directly with application business owners by, for example, using questionnaire-style data gathering is a good start. It's useful to focus on applications (with owners who are responsible for risk management) as the initial unit of inventory measure, but remember that many vital software components aren't applications (e.g., libraries, APIs, scripts, pipeline tests, infrastructure-as-code). In addition to understanding application characteristics (e.g., programming language, architecture type such as web or mobile, the revenue generated) as a view into risk, you'll want to capture and maintain the same information for all software, focusing on understanding where sensitive data resides and flows (e.g., PII inventory) [CP2.1] along with the status of active development projects.

Rather than taking an organizational structure and owner-based view, engineering teams usually attempt to understand software inventory by extracting it from the same tools they use to manage their IT assets. They usually combine two or more of the following approaches to software inventory creation:

- Discovery, import, and visualization of assets managed by the organization's cloud and data center virtualization management consoles
- Scraping and extracting assets and tags from infrastructure-as-code held in code repositories, as well as processing metadata from container and other artifact registries

Essentially, the SSG is a group of people—whether one person, 10, or 100—who must improve the security posture of the software portfolio.

- Outside-in web and network scanning for publicly discoverable assets, connectivity to known organizational assets, and related ownership and administrative information

With a software inventory in hand, the SSG will impose security requirements using formalized risk-based approaches to cover as much of the software portfolio as possible. Using simple criteria (e.g., software size, regulatory constraints, internal-facing vs. external-facing, data classification), they'll assign a risk classification (e.g., high, medium, low) to each application [AA1.4] and define the initial set of software and project teams with which to prototype security activities. Although application risk classifications are often the primary driver, we have observed firms using other information, such as whether a major change in application architecture is being undertaken (e.g., shift to a cloud-native architecture) or whether the software contains critical code (e.g., cryptography, proprietary business logic). Firms find it beneficial to include in the selection process some engineering teams that are already doing some security activity organically.

Engineering teams might have a different idea of what in-scope software means relative to the security efforts they already have underway—if they're working on one application, then that application is likely to be their scope. When required to prioritize specific applications' components, we observe engineering teams using the following as input:

- Teams conducting active new development or major refactoring (velocity)
- Those services or data repositories to which specific development or configuration requirements for security or privacy apply [CP1.1, CP1.2] (regulation)
- Software that solves critical technical challenges or that adopts key technologies (opportunity)

Prioritized software is then usually the target for test automation [ST2.5], vulnerability discovery tooling, or security features [SFD1.1].

ENFORCE SECURITY BASICS EVERYWHERE

Commonly observed today regardless of SSG age are basic security controls enforced in hosts and networks [SE1.2] and in cloud environments [SE1.3]. A common strength for organizations that have good controls over the infrastructure assets they manage, these basics are accomplished through a combination of IT provisioning controls, written policy, prebuilt and tested golden images, sensors and monitoring capabilities, server hardening and configuration standards, infrastructure-as-code, and entire groups dedicated to patching. As firms migrate private infrastructure to cloud environments, organizations must carefully reestablish their assurance-based controls to maintain and verify adherence to security policy. To keep tabs on the growing number of virtual assets created by engineering groups and their automation, organizations often must deploy custom solutions [AM2.9] to overcome limitations in a cloud provider's ability to meet desired policy.

Governance and engineering teams often cooperate to build in enforced security basics for infrastructure and cloud

environments, leveraging containers [SE2.5], infrastructure-as-code [SE1.4], and orchestration [SE2.7]. Over time, these security basics expand to include internal development environments, toolchains, deployment automation, code repositories, and other important infrastructure.

INTEGRATE DEFECT DISCOVERY AND PREVENTION

Initial defect discovery efforts tend to be one-off (by using centralized commercial tools [CR1.2]) and to target the most critical applications, with a plan to scale efforts over time. Scaling prioritization might be selected for compliance or contractual reasons or because it applies to a phase of the software lifecycle (e.g., shift everywhere to do threat modeling at design time [AA1.1], composition analysis on software repositories [SE3.8], SAST during development [CR1.4], DAST in preproduction [ST1.4], and penetration testing on deployed software [PT1.1, PT1.3]). The point is to automate and scale the chosen defect discovery activities, but scaling through automation and integration must come without disrupting CI/CD pipelines (e.g., due to tools having long execution times), without generating large volumes of perceived false positives, and without impeding delivery velocity (e.g., through opaquely breaking builds or denying software promotion) except under clear or agreed-upon circumstances.

In addition to defect discovery, engineering teams might favor prevention controls they can apply to software directly in the form of security features [SFD1.1]. These controls can take the form of microservices (e.g., authentication or other identity and access management) [SE2.5], common product libraries (e.g., encryption) [SFD2.1], or even infrastructure security controls (e.g., controlling scope of access to production secrets through vault technologies).

Some engineering groups have taken steps to tackle the prevention of certain classes of vulnerability in a wholesale manner [CMVM3.1], using development frameworks that preclude them. Organizations will want to ask security-minded engineers for their opinion about framework choices and empower them to incorporate their understanding of security features and security posture tradeoffs.

UPGRADE INCIDENT RESPONSE

Ensuring that defined incident response processes include SSG representation [CMVM1.1] and determining whether an incident has software security roots requires specific skills that are not often found in traditional IT groups. The organization will want to work with engineering teams, especially DevOps engineers, to help make the connections between those events and alerts raised in production and the associated artifacts, pipelines, repositories, and responsible teams. This traceability allows these groups to effectively prioritize security issues on which the SSG will focus. Feedback from the field on what is happening greatly enhances the top N lists ([AM3.5, CR2.7]) that many organizations use to help establish priorities.

Security engineers who are in development teams and more familiar with application logic might be able to facilitate instructive monitoring and logging. They can coordinate with DevOps engineers to generate in-application defenses that are tailored for business logic and expected behavior,

therefore, they are likely more effective than, for example, WAF rules. Introducing such functionality will in turn provide richer feedback and allow a more tailored response to application behavior [SE3.3].

Organizations deploying cloud-native applications using orchestration might respond to incidents (or to data indicating imminent incidents) with an increase in logging, perhaps by adjusting traffic to the distribution of image types in production. Much of this is possible only with embedded security engineers who are steeped in the business context of a development team and have good relationships with that team's DevOps engineers, or security champions can be a good resource for these individuals. Under these circumstances, incident response moves at the speed of a well-practiced single team [CMVM1.4] rather than that of an interdepartmental playbook.

REPEAT AND IMPROVE

As noted earlier, working through activity growth for emerging and maturing SSIs probably won't happen in a straight line. There'll be changes in priorities, resources, and responsibilities, along with changes in attackers, attacks, technologies, and everything else. It's necessary to take time periodically to determine how well the SSI is performing against business objectives and adjust as necessary.

As a reminder, organizations rarely move their entire SSI from emerging to maturing to enabling all at once. Different parts of the SSI will shift between emerging, maturing, and enabling a few times over the years, with different timing that SSG leaders will need to plan for.

FOR AN ENABLING SSI: DATA-DRIVEN IMPROVEMENTS

Achieving software security scale—of expertise, portfolio coverage, tool integration, vulnerability discovery accuracy, process consistency, etc.—remains a top priority. However, firms often scale one or two capabilities (e.g., defect discovery, training) and fail to scale others (e.g., AA, vendor management). Given mature activities, there's a treasure trove of data to be harvested and included in KPI and KRI reporting dashboards, but executives will start asking the very difficult questions: Are we getting better? Is our implementation working well? Where are we lagging? How can we go faster with less overhead? What's our message to the board? The efficacy of an SSI will be supported by ongoing data collection and metrics reporting that seeks to answer such questions [SM3.3].

PROGRESS ISN'T A STRAIGHT LINE

As mentioned earlier, organizations don't always progress from emerging to maturing to enabling in one try or on a straight path, some SSI capabilities might be enabling while others are still emerging or maturing. Based on our experience, firms with some portion of their SSI operating in an enabling state have likely been in existence for longer than three years. Although we don't have enough data to generalize enabling SSIs, we do see common themes for those that strive to reach this state:

- **Top N risk reduction.** Everyone relentlessly identifies and closes top N weaknesses, placing emphasis on obtaining visibility into all sources of vulnerability, whether in-house developed code, open source code [SR2.7], vendor code [SR3.2], toolchains, or any associated environments and processes [SE1.2, SE1.3]. These top N weaknesses are most useful when specific to the organization, evaluated at least annually, and tied to metrics to prioritize SSI efforts that improve risk posture.
- **Tool customization.** Security leaders place a concerted effort on tuning tools (e.g., customization for static analysis, fuzzing, penetration testing) to improve integration, accuracy, consistency, and depth of analysis [CR2.6, ST2.6, AM3.2, PT3.2]. Customization focuses not only on improving result fidelity and applicability across the portfolio but also on pipeline integration and timely execution, improving ease of use for everyone.
- **Feedback loops.** Loops are created between SSDL activities to improve effectiveness as deliverables from activities ebb and flow with each other. For example, an expert in QA might leverage AA results when creating security test cases [ST3.3]. Likewise, feedback from the field might be used to drive SSDL improvement through enhancements to a hardening standard [CMVM3.2]. The concept of routinely conducting blameless postmortems to find root causes and drive remediation seems to be gaining ground in some firms.
- **Data-driven governance.** The more mature groups instrument everything to collect data that in turn becomes metrics for measuring SSI efficiency and effectiveness against KRIs and KPIs [SM3.3]. As an example, a metric such as defect density might be leveraged to track performance of individual business units and application teams. Metrics choices are very specific to each organization and also evolve over time.

Achieving software security scale—of expertise, portfolio coverage, tool integration, vulnerability discovery accuracy, process consistency, etc.— remains a top priority.


PUSH FOR AGILE-FRIENDLY SSIs

In recent years, we've observed governance-oriented teams—often out of necessity to remain in sync with engineering teams—evolving to become more Agile-friendly:


- Putting “Sec” in DevOps is becoming a mission-critical objective. SSG leadership routinely partners with IT, cloud, development, QA, and operations leadership to ensure that the SSI mission aligns with DevOps values and principles.
- SSG leaders realize they need in-house talent with coding expertise to improve not only their credibility with engineering but also their understanding of modern software delivery practices. Job descriptions for SSG roles now mention experience and qualification requirements such as cloud, mobile, containers, and orchestration security, as well as coding. We expect this list to grow as other topics become more mainstream, such as architecture and testing requirements around serverless computing and single-page application approaches.
- To align better with DevOps values (e.g., agility, collaboration, responsiveness), SSG leaders are beginning to replace traditional people-driven activities with people-optional, pipeline-driven automated tasks. This often comes in the form of automated security tool execution, bugs filed automatically to defect notification channels, builds flagged for critical issues, and automated triggers to respond to real-time operational events.

- Scaling outreach and expertise through the implementation of an ever-growing security champions program is viewed as a short-term rather than long-term goal. Organizations report improved responsiveness and engagement as part of DevOps initiatives when they've localized security expertise in the engineering teams. Champions are also becoming increasingly sophisticated in building reusable artifacts (e.g., security sensors) in development and deployment streams to directly support SSI activities.
- SSG leaders are partnering with operations to implement application-layer production monitoring and automated mechanisms for responding to security events. There is a high degree of interest in consuming real-time security events for data collection and analysis to produce useful metrics.

In summary, engineering teams have likely taken an enabling approach from the beginning. Their security efforts are contributions from engineers who deliver software early and often, constantly improving it rather than relying on explicit strategy backed by top-down policies. They make their software available to everyone to prevent future issues and use evangelism to encourage uptake. They review production failures and make changes, often with automation, to their toolchains and processes. That said, perceptions of business and technical risk between corporate and engineering groups often differ in substantial ways. Bringing the groups together to share responsibilities for software security, as well as definitions of and goals for needed risk management, while enabling broad stakeholder productivity is a primary goal for any SSI.




PART 4: BSIMM PARTICIPANTS



PART 4: BSIMM PARTICIPANTS

BSIMM participants comprise software security leaders and team members from around the globe. They have a common mission to continuously improve their SSIs in light of changes in the world around them. You can use the information they've provided to learn from their efforts.

This 2024 edition of the BSIMM report—BSIMM15—examines anonymized data from the software security activities of 121 organizations. This diverse group spans multiple sizes of security teams, development teams, and software portfolios, as well as regions, vertical markets, and security team ages.



PARTICIPANTS

The participating organizations fall across various verticals, including cloud, financial services, FinTech, ISVs, insurance, IoT, healthcare, and technology organizations (see Figure 7).

Unique in the software security industry, the BSIMM project has grown from nine participating companies in 2008 to 121 in 2024, currently with about 3,400 software security group members and more than 7,700 security champions. Today, the average age of participants' SSIs is 5.4 years (see Table 5).




FIGURE 7. BSIMM15 PARTICIPANTS. Participant percentages per tracked region and vertical.

| BSIMM PARTICIPANT NUMBERS OVER TIME | | | | | | | | |
|-------------------------------------|-----------|------------|------------|------------|------------|------------|------------|------------|
| | BSIMM15 | BSIMM14 | BSIMM13 | BSIMM12 | BSIMM11 | BSIMM10 | BSIMM9 | BSIMM8 |
| FIRMS | 121 | 130 | 130 | 128 | 130 | 122 | 120 | 9 |
| SSG MEMBERS | 3,428 | 3,527 | 3,342 | 2,837 | 1,801 | 1,596 | 1,600 | 370 |
| SECURITY CHAMPIONS | 7,703 | 7,427 | 8,508 | 6,448 | 6,656 | 6,298 | 6,291 | 710 |
| DEVELOPERS | 259,326 | 267,731 | 408,999 | 398,544 | 490,167 | 468,500 | 415,598 | 67,950 |
| APPLICATIONS | 90,747 | 96,361 | 145,303 | 153,519 | 176,269 | 173,233 | 135,881 | 3,970 |
| AVG. SSG AGE (YEARS) | 5.44 | 5.24 | 5.00 | 4.41 | 4.32 | 4.53 | 4.13 | 5.32 |
| SSG MEMBERS TO DEVELOPERS | 4.4 / 100 | 3.87 / 100 | 3.01 / 100 | 2.59 / 100 | 2.01 / 100 | 1.37 / 100 | 1.33 / 100 | 1.13 / 100 |

TABLE 5. BSIMM PARTICIPANT NUMBERS OVER TIME. The chart shows how the BSIMM study has grown over the years.

ACKNOWLEDGEMENTS

Our thanks to the 121 executives, including those who wish to remain anonymous, from the SSIs we studied to create BSIMM15.

Our thanks also to the nearly 165 individuals who helped gather the data for the BSIMM data pool over time.

In particular, we thank Adam Brown, Aditi Gupta, Akhil Mittal, Akira Watanabe, Akshay Sawant, Alex Jupp, Alexios Fakos, Alistair Nash, Anders Stadum, Anil Gajawada, Aseem Lodha, Avi Sambira, Balaji Padmanabhan, Ben Hutchison, Bohuai Liu, Brendan Sheairs, Cem Nisanoglu, Chandu Ketkar, Daniel Cohen, David Johansson, Denis Sheridan, Derek Evans, Devaraj Munuswamy, Diya Ghosh, Don Pollicino, Durai G, Eason Yu, Eli Erlikhman, Faheem Sultan, Grant Van Gorder, Harshad Janorkar, Ibrahim Khan, Iman Louis, Jacob Ewers, Jamie Boote, Jas Alsuwailam, Jatin Virmani, John Tapp, Jonathan Dunfee,

Josh Brown, Kayalvizhi Devakumaran, Kevin Nassery, Lahiru Pinnaduwwage, Lekshmi Nair, Leo Berrun, Li Zhao, Manik Virmani, Matt Chartrand, Matt Maddox, Melih Tas, Michael Fabian, Mike Fabian, Mike Lyman, Nivedita Murthy, Rajiv Harish, Rakshitha Rao, Ravinder Reddy Amireddy, Rehan Bashir, Sachin Shetty, Sam Schueller, Smith Kaneria, Stanislav Sivak, Stephen Gardner, Surya Uddhi Nagaraj, Takeshi Ohmori, Thaddeus Bender, Thomas Madden, Tom Stripling, Tony Blakemore, Uzear Ahmed, Vijay Sharma, Warrie Proffitt, and Zhihao Yu.


We also thank David Johansson and Surya Uddhi Nagaraj for their work managing the BSIMM tooling and data and creating the extracts used in this report. In addition, we thank Jennifer Stout and Amy La Russa for their work on various aspects of this report.

BSIMM15 was authored by Jamie Boote, Ben Hutchison, Mike Lyman, Sammy Migués, and Sam Schueller with special guest authors Tim Mackey and David Benas.

| | | |
|------------------------------|-----------------------------------|----------------------------|
| AARP | Honeywell | Phison |
| Aetna | HUMAN Security | QlikTech |
| Airoha | Imperva | Realtek |
| AON | Inspur Software Intralinks | Reckitt |
| Arlo | International AB | Sammons Financial |
| Axway | iPipeline | ServiceNow |
| Bank of America | Johnson & Johnson | Signify |
| Bell Network | Landis+Gyr | SonicWall |
| CIBC | Lenovo | Synchrony Financial |
| Citi | MassMutual | TD Ameritrade |
| Diebold Nixdorf | MediaTek | Teradata |
| Egis Technology | Medtronic | U.S. Bank |
| Eli Lilly and Company | MiTAC | Unisoc |
| EQ Bank | Navient | Vanguard |
| Fidelity | Navy Federal Credit Union | Veritas |
| Finastra | NetApp | Vivo |
| Genetec | Oppo | ZoomInfo |
| HCA Healthcare | Pegasystems | |

This work is licensed under the Creative Commons Attribution-Share Alike 3.0 License. To view a copy of this license, visit <http://creativecommons.org/licenses/by-sa/3.0/legalcode> or send a letter to Creative Commons, 171 Second Street, Suite 300, San Francisco, California, 94105, USA.

PART 5: **QUICK GUIDE TO** **SSI MATURITY**



PART 5: QUICK GUIDE TO SSI MATURITY

Twelve questions can help clarify where your SSI is today. Combined with a detailed software security scorecard (see below on how to measure your own program) and knowledge about roles and responsibilities, you can use this information to plan strategic changes for ongoing success.

SSI maturity is a complex thing. Each organization will apply different values to efforts and progress in people, process, technology, and culture, but they will also evolve differently in their vision for success as well as how they spend resources, grow the program, and manage risk. This section provides an approach to organizing, growing, and maturing an SSI that works for everyone. Refer to Part 3 for more details.

A BASELINE FOR SSI LEADERS

All program leaders require a detailed understanding of their efforts and whether those efforts align with business objectives. A good start here is to understand whether organizational SSI efforts align well with changes in the software security landscape driven by global events, digital transformation, and engineering evolution, as well as with how software is made today. Use your answers to the questions below to determine whether it's time to invest in new growth—and if you don't know all these answers, use the list to gather information from each SSI stakeholder responsible for aspects of software security risk management in your organization.

KEEPING PACE WITH CHANGE IN YOUR SOFTWARE PORTFOLIO

- Do you maintain at least a near-current view of all your software and development assets, including internal code, third-party code, open source, development environments and toolchains, infrastructure-as-code, and other software assets?
- Are you creating and using in your risk management processes SBOMs that detail all the components in the SSI's software portfolio?
- Do you have a near-real-time view of your operations environments, along with a view into their aggregate attack surface and aggregate risk?

CREATING THE DEVSECOPS CULTURE YOU NEED

- Are you building bridges between the various software security stakeholders in your organization—governance, technical, audit, vendor management, cloud, etc.—to align culture, approach, technology stacks, and testing strategies?
- Have you scaled your security champions program across your software portfolio, including skills specific to automation, technology stacks, application architectures, cloud-native development, and other important DevOps needs?
- Are you delivering important security policy, standards, and guidelines as-code that run in engineering and operations toolchains?

SHIFTING SECURITY EFFORTS EVERYWHERE IN THE ENGINEERING LIFECYCLE

- Are you automating security decisions to remove time-consuming manual review and moving toward a secure, auditable, governance-as-code-driven SDLC?
- Are you following a shift everywhere strategy to move from large, time-consuming security tests to smaller, faster, timelier, pipeline-driven security tests conducted to improve engineering team performance?
- Are you managing supply chain risk through vendor software assurance, governance-driven access and usage controls, maintenance standards, and collected provenance data?

MEASURING YOUR SSI

- Do you routinely use telemetry from security testing, operations events, risk management processes, event postmortems, and other efforts to drive process and automation improvements in your DevOps toolchain or governance improvements in your policies and standards?
- Does your SSI strategy include security efforts needed specifically for modern technologies, such as cloud, container, orchestration, open source management, development pipeline, etc.?
- Are you actively experimenting with new technologies, such as AI and LLMs, that have the opportunity to integrate security and engineering functions while also reducing engineering friction?

Most organizations have already covered the basics of software security policy, testing, and outreach, but it takes a concerted effort to scale an SSI to address changes in portfolio size, technology, infrastructure, regulation, laws, attackers, attacks, and more. Internal review of efforts vs. needs is always a good way to move forward.

USING A BSIMM SCORECARD TO MAKE PROGRESS

A BSIMM scorecard is a management tool that allows your SSI and SSG leadership to:

- Assess your level of maturity so you can evolve your software security journey in stages, first building a strong emerging foundation, then scaling and maturing the more complex activities over time
- Communicate your software security posture to customers, partners, executives, and regulators. A scorecard helps everyone understand where you are and where you want to go in your journey when you're explaining your strategic plan and budgets
- See actual measurement data from the field, which helps in building a long-term plan for an SSI and in tracking progress against that plan

In addition to being a lens on the state of software security, the BSIMM serves as a measuring stick to determine where your SSI currently stands relative to our participants, whether as a whole or for specific verticals. A direct comparison of your efforts to the BSIMM15 scorecard for the entire data pool (see

Part 7) is probably the best first step. Follow the steps below to use the BSIMM to create your own SSI scorecard (see Figure 8 for an example).

UNDERSTAND YOUR ORGANIZATIONAL MANDATE

- Decide what the SSI is expected to accomplish. Who are the executive sponsors, and what resources are they expected to provide? From a RACI perspective, who are the responsible and accountable stakeholders? What metrics must be provided to executive management to demonstrate acceptable progress?
- Set the proper scope for the SSI. At a high level, describe the applicable software portfolio and the associated software ownership (e.g., risk managers). Ensure that you include all applications and related software in the SSG's remit.

BUILD THE SCORECARD

- Make a list of stakeholders to interview. No single person knows everything about a modern SSI, so ensure that you have broad coverage across the SSG, security champions, engineering, QA, operations, and security testing. As needed, extend the stakeholder list to include teams from reliability, cloud, privacy, training, infrastructure, resilience, AI/ML, and others whose efforts have a direct impact on software security.
- Understand the BSIMM. Review the BSIMM activities and gain an understanding of the practices, the individual activities, and the connected themes that run through them. For example, the activities for software security testing appear across multiple BSIMM practices.
- Interview everyone and consolidate the results. Keep interviews brief and focused on the intersection of the interviewee's role and specific BSIMM activities. Ensure that you get the data and artifacts that demonstrate the organization is sufficiently—in both depth and breadth—performing each activity before you award credit.
- Create your scorecard. Use a binary 1 or 0, a scale of low, medium, and high, or even a graduated scale such as a percentage to combine aspects of depth, breadth, and maturity.

MAKE A STRATEGIC PLAN AND EXECUTE

- Compare your scorecard first to your stakeholders' realistic expectations and then also to what's common in the data pool. Prioritize effort on the important gaps as well as those gaps with a long lead time. See Part 3 for more details on how to build an execution plan. Mark your calendar to revisit the scorecard in 12 to 18 months, document your progress, and create a new scorecard.

- Define and use metrics to gauge progress. Every program needs a barometer for success, and each organization finds different things to be the best indicators for them. Whether described as metrics, KPIs, KRIs, SLOs, or something else, use what works best for you, your executive team, and your board (with each likely needing different metrics).

For most organizations, a single aggregated scorecard covering the entire SSI will suffice to inform future planning. In some cases, however, it will be beneficial to create individual scorecards for the SSG and for business units or application teams that have varying software security approaches or maturity levels.

Figure 8 depicts an example firm that performs 41 BSIMM15 activities (noted as 1s in its EXAMPLEFIRM scorecard columns, e.g., SM1.1), including nine activities that are the most common in their respective practices (orange, e.g., CP1.2). Note the firm does not perform the most observed activities in the other three practices (gray boxes, e.g., SM1.4) and should take some time to determine whether these are necessary or useful to its overall SSI. The BSIMM15 FIRMS columns show the number of observations (currently out of 121) for each activity, allowing the firm to understand the activity's general popularity within the current data pool. If you want to evaluate your scorecard against a particular vertical, refer to Part 9.

Once you have determined where you stand with activity efforts compared to your expectations, you can devise a plan for improvement. Organizations almost always choose some hybrid of expanding their SSI with new activities and scaling some existing activities across more of the software portfolio and stakeholder teams.


Note that there's no inherent reason to adopt all activities in each practice. Prioritize the ones that make sense for your organization today and set aside those that don't—but revisit those choices periodically. Once they've adopted an activity set, most organizations strategically work on the depth, breadth, and cost-effectiveness (e.g., via automation) of each activity in accordance with their view of the risk management efforts required in their environments for their business objectives.

To help refine the current and future activity prioritization for your SSI, you can go beyond the BSIMM15 FIRMS data in Part 7 to Figure 18 and analyze how SSIs evolve with remeasurements (see DATA ANALYSIS: LONGITUDINAL in Part 9) and with age (see DATA ANALYSIS: SSG in Part 9). You can also examine what's different about your vertical or verticals (see DATA ANALYSIS: VERTICALS AND PRACTICES in Part 9) and understand the impact of a champions program (see DATA ANALYSIS: SECURITY CHAMPIONS in Part 9) on SSIs.

| GOVERNANCE | | | INTELLIGENCE | | | SSDL TOUCHPOINTS | | | DEPLOYMENT | | |
|---------------------|----------------------------|--------------|----------------------------|----------------------------|--------------|-----------------------|----------------------------|--------------|-----------------------------|----------------------------|--------------|
| ACTIVITY | BSIMM15 FIRMS (OUT OF 121) | EXAMPLE FIRM | ACTIVITY | BSIMM15 FIRMS (OUT OF 121) | EXAMPLE FIRM | ACTIVITY | BSIMM15 FIRMS (OUT OF 121) | EXAMPLE FIRM | ACTIVITY | BSIMM15 FIRMS (OUT OF 121) | EXAMPLE FIRM |
| STRATEGY & METRICS | | | ATTACK MODELS | | | ARCHITECTURE ANALYSIS | | | PENETRATION TESTING | | |
| [SM1.1] | 90 | 1 | [AM1.2] | 64 | | [AA1.1] | 99 | 1 | [PT1.1] | 104 | |
| [SM1.3] | 72 | | [AM1.3] | 49 | 1 | [AA1.2] | 56 | 1 | [PT1.2] | 95 | 1 |
| [SM1.4] | 109 | | [AM1.5] | 77 | | [AA1.4] | 55 | | [PT1.3] | 76 | 1 |
| [SM1.7] | 72 | | [AM2.1] | 18 | | [AA2.1] | 37 | | [PT2.2] | 39 | |
| [SM2.1] | 65 | | [AM2.6] | 12 | 1 | [AA2.2] | 38 | 1 | [PT2.3] | 51 | |
| [SM2.3] | 67 | | [AM2.7] | 16 | 1 | [AA2.4] | 40 | 1 | [PT3.1] | 26 | 1 |
| [SM2.6] | 69 | | [AM2.8] | 24 | | [AA3.1] | 20 | | [PT3.2] | 21 | |
| [SM2.7] | 52 | 1 | [AM2.9] | 18 | | [AA3.2] | 8 | | | | |
| [SM3.1] | 30 | | [AM3.2] | 8 | | [AA3.3] | 18 | | | | |
| [SM3.2] | 23 | | [AM3.4] | 11 | | | | | | | |
| [SM3.3] | 32 | | [AM3.5] | 10 | | | | | | | |
| [SM3.4] | 11 | | | | | | | | | | |
| [SM3.5] | 1 | | | | | | | | | | |
| COMPLIANCE & POLICY | | | SECURITY FEATURES & DESIGN | | | CODE REVIEW | | | SOFTWARE ENVIRONMENT | | |
| [CP1.1] | 98 | 1 | [SFD1.1] | 93 | 1 | [CR1.2] | 80 | 1 | [SE1.1] | 76 | |
| [CP1.2] | 105 | 1 | [SFD1.2] | 83 | 1 | [CR1.4] | 106 | 1 | [SE1.2] | 102 | 1 |
| [CP1.3] | 94 | 1 | [SFD2.1] | 42 | | [CR1.5] | 75 | | [SE1.3] | 87 | 1 |
| [CP2.1] | 49 | | [SFD2.2] | 68 | | [CR1.7] | 51 | | [SE1.4] | 74 | 1 |
| [CP2.2] | 58 | | [SFD3.1] | 18 | | [CR2.6] | 24 | 1 | [SE2.4] | 51 | |
| [CP2.3] | 69 | | [SFD3.2] | 21 | | [CR2.7] | 19 | | [SE2.5] | 64 | 1 |
| [CP2.4] | 60 | | [SFD3.3] | 12 | | [CR2.8] | 27 | 1 | [SE2.7] | 42 | 1 |
| [CP2.5] | 70 | 1 | | | | [CR3.2] | 16 | | [SE3.2] | 24 | |
| [CP3.1] | 36 | | | | | [CR3.3] | 6 | | [SE3.3] | 17 | |
| [CP3.2] | 39 | | | | | [CR3.4] | 3 | | [SE3.6] | 25 | |
| [CP3.3] | 13 | | | | | [CR3.5] | 6 | | [SE3.8] | 3 | |
| | | | | | | | | | [SE3.9] | 3 | |
| | | | | | | | | | [SE3.10] | 0 | |
| TRAINING | | | STANDARDS & REQUIREMENTS | | | SECURITY TESTING | | | CONFIG. MGMT. & VULN. MGMT. | | |
| [T1.1] | 62 | 1 | [SR1.1] | 84 | 1 | [ST1.1] | 102 | 1 | [CMVM1.1] | 111 | 1 |
| [T1.7] | 61 | 1 | [SR1.2] | 96 | 1 | [ST1.3] | 79 | 1 | [CMVM1.2] | 85 | |
| [T1.8] | 50 | | [SR1.3] | 86 | | [ST1.4] | 54 | | [CMVM1.3] | 89 | 1 |
| [T2.5] | 41 | | [SR1.5] | 96 | 1 | [ST2.4] | 20 | | [CMVM1.4] | 88 | |
| [T2.8] | 25 | 1 | [SR2.2] | 70 | | [ST2.5] | 34 | | [CMVM2.3] | 46 | |
| [T2.9] | 31 | 1 | [SR2.5] | 62 | 1 | [ST2.6] | 24 | | [CMVM2.4] | 41 | |
| [T2.10] | 25 | | [SR2.7] | 55 | | [ST3.3] | 16 | | [CMVM3.1] | 13 | |
| [T2.11] | 29 | | [SR3.2] | 19 | | [ST3.4] | 5 | | [CMVM3.2] | 24 | |
| [T2.12] | 38 | | [SR3.3] | 17 | | [ST3.5] | 6 | | [CMVM3.3] | 22 | 1 |
| [T3.1] | 8 | | [SR3.4] | 20 | | [ST3.6] | 8 | | [CMVM3.4] | 31 | 1 |
| [T3.2] | 14 | | [SR3.5] | 0 | | | | | [CMVM3.5] | 17 | |
| [T3.6] | 9 | | | | | | | | [CMVM3.6] | 4 | |
| | | | | | | | | | [CMVM3.8] | 1 | |

FIGURE 8. BSIMM15 EXAMPLE FIRM SCORECARD. A scorecard helps everyone understand the software security efforts that are currently underway. It also helps organizations make comparisons to participants and serves as a guide on where to focus next.

PART 6: **SOFTWARE** **SECURITY** **INITIATIVE BY** **THE NUMBERS**



PART 6: SOFTWARE SECURITY INITIATIVE BY THE NUMBERS

ROLES IN A SOFTWARE SECURITY INITIATIVE

An SSI requires thoughtful staffing with both full-time and dotted-line people. You can use the descriptions below to help define roles and responsibilities that accommodate your needs for execution and growth.

EXECUTIVE LEADERSHIP

Historically, security initiatives that achieve firm-wide impact are sponsored by a senior executive who creates an SSG

where software security governance and testing are distinctly separate from software delivery (even when the groups have many shared responsibilities). Security initiatives without that executive sponsorship, by comparison, have historically had little lasting impact across the firm. By identifying a senior executive and putting them in charge of software security, the organization can address two “Management 101” concerns: accountability and empowerment.

In BSIMM-V, we saw CISOs as the nearest executive in 21 of 67 firms, which grew in BSIMM6 to 31 of 78, and again for BSIMM7 with 52 of 95. Since then, the percentage has remained relatively flat, even as BSIMM participation has grown, as shown in Figure 9.

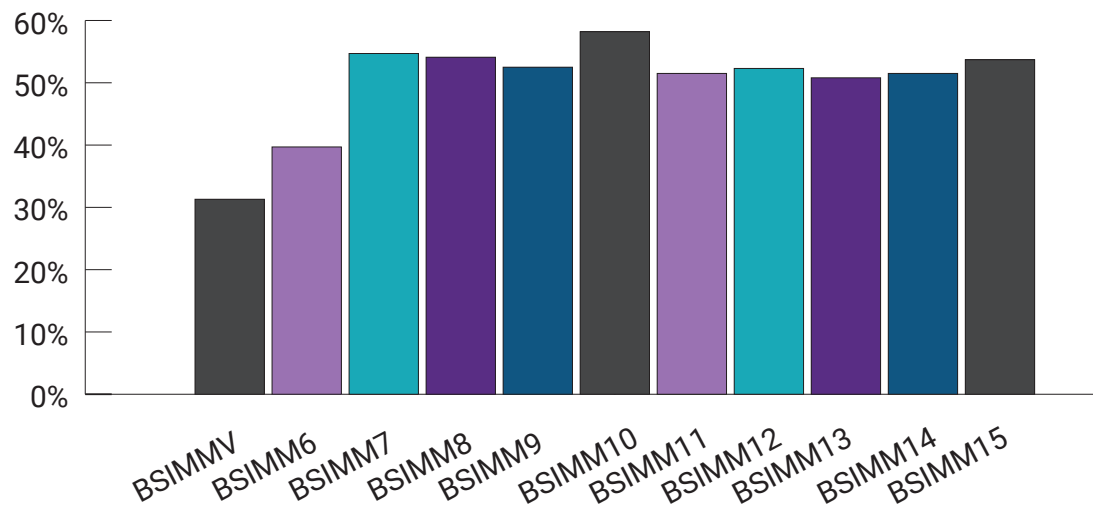


FIGURE 9. PERCENTAGE OF SSGs WITH A CISO AS THEIR NEAREST EXECUTIVE. Assuming new CISOs generally receive responsibilities for SSIs, this data suggests that CISO role creation is also flattening out.

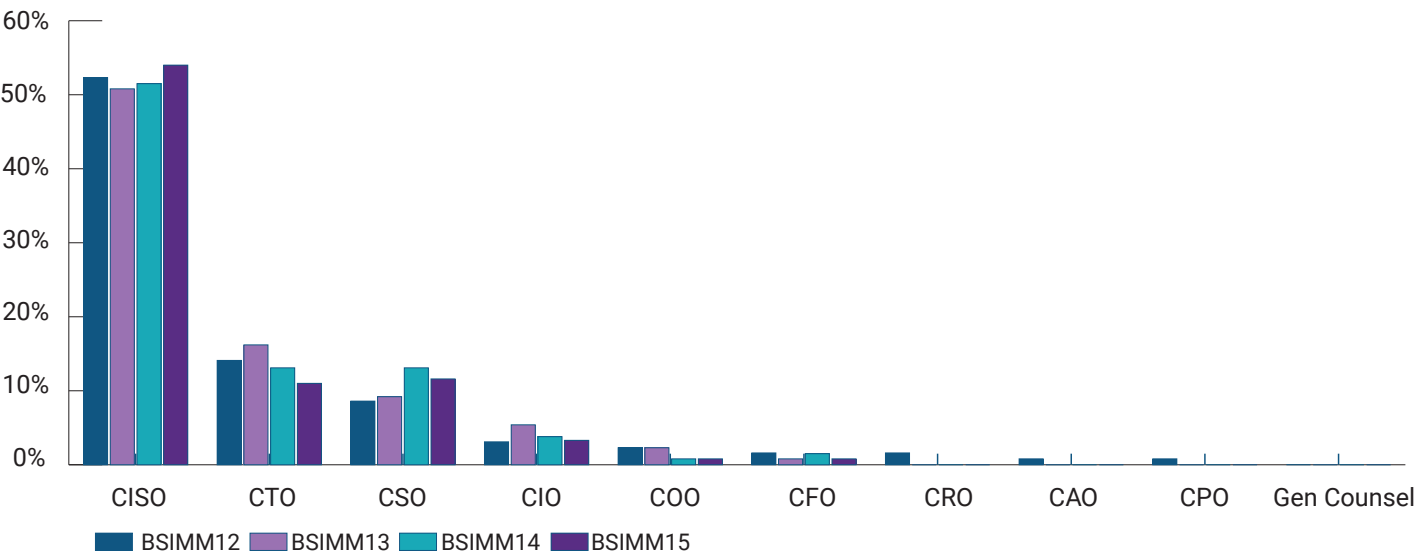


FIGURE 10. NEAREST EXECUTIVE TO SSG. Although many SSGs have a CISO as their nearest executive, we see a variety of executives overseeing software security efforts in the 121 BSIMM15 firms.

If we look across all the executives nearest to SSG owners, not just CISOs, we observe a large spread in the reporting path to executive leadership for BSIMM12 through BSIMM15, as shown in Figure 10.

The dark purple columns show by percentage the SSG leader's nearest executive in the BSIMM15 data pool, while the other columns show the percentages for previous BSIMMs. For example, a CISO is the closest executive in 54% of organizations (65 of 121) in the BSIMM15 data pool, and that percentage ranged from 50% to 58% in BSIMM7 through BSIMM14. Starting with the BSIMM13 data pool, we no longer see SSGs reporting directly to CRO (risk), CAO (assurance), CPO (privacy), and General Counsel roles. Note that for BSIMM15, we added nine firms and removed 18 others, which also affects analysis of reporting chains. Of course, across various organizations, not all people with the same title perform, prioritize, enforce, or otherwise provide resources for the same efforts in the same way.

CISOs, in turn, report to different executives among the 121 BSIMM15 firms. Figure 11 shows that CISOs report most commonly to CIOs (20 of 65, or almost 31% of the time) and report directly to the CEO about 15% of the time (10 of 65).




FIGURE 11. TO WHOM THE CISO REPORTS. For BSIMM15 participants, the CISO reports to a variety of roles, with the most common being the CIO, CTO, and a technology executive (e.g., head of engineering, architecture, or software).

SOFTWARE SECURITY GROUP LEADERS

SSG leaders are individuals in charge of day-to-day efforts in the 121 SSIs we studied for BSIMM15. They have a variety of titles, such as the following:

- Associate Vice President Enterprise Information Protection Product Security
- Chief Product Security Officer
- Cybersecurity Advisor
- Director (AST Lead)
- Director of Engineering Security
- Director Software Security
- Director, Application Security
- Head of Product & Data Security
- Head, Enterprise Security Architecture
- Lead Security Architect
- Managing Director, SSG Lead
- Product Security Group Director
- Product Security Manager
- Senior Manager, Information Security | Appsec Engineering
- Sr. Director of Security Engineering
- Sr. Director, Solutions Engineering
- SSG Chair
- VP of Cybersecurity
- VP Product & Application Security
- VP SecArch

When the SSG leader is an executive themselves, which happens 12% of the time (15 out of 121), they are CISOs almost 60% of the time (nine out of 15), with other titles being CTO, CPSO (Chief Product Security Officer), and CSO. As shown in Figure 12, SSG leaders are typically one or two hops from their nearest executive (e.g., a CxO or related technology organization title). In addition, we observed that this nearest executive is usually a further two hops away from the CEO.




FIGURE 12. SSG LEADERSHIP REPORTING CHAINS. SSG leaders are typically three or four hops away from the CEO.

SOFTWARE SECURITY GROUP (SSG)

Each of the 121 initiatives in BSIMM15 has an SSG—an organizational group or person dedicated to software security. In fact, without an SSG, successfully carrying out BSIMM activities across a software portfolio is very unlikely, so the creation of such a group is a crucial first step. The SSG might start as a team of one—just the SSG leader—and expand over time. The SSG might be entirely a corporate team, entirely an engineering team, or an appropriate hybrid. The team's name might also have an appropriate organizational focus, such as application security group or product security group, or perhaps DevSecOps.

Some SSGs are highly distributed across a firm whereas others are centralized. Even within the most distributed organizations, we find that software security activities are almost always coordinated by an SSG.

Although no two of the 121 firms we examined had exactly the same SSG structure, we did observe some commonalities. At the highest level, SSGs seem to come in five overlapping structures:

- Organized to provide software security services
- Organized around setting and verifying adherence to policy
- Designed to mirror business unit organizations
- Organized with a hybrid policy and services approach
- Structured around managing a matrixed team of experts doing software security work across the development or engineering organizations

Table 6 shows SSG-related statistics across the 121 BSIMM15 firms but note that large outliers affect the numbers this year. The “Notes” column shows the effect of removing outliers, or the top 10 firms, for that SSG characteristic. When planning the size and structure of your own SSG, consider the number of developers and applications to determine what resources you need to scale the SSI. Refer to DATA ANALYSIS: SSG in Part 9 for more details on how SSGs evolve over time.

SECURITY CHAMPIONS (SATELLITE)

In addition to the SSG, many SSIs have identified individuals (often developers, testers, architects, cloud and DevOps engineers, and other SDLC roles) who are a driving force in improving software security but are (likely) not directly employed in the SSG. We historically refer to this group as the satellite, while many organizations today refer to them as their software security champions. Champions can enable an SSI to scale its efforts while reducing dependency on the SSG team, and there appears to be a correlation between a higher BSIMM score and the presence of champions, as shown in Figure 13. Having security champions carry out software security activities removes SSG members from the engineering-critical path and empowers engineering teams to own their software security deliverables and share responsibility for software security objectives.

| THE SOFTWARE SECURITY GROUP | | | | | |
|---|----------|--------|---------|----------|--|
| STATISTICS | AVERAGE | MEDIAN | LARGEST | SMALLEST | NOTES |
| SSG Size | 28.3 | 8 | 892 | 1 | Average drops to 21.1 (one outlier) |
| SSG Members to Developer Ratio (per 100 developers) | 4.4 | 1.6 | 100 | 0.02 | Average drops to 3.6 (one outlier) Average drops to 2.1 (no top 10) |
| SSG to Developer Ratio (700+ Developers) – 62 Firms | 1.64 | 0.67 | 14.87 | 0.02 | |
| SSG to Developer Ratio (less than 700 Developers) – 59 Firms | 7.29 | 2.67 | 100 | 0.33 | |
| Number of Developers | 2,143.19 | 700 | 30,000 | 1 | |
| Number of Applications | 749.98 | 140 | 8,000 | 1 | |
| SSG Age | 5.44 | 4.5 | 23 | 0.1 | |
| Champions to Developer Ratio (per 100 developers) | 6.21 | 1.82 | 102.2 | 0 | Average drops to 4.6 (two outliers) Average drops to 2.9 (no top 10) |
| Champions to Developer Ratio (700+ Developers) – 62 Firms (per 100 developers) | 5 | 2.42 | 57.14 | 0 | |
| Champions to Developer Ratio (less than 700 Developers) – 59 Firms (per 100 developers) | 7.49 | 0 | 102.2 | 0 | |
| SSG to Application Ratio (per 100 applications) | 26.08 | 2.5 | 600 | 0.02 | Average drops to 16.4 (two outliers) Average drops to 9.8 (no top 10) |

TABLE 6. THE SOFTWARE SECURITY GROUP. We calculated the ratio of full-time SSG members to developers for the entire data pool by averaging the individual ratio for each participating firm. In the Notes column, we show the impact of removing outliers in the data.

Security champions are often chosen for software portfolio coverage (with one or two members in each engineering group) and sometimes for reasons such as technology stack coverage or geographical reach. The champions can act as a sounding board for the feasibility and practicality of proposed software security changes and improvements. Understanding how SSI governance changes might affect project timelines and budgets helps the champions proactively identify potential frictions and minimize them.

A successful security champions programs gets together regularly to compare notes, learn new technologies, and expand stakeholder understanding of the organization’s software security challenges. Motivated individuals often share digital work products, such as sensors, code, scripts, tools, and security features, rather than, for example, getting together to discuss enacting a new policy. Specifically, these proactive champions are working bottom-up and delivering software security features and awareness through implementation.

For more information about security champions, refer to DATA ANALYSIS: SECURITY CHAMPIONS in Part 9.

OTHER KEY STAKEHOLDERS

SSIs are truly cross-departmental efforts that involve a variety of stakeholders:

- Builders, including developers, architects, and their managers, must practice security engineering, taking some responsibility for both the definition of “secure enough” as well as ensuring that what’s delivered achieves the desired posture. An SSI requires collaboration between the SSG and these engineering teams to carry out the activities described in the BSIMM.
- Testers typically conduct functional and feature testing, but moving on to include security testing is very useful. Some testers are beginning to anticipate how software architectures and infrastructures can be attacked and are working to find an appropriate balance between automated and manual testing to ensure adequate security testing coverage.
- Operations teams must continue to design, defend, and maintain resilient environments because software security doesn’t end when software is “shipped.” In accelerating trends, development and operations are collapsing into

one or more DevOps teams, and the business functionality delivered is becoming very dynamic. This means that an increasing amount of security effort, including infrastructure controls and security configuration, is becoming software defined (and that software should also be secure).

- Administrators must understand the distributed nature of modern systems, create and maintain secure configurations, and practice the principle of least privilege, especially when it comes to host, network, infrastructure, and cloud services for deployed applications.
- Executives and middle management, including business owners and product managers, must understand how early investment in security design and analysis affects the degree to which users will trust their products. Business requirements should explicitly address security needs, including security-related compliance. Any sizable business today depends on software to work; thus, software security is a business necessity. Executives are also the group that must provide resources for new efforts that directly improve software security and must actively support digital transformation efforts related to infrastructure- and governance-as-code.
- GRC, legal, and data privacy specialists form an integral part of the software security effort in some firms, combining forces with security specialists when engaging with engineering. They might be responsible for analysis of contract terms, regulatory and compliance requirements including privacy regulations, definition of privacy requirements, and tracking of PII and other regulated data categories. This has become increasingly common in response to requirements such as GDPR, CCPA, and other regulations.
- Procurement and vendor management need to communicate and enforce security requirements with vendors, including those who supply on-premises products, custom software, and SaaS. Software supply chain vendors are increasingly subjected to software security SLAs and reviews (such as the PCI SSF and NIST’s SSDF) to help ensure that their products are the result of an SSDL. Of course, not all software (e.g., open source) comes from a vendor. Procurement and vendor management play a vital role but aren’t the only stakeholders responsible for managing software supply chain risk.





FIGURE 13. THE CHAMPIONS AND THE BSIMM SCORE. 92% of the top-scoring firms in the BSIMM15 data pool have a security champions. In contrast, fewer than 35% of bottom-scoring firms have one.

PART 7: **THE BSIMM** **FRAMEWORK**



PART 7: THE BSIMM FRAMEWORK

Most of the BSIMM will likely fit perfectly for your SSI, but some parts might feel a little less applicable. Understanding the model allows you to both learn from others and ensure that your program is right for your organization.

We built the first version of the BSIMM 16 years ago (late 2008) as follows:

- We relied on our own knowledge of software security practices to create the initial SSF.
- We conducted a series of in-person interviews with nine executives in charge of SSIs. From these interviews, we identified a set of 110 software security activities that we organized according to the SSF.
- We then created scorecards for each of the nine initiatives that showed which of the activities each initiative carried out. To validate our work, we asked each participating firm to review the SSF, practices, activities, and the scorecard we created for their initiative, making the necessary adjustments based on their feedback.

Today, we continue to do BSIMM assessments with in-person interviews whenever possible, which we've done with a total of 283 firms so far. Increasingly, we assessed both the SSG and one or more business units as part of creating an aggregated SSI view for a firm. We evolve the model by digging for new kinds of efforts during assessments, both as new participants join and as current participants are remeasured, then by adding new activities when warranted. We've added 19 activities and dropped one original activity that was effectively a duplicate of another since 2008 to give us a total of 128 activities today. We also adjust the positioning of activities in the model practices according to their observation rates.

ONE DROPPED ACTIVITY

In the history of BSIMM, only one activity has been dropped from the framework. Between BSIMM1 and BSIMM2, we noticed that [CR1.3] establish coding labs or office hours focused on review and [T1.3] establish SSG office hours were always scored together. Firms either got both points or neither, but neither activity required the other, and they were effectively the same. So CR1.3 was dropped in favor of T1.3.

CORE KNOWLEDGE

The BSIMM core knowledge encompasses the activities we have directly observed in BSIMM participants. We organize that core knowledge into an SSF, represented in Table 7, that is organized into four domains—Governance, Intelligence, SSDL Touchpoints, and Deployment—with those domains containing the 128 BSIMM15 activities.

From an executive perspective, you can view BSIMM activities as controls implemented in a software security risk management framework. The implemented activities might function as preventive, detective, corrective, or compensating controls in your SSI. Positioning the activities as controls allows for easier understanding of the BSIMM's value by governance, risk, compliance, legal, audit, and other risk management groups.

We divide activities into levels per practice based on the frequency with which they're observed in our participants. Doing this helps organizations quickly understand whether the activity they're contemplating is common or uncommon across other organizations. Level 1 activities (often straightforward and universally applicable) are those that are most observed across the data pool of 121 firms, level 2 activities (often more difficult to implement and requiring more coordination) are less frequently observed, and level 3 activities (usually more difficult to implement and not always applicable) are more rarely observed. Note that new activities are added at level 3 because we don't yet know how common they are, so they start with zero observations.

| DOMAINS | | | |
|--|---|---|---|
| GOVERNANCE | INTELLIGENCE | SSDL TOUCHPOINTS | DEPLOYMENT |
| Practices that help organize, manage, and measure a software security initiative. Staff development is also a central governance practice. | Practices that result in collections of corporate knowledge used in carrying out software security activities throughout the organization. Collections include both proactive security guidance and organizational threat modeling. | Practices associated with analysis and assurance of particular software development artifacts and processes. All software security methodologies include these practices. | Practices that interface with traditional network security and software maintenance organizations. Software configuration, maintenance, and other environment issues have direct impact on software security. |
| PRACTICES | | | |
| GOVERNANCE | INTELLIGENCE | SSDL TOUCHPOINTS | DEPLOYMENT |
| 1. Strategy & Metrics (SM) 2. Compliance & Policy (CP) 3. Training (T) | 4. Attack Models (AM) 5. Security Features & Design (SFD) 6. Standards & Requirements (SR) | 7. Architecture Analysis (AA) 8. Code Review (CR) 9. Security Testing (ST) | 10. Penetration Testing (PT) 11. Software Environment (SE) 12. Configuration Management & Vulnerability Management (CMVM) |

TABLE 7. THE SOFTWARE SECURITY FRAMEWORK. Twelve practices align with the four high-level domains and contain the 128 BSIMM15 activities.

UNDERSTANDING THE MODEL

A domain, such as Governance, contains practices, such as Strategy & Metrics, each of which contains activities that each have a detailed description. Creating a scorecard (e.g., activity SM1.1 was observed and is marked with a "1") informs decisions about strategic change (see Figure 14).

| GOVERNANCE | |
|-----------------------------|--|
| 1. Strategy & Metrics (SM) | |
| 2. Compliance & Policy (CP) | |
| 3. Training (T) | |

| GOVERNANCE | | | |
|--------------------|---|---------|--|
| STRATEGY & METRICS | | | |
| [SM1.1] | Publish process and evolve as necessary. | [SM2.7] | Create evangelism role and perform internal marketing. |
| [SM1.3] | Educate executives on software security. | [SM3.1] | Use a software asset tracking application with portfolio view. |
| [SM1.4] | Implement security checkpoints and associated governance. | [SM3.2] | Make SSL efforts part of external marketing. |
| [SM1.7] | Enforce security checkpoints and track exceptions. | [SM3.3] | Identify metrics and use them to drive resourcing. |
| [SM2.1] | Publish data about software security internally and use it to drive change. | [SM3.4] | Integrate software-defined lifecycle governance. |
| [SM2.3] | Create or grow a security champions program. | [SM3.5] | Integrate software supply chain risk management. |
| [SM2.6] | Require security sign-off prior to software release. | | |

[SM2.7: 62] CREATE EVANGELISM ROLE AND PERFORM INTERNAL MARKETING.

Build support for software security throughout the organization via ongoing evangelism and ensure that everyone aligns on security objectives. This internal marketing function, often performed by a variety of stakeholder roles, keeps executives and others up to date on the magnitude of the software security problem and the elements of its solution. A champion or a scrum master familiar with security, for example, could help teams adopt better software security practices as they transform to Agile and DevOps methods. Similarly, a cloud expert could demonstrate the changes needed in security architecture and testing for serverless applications. Evangelists can increase understanding and build credibility by giving talks to internal groups (including executives), publishing roadmaps, authoring technical papers for internal consumption, or creating a collection of papers, books, and other resources on an internal website (see [SR1.2]) and promoting its use. In turn, organizational feedback becomes a useful source of improvement ideas.

| GOVERNANCE | | |
|--------------------|------------------------------|--------------|
| ACTIVITY | "BSIMMI5 FIRMS (OUT OF 121)" | EXAMPLE FIRM |
| STRATEGY & METRICS | | |
| [SM1.1] | 90 | 1 |
| [SM1.3] | 72 | |
| [SM1.4] | 109 | |
| [SM1.7] | 72 | |
| [SM2.1] | 65 | |
| [SM2.3] | 67 | |
| [SM2.6] | 69 | |
| [SM2.7] | 52 | 1 |
| [SM3.1] | 30 | |
| [SM3.2] | 23 | |
| [SM3.3] | 32 | |
| [SM3.4] | 11 | |
| [SM3.5] | 1 | |

FIGURE 14. THE PIECES OF THE BSIMM MODEL.

DETAILED VIEW OF THE BSIMM FRAMEWORK

The BSIMM framework and data model evolve over time to accurately represent actual software security practices. Understanding these changes will help you set strategic directions for your own SSI.

Here, we explore the BSIMM framework in more detail, including the methodology of how we created the model, how it evolved over time, and how we updated it for BSIMM15.

As a descriptive model, the only goal of the BSIMM is to observe and report. We like to say we visited many restaurants to see what was happening and observed that “there are three chicken eggs in an omelet.” Note that the BSIMM does not extrapolate to say, “all omelets must have three eggs,” “only chicken eggs make acceptable omelets,” “omelets must be eaten every day,” or any other value judgments. We offer simple observations, simply reported.


Of course, during our assessment efforts across hundreds of organizations, we also make qualitative observations about how SSIs are evolving and report many of those as trends, insights, analysis, and other topical discussions both in this document and among BSIMM participants.

Our “just the facts” approach is hardly novel in science and engineering, but in the realm of software security, it has not previously been applied at this scale. Other work around SSI modeling has either described the experience of a single organization or offered prescriptive guidance based on a combination of personal experience and opinion.

During our assessment efforts across hundreds of organizations, we make qualitative observations about how SSIs are evolving and report many of those as insights, analysis, and other discussions in this document.

THE DETAILED BSIMM SKELETON

The detailed BSIMM skeleton provides a way to view the model at a glance and is useful when assessing an SSI. The skeleton is shown in Figure 15 and contains the activity references organized by levels, domains, and practices. More complete descriptions of the activities and examples are available in Part 8 of this document.



| GOVERNANCE | | | | | |
|--------------------|---|---------------------|---|----------|---|
| STRATEGY & METRICS | | COMPLIANCE & POLICY | | TRAINING | |
| [SM1.1] | Publish process and evolve as necessary. | [CP1.1] | Unify regulatory pressures. | [T1.1] | Conduct software security awareness training. |
| [SM1.3] | Educate executives on software security. | [CP1.2] | Identify privacy obligations. | [T1.7] | Deliver on-demand individual training. |
| [SM1.4] | Implement security checkpoints and associated governance. | [CP1.3] | Create policy. | [T1.8] | Include security resources in onboarding. |
| [SM1.7] | Enforce security checkpoints and track exceptions. | [CP2.1] | Build a PII inventory. | [T2.5] | Enhance security champions through training and events. |
| [SM2.1] | Publish data about software security internally and use it to drive change. | [CP2.2] | Require security sign-off for compliance-related risk. | [T2.8] | Create and use material specific to company history. |
| [SM2.3] | Create or grow a security champions program. | [CP2.3] | Implement and track controls for compliance. | [T2.9] | Deliver role-specific advanced curriculum. |
| [SM2.6] | Require security sign-off prior to software release. | [CP2.4] | Include software security SLAs in all vendor contracts. | [T2.10] | Host software security events. |
| [SM2.7] | Create evangelism role and perform internal marketing. | [CP2.5] | Ensure executive awareness of compliance and privacy obligations. | [T2.11] | Require an annual refresher. |
| [SM3.1] | Use a software asset tracking application with portfolio view. | [CP3.1] | Document a software compliance story. | [T2.12] | Provide expertise via open collaboration channels. |
| [SM3.2] | Make SSI efforts part of external marketing. | [CP3.2] | Ensure compatible vendor policies. | [T3.1] | Reward progression through curriculum. |
| [SM3.3] | Identify metrics and use them to drive resourcing. | [CP3.3] | Drive feedback from software lifecycle data back to policy. | [T3.2] | Provide training for vendors and outsourced workers. |
| [SM3.4] | Integrate software-defined lifecycle governance. | | | [T3.6] | Identify new security champions through observation. |
| [SM3.5] | Integrate software supply chain risk management. | | | | |

| INTELLIGENCE | | | | | |
|---------------|--|----------------------------|---|--------------------------|--|
| ATTACK MODELS | | SECURITY FEATURES & DESIGN | | STANDARDS & REQUIREMENTS | |
| [AM1.2] | Use a data classification scheme for software inventory. | [SFD1.1] | Integrate and deliver security features. | [SR1.1] | Create security standards. |
| [AM1.3] | Identify potential attackers. | [SFD1.2] | Application architecture teams engage with the SSG. | [SR1.2] | Create a security portal. |
| [AM1.5] | Gather and use attack intelligence. | [SFD2.1] | Leverage secure-by-design components and services. | [SR1.3] | Translate compliance constraints to requirements. |
| [AM2.1] | Build attack patterns and abuse cases tied to potential attackers. | [SFD2.2] | Create capability to solve difficult design problems. | [SR1.5] | Identify open source. |
| [AM2.6] | Collect and publish attack stories. | [SFD3.1] | Form a review board to approve and maintain secure design patterns. | [SR2.2] | Create a standards review process. |
| [AM2.7] | Build an internal forum to discuss attacks. | [SFD3.2] | Require use of approved security features and frameworks. | [SR2.5] | Create SLA boilerplate. |
| [AM2.8] | Have a research group that develops new attack methods. | [SFD3.3] | Find and publish secure design patterns from the organization. | [SR2.7] | Control open source risk. |
| [AM2.9] | Monitor automated asset creation. | | | [SR3.2] | Communicate standards to vendors. |
| [AM3.2] | Create and use automation to mimic attackers. | | | [SR3.3] | Use secure coding standards. |
| [AM3.4] | Create technology-specific attack patterns. | | | [SR3.4] | Create standards for technology stacks. |
| [AM3.5] | Maintain and use a top N possible attacks list. | | | [SR3.5] | Create standards controlling and guiding the adoption of new technologies. |

| SSDL TOUCHPOINTS | | | | | |
|-----------------------|--|-------------|---|------------------|--|
| ARCHITECTURE ANALYSIS | | CODE REVIEW | | SECURITY TESTING | |
| [AA1.1] | Perform security feature review. | [CR1.2] | Perform opportunistic code review. | [ST1.1] | Perform edge/boundary value condition testing during QA. |
| [AA1.2] | Perform design review for high-risk applications. | [CR1.4] | Use automated code review tools. | [ST1.3] | Drive tests with security requirements and security features. |
| [AA1.4] | Use a risk methodology to rank applications. | [CR1.5] | Make code review mandatory for all projects. | [ST1.4] | Integrate opaque-box security tools into the QA process. |
| [AA2.1] | Perform architecture analysis using a defined process. | [CR1.7] | Assign code review tool mentors. | [ST2.4] | Drive QA tests with AST results. |
| [AA2.2] | Standardize architectural descriptions. | [CR2.6] | Use custom rules with automated code review tools. | [ST2.5] | Include security tests in QA automation. |
| [AA2.4] | Have SSG lead design review efforts. | [CR2.7] | Use a top N bugs list (real data preferred). | [ST2.6] | Perform fuzz testing customized to application APIs. |
| [AA3.1] | Have engineering teams lead AA process. | [CR2.8] | Use centralized defect reporting to close the knowledge loop. | [ST3.3] | Drive tests with design review results. |
| [AA3.2] | Drive analysis results into standard design patterns. | [CR3.2] | Build a capability to combine AST results. | [ST3.4] | Leverage code coverage analysis. |
| [AA3.3] | Make the SSG available as an AA resource or mentor. | [CR3.3] | Create capability to eradicate bugs. | [ST3.5] | Begin to build and apply adversarial security tests (abuse cases). |
| | | [CR3.4] | Automate malicious code detection. | [ST3.6] | Implement event-driven security testing in automation. |
| | | [CR3.5] | Enforce secure coding standards. | | |

| DEPLOYMENT | | | | | |
|---------------------|---|----------------------|--|---|---|
| PENETRATION TESTING | | SOFTWARE ENVIRONMENT | | CONFIGURATION MANAGEMENT & VULNERABILITY MANAGEMENT | |
| [PT1.1] | Use external penetration testers to find problems. | [SE1.1] | Use application input monitoring for security purposes. | [CMVM1.1] | Create or interface with incident response. |
| [PT1.2] | Feed results to the defect management and mitigation system. | [SE1.2] | Ensure host and network security basics are in place. | [CMVM1.2] | Identify software defects found in operations monitoring and feed them back to engineering. |
| [PT1.3] | Use penetration testing tools internally. | [SE1.3] | Implement cloud security controls. | [CMVM1.3] | Track software defects found in operations through the fix process. |
| [PT2.2] | Penetration testers use all available information. | [SE1.4] | Define secure deployment parameters and configurations. | [CMVM1.4] | Have emergency response. |
| [PT2.3] | Schedule periodic penetration tests for application coverage. | [SE2.4] | Protect code integrity. | [CMVM2.3] | Develop an operations software inventory. |
| [PT3.1] | Use external penetration testers to perform deep-dive analysis. | [SE2.5] | Use application containers to support security goals. | [CMVM2.4] | Streamline incoming responsible vulnerability disclosure. |
| [PT3.2] | Customize penetration testing tools. | [SE2.7] | Use orchestration for containers and virtualized environments. | [CMVM3.1] | Fix all occurrences of software defects found in operations. |
| | | [SE3.2] | Use code protection. | [CMVM3.2] | Enhance the SSDL to prevent software defects found in operations. |
| | | [SE3.3] | Use application behavior monitoring and diagnostics. | [CMVM3.3] | Simulate software crises. |
| | | [SE3.6] | Create bills of materials for deployed software. | [CMVM3.4] | Operate a bug bounty program. |
| | | [SE3.8] | Perform application composition analysis on code repositories. | [CMVM3.5] | Automate verification of operational infrastructure security. |
| | | [SE3.9] | Protect integrity of development toolchains. | [CMVM3.6] | Publish risk data for deployable artifacts. |
| | | [SE3.10] | Protect the integrity of development endpoints. | [CMVM3.8] | Do attack surface management for deployed applications. |

FIGURE 15. THE DETAILED BSIMM SKELETON. Within the SSF, the 128 activities are organized into three levels, across 12 practices, and within the four BSIMM domains.

CREATING BSIMM15 FROM BSIMM14

BSIMM15 includes updated activity descriptions, data from firms in multiple vertical markets, and a longitudinal study. For BSIMM15, we added nine firms and removed 18, resulting in a data pool of 121 firms. In addition, in the time since we launched BSIMM14, six firms conducted reassessments to update their scorecards, and we assessed 14 additional business units for six firms. (These business unit assessments do not appear in the BSIMM15 data pool.)

As shown below, we used the resulting observation counts to refine activity placement in the framework, which resulted in moving four activities to different levels. In addition, we added two newly observed activities, resulting in a total of 128 activities in BSIMM15:

- [SM2.2] **Enforce security checkpoints and track exceptions** became [SM1.7]
- [SE2.2] **Define secure deployment parameters and configurations** became [SE1.4]
- [CMVM2.1] **Have emergency response** became [CMVM1.4]
- [CMVM3.7] **Streamline incoming responsible vulnerability disclosure** became [CMVM2.4]
- [SR3.5] **Create standards controlling and guiding the adoption of new technologies** was added to the model
- [SE3.10] **Protect the integrity of development endpoints** was added to the model

As an example of how the BSIMM functions as an observational model, consider the activities that are now [AM3.5] **maintain and use a top N possible attacks list** and [T3.6] **identify new security champions through observation**, both of which started as level 1 activities. The BSIMM1 activity [AM1.1] **build and maintain a top N possible attacks list** became AM2.5 in BSIMM7 and then [AM3.5] in BSIMM14 as observation rates declined relative to other Attack Model activities. [T1.4] **identify new security champions through observation**, also a BSIMM1 activity, became [T2.7] in BSIMM4 and [T3.6] in BSIMM8 as organizations adopted other ways of identifying new security champion candidates. [T3.6] also demonstrates how activities evolve over time, changing from **identify satellite through training** to **identify new satellite members through observation** with BSIMM11, **identify new satellite members (security champions) through observation** with BSIMM13, and finally **identify new security champions through observation** to reflect both where the identifying was occurring as well as more common terminology in the industry.

In BSIMM13, we had the first activity that migrated from level 3 to level 1—[SE1.3] **implement cloud security controls**, which was introduced in BSIMM9. While the relative growth of [SE2.5] **use application containers to support security goals** has slowed down, it is one of the potential candidates to migrate from level 3 to level 1 over the next couple of years. See Table 8 for the observation growth in activities that were added since BSIMM7.

| OBSERVATIONS | | | | | | | | | |
|-----------------------|--------|--------|--------|---------|---------|---------|---------|---------|---------|
| ACTIVITY | BSIMM7 | BSIMM8 | BSIMM9 | BSIMM10 | BSIMM11 | BSIMM12 | BSIMM13 | BSIMM14 | BSIMM15 |
| SE3.4 (now SE2.5) | 0 | 4 | 11 | 14 | 31 | 44 | 52 | 63 | 64 |
| SE3.5 (now SE2.7) | | | 0 | 5 | 22 | 33 | 42 | 47 | 42 |
| SE3.6 | | | 0 | 3 | 12 | 14 | 18 | 22 | 25 |
| SE3.7 (now SE1.3) | | | 0 | 9 | 36 | 59 | 79 | 92 | 87 |
| SM3.4 | | | | 0 | 1 | 6 | 5 | 8 | 11 |
| AM3.3 (now AM2.9) | | | | 0 | 4 | 6 | 11 | 17 | 18 |
| CMVM3.5 | | | | 0 | 8 | 10 | 13 | 16 | 17 |
| ST3.6 | | | | | 0 | 2 | 3 | 6 | 8 |
| CMVM3.6 | | | | | 0 | 0 | 3 | 3 | 4 |
| CMVM3.7 (now CMVM2.4) | | | | | | 0 | 20 | 35 | 41 |
| SM3.5 | | | | | | | 0 | 0 | 1 |
| SE3.8 | | | | | | | 0 | 2 | 3 |
| CMVM3.8 | | | | | | | 0 | 0 | 1 |
| SE3.9 | | | | | | | | 0 | 3 |
| SR3.5 | | | | | | | | | 0 |
| SE3.10 | | | | | | | | | 0 |

TABLE 8. NEW ACTIVITIES. Some activities have seen exceptional growth (highlighted in orange) in observation counts, likely demonstrating their widespread utility. [SE3.7], highlighted in gray, is the first activity to migrate from level 3 (very uncommon) to level 1 (common).

MODEL CHANGES OVER TIME

WHERE DO OLD ACTIVITIES GO?

We continue to ponder the question, “Where do activities go when no one does them anymore?” We’ve noticed that the observation rate for other seemingly useful activities has decreased significantly in recent years:

- [CR3.5] **Enforce secure coding standards** declined from a high of 55.6% of firms in BSIMM1 to a low of 0% of firms in BSIMM12, although it has climbed back to 5% with BSIMM15
- [T3.6] **Identify new security champions through observation** observed in 11 of 51 firms in BSIMM4 but only in nine of 121 firms in BSIMM15
- [SFD3.3] **Find and publish secure design patterns from the organization** observed in 14 of 51 firms in BSIMM4 but only in 12 of 121 firms in BSIMM15
- [SR3.3] **Use secure coding standards** observed in 23 of 78 firms in BSIMM6 but only in 17 of 121 firms in BSIMM15

We believe there are two primary reasons why observations for some activities have decreased toward zero over time. First, some activities have become part of the culture and drive different behavior—for example, choosing security champions members might become a more organic part of the SSDL without requiring extra effort in identifying champions members ([T3.6] **identify new security champions through observation**) to grow that team ([SM2.3] **create or**

grow a security champions program). Second, some activities don’t yet fit tightly with the evolving engineering culture, and the activity effort currently causes too much friction. For example, continuously going to engineering teams to find secure design patterns, [SFD3.3] **find and publish secure design patterns from the organization**, might unacceptably delay key development processes.

It might also be the case that evolving SSI and DevOps architectures are changing the way some activities are getting done. If an organization’s use of purpose-built architectures, development kits, and libraries is sufficiently consistent, perhaps it’s less necessary to lean on prescriptive coding standards ([CR3.5] **enforce secure coding standards**) as a measure of acceptable code.

As a point of culture-driven contrast, we see significant increases in observation counts for activities such as [SE1.3] **implement cloud security controls**, [SE2.5] **use application containers to support security goals**, and [SE2.7] **Use orchestration for containers and virtualized environments**, likely for similar reasons that we see lower counts for the other activities above. The engineering culture has shifted to be more self-service and to include increased telemetry that produces more data for everyone to use. We keep a close watch on the BSIMM data pool and will make adjustments as needed, which might include dropping an activity from the model.

Being a unique, real-world reflection of actual software security practices, the BSIMM naturally changes over time. While each release of the BSIMM captures the current dataset and provides the most current guidance, reflection upon past changes can help clarify the ebb and flow of specific activities. Table 9 shows the activity moves, adds, and deletes that have occurred since the BSIMM’s creation.

| | |
|---|---|
| CHANGES FOR BSIMM15 (128 ACTIVITIES) | <ul style="list-style-type: none"> • [SM2.2] Enforce security checkpoints and track exceptions became [SM1.7] • [SE2.2] Define secure deployment parameters and configurations became [SE1.4] • [CMVM2.1] Have emergency response became [CMVM1.4] • [CMVM3.7] Streamline incoming responsible vulnerability disclosure became [CMVM2.4] • [SR3.5] Create standards controlling and guiding the adoption of new technologies was added to the model • [SE3.10] Protect the integrity of development endpoints was added to the model |
| CHANGES FOR BSIMM14 (126 ACTIVITIES) | <ul style="list-style-type: none"> • [T3.5] Provide expertise via open collaboration channels became [T2.12] • [AM2.2] Create technology-specific attack patterns became [AM3.4] • [AM2.5] Maintain and use a top N possible attacks list became [AM3.5] • [AM3.1] Have a research group that develops new attack methods became [AM2.8] • [AM3.3] Monitor automated asset creation became [AM2.9] • [SR2.4] Identify open source became [SR1.5] • [CMVM2.2] Track software bugs found in operations through the fix process became [CMVM1.3] • [SE3.9] Protect integrity of SDLC toolchains was added to the model |
| CHANGES FOR BSIMM13 (125 ACTIVITIES) | <ul style="list-style-type: none"> • [T3.3] Host software security events became [T2.10] • [T3.4] Require an annual refresher became [T2.11] • [SR3.1] Control open source risk became [SR2.7] • [AA1.3] Have SSG lead design review efforts became [AA2.4] • [CR1.6] Use centralized defect reporting to close the knowledge loop became [CR2.8] • [SE2.6] Implement cloud security controls became [SE1.3] • [SM3.5] Integrate software supply chain risk management added to the model • [SE3.8] Perform application composition analysis on code repositories added to the model • [CMVM3.8] Do attack surface management for deployed applications added to the model |
| CHANGES FOR BSIMM12 (122 ACTIVITIES) | <ul style="list-style-type: none"> • [SM1.2] Create evangelism role and perform internal marketing became [SM2.7] • [T1.5] Deliver role-specific advanced curriculum became [T2.9] • [ST2.1] Integrate black-box security tools into the QA process became [ST1.4] • [SE3.5] Use orchestration for containers and virtualized environments became [SE2.7] • [CMVM3.7] Streamline incoming responsible vulnerability disclosure added to the model |

| | |
|---|--|
| CHANGES FOR BSIMM11 (121 ACTIVITIES) | <ul style="list-style-type: none"> • [T2.6] Include security resources in onboarding became [T1.8] • [CR2.5] Assign tool mentors became [CR1.7] • [SE3.4] Use application containers to support security goals became [SE2.5] • [SE3.7] Ensure cloud security basics became [SE2.6] • [ST3.6] Implement event-driven security testing in automation added to the model • [CMVM3.6] Publish risk data for deployable artifacts added to the model |
| CHANGES FOR BSIMM10 (119 ACTIVITIES) | <ul style="list-style-type: none"> • [T1.6] Create and use material specific to company history became [T2.8] • [SR2.3] Create standards for technology stacks moves to become [SR3.4] • [SM3.4] Integrate software-defined lifecycle governance added to the model • [AM3.3] Monitor automated asset creation added to the model • [CMVM3.5] Automate verification of operational infrastructure security added to the model |
| CHANGES FOR BSIMM9 (116 ACTIVITIES) | <ul style="list-style-type: none"> • [SM2.5] Identify metrics and use them to drive resourcing became [SM3.3] • [SR2.6] Use secure coding standards became [SR3.3] • [SE3.5] Use orchestration for containers and virtualized environments added to the model • [SE3.6] Enhance application inventory with operations bill of materials added to the model • [SE3.7] Ensure cloud security basics added to the model |
| CHANGES FOR BSIMM8 (113 ACTIVITIES) | <ul style="list-style-type: none"> • [T2.7] Identify new satellite through training became [T3.6] • [AA2.3] Make SSG available as AA resource or mentor became [AA3.3] |
| CHANGES FOR BSIMM7 (113 ACTIVITIES) | <ul style="list-style-type: none"> • [AM1.1] Maintain and use a top N possible attacks list became [AM2.5] • [AM1.4] Collect and publish attack stories became [AM2.6] • [AM1.6] Build an internal forum to discuss attacks became [AM2.7] • [CR1.1] Use a top N bugs list became [CR2.7] • [CR2.2] Enforce coding standards became [CR3.5] • [SE3.4] Use application containers to support security goals added to model |
| CHANGES FOR BSIMM6 (112 ACTIVITIES) | <ul style="list-style-type: none"> • [SM1.6] Require security sign-off prior to software release became [SM2.6] • [SR1.4] Use secure coding standards became [SR2.6] • [ST3.1] Include security tests in QA automation became [ST2.5] • [ST3.2] Perform fuzz testing customized to application APIs became [ST2.6] |

| | |
|---|---|
| CHANGES FOR BSIMM-V (112 ACTIVITIES) | <ul style="list-style-type: none"> • [SFD2.3] Find and publish mature design patterns from the organization became [SFD3.3] • [SR2.1] Communicate standards to vendors became [SR3.2] • [CR3.1] Use automated tools with tailored rules became [CR2.6] • [ST2.3] Begin to build and apply adversarial security tests (abuse cases) became [ST3.5] • [CMVM3.4] Operate a bug bounty program added to model |
| CHANGES FOR BSIMM4 (111 ACTIVITIES) | <ul style="list-style-type: none"> • [T2.1] Deliver role-specific advanced curriculum became [T1.5] • [T2.2] Company history in training became [T1.6] • [T2.4] Deliver on-demand individual training became [T1.7] • [T1.2] Include security resources in onboarding became [T2.6] • [T1.4] Identify new satellite members through training became [T2.7] • [T1.3] Establish SSG office hours became [T3.5] • [AM2.4] Build an internal forum to discuss attacks became [AM1.6] • [CR2.3] Make code review mandatory for all projects became [CR1.5] • [CR2.4] Use centralized reporting to close the knowledge loop became [CR1.6] • [ST1.2] Share security results with QA became [ST2.4] • [SE2.3] Use application behavior monitoring and diagnostics became [SE3.3] • [CR3.4] Automate malicious code detection added to model • [CMVM3.3] Simulate software crises added to model |
| CHANGES FOR BSIMM3 (109 ACTIVITIES) | <ul style="list-style-type: none"> • [SM1.7] Identify metrics and use them to drive resourcing became [SM2.5] • [SM2.4] Require security sign-off became [SM1.6] • [AM2.3] Gather and use attack intelligence became [AM1.5] • [ST2.2] Drive tests with security requirements and security features became [ST1.3] • [PT2.1] Use pen testing tools internally became [PT1.3] |
| CHANGES FOR BSIMM2 (109 ACTIVITIES) | <ul style="list-style-type: none"> • [T2.3] Require an annual refresher became [T3.4] • [CR2.1] Use automated tools became [CR1.4] • [SE2.1] Use code protection became [SE3.2] • [SE3.1] Use code signing became [SE2.4] • [CR1.3] removed from the model |
| CHANGES FOR BSIMM1 (110 ACTIVITIES) | <ul style="list-style-type: none"> • Added 110 activities |

TABLE 9. ACTIVITY CHANGES OVER TIME. This table allows for historical review of how BSIMM activities have been added, moved, and deleted since inception.

DATA: BSIMM15

Every organization wants to do software security more effectively and efficiently. You can use this information to understand what BSIMM participants are doing today and how those efforts have evolved over time, then plan your own SSI changes.

The BSIMM data yields very interesting analytical results, as shown throughout this document. Figure 18 shows the highest-resolution observation data that is published, and organizations can use this information to note how often we observe each activity across all 121 participants to help plan their next areas of focus. Activities that are broadly popular will likely benefit your organization as well.

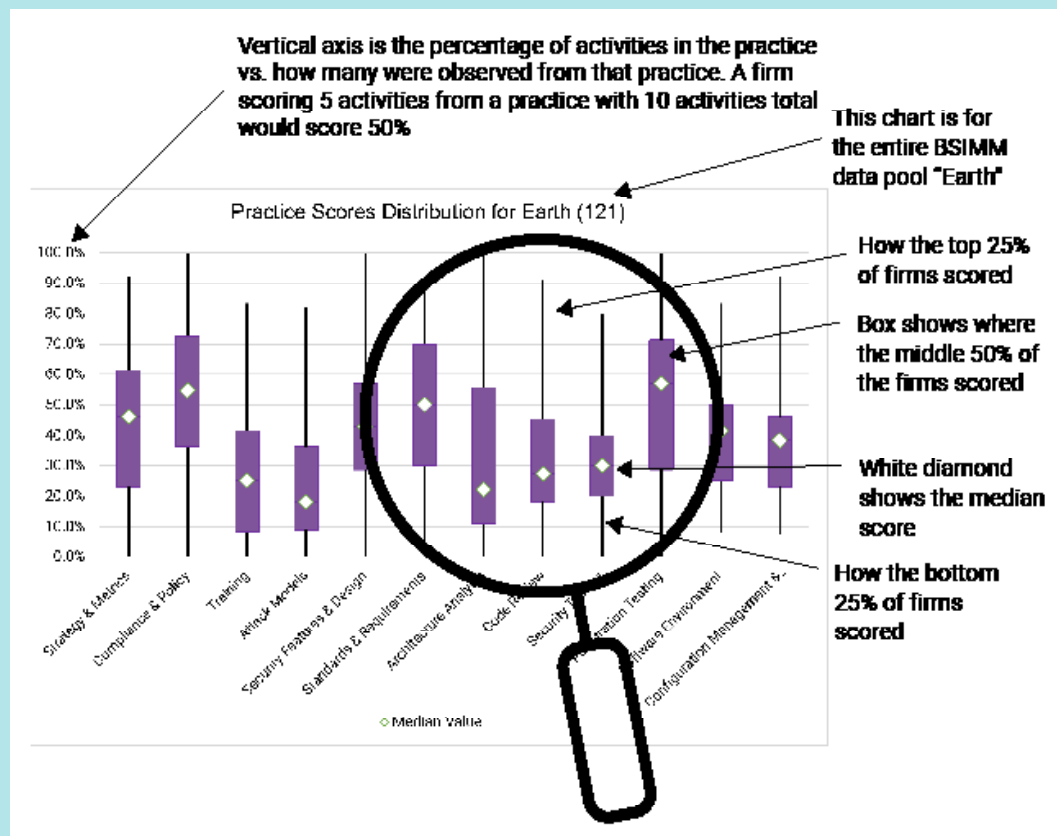
In Figure 18, we also identified the most common activity in each practice (highlighted in orange). To provide some perspective on what “most common” means, although [T1.1] is the most common activity in the Training practice with 62 observations, Table 11 shows that it isn't in the top 20 activities across all practices.

To provide another view into this data, we created a boxplot chart by taking the percentage of activities observed in each practice and showing the distribution of those percentages for all 121 firms in the data pool (see Figure 16). Looking at the score distributions provides more data than the simple average we used to use, for example, four practices see some companies doing 100% of the activities, Compliance & Policy, Security Features & Design, Architecture Analysis, and Penetration Testing, while the interquartile range (the middle 50% of the scores) are significantly lower for Security Features & Design and Architecture Analysis than they are for the other two. Only two practices, Software Environment and Configuration Management & Vulnerability Management, show all 121 firms doing at least something in those practices. The other 10 practices all have at least one firm doing nothing in that practice. Looking at the median scores, Compliance & Policy and Penetration Testing have the highest median scores while Attack Models the lowest.

The range of observed scores in the current data pool is 12 for the lower score and 100 for the higher score, indicating a wide range of SSI maturity levels in the BSIMM15 data.

READING A BOXPLOT

A boxplot, also known as a box and whisker plot, provides a visual summary of a range of data. It shows the minimum and maximum, the 1st quartile (bottom 25% of the values), 3rd quartile (top 25% of the values), the interquartile range (the middle 50% of the values), and the median.



25% of the values will fall between the 1st quartile and the minimum, 25% of the values will fall between the 3rd quartile and the maximum, and 50% will fall between the 1st and 3rd quartiles. The median is the middle value in a set of numbers when they are sorted in ascending or descending order.

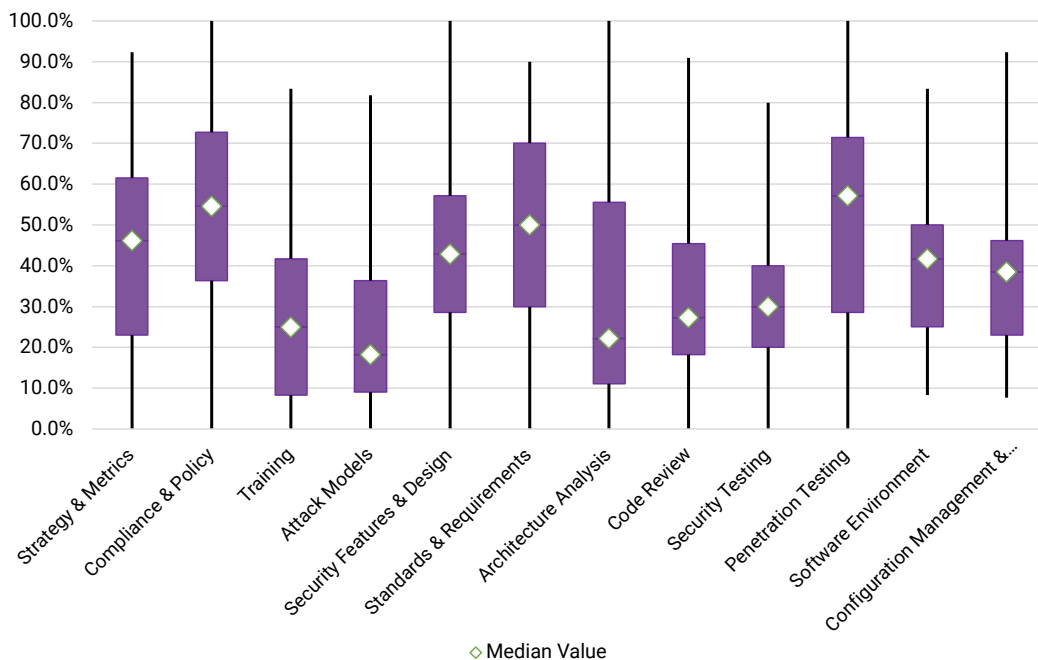


FIGURE 16. ALL FIRMS SCORE DISTRIBUTION. This diagram shows the range of the normalized observations collectively reached in each practice by the 121 BSIMM15 firms. Across these firms, the normalized observations are higher in the Compliance & Policy, Standards & Requirements, and Penetration Testing practices compared to Training, Attack Models, and Security Testing.

AGE-BASED PROGRAM CHANGES

Figure 17 shows the distribution of scores of all 121 participating firms. To create this graph, we divided the scores into eight bins that are then further divided by only one BSIMM assessment having been done (iteration 1), one reassessment (iteration 2), and two or more reassessments done (iteration 3+). We also plotted the average age of the firms' SSIs in each bin as a horizontal line. In general, firms where more BSIMM activities were observed have older SSIs and are more likely to have performed multiple BSIMM measurements.

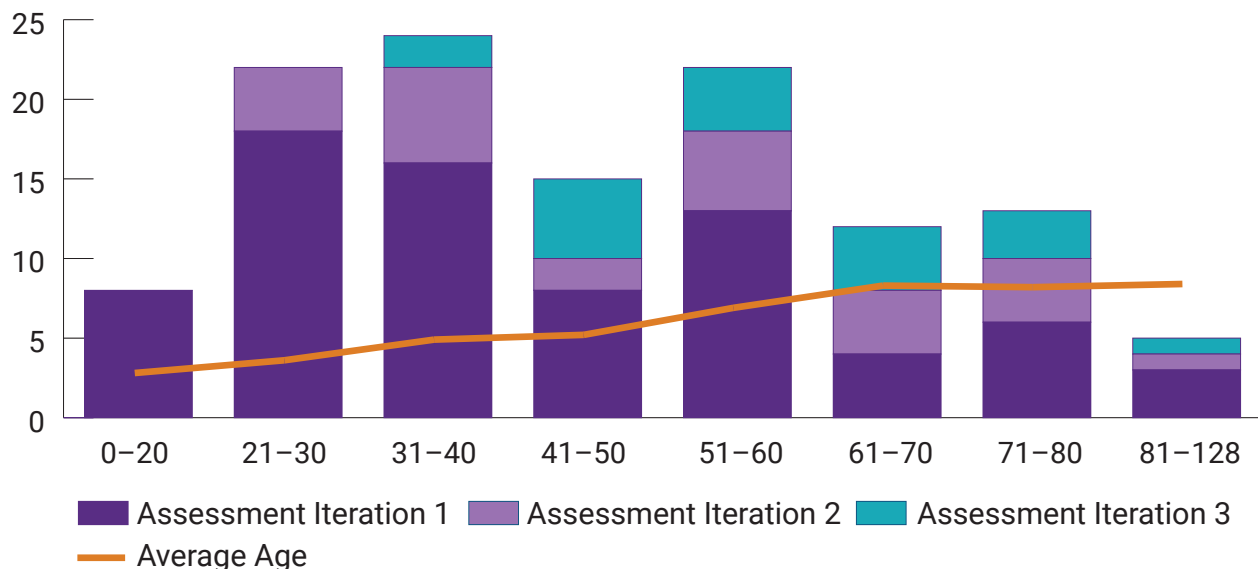


FIGURE 17. BSIMM SCORE DISTRIBUTION. Assessment scores most frequently fall into the 31 to 40 range in BSIMM15 with an average SSG age of 4.6 years. In general, firms that mature and continue to use the BSIMM as a measurement tool over time (e.g., iteration 2, iteration 3+) tend to have higher scores. Refer to DATA ANALYSIS: LONGITUDINAL in Part 9 for more details on how SSIs evolve over multiple measurements.

| GOVERNANCE | | | INTELLIGENCE | | | SSDL TOUCHPOINTS | | | DEPLOYMENT | | |
|---------------------|----------------------------|----------------------------|----------------------------|----------------------------|----------------------------|-----------------------|----------------------------|----------------------------|-----------------------------|----------------------------|----------------------------|
| ACTIVITY | BSIMM15 FIRMS (OUT OF 121) | BSIMM15 FIRMS (PERCENTAGE) | ACTIVITY | BSIMM15 FIRMS (OUT OF 121) | BSIMM15 FIRMS (PERCENTAGE) | ACTIVITY | BSIMM15 FIRMS (OUT OF 121) | BSIMM15 FIRMS (PERCENTAGE) | ACTIVITY | BSIMM15 FIRMS (OUT OF 121) | BSIMM15 FIRMS (PERCENTAGE) |
| STRATEGY & METRICS | | | ATTACK MODELS | | | ARCHITECTURE ANALYSIS | | | PENETRATION TESTING | | |
| [SM1.1] | 90 | 74.38% | [AM1.2] | 64 | 52.89% | [AA1.1] | 99 | 81.82% | [PT1.1] | 104 | 85.95% |
| [SM1.3] | 72 | 59.50% | [AM1.3] | 49 | 40.50% | [AA1.2] | 56 | 46.28% | [PT1.2] | 95 | 78.51% |
| [SM1.4] | 109 | 90.08% | [AM1.5] | 77 | 63.64% | [AA1.4] | 55 | 45.45% | [PT1.3] | 76 | 62.81% |
| [SM1.7] | 72 | 59.50% | [AM2.1] | 18 | 14.88% | [AA2.1] | 37 | 30.58% | [PT2.2] | 39 | 32.23% |
| [SM2.1] | 65 | 53.72% | [AM2.6] | 12 | 9.92% | [AA2.2] | 38 | 31.40% | [PT2.3] | 51 | 42.15% |
| [SM2.3] | 67 | 55.37% | [AM2.7] | 16 | 13.22% | [AA2.4] | 40 | 33.06% | [PT3.1] | 26 | 21.49% |
| [SM2.6] | 69 | 57.02% | [AM2.8] | 24 | 19.83% | [AA3.1] | 20 | 16.53% | [PT3.2] | 21 | 17.36% |
| [SM2.7] | 52 | 42.98% | [AM2.9] | 18 | 14.88% | [AA3.2] | 8 | 6.61% | | | |
| [SM3.1] | 30 | 24.79% | [AM3.2] | 8 | 6.61% | [AA3.3] | 18 | 14.88% | | | |
| [SM3.2] | 23 | 19.01% | [AM3.4] | 11 | 9.09% | | | | | | |
| [SM3.3] | 32 | 26.45% | [AM3.5] | 10 | 8.26% | | | | | | |
| [SM3.4] | 11 | 9.09% | | | | | | | | | |
| [SM3.5] | 1 | 0.83% | | | | | | | | | |
| COMPLIANCE & POLICY | | | SECURITY FEATURES & DESIGN | | | CODE REVIEW | | | SOFTWARE ENVIRONMENT | | |
| [CP1.1] | 98 | 80.99% | [SFD1.1] | 93 | 76.86% | [CR1.2] | 80 | 66.12% | [SE1.1] | 76 | 62.81% |
| [CP1.2] | 105 | 86.78% | [SFD1.2] | 83 | 68.60% | [CR1.4] | 106 | 87.60% | [SE1.2] | 102 | 84.30% |
| [CP1.3] | 94 | 77.69% | [SFD2.1] | 42 | 34.71% | [CR1.5] | 75 | 61.98% | [SE1.3] | 87 | 71.90% |
| [CP2.1] | 49 | 40.50% | [SFD2.2] | 68 | 56.20% | [CR1.7] | 51 | 42.15% | [SE1.4] | 74 | 61.16% |
| [CP2.2] | 58 | 47.93% | [SFD3.1] | 18 | 14.88% | [CR2.6] | 24 | 19.83% | [SE2.4] | 51 | 42.15% |
| [CP2.3] | 69 | 57.02% | [SFD3.2] | 21 | 17.36% | [CR2.7] | 19 | 15.70% | [SE2.5] | 64 | 52.89% |
| [CP2.4] | 60 | 49.59% | [SFD3.3] | 12 | 9.92% | [CR2.8] | 27 | 22.31% | [SE2.7] | 42 | 34.71% |
| [CP2.5] | 70 | 57.85% | | | | [CR3.2] | 16 | 13.22% | [SE3.2] | 24 | 19.83% |
| [CP3.1] | 36 | 29.75% | | | | [CR3.3] | 6 | 4.96% | [SE3.3] | 17 | 14.05% |
| [CP3.2] | 39 | 32.23% | | | | [CR3.4] | 3 | 2.48% | [SE3.6] | 25 | 20.66% |
| [CP3.3] | 13 | 10.74% | | | | [CR3.5] | 6 | 4.96% | [SE3.8] | 3 | 2.48% |
| | | | | | | | | | [SE3.9] | 3 | 2.48% |
| | | | | | | | | | [SE3.10] | 0 | 0.00% |
| TRAINING | | | STANDARDS & REQUIREMENTS | | | SECURITY TESTING | | | CONFIG. MGMT. & VULN. MGMT. | | |
| [T1.1] | 62 | 51.24% | [SR1.1] | 84 | 69.42% | [ST1.1] | 102 | 84.30% | [CMVM1.1] | 111 | 91.74% |
| [T1.7] | 61 | 50.41% | [SR1.2] | 96 | 79.34% | [ST1.3] | 79 | 65.29% | [CMVM1.2] | 85 | 70.25% |
| [T1.8] | 50 | 41.32% | [SR1.3] | 86 | 71.07% | [ST1.4] | 54 | 44.63% | [CMVM1.3] | 89 | 73.55% |
| [T2.5] | 41 | 33.88% | [SR1.5] | 96 | 79.34% | [ST2.4] | 20 | 16.53% | [CMVM1.4] | 88 | 72.73% |
| [T2.8] | 25 | 20.66% | [SR2.2] | 70 | 57.85% | [ST2.5] | 34 | 28.10% | [CMVM2.3] | 46 | 38.02% |
| [T2.9] | 31 | 25.62% | [SR2.5] | 62 | 51.24% | [ST2.6] | 24 | 19.83% | [CMVM2.4] | 41 | 33.88% |
| [T2.10] | 25 | 20.66% | [SR2.7] | 55 | 45.45% | [ST3.3] | 16 | 13.22% | [CMVM3.1] | 13 | 10.74% |
| [T2.11] | 29 | 23.97% | [SR3.2] | 19 | 15.70% | [ST3.4] | 5 | 4.13% | [CMVM3.2] | 24 | 19.83% |
| [T2.12] | 38 | 31.40% | [SR3.3] | 17 | 14.05% | [ST3.5] | 6 | 4.96% | [CMVM3.3] | 22 | 18.18% |
| [T3.1] | 8 | 6.61% | [SR3.4] | 20 | 16.53% | [ST3.6] | 8 | 6.61% | [CMVM3.4] | 31 | 25.62% |
| [T3.2] | 14 | 11.57% | [SR3.5] | 0 | 0.00% | | | | [CMVM3.5] | 17 | 14.05% |
| [T3.6] | 9 | 7.44% | | | | | | | [CMVM3.6] | 4 | 3.31% |
| | | | | | | | | | [CMVM3.8] | 1 | 0.83% |

FIGURE 18. BSIMM15 SCORECARD. This scorecard shows how often we observed each of the BSIMM15 activities in the data pool of 121 firms.

ACTIVITY CHANGES OVER TIME

The popular business book, *The 7 Habits of Highly Effective People*, explores the theory that successful individuals share common qualities in achieving their goals and that these qualities can be identified and applied by others. The same premise can also be applied to SSIs. Table 11 lists the 20 most observed activities in the BSIMM15 data pool. The data suggests that if your organization is working on its own SSI, you should consider implementing these activities. As a reminder of how practices and activity labeling works, activity [SM1.4] is from the Strategy & Metrics practice, and it was observed in 90.1% of the 121 BSIMM15 participant organizations.

Instead of the top 20 activities overall, Table 10 shows the most common activity in each BSIMM practice. Although we can't directly conclude that these 12 activities are necessary for all SSIs, we can say with confidence that they're commonly found in initiatives whose efforts span all 12 practices. This suggests that if an organization is working on an initiative of its own, its efforts will likely include the majority of these 12 activities over time. Simply put, Table 10 and Table 11 can help you understand what most firms are already doing and discover potential gaps in your program.

In addition to looking at the most common activities, we can also analyze the fastest-growing activity observation rates between BSIMM14 and BSIMM15. Level 1 BSIMM activities are the most common activities observed in each practice, and in BSIMM15, three of the level 1 activities still saw positive growth despite a decline in the BSIMM data pool from 130 firms to 121. Table 12 shows these three activities.

| BSIMM15 TOP ACTIVITES BY PRACTICE | | |
|-----------------------------------|---------|---|
| ACTIVITY | PERCENT | DESCRIPTION |
| [SM1.4] | 90.1 | Implement security checkpoints and associated governance. |
| [CP1.2] | 86.8 | Identify privacy obligations. |
| [T1.1] | 51.2 | Conduct software security awareness training. |
| [AM1.5] | 63.6 | Gather and use attack intelligence. |
| [SFD1.1] | 76.9 | Integrate and deliver security features. |
| [SR1.5] | 79.3 | Identify open source. |
| [AA1.1] | 81.8 | Perform security feature review. |
| [CR1.4] | 87.6 | Use automated code review tools. |
| [ST1.1] | 84.3 | Perform edge/boundary value condition testing during QA. |
| [PT1.1] | 86.0 | Use external penetration testers to find problems. |
| [SE1.2] | 84.3 | Ensure host and network security basics are in place. |
| [CMVM1.1] | 91.7 | Create or interface with incident response. |

TABLE 10. MOST COMMON ACTIVITY PER PRACTICE. This table shows the most observed activity in each of the 12 BSIMM practices for the entire data pool of 121 participant firms.

| BSIMM15 TOP 20 ACTIVITIES BY OBSERVATION PERCENTAGE | | |
|---|---------|---|
| ACTIVITY | PERCENT | DESCRIPTION |
| [CMVM1.1] | 91.7 | Create or interface with incident response. |
| [SM1.4] | 90.1 | Implement security checkpoints and associated governance. |
| [CR1.4] | 87.6 | Use automated code review tools. |
| [CP1.2] | 86.8 | Identify privacy obligations. |
| [PT1.1] | 86.0 | Use external penetration testers to find problems. |
| [ST1.1] | 84.3 | Perform edge/boundary value condition testing during QA. |
| [SE1.2] | 84.3 | Ensure host and network security basics are in place. |
| [AA1.1] | 81.8 | Perform security feature review. |
| [CP1.1] | 81.0 | Unify regulatory pressures. |
| [SR1.2] | 79.3 | Create a security portal. |
| [SR1.5] | 79.3 | Identify open source. |
| [PT1.2] | 78.5 | Feed results to the defect management and mitigation system. |
| [CP1.3] | 77.7 | Create policy. |
| [SFD1.1] | 76.9 | Integrate and deliver security features. |
| [SM1.1] | 74.4 | Publish process and evolve as necessary. |
| [CMVM1.3] | 73.6 | Track software defects found in operations through the fix process. |
| [CMVM1.4] | 72.7 | Have emergency response. |
| [SE1.3] | 71.9 | Implement cloud security controls. |
| [SR1.3] | 71.1 | Translate compliance constraints to requirements. |
| [CMVM1.2] | 70.2 | Identify software defects found in operations monitoring and feed them back to engineering. |

TABLE 11. TOP 20 ACTIVITIES BY OBSERVATION PERCENTAGE. Shown here are the most observed activities in the BSIMM15 data pool of 121 firms. This frequent observation means that each activity has broad applicability across a wide variety of SSIs.

Another way to look at the growth of activities between BSIMM14 and BSIMM15 is to look for trends, such as a high growth in observation rates among common controls. There were 39 activities in BSIMM15 with observations in the range of 40 to 79. The observation rate for three of these activities, shown in Table 13, saw growth rates greater than 8%. In addition, there were 29 activities with observations in the 20 to 39 range, and four of them grew at 20% or more (see Table 14). Finally, three level 1 activities, by definition the generally most common activities within each practice, still saw positive growth rates despite the overall BSIMM data pool shrinking by nine firms, indicating firms are continuing to invest in the fundamentals (see Table 12).

If we analyze these fast-growing activities, we observe a few areas of interest to consider in your SSI:

- Now that [CR1.5] **use automated code review tools** is observed in more than 87% of all firms, SSGs are starting to enforce code reviews for all projects. In addition, some are expanding beyond doing DAST to include **security testing in QA automation** [ST2.5]. This might highlight that more firms are moving to the maturing phase of their SSIs (see Part 3) and are now working on the scalability, efficiency, and effectiveness aspects of their programs.
- Firms are continuing to invest in [CMVM2.4] **streamline incoming responsible vulnerability disclosure**, [SE2.4] **protect code integrity**, and [SE1.4] **define secure deployment parameters and configurations**, all of which are requirements in emerging regulatory activities such as self-attestation to meeting standards from NIST SP 800-218 v1.1 in accordance with Executive Order 14028, Office of Management and Budget (OMB) Memorandum M-22-18, and OMB Memorandum M-23-16 for firms that sell products with software to the United States government.
- Firms are making it easier for developers and security champions to collaborate with security subject matter experts through [T2.12] **provide expertise via open collaboration channels** rather than relying on informal one-on-one communications. This allows all participants to learn from the discussions of others as they occur.

| BSIMM HIGH-GROWTH ACTIVITIES (1) | | |
|----------------------------------|--------|---|
| ACTIVITY | GROWTH | DESCRIPTION |
| [SE1.4] | 8.8% | Define secure deployment parameters and configurations. |
| [SM1.7] | 1.4% | Enforce security checkpoints and track exceptions. |
| [CR1.5] | 1.4% | Make code review mandatory for all projects. |

TABLE 12. VERY COMMON ACTIVITIES WITH CONTINUED POSITIVE GROWTH. This table shows that firms, including those just starting their SSIs, continue to invest in fundamental activities.


| BSIMM HIGH-GROWTH ACTIVITIES (2) | | |
|----------------------------------|--------|---|
| ACTIVITY | GROWTH | DESCRIPTION |
| [CMVM2.4] | 17.1% | Streamline incoming responsible vulnerability disclosure. |
| [SE2.4] | 13.3% | Protect code integrity. |
| [SE1.4] | 8.8% | Define secure deployment parameters and configurations. |

TABLE 13. COMMON ACTIVITIES WITH HIGH GROWTH IN OBSERVATION RATES. This table shows an ongoing trend of investment in common activities. If you are not currently performing or planning to perform these activities, consider them during your next planning cycle.

| BSIMM HIGH-GROWTH ACTIVITIES (3) | | |
|----------------------------------|--------|---|
| ACTIVITY | GROWTH | DESCRIPTION |
| [T2.12] | 35.7% | Provide expertise via open collaboration channels. |
| [SE3.2] | 33.3% | Use code protection. |
| [CMVM3.3] | 22.2% | Simulate software crises. |
| [AM2.8] | 20.0% | Have a research group that develops new attack methods. |

TABLE 14. ACTIVITIES WITH HIGH GROWTH IN OBSERVATION RATES. This table shows potential new trends in the BSIMM15 data pool.

PART 8: THE BSIMM ACTIVITIES



PART 8: THE BSIMM ACTIVITIES

The BSIMM activities are the individual controls used to construct or improve an SSI. They range through people, process, technology, and culture. You can use this information to choose which controls to apply within your initiative, then align your implementation strategy and metrics with your desired outcomes.

The BSIMM framework comprises four domains—Governance, Intelligence, SSDL Touchpoints, and Deployment—and these domains contain 12 practices, such as Strategy & Metrics, Attack Models, and Code Review, which themselves contain activities. These activities are the BSIMM building blocks, the smallest unit of software security granularity implemented to build SSIs. Rather than prescriptively dictating a set of best practices, the BSIMM descriptively observes, quantifies, and documents the actual activities carried out in SSIs across diverse organizations.

ACTIVITIES IN THE BSIMM

The BSIMM is a data-driven model that evolves over time. Over the years, we have added, deleted, and adjusted the levels of various activities based on the data observed throughout the BSIMM's evolution. When considering whether to add a new activity, we analyze whether the effort we're observing is truly new to the model or simply a variation on an existing activity. Similarly, for deciding whether to move an activity between levels within a practice, we use the results of an intra-level standard deviation analysis and the trend in observation counts.

Each activity has a unique label and name—e.g., activity [SM1.4] is in the Strategy & Metrics practice and is named **implement security checkpoints and associated governance**. To preserve backward compatibility, we make all changes by adding new activity labels to the model, even when an activity has simply changed levels within a practice (as an example, we would add a new CR#.# label for both new and moved activities in the Code Review practice).

BSIMM activity levels distinguish the frequency with which activities are observed in the participating organizations. As seen in Part 7, frequently observed activities are designated level 1, with less frequent and infrequently observed activities designated as levels 2 and 3, respectively. Using [SM1.4] as an example again, we see that it is a frequently observed activity in the Strategy & Metrics practice. Note that the new activities we add to the model start with zero observations and are therefore always added at level 3.



Top 10 Activity
in BSIMM15



New Activity
in BSIMM15



Assists With
Adopting AI/ML

To help understand the “Assists With Adopting” AI/ML activities, see AI/ML and BSIMM15 in Part 3.

GOVERNANCE

GOVERNANCE: STRATEGY & METRICS (SM)

The Strategy & Metrics practice encompasses planning, assigning roles and responsibilities, identifying software security goals, determining budgets, and identifying metrics and software release conditions.

[SM1.1: 90] PUBLISH PROCESS AND EVOLVE AS NECESSARY.

The process for addressing software security is defined, published internally, and broadcast to all stakeholders so that everyone knows the plan. Goals, roles, responsibilities, and activities are explicitly defined. Most organizations examine existing methodologies, such as the NIST SSDF, Microsoft SDL, or Black Duck Touchpoints, then tailor them to meet their needs. Security activities will be adapted to software lifecycle processes (e.g., waterfall, Agile, CI/CD, DevOps), so activities will evolve with both the organization and the security landscape. The process doesn't need to be publicly promoted outside the firm to have the desired impact (see [SM3.2]). In addition to publishing the written process, some firms also automate parts (e.g., a testing strategy) as governance-as-code (see [SM3.4]).

[SM1.3: 72] EDUCATE EXECUTIVES ON SOFTWARE SECURITY.

Executives are regularly shown the ways malicious actors attack software and the negative business impacts those attacks can have on the organization. Go beyond reporting of open and closed defects to educate executives on the business risks, including risks of adopting emerging engineering technologies and methodologies without security oversight. Demonstrate a worst-case scenario in a controlled environment with the permission of all involved (e.g., by showing attacks and their business impact). Presentation to the Board can help garner resources for new or ongoing SSI efforts. Demonstrating the need for new skill-building training in evolving areas, such as DevOps groups using cloud-native technologies, can help convince leadership to accept SSG recommendations when they might otherwise be ignored in favor of faster release dates or other priorities. Bring in an outside expert when necessary to bolster executive attention.



[SM1.4: 109] IMPLEMENT SECURITY CHECKPOINTS AND ASSOCIATED GOVERNANCE.

The software security process includes checkpoints (such as gates, release conditions, guardrails, milestones, etc.) at one or more points in a software lifecycle. The first two steps toward establishing security-specific checkpoint conditions are to identify process locations that are compatible with existing development practices and to then begin gathering the information necessary, such as risk-ranking thresholds or defect data, to make a go/no-go decision. Importantly, the conditions need not be enforced at this stage—e.g., the SSG can collect security testing results

for each project prior to release, then provide an informed opinion on what constitutes sufficient testing or acceptable test results without trying to stop a project from moving forward (see [SM1.7]). Shorter release cycles might require creative approaches to collecting the right evidence and rely heavily on automation. Socializing the conditions and then enforcing them once most project teams already know how to succeed is a gradual approach that motivates good behavior without introducing unnecessary friction.

[SM1.7: 72] ENFORCE SECURITY CHECKPOINTS AND TRACK EXCEPTIONS.

Enforce security release conditions at each checkpoint (gate, guardrail, milestone, etc.) for every project, so that each project must either meet an established measure or follow a defined process for obtaining an exception to move forward. Use internal policies and standards, regulations, contractual agreements, and other obligations to define release conditions, then track all exceptions. Verifying conditions yields data that informs the KRIs and any other metrics used to govern the process. Automatically giving software a passing grade or granting exceptions without due consideration defeats the purpose of verifying conditions. Even seemingly innocuous software projects (e.g., small code changes, infrastructure access control changes, deployment blueprints) must successfully satisfy the prescribed security conditions as they progress through the software lifecycle. Similarly, APIs, frameworks, libraries, bespoke code, microservices, container configurations, etc., are all software that must satisfy security release conditions. It's possible, and often very useful, to have verified the conditions both before and after the development process itself. In modern development environments, the verification process will increasingly become automated (see [SM3.4]).

[SM2.1: 65] PUBLISH DATA ABOUT SOFTWARE SECURITY INTERNALLY AND USE IT TO DRIVE CHANGE.

To facilitate improvement, data is published internally about the state of software security within the organization. Produce security or development dashboards with metrics for executives and software development management. Dashboards can be part of pipeline toolchains to enable developer self-improvement. Sometimes, this published data won't be shared with everyone in the firm but only with the stakeholders who are tasked to drive change. In other cases, open book management and data published to all stakeholders helps everyone know what's going on. If the organization's culture promotes internal competition between groups, use this information to add a security dimension. Integrate automated security telemetry to gather measurements quickly and accurately to increase timeliness of security data in areas such as speed (e.g., time to fix) and quality (e.g., defect density). Publishing data about new technologies (e.g., security and risk in cloud-native architectures) is important for identifying needed improvements.

[SM2.3: 67] CREATE OR GROW A SECURITY CHAMPIONS PROGRAM.

Form a collection of people scattered across the organization—often called security champions—who show an above-average level of security interest or skill and who contribute software security expertise to development, QA, and operations teams. Forming this social network of advocates is a good step toward scaling security into software engineering. One way to build the initial group is to track the people who stand out during introductory training courses (see [T3.6]). Another way is to ask for volunteers. In a more top-down approach, initial champions membership is assigned to ensure good coverage of development groups, but ongoing membership is based on actual performance. The champions can act as a sounding board for new projects and, in new or fast-moving technology areas, can help combine software security skills with domain knowledge that might be under-represented in the SSG or engineering teams. Agile coaches, scrum masters, and DevOps engineers can make particularly useful champions members, especially for detecting and removing process friction. In some environments, champions-led efforts are delivered via automation (e.g., as-code).

[SM2.6: 69] REQUIRE SECURITY SIGN-OFF PRIOR TO SOFTWARE RELEASE.

The organization has an initiative-wide process for documenting accountability and accepting security risk by having a risk owner use SSG-approved criteria to sign off on the state of all software prior to release. The sign-off policy might also require the accountable person to, e.g., acknowledge critical vulnerabilities that have not been mitigated or SSDL steps that have been skipped. Informal or uninformed risk acceptance alone isn't a security sign-off because the act of accepting risk is more effective when it's formalized (e.g., with a signature, a form submission, or something similar) and captured for future reference. Similarly, simply stating that certain projects don't need sign-off at all won't achieve the desired risk management results. In some cases, however, the risk owner can provide the sign-off on a particular set of software project acceptance criteria, which are then implemented in automation to provide governance-as-code (see [SM3.4]), but there must be an ongoing verification that the criteria remain accurate and the automation is working.

[SM2.7: 52] CREATE EVANGELISM ROLE AND PERFORM INTERNAL MARKETING.

Build support for software security throughout the organization via ongoing evangelism and ensure that everyone aligns on security objectives. This internal marketing function, often performed by a variety of stakeholder roles, keeps executives and others up to date on the magnitude of the software security problem and the elements of its solution. A champion or a scrum master familiar with security, for example, could help teams adopt better software security practices as they transform to Agile and DevOps methods. Similarly, a cloud expert could demonstrate the changes needed in security architecture and testing for serverless applications. Evangelists can increase understanding and build credibility by giving

talks to internal groups (including executives), publishing roadmaps, authoring technical papers for internal consumption, or creating a collection of papers, books, and other resources on an internal website (see [SR1.2]) and promoting its use. In turn, organizational feedback becomes a useful source of improvement ideas.

[SM3.1: 30] USE A SOFTWARE ASSET TRACKING APPLICATION WITH PORTFOLIO VIEW.

The SSG uses centralized tracking automation to chart the progress of every piece of software and deployable artifact, from creation to decommissioning, regardless of development methodology. The automation records the security activities scheduled, in progress, and completed, incorporating results from SSDL activities even when they happen in a tight loop or during deployment. The combined inventory and security posture view enables timely decision-making. The SSG uses the automation to generate portfolio reports for multiple metrics and, in many cases, publishes this data at least among executives. As an initiative matures and activities become more distributed, the SSG uses the centralized reporting system to keep track of all the moving parts.

[SM3.2: 23] MAKE SSI EFFORTS PART OF EXTERNAL MARKETING.

To build external awareness, the SSG helps market the SSI beyond internal teams. The process of sharing details externally and inviting critique is used to bring new perspectives into the firm. Promoting the SSDL externally can turn security efforts into a market differentiator, and feedback from external marketing can grow an SSI's risk reduction exercises into a competitive advantage. The SSG might provide details at external conferences or trade shows. In some cases, a complete SSDL methodology can be published and promoted outside the firm, and governance-as-code concepts can make interesting case studies.

[SM3.3: 32] IDENTIFY METRICS AND USE THEM TO DRIVE RESOURCING.

The SSG and its management identify metrics that define and measure SSI progress in quantitative terms. These metrics are reviewed on a regular basis and drive the initiative's budgeting and resource allocations, so simple counts and out-of-context measurements won't suffice here. On the technical side, one such metric could be defect density, a reduction of which could be used to show a decreasing cost of remediation over time, assuming, of course, that testing depth has kept pace with software changes. Data for metrics is best collected early and often using event-driven processes with telemetry rather than relying on calendar-driven data collection. The key is to tie security results to business objectives in a clear and obvious fashion to justify resourcing. Because the concept of security is already tenuous to many businesspeople, make the tie-in explicit.

[SM3.4: 11] INTEGRATE SOFTWARE-DEFINED LIFECYCLE GOVERNANCE.

Organizations begin replacing traditional document-, presentation-, and spreadsheet-based lifecycle management

with software-based delivery platforms. For some software lifecycle phases, humans are no longer the primary drivers of progression from one phase to the next. Instead, organizations rely on automation to drive the management and delivery process with software such as Spinnaker or GitHub, and humans participate asynchronously (and often optionally). Automation often extends beyond the scope of CI/ CD to include functional and nonfunctional aspects of delivery, such as health checks, cut-over on failure, rollback to known-good state, defect discovery and management, compliance verification, and a way to ensure adherence to policies and standards. Some organizations are also evolving their lifecycle management approach by integrating their compliance and defect discovery data, perhaps augmented by intelligence feeds and other external data, to begin moving from a series of point-in-time go/no-go decisions (e.g., release conditions) to a future state of continuous accumulation of assurance data (see [CMVM3.6]).



[SM3.5: 1] INTEGRATE SOFTWARE SUPPLY CHAIN RISK MANAGEMENT.

Organizational risk management processes ensure that important software created by and entering the organization is managed through policy-driven access and usage controls, maintenance standards (see [SE3.9]), and captured software provenance data (see [SE2.4]). Apply these processes to external (see [SR2.7]), bespoke, and internally developed software (see [SE3.9]) to help ensure that deployed code has the expected components (see [SE3.8]). The lifecycle management for all software, from creation or importation through secure deployment, ensures that all access, usage, and modifications are done in accordance with policy. This assurance is easier to implement at scale using automation in software lifecycle processes (see [SM3.4]).

GOVERNANCE: COMPLIANCE & POLICY (CP)

The Compliance & Policy practice is focused on identifying controls for compliance regimens such as PCI DSS and GDPR, developing contractual controls such as SLAs to help manage COTS risk, setting organizational software security policy, and auditing against that policy.



[CPI.1: 98] UNIFY REGULATORY PRESSURES.

Have a cross-functional team that understands the constraints imposed on software security by regulatory or compliance drivers that are applicable to the organization and its customers. The team takes a common approach that removes redundancy and conflicts to unify compliance requirements, such as from PCI security standards; GLBA, SOX, and HIPAA in the US; or GDPR in the EU. A formal approach will map applicable portions of regulations to controls (see [CP2.3]) applied to software to explain how the organization complies. Existing business processes run by legal, product management, or other risk and compliance groups outside the SSG could serve as the regulatory focal point, with the SSG providing software security knowledge. A unified set of software

security guidance for meeting regulatory pressures ensures that compliance work is completed as efficiently as possible.

TOP 10

[CPI.2: 105] IDENTIFY PRIVACY OBLIGATIONS.

The SSG identifies privacy obligations stemming from regulation and customer expectations, then translates these obligations into both software requirements and privacy best practices. The way software handles PII might be explicitly regulated, but even if it isn't, privacy is an important topic. For example, if the organization processes credit card transactions, the SSG will help in identifying the privacy constraints that the PCI DSS places on the handling of cardholder data and will inform all stakeholders (see [SR1.3]). Note that outsourcing to hosted environments (e.g., the cloud) doesn't relax privacy obligations and can even increase the difficulty of recognizing and meeting all associated needs. Also, note that firms creating software products that process PII when deployed in customer environments might meet this need by providing privacy controls and guidance for their customers. Evolving consumer privacy expectations, the proliferation of "software is in everything," and data scraping and correlation (e.g., social media) add additional expectations and complexities for PII protection.

[CPI.3: 94] CREATE POLICY.

The SSG guides the organization by creating or contributing to software security policies that satisfy internal, regulatory, and customer-driven security requirements. This policy is what is permitted and denied at the initiative level—if it's not mandatory and enforced, it's not policy. The policies include a unified approach for satisfying the (potentially lengthy) list of security drivers at the governance level so that project teams can avoid keeping up with the details involved in complying with all applicable regulations or other mandates. Likewise, project teams won't need to relearn customer security requirements on their own. Architecture standards and coding guidelines aren't examples of policy, but policy that prescribes and mandates their use for certain software categories falls under this umbrella. In many cases, policy statements are translated into automation to provide governance-as-code. Even if not enforced by humans, policy that's been automated must still be mandatory. In some cases, policy will be documented exclusively as governance as-code (see [SM3.4]), often as tool configuration, but it must still be readily readable, auditable, and editable by humans.

[CP2.1: 49] BUILD A PII INVENTORY.

The organization identifies and tracks the kinds of PII processed or stored by each of its systems, along with their associated data repositories. In general, simply noting which applications process PII isn't enough—the type of PII (e.g., PHI, PFI, PI) and where it's stored are necessary so that the inventory can be easily referenced in critical situations. This usually includes making a list of databases that would require customer notification if breached or a list to use in crisis simulations (see [CMVM3.3]). Build the PII inventory by starting with each individual application and noting its PII use or by starting with PII types and noting the applications that touch each one.

System architectures have evolved such that PII will often flow into cloud-based service and endpoint device ecosystems, then come to rest there (e.g., content delivery networks, workflow systems, mobile devices, IoT devices), making it tricky to keep an accurate PII inventory.

[CP2.2: 58] REQUIRE SECURITY SIGN-OFF FOR COMPLIANCE-RELATED RISK.

The organization has a formal compliance risk acceptance sign-off and accountability process that addresses all software development projects. In this process, the SSG acts as an advisor while the risk owner signs off on the software's compliance state prior to release based on its adherence to documented criteria. The sign-off policy might also require the head of the business unit to, e.g., acknowledge compliance issues that haven't been mitigated or compliance-related SSDL steps that have been skipped, but sign-off is required even when no compliance-related risk is present. Sign-off is explicit and captured for future reference, with any exceptions tracked, even in automated application lifecycle methodologies. Note that an application without security defects might still be noncompliant, so clean security testing results are not a substitute for a compliance sign-off. Even in DevOps organizations where engineers have the technical ability to release software, there is still a need for a deliberate risk acceptance step even if the compliance criteria are embedded in automation (see [SM3.4]). In cases where the risk owner signs off on a particular set of compliance acceptance criteria that are then implemented in automation to provide governance as-code, there must be ongoing verification that the criteria remain accurate and the automation is actually working.

[CP2.3: 69] IMPLEMENT AND TRACK CONTROLS FOR COMPLIANCE.

The organization can demonstrate compliance with applicable requirements because its SSDL is aligned with the control statements that were developed by the SSG in collaboration with compliance stakeholders (see [CP1.1]). The SSG collaborates with stakeholders to track controls, navigate problem areas, and ensure that auditors and regulators are satisfied. The SSG can then remain in the background when the act of following the SSDL automatically generates the desired compliance evidence predictably and reliably. Increasingly, the DevOps approach embeds compliance controls in automation, such as in software-defined infrastructure and networks, rather than in human process and manual intervention. A firm doing this properly can explicitly associate satisfying its compliance concerns with following its SSDL.

[CP2.4: 60] INCLUDE SOFTWARE SECURITY SLAS IN ALL VENDOR CONTRACTS.

Software vendor contracts include an SLA to ensure that the vendor's security efforts align with the organization's security and compliance story. Each new or renewed contract contains provisions requiring the vendor to address software security and deliver a product or service compatible with the organization's security policy. In some cases, open source licensing concerns initiate the vendor management process, which can open the door for additional software security language in the

SLA (see [SR2.5]). Typical provisions set requirements for policy conformance, incident management, training, defect management, and response times for addressing software security issues. Traditional IT security requirements and a simple agreement to allow penetration testing or another defect discovery method aren't sufficient here.

[CP2.5: 70] ENSURE EXECUTIVE AWARENESS OF COMPLIANCE AND PRIVACY OBLIGATIONS.

Gain buy-in around compliance and privacy obligations by providing executives with plain-language explanations of both the organization's compliance and privacy requirements and the potential consequences of failing to meet those requirements. For some organizations, explaining the direct cost and likely fallout from a compliance failure or data breach can be an effective way to broach the subject. For others, having an outside expert address the Board works because some executives value an outside perspective more than an internal one. A sure sign of proper executive buy-in is an acknowledgment of the need along with adequate allocation of resources to meet those obligations. Use the sense of urgency that typically follows a compliance or privacy failure to build additional awareness and bootstrap new efforts.

[CP3.1: 36] DOCUMENT A SOFTWARE COMPLIANCE STORY.

The SSG can demonstrate the organization's up-to-date software security compliance story on demand. A compliance story is a collection of data, artifacts, policy controls, or other documentation that shows the compliance state of the organization's software and processes. Often, senior management, auditors, and regulators—whether government or other—will be satisfied with the same kinds of reports that can be generated directly from various tools. In some cases, particularly where organizations leverage shared responsibility through cloud services, the organization will require additional information from vendors about how that vendor's controls support organizational compliance needs. It will often be necessary to normalize information that comes from disparate sources.

[CP3.2: 39] ENSURE COMPATIBLE VENDOR POLICIES.

Ensure that vendor software security policies and SSDL processes are compatible with internal policies. Vendors likely comprise a diverse group—cloud providers, middleware providers, virtualization providers, container and orchestration providers, bespoke software creators, contractors, and many more—and each might be held to different policy requirements. Policy adherence enforcement might be through a point-in-time review (such as ensuring acceptance criteria), automated checks (such as those applied to pull requests, committed artifacts like containers, or similar), or convention and protocol (such as preventing services connection unless security settings are correct and expected certificates are present). Evidence of vendor adherence could include results from SSDL activities, from manual tests or tests built directly into automation or infrastructure, or from other software lifecycle instrumentation. For some policies or SSDL processes, vendor questionnaire responses and attestation alone might be sufficient.

[CP3.3: 13] DRIVE FEEDBACK FROM SOFTWARE LIFECYCLE DATA BACK TO POLICY.

Feed information from the software lifecycle into the policy creation and maintenance process to drive improvements, such as in defect prevention and strengthening governance-as-code practices (see [SM3.4]). With this feedback as a routine process, blind spots can be eliminated by mapping them to trends in SSDL failures. Events such as the regular appearance of inadequate architecture analysis, recurring vulnerabilities, ignored security release conditions, or the wrong vendor choice for carrying out a penetration test can expose policy weakness (see [CP1.3]). As an example, lifecycle data including KPIs, OKRs, KRIs, SLIs, SLOs, or other organizational metrics can indicate where policies impose too much bureaucracy by introducing friction that prevents engineering from meeting the expected delivery cadence. Rapid technology evolution might also create policy gaps that must be addressed. Over time, policies become more practical and easier to carry out (see [SM1.1]). Ultimately, policies are refined with SSDL data to enhance and improve effectiveness.

GOVERNANCE: TRAINING (T)

Training has always played a critical role in software security because organizational stakeholders across governance, risk, and compliance (GRC), legal, engineering, operations, and other groups often start with little security knowledge.

[TI.1: 62] CONDUCT SOFTWARE SECURITY AWARENESS TRAINING.

To promote a culture of software security throughout the organization, the SSG conducts periodic software security awareness training. This training might be delivered via SSG members, security champions, an outside firm, the internal training organization, or e-learning, but course content isn't necessarily tailored for a specific audience—developers, QA engineers, and project managers could attend the same "Introduction to Software Security" course, for example. Augment this content with a tailored approach that addresses the firm's culture explicitly, which might include the process for building security in, avoiding common mistakes, and technology topics such as CI/CD and DevSecOps. Generic introductory courses that only cover basic IT or high-level security concepts don't generate satisfactory results. Likewise, awareness training aimed only at developers and not at other roles in the organization is insufficient.

[TI.7: 61] DELIVER ON-DEMAND INDIVIDUAL TRAINING.

The organization lowers the burden on students and reduces the cost of delivering software security training by offering on-demand training for SSDL stakeholders. The most obvious choice, e-learning, can be kept up to date through a subscription model, but an online curriculum must be engaging and relevant to students in various roles (e.g., developer, QA, cloud, ops) to achieve its intended purpose. Ineffective (e.g., aged, off-topic) training or training that isn't used won't create any change. Hot engineering topics like containerization and security orchestration, and new training delivery styles such

as gamification, will attract more interest than boring policy discussions. For developers, it's possible to provide training directly through the IDE right when it's needed, but in some cases, building a new skill (such as cloud security or threat modeling) might be better suited for instructor-led training, which can also be provided on demand.

[T1.8: 50] INCLUDE SECURITY RESOURCES IN ONBOARDING.

The process for bringing new hires into a software engineering organization requires timely completion of a training module about software security. While the generic new hire process usually covers topics like picking a good password and avoiding phishing, this orientation period is enhanced to cover topics such as how to create, deploy, and operate secure code, the SSDL, security standards (see [SR1.1]), and internal security resources (see [SR1.2]). The objective is to ensure that new hires contribute to the security culture as soon as possible. Although a generic onboarding module is useful, it doesn't take the place of a timely and more complete introductory software security course.

[T2.5: 41] ENHANCE SECURITY CHAMPIONS THROUGH TRAINING AND EVENTS.

Strengthen the security champions network (see [SM2.3]) by inviting guest speakers or holding special events about advanced software security topics. This effort is about providing to the champions customized training (e.g., the latest software security techniques for DevOps or serverless technologies or on the implications of new policies and standards) so that it can fulfill its assigned responsibilities—it's not about inviting champions members to routine brown bags or signing them up for standard computer-based training. Similarly, a standing conference call with voluntary attendance won't get the desired results, which are as much about building camaraderie as they are about sharing knowledge and organizational efficiency. Regular events build community and facilitate collaboration and collective problem-solving. Face-to-face meetings are by far the most effective, even if they happen only once or twice a year and even if some participants must attend by videoconferencing. In teams with many geographically dispersed and work-from-home members, simply turning on cameras and ensuring that everyone gets a chance to speak makes a substantial difference.

[T2.8: 25] CREATE AND USE MATERIAL SPECIFIC TO COMPANY HISTORY.

To make a strong and lasting change in behavior, training includes material specific to the company's history of software security challenges. When participants can see themselves in a problem, they're more likely to understand how the material is relevant to their work as well as when and how to apply what they've learned. One way to do this is to use noteworthy attacks on the company's software as examples in the training curriculum. Both successful and unsuccessful attacks, as well as notable results from penetration tests, design review, and red team exercises, can make good teachable moments. Stories from company history can help steer training in the right direction but only if those stories are still relevant and not

overly censored. This training should cover platforms used by developers (developers orchestrating containers probably won't care about old virtualization problems) and problems relevant to languages in common use.

[T2.9: 31] DELIVER ROLE-SPECIFIC ADVANCED CURRICULUM.

Software security training goes beyond building awareness (see [T1.1]) to enabling students to incorporate security practices into their work. This training is tailored to cover the tools, technology stacks, development methodologies, and issues that are most relevant to the students. An organization could offer tracks for its engineers, for example, supplying one each for architects, developers, operations, DevOps, site reliability engineers, and testers. Tool-specific training is also commonly needed in such a curriculum. While it might be more concise than engineering training, role-specific training is also necessary for many other stakeholders within an organization, including product management, executives, and others. In any case, the training must be taken by a broad enough audience to build the collective skillsets required.

[T2.10: 25] HOST SOFTWARE SECURITY EVENTS.

The organization hosts security events featuring external speakers and content in order to strengthen its security culture. Good examples of such events are Intel iSecCon and AWS re:Inforce, which invite all employees, feature external presenters, and focus on helping engineering create, deploy, and operate better code. Employees benefit from hearing outside perspectives, especially those related to fast-moving technology areas with software security ramifications, and the organization benefits from putting its security credentials on display (see [SM3.2]). Events open only to small, select groups or simply putting recordings on an internal portal, won't result in the desired culture change across the organization.

[T2.11: 29] REQUIRE AN ANNUAL REFRESHER.

Everyone involved in the SSDL is required to take an annual software security refresher course. This course keeps the staff up to date on the organization's security approach and ensures that the organization doesn't lose focus due to turnover, evolving methodologies, or changing deployment models. The SSG might give an update on the security landscape and explain changes to policies and standards. A refresher could also be rolled out as part of a firmwide security day or in concert with an internal security conference. While one refresher module can be used for multiple roles (see [T2.9]), coverage of new topics and changes to the previous year's content should result in a significant amount of fresh content.

[T2.12: 38] PROVIDE EXPERTISE VIA OPEN COLLABORATION CHANNELS.

Software security experts offer help to anyone in an open manner during regularly scheduled office hours or openly accessible channels on Slack, Jira, or similar. By acting as an informal resource for people who want to solve security

problems, the SSG leverages teachable moments and emphasizes the carrot over the stick approach to security best practices. Office hours might be hosted one afternoon per week by a senior SSG member, perhaps inviting briefings from product or application groups working on hard security problems. Slack and other messaging applications can capture questions 24x7, functioning as an office hours platform when appropriate subject matter experts are consistently part of the conversation and are ensuring that the answers generated align with SSG expectations. An online approach has the added benefit of discussions being recorded and searchable.

[T3.1: 8] REWARD PROGRESSION THROUGH CURRICULUM.

Progression through the security curriculum brings personal benefits, such as public acknowledgement or career advancement. The reward system can be formal and lead to a certification or an official mark in the human resources system, or it can be less formal and include motivators such as documented praise at annual review time. Involving a corporate training department and human resources team can make the impact of improving security skills on career progression more obvious, but the SSG should continue to monitor security knowledge in the firm and not cede complete control or oversight. Coffee mugs and t-shirts can build morale, but it usually takes the possibility of real career progression to change behavior.

[T3.2: 14] PROVIDE TRAINING FOR VENDORS AND OUTSOURCED WORKERS.

Vendors and outsourced workers receive appropriate software security training, comparable to the level of training given to employees. Spending time and effort helping suppliers get security right at the outset is much easier than trying to determine what went wrong later, especially if the development team has moved on to other projects. Training individual contractors is much more natural than training entire outsourced firms and is a reasonable place to start. It's important that everyone who works on the firm's software has an appropriate level of training that increases their capability of meeting the software security expectations for their role, regardless of their employment status. Of course, some vendors and outsourced workers might have received adequate training from their own firms, but that should always be verified.

[T3.6: 9] IDENTIFY NEW SECURITY CHAMPIONS THROUGH OBSERVATION.

Future security champions are recruited by noting people who stand out during opportunities that show skill and enthusiasm, such as training courses, office hours, capture-the-flag exercises, hack-a-thons, etc. and then encouraging them to join the champions. Pay particular attention to practitioners who are contributing things such as code, security configurations, or defect discovery rules. The champions program often begins as an assigned collection of people scattered across the organization who show an above-average level of security interest or advanced knowledge of new technology stacks and development methodologies (see [SM2.3]). Identifying future members proactively is a step toward creating a social network

that speeds the adoption of security into software development and operations. A group of enthusiastic and skilled volunteers will be easier to lead than a group that is drafted.

INTELLIGENCE

INTELLIGENCE: ATTACK MODELS (AM)

Attack Models capture information used to think like an attacker, including threat modeling inputs, abuse cases, data classification, and technology-specific attack patterns.

[AMI.2: 64] USE A DATA CLASSIFICATION SCHEME FOR SOFTWARE INVENTORY.

Security stakeholders in an organization agree on a data classification scheme and use it to inventory software, delivery artifacts (e.g., containers), and associated persistent data stores according to the kinds of data processed or services called, regardless of deployment model (e.g., on- or off-premises). Many classification schemes are possible—one approach is to focus on PII, for example. Depending on the scheme and the software involved, it could be easiest to first classify data repositories (see [CP2.1]), then derive classifications for applications according to the repositories they use. Other approaches include data classification according to protection of intellectual property, impact of disclosure, exposure to attack, relevance to GDPR, and geographic boundaries.

[AMI.3: 49] IDENTIFY POTENTIAL ATTACKERS.

The SSG identifies potential attackers in order to understand and begin documenting their motivations and abilities. The outcome of this periodic exercise could be a set of attacker profiles that includes outlines for categories of attackers, and more detailed descriptions for noteworthy individuals, that are used in end-to-end design review (see [AA1.2]). In some cases, a third-party vendor might be contracted to provide this information. Specific and contextual attacker information is almost always more useful than generic information copied from someone else's list. Moreover, a list that simply divides the world into insiders and outsiders won't drive useful results. Identification of attackers should also consider the organization's evolving software supply chain, attack surface, theoretical internal attackers, and contract staff.



[AMI.5: 77] GATHER AND USE ATTACK INTELLIGENCE.

The SSG ensures the organization stays ahead of the curve by learning about new types of attacks and vulnerabilities, then adapts that information to the organization's needs. Attack intelligence must be made actionable and useful for a variety of consumers, which might include developers, testers, DevOps, security operations, and reliability engineers, among others. In many cases, a subscription to a commercial service can provide a reasonable way of gathering basic attack intelligence related to applications, APIs, containerization, orchestration, cloud environments, etc. Attending technical conferences and monitoring attacker forums, then correlating that information with what's happening in the organization (perhaps by

leveraging automation to mine operational logs and telemetry) helps everyone learn more about emerging vulnerability exploitation.

[AM2.1: I8] BUILD ATTACK PATTERNS AND ABUSE CASES TIED TO POTENTIAL ATTACKERS.

The SSG works with stakeholders to build attack patterns and abuse cases tied to potential attackers (see [AM1.3]). Attack patterns frequently contain details of the targeted asset, attackers, goals, and the techniques used. These resources can be built from scratch or from standard sets, such as the MITRE ATT&CK framework, with the SSG adding to the pile based on its own attack stories to prepare the organization for SSDL activities such as design review and penetration testing. For example, a story about an attack against a poorly designed cloud-native application could lead to a containerization attack pattern that drives a new type of testing (see [ST3.5]). If a firm tracks the fraud and monetary costs associated with specific attacks, this information can in turn be used to prioritize the process of building attack patterns and abuse cases. Organizations will likely need to evolve both their attack pattern and abuse case creation prioritization and their content over time due to changing software architectures (e.g., zero trust, cloud native, serverless), attackers, and technologies.

[AM2.6: I2] COLLECT AND PUBLISH ATTACK STORIES.

To maximize the benefit from lessons that don't always come cheap, the SSG collects and publishes stories about attacks against the organization's software. Both successful and unsuccessful attacks can be noteworthy, and discussing historical information about software attacks has the added effect of grounding software security in a firm's reality. This is particularly useful in training classes (see [T2.8]) to help counter a generic approach that might be overly focused on other organizations' most common bug lists or outdated platform attacks. Hiding or overly sanitizing information about attacks from people building new systems fails to garner any positive benefits from a negative event.

[AM2.7: I6] BUILD AN INTERNAL FORUM TO DISCUSS ATTACKS.

The organization has an internal, interactive forum where the SSG, champions, incident response, and others discuss attacks and attack methods. The discussion serves to communicate the attacker perspective to everyone, so it's useful to include all successful attacks here, regardless of attack source, such as supply chain, internal, consultants, or bug bounty contributors. The SSG augments the forum with an internal communication channel (see [T2.12]) that encourages subscribers to discuss the latest information on publicly known incidents. Dissection of attacks and exploits that are relevant to a firm are particularly helpful when they spur discussion of software, infrastructure, and other mitigations. Simply republishing items from public mailing lists doesn't achieve the same benefits as active and ongoing discussions, nor does a closed discussion hidden from those creating code and configurations. Everyone should feel free to ask questions and learn about vulnerabilities and exploits.

[AM2.8: 24] HAVE A RESEARCH GROUP THAT DEVELOPS NEW ATTACK METHODS.

A research group works to identify and mitigate the impact of new classes of attacks and shares their knowledge with stakeholders. Identification does not always require original research—the group might expand on an idea discovered by others. Doing this research inhouse is especially important for early adopters of new technologies and configurations so that they can discover potential weaknesses before attackers do. One approach is to create new attack methods that simulate persistent attackers during goal-oriented red team exercises (see [PT3.1]). This isn't a penetration testing team finding new instances of known types of weaknesses, it's a research group that innovates attack methods and mitigation approaches. Example mitigation approaches include test cases, static analysis rules, attack patterns, standards, and policy changes. Some firms provide researchers time to follow through on their discoveries by using bug bounty programs or other means of coordinated disclosure (see [CMVM2.4]). Others allow researchers to publish their findings at conferences like DEF CON to benefit everyone.

[AM2.9: I8] MONITOR AUTOMATED ASSET CREATION.

Implement technology controls that provide a continuously updated view of the various network, machine, software, and related infrastructure assets being instantiated by engineering teams. To help ensure proper coverage, the SSG works with engineering teams (including potential shadow IT teams) to understand orchestration, cloud configuration, and other self-service means of software delivery to ensure proper monitoring. This monitoring requires a specialized effort—normal system, network, and application logging and analysis won't suffice. Success might require a multi-pronged approach, including consuming orchestration and virtualization metadata, querying cloud service provider APIs, and outside-in crawling and scraping.

[AM3.2: 8] CREATE AND USE AUTOMATION TO MIMIC ATTACKERS.

The SSG arms engineers, testers, and incident response with automation to mimic what attackers are going to do. For example, a new attack method identified by an internal research group (see [AM2.8]) or a disclosing third party could require a new tool, so the SSG, perhaps through the security champions, could package the tool and distribute it to testers. The idea here is to push attack capability past what typical commercial tools and offerings encompass, then make that knowledge and technology easy for others to use. Mimicking attackers, especially attack chains, almost always requires tailoring tools to a firm's particular technology stacks, infrastructure, and configurations. When technology stacks and coding languages evolve faster than vendors can innovate, creating tools and automation in-house might be the best way forward. In the DevOps world, these tools might be created by engineering and embedded directly into toolchains and automation (see [ST3.6]).



[AM3.4: II] CREATE TECHNOLOGY-SPECIFIC ATTACK PATTERNS.

The SSG facilitates technology-specific attack pattern creation by collecting and providing knowledge about attacks relevant to the organization's technologies. For example, if the organization's cloud software relies on a cloud vendor's security apparatus (e.g., key and secrets management), the SSG or appropriate SMEs can help catalog the quirks of the crypto package and how it might be exploited. Attack patterns directly related to the security frontier (e.g., AI, serverless) can be useful here as well. It's often easiest to start with existing generalized attack patterns to create the needed technology specific ones, but simply adding "for microservices" at the end of a generalized pattern name, for example, won't suffice.

[AM3.5: IO] MAINTAIN AND USE A TOP N POSSIBLE ATTACKS LIST.

The SSG periodically digests the ever-growing list of applicable attack types, creates a prioritized short list—the top N—and then uses the list to drive change. This initial list almost always combines input from multiple sources, both inside and outside the organization. Some organizations prioritize their list according to a perception of potential business loss while others might prioritize according to preventing successful attacks against their software. The top N list doesn't need to be updated with great frequency, and attacks can be coarsely sorted. For example, the SSG might brainstorm twice a year to create lists of attacks the organization should be prepared to counter "now," "soon," and "someday."

INTELLIGENCE: SECURITY FEATURES & DESIGN (SFD)

The Security Features & Design practice is charged with creating usable security patterns for major security controls (meeting the standards defined in the Standards & Requirements practice), building components and services for those controls, and establishing collaboration during security design efforts.

[SFDI.I: 93] INTEGRATE AND DELIVER SECURITY FEATURES.

Provide proactive guidance on preapproved security features for engineering groups to use rather than each group implementing its own security features. Engineering groups benefit from implementations that come preapproved, and the SSG benefits by not having to repeatedly track down the kinds of subtle errors that often creep into security features (e.g., authentication, role management, key management, logging, cryptography, protocols). These security features might be discovered during SSDL activities, created by the SSG or specialized development teams, or defined in configuration templates (e.g., cloud blueprints) and delivered via mechanisms such as SDKs, containers, microservices, and APIs. Generic security features often must be tailored for specific platforms. For example, each mobile and cloud platform might need its own means by which users are authenticated and authorized, secrets are managed, and user actions are centrally logged and monitored. It's

implementing and disseminating these defined security features that generates real progress, not simply making a list of them.

[SFDI.2: 83] APPLICATION ARCHITECTURE TEAMS ENGAGE WITH THE SSG.

Application architecture teams take responsibility for security in the same way they take responsibility for performance, availability, scalability, and resiliency. One way to keep security from falling out of these architecture discussions is to have secure design experts (from the SSG, a vendor, etc.) participate. Increasingly, architecture discussions include developers and site reliability engineers who are governing all types of software components, such as open source, APIs, containers, and cloud services. In other cases, enterprise architecture teams have the knowledge to help the experts create secure designs that integrate properly into corporate design standards. Proactive engagement with experts is key to success here. In addition, it's never safe for one team to assume another team has addressed security requirements—even moving a well-known system to the cloud means reengaging the experts.

[SFD2.I: 42] LEVERAGE SECURE-BY-DESIGN COMPONENTS AND SERVICES.

Build or provide approved secure-by-design software components and services for use by engineering teams. Prior to approving and publishing secure-by-design software components and services, including open source and cloud services, the SSG must carefully assess them for security. This assessment process to declare a component secure-by-design is usually more rigorous and in-depth than that for typical projects. In addition to teaching by example, these resilient and reusable building blocks aid important efforts such as architecture analysis and code review by making it easier to avoid mistakes. These components and services also often have features (e.g., application identity, RBAC) that enable uniform usage across disparate environments. Similarly, the SSG might further take advantage of this defined list by tailoring static analysis rules specifically for the components it offers (see [CR2.6]).

[SFD2.2: 68] CREATE CAPABILITY TO SOLVE DIFFICULT DESIGN PROBLEMS.

Contribute to building resilient architectures by solving design problems unaddressed by organizational security components or services, or by cloud service providers, thus minimizing the negative impact that security has on other constraints, such as feature velocity. Involving the SSG and secure design experts in application refactoring or in the design of a new protocol, microservice, or architecture feature (e.g., containerization) enables timely analysis of the security implications of existing defenses and identifies elements to be improved. Designing for security early in the new project process is more efficient than analyzing an existing design for security and then refactoring when flaws are uncovered (see [AA1.1], [AA1.2], [AA2.1]). The SSG could also get involved in what would have historically been purely engineering discussions, as even rudimentary use of cloud-native technologies (e.g., "Hello, world!") requires proper

use of configurations and other capabilities that have direct implications on security posture.



[SFD3.1: 18] FORM A REVIEW BOARD TO APPROVE AND MAINTAIN SECURE DESIGN PATTERNS.

A review board formalizes the process of reaching and maintaining consensus on security tradeoffs in design needs. Unlike a typical architecture committee focused on functions, this group focuses on providing security guidance, preferably in the form of patterns, standards, features, or frameworks. It also periodically reviews already published design guidance (especially around authentication, authorization, and cryptography) to ensure that design decisions don't become stale or out of date. This review board helps control the chaos associated with adoption of new technologies when development groups might otherwise make decisions on their own without engaging the SSG or champions. Review board security guidance can also serve to inform outsourced software providers about security expectations (see [CP3.2]).

[SFD3.2: 21] REQUIRE USE OF APPROVED SECURITY FEATURES AND FRAMEWORKS.

Implementers must take their security features and frameworks from an approved list or repository (see [SFD1.1], [SFD2.1], [SFD3.1]). There are two benefits to this activity—developers don't spend time reinventing existing capabilities, and review teams don't have to contend with finding the same old defects in new projects or when new platforms are adopted. Reusing proven components eases testing, code review, and threat modeling (see [AA1.1]). Reuse is a major advantage of consistent software architecture and is particularly helpful for Agile development and velocity maintenance in CI/CD pipelines. Packaging and applying required components, such as via containerization (see [SE2.5]), makes it especially easy to reuse approved features and frameworks.

[SFD3.3: 12] FIND AND PUBLISH SECURE DESIGN PATTERNS FROM THE ORGANIZATION.

Foster centralized design reuse by collecting secure design patterns (sometimes referred to as security blueprints) from across the organization and publishing them for everyone to use. A section of the SSG website (see [SR1.2]) could promote positive elements identified during threat modeling or architecture analysis so that good ideas spread widely. This process is formalized—an ad hoc, accidental noticing isn't sufficient. Common design patterns accelerate development, so it's important to use secure design patterns, and not just for applications but for all software assets (e.g., microservices, APIs, containers, infrastructure, and automation).

INTELLIGENCE: STANDARDS & REQUIREMENTS (SR)

The Standards & Requirements practice involves eliciting explicit software security requirements from the organization, determining which COTS tools to recommend, building

standards for major security controls (such as authentication and input validation), creating security standards for technologies in use, and creating a standards review process.

[SRI.1: 84] CREATE SECURITY STANDARDS.

The organization meets the demand for security guidance by creating standards that explain the required way to adhere to policy and carry out security-centric design, development, and operations. A standard might mandate how to perform identity-based application authentication or how to implement transport-level security, perhaps with the SSG ensuring the availability of a reference implementation. Standards often apply to software beyond the scope of an application's code, including container construction, orchestration, infrastructure-as-code, and cloud security configuration. Standards can be deployed in a variety of ways to keep them actionable and relevant. For example, they can be automated into development environments (such as an IDE or toolchain) or explicitly linked to code examples and deployment artifacts (e.g., containers). In any case, to be considered standards, they must be adopted and enforced. Standards for technology stacks [SR3.4] and standards for incorporating new technologies [SR3.5] can be expected to aid in the creation of these standards but are not required.



[SRI.2: 96] CREATE A SECURITY PORTAL.

The organization has a well-known central location for information about software security. Typically, this is an internal website maintained by the SSG and security champions that people refer to for current information on security policies, standards, and requirements, as well as for other resources (such as training). An interactive portal is better than a static portal with guideline documents that rarely change. Organizations often supplement these materials with mailing lists, chat channels (see [T2.12]), and face-to-face meetings. Development teams are increasingly putting software security knowledge directly into toolchains and automation that are outside the organization (e.g., GitHub), but that does not remove the need for SSG-led knowledge management.

[SRI.3: 86] TRANSLATE COMPLIANCE CONSTRAINTS TO REQUIREMENTS.

Compliance constraints are translated into security requirements for individual projects and communicated to the engineering teams. This is a linchpin in the organization's compliance strategy—by representing compliance constraints explicitly with requirements and informing stakeholders, the organization demonstrates that compliance is a manageable task. For example, if the organization builds software that processes credit card transactions, PCI DSS compliance plays a role during the security requirements phase. In other cases, technology standards built for international interoperability can include security guidance on compliance needs. Representing these standards as requirements also helps with traceability and visibility in the event of an audit. It's particularly useful to codify the requirements into reusable code (see [SFD2.1]) or artifact deployment specifications (see [SE1.4]).

[SRI.5: 96] IDENTIFY OPEN SOURCE.

Identify open source components and dependencies included in the organization's code repositories and built software, then review them to understand their security posture. Organizations use a variety of tools and metadata provided by delivery pipelines to discover old versions of open source components with known vulnerabilities or that their software relies on multiple versions of the same component. Scale efforts by using automated tools to find open source, whether whole components or perhaps large chunks of borrowed code. Some software development pipeline platforms, container registries, and middleware platforms have begun to provide this visibility as metadata (e.g., SBOMs [SE3.6]) resulting from behind-the-scenes artifact scanning. Some organizations combine composition analysis results from multiple phases of the software lifecycle to get a more complete and accurate list of the open source being included in production software.

[SR2.2: 70] CREATE A STANDARDS REVIEW PROCESS.

Create a process to develop software security standards and ensure that all stakeholders have a chance to weigh in. This review process could operate by appointing a spokesperson for any proposed security standard, putting the onus on the person to demonstrate that the standard meets its goals and to get buy-in and approval from stakeholders. Enterprise architecture or enterprise risk groups sometimes take on the responsibility of creating and managing standards review processes. When the standards are implemented directly as software, the responsible person might be a DevOps manager, release engineer, or whoever owns the associated deployment artifact (e.g., the orchestration code). Common triggers for standards review processes include periodic updates, security incidents, major vulnerabilities discovered, adoption of new technologies, acquisition, etc.

[SR2.5: 62] CREATE SLA BOILERPLATE.

The SSG works with the legal department to create standard SLA boilerplate for use in contracts with vendors and outsource providers, including cloud providers, to require software security efforts on their part. The legal department might also leverage the boilerplate to help prevent compliance and privacy problems. Under the agreement, vendors and outsource providers must meet company-mandated software security SLAs (see [CP2.4]). Boilerplate language might call for objective third-party insight into software security efforts, such as SSDF gap analysis (<https://csrc.nist.gov/Projects/ssdf>), BSIMMsc measurements, or BSIMM scores.

[SR2.7: 55] CONTROL OPEN SOURCE RISK.

The organization has control over its exposure to the risks that come along with using open source components and all the involved dependencies, including dependencies integrated at runtime. Controlling exposure usually includes multiple efforts, with one example being responding to known vulnerabilities in identified open source (see [SR1.5]). The use of open source could also be restricted to predefined projects or to a short list of versions that have been through an approved security screening process, have had unacceptable vulnerabilities

remediated, and are made available only through approved internal repositories and containers. For some use cases, policy might preclude any use of open source. The legal department often spearheads additional open source controls due to license compliance objectives and the viral license problem associated with GPL code. SSGs that partner with and educate the legal department can help move an organization to improve its open source risk management practices, which must be applied across the software portfolio to be effective.

[SR3.2: 19] COMMUNICATE STANDARDS TO VENDORS.

Work with vendors to educate them and promote the organization's security standards. A healthy relationship with a vendor often starts with contract language (see [CP2.4]), but the SSG should engage with vendors, discuss vendor security practices, and explain in simple terms (rather than legalese) what the organization expects. Any time a vendor adopts the organization's security standards, it's a clear sign of progress. Note that standards implemented as security features or infrastructure configuration could be a requirement to services integration with a vendor (see [SFD1.1], [SE1.4]). When the firm's SSDL is publicly available, communication regarding software security expectations is easier. Likewise, sharing internal practices and measures can make expectations clear.

[SR3.3: 17] USE SECURE CODING STANDARDS.

Developers use secure coding standards to avoid the most obvious bugs and as ground rules for code review. These standards are necessarily specific to a programming language, and they can address the use of popular frameworks, APIs, libraries, and infrastructure automation. Secure coding standards can also be for low- or no-code platforms (e.g., Microsoft Power Apps, Salesforce Lightning). While enforcement isn't the point at this stage (see [CR3.5]), violation of standards is a teachable moment for all stakeholders. Other useful coding standards topics include proper use of cloud APIs, use of approved cryptography, memory sanitization, banned functions, open source use, and many others. If the organization already has coding standards for other purposes (e.g., style), its secure coding standards should build upon them. A clear set of secure coding standards is a good way to guide both manual and automated code review, as well as to provide relevant examples for security training. Some groups might choose to integrate their secure coding standards directly into automation. Socializing the benefits of following standards is also a good first step to gaining widespread acceptance (see [SM2.7]).

[SR3.4: 20] CREATE STANDARDS FOR TECHNOLOGY STACKS.

The organization standardizes on the use of specific technology stacks, which translates into a reduced workload because teams don't have to explore new technology risks for every new project. The organization might create a secure base configuration (commonly in the form of golden images, Terraform definitions, etc.) for each technology stack, further reducing the amount of work required to use the stack safely. In cloud environments, hardened configurations likely include up-to-date security patches, configurations, and

services, such as logging and monitoring. In traditional on-premises IT deployments, a stack might include an operating system, a database, an application server, and a runtime environment (e.g., a MEAN stack). Standards for secure use of reusable technologies, such as containers, microservices, or orchestration code, means that getting security right in one place positively impacts the security posture of all downstream efforts (see [SE2.5]).

[SR3.5: 0] CREATE STANDARDS CONTROLLING AND GUIDING THE ADOPTION OF NEW TECHNOLOGIES.



The SSG is involved in efforts to provide internal practices for technologies so new that industry best practices have not yet been

codified. Involving the SSG in exploration efforts to understand

and plan for new technology minimizes the negative impacts that insecure implementations will have by proactively accounting for potential security pitfalls. The SSG's involvement can result in updates to policies and standards [SR1.1], new security requirements for technology stacks [SR3.4], secure-by-design components and services [SFD2.1, SFD3.2], or coding guidelines [SR3.3]. The SSG must be involved in proactive efforts surrounding the adoption of new technologies rather than merely retroactively securing existing integrations [SFD2.2] or updating policy and standards in response to changing regulations [CP1.1] or emerging threat intelligence [AM1.5].

This effort helps control the chaos associated with adoption of new technologies (such as the rise of AI and LLMs) when development groups might otherwise make decisions on their own without engaging the SSG or champions. It is all about ensuring that security is considered from the beginning instead of having to be bolted on after the fact.

SDLC TOUCHPOINTS

SDLC TOUCHPOINTS: ARCHITECTURE ANALYSIS (AA)

Architecture analysis encompasses capturing software architecture in concise diagrams, applying lists of risks and threats, adopting a process for review (such as Microsoft Threat Modeling [STRIDE] or Architecture Risk Analysis [ARA]), building an assessment and remediation plan for the organization, and using a risk methodology to rank applications.

**TOP
10**

[AA1.1: 99] PERFORM SECURITY FEATURE REVIEW.

Security-aware reviewers identify application security features, review these features against application security requirements and runtime parameters, and determine if each feature can adequately perform its intended function—usually collectively referred to as threat modeling. The goal is to quickly identify missing security features and requirements, or bad deployment configuration (authentication, access control, use of cryptography, etc.), and address them. For example, threat modeling would identify both a system that was subject to escalation of privilege attacks because of broken access control as well as a mobile application that incorrectly puts PII in local storage. Use of the firm's secure-by-design components often streamlines this process (see [SFD2.1]). Many modern applications are no longer simply "3-tier" but instead involve components architected to interact across a variety of tiers—browser/endpoint, embedded, web, microservices, orchestration engines, deployment pipelines, third-party SaaS, etc. Some of these environments might provide robust security feature sets, whereas others might have key capability gaps that require careful analysis, so organizations should consider the applicability and correct use of security features across all tiers that constitute the architecture and operational environment.

[AA1.2: 56] PERFORM DESIGN REVIEW FOR HIGH-RISK APPLICATIONS.

Perform a design review to determine whether the security features and deployment configuration are resistant to attack in an attempt to break the design. The goal is to extend the more formulaic approach of a security feature review (see [AA1.1]) to model application behavior in the context of real-world attackers and attacks. Reviewers must have some experience beyond simple threat modeling to include performing detailed design reviews and breaking the design under consideration. Rather than security feature guidance, a design review should produce a set of flaws and a plan to mitigate them. An organization can use consultants to do this work, but it should participate actively. A review focused only on whether a software project has performed the right process steps won't generate useful results about flaws. Note that a sufficiently robust design review process can't be executed at CI/CD speed, so organizations should focus on a few high-risk applications to start (see [AA1.4]).

[AA1.4: 55] USE A RISK METHODOLOGY TO RANK APPLICATIONS.

Use a defined risk methodology to collect information about each application in order to assign a risk classification and associated prioritization. It is important to use this information in prioritizing what applications or projects are in scope for testing, including security feature and design reviews. Information collection can be implemented via questionnaire or similar method, whether manual or automated. Information needed for classification might include, "Which programming languages is the application written in?" or "Who uses the application?" or "Is the application's deployment software-orchestrated?" Typically, a qualified member of the application team provides the information, but the process should be short enough to take only a few minutes. The SSG can then use the answers to categorize the application as, e.g., high, medium, or low risk. Because a risk questionnaire can be easy to game, it's important to put into place some spot-checking for validity and accuracy—an overreliance on self-reporting can render this activity useless.

[AA2.1: 37] PERFORM ARCHITECTURE ANALYSIS USING A DEFINED PROCESS.

Define and use a process for AA that extends the design review (see [AA1.2]) to also document business risk in addition to technical flaws. The goal is to identify application design flaws as well as the associated risk (e.g., impact of exploitation), such as through frequency or probability analysis, to more completely inform stakeholder risk management efforts. The AA process includes a standardized approach for thinking about attacks, vulnerabilities, and various security properties. The process is defined well enough that people outside the SSG can carry it out. It's important to document both the architecture under review and any security flaws uncovered, as well as risk information that people can understand and use. Microsoft Threat Modeling, Versprite PASTA, and Black Duck ARA are examples of such a process, although these will likely need to be tailored to a given environment. In some cases, performing AA and documenting business risk is done by different teams working together in a single process. Uncalibrated or ad hoc AA approaches don't count as a defined process.

[AA2.2: 38] STANDARDIZE ARCHITECTURAL DESCRIPTIONS.

Threat modeling, design review, or AA processes use an agreed upon format (e.g., diagramming language and icons, not simply a text description) to describe architecture, including a means for representing data flow. Standardizing architecture descriptions between those who generate the models and those who analyze and annotate them makes analysis more tractable and scalable. High-level network diagrams, data flow, and authorization flows are always useful, but the model should also go into detail about how the software itself is structured. A standard architecture description can be enhanced to provide an explicit picture of information assets that require protection, including useful metadata. Standardized icons that are consistently used in diagrams, templates, and dry-erase board squiggles are especially useful, too.

[AA2.4: 40] HAVE SSG LEAD DESIGN REVIEW EFFORTS.

The SSG takes a lead role in performing design review (see [AA1.2]) to uncover flaws. Breaking down an architecture is enough of an art that the SSG, or other reviewers outside the application team, must be proficient, and proficiency requires practice. This practice might then enable, e.g., champions to take the day-to-day lead while the SSG maintains leadership around knowledge and process. The SSG can't be successful on its own either—it will likely need help from architects or implementers to understand the design. With a clear design in hand, the SSG might be able to carry out a detailed review with a minimum of interaction with the project team. Approaches to design review evolve over time, so don't expect to set a process and use it forever. Outsourcing design review might be necessary, but it's also an opportunity to participate and learn.

[AA3.1: 20] HAVE ENGINEERING TEAMS LEAD AA PROCESS.

Engineering teams lead AA to uncover technical flaws and document business risk. This effort requires a well-understood and well-documented process (see [AA2.1]). But even with a good process, consistency is difficult to attain because breaking architecture requires experience, so provide architects with SSG or outside expertise in an advisory capacity. Engineering teams performing AA might normally have responsibilities such as development, DevOps, cloud security, operations security, security architecture, or a variety of similar roles. The process is more useful if the AA team is different from the design team.

[AA3.2: 8] DRIVE ANALYSIS RESULTS INTO STANDARD DESIGN PATTERNS.

Failures identified during threat modeling, design review, or AA are fed back to security and engineering teams so that similar mistakes can be prevented in the future through improved design patterns, whether local to a team or formally approved for everyone (see [SFD3.1]). This typically requires a root-cause analysis process that determines the origin of security flaws, searches for what should have prevented the flaw, and makes the necessary improvements in documented security design patterns. Note that security design patterns can interact in surprising ways that break security, so apply analysis processes even when vetted design patterns are in standard use. For cloud services, providers have learned a lot about how their platforms and services fail to resist attack and have codified this experience into patterns for secure use. Organizations that heavily rely on these services might base their application-layer patterns on those building blocks provided by the cloud service provider (for example, AWS CloudFormation and Azure Blueprints) when making their own.

[AA3.3: 18] MAKE THE SSG AVAILABLE AS AN AA RESOURCE OR MENTOR.

To build organizational AA capability, the SSG advertises experts as resources or mentors for teams using the AA process (see [AA2.1]). This effort might enable, e.g., security champions, site reliability engineers, DevSecOps engineers, and others to take the lead while the SSG offers advice. As one example, mentors help tailor AA process inputs (such as design or attack patterns) to make them more actionable for specific technology stacks. This reusable guidance helps protect the team's time so they can focus on the problems that require creative solutions rather than enumerating known bad habits. While the SSG might answer AA questions during office hours (see [T2.12]), they will often assign a mentor to work with a team, perhaps comprising both security-aware engineers and risk analysts, for the duration of the analysis. In the case of high-risk software, the SSG should play a more active mentorship role in applying the AA process.

SDLC TOUCHPOINTS: CODE REVIEW (CR)

The Code Review practice includes use of code review tools (e.g., static analysis), development of tailored rules, customized profiles for tool use by different roles (e.g., developers vs. auditors), manual analysis, and tracking and measuring results.

[CRI.2: 80] PERFORM OPPORTUNISTIC CODE REVIEW.

Perform code review for high-risk applications in an opportunistic fashion. For example, organizations can follow up a design review with a code review looking for security issues in source code and dependencies and perhaps also in deployment artifact configuration (e.g., containers) and automation metadata (e.g., infrastructure-as-code). This informal targeting often evolves into a systematic approach (see [CR1.4]). Manual code review could be augmented with the use of specific tools and services, but it has to be part of a proactive process. When new technologies pop up, new approaches to code review might become necessary.

**TOP
10**

[CRI.4: 106] USE AUTOMATED CODE REVIEW TOOLS.

Incorporate static analysis into the code review process to make the review more efficient and consistent. Automation won't replace human judgment, but it does bring definition to the review process and security expertise to reviewers who typically aren't security experts. Note that a specific tool might not cover an entire portfolio, especially when new languages are involved, so additional local effort might be useful. Some organizations might progress to automating tool use by instrumenting static analysis into source code management workflows (e.g., pull requests) and delivery pipeline workflows (build, package, and deploy) to make the review more efficient, consistent, and aligned with release cadence. Whether use of automated tools is to review a portion of the source code incrementally, such as a developer committing new code or small changes, or to conduct full analysis by scanning the entire codebase, this service should be explicitly connected to a larger SSDL defect management process applied during software development. This effort is not useful when done just to "check the security box" on the path to deployment.

AI

[CRI.5: 75] MAKE CODE REVIEW MANDATORY FOR ALL PROJECTS.

A security-focused code review is mandatory for all software projects, with a lack of code review or unacceptable results stopping a release, slowing it down, or causing it to be recalled. While all projects must undergo code review, the process might be different for different kinds of projects. The review for low-risk projects might rely more heavily on automation (see [CR1.4]), for example, whereas high-risk projects might have no upper bound on the amount of time spent by reviewers. Having a minimum acceptable standard forces projects that don't pass to be fixed and reevaluated. A code review tool with nearly all the rules turned off (so it can run at CI/ CD automation speeds, for example) won't provide

sufficient defect coverage. Similarly, peer code review or tools focused on quality and style won't provide useful security results.

[CRI.7: 51] ASSIGN CODE REVIEW TOOL MENTORS.

Mentors show developers how to get the most out of code review tools, including configuration, triage, and remediation. Security champions, DevOps and site reliability engineers, and SSG members often make good mentors. Mentors could use office hours or other outreach to help developers establish the right configuration and get started on interpreting and remediating results. Alternatively, mentors might work with a development team for the duration of the first review they perform. Centralized use of a tool can be distributed into the development organization or toolchains over time through the use of tool mentors, but providing installation instructions and URLs to centralized tool downloads isn't the same as mentoring. Increasingly, mentorship extends to code review tools associated with deployment artifacts (e.g., container security) and infrastructure (e.g., cloud configuration). While AI is becoming useful to augment human code review guidance, it likely doesn't have the context necessary to replace it.

[CR2.6: 24] USE CUSTOM RULES WITH AUTOMATED CODE REVIEW TOOLS.

Create and use custom rules in code review tools to help uncover security defects specific to the organization's coding standards or to the framework-based or cloud-provided middleware the organization uses. The same group that provides tool mentoring (see [CR1.7]) will likely spearhead this customization. Custom rules are often explicitly tied to proper usage of technology stacks in a positive sense and avoidance of errors commonly encountered in a firm's codebase in a negative sense. Custom rules are also an easy way to check for adherence to coding standards (see [CR3.5]). To reduce the workload for everyone, many organizations also create rules to remove repeated false positives and to turn off checks that aren't relevant.

[CR2.7: 19] USE A TOP N BUGS LIST (REAL DATA PREFERRED).

Maintain a living list of the most important kinds of bugs the organization wants to eliminate from its code and use it to drive change. Many organizations start with a generic list pulled from public sources, but broad-based lists such as the OWASP Top 10 rarely reflect an organization's bug priorities. Build a valuable list by using real data gathered from code review (see [CR2.8]), testing (see [PT1.2]), software composition analysis (see [SE3.8]), and actual incidents (see [CMVM1.1]), then prioritize it for prevention efforts. Simply sorting the day's bug data by number of occurrences won't produce a satisfactory list because the data changes so often. To increase interest, the SSG can periodically publish a "most wanted" report after updating the list. One potential pitfall with a top N list is that it tends to include only known problems. Of course, just building the list won't accomplish anything—everyone has to use it to find and fix bugs.

[CR2.8: 27] USE CENTRALIZED DEFECT REPORTING TO CLOSE THE KNOWLEDGE LOOP.

The defects found during code review are tracked in a centralized repository that makes it possible to do both summary and trend reporting for the organization. Reported defects drive engineering improvements such as enhancing processes, updating standards, adopting reusable frameworks, etc. For example, code review information is usually incorporated into a CISO-level dashboard that can include feeds from other security testing efforts (e.g., penetration testing, composition analysis, threat modeling). Given the historical code review data, the SSG can also use the reports to demonstrate progress (see [SM3.3]) or drive the training curriculum. Individual bugs make excellent training examples (see [T2.8]). Some organizations have moved toward analyzing this data and using the results to drive automation (see [ST3.6]).

[CR3.2: 16] BUILD A CAPABILITY TO COMBINE AST RESULTS.

Combine application security testing (AST) results so that multiple testing techniques feed into one reporting and remediation process. In addition to code review, testing techniques often include dynamic analysis, software composition analysis, container scanning, cloud services configuration review, etc. The SSG might write scripts or acquire software to gather data automatically and combine the results into a format that can be consumed by a single downstream review and reporting solution. The tricky part of this activity is normalizing vulnerability information from disparate sources that might use conflicting terminology or scoring. In some cases, using a standardized taxonomy (e.g., a CWE-like approach) can help with normalization. Combining multiple sources helps drive better-informed risk mitigation decisions and identify engineering improvements.

[CR3.3: 6] CREATE CAPABILITY TO ERADICATE BUGS.

When a security bug is found during code review (see [CR1.2], [CR1.4]), the organization searches for and then fixes all occurrences of the bug, not just the instance originally discovered. Searching with custom rules (see [CR2.6]) makes it possible to eradicate the specific bug entirely without waiting for every project to reach the code review portion of its lifecycle. This doesn't mean finding every instance of every kind of cross-site scripting bug when a specific example is found—it means going after that specific example everywhere. A firm with only a handful of software applications built on a single technology stack will have an easier time with this activity than firms with many large applications built on a diverse set of technology stacks. A new development framework or library, rules in RASP or a next-generation firewall, or cloud configuration tools that provide guardrails can often help in (but not replace) eradication efforts.

[CR3.4: 3] AUTOMATE MALICIOUS CODE DETECTION.

Use automated code review to identify malicious code written by in-house developers or outsource providers. Examples of malicious code include backdoors, logic bombs, time bombs, nefarious communication channels, obfuscated program logic, and dynamic code injection. Although out-of-the-box automation might identify some generic malicious-looking constructs, custom rules for the static analysis tools used to codify acceptable and unacceptable patterns in the organization's codebase will likely become a necessity. Manual review for malicious code is a good start but insufficient to complete this activity at scale. While not all backdoors or similar code were meant to be malicious when they were written (e.g., a developer's feature to bypass authentication during testing), such things tend to stay in deployed code and should be treated as malicious until proven otherwise. Discovering some types of malicious code will require dynamic testing techniques.

[CR3.5: 6] ENFORCE SECURE CODING STANDARDS.

A violation of secure coding standards is sufficient grounds for rejecting a piece of code. This rejection can take one or more forms, such as denying a pull request, breaking a build, failing quality assurance, removing from production, or moving the code into a different development workstream where repairs or exceptions can be worked out. The enforced portions of an organization's secure coding standards (see [SR3.3]) often start out as a simple list of banned functions or required frameworks. Code review against standards must be objective—it shouldn't become a debate about whether the noncompliant code is exploitable. In some cases, coding standards are specific to language constructs and enforced with tools (e.g., codified into SAST rules). In other cases, published coding standards are specific to technology stacks and enforced during the code review process or by using automation. Standards can be positive ("do it this way") or negative ("do not use this API"), but they must be enforced.

SDLC TOUCHPOINTS: SECURITY TESTING (ST)

The Security Testing practice is concerned with prerelease defect discovery as well as integrating security into standard QA processes. The practice includes the use of opaque-box AST tools (including fuzz testing) as a smoke test in QA, risk-driven crystal-box test suites, application of the attack model, and code coverage analysis. Security testing focuses on vulnerabilities in construction.

**TOP
10**

[ST1.1: 102] PERFORM EDGE/ BOUNDARY VALUE CONDITION TESTING DURING QA.

QA efforts go beyond functional testing to perform basic adversarial tests and probe simple edge cases and boundary conditions, with no particular attacker skills required. When QA pushes past standard functional testing that uses expected input, it begins to move toward thinking like

an adversary. Boundary value testing, whether automated or manual, can lead naturally to the notion of an attacker probing the edges on purpose (e.g., determining what happens when someone enters the wrong password over and over).

[ST1.3: 79] DRIVE TESTS WITH SECURITY REQUIREMENTS AND SECURITY FEATURES.

QA targets declarative security mechanisms with tests derived from security requirements and features. A test could try to access administrative functionality as an unprivileged user, for example, or verify that a user account becomes locked after some number of failed authentication attempts. For the most part, security features can be tested in a fashion similar to other software features—security mechanisms such as account lockout, transaction limitations, entitlements, etc., are tested with both expected and unexpected input as derived from security requirements. Software security isn't security software, but testing security features is an easy way to get started. New software architectures and deployment automation, such as with container and cloud infrastructure orchestration, might require novel test approaches.

[ST1.4: 54] INTEGRATE OPAQUE-BOX SECURITY TOOLS INTO THE QA PROCESS.

The organization uses one or more opaque-box security testing tools as part of the QA process. Such tools are valuable because they encapsulate an attacker's perspective, albeit generically. Traditional dynamic analysis scanners are relevant for web applications, while similar tools exist for cloud environments, containers, mobile applications, embedded systems, APIs, etc. In some situations, other groups might collaborate with the SSG to apply the tools. For example, a testing team could run the tool but come to the SSG for help with interpreting the results. When testing is integrated into Agile development approaches, opaque-box tools might be hooked into internal toolchains, provided by cloud-based toolchains, or used directly by engineering. Regardless of who runs the opaque-box tool, the testing should be properly integrated into a QA cycle of the SSDL and will often include both authenticated and unauthenticated reviews.

[ST2.4: 20] DRIVE QA TESTS WITH AST RESULTS.

Share results from application security testing, such as penetration testing, threat modeling, composition analysis, code reviews, etc., with QA teams to evangelize the security mindset. Using security defects as the basis for a conversation about common attack patterns or the underlying causes for them allows QA teams to generalize this information into new test approaches. Organizations that leverage software pipeline platforms such as GitHub, or CI/ CD platforms such as the Atlassian stack, can benefit from teams receiving various testing results automatically, which should then facilitate timely stakeholder conversations—emailing security reports to QA teams will not generate the desired results. Over time, QA teams learn the security mindset, and the organization benefits from an improved ability to create security tests tailored to the organization's code.

[ST2.5: 34] INCLUDE SECURITY TESTS IN QA AUTOMATION.

Security tests are included in an automation framework and run alongside functional, performance, and other QA test suites. Executing this automation framework can be triggered manually or through additional automation (e.g., as part of pipeline tooling). When test creators who understand the software create security tests, they can uncover more specialized or more relevant defects than commercial tools might (see [ST1.4]). Security tests might be derived from typical failures of security features (see [SFD1.1]), from creative tweaks of functional and developer tests, or even from guidance provided by penetration testers on how to reproduce an issue. Tests that are performed manually or out-of-band likely will not provide timely feedback.

[ST2.6: 24] PERFORM FUZZ TESTING CUSTOMIZED TO APPLICATION APIS.

QA efforts include running a customized fuzzing framework against APIs critical to the organization. An API might be software that allows two applications to communicate or even software that allows a human to interact with an application (e.g., a webform). Testers could begin from scratch or use an existing fuzzing toolkit, but the necessary customization often goes beyond creating custom protocol descriptions or file format templates to giving the fuzzing framework a built-in understanding of application interfaces and business logic. Test harnesses developed explicitly for specific applications make good places to integrate fuzz testing.

[ST3.3: 16] DRIVE TESTS WITH DESIGN REVIEW RESULTS.

Use design review or architecture analysis results to direct QA test creation. For example, if the results of attempting to break a design determine that “the security of the system hinges on the transactions being atomic and not being interrupted partway through,” then torn transactions will become a primary target in adversarial testing. Adversarial tests like these can be developed according to a risk profile, with high-risk flaws at the top of the list. Security defect data shared with QA (see [ST2.4]) can help focus test creation on areas of potential vulnerability that can, in turn, help prove the existence of identified high-risk flaws.

[ST3.4: 5] LEVERAGE CODE COVERAGE ANALYSIS.

Testers measure the code coverage of their application security testing to identify code that isn't being exercised and then adjust test cases to incrementally improve coverage. AST can include automated testing (see [ST2.5], [ST2.6]) and manual testing (see [ST1.1], [ST1.3]). In turn, code coverage analysis drives increased security testing depth. Coverage analysis is easier when using standard measurements, such as function coverage, line coverage, or multiple condition coverage. The point is to measure how broadly the test cases cover the security requirements, which is not the same as measuring how broadly the test cases exercise the code.

[ST3.5: 6] BEGIN TO BUILD AND APPLY ADVERSARIAL SECURITY TESTS (ABUSE CASES).

QA teams incorporate test cases based on abuse cases (see [AM2.1]) as testers move beyond verifying functionality and take on the attacker's perspective. One way to do this is to systematically attempt to replicate incidents from the organization's history. Abuse and misuse cases based on the attacker's perspective can also be derived from security policies, attack intelligence, standards, and the organization's top N attacks list (see [AM3.5]). This effort turns the corner in QA from testing features to attempting to break the software under test.

[ST3.6: 8] IMPLEMENT EVENT-DRIVEN SECURITY TESTING IN AUTOMATION.

The SSG guides implementation of automation for continuous, event-driven application security testing. An event here is simply a noteworthy occurrence, such as dropping new code in a repository, a pull request, a build request, or a push to deployment. Event-driven testing implemented in pipeline automation (rather than testing only in production) typically moves the testing closer to the conditions driving the testing requirement (whether shift left toward design or shift right toward operations), repeats the testing as often as the event is triggered, and helps ensure that the right testing is executed for a given set of conditions. Success with this approach depends on the broad use of sensors (e.g., agents, bots) that monitor engineering processes, execute contextual rules, and provide telemetry to automation that initiates the specified testing whenever event conditions are met. More mature configurations typically include risk driven conditions (e.g., size of change, provenance, function, team).

DEPLOYMENT

DEPLOYMENT: PENETRATION TESTING (PT)

The Penetration Testing practice involves standard outside-in testing of the sort carried out by security specialists. Penetration testing focuses on vulnerabilities in preproduction and production code, providing direct feeds to defect management and mitigation.

**TOP
10**

[PTI.I: 104] USE EXTERNAL PENETRATION TESTERS TO FIND PROBLEMS.

External penetration testers are used to demonstrate that the organization's software needs help. Finding critical vulnerabilities in high-profile applications provides the evidence that executives often require. Over time, the focus of penetration testing moves from trying to determine if the code is broken in some areas to a sanity check done before shipping or on a periodic basis. In addition to breaking code, this sanity check can also be an effective way to ensure that vulnerability prevention techniques are both used and effective. External penetration testers who bring a new set of experiences and skills to the problem are the most useful.

[PTI.2: 95] FEED RESULTS TO THE DEFECT MANAGEMENT AND MITIGATION SYSTEM.

All penetration testing results are fed back to engineering through established defect management or mitigation channels, with development and operations responding via a defect management and release process. In addition to application vulnerabilities, also track results from testing other software such as containers and infrastructure configuration. Properly done, this exercise demonstrates the organization's ability to improve the state of security and emphasizes the importance of not just identifying but actually fixing security problems. One way to ensure attention is to add a security flag to the bug-tracking and defect management system. The organization might leverage developer workflow or social tooling (e.g., JIRA or Slack) to communicate change requests, but these requests are still tracked explicitly as part of a vulnerability management process.

[PTI.3: 76] USE PENETRATION TESTING TOOLS INTERNALLY.

The organization creates an internal penetration testing capability that uses tools as part of an established process. Execution can rest with the SSG or be part of a specialized team elsewhere in the organization, with the tools complementing manual efforts to improve the efficiency and repeatability of the testing process. The tools used will usually include off-the-shelf products built specifically for application penetration testing, network penetration tools that specifically understand the application layer, container and cloud configuration testing tools, and custom scripts. Free-time or crisis-driven efforts aren't the same as an internal capability.

[PT2.2: 39] PENETRATION TESTERS USE ALL AVAILABLE INFORMATION.

Penetration testers, whether internal or external, routinely make use of artifacts created throughout the SSDL to do more comprehensive analysis and find more problems. Example artifacts include design documents, architecture analysis results, misuse and abuse cases, code review results, cloud environment and other deployment configurations, and source code if applicable. Focusing on high-risk applications is a good way to start. Note that having access to SSDL artifacts is not the same as using them.

[PT2.3: 51] SCHEDULE PERIODIC PENETRATION TESTS FOR APPLICATION COVERAGE.

All applications are tested periodically, which could be tied to a calendar or a release cycle. High-risk applications could get a penetration test at least once per year, for example, even if there have not been substantive code changes, while other applications might receive different kinds of security testing on a similar schedule. Any security testing performed must focus on discovering vulnerabilities, not just checking a process or compliance box. This testing serves as a sanity check and helps ensure that yesterday's software isn't vulnerable to today's attacks. The testing can also help maintain the security of software configurations and environments, especially for containers and components in the cloud. One important aspect

of periodic security testing across the portfolio is to make sure that the problems identified are actually fixed. Software that isn't an application, such as automation created for CI/CD, infrastructure-as-code, etc., deserves some security testing as well.

[PT3.1: 26] USE EXTERNAL PENETRATION TESTERS TO PERFORM DEEP-DIVE ANALYSIS.

The SSG uses external penetration testers to do a deep-dive analysis on critical software systems or technologies and to introduce fresh thinking. One way to do this is to simulate persistent attackers using goal-oriented red team exercises. These testers are domain experts and specialists who keep the organization up to speed with the latest version of the attacker's perspective and have a track record for breaking the type of software being tested. When attacking the organization's software, these testers demonstrate a creative approach that provides useful knowledge to the people designing, implementing, and hardening new systems. Creating new types of attacks from threat intelligence and abuse cases typically requires extended timelines, which is essential when it comes to new technologies, and prevents checklist-driven approaches that look only for known types of problems.

[PT3.2: 21] CUSTOMIZE PENETRATION TESTING TOOLS.

Build a capability to create penetration testing tools, or to adapt publicly available ones, to attack the organization's software more efficiently and comprehensively. Creating penetration testing tools requires a deep understanding of attacks (see [AM2.1], [AM2.8]) and technology stacks (see [AM3.4]). Customizing existing tools goes beyond configuration changes and extends tool functionality to find new issues. Tools will improve the efficiency of the penetration testing process without sacrificing the depth of problems that the SSG can identify. Automation can be particularly valuable in organizations using Agile methodologies because it helps teams go faster. Tools that can be tailored are always preferable to generic tools. Success here is often dependent on both the depth and scope of tests enabled through customized tools.

DEPLOYMENT: SOFTWARE ENVIRONMENT (SE)

The Software Environment practice deals with OS and platform patching (including in the cloud), WAFs (web application firewalls), installation and configuration documentation, containerization, orchestration, application monitoring, change management, and code signing.

[SEI.1: 76] USE APPLICATION INPUT MONITORING FOR SECURITY PURPOSES.

The organization monitors input to the software that it runs in order to spot attacks. Monitoring systems that write log files are useful only if humans or bots periodically review the logs and take action. For web applications, RASP or a WAF can do this monitoring, while other kinds of software likely require other approaches, such as custom runtime instrumentation. Software and technology stacks, such as mobile and IoT, likely

require their own input monitoring solutions. Serverless and containerized software can require interaction with vendor software to get the appropriate logs and monitoring data. Cloud deployments and platform-as-a-service usage can add another level of difficulty to the monitoring, collection, and aggregation approach.

**TOP
10**

[SEI.2: 102] ENSURE HOST AND NETWORK SECURITY BASICS ARE IN PLACE.

The organization provides a solid foundation for its software in operation by ensuring that host (whether bare metal or virtual machine) and network security basics are in place across its data centers and networks and that these basics remain in place during new releases. Host and network security basics must account for evolving network perimeters, increased connectivity and data sharing, software-defined networking, and increasing dependence on vendors (e.g., content delivery, load balancing, and content inspection services). In addition to securing your production environment, the organization should consider securing their development endpoints [SE3.10] and tool chains [SE3.9]. Doing software security before getting host and network security in place is like putting on shoes before putting on socks.

[SEI.3: 87] IMPLEMENT CLOUD SECURITY CONTROLS.

Organizations ensure that cloud security controls are in place and working for both public and private clouds. Industry best practices are a good starting point for local policy and standards to drive controls and configurations. Of course, cloud-based assets often have public-facing services that create an attack surface (e.g., cloud-based storage) that is different from the one in a private data center, so these assets require customized security configuration and administration. In the increasingly software-defined world, the SSG has to help everyone explicitly configure cloud-specific security features and controls (e.g., through cloud provider administration consoles) comparable to those built with cables and physical hardware in private data centers. Detailed knowledge about cloud provider shared responsibility security models is always necessary to ensure that the right cloud security controls remain in place.

[SEI.4: 74] DEFINE SECURE DEPLOYMENT PARAMETERS AND CONFIGURATIONS.

Create deployment automation or installation guides (e.g., standard operating procedures) to help teams and customers install and configure software securely. Software here includes applications, products, scripts, images, firmware, and other forms of code. Deployment automation usually includes a clearly described configuration for software artifacts and the infrastructure-as-code (e.g., Terraform, CloudFormation, ARM templates, Helm Charts) necessary to deploy them, including details on COTS, open source, vendor, and cloud services components. All deployment automation should be understandable by humans, not just by machines, especially when distributed to customers. Where deployment automation is not applicable, customers or deployment teams need installation guides that include hardening guidance and secure configurations.

[SE2.4: 5I] PROTECT CODE INTEGRITY.

Use code protection mechanisms (e.g., code signing) that allow the organization to attest to the provenance, integrity, and authorization of important code. While legacy and mobile platforms accomplished this with point-in-time code signing and permissions activity, protecting modern containerized software demands actions in various lifecycle phases. Organizations can use build systems to verify sources and manifests of dependencies, creating their own cryptographic attestation of both. Packaging and deployment systems can sign and verify binary packages, including code, configuration, metadata, code identity, and authorization to release material. In some cases, organizations allow only code from their own registries to execute in certain environments. Protecting code integrity can also include securing development infrastructure, using permissions and peer review to govern code contributions, and limiting code access to help protect integrity (see [SE3.9]).

[SE2.5: 64] USE APPLICATION CONTAINERS TO SUPPORT SECURITY GOALS.

The organization uses application containers to support its software security goals. Simply deploying containers isn't sufficient to gain security benefits, but their planned use can support a tighter coupling of applications with their dependencies, immutability, integrity (see [SE2.4]), and some isolation benefits without the overhead of deploying a full operating system on a virtual machine. Containers are a convenient place for security controls to be applied and updated consistently (see [SFD3.2]), and while they are useful in development and test environments, their use in production provides the needed security benefits.

[SE2.7: 42] USE ORCHESTRATION FOR CONTAINERS AND VIRTUALIZED ENVIRONMENTS.

The organization uses automation to scale service, container, and virtualized environments in a disciplined way. Orchestration processes take advantage of built-in and add-on security features (see [SFD2.1]), such as hardening against drift, secrets management, RBAC, and rollbacks, to ensure that each deployed workload meets predetermined security requirements. Setting security behaviors in aggregate allows for rapid change when the need arises. Orchestration platforms are themselves software that becomes part of your production environment, which in turn requires hardening and security patching and configuration—in other words, if you use Kubernetes, make sure you patch Kubernetes.

[SE3.2: 24] USE CODE PROTECTION.

To protect intellectual property and make exploit development harder, the organization erects barriers to reverse engineering its software (e.g., anti-tamper, debug protection, anti-piracy features, runtime integrity). For some software, obfuscation techniques could be applied as part of the production build and release process. In other cases, these protections could be applied at the software-defined network or software orchestration layer when applications are being dynamically regenerated post-deployment. Code protection is particularly

important for widely distributed code, such as mobile applications and JavaScript distributed to browsers. On some platforms, employing Data Execution Prevention (DEP), Safe Structured Exception Handling (SafeSEH), and Address Space Layout Randomization (ASLR) can be a good start at making exploit development more difficult, but be aware that yesterday's protection mechanisms might not hold up to today's attacks.

[SE3.3: 17] USE APPLICATION BEHAVIOR MONITORING AND DIAGNOSTICS.

The organization monitors production software to look for misbehavior or signs of attack. Go beyond host and network monitoring to look for software-specific problems, such as indications of malicious behavior, fraud, and related issues. Application-level intrusion detection and anomaly detection systems might focus on an application's interaction with the operating system (through system calls) or with the kinds of data that an application consumes, originates, and manipulates. Signs that an application isn't behaving as expected will be specific to the software business logic and its environment, so one-size-fits-all solutions probably won't generate satisfactory results. In some types of environments (e.g., platform-as-a-service), some of this data and the associated predictive analytics might come from a vendor.

[SE3.6: 25] CREATE BILLS OF MATERIALS FOR DEPLOYED SOFTWARE.

Create a BOM detailing the components, dependencies, and other metadata for important production software. Use this BOM to help the organization tighten its security posture, i.e., to react with agility as attackers and attacks evolve, compliance requirements change, and the number of items to patch grows quite large. Knowing where all the components live in running software—and whether they're in private data centers, in clouds, or sold as box products (see [CMVM2.3])—allows for timely response when unfortunate events occur.

[SE3.8: 3] PERFORM APPLICATION COMPOSITION ANALYSIS ON CODE REPOSITORIES.

Use composition analysis results to augment software asset inventory information with data on all components comprising important applications. Beyond open source (see [SR1.5]), inventory information (see [SM3.1]) includes component and dependency information for internally developed (first-party), commissioned code (second-party), and external (third-party) software, whether that software exists as source code or binary. One common way of documenting this information is to build SBOMs. Doing this manually is probably not an option—keeping up with software changes likely requires toolchain integration rather than carrying this out as a point-in-time activity. This information is extremely useful in supply chain security efforts (see [SM3.5]).

[SE3.9: 3] PROTECT INTEGRITY OF DEVELOPMENT TOOLCHAINS.

The organization ensures the integrity of software it builds and integrates by maintaining and securing all development

infrastructure and preventing unauthorized changes to source code and other software lifecycle artifacts. Development infrastructure includes code and artifact repositories, build pipelines, and deployment automation. Secure the development infrastructure by safely handling and storing secrets, following pipeline configuration requirements, patching tools and build environments, limiting access to pipeline settings, and auditing changes to configurations. Preventing unauthorized changes typically includes enforcing least privilege access to code repositories and requiring approval for code commits. Automatically granting access for all project team members isn't sufficient to adequately protect software integrity.



[SE3.I0: 0] PROTECT THE INTEGRITY OF DEVELOPMENT ENDPOINTS.

The organization maintains the integrity of the software it builds by applying security basics to the workstations used by development stakeholders who interact with the development toolchain. Development endpoints are the workstations used for writing source code, configuring the development toolchain, testing the software's functionality, or modifying data in the code or artifact repositories. Organizations can protect development endpoints by limiting or monitoring privileged actions, ensuring that the operating system and anti-virus definitions are up to date, vetting installed software, or by providing a separate, secured workstation for development that is not used for administrative tasks. Establishing and applying a development endpoint security baseline allows for stakeholders to perform the technical tasks required by software development, but also provides another layer of defense to the development toolchain [SE3.9].

DEPLOYMENT: CONFIGURATION MANAGEMENT & VULNERABILITY MANAGEMENT (CMVM)

The Configuration Management & Vulnerability Management practice concerns itself with operations processes, patching and updating applications, version control, defect tracking and remediation, and incident handling.



[CMVMI.1: III] CREATE OR INTERFACE WITH INCIDENT RESPONSE.

The SSG is prepared to respond to an event or alert and is regularly included in the incident response process, either by creating its own incident response capability or by regularly interfacing with the organization's existing team. A standing meeting between the SSG and the incident response team keeps information flowing in both directions. Having prebuilt communication channels with critical vendors (e.g., ISP, monitoring, IaaS, SaaS, PaaS) is also very important.

[CMVMI.2: 85] IDENTIFY SOFTWARE DEFECTS FOUND IN OPERATIONS MONITORING AND FEED THEM BACK TO ENGINEERING.

Defects identified in production through operations monitoring are fed back to development and used to change engineering behavior. Useful sources of production defects include

incidents, bug bounty (see [CMVM3.4]), responsible disclosure (see [CMVM2.4]), SIEMs, production logs, customer feedback, and telemetry from cloud security posture monitoring, container configuration monitoring, RASP, and similar technologies. Entering production defect data into an existing bug-tracking system (perhaps by making use of a special security flag) can close the information loop and make sure that security issues get fixed. In addition, it's important to capture lessons learned from production defects and use these lessons to change the organization's behavior. In the best of cases, processes in the SSDL can be improved based on operations data (see [CMVM3.2]).

[CMVMI.3: 89] TRACK SOFTWARE DEFECTS FOUND IN OPERATIONS THROUGH THE FIX PROCESS.

Defects found in operations (see [CMVM1.2]) are entered into established defect management systems and tracked through the fix process. This tracking ability could come in the form of a two-way bridge between defect finders and defect fixers or possibly through intermediaries (e.g., the vulnerability management team), but make sure the loop is closed completely. Defects can appear in all types of deployable artifacts, deployment automation, and infrastructure configuration. Setting a security flag in the defect tracking system can help facilitate tracking.

[CMVMI.4: 88] HAVE EMERGENCY RESPONSE.

The organization can make quick code and configuration changes when software (e.g., application, API, microservice, infrastructure) is under attack. An emergency response team works in conjunction with stakeholders such as application owners, engineering, operations, and the SSG to study the code and the attack, find a resolution, and fix the production code (e.g., push a patch into production, rollback to a known-good state, deploy a new container). Often, the emergency response team is the engineering team itself. A well-defined process is a must here, a process that has never been used might not actually work.

[CMVM2.3: 46] DEVELOP AN OPERATIONS SOFTWARE INVENTORY.

The organization has a map of its software deployments and related containerization, orchestration, and deployment automation code, along with the respective owners. If a software asset needs to be changed or decommissioned, operations or DevOps teams can reliably identify both the stakeholders and all the places where the change needs to occur. Common components can be noted so that when an error occurs in one application, other applications sharing the same components can be fixed as well. Building an accurate representation of an inventory will likely involve enumerating at least the source code, the open source incorporated both during the build and during dynamic production updates, the orchestration software incorporated into production images, and any service discovery or invocation that occurs in production.

[CMVM2.4: 4I] STREAMLINE INCOMING RESPONSIBLE VULNERABILITY DISCLOSURE.

Provide external bug reporters with a line of communication to internal security experts through a low-friction, public entry point. These experts work with bug reporters to invoke any necessary organizational responses and to coordinate with external entities throughout the defect management lifecycle. Successful disclosure processes require insight from internal stakeholders, such as legal, marketing, and public relations roles, to simplify and expedite decision-making during software security crises (see [CMVM3.3]). Although bug bounties might be important to motivate some researchers (see [CMVM3.4]), proper public attribution and a low-friction reporting process is often sufficient motivation for researchers to participate in a coordinated disclosure. Most organizations will use a combination of easy-to-find landing pages, common email addresses (security@), and embedded product documentation when appropriate (security.txt) as an entry point for external reporters to invoke this process.

[CMVM3.1: 13] FIX ALL OCCURRENCES OF SOFTWARE DEFECTS FOUND IN OPERATIONS.

When a security defect is found in operations (see [CMVM1.2]), the organization searches for and fixes all occurrences of the defect in operations, not just the one originally reported. Doing this proactively requires the ability to reexamine the entire operations software inventory (see [CMVM2.3]) when new kinds of defects come to light. One way to approach reexamination is to create a ruleset that generalizes deployed defects into something that can be scanned for via automated code review. In some environments, addressing a defect might involve removing it from production immediately and making the actual fix in some priority order before redeployment. Use of orchestration can greatly simplify deploying the fix for all occurrences of a software defect (see [SE2.7]).

[CMVM3.2: 24] ENHANCE THE SSDL TO PREVENT SOFTWARE DEFECTS FOUND IN OPERATIONS.

Experience from operations leads to changes in the SSDL (see [SM1.1]), which can in turn be strengthened to prevent the reintroduction of defects. To make this process systematic, the incident response postmortem includes a feedback-to-SSDL step. The outcomes of the postmortem might result in changes such as to tool-based policy rulesets in a CI/CD pipeline and adjustments to automated deployment configuration (see [SE1.4]). This works best when root-cause analysis pinpoints where in the software lifecycle an error could have been introduced or slipped by uncaught (e.g., a defect escape). DevOps engineers might have an easier time with this because all the players are likely involved in the discussion and the solution. An ad hoc approach to SSDL improvement isn't sufficient for prevention.

[CMVM3.3: 22] SIMULATE SOFTWARE CRISES.

The SSG simulates high-impact software security crises to ensure that software incident detection and response capabilities minimize damage. Simulations could test for

the ability to identify and mitigate specific threats or could begin with the assumption that a critical system or service is already compromised and evaluate the organization's ability to respond. Planned chaos engineering can be effective at triggering unexpected conditions during simulations. The exercises must include attacks or other software security crises at the appropriate software layer to generate useful results (e.g., at the application layer for web applications and at lower layers for IoT devices). When simulations model successful attacks, an important question to consider is the time required for cleanup. Regardless, simulations must focus on security-relevant software failure, not on natural disasters or other types of emergency response drills. Organizations that are highly dependent on vendor infrastructure (e.g., cloud service providers, SaaS, PaaS) and security features will naturally include those things in crisis simulations.

[CMVM3.4: 3I] OPERATE A BUG BOUNTY PROGRAM.

The organization solicits vulnerability reports from external researchers and pays a bounty for each verified and accepted vulnerability received. Payouts typically follow a sliding scale linked to multiple factors, such as vulnerability type (e.g., remote code execution is worth \$10,000 vs. CSRF is worth \$750), exploitability (demonstrable exploits command much higher payouts), or specific service and software versions (widely deployed or critical services warrant higher payouts). Ad hoc or short-duration activities, such as capture-the-flag contests or informal crowdsourced efforts, don't constitute a bug bounty program.

[CMVM3.5: 17] AUTOMATE VERIFICATION OF OPERATIONAL INFRASTRUCTURE SECURITY.

The SSG works with engineering teams to verify with automation the security properties (e.g., adherence to agreed-upon security hardening) of infrastructure generated from controlled self-service processes. Engineers use self-service processes to create networks, storage, containers, and machine instances, to orchestrate deployments, and to perform other tasks that were once IT's sole responsibility. In facilitating verification, the organization uses machine-readable policies and configuration standards (see [SE1.4]) to automatically detect issues and report on infrastructure that does not meet expectations. In some cases, the automation makes changes to running environments to bring them into compliance, but in many cases, organizations use a single policy to manage automation in different environments, such as in multi- and hybrid-cloud environments.

[CMVM3.6: 4] PUBLISH RISK DATA FOR DEPLOYABLE ARTIFACTS.


The organization collects and publishes risk information about the applications, services, APIs, containers, and other software it deploys. Whether captured through manual processes or telemetry automation, published information extends beyond basic software security (see [SM2.1]) and inventory data (see [CMVM2.3]) to include risk information. This information usually includes constituency of the software (e.g., BOMs [SE3.6]), provenance data about what group created it and how, and the

risks associated with known vulnerabilities, deployment models, security controls, or other security characteristics intrinsic to each artifact. This approach stimulates cross-functional coordination and helps stakeholders take informed risk management action. Making a list of risks that aren't used for decision support won't achieve useful results.


[CMVM3.8: I] DO ATTACK SURFACE MANAGEMENT FOR DEPLOYED APPLICATIONS.

Operations standards and procedures proactively minimize application attack surfaces by using attack intelligence and application weakness data to limit vulnerable conditions. Finding and fixing software defects in operations is important (see [CMVM1.2]) but so is finding and fixing errors in cloud security models, VPNs, segmentation, security configurations

for networks, hosts, and applications, etc., to limit the ability to successfully attack deployed applications. Combining attack intelligence (see [AM1.5]) with information about software assets (see [AM2.9]) and a continuous view of application weaknesses helps ensure that attack surface management keeps pace with attacker methods. SBOMs (see [SE3.6]) are also an important information source when doing attack surface management in a crisis.



PART 9: WHAT CAN THE DATA TELL US



PART 9: WHAT CAN THE DATA TELL US

DATA ANALYSIS: SSG

SSGs are the primary implementers of an SSI, responsible for governance, enablement, productivity, and continuous growth. You can use this information to put your SSI and SSG on a growth path.

This section analyzes how SSIs evolve over time by analyzing SSG age, SSG score, and other relevant data.

SSG CHARACTERISTICS

As our BSIMM participants changed, we added a greater number of firms with newer SSIs and began to track new verticals that have less software security experience (see Table 15 DATA ANALYSIS: VERTICALS AND PRACTICES in Part 7). Thus, we expected a decrease in participant scores, which is easily seen in Figure 19 for BSIMM7 through BSIMM8.

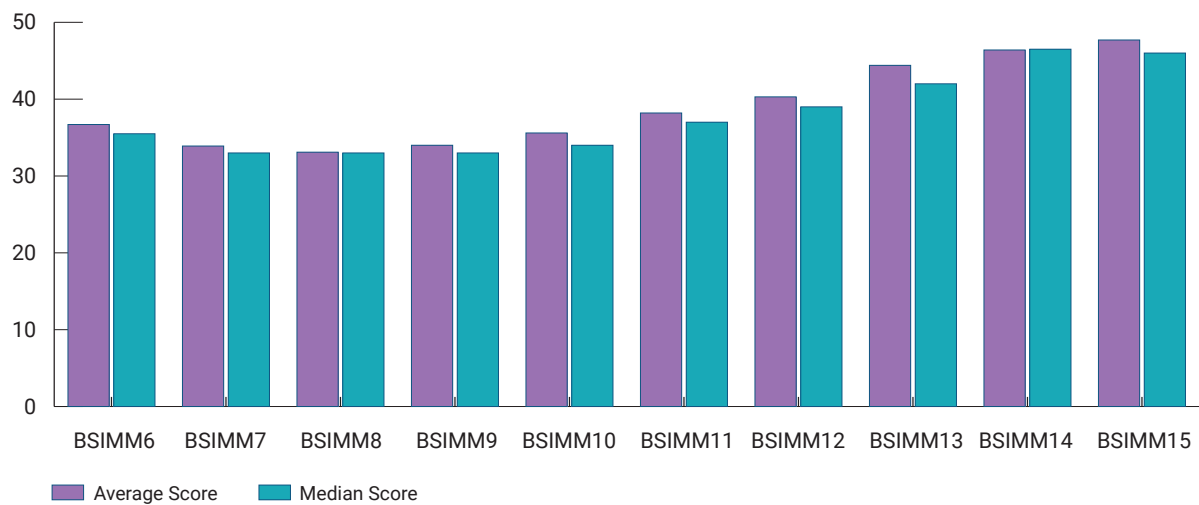


FIGURE 19. AVERAGE BSIMM PARTICIPANT SCORE. Adding firms with less experience decreased the average score from BSIMM7 through BSIMM8, even as remeasurements have shown that individual firm maturity increases over time.

In BSIMM9, the average and the median scores started to increase. We saw the largest increase in BSIMM13 when the average and median scores increased by 4.1 and 3, respectively. One reason for this change in average data pool score appears to be the mix of firms using the BSIMM as part of their SSI journey. For example, Figure 20 shows how the SSG age of firms entering the BSIMM data pool changed over time. In BSIMM15, and in concert with the increase in average scores seen for BSIMM13 and BSIMM14 in Figure 20, we saw a significantly higher average and median SSG age of new firms vs. what was seen in previous years.




FIGURE 20. AVERAGE AND MEDIAN SSG AGE FOR NEW FIRMS ENTERING THE BSIMM DATA POOL. The median SSG age of firms entering BSIMM7 through BSIMM8 was declining and so did the average BSIMM score, while outliers in BSIMM7 and BSIMM8 resulted in a high average SSG age. Starting with BSIMM9, the median age of firms entering the BSIMM was higher again, which tracks with the increase of average BSIMM scores.

A second reason appears to be firms continuing to use the BSIMM to guide their initiatives. Firms using the BSIMM as an ongoing measurement tool are likely also making sufficient improvements to justify the ongoing creation of SSI scorecards. See DATA ANALYSIS: LONGITUDINAL in Part 9 for more details on how SSIs evolve as seen through remeasurement data.




FIGURE 21. NUMBER OF FIRMS AGED OUT OF THE BSIMM DATA POOL.

Given their importance to overall SSI efforts, we also closely monitor champions trends. Many firms with no champions continue to exist in the data pool, which causes the median champions size to be 10 (50 of 121 firms had no champions at the time of their current assessment); 57% of the 23 firms added for BSIMM15 had no champions at assessment time, as seen in Figure 22.



FIGURE 22. STATISTICS FOR FIRMS WITH AND WITHOUT CHAMPIONS. This data appears to validate the notion that having more people, both centralized and distributed into engineering teams, helps SSIs achieve higher scores. For the 71 BSIMM15 firms with champions at last assessment time, the average champions size was 105 with a median of 40 (not shown). We present the average and median SSG size to remove the impact of a few significant outliers.

SSG CHANGES BASED ON AGE

We've mentioned a trend that older SSIs generally achieve higher scores, and we showed this trend in Figure 17. Here, we analyze the data in more detail to identify additional trends related to SSG age.




FIGURE 23. AVERAGE TOTAL SCORES BY AGE. By notionally organizing SSIs into emerging, maturing, and enabling phases by age in years, we see a steady growth in score as SSIs mature.

For this analysis, we put the 121 BSIMM15 SSIs into six groups based on SSG age. Figure 23 shows the trend discussed earlier: the older the SSI, the higher its BSIMM score. While the journey through emerging, maturing, and enabling phases is not a straight line (see Part 3), here we equate the emerging phase with the first two bars from the left (0-1.5 and 1.5-3.0 years of age), maturing phase with the next two bars, and enabling phase with the last two.

While Figure 23 provides a low-resolution view into how SSIs change with SSG age, the following five figures increase the resolution and compare the normalized spiders for SSIs organized by their age. Figure 24 shows, on a percentage scale, how the SSI is changing through its emerging phase. The purple line shows what the program looks like when SSIs are initially organizing themselves and discovering what activities are already happening in the organization. At this point in the journey, we typically see a relatively high effort in Compliance & Policy, Standards & Requirements, Penetration Testing, and Configuration Management & Vulnerability Management. Likely, these efforts are already in place due to compliance obligations, an existing cybersecurity program and its focus on standards, and quick wins in defect discovery by leveraging penetration testing.

Over the next 18 months (teal line), SSIs build some capability around documenting and socializing the SSDL, publishing and promoting the process, and defect discovery for high-priority applications. The differences between two spiders in Strategy & Metrics, Compliance & Policy, Security Features & Design, Standards & Requirements, and Architecture Analysis result from these efforts. Firms also increase their investment in Penetration Testing and securing their Software Environment.




FIGURE 24. COMPARING EMERGING SSGs. As emerging SSGs move from initial discovery steps (purple line) toward defining and rolling out the program (teal line), they invest in Strategy & Metrics, Compliance & Policy, Standards & Requirements, Penetration Testing, and Architecture Analysis. This tracks with recommendations in Part 3 on how to start an SSI, where almost 45% of all recommended activities in Figure 5 are from these four practices.

As SSGs move toward the maturing phase, they start focusing on improving the efficiency, effectiveness, and scale of existing efforts, see the “FOR A MATURING SSI: Harmonizing Objectives” section of Part 3. This push typically involves getting more value out of existing activities rather than doing more activities. Figure 25 shows the difference in normalized spiders for organizations toward the end of their emerging phase (teal line) and the beginning of their maturing phase (dark blue line).

The lack of any large differences between the spiders in Figure 25 shows that firms at this stage focus on tweaking the existing program as they improve scale, efficiency, and effectiveness—changes are often an investment in automation (such as code reviews). As shown in the diagram, when these SSGs look to improve scale and efficiency, they appear to have less time for manual efforts in the Architecture Analysis practice.



FIGURE 25. COMPARING LATE EMERGING TO EARLY MATURING SSGs. As firms move from emerging to maturing, the average score increase is relatively small. This aligns with our qualitative observations in Part 3 that these firms often focus more on the scale, efficiency, and effectiveness of existing activities in their SSGs vs. working on implementing new activities.

As SSGs move toward the end of their maturing phase, they start investing again in improving policies, standards, requirements, processes, metrics, and evangelism as shown by significant differences in the spiders in Figure 26. The increase in observation rates in the Strategy & Metrics, Compliance & Policy, and Standards & Requirements practices demonstrate this trend. If firms are not careful, they can begin to lose focus on some security activities like we see in Training and Code Review in Figure 26.




FIGURE 26. COMPARING MATURING SSIs. As firms move toward the end of their maturing journey, SSGs start focusing again on implementing new activities resulting in some incremental improvements.

Some factors specific to verticals might significantly influence the overall shape of the spiders. For example, 20% of firms with an SSG age between six and nine years are in the Insurance vertical as compared to 6.6% in the entire BSIMM15 data pool. Similarly, 27% of the firms with an SSG age above nine years are in the Financial vertical versus an average of 15% across the other age brackets. As we analyze the next two figures, we keep these facts in mind. Refer to DATA ANALYSIS: VERTICALS AND PRACTICES here in Part 9 for more analysis of how the verticals compare to each other.

One potential explanation for the dip in Security Testing shown in Figure 27 is that the Insurance vertical has some of the lowest observation rates for this practice, and Insurance is overrepresented in this age bracket, 20%, while overall insurance represents just 6.6% of the firms. For the spike in Penetration Testing, we saw strong growth in level 2 and level 3 Penetration Testing activities between the 4.5 – 6 age bracket and the 6 – 9 age bracket across all industries as firms increasingly adopted enabling activities like [PT2.2] **penetration testers use all available information**, [PT2.3] **schedule periodic penetration tests for application coverage**, [PT1.1] **use external penetration testers to perform deep-dive analysis**, and [PT3.2] **customize penetration testing tools**. Outside of the outliers mentioned above, SSIs gradually increase their effort in all other practices as they start their enabling journey.




FIGURE 27. COMPARING LATE MATURING TO EARLY ENABLING SSIs. As firms move from the maturing to the enabling stage, SSIs continue to invest in Compliance & Policy. Overall, this comparative growth aligns with concepts such as putting “Sec” in DevOps as well as scaling outreach and expertise, which are discussed in the “Enabling SSIs” section of Part 3.

In Figure 28, we see significant gains in all practices, with the smallest gain in Penetration Testing, which spiked in the 4.5 – 6 age range. There are large gains in the Security Features & Design practice largely driven by strong gains in the enabling level 3 activities [SFD3.1] **form a review board to approve and maintain secure design patterns**, [SFD3.2] **require use of approved security features and frameworks**, and [SFD3.3] **find and publish secure design patterns from the organization**, all of which can enhance the security of a firm's software and speed up the development process at the same time. Attack Models and Standards & Requirements both saw strong growth among the level 2 and level 3 activities as SSGs are able to expand their reach.




FIGURE 28. COMPARING ENABLING SSGs. As SSGs continue their enabling phase, they invest significant effort in reusable and pre-baked security controls (e.g., from the Security Features & Design practice) and learning from the attacker's perspective (e.g., from the Attack Models practice). In fact, the increase in observation rate of activities in Security Features & Design is the highest increase in observation rates among all practices across all age buckets.

DATA ANALYSIS: SECURITY CHAMPIONS

A security champions program allows an SSI and SSG to scale their reach throughout the organization and harmonize everyone's approach to software security. You can use this information to help justify your own outreach program.

A security champions program is an organized effort to deputize members of the development community into being software security leaders for their geographies, application teams, or technology groups. Once they are inducted into the program, the SSI provides the champions with training, support, and the access needed to answer security questions.

A security champions program is an effective way to address the people and culture portions of the people, process, technology, and culture view of an SSI's scope. Firms typically rely on their security champions to lead the ground-level security push among developers, architects, QA, operations, and other stakeholders such as cloud and site reliability. A strong security champions program enables an SSI to scale people-driven activities, tune automated activities, and prioritize remediation tracking activities within an organization. In Figure 29, the orange line shows that firms can achieve higher scores even with a lower ratio of SSG to developers (e.g., the bottom 20%

have an average SSG-to-developer ratio of 3.8 while the top 20% have an average SSG-to-developer ratio of 2.1). One way these firms are able to scale is by increasing the ratio of champions to developers, as shown by the teal bars (e.g., the bottom 20% have an average champion-to-developer ratio of 5.5 while the top 20% have an average champions-to-developer ratio of 14.6).




FIGURE 29. AVERAGE RATIO OF SSG AND CHAMPIONS SIZE TO DEVELOPERS FOR THREE SCORE BUCKETS. There is a strong correlation between security champions' support and overall BSIMM score (scale on the right).

While the presence of a champions program doesn't guarantee a high number of activity observations, there is a correlation that appears when grouping BSIMM firms by scores. Nearly 90% of firms in the highest scoring group have a champions program as compared to 32% in the lowest scoring group. Figure 30 shows the score increases from an average of 22.4 activities in the lowest scoring group (shown on the orange line), up to an average of 76.9 activities in the highest scoring group (shown here as the top 20%).




FIGURE 30. PERCENTAGE OF FIRMS THAT HAVE A CHAMPIONS PROGRAM, ORGANIZED IN THREE BUCKETS BY BSIMM SCORE. Presence of a champions program and average score (scale on the right) appear to be correlated, but we don't have enough data to say which is the cause and which is the effect. Here we see, for example, that in the bottom scoring 20% (about eight firms) of the 73 (of 121) firms with champions, the average score was just over 20 compared to an average score of over 75 for the top scoring 20% with champions.

When separating firms into groups with and without champions, the activity observation rate increases in nearly every practice (see Figure 31). While the biggest differences between the two spiders are in Strategy & Metrics, Training, Architecture Analysis, and Standards & Requirements, the firms with champions also spend consistently more effort on defect discovery in the Code Review, Security Testing, and Penetration Testing practices.




FIGURE 31. COMPARING FIRMS WITH AND WITHOUT CHAMPIONS. The presence of a champions program seems to correlate strongly with an increase in program maturity as evidenced by increased scores by practice on a percentage scale.

Figure 32 shows that as SSIs get older, they have higher average scores and are more likely to have a champions program, so is the presence of champions the reason for higher scores or the consequence of older SSIs? One way to answer this question is to look at the average ratio of SSG size to number of developers, shown in Figure 29, which might indicate that there is a correlation between SSI reach and the size of the security champions team.




FIGURE 32. BSIMM SCORE DISTRIBUTION RELATIVE TO CHAMPIONS SIZE AND SSG AGE. Older SSIs (light purple line) not only tend to have a higher BSIMM score (buckets 0-20, 21-30, etc.), they are also more likely to have a champions program (orange line).

90% of firms in the highest scoring group have a champions program, compared to 32% in the lowest scoring group.

Eighty percent of the 45 BSIMM15 firms that have been assessed more than once have a champions program, while 49% of the firms on their first assessment do not. Many firms that are new to software security take some time to identify and develop a champions program. This data suggests that as an SSI matures, its activities become distributed and institutionalized into the organizational structure, perhaps even into engineering automation as well, requiring an expanded champions program to provide expertise and be the local voice of the SSG.

DATA ANALYSIS: VERTICALS AND PRACTICES

While every company is a software company these days, there are differences in SSI implementation. You can use this information on how vertical markets approach software security to inform your own strategy.

An important use of the BSIMM data is to help everyone see how different groups of organizations approach the implementation of software security activities. Do certain groups focus more on governance than testing? Or perhaps architecture and secure-by-design components vs. operational maintenance? What about training? Or vendor management? While it seems true that “every company is becoming a software company,” different verticals still have their own priorities. The BSIMM data helps us to observe and analyze this.

BSIMM15 currently reports eight verticals where there is sufficient data to keep the data reasonably anonymous:

- Financial
- Financial Technology (FinTech)
- Independent Software Vendor (ISV)
- Technology (Tech)

- Healthcare
- Internet of Things (IoT)
- Cloud
- Insurance

Table 15 shows how the representation of different verticals has grown and evolved over the history of the BSIMM. Financial, ISV, and technology firms were early adopters of the BSIMM. Several firms fall into more than one vertical, so the numbers in Table 15 do not add up to the 121 firms in the BSIMM15 data pool, nor to the number of firms in the previous BSIMM data pools. There are several other verticals we track, and when there are sufficient firms in those verticals, we will begin to or resume reporting on those verticals.

An important use of the BSIMM data helps everyone see how different organizations approach implementing software security activities.

| BSIMM VERTICAL PARTICIPANTS OVER TIME | | | | | | | | |
|---------------------------------------|-----------|---------|-----|------|------------|--------------------|-------|-----------|
| | FINANCIAL | FINTECH | ISV | TECH | HEALTHCARE | INTERNET OF THINGS | CLOUD | INSURANCE |
| BSIMM15 | 35 | 9 | 32 | 43 | 9 | 19 | 26 | 14 |
| BSIMM14 | 43 | 12 | 33 | 39 | 10 | 21 | 32 | 15 |
| BSIMM13 | 44 | 15 | 38 | 33 | 11 | 19 | 35 | 15 |
| BSIMM12 | 38 | 21 | 42 | 28 | 14 | 18 | 26 | 13 |
| BSIMM11 | 42 | 21 | 46 | 27 | 14 | 17 | 30 | 14 |
| BSIMM10 | 57 | | 43 | 20 | 16 | 13 | 20 | 11 |
| BSIMM9 | 50 | | 42 | 22 | 19 | 16 | 17 | 10 |
| BSIMM8 | 47 | | 38 | 16 | 17 | 12 | 16 | 11 |
| BSIMM7 | 42 | | 30 | 14 | 15 | 12 | 15 | 10 |
| BSIMM6 | 33 | | 27 | 17 | 10 | 13 | | |
| BSIMM-V | 26 | | 25 | 14 | | | | |
| BSIMM4 | 19 | | 19 | 13 | | | | |
| BSIMM3 | 17 | | 15 | 10 | | | | |
| BSIMM2 | 12 | | 7 | 7 | | | | |
| BSIMM1 | 4 | | 4 | 2 | | | | |

TABLE 15. BSIMM VERTICALS OVER TIME. The BSIMM data pool has grown over the years as shown by growth in vertical representation. Remember that a firm can appear in more than one vertical. Note also that FinTech became a separate vertical from Financial in BSIMM11.

VERTICAL SCORECARDS

Figure 33 shows the BSIMM scorecards for the eight verticals compared side by side, allowing for discovery of differences and similarities between verticals. This report includes some new information for the vertical scorecards:


- For each activity per vertical, we present the observation rate as a percentage (e.g., 71% of firms in the Cloud vertical are performing [SM1.1]).
- To show the biggest outliers within each vertical, we highlighted activities where observation rates are either at least 1.75 standard deviations above average (highlighted in teal) or at least 1.75 standard deviations below average (highlighted in gold). Use these highlighted differences to identify apparently higher- and lower-value activities unique to a vertical.
- We also highlighted five activities (see the activity column) with the least differences between verticals (light gold) and five activities with the largest differences between verticals (light teal). The activities in light gold appear to be uniformly applicable across all verticals, while those in teal appear to be more vertical-specific.

The following are observations from Figure 33:

- The five activities with the least variation in observation rate between verticals, not surprisingly, are some of the most common activities in BSIMM15. These are [SM1.4] **implement security checkpoints and associated governance**, [CP1.2] **identify privacy obligations**, [CR1.4] **use automated code review tools**, [PT1.1] **use external penetration testers to find problems**, and [CMVM1.1] **create**

or interface with incident response. This is another indicator that these activities are applicable to all SSIs, independent of what vertical the firm is in.

- Activities [SM3.5] **integrate software supply chain risk management**, [SE3.8] **perform application composition analysis on code repositories**, and [CMVM3.8] **do attack surface management for deployed applications** were all introduced with BSIMM13 and show some of the greatest differences among verticals. All three show significantly higher observation rates among the FinTech companies than other verticals, indicating they are early adopters of these three activities and consider them important to invest in.
- [ST3.5] **Begin to build and apply adversarial security tests (abuse cases)** was an original activity in BSIMM1 as [ST2.3] but became [ST3.5] with BSIMM-V as other activities saw higher observation rates. [CR3.4] **Automate malicious code detection** was introduced in BSIMM4 and has remained a level 3 activity ever since. Both show little investment across most verticals, and most observations are in the Technology vertical. Technology vertical products are generally shipped to their customers and may exist in hostile environments. The Technology vertical's significantly higher investment in these two activities indicate that firms in this vertical value understanding how their products could be abused in potentially hostile environments, and they want to take steps to help ensure no malicious code is delivered to the customer.



| GOVERNANCE | | | | | | | | |
|---------------------|---------------|-------------------|----------------|-------------------|-------------------|-------------|-------------|--------------|
| ACTIVITY | CLOUD (OF 26) | FINANCIAL (OF 35) | FINTECH (OF 9) | HEALTHCARE (OF 9) | INSURANCE (OF 14) | IOT (OF 19) | ISV (OF 32) | TECH (OF 43) |
| STRATEGY & METRICS | | | | | | | | |
| [SM1.1] | 77% | 69% | 78% | 89% | 71% | 100% | 69% | 88% |
| [SM1.3] | 58% | 63% | 44% | 56% | 64% | 53% | 59% | 56% |
| [SM1.4] | 88% | 94% | 89% | 89% | 79% | 95% | 84% | 93% |
| [SM1.7] | 62% | 66% | 67% | 22% | 50% | 58% | 59% | 67% |
| [SM2.1] | 65% | 60% | 67% | 44% | 64% | 47% | 44% | 58% |
| [SM2.3] | 73% | 37% | 89% | 56% | 43% | 68% | 78% | 56% |
| [SM2.6] | 62% | 63% | 67% | 33% | 43% | 58% | 53% | 60% |
| [SM2.7] | 46% | 43% | 44% | 67% | 57% | 37% | 44% | 47% |
| [SM3.1] | 31% | 20% | 33% | 0% | 7% | 37% | 25% | 37% |
| [SM3.2] | 12% | 20% | 22% | 11% | 21% | 37% | 13% | 23% |
| [SM3.3] | 19% | 34% | 44% | 11% | 21% | 32% | 22% | 26% |
| [SM3.4] | 8% | 11% | 11% | 11% | 14% | 5% | 9% | 7% |
| [SM3.5] | 0% | 3% | 11% | 0% | 0% | 0% | 0% | 0% |
| COMPLIANCE & POLICY | | | | | | | | |
| [CP1.1] | 69% | 83% | 78% | 100% | 79% | 84% | 88% | 77% |
| [CP1.2] | 85% | 94% | 100% | 100% | 100% | 100% | 84% | 84% |
| [CP1.3] | 62% | 86% | 67% | 67% | 79% | 84% | 72% | 81% |
| [CP2.1] | 50% | 37% | 56% | 44% | 29% | 53% | 44% | 44% |
| [CP2.2] | 46% | 60% | 33% | 33% | 43% | 63% | 38% | 49% |
| [CP2.3] | 54% | 57% | 56% | 67% | 50% | 53% | 56% | 63% |
| [CP2.4] | 46% | 57% | 44% | 33% | 43% | 42% | 50% | 56% |
| [CP2.5] | 54% | 66% | 56% | 56% | 43% | 47% | 66% | 49% |
| [CP3.1] | 19% | 37% | 44% | 22% | 36% | 26% | 16% | 30% |
| [CP3.2] | 27% | 34% | 22% | 33% | 21% | 37% | 31% | 40% |
| [CP3.3] | 15% | 9% | 11% | 0% | 7% | 26% | 9% | 19% |
| TRAINING | | | | | | | | |
| [T1.1] | 50% | 49% | 44% | 33% | 50% | 68% | 47% | 63% |
| [T1.7] | 54% | 54% | 67% | 44% | 50% | 53% | 50% | 58% |
| [T1.8] | 35% | 54% | 22% | 33% | 50% | 32% | 28% | 49% |
| [T2.5] | 46% | 14% | 44% | 33% | 21% | 42% | 44% | 44% |
| [T2.8] | 23% | 11% | 0% | 22% | 14% | 32% | 19% | 30% |
| [T2.9] | 19% | 29% | 22% | 22% | 29% | 42% | 13% | 37% |
| [T2.10] | 19% | 26% | 33% | 11% | 21% | 26% | 19% | 19% |
| [T2.11] | 19% | 23% | 0% | 33% | 36% | 37% | 19% | 28% |
| [T2.12] | 31% | 34% | 44% | 0% | 21% | 26% | 41% | 35% |
| [T3.1] | 4% | 9% | 22% | 0% | 7% | 0% | 6% | 9% |
| [T3.2] | 4% | 17% | 11% | 11% | 7% | 11% | 3% | 12% |
| [T3.6] | 8% | 9% | 22% | 0% | 0% | 16% | 3% | 12% |

| | |
|--|---|
| | Score is at least 1.75 standard deviations above average |
| | Score is at least 1.75 standard deviations below average |
| | 5 activities with the most difference between verticals |
| | 5 activities with the least differences between verticals |

| INTELLIGENCE | | | | | | | | |
|----------------------------|---------------|-------------------|----------------|-------------------|-------------------|-------------|-------------|--------------|
| ACTIVITY | CLOUD (OF 26) | FINANCIAL (OF 35) | FINTECH (OF 9) | HEALTHCARE (OF 9) | INSURANCE (OF 14) | IOT (OF 19) | ISV (OF 32) | TECH (OF 43) |
| ATTACK MODELS | | | | | | | | |
| [AM1.2] | 46% | 74% | 78% | 100% | 93% | 26% | 38% | 40% |
| [AM1.3] | 27% | 51% | 33% | 56% | 71% | 37% | 19% | 37% |
| [AM1.5] | 58% | 80% | 44% | 100% | 79% | 68% | 44% | 65% |
| [AM2.1] | 15% | 20% | 0% | 11% | 29% | 11% | 6% | 14% |
| [AM2.6] | 15% | 9% | 0% | 0% | 0% | 5% | 6% | 16% |
| [AM2.7] | 15% | 17% | 11% | 11% | 21% | 5% | 16% | 12% |
| [AM2.8] | 27% | 26% | 22% | 11% | 14% | 32% | 13% | 21% |
| [AM2.9] | 19% | 20% | 0% | 0% | 7% | 16% | 19% | 9% |
| [AM3.2] | 8% | 9% | 0% | 11% | 7% | 5% | 3% | 5% |
| [AM3.4] | 4% | 11% | 0% | 11% | 7% | 5% | 3% | 12% |
| [AM3.5] | 4% | 11% | 0% | 11% | 29% | 0% | 6% | 9% |
| SECURITY FEATURES & DESIGN | | | | | | | | |
| [SFD1.1] | 69% | 74% | 89% | 78% | 71% | 68% | 69% | 84% |
| [SFD1.2] | 73% | 69% | 56% | 78% | 71% | 89% | 72% | 72% |
| [SFD2.1] | 38% | 34% | 56% | 22% | 7% | 37% | 28% | 44% |
| [SFD2.2] | 73% | 49% | 56% | 33% | 50% | 79% | 63% | 67% |
| [SFD3.1] | 12% | 20% | 22% | 22% | 14% | 16% | 16% | 14% |
| [SFD3.2] | 12% | 17% | 22% | 11% | 7% | 11% | 16% | 23% |
| [SFD3.3] | 4% | 14% | 22% | 0% | 14% | 11% | 9% | 12% |
| STANDARDS & REQUIREMENTS | | | | | | | | |
| [SR1.1] | 62% | 74% | 56% | 78% | 79% | 74% | 56% | 74% |
| [SR1.2] | 81% | 74% | 67% | 78% | 71% | 95% | 81% | 88% |
| [SR1.3] | 58% | 80% | 89% | 78% | 64% | 79% | 75% | 67% |
| [SR1.5] | 77% | 71% | 100% | 100% | 71% | 89% | 84% | 84% |
| [SR2.2] | 50% | 66% | 44% | 67% | 79% | 58% | 44% | 60% |
| [SR2.5] | 46% | 57% | 56% | 44% | 43% | 53% | 47% | 63% |
| [SR2.7] | 54% | 43% | 89% | 33% | 43% | 37% | 44% | 53% |
| [SR3.2] | 0% | 17% | 0% | 33% | 21% | 26% | 13% | 16% |
| [SR3.3] | 19% | 11% | 22% | 11% | 0% | 21% | 6% | 21% |
| [SR3.4] | 15% | 23% | 22% | 11% | 14% | 21% | 9% | 19% |
| [SR3.5] | 0% | 0% | 0% | 0% | 0% | 0% | 0% | 0% |

| | |
|--|---|
| | Score is at least 1.75 standard deviations above average |
| | Score is at least 1.75 standard deviations below average |
| | 5 activities with the most difference between verticals |
| | 5 activities with the least differences between verticals |

| SDL TOUCHPOINTS | | | | | | | | |
|-----------------------|---------------|-------------------|----------------|-------------------|-------------------|-------------|-------------|--------------|
| ACTIVITY | CLOUD (OF 26) | FINANCIAL (OF 35) | FINTECH (OF 9) | HEALTHCARE (OF 9) | INSURANCE (OF 14) | IOT (OF 19) | ISV (OF 32) | TECH (OF 43) |
| ARCHITECTURE ANALYSIS | | | | | | | | |
| [AA1.1] | 88% | 74% | 89% | 67% | 86% | 89% | 94% | 86% |
| [AA1.2] | 46% | 34% | 33% | 78% | 43% | 68% | 44% | 60% |
| [AA1.4] | 35% | 71% | 67% | 67% | 64% | 32% | 16% | 40% |
| [AA2.1] | 27% | 17% | 22% | 44% | 14% | 63% | 31% | 49% |
| [AA2.2] | 23% | 17% | 22% | 44% | 14% | 58% | 34% | 49% |
| [AA2.4] | 31% | 29% | 33% | 56% | 29% | 42% | 31% | 42% |
| [AA3.1] | 15% | 9% | 11% | 22% | 14% | 21% | 22% | 30% |
| [AA3.2] | 4% | 9% | 0% | 11% | 7% | 16% | 3% | 7% |
| [AA3.3] | 15% | 9% | 11% | 11% | 7% | 21% | 16% | 23% |
| CODE REVIEW | | | | | | | | |
| [CR1.2] | 69% | 69% | 78% | 67% | 57% | 84% | 53% | 65% |
| [CR1.4] | 81% | 89% | 100% | 100% | 86% | 84% | 91% | 88% |
| [CR1.5] | 50% | 60% | 78% | 56% | 57% | 68% | 59% | 72% |
| [CR1.7] | 50% | 34% | 44% | 44% | 36% | 58% | 44% | 51% |
| [CR2.6] | 27% | 17% | 33% | 11% | 7% | 16% | 19% | 23% |
| [CR2.7] | 19% | 11% | 22% | 11% | 14% | 11% | 6% | 26% |
| [CR2.8] | 15% | 34% | 22% | 33% | 21% | 5% | 16% | 16% |
| [CR3.2] | 12% | 17% | 11% | 0% | 0% | 21% | 6% | 19% |
| [CR3.3] | 8% | 3% | 11% | 11% | 7% | 0% | 0% | 7% |
| [CR3.4] | 0% | 3% | 0% | 0% | 0% | 0% | 0% | 5% |
| [CR3.5] | 4% | 6% | 11% | 0% | 0% | 11% | 3% | 7% |
| SECURITY TESTING | | | | | | | | |
| [ST1.1] | 92% | 71% | 78% | 100% | 79% | 95% | 84% | 93% |
| [ST1.3] | 69% | 46% | 67% | 67% | 57% | 84% | 72% | 79% |
| [ST1.4] | 46% | 31% | 78% | 67% | 36% | 58% | 47% | 56% |
| [ST2.4] | 15% | 9% | 22% | 0% | 0% | 21% | 16% | 26% |
| [ST2.5] | 42% | 11% | 33% | 11% | 14% | 32% | 38% | 40% |
| [ST2.6] | 12% | 9% | 11% | 0% | 0% | 42% | 16% | 42% |
| [ST3.3] | 8% | 3% | 0% | 11% | 7% | 26% | 6% | 30% |
| [ST3.4] | 8% | 3% | 11% | 0% | 0% | 0% | 3% | 9% |
| [ST3.5] | 0% | 0% | 0% | 0% | 0% | 0% | 3% | 12% |
| [ST3.6] | 19% | 6% | 22% | 0% | 7% | 0% | 9% | 5% |

| | |
|--|---|
| | Score is at least 1.75 standard deviations above average |
| | Score is at least 1.75 standard deviations below average |
| | 5 activities with the most difference between verticals |
| | 5 activities with the least differences between verticals |

| DEPLOYMENT | | | | | | | | |
|---|---------------|-------------------|----------------|-------------------|-------------------|-------------|-------------|--------------|
| ACTIVITY | CLOUD (OF 26) | FINANCIAL (OF 35) | FINTECH (OF 9) | HEALTHCARE (OF 9) | INSURANCE (OF 14) | IOT (OF 19) | ISV (OF 32) | TECH (OF 43) |
| PENETRATION TESTING | | | | | | | | |
| [PT1.1] | 96% | 89% | 100% | 89% | 86% | 84% | 91% | 79% |
| [PT1.2] | 85% | 74% | 100% | 67% | 57% | 68% | 88% | 77% |
| [PT1.3] | 62% | 69% | 67% | 67% | 71% | 63% | 66% | 53% |
| [PT2.2] | 38% | 31% | 56% | 11% | 14% | 47% | 38% | 40% |
| [PT2.3] | 62% | 54% | 67% | 11% | 43% | 37% | 53% | 33% |
| [PT3.1] | 35% | 26% | 33% | 11% | 7% | 26% | 16% | 28% |
| [PT3.2] | 19% | 17% | 22% | 11% | 14% | 32% | 6% | 28% |
| SOFTWARE ENVIRONMENT | | | | | | | | |
| [SE1.1] | 58% | 83% | 67% | 89% | 79% | 58% | 53% | 44% |
| [SE1.2] | 85% | 89% | 100% | 89% | 100% | 95% | 72% | 88% |
| [SE1.3] | 81% | 83% | 89% | 89% | 93% | 68% | 81% | 53% |
| [SE1.4] | 62% | 66% | 67% | 11% | 50% | 63% | 63% | 67% |
| [SE2.4] | 50% | 11% | 22% | 33% | 7% | 74% | 47% | 74% |
| [SE2.5] | 62% | 63% | 89% | 56% | 64% | 58% | 56% | 42% |
| [SE2.7] | 50% | 40% | 67% | 33% | 36% | 26% | 41% | 23% |
| [SE3.2] | 8% | 11% | 22% | 0% | 0% | 26% | 9% | 40% |
| [SE3.3] | 15% | 23% | 44% | 11% | 14% | 5% | 13% | 5% |
| [SE3.6] | 19% | 11% | 11% | 0% | 0% | 32% | 16% | 40% |
| [SE3.8] | 8% | 3% | 22% | 0% | 0% | 0% | 3% | 0% |
| [SE3.9] | 0% | 3% | 0% | 0% | 0% | 0% | 3% | 2% |
| [SE3.10] | 0% | 0% | 0% | 0% | 0% | 0% | 0% | 0% |
| CONFIGURATION MANAGEMENT & VULNERABILITY MANAGEMENT | | | | | | | | |
| [CMVM1.1] | 85% | 94% | 89% | 89% | 86% | 89% | 91% | 95% |
| [CMVM1.2] | 77% | 69% | 67% | 67% | 64% | 84% | 72% | 72% |
| [CMVM1.3] | 77% | 69% | 78% | 78% | 50% | 84% | 75% | 84% |
| [CMVM1.4] | 69% | 80% | 56% | 67% | 71% | 74% | 78% | 67% |
| [CMVM2.3] | 42% | 43% | 44% | 44% | 36% | 32% | 34% | 37% |
| [CMVM2.4] | 38% | 29% | 22% | 0% | 14% | 58% | 38% | 47% |
| [CMVM3.1] | 12% | 9% | 22% | 0% | 0% | 11% | 13% | 16% |
| [CMVM3.2] | 15% | 23% | 11% | 0% | 7% | 26% | 9% | 30% |
| [CMVM3.3] | 23% | 29% | 33% | 22% | 29% | 11% | 13% | 19% |
| [CMVM3.4] | 31% | 23% | 44% | 11% | 29% | 21% | 28% | 21% |
| [CMVM3.5] | 15% | 20% | 11% | 11% | 14% | 11% | 9% | 14% |
| [CMVM3.6] | 4% | 6% | 11% | 0% | 0% | 5% | 3% | 5% |
| [CMVM3.8] | 0% | 3% | 11% | 0% | 0% | 0% | 0% | 0% |

| | |
|--|---|
| | Score is at least 1.75 standard deviations above average |
| | Score is at least 1.75 standard deviations below average |
| | 5 activities with the most difference between verticals |
| | 5 activities with the least differences between verticals |

FIGURE 33. VERTICAL COMPARISON SCORECARD. This table allows for easy comparisons of observation rates for the eight verticals tracked in BSIMM15. A light gold color in the Activity column shows the five activities with the least differences in observation rates between verticals, whereas a light teal color shows the five activities with the most differences. Dark teal and gold in the remaining columns show observation rates that are significantly different from the average, either above or below.

PENETRATION TESTING

When comparing the percentage of activities observed with practices, Penetration Testing shows the highest median value of all 12 practices, with a median score of 57.1% observed for the entire data pool. This trend continues across most of the eight verticals (see Figure 34).




FIGURE 34. SCORE RANGES FOR PENETRATION TESTING.

The median value remains 57.1% across six of the eight verticals, and the range of the percentage of the observations is remarkably similar across most of the verticals. This is not surprising as penetration testing is one of the earliest recognized ways of finding vulnerabilities in software and is required by a number of regulations like PCI DSS.

The standout is the FinTech vertical, where all nine of the firms are doing at least 28.6% of the activities in the practice, and the interquartile range, the middle 50% of the scores, is significantly higher than all the other verticals. PCI DSS is likely a major driver for the FinTech vertical, although the Financial vertical has not adopted as many Penetration Testing activities as their FinTech peers and instead are similar to most of the other verticals and the entire BSIMM15 data pool.

Figure 33 allows us to take a deeper look at individual activities in the Penetration Testing practice. As noted above, [PT1.1] *use external penetration testers to find problems* is one of five

activities with very little differences noted between verticals. Observation rates are extremely high for [PT1.1], ranging from a low of 84% of the firms in the IoT vertical to 100% of the firms in the FinTech vertical. In terms of level 3 activities, there are generally higher observation rates for level 3 activities than there are in other practices. Each of the level 3 activities has at least some firms doing them across all eight verticals, something not seen in any other practice, even if activities introduced in the last few years are discounted.

COMPLIANCE & POLICY

Another strong practice in terms of median scores is Compliance & Policy, with an overall median score of 54.4%. In fact, six of the eight verticals have a median score of 54.4%. Financial achieves the highest median score of 63.6%, and Insurance achieves the lowest median score of 45.5%.




FIGURE 35. SCORE RANGES FOR COMPLIANCE & POLICY.

In five of the verticals, Cloud, IoT, FinTech, Financial, Insurance, and Healthcare, every firm is doing at least some of the Compliance & Policy activities, however, only IoT has firms doing 100% of the activities in the practice. Technology has firms doing 100% of the activities, but it also has firms doing none of the activities in the practice.

Figure 35 allows us to take a deeper look at individual activities in the Compliance & Policy practice. As noted above, [CP1.2] **identify privacy obligations** is one of five activities with very little differences noted between verticals. Observation rates are extremely high for [CP1.2] from a high of 100% for FinTech, Healthcare, Insurance, and IoT to a low of 84% for ISV and Tech. Considering ISV and Tech firms typically ship products that reside in a customer’s environment, these privacy obligations generally belong to the customer, but 84% of the firms in these verticals have still taken the time to identify their own obligations in terms of privacy.

Healthcare has two standout activities, [CP1.1] **unify regulatory pressures** and [CP2.3] **implement and track controls for**

compliance, where it does significantly more (at least 1.75 standard deviations above average) than its peers in the other verticals. This is not surprising, given the regulated nature of the healthcare industry.

IoT stands out from its peers on [CP3.3] **drive feedback from software lifecycle data back to policy**, where it does significantly more (again, greater than 1.75 standard deviations more) than peers in other verticals. This would indicate it is investing more than its peers to routinely learning from what is happening in their software development lifecycles.

ARCHITECTURE ANALYSIS

The practice with the greatest difference between median values is Architecture Analysis, with the highest median score of 55.6% (IoT) and lowest score of 22.2% (FinTech, Financial, and Insurance.) Other than those three low-scoring verticals, the median scores and interquartile ranges for other verticals show wide variance, indicating there is a difference in how valuable different industries view the activities in this practice (see Figure 36).




FIGURE 36. SCORE RANGES FOR ARCHITECTURE ANALYSIS.

While IoT shows the highest median value among the eight verticals, it does not have any firms performing all the activities, unlike the Financial vertical, where 100% is achieved. IoT also does not set the bar for minimum number of observed activities since it has firms that are doing no activities in the practice. On the other hand, Insurance, one of the low median score verticals, shows every firm is doing at least something in the Architecture Analysis practice.

ATTACK MODELS

The Attack Models practice shows the lowest range of scores and median scores of the 12 practices. Financial, Insurance, and Healthcare all show the highest overall scores among the eight verticals and the highest median scores. Only Insurance and Healthcare have all their firms doing at least something in the practice (see Figure 37).




FIGURE 37. SCORE RANGES FOR ATTACK MODELS.

Looking at Figure 33, VERTICAL COMPARISON SCORECARD, Insurance has three activities in which it has invested more heavily than the other verticals have. These are [AM1.3] **identify potential attackers**, [AM2.1] **build attack patterns and abuse cases tied to potential attackers**, and [AM3.5] **maintain and use a top N possible attacks list**.

Few of the activities in the Attack Models lend themselves to automation and tend to be hands-on efforts of SMEs. These relatively low scores mirror the general trends in seeing activities that allow automation increase in observation rates while SME-driven activities are seeing a relative decline.

TRAINING

While the BSIMM generally tries to avoid judgment on data, the Training practice shows disappointingly low scores. Teaching the people involved in the creation of software how to do so securely is one of the best defenses we have in securing our software. Beyond that, looking at this practice provides an excellent window into how business decisions drive how firms see if doing certain activities makes sense for them.

The lower upper end of the score ranges is understandable: some of the level 2 and level 3 activities only apply to a certain type business, for example, if a company does not hire vendors or outsource developers, there is no reason to invest in [T3.2] **provide training for vendors and outsourced workers**. In fact, some firms who do have vendors or outsourced workers developing and testing code have decided for budgetary or even legal reasons to not provide training for those workers. [T3.6] **Identify new security champions through observation** and [T2.5] **enhance security champions through training and events** do not apply to any firm that does not have a security champions program. [T2.10] **Host software security events** is about having significant events with large participation and outside speakers, all of which comes with significant costs. [T3.1] **Reward progression through curriculum** is about having real rewards and not just trinkets, and things like increased salaries and bonuses have significant budgetary impacts. It is easy to understand why nobody is doing 100% of the Training practice activities and for very legitimate reasons business reasons (see Figure 38).




FIGURE 38. SCORE RANGES FOR TRAINING.

On the other end of the score ranges, only the FinTech vertical has every firm doing at least some sort of training activity from the practice. The other seven verticals include firms that do not do any training activity sufficient to score in any of the 12 Training practice activities. The Healthcare vertical shows a median score of one activity out of the 12.

IOT, TECH, ISV, AND FINTECH

Besides looking at individual practices, looking at all 12 practices score distributions by vertical can provide some interesting insights.

Figure 39, Figure 40, and Figure 41 show the score distributions for the Tech, IoT, and ISV verticals, all of which are likely to see their products going into a customer's environment outside the control of the firm that produced the product. While there are some differences, the score distributions look remarkably similar, especially in the interquartile range area where 50% of the scores fall very similarly in each practice. These similarities suggest that the three verticals see the priorities between the 12 practices in roughly the same way.




FIGURE 39. SCORE DISTRIBUTIONS FOR TECH.




FIGURE 40. SCORE DISTRIBUTIONS FOR IOT.




FIGURE 41. SCORE DISTRIBUTIONS FOR ISV.

Isolating the comparison to just the median scores shows similarities in scores in an even clearer way (see Figure 42).




FIGURE 42. MEDIAN SCORES FOR TECH, IOT, AND ISV VERTICALS.

The median scores for all three verticals are very close, while the median scores for Tech and ISV are the same across the board. There are slight differences for IoT in Attack Models, Security Features & Design, Standards & Requirements, Architecture Analysis, and Software Environment, representing five of the 12 practices.

To illustrate different priorities, consider the FinTech vertical in Figure 43. The score differences show a different set of priorities than Tech, IoT, and ISV.




FIGURE 43. SCORE DISTRIBUTIONS FOR FINTECH.




FIGURE 44. MEDIAN SCORES FOR TECH, IOT, ISV, AND FINTECH VERTICALS.

Comparing the median scores, Figure 44 shows some similar priorities between FinTech, Tech, IoT, and ISV verticals but greater variability than seen between just Tech, IoT, and ISV. FinTech shows higher interest in Strategy & Metrics, Code Review, and Security Testing while not focusing as much on Training or Architecture Analysis.

FINTECH AND FINANCIAL

On the surface, FinTech and Financial might have similar priorities in software security, but that does not appear to be the case, as seen in Figure 43 and Figure 45.

There are some clear differences between the two industries, and the relative sizes of the various interquartile ranges, the ranges that contain 50% of the scores, is interesting. FinTech has relatively small interquartile ranges compared to Financial for Training, Attack Models, and Security Testing. This could be an artifact of the relatively small sample size of the FinTech vertical, but it could also be an indication that all nine of the firms see these practices with very similar priorities. FinTech has larger interquartile ranges for Compliance & Policy and Security Features & Design than does Financial. Even in a relatively small vertical like FinTech, a large range of scores shows a wide variety of prioritization of activities in practices where the range is large.




FIGURE 45. SCORE DISTRIBUTIONS FOR FINANCIAL.




FIGURE 46. MEDIAN SCORE COMPARISON FOR FINTECH AND FINANCIAL.

The median scores of the FinTech and Financial verticals (see Figure 46) show some different priorities. While both score well in Compliance & Policy, Financial wins out by a bit, which is perhaps a reflection of the heavily regulated nature of the banking industry. Financial also does better in Training and Attack Models. The difference in Attack Models is likely a result in the difference in the industry: Financial is largely focused on securing its own assets in its own software, while

FinTech is largely focused on providing secure software for others. This difference between protecting your own assets and protecting software others are using likely drives FinTech's better performance in Security Features & Design, Standards & Requirements, Code Review, and Security Testing. FinTech's better performance in Software Environment reflect how its software is often delivered as SaaS.



DATA ANALYSIS: LONGITUDINAL

Every SSI changes over time as technologies, attackers, attacks, budgets, and everything else also changes. You can use this information to see whether your SSI's trajectory is similar to that of other programs.

The BSIMM captures real-world data about how organizations approach software security across their portfolio. Given the BSIMM's longevity, this data provides a unique snapshot of how participant SSIs have evolved over the past 16 years, as well as how individual programs have changed from assessment to assessment.

BUILDING A MODEL FOR SOFTWARE SECURITY

In the late 1990s, software security began to flourish as a discipline separate from computer and network security. Researchers began to put more emphasis on studying the ways in which a developer can contribute to or unintentionally

undermine the security of an application and started asking some specific questions: What kinds of bugs and flaws lead to security problems? How can we identify these problems systematically?

Within a few years, there was an emerging consensus that building secure software required more than smart individuals toiling away on guidance and training. Getting security right, especially across a software portfolio, meant being directly involved in the software development process, guiding it even as the process evolves. Since then, practitioners have come to learn that process, testing, and developer tools alone are insufficient: software security encompasses business, social, and organizational aspects as well.

Table 16 shows how the BSIMM has grown over the years. (Recall that our data freshness constraints, introduced with BSIMM-V and later tightened, cause data from firms with aging measurements to be removed.) BSIMM15 describes the work of 11,100 SSG and champions working directly in software security, impacting the security efforts of almost 260,000 developers.

| BSIMM VERTICAL PARTICIPANTS OVER TIME | | | | | |
|---------------------------------------|-------|-----------------|-----------------|-----------------|------------------|
| | FIRMS | 1ST MEASUREMENT | 2ND MEASUREMENT | 3RD MEASUREMENT | 4TH+ MEASUREMENT |
| BSIMMI5 | 121 | 76 | 26 | 11 | 8 |
| BSIMMI4 | 130 | 81 | 31 | 10 | 8 |
| BSIMMI3 | 130 | 76 | 35 | 11 | 8 |
| BSIMMI2 | 128 | 76 | 31 | 14 | 7 |
| BSIMMI1 | 130 | 77 | 32 | 12 | 9 |
| BSIMMI0 | 122 | 72 | 29 | 13 | 8 |
| BSIMM9 | 120 | 78 | 22 | 13 | 7 |
| BSIMM8 | 109 | 73 | 20 | 11 | 5 |
| BSIMM7 | 95 | 65 | 15 | 13 | 2 |
| BSIMM6 | 78 | 52 | 16 | 8 | 2 |
| BSIMM-V | 67 | 46 | 17 | 4 | 0 |
| BSIMM4 | 51 | 38 | 12 | 1 | 0 |
| BSIMM3 | 42 | 31 | 11 | 0 | 0 |
| BSIMM2 | 30 | 30 | 0 | 0 | 0 |
| BSIMMI | 9 | 9 | 0 | 0 | 0 |

TABLE 16. BSIMM ASSESSMENTS DONE OVER TIME. The chart shows how the BSIMM study has grown over the years, including how some firms have received multiple measurements.

Forty-five of the current participating firms have been through at least two assessments, allowing us to study how their initiatives changed over time. Across North America, EMEA, and APAC, 26 firms are on their second assessment, 11 firms are on their third, five are on their fourth, and two are on their fifth. One North America firm has undertaken its sixth assessment, continuing its use of the BSIMM as an SSI planning and management tool. Figure 47 shows these firms by percentages across three major BSIMM regions.




FIGURE 47. ONGOING USE OF THE BSIMM IN DRIVING ORGANIZATIONAL MATURITY. Organizations are continuing to do remeasurements to show that their efforts are achieving the desired results (e.g., about 31% of North America participants have been assessed more than once).

CHANGES BETWEEN FIRST AND SECOND ASSESSMENTS

Forty-five of the 121 firms in BSIMM15 have been measured at least twice. On average, the time between first and second measurements for those 45 firms was 33.5 months, and overall, the average time between assessments is 32.7 months. Although observations of individual activities among the 12 practices come and go (as shown in Figure 48), in general, remeasurement over time shows a clear trend of increased maturity. Simply put, SSIs mature over time.

As shown in Figure 48, firms moving from their first assessment to their second tend to invest in:

- Defining their program ([SM1.1] **publish process and evolve as necessary**, [SM2.1] **publish data about software security internally and use it to drive change**), scaling the program using champions ([SM2.3] **create or grow a security champions program**), and evangelizing the secure SDLC as well
- Defining and enforcing policy and standards ([CP1.3] **create policy**)
- Managing vendors through boilerplate security SLAs ([CP2.4] **include software security SLAs in all vendor contracts**, [SR2.5] **create SLA boilerplate**)
- Identifying open source components ([SR1.5] **identify open source**)

Figure 49 shows the average normalized observation rate per practice for the 45 firms that have had a second assessment. Over the average of about 33.5 months between the two assessments, we see clear growth in every practice, especially in Strategy & Metrics, Compliance & Policy, and Standards & Requirements. The practices with the highest overall growth align with the individual activities identified in Figure 48. The changes indicate that firms feel prepared for their first assessment after focusing on foundational and technical activities such as training and testing but then expand into governance as they mature their SSIs.

There are two factors causing the numerical changes seen in the longitudinal scorecard (Figure 48, showing 45 BSIMM15 firms moving from their first to second assessments). The first factor is that more firms have now done their second assessment (adding firms to this group), and the second is that we drop old data (removing firms from this group). Grouped together, the two factors can cause a significant amount of change in the group of firms that have had a second assessment, even if the change isn't directly visible in the scorecard.

| GOVERNANCE | | | INTELLIGENCE | | | SSDL TOUCHPOINTS | | | DEPLOYMENT | | |
|---------------------|-----------------------------|-----------------------------|----------------------------|-----------------------------|-----------------------------|-----------------------|-----------------------------|-----------------------------|-----------------------------|-----------------------------|-----------------------------|
| ACTIVITY | BSIMM ROUND 1 (OF 45 FIRMS) | BSIMM ROUND 2 (OF 45 FIRMS) | ACTIVITY | BSIMM ROUND 1 (OF 45 FIRMS) | BSIMM ROUND 2 (OF 45 FIRMS) | ACTIVITY | BSIMM ROUND 1 (OF 45 FIRMS) | BSIMM ROUND 2 (OF 45 FIRMS) | ACTIVITY | BSIMM ROUND 1 (OF 45 FIRMS) | BSIMM ROUND 2 (OF 45 FIRMS) |
| STRATEGY & METRICS | | | ATTACK MODELS | | | ARCHITECTURE ANALYSIS | | | PENETRATION TESTING | | |
| [SM1.1] | 26 | 39 | [AM1.2] | 27 | 30 | [AA1.1] | 41 | 41 | [PT1.1] | 40 | 40 |
| [SM1.3] | 24 | 31 | [AM1.3] | 10 | 20 | [AA1.2] | 12 | 21 | [PT1.2] | 34 | 33 |
| [SM1.4] | 39 | 41 | [AM1.5] | 19 | 30 | [AA1.4] | 22 | 26 | [PT1.3] | 26 | 33 |
| [SM1.7] | 21 | 28 | [AM2.1] | 3 | 7 | [AA2.1] | 8 | 18 | [PT2.2] | 6 | 10 |
| [SM2.1] | 14 | 28 | [AM2.6] | 3 | 5 | [AA2.2] | 7 | 18 | [PT2.3] | 11 | 16 |
| [SM2.3] | 13 | 35 | [AM2.7] | 3 | 5 | [AA2.4] | 11 | 15 | [PT3.1] | 4 | 6 |
| [SM2.6] | 23 | 26 | [AM2.8] | 0 | 7 | [AA3.1] | 2 | 7 | [PT3.2] | 4 | 5 |
| [SM2.7] | 17 | 25 | [AM2.9] | 3 | 4 | [AA3.2] | 0 | 1 | | | |
| [SM3.1] | 10 | 11 | [AM3.2] | 1 | 0 | [AA3.3] | 2 | 5 | | | |
| [SM3.2] | 0 | 5 | [AM3.4] | 1 | 4 | | | | | | |
| [SM3.3] | 5 | 13 | [AM3.5] | 5 | 4 | | | | | | |
| [SM3.4] | 0 | 4 | | | | | | | | | |
| [SM3.5] | 0 | 0 | | | | | | | | | |
| COMPLIANCE & POLICY | | | SECURITY FEATURES & DESIGN | | | CODE REVIEW | | | SOFTWARE ENVIRONMENT | | |
| [CP1.1] | 32 | 40 | [SFD1.1] | 35 | 38 | [CR1.2] | 29 | 29 | [SE1.1] | 20 | 33 |
| [CP1.2] | 36 | 42 | [SFD1.2] | 29 | 34 | [CR1.4] | 27 | 42 | [SE1.2] | 40 | 42 |
| [CP1.3] | 23 | 39 | [SFD2.1] | 9 | 18 | [CR1.5] | 14 | 25 | [SE1.3] | 9 | 25 |
| [CP2.1] | 16 | 23 | [SFD2.2] | 14 | 20 | [CR1.7] | 10 | 22 | [SE1.4] | 18 | 23 |
| [CP2.2] | 19 | 19 | [SFD3.1] | 2 | 6 | [CR2.6] | 5 | 11 | [SE2.4] | 11 | 16 |
| [CP2.3] | 23 | 27 | [SFD3.2] | 4 | 6 | [CR2.7] | 7 | 9 | [SE2.5] | 8 | 17 |
| [CP2.4] | 12 | 27 | [SFD3.3] | 2 | 0 | [CR2.8] | 12 | 16 | [SE2.7] | 7 | 11 |
| [CP2.5] | 31 | 26 | | | | [CR3.2] | 1 | 6 | [SE3.2] | 6 | 5 |
| [CP3.1] | 8 | 14 | | | | [CR3.3] | 1 | 1 | [SE3.3] | 3 | 3 |
| [CP3.2] | 6 | 16 | | | | [CR3.4] | 0 | 0 | [SE3.6] | 2 | 5 |
| [CP3.3] | 3 | 7 | | | | [CR3.5] | 0 | 1 | [SE3.8] | 0 | 0 |
| | | | | | | | | | [SE3.9] | 0 | 1 |
| | | | | | | | | | [SE3.10] | 0 | 0 |
| TRAINING | | | STANDARDS & REQUIREMENTS | | | SECURITY TESTING | | | CONFIG. MGMT. & VULN. MGMT. | | |
| [T1.1] | 26 | 27 | [SR1.1] | 31 | 35 | [ST1.1] | 36 | 41 | [CMVM1.1] | 37 | 41 |
| [T1.7] | 12 | 28 | [SR1.2] | 29 | 37 | [ST1.3] | 35 | 33 | [CMVM1.2] | 36 | 34 |
| [T1.8] | 9 | 18 | [SR1.3] | 35 | 37 | [ST1.4] | 11 | 24 | [CMVM1.3] | 32 | 34 |
| [T2.5] | 8 | 19 | [SR1.5] | 19 | 34 | [ST2.4] | 4 | 7 | [CMVM1.4] | 31 | 34 |
| [T2.8] | 9 | 9 | [SR2.2] | 17 | 28 | [ST2.5] | 4 | 9 | [CMVM2.3] | 20 | 24 |
| [T2.9] | 5 | 14 | [SR2.5] | 11 | 28 | [ST2.6] | 8 | 9 | [CMVM2.4] | 3 | 9 |
| [T2.10] | 3 | 10 | [SR2.7] | 11 | 6 | [ST3.3] | 0 | 5 | [CMVM3.1] | 1 | 3 |
| [T2.11] | 2 | 13 | [SR3.2] | 5 | 10 | [ST3.4] | 1 | 2 | [CMVM3.2] | 2 | 6 |
| [T2.12] | 1 | 10 | [SR3.3] | 4 | 6 | [ST3.5] | 0 | 1 | [CMVM3.3] | 3 | 8 |
| [T3.1] | 0 | 3 | [SR3.4] | 10 | 9 | [ST3.6] | 0 | 1 | [CMVM3.4] | 4 | 12 |
| [T3.2] | 5 | 9 | [SR3.5] | 0 | 0 | | | | [CMVM3.5] | 3 | 4 |
| [T3.6] | 0 | 2 | | | | | | | [CMVM3.6] | 0 | 0 |
| | | | | | | | | | [CMVM3.8] | 0 | 0 |

FIGURE 48. BSIMM15 REASSESSMENTS SCORECARD ROUND 1 VS. ROUND 2. This chart shows how 45 SSIs changed between their first and second assessments. Dark gold shows the top five activities with the most increase in observations by count. Light gold shows the next five activities with the most increase in observations by count.




FIGURE 49. FIRMS ROUND 1 VS. FIRMS ROUND 2. This diagram illustrates the normalized observation rate change, on a percentage scale, in 49 firms between their first and second BSIMM assessments.

CHANGES BETWEEN FIRST AND THIRD ASSESSMENTS

Nineteen of the 121 firms in BSIMM15 have been measured at least three times, and on average, the time between first and third measurements for those 19 firms was 57.8 months. Although individual activities among the 12 practices come and go, in general, remeasurement over time shows a clear trend of increased maturity.

As shown in Figure 50, firms that move from their first assessment to their third over the course of about 57.8 months, in addition to changes shown previously, tend to invest in:

- Enabling self-sufficient engineering teams by leveraging investments in training and mentoring, ([T1.7] **deliver on-demand individual training**, [T1.8] **include security resources in onboarding**, and ([CR1.7] **assign code review tool mentors**)
- Securing cloud environments ([SE1.3]) and using containers to support security efforts ([SE2.5])
- Formalizing the creation of standards ([SR2.2]) and marketing security efforts internally ([SM2.7])

Interestingly, while Figure 51 shows growth in every practice, it shows only a slight increase in the Security Testing and Configuration Management & Vulnerability Management practices. This could mean that most organizations do a variety of Security Testing and Configuration Management & Vulnerability Management activities earlier on in their journeys.

Although individual activities in the 12 practices come and go, in general, remeasurement over time shows a clear trend of increased maturity.

| GOVERNANCE | | | INTELLIGENCE | | | SSDL TOUCHPOINTS | | | DEPLOYMENT | | |
|---------------------|--------------------------------------|-----------------------------------|----------------------------|--------------------------------------|-----------------------------------|-----------------------|--------------------------------------|-----------------------------------|-----------------------------|--------------------------------------|-----------------------------------|
| ACTIVITY | BSIMM ROUND 1 (OF 19 FIRMS) | BSIMM ROUND 3 (OF 19 FIRMS) | ACTIVITY | BSIMM ROUND 1 (OF 19 FIRMS) | BSIMM ROUND 3 (OF 19 FIRMS) | ACTIVITY | BSIMM ROUND 1 (OF 19 FIRMS) | BSIMM ROUND 3 (OF 19 FIRMS) | ACTIVITY | BSIMM ROUND 1 (OF 19 FIRMS) | BSIMM ROUND 3 (OF 19 FIRMS) |
| STRATEGY & METRICS | | | ATTACK MODELS | | | ARCHITECTURE ANALYSIS | | | PENETRATION TESTING | | |
| [SM1.1] | 7 | 18 | [AM1.2] | 15 | 16 | [AA1.1] | 16 | 19 | [PT1.1] | 16 | 18 |
| [SM1.3] | 9 | 14 | [AM1.3] | 4 | 12 | [AA1.2] | 2 | 8 | [PT1.2] | 12 | 16 |
| [SM1.4] | 16 | 18 | [AM1.5] | 10 | 14 | [AA1.4] | 10 | 12 | [PT1.3] | 11 | 14 |
| [SM1.7] | 6 | 13 | [AM2.1] | 1 | 8 | [AA2.1] | 0 | 7 | [PT2.2] | 2 | 8 |
| [SM2.1] | 6 | 14 | [AM2.6] | 0 | 1 | [AA2.2] | 1 | 7 | [PT2.3] | 4 | 7 |
| [SM2.3] | 4 | 13 | [AM2.7] | 1 | 3 | [AA2.4] | 2 | 6 | [PT3.1] | 1 | 3 |
| [SM2.6] | 8 | 13 | [AM2.8] | 0 | 1 | [AA3.1] | 0 | 3 | [PT3.2] | 1 | 4 |
| [SM2.7] | 3 | 15 | [AM2.9] | 1 | 3 | [AA3.2] | 0 | 0 | | | |
| [SM3.1] | 5 | 7 | [AM3.2] | 0 | 1 | [AA3.3] | 0 | 1 | | | |
| [SM3.2] | 0 | 4 | [AM3.4] | 0 | 2 | | | | | | |
| [SM3.3] | 3 | 5 | [AM3.5] | 2 | 5 | | | | | | |
| [SM3.4] | 0 | 1 | | | | | | | | | |
| [SM3.5] | 0 | 0 | | | | | | | | | |
| COMPLIANCE & POLICY | | | SECURITY FEATURES & DESIGN | | | CODE REVIEW | | | SOFTWARE ENVIRONMENT | | |
| [CP1.1] | 12 | 19 | [SFD1.1] | 17 | 18 | [CR1.2] | 12 | 16 | [SE1.1] | 8 | 15 |
| [CP1.2] | 15 | 18 | [SFD1.2] | 13 | 14 | [CR1.4] | 13 | 19 | [SE1.2] | 16 | 17 |
| [CP1.3] | 9 | 15 | [SFD2.1] | 3 | 7 | [CR1.5] | 3 | 12 | [SE1.3] | 2 | 12 |
| [CP2.1] | 7 | 10 | [SFD2.2] | 6 | 12 | [CR1.7] | 3 | 14 | [SE1.4] | 5 | 7 |
| [CP2.2] | 5 | 11 | [SFD3.1] | 0 | 4 | [CR2.6] | 2 | 5 | [SE2.4] | 3 | 4 |
| [CP2.3] | 9 | 13 | [SFD3.2] | 3 | 7 | [CR2.7] | 5 | 6 | [SE2.5] | 1 | 11 |
| [CP2.4] | 6 | 10 | [SFD3.3] | 0 | 2 | [CR2.8] | 7 | 10 | [SE2.7] | 0 | 6 |
| [CP2.5] | 12 | 13 | | | | [CR3.2] | 0 | 2 | [SE3.2] | 2 | 2 |
| [CP3.1] | 5 | 9 | | | | [CR3.3] | 1 | 3 | [SE3.3] | 2 | 2 |
| [CP3.2] | 5 | 5 | | | | [CR3.4] | 0 | 0 | [SE3.6] | 1 | 3 |
| [CP3.3] | 1 | 1 | | | | [CR3.5] | 0 | 0 | [SE3.8] | 0 | 0 |
| | | | | | | | | | [SE3.9] | 0 | 0 |
| | | | | | | | | | [SE3.10] | 0 | 0 |
| TRAINING | | | STANDARDS & REQUIREMENTS | | | SECURITY TESTING | | | CONFIG. MGMT. & VULN. MGMT. | | |
| [T1.1] | 10 | 12 | [SR1.1] | 13 | 16 | [ST1.1] | 16 | 17 | [CMVM1.1] | 17 | 18 |
| [T1.7] | 5 | 16 | [SR1.2] | 12 | 18 | [ST1.3] | 15 | 16 | [CMVM1.2] | 16 | 15 |
| [T1.8] | 1 | 12 | [SR1.3] | 14 | 17 | [ST1.4] | 4 | 13 | [CMVM1.3] | 14 | 13 |
| [T2.5] | 3 | 10 | [SR1.5] | 6 | 15 | [ST2.4] | 0 | 2 | [CMVM1.4] | 15 | 18 |
| [T2.8] | 3 | 5 | [SR2.2] | 6 | 15 | [ST2.5] | 0 | 4 | [CMVM2.3] | 12 | 13 |
| [T2.9] | 0 | 8 | [SR2.5] | 4 | 10 | [ST2.6] | 2 | 1 | [CMVM2.4] | 0 | 2 |
| [T2.10] | 0 | 1 | [SR2.7] | 4 | 10 | [ST3.3] | 0 | 1 | [CMVM3.1] | 1 | 0 |
| [T2.11] | 0 | 5 | [SR3.2] | 4 | 5 | [ST3.4] | 0 | 1 | [CMVM3.2] | 0 | 0 |
| [T2.12] | 0 | 5 | [SR3.3] | 1 | 3 | [ST3.5] | 0 | 1 | [CMVM3.3] | 2 | 4 |
| [T3.1] | 0 | 1 | [SR3.4] | 6 | 5 | [ST3.6] | 0 | 0 | [CMVM3.4] | 2 | 8 |
| [T3.2] | 0 | 5 | [SR3.5] | 0 | 0 | | | | [CMVM3.5] | 2 | 4 |
| [T3.6] | 0 | 0 | | | | | | | [CMVM3.6] | 0 | 0 |
| | | | | | | | | | [CMVM3.8] | 0 | 0 |

FIGURE 50. BSIMM15 REASSESSMENTS SCORECARD ROUND 1 VS. ROUND 3. This chart shows how 18 SSIs changed between their first and third assessments. Dark gold shows the top five activities with the most increase in observations by count. Light gold shows the next five activities with the most increase in observations by count.




FIGURE 51. FIRMS ROUND 1 VS. FIRMS ROUND 3 SPIDER CHART. This diagram illustrates the normalized observation rate change, on a percentage scale, in 18 firms between their first and third BSIMM assessments.

About Black Duck

Black Duck® meets the board-level risks of modern software with True Scale Application Security, ensuring uncompromised trust in software for the regulated, AI-powered world. Only Black Duck solutions free organizations from tradeoffs between speed, accuracy, and compliance at scale while eliminating security, regulatory, and licensing risks. Whether in the cloud or on premises, Black Duck is the only choice for securing mission-critical software everywhere code happens. With Black Duck, security leaders can make smarter decisions and unleash business innovation with confidence. Learn more at www.blackduck.com.

