

Getting Started with Static Analysis

 **PARASOFT.**



Table of Contents

What is Static Analysis?	2	Dashboards for High Level Analysis	9
How does Static Analysis work?.....	2	Managing Violations of Coding Standards	11
Introducing Static Analysis into Your Project.....	3	Managing Backlog of Bug and Security Warnings	12
Start with the End in Mind.....	3	Optimizing Static Analysis	13
Static Analysis at Each Stage of Product Maturity.....	4	Define the set of rules to be used	13
Existing Project – in Current Development.....	4	Modify settings for data flow analysis.....	13
Existing Project – Product on the Market in Maintenance.....	4	Customize rule groups, identifiers, titles and severities	13
Greenfield Project.....	4	Customize existing rules	13
		Create new static analysis rules.....	13
Managing Early Static Analysis Reports	4	Summary	14
Line in the sand approach	4		
Acknowledge and defer approach.....	5		
Greenfield approach.....	5		
Configuration versus Filtering.....	5		
Integrating Static Analysis into Everyday Workflows.....	6		
Integration into Build Systems and Continuous Integration Pipelines	7		
Tackling the Backlog of Warnings and Technical Debt	9		

What is Static Analysis?

In simple terms, static analysis is the process of examining source and binary code without execution, usually for the purposes of finding bugs or evaluating quality. Unlike dynamic analysis (e.g. Parasoft Insure++), which requires a running program to work, static analysis can be run on source without the need for an executable.

This means static analysis can be used on partially-complete code, libraries, and third-party source code. Static analysis is accessible to the developer, to be used as code is being written or modified, or to be applied on any arbitrary code base. In the application security domain, static analysis goes by the term Static Application Security Testing (SAST). Many commercial tools support both security vulnerability detection alongside bug detection, quality metrics and coding standard conformance.

Static analysis tools are mandated or highly recommended by safety standards such as ISO 26262 and EN 50128, for their ability to detect hard-to-find defects and improve security of software. Static analysis tools also help software teams conform to coding standards such as MISRA or CERT.

How does Static Analysis work?

Static code analyzers use a compiler-like front-end to build a syntactic and semantic model of the software. The syntactic model is then analyzed against a set of rules or “checkers” to see if the code is in violation. These checkers use pattern-matching algorithms to detect errors such as poor use of language constructs, use of insecure functions, and violations of coding guidelines. The specific set of checkers used is configurable by the user. Pre-set configurations are provided for convenience, for instance for coding standards such as MISRA C.

More sophisticated checkers employ semantic analysis that uses data and control flow to detect complex bugs and security vulnerabilities. To do this, the static analyzer builds an execution model of the software, considers possible paths through the code, and evaluates use of data as it flows from source (e.g. user input) to its destination (e.g. an API call such as a system call). Analyzing every single possible condition and path would be too time consuming, so the analyzer uses heuristics to detect the most likely paths for evaluation. Types of errors detected by these checkers include null pointer dereference, buffer overflows, and security vulnerabilities like command and SQL injections.

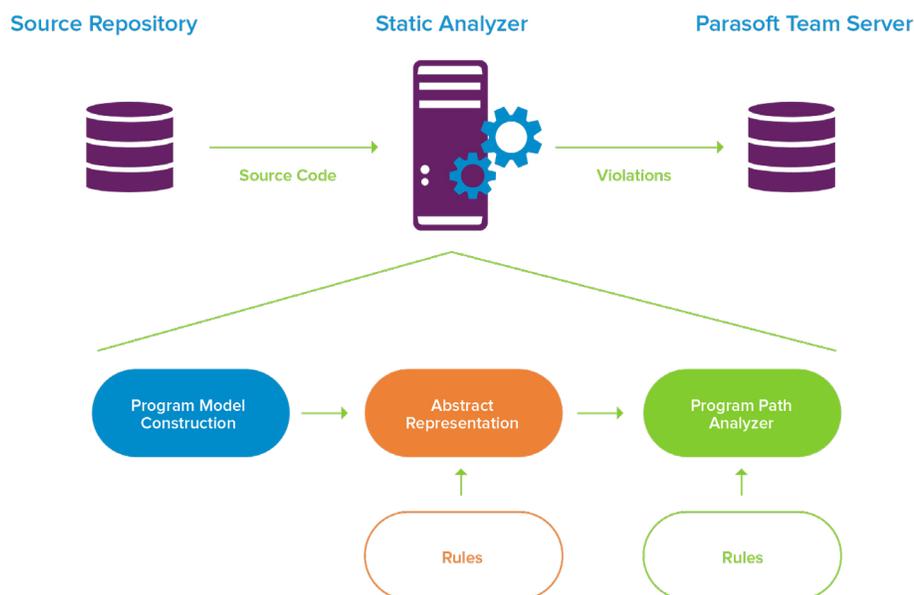


Figure 1:

A high level view of the static analysis process.

Static analysis is most often used in two ways:

- At the developer desktop integrated into their development environment (IDE)
- At the command line as part of a build or continuous integration process

Commercial tools like Parasoft C/C++test and Jtest integrate with the leading IDEs, and provide tools in both environments so users can effectively manage analysis results. Integrating tools and processes to manage static analysis results enables users to move beyond just a list of warnings.

So where do you start? This paper will guide you through the introduction of an advanced static analysis tool to your project. The primary assumption is that software already exists and the team is actively working on developing or maintaining a product. Greenfield (a completely new project) is also considered, however, many of the processes and techniques do overlap.

Introducing Static Analysis into Your Project

A great thing about static analysis tools is that they can be introduced and used at any stage of a project, effective even if a project is incomplete and partially coded. The biggest challenge with introducing static analysis is that a large amount of code can produce a large number of warnings. Therefore, the focus when integrating static analysis into a project, should be on getting the team productive as soon as possible, and minimizing the opportunity for the team to get overwhelmed by all the static analysis warnings. This is not to diminish the importance of these warnings, but most developers don't have the luxury of fixing existing or legacy code, at least not immediately. The focus of the approach outlined below is to first integrate the tools into everyday processes so that access and usability is maximized, and then deal with the most critical bugs and security vulnerabilities. Finally, as the team becomes more proficient, optimize the tools and processes to increase the return on investment.

Start with the End in Mind

In order to get the most out of static analysis tools, it's important to understand the end goal. If the goal is better security, for example, that will shape the focus of analysis and remediation, or if the goal is complying a coding standard such as MISRA C, the focus will become satisfying the coding standard and proving it to certification entities as required.

When first adopting static analysis, it is especially important to keep the end goal in mind. It is easy to fall into the trap that more is better, or in other words, that more analysis and more warnings means getting the most value out of the tool. This is a common trap. Instead, stay focused on the goal. If security is the focus, keep the focus on improving security and reduce the distraction of other types of warnings. Of course critical bugs are always important to track down, but they shouldn't distract from the main goal. Over time, as the team becomes more proficient, you will be able to incorporate other secondary goals such as improving overall quality and enforcing coding standards. As static analysis becomes part of each developer's daily routine, they will be able to analyze results more quickly and fix bugs more efficiently. At this time, the secondary goals will more effectively be achieved, instead of simply being overwhelming.

Static Analysis at Each Stage of Product Maturity

It's incredibly important to keep these goals in mind and ease introduction of static analysis, in order to reduce the likelihood of overwhelming the team. Also important to consider is the maturity of the product under development, as it impacts the way static analysis can be adopted. Consider the major stages of development below.

Existing Project – in Current Development

The most common scenario is a software organization that decides to use static analysis, and rolls it out to their current projects. Each project may choose to adopt the tools at the start of a sprint or at the beginning of a major new feature update. Realistically, software teams are always working – even as one product is “finished,” another version or variant is underway. The key aspect of this adoption scenario is that there is a significant body of existing code and new code being developed daily. The recommended approach to integration is called **“a line in the sand”** approach, which we will discuss in more detail in just a moment. At a high level, this approach means improving new code as it's developed, while deferring less critical warnings as technical debt.

Existing Project – Product on the Market in Maintenance

Adopting static analysis for a mature product may have different goals than a project still under development. This is a product that is in the elder years of the software development lifecycle, in which little new code is being written, only to fix lingering bugs and security vulnerabilities. The primary approach to adopting static analysis for these projects is called **“acknowledge and defer.”** Since little new code is being developed, all of the discovered bugs and security vulnerabilities are added to the existing technical debt.

Greenfield Project

Although it's not often that software teams get to have a fresh start, a new product and project is the ideal point to integrate new tools and techniques into the development process. In these projects, little existing code specific to the project exists, but it still may rely on third-party and open-source software. Developers can integrate static analysis in their development environments from the start, ensuring a high standard of quality as code is being written. This allows for the adoption of coding standards and ensuring critical static analysis warnings are dealt with as they arise, thus adding less bugs and vulnerabilities to the technical debt pile. The approach to adoption in this case is aptly named **“greenfield.”**

Managing Early Static Analysis Reports

Once a static analysis tool has been installed and configured in a project and any dependency issues have been sorted out, there is usually a fairly lengthy report of violations and warnings reported by the tool. This can be overwhelming, especially in a large codebase, so how these initial reports are managed will directly influence the success of integrating the tool into the project. Not all warnings are critical, so everything doesn't need to be dealt with immediately. Learning what to address immediately and what to defer is the key to success. As mentioned above, the maturity and size of the product has a direct influence on approach, outlined below in more detail.

Line in the sand approach

As the name implies, in this approach, developers decide that after the initial analysis, they won't let any more critical warnings and violations enter the code base. In other words, they make a commitment to analyze each critical warning to decide its veracity, and implement a timely fix, if it's indeed a bug. The team may also decide to add critical warnings already discovered in existing code to be added to the list of bugs in their reporting tool. Examples of these types of warnings might be critical security vulnerabilities such as SQL injections, or serious memory errors such as buffer overflows. In most cases, the less serious warnings can be deferred for later analysis. You might be thinking, “doesn't this just add to our technical debt?” And if you are, you're right! But at this stage, we're ok with that. Any potential bugs within these warnings were already in the technical debt pile. At least now, they are identified and much easier to fix at a later time.

Learning what to address immediately and what to defer is the key to success.

Acknowledge and defer approach

In the case where a product is already on the market and under maintenance, it is still beneficial to identify any lingering bugs and security vulnerabilities in the code, but it's not feasible for developers to analyze (let alone fix) all these warnings. In such a case, it makes sense to look at the top most critical reports and decide a course of action. The rest of the warnings are acknowledged, as in the software team recognizes they exist, but they are mostly deferred for a later time. (This again adds to technical debt of the organization, but as mentioned above, these bugs technically exist already as technical debt.) This approach differs from the line-in-the-sand approach in that after identifying the key warnings, you simply defer the rest, without necessarily any analysis.

Greenfield approach

A project with little existing code is an ideal starting point for static analysis. In this case, the software teams can investigate all warnings that arise and fix found bugs. Unlike the other approaches, there are only a few warnings to manage, so developers can tackle the additional workload. This is also an ideal time to implement and enforce a coding standard through the tools, since violations can be identified and fixed right within the IDE and before any code is submitted to version control (which you could also do in the other scenarios described here).

The adoption of static analysis in the three major stages of maturity are differentiated by how they deal with the backlog of warnings, as illustrated below:

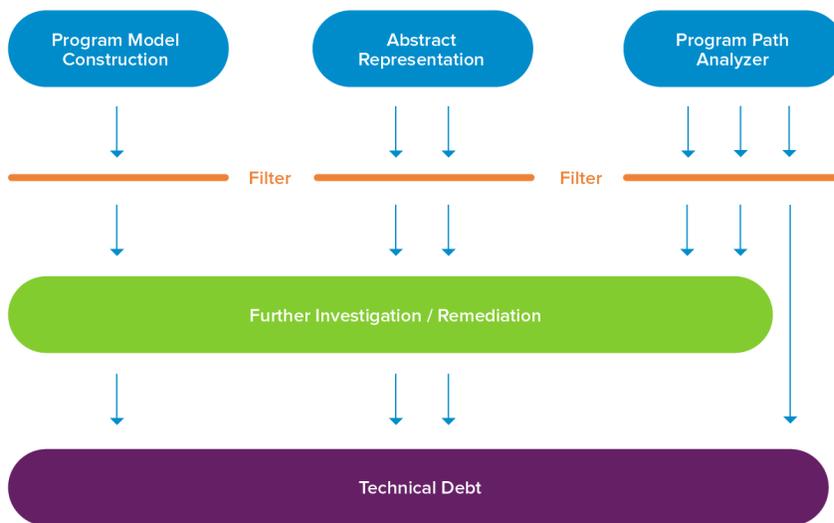


Figure 2:

In a greenfield project, most reported warnings are investigated and fixed with little going into the technical debt pile. Projects under development tend to have a backlog of warnings that are mostly deferred with only critical warnings being dealt with, and products under maintenance tend to have most warnings deferred.

Configuration versus Filtering

One of the main differences between open source or lightweight static analysis tools and commercial advanced static analysis tools is the ability to configure which set of checkers are enabled for the analysis, and filter out reported results based on warning category, file name, severity and other attributes. This helps with the goal of not getting overwhelmed – developers can focus on just the types of warnings that they are interested in, and reduce the amount of information provided at any given time.

There is also a difference worth noting between configuring checkers and filtering results. Although initially it might seem better to limit the number of rules in the global configuration, filtering should often be used instead, to limit the scope of reporting rather than eliminate the checker entirely. If a rule that later turns out to be important is turned off in the configuration, there will be no history in the warning repository, so you won't be able to find out if the error was introduced by recent changes or already in the code before static analysis was adopted.

A recommended approach is to use configuration to simply limit the set of rules to those that are foreseeable as useful for the software team. Again, start with the end goal in mind: if improving security is the key goal, it makes sense to enable all security-related rules, disable less important rules, and enable one of the built-in secure coding standards such as CERT C. Then, if you're using an advanced static analysis solution like Parasoft C/C++test, you can leverage its built-in management tools to deal with the data produced from the static analysis reports and drive future development focus.

Integrating Static Analysis into Everyday Workflows

The key to making static analysis a success in a project is to make sure the tools are easy to use and accessible by developers, so the tool must provide useful, actionable information upfront without overwhelming users with information. This is best achieved within the environment the developer is working in, such as Eclipse or Visual Studio. Static analysis warnings are delivered in the same manner as compiler errors in the IDE, and these warnings are highlighted in the code to make analysis and fixing much easier. See the following example of Parasoft Jtest integrated in Eclipse:

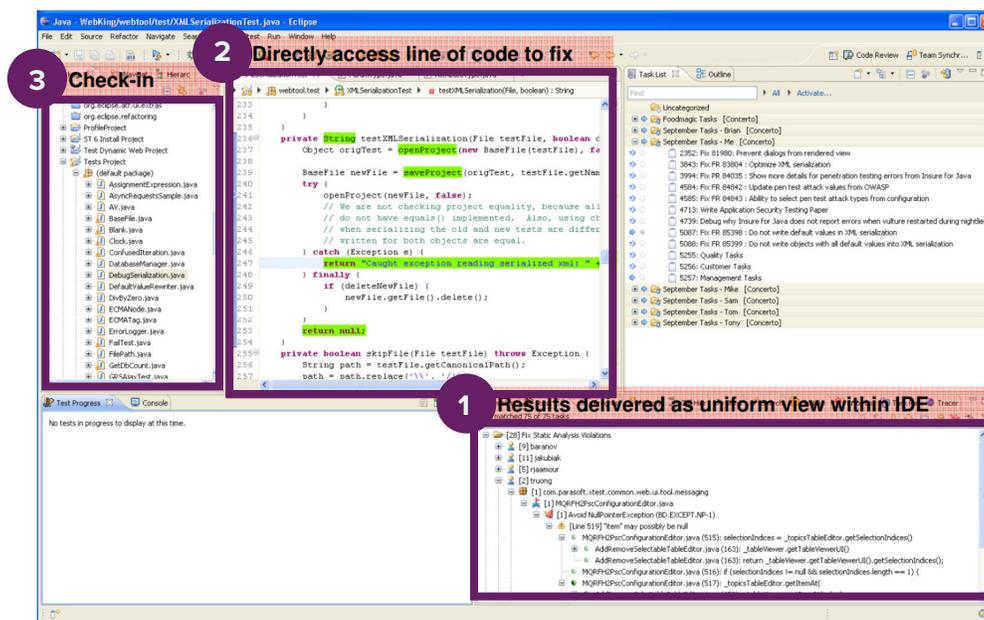


Figure 3:

Example of static analysis tool integration into an IDE

So what does a developer do when analyzing each warning? This is best described in the process diagram shown below:

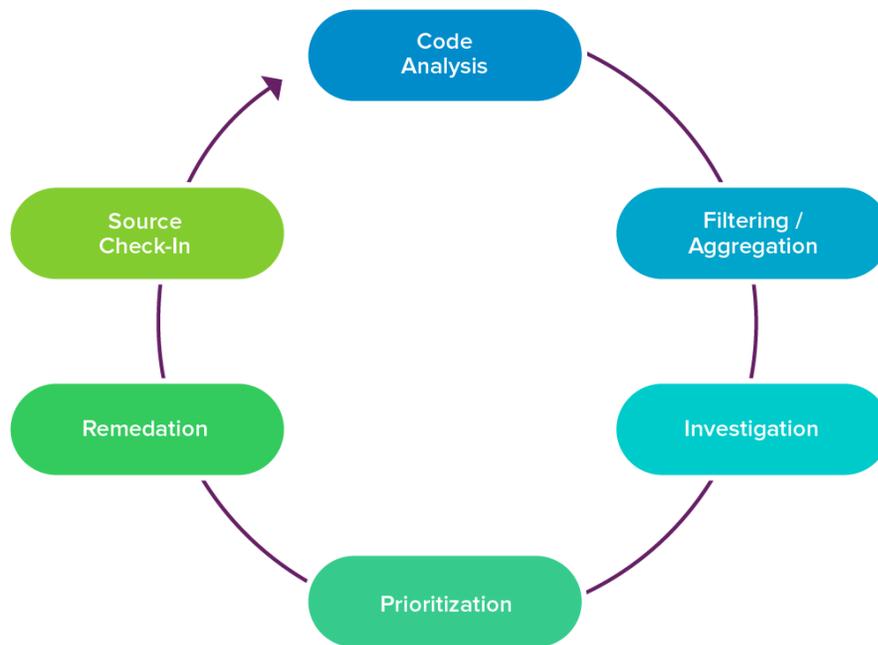


Figure 4:

The static analysis warning lifecycle.

It starts with aggregating and filtering the static analysis results, a critical first step to prioritize and focus on key warnings. Usually, QA and team leaders decide on the quality goal in mind and structure the configuration of the analysis around this goal. To improve security, for example, checkers related to security would be enabled along with a secure coding standard such as CERT C.

Developers then investigate and fix the warnings they find, based on the policies the team has in place and the maturity of the product. As explained above, in a greenfield project, most warnings would be investigated and prioritized since the amount of code would be relatively small, while at later stages of maturity the filtering and prioritization would be stricter so developers would only address truly critical warnings. In either case, the process is the same.

After changes are made in response to a static analysis warning, code is checked into version control and analyzed again during the next build. This short and tight feedback loop greatly improves the quality and security of the code right at the time it's being written and modified.

Integration into Build Systems and Continuous Integration Pipelines

The main integration point for static analysis tools and build systems is through a command line interface. Static analysis used in this fashion acts somewhat like a compiler would in the build structure. Files are processed in the same manner, although the output isn't an executable but rather results that are stored in a repository, indexed by file and build number. With Parasoft C/C++test, this is handled by Parasoft's reporting and analytics system (Parasoft DTP), which is both a repository and an intelligence engine that analyzes the results. This analysis and accompanying information is fed back to each developer to their IDE and made available through Parasoft's web portal for managers and team leads.

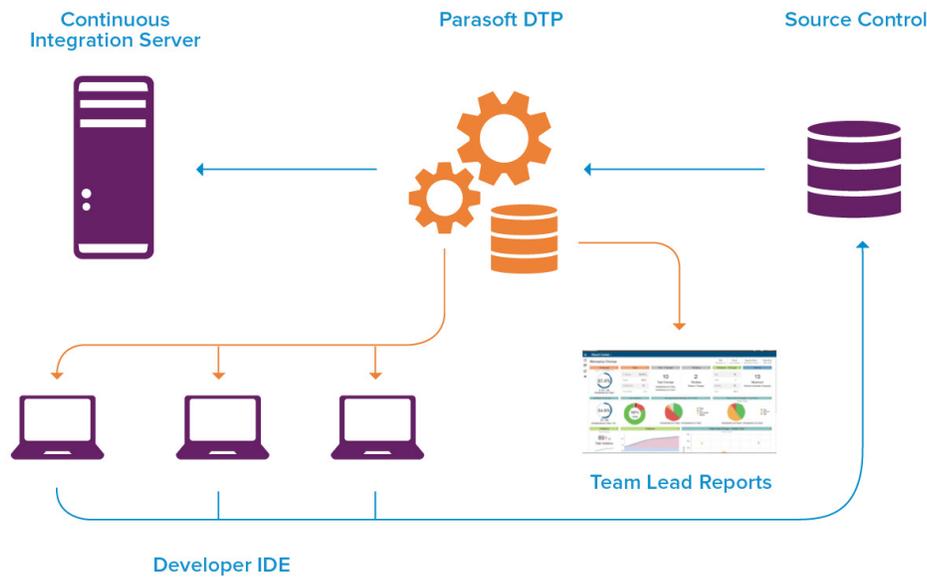


Figure 5: Integrating Parasoft static analysis into a continuous integration pipeline. Parasoft DTP is the central repository and analysis engine for warnings.

Parasoft's automation capabilities are designed to be integrated into any kind of build system used at a customer premises, including continuous integration processes. Parasoft tools integrate with popular CI tools such as Jenkins, Bamboo and Visual Studio Team Services making it easy to include automated unit testing and static analysis as part of an integration, build and deploy cycle. For example, the Parasoft plug in for Jenkins provides a build settings menu to enable the analysis and location for reports.

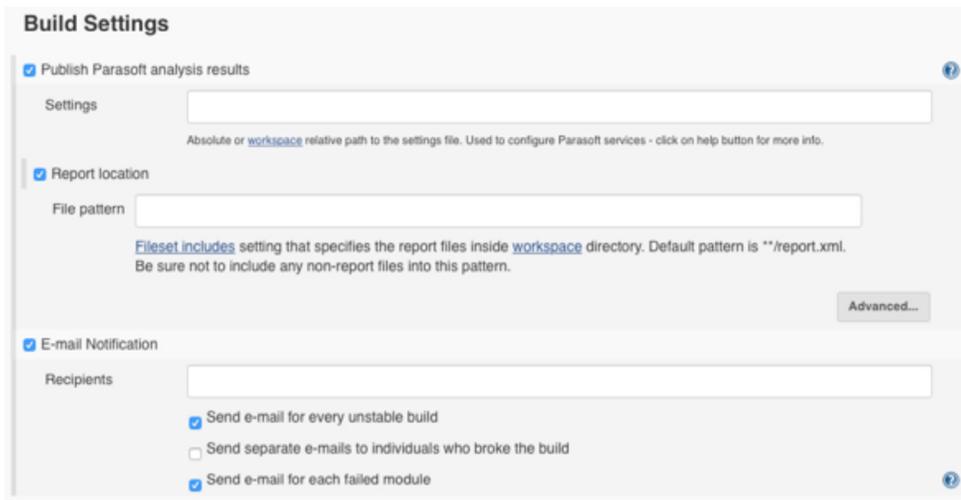


Figure 6: Build settings specifically for Parasoft static analysis within Jenkins.

Also, the plugin includes a custom build environment for Parasoft analysis, integrating the analysis process into an existing pipeline setup. See the following example from Jenkins build menu. Integrating directly with the common tools used in the industry makes including full-project testing and static analysis straightforward.

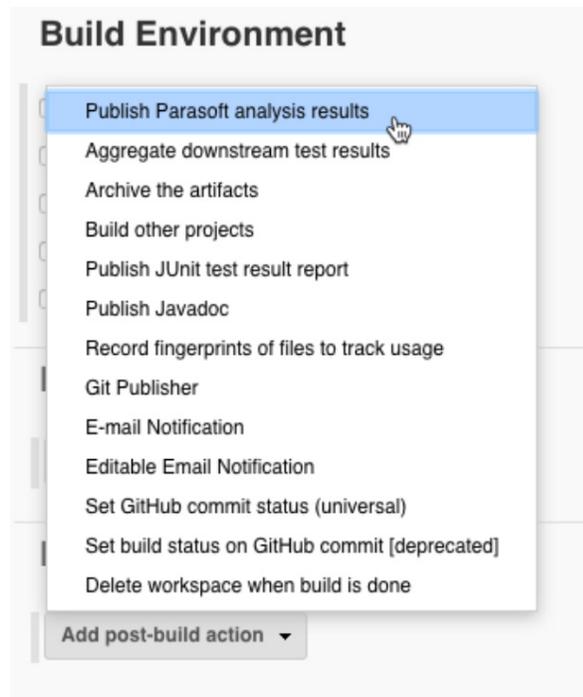


Figure 7:

Custom build environment for Parasoft analysis.

Tackling the Backlog of Warnings and Technical Debt

The next phase of adopting static analysis tools is working to reduce the backlog of warnings and technical debt in a project. In the case of a product under maintenance or in development, there is likely a sizeable backlog to deal with. In the case of a greenfield project, there is less backlog, although the recommendations remain the same for each stage of maturity.

The best starting point for dealing with a backlog of warnings is to prioritize and filter the results based on the desired outcome. Using security as an example, it makes sense to prioritize the backlog in terms of security warnings, ranked by criticality. This is a continuation of the approach used when first introducing static analysis into a project, but now the focus is on the next digestible set of warnings for the team to analyze. This is handled in several ways by Parasoft C/C++test, as we'll get into below.

Dashboards for High Level Analysis

Parasoft's centralized reporting and analytics dashboards provide developers and managers with the ability to see the current state of a project from various viewpoints, and further navigate into more detail where needed to establish a set of warnings for further investigation. Consider the following dashboard showing a project's current compliance with the CERT C coding standard:

CERT C

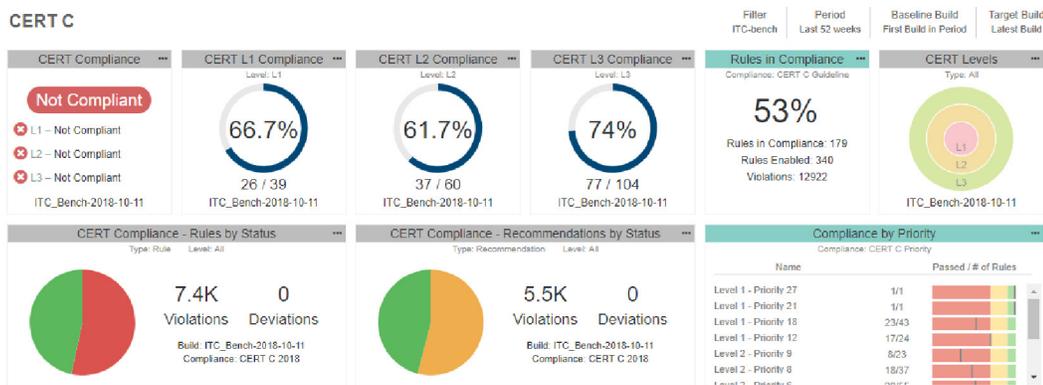


Figure 8:

An example of a web portal dashboard. In this case, a summary of CERT C compliance.

Using this web portal, users can dig deeper into the analysis, down to the file and code level if needed, and investigate warnings entirely within the web portal or in an IDE. Warnings can be prioritized, assigned to developers, suppressed, or marked as a false positive at this stage. See the following example from the Parasoft explorer:

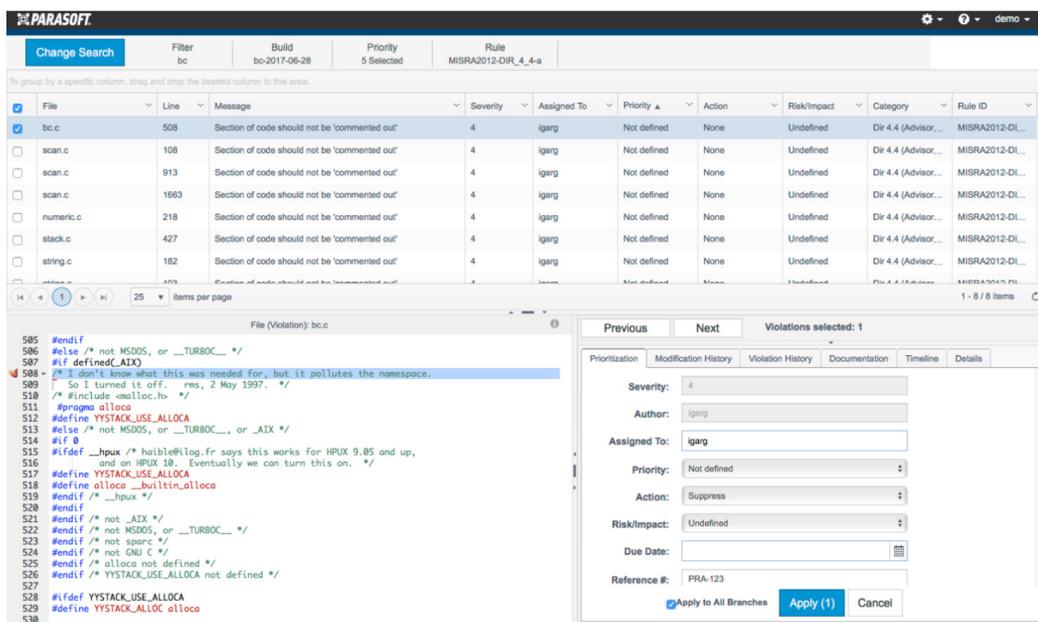


Figure 9:

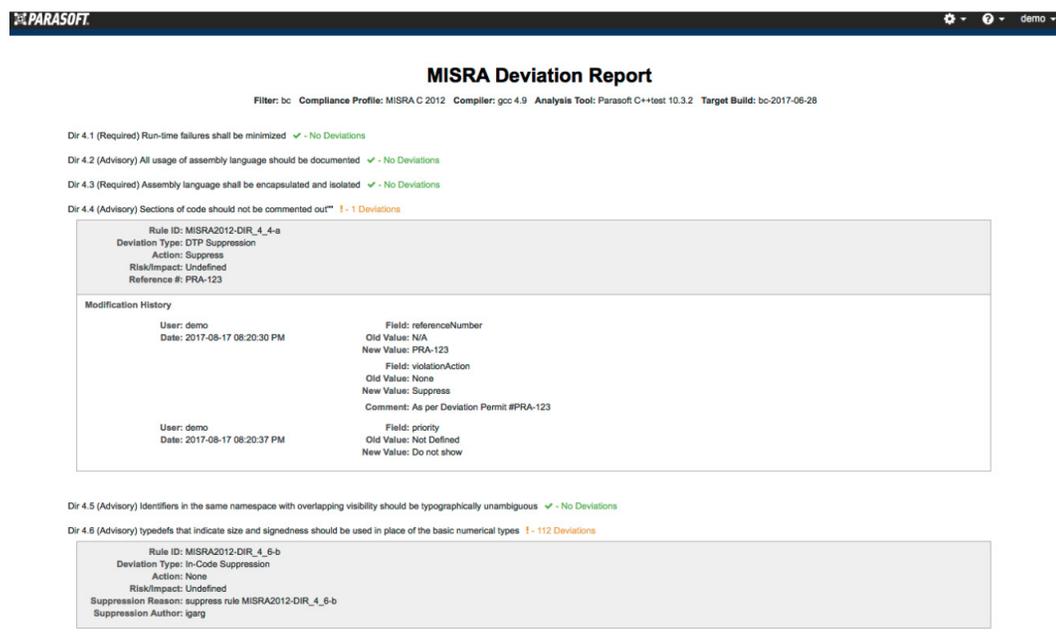
The Parasoft violation explorer.

Managing Violations of Coding Standards

In most cases when analyzing source for coding standard compliance, violations are reported as static analysis warnings. In a large project, there are initially going to be lots of warnings, and managing them quickly and efficiently is critical. Parasoft's violation explorer is the key tool to navigate, evaluate, prioritize, and assign reported errors for remediation. If a static analysis rule violation turns out to be valid but justifiable, considered harmless, or not applicable, a developer can suppress the error and a deviation can be documented. These deviations are reported up through each level of the project to the dashboard and compliance documentation.

To make coding standard compliance feasible for existing projects, it's critical that teams focus on the rules that are considered mandatory first. Compliance is often based on meeting the mandatory requirements with violations of recommended rules if they are documented appropriately. Standards allow rules to be re-categorized, if non-mandatory, allowing for violations if justifiable and documented. Without this, trying to correct every violation becomes infeasible.

Parasoft saves its users many extra hours of work by providing management with a navigable interface to explore violations and automatically generate reports for certification evidence, if needed. An example of a MISRA C deviation report is shown below:



MISRA Deviation Report

Filter: bc: Compliance Profile: MISRA C 2012 Compiler: gcc 4.9 Analysis Tool: Parasoft C++test 10.3.2 Target Build: bc-2017-06-28

Dir 4.1 (Required) Run-time failures shall be minimized ✓ - No Deviations

Dir 4.2 (Advisory) All usage of assembly language should be documented ✓ - No Deviations

Dir 4.3 (Required) Assembly language shall be encapsulated and isolated ✓ - No Deviations

Dir 4.4 (Advisory) Sections of code should not be commented out¹ 1 - 1 Deviations

Rule ID: MISRA2012-DIR_4_4-a
 Deviation Type: DTP Suppression
 Action: Suppress
 Risk/Impact: Undefined
 Reference #: PRA-123

Modification History

User: demo Date: 2017-08-17 08:20:30 PM	Field: referenceNumber Old Value: N/A New Value: PRA-123
	Field: violationAction Old Value: None New Value: Suppress Comment: As per Deviation Permit #PRA-123
User: demo Date: 2017-08-17 08:20:37 PM	Field: priority Old Value: Not Defined New Value: Do not show

Dir 4.5 (Advisory) Identifiers in the same namespace with overlapping visibility should be typographically unambiguous ✓ - No Deviations

Dir 4.6 (Advisory) typedefs that indicate size and signedness should be used in place of the basic numerical types 1 - 112 Deviations

Rule ID: MISRA2012-DIR_4_6-b
 Deviation Type: In-Code Suppression
 Action: None
 Risk/Impact: Undefined
 Suppression Reason: suppress rule MISRA2012-DIR_4_6-b
 Suppression Author: igang

Figure 10:

An example Parasoft MISRA C Deviation Report.

Managing Backlog of Bug and Security Warnings

It's important for teams adopting static analysis to understand that it isn't necessary to fix or analyze all warnings. All warnings are not created equal, so the severity level is the best indicator of how much effort should be placed on investigating and fixing a warning. Continuing the "line in the sand" approach discussed above, when digging into the backlog of warnings, we are effectively moving the line in the sand a bit farther each time.

Parasoft Jtest and Parasoft C/C++test enable users to prioritize and filter warnings within the IDE using configurations. For example, severity and category (type of warning, e.g. security-related) can be used to create a set of warnings suitable for analysis. An example new user configuration is shown below:

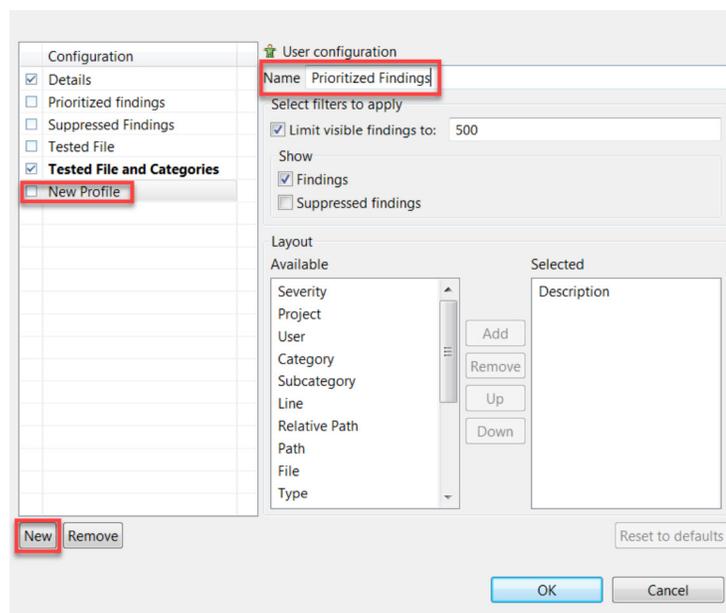


Figure 11:

Custom test configuration settings inside the IDE.

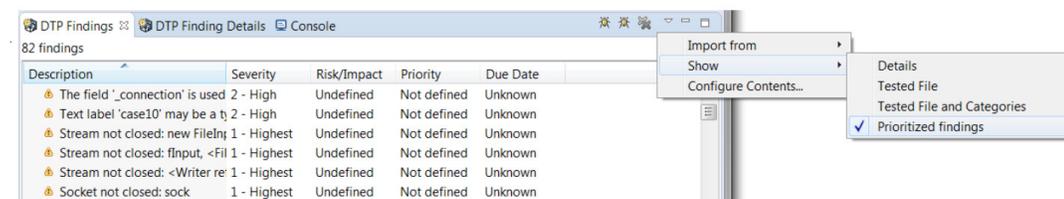


Figure 12:

Configurations can be selected in the DTP Findings view in the IDE.

Incrementally moving the "line in the sand" to tackle the next highest priority and category is the best approach for dealing with a large backlog of warnings. Eventually, a cut off point is reached due to time and budget, but software teams should feel comfortable that they have made significant improvement in quality and security despite any remaining backlog of warnings.

Optimizing Static Analysis

At some point, when a software team has integrated static analysis into their day-to-day activities and everyone is comfortably working with the results, the team will want to customize the tools to better fit their project and organization. This is an important stage of static analysis adoption: optimization. Optimizing static analysis means adapting the rules and the way violations are reported to improve the efficiency and ROI for the organization. Parasoft's static analysis tools support a variety of configuration and customization options to give complete control, if needed, over the static analysis rules and reports.

Define the set of rules to be used: Developers can decide which set of rules are to be enabled, globally, for an entire static analysis run. As mentioned above, it's important to only disable rules that the team doesn't see as useful in the foreseeable future.

Modify settings for data flow analysis: In Parasoft's static analysis tools, flow analysis rules are handled as a special category since they involve more complex analysis. Developers can enable and disable which rules that want to use in this category. There is also the ability to report on un-validated violations (cases where bugs are detected but path analysis has not validated the possibility of this error being possible). This means more potential bugs found (lower false negatives) at the cost of more incorrect findings (false positives).

There are several options for data flow analysis to control the depth of the analysis which impacts performance and accuracy. There are also options for detecting multithreading issues, verbosity of the reports, and custom terminations (functions that terminate a flow of execution). There are also options to control the list of resource functions such as memory creation and deletion (to detect memory leaks, for example). This is also where data sources and sinks are defined for tainted data flow analysis.

Customize rule groups, identifiers, titles and severities: In organizations that have already established coding guidelines, it's possible to customize the warnings reported by Parasoft's static analysis tools to conform to a completely custom set of warnings, descriptions, and categories. These custom configurations can be shared amongst the entire team through the team server as well as directly into each IDE.

Customize existing rules: An important feature of Parasoft's static analysis is the ability to customize both the set of rules to check and the rules themselves — including creating new custom rules. The rules provided with Parasoft tools are a comprehensive set to select from, but specific coding guidelines may differ slightly, or require well-defined exceptions to the rules. Since Parasoft tools use both pattern-based analysis and the more complex flow analysis, customizing each type of rule is different.

Pattern-based rules have checks for pattern matching built into each rule. These are usually modified by parameterizing the rule to suit a user-specific situation. In Parasoft tools, this is easily done with the RuleWizard widget, which provides an easy-to-use interface for customizing existing and new rules (explained in more detail below).

Flow-based rules are more complex and aren't fully customizable, but are parametrized instead, allowing for new applications of the rule logic. An example of this would be modifying the memory leak detection rule to detect leaks from project-specific memory allocation and deallocation functions, or even resource leaks from other resources such as files.

Create new static analysis rules: Software teams can create whole new pattern-based rules, if needed, or base their custom rule on an existing rule. This customization is achieved in the user-friendly RuleWizard, as mentioned above. Rules are created graphically (by creating a flow-chart-like representation of the rule) or automatically (by providing code that demonstrates a sample rule violation).

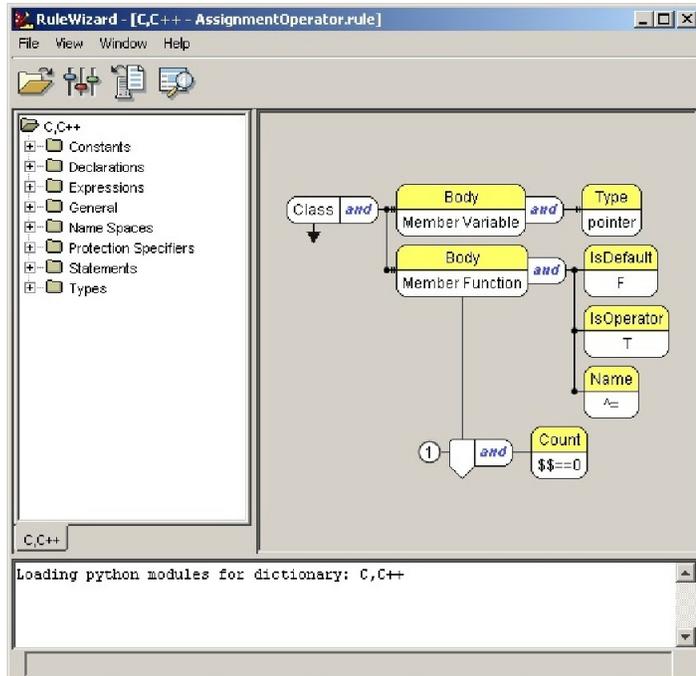


Figure 13:

The RuleWizard provides a visual way to customize static analysis rules.

There are many other customization opportunities with Parasoft static analysis tools, to meet the specific demands of every unique software organization. As adoption of a tool matures, teams can make the most of their investment with customized settings and rules.

Summary

Static analysis tools provide software organizations with the ability to detect and track bugs and security vulnerabilities without needing to execute code. These tools can be applied to existing, legacy, and third-party code, and is also useful for incomplete code, to enable bug detection as soon as code is written, right in the developer's IDE. Parasoft's static analysis tools support industry-leading IDEs and integrate into existing build systems and continuous integration pipelines, while static analysis warnings are stored on a file and build basis alongside test results for prioritization and management.

The adoption of static analysis varies to some extent based on the maturity of the project. A large body of code does result in numerous warnings, but this is completely manageable and the success of adoption depends on how teams decide to tackle the results. Various techniques are introduced for each major maturity level of a project and how these tools can be integrated into the day-to-day workflow for developers, team leads, and managers.

As adoption of Parasoft static analysis tools matures, software teams can optimize their tool chain by customizing the tools themselves to meet their specific needs, increasing their ROI and achieving their project's end goals.



Parasoft Corporation

101 E Huntington Drive Monrovia, CA 91016 USA

Sales: 1-888-305-0041 | International Sales: +1-626-256-3680

www.parasoft.com