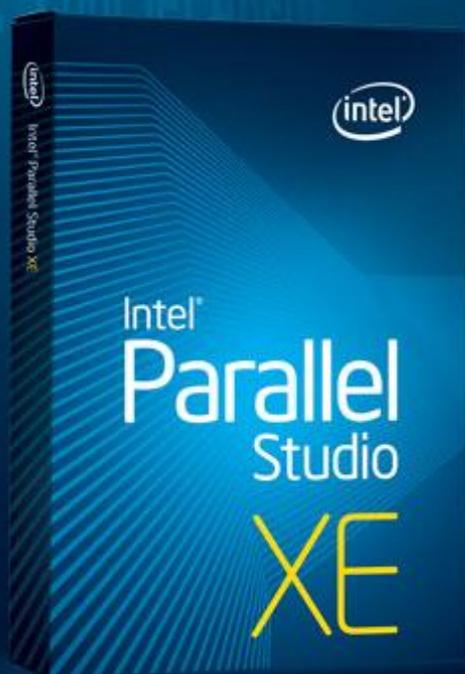




# Improve C++ Code Quality with Static Analysis



This document is an introductory tutorial describing how to use static analysis on C++ code to improve software quality, either by eliminating bugs or making the code more secure, or both. It provides a brief description of the goals of the product feature and walks through an end-to-end example showing how it is used.

Static Analysis is available in the Intel® Parallel Studio XE and Intel® Cluster Studio XE suites.



# Improve C++ Code Quality with Static Analysis

## Introduction

This document is an introductory tutorial describing the static analysis feature of the Intel® Parallel Studio XE. It provides a brief description of the goals of the product feature and walks through an end-to-end example showing how it is used.

## What is Static Analysis?

Static analysis attempts to identify errors and security weaknesses through deep analysis of source code. Static Analysis is primarily aimed at developers and QA engineers who wish detect software defects early in the development cycle in order to reduce time, cost and increase ROI. It also assists developers in hardening their application against security attack. Static analysis provides an effective way to discover defects, especially in code that is hard to exercise thoroughly with tests. Static analysis can also detect race conditions resulting from misuse of parallel programming frameworks such as OpenMP and Intel Cilk Plus.

Many coding errors and patterns of unsafe usage can be discovered through static analysis. The main advantage of static analysis over dynamic analysis is that it examines all possible execution paths and variable values, not just those that are provoked during testing. This aspect of static analysis is especially valuable in security assurance, since security attacks often exercise an application in unforeseen and untested ways.

Static analysis can detect over 250 different error conditions. They include buffer overflows and boundary violations, misuse of pointers and heap storage, memory leaks, use of uninitialized variables and objects, unsafe/incorrect use of C/C++ or Fortran features and libraries.

## How does Static Analysis Work?

Static analysis is performed by the Intel® C++ and the Intel® Fortran Compilers operating in a special mode. In this mode the compiler dedicates more time to error analysis and bypasses the instruction generation process entirely. This allows it to find errors that go undetected during ordinary compilation. Static analysis requires your code to compile without serious errors using the Intel Compiler, but you do not execute the results of this compilation. You do not need to use Intel compiler to create your production binaries in order to take advantage of static analysis. We will begin the tutorial with showing how to prepare an application for static analysis.

Static analysis can be done on a partial program or a library, but it works best when operating on an entire program. Each individual file is compiled into an object module, and the analysis results are produced at the link step. The results are then viewed in Intel® Inspector XE.

The results of static analysis are often inconclusive. The tool results are best thought of as potential problems deserving investigation. You will have to determine whether a fix is required or not. The Intel Inspector XE GUI is designed to facilitate this process. You indicate the results of your analysis by assigning a state to each problem in the result. Intel Inspector XE saves the state information in its result files.

Over time the source will change and you will want to repeat the analysis. The first time you open your new analysis result in Intel Inspector XE, it automatically calculates a correspondence between the previous result and the new result. It uses this correspondence to initialize the state information for the new result. This means you do not have to investigate the same issues over again.

When you decide a problem does require fixing you should report it into your normal bug tracking system, just as you would report a defect detected by an executable test. Intel Inspector XE is not a bug tracking system. The state information tracks your progress in investigating the result of analysis. What you do with your conclusions is outside the scope of static analysis and the Intel Parallel Studio XE.



## Improve C++ Code Quality with Static Analysis

### Setting up for Static Analysis

The set up process is usually fairly simple, but it can be quite complicated, depending on the situation. For Microsoft\* Windows\* when using Microsoft\* Visual Studio\*, set up is automated with a simple menu/dialog driven approach that automatically configures your solution for static analysis. Details on how to use this and how it works are in the tutorial that follows.

In Linux\* OS or when building with a make file or command line script in Windows we recommend you modify your build procedure (whatever it is) to create a new build configuration for static analysis. The term build configuration refers to a mode of building your application, using specific compiler options and directing the intermediate files to specific directories. Most applications have at least two build configurations: debug and release. You will want to create one more for static analysis. The static analysis configuration must build with the Intel compiler with additional options set to enable static analysis.

Please note the distinction between creating a new build configuration and building an existing configuration with additional options. If you build, say, your debug configuration with the additional options that enable static analysis then you will get analysis results (as long as your debug configuration builds with the Intel compiler). This is a perfectly good way to do an initial product evaluation. However, it's awkward to work this way on an ongoing basis, because the object modules produced by static analysis will overwrite your debug-mode object modules every time you run static analysis. Just as you want to keep debug object modules separate from release object modules, you will want to keep debug object modules separate from static analysis object modules. If you are going to use static analysis on an ongoing basis, you will want to get set up properly.

The process for creating a new build configuration will be different for each application. The sample applications used in this tutorial can be built under Visual Studio on Microsoft Windows OS or with a make file on Linux OS. We will walk through the steps needed to update this make file for static analysis.

If your application build is very complex and you don't feel confident that it can be modified safely, there is an alternative set up method. You can execute your normal build under a "watcher" utility called `inspxe-inject`. This application intercepts process creations and recognizes all the compilation and link steps performed during your build. It records this information in a build specification file. This file can be supplied as input to another utility, `inspxe-runsc`, which invokes the Intel compiler to repeat the same build steps as your original build did. These utilities are not explained in this tutorial.

# Improve C++ Code Quality with Static Analysis

## Start of C++ Tutorial

This tutorial can be executed on Windows or on Linux OS. It can also be done using a C++ sample application. The differences between Linux OS and Windows OS are small and these cases are described together. When we come to a step that is done differently on Windows and Linux OS, we will say something like “On Windows OS, do ...”

## Download and Install Intel® Parallel Studio XE

If you don't have Intel Parallel Studio XE installed, do that first. Estimated completion time: 15-30 minutes

1. [Download](#) an evaluation copy of Intel Parallel Studio XE.
2. Install Intel Parallel Studio XE by clicking on the `parallel_studio_xe_2013_setup.exe` (can take 15 to 30 minutes depending on your system).

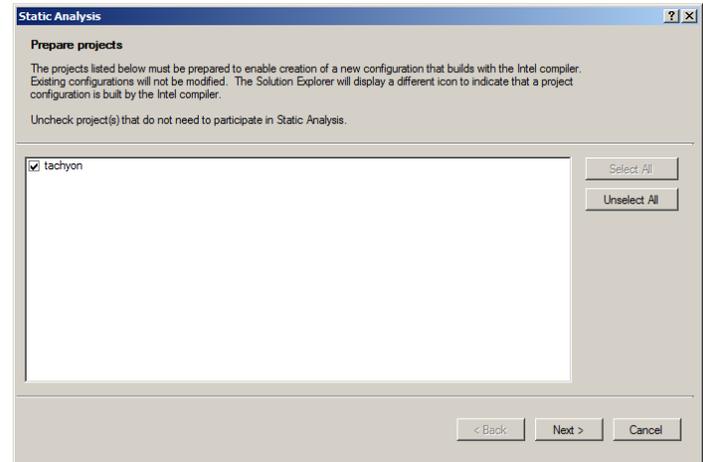
We will start by getting set up for static analysis. This is done differently on Windows and Linux OS. For Windows OS read the following section. For Linux OS, go to [Setting up and Running Static Analysis Using a make File on Linux\\* OS](#).

### Setting up and Running Static Analysis Using Microsoft\* Visual Studio\* on Windows\* OS

We will go through the setup process using Visual Studio solution `tachyon.sln` that is supplied with the sample program. Start by opening this solution in Visual Studio. The process for setting up for static analysis is quite simple. If you are using Visual Studio\* 2008 or Visual Studio 2010 then Visual Studio will need to convert for solution file from Visual Studio 2005 format. Let it go ahead and do this conversion.

Set up for static analysis by selecting the **Build > Build Solution for Intel Static Analysis** menu item.

1. For Visual Studio 2005 and 2008 the 'Prepare projects' dialog appears as shown in the figure below.

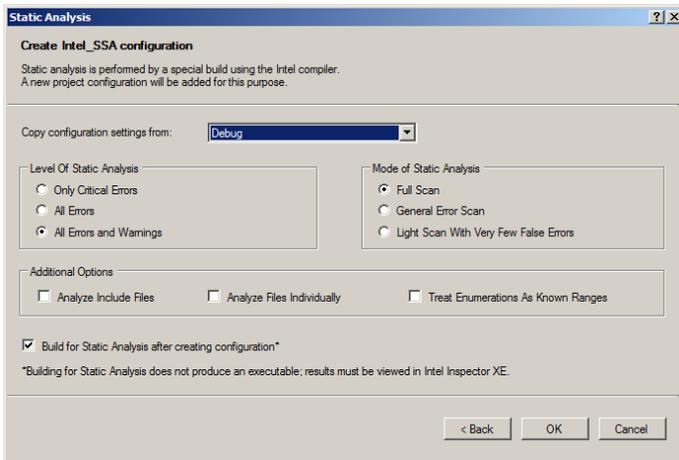


2. Click on 'Next'. When we're finished, this will convert the project from Visual C++ format to Intel format. You can tell what projects are in Visual C++ versus Intel format by looking at their icon in the solution explorer. Intel format projects have the purple Composer icon. This conversion is done in such a way that existing configurations continue to build with the Visual C++ compiler. This step is not needed for Visual Studio 2010.
3. The 'Create configuration' dialog is used to choose an existing baseline configuration and specify the static analysis options.

The 'Level Of Static Analysis' can be set to one of three values:

- a) 'Only Critical Errors' reports only diagnostics with a severity level of "critical". Diagnostics with a severity level of "error" or "warning" are not displayed.
- b) 'All Errors' Reports all diagnostics with a severity level of "critical" and "error". Diagnostics with a severity level of "warning" are not reported.

## Improve C++ Code Quality with Static Analysis



- c) 'All Errors and Warnings' reports all diagnostics regardless of the severity level.

The 'Mode of Static Analysis' can be set as follows:

- 'Full Scan' mode attempts to find all program weaknesses, even at the expense of more false positives. This mode is recommended when using the tool for security assurance. This is the default.
- 'General Error Scan' mode attempts to reduce false positives somewhat more than reducing false negatives. This mode is recommended when using the tool for general error detection.
- 'Light Scan With Very Few False Errors' mode attempts to avoid all false positives, even at the expense of substantially reduced error checking. This mode is recommended when using the tool for acceptance screening.

Click 'OK' to accept the default settings. This creates a new Intel\_SSA configuration. It copies the properties from the baseline (Debug) configuration and then changes it to 1) build with the Intel compiler and 2) enable the selected static analysis related options. Since the checkbox next to "Build for Static Analysis after creating configuration" is checked, it will immediately build this configuration, which is equivalent to running static analysis.

That's it; you are now set up for static analysis and have launched your first analysis.

For information on setting up applications built using a command line script or make file, see the following section explaining the setup process on Linux OS.

### Setting up and Running Static Analysis Using a make File on Linux\* OS

The sample application comes with a make file that can be used to build the application on Linux OS. To save time, we have included two versions of the make file. One is the original make file, **tachyon.mk**. This make file builds the application using the gcc compiler. The other is the updated make file, **tachyon\_ssa.mk**, that adds a new build target, SSA. The SSA build target is just like the debug build target except for these changes:

- It builds using the Intel compiler
- It adds the additional option "-diag-enable sc3". This will cause static analysis to report all diagnostics regardless of the severity level. This is equivalent to specifying 'All Errors and Warnings' in Windows. Alternatively, "-diag-enable sc1" or "-diag-enable sc2" could be specified instead. "-diag-enable sc1" reports only diagnostics with a severity level of "critical". Diagnostics with a severity level of "error" or "warning" are not displayed. "-diag-enable sc2" reports all diagnostics with a severity level of "critical" and "error". Diagnostics with a severity level of "warning" are not reported.
- It puts the intermediate files in the SSA subdirectory.

You can compare the two make files to see these differences.

Additionally, -diag-enable sc-mode can be added to the command line to specify the analysis mode. The mode can be any of the following: full, concise, or precise.

- 'full' mode attempts to find all program weaknesses, even at the expense of more false positives. This mode is recommended when using the tool for security assurance. This is the default when -diag-enable sc[1|2|3] is present.

# Improve C++ Code Quality with Static Analysis

- 'concise' mode attempts to reduce false positives somewhat more than reducing false negatives. This mode is recommended when using the tool for general error detection.
- 'precise' mode attempts to avoid all false positives, even at the expense of substantially reduced error checking. This mode is recommended when using the tool for acceptance screening.

This illustrates the changes that are needed to update a make file to prepare for static analysis.

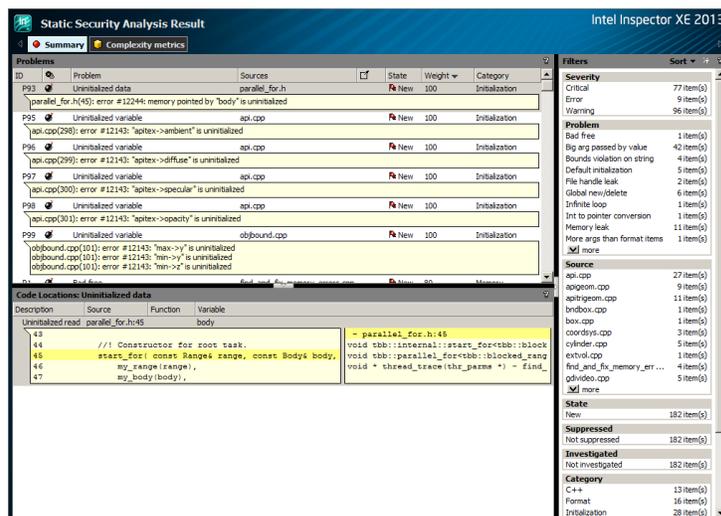
Now that you are set up, all you have to do is build your static analysis build configuration to perform analysis. You can do this by performing the following steps:

- Open a command shell
  - Set the environment variables for the Intel compiler by executing the `iclvars.sh` script in the compiler bin directory, supplying the `ia32` option.
  - Execute `make -f tachyon_ssa.mk`.
- TIP: keep this command window open for later operations.

## Interacting with the Analysis Results

On Windows OS, the Intel® Inspector XE automatically opens a new result as soon as the build completes. On Linux OS, type `inspxe-gui` to invoke the Intel Inspector XE and then use the **File > Open** menu to open the result. By default, the file you want to open is called `r000sc.inspxe`, and is contained in a directory named `r000sc` below the root directory of the tachyon project. The remainder of this tutorial is almost identical for Windows and Linux OS. The main difference is that on Windows OS, the Intel Inspector XE GUI is embedded within Visual Studio, while on Linux OS the GUI runs as a stand-alone program. The look of the individual Intel Inspector XE windows is almost identical on Windows and Linux OS.

The initial window you will see will look something like this.



This window consists of three main areas. The upper left pane is the table of Problem Sets. This is your "to do" list, the things you need to investigate. The lower left pane shows the code locations corresponding to the currently selected problem set. The right pane shows the filters. It controls what problem sets are displayed and which are hidden.

You can sort the table of problem sets by clicking on any of the column headers. By default the problem set table is sorted by **weight**. The weight is a value between 1 and 100 which reflects how interesting a problem is. Problems that can do more damage have higher weight. Problems that are more likely to be true problems (as opposed to false positives) are also given higher weight. So the weight provides a natural guidance for your order of investigation.

Let's start with an easy one. P67 is an unsafe format specifier. Let's see what we can learn about this problem.

# Improve C++ Code Quality with Static Analysis

## Investigating a Problem

Initially, all we know about the problem is summarized in the table entry:

ID	Problem Type	File	Line	Weight	Category
P67	Unsafe format specifier	parse.cpp	New	70	Format

**Unsafe format specifier** is the short description of the problem. The full description is shown in the shaded area. The “New” entry in the state column indicates that this problem was discovered for the first time in this analysis run and has not yet been investigated. The 70 weight value indicates the weight. The category of problem is “Format”, which means it is related to misuse for printf-style format specifications.

Click on this problem to select it. The lower pane refreshes itself to show the source code locations related to this problem. Here’s what it says:

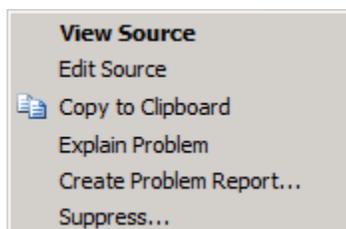
Description	Source	Function	Variable
Format mismatch	parse.cpp:187	unsigned int GetString(_iobuf *,char const *)	

```

185 char data[255];
186
187 fscanf(dfile, "%s", data);
188 if (strcmp(data, "string") != 0) {
189     fprintf(stderr, "parse: Expected %s, got %s \n", strl
  
```

Here we see the source location (file parse.cpp, line 187) and a small snippet at the referenced location. We also see the role that this source reference plays in the problem. **Format mismatch** indicates that this is a place where a format string was used. It also shows the name of the function that contains the line (GetString) and its parameter signature.

One way we could get more information about a problem is to read an explanation of what that problem type means. Some static analysis errors are pretty technical and require explanation. To see the explanation for this problem type, right click the problem. That brings up this pop-up menu:



Select **Explain Problem** to open a help topic that explains this problem in detail. This is the help topic for this particular problem type:

### Unsafe format specifier

Some forms of formatted input can cause buffer overflow and should not be used.

Care must be taken on formatted input to avoid buffer overflow. In particular, the “%s” input format is inherently unsafe. A better alternative is “%*ddd*s” or “%\*s”, where *ddd* is the sized of the destination buffer, for example “%24s”. If the buffer size is not a compile time constant, then “%\*s” can be used, where “\*” obtains the maximum size from the next input argument, for example, `scanf("%*s", sizeof(buffer), buffer);`

ID	Observation	Description
1	Format mismatch	The unsafe formatted input statement

#### Example

```

#include <stdio.h>
char buffer[1024];
int main(int argc, char **argv)
{
    scanf("%s", buffer); // unsafe: could overflow buffer
    // better is scanf("%24s", sizeof(buffer), buffer);
    printf("read %s\n", buffer);
    return 0;
}
  
```

Copyright © 2010, Intel Corporation. All rights reserved.

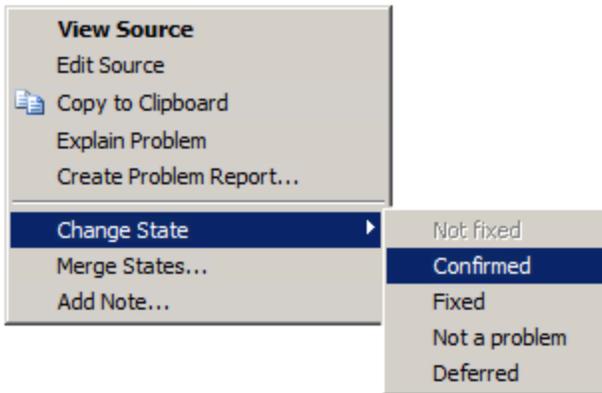
As you can see, the problem type reference material explains more fully what the precise error condition is and its potential consequences. It explains more fully the role the various code locations play in creating the problem, and provides an example that demonstrates the problem. Where possible, it also provides better alternatives to avoid this issue

The next thing we need to do is determine whether this problem is really present in our application. To do that, we need to look at the source code. The small snippet is pretty clear. The highlighted call to `fscanf` has a format string with a “%s” format specifier. This reads input characters up to the next newline and stores the data in the array “data”. There is no guarantee that the number of characters read will not overflow the bounds of the array, so this statement could corrupt memory. We got lucky here, because the code snippet contained all we needed to know about this problem. We will see later how we see more of the source when we need to.

Since we have confirmed that this is a real error, let’s record our conclusion. Right click on the problem and select **Change State > Confirmed** from the pop-up menu.



# Improve C++ Code Quality with Static Analysis

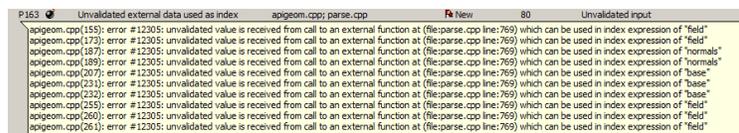
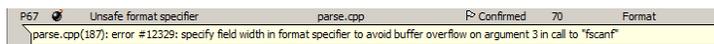


Notice how the problem set table has also redrawn to show only problems in this source file. You can turn off a filter by clicking on the “All” box, but leave it turned on for now.

## Investigating a Second Problem

Take a look at another problem in the solution. This time pick problem P163, which should be the first one you now see in the table of problem sets:

Now we’re done with that problem. You can see the state is updated in the problem set table:



## Reducing Clutter with Filtering

Filters allow you to focus on the problems you are interested in and hide the problems you want to ignore.

Once a problem has been investigated there is usually no reason to look at it again. One of the nicest uses for filters to hide all the problems you are finished investigating. Go all the way to the bottom of the filter window and click on the “Not investigated” item inside the “Investigated” filter.



When you do that, the filter item redraws to indicate that it is active:



Notice how that problem we marked as “Confirmed” disappeared from the table of problem sets when we did that. It a good idea to keep this filter set like this while you work on analyzing results.

The first several filters, Severity, Problem, Source, State, and Category, correspond to columns in the table of problem sets. You can hide all rows in the table that do not match a specific value in some column. Click on the second line in the Source filter (apigeom.cpp). It will look like this:

This problem is a little more interesting. As you can see, it has several source locations in different files. The problem type is “Unvalidated external data used as index”. The full description describes several places where an unvalidated variable is used to compute and array index possibly leading to an out of bounds condition. In all cases the place where the variable was set is the same (parse.cpp, line 769).

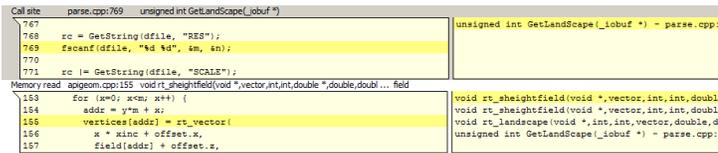
This illustrates why we call this the table of **problem sets** instead of **problems**. Here static analysis has combined several related problems into one problem set because the problems may have a common solution.

Let’s investigate to see what is happening here. As before, select this problem set and look at the lower pane to see the related source code. Let’s focus on first error listed in the problem set:

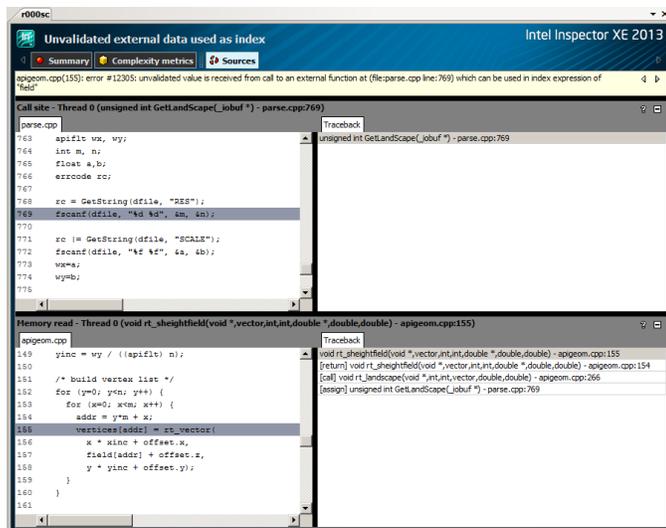


Here on the left we see two code locations and on the right corresponding traceback. The first code location is the place where the ‘m’ which is used to compute the array index was assigned through a call to fscanf (Call site) and the other where the index used in array ‘field’ is computed.

# Improve C++ Code Quality with Static Analysis



To fully view the source, double click on the snippet containing apigeom.cpp:155.



By examining the source and the corresponding traceback it can be seen that 'm' is not validated before it is used to compute the index 'addr' used by the array 'field'.

## Fixing a Problem and Rescanning

Now go ahead and fix this problem. To fix this problem, add the following to the beginning of the function 'rt\_sheightfield' to perform range checking on 'm'.

```

if (m < 1) {
    m = 1;}
else if (m > 50) {
    m = 50;}
  
```

To modify the source we need to get into a real source editor on line 140 in apigeom.cpp, just before the malloc call. You can enter your normal source editor by double-clicking the line in the sources view, or by right-clicking that line and selecting **Edit Source** from the pop-up menu. On Linux OS this will activate whatever editor is selected by the EDITOR environment variable. On Windows OS this opens the Visual Studio source editor.

On Windows OS the editor window will open a new tab in the same tab group window that contains the results. So opening the source for editing will hide the result. To view the results again, click the **r000sc** tab.



If you have a big screen, you can put the result and the source files you are editing in different tab groups.

You can see the following in the editor:

```

...
apiflt xinc, yinc;

vertices = (vector *) malloc(m*n*sizeof(vector));
normals = (vector *) malloc(m*n*sizeof(vector));
...
  
```

Rewrite this code like this:

```

...
apiflt xinc, yinc;
int x, y, addr;

if (m < 1) {
    m = 1;}
else if (m > 50) {
    m = 50;}

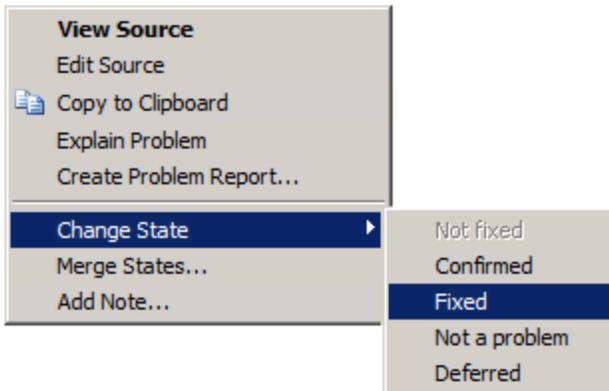
vertices = (vector *) malloc(m*n*sizeof(vector));
normals = (vector *) malloc(m*n*sizeof(vector));
...
  
```

Type or copy-paste these four code lines into the source exactly as shown, save the result, and close the editor but do not rebuild the application yet. In fact this change does not entirely fix the problem set, but make the change anyway.

Now go back into the summary view. In Visual Studio, click on the **r000sc** tab to get back to the Intel Inspector XE GUI. Click on the box that says **Summary** at the upper left of the Sources view.

Look at the problem set table and look at P163 again. The fix applied only addresses the errors found on lines 155, 173, 187 and 189 of apigeom.cpp, but does address the remaining errors listed. Still, go ahead and use the right-click pop-up menu to change the state of P163 to **Fixed**.

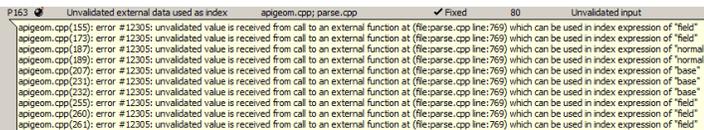
# Improve C++ Code Quality with Static Analysis



Notice that as soon as you do this, the problem disappears, since you set the filter set to show only uninvestigated problems. To see that problem again, go over to **Filters** pane, go down to the bottom and turn off the Investigated filter by clicking on All:



Now you can see the problem again, and see that it really is marked as Fixed.

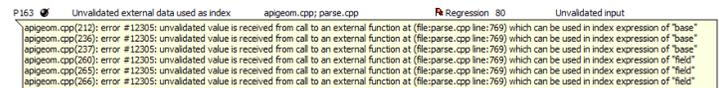


Now, go ahead and rebuild the application to create a new analysis result and open the new result, r001sc. On Windows OS, you can do all this by selecting the **Build>Rebuild solution for Static Analysis...** menu item. You will want to close the Visual Studio Output window when the build completes. On Linux OS you should repeat the steps you used earlier to build (type `make -f "tachyon_ssa.mk"` in your command window, then use **File > Open** in the Intel Inspector GUI to open r001sc).

The first thing you will notice is that most problems now say "Not fixed" instead of "New" in their state. This demonstrates the way that Intel Inspector XE automatically initializes the state in the new result, r001sc, from the previous result, r000sc. Problems that were seen before are no longer considered "New". They are simply not investigated yet, which is what the "Not fixed" state indicates.

You will also notice that the problem we investigated earlier, the unsafe format problem, is still in a "Confirmed" state, just as we marked it in r000sc.

Use the filter to select only the problems in file apigeom.cpp. Look at the problem that is the top of the list. Here's what it looks like now:



Intel Inspector XE set this problem set to a **Regression** instead **Fixed**, since not all the problems in the problem set were addressed. **Regression** is considered an uninvestigated state, so this problem would still be seen if you set the filter to show only problems whose investigation state is **Not investigated**.

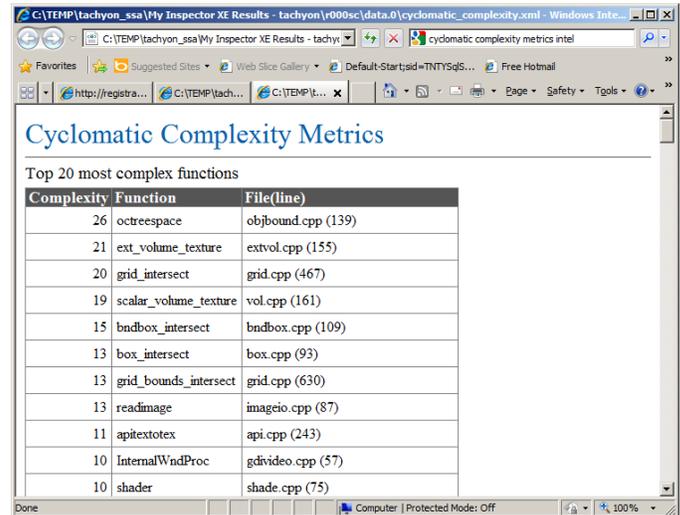
## Improve C++ Code Quality with Static Analysis

### Cyclomatic Complexity Metrics Report

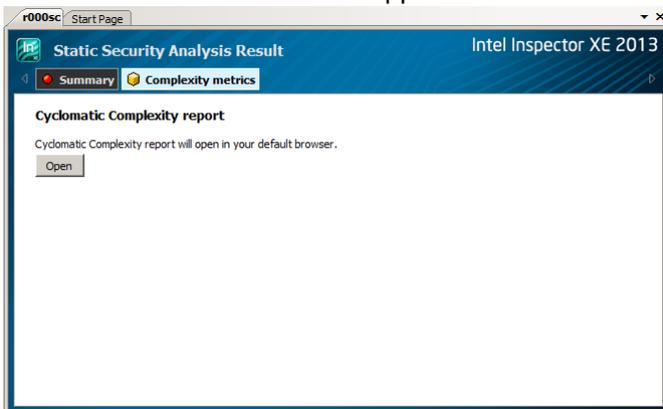
Intel® Composer XE 2013 is able to automatically generate Cyclomatic Complexity metrics as part of static analysis.

Cyclomatic Complexity metrics (CC) have two recommended uses in software development. First, controlling complexity can make code more maintainable to reduce bugs. Code with CC measurements that exceeds a threshold value could be reviewed and refactored to reduce complexity when appropriate. Code whose measurements exceed a higher threshold could be required to be refactored unless a formal review determines that the level of complexity is acceptable. Secondly, the CC value can be used to estimate the number of tests required to thoroughly exercise all application paths.

Intel® Inspector XE 2013 provides an ability to visualize Cyclomatic Complexity metrics. A new tab “Complexity metrics” is added to Static Analysis result view. The ‘Open’ button shows the Cyclomatic Complexity HTML report in the default browser. The report shows the top 20 most complex functions as well as the complexity of the full list of functions for the application.



Complexity	Function	File(line)
26	octreespace	objbound.cpp (139)
21	ext_volume_texture	extvol.cpp (155)
20	grid_intersect	grid.cpp (467)
19	scalar_volume_texture	vol.cpp (161)
15	bndbox_intersect	bndbox.cpp (109)
13	box_intersect	box.cpp (93)
13	grid_bounds_intersect	grid.cpp (630)
13	readimage	imageio.cpp (87)
11	apitextotex	api.cpp (243)
10	InternalWndProc	gdvideo.cpp (57)
10	shader	shade.cpp (75)





# Improve C++ Code Quality with Static Analysis

## Summary and Review

To review, we have seen:

- 1) The table of problem sets summarizes the problems found by static analysis.
- 2) Selecting a problem set shows the associated code locations in the lower pane.
- 3) Dig into a problem by viewing code snippets, going to the **Sources** view (where you can see tracebacks), or going to your source editor.
- 4) Get a full explanation of what a problem type means by using the "Explain Problem" menu item to access the problem type reference material.
- 5) Set the state of a problem set to record the results of your investigation.
- 6) Use filters to hide problems that you are done investigating, or to focus on particular source files or particular classes of problems.
- 7) Static analysis sometimes combines related problems into problem sets to reduce investigation time.
- 8) Static analysis automatically carries state information forward from result to result. It keeps track of what problems are new, what problems were investigated, and verifies that fixed problems are really fixed.
- 9) Intel® Composer XE 2013 through static analysis is able to generate Cyclomatic Complexity metrics.

## Additional Resources

[Intel Learning Lab](#) - Technical videos, whitepapers, webinar replays and more.

[Intel Parallel Studio XE product page](#) - How to videos, getting started guides, documentation, product details, support and more.

[Evaluation Guide Portal](#) - Additional evaluation guides that show how to use various powerful capabilities.

[Intel® Software Network Forums](#) - A community for developers.

[Intel® Software Products Knowledge Base](#) - Access to information about products and licensing,

[Download a free 30 day evaluation](#)



# Improve C++ Code Quality with Static Analysis

## Purchase Options: Language Specific Suites

Several suites are available combining the tools to build, verify and tune your application. Products that include static analysis are highlighted in blue. Single or multi-user licenses along with volume, academic, and student discounts are available.

Suites >>		Intel® Cluster Studio XE	Intel® Parallel Studio XE	Intel® C++ Studio XE	Intel® Fortran Studio XE	Intel® Composer XE	Intel® C++ Composer XE	Intel® Fortran Composer XE
Components	Intel® C / C++ Compiler	●	●	●		●	●	
	Intel® Fortran Compiler	●	●		●	●		●
	Intel® Integrated Performance Primitives <sup>3</sup>	●	●	●		●	●	
	Intel® Math Kernel Library <sup>3</sup>	●	●	●	●	●	●	●
	Intel® Cilk™ Plus	●	●	●		●	●	
	Intel® Threading Building Blocks	●	●	●		●	●	
	Intel® Inspector XE	●	●	●	●			
	Intel® VTune™ Amplifier XE	●	●	●	●			
	Intel® Advisor XE	●	●	●	●			
	Static Analysis	●	●	●	●			
	Intel® MPI Library	●						
	Intel® Trace Analyzer & Collector	●						
	Rogue Wave IMSL* Library <sup>2</sup>							●
Operating System <sup>1</sup>	W, L	W, L	W, L	W, L	W, L	W, L, O	W, L, O	

Note: <sup>1</sup> Operating System: W=Windows\*, L= Linux\*, O= OS X\*. <sup>2</sup> Available in Intel® Visual Fortran Composer XE for Windows with IMSL\*

<sup>3</sup> Not available individually on OS X, it is included in Intel® C++ & Fortran Composer XE suites for OS X



Learn more about Intel Parallel Studio XE

- Click or enter the link below:  
<http://intel.ly/parallel-studio-xe>
- Or scan the QR code on the left



Download a free 30-day evaluation

- Click or enter the link below:  
<http://intel.ly/sw-tools-eval>
- Click on 'Product Suites' link

## Notices

INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF INTEL PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT.

### Optimization Notice

Notice revision #20110804

Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice.

