

# C++ Software Quality in the ATLAS Experiment

Stefan Kluth, Stewart Martin-Haugh, Emil Obreshkov,  
*Shaun Roe*, Rolf Seuster, Scott Snyder, **Graeme Stewart**

Tools in use

Experience


Progress, Pleasures and Pitfalls

# TOOLS

Gcc-plugin | tctoolkit  
lwyu | PRQA  
**Coverity**  
**Cppcheck**  
Clang | Ubsan


Compilation \* Static Analysis \* run-time behaviour  
continuous testing

# ATLAS code



**~420**

**developers**



correspondence →



~140 'domains'  
e.g. Tracking-Fitting  
2400 individual packages



**(gcc 4.9)**

# Compiler checks

Code should always, where possible, compile without warnings.

It is beneficial to expose the code to different compilers; e.g. Clang builds give additional info:

- Mismatched struct/class; Unused private class members; Mismatched header guard.
- `std::abs(eta < 2.5)`
- `if ((bitfield && 0x0011) != 0) ...`
- `!val==10`

## gcc plug-ins

- Check for inheritance structure already in place
- Coming soon: naming convention checks

# Static Analysis

Attempts to follow your program and points out possible errors; some examples:

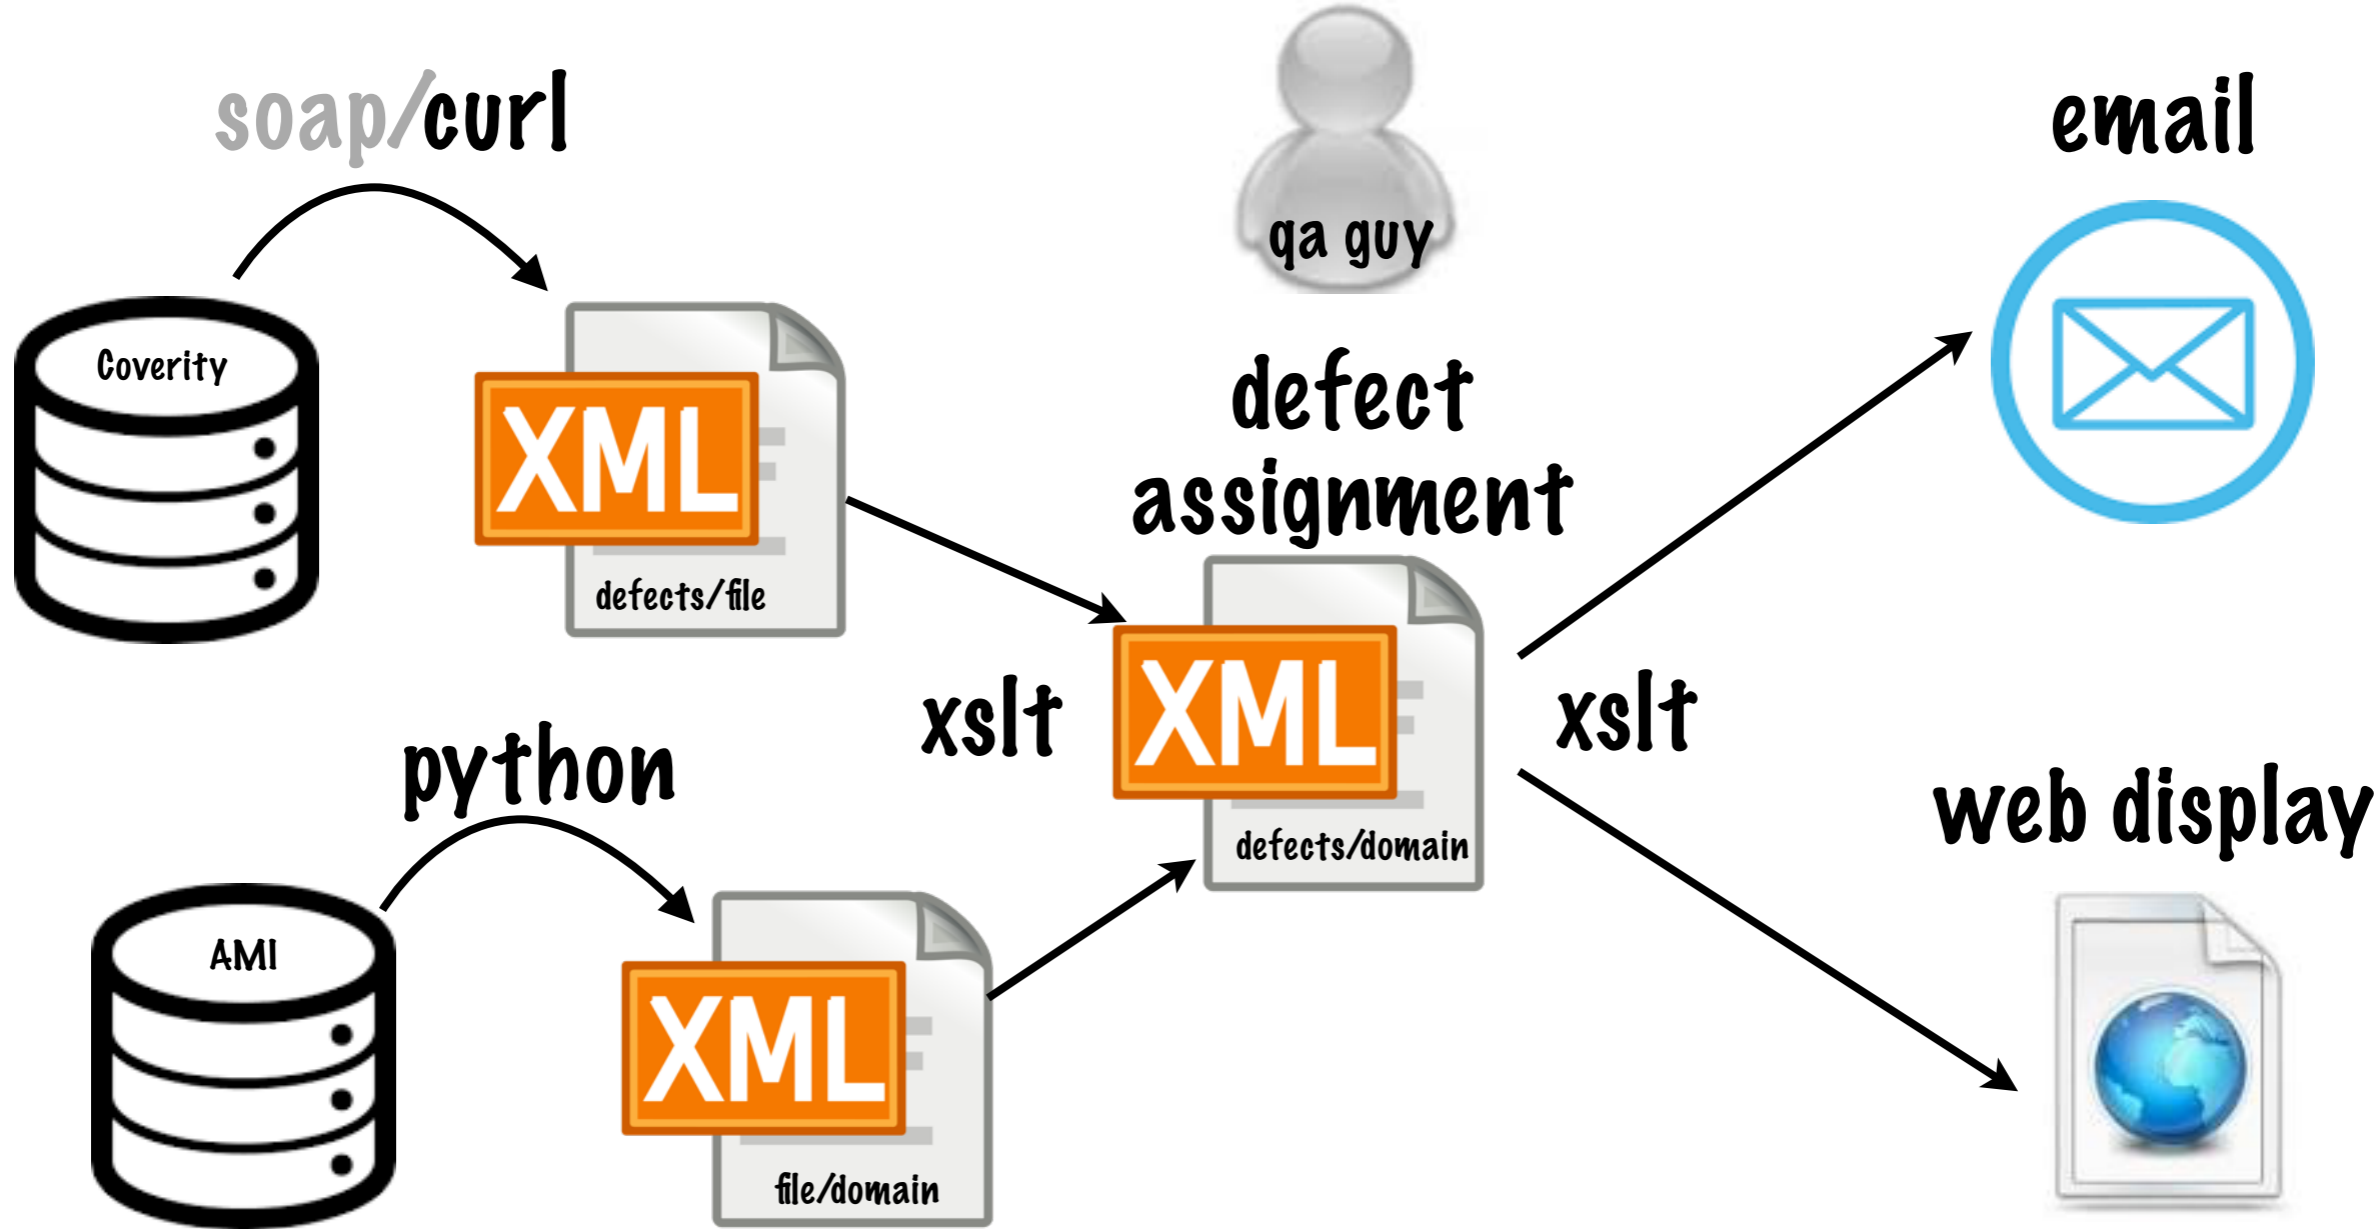
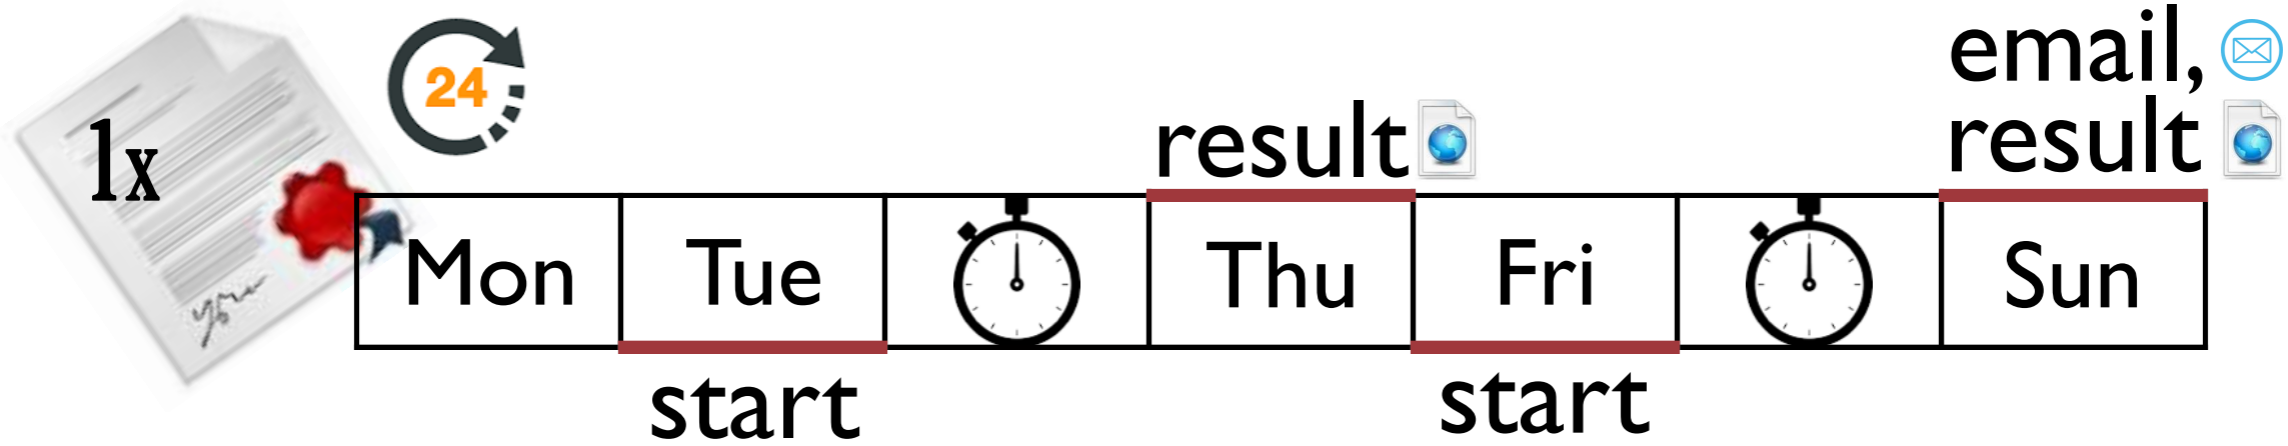
- 'new' without delete
- scalar 'delete' on array object
- array indices out-of-bounds
- possible cut-and-paste errors
- suspect indentation

```
if (m_minNUsedHitsdEdx > 0) {  
    original: this->m_minNUsedHitsdEdx looks like the original copy.  
    ATH_MSG_INFO( " Minimum used hits for dEdx: " << m_minNUsedHitsdEdx );  
    auto dEdx = std::make_shared<MinUsedHitsdEdxCut>(this, m_minNUsedHitsdEdx);  
    ATH_CHECK( dEdx->initialize() );  
    m_trackCuts["dEdxHits"].push_back(dEdx);  
}  
if (m_minNOverflowHitsdEdx > 0) {  
    CID 28963 (#1 of 1): Copy-paste error (COPY_PASTE_ERROR)  
    copy_paste_error: m_minNUsedHitsdEdx in this->m_minNUsedHitsdEdx looks like a copy-paste error.  
    Should it say m_minNOverflowHitsdEdx instead?  
    ATH_MSG_INFO( " Minimum IBL overflow hits for dEdx: " << m_minNUsedHitsdEdx );  
}
```

```
64 if (m_fullMaterial)  
65     sl << " - fullMaterial : " << *m_fullMaterial << std::endl;  
66     sl << " - split factor : " << m_splitFactor << std::endl;  
67 return sl;  
  
parent: This 'if' statement is the parent, indented to column 3.  
nephew: This statement is nested within its parent, indented to column 8.  
CID 11566 (#1 of 1): Nesting level does not match indentation (NESTING_INDENT_MISMATCH)  
uncle: This statement is indented to column 8, as if it were nested within the preceding parent statement, but it is not.
```

The examples are drawn from Coverity<sup>®</sup>, one of the most widely used static analysers.

# Coverity®





- 
- 

**12598 09/07/2014 (High) Resource leak : /ForwardDetectors/ALFA/ALFA\_CLinkAlg/src/ALFA\_CLinkAlg.cxx in function "execute"**



AtlasEvent Configuration Help Shaun Roe Enter CID(s)

12598	Resource leak	High	New	1	09/07/2014	Unassigned	Unclassified
-------	---------------	------	-----	---	------------	------------	--------------

All 1 issue selected Page 1 of 1

/ForwardDetectors/ALFA/ALFA\_CLinkAlg/src/ALFA\_CLinkAlg.cxx

```
51 }
52
53 StatusCode ALFA_CLinkAlg::execute()
54 {
55     1. Condition !!this->msgLvl(MSG::DEBUG) , taking false branch
56     2. Condition !!this->msgLvl(MSG::DEBUG) , taking false branch
57     ATH_MSG_DEBUG ("ALFA_CLinkAlg::execute()");
58     3. alloc_fn: Storage is returned from allocation function operator new .
59     4. var_assign: Assigning: pDataEvent = storage returned from new ALFA_CLinkEvent .
60     ALFA_CLinkEvent* pDataEvent=new ALFA_CLinkEvent ();
61     5. noescape: Resource pDataEvent is not freed or pointed-to in LoadAllEventData . [show details]
62     6. Condition !sc_.isSuccess() , taking true branch
63     CID 12598 (#1 of 1): Resource leak (RESOURCE_LEAK)
64     7. leaked_storage: Variable pDataEvent going out of scope leaks the storage it points to.
65     CHECK (LoadAllEventData (pDataEvent));
66     if (m_nDataType==1) pDataEvent->SetDCSFolderIDs(&m_CurrentDCSId);
67     CHECK (evtStore()->record(pDataEvent, "ALFA_CLinkEvent"));
```

**12598 Resource leak**

The system resource will not be reclaimed and reused, reducing the future availability of the resource.

In ALFA\_CLinkAlg::execute(): Leak of memory or pointers to system resources (CWE-404)

▼ Triage

Classification:

Severity:

Action:

Ext. Reference:

Owner:

Enter comments (See the History section below for previous comments)


► Projects and Streams

► Detection History

► Triage History

▼ Occurrences

1: AtlasEvent

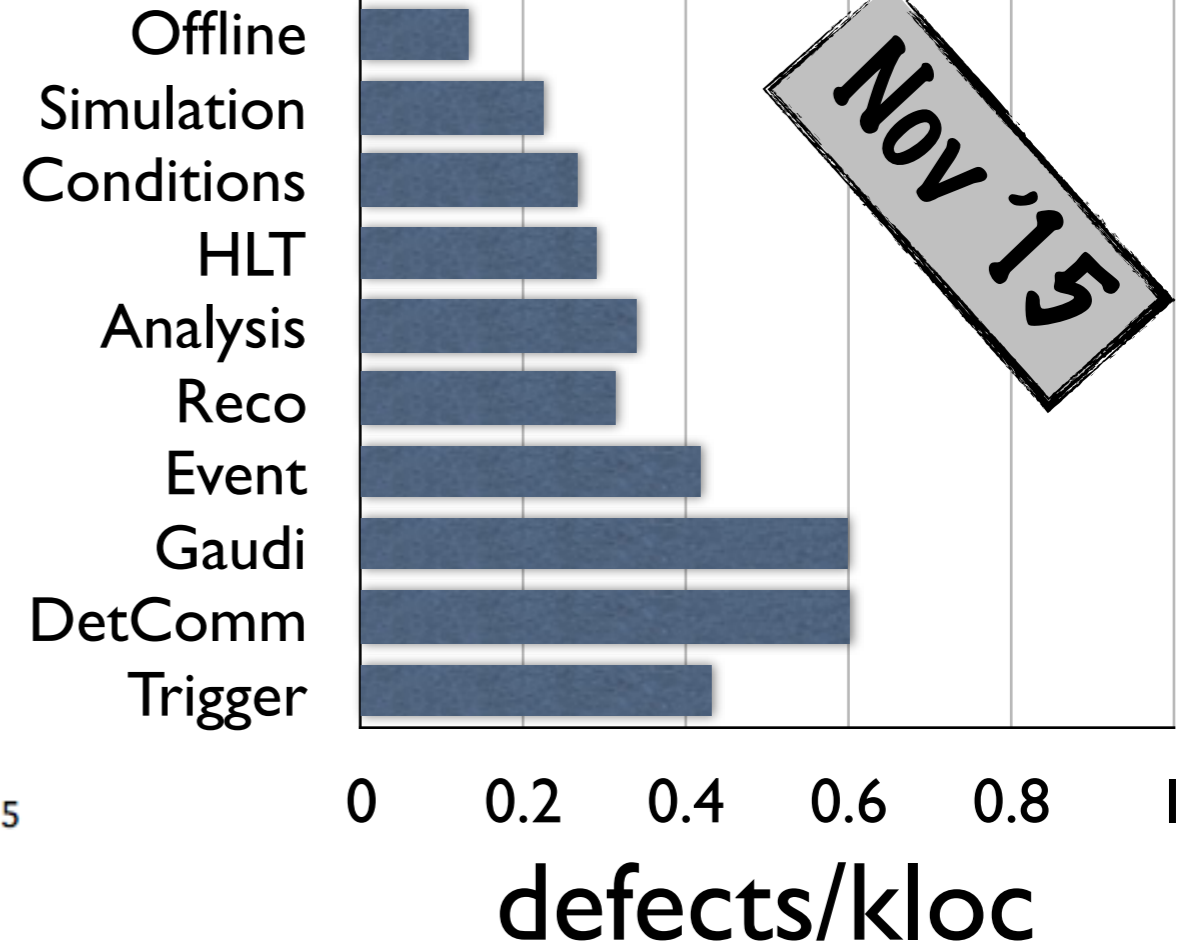
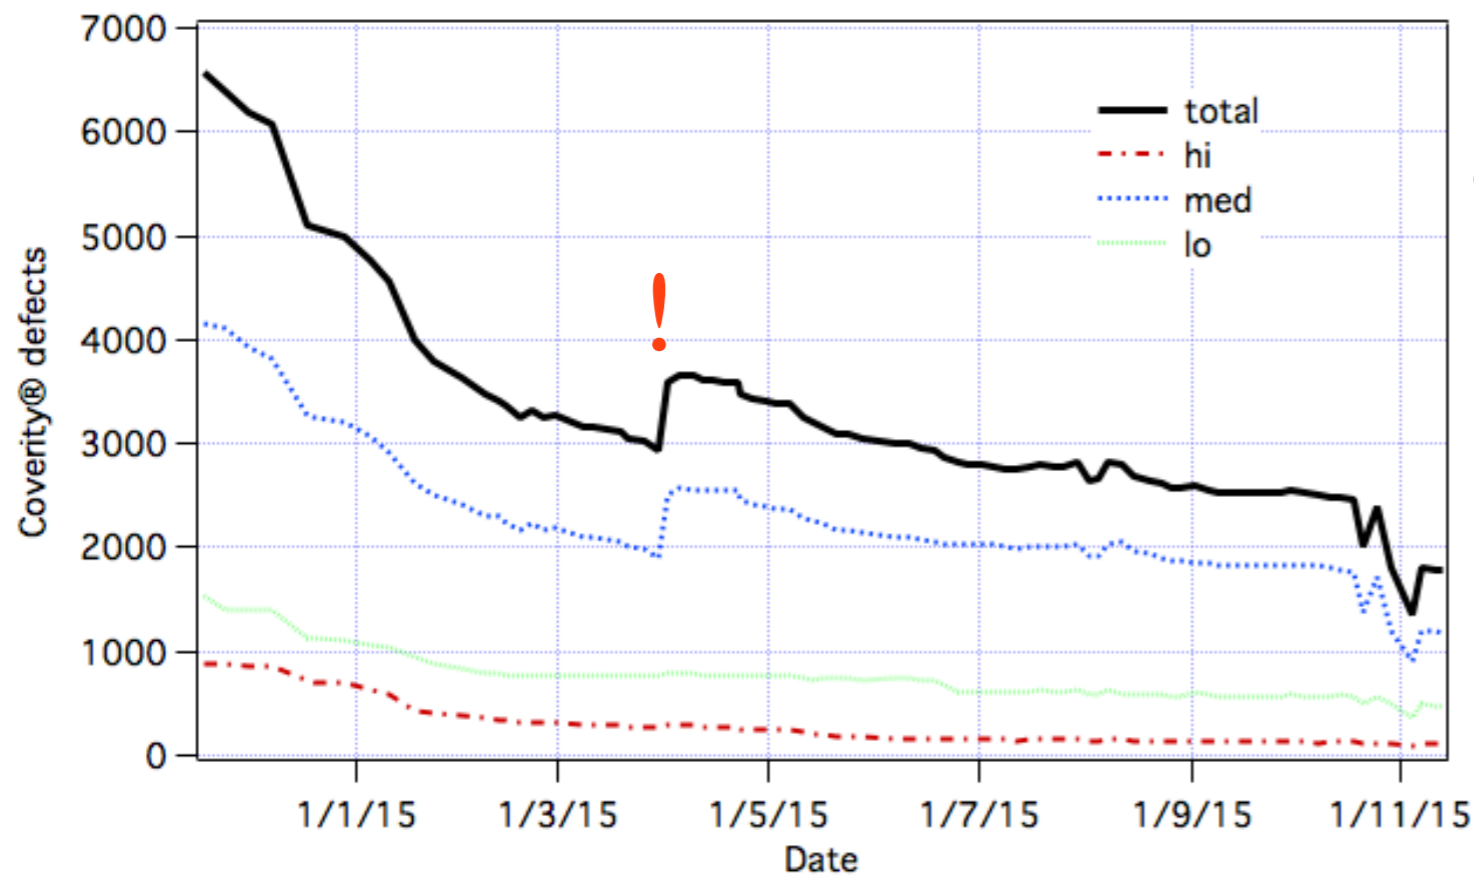


A command line tool, 'issues' will query the web page and list issues for a specified file

Pitfalls: Coverity's classification may not correspond to your own; e.g. **Uninitialised members** are classified 'high', as are obvious **resource leaks**; however a **faulty assignment operator** which can also leak resources will be 'low'. **Parse errors** can occur and are classified 'low' but may mask many other defects.



# Coverity® progress



## Industry report 2012

“Coverity’s analysis found an average defect density of .69 for open source software projects that leverage the Coverity Scan service, and an average defect density of .68 for proprietary code developed by Coverity enterprise customers. Both have better quality as compared to the **accepted industry standard defect density for good quality software of 1.0** [defects/kloc].”

# Coverity<sup>®</sup> pleasures and pitfalls



“spending hours with a memory profiler could save you 5 minutes looking at the coverity report” 😊

False positives (very few): 😞  
e.g. “restoring ostream format”

“parse errors” (particularly complicated templates)

Near misses:

```
792 //set the foreign key
6. return_constant: Function call variableType() returns 4.
CID 11595 (#2-1 of 2): Out-of-bounds read (OVERRUN)
7. overrun-local: Overrunning array of 4 bytes at byte offset 4 by dereferencing pointer &"_FK"[variableType()].
793 pixel_columns.createForeignKey(variableType<T>()+"_FK", "FK", m_pixeltable, "FK");
794 // create indices
```

Misses: `if (m_Analogue[1][strip]<0.0000001 || m_Analogue[1][strip]>-0.0000001){ //`

# Other Static Analysers

Installed, to be integrated into reporting system:

**cppcheck** (running routinely since Nov '15)

Open source, easily configurable, quick (~1 hr)  
somewhat noisy, more false positives, but *does find additional defects*. Output to XML/static web page. Options for performance (e.g. “use pre-increment”), style.

**Include-what-you-use**

Easy to set up; scope is limited to tidying up #includes; web interface available for reports (R. Seuster).

**OCLint**

Open source, based on Clang static analysis tools; ~12 hr scan with integrated result-to-web generation.

**Investigated:**

**PRQA:** Nice system, flexible, industry standards, can enforce naming conventions; \$\$, difficult to integrate in our build system

## Cppcheck report - AtlasReconstruction:

Defect summary;	Line	Id	Severity	Message
440 total		missingInclude	information	Cppcheck cannot find all the include files (use --check-config for details)
140 memleak	<a href="#">399</a>	coutCerrMisusage	error	Invalid usage of output stream: '<< std::cout'.
62 literalWithCharPtrCompare				
51 uninitMemberVar	<a href="#">258</a>	nullPointer	warning	Possible null pointer dereference: vot - otherwise it is redundant to check it against null.
42 nullPointer				
33 uninitStructMember	<a href="#">161</a>	mismatchAllocDealloc	error	Mismatching allocation and deallocation: disto
25 invalidscanf_libc				
17 deallocuse	<a href="#">586</a>	mismatchAllocDealloc	error	Mismatching allocation and deallocation: values
11 oppositeInnerCondition				
10 derefInvalidIterator	<a href="#">75</a>	uninitMemberVar	warning	Member variable 'Calibrator::resHist' is not initialized in the constructor.
9 eraseDereference				
8 zerodiv	<a href="#">493</a>	memleak	error	Memory leak: hist
7 mismatchAllocDealloc	<a href="#">650</a>	invalidscanf_libc	portability	scanf without field width limits can crash with huge input data on some versions of libc.
4 doubleFree	<a href="#">706</a>	invalidscanf_libc	portability	scanf without field width limits can crash with huge input data on some versions of libc.
4 invalidPrintfArgType_sint	<a href="#">707</a>	invalidscanf_libc	portability	scanf without field width limits can crash with huge input data on some versions of libc.
3 bufferAccessOutOfBounds	<a href="#">708</a>	invalidscanf_libc	portability	scanf without field width limits can crash with huge input data on some versions of libc.
2 StlMissingComparison	<a href="#">898</a>	uninitStructMember	error	Uninitialized struct member: startdata.det
2 compareBoolExpressionWithInt				
2 uninitvar				
1 arrayIndexOutOfBoundsCond				

..linked to annotated code

```

297     this->cd(1);
298     Moduletuple->Draw("y:x>>reshist10(40,-1200,1200,40,-1200,1200)",selectionA,"colz");
299     TH2F *reshist10 = (TH2F*)gPad->GetPrimitive("reshist10");
300     if (!reshist10) throw string("Variable not found!");
301     if (variable == "res"){ <--- String literal compared with variable 'variable'. Did you intend to use strcmp() instead?
302         reshist10->GetZaxis()->SetRangeUser(0.135, 0.150);
303     }
  
```

# Include-what-you-use

Choose first the project, then follow path to your favourite package

The screenshot shows the IWYU web interface. At the top, a blue box contains the text "Choose first the project, then follow path to your favourite package". Below this, there are several dropdown menus for project selection: "AtlasReconstruction", "Reconstruction", "MuonIdentification", and "[ remaining ]". The "Reconstruction" dropdown is currently selected. Below the dropdowns, there are two columns: "remove includes" (highlighted in orange) and "add includes" (highlighted in green). The "add includes" column contains a list of include statements with comments:

```
#include "AthenaBaseComps/AthCheckMacros.h" // for ATH_CHECK
#include "AthenaBaseComps/AthMsgStreamMacros.h" // for ATH_MSG_DEBUG, etc
#include "AthenaKernel/errorcheck.h" // for CHECK
#include "GaudiKernel/GaudiHandle.h" // for GaudiHandle
#include "GaudiKernel/IAlgorithm.h" // for IAlgorithm
#include "GaudiKernel/IProperty.h" // for IProperty
#include "GaudiKernel/IStateful.h" // for IStateful
#include "GaudiKernel/MsgStream.h" // for MsgStream, operator<<
```

...advises on which 'includes' to add, remove or can be replaced by a forward class declaration.

# Run-time Sanitizers

Implemented in nightly tests of DBG builds:

- **UBSan** (undefined behaviour)

Clear warnings; easy to implement :

```
HepMcParticleLink.h:72:51: runtime error: left shift of 200001 by 16 places cannot be represented in type 'int'
```

Investigated:

- **ASan** (address)

Similar to the debug tool Valgrind but much faster and catches more; output maybe a bit intimidating...

ABI incompatible.

```
==7552==ERROR: AddressSanitizer: heap-buffer-overflow on address 0x6140000ffd4 at pc 0x400825 bp 0x7fff248f7d80 sp 0x7fff248f7d78...
```

- **TSan** (thread)

Currently not useful 'out of the box' for ATLAS



"If all you have is a hammer,  
everything looks like a nail"  
- Psychology of Science (Maslow, 1966)

Coding to solve the line-by-line problems revealed by these tools *might* produce better code, but it won't forcibly produce good code.

“Code smell” detectors (tentative results so far):

- **TCToolkit**

- ~1 hr to scan release, produces >50Mb html.

- code duplication detector (looks useful)
- token tag cloud (“how noisy” the code is)
- class co-occurrence matrix (interdependencies)

# Continuous testing

‘ATN’ tests run nightly in all releases;

- Unit tests configured in a dedicated SVN test directory for the package
- XML file describes expected output; maybe simple return code, or comparison with a reference file.

```
1 <?xml version="1.0"?>
2 <atn>
3   <TEST name="ISF_EventTests" type="makecheck">
4     <package>Simulation/ISF/ISF_Core/ISF_Event</package>
5     <author> Elmar Ritsch </author>
6     <mailto> elmar.ritsch@cern.ch </mailto>
7     <expectations>
8       <errorMessage>Athena exited abnormally</errorMessage>
9       <errorMessage>differ</errorMessage>
10      <warningMessage> # WARNING_MESSAGE : post.sh> ERROR</warningMessage>
11      <successMessage>check ok</successMessage>
12      <returnValue>0</returnValue>
13    </expectations>
14  </TEST>
15 </atn>
```

‘RTT’ tests

- Produce plots of physics quantities (momenta, track parameters etc) and compared with reference plots
- Reviewed by shifters, who report each week.



# Summary

- All of these tools are useful as developer aids;
  - Coverity in particular reveals many programming errors and maintainability issues
- Public league tables help motivate groups.
  - None of these will replace peer review or make up for a lack of education (another important topic) or experience.
  - ATLAS will continue to use these tools, integrating the reports from different analyses, and keep an eye on new tools as they become