



How to Build Your Own Static Analyzer For Fun and Profit

Dr. Andy Chou
Chief Scientist
Coverity Inc.

© Coverity, Inc. 2004. All Rights Reserved. This publication, in whole or in part, may not be reproduced, stored in a computerized, or other retrieval system or transmitted in any form, or by any means whatsoever without the prior written permission of Coverity, Inc.



Game Plan

- Market research
- Gather requirements
- Build it
- Sell it
- Take over the world

2



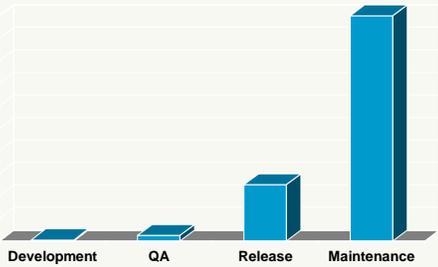
Is it worth doing?

- NIST Report: bad software costs \$59.5 billion/yr
- 0.5% of US GDP
- Ever stayed up all night to find a bug in your code?
- ... someone else's code?

3



Cost of Fixing a Defect

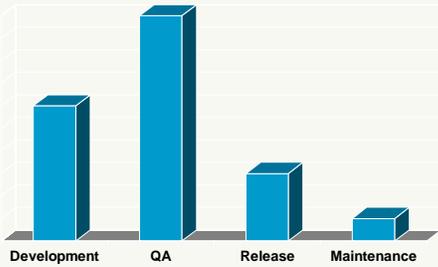


Stage	Relative Cost
Development	Low
QA	Low-Medium
Release	Medium-High
Maintenance	Very High

4



Number of Bugs Found



Stage	Number of Bugs Found
Development	Medium
QA	High
Release	Low-Medium
Maintenance	Low

5



Security vulnerabilities

Source: Robert Seacord, CERT



Reacting to vulnerabilities in existing systems is not working

Total vulnerabilities reported (1995-2004): 16,726

Year	Total Vulnerabilities Reported
97	311
98	262
99	417
2000	1090
01	2,437
02	4,129
03	3,784
04	3,780

6

Why Static Analysis?

- No test cases
- Does not require a “finished product”
- Explores all paths through the code
- Can check a wide variety of program properties
- Applies very early in development cycle

Game Plan

- Market research
- **Gather requirements**
- Build it
- Sell it
- Take over the world

Theoretical Properties

- Soundness
 - No false negatives
- Completeness
 - No false positives

Soundness & Completeness

Soundness & Completeness

Soundness & Completeness

- Neither is necessary nor sufficient to be useful for bug-finding
- Executive decision:
 - Live with some false positives and negatives
 - Bonus round if soundness or completeness can be achieved without sacrificing other requirements

Some Real-World Requirements 

- Parse the code
- Find a large number of good bugs
- Few false positives
- Scale to large code bases
- Explain bugs to the user

13

Two Cardinal Rules 

- The First Two Cardinal Rules of Static Analysis:
 - (1) You can't analyze what you can't find.
 - (2) You can't analyze what you can't parse.
- Coverity's first year was a crash course in learning these lessons

14

Parsing Code 

- Myth: The C language exists
- Bigger Myth: The C++ language exists
- What actually exists:
 - Compiler = a specific version of a "native" compiler
 - f = file.c, plus headers
 - o = ordered list of command line options and response files
 - e = environment
 - c = compiler-specific configuration files

$Lang(Compiler) = \{ f \mid \exists o, e, c \text{ Accepts}(Compiler, f, o, e, c) \}$

15

Parsing Code 

- Need to accept a reasonable subset of:
 $Lang(Compiler_1) \cup Lang(Compiler_2) \cup \dots$
- Some of the compilers with significant market share:
 

16

Parsing Horrors 

- Banal. But take more time than you can believe:


```
void x;      typedef char int;      int foo(int a, int a);
unsigned x @ "TEXT";      int c = L";
unsigned x = Oxdead_beef;
End lines with "\r" rather than "\n"
#pragma asm      asm foo() {      // newline = end
  mov eax, eab      mov eax, eab;      __asm mov eax, eab
#pragma end_asm      }      // "]" = end
                    __asm [      mov eax, eab
                    ]      ]
```
- And, of course, asm:

17

Punting on Parsing 

- Use Edison Design Group (EDG) frontend
- Still need custom rewriter for many supported compilers:

205 hpux_compilers.c	453 sun_compilers.c
215 iar_compiler.c	485 arm_compilers.c
240 ti_compiler.c	617 gnu_compilers.c
251 green_hills_compiler.c	748 microsoft_compilers.c
377 intel_compilers.c	1587 metrowerks_compilers.c
453 diab_compilers.c	...

18

Some Real-World Requirements 

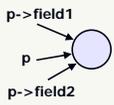
- Parse the code
- Find a large number of good bugs
- Few false positives
- Scale to large code bases
- Explain bugs to the user

19

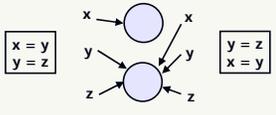
What About Alias Analysis? 

- Traditional static analysis for optimizations
- Scalable algorithms are very coarse

Field insensitivity



Flow insensitivity



20

Bugs to Detect 

- Some examples

<ul style="list-style-type: none"> • Crash Causing Defects • Null pointer dereference • Use after free • Double free • Array indexing errors • Mismatched array new/delete • Potential stack overrun • Potential heap overrun • Return pointers to local variables • Logically inconsistent code 	<ul style="list-style-type: none"> • Uninitialized variables • Invalid use of negative values • Passing large parameters by value • Underallocations of dynamic data • Memory leaks • File handle leaks • Network resource leaks • Unused values • Unhandled return codes • Use of invalid iterators
--	--

21

Open_args 

- Prototype for open() syscall:

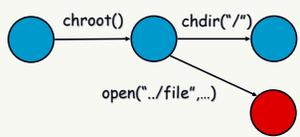

```
int open(const char *path, int oflag, /* mode_t mode */...);
```
- Typical mistake:


```
fd = open("file", O_CREAT);
```
- Result: file has random permissions
- Check: Look for oflags == O_CREAT without mode argument

22

Chroot checker 

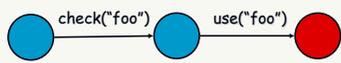
- chroot() changes filesystem root for a process
 - Confine process to a "jail" on the filesystem
- Doesn't change current working directory



23

TOCTOU 

- Race condition between time of check and use
- Not applicable to all programs



24

Tainting checkers

Tainted data accepted from source
 ↓
 Unvetted data taints other data transitively
 ↓
 Tainted data is used in an operator or function

Example Sinks:	system()	printf()	malloc()	strcpy()	Sent to HDDMS	Included in HTML
Resultant Vulnerability:	command injection	format string manip.	integer/buffer overflow	buffer overflow	SQL injection	cross site scripting

25

Demo

26

Example Code

```

#include <stdlib.h>
#include <stdio.h>

void say_hello(char * name, int size) {
    printf("Enter your name: ");
    fgets(name, size, stdin);
    printf("Hello %s.\n", name);
}

int main(int argc, char *argv[]) {
    if (argc != 2) {
        printf("Error, must provide an input buffer size.\n");
        exit(-1);
    }
    int size = atoi(argv[1]);
    char * name = (char*)malloc(size);
    if (name) {
        say_hello(name, size);
        free(name);
    } else {
        printf("Failed to allocate %d bytes.\n", size);
    }
}

```

27

Callgraph

```

graph TD
    main((main)) --> atoi((atoi))
    main --> exit((exit))
    main --> free((free))
    main --> malloc((malloc))
    main --> say_hello((say_hello))
    say_hello --> fgets((fgets))
    say_hello --> printf((printf))
    main --> main

```

28

Reverse Topological Sort

```

graph TD
    main((main)) --> atoi((atoi))
    main --> exit((exit))
    main --> free((free))
    main --> malloc((malloc))
    main --> say_hello((say_hello))
    say_hello --> fgets((fgets))
    say_hello --> printf((printf))
    main --> main

```

29

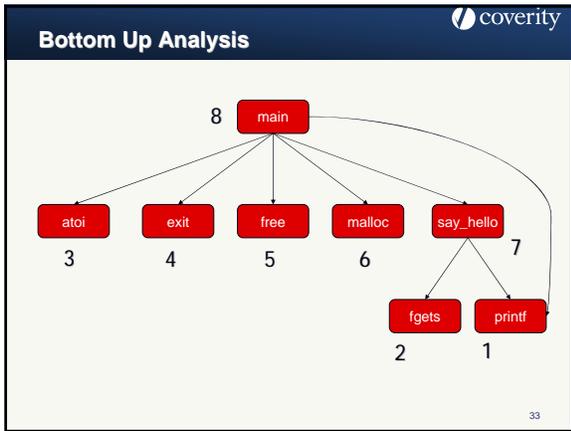
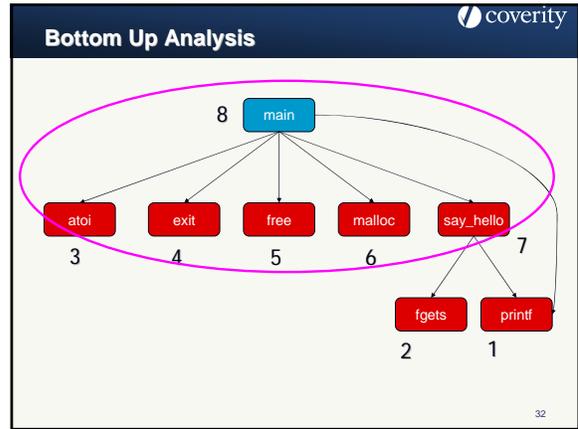
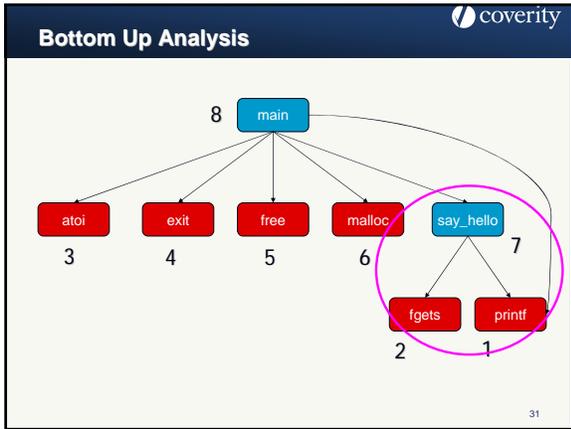
Apply Library Models

```

graph TD
    main((main)) --> atoi((atoi))
    main --> exit((exit))
    main --> free((free))
    main --> malloc((malloc))
    main --> say_hello((say_hello))
    say_hello --> fgets((fgets))
    say_hello --> printf((printf))
    main --> main

```

30

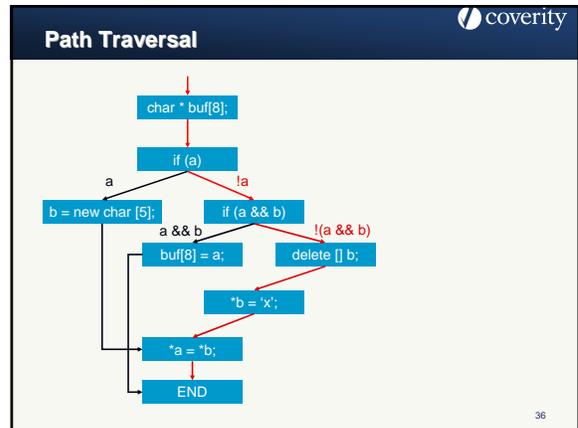
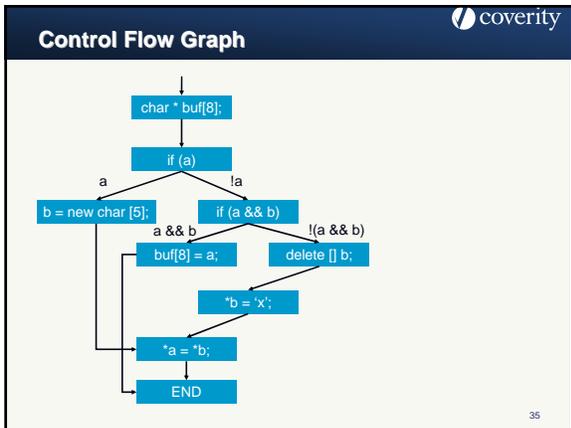


Finding Local Bugs

```

#define SIZE 8
void set_a_b(char * a, char * b) {
    char * buf[SIZE];
    if (a) {
        b = new char[5];
    } else {
        if (a && b) {
            buf[SIZE] = a;
            return;
        } else {
            delete [] b;
        }
        *b = 'x';
    }
    *a = *b;
}

```



Apply Checking coverity

Null pointers Use after free Array overrun

```

char * buf[8];
if (a)
  !a
  if (a && b)
    !(a && b)
    delete [] b;
    *b = 'x';
    *a = *b;
END

```

37

Apply Checking coverity

Null pointers Use after free Array overrun

"buf is 8 bytes"

```

char * buf[8];
if (a)
  !a
  if (a && b)
    !(a && b)
    delete [] b;
    *b = 'x';
    *a = *b;
END

```

38

Apply Checking coverity

Null pointers Use after free Array overrun

"buf is 8 bytes"

```

char * buf[8];
if (a)
  !a
  if (a && b)
    !(a && b)
    delete [] b;
    *b = 'x';
    *a = *b;
END

```

"a is null"

39

Apply Checking coverity

Null pointers Use after free Array overrun

"buf is 8 bytes"

```

char * buf[8];
if (a)
  !a
  if (a && b)
    !(a && b)
    delete [] b;
    *b = 'x';
    *a = *b;
END

```

"a is null"

Already knew a was null

40

Apply Checking coverity

Null pointers Use after free Array overrun

"buf is 8 bytes"

```

char * buf[8];
if (a)
  !a
  if (a && b)
    !(a && b)
    delete [] b;
    *b = 'x';
    *a = *b;
END

```

"a is null"

"b is deleted"

41

Apply Checking coverity

Null pointers Use after free Array overrun

"buf is 8 bytes"

```

char * buf[8];
if (a)
  !a
  if (a && b)
    !(a && b)
    delete [] b;
    *b = 'x';
    *a = *b;
END

```

"a is null"

"b is deleted"

"b dereferenced!"

42

Apply Checking

Null pointers Use after free Array overrun
"buf is 8 bytes"

```

char * buf[8];
if (a)
  !a
  if (a && b)
    !(a && b)
    delete [] b;
    *b = 'x';
    *a = *b;
END

```

"a is null"

"b is deleted"

"b dereferenced!"

No more errors reported for b

43

Some Real-World Requirements

- Parse the code
- Find a large number of good bugs
- **Few false positives**
- Scale to large code bases
- Explain bugs to the user

44

False Positives

- What is a bug? Something the user will fix.
- Many sources of false positives
 - False paths
 - Idioms
 - Execution environment assumptions
 - Killpaths
 - Conditional compilation
 - "third party code"
 - Analysis imprecision
 - ...

$p = 0; *p;$
`fatal("ouch");`

45

A False Path

```

char * buf[8];
if (a)
  a
  b = new char [5];
  if (a && b)
    a && b
    buf[8] = a;
    delete [] b;
    *b = 'x';
  !!(a && b)
  *a = *b;
END

```

46

False Path Pruning

Integer Range Disequality Branch

```

char * buf[8];
if (a)
  !a
  if (a && b)
    a && b
    buf[8] = a;
END

```

47

False Path Pruning

Integer Range Disequality Branch

```

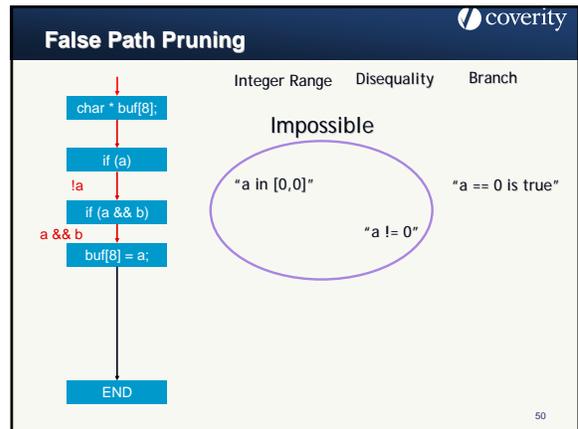
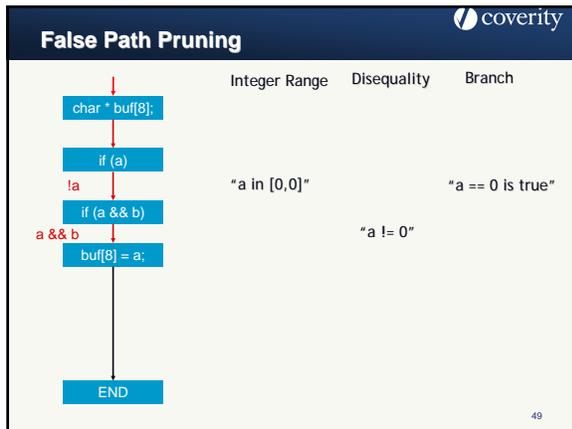
char * buf[8];
if (a)
  !a
  if (a && b)
    a && b
    buf[8] = a;
END

```

"a in [0,0]"

"a == 0 is true"

48



Environment Assumptions

- Should the return value of malloc() be checked?

```
int *p = malloc(sizeof(int));
*p = 42;
```

OS Kernel: Crash machine.	File server: Pause filesystem.	Web application: 200ms downtime
Spreadsheet: Lose unsaved changes.	Game: Annoy user.	IP Phone: Annoy user.
Library: ?	Medical device: malloc?!	

51

Statistical Analysis

- Assume the code is usually right

3/4 deref	}	int *p = malloc(sizeof(int)); *p = 42;	int *p = malloc(sizeof(int)); if(p) *p = 42;	1/4 deref
		int *p = malloc(sizeof(int)); *p = 42;	int *p = malloc(sizeof(int)); if(p) *p = 42;	
		int *p = malloc(sizeof(int)); *p = 42;	int *p = malloc(sizeof(int)); if(p) *p = 42;	
		int *p = malloc(sizeof(int)); if(p) *p = 42;	int *p = malloc(sizeof(int)); *p = 42;	

52

- ### Some Real-World Requirements
- Parse the code
 - Find a large number of good bugs
 - Few false positives
 - Scale to large code bases
 - Explain bugs to the user
- 53

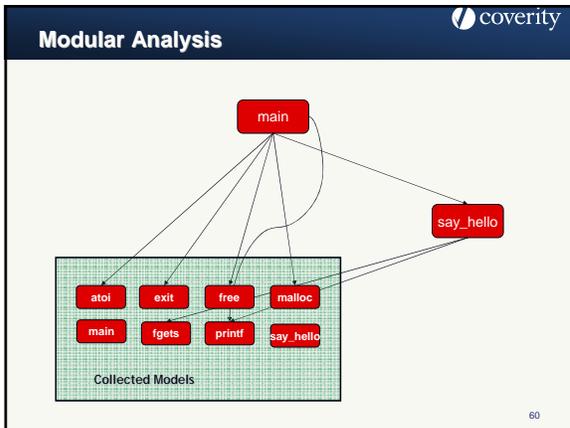
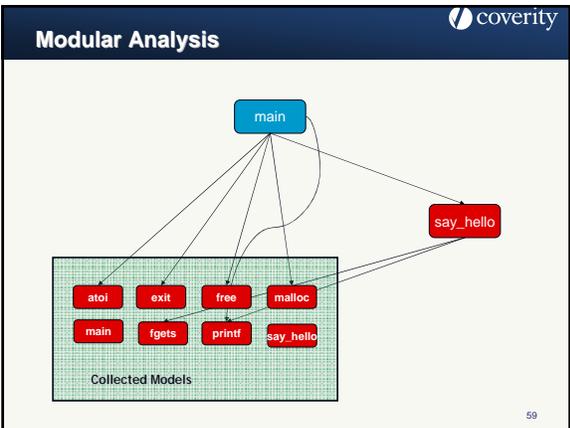
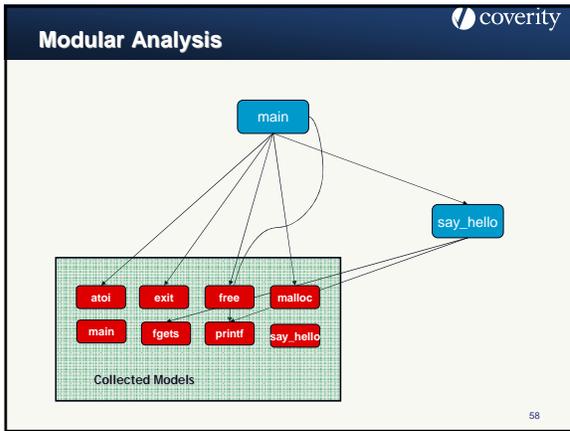
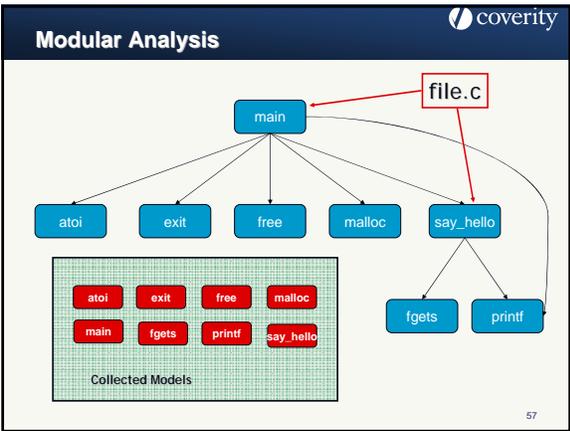
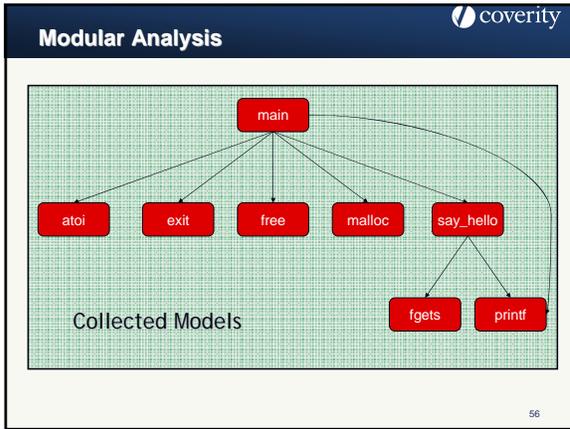
- ### How Large is "Large"?
- Various definitions of what a LOC is
 - Some numbers (C/C++ code only)
 - Linux kernel + drivers: ~3.2 MLOC
 - Firefox browser: ~1.9 MLOC
 - Gcc: 0.75 MLOC
 - OpenSSL: 0.21 MLOC
 - XMMS: 0.12 MLOC
 - Commercial code bases: 1 MLOC+ is common
 - Large sites: 30 MLOC+
- 54

Scalability 

- Usual restriction: nightly build cycle
- Max: ~12 hours (~700 lines/sec)
- Solutions
 - Parallel analysis
 - Incremental analysis
 - Modular analysis

} (hard)

55



Some Real-World Requirements

- Parse the code
- Find a large number of good bugs
- Few false positives
- Scale to large code bases
- Explain bugs to the user

61

Example from Linux

Example from Linux 2.6.10 (drivers/scsi/aic7xxx_old.c)

```

Added 'host' due to comparison 'host != 0'
Also see events: [var_deref_model]
At conditional (1) 'host != 0' taking false path

0427  if (host != NULL)
0428  {
0429      p = (struct aic7xxx_host *) host->boarddata;
0430      memset(p, 0, sizeof(struct aic7xxx_host));
0431      p = temp;
0432      p->host = host;
0433
0434      p->sub_data = malloc(sizeof(sub_data_type), GFP_ATOMIC);
0435      if (!p->sub_data) return NULL;
0436
0437      memset(p->sub_data, 0, sizeof(sub_data_type));
0438      mutex_lock(&p->sub_data->mutex);
0439      }
0440      else
0441      {
0442          /*
0443           * For some reason we don't have enough memory. Free the
0444           * allocated memory for the aic7xxx_host struct, and return NULL.
0445           */
0446          cleanup_region(p->base, MAXREG - MINREG);
0447          return cleanup(host);
0448          return NULL;
0449      }
0450      p->host_no = host->host_no;
0451      }

```

Variable 'host' tracked as NULL was passed to a function that dereferences it [model]
Also see events: [var_compare_op]

```

0452  scsi_set_device(host, &p->pdev->dev);
0453  return (p);
0454  }
0455

```

62

Explaining Bugs

```

Added 'host' due to comparison 'host != 0'
Also see events: [var_deref_model]
At conditional (1) 'host != 0' taking false path

0427  if (host != NULL)
0428  {

```

- The red messages are *events*, which are the places in the code where the analysis makes a decision
- The green messages indicate the path the analysis followed
- In this case, the analysis indicates that when following the false path from the comparison on line 0427, the variable `host` is NULL

63

Explaining Interprocedural Bugs

```

Variable 'host' tracked as NULL was passed to a function that dereferences it [model]
Also see events: [var_compare_op]

0452  scsi_set_device(host, &p->pdev->dev);
0453  return (p);
0454  }
0455

```

- When the analysis reports a message that crosses procedure boundaries, a "model" link appears
 - Following the model link will show the event in the callee function
- Each variable, field, function call is linked to its declaration through the cross referencing links

64

Explaining Interprocedural Bugs

```

542  static inline void scsi_set_device(struct Scsi_Host *shost,
543  struct device *dev)
544  {
545      shost->shost_gendev.pakent = dev;
546  }

```

Potentially NULL pointer

Directly dereferenced parameter 'shost'

65

Game Plan

- Market research
- Gather requirements
- Build it
- Sell it
- Take over the world

66

Coverity's Commercial History

Breakthrough technology out of Stanford	Company incorporated	Achieved profitability	Product growth and proliferation
2000-2002 • Meta-level compilation checker ("Stanford Checker") detects 2000+ bugs in Linux. 	2002 • Deluge of requests from companies wanting access to the new technology. • First customer signs: Sanera systems	2003 • 7 early adopter customers, including VMWare, SUN, Handspring. • Coverity achieves profitability.	2004-05 • Version 2.0 product released. • Company quadruples • 110+ customers including Juniper, Synopsys, Veritas, nVidia, palmOne. • Self funded

67

VBE (Virtual Build Environment)

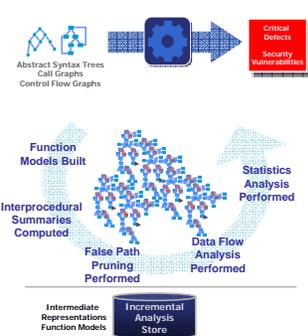


No changes required to your build environment

- Utilizes automated configuration
- Leverages existing build system
- Builds detailed and accurate representation of whole code:
 - Abstract syntax trees
 - Call graphs
 - Control flow graphs

Supported Compilers: Windows, Solaris, GCC, Green Hills, HP-UX, NetBSD, FreeBSD, ARM, Intel, Texas Instruments, WIND RIVER, ANSI.

Analysis Engine



Deep and comprehensive Analysis and Detection

- Sophisticated dataflow analysis:
 - Cross-functional
 - Context aware
 - Value tracking
- 100% of all possible paths
- Deep inspection:
 - Uncovers more defects
 - Better accuracy
 - 20% false positive rate
- Scalable to tens of millions of lines of code
- Unparalleled performance:
 - Subsequent runs
 - No loss of accuracy
 - Full analysis on desktop

Function Models Built, Interprocedural Summaries Computed, False Path Pruning Performed, Data Flow Analysis Performed, Statistics Analysis Performed, Incremental Analysis Store.

Prevent Library of Checkers

Quality, Concurrency, Security

- System and Process Crashes**
 - Memory/Resource Leaks
 - Data, Memory, File Corruption
 - Performance Degradation
 - Unpredictable Behavior
- Deadlocks**
 - Lock Contention
 - Unpredictable performance
 - Performance degradation
- Denial of Service**
 - Privilege Escalation
 - Malicious Code

Checker Library

Resource Problems <ul style="list-style-type: none"> Resource Leak <ul style="list-style-type: none"> Memory leak File pointer leak System resource leak Double Lock Missing unlock Incorrect lock acquisition Sleeping while locked 	Pointer Errors <ul style="list-style-type: none"> Use of uninitialized data Uninitialized memory Uninitialized variable Mismatched allocation operators Dereferencing invalid pointers Null pointer dereference Wrong address space Accessing freed pointers Dangling stack references Use of freed resource Double Free (memory, file pointers, system resources) Use after free (memory, file pointers, system resources) 	Logic Errors <ul style="list-style-type: none"> Flawed branch logic Use of invalid STL iterators Unusual operators Inconsistent error handling Security logic errors Time of check, time of use Insecure file creation Improper chroot Improper privilege inheritance Security Warnings Potentially insecure coding practices 	Bounds Errors <ul style="list-style-type: none"> Out of bounds array access Buffer underflow Stack smashing Stack overflow Stack buffer overrun Stack string overrun Bad negative integer cast Incorrect allocation size Non-null terminated strings
Concurrency Problems <ul style="list-style-type: none"> Double Lock Missing unlock Incorrect lock acquisition Sleeping while locked 	External Data Handling <ul style="list-style-type: none"> Integer Loop bound Array access Allocation size Strings Buffer overflow SQL injection Format string errors Cross-site scripting 	API Usage Errors <ul style="list-style-type: none"> Passing large parameters Inspecting temp file creation Improper method override 	

Game Plan

- Market research
- Gather requirements
- Build it
- Sell it**
- Take over the world

71

Sales Model

- Offer prospects a free trial
- Run over tons of code, find tons of bugs
- Have a meeting showing bugs to developers and their managers
- Delete most of the results and go home
- Wait for them to call back

72

Academics don't understand money

- "We'll just charge per seat like everyone else"
 - Finish the story: "Company X buys three Purify seats, one for Asia, one for Europe and one for the US..."
- Try #2: "we'll charge per lines of code"
 - "That is a really stupid idea: (1) ..., (2) ... , ... (n) ..."
 - Actually works. I'm still in shock. Would recommend it.
- Good feature for seller:
 - No seat games. Revenue grows with code size. Run on another code base = new sale.
- Good feature for buyer: No seat-model problems
 - Buy once for project, then done. No per-seat or per-usage cost; no node lock problems; no problems adding, removing or renaming developers (or machines)
 - People actually seem to like this pitch.

73

Some experience

- Surprise: Sales guys are great
 - Easy to evaluate. Modular.
- Company X buys tool, then downsizes. Good or bad?
 - For sales, very good: X fires 110 people. They get jobs elsewhere. Recommend coverity. 4 closed deals.
- Large companies "want" to be honest
 - Veritas: want monitoring so don't accidentally violate!
- What can you sell?
 - User not same as tool builder. Naive. Inattentive. Cruel.
 - Makes it difficult to deploy anything sophisticated.
 - Example: statistical inference.
 - Some ways, checkers lag much behind our research ones.

74

"No, your tool is broken: that's not a bug"

- "No, the loop will go through once!"


```
for(i=1; i < 0; i++) {
  ...deadcode...
}
```
- "No, && is 'or'!"


```
for(s=0; s < n; s++) {
  ...
  switch(s) {
    case 0: assert(0); return;
    ...
  }
  ...dead code...
}
```
- "No, ANSI lets you write 1 past end of the array!"
 - ("We'll have to agree to disagree." !!!!!)

```
unsigned p[4]; p[4] = 1;
```

75

Game Plan

- Market research
- Gather requirements
- Build it
- Sell it
- **Take over the world**

76

A Partial List of 110+ Customers

EDA 	Storage 	Security
Networking 	Government 	Embedded
Biz Applications 	OS 	Open Source

77

Open Source Quality Project

Vulnerability Discovery and Remediation Open Source Hardening Project

Purpose

- Research latest source analysis techniques for security and quality
- Establish new baseline for software quality and security in open source
- Establish a metric for software quality

78

Slide 73

DE1 whole bunch of options: razor blade model, where we give away checkers for free and charge for system. or inverse razor where we give away system and charge for checkers. or charge per seat, or charge per lines of code (prefix). get a lot of pushback on the last one. prefix worked ok, but not what we would consider a success.

i argued very strongly against per line model. completely wrong.

Dawson Engler, 8/23/2005

Slide 74

DE2 count how often something true versus not. sort in decreasing deviance. inspect until hit fp. developer inspect all, mark as FP. say tool sucks to everyone.

Dawson Engler, 8/23/2005

Slide 75

DE3 we know about these since happen in results meetings. a bit dangerous, but usually other developers will laugh at the confused one. scary to think of other times when things just marked as FP.

Dawson Engler, 8/23/2005

Preliminary Findings

- Over 50 commonly used open source packages
- Over 32 MLOC analyzed nightly with 3 standard boxes
- More packages added regularly

- .434 defects/KLOC open source average
- .290 defects/KLOC LAMP stack

79

FreeBSD

"Coverity's Prevent is an invaluable tool that we've now been able to integrate into the FreeBSD Project development process with nightly source code scans.

Eighty-five Free BSD developers are now registered to review Coverity-generated bug reports, resulting in hundreds of important bug fixes, one leading to a security advisory.

Coverity's contributions have significantly improved the quality of Free BSD source codebase, which is greatly appreciated by both FreeBSD developers and users"

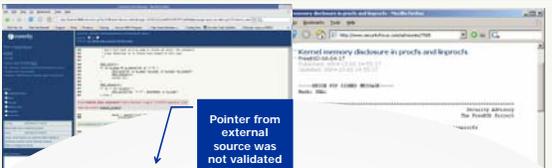
-- Robert Watson,
President FreeBSD Foundation



- 1,682,166 LOC
- 635 defects (.377 defects /KLOC)

80

Free BSD Security Alert



Pointer from external source was not validated before being dereferenced

Malicious local user could perform a local denial of service attack by causing a system panic. A user could read parts of kernel memory such as the file cache or terminal buffers which can include user-entered passwords.

X.Org Case study



- X Windows Systems (also known as X11)
- Originated at MIT in 1984
- Open source
- Primary GUI for Unix operating systems including Linux, Solaris and others
- Bundled in Open VMS and Mac OS X and many other operating systems environments

82

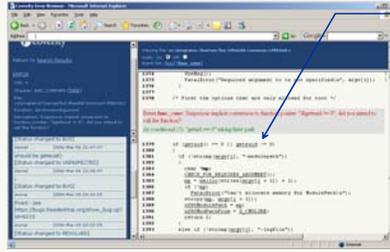
X.Org Case study



- Scan found 1681 potential defects in 2.3 MLOC
- Defect rate is .71 defects /KLOC
 - Above average
- First 11 days inspected 879 with an out-of-the-box false positive rate of **21%**
- Half way through fixing reported bugs

83

Example Security Violation (Bug #4)



Missing () to libc function geteuid detected by BAD_COMPARE checker

Since without the parentheses, the code is simply checking to see if the geteuid function in libc was loaded somewhere other than address 0 (which is pretty much guaranteed to be true), it was reporting it was safe to allow risky options for all users, and thus a security hole was born.

- Alan Coopersmith, Sun Developer

X.Org Privilege Escalation

```

/* First the options that are only allowed for root */
if (getuid() == 0 || geteuid() != 0)
...

```

85

X.Org Privilege Escalation

```

/* First the options that are only allowed for root */
if (getuid() == 0 || geteuid() != 0)
...

```

- The vulnerability allows ordinary users on a system to perform the following
 - Run arbitrary code as root administrator
 - Create and overwrite files as root administrator

86

Security Advisory

First exploit was published 5 hours after the bug was publicly reported on 3/20/2006

87

Ethernet 0.99.0

- Released 4/24/06

Under a grant funded by the U.S. Department of Homeland Security, Coverity has uncovered a number of vulnerabilities in Ethernet:

- The statistics counter could crash Ethernet. Versions affected: 0.10.10 - 0.10.14. CVE: [CVE-2006-1927](#)
- Ethernet could crash while reading a malformed Sniffer capture. Versions affected: 0.8.12 - 0.10.14. CVE: [CVE-2006-1928](#)
- An invalid display filter could crash Ethernet. Versions affected: 0.9.16 - 0.10.14. CVE: [CVE-2006-1929](#)
- The general packet dissector could crash Ethernet. Versions affected: 0.10.9 - 0.10.14. CVE: [CVE-2006-1927](#)
- The AMM dissector could crash Ethernet. Versions affected: 0.10.7 - 0.10.14. CVE: [CVE-2006-1927](#)
- The SPC dissector could crash Ethernet. Versions affected: 0.9.8 - 0.10.14. CVE: [CVE-2006-1929](#)
- The DCEBSP dissector could crash Ethernet. Versions affected: 0.9.16 - 0.10.14. CVE: [CVE-2006-1930](#)
- The ASN-1 dissector could crash Ethernet. Versions affected: 0.9.8 - 0.10.14. CVE: [CVE-2006-1929](#)
- The SMB PIPE dissector could crash Ethernet. Versions affected: 0.8.20 - 0.10.14. CVE: [CVE-2006-1938](#)
- The BER dissector could loop excessively. Versions affected: 0.10.4 - 0.10.14. CVE: [CVE-2006-1933](#)
- The SMDP dissector could abort. Versions affected: 0.10.4 - 0.10.14. CVE: [CVE-2006-1934](#)
- The Network Instruments file code could overrun a buffer. Versions affected: 0.10.0 - 0.10.14. CVE: [CVE-2006-1934](#)
- The NetScout/Windows Sniffer file code could overrun a buffer. Versions affected: 0.10.13 - 0.10.14. CVE: [CVE-2006-1934](#)
- The GSM GMM dissector could crash Ethernet. Versions affected: 0.9.16 - 0.10.14. CVE: [CVE-2006-1932](#)
- The ALCAP dissector could overrun a buffer. Versions affected: 0.10.14. CVE: [CVE-2006-1934](#)
- The labnet dissector could overrun a buffer. Versions affected: 0.8.6 - 0.10.14. CVE: [CVE-2006-1936](#)
- ASN-1 based dissectors could crash Ethernet. Versions affected: 0.9.10 - 0.10.14. CVE: [CVE-2006-1929](#)
- The H.264 dissector could crash Ethernet. Versions affected: 0.10.11 - 0.10.14. CVE: [CVE-2006-1932](#)
- The DCEBSP NT dissector could crash Ethernet. Versions affected: 0.9.14 - 0.10.14. CVE: [CVE-2006-1932](#)
- The PER dissector could crash Ethernet. Versions affected: 0.9.14 - 0.10.14. CVE: [CVE-2006-1930](#)

Under Windows, Unicode characters in profile and configuration file paths could cause problems. Versions affected: 0.10.14. The Coverity audit turned up several UI-related bugs that could make Ethernet crash.

88

More Experiences

- Security is dismal in the real world
- For many companies, source code analysis is still #4 or #5 on the priority list
 - Patching process (get customers to react to the problem)
 - Marketing (convince the customer there is no problem)
 - Training (teach customers how to turn off the parts with problems)
 - Tiger teams (guess what the problems might be)

"There are a thousand hacking at the branches of evil to one who is striking at the root"

- Henry David Thoreau, *Walden*, 1854

89

Conclusion

- Bugs are everywhere
- No bug is too stupid to look for
- Finding errors often easy, saying why is hard
- False positives kill usage, credibility, ROI
- Static analysis will become mainstream

90

Backup slides



91

Myth: More analysis is always better



- Does not always improve results, and can make worse
- The best error:
 - Easy to diagnose
 - True error
- More analysis used, the worse it is for both
 - More analysis = the harder error is to reason about, since user has to manually emulate each analysis step.
 - Number of steps increase, so does the chance that one went wrong.
No analysis = no mistake.
- In practice:
 - Demote errors based on how much analysis required
 - Revert to weaker analysis to cherry pick easy bugs
 - Give up on errors that are too hard to diagnose.

92

Myth: Soundness is a virtue



- Soundness: Find all bugs of type X.
 - Not a bad thing. More bugs good.
 - BUT: can only do if you check weak properties.
- What soundness really wants to be when it grows up:
 - Total correctness: Find all bugs.
 - Most direct approximation: find as many bugs as possible.
- Opportunity cost:
 - Diminishing returns: Initial analysis finds most bugs
 - Spend on what gets the next biggest set of bugs
 - Easy experiment: bug counts for sound vs unsound tools.
- End-to-end argument:
 - "It generally does not make much sense to reduce the residual error rate of one system component (property) much below that of the others."

93