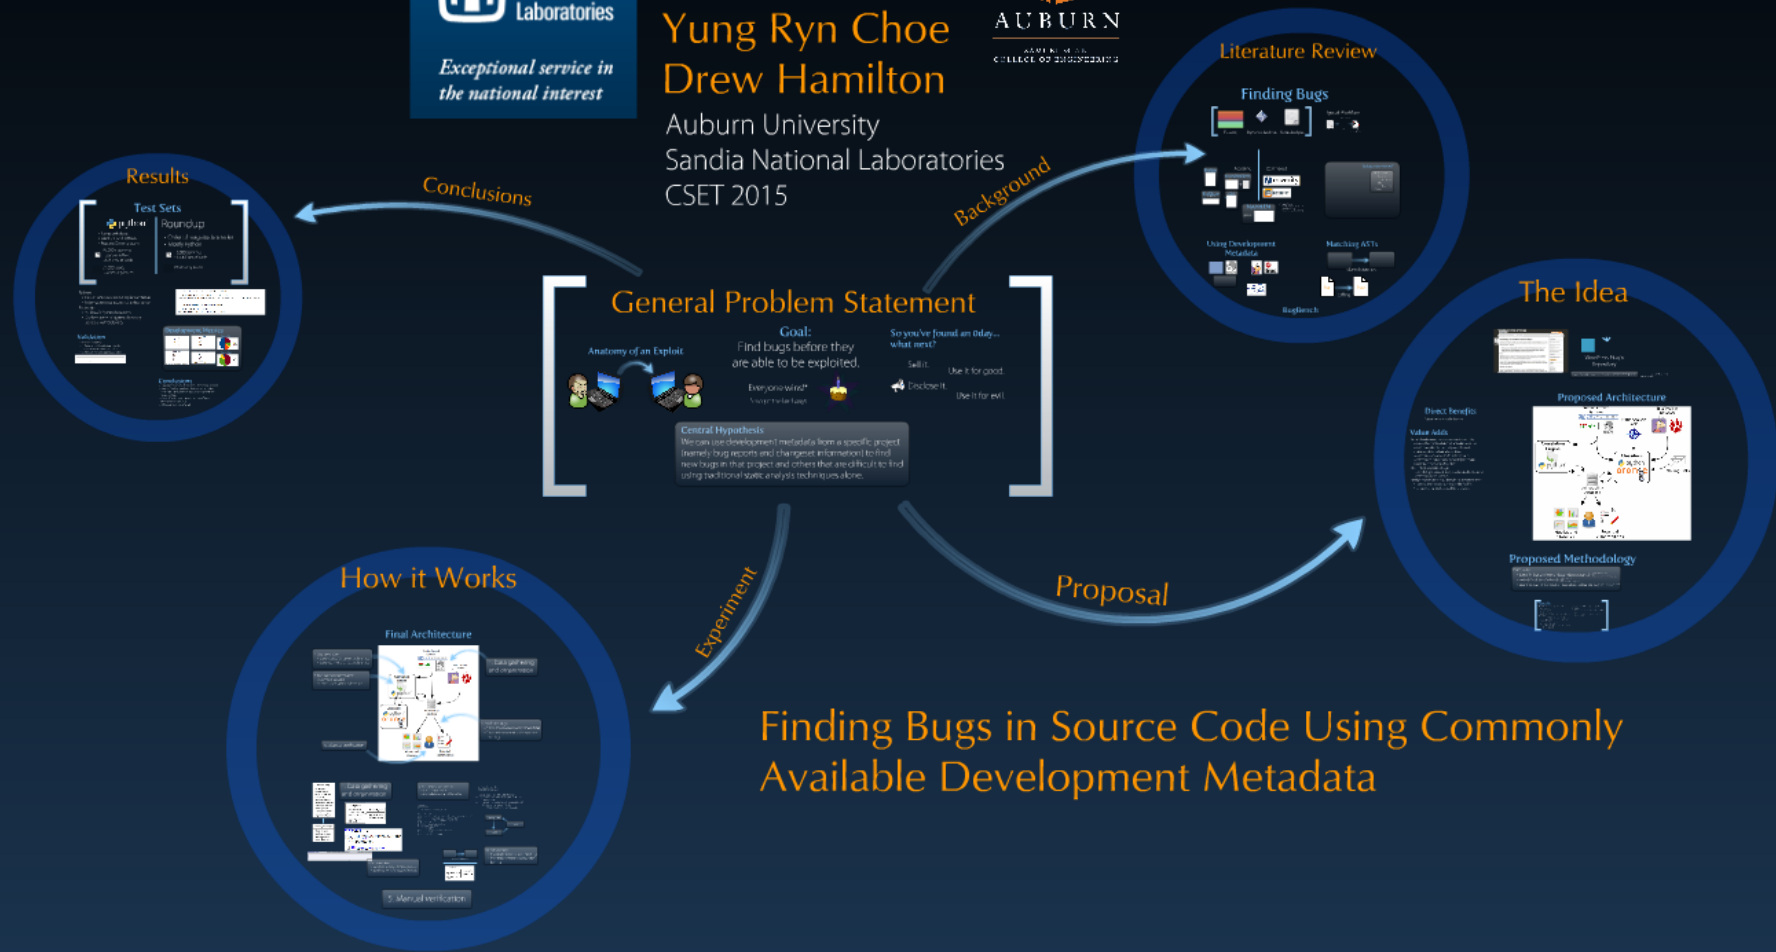




Devin Cook
Yung Ryn Choe
Drew Hamilton

Auburn University
Sandia National Laboratories
CSET 2015



Finding Bugs in Source Code Using Commonly Available Development Metadata

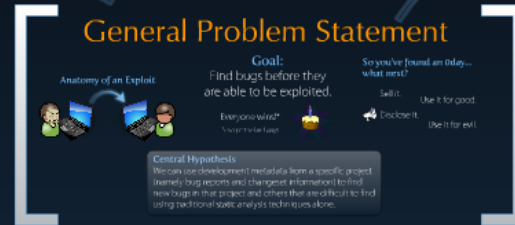


Devin Cook
Yung Ryn Choe
Drew Hamilton

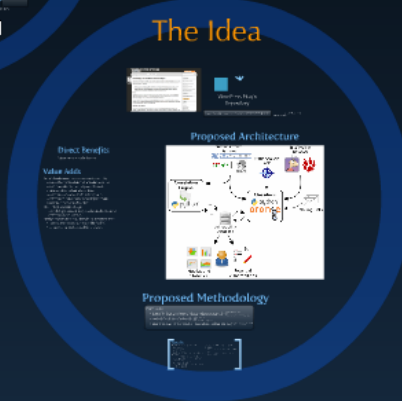
Auburn University
Sandia National Laboratories
CSET 2015




Conclusions



Background



Proposal



Experiment

Finding Bugs in Source Code Using Commonly Available Development Metadata

General Problem Statement

Anatomy of an Exploit



Goal:

Find bugs before they are able to be exploited.

Everyone wins!*

*except the bad guys



So you've found an Oday... what next?



Sell it.



Disclose it.

Use it for good.



Use it for evil.



Central Hypothesis

We can use development metadata from a specific project (namely bug reports and changeset information) to find new bugs in that project and others that are difficult to find using traditional static analysis techniques alone.

Anatomy of an Exploit



Goal:

Find bugs before they
are able to be exploited.

Everyone wins!*

*except the bad guys



So you've found an 0day... what next?



Sell it.

Use it for good.



Disclose it.

Use it for evil.





Everyone wins!*

*except the bad guys



Central Hypothesis

We can use development metadata from a specific project (namely bug reports and changeset information) to find new bugs in that project and others that are difficult to find using traditional static analysis techniques alone.



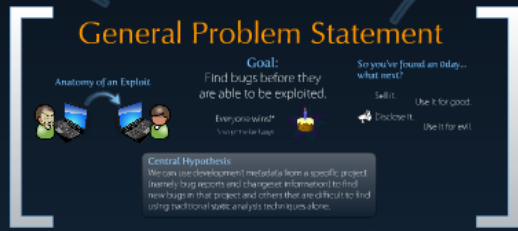
Devin Cook
Yung Ryn Choe
Drew Hamilton



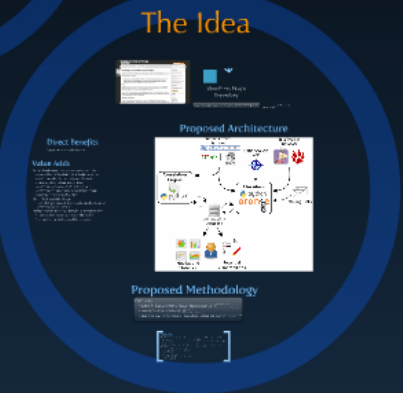
Auburn University
Sandia National Laboratories
CSET 2015



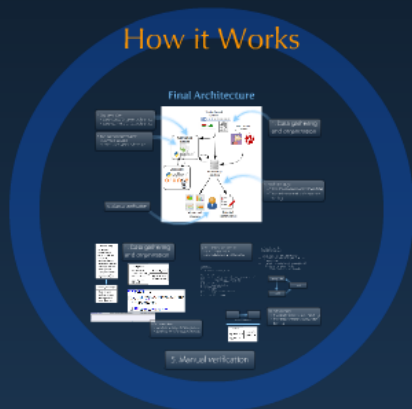
Conclusions



Background



Proposal



Experiment

Finding Bugs in Source Code Using Commonly Available Development Metadata

Literature Review

Finding Bugs



Typical Workflow



Academic

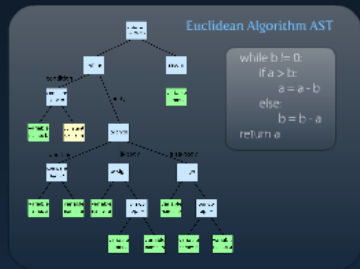
- RASAT
- BitBlaze
- rosecheckers
- ISA
- MAYHEM

Commercial

- coverity
- FORTIFY

Common Themes:


- Modular front and back end
- Some use ASTs
- Done! Song and David Bunney



Using Development Metadata


Subversion, mercurial, git, trac, Maven

Matching ASTs



BugBench

Finding Bu



Fuzzing



Dynamic Analysis

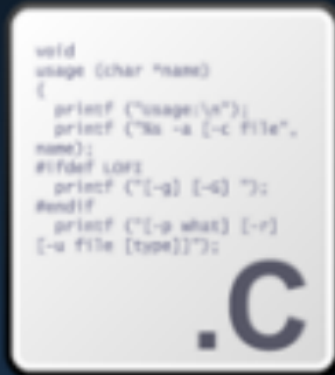


Static Analysis



Typical Workflow

Rules/Plugins



Source Code



Tool

Potential
Vulnerabilities



Academic




Commercial

RASAr




rosecheckers



BitBlaze

(Binary Only)

Vine




ISA



MAYHEM

Fairly standard detection techniques (although optimized), but generates working exploits.



Common Themes:

- Modular (front end/back end)
- Some kind of AST/IR
- Dawn Song and David Brumley

Literature Review

Finding Bugs



Typical Workflow



Academic

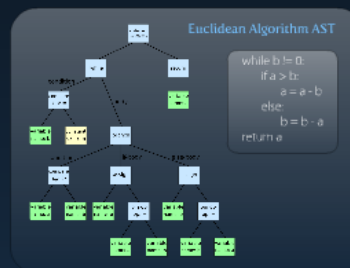
- RASAT
- BitBlaze
- rosecheckers
- ISA
- MAYHEM

Commercial

- coverity
- FORTIFY

Common Themes:


- Modular front and back end
- Some use ASTs
- Owen Song and David Bunney



Using Development Metadata

mercurial git trac

Matching ASTs



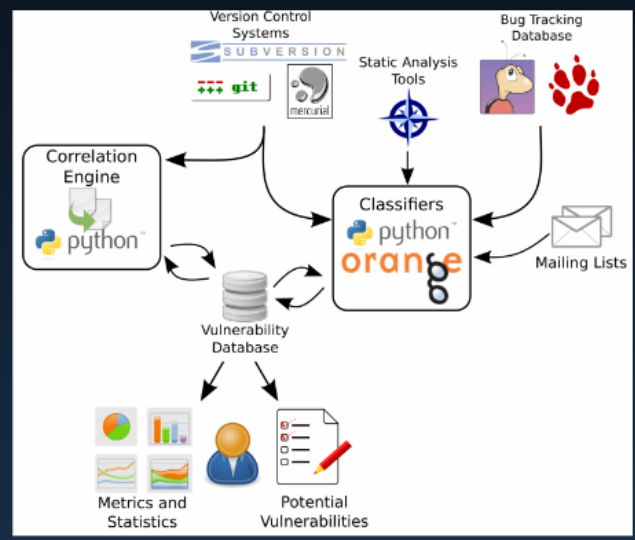
BugBench

The Idea



WordPress Plugin Repository

Proposed Architecture



Direct Benefits

More secure code. Easier.

Value Adds



- Better development metrics related to security
- committers with a habit of introducing vulns
- common vulns found in a project's code
- number of vulns found over time
- committers good at finding/fixing vulns
- common mistakes certain developers make
- average time to find/fix vulns

Not limited to security bugs
Everything discussed may also be applied to other common classes of bugs

- Interfacing with existing vulnerability detection tools
- incorporation in to larger tools (like BASAr)
 - using existing tools as additional signals

Proposed Methodology

- Core tasks:
- Identify that a piece of data relates to a vuln
 - correlate all related data
 - use that new information to detect other vulns

Details

- Identify all data in repo
- Correlate data
- Analyze data
- Identify vulns
- Correlate vulns
- Use vulns to detect other vulns

SPARE CLOCK CYCLES

HACKING IS FREEDOM.

Subscribe via RSS

HOME

ABOUT ME (AND THIS BLOG)

CODE AND SUCH

18

SEP/11

14

Explo(it)r)ing the WordPress Extension Repos

Today's post is kind of long, so I thought I should warn you in advance by adding an additional paragraph for you to read. I also wanted to provide download links for those who'd rather just read the code. It isn't the cleanest code in the world, so I apologize in advance. I discuss what all of these are for and how they work later on in the post, so if you're confused and/or curious, read on. Downloads:

- Copies of the WordPress theme and plugin repositories can be grabbed via [torrent](#) (Please note that the plugin repo has a few directories incomplete/missing; this can be fixed by running my checkout code)
- A new WordPress plugin fingerprinting tool, [wpfinger](#) ([download](#)). This tool can infer detailed version information on just about every plugin in the WordPress repository. This package also contains some useful libraries for checking out the repositories and scraping plugin rankings, as this is used in the fingerprinting tool.

Intro

After finding an [arbitrary file upload vulnerability in 1 Flash Gallery](#), I became curious as to how many other WordPress plugins made basic security mistakes. The 1 Flash Gallery plugin issue, it seems, is that they CTRL-C-V'd code from a project called [Uploadify](#), which has been known to be vulnerable [for quite awhile](#).

After realizing this, I became curious as to how many plugins make easy-to-spot security mistakes, such as reusing vulnerable libraries or doing such things as `include($_REQUEST['lulz'])`. However, my curiosity was initially somewhat hampered by the fact that downloading and auditing every WordPress plugin one at a time is not only a mind numbing task, but a herculean one as well. And, well, I'm incredibly lazy.

Getting the Repos

So what to do? Well, it turns out that WordPress is nice enough to have public repositories (<http://plugins.svn.wordpress.org> and <http://themes.svn.wordpress.org>) containing all plugins that have ever been submitted, as well as every theme. This, of course, was exciting: I could just check this out, whip out some grep-fu, and have my answers.

Recent Posts

[Stack Necromancy: Defeating Debuggers By Raising the Dead](#)

[Exploiting an IP Camera Control Protocol: Redux](#)

[Explo\(it\)r\)ing the WordPress Extension Repos](#)

[1 Flash Gallery: Arbitrary File Upload](#)

[Sergio Proxy v0.2 Released](#)

Beer Fund!

Find something particularly useful or interesting? Buy me a beer!

[Donate](#)

Blogroll

[Cujef.com](#)

[ha.ckers.org](#)

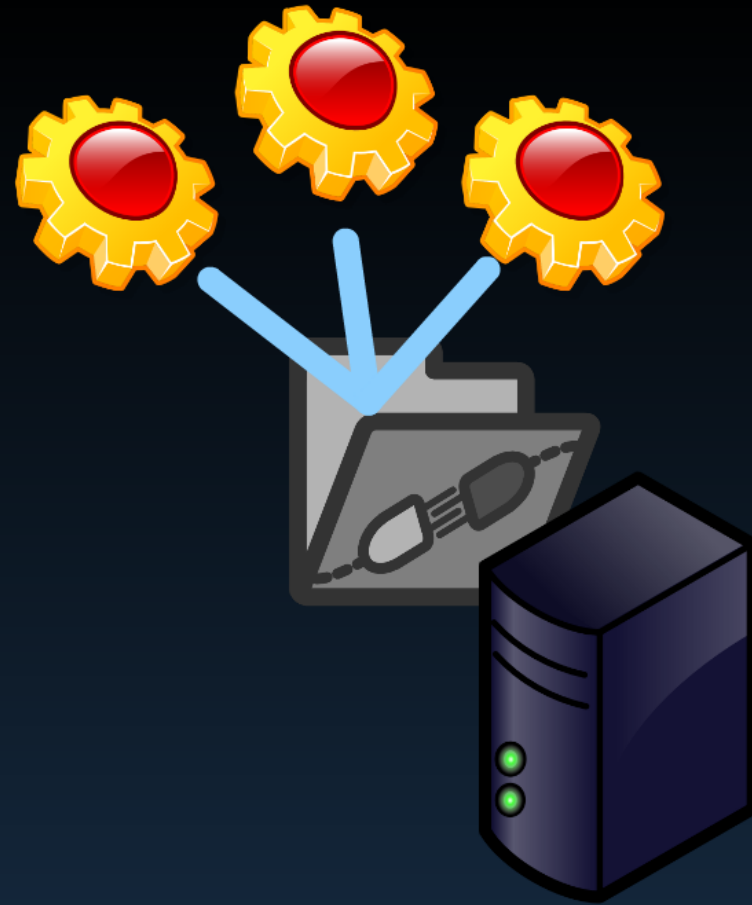
[Mark Adam's Blog](#)

[Metasploit Blog](#)

[Schneier on Security](#)



Prezi



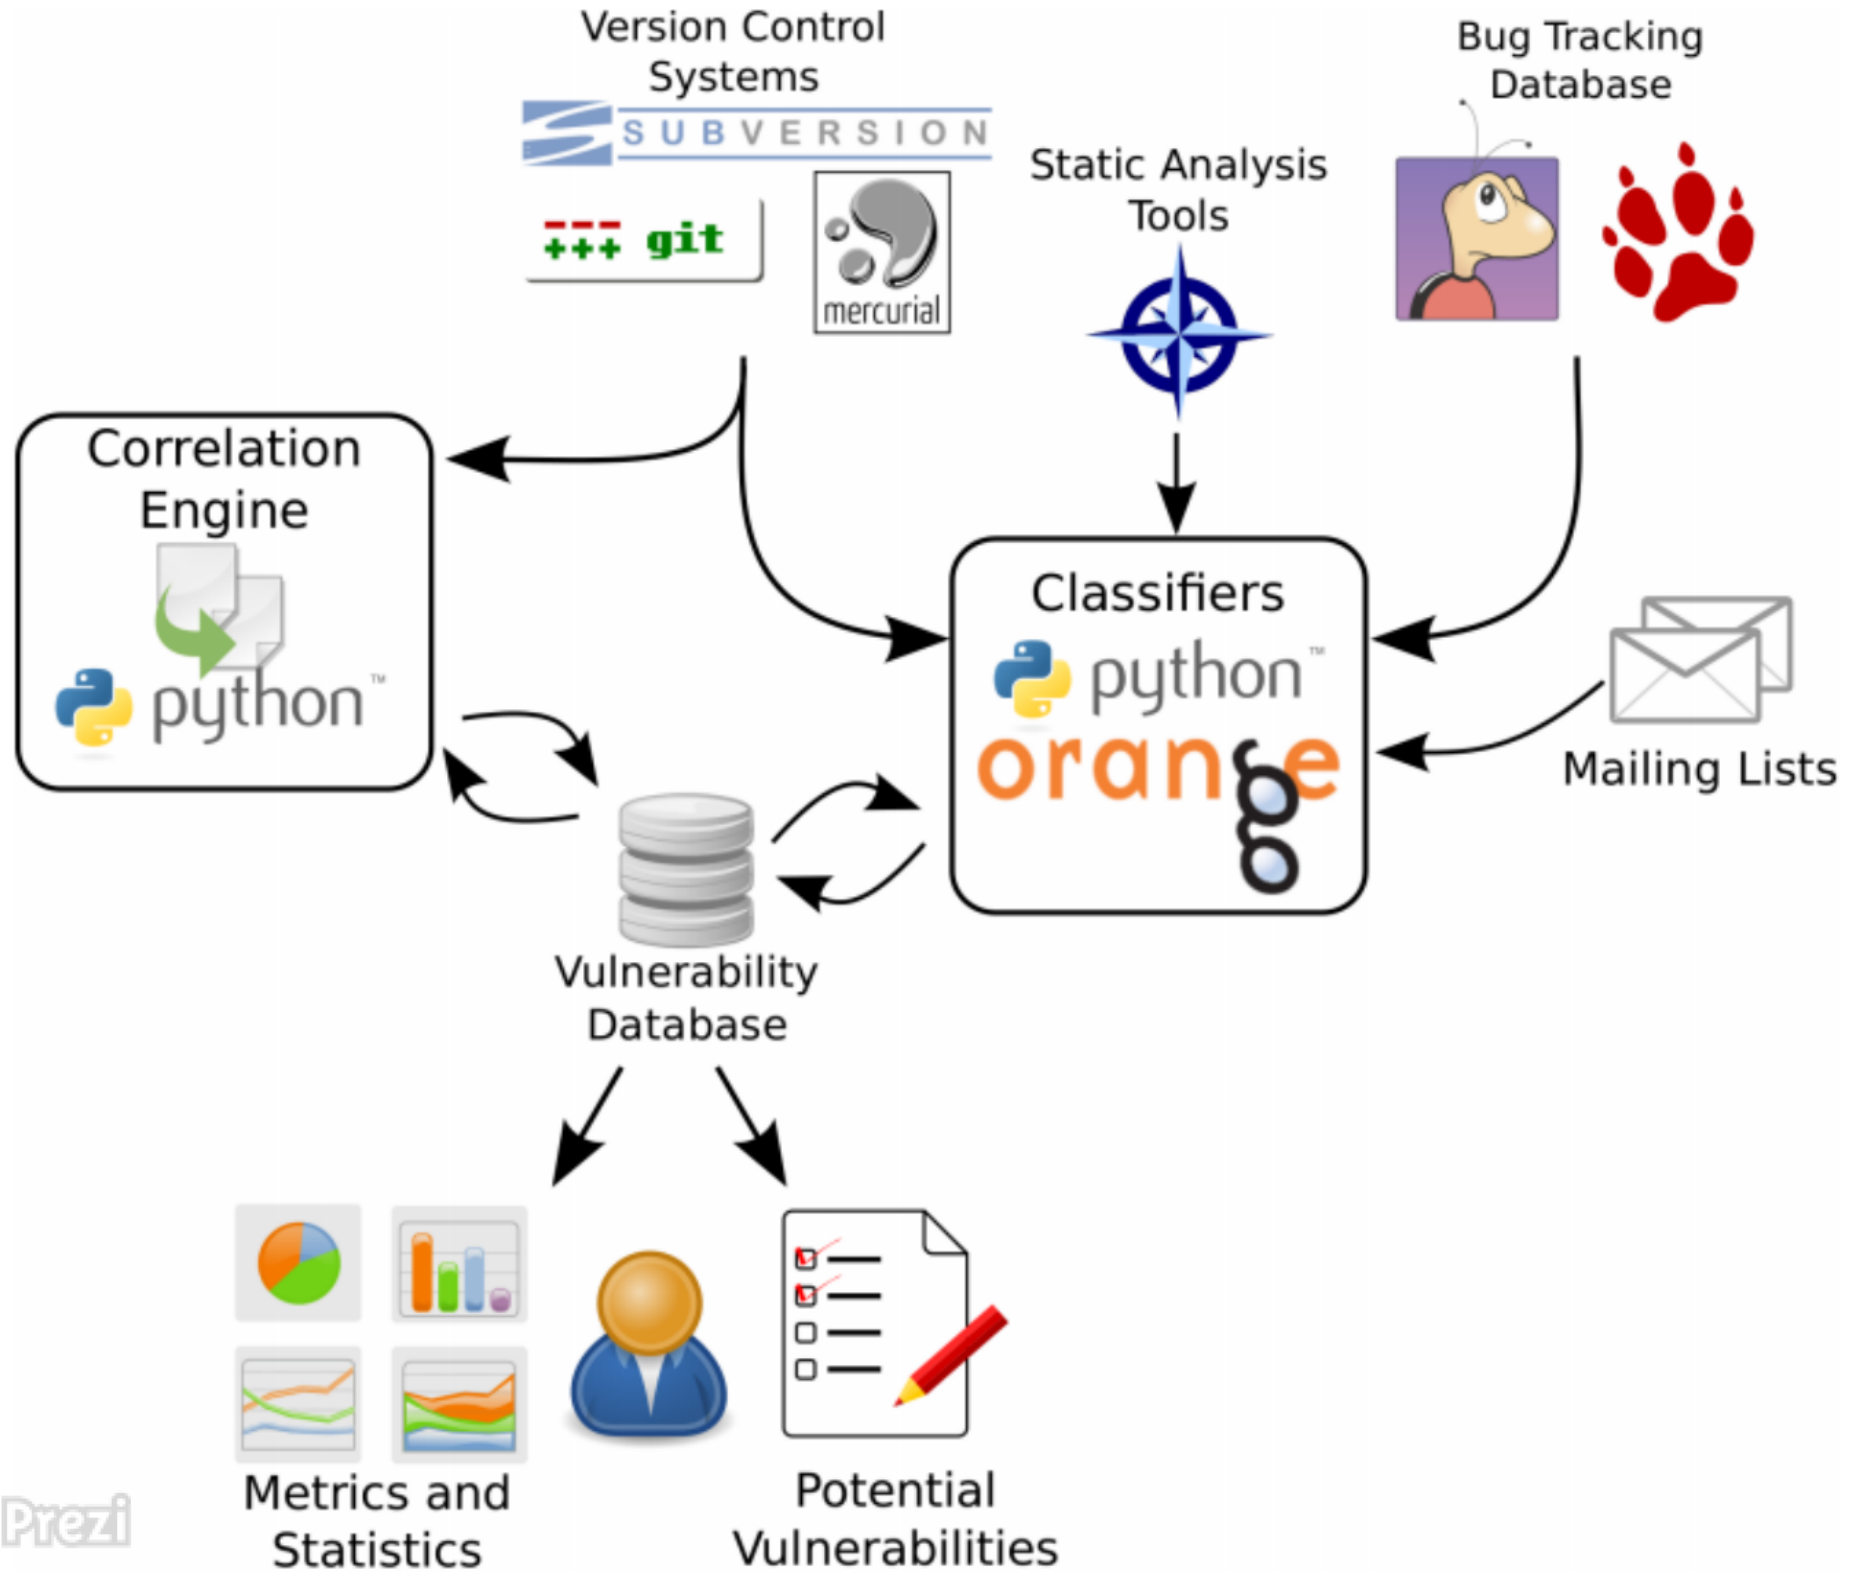
WordPress Plugin Repository



WordPress Plugin Repository

```
$ egrep -i '(include|require)(_once)?(\(|\s+)[^[:;]]*\$_(REQUEST|GET|POST|COOKIE)'
```

= 11 unauthenticated RFI vulnerabilities
...little to no effort



Proposed Methodology

Core tasks:

- Identify that a piece of data relates to a vuln } Explicit metadata tags
AI classifiers
relation to other data already identified
- correlate all related data } Dereference links
Recurse
Build DB of indicators/signals
- use that new information to detect other vulns } assign "security score" to commits
examine existing code

Details

Use additional data as input - more "signals"

- bugtrackers
- commit messages
- mailing lists
- public vuln DBs (known-vuln code/packages)

Find patterns

- certain committers that make certain mistakes
- certain files or areas that often contain bugs
- vuln code reuse
- more vulns late at night?
- devs with short names introduce more vulns? - could be anything!

Generate a DB of vulnerable patterns/code

Look for bugs

- Given a new commit, compute the "security score"
- Given a new repo (or just some source), go find some 0days!

Proposed Methodology

Core tasks:

- Identify that a piece of data relates to a vuln } Explicit metadata tags
AI classifiers
relation to other data already identified
- correlate all related data } Dereference links
Recurse
Build DB of indicators/signals
- use that new information to detect other vulns } assign "security score" to commits
examine existing code

Details

Use additional data as input - more "signals"

- bugtrackers
- commit messages
- mailing lists
- public vuln DBs (known-vuln code/packages)

Find patterns

- certain committers that make certain mistakes

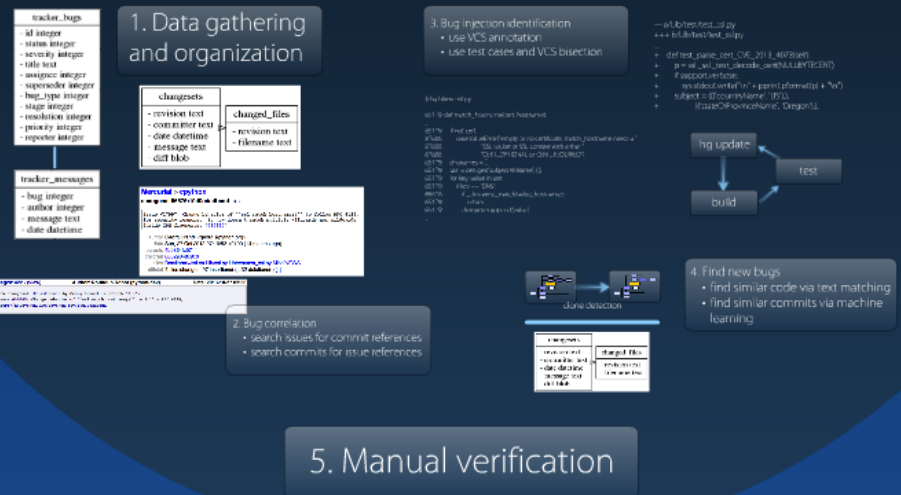
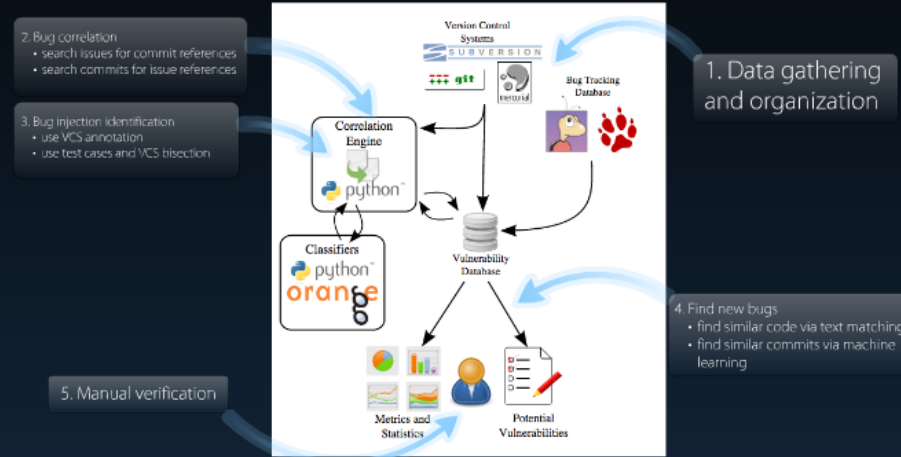
Generate a DB of vulnerable patterns/code

Look for bugs

- Given a new commit, compute the "security score"
- Given a new repo (or just some source), go find some 0days!

How it Works

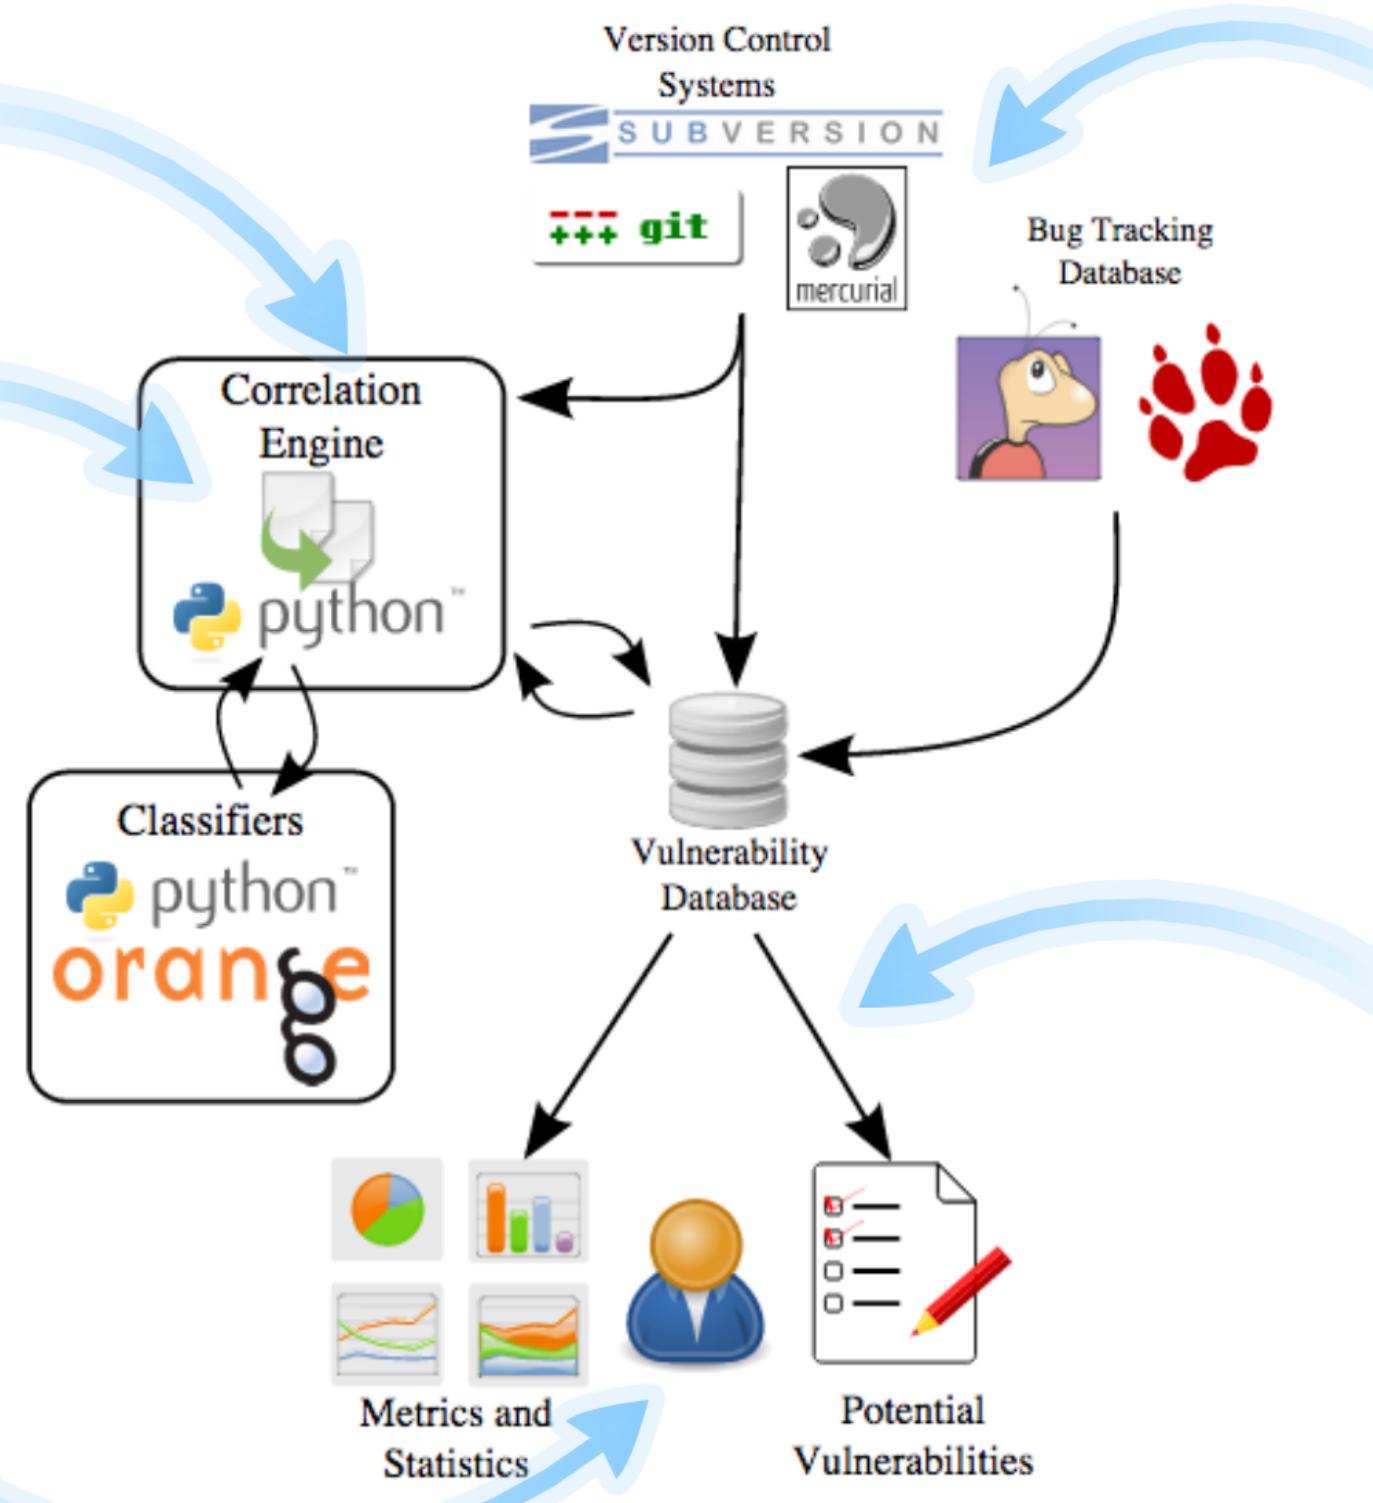
Final Architecture



commit references
value references

on
dissection

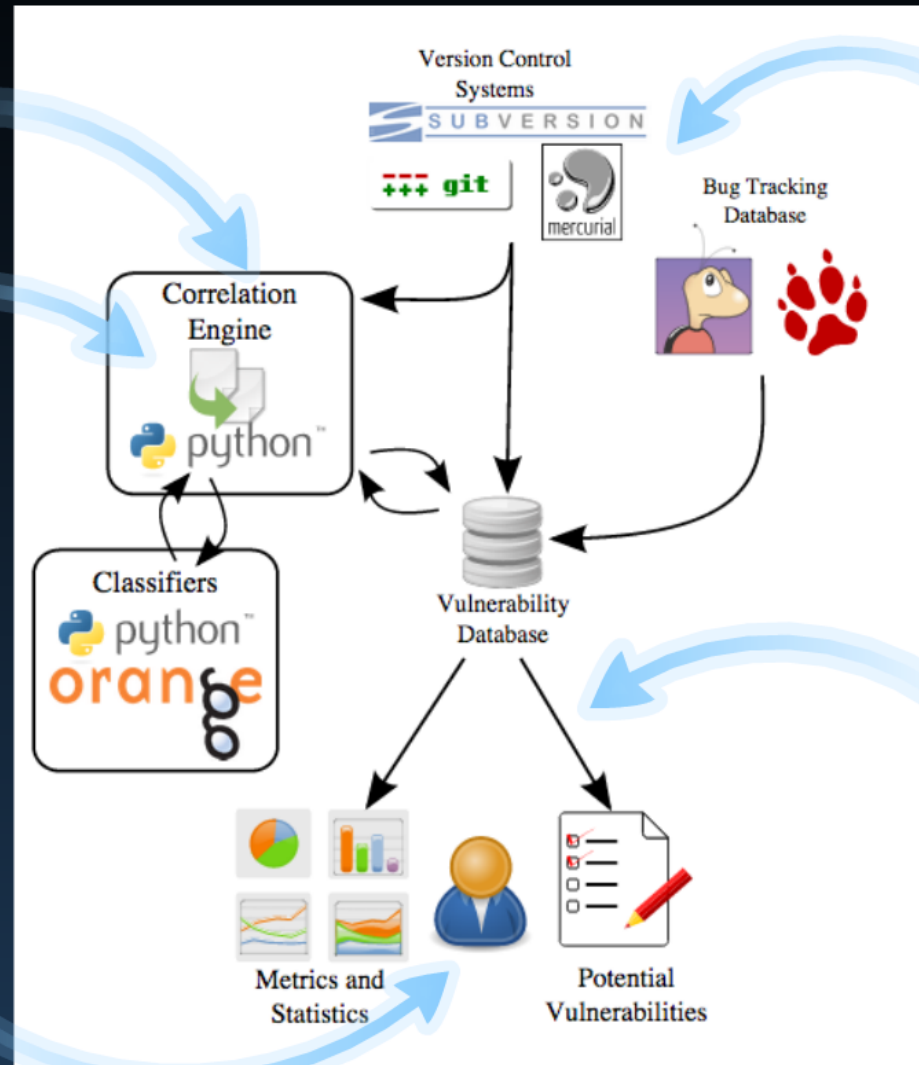
ification



1. Data
and or

4. Find new bugs
- find similar c
 - find similar c
- learning

Final Architecture



1. Data gathering and organization

4. Find new bugs

- find similar code via text matching
- find similar commits via machine learning

5. Manual verification

2. Bug correlation

- search issues for commit references
- search commits for issue references

3. Bug injection identification

- use VCS annotation
- use test cases and VCS bisection

1. Data gathering

3. Bug injection identification

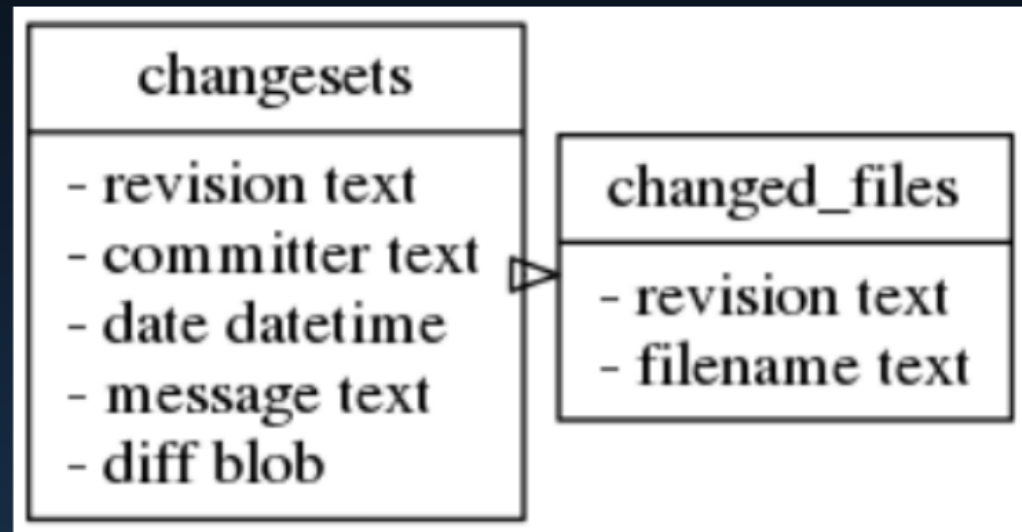
- use VCS annotation

```
-- a/Lib/test/test_ssl.py  
+++ b/Lib/test/test_ssl.py
```


1. Data gathering and organization

tracker_bugs
- id integer
- status integer
- severity integer
- title text
- assignee integer
- superseder integer
- bug_type integer
- stage integer
- resolution integer
- priority integer
- reporter integer

tracker_messages
- bug integer
- author integer
- message text
- date datetime



Mercurial > cpython

changeset 86676:10d0edadbccd 3.3

```
Issue #17997: Change behavior of ``ssl.match_hostname()`` to follow RFC 4343 for security reasons. It now doesn't match multiple wildcards nor wildcards inside IDN fragments. [#17997]
```

author Georg Brandl <georg@python.org>

date Sun, 27 Oct 2013 07:16:53 +0100 (14 months ago)



- stage integer
- resolution integer
- priority integer
- reporter integer

tracker_messages

- bug integer
- author integer
- message text
- date datetime

- revision text
- committer text
- date datetime
- message text
- diff blob

changed_files

- revision text
- filename text

Mercurial > cpython

changeset 86676:10d0edadbcdd 3.3

Issue #17997: Change behavior of `ssl.match_hostname()` to follow RFC 6125, for security reasons. It now doesn't match multiple wildcards nor wildcards inside IDN fragments. [#17997]

author Georg Brandl <georg@python.org>
 date Sun, 27 Oct 2013 07:16:53 +0100 (14 months ago)
 parents 4b0364fc5711
 children 68029048c9c6
 files Doc/library/ssl.rst Lib/ssl.py Lib/test/test_ssl.py Misc/NEWS
 diffstat 4 files changed, 97 insertions(+), 32 deletions(-) [+]

msg201422 - (view)

Author: Roundup Robot (python-dev)

Date: 2013-10-27 01:38

New changeset 10d0edadbcdd by Georg Brandl in branch '3.3':
 Issue #17997: Change behavior of `ssl.match_hostname()` to follow RFC 6125,
<http://hg.python.org/cpython/rev/10d0edadbcdd>

2. Bug correlation

- search issues for commit references
- search commits for issue references

...
65179: c
...
65179:
87688:
87688:
87688:
65179:
65179:
65179:
65179:
65179:
86676:
65179:
65179:
...

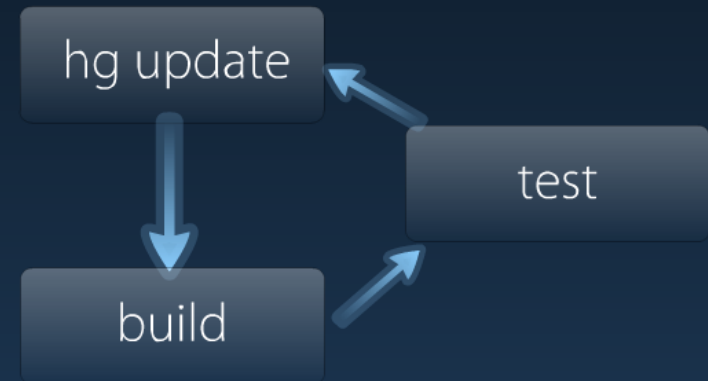
3. Bug injection identification

- use VCS annotation
- use test cases and VCS bisection

```
$ hg blame ssl.py
```

```
...
65179: def match_hostname(cert, hostname):
...
65179:     if not cert:
87688:         raise ValueError("empty or no certificate, match_hostname needs a "
87688:             "SSL socket or SSL context with either "
87688:             "CERT_OPTIONAL or CERT_REQUIRED")
65179:     dnsnames = []
65179:     san = cert.get('subjectAltName', ())
65179:     for key, value in san:
65179:         if key == 'DNS':
86676:             if _dnsname_match(value, hostname):
65179:                 return
65179:     dnsnames.append(value)
...
```

```
--- a/Lib/test/test_ssl.py
+++ b/Lib/test/test_ssl.py
...
+ def test_parse_cert_CVE_2013_4073(self):
+     p = ssl._ssl._test_decode_cert(NULLBYTECERT)
+     if support.verbose:
+         sys.stdout.write("\n" + pprint.pformat(p) + "\n")
+     subject = (('countryName', 'US'),),
+         (('stateOrProvinceName', 'Oregon'),),
```





4. Find new bugs

- find similar code via text match

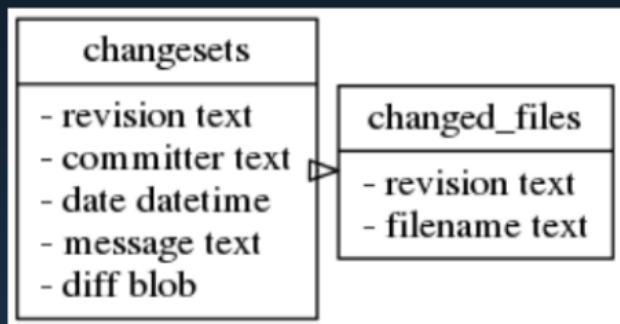
h(value, hostname):

(value)



4. Find new bugs


- find similar code via text matching
- find similar commits via machine learning



1. Data gathering and organization

tracker_bugs
- id integer
- status integer
- severity integer
- title text
- assignee integer
- superseder integer
- bug_type integer
- stage integer
- resolution integer
- priority integer
- reporter integer

tracker_messages
- bug integer
- author integer
- message text
- date datetime



```

Mercurial > cpython
changeset 06676:10d0dadbedd 33

Issue #1997: Change behavior of ``ssl.match_hostname`` to follow RFC 6125,
for security reasons. It now doesn't match multiple wildcards nor wildcards
inside CN fragments. [#17997]

author Georg Brandt <georg@python.org>
date Sun, 27 Oct 2013 07:16:53 +0100 (14 months ago)
parents 45c0641c5711
children 03022918-9a98
files Doc/Library/ssl.rst Lib/ssl.py Lib/test/ssl.py Misc/NEWS
diffstat: 4 files changed, 97 insertions(+), 32 deletions(-) [1]
  
```

3. Bug injection identification


- use VCS annotation
- use test cases and VCS bisection

```

$ hg blame ssl.py
65179: def match_hostname(cert, hostname):
...
65179:     if not cert:
87688:         raise ValueError("empty or no certificate, match_hostname needs a "
87688:             "SSL socket or SSL context with either "
87688:             "'CERT_OPTIONAL' or 'CERT_REQUIRED'")
65179:     dnsnames = []
65179:     san = cert.get('subjectAltName', ())
65179:     for key, value in san:
65179:         if key == 'DNS':
86676:             if _dnsname_match(value, hostname):
65179:                 return
65179:             dnsnames.append(value)
  
```

```

--- a/Lib/test/test_ssl.py
+++ b/Lib/test/test_ssl.py
...
+ def test_parse_cert_CVE_2013_4073(self):
+     p = ssl_ssl_test_decode_cert(NULLBYTECERT)
+     if support.verbose:
+         sys.stdout.write("\n" + pprint.pformat(p) + "\n")
+     subject = (('countryName', 'US'),),
+             (('stateOrProvinceName', 'Oregon'),),
  
```

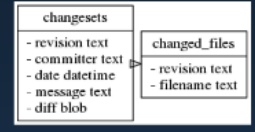



```

msg201422 - (view) Author: Roundup Robot (python-dev) Date: 2013-10-27 01:38
New changeset 10d0dadbedd by Georg Brandt in branch '3.3':
Issue #1997: change behavior of ``ssl.match_hostname`` to follow RFC 6125,
https://hg.python.org/eggsbot/rev/10d0dadbedd
  
```

2. Bug correlation

- search issues for commit references
- search commits for issue references



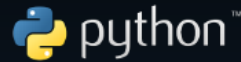
4. Find new bugs

- find similar code via text matching
- find similar commits via machine learning

5. Manual verification

Results

Test Sets



- Large codebase
- Mostly C and Python
- Routine Coverity scans

93,000+ commits
 192 committers
 ~1M lines of code
 34,000 issues
 162 security issues

Roundup

- Order of magnitude smaller
- Mostly Python

<5,000 commits
 ~100K lines of code
 22 security issues

Python:

- DoS in reference WSGI implementation
- Information disclosure via pydoc server

Roundup:

- 2x Unsafe Python functions
- 2 others were mitigated directory traversal vulnerabilities

```
[*] issue1071(closed) security(msg:1) related(0)
[] | self.raw.requestline = self.rfile.readline()
cpython/Lib/wsgref/simple.server.py: self.raw.requestline = self.rfile.readline()


[*] issue679056(closed) security(msg:0) related(2)
[] | self.address = ("", port)
cpython/Lib/pydoc.py: self.address = ("", port)
```

Validation

- Can we find 0days?
- Is the false positive rate reasonable?
 - Coverity developers say <30%
- Is the false negative rate reasonable?

Project	Issues Reported	Total Issues	Issues Dismissed	Outliers	True Positives
Python	20/22	22	2	0	2
Roundup	6/22	22	9	4	2

Development Metrics




Conclusions

- Novel approach for identifying vulnerability injection
- Good at finding new bugs that are similar to bugs that have already been fixed (and can sometimes reuse old fixes)
- Can still find bugs in projects with insufficient development metadata
- Almost entirely automated



Test Sets



- Large codebase
- Mostly C and Python
- Routine Coverity scans

93,000+ commits



192 committers

~1M lines of code



34,000 issues

162 security issues

Roundup

- Order of magnitude smaller
- Mostly Python

<5,000 commits



~100K lines of code



22 security issues

Python:



WSGI in reference WSGI implementation


information disclosure via pydoc server

```
[*] issue10714{closed|security|msg:4|related:3}
```

```
[ ] self.raw_requestline = self.rfile.readline()
```

```
cpython/Lib/wsgiref/simple_server.py: self.raw_requestline = self.rfile
```

```
[*] issue672656{closed|security|msg:5|related:2}
```



- Large codebase
- Mostly C and Python
- Routine Coverity scans

93,000+ commits



192 committers

~1M lines of code



34,000 issues

162 security issues



Roundup

- Order of magnitude smaller
- Mostly Python



<5,000 commits
~100K lines of code



22 security issues

Python:

- DoS in reference WSGI implementation
- Information disclosure via pydoc server

Roundup:

- 2x Unsafe Python functions
- 2 others were mitigated directory traversal vulnerabilities

```
[*] issue10714{closed | security | msg:4 | related:3}
```

```
[ ] self.raw_requestline = self.rfile.readline()
```

```
cpython/Lib/wsgiref/simple_server.py: self.raw_requestline = self.rfile.readline()
```

```
[*] issue672656{closed | security | msg:5 | related:2}
```

```
[ ] self.address = ('', port)
```

```
cpython/Lib/pydoc.py: self.address = ('', port)
```

Validation

- Can we find 0days?
- Is the false positive rate reasonable?
 - Coverity developers say <30%
- Is the false negative rate reasonable?

Project	LOC (Approx.)	Total Issues	Code Signatures	Candidates	True Positives
Python	976,413	162	38	21	2
Roundup	88,578	22	6	4	2

Table 4.5: Vulnerability Detection Rates

Development Metrics

Developer	Rate
Dev00	14.29%
Dev01	4.35%
Dev02	3.23%
Dev03	2.78%
Dev04	2.53%
Dev05	2.08%
Dev06	1.59%
Dev07	1.59%
Dev08	1.19%
Dev09	1.15%

Table 4.1: Python Developer Vulnerability Injection Rate (Top 10)

Developer	Resolutions
Dev14	37
Dev12	34
Dev31	24
Dev37	24
Dev36	21
Dev42	20
Dev18	14
Dev43	11
Dev24	11
Dev30	10

Table 4.3: Python Developer Vulnerability Resolutions (Top 10)


Developer	Reports
Dev12	18
Dev44	10
Dev18	5
Dev45	3
Dev46	3
Dev47	3
Dev14	3
Dev24	3
Dev48	2
Dev30	2

Table 4.2: Python Developer Vulnerability Reports (Top 10)


File	Issues
Objects/unicodeobject.c	16
Modules/_ssl.c	15
Lib/ssl.py	13
Lib/zipfile.py	11
Lib/platform.py	10
Lib/netrc.py	10
Lib/smtpd.py	9
Python/sysmodule.c	9
Lib/os.py	8
Python/mysnprintf.c	8

Table 4.4: Files in the Python Project Containing Security Issues (Top 10)

Time to Fix



Time to Find



Developer	Rate
Dev00	14.29%
Dev01	4.35%
Dev02	3.23%
Dev03	2.78%
Dev04	2.53%
Dev05	2.08%
Dev06	1.59%
Dev07	1.59%
Dev08	1.19%
Dev09	1.15%

Table 4.1: Python Developer Vulnerability Injection Rate (Top 10)

Developer	Reports
Dev12	18
Dev44	10
Dev18	5
Dev45	3
Dev46	3
Dev47	3
Dev14	3
Dev24	3
Dev48	2
Dev30	2

Table 4.2: Python Developer Vulnerability Reports (Top 10)

Table 4.3: Python Developer Vulnerability Resolutions (Top 10)


File	Issues
Objects/unicodeobject.c	16
Modules/_ssl.c	15
Lib/ssl.py	13
Lib/zipfile.py	11
Lib/platform.py	10
Lib/netrc.py	10
Lib/smtpd.py	9
Python/sysmodule.c	9
Lib/os.py	8
Python/mysnprintf.c	8

Table 4.4: Files in the Python Project Containing Security Issues (Top 10)


Developer	Resolutions
Dev14	37
Dev12	34
Dev31	24
Dev37	24
Dev36	21
Dev42	20
Dev18	14
Dev43	11
Dev24	11
Dev30	10

Table 4.3: Python Developer Vulnerability Resolutions (Top 10)

Time to Fix



Time to Find



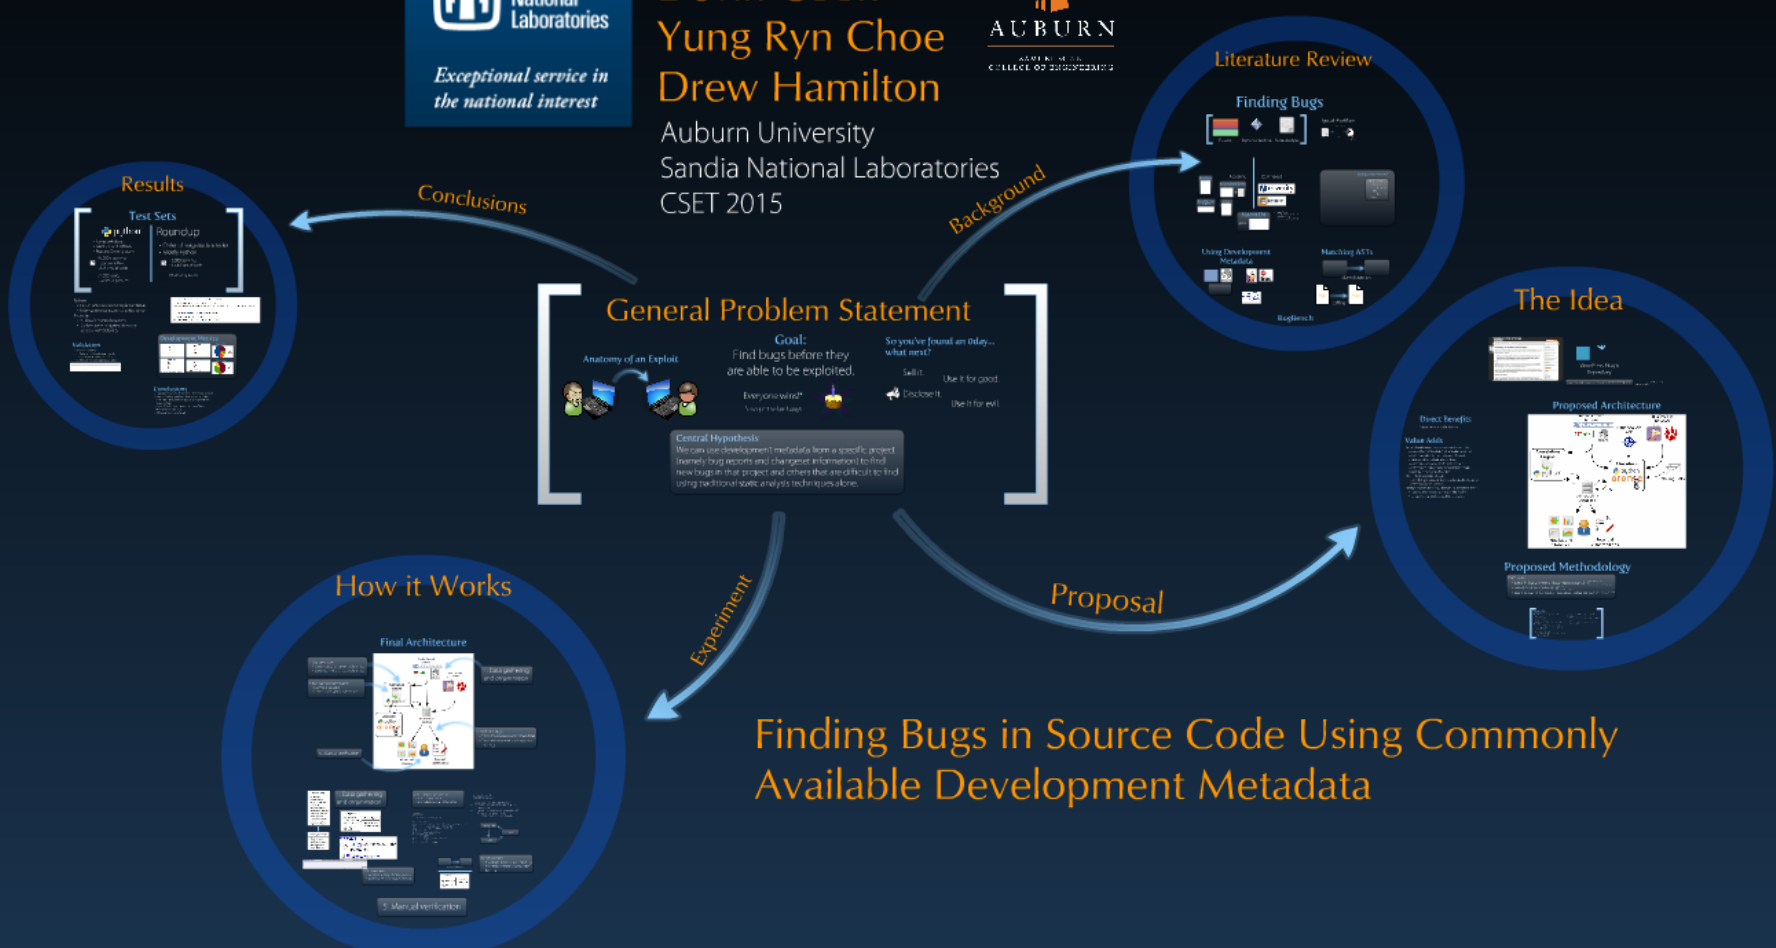
Conclusions

- Novel approach for identifying vulnerability injection
- Good at finding new bugs that are similar to bugs that have already been fixed (and can sometimes reuse old fixes)
- Can still find bugs in projects with insufficient development metadata
- Almost entirely automated



Devin Cook
Yung Ryn Choe
Drew Hamilton

Auburn University
Sandia National Laboratories
CSET 2015



Finding Bugs in Source Code Using Commonly Available Development Metadata