

# Make your tests fail

How machines can help find corner cases that humans ignore

Isabel Drost-Fromm



elastic

# Isabel Drost-Fromm



Software Engineer Elasticsearch

Member Apache Software Foundation

Co-founder Apache Mahout

Co-founder Berlin Buzzwords

“Not good enough an engineer to write code w/o tests.”



elastic

```
public void testReclaimedDeletes() throws Exception {
    setIndexWriterMaxDocs(10);
    try {
        Directory dir = newDirectory();
        IndexWriter w = new IndexWriter(dir, new IndexWriterConfig(null));
        for(int i=0;i<10;i++) {
            Document doc = new Document();
            doc.add(newStringField("id", ""+i, Field.Store.NO));
            w.addDocument(doc);
        }

        // Delete 5 of them:
        for(int i=0;i<5;i++) {
            w.deleteDocuments(new Term("id", ""+i));
        }

        w.forceMerge(1);

        assertEquals(5, w.maxDoc());

        // Add 5 more docs
        for(int i=0;i<5;i++) {
            w.addDocument(new Document());
        }

        // 11th document should fail:
        try {

```

You do write tests – ideally before writing the code itself?



```
>>
743 9417  public IndexWriter(Directory d, IndexWriterConfig conf) throws IOException
744 9417    conf.setIndexWriter(this); // prevent reuse by other instances
745 9416    config = conf;
746 9416    directory = d;
747 9416    analyzer = config.getAnalyzer();
748 9416    infoStream = config.getInfoStream();
749 9416    mergeScheduler = config.getMergeScheduler();
750 9416    codec = config.getCodec();
751
752 9416    bufferedUpdatesStream = new BufferedUpdatesStream(infoStream);
753 9416    poolReaders = config.getReaderPooling();
754
755 9416    writeLock = directory.makeLock(WRITE_LOCK_NAME);
756
757 9414    if (!writeLock.obtain(config.getWriteLockTimeout())) // obtain write lock
758 0      throw new LockObtainFailedException("Index locked for write: " + writ
759
760 9414    boolean success = false;
761 9414    try {
762 9414      OpenMode mode = config.getOpenMode();
763 9414      boolean create;
764 9414      if (mode == OpenMode.CREATE) {
765 1963        create = true;
766 7451      } else if (mode == OpenMode.APPEND) {
767 802        create = false;
768    } else {
```

You regularly check code coverage?



 [Back to Dashboard](#)

 [Status](#)

 [Changes](#)

 [Dependency Graph](#)

 [Document](#)

 [Email Template Testing](#)

 **Build History** [\(trend\) ▾](#)

-  #710 [Dec 15, 2014 8:37:25 AM](#)
-  #709 [Dec 13, 2014 9:51:25 PM](#)
-  #708 [Dec 12, 2014 11:11:03 PM](#)
-  #707 [Dec 11, 2014 3:16:33 PM](#)
-  #706 [Dec 10, 2014 10:27:58 PM](#)
-  #705 [Dec 9, 2014 11:14:25 PM](#)
-  #704 [Dec 8, 2014 3:55:42 AM](#)

## Project Lucene-Artifa



[Document](#)



[Last Successful Artifacts](#)



[Recent Changes](#)

### Upstream Projects

 [Lucene-Solr-NightlyTests-5.x](#)

### Downstream Projects

 [Solr-Artifacts-5.x](#)

### Permalinks

You run your tests regularly?





Java:

```
int input = generateInteger(  
    Integer.MIN_VALUE,  
    Integer.MAX_VALUE);  
  
int output = Math.abs(input);
```



Java:

```
int input = generateInteger(  
    Integer.MIN_VALUE,  
    Integer.MAX_VALUE);  
  
int output = Math.abs(input);
```

output == -2147483648



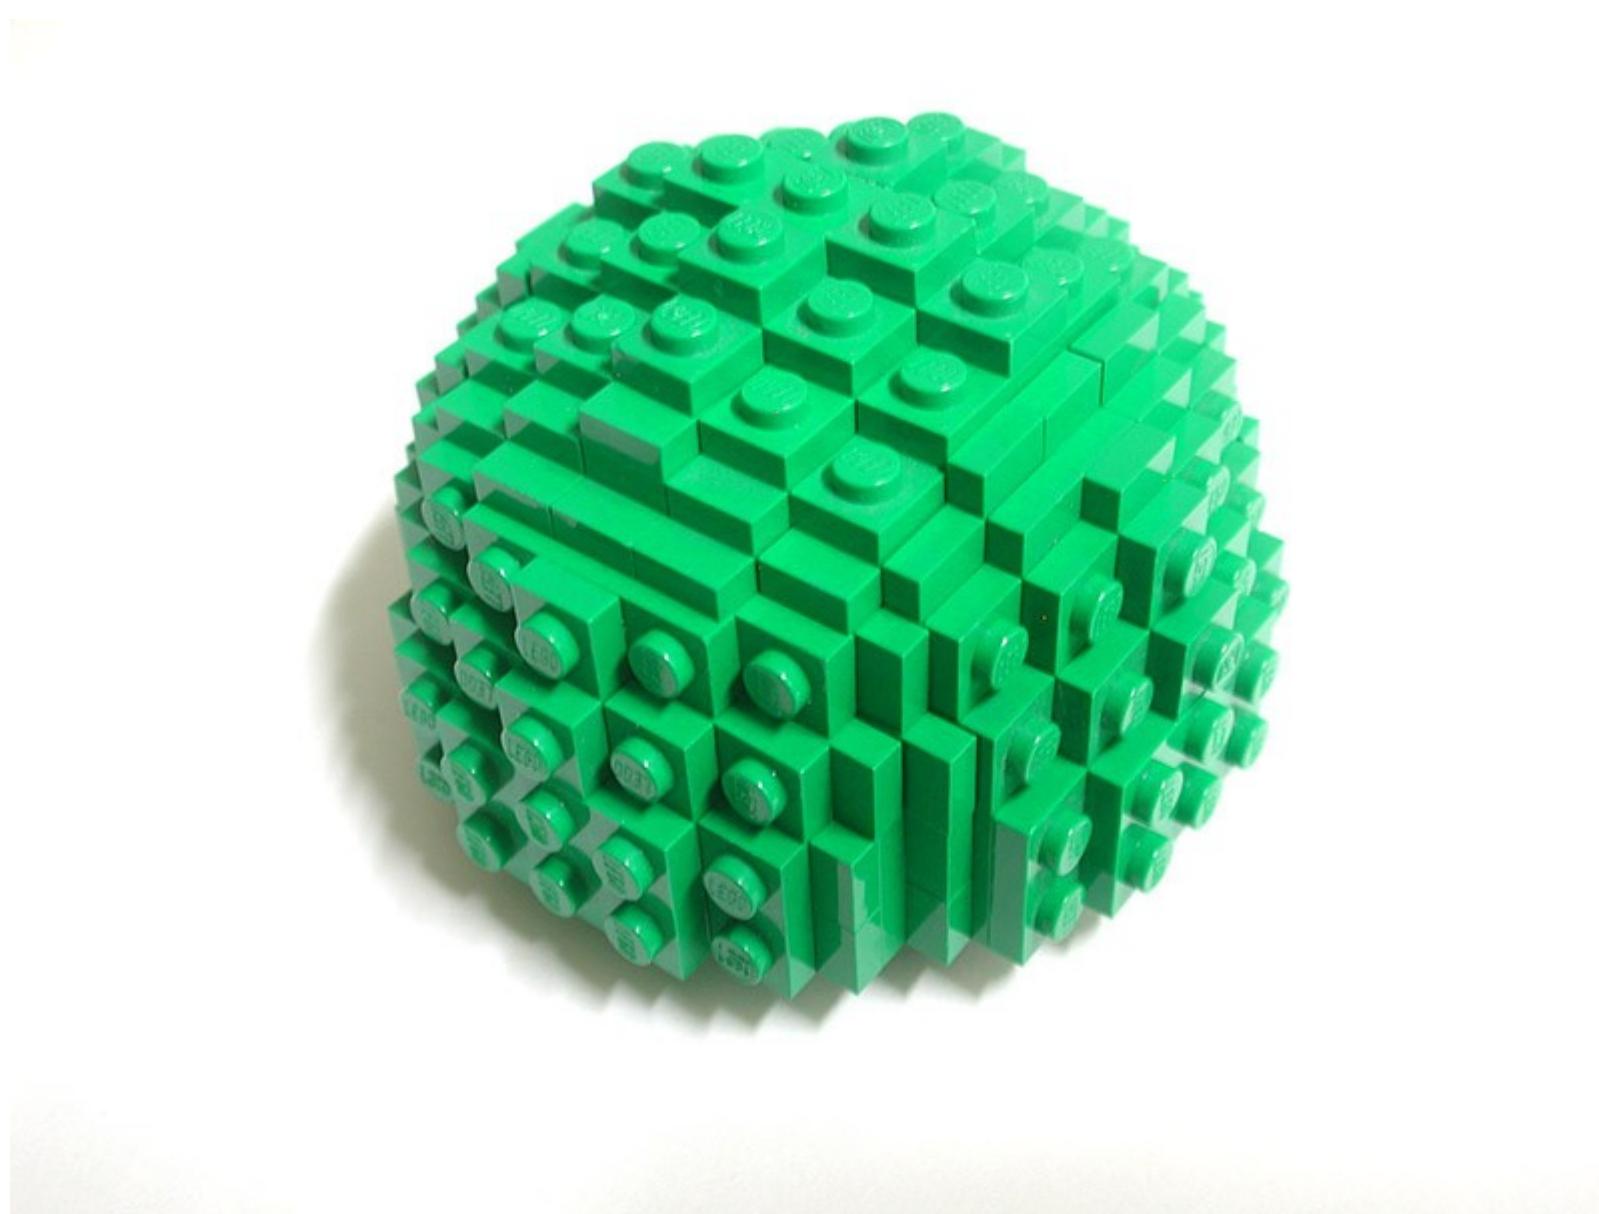


Image by Bram <https://flic.kr/p/8XKz1r>



elastic



Image by Kārlis Dambrāns <https://flic.kr/p/kTYy5V>





carrotsearch randomized testing

Image by David <https://flic.kr/p/4ubXx>



Input data

numbers

strings

specific  
charset only

Environment

change  
Locale



elastic

## Input data

numbers

strings

specific  
charset only

## Environment

change  
Locale

Reproducible due to specific seed.

Repeating tests runs them on different data.

Support for test timeouts.

Detects leaking threads.

Annotate tests to e.g. run only nightly.



elastic

Checking result cheaper than  
computing it.

Use slower but simpler algorithm  
for comparison.



Cheap upper/lower bound  
on result.



```
01 @Test
02 @Repeat (iterations = 100)
03 public void testSorting {
04     int length = randomIntBetween(10, 100);
05     short[] list = new short[length];
06     for (int i = 0; i < length; i++) {
07         list[i] = randomShort();
08     }
09     short[] result = Arrays.sort(list);
10     assertTrue(isSorted(result));
11 }
```



elastic

```

@ThreadLeakScope(Scope.TEST)
@ThreadLeakAction({Action.WARN, Action.INTERRUPT})
@ThreadLeakLinger(linger = 1000)
public static class TheGoodBadAndUgly extends RandomizedTest {
    @Test
    public void good() {
        // I do nothing and I'm good.
    }

    @Test
    @Repeat(iterations = 10)
    public void bad() {
        // I fail randomly, about 20% of the time.
        assertFalse(randomIntBetween(1, 100) <= 20);
    }

    @Test
    public void ugly() {
        // I start and abandon a thread which falls
        // outside the test's scope. The test will fail.
        new Thread() {
            public void run() {
                RandomizedTest.sleep(5000);
            }
        }.start();
    }
}

@Seed("AF567B2B9F8A8F1C")
public class Test005RecoveringRandomSeed
{
    @BeforeClass
    public static void printMasterContext()
    {
        System.out.println("# Static context");
        System.out.println(getContext().getR

```



elastic

The screenshot shows the Eclipse IDE interface with the JUnit perspective selected. The top bar includes tabs for JUnit, Problems, and TestNG, along with various toolbars and icons. A message at the top states "Finished after 11.66 seconds". Below this, summary statistics are displayed: "Runs: 12/12", "Errors: 1", and "Failures: 1".

The main view displays a tree structure of test cases under the package `com.carrotsearch.randomizedtesting.snippets`. The tree includes nodes for "good (0.001 s)", "bad (0.007 s)", and multiple "bad" nodes with specific seed values (#0 to #9). One "bad" node is expanded to show its failure trace.

The "Failure Trace" panel on the right shows the stack trace for the failed test:

```
java.lang.AssertionError
at __randomizedtesting.SeedInfo.seed([6C6E22495E089CC9:D9F...])
at org.junit.Assert.fail(Assert.java:92)
at org.junit.Assert.assertTrue(Assert.java:43)
at org.junit.Assert.assertFalse(Assert.java:68)
at org.junit.Assert.assertFalse(Assert.java:79)
at com.carrotsearch.randomizedtesting.snippets...
at sun.reflect.NativeMethodAccessorImpl.invoke...
at sun.reflect.NativeMethodAccessorImpl.invoke...
at sun.reflect.DelegatingMethodAccessorImpl.invoke...
at java.lang.reflect.Method.invoke(Method.java:5...
at com.carrotsearch.randomizedtesting.Randomi...
at com.carrotsearch.randomizedtesting.Randomi...
```



elastic

REPRODUCE WITH: gradle :core:integTest -Dtests.seed=6CB9C3FB4D9E4867  
-Dtests.class=org.elasticsearch.indices.mapping.UpdateMappingIntegrationIT  
-Dtests.method="testDynamicUpdates"  
-Des.logger.level=WARN  
-Dtests.security.manager=true  
-Dtests.locale=es\_CL  
-Dtests.timezone=America/Jamaica



## On the unit test level:

Random input w/ fixed seed to increase search space.

In CI re-running tests covers adds to coverage increase.



ORGANISATION EUROPÉENNE POUR LA RECHERCHE  
**CERN** EUROPEAN ORGANIZATION FOR NUCLEAR

1211 GENÈVE 23 (SUISSE)

This machine is a server  
DO NOT POWER DOWN!!



Image by Robert Scoble - <https://flic.kr/p/4qZaVM>  
Tag on Tim Berners-Lee's original NeXT machine -- first Web server

For distributed systems:

Vary the number of nodes.

Vary the number of processes on a node.



For systems running on the JVM:

- Vary JVM optimisation parameters.

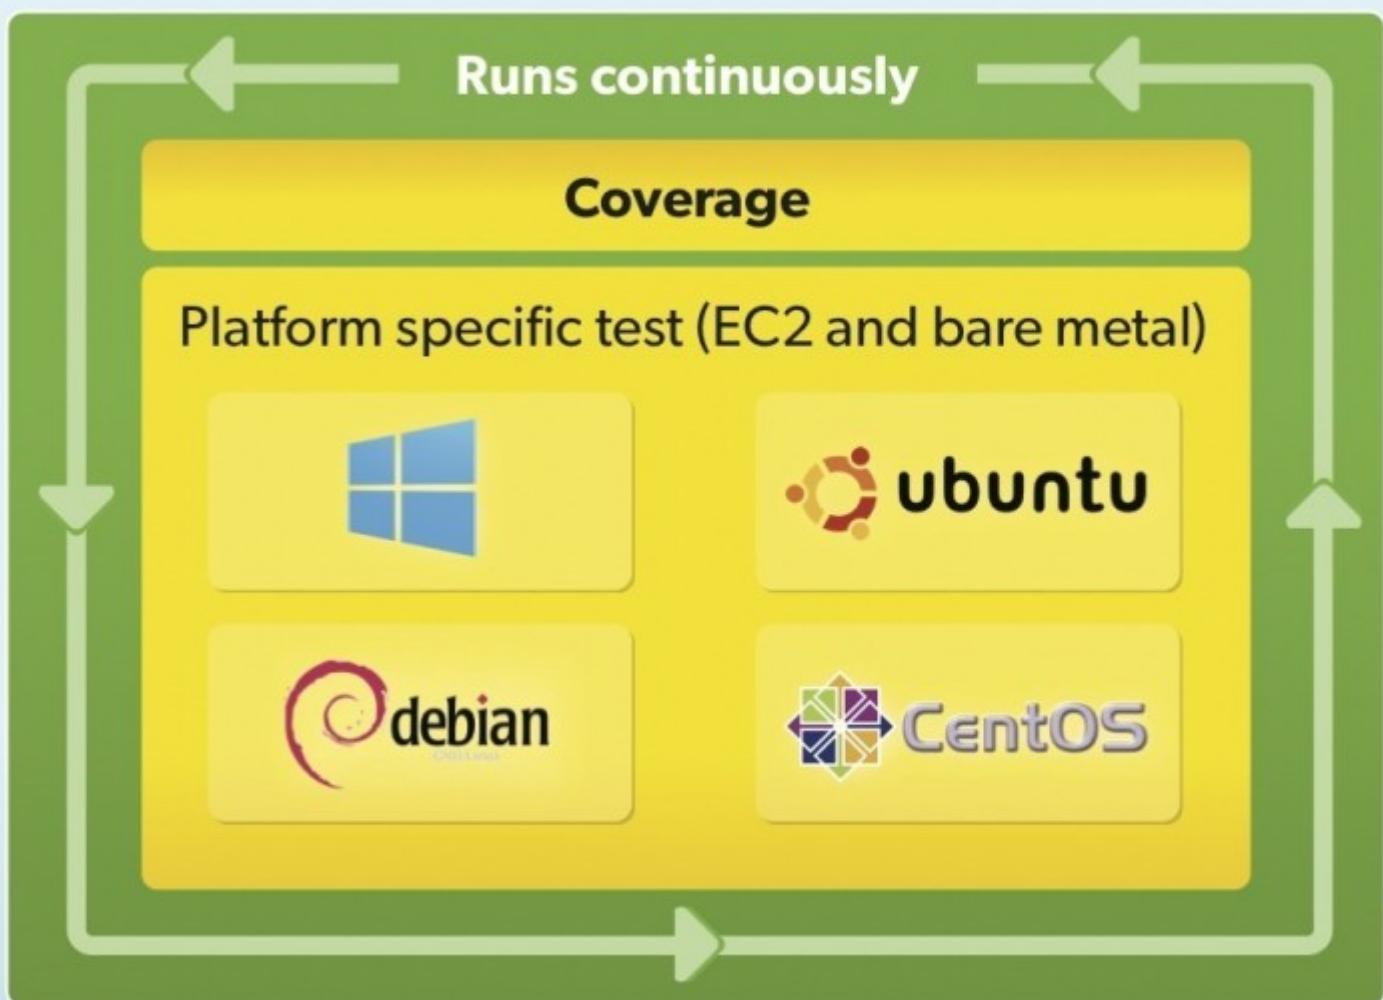
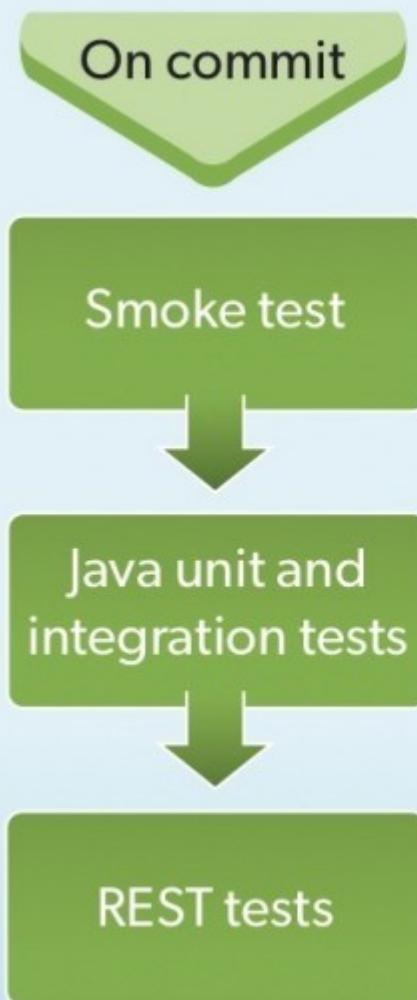
- Vary the JVM distribution used for testing.

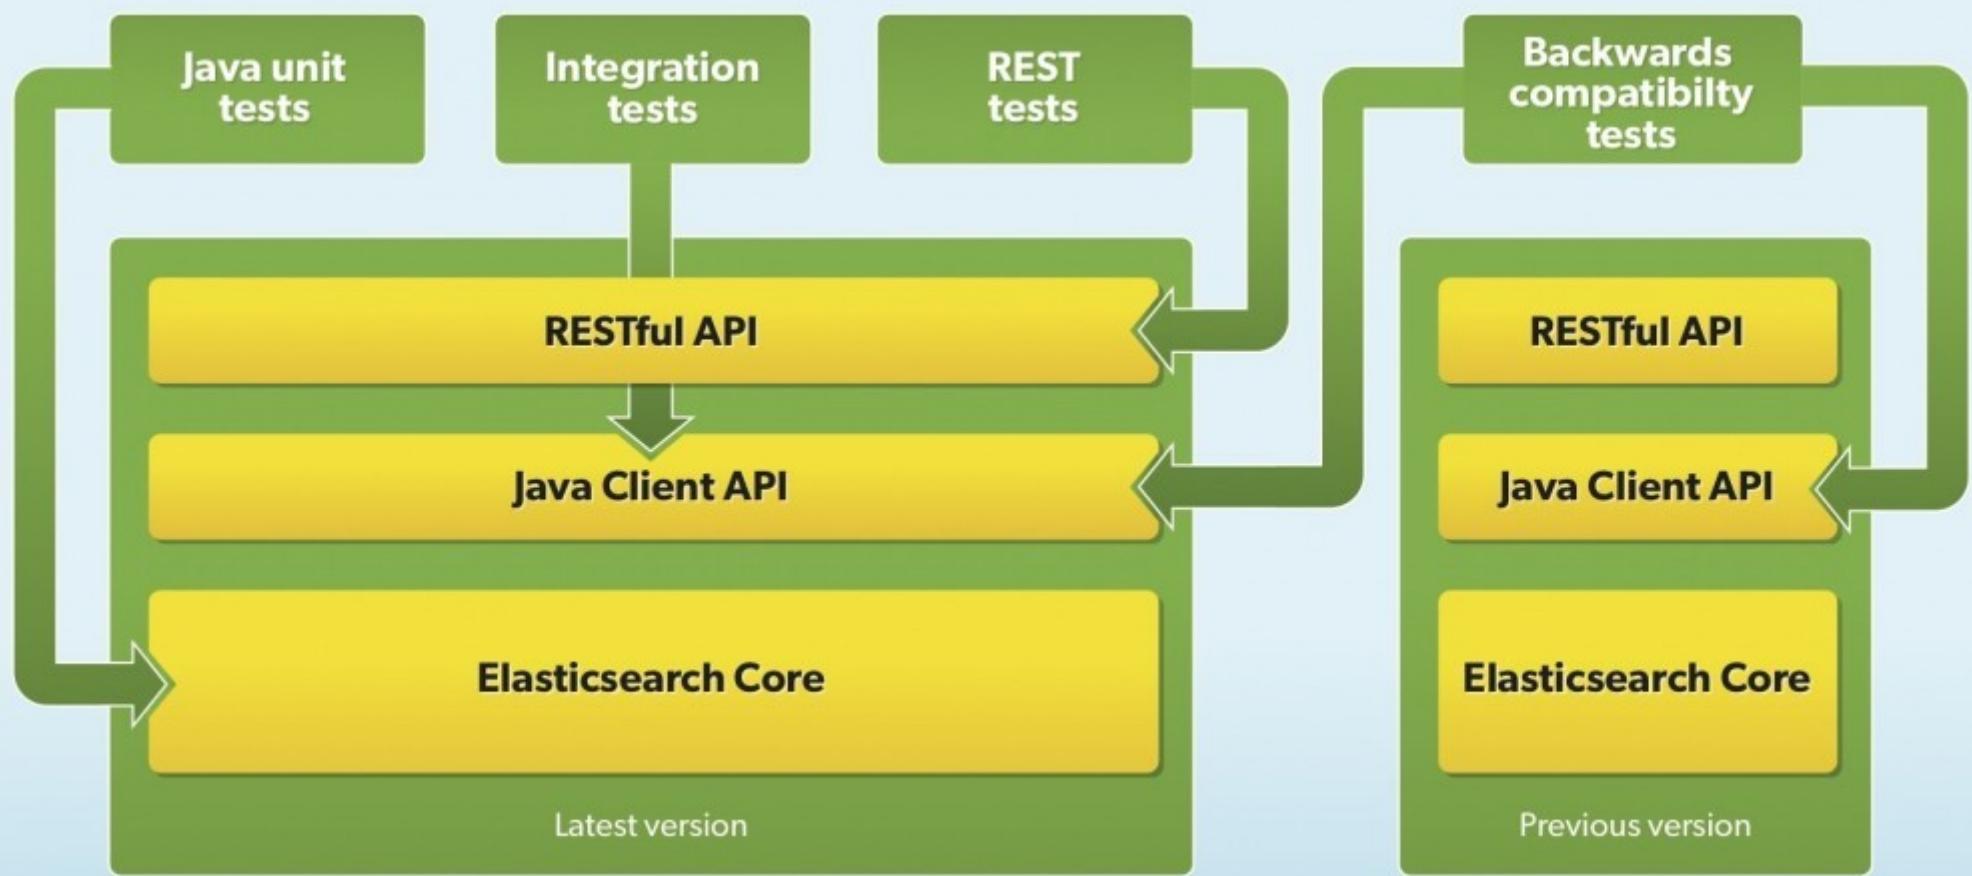




Copyright Elasticsearch 2015.

Image of Herbert taken 2013 by Thilo Fromm







- #2323 [Dec 13, 2014 4:35:08 PM](#)
- #2322 [Dec 13, 2014 8:22:57 AM](#)
- #2321 [Dec 13, 2014 4:23:59 AM](#)
- #2320 [Dec 13, 2014 12:23:57 AM](#)
- #2319 [Dec 12, 2014 9:22:30 PM](#) ▾
- #2318 [Dec 12, 2014 4:31:24 PM](#)
- #2317 [Dec 12, 2014 8:22:57 AM](#)
- #2316 [Dec 12, 2014 4:22:59 AM](#)
- #2315 [Dec 12, 2014 12:22:57 AM](#)
- #2314 [Dec 11, 2014 8:22:57 PM](#)
- #2313 [Dec 11, 2014 4:23:57 PM](#)
- #2312 [Dec 11, 2014 12:23:58 PM](#)
- #2311 [Dec 11, 2014 8:23:58 AM](#)

# Who broke the build?



## Oracle Java / Sun Java / OpenJDK Bugs

If you are affected by one of these issues that Oracle's Java has yet to accept or resolve

Oracle bug	Lucene mirror issue	Impact to Lucene
<a href="#">6265734</a>	<a href="#">LUCENE-573</a>	NIOFSDirectory has very poor performance
<a href="#">6478546</a>	<a href="#">LUCENE-1566</a>	You hit a false OutOfMemoryException when
<a href="#">6707044</a>	<a href="#">LUCENE-1282</a>	Index corruption
<a href="#">4724038</a>	<a href="#">LUCENE-1658</a> ( <a href="#">LUCENE-1669</a> )	MMapDirectory won't close files until GC (high bytes when used for read/write access in re
(No Oracle bug yet)	<a href="#">LUCENE-1342</a>	SEGV during indexing, with Java 1.6 64 bit
(No Oracle bug yet)	Discussed <a href="#">here</a> and <a href="#">here</a>	On 64 bit JREs, reading from files may han
<a href="#">6588260</a>	<a href="#">LUCENE-2449</a>	No impact as of 3.1/4.0 -- we now avoid cal
<a href="#">6240963</a>	<a href="#">LUCENE-2685</a>	Fix Delivered, High
<a href="#">5091921</a>	<a href="#">LUCENE-2975</a>	readVInt() returns wrong results
<a href="#">7044738</a>	<a href="#">LUCENE-3346</a>	readVInt() returns wrong results
.....	<a href="#">IIICFNF-3335</a>	Porter Stemmer crashes JRE

Uncover bugs in your environment.





# GOTOber - where's your testing pain?



elastic