

Use at Your Own Risk: The Java Unsafe API in the Wild

Luis Mastrangelo, Luca Ponzanelli, Andrea Mocci,
Michele Lanza, Matthias Hauswirth, Nathaniel Nystrom

Faculty of Informatics
Università della Svizzera Italiana (USI)
Switzerland

Java is a safe language.

```
$ java ch.usi.inf.MySafeJavaProgram
```

```
$ java ch.usi.inf.MySafeJavaProgram
#
# A fatal error has been detected by the Java Runtime Environment:
#
# SIGSEGV (0xb) at pc=0x000000011170bf04, pid=43948, tid=3591
#
# JRE version: Java(TM) SE Runtime Environment (7.0_79-b15) (build 1
# Java VM: Java HotSpot(TM) 64-Bit Server VM (24.79-b02 mixed mode b
# Problematic frame:
# V [libjvm.dylib+0x50bf04] Unsafe_GetNativeLong+0x4c
```

```
$ java ch.usi.inf.MySafeJavaProgram
#
# A fatal error has been detected by the Java Runtime Environment:
#
# SIGSEGV (0xb) at pc=0x000000011170bf04, pid=43948, tid=3591
#
# JRE version: Java(TM) SE Runtime Environment (7.0_79-b15) (build 1
# Java VM: Java HotSpot(TM) 64-Bit Server VM (24.79-b02 mixed mode b
# Problematic frame:
# V [libjvm.dylib+0x50bf04] Unsafe_GetNativeLong+0x4c
```

What is
sun.misc.Unsafe?

Why Developers Should Not Write Programs That Call 'sun' Packages

The classes that JavaSoft includes with the JDK fall into at least two packages: `java.*` and `sun.*`. Only classes in `java.*` packages are a standard part of the Java Platform and will be supported into the future. In general, API outside of `java.*` can change at any time without notice, and so cannot be counted on either across OS platforms (Sun, Microsoft, Netscape, Apple, etc.) or across Java versions. Programs that contain direct calls to the `sun.*` API are not 100% Pure Java. In other words:

The `java.*` packages make up the official, supported, public Java interface.

If a Java program directly calls only API in `java.*` packages, it will operate on all Java-compatible platforms, regardless of the underlying OS platform.

The `sun.*` packages are *not* part of the supported, public Java interface.

A Java program that directly calls any API in `sun.*` packages is *not* guaranteed to work on all Java-compatible platforms. In fact, such a program is not guaranteed to work even in future versions on the same platform.

For these reasons, there is no documentation available for the `sun.*` classes. Platform-independence is one of the great advantages of developing in Java. Furthermore, JavaSoft, and our licensees of Java technology, are committed to maintaining the APIs in `java.*` for future versions of the Java platform. (Except for code that relies on bugs that we later fix, or APIs that we deprecate and eventually remove.) This means that once your program is written, the binary will work in future releases. That is, future implementations of the Java platform will be backward compatible.

Each company that implements the Java platform will do so in their own private way. The classes in `sun.*` are present in the JDK to support the JavaSoft implementation of the Java platform: the `sun.*` classes are what make the classes in `java.*` work "under the covers" for the JavaSoft JDK. These classes will not in general be present on another vendor's Java platform. If your Java program asks for a class `sun.package.Foo` by name, it will likely fail with `ClassNotFoundException`, and you will have lost a major advantage of developing in Java.

Technically, nothing prevents your program from calling API in `sun.*` by name, but these classes are unsupported APIs, and we are not committed to maintaining backward compatibility for them. From one release to another, these classes may be removed, or they may be moved from one package to another, and it's fairly likely that the API (method names and signatures) will change. (From the JavaSoft point of view, since we are committed to maintaining the `java.*` APIs, we need to be able to change `sun.*` to enhance our products.) In this case, even if you are willing to run only on the JavaSoft implementation, you run the risk of a new version of the implementation breaking your program.

In general, writing Java programs that rely on `sun.*` is risky: they are not portable, and the APIs are not supported.

Unsafe is hidden

Unsafe is hidden

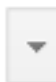
```
1 Constructor<Unsafe> c = Unsafe.class.  
    getDeclaredConstructor();  
2 c.setAccessible(true);  
3 Unsafe unsafe = c.newInstance();
```

JVM crash on 64 bit SPARC with Elasticsearch 1.2.2 due to unaligned memory access

26 posts by 6 authors  



David Roberts

22/07/2014 

★ Hello,

After upgrading from Elasticsearch 1.0.1 to 1.2.2 I'm getting JVM core dumps on Solaris 10 on SPARC.

```
# A fatal error has been detected by the Java Runtime Environment:
#
# SIGBUS (0xa) at pc=0xffffffff7e452d78, pid=15483, tid=263
#
# JRE version: Java(TM) SE Runtime Environment (7.0_55-b13) (build 1.7.0_55-b13)
# Java VM: Java HotSpot(TM) Server VM (24.55-b03 mixed mode solaris-sparc compressed oops)
# Problematic frame(s):
# V [libjvm.so+0x7e452d78] Unsafe_GetLong+0x10
```

I'm pretty sure the problem here is that Elasticsearch is making increasing use of "unsafe" functions in Java, presumably to speed things up, and some CPUs are more picky than others about memory alignment. In particular, x86 will tolerate misaligned memory access whereas SPARC won't.

Somebody has tried to report this to Oracle in the past and (understandably) Oracle has said that if you're going to use unsafe functions you need to understand what you're doing: http://bugs.java.com/bugdatabase/view_bug.do?bug_id=8021574

A quick grep through the code of the two versions of Elasticsearch shows that the new use of "unsafe" memory access functions is in the BytesReference, MurmurHash3 and HyperLogLogPlusPlus classes:

```
bash-3.2$ git checkout v1.0.1
Checking out files: 100% (2904/2904), done.
```

```
bash-3.2$ find . -name '*.java' | xargs grep UnsafeUtils
./src/main/java/org/elasticsearch/common/util/UnsafeUtils.java:public enum UnsafeUtils {
./src/main/java/org/elasticsearch/search/aggregations/bucket/BytesRefHash.java:    if (id == -1L || UnsafeUtils.equals(key, get(id, spare))) {
./src/main/java/org/elasticsearch/search/aggregations/bucket/BytesRefHash.java:    } else if (UnsafeUtils.equals(key, get(curId, spare))) {
./src/test/java/org/elasticsearch/benchmark/common/util/BytesRefComparisonsBenchmark.java:import org.elasticsearch.common.util.UnsafeUtils;
./src/test/java/org/elasticsearch/benchmark/common/util/BytesRefComparisonsBenchmark.java:    return UnsafeUtils.equals(b1, b2);
```



David Roberts

22/07/2014

★ Hello,

After upgrading from Elasticsearch 1.0.1 to 1.2.2 I'm getting JVM core dumps on Solaris 10 on SPARC.

```
# A fatal error has been detected
# SIGBUS (0xa) at pc=0xfffffff
# JRE version: Java(TM) SE R
# Java VM: Java HotSpot(TM)
# Problematic frame(s):
# V [libjvm.so+0x10052d78] Uns
```

I'm pretty sure the problem is... picky than others about memor

Somebody has tried to report th... you're doing: <http://bugs.java.o>

A quick grep through the code... and HyperLogLogPlusPlus cla

```
bash-3.2$ git checkout v1.0.1
Checking out files: 100% (2904
```

```
bash-3.2$ find . -name '*.java' |
./src/main/java/org/elasticsearch
./src/main/java/org/elasticsearch
./src/main/java/org/elasticsearch
./src/test/java/org/elasticsearch
./src/test/java/org/elasticsearch
```

Hazelcast cust serialization with kryo fail in jdk 64bit in solaris-sparc #219



bwzhang2011 opened this issue on May 14, 2014 · 11 comments



bwzhang2011 commented on May 14, 2014

please take a look at [hazelcast/hazelcast#2453](#)

the kryo serialization seems fail in solaris sparc and I don't know whether kryo test cases has been completely tested in jdk solaris edition. and I hope someone could help me point out my mistake or something wrong with kryo itself for different o.s



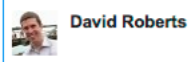
romix commented on May 14, 2014

Collaborator

It seems like it has problems using Unsafe-backend on Solaris. But you can easily switch to **asm-based** one, using

```
kryo.setAsmEnabled(true);
```

ASM-based backend should work on any platform.



David Roberts

Hello,

After upgrading from Elasticsearch 1.0.1 to 1.2.2 I'm getting

```
# A fatal error has been detected by the Java Runtime Envir
#
# SIGBUS (0xa) at pc=0xfffffff7e452d78, pid=15483, tid=26
#
# JRE version: Java(TM) SE Runtime Environment (7.0_55-
# Java VM: Java HotSpot(TM) 64-Bit Server VM (24.55-b03
# Problematic frame(s):
# V [libjvm.so+0x152d78] Unsafe_GetLong+0x1
```

I'm pretty sure the problem is that Elasticsearch is making me picky than others about memory alignment. In particular, x8

Somebody has tried to report this to Oracle in the past and (you're doing: http://bugs.java.com/bugdatabase/view_bug.do

A quick grep through the code of the two versions of Elasticsearch and HyperLogLogPlusPlus classes:

```
bash-3.2$ git checkout v1.0.1
Checking out files: 100% (2904/2904), done.
```

```
bash-3.2$ find . -name '*.java' | xargs grep UnsafeUtils
./src/main/java/org/elasticsearch/common/util/UnsafeUtils.java
./src/main/java/org/elasticsearch/search/aggregations/bucket
./src/main/java/org/elasticsearch/search/aggregations/bucket
./src/test/java/org/elasticsearch/benchmark/common/util/Byte
./src/test/java/org/elasticsearch/benchmark/common/util/Byte
```

Snappy causes SIGBUS jvm crasher on HP-UX via UnsafeMemory.loadInt() #24

ekeithkw opened this issue on Jan 14 · 1 comment



ekeithkw commented on Jan 14

Steps to reproduce:

1. Install HP-UX Java 7 (e.g. from "Itanium_JDK_JRE_7.0.09_-_Feb_2014_Z7550-01326_java_17009_ia.depot") to a local folder.
2. Set JDK_HOME to the above folder.
3. Download CLion for Linux, from jetbrains.com > Products > CLion > Early Access Program (EAP); gunzip and "tar xvf" to unpack (e.g. to /opt).
4. Run "/opt/clion-140.569.17/bin/clion.sh". Deselect all plug-ins.
5. It appears necessary to create or open a project (failure occurs when opening a project).

Snippet from hs_err.log file:

Current thread (01370200): JavaThread *ApplicationImpl pooled thread 3" [_thread_in_vm, id=30, lwp_id=4383651, stack(2d601000,2d641000)]

```

signinfo:si_signo=SIGBUS: si_errno=0, si_code=1 (BUS_ADRALN), si_addr=2f4b284d
...
j sun.misc.Unsafe.getInt(Ljava/lang/Object;)I+0
j org.iq80.snappy.SnappyInternalUtils.loadInt([BI)I+50
j org.iq80.snappy.SnappyInternalUtils.loadInt([BI)I+5
j org.iq80.snappy.SnappyCompressor.findCandidate([BIIII[SI][I+19
j org.iq80.snappy.SnappyCompressor.compressFragment([BII[BI[S]I+157
j org.iq80.snappy.SnappyCompressor.compress([BII[BI]I+63
j org.iq80.snappy.Snappy.compress([BII[BI]I+6
j org.iq80.snappy.SnappyOutputStream.writeCompressed([BII)V+27
j org.iq80.snappy.SnappyOutputStream.flushBuffer()V+17
j org.iq80.snappy.SnappyOutputStream.flush()V+18
j sun.nio.cs.StreamEncoder.implFlush()V+15
j sun.nio.cs.StreamEncoder.flush()V+12
j java.io.OutputStreamWriter.flush()V+4
j org.jdom.output.XMLOutputter.output(Lorg/jdom/Element;Ljava/io/Writer;)V+12
j com.intellij.openapi.components.impl.stores.StateMap.a(Lorg/jdom/Element;)[B+85
j com.intellij.openapi.components.impl.stores.StateMap.getStateAndArchive(Ljava/lang/String;)Lorg/j

```

1 comment

on't know whether kryo test cases has been someone could help me point out my mistake or

Collaborator

Solaris. used one, using



David Roberts

Hello,

After upgrading from Elasticsearch 1.0.1 to 1.2.2 I'm getting:

```
# A fatal error has been detected by the Java Runtime Env
#
# SIGBUS (0xa) at pc=0xfffffff7e452d78, pid=15483, tid=
#
# JRE version: Java(TM) SE Runtime Environment (7.0_51)
# Java VM: Java HotSpot(TM) 64-Bit Server VM (24.55-b0)
# Problematic frame:
# V [libjvm.so+0x1000000000000000] Unsafe_GetLong+0x1000000000000000
```

I'm pretty sure the problem is that Elasticsearch is more picky than others about memory alignment. In particular, x

Somebody has tried to report this to Oracle in the past and you're doing: http://bugs.java.com/bugdatabase/view_bug.do?bug_id=7148283

A quick grep through the code of the two versions of Elastic and HyperLogLogPlusPlus classes:

```
bash-3.2$ git checkout v1.0.1
Checking out files: 100% (2904/2904), done.
```

```
bash-3.2$ find . -name '*.java' | xargs grep UnsafeUtils
./src/main/java/org/elasticsearch/common/util/UnsafeUtils.java
./src/main/java/org/elasticsearch/search/aggregations/bucket/
./src/main/java/org/elasticsearch/search/aggregations/bucket/
./src/test/java/org/elasticsearch/benchmark/common/util/By
./src/test/java/org/elasticsearch/benchmark/common/util/By
```

Snappy causes SIGBUS j UnsafeMemory.loadInt()

Open ekeithkw opened this issue on Jan 14 · 1



ekeithkw commented on Jan 14

Steps to reproduce:

1. Install HP-UX Java 7 (e.g. from "Itanium 01326_java_17009_ia.depot") to a local folder.
2. Set JDK_HOME to the above folder.
3. Download CLion for Linux, from jetbrat gunzip and "tar xvf" to unpack (e.g. to /opt).
4. Run "/opt/clion-140.569.17/bin/clion.sh
5. It appears necessary to create or open

Snippet from hs_err.log file:

```
Current thread (01370200): JavaThread "IWP"
lwp_id=4383651, stack(2d601000,2d641000)
```

```
siginfo: si_signo=SIGBUS: si_errno=0, si_code=SEGV_MAPERR
at sun.misc.Unsafe.defineClass(
at sun.misc.Unsafe.defineClass(
at sun.reflect.MethodAccessorGenerator$1.run(MethodAccessorGenerator.java:399)
at sun.reflect.MethodAccessorGenerator$1.run(MethodAccessorGenerator.java:396)
at java.security.AccessController.doPrivileged(AccessController.java:0)
at sun.reflect.MethodAccessorGenerator.generate(MethodAccessorGenerator.java:395)
```

Bug 229655 - StackOverflowError at sun.misc.Unsafe.defineClass

Status: RESOLVED FIXED

Product: debugger
 Component: Java
 Version: 7.4
 Hardware: All All

Priority: P3 (vote)
 Target Milestone: 7.4

Assigned To: Martin Entlicher
 QA Contact: issues@debugger

URL:
 Whiteboard: EXCEPTIONS_REPORT
 Keywords:

Depends on:
 Blocks:

Show dependency tree / graph

Reported: 2013-05-13 12:56 UTC by Ondrej Vrabec

Modified: 2013-05-23 08:17 UTC (History)

CC List: 1 user (show)

See Also:

Issue Type: DEFECT

Exception Report : 20060

comments

n't know whether kryo test cases has been someone could help me point out my mistake or

Collaborator

Solaris. done, using

Ondrej Vrabec 2013-05-13 12:56:07 UTC

```
Build: NetBeans IDE Dev (Build 20130510-fb9d6c10dcaa)
VM: Java HotSpot(TM) 64-Bit Server VM, 23.6-b04, Java(TM) SE Runtime Environment, 1.7.0_10-b18
OS: Linux
```

User Comments:
 ovrabec: applying code changes

```
Stacktrace:
java.lang.StackOverflowError
at sun.misc.Unsafe.defineClass(
at sun.misc.Unsafe.defineClass(ClassDefiner.java:63)
at sun.reflect.MethodAccessorGenerator$1.run(MethodAccessorGenerator.java:399)
at sun.reflect.MethodAccessorGenerator$1.run(MethodAccessorGenerator.java:396)
at java.security.AccessController.doPrivileged(AccessController.java:0)
at sun.reflect.MethodAccessorGenerator.generate(MethodAccessorGenerator.java:395)
```


JVM crash on 64 bit SPARC with Elasticsearch 1.2.2 due to unaligned memory access

26 posts by 6 authors

David Roberts 22/07/2014

Hello,

After upgrading from Elasticsearch 1.0.1 to 1.2.2 I'm getting JVM core dumps on Solaris 10 on SPARC.

```
# A fatal error has been detected by the Java Runtime Environment:
#
# SIGBUS (0xa) at pc=0xfffffff7e452d78, pid=15483, tid=263
#
# JRE version: Java(TM) SE Runtime Environment (7.0_55-b13) (build 1.7.0_55-b13)
# Java VM: Java HotSpot(TM) 64-Bit Server VM (24.55-b03 mixed mode solaris-sparc compressed oops)
# Problematic frame(s):
# V [libjvm.so+0x10052d78] Unsafe_GetLong+0x10052d78
```

I'm pretty sure the problem is that Elasticsearch is making increasing use of "unsafe" functions in Java, presumably to speed things up, and some CPUs are more picky than others about memory alignment. In particular, x86 will tolerate misaligned memory access whereas SPARC won't.

Somebody has tried to report this to Oracle in the past and (understandably) Oracle has said that if you're going to use unsafe functions you need to understand what you're doing: http://bugs.java.com/bugdatabase/view_bug.do?bug_id=8021574

A quick grep through the code of the two versions of Elasticsearch shows that the new use of "unsafe" memory access functions is in the BytesReference, MurmurHash3 and HyperLogLogPlusPlus classes:

```
bash-3.2$ git checkout v1.0.1
Checking out files: 100% (2904/2904), done.

bash-3.2$ find . -name '*.java' | xargs grep UnsafeUtils
./src/main/java/org/elasticsearch/common/util/UnsafeUtils.java:public enum UnsafeUtils {
./src/main/java/org/elasticsearch/search/aggregations/bucket/BytesRefHash.java:    if (id == -1L || UnsafeUtils.equals(key, get(id, spare))) {
./src/main/java/org/elasticsearch/search/aggregations/bucket/BytesRefHash.java:    } else if (UnsafeUtils.equals(key, get(curlid, spare))) {
./src/test/java/org/elasticsearch/benchmark/common/util/BytesRefComparisonsBenchmark.java:import org.elasticsearch.common.util.UnsafeUtils;
./src/test/java/org/elasticsearch/benchmark/common/util/BytesRefComparisonsBenchmark.java:    return UnsafeUtils.equals(b1, b2);
```

Snappy causes SIGBUS jvm crasher on HP-UX via UnsafeMemory.loadInt() #24

Open ekeithkw opened this issue on Jan 14 · 1 comment

ekeithkw commented on Jan 14

Steps to reproduce:

1. Install HP-UX Java 7 (e.g. from "Itanium_JDK_JRE_7.0.09_-_Feb_2014_Z7550-01326_java_17009_ia.depot") to a local folder.
2. Set JDK_HOME to the above folder.
3. Download CLion for Linux, from jetbrains.com > Products > CLion > Early Access Program (EAP); gunzip and "tar xv" to unpack (e.g. to /opt).
4. Run "/opt/clion-140.569.17/bin/clion.sh". Deselect all plug-ins.
5. It appears necessary to create or open a project (failure occurs when opening a project).

Snippet from hs_err.log file:

```
Current thread (01370200): JavaThread "ApplicationImpl pooled thread 3" [_thread_in_vm, id=30, lwp_id=4383651, stack(2d601000,2d641000)]

siginfo:si_signo=SIGBUS: si_errno=0, si_code=1 (BUS_ADRALN), si_addr=2f4b284d
...
at sun.misc.Unsafe.getInt(Ljava/lang/Object;J)I+0
at org.elasticsearch.snappy.SnappyInternalUtils.loadInt([B]I+50
at org.elasticsearch.snappy.SnappyInternalUtils.loadInt([B]I+5
at org.iq80.snappy.SnappyCompressor.findCandidate([BIII[SI]I+19
at org.iq80.snappy.SnappyCompressor.compressFragment([BII[BI[SI]I+157
at org.iq80.snappy.SnappyCompressor.compress([BII[BI]I+63
at org.iq80.snappy.Snappy.compress([BII[BI]I+6
at org.iq80.snappy.SnappyOutputStream.writeCompressed([BII]V+27
at org.iq80.snappy.SnappyOutputStream.flushBuffer()V+17
at org.iq80.snappy.SnappyOutputStream.flush()V+18
at sun.nio.cs.StreamEncoder.implFlush()V+15
at sun.nio.cs.StreamEncoder.flush()V+12
at java.io.OutputStreamWriter.flush()V+4
at org.jdom.output.XMLOutputter.output(Lorg/jdom/Element;Ljava/io/Writer;)V+12
at com.intellij.openapi.components.impl.stores.StateMap.a(Lorg/jdom/Element;)[B+85
at com.intellij.openapi.components.impl.stores.StateMap.getStateAndArchive(Ljava/lang/String;)Lorg/j
```

Hazelcast cust serialization with kryo fail in jdk 64bit in solaris-sparc #219

Closed bwzhang2011 opened this issue on May 14, 2014 · 11 comments

bwzhang2011 commented on May 14, 2014

please take a look at [hazelcast/hazelcast#2453](#)

the kryo serialization seems fail in solaris sparc and I don't know whether kryo test cases has been completely tested in jdk solaris edition. and I hope someone could help me point out my mistake or something wrong with kryo itself for different o.s

romix commented on May 14, 2014 Collaborator

It seems like it has problems using Unsafe-backend on Solaris. But you can easily switch to ASM-based one, using

```
kryo.setAsmEnabled(true);
```

ASM-based backend should work on any platform.

Bug 229655 - StackOverflowError at sun.misc.Unsafe.defineClass

Status: RESOLVED FIXED

Product: debugger

Component: Java

Version: 7.4

Hardware: All All

Priority: P3 (vote)

Target Milestone: 7.4

Assigned To: Martin Entlicher

QA Contact: issues@debugger

URL: [Whiteboard: EXCEPTIONS_REPORT](#)

Keywords:

Depends on:

Blocks: [Show dependency tree / graph](#)

Reported: 2013-05-13 12:56 UTC by Ondrej Vrabec

Modified: 2013-05-23 08:17 UTC (History)

CC List: 1 user (show)

See Also:

Issue Type: DEFECT

Exception Report: [20060](#)

Ondrej Vrabec 2013-05-13 12:56:07 UTC

Build: NetBeans IDE Dev (Build 20130510-fb9d6c10dcaa)

VM: Java HotSpot(TM) 64-Bit Server VM, 23.6-b04, Java(TM) SE Runtime Environment, 1.7.0_10-b18

OS: Linux

User Comments:

ovrabec: applying code changes

Stacktrace:

```
java.lang.StackOverflowError
at sun.misc.Unsafe.defineClass0(Ljava/lang/ClassDefiner.java:0)
at sun.misc.Unsafe.defineClass(Ljava/lang/ClassDefiner.java:63)
at sun.reflect.MethodAccessorGenerator$1.run(MethodAccessorGenerator.java:399)
at sun.reflect.MethodAccessorGenerator$1.run(MethodAccessorGenerator.java:396)
at java.security.AccessController.doPrivileged(AccessController.java:0)
at sun.reflect.MethodAccessorGenerator.generate(MethodAccessorGenerator.java:395)
```

Research Question 1

Does *Unsafe* **impact**
common application
code?

959,300 artifacts

1.7 TB

959,300 artifacts

1.7 TB



Last version

959,300 artifacts

1.7 TB



Last version

106,574 artifacts

959,300 artifacts

1.7 TB



Last version

106,574 artifacts



.jar, .war, .ear, .ejb

959,300 artifacts

1.7 TB



Last version

106,574 artifacts



.jar, .war, .ear, .ejb

86,479 artifacts

74 GB

959,300 artifacts

1.7 TB



Last version

106,574 artifacts



.jar, .war, .ear, .ejb

86,479 artifacts

74 GB



959,300 artifacts

1.7 TB



Last version

106,574 artifacts



.jar, .war, .ear, .ejb

86,479 artifacts

74 GB



817 *Unsafe* artifacts (1%)

959,300 artifacts

1.7 TB



Last version

106,574 artifacts



.jar, .war, .ear, .ejb

86,479 artifacts

74 GB



817 *Unsafe* artifacts (1%)

48,139 call sites
351 field accesses

 The Central Repository



Languages



 The Central Repository



Languages



Frameworks



 The Central Repository



Languages



Frameworks



Serialization



 The Central Repository



Languages



Frameworks



Serialization



Database



 The Central Repository



Languages



Frameworks



Serialization



Database



Concurrency



 The Central Repository



Languages



Frameworks



Serialization



Database



Concurrency



Big Data



959,300 artifacts

1.7 TB



Last version

106,574 artifacts



.jar, .war, .ear, .ejb

86,479 artifacts

74 GB



817 *Unsafe* artifacts (1%)

48,139 call sites
351 field accesses

959,300 artifacts

1.7 TB



Last version

106,574 artifacts



.jar, .war, .ear, .ejb

86,479 artifacts

74 GB



48,139 call sites
351 field accesses

817 *Unsafe* artifacts (1%)

959,300 artifacts

1.7 TB



Last version

106,574 artifacts



.jar, .war, .ear, .ejb

86,479 artifacts

74 GB



48,139 call sites
351 field accesses

817 *Unsafe* artifacts (1%)

Artifact dependencies

959,300 artifacts

1.7 TB



Last version

106,574 artifacts



.jar, .war, .ear, .ejb

86,479 artifacts

74 GB



48,139 call sites
351 field accesses

817 *Unsafe* artifacts (1%)

Artifact dependencies



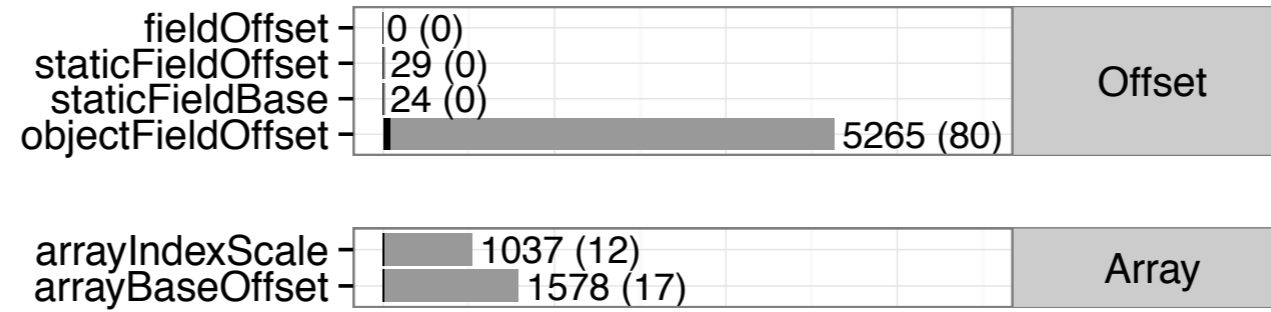
21,297 artifacts depend on *Unsafe* (25%)

Research Question 2

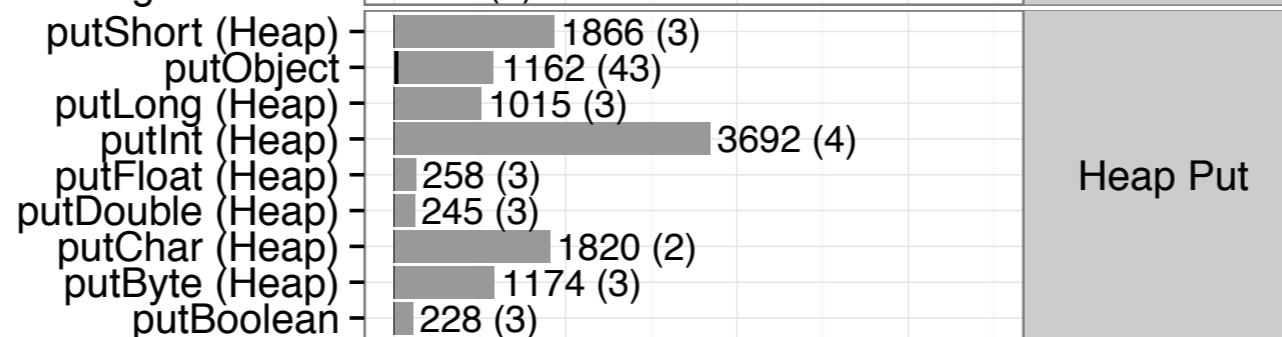
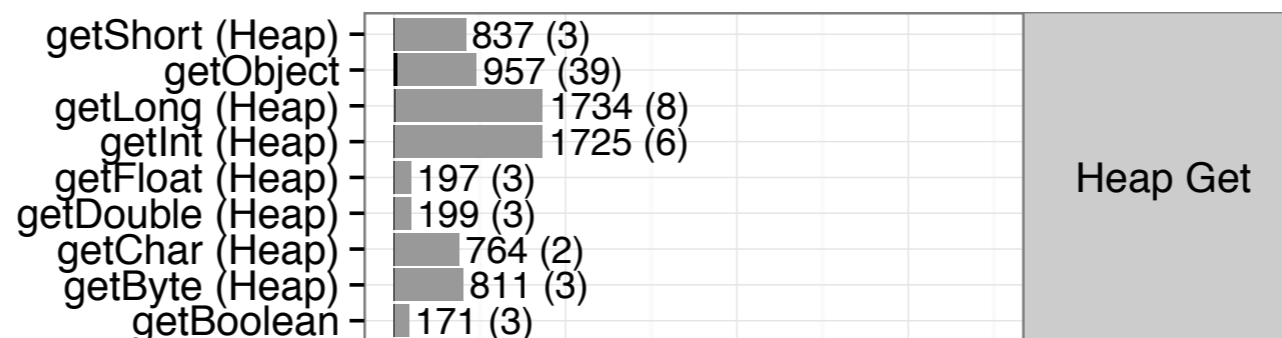
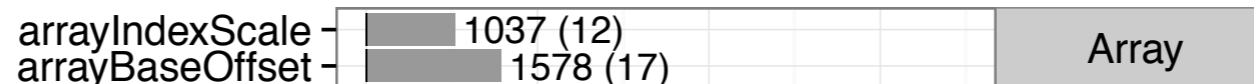
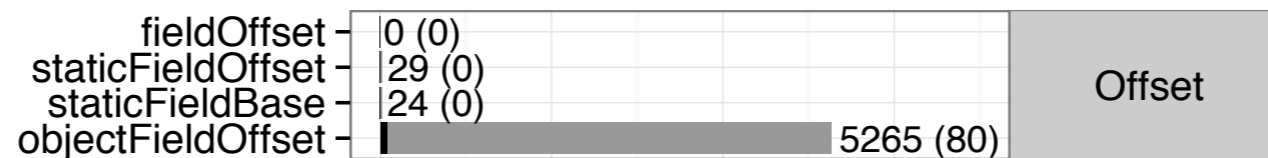
Which **features** of *Unsafe* are used?

Call Sites / sun.misc.Unsafe methods

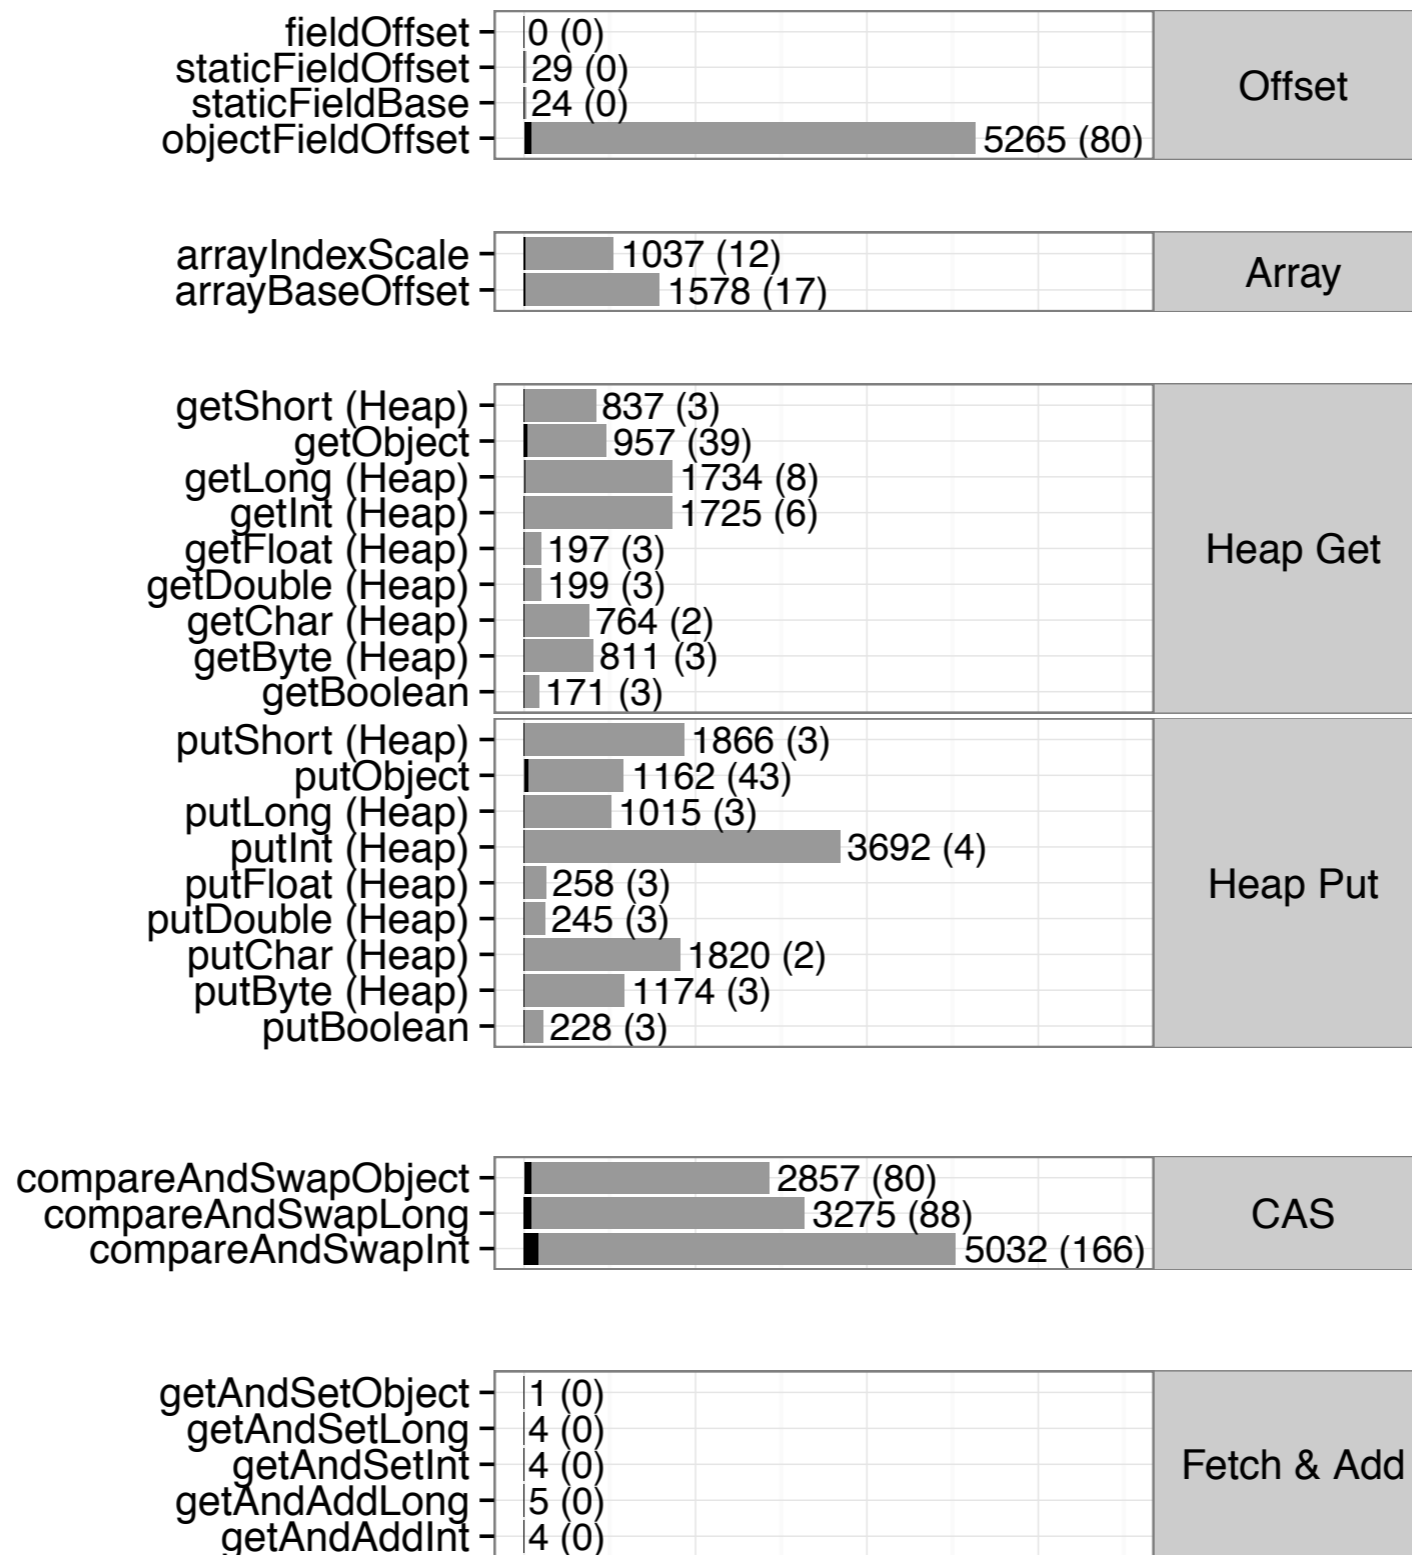
Call Sites / sun.misc.Unsafe methods



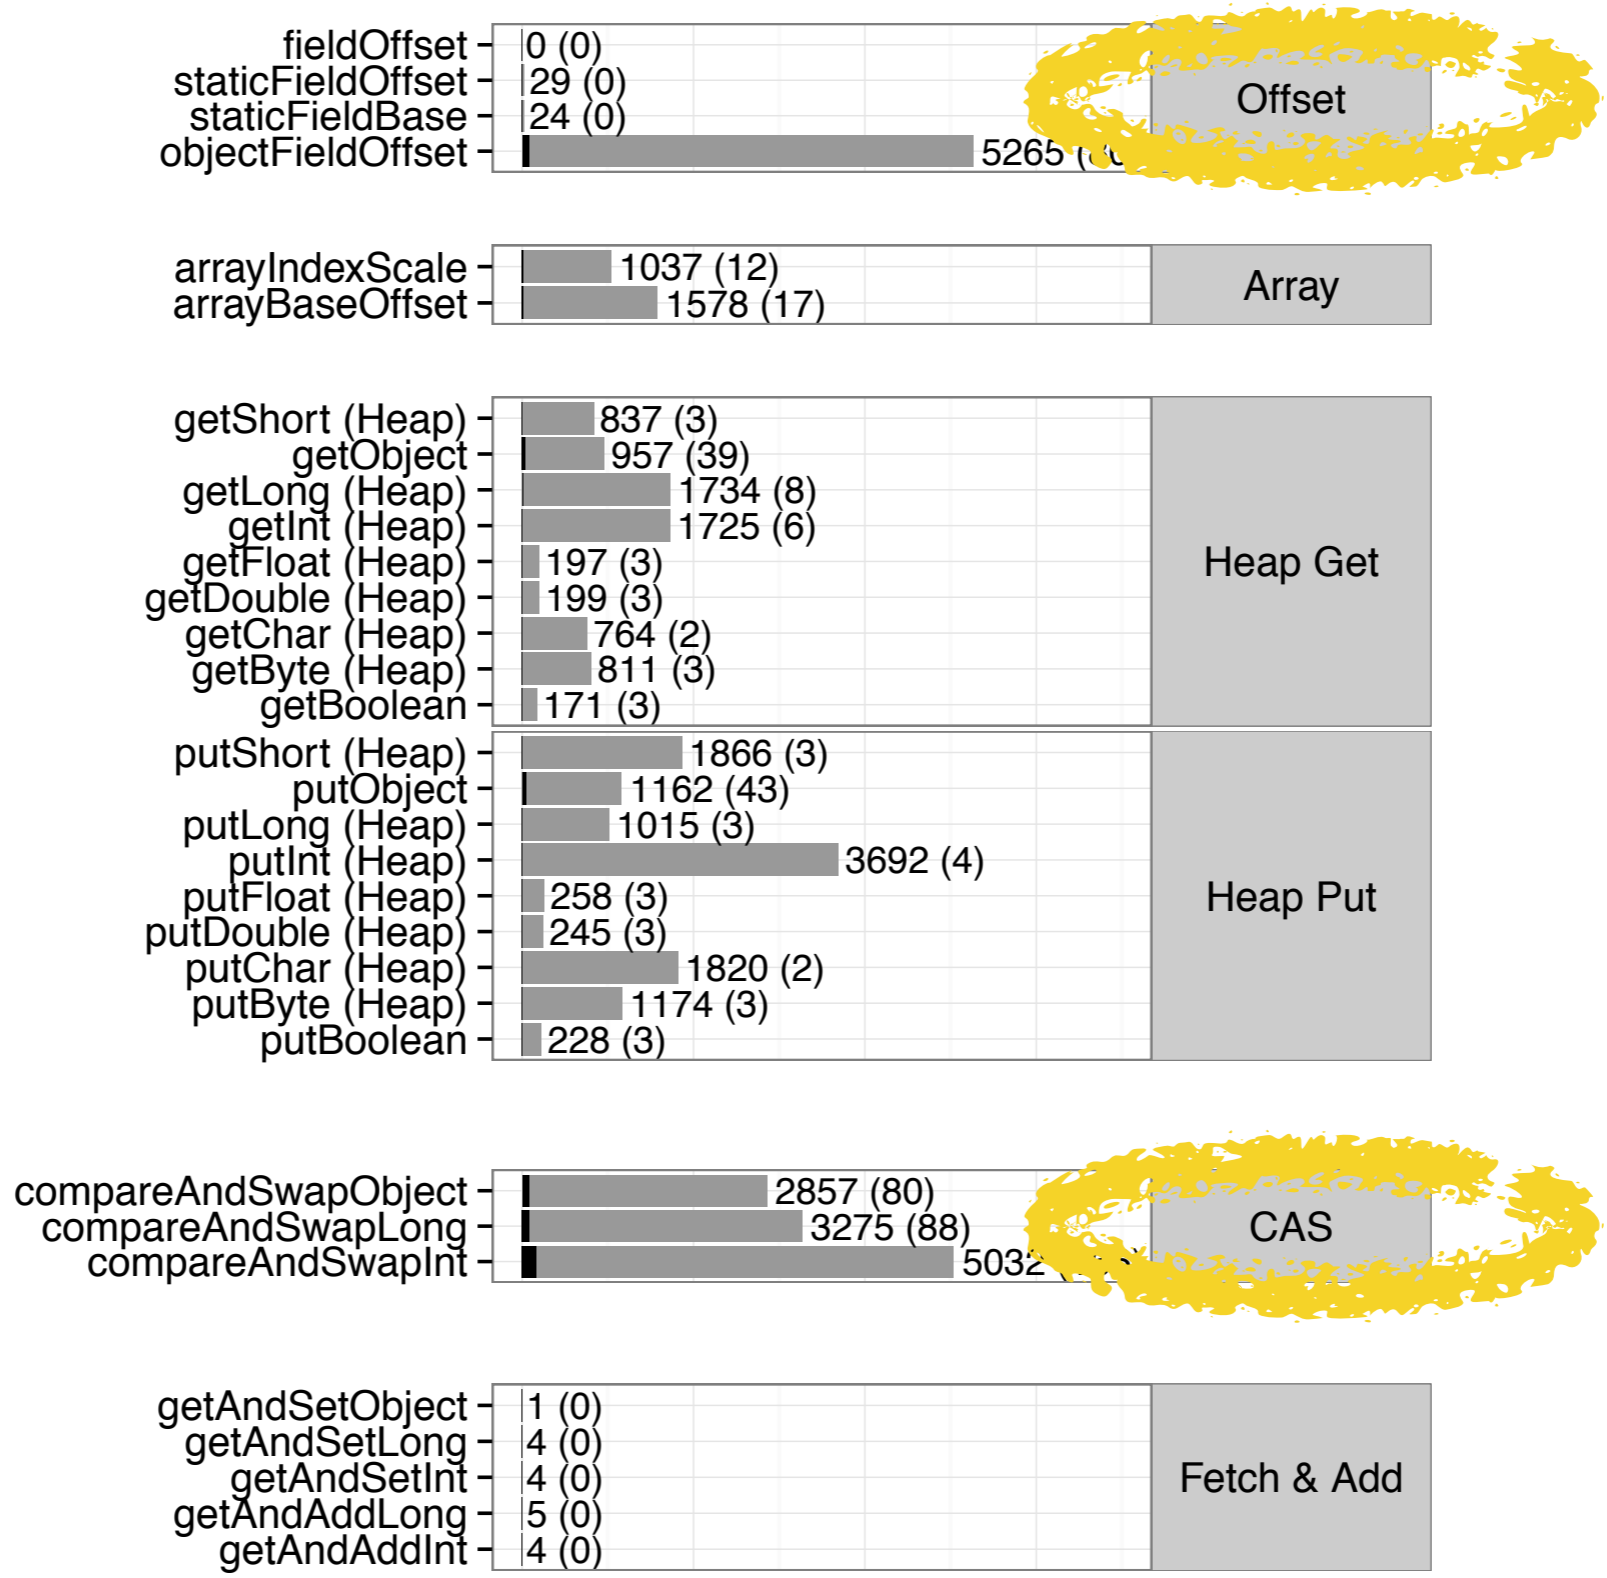
Call Sites / sun.misc.Unsafe methods



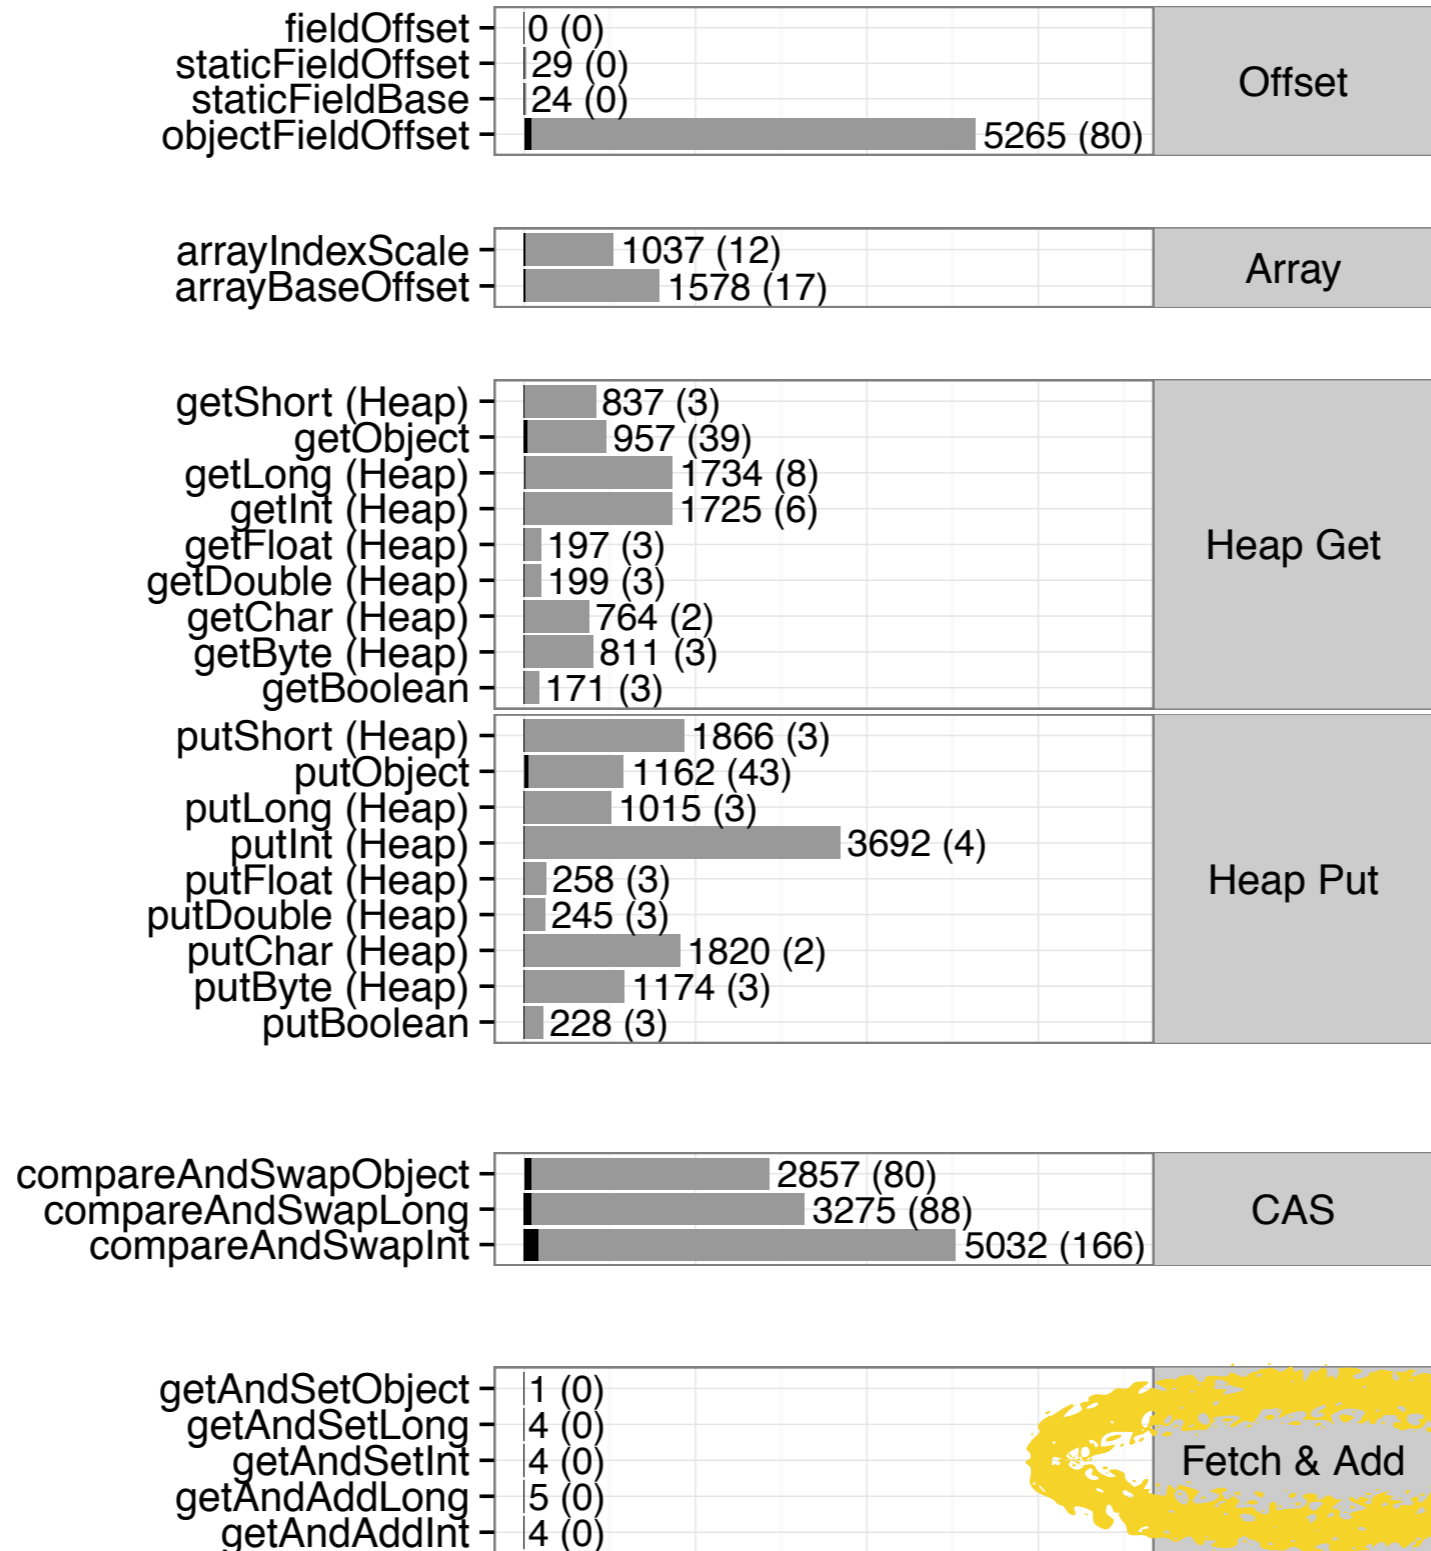
Call Sites / sun.misc.Unsafe methods



Call Sites / sun.misc.Unsafe methods

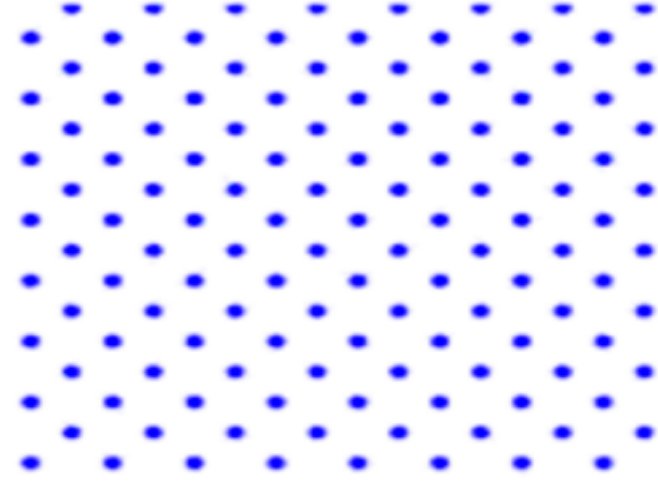
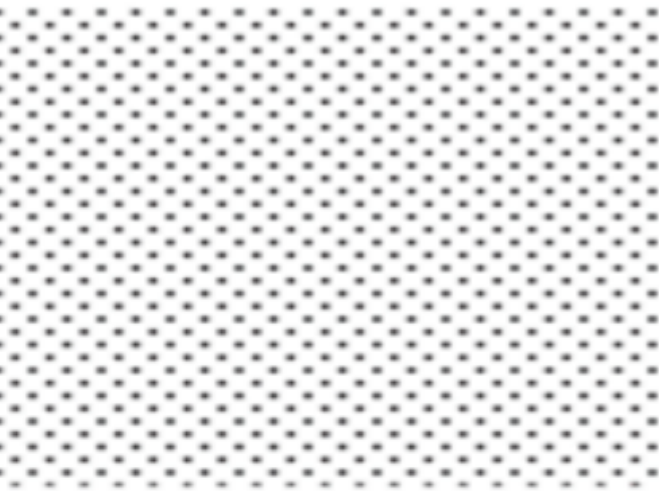
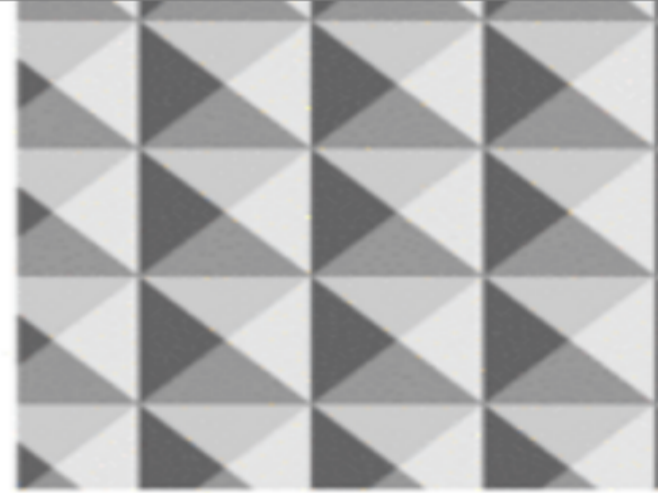
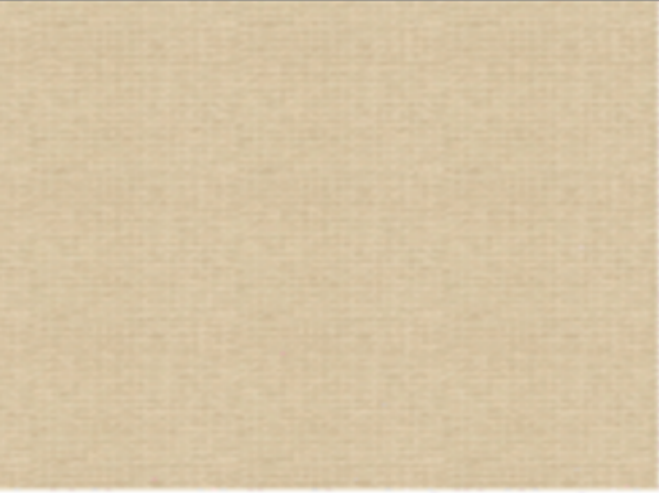


Call Sites / sun.misc.Unsafe methods

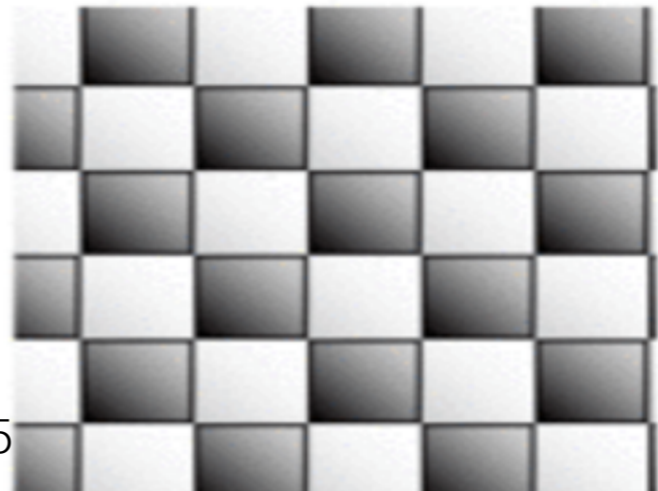
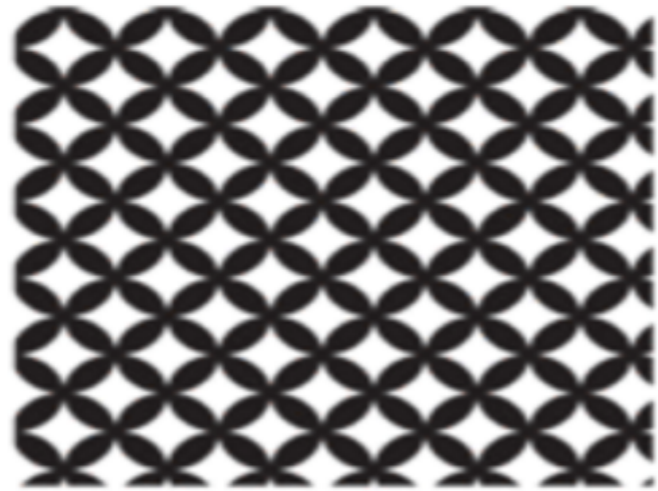
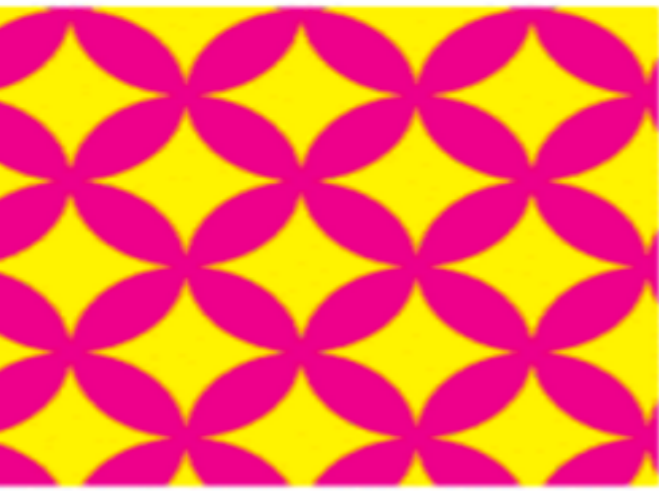


Research Question 3

Why are *Unsafe* features used?



14 patterns



Process Byte Arrays in Block

Process Byte Arrays in Block

```
for (int i = 0; i < minWords * Longs.BYTES; i += Longs.BYTES) {
    long lw = theUnsafe.getLong(left, BYTE_ARRAY_BASE_OFFSET + (long) i);
    long rw = theUnsafe.getLong(right, BYTE_ARRAY_BASE_OFFSET + (long) i);
    if (lw != rw) {
        if (BIG_ENDIAN) {
            return UnsignedLongs.compare(lw, rw);
        }

        /*
         * We want to compare only the first index where left[index] != right[index].
         * This corresponds to the least significant nonzero byte in lw ^ rw, since lw
         * and rw are little-endian. Long.numberOfTrailingZeros(diff) tells us the least
         * significant nonzero bit, and zeroing out the first three bits of L.nTZ gives us the
         * shift to get that least significant nonzero byte.
         */
        int n = Long.numberOfTrailingZeros(lw ^ rw) & ~0x7;
        return (int) (((lw >>> n) & UNSIGNED_MASK) - ((rw >>> n) & UNSIGNED_MASK));
    }
}
```

Process Byte Arrays in Block

```
for (int i = 0; i < minWords * Longs.BYTES; i += Longs.BYTES) {
    long lw = theUnsafe.getLong(left, BYTE_ARRAY_BASE_OFFSET + (long) i);
    long rw = theUnsafe.getLong(right, BYTE_ARRAY_BASE_OFFSET + (long) i);
    if (lw != rw) {
        if (BIG_ENDIAN) {
            return UnsignedLongs.compare(lw, rw);
        }

        /*
         * We want to compare only the first index where left[index] != right[index].
         * This corresponds to the least significant nonzero byte in lw ^ rw, since lw
         * and rw are little-endian. Long.numberOfTrailingZeros(diff) tells us the least
         * significant nonzero bit, and zeroing out the first three bits of L.nTZ gives us the
         * shift to get that least significant nonzero byte.
         */
        int n = Long.numberOfTrailingZeros(lw ^ rw) & ~0x7;
        return (int) (((lw >>> n) & UNSIGNED_MASK) - ((rw >>> n) & UNSIGNED_MASK));
    }
}
```


Process Byte Arrays in Block

```
for (int i = 0; i < minWords * Long.numberOfTrailingZeros(diff); i++) {
    long lw = theUnsafe.getLong(left, BYTE_ARRAY_BASE_OFFSET + (long) i);
    long rw = theUnsafe.getLong(right, BYTE_ARRAY_BASE_OFFSET + (long) i);
    if (lw != rw) {
        if (BIG_ENDIAN) {
            return UnsignedLongs.compare(lw, rw);
        }

        /*
         * We want to compare only the first index where left[index] != right[index].
         * This corresponds to the least significant nonzero byte in lw ^ rw, since lw
         * and rw are little-endian. Long.numberOfTrailingZeros(diff) tells us the least
         * significant nonzero bit, and zeroing out the first three bits of L.nTZ gives us the
         * shift to get that least significant nonzero byte.
         */
        int n = Long.numberOfTrailingZeros(lw ^ rw) & ~0x7;
        return (int) (((lw >>> n) & UNSIGNED_MASK) - ((rw >>> n) & UNSIGNED_MASK));
    }
}
```

Process Byte Arrays in Block

```
for (int i = 0; i < UnsignedLongs.BYTES; i += Longs.BYTES) {
    long lw = theUnsafe.getLong(left, BYTE_ARRAY_BASE_OFFSET + (long) i);
    long rw = theUnsafe.getLong(right, BYTE_ARRAY_BASE_OFFSET + (long) i);
    if (lw != rw) {
        if (BIG_ENDIAN) {
            return UnsignedLongs.compare(lw, rw);
        }

        /*
         * We want to compare only the first index where left[index] != right[index].
         * This corresponds to the least significant nonzero byte in lw ^ rw, since lw
         * and rw are little-endian. Long.numberOfTrailingZeros(diff) tells us the least
         * significant nonzero bit, and zeroing out the first three bits of L.nTZ gives us the
         * shift to get that least significant nonzero byte.
         */
        int n = Long.numberOfTrailingZeros(lw ^ rw) & ~0x7;
        return (int) (((lw >>> n) & UNSIGNED_MASK) - ((rw >>> n) & UNSIGNED_MASK));
    }
}
```

Process Byte Arrays in Block

```
for (int i = 0; i < minWords * Longs.BYTES; i += Longs.BYTES) {
    long lw = theUnsafe.getLong(left, BYTE_ARRAY_BASE_OFFSET + (long) i);
    long rw = theUnsafe.getLong(right, BYTE_ARRAY_BASE_OFFSET + (long) i);
    if (lw != rw) {
        if (BIG_ENDIAN) {
            return UnsignedLongs.compare(lw, rw);
        }
        // sanity check - this should never fail
        /* if (theUnsafe.arrayIndexScale(byte[].class) != 1) {
            * W
            * T    throw new AssertionError();
            * a }
            * s
            * shift to get that least significant nonzero byte.
            */
        int n = Long.numberOfTrailingZeros(lw ^ rw) & ~0x7;
        return (int) (((lw >>> n) & UNSIGNED_MASK) - ((rw >>> n) & UNSIGNED_MASK));
    }
}
```

Process Byte Arrays in Block

```
for (int i = 0; i < minWords * Longs.BYTES; i += Longs.BYTES) {
    long lw = theUnsafe.getLong(left, BYTE_ARRAY_BASE_OFFSET + (long) i);
    long rw = theUnsafe.getLong(right, BYTE_ARRAY_BASE_OFFSET + (long) i);
    if (lw != rw) {
        if (BIG_ENDIAN) {
            return UnsignedLongs.compare(lw, rw);
        }
        // sanity check: this should never fail
        /* if (theUnsafe.arrayIndexScale(byte[].class) != 1) {
            * W
            * T
            * a }
            * s
            * shift to get that least significant nonzero byte.
            */
        int n = Long.numberOfTrailingZeros(lw ^ rw) & ~0x7;
        return (int) (((lw >>> n) & UNSIGNED_MASK) - ((rw >>> n) & UNSIGNED_MASK));
    }
}
```


Process Byte Arrays in Block

```
for (int i = 0; i < minWords * Longs.BYTES; i += Longs.BYTES) {
    long lw = theUnsafe.getLong(left, BYTE_ARRAY_BASE_OFFSET + (long) i);
    long rw = theUnsafe.getLong(right, BYTE_ARRAY_BASE_OFFSET + (long) i);
    if (lw != rw) {
        if (BIG_ENDIAN) {
            return UnsignedLongs.reverseBytes(lw);
        }
        // sanity
        /* if (theUnsafe.getLong(right, BYTE_ARRAY_BASE_OFFSET + (long) i) != 1) {
            * W
            * T throw new IllegalArgumentException("Mismatched values");
            * a }
            * s
            * shift to get the
            */
        int n = Long.numberOfTrailingZeros(lw ^ rw);
        return (int) (((lw >>> n) & UNSIGNED_MASK) - ((rw >>> n) & UNSIGNED_MASK));
    }
}
```

- Stored contiguously

Process Byte Arrays in Block

```
for (int i = 0; i < minWords * Longs.BYTES; i += Longs.BYTES) {
    long lw = theUnsafe.getLong(left, BYTE_ARRAY_BASE_OFFSET + (long) i);
    long rw = theUnsafe.getLong(right, BYTE_ARRAY_BASE_OFFSET + (long) i);
    if (lw != rw) {
        if (BIG_ENDIAN) {
            return UnsignedLongs.compare(lw, rw);
        }
        // sanity
        /* if (theUnsafe.getByte(right, BYTE_ARRAY_BASE_OFFSET + (long) i) != 1) {
            * W
            * T throw new AssertionError("Mismatch at " + i);
            * a }
            * s
            * shift to get the
            */
        int n = Long.numberOfTrailingZeros(lw ^ rw);
        return (int) (((lw >>> n) & UNSIGNED_MASK) - ((rw >>> n) & UNSIGNED_MASK));
    }
}
```

- Stored contiguously
- Endianness

Process Byte Arrays in Block

```
for (int i = 0; i < minWords * Longs.BYTES; i += Longs.BYTES) {
    long lw = theUnsafe.getLong(left, BYTE_ARRAY_BASE_OFFSET + (long) i);
    long rw = theUnsafe.getLong(right, BYTE_ARRAY_BASE_OFFSET + (long) i);
    if (lw != rw) {
        if (BIG_ENDIAN) {
            return UnsignedLongs.compare(lw, rw);
        }
        // sanity
        /* if (theUnsafe.getLong(right, BYTE_ARRAY_BASE_OFFSET + (long) i) != 1) {
            * W
            * T throw new AssertionError("Mismatch");
            * a }
            * s
            * shift to get the
            */
        int n = Long.numberOfTrailingZeros(lw ^ rw);
        return (int) (((lw >>> n) & UNSIGNED_MASK) - ((rw >>> n) & UNSIGNED_MASK));
    }
}
```

- Stored contiguously
- Endianness
- Performance issues

Atomic Operations

```
do
{
    currentValue = get();
    newValue = currentValue + increment;
}
while (!compareAndSet(currentValue, newValue));
```

```
do
{
    currentValue = get();
    newValue = currentValue + increment;
}
while (!compareAndSet(currentValue, newValue));
```

```
do
{
    currentValue = get();
    newValue = currentValue + increment;
}
while (!compareAndSet(currentValue, newValue));
```

```
do
{
    currentValue = get\(\);
    newValue = currentValue + increment;
}
while (!compareAndSet(currentValue, newValue));
```

```
return UNSAFE.compareAndSwapLong(this, VALUE_OFFSET, expectedValue, newValue);
```



```
do
{
    currentValue = get();
    newValue = currentValue + increment;
}
while (!compareAndSet(currentValue, newValue));
```

```
return UNSAFE.compareAndSwapLong(this, VALUE_OFFSET, expectedValue, newValue);
```

```
do
{
    currentValue = get();
    newValue = currentValue + increment;
}
while (!compareAndSet(currentValue, newValue));
```

```
return UNSAFE.compareAndSwapLong(this, VALUE_OFFSET, expectedValue, newValue);
```

```
do
{
    currentValue = get\(\);
    newValue = currentValue + increment;
}
while (!compareAndSet(currentValue, newValue));
```

```
return UNSAFE.compareAndSwapLong(this, VALUE_OFFSET, expectedValue, newValue);
```

```
VALUE_OFFSET = UNSAFE.objectFieldOffset(Value.class.getDeclaredField("value"));
```

LMAX Disruptor

com.lmax.disruptor.Sequence

Atomic Operations

```
do
{
    currentValue = get();
    newValue = currentValue + increment;
}
while (!compareAndSet(currentValue, newValue));
```

```
return UNSAFE.compareAndSwapLong(this, VALUE_OFFSET, expectedValue, newValue);
```

```
VALUE_OFFSET = UNSAFE.objectFieldOffset(Value.class.getDeclaredField("value"));
```

LMAX Disruptor

com.lmax.disruptor.Sequence

```
do
{
    currentValue = get();
    newValue = currentValue + increment;
}
while ( ! compareAndSet(currentValue, newValue));
```

• Hard to get it right

```
return UNSAFE.compareAndSet(currentValue, newValue);
```

```
VALUE_OFFSET = UNSAFE.objectFieldOffset(Value.class.getDeclaredField("value"));
```

LMAX Disruptor

com.lmax.disruptor.Sequence

```
do
{
    currentValue = get();
    newValue = currentValue + increment;
}
while ( ! compareAndSet(currentValue, newValue));
```

- Hard to get it right
- Java Memory Model

```
return UNSAFE.compareAndSet(currentValue, newValue);
```

```
VALUE_OFFSET = UNSAFE.objectFieldOffset(Value.class.getDeclaredField("value"));
```

LMAX Disruptor

com.lmax.disruptor.Sequence

Update Final Fields

Update Final Fields

```
if (WRITE_TO_FINAL_FIELDS) {  
    String string = new String();  
    UNSAFE.putObject(string, STRING_VALUE_FIELD_OFFSET, chars);  
    UNSAFE.putInt(string, STRING_COUNT_FIELD_OFFSET, chars.length);  
    return string;  
} else {  
    return new String(chars);  
}
```

```
if (WRITE_TO_FINAL_FIELDS) {
    String string = new String();
    UNSAFE.putObject(string, STRING_VALUE_FIELD_OFFSET, chars);
    UNSAFE.putInt(string, STRING_COUNT_FIELD_OFFSET, chars.length);
    return string;
} else {
    return new String(chars);
}
```



groovy.json.internal.FastStringUtils

Update Final Fields

```
if (WRITE_TO_FINAL_FIELDS) {  
    String string = new String();  
    UNSAFE.putObject(string, STRING_VALUE_FIELD_OFFSET, chars);  
    UNSAFE.putInt(string, STRING_COUNT_FIELD_OFFSET, chars.length);  
    return string;  
} else {  
    return new String(chars);  
}
```



groovy.json.internal.FastStringUtils

Update Final Fields

```
if (WRITE_TO_FINAL_FIELDS) {  
    String string = new String(chars);  
    UNSAFE.putObject(string, STRING_VALUE_FIELD_OFFSET, chars);  
    UNSAFE.putInt(string, STRING_COUNT_FIELD_OFFSET, chars.length);  
    return string;  
} else {  
    return new String(chars);  
}
```



groovy.json.internal.FastStringUtils

```
if (WRITE_TO_FINAL_FIELDS) {  
    String string = new String();  
    UNSAFE.putObject(string, STRING_VALUE_FIELD_OFFSET, chars);  
    UNSAFE.putInt(string, STRING_COUNT_FIELD_OFFSET, chars.length);  
    return string;  
} else {  
    return new String(chars);  
}
```

```
STRING_VALUE_FIELD_OFFSET = getFieldOffset("value");  
STRING_COUNT_FIELD_OFFSET = getFieldOffset("count");
```



groovy.json.internal.FastStringUtils


```
if (WRITE_TO_FINAL_FIELDS) {  
    String string = new String();  
    UNSAFE.putObject(string, STRING_VALUE_FIELD_OFFSET, chars);  
    UNSAFE.putInt(string, STRING_COUNT_FIELD_OFFSET, chars.length);  
    return string;  
} else {  
    return new String(chars);  
}  
  
STRING_VALUE_FIELD_OFFSET = getFieldOffset("value");  
STRING_COUNT_FIELD_OFFSET = getFieldOffset("count");  
  
return UNSAFE.objectFieldOffset(String.class.getDeclaredField(fieldName));
```



groovy.json.internal.FastStringUtils

Update Final Fields

```
if (WRITE_TO_FINAL_FIELDS) {
    String string = new String();
    UNSAFE.putObject(string, STRING_VALUE_FIELD_OFFSET, chars);
    UNSAFE.putInt(string, STRING_COUNT_FIELD_OFFSET, chars.length);
    return string;
} else {
    return new String(chars);
}

STRING_VALUE_FIELD_OFFSET = getFieldOffset("value");
STRING_COUNT_FIELD_OFFSET = getFieldOffset("count");

return UNSAFE.objectFieldOffset(String.class.getDeclaredField(fieldName));
```



groovy.json.internal.FastStringUtils

```
if (WRITE_TO_FINAL_FIELDS) {  
    String string = new String();  
    UNSAFE.putObject(string, STRING_VALUE_FIELD_OFFSET, chars);  
    UNSAFE.putInt(string, STRING_COUNT_FIELD_OFFSET, chars.length);  
    return string;  
} else {  
    return new St  
}  
  
STRING_VALUE_FIELD_OFFSET = getFieldOffset("value");  
STRING_COUNT_FIELD_OFFSET = getFieldOffset("count");  
  
return UNSAFE.objectFieldOffset(String.class.getDeclaredField(fieldName));
```

- Compiler optimizations



groovy.json.internal.FastStringUtils

```
if (WRITE_TO_FINAL_FIELDS) {
    String string = new String();
    UNSAFE.putObject(string, STRING_VALUE_FIELD_OFFSET, chars);
    UNSAFE.putInt(string, STRING_COUNT_FIELD_OFFSET, chars.length);
    return string;
} else {
    return new String(chars);
}

STRING_VALUE_FIELD_OFFSET = getFieldOffset("value");
STRING_COUNT_FIELD_OFFSET = getFieldOffset("count");

return UNSAFE.objectFieldOffset(String.class.getDeclaredField(fieldName));
```

- Compiler optimizations
- Thread visibility



groovy.json.internal.FastStringUtils

Research Question 4

What **problems** do developers who use *Unsafe* encounter?

Using sun.misc.Unsafe in real world [closed]



Of course the Unsafe class is undocumented, but how can I use it in a real world scenario

141

java

memcpy function in C++ to Java equivalent

▲ I have in C++

5

```
memcpy (&wkpm, (PMSK *)pr_masks + (long)(x - 1), sizeof(PMSK));
```

▼ where `PMSK` is a struct. It will be a class in Java.

★ Now assuming that here I am copying the whole chunk of memory into `pr_masks` i.e creating an additional instance of the `PMSK` class. How to do this in Java.

Example: In a java code at line 20 I want capture the class instance and then again use that same instance in line 100. In between there may be many modifications.

Hope I am clear with my question.

Thanks

java

c++

How to free memory using Java Unsafe, using a Java reference?



0

Java Unsafe class allows you to allocate memory for an object as follows, but using this method how would you free up the memory allocated when finished, as it does not provide the memory address...



1

```
Field f = Unsafe.class.getDeclaredField("theUnsafe"); //Internal reference
f.setAccessible(true);
Unsafe unsafe = (Unsafe) f.get(null);

//This creates an instance of player class without any initialization
Player p = (Player) unsafe.allocateInstance(Player.class);
```

Is there a way of accessing the memory address from the object reference, maybe the integer returned by the default hashCode implementation will work, so you could do...

```
unsafe.freeMemory(p.hashCode());
```

doesn't seem right some how...

java

memory-management

jvm

directmemory

Process Byte Arrays in Block

****BUSTED**** How to speed up a byte[] lookup to be faster using sun.misc.Unsafe?



I am experimenting with Unsafe to iterate over memory instead of iterating over the values in a byte[]. A memory block is allocated using unsafe. The memory is sufficient to hold 65536 byte values.

asked 3 years ago

viewed 2396 times

active 2 years ago



I AM TRYING THIS:



1

```
char aChar = some character

if ((byte) 0 == (unsafe.getBytes(base_address + aChar) & mask)){
    // do something
}
```

INSTEAD OF:

```
char aChar = some character

if ((byte) 0 == ( lookup[aChar] & mask )){
    // do something
}
```

Jobs near

Senior Java Developer
lastminute.com
Chiasso, Switzerland

java

Agile Coach - Java
lastminute.com
Chiasso, Switzerland

[More jobs near L](#)

<https://bitbucket.org/acuarica/java-unsafe-analysis>

luis.mastrangelo@usi.ch



959,300 artifacts 1.7 TB

Last version

106,574 artifacts

.jar, .war, .ear, .ejb

Impact

48,139 call sites
351 field accesses

817 *Unsafe* artifacts (1%)

Artifact dependencies

21,297 artifacts depend on *Unsafe* (25%)



959,300 artifacts 1.7 TB

Last version

106,574 artifacts

.jar, .war, .ear, .ejb

Impact

48,139 call sites
351 field accesses

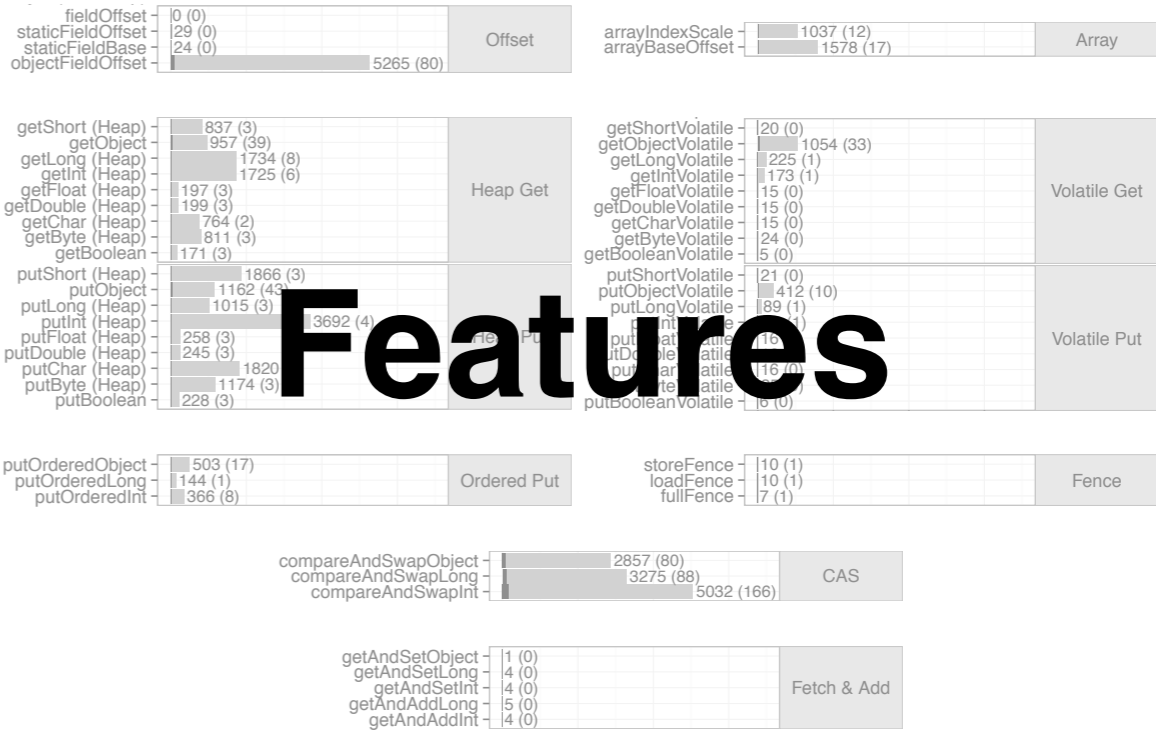
817 *Unsafe* artifacts (1%)

Artifact dependencies

21,297 artifacts depend on *Unsafe* (25%)

Application () Language (■)

Call Sites / sun.misc.Unsafe methods





959,300 artifacts 1.7 TB

Last version

106,574 artifacts

.jar, .war, .ear, .ejb

Impact

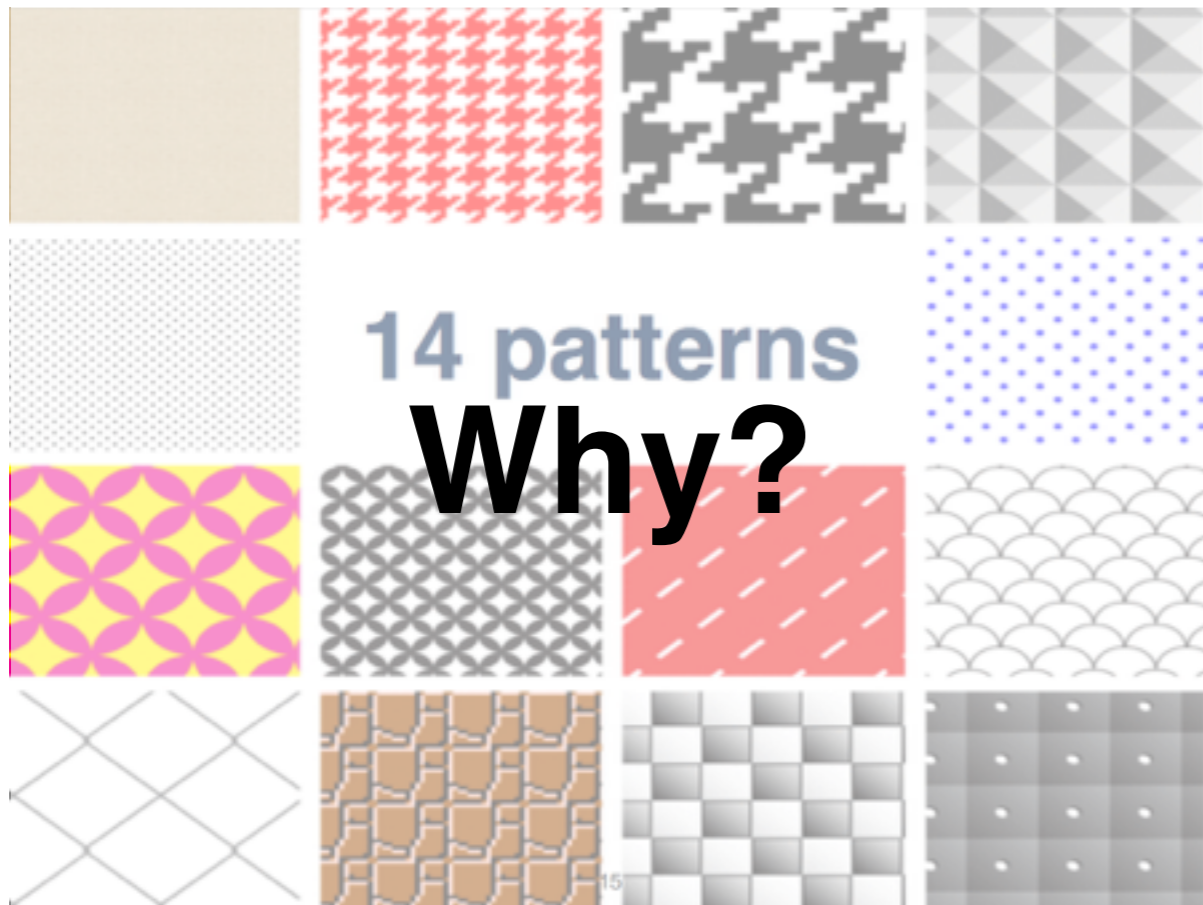
48,139 call sites
351 field accesses

817 Unsafe artifacts (1%)

Artifact dependencies

21,297 artifacts depend on Unsafe (25%)

11



15

Application (Language)

Call Sites / sun.misc.Unsafe methods



13



959,300 artifacts 1.7 TB

Last version

106,574 artifacts

.jar, .war, .ear, .ejb

Impact

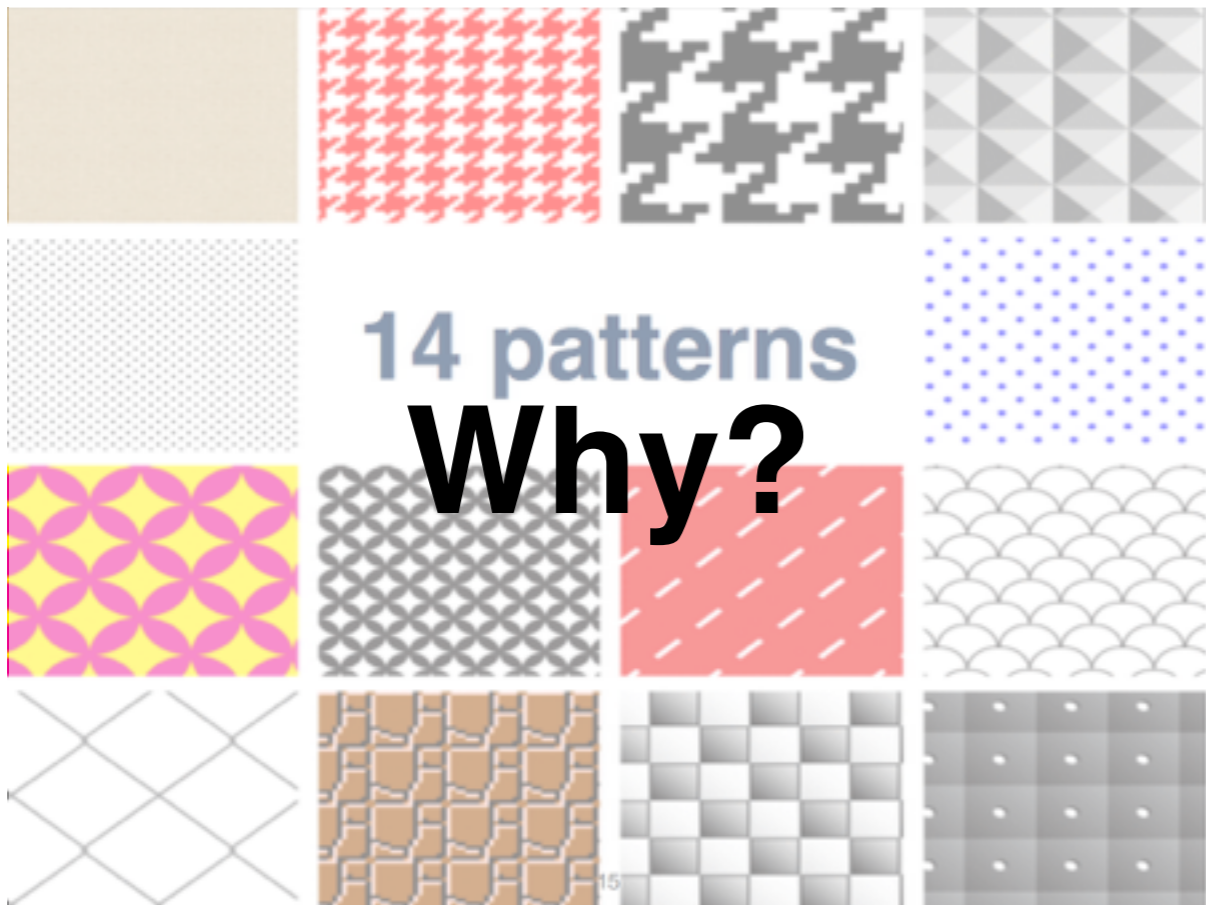
48,139 call sites
351 field accesses

817 *Unsafe* artifacts (1%)

Artifact dependencies

21,297 artifacts depend on *Unsafe* (25%)

11



Features

13

Misdirected uses

How to free memory using Java Unsafe, using a Java reference?

```

Java Unsafe class allows you to allocate memory for an object as follows, but using this method how would you free up the memory allocated when finished, as it does not provide the memory address...

Field f = Unsafe.class.getDeclaredField("theUnsafe"); //Internal reference
f.setAccessible(true);
Unsafe unsafe = (Unsafe) f.get(null);
//The Unsafe class has a method to free memory, but it requires the memory address.
Play p = (Player) unsafe.allocateInstance(Player.class);

unsafe.freeMemory(p.hashCode());

```

Problems

doesn't seem right some how...

java memory-management jvm directmemory



959,300 artifacts 1.7 TB

Last version

106,574 artifacts

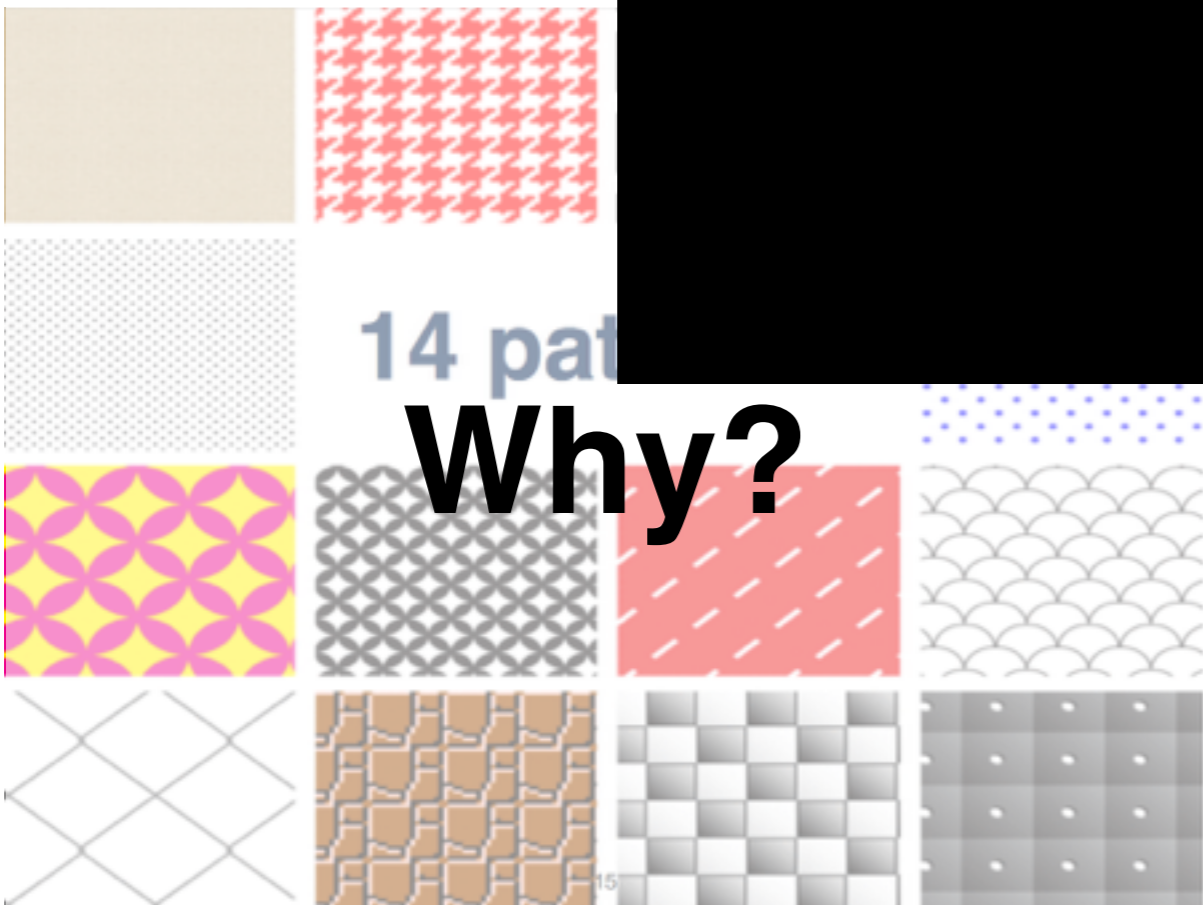
.jar, .war, .ear, .ejb

Impact

48,139 call sites
351 field accesses

817 Unsafe artifacts (1%)

21,297 artifacts depe



14 pat

Why?



Application (Language)

Call Sites / sun.misc.Unsafe methods

Table with columns for method names, counts, and categories like 'Offset', 'Heap Get', 'Volatile Get', 'Volatile Put', 'Fence', 'CAS', and 'Fetch & Add'.

Evidence

Problems

```
f.setAccessible(true);
Unsafe unsafe = (Unsafe) f.get(null);
//This uses Integer and plays with the hashCode method
Player p = (Player) unsafe.allocateInstance(Player.class);
unsafe.freeMemory(p.hashCode());
```

Is there a way of accessing the memory address from the object reference, maybe the integer returned by the default hashCode implementation will work, so you could do...

tags: java, memory-management, jvm, directmemory



959,300 artifacts 1.7 TB

Last version

106,574 artifacts

.jar, .war, .ear, .ejb

Impact

48,139 call sites
351 field accesses

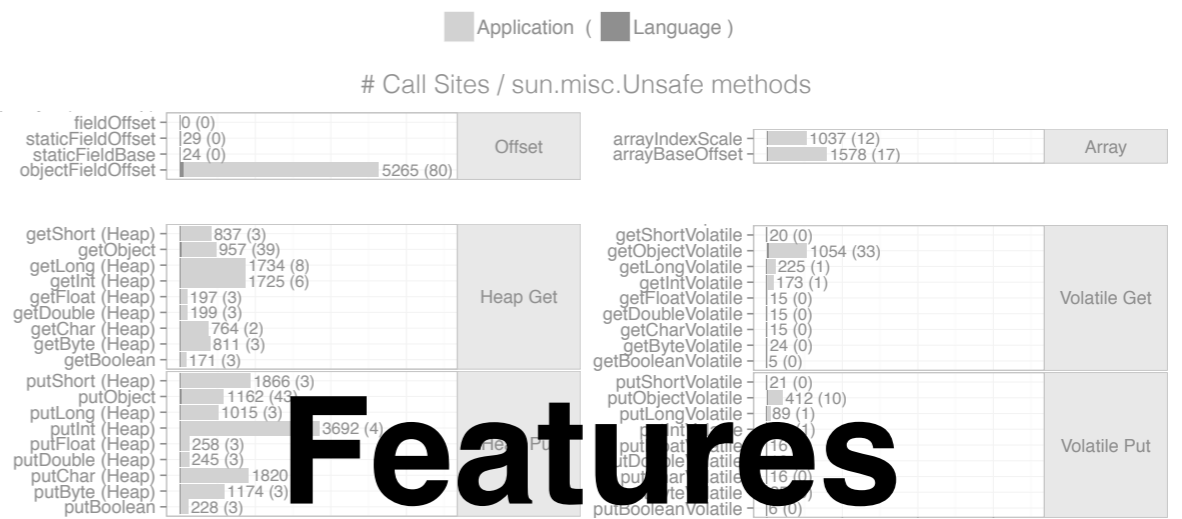
817 Unsafe artifacts (1%)

21,297 artifacts depe



14 pat

Why?



Features

Evidence Evolution

Problems

```
f.setAccessible(true);
Unsafe unsafe = (Unsafe) f.getDeclaredField("theUnsafe").get(null);
//This uses the Unsafe class to get the memory address of the object
Player p = (Player) unsafe.allocateInstance(Player.class);
unsafe.freeMemory(p.hashCode());
```

Is there a way of accessing the memory address from the object reference, maybe the integer returned by the default hashCode implementation will work, so you could do...

java memory-management jvm directmemory



959,300 artifacts 1.7 TB

Last version

106,574 artifacts

.jar, .war, .ear, .ejb

Impact

48,139 call sites
351 field accesses

817 Unsafe artifacts (1%)

21,297 artifacts depe



14 pat

Why?

Evidence
Evolution
Unsafe

Problems



Misdirected uses

Unsafe, using a Java reference?

memory for an object as follows, but using this method
ed when finished, as it does not provide the memory

```
f.setAccessible(true);
Unsafe unsafe = (Unsafe) f.get("Unsafe");
//This class is used to free memory for an object as follows, but using this method
//when finished, as it does not provide the memory
Play p = (Player) unsafe.allocateInstance(Player.class);
unsafe.freeMemory(p.hashCode());
```

doesn't seem right some how...

java memory-management jvm directmemory