

## HOW TO GET STARTED IN HACKING OPENJDK?

*Vineel Kumar Reddy Kovvuri*  
<http://vineelkumarreddy.com>

# Contents

<b>1</b>	<b>Introduction</b>	<b>2</b>
<b>2</b>	<b>How do I get the source code?</b>	<b>2</b>
<b>3</b>	<b>How do I build it?</b>	<b>3</b>
<b>4</b>	<b>How do I hack it?</b>	<b>5</b>
<b>5</b>	<b>References</b>	<b>10</b>

# 1 Introduction

This blog post is not about internals of JDK(Java Development Kit), But a mere documentation for people who want to get started. JDK is open sourced a long time ago. The reference implementation of Java is now based on [OpenJDK](#). So any one interested in it can get the source code and play with it. Here on, when I refer to JDK it means all the components including Java class libraries/JVM(Hotspot)/Java Compiler. In this post, I will walk you in getting and building the latest JDK 9 sources.

## 2 How do I get the source code?

All the components of JDK 9 are from <http://hg.openjdk.java.net/jdk9/> . OpenJDK is under Mercurial. For any one familiar with Git, this is no different. Below is the layout of repositories in JDK 9 OpenJDK forest.

1. dev - top-level umbrella project for all the sub projects
  - (a) corba - Not my interest
  - (b) hotspot - The Java Virtual Machine implementation. JIT/GC go here
  - (c) jaxp - Not my interest
  - (d) jaxws - Not my interest
  - (e) jdk - The Java class libraries including native JNI implementation go here
  - (f) langtools - The Java compiler and other tools go here
  - (g) nashorn - Not my interest

There are multiple ways to get the code. The easiest being, first clone the umbrella 'dev' project and run `get_source.sh` to pull all other repositories.

```
hg clone http://hg.openjdk.java.net/jdk9/dev 9dev
cd 9dev
sh ./get_source.sh
```

Once done you should see below directory structure.

```
vineel@Z97X:~/9dev$ tree -L 1
.
├── [1.5K] ASSEMBLY_EXCEPTION
├── [4.0K] build
├── [4.0K] common
├── [1.6K] configure
├── [4.0K] corba
├── [3.0K] get_source.sh
├── [4.0K] hacking
├── [4.0K] .hg
├── [ 103] .hgignore
├── [ 17K] .hgtags
├── [4.0K] hotspot
├── [4.0K] jaxp
├── [4.0K] jaxws
├── [4.0K] .jcheck
├── [4.0K] jdk
├── [4.0K] langtools
├── [ 19K] LICENSE
├── [4.0K] make
├── [2.7K] Makefile
├── [ 53K] modules.xml
├── [4.0K] nashorn
├── [1.5K] README
├── [ 57K] README-builds.html
├── [ 50K] README-builds.md
├── [4.0K] test
└── [174K] THIRD_PARTY_README
```

Figure 1

### 3 How do I build it?

Building OpenJDK is straight forward too, The only precursor is to get the dependencies. The OpenJDKs Adopt OpenJDK project has already documented the dependencies here, Its a matter of sudo apt-get install . Run below commands to build the source code.

bash ./configure #Setup the environment  
make all #build the entire forest

```
vineel@Z97X:~/9dev$ tree -L 2 build/  
build/  
├── [4.0K] linux-x86_64-normal-server-release → Generated by bash ./configure  
├── [2.7K] bootcycle-spec.gmk  
├── [2.0K] build.log  
├── [ 24K] build.log.old  
├── [4.0K] buildtools  
├── [3.1K] compare.sh  
├── [ 11K] configure.log  
├── [ 11K] configure.log.old  
├── [4.0K] configure-support  
├── [4.0K] hotspot  
├── [4.5K] hotspot-spec.gmk  
├── [4.0K] images → Generated by make all  
├── [4.0K] jdk  
├── [1.2K] Makefile  
├── [4.0K] make-support  
├── [ 26K] spec.gmk  
└── [4.0K] support
```

Figure 2: OpenJDK build directory after running bash ./configure

```
vineel@Z97X:~/9dev$ tree -L 2 build/*/images
build/linux-x86_64-normal-server-release/images
├── [4.0K] docs
│   ├── [4.0K] api
│   ├── [4.0K] jdk
│   ├── [4.0K] jre
│   └── [4.0K] platform
├── [4.0K] jdk
│   ├── [1.5K] ASSEMBLY_EXCEPTION
│   ├── [4.0K] bin
│   ├── [4.0K] conf
│   ├── [4.0K] demo
│   ├── [4.0K] include
│   ├── [118K] jrt-fs.jar
│   ├── [4.0K] lib
│   ├── [ 19K] LICENSE
│   ├── [4.0K] man
│   ├── [1.2K] release
│   ├── [4.0K] sample
│   ├── [ 51M] src.zip
│   └── [174K] THIRD PARTY README
├── [4.0K] jre
│   ├── [1.5K] ASSEMBLY_EXCEPTION
│   ├── [4.0K] bin
│   ├── [4.0K] conf
│   ├── [4.0K] lib
│   ├── [ 19K] LICENSE
│   ├── [4.0K] man
│   ├── [1.0K] release
│   └── [174K] THIRD PARTY README
├── [1.4M] sec-bin.zip
├── [4.0K] test
│   ├── [4.0K] hotspot
│   ├── [4.0K] jdk
│   └── [ 15] Readme.txt
└── [ 610] _the.nashorn.jar.vardeps
```

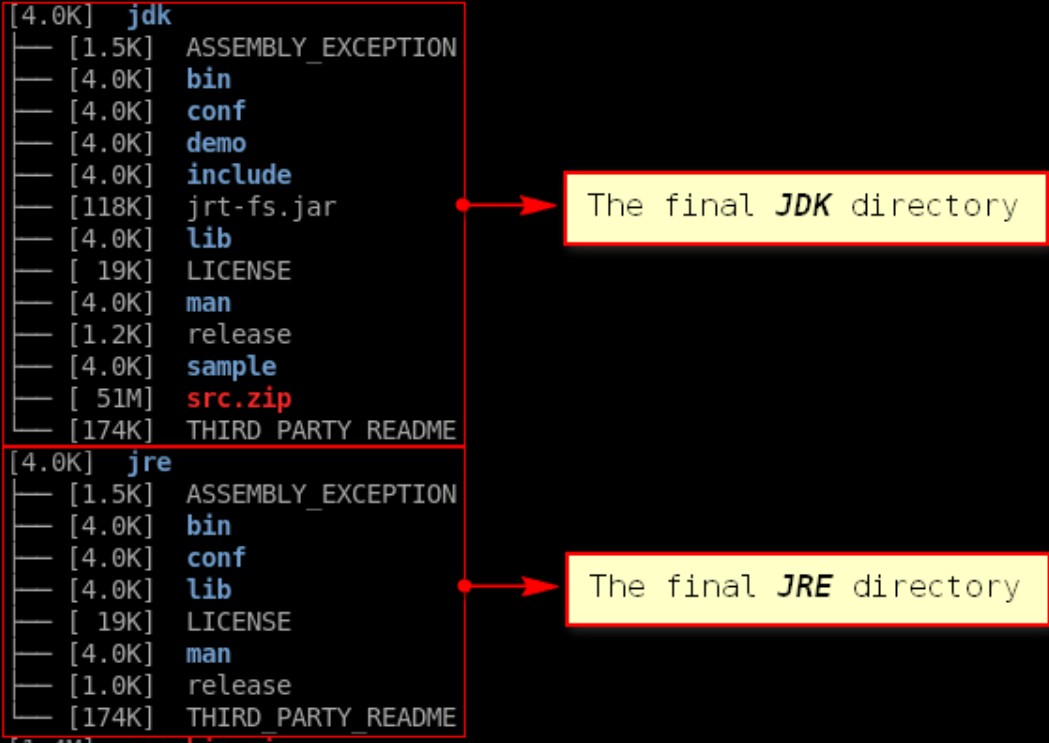


Figure 3: OpenJDK images directory after running make all

## 4 How do I hack it?

Now its all yours! This is where the fun begins, Even though OpenJDK sources are highly organized, you might need the help of some tools to get around with it. My favorite tool of choice for any source code exploration is OpenGrok, I have my local OpenGrok instance setup.

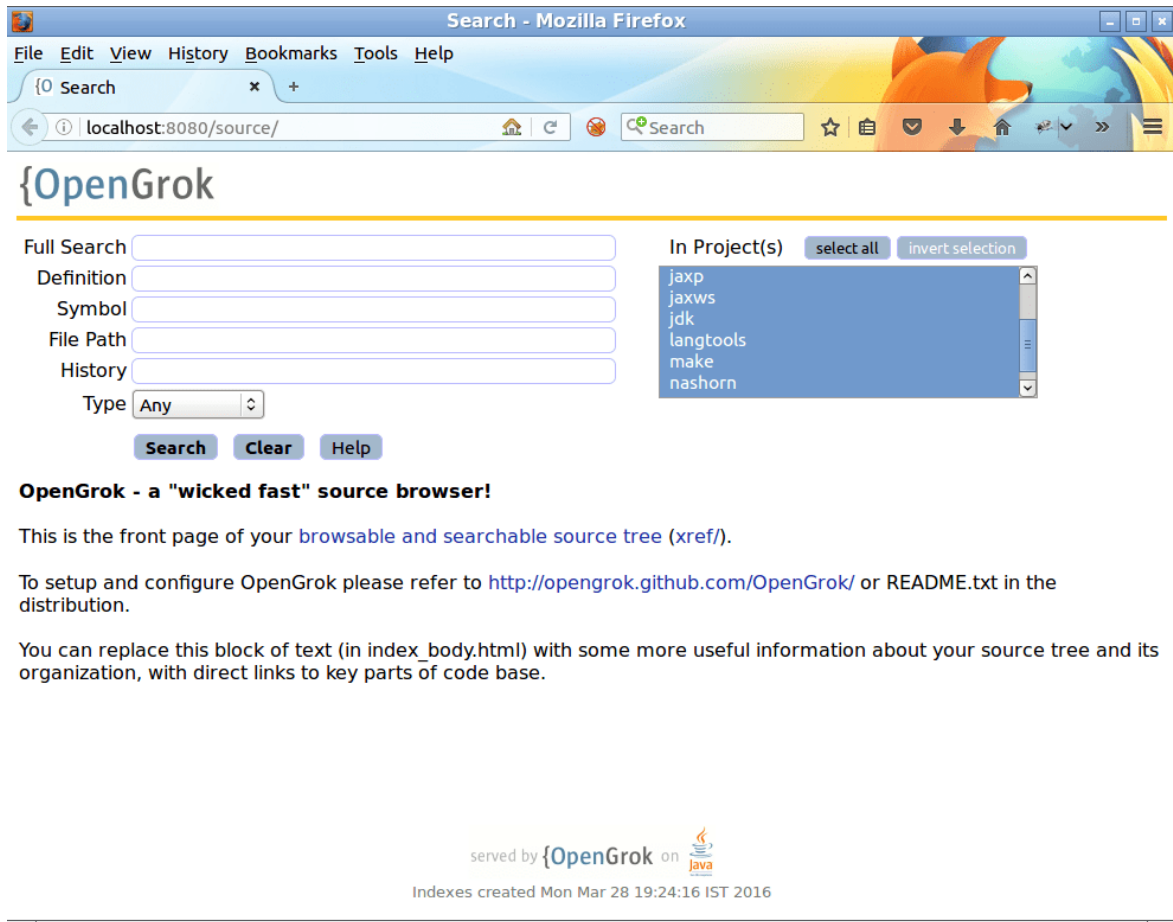


Figure 4

Since `images/jdk/bin/java` is the binary which executes your class files. I wanted to locate the code for it, so that, I can see where the rabbit hole is leading to (essentially in to hotspot :), The right way to find out is to decipher the build system to understand how things are getting build and from where they are getting build. I really really! dont want to explore the build system now. So, Its time to wear my black hat and roll up some reverse engineering skills!

With the help of `objdump` I disassembled the `jdk/bin/java` binary and looked for symbols, specifically in its main function. Once I know some candidate symbols, I can get to the code from OpenGrok. So below is the disassembled view of main function of the java binary. This function is in turn making calls to `JLL_*` apis. So I selected `JLI_PreprocessArg` as my candidate symbol.

```

vineel@Z97X:~/9dev$ objdump -d build/*/images/jdk/bin/java | grep -i -A30 "<main>"
000000000400880: <main>:
400880: 55                push  %rbp
400881: 48 89 e5          mov   %rsp,%rbp
400884: 41 57            push  %r15
400886: 41 56            push  %r14
400888: 41 55            push  %r13
40088a: 41 54            push  %r12
40088c: 49 89 f6          mov   %rsi,%r14
40088f: 53                push  %rbx
400890: 31 f6            xor   %esi,%esi
400892: 89 fb            mov   %edi,%ebx
400894: bf 01 00 00 00    mov   $0x1,%edi
400899: 48 83 ec 18      sub   $0x18,%rsp
40089d: e8 3e ff ff ff   callq 4007e0 <JLI_InitArgProcessing@plt>
4008a2: 8d 7b 01         lea  0x1(%rbx),%edi
4008a5: 48 63 ff         movslq %edi,%rdi
4008a8: e8 63 ff ff ff   callq 400810 <JLI_List_new@plt>
4008ad: 85 db            test  %ebx,%ebx
4008af: 49 89 c5          mov   %rax,%r13
4008b2: 7e 0e            jle  400922 <main+0xa2>
4008b4: 8d 43 ff         lea  -0x1(%rbx),%eax
4008b7: 49 8d 44 c6 08    lea  0x8(%r14,%rax,8),%rax
4008bc: 48 89 45 c8      mov   %rax,-0x38(%rbp)
4008c0: 49 8b 3e          mov   (%r14),%rdi
4008c3: e8 28 ff ff ff   callq 4007f0 <JLI_PreprocessArg@plt>
4008c8: 48 85 c0          test  %rax,%rax
4008cb: 49 89 c7          mov   %rax,%r15
4008ce: 0f 84 bd 00 00 00 je   400991 <main+0x111>
4008d4: 48 8b 50 08      mov   0x8(%rax),%rdx
4008d8: 85 d2            test  %edx,%edx
4008da: 7e 2c            jle  400908 <main+0x88>
vineel@Z97X:~/9dev$ █

```

Disassembly of *main* function in *images/jdk/bin/java* binary

This symbol looks promising!

Figure 5

Now searching for the symbol `JLI_PreprocessArg` in OpenGrok found the source code for `jdk/bin/java`!



Search - Mozilla Firefox

File Edit View History Bookmarks Tools Help

{0 Search

localhost:8080/source/search?q=&defs=&refs=

# {OpenGrok

Home Sort by: last modified time | **relevance** | path

Full Search

Definition

Symbol **JLI\_PreprocessArg**

File Path

History

Type Any

**Search** **Clear** **Help**

In Project(s) **select all** **invert selection**

- common
- corba
- hotspot
- jaxp
- jaxws
- jdk

Searched **refs:JLI\_PreprocessArg** (Results **1 - 4** of **4**) sorted by relevance

**/jdk/src/java.base/share/native/launcher/**

HAD **main.c** 131 JLI\_List argsInFile = JLI\_PreprocessArg(argv[i]);

**/jdk/src/java.base/share/native/libjli/**


HAD **jli\_util.h** 137 JLI\_List JLI\_PreprocessArg(const char \*arg);

HAD **args.c** 369 JLI\_List JLI\_PreprocessArg(const char \*arg) *function*  
522 JLI\_List tokens = JLI\_PreprocessArg(argv[i]);

**/jdk/src/java.base/windows/native/libjli/**

HAD **cmdtoargs.c** 288 argsInFile = JLI\_PreprocessArg(arg);

**Completed in 349 milliseconds**

served by {OpenGrok on 

Indexes created Mon Mar 28 19:24:16 IST 2016

Figure 6

```
94 main(int argc, char **argv)
95 {
96     int margc;
97     char** margv;
98     const jboolean const_javaw = JNI_FALSE;
99 #endif /* JAVAW */
100     JLI_InitArgProcessing(!HAS_JAVA_ARGS, const_disable_argfile);
101
102 #ifdef _WIN32
103 {
104     int i = 0;
105     if (getenv(JLDEBUG_ENV_ENTRY) != NULL) {
106         printf("Windows original main args:\n");
107         for (i = 0; i < __argc; i++) {
108             printf("wwd_args[%d] = %s\n", i, __argv[i]);
109         }
110     }
111 }
112 }
113 JLI_CmdToArgs(GetCommandLine());
114 margc = JLI_GetStdArgc();
115 // add one more to mark the end
116 margv = (char **)JLI_MemAlloc((margc + 1) * (sizeof(char *)));
117 {
118     int i = 0;
119     StdArg *stdargs = JLI_GetStdArgs();
120     for (i = 0; i < margc; i++) {
121         margv[i] = stdargs[i].arg;
122     }
123     margv[i] = NULL;
124 }
125 #else /* *NIXES */
126 {
127     // accommodate the NULL at the end
128     JLI_List args = JLI_List_new(argc + 1);
129     int i = 0;
130     for (i = 0; i < argc; i++) {
131         JLI_List argsInFile = JLI_PreprocessArg(argv[i]);
132         if (NULL == argsInFile) {
133             JLI_List_add(args, JLI_StringDup(argv[i]));
134         } else {
135             int cnt, idx;
136             cnt = argsInFile->size;
137             for (idx = 0; idx < cnt; idx++) {
138                 JLI_List_add(args, argsInFile->elements[idx]);
139             }
140             // Shallow free, we reuse the string to avoid copy
141             JLI_MemFree(argsInFile->elements);
142             JLI_MemFree(argsInFile);
143         }
144     }
145     margc = args->size;
146     // add the NULL pointer at argv[argc]
147     JLI_List_add(args, NULL);
148     margv = args->elements;
149 }
150 #endif /* WIN32 */
151 return JLI_Launch(margc, margv,
152                 sizeof(const_jargs) / sizeof(char *), const_jargs,
153                 sizeof(const_appclasspath) / sizeof(char *), const_appclasspath,
154                 VERSION_STRING,
155                 DOT_VERSION,
156                 (const_progname != NULL) ? const_progname : *margv,
```

Figure 7: Inside main.c

To double confirm, I have added a print message and did the make all and it indeed worked as expected! You could see my print when I execute java from the command line in the right side window.

```

main.c (-/9dev/jdk/src/java.base/share/native/launcher) - VIM 112x59
105 int i = 0;
106 if (getenv("LD_DEBUG_ENV_DUMP") != NULL) {
107     printf("Windows original main args:\n");
108     for (i = 0; i < argc; i++) {
109         printf("main_arg[%d] = %s\n", i, argv[i]);
110     }
111 }
112
113 JLI_CmdToArgs(GetCommandLine());
114 margc = JLI_GetStdArgs();
115 // add one more to mark the end
116 margv = (char **)JLI_MemAlloc((margc + 1) * (sizeof(char *)));
117 {
118     int i = 0;
119     StdArg *stdargs = JLI_GetStdArgs();
120     for (i = 0; i < margc; i++) {
121         margv[i] = stdargs[i].arg;
122     }
123     margv[i] = NULL;
124 }
125 #else /* *NIXES */
126 {
127     // accommodate the NLL at the end
128     JLI_List args = JLI_List_new(margc + 1);
129     int i = 0;
130     for (i = 0; i < margc; i++) {
131         JLI_List argsInFile = JLI_PreprocessArg(argv[i]);
132         if (argsInFile != NULL) {
133             JLI_List_addargs, JLI_StringDup(argv[i]);
134         } else {
135             int cnt, idx;
136             cnt = argsInFile->size;
137             for (idx = 0; idx < cnt; idx++) {
138                 JLI_List_addargs, argsInFile->elements[idx];
139             }
140             // Shallow free, we reuse the string to avoid copy
141             JLI_MemFree(argsInFile->elements);
142             JLI_MemFree(argsInFile);
143         }
144     }
145     margc = args->size;
146     // add the NLL pointer at argv[margc]
147     JLI_List_add(args, NULL);
148     margv = args->elements;
149 }
150 #endif /* *NIXES */
151 printf("line1 from the test_java main()\n");
152 return JLI_launch(margc, margv,
153                sizeof(const_jargs) / sizeof(char *), const_jargs,
154                sizeof(const_appclasspath) / sizeof(char *), const_appclasspath,
155                VERSION_STRING,
156                DOT_VERSION,
157                (const_progname != NULL) ? const_progname : "javaw",
158                (const_launcher != NULL) ? const_launcher : "javaw",
159                HAS_JAVA_ARGS,
160                const_cpwildcard, const_javaw, const_ergo_class);
161 }
main.c 161L, 5470C

```

```

/bin/bash 122x59
Updating javac due to 1 file(s)
Updating javah due to 1 file(s)
Updating serialver due to 1 file(s)
Updating jimage due to 1 file(s)
Updating jshell due to 1 file(s)
Updating jhsdb due to 1 file(s)
Updating jar due to 1 file(s)
Updating jarsigner due to 1 file(s)
Updating javadoc due to 1 file(s)
Updating jinfo due to 1 file(s)
Updating jmap due to 1 file(s)
Updating jps due to 1 file(s)
Updating jstack due to 1 file(s)
Updating jstat due to 1 file(s)
Updating jconsole due to 1 file(s)
** Starting verify-modules
Checking dependencies across JDK modules
Updating support/src.zip
Creating jre image
Creating jdk image
duplicate resource "META-INF/services/sun.jvmstat.monitor.MonitoredHostService", skipping
duplicate resource "META-INF/services/com.sun.jdi.connect.Connector", skipping
duplicate resource "META-INF/services/sun.jvmstat.monitor.MonitoredHostService", skipping
Access verification succeeded.
## Finished verify-modules (build time 00:00:13)
Finished building target 'all' in configuration 'linux-x86_64-normal-server-release'
vmineel@297x:~/9dev$ cd build/linux-x86_64-normal-server-release/
buildtools/ hotspot/ jdk/ support/
configure-support/ images/ make-support/
vmineel@297x:~/9dev$ cd build/linux-x86_64-normal-server-release/
buildtools/ hotspot/ jdk/ support/
configure-support/ images/ make-support/
vmineel@297x:~/9dev$ cd build/linux-x86_64-normal-server-release/images/
dcsx_jdk/ jre/ test/
vmineel@297x:~/9dev$ cd build/linux-x86_64-normal-server-release/images/jdk/bin/
vmineel@297x:~/9dev/build/linux-x86_64-normal-server-release/images/jdk/bin$ ./java
Usage: java [options] class [args...]
       or java [options] -jar jarfile [args...]
       (to execute a class)
       or java [options] -jar jarfile [args...]
       (to execute a jar file)
where options include:
-d32      use a 32-bit data model if available
-d64      use a 64-bit data model if available
-server   to select the "server" VM
           The default VM is server,
           because you are running on a server-class machine.
-cp <class search path of directories and zip/jar files>
-classpath <class search path of directories and zip/jar files>
           A : separated list of directories, JAR archives,
           and ZIP archives to search for class files.
-Dname=value set a system property
-verbose:[class|gc|jni]
           enable verbose output

```

Figure 8: Inside main.c

## 5 References

1. The OpenJDK Developers' Guide - <http://openjdk.java.net/guide/>