



HAL
open science

I-JVM: a Java Virtual Machine for Component Isolation in OSGi

Nicolas Geoffray, Gaël Thomas, Gilles Muller, Pierre Parrend, Stéphane
Frénot, Bertil Folliot

► **To cite this version:**

Nicolas Geoffray, Gaël Thomas, Gilles Muller, Pierre Parrend, Stéphane Frénot, et al.. I-JVM: a Java Virtual Machine for Component Isolation in OSGi. [Research Report] RR-6801, INRIA. 2009, pp.21. inria-00354580

HAL Id: inria-00354580

<https://inria.hal.science/inria-00354580>

Submitted on 20 Jan 2009

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET EN AUTOMATIQUE

I-JVM: a Java Virtual Machine for Component Isolation in OSGi

Nicolas Geoffray — Gaël Thomas — Gilles Muller — Pierre Parrend — Stephane Frénot —
Bertil Folliot

N° 6801

Janvier 2009

Thème COM



*R*apport
de recherche

ISRN INRIA/RR--6801--FR+ENG

ISSN 0249-6399

I-JVM: a Java Virtual Machine for Component Isolation in OSGi

Nicolas Geoffray, Gaël Thomas, Gilles Muller , Pierre Parrend ,
Stephane Frénot , Bertil Folliot

Thème COM — Systèmes communicants
Équipes-Projets Regal, Ares

Rapport de recherche n° 6801 — Janvier 2009 — 21 pages

Abstract: The OSGi framework is a Java-based, centralized, component oriented platform. It is being widely adopted as an execution environment for the development of extensible applications. However, current Java Virtual Machines are unable to isolate components from each other. For instance, a malicious component can freeze the complete platform by allocating too much memory or alter the behavior of other components by modifying shared variables.

This paper presents I-JVM, a Java Virtual Machine that provides a lightweight approach to isolation while preserving compatibility with legacy OSGi applications. Our evaluation of I-JVM shows that it solves the 8 known OSGi vulnerabilities that are due to the Java Virtual Machine. Overall, the overhead of I-JVM compared to the JVM on which it is based is below 20%.

Key-words: Java, Isolation, Thread Migration, Resource Accounting, OSGi

I-JVM: une machine virtuelle Java pour l'isolation de composants dans OSGi

Résumé : OSGi est une plateforme orientée service implémentée en Java qui est de plus en plus utilisée pour le développement d'applications extensibles. Cependant, les machines virtuelles Java existantes ne sont pas capables d'isoler des composants entre eux. Par exemple, un composant malicieux peut bloquer l'exécution de la plateforme en allouant trop de mémoire ou modifier le comportement d'autres composants en modifiant des variables globales.

Nous présentons I-JVM, une machine virtuelle Java qui offre une isolation légère entre composants tout en préservant la compatibilité avec les applications OSGi existantes. I-JVM résout les 8 vulnérabilités connues sur la plateforme OSGi liées à la machine virtuelle, et ne diminue que de 20% les performances des applications en comparaison avec la machine virtuelle sur laquelle elle est implémentée.

Mots-clés : Java, Isolation, Migration de threads, Comptage de ressources, OSGi

1 Introduction

The OSGi framework [24] is a Java-based component platform which is being widely adopted as an execution environment for the development of extensible applications, such as Eclipse [2] or Java Enterprise Servers [3]. Extensibility in the OSGi platform is provided through a deployment unit called a *bundle* which groups together a set of components that are loaded through a specific Java class loader. The OSGi platform is popular because it provides modularity while still providing efficient communication through direct method calls between components.

Initially, the OSGi platform was designed for environments where all bundles trust each other. Nowadays, it is also promoted for systems such as next generation Internet home gateways where third party services can be downloaded dynamically [28]. However, the OSGi platform cannot protect a bundle against another malicious or buggy bundle. First, `java.lang.Class` objects, strings and static variables are shared in the Java Virtual Machine (JVM). The corruption of any one of these entities by a malicious or buggy bundle will impact all bundles. Second, a thread can freeze the JVM and deny service by exhausting memory or monopolizing the CPU. Third, it may be impossible to stop a bundle denying service, which makes a shut down of the entire platform the only solution. Not surprisingly, a recent work has identified 25 different vulnerabilities in current implementations of the Java/OSGi platform that may either lead to a violation of data integrity or a freeze of the platform [27]. While 17 of these vulnerabilities are due to a weak implementation of the OSGi framework itself and can be solved by adding suitable security checks, the remaining 8 originate in isolation issues and need to be solved at the JVM level.

Several approaches to providing isolation in a single JVM, through *Isolates* (or Java processes) have been recently introduced [5, 10, 16]. All of these solutions rely on duplicating the `java.lang.Class` objects, separating strings and static variables between isolates, and confining a thread to a single isolate. As a consequence, a communication between two isolates must be done using an RPC-like mechanism, which involves parameter copying and in some cases thread synchronization. Since the OSGi platform uses communication between bundles heavily, using RPCs would induce a non negligible overhead on the whole system performance.¹ Additionally, copying parameters implies modifying legacy bundles either at the source or bytecode level [31], that could compromise compatibility with legacy bundles.

This paper presents I-JVM, a Java Virtual Machine with lightweight isolates that is specifically designed to support the needs of the OSGi platform by associating each bundle with a separate isolate. The key contribution of I-JVM is to permit thread migration between isolates in order to keep the cost of an inter-isolate method call low. This enables complete bytecode compatibility with legacy OSGi bundles by avoiding the need to rewrite inter-bundle method calls. The main features of I-JVM are:

- Memory isolation. As shown by the Multi-Tasking Virtual Machine (MVM) [10], making `java.lang.Class` objects, strings and static variables pri-

¹Despite using a highly optimized virtual machine, the developers of Singularity [13] report that a simple local remote procedure call through a shared heap takes about 2500 cycles while a direct call takes only about 10 cycles.

vate to an isolate is sufficient to ensure memory isolation in a single JVM. Therefore, an isolate cannot access an object from another isolate unless a reference is given explicitly through method invocation.

- Resource accounting. I-JVM keeps track of the current isolate in which a thread is running. This allows recording the amount of memory and CPU time spent within a isolate. These statistics enable an administrator to detect denial of service attacks from malicious bundles.
- Termination of isolates. When an isolate terminates, its classes should not be invoked anymore. In case a thread returns back to the terminating isolate, I-JVM modifies the stack so that an exception is raised and trapped at a lower stack level. All the objects referenced by the terminating isolate are reclaimed by the garbage collector, with the exception of objects shared with other bundles.

I-JVM has been developed by modifying the LadyVM Java Virtual Machine [15], a JVM specifically designed for easing experiments in Java. We have used I-JVM to run two legacy OSGi platforms: Felix [1] of the Apache Community and Equinox [2] of the Eclipse Project.

Overall the results of this study are:

- I-JVM solves the OSGi JVM-related weaknesses identified in [27]. We present 8 attacks that cover these weaknesses.
- I-JVM has a 16% overhead on inter-bundle calls. This is an order of magnitude better than the cost of an RPC call between two processes. Overall, the I-JVM slowdown is between 1% and 20% on a representative suite of macrobenchmarks.
- I-JVM requires the addition of only 650 lines of code to LadyVM. Implementing I-JVM features in a legacy JVM should not be much more complex.

The rest of the paper is structured as follows. Section 2 describes the vulnerabilities of the OSGi platform. Section 3 explains and discusses the design and implementation of I-JVM. Section 4 provides performance measurements of I-JVM, and evaluates its robustness against denial of service attacks. Section 5 describes related work. Section 6 concludes the paper.

2 Vulnerabilities of OSGi

Vulnerabilities in OSGi have been identified at three sources [27]: (i) at the underlying operating system level, (ii) at the OSGi platform level and (iii) at the JVM level. The first kind of vulnerability is due to the possibility of running native code either inside the JVM process or as a separate process. These vulnerabilities are enabled by JNI or the `Runtime.exec` Java call. The second kind of vulnerability is related to weaknesses in the OSGi run-time and can be solved by adding security checks in the OSGi implementation [27]. In this paper, we attempt to solve the third kind of vulnerability which targets the JVM platform.

JVM vulnerabilities can themselves be subdivided into three categories: (i) lack of isolation, (ii) lack of resource accounting, (iii) failure to terminate a bundle. In the rest of this section, we present a suite of 8 attacks that cover the previously reported JVM vulnerabilities in [26, 27]. Our experiments in Section 4 show that all attacks may corrupt, freeze or abort unprotected OSGi platforms.

Lack of isolation. As mentioned previously, *java.lang.Class* objects, strings and static variables are shared in the JVM by all bundles. A malicious bundle can alter static variables or lock shared objects, and therefore interfere with the execution of other bundles. We consider two representative attacks:

- A1 - Modification of a static variable: All bundles share static variables. Therefore a bundle can modify a public non-final static variable defined by either other bundles, the OSGi platform or the core Java System Library. For example, a malicious bundle can set a shared variable to null, thus prevent the correct execution of other bundles. Bundles can discover static variables from other bundles either at compilation, or at runtime with the reflection API of Java.
- A2 - Synchronized method or synchronized call block: a bundle can lock shared strings, *java.lang.Class* objects or static variables, which can eventually freeze the system.

Lack of resource accounting. JVMs implement a bundle by using a specific class loader. However, JVMs do not perform resource accounting on a per class loader basis. In case of the over-use of resources, it is impossible to identify the faulty bundle and stop execution of its code. Resource accounting would help detecting the following five denial of service attacks:

- A3 - Memory exhaustion: a malicious bundle consumes most of the memory by holding references to many or large objects. This leads to an *OutOfMemoryError* for other bundles.
- A4 - Standalone infinite loop: a malicious bundle consumes all CPU resources by entering an infinite loop.
- A5 - Excessive object creation: a malicious bundle repetitively allocates objects without referencing them, thus triggering garbage collection and object finalization, which monopolize the CPU.
- A6 - Excessive thread creation: a malicious bundle crashes the platform by exceeding the number of threads supported.
- A7 - Hanging thread: a malicious bundle blocks when being called, thus never returning to the caller.

Bundle termination. In some situations, the JVM is unable to unload a bundle and deallocate all its allocated objects. First, other bundles may continue to reference the bundle, thus preventing the JVM from unloading the bundle classes, because a call to methods defined by the classes of the bundle can still occur. Second, if the OSGi runtime recognizes a bundle as misbehaving and wants to stop its execution, methods of the bundle may be executing or be in the call stack of running threads.

- A8 - Lack of termination support: a malicious bundle continues execution even if the OSGi platform tries to unload it.

3 I-JVM Design and Implementation

The design goal of I-JVM is to provide bundle isolation while preserving the communication model of the OSGi platform, which relies on direct method calls. An OSGi application is composed of a set of dynamically loaded bundles and of the OSGi runtime itself. To implement isolation, each bundle is executed within a separate isolate. Additionally, the OSGi runtime runs in a specific isolate, `Isolate0`, which has higher rights than standard bundles.

I-JVM permits explicit object sharing between isolates by passing an object reference in an inter-isolate method call. The key point of I-JVM is to provide thread migration between isolates in a single address space, which is a prerequisite for object sharing. Different isolates execute therefore on the same execution stack, which impacts isolation, resource accounting and termination. In this Section, we present I-JVM in detail. We focus on the main issues: isolation, resource accounting and termination of isolates. Then, we describe how to run the OSGi platform using I-JVM. Finally, we report the implementation of I-JVM in LadyVM [15].

3.1 Isolation and Thread Migration

I-JVM runs isolates in the same address space. Isolates provide a lightweight protection mechanism integrated in the JVM, so that the classes running in one isolate cannot crash classes running in another independent isolate. In I-JVM, an isolate is built from a class loader, so its scope is the classes loaded by the class loader. There is a specific isolate, `Isolate0`, which has higher rights on the platform than standard isolates. These rights are the permissions to start a new isolate, to terminate an isolate and to shut down the entire Java platform. The first Java class loader created becomes `Isolate0`. The subsequent class loaders are standard isolates.

An inter-isolate method call induces a thread migration. Each thread possesses a reference to the isolate in which it is currently running. When a thread calls a method in another isolate, I-JVM sets the thread's isolate reference to the called isolate. When the thread leaves the callee, I-JVM sets the reference back to the caller's isolate. When an isolate calls a method in the same isolate or in the core Java System Library, the isolate reference is not changed. Classes from the Java System Library are not executed in a special isolate but in the isolate that called it.

To implement isolates, the main change from the JVM specification is to have a per-isolate private copy of static variables, strings and *java.lang.Class* objects. This is sufficient to ensure that an isolate does not have access to the internal state of another isolate since: (i) an isolate cannot construct a foreign reference thanks to the type safety of the Java bytecode; (ii) an isolate cannot access an isolate private field or method thanks to the scope of fields and methods in the Java bytecode.

Isolation of static variables is done by associating a task class mirror [10] per class. The task class mirror of a class contains the initialization state of

the class, the static variables and the associated *java.lang.Class* object. I-JVM uses the current isolate reference of the thread as an index into the array of task class mirrors of a class. Accessing a static variable requires fetching the execution environment, loading the isolate reference from the execution environment, loading the task class mirror and finally loading the static variable. Compared to simply loading a static variable, the task class mirror approach requires two additional loads. An isolate also always has to check the initialization state of a class before accessing one of its static variable or before calling a static method. Like in MVM, the just in time compiler cannot remove all of the class initialization checks, because the code compiled must be reentrant [10].

3.2 Resource Accounting

I-JVM monitors the execution of isolates. Resources consumed in the code of an isolate are charged to the isolate. Resources consumed in a method of the Java System Library are charged to the caller of the method. I-JVM counts the CPU time consumed, the objects allocated, the number of threads created, the number of connections used, the number of bytes read or written through I/O connections, and the number of garbage collection activations.

Memory and connections: When an isolate allocates an object, I-JVM charges the object to the isolate. If the object is a connection (FileDescriptor or Socket), the connection is also charged to the creator. If the isolate gives this object to another isolate through an inter-isolate call (i.e. the object becomes shared), I-JVM does not immediately update the memory usage of the callee. An update would require call barriers and write barriers, which would drastically degrade performance. We also do not divide the resource consumption of shared objects or connections between isolates because doing so would require maintaining a list of isolates that use the shared object, thus would introduce a new list traversal for all objects during garbage collection.

To update the amount of memory held by an isolate, I-JVM relies on the Garbage Collector (GC). Besides collecting unreferenced objects, the GC runs the following algorithm for memory and connection accounting:

1. The amount of memory and connections used per isolate is reinitialized to zero.
2. For each isolate, the GC adds its list of strings, static objects and *java.lang.Class* objects to its root objects.
3. For each thread, the GC inspects the execution frames on their stack. For an execution frame, the GC knows the Java method that created the frame, hence its isolate. The GC adds therefore all the object referenced by the frame to the isolate in which the frame is executed. It does not inspect frames from the core Java System Libraries methods because the objects referenced in the frame are also referenced by the isolate that called the library.
4. The GC traces the roots of isolates (i.e., the objects identified in steps 2. and 3.). An object is charged to the first isolate that references it.

While unprecise, we think the approach is sufficient because it gives a rough estimation of memory consumption per bundle in the presence of object sharing. We leave as future work improvements on the precision of memory accounting.

Threads: Threads are created within an isolate. I-JVM counts the number of threads an isolate creates. Note that threads are charged to their creator, but may execute code from any isolate via inter-bundle calls.

I/O reads and writes: Any read or write to a connection is instrumented in order to charge the isolate performing the operation. The approach is similar to JRes' accounting of network resources [11]: there are few classes that perform read and writes on connections, and instrumenting them is straightforward.

CPU time: For CPU accounting, we have studied two solutions. The first one is to insert per-isolate time updates during an inter-isolate call. This induces a significant performance penalty because it requires: (i) two system calls to fetch the current time, one when entering and one when leaving the isolate and (ii) a lock acquisition to update the CPU time when leaving the call. Therefore, we chose a second solution in which I-JVM counts the CPU time spent in isolates by regularly sampling the value of the isolate reference of a running thread.

Garbage collection: I-JVM uses a single GC for all isolates. To detect attacks on garbage collection activations, I-JVM counts the number of times an isolate triggers the GC.

3.3 Isolate Termination

There are two main problems when terminating an isolate. First, threads migrate between isolates; therefore I-JVM cannot just kill the threads created by the isolate. Also, a thread created by another isolate may be executing code from the terminating isolate. Second, shared objects referenced by other isolates cannot be released.

When terminating an isolate, I-JVM stops the execution of all threads by sending them a signal. The handler of the signal is defined at startup and cannot be modified by isolates, thanks to the Java language. Upon receiving the signal, a thread inspects and modifies its stack as follows. For each frame, if it is called from a frame that belongs to the terminating isolate, the thread changes the return pointer to throw a `StoppedIsolateException` exception. The terminating isolate cannot catch this exception: even if the isolate tries to catch it in the Java code, I-JVM will ignore it. I-JVM also takes a special action for the last frame:

- If it belongs to the Java System Library, I-JVM sets the interrupted flag of the thread so that I/O or sleep calls are interrupted. This approach is similar to protection domain termination in Spring [17].
- If it belongs to the terminating isolate, the thread throws the `StoppedIsolateException`.

Moreover, code from the terminating isolate should not be called anymore, so that the isolate does not attempt to deny service once again. I-JVM prevents execution of the isolate by (i) not JIT compiling the methods not JITed yet and (ii) modifying the code of the already compiled methods. At the beginning of each of these methods, I-JVM inserts a branch to a function that throws the `StoppedIsolateException`.

An isolate is only removed from memory when there is no remaining object whose class is defined by the isolate.

3.4 Running OSGi on I-JVM

An OSGi application is made of a set of dynamically loaded bundles and the OSGi runtime itself. To implement protection, each bundle is executed within a separate isolate; the OSGi runtime runs in the `Isolate0`.

`Isolate0` is associated with the applicative class loader that loads the `main` function of the OSGi framework. When OSGi loads a new bundle, it allocates a new class loader. I-JVM associates therefore a standard isolate to this class loader. This makes it possible to run a legacy OSGi runtime on I-JVM without any modification. The `start` method of a bundle receives an object that represents OSGi. This object is the first shared object between bundles. It is used in OSGi to register object references in a name service and to find foreign references. Hence, at startup, a bundle can only access foreign objects in this name service.

We additionally define a few rules that an OSGi runtime must follow:

1. It should create a new thread when calling the `start` and `stop` methods of a bundle, in order to prevent a malicious bundle from freezing the OSGi runtime.
2. It must use Java permissions to deny access of privileged resources to bundles. For example, the JVM allows Java applications to run non-Java code through the use of the JNI interface or the `Runtime.exec` call. This gives a bundle the possibility to run unverified code that could destroy the OSGi platform. A second example is the `System.exit` method which shuts down the JVM.
3. It should send a `StoppedBundleEvent` to all bundles when a bundle is being killed. A bundle can then take any action it desires: it can ignore the event or may release the references it had on the terminating bundle objects. If the bundle does not release the references, I-JVM may charge the objects to the bundle. The key point is that resources from the terminating bundle will not be released until all bundles release their references to them.

We also define one rule for writing bundles:

1. Bundles should be prepared to catch any kind of exceptions when calling other bundle methods. Like any regular JVM, I-JVM uses exceptions to signal an error during the execution of an isolate. Such errors include regular errors such as erroneous class files or null pointer exceptions, but also I-JVM specific errors when the isolate is being killed.

3.5 I-JVM Implementation

We have implemented I-JVM in the LadyVM virtual machine [15] which is specification compliant with the JVM. Overall, I-JVM required the addition of 650 lines of code to LadyVM which are distributed as follows:

- Static variables, strings and `java.lang.Class` objects: 200 lines of code for implementing the task class mirror in each class. The changes are done at two levels: (i) in the Java class representation, which contains the task class mirror, (ii) in the bytecode translator to modify the accesses to static variables and `java.lang.Class` objects to reference the task class mirror.

- Method call: 150 lines of code for the update of the isolate reference in each isolate method compiled by LadyVM.
- Resource accounting: 100 lines of code for accounting, for CPU, memory, I/O and threads.
- Isolate per bundle: 50 lines of code to create and attach a new isolate to a class loader when the latter first loads a class.
- Isolate termination: 150 lines of code for the termination of isolates.

These numbers are quite low and show that implementing I-JVM within a legacy JVM should be quite easy.

While we rely on the Java language for memory protection between bundles, we changed the equality feature of strings in the JVM for the purpose of isolation between bundles. The Java language assumes that strings in class files are hashed. In I-JVM, each bundle has its map of strings, therefore the `==` operator does not work for strings allocated by different bundles. Programmers should use the `equals` function instead. The issue is also raised in KaffeOS [5].

4 Evaluation

In this section, we first motivate the need for fast inter-bundle calls in OSGi by benchmarking a simple OSGi application. Then, we evaluate I-JVM, in terms of performance and memory overhead, and in terms of robustness against attacks. Finally, we discuss the limitations of our resource accounting scheme.

4.1 Inter-bundle Calls

To motivate the approach of OSGi and the need for fast inter-bundle calls, we evaluated the application demo provided by Felix [1]. The application is a paint program architected with bundles. The drawing area, as well as the shapes that can be drawn are implemented as bundles. The user can add shapes dynamically or remove them through the OSGi bundle manager.

We measured the number of inter-bundle calls when I-JVM executes the application. Each time a shape is dragged on the drawing area, an inter-bundle call happens between the drawing area and the shape. When the user moves the shape in the drawing area, dragging and moving the shape from upper-left to the bottom-right makes roughly two hundred inter-bundle calls. Table 1 shows the time for performing two hundred inter-bundle calls depending on the communication implementation. We evaluate four kinds of implementations: (i) local call, (ii) RMI call, which is the standard inter-application communication in Java, (iii) Incommunicado [25], the communication model of Isolates, and (iv) I-JVM. I-JVM is an order of magnitude faster than other approaches for inter-isolate communication.

4.2 Performance Overhead of I-JVM

Since LadyVM has been designed for easing experimenting with virtual machines, it does not compete with industrial JVMs. We thus report our ex-

	Local method	RMI local call	Incommunicado	I-JVM
Time	20 μ s	90ms	9ms	24 μ s

Table 1: Cost of 200 inter-bundle calls, depending on the communication model. The benchmarks were measured on a Pentium D 3.0GHz with 3GB of memory. Incommunicado is reported to be ten times faster than RMI [25].




Figure 1: Performance of I-JVM for the micro-benchmarks, relative to LadyVM.

periments in terms of relative performance and memory overheads of I-JVM compared to LadyVM.

All experiments were done on a Pentium D 3.0GHz with 3GB of memory running Mandriva Linux 2.6.23.

To evaluate the overhead of isolation and resource accounting in I-JVM, we have run the following set of micro-benchmarks: intra-isolate and inter-isolate calls, object allocation and access to static variables. We measured the overhead performing the same operation a million times. We also ran the SPEC JVM98 benchmark in an isolate to measure the overall runtime cost of I-JVM. Finally, we evaluated the memory overhead induced by running two legacy implementations of OSGi, Felix [1] and Equinox [2] on top of I-JVM.

Figure 1 shows the relative performance of I-JVM compared to LadyVM for the micro-benchmarks. I-JVM adds two test instructions when executing an intra-isolate method call. For an inter-isolate call, it also updates the current isolate of the thread, thus adding four more store operations. Overall, an intra-bundle call induces a 14% overhead and an inter-bundle call induces a 16% overhead.

We also benchmarked the performance of object allocation by repetitively allocating a *java.lang.Object* object. In LadyVM and I-JVM, the size of such an object is 28 bytes. The results show that there is an 18% overhead compared

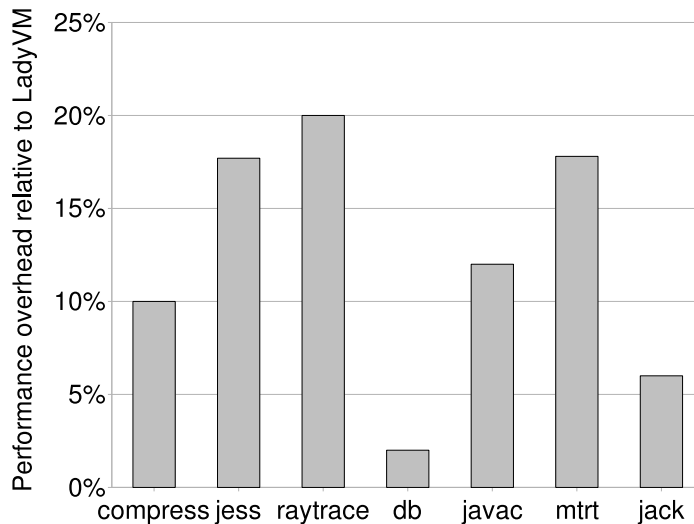


Figure 2: Overhead of I-JVM for the Spec JVM98 benchmarks, relative to LadyVM.

to LadyVM, due to resource accounting, testing the memory limit when an isolate allocates an object and the intra-bundle cost of calling the *java.lang.Object* constructor. Finally, we measured static variable access. We removed all compilation optimizations in I-JVM to exhibit the cost of one access to the static variable. The benchmark shows that accessing a static variable gives a 46% overhead penalty on I-JVM. This is due to the task class mirror requiring two loads more than a simple load of a static variable plus an initialization check to verify that the static variable has been allocated. When I-JVM runs with all compilation optimizations, the overhead of accessing a static variable a million times is below 1% because the extra instructions execute only once.

We measured the execution time of SPEC JVM98 benchmarks [4] running within Isolate0, so as to evaluate the overhead of I-JVM for standard Java programs. Figure 2 shows that the overhead of I-JVM is below 20% for all benchmarks. By comparison with MVM that reports a maximum overhead of 10% [10], I-JVM is less efficient. The main reason is that I-JVM induces an overhead for resource accounting and the test during intra-isolate calls. However, the cost of inter-isolate communication in MVM is an order of magnitude higher than a simple method call [25].

Finally, we measured the memory overhead induced by I-JVM when running an OSGi implementation. There are two places where I-JVM requires more memory than a standard JVM: (i) the array of task class mirrors for each class and (ii) a per-isolate set of strings and statistics information. Figure 3 shows the memory used when running the base configurations of Felix and Equinox. Felix runs the OSGi runtime and three management bundles (administration, shell, repository). Equinox runs the OSGi runtime and twenty-two management bundles. Overall, the memory overhead for both OSGi implementations is below 16%.




Figure 3: Memory consumption of I-JVM and LadyVM on OSGi implementations.

4.3 Robustness Evaluation

In Section 2, we presented a set of attacks to measure the robustness of JVM implementations running an OSGi platform. In this section, we compare the result of executing these attacks on I-JVM and on the Sun JVM. I-JVM prevents the eight kinds of attacks by relying on an administrator. With the Sun JVM, the administrator loses control of the platform and has no possibility of stopping the execution of bundles, even if he detects the offending bundles. With the Sun JVM, the platform freezes or aborts under denial of service attacks.

A1 - Store mutable object in static variable. Bundle A defines an array as a static variable and works on the elements of the array. Bundle B finds the static variable either at compile-time, or at runtime through the reflection B. B sets to null the contents of the array.

Result with Sun JVM: Bundle A throws a *NullPointerException*.

Result with I-JVM: I-JVM isolates bundles so that they cannot access another bundle's static variables. The array is duplicated, therefore the modifications made by bundle B are local to bundle B. Bundle A continues to work on the array.

A2 - Synchronized method or synchronized call block. Bundle A calls a static *synchronize* method defined in a class that belongs to the bundle. Bundle A therefore synchronizes on the *java.lang.Class* object of the class. Another bundle B explicitly synchronizes on the object, holding the object forever.

Result with Sun JVM: Bundle A is blocked.

Result with I-JVM: I-JVM disallows the sharing of strings, *java.lang.Class* objects and static variables, therefore there is no interference between bundles

that do not communicate with each other. Therefore, bundle A continues to run.

A3 - Memory exhaustion : A set of bundles are running on the platform. The OSGi runtime dynamically installs a new one that allocates many objects and stores them in an array, thus preventing the GC from deallocating the objects.

Result with Sun JVM: All bundles get a *OutOfMemoryError* when allocating a new object.

Result with I-JVM: I-JVM counts the memory used by each bundle. Based on this information, the administrator kills the offending bundle and all other bundles continue to run.

A4 - Exponential object creation. A set of bundles are running on the platform. The OSGi runtime dynamically installs a new one that allocates many objects but does not keep a reference to them, thus triggering the GC many times.

Result with Sun JVM: The JVM spends its time garbage collecting. The non-offending bundles make progress slowly.

Result with I-JVM: I-JVM counts the number of times a bundle runs a GC. Based on this information, the administrator kills the offending bundle and all other bundles continue to run.

A5 - Recursive thread creation. A set of bundles are running on the platform. The OSGi runtime dynamically installs a new one that endlessly creates threads.

Result with Sun JVM: All bundles get a *OutOfMemoryError* when allocating a new object or a new thread.

Result with I-JVM: I-JVM counts the number of threads a bundle creates. Based on this information, the administrator kills the offending bundle and all other bundles continue to run.

A6 - Standalone infinite loop. A set of bundles are running on the platform. The OSGi runtime dynamically installs a new one that runs an infinite loop.

Result with Sun JVM: The non-malicious bundles make progress slowly.

Result with I-JVM: I-JVM counts the CPU usage of each bundle. Based on this information, the administrator kills the offending bundle and all other bundles continue to run.

A7 - Hanging thread. Bundle A calls a method of bundle B and bundle B calls *Thread.sleep(0)*.

Result with Sun JVM: Execution never returns to bundle A.

Result with I-JVM: I-JVM inspects the current bundle of each thread and counts the number of sleeping threads in a bundle. Based on this information, the administrator kills the bundle that called *Thread.sleep*. If bundle A was prepared to catch the *StoppedIsolateException*, execution returns to A. Otherwise, the exception is caught at a lower stack level.

A8 - Lack of termination support. Bundle A calls bundle B and is waiting for an object value. Bundle B returns an object that points to the internal representation of bundle B in the OSGi platform. Bundle A stores the reference in one of its variable. Bundle B then makes a denial of service attack. Therefore, the administrator tries to unload bundle B.

Result with Sun JVM: The OSGi platform is unable to unload the bundle, and the attack continues to run.

Result with I-JVM: All threads that execute code from bundle B throw an exception and execution never returns to bundle B.

4.4 Automatic Denial of Service Attacks Detection

We have shown that the resource accounting in I-JVM allows an administrator to locate misbehaving bundles. However CPU time, the number of collection activations and the amount of memory cannot in the current design be used to automatically kill these bundles. The following experiments show the limits of our approach:

1. CPU time. A malicious bundle M calls a function of another bundle A a million times. Since I-JVM regularly samples the current isolate of a thread, it randomly charges the CPU to the caller or the callee. At the end of the experiment, I-JVM charged roughly 75% of the CPU to A and 25% to M . Since the callee updates the current isolate, it executes more code than the caller, which explains the CPU distribution of our experiment.
2. Garbage collection. We changed the function implemented in A to allocate and return a new object. Since, M is calling A a million times, a garbage collection is triggered on behalf of A .
3. Memory accounting. Now bundle M implements a function that returns a large object (100M). The bundle is supposed to implement a well-defined interface (in our experiment a dictionary service) and is called by other bundles. When a garbage collection happens, the garbage collector does not charge the large objects to M but to the callers of M .

These examples show that our resource accounting is not as precise as process-based resource accounting. However, resource accounting in the presence of object sharing and thread migration requires a trade-off between preciseness and efficiency. We leave as future work improvements on better resource accounting.

5 Related Work

Our approach combines ideas from different research areas. In this section, we report how our work relates to (i) operating system structure for resource management and communications and (ii) resource accounting and isolation in JVMs.

5.1 Operating Systems

Resource management and isolation are usually the responsibility of operating systems. The process abstraction is the means of isolation and resource accounting. One can set resource limits for a process and for its children. Processes communicate with each other and with the kernel through system calls and arguments are copied between user space and kernel space.

Micro-kernel operating systems differ from monolithic operating systems by placing the various modules of a kernel in different protection domains. Each module has its own page table and does not see other modules (except the kernel itself). The kernel and modules communicate through Inter-Process Communication (IPC) primitives. Micro-kernels require fast IPCs between processes in order to achieve competitive performance compared to monolithic kernels. Techniques like LRPC [7], thread migration [14] and continuations [12] reduce the cost of IPCs by using the same thread: they jump directly between the caller and the callee without involving the scheduler. Also, projects like L4 [23] achieve a high level of IPC performance by fine tuning IPCs with hardware techniques.

Single address space operating systems like Opal [9] and Mungi [19] share one global virtual address space among processes. But each process can only access its own memory thanks to memory protection, hence parameters are copied from the caller to the callee. Other single address space operating systems such as Spin [8], JX [16], JavaOS [29], or Singularity [20] base their protection on the type-safety of the language. However all these operating systems enforce isolation by avoiding direct procedure calls and by using IPCs through shared heaps or portals. Both cases break compatibility with existing OSGi applications.

The Scout operating system [30], Rialto [21], and resource containers [6] introduce new approaches for resource accounting in operating systems. Threads in these systems are not bound to a protection domain but migrate between protection domains while charging resources to a single resource management entity. These systems differentiate users and protection domains. In OSGi a protection domain (ie a bundle) *is* a user. Moreover, a protection domain switch has the same cost than a standard IPC.

5.2 Isolation in Java Virtual Machines

The standard isolation mechanism in current JVMs is based on class loaders [22]. Class loaders provide name space isolation, i.e. there is no collision between classes with the same name but loaded by two different class loaders. Class loaders have weak isolation guarantees, as they still share static variables, interned strings and *java.lang.Class* objects.

There are many projects that attempt to give operating system features to the JVM. The J-Kernel [18] provides multiple protection domains in Java in which domains communicate through capabilities, which have similar costs to regular IPCs. Its sister project, JRes [11] adds resource accounting in the Java platform on a per process basis. KaffeOS [5] is an extended JVM that executes multiple applications in the same virtual machine. Applications communicate through the use of a shared heap which prevents references to local heaps. The Multi-Tasking Virtual Machine (MVM) [10] is also an extended JVM that executes multiple applications, but without a shared heap. Isolate communication

in MVM is implemented with Links or Incomunicado [25] which are an order of magnitude less efficient than simple method calls. While the J-Kernel, KaffeOS and MVM provide safe isolate termination, they rely on expensive inter-isolate communications.

OSGi uses Java permissions to enhance security and limit the rights of a bundle. When a method wants to know the protection level of its caller, the JVM inspects the stack trace to find which class loader loaded the class of the caller. A class loader can be given a security policy, which indicates what kind of privileged operations the classes loaded by this class loader can perform. Inspection of a stack is an expensive operation, which would dramatically reduce the performance of our approach if it were performed on each access to a static variable, string or `java.lang.Class`.

6 Conclusion

We have described the design and implementation of I-JVM, a Java Virtual Machine extended with component isolation and termination in OSGi. I-JVM enables a lightweight isolation of OSGi bundles while still providing fast communication through thread migration across bundles and direct sharing of objects. The isolate architecture of I-JVM allows a per-bundle resource accounting that an administrator can use to terminate a misbehaving bundle. Even though isolation and resource accounting has a small overhead compared to a regular JVM, our evaluation shows that I-JVM is able to inform denial of service attacks to an administrator and stop their execution.

In this paper, we considered resource accounting as an assistance for an administrator to locate possible resource problems and kill the bundles he thinks are malicious. We plan as future work to improve the preciseness of resource usage. Still, we believe that the bundle isolation and termination features of I-JVM are essential features for the management of current and future OSGi platforms.

7 Availability

I-JVM is publicly available via an open-source license at the URL:

<http://vmlkit.llvm.org>

References

- [1] Apache felix. <http://felix.apache.org/site/index.html>.
- [2] Equinox. <http://www.eclipse.org/equinox>.
- [3] Jonas J2EE Server. <http://jonas.objectweb.org>.
- [4] SPECjvm98. <http://www.spec.org/jvm98/>.

-
- [5] G. Back, W. H. Hsieh, and J. Lepreau. Processes in KaffeOS: isolation, resource management, and sharing in Java. In *Proceedings of the Operating Systems Design and Implementation Symposium*, pages 333–346, San Diego, USA, October 2000. USENIX Association.
 - [6] G. Banga, P. Druschel, and J. Mogul. Resource containers: a new facility for resource management in server systems. In *Proceedings of the Operating Systems Design and Implementation Symposium*, pages 45–58, New Orleans, USA, February 1999. USENIX Association.
 - [7] B. Bershad, T. Anderson, E. Lazowska, and H. Levy. Lightweight remote procedure call. *Transactions on Computer Systems*, 8(1):37–55, 1990.
 - [8] B. Bershad, S. Savage, P. Pardyak, E. Sirer, M. Fiuczynski, D. Becker, C. Chambers, and S. Eggers. Extensibility safety and performance in the SPIN operating system. In *Proceedings of the Symposium on Operating Systems Principles*, pages 267–283, Copper Mountain, USA, December 1995. ACM.
 - [9] J. Chase, H. Levy, M. Feeley, and E. Lazowska. Sharing and protection in a single-address-space operating system. *ACM Transactions on Computer Systems*, 12(4):271–307, 1994.
 - [10] G. Czajkowski and L. Daynès. Multitasking without compromise: a virtual machine evolution. In *Proceedings of the Object Oriented Programming, Systems, Languages, and Applications Conference*, pages 125–138, Tampa Bay, USA, October 2001. ACM.
 - [11] G. Czajkowski and T. Eicken. JRes: a resource accounting interface for Java. In *Proceedings of the Object Oriented Programming, Systems, Languages, and Applications Conference*, pages 21–35, Vancouver, Canada, October 1998. ACM.
 - [12] R. Draves, B. Bershad, R. Rashid, and R. Dean. Using continuations to implement thread management and communication in operating systems. In *Proceedings of the Symposium on Operating Systems Principles*, pages 122–136, Pacific Grove, USA, October 1991. ACM.
 - [13] M. Fähndrich, M. Aiken, C. Hawblitzel, O. Hodson, G. Hunt, J. Larus, and S. Levi. Language support for fast and reliable message-based communication in Singularity OS. In *Proceedings of the EuroSys Conference*, pages 177–190, Leuven, Belgium, April 2006. ACM.
 - [14] B. Ford and J. Lepreau. Evolving Mach 3.0 to a migrating thread model. In *Proceedings of the USENIX Winter Technical Conference*, pages 97–114, San Francisco, USA, January 1994. USENIX Association.
 - [15] N. Geoffray, G. Thomas, C. Clément, and B. Folliot. A lazy developer approach: building a JVM with third party software. In *Proceedings of the International Symposium on Principles and Practice of Programming In Java*, pages 73–82, Modena, Italy, September 2008. ACM.

-
- [16] M. Golm, M. Felsera, C. Wawersich, and J. Kleinoeder. The JX operating system. In *Proceedings of the Usenix Annual Technical Conference*, pages 45–58, Monterey, USA, June 2002. USENIX Association.
- [17] G. Hamilton and P. Kougiouris. The Spring Nucleus: a microkernel for objects. In *Proceedings of the USENIX Summer Technical Conference*, pages 1–15, Cincinnati, USA, June 1993. USENIX Association.
- [18] C. Hawblitzel, C. Chang, G. Czajkowski, D. Hu, and T. von Eicken. Implementing multiple protection domains in Java. In *Proceedings of the USENIX Annual Technical Conference*, pages 259–270, New Orleans, USA, June 1998. USENIX Association.
- [19] G. Heiser, K. Elphinstone, J. Vochtelo, S. Russell, and J. Liedtke. The Mungi single-address-space operating system. *Software: Practice and Experience*, 28(9):901–928, 1998.
- [20] G. Hunt, M. Aiken, M. Fähndrich, C. Hawblitzel, O. Hodson, J. Larus, B. Steensgaard, D. Tarditi, and T. Wobber. Sealing OS processes to improve dependability and safety. In *Proceedings of the Eurosys Conference*, pages 341–354, Lisboa, Portugal, April 2007. ACM.
- [21] M. Jones, P. Leach, R. Draves, and J. Barrera. Modular real-time resource management in the Rialto operating system. In *Proceedings of the Hot Topics in Operating Systems Workshop*, pages 12–17, Washington, USA, May 1995. IEEE Computer Society.
- [22] S. Liang and G. Bracha. Dynamic class loading in the Java virtual machine. In *Proceedings of the Object Oriented Programming, Systems, Languages, and Applications Conference*, pages 36–44, Vancouver, Canada, October 1998. ACM.
- [23] J. Liedtke, K. Elphinstone, S. Schinberg, H. Hartig, G. Heiser, N. Islam, and T. Jaeger. Achieved IPC performance. In *Proceedings of the Hot Topics in Operating Systems Workshop*, pages 28–31, Washington, USA, May 1997. IEEE Computer Society.
- [24] OSGi Alliance. OSGi service platform, core specification release 4.1. Draft, 05 2007.
- [25] K. Palacz, J. Vitek, G. Czajkowski, and L. Daynès. Incommunicado: efficient communication for isolates. In *Proceedings of the Object-Oriented Programming, Systems, Languages, and Applications Conference*, pages 262–274, Seattle, USA, November 2002. ACM.
- [26] P. Parrend and S. Frénot. Classification of component vulnerabilities in Java service oriented programming platforms. In *Proceedings of the Component-Based Software Engineering Symposium*, pages 80–96, Karlsruhe, Germany, October 2008. Springer-Verlag.
- [27] P. Parrend and S. Frénot. Security benchmarks of OSGi platforms: toward hardened OSGi. *Software: Practice and Experience*, 2008. Accepted for publication.

- [28] Y. Royon and S. Frénot. Multiservice home gateways: business model, execution environment, management infrastructure. *IEEE Communications Magazine*, 45(10):122–128, October 2007.
- [29] T. Saulpaugh and C. Mirho. *Inside the JavaOS operating system*. Addison Wesley Longman, 1999.
- [30] O. Spatscheck and L. Peterson. Defending against denial of service attacks in Scout. In *Proceedings of the Operating Systems Design and Implementation Symposium*, pages 59–72, New Orleans, USA, February 1999. USENIX Association.
- [31] E. Tilevich and Y. Smaragdakis. J-Orchestra: automatic Java application partitioning. In *Proceedings of the European Conference on Object-Oriented Programming*, pages 178–204, Malaga, Spain, June 2002. Springer-Verlag.

Contents

1	Introduction	3
2	Vulnerabilities of OSGi	4
3	I-JVM Design and Implementation	6
3.1	Isolation and Thread Migration	6
3.2	Resource Accounting	7
3.3	Isolate Termination	8
3.4	Running OSGi on I-JVM	9
3.5	I-JVM Implementation	9
4	Evaluation	10
4.1	Inter-bundle Calls	10
4.2	Performance Overhead of I-JVM	10
4.3	Robustness Evaluation	13
4.4	Automatic Denial of Service Attacks Detection	15
5	Related Work	15
5.1	Operating Systems	16
5.2	Isolation in Java Virtual Machines	16
6	Conclusion	17
7	Availability	17



Centre de recherche INRIA Paris – Rocquencourt
Domaine de Voluceau - Rocquencourt - BP 105 - 78153 Le Chesnay Cedex (France)

Centre de recherche INRIA Bordeaux – Sud Ouest : Domaine Universitaire - 351, cours de la Libération - 33405 Talence Cedex
Centre de recherche INRIA Grenoble – Rhône-Alpes : 655, avenue de l'Europe - 38334 Montbonnot Saint-Ismier
Centre de recherche INRIA Lille – Nord Europe : Parc Scientifique de la Haute Borne - 40, avenue Halley - 59650 Villeneuve d'Ascq
Centre de recherche INRIA Nancy – Grand Est : LORIA, Technopôle de Nancy-Brabois - Campus scientifique
615, rue du Jardin Botanique - BP 101 - 54602 Villers-lès-Nancy Cedex
Centre de recherche INRIA Rennes – Bretagne Atlantique : IRISA, Campus universitaire de Beaulieu - 35042 Rennes Cedex
Centre de recherche INRIA Saclay – Île-de-France : Parc Orsay Université - ZAC des Vignes : 4, rue Jacques Monod - 91893 Orsay Cedex
Centre de recherche INRIA Sophia Antipolis – Méditerranée : 2004, route des Lucioles - BP 93 - 06902 Sophia Antipolis Cedex

Éditeur
INRIA - Domaine de Voluceau - Rocquencourt, BP 105 - 78153 Le Chesnay Cedex (France)
<http://www.inria.fr>
ISSN 0249-6399