

# Java虚拟机规范

(英文版)


## (Java SE 8版)

The Java Virtual Machine Specification (Java SE 8 Edition)

[美]

Tim Lindholm  
Frank Yellin  
Gilad Bracha  
Alex Buckley

著



· 原味精品书系 ·

# Java虚拟机规范 (英文版)

## (Java SE 8版)

The Java Virtual Machine Specification (Java SE 8 Edition)

[美] Tim Lindholm  
Frank Yellin 著  
Gilad Bracha  
Alex Buckley



电子工业出版社  
Publishing House of Electronics Industry  
北京·BEIJING

## 内容简介

本书由 Java 虚拟机技术创建人撰写,全面而准确地阐释了 Java 虚拟机各方面的细节,围绕 Java 虚拟机整体架构、编译器、class 文件格式、加载、链接与初始化、指令集等核心主题对 Java 虚拟机进行全面而深入的分析,深刻揭示 Java 虚拟机的工作原理。书中完整地讲述了由 Java SE 8 所引入的新特性,例如对包含默认实现代码的接口方法所做的调用,以及为支持类型注解及方法参数注解而对 class 文件格式所做的扩展等,还阐明了 class 文件中各属性的含义及字节码验证的规则。

本书基于 Java SE 8,是深度了解 Java 虚拟机和 Java 语言实现细节的极佳选择。

Original edition, entitled Java Virtual Machine Specification, Java SE 8 Edition, 9780133905908 by Tim Lindholm, Frank Yellin, Gilad Bracha, Alex Buckley, published by Pearson Education, Inc., publishing as Addison-Wesley, Copyright © 2014 Pearson Education, Inc.

All rights reserved. No part of this book may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording or by any information storage retrieval system, without permission from Pearson Education, Inc.

China edition published by Pearson Education Asia Ltd. and Publishing House of Electronics Industry Copyright © 2016. The edition is manufactured in the People's Republic of China, and is authorized for sale and distribution only in the mainland of China exclusively (except Hong Kong SAR, Macau SAR, and Taiwan).

本书英文影印版专有出版权由 Pearson Education 培生教育出版亚洲有限公司授予电子工业出版社。未经出版者预先书面许可,不得以任何方式复制或抄袭本书的任何部分。

本书仅限中国大陆境内(不包括中国香港、澳门特别行政区和中国台湾地区)销售发行。

本书英文影印版贴有 Pearson Education 培生教育出版集团激光防伪标签,无标签者不得销售。

版权贸易合同登记号 图字:01-2015-6102

### 图书在版编目(CIP)数据

Java 虚拟机规范:Java SE 8 版=The Java Virtual Machine Specification, Java SE 8 Edition:英文/(美)林霍尔姆(Lindholm,T.)等著.—北京:电子工业出版社,2016.4

(原味精品书系)

ISBN 978-7-121-27305-6

I. ① J…II. ① 林…III. ① JAVA 语言—程序设计—英文 IV. ① TP312

中国版本图书馆 CIP 数据核字(2015)第 231508 号

策划编辑:张春雨 刘 芸

责任编辑:徐津平

印 刷:三河市兴达印务有限公司

装 订:三河市兴达印务有限公司

出版发行:电子工业出版社

北京市海淀区万寿路 173 信箱 邮编:100036

开 本:787×980 1/16 印张:37.5 字数:720 千字

版 次:2016 年 4 月第 1 版

印 次:2016 年 4 月第 1 次印刷

定 价:108.00 元

凡所购买电子工业出版社图书有缺损问题,请向购买书店调换。若书店售缺,请与本社发行部联系,联系及邮购电话:(010) 88254888。

质量投诉请发邮件至 zlts@phei.com.cn,盗版侵权举报请发邮件至 dbqq@phei.com.cn。

服务热线:(010) 88258888。

向索菲亚和苏珊致以最深的谢意!



# 前言

本书涵盖了自 2011 年发布 Java SE 7 版之后所发生的所有变化。此外，为了与常见的 Java 虚拟机实现相匹配，本书还添加了大量修订及说明。

本版与前面各版一样，仅仅描述了抽象的 Java 虚拟机，而在实现具体的 Java 虚拟机时，本书仅指出了设计规划。Java 虚拟机的实现必须体现出本书所列规范，但仅在确有必要时才受限。

对 Java SE 8 而言，Java 编程语言里的一些重要变化在 Java 虚拟机中都有相应的体现。为了尽量保持二进制兼容性，我们应该直接在 Java 虚拟机里指定令人满意的默认方法，而不应该依赖于编译器，因为那样做将无法在不同厂商、版本的产品之间移植。此外，那种做法也不可能适用于已有的 class 文件。在设计 JSR 335——*Lambda Expressions for the Java Programming Language*（《Java 编程语言的 lambda 表达式》）时，Oracle 公司的 Dan Smith 向虚拟机实现者咨询了将默认方法集成到常量池和方法结构、方法与接口方法解析算法，以及字节码指令集中的上佳方式。JSR 335 也允许在 class 文件级别的接口里出现 private 方法与 static 方法，而这些方法也同接口方法解析算法紧密地结合起来了。

Java SE 8 的特点之一是：Java SE 平台的程序库也伴随着 Java 虚拟机一起进化。有个小例子可以很好地说明这一特点：在运行程序的时候，Java SE 8 可以获取方法的参数名，虚拟机会把这些名字存放在 class 文件结构中，而与此同时，`java.lang.reflect.Parameter` 里也有个标准的 API 能够查询这些名字。另外，我们也可以通过 class 文件结构中一项有趣的统计数据来说明这个特点：本规范的第 1 版中定义了 6 个属性，其中有 3 个属性对 Java 虚拟机至关重要，而 Java SE 8 版的规范则定义了 23 个属性，其中只有 5 个属性对 Java 虚拟机很重要。换句话说，在新版规范中，属性主要是为了支持程序库而设计的，其次才是为了支持 Java 虚拟机本身。为了帮助读者理解 class 文件结构，本规范会更为清晰地描述出每项属性的角色及使用限制。

在 Oracle 公司的 Java Platform 团队里，有多位同事都对这份规范提供了很大的支持，他们包括：Mandy Chung、Joe Darcy、Joel Franck、Staffan Friberg、Yuri Gaevsky、Jon Gibbons、Jeannette Hung、Eric McCorkle、Matherey Nunez、Mark Reinhold、John Rose、Georges Saab、Steve Sides、Bernard Traversat、Michel Trudeau 和 Mikael Vidstedt。特别感谢 Dan Heidinga (IBM)、Karen Kinnear、Keith McGuigan 及 Harold Seigel 对常见的 Java 虚拟机实现中的兼容性及安全性问题做出的贡献。

Alex Buckley

于加利福尼亚州圣克拉拉

2014 年 3 月



## 前言 xiii

## 1 Introduction 1

- 1.1 A Bit of History 1
- 1.2 The Java Virtual Machine 2
- 1.3 Organization of the Specification 3
- 1.4 Notation 4
- 1.5 Feedback 4

## 2 The Structure of the Java Virtual Machine 5

- 2.1 The `class` File Format 5
- 2.2 Data Types 6
- 2.3 Primitive Types and Values 6
  - 2.3.1 Integral Types and Values 7
  - 2.3.2 Floating-Point Types, Value Sets, and Values 8
  - 2.3.3 The `returnAddress` Type and Values 10
  - 2.3.4 The `boolean` Type 10
- 2.4 Reference Types and Values 11
- 2.5 Run-Time Data Areas 11
  - 2.5.1 The `pc` Register 12
  - 2.5.2 Java Virtual Machine Stacks 12
  - 2.5.3 Heap 13
  - 2.5.4 Method Area 13
  - 2.5.5 Run-Time Constant Pool 14
  - 2.5.6 Native Method Stacks 14
- 2.6 Frames 15
  - 2.6.1 Local Variables 16
  - 2.6.2 Operand Stacks 17
  - 2.6.3 Dynamic Linking 18
  - 2.6.4 Normal Method Invocation Completion 18
  - 2.6.5 Abrupt Method Invocation Completion 18
- 2.7 Representation of Objects 19
- 2.8 Floating-Point Arithmetic 19
  - 2.8.1 Java Virtual Machine Floating-Point Arithmetic and IEEE 754 19
  - 2.8.2 Floating-Point Modes 20
  - 2.8.3 Value Set Conversion 20
- 2.9 Special Methods 22
- 2.10 Exceptions 23
- 2.11 Instruction Set Summary 25



- 2.11.1 Types and the Java Virtual Machine 26
- 2.11.2 Load and Store Instructions 29
- 2.11.3 Arithmetic Instructions 30
- 2.11.4 Type Conversion Instructions 32
- 2.11.5 Object Creation and Manipulation 34
- 2.11.6 Operand Stack Management Instructions 34
- 2.11.7 Control Transfer Instructions 34
- 2.11.8 Method Invocation and Return Instructions 35
- 2.11.9 Throwing Exceptions 36
- 2.11.10 Synchronization 36
- 2.12 Class Libraries 37
- 2.13 Public Design, Private Implementation 37

### **3 Compiling for the Java Virtual Machine 39**

- 3.1 Format of Examples 39
- 3.2 Use of Constants, Local Variables, and Control Constructs 40
- 3.3 Arithmetic 45
- 3.4 Accessing the Run-Time Constant Pool 46
- 3.5 More Control Examples 47
- 3.6 Receiving Arguments 50
- 3.7 Invoking Methods 51
- 3.8 Working with Class Instances 53
- 3.9 Arrays 55
- 3.10 Compiling Switches 57
- 3.11 Operations on the Operand Stack 59
- 3.12 Throwing and Handling Exceptions 60
- 3.13 Compiling `finally` 63
- 3.14 Synchronization 66
- 3.15 Annotations 67

### **4 The `class` File Format 69**

- 4.1 The `ClassFile` Structure 70
- 4.2 The Internal Form of Names 74
  - 4.2.1 Binary Class and Interface Names 74
  - 4.2.2 Unqualified Names 75
- 4.3 Descriptors 75
  - 4.3.1 Grammar Notation 75
  - 4.3.2 Field Descriptors 76
  - 4.3.3 Method Descriptors 77
- 4.4 The Constant Pool 78
  - 4.4.1 The `CONSTANT_Class_info` Structure 79
  - 4.4.2 The `CONSTANT_Fieldref_info`, `CONSTANT_Methodref_info`, and `CONSTANT_InterfaceMethodref_info` Structures 80
  - 4.4.3 The `CONSTANT_String_info` Structure 81
  - 4.4.4 The `CONSTANT_Integer_info` and `CONSTANT_Float_info` Structures 82

- 4.4.5 The `CONSTANT_Long_info` and `CONSTANT_Double_info` Structures 83
- 4.4.6 The `CONSTANT_NameAndType_info` Structure 85
- 4.4.7 The `CONSTANT_Utf8_info` Structure 85
- 4.4.8 The `CONSTANT_MethodHandle_info` Structure 87
- 4.4.9 The `CONSTANT_MethodType_info` Structure 89
- 4.4.10 The `CONSTANT_InvokeDynamic_info` Structure 89
- 4.5 Fields 90
- 4.6 Methods 92
- 4.7 Attributes 95
  - 4.7.1 Defining and Naming New Attributes 101
  - 4.7.2 The `ConstantValue` Attribute 101
  - 4.7.3 The `Code` Attribute 102
  - 4.7.4 The `StackMapTable` Attribute 106
  - 4.7.5 The `Exceptions` Attribute 113
  - 4.7.6 The `InnerClasses` Attribute 114
  - 4.7.7 The `EnclosingMethod` Attribute 116
  - 4.7.8 The `Synthetic` Attribute 118
  - 4.7.9 The `Signature` Attribute 118
    - 4.7.9.1 Signatures 119
  - 4.7.10 The `SourceFile` Attribute 123
  - 4.7.11 The `SourceDebugExtension` Attribute 124
  - 4.7.12 The `LineNumberTable` Attribute 124
  - 4.7.13 The `LocalVariableTable` Attribute 126
  - 4.7.14 The `LocalVariableTypeTable` Attribute 128
  - 4.7.15 The `Deprecated` Attribute 129
  - 4.7.16 The `RuntimeVisibleAnnotations` Attribute 130
    - 4.7.16.1 The `element_value` structure 132
  - 4.7.17 The `RuntimeInvisibleAnnotations` Attribute 135
  - 4.7.18 The `RuntimeVisibleParameterAnnotations` Attribute 136
  - 4.7.19 The `RuntimeInvisibleParameterAnnotations` Attribute 137
  - 4.7.20 The `RuntimeVisibleTypeAnnotations` Attribute 139
    - 4.7.20.1 The `target_info` union 144
    - 4.7.20.2 The `type_path` structure 148
  - 4.7.21 The `RuntimeInvisibleTypeAnnotations` Attribute 152
  - 4.7.22 The `AnnotationDefault` Attribute 153
  - 4.7.23 The `BootstrapMethods` Attribute 154
  - 4.7.24 The `MethodParameters` Attribute 156
- 4.8 Format Checking 158
- 4.9 Constraints on Java Virtual Machine Code 159
  - 4.9.1 Static Constraints 159
  - 4.9.2 Structural Constraints 163
- 4.10 Verification of `class` Files 166
  - 4.10.1 Verification by Type Checking 167
    - 4.10.1.1 Accessors for Java Virtual Machine Artifacts 169
    - 4.10.1.2 Verification Type System 173
    - 4.10.1.3 Instruction Representation 177
    - 4.10.1.4 Stack Map Frame Representation 178

- 4.10.1.5 Type Checking Abstract and Native Methods 184
- 4.10.1.6 Type Checking Methods with Code 187
- 4.10.1.7 Type Checking Load and Store Instructions 194
- 4.10.1.8 Type Checking for `protected` Members 196
- 4.10.1.9 Type Checking Instructions 199
- 4.10.2 Verification by Type Inference 319
  - 4.10.2.1 The Process of Verification by Type Inference 319
  - 4.10.2.2 The Bytecode Verifier 319
  - 4.10.2.3 Values of Types `long` and `double` 323
  - 4.10.2.4 Instance Initialization Methods and Newly Created Objects 323
  - 4.10.2.5 Exceptions and `finally` 325
- 4.11 Limitations of the Java Virtual Machine 327

## 5 Loading, Linking, and Initializing 329

- 5.1 The Run-Time Constant Pool 329
- 5.2 Java Virtual Machine Startup 332
- 5.3 Creation and Loading 332
  - 5.3.1 Loading Using the Bootstrap Class Loader 334
  - 5.3.2 Loading Using a User-defined Class Loader 335
  - 5.3.3 Creating Array Classes 336
  - 5.3.4 Loading Constraints 336
  - 5.3.5 Deriving a Class from a `class` File Representation 338
- 5.4 Linking 339
  - 5.4.1 Verification 340
  - 5.4.2 Preparation 340
  - 5.4.3 Resolution 341
    - 5.4.3.1 Class and Interface Resolution 342
    - 5.4.3.2 Field Resolution 343
    - 5.4.3.3 Method Resolution 344
    - 5.4.3.4 Interface Method Resolution 346
    - 5.4.3.5 Method Type and Method Handle Resolution 347
    - 5.4.3.6 Call Site Specifier Resolution 350
  - 5.4.4 Access Control 351
  - 5.4.5 Overriding 352
- 5.5 Initialization 352
- 5.6 Binding Native Method Implementations 355
- 5.7 Java Virtual Machine Exit 355

## 6 The Java Virtual Machine Instruction Set 357

- 6.1 Assumptions: The Meaning of "Must" 357
- 6.2 Reserved Opcodes 358
- 6.3 Virtual Machine Errors 358
- 6.4 Format of Instruction Descriptions 359
  - mnemonic 360
- 6.5 Instructions 362
  - aaload* 363

*aastore* 364  
*aconst\_null* 366  
*aload* 367  
*aload\_<n>* 368  
*anewarray* 369  
*areturn* 370  
*arraylength* 371  
*astore* 372  
*astore\_<n>* 373  
*athrow* 374  
*baload* 376  
*bastore* 377  
*bipush* 378  
*caload* 379  
*castore* 380  
*checkcast* 381  
*d2f* 383  
*d2i* 384  
*d2l* 385  
*dadd* 386  
*daload* 388  
*dastore* 389  
*dcmp<op>* 390  
*dconst\_<d>* 392  
*ddiv* 393  
*dload* 395  
*dload\_<n>* 396  
*dmul* 397  
*dneg* 399  
*drem* 400  
*dreturn* 402  
*dstore* 403  
*dstore\_<n>* 404  
*dsub* 405  
*dup* 406  
*dup\_x1* 407  
*dup\_x2* 408  
*dup2* 409  
*dup2\_x1* 410  
*dup2\_x2* 411  
*f2d* 413  
*f2i* 414  
*f2l* 415  
*fadd* 416  
*faload* 418  
*fastore* 419  
*fcmp<op>* 420  
*fconst\_<f>* 422

## 目录

*fdiv* 423  
*fload* 425  
*fload\_<n>* 426  
*fmul* 427  
*fneg* 429  
*frem* 430  
*freturn* 432  
*fstore* 433  
*fstore\_<n>* 434  
*fsub* 435  
*getfield* 436  
*getstatic* 438  
*goto* 440  
*goto\_w* 441  
*i2b* 442  
*i2c* 443  
*i2d* 444  
*i2f* 445  
*i2l* 446  
*i2s* 447  
*iadd* 448  
*iaload* 449  
*iand* 450  
*iastore* 451  
*iconst\_<i>* 452  
*idiv* 453  
*if\_acmp<cond>* 454  
*if\_icmp<cond>* 455  
*if<cond>* 457  
*ifnonnull* 459  
*ifnull* 460  
*iinc* 461  
*iload* 462  
*iload\_<n>* 463  
*imul* 464  
*ineg* 465  
*instanceof* 466  
*invokedynamic* 468  
*invokeinterface* 473  
*invokespecial* 477  
*invokestatic* 481  
*invokevirtual* 484  
*ior* 489  
*irem* 490  
*ireturn* 491  
*ishl* 492  
*ishr* 493  
*istore* 494

*istore\_<n>* 495  
*isub* 496  
*iushr* 497  
*ixor* 498  
*jsr* 499  
*jsr\_w* 500  
*l2d* 501  
*l2f* 502  
*l2i* 503  
*ladd* 504  
*laload* 505  
*land* 506  
*lastore* 507  
*lcmp* 508  
*lconst\_<l>* 509  
*ldc* 510  
*ldc\_w* 512  
*ldc2\_w* 514  
*ldiv* 515  
*lload* 516  
*lload\_<n>* 517  
*lmul* 518  
*lneg* 519  
*lookupswitch* 520  
*lor* 522  
*lrem* 523  
*lreturn* 524  
*lshl* 525  
*lshr* 526  
*lstore* 527  
*lstore\_<n>* 528  
*lsub* 529  
*lushr* 530  
*lxor* 531  
*monitorenter* 532  
*monitorexit* 534  
*multianewarray* 536  
*new* 538  
*newarray* 540  
*nop* 542  
*pop* 543  
*pop2* 544  
*putfield* 545  
*putstatic* 547  
*ret* 549  
*return* 550  
*saload* 551  
*sastore* 552

目录

*sipush* 553  
*swap* 554  
*tableswitch* 555  
*wide* 557

**7 Opcode Mnemonics by Opcode 559**

**Index 563**

**A Limited License Grant 581**

---

# Introduction

## 1.1 A Bit of History

The Java® programming language is a general-purpose, concurrent, object-oriented language. Its syntax is similar to C and C++, but it omits many of the features that make C and C++ complex, confusing, and unsafe. The Java platform was initially developed to address the problems of building software for networked consumer devices. It was designed to support multiple host architectures and to allow secure delivery of software components. To meet these requirements, compiled code had to survive transport across networks, operate on any client, and assure the client that it was safe to run.

The popularization of the World Wide Web made these attributes much more interesting. Web browsers enabled millions of people to surf the Net and access media-rich content in simple ways. At last there was a medium where what you saw and heard was essentially the same regardless of the machine you were using and whether it was connected to a fast network or a slow modem.

Web enthusiasts soon discovered that the content supported by the Web's HTML document format was too limited. HTML extensions, such as forms, only highlighted those limitations, while making it clear that no browser could include all the features users wanted. Extensibility was the answer.

The HotJava browser first showcased the interesting properties of the Java programming language and platform by making it possible to embed programs inside HTML pages. Programs are transparently downloaded into the browser along with the HTML pages in which they appear. Before being accepted by the browser, programs are carefully checked to make sure they are safe. Like HTML pages, compiled programs are network- and host-independent. The programs behave the same way regardless of where they come from or what kind of machine they are being loaded into and run on.



A Web browser incorporating the Java platform is no longer limited to a predetermined set of capabilities. Visitors to Web pages incorporating dynamic content can be assured that their machines cannot be damaged by that content. Programmers can write a program once, and it will run on any machine supplying a Java run-time environment.

## 1.2 The Java Virtual Machine

The Java Virtual Machine is the cornerstone of the Java platform. It is the component of the technology responsible for its hardware- and operating system-independence, the small size of its compiled code, and its ability to protect users from malicious programs.

The Java Virtual Machine is an abstract computing machine. Like a real computing machine, it has an instruction set and manipulates various memory areas at run time. It is reasonably common to implement a programming language using a virtual machine; the best-known virtual machine may be the P-Code machine of UCSD Pascal.

The first prototype implementation of the Java Virtual Machine, done at Sun Microsystems, Inc., emulated the Java Virtual Machine instruction set in software hosted by a handheld device that resembled a contemporary Personal Digital Assistant (PDA). Oracle's current implementations emulate the Java Virtual Machine on mobile, desktop and server devices, but the Java Virtual Machine does not assume any particular implementation technology, host hardware, or host operating system. It is not inherently interpreted, but can just as well be implemented by compiling its instruction set to that of a silicon CPU. It may also be implemented in microcode or directly in silicon.

The Java Virtual Machine knows nothing of the Java programming language, only of a particular binary format, the `class` file format. A `class` file contains Java Virtual Machine instructions (or *bytecodes*) and a symbol table, as well as other ancillary information.

For the sake of security, the Java Virtual Machine imposes strong syntactic and structural constraints on the code in a `class` file. However, any language with functionality that can be expressed in terms of a valid `class` file can be hosted by the Java Virtual Machine. Attracted by a generally available, machine-independent platform, implementors of other languages can turn to the Java Virtual Machine as a delivery vehicle for their languages.