

### New JVM Features – Multitenancy and Packed Objects





Important Disclaimers

- THE INFORMATION CONTAINED IN THIS PRESENTATION IS PROVIDED FOR INFORMATIONAL PURPOSES ONLY.
- WHILST EFFORTS WERE MADE TO VERIFY THE COMPLETENESS AND ACCURACY OF THE INFORMATION CONTAINED IN THIS PRESENTATION, IT IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED.
- ALL PERFORMANCE DATA INCLUDED IN THIS PRESENTATION HAVE BEEN GATHERED IN A CONTROLLED ENVIRONMENT. YOUR OWN TEST RESULTS MAY VARY BASED ON HARDWARE, SOFTWARE OR INFRASTRUCTURE DIFFERENCES.
- ALL DATA INCLUDED IN THIS PRESENTATION ARE MEANT TO BE USED ONLY AS A GUIDE.
- IN ADDITION, THE INFORMATION CONTAINED IN THIS PRESENTATION IS BASED ON IBM'S CURRENT PRODUCT PLANS AND STRATEGY, WHICH ARE SUBJECT TO CHANGE BY IBM, WITHOUT NOTICE.
- IBM AND ITS AFFILIATED COMPANIES SHALL NOT BE RESPONSIBLE FOR ANY DAMAGES ARISING OUT OF THE USE OF, OR OTHERWISE RELATED TO, THIS PRESENTATION OR ANY OTHER DOCUMENTATION.
- NOTHING CONTAINED IN THIS PRESENTATION IS INTENDED TO, OR SHALL HAVE THE EFFECT OF:
- CREATING ANY WARRANT OR REPRESENTATION FROM IBM, ITS AFFILIATED COMPANIES OR ITS OR THEIR SUPPLIERS AND/OR LICENSORS

# A smarter planet C Solar Solar



### Introduction to the speaker

- Neil Masson
- Many (many) years experience in Java
- IBM Service Core L3 Team





### What should you get from this talk?

- JVM proving to be a fertile ecosystem for languages
- Plenty of opportunity to innovate in other spaces
- Runtime is the gateway to this innovation
- Largely ignored the last few years, but this is where the core inventions can occur



# The runtime isn't boring!

Multitenancy

### Just what do you mean by "multitenancy"?

With a multitenant architecture, a software application is designed to virtually partition its data and configuration, and each client organization works with a customized virtual application instance.

oftware for a smarter planet

#### Working Definition

- -A single instance of a software application that serves multiple customers  $\rightarrow$  Each customer is a tenant.
- -Tenants can customize some parts of the application (look and feel) but not the code.
- Infrastructure usually opaque
  - $\rightarrow$  opportunity for provider

Why? <u>Cost Savings</u>: As compared to single-tenant deployment model





# a smarter planet



#### JDK Support for Spectrum of Sharing / Multitenancy (Level 1-5)

![](_page_7_Figure_3.jpeg)

data center

![](_page_7_Figure_4.jpeg)

Changes

S2. Shared

Hardware

![](_page_7_Figure_5.jpeg)

![](_page_7_Figure_6.jpeg)

![](_page_7_Figure_7.jpeg)

**Application** Changes

![](_page_8_Figure_1.jpeg)

### A Second Contract of the second contract of t

![](_page_9_Picture_1.jpeg)

### Hardware Virtualization

![](_page_9_Figure_3.jpeg)

- Hypervisors run multiple applications side-by-side safely
- Advantages
  - Capture idle CPU cycles
  - Automatic de-duplication (RAM)
  - Ability to meter and shift resource toward demand
  - No need to change tenant applications

### A marter planet C S O

![](_page_10_Picture_1.jpeg)

### Hardware Virtualization

![](_page_10_Figure_3.jpeg)

- Hypervisors JVMs can run multiple applications side-by-side safely
- Advantages
  - Capture idle CPU cycles
  - Automatic de-duplication (ability to share Java artifacts)
  - Ability to meter and shift resource toward demand
  - No need to change tenant applications

![](_page_11_Picture_0.jpeg)

#### Multitenancy: Low (or no) barrier to entry

- Multitenancy is all about reducing duplication by transparently sharing a JVM
  - -1 GC, 1 JIT, shared heap objects
  - -plus: JVM-enforced resource constraints to meter and limit consumption
- Ergonomics: Opt-in to multitenancy with a single flag: -Xmt (multitenancy) – no application changes required

![](_page_11_Figure_6.jpeg)

![](_page_12_Picture_0.jpeg)

#### JVM: Separating State

- Static variables are a problem for sharing
- Consider the effect of global defaults

java.util.TimeZone

private static volatile TimeZone defaultTimeZone;

![](_page_13_Picture_0.jpeg)

### JVM: Separating State

- Use @TenantScope annotation
- Each tenant has its own version of defaultTimeZone
- All JVM classes annotated

java.util.TimeZone

@TenantScope
private static volatile TimeZone defaultTimeZone;

A sparter planet C S O

![](_page_14_Picture_1.jpeg)

#### **Tenants: Separating State**

![](_page_14_Figure_3.jpeg)

# A marter planet C Solares

![](_page_15_Picture_1.jpeg)

#### **Tenants: Separating State**

![](_page_15_Figure_3.jpeg)

![](_page_16_Picture_0.jpeg)

![](_page_16_Picture_1.jpeg)

More that just JVM State...

- Throttling of resources
  - -Threads, GC, sockets, files (IO in general), native memory
- Past and existing examples do exist!
  - -Commercial / In house custom solutions
  - -JSR 181 Isolates / 284 Resource Management
- Security is of course huge

![](_page_17_Picture_1.jpeg)

### Other Thoughts – Native Libraries and shared state

![](_page_17_Figure_3.jpeg)

Native libraries contain state that may not be shareable across tenants

![](_page_17_Figure_5.jpeg)

- Use separate processes to manage different state
- Each process now holds the context
- Challenges: Latency

Questions?

Packed Objects

![](_page_20_Picture_1.jpeg)

#### Problem? What problem?

- JNI just isn't a great way to marshal data
- Locality in Java can matter
- Existing native and data placement stories aren't very good
- In many cases, legacy systems exist the interop is just terrible
- So we want something that integrates well with the Java language and helps us...

# A analysis of tware for a smarter planet C Software for a smar

![](_page_21_Picture_1.jpeg)

### What are we trying to solve?

![](_page_21_Figure_3.jpeg)

Simple enough...

# A stanter planet C Software for a smarter planet C Software fo

![](_page_22_Picture_1.jpeg)

#### What are we trying to solve?

![](_page_22_Figure_3.jpeg)

Simple enough...

Header overhead

# A an arter planet a smarter planet a solution of the second secon

![](_page_23_Picture_1.jpeg)

#### What are we trying to solve?

![](_page_23_Figure_3.jpeg)

Simple enough...

- Header overhead
- Pointer chasing

# A sparter planet C Software for a smarter planet C Software

![](_page_24_Picture_1.jpeg)

#### What are we trying to solve?

![](_page_24_Figure_3.jpeg)

Simple enough...

- Header overhead
- Pointer chasing
- Locality

A sparter planet C

IBM

#### What are we trying to solve?

![](_page_25_Figure_3.jpeg)

![](_page_25_Figure_4.jpeg)

![](_page_25_Picture_5.jpeg)

Fighting the Java/Native interface

![](_page_26_Picture_1.jpeg)

Ok so we have some criteria...

- Ability to do away with headers
- Ability to bring multiple objects close together
- On heap / off heap seamless referencing of data
- This actually sounds a lot like C structure types

```
struct Address {
    char[4] addr;
    short port;
}
struct Header {
    struct Address src;
    struct Address dst;
}
```

struct Header

![](_page_26_Picture_9.jpeg)

Packed Objects!

### A Software for a smarter planet C Software

![](_page_27_Picture_1.jpeg)

### Packed Objects: Under the covers

![](_page_27_Picture_3.jpeg)

![](_page_27_Picture_4.jpeg)

#### aPoint

int	Х
int	V

# A sparter planet C Solar of the second secon

![](_page_28_Figure_1.jpeg)

### Packed Objects: Under the covers

Object header

![](_page_28_Picture_4.jpeg)

![](_page_28_Figure_5.jpeg)

# A smarter planet C Solar

![](_page_29_Figure_1.jpeg)

### Packed Objects: Under the covers

Object header

![](_page_29_Picture_4.jpeg)

![](_page_29_Figure_5.jpeg)

A stanter planet C S C

![](_page_30_Figure_1.jpeg)

### **Packed Objects: In Practice**

Object header

![](_page_30_Figure_4.jpeg)

![](_page_30_Figure_5.jpeg)

# A sparter planet C S O t

![](_page_31_Figure_1.jpeg)

![](_page_31_Figure_3.jpeg)

![](_page_31_Figure_4.jpeg)

![](_page_31_Picture_5.jpeg)

# A an arter planet a smarter planet a sma

![](_page_32_Figure_1.jpeg)

![](_page_32_Figure_3.jpeg)

![](_page_32_Picture_4.jpeg)

![](_page_32_Picture_5.jpeg)

# A an arter planet C S O

![](_page_33_Figure_1.jpeg)

![](_page_33_Figure_3.jpeg)

# A smarter planet C S O C

![](_page_34_Figure_1.jpeg)

![](_page_34_Figure_3.jpeg)

PackedPoint p = aPackedLine.e

![](_page_35_Picture_1.jpeg)

### Packed Objects: In Practice with Native Access

![](_page_35_Figure_3.jpeg)

![](_page_36_Figure_1.jpeg)

#### Packed Objects: In Practice with Native Access

![](_page_36_Figure_3.jpeg)

![](_page_36_Picture_4.jpeg)

![](_page_37_Picture_1.jpeg)

#### Packed Objects: In Practice with Native Access

![](_page_37_Figure_3.jpeg)

![](_page_37_Picture_4.jpeg)

A software for a smarter planet C Software for a smarter plane

![](_page_38_Picture_1.jpeg)

#### Lets Build Something in C!

![](_page_38_Figure_3.jpeg)

- Nested substructures
- Compact representation
- Alignment aspects

# A analysis of tware for a smarter planet C S O T

![](_page_39_Picture_1.jpeg)

### Let's Build the Same "Something" in Java!

![](_page_39_Figure_3.jpeg)

- Headers
- No locality
- Alignment

# A sparter planet C Solar P

![](_page_40_Picture_1.jpeg)

### What if we did this with Packed Objects?

![](_page_40_Picture_3.jpeg)

• The Java code is pretty clean... and a pretty good result!

Address src

Address dst

a marter planet C Solar

![](_page_41_Picture_1.jpeg)

#### What about native access?

![](_page_41_Figure_3.jpeg)

How do we implement this normally?

# A software for a smarter planet C S O T Z O

![](_page_42_Picture_1.jpeg)

#### **JNI** implementation

```
public class PackedHeader {
    private long pointer;
    public byte[] getSourceAddress() { return getSourceAddressImpl(pointer); }
    public short getSourcePort() { return getSourcePortImpl(pointer); }
}
JNICALL jshort Java_pkg_PackedHeader_getSourcePort(JNIEnv* env, jobject recv, jlong pointer) {
    struct PacketHeader* header = (struct PacketHeader*)pointer;
    return (jshort)header->src.port;
}
JNICALL jbyteArray Java_pkg_PackedHeader_getSourceAddress(JNIEnv* env, jobject recv, jlong pointer) {
    struct PacketHeader* header = (struct PacketHeader*)pointer;
    jbyteArray result = (*env)->NewByteArray(env, 4);
    (*env)->SetByteArrayRegion(env, result, 0, 4, &(header->src.addr));
    return result;
}
```

- Usual "stash pointers in long types" tricks
- JNI costs tend to be high

# A sparter planet C Software for a smarter planet C Software fo

![](_page_43_Picture_1.jpeg)

DirectByteBuffer implementation

```
class PackedHeader {
    private ByteBuffer buffer;
    private static final int SRC_ADDR_OFFSET = 0;
    private static final int SRC_PORT_OFFSET = 4;
    private static final int DEST_ADDR_OFFSET = 8;
    private static final int DEST_PORT_OFFSET = 12;

    public short getSourcePort() { return buffer.getShort(SRC_PORT_OFFSET); }
    public byte[] getSourceAddress() {
        byte[] result = new byte[4];
        buffer.get(result, SRC_ADDR_OFFSET, 4);
        return result;
     }
}
```

- No extra JNI to write (this is good)
- Keeping your indices straight is never fun

# and the second let a second let a

![](_page_44_Picture_1.jpeg)

### Unsafe implementation

```
class PackedHeader {
    private Unsafe unsafe;
    private long pointer;
    private static final int SRC_ADDR_OFFSET = 0;
    private static final int SRC_PORT_OFFSET = 4;
    private static final int DEST_ADDR_OFFSET = 8;
    private static final int DEST_PORT_OFFSET = 12;

    public short getSourcePort() { return unsafe.getShort(pointer + SRC_PORT_OFFSET); }
    public byte[] getSourceAddress() {
        byte[] result = new byte[4];
        unsafe.copyMemory(null, pointer + SRC_ADDR_OFFSET, result, 0, 4);
        return result;
     }
}
```

- You shouldn't be here
- Still playing the indices game

#### PackedObject answer

```
final class PacketHeader extends PackedObject {
   Address src;
   Address dest;

   public short getSourcePort() { return src.port; }
   public PackedByte[] getSourceAddress() { return src.addr; }
}
```

![](_page_45_Figure_3.jpeg)

- Looks like natural Java code
- Foregoes JNI
- Same type capable of on-heap representation

![](_page_46_Picture_0.jpeg)

![](_page_46_Picture_1.jpeg)

#### Active work and next steps

- Experimenting with this now
- Yes, there are security aspects to be aware of here
- This is potentially part of a larger look at Java / Platform interop
- Not specifically viewed as a cure to GC problems
- This forms the basis for many other solutions to existing problems...

Questions?

and the second s

![](_page_48_Picture_1.jpeg)

#### References

#### Get Products and Technologies:

- -IBM Java Runtimes and SDKs:
  - https://www.ibm.com/developerworks/java/jdk/
- -IBM Monitoring and Diagnostic Tools for Java:
  - https://www.ibm.com/developerworks/java/jdk/tools/

#### Learn:

- -IBM Java InfoCenter:
  - http://publib.boulder.ibm.com/infocenter/java7sdk/v7r0/index.jsp

#### Discuss:

- -IBM Java Runtimes and SDKs Forum:
  - http://www.ibm.com/developerworks/forums/forum.jspa?forumID=367&start=0

![](_page_49_Picture_0.jpeg)

**Copyright and Trademarks** 

- © IBM Corporation 2013. All Rights Reserved.
- IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corp., and registered in many jurisdictions worldwide.
- Other product and service names might be trademarks of IBM or other companies.
- A current list of IBM trademarks is available on the Web see the IBM "Copyright and trademark information" page at URL: www.ibm.com/legal/copytrade.shtml