

SILICON VALLEY



IN-PERSISTENT-MEMORY COMPUTING WITH JAVA

ERIC KACZMAREK
INTEL CORPORATION

LEGAL DISCLAIMER & OPTIMIZATION NOTICE

INFORMATION IN THIS DOCUMENT IS PROVIDED “AS IS”. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. INTEL ASSUMES NO LIABILITY WHATSOEVER AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO THIS INFORMATION INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT.

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products.

Copyright © 2016, Intel Corporation. All rights reserved. Intel, Pentium, Xeon, Xeon Phi, Core, VTune, Cilk, and the Intel logo are trademarks of Intel Corporation in the U.S. and other countries.

Optimization Notice


Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice.

Notice revision #20110804

IN-MEMORY COMPUTING

- Many sessions to discuss its value and approaches 😊
- Definition from a quick search

“In-memory computing is the storage of information in the main random access memory (RAM) of dedicated servers rather than in complicated relational databases operating on comparatively slow disk drives”** - Techopedia.com



Is this the name/title of the interface?
If not, "interface" should be lower case. My guess is that it should read "Behind a SATA/SAS/PCIE interface"

- Behind a SATA/SAS/PCIE interface
- High access latencies
- Low read/write bandwidth
- Data stored as a stream of bytes

TRADITIONAL STORAGE STILL RELEVANT

- Cheap
- High capacity
- Durable

SIDE BY SIDE STORAGE AND MEMORY

	Storage (HDD, SSD, NVMe)	Memory (DRAM)
Capacity	Terabytes	Gigabytes
Durability	Yes	No (through software)
Access	Stream of bytes	Random data access
Bandwidth	~3GBps	~60GBps
Latencies	Microseconds	Nanoseconds

Complement each other


WHAT IF WE CAN HAVE THE BEST OF BOTH

	Storage (HDD, SSD, NVMe)	Memory (DRAM)
Capacity	Terabytes	Gigabytes
Durability	Yes	No (through software)
Access	Stream of bytes	Random data access
Bandwidth	~3GBps	~60GBps
Latencies	Microseconds	Nanoseconds

There was a "click to add" still showing in the background. I removed the blank text box that was housing it.

Vincent, Amber

WITH A TRADITIONAL STORAGE



There was a "click to add" still showing in the background. I removed the blank text box that was housing it.

Vincent, Amber

WITH A TRADITIONAL STORAGE


- NVM Tread
- NVM xfer
- Controller ASIC
- Controller Firmware

I/O performance determined by more than the NVM media, factors like controller latency, drivers, PCI-E performance and software stacks. Application performance will not equal the media performance




I removed the text box with "click to add content"

PERSISTENT MEMORY



PERSISTENT MEMORY DEFINITION AND VALUE

- Memory-like performance – not a DRAM replacement
- Byte addressable – no DRAM footprint
- Durable across applications or system restarts
- Large capacity (terabytes)
- Direct user mode access – no kernel code in data path



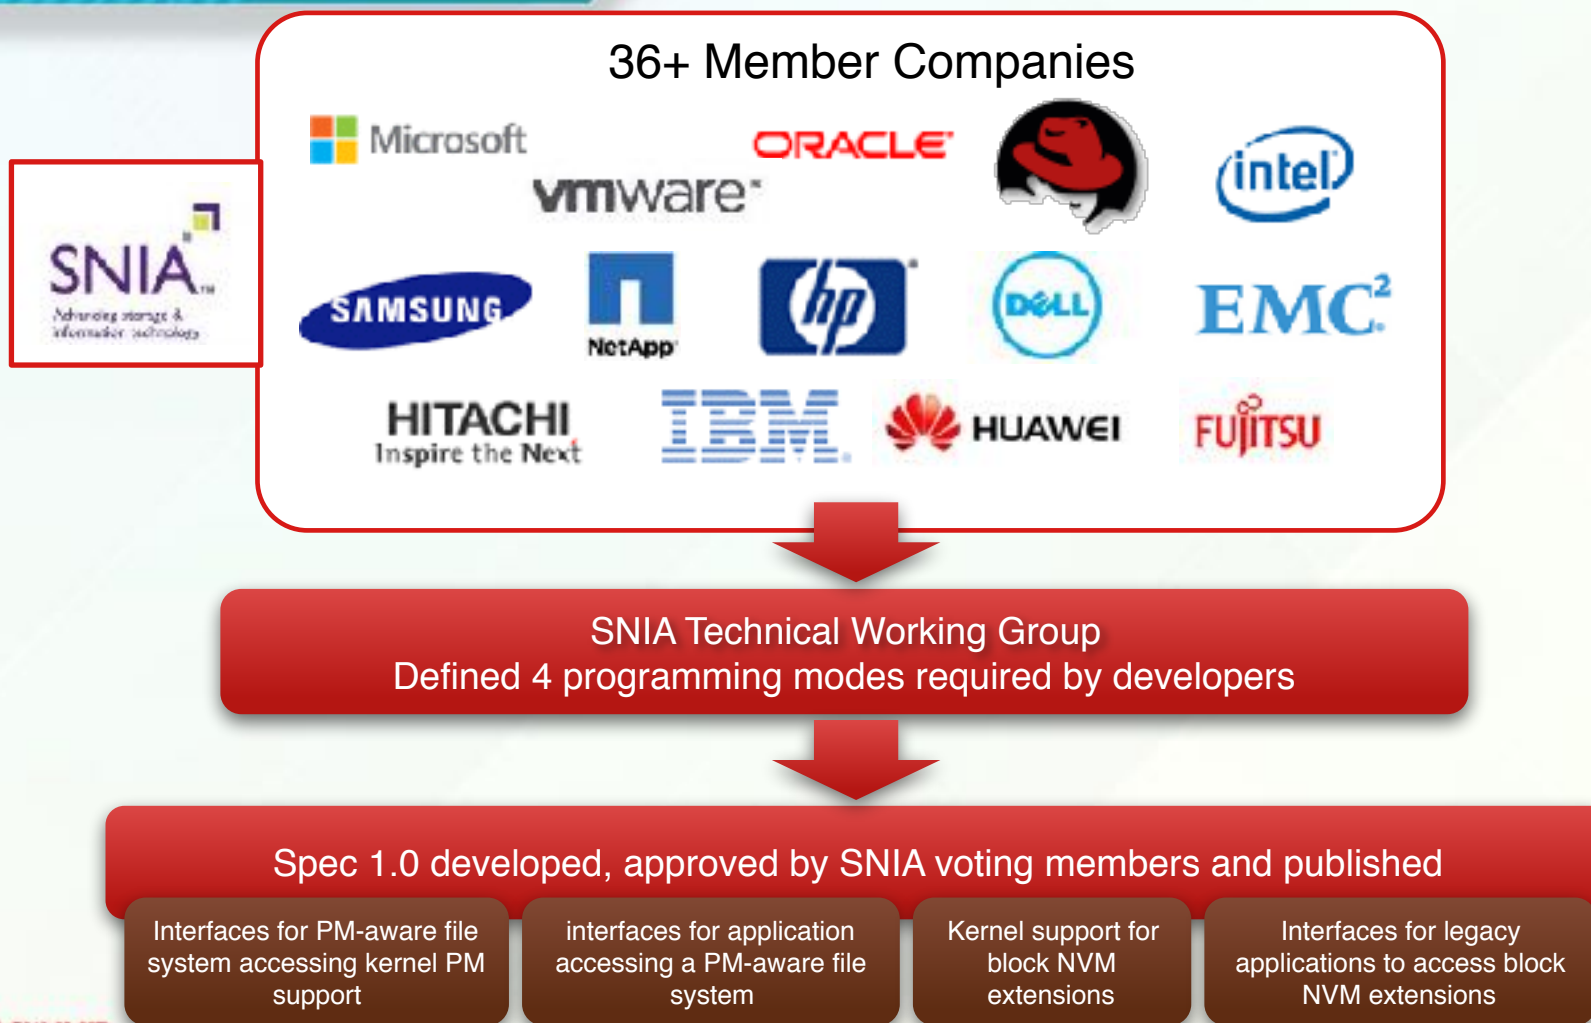
I removed the text box with "click to add content"

Vincent, Amber
Made grey box slightly bigger for

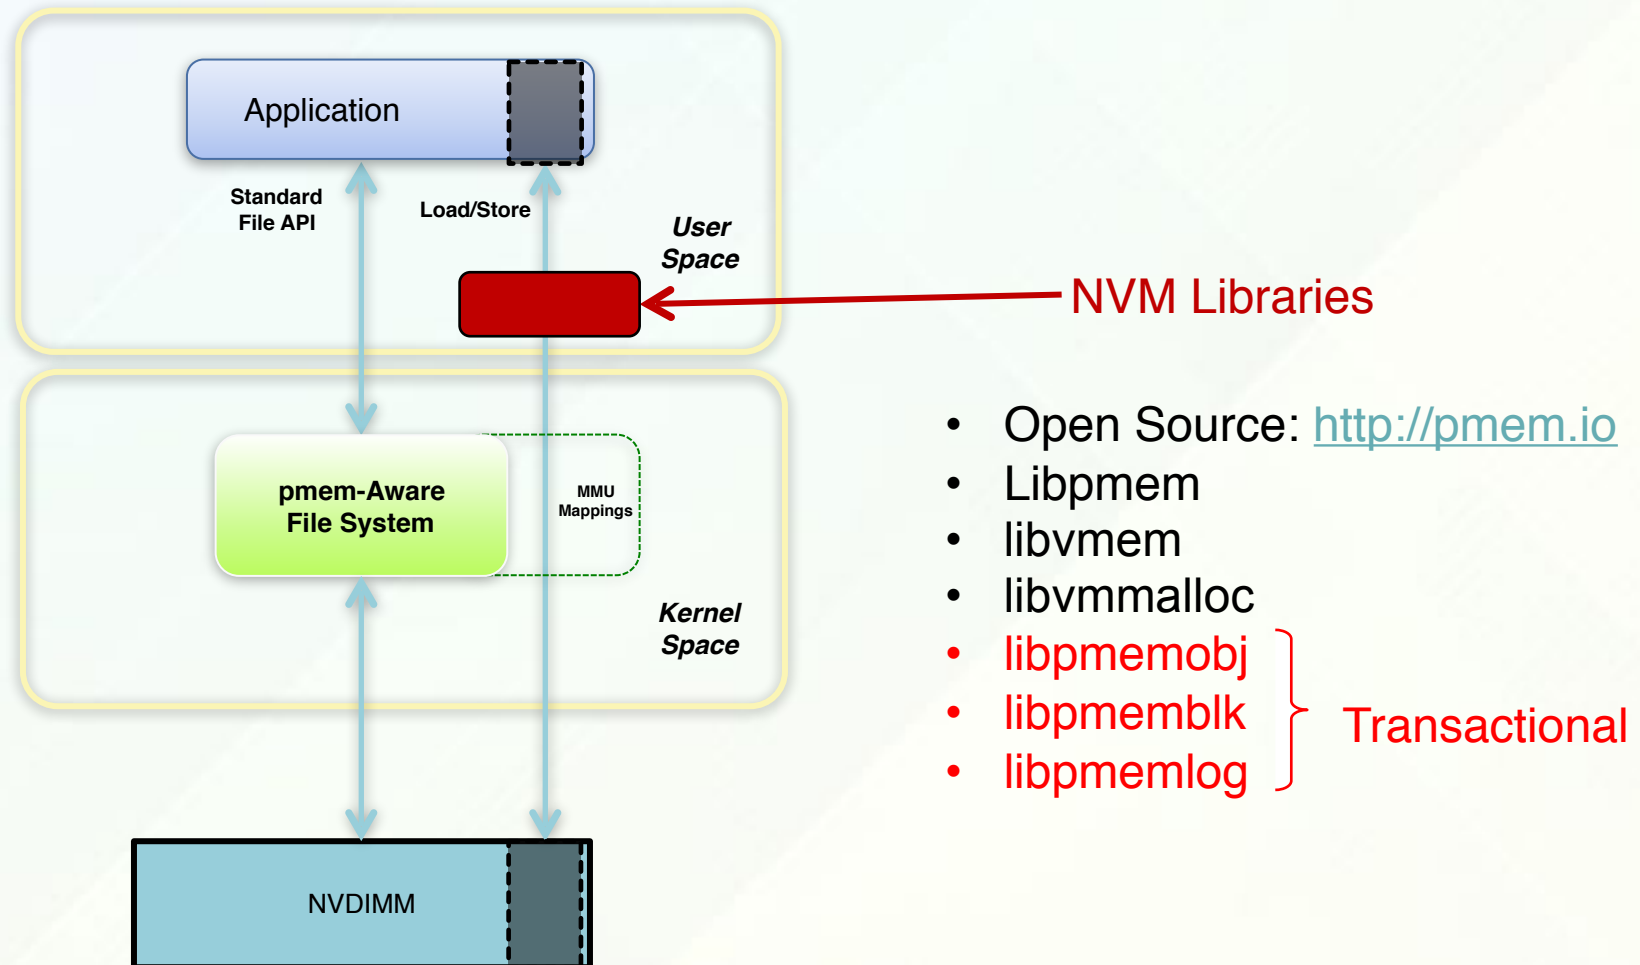
MEMORY



OPEN INDUSTRY PROGRAMMING MODEL



C/C++ OPEN SOURCE NVM LIBRARY



WHAT ABOUT JAVA

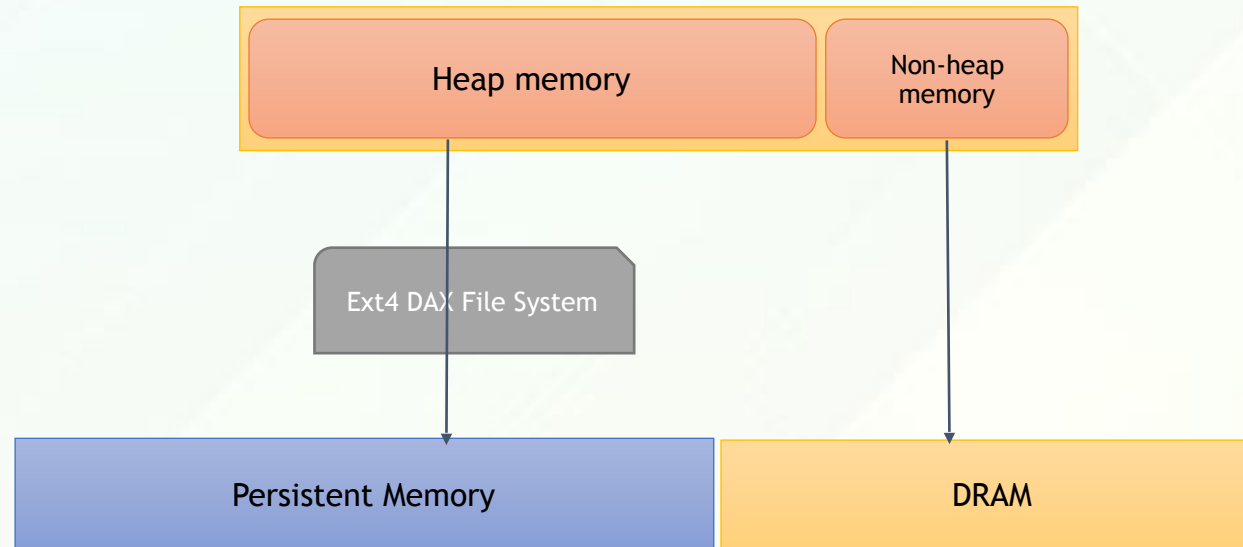
- Broad range of middleware developed in Java or Scala (JVM languages)
 - Ex: Apache Spark, Apache Cassandra, Apache Ignite, Hazelcast IMDG, etc.
- Java abstracts hardware from the developer
- Hooks or hardware access possible through JNI, Unsafe, etc. but
 - Not portable
 - Performance overhead (data marshaling, thread safety, etc.)
 - Might not be supported in future releases

EXPOSING PERSISTENT MEMORY TO JAVA

- Entire Java heap in Persistent Memory
- Heterogeneous Java Heap
- Persistent Collections for Java (PCJ)

ENTIRE JAVA HEAP IN PERSISTENT MEMORY

- Heap memory allocation on persistent memory
- No code changes
- `java -Xmx32g -Xms16g -XX:HeapDir=/XPointFS/heap ApplicationClass`
JVM

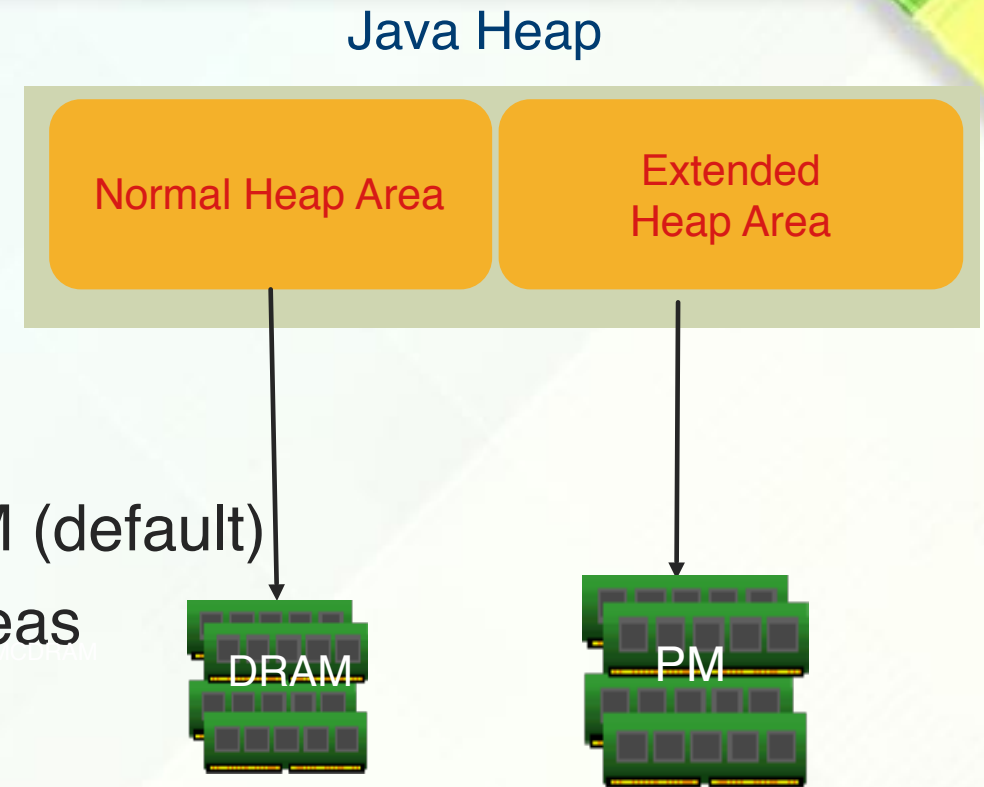


ENTIRE JAVA HEAP IN PERSISTENT MEMORY USE CASES

- In multi-JVM deployments to prioritize Java VMs. (ex: Oracle Fusion Apps)
- Applications which can benefit from large memory
- OpenJDK JEP: <https://bugs.openjdk.java.net/browse/JDK-8171181>

HETEROGENEOUS JAVA HEAP


- User directed allocation
- Frequently accessed objects reside in DRAM (default)
- Garbage collection (G1 GC) collects both areas



HETEROGENEOUS JAVA HEAP INTERFACE

- Extended HeapArea size specified through flag `-Xmp` (ex: `-Xmx60g -Xmp500g`)
- APIs for setting allocation context

```
// Set the allocation target to PM
Heap.setAllocationTarget(AllocationTarget.PM);
HashMap user_records = new HashMap(1000000);
// reset allocation target to DRAM.
Heap.resetAllocationTarget();
```



Where is data persistence?

PERSISTENT COLLECTIONS FOR JAVA

- Library of persistent classes
- Custom persistent classes
- Low-level Memory Regions

<https://github.com/pmem/pcj>

LIBRARY OF PERSISTENT CLASSES

Primitive arrays (e.g. `PersistentByteArray`, mutable and immutable)

`PersistentArray<E extends PersistentObject>` (mutable and immutable)

`PersistentTuple<T1 extends PersistentObject, ...>` (mutable and immutable)

`PersistentArrayList<E extends PersistentObject>`

`PersistentHashMap<K extends PersistentObject, V extends PersistentObject>`

`PersistentLinkedList<E extends PersistentObject>`

`PersistentLinkedQueue<E extends PersistentObject>`

`PersistentSkipListMap<K extends PersistentObject, V extends PersistentObject>`

`PersistentFPTreeMap<K extends PersistentObject, V extends PersistentObject>`

`PersistentSIHashMap<K extends PersistentObject, V extends PersistentObject>`

`ObjectDirectory` - indefinitely reachable root map of `<String, T extends PersistentObject>`

Primitive types (as field and array element values, no separate classes)

Boxed primitives (e.g. `PersistentLong`)

`PersistentString`

`PersistentByteBuffer`

`PersistentAtomicReference<T extends PersistentObject>`

LIBRARY OF PERSISTENT CLASSES

- State stored on persistent heap
- Instances behave like regular Java objects, just longer-lived
- Reachability-based lifetime
- Easy-to-understand data consistency model (transactional)

USING PERSISTENT COLLECTIONS

```
PersistentIntArray data = new PersistentIntArray(1024);
ObjectDirectory.put("MyApplicationData", data);
// no serialization, reference to array is written
data.set(0, 123);

// Restart JVM or system
PersistentIntArray data1 =
ObjectDirectory.get("MyApplicationData", PersistentIntArray.class);
assert(data.get(0) == 123);
```

SUPPORT FOR CUSTOM PERSISTENT CLASSES

- Extending built-in persistent class
- Creating a new persistent class

EXTENDING BUILT-IN PERSISTENT CLASS

```
public class Employee extends PersistentTuple2<PersistentLong, PersistentString> {  
    public Employee(PersistentLong id, PersistentString name) {  
        setId(id);  
        setName(name);  
    }  
    public PersistentLong getId() {  
        return _1();  
    }  
    public void setId(PersistentLong id) {  
        _1(id);  
    }  
}
```

```
    public PersistentString getName() {  
        return _2();  
    }  
    public void setName(PersistentString name) {  
        _2(name);  
    }  
    public String toString() {  
        return String.format("Employee(%s, %s)", getId(), getName());  
    }  
}
```

CREATING A PERSISTENT CLASS

Non-persistent

```
01 public final class Employee {
02     private final long id;
03     private String name;
04
05
06     public Employee(long id, String name) {
07
08         this.id = id;
09         setName(name);
10     }
11
12
13
14     public long getId() {return id;}
15
16     public String getName() {return name;}
17
18     public void setName(String name) {this.name = name;}
19
20     public int hashCode() {return Long.hashCode(getId());}
21
22     public boolean equals(Object obj) {
23         if (!obj instanceof Employee) return false;
24         Employee emp = (Employee)obj;
25         return emp.getId() == getId() && emp.getName().equals(getName());
26     }
27
28     public String toString() {
29         return String.format("Employee(%d, %s)", getId(), getName());
30     }
31 }
```


CREATING A PERSISTENT CLASS

Persistent

```
01 public final class Employee extends PersistentObject {
02     private static final LongField ID = new LongField();
03     private static final StringField NAME = new StringField();
04     private static final ObjectType<Employee> TYPE = ObjectType.withFields(Employee.class, ID, NAME);
05
06     public Employee(long id, PersistentString name) {
07         super(TYPE);
08         setLongField(ID, id);
09         setName(name);
10     }
11
12     private Employee(ObjectPointer<Employee> p) {super(p);}
13
14     public long getId() {return getLongField(ID);}
15
16     public PersistentString getName() {return getObjectField(NAME);}
17
18     public void setName(PersistentString name) {setObjectField(NAME, name);}
19
20     public int hashCode() {return Long.hashCode(getId());}
21
22     public boolean equals(Object obj) {
23         if (!(obj instanceof Employee)) return false;
24         Employee emp = (Employee)obj;
25         return emp.getId() == getId() && emp.getName().equals(getName());
26     }
27
28     public String toString() {
29         return String.format("Employee{id, %s}", getId(), getName());
30     }
31 }
```

LOW-LEVEL MEMORY REGIONS

- Interface from OpenJDK Panama project
- Get and set for byte, short, int, long (on persistent memory)
- Heap API to allocate and free MemoryRegions
- Developers can
 - Retrofit existing code at low-level
 - Create their own abstractions
- Three versions
 - RawMemoryRegion -- useful for volatile use or when caller provides data consistency externally
 - FlushableMemoryRegion -- includes flush() method and fail-safe isFlushed() state
 - TransactionalMemoryRegion -- writes are transactional

APPLICATIONS OF PERSISTENT MEMORY

APPLICATIONS OF PERSISTENT MEMORY

All of In-Memory Computing Applications?

EVERYTHING SOUNDS SO EASY...

- Not so...
- Software innovation – new programming paradigm: “To persist or not to persist”
 - Think early days of the smart phone
 - Any write could be your last write – do you need the data when the application restarts?
 - More than just large memory
- Existing software needs to be re-architected – to unlock features and performance
 - Apache Cassandra, Apache Spark
- Traditional memory still in every system – applications need to be aware

CALL TO ACTION

- Innovate on persistence – discover usages!!
- Feedback on Java persistent programming model

JOIN THE DISCUSSION

- Learn about the Persistent Memory programming model - <http://www.snia.org/forums/sssi/nvmp>
- Join the pmem NVM Libraries Open Source project - <http://pmem.io>
- Read the documents and code supporting ACPI 6.1 and Linux NFIT drivers
 - http://www.uefi.org/sites/default/files/resources/ACPI_6.1.pdf
 - <https://github.com/pmem/ndctl>
 - <http://pmem.io/documents/>
 - <https://github.com/01org/prd>
- Intel Architecture Instruction Set Extensions Programming Reference
 - <https://software.intel.com/en-us/intel-isa-extensions>
- Intel 3D XPoint™ Memory
 - <https://software.intel.com/en-us/persistent-memory>

Thank You