

# Virtualization Aware JVM

Making the most of a virtualized environment

Tim Ellison – Hursley labs



## 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.

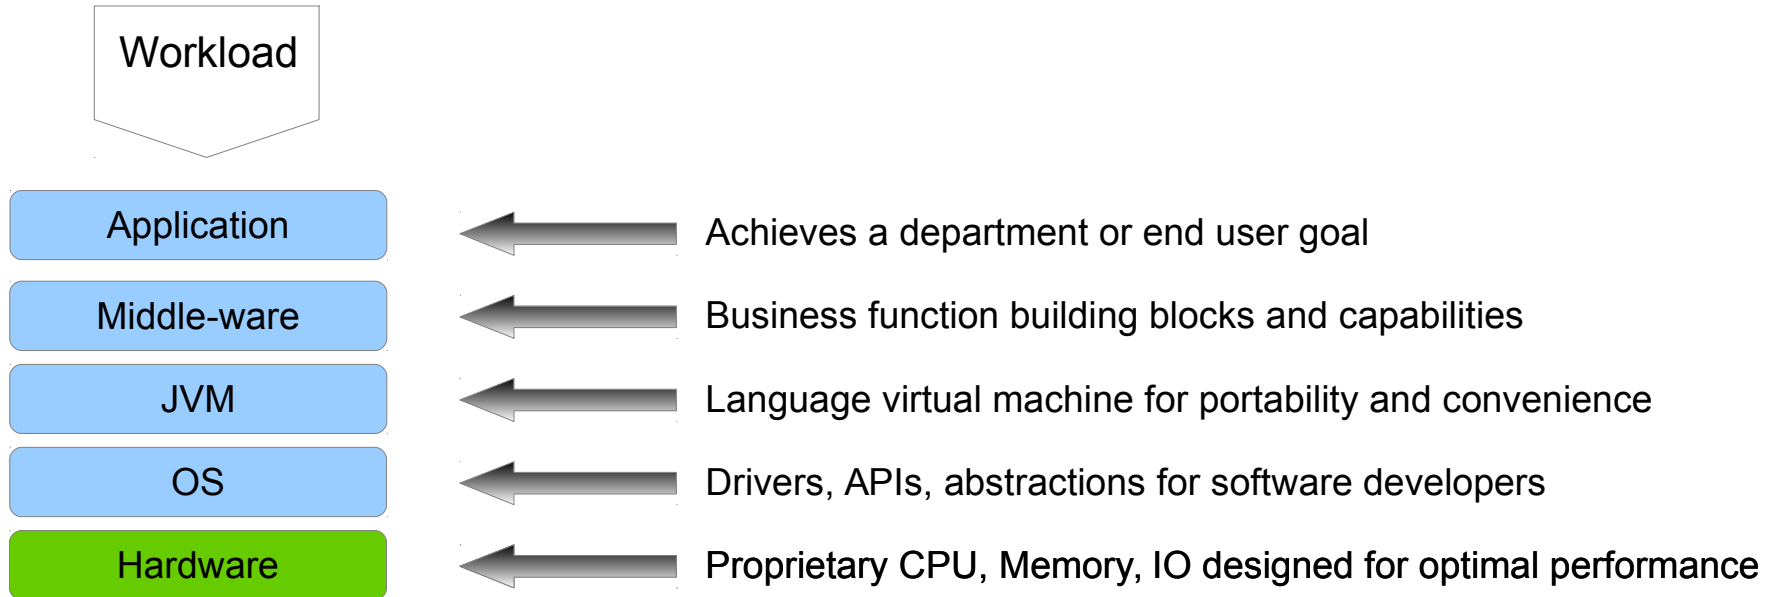
NOTHING CONTAINED IN THIS PRESENTATION IS INTENDED TO, OR SHALL HAVE THE EFFECT OF:

- CREATING ANY WARRANT OR REPRESENTATION FROM IBM, ITS AFFILIATED COMPANIES OR ITS OR THEIR SUPPLIERS AND/OR LICENSORS

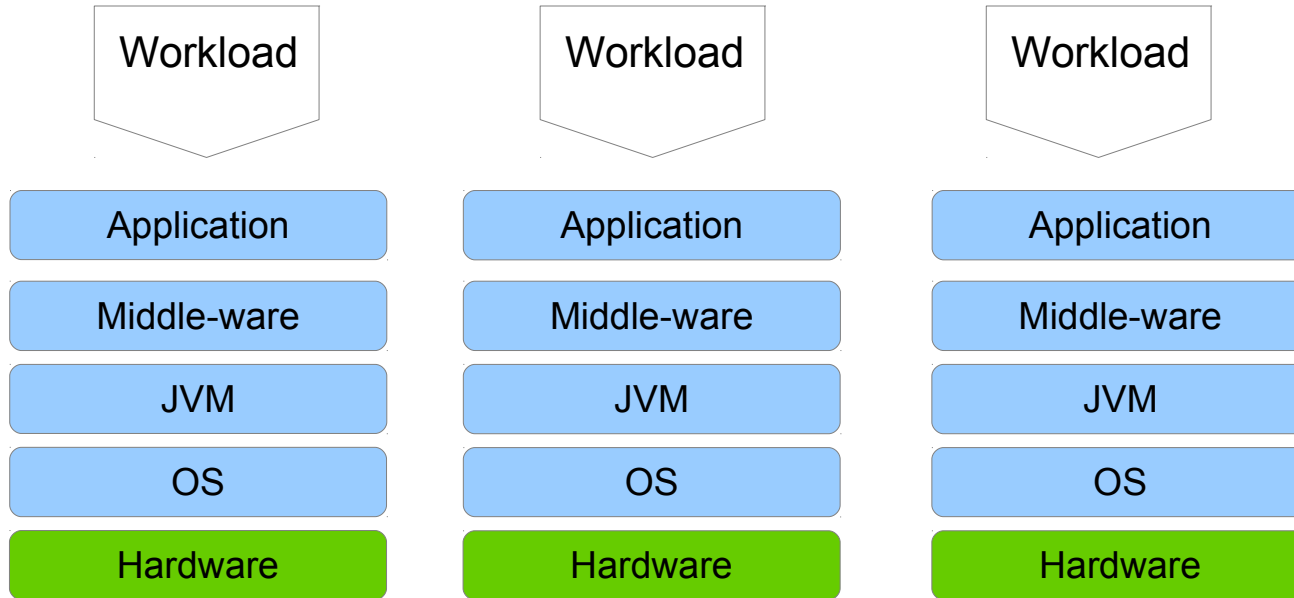
## Goals of this talk

- Describe the impact of deep systems virtualization on the Java Virtual Machine (JVM) and Java Applications.
- Explain the benefits of making the JVM virtualization aware.
- Discuss techniques to optimize implementations for virtualized environments.

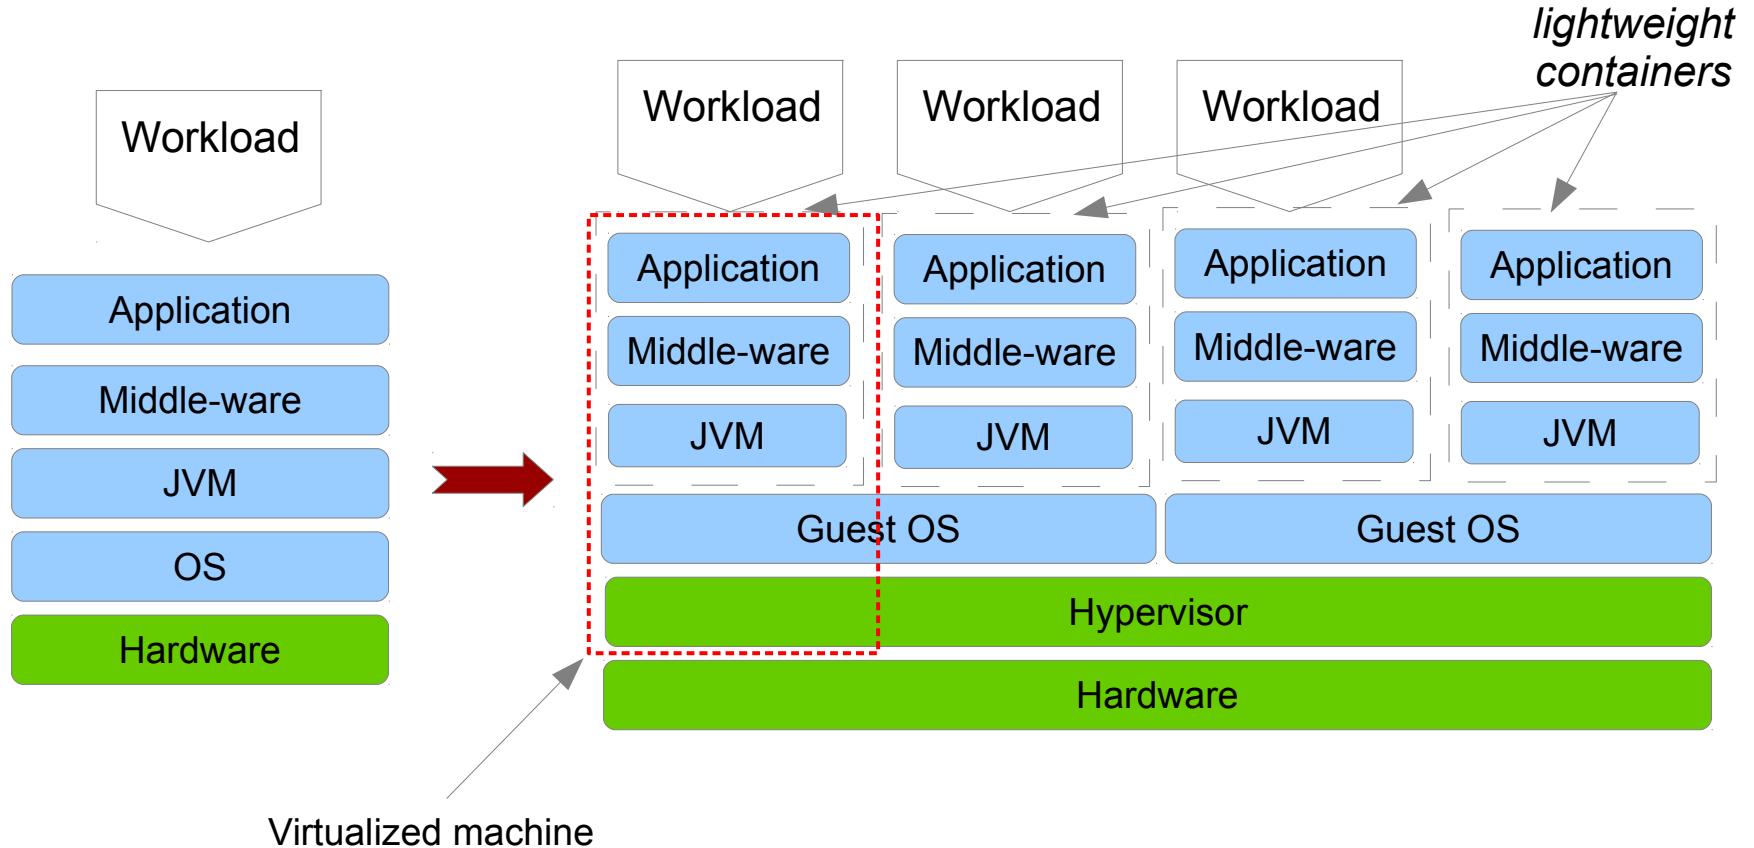
## A stylized system environment



# A data centre environment is no longer what it was...





# A data centre environment is no longer what it was...



## Virtualization available at every layer

- **Hardware Virtualization**
  - Helps pack in more applications compared to bare metal
  - These provide very good isolation at a cost of duplication of services → Higher Overhead
  - Usecase: Multiple customers can be hosted on the same infrastructure.
- **OS Virtualization - Containers**
  - New light weight containers such as Docker, Warden, AIX WPAR, Solaris Containers provide basic isolation mechanism with very little overhead.
  - Usecase: Multiple applications each running inside a container of its own for a single customer.
- **Runtime Virtualization – Virtualization aware Java VM**
  - Runtime Virtualization increases the resources that can be shared while still offering reduced isolation.
  - Usecase: Several instances of an application hosted on a single instance of Java for a single customer.

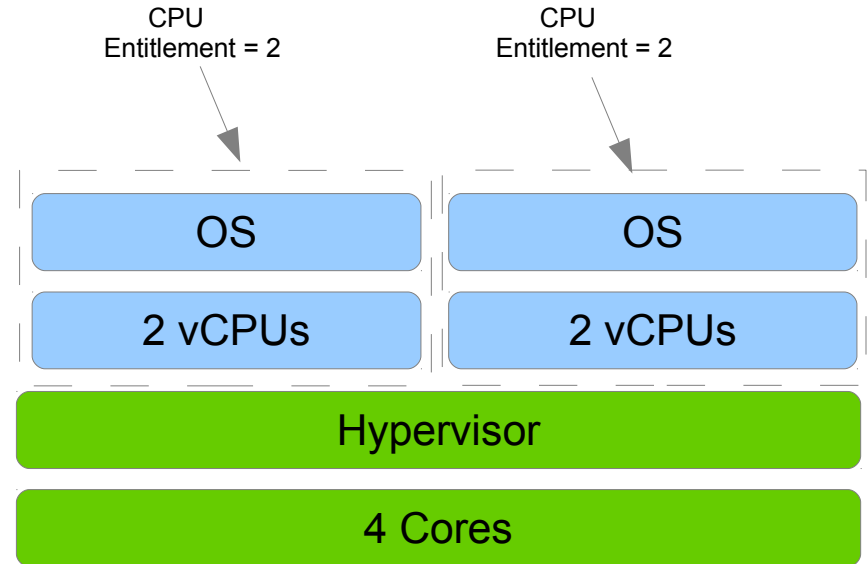
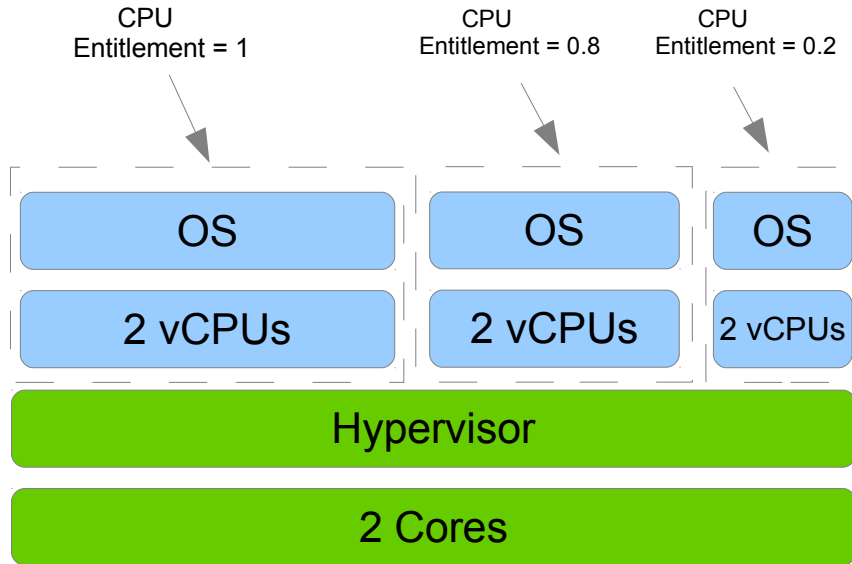
# What does this virtualization story mean to the JVM ?



- Fixed view of hardware resources.
- JVM resources do not change during application life-cycle.
- Configuration usually determined at start-up, and doesn't change during the entire run.
- Resources are elastic. Hypervisors might “re-assign” resources based on need.
- Resource reporting APIs are unreliable !
- Resources are shared, and limits can change dynamically.



# Physical View vs. Logical View – “How many CPUs do I have?”



## Am I Real or Am I Virtual ?!

Q. Do Java **applications** need to know if they are running inside a Guest OS ?

A. Mostly No. Only a small class of applications benefit

- ◆ Resource Orchestrators and Load Balancers
- ◆ Monitoring tools
- ◆ Debuggers and RAS Tools




Q. Does the **JVM implementation** need to know if it is running inside a Guest OS ?

A. Absolutely Yes !

- ◆ Manage Java resources based upon real physical resources.
- ◆ Provide virtualization info to applications.

## Benefits of a virtualization aware JVM




- Make the most efficient use of resources in a virtualized environment.
- Propagate the knowledge up the software stack to enable load balancers to take appropriate decisions.
- Remove necessity for multiple middle-ware products to understand intricacies of hypervisors.
- Provide unified interface to be able to deal with a multitude of hypervisors.

## IBM's JMX Beans for Virtualization

- `com.ibm.lang.management.OperatingSystemMXBean` – with IBM Extensions
  - Extended OS usage statistics – Logical view
    - Processor
    - Memory
  - OS Support: AIX, Linux, Windows and z/OS
- `com.ibm.virtualization.management.HypervisorMXBean`
  - Detect if we are running on a hypervisor.
  - Hypervisor Details (Currently only Vendor Name string)
  - Hypervisors Supported: z/VM, PR/SM, PowerVM, VMWare, KVM (x86 and Power), Hyper-V
- `com.ibm.virtualization.management.GuestOSMXBean`
  - Guest OS usage statistics as seen from the hypervisor – Physical View
  - Usage Statistics
    - Processor
    - Memory
  - AIX & Linux on PowerVM, Linux and Windows on VMWare, z/OS & zLinux on z/VM

# JMX Beans for Virtualization



# OperatingSystemMXBean – New APIs

com.ibm.lang.management

## Interface `OperatingSystemMXBean`

<code>MemoryUsage</code>	<code>retrieveMemoryUsage()</code> Instantiates and returns an instance of <code>MemoryUsage</code> object that represents the current snapshot of Memory usage statistics.
<code>MemoryUsage</code>	<code>retrieveMemoryUsage(MemoryUsage memObj)</code> Returns an updated <code>MemoryUsage</code> object that represents the current snapshot of Memory usage statistics.
<code>ProcessorUsage[]</code>	<code>retrieveProcessorUsage()</code> Instantiates and returns an array of <code>ProcessorUsage</code> objects that represent the current snapshot of individual Processor usage times.
<code>ProcessorUsage[]</code>	<code>retrieveProcessorUsage(ProcessorUsage[] procArray)</code> Returns an updated array of <code>ProcessorUsage</code> objects that represent the current snapshot of individual Processor usage times.
<code>ProcessorUsage</code>	<code>retrieveTotalProcessorUsage()</code> Instantiate and return a new <code>ProcessorUsage</code> object that represents the current snapshot of Processor usage statistics.
<code>ProcessorUsage</code>	<code>retrieveTotalProcessorUsage(ProcessorUsage procObj)</code> Returns an updated <code>ProcessorUsage</code> object that represents the current snapshot of Processor usage statistics.

## Method Summary

### Methods

Modifier and Type	Method and Description
long	<b>getBuffered()</b> The physical RAM used for file buffers (in bytes) or -1 if not available on the underlying machine.
long	<b>getCached()</b> The physical RAM used as cache memory (in bytes) or -1 if not available on the underlying machine.
long	<b>getFree()</b> Returns the physical memory unused on the system (in bytes) or -1 if this is not available on the underlying machine.
long	<b>getSwapFree()</b> Total amount of free swap memory (in bytes) or -1 if this is not available on the underlying machine.
long	<b>getSwapTotal()</b> Total swap memory on the machine (in bytes) or -1 if this is not available on the underlying machine.
long	<b>getTimestamp()</b> Returns the timestamp when memory usage statistics were last sampled (in microseconds).
long	<b>getTotal()</b> Returns the total usable physical memory installed on the system (in bytes).

## Method Summary

### Methods

Modifier and Type	Method and Description
long	<b>getBusy()</b> Returns the time spent by the current processor in executing a non-idle thread (in microseconds) or else -1, if this is not available on the underlying machine.
int	<b>getId()</b> Returns a unique identifier assigned to the current processor.
long	<b>getIdle()</b> Returns the time spent by the processor sitting idle (in microseconds) or else -1, if this is not available on the underlying machine.
boolean	<b>getOnline()</b> Tells whether the current processor is online or not.
long	<b>getSystem()</b> Returns the time spent in system mode (in microseconds) or else -1, if this is not available on the underlying machine.
long	<b>getTimestamp()</b> The timestamp when processor usage statistics were last sampled.
long	<b>getUser()</b> Returns the time spent in user mode (in microseconds) or else -1, if this is not available on the underlying machine.
long	<b>getWait()</b> Returns the time spent by the current processor over IO wait (in microseconds) or else -1, if this is not available on the underlying machine.



## Interface HypervisorMXBean

### All Known Implementing Classes:

HypervisorMXBeanImpl

### public interface **HypervisorMXBean**

The IBM-specific interface for finding out if the current system is running in a virtualized environment, and if so, information about that environment

## Method Summary

### Methods

Modifier and Type	Method and Description
java.lang.String	<b>getVendor()</b> Returns the vendor of the hypervisor if running in a virtualized environment
boolean	<b>isEnvironmentVirtual()</b> Indicates if the system is running in a virtualized environment or not

## Interface GuestOSMXBean

### All Superinterfaces:

java.lang.management.PlatformManagedObject

### All Known Implementing Classes:

GuestOS

```
public interface GuestOSMXBean
extends java.lang.management.PlatformManagedObject
```

The IBM Specific MXBean interface that provides Guest OS statistics as seen by the hypervisor host

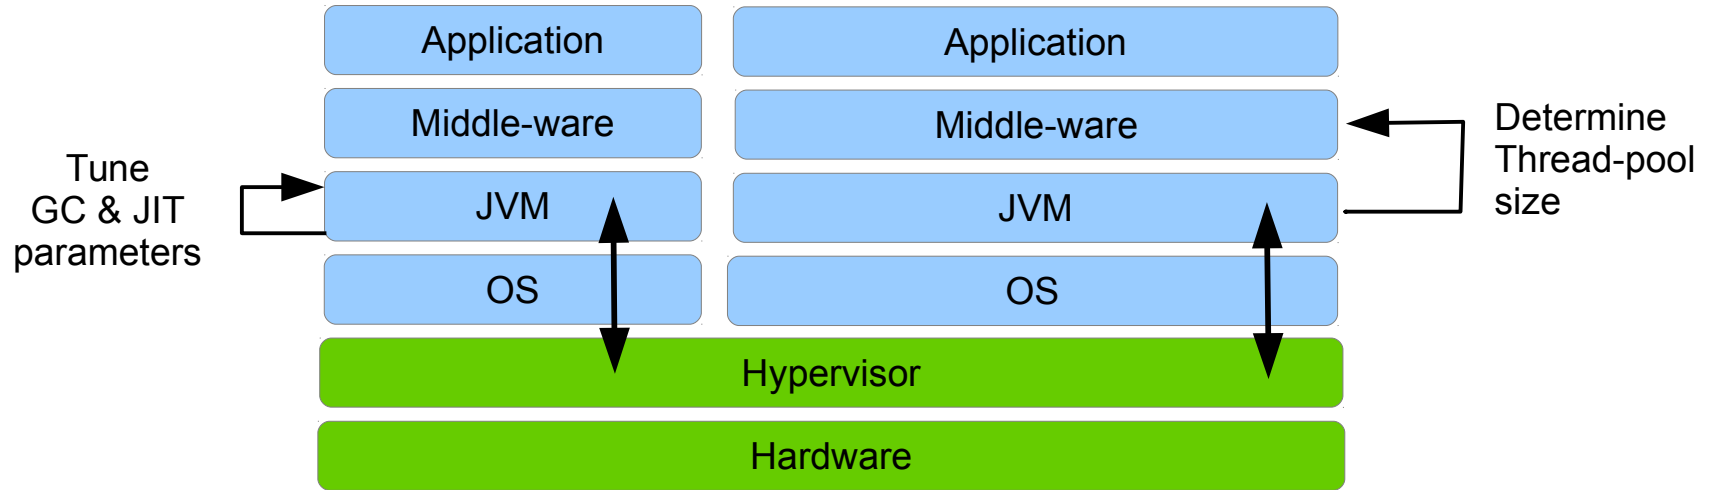
### See Also:

GuestOS


## Method Summary

Modifier and Type	Method and Description
<b>GuestOSMemoryUsage</b>	<b>retrieveMemoryUsage()</b> Function instantiates a GuestOSMemoryUsage object with the current snapshot of Memory Usage statistics of the guest as seen by the hypervisor.
<b>GuestOSMemoryUsage</b>	<b>retrieveMemoryUsage(GuestOSMemoryUsage gmUsage)</b> Function updates the user provided GuestOSMemoryUsage object with Memory Usage statistics of the guest as seen by the hypervisor.
<b>GuestOSProcessorUsage</b>	<b>retrieveProcessorUsage()</b> Function instantiates a GuestOSProcessorUsage object with the current snapshot of Processor Usage statistics of the guest as seen by the hypervisor.
<b>GuestOSProcessorUsage</b>	<b>retrieveProcessorUsage(GuestOSProcessorUsage gpUsage)</b> Function updates the user provided GuestOSProcessorUsage object with Processor Usage statistics of the guest as seen by the hypervisor.

# GuestOSMXBean use-case – Tune internals of JVM & Websphere



# OperatingSystemMXBean use-case – Load Balance JVM Instances



## Implementation Notes

- The Beans are available in IBM Java SDK version 7.1 onwards
- Hypervisor specific setup may be required to obtain usage data.
  - VMWare requires VMGuestLib to be installed (Is part of VMware tools)
  - Hypfs needs to be mounted on zLinux
- Hypervisor detection known to fail on certain older Intel processors
  - IBM\_JAVA\_HYPERVISOR\_SETTINGS environment variable can be used as a workaround
  - See javadoc for HypervisorMXBean for more details.
- Some data may not be available on specific OS / Hypervisor combinations.
  - Javadoc should have all relevant info.
- Currently only one level of hypervisor is supported. In case of multiple layers of hypervisor (e.g. on zSystems), only the top level hypervisor info is returned.
- Javacore now contains virtualization info.


## Additional virtualization optimizations

# Dynamic Heap Adjustment (softmx)

- -Xmx
  - absolute limit, fixed at startup
- -Xsoftmx
  - soft limit  $\leq$  -Xmx, set dynamically through JMX
  - Garbage collector tries to shrink to softmx over time
  - Once at target will not expand beyond it
- OS Interaction
  - JVM advises OS when memory freed
  - Effectiveness depends on OS support
- Use cases
  - Cap early, grow later
  - Shrink to free unused memory

The screenshot shows the JMX console for a JVM process with pid 3716. The 'VM Summary' tab is active, displaying various JVM attributes and their values. The 'Attribute values' table is as follows:

Name	Value
GCMode	-Xgcpolicy:gencon
HeapMemoryUsage	javax.management.openmbean....
MaxHeapSize	268435456
MaxHeapSizeLimit	<b>536870912</b>
MinHeapSize	<b>4194304</b>
NonHeapMemoryUsage	javax.management.openmbean....
ObjectPendingFinalizationCount	0
SetMaxHeapSizeSupported	true
SharedClassCacheFreeSpace	0
SharedClassCacheSize	0
TotalPhysicalMemory	<b>8118067200</b>
UsedPhysicalMemory	<b>4163928064</b>
Verbose	false



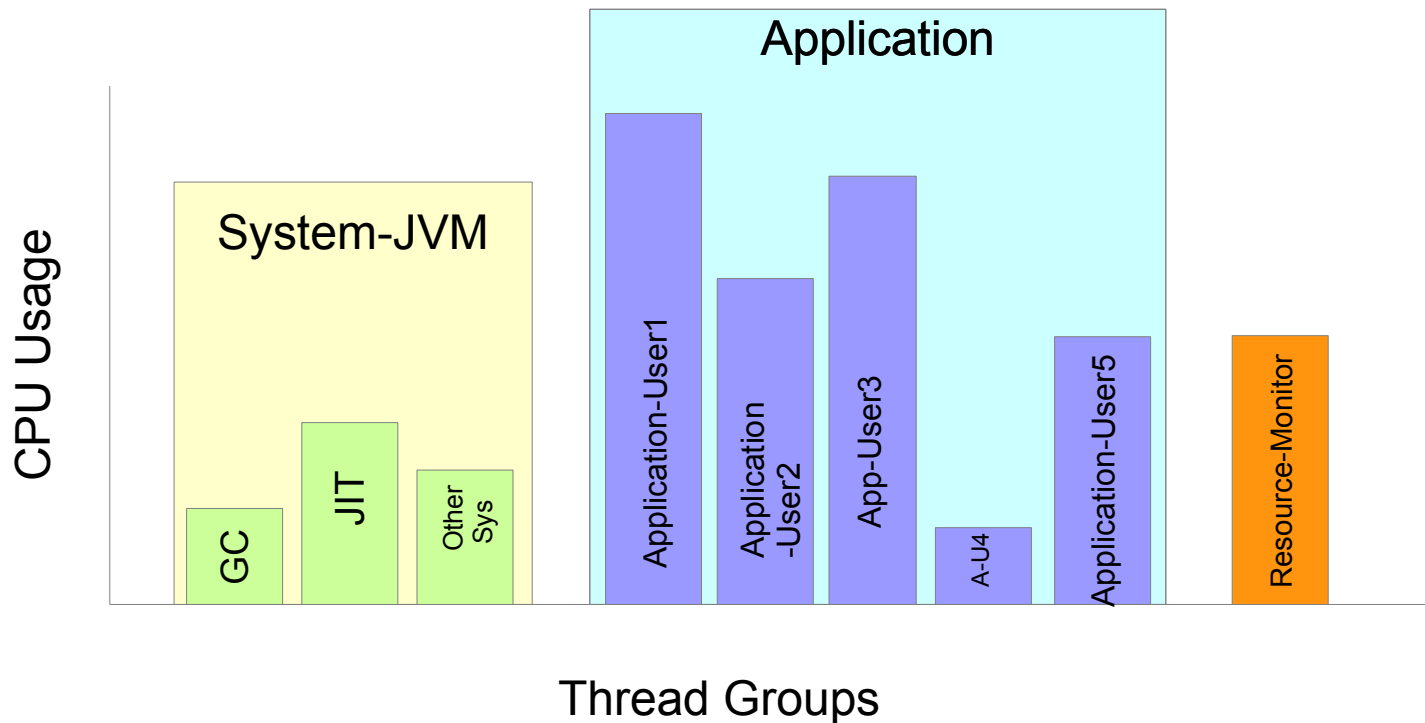
## What is your application doing when Idle ?

- Does your application use lots of CPU even when Idle ?
- Historically poor Java idle behavior causes CPU burn
  - Inefficient CPU usage in the Cloud
    - Starves other VM's / JVM instances
  - Increased client costs for CPU Usage
    - Especially so on zSystems
- Memory hoarding of Idle JVM's.
  - Reduce memory footprint when idle.
    - “Boilerplate” applications in Bluemix that are never used, but are taking up lots of memory.
- First step is to be able to measure precise JVM CPU usage





## Ability to categorize JVM threads for CPU usage



# JvmCpuMonitorMXBean

com.ibm.lang.management

## Interface JvmCpuMonitorMXBean

### All Superinterfaces:

java.lang.management.PlatformManagedObject

```
public interface JvmCpuMonitorMXBean
extends java.lang.management.PlatformManagedObject
```

This interface provides APIs to obtain JVM CPU usage information in terms of thread categories. APIs are also available to get and set the thread category.

### Method Summary

#### All Methods

#### Instance Methods

#### Abstract Methods

#### Modifier and Type

#### Method and Description

java.lang.String

`getThreadCategory(long id)`

This function gets the current value of the thread category for the target thread.

JvmCpuMonitorInfo

`getThreadsCpuUsage(JvmCpuMonitorInfo jcmInfo)`

This function updates the user provided JvmCpuMonitorInfo object with CPU usage statistics of the various thread categories.

int

`setThreadCategory(long id, java.lang.String category)`

This function sets the thread category of the target thread.

# JvmCpuMonitorInfo

## Method Summary

All Methods	Static Methods	Instance Methods	Concrete Methods
Modifier and Type	Method and Description		
boolean		<code>equals(java.lang.Object obj)</code>	
static <code>JvmCpuMonitorInfo</code>		<code>from(javax.management.openmbean.CompositeData cd)</code> Receives a <code>CompositeData</code> representing a <code>JvmCpuMonitorInfo</code> object and attempts to return the root <code>JvmCpuMonitorInfo</code> instance.	
long		<code>getApplicationCpuTime()</code> This method returns the total CPU usage for all application threads.	
long[]		<code>getApplicationUserCpuTime()</code> This method returns an array of CPU usage for all user defined thread categories.	
long		<code>getGcCpuTime()</code> This method returns the total CPU usage of all GC threads.	
long		<code>getJitCpuTime()</code> This method returns the total CPU usage of all JIT Threads.	
long		<code>getResourceMonitorCpuTime()</code> This method returns the total CPU usage for all threads of the "Resource-Monitor" category.	
long		<code>getSystemJvmCpuTime()</code> This method returns the total CPU usage of the "System-JVM" category, which includes GC, JIT and other JVM daemon threads.	
long		<code>getTimestamp()</code> This method returns the last sampling time stamp.	
int		<code>hashCode()</code>	
<code>java.lang.String</code>		<code>toString()</code> Text description of this <code>JvmCpuMonitorInfo</code> object.	

## Javacore info

```

^1XMTHDSUMMARY  Threads CPU Usage Summary
NULL          =====
NULL
1XMTHDCATINFO Warning: to get more accurate CPU times for the GC, the option -XX:-ReduceCPUMonitorOverhead can be used
NULL
1XMTHDCATEGORY All JVM attached threads: 50.129351000 secs
1XMTHDCATEGORY |
2XMTHDCATEGORY +--Resource-Monitor: 9.637264000 secs
1XMTHDCATEGORY |
2XMTHDCATEGORY +--System-JVM: 1.104925000 secs
2XMTHDCATEGORY | |
3XMTHDCATEGORY | +--GC: 0.010265000 secs
2XMTHDCATEGORY | |
3XMTHDCATEGORY | +--JIT: 0.486622000 secs
1XMTHDCATEGORY |
2XMTHDCATEGORY +--Application: 39.387162000 secs
2XMTHDCATEGORY | |
3XMTHDCATEGORY | +--Application-User1: 9.337640000 secs
2XMTHDCATEGORY | |
3XMTHDCATEGORY | +--Application-User2: 9.450345000 secs
2XMTHDCATEGORY | |
3XMTHDCATEGORY | +--Application-User3: 9.443287000 secs
2XMTHDCATEGORY | |
3XMTHDCATEGORY | +--Application-User4: 4.235790000 secs
2XMTHDCATEGORY | |
3XMTHDCATEGORY | +--Application-User5: 4.413613000 secs
NULL
NULL          -----

```

## Use Cases for Thread-based CPU Usage

- Reporting transaction CPU usage
  - Identifying "expensive" transactions
  - Reporting JVM overhead over specific intervals
  - Foundation for future work on tracking idle behaviour
  - ...
- 
- Available in IBM Java SDK version 8 onwards

### References

- [http://www-01.ibm.com/support/knowledgecenter/SSYKE2\\_8.0.0/com.ibm.java.api.80.doc/com.ibm.lang.management/com.ibm.lang.management/JvmCpuMonitorMXBean.html](http://www-01.ibm.com/support/knowledgecenter/SSYKE2_8.0.0/com.ibm.java.api.80.doc/com.ibm.lang.management/com.ibm.lang.management/JvmCpuMonitorMXBean.html)
- [http://www-01.ibm.com/support/knowledgecenter/SSYKE2\\_8.0.0/com.ibm.java.lnx.80.doc/diag/appendixes/cmdline/xxreducecpumonitoroverhead.html?lang=en](http://www-01.ibm.com/support/knowledgecenter/SSYKE2_8.0.0/com.ibm.java.lnx.80.doc/diag/appendixes/cmdline/xxreducecpumonitoroverhead.html?lang=en)

## -Xtune:virtualized

- Available from Java 7 SR4 onwards.
- Reduces JVM CPU consumption when Idle (mostly JIT).
- Needs a large shared class cache to maintain peak performance.
- AOT space in the Shared Class Cache (SCC) must not be capped.

## Summary

- Virtualization layers hide “real” resource information.
- JVM needs to know the underlying architecture at start time.
  - It then needs to periodically monitor for any changes.
- Applications can make use of the MBeans to do the same.
- Use only as much memory as you need.
  - Make use of SoftMX to reduce the heap dynamically.
- Idle detection and deep sleep helps cut costs.
  - Use MXBean to monitor Idle behaviour.

---

## Copyright and Trademarks

© IBM Corporation 2015. All Rights Reserved.

IBM, the IBM logo, and [ibm.com](http://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](http://www.ibm.com/legal/copytrade.shtml)