

Zero Copy: A Green Approach to Data Transfer

11, 12 Dec. 2009

Pune, India

IndicThreads.com Conference On Java Technology

Pramod B Nagaraja

IBM India Pvt LTD

Agenda

- Zero Copy
- Data Transfer : Age-Old Practice
- Data Transfer : Best Practice
- Performance Comparison
- Application Scenario
- Limitations
- Conclusion



What is Zero Copy ?

- Data transfer from disk to user buffers without an intermediate copy into OS Kernel buffers
- Computer operations in which CPU does not perform the task of copying data from one memory area to another
 - DMA
 - MMU



What is Zero Copy ?

- Fewer intermediate Data Copies, fewer context switches
- Linux :: sendfile, sendfile64




Agenda

- ✓ Zero Copy
- Data Transfer : Age-Old Practice
- Data Transfer : Best Practice
- Performance Comparison
- Application Scenario
- Limitations
- Conclusion



Application Scenario



Data Transfer : Age-Old Practice

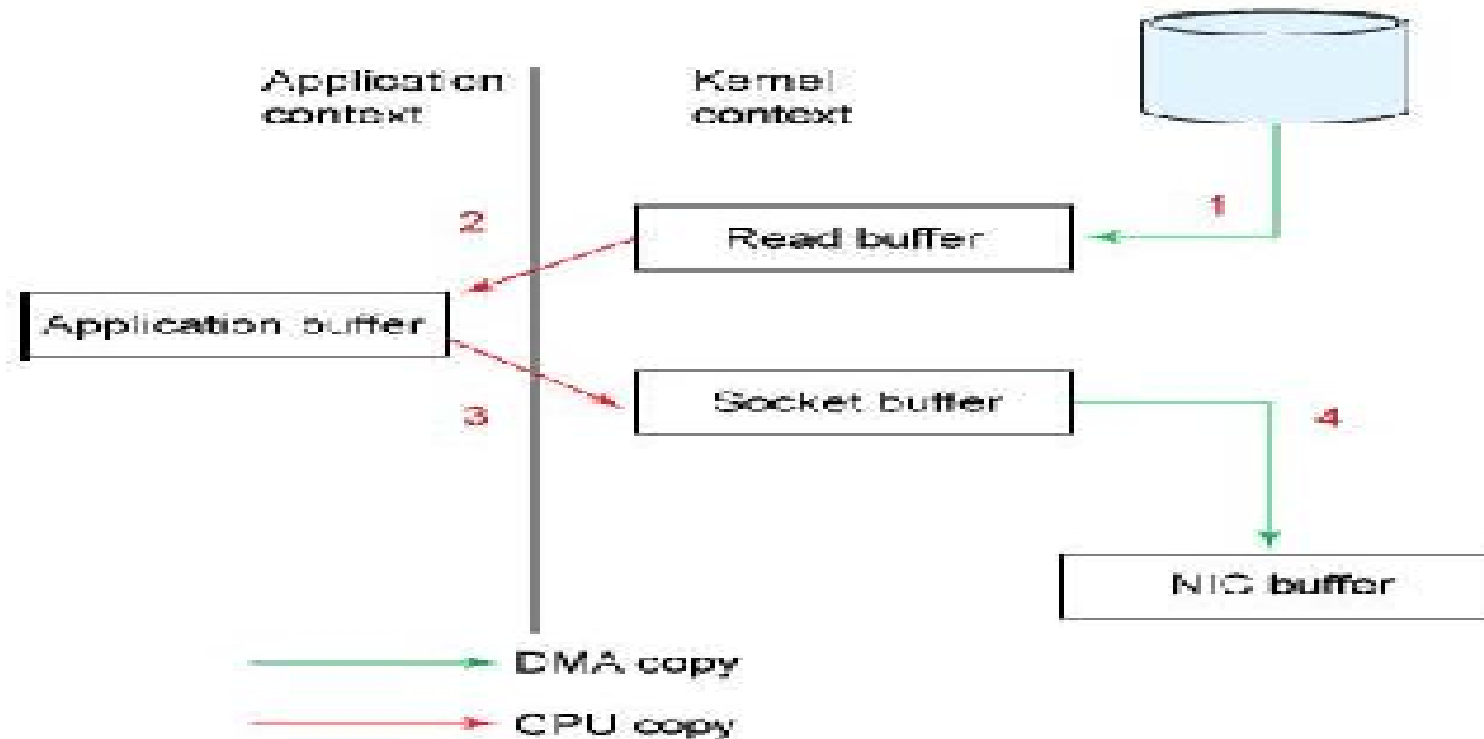
```
File.read(fd, buf, len);  
Socket.send(socket, buf, len);
```

```
InputStream = new FileInputStream(fname);  
output = new  
DataOutputStream(socket.getOutputStream());  
byte[] b = new byte[4096];  
long read = 0, total = 0;  
while((read = InputStream.read(b))>=0) {  
    total = total + read;  
    output.write(b);  
}
```




Data Transfer : Age-Old Practice

- 4 Data Copies



Data Transfer : Age-Old Practice

- 4 Context Switches



Agenda

- ✓ Zero Copy
- ✓ Data Transfer : Age-Old Practice
- Data Transfer : Best Practice
- Performance Comparison
- Application Scenario
- Limitations
- Conclusion



Data Transfer : Best Practice

- **FileChannel Approach**


```
java.nio.channels.FileChannel.transferTo(pos, cnt, writableChannel);
```

```
SocketChannel sc = SocketChannel.open();  
sc.connect(sad);  
FileChannel fc = new  
FileInputStream(fname).getChannel();  
long nsent = 0, curnset = 0;  
curnset = fc.transferTo(0, fsize, sc);
```




Data Transfer : Best Practice

✓ 3 Data Copies



Data Transfer : Best Practice

✓ 2 Context Switches




Agenda

- ✓ Zero Copy
- ✓ Data Transfer : Age-Old Practice
- ✓ Data Transfer : Best Practice
- Performance Comparison
- Application Scenario
- Limitations
- Conclusion



Performance Comparison

- ~ 65% Performance Boost up !



Measurements take on Linux (xi32) with 2.6 kernel



Agenda

- ✓ Zero Copy
- ✓ Data Transfer : Age-Old Practice
- ✓ Data Transfer : Best Practice
- ✓ Performance Comparison
- Application Scenario
- Limitations
- Conclusion



Application Scenario

- FTP Servers
- Web Servers
- Mail Servers
- Any Applications serving static content



Agenda

- ✓ Zero Copy
- ✓ Data Transfer : Age-Old Practice
- ✓ Data Transfer : Best Practice
- ✓ Performance Comparison
- ✓ Application Scenario
- Limitations
- Conclusion



Theoretically Speaking there is no difference
b/w Theory and Practice .. but in Practice
there is !!



Limitations

- Suitable only for *static* content
- Not suitable if Data to be transferred needs to be modified/appended with/appended to, etc
- FileChannels can only be used to transfer data from File to File, File to Socket objects.



References

- More details can be found @ <http://www.ibm.com/developerworks/library/j-zero-copy/index.html>
- <http://www.linuxjournal.com/article/6345>



Q&A



Thank You !!!!

