

BlueArc- Virtual Machine Performance Problem Summary

Edward Simmonds
FEF-SLSMS
01/03/2011

Problem Summary

After deploying a number of virtual machines using OracleVM, Nova users began reporting performance issues when accessing filesystems mounted from BlueArc servers. Initial troubleshooting suggested that the problem occurred when accessing storage served from Blue3, one of the BlueArc storage servers. More testing showed that performance was degraded, but only when reading from the BlueArc server (via NFS).

Further testing revealed a number of other things:

- Problem only occurs when accessing BlueArc-served storage. Accessing other NFS servers does not seem to be an issue (NetApp filers and Linux NFS servers).
- Problem exists only when reading data. Writing data achieves near-native performance regardless of the source writing to the BlueArc.
- Problem is approximately three times worse when accessing the BlueArc from systems in GCC compared with systems in FCC. The BlueArc servers are in FCC.
- Problem exists in a number of different virtualization environments, not just OracleVM, which is Xen-based. Every type of virtual machine we tested (KVM, Xen HVM, Xen PV) showed problems with one exception noted below.
- Performance is slower not only with NFS reads, but also via FTP. (BlueArc can also act as an FTP server.)

Performance Comparison

A sample of performance observations is represented in the following chart. It is important to note that these numbers were compiled during the Christmas holiday period, so they are higher than normal due to the lower load on the BlueArc and the network. The 97.5 MB/s number being reported by the SLF 5.5 XEN HVM in particular is not typical.

	OVM (XEN HVM)	SLF 5.5 (XEN HVM)	SLF 5.5 (KVM)	SLF 5.5 (XEN SR-IOV HVM)	Physical Host
GCC (TCP)	11.1 MB/s	36.6 MB/s	4.8 MB/s	89.9 MB/s	104 MB/s
GCC (UDP)	42.7 MB/s	97.5 MB/s	15.7 MB/s	Test not needed	118 MB/s
FCC (TCP)	37.4 MB/s	Not Tested	Not Tested	Test not needed	118 MB/s
FCC (UDP)	118 MB/s	Not Tested	Not Tested	Test not needed	118 MB/s

All tests were performed against Blue3. All tests were performed using default network settings, and with a block size of 1024. We have tried numerous other variations of network settings, NFS mount settings, and block sizes, none of which improve the performance. Some tests were not performed because a virtualized environment of that type is not available in that location to test with. For the

“Test not needed” cells, the base performance is so high (i.e., 89.9 MB/s), that no further testing is needed.

SR-IOV is a specification that allows a PCIe device to appear as multiple separate physical PCIe devices. These “virtual” PCIe devices can then be attached to individual virtual machines, which improved I/O performance. However, SR-IOV is not supported by the vendor, and therefore cannot be implemented as a solution.

A couple of further notes from the table:

- There is a two to three times performance improvement using UDP.
- There is a three times performance improvement from servers in FCC over nodes in GCC using the same virtualization platform.

Discussion

There are several confounding variables which make further isolation difficult. The first is the location being used to access the BlueArc. FCC outperforms GCC by a factor of three. This performance difference has not been observed in physical nodes.

The second variable is the virtualization platform being used. All VM's we have tried show performance degradation to some degree, with the exception of Xen SR-IOV HVM, which had near line-speed performance. Unfortunately, Xen SR-IOV is not a supported configuration, so we do not believe it is a viable option at this time. Steve Timm, at our request, has tested his virtual machines against the same storage, and has observed similar results-- some of the numbers above are reported from his virtual servers.

The third variable is the storage device providing the data, in this case Blue3, a BlueArc server. NFS mounts to other NFS servers do not exhibit this performance problem. We have been working with the Network and Virtual Services Department, who also contacted BlueArc, but we have been unsuccessful in finding a solution.

Recommendations

At this time, we believe that we have demonstrated that virtual machines will likely experience decreased read performance when accessing BlueArc storage servers, particularly from GCC. It is very difficult to further isolate this because of the confounding variables discussed previously. Therefore, we believe that a higher level discussion needs to take place to answer the following question:

- What is “acceptable performance” for a virtual machine accessing the BlueArc? If the current

speeds being achieved are adequate, we can take the expected performance difference into account when capacity planning and designing future virtualized environments.