Using Virtual Lustre Clients on the WAN for Analysis of Data from High Energy Physics Experiments

CHEP 2012, New York, USA, 2012

D. Bourilkov, P. Avery, M. Cheng, Y. Fu, B. Kim University of Florida
J. Palencia, R. Budden, K. Benninger Pittsburgh Supercomputing Center
D. Shrum, J. Wilgenbusch Florida State University

We describe the work on creating system images of Lustre virtual clients in the ExTENCI project, using several virtual technologies: Xen, VirtualBox, VMware, KVM/QEMU. These virtual machines can be built at several levels, from a basic Linux installation (we use Scientific Linux 5 as an example), adding a Lustre client with Kerberos authentication, and up to complete clients including local or distributed (based on CernVM-FS) installations of the full CERN and project specific software stack for typical LHC experiments. The level, and size, of the images are determined by the users on demand. Various sites and individual users can just download and use them out of the box on Linux/UNIX, Windows and Mac OS X based hosts. We compare the performance of virtual clients with that of real physical systems for typical high energy physics applications like Monte Carlo simulations or analysis of data stored in ROOT trees.

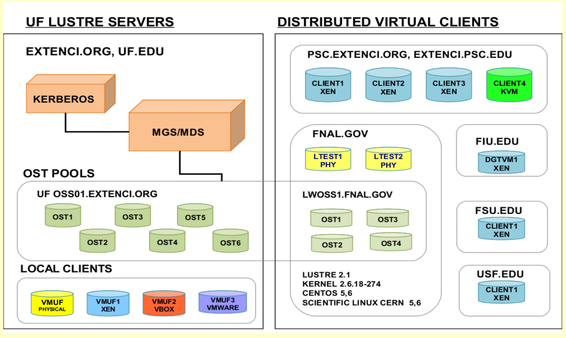

	Xen	VirtualBox	KVM	VMplayer
License	GPL	Oracle	GPL	VMware
Guest SMP	Yes	Yes	Yes	Yes
AMD-V/Intel-VT	Support	Support	Require	Support
Host OS	Modified kernel	Load modules	Load modules	Load modules
Guest OS	modified	native	native	native
Guest I/O	para	emulation	emulation	emulation
GUI	No	Yes	No	Yes
Running snapshot	Yes	Yes	Yes	No
Image format	raw	vdi	qcow2	vmdk

Physical host Configuration (vmuf.extenci.org)
CPU: Quad-Core AMD Opteron™ Processor 2378
# of Cores: 8
Memory: 16 GB
Operating system: Centos 5
Kernel: 2.6.18-274.18.1.el5 for VBox, KVM, and VMplayer
2.6.18-274.17.1.el5xen for Xen
Network 10 Gb/s

Virtual client Configuration
CPU: Quad-Core AMD Opteron ™ Processor 2378
# of Cores: 1
Memory: 2 GB
Operating system: Centos 5
Kernel: 2.6.18-274.17.1.el5xen

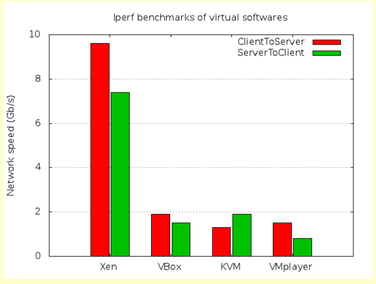

I/O performance of virtual clients using different technologies: Iperf network tests (left) and IOzone tests for access to the Lustre file system (right). Xen gives the best results in all cases.

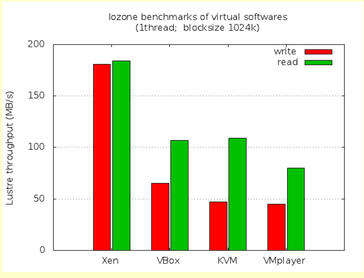

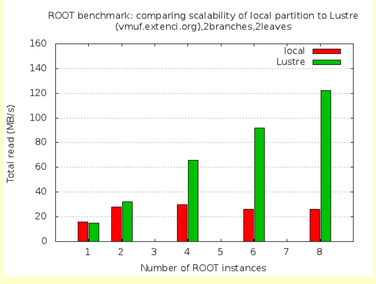

ROOT benchmarks: comparison of reading ROOT trees with different branch/leave structures from local files or from the distributed Lustre file system. Lustre shows good scaling.

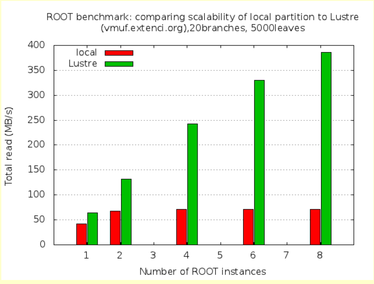

Scalability of our system: for a realistic mix of clients accessing data over the LAN and the WAN, we ran sets of tests first sequentially (stacked), then all clients in parallel. As shown on the left, the simultaneous run produces I/O rates close to the stacked rates.

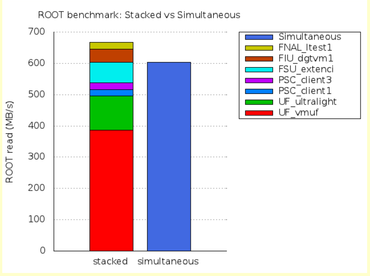

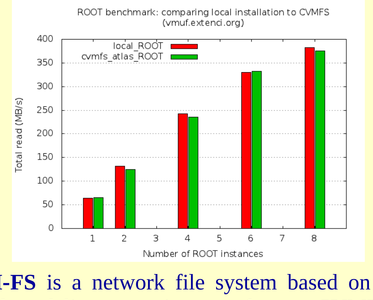

CernVM-FS is a network file system based on HTTP and optimized to deliver experiment software. CVMFS provides complete CMS and ATLAS software installations, enabling the building of "light" virtual clients. Here we compare the **performance of ROOT installed locally, or taken from CVMFS.**

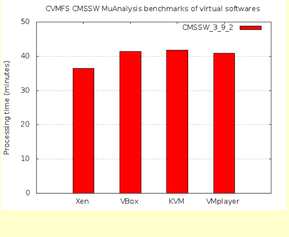

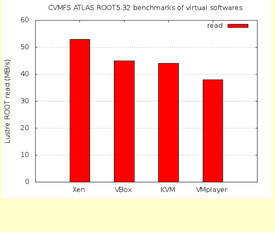

Examples of **virtualized CMS and ATLAS applications.** As the applications are CPU intensive, the difference in I/O rates is less noticeable. Still we get the best performance with Xen.