Preparation of openlab flyers

CERN openlab II Quarterly Review
9 October 2007

Sverre Jarp
CERN openlab CTO

Current prototype

- Elaborated by H.Bjerke with assistance from R.Mondardini

Virtualization

Various virtualization developments in the CERN openlab, testing virtualization technologies and developing solutions to enable virtualization in grid environments.

Benchmarks

Different virtualization technologies behave differently when they are put under stress. A comprehensive benchmarking suite, consisting of both synthetic benchmarks and a set of CERN real-world applications, has been used to assess the performance characteristics of many of the various virtualization technologies in the current virtualization landscape.

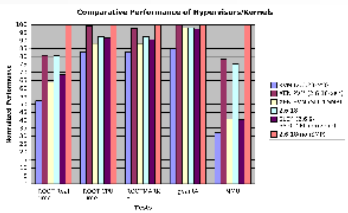

Our measurements show that using Xen guest domains often has a negligible impact on performance compared to a native domain. Sensitive operations - operations that need the intervention of the hypervisor - however, can have a moderate to significant impact on performance. For example, our measurements show that page faults incur a 20% overhead in paravirtualized guest domains and 60% in hardware assisted fully virtualized (HVM) guest domains.

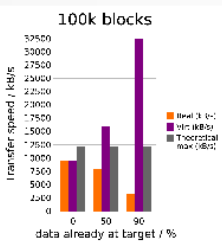

Content-Based Transfer

VM images are big - often several hundred megabytes. The waiting time to transfer images takes away some of the benefits of using VMs and can congest the network. However, most images are relatively similar, so only the delta between the new image and already existing images needs to be transferred.

Identical blocks between images can be identified with a hash checksum. A comparison of the filesystems of two batch machines in the computer center showed that between them, 84% of the filesystem blocks were identical.

The development of a Content-Based Transfer tool has shown that a speedup of about 2.6 can be achieved for the transfer of a 400 MB image to a target machine with 90% identical blocks.

OS Farm

Various projects at CERN need quick deployment of clean machines. For this, a rich set of Linux flavours is needed. The default distribution for CERN is Scientific Linux CERN 4 (SLC4), but in transitional periods, different versions of SLC are needed. Also, Debian, Ubuntu, CentOS and Scientific are desirable.

OS Farm is a complementary tool for creating VM images for Xen, on demand. A web interface provides the user with image configuration options, such as a selection of yum packages. Also, some virtual appliances are provided, e.g. for the gLite grid middleware.

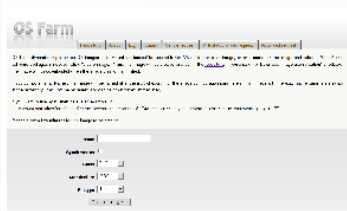

OS Farm uses LVM snapshots for instantaneous copy-on-write base stages that can be reused for further package configurations.

Outlook

Will continue measuring the performance of different virtualization technologies. Further optimizations to Content-Based Transfer will be investigated with Perfmon. Content-Based Transfer will be integrated with OS Farm.

Project name	Virtualization
Project number	001
Participants	Havard Bjerke, Irfan Habib, Dharminder Rattu
Partner participation	Intel
CERN interested parties	Grid Deployment Group, Fabric Infrastructure and Operations, Physics Experiments
Documents	www.cern.ch/osfarm twiki.cern.ch/twiki/bin/view/Virtualization/WebHome

www.cern.ch/openlab

Aim

- Provide portfolio of flyers
  - Intended as recto/verso A4
    - But scalable up and down
      - Also to poster size
- One flyer per main project
  - Not per activity!
- Description of current activities and outlook
- Key information in box
  - Project name/number
  - Participants (CERN and partners)
  - Interested parties
  - Documentation pointers
- Regular updates
  - Couple of times per year

Simple recipe

- Based on Scribus
  - Install the product (windows: http://prdownloads.sourceforge.net/scribus/scribus-1.3.3.9-win32-install.exe?download)
  - Download the project template: http://hbjerke.web.cern.ch/hbjerke/project.sla.gz
  - Open the template in Scribus

Q & A