Linux-VServer

DI Herbert Pötzl

1 Introduction

Computers have become sufficiently powerful to use virtualization to create the illusion of many smaller virtual machines, each running a separate operating system instance.

- Virtual Machines
- System Emulators
- Partitioning

2 The Concept

Virtual Servers do not necessarily require a separate operating system for each instance – many resources can be shared.

Context Isolation allows to put several *Servers* on a Host, which will then share the available resources in an efficient manner.

2.1 Advantages

- ✘ Minimal Overhead
- ✘ Hardware Abstraction
- ✘ Shared Resources

2.2 Possible Drawbacks

- ✘ A single Kernel
- ✘ Security Issues?

3 Nomenclature

Host is the real or virtual machine running the Linux-VServer enabled Kernel

Guest is the virtual private server (or short VPS) composed of a chrooted environment, isolated processes, and restricted IP ranges.

Context is the isolated and partially virtualized *environment* to which processes are *confined*

4 Existing Infrastructure

- ✘ Linux Capability System
- ✘ Resource Limits (ulimit)
- ✘ File Attributes (xattr)
- ✘ The chroot(1) Command

5 Required Modifications

- ✘ Context Separation
- ✘ Network Separation
- ✘ The Chroot Barrier
- ✘ Upper Bound for Caps
- ✘ Resource Isolation
- ✘ Filesystem XID Tagging

6 Additional Modifications

- ✘ Context Flags
- ✘ Context Capabilities
- ✘ Context Accounting
- ✘ Context Limits
- ✘ Virtualization
- ✘ Improved Security
- ✘ Kernel Helper

7 Features and Bonus Material

- ✘ Unification
- ✘ Private Namespaces
- ✘ The Linux-VServer Proc-FS
- ✘ Token Bucket Extensions
- ✘ Context Disk Limits
- ✘ Per-Context Quota
- ✘ The VRoot Proxy Device

8 Field of Application

- ✘ Administrative Separation
- ✘ Service Separation
- ✘ Enhancing Security
- ✘ Easy Maintenance
- ✘ Fail-over Scenarios
- ✘ Simplified Testing

9 Performance and Stability

- ✘ Impact of Linux-VServer on the Host
- ✘ Overhead inside a Context
- ✘ Size of the Kernel Patch

patch	lines	chars	hunks	new
vs1.00	2845	95567	178	997
vs1.20	4305	131922	216	1857
vs1.2.10	4820	149187	252	1968
vs1.9.0	11382	326042	494	5396
vs2.00	19673	557988	856	8987
2.6.11Δ	682454	21905964	23506	108544

10 Non Intel x86 Hardware

- ✔ ia64, x86_64
- ✔ alpha, arm
- ✔ hppa, hppa64
- ✔ ppc, ppc64
- ✔ sparc, sparc64
- ✔ mips o/n32, mips64
- ✔ s390, s390x
- ✔ um, xen

Examples ...

Q & A

www: http://linux-vserver.org
irc: #vserver @ irc.oftc.net