
Subject: [ANNOUNCE] OpenVZ patch for 2.6.15 kernel
Posted by [Kirill Korotaev](#) on Thu, 02 Feb 2006 11:59:06 GMT
[View Forum Message](#) <> [Reply to Message](#)

OpenVZ team is happy to announce the release of its virtualization solution based on 2.6.15 kernel.

As in previous releases, OpenVZ 2.6.15 kernel patch includes:

- virtualization
- fine grained resource management (user beancounters)
- 2 level disk quota
- a number of mainstream fixes (umount race, proc locking etc.).

More information about OpenVZ project is available at <http://openvz.org/>

Fine grained broken-out patch set can be found at
<http://download.openvz.org/kernel/broken-out/2.6.15-025stab012.1/>

About OpenVZ software

~~~~~

OpenVZ is a kernel virtualization solution which can be considered as a natural step in the OS kernel evolution: after multiuser and multitasking functionality there comes an OpenVZ feature of having multiple environments.

Virtualization lets you divide a system into separate isolated execution environments (called VPSs - Virtual Private Servers). From the point of view of the VPS owner (root), it looks like a stand-alone server. Each VPS has its own filesystem tree, process tree (starting from init as in a real system) and so on. The single-kernel approach makes it possible to virtualize with very little overhead, if any.

OpenVZ in-kernel modifications can be divided into several components:

#### 1. Virtualization and isolation.

Many Linux kernel subsystems are virtualized, so each VPS has its own:

- process tree (featuring virtualized pids, so that the init pid is 1);
- filesystems (including virtualized /proc and /sys);
- network (virtual network device, its own ip addresses, set of netfilter and routing rules);
- devices (if needed, any VPS can be granted access to real devices like network interfaces, serial ports, disk partitions, etc);
- IPC objects.

#### 2. Resource Management.

This subsystem enables multiple VPSs to coexist, providing managed resource sharing and limiting.

- User Beancounters is a set of per-VPS resource counters, limits, and guarantees (kernel memory, network buffers, phys pages, etc.).
- Two-level disk quota (first-level: per-VPS quota; second-level: ordinary user/group quota inside a VPS)

Resource management is what makes OpenVZ different from other solutions of this kind (like Linux VServer or FreeBSD jails). There are a few resources that can be abused from inside a VPS (such as files, IPC objects, ...) leading to a DoS attack. User Beancounters prevent such abuses.

As virtualization solution OpenVZ makes it possible to do the same things for which people use UML, Xen, QEmu or VMware, but there are differences:

- (a) there is no ability to run other operating systems (although different Linux distros can happily coexist);
- (b) performance loss is negligible due to absence of any kind of emulation;
- (c) resource utilization is much better.

Thanks,  
OpenVZ team.

---