


λ vps AdminOS

Pavel Šnajdr


InstallFest 2018



- Started with OpenVZ in 2009

Since day 1

containers are viewed as the ultimate future
of shared computing on x86



- **1753 containers** (vcpu/4G/120G as minimum)
 - 3506 vcores + 7T RAM allocated
- **only 17 servers!**
 - 216 cores + HT, 4.5T RAM, ~4.25 kW
 - 1.8T CTs, 1.3T in ZFS ARC, 1.4T reported free

Node list

	#	Name	VPS	Up	Load	%iowait	%idle	Free mem	ARC	%hit	Version	Kernel	
	5	vpsadmin.prg	0	21.0	1.5	0.00	97.08	-10.81	0.00	0.00	2.9.0	042stab126.666	
	101	node1.prg	0	78.0	0.43	0.00	100.00	202.42	46.75	95.17	2.4.0	042stab112.15	
	102	node2.prg	94	54.8	25	3.49	6.61	35.46	63.76	97.39	2.9.0	042stab126.666	
	103	node3.prg	95	54.8	8.64	0.11	70.96	50.92	60.85	99.70	2.9.0	042stab126.666	
	104	node4.prg	83	48.9	21.79	2.59	20.99	70.67	38.11	99.22	2.9.0	042stab126.666	
	105	node5.prg	93	54.8	10.6	2.74	59.92	41.62	62.78	99.85	2.9.0	042stab126.666	
	106	node6.prg	97	54.8	12.54	0.14	59.60	58.98	57.74	99.90	2.9.0	042stab126.666	
	108	node7.prg	70	54.8	148.11	0.99	45.69	68.77	66.46	99.51	2.9.0	042stab126.666	
	109	node8.prg	95	39.9	6.02	0.14	74.12	56.38	68.21	98.61	2.9.0	042stab126.666	
	110	node9.prg	99	54.8	8.48	7.78	57.99	45.69	72.14	98.82	2.9.0	042stab126.666	
	111	node10.prg	98	54.9	11.52	0.80	28.97	31.33	68.46	99.85	2.9.0	042stab126.666	
	112	node11.prg	94	54.9	4.52	1.37	70.81	47.13	74.91	98.55	2.9.0	042stab126.666	
	113	node12.prg	97	54.8	10.89	0.14	61.90	41.64	39.54	99.09	2.9.0	042stab126.666	
	114	node13.prg	94	54.9	22.45	2.52	22.58	38.36	66.44	99.32	2.9.0	042stab126.666	
	115	node14.prg	105	21.0	37.76	1.17	11.33	54.55	251.50	99.30	2.9.0	042stab126.666	
	210	node1.brq	92	54.9	4.93	0.99	76.03	24.14	76.65	98.93	2.9.0	042stab126.666	
	211	node2.brq	133	41.3	13.92	0.53	28.24	18.78	37.69	99.60	2.9.0	042stab126.666	
	212	node3.brq	132	43.5	12.13	2.50	72.18	36.22	46.45	98.94	2.9.0	042stab126.666	
	213	node4.brq	8	121.4	0.15	0.00	100.00	237.61	7.12	55.59	2.9.0	042stab124.2	
	300	node1.pgnd	68	55.6	2.86	0.00	85.15	47.54	125.46	99.90	2.9.0	042stab126.666	
	301	node2.pgnd	0	0.0	0.21	1.39	91.92	218.91	0.00	2.57	2.8.0	042stab125.5	
	160	backuper.prg	0	61.1	1.48	0.42	81.09	-70.21	62.87	77.53	2.9.0	042stab124.2	
	170	nasbox.prg	0	110.9	3.22	1.27	89.03	-167.39	109.19	75.54	2.9.0	042stab125.5	



- OpenVZ 6 “Legacy”
 - vzctl ~abandoned (<https://github.com/vpsfreecz/vzctl>)
 - Kernel EOL Nov 2019
- OpenVZ 7
 - Integrated as distro atop RHEL7, no process transparency
No community, “just use Virtuozzo 7 images” kind-of-approach.

Searching for replacement...

- Requirements
 - “Full VM look & feel”
 - Reliable isolation
 - Security aspect
 - Resource isolation
 - Powerful storage
 - Easy administration

Requirements for replacement

- “Full VM look & feel”
 - LXC knows how to start such a CT well
 - But LXC alone is management hell...
 - Can we use LXD?

Requirements for replacement

- Reliable isolation
 - User name-space is a must
 - UID/GID offsets remapping edge case?
 - Our members have sometimes 100M+ files on /
 - With opinionated upstream, LXD is a no-go for us
 - Our networking would mean out-of-tree patches for LXD
 - Storage ditto

Requirements for replacement

- Powerful storage
 - ZFSonLinux proven, most IOPS reducing solution
 - Data securely stored (native encryption!)
 - send/receive solved the backup problem

Requirements for replacement

- Easy administration

- Nodes dedicated to containerization

- Custom OS!

- Why not live system?

- No local state → No local surprises.

- Solid foundation to build upon?

- If only there was something pure... and functional...

λ **vp**s AdminOS



NixOS

Nix

- Purely functional package manager
 - Reliable
 - Reproducible
 - Source/binary model
 - Multi-version
 - Rollback

NixOS

- Purely Functional Linux Distribution

```
{
```

```
  boot.loader.grub.device = "/dev/sda";
```

```
  fileSystems."/".device = "/dev/sda1";
```

```
  services.sshd.enable = true;
```

```
}
```

vpsAdminOS

- NixOS without fluff
 - Linux 4.15+, runit, LXC/LXCFS, ZFSonLinux
- Bootable images built per-node from git config
 - PXE, USB, CDROM boot supported, UEFI + Legacy
- NixOps for assuming full control
 - Large scale deployments made easy

vpsAdminOS

- osctl

The most admin-friendly container management tool™

- Manages users for user name-space support
- Control groups management
- Container management
- Template repositories
- # osctl ct top

vpsAdminOS

Quick start

- First install and get Nix running on your workstation, then follow these steps:

```
export NIX_PATH=`pwd`
```

```
git clone https://github.com/sorki/nixpkgs --branch vpsadminos
```

```
git clone https://github.com/vpsfreecz/vpsadminos/ && cd vpsadminos/os
```

```
# Cook up your own configuration
```

```
cp conf_local.nix{.sample,} && vim conf_local.nix
```

```
# build your first vpsAdminOS
```

```
make
```

```
# to test run under qemu
```

```
make qemu
```

vpsAdminOS

- Links

<https://vpsadminos.org/>

<https://github.com/vpsfreecz/vpsadminos>

- IRC

#vpsadminos @ freenode

vpsAdminOS

- Demo

vpsAdminOS

- QA

λ vps AdminOS

\EOF