

Cloud services on top of uKernel

Vasily A. Sartakov

FOSDEM'15

About

- Ksys labs
- Small RnD company
 - Technology transfer (40%)
 - Experimental developments (40%)
 - Applied research (20%)
- Security
 - Embedded & Cyberphysics
 - Networks

About

- Since 2011 working with L4 microkernels
 - Secure screen (2011-2015)
 - Gen.1 OMAP3, L4Linux, L4Re, Fiasco.OC, GUI (Disko)
 - Gen.2 OMAP3/4, Fiasco.OC, Genode
 - Gen.3 IMX6, TZ, Genode
 - Network gateways (2012-2013)
 - X86, L4Linux, Genode/L4Re
 - Stack protection (2013-2014)
 - Non-executable memory support for L4Re
 - X86-based tablet:
 - NOVA, VirtualBox, Genode, Android

Technology Push

- Genode FW goes to general purpose OS, rapidly increase software support, sometimes in detriment of performance.
- SEL4 is just published (half year ago)
- L4Re – another set of libraries, plus L4Linux.
- More and more experimental projects...

Market Pull

- Performance
 - Let's say 10% degradation for any kind of test (network, calculation)
- Security
 - Microkernel is good as concept, Capability – nice, but:
 - Stack protection, ROP protection, IDS
 - Denial-of-Service attacks
 - Resilience

Gap

- Security
- Performance
- Libraries (functionality)

- Something useful, that could become of interested of somebody..

Bridging the gap...


- Network cloud services
 - Key-value store (memcached, redis, leveldb)
 - http (nginx)
 - HAProxy (DevOps way)
 - Distributed filesystems (GlusterFS)
- Embedded
 - Multi-agent platform

Bridging the gap...

- Distributed as configurable VMs. i.e. do not need to build whole Environment tree
- Two applications – one VM
 - REST interface + something
 - REST-based KVS
 - HTTP
 - Filesystem? REST interface to GlusterFS?
- Aim: provide comparable (performance, security, functionality) to Linux-based cloud solution.


Work in progress

- First example (Mid 2014): Redis, NOVA, Genode, real hardware:



Work in progress

- Now (30.01.2015): Redis, NOVA, Genode



Tuning...

- Lwip is faster
- One thread
- Pre allocation
- Notifications, queues, etc..
- Good for start, but not enough for real life

Same for

- Nginx
- Memcached
- HAProxy


Now:

- Benchmarking kernels and environments with custom KVS (without network)

Should we think about Linux

- No, of course
- Our main concurrent: OSv
 - Unikernel
 - Guest VMs
 - Less security, but much more performance:

Redis: OSv vs Linux



Thank you.

sartakov@ksyslabs.org

<http://ksyslabs.org>

<http://ksyslabs.ru>