

Pre-Virtualization: Uniting Two Worlds

Joshua LeVasseur

Universität Karlsruhe (TH), Germany

Volkmar Uhlig

IBM Watson Research Center, NY

Ben Leslie

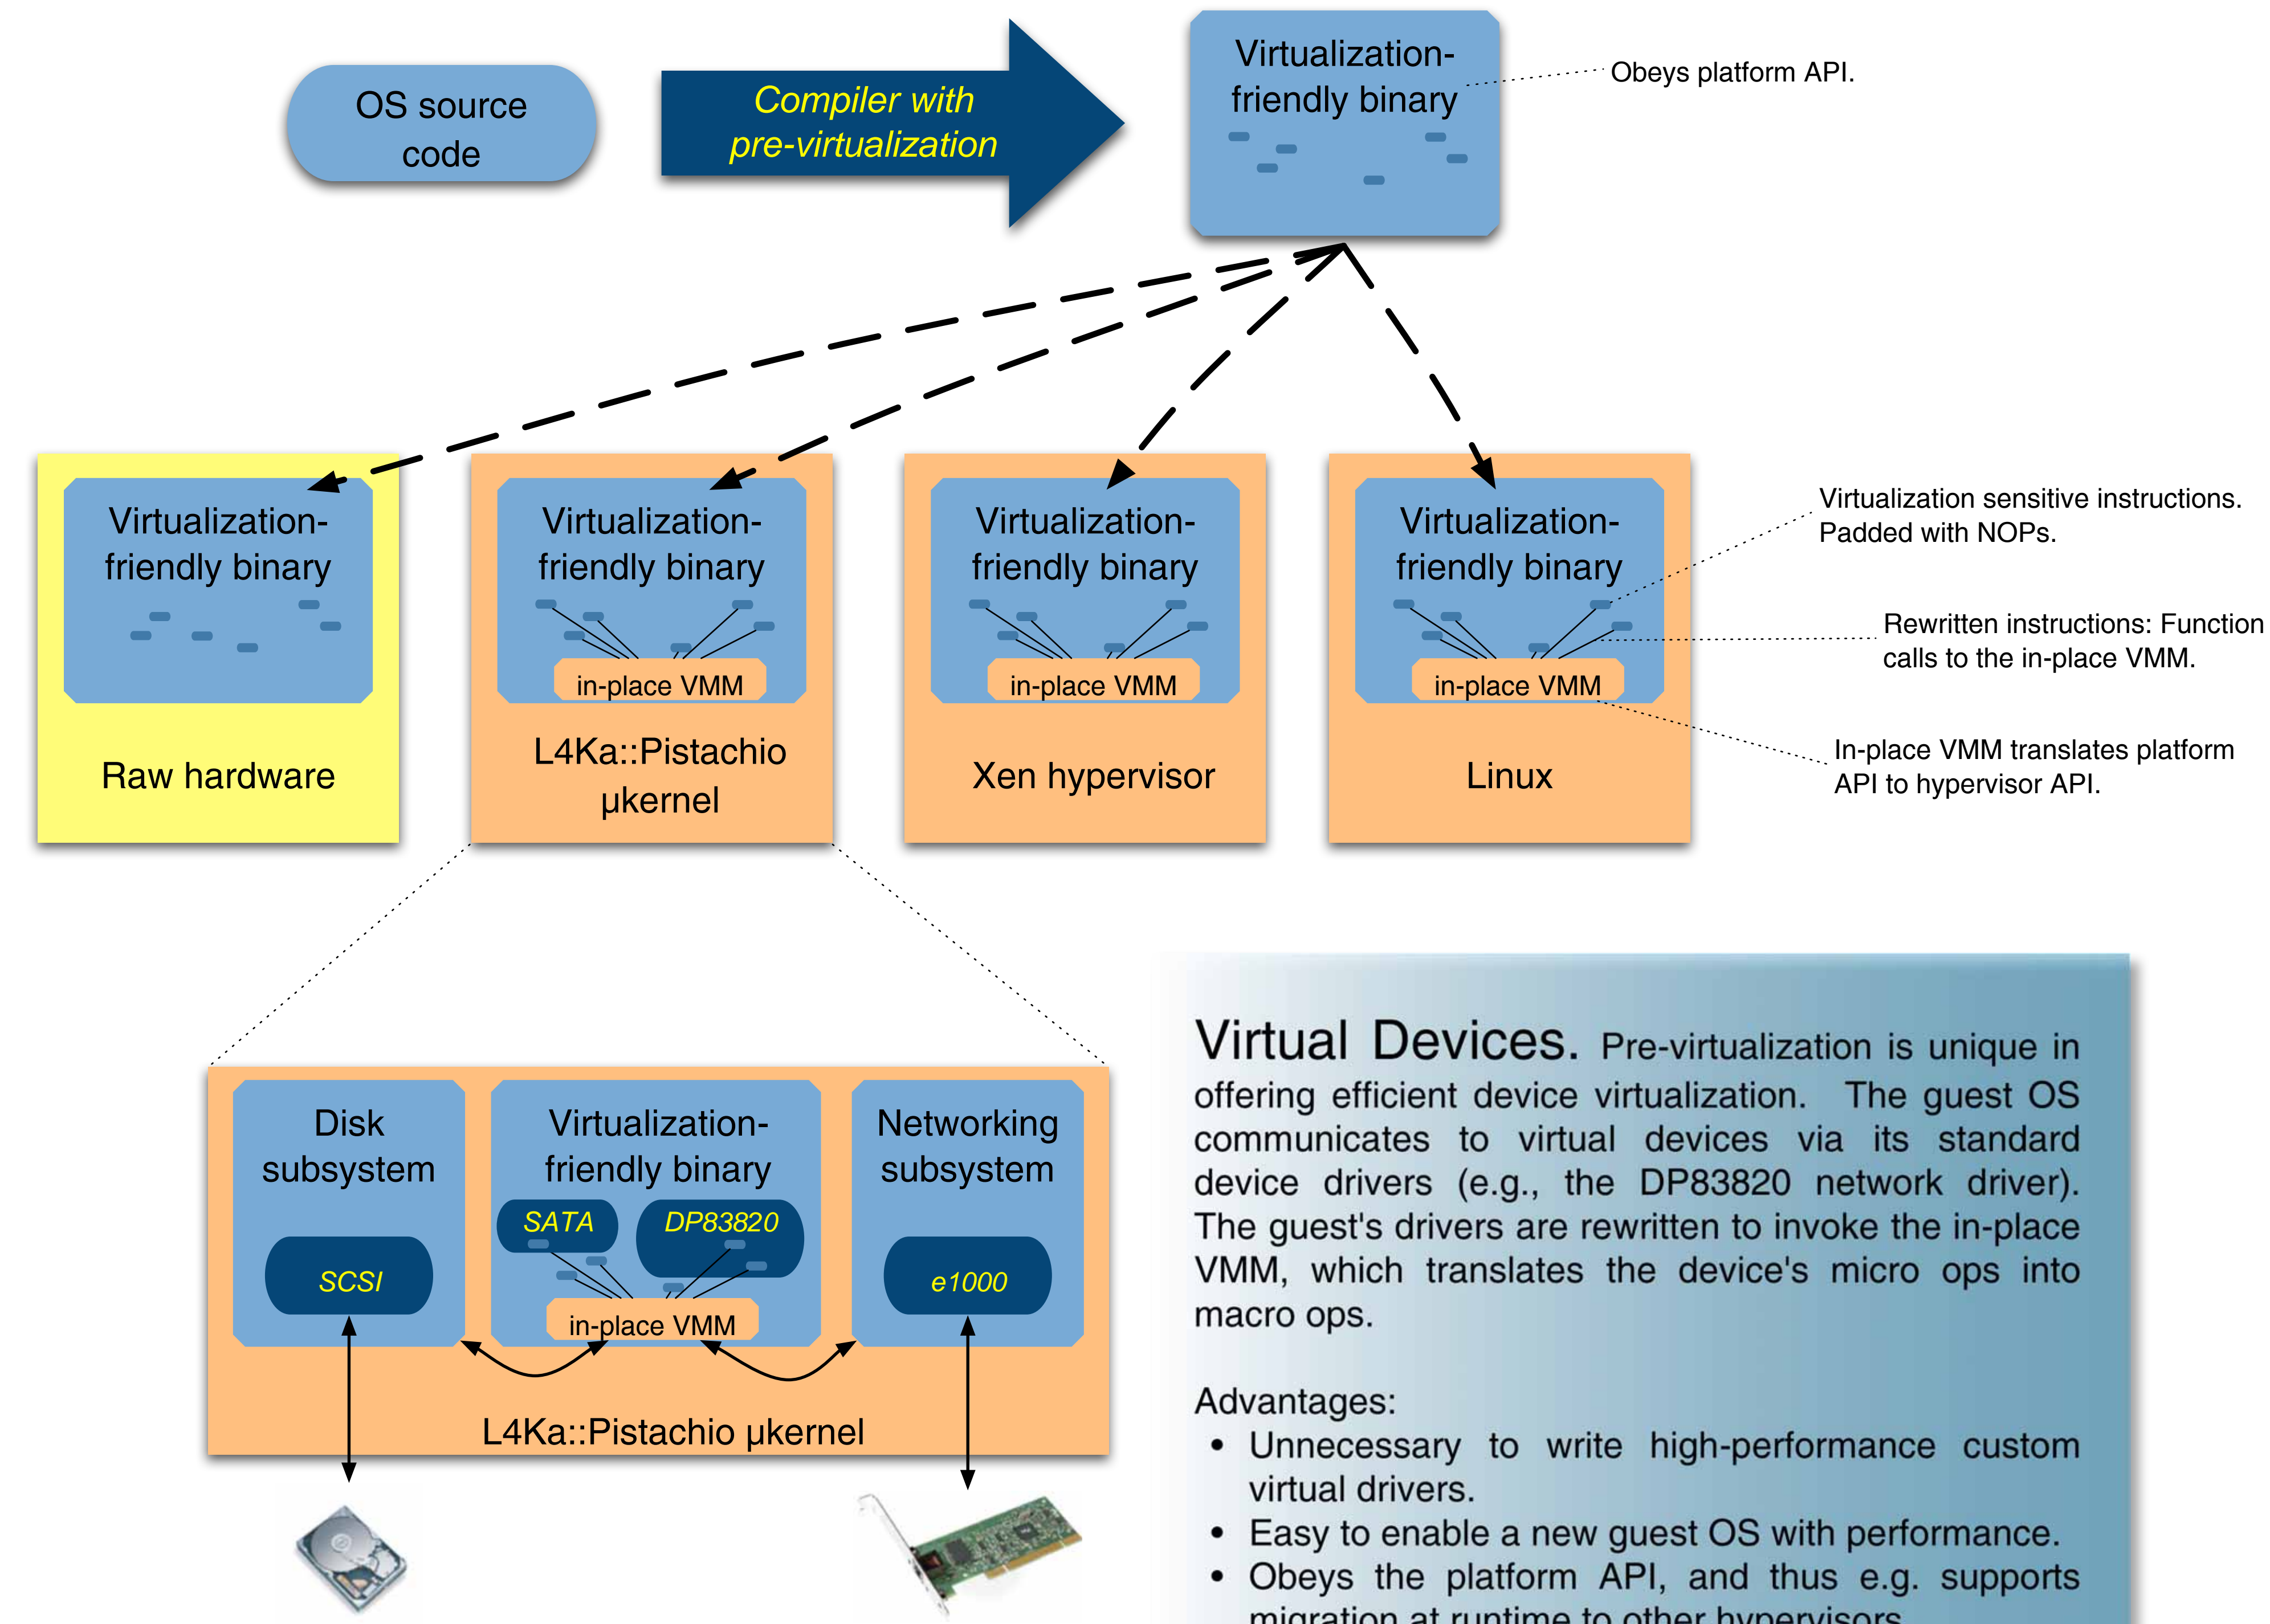

National ICT Australia
University of New South Wales

Matthew Chapman

National ICT Australia
University of New South Wales

Gernot Heiser

National ICT Australia
University of New South Wales



System	Xput [Mb/s]	CPU util	cycles/byte
native, raw	780.9	35.2%	9.64
NOPs, raw	780.2	33.5%	9.17
L4Ka::Linux	780.1	35.7%	9.77
L4Ka in-place VMM	779.8	37.3%	10.22
XenoLinux	780.7	41.3%	11.29
Xen in-place VMM	778.7	41.1%	11.28

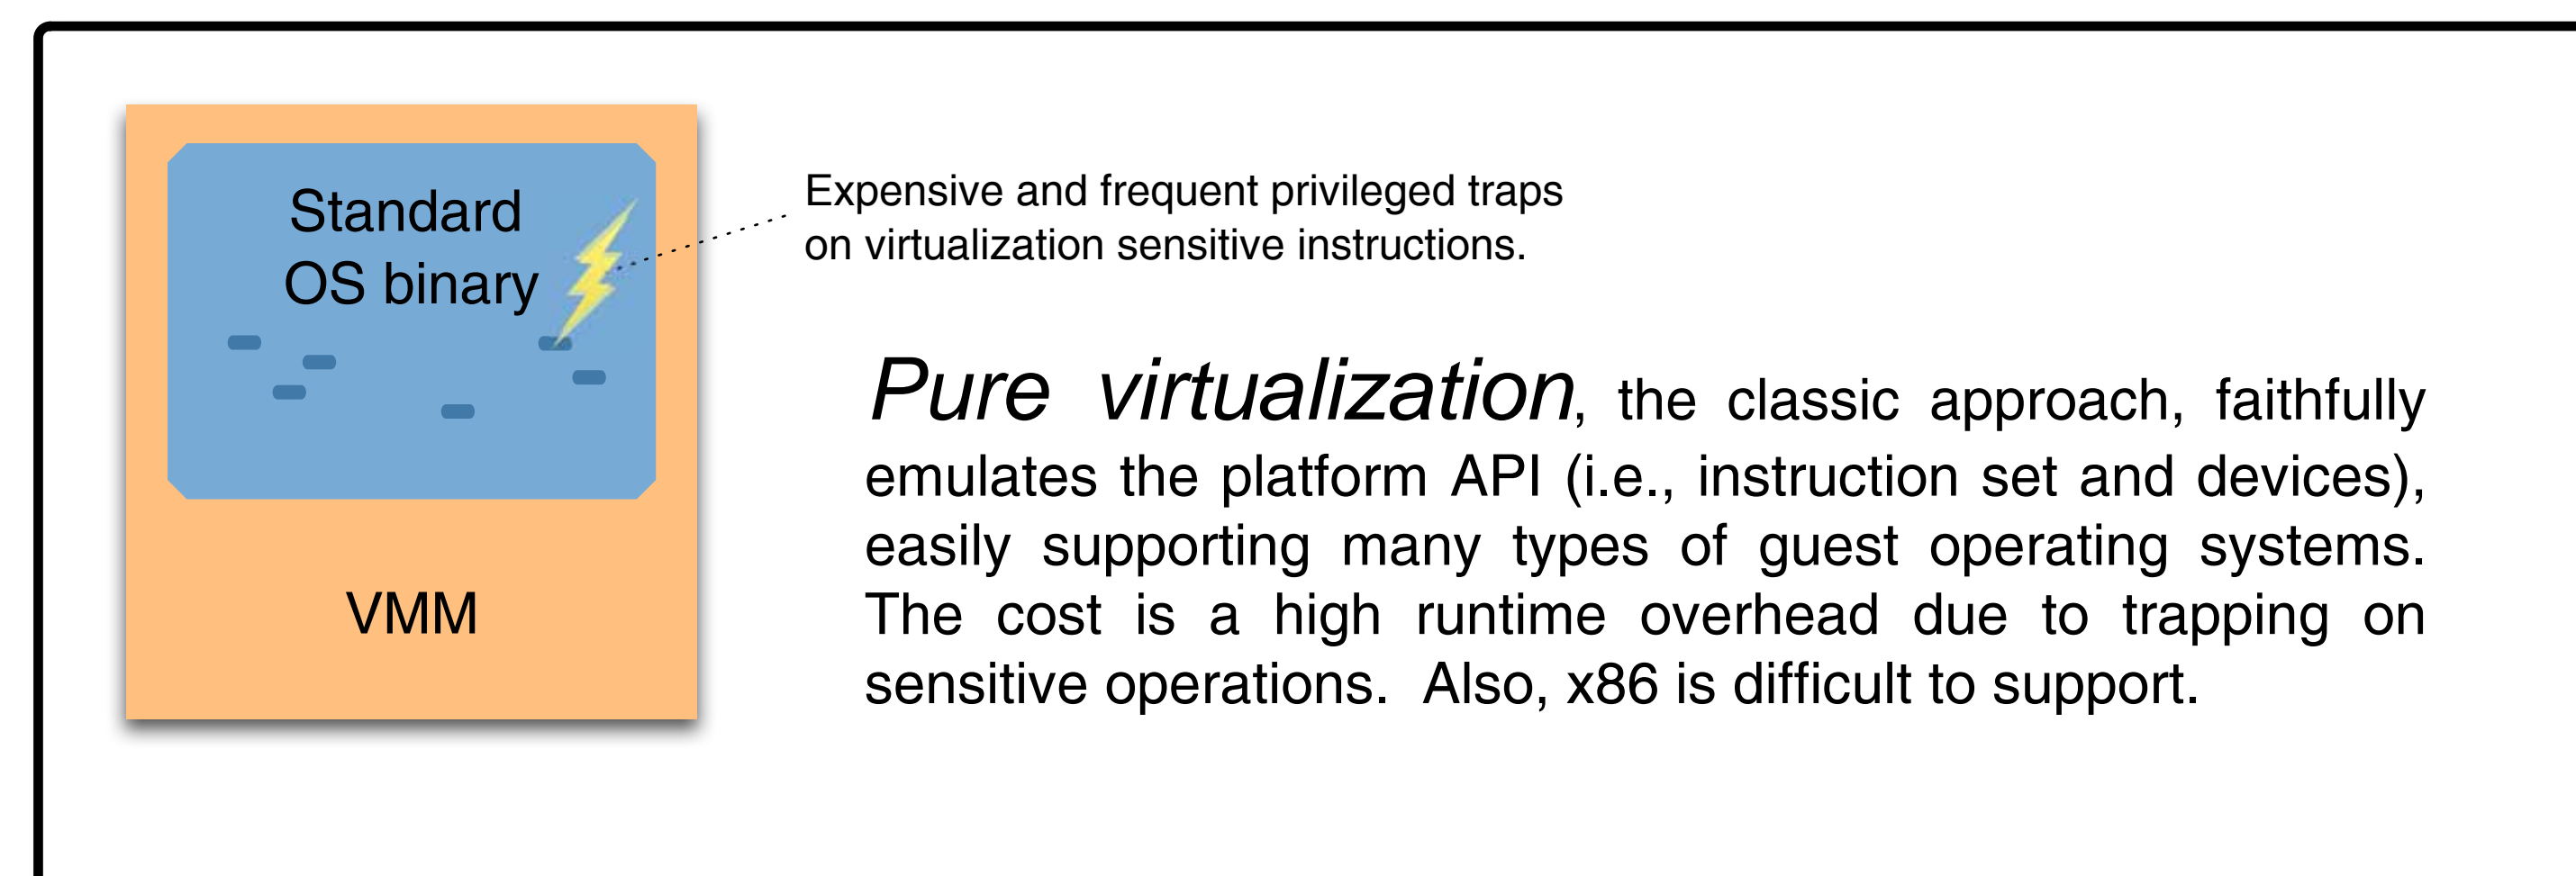
System	Xput [Mb/s]	CPU util	cycles/byte
native, raw	740.8	36.0%	10.39
NOPs, raw	740.8	36.4%	10.49
XenoLinux	739.6	43.2%	12.48

Performance. Netperf receive benchmark that transferred 1GB of data. Test machine: 2.8GHz Pentium 4, configured for 256MB, XT-PIC, direct device access, running Debian 3.1 from local SATA. Client machine: 1.4GHz Pentium 4. Gigabit Ethernet connection.

System	Xput [Mb/s]	CPU util	cycles/byte
Custom driver	707.5	60.3%	18.21
DP83820 emulation	707.1	59.8%	18.06

In device driver reuse, we used two VMs: one with indirect network access, the other with direct access. The indirect VM either used a pre-virtualized DP83820 driver, or para-virtualization with a custom virtual driver.

Conclusion. Pre-virtualization offers performance rivaling para-virtualization, and by using the platform API, pre-virtualization also offers many of the advantages of pure virtualization. Pre-virtualization's automation substantially reduces the engineering effort to build a high-performance virtual machine.



Comparison	Pre-virtualization	Pure virtualization	Para-virtualization
Mostly automated	✓	✓	-
Single guest binary, runs on all hypervisors	✓	✓	-
High performance	✓	✓	✓
Quickly enable a new hypervisor	✓	?	-
Migrate between incompatible hypervisors at runtime	✓	✓	-
Quickly port a new guest OS	✓	✓	-
High performance virtual devices	✓	-	-
Compatible with x86	✓	-	✓
Follows platform API	✓	✓	-

