


Void Linux

We heard you like little endian



Firstly, who am I?

- A programmer from the Czech Republic
- Involved in open source since 2007
- Background in graphics toolkits (EFL), programming languages, game development
- Currently the primary maintainer for POWER architecture support in Void (and other several things)

What is Void Linux?

- Independent Linux distro with its own package manager
- Focus on simplicity and pragmatism
- Rolling release, stable software, binary based
- Portable (also on x86, ARM, MIPS), choice of glibc and musl, runit, “ports” (xbps-src), cross-compilation
- Low barrier of entry, open and inclusive, informal

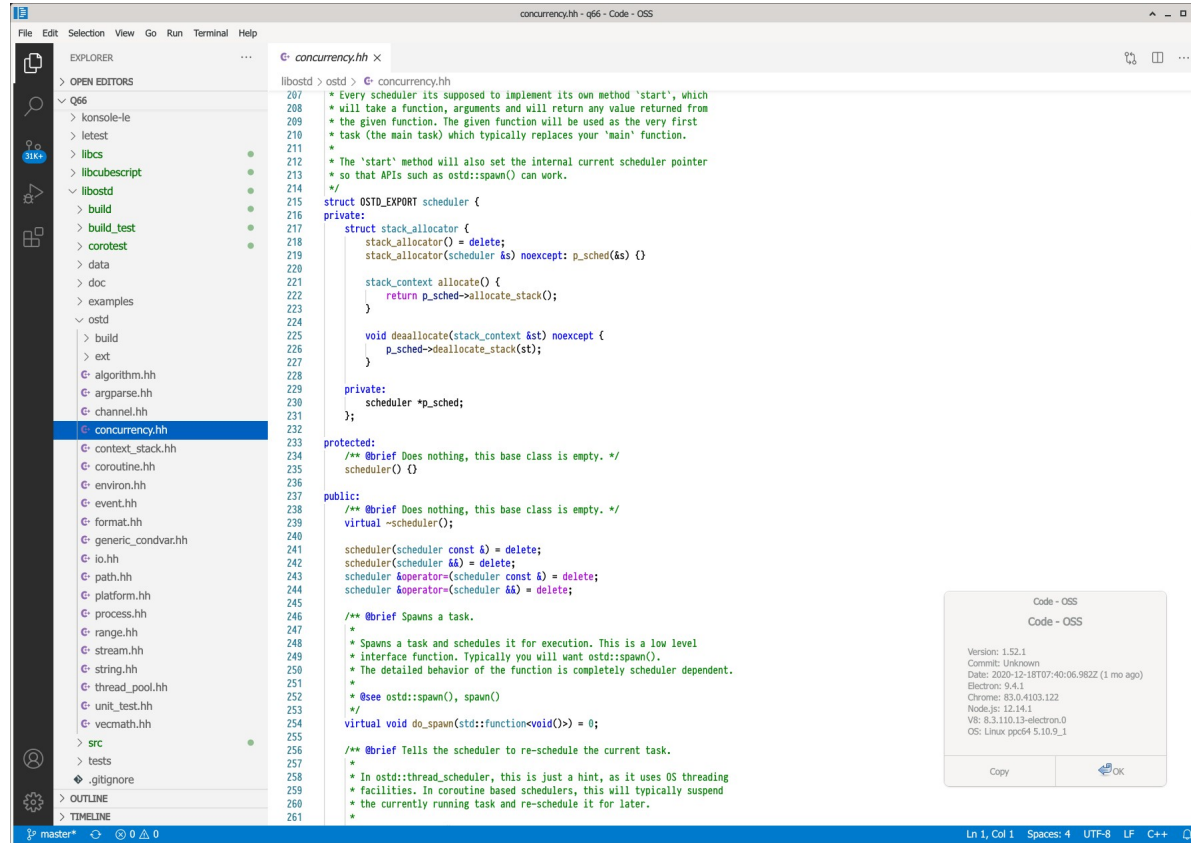
Void on POWER vs other distros

- Desktop/workstation focus, not server
- Both little and big endian support
- ELFv2 ABI on big endian glibc
- 4K pages kernel by default
- Wider desktop software support

Quick history of Void on POWER

- Late 2018 – initial port on Talos 2 Lite (ppc64le)
- Early 2019 – gradual expansion of package coverage
- 64-bit BE musl port, 32-bit ports
- April 2019 – 64-bit BE glibc port (ELFv2 ABI)
- Late 2019 – complete repository coverage
- OpenPOWER Summit EU talk
- November 2019 – commit bit in upstream

News since then



Java bootstrap

- OpenJDK 8 and 11 in repositories
- Needs itself to build (one version older)
- Re-introduction of GCC6 (last version with GCJ)
- Bootstrap packages for JDK 7, 9, 10
- JITed on ppc64(le), slow on ppc32 (zero backend)

Go bootstrap

- Only on ppc64le
- Using binary bootstrap (official binaries)
- gccgo unreliable (gcompat on musl)
- BE support – wrong ABI and needs POWER8
- 32-bit support – missing entirely

Haskell bootstrap

- Self-built binary bootstrap snapshots
- Includes musl (figured out cross-builds)
- BE missing for now (ELFv2 patch made though)
- Big endian bugs prevent functional bindist
- More target support later for 8.10

Other languages

- Greater Common Lisp support
- SBCL finally in repos (for now LE only)
- Bootstrapped with ECL
- Preparing a Clisp update
- D, Zig and some others still missing

LibreSSL performance

- Void uses LibreSSL instead of OpenSSL
- Assembly only on x86 and some 32-bit ARM
- Inferior performance and no hardware crypto support
- Solution: import assembly bits from OpenSSL upstream

LibreSSL performance

- <https://github.com/q66/libressl-portable-asm>
- Now shipping in Void
- Also adds Aarch64 assembly support and improves ARM
- Up to 20x performance increase
- Especially in hw crypto (ghash/AES etc.) on POWER9

Chromium support

- No official upstream support
- Downstream patches available
- Shipping in Void since 84
- Currently only browser engine with JIT support
- Works on musl as well

Electron support

- We have systemwide Electron now
- For now enables VS Code, Element, Rocket.Chat-Desktop
- Only LE, musl works
- Ugly build system workarounds

AMD GPU support

- Problem: AMDGPU DC
- Haphazard hardware floating point usage
- Page size dependency
- Upstream support only in recent kernels
- Need backports – provided down to 5.4

Cross-toolchain rework

- Code duplication for cross-toolchain build templates
- Solution: introduce a common build-style
- Each cross-toolchain template is now ~50 lines
- Extra goodies: glibc cross on musl hosts, etc.
- Unified configure args and so on for all
- No more dirty masterdirs

Infrastructure status

```
Terminal -
1[|||||100.0%] 19[|||||100.0%] 37[|||||100.0%] 55[|||||100.0%]
2[|||||100.0%] 20[|||||100.0%] 38[|||||100.0%] 56[|||||100.0%]
3[|||||100.0%] 21[|||||100.0%] 39[|||||100.0%] 57[|||||100.0%]
4[|||||100.0%] 22[|||||100.0%] 40[|||||98.1%] 58[|||||100.0%]
5[|||||100.0%] 23[|||||100.0%] 41[|||||100.0%] 59[|||||100.0%]
6[|||||99.4%] 24[|||||100.0%] 42[|||||100.0%] 60[|||||100.0%]
7[|||||100.0%] 25[|||||100.0%] 43[|||||100.0%] 61[|||||91.0%]
8[|||||100.0%] 26[|||||100.0%] 44[|||||100.0%] 62[|||||100.0%]
9[|||||100.0%] 27[|||||100.0%] 45[|||||100.0%] 63[|||||100.0%]
10[|||||100.0%] 28[|||||98.1%] 46[|||||100.0%] 64[|||||100.0%]
11[|||||100.0%] 29[|||||98.1%] 47[|||||100.0%] 65[|||||100.0%]
12[|||||100.0%] 30[|||||100.0%] 48[|||||99.4%] 66[|||||100.0%]
13[|||||100.0%] 31[|||||100.0%] 49[|||||100.0%] 67[|||||100.0%]
14[|||||100.0%] 32[|||||99.4%] 50[|||||100.0%] 68[|||||100.0%]
15[|||||100.0%] 33[|||||100.0%] 51[|||||100.0%] 69[|||||100.0%]
16[|||||100.0%] 34[|||||100.0%] 52[|||||100.0%] 70[|||||100.0%]
17[|||||96.8%] 35[|||||100.0%] 53[|||||100.0%] 71[|||||99.4%]
18[|||||100.0%] 36[|||||100.0%] 54[|||||100.0%] 72[|||||100.0%]
Mem[|||||58.96/126] Tasks: 369, 37 thr; 72 running
Swp[|||||0K/0K] Load average: 7.35 2.19 0.87
Uptime: 38 days, 07:15:23

PID USER PRI NI VIRT RES SHR S CPU% MEM% TIME+ Command
24639 q66 20 0 56.4G 49.2G 28128 S 0.6 38.3 1669h qemu-system-ppc64 -drive file=/media/vms/builder_ppc64_be.qcow2,if=v
24647 q66 20 0 56.4G 49.2G 28128 S 0.0 38.3 49h18:04 qemu-system-ppc64 -drive file=/media/vms/builder_ppc64_be.qcow2,if=v
24651 q66 20 0 56.4G 49.2G 28128 S 0.0 38.3 49h12:21 qemu-system-ppc64 -drive file=/media/vms/builder_ppc64_be.qcow2,if=v
24653 q66 20 0 56.4G 49.2G 28128 S 0.0 38.3 48h31:12 qemu-system-ppc64 -drive file=/media/vms/builder_ppc64_be.qcow2,if=v
24649 q66 20 0 56.4G 49.2G 28128 S 0.0 38.3 48h20:22 qemu-system-ppc64 -drive file=/media/vms/builder_ppc64_be.qcow2,if=v
24657 q66 20 0 56.4G 49.2G 28128 S 0.0 38.3 48h08:40 qemu-system-ppc64 -drive file=/media/vms/builder_ppc64_be.qcow2,if=v
24655 q66 20 0 56.4G 49.2G 28128 S 0.0 38.3 48h07:48 qemu-system-ppc64 -drive file=/media/vms/builder_ppc64_be.qcow2,if=v
24659 q66 20 0 56.4G 49.2G 28128 S 0.0 38.3 48h03:52 qemu-system-ppc64 -drive file=/media/vms/builder_ppc64_be.qcow2,if=v
24661 q66 20 0 56.4G 49.2G 28128 S 0.0 38.3 48h01:00 qemu-system-ppc64 -drive file=/media/vms/builder_ppc64_be.qcow2,if=v
24663 q66 20 0 56.4G 49.2G 28128 S 0.0 38.3 47h44:23 qemu-system-ppc64 -drive file=/media/vms/builder_ppc64_be.qcow2,if=v
F1Help F2Setup F3Search F4Filter F5Tree F6SortBy F7Nice F8Nice F9Kill F10Quit
```

Infrastructure status

- Still a Talos 2 Lite (18C, 128G RAM) in my bedroom
- Big endian builds in a VM (KVM-HV)
- Little endian builds on bare metal
- Primary mirror – separate server (in Chicago, IL)
- 10G network, 2TB storage

Infrastructure status

- Several other mirrors provided by community
- Could use faster build hardware
- Could use hardware not in my bedroom
- Primary mirror should migrate to build machine
- More build automation necessary – WIP

32-bit little endian

```
Terminal -
Untitled x Untitled x Untitled x

bash-5.1# neofetch

      _.-:;=====;._
     _-+==+==+==+==+==+;
    -+==+==+==+==+==+==+
   -+==+==+==+==+==+==+
  _vi,                --+==+==+:
 .uvnvi.              -+==+==+
.vvnvni`             .+==+==+
+QmQMpvpvvnv;_yYsyQQWUUQQm #QmQ#:QQQWUVSQQm.
-QQWQWpvpvowZ?..wQQQE<=QWQW/QWQW.QQWQ(: jQWQE
-$QQQqmmu!` jQQQ@<=QWQQ)mQQQ.mQQQC+;jWQQ@'
-$WQ8YnI:  QWQQwgQQQW`mWQQ.jQWQQgyyWW@!
-1vvvnvv.  `~+++`      ++|+++
+vnvnvv,    `~|==`      `~|==
+vnvnvnvs.      :=-`
-Invnvnvnsi.....=sv=.
+Invnvnvnvnvnvnvnvnvn;
~|Invnvnvnvnvnvnvnv}+`
~|(*1)*|~

root@rimi
-----
OS: Void Linux ppcle
Host: C1P9S01 REV 1.01
Kernel: 5.10.9_1
Uptime: 3 days, 7 hours, 15 mins
Packages: 245 (xbps-query)
Shell: sh
Resolution: 3840x2160
Terminal: xfce4-terminal
Terminal Font: Monospace 12
CPU: POWER9 (32) @ 3.800GHz
GPU: AMD ATI Radeon RX 5600 OEM/5600 XT / 5700/5700 XT
Memory: 24127MiB / 31803MiB

bash-5.1# file /bin/echo
/bin/echo: ELF 32-bit LSB pie executable, PowerPC or cisco 4500, version 1 (SYSV), dynamically linked, interpreter /lib/ld.so.1, BuildID[sha1]=a205e8f2afc95f8186c6030a18e680633ba1ca41, for GNU/Linux 3.2.0, stripped
bash-5.1# uname -m
ppcle
bash-5.1#
```

32-bit little endian

- Why not?
- Officially does not exist
- In practice, this just works though
- Pretty basic bringup process
- Made easy by xbps-src

32-bit little endian

- No native support for 32-bit LE in Linux
- Works as a chroot on 64-bit kernel though
- But not out of box: fixes are needed
- Fix 32-bit compat in kernel first
- Several issues

32-bit LE kernel compat

- Entering signal handlers clears MSR_LE
- vDSO32 disabled by default
- Certain vDSOs broken on LE – disabled those
- Syscalls with 64-bit args: wrong byte order
- All very trivial patches
- Things “run” without them

32-bit LE kernel compat

- But they crash at random points
- After patching, no more known issues
- Fixes are being upstreamed
- VDSO32 will be fixed by moving it to C
- Backports in all Void kernels (4.19+)

Bootstrapping 32-bit LE

- First: create build and cross xbps-src profiles
- Second: create cross-toolchains for 32-bit LE targets
- Minor patching in glibc and musl needed
- Adjustments around xbps-src needed

Bootstrapping 32-bit LE

- Try cross-compiling base-chroot
- Minimal set of packages for the build container
- Fix any errors along the way
- Relatively few and mostly trivial
- Binary-bootstrap a build container after that
- Then rebuild every package (no cross)

Bootstrapping 32-bit LE

- Now we can build the other base metapackages
- Afterwards, other software can be built
- Usual “fix whatever pops up” approach applies
- Most things generally just work
- Could probably function on 32-bit userland

32-bit LE: the Buts

- No official support in glibc
- Dynlinker, symvers etc. shared with 32-bit BE
- Going to propose an official port later
- Same SVR4 ABI as 32-bit BE
- Musl pretty much just works

32-bit LE: the Why

- Because we could, of course
- Easier emulation of 32-bit x86
- Port LE Linux to G3s/G4s? :)
- Or use BE Linux and swizzle syscall data?
- For now, strictly POWER8+ though

32-bit LE: the obstacles

- LLVM support
- Recently upstreamed by the FreeBSD folks
- Backported to Void
- Rust support
- Blocks a fair amount of userland right now
- WIP patches available

32-bit LE: the obstacles

- Proper glibc port needs to be made
- Use the opportunity to improve the ABI?
- Use 64-bit long doubles later (ditch ibm128)
- Choose a unique dynlinker name and fresh symvers
- Fix up our 32-bit BE ports too?

32-bit LE: the results

- Testing repos now available (glibc and musl)
- Can fully compile itself
- A portion of the repo packaged
- Core userland, dev tools, LLVM, Mesa, SDL...
- Initial port of box86

Box86

- Linux x86 emulator developed for ARM devices
- Needs 32-bit little endian host
- Initial PowerPC port now
- Can run glxgears
- Can run Unreal Tournament
- Fairly slow though... no dynarec

The future

[illegible]

Infrastructure

- Official repo at some point?
- 64LE should be easy enough
- 64BE and 32-bit have some issues with server software
- They also need clearing up the ABI situation

ELFv2 BE ABI

- We will need official support in glibc
- Right now things work “by accident”
- Need unique dynlinker name and symvers
- Upstream said they wouldn’t be against
- But, need to formalize the ABI

ELFv2 BE ABI

- IBM thinks VSX requirement is a part of the ABI
- Solution: Formalize “ELFv2 Legacy-Compatible” ABI
- To be shared by the new glibc port, musl and FreeBSD
- ELFv2 minus POWER8/VSX requirement

ELFv2 BE ABI

- Extension of the official spec
- No VSX nor VMX requirement
- 64-bit long doubles – ditch ibm128, and do not use ieee128 (needs vector registers for passing)
- Coming during 2021

Other things

- Upstream as many patches as possible
- Write a new installer
- Work on enablement of new software
- Work with upstreams on adding support

Conclusion

- Lots of work has been done in a year and half
- There is still plenty more to be done
- Join us on IRC: #voidlinux-ppc @ freenode
- For detailed statistics, visit
<https://repo.voidlinux-ppc.org/stats.html>

Questions time

And thanks for listening!



Daniel Kolesa
Twitter: [@octaforge](https://twitter.com/octaforge)
daniel@octaforge.org
<https://voidlinux-ppc.org>
FOSDEM 2021