



FreeBSD on PowerNV

POWER8 and POWER9

Kevin Bowling - Systems Software Team
In collaboration with QCM Technologies,
IBM, and Semihalf

11 Nov 2017

What is OpenPOWER?

- Commercial RISC arch launched ~1990
 - AIM, POWER, FreeScale, etc
- Speeds and Feeds
 - Short pipeline, high clock, high core, SMT
 - P9 mem BW: 120GB/s (SO) 230GB/s (SU)
- I/O
 - PCIe 4.0, OpenCAPI, NVLink
- 1.5yr tick/tock
- Anchored by HPC
 - Summit and Sierra supercomputers
 - <https://www.nextplatform.com/2017/09/19/power9-rollout-begins-summit-sierra/>
- Open ISA, platform
 - ALL firmware on GitHub
 - <https://github.com/open-power>

Cast of Characters

- Semihalf
 - Development leader
- QCM Technologies
 - QCM Technologies is based in Scottsdale AZ, and has been an IBM Premier Business Partner since 2001
 - Core competencies include IBM Power Systems and Storage with key relationships in POWER development
- IBM
 - Arch and reference material, relicensing
- LLNW
 - CDN benchmark target, sales lead
- Special thanks to nwhitehorn@ and jhibbits@

OpenPOWER S821LC

Initial porting target




OpenPOWER S821LC

Initial porting target



OpenPOWER S821LC

Initial porting target



Port Status

- Basic architecture support
 - Memory Management Unit
 - Interrupts
 - System timers
 - OPAL console
 - PCIe initial support
- RAM root filesystem embedded into the kernel

Booting up

- Store kernel and kexec files in Petitboot accessible location (TFTP server in this example)
- Enter Petitboot shell
- Copy the modified kexec binary into Petitboot Linux
 - `tftp -g -r <tftp_directory>/kexec <server_ip>`
- Copy the FreeBSD kernel into Petitboot Linux
 - `tftp -g -r <tftp_directory>/kernel <server_ip>`
- Run the FreeBSD
 - `chmod +x kexec`
 - `./kexec -l kernel`
 - `./kexec -e -f`

Elided dmesg

```
FreeBSD 12.0-CURRENT #308 b50f6f60c1d(tst)-dirty: Mon Sep
25 14:18:07 CEST 2017
wma@ppc64-prime:/home/wma/ppc64-freebsd/obj/powerpc.powerp
c64/home/wma/ppc64-freebsd/sys/GENERIC64 powerpc
...
```

```
FreeBSD/SMP: Multiprocessor System Detected: 128 CPUs
```

```
...
cpu0: IBM POWER8 revision 2.0, 2328.00 MHz
cpu0: Features
dc005180<PPC32,PPC64,ALTIVEC,FPU,MMU,SMT,ARCH205,ARCH206,V
SX>
cpu0: Features2 c2000000<ARCH207,HTM,VCRYPTO>
...
```

Current Blockers

- SMP boot does not work on boot reliably
- AHCI does not work due to bugs in hardware and lack of quirks in FreeBSD
- Intel IGB e1000 driver stalls under a heavy traffic
- Console drops some output if big chunks of text are displayed

Long term risks/needs

- Toolchain
 - Want external GCC and clang, LLD
- Big Endian vs Little Endian
 - Linux ppc switching to ppc64le
- Port drift..
 - Drivers, atomics, pmap, etc
- Architectural features
 - High core/thread count
 - Dynamic SMT
 - Crypto
 - HTM
- bhyve

Following along & getting involved

- Meet with QCM and IBM to gauge interest today
 - Sam, Michael, Dan
- freebsd-ppc@ mailing list
- Ongoing reviews in phabricator, pettiboot

If you want more info or to make introductions:
kbowling@{llnw.com,freebsd.org}