

A blue car is driving away on a road that curves into the distance. The road is flanked by green hills and a body of water. A large hand is holding a lit cigarette, with the cigarette butt extending across the road. The car's headlights illuminate the road ahead.

Linux on Power and the IBM Linux Technology Centre

Hugh Blemings, Open Source Hacker, IBM LTC
May 2008



Introductions & Agenda

» Who am I ?

» Agenda

- About the IBM Linux Technology Centre (LTC)
- Linux on Power History & Status
- LTC Development areas for Power
- Overview of Device Trees
- Overview of ePAPR
- Future Directions for Linux on Power

IBM Linux Technology Centre (LTC)

- » Main Linux development group within IBM
- » Part of Systems & Technology Group
 - Close links with both Server and Microelectronics development teams
 - Gives Linux unique level of influence into products
- » ~600+ Staff, 40+ Cities, 7+ Countries
- » Work closely with, indeed as part of, Open Source Community
 - Code accepted/rejected on technical merit
 - Work in both *Spirit* and *Letter* of Open Source Community norms
- » Team includes some very well known OSS developers

IBM Linux Technology Centre (LTC) (cont'd)

» Simple Mission Statement

- Make Linux Better
 - Test, Security, Standards, RAS
- Expand Linux's Reach
 - Telcos, Realtime, Embedded
- Bring IBM products to Linux
 - Hardware, Software, Services
- Customer Collaboration
 - Exploring variety of areas

Linux on Power – a brief history

- » One of the earlier non-x86 architecture ports
 - November 1995 in kernel 1.3.45
- » 32 Bit PowerPC first
 - Apple PowerMacs, IBM RS/6000 & Embedded boards
 - arch/ppc
- » 64 Bit PowerPC c. 2001
 - IBM pSeries/iSeries, Apple G5 & Embedded boards
 - arch/ppc64
- » Decision to consolidate 32 & 64 bit to single architecture tree c. 2005
 - arch/powerpc
 - x86 & x86-64 followed just recently...

Linux on Power – a brief history (continued)

- » 64 bit kernels handle 32 bit userspace with no overhead
 - Useful code size advantages
- » Lot of common code
 - Across Server, Embedded and CELL microprocessors
 - Across Platforms / Vendors
- » OzLabs/LTC team members involved from earliest days
- » LoP credited with help keeping broader codebase endian and word-size independent
- » Along with System Z (Mainframe) port, helped drive virtualisation support in kernel

Linux on Power – current status

- » Kernel port is actively maintained
 - Mixture of corporate and community contributors
 - Freescale, AMCC, PASEmi & IBM all involved
 - Over 25 full time contributors
 - More than 30 part time contributors
 - Remains very well regarded by broader kernel community
 - A “model” port
 - e.g. Microblaze port closely modelled on Power
 - Consolidation to arch/powerpc nearly complete
 - Strongly encourage new ports to be under arch/powerpc
- » Compiler/Toolchain actively developed

Linux on Power – current status (continued)

- » Community outside corporates remains vibrant
 - Slight decline in numbers since Apple moved to Intel
 - Increase in embedded developers
- » power.org addressing lack of commodity workstation with 4 Way 970MP based “Developer Toolkit”
- » linuxppc-dev mailing list
 - Average 400 messages/week
 - <https://ozlabs.org/mailman/listinfo/linuxppc-dev>
- » #mklinux and #ppc64 IRC channels
 - On freenode.net

LTC Development areas for Power

» Kernel

- Enabling Linux on IBM Servers, Blades
- Enabling embedded SoC
- Combination gives coverage large system to small
 - All systems benefit from enterprise class testing regime
- Arguing Power viewpoint in the community
 - Overall Linux on Power “Maintainer” is an IBMer
 - Thus works directly w/Linus Torvalds, Andrew Morton etc.

LTC Development areas for Power (continued)

» Compiler/Toolchain/Libraries

- Developing/maintaining gcc for Power
- Linker and other toolchain applications
- Some library work
 - Including some specialised libraries for SoC

» Virtualisation

- Mature virtualisation support atop server hypervisor
- Xen & kvm support

Overview of Device Trees

- » Provide a (mostly!) consistent way of representing hardware properties to OS at boot time
- » Origins in Open Firmware
- » Typically passed from firmware to bootloader as a binary “blob”
- » Support by Linux on Power from early days
- » Work by Gibson et. al. to create Device Tree Compiler and associated tools
 - Allows device tree creation from high level text representation
 - Library allows boot loader, kernel and/or firmware to get properties in consistent manner
- » <http://www.jdl.com/software>

Overview of Device Trees (continued)

» Example CPU entry

```
cpus {
    PowerPC,440GX@0 {
        device-type = "cpu";
        clock-frequency = <800000000>; // 800MHz
        timebase-frequency = <0>; // Filled in by boot code
        i-cache-line-size = <32>;
        d-cache-line-size = <32>;
        i-cache-size = <0x8000>; /* 32 kB */
        d-cache-size = <0x8000>; /* 32 kB */
        dcr-controller;
        dcr-access-method = "native";
    };
};
```

Overview of Device Trees (continued)

» Example of accessing Device Tree properties

```
/* From arch/powerpc/boot/devtree.c */

void dt_fixup_cpu_clocks(u32 cpu, u32 tb, u32 bus)
{
    void *devp = NULL;

    /* ... */

    while ((devp = find_node_by_devtype(devp, "cpu"))) {
        setprop_val(devp, "clock-frequency", cpu);
        setprop_val(devp, "timebase-frequency", tb);
        if (bus > 0)
            setprop_val(devp, "bus-frequency", bus);
    }

    /* ... */
}
```

Overview of ePAPR

- » Simplified version of the Power Architecture Platform Reference targeted at embedded applications
- » Created under the auspices of power.org with Freescale & IBM being major contributors
- » Provides formal definitions of
 - Device tree
 - Binary image formats
 - Passing of control from firmware to OS (“Application”)
- » First draft being ratified at present
 - No RTAS definitions, may include in later revisions
 - SMT/Hyperthreaded CPUs may need definition

Linux on Power – The Future

- » LoP will remain key part of IBM's Linux strategy
 - Investment and expertise will remain strong too
- » Contributions of LoP community will remain key part of broader Linux kernel ecosystem
- » For people working on Linux on Power
 - Thankyou!
 - Please work with arch/powerpc whenever you can
 - arch/ppc will go away by end of year
 - Maintainers happy to help
- » Expect influence/use of Device Trees & ePAPR to grow

Close

- » Questions
- » Thankyou :)

Power.ORG ™

Munich