

# **Porting Linux to the Digital Alpha microprocessor**

Jim Paradis  
Principal Software Engineer  
Alpha Migration Tools Group  
Digital Equipment Corporation

*paradis@sousa.amt.tay1.dec.com*

# Linux/Alpha

## What is Linux/Alpha?

- 32-bit port of Linux to Digital's Alpha CPU architecture
- Currently based on Linux 1.0.0 sources

## Why Linux/Alpha?

- Small, cheap, high-performance operating system for entry-level Alpha systems (less \$\$\$)
- Freely-redistributable reference OS for Digital's OEM (chip & board) customers
- Establish Digital's presence in "Super-PC" marketplace (Alpha, PowerPC, P6/P7, etc).
- Stimulate "impulse" purchase of Alpha Chips [ Don't worry... we'll make more 8-) ]

# Project Goals

## Linux/Alpha should be:

### Small

- Minimal runtime on 1-2 floppies, full system on 340Mb disk
- Text-only system in 8Mb memory, X-based system in 16Mb memory

### Cheap

- Linux/Alpha will be distributed under GPL (<ftp://ftp.dec.com/pub/DEC/Linux-Alpha>)
- DEC may develop added-value components for sale (e.g. high-performance X servers), but a fully functional system will be free.

### Compatible

- Source-compatible with Linux/Intel
- Investigating **binary** compatibility options with Linux/Intel

### Fast

- Should not unduly hamper CPU performance
- It's an Alpha; what can I say?

# Our port vs. Linus' port

- Complementary efforts
  - Us: Fast time-to-market, 32-bit port, minimal changes to existing code structure.
  - Linus: 64-bit port, extensive changes to existing code structure to accommodate multiple architectures
- Different emphasis means each of us solves different problems and can save the other problems later on.
- Once Linux/Alpha 1.0 is self-hosting (5/95 timeframe), it can be used for 1.2 development.
- Eventual unification of code streams.
  - We will contribute what we learned from 1.0 to Linus for the 1.2 codebase
  - Will eventually have 64-bit 1.2 kernel with both 32 and 64-bit APIs and ABIs
    - \* Compatibility with Linux/Alpha 1.0 binaries
    - \* Useful for any possible Intel binary-compatibility strategy.

# Hardware Support

- DEC 2000 AXP/DECpc 150 AXP (aka Jensen)
  - Adaptec 1742 SCSI
  - Floppy, keyboard, text-mode VGA
- DECchip Evaluation Boards
  - Floppy, keyboard, text-mode VGA
- Digital AXPpci/33 Motherboard
  - Floppy, keyboard, text-mode VGA
- Coming Soon:
  - Adaptec 1542 SCSI
  - NCR 53C810 SCSI
- **Care to contribute? 8-)**

# Alpha vs. Intel

## Similarities:

- Little-endian
- Support 32-bit integers & pointers
- Paged memory management

## Differences:

- RISC vs. CISC
- Alpha does not have byte/word memory access instructions
- OSF PAL supports 7 interrupt levels, Intel has two (cli/sti)
- Intel has segmented as well as paged memory management
- Alpha has lighter-weight procedure-call mechanism than Intel
- Alpha has more register context to save than Intel. Intel has more non-register context to save than Alpha.

# Linux/Alpha: Design Decisions

- Make it *work* first, make it *fast* later
- 32-bit system: Default integer word size is 32 bits, and all addresses fit in 32 bits.
  - Eliminates 64-bit pitfalls in existing 32-bit code
  - Uses less memory & disk than 64-bit code
  - PC-class systems do not need 64-bit addressing capability (yet!)
  - 64-bit quantities still available for computation via **long long** and **double** datatypes.
- Console and PALcode
  - SRM console on all platforms that support it.
  - Linux Mini-loader for most PC-class platforms (except DEC 2000)
  - OSF PALcode
    - \* Writing new PAL is a pain
    - \* OSF PAL generally applicable for most UNIX ports
    - \* Freely-distributable version available

# Design Decisions (cont'd)

- Memory Layout

- 0x00000000-0x3fffffff: User Code/Data/Stack

- 0x40000000-0x6fffffff: Unused

- 0x70000000-0x7bfffffff: 32-bit KSEG

- 0x7c000000-0x7fffffff: Kernel code/data/stk

- Addresses cannot exceed **0x7fffffff** due to sign-extension considerations
  - Kernel and user in same memory space to allow easy kernel access to user memory
  - **KSEG**: virtual = physical + constant
    - 64-bit: constant = 0xffffc00000000000 (provided by PALcode)
    - 32-bit: constant = 0x70000000 (We map this ourselves)
  - Current 32-bit KSEG boundaries limit us to 192Mb of physical memory. Shouldn't be a problem for the first release...
  - Kernel pages protected: KRW or KRO, UNA
  - User pages protected: KRW, URW or URO
  - Kernel can access user memory, user cannot access kernel memory.



# Technical Challenges

- **Kernel - context switching**
  - Intel gets complete switch by calling task gate
  - OSF PAL switches minimal context, OS must switch rest. Most context lives in task\_struct or on kernel stack.
- **Memory Management**
  - Alpha 3-level page table "wastes" 2 pages per process; significant with big (8K) pages and small memory (8Mb system has only 1024 pages!) We have an answer 8-)
  - Memory-management *algorithms* were portable, but implementations were not. We had to "rewrite" much of memory.c, keeping the algorithms but varying the details (*Linus redid the memory management in 1.2 to make the code more general*).
- **File System**
  - Executable loading, mostly detail work
  - Other than that, filesystem worked flawlessly from the beginning!

# Technical Challenges (cont'd)

- **Device Drivers**

- Bus access on Intel is straightforward (inb/outb, bus memory is physical memory)
- Bus access on Alpha requires glue logic in both hardware and software ("magic swizzles")
- Many PC devices come with Intel initialization code in onboard BIOS. Intel PCs, of course, will run this code before Linux boots. Alpha PCs won't.
- DMA engines operate differently on PC-class systems and on Alpha systems
- Problems need to be solved only once; other drivers "should" fall into place

# Project Status

- Developers' Kit publically available at ftp.dec.com.
- Kernel boots to shell prompt, runs numerous utilities
  - Current shell is **rc**; **tcsh** and **bash** nearly working.
  - Can mount filesystems on floppy, SCSI disk, and CD-ROM.
  - Network taking shape (**ping** gets a response!)
  - Source-level kernel debug available on DEC 2000 and on ISP simulator.
  - Utilities known to work:  
arch banner basename cal cat chattr chgrp  
chmod chown cksum cmp column comm cp diff  
dir dirname du dump echo egrep **elvis** env  
expand fgrep fmt **ftp** grep head hexdump id  
ln look ls lsattr mesg mkdir mke2fs mknod  
more mount mv od **ping** printenv pwd rc rm  
rmdir sort strings sum sync tac tail time  
touch wc whereis haltit *and more!*
  - Many utilities compile and run the first time!
  - Plan to be self-hosting by DECUS (5/95)

# Future Directions

- Self-hosting version of Linux/Alpha 1.0 by 5/95
- Unification with Linux/Alpha 1.2 code base by summer '95
- Work with XFree86 on X server
- Increased platform and hardware support (code contributions more than welcome!)
- Investigate Intel binary compatibility
- Open to suggestions!