

# Linux ARM®-Server im Unternehmenseinsatz

Daniel Gollub <gollub@b1-systems.de>  
B1 Systems GmbH

<http://www.b1-systems.de>



# Outline

- Quick and short intro about ARM®
- About SoC
- About Micro-Servers
- State of non-embedded ARM® Linux Distributions
- Usecases, Obstacles and opportunities of ARM® in Linux Enterprise Server Space

Quick and short intro about  
ARM® architecture

# Quick and short intro about ARM® architecture

- ARM® architecture is licensable
- RISC-based
- lower-cost
- lower power-consumption (less heat -> less cooling)

# Quick and short intro about ARM® architecture

- Example ARM® Instruction sets: ARM®v7, ARM®v7a, ARM®v8
- Example ARM® processors: Cortex™-A9, Cortex™-A15, Cortex™-A53
- Debian architecture name: (armel), armhf, arm64
- openSUSE/Fedora: ???, armv7hl, aarch64?

# SoC (System on Chip)

- System on Chip (SoC)
- complete system on a chip
  - (multi-core) processor
  - RAM
  - NIC
  - (GPU)
  - I/O Subsystems (SATA, PCIe, ...)
  - (Server: Mgmt Board - IPMI!)

## About Micro Servers



# About Micro Servers

- There are not only ARM®-based Micro-Servers ...
- Density
- low-power server systems
- less heat -> higher density
- best performance per watt

## Example ARM® SoC HW Specification

# Example ARM® SoC HW Specification

- Quad-Core ARM® Cortex™-A9 cores (1.1 - 1.4 GHz)
  - 32 KB L1 instruction cache per core
  - 32 KB L1 data cache per core
  - 4 MB shared L2 cache with ECC
- Management Board: IPMI 2.0 and DCMI, Serial-over-Lan
- PCIe x8, x4, x2, x1 lanes
- up to 5 SATA disks, SATA 1.5 and 3.0 GB/s supported
- Max-utilization: **5 Watts**, Idle: **0.5 Watts** ...

- Node Cards with **four** SoCs

# Example ARM® SoC based "Hyperscale Server"

- 2U enclosure
- 48 SoCs (12 Node Cards with **four** SoC each) with 4 GB RAM
- = 48 independent Bare Metal Systems
- = 192 CPU Cores
- = 192 GB RAM
- "< 300 Watts" (disks excluded?)
- (2U only fits 24 drives. 4U: 36x 3.5", 72x 2.5")

# Example ARM® SoC based "Hyperscale Server"

- 2U enclosure
- 48 SoCs (12 Node Cards with **four** SoC each) with 4 GB RAM
- = 48 independent Bare Metal Systems
- = 192 CPU Cores
- = 192 GB RAM
- "< 300 Watts" (disks excluded?)
- (2U only fits 24 drives. 4U: 36x 3.5", 72x 2.5")

## State of non-embedded ARM® Linux Distributions

- Ubuntu released ARM® (32bit) Linux Distro targeted for servers
- Fedora and openSUSE also provided ARM® 32bit distro releases
- Very high activity in preparing for aarch64
- Kernel: support for multi-platform ARM® kernel



## Aarch64 and Linux

- support got merged in Linux Kernel with 3.7 release
- glibc 2.17 supports aarch64
- gcc (cross) toolchain provided by Linaro
- OpenJDK development for Aarch64

# Usecases, Obstacles and opportunities of ARM® in Linux Enterprise Server Space

## Usecases, Obstacles and opportunities of ARM® in Linux Enterprise Server Space

# Usecases, Obstacles and opportunities of ARM® in Linux Enterprise Server Space

- How is your average utilization?
- Do you have to scale for peaks only?

# Usecases, Obstacles and opportunities of ARM® in Linux Enterprise Server Space

- How is your average utilization?
- Do you have to scale for peaks only?
- Do you know your data center power-bill?

# Usecases, Obstacles and opportunities of ARM® in Linux Enterprise Server Space

- How is your average utilization?
- Do you have to scale for peaks only?
- Do you know your data center power-bill?

# Usecases, Obstacles and opportunities of ARM® in Linux Enterprise Server Space

- Front web tier (Apache, memcache, ...)
- load balacing - handling the peaks

# Usecases, Obstacles and opportunities of ARM® in Linux Enterprise Server Space

- Front web tier (Apache, memcache, ...)
- load balacing - handling the peaks
- HPC, ...



# Usecases, Obstacles and opportunities of ARM® in Linux Enterprise Server Space

- Front web tier (Apache, memcache, ...)
- load balacing - handling the peaks
- HPC, ...
- HA: byzantine fault-tolerant (BFT):  $k$  fault tolerant, requires  $3k + 1$  systems, where  $2k + 1$  work correctly.
- Big Data? Hadoop?

# Usecases, Obstacles and opportunities of ARM® in Linux Enterprise Server Space

- Front web tier (Apache, memcache, ...)
- load balacing - handling the peaks
- HPC, ...
- HA: byzantine fault-tolerant (BFT):  $k$  fault tolerant, requires  $3k + 1$  systems, where  $2k + 1$  work correctly.
- Big Data? Hadoop?
- Storage (CePH, Gluster, ...)

# Usecases, Obstacles and opportunities of ARM® in Linux Enterprise Server Space

- Front web tier (Apache, memcache, ...)
- load balacing - handling the peaks
- HPC, ...
- HA: byzantine fault-tolerant (BFT):  $k$  fault tolerant, requires  $3k + 1$  systems, where  $2k + 1$  work correctly.
- Big Data? Hadoop?
- Storage (CePH, Gluster, ...)
- scale-out - not up.

# Usecases, Obstacles and opportunities of ARM® in Linux Enterprise Server Space

- Front web tier (Apache, memcache, ...)
- load balacing - handling the peaks
- HPC, ...
- HA: byzantine fault-tolerant (BFT):  $k$  fault tolerant, requires  $3k + 1$  systems, where  $2k + 1$  work correctly.
- Big Data? Hadoop?
- Storage (CePH, Gluster, ...)
- scale-out - not up.

# Usecases, Obstacles and opportunities of ARM® in Linux Enterprise Server Space

- Front web tier (Apache, memcache, ...)
- load balacing - handling the peaks
- HPC, ...
- HA: byzantine fault-tolerant (BFT):  $k$  fault tolerant, requires  $3k + 1$  systems, where  $2k + 1$  work correctly.
- Big Data? Hadoop?
- Storage (CePH, Gluster, ...)
- scale-out - not up.

# Usecases, Obstacles and opportunities of ARM® in Linux Enterprise Server Space

- IaaS?
  - Virtualization: yes, with Cortex™-A15 (KVM and Xen?) or option LXC
  - x86 is neither good for efficient virtualization (Popek and Goldberg)
- SaaS / PaaS
  - Java? Yes, JVM is also available for ARM® 32bit
- x86 emulation?
  - CISC is hard to emulate (Code Discovery Problem)
  - but someone is working on that
- 32bit Limits - 4 GB only?
  - LPAE part of ARM®v7-a
- Java? Yes, JVM is also available for ARM® 32bit

# Usecases, Obstacles and opportunities of ARM® in Linux Enterprise Server Space

- Obstacles:
  - Disk-less server?
  - Management of disk-less server? (Provisioning, System Management, IaaS, ...)
  - Different handling for many many tiny bare-metal system required
  - Scale-out technology comes more important for Micro-Servers

## Dealing with Multiple Architectures



# Dealing with Multiple Architectures

- Open Source software should cause not much headache
- property would have to run on the intended target platform
- build engineering tips: use Distro build-systems. E.g: Open Build Service

ARM® Servers era begins in 2013?

ARM® Servers era begins in 2013?

# ARM® Servers era begins in 2013?

- Yes.
- Job Offerings: Social Network Giant looks for ARM® Server engineers

# ARM® Servers era begins in 2013?

- Yes.
- Job Offerings: Social Network Giant looks for ARM® Server engineers
- 2013-02-25 - Baidu
  - Personal cloud storage application: Baidu Cloud
  - "lower the total cost of ownership by 25% compared to x86-based solution"
  - 96 TB storage in 2U

# ARM® Servers era begins in 2013?

- Yes.
- Job Offerings: Social Network Giant looks for ARM® Server engineers
- 2013-02-25 - Baidu
  - Personal cloud storage application: Baidu Cloud
  - "lower the total cost of ownership by 25% compared to x86-based solution"
  - 96 TB storage in 2U

# Data Center of tomorrow?

- Hybrid platform solutions
- (again) heterogen server landscape
- Large scale setups
- ... maybe in the cloud

Danke

Danke!