Modern network servers

FreeBSD
FOUNDATION

Example network service: dk-hostmaster.dk

SMALL COUNTRY. SMALL AMOUNT OF TRAFFIC.

REPRESENTATIVE CCTLD DEPLOYMENT.

Traffic as of June 2018

- Per (Danish) unicast server: 1,000 qps
- 1.3M domain names
- 23,000 DNSSEC signed domains (zone size!)
- Total zone file size: ~120MByte unsigned, ~175MByte signed
- Depending on the nameserver: ~1200Mbyte in RAM (NSD)
- CPU load approximately 0 (except when loading a new zone)
- >10Mbps continuous traffic

Data provided by dk-hostmaster.dk

Operational structure

Unicast servers

- a.nik.dk (Aalborg)
- b.nik.dk (Ballerup)
- l.nik.dk (Lyngby)
- s.nik.dk (Amsterdam)

Global anycast

- c.nik.dk (7 nodes)
- p.nik.dk (50+ nodes)

Implementation

- Off the shelf Intel servers
- Running FreeBSD
- Resource utilisation: 2–5%

Many of these stats are public:

http://stats.dk-hostmaster.dk/

Network servers

HARDWARE CONSIDERATIONS FOR NETWORK SERVERS VS. GENERIC-PURPOSE SERVERS

Resources network servers care about

- CPU: lots of trivial operations running in parallel
- Memory and bus bandwidth
- Network
- Disk I/O

Previous attempts at low-power, high-bandwidth CPUs

- Sun T1/… (SPARCv9)
- Cavium Octeon (MIPS)
- Tilera, Quanta Computer, → EZChip → Mellanox
- Intel Atom (x86)
- 32-bit ARM (ARMv7)

FreeBSD FOUNDATION

ARM64

A NEW ATTEMPT AT A SERVER-GRADE ARCHITECTURE

ARM64: a new attempt at server-grade CPUs

Main architecture highlights

- ARMv8 – a 64-bit Instruction Set Architecture
- AES instructions similar to Intel AES-NI (2GB/s AES in block mode on FreeBSD)
- Standardised on-core peripherals: GIC, MSI(-X) interrupts, timer, IOMMU
- Platform approach – attempt at creating unified configuration mechanisms: ACPI, UEFI
- Standard server peripherals: AHCI for SATA, NVMe, Intel/Chelsio Ethernet …

Commercial off-the-shelf hardware: Cavium

Cavium ThunderX CPU

- 48 cores (scales to 96)
- 16x PCIe 3 lanes
- 3x 40GbE or 12x 10GbE (`vnic(4)`) with programmable queue processing

Gigabyte server board (~$3000 USD)

- 1x Cavium ThunderX ARM processor
- 8x DDR4 DIMM slots
- 1x 40GbE QSFP+ LAN port and 4x 10GbE SFP+ LAN ports
- 4x 3.5. hot-swappable HDD/SSD bays

FreeBSD FOUNDATION

Commercial off-the-shelf hardware: AMD

AMD A1100 CPU

- AMD update to the ARM Cortex A57

SoftIron (~$600 USD)

- AMD Opteron A1100 series processor (4 cores)
- 2x RDIMM with 8GB DDR4 DRAM
- 1x 1Gbit Ethernet
- 2x USB 3.0 ports
- 2x SATA 3.0 ports
- 1x 1TByte HDD

FreeBSD FOUNDATION

Operating system support

ARM64 is a Tier-1 supported platform on FreeBSD

- Fully supported by binary updates (freebsd-update and pkg)
- Supported by security-officer and release engineering
- >20,000 third-party binary software packages work out of the box

DNS on FreeBSD on ARM64

- Full support by all common authoritative and recursive DNS implementations
- Binary packages available for BIND9, NSD, Knot, PowerDNS and Unbound
- Direct vendor support (NLNetLabs, ISC) or maintainer working closely with vendor

Performance comparison

Intel Xeon

E5-2630 v4 (10 cores + HT), 128GByte RAM, 2x240GByte SSD, 2x1TByte HDD

Cost: ~$3,100 USD

Cavium ThunderX

ThunderX, 1 socket, 128GByte RAM, 4x2TByte HDD

Cost: ~$3,000 USD

LLVM clang build on FreeBSD-HEAD (pessimal workload)

ThunderX: 32 minutes total, 74,000 minutes CPU time (20h)

Xeon: 10 minutes total, 1h CPU time

DNS servers spend very little timing compiling

LLVM build workload

- Lots of continuous disk I/O
- Gigantic memory footprint and churn
- CPU load can only be parallelised to a limited extent

Authoritative DNS workload

- Modest amount of disk I/O when loading zone files
- Entire working data set fits in memory (true for most ccTLDs and nTLDs)
- Easy to parallelise CPU load

FreeBSD FOUNDATION

Power consumption and cooling

Traditional x86 servers

- Very warm CPUs – Haswell: 135W TDP (specification), ~250W under load
- Need lots of active cooling – consumes even more power

ARM64 servers

- Much cooler CPUs – ThunderX: 120W TDP (specification), ~200W under load
- Several vendors sell much lower power CPUs
- Still need cooling – but a lot less of it!

A slight digression on security

Speculative execution vulnerabilities (Meltdown / Spectre)

- ARM64 is vulnerable to many of the same vulnerabilities as x86-64
- On FreeBSD, both vulnerabilities have been mitigated in supported releases
- Probably not a huge concern on DNS servers

FreeBSD FOUNDATION

Conclusions

ARM64 is a viable platform for network workloads

- Comparatively cheap hardware
- Great for parallel network-bound workloads
- Well-supported by FreeBSD and popular DNS software

FreeBSD FOUNDATION