Modern Network Servers

Konstantin Belousov kib@freebsd.org

December 1, 2016

git src: 2016-11-30 16:05:58 +0200 7ab766e

The FreeBSD FOUNDATION

Traffic

Small country, small traffic.

- per unicast server 1,000 qps
- 1.3M domains names
- 11,000 DNSSEC signed domains (zone size !)
- Total zone file size 190Mb as raw text
- Depending on the nameserver: 800-1000Mb in RAM
- CPU load of about 0 (except when loading a new zone)
- >10Mbps outgoing

Op structure

- Unicast servers
  1. a.nik.dk
  2. b.nik.dk
  3. c.nik.dk
  4. l.nik.dk
- Anycast clouds:
  1. 7 nodes
  2. >120 nodes

Implementation

1. Off-the-shelf Intel servers
2. Running FreeBSD
3. Resource utilization 2-5%

Resources

- CPU usage - lots of parallel trivial transactions
- Memory/Bus bandwidth
- Network
- disk io

Other attempts

Low power, slow single-thread, high bandwidth CPUs

Architectures

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

Main points

- ARMv8-A 64bit Instruction Set Architecture
- ARMv8.1 AES instructions similar to Intel AES-NI (2GB/s AES in block mode on FreeBSD)
- standartized uncore: GIC, MSI(-X) interrupts, timer, IOMMU
- platform approach - attempt of creating unified configuration mechanisms: ACPI, UEFI
- standard server peripherals: AHCI for SATA, NVMe, Intel/Chelsio ethernet ...

Software installation

Single OS image for the platform

Cavium ThunderX CPU

top config: 48 cores (scales to 96), 16x PCIe 3 lanes, 3x40Gbe or 12x10Gbe (vnic(4)) with programmable queue processing

Gigabyte

http://b2b.gigabyte.com/products/product-page.aspx?pid=5864

```
1 x Cavium. ThunderX. ARM processor
8 x DDR4 DIMM slots
1 x 40GbE QSFP+ LAN port
4 x 10GbE SFP+ LAN ports
4 x 3.5. hot-swappable HDD/SSD bays
400W 80 PLUS Gold single PSU
```

AMD A1100 CPU

AMD update to the ARM Cortex A57

SoftIron

https://shop.softiron.co.uk/product/overdrive-1000/

AMD Opteron A1100 series processor
4 x 64-bit ARM Cortex A57 Cores
2 x RDIMM with 8GB DDR4 DRAM
1 x 1GBase-T Ethernet
2 x USB 3.0 ports
2 x SATA 3.0 ports
1 x 1TB HDD

$600

Intel Xeon

E5-2630 v4 (10 cores + HT), 128gb ram, 2x240gb ssd, 2x1TB hdd
$3,100

Cavium ThunderX CPU

ThunderX, 1 socket, 128gb ram, 4x2Tb hdd
$3,000

Clang build

ThunderX: 32 minutes total, 74,000 minutes CPU time (20h)
Intel: 10 minutes total, 1h CPU time

Tier-1 ARM64 platform

1. binary updates of the base system
2. support from security and release engineering teams
3. ports: pre-built packages 20,000 out of 26,500

DNS software

All important DNS software is there: Knot, NSD, BIND, Powerdns, Unbound, OpenDNSSEC.
Supported either directly by vendor (ISC, NLNetLabs), or have a FreeBSD port maintainer working closely with the vendor.

Summary

- Watch out ARM64 hardware.
- Watch out Intel as well: D1500 SoCs.