

ISC DHCP Update

Measuring Performance and BIND10 Integration

ENOG3, Odessa

23 May 2012

Tomek Mrugalski <tomasz@isc.org>, ISC



DHCP Performance Problem Space

- Vendors often provide performance results, why measure it again?
 - Marketing data is always trustworthy, right?
 - Your HW may differ from reference HW(CPU, disk, fs, OS,...)
 - Your traffic model may differ
- The most reliable way is measure it on your own.
- No suitable solutions on market
 - Outdated, non-portable, expensive and/or closed

Perfdhcp :: Status

- Open source (BSD), currently Linux, but BSD and Solaris planned
- DHCPv4 & DHCPv6 (2-way & 4-way exchanges)
- Support for packet template files (optional)
- Server/interface selection (multicast/unicast)

- Parametrized traffic and test

- # of clients,
- # of transactions/sec,
- best effort test,
- test duration,
- number of requests,
- max number/percentage of drops ...

perfdhcp

```
[-hv] [-4|-6] [-r<rate>] [-t<report>] [-R<range>] [-b<base>]
[-n<num-request>] [-p<test-period>] [-d<drop-time>] [-D<max-drop>]
[-l<local-addr|interface>] [-P<preload>] [-a<aggressivity>]
[-L<local-port>] [-s<seed>] [-i] [-B] [-c] [-1]
[-T<template-file>] [-X<xid-offset>] [-O<random-offset>]
[-E<time-offset>] [-S<srvid-offset>] [-l<ip-offset>]
[-x<diagnostic-selector>] [-w<wrapped>] [server]
```

- Diagnostics selector

- Measurements:

```
sent: 2092678/51306, received: 51306/42993 (drops: 2041372/8313)
tooshort: 0, orphans: 0, local limits: 0
rate: 398.620005
RTT0: min/avg/max/stddev: 0.718/88.420/93804.865/946.182 ms
RTT2: min/avg/max/stddev: 0.614/86.562/11609.589/395.819 ms
```

- Plans: develop DHCPv6 capabilities first (2012):

- Prefix Delegation, Relays, DOCSIS3.0, customization, validation

Why DHCP rewrite?

- Existing code is 12+ years old
- Hardware changed (many cores)
- Networks changed
- DHCP landscape changed
- New software development techniques
- Performance
- Monolithic
- Documentation is lacking

BIND10 DHCP

Codename Kea

- Common infrastructure with BIND10 DNS
 - On-line configuration
 - Logging
 - Statistics
- Performance is essential
- IPv6 is a first class citizen, not add-on
- C++ as a language of choice
- Multi-core support
- Switchable backends (flat file, SQLite, MySQL, ...)
- Hooks
 - Modular
 - Resiliency (fault isolation and recovery)



Work to Date

DHCPv4 server

(b10-dhcp4)

- Supports DORA
- One, fixed lease
- Netmask, router, DNS
- Relayed traffic only

DHCPv6 server

(b10-dhcp6)

- Supports SARR
- One, fixed lease
- DNS
- Direct traffic only

libdhcp++

- packet, options parsing/assembly
- socket/interface detection (Linux)

- Documentation: Man pages, Admin Guide, Developer's Guide
- Designs: Hooks; Lease/database; Option Definition Design

<http://bind10.isc.org/docs/bind10-guide.html>

<http://bind10.isc.org/wiki/Kea>

Kea Roadmap

2012

- Abstract Pool/Lease Storage Interface
- SQL-based back-end (SQLite vs MySQL vs flat-file)
- Option definition framework
- Relay support (?)

2013 and beyond

- Hooks
- Prefix Delegation
- DDNS
- Failover

Interested?

Fully open source model

- Source code (GIT repo)
- Bug/Suggestion/Proposals (Trac tickets available)
- Test reports
- Roadmap (agile development, sprints)

Contribute

- We are looking for sponsors (money and developers)
- Development contracts
- Review design documents (e.g. requirements)
- Submit patches

Unfunded ideas

- DHCPv6 failover
- Multi-master database
- Different backends (PostgreSQL? Cassandra?)
- CPE market



Questions?



Thank you

isc.org