

CERN central DHCP service: Migration from ISC DHCP to Kea

Maria Hrabosova

CERN IT Department

March 16, 2021

Outline

- 1 Introduction
 - Motivation
 - Kea DHCP
- 2 Configuration translation
 - Current setup
 - ISC DHCP configuration
 - Kea configuration
 - Problems
 - Kea features
- 3 Testing
 - Validation of Kea configuration
 - Test environment
- 4 Configuration generation
 - DHCP Updater
- 5 Migration status and plans

Section 1

Introduction

Motivation

- Foreseen end of life of ISC DHCP (dhcpd)
 - Date not announced yet
- No new features, only security fixes (few customers left)

Kea DHCP

- Developed by ISC
 - <https://www.isc.org/>
- Used for example by [Facebook](#)
- Still under development
- New features
- Some features from ISC DHCP dropped on purpose
- Can be customized using “hooks” libraries
 - Free: High Availability, MySQL Configuration Backend, BOOTP, ...
 - Premium: Forensic Logging, Flexible Identifier, additional management API commands, ...

Section 2

Configuration translation

Current setup

Current setup

- ISC DHCP
- 3 separate network domains
 - 2 DHCP servers in each
- DHCPv4, DHCPv6, DDNS
- ~7000 subnets
- ~400 000 host reservations
- Configuration re-generated every 5 minutes
 - Client information and topology extracted mainly from the CERN central network database

ISC DHCP configuration

ISC DHCP configuration

- Topology
 - Shared networks
 - Subnets
 - IP address pools
- Client classes
 - E.g., list of MAC addresses blocked by the security team
 - Allow/deny statements in pool configuration
- Client groups
 - DHCP options (next-server, boot-file-name, ...)
 - Host reservations (hostname, IP address, options)
- Other
 - Vendor spaces
 - Lease times
 - BOOTP
 - ...

Kea configuration

Client classes in Kea (1)

- Configure 3 different elements from ISC DHCP configuration
 - Client groups
 - Client classes
 - Allow/deny statements in IP address pools

Kea configuration

Client classes in Kea (2)

```
"client-classes": [
{
    "name": "LINUX",
    "test": "substring(option[vendor-class-identifier].text, 0, 9) == 'PXEClient'",
    "option-data": [
        {
            "name": "vendor-class-identifier",
            "csv-format": true,
            "data": "PXEClient"
        }
    ],
    "next-server": "10.0.0.99"
}
]
```

Kea configuration

Client classes in Kea (3)

```
"client-classes": [
    {
        "name": "blocked-list",
        "test": "pkt4.mac == 0xaabbccddeeff or pkt4.mac == 0x112233445566"
    },
    {
        "name": "pool#!blocked-list",
        "test": "not member('blocked-list')"
    }
],
"subnet4": [
    {
        "subnet": "10.0.0.0/24",
        "pools": [
            {
                "pool": "10.0.0.100 - 10.0.0.200",
                "client-class": "pool#!blocked-list"
            }
        ]
    }
]
```



Kea configuration

Host reservations in Kea

- Not in client classes
- Client class specified in the reservation
- Host reservations with a fixed IP address
 - Inside subnet configuration
- Host reservations without a fixed IP address
 - Global reservations

Kea configuration

Host reservations in ISC DHCP

```
group { # LINUX Clients
    host HOST1-AA-BB-CC-DD-EE-FF {
        hardware ethernet AA:BB:CC:DD:EE:FF;
        option host-name "host1";
        fixed-address 10.0.0.99;
    }
    host HOST2-11-22-33-44-55-66 {
        hardware ethernet 11:22:33:44:55:66;
        option host-name "host2";
    }
}
```

Kea configuration

Subnet reservations in Kea

```
"subnet4": [
    {
        "subnet": "10.0.0.0/24",
        "reservations-global": true,
        "reservations-in-subnet": true,
        "reservations": [
            {
                "hw-address": "aa:bb:cc:dd:ee:ff",
                "ip-address": "10.0.0.99",
                "hostname": "host1",
                "client-classes": [
                    "LINUX"
                ]
            }
        ]
    }
]
```



Kea configuration

Global reservations in Kea

```
"Dhcp4": [
{
    "reservations": [
        {
            "hw-address": "11:22:33:44:55:66",
            "hostname": "host2",
            "client-classes": [
                "LINUX"
            ]
        }
    ]
}
```



Problems

- Cannot redefine standard DHCP options (e.g., *ntp-servers*)
 - Workaround: set the value in hexadecimal regardless of the type of the option
- String type was not supported in vendor-encapsulated-options
 - Added in 1.3.0
- BOOTP support added in Kea 1.7.2
- Multiple host reservations for the same IP address were not allowed in the same subnet
 - Possibility to disable the uniqueness test since 1.9.1
- Client hardware addresses needed in class configuration
 - It was not enough to specify the class in the host reservations
 - Fixed in 1.9.5
- Multiple IP addresses not allowed in DHCPv4 host reservations
- Authoritative clause not supported in DHCPv6

Kea features

- High Availability
 - Multiple cooperating server instances
 - Load balancing (each server responds to 1/2 of the requests)
 - If any of these instances becomes unavailable a surviving server instance can continue providing the reliable service to the clients
- Management API
 - Online reconfiguration without requiring server shutdown
 - Modify configuration, enable/disable DHCP service, get service status, ...
- Monitoring with Stork
 - See up-to-date details regarding pool utilization, monitor the state of the server instances, ...
 - Can be integrated with Prometheus and Grafana

Section 3

Testing

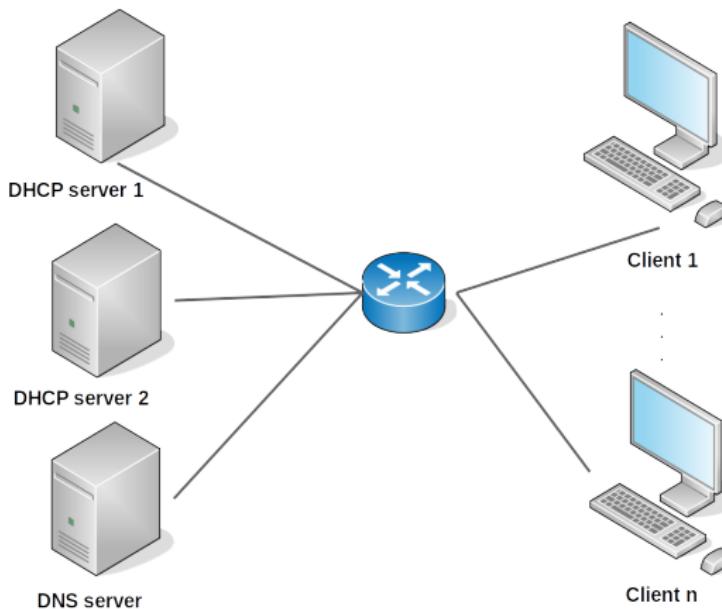
Validation of Kea configuration

Validation of Kea configuration

- Goal: make sure that the new Kea configuration behaves the same way as the previous ISC DHCP configuration
- Generate a sample ISC DHCP configuration that covers all needed use cases
- Generate a matching Kea configuration
- Send identical queries to both servers and compare the replies
 - ~17000 DHCPv4 requests

Test environment

- Virtual network in Podman



Section 4

Configuration generation

DHCP Updater

- Golang
 - Used Perl for updating ISC DHCP configuration
- Generate Kea configuration using data from the CERN central network database
- Update the DHCP configuration on the servers every 5 minutes

Section 5

Migration status and plans

Migration status and plans

- Translated and tested DHCPv4 configuration
 - Except for dhcp-server-identifier and BOOTP
- Translated DHCPv6 and DDNS configuration, testing in progress
- TODO
 - Implement regular configuration updates
 - Setup monitoring and alarms
 - Solve last problems
 - Pilot in Q2

Introduction
ooo

Configuration translation
oooooooooooo

Testing
ooo

Configuration generation
oo

Migration status and plans
ooo●

Questions?



maria.hrabosova@cern.ch

CERN central DHCP service: Migration from ISC DHCP to Kea

25 / 24