Kea 2.0
A modern DHCP

Tomek Mrugalski,
Director of DHCP Engineering, ISC
2022-May-6, DKNOG'12, Copenhagen

Tomek Mrugalski

- MSc (2003), PhD (2010), both about DHCPv6
- Started Dibbler in 2003 (complete DHCPv6 solution)
- 7 years at Intel
- IETF (since 2009)
  - DHC WG co-chair at IETF (till 2020)
  - 13 RFCs published
  - DHCPv6bis (RFC8415) as primary author
- ISC (since 2011)
  - First engineer working on Kea
  - Currently Director of DHCP engineering

Tomek Mrugalski

ISC DHCP Legacy

- Provided in many major operating systems
- Development started in 1995
- widely used, but not aging well
- ISC DHCP “development” is in maintenance mode only
- Kea is a replacement for the ISC DHCP server
- 4.4.3 released in Mar 2022. **Last release for client and relay.**
- Upcoming 4.5.0 will be server only.
- If you are running this in your network today - consider it technical debt

Time to Migrate to Kea

- Run the migration assistant
- Fix up 20 - 30% this doesn't cover
- Migrate leases if desired

  https://www.isc.org/presentations/
- NANOG'76 talk

  https://pc.nanog.org/static/published/meetings//NANOG76/daily/day_2.html#talk_1998

Kea Differences from ISC DHCP

- Extensive **REST Management API**
- Separate **‘backends’** leveraging popular open source DBs
  - Leases
  - Reservations
  - Server configurations
- Extensible with optional hooks libraries, including many from ISC
- Open source (MPL2), **with commercial add-ons**
- Available as source, or as **ISC packages** for popular OSes
- Both **stable and development** branches available

Kea's REST API

- Allows on-line reconfiguration of DHCPv4, DHCPv6 and DDNS servers without restarting
- Kea configuration AND the REST api, use JSON syntax (comments allowed)
- API commands are fully documented in

  https://kea.readthedocs.io/en/kea-2.1.4/api.html

  184 commands available and growing

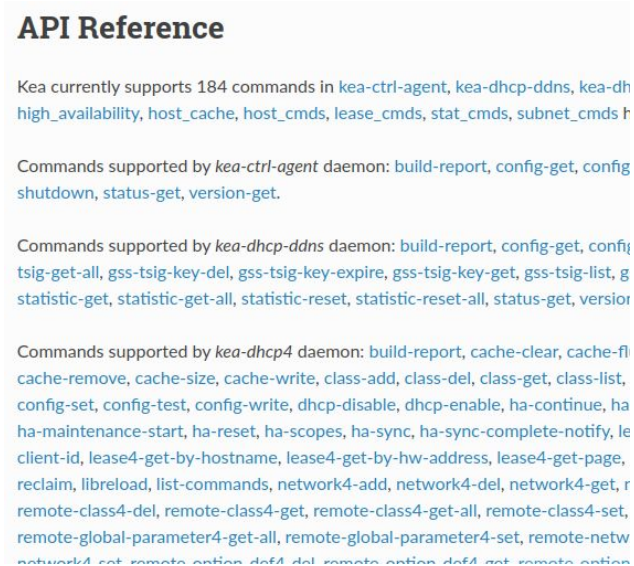

API Reference

Kea currently supports 184 commands in kea-ctrl-agent, kea-dhcp-ddns, kea-dh
high_availability, host_cache, host_cmds, lease_cmds, stat_cmds, subnet_cmds h

Commands supported by *kea-ctrl-agent* daemon: build-report, config-get, config
shutdown, status-get, version-get.

Commands supported by *kea-dhcp-ddns* daemon: build-report, config-get, confi
tsig-get-all, gss-tsig-key-del, gss-tsig-key-expire, gss-tsig-key-get, gss-tsig-list, g
statistic-get, statistic-get-all, statistic-reset, statistic-reset-all, status-get, versio

Commands supported by *kea-dhcp4* daemon: build-report, cache-clear, cache-fl
cache-remove, cache-size, cache-write, class-add, class-del, class-get, class-list,
config-set, config-test, config-write, dhcp-disable, dhcp-enable, ha-continue, ha
ha-maintenance-start, ha-reset, ha-scopes, ha-sync, ha-sync-complete-notify, le
client-id, lease4-get-by-hostname, lease4-get-by-hw-address, lease4-get-page,
reclaim, libreload, list-commands, network4-add, network4-del, network4-get, r
remote-class4-del, remote-class4-get, remote-class4-get-all, remote-class4-set,
remote-global-parameter4-get-all, remote-global-parameter4-set, remote-netw

API Basics

1. Send **`list-commands`** command:

```
# kea-shell --host ::1  --port 8080 --service dhcp6 list-commands
^D
```

API Basics

2. Get list of currently supported commands in return:

```
{
    “command”: “list-commands”,
    “service”: [ “dhcp6” ]
}
```

```
{
  "arguments": [
    "build-report",
    "config-get",
    "config-set",
    "config-test",
    "remote-global-parameter4-del",
    "remote-global-parameter4-get",
    "remote-global-parameter4-get-all”,
    . . .
    “remote-subnet6-list”,
    "server-tag-get",
    "shutdown",
    “statistic-{get,remove,reset}“,
    “statistic-{get,remove,reset}-all”,
    "version-get"
  ],
  "result": 0
}
```

Why use database ‘backends’?

- SQL data can be modified any time
- All changes applied instantly (no restart)
- Adapt your provisioning systems to write directly to the database or
- Use the API (some of these require premium hooks libraries)
- More complicated deployment, more things to install and manage (the db)
- CSV, MySQL, PostgreSQL
    - Cassandra deprecated, being removed in upcoming 2.2

Available Backends

DHCPv4, DHCPv6 server

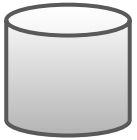

MySQL
Postgres (2.1)

Changing

Often

Rarely

- Leases (addresses, prefixes)
- Host reservations (per host details)
- Options
- Pools
- Subnets
- Shared networks
- Option definitions
- Global parameters

Lease backend

Hosts backend

Configuration backend

Recent changes in Kea 2.0

1. Significant performance boost with multi-threading
2. Addition of TLS security for connections
    - Kea - db backends
    - Kea - stork
    - Kea - api clients
3. New features
    - Cache threshold
    - Script hook
4. Stork graphical dashboard

Multi-threading (Kea 1.8)

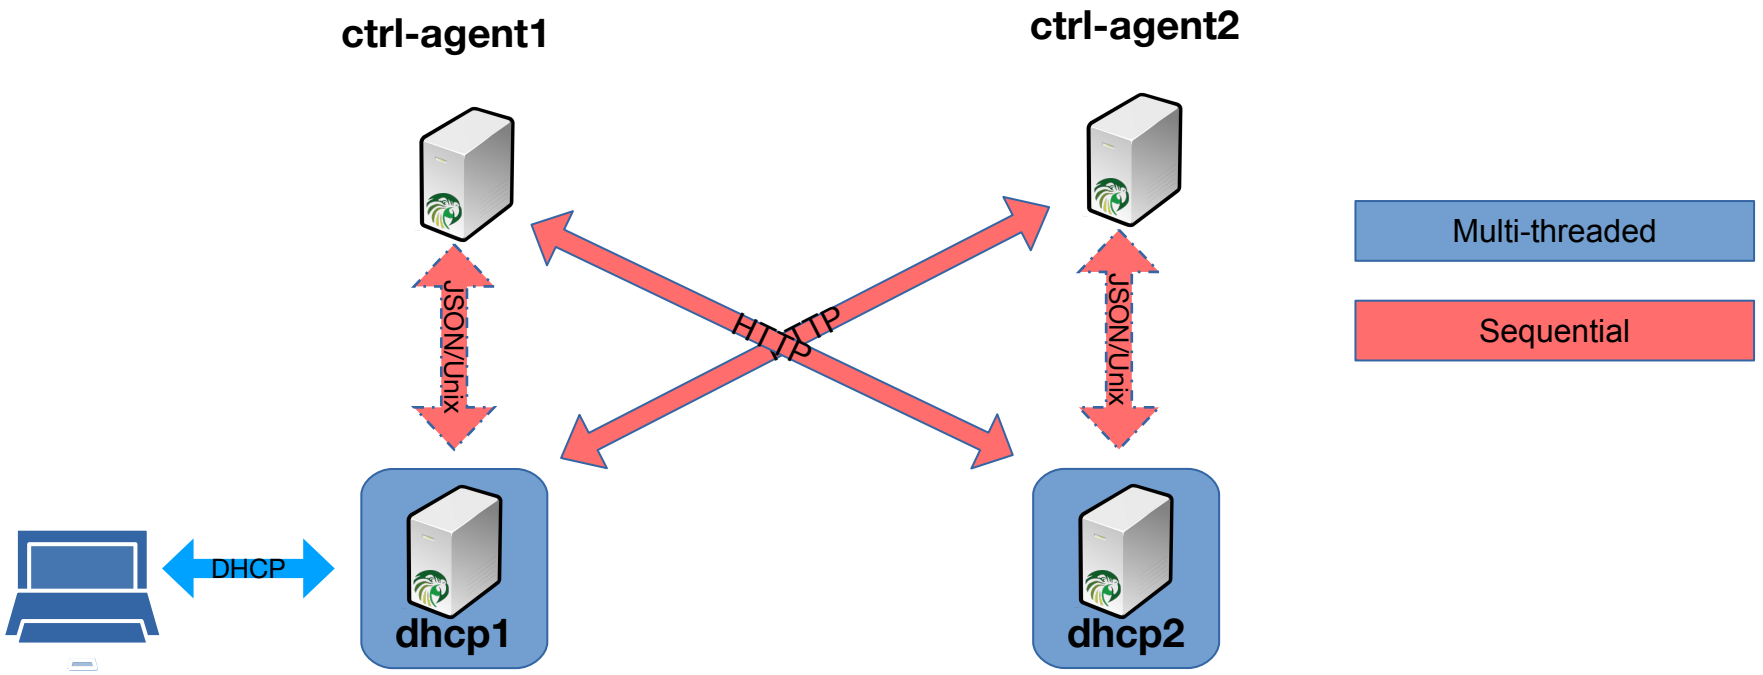

High Availability with Multi-threading (Kea 2.0)

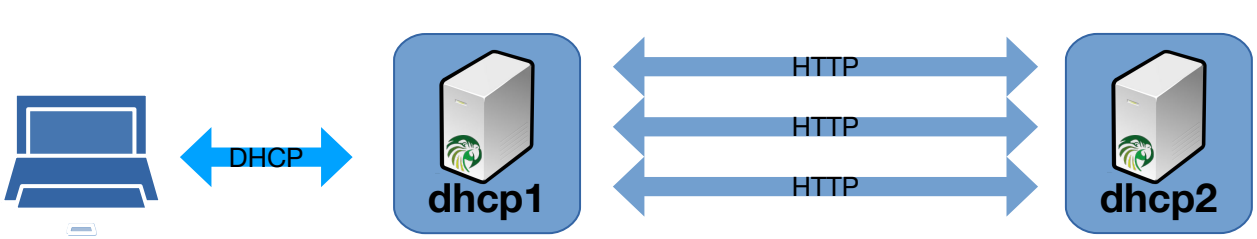

Multi-threading performance boost (Kea 2.0)

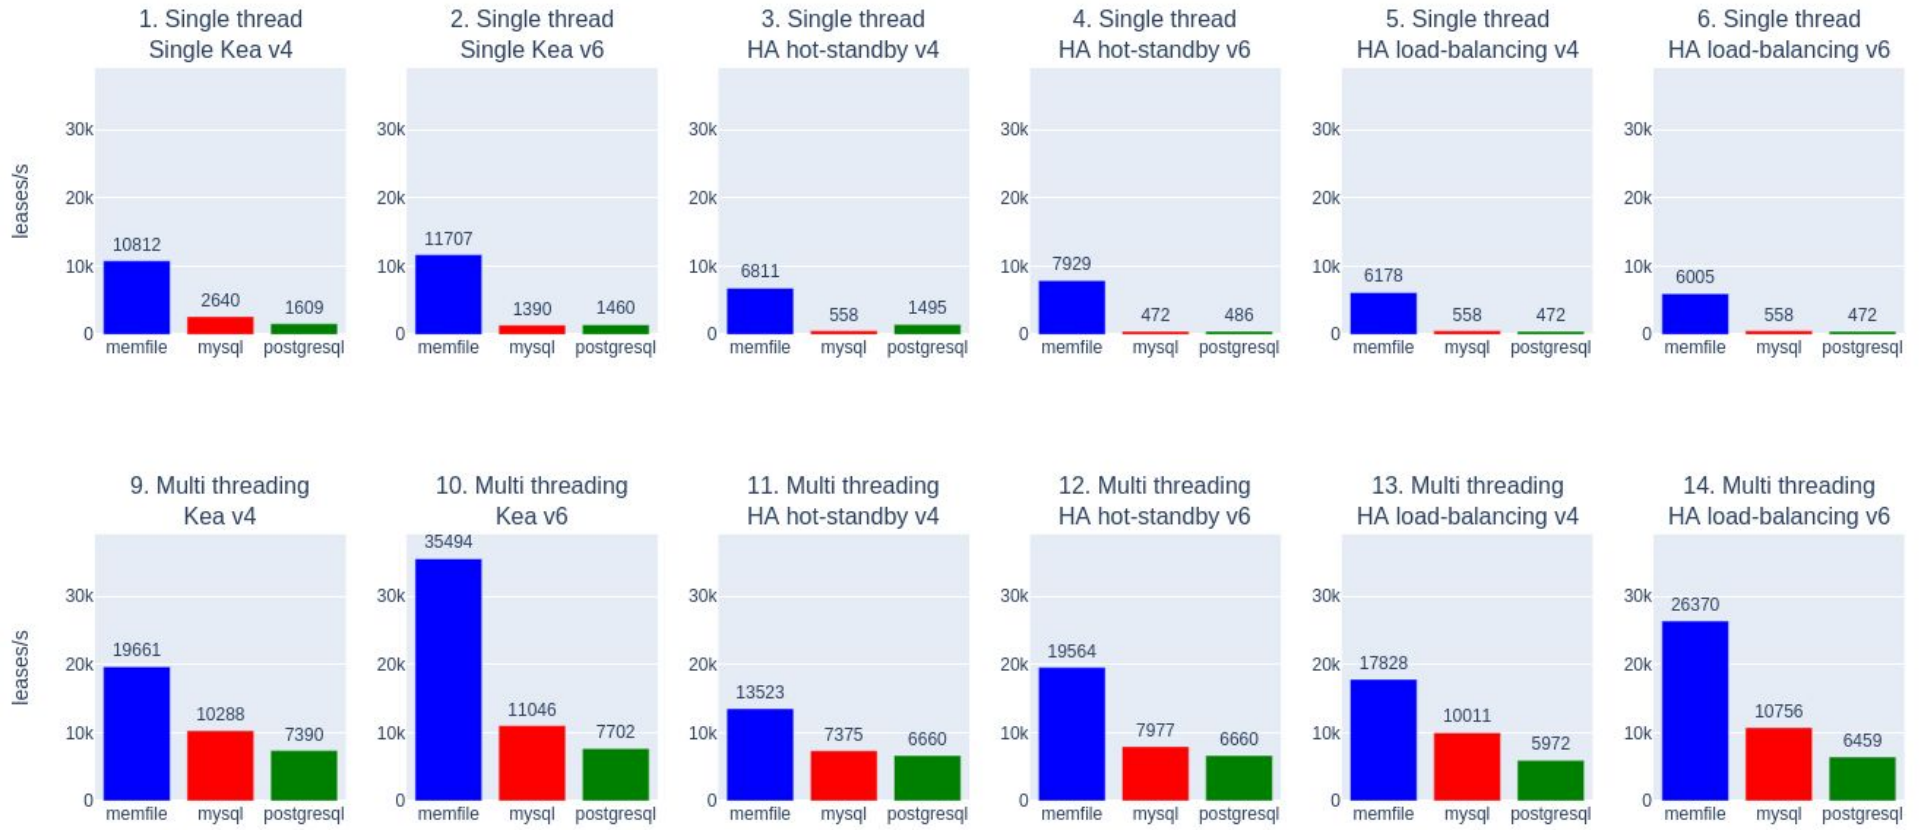

TLS support

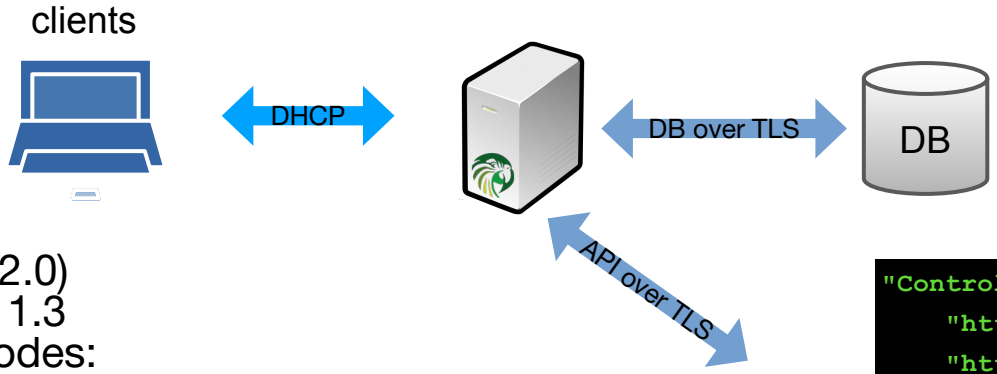

API (Kea 2.0)

- TLS 1.3
- 3 modes:
  - Disabled
  - Encryption (client verifies server)
  - Mutual (both sides verifies each other)

Secure database connection

- MySQL (Kea 2.1)
- PostgreSQL (Kea 2.0)

More security enhancements coming!

```
"Control-agent": {
    "http-host": "127.0.0.1",
    "http-port": 8000,

    // TLS trust anchor (Certificate Authority).
    "trust-anchor": "my-ca",

    // TLS server certificate file name.
    "cert-file": "my-cert",

    // TLS server private key file name.
    "key-file": "my-key",

    "cert-required": true
}
```

Cache Threshold (2.0)

- Problem: Buggy clients renewing early
- Each renewal:
  - Host reservation lookup
  - Lease lookup
  - Logging*
  - HA: partner update*
  - DNS Update*
- Solution: cache replies
- IPv4 and IPv6

```
"subnet6": [
    {
        "subnet": "2001:db8::/64",
        "pools": [ { "pool": "2001:db8::/64" } ],
        "renew-timer": :1000,
        "valid-lifetime": 2000,

        "cache-threshold": .25,
        "cache-max-age": 600,
        ...
    }
],
```

Script Hook (Kea 2.0)

But I want to … <your secret voodoo here>

```
{
  "hooks-libraries": [
  {
    "library": "libdhcp_run_script.so",
    "parameters": {
      "name": "/path/script.sh",
      "sync": false
    }
  },
  ... // other hooks
  ]
}
```

```
#!/bin/bash

lease4_renew () {
    ...
}

case "$1" in
    "lease4_renew")
        lease4_renew
        ;;
    *)
        unknown_handle "${@}"
        ;;
esac
```

Stork

Kea (and BIND9) Dashboard/GUI/IPAM

Stork Dashboard

One Stork server + one or more agents

- Collects data from Kea/BIND9 services
- Aggregates data
- Web interface
- Export to Prometheus/Grafana

- Server details: version, build, installed hooks, cpu, memory
- Fault monitoring: subnet utilization, HA failures, log viewer
- Statistics: DORAs, QPS, NAKs
- Config viewer: file locations, database backends, etc

- Monthly* releases
- Dashboard for now, configuration management coming up in 1.3

Stork Dashboard - Subnet Utilization

Stork DHCP Services Monitoring Configuration Help Search Logout (admin)

DHCP > Subnets

Filter subnets: subnet or any other field Protocol: any

Subnet ID	Subnet	Addresses			Pools	Shared Network	App Name
		Total	Assigned	Used %			
1	192.0.5.0/24	50	42	84 %	192.0.5.1-192.0.5.50	frog	kea@agent-kea
2	192.0.6.0/24	110	0	0 %	192.0.6.1-192.0.6.40 192.0.6.61-192.0.6.90 192.0.6.111-192.0.6.150	frog	kea@agent-kea
3	192.0.7.0/24	50	50	100 %	192.0.7.1-192.0.7.50	frog	kea@agent-kea
4	192.0.8.0/24	50	0	0 %	192.0.8.1-192.0.8.50	frog	kea@agent-kea
5	192.0.9.0/24	50	0	0 %	192.0.9.1-192.0.9.50	frog	kea@agent-kea
6	192.1.15.0/24	50	20	40 %	192.1.15.1-192.1.15.50	mouse	kea@agent-kea
7	192.1.16.0/24	150	39	26 %	192.1.16.1-192.1.16.50 192.1.16.51-192.1.16.100 192.1.16.101-192.1.16.150	mouse	kea@agent-kea
8	192.1.17.0/24	245	0	0 %	192.1.17.1-192.1.17.20 192.1.17.21-192.1.17.40 192.1.17.41-192.1.17.60 192.1.17.66-192.1.17.80 192.1.17.81-192.1.17.100 192.1.17.101-192.1.17.120 192.1.17.121-192.1.17.140 192.1.17.141-192.1.17.160 192.1.17.161-192.1.17.180 192.1.17.181-192.1.17.200 192.1.17.201-192.1.17.220 192.1.17.221-192.1.17.240 192.1.17.241-192.1.17.243 192.1.17.244-192.1.17.246 192.1.17.247-192.1.17.250	mouse	kea@agent-kea
9	192.0.2.0/24	200	1	0.5 %	192.0.2.1-192.0.2.50 192.0.2.51-192.0.2.100 192.0.2.101-192.0.2.150 192.0.2.151-192.0.2.200		kea@agent-kea
10	1.0.0.0/16	65,531	0	0 %	1.0.0.4-1.0.255.254		kea@agent-kea-many-subnets

1 of 694 pages « ‹ 1 2 3 4 5 › » 10 Total: 6931 subnets

Stork GUI - Monitoring HA Status

- Groups HA pairs
- Displays roles
  - Primary/standby
  - Load balancing
- Heartbeat status
- HA States
- Scopes served
- Last outage

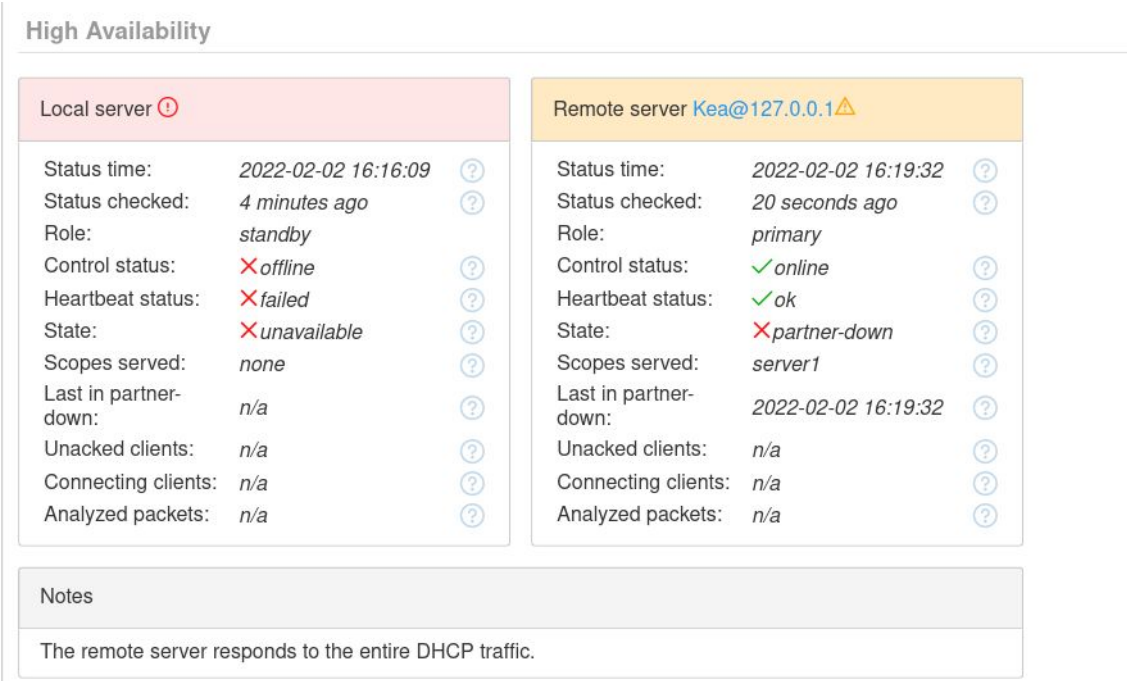

Prometheus / Grafana export

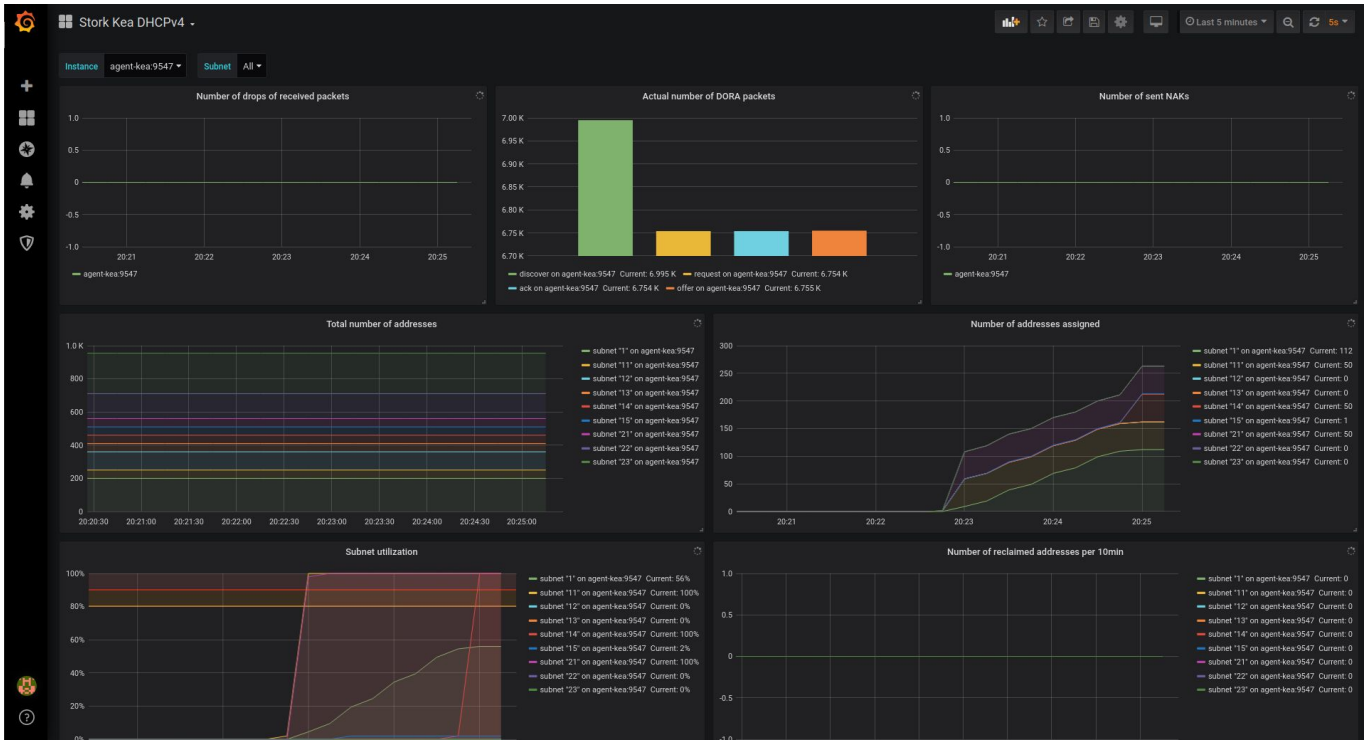

Participation is Welcome!

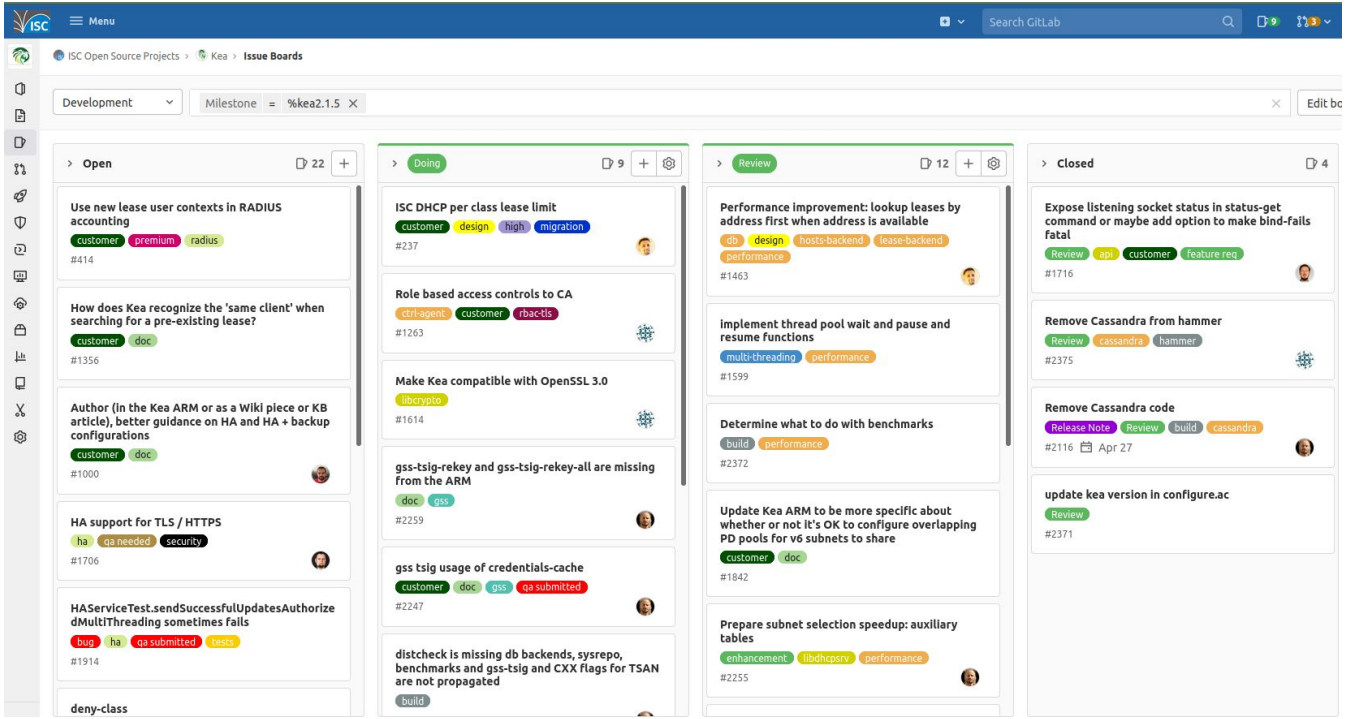

https://gitlab.isc.org/isc-projects/kea/

https://gitlab.isc.org/isc-projects/stork/

Questions?

isc.org/kea
gitlab.isc.org/isc-projects/kea
gitlab.isc.org/isc-projects/stork