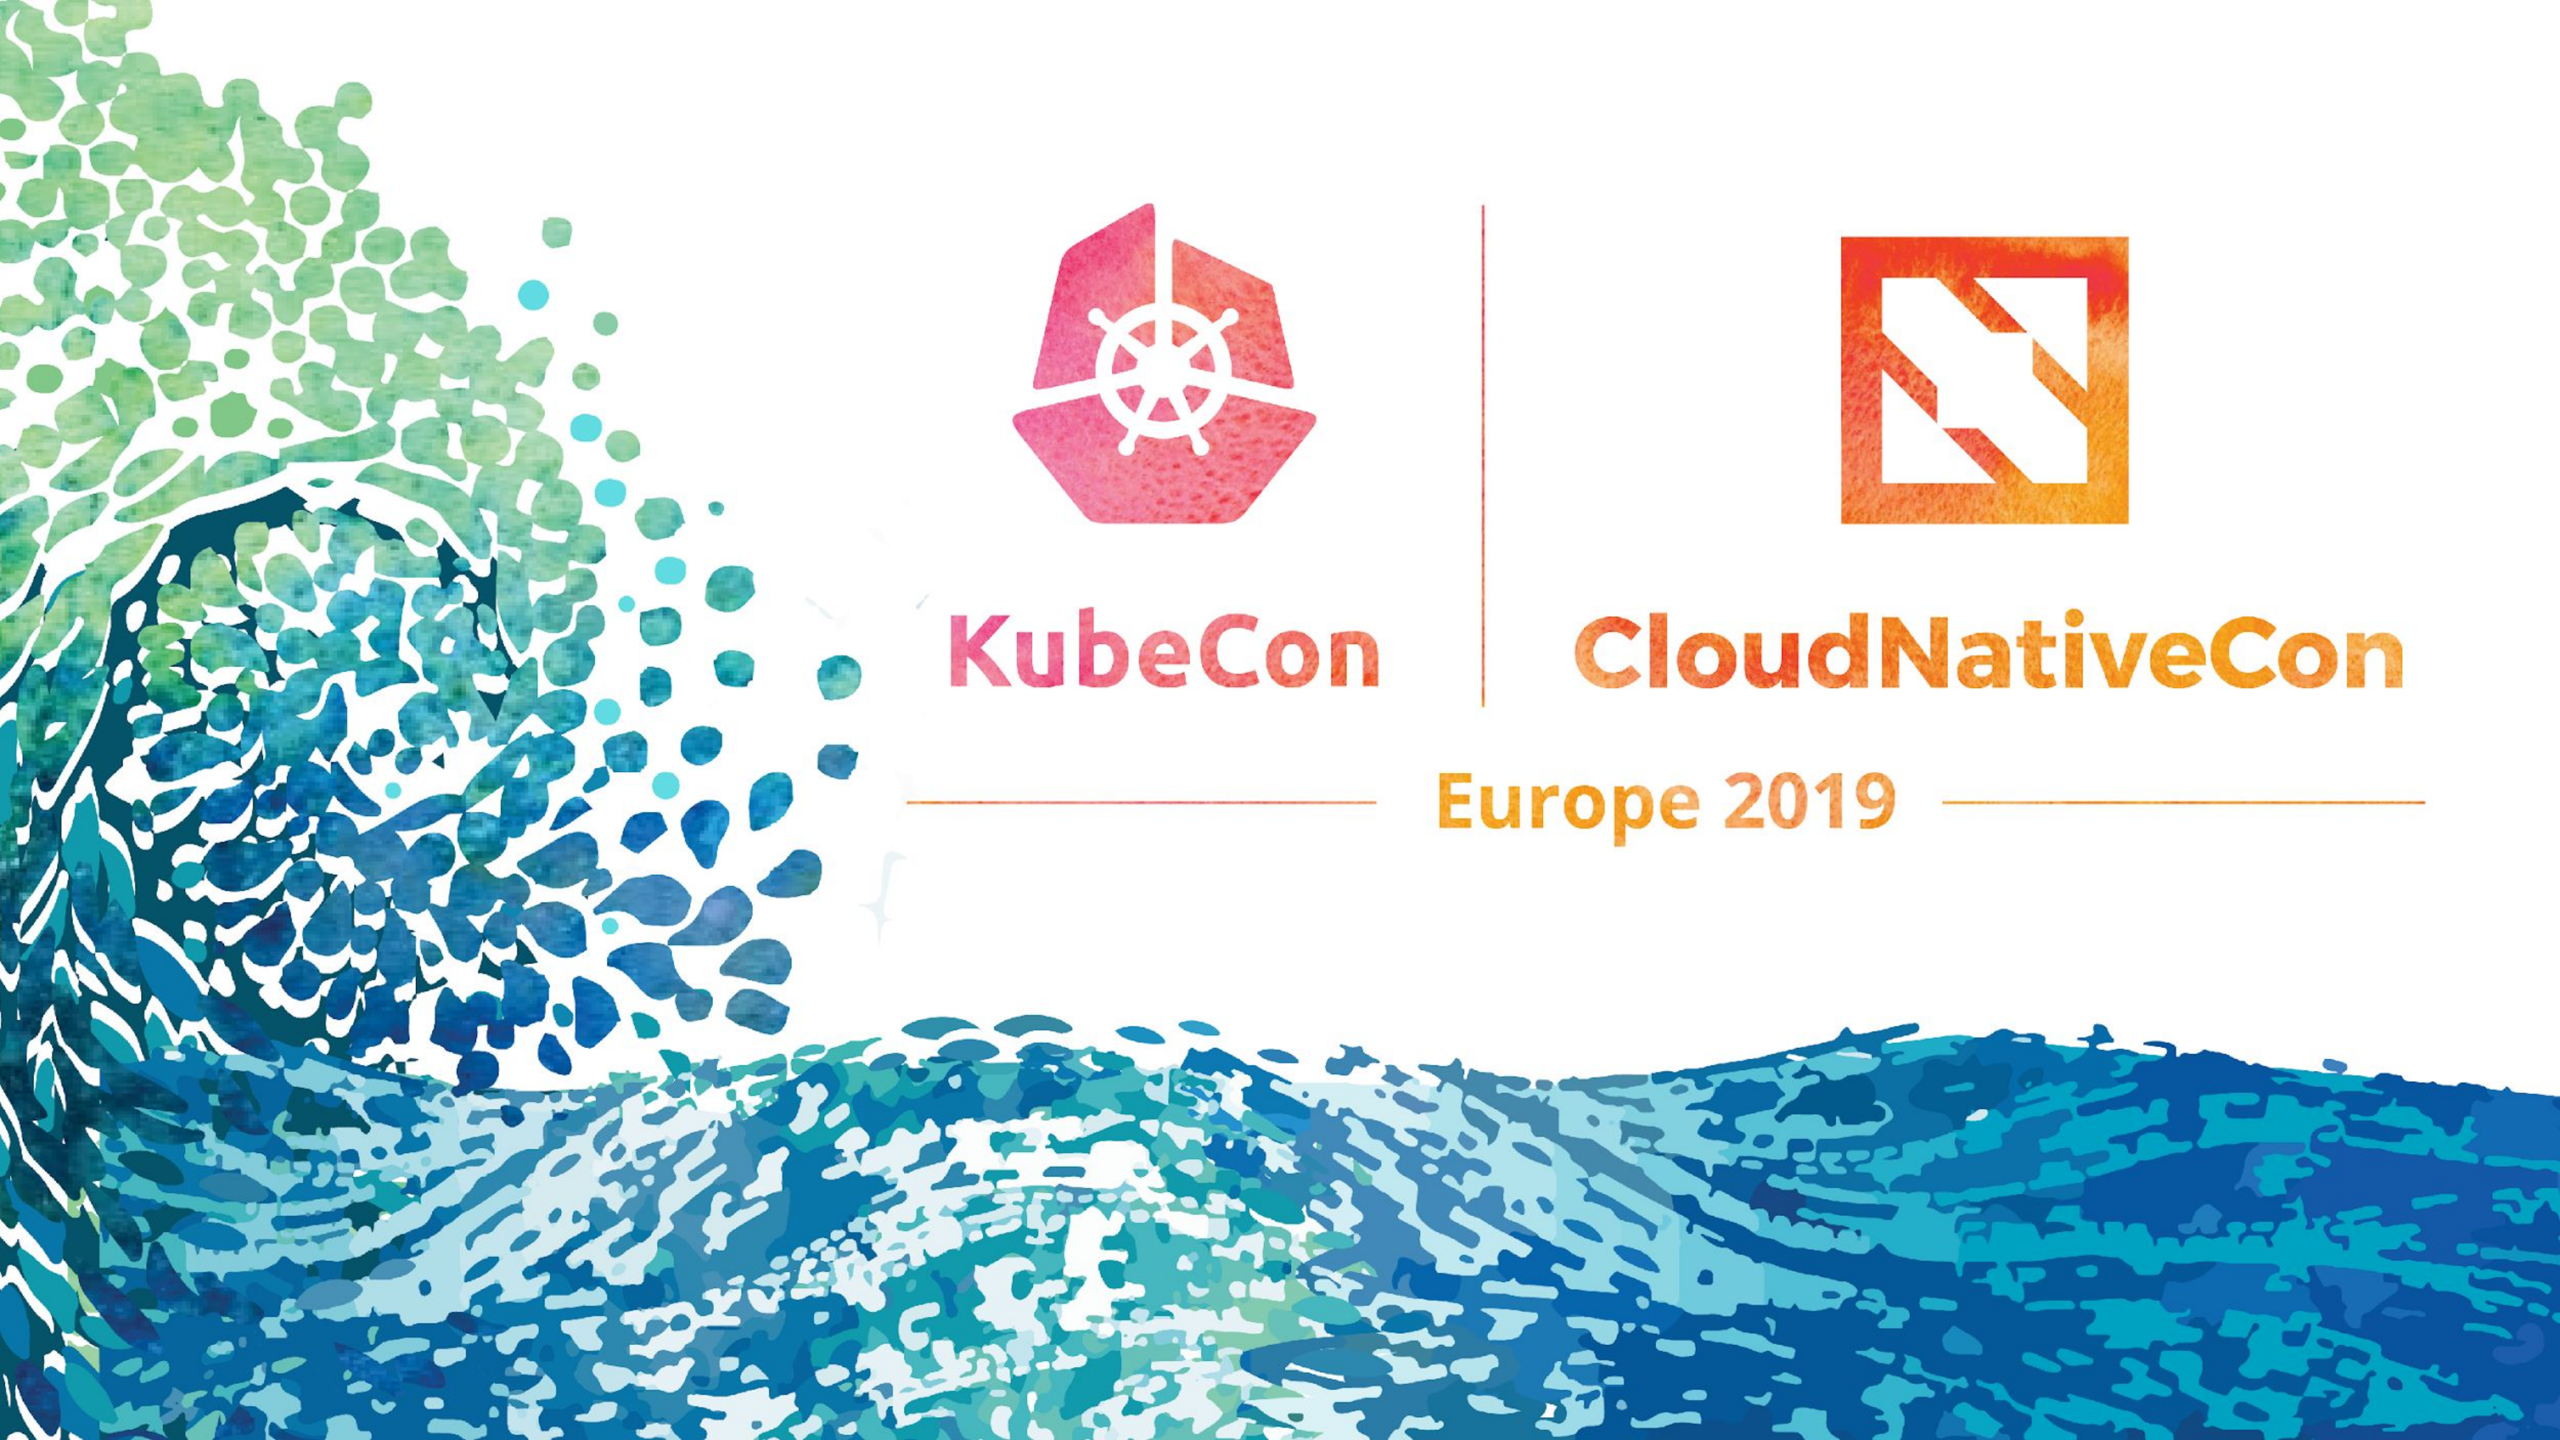

KubeCon | CloudNativeCon

Europe 2019

KubeCon | CloudNativeCon

Europe 2019

NATS Deep Dive

Waldemar Quevedo
Synadia Communications, Inc

About the speaker

- Waldemar Quevedo / @wallyqs
- Software Engineer at **Synadia Communications, Inc**
- Using NATS based systems since 2012
- Author of *Practical NATS* (Apress, 2018)

  https://www.apress.com/gp/book/9781484235690

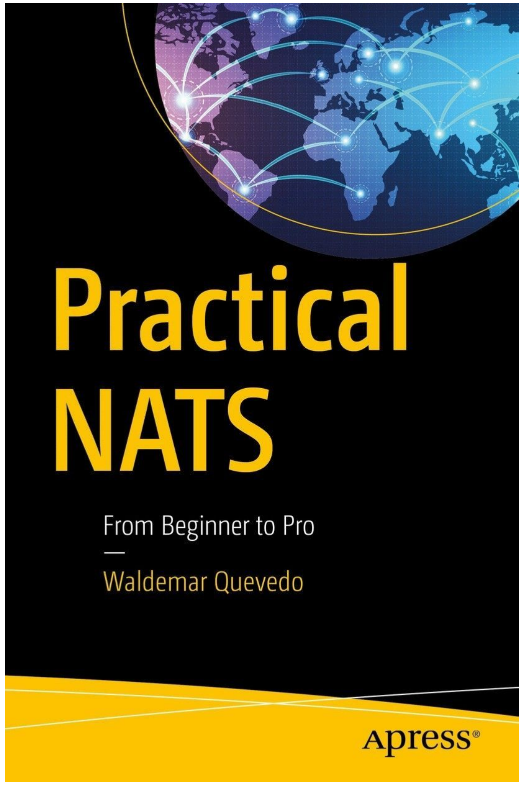

Practical NATS

From Beginner to Pro

Waldemar Quevedo

Apress®

Agenda

- New features in NATS v2
- Deploying NATS on Kubernetes

NATS v2

To be released June 10th

Questions?

NATS v2

- NATS as a core component that can be used to build a global communication network.
- An always available dial-tone, a global utility.
- telnet connect.ngs.global 4222

https://synadia.com/ngs

NATS v2 Features

- Multi-Tenancy and Accounts
- New Clustering Protocol
  - Client protocol is 100% backward compatible
- Streams and Services (Account sharing)
  - Imports/exports
- Security and Auth
  - NKEYS (ed25519 based keys)
  - TLS certs DN/SAN based auth
- Gateways, Super clusters & Leafnodes
- System Accounts
- Decentralized Management
- Graceful shutdown

KubeCon | CloudNativeCon

Europe 2019

Multi-Tenancy & Accounts

Containers for Messaging

NATS v1 Authorization

In NATS v1, all the subjects exists in the same shared space for all users.

→ All servers have to share the same configuration too.

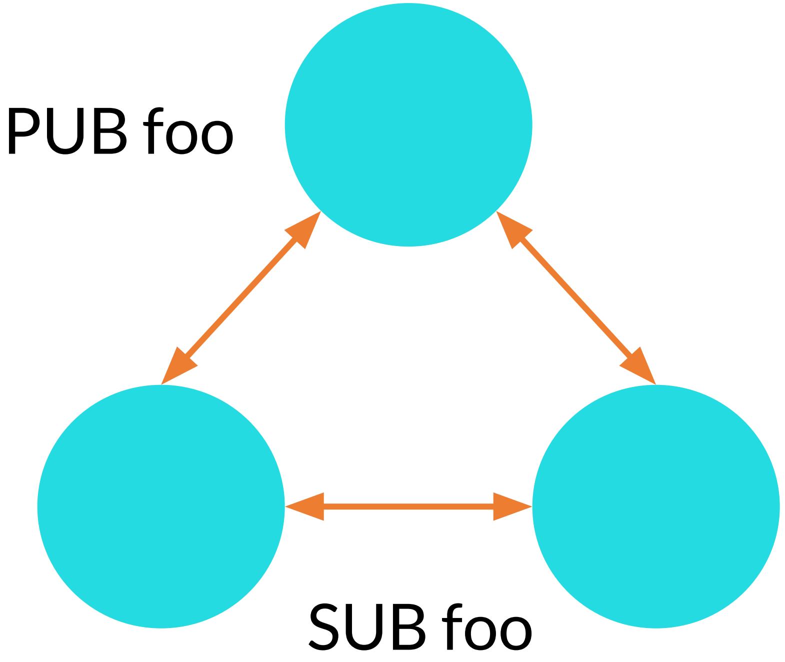

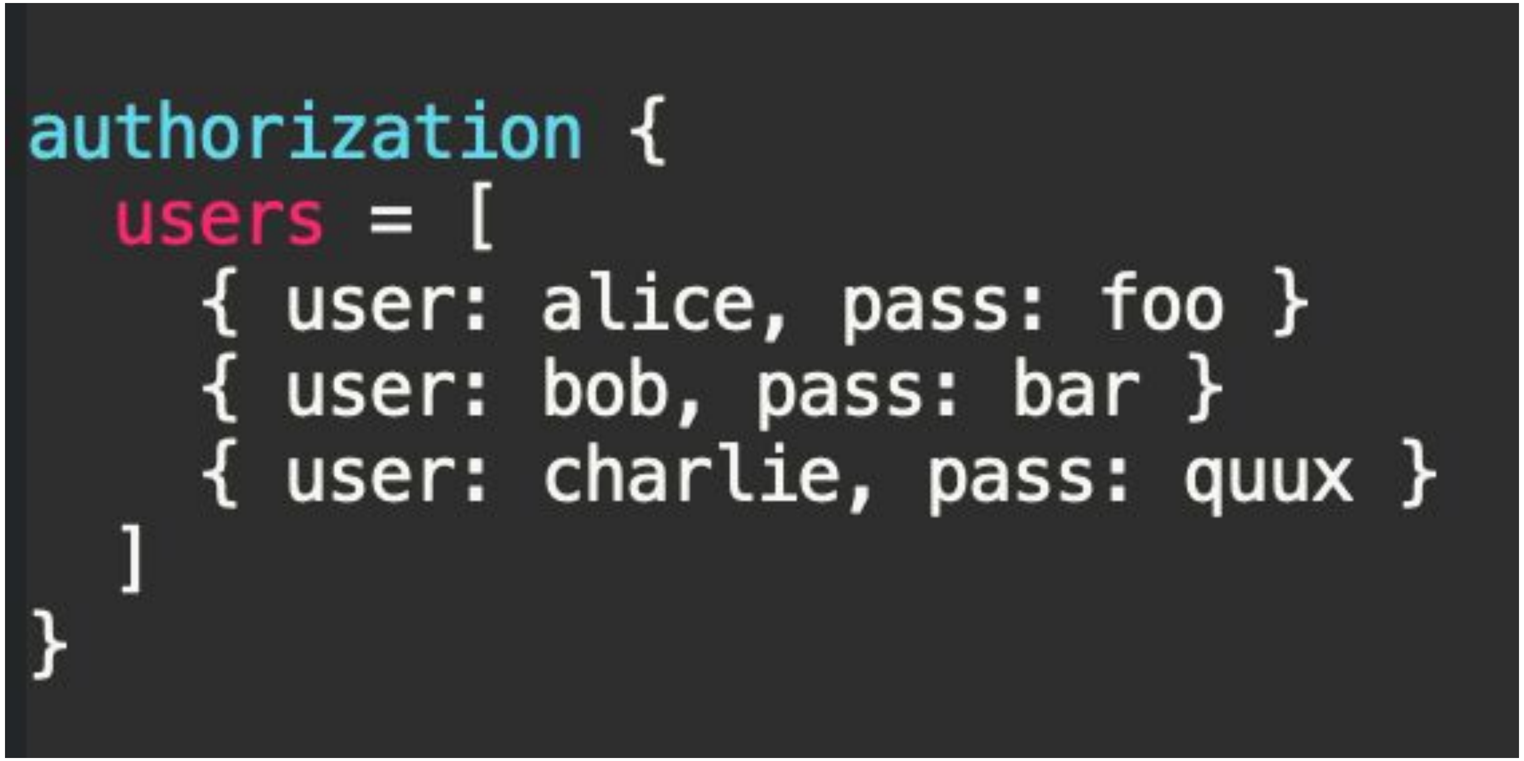

```
authorization {
  users = [
    { user: alice, pass: foo }
    { user: bob, pass: bar }
    { user: charlie, pass: quux }
  ]
}
```

NATS v1 Authorization

- Permissions can be applied for a user in order to prevent or allow sharing data between multiple users.

```
CONNECT {"user": "alice", "pass": "foo"}
+OK
PUB hello 5
world
-ERR 'Permissions Violation for Publish to "hello"'
```

```
CONNECT {"user": "bob", "pass": "bar"}
+OK
SUB greetings 5
-ERR 'Permissions Violation for Subscription to "greetings"'
PUB greetings 5
hello
```

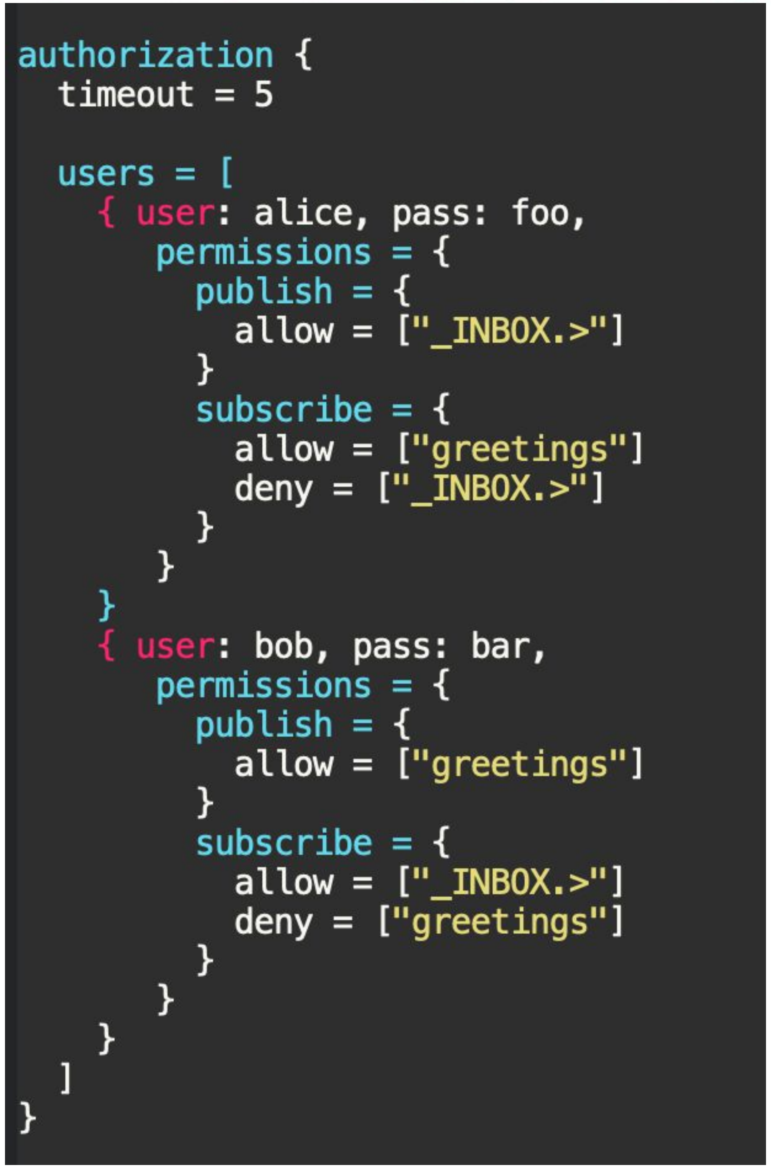

```
authorization {
  timeout = 5

  users = [
    { user: alice, pass: foo,
      permissions = {
        publish = {
          allow = ["_INBOX.>"]
        }
        subscribe = {
          allow = ["greetings"]
          deny = ["_INBOX.>"]
        }
      }
    }
    { user: bob, pass: bar,
      permissions = {
        publish = {
          allow = ["greetings"]
        }
        subscribe = {
          allow = ["_INBOX.>"]
          deny = ["greetings"]
        }
      }
    }
  ]
}
```

Introducing Accounts

In NATS v2, it is now possible to isolate subject spaces.

```
CONNECT {"user": "alice", "pass": "foo"}
+OK
SUB greetings 1
+OK


CONNECT {"user": "bob", "pass": "bar"}
+OK
PUB greetings 5
hello
```

Message is not forwarded since users isolated into different accounts

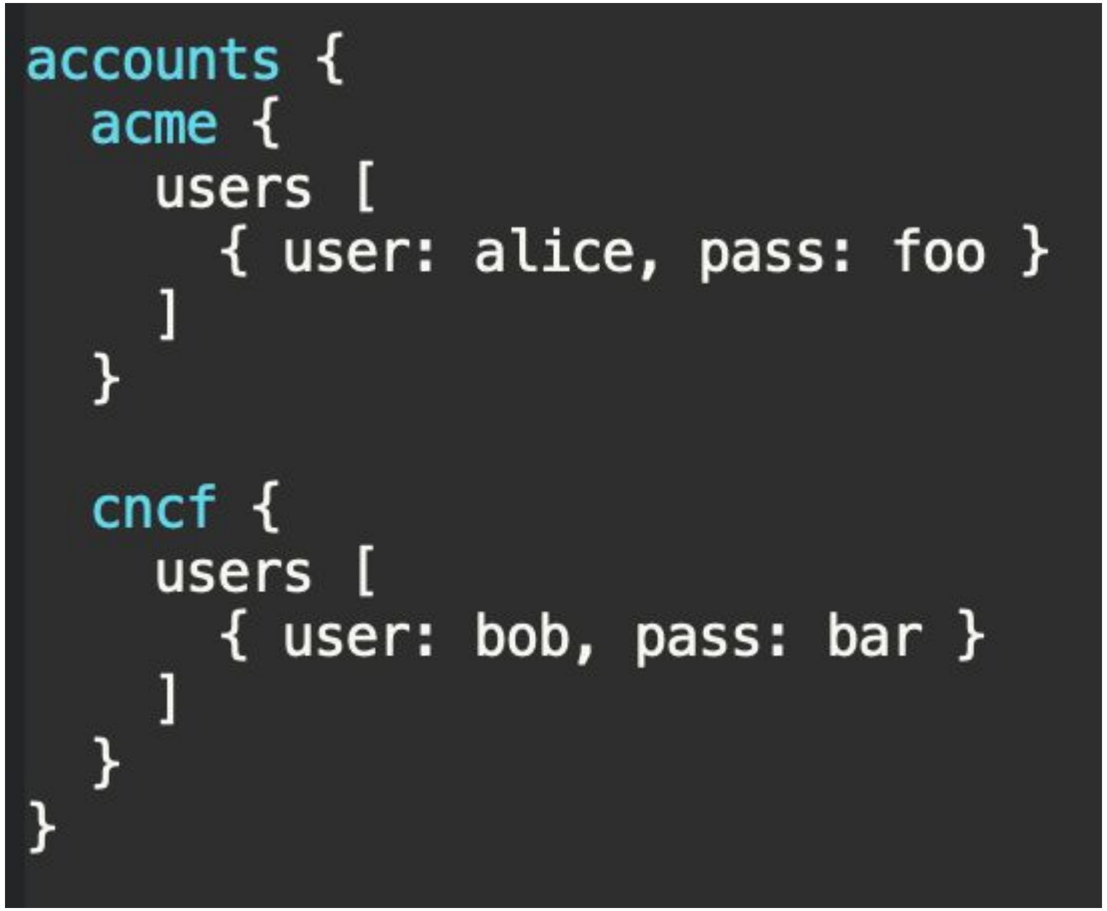

```
accounts {
  acme {
    users [
      { user: alice, pass: foo }
    ]
  }

  cncf {
    users [
      { user: bob, pass: bar }
    ]
  }
}
```

Accounts

- Accounts are isolated communication contexts allowing secure multi-tenancy
  - ✓ Containers for messaging! 📦
- Bifurcate technology from business driven use cases
  - ✓ Data silos are created by design, not software limitations
- Easy, Secure and Cost Effective
  - ✓ One NATS deployment for operators to manage
  - ✓ Decentralized - organizations can self-manage
- Share data between accounts
  - ✓ Secure Streams and Services
  - ✓ Only mutual agreement will permit data flow

Accounts + Clustering

- **Clustering protocol** reworked to support account isolation and multi tenancy use cases. Not compatible with NATS v1.X

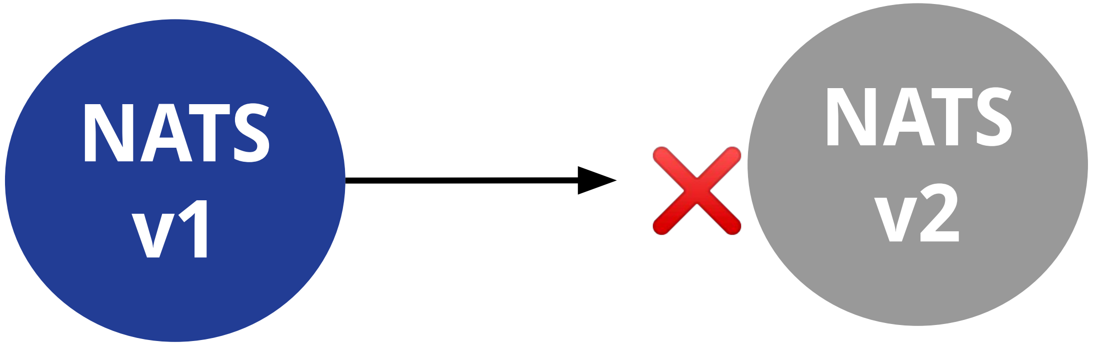

- **Clients protocol** is the same.
  - No changes to applications required

The Global Account

In NATS v2, even if no accounts are configured they are actually being part of a reserved global account named **$G**

```
- cid:2 - <<- [SUB hello 90]
- cid:2 - ->> [OK]
- rid:1 - ->> [RS+ $G hello]
- rid:1 - <<- [RMSG $G hello 5]
- rid:1 - <<- MSG_PAYLOAD: ["world"]
- cid:2 - ->> [MSG hello 90 5]
```

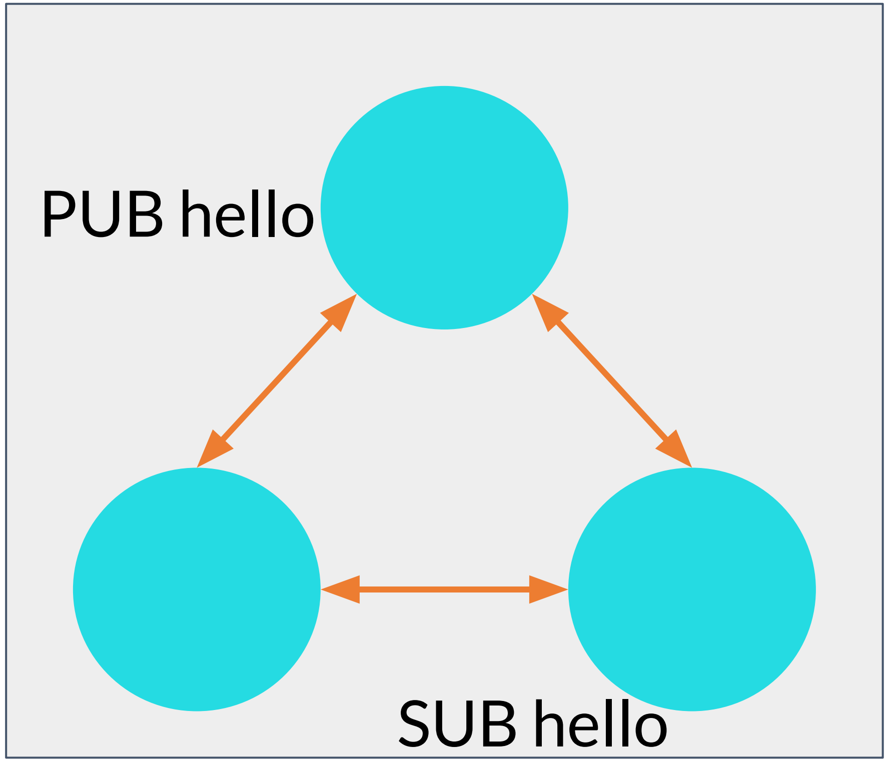

KubeCon | CloudNativeCon

Europe 2019

Accounts Sharing

Account Sharing

In order to introduce the concept of account sharing, we need to define two new abstract concepts: **Streams & Services**

- **Services** represent secure RPC endpoints

- **Streams** represent data flow between accounts

Services

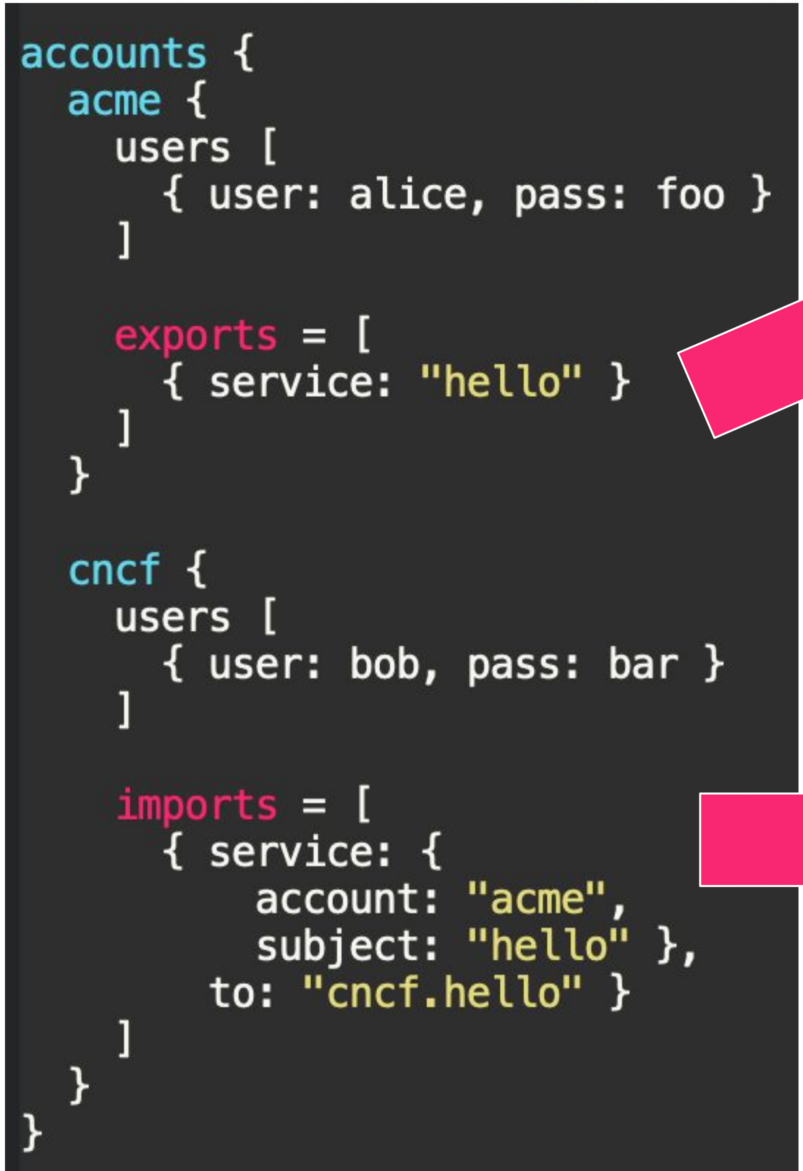

```
accounts {
  acme {
    users [
      { user: alice, pass: foo }
    ]

    exports = [
      { service: "hello" }
    ]
  }

  cncf {
    users [
      { user: bob, pass: bar }
    ]

    imports = [
      { service: {
            account: "acme",
            subject: "hello" },
        to: "cncf.hello" }
    ]
  }
}
```

Export

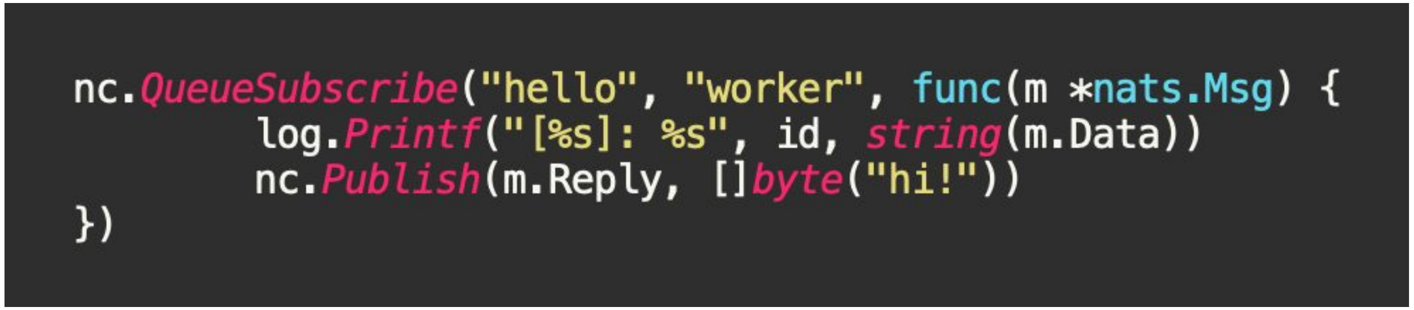

```
nc.QueueSubscribe("hello", "worker", func(m *nats.Msg) {
        log.Printf("[%s]: %s", id, string(m.Data))
        nc.Publish(m.Reply, []byte("hi!"))
})
```

Import

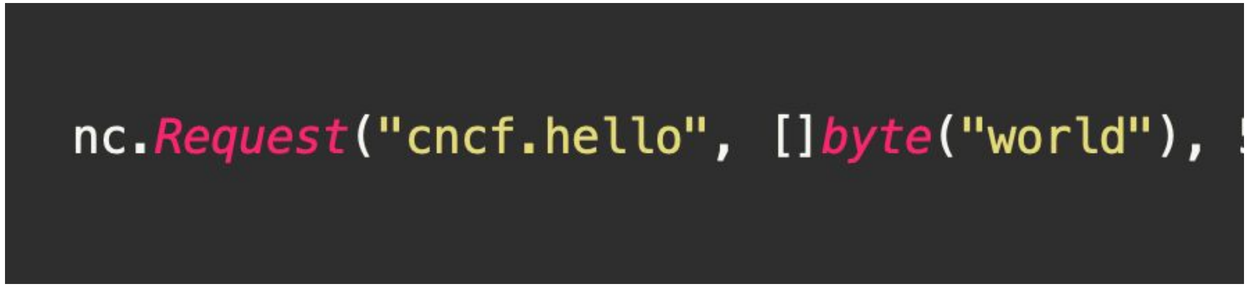

```
nc.Request("cncf.hello", []byte("world"),
```

Requests across accounts

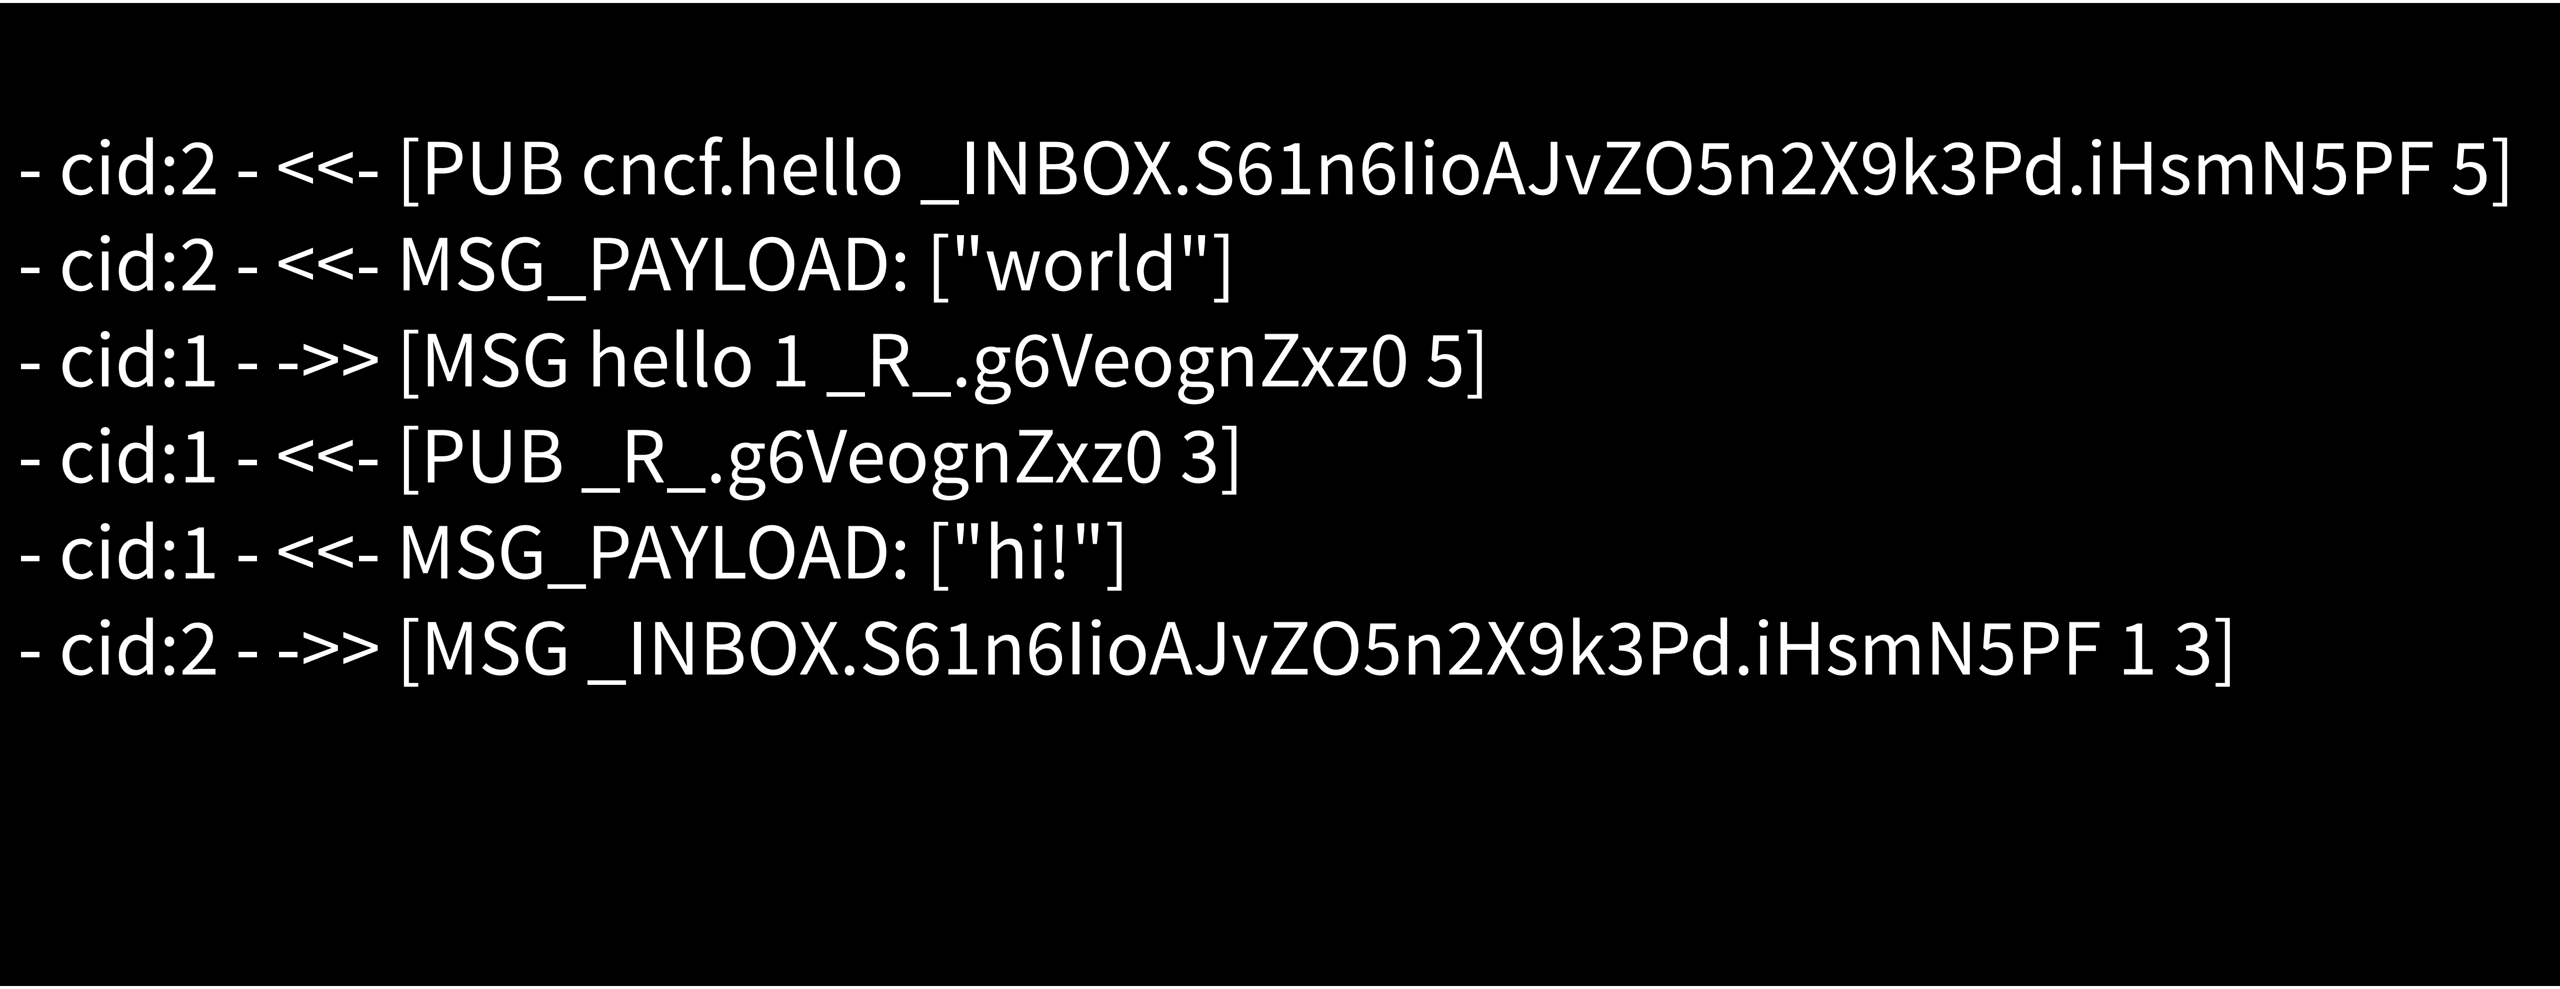

```
- cid:2 - <<- [PUB cncf.hello _INBOX.S61n6IioAJvZO5n2X9k3Pd.iHsmN5PF 5]
- cid:2 - <<- MSG_PAYLOAD: ["world"]
- cid:1 - ->> [MSG hello 1 _R_.g6VeognZxz0 5]
- cid:1 - <<- [PUB _R_.g6VeognZxz0 3]
- cid:1 - <<- MSG_PAYLOAD: ["hi!"]
- cid:2 - ->> [MSG _INBOX.S61n6IioAJvZO5n2X9k3Pd.iHsmN5PF 1 3]
```

Requests across accounts

```
- cid:2 - <<- [PUB cncf.hello _INBOX.S61n6IioAJvZO5n2X9k3Pd.iHsmN5PF 5]
- cid:2 - <<- MSG_PAYLOAD: ["world"]
- cid:1 - ->> [MSG hello 1 _R_.g6VeognZxz0 5]
- cid:1 - <<- [PUB _R_.g6VeognZxz0 3]
- cid:1 - <<- MSG_PAYLOAD: ["hi!"]
- cid:2 - ->> [MSG _INBOX.S61n6IioAJvZO5n2X9k3Pd.iHsmN5PF 1 3]
```

Streams

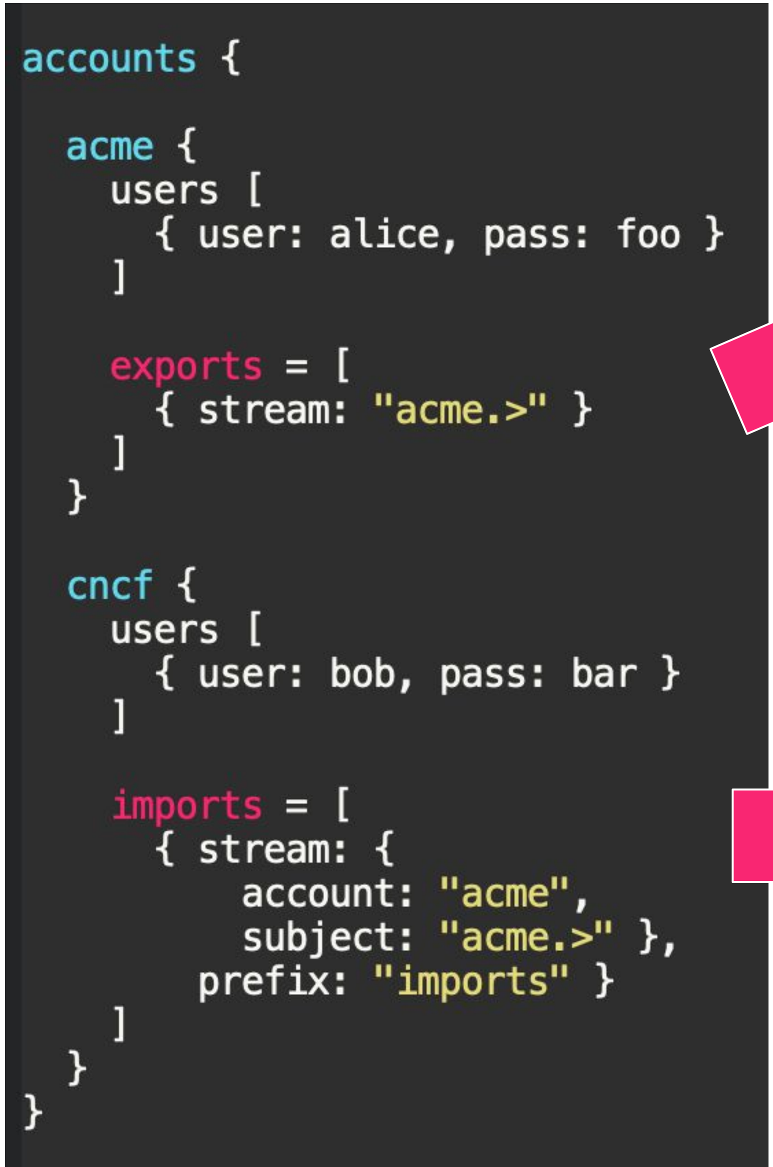

```
accounts {

  acme {
    users [
      { user: alice, pass: foo }
    ]

    exports = [
      { stream: "acme.>" }
    ]
  }

  cncf {
    users [
      { user: bob, pass: bar }
    ]

    imports = [
      { stream: {
          account: "acme",
          subject: "acme.>" },
        prefix: "imports" }
    ]
  }
}
```

Export

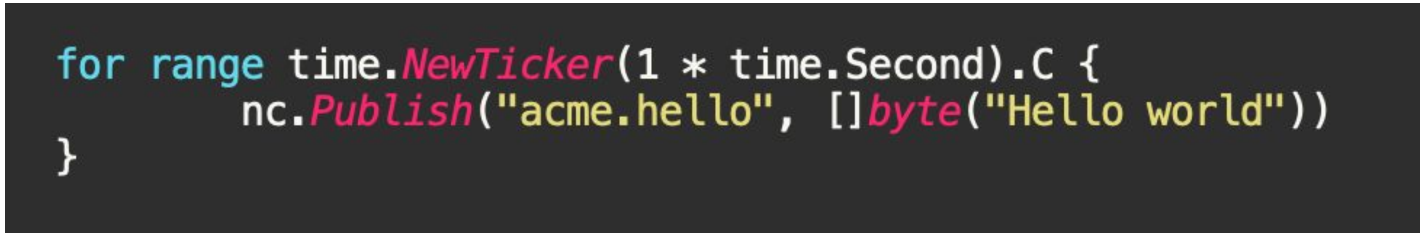

```
for range time.NewTicker(1 * time.Second).C {
        nc.Publish("acme.hello", []byte("Hello world"))
}
```

Import

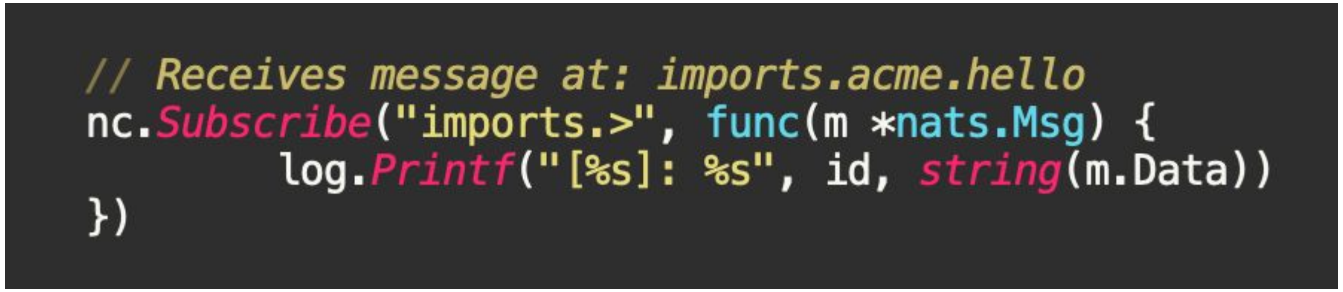

```
// Receives message at: imports.acme.hello
nc.Subscribe("imports.>", func(m *nats.Msg) {
        log.Printf("[%s]: %s", id, string(m.Data))
})
```

Security & Authentication

NKEYS

https://github.com/nats-io/nkeys

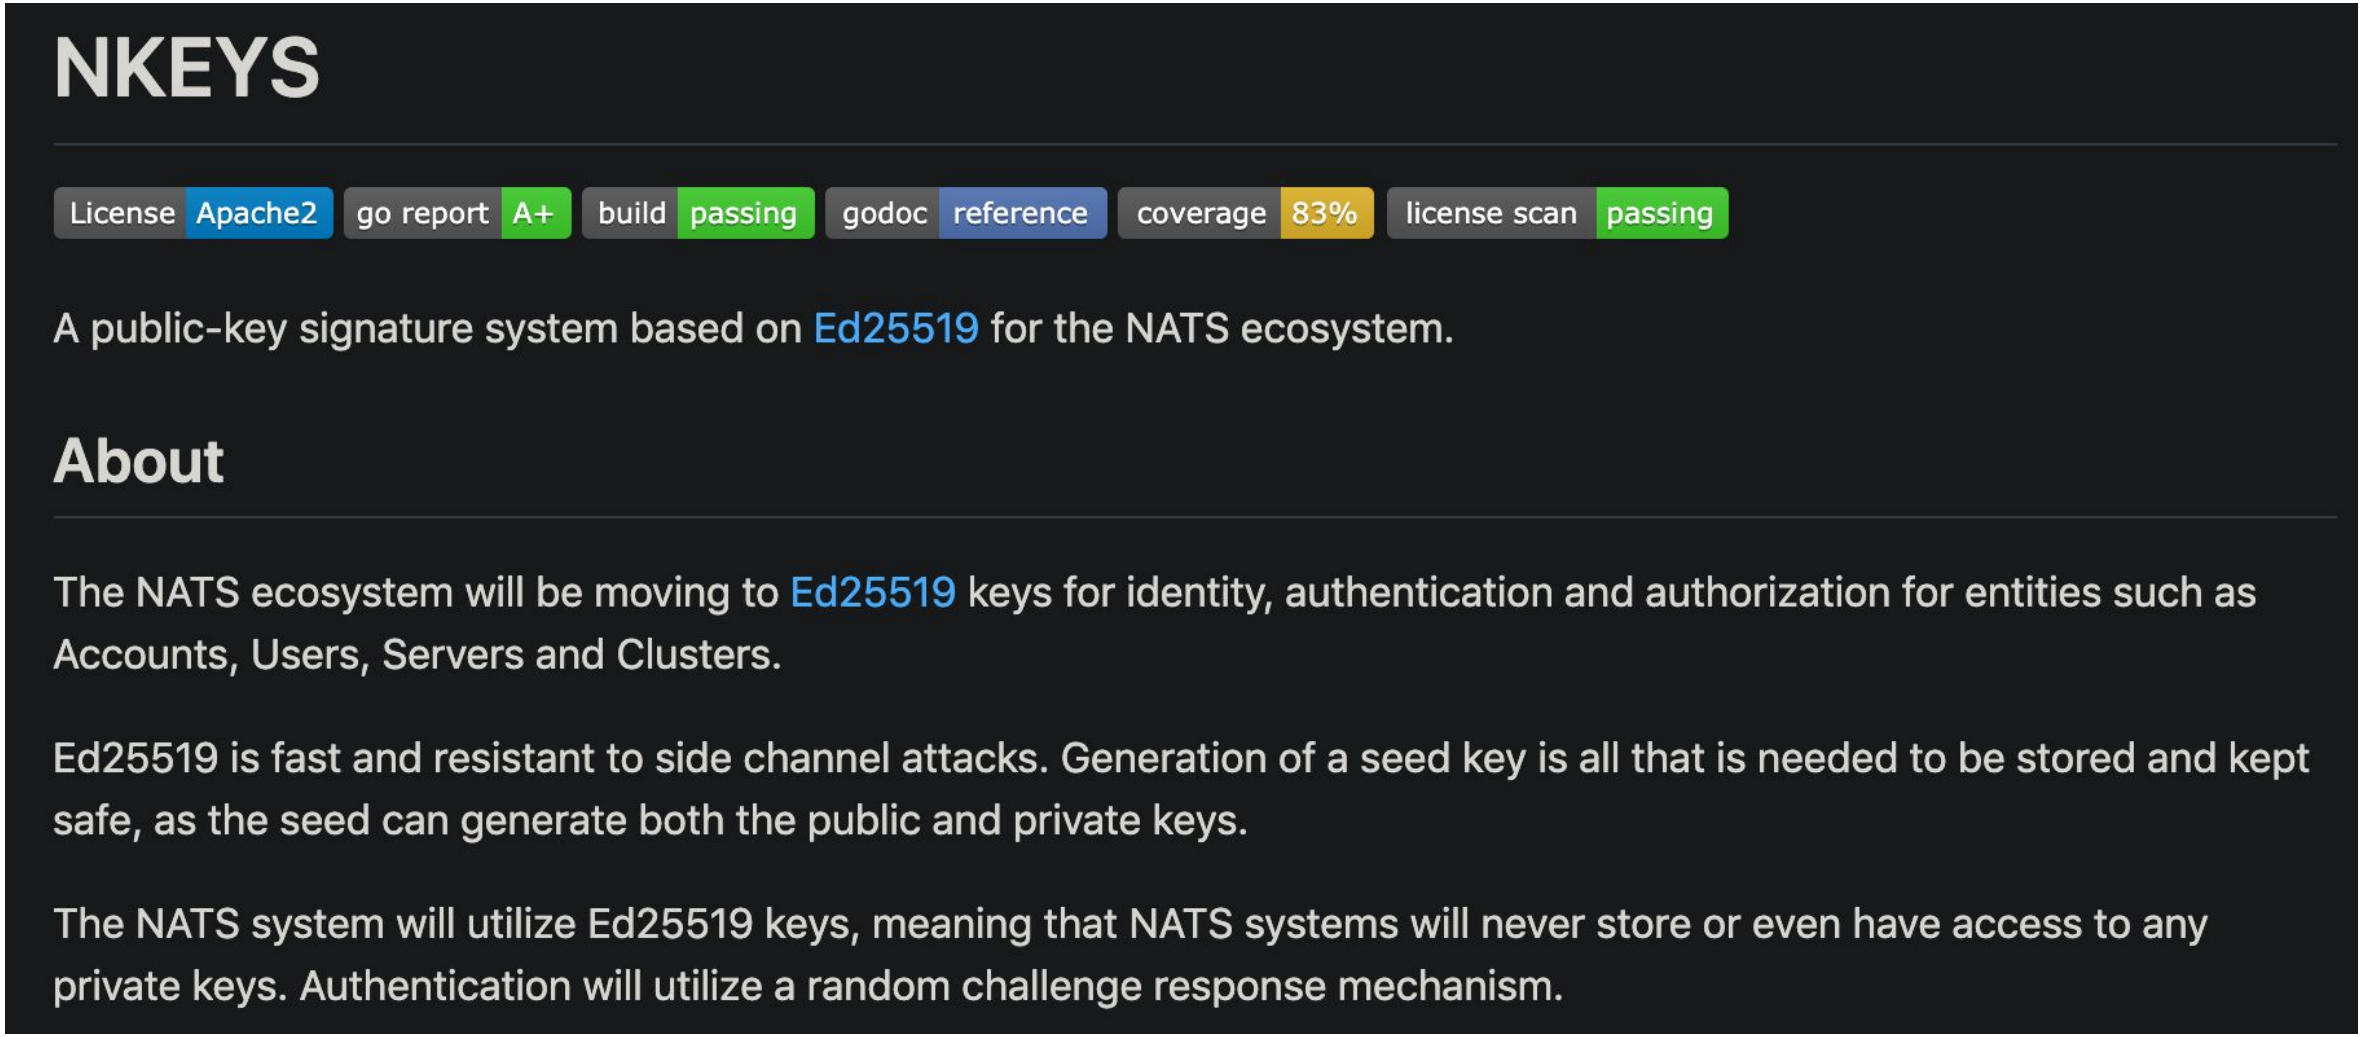

NKEYS

License Apache2 | go report A+ | build passing | godoc reference | coverage 83% | license scan passing

A public-key signature system based on Ed25519 for the NATS ecosystem.

About

The NATS ecosystem will be moving to Ed25519 keys for identity, authentication and authorization for entities such as Accounts, Users, Servers and Clusters.

Ed25519 is fast and resistant to side channel attacks. Generation of a seed key is all that is needed to be stored and kept safe, as the seed can generate both the public and private keys.

The NATS system will utilize Ed25519 keys, meaning that NATS systems will never store or even have access to any private keys. Authentication will utilize a random challenge response mechanism.

NKEYS and JWTs

A new NATS Identity authentication and authorization system.

- ED25519 based encoded keys made simple
  - Fast and resistant to side-channel attacks
  - Sign and Verify
- NATS servers **never see private keys**
  - Server sends nonce during connect, verifies client signatures
- JWT associate users with accounts and permission sets

NKEYS and JWTs

- `nk` tool

  ```
  GO111MODULE=on go get github.com/nats-io/nkeys/nk
  ```

- Generate a User NKEY

  ```
  nk --gen user
  SUAD2IRGV5JSPYFDNDAAKCMVO4UCEUF3D24NJZD6JQDXVALJT5JUT67GSA
  ```

- Generate a User NKEY Public key

  ```
  nk -inkey user.nkey -pubout
  UARTETAADVR7EFHQTEQEG4CW6QZA6O5SVCI3PHRVAJ2OLJTNMVBSKUU6
  ```

- Generate a signature, and verify with the signature.

  ```
  nk -sign test.txt -inkey alice.nkey > alice.sig
  nk -verify test.txt -sigfile alice.sig -pubin alice.pub
  ```

Using NKEYS users

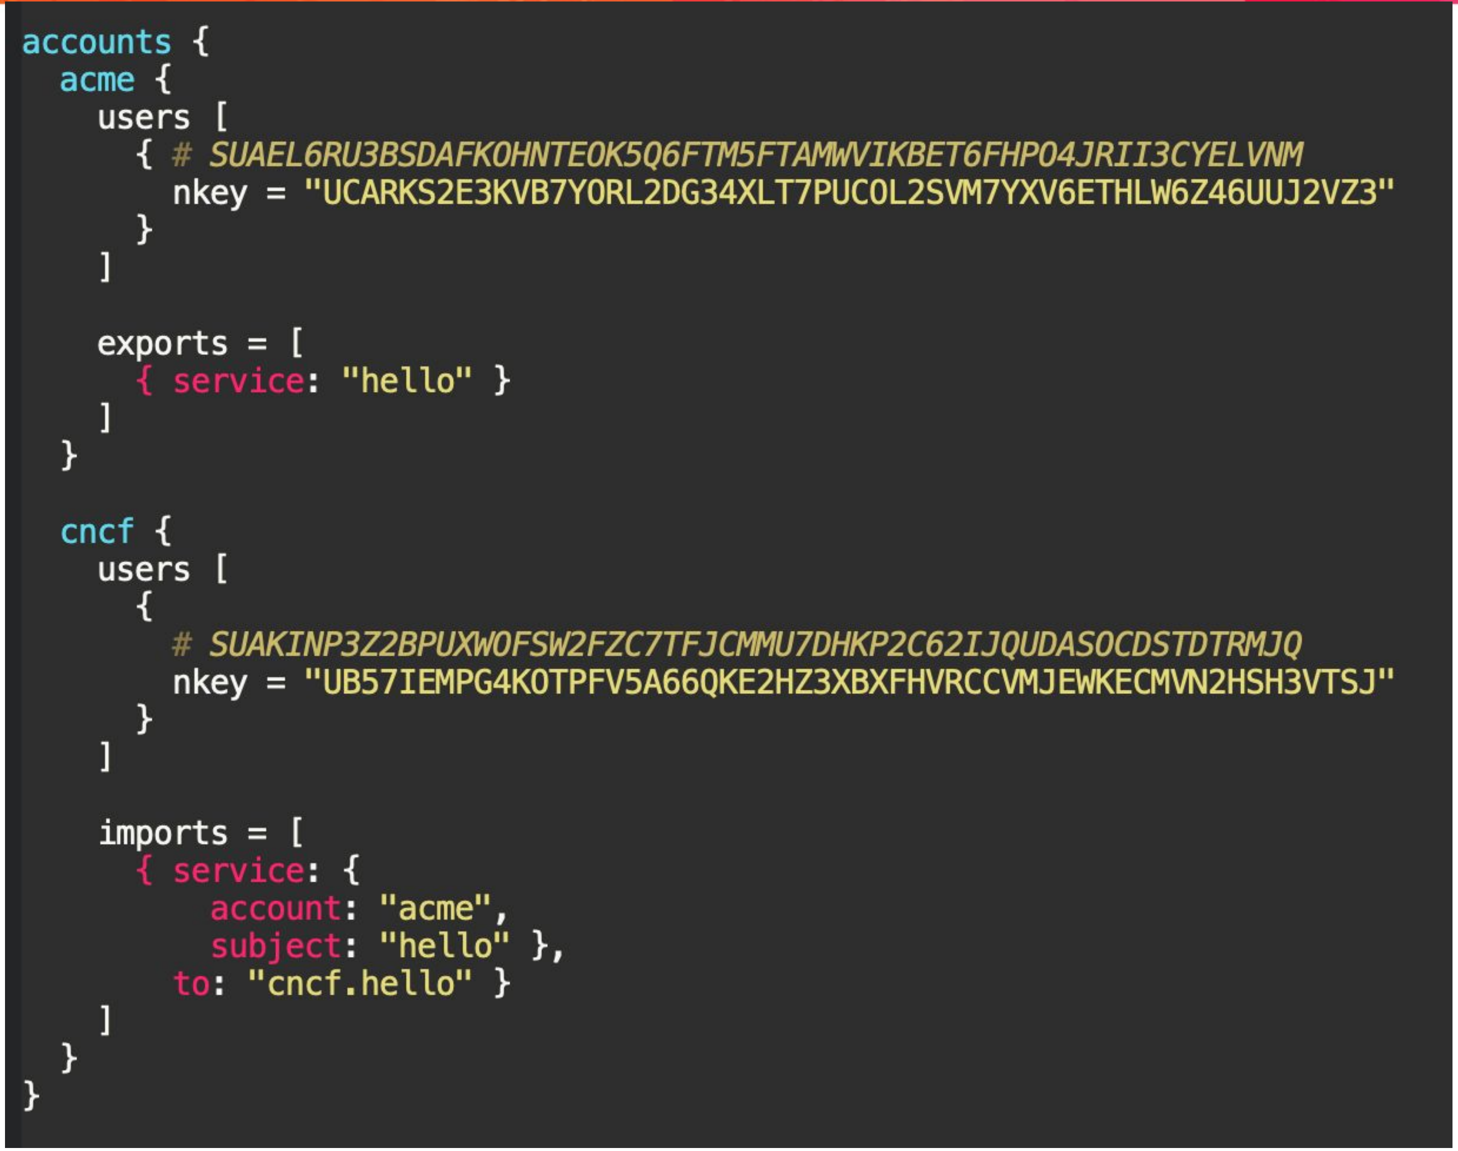

```
accounts {
  acme {
    users [
      { # SUAEL6RU3BSDAFKOHNTEOK5Q6FTM5FTAMWVIKBET6FHPO4JRII3CYELVNM
        nkey = "UCARKS2E3KVB7YORL2DG34XLT7PUCOL2SVM7YXV6ETHLW6Z46UUJ2VZ3"
      }
    ]

    exports = [
      { service: "hello" }
    ]
  }

  cncf {
    users [
      {
        # SUAKINP3Z2BPUXWOFSW2FZC7TFJCMMU7DHKP2C62IJQUDASOCDSTDTRMJQ
        nkey = "UB57IEMPG4KOTPFV5A66QKE2HZ3XBXFHVRCCVMJEWKECMVN2HSH3VTSJ"
      }
    ]

    imports = [
      { service: {
          account: "acme",
          subject: "hello" },
        to: "cncf.hello" }
    ]
  }
}
```

KubeCon | CloudNativeCon

Europe 2019

Decentralized Management

JWTs

JWTs are used to represent identities in NATS

- User, Account, Cluster, or Server

User JWTs Contain

- Account NKey (Issuer)
- Public NKey (Subject)
- Friendly Name
- Permissions
- Limits
- Not Before and Expiration

JWTs

```
{
 "jti": "3Y2OIRCSQLHOZI2KWXPS7JCRIR5BT5ZGZ5G74VHFCMUJAZUPCYCA",
 "iat": 1544140248,
 "iss": "ADQO262SKHLYIQTIBU3VG2K4GWRVO4TXYYJDHKI7QBMWYW6HACLQZIVB",
 "name": "Waldemar",
 "sub": "UCZRG6WDXWMIKDPLUMMRS2UAO2NSA5GOU2WCTXQLK7TRUWLLQ2CAXY7M",
 "type": "user",
 "nats": {
  "pub": {
   "allow": [
    "public.>"
   ]
  },
  "sub": {
   "deny": [
    "private.>"
   ]
  }
 }
}
```

NATS Account Server

https://github.com/nats-io/nats-account-server

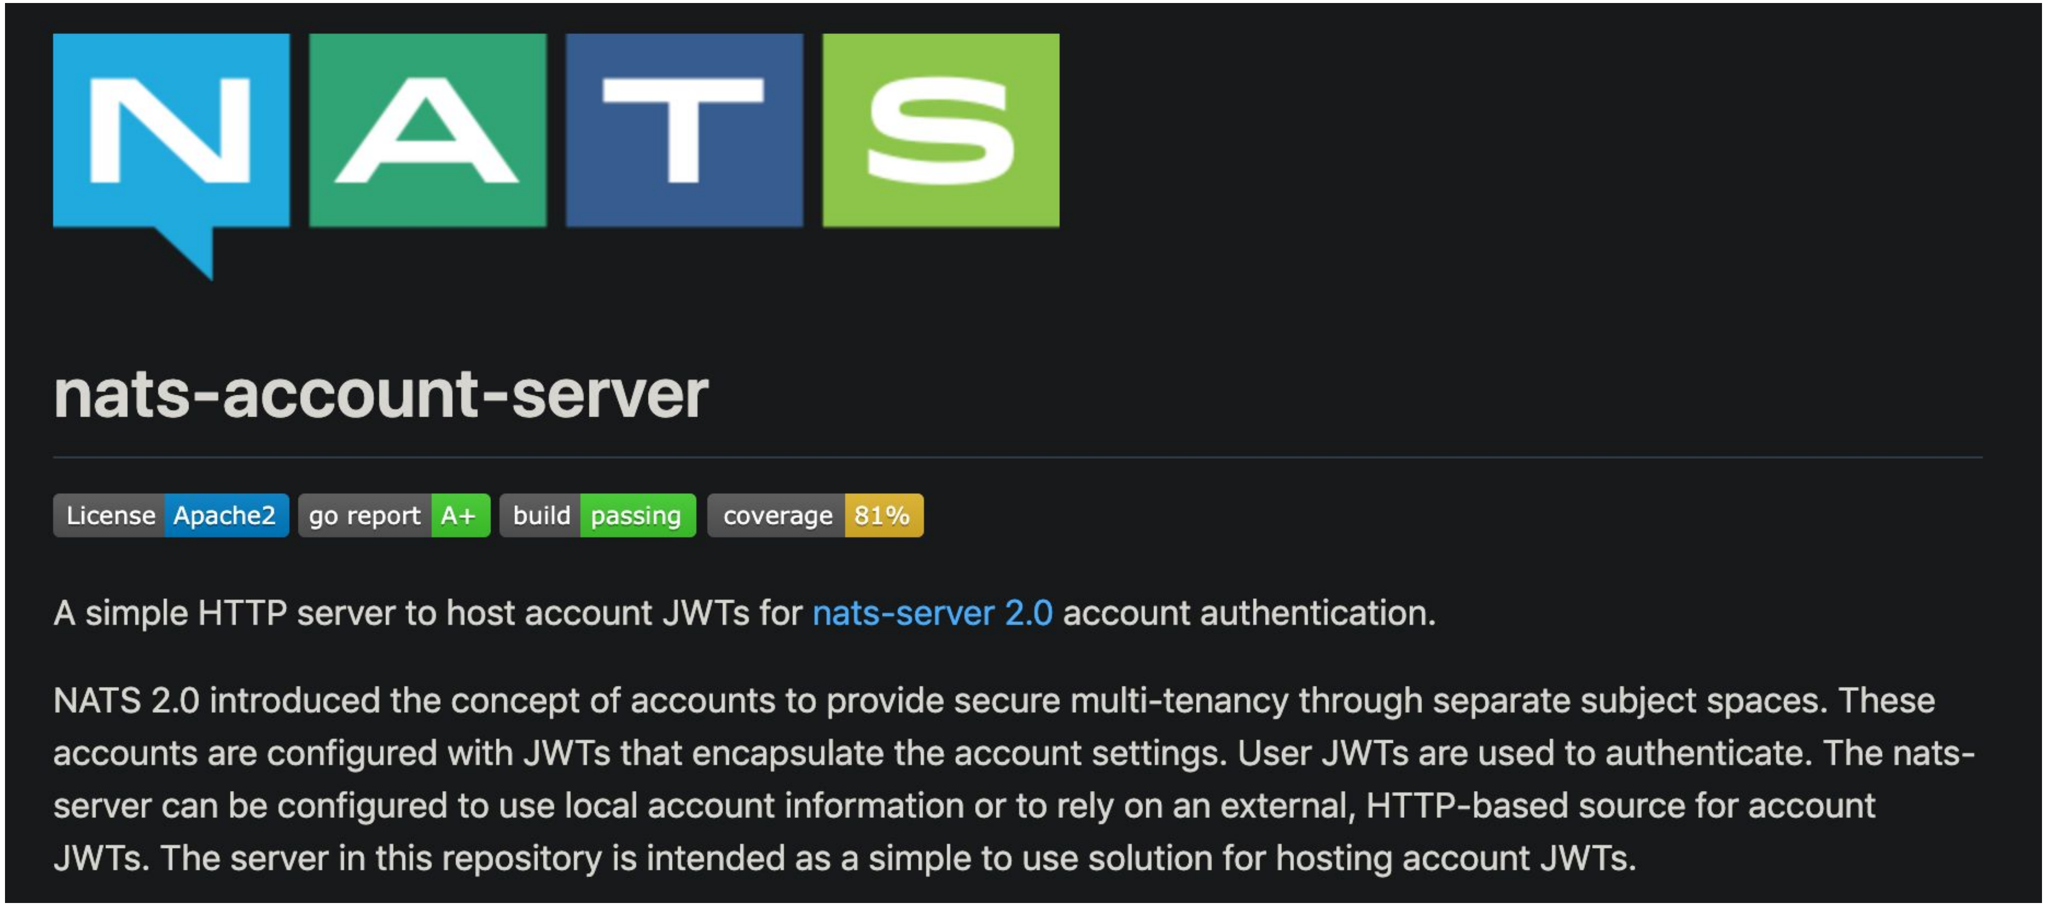

nats-account-server

License Apache2 | go report A+ | build passing | coverage 81%

A simple HTTP server to host account JWTs for nats-server 2.0 account authentication.

NATS 2.0 introduced the concept of accounts to provide secure multi-tenancy through separate subject spaces. These accounts are configured with JWTs that encapsulate the account settings. User JWTs are used to authenticate. The nats-server can be configured to use local account information or to rely on an external, HTTP-based source for account JWTs. The server in this repository is intended as a simple to use solution for hosting account JWTs.

NSC

https://github.com/nats-io/nsc

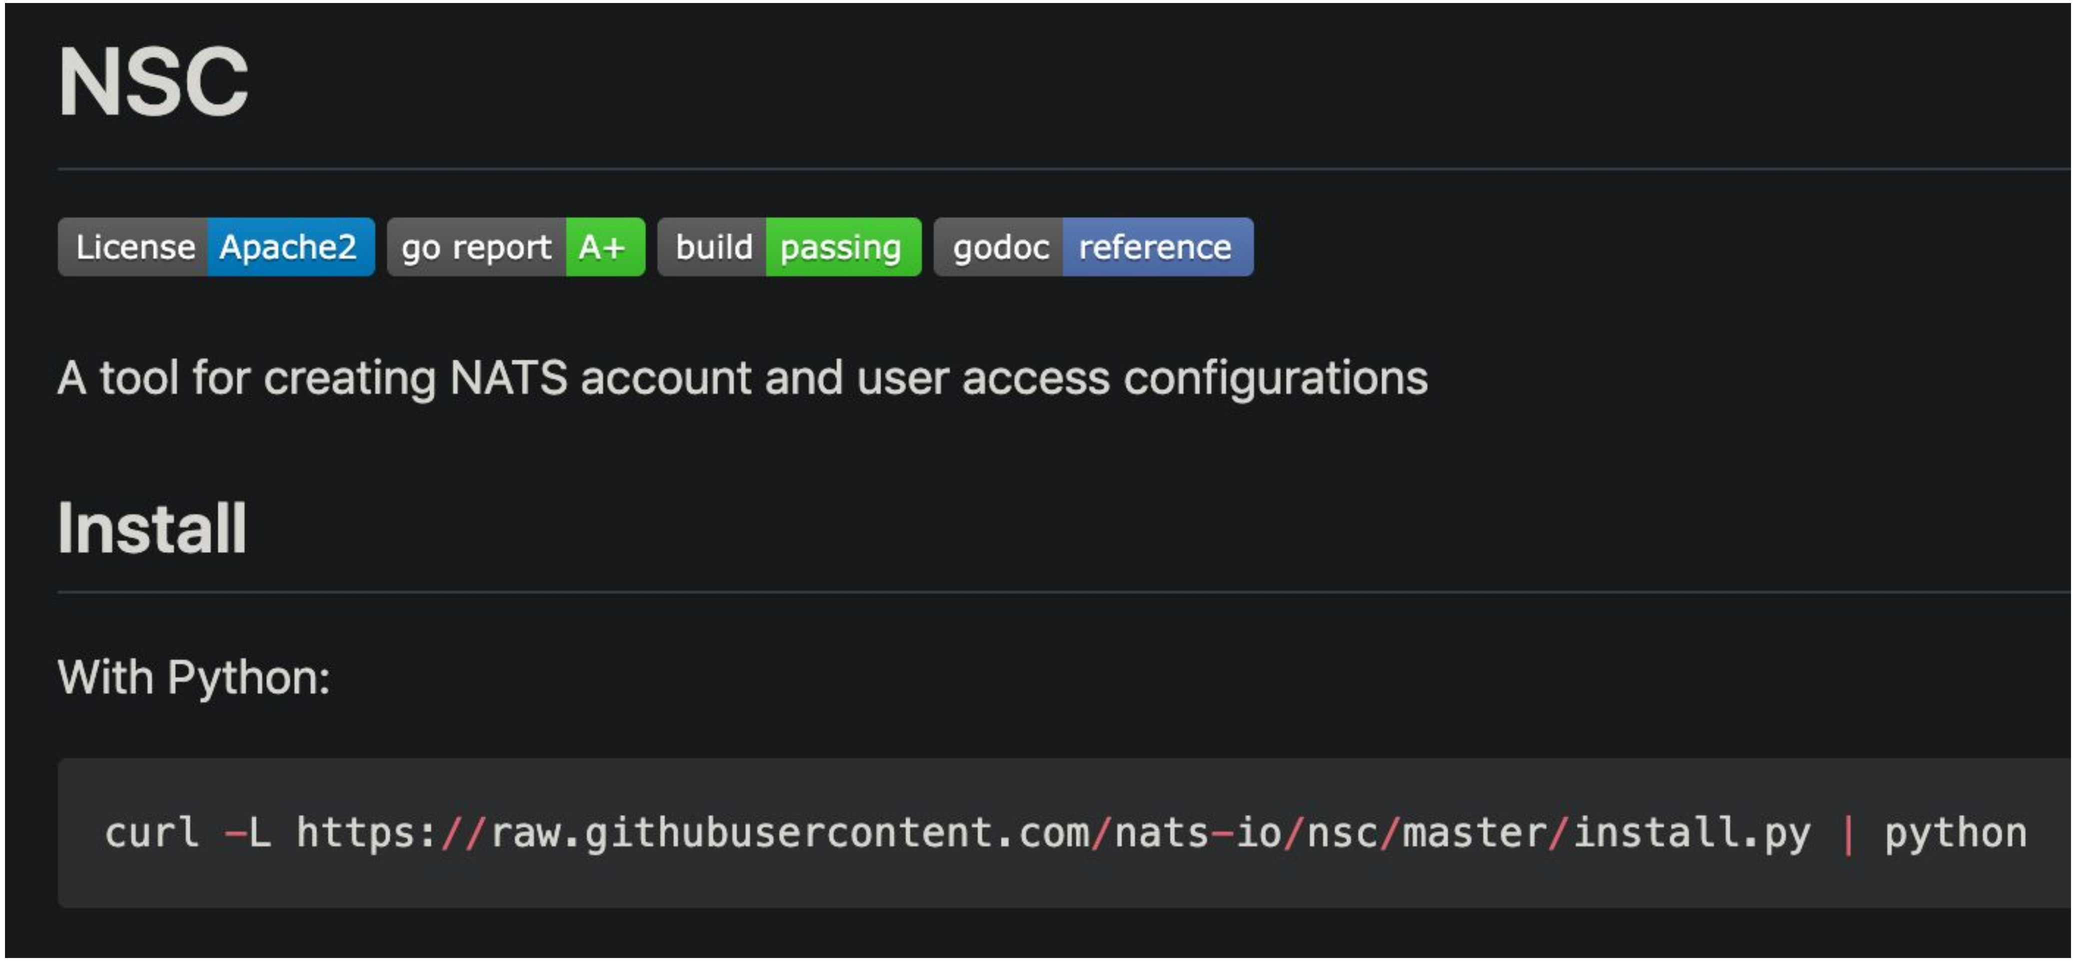

NSC

License Apache2 | go report A+ | build passing | godoc reference

A tool for creating NATS account and user access configurations

Install

With Python:

```
curl -L https://raw.githubusercontent.com/nats-io/nsc/master/install.py | python
```

NKEYS and JWTs

1) Client initiates connection
2) Server sends an INFO with a nonce
3) Client sends CONNECT
   - ✓ Signs the nonce with private nkey seed
   - ✓ Provides public nkey
4) Server verifies
   - ✓ Key
   - ✓ Signature
   - ✓ Nonce

The Server Never stores or even accesses the private key!

NKEYS Connect

```
telnet connect.ngs.global 4222
INFO
{"server_id":"NCX...","version":"2.0.0-RC12","proto":1,"git_commit":"25cd64b","go":"
go1.11.2","host":"XX.XX.XX.XX","port":4222,"auth_required":true,"tls_required":true
,"max_payload":1048576,"client_id":222,"nonce":"NF3ll3XLz9q1j8Y","cluster":"...","c
onnect_urls":["XX.XX.XX.XX:4222","XX.XX.XX.XX:4222","XX.XX.XX.XX:4222"]}
```

NKEYS Connect

```
telnet connect.ngs.global 4222
INFO
{"server_id":"NCX...","version":"2.0.0-RC12","proto":1,"git_commit":"25cd64b","go":"
go1.11.2","host":"XX.XX.XX.XX","port":4222,"auth_required":true,"tls_required":true
,"max_payload":1048576,"client_id":222,"nonce":"NF3ll3XLz9q1j8Y","cluster":"...",
"connect_urls":["XX.XX.XX.XX:4222","XX.XX.XX.XX:4222","XX.XX.XX.XX:4222"]}
```

NKEYS Connect

```
telnet connect.ngs.global 4222
...
CONNECT
{"verbose":false,"pedantic":false,"jwt":"eyJ...","sig":"-t72B9ZyDGJQuOOyIF7R9RowQBuQxk_r9XqFoUs2VwDtjPZorlUTvecL36HIZOT_-cges4g0kteDQIQuVtsSBQ","tls_required":false,"name":"","lang":"go","version":"1.12.0","protocol":1,"echo":true}
```

NKEYS Connect

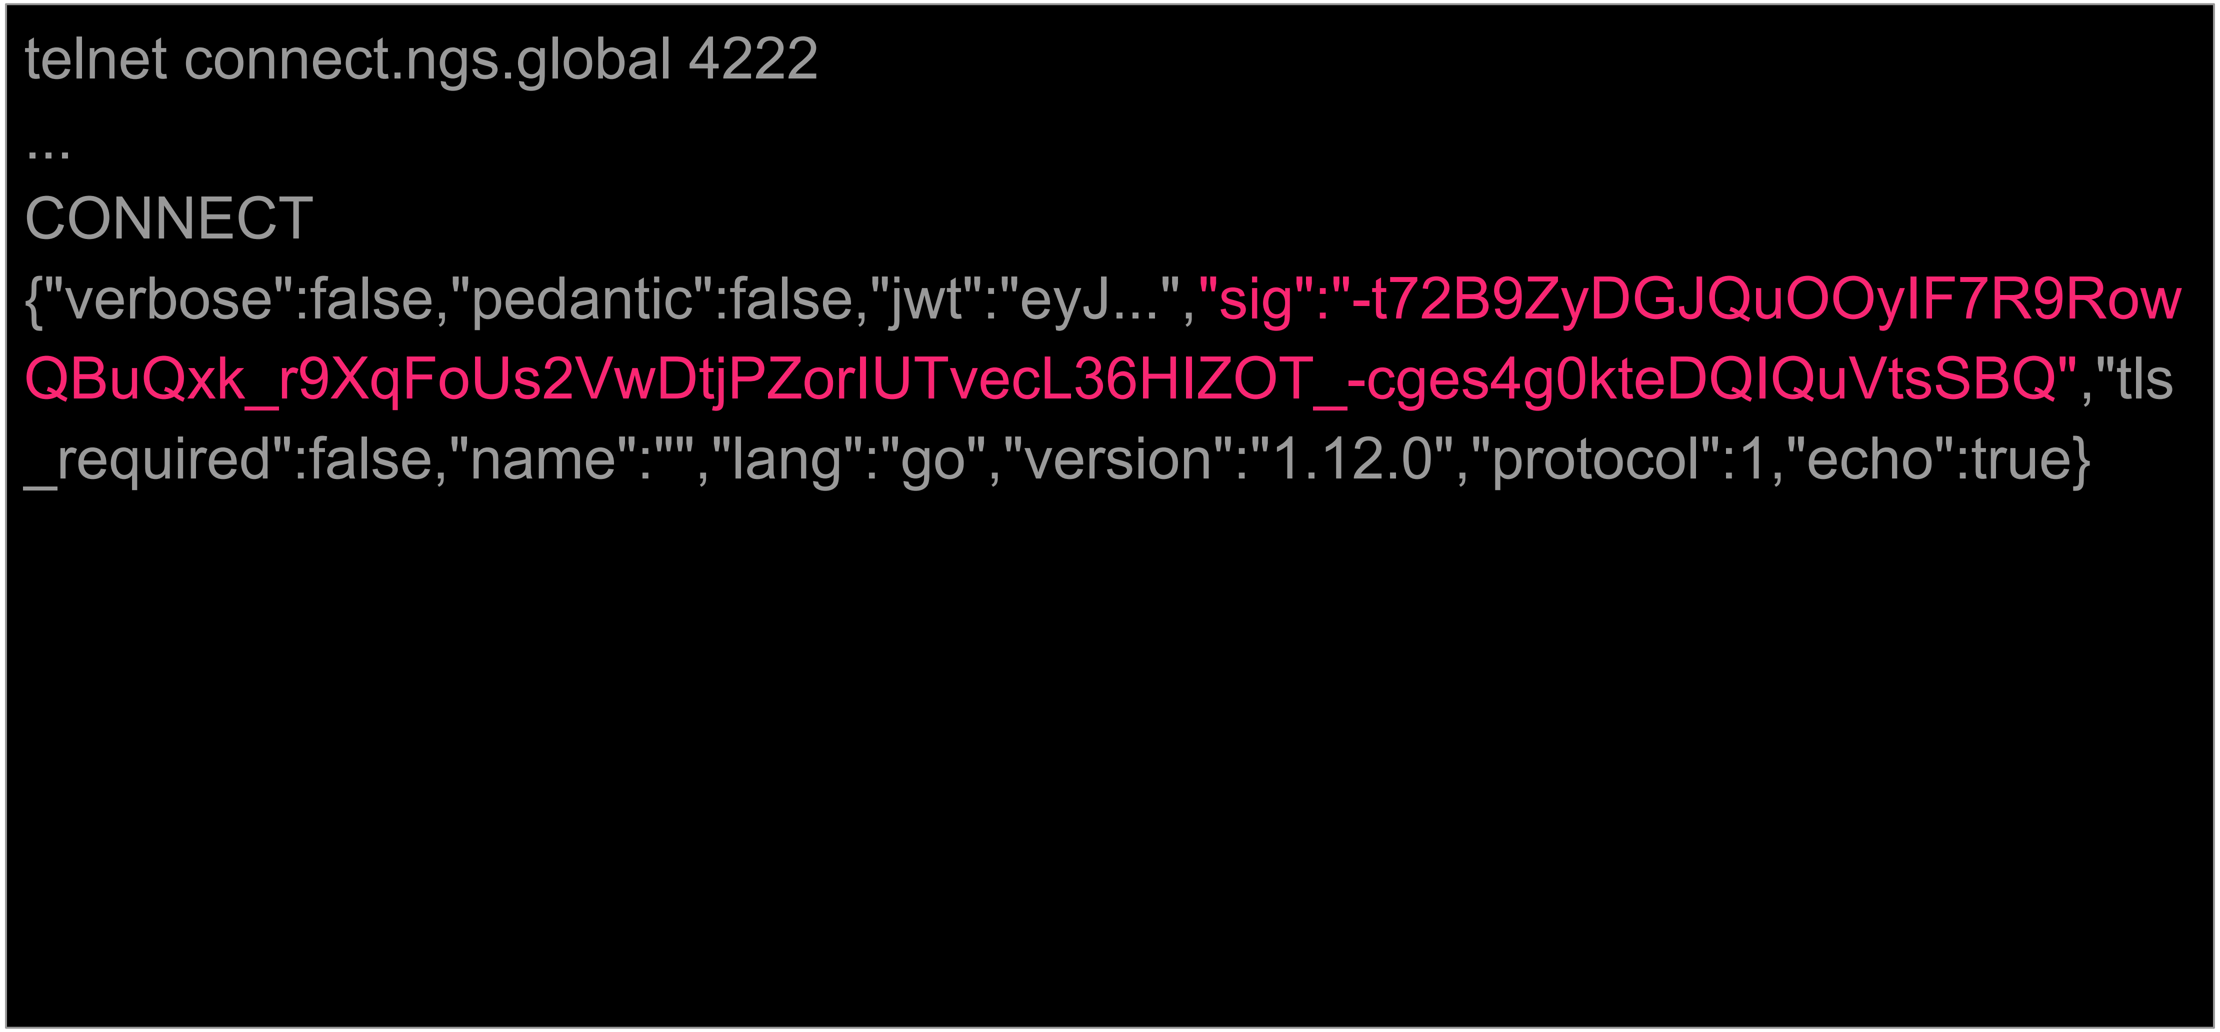

```
telnet connect.ngs.global 4222
...
CONNECT
{"verbose":false,"pedantic":false,"jwt":"eyJ...","sig":"-t72B9ZyDGJQuOOyIF7R9Row
QBuQxk_r9XqFoUs2VwDtjPZorlUTvecL36HIZOT_-cges4g0kteDQIQuVtsSBQ","tls
_required":false,"name":"","lang":"go","version":"1.12.0","protocol":1,"echo":true}
```

Super Clusters

Gateways + Leaf nodes

NATS Super Cluster

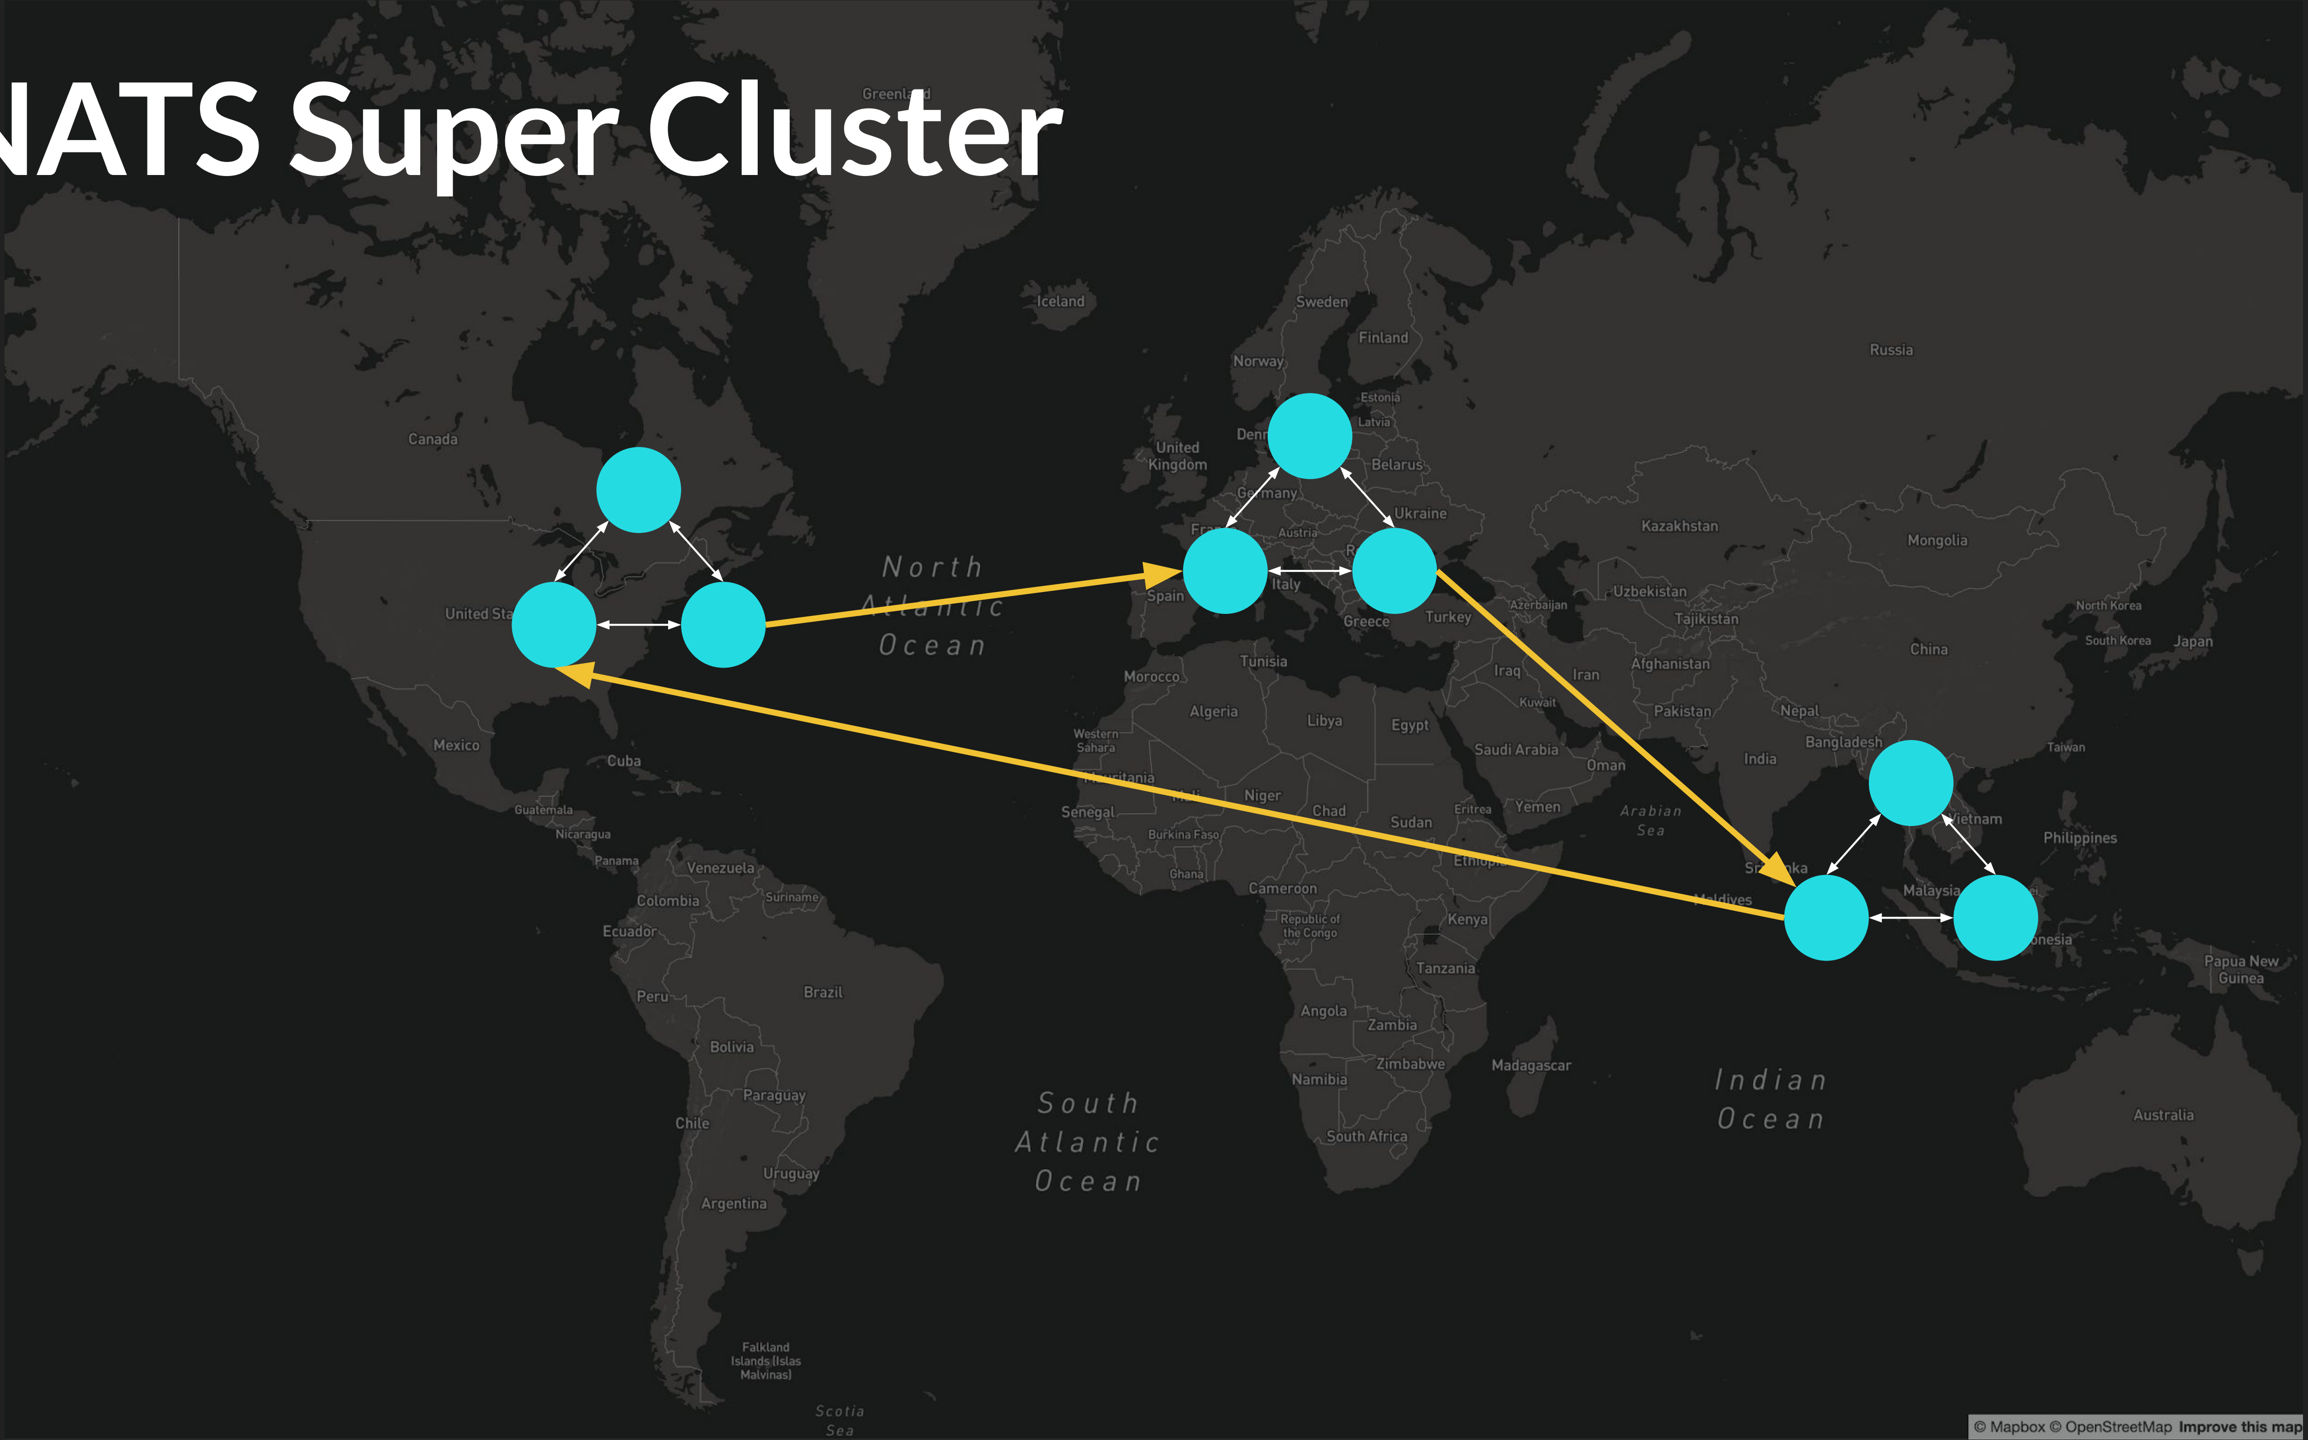

NATS Super Clusters

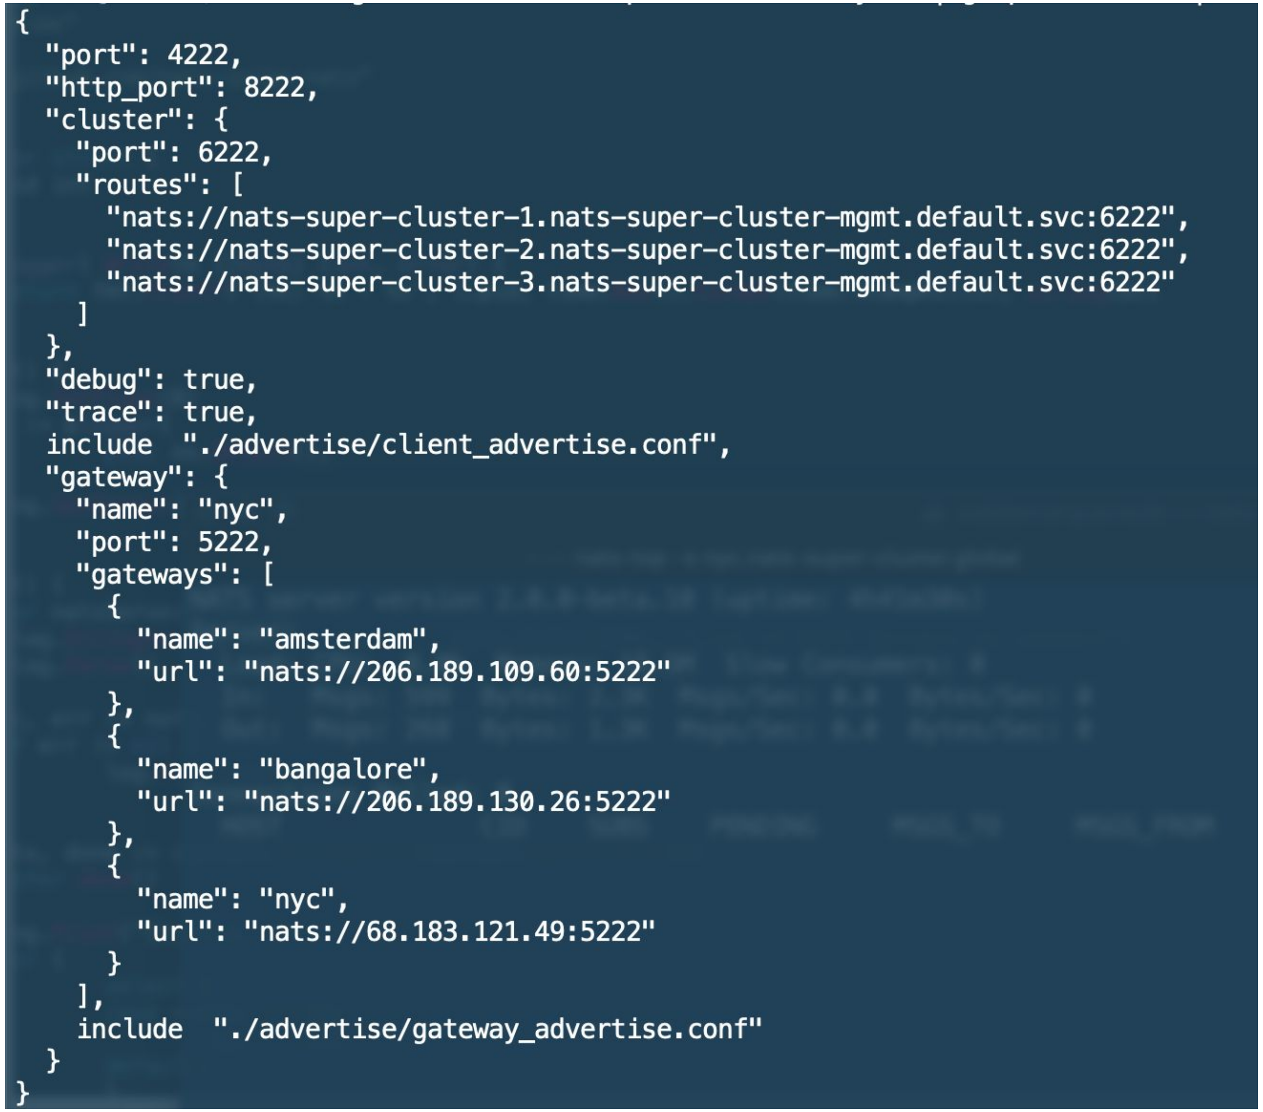

```
{
  "port": 4222,
  "http_port": 8222,
  "cluster": {
    "port": 6222,
    "routes": [
      "nats://nats-super-cluster-1.nats-super-cluster-mgmt.default.svc:6222",
      "nats://nats-super-cluster-2.nats-super-cluster-mgmt.default.svc:6222",
      "nats://nats-super-cluster-3.nats-super-cluster-mgmt.default.svc:6222"
    ]
  },
  "debug": true,
  "trace": true,
  include  "./advertise/client_advertise.conf",
  "gateway": {
    "name": "nyc",
    "port": 5222,
    "gateways": [
      {
        "name": "amsterdam",
        "url": "nats://206.189.109.60:5222"
      },
      {
        "name": "bangalore",
        "url": "nats://206.189.130.26:5222"
      },
      {
        "name": "nyc",
        "url": "nats://68.183.121.49:5222"
      }
    ],
    include  "./advertise/gateway_advertise.conf"
  }
}
```

NATS Super Cluster

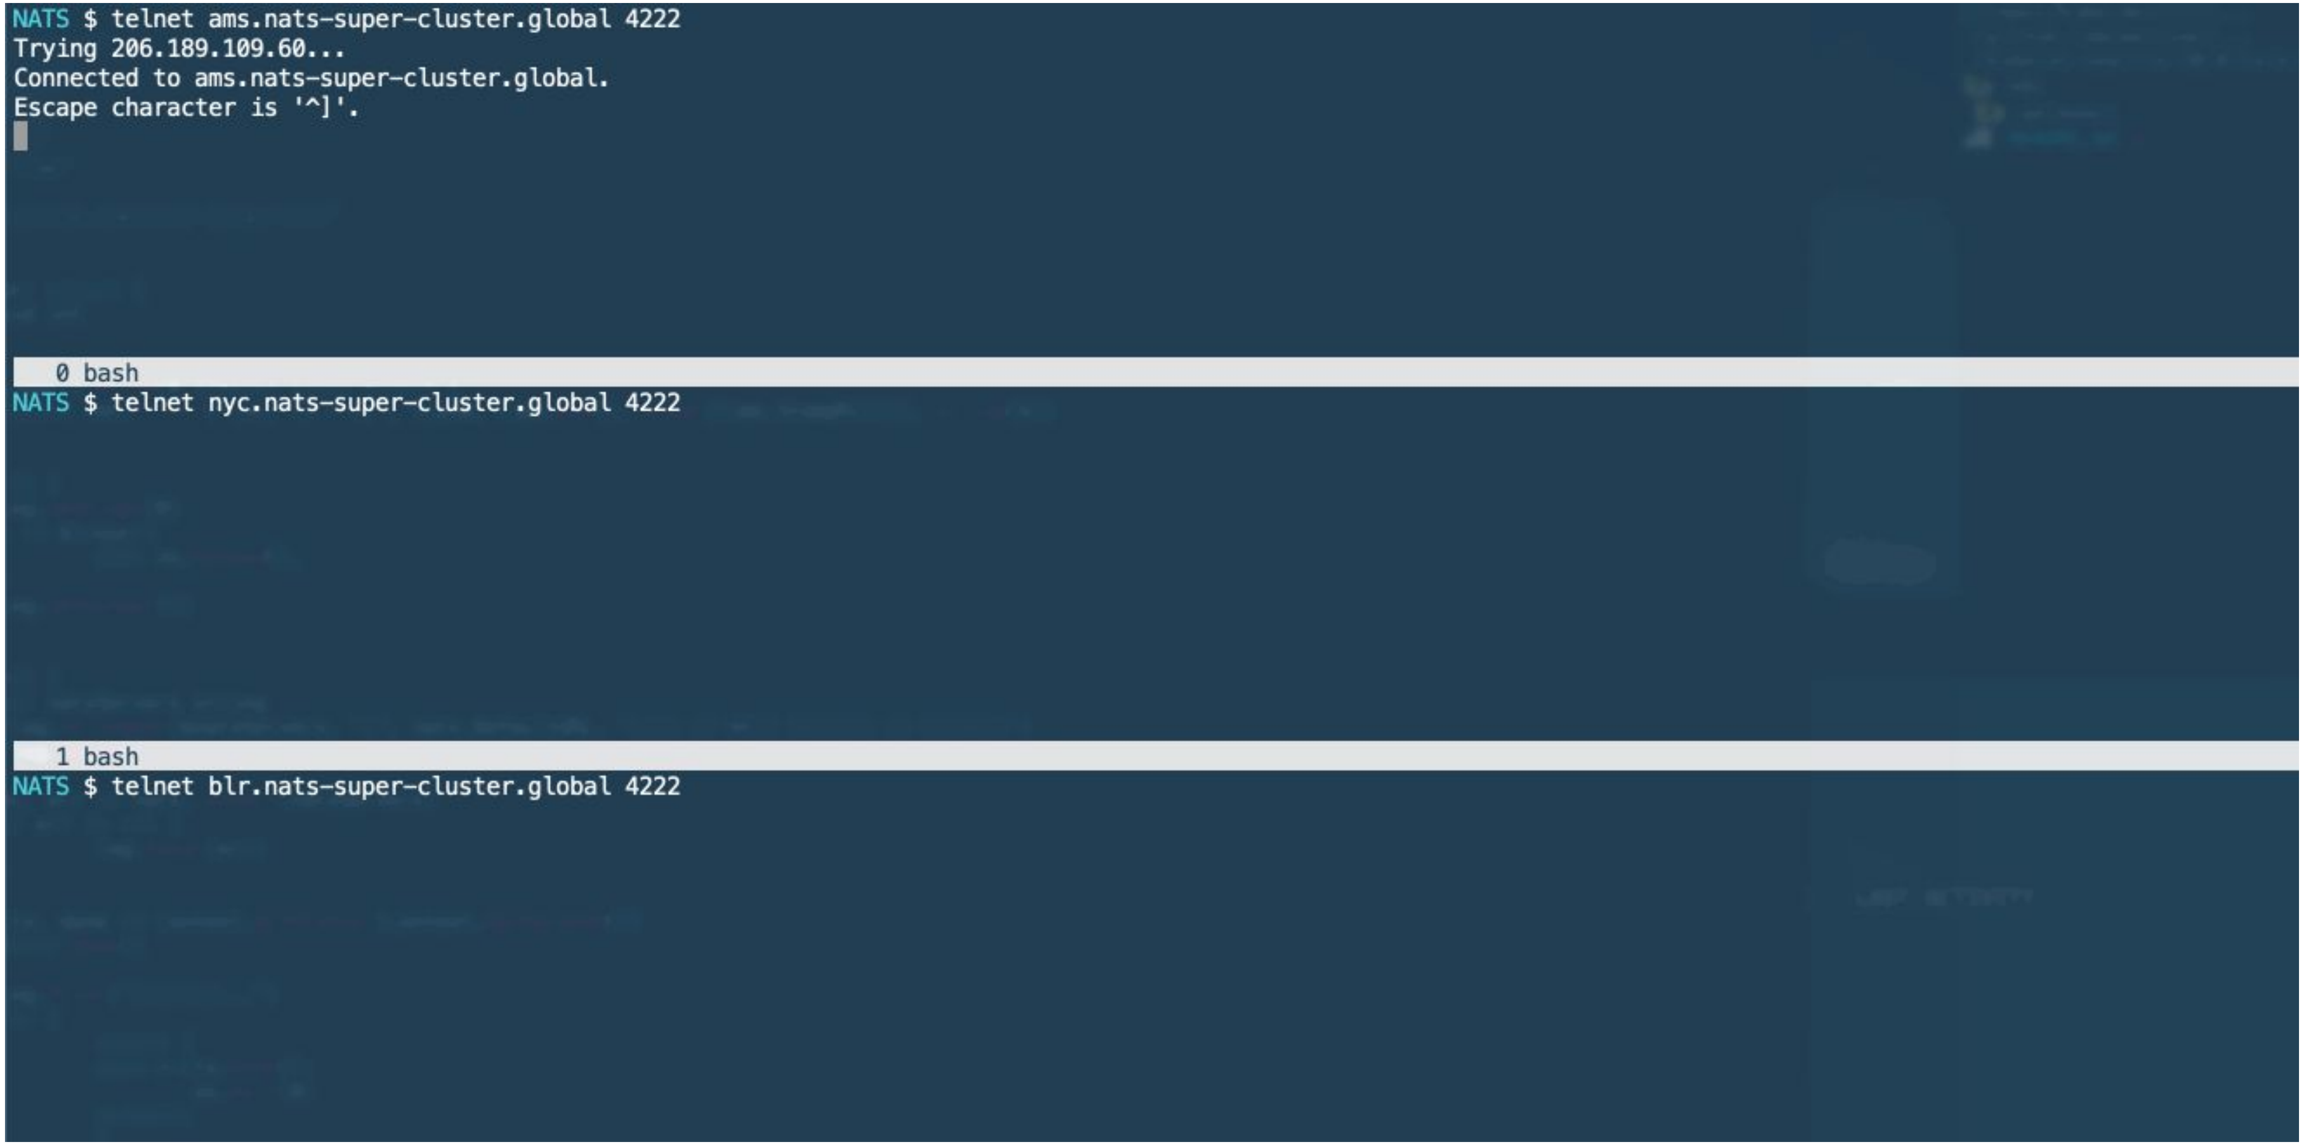

```
NATS $ telnet ams.nats-super-cluster.global 4222
Trying 206.189.109.60...
Connected to ams.nats-super-cluster.global.
Escape character is '^]'.
```

```
  0 bash
NATS $ telnet nyc.nats-super-cluster.global 4222
```

```
  1 bash
NATS $ telnet blr.nats-super-cluster.global 4222
```

Geo-Aware Queue Subscribers

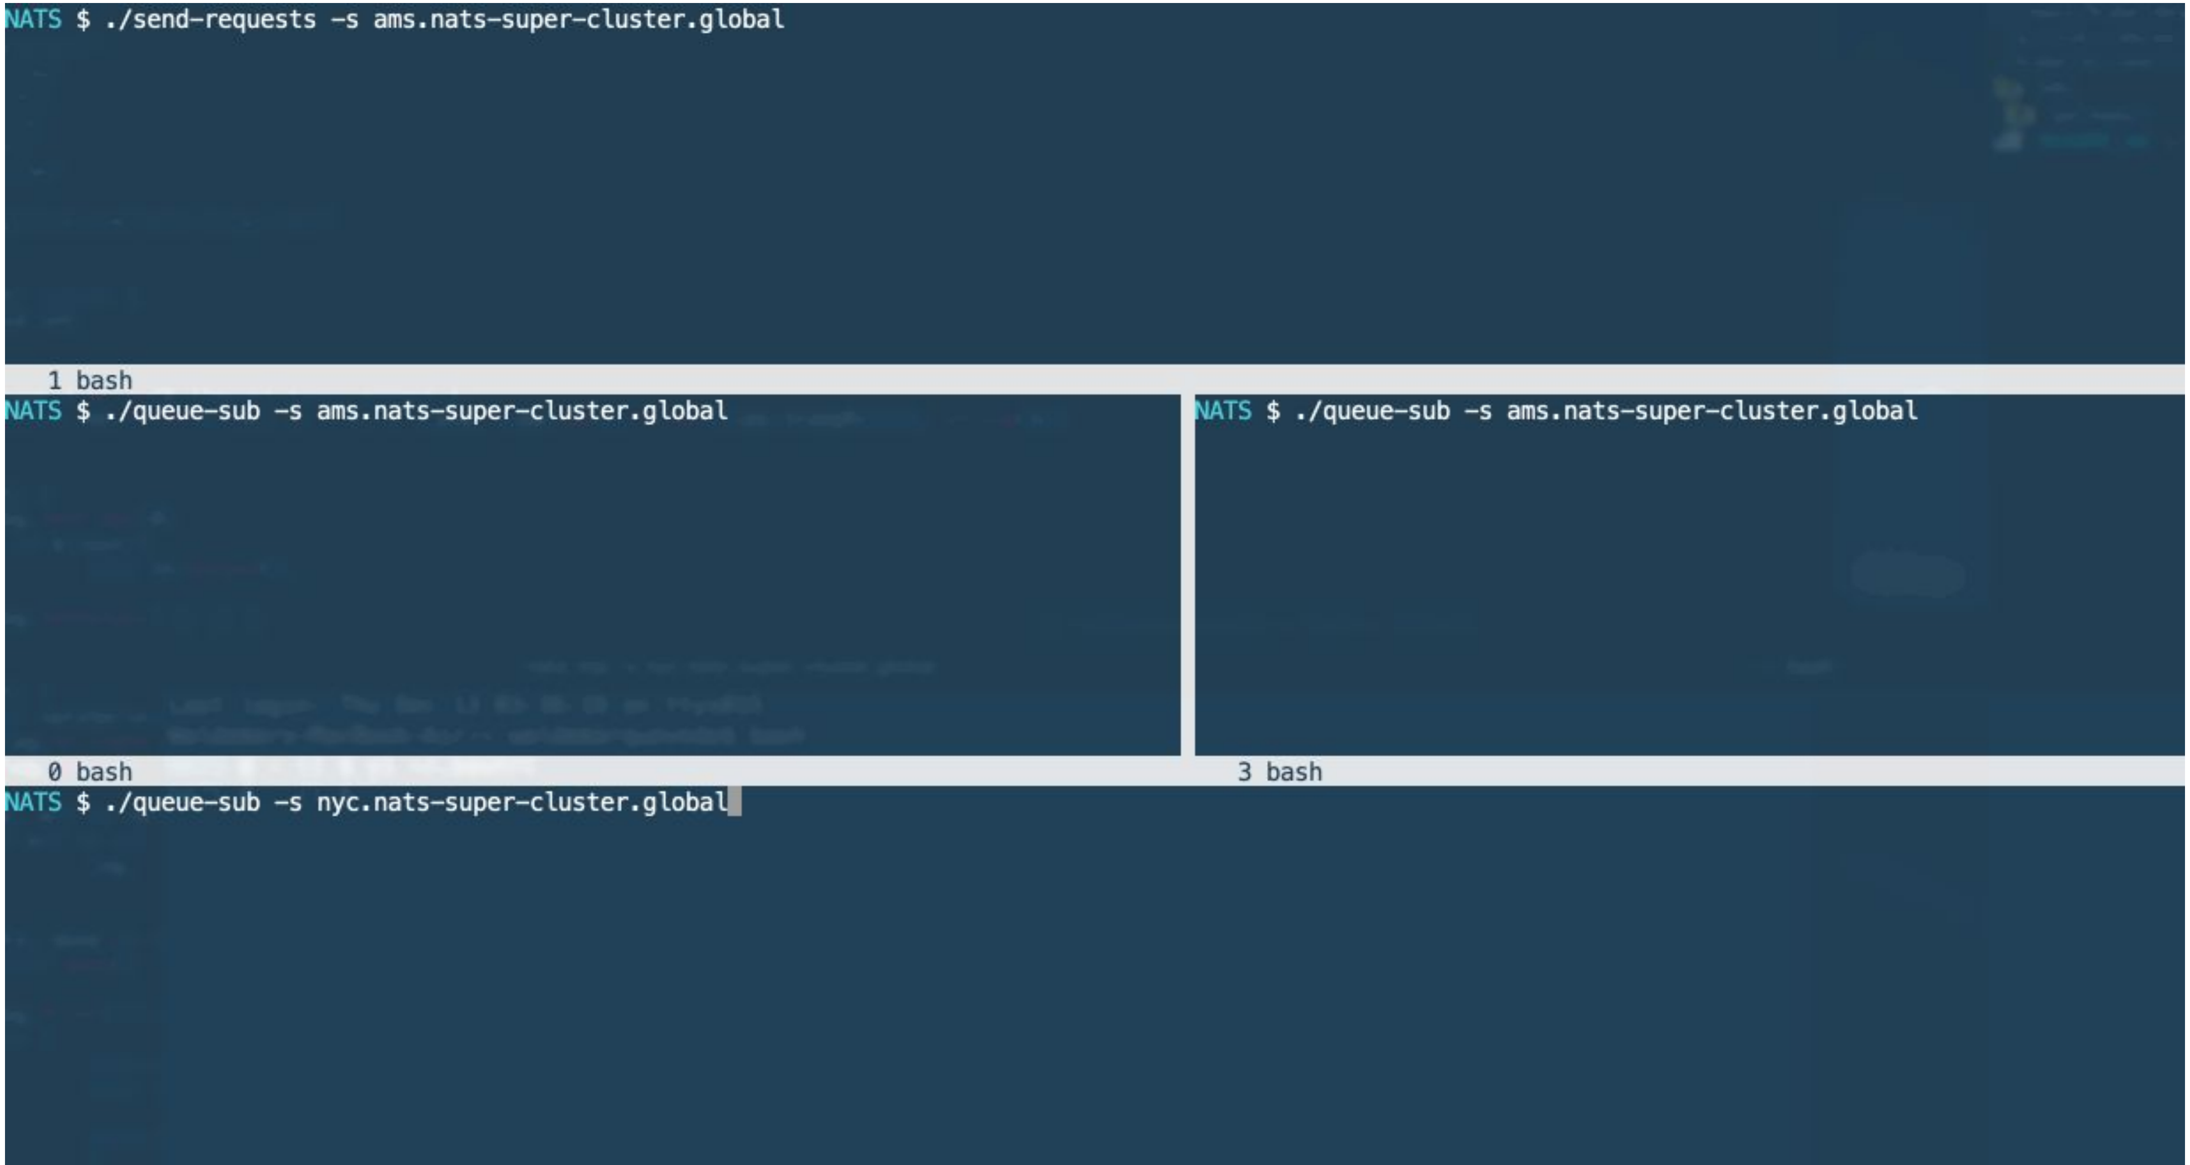

```
NATS $ ./send-requests -s ams.nats-super-cluster.global
```

```
  1 bash
NATS $ ./queue-sub -s ams.nats-super-cluster.global
```

```
NATS $ ./queue-sub -s ams.nats-super-cluster.global
```

```
  0 bash                                                  3 bash
NATS $ ./queue-sub -s nyc.nats-super-cluster.global
```

NATS Super Cluster

At most 3 hops

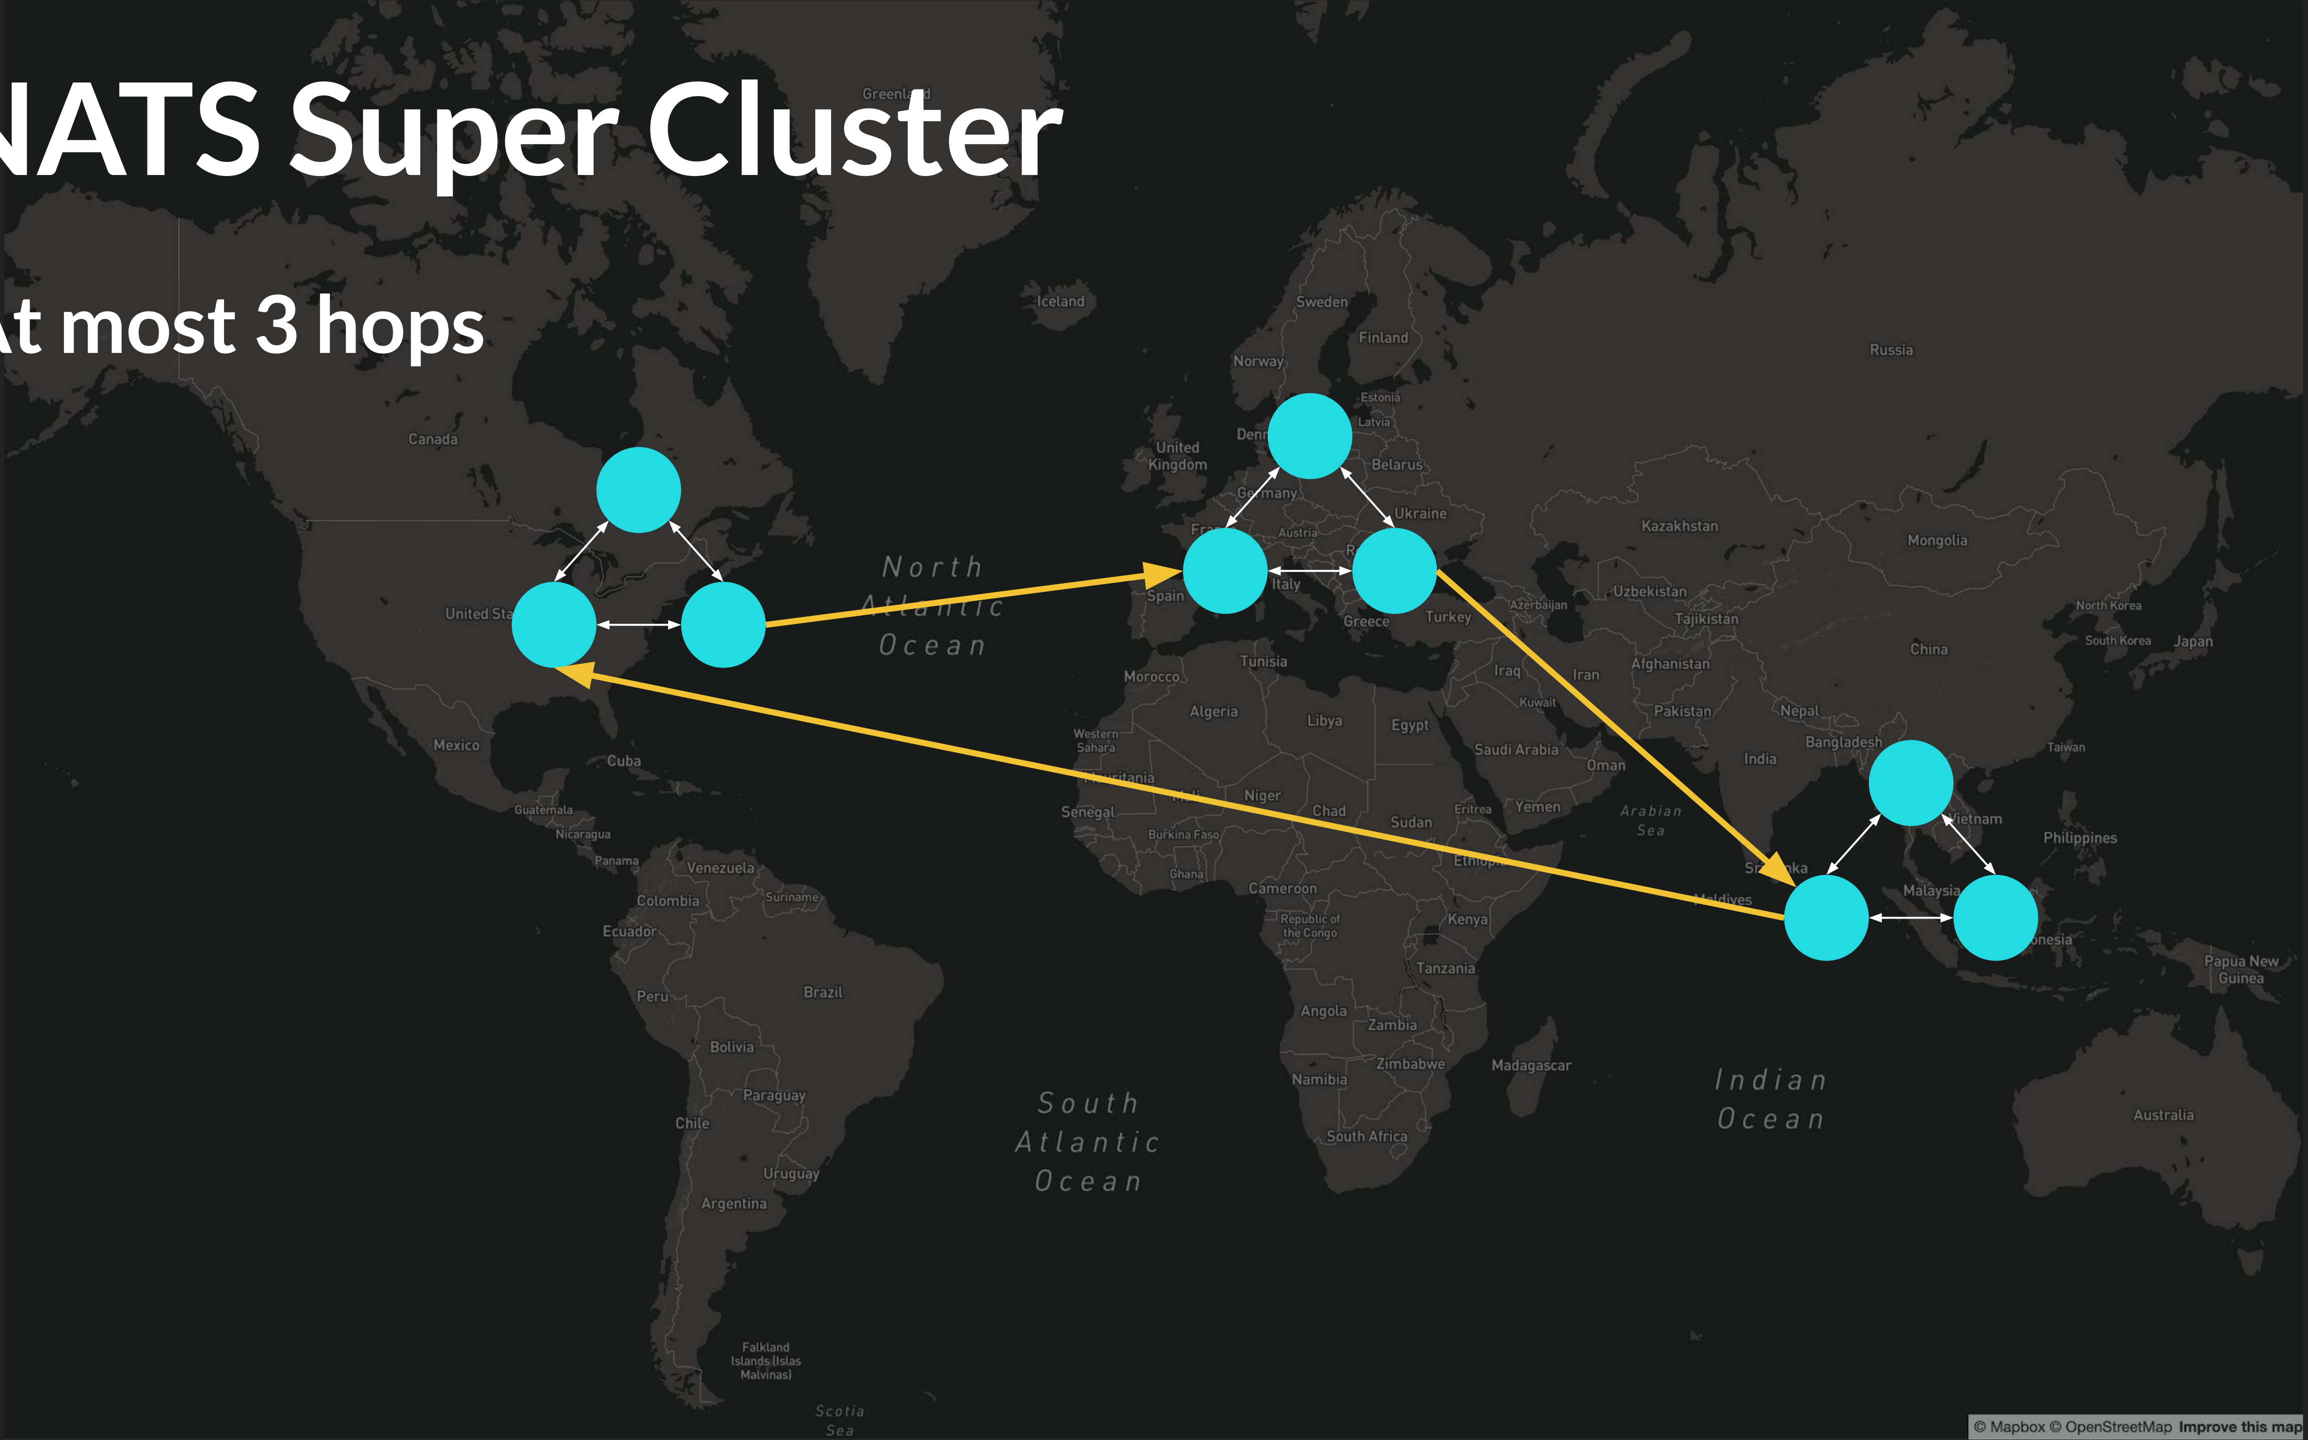

NATS Super Cluster

At most 3 hops

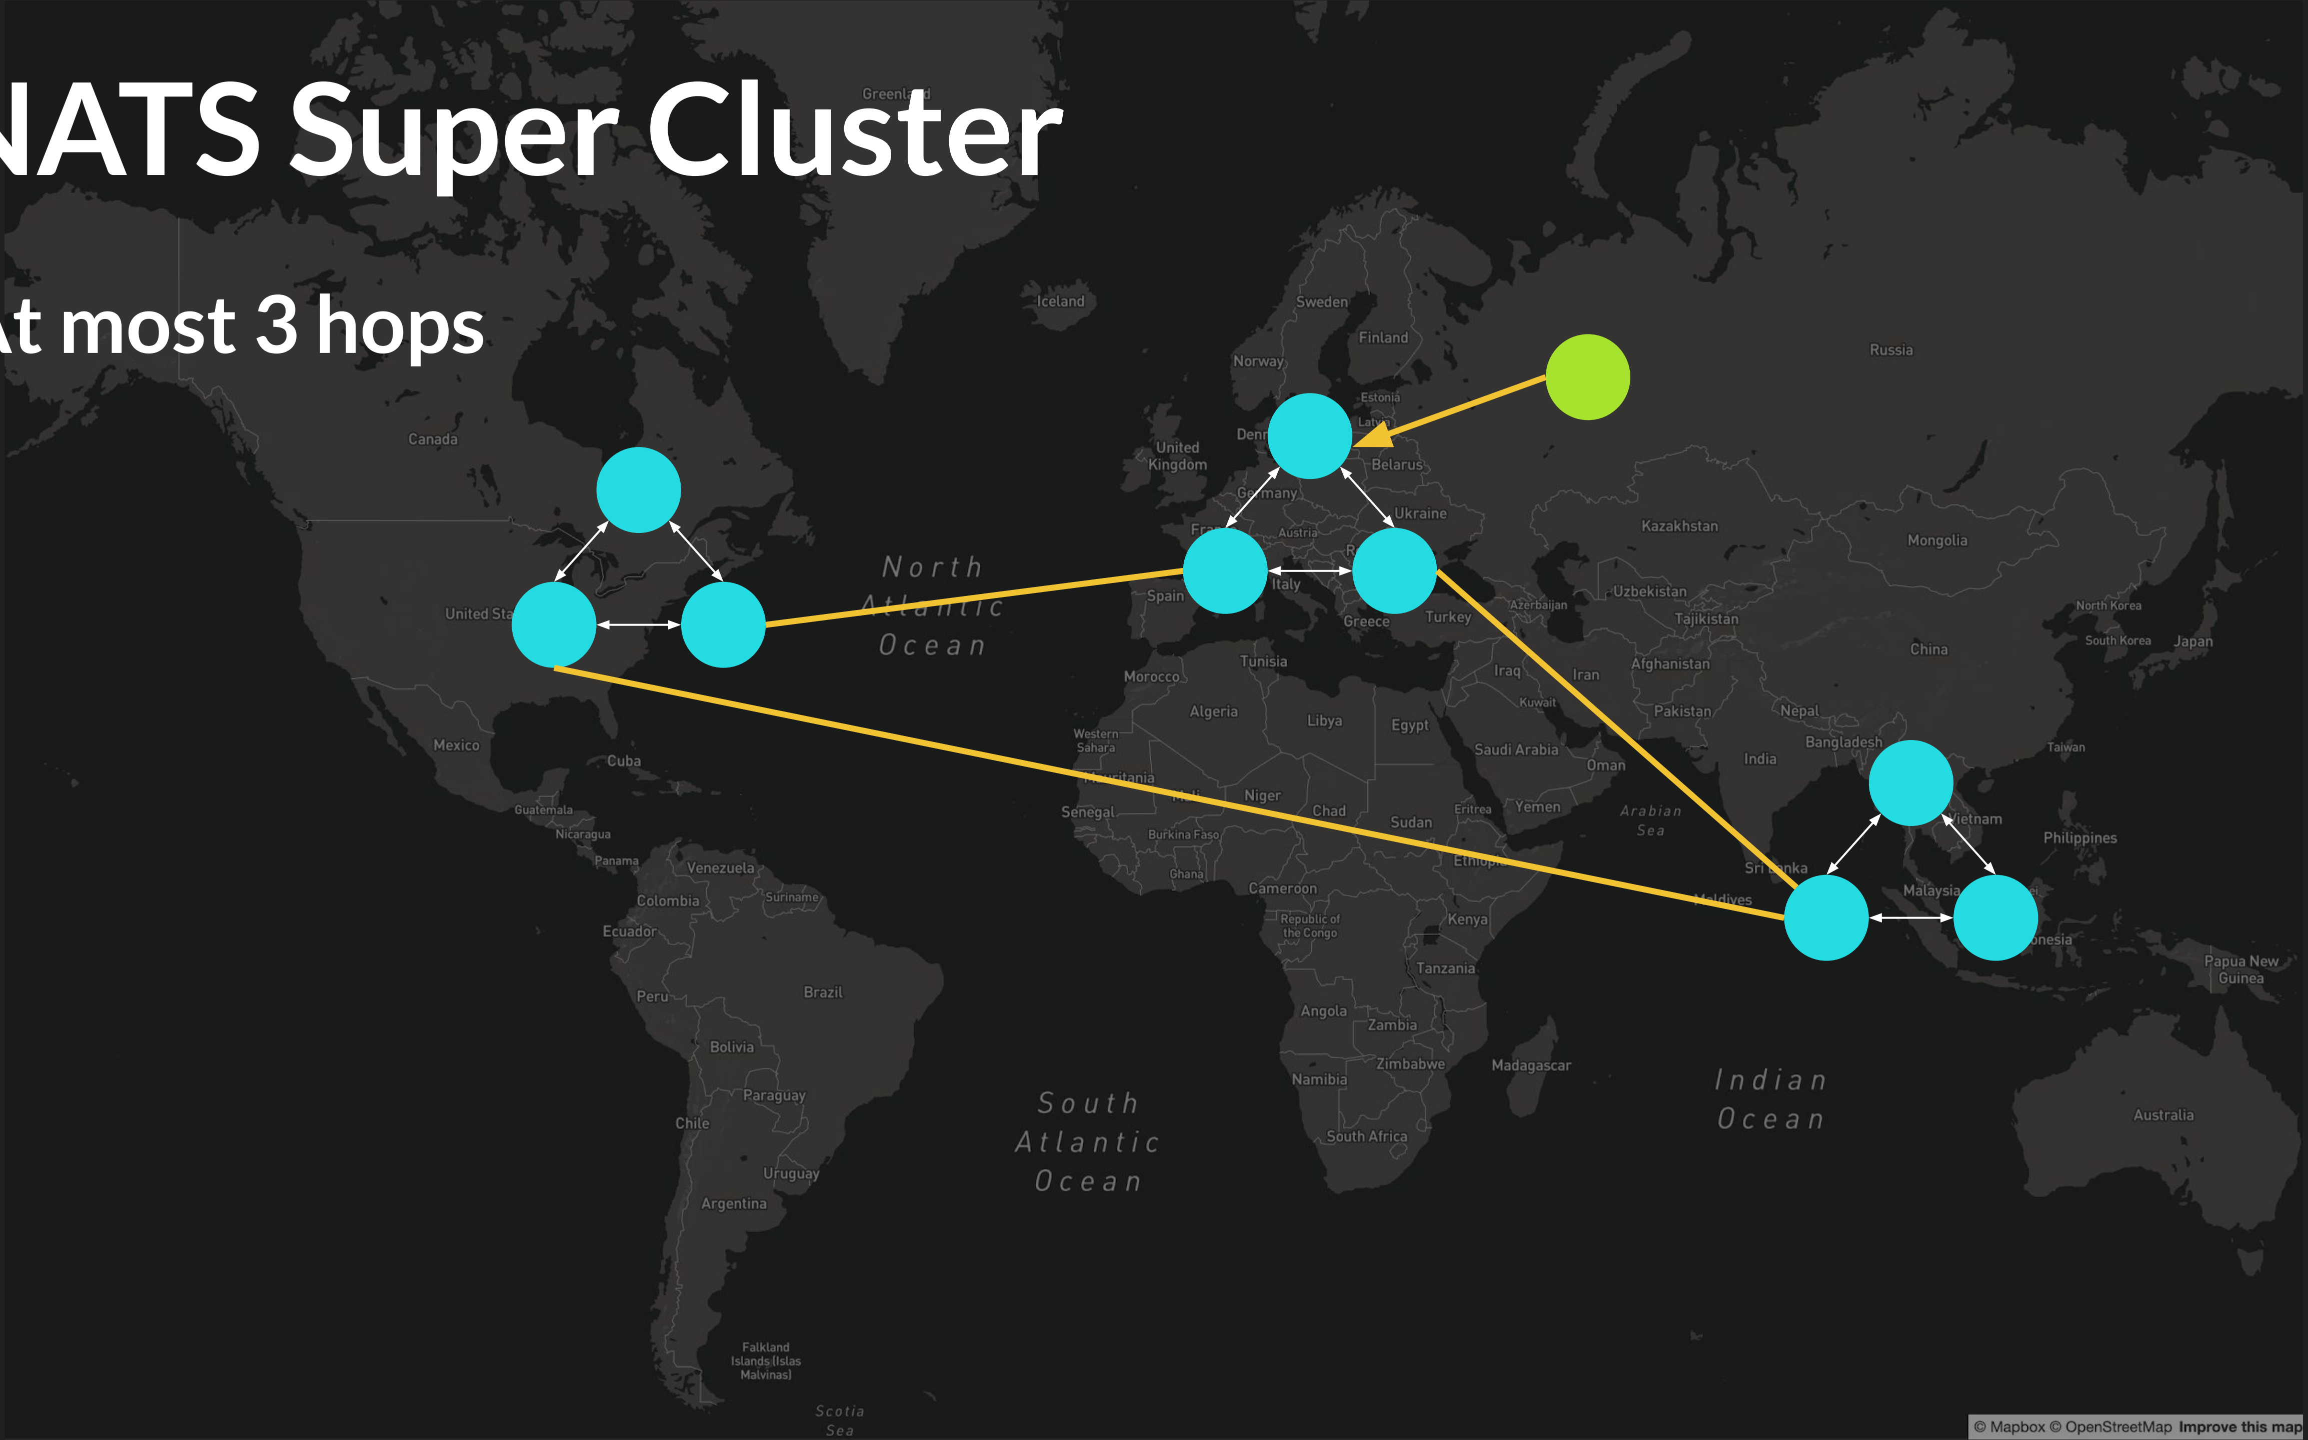

NATS Super Cluster

At most 3 hops

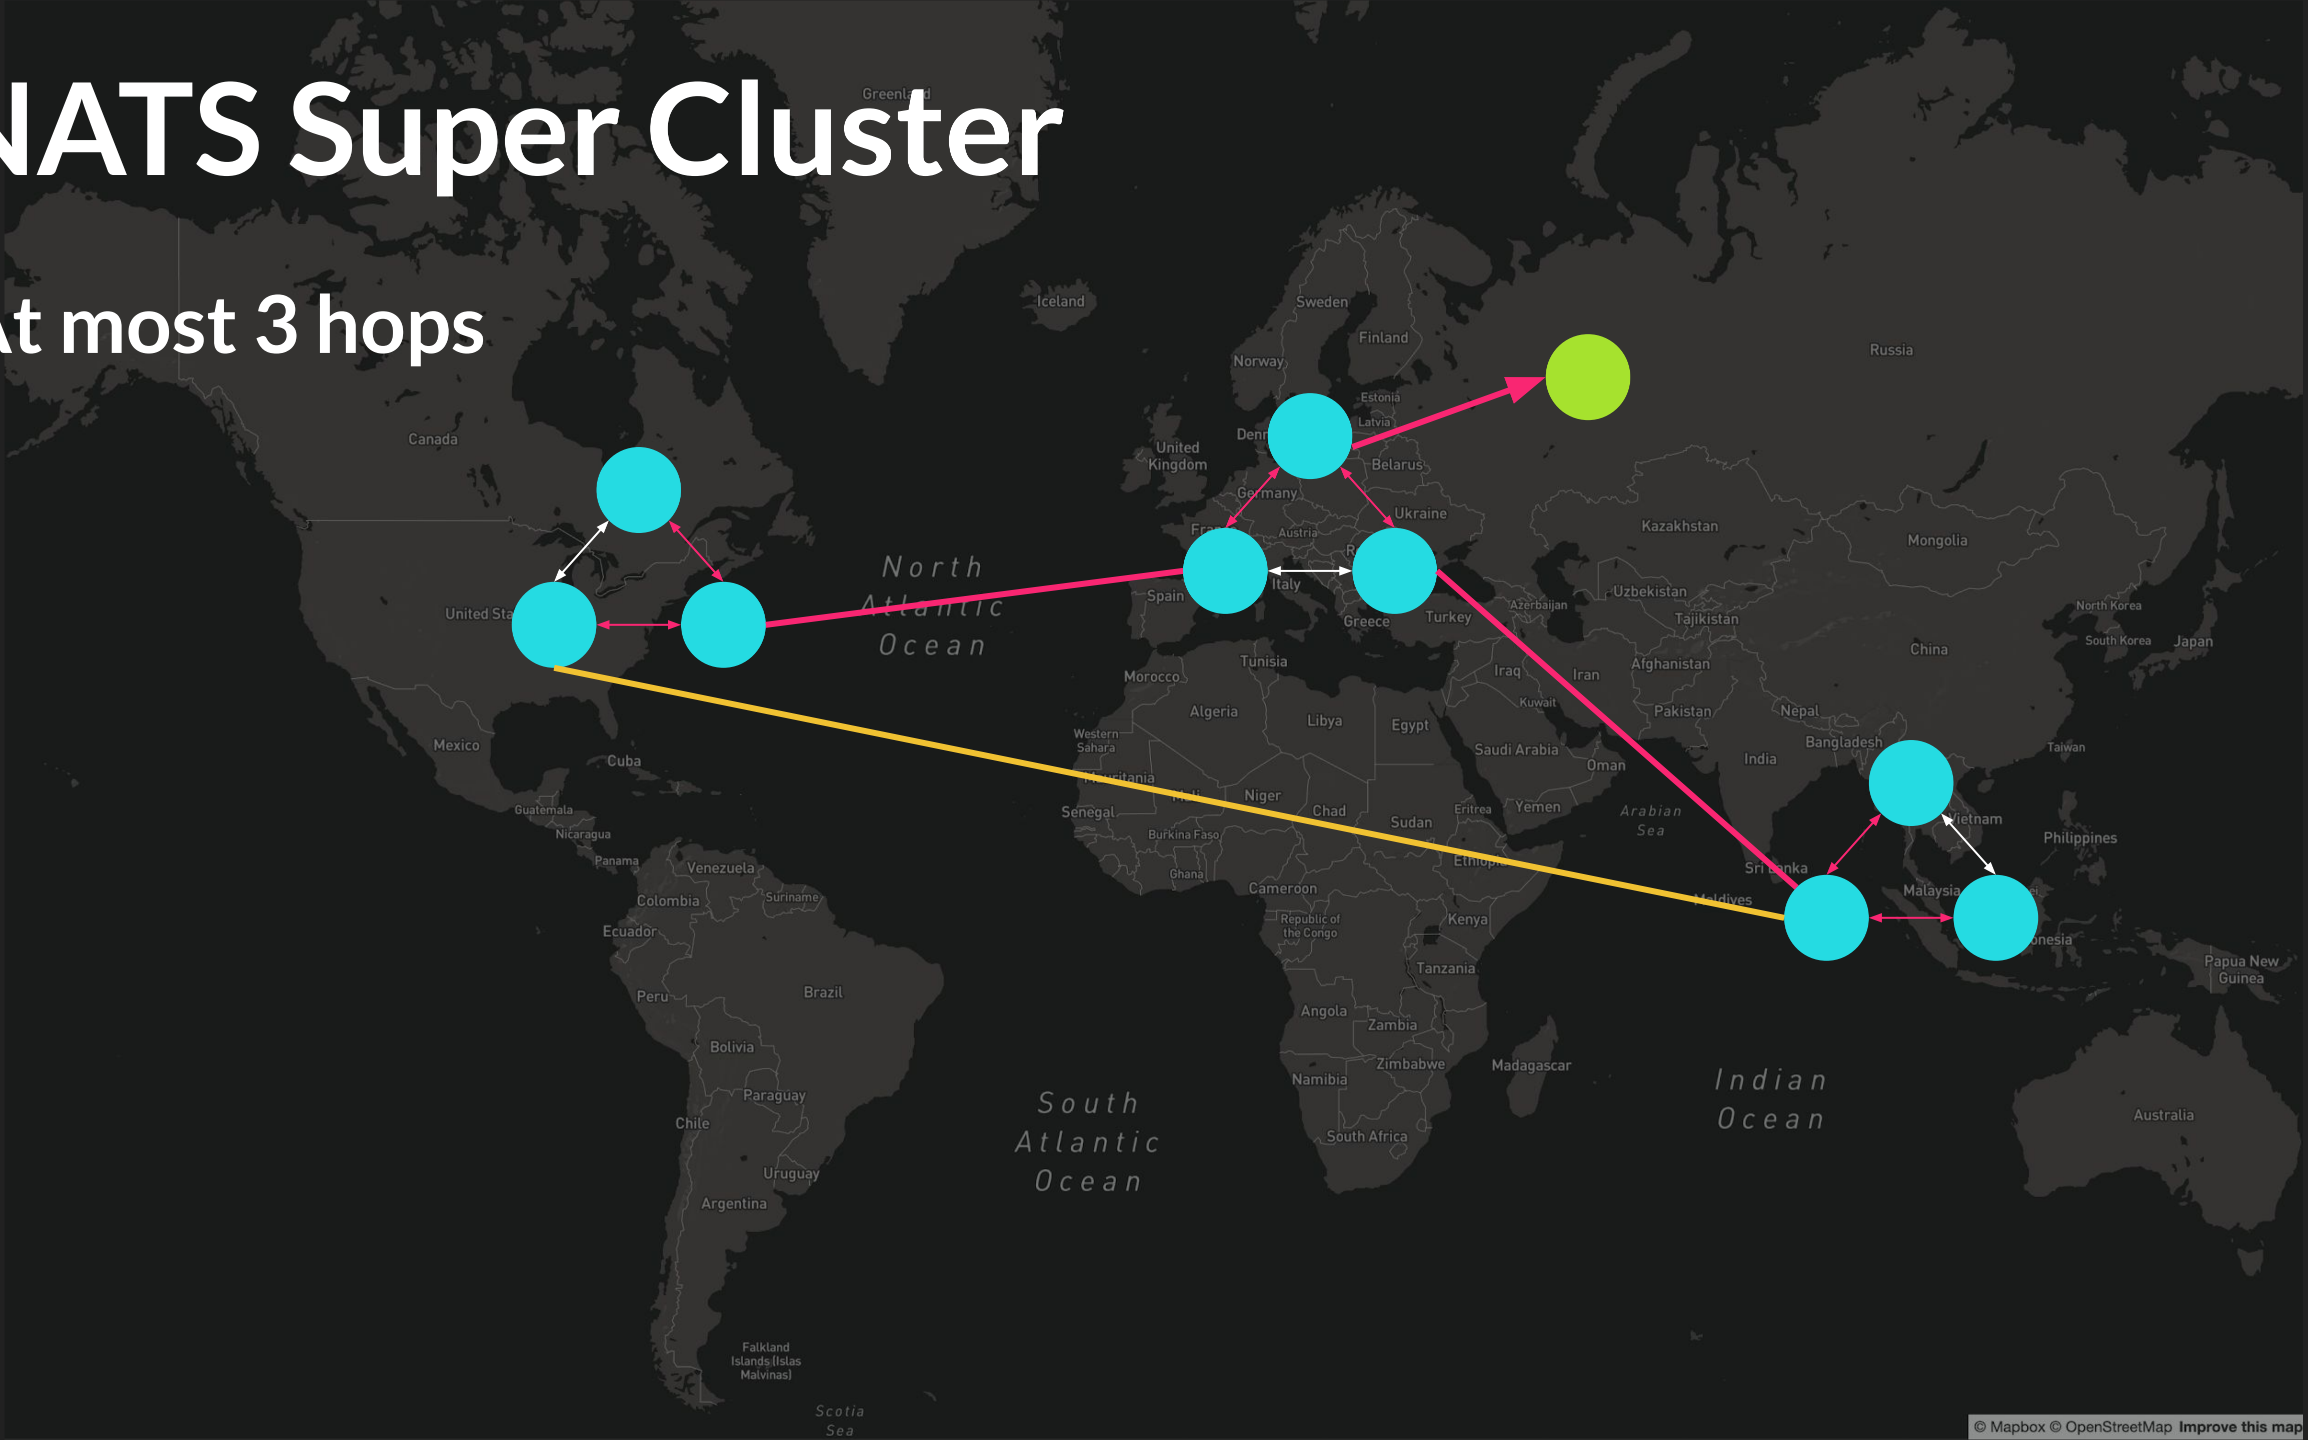

NATS Super Cluster + Leafnodes

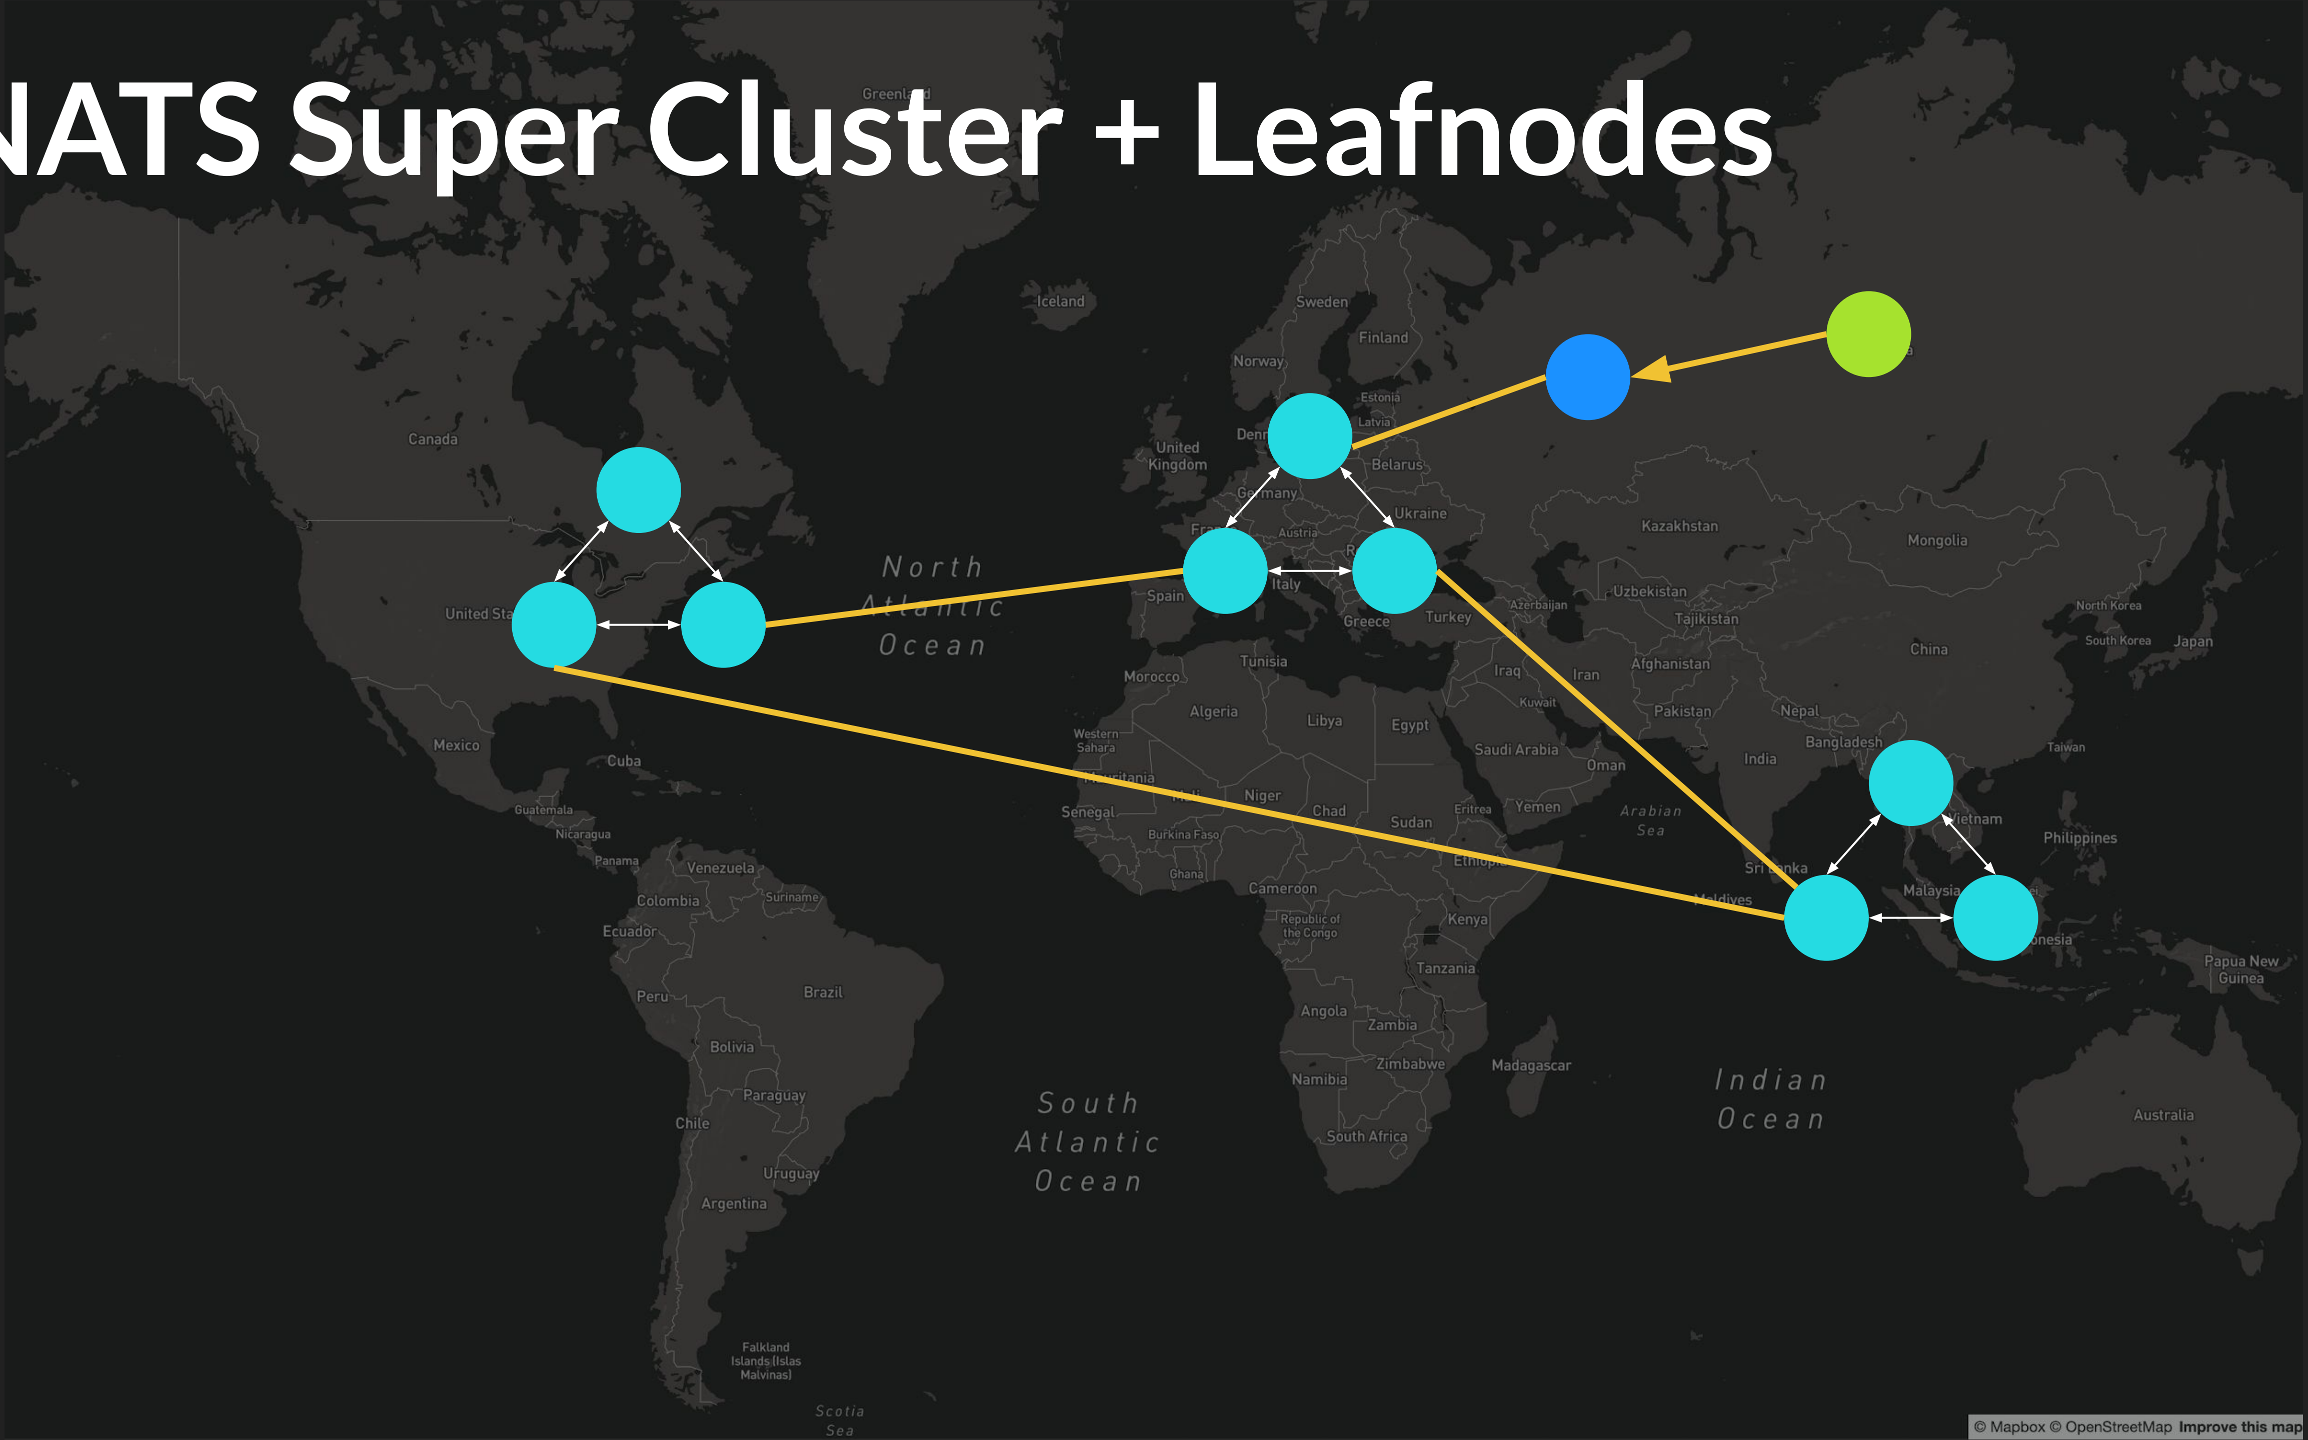

NATS Super Cluster + Leafnodes

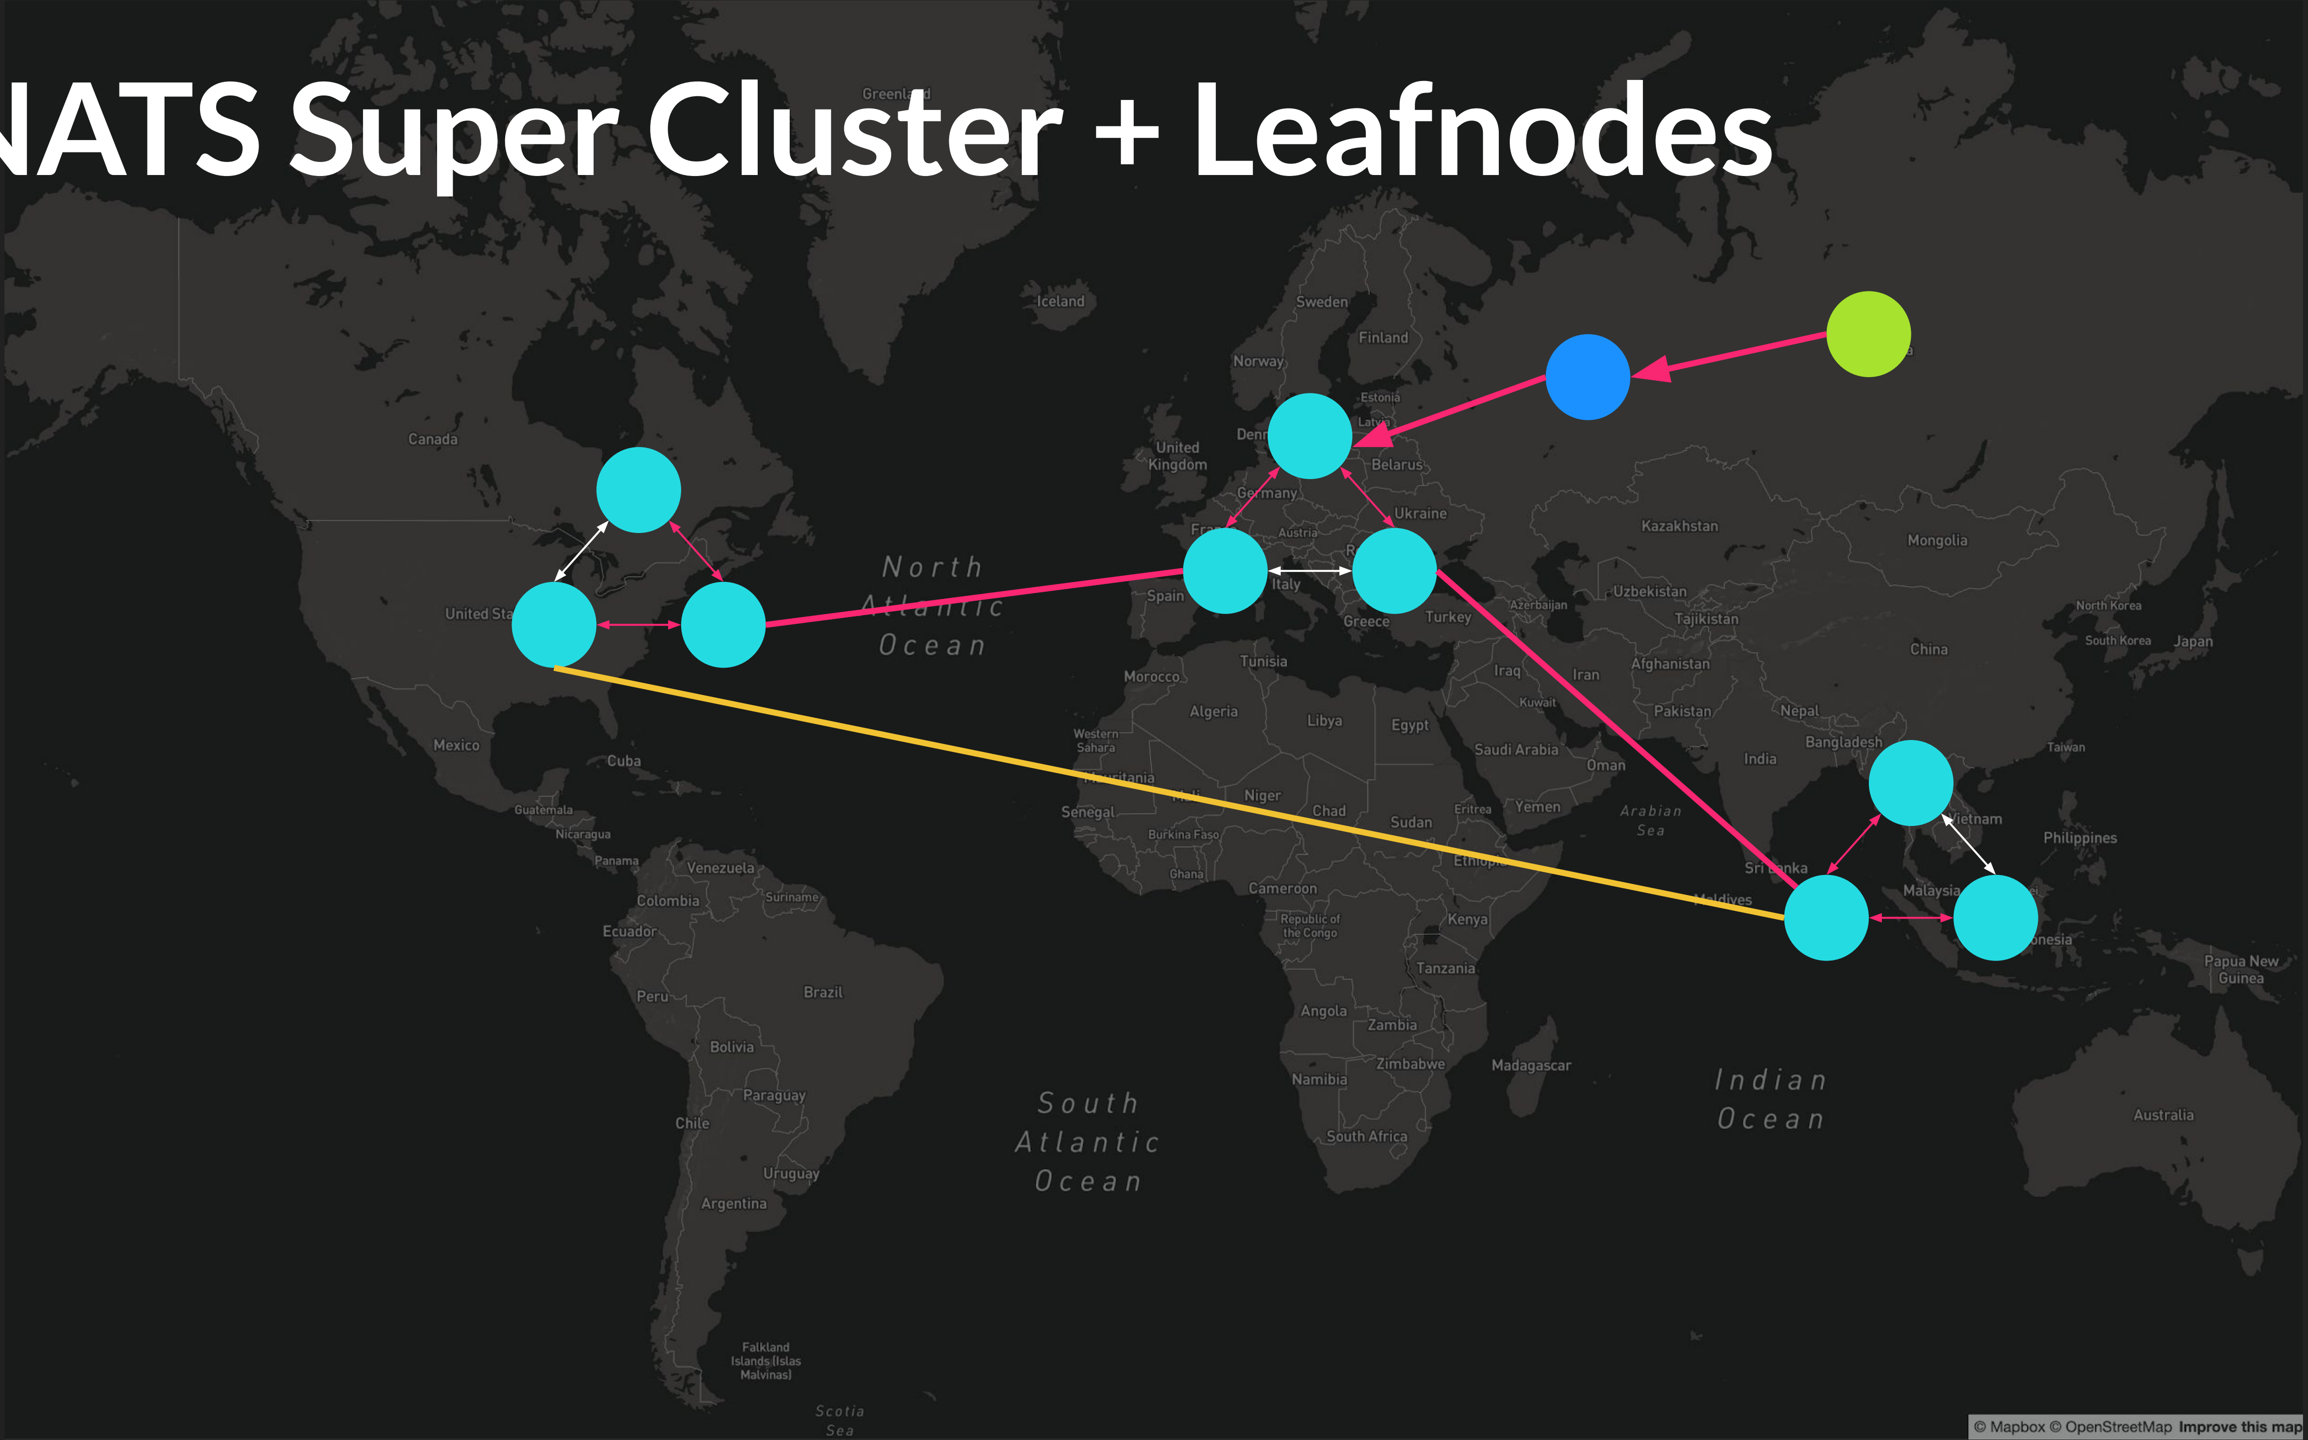

NATS Super Cluster + Leafnodes

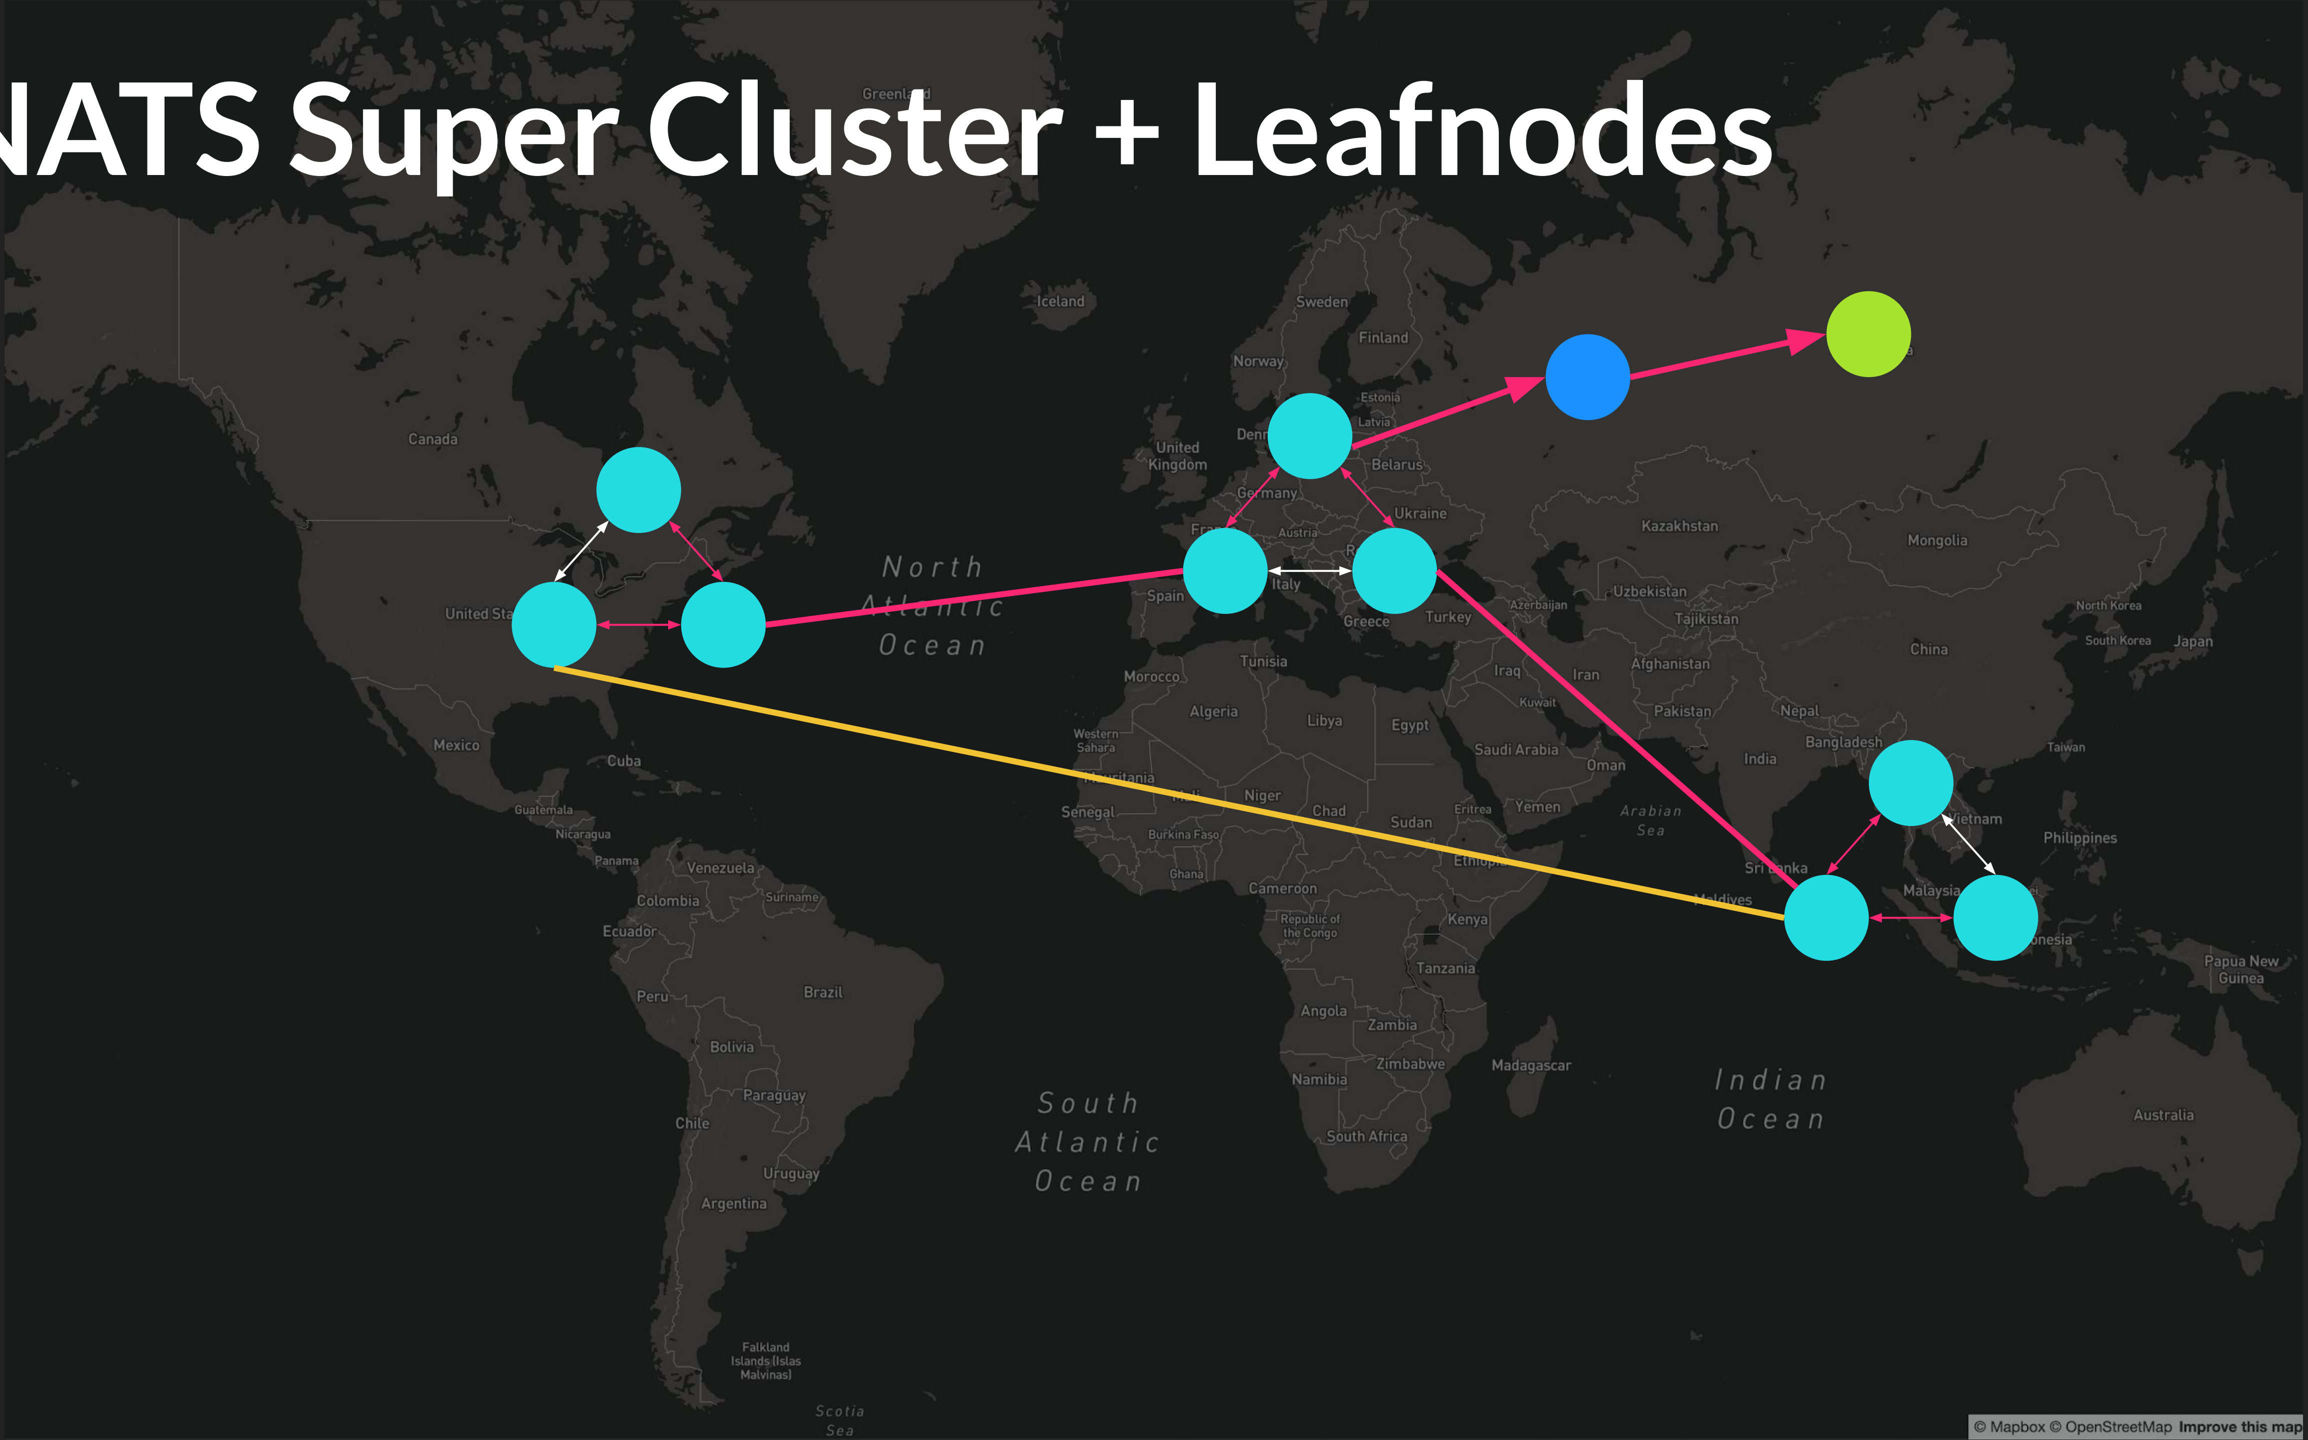

System Events

System Events

$SYS.ACCOUNT.<id>.CONNECT (client connects)

$SYS.ACCOUNT.<id>.DISCONNECT (client disconnects)

$SYS.SERVER.ACCOUNT.<id>.CONNS (account connections status)

$SYS.SERVER.<id>.CLIENT.AUTH.ERR (authentication error)

$SYS.ACCOUNT.<id>.LEAFNODE.CONNECT (leaf node connects)

$SYS.ACCOUNT.<id>.LEAFNODE.DISCONNECT (leaf node disconnects)

$SYS.SERVER.<id>.STATSZ (stats summary)

$SYS.REQ.SERVER.<id>.STATSZ (request server stat summary)

$SYS.REQ.SERVER.PING (discover all servers and metrics)

System Events

```
system_account = "AASYS..."

[62503] 2018/12/07 09:17:33.940827 [INF] Starting nats-server version 2.0.0-beta.2
[62503] 2018/12/07 09:17:33.940967 [DBG] Go build version go1.11.2
[62503] 2018/12/07 09:17:33.940975 [INF] Git commit [not set]
[62503] 2018/12/07 09:17:33.940987 [INF] Trusted Operators
[62503] 2018/12/07 09:17:33.941002 [INF]   System   : "NGS"
[62503] 2018/12/07 09:17:33.941009 [INF]   Operator: "Synadia Communications Inc."
[62503] 2018/12/07 09:17:33.941041 [INF]   Issued   : 2018-12-02 05:51:13 -0800 PST
[62503] 2018/12/07 09:17:33.941050 [INF]   Expires : 2019-12-02 05:51:13 -0800 PST
[62503] 2018/12/07 09:17:33.941197 [TRC] SYSTEM - <<- [SUB $SYS.SERVER.ACCOUNT.*.CONNS 1]
[62503] 2018/12/07 09:17:33.941262 [TRC] SYSTEM - <<- [SUB
$SYS._INBOX_.NCXTMN66MO2LRLK7EVWIO7UZPMK2Z2JSSLBL6TQ3JRXTKDP7RKPXE4TQ 2]
[62503] 2018/12/07 09:17:33.941285 [TRC] SYSTEM - <<- [SUB $SYS.REQ.ACCOUNT.*.CONNS 3]
[62503] 2018/12/07 09:17:33.941304 [TRC] SYSTEM - <<- [SUB $SYS.SERVER.*.SHUTDOWN 4]
[62503] 2018/12/07 09:17:33.941320 [TRC] SYSTEM - <<- [SUB $SYS.ACCOUNT.*.CLAIMS.UPDATE 5]
[62503] 2018/12/07 09:17:33.941347 [TRC] SYSTEM - <<- [SUB
$SYS.REQ.SERVER.NCXTMN66MO2LRLK7EVWIO7UZPMK2Z2JSSLBL6TQ3JRXTKDP7RKPXE4TQ.STATSZ 6]
[62503] 2018/12/07 09:17:33.941379 [TRC] SYSTEM - <<- MSG_PAYLOAD: ["{\n  \"server\": {\n    \"host\":
\"0.0.0.0\",\n    \"id\": \"NCXTMN66MO2LRLK7EVWIO7UZPMK2Z2JSSLBL6TQ3JRXTKDP7RKPXE4TQ\",\n    \"ver\":
\"2.0.0-beta.2\",\n    \"seq\": 2\n  },\n  \"acc\": \"AASYS...\",\n  \"conns\": 0\n}"]
```

System Events

```
$ nats-sub -creds ngs_sys.chain -s nats://nyc.nats-super-cluster.global"\$SYS.SERVER.>"
Listening on [$SYS.SERVER.>]
[#1] Received on [$SYS.SERVER.NAEVSLYZDRBITMXFHAV4DT3J7BA6ZN27NHWFH2272.STATSZ]: '{
  "server": {
    "id": "NAEVSLYZDRBITMXFHAV4DT3J7BA6ZN27NHWFH2”,
    "cluster": "bangalore",
    "ver": "2.0.0-beta.9",
    "seq": 28,
    "time": "2018-12-07T17:45:26.016398423Z"
  },
  "statsz": {
    "mem": 26882048,
    "cores": 2,
    "cpu": 0,
    "connections": 14,
    "total_connections": 435,
    "active_accounts": 18,
    "subscriptions": 73,
    "sent": {
      "msgs": 2395,
      "bytes": 698170
    },...
}'
```

KubeCon | CloudNativeCon

Europe 2019

Graceful Shutdown

aka. Lame Duck Mode

Lame Duck Shutdown ~~🦆~~

- Server now traps USR2 signal which will slowly disconnect clients.

- Avoids *thundering herd* issue, letting them reconnect to another server in the pool at a better pace.

Lame Duck Shutdown

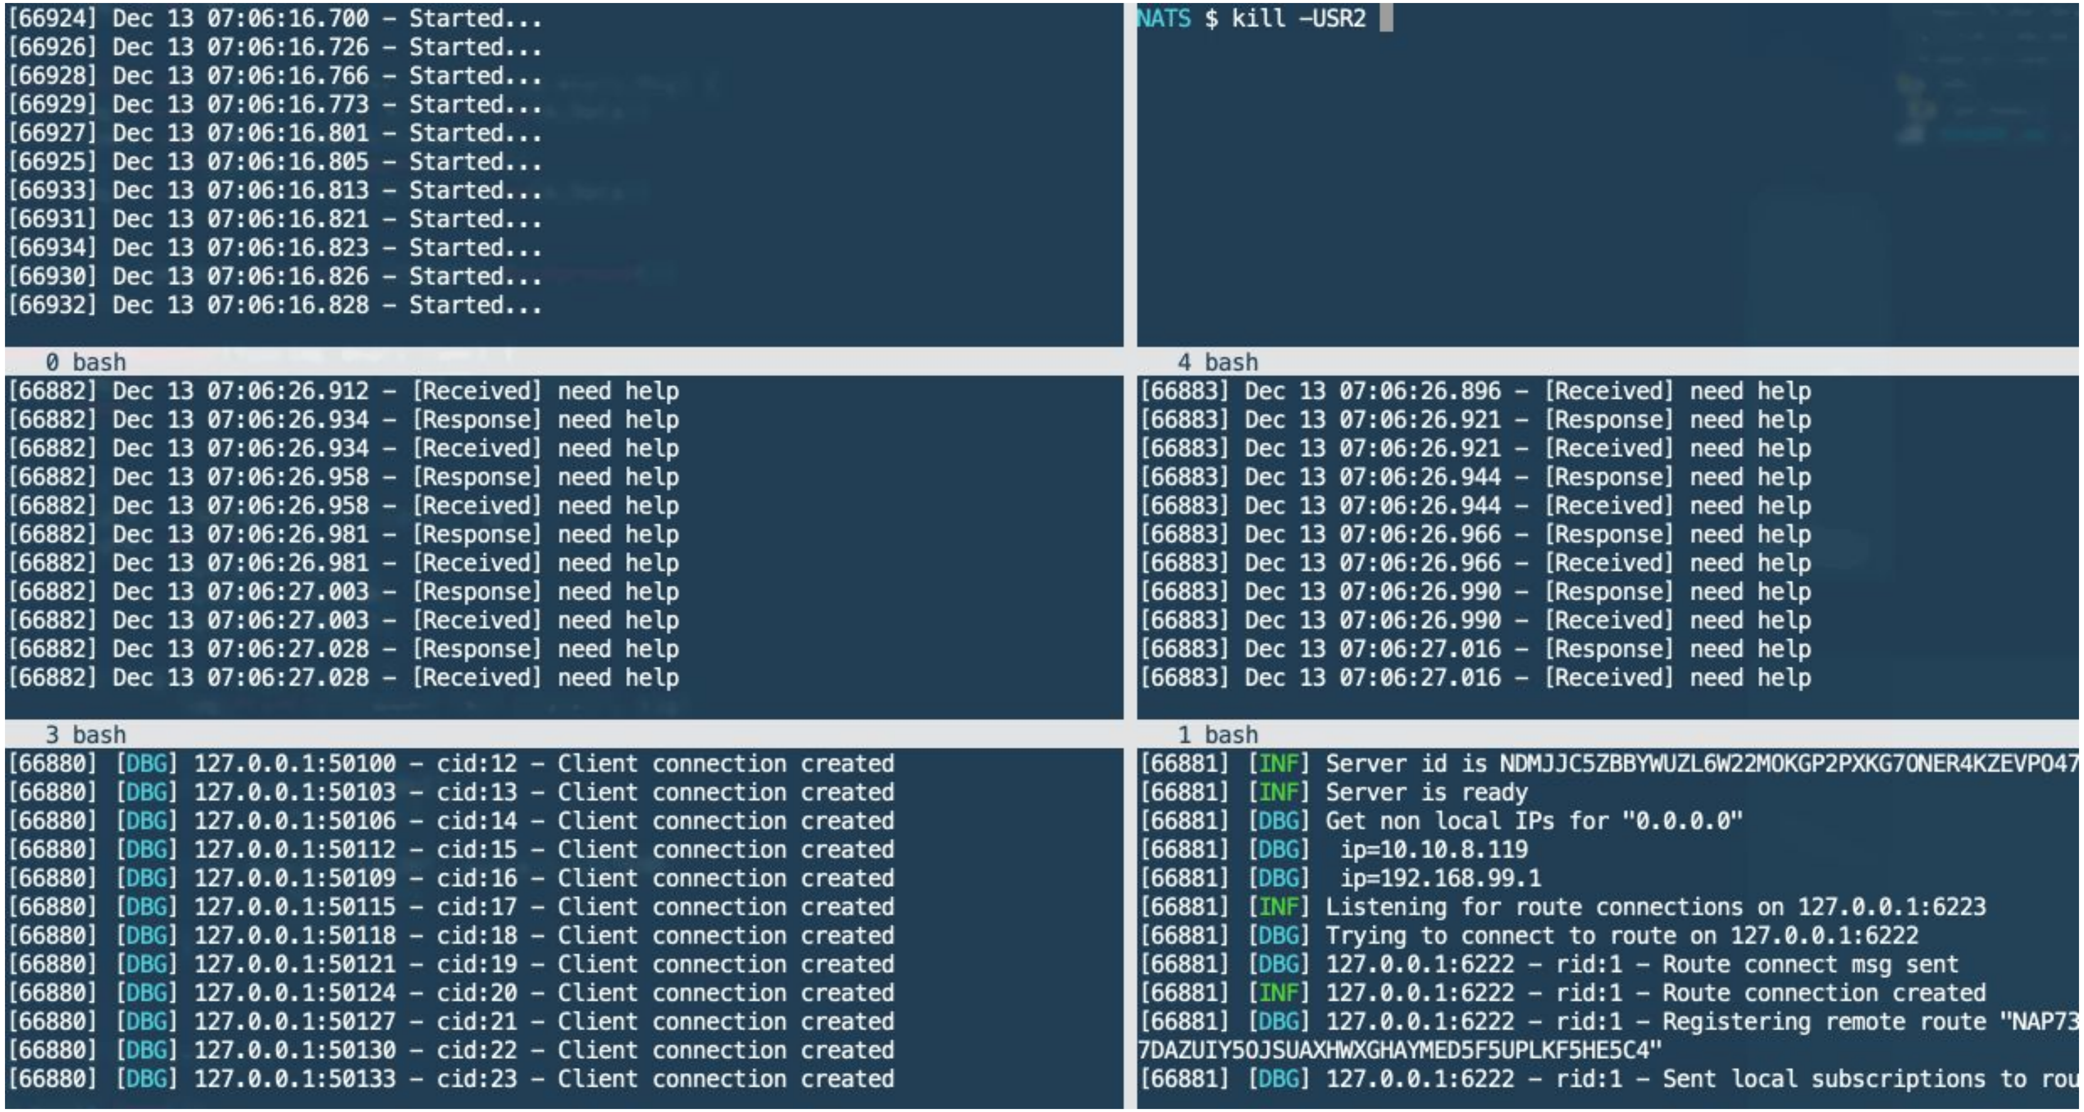

```
[66924] Dec 13 07:06:16.700 - Started...
[66926] Dec 13 07:06:16.726 - Started...
[66928] Dec 13 07:06:16.766 - Started...
[66929] Dec 13 07:06:16.773 - Started...
[66927] Dec 13 07:06:16.801 - Started...
[66925] Dec 13 07:06:16.805 - Started...
[66933] Dec 13 07:06:16.813 - Started...
[66931] Dec 13 07:06:16.821 - Started...
[66934] Dec 13 07:06:16.823 - Started...
[66930] Dec 13 07:06:16.826 - Started...
[66932] Dec 13 07:06:16.828 - Started...
```

```
NATS $ kill -USR2
```

```
 0 bash
[66882] Dec 13 07:06:26.912 - [Received] need help
[66882] Dec 13 07:06:26.934 - [Response] need help
[66882] Dec 13 07:06:26.934 - [Received] need help
[66882] Dec 13 07:06:26.958 - [Response] need help
[66882] Dec 13 07:06:26.958 - [Received] need help
[66882] Dec 13 07:06:26.981 - [Response] need help
[66882] Dec 13 07:06:26.981 - [Received] need help
[66882] Dec 13 07:06:27.003 - [Response] need help
[66882] Dec 13 07:06:27.003 - [Received] need help
[66882] Dec 13 07:06:27.028 - [Response] need help
[66882] Dec 13 07:06:27.028 - [Received] need help
```

```
 4 bash
[66883] Dec 13 07:06:26.896 - [Received] need help
[66883] Dec 13 07:06:26.921 - [Response] need help
[66883] Dec 13 07:06:26.921 - [Received] need help
[66883] Dec 13 07:06:26.944 - [Response] need help
[66883] Dec 13 07:06:26.944 - [Received] need help
[66883] Dec 13 07:06:26.966 - [Response] need help
[66883] Dec 13 07:06:26.966 - [Received] need help
[66883] Dec 13 07:06:26.990 - [Response] need help
[66883] Dec 13 07:06:26.990 - [Received] need help
[66883] Dec 13 07:06:27.016 - [Response] need help
[66883] Dec 13 07:06:27.016 - [Received] need help
```

```
 3 bash
[66880] [DBG] 127.0.0.1:50100 - cid:12 - Client connection created
[66880] [DBG] 127.0.0.1:50103 - cid:13 - Client connection created
[66880] [DBG] 127.0.0.1:50106 - cid:14 - Client connection created
[66880] [DBG] 127.0.0.1:50112 - cid:15 - Client connection created
[66880] [DBG] 127.0.0.1:50109 - cid:16 - Client connection created
[66880] [DBG] 127.0.0.1:50115 - cid:17 - Client connection created
[66880] [DBG] 127.0.0.1:50118 - cid:18 - Client connection created
[66880] [DBG] 127.0.0.1:50121 - cid:19 - Client connection created
[66880] [DBG] 127.0.0.1:50124 - cid:20 - Client connection created
[66880] [DBG] 127.0.0.1:50127 - cid:21 - Client connection created
[66880] [DBG] 127.0.0.1:50130 - cid:22 - Client connection created
[66880] [DBG] 127.0.0.1:50133 - cid:23 - Client connection created
```

```
 1 bash
[66881] [INF] Server id is NDMJJC5ZBBYWUZL6W22MOKGP2PXKG7ONER4KZEVPO47
[66881] [INF] Server is ready
[66881] [DBG] Get non local IPs for "0.0.0.0"
[66881] [DBG]  ip=10.10.8.119
[66881] [DBG]  ip=192.168.99.1
[66881] [INF] Listening for route connections on 127.0.0.1:6223
[66881] [DBG] Trying to connect to route on 127.0.0.1:6222
[66881] [DBG] 127.0.0.1:6222 - rid:1 - Route connect msg sent
[66881] [INF] 127.0.0.1:6222 - rid:1 - Route connection created
[66881] [DBG] 127.0.0.1:6222 - rid:1 - Registering remote route "NAP73
7DAZUIY5OJSUAXHWXGHAYMED5F5UPLKF5HE5C4"
[66881] [DBG] 127.0.0.1:6222 - rid:1 - Sent local subscriptions to rou
```

KubeCon | CloudNativeCon

Europe 2019

Deploying NATS v2 on Kubernetes

Using the NATS Operator

NATS Operator

The recommended way of running NATS on Kubernetes

https://github.com/nats-io/nats-operator

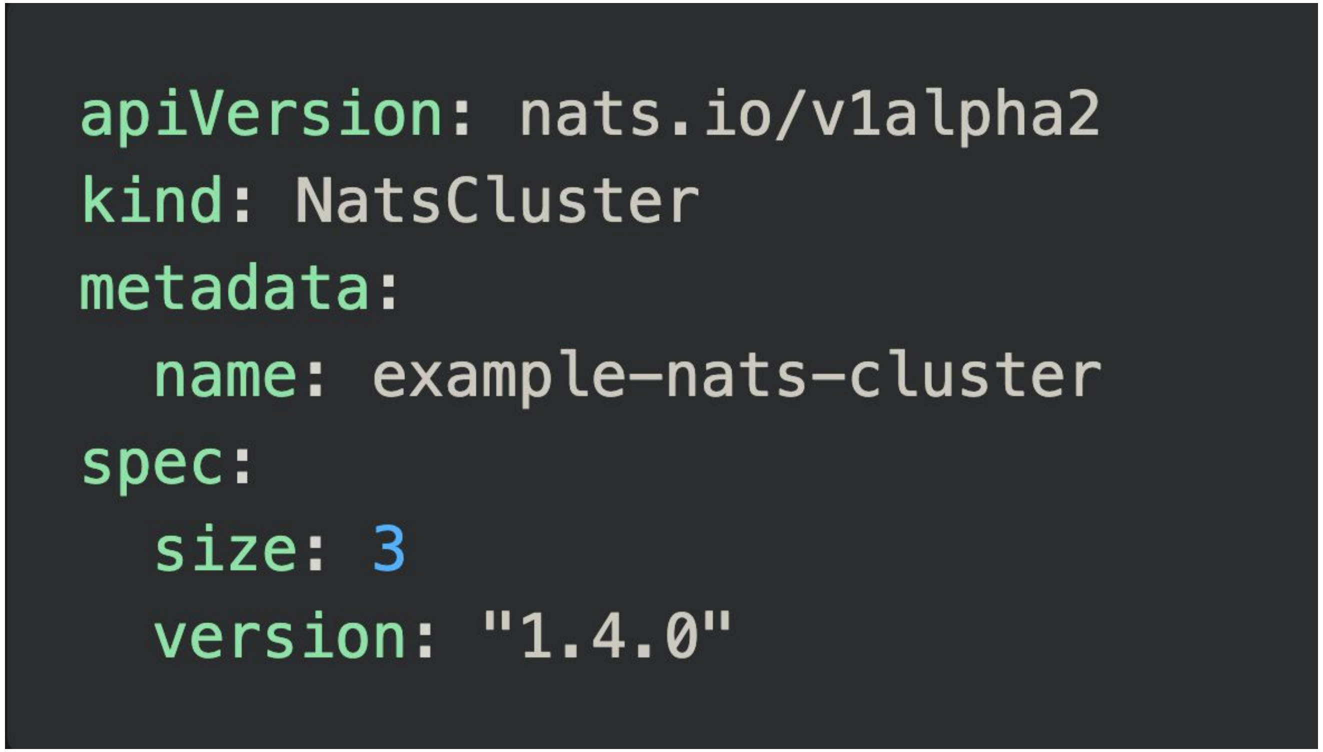

```yaml
apiVersion: nats.io/v1alpha2
kind: NatsCluster
metadata:
  name: example-nats-cluster
spec:
  size: 3
  version: "1.4.0"
```

NATS Operator

- Works best on Kubernetes v1.12
- Clustering support
- Supports Configuration Reload
- Cluster-scoped
- Dynamic Authentication using bound tokens
- TLS support using cert-manager
- Prometheus Sidecar
- NATS v2 ready
    - Gateways, Leafnodes and Trusted Operator mode

NATS v2 on K8S

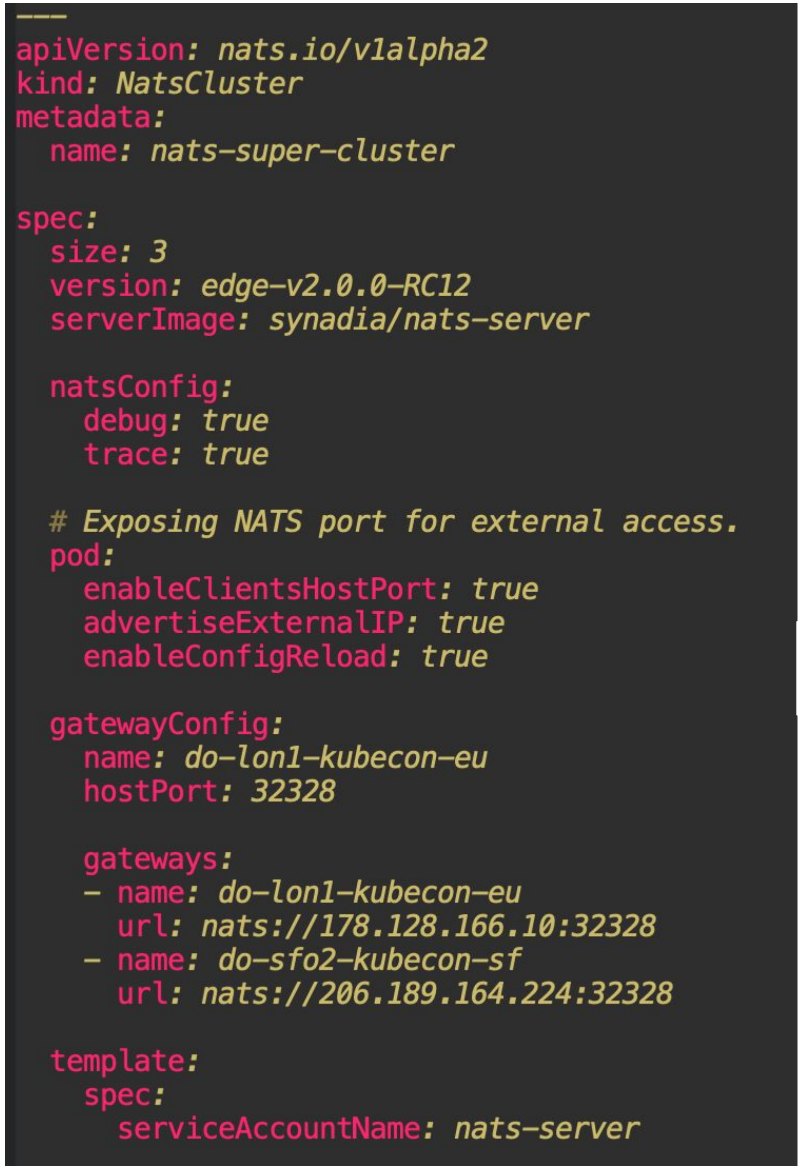

```yaml
---
apiVersion: nats.io/v1alpha2
kind: NatsCluster
metadata:
  name: nats-super-cluster

spec:
  size: 3
  version: edge-v2.0.0-RC12
  serverImage: synadia/nats-server

  natsConfig:
    debug: true
    trace: true

  # Exposing NATS port for external access.
  pod:
    enableClientsHostPort: true
    advertiseExternalIP: true
    enableConfigReload: true

  gatewayConfig:
    name: do-lon1-kubecon-eu
    hostPort: 32328

    gateways:
    - name: do-lon1-kubecon-eu
      url: nats://178.128.166.10:32328
    - name: do-sfo2-kubecon-sf
      url: nats://206.189.164.224:32328

  template:
    spec:
      serviceAccountName: nats-server
```

Result

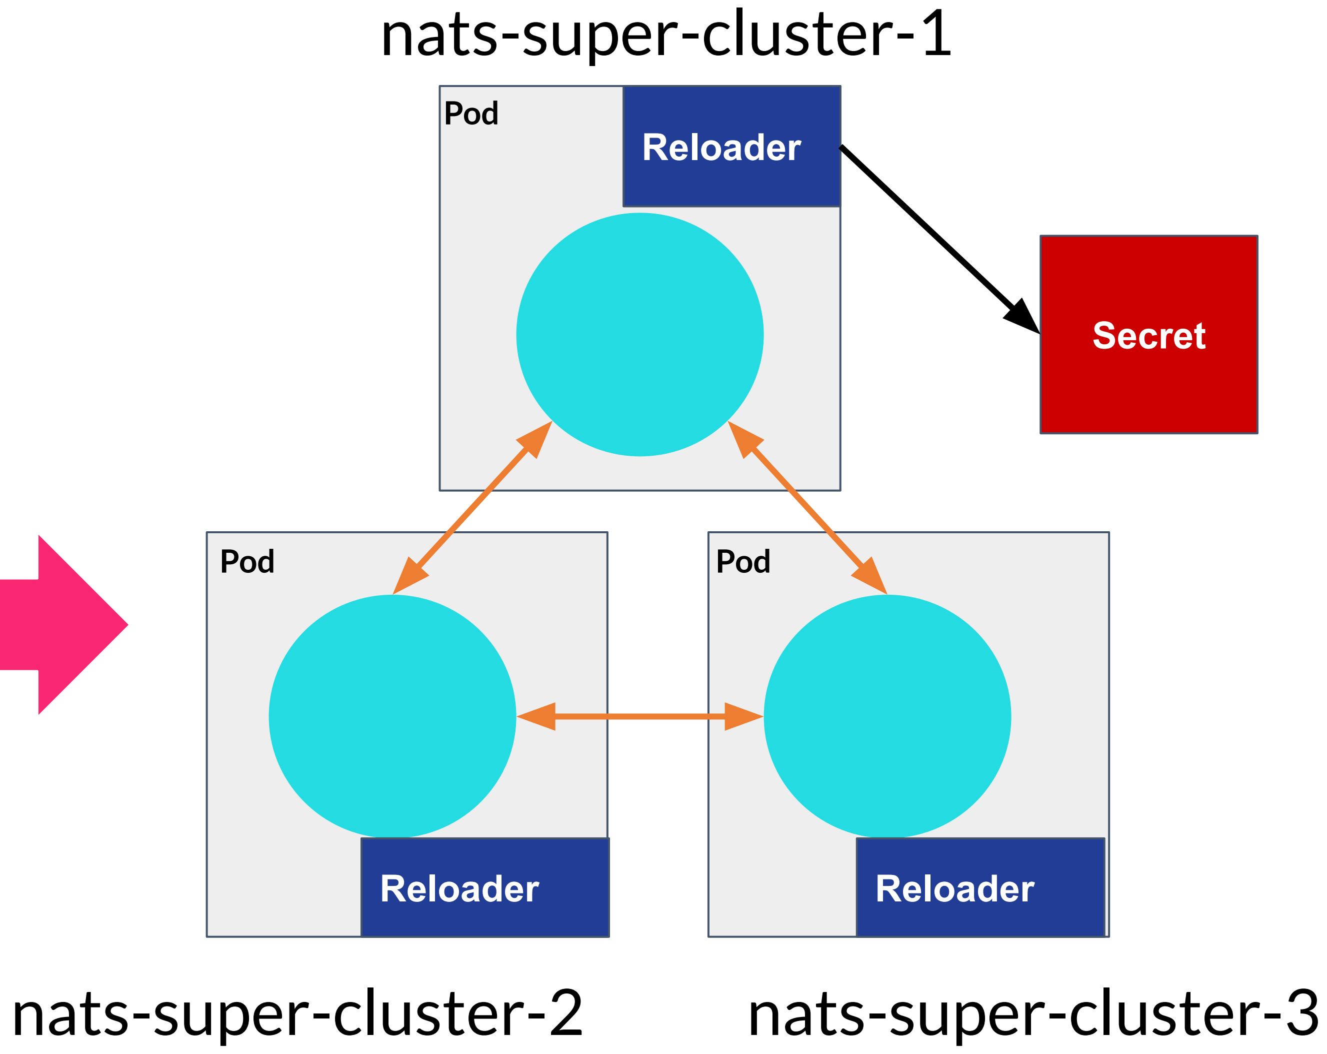

Clustering using Headless svc

An A record for each one of the pods to create the full mesh cluster.

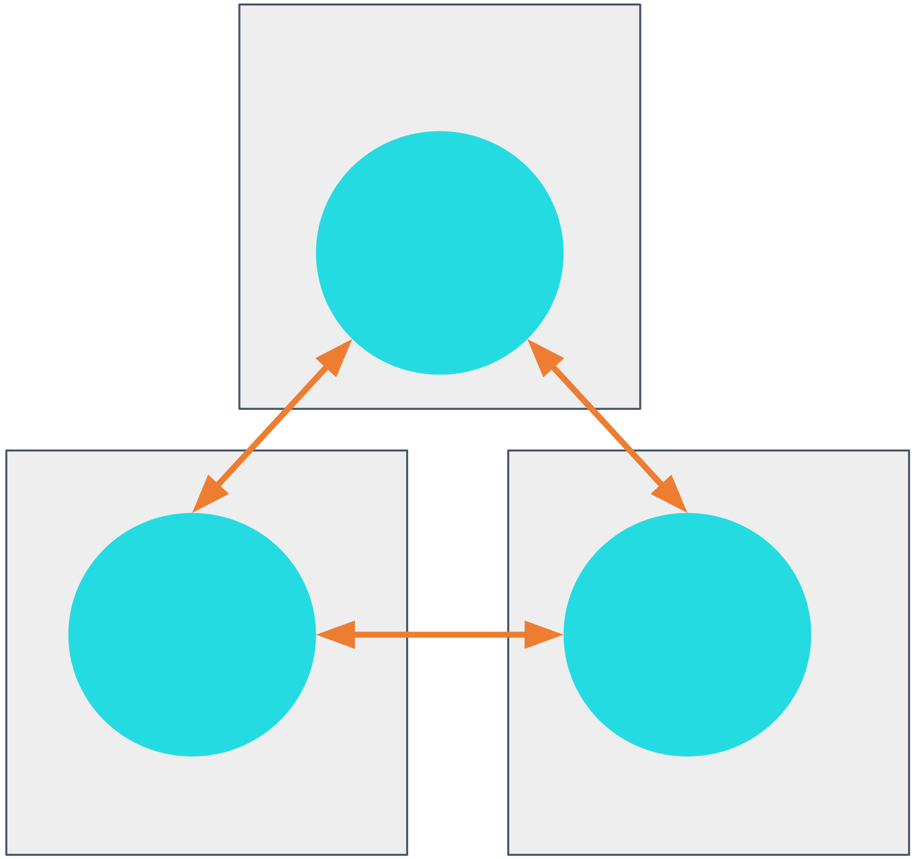

nats-super-cluster-1.nats-super-cluster-mgmt

nats-super-cluster-2.nats-super-cluster-mgmt

nats-super-cluster-3.nats-super-cluster-mgmt

Advertising external IP

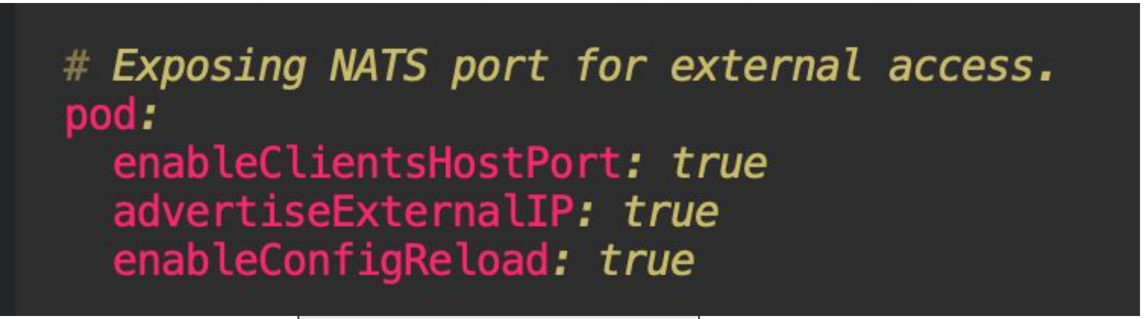

```
# Exposing NATS port for external access.
pod:
  enableClientsHostPort: true
  advertiseExternalIP: true
  enableConfigReload: true
```

If the kubelet nodes have an ExternalIP in the metadata then this option makes it available for the pod which can help on failover.

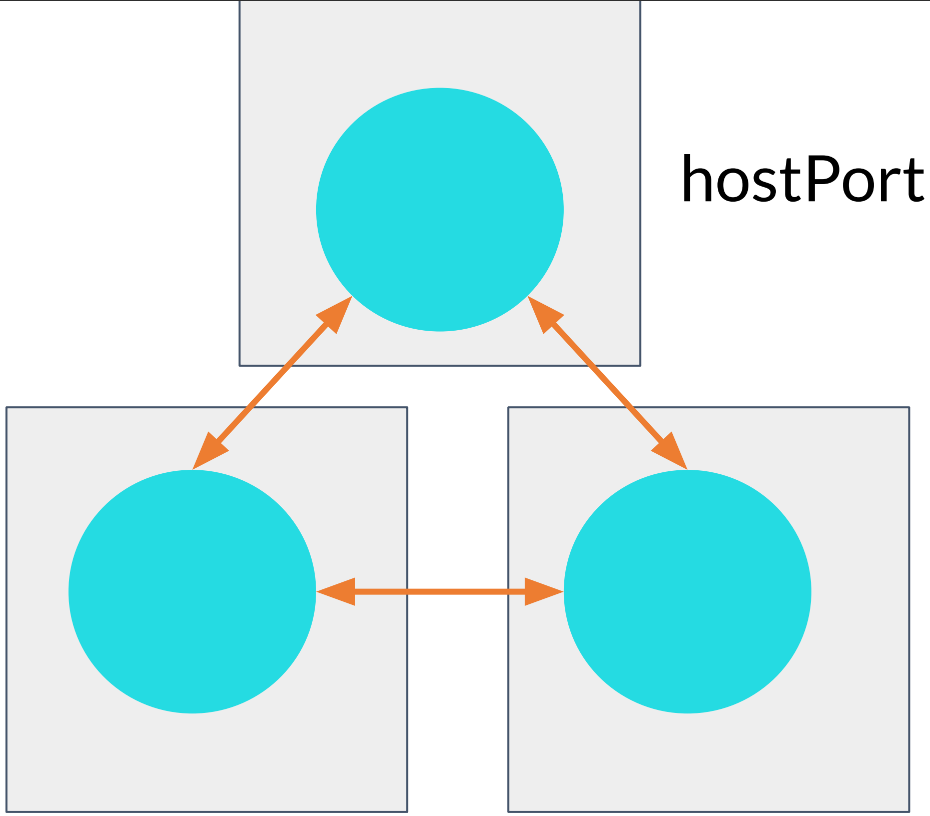

hostPort: 4222

telnet 178.128.164.94 4222
INFO {...
connect_urls":["178.128.164.94:4222","178.128.166.10:4222","178.128.164.132:4222"]}

Operator Config

- A system account for generating system events
- The operator JWT
- The resolver configuration (location of the nats account server)

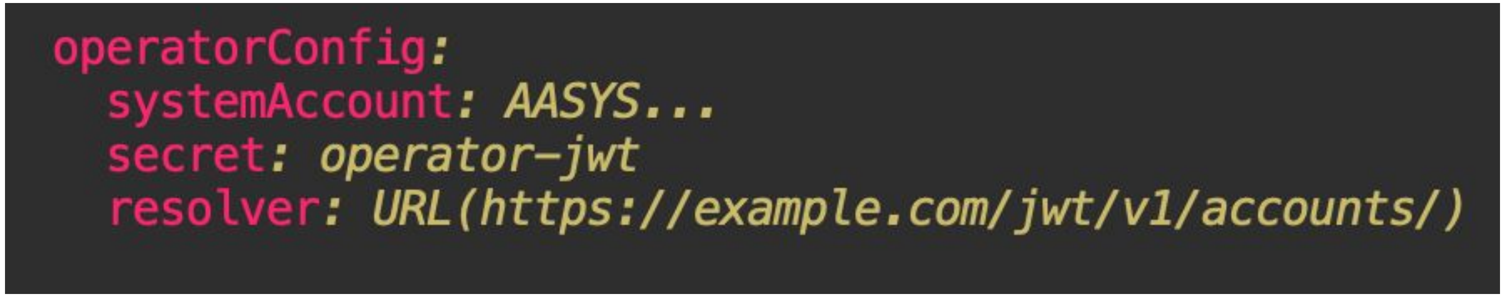

```
operatorConfig:
  systemAccount: AASYS...
  secret: operator-jwt
  resolver: URL(https://example.com/jwt/v1/accounts/)
```

Using the System Account

- Collect all the available servers across datacenters by sending a NATS request that expects many responses (very useful!)

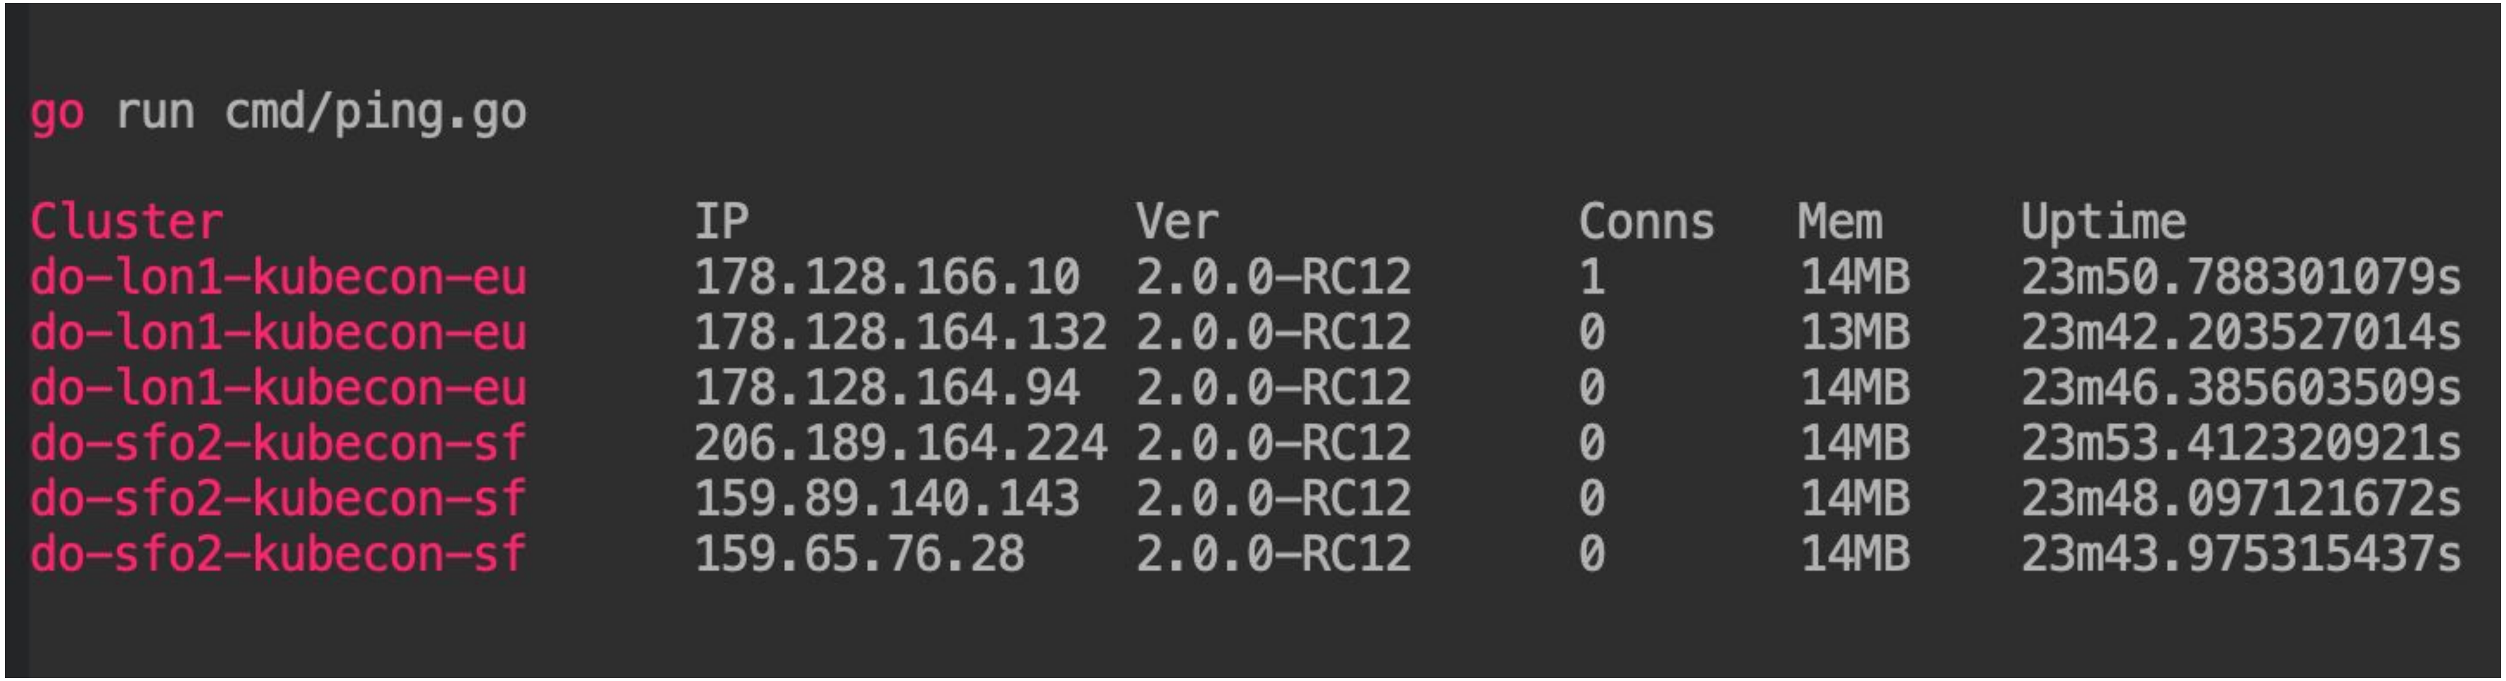

```
go run cmd/ping.go

Cluster                 IP              Ver             Conns   Mem     Uptime
do-lon1-kubecon-eu      178.128.166.10  2.0.0-RC12      1       14MB    23m50.788301079s
do-lon1-kubecon-eu      178.128.164.132 2.0.0-RC12      0       13MB    23m42.203527014s
do-lon1-kubecon-eu      178.128.164.94  2.0.0-RC12      0       14MB    23m46.385603509s
do-sfo2-kubecon-sf      206.189.164.224 2.0.0-RC12      0       14MB    23m53.412320921s
do-sfo2-kubecon-sf      159.89.140.143  2.0.0-RC12      0       14MB    23m48.097121672s
do-sfo2-kubecon-sf      159.65.76.28    2.0.0-RC12      0       14MB    23m43.975315437s
```

Questions?

KubeCon | CloudNativeCon

Europe 2019

Thanks!

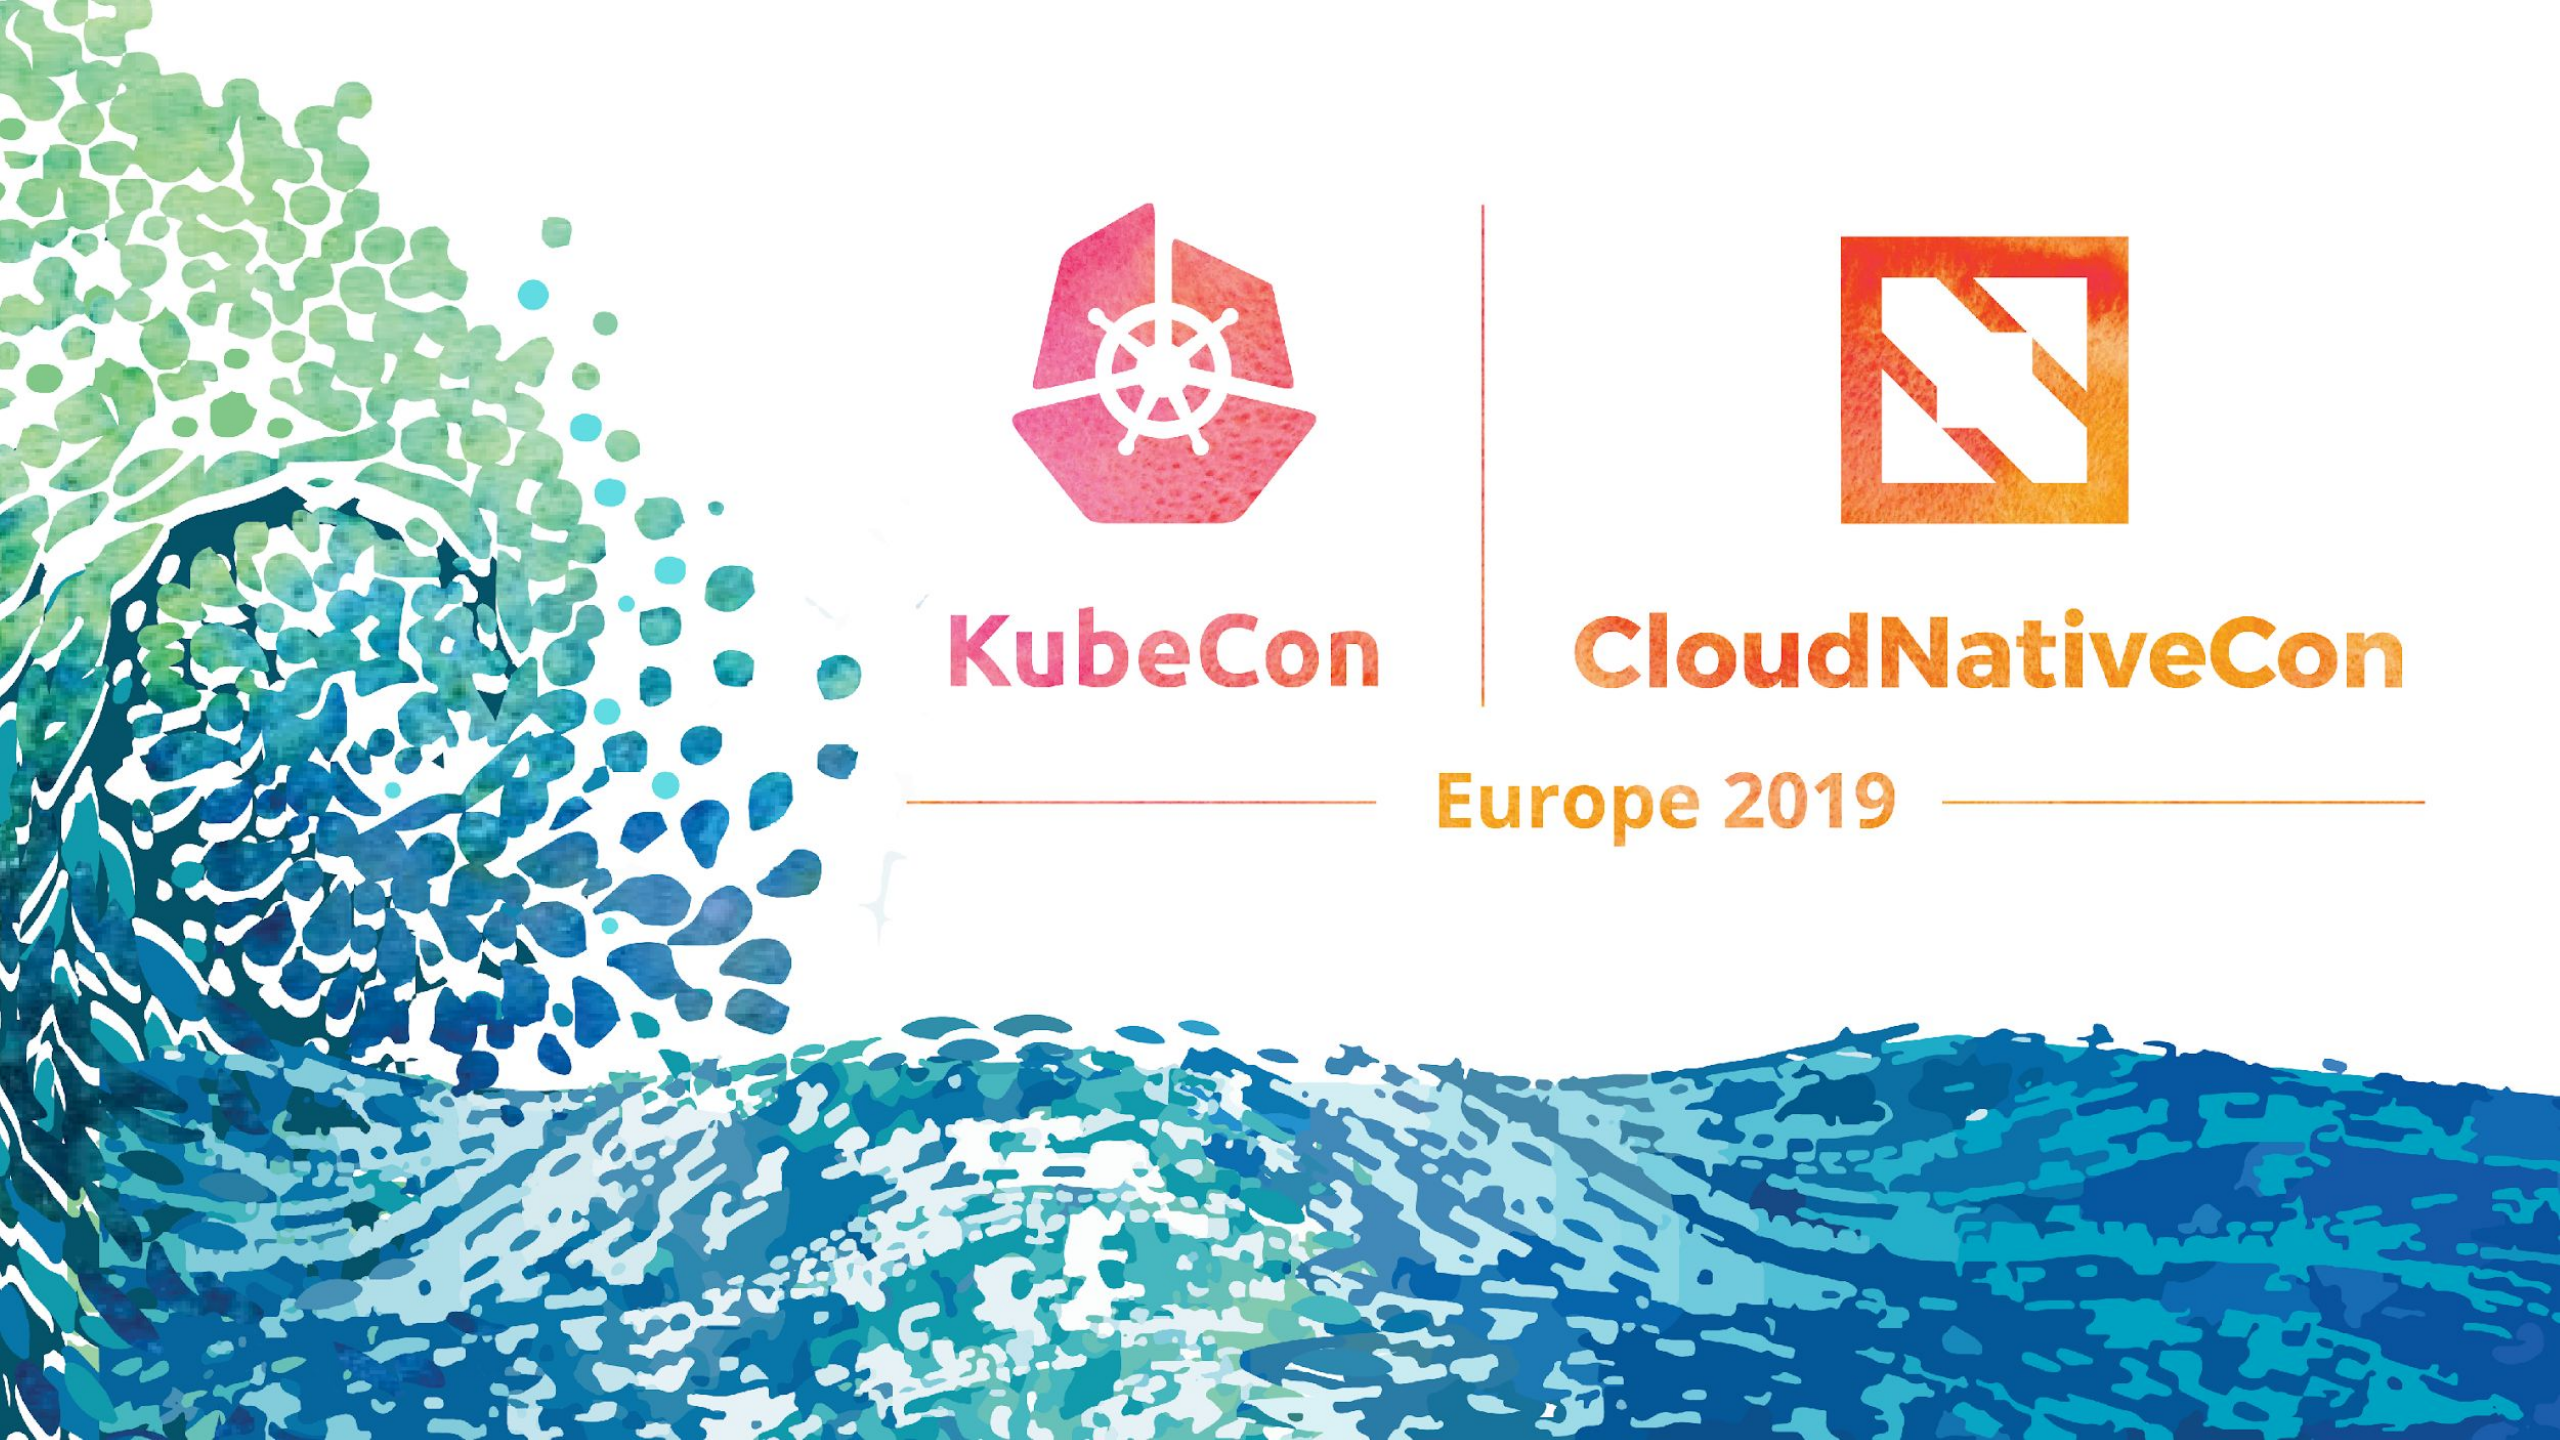

KubeCon | CloudNativeCon

Europe 2019