



KubeCon



CloudNativeCon

Europe 2018

# Introducing

# NATS

Thursday, May 3 14:45 - 15:20

Colin Sullivan / [colin@nats.io](mailto:colin@nats.io)

Waldemar Quevedo / [wally@nats.io](mailto:wally@nats.io)





# What is messaging?



KubeCon



CloudNativeCon

Europe 2018

Messaging here means “Message Oriented Middleware”

- Wikipedia
  - “*Message-oriented middleware (MOM) is software or hardware infrastructure supporting sending and receiving messages between distributed systems.*”
- Ken Klingenstein
  - “*Middleware is the intersection of the stuff that network engineers don't want to do with the stuff that applications developers don't want to do.*”



# Why does CNCF need messaging?



KubeCon



CloudNativeCon

Europe 2018

- Distributed systems architecture has been disrupted via decomposition as cloud technology matures.
- Communications is lagging and a technology is needed to address various use cases:
  - ✓ Multiple messaging patterns bundled into one technology
  - ✓ Location transparency
  - ✓ Decoupling of data producers and consumers
  - ✓ Built-in load balancing
- NATS was created specifically to meet these communication needs for next generation cloud native applications.



# What is NATS?



KubeCon



CloudNativeCon

Europe 2018

NATS is a seven year old, production tested, cloud-native messaging system made for developers and operators who want to spend more time doing their work and less time worrying about how to do messaging.

- DNA: Performance, simplicity, security, and availability.
- Built from the ground up to be cloud native
- Common pattern support including request/reply, pub/sub, and load balanced queue subscribers



# History



KubeCon



CloudNativeCon

Europe 2018



## Derek Collison

*Founder and CEO at Synadia*

Founder and former CEO at Apcera  
CTO, Chief Architect at VMware  
Architected CloudFoundry  
Technical Director at Google  
SVP and Chief Architect at TIBCO

Created by Derek Collison

Derek has been building messaging systems and solutions > 25 yrs

Maintained by a highly experienced messaging team at Synadia

Engaged User Community



# Overview



KubeCon



CloudNativeCon

Europe 2018

- Pure publish/subscribe with common messaging patterns built on top
- Clustering of servers with built-in auto discovery
- Clients available for over 30 programming languages
- Log based streaming with guaranteed delivery
- Prometheus NATS Exporter
- Kubernetes Operator



# Use Cases



KubeCon



CloudNativeCon

Europe 2018

NATS is used in a range of technical use cases

- Messaging in the cloud
  - ✓ General messaging
  - ✓ Microservices Transport
  - ✓ Control Planes
  - ✓ Service Discovery
- IoT Messaging
- Mobile and Big Data
- High Fan-out Messaging
- Replacing or Augmenting Legacy Messaging



# Users and Adopters



KubeCon



CloudNativeCon

Europe 2018



Acadian | Apcera | Apporeto | Baidu | Bridgevine | Capital One | Clarifai | Cloud Foundry | Comcast | Ericsson | Faber | Fission |  
General Electric | Greta | HTC | Logimethods | Netlify | Pex | Pivotal | Platform9 | Rapidloop | Samsung | Sendify | Sensay |  
StorageOS | VMware | Weaveworks | Workiva



# Contribution Statistics



KubeCon



CloudNativeCon

Europe 2018

- 200+ contributors
- 30 public repos
  - 50+ releases
  - 8000+ GitHub stars across repos
- ~25M NATS server Docker Hub pulls
- ~7M NATS streaming server pulls
- 680+ Slack members
- 20+ releases of the NATS server since June 2014, ~= 5/year



**KubeCon**



**CloudNativeCon**

Europe 2018

# Messaging Patterns





# Messaging Patterns



KubeCon



CloudNativeCon

Europe 2018

NATS has three built-in patterns:

- Publish/Subscribe
  - Publish data to a subject where 1-N subscribers will receive the data.
- Request/Reply
  - Send a request message containing a unique reply subject and responders can send a reply message back only to the responder.
- Load balanced queue subscribers
  - Subscribers belong to a group, and the server sends a message to only one subscriber in the group, effectively load balancing.



**KubeCon**



**CloudNativeCon**

Europe 2018

# NATS DNA





DNA



KubeCon



CloudNativeCon

Europe 2018

- Performance and Scalability
- Simplicity
- Security
- Availability

***“Just what you need and none of what you don’t.”***



KubeCon



CloudNativeCon

Europe 2018

# Performance and Scalability

Simplicity  
Security  
Availability





# Performance



KubeCon



CloudNativeCon

Europe 2018

- Performance is a part of **every decision** we make
- NATS is optimized to perform and scale
  - ✓ Design for scale
  - ✓ Careful analysis of the fastpath
- Just as important is what **NOT** to implement
  - ✗ Message guarantees in core NATS
  - ✗ Transactions
  - ✗ Message Schemas
  - ✗ Last Will and Testament
  - ✗ Message Groups



# Performance



KubeCon



CloudNativeCon

Europe 2018

## 18 Million msgs/sec

Benchmark____Pub0b_Payload-20	30000000	55.1 ns/op	199.78 MB/s
Benchmark____Pub8b_Payload-20	30000000	55.8 ns/op	340.21 MB/s
Benchmark____Pub32b_Payload-20	20000000	63.4 ns/op	694.34 MB/s
Benchmark____Pub128B_Payload-20	20000000	79.8 ns/op	1766.47 MB/s
Benchmark____Pub256B_Payload-20	20000000	98.1 ns/op	2741.51 MB/s
Benchmark____Pub1K_Payload-20	5000000	283 ns/op	3660.72 MB/s
Benchmark____Pub4K_Payload-20	1000000	1395 ns/op	2945.30 MB/s
Benchmark____Pub8K_Payload-20	500000	2846 ns/op	2882.35 MB/s
Benchmark_AuthPub0b_Payload-20	10000000	126 ns/op	86.82 MB/s
Benchmark_____PubSub-20	10000000	135 ns/op	
Benchmark____PubSubTwoConns-20	10000000	136 ns/op	
Benchmark____PubTwoQueueSub-20	10000000	152 ns/op	
Benchmark____PubFourQueueSub-20	10000000	152 ns/op	
Benchmark__PubEightQueueSub-20	10000000	152 ns/op	



KubeCon



CloudNativeCon

Europe 2018

Performance and Scalability

**Simplicity**

Security

Availability





# Simplicity



KubeCon



CloudNativeCon

Europe 2018

NATS aims for ease of use from installation through operation

- Single binary
  - 7 MB docker image
  - No external dependencies
- Text-based protocol with only a handful of verbs
- Easy to parse JSON returned from monitoring endpoints
- Straightforward configuration of TLS and credentials



# Simplicity (Continued)



KubeCon



CloudNativeCon

Europe 2018

- Little configuration
  - Clients only require a url and credentials
  - Servers use a few command parameters or a single configuration file
- Clustering is automatic
  - Configure seed servers and just add servers to grow your cluster
  - Supported clients will be notified of cluster topology changes
- Simple and Straightforward API



# Simplicity (Go API)



KubeCon



CloudNativeCon

Europe 2018

```
package main

import (
    "log"

    "github.com/nats-io/go-nats"
)

func main() {
    nc, _ := nats.Connect("nats://demo.nats.io:4222")
    defer nc.Close()

    nc.Publish("nats.demo.simple", []byte("Hello!"))
    nc.Flush()

    if err := nc.LastError(); err != nil {
        log.Fatal(err)
    } else {
        log.Printf("Published message.")
    }
}
```

Simple application written in go to publish a message.

Connect, Publish, and check for errors.



# Simplicity (Java API)



KubeCon



CloudNativeCon

Europe 2018

```
import io.nats.client.*;

public class SimpleSub {

    static public void main(String args[]) throws Exception {
        try (Connection nc = Nats.connect(url: "nats://demo.nats.io:4222")) {
            nc.subscribe(subject: "nats.demo.simple", message -> {
                System.out.println("Received message: " + new String(message.getData()));
            });

            System.out.println("Waiting for messages...");

            Thread.sleep(Long.MAX_VALUE);
        }
    }
}
```

Java is just as simple.



KubeCon



CloudNativeCon

Europe 2018

Performance and Scalability  
Simplicity

**Security**

Availability





# Security



KubeCon



CloudNativeCon

Europe 2018

NATS secures distributed systems through...

- ✓ TLS
- ✓ User/Token Based Authentication
- ✓ Authorization

...update these through configuration reload with zero downtime.



# Security



KubeCon



CloudNativeCon

Europe 2018

## TLS Support

- Configured in the NATS server
- Client Side Certificates
- CA Certificate Support
- Defaults to most secure ciphers
- Override to set ciphers and curve preferences

## User Authentication

- Credentials with username/password or token are associated with a connection
- Bcrypt password protection in configuration files



## Authorization

- Permission based roles
- Publish/Subscribe
- Configuration Reload
  - Update / Add / Remove TLS Certificates
  - Add or remove credentials
  - Grant or revoke permissions

The Synadia team is designing a new architecture where no private keys/passwords are stored on the system.



KubeCon



CloudNativeCon

Europe 2018

Performance and Scalability  
Simplicity  
Security

**Availability**





# Availability



KubeCon



CloudNativeCon

Europe 2018

NATS prioritizes the health and availability of the system as a whole rather than attempting to service any individual client or server, creating a foundation for stable and resilient systems.

- The NATS dial-tone, always on
- NATS server “selfish optimization”
  - Slow consumers and other badly behaving clients are disconnected
- Full Mesh Clustering of NATS servers
- Self Healing
  - Clients and servers automatically reconnect or rejoin a cluster after failures



# Clustering



KubeCon



CloudNativeCon

Europe 2018

NATS can grow or shrink a server cluster through auto-discovery

- Dynamically scale a cluster
  - Start with a few “seed” servers and add as necessary to scale
  - Topology changes are propagated to other servers in the cluster allowing servers to join the full mesh
- Supported clients are **also aware of topology changes**
  - Internal knowledge of cluster topology is updated
  - Automatically fail over to new servers in a cluster

This results in no configuration changes for clients when the cluster topology changes



# Resilience



KubeCon



CloudNativeCon

Europe 2018

*“Simplicity is prerequisite for reliability” - Edsger Dijkstra*



**Christopher Watford**

@ecgwatford

Following

Replying to @stevedischinger

122 days uptime on 0.8.0 right now. ~130M messages. @stevedischinger @fathmanjim @nats\_io

Self healing clusters and automatic reconnection allow for resilience at scale



**Jim Fathman**

@FathmanJim

Following

Beginning to think my #NATSio T-shirt will be worn out before I restart NATS for the first time. Uptime 120 days, using #Nodejs client. #IoT



KubeCon



CloudNativeCon

Europe 2018

# How badly do I need a message?

Delivery Modes





# Delivery Modes



KubeCon



CloudNativeCon

Europe 2018

## What are delivery modes?

- At most once
  - No guarantee of delivery - messages can be lost - applications must detect and handle lost messages
- At least once
  - A message will always be delivered, but in certain cases may be delivered more than once
- Exactly once
  - Arguably unnecessary, complex and slow



# NATS Delivery Modes



KubeCon



CloudNativeCon

Europe 2018

Core NATS provides **at most once** delivery guarantees

NATS Streaming provides **at least once** delivery guarantees



# NATS Streaming



KubeCon



CloudNativeCon

Europe 2018

NATS Streaming is a data streaming system powered by NATS

Features include

- **At-least-once** delivery
- Replay by time or seqno offset
- Last/initial value caching
- Durable subscribers
- Rate matching per subscriber
- Memory, File, or Database storage
- High Availability through fault tolerant or clustered configurations
- Scale through partitioning



**KubeCon**



**CloudNativeCon**

Europe 2018

# Monitoring





# Monitoring NATS Servers



KubeCon



CloudNativeCon

Europe 2018

- Each server monitored independently
- Simply enable a monitoring port
- Poll an endpoint for metrics:
  - <http://demo.nats.io:8222/varz>
  - <http://demo.nats.io:8222/connz>
  - <http://demo.nats.io:8222/subsz>
  - <http://demo.nats.io:8222/routez>
- Simplicity allows for building complex tooling



# Monitoring Data



KubeCon



CloudNativeCon

Europe 2018

Server monitoring information is JSON formatted:

```
{
  "server_id": "EiRJABZmVpWQDpriVqbbtw",
  "version": "1.0.6",
  "go": "go1.9.2",
  "host": "0.0.0.0",
  ....
}
```



# Monitoring with nats-top



KubeCon



CloudNativeCon

Europe 2018

```
NATS server version 0.7.3 (uptime: 1m35s)
Server:
  Load: CPU: 61.1% Memory: 19.4M Slow Consumers: 0
  In:   Msgs: 5.3M Bytes: 16.0M Msgs/Sec: 142738.7 Bytes/Sec: 418.2K
  Out:  Msgs: 53.5M Bytes: 160.4M Msgs/Sec: 1427340.8 Bytes/Sec: 4.1M

Connections: 110
HOST          CID   NAME      SUBS  PENDING  MSGS_TO  MSGS_FROM  BYTES_TO  BYTES_FROM  LANG  VERSION  UPTIME  LAST ACTIVITY
127.0.0.1:58041 47    0         0     0         58.1K    0          174.2K     0          go    1.1.7    32s     2016-02-09 00:2
127.0.0.1:58058 64    0         0     0         33.5K    0          100.5K     0          go    1.1.7    32s     2016-02-09 00:2
127.0.0.1:58050 56    0         0     0         33.7K    0          101.1K     0          go    1.1.7    32s     2016-02-09 00:2
127.0.0.1:58070 76    0         0     0         91.6K    0          274.9K     0          go    1.1.7    32s     2016-02-09 00:2
127.0.0.1:58024 30    0         0     0         121.9K   0          365.8K     0          go    1.1.7    32s     2016-02-09 00:2
127.0.0.1:58016 22    0         0     0         32.0K    0          95.9K      0          go    1.1.7    32s     2016-02-09 00:2
127.0.0.1:58073 79    0         0     0         25.9K    0          77.7K      0          go    1.1.7    32s     2016-02-09 00:2
127.0.0.1:58007 13    0         0     0         65.6K    0          196.7K     0          go    1.1.7    32s     2016-02-09 00:2
127.0.0.1:58023 29    0         0     0         62.1K    0          186.2K     0          go    1.1.7    32s     2016-02-09 00:2
127.0.0.1:58015 21    0         0     0         39.2K    0          117.6K     0          go    1.1.7    32s     2016-02-09 00:2
127.0.0.1:58048 54    0         0     0         24.4K    0          73.3K      0          go    1.1.7    32s     2016-02-09 00:2
127.0.0.1:58085 91    0         0     0         115.4K   0          346.2K     0          go    1.1.7    32s     2016-02-09 00:2
127.0.0.1:58008 14    0         0     0         44.3K    0          132.9K     0          go    1.1.7    32s     2016-02-09 00:2
127.0.0.1:58027 33    0         0     0         62.1K    0          186.2K     0          go    1.1.7    32s     2016-02-09 00:2
127.0.0.1:58047 53    0         0     0         69.7K    0          209.1K     0          go    1.1.7    32s     2016-02-09 00:2
127.0.0.1:58013 19    0         0     0         21.2K    0          63.7K      0          go    1.1.7    32s     2016-02-09 00:2
127.0.0.1:58062 68    0         0     0         56.8K    0          170.3K     0          go    1.1.7    32s     2016-02-09 00:2
127.0.0.1:58087 93    0         0     0         23.1K    0          69.2K      0          go    1.1.7    32s     2016-02-09 00:2
127.0.0.1:58038 44    0         0     0         62.8K    0          188.5K     0          go    1.1.7    32s     2016-02-09 00:2
127.0.0.1:58081 87    0         0     0         9.8K     0          29.3K      0          go    1.1.7    32s     2016-02-09 00:2
127.0.0.1:58064 70    0         0     0         39.4K    0          118.1K     0          go    1.1.7    32s     2016-02-09 00:2
127.0.0.1:58001 07    0         0     0         10.8K    0          50.5K      0          go    1.1.7    32s     2016-02-09 00:2
```

<https://github.com/nats-io/nats-top>



# NATS and Prometheus

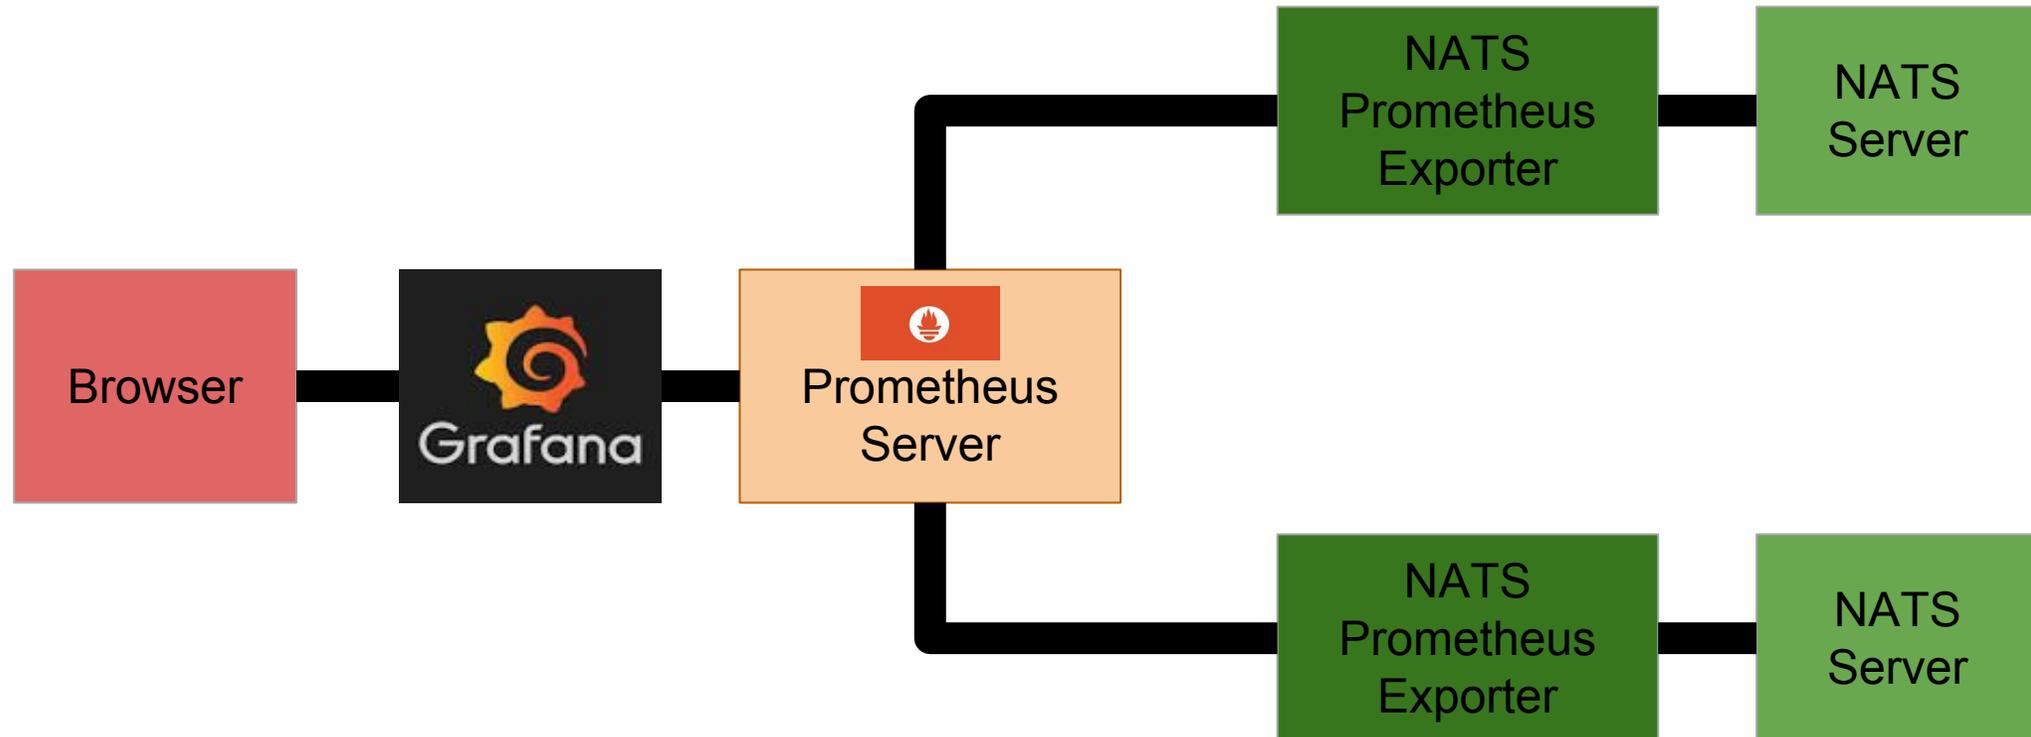


KubeCon



CloudNativeCon

Europe 2018





# NATS / Prometheus / Grafana



KubeCon



CloudNativeCon

Europe 2018

Use a visualization tool with Prometheus and enable alerts or create rules to take action on NATS monitoring data.



<https://github.com/nats-io/prometheus-nats-exporter>



**KubeCon**



**CloudNativeCon**

Europe 2018

# The Future





# Future Plans



KubeCon



CloudNativeCon

Europe 2018

- Extended client support for **nkeys**, public key signatures utilizing Ed25519
- Cloud provider integrations
- **Extensible Security** to plug in common or custom authentication and authorization systems
- Features to connect and bridge clusters of NATS servers to create a NATS deployment at a global scale, providing connectivity between millions, if not **billions of clients**
- Continued work on integrations with other CNCF projects including additional support for Prometheus and Kubernetes



# Contributing



KubeCon



CloudNativeCon

Europe 2018

We welcome contributions of all kinds. Some ways to contribute include:

- Highlight your NATS usage or insights on the NATS blog
- Fix a bug
- Add, fix, or clarify documentation
- Propose or add a feature through a Github PR
- Present your NATS project at meetups

Read more at <https://nats.io/documentation/contributing>



**KubeCon**



**CloudNativeCon**

Europe 2018

**Thank You!**





**KubeCon**



**CloudNativeCon**

Europe 2018

# Questions?

