

Weave Net

Five years with no central control.

FOSDEM 2020

Bryan Boreham
<https://weave.works>



@bboreham
@weaveworks

Bryan Boreham

Lead on Weave Net since 2015.

Project member of Kubernetes,
CNI, Cortex, Scope, ...

Not a networking expert.



Weave Net

- Open Source container network
- Easy to install; runs anywhere*
- No “Enterprise Version”



What is a “container network”?

**“There’s 🖐️ no 🖐️ such 🖐️ thing 🖐️ as 🖐️
container 🖐️ networking”**

It’s just networking. Anyone who tells you differently is either trying to sell you something or doesn’t understand how networking works.



Justin Garrison

Sep 26, 2016



Justin Garrison

@rothgar

Replying to @capileigh and @resouer

I use weave (via kubeadm) in my home metal cluster. I don't have a ton of experience debugging it but that's because it has just worked for me.

11:42 PM · May 20, 2018



What is a “container network”

Containers give you isolation.

- Each container runs in its own network namespace.

How do these network namespaces talk to each other?

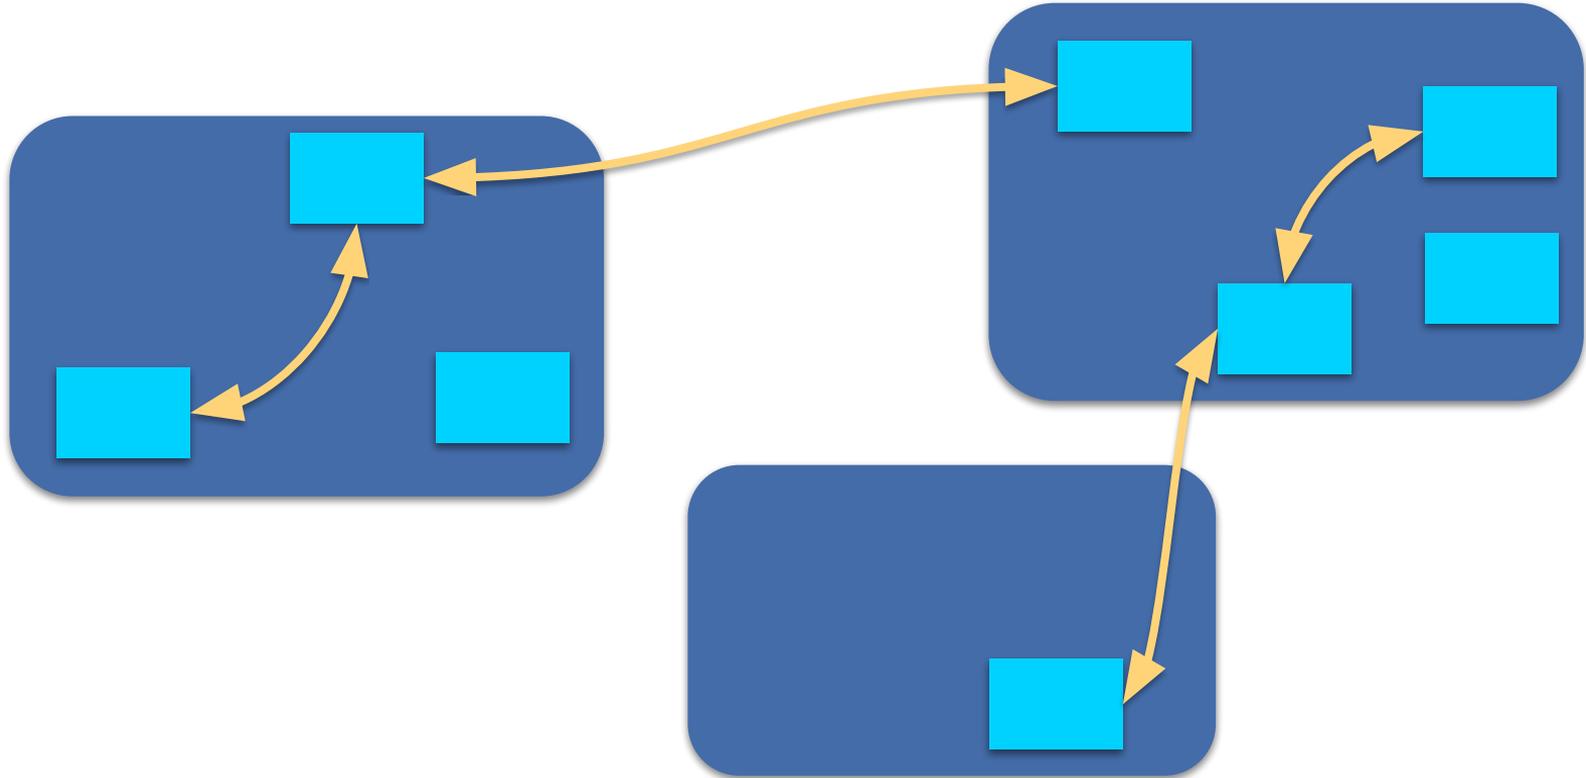
- That’s a container network.



Let's look at how it works



Container network model



Matthew Sackman

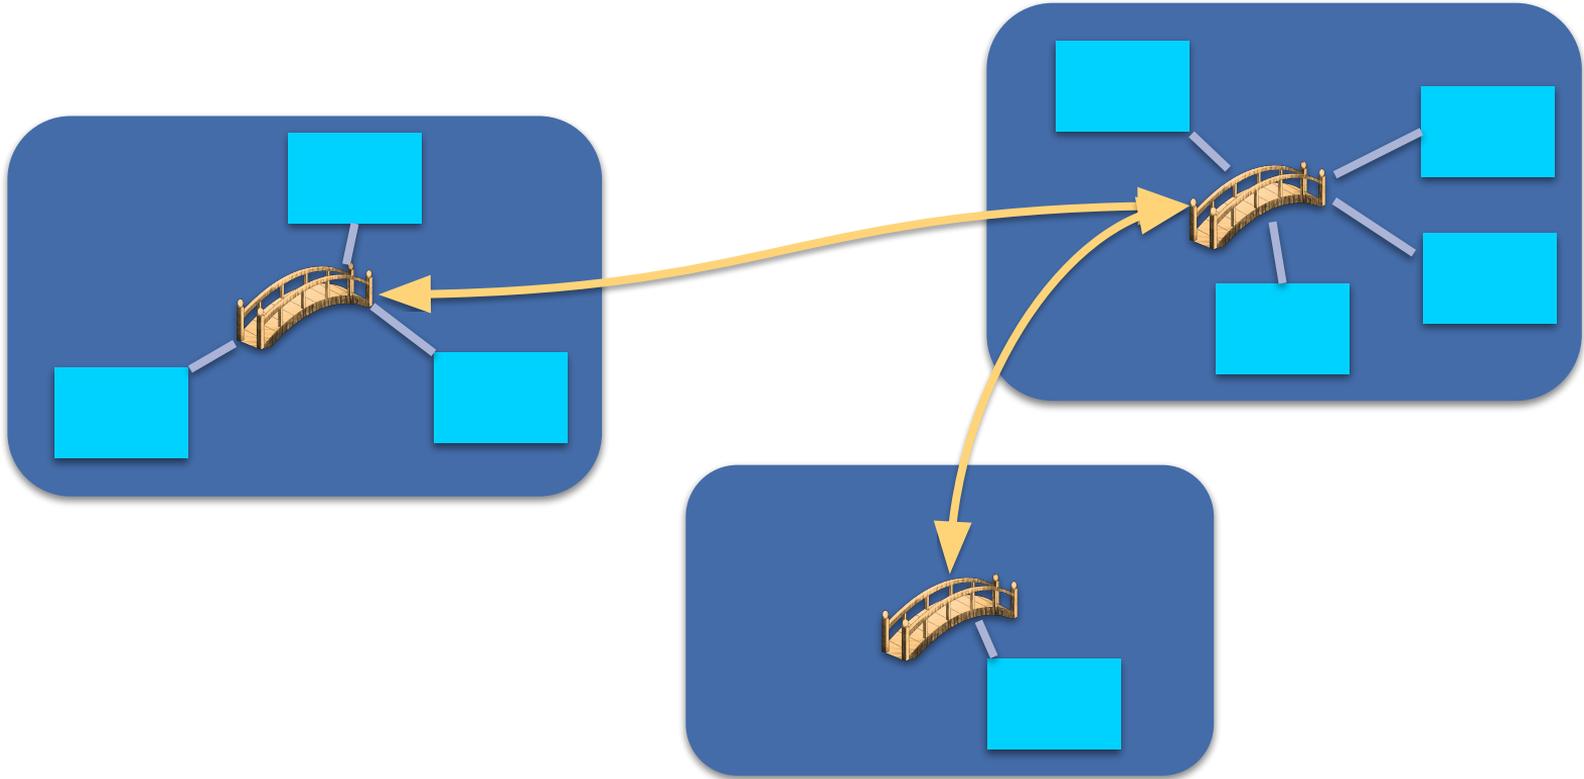
Ex-RabbitMQ, Erlang expert.

Wrote the first version of Weave Net.

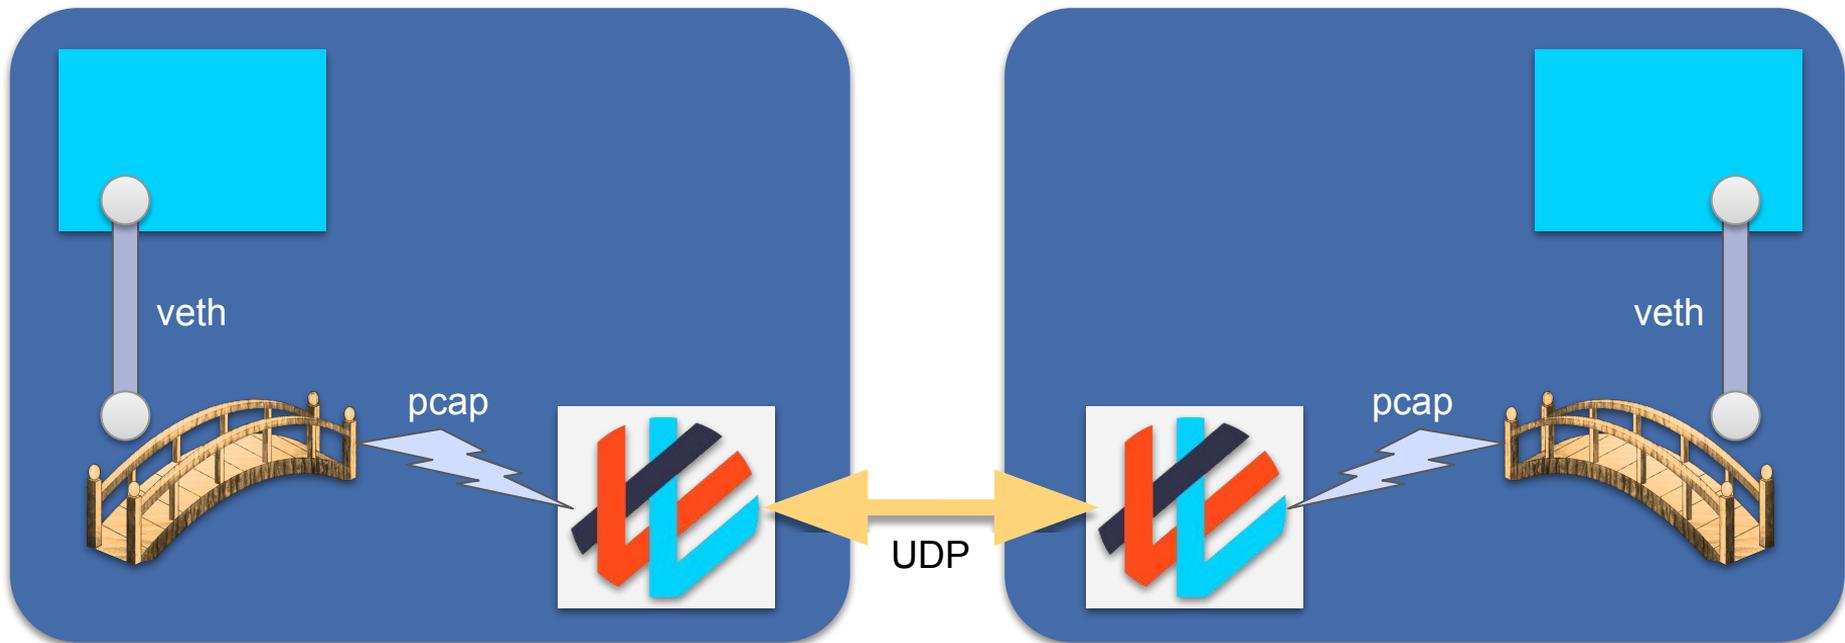
3,400 lines of Go



Containers with bridges



Weave Net 1.0



Distributed Ethernet Switch*

Weave Net daemon learns where MACs come from

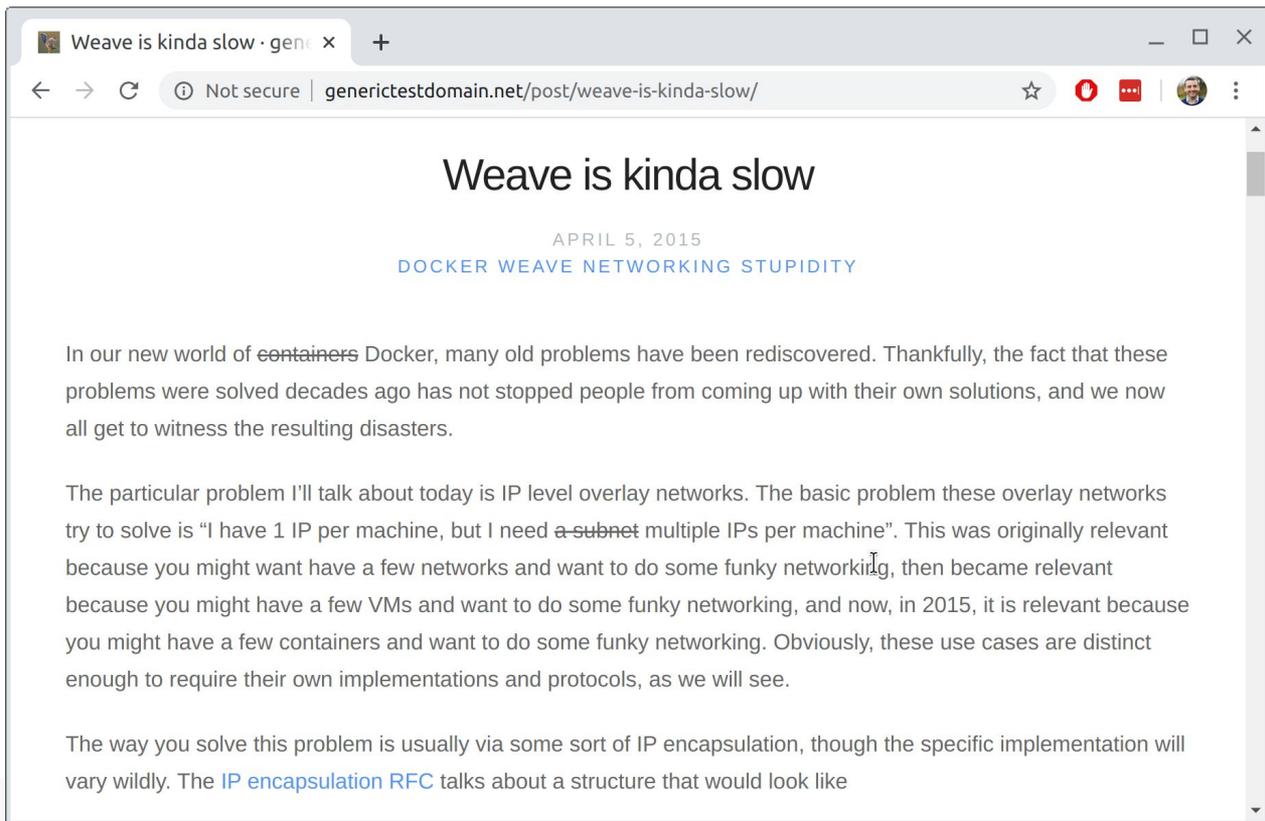
- when it sees the first packet from that MAC.

Thus, it knows where to send each packet**.

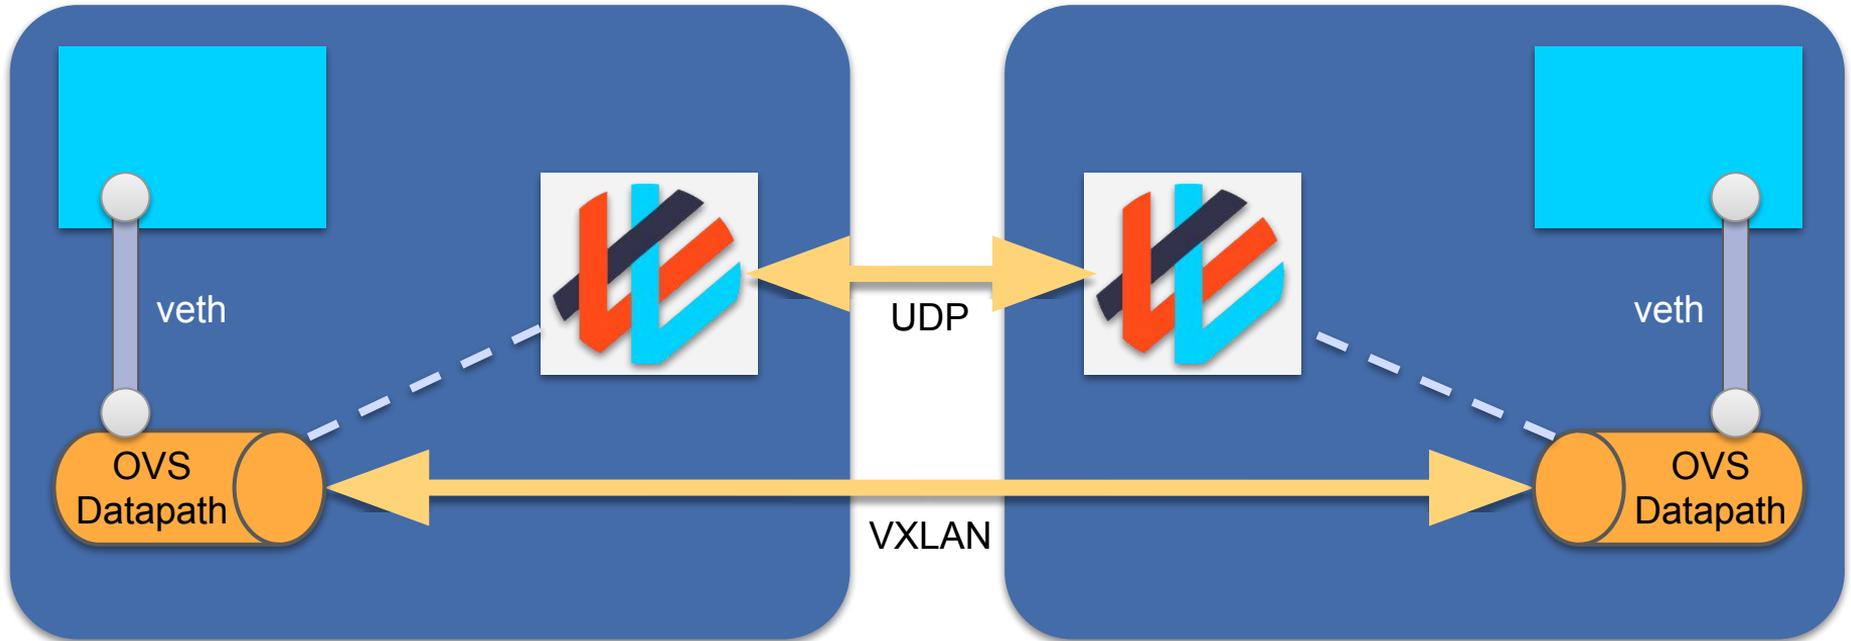
If it doesn't know where a MAC comes from?

- send it everywhere!

“Weave is kinda slow”



Weave Net 1.2 “Fast Data Path”



David Wragg

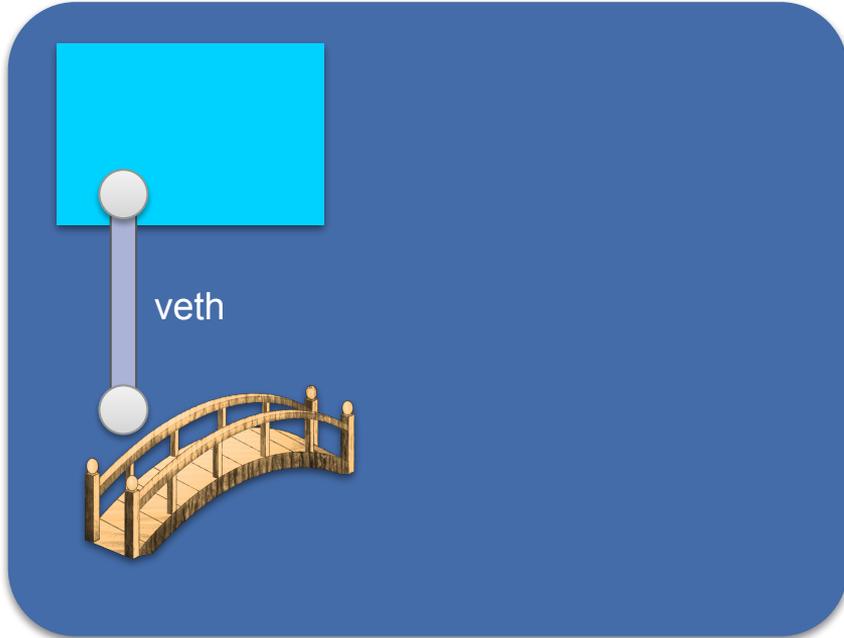
Ex-Pivotal

Implementer of the “fast data path”

Now at Cloudflare



How to set up all the devices?



Jérôme Petazzoni

Software-Defined Networking tools for LXC (Linux Containers)

86 commits

1 branch

0 packages

0 releases

38 contributors

Tree: 08d183f6c3

New pull request

Create new file

Upload files

Find file



Jérôme Petazzoni Merge pull request #93 from hugoduncan/patch-1

Latest commit 08

- .gitignore Ignore *- files.
- LICENSE Add Apache 2.0 license. (I hope it is OK with everyone!)
- README.md Update README to not suggest an lxc dependency
- pipework fix invalid sh syntax

5 years ago

README.md

Pipework

Software-Defined Networking for Linux Containers

Pipework lets you connect together containers in arbitrarily complex scenarios. Pipework uses cgroups and namespace and works with "plain" LXC containers (created with `lxc-start`), and with the awesome [Docker](#).



The weave script

Tree: ce631e7f28 **weave / weave** Find file Copy path

 **rade** stricter param count checking ce631e7 on 24 Oct 2014

1 contributor

Executable File **350 lines (316 sloc)** 10.9 KB Raw Blame History

```
1 #!/bin/sh
2 set -e
3
4 usage() {
5     echo "Usage:"
6     echo "weave launch <cidr> [-password <password>]"
7     echo "weave run <cidr> <docker run args> ..."
8     echo "weave start <cidr> <container_id>"
9     echo "weave attach <cidr> <container_id>"
10    echo "weave detach <cidr> <container_id>"
11    echo "weave expose <cidr>"

```

Branch: 1.9 **weave / weave**

 **deitch** First cut of adding hairin functionality to weave

19 contributors

Executable File **2329 L**

Branch: 2.6 **weave / weave**

 **brb** Merge pull request #3388 from myENA/feature/skip-nat-via-weave-expose

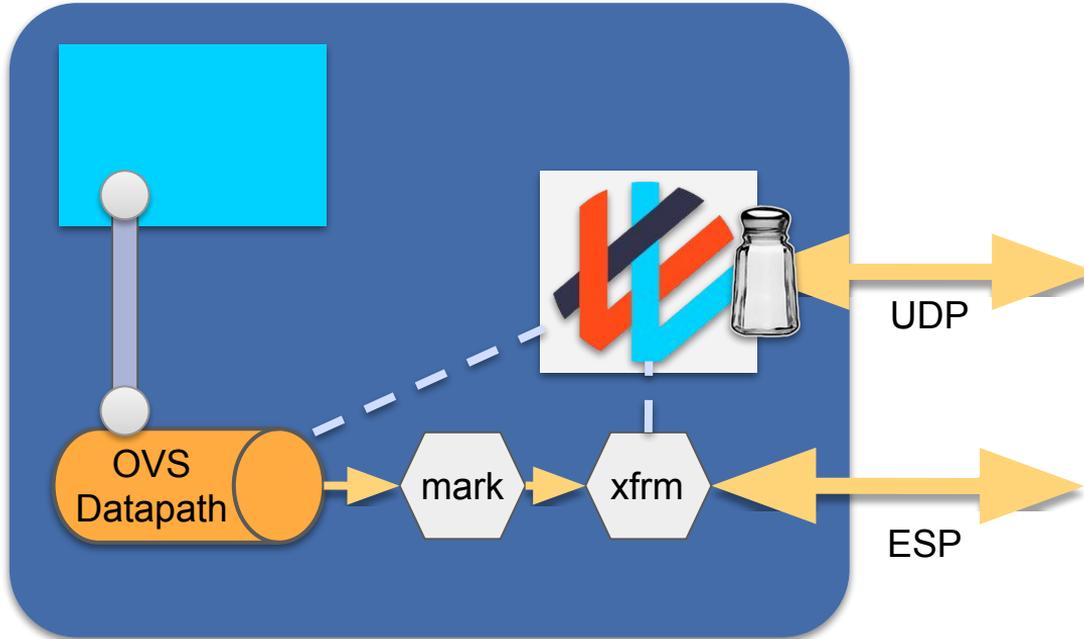
25 contributors

Executable File **1591 lines (1435 sloc)** 51 KB

```
1 #!/bin/sh
2 set -e
3

```

Encryption



Martynas Pumputis

Implementer of Weave Net XFRM encryption.

Kernel fixes for conntrack race conditions, etc.

Now at Isovalent (Cillium)



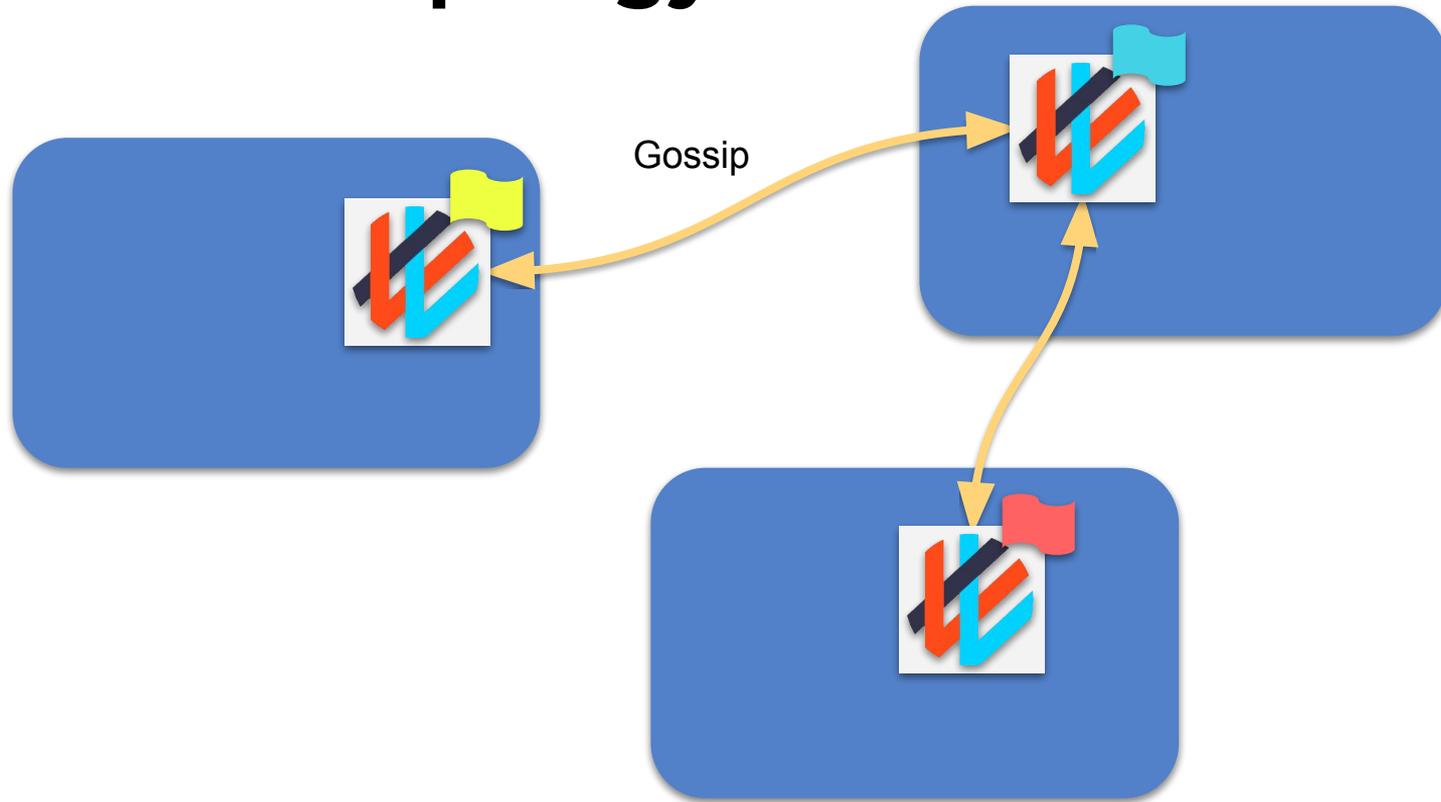
Multicast

Weave Net handles multicast

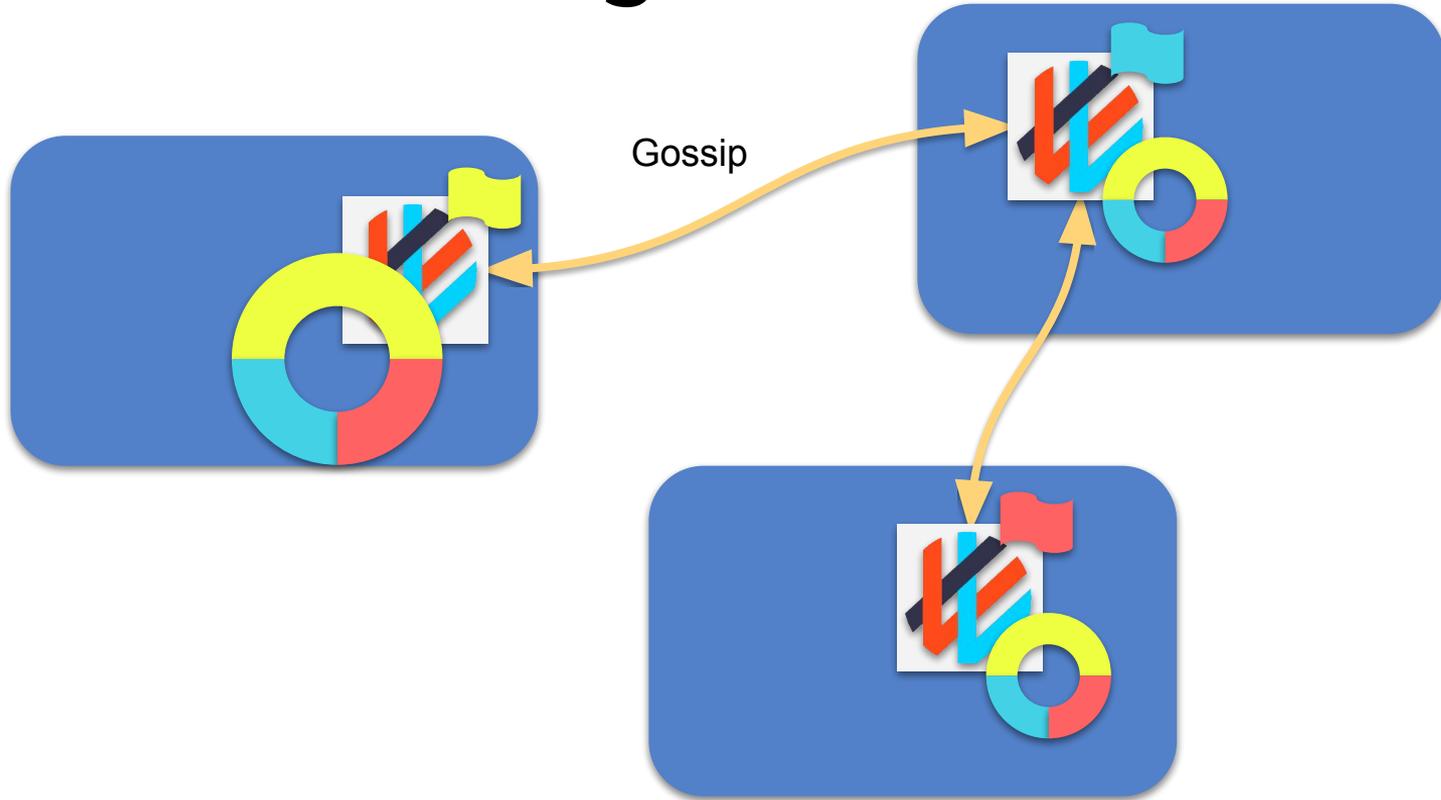
- via the “send the packet everywhere” logic.

Peers and Topology

Peers and topology



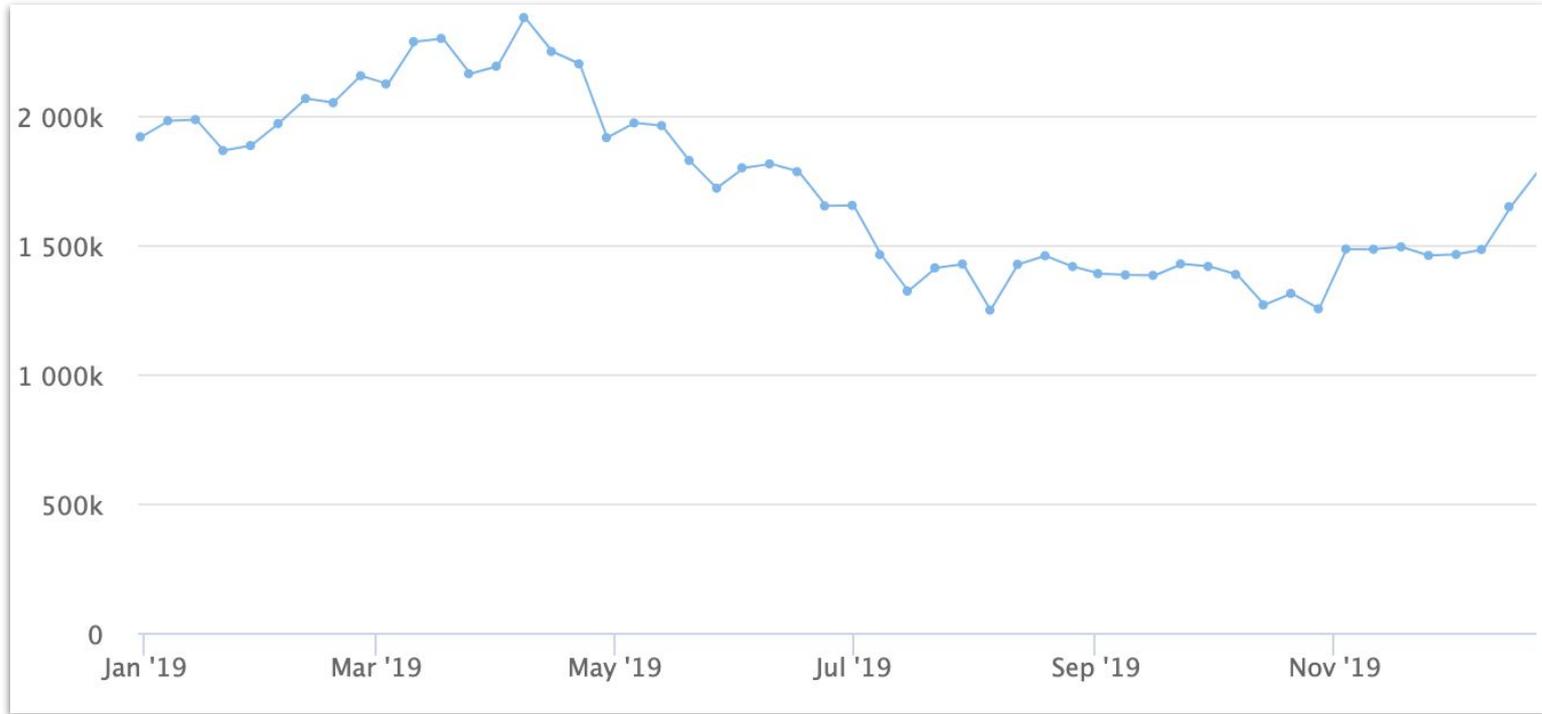
IP Address Management



Community



Weave Net installs per week



Lots of requests, very few PRs

[Profile Picture] commented on 31 Aug 2018 • edited ▾

Please make `NO_MASQ_LOCAL=1` the default setting.

unavie commented on 22 Apr 2019

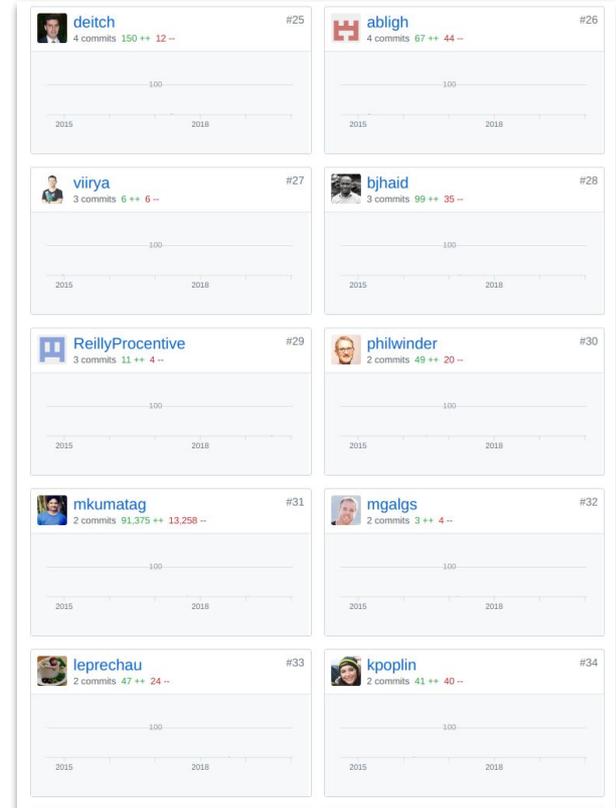
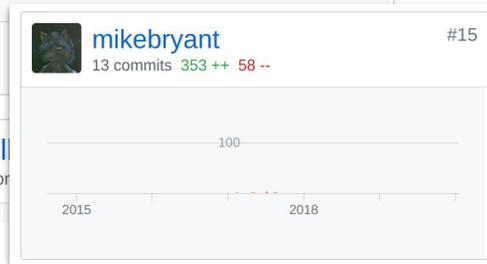
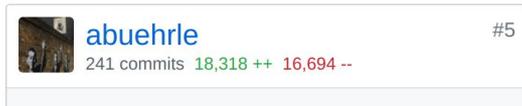
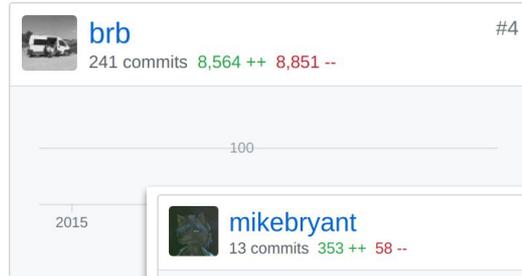
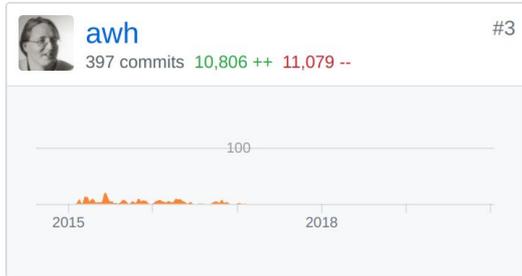
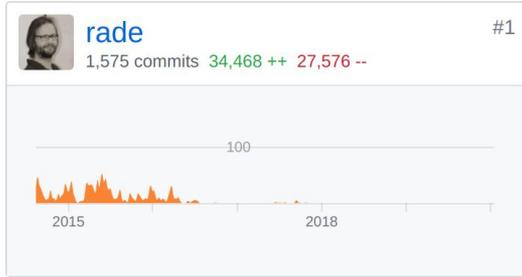
Lost 2 days on this stuff. Indeed, please make this the default.

👍 1

virtuallxs commented 9 days ago

I'm wondering, is there any specific reason this isn't the default? Does this being on break something else, or cause something to behave unexpectedly?

Mostly paid contributors



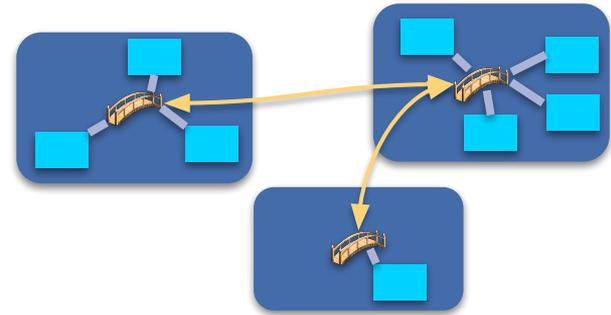
Kubernetes

Kubernetes

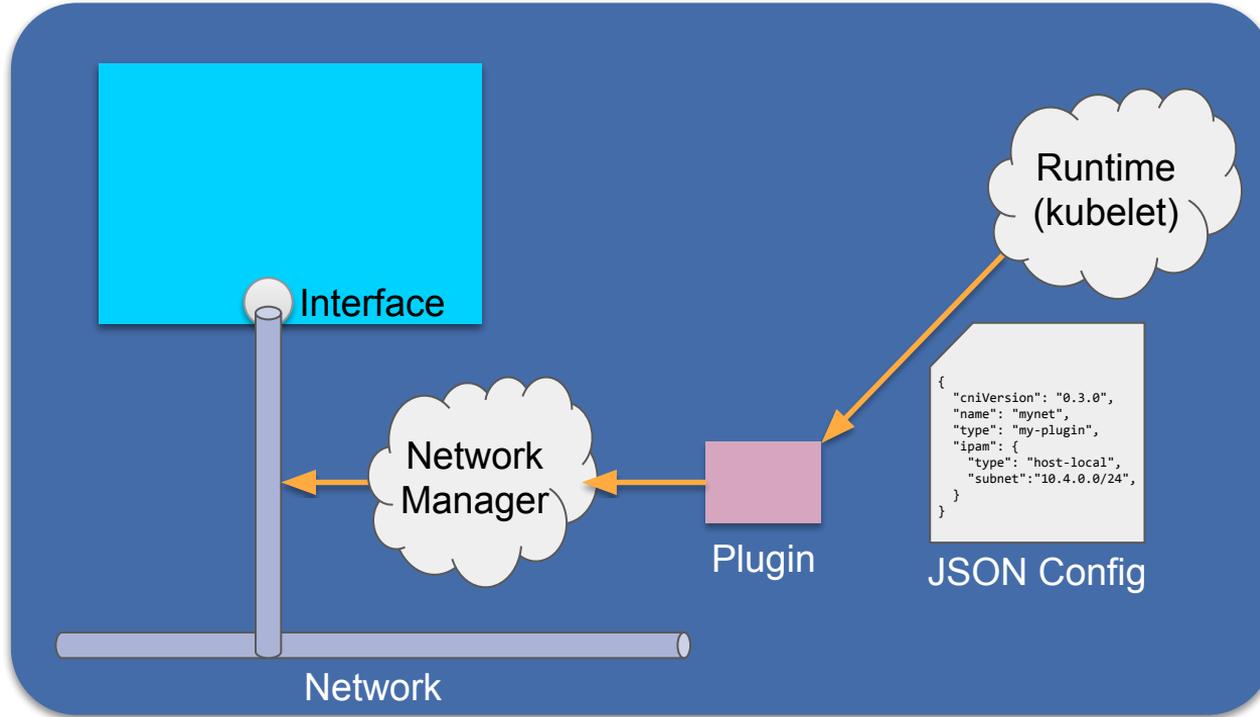
Mandates NAT-free network between “pods”.

3rd-party pod networks.

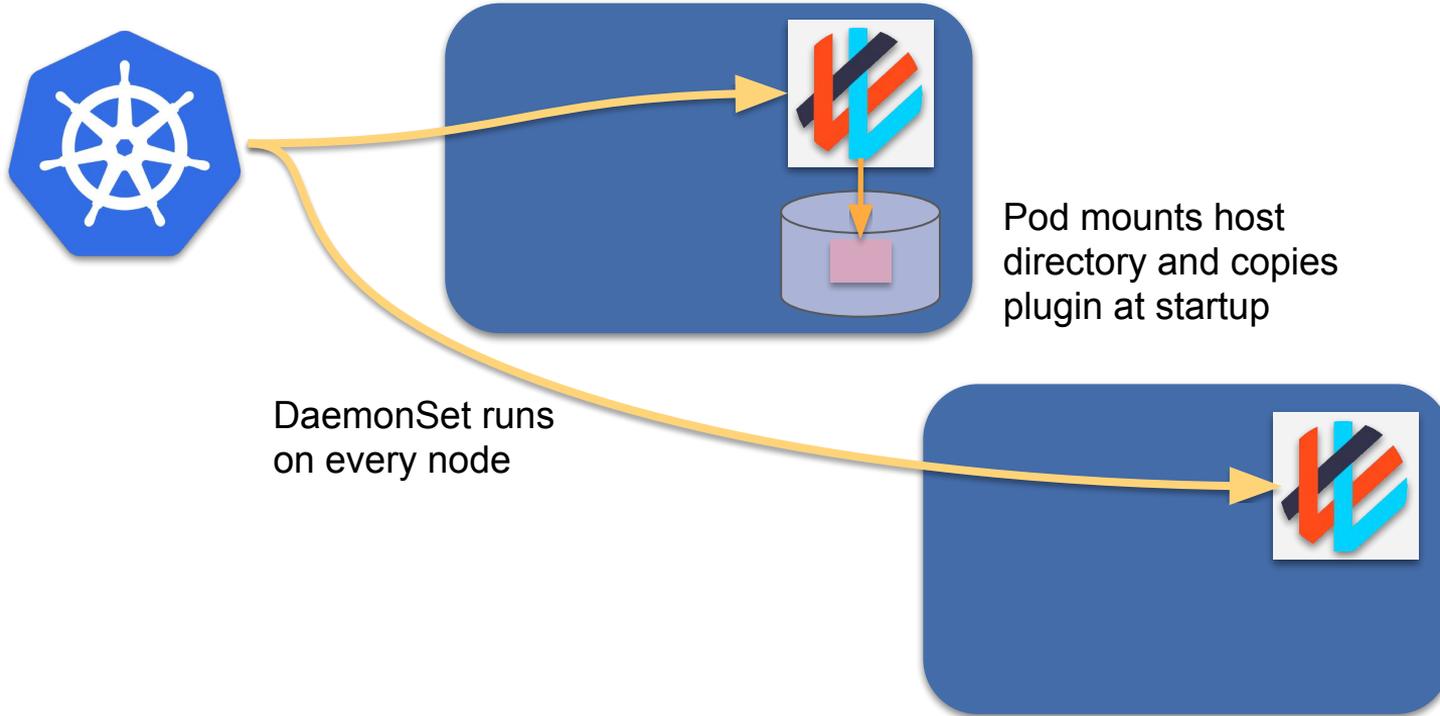
Rkt, from CoreOS, has a simple ‘exec’ model to add a network.



CNI - the Container Network Interface



Installing via DaemonSet



End of main content

Things I didn't cover

Kubernetes Network Policy

Launch modes

Scalability

Service Management / Service Discovery

Bug bounty programme

Questions?

Bryan Boreham
<https://weave.works>



@bboreham
@weaveworks

