

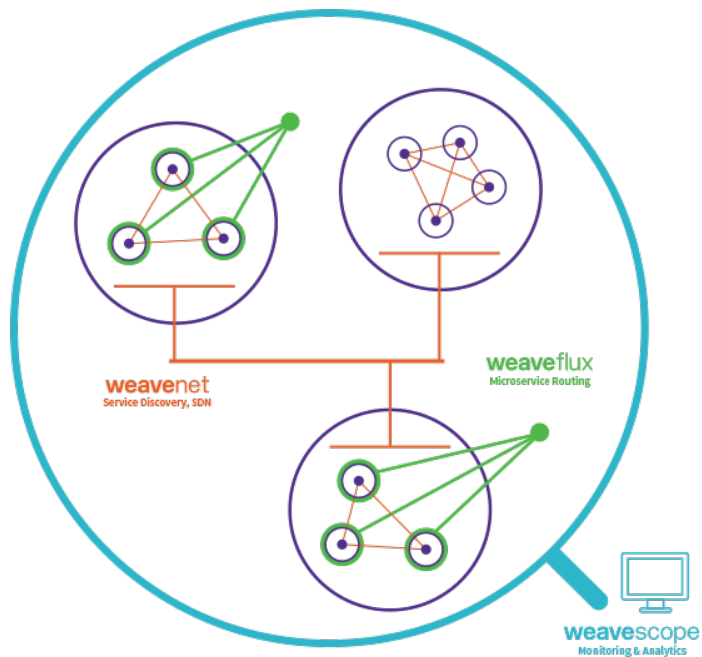


Open vSwitch

Martynas Pumputis

Weaveworks

OVS for Containers with Weave Net



<https://weave.works>

<https://github.com/weaveworks>



Problem Statement

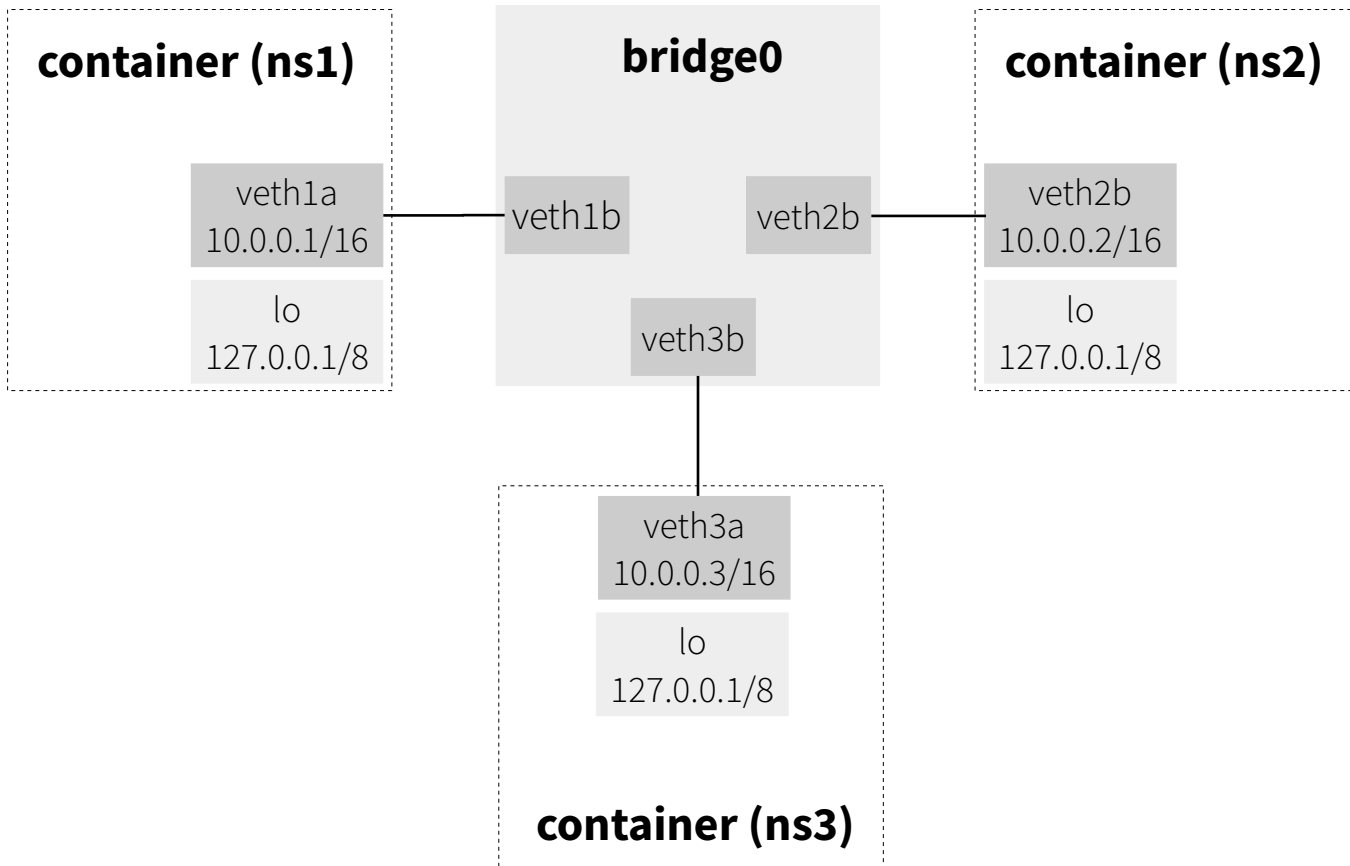
container (ns1)

```
lo  
127.0.0.1/8
```

```
// include/net/net_namespace.h  
struct net {  
    struct user_namespace *user_ns;  
    struct netns_ipv4      ipv4;  
    struct net_device      *loopback_dev;  
    ...  
}
```


```
// include/linux/netdevice.h  
struct net_device {  
    struct net *nd_net;  
    ...  
}
```

```
// include/net/sock.h  
struct sock {  
    struct net *skc_net;  
    ...  
}
```





```
C1 $ curl http://192.168.0.2:80 # C4  
C1 $ curl http://192.168.0.2:81 # C5
```

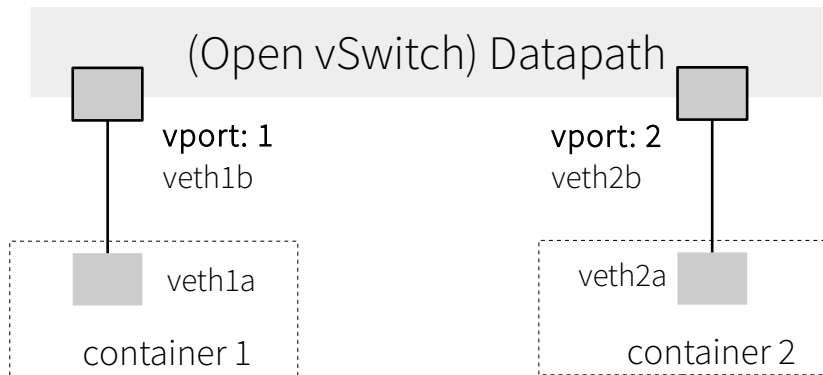


```
C1 $ curl http://10.0.0.4:80 # C4  
C1 $ curl http://10.0.0.5:80 # C5
```

Fast multi-host overlay network for containers

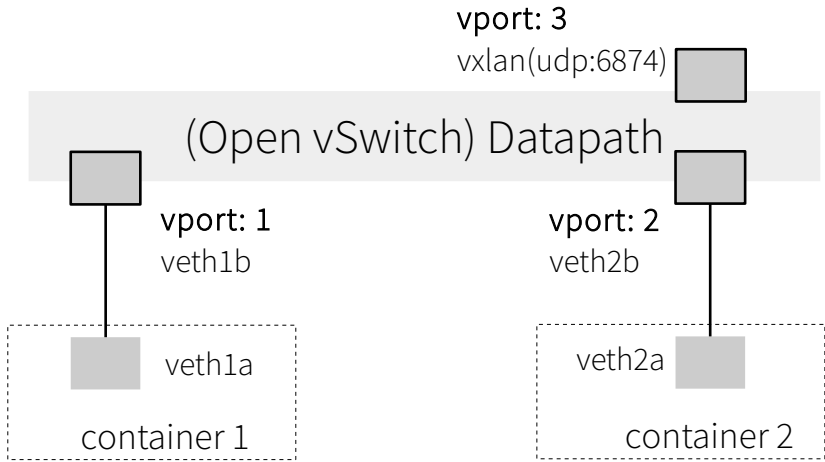
Approach

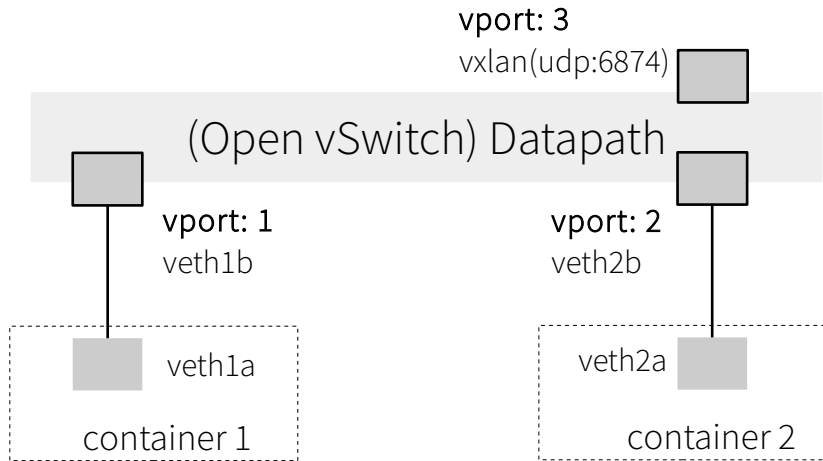
Weave Net Router



```
func handleMiss(packet []byte, k odp.FlowKeys) {  
    if sendToMAC[k.SrcMAC] == nil {  
        sendToMAC[k.SrcMAC] = k.inVportID  
    }  
    if outVport := sendToMAC[k.DstMAC]; outVport != nil {  
        send(outVport, packet)  
    } else {  
        broadcast(packet, k)  
    }  
}
```

flow key: in_port(1), eth(src=veth1a, dst=veth2a)
action: out_port(2)





```

func handleMiss(packet []byte, k odp.FlowKeys) {
    srcPeer, dstPeer := extractPeers(k)
    if dstPeer != ourself {
        relay(packet, k)
        return
    }
    if getPeer[k.DstMAC] == ourself {
        inject(packet, k)
        return
    }
    broadcast(packet, k)
}


```

flow key: in_port(3), eth(src=veth1a, dst=veth5a),
tunnel(id=..., ipv4src=192.168.1.2, ipv4dst=192.168.1.1)
action: out_port(1)

```
import "github.com/weaveworks/go-odp"  
  
<..>  
  
dpif, err := odp.NewDpif()  
dp, err := dpif.LookupDatapath(name)  
vport, err := dp.CreateVport(odp.VxlanVportSpec{"foobar-vxlan", 6785})  
flow := odp.NewFlowSpec()  
flow.AddKey(..)  
flow.AddAction(..)  
err = dp.CreateFlow(flow)
```

github.com/weaveworks/go-odp

Effortless Eventual Consistency with Weave Mesh – Peter Bourgon, Matthias Radestock



github.com/weaveworks/mesh

- * Docker Plugin (libnetwork) and Docker Proxy
- * CNI (Kubernetes, Mesos)

- * DNS, IPAM w/o a consistent store
- * Encryption
- * Multicast

github.com/weaveworks/weave

Problems

- Missing contrack support (fixed in 4.2)
- Limited MTU of vxlan devices w/o an underlying device (fixed in 4.5)

The image shows the WeaveScope web interface. At the top left is the WeaveScope logo. To its right are three tabs: 'APPLICATIONS BY NAME', 'CONTAINERS BY HOSTNAME BY IMAGE', and 'HOSTS'. The main area displays a network diagram with nodes representing containers. The nodes are labeled: 'docker_app1_1 weave-scope-demo', 'docker_db1_1 weave-scope-demo', 'docker_app2_1 weave-scope-demo', 'docker_lb1_1 weave-scope-demo', and 'docker_db3_1 weave-scope-demo'. Lines connect these nodes, representing network links. In the bottom left corner, there are two buttons: 'Stopped containers hidden SHOW' and 'System containers hidden SHOW'. Below these buttons, it says '7 nodes (10 filtered)'. On the right side, a detailed view of the 'docker_db1_1' container is shown. The title bar is green and contains the container name 'docker_db1_1' and the namespace 'weave-scope-demo'. Below the title bar are three icons: a play button, a refresh button, and a close button. The main content area of the container view is white and contains the following information:


CONTAINER IMAGE "PETERBOURGON/TNS-DB:LATEST"
Image ID e4fb99a3c9adf3f1cbb1cfc07...

CONTAINER "DOCKER_DB1_1"
State running
ID b53e95840686276bbf056f80...
Image ID e4fb99a3c9adf3f1cbb1cfc07...
Created 18 Nov 15 22:47 UTC
Command /db db1 db2 db3
IP Address 172.17.0.1
Label "com.docker.co..." 586b47de98dc6480a1c85646...
Label "com.docker.co..." 1
Label "com.docker.co..." False
Label "com.docker.co..." docker
Label "com.docker.co..." db1
Label "com.docker.co..." 1.5.1
Memory Usage 1.33 MB
CPU Usage 0.11%

At the bottom right of the interface, it says 'Version 0.10.0 [Report an issue](#)'.

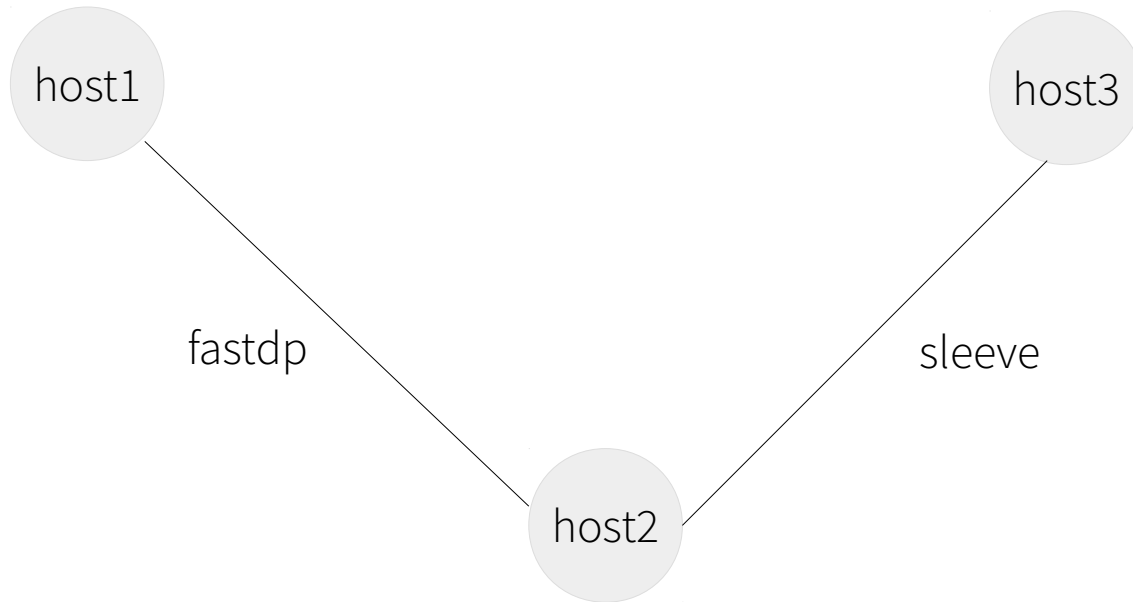
Evaluation

Throughput (Gbits/sec)



iperf3 on AWS c3.8xlarge, Ubuntu 16.04 LTS (4.4.0), Weave Net 1.8.0

Demo



Questions?

@weaveworks / @martyns