


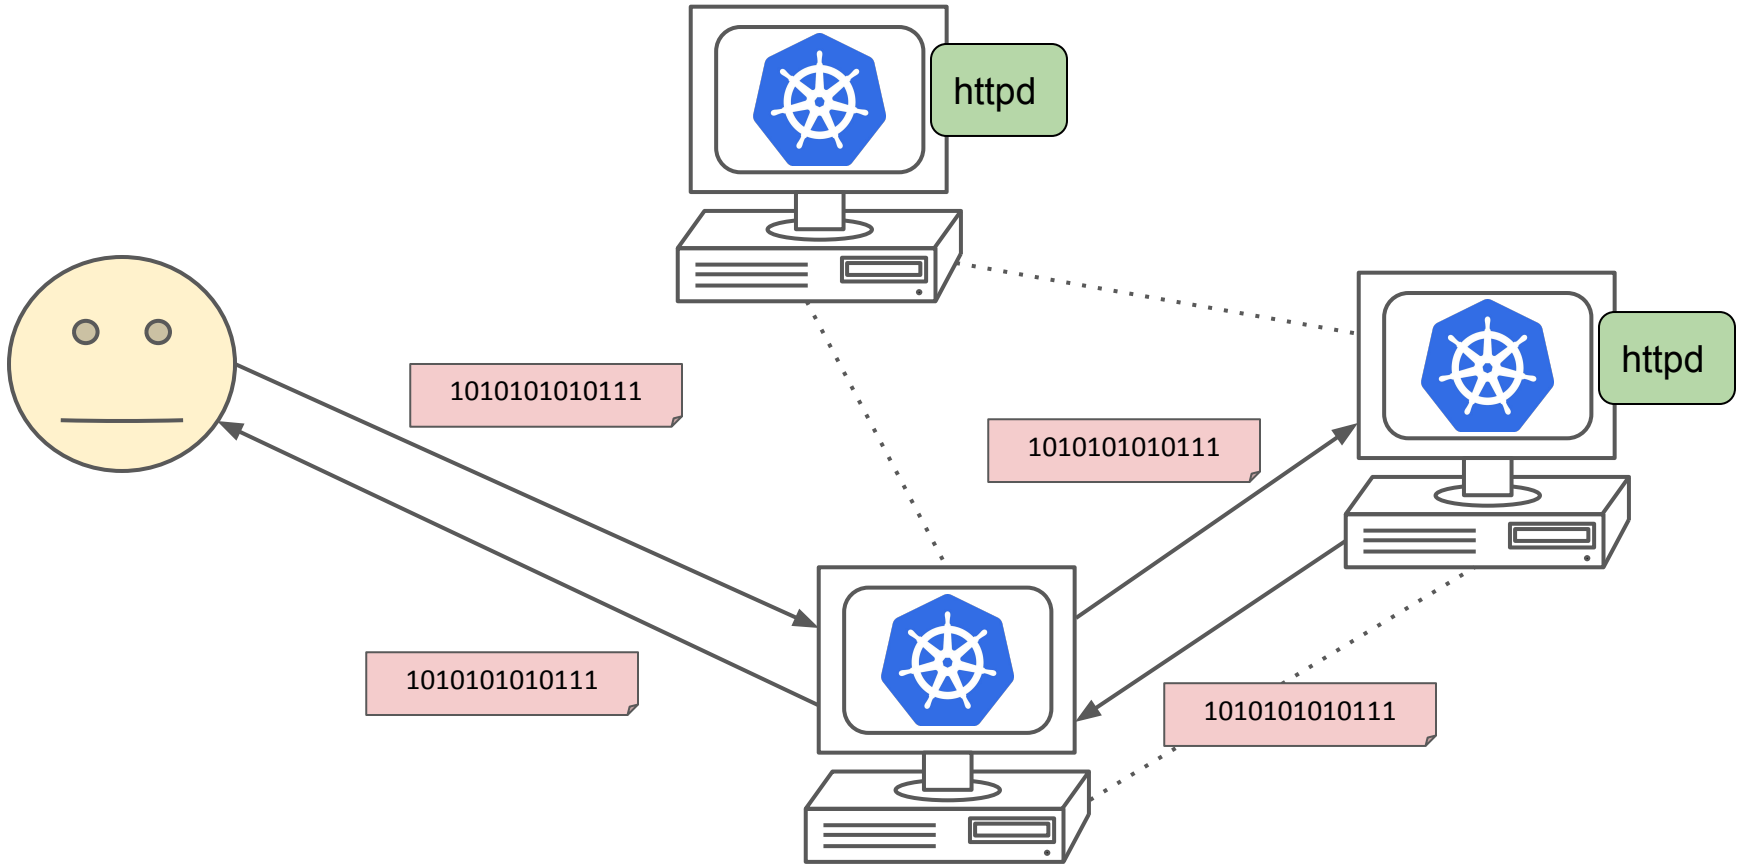
October 28, 2020

# North-South Load Balancing of Kubernetes Services with eBPF/XDP



Martynas Pumputis (Isovalent)





-A KUBE-FORWARD -d 10.217.0.0/16 -m comment --comment "kubernetes forwarding contract pod destination rule" -m conntrack --cstate RELATED,ESTABLISHED -j ACCEPT

-A KUBE-SERVICES -d 10.99.38.155/32 -p tcp -m comment --comment "default/nginx-59: has no endpoints" -m tcp --dport 80 -j REJECT --reject-with icmp-port-unreachable

-A KUBE-SERVICES -d 10.96.61.252/32 -p tcp -m comment --comment "default/nginx-64: has no endpoints" -m tcp --dport 80 -j REJECT --reject-with icmp-port-unreachable

-A KUBE-SERVICES -d 10.104.166.10/32 -p tcp -m comment --comment "default/nginx-67: has no endpoints" -m tcp --dport 80 -j REJECT --reject-with icmp-port-unreachable

-A KUBE-SERVICES -d 10.98.85.41/32 -p tcp -m comment --comment "default/nginx-9: has no endpoints" -m tcp --dport 80 -j REJECT --reject-with icmp-port-unreachable

-A KUBE-SERVICES -d 10.97.138.144/32 -p tcp -m comment --comment "default/nginx-17: has no endpoints" -m tcp --dport 80 -j REJECT --reject-with icmp-port-unreachable

-A KUBE-SERVICES -d 10.106.49.80/32 -p tcp -m comment --comment "default/nginx-37: has no endpoints" -m tcp --dport 80 -j REJECT --reject-with icmp-port-unreachable

-A KUBE-SERVICES -d 10.104.164.205/32 -p tcp -m comment --comment "default/nginx-5: has no endpoints" -m tcp --dport 80 -j REJECT --reject-with icmp-port-unreachable

-A KUBE-SERVICES -d 10.104.25.150/32 -p tcp -m comment --comment "default/nginx-19: has no endpoints" -m tcp --dport 80 -j REJECT --reject-with icmp-port-unreachable

-A KUBE-SERVICES -d 10.106.234.213/32 -p tcp -m comment --comment "default/nginx-88: has no endpoints" -m tcp --dport 80 -j REJECT --reject-with icmp-port-unreachable

-A KUBE-SERVICES -d 10.109.209.136/32 -p tcp -m comment --comment "default/nginx-33: has no endpoints" -m tcp --dport 80 -j REJECT --reject-with icmp-port-unreachable

-A KUBE-SERVICES -d 10.106.196.105/32 -p tcp -m comment --comment "default/nginx-49: has no endpoints" -m tcp --dport 80 -j REJECT --reject-with icmp-port-unreachable

-A KUBE-SERVICES -d 10.111.101.6/32 -p tcp -m comment --comment "default/nginx-53: has no endpoints" -m tcp --dport 80 -j REJECT --reject-with icmp-port-unreachable

-A KUBE-SERVICES -d 10.110.226.230/32 -p tcp -m comment --comment "default/nginx-79: has no endpoints" -m tcp --dport 80 -j REJECT --reject-with icmp-port-unreachable

-A KUBE-SERVICES -d 10.98.99.136/32 -p tcp -m comment --comment "default/nginx-6: has no endpoints" -m tcp --dport 80 -j REJECT --reject-with icmp-port-unreachable

-A KUBE-SERVICES -d 10.99.75.233/32 -p tcp -m comment --comment "default/nginx-7: has no endpoints" -m tcp --dport 80 -j REJECT --reject-with icmp-port-unreachable

-A KUBE-SERVICES -d 10.108.41.202/32 -p tcp -m comment --comment "default/nginx-14: has no endpoints" -m tcp --dport 80 -j REJECT --reject-with icmp-port-unreachable

-A KUBE-SERVICES -d 10.97.36.249/32 -p tcp -m comment --comment "default/nginx-99: has no endpoints" -m tcp --dport 80 -j REJECT --reject-with icmp-port-unreachable

-A KUBE-SERVICES -d 10.98.213.37/32 -p tcp -m comment --comment "default/nginx-77: has no endpoints" -m tcp --dport 80 -j REJECT --reject-with icmp-port-unreachable

-A KUBE-SERVICES -d 10.107.229.31/32 -p tcp -m comment --comment "default/nginx-92: has no endpoints" -m tcp --dport 80 -j REJECT --reject-with icmp-port-unreachable

-A KUBE-SERVICES -d 10.98.64.251/32 -p tcp -m comment --comment "default/nginx-16: has no endpoints" -m tcp --dport 80 -j REJECT --reject-with icmp-port-unreachable

-A KUBE-SERVICES -d 10.101.88.159/32 -p tcp -m comment --comment "default/nginx-31: has no endpoints" -m tcp --dport 80 -j REJECT --reject-with icmp-port-unreachable

-A KUBE-SERVICES -d 10.105.71.74/32 -p tcp -m comment --comment "default/nginx-41: has no endpoints" -m tcp --dport 80 -j REJECT --reject-with icmp-port-unreachable

-A KUBE-SERVICES -d 10.108.92.226/32 -p tcp -m comment --comment "default/nginx-63: has no endpoints" -m tcp --dport 80 -j REJECT --reject-with icmp-port-unreachable

-A KUBE-SERVICES -d 10.109.252.234/32 -p tcp -m comment --comment "default/nginx-18: has no endpoints" -m tcp --dport 80 -j REJECT --reject-with icmp-port-unreachable

-A KUBE-SERVICES -d 10.104.118.66/32 -p tcp -m comment --comment "default/nginx-30: has no endpoints" -m tcp --dport 80 -j REJECT --reject-with icmp-port-unreachable

-A KUBE-SERVICES -d 10.106.224.55/32 -p tcp -m comment --comment "default/nginx-83: has no endpoints" -m tcp --dport 80 -j REJECT --reject-with icmp-port-unreachable

-A KUBE-SERVICES -d 10.109.16.199/32 -p tcp -m comment --comment "default/nginx-100: has no endpoints" -m tcp --dport 80 -j REJECT --reject-with icmp-port-unreachable

-A KUBE-SERVICES -d 10.109.231.213/32 -p tcp -m comment --comment "default/nginx-61: has no endpoints" -m tcp --dport 80 -j REJECT --reject-with icmp-port-unreachable

-A KUBE-SERVICES -d 10.98.27.250/32 -p tcp -m comment --comment "default/nginx-95: has no endpoints" -m tcp --dport 80 -j REJECT --reject-with icmp-port-unreachable

-A KUBE-SERVICES -d 10.105.42.108/32 -p tcp -m comment --comment "default/nginx-12: has no endpoints" -m tcp --dport 80 -j REJECT --reject-with icmp-port-unreachable

-A KUBE-SERVICES -d 10.99.35.236/32 -p tcp -m comment --comment "default/nginx-20: has no endpoints" -m tcp --dport 80 -j REJECT --reject-with icmp-port-unreachable

-A KUBE-SERVICES -d 10.111.42.123/32 -p tcp -m comment --comment "default/nginx-21: has no endpoints" -m tcp --dport 80 -j REJECT --reject-with icmp-port-unreachable

-A KUBE-SERVICES -d 10.99.47.225/32 -p tcp -m comment --comment "default/nginx-22: has no endpoints" -m tcp --dport 80 -j REJECT --reject-with icmp-port-unreachable

-A KUBE-SERVICES -d 10.104.184.242/32 -p tcp -m comment --comment "default/nginx-51: has no endpoints" -m tcp --dport 80 -j REJECT --reject-with icmp-port-unreachable

-A KUBE-SERVICES -d 10.98.77.93/32 -p tcp -m comment --comment "default/nginx-68: has no endpoints" -m tcp --dport 80 -j REJECT --reject-with icmp-port-unreachable

-A KUBE-SERVICES -d 10.110.169.113/32 -p tcp -m comment --comment "default/nginx-72: has no endpoints" -m tcp --dport 80 -j REJECT --reject-with icmp-port-unreachable

-A KUBE-SERVICES -d 10.100.231.169/32 -p tcp -m comment --comment "default/nginx-90: has no endpoints" -m tcp --dport 80 -j REJECT --reject-with icmp-port-unreachable

-A KUBE-SERVICES -d 10.96.58.51/32 -p tcp -m comment --comment "default/nginx-4: has no endpoints" -m tcp --dport 80 -j REJECT --reject-with icmp-port-unreachable

-A KUBE-SERVICES -d 10.101.132.61/32 -p tcp -m comment --comment "default/nginx-25: has no endpoints" -m tcp --dport 80 -j REJECT --reject-with icmp-port-unreachable

-A KUBE-SERVICES -d 10.100.64.242/32 -p tcp -m comment --comment "default/nginx-39: has no endpoints" -m tcp --dport 80 -j REJECT --reject-with icmp-port-unreachable

-A KUBE-SERVICES -d 10.111.154.81/32 -p tcp -m comment --comment "default/nginx-50: has no endpoints" -m tcp --dport 80 -j REJECT --reject-with icmp-port-unreachable

-A KUBE-SERVICES -d 10.100.179.151/32 -p tcp -m comment --comment "default/nginx-96: has no endpoints" -m tcp --dport 80 -j REJECT --reject-with icmp-port-unreachable

-A KUBE-SERVICES -d 10.111.69.30/32 -p tcp -m comment --comment "default/nginx-35: has no endpoints" -m tcp --dport 80 -j REJECT --reject-with icmp-port-unreachable

-A KUBE-SERVICES -d 10.96.35.212/32 -p tcp -m comment --comment "default/nginx-38: has no endpoints" -m tcp --dport 80 -j REJECT --reject-with icmp-port-unreachable

-A KUBE-SERVICES -d 10.96.61.26/32 -p tcp -m comment --comment "default/nginx-84: has no endpoints" -m tcp --dport 80 -j REJECT --reject-with icmp-port-unreachable

-A KUBE-SERVICES -d 10.96.229.244/32 -p tcp -m comment --comment "default/nginx-87: has no endpoints" -m tcp --dport 80 -j REJECT --reject-with icmp-port-unreachable

-A KUBE-SERVICES -d 10.104.247.138/32 -p tcp -m comment --comment "default/nginx-66: has no endpoints" -m tcp --dport 80 -j REJECT --reject-with icmp-port-unreachable

-A KUBE-SERVICES -d 10.96.214.153/32 -p tcp -m comment --comment "default/nginx-11: has no endpoints" -m tcp --dport 80 -j REJECT --reject-with icmp-port-unreachable


-A KUBE-SERVICES -d 10.102.208.205/32 -p tcp -m comment --comment "default/nginx-55: has no endpoints" -m tcp --dport 80 -j REJECT --reject-with icmp-port-unreachable


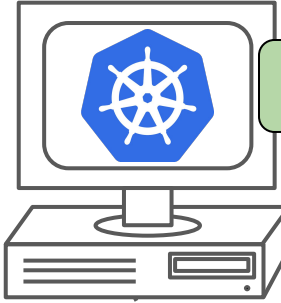

-A KUBE-SERVICES -d 10.106.35.32/32 -p tcp -m comment --comment "default/nginx-58: has no endpoints" -m tcp --dport 80 -j REJECT --reject-with icmp-port-unreachable

-A KUBE-SERVICES -d 10.107.174.56/32 -p tcp -m comment --comment "default/nginx-65: has no endpoints" -m tcp --dport 80 -j REJECT --reject-with icmp-port-unreachable

-A KUBE-SERVICES -d 10.98.142.83/32 -p tcp -m comment --comment "default/nginx-2: has no endpoints" -m tcp --dport 80 -j REJECT --reject-with icmp-port-unreachable

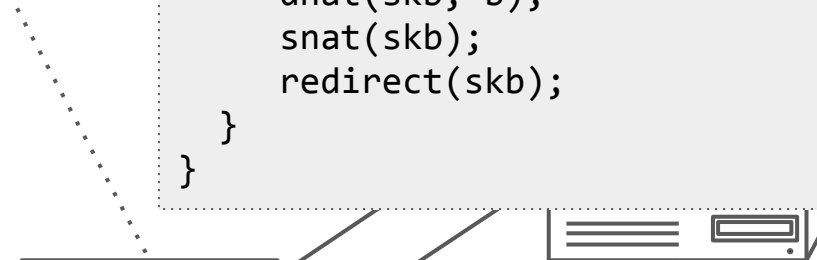

-A KUBE-SERVICES -d 10.106.248.222/32 -p tcp -m comment --comment "default/nginx-15: has no endpoints" -m tcp --dport 80 -j REJECT --reject-with icmp-port-unreachable

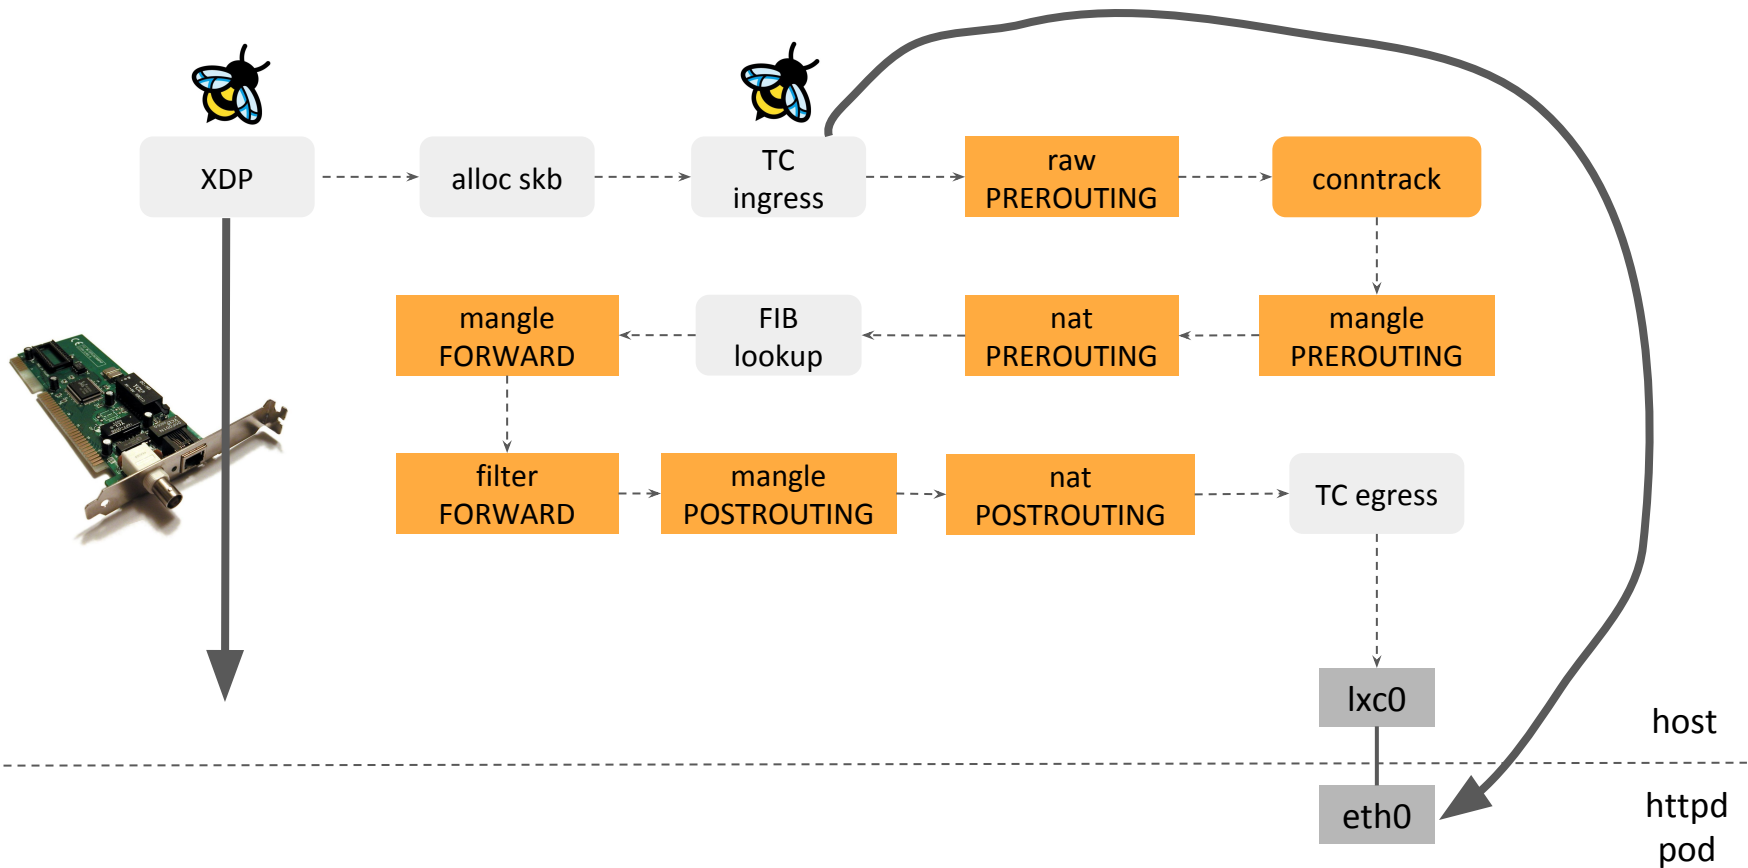


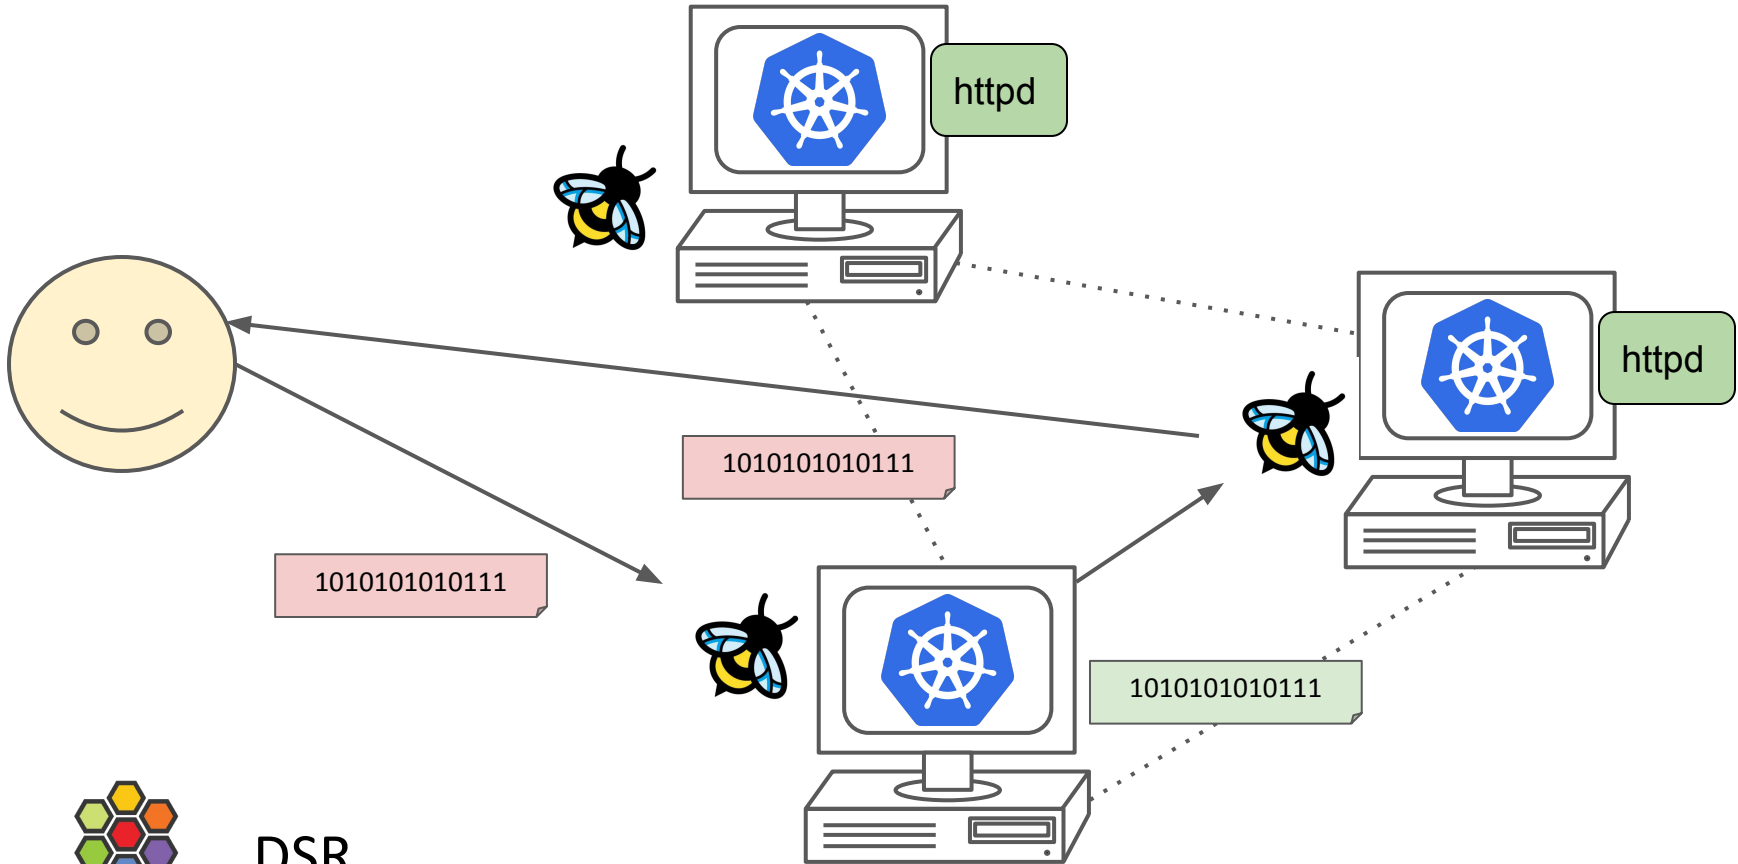


```
void loadbalance(skb) {  
    svc = bpf_map_lookup_elem(..);  
    if (svc) {  
        b = select_backend(svc);  
        dnat(skb, b);  
        snat(skb);  
        redirect(skb);  
    }  
}
```


pd








DSR





Consistent hashing  
(Maglev)





## Forwarding performance of tested Kubernetes node (higher is better)





# Questions?

<https://cilium.link/kubeproxy-free>



@martyns

