



**DYNAMIC NETWORK  
SERVICES INCORPORATED**

# Introducing OpenBFDD

Tom Daly, CTO

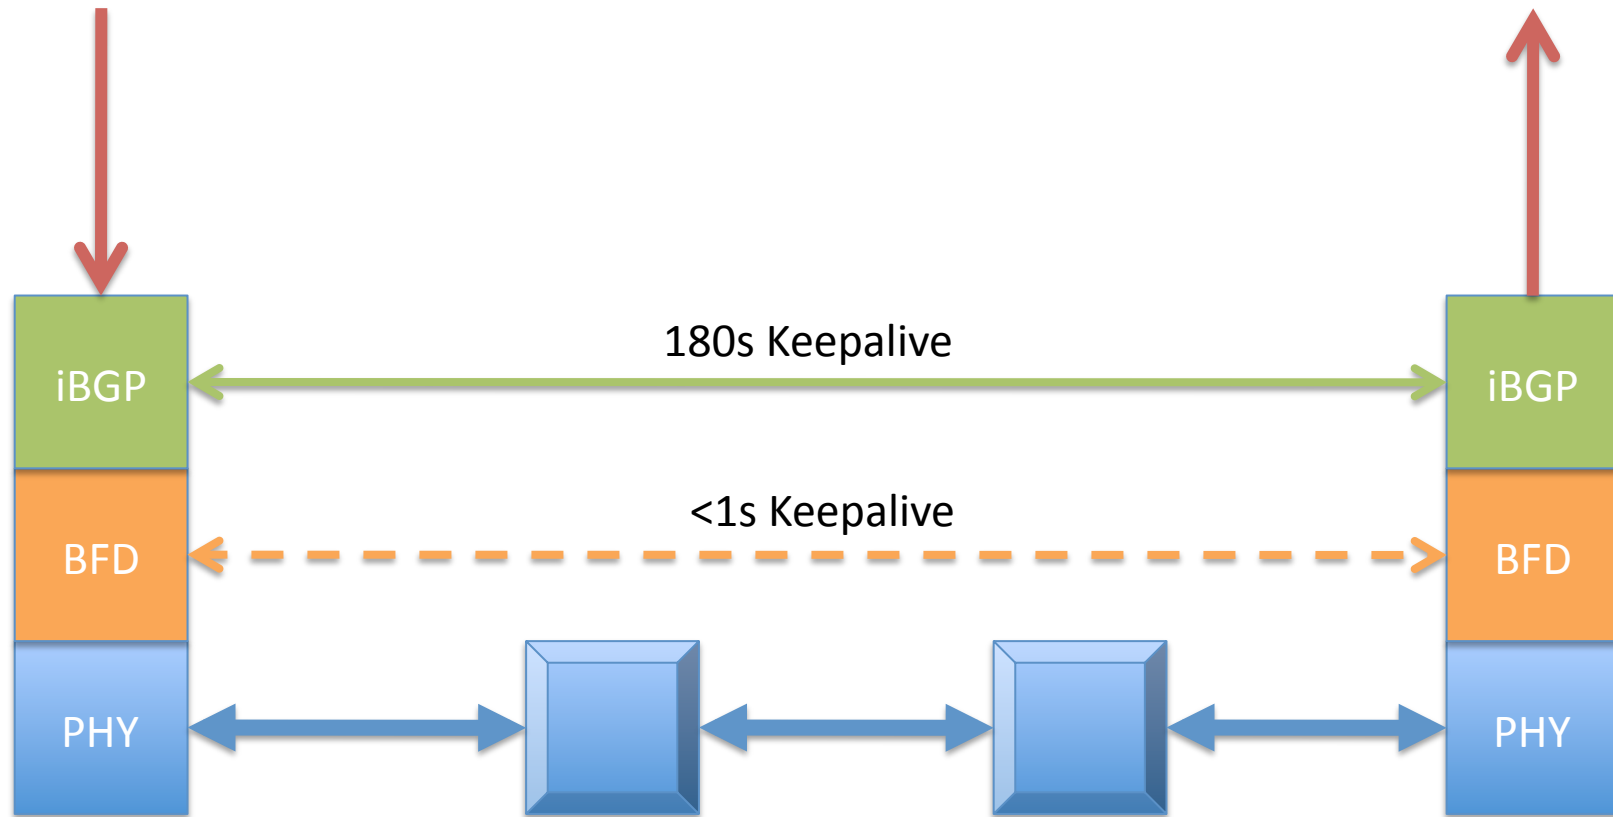
<http://dyn.com>

“Uptime is the Bottom Line”

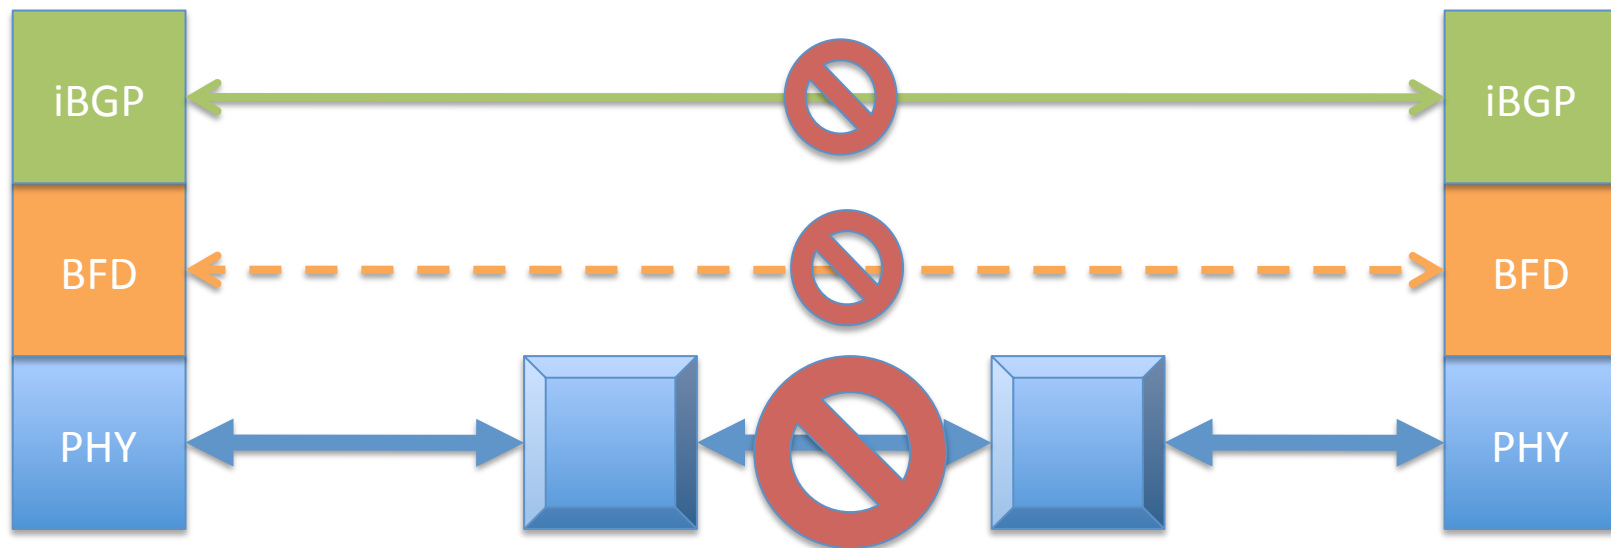
# Bidirectional Forwarding Detection

- BFD - RFC 5880
- Used to trigger more timely reconvergence for routing protocols running across shared media paths
- Can be pinned under OSPF, BGP, and other routing protocols and used a trigger

# Routing / Forwarding with BFP



# Routing / Forwarding with BFP

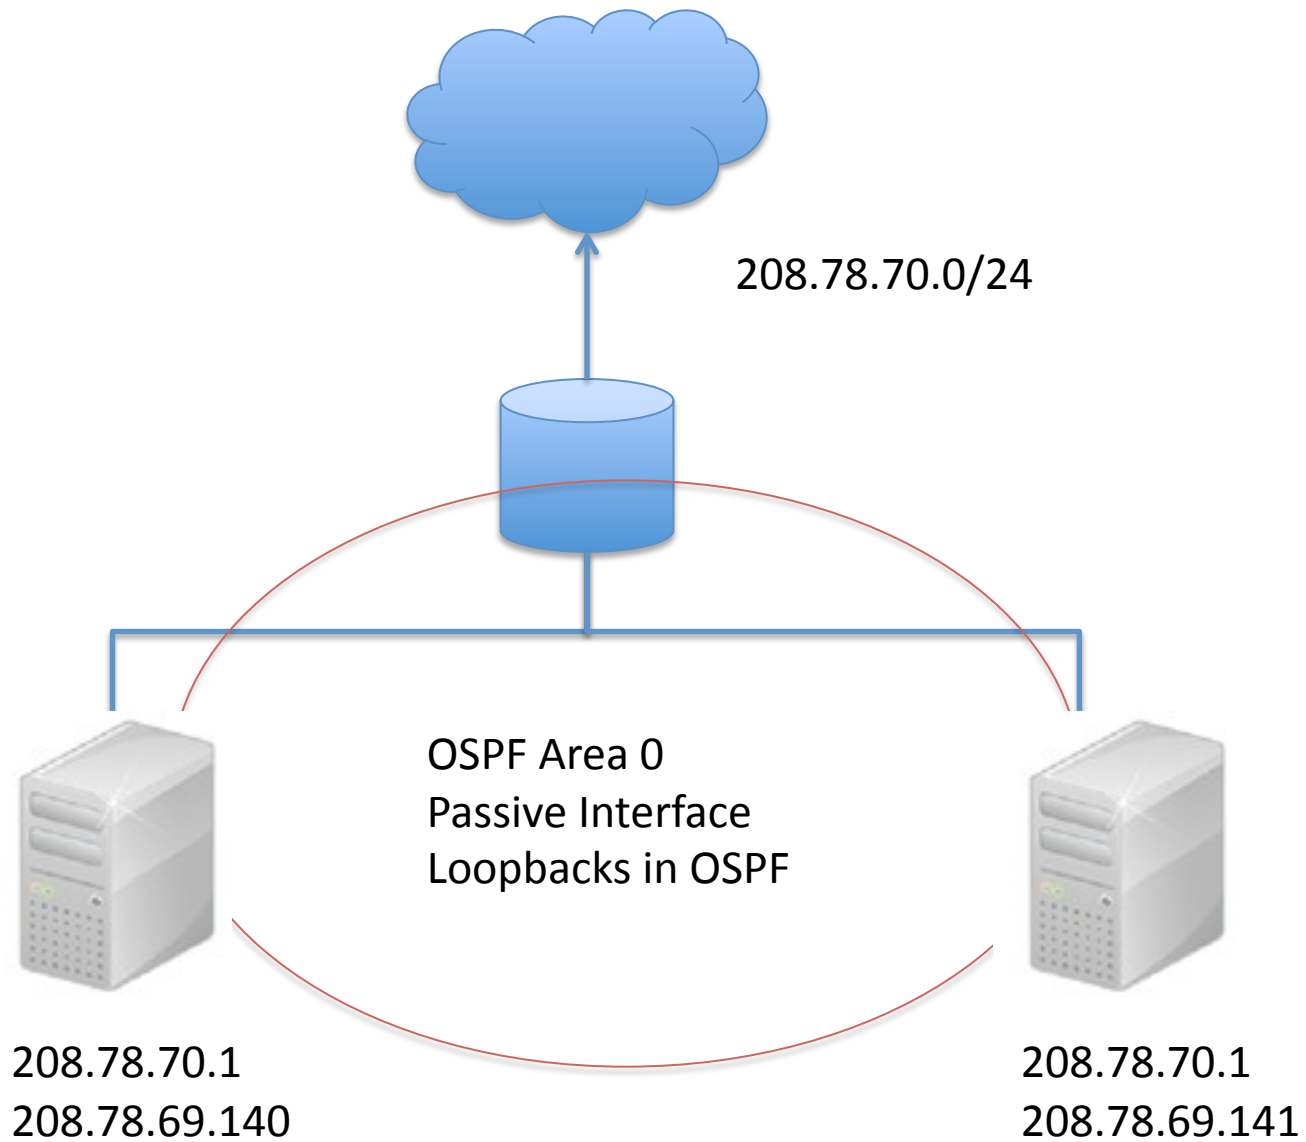


# A Novel Use Case

- Anycast can be used within the datacenter for low latency, low cost load balancing
- To accomplish this, we use an IGP + ECMP
- IGP running on application servers signals next hop availability to router
- Router sees multiple paths of same cost and load balances to them use L4 hashing

# OSPF + ECMP

- Router speaking OSPF to internal network
- Quagga + OSPFD running on every box
- Service VIP is on loopback, announced into OSPF
- Router sees equal cost paths to next-hops, and load shares
- Routers can be configured to hash on 2-tuple (src / dst IP) or 4-tuple (IPs + Ports)
- If something goes wrong on a box, shut down quagga





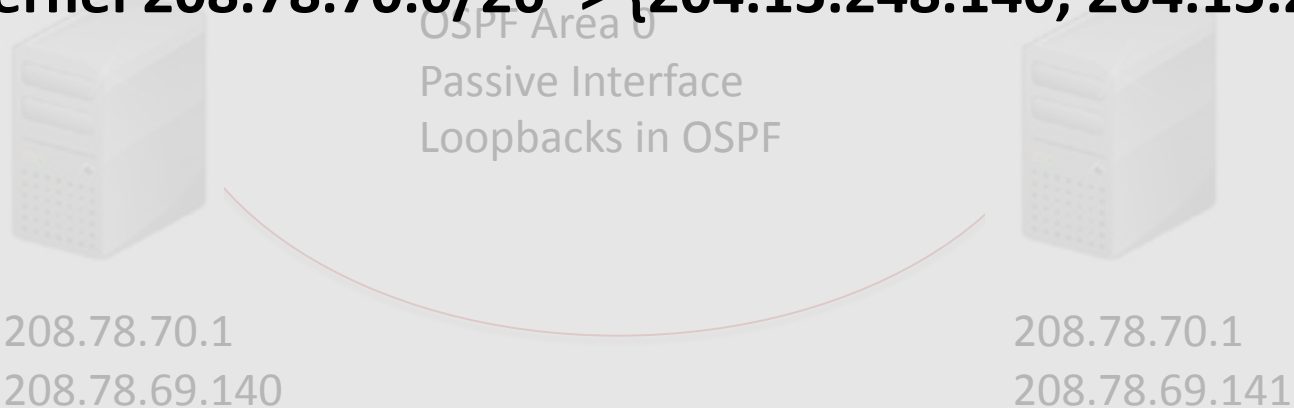
208.78.70.0/24

```
tom@core-01-ewr> show route 208.78.70.1 extensive  
{...}
```

```
208.78.70.0/26 (1 entry, 1 announced)
```

```
TSI:
```

```
KRT in-kernel 208.78.70.0/26 -> {204.13.248.140, 204.13.248.141}
```






# Scaling OSPF

- OSPF got too heavy 16-way (or 64-way) ECMP
- Service daemon crashes or machine failures result in OSPF topology changes
- Topology changes cause undue strain on router CPU
- Needed a new way to enumerate available next hops for ECMP

# Enter OpenBFDD

- Open Source implementation of the BFD protocol for FreeBSD
- Used to enumerate next-hop availability between DC edge routers and internal service machines.





208.78.70.0/24

tom@core-01-pao.dyndns.com> show bfd session

Detect Transmit

Address	State	Interface	Time	Interval	Multiplier
208.78.69.140	Up	ge-1/3/0.20	1.500	0.500	3
208.78.69.141	Up	ge-1/3/0.20	1.500	0.500	3

2 sessions, 2 clients

Cumulative transmit rate 6.0 pps, cumulative receive rate 6.0 pps

208.78.70.1  
208.78.69.140

208.78.70.1  
208.78.69.141



208.78.70.0/24

tom@core-01-pao.dyndns.com> show route 208.78.70.1

{...}

208.78.70.0/26 \*[Static/5] 5d 22:09:55  
to 208.78.69.140 via ge-1/3/0.20  
> to 208.78.69.141 via ge-1/3/0.20



208.78.70.1  
208.78.69.140



208.78.70.1  
208.78.69.141



# Current Status

- bfd-beacon: Able to negotiate and turn up an asynchronous BFD session.
- Swapping out event libraries to support other operating systems (currently kevent)
- Future support for Demand Mode and Echo

<http://github.com/dyninc/OpenBFDD>