

# Dynamic P2P with BGP Route Servers

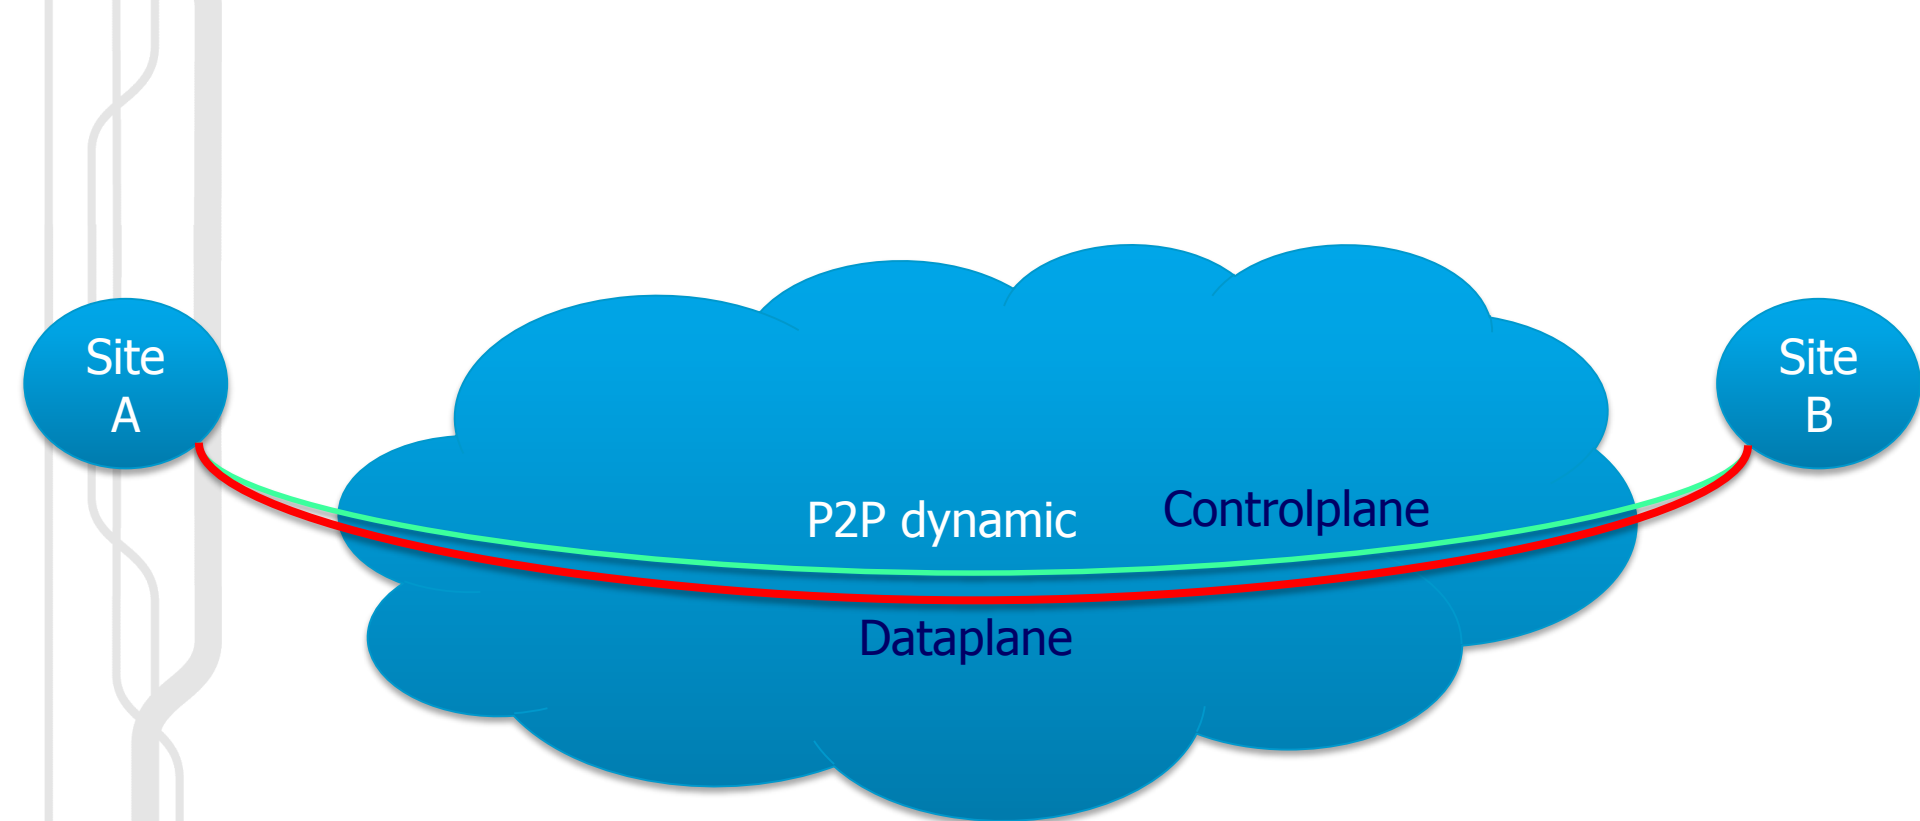
## BFD for data-plane verification

Magnus Bergroth

NORDUnet

- Dynamic P2P links has two end points that normally terminates in a aggregation router at each site.
- On logical interface per destination site.
- eBGP are configured over the logical interface to each site.
- Reachability is advertised after the P2P link is up and BGP is established.

- Full mesh of BGP sessions.
- Extensive amount of configuration.
- BGP sessions over short lived P2P links are most of the time down and causes alarms.



- Controlplane shared with dataplane
- Dataplane reachability detected when controlplane goes down

eBGP between loopbacks

130.242.0/24  
Next-hop  
10.0.0.0

193.10.0/24  
Next-hop  
10.0.0.1

Site  
A

Site  
B

10.0.0.0/31

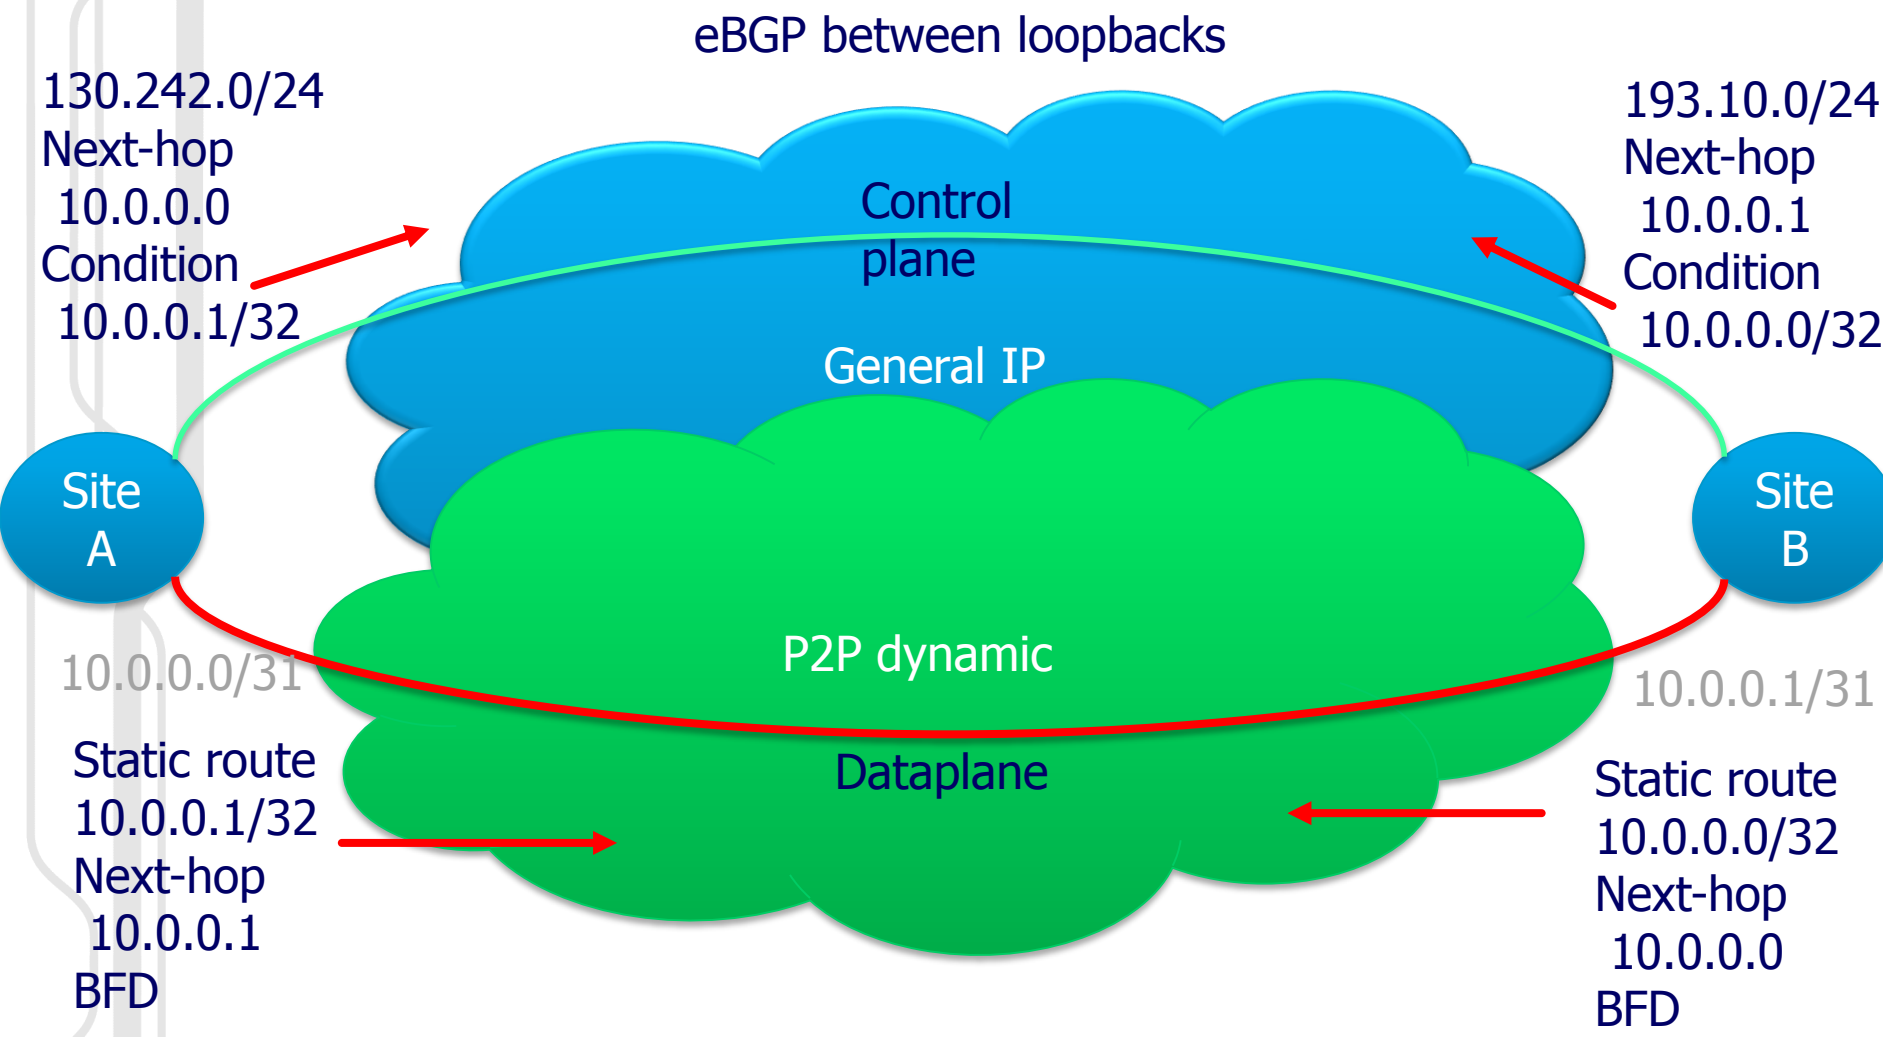
10.0.0.1/31

Control  
plane


General IP

P2P dynamic

Dataplane



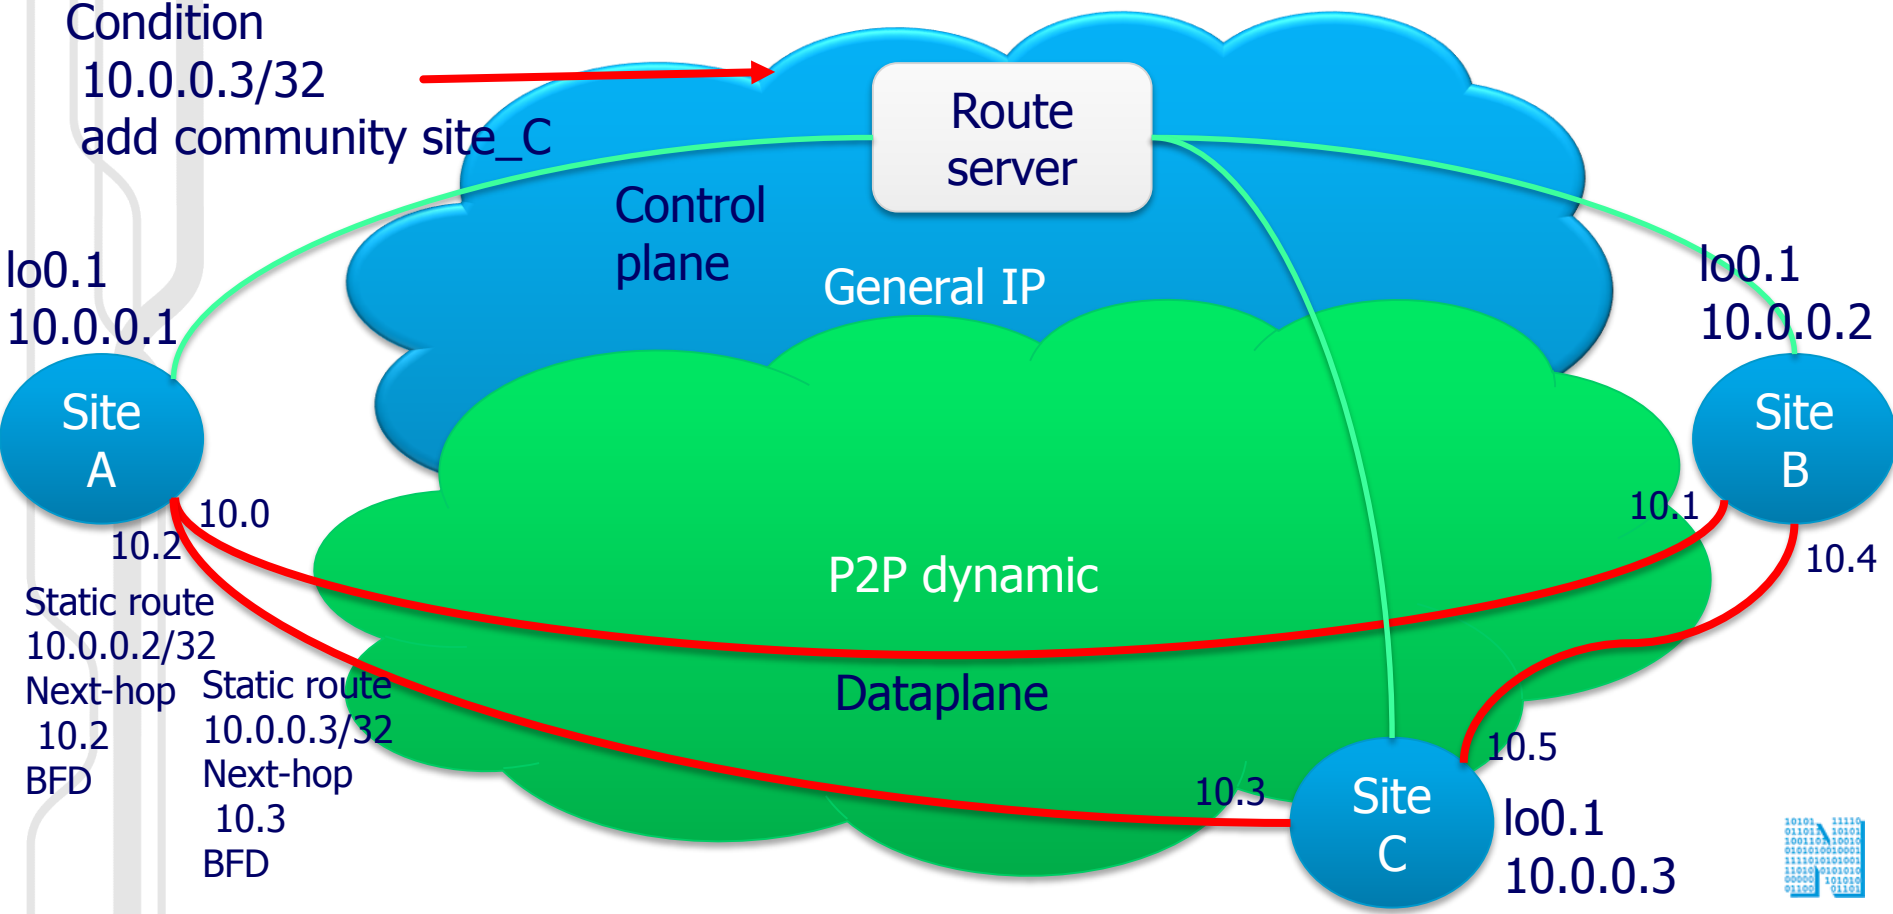
- Simplify the BGP setup
- Only one BGP session per site
- Route server with one outgoing RIB per site, steering using communities



- Logical interface per site
- BFD per site
- Community per site

```

130.242.0/24
Next-hop
10.0.0.1
Condition
10.0.0.2/32
add community site_B
Condition
10.0.0.3/32
add community site_C
    
```



lo0.1  
10.0.0.1  
Site A

lo0.1  
10.0.0.2  
Site B

lo0.1  
10.0.0.3  
Site C

```

Static route
10.0.0.2/32
Next-hop
10.2
BFD
Next-hop
10.3
    
```

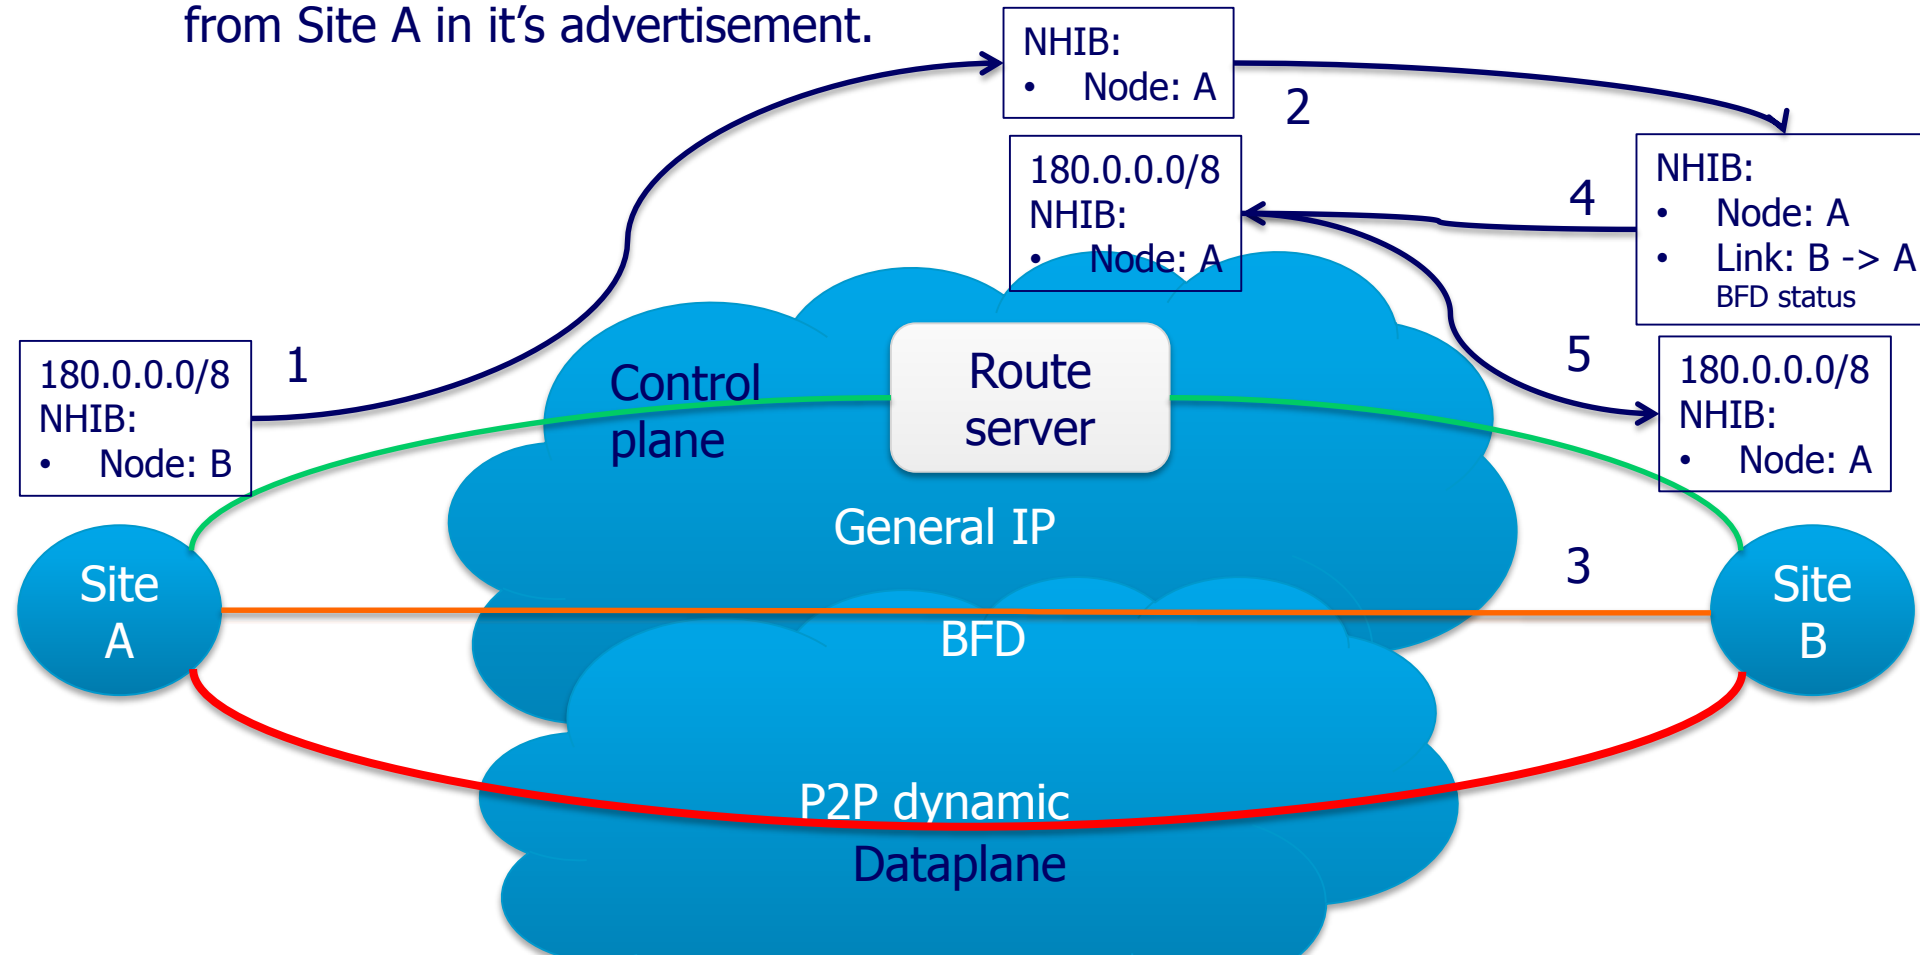




Edoardo found

- Created for route-servers at IXP
- Client routers verifies connectivity of dataplane with BFD Bidirectional Forwarding Detection, RFC5880.
- Client routers communicates reachability northbound to the Routeserver, via Link-State and TE Information using BGP (BGP-LS):
- Routeserver stores Per peer: Next-Hop Information Base (NHIB) reachability for all next-hops

1. Site A send that it will communicate with Site B.
2. Route server updates NHIB with nexthop from Site A and send to Site B via BGP-LS.
3. Site B automatically setup a BFD session towards Site A
4. Site B updates it's NHIB and advertise to Route-server via BGP-LS
5. If Site B advertise reachability to Site A will the Route server include it's routes from Site A in the BGP advertise to Site B.
6. If Site B advertise `_no_` reachability will the route server remove the prefixes from Site A in it's advertisement.



- Only one BGP session with the Router-server. Minimize the router configuration.
- The BGP session will always be up.
- Connectivity to new sites are added by adding communities to advertised prefixes.
- Prefixes are learned via BGP only when the Data-plane(dynamic P2P) is up.
- Fast detection when the dynamic P2P are teared down. No risk of black holing

- Still a draft
- Quagga and Bird are working on implementation.
- The draft suggest a dataplane to be consider permanent down after 24hours. This needs to be configurable to be tested for ever.
- The Routeserver init state is that the dataplane is working. For dynamic P2P is this not the case.
- Requeste has been made for the two above to be changed.

- Questions?