

# Taking the Edge off with Espresso: Scale, Reliability and Programmability for Global Internet Peering

Kok-Kiong Yap, Murtaza Motiwala, Jeremy Rahe, Steve Padgett, Matthew Holliman, Gary Baldus, Marcus Hines, Tae-eun Kim, Ashok Narayanan, Ankur Jain, Victor Lin, Colin Rice, Brian Rogan, Arjun Singh, Bert Tanaka, Manish Verma, Puneet Sood, Mukarram Tariq, Matt Tierney, Dzevad Trumic, Vytautas Valancius, Calvin Ying, Mahesh Kallahalla, Bikash Koley, Amin Vahdat

Google  
ACM SIGCOM 2017

*Presented by:*  
*Rafael Lourenço*  
*July 13, 2018*  
*NetLab Weekly Meetings*

**UC DAVIS**  
UNIVERSITY OF CALIFORNIA



# Taking the Edge off with Espresso

Scale, Reliability and Programmability for Global Internet Peering

**KK Yap**, Murtaza Motiwala, Jeremy Rahe, Steve Padgett, Matthew Holliman, Gary Baldus, Marcus Hines, Taeun Kim, Ashok Narayanan, Ankur Jain, Victor Lin, Colin Rice, Brian Rogan, Arjun Singh, Bert Tanaka, Manish Verma, Puneet Sood, Mukarram Tariq, Matt Tierney, Dzevad Trumic, Vytautas Valancius, Calvin Ying, Mahesh Kallahalla, Bikash Koley, Amin Vahdat and many others.



# Problem Statement

Egress Terabits/sec of traffic to our Internet peers

- High-def video, cloud traffic, etc.

# Problem Statement

Egress Terabits/sec of traffic to our Internet peers

- High-def video, cloud traffic, etc.

## **1. Optimize traffic per-customer and per-application**

- e.g., optimal video quality, or differentiated service for cloud

# Problem Statement

Egress Terabits/sec of traffic to our Internet peers

- High-def video, cloud traffic, etc.

## **2. Deliver new features quickly**

# Espresso: Google's SDN Peering Edge


Our previous experience with SDN

- B4 [SIGCOMM 2013] and Jupiter [SIGCOMM 2015]
- Enable flexible traffic engineering
- Increase feature velocity

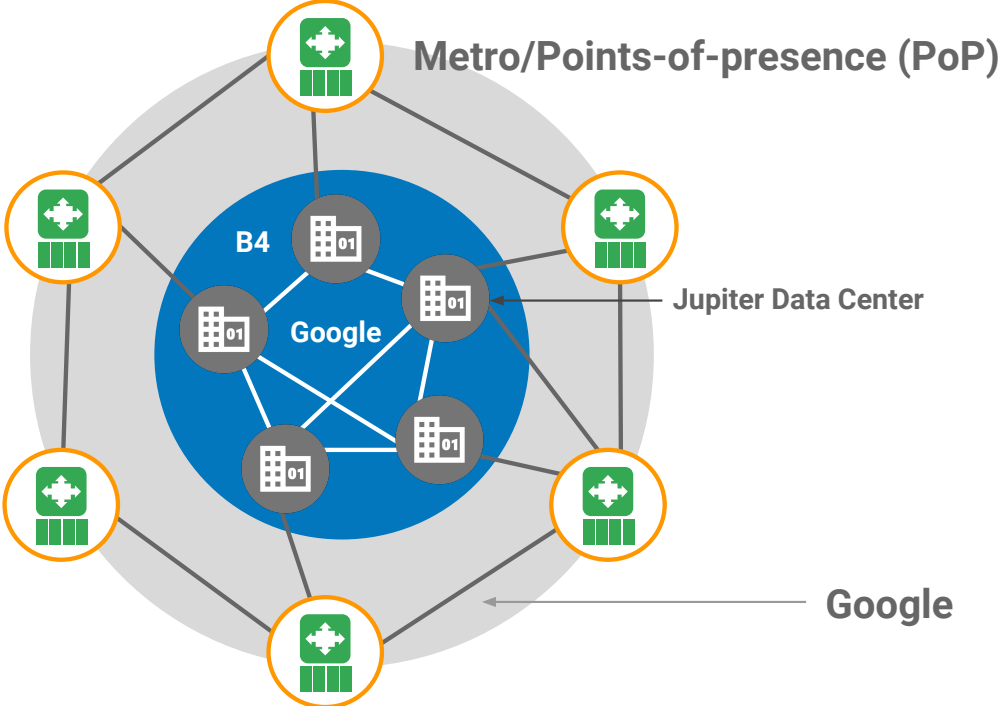
# Agenda

- Problem Statement
- Espresso in Context
- Design Principles
- Architecture Overview
- Results
- Conclusion

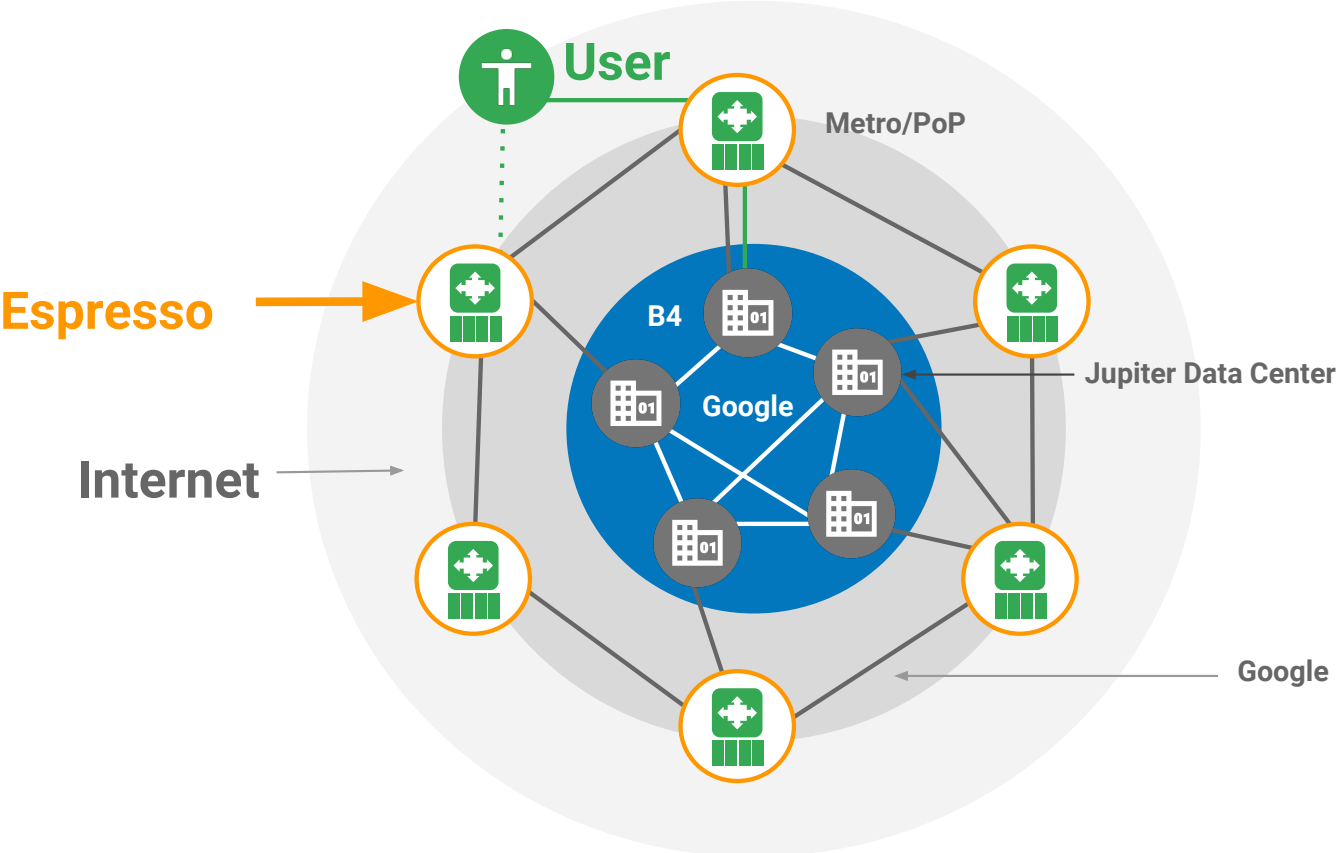
# Espresso in Context



# Espresso in Context

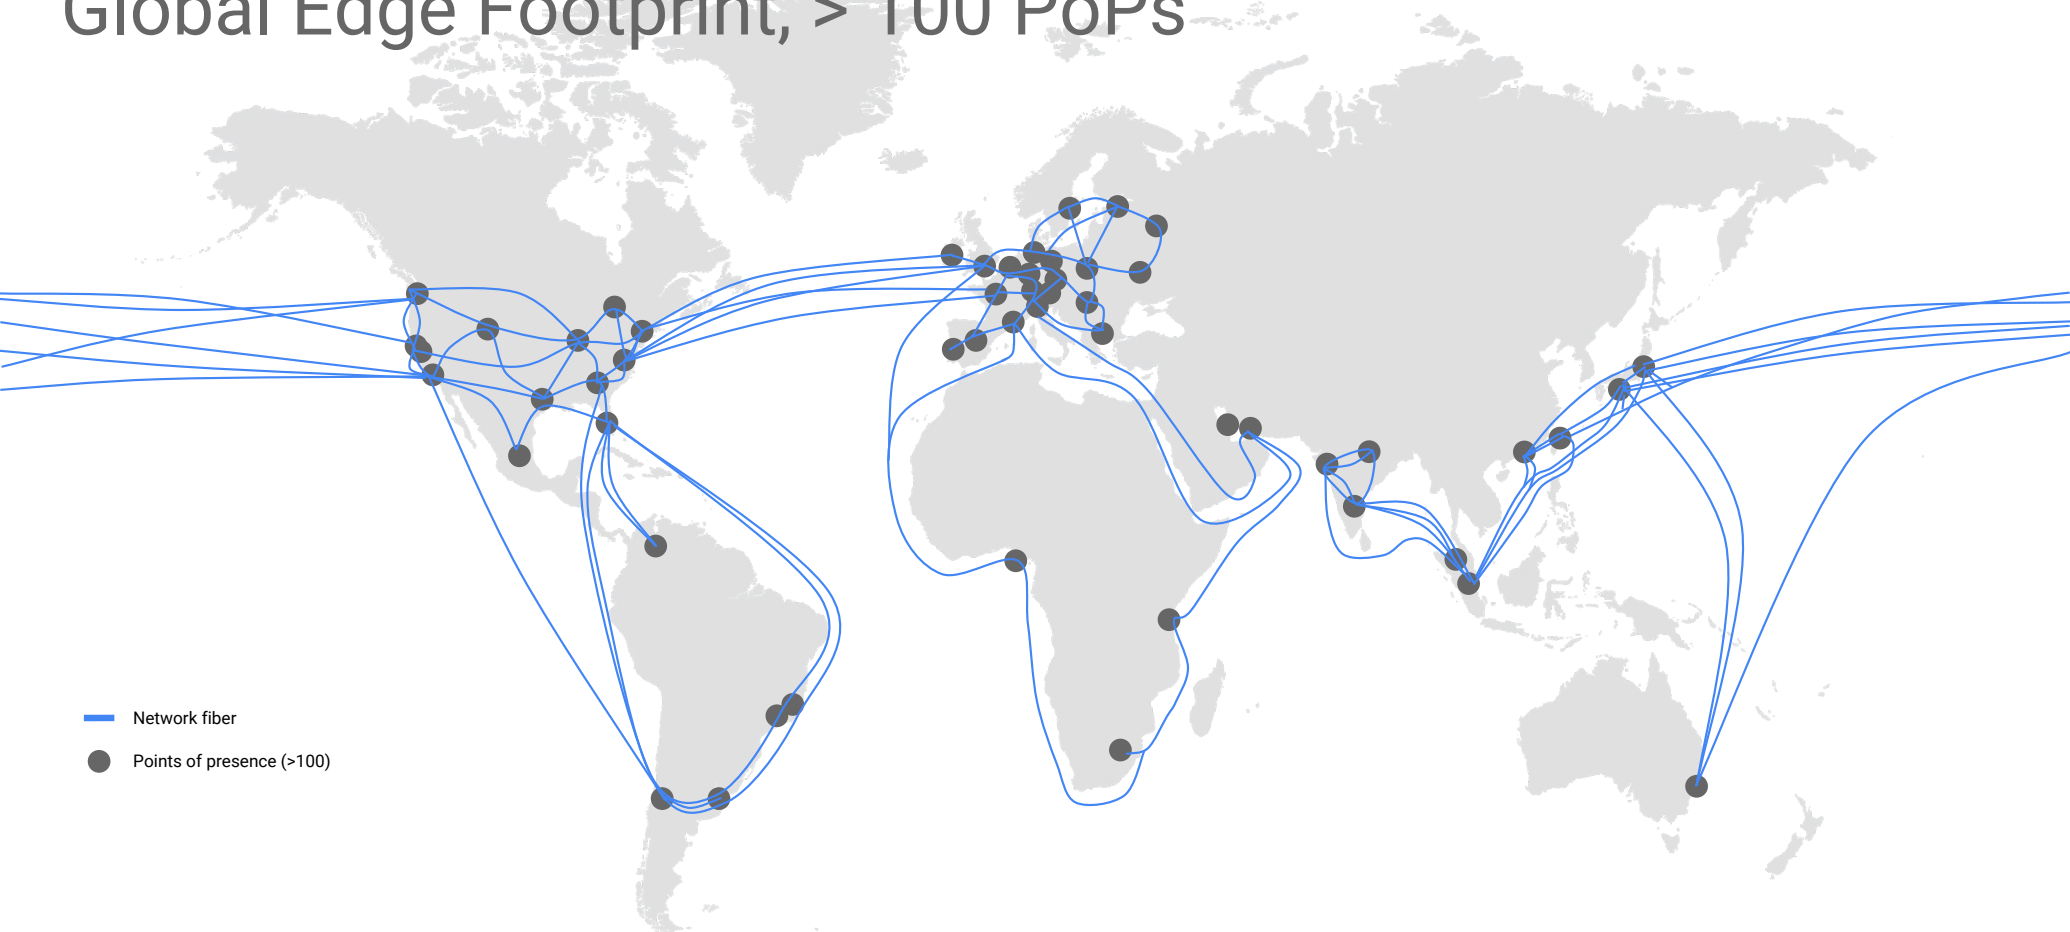


# Espresso in Context





# Global Edge Footprint, > 100 PoPs




# Espresso's Design Principles

1. Hierarchical control plane
  - Global optimization while local control plane provide fast reaction.
2. Fail static
  - Local control plane continues to function without global controller failure.
3. Software programmability
  - Externalize features into software to exploit commodity servers for scale.
4. Testability
5. Manageability

*Hierarchical control plane*  
*Fail static*  
**Software programmability**

# Architecture: Externalizing BGP

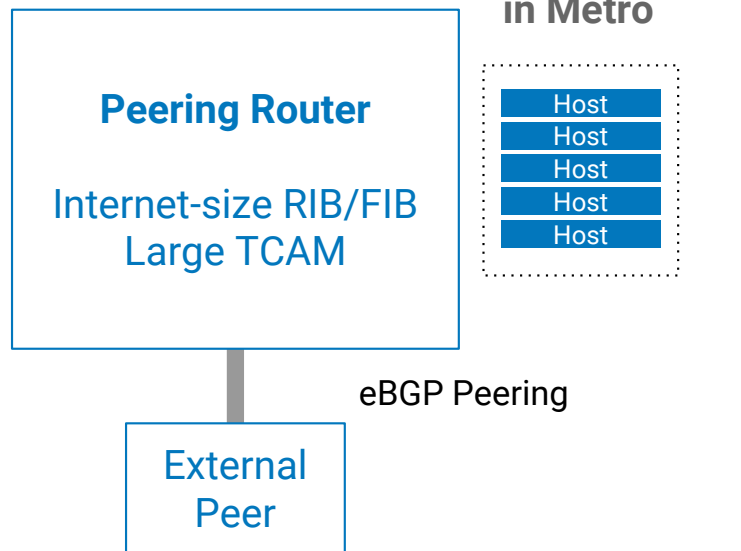
## Traditional Peering Router



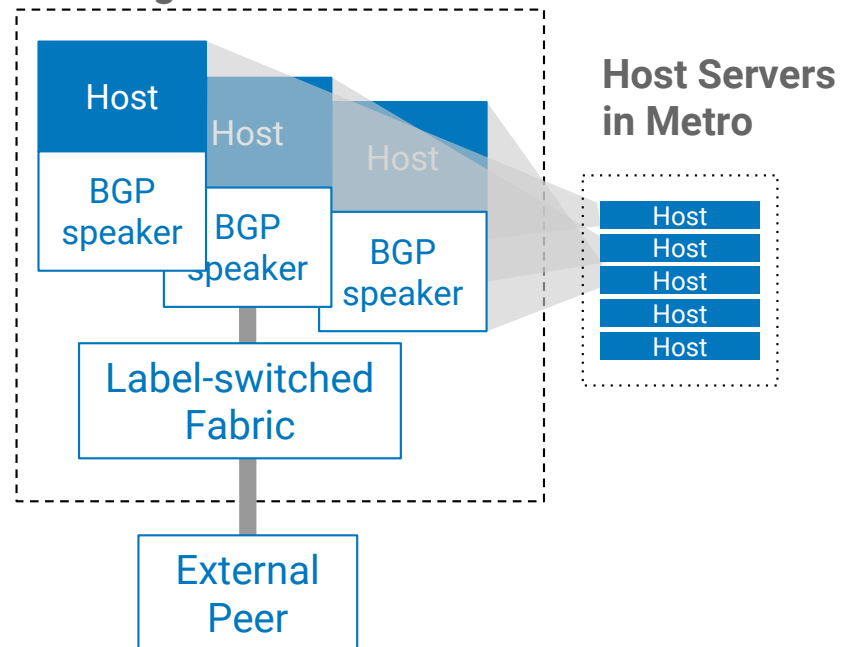
# Architecture: Reliability and Scale of BGP

*Hierarchical control plane*  
*Fail static*  
**Software programmability**

## Traditional Peering Router



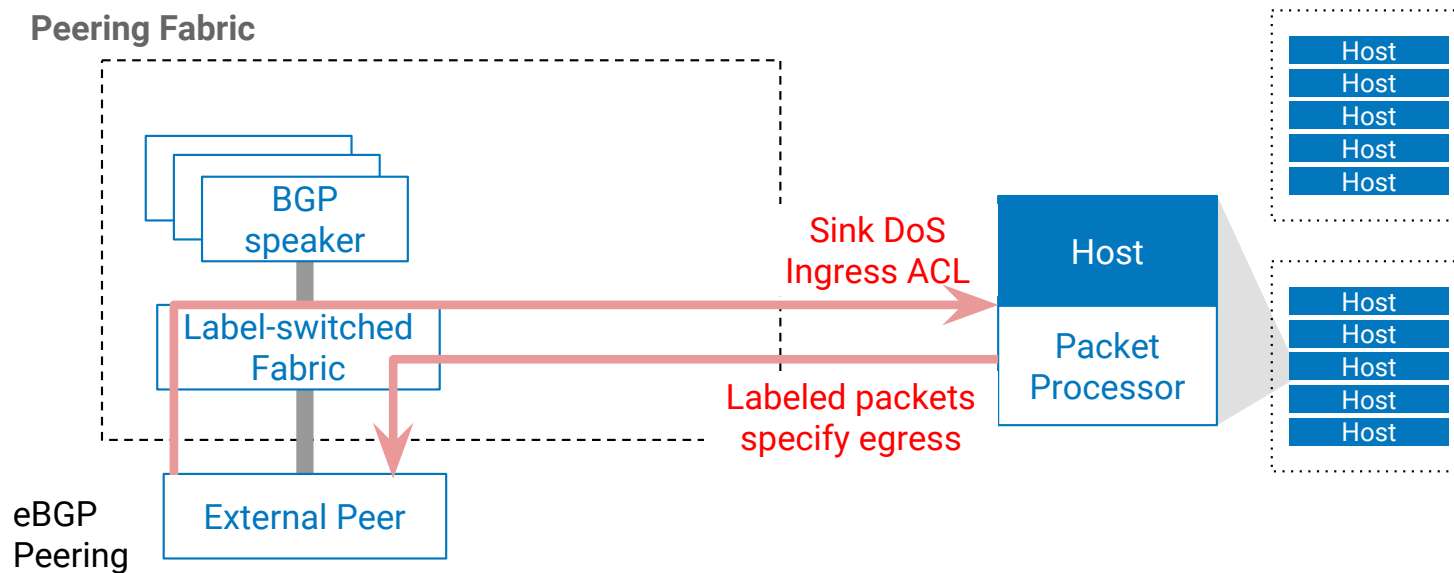
## Espresso Peering Fabric



# Architecture: Externalize Packet Processing

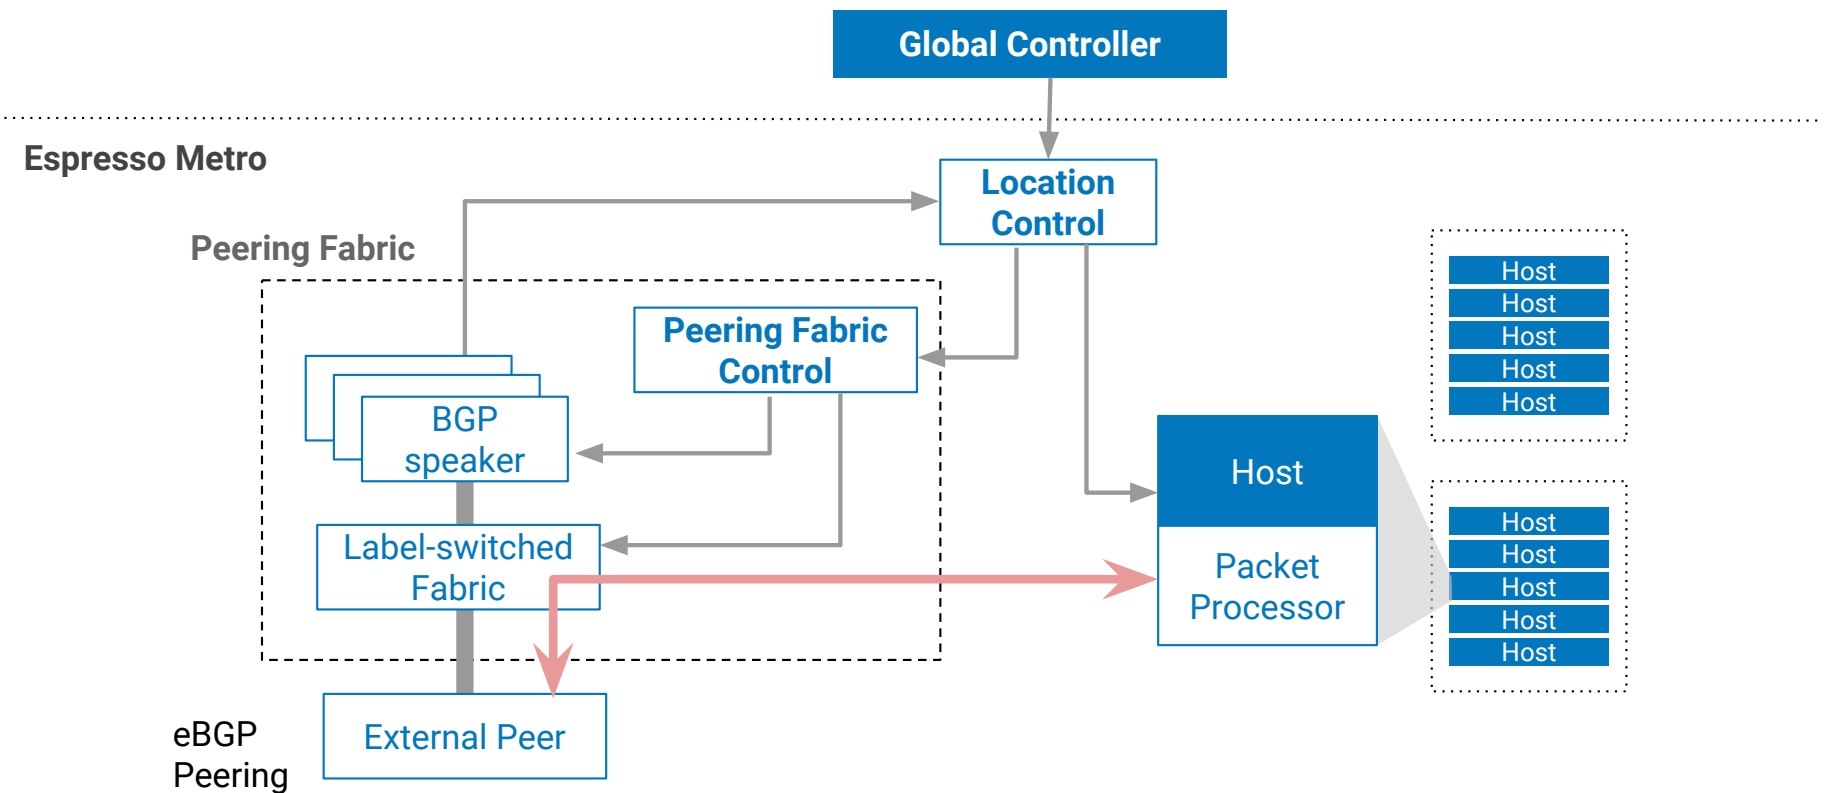
*Hierarchical control plane*  
*Fail static*  
**Software programmability**

Host-based packet processor allows flexible packet processing, including ACL and handling of DoS.



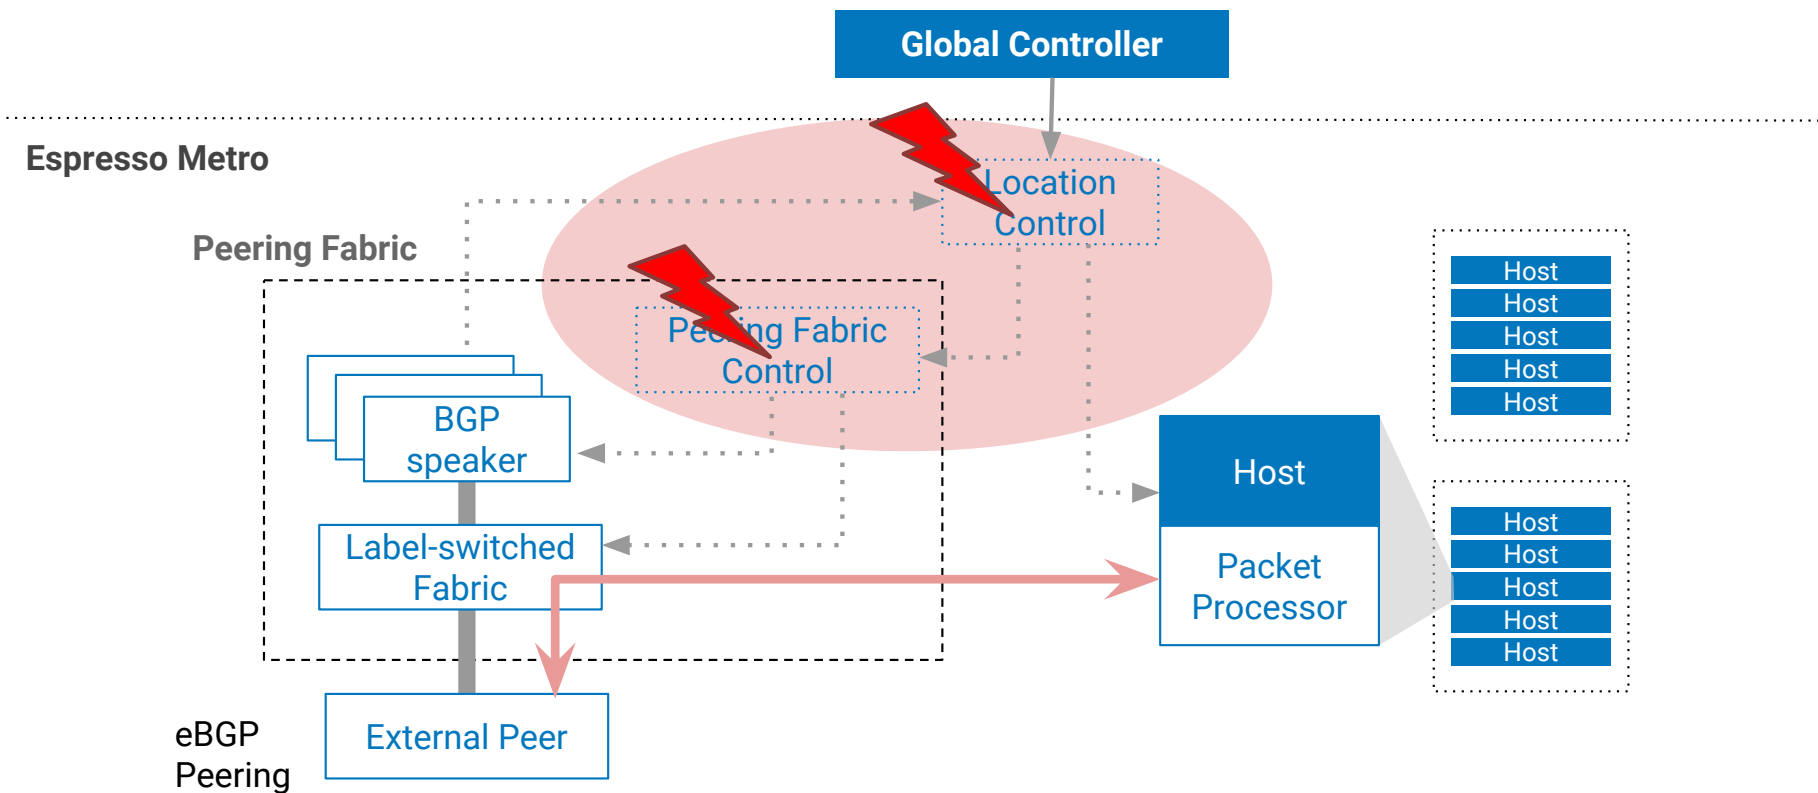
# Architecture: Hierarchical Control

**Hierarchical control plane**  
Fail static  
Software programmability



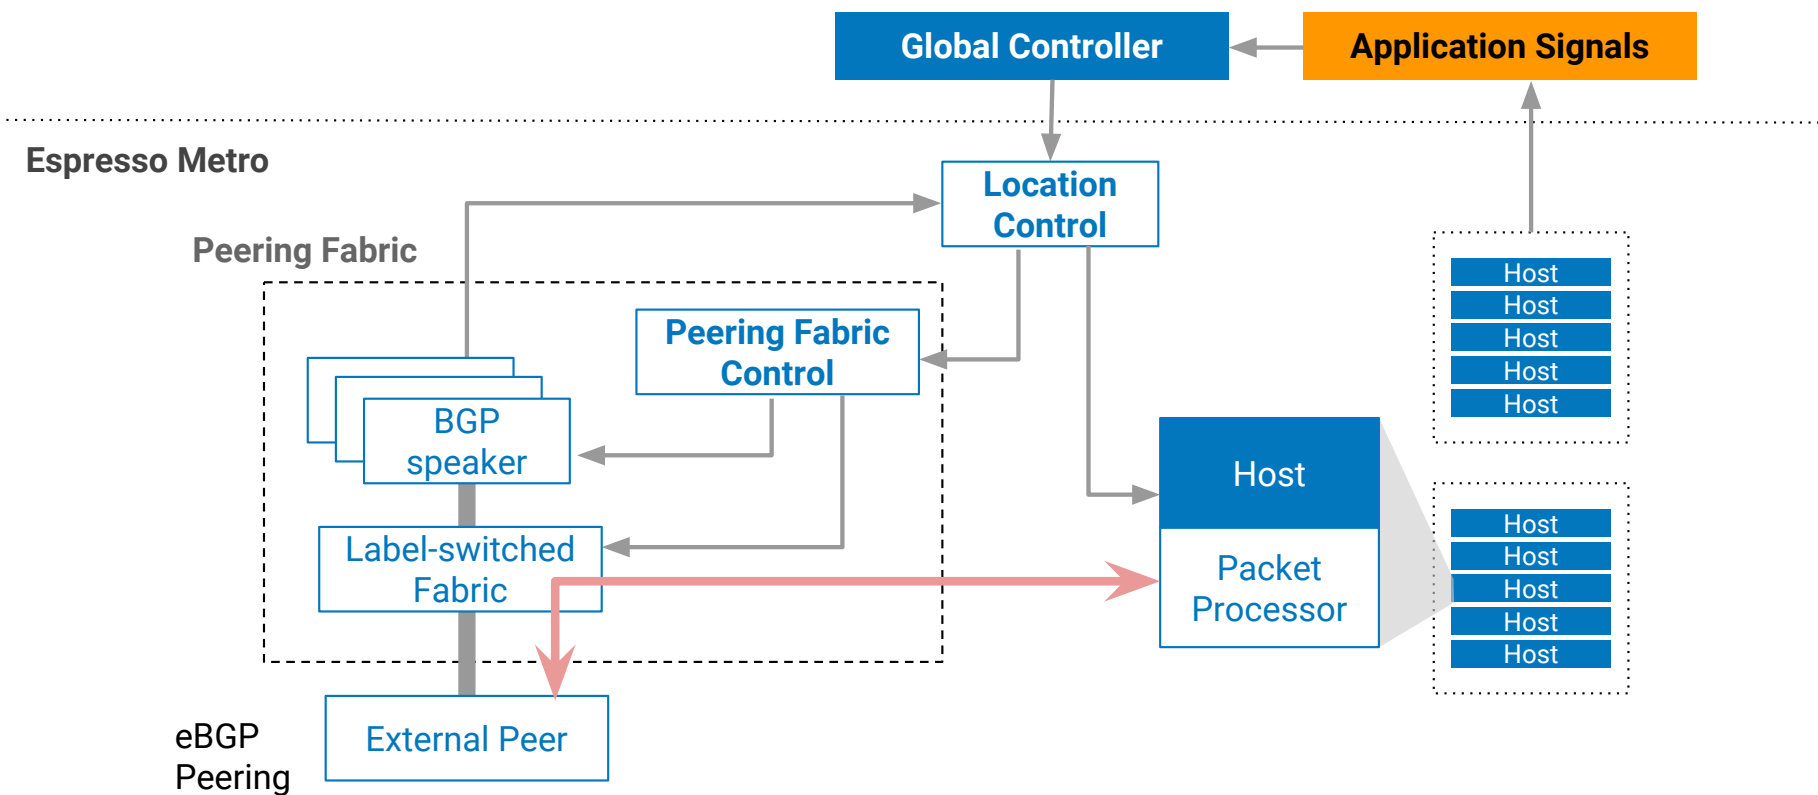
# Architecture: Fail Static

**Hierarchical control plane**  
**Fail static**  
Software programmability



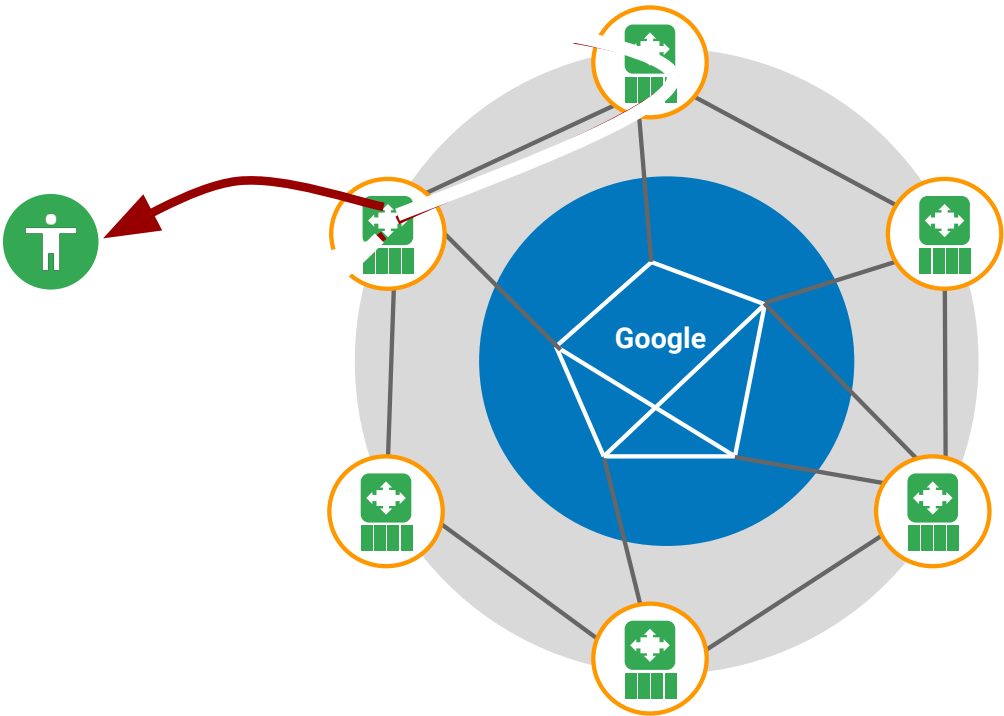
# Architecture: Application Aware Routing

*Hierarchical control plane*  
*Fail static*  
**Software programmability**





# Using User's Best Path, not BGP's



# Improvements in End User Experience

<b>Client ISP</b>	<b>Change in mean time between rebuffers (MTBR)</b>	<b>Change in Mean Goodput</b>
A	10 → 20 min	2.25 → 4.5 Mbps
B	4.6 → 12.5 min	2.75 → 4.9 Mbps
C	14 → 19 min	3.2 → 4.2 Mbps


# Release Velocity

<b>Component</b>	<b>Average Velocity (days)</b>
Local Controller	11.2
BGP speaker	12.6
Peering Fabric Controller	15.8

# Conclusion

~~SDN is only suited for walled gardens.~~

# Conclusion



Local view  
Connectivity based optimization  
Slow evolution  
Costly

# What I learned during my time in Davis

- Focus on the problem, not only on the solution
- Try to not make it just one more incremental work
- The following are (at least in part) out of your control:
  1. Whether the problem being investigated will be very important
  2. Whether the methods you use to solve will be the best
  3. Whether you will have the insight to elegantly solve the problem
- **But:** describing and discussing your problem well, writing well, doing good job with examples, results, and presentation **are all within your control**
- Who is the client? Who will benefit?
- Frequently question all your assumptions (in life, also)
- Respect other cultures, specially the one of the country where you are currently living in
- Make friends
- Have fun and be happy