

# CAP for Networks

Aurojit Panda, Colin Scott, Ali Ghodsi, Teemu Koponen, Scott Shenker  
UC Berkeley, KTH, VMware, ICSI

# CAP Theorem

When distributed systems face network **P**artitions  
pick one of

- Service **C**orrectness (Consistency)
- Service **A**vailability

# CAP Theorem: Impact

Divides the database community (even today)

SQL

Correctness above all



NoSQL

Availability above all



How does the CAP  
theorem apply to  
networks?

Our goal: articulate  
fundamental  
tradeoffs networking  
will face

# Traditional Networks

When intradomain routing was the main concern

- **Correctness:** Deliver packets to destination
- **Availability:** Deliver packets to destination
- **Correctness** is the same as **Availability**

# The move to SDN

SDN provides more sophisticated **functionality**:

- Tenant isolation (ACL enforcement)
- Traffic Engineering
- Virtualization

Control plane partitions no longer imply data plane partitions

- Control traffic often does not use data plane network

# Availability $\neq$ Correctness

During control plane partitions

- Data plane connected => Deliver packets (**Availability**)
- Inconsistent control plane data (~~**Correctness**~~)
- **Availability** does not imply **Correctness**




How does the CAP  
theorem apply to  
~~networks~~ SDN?

# How does the CAP theorem apply to ~~networks~~ SDN?


Can one provide isolation and availability in the presence of control plane partitions?

# Network Model



- Out-of-band control network.
- Routing and forwarding based on addresses.
- Policy specification using end-host names.
- Controller responsible for local name-address bindings.

# Isolation Result



- Consider policy isolating A from B.
- A control network partition occurs.
- Only possible choices
  - Let all packets through (including from A to B) (~~Correctness~~)
  - Drop all packets (including from A to D) (~~Availability~~)

# Value of Model

Practical workarounds follow directly

# Workarounds for Isolation

- Identity-Address disconnect underlies isolation result

# Workarounds for Isolation

- Identity-Address disconnect underlies isolation result
- Network can constrain address allocation

# Workarounds for Isolation

- Identity-Address disconnect underlies isolation result
- Network can route on identity rather than addresses




# Workarounds for Isolation

- Identity-Address disconnect underlies isolation result
- Use in-band control networks rather than out-of-band

# Workarounds not General


## Edge Disjoint Traffic Engineering

- Two flows must traverse disjoint links



# Workarounds not General

Previous workarounds not applicable!



Can one provide correctness  
and availability in the  
presence of partitions?

**Not in general**

# In the Paper

“CAP for Networks”, HotSDN ‘13

- More policies and proofs
- More details on workarounds
- Other ways to model the network

# CAP for Networks

Choices for network architects

Correctness above all

**VMware NSX**

**ICING**

Availability above all

**Traditional Routing**

**BGP**

---

Policy-Specific Workarounds

**Packet Labeling**

**In-Band Control**

# Backup Slides

# Host Migration

- Our model assumes host migrations without controller involvement.
- In part this is because host migrations are common
  - Soundararajan and Govil 2010: **6** migrations/day/VM
  - In a datacenter **~480,000** migrations/day
  - **5.5** migrations per second
- Controller involvement is too expensive in datacenters
  - NSX and BSC work in a similar manner
- In enterprises controller involvement complicated by mobility.