

---

# Rethinking Centralization

---

RUSS WHITE



---

Note Well: This is an experimental presentation


---





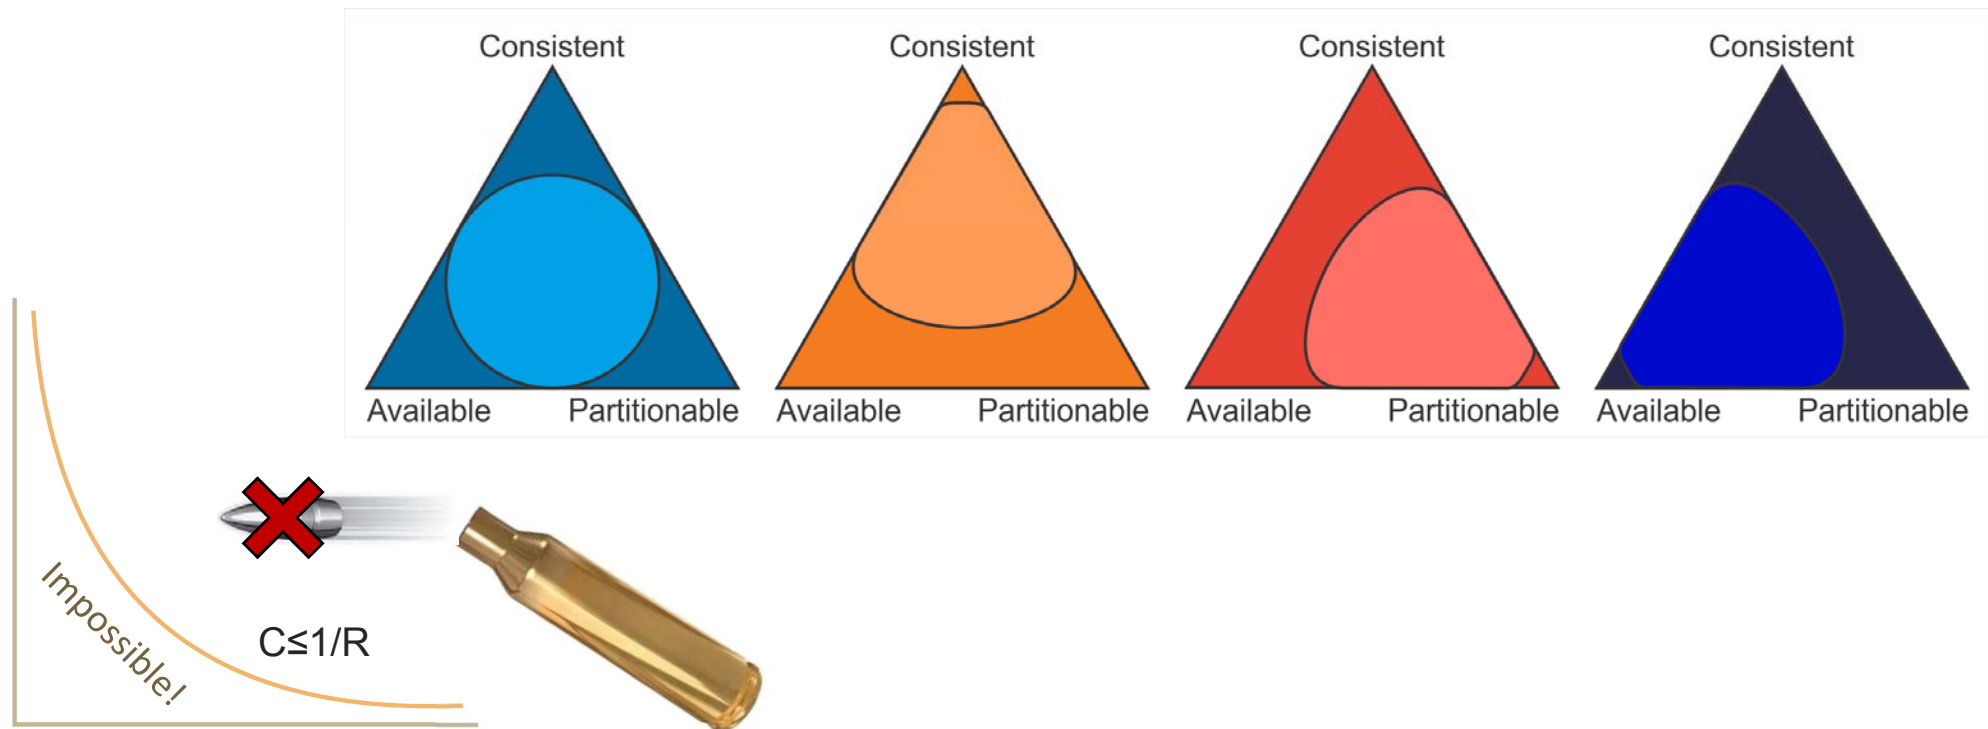
# Why do we do this?

- This centralized stuff is too slow/hard/complex – wouldn't it be cool if we could distribute route calculation?
- This distributed stuff is too hard/complex – wouldn't it be cool if we could centralize all of our policy?
- Can we stop this CAPEX waste?



via Greg Ferro | <http://etherealmind.com>

# There is *no* Silver Bullet for Complexity





There is no silver bullet for complexity — but... We're certainly riding this pendulum looking for one...



---

So here's a radical thought:

STOP IT!





Where do  
we go  
from  
here?



---

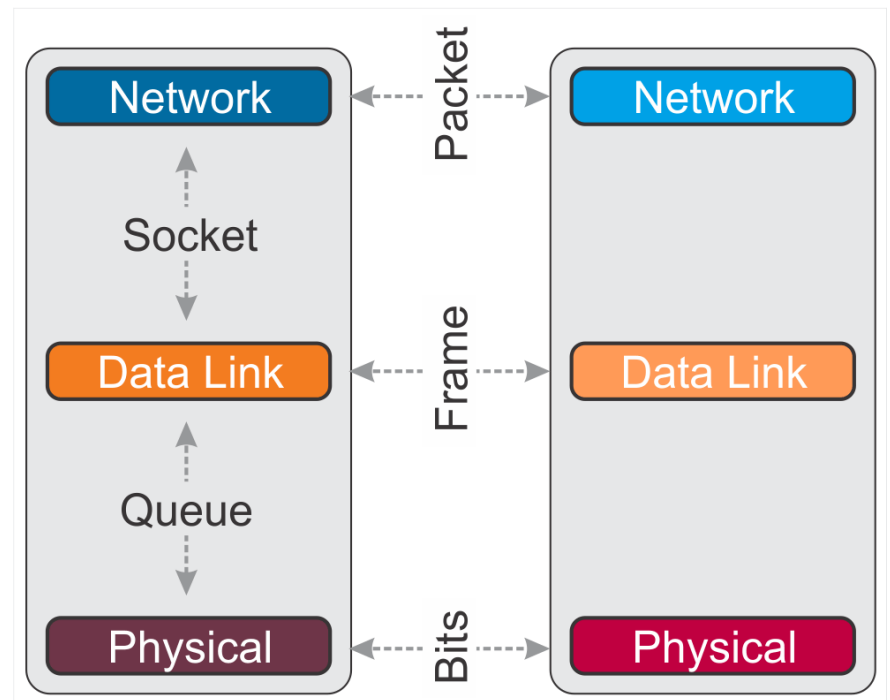
How do we solve this in the rest of the network?

---



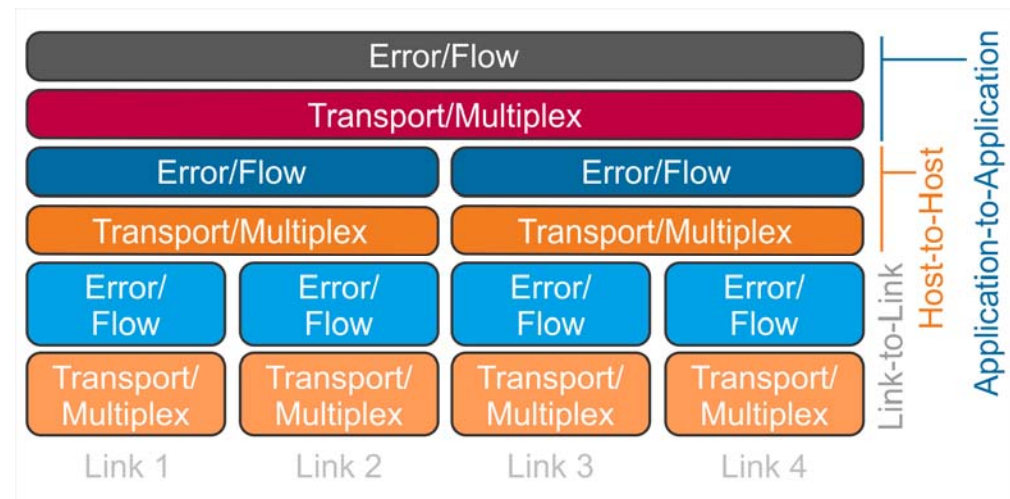
# We Layer Protocols...

- Simplifies protocol design
- Consistent/understandable API's
- Splits failure domains



# We Layer Protocols...

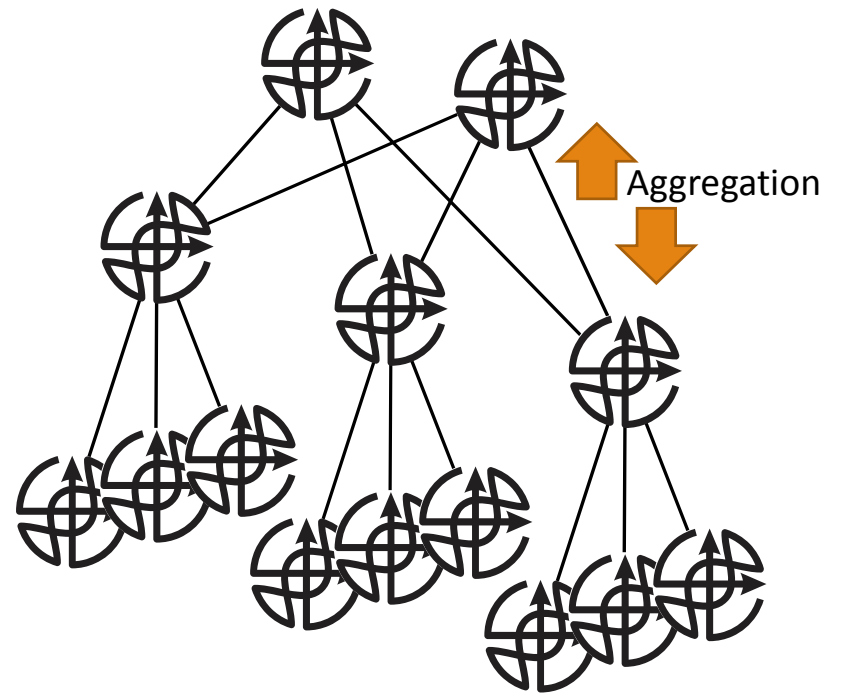
- An iterated model (RINA)
- “Lower” Protocol
  - Transport (Marshaling)
  - Multiplex
- “Upper” Protocol
  - Error Correction
  - Flow Control



# We Layer Topologies....

---

- Simplifies problem spaces
- Breaks up failure domains
- Focuses problem/solution sets



# We Layer Lots of Other Things...

---

- IGP/EGP
  - EGP – customer routes, policy
  - IGP – internal connectivity, management
- Hierarchy
  - Reachability aggregation
  - Repeatable configurations
  - Failure domain management
- Flooding Domains
  - Topology aggregation
  - Failure domain management
- Layer 2/Layer 3
  - Aggregation
  - Failure domain management
  - Broadcast domain management
- Virtualization
  - Separate customers, applications, etc.
  - Failure domain management
  - Broadcast domain management
- Etc.

---


Why not layer the control plane?

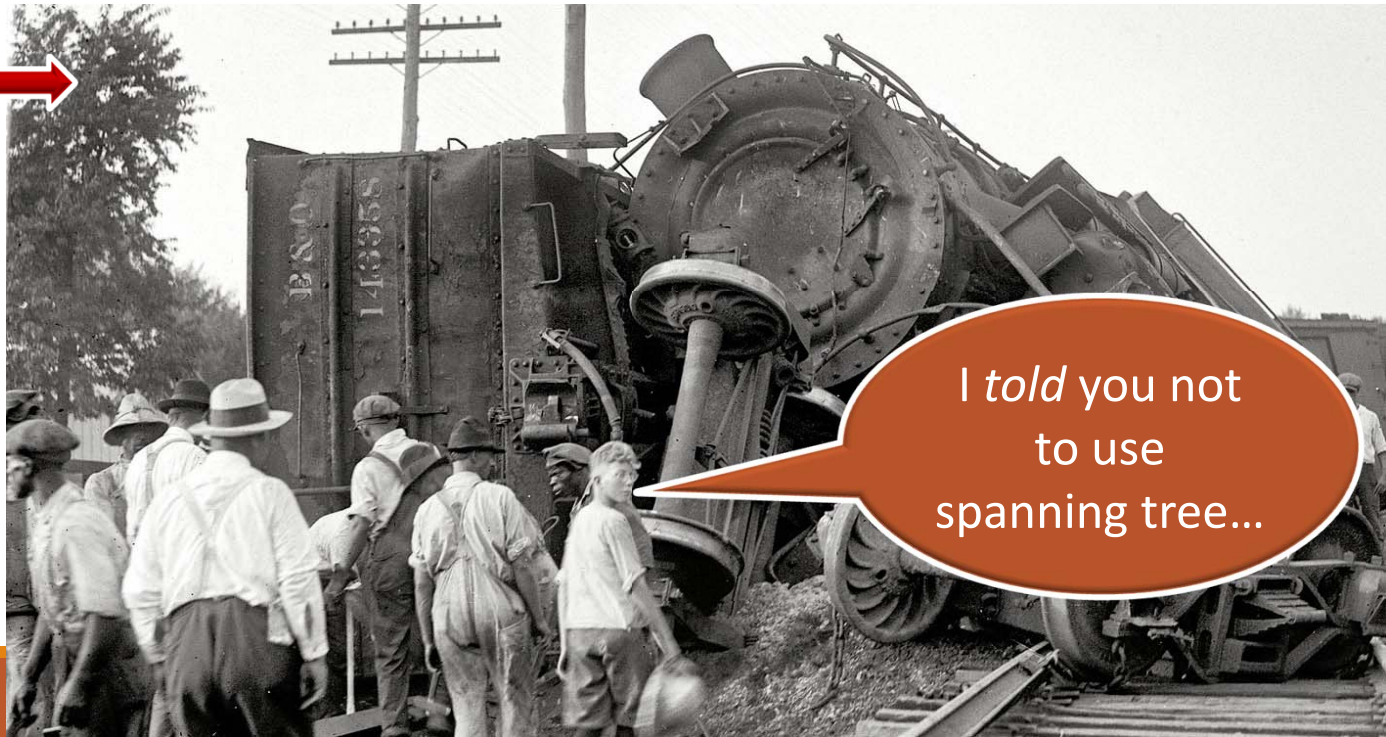
---



# What do I Expect a Control Plane to do?

---

- A “handle” to implement policy
- *Not* do this 
- Shortest Path
- React to Change



*I told you not  
to use  
spanning tree...*

# Data Plane Policy

---

- Quality of Service/Queue management
  - Deep Packet Inspection
  - Filtering (stateful or not)
  - Load balancing (some instances)
  - ...
- *Data plane policy as anything which modifies the handling of packets on a per hop basis to improve or control service*



# Control Plane Policy

---

- Aggregation
  - Traffic Engineering
  - Service Chaining
  - Tunneling (Virtualization)
  - Load Balancing (some instances)
  - ...
- What one thing do all these have in common?
  - They each have the potential to *increase stretch*
  - *Control plane policy is anything which modifies the forwarding path off the shortest path to achieve a specific goal*

# Control Plane Policy != Reachability

---



These...  
are never  
printed on  
this...





- Can we Split Policy out as a separate layer?
- If policy  $\neq$  reachability, why not?
- An experimental proposal...

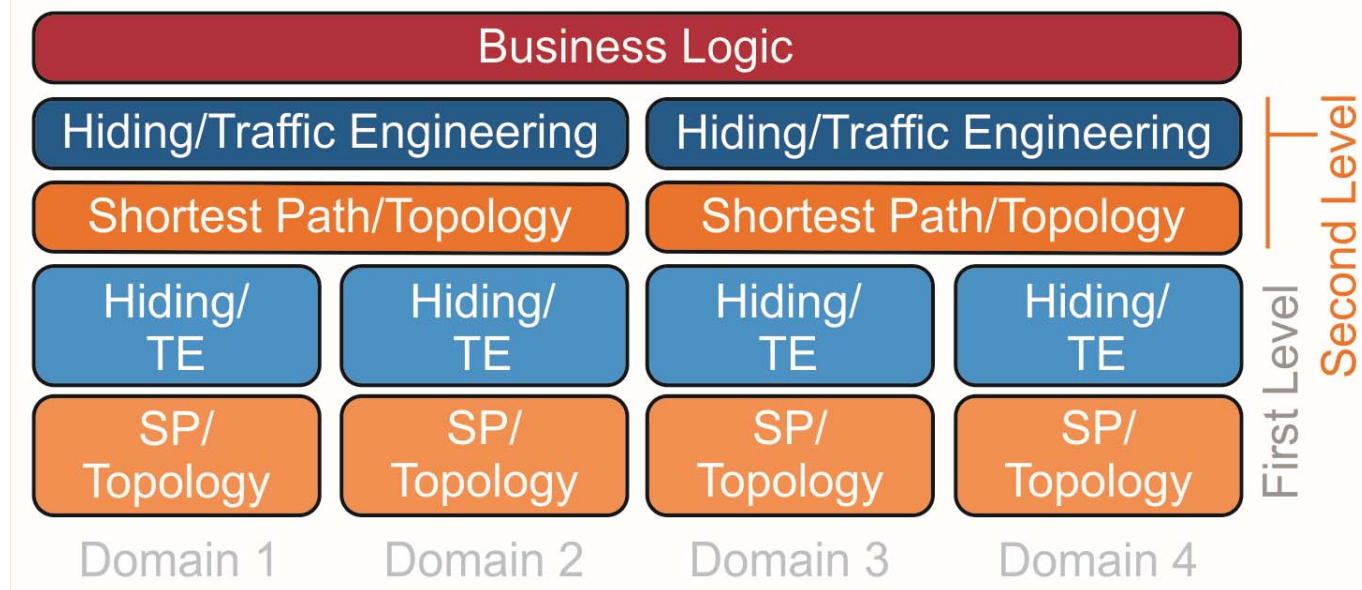
# Centralized vs Decentralized

- A “handle” to implement policy
- *Not* fail
- Shortest Path
- React to Change



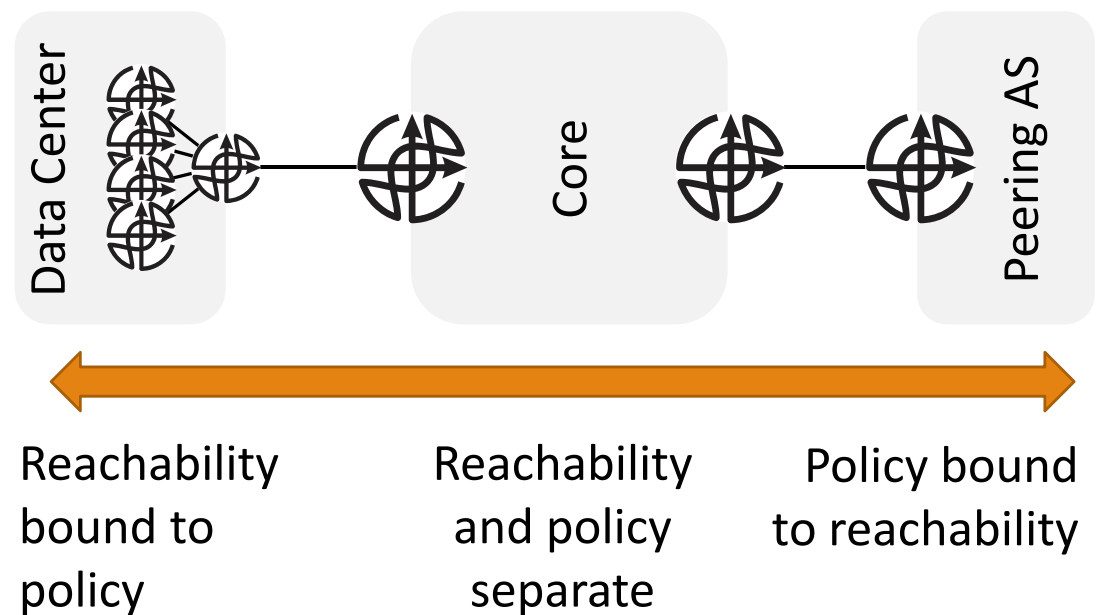
# An Experimental Proposal (1)

- The control plane can (potentially) fall into two/three layers, then...
- “Ultimate” layer
  - Business logic
- “Upper” layer
  - Hide information
  - Engineer traffic & flows
- “Lower” layer
  - Determine reachability
  - React to topology change



# An Experimental Proposal (2)

- Consider each slice of the network in terms of *problem bounding* and *APIs*
- The “right solution” might vary



# Whither? (Ask Hard Questions)

---

- Is this the right model to move forward?
- Ask hard questions
  - *Will we really be better off troubleshooting an API and code than a protocol?*
  - *Can polling ever outrun local reaction to events?*
  - *Should policy always (or ever) be expressed in terms of forwarding?*
  - *Is a unified view of the network actually possible?*
  - *What should we do about failure domains?*
  - *Are we ready to throw away 20+ years of experience to “try something new” across the board?*

# Whither? (Community Action)

---

- Support community efforts
  - YANG/NETCONF
  - I2RS
- Think about existing protocols for southbound
  - BGP, PCEP, ??
  - What are the gaps?
  - What needs to be done?
- How does this fit into open source?



---

# Questions? Thoughts? Tomatoes?

---

Russ White

<http://ntwrk.guru>


<http://packetpushers.net>

<http://www.ericsson.com/spotlight/cloud/blog/>

<http://ipj.dreamhosters.com/>

russ (at) riw (dot) us

@rtggeek



Where's the  
squirrel? I was  
*promised* a  
squirrel...