

# How to Bring Calm to an Otherwise Too Exciting Networking World



Amin Vahdat  
ONS 2011

# Networks Today Are Very Exciting



# Networking State of the Art



- I am on the edge of my seat when:
  - A link fails
  - I perform a switch software upgrade
  - Multiple applications try to share the same network
  - I add a new customer to an existing network
  - I try to expand my network fabric
  - I set security/isolation policy
  - I migrate a virtual machine
  - I need performance isolation
  - I define a new traffic engineering policy
  - I introduce a product from a new vendor

# Networking Reality for the Masses



- I am on the edge of my seat when:
  - A link fails
  - I perform a switch software upgrade
- I pray to my favorite deity when:
  - Multiple applications try to share the same network
  - I add a new customer to an existing network
- I don't even bother trying to:
  - Expand my network fabric
  - Set a security/isolation policy
  - Migrate a virtual machine
  - Enforce performance isolation
  - Use traffic engineering
  - Introduce a product from a new vendor

# What We're Left With



- Overriding philosophy: if we don't touch the network, maybe it won't break
- Fundamental *under-investment* in network
  - Unbalanced systems\*
  - Millions of \$\$ of compute and storage systematically under-utilized because network can't keep them fed
- A network that gets in my way rather than enables computation and data storage

\* See "TritonSort: A Balanced Large-Scale Sorting System", Alexander Rasmussen, George Porter, Michael Conley, Harsha V. Madhyastha, Radhika Niranjan Mysore, Alexander Pucher, and Amin Vahdat, Proceedings of the 8th ACM/USENIX Symposium on Networked Systems Design and Implementation, March 2011.

# What We're Left With



- Overriding philosophy: if we don't touch the network, maybe it won't break
- Fundamental *under-investment* in network
  - Unbalanced systems\*
  - Millions of \$\$ of compute and storage systematically under-utilized because network can't keep them fed
- A network that gets in my way rather than enables computation and data storage

**Companies would pay a lot of money for a network that adds to their ROI on their compute/storage infrastructure**

\* See "TritonSort: A Balanced Large-Scale Sorting System", Alexander Rasmussen, George Porter, Michael Conley, Harsha V. Madhyastha, Radhika Niranjan Mysore, Alexander Pucher, and Amin Vahdat, Proceedings of the 8th ACM/USENIX Symposium on Networked Systems Design and Implementation, March 2011.



# What We Should be Working Toward in Networking



# How to Get There From Here



- The driving philosophy behind networking has been:
  - Approximate the global optimum based on local view of events
  - Network fabrics are just a collection of individual boxes, each acting autonomously
- Full decentralization is beautiful and pure but also impossible to get right




# How to Get There From Here



- Leverage all of the progress made over the past 30 years in other disciplines of computer systems
  - Distributed systems for fault tolerance, scale, and logically centralized control for global optimum\*
  - High-level Programming Languages to specify network-wide intent
  - Correctness proofs to enforce invariants
  - Extensible databases to maintain network schema
  - Staged, coordinated fabric-wide upgrades for hitless reconfiguration

\* One starting point: Sanjay Ghemawat, Howard Gobioff, Shun-Tak Leung, "The Google File System." Proceedings of ACM SOSP, October 2003

# How to Get There From Here



## *OpenFlow and Software Defined Networking*

Word of caution: tectonic re-imaginings will mean a few steps back before moving forward