

Jupiter Rising: A Decade of Clos Topologies and Centralized Control in Google's Datacenter Network – Public Review

Ming Zhang
Microsoft Research
mzh@microsoft.com

This paper provides a retrospective view on the evolution of Google's datacenter networks across five generations over the past decade. The network design is largely driven by three key decisions: 1) adopting Clos topology to achieve high scalability and failure resilience; 2) leveraging merchant silicon switches to cut down network cost; 3) replacing conventional distributed control protocols with centralized ones to reduce management complexity. The results are impressive: the network bisection bandwidth grows by three orders of magnitude (from Tbps to Pbps) within just ten years.

The PC agreed that many of the design decisions described in the paper are considered best practices nowadays and have been studied in earlier research papers. However, the core value of the paper lies in the unique experiences in building, deploying and operating real-world large-scale network infrastructure, which is very difficult to duplicate by the networking community. It will serve as an excellent reference paper for any researchers who work on datacenter networking and establish a baseline for future research proposals in the field.

One notable feature of the paper is that it covers nearly every important topic in datacenter networking, such as topology, switch design, routing, inter-cluster connectivity, incremental deployment, switch management, configurations, upgrade, and troubleshooting, many of which are hard to study in academic settings. Despite of the breadth of the topics, the authors did a fantastic job in maintaining the paper coherence while still providing sufficient level of technical details.

It is interesting to observe that Google's networks evolve almost in parallel with the advances in datacenter networking research in the community. For instance, Firepath, the routing protocol for Firehose, Watchtower, and Saturn fabrics, adopted centralized network control, which is also a centerpiece in broader software defined networking (SDN) efforts. The fact that Google is able to run centralized control protocols in production makes a strong case for the applicability of SDN in large-scale networks.

Besides documenting things that work, this paper also describes things that do not work and the lessons learned from their decade-long experience in building and operating datacenter networks. Mindful readers may want to pay attention to why Firehose 1.0 failed and how things were turned around in Firehose 1.1, e.g., by replacing regular servers with dedicated single-board computers to improve stability. The network fault scenarios described in Section 6 represent yet another attractive part of the paper, as they shed light on areas where future research may generate huge practical impact.