Wireless Engineering Research and Education Center

# A Transport Layer Solution to Address TCP Incast in Data Center Networks

**Santosh Kulkarni**

Graduate Student, Department of Computer Science and Software Engineering
Auburn University, Auburn, AL

## Abstract

Modern data centers host tens of thousands of servers that communicate mutually using high speed network interconnects. Cost and compatibility reasons persuade many such centers to consider Ethernet for their baseline communication fabric. Until recently, Ethernet speeds inside data centers averaged 100 Mbps. However, evolution of IEEE 802.3 led to the development of 1 Gbps and 10 Gbps Ethernet. Such high speed Ethernet requires proportional scaling of TCP/IP processing for network intensive applications to really benefit from the increased bandwidth. While IP is expected to scale well in this context, TCP is known to have problems supporting very high data rates at very low latencies.

One such shortcoming is the catastrophic collapse in TCP's throughput that occurs when the number of servers sending data to a client increases past the ability of an Ethernet switch to buffer packets. In this talk we investigate the root causes for TCP's throughput collapse in data centers and examine a transport layer fix that is effective in addressing this problem.

## Bio

Santosh Kulkarni is a graduate student in the Department of Computer Science and Software Engineering, Auburn University. He received his Masters degree from AU in May, 2009 and is currently pursuing his doctoral degree under the guidance of his advisor, Dr. Prathima Agrawal. His active areas of research include Data Center Networks, Cloud Computing, Wireless LANs, Cooperative Networks and Distributed Computing.

## MONDAY, OCTOBER 24, 2011, 4:00 P.M.
## 235 BROUN HALL

http://www.eng.auburn.edu/~pagrawal/seminar/2011