



Thoughts about TCP offloading

Jeff Mogul (HP Labs – not speaking for HP!)

My 2003 HotOS paper title:

TCP Offload is a dumb idea whose time has come

Is TCP offloading still a dumb idea?



- **Yes, if you misunderstand the problem:**
 - Don't count instructions; worry about the bus
 - Offloading is good if it eliminates lots of bus traffic
 - Short connections don't matter
 - You won't save anything; you'll run into NIC resource limits
- **No, if offload enables something really useful**
 - RDMA, since it avoids buffer copies & bus traffic
 - Or if you can put encryption on NIC, as well
- **Think about connection handoff**
 - Rather than offloading entire life-cycle of connection
 - A way to avoid spending overhead on short conns
 - Policies matter [e.g. Kim & Rixner, OSDI 2006]

Trends and challenges

- Moore's Law (if it holds) works against offloading
 - General-purpose instructions are a better deal
 - Until the silicon on a NIC costs less than the package
- NICs moving into the CPU
 - No more PCI bus? Much less need for RDMA
- CPU "I/O Assistance" functions might suffice
 - Async bcopy, TCP segmentation/checksum
- Offload vs. CPU virtualization
 - We really don't want VMM to know about connections
 - But RNICs/TOEs have hard resource limits
 - How does VMM allocate connections among VMs?