

June 17th 2013

Workshop on Reducing Internet Latency Position Paper

Lucien Avramov and Jacob Rapp
Cisco Systems Inc.

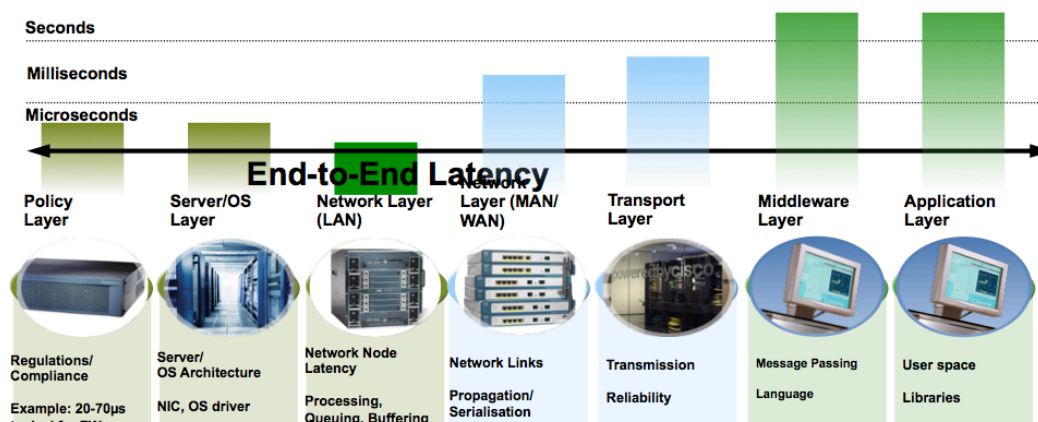
1) Introduction: about us

Reducing the latency on networks has been a primary focus for us in the data-center space. We work in the Cisco Nexus Business Unit as Engineer and Manager. In our product line we essentially focus on data center switching and look across the entire stack including application and compute. We have been designing Ultra Low Latency networks for High Performance Compute [HPC] environments and High Frequency Trading [HFT] where reducing the end-to-end latency is key. Presenting at HFT and HPC conferences on a regular basis and bringing a network centric point of view of the latency reduction of the Internet is our areas of expertise.

We have built a 200+ server cluster to model applications and what effect various latency characteristics have on the application. One example is the effects of Big Data applications, such as Hadoop, in the data center. What happens with large throughput sensitive flows have to compete with latency sensitive applications?

Today at Cisco we are capable of building with our products networks of 10GE and 40GE port to the host up to 9216 ports with a end-to-end any port to any-port latency of 1.4 usec. This is the most dense and lowest latency that can be build today on Ethernet networks.

We believe this workshop is very relevant to our skillset and knowledge in the industry build with solutions for all our customers in different sectors.



Many variables involved—key is to profile end-to-end trading process
Figure: showing the latency order of magnitude across the OSI stacks

2) Benchmarking network equipment and its relevance to latency

To be able to reduce the latency, it's key to understand how to define and measure it. We are actively working on IETF drafts for data center benchmarking and latency is a strong element: how to compare different vendor equipment or different network architecture DUT latency's in

order to provide an apple to apple comparison to the end user is what we strive to accomplish with our IETF contribution and publications.

For example we recently submitted after feedback from the BMWG group the following draft:

<http://www.ietf.org/id/draft-dcbench-def-00.txt>

3) 10 internetwork design principles to reduce the latency

3.1) The speed: the faster the network the lower serialization delay and latency

3.2) Physical media type: copper twinax cables are faster today than fiber optic, and microwave is a faster media than fiber for interconnects, for example Chicago / NYC with microwave saves a considerable latency compared to the traditional dark fiber between the two cities.

3.3) Switching mode: cut through provides a predictable performance across packet size, versus store forward

3.4) Buffer amount in the network: what is the right buffer amount to provide high performance? Buffer bloat is affecting data center latency numbers. Large Throughput sensitive TCP flows are building up queue depth and causing delays for smaller latency sensitive flows.

3.5) The feature-set used on the network equipment has a direct impact on the end to end latency, for example protocols such as CDP, STP, LLDP etc. contribute up to 2.5x more latency than when they are not used.

3.6) Rack mount servers are lower latency than blades, and non-virtualized OS also saves latency

3.7) The CPU / Memory choice in the server does matter as it will dictate the performance of the compute.

3.8) The network adapter card and protocol used will have impact that can bring latency up to 4x time less [from 20 usec to 5 usec]

3.9) Visibility and analytics are key in understanding the effects of latency. Precision Time Protocol, PTP IEEE 1588 v2, helps provide a precise clock across networking and compute devices to gain insight.

3.10) Security increases latency drastically, there are ways to get around this in the network, and we can discuss it.

4) Design case studies that reduce the latency

-the HFT and HPC case study, where we can explain how to design an ultra low latency network in this sector

-Big data case study, where we can explain the effects of throughput sensitive traffic vs. latency sensitive traffic for applications such as Hadoop

Contact us:

Lucien Avramov: lavramov@cisco.com

Jacob Rapp: jarapp@cisco.com