

Notes from TIP 2013 in Honolulu

Performance Working Group at TIP 2013 in Honolulu

<http://events.internet2.edu/2013/tip/agenda.cfm?go=session&id=10002739&event=1261>

January 14, 2013

AGENDA

Welcome – Aaron Brown
Performance and Virtualization - Alessandra Scicchitano
A Performance Potpourri – Brian Tierney
Centrally Managing perfSONAR Tests for Organizations and VOs - Aaron Brown
The Effects Of GSO/TSO On TCP Performance – Andy Germain
Community Updates / Open Forum

DISCUSSION

Performance & Virtual Machines,

speaker: Alessandra Scicchitano, SWITCH

eduPERT is the (GEANT) European group for the performance people in the European NRENs.

When you have performance issues on virtual machines, it can be very hard to know how to start. Performance measurements are skewed because the real machine isn't accessible. The speaker presented a use case, and the lesson learned was that the choice of virtualization technique impacts performance debugging. For these kinds of problems, where the issue is between the mother machine and the virtual machine(s), there is almost no help available.

Update of ESNet perfSONAR Activities

speaker: Brian Tierney, ESnet

slides: <https://spaces.at.internet2.edu/download/attachments/10402/20130114-tierny-perfsonarupdate.pdf>

Recently upgraded ESnet backbone to 100G. ESnet runs perfSONAR everywhere. Just released beta perfSONAR-PS Toolkit 3.3-rc1, a major overhaul. More to come in next few weeks, including integration of Web10G. ESnet has 1.5 FTEs dedicated to perfSONAR support. ESnet has 80 machines deployed. They run BWCTL and OWAMP on separate hosts so that throughput tests (BWCTL) don't mess up latency tests (OWAMP) or vice-versa.

New MaDDash dashboard

New lookup dashboard

Centrally Managing perfSONAR Tests for Organizations and VOs

speaker: Aaron Brown, Internet2

slides: <https://spaces.at.internet2.edu/download/attachments/10402/20130114-brown-meshconfig.pdf>

Problem: each perfSONAR server is a testing island. How do we avoid doing the same test between servers, or missing a path in a mesh? Solution: a centrally-managed description of the "mesh", in JSON format. Each server downloads the the shared description periodically. This is integrated with MaDDash.

Looking for interested beta-testers.

The Effects Of GSO/TSO On TCP Performance

speaker: Andy Germain, NASA

The TSO feature of modern NICs will screw up tcpdump because a large packet is sent to the NIC and is then fragmented by the NIC. It makes sense to offload the CPU.

Problem: it makes tcpdump harder to use. Comparing a tcpdump to a dump from an external sniffer won't work.

Bigger Problem: the NIC can fragment a 64k "packet" into 46 MSS-sized packets and send them at wire speed. This can overload downstream devices, causing dropped packets, which causes very slow performance. Best guess as to why is that an input buffer isn't big enough to handle the burst of packets.

Recommendation: turn off TSO. Andy saw these problems at 4 separate sites. Turning TSO off improved performance: packet loss went down.

Next Meeting

Next Meeting of the Performance WG will be at the Internet2 Annual Meeting. <http://events.internet2.edu/2013/spring-mm/>

Hope to see you there.