Sunday, January 11. 2009

### Routing fish

At first ... i´m not aware of an upcoming OpenRouter product. But that means nothing. I´m not aware of all project, albeit i think i´m reasonably well informed. Enough disclamer stuff. The discussion on the fishwork article lead me to an interesting thought. How about building a high performance router out of a 5120? Most of the stuff is already there ... The system has two 10GBe interfaces and four 1 GBe interfaces. The 10GBe interfaces are on die, directly coupled to the crossbar, so you have a low latency there. A nice configuration to start with. For more interfaces i would suggest the usage of a "router-on-a-stick" configuration: Activating VLAN-tagging on the 10GBe-Interfaces and use a normal switch as a fan-out unit.

So what can we do software-wise with the system: The standards ... obviously ... like routing or channel bundling. Furthermore: We have Crossbow with it´s virtual nics, it´s virtual switches and so on. We have Zones. So we could divide the router in seperate administrative domains. With Crossbow you can do traffic management as well and controlling bandwidth consumption.

There is Quagga as an open-source routing protocol suite available in standard Solaris: OSPFv2, OSPFv3, RIP v1 and v2, RIPng and BGP-4. There is a full-fledged IPsec suite in the OS and 8 crypto units in the CPU. It should be possible to build a decent highspeed SSL-Accelerator with kssl or apache2. There is a IPfilter firewall included in Solaris. The project to extend the networking subsystem with more dtrace probes is in it´s full run. And we have this excellent GUI with Fishworks. Instead of measuring disk latency you would measure link latency.

Combining all this stuff should result in an really interesting package. Sound like a decent idea .. i thini i will invest some further thoughts on it ...

Posted by Joerg Moellenkamp in English, Solaris at 22:12

Too bad the the decision for a firewall in Solaris was to IPFlter and not the much better PF from the OpenBSD project. NetBSD and FreeBSD switched to PF a long time ago. PF is much better to configure and offers more features than IPFilter.
Anonymous on Jan 12 2009, 00:36

s/5140/5120. The 5140 is based on Niagara T2+ CPUs without on-die 10GBit. This was removed from T2 to T2+ CPU for the Cache Coherency Link
Anonymous on Jan 13 2009, 17:57

Sorry, you are correct .... typed 5240 all day long for a quote ... freudian typo ...
Anonymous on Jan 13 2009, 18:01