

Thursday, January 21. 2010

Less known Solaris features - IP Multipathing (Part 1): Introduction

root was sitting in SuperUser castle and everything was fine in the kingdom. But then a loud squeaking and creaking noise found root's attention. The demons of hypertext wrote into the Scrolls of Log, that they couldn't fulfill their work any longer, as the sole bridge into the vast kingdom of root lowered at this moment was broken.

Root spoke: "Use the other bridge ... there are two for a reason. Do I have to think for you all?". But the demons replied: "We can't do that ... only the infinite power of root can lower the bridge". Thus root lowered the bridge but thought "I have to do more important things than lowering bridges".

Thus root spoke a chant of infinite power and a daemon was spawned from the ether. root told the daemon "You are the guardian of the link! Protect it. Guard it. And when everything else fails, you are allowed to lower the second bridge to SuperUser castle."

IntroductionBefore people start to think about clusters and load balancers to ensure the availability, they should start with the low hanging fruits. Such a low hanging fruit is the protection of the availability of the network connection. Solaris has an integrated mechanism to ensure this availability. It's called IP Multipathing.

IP Multipathing is an important part of the solution for an ever reoccurring problem, as almost all applications interact with the outside world on one way or the other. Thus ensuring the mechanisms of communication is a part of almost all architectures.

Even when you have other availability mechanisms like balancers, you want to use a protection of the IP connection out of a simple reason: Many applications have a session context and not all software architectures can replicate those session contexts to another system to enable a failover without losing the session. So do you really want to lose this context just because of a failing network card or because of an admin unplugging a cable? Or do you really want to provoke a cluster failover because of a failing network card? IPMP can keep such failures on a low level without needing high-availability mechanisms with a much larger impact.

Out of this reason IP Multipathing is an important part for most HA infrastructures. This tutorial wants to give you an introduction in this topic. It's not really an "less known feature" because for many people working with Solaris, IPMP is a daily part of their work. But many people new to Solaris or OpenSolaris aren't aware of the fact that Solaris has an integrated mechanism for IP Multipathing (As well as most newbies to Solaris aren't aware of MPxIO ... the counterpart of IPMP for storage). Furthermore this tutorial wants to give some insights into new developments in the field of IP Multipathing.

Where should I start?This tutorial will explain two mechanisms, because the realm of "IP Multipathing" is a topic in flux at the moment. The implementation in Solaris 10 and older releases of OpenSolaris (before Build 107) is vastly different to the implementation in current releases of OpenSolaris(Build 107 and up).

I thought a while about the problem, what method should make the start in this tutorial. At the end I decided to explain the new IPMP mechanism first as the concepts of multipathing are a little bit more obvious in the new implementation.

Posted by Joerg Moellenkamp in English, Oracle, Solaris at 20:09

great introduction!

Anonymous on Jan 21 2010, 20:33

sehr geiler anfang!!!!

Anonymous on Mar 6 2010, 14:04