2nd SoC workshop CERN - Minutes – June 11, 2021 – SysAdmin

https://indico.cern.ch/event/996093

Notes:

- These minutes were collected by P. Zejdl.
- They provide a record of the questions and answers following the presentations. For the presentations themselves including recordings, please, refer to the indico timetable.

Friday – June, 11th:

- M. Dobson, System and Network Administration Aspects of Zynq MPSoC Devices in the Experiments:
    - Q: In Atlas phase 1, my group and students would benefit from the VM (with the basic services) for SoC. There is a lot of processing power on SoC, in most cases at the moment it is underutilized. I wonder if part of the planning is to use them, in addition, for data aggregation, network monitoring, or to send summary information?
      A: I've noted down your request for VM / Docker for running the services. Indeed, it would ease, how much people have to learn, and clearly, it's not a hands off, because there is still some configuration needed, but it should make life a little bit easier for sure.
      A: There is potentially room for running quite a lot of things, but I think most people want to reserve that in case they need it for what they do in the crate. But I agree that there is room there, I'm not sure that we're the best people to say "Can we do some aggregation" because it really does depend on the traffic. So I think this is something that should be discussed more widely about how we run the systems in general (system administration monitoring information versus run control level monitoring, etc.)
    - Q: For Atlas phase 1 ATCA: Atlas is not using naming convention today, would be very nice to agree on ATCA shelf address naming scheme so IPMC can be programmed etc. We have IP BUS that does DHCP! (Not yet with DHCP Client ID, but can be done).
      A: Sysadmins of both experiments are busy with preparations for RUN3. But one can use an DNS alias which uses a naming convention and keep the names which Diana has today in the scheme. This will allow you to allocate today the right names.
    - Q: We are bootstrapping off the shelf address to establish the addressability. So it is the shelf address that we need to sort out now.
      A (D. A. Scannicchio): For the naming convention, we will have a transition period, and we need to agree on the naming convention to be adopted. I would like to mention that running a (custom) DHCP server in ATCN (ATLAS Technical and Control Network) is forbidden. Any DHCP server would be illegal.
    - Q: We are using a DHCP server in GPN.
      A (D. A. Scannicchio): So, we have make this is clear… we can have a chat.

Later on, of course. But yes, keep in mind that things can be different between GPN and ATCN, because of all the rules.
A (M. Dobson): I will just to say I offer my help to Diana to try and progress on this also for client ID, but as I said, we're, we're both stuck with RUN3 stuff, which is maybe more urgent. But certainly if we can we help. We already have ideas of how to do some of this.

- Q: Is it possible to have DHCP Client ID early, before the kernel needs to bind to the NFS? (Kernel also needs to get IP address in order to speak with NFS)
  A: Yes that's possible. I believe it was shown in Petr's talk on Wednesday. The SIPL implementation written by R. Spiwoks provides functions in FSBL and U-Boot. In U-Boot we use the connectivity to the IPMC to read out the shelf address, and the slot number, and then construct the client ID, and use it for the DHCP. Then pass all info to the kernel as it boots, and it can use it.
- Q: We currently still intend to avoid an NFS booting on our systems, although we would certainly be interested in any read only root FS… But our reasoning for this is that the NFS server introduces a strong single point of failure. Is there any capacity, to have redundancy there, or automatic fail over of some kind that would basically remove that single point of failure from the equation?
  A: We use NetApp Filer, which serves everything over NFS and that is fully redundant. So this is something that we could clearly use. This is something that we have in the back of our mind. We're not sure if that's the way we want to go but certainly it's an option.
  A (D. A. Scannicchio): Already today, yes, local file server could be considering a single point of failure, but we have monitoring place so in case something happens we can very easily move the network device to another local file server. So it is just a time to reboot, netboot, and they come back.
- Q: It makes sense. My main concern is just, well, during a run. At the best that's going to take a few minutes and that's a few minutes of loss data.
  A (D. A. Scannicchio): If we have such a failure though we have a kind of bigger issue, all around the experiment, of course, and I want to stress out, that we are closely monitoring, or the device exactly to pick up any issue before it becomes a real issue.
  A (M. Dobson): At CMS as well. … We have devices (servers) that are already passing run control information to the boards. These devices are anyway needed in the run. If the device like would go down today, we'd have the same problem. So, I'm not sure we're introducing necessarily more single points of failure. So again I am not particularly worried. But I do understand where you're coming from, and netboot is there mainly to limit the SD card ageing.