

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

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

Notes:

- These minutes were collected by M. Twomey.
- 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:

- A. Kazarov, ATLAS TDAQ Online Software Plans for SoC:
 - Q (D. Miller): Right, so the multi processor versions of those Xilinx SoC's for example have a real time processor have the ARM processor, have the GPU. And even more modern devices, or more newer devices have, you know, a variety of, you know, (ai cars and engines and different, different devices on the chip), of course. So, Does, does that change any of the picture that you described here does that the potential need to execute? You know code on a GPU on a SoC, or to utilize the real time processor does? Does that, does that change any of the technicalities that you've described here? I apologize, I'm just not fluent enough in some of this to understand whether there's an impact there or are you really focused on the ARM side and as long as that can control the other components on the chip It's okay? So I hope my question is understandable, in some sense. Thank you.
A: Well from from this perspective for this interface, as a SoC system, is just a Linux box for us, whatever you can do from standard C++ invocation, whatever libraries. You can use to access your GPU, as a process also as an admin console. You can just use it in, in your normal way... shouldn't affect the design of this communication layer.
 - Q (D. Gastler): You talked about not having ROOT and having a lightweight, histogram plotter, I was just wondering, if you do have that, could you possibly split it off so we can put it in the sub module? Because I don't really want to have root installed, but it would be nice to have some histogramming capabilities. I don't want to rewrite it myself. So, is that possible?
A: You mean on the client side or on the server side?
Q: I just mean the code itself. I mean did you split off you histogramming code into a submodule? Did you make the plotting software at the library itself in a standalone repository, so other people can use it outside of your framework?
A: Well, the library itself doesn't define any histogramming capabilities right. It may be just a vector of characters, which you can convert to your histogram yourself. If it will be an additional library which can help you with doing that, without ROOT.
Q: Ok, so it's not really a histogramming library? It's just any D-array library.
A: It's not a histogram library at all. No. It's not a replacement of ROOT or anything like that, no.

- Q (W. Vandelli): Well, maybe more than a question is a suggestion or, an idea I mean, yes you went through these requirement collection procedures, there is a document that is all very good you, to know it's good software, good software processes, but we all know that the collection requirement is not an ideal process and we know from experience it's not your fault that will be missing features in this document, people didn't think about some missing requirements. Now I appreciate that you already have a test infrastructure running the ARM in a VM, that's very good. I was wondering if you would be interested in having early customers testing your products, among the people, potentially among the people in the audience?

A: Well, what you imply is that it should be published somewhere in the repository with basic instructions how to, how to run it.

Q: I'm not implying only that mean, I am effectively... the question isn't only, I guess it's not only to you But more generally to people if they ...people listening these these words, if they are interested in this kind of development that you're doing they should already get in touch with you as early as possible. And I'm asking if you would be interested in getting that.

A: Of course. Yes.

A (G. Avolio): I maybe I can add a few beats in the context of the Level II scheduling and planning so this is a, this is exactly the one of the goals to achieve in the prototype so to have something that has some functionality that works in a reasonable way. And then to expose to the, to the people that need to use it at the end. And then, and then eventually move to the next phase because we are still a very at the very first phase, we don't have a design, we don't have anything so this is just the attempts and and that's and again that's actually the idea of having these, these kind of prototype available by, by October. of something that people can start using, and to get feedback about functionalities and all the rest from people on the field actually, yes.

- Q (R. Kopeliansky): Just following up to what Giuseppe and Andrei said, and does it mean that perhaps for phase one it will be possible for some systems to use it during phase one to test it or they are in production.

A (G. Avoli): Not in production, of course. The timeline is not the plot that one. No.

- Q (H. Sakulin): You mentioned in one of the earlier slides that you're planning that the server running on the SOC system should be able to execute arbitrary c++ code that you sent from your application. Can you say a bit more about this?

A: Sorry, no it's not the code which will be sent. And I meant just the code legs that that upon receiving a command. But you should be able just to add your code in response to the command and that's it.

A (W. Vandelli): Though sending arbitrary code; that would be cool. I agree with Hannes. If we could do it would be good.

- A. Iribarren, ARM64 Linux Support and Future CERN Linux:

- Q (Wainer Vandelli): I am wondering, you know, in the context of experiments where we work with, with many external institutions, that are colleagues that are

not CERN based. I guess I know the answer, but I guess is worth asking if any way I mean, I suppose, most of these developments that you are doing to support ARM at CERN, can be exported then to, to our collaborating Institutes?

A: Yeah all the repositories are publicly available. All of the software is there.

- Q (S. La Frasse): Is the Raspberry Pi considered as another one of the hardware solutions to test your release?

A: A RaspberryPi Four as an ARM 64 bit, that one would work. So it's not something that we are going to use at this because what we, I mean we want something that is installed in the computer center and so fully managed by all of the rest of our tools.

Q: Of course, of course, but I meant folder for the experiences; to have a small hardware, cheap hardware through, to test your release and try and begin working with it. For the experiences to test it.

A: Yeah, yeah, and you could definitely do it, I haven't tried I don't know it doesn't work. I did by once the very very beginning the Docker image that we were producing the Raspberry Pi at home, and it worked. I didn't try the full OS and everything else. But yeah, let us know if it works well.

Q (M. Dobson): CMS certainly, one of the groups I was tested using the Raspberry Pi 4. For indeed in storing the OS, building software, testing the builds that are done, a different way and things like that so yes, It has been done.

- Q (J. Tikalsky): I guess I was wondering if you have any plans to support 32 bit, ARM architectures? I know there's been some mention of the Zynx 7000 in this workshop, and we are running some ourselves. Are you planning to bring support for that or are you not intending to help those of us with Zynx 7000 devices.

A: No that...I'm afraid that's not possible for us at the moment, our upstream only provides ARM 64. And we essentially make all of the upstream repositories and then add our up on top. Adding arm 32 would be a very big undertaking.

Q: Can you remind we what you build source, or what your upstream source is? I know CENTOS has ARM 64 and ARM32 repositories, but I guess I just don't necessarily know what source you are using that only has ARM64?

A: For Stream 8 so when we checked initially it was not available. I don't know now if they have now started to build for it.

- Q (F. Vaga): So my question is not directly related to to this talk about a bit. So I'm wondering, since you are all having your system on to base the card Sir, if you measure the, the performance difference between a binary distribution like it could be CENTOS stream 8 and building from scratch from sources using the build ROOT or Yocto because I'm explaining this version like Red Hat Enterprise Linux 8 or DEBIAN or whatever it's a binary distribution that means that it transfers to any ARM 64 and to do that, they setup the compiler options for performances are set up to generic, meaning that they do not optimize for the specific system on chip that specific ARM that you use Cortex A9 or whatever. So they use a generic set of instructions so that can work on everything. So I'm wondering if some actually measure that this performance difference?

A (M. Dobson): Maybe someone can correct me but I think this has been looked

at, at least for x86 in quite a number of the experiments I suspect, at some point, to basically compare whether it is worth putting any specific flags on there or not. And if I remember correctly, the, the series is relatively low these days, at least on the x86 clearly for ARM, I don't know. So yeah, something maybe to be tested. I'm not sure we get more than a few percent difference.

Q (F. Vaga): Yeah for INTEL, for sure. It could make a difference. I don't know, I'm asking.

A (W. Vandelli): I think I think is a valid question for equal, and clearly seems nobody made the measurement so you're welcome to do it and presented at the next workshop. No I love what Mark was saying you know my background is more from software, software, etc. the reality finally is Yes, You may be right there may be a performance on impact, but if you want the centralized support for your operating system, and your software, there is no choice.

A (A. Iribarren): You would be seeing operating system is generic arm built, but the software that you run on in your software is tailor the software that runs on. But maybe that already gets you a quite a bit of the performance in fact.

A (R. Spiwoks): I just wanted to add that I agree with the previous speakers, actually, that it's a very valid and interesting question that I'm not aware of any benchmarks, comparing CentOS and purpose built operating system. And I would like to add that so far actually we've not been concerned too much about the pure processing performance of the SoC's we've been more concerned about the communication with our FPGA, where we've actually optimize the, the speed, and also getting data out of the S0C, using the Ethernet link in a very efficient way but for the processing, I think for what we are using. We are far from being at the limit of the processing power so it doesn't really affect us a lot.

A (L. Ardila): Yes, last, last year. Last year, in, in our contribution to this workshop, there was one slide, comparing various different processes, the sink, the regular sync with the dual arm. The single just kill against the some Express modules based on atom processors and also a common Express module based on Intel i5. And in thesis light there is some of the results for the for four different measurements, and there are two links with a bit more information. Perhaps it is not at the granularity of comparing different operating systems but the word some different versions of GCC compared they're like six, and eight. I think also for, which is the one that comes standard with central seven I believe I can put the link here.

- M. Docson, System and Network Administration Aspects of Zynq MPSoC Devices in the Experiments
 - Q (D. Miller): But the point being that I personally think you know my group, you know students and postdocs who also aren't super experts in this kind of thing would benefit a lot from the VM suggestion that you made, you know, two slides ago so I would just like, you know, to for what it's worth to put in my vote for that. And, you know, one of the thoughts that I had here again very perhaps naive and uninformed is, there's a lot of processing power on a lot of these chips and, and as many people have mentioned it in most cases at the moment it's

underutilized. Now, of course, I'm sure a lot of people like me want to reserve that for doing, you know, physics, but perhaps in addition, there's a possibility that there are processes that could be run on the ARM, as you say, kind of aggregation or or local metadata production that could then be distributed and I wonder if that is it all part of the planning, you know uniform kind of aggregation monitoring network control on the chip that could then send summary information, and perhaps keep local logs to be pulled you know at the points of downtime or something later on. I guess I'll I'm trying to say is that generally it seems like there's a lot that these chips can do that, you know, industry makes use of you see all these you know statements from Xilinx or Intel about what they're doing and, you know, I feel like maybe even though it's complicated in the space of possibilities as large we could try to investigate some of that as well so anyway thanks a lot.

A: Thank you. Yeah. So, certainly, I've noted down your, your request for VM slash Docker for 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. So I've noted it down and we'll think about having something available. And we'll post news as it comes to the various channels. Coming back to running things on the zinc so I agree potentially there's room for for running quite a lot of things. But I think your, your other side was the one which most people will take which is, I want to reserve that in case I need it for for what I'm doing 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 indeed that we should, that should be discussed more widely about how we run the systems in general

- Q (D. Sankey): We are currently installing shelves, where our modules are shelf address aware. So really on your slide 18. We need to decide, we would really like to decide now men because what I'm doing with. Send it not supporting client ID isms associating CERN OUI MAC addresses. With each slot in my shelf. But we need to me. This is the last workshop we had the same conversation, and we did. We really need to decide. Going forward, and notice is this right at the moment. Yeah, and it's not using this kind of name naming convention today. So this probably isn't so much question to you because CMS using his name convention. But we really need to certainly be very nice by obviously we can change, but we're very nice to agree a ATCA shelf address naming scheme. So that I can program like MCs to correct to recognize the correct shelf addresses the bonus for you. Is that I now have a version of IP bus that does DHCP.

A: Yes. So, Diana may want to answer maybe I give a forward question. I mean, we have discussed with Diana no later than this morning on this. And as I've mentioned. Both she and and us busy with run three stuff at the moment. And and this is unfortunate, because we, we don't have too much time to spend on on these things which CMS is phase two. Atlas is phase one but again it's very big changes if you will. Now what I can suggest maybe as a, as a potential way

forward is, as I wrote down there, you can have DNS aliases to have a more readable name. What you could also do is today. Use an alias which uses a naming convention and keep the names which Diana has today in the, in the scheme for example. So, this will allow you to allocate today the right names. And when you enter is ready with client ID, then, and this naming scheme then we could just move to it, because today it's defined it in ages to what she normally registers for you know i mean the module names are no problem because obviously I can use the current I was naming scheme for that it's, it's ready. Some bootstrapping off the shelf address to establish my credibility. So, yeah, then it is the shelf address that already need to sort out now. But I think there's the onus is no is if we think we're going for the sort of building slot. You name for the shelf. Then, I just need to pick one or the other two at a time we can run Yeah, actually, we understand Yep. Yep. Yeah. And this gentleman's got her hand up so I'll let her answer. Yeah. I think because at the naming convention. You already answer you, we will have a transition period. And besides this we need to agree. First of all, on the naming convention to be adopted.

A (D. A. Scannicchio): I think that for the naming convention. You already answer you, we will have a transition period. And besides this we need to agree. First of all, on the naming convention to be adopted. And I would like to mention as you have a failed to DHCP server that you have running on your device that are in GPN, in ATCN this is forbidden. So, any DHCP server will be illegal.

A (M. Dobson): I will just to say I I offer my help to Diana to try and progress on this also for client ID, but as I said, we're both stuck with run three stuff, which. Yeah, which is maybe more urgent, but certainly if we can help. We will help you to try and get something. Something working for you.

- Q (J. Tikalsky): So my question is a little bit along slightly different lines, you were mentioning the possibility of network based booting for the Linux systems on the boards and that is something that various people have been interested in. And my question is where two parts. I had the thought on the subject of these geographic IP addressing and naming schemes of the question of the kernel's super early, before any user, LAN IP address. Is it possible to have client ID interoperate with that to get the address found before the kernel's needs to bind to the NFS server?

A: Yes, yes, yes that's possible. It was, I believe shown in Petr's talk on Wednesday. The simple implementation that we used was written by Ralph. And it provides functions at various levels. One of them is the first stage boot loader. The other is the second stage boot loader and u-boot. And that's where we use it. 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 .

- Q (J. Tikalsky): My second question was along those lines. Um, we currently still intend to avoid an NFS booting on our systems, although we would certainly be interested in any read only request work is strong, sort of, preserve life's like shopping cart ID cards we intend to do that anyway. But our reasoning for this is that the NFS server introduces a strong single point of failure. Is there any

capacity, I'm not sure I'm aware of any to have redundancy there, or automatic fail over of some kind that would basically remove that single point of failure from the equation?

A: So, the Linux point of view, you could try and do something, the server level, although in NFS have, to some degree, a statefull. So you may not be able to achieve what you want. However, what we are sure about is that we use a NetAPP Filer, which serves everything over in NFs and that is fully redundant. So this is something that we could clearly use for this kind of thing. 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): If I may, I'm additionally already today Yes, local file server could be considered a single point of failure, but we have monitoring place so in case something happens we can very easily move the netbooted device to another local file server. So it was just a time to reboot the, the net booted, and they come back. Yes. 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 an issue, a real issue.

A (M. Dobson): The other thing I wanted to say is that very often welcome, at least in in one of the other ideas that we had for an Fs was that this would run on a, on a machine that's serving a different function as well. So for example, passing along run control information to the boards. So, it's in CMS terms, it's one of the function managers for part of the detector or something like that. So, these devices are anyway needed in the run. So, if the device by that would go down today, we'd have the same, same problem. So, I'm not sure we're introducing necessarily more single points of failure. So so again I am not particularly worried. But I do understand where you're coming from, and the net booties there mainly so that we limit the SD card aging

- S. Lueders, CERN-IT Security:

- Q (D. Miller): So, one of the issues that that I think we have faced is, um, you know the identical treatment of some of the system on chips, as a programmable logic controllers I mean the CERN rules, kind of, at least implicitly treated them on very equal footing, but I think there's some very unequal aspects of them could, could you just comment on on whether you know you think that's correct and what you see as some of the differences that we might as users of the system on chip devices. Think about specifically with respect to rules that are historically more targeted at, you know, simpler programmable logic controller type devices.

A: The rules are not really different between the PLC and the server if you just look at the basics Yeah, keep it up to date, control the input control what's running making sure that default passwords are removed and so on. So there's no no big difference. I can give you also the baseline for service for example. Yeah, and I'm not, I don't want you to go through all the points and, and it's more discussion document we can see what you can do. Its appointed to the document

Zynx Ulltrascale+ MP ?%

Other ?%

- Poll1b: How do you implement it?

commercial SoM ?5

custom SoM ?%

PCB ?%

- Q (R. Spiwoks): Can I just ask what those others are?

A (M. Dobson): So I answered other, this is, COMExpress.

A (S. Schlenker): Also, the HAP epsilon, the switch part ASIC has a MIPS CPU, running Linux.

- Q (W. Vandelli): I guess my question would be, you know, given the majority, the majority are using think my question would be why to the others have choose the other. I mean, what was the feature that made you go for something that seems to be at least in this community.

A (S. Schlenker): Maybe I can comment from my experience is also that certain engineers have certain experience with certain brands, so that also then you know defines their choice. If you get used to some tools you don't want to change.

- Q (R. Spiwoks): It's actually quite well distributed between the three possibilities directly on the printed circuit board or using a commercial system on module or a custom system on module.

Q (M. Dobson): Maybe those with the directly on the PCB could could give us some illustrations of why you've chosen that way more than, then going for, for, for a SoM that you place on to your PCB.

A (R. Spiwoks): I can answer for one of them for the MuCTPI because actually we started out very early, and we weren't really aware of system on modules, and some of them weren't probably available at that time. So, long time ago this was the only choice one had or least the only choice we saw. But we are changing, and we will move to a system on module.

A (S. Tang): Like a budget. We want to use a lot of transceivers, at high bandwidth for data transfers.

A (D. Miller): And just to just to add to that, I mean, part of it is, you know, and I've come and kind of made this comment in a few of my questions or follow ups on people's talks is that, you know, we're even despite what Stephen says about the risks, you know we're doing quite a number of different things on the, on the Zynx, it's included in the real time path, it's used for, you know, essentially crop distribution, but we're also using the ARM side for DCs and the OPC/UA server implementation, etc etc so it's basically just, you know, a couple of bullet points to follow on Shaochun's top level comment.

A (J. Tikalsky): We actually have two different ones for the purposes with our IPMC, where we had a very specific design goal in mind, we need fast response from the same Soc, so we can go with it. Like standard sort of just pure processor goal. For our main board, The reason we didn't go directly on the PCB in any way is because we have three boards with more on the way that all the

effectively the same thing. system for their control infrastructure and duplicating that circuit, and getting it, making sure that nothing went wrong on all the different cards, it was just far easier to have a module that we can build. However, many of them. Well, this doesn't apply to our Zynx, or modules, because we're not using processing power for for our FPGA's we actually decided against our workhorse FPGA switches decided against this and put them directly on the boards, because of better connectivity on the copper for the high rates versus having to go through an extra set of connectors. So those are some of the considerations involved.

A (L. Ardila): I can also answer from our side we use a custom system on module for safety purposes, for reliability actually there is nothing available on the market that which. The requirement, In terms of failure rates for example...It's mainly for Functional Safety. We, we needed actually a probability of failure, which was quite low. And nothing, according to the study that you have performed. five years ago, was fulfilling this requirement so we have moved with AVNET. And we have used actually, a commercially available one and you have modified it by upgrading the component.

- Poll 2: I would like to know what operating system you're using:

CentOS	50%
Peta- Linux,	21%
Yocto	14%
RealTime	9%
Bare Metal	5%

A (J. Tikalsky): I voted CentOS because that is where we are going to be interested in this part was, we are using real time for our Zynx IPMC.

- (R. Spiwox): So only for those that actually have a Zynx 7000, so maybe the others can just say that they're not using it. Do you still need 32-bit support.

Yes	18%
No	27%
Not Using	50%

A (J. Yikalsky): I actually ended up voting no on this because I'm not completely sure. We do have a mix of 7000's and MP SoC's there on those interchangeable mezzanines that I have mentioned. I don't know how long the 7000s will be in use, so the choice will be there. I expect it is relatively possible that we may phase them out by 2024, 2025 in favor of the things but just don't know.

A (H. Boukabache): I have voted yes because in our case we have invested so much money in this production right. and we would like to keep it as much as possible so it was very difficult to get any commitment from Xilinx and from AVNET concerning the long term availability of the of the Zynx actually 7000 so four years ago I got an official letter saying that they will keep it at least 10 years, which actually. can you see me? Let's say bring us to 2000 let's say maybe, I hope, 27. If we cannot produce it anymore, actually we will very probably upgrade also to the UltraScale. This is also the advantage of using a SoM, the factors also that you're working with AVNET actually I'm not doing any

advertisement prevalent. Just the fact that we are having actually kind of standardized, even if it's not the word for that, for the footprint we can upgrade it with another, another board that she was actually the same footprint.

A (E. Hazen): I can maybe address that. We expect no 7000 series Zynx's in phase two. So, but, if there is development work going on over the next few years, it would be very helpful to have 32 bit versions of generic software that we are required to use because it's expensive to go back and replace all those seven series, so we're in probably in the same boat in as many other people in that regard but to be sure in the end I think we use only as in culture skill plus in phase two installation.

- Poll3: SoC IP addresses?

DHCP w/MAC 63%

DHCP w/ClientID 37%

Static ?%

A (J. Tikalsky): We have Client identifier for the working ARM in our lab. and the latest version of our system is using it.

A (P. Zejdl): I was convinced that Mark was wrong with what he said in his talk but then I went up and read the hpm dot three, and the ATCA specification. Because when I did it I thought I was convinced I did it on logical start address. But then I realized the shop that they didn't at the time, there was a sense for mapping between physical slot and logical slot. And when I read the manual. Mark was right. It wasn't what I expected that.

- Q (M. Dobson): Maybe reason for using static and reason. Well, are you planning on moving to one of the other options?

A (R. Spiwoks): I can answer for the MuCTPI. It's static, but only for convenience. We will move forward to the DHCP.

A (H. Boukabache): Actually, in our case we were using static IP at the beginning and it's a users that's asked for the DHCP because it's easier to swap between instruments.

- Q (M. Dobson): Is MAC address sufficient for you or would you like to use some kind of client ID to actually uniquely identify a device in its location, rather than physical entity?

A (H. Boukabache): I think clientID would be very good idea. Now actually we're are using the MAC Address and a serial because it's the only way or the only method that are available, but if you have a ClientID...

- All, Discussion: questions previously submitted to Marc:

- Q: Mac Adresses - For MAC address assignment, user would prefer to be assigned a block (maybe a total of 4 per blade) of contiguous ones that can be programmed in non-volatile memory. How to go about this?

A: After discussion with relevant people, MAC addresses with a CERN prefix can be obtained from CERN through the service desk for any boards which are going to be used at CERN. It is however recommended that the board manufacturers ask for the clocks they require, and program them into the memory of those boards, rather than the users of the boards making separate requests.

- Q: IP Addresses - DHCP or a similar solution is fine (for CMS, we use a Host ID scheme based on the crate/slot). Will this Be available for ATLAS and SMC in the future?
A: Yes the idea is that ATLAS and CS would go to the same naming convention (still to be refined and agreed) and the same proposal for ClientID Usage in DHCP. But of course this as I said, is the aim there is phase two, not necessarily phase one, mainly because of manpower restrictions.
- Network Boot - network boot and a network disk solution seems best to us. We do have the option of an eMMC on the blade.
A: The proposal is indeed to do Network boot and use core, high availability servers for this, with the minimal amount of things on the local storage (be it SD card, QSPI flash, or eMMC).
- Q (M. Twomey): I know it would have been good for a poll, but how many people are looking to put QSPI flash, either directly on their boards, but having some sort of a flash connected to their processor?
A (J. Tiklasky): I could say that we do have QSPI flash on all of our system our modules for our IPMC's it's the only media we have on them. And for our embedded Linux mezzanine, the control processors. It is our recovery from a storage location.
A (H. Boukabache): Well actually I have presented it to during my presentation we we have a, we have the QSPI on the, on the SD card and you have also the eMMC that she's available but she was on the the QSPI on the, on the SD card. Actually we primarily boot on the SD card if something doesn't go well we go to the QSPI where there is a minimum image there.
A (N.N.): So very similar to some of the other schemes that have been proposed Yeah, thanks. Okay for Apollo, will we don't currently use the QSPI but I think when we have time we're going to make the QSPI another fallback mechanism, because we can switch the boot modes via the IPMC.
- Q: Where to get or store MAC addresses?
A (J. Tikalsky): I think that's probably just one of the interests along the same lines, is probably what the best arguments for getting one of these commercially EPROMS for they just guarantee that the EPROMS address is, you mean by default, it's right protected as well, you can't screw it up. So, from there all you have to do is make sure that your images, reading the MAC address on the EPROM.
A (R. Spiwoks): On the MUCTPI we actually have those EPROMS with the unique MAC address, but then we are reprogramming the programmable part with the CERN one.