

Hypervisor Project - minutes

★ This contains meeting minutes. If meetings are cancelled this would be noted here. Some dates are missing minutes but ideas are always captured on the main project pages [here](#).

Minute takers according to [this page](#)

July 12, 2021

Participants:

- Stephen (Renesas)
- Adam (Kernkonzept)
- Matti (Opensynergy)
- Gunnar (GENIVI)
- Oleksandr (EPAM)
- Kai (EB)

Minutes

- Gunnar: I met with Arm (Bernhard) and Francois Ozog (active in Linaro). We discussed virtual platform specification, if there are additional automotive-specific hardware that need VIRTIO support (e.g. do we have a proposal for FM tuner? (No, not currently))

Discussion on the usage of virtual platforms as a hardware-portability solution:

- Gunnar: We see this trend. I'd like to analyze this idea further, compared to for example just making stable user-space APIs and creating new kernel drivers for new hardware.
 - => Is this not just shifting portability work "somewhere else".
 - Let's analyze how "virtual platform API is stable" is different from "user-space API is stable" (i.e. just port drivers to new hardware..)
- Adam: BSPs that add a lot of patches to Linux get stuck on older versions. Not using mainline. Updating to new kernel drivers difficult (N.B. officially no APIs are stable inside Linux kernel)
- Matti: Linux kernel in particular, is an operating system kernel and a hardware-abstraction at the same time, which is challenging.
 - Future challenges: Value-add for SoCs could be difficult if the [hardware abstraction] API is limiting.
- The virtual platform / VIRTIO / Trout and others. It seems the "new" distinction is a clear focus on a hardware-abstraction layer?
 - (partly agreed, discussed the existence of firmware as one abstraction already used, and on PC side the hardware it self was standardized, etc. Overall, a complex reality)
- Summary:
 - Kernel updates are driven by fast development, security issues, etc.
 - Fast changes, no internal API stability guarantee etc. = updating drivers is a major issue.
 - It seems this situation can't be avoided, and the industry is looking for a solution to the challenge.
 - The theory seems to be that keeping API stable and efficiently port to new hardware, could be more feasible to do in the virtual platform.
- Discussion turned to evaluating the effort (that the industry is also considering) to create new kernel(s) instead, as a Linux alternative (for safety, and development concerns as shown above).
E.g. if they have a POSIX API, is that "enough". => No... could still be a large effort to port existing functionality.
 - Gunnar: As an example, consider systemd which early on used Linux-only APIs. And then a lot of functionality built on systemd. Such examples makes porting software a challenge.
- Also discussed how large part of Linux' functionality is actually used in practice in the industry?
- Matti: (Explicit example that shows the current situation)
 - [io_uring](#) recently added to Linux. It accelerates applications a lot. A hardware vendor that is able to use these latest features could have an advantage (hardware itself not faster, but the software runs faster).
 - => With a virtual platform abstraction, updating to a new kernel with the new features should be feasible.

June 2021

- Regular meetings held every week. No specific meeting minutes written. Results in AVPS and topic planning ([link below](#))
- AVPS v2.0 release.
- Working on [deep-dive topic planning](#)

December 14 2020 May 2021

- Regular meetings held every week and results tracked in [AVPS document](#) and [JIRA tickets](#).

December 7

- **Participants:** Stephen, Gunnar, Dmitry, Adam, Peter
- GPIO - clarified misunderstanding that referenced SCMI even though SCMI does not cover GPIO. Closed open issues and finalized requirements. Chapter will be *done*, after final polishing.
- Initial discussion on the structure of final GPU chapter (bringing together all input from previous hardware studies)

- We might need a plausible product design scenario to support the inclusion of VIRTIO-GPU / VirGL based requirements. Also consider how to handle a non-Linux OS, if that is included in the scenario.
Most other requirements will lean on hardware virtualization support, but the group did not want to lose the idea of support simpler hardware too.
- Dmitry has done some updates on Camera/Media – postpone discussion to next time.

November 30

FOCUS : GPU virtualization – Features of the PowerVR/Imagination GPUs (used in Renesas R-Car SoCs).

- **Presentation from Thomas Bruss, Renesas.**

- Discussing the opportunity for a common "model" of modern GPUs that can be the basis for shared requirements in the virtual platform standard.

November 16 - 23

Discussions captured in JIRA tickets and in the document.

November 9

Minutes pending

Participation planning:

- **October 19** - GPU focus? Discussion of new Mali features?
 - ✓ Bernhard + Arm colleagues, wants to present Mali GPU
 - ✓ Adam
 - ✗ Dmitry still vacation
 - ? Oleksandr - vacation?
 - ? Daniel Stone? Emailed last week. Send calendar invite also.
 - ? Eugen - *Emailed*
 - ★ Stephen likely interested
 - ✓ Ozgur - *Emailed*
 - ✓ Peter
- ~~October 26~~ No meeting (AMM week)
- **November 2**
 - ✓ Adam
 - ✓ Oleksandr
 - ? Dmitry
 - ? Peter ...
- **November 9** ✓ Adam ✓ Oleksandr ? Dmitry ? Peter ...
- **November 16** ✓ Adam ✓ Oleksandr ? Dmitry ? Peter ...
- **November 23** ✓ Adam ✓ Oleksandr ? Dmitry ? Peter ...

October 12

- ✗ Dmitry vacation, ✗ Adam has a conflict this week only
- ✓ Oleksandr
- ✓ Stephen join if possible (ad-hoc)
- ✓ Peter

October 5

Participants:

✓ Adam ✓ Stephen ✗ Dmitry busy ✗ Peter busy ✗ Oleksandr busy ⚠ Bernhard tried joining (missed by Gunnar)

August 31 - September 28

- Another 1-2 discussions on Graphics
- Continued discussion and work on the specification – results are in [working document](#) and in [JIRA](#).

Participants:

- Primarily Adam, Dmitry, Oleksander and Kai. Occasionally Bernhard. For graphics also Eugen. Peter had vacation.

August 24

Participants:

- Peter (BMW)
- Daniel Stone (Collabora)
- Adam (Kernkonzept)
- Gunnar (GENIVI)
- Dmitry (OpenSynergy)

Apologies:


- ~~Adam noted he has a conflicting meeting this week~~
- Oleksandr / EPAM colleagues – there is a public holiday

Minutes

- Producing the generic model of future GPUs. Adam started writing inside the AVPS as a place to hold the information for now. Daniel can view and contribute.
- Gunnar: That's fine for now. We should work on cleaning up the AVPS for release in the months to come.
- Bernhard: Next generation is fairly fixed but a wish list is useful.
- Peter: It is very valuable to know how Arm engineers think about this. We can discuss common parts of future designs
Gunnar: ... even if Silicon Vendors have their own unique details and/or use some other GPU cores.

Switching to discuss AVPS planning and open JIRA tickets.

Gunnar: A least the "Evaluate Status" column is getting fairly empty. This is good. There are some big ones there still, e.g. DSP/TSP.

- We discussed DSPs primarily. - see notes in JIRA ticket.  **HV-17 - AVPS:DSPs / TSP** CHAPTER DONE
- Gunnar: There is a lot to sort out here
 - 1. Technical organization (like defining the potential software/hardware stack, similar to the GPU discussion)
 - 2. ...and as far as the technology is still unclear - understand current project organization (how are projects solving this today, what conversations need to be had between product owner (e.g. OEM), Hypervisor implementor, and the silicon vendor?)
The idea like for any hardware chapter in AVPS, is *at minimum*, to try to have a few of those conversations proactively, so that there is a better basis to start from when going into a production project, and *as far as possible*, create some generic solutions/standards that remove obstacles.

August 17

Apologies:

- Dmitry

Participants:

- Peter, Adam, Oleksandr, Gunnar

Graphics discussion not fully successful today due to missing key participants, but we collected the thoughts of those who were here.

Graphics

Let's recap our thoughts and conclusions from last week's discussion

Adam: I have always thought you need hardware support to do this properly. But of course if you are already in another situation you might have to do something different.
But [the hardware-assisted solutions] involve NDAs and binary blobs (that are more difficult to analyze generally).

Peter: We discussed some similarities and a working model last week, the idea of the command queues that seems to be similar in future solutions. There are some special support for video and others. I would like to achieve this model that can progress this without being too specific on one solution. It may also avoid NDA issues. Find these common abstractions to work on. Multiple hardware should be able to be mapped into this.

Peter: We need to consider about composing images also - the display output (2D). When this is combined with other video feeds, there are safety concerns that need solutions.

Gunnar: There are some patterns for handling safety (e.g. fallback to simpler HMI if the advanced one is fails (recognizing failure in various ways)), and these patterns may drive our standardization approach.

Peter: These patterns are useful. Also considering that there are different needs - not all platforms require ASIL-B, for example.

Oleksandr: I'm not as focused on the graphics subsystem, so it's hard to say that much and what we are doing concretely is covered by NDAs. We have experience with hardware that has hardware-assisted virtualization support and we are implementing such things in Xen.

Some hardware has implemented more than one GPU and can use simple pass-through of each of these to for example 2 VMs.

Gunnar: Discussion last week with Daniel and others led to the conclusion we should start this "general model" discussion with the implementers of future graphics hardware. I feel like we need to produce a description that we can bring with us to start that conversation.

Adam: I would like to give it a try. (This week)

Gunnar: Great, and then Peter may want to add to this.

Power Management

Oleksandr: I have made some proposed edits. SCPI should be deprecated because it is now covered in SCMI and we should require the 2.0 version of SCMI.

IOMMU

Oleksandr shares some experience from Xen implementation.

Further discussion how to bring the chapter to a close, in particular the 2-stage IOMMU question that is still not brought to full conclusion.

August 10 – Graphics/GPU focus

Participants

- Daniel Stone, Ozgur Bulkan, Adam, Dmitry, Peter, Gunnar, Oleksandr, ... ? (I missed recording the full list)

Gunnar introduced the AVPS goals a bit and mentioned that in some hardware areas (and almost surely for GPU), the AVPS specification can/should not reach the point where it requires only one way of doing things. It can still be useful to dig deeper into a few things, and then describe a few different choices through optional or conditional requirements. It is still much more useful to have that analysis done when a production project starts than to have a blank page.

Daniel: Requiring Vulkan would make sense but will not be supported everywhere so unfortunately it cannot be the only choice. GL API is stable.

Peter: *[We need to discuss]* technology roadmap for next generation GPUs (HW support for virtualization)

Daniel:

ARM, Renesas, Qualcomm, others would need to discuss this to get the information, but everything might not be shared.

...HV vendors implement this when the technology is ready (NDAs)

Every VM has its own virtualized GPU. Dedicated command queue per VM. Buffer handling can be challenging (sec/safety, etc.) It gives you a "direct access" to the hardware. Normal drivers and software can be used, because it looks like a GPU.

Peter: Can a standard model for GPU be defined which can help development (of specification and code/solutions)? Command queue is common for all.

Daniel: Concentrating on the API definition which can be generic, and over time hardware support could make the implementation of it easier. This is the VIRTIO queue API level.

Ozgur: What about the original VIRTIO-GPU 3D specification

Daniel: It is similar to VirGL but was not quite the same.

Discussion would be needed with the real graphics engineers (within ARM, Imagination, Qualcomm....) to see if they are open to discuss standard APIs in one level.

Ozgur: Vulkan path is yet another choice?

Daniel: Conceptually similar. Host/HV would still need to manage resources, track all the Vulkan objects, buffers, textures, targets. This is the same for virtualization of Vulkan and OpenGL.

Ozgur: SoC providers do not want to rewrite their implementation (e.g. to support VIRTIO queues *[or other standard virtualization API]* Basically they provide everything from user space APIs to the hardware...

...

Daniel: Yes, I agree. Hardware independent and hardware dependent APIs will likely be different APIs. We can (only) expect a few basic things about HW dependent (using channels/queues and buffer handling).

Ozgur: VIRTIO is only concerned with multi client rendering?

Daniel: Yes, it does not handle multi leases and other more complex situations.

Ozgur: Can we start thinking about things like nested compositors?

Daniel: DRM/KMS provides some capability for HW compositors but not the constraints you need (to match the (virtual) hardware capabilities).

Ozgur: Will it need kernel/driver changes?

Daniel: It's primarily the API side for virtualization.

Daniel: Way forward. Split the problem.

1. Write how we think a fully hardware specific path will work. (a generic model for how we think HW assisted virtualization is done)

2. The fully generic VirGL / vulkan path. We can start encoding more specific requirements.

Ozgur: What can be implemented to test this now?

Daniel: VirGL already works on top of any OpenGL (GL ES) implementation. Mesa, and also others.

Way forward

Daniel: I'm happy to encode the requirements that we have already today in VirGL. Over time I can try to do the same for the Vulkan path. For the HW specific approach, can we reach out to the platform vendors? or HV implementing companies that are likely implementing some of these things now [*with the latest hardwares*].

Gunnar asks Dmitry to see if there is a way to have a (private) conversation with potentially ongoing commercial implementation projects, to draw the boundaries between what could drive the specification content (and what could not).

Peter: The model approach is very good because we would then be less dependent on the hardware vendor proprietary information while this is being built out.

Gunnar mentions that ongoing work should be complementary and can together improve the situation, from specification side, and from trying out implementations on AGL and other projects. The AVPS work provides the API specification, ensures multi-OS aspect is not lost, and aims to influence AUTOSAR and other industry initiatives to align on a common way forward [*based on the open development of AVPS*].

Daniel: It would be useful to get Panasonic's input since they drive a lot of implementation in AGL for virtualization.

Gunnar: Yes. Both the implementation and AVPS is VIRTIO focused which should help compatibility.

July 20

Participants

- Dmitry
- Adam
- Oleksandr
- Kai
- Gunnar

Introduction to Oleksandr (EPAM), working on VIRTIO support in Xen. He has lately worked specifically on implementation on IOMMU on Renesas platforms.

Discussed again how to rewrite Camera + IOMMU from scratch with all the new information we have gathered (It is currently disorganized, but has a lot of good knowledge).

Actions for Dmitry and Adam to do this rewrite.

Gunnar on vacation next week. Do you want Philippe to open the call?

Adam: I will work on the chapter, I don't need the phone call for that.

=> Next week will be skipped.

Absence planning

See July 6 and

+ Gunnar, away (this week) and next, and probably again end of August (1 week)

+ Kai - mid August 17 August 6 September

+ Oleksandr - possibly starting last week of August 2 weeks

Kai found an interesting paper from Strategy Analytics. Plan discussion about it later on.

July 13

Additional work on Camera and IOMMU with Dmitry and Adam - results in document.

AI: (Media and Cameras) Gunnar to work on combining original and new text to more logical flow.

AI: (Media and Cameras) Dmitry to also review red text in the same chapter

Info: Let's plan for Graphics oriented meeting some time in the future, inviting Daniel Stone.

Be aware of presentation that Daniel also posted to the [GPU related Wiki page](#).

Info: Gunnar likely to have vacation some time after next week's meeting.

July 6

In-depth review and improvement of Cameras chapter results are in the working doc.

EPAM colleagues missing, will follow up those topics next week.

Vacation planning

Peter - No vacation until Mid September

Dmitry - same for me, probably October at the earliest.

Adam - similar, nothing planned at the moment...

Gunnar - end of July / beginning August. Will specify better soon.

June 29

Apologies from **Artem** + colleagues (public holiday), Peter (unexpected mgmt escalation), **Adam**, etc.

Have no info from **Dmitry** or **Matti** about attendance today. Also not from **Kai**.

Other participants (Matt, Bernhard) are more attending on-demand / occasionally and were not expected today.

=> Meeting is skipped.

June 22

AVPS v2 work:

Continue discussing with Artem and Adam on IOMMU including programmability from guest, and requirements for coprocessors. We are getting there. Adam will try to rewrite once more.

Artem: I am sending over details that might be interesting, about the implementation for co-processors on Xen:

- [Co-processor_management_under_xen.pdf](#) (XenProject website)
- [XPDDS 17 keynote Shared Coprocessor Framework on ARM](#) (Slideshare.net)

Went through comments and minor changes in the whole document and accepted all minor changes. Open points kept for discussion with a few others.

Need time to discuss with Dmitry on changes in Media chapter and more.

Artem: I intend to involve 2 new colleagues into this work.

April 27 - June 15

Meetings held as usual, except for some bank holiday.

Minutes have not been recorded separately. Instead, discussions and updates on the [Automotive Virtual Platform Specification](#) have been recorded directly in the working document as text and comments. The [JIRA Tracker](#) is also used to track updates and assign tasks.

April 20

Participants

- Bernhard Rill
- Adam Lackorzynski
- Dmitry Sepp

- Matti Möll
- Gunnar
- Alex Agizim (EPAM)
- Kai Lampka

Apologies

- Peter (Zoom was blocked?)

Minutes

Went through tickets.

Discussed development support ticket DLT discussed at some length. Various "known" discussion points (ID uniqueness, on-disk storage format, common timestamp references) but they become more relevant in a Hypervisor setup.

Kai reviewed Booting chapter and found a few issues/unclear direction. Moved chapter to "Improve" column.

April 13 – no meeting (Easter Monday)

April 6

Minutes pending

March 30

Participants

- Bernhard Rill
- Peter
- Adam Lackorzynski
- Dmitry Sepp
- Matti Möll
- Gunnar

Apologies

- Artem (conflicting meeting)

Minutes

- Peter: Defining the platform better. The Arm specification on secure partitioning can also be useful to look at.
 - This also includes use-cases, to better define...
- Bernhard: Yes, and a few other specifications might be useful also. Others are not so applicable. I can look at this.
- Peter: Starting with VM scheduling, for example.
 - Matti: We can only define the interface... what other "surfaces" might be described?
 - Peter: *Just thinking that going from the top we realise VMs need to be scheduled.... Example - Android + RTOS, how is switching done and what does it lead to?*
 - Gunnar: This goes into a kind of system design. If you are speaking about cores etc.
- Bernhard: Two prominent examples of interfaces...
- Matti: Define what is a platform in general, and is it focused on components or interfaces?
 - A platform allows things to be built on top of it. It cannot be too hard to implement, because then it will not be implemented to the degree that it becomes really usable.
 - Think about the right level. Not too complicated. Not narrowing the solution space too much. And not too broad so that portability is not achieved.
- Peter: We are quite well defined here. Starting from defining a... hypervisor based platform
- Matti: If it is an automatic (preemptive) scheduler, it cannot be defined using interfaces.
- Gunnar: You can set priorities. Not from the VM, but when configuring the platform. (Therefore the AVPS could set requirements on it).
- Peter: We can still write requirements for this. The behavior can be defined.
- Peter: Variants of the platform advanced/basic could be an example.
- Matti: Currently covered in AVPS by "If the platform includes feature X, then it must..."

Short review (on overtime) of the [AVPS v2 ticket priorities](#).

March 24

Whitepaper work and AVPS prioritization

Minutes pending

March 16

Whitepaper work and AVPS prioritization

Minutes pending

March 9

Participants:

- Dmitry
- Kai
- Adam
- Peter
- Michael
- Gunnar
- Philippe

Minutes

Gunnar Presented [Webinar-Report draft](#). Purpose and content. Recommended reading for all, but especially those who did not attend.

AMM plans: **GENIVI All-Member Meeting** planned for **May 12-14 in Leipzig, Germany**

AI (all): Block 12-14 May (+ possible travel time) in your calendars.

- Philippe: Make sure to separate content preparation from any concerns about virus outbreak.
- The conference is planned to be F2F as of now. In any case we do not stop our preparation since content is reusable. Evaluation of the situation of course continues in parallel.
Preliminary attendance possibility:
 - Kai - Not travelling – family plans
 - Adam - **OK**. Leipzig is near.
 - Dmitry - **OK**.
 - Peter - **OK**, most likely, company will clarify attendance and likely to come, also bringing other colleagues.

Discussing the comments given in whitepaper editable document, and many general discussions about the direction and content.

- Possibly split into more than one document. Why Virtualization in general vs. deep dive into technical descriptions.
- or split one document into two sections
- Need more structure, more headings, more sub-chapters, more summary and leading the reader from intro to details.

AI (all): Get your (Google) account connected to the document, with edit capability **AND everyone expected to give **some** input until next time.**

March 2

Participants:

- Artem
- Peter
- Gunnar
- Philippe

Minutes

- Xen project status & current state from Artem
- Walkthrough AVPS 2.0 prioritization tickets
- Lots of discussions on plans / purpose / and various input to the Hypervisor Project overall – initiated by Peter, filled in by Artem.

February 17, 21, 28

Minutes pending

February 10

Participants:

- Kai
- Dmitry
- Adam
- Gunnar

Whitepaper:

- Main subject as agreed: "Why is Virtualization needed?"
- But should we also include challenges that are added as a result of choosing virtualization?
For example shared-memory as discussed in the comment by Michael on [Whitepaper - first private draft text](#)

Discussing shared memory challenges:

Locking implies timing interference. Lock-free implementations are complex but important. Can hardware support be provided to improve the implementation of shared memory implementation?

Kai: You would sacrifice isolation if different criticality functions access the same memory.

Adam: For video the approach be fine as long if data corruption is less of a concern. Depends on the use cases of course.

For sharing it is required to consider hardware support for making some access read-only to guarantee that non-critical software cannot change the memory values of critical software.

Gunnar: What about if there is only one writer (critical task) and multiple readers. Can non-critical readers negatively affect critical software by starving the memory bandwidth or similar?

Adam: Generally starving is not possible due to hardware setup. A worst case analysis would be needed.

Adam: Hardware counters usually exist that can enforce budgets on the memory usage. Usually two counters or so, the number of counters might not be enough.

Kai: How do the lock-free implementations perform in worst-case situations?

Dmitry, Kai: Locking (handling multiple writers or handling that writing is complete before reading starts) is orthogonal to the performance issue.

Gunnar: One detail, don't forget buffer sealing (as Linux calls it). It is a security feature (needed for some use cases) to set a buffer read-only as it is being handed over. This is to be able to guarantee that a writer does not modify the buffer again, *after* it has been handed over to the reader(s). Such late modification could be used by a malicious writer to exploit bugs in the reader implementation.

Conclusion: Yes, it might be useful to include, in the white paper, some challenges topics and how to solve them (e.g. shared memory).

Gunnar to send out links to webinar. Participants invite their colleagues.

Weekly meeting, February 3, 2020

Minutes

- We went through the slides for upcoming webinar. Feedback on content.

Weekly meeting, February 27, 2020

No minutes

Weekly meeting, January 20, 2019

Restart after new-year, CES, etc.

Spec has been released. Dissemination/information given to various automotive industry individuals during CES.

Matti is away from project for a while.

New JIRA items, one per "chapter" to track which parts need update.

Some additional edits will be done by OpenSynergy tech-writer (Susan)

Action to all: Each participant finds at least one colleague/friend to show specification to for feedback.

Webinar is being planned beginning of February.

Artem: For Xen we are quite busy atm, working on safety certification, some ELISA project participation, and of course virtio implementation...

No progress on whitepaper - need leadership to assign tasks/chapters in practice.

Weekly meeting, December 9, 2019

Rough minutes: Went through spec from the last page in reverse. Discussed open points.

Actions taken by Adam to update

Weekly meeting, week of December 2, 2019

Participants

- Gunnar Andersson
- Eugen Friedrich, ADIT
- Michael Doering, Bosch
- Artem Mygaiev (until 10:45)
- Kai Lampka
- Adam Lackorzynski
- Matti Moell
- Philippe Robin (part)

Apologies/Absent

- Michele
- Dmitry

Agenda

- Memory/buffer sharing standards for graphics applications (with Eugen Friedrich)
- AVPS completion work

Discussing graphics/memory buffer sharing.

libvirgilenderer implements the API, including reference implementation of VULKAN API

In practice there will be only one implementation. It might make sense to require using libvirgilrenderer quite simply.

The driver is part of MESA project.

Non-virtual operating systems? QNX apparently ported both driver & renderer part.

Integrity? MESA driver is a user space library - creates a command stream - should be easily portable (paravirtualization assumed).

Eugen: There was a proprietary API for memory sharing proposed in...

...Other APIs just give you a handle, and you don't really control what happens below.

...Should there be a defined way how memory can be represented in a generic way?

Matti: VIRTIO basically is this in fact. The scatter-gather lists are providing this. Managing the lifecycle of the buffers is the challenge. Gerd Hoffman / RedHat has proposed a standard - see virtio list. (<https://lists.oasis-open.org/archives/virtio-dev/201911/msg00149.html>)

Need to build in handling of the particular characteristics of the hardware. E.g. special alignment or size restrictions. Usually VIRTIO has per-device handling today (GPU, block-device, ...) because of this.

Intel created a proposal to require host to allocate and give to the guest. Memory accounting troubles follow. This was not accepted in the community.

recent proposal VIRTIO-GPU resource attachment. patches to qemu to get a DMA buffer from a guest target buffer. udma (in qemu). This could be applied to embedded HVs.

create_memfd - filedescriptor. udma driver - ioctls to control it. See DMAbuf in Linux source. drivers/dmabuf/udmabuf. Note memfds are not part of POSIX standard.

Gunnar: These are Linux specific APIs then. Any different consideration for non-Linux OS (in a virtualization environment).

Matti: Should even be easier to implement in simpler OSes (since user space code may be more privileged to access the details).

Matti: One answer is that the GPU 2D specification already allows sharing buffers. You have to keep into account some details about the memory model of the HV.

VMs can tell the HV environment about buffers they would like to show, including giving this buffer to another VM.

Some kind of global compositor and it is given the buffers, e.g. that compositor might be within one VM, but it could be a hardware display device...

How to handle the lifecycle of the buffers?

Matti: VIRTIO specifies a low level way to communicate between VMs with virtqueues. The other is the description of the device implementation. How to handle the scatter gather lists of buffers.

Gunnar: The discussion involves different standards, in addition to standard Linux/Wayland, there is for example the Android graphics stack, potentially others...

Wayland can display a DMA-buf so if only you can get a handle to that, it's one way.

AVPS completion work

- Went through more comments. Some parts have been proposed as resolved after rewriting the section (Gunnar)
- Matti Technical writer input, can the document be freely changed?
- Agreed to welcome technical writer changes, but with some track-changes capability so we can see if any part was changed in a way that makes it misleading.
- Gunnar: ...it sometimes happens when unclear explanations are clarified that they end up to say the wrong thing instead.

Weekly meeting, week of November 18, 2019

- Stephen Lawrence, Renesas
- Gaurav Sinha (Micron)
- Michael Doering, Bosch
- Adam Lackorzynski
- Dmitry Sepp
- Philippe Robin (part)
- Gunnar Andersson

Minutes

Intro and discussion with Gaurav. AVPS work - documented in the google doc itself.

Weekly meeting, week of November 18, 2019

Participants

- Artem Mygaiev (part)
- Gunnar Andersson
- Michael Doering, Bosch
- Kai Lampka

Apologies/Absent

- Michele
- OpenSynergy x 2
- Adam

Minutes

- Introduction Michael Doering, Bosch
 - Proficiency in networking especially TSN, and other things
- Introduction of work, purpose, results for new participant Michael
- Haven't heard from OpenSynergy this week. GA will follow up.
- Artem: Note recent feature list for Android R – VIRTIO mentioned. Follow-up with Google.
- Gunnar: Yes, following up on HV standardization is a good idea when we discuss Android-SIG activities with Google
- Showing comments in draft [document](#) from Michele
- Power management text under Hardware considerations sparked discussion on if we have concluded the situation for Power Management feature requirements. Actual chapter is not written?
- Michael will read through whole document and get familiar with the work in general
 - Possibility to pick any chapter and improve it.
 - Gunnar: We appreciate input on TSN but it is not critical for first draft so focus might be elsewhere first
- Kai has limited time next week/weeks due to teaching obligations, as well as planned personal training

Weekly meeting, week of November 11, 2019

Cancelled due to Tech Summit. Not everyone was aware though.

Weekly meeting, week of November 4, 2019

Participants

- Artem Mygaiev
- Michele Paolino
- Philippe Robin
- Dmitry Sepp
- Gunnar Andersson

Apologies

- Matti Möll (office guests)
- Kai Lampka (other obligations)
- Adam Lackorzynski (other obligations)

Minutes:

- Working on draft spec. Most discussion was on IOMMU with requested input from Artem
- Michele to look over web pages, project introduction, and Google Doc (editing, commenting...). Consider hardware specifics and ensuring t
- Artem to look over draft spec, perhaps mostly IOMMU. Focus on narrowing spec into a draft release.

Weekly meeting, week of October 27, 2019

Weekly meeting, week of October 20, 2019

Weekly meeting, week of October 14, 2019

New meeting time used today, Monday 10:00 CEST

Participants

- Adam Lackorzynski
- Artem Mygaiev
- Matti Möll
- Jens Uwe Schaefer
- Kai Lampka
- Michele Paolino
- Gunnar Andersson

Minutes

- Introduction and welcome to new participants
- Review current work topics and purpose with new participants
- We went through the status of the the Virtual Platform definition work, updating [the table](#) and some few changes in [spec text](#)

Weekly meeting, week of October 9, 2019

Meeting missed/skipped due to ongoing rescheduling

Weekly meeting, October 1, 2019

Participants

- Artem
- Matti
- Dmitry
- Gunnar
- Kai

Apologies

- Adam

Minutes

Tech Summit, November 12-13

- Matti: Someone from OpenSynergy based in USA (but not Detroit) might be interested to join, but not familiar with the GENIVI work so far.
- Artem: Checking my november planning - I might have a trip to USA
- Kai: I could check with EB subsidiaries in USA. There are people with HV expertise.

New meeting time

- Kai free all Tuesday but not 10.00 around 11.
- Matti busy Tuesday 11-11.30
- Adam wants to change day. (Kai says lunch around 1200)
- Monday 10.00 is an option
- Afternoons: Matti more likely to be busy (also at short notice). 3PM maybe...(any day)
- Kai: around 3PM should be OK. Most options sound OK.
- Most options seem to work for Artem.
- (AI) Gunnar: Set up a doodle poll with these constraints.

Virtual Platform Status

- Short discussion: What is "networking"? Does vsock apply in that concept, or elsewhere?
- Artem: Look at [ULSNet](#) for another view on what fits into OSI model. (joke) 😊
- We went through the entire [table of status](#) for information and for updating, looking into the [draft specification chapters](#) to confirm the current content.
- Last part of table needs to continue next week.

Other

- Artem: I would like to connect the safety-critical Linux community with the work we are doing here. There are a lot of discussions of when hypervisors are needed and when safety-critical Linux can suffice.
- Gunnar: Sounds good, let's make the connections. The whitepaper subject ("Why Hypervisors are needed") seems related.

Adjourned (late) at ~11.35

F2F meeting, September 24-25 2019

⚠ Some rough notes from F2F were captured here. Cleanup in progress

F2F September 24 - notes to be copied to the right place...

<https://etherpad.net/p/hvws>

Motivation: Why to use HV:

WHAT IS THE PURPOSE OF THE WHITE PAPER?

1) *...We need to explain why virtualization is actually needed. (It is still not fully accepted as necessary by all)*

This is the agreed main topic.

2) Explaining concepts -> 1 chapter pro/cons?

Microkernel, Monolithic <- multiple privilege level of modern CPUs, Type 1, Type 2, Paravirt, HW emulation, Linux Container
try not to rehash too much of existing data on this, but make introduction to the reader who does not know it and needs the basics to understand the rest of the paper.

3) *How do improve the usage of virtualization to meet the goals previously stated.*

E.g. chapter on needed hardware support --> At least 1 chapter
standardization of interfaces --> At least 1 paragraph
training, education, explanation... --> This is basically just a sentence We need more training.
Sample implementations --> ... also just a sentence.

HERE ARE THE REASONS:

*Certain concrete security/safety issues that can be shown clearly and that HV can counteract
System flexibility is another very important point.*

Consolidation of functions into a single SoC (multiple ECUs on a single platform), this includes legacy systems and new SWC
This calls for a clear interface between SW stacks of different vendors and shall allow provider of SWC to make their choices on the execution env./OS, examples to this are different combinations of OS/adaptive Autosar implementations or combinations of classic Autosar modules, e.g., SafetyOS / Com Module / RTE provider.

To this and, certification, let it be for security or safety is lost, as soon as components are re-arranged into new setups.

This clearly calls for flexibility in the sense allowing SWC provider to integrate their complete SW stack into the platform and still guarantee fundamental isolation properties.

Technical:

- * Single-kernel systems provide isolation but there are limits.

- > Although capability of configuring for example the Linux kernel for isolation is ever improving, there are counter examples (see other chapter).

- In addition, a single Linux kernel solves only some scenarios since putting all SWCs on top of a single Linux kernel is still limited by the requested OS diversity (see next chapter).

Other realities:

- * More freedom in choosing operating system kernel. (The automotive industry has a reality of using multiple OS choices)

- * Multiple software component parts coming from different vendors combined on common hardware. (The automotive industry has a reality of business setups with multiple vendors, each wishing for a defined and limited responsibility)

- > Virtualization enables the use of common hardware but have freedom from interference

VMs provide a small, clear interface:

The virtual platform is a relatively speaking small and simple interface to separate different software deliveries from different vendors.

- > Compare scenario where different vendors deliver programs to be on a single kernel -> containers -> messy chapter, try to avoid mess...

- > Other options include single kernel + isolation (container approach)

- or advanced runtime/platform definition for software component interoperability (Classic AUTOSAR)

- *Standardization of this interface will improve this further (reference Virtual Platform definition)

- write a bit about how this improves things*

Containers vs VMs

Are there concrete examples of isolation (resource sharing/interference) scenarios not handled by containers.

Avoid the either-or discussion since the future may be in container + VMs combinations.

Kai: Performance counters required to manage resource interference <-- details needed

HVs that can do resource management on a more detailed level than any kernel (currently) can.

As an example: A hypervisor implements limits on the rate of cache misses allowed for a particular VM. This might be implemented to protect another critical part from being cache starved (e.g. a real time core which shares memory cache). Excessive cache misses would cause the offending VM to be throttled.

Are such requirements possible or likely to be accepted (into Linux mainline kernel)?

Contention on implicitly shared resources (memory buses, caches, interconnects...)

Opinions on the high level purpose of the paper.

Certain concrete security/safety issues that can be shown clearly and that HV can

System flexibility is another very important point.

Fundamental system components can be replaceable and pluggable.

E.g. it is possible to insert (delivered by a separate vendor) a virtual network switch with advanced features such as traffic shaping, scheduling and optimization among VMs sharing a network. While it would be technically possible to add this to the Linux kernel, it is less likely to be accepted as a separately delivered independent software part. (Of course this is indirectly a consequence of the purely technical aspect that Linux is a monolithic kernel, since a microkernel + user-space implementation of system features would yield a similar effect). Along the same lines, there are license requirements that are in effect. Adding code to the Linux kernel requires them to be GPLv2 licensed, whereas independent parts (as VMs or in user-space) do not. It is easier to assign responsibility to the vendor of this component if it is isolated from the rest of the kernel.

Interaction between general-purpose and dedicated cores is poorly understood.

Can we explain and give examples or is it enough to state this?

3.2 Network Device

Standard networks

Standard networks include those that are not automotive specific, but instead frequently used in the computing world. In other words these are typically IP based networks, but some of them simulate this level through other means (e.g. vsock which does not use IP addressing). The physical layer is normally some variation of the Ethernet/WiFi standard(s) (according to standards 802.*) or other transport that transparently exposes a similar network socket interface

virtio-net = Layer 2 (Ethernet / MAC addresses)

virtio-vsock = Layer 4. Has its own socket type. Optimized by stripping away the IP stack. Possibility to address VMs without using IP addresses. Primary function is Host (HV) to VM communication.

Discussion:

Vsock: Each VM has logical ID but the VM normally does not know about it. Example usage: Running a particular agent in the VM that does something on behalf of the HV. There is also the possibility to use this for VM-to-VM communication, but since this is a special socket type it would involve writing code that is custom for the virtualization case, as opposed to native.

vsock is the application API. Multiple different named transport variations exist in different hypervisors which means the driver implementation differs depending on chosen hypervisor. Virtio-vsock however locks this down to one chosen method.

Requirements:

- If the platform implements virtual networking, it shall also use the virtio-net required interface between drivers and Hypervisor.
- If the platform implements vsock, it shall also use the virtio-vsock required API between drivers and Hypervisor.
- Virtual network interfaces shall be exposed as the operating system's standard network interface concept, i.e. they should show up as a normal network device.
- The hypervisor/equivalent shall provide the ability to dedicate and expose any hardware network interface to one virtual machine.
- The hypervisor/equivalent *shall(?)* be able to configure virtual inter-vm networking interfaces.
- Implementations of virtio-net shall support the

Discussion:

Virtual network interfaces ought to be exposed to user space code in the guest OS as standard network interfaces. This minimizes custom code appearing because of the usage of virtualization is minimized.

MTU may differ on the actual network being used. There is a feature flag that a network device can state its maximum (advised) MTU and the guest application code might make use of this to avoid segmented messages.

Parked questions

- What could a common test suite look like?
- Google virtual ethernet driver - merged into mainline?

1. Introduction

Automotive requirements lead to particular choices and needs from the underlying software stack. Existing standards for device drivers in virtualization need to be augmented because they are often not focused on automotive or even embedded, systems. Much of the progression comes from the IT/server consolidation and in the Linux world, some come from virtualization of workstation/desktop systems.

A collection of virtual device driver APIs constitute the defined interface between virtual machines and the virtualization layer, i.e. the hypervisor or virtualization "host system". Together they make up a definition of a virtual platform.

This has a number of advantages:

- Device drivers (for paravirtualization) for the kernel (Linux in particular), don't need to be maintained uniquely for different hypervisors
- Simplify moving hypervisor guests between different hypervisor environments
- Support guest and host interoperability:
- Programming against a specification simplifies developing local/custom variants of systems (e.g. Infotainment systems tailored for a certain geographical market)
- Simplify moving legacy systems to a virtual environment
- Some potential for shared implementation across guest operating systems
- Some potential for shared implementation across hypervisors with different license models
- Industry shared requirements and potential for shared test-suites
- A common vocabulary and understanding to reduce complexity of virtualization.
- Similarly, guests VMs can be engineered to match the specification.

In comparison, the OCI initiative for containers serves a similar purpose. There are many compatible container runtimes and the opportunity enabled by virtual platform definition is to have standardized "hypervisor runtime environments" that allow a standards compliant virtual (guest) machine to run with less integration efforts.

The specification shall enable all of the above while still enabling the ability for different implementations to differentiate, add additional features, optimize the implementation and put focus on different topics.

2. Architecture

Parked questions

- What could a common test suite look like?
- Google virtual ethernet driver - merged into mainline?

September 17, 2019

Participants:

- Adam
- Kai
- Gunnar

Minutes

Short meeting because of few attendees. We expect some are preparing for next week F2F ;-)

Sync up on [whitepaper work and review](#) with Kai after absence.

Discussing F2F logistics and whitepaper planning.

September 10, 2019

Participants:

- Adam
- Bernhard
- Gunnar
- Philippe
- Matti
- Artem

Apologies:

- Kai (vacation / out of office)

Minutes:

Matti: I think the draft covers a good outline fairly well

Gunnar: Conclusion chapter

point out what is out of scope

Gunnar: Reasonable point not to include everything

Bernhard: Several ARM features supporting this, we have provided

Currently, includes when to use paravirt, hardware support, isolation generally.

Artem: Useful to include some example of what different HW vendors include (e.g. for inter-core communication as discussed last week)

Artem: Power management is also important

Bernhard: Also S-MMU is such a related feature.

Matti: Travelling next week, 1AM time slot. Will try to join.

Discussion on F2F agenda.

Action (Gunnar): It is time to send out invitation/information about the F2F to mailing list.

More discussion on whitepaper content and plans for who will write which chapters.

- captured in [Review of content/outline for the early draft of HV whitepaper](#)

September 03, 2019

Participants:

- Dmitry
- Adam
- Bernhard
- Gunnar
- Philippe
- Matti
- Artem

Apologies:

- Kai (vacation)
- Bruno (conflict)

Minutes

- Discussed the [Whitepaper draft content](#) and the result was added **using green text**.
- Outline/content of the inter-core communication was added with Matti's input
- The document purpose was further discussed and ideas were **noted in the beginning of the document**

August 27, 2019

Participants:

- Dmitry
- Adam
- Bernhard
- Gunnar
- Philippe

Apologies:

- Kai (vacation)
- Matti (busy)

Minutes

Discussing the "outline" (general content and order of content) of the first draft in the page [Review of content/outline for the early draft of HV whitepaper](#)
Most notes taken on that page directly...

Action: Dmitry to try to find time to draft some I/O challenges leading to need for IOMMU and eventually need of Virtual IOMMU... (as needed by whitepaper)
... *but this week is very busy...*

Action: All to look over/adjust F2F agenda

Bernhard: Still considering what we can contribute to chapter 2 in terms of new coming hardware chapters.

August 20, 2019

Minutes pending

August 13, 2019

Apologies

- Matti
- Dmitry

Participants:

- Bernhard
- Gunnar
- Philippe
- Kai
- Bruno

Minutes

- Bernhard/Gunnar: ARM is still waiting for input what ARM can bring to the F2F.
- Bernhard: On the whitepaper side, Chapter 2 seems like we can help. But it is very hardware dependent. There should be some info how hardware features evolve. In this ARM can contribute.
- Gunnar/Bernhard: Open question for team/writers to answer: HW needs could be reflected differently in the architecture - more generalized, or specific. Likely there might be a few specific examples from ARM. Sometimes examples are even more specific than what is described in generic-ARM (in other words a specific SoC might have added a unique feature).
- Bernhard: Example SCMI presented before, we can extend on that and/or TrustZone
- Gunnar: Start Attendance list. I will add it on F2F agenda.
- Bruno: I cannot join the F2F in September - travelling for 3 months.
- Gunnar: Please consider if another engineer near you can support networking chapter finalization in [Virtual Platform specification](#)
- Al:(all) Review [F2F agenda](#) and give feedback on the content and allotted time for each topic.

August 6, 2019

Participants

- Dmitry
- Bernhard
- Kai
- Adam
- Philippe

Apologies

- Matti
- Gunnar

Agenda:

- whitepaper projectization
- Discuss F2F agenda

Minutes

- [Whitepaper](#) projectization in Jira
 - Philippe: points to the proposed document outline available at the bottom of [\[HV\] Next Generation Multi-OS System Design Whitepaper](#)
 - sections are assigned as follows
 - section 1: introduction - Kai has already produced a 2-page intro and will add it to the wiki
 - section 2: What does the HW need to fulfill to support unmodified execution of complete SW stacks,
 - what are the problems ? what is in the pipe ? might reuse elements of Adam's talk at the AMM, assigned to Bernhard
 - section 3: isolation & partitioning : Bernhard can provide high levels definition a,d an "academic view", which will have to cross-checked with reality later, Adam & Kai can provide a text on isolation & timing
 - section 4: Inter-core communication - assigned to Dmitry and Matti
 - section 5: summary - assigned to Kai then Gunnar
 - upcoming F2F is scheduled on 23-24 September, it would be good to have a version 1 of the paper at the beginning of September in 4-5 week time
 - Philippe: will create Jira tickets to track and initiate to those tickets with a 4-5 week sprint timeline
- Virtual platform: projectization in Jira - assignment of sections to contributors, doc ownership
 - Matt has populated Jira with a set of tickets correspond to the sections of the specs
 - discussion on the assignement of those tickets shifted to next week's call
- OEM feedback on HV work
 - Philippe shared Renault software expert feedback on HV work in GENIVI
 - background information: slidedeck on HV work prepared by Gunnar
 - Multi-OS System design whitepaper: "I guess everybody has already produced similar reports internally, but the landscape moves fast, and generating something @Genivi is certainly a good idea. I remain convinced that there is no serious use case

for hypervisors proper until AD (Autonomous Driving), but study of alternative paths for multicore Socs, especially asymmetric ones, is certainly useful."

- Automotive virtual platform standardization "This one looks much more ambitious, expensive to achieve and also very useful if it works, assuming convincing use cases show up"
- the interesting point is that we should be able to ask Renault to review the docs in the future
- F2F agenda preparation
 - Kai: IMHO 60-70% time should be assigned to virtio => to be checked with Matti (Philippe will send a note to Matti)
 - 1 hour each in each section
 - how ambitious can the project be ?
 - 30-40 % time should be devoted to the white paper, go chapter by chapter
 - Philippe: will create a draft agenda in the wiki
 - Philippe: logistics details need to be added to the project front page
 - /TODO/ Dmitry add logistics details to the wiki (meeting location, direction to opensynergy, hotel recommendation and deal if any)
- Next week agenda
 - review and assignment of backlog and sprint tickets
 - F2F agenda finalization

July 30, 2019

Apologies

Participants

- Matti
- Bernhard
- Bruno
- Dmitry
- Gunnar
- Philippe
- Kai (last 5 minutes)

Minutes

- JIRA project work
- Matti convert the [virtual platform](#) "table" to tickets
- Kai still interested to write but ran out of time. New attempt for next week.
- Break down the work by writing some JIRA tickets (Kai).
- Next week: also Discuss F2F agenda

July 23, 2019

Apologies

- Matti (vacation)

Participants

- Dmitry
- Bernhard
- Vasco
- Philippe
- Gunnar

Minutes

- Link to whitepaper planning (scroll down to see proposed document outline) [\[HV\] Next Generation Multi-OS System Design Whitepaper](#)
- Vasco offered to start writing a chapter, or much more detailed outline of the chapter. I.e. what's the actual content.
- Proposal to switch meeting time came up again

July 16th, 2019

Participants

- Dmitry
- Adam
- Gunnar
- (Kai Lampka - see end of minutes)

Apologies

- (Kai Lampka)

- Bernhard

Minutes

Video virtualization standards update (Dmitry):

V4L2 is stable but Android is the main challenge because of a lot of changes and new development. Changing from OMX (older API) to Codec2 - new HAL /API, operating on top of OMX (backward compatible) and V4L2 and other.
Some video accelerator technology coming in from ChromeOS / Chromium.

Dmitry: We haven't approached VIRTIO with any proposals yet. Need a bit more stability first.

Gunnar: In your opinion how is Android regarded in the upstream VIRTIO work? Or are most of the participants primarily focused on Linux / QEMU ?
... seems primarily Linux / QEMU is driving still.

Dmitry/Gunnar: CrosVM is a hypervisor monitor in ChromeOS. It uses KVM as the control interface to the kernel.

Discussion about project organization. Should we get our ongoing work up on JIRA tickets for more clarity?

Gunnar: We started this a bit more ad-hoc and driving forward with [tables](#) and minutes. [...but we have not so good delivery accuracy...]

Most other GENIVI collaboration projects use a lightweight SCRUM with some defined sprint content, depending on people availability.

JIRA is successful in several other projects ([Android SIG](#), [Connected Services](#), [GPRO/Franca-ARA:COM](#) and others) (**NOTE**: some links might require you to log into JIRA)

Discussing F2F agenda

Gunnar: Originally, we said to focus on the virtual platform specification.

Dmitry: Agreed, primarily specification. But what do the others think?

Gunnar: Note that Participation from Kai will be not every week - Adam is supporting on similar topics.

Adam: My main area of knowledge around VIRTIO is also block device and storage, like discussed with Kai before.

Gunnar: Please consider if you can help out to write the Block device chapter into the specification.

AI (Gunnar): Contact Artem again.

Meeting adjourned 10:40.

...then Kai joined. 😊

[We had a quick sync up between Kai & Gunnar]

Kai: I have a conflicting meeting every week but sometimes I can join a bit late like today.

Gunnar: OK, let's try to work with that somehow.

... mentioned JIRA idea as above

Kai: ... OK, maybe useful for whitepaper delivery.

AI (Gunnar): Look at setting up JIRA project for HWWS project

F2F agenda

Gunnar: What do you think should be the main agenda for the F2F? Virtual Platform Specification?

Kai: Yes, primary focus on specification, some whitepaper sync up (assuming draft is done before September)

Kai: I think everyone can write a chapter in the whitepaper, we could get a draft ready by September. Let's agree on the

Gunnar: Currently the primary outline is based on your combined proposal. I only left the other chapters at the end as a kind of history and possible to select some ideas from.

It also includes our brainstorm at the top, so the page is a bit messy now. Let's move the agreed outline to a separate page/document for clarity.

Kai: I will work on this.

Kai: Next week I expect to be able to join also, perhaps a little earlier.

Future absences:

Kai: Last week August, first of September

July 9, 2019

Participants

- Bernhard
- Dmitry
- Franz
- Gunnar

Apologies

- Kai Lampka
- Matti (vacation)
- Philippe

Minutes

- Recap of previous minutes = Informing Bernhard and Franz about previous meetings.
- Informing about last week's discussion on Whitepaper scope and outline. Useful, but not perfect since we could not get Kai to repeat his thoughts directly instead only discussing what the rest of us remembered. The participants looked forward to hear from Kai more directly at a later time.
- Dmitry: I was looking forward to continue/complete that discussion...
- Further discussion about whitepaper outline and content.
- Gunnar: The agenda for F2F needs to be set. As far as I know we decided a F2F was needed primarily to finalize the virtual platform specification.
- Bernhard: Interested in the hardware requirements part of course (from whitepaper & virtual platform discussion)
- Gunnar: When can we get an update on audio interface standardization? (from OpenSynergy)
- Dmitry: Final approval is stuck on a few details. Should be ready soon.
- Planning absences and summer-time meeting schedule.
- Gunnar: We might need to check availability and see if some slot should be cancelled.

Known absences as of today:

Franz - last week of August and also first week.

Bernhard - much later (November)

Dmitry - no plans so far

Gunnar - TBC. Likely last week July and beginning/mid of August

July 2, 2019

Apologies

- Bernhard (conference)

F2F workshop decided : Sept 24-25.

Kai presenting starting point for Outline

- Some discussions about MCAL. Vendors deliver MCAL with a quality statement and applicability for specific safety requirements.
- On Linux the drivers are not given with such quality statements.
- Costs of qualifying final systems not always considered.

Matti: Some parts of the system such as clock control need to be isolated from... E.g. clock controller for Ethernet network needs to be under the control of an equally safe part of the system. Use a safety island or a VM responsible for this. Some tension between hardware vendors providing such features and the proponents of hypervisors.

Adam: You can also mix this stuff. Lay it out as you need it.

Gunnar: This is what I mean about design guidance. Present choices, present consequences of choosing, and then

Lots of discussion on scope and possibility to include the Design Guidance (mostly between Gunnar and Kai)

Kai wrote some additional points down during discussion and will send them over. Most likely these will be integrated into the whitepaper guideline.

June 25th, 2019

Participants

- Philippe Robin
- Dmitry Morozov
- Gunnar Andersson
- Deventra T
- Kai

Apologies

- Matti (vacation)
- Vasco (travel)
- Bernhard (conference, also next week)

Minutes

Dmitry: We follow VIRTIO block device standard.

Adam: Other than the trim/discard stuff we have noted is missing, we have no issue with VIRTIO block device standard. It is fairly small after all. There's a patch for Linux, it should be merged now. Eventually it should show up in VIRTIO

Adam: The trim/discard has been added to VIRTIO 1.1

Gunnar: Let's complete the spec - write a few requirements into the block device chapter.

Kai: We can do that

Dmitry: We don't really use it but VIRTIO should cover this. It is quite mature.

Gunnar: Let's review the crypto support chapter in VIRTIO

Dmitry: 2D is fine.

Dmitry: 3D is still changing. Android will require Vulkan. New versions should be based on Vulkan. Someone needs to introduce Vulkan support in EGL renderer or everything move to Vulkan.

Gunnar: Not moving fast enough to get to a stable point yet, then?

Gunnar: Vulkan support on the driver side?

Dmitry: Android Emulator should need it. Google might be working on it?

Vulkan support exists on bare metal hardware (GPU vendors provide it) but not yet for virtualization. This is a kind of showstopper for Android in virtualization in the future.

Gunnar: Is there a minimal set of requirements to write down today?

Dmitry: 3D part is still a big question. It's hard to decide on the requirement set.

Input

Dmitry: We have some implementation of this spec.

Adam: vsock can be used for VM to VM communication.

Adam: The user level APIs are normally standard socket APIs so that is convenient

Gunnar: But can you assume all features work? Let's say I select/poll on the vsock, and other file descriptors, will it work:

Adam: Impl by kernel...

...Alternative: VIRTIO console/character device. But that's a different interface.

Dmitry: One of the stty needed to change to RAW data transfer and then it's available.

Gunnar:

9pfs

IOMMU -

Dmitry: Lots of updates still, upstreaming. When that is done I will look at the final specification and update chapter, should be about 1 month or so.

Gunnar: Can you provide links to upstream / blog etc.?

Dmitry: Already in Wiki see [IOMMU Summary](#)

Dmitry: Matti will be in office next week, then away for 3 weeks. Planning to upstream more patches.

Adam:

Sensors

Dmitry: From mailing list: OASIS don't want to accept any sensors. In the end it's just a byte transfer.

June 18th, 2019

Participants

- Philippe Robin
- Adam Lackorzynski
- Dmitry Morozov

- Matti Möll
- Stephen Lawrence
- Vasco Fachin
- Bernhard Rill
- Gunnar Andersson

Apologies

- Kai Lampka

Minutes

Whitepaper scope, followup

- Whitepaper focus: Explain what can be done with current hardware, vs. wish list for the future
- Adam will discuss with Kai when he is back

F2F

- 20-22 Sept. is All Systems Go Conference in Berlin, Fri-Sun
- HV Conference – see Doodle.

Doodle for F2F:

https://dudle.inf.tu-dresden.de/Genivi_HVWS_F2F_Workshop_September_2019/

Number of days? 2 or 3, let's create another poll

<https://dudle.inf.tu-dresden.de/v5fxnz31/>

AI(all): Fill in **both** of the 2 Doodles

MCU Hypervisors

Bernhard showing 2 slides (taken from a 162 slide presentation - there is more info of course)

Cores The R7 well known, R52, brand new.

Can control who has access to physical memory

RTOS1 & 2 in the picture accesses physical memory directly. Note still NO address translation.

Multiple RTOS, multiple classic AUTOSAR stacks, for example.

EL-2 MPU is the new one.

You could integrate a rich OS without letting them know they run on the HV.

Applicable only for Cortex A profile only

Hypervisor could be used but trapping accesses would be costly

Better to have operating systems are fully aware/designed.

Separation Kernel might be a more apt name for this simple partitioning (academic discussion)

Changing timing of RTE on a classic autosar stack may need recertification (for critical functions).

With this add on, a safety-critical (ASIL B) can be isolated and guaranteed its resources. – by running more than one complete AUTOSAR stack, in partitions.

Another case: Software updatability – some parts are updatable through SOTA and others cannot be affected.

Note Double or triple memory requirements. (because multiple AUTOSAR stacks) but it might be worth it.

Another use case: Heterogeneous designs. Safety Islands (often implemented in R7).

Known/publicly available info about licensees: NXP, ST and DENSO.

(i.e. R52 silicon is available now)

Adam: On the term "Hypervisor" Some call it Virtualization even when there are only very simple hardware separation features built in.

Note that Other MCU vendors with even less capability built into hardware are using the term "Virtualization support".

From Matti Möll to Everyone: 01:48 AM

<http://www.projekt-aramis.de/>

From Bernhard Rill to Everyone: 01:49 AM

<https://www.aramis2.org/>

Matti: A related info from OpenSynergy

<https://www.opensynergy.com/wp-content/uploads/2018/06/Hypervisor-for-latest-NXP-microcontroller.pdf>

Stephen: Renesas related info: Trustzone security extensions, a similar concept was applied in the R7. See documentation in the Lifecycle documentation for the SoC.

June 11th, 2019

Participants

- Gunnar
- Adam
- Alex
- Artem
- Bernhard
- Dmitry
- Stephen
- Vasco
- Phillippe
- Franz

Apologies

- Kai Lampka

Whitepaper Discussion

- Recap from the previous week
 - See outline here: [HV](#)
 - Multi OS system design on heterogeneous multi-cores as a general topic
 - System wide QoS
 - Protocols
 - The need for hardware device sharing ()
 - Hardware wish list for virtualization friendly SoCs
 - BSP drivers in general
 - firmware wishes
 - Hardware wishlists could become outdated at some point
 - We should avoid scope creep
 - Finding/Using the right terms
 - Target audiences...
- Hardware requirements
 - System design is unique and can be used in so many different ways
 - Hard to find a wish list of features
 - Make sure that mandatory and optional hardware features
 - How hard can the requirements be, what is the guiding function behind a certain hardware feature
 - Hardware requirement scope?
 - Trustzone/firmware interfaces
 - Architecture coverage (arm, x86...)
 - hardware virtualization support
 - There is a huge need to a virtualization hardware wish list, maybe is also makes sense to start a separate wish list already?
 - If hardware doesn't behave nicely, software needs to do more work
 - Expand the target audience to IP/HW vendors
 - Stephen: We might have two different topics here
 - Gunnar: The wish list is probably going to be spread out in the whole document
 - **The whitepaper should convey the thinking behind the wish list**
 - Al: Create a wish list, not a content provider
 - Bernhard can provide a description of heterogeneous system design
- Gunnar can check, where the doodle for the working session is

June 4, 2019

Minutes by Kai Lampka

a) White paper: discussion on potential sections. -everybody is aksed to llok again at https://www.automotivelinux.org/wp-content/uploads/sites/4/2018/06/GoogleDrive_The-AGL-software-defined-connected-car-architecture.pdf

For inspiration.

- (i) Why are we doing this motivation also addressing heterogenous multicores:
- (ii) Use-case of SoC partitioning into safety- and security islands.
- (iii) Clarification of terminologies, para-virtualization, TCB, microkernel-based approaches, monolithic HV, type-1 and type-2 and embedded HV
 - a. what is needed in HW to achieve this. Detail on a "wish-list" for HW-vendors to support SoC virtualization
 - b. Differentiation to containers and drawback, a critical view on containers do's and don't, same holds for HV.
- (iv) What HW can do for isolation resp. platform partitioning
 - a. Spatial isolation
 - b. Timing Isolation: Contention on (implicitly shared infra-structure) and explicitly shared devices.
 - c. Coming to future HW-based solutions, e.g., MPAM

Please consult https://static.docs.arm.com/ddi0598/a/DDI0598_MPAM_supp_armv8a.pdf for inspiration.

- (v) VirtIO as mean of
 - a. interaction of VM to VM, HV to VM and HV-off partitions to HV/VMs
 - b. Sharing of devices in the above setup. Define also different capabilities of devices, vfunctions and "virtualization-ignorant" devices

b) Meeting on Virtio in Berlin AI: Kai sends doodle link to Gunnar):

- a. Planned to meeting CW 38 for addressing on
 - i. VirIO spec contribution
 - ii. White paper as discussed above

c) Status of technical discussion (needs attendance of Artem), we defer this.

May 28, 2019

Participants

- Gunnar
- Phillipe
- Dmitry
- Vasco
- Matti

Agenda

- Discussion on white paper
 - Recap: Idea to write a white paper about the system design implications of having a hypervisor
 - Possible white paper could also better explain what virtio is?
 - How do containers play into system design?
 - We need Kai's input on this
 - Containers in the cloud usually use virtualization to enhance the separation
 - Hypervisors and containers and how they interact with system design for ECUs
 - Communication in asymmetric multi-core designs
 - Distributed systems and hypervisors
 - Software defined system architectures
 - Similar to what classic autosar was meant to achieve
 - Possible future outlook
 - possible goal for all the hypervisor standardization issues
 - How to run with and without hypervisors?
 - We should start writing a simple outline and see where it takes us
- Feedback from AMM
 - Workshop was very interesting
 - Maybe a little bit rushed through the topics created a lot of food for thought
- Possible Workshop for working on the spec & whitepaper
 - Good date would be week 38, just before "All systems Go" conference in Berlin or week 37
 - Next step, send out a poll for date options (Gunnar)
- Meeting next week
 - Vasco will be there, Matti hopes, Dmitry is off

May 21 30, 2019

Participants

Agenda

- **AMM recap**
- **Planning**

Minutes

Some notes from workshop:

Artem: Virtual audio for Xen has been in kernel about 1 year.

Artem: We created a HAL in Xen (Xen-PV), it is open-source.

Matti: New audio mode in Android, parallel to Audioflinger mode. Similar to ALSA...?

Long buffers: Virtual process dies - does not fill buffer...

Hardware

Souvik presenting for ARM

- SCMI support is in mainline kernel

May 16 - AMM workshop

Participants

- Artem, Kai, Matti, Dmitry, Gunnar, Adam Lackorzynski, ...
- + about 20-25? workshop participants

Minutes

- Lots of topics (5 presentations), lots of discussion
- Little time for Q&A on each - basically every topic needs deeper investigation and discussions
- See [presentations](#)

May 14 - cancelled due to AMM

May 7

Few participants

- AMM preparation, checking final workshop timing

Minutes TBD

April 30, 2019

Participants

- Franz Walkembach
- Gunnar Andersson
- Bernhard Rill
- Steve Liu (Neusoft)
- Philippe Robin

Apologies

None

Minutes

- Introduction, Steve Liu
- Recap of USB discussion. Findings and still open questions.
- Discussing adoption of standards. Do OEMs need to require this in RFQs before it takes off? (Similar: Do OEMs need to require VIRTIO, for non-VIRTIO hypervisors to implement it?)
- Planning [Workshop content](#)

April 23, 2019

Participants

- Gunnar
- Artem
- Dmitry
- Matti
- Philippe
- Franz
- Bernhard

Apologies

- Matt Spencer (ARM)

Minutes

USB Virtualization

- Android Debug (ADB): Google demands access to this interface for production programs
 - We have previously seen almost exclusively Host role requirements in Automotive, i.e. the car/ECU has the USB Host Role and furthermore assumed that there is little need for USB On-the-go (OTG)
 - ... but this suggests the device role will be required (through the Google requirements), and also that OTG becomes a bit more likely to be included (to "solve" the need for both Host and Device roles)
 - If hardware allows, this could of course be a single device port, and passed through to Android OS
- New Input documents
 - Hardware manual for Intel which includes description of hardware assisted virtualization.
 - ACRN implementation for USB
- USB device virtualization discussion: (i.e. the ECU implements device role)
 - In general a USB device is considered as 1 hardware unit, thus we need to document how a pass through implementation can be implemented
 - (NOTE: In this discussion we were jumping a bit back and forth between Host and Device virtual discussion)
 - We can (for both Host and Device) look at the ACRN implementation, Gunnar showed a picture which Artem has sent to him
 - Its USB architecture shows that you have a Service OS which talks to the phy mux control and the xHCI controller
 - The USER Operating Systems have only access to the xDevice Controller Interface (DCI)
 - OTG needs to be covered as well
 - Discussion: Re-use of interfaces from the Intel ACRN hypervisor can be considered for a standards attempt?
 - TBD: if this interface does cover other architectures e.g. ARM as well?
 - Rensas and Qualcomm have different implementations
- Host role virtualization...
- Need a way to assign policies for sharing
- Discussing EHCI, XHCI, Is assuming XHCI transport a way forward?
- Matti: But we need some kind of hardware neutral abstraction
- Artem: Agreed, but XHCI seems to be a kind of superset - could it be the HW independent (superset) interface?
- Artem: VMM enumerates device in DOM0
- Gunnar: This is what I mean by Hypervisor must implement some part of stack, handle attach and initial reading of device IDs for example, if they are to be filtered and passed to only one OS. But what we mean by "in the Hypervisor" might differ in different HV architectures.
- Use-cases which need to be considered for USB host mode virtualization: USB tethering, and phone "projection" / Apple Car Play, USB Switch support, mass-storage-device, Traffic Signal & Toll-booth beacon external device.
- Limitations, often only one USB controller, therefore switch.
- How does synchronizations between the host/service OS (Xen: Dom0) and User OSes (VMs) work?
- We need more use cases...
- Gunnar: Yes but maybe enough to agree the need is there (for Host virt). E.g. those device types that were mentioned already. E.g.: Mass storage to one OS, toll-booth device to another.
- ...
- Device controller
- Matti: How does XHCI map to MMIO based controllers?
- ... Right we must consider how IOMMU virtualisation affects this
- Artem: Yes it depends if PCIe is used or not (XHCI usually on PCIe)
- USB over IP implementation in Linux - is this an appropriate/reusable protocol?

AMM

- Arm SCMI presentation at GENIVI AMM (HV workshop), Arm experts will be there. Gunnar to check if there is an space for a speaking slot

April 16, 2019

Participants

- Gunnar
- Dmitry
- Artem
- Philippe

Apologies

- Franz

Minutes

Booting

April 9, 2019

Participants

- Gunnar
- Matti
- Dmitry
- Franz
- Artem
- Bernard
- Philippe

Minutes

Booting

- OpenSynergy proposes AEM EBBR (<https://github.com/ARM-software/ebbr>). Gunnar is OK and asks about VirtIO support. Matti points to the Tianocore EDK2 project that has support for both VirtIO and XenBus (<https://github.com/tianocore/edk2>).
- Gunnar creates a new "Booting" chapter in the specification (to be filled).
- Matti: Linux and Android have different requirements for booting: Android needs specific parameters to be passed. U-Boot has a special vboot (verified boot) utility. With EBBR we can pass those as UEFI vars (<https://github.com/tianocore/tianocore.github.io/wiki/ArmPlatformPkg-AArch64>).
- Franz: we should not specify strict requirements for booting and any specific APIs. Gunnar: we are not specifying how a hypervisor must boot, just that a hypervisor supports booting in this way as well. For some devices we need to have strict rules as we are working on the specification. Matti: the booting procedure description tells how to boot a guest under a hypervisor, not about booting a hypervisor on some particular hardware.

ARM SBSA Debrief Contd

- Gunnar continues on power states. Matti: PSCI is very arch pecific. Gunnar: we might place a requirement: if ARM, then use PSCI. Matti 2 Artem: does Xen use PSCI? Artem: On ARM - yes. There are new power management functions yet to be merged. Those will be covered by a new SCMI spec. Franz: PikeOS also uses PSCI.
- Gunnar: is PSCI supported on all ARM architectures? Matti: no. Custom interfaces can be implemented for legacy solutions. Franz has concerns about performance.
- Gunnar: EPAM suggests SCMI for sensors. Matti agrees with the proposal. Matti: not all HW platforms support SCMI. Pure virtualized SCMI solutions and solutions with HW assisted SCMI virtualization might exist. SCMI is also not really ARM specific. Bernard can check for available ARM experts to discuss this on AMM.
- Gunnar wants to discuss platform independent

April 2, 2019

Participants

- Kai
- Gunnar
- Matti
- Dmitry
- Franz
- Kai
- Matt
- Philippe

Minutes

ARM SBSA Debrief

- Gunnar: Walking through the ARM SBSA presentation to pick up comments / feedback / further discussion points

raw notes (to be checked by Gunnar)

- discussion on HV booting
- Matti: UEFI - bootloader less flexible than uboot, it might make sense to reference UEFI work, but UEFI is "a beast"
- Kai: @Matti you put the complexity in the device tree, what do you think ?
- Matti: mentions ACPI
- (+Some reference to Microsoft Windows... basically common ways of booting operating systems on computers)
- Gunnar: what about the Embedded/Real-Time OSes? [They don't support the same booting methods?]
- Matti: Basically, Linux and Android can be booted using UEFI
- Matt: makes comment about EBBR spec (Embedded Base Boot Requirement)s, it is time to talk to linaro guys IMHO
- Artem: there are some [embedded] platforms that use already UEFI
- Artem: we can summarize some of the requirements from EBBR
- /TODO/ Artem summarize EBBR requirements in the automotive virtualization platform spec
- Discussion on booting AUTOSAR --> (the exact kernel is not specified)
- unidentified: important to be agnostic to which AUTOSAR OS is booted
- Matti: very valid point to consider, a couple of years ago, boot was project specific, separate boot protocol per ^project, this has changed now !!
- short discussion on how AUTOSAR might engage in the booting process specification
- Back to SBSA - power state coordination interface PSCI
- Matti: on arm systems, all hw and hv vendors support PSCI, note that PSCI is for cores only, SCMI is for peripherals
- Artem: discussion on SCMI specifications (not captured), how to write specs for this ?
- Matti: would be very difficult to harmonize this, automotive is a niche at the end, and cannot be maintained elsewhere
- Matt: do we think this is a requirement in the automotive space? [is what a requirement? sorry missed the context]
- Matt: mentions the thermal constraint, we might have to shutdown some VMs due to overheat

HV Workshop topics

- Matti: 1) Audio virtualisation - update on our VIRTIO proposal, and discussion
2) Video (codecs and camera) - present and talk about it
- Artem: 1) Xen-PV as option to VIRTIO protocol
2) Native TrustZone access for guests, including our recent OPTTEE changes
(Matti requests also this)
- Gunnar: We also have Kai's presentation on VIRTIO "natively"

March 26, 2019

todo

March 19, 2019

Cancelled

March 12, 2019

ARM Presentation on Server Base System Architecture

- Recording
- [Presentation slides](#)

March 5, 2019

Minute taker: Vasco Fachin

Hypervisor Genivi Meeting Meeting Minutes 05.03.2019

Participants :

Vasco Fachin

Gunnar Andersson

Philippe R

Dmitry Morozov

Franz Walkembach

Bernhard Rill

Later:

Sang-bum Suh

Much Later:

Stephen Lawrence (Renesas)

Notes

Gunnar: Sang-bum Suh contact over email, should join today

Franz will join the Renesas Meeting, Philippe will show the Demo done at CES

Summary of last week

Gunnar: Review of last week. Kai's proposal for the new topic, then talk about physical and virtual Audio Driver. Input from Matti. Information is present in the wiki page. Open Sinergy has a proposal for a VIRTIO Audio.

Franz: Shared audio from Hypervisor perspective.

Gunnar: Shall be similar to what Matti presented. Maybe you (Franz) can take a look at what's been discussed.

AMM Participation

Gunnar: Waiting for Sang-bum Suh, looking at the AMM attendance.

AMM company overview/experience on how what to deal with hypervisors: SysGo, EPAM, OpenSynergy have been reached out.

Dmitry: Have to check if this is possible

Gunnar: Artem and Lars (Xen) have already planned the session.

Sang-bum, Bernhard (maybe other colleagues) joins AMM, Dmitry and Matti probably not, Franz will probably not join. Vasco has to clarify.

SBSA

Bernhard: Next week we shall have the SBSA presentation at 11 CET. If there is interest, some people from ARM could join the AMM meeting. Needs to be clarified until next Tuesday

Security

Sang-bum: Hypervisor System Architecture includes Cryptography.

Root of trust: Hardware provides unique random number used for crypto modules. The Crypto Function can be used among different OSs. Depending on how many OSs are depending on the underlying crypto modules, these might be virtualised.

- TPM service cannot be multitask, single threaded. Cannot handle multiple requests.

- virtual Crypto Modules shall be included as part of the Guest OS stack, replacing hardware with virtual modules, maybe using VIRTIO

Hardware will initiate some unique random number, that cannot be copied. Needed in multiple VMs

Sang-bum: we need to get feedback from Car manufacturers what is the desired configuration.

Gunnar: we might need to anticipate it, with some use cases.

Dmitry: not my area of expertise

Gunnar: we might need to reach out for experts from each company. Or what is the proper way to deal with this.

Franz: no feedback to the topic. Could probably try to get some internal feedback. Security is a big topic as of now. Every customer seems to have different ideas.

Gunnar: This should be the case, since there seems to be a consensus on which is a suitable security feature. But we don't want to go around in circles.

Franz: Security Group within Genivi?

Gunnar: not at this stage.

Franz: we need to summarise the discussion in order to be presented to other company members.

Vasco: also from me, no expertise in this area.

Sang-bum: We need to refer to the OEMs

Gunnar: This is important but we should not allow it to be a blocker. (Each company here has experience about systems] ; blocking might delay the discussion.

Sang-bum: we need to look at the whole system architecture, not stay limited to some component, because we are not experts in all the components.

Updates from Embedded World

Stephen: Renesas announced RH850 microcontrollers with hypervisor support:

www.renesas.com/eu/en/about/press-center/news/2019/news20190225.html. Focusing on classic or adaptive Autosar.

Block Diagram www.renesas.com/eu/en...6.html

Bernhard: ST presented something similar at Embedded World. www.st.com/content/st_com/en/about/media-center/press-item.html/p4141.html

February 26th, 2019

Participants:

- Matti (OpenSynergy)
- Guru (Bosch)
- Kai (Elektrobit)
- Matt (ARM)
- Dmitry (OpenSynergy)
- Philippe
- Gunnar
- Albert (Continental), last minutes of meeting

Apologies:

- Artem (EPAM)
- Anup (this time slot always a problem)

Minutes:

VIRTIO for "non-virtual" interfaces

This idea concerns communication between operating systems, especially running on different cores in an soc, with or without a hypervisor.

Kai: for example classic autosar rtos owning an ethernet connection running on smaller dedicated cpu core. ethernet connection can be exposed to multimedia cores using virtio network queues.

...idea extended? block devices? what else?

...needs an interface that exposes addresses, so that knows where to fetch the ethernet frames. i.e. shared memory, i.e. physical addresses

Matti: are you using openAMP for this?

...you could use virtio protocol also.

Matti: (previous work) Xilinx + mentor published openAMP for this purpose? I think TI showed something like it on Jacinto 6 .

...VIRTIO network device exposed to multimedia processor.

Kai: It ties into Multi-OS integration focus of the AMM

Audio in virtualization

Gunnar: I'd like to get to a conclusion on this topic, based on previous discussions but hopefully new info

...Let's start with, Hypervisor exposes an interface to VMs. What does it look like?

Matti: Let's start with how hardware works. Audio is a continuous bit stream. Hardware usually has multiple channels. Different hardware knows different buffer layouts. The HW walk through buffers in a timed manner. Take out the data and convert to analog audio and play it.

Gunnar: We got pointers before, quite a while ago, from Artem on the implementation support in Xen - including some parts that are in the kernel tree.

...(But we need to bring this up a level, to standards)

Matti: Xen, as far as I understood. It sets buffers, allows you to access PCM streams. Real hardware works typically with plain PCM format. Also virtual hardware usually accept PCM.

The most common model in virtualization:

Emulating **Intel sound card, HD97 audio controller**. Specify PCM, bit, length, sample rate. Then dump stuff (audio data) in there.

...It's similar for USB audio. You register a couple of buffers. You send data over. (To summarize, meta data + content, much like any other device to be honest). Not that complicated. The only tricky thing is it is timing sensitive.

Gunnar: Yes let's talk about the real time demands, and virtual machines?

Matti: Hardware signals overflow or underflow. Sends interrupts, so that software starts filling buffers.

...But there is often a lot of slack in audio. Fast microcontrollers can nowadays keep up with what human ear allows for.

Gunnar/Matti: ... Human ear allows for up to several tens of milliseconds, at least.

(Brief sidebar on audio/video sync, a.k.a. lip-sync)

Matti: On things like latency, note for example Android compatibility document CDD, has strict audio (timing) requirements.

...(The latencies) are hard to measure and evaluate.

Gunnar: And this would apply also to "virtual hardware"?

Matti: Yes, if you're running Android as a guest, the requirements are there - they are more about pipeline latency all the way to the end user. They are there to guarantee everything.

...you may have an external amplifier on a network like MOST. The whole chain must fulfil the requirements.

... Once all the network delays are considered, the remaining time for virtualization delays is very small.

Gunnar: (OK... standards)

Matti: So, we built a PoC for VIRTIO. We want to send this upstream. Talked to the "guardians" of the specification. We are still finalizing details. Expect to send out first RFCs (to VIRTIO mailing list) in the coming months or so.

...Prepared this with Michael Tsirkin. from RedHat, and others. Relatively positive.

Matti: The PoC has good enough timing for video calls.. We have not measured exact delay.

It is a protocol specification for VIRTIO Audio.

...(From VIRTIO mailing list) Note that Google emulator guys are also looking into audio virtualization. Used for Android emulator (development environment). Want to add support to mainline QEMU, to reduce their own maintenance.

(more details, what's in it?)

Matti: Buffer exchange. Define which audio streams can be supported? Discoverable, number of channels, direction, layout, etc. Some hardware prefers certain buffer layouts, certain sample rates, etc. Application should send the right formats.

Gunnar: ... If the hardware cannot process a certain format, it must get what it needs. It is better to do this conversion close to application then? Hypervisor layer should not do processing

Matti: Upsampling is easy - so that could be done. Sample rate conversion with good quality takes more.

...That is often offloaded to DSP or even more specialized hardware.

... So (that hardware) could be used by a Hypervisor to implement some such features.

Gunnar: So we should leave it open how advanced Hypervisor implements processing of audio. This could be unique feature per implementation?

... (Virtual) Device can report its capabilities and it will either be rudimentary or advanced, the guest have to adapt...

Matti: (I'd say so).

...Better hardware support means less interrupts, less waking up the CPU. So hardware could support a deep pipeline, e.g. take MP3 input instead of

PCM. But it is complicated and doesn't even save that much battery any longer, so it has mostly been abandoned. It doesn't work well with content protection (DRM) .

Matti/Gunnar: Summarizing: DRM/security is too complex in hardware pipeline, and it gets out-of-date).

Matt S (ARM): ... let's talk a bit more about DRM. How is it covered?

Matti: Yeah, I think those questions are out of scope. ...It's a pure buffer (PCM)
... (and the other details) Precedence, priority, loudness, mixing...

Matti: (Note that) Mixing doesn't fit into the device protocol.

Matt S: Any Focus on Functional safety?

Matti: Yes, but

Gunnar: But DRM needs to be in scope, right. At least to protect the unencrypted buffers?

Matti: Audio DRM systems seem to go (*protect*) until they reach a PCM buffer, because on most systems it is possible to record this final stream in some way anyway.

Gunnar: Guess most protection today focuses on video content.

Matt S: I'd think that Premium Dolby sound is still considered important (to protect)
... I'd like to reach out to Dolby to get their input on the whole discussion.

Matti: That would surely be very appreciated input.

Matt: (for such protection) Key exchange - does it need to be part of the protocol?
...If we have a mean to talk to an end point - we need to make sure the keys are in place. Dolby will probably have a special codec, probably an encrypted stream. When does key exchange happen?

Gunnar: (*Agree, at least the occurrence of this exchange should be in the protocol, (or explained when it happens).*)

Gunnar: Then method itself I guess uses standard methods, Diffie-Hellman, etc.

...

Gunnar: OK how to define how we share or pass through DSPs and custom hardware (like codec support and Sample Rate Conversion)

Matti: It is also very hardware specific. What can be done. Some (hardware) allows pass-through of these decoders, some make it very difficult to do so. Some discourage it because they don't have an IOMMU.
...DSP is more likely to be used for post-processing (*than for MP3 decoding*). It's a power saving feature and not a performance thing.

Gunnar: OK. We discussed earlier meeting if (*at least*) the method of configuration for pass-through devices can be standardized (*across HVs*)
... *Although it's clear and open in Xen (standard Linux format device trees, for example).*

Matti: I agree/confirm that how that device tree is created is proprietary. Until it reaches the guest, things are highly design/implementation specific.
...I think most OSes use a device-tree approach. e.g. some RTOS are using device trees too.

Gunnar: But other OSes have different formats (of DT)?

Matti: Yes. But I think basic structures are fairly similar at least, but I'm not too familiar with other OSes.

Gunnar: So a type of device tree *is* the specification for pass-through devices? And the creation of it is proprietary.
(*consensus*)

Wrap-up.

February 19th, 2019

Participants:

- Bernhard Rill
- Gunnar Andersson
- Dimitri Morozov
- Vasco Fachin
- Kai Lampka
- Sang-Bum Suh
- Artem Mygaiev

Apologies:

- Philippe Robin

Review plan for next meetings

Gunnar: Table has updated status. Let's review...

Crypto -

Input from Sang-Bum?

S-B: Should be included in VIRTIO. There is one chapter...

Artem: Let me share what I know:

...That component would be utilized by the guest operating system. The cryptography accelerator components.

...When we are thinking about TPM chips. Those are single-threaded machines.

...so that's why VIRTIO proposed virtual TPM.

...We haven't seen in the past that this...

...Security cannot be provided for threaded programming support. Virtual crypto, should be discussed.

Artem: For ARM it works like this.

...The execution level of security is higher than the hypervisor. If there is a task running in EL3 it will always have higher priority than hypervisor.

...Therefore, Hypervisor cannot influence what's happening on the secure world.

...Note of course SEL2 implementations coming in next ARM architecture which we talked about before.

...To implement virtualization support, we have upstreams VM IDs to OPTEE. This is already included. R-Car is supported.

...For the guest on ARM, if OPTEE is used (or any other TEE supporting virt) they will be transparent operationally hardware sharing inside trustzone has to be implemented, again out of control for HV.

...Architecturally I can only explain how that work on current ARM architecture, not Intel.

V8.4 There is the alpha draft of the specification available.

Bernhard: Yes, the one we discussed before. NDA only -- you need to send me an email to request. (ongoing)

Artem: Thus on ARM even with current implementation, HV is out of picture.

...You need to support it in TEE software. You also need to implement sharing on the secure device drivers level.

...The guest will be able to access it directly without virtualization support

... HV needs to pass ??? VMID

.. with security policies you may allow or disallow security, trapping some syscalls.

... Intel is probably different

Sang-Bum: It might be true, if sec controlled by ARM SoC hardware

...Running trust zone area. HV only running at normal.

...Before HV applied to ARM, operating system utilize ARM hardware, should be embedded within that operating system.

...Applying HV at normal domain, the operating system manage the ARM hardware, also should be run on hardware in some cases.

...HV only has control over ARM hardware, if the other guest operating system

Gunnar: Sang-bum can you put some of these thoughts into writing. (I cannot take notes fast enough)

Sang-Bum: Yes.

Artem: for Embedded, there is a standard called Global Platform. Many trusted execution environments make use of this.

Sang-bum: This is covering trusted execution environments. We can use global platform APIs.

AI(Artem): Provide references go global platform

Planning SBSA timing

12 March: CET 11-12 is OK?

Sang-Bum: Yes, 7PM for me, it is OK

Artem: I might be travelling on 12th... I will know in a few days.

Gunnar: I think this time was generally OK with Anup. I will check.

Sang-Bum: I'm not available next week tuesday, but 5th March is OK.

AI: Where is Global Platform defined?

AI: Gunnar - confirm 11.00 for 12 March

USB Virtualization

Gunnar: Let's go through the proposal I wrote based on previous discussions

Artem: Actually I have found more info. Some solutions for virtualized USB exist. I don't yet know if they make sense and are useful or not.

Gunnar: Please provide links.

Artem: I will check (until next week) what is possible with respect to USB virtualization.s

Gunnar: I then ask if the *configuration* of the pass-through can be standardized

Artem: Well at least on Xen it is all defined using device-tree (according to Linux kernel standard).

... Some patches for partial device tree. In other words, HV gets a big one, and splits the pass-through according to some, configuration and injects the results into the guest DTs.

Gunnar: So that is Linux (guest) specific? Just let me ask, for example, Xen runs Windows - how would that work?

Artem: Windows - done with ACPI on Intel at least... I have not looked because I work on ARM and Windows on ARM is...well has a difficult history.

Gunnar: OK, sure, but let's say FreeRTOS...

Artem: Everything has to be statically configured. We do the configuration at compile-time in that case (compiling the RTOS hardware part)
 ...Configuration can be passed to other domains
 ...In smaller RTOS they don't always support injecting a hardware description like device tree so we need to recompile.

Gunnar: Other vendors? Are there different ways to configure this in every hypervisor?

Sang-bum: I think it's vendor specific. It's hard to follow for example device-tree without licensing concerns (if your software is not GPL licensed)

Gunnar: I think I understand... it might be difficult to reuse major things like device-tree without also reusing the implementation of it.
 ... but you could write your own specification and implement it?

Sang-bum: I don't think we need to discuss a standard specification for how hardware is integrated. It is a kind of implementation detail.

Gunnar: How is it done in OpenSynergy, if you want to share that detail...

Dmitry: I think some XML based configuration files. And special compiler to capture this input. I agree this is implementation specific

Gunnar: OK, statement about this being proprietary is basically correct. I might still rewrite it a bit based on this input.

Kai: I have sent a proposal for talk at AMM. Idea: VIRTIO might be used to get operating system interaction also when running on bare hardware.

(What follows is a discussion (on overtime) between Artem and Kai about this idea. We expect more interesting discussion on this.)

Kai: consider different operating systems running on different (dedicated) cores. E.g. RTOS on a small core.

Artem: With this idea, sharing hardware between real-time and non-real time domain worries me.

Kai: I meant this, you can use VIRTIO protocol to exchange between domains.

Artem: OK, but you don't need the whole VIRTIO then. There are often hardware mailboxes for example, as alternative (to virtqueues)

Kai: Think about for example VIRTIO NET.

February 12th, 2019

Participants:

- Bernhard Rill (taking minutes)
- Philippe Robin
- Gunnar Andersson
- Dimitri Morozov
- Vasco Fachin
- Kai Lampka
- Erik Jaegerwall

Agenda:

- Date for SBSA presentation from Arm
- API standardization / virtual platform definition / Backlog refinement and planning
 - [\(Topics intro and table summary here\)](#)
 - [\(draft spec work here\)](#)
- different timeslot for a call that Anup can join for the watchdog discussion.
- Access to ArmV8.4 Trustzone API specification (Secure Partition Interface specification)
- Events

Minutes:

- Date for the SBSA presentation is 05th of March. Charles Garcia from Arm is going to execute this presentation. Exact timeslot will be defined based on the doodle from Gunnar (note below)
- API standardization:
 - Dmitry mentioned that OpenSynergy is working on an Audio implementation (they evaluate if they can transfer Audio frames via VirtIO)
 - It might be that OASIS is working in this area. It might be good that OpenSynergy gets in contact with them.
 - EPAM Artem has summarized the sensors section. Feedback from the team is needed
 - USB: Vasco mentioned that a full virtualization of an USB is very hard; it might be desirable to setup a pass-through device
 - What are the use-cases which would lead towards a need to fully virtualize an USB? TBD
 - Limiting factors for a full virtualized USB are
 - complexity of the USB (DMA support)
 - Security
 - **Action item Vasco** is going to summarize his view on USB in the summary table
 - Bluetooth virtualization is most likely not needed
 - Memory ballooning is as well out of scope
 - Gunnar has updated the chapters network, cryptography and text console. Feedback is welcome!
- Different timeslot for a call that Anup can join for the watchdog discussion.
 - **Action item Gunnar** will send a Doodle link
- Access to ArmV8.4 Trustzone API specification (Secure Partition Interface specification)
 - let Bernhard know if you would like have access to this specification (under NDA)
- Events:
 - Embedded world

February 5, 2019

Participants:

- Franz Walkembach (taking minutes)
- Philippe Robin
- Gunnar Andersson
- Dimitri Morozov
- Vasco Fachin
- Andreas Prosell
- Artem Mygaiev
- Erik Jaegervall
- Vivek Galatage
- Matt Spencer
- Stephen Lawrence

apologies

- Kai

Agenda:

- Review previous week's minutes
- Virtualization talks at the [FOSDEM conference](#)
- API standardization / virtual platform definition
- Backlog refinement and planning
- [\(Topics intro and table summary here\)](#)
- [\(draft spec work here\)](#)

Minutes:

- Introduction of workgroup members and available content to the new team members
- Review of last meeting minutes
- Looking for RISC-V knowledgeable persons at Genivi
- It was stated that RISC-V might be a good member for Genivi, but currently not important for the HV workgroup
- AMM in Munich in May is approaching quickly
- Gunnar gave feedback on the FOSDEM conference (link in Wiki) on 2 and 3 February in Brussels with some presentations on virtualization topics and tools, secure boot, virtual IOMMU implementation using HW nested paging. Virt audio by Red Hat
 - Please have a look at the videos and content where interested.
 - Fosdem tracks - recommended readings
 - [virtualization and IaaS](#) track - look in particular at
 - [Getting To Blinky: Virt Edition](#)
Making device pass-through work on embedded ARM
 - [Virtual IOMMU Implementation using HW Nested Paging](#)
 - [Software Defined Networking](#) track - look in particular at
 - [An update on Virtio](#)
What happened recently and what's coming soon?
- Artem: one update I wanted to share - finally Linaro OP-TEE (FOSS implementation of secure OS and environment for ARM Trust Zone) have just merged support for virtualization of guests running Trusted Applications that we've developed @EPAM & Xen Project: https://github.com/OP-TEE/optee_os/pull/2370, this allows native execution of TAs and can be included in spec
- Draft of virtual platform definition, API and synergy document created by Gunnar, incl. auto requirements and scoping. Gunnar will post it on the Wiki for people that are interested
- Erik is working on a group at Bosch focussing on security and was wondering how security influences the HV work group. We should work on a top level feature overview.
- Andreas at Luxoft with PELUX open source platform that might include hypervisor technology including shared graphic functionality.
- Vivek will investigate the work group content and get back with topics
- Gunnar highlighted the contribution need to make the group successful
- Philippe reminded all to propose topics for the AMM in May

January 29, 2018

Participants:

- Gunnar Andersson
- Artem Mygaiev
- Bernhard Rill
- Christian Schulenberg
- Dmitry Morozov

- Franz Walkembach
- Kai Lampka
- Matt Spencer
- Philippe Robin

Minutes

- Last week's meeting review. Franz: RISC-V consortium and GENIVI board cooperation should be discussed. Gunnar sees no issue in trying new hardware architecture in the technical project, and GENIVI organization is generally hardware-independent. Gunnar requests appropriate contact (from Franz), especially if someone works on Virtualization technologies.
- Participation. No updates from Green Hills regarding a new participant *We have not reached out to our contacts, Philippe will do it.* Anup has a conflict on this time slot but wants to participate via email. Gunnar proposes to shift the meeting (some of the time, not all of the time), for Anup, but also (maybe a different shift) to allow people from the US to participate.
- SEL2.
 - Bernard: API specification (alpha revision) is about to be released. It will be accessible only under NDA. Artem proposes to create an internal review team.
 - Bernard could invite colleagues to discuss SBSA (possibility to join in two weeks or so).
- Draft specification. Gunnar: Last week I felt we did not have full agreement on what we want to achieve. The project description I will show later is one input to that discussion.
- Come up with ideas for tech presentations for AMM meeting in Munich (May 14-16). Two weeks deadline has been proposed.
- PV IO interface by Xen as an alternative to VirtIO. Artem proposes the topic for discussion. No documents with description so far. Artem is to share related links.
- Reviewing Virtual Device Categories table. No progress. We need to clarify VM responsibilities.
- Also we need to convince people from OEMs to participate. We need requirements in general regarding usage of hypervisors in automotive.
 - *Gunnar agrees, but proposes that we also know fairly well (as suppliers) the generalized use-cases and needs on a high level. It is unfortunate but should not make us stop completely.*
- From interest to participation. Gunnar introduces the Project description: a high level idea of requirements and respective solutions. OCI as an example. Everyone is welcome to read and give some feedback. Franz asks about the list of VMs (OSes) we want to support. Gunnar: Any OSes that are in the scope of automotive industry (for example Linux (multiple variations), Android, QNX, INTEGRITY, AUTOSAR Adaptive).

January 22, 2019

Participants:

- Gunnar Andersson
- Artem Mygaiev
- Kai Lampka
- Philippe Robin
- Dmitry Morozov
- Franz Walkembach
- Bernhard Rill

Minutes

1. Franz suggested GENIVI community looking at RISC V architecture – this might become a production grade at some point
2. Gunnar goes over last call's minutes
3. GHS: Nikola cannot join anymore because he changed position, GHS said they are interested in continuing their participation to the work
4. Bernhard don't have updates on SEL2 / SP805 yet
 - a. Mentioned that discussion can happen with ARM engineers on WDT – but this is part of "supervisor" not hypervisor. This can be reviewed in scope of SBSA
 - b. SEL2 – maybe some spec updates? Not yet
5. Gunnar pointed out the need to take time schedules more aggressively
 - a. Philippe suggested to use May 14-16 AMM in Munich as a milestone for some 1st version release and holding technical workshop on virtualization & spec review
 - b. Gunnar requests for topics/proposals for presentations on GENIVI AMM (Artem wants to suggest reviewing pvif as alternative to virtio)
6. Reviewing Virtual Device Categories table yet
 - a. Block storage – no updates yet from Kai, Gunnar is reviewing goals and approach. One of the main un-covered topics is "automotive-grade persistency management" – does this imply some requirements on hypervisor level? Artem – hypervisor just shall not interfere. Gunnar – there may be 2 approaches: write-through with "privileged" processes or going through all system layers as usual. Kai – write down the requirements. Discussion continues into discussing API stability for all scenarios. Gunnar – does virtio provide everything needed for technical implementation of abovementioned? Kai – no, probably not, it does not allow applying policy, etc. Need to define what is missing, requirements, etc. From example of 2D GPU, Kai is asking – from this table it is not clear if virtio fits requirements. Gunnar – requirements must be checked against virtio spec. Kai – requirements as of now implying some design choices which may not be correct. Gunnar – yes, but requirements are there to support different implied designs, not to enforce them. Goal is to make a generic specification which can be tailored for a target system.

January 15, 2019

Participants:

- Dmitry Morozov (OpenSynergy)
- Artem (EPAM)
- Gunnar (GENIVI)
- Bernhard Rill (ARM)

Minutes

Fairly short sync up on a few topics and planning the work going forward.

Gunnar: I'm thinking of writing a short introduction chapter for SCMI to make the proposal more concrete. Hopefully this sparks feedback and it's an important chapter.

Artem: SCMI is being updated by ARM. So we should take that into account.
I think we should participate in that kind of spec update.

Bernhard joining...

Bernhard: I can find the right contact people for this.

Artem: Is there any updates on security?
...like SEL2 implementation for ARMv8

Artem: For your information we just upstreamed to OPTEE and to TrustZone support for previous architectures.

Bernhard: On the Watchdog topic:
...The SBSA, Server Based Systems Architecture, includes Watchdog info that could be relevant.
...I can find a team from ARM to give more information about this.

Artem: Is SBSA only for enterprise domain?

Bernhard: No we write it independently. It is part of the architecture itself.

Gunnar: I notice Anup's intro text links to an ARM spec on the watchdog part but that linked document seems chip-specific (I mean ARM-architecture specific). I think for a cross-platform specification we're looking for something slightly different.

Gunnar: Of course wherever we end up seeing platform differences for a feature, that may be an outcome as well
...if common abstraction is not appropriate for some reason. The important thing is to clarify this and document it.

CES:

Some reflections... What was most shown and discussed at CES? (Autonomous vehicles)
and what was the most interesting for you (Gunnar: For me 1000+ people in GENIVI/NAB showcase event. Not the easiest place for good business discussions (loud) but it's a fun big meetup place at least.). Also, LGE keynote.

December 11, 2018

Participants:

- Kai Lampka (Elektrobit)
- Alex Agizim (EPAM)
- Kamalakkannan Natar... (Bosch)
- Anup Patel (XVisor)
- Dmitry Morozov (OpenSynergy)
- Philippe (GENIVI)
- Gunnar (GENIVI)
- Matt Spencer (ARM)
- Vasco Fachin (Bosch)

Notes from HV meeting 2018-12-11

Gunnar: Did you perform the action to decide on support (or not) for SCMI-for-sensors proposal?

... SCMI being discussed in OpenSynergy

Dmitry: ...company working to form an opinion, will come back.

Dmitry: It seems to be more features on backends side clocks and power domains

...Is SCMI an appropriate protocol

Gunnar: Yes, the specification itself highlights those features. The proposal that this will be repurposed as a standard sensors API is the new idea.

Kai: Not so much focus on that right now

Kai: feedback and a review should be possible

Alex: ... need to give feedback

Alex - noo need for input because EPAM

Gunnar: SCMI - more comments?

Matt: I'd rather have Bernhard comment on it.

Vasco: I need to look into this. Kind of a new topic.

Gunnar: Still open question for me is Intel suppt? I'd consider contacting ACRN project as a start.

Matt: I'd like to think that it can be adopted by Intel, will check if we have had any feedback.
Will check with Bernhard Rill.

Virtualization spec

Dmitry's input:

SHALL/MUST Reuse definitions from VIRTIO

Gunnar: I notice a mix of terms here.

Dmitry: Yes, shall and must are the same thing.

MUST seems to be used as positive and SHALL NOT used as negative (in VIRTIO)

Discussion with Kai on block device

Some features might not be covered by VIRTIO in enough detail.

Kai: trim command is worthwhile to describe.

Gunnar: OK can we then try this as a start - describe trim command in detail, specification, could be requirements... It should be concrete - we're writing a specification...

Gunnar: I want to get the AGL review off the backlog, it's been open long enough. Please read the paper, figure out what is useful / not useful in the context we are working. Just like source code we should reuse already done work.... (Just like others may reuse what we produce in the virt platform spec).

Gunnar: More from backlog - the architecture & requirements work thread has stopped because of lack of input and previous driver Sriram is busy.

Gunnar: Have you looked at this within Bosch? Vasco?

Vasco: I will consider this

- Architecture
- Use case
- Requirements

Gunnar: You can address it from any angle you like of the above 3, it will lead to the kind of results we mean here.

Action (all): AGL paper read through (again), see links

December 4, 2018

Participants:

- Kai Lampka (Elektrobit)
- Kamalakkannan Natar... (Bosch)
- Anup Patel (XVisor)
- Dmitry Morozov (OpenSynergy)
- Philippe
- Gunnar
- Bernhard Rill (ARM)

Minutes (TBD)

- ...Lots of discussion on block devices and realization of automotive persistence requirements in a full system.
-

November 27, 2018

Participants

- Bernhard
- Franz
- Artem
- Alex

- Vasco
- Dmitry
- Gunnar

Dmitry: I have written 2D and 3D graphics requirements on the [Specification draft page](#).

Reviewing Dmitry's input...

Dmitry: 2D is quite clear

Gunnar: The 2D chapter references VIRTIO 1.0 from 2015. Is it up to date?

Dmitry: Yes, I think so. Not much change (in 2D area)

Dmitry: VIRTIO-VIRGL - empty reference for now. There is no spec to point to yet. I can only refer to kernel source etc.

...hope to see something coming here.

Gunnar: Great start, comprehensive and detailed work!

Gunnar: But it is very minimalistic, which we want of course, but maybe too much.

... We need to write each topic/chapter names so that people can see what each requirement is about.

Dmitry: OK, I will fix that

Gunnar: We should also write some introduction text to topics. Put in a paragraph or so at least. If a particular topic is not covered in VIRTIO remember that we still need to put a placeholder title in there.

Gunnar: Asking the group generally: usefulness of specification? Are we still doing the right thing?

Franz: In some areas PikeOS has its own APIs and we do not use VIRTIO

Gunnar: Understood, and that's a choice anyone is free to make. To be clear we're not saying everything must be exactly according to VIRTIO, but of course other proposals must be laid out if they are going to affect standards work.

...Going through the table to review last week's discussion.

Gunnar: USB (repeat) Any device role? Device role? Anyone?

Gunnar: What are the use cases, you all work in real projects so you are the ones that might know of any.

Franz: I see mostly connection mobile phone, provide streaming video via hone.

Gunnar: That is normally still car in host role and phone in device role I believe.

Gunnar: Anyone else know of any case where car has device-role?

(none answered)

Gunnar: Artem, Updates?

Artem: Sorry been travelling. Will be more up to date next week.

Vasco (chat): Have to drop off...

Going through the table a bit more. Mostly we need to collect and consolidate what has already been created.

Gunnar: Bernhard, updates or thoughts? Your ARM colleagues have provided good input the last few weeks, for example about what lies in the future, i.e. what's still open topics in virtualization.

Bernhard: Sorry, back from vacation, also need to catch up on notes first.

Gunnar: 😊 OK, no problem. Let's reconvene next week with more energy. We now have a good draft that will lead the way for other topics.

From Tuesday, October 23 to date

iteration on the following wiki page content

- [API standardization / virtual platform definition](#)

Tuesday October 16, 2018

Participants

- Lars Kurth
- Dmitry Morozov)
- Anup Patel
- Artem Mygaiev
- Kai Lampka
- Dmitry Morozov
- Gunnar
- Franz Walkembach
- Matt Spencer
- Vasco Fachin
- Alex Agizim

Apologies

- Bernhard Rill
- Philippe
- Sang-Bum Suh

Minutes

Quick de-brief

Gunnar: Questions coming up after ARM presentation?

Artem: No immediate questions from me. Julien is main ARM Xen maintainer - he is reviewing the material.
...but it would be interesting if we can come up with a bit more use-cases than the one mentioned in the presentation.

Gunnar: Agreed, I'd like to extend that question to all the HV vendors. Please give your opinion at a later meeting (on where *you* might like to use the secure execution modes).

Franz: It's generally a good whitepaper. I believe it is quite heavily downloaded also.

Gunnar: Please discuss with all your technical experts in the companies about useage of this, since when we start thinking of usage maybe the ideas we then have will impact the API standards work.

Gunnar: I guess we have crypto listed in our table. That's one, but maybe more.

Artem: Yes, and that is our position [...that cryptography should be implemented in trusted execution environment]. Specifically TEE, should run in the TrustZone mode(s).

Gunnar: Going through the table to see where we stand. Are all topics being covered.
...can we bring any to some final conclusion.

... vIOMMU.

Dmitry: some code proposals but for various reasons it seems it will not make it into mainline Linux. Performance is slow. But there are no known use-cases in Automotive?

... CAN

Franz has linked a virtio-can driver

Anup: This is the frontend driver. I believe the backend was implemented in XVisor but was never sent upstream to me.

(more discussion)

Artem: I think we found little need to virtualize. Actual CAN access is typically implemented in another CPU. Perhaps for sniffing / logging purposes [but that's so simple that you don't need a full stack]

Gunnar: Yes, I have also mostly seen designs where there is a separate Vehicle Interface Processor, or at least a separate core on SoC.

Artem / others: The conclusion might be that there is little need to virtualize CAN. *USB might be similar but on the other hand it supports virtualization.*

Gunnar: Sure and this might be the conclusion... I can imagine some chapters [in a virtual platform specification] would just make this conclusion and perhaps point to some reference (in this case virtio-can) if someone feels the need to go beyond that.

Continued general discussion

Lars: We should pick one or two easy ones and not try to reach the answer for each.

Gunnar: Agree, . I'm asking about them here but mostly it is the intention of going through the list, to see where we stand... to find simpler ones to start with.

Anup: I think watchdog is important and also Random Number Generator. Virtio has a proposal for RNG but not watchdog.

Gunnar: We might discuss RNG under "crypto" but it's not the only usage so let's just add it separately. Everyone, feel free to add to the list!

GPU...

Gunnar: I think we need to get Matti and Nikola together to finalize discussion on the feasibility of 3D API standards. For 2D everyone seemed to agree that VIRTIO should work. For 3D, I think there are nuances we need to cover. It's never all or nothing - we should be able to find *some* common parts (API and/or code).

9pfs...

Gunnar / Lars / other discussing. It seems we can wrap it up with the conclusion that we don't see a strong use in Auto/Embedded. Gunnar: I'd be fine with that - we should cover the most common systems. I would write an initial "chapter" on this as an example. But that's mostly a "negative" example [i.e. documenting that it is out of scope]. Now we need also find a positive one, which is needed and where the API standard is decided.

Lars: (For Xen) we only needed it to support running containers. I don't think it plays a part in server virtualization since there are so many other network protocols like NFS (and the VMs communicate between each other using those). As soon as you set up networking, any network filesystem protocol works.

Gunnar: My perception is that [in relation to virtio] this is from VM to hypervisor/host, and that only makes sense in Desktop - VirtualBox/VMWare Workstation, etc. As a standard I imagined NFS would be too big/complicated (to use as API to/from a hypervisor)

AI: All participants asked to:

1. Come to a (personal) proposal for your section and document this (is VIRTIO adequate, what else is needed, etc. The process that is mentioned on [working page](#)
2. If we feel uncertain, e.g. must have more use-cases, write that down. What is required for you to reach the point of 1.

Gunnar adjourned the meeting with the idea that today's discussion was preparing us for getting this done (starting with one or two simple

AI (Anup): Pick a topic to lead. A free one, or you can also add to one that already has a name.

October 11, 2018 - [Tech summit working session](#) in Bangalore, with phone conference

October 9, 2018 - No phone conference because of [tech summit](#)

Tuesday October 2, 2018

Participants

- Bernhard Rill
- Philippe
- Sang-Bum Suh
- Kai Lampka
- Anup Patel
- Artem Mygaiev
- Dmitry Morozov
- Gunnar

Discussing the tech summit

The overall plan for working session is at:

[HVWS Workshop Schedule at Bangalore Tech Summit](#)

A few people asked for a more exact agenda...

The first hour (10.30 CEST, if we have calculated correctly) will be on GPU sharing.

At 11.30 approximately, switch to Security Block with ARM leading. Up to 45 minutes ARM intro/presentation, followed by Q&A / discussion

After that, follow-up topics.

(later in meeting) Kai offered to prepare some presentation on block devices.

Bernhard: What will be the planned for conferencing?

Philippe/Gunnar: Zoom is our assumed default. If we need to change to Skype or Hangout we will let you know. Details will be sent by Philippe.

Bernhard: We would like to have more questions in advance to cover in presentation.

Gunnar: Collecting at the bottom of this page - please add to it!

Artem: We could discuss my questions on implementation of OPTee and support for virtualization ARMv8 and prior...

Virtual Platform Definition:

Crypto

Gunnar: Can we start HSM discussion by evaluating crypto support topic listed in table? Sang-Bum?

Sang-Bum: I have too many other engagements. I will be busy until 20th October at least. Maybe after that.

Block devices

Gunnar: We have discussed various aspect I'd like to bring down to some concrete results. Let's document it. Is VIRTIO good enough or what is the remaining gap?

Kai: Still investigating more

... but I think VIRTIO is not concrete enough. I think it needs more specialized descriptions

...That's why I want to investigate the TRIM command to see how well it works.

Kai: Persistence mgmt (stack) is not fully sorted out. (Not clear where in the software stack you should do what). What about Transaction/Commit semantics for storage? Should that (API) be defined?

Gunnar: Experience from XVisor?

Anup: XVisor provides the VIRTIO block device standard and that's about it.
...It's up to each system to decide
... Where/when can dedicated memory per VM be used (pass-through), and where not?

Kai: For cost reason, one chip per VM is usually not realistic

Anup: Agreed.

General

Anup: The virtual platform specification, will it need to define the exact memory layout [for memory-mapped devices]? QEMU basically does this...

... some discussion

... conclusion that presumably yes this is needed if a VM is going to be fully portable? Let's return to this question.

Next meeting in 2 weeks, (with the Tech Summit working session in between).

September 25, 2018

Minutes TBD

September 18, 2018

Participants:

- Anup Patel (XVisor)
- Nikola Velinov (Green Hills)
- Subramanian (Alpine)
- Artem Mygaiev (EPAM)
- Vasco Fachin (ADIT)
- Gunnar (GENIVI)
- Alex Agizim (EPAM)
- Bernhard Rill (ARM)
- Dmitry Morozov (OpenSynergy)
- Devendra T (KPIT)

Apologies:

- Philippe R
- Kai Lampka
- Matti (Opensynergy)

Agenda:

- XVisor introduction
- Tech Summit
- ~~Sensors (if time)~~

Minutes:

- Some introductions.
 - Devendra joining from KPIT.
 - (later) Alex from EPAM
 - Anup: I work at Qualcomm as a server virtualization engineer. XVisor is a side project.

XVisor introduction:

- Anup: Hobby project at first (~2010). Could we run more than one RTOS on the CPU? Played around with minimal scheduler and were excited when it started to work. It grew into this project.
 - ... Interested in providing the minimal possible scheduling, avoiding complexity
 - ... PowerPC first, then tried also MIPS
 - ...Availability of hardware was a concern. Started with ARM eventually.
 - ...Now we have ports for supporting the virtualization extensions.
- Gunnar: considered Intel support?
- Anup: Yes we have it - another person working on that but not me personally

- Anup: We also have some support for ARM architectures without built-in virtualization features. We patch the kernel, to remove certain "problematic" instructions with a hyper-call. So with some modification we can run RTOS and Linux on hardware w/o virt support. Not as good performance of course but still cool!
- Anup: Most people are familiar with Type-1/Type-2. But the word "Monolithic" might need clarification.
- Anup: Walking through New/extended classification scheme. Details beyond the "Type 1/2" definition.
- Gunnar: Do the rest of you agree with the classification? It would be interesting to see other missing hypervisors categorized as well.
- Anup: We support VIRTIO for paravirt. Can remove drivers if not required, or create hybrid approach where some are pass-through
- Anup (presenting Xen slide)
- Artem: Not quite, backends will be in dom0 and in this case Qemu is not needed. But if you really want to use emulation then emulators will run on the hypervisor level.
- ...There are also peripheral emulators that can run on Xen level. Information in research paper from 2 years ago: running device drivers in Xen as a privileged application.
- ...You can go into privileged mode, effectively switching back to ... directly run on execution level zero. Any type of emulator. Potentially unsafe device drivers can be run here. In most cases if you talk about PV model only then the picture is correct.
- ...Alternative, backend talks to backend with its own device driver model then QEmu not needed.
- Gunnar, discussing working session
- ... XVisor details and Q&A?
- Anup: I could prepare a few interesting questions. Memory handling...
- Gunnar: Security topic. I also think the task scheduling challenge is (forever) interesting.
- ... ARM TrustZone whitepaper
- Bernhard: There are ARM engineers that can introduce this. Just give me a more precise time slot.
- Gunnar: Good. Since we have many virtual participants and the logistics of a large room can be difficult, it's important that we organize the working sessions around a few presentations.
- Nikola: I would also like a more detailed time schedule
- Gunnar: Sure I can set that up, but remember you all need to provide input (topics) to actually create it.
- Gunnar: Are you familiar with [ACRN](#)? We have not done any deep-dive into it on this project yet but I think we should at some point. I think it's Intel only, but also follows a minimalist approach. Might be worth for you to look at.
- Anup mentioning paper(s) behind XVisor. Links will be sent to mailing list.
- ...There is also a paper on TrustZone usage.
- Gunnar: That fits well since ARM presented TrustZone information to us recently.
- Gunnar: If you have additional topics for working session you can make a short intro - send me that proposal!

September 11, 2018

Participants:

- Kai (Elektrobit)
- Nikola (GHS)
- Sang Bum (Perseus)
- Subramanian (Alpine)
- Matti (Opensynergy)
- Christoph (ADIT)
- Vasco Fachin (ADIT)
- Gunnar (GENIVI)
- Philippe (GENIVI)

Agenda:

- Technical Summit sessions (Bangalore, 10-11 October)
- progress on actions: [API standardization](#) / [virtual platform definition](#)

Minutes

- Technical Summit sessions (Bangalore, 10-11 October)
 - calendar invite for the HV working session sent: session is scheduled on Thu 11 October 2pm-4:30pm local time / 10:30am-2pm CEST
- New participant
 - Christoph introduces his colleague Fachin who will take over ADIT participation to the HV project
- Progress on actions: [API standardization](#) / [virtual platform definition](#)
 - gpu sharing

September 4, 2018

Participants:

- Bernhard Rill (ARM)
- Artem (EPAM)
- Dmitry (Opensynergy)
- Kai (Elektrobit)
- Nikola (GHS)
- Sang Bum (Perseus)
- Guru (Bosch)
- Gunnar (GENIVI)
- Philippe (GENIVI)

Apologies:

Agenda:

- Technical Summit sessions (Bangalore, 10-11 October)
- Updates from Dmitry on GPU Sharing ([operation](#))
- Discuss reference received from Bernhard (email)
- "New" (actually old) Hypervisor: [XVisor project](#)
- ~~Discussing the format of the Automotive Virtual Platform specification (GA) (postponed)~~

Minutes

- tech summit
 - discussion on deep-dive topics for discussion in the HV working session: gpu sharing is an important topic to discuss
 - HV working session will be scheduled on Day 2 afternoon (11 October) to allow European participants to join by telco
 - Bernhard has sent an email introducing a white paper released by Arm which gives an insight into the architectural updates in Armv8.4 in the Trustzone:
<https://community.arm.com/processors/b/blog/posts/architecting-more-secure-world-with-isolation-and-virtualization>

August 04, 2018

Participants:

-

August 28, 2018

Participants:

- Bernhard Rill (ARM)
- Artem (EPAM)
- Christoph (ADIT)
- Matti & Dmitry (Opensynergy)
- Kai (Elektrobit)
- Gunnar (GENIVI)
- Philippe (GENIVI)

Apologies:

- Nikola (GHS)

Agenda:

- Bangalore tech summit
- Device standards / VIRTIO continued

Minutes

- tech summit
 - discussion on tech summit content, we envisage doing remote presentations / discussions via video link
 - Artem: will bring some topics for discussion on the table
 - Opensynergy: Matti confirms he cannot join F2F, is asked to provide inputs
 - **TODO All provide your inputs for deep-dive discussion by the end of this week**
 - Gunnar: remind about the working session structuring: quick intro by topic owner and items for discussion
- technical work
 - review of device drivers work items
 - networking: Gunnar - we have an opening for it (i.e. we need someone)
 - discussion on gpu sharing
 - need to support Vulkan and opengl
 - Android space is getting the whole traction
 - Google is sponsoring the Vulkan work

August 14, 2018

Participants:

- Nikola (GHS)
- George Dunlap (Xen Project / Citrix)
- Lars Kurth (Xen Project / Citrix)
- Bernhard Rill (ARM)
- Artem (EPAM)
- Christoph (ADIT)
- Gunnar (GENIVI)
- Sriram (KPIT), second half of meeting

Apologies:

- Philippe, Sang-Bum, ...others

Agenda:

- **Device standards / VIRTIO continued**
- **Bangalore tech summit**

Minutes

Gunnar: Let's Re-introduce the VIRTIO/Device Driver/Virtual Platform definition project because of new/returning participants.

The intention is to write a virtual platform definition that can encompass the whole Automotive Industry. So, supporting Linux & non-Linux operating systems (according to Industry wishes). It would ideally support hypervisors developed with FOSS licenses and other models.

VIRTIO has been proposed as starting point. We're now evaluating each device type / topic:

- What is defined by VIRTIO
- What are the automotive Requirements
- Evaluate applicability and completeness. Clarify the gap.

Gunnar: First study has been on VIRTIO 1.0 but I have seen that there is additional work ongoing. There is a git master...

Lars: Yes, a version 1.1 is planned. You should make sure to cover the latest.

Gunnar: Agreed, action taken to steer our evaluation towards git-repo master - if that's what is most appropriate?

Lars: I assume that's it, but I or George, could look into that.

George shares various experience from Xen project:

...VIRTIO was designed with KVM in mind first

...also for Xen we have found this to be a problem in some areas

...For example, it is assumed that QEMU (which provides the VIRTIO implementation when using KVM) has full access to all of the guest memory all the time.

...it is stated that VIRTIO devices bypass IOMMU completely.

George: In Xen we want to build features that do not match this, such as VM having control over which backends can write into its memory. We have a concept of Driver Domains, which adds security. A layer of security in case of bugs/vulnerability in implementation. For example something like a network card driver may be run in its own VM, with a well defined communication interface to the client VMs that use it.

Lars: Should we write down driver isolation as an automotive requirement?

Bernhard: Also, looking at the list of general considerations, please make sure to add Functional Safety

Artem: ... and Security

Artem: I see comments about implementation dependent things. Isn't the goal for GENIVI to implement standard implementations [that can be used by multiple parties?]

Gunnar: (paraphrased): Yes, this is a likely goal but it remains to be seen how this project progresses. We start with analyzing and defining requirements and specification. However, a specification needs implementation to prove viability. This is GENIVI's experience since the beginning. Previous compliance programs, have always required _some_ software to prove for example an API specification is appropriate, before it becomes part of the specification.

Gunnar: For this project we have received input from, for example Green Hills, that [even independent of the question of porting VMs across different HVs], at implementation and quality maintenance of drivers is a significant effort. So it seems many, including commercial HV vendors would benefit from more shared implementations too if it's feasible.

Nikola: Agreed. We will have to see [how much implementation can be shared] - such as... how much work is required to make VIRTIO implementation have high enough performance?

Comments From Nikola Velinov on the meeting notes: The shared implementation can serve as a good reference for identifying the 'Virtual Platform' requirements on the actual virtualized device. It might not be the best approach to have the platform demand the usage of virtio. Rather it would be better if the 'Virtual Platform' defines requirements on the actual virtualized devices and points to the virtio standard as a reference of such in terms of the guest component. Performance for virtio in this case would not be as critical.

Bernard Rill: Have you [Xen Project] evaluated portability across architectures? ... I mean SW layers etc.

[Discussion to understand how/if such standards are easy or hard to implement in diverse software.]

Nikola: I would also agree. It is clear that VIRTIO came from a non-embedded starting point. Therefore need to figure out if it can be transformed towards better supporting embedded.

(... also some answer from Xen Project)

Gunnar: Interesting and important - please share such experiences by documenting/linking in the Wiki. We need to collect evidence and information to see the full picture. But I would like to steer the conversation back from "is this possible" more towards actually doing the required work now. (looking at table again)

Gunnar: Please volunteer for the topics that have no Champion yet.

Artem: Looking at Sensors... Aren't most sensors just providing an interface using some standard device class, such as serial? They rarely provide any particular HW support, so it's surely para-virt. So it is more about defining a protocol. We have in fact defined some protocols, as part of XenPV work.

Gunnar: That might very well be the conclusion. Seems you have done half the work now 😊 - can you add these thoughts to the Wiki, and then we check consensus later? I'll put your name down on Sensors 😊

Artem is assigned to Sensors. He also *volunteers* EPAM for media acceleration topic.

... Also what about "data-intensive devices". Fast DMA/memory implementation.

Gunnar: I don't know. I guess the IOMMU topic will branch out into a wide discussion (It's all about memory handling).

Lars: I can't volunteer me or George at this time - need to check availability.

Lars: I saw no info about mailing list...

Gunnar: At current we use genivi-projects. But we can set up a dedicated list. What would be the group desire - to have a smaller list for intense discussion in the core group? Because I think to keep others informed, it would just be yet another list for them to subscribe to.

Lars: Genivi-projects should be OK. It won't be too high volume.

Gunnar: OK, I will add clearer info to project page.

Lars: We might be interested in smaller focused meetings around some topics, bringing together for example Matti, George and perhaps Stefano.

Gunnar: No problem, we just arrange the meeting time for particular meetings.

Lars: OK, I will use the mailing list.

Lars: I think (VIRTIO) v1.1 has a deadline close to end of the year. We should check the window of opportunity to affect it.

Sriram: I have joined. I will study the Wiki page and VIRTIO specifications.

Lars: ...will be busy for the next 3 weeks or so. Open Source Summit and other things.

Summary of meeting and housekeeping.
Meeting adjourned.

August 7, 2018

Cancelled due to vacations

July 31, 2018

- Matt Spencer
- Bernhard Rill
- Kai Lampka
- Christian Schulenberg
- Stephen Lawrence
- Philippe

Agenda:

- preparation of the technical summit in Bangalore - 9-10 October
- Focus will be on HV project and Graphics Sharing project
- We intend to have a HV project readout on Day 1 afternoon (duration: 1h30) and a working session on Day 2 (duration: half-day). There will be a networking event on Day 1 evening.
- We need a critical mass of attendees already involved in the HV project to have a productive working session
- We would like to check who intends to attend the event (either personally or through one of your colleagues based in India or elsewhere) and gather topics for discussion for the working session.
- there will be no big showcase at the technical summit (next big showcase will be at CES 2019), some demos might be shown though at the technical summit upon request
- Main topics for the working session will
 - [device drivers standardization](#): it is important to make sure the virtio community is aware of the work done by GENIVI, for the next two weeks, it is expected that assignees (in the table at the bottom of [API standardization / virtual platform definition](#)) will add description of device drivers, then starting mid-August, the HV project should start discussing how to coordinate with upstream projects (please look at Opensynergy presentation at ALS in June)
 - [automotive use cases](#): for the time being we have been only through one use case (rear view camera), analysing a second use case should be at the agenda of the HV project after mid-August so that some results might be shown and discussed at the technical summit
 - **TODO** @Sriram any feedback on this ?

- in addition, participants are welcome to propose topics / problems to solve that could be debated at the working session
- next call is scheduled on Tuesday 14 August

July 24, 2018

- Frank
- Gunnar
- Matti
- Sang-bum
- Philippe

Agenda:

- [API standardization / virtual platform definition](#)
 - Action for all: [Look at list of devices](#) (bottom). Continuation of allocation of device drivers description to participants
 - discussion on criteria
 - Matti: more detailed criteria for applicability of device drivers to automotive are needed, we need to formalize the criteria
 - Matti: the idea for FUSE came from 9pfs, FUSE is a modern linux version of a pseudo file system like 9pfs, it is a very interesting technology
 - Matti: one criterion is : is it standardized already ? is it implemented ?
 - example: CAN device driver is implemented but not standardized yet
 - another criterion is complexity (i.e. complexity estimate of the device implementation: for CAN all you need is a kernel implementing the CAN device class
- next two weeks will be dedicated to filling the wiki page [API standardization / virtual platform definition](#) with content

July 17, 2018

- Frank
- Nikola
- Kai
- Philippe
- Gunnar
- Ralph
- Sang-bum

Apologies:

- Sriram Gopalan
- ...

Minutes

- Additional discussion about [Virtual Platform Definition](#)
- Kai has previously been "assigned as a volunteer" 😊 for block storage. And it was confirmed.
- Kai: Can we write down more than one person, a team of interested people?
- Gunnar: We can add more, but in my experience we need a topic lead to move things forward. The lead can of course be involved in "recruiting" more people also...
- Gunnar asked Nikola, who selected Network devices & Ethernet (+AVB)
- Gunnar asked Frank, who accepted to look into USB - will discuss with other SysGo engineers.
- Gunnar asked Ralph.... Ralph will ask Matti to select some topics.
- Gunnar asked Sang-bum to cover some topic(s) based on his experience. Sang-bum does not have time right now to evaluate a topic.
- Everyone needs to do the evaluation of [their selected topic](#)
 1. Write down / figure out the automotive requirements
 2. Read VIRTIO chapter
 3. Decide if VIRTIO is appropriate and complete for requirements (Gap Analysis)
 4. Write down steps to close any gap
- **Action(All):** [Assign yourself as "lead" for 1-2 topics and write it on the page. Then perform the steps above.](#)

July 03, 2018

Participants:

- Christoph Lipka
- Guru R
- Sang bum Suh
- Sriram
- Nikola
- Gunnar
- Philippe

Agenda

- *technical summit*, 10-11 October, Bangalore, India
- Milestones, deliverables, and workplan.
- API standardization – [VIRTIO intro!](#)

Minutes

- Plan for **GENIVI Tech Summit in Bangalore**
 - 10-11 October, Bangalore, India
- HV track: Main setup.
 - 1) Present selected challenges and info from input material from HV workshop at AMM
 - 2) Present results so far from workstreams - Use-case/RequirementArchitecture, and API standards
 - 3) Planned future work in workstreams.
 - 4) Collaborative session / workshop (new participants)
- Gunnar: On 1), there were a lot of really good and thoughtful challenges discussed at the workshop. We should inform our Tech Summit audience (many new people) about the best parts.
- How many participants are expected?
- PR: This is the first of the Tech Summit format, previously AMMs. Working target is around 300 people.
- Where can I find out more?
- PR: It was announced in the most recent Member's Newsletter. You can expect a Wiki page to appear soon.

Text below taken from the newsletter:

GENIVI Announces Schedule for Fall Technical Summit in India

Many will remember that in 2018, GENIVI moved from a two member meeting per year model to a single, large member meeting in the spring and 1-2 more regional technical summits in the fall. The details for one of those summits are nearing completion and GENIVI wants to get this important event into your calendars immediately.

On 10-11 October, GENIVI will hold a technical summit in Bangalore, India. The summit will expand on two active projects within the vehicle domain interaction strategy, notably Graphics Sharing and Hypervisors. The agenda will be finalized during coming weeks; however, GENIVI has in mind three primary goals:

** **Provide an overview of the GENIVI Alliance**, its projects, and recent deliverables, **to an audience** that may have not been able to attend recent member meetings*

** **To inform and engage a technical audience** in the work of our domain interaction projects starting with **Distributed Graphics** and **Hypervisors***

** **To equip developers with hands-on experience** using APIs, reference code and **supporting documentation** so that they can produce software that delivers **solutions needed for domain interaction challenges**.*

*The summit will be held at the **Sheraton Grand Hotel** at Brigade Gateway and will begin at **9:00 am on 10 October** and **end at 4:00 pm on 11 October**. A **networking reception** will be held at the end of the first day.*

***Registration** for the event will open on **1 August**.*

*Please consider attending this important technical event and should your organization be interested in **sponsoring the event**, please contact **Karin Hanson**, **GENIVI Event Manager** for more information on opportunities available*

- AI:(All): Work to secure participation in the tech summit (discuss within your companies).

virtio

- **Matti**: presents the wiki page he created. Some progress at the ALS Summit in Japan where Ralf made a presentation. (As originally discussed in GENIVI workshop, and thereafter), the idea is to have a platform runtime for the virtio guests that are OS independent, this will enable customers to move the guests from platform to platform (ie a vehicle ecu to a cloud environment)
- ..discussion on AGL way of positioning their work on HV
- **Sang bum**: imho AGL project is trying to create to an equivalent of a Android platform, (because it was the goal of Tizen also)
- **Gunnar**: Perhaps. A Linux distribution, I would call it.
- **Sang bum**: GENIVI can rather provide a framework and architectural vision discussion on standardization work, where AGL can be one possible Linux system.
- **Gunnar**: AGL is one possible Linux system. As such it can be run in what we will describe.
- **Gunnar**: (I don't understand the relationship between AGL project and proprietary hypervisors and proprietary operating systems. At least now GENIVI is covering the scope of interaction between whatever the industry wants to run, including possibly non-Linux operating systems, and commercial hypervisors. It seems an appropriate scope of work for GENIVI. The defined virtual platform should allow for any such combination (that the industry needs).

...

Process improvement

- having a dedicated HV project mailing list ?

Planned absences

26-7/13/08 christoph

24/7-31/7 nikola not available, same for 7/8 TBC

end of September - sang bum

15/7-21/7 Guru

absence a few weeks, probably some time in July-Aug Gunnar

9/7-13/7 & 1/8-15/8 Philippe

no upcoming holiday before mid-October Matti

June 26, 2018

Participants:

- Christoph
- Sang bum
- Sriram
- Nikola
- Gunnar
- Philippe

Minutes

- Discussion on boot times...
 - Just for reference: [1 second Linux Boot example \(old\)](#) (not snapshot or hibernation, if I recall correctly)
"Real-world" might be different.
- Sang-Bum: <100ms hypervisor load is generally possible. What is the requirement?
- /TODO/ Add Sriram's notes
- [architecture](#) proposed by Sriram at the bottom of this page is aligned with the rear camera use case
- discussion on EVS (External Video System) on top of android
- discussion on Green Hills solution
- /TODO/ Nikola upload the diagram with a micro-kernel based RTOS architecture (with HV as a thin layer on top of RTOS)
- RTOS is responsible for the scheduling, HV is there to enable Linux
- Nikola: goto GHS website the GHS solution is fully described there
- Nikola: you do not need to boot the whole linux to get the rearview camera up & running, you need to boot only the rtos part
- Sang bum: would you require two video streams, one coming from rtos and the later another one coming from android or linux ?
- discussion on the handover of video streams

Sriram's notes - minor edits and formatting

(Camera use-case an architecture)

- Birds Eye View / Surround View could be possibly only implemented in the Android Partition and may not be available as part of early RVC
- Following are options for supporting such a use case through out the lifetime of a IVI session (from startup to shutdown)
 - Everything in Vehicle domain (Basic as well as advanced view) => No handover / arbitration required with any other domain
 - Basic implementation in Vehicle domain + Advanced implementation in Linux/Android domain => Handover / arbitration of camera stream is required
 - Everything in Linux/Android domain (Basic as well as advanced) =>
 - Dedicated partition (potentially a third domain besides Vehicle and Android/Linux domain) that does everything
- What are the implications for Hypervisor?
 - Constraint as of today in Android : Cannot do early RVC within 2 seconds.. however that is on the roadmap
 - Linux does not have this constraint as it is completely Open Architecture and free from Compatibility aspects
- Solutions for early camera systems
 - Snapshot / Hibernation speeds up boot and production solutions with < 2 seconds camera is possible
 - Example from a camera product(high end camera)
 - No safety aspects in consideration in the above example; non-automotive
- Load time of hypervisor ~ 100 ms

June 19, 2018

Agenda:

- VIRTIO intro
- Use Case Architecture Problem Statement Results

Participants:

- Nikola
- Sriram
- Subramanian
- Christoph
- Gunnar
- ...

Apologies:

- Sang-Bum
- Matti
- Horst

Minutes

- The need for a shared vocabulary again
- Nikola: API standards like VIRTIO should apply to Type1/Type2 equally
 - Link [VIRTIO 1.0](#)
- Is libvirt implementation a useful addition to the plain specification? - is it really widely used?
- Discussing project scope & results
 - ...
- Plan: Discuss use case per meeting
- Discussing rear-view camera case once more
 - E.g. long-booting kernel require another solution for camera
 - Common solution: Bootloader early-boot program runs camera later on there is some type of hand-over
 - Gunnar: As discussed last week, a small core (on SoC) usually handle the wake-up scenario. That in itself is not requiring a HV. But what runs next is what we should discuss.
 - Nikola: A Type 2 HV can run processes (e.g. camera) separate from a large kernel. It can continue running that in fact, so no hand-over required.
 - Also possible with Type-1 (which is assumed to be small), which runs a small partition for camera
 - Safety requirements.
 - Android have now published some of their early camera proposals link?

June 12, 2018

Agenda:

- getting everybody updated
- review of critical use cases.

Participants:

- Kai
- Sang bum
- Bernhard (arm)
- Kevin
- Nikola
- Sriram
- Matti
- Franz
- Gunnar
- Philippe

Apologies:

- Christoph Lipka (ADIT), Albert (Conti)

Minutes

- question on cpu usage following a discussion with a car OEM
 - Sang bum: reports on a discussion on cpu consumption he had with an OEM
 - Gunnar: depends on types of ECU : ic or infotainment
 - Matti: most of systems are very much loaded because as soon as a system has spare capacity OEMs swap to a smaller chip because it costs less
- device driver virtualization
 - wiki page: <https://at.projects.genivi.org/wiki/x/ZpT0>
 - Matti: posted one hour ago a short intro on virtio in the above page
 - report on review comments on the wiki page postponed to next week
- use cases discussion
 - Sang bum: starts with a question about the use of Android os environment
 - Gunnar: Android has been used for several years w/o support asked from google and w/o certification in mind, now some oems are engaging more closely with Google and look for approval
 - Matti: agrees with Gunnar, Google is heavily pushing their techno into automotive and they will likely agree with the utilization of hypervisors
 - Sang bum: asked whether there is a need to dynamically load an OS (i.e. load android when you want to run an android application)
 - Sriram: this is an interesting requirement
 - Franz: multiplying OSes make things more complex when executing the project (3/4 parties need to manage the project which makes it difficult)
 - Matti: this cannot happen with Apple iOS because Apple has never offered iOS on an hw not designed by them
 - Matti: all other OSes could be loaded on demand, Windows might become something people would like to have, a rear set entertainment is not really automotive, it is rather a tablet, this is why Windows could come back in the landscape
 - [not captured for a couple of minutes]
- back on use cases
 - link: <https://at.projects.genivi.org/wiki/x/WpP0>

- Nikola: comments on the challenges column of the table, some challenges are relevant to architecture rather HV technologies, would it be worth identifying them better ?
- Sriram: yes, this would help
- Sriram: shows diagram on architecture at the bottom of the page
- Sang bum: how can SoCrun both linux and rtos w/o virtualization ?
- Sriram: this is because modern silicon offered the ability to specialize cores [capture is TBC], almost all infotainment systems were following this architecture up to now, but this is changing
- Sriram will share more diagrams
- AOB
- Technical summit - Fall 2018 (week starting on October 8, location: Bangalore, India)
 - Philippe: informs the group that the topics selected for discussion in the summit are related to graphics sharing and hypervisors
 - Sriram: coins the idea of inviting xviser, an HV open source project whose maintainer is based in Bangalore, can reach out to him for a possible presentation, link: <http://xhypervisor.org/>
 - Franz: what are the results we need to talk about ?
 - Gunnar: [not captured]
 - Kai: what would be the prerequisites to deliver a demo there ?
 - /TODO/ Philippe gather and provide info about the session schedule and space for demos at the Technical Summit
- Adjourned: 11:05am CET
- next week
 - review of Opensynergy inputs on virtual device drivers

June 5, 2018

Agenda:

- getting everybody updated
- review of critical use cases.

Participants:

- Albert Kos (Conti)
- Kai Lampka (EB)
- Christoph Lipka (ADIT)
- Franz Walkembach (Sysgo)
- Sriram (KPIT)
- Sang bum Suh (Perseus)
- Gunnar (GENIVI)
- Philippe (GENIVI)

apologies: Matti (Opensynergy)

Minutes

- getting everybody updated: we have currently 3 threads active
 - device driver virtualization (prepared by Matti)
 - AGL paper review (prepared by Nikola)
 - critical use cases (prepared by Sriram)
- device driver virtualization
 - Matti cannot attend today's call, got the clearance to provide his inputs, will upload them in the wiki today, topic for next week
- critical use cases
 - Sriram presents the [critical use cases wiki page](#) and his views on how APIs can be identified allocated in an instrument cluster context
 - Sriram: this is WIP, will add block diagrams
 - Sriram: details the assumptions and the use cases
 - Gunnar: we need to check how these use cases fit with the domain interacti
- more on use case / scenario #1 - Rear view Camera at Startup
 - Sriram: the rear view camera can be Ethernet connected, or LVDS connected, rather than analog cameras (converted to digital)
 - you still need to compose the camera output with the park assistant
 - it is a data stream available on a certain port
 - Ethernet based cameras would be a good use case to start with
- discussion
 - Sang-bum: AGL paper talks about resource allocation on top of hypervisors (linux VM, Android VM)
 - Albert: we need to make a distinction between various partitions (safe, secure, others), it is a very complex landscape, IMHO that cannot be solved with one HV, first of all we have to focus on the safety critical part (Autosar based) based on type 1 HV and look at type 2 HV for other applications (e.g. linux containers, etc.)
 - Gunnar: we are not at this step of describing a solution using type 1 and/or type 2
 - Sang bum: we need to make a decision on our scope type 1 only and/or upper layers with sw virtualization
 - **/TODO/ Albert provide a problem statement for next week (for instance on the rear view camera)**
 - Kai: IMHO in order to be concrete with safety, we need to address ASIL levels which is a far-reaching goal, we should first focus on requirements like boot time and data throughput and latency bounds which current HV technologies have problem to comply with
 - Kai: most OEMS are a bit hesitating because they are not convinced HV will meet performance requirements
 - Sang bum: what about VMs on top of an HV ?
 - Kai: first need is to run legacy code and to have different sw islands to run them
 - Sang bum: IMHO para-virtualization is included in the scope of solutions
 - Sang bum: state of the art is that silicon supports virtualization, we do not need para-virtualization
 - Sang bum: in GENIVI we could runLMBench on linux on top of HVs and make performance measurements
 - discussion on performance benchmarking (for audio for instance)
 - Gunnar: this benchmarking is an industry effort, would start this by having this group agree on the method to do the measures
 - Albert: not sure, we did this benchmarking already

- Sriram: I will add KPI to the use cases I have provided, we need to treat each problem statement, we have the right group to do the job
- Sang bum: Android ? running on the fly of Android applications is another use case, using this approach OEMs could avoid Google certification
- Gunnar: asks every participant to provide his views in the wiki
- **/TODO/ All provide inputs on the various problem statements in the wiki page**
- next week
 - review of Opensynergy inputs on virtual device drivers

May 29, 2018

Agenda:

- review of the [summary of AGL paper](#) prepared by Nikola .

Participants:

- Nikola Velinov (GHS)
- Sriram Gopalan (KPIT)
- Christoph Lipka (ADIT)
- Philippe (GENIVI)
- Gunnar (GENIVI)
- Sang-Bum Suh (Perseus)

Minutes

review of the [summary of AGL paper](#) prepared by Nikola
summary is short, everybody invited to read it

sections 3, 5 & 6 are the most relevant for HV project

section 3 can be adopted as a set of reference use cases

discussion on how we (as a GENIVI project) can build on the inputs from the AGL paper

Certain chapter should be reusable as they are, others with some modification. Some chapters have useful content which should be quite widely applicable (multiple Linux, Yocto-based systems etc.) but the text cannot be used as-is since it uses a lot of specific language referring to "the AGL system", and similar expressions.

next week

- more feedback from other participants who did not attend today's call
- Opensynergy inputs on virtual device drivers (if internal go was given)

May 22, 2018

Participants:

- Kai Lampka (Elektrobit)
- Gayathri PP (Tata Elxsi)
- Matti Möll (OpenSynergy)
- Nikola Velinov (GHS)
- Sriram Gopalan (KPIT)
- Philippe (GENIVI)
- Gunnar (GENIVI)

Apologies

- Sang-Bum Suh (Perseus)

Minutes

API standardization

Matti started writing down thoughts. Needs some approval on content

Matti introducing:

- Definition of the virtual platform...
- virtio as specified, and/or looking at defacto implementations e.g. in Linux
- there are feature flags in the virtio spec - define which ones should be mandatory for automotive use case
- Then guests can be developed against this virtual platform – possible to run guest on any HV that fulfils the platform
- It's like "Virtual appliances" used to be a popular concept. Now containers/applications-containers, same principle.
 - Package certain software as a machine image. Deploy to cloud.
 - Before virtualization "appliance" is a box of where the software is running, then virtual-appl.
- This work in GENIVI could specify such a baseline environment that describes the set of devices.
- Can be done by referencing existing standards.
- Agreeing on the subset of features that automotive requires - what is mandatory instead of optional.
- Example: The discard capability of the block device. Is optional in virtio, but for embedded, this should be mandatory.
- Discard = Hinting to a flash device that something is not being used anymore.

Discussion:

- Nikola: Virtio standardizes a restricted set of devices only, e.g. GPU(?), what about the devices that are still in draft status.
- Matti: One option is to drive the virtio standards forward (to fill those gaps)
... the other is to just say we shall be compatible with current implementation (in Linux)
- Nikola: agrees with Matti's proposal
- Nikola: I'm interested in how to ensure compatibility/compliance to what we define.
- Gunnar: Driving virtio from draft to specific ("work upstream"), specify gaps, and promote that we support a particular defacto definition. All of those activities help drive standards forward.
- Matti: Plugfests is an option, such as has been done for Bluetooth. (for ensuring compatibility, and also driving standards forward)
- Kai: How much functionality should go into (virtio) driver vs (low-level) device driver part.
- Kai: How to handle non-open implementations (hardware specific low-level driver) etc?
- Matti: VirtIO should abstract those details typically. We should help to define that.
- Kai: HVs allow for Central device driver management which is useful, it can enforce certain policies.
- Matti: I agree. Example: block-write rate limiting
- Sriram: What about Type 1 vs Type 2. For example using KVM, having a full Linux kernel can be very useful. Type 2 approach enables more possibilities that bare metal implementation. Also look at Xen. Why don't we look at those. Maybe also go back to requirements first.
- Gunnar: All of those technologies are also represented in our group, (and some might even consider themselves to be their own type). [Thin layer HV and RTOS based, etc.]
Virtual Open Systems work on KVM for example, and others are strong on Xen. I don't right now see a reason to limit to one type.
- Several: Requirements can help clarify what is needed.
- Gunnar: I think we should drive those things forward also. I called it a separate workstream but I don't want to create boundaries. I think these things are very related. The only thing is, we won't stop one track (API standards) that we already started, in order to go back to the foundation. I think we can progress this in parallel now.
- Sriram: I have some use cases I'm thinking about, I could write them down.

Additional thoughts from participants.

Wrapping up

Action Items

- Nikola study and summarize AGL whitepaper, in particular looking for basis for defining "the virtual platform", and of course requirements, use-cases, etc. Use Wiki page to summarize.
- Sriram - Write down cases / concrete architecture examples / specific challenge (such as, audio startup&latency in a particular architecture)

May 15, 2018

No posted minutes

May 8, 2018

Minutes

Discussion on – see the title Device Standardization on main page: [Hypervisor Project](#)

Sang-Bum: Hypervisors need to include a mandatory access control features

Matti: But in theory guests can run without ever speaking to a hypervisor.

Matti: It is difficult to standardize APIs to speak to the hypervisor itself - easier to standardize device driver layer.

Sang-Bum: We need to add a security architecture to control (negative) impact from one guest to another. We need MAC support APIs to achieve that.

Matti: I would like to start the standardization topic by writing down a proposal.

April 19, 2018 (Full-day All Member Meeting Workshop)

Please see the [Hypervisor Workshop Schedule at Munich AMM 2018](#) page for schedule, speakers, participants and meeting minutes.

April 10, 2018

Further preparations of AMM agenda

April 3, 2018

Participants

- Ajmal (Tata Elxsi)
- Artem Mygaiev (EPAM)
- Gayathri PP (Tata Elxsi)
- Ralph Sasse (Opensynergy)

- Guru (Bosch)
- Christian (BMW)
- Nikola (Green Hills Software)
- Franz (Sysgo)
- Gunnar Andersson (GENIVI)
- Philippe Robin (GENIVI)

Apologies

- Sang-bum Suh (Perseus)

Minutes

- Notice / presentation of new participants: Nikola (Green Hills Software). And previous participants now with opportunity to present themselves: Ajmal and Gayathri PP (TataElxsi)
- *Continuation of the assignment of topics preparation with the participants, the topics and assignments for the workshop are listed under [Hypervisor Project](#).*

March 27, 2018

Participants

- Ajmal (Tata Elxsi)
- Alex Agizim (EPAM)
- Artem Mygaiev (EPAM)
- Gayathri PP (Tata Elxsi)
- Horst Saier (Mentor)
- Matti Möll (OpenSynergy)
- Raj
- Sean Park (LGE)
- jithin (Tata Elxsi) ?
- Subramanian (Alpine)
- Stephen Lawrence (Renesas)
- Gunnar Andersson (GENIVI)
- Philippe Robin (GENIVI)

Apologies

- Sang-bum Suh (Perseus)

Minutes

- Gunnar tried to get contact with TataElxsi participants but still no audio coming through (microphone not working)
- *Continuation of the assignment of topics preparation with the participants, the topics and assignments for the workshop are listed under [Hypervisor Project](#).*

March 20, 2018

Participants

Philippe Robin (GENIVI)
 Sang-bum Suh (Perseus)
 Matti Möll (Opensynergy)
 jithin (TataElxsi)
 Gunnar Andersson (GENIVI)
 Christian Schulenberg (BMW)
 Subramanian (Alpine)
 Gayathri PP (Tata Elxsi)
 Stephen Lawrence (Renesas)
 Ajmal (Tata Elxsi)

Minutes

- Gunnar tried to get contact with TataElxsi participants but still no audio coming through (microphone not working).
- *Gunnar reviewed the assignment of topics preparation with the participants, the topics and first assignments for the workshop are listed under [Hypervisor Project](#).*

March 12, 2018

Participants

Philippe Robin (GENIVI)
 Gunnar Andersson (GENIVI)
 Sang-bum Suh (Perseus)
 Christian Schulenberg (BMW)
 Horst Saier (Mentor)

Subramanian (Alpine)
Guru (Bosch)

Minutes

Gunnar highlighted some of the topics for the workshop listed under [Hypervisor Project](#).

Sang bum: introduced the workshop to LGe, Hyundai and Access in the recently hold Korea REG F2F, would like to collect their opinion so that we can share at the workshop
trying to contact xen so that they give a presentation at the workshop on their automotive projects, intends to contact redhat with is leading virtio
Sang bum: contact with car oem and tiers 1, my personal opinion is they do not know yet what it is the exact case to usefully apply HVs to a vehicle, in the process of trying to convince car OEMs to deliver market scenarios, coins the idea of sending a questionnaire to car oems
Sang bum will share an initial questionnaire with us at the next meeting (20 March)
Horst: my interest lies rather in graphics sharing
Gunnar: graphics will be one of the topics of the workshop
Horst: how to share graphic buffers, is a solution available in the open ? it is currently very silicon vendor dependent
Gunnar: [Horst Saier](#) can you a short intro in the workshop about it ?
Gunnar: @sang bum: are you familiar with gpu sharing ?
Sang bum: yes, I am very familiar, the problem is that silicon vendors except Intel do not publish the code of the drivers for gpu sharing
short discussion on audio virtualization
Sang-bum: would like to discuss device driver architecture at the workshop
Sang-bum: will propose a list of topics for Wed 14 March EOB
Christian: we are very interested in the market survey and what is available from vendors

March 6, 2018

Participants

- Guru (Bosch)
- Albert (Continental)
- Christoph (ADIT)
- Christian (BMW)
- Gunnar (GENIVI)
- ... maybe someone else - not sure

Minutes

We simply discussed and filled in the topics under [Hypervisor Project](#). Discussion much driven by Albert and Christoph.

February 27, 2018

Participants

- Albert / Continental
- C. Gouma / Sysgo
- Fabien H / Valeo
- Gayathri PP
- George
- Karthikeyan R
- Marco R / Mentor
- Matti Möll / Opensynergy
- Philippe / GENIVI
- sbsuh
- Stephen L / Renesas
- Swaminathan G
- Gunnar / GENIVI

Minutes

- Introductions
 - Gunnar tried to get contact with Ajmal, Gayathri, Karthikeyan, George, and some others but no answer.
 - Apparently there were microphone troubles all around
- Purpose/idea and organization history of Hypervisor initiative
 - Seoul AMM activities October 2017
 - Discussion & Preparing
 - "Do it right" - i.e. start only when/if significant interest
 - New developments – new lead (Sang-Bum)
 - Workshop at AMM lots of interest - time to start planning agenda
- Started the topic list under [Hypervisor Project - minutes](#)
- Requirements need - OEMs usually don't give HV related requirements - they are more functional.
- Could we create some kind of (shared) Baseline requirements for HV vendors?
- (Sysgo): I do not believe in any kind of open source solution...
 - ... to solve Level 2 [or higher] autonomous driving...
 - ... the code base is too big
- Gunnar: ... making a note that code base size is independent of code license
- (someone): We need to know which devices are important (for OEMs)

- (someone): We expect safety requirements to be provided
- Albert:
 - Need to listen to the voice of customers
 - case studies
 - safety issues
 - 26262 - many are not familiar
 - difference in needs between IVI (not safety critical) vs cluster (is safety critical)
- accelerate towards (some kind of) safety approval
- HV solutions have been deployed (to solve this) not only in automotive - in aviation at least 5 years ago

an white paper released by Arm which gives and insight into the architectural updates in Armv8.4 in the Trustzone:
<https://community.arm.com/processors/b/blog/posts/architecting-more-secure-world-with-isolation-and-virtualization>

- Gunnar wants to discuss platform independent