Minutes

From the agenda:

We don't have any design and problem related items this meeting. This means that we will only cover discussions specific to some series. Note that the meeting will probably not be very interesting for people whose series are not on the agenda. Feel free to join and observe the meeting, but it's also OK to drop out.

For series on the agenda: we will only discuss your series if the originator is on the call. For each series, I will call out the owner: if the owner is not there, I will move to the next one.

Intel has sent me an updated list based on their priorities. I pushed items which have no issues down the priority list. Also, I tried to order based on priority and vendor.

Formatting notes:

Everything discussed in the meeting is marked in blue. Meta information such as attendees and the original text from the agenda are black. Actions are marked with **ACTION**. The following people have actions on themselves.

- Lars Kurth (Citrix)
- John Ji (Intel)
- Paul Durrant (Citrix)
- Roger Pau Monne, George Dunlap (Citrix)
- Janakarajan Natarajan (AMD)
- Haozhong Zhang (intel)

Attendees

- Lars Kurth (Citrix)
- Janakarajan Natarajan (AMD)
- Daniel Kiper (Oracle)
- Juergen Gross, Jan Beulich (Suse)
- Christopher Clark (OpenXT)
- John Ji, Chao Peng, Haozhong Zhang, Yi Zhang, Chao Gao, Boqun Feng, Luwei Kang, Yu Zhang (Intel)
- George Dunlap, Wei Lui, Roger Pau Monne, Andrew Cooper (Citrix)

Covered Agenda

- Quick round the table: name, company
- General Items
- Series for 4.11
- Other Series with Issues
- AOB

Not covered due to lack of time

- Other Series with no Technical Issues which had no review - not covered
- Other Series - Progressing or Waiting (we will probably not get to these) - just there for reference, not covered

General Items: RFCs

Jan: Generally reviewers prioritize RFCs lower than other non-RFC patch series. Jan's view is that complex series and series above a certain version number should not be marked as RFCs.

For 4.11

[PATCH v4 00/10] x86: emulator enhancements

Sent in for meeting agenda by George
https://marc.info/?l=xen-devel&m=151982229407799
https://xen.markmail.org/thread/roukz6r3gcuhxinn

Notes: v4 posted by Jan Beulich on 28 Feb 2018. Most patches seem to have acks or r-bs, but I know this one has been around a long time, so it might be worth making sure we can get it in before the feature freeze.

Subject	AB/RB	Review
[PATCH v4 01/20] x86emul: extend vbroadcasts{s, d} to AVX2	**AC**	
[PATCH v4 02/20] x86emul: support most remaining AVX2 insns	**AC**	
[PATCH v4 03/20] x86emul: support AVX2 gather insns	**AC**	
[PATCH v4 04/20] x86emul: support XOP insns	**AC**	
[PATCH v4 05/20] x86emul: support 3DNow! insns		AC, Minor
[PATCH v4 06/20] x86emul: place test blobs in executable section	**AC**, PD	AC, Minor
[PATCH v4 07/20] x86: move and rename XSTATE_*	**AC**, PD	
[PATCH v4 08/20] x86emul: abstract out XCRn accesses	PD	AC, Dispute
[PATCH v4 09/20] x86emul: adjust_bnd() should check XCR0	AC	
[PATCH v4 10/20] x86emul: make all FPU emulation use the stub	AC	
[PATCH v4 11/20] x86/HVM: eliminate custom #MF/#XM handling	**AC**	

[PATCH v4 12/20] x86emul: support SWAPGS	AC	
[PATCH v4 13/20] x86emul: tell cmpxchg hook whether LOCK is in effect	AC, PD	
[PATCH v4 14/20] x86/PV: convert page table emulation code from paddr_t to intpte_t	**AC**	
[PATCH v4 15/20] x86emul: correctly handle CMPXCHG* comparison failures	**AC**, **TD**	
[PATCH v4 16/20] x86emul: add read-modify-write hook		None
[PATCH v4 17/20] x86/HVM: do actual CMPXCHG in hvmemul_cmpxchg()	PD	Probably needs AC
[PATCH v4 18/20] x86/HVM: make use of new read-modify-write emulator hook	AC, PD	
[PATCH v4 19/20] x86/shadow: fully move unmap-dest into common code	AC	
[PATCH v4 20/20] x86/shadow: fold sh_x86_emulate_{write, cmpxchg}() into their only callers	AC	

All agreed that we should get this series into 4.11

[PATCH v4 08/20] x86emul: abstract out XCRn accesses
Jan and Andrew have been discussing a way forward during the day with Andrew: Jan can agree to the latest proposal, then submit v5.

[PATCH v4 16/20] x86emul: add read-modify-write hook:
[PATCH v4 17/20] x86/HVM: do actual CMPXCHG in hvmemul_cmpxchg():
Andy only skimmed these. Will look at these in detail once in detail once the patch has been re-based for patch 8

George/Andy: Need to do an AFL fuzzer run after RC1 as this used to be a problem area for XSAs.

George was working on AFL fuzzer improvements which are ongoing, but this should all be in place for RC1

[PATCH v17 00/11] x86: guest resource mapping

Sent in for meeting agenda by George
https://xen.markmail.org/thread/ge2hlgljac3uqepe

v17 posted by Paul Durrant on 3 January 2018
This series is a prerequisite for “[RFC Patch v4 0/8] Extend resources to support more vcpus in single VM”

Notes: All but 6/11 have a fair amount of A-b's or R-b’s

Subject	AC / RB	Comments
[PATCH v17 01/11] x86/hvm/ioreq: maintain an array of ioreq servers rather than a list	RPM, JB	
[PATCH v17 02/11] x86/hvm/ioreq: simplify code and use consistent naming	RPM, WL, **JB**	
[PATCH v17 03/11] x86/hvm/ioreq: use gfn_t in struct hvm_ioreq_page	RPM, WL, **JB**	
[PATCH v17 04/11] x86/hvm/ioreq: defer mapping gfns until they are actually requested	RPM, **WL**, JB	
[PATCH v17 05/11] x86/mm: add HYPERVISOR_memory_op to acquire guest resources	JB, **DDG**	
[PATCH v17 06/11] x86/hvm/ioreq: add a new mappable resource type...		**JB - not sure of status**
[PATCH v17 07/11] x86/mm: add an extra command to HYPERVISOR_mmu_update...	JB	
[PATCH v17 08/11] tools/libxenforeignmemory: add support for resource mapping	RPM, **WL**	
[PATCH v17 09/11] tools/libxenforeignmemory: reduce xenforeignmemory_restrict code footprint	RPM, **WL**	
[PATCH v17 10/11] common: add a new mappable resource type: XENMEM_resource_grant_table	JB	
[PATCH v17 11/11] tools/libxenctrl: use new xenforeignmemory API to seed grant table	**Marek**, **WL**, RPM	

All agreed that we should get this series into 4.11

Paul has to verify whether new series work after rebasing
Not blocked on anyone: No actions on anyone but Paul as far as aware

For those not aware: patch 6 - has caused 2 XSAs, which is why this is getting extra review

Longer Term - Issues

[RFC XEN PATCH v4 00/41] Add vNVDIMM support to HVM domains

Sent in for meeting agenda by George
https://marc.info/?l=xen-devel&m=151264150712808
https://xen.markmail.org/thread/6uzmarrlws73mq5d

RFC posted by Haozhong Zhang on 7 December 2017. A few messages about the overall architecture; some more detailed comments by Anthony on the integration with the toolstack. Otherwise feedback by Roger & Jan.

Issues: Lack reviews for memory management part

Zhang: Would like to get review on Memory Management (which would be George)

Andy: This series has many non-trivial changes in many different areas
E.g. how can we avoid overwriting data when it is mapped map and accessed NVDIMM without clobbering data

Lars:
Besides that: what are the open issues?
Can we divide into smaller pieces or logical areas?

Andy: There are 2 large complicated areas

- Have to come up with a new way of managing NVDIMMs in the Hypervisor (majority of the complicated work) … We have to get that right
- How do we give NVDIMMS to the guest: open question whether we should rely on QEMU or not

One suggestion was to split the series into two halves: we had a discussion about this but didn't get to a resolution due to the complexities. But we all agreed that we need to break the discussion into logical chunks to be able to move this series forward and focus on highlighting the unresolved issues.

Juergen pointed out that there is no up-to-date design for the series. There was an early version, but the code and design are out-of-sync. It would make sense to do something similar as we did for Migration stream v2.

George also agreed that he is struggling to review the series, due to the outdated design.

Andy: The guest interface should be straightforward. The bigger issue is the interface between Xen and Dom0.

Juergen: we will need a design document as a basis for an interface anyway, including the Xen - Dom0 interface.

Zhang: The first patch had a design document. However, every version changed the design a bit, but I didn't update the design document.

We discussed a little whether it would not be to much of a burden to do this, but Intel offered to update the design and include it into the next iteration of the series.

We also discussed whether it would make sense for someone not Jan or Andy to help with this series to make sure it moves forward. Royger (as he is already reviewing the series) with help from George can pick this up (George will need to review the memory management parts). Juergen may also get involved.

Wei: Need to get general understanding on the architecture. Wei highlighted that the responsibility of the developer to drive the conversation. We have a few people to help out resolve architectural questions.

Lars: It would help if we had someone from a team based in Europe to help drive this, as this will help with timezone issues.

Plan forward
Royger to work with Zhang: write down the updated design first. Then resolve the difficult outstanding issues either by mail or if this doesn't work in a meeting.
ACTION: Haozhong Zhang to update the design doc and include it into the next version of the series (1st patch of series).
ACTION: Haozhong Zhang to drop the RFC and CC George and Roger
ACTION: Royger will help and give feedback. George will also be involved as he needs to review the memory side of the series. He will
ACTION: If needed - we can set up a meeting between Zhang and other stakeholders. Lars and John to take over an admin role to make sure developers can focus on the substance.

[PATCH RFC 00/10] x86 passthrough code cleanup

Sent in for meeting agenda by Wei

https://lists.xenproject.org/archives/html/xen-devel/2018-02/msg01939.html

Wei wanted to get the maintainers opinions on what is required make passthrough code cleaner.

Wants to get feedback from AMD to see whether the clean-up as proposed is going into the right direction. Kevin has responded: but has not given a clear yes or no on the direction.

ACTION: John - ask Kevin Tian to give a clear go/no-go decision about the direction of this series
ACTION: Janakarajan Natarajan (AMD) to follow up within AMD

[PATCH 0/7] paravirtual IOMMU interface

Sent in for meeting agenda by George
https://marc.info/?l=xen-devel&m=151843249327749
https://xen.markmail.org/thread/kmxk4hoj2ao65qsa

v1 posted by Paul Durrant on 12 Feb 2018.

Seems to have had a lot of feedback from Kevin Tian.

Paul: Not entirely sure about IOMAPPING of pages. We need to address the CPU side mappings in the grant tables: to implement this we have to relax one of the permissions checks.

George: Was that the check that saved us from some XSAs?

Andrew: The problem is that “page ownership” is rather complicated and x86 & arm are completely different.

Andrew believes we have to address this area. But it would be nice not to add another band-aid.

ACTION: Paul to resend the series with a clear problem statement. It may also make sense for Andy, Paul and George to sit together

[PATCH v4 0/4] x86/cpuid: enable new cpu features

Latest Posting Date: Wed, 3 Jan 2018
Link: https://lists.xen.org/archives/html/xen-devel/2018-01/msg00049.html
From: Yang Zhong
Number of ACKs: 0
Dependencies: Test cases and blowfish test

Issues: Jan thought those patches were okay for him, but he asked Yang to implement test cases for GFNIs and use blowfish tool to check other encryption related CPU features

Jan: In reviewing these first four issues patches it appeared that the implementation was not done without testing. Thus I was asking about testing. There is a tool in the tests directory that compiles the x86 emulator in user space and runs with it. One of the tests is a compiled version of blowfish. We can and should do similar things for every new x86 emulator features, as the emulator is complex and an area which has in the past created many XSAs.

We would also run the AFL fuzzer over it: George can point to it.

ACTION: Lars to point to the existing tool
See http://xenbits.xen.org/gitweb/?p=xen.git;a=tree;f=tools/tests/x86_emulator
Also, if you look in Jan's emulator series, most patches touch both the hypervisor and that test logic, e.g. many patches in https://xen.markmail.org/thread/roukz6r3gcuhxinn

ACTION: John will make sure that Yang is following up on this.

AOB

Meeting format

Andy: no suggestions to change
Lars: the only issue I noticed that we had people
Video conference: do this as needed (most conference services have reasonably worming html5)

XPTI Status

Intel asked about XPTI status: the background is whether many more changes are expected and whether it is safe to rebase series.

XPTI functionally works. Some performance issues will need to be addressed, but these should have a fairly small impact on series. 5 level paging came up in particular: XPTI should not have much impact on this and hardly any going forward. So rebasing now should be safe.

ACTION: Lars to add XPTI / PVH update sections to the next meeting

Not Discussed at this meeting

Longer Term - Issues

[PATCH RFC 00/14] EPT-Based Sub-page Write Protection Support

Sent in for meeting agenda by George
https://marc.info/?l=xen-devel&m=150840502417156
https://xen.markmail.org/thread/m75h6b2aiwk5h7fx
RFC posted by Zhang Yi Oct 19, 2017

No acks, reviews only by memaccess maintainers / developers

Issues: Use case for the feature is still not clear and needs discussion

No time to discuss: we spent a lot of time on “Add vNVDIMM support to HVM domains” and I felt that this was another lengthy discussion and prioritized others first.

Longer Term - No Code Reviews yet

[PATCH RESEND v1 0/7] Intel Processor Trace virtulization enabling

Sent in for meeting agenda by George
https://marc.info/?l=xen-devel&m=151608947805423

https://xen.markmail.org/thread/rbaf7cxh2a7wwchf

v1.1 Posted by Kang, Luwei on 15 January 2018.
Issue: No feedback.

[RFC PATCH v2 00/17] RFC: SGX Virtualization design and draft patches

Latest Posting Date: Mon, 4 Dec 2017
Link: https://lists.xen.org/archives/html/xen-devel/2017-12/msg00104.html
From: Boqun Feng
Number of ACKs: 0

Issue: No feedback.

Longer Term - Progressing or Waiting

[PATCH v4 00/28] add vIOMMU support with irq remapping function of virtual VT-d

Sent in for meeting agenda by George
v3 posted by Lan Tianyu on 22 September 2017: marc.info/?l=xen-devel&m=150607140722407
v4 posted by Chao Gao: https://xen.markmail.org/thread/wfyorbn3nzsio6s7

Seems to have had review by Roger Pau Monne (1 ACK)
No issues

[RFC Patch v4 0/8] Extend resources to support more vcpus in single VM

Sent in by George
RFC v3 by Lan Tianyu: https://marc.info/?l=xen-devel&m=150530044827940 (Sep 17)
RFC v4 re-posted by Chao Gao: https://xen.markmail.org/thread/tlto7b3fadp7kkw6 (Dec 17)

From: Chao Gao
Number of ACKs: 2
Quite a bit of feedback on v4 from a few people up to Feb 28th

Dependencies: Virtual interrupt remapping of virtual VT-d and Changes to IOREQ server is based on Paul Durrant's "x86: guest resource mapping".

[RFC PATCH 0/8] Add guest CPU topology support

Sent in for meeting agenda by George
https://marc.info/?l=xen-devel&m=151538433419631
https://xen.markmail.org/thread/od46uc5nwhshnluz

Some feedback from Andrew Cooper and Daniel De Graaf

Dependencies: Andrew's CPUID work. Currently, this version doesn't have any dependency. But Andrew thought it was on the wrong direction. So Chao decided to wait for Andrew’s work to finish and rework based on CPUID.