



Nutanix, AHV and Performance

Dr Felipe Franciosi

AHV Engineering Lead

OCTOBER 2018 | OPEN SOURCE SUMMIT EUROPE

> Disclaimer

This presentation and the accompanying oral commentary may include express and implied forward-looking statements, including but not limited to statements concerning our business plans and objectives, product features and technology that are under development or in process and capabilities of such product features and technology, our plans to introduce product features in future releases, the implementation of our products on additional hardware platforms, strategic partnerships that are in process, product performance, competitive position, industry environment, and potential market opportunities. These forward-looking statements are not historical facts, and instead are based on our current expectations, estimates, opinions and beliefs. The accuracy of such forward-looking statements depends upon future events, and involves risks, uncertainties and other factors beyond our control that may cause these statements to be inaccurate and cause our actual results, performance or achievements to differ materially and adversely from those anticipated or implied by such statements, including, among others: failure to develop, or unexpected difficulties or delays in developing, new product features or technology on a timely or cost-effective basis; delays in or lack of customer or market acceptance of our new product features or technology; the failure of our software to interoperate on different hardware platforms; failure to form, or delays in the formation of, new strategic partnerships and the possibility that we may not receive anticipated results from forming such strategic partnerships; the introduction, or acceleration of adoption of, competing solutions, including public cloud infrastructure; a shift in industry or competitive dynamics or customer demand; and other risks detailed in our Form 10-Q for the fiscal quarter ended April 30, 2017, filed with the Securities and Exchange Commission. These forward-looking statements speak only as of the date of this presentation and, except as required by law, we assume no obligation to update forward-looking statements to reflect actual results or subsequent events or circumstances. Any future product or roadmap information is intended to outline general product directions, and is not a commitment, promise or legal obligation for Nutanix to deliver any material, code, or functionality. This information should not be used when making a purchasing decision. Further, note that Nutanix has made no determination as to if separate fees will be charged for any future product enhancements or functionality which may ultimately be made available. Nutanix may, in its own discretion, choose to charge separate fees for the delivery of any product enhancements or functionality which are ultimately made available.

Certain information contained in this presentation and the accompanying oral commentary may relate to or be based on studies, publications, surveys and other data obtained from third-party sources and our own internal estimates and research. While we believe these third-party studies, publications, surveys and other data are reliable as of the date of this presentation, they have not independently verified, and we make no representation as to the adequacy, fairness, accuracy, or completeness of any information obtained from third-party sources.

Agenda

- 1 Nutanix Enterprise Cloud Platform
- 2 Open Source at Nutanix
- 3 Acropolis Hypervisor
- 4 Next Steps in Performance



Nutanix

Nutanix delivers software-defined hyperconverged infrastructures


Fast Facts

- Founded in 2009
- 4,000+ Employees (1,000+ Engineers)
- 10,000+ Customers
- Presence in 145 Countries
- HCI Leader (Gartner MQ 2018, Forrester Wave 2018)
- NASDAQ: NTNX
- Net Promoter Score: + 90

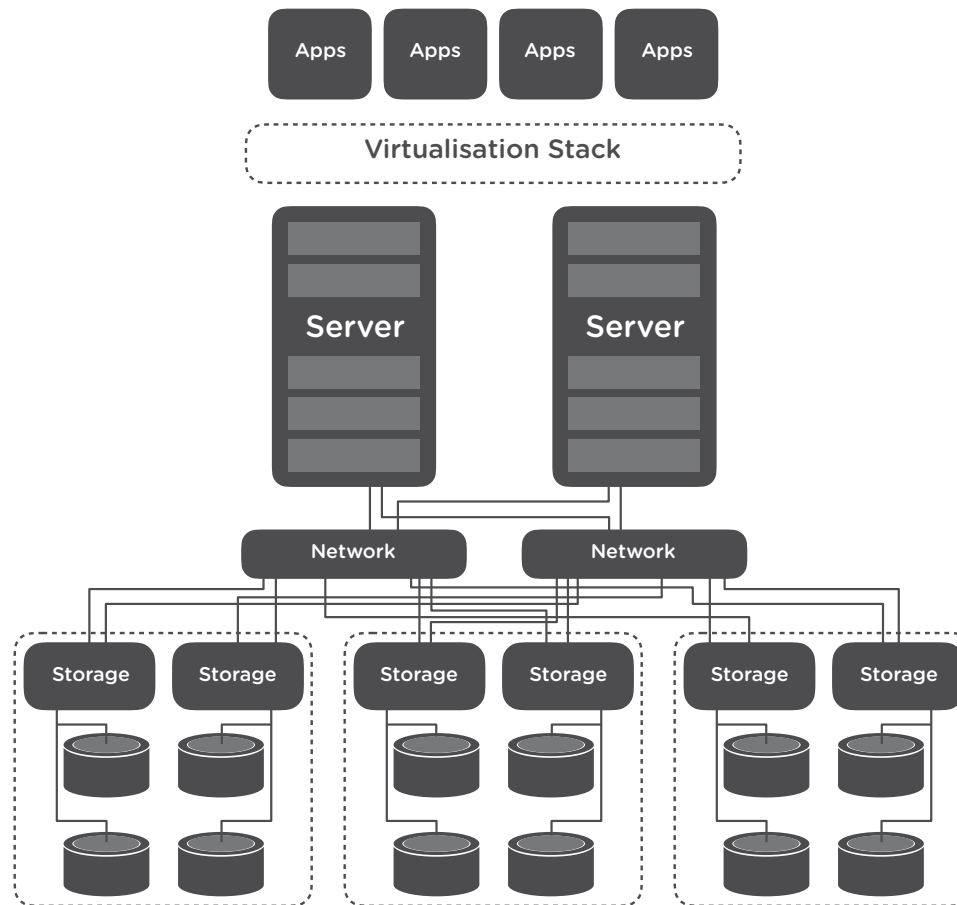


Nutanix Headquarters in San Jose, California

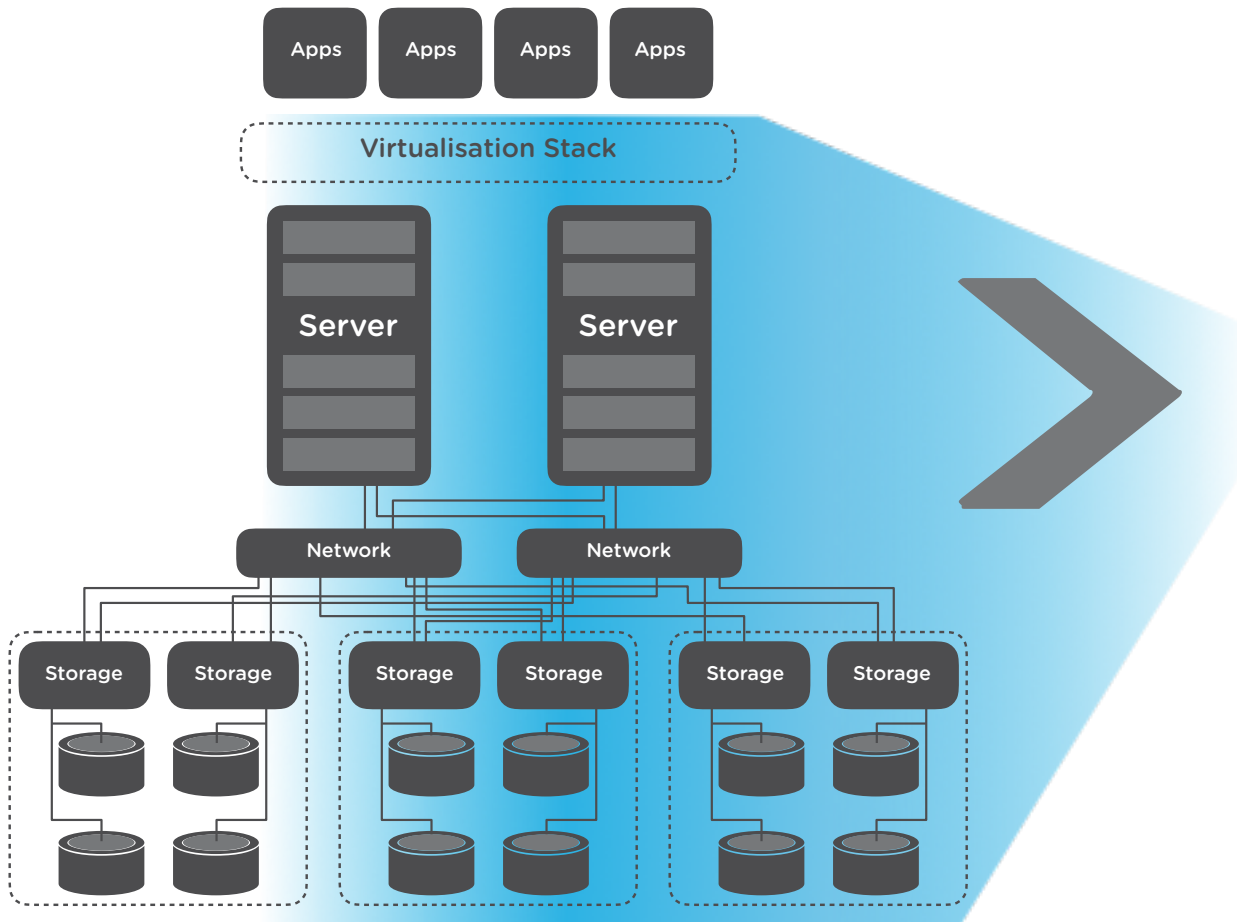
Data Centre Management is Painful



Traditional Data Centre



Modern Data Centre



**Integrated, scale-out
compute and storage
with built-in virtualisation
and management**



Traditional vs Modern Data Centre

Function-specific Hardware
(Hardware-Defined Data Centre)
Expensive, Inflexible



source: Invisible Infra 10/02/2013

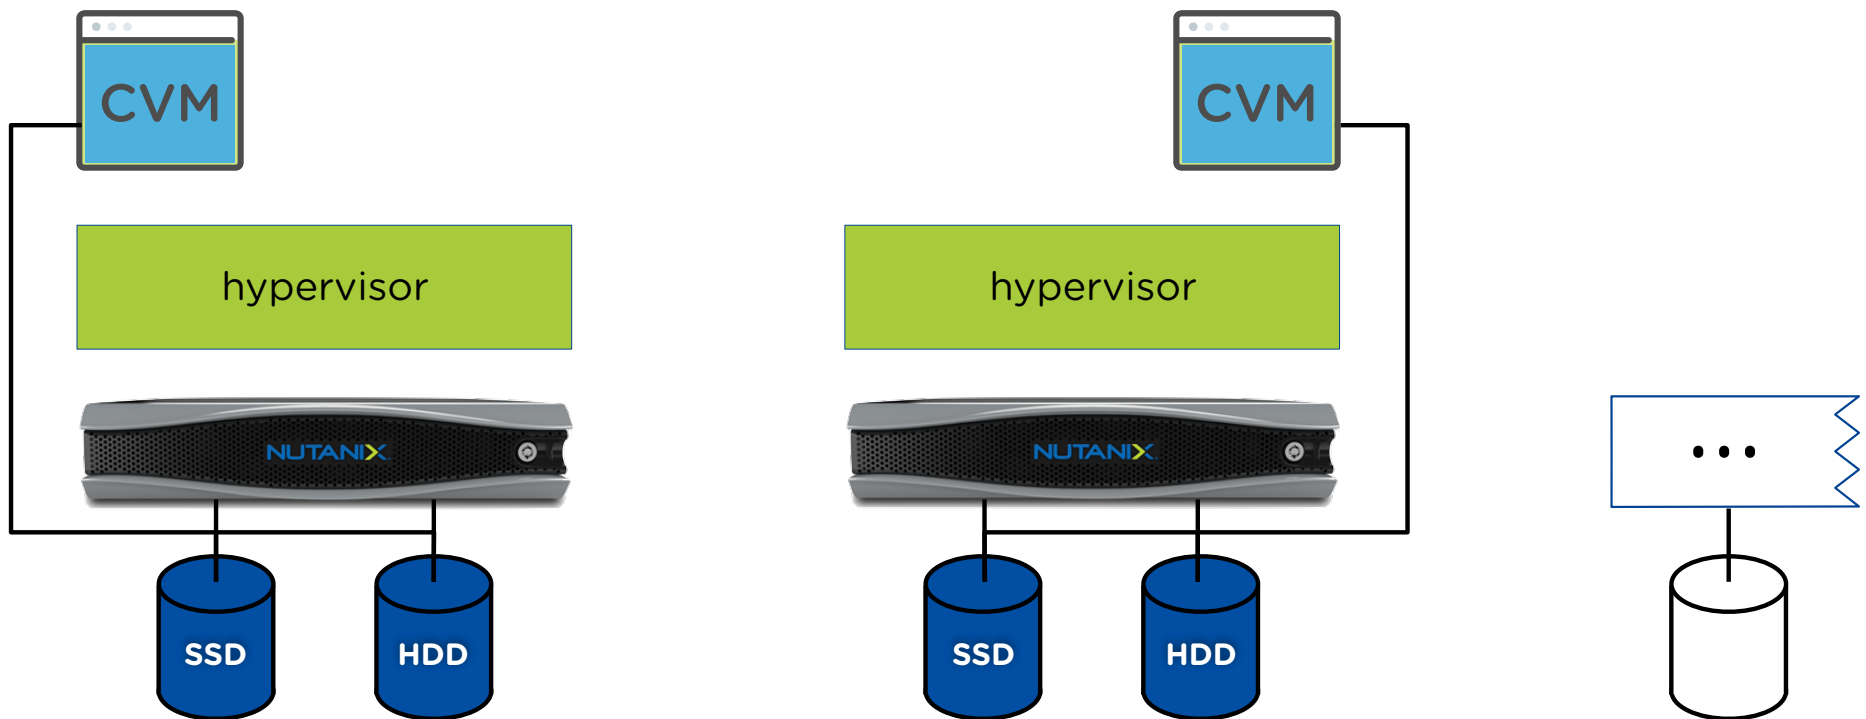
Commodity Hardware
(Software-Defined Data Centre)
Inexpensive, Flexible



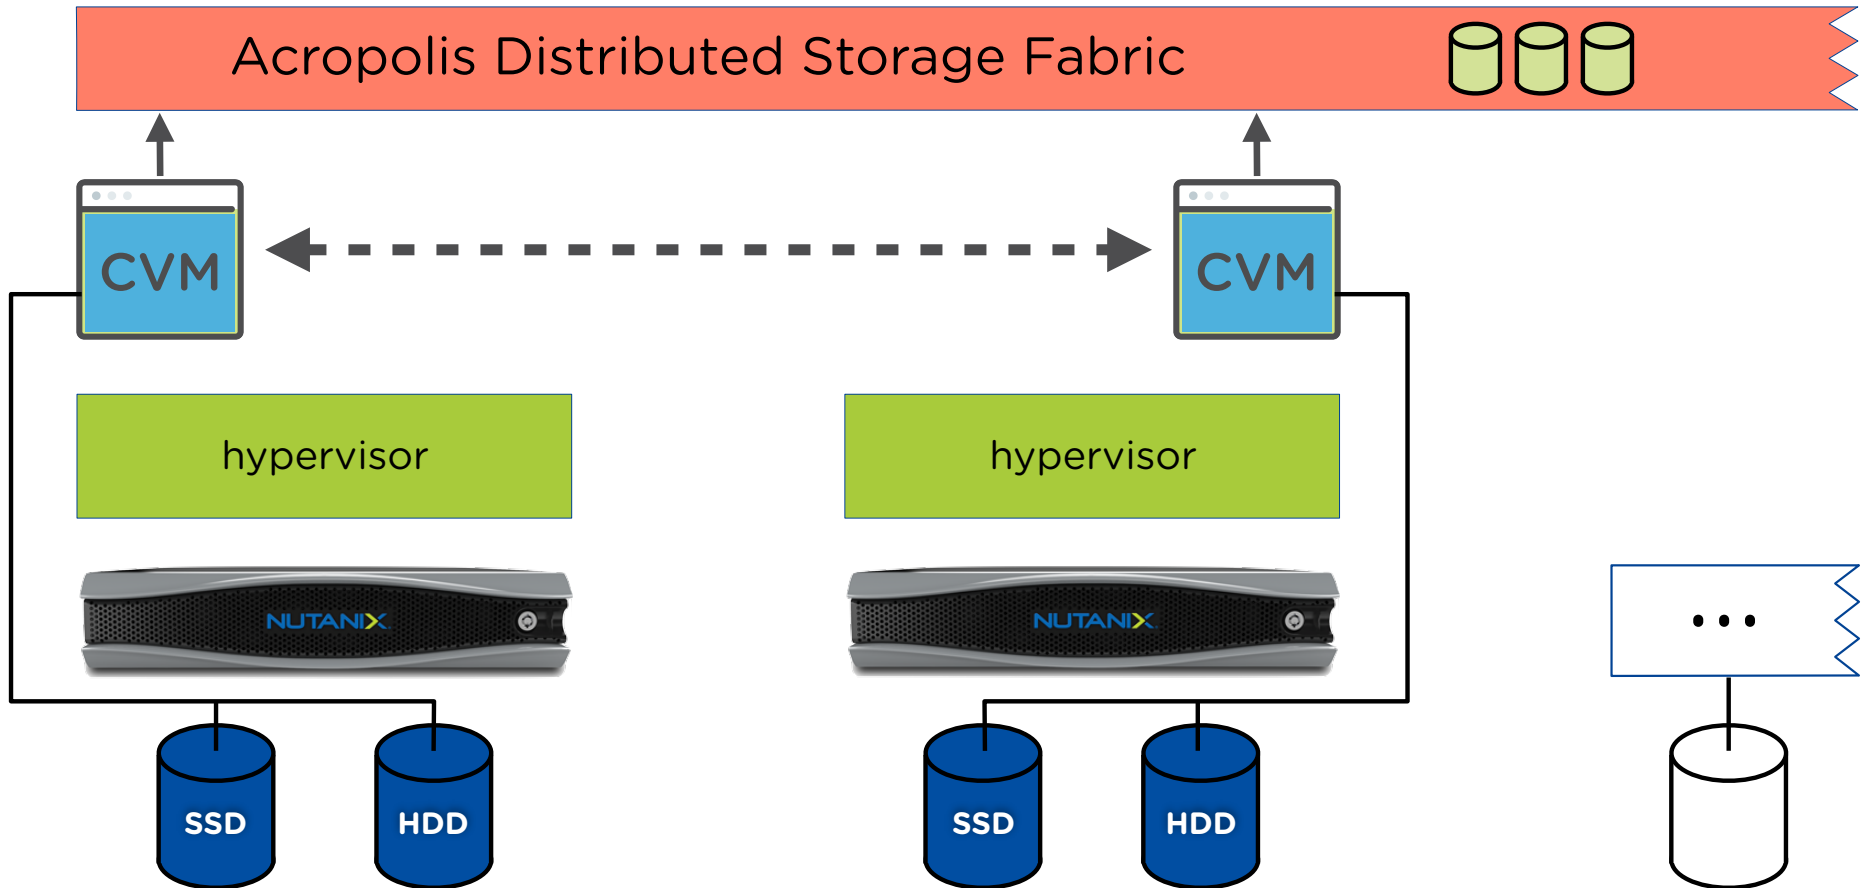
source: Le Monde Informatique 13/10/2015



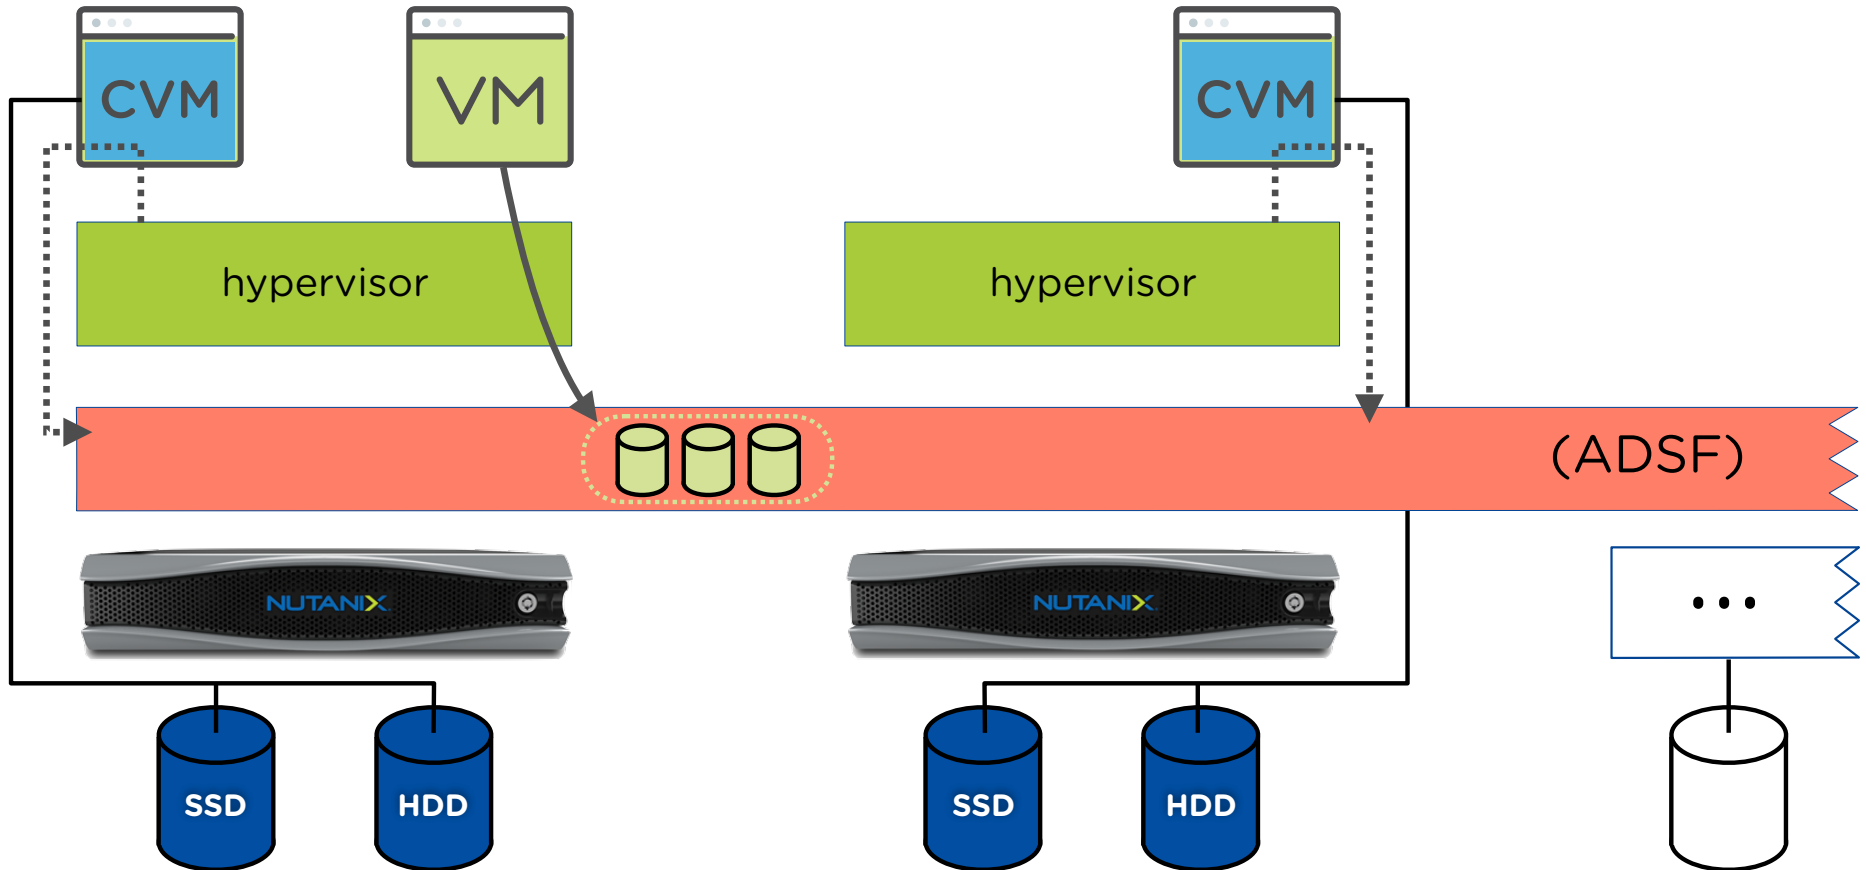
Nutanix: Behind the Scenes




Nutanix: Behind the Scenes



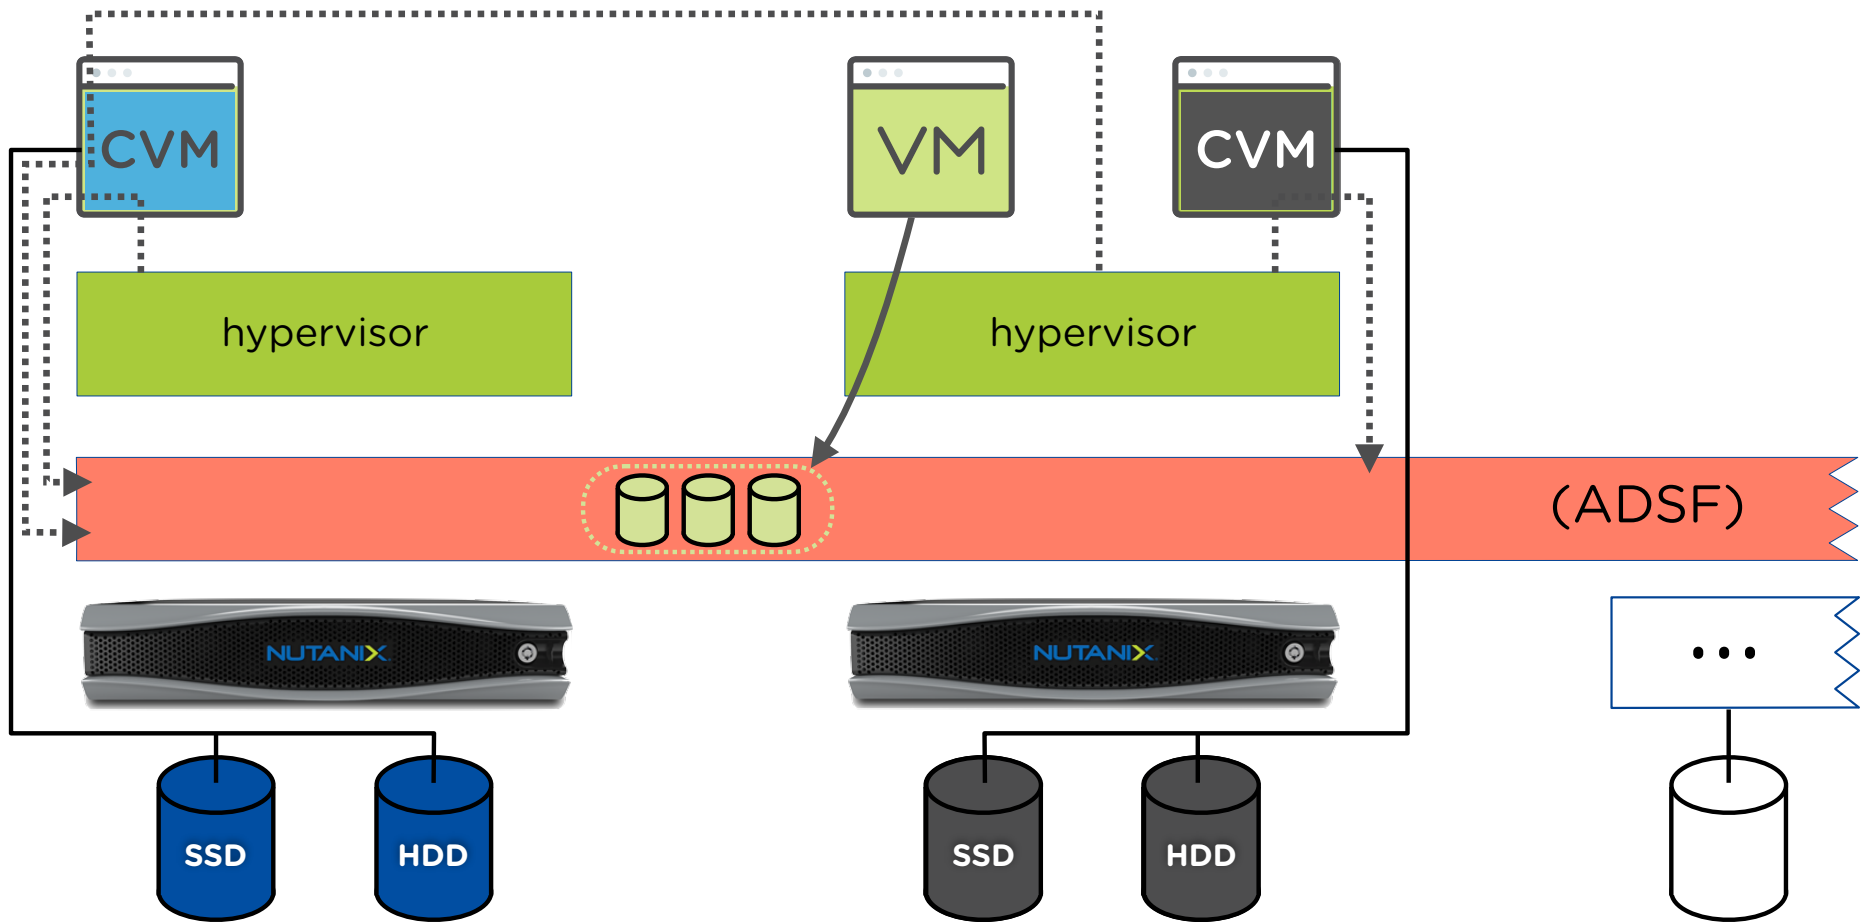
Nutanix: Behind the Scenes




Nutanix: Behind the Scenes



Nutanix: Behind the Scenes



xkcd: 1737



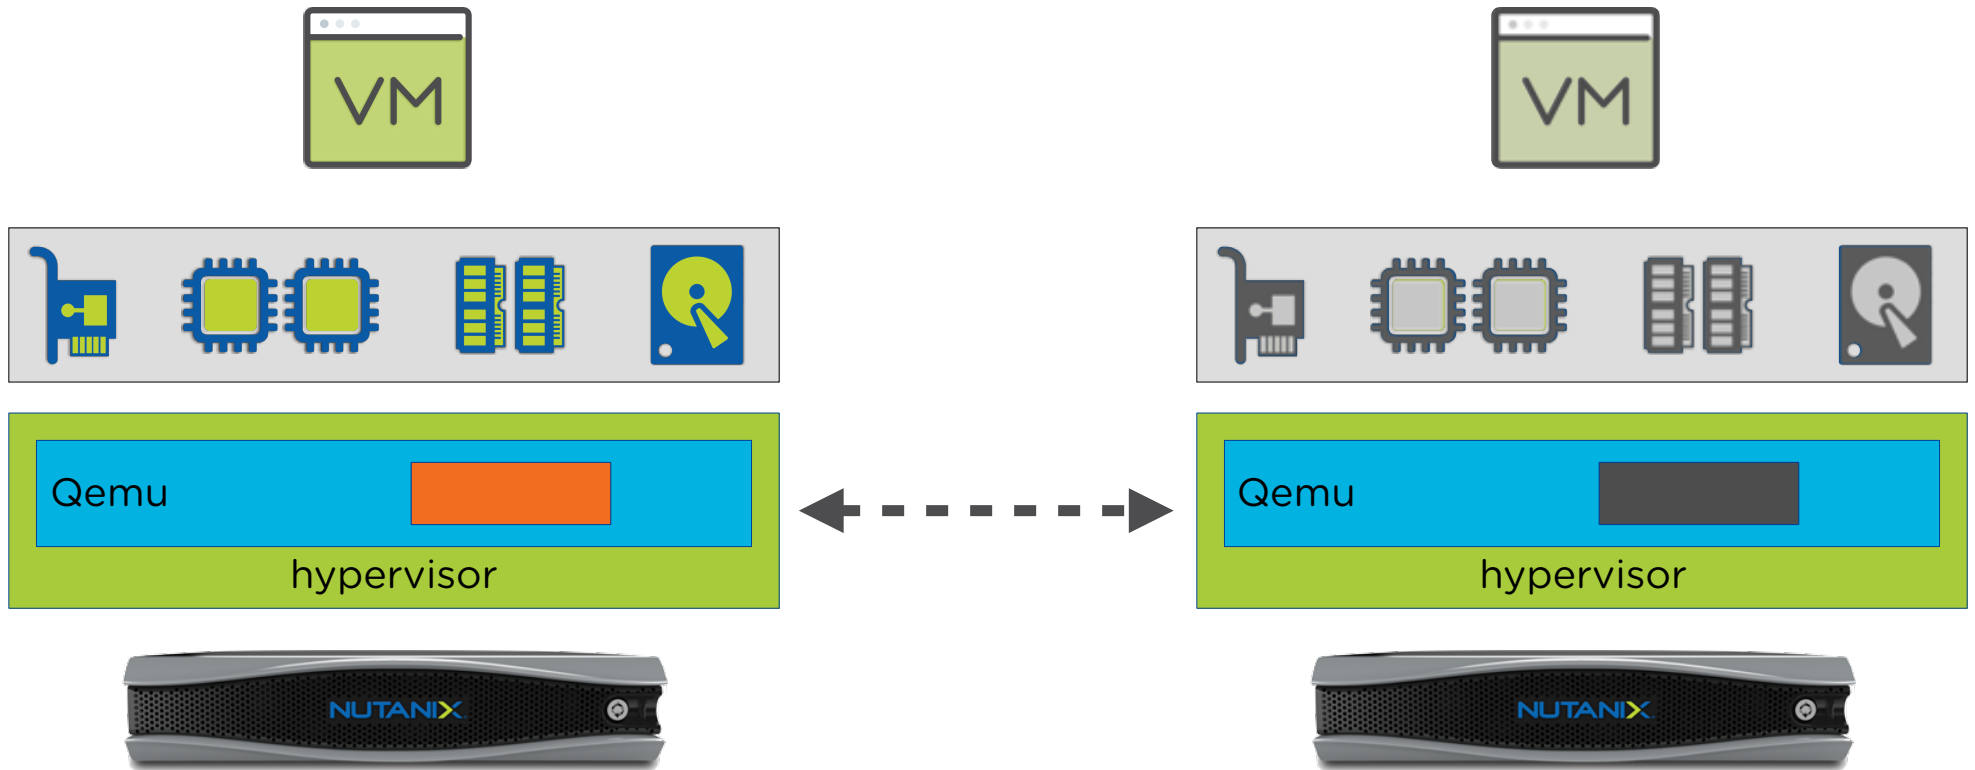
Nutanix: The Acropolis Hypervisor (AHV)

Acropolis Hypervisor

- **Virtualisation Appliance** based on CentOS 6.9
- Uses Various **Open Source** Projects
 - Linux **Kernel 4.4 LTS (KVM)**
 - **Qemu 2.6**
 - **Libvirt 2.0**
 - **Open vSwitch 2.5.2**
- Nutanix **contributes to Open Source**
 - **Live Migration** Improvements
 - Virtualisation **Features**
 - Security/Functional/Test **Fixes**

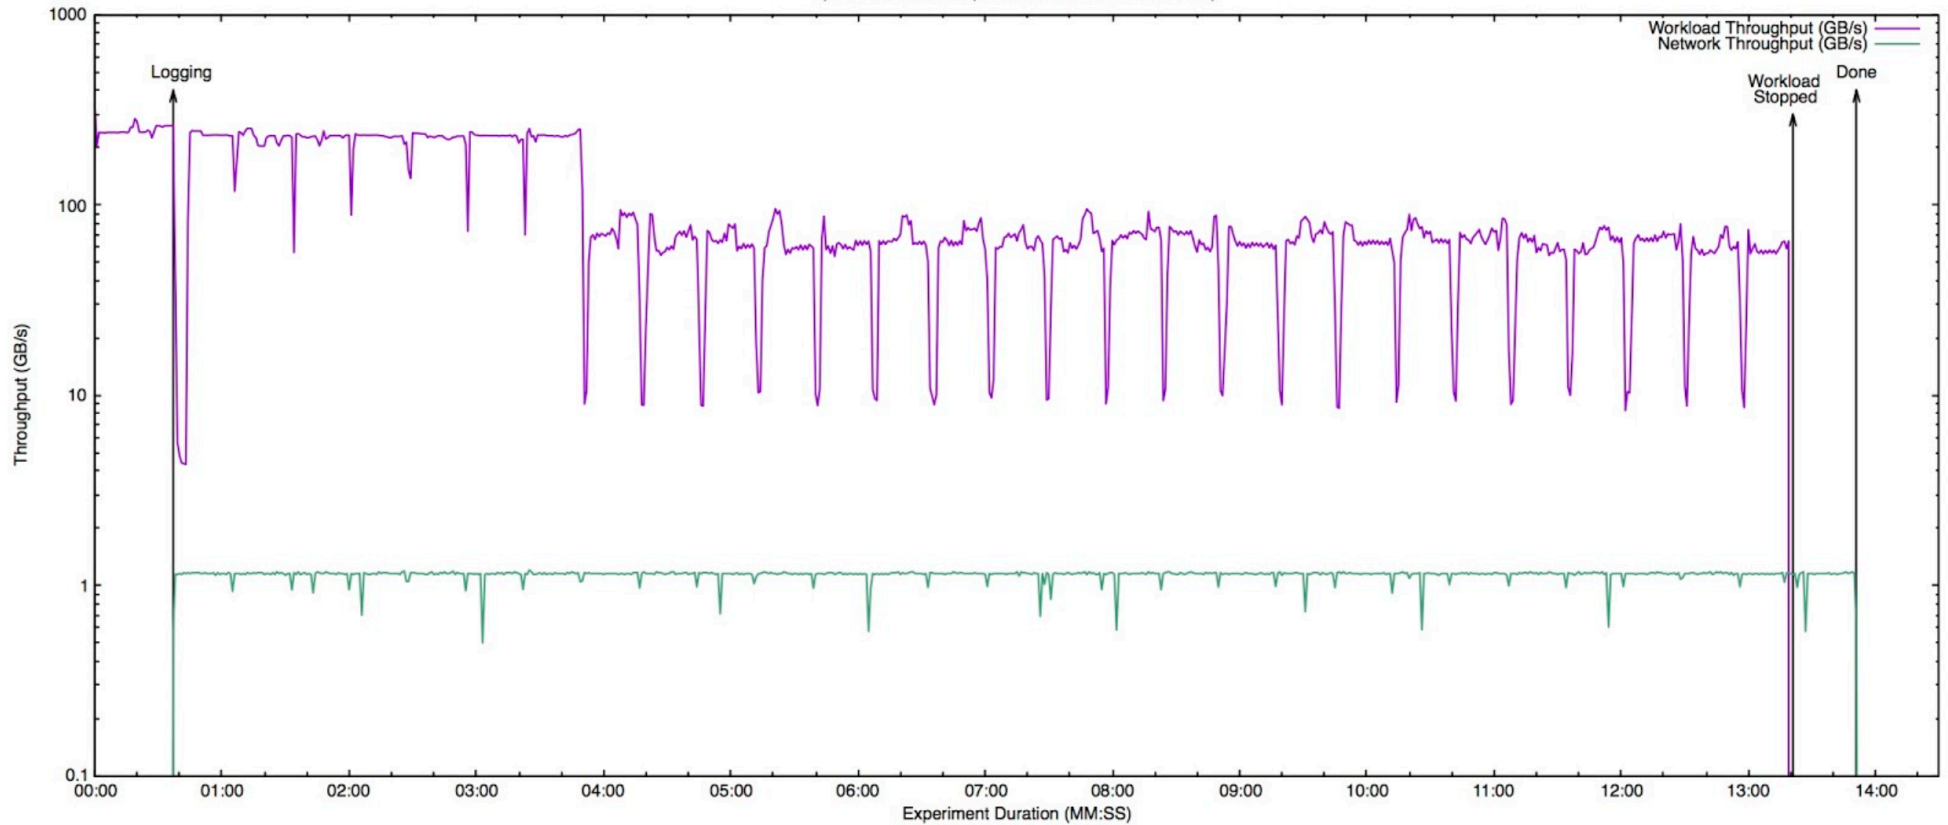


AHV: Live Migration




AHV: Live Migration Performance

Migration of a 'Dirty Harry' VM with 32GB of RAM and 8 vCPUs
(Workload using 30GB Dataset and 8 Dirty Threads)
qemu-2.3.0-nutanix (from AHV Asterix 20160925.50)



AHV: Live Migration Performance

Migration of a 'Dirty Harry' VM with 32GB of RAM and 8 vCPUs
(Workload using 30GB Dataset and 8 Dirty Threads)
qemu-2.6.0-nutanix (from AHV Master)

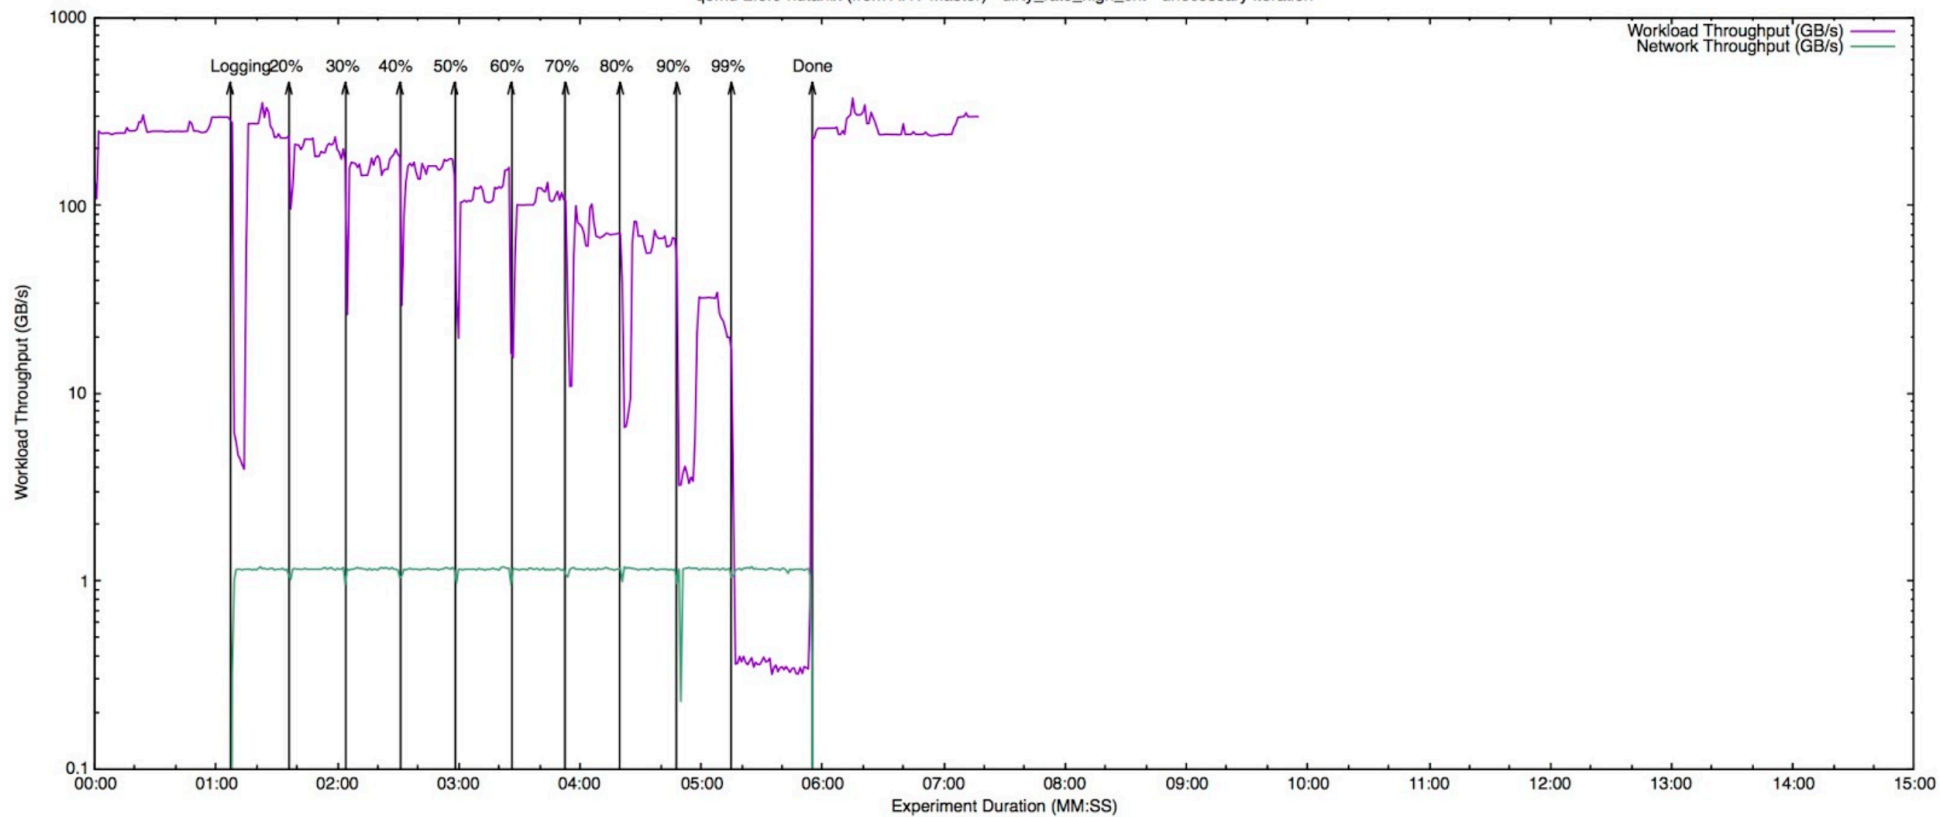


AHV: Live Migration Performance

Migration of a 'Dirty Harry' VM with 32GB of RAM and 8 vCPUs


(Workload using 30GB Dataset and 8 Dirty Threads)

qemu-2.6.0-nutanix (from AHV Master) - dirty_rate_high_cnt - unnecessary iteration



AHV: Live Migration Performance

Migration of a 'Dirty Harry' VM with 32GB of RAM and 8 vCPUs
(Workload using 30GB Dataset and 8 Dirty Threads)
qemu-2.3.0-nutanix (from AHV Asterix 20160925.66 - ca77019a)



AHV: Storage Performance

NVMe drives are highly parallel and very fast

- Require high request size for top throughput
- Require high queue depth for top IOPS
- Require multiple cores to drive top performance

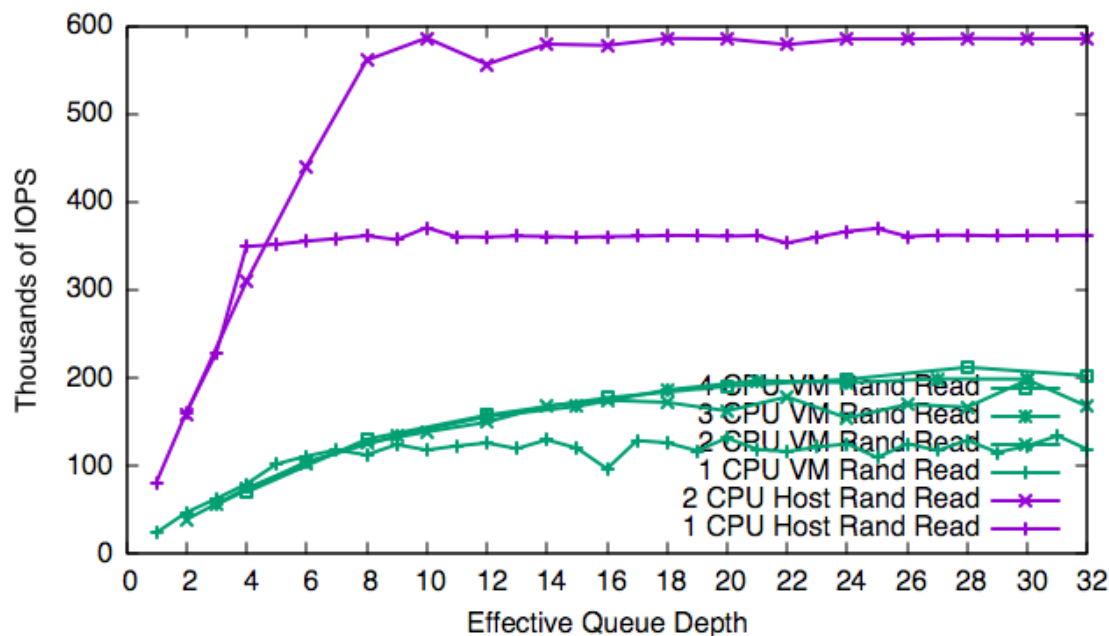
Very challenging for hypervisors

- For CPU efficiency, require either:
 - Entire userspace stacks (eg. SPDK)
 - Hardware offload (eg. SmartNICs)
- Both methods need "datapath transfer"

Intel P4800 SSDPE21K375GA (FW E2010324)

2 x Intel(R) Xeon(R) CPU E5-2667 v4 3.20GHz

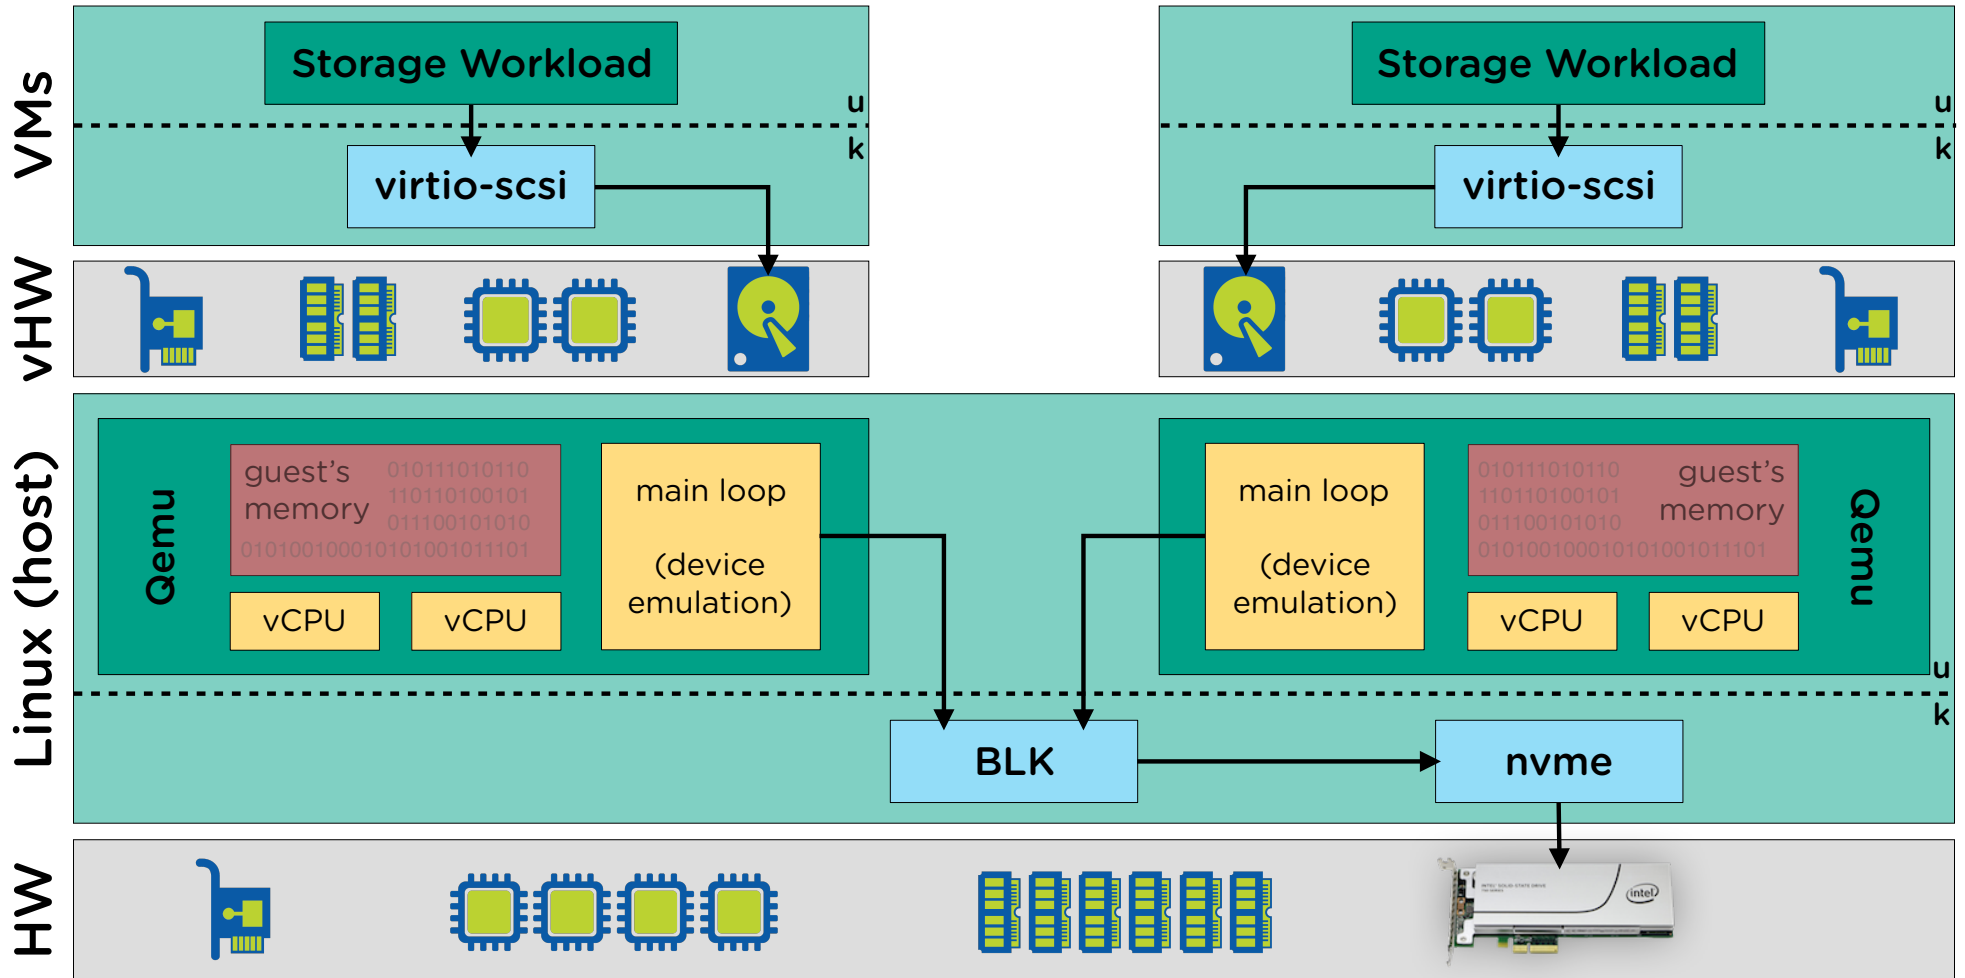
AHV 20170830 (Off EL6 and 4.4.77), FIO 2.0.13



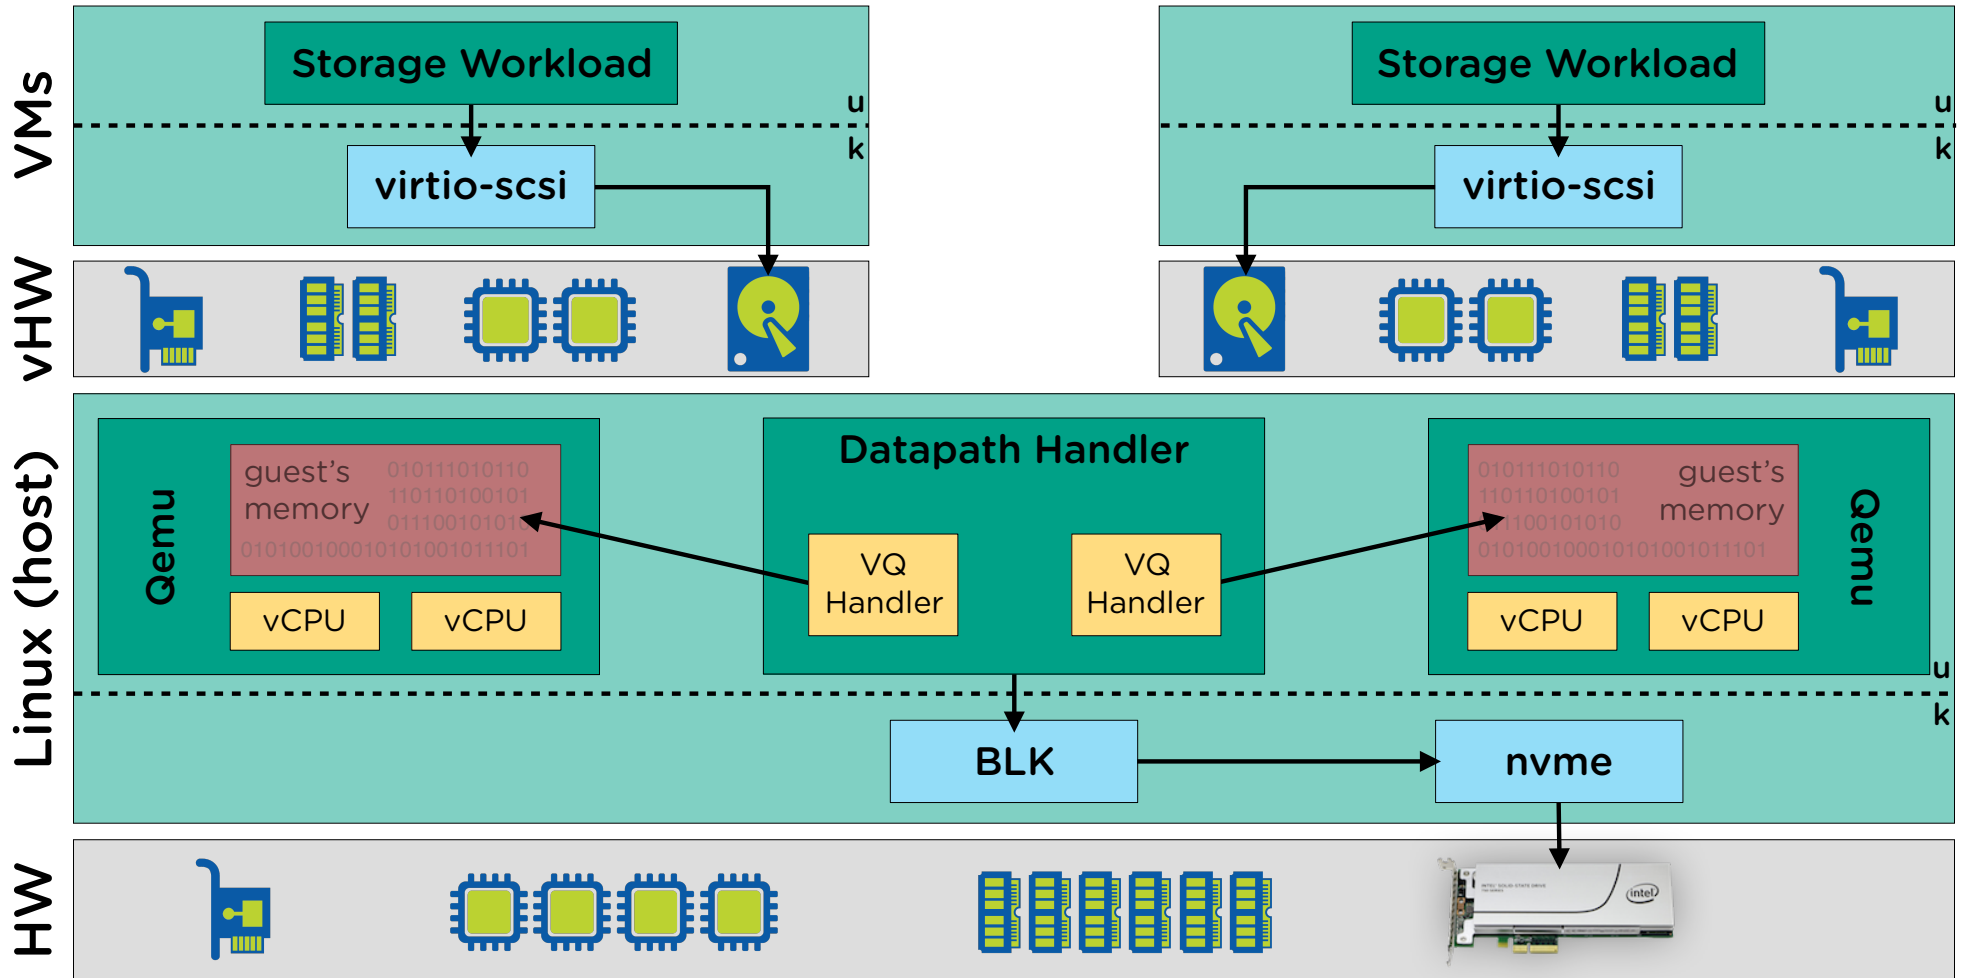
VM using Debian 9.4 guest, 8 vCPUs, 32GB vRAM, FIO 2.16
Host using Qemu 2.6 with virtio-scsi and AHV 20170830



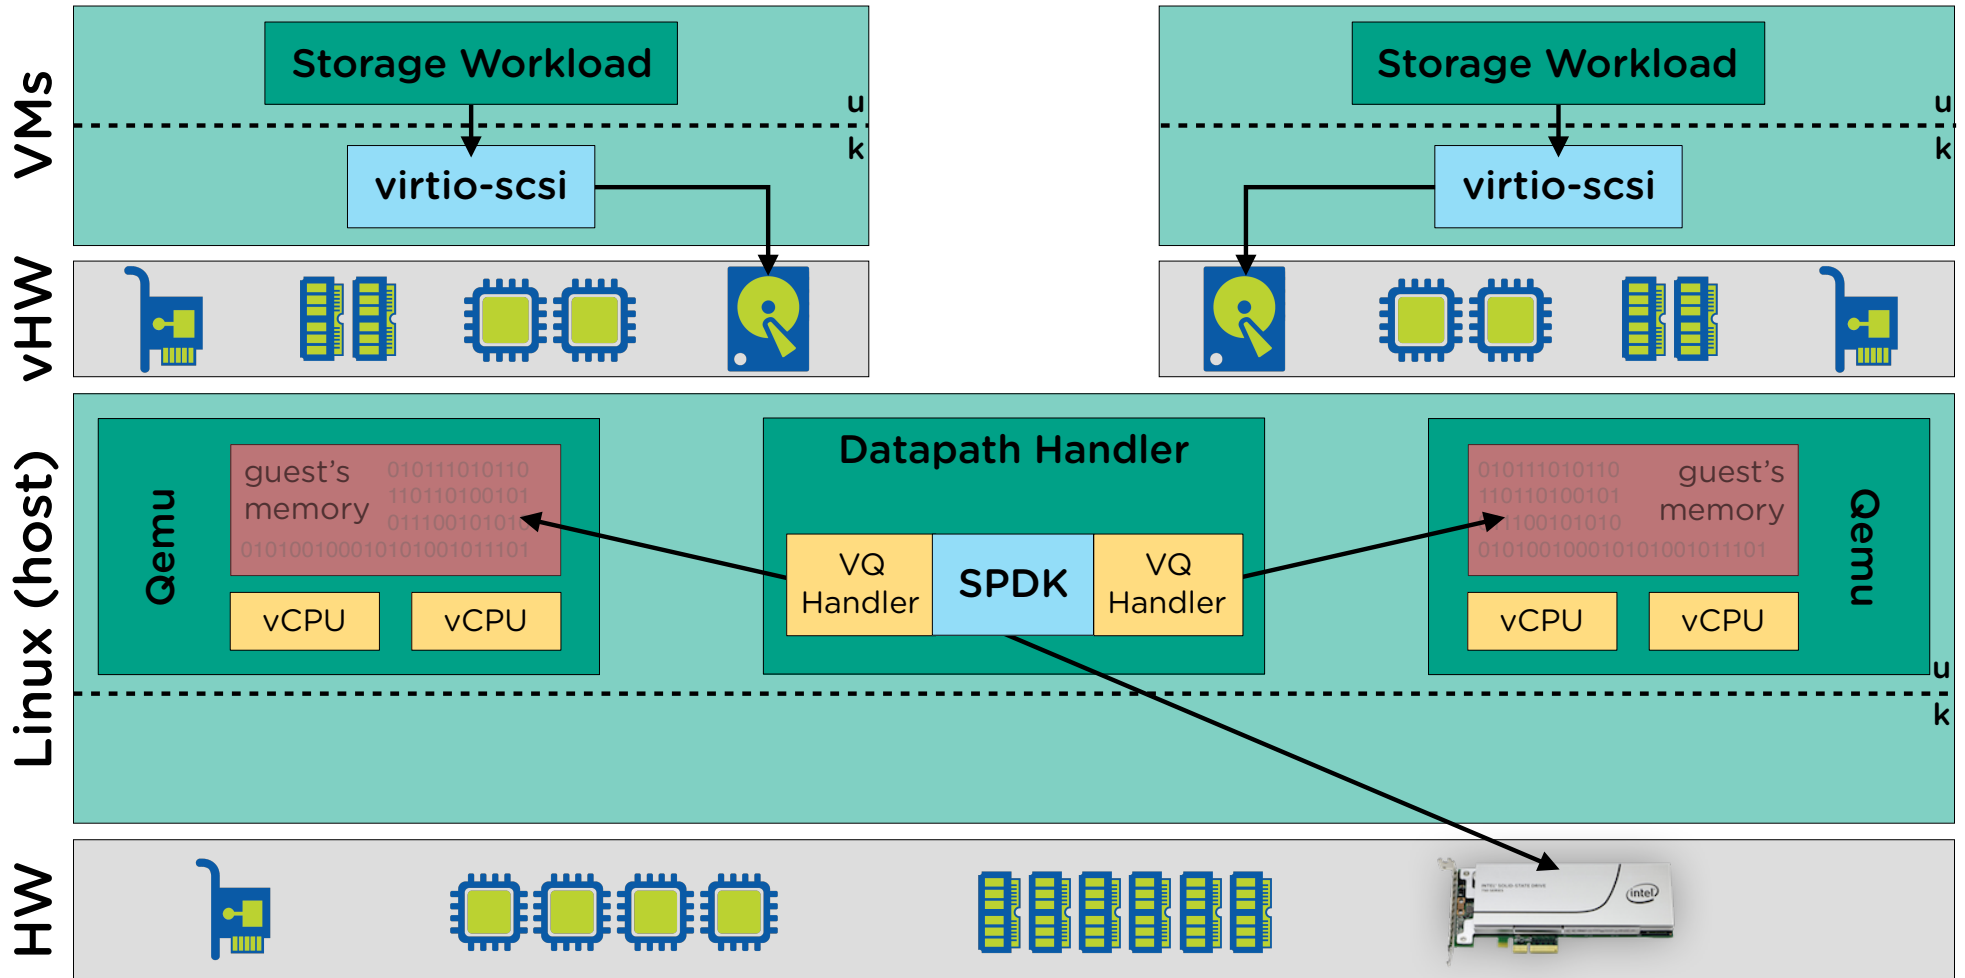
AHV: Storage Performance



AHV: Storage Performance



AHV: Storage Performance



AHV: Storage Performance


Does it really make a difference?

- Datapath offload enables various optimisations
- Separate process can handle IO more efficiently
- SPDK backend achieves bare metal performance
- Require multiple VM cores to drive top performance
 - Inefficient virtio-scsi implementation (kernel)

Intel P4800 SSDPE21K375GA (FW E2010324)

2 x Intel(R) Xeon(R) CPU E5-2667 v4 3.20GHz

AHV 20170830 (Off EL6 and 4.4.77), FIO 2.0.13



VM using Debian 9.4 guest, 8 vCPUs, 32GB vRAM, FIO 2.16
Host using Qemu 2.6 with virtio-scsi and AHV 20170830




AHV: Storage Performance

Does it really make a difference?

- Datapath offload enables various optimisations
- Separate process can handle IO more efficiently
- SPDK backend achieves bare metal performance
- Require multiple VM cores to drive top performance
 - Inefficient virtio-scsi implementation (kernel)

Driving virtio-scsi from guest

- Entire userspace stack all the way from the guest
- Bypasses kernel in the guest and hypervisor



VM using Debian 9.4 guest, 8 vCPUs, 32GB vRAM, FIO 2.16
 Host using Qemu 2.6 with virtio-scsi and AHV 20170830



We  Open Source

F R E E D O M
T O
Invent



The background features a large, stylized 'X' shape composed of four overlapping, rounded rectangular segments. The segments are colored in a gradient from light blue to light green. The Nutanix logo is centered within this 'X' shape.

NUTANIX™

Thank you