HPVM experiences

Lars “Zinken” Sundqvist
Mission Critical Ambassador (OpenVMS and HP-UX)

Agenda

1. **Overview**
   - What is HPVM
2. **Todays HPVM versions**
   - HPVM 4.3 vs HPVM 6.1.5
3. **Customer usage, experiences and recommendations**
   - Including performance
4. **Futures and Q&A**

Overview

What is HPVM

HP Integrity Virtual Machines Technology Overview

Virtual hardware

Virtual Machine 1

app1 app2

Guest OS (OpenVMS)

I/O

Virtual Machine 2

app2

Guest OS (OpenVMS)

I/O

Virtual Machine 3

app1 app2

Guest OS (HP-UX)

I/O

host Virtual Machine Monitor (VM Monitor)

vmm driver PRM

Platform Manager (HP-UX)

HP-UX drivers

vm1 vm2 vm3 WLM

Memory

I/O

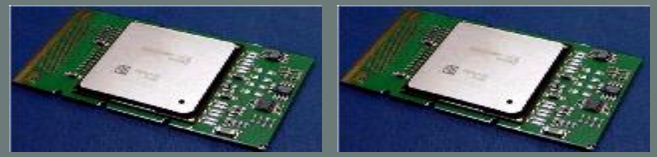

Dynamic I/O Sharing, networking

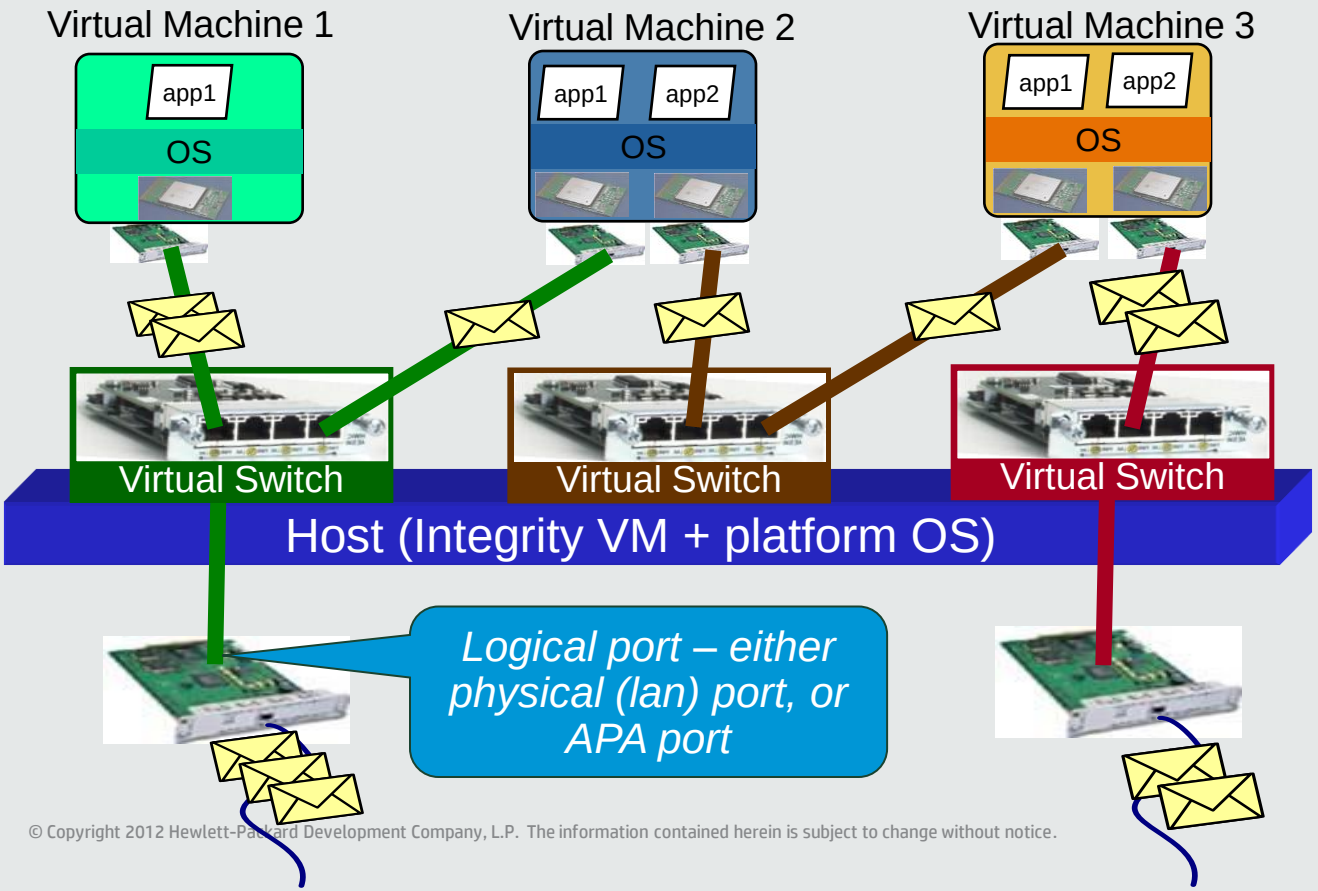

Virtual machine's network packets directed to physical NIC by the Integrity VM Host

Virtual NIC may be defined without a logical port for guest-to-guest communication

NIC can be isolated to a virtual machine

Storage I/O Virtualization

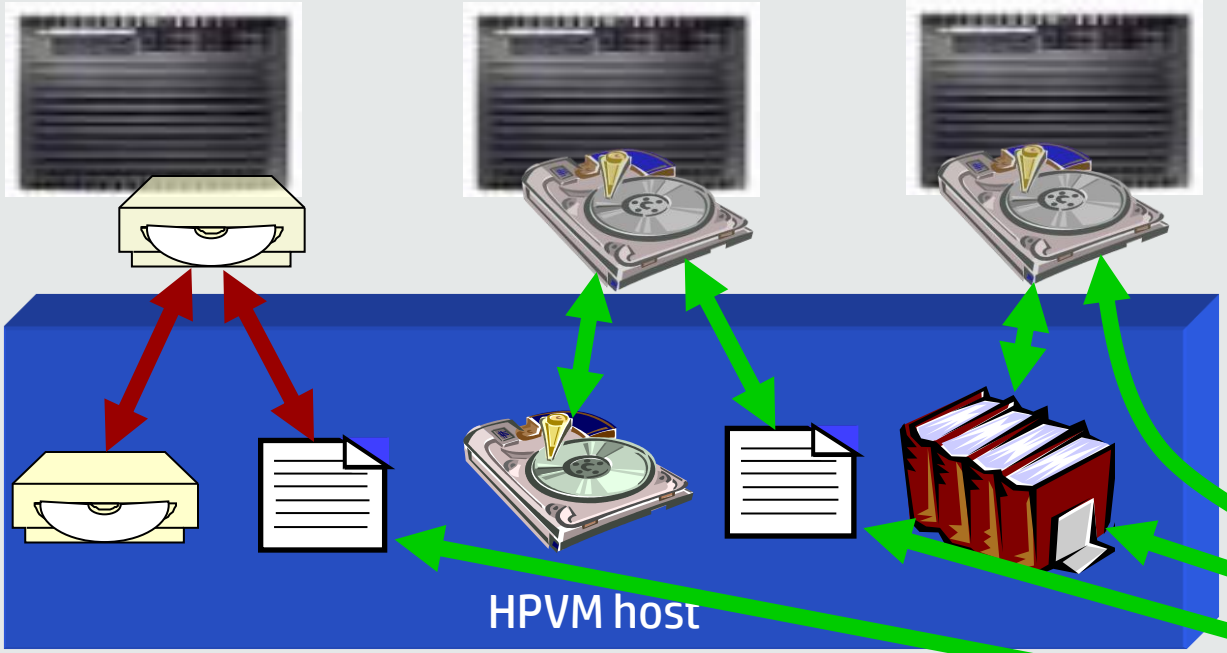

DVD virtualized on host by:
- -Physical DVD
- -File

Disk virtualized on host by:
- - Physical disk
- - File
- - Logical volume
- - SAN

Todays HPVM versions

HPVM 4.3 vs HPVM 6.1.5

Versions to use and not to use

- **Use:**
- ❑ **HP-UX 11iv3 September 2011 + HPVM 4.3 + PK2 + OpenVMS 8.4 + Update 500**
- ❑ **HP-UX 11iv3 September 2011 + HPVM 4.3 + PK2 + OpenVMS 8.4 + Update 600**
- ☒ **HP-UX 11iv3 September 2011 + HPVM 4.3 + PK2 + OpenVMS 8.4 + Update 700**

- **Dont use:**
- ❑ **HP-UX 11i v3 March 2012 + HPVM 4.3 + PK2 + OpenVMS 8.4 + any patchkit**
  - Not qualified which means not supported
- ❑ **HP-UX 11i v3 September 2012+ HPVM 6.1.x + OpenVMS 8.4 + any patchkit**
  - Not supported

Anyway, what's new with HPVM 6.1.5?

And why is not HPVM 6.x interesting to OpenVMS so far?

What have HPVM engineering been working with lately?

HP Virtualization Continuum for HP-UX (and OpenVMS)

HP nPartitions

Partitioning with electrical and security isolation

HP Virtual Partitions

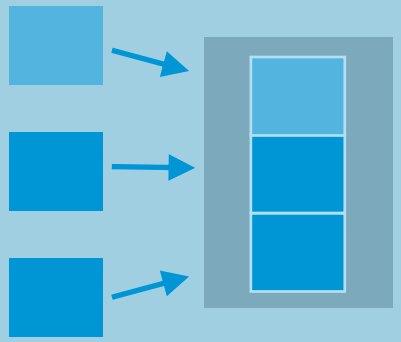

More granular partitioning at core level for additional flexibility

HP Integrity Virtual Machines

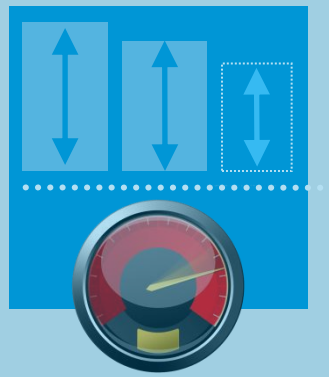

Software virtualization for dynamic, shared resource allocation and mobility

HP-UX Containers

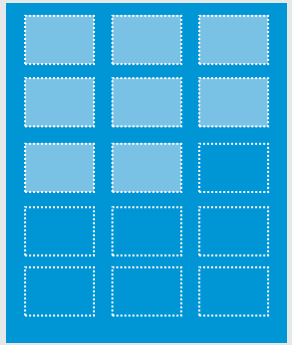

Shared OS virtualization to consolidate workloads within a single HP-UX instance

More Isolation **More Flexibility**

Offline vPars/VM v6.1.5 Transformation

- **Easily convert v6.1.5 vPars to VMs or vice versa**
- CPU Entitlement Implications
  - VM → vPar = CPUs become dedicated to the vPar
  - vPar → VM = VM vCPUs entitlements default to 100%
- Transformation Caveats
  - A v6.1.5 VSP can run either VMs or vPars – *not both simultaneously*

VSP: Virtualization Services Platform (aka HPVM Host, Hypervisor)

VM to vPar Transformation Example

```
VSP> hpvmstatus -s | grep supported
        vPar/VM types supported by this VSP = Shared

VSP> hpvmstop -p2
VSP> hpvmstatus
[Virtual Machines]
Virtual Machine Name VM #  Type OS Type State     #VCPUs  #Devs #Nets Memory
==================== ===== ==== ======= ========= ====== ===== ===== =======
atcuxvm2                 1 SH   HPUX    Off            8     4     1   64 GB
atcuxvm5                 2 SH   HPUX    Off            8     4     2   64 GB

VSP> hpvmstatus -s | grep supported
        vPar/VM types supported by this VSP = vPar, Shared

VSP> hpvmmodify -p2 -x vm_type=vpar
VSP> hpvmstart -p2

VSP> hpvmstatus -s | grep supported
        vPar/VM types supported by this VSP = vPar
```

VSP is in VM mode (i.e. Shared guests)

All VM guests stopped

VSP can run either vPars or VMs

Convert VM to vPar and start vPar

VSP is in vPar mode

Offline vPars v6.1.5 Migration

Move vPars between different physical servers

- **Distinguishing feature from previous versions of vPars!**
- Evacuate a VSP server for upgrades, repairs, etc.
- Rebalance VSP resources

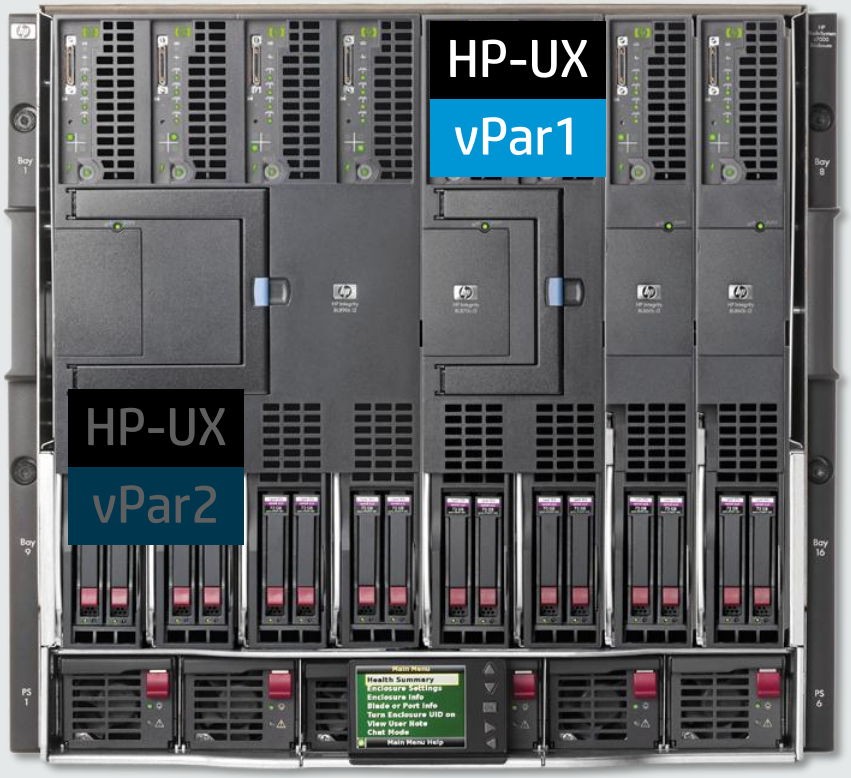

Customer usage, experience and recommendations

Some customer experiences

- **Usage is typical test and development environments**
  - Still is high availability needed
- **Use whole disks/LUNs as backend storage**
  - Decreases need for HP-UX knowledge
  - Better I/O performance

Integrity VM Guest Backing Store Type Comparison

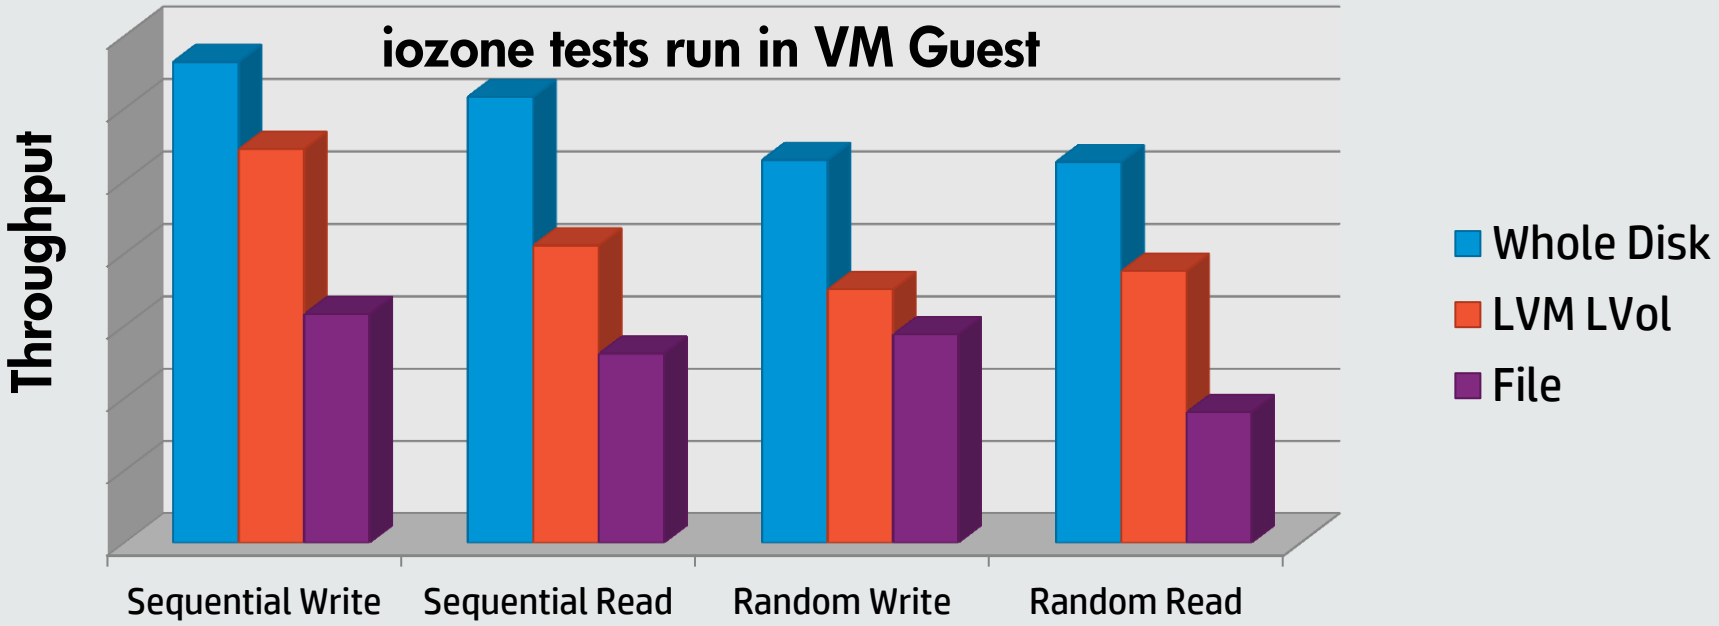

I/O Request Size Comparison: Larger is Better

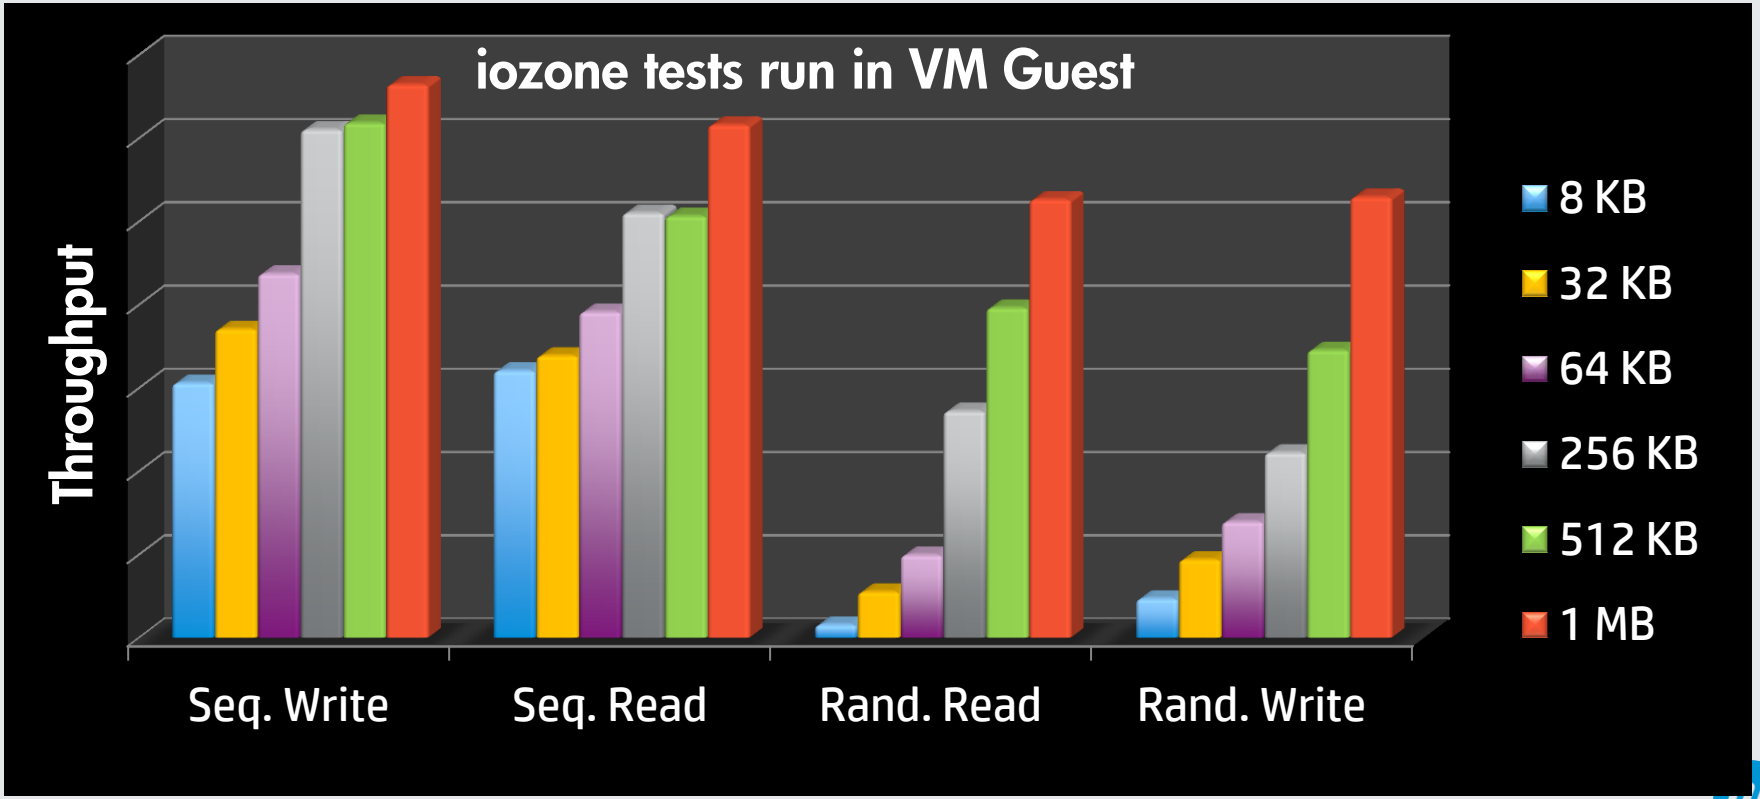

Some customer experiences continued.

- **Oldest installation in Sweden with OpenVMS guests?** October 2010
- **Oldest installation in Sweden with HP-UX guests?** July 2006
- **Dont be afraid on install to much memory**
  - Typically is more guests created then planned from start
- **On-line migrations actually helped to workaround a HPVM problem**
  - Possible by using HPVM + Online Migration licens, or better use VSE-OE or DC-OE

Some customer experiences continued..

- **Guests as cluster quorum nodes in a third location**
  - Why use a physical system?
- **Have a few ”your own” test guest configured**
  - Test modifications/upgrades on them first
  - For reference
- **Dont be affraid to test things**
  - Seeing is believing, done it yourself is good for health ☺
- **If unsure of HP-UX and patching, take help from someone who knows**
  - Be sure to document how to

Some customer experiences continued…

- **Document the installation and create your own short how-to guides**
  - Always helpful if needed or as reference
- **Treat the guest as a normal physical system**
  - But understand the MP (iLO3) versus vMP
- **Use GUI to creating/managing guests (vmmgr -> vsemgr)**
  - And copy the ”Command Priview” information as reference or putting into scripts
  - Or for better control, run the scripts manually
- **Plan for use of NPIV in future**
  - Use Volume Shadowing to add NPIV disks, remove AVIO disks
- **Keep one NIC on host not used by a vSwitch**
  - Best Online migration performance

Some customer experiences continued….

- Use Dedicated Network for Online VM Migrations
- Do not configure a vSwitch on OVMM NICs
- Activating a vSwitch on a NIC causes:
  - Disable Checksum Offload (**CKO**)
  - Disable TCP Segmentation Offload (**TSO**)
  - Enable **Promiscuous Mode**
- Potentially affects Online Migration performance as well as any other Host –> Host network traffic
- Disabling the vSwitch does not automatically re-enable CKO/TSO on the NIC

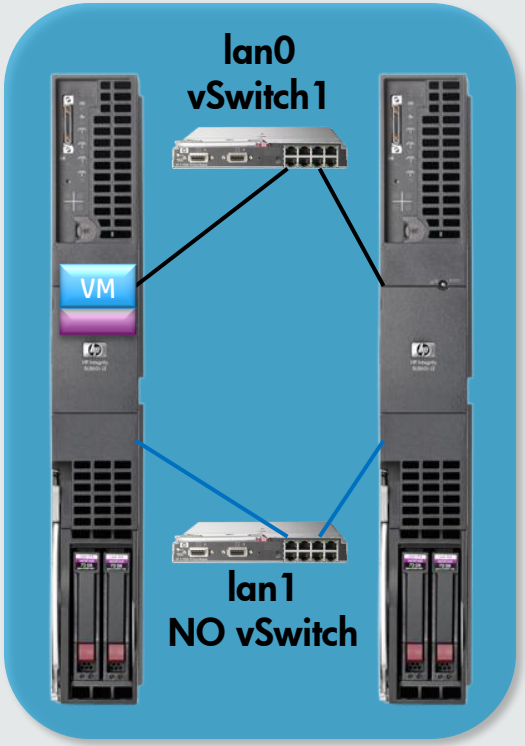

Some customer experiences continued….

- Use larger I/O requests if possible

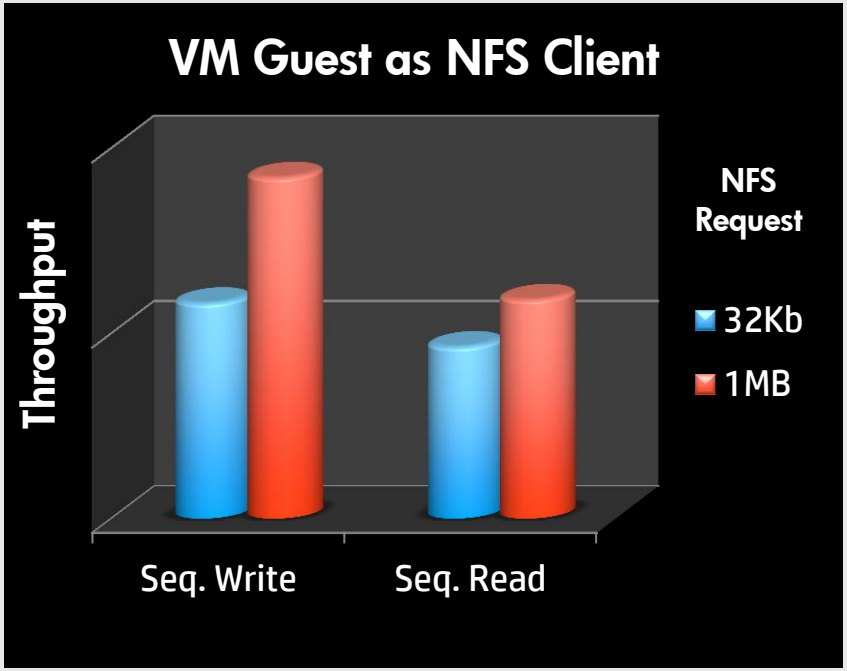

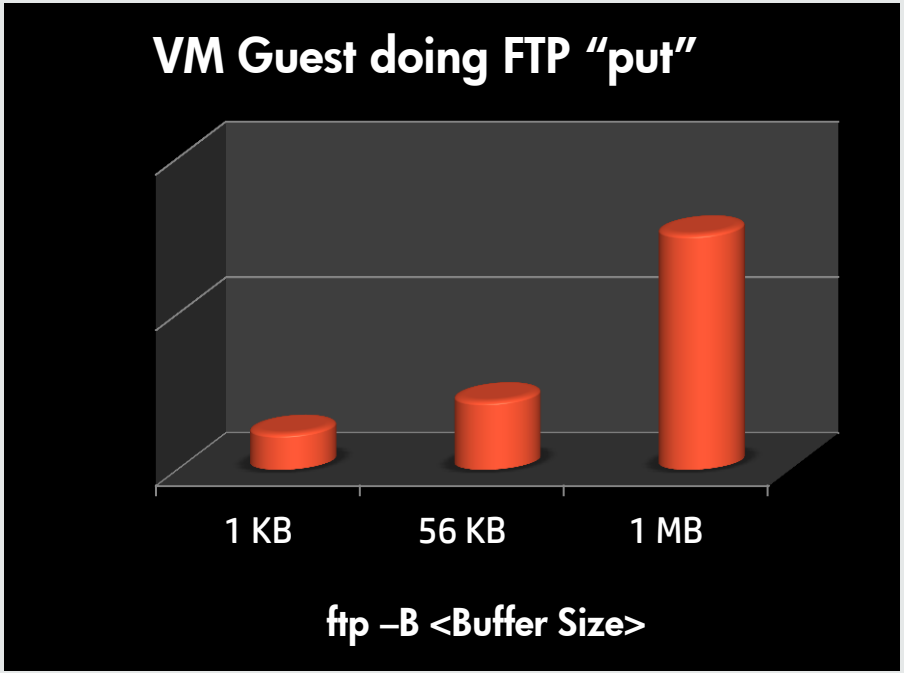

Some customer experiences continued.....

- Be careful with vCPU entitlement capping (as an example -e 20:50) since it will effect network performance

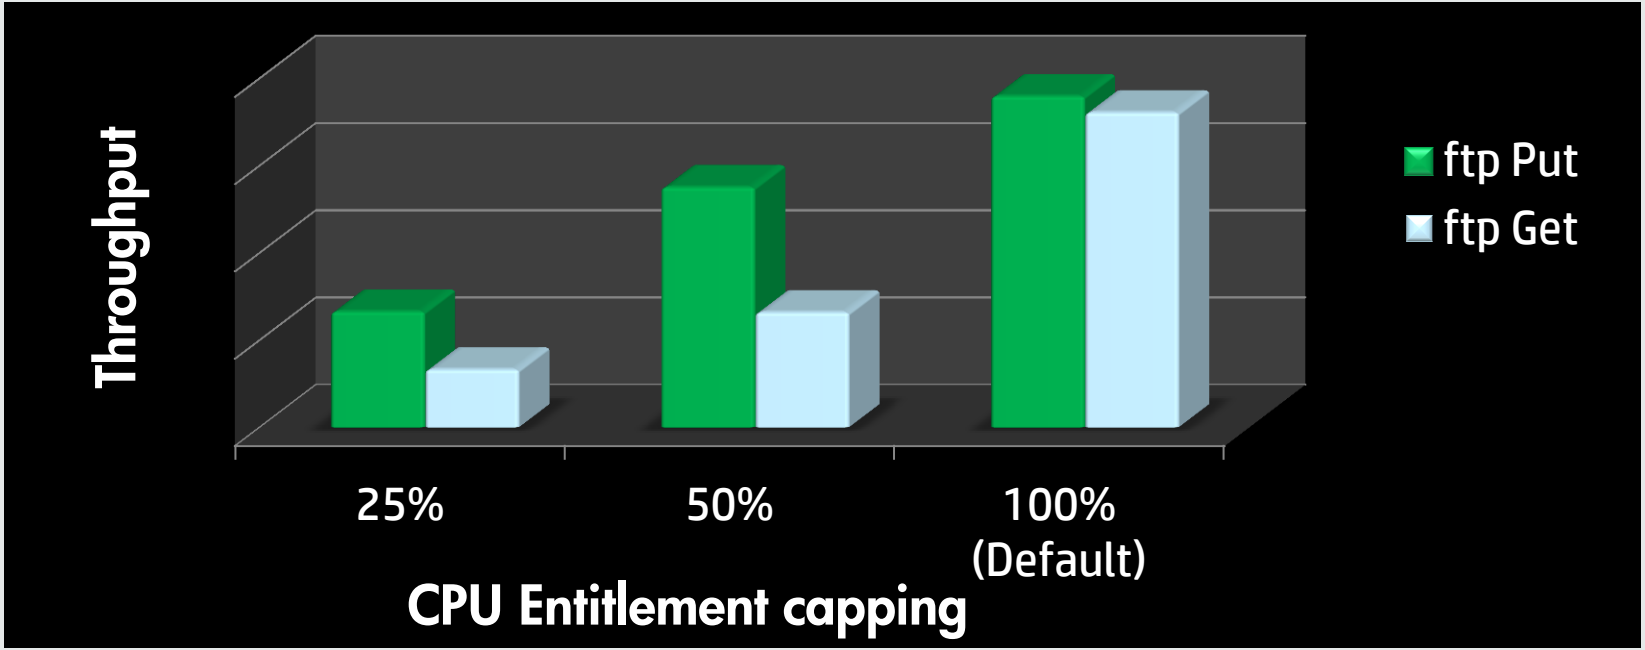

Futures and Q&A

HPVM 6.x support for OpenVMS guests

Native OpenVMS support for Poulson processor based systems like rx2800 i4 or BL8x0c i4 is needed to be supported to run as guest

NPIV and OpenVMS guests

N-Port ID Virtualization (NPIV)

Direct storage visibility for vPars and VMs

- **Virtual FC ports share a Physical FC port**
- **Improved storage security**
  - Provision storage directly to the vPar or VM
  - HPVM host/VSP doesn't see storage provisioned to vPar or VM
- **Enables enhanced storage features**
  - Applications requiring physical disk/array access
  - OpenVMS Multi-Pathing inside the VM

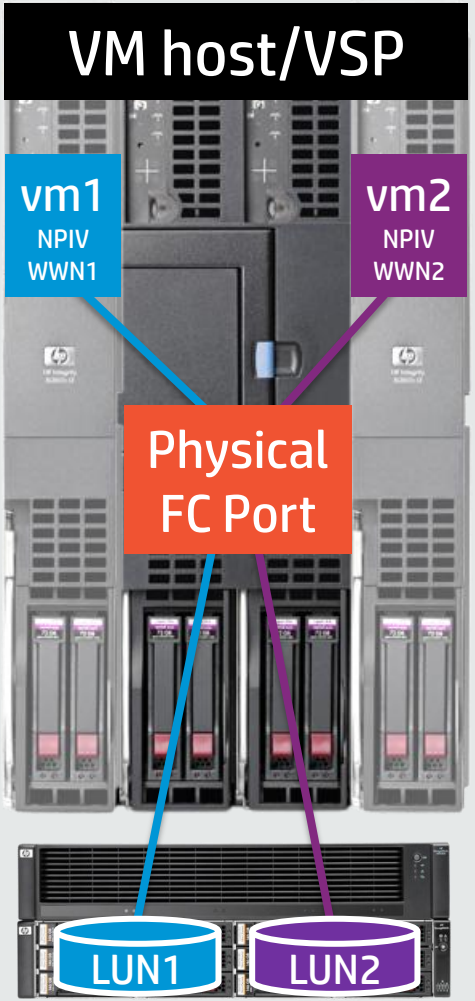

Globally Unique Identifier (GUID) Manager

Repository to allocate and manage unique World Wide Names for NPIV HBAs

Goals

- Avoid configuring VMs and vPars with duplicate NPIV WWNs
- Reduce risk of data corruption

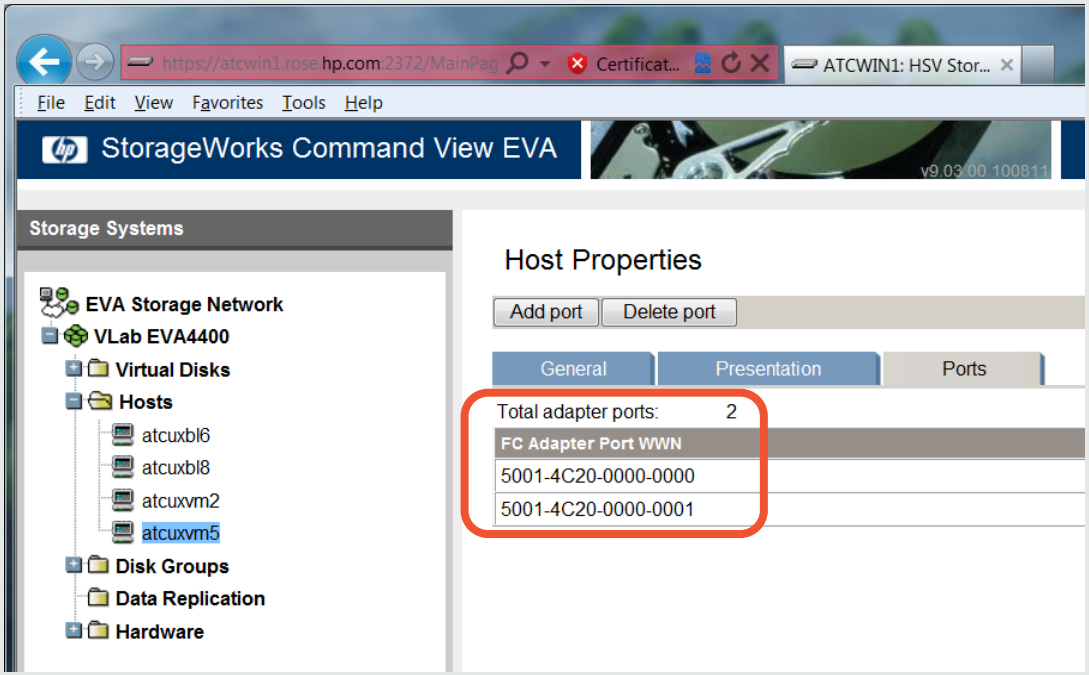

GUID Manager Integration with vPars & Integrity VM

Several vPars and VM commands are GUID Manager “aware”

- hpvmcreate(1M), vparcreate(1M)
- hpvmmodify(1M), vparmodify(1M)
- hpvmremove(1M), vparremove(1M)
- hpvmstatus(1M), vparstatus(1M)

Example: request WWNs from the GUID Manager and assign them to vHBAs

```
# hpvmcreate -p atcuxvm5 -a hba:avio_stor:,,,:npiv:/dev/fcd0
# hpvmmodify -p atcuxvm5 -a hba:avio_stor:,,,:npiv:/dev/fcd1
# hpvmstatus -v atcuxvm5

[IO Details]
hba:avio_stor:0,5,0x50014C2000000000,0x50014C2800000000:npiv:/dev/fcd0
hba:avio_stor:0,6,0x50014C2000000001,0x50014C2800000001:npiv:/dev/fcd1
```

NPIV Supported Limits for vPars & Integrity VM v6.1.5

Description	Limit
NPIV HBAs per vPar/VM	8
NPIV HBAs per Physical FC Port	32
LUN Paths per NPIV HBA	2048
Paths per NPIV LUN (Multipath)	8
LUNs per NPIV HBA	2048
NPIV LUNs per vPar/VM	2048
LUN Paths per vPar/VM	16384

NPIV Notes

- NPIV is now available for **Integrity VM v4.3 w/ PK2!**
- Only HP-UX **11i v3** VMs support NPIV
- OpenVMS 8.4 NPIV support in plan
- Install **latest AVIO** Storage drivers for NPIV fixes
- NPIV is supported on **QLogic** cards only. **Emulex** support is planned for a future OE.

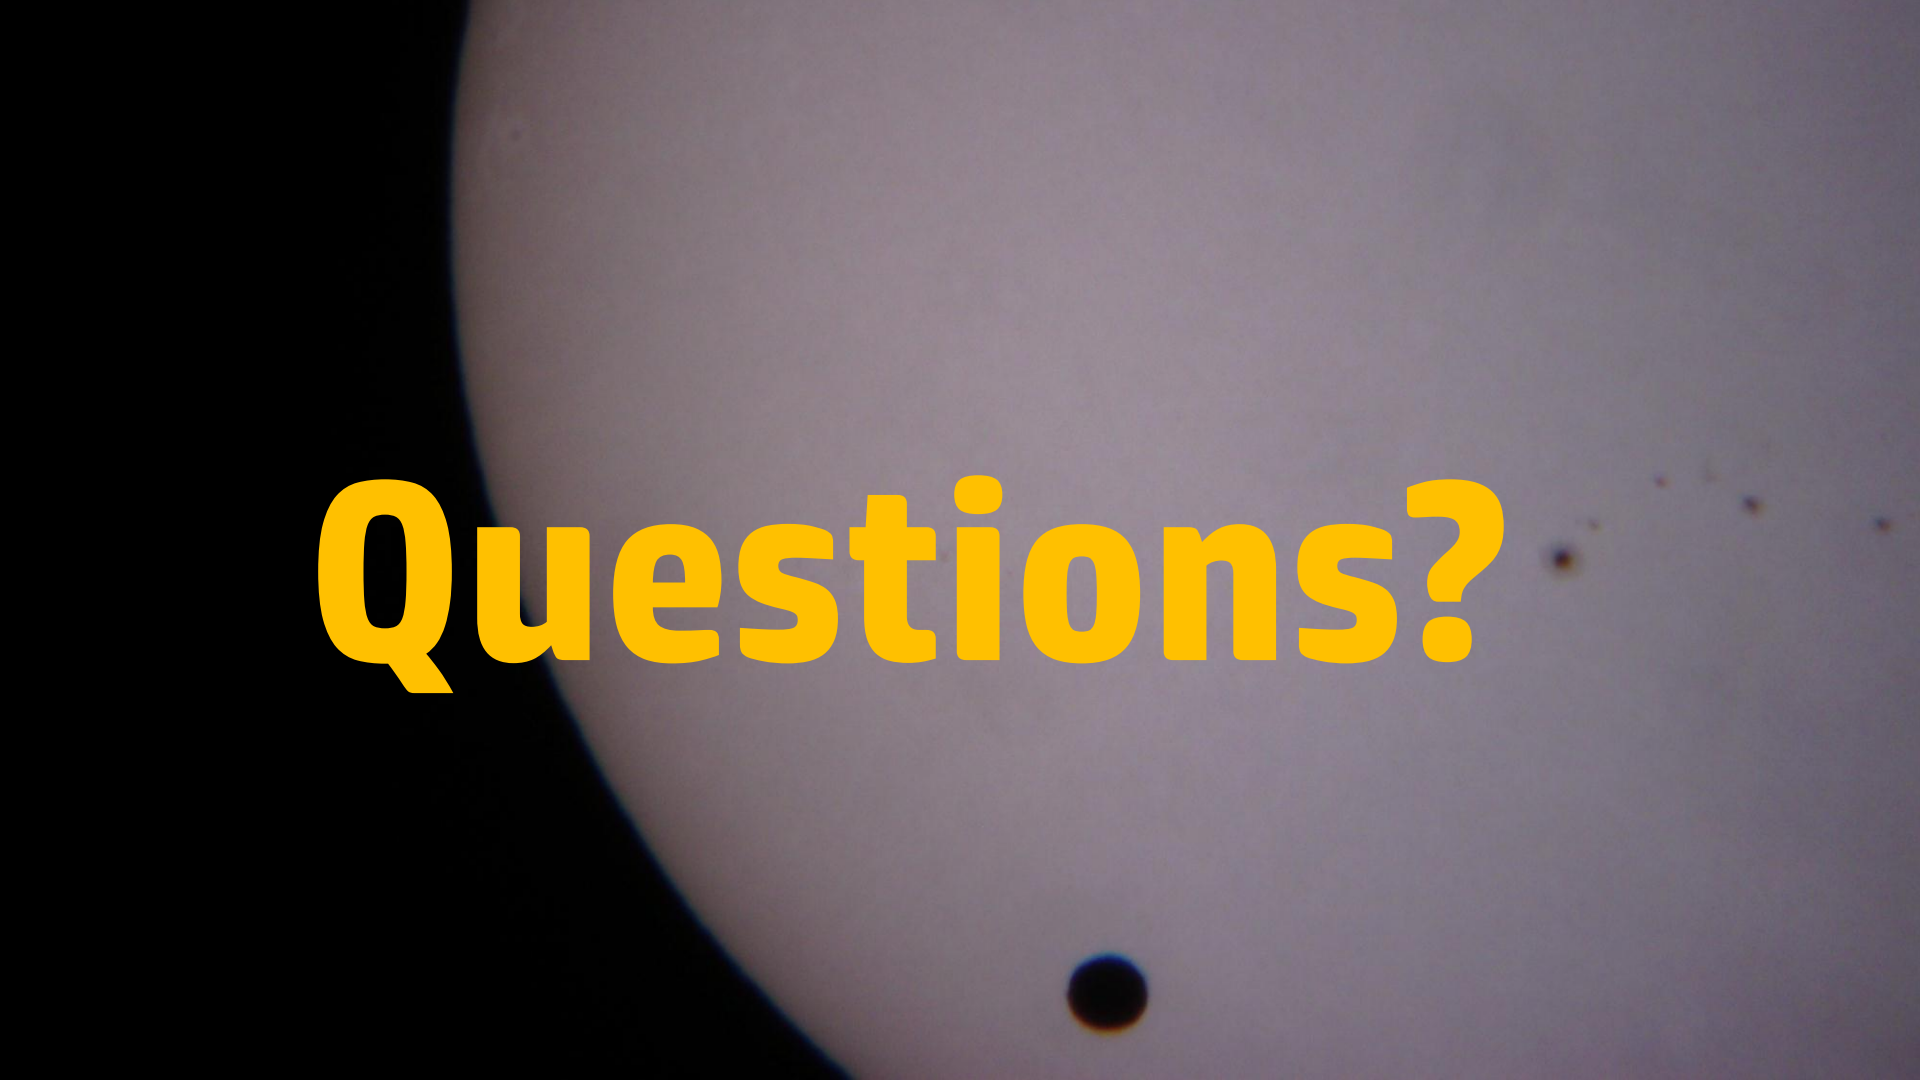
Questions?

The Venus passage 20120606 04:52 CET

Thank you