

# **Virtualisation**

## **M2 – P7**

### **2015**

**François Armand**

Université Paris Diderot

# A enda !1/2"


- Use #ases
- Multi-#ore # \$allen es
- Virtuali%ation &a' ono( )
- Virtuali%ation\* \$o+,
  - &rans-arent virtuali%ation\*
    - . ard+are assisted virtuali%ation
    - D)na( i# /inar) &ranslation
  - Para-virtuali%ation
- Virtuali%ation and Devi#es

# A enda !2/2"

- . ard+are 0volution
- . osted versus 1ative Virtuli%ation
  - 0' a( -les \* 2VM3 4en3 OS4
- Advan#ed Virtuali%ation #a-a5ilities
- I( -a#t on 6S and a--li#ations -a#7a in
- Standardi%ation


# Use 8ase\* Data 8enters

- Consolidate + or 7 loads runnin on inde-endent - \$) si#al ( a#\$ines on a sin le server} + \$ile ( aintainin inde-enden#e
- Sto- Ma#\$ine s-ra+l: 😊
- S-lit !virtuali%e" - \$) si#al servers into 9s( aller: Virtual Ma#\$ines
- Allo#ate VM d)na( i#all)

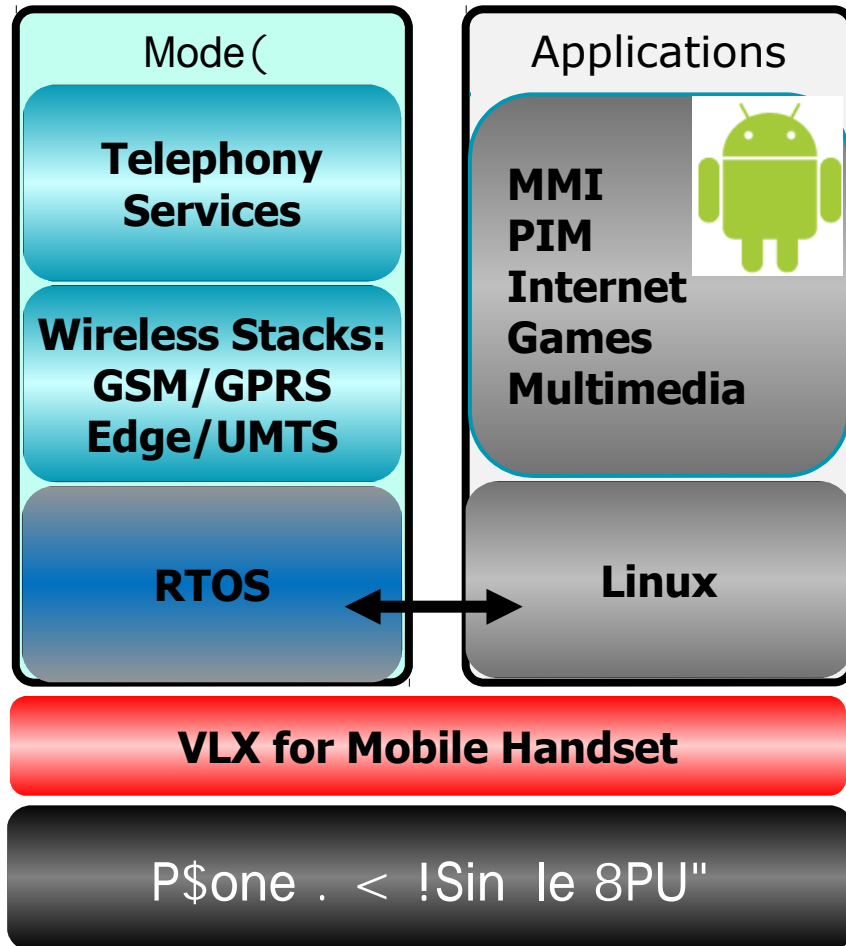


# Use 8base\* < or 7station

- Similar to data #enters
- Run (ore t\$an a  
single environ( ent at  
once on a single  
( a#\$ine



# Use Base\* Mobile . andsets



- Run =inu' a--li#ations on 5ase5and -ro#essor
- Re-use e' istin ( ode( so>t+are sta#7 +it\$ its R&6S !no # \$an es"
- Su--ort o> =inu' at a ( ini( al develo-( ent #ost
- 6-eratin S)ste( inde-enden#e >or >uture evolutions
- See also 1&& Do#o( o 6S&I

# Use 8ase\* Mo5ile . andsets

## ● Virtual=0 i' V=4 *!-revious slide"*

- 14P -lat>or( 3 2 A#er ( o5ile -\$ones




## ● VM+are Mo5ile Virtual Plat>or( !MVP"

- &ran o Virtual Pro#essor a#?uisition

## ● 6-en 2ernel =a5

- =; Mi#rovisor
- Used 5) @ual#o( ( and ot\$ers !. &83 MotorolaA"

# Virtualisation in . i \$-&\$rou \$-ut 1et+or7 0?ui- ( ent

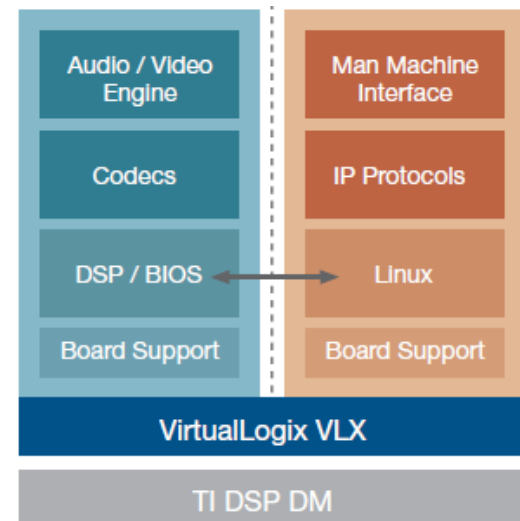




# Use #ase\*

## Virtualisation in Multi Media Devices

- Reduction of /6M or \$i \$-volume (e lo+-end -rodu#ts
  - No need for a General Purpose Processor
    - C 20 to 25 D /6M reduction
  - Run =inu' to et\$er +it\$ 6S su--ortin 8ode#s on a sin le &I DSP
  - =evera e =inu' environ( ent
  - Reuse e' istin DSP so>t+are



# 6t\$er Use 8ases

## ● Instru( entation? Auto( ation

- Run a R&6S and a BP6S >or Bra-\$i#al Inter>a#e

## ● Points o> Sales

- Run t\$e UI and t\$e se#ure transa#tion environ( ent on t\$e sa( e -ro#essor

## ● Mil / Aero

- Run se#urel) isolated / #erti>ied environ( ents si( ultaneousl)

## ● MoreA

# A enda !1/2"

- Use #ases
- Multi-#ore # \$allen es
- Virtuali%ation &a' ono( )
  - S)ste( level versus Pro#ess level virtuali%ation
  - 1ative versus . osted Virtuali%ation
  - &rans-arent virtuali%ation versus -ara-virtuali%ation
- 0( 5edded E Real-&i( e Virtuali%ation Re?uire( ents
- . ard+are 0volution

# Multi-core challenges

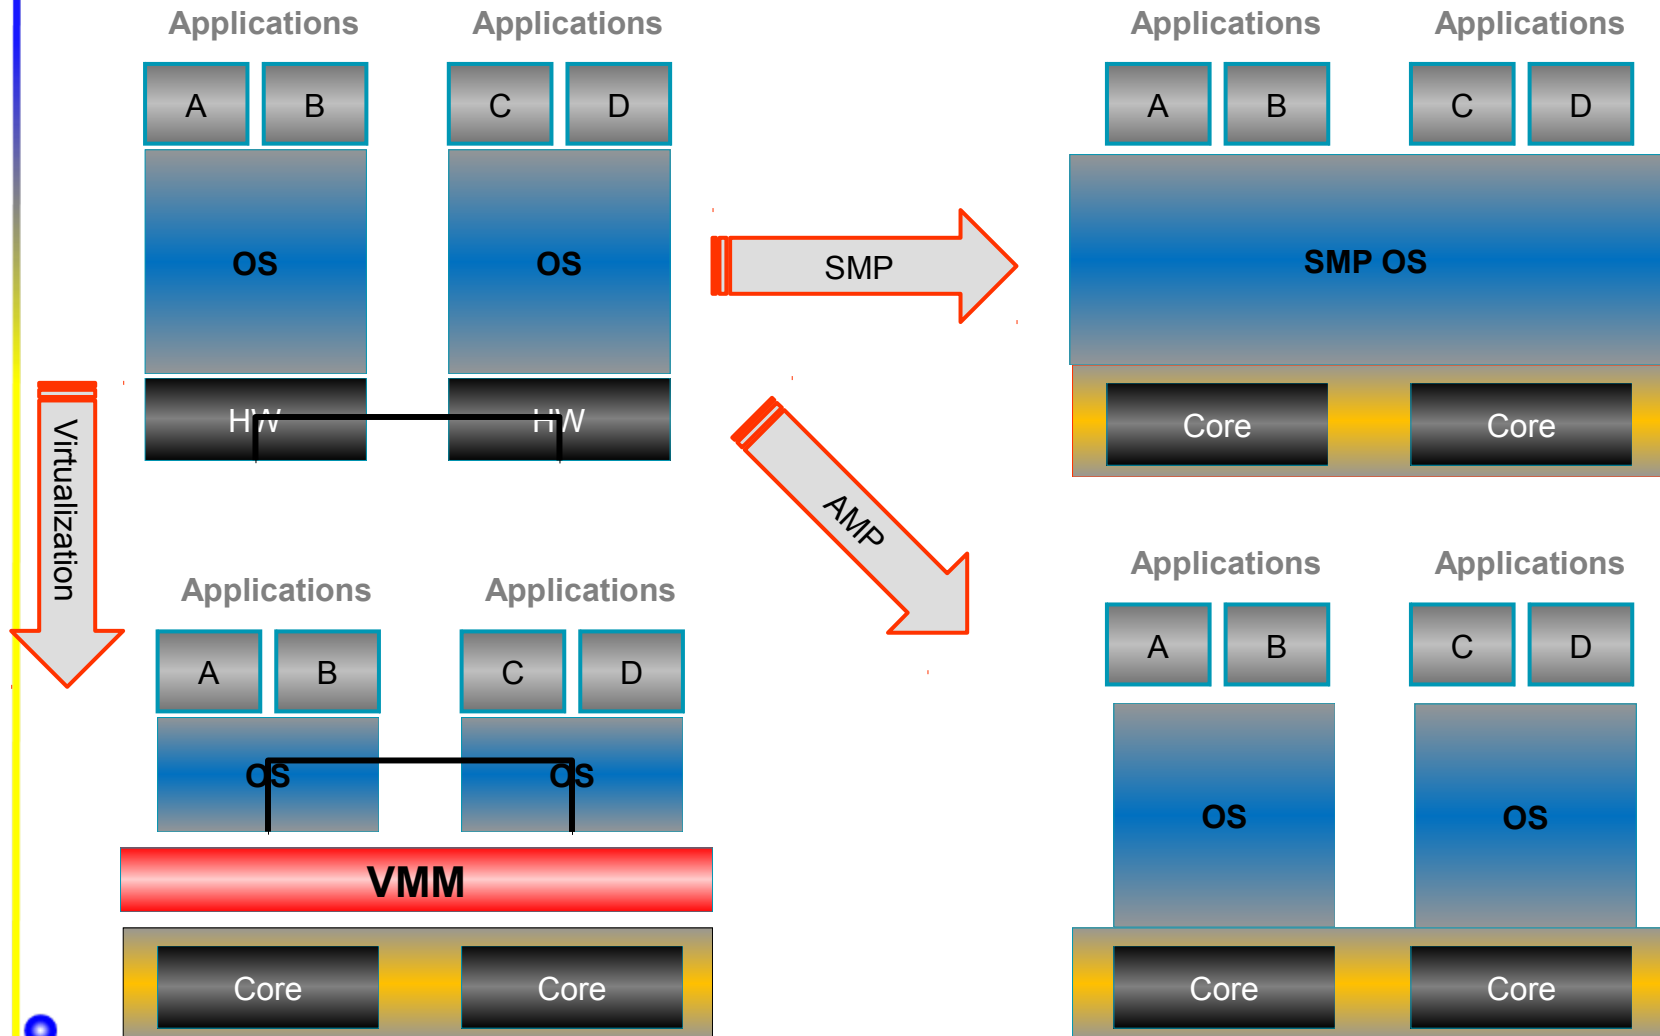
- Increased number of processors provided by servers but throughput per server density and (multi-core systems)
- Many legacy RISC have been designed based on uni-processor assumption
  - Scaling on (multi-core requires time)
  - "as seen" is still a major BIOS compatibility issues for 6S and Applications"
- Legacy RISC architectures designed for uni-processor as well as need to (move to true (multi-threaded + true parallelism) AF
  - Over (core density) ultimate than 6Ss

# Multi-processor systems

- Virtualisation enables to run multiple instances of a *single real-time* software load designed for uni-processor systems on a multi-processor
- Possibility to run in uni-processor Virtual Machines assigned to a shared CPU
- For BP6S\* run multiple SMP instances on a larger (single line) on a 1K cores (single line)
- One modification of the code to increase the code
- Availability provided by Virtualisation layer

# Virtualisation enables consolidation

- In addition to the core server hardware



# A enda !1/2"

- Use #ases
- Multi-#ore # \$allen es
- Virtuali%ation &a' ono( )
  - S)ste( level3 Pro#ess level virtuali%ation3 6S virtuali%ation
  - 1ative versus . osted Virtuali%ation
  - &rans-arent virtuali%ation versus -ara-virtuali%ation
- 0( 5edded E Real-&i( e Virtuali%ation Re?uire( ents
- . ard+are 0volution

# Virtualisation is + as getting interest





# Virtualisation, 6\$3 VirtualisationL

Virtual Networking? Intel VT? UML? IBM/VM?  
 VMware? Virtual Server? Transitive QuickTransit?  
 Java, JVM? Application Virtualization? Pascal Pcode?  
 AMD SVM? Platespin?  
 TransMeta Crusoe? Virtual Solutions?  
 SIMICS? Softricity? QEMU? KVM?  
 Dynamo?  
 Virtual Reality? VirtualIPC? Virtual Storage?  
 FX!32?

# Different classes of virtualisation

## • Abstraction

- Make 1 resource appear as 1 clusters / vms/disks

## • Partition / Relocation \*

- Make a resource appear as 1 VMs/disks
- Make no(5ined + it\$ 9a relation: !disks"

## • Translation ( emulation"

- Make 4 appear as M !so( etic( es 4 is identi#al to M"
- Make no(5ined + it\$ 9-partition:

Mostly interested in 9-artition:


# . istor)

- A--eared durin t\$e K0ls I/M
- Po-ular to-i# durin t\$e K0ls and t\$e 70ls
- 9Surve) o> Virtual Ma#\$ine Resear#\$: !Bold5er 3 1N7; "
  - 9Virtual Ma#\$ines \$ave >inall) arrivedF Dis( issed >or a nu( 5er o> )ears as ( erel) a#ade( i# #uriosities3 t\$e) are no+ seen as #ost-e>>e#tive te#\$ni?ues >or or ani%in #o( -uter s)ste( s resour#es to -rovide e' traordinar) s)ste( >le' i5ilit) and su--ort >or #ertain uni?ue a--li#ations:
- 72 -a-ers -u5lis\$ed >ro( 1NKK to 1N70



# Virtual Machine Interfaces

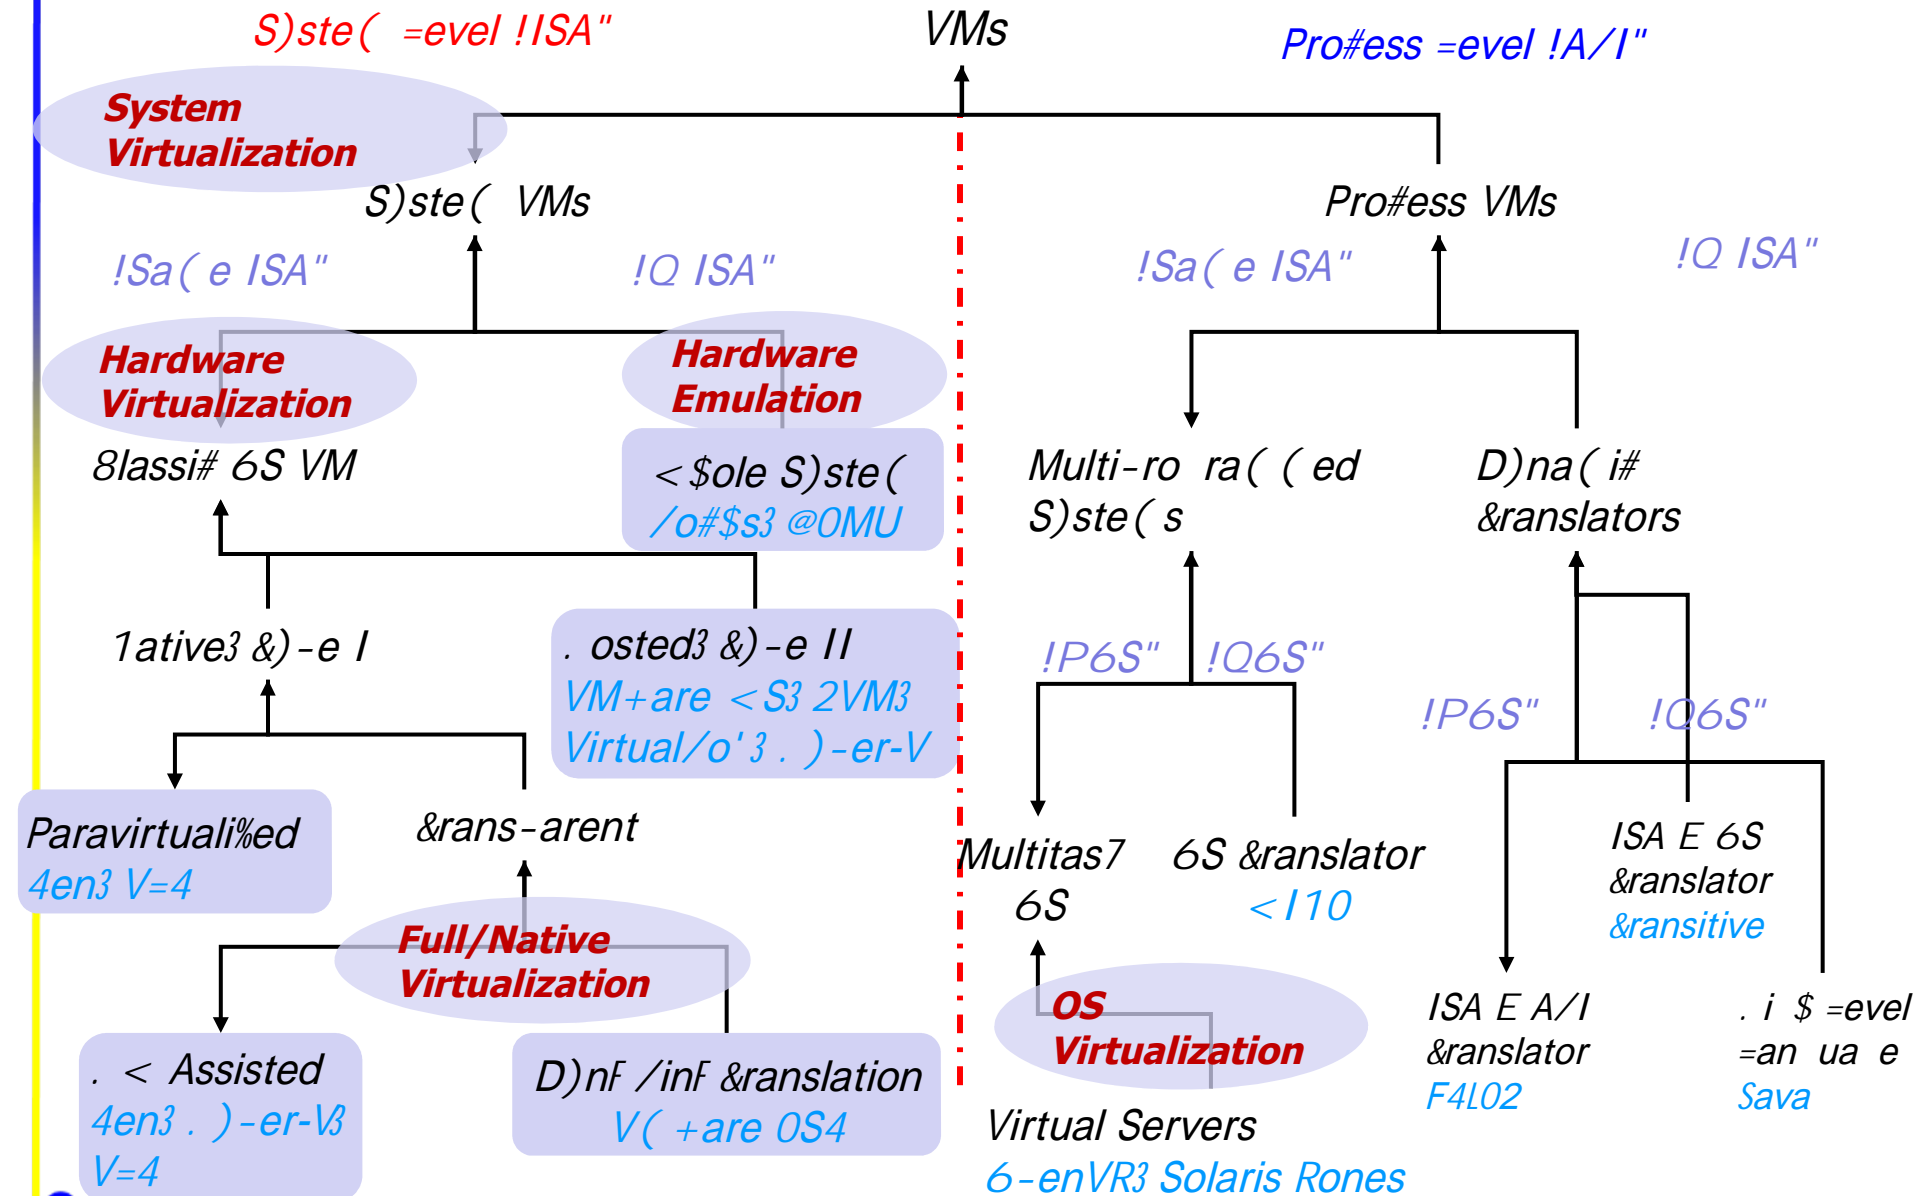
- System built on hardware architectures
- 2 (main interfaces) ISA / ABI



- System level VM  
- provide an ISA interface

- Process level VM  
- provide an ABI interface

# &a' ono( ) !derived >ro( OF S( it\$ E 1air"



# 6S Virtualization or Virtual Servers

## ● What is it?


- Virtualization / 6-enVR

## ● Pros\*

- Supports distribution heterogeneity
- High utilization of hardware
  - Low (cost) and error (management overhead)
- Scales to (any) instances

## ● Cons\*

- Single 6S instance is not (a) joint operation
- Modified 6S is intrusive\* as to pollute 6S evolution
- Does not support 6S heterogeneity! i.e. no R&6S / BP6S innovation



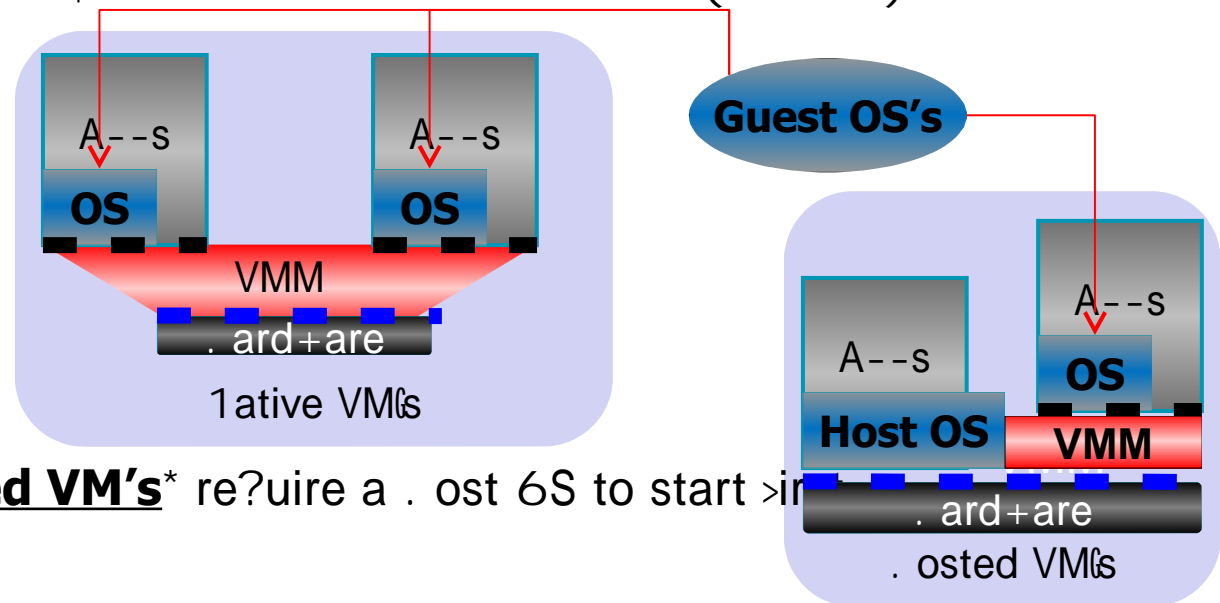
# 6S Virtualisation \* 0' a( -les

- AI4 < or 7in Partitions
- Solaris Zones & Containers
- FreeBSD Jail
- =inu' -Vserver & FreeVPS
- Virtuo%%o / 6-enVR
- i8ore Virtual Accounts ! < indo + s" ,



# 8classi# 6S VMs

- On a 5le to run (ulti-le Jinde-endent- 6Ss si(ultaneousl) on t\$e sa(e -ro#essor! **guest os**"3 ea#\$ in its o+n **"Virtual Machine:**
- &+o (ain a--roa#s\$es\*
  - **Native VM's\*** Introdu#e a so>t+are la)er 5et+een t\$e \$ard+are and t\$e 6S\* Virtual Ma#\$ine Monitor !VMM" or 9/are (etal: . )-ervisor



- **Hosted VM's\*** re?uire a . ost 6S to start >ir

# 8lassi# 6S VMs \* Issues

## ● Run (ulti-le 6Ss

- 6S desi ned assu( in t\$e) are t\$e onl) so>t+are #ontrollin t\$e -)\$si#al resour#es o> t\$e ( a#\$ineF

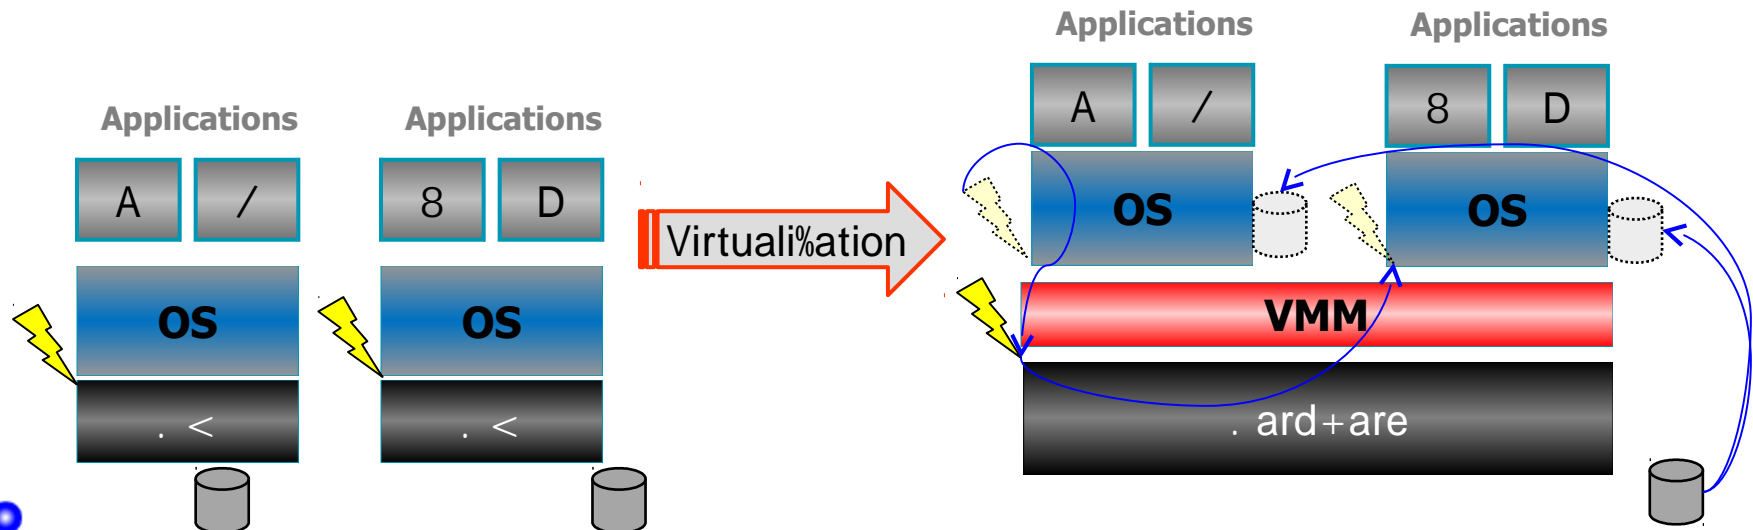
## ● 8on#urrent a##ess to -)\$si#al resour#es \*

- 8PU PU s#\$eduler3 Me( or) PU -artition3 MMU , Devi#es ,

## ● 1eed to dete#t and resolve #on>li#ts su#\$ as

- ( as7in interrups3 initiatin an I/O3 or -ro ra( ( in t\$e MMU

## ● Provide t\$e e' -e#ted 5e\$avior +it\$in t\$e Virtual Ma#\$ineF



# 8lassi# 6S VMs \* Issues

● Detect sensitive instructions and 9>a7e: t\$( F

● &rans-arent +a)s

- Bood \$ard+are su--ort\*
  - 0' e#ute 6S at lo+er \$ard+are -rivile e !ifeF \* user ( ode" to tra- su#\$ instru#tions u-on e' e#ution
- <ea7 \$ard+are su--ort\*
  - 8ir#u( vent non tra--in instru#tions

● 1on &rans-arent +a)

- Modi>) 6S a\$ead o> ti( e

# Class# 6S VMs

## Goal\*

- Run the binary) guest 6S + it's lower hardware  
- privilege as the supervisor (code) it's been designed  
for

## Means\*

- Trans-arent Virtualisation\* *!>ull or native"*
  - 1o (odi>i#ation o> t\$e 6S i( a e
  - Full) Virtuali%a5le Pro#essors !V&-' 3 AMD-V3 I/M PP8"
  - D)na( i# /inar) &ranslation !VM+are"
- Para-virtuali%ation\*
  - Modi>i#ation o> so( e o> t\$e 6S . A= sour#e >iles
  - *!#an 5e seen as a -ort to a ne+ -ro#essor ver) si( ilar to t\$e real one"*

# Full Virtualisable Processors

## ● Issue\*

- Is there any instructions + those 5e\$avior di>>ers or should differ in these 2 (odes,


## ● Sensitive Instructions\*

- &ra- to su-ervisor (ode + \$en e' e#uted >ro( user (ode \* 62
  - 0' \* #li3 sti !Intel 'VK" tra- + \$en e' e#uted >ro( user-(ode
- 1o-o-\* not 62\*
  - 0' \* P6PF !Intel 'VK"\* Interru-t-ena5le >la re( ains una>>e#ted in user (ode
- Bet s)ste( / \$ard+are status\* not 62
  - 0' \* Read 8R0 !Intel 'VK" +ould return true - \$)si#al in>o3 instead o> virtuali%ed in>oF

# Hard+are assisted virtualization

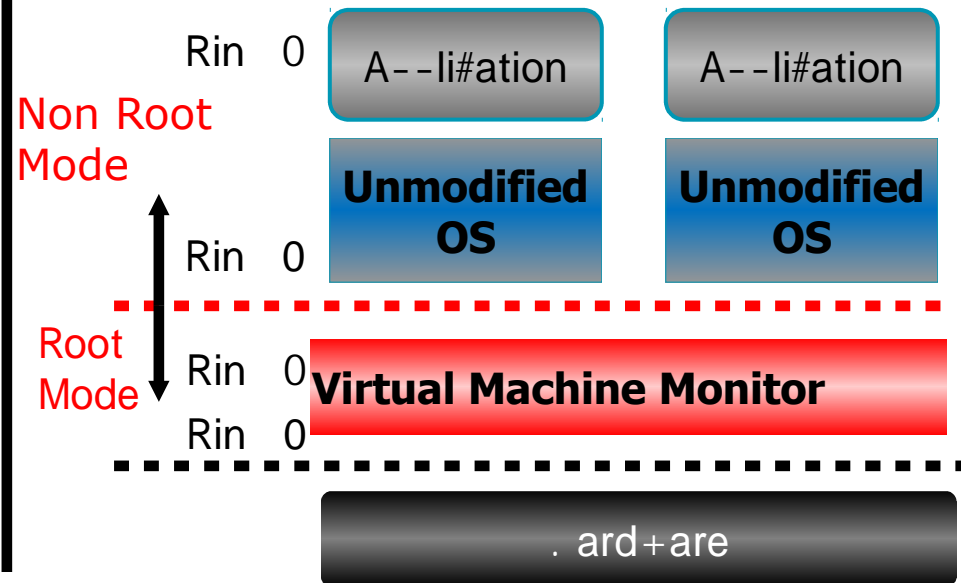
## Non Hardware assisted

- Binary translation
- Para-virtualization
- Software-only Source code




## Virtualization Technology

- 2 1e+ e' e#ution (odes3
- And a set o> \$ard+are-  
5ased tri ers to s+it#\$  
5et+een t\$e(



# . ard+are assisted virtuali%ation

- &\$e 6S \$as not a##ess to t\$e >ull ISA + \$en  
runnin in a VMF



# D)na( i# /inar) &ranslation

- 1on . ard+are assisted virtuali%ation\* !' VK3A"
- Run un( odi>ied 5inar) uest in a less--rivile ed ( ode !e' \*  
rin 0 instead o> rin 0 on ' VK"
- D)na( i# 5inar) translation !eF F\* VM+are on ' VK"
  - VMM d)na( i#all) 9re-+rites: -rivile ed instru#tions +\$i#\$ +ould  
5e silentl) e' e#uted in user ( odeF
  - 6n de( and3 #a#\$ed
  - Me( or) #onsu( -tion i( -a#t
  - &i( in deter( inis( i( -a#t
- &rans-arent 5ut #o( -le' solution
  - #o( -le' it) resides in t\$e VMM



# Paravirtualisation

- Modified guest OS
  - Run guest OS in less privileged mode
    - 0 \* ring 1 instead of ring 0 on x86
    - 0 \* user mode instead of supervisor mode on ARM
- Better performance (an extra level of translation, but intrusive)
  - Solution of sharing hardware support (PP83 ARM) and OS modifications are possible

# Virtualisation and devices

## Shared devices

- Accessed by (more than) one VM
- OS partitions are not
- OS network interface in /dev/vnetX set between virtual and physical

## On shared devices

- Devices used exclusively by a single VM
- OS interface

## Virtualised VMM

## Virtualised within a dedicated VM


- Do(03 Do(1/6 in 4en3 9An): VM in V=4

## Direct physical device access to VM

- VMM P8I support / extensions 3 VMD@3 / V=43 A

# Virtualisation and Devices

- Different ways to provide access to devices\*
- Trans-arent I/Os or 9-ara-virtuali%ed: I/Os
  - Profs and 8onfs in 5ot\$ #ases



# Virtualisation and Devices

## !8ont@d"


### • /etter \$ard+are su--ort\*

- P8I SRI6V3 MRI6V3
- Intel V&-d3
- S-e#i>i# #ontrollers !ef F\* VMD@"

### • 6r S-e#i>i# VMM i( -le( entations

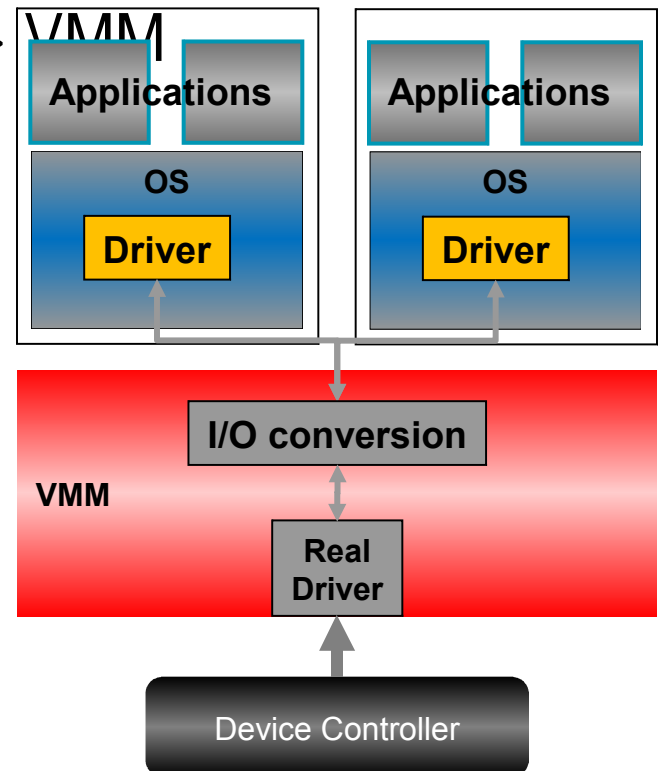
- V=4
- SU1 I/6 =D6M !li( ited nu( 5er"

### • Un( odi>ied drivers3 5etter -er>or( an#e




# S\$arin Devi#es

- S\$ared devi#es are a #on#ern >or >ailure resilien#e
- S\$ared devi#es -rovided 5) VMM\*
  - Failure o> driver i( -lies >ailure o>
  - And >ailures o> all VMs




# S\$arin Devi#es

- S\$arin -rovided 5) a VM3 t\$rou \$ 5a#7-end driver
- Failure o> driver PU >ailure o> VM
- 6nl) #lient VMs are i( -a#ted
- Restart under #ondition



# Not Sharing Devices

- Multiple VMs would solve the de-duplication issue
- At the host OS (core devices)



# A enda !1/2"

- Use #ases
- Multi-#ore # \$allen es
- Virtuali%ation &a' ono( )
  - S)ste( level3 Pro#ess level virtuali%ation3 6S virtuali%ation
  - 1ative versus . osted Virtuali%ation
  - &rans-arent virtuali%ation versus -ara-virtuali%ation
- 0( 5edded E Real-&i( e Virtuali%ation Re?uire( ents
- . ard+are 0volution



# Hard+are Evolution

## 4VK

- Multi-core
- Hard+are assisted virtualisation\*
  - vt-' 3 vt-i3 sv( 3
  - Pa e ta5le su--ort
  - I/6s\* vt-d3 io( ( u3 !trans-arent virtuali%ed DMA"
  - H P8I\* SRI6V3 MRI6VI3 !VMD@"

## Power Architecture for 5edded !@orl@"

- Multi-core -ro#essors
- 0 rin s\* user3 su-ervisor3 \$)-ervisor
- Address s-a#es ta ed +it\$ ID
- I/6 su--ort via PAMU !si( ilar to vt-d / io( ( u"

## Power for Servers

# . ard+are 0volution

## ● ARM

- Road( a- to ( ulti-#ore
- 1o #urrent su--ort >or virtuali%ation
- . o+ever3 &rustRone -er( its to isolate 2  
environ( ents

## ● SPAR8

- 1ia ara\* ( ulti-#ore / ( ulti-t\$readed
- . ard+are/Fir( +are su--ort >or -ara-virtuali%ation

# Agenda !2/2"

- Desktop Data Center solutions

## • Server Consolidation solutions\*

- Hybrid Storage Solutions
- Asynchronous Multi-Processing
- Micro-kernel based solutions
- Consolidated Real-time - servers

## • Miscellaneous

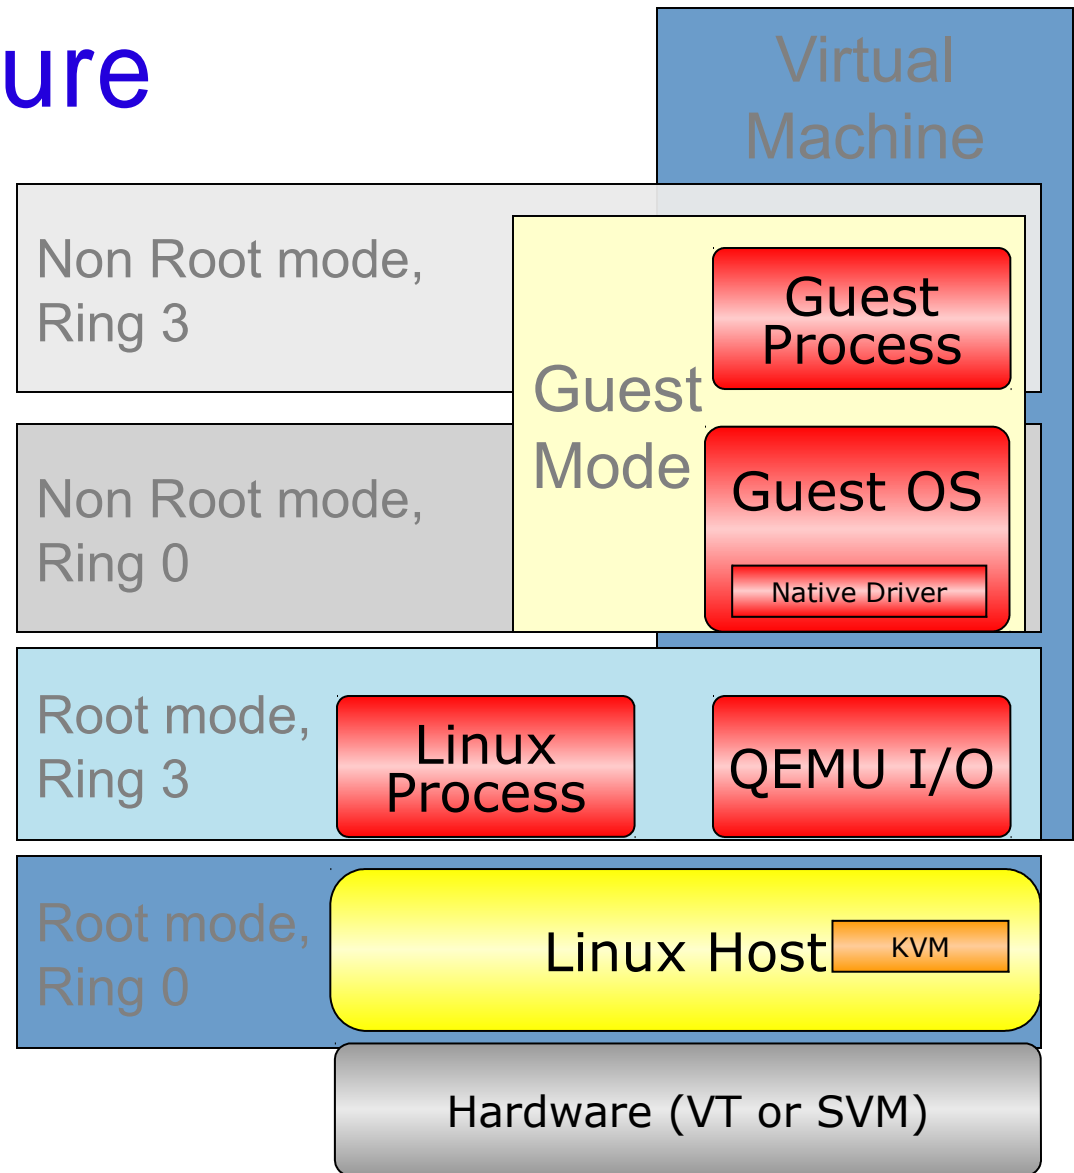
- Tools
- Availability
- Standardisation

# . osted Virtuali%ation\*

- Start a \$ost 6S >irst
- D)na( i#all) e' tends t\$e 6S +it\$ 7ernel (odule and a--li#ations -rovidin virtuali%ation
- Virtuali%ation relies on . ost 6S servi#es
- &)-i#all) 5inar) Buest 6S su--orted !&rans-arent Virtuali%ation"
- 1o isolation 5et+een . ost 6S and VMM
- 0' a( -les\*
  - VM+are < S3 SVista3 Parallels3 Virtual/o' 3 . )-er-V3 2VM
  - HUser-Mode =inu' I

# 2VM Architecture

- Guest I/O are translated to VM and redirected to MMIO
- VMs are run and scheduled as user-space processes



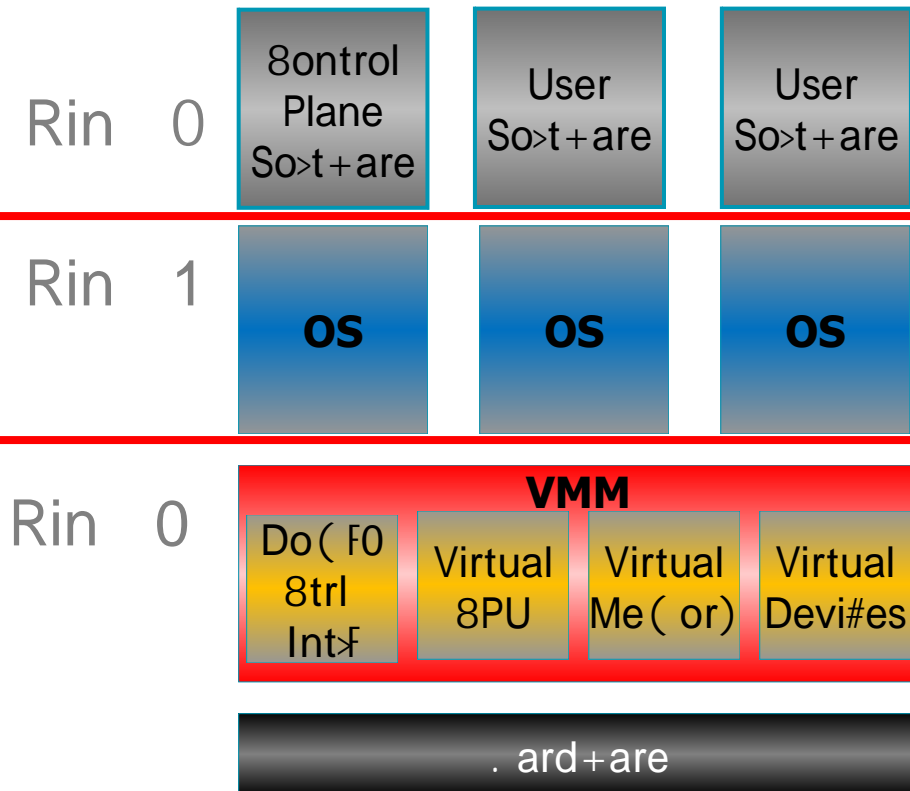
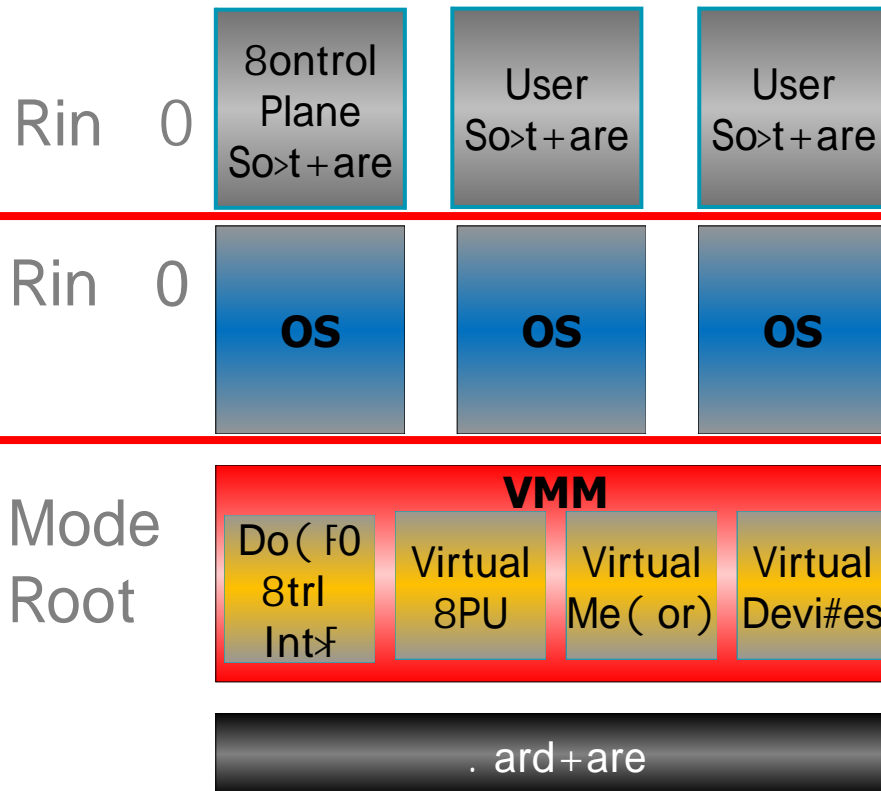
# 1ative Virtuali%ation

- Start a \$) -ervisor / VMM on 5are ( etal3
- &\$en3 t) -i#all) start a 9#ontrol do( ain: !4en" or a Servi#e  
8onsole !VM+are"
- 8an t\$en Jd)na( i#all) - #reate VMs
- Devi#e Virtuali%ation ( a) 5e -rovided
  - /) VMM !VM+are OS4"
  - /) ot\$er VMs !4en"
- 0' a( -les\*
  - VM+are OS4!i"3 4en3
  - V=4

# /are Metal . )-ervisors\* 401

. ard+are Assisted  
&rans-arent Virtuali%ation !V&"

Para-Virtuali%ation  
Pre-V& ' VK

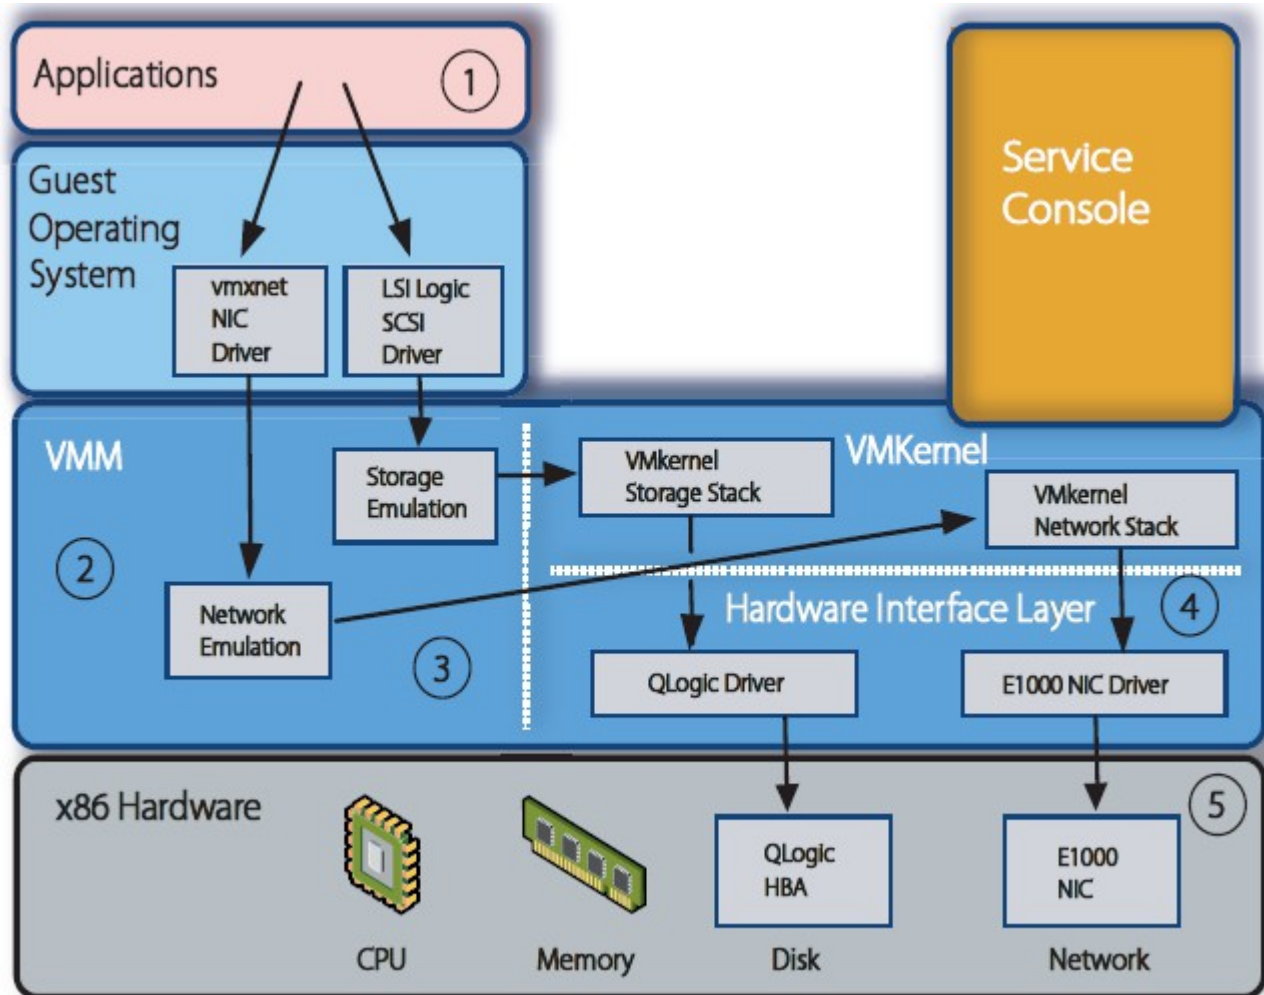


# 401\*

- Most dedicated to Data Centers\* - VK3 Itanium and I/M PP8
  - Paravirtualisation or . < assisted virtualisation eliminates the need for binary translation
  - Supports x86-64 Free/SD3 Plan3 A
  - Supports x86-64 on Intel V&3 AMD SVMF
  - Tools for data center (analysis, configuration, provisioning, etc.)
- Limited support for other processors
  - ARM in progress
- No support for embedded configurations
  - All VMs created on DoC 0
  - Need to wait for hardware initialization on DoC 0



# Device Emulation & Translation\* VMware ES4



# /are-Metal . )-ervisors\* OS4

- D)na( i# /inar) &ranslation !6S un( odi>ied in rin 0"
  - So >ar3 even on !02 5its" Intel V&-i3 AMD SVM
  - . ard+are virtuali%ation used on K; 5its ( a#\$ines
- 9/i :\*
  - In#ludes Virtual Me( or) Mana e( ent3 File s)ste( !VMFS"3 1et+or7 sta#73 driversA
- Me( or) Mana e( ent\*
  - P\$)si#al Me( or) ( a) 5e over #o( ( itted !-a in out uest 6S6s"
  - D)na( i# #ontent 5ased -a e s\$ar in 3 Me( or) 5alloonin
- Dedi#ated to Data 8enters !' VK / Itaniu( onl)"



# Live Migration of VM

- Migrate a running VM to another server to avoid downtime + it's out of service! (initially do not migrate) =oad / alan#in 3 Preventive (maintenance) of hardware #0(-attrib) #constraints

# Virtualisation and #clusters

## ● Central ( ana e( ent tools

- 0' \* VM+are virtual 8enter3 si( ilar tools >or ot\$ers

## ● Resource #ontrols on VMs

- . ierar#\$i#al resour#e -ools !( e( or)3 8PU3A"
  - Ma) s-an -\$)si#al ( a#\$ine 5oundaries
- Si( ilar to resour#e ( ana e( ent >or -ro#ess

## ● Auto( ati# restart o> >ailed VMs A so( e+ \$ere

- 9. i \$-Availa5ilit): solutions

## ● =ive Mi ration

## ● 9Sna-s\$ot li5raries:


# Virtualisation leads to Virtualisation

- Deliver Virtualisation to its containing Virtual Machine
  - 1st version issue and core
  - 1st (initial) libraries tools and other applications
  - Virtualisation already installed and configured
  - VMs are as a large directory of Virtual Machines or VMPlayer
  - Benefit of physical devices dependencies deal only with virtual devices

# 1e+ +a)s to loo7 at a--li#ations

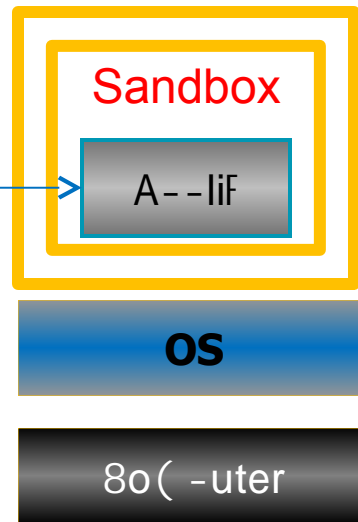
Virtual  
A--lian#es

Virtuali%ed  
A--li#ations





Promoted by VMware

Pushed/ pulled  
from server



Provided by  
MSFT/Softgrid



Provided by Citrix

# Evolution of Virtual Appliances

## Access to devices\* virtualised

- Older through a dedicated VM (e.g. 3 or 5) VMM (OS4)

