Virtualization

&& (and)

|| (or)

^^ (exor)

Security

Usual OS diagram

- Traditional kernel:
  - Control HW access
    - CPU (timesharing)
    - (virtual) memory
    - Disk (File system)
    - Devices (net, console, video, audio etc)
  - IPC/signals
  - Low-level network
  - Identity mgmt (UID/GID)

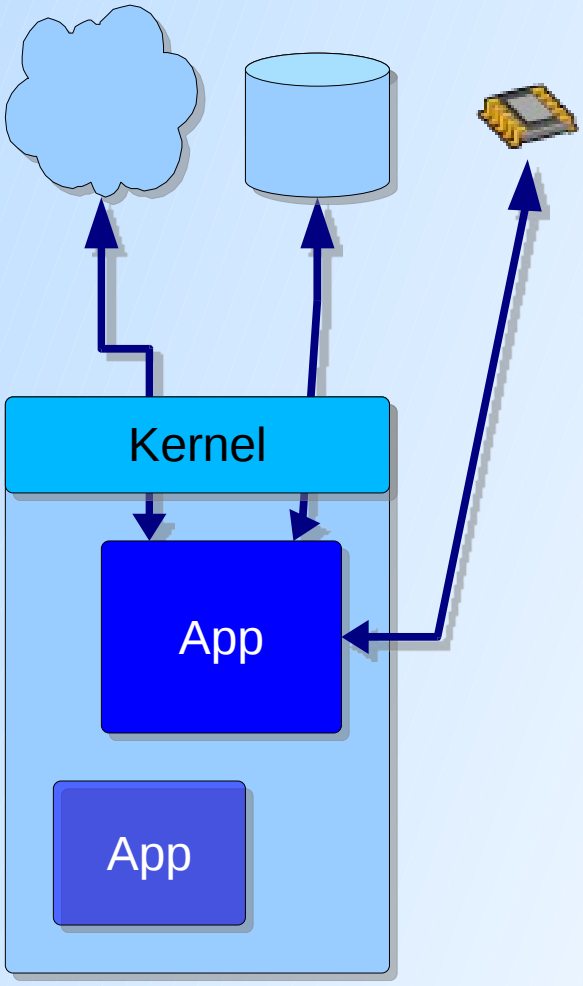

VM-in-a-box

Add:

- Hypervisor (HW access)
- Network (virtual)
- “Controlling” dom0

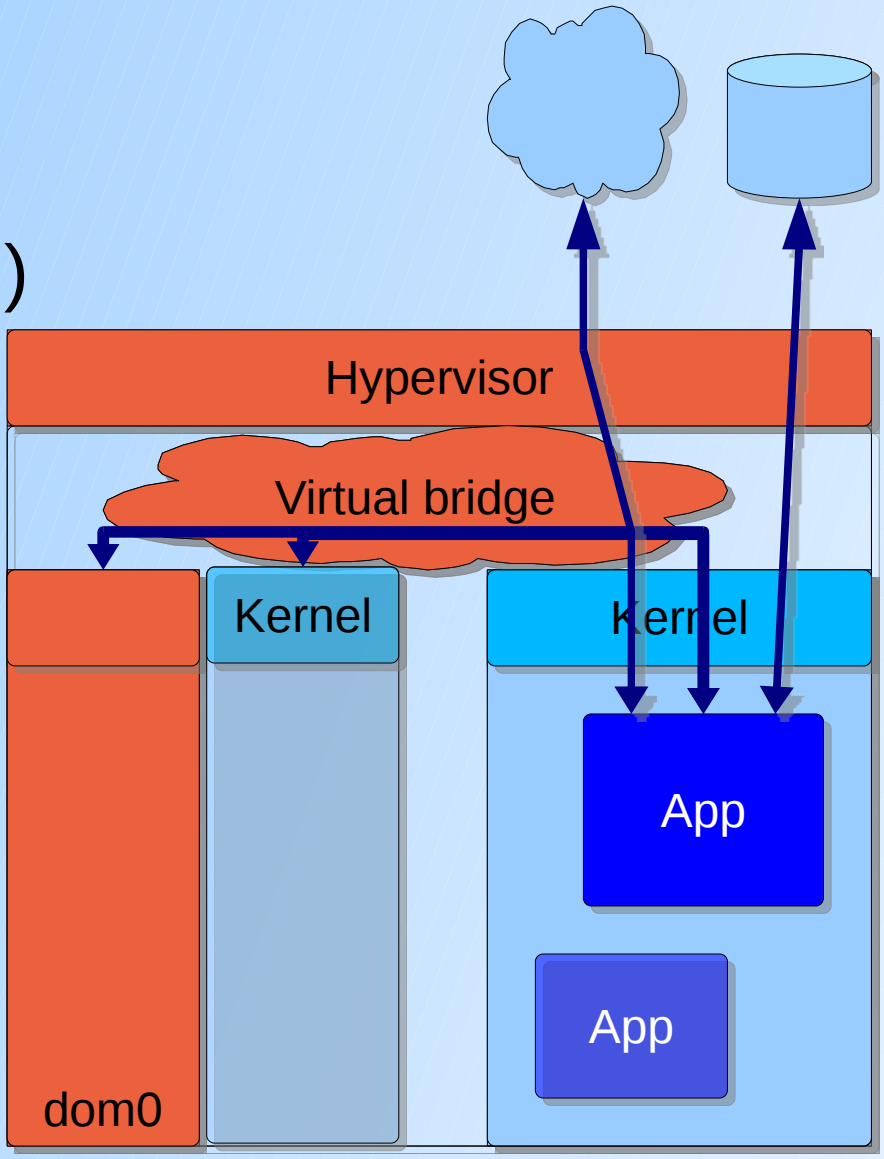

Wait, there is more..

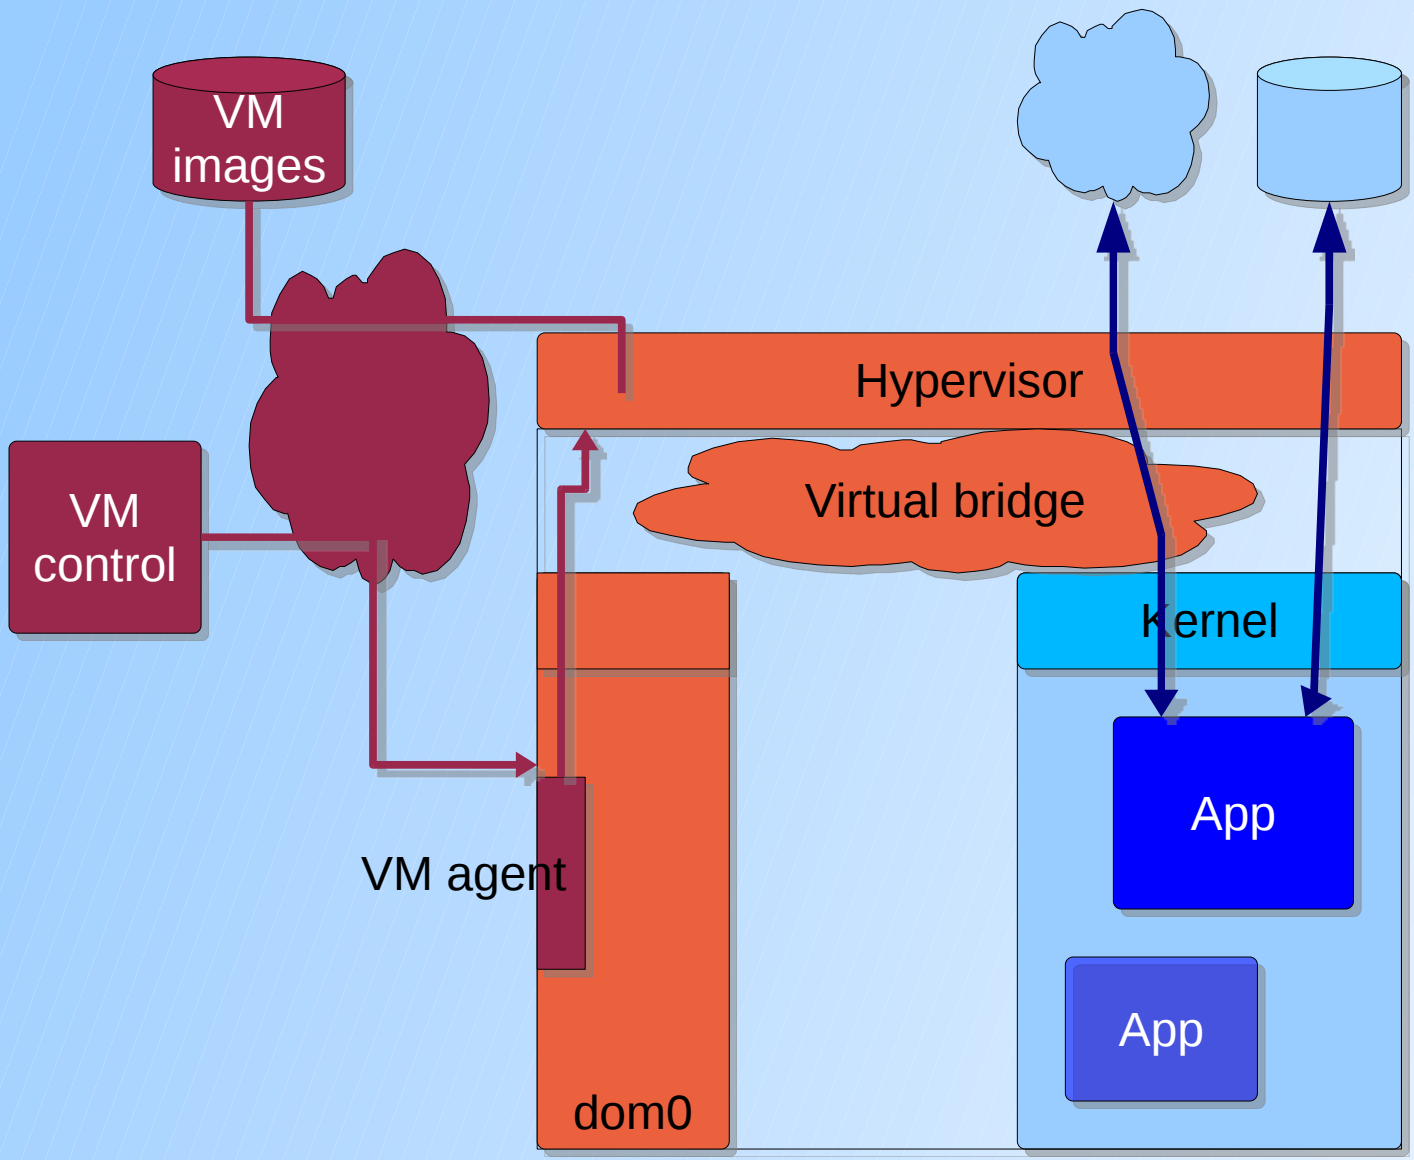

.. and more

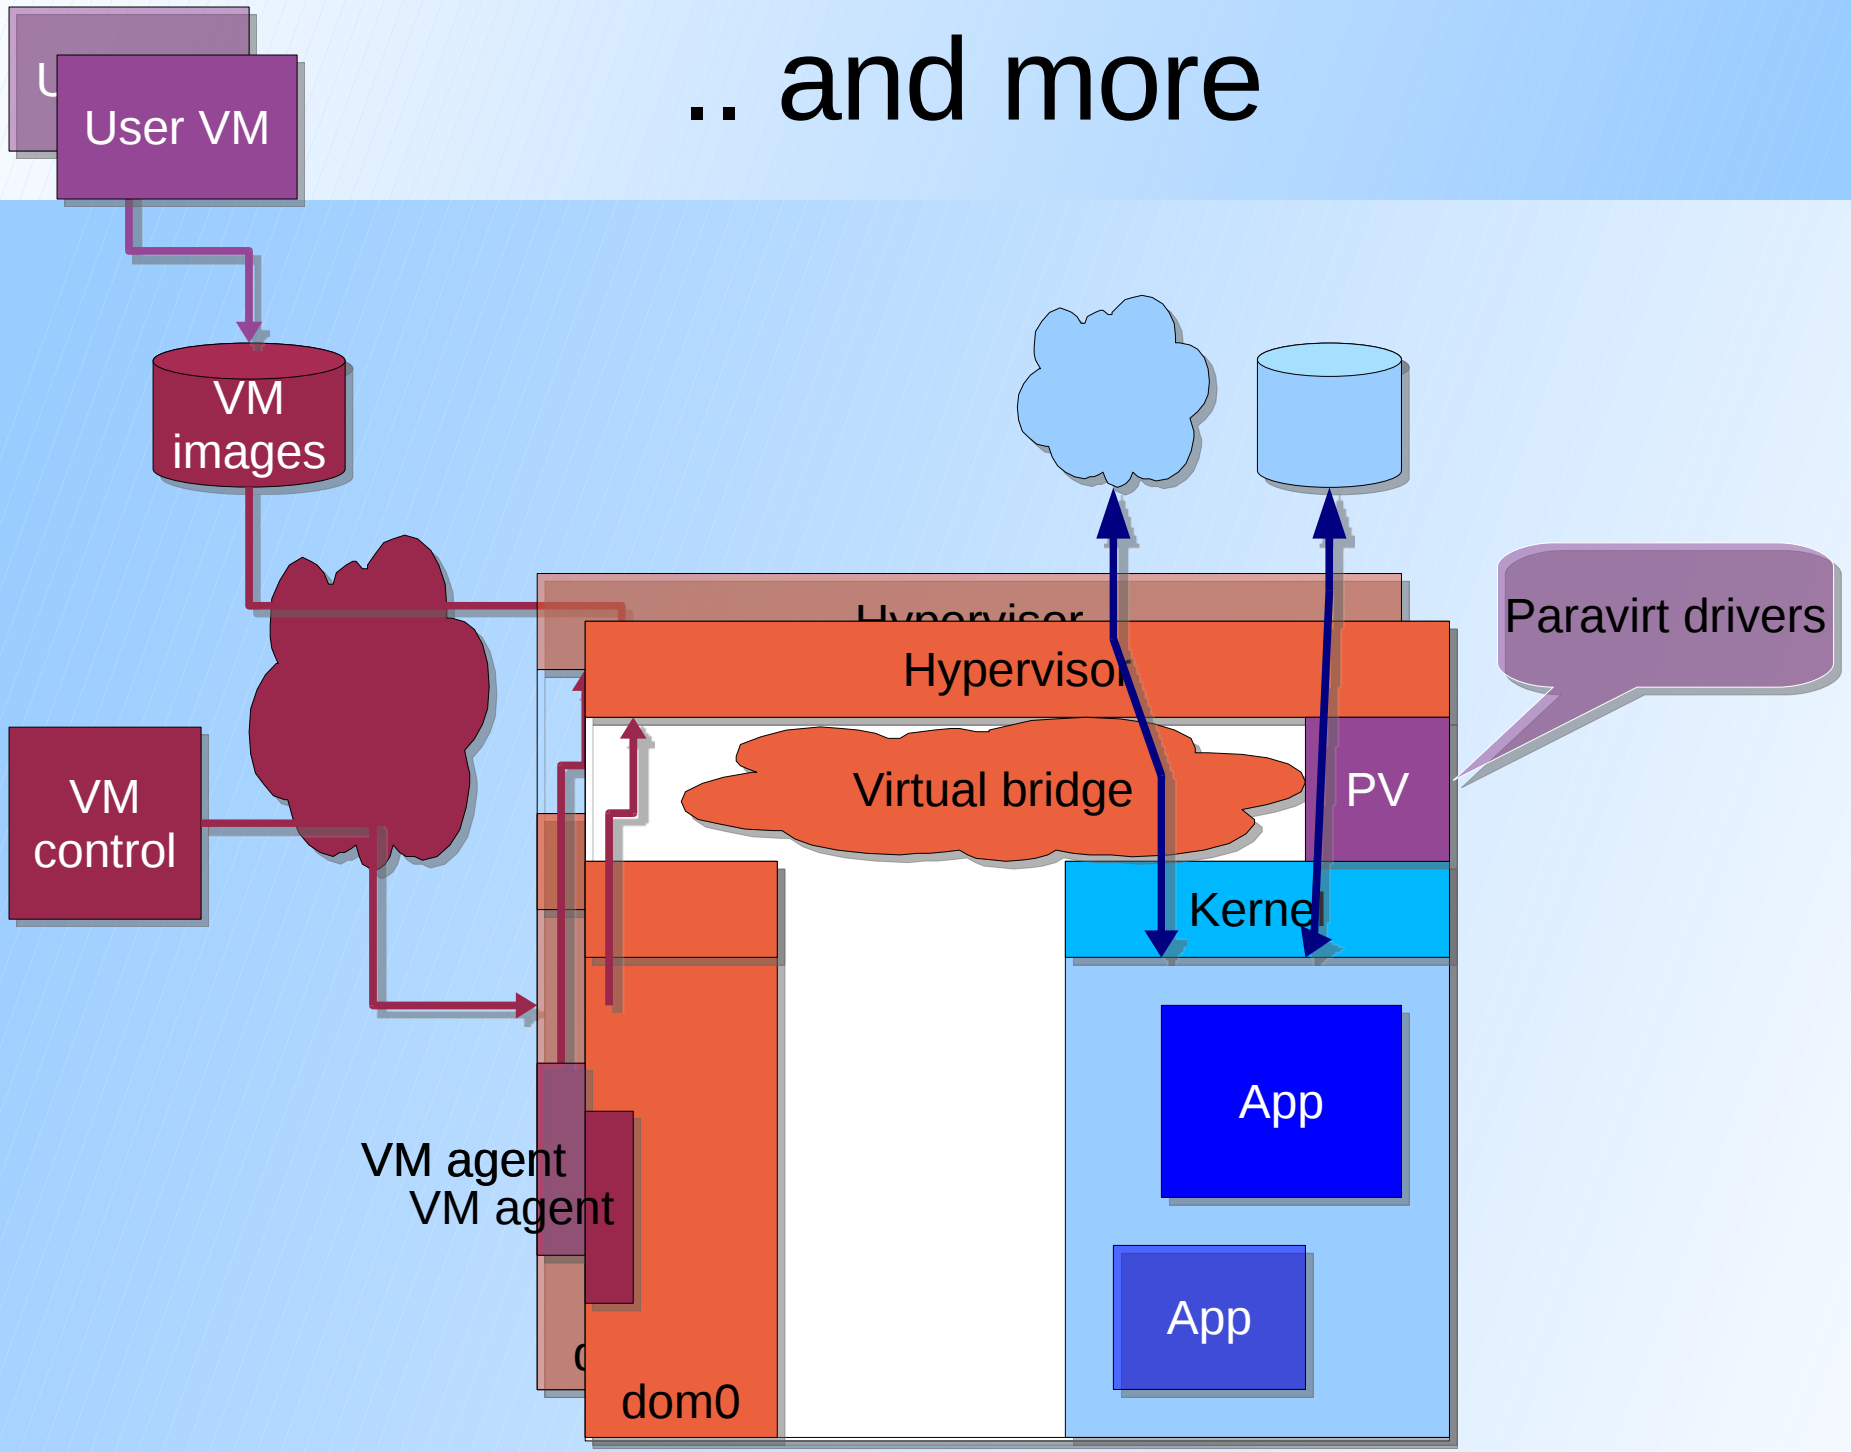
User VM
VM images
VM control
Hypervisor
Virtual bridge
PV
Paravirt drivers
Kernel
App
App
VM agent
VM agent
dom0

VMs: “virtual” security

General: new code, new interfaces = new attacks.

- VMs are inherently **less** secure --- all other things being equal
- More code, more interfaces = more “attack surface”
- Typically “all other things” *not* equal.

VMs provide scope for more secure configurations

- Potentially *less impact* of a vulnerability/exploit – containment might offset complexity
- So need to compare before/after ..

Examples

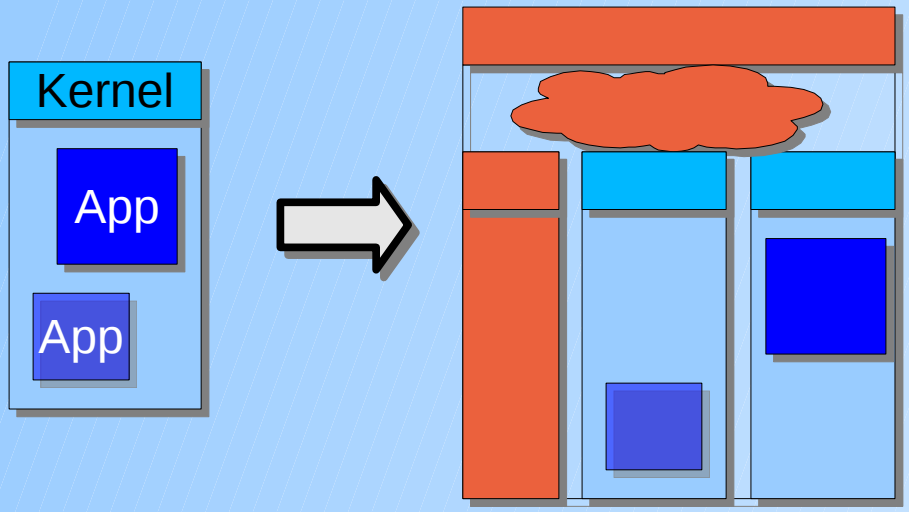

applications from same machine (or trust zone) go into separate VMs

→ ~~more secure~~

applications from different trust zones go into same hypervisor

→ less secure

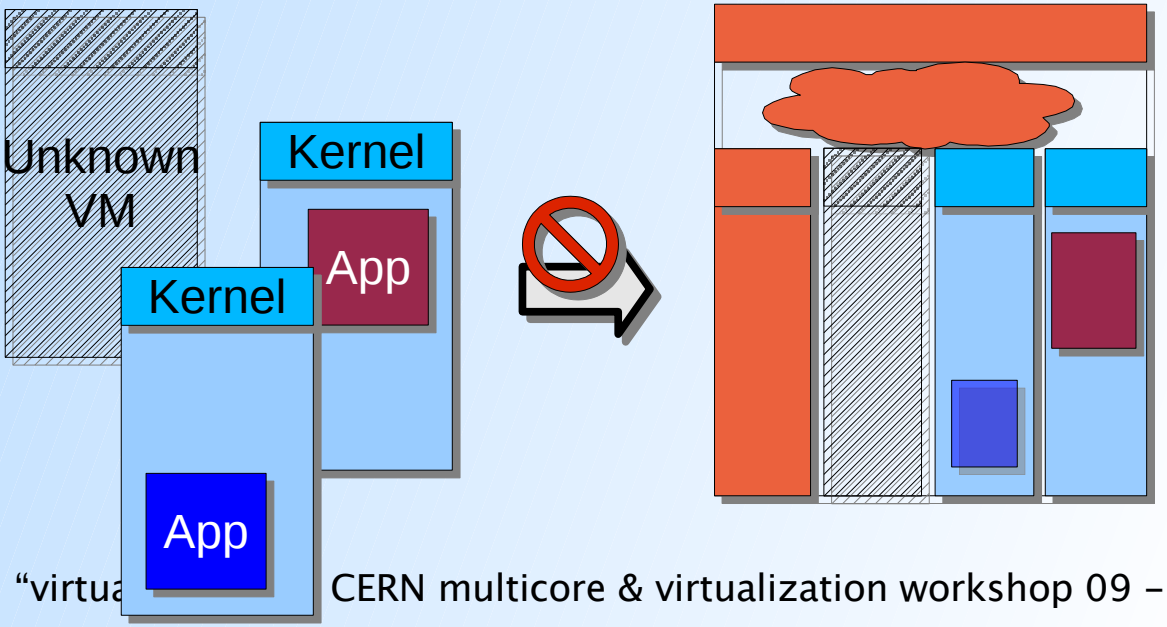

VM-specific attack points

- **VM escape**: break out of VM, into hypervisor or host OS = access to other VMs. (rare)
- "Paravirtual" drivers: (tradeoff: speed vs "virtual")

  Ex. Xen PVFB, Vmware tools, ..
- "VM network":
  - sniffing, IP-spoofing, ARP-spoofing, DoS...
- Controlling infrastructure
  - homegrown or vendor-supplied – remote reset, "elastic" things... Ex: HyperVM/VAserv 100k sites deleted.

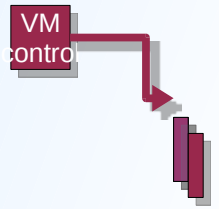

Other security issues around VMs

- Access to VM images

  - Storage/transit: secure? encrypt? (who trusts whom?)

  - Persistent changes to image (or not?) - sandbox?
  - Patching – when, by whom?
    - Potentially easy on trusted guests
    - Unknown on user-supplied = “museum system”

    (Difficult on hypervisors - migration or downtime)
- Snapshot images – full memory dump, incl password/certs..
- “Privileged” network access via host?

“*Nobody is using IP-based access controls nowadays!*”

Wrong.. LHC-OPN, NFS et.al., site firewalls - & Reputation!

Other security considerations

- (user/job) credentials in the VM?
  - Easy for “trusted” images
    - re-use existing tools, get at runtime from LSF/Batch/MyProxy/..
  - Harder for “clouds” - no standard
    - (typically pushed to VM owner/user)

Crucial point: “**trust**” the VM?

- “**Trusted**” VM: → minimal operational changes

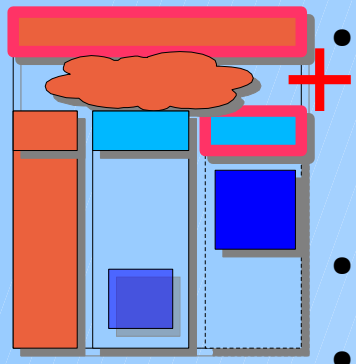

  - Admin access restricted, similar to physical machines
    - Can trust kernel access control, UID/GIDs, local logging, network config.. = additional layer
  - operating system is maintained/patched
  - Inventory control/tracking
  - Batch, Accounting, resource quotas work (machine-local tools)
  - (Might contain 3rd-party SW)

- untrusted/**user-supplied** VM:

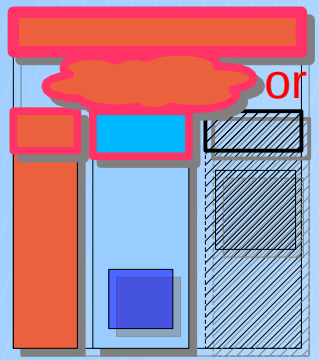

  - Trust boundary shifts to hypervisor ↔ guest kernel
  - Full access to “virtual network” (bridge/NAT)
  - Batch/accounting/quota tasks shift to hypervisor

Trust, cont.

- trusted-third-party VM? ("experiment")
  - "Historically difficult" to get single SW version from experiments (internal fragmentation = no single "production" version)
  - Different focus: functionality+stability vs security
    - Series of "frozen" releases, patching unlikely
    - Expect "service creep" - maintainer cannot refuse new things
    - Deployment is open issues – push to sites, cache ,..
  - CERNVM-style might be a solution - pull in exp sw as needed (but keep OS under site control)

Choice depends on "virtual" focus:

- Efficiency vs Flexibility vs Control

Corollary: how to track even "trusted" images over time?
Sign+Expire? image+mod-on-boot?

“HEP sites” stance

- Question at HEPIX (May 2009, Umeå/SE):
  - What is the “HEP site” stance on user-supplied images?

  (“cloud”-style image vs grid-job-style computing)

  - Discussion summary:
    - Pure “user-images” seem to be out – no real use case so far, huge impact on site networking setup. Eventually: perhaps in fully-contained sandbox (i.e. no network at all)
    - Experiment-images: might come but *lots* of open questions
    - Site-internal images – lots of interest from sites

What are others doing

or: are sites too paranoid?

- AMAZON EC2, FlexiScale, Rackspace etc
  - All take user-supplied images..
  - Well-known AMIs configure “ownership on boot”. Nothing more.
- Model: “hosting service” != HEP computing
  - Usually: passive protection of guest
    - Xen (customized) + paravirt guests; effort on virtual net
    - Inbound firewall on hypervisor (but can open)
    - Good recommendations but “VM **security not their problem**”.
  - Economic incentive against active attacks:
    - $$/network, $$/CPU, take offline, owner's credit card.

(commercial) cloudy security

- see D.Kelsey talk at HEPIX09 - http://indico.cern.ch/contributionDisplay.py?contribId=60&sessionId=13&confId=45282

  Goes beyond purely technical arguments
- Trust:
  - "Predictable, dependable, responsible, honest, reliable behaviour"
  - From experience/collaboration or policies
- Policy (provider vs Grid):
  - usage terms/license vs various Grid policies
    - (incident reaction, notification procedures etc)
  - Legalese – authority to sign contract for your organization?

Summary

Lots of benefits from virtualization..

- Security perhaps isn't the strongest.

But: also lots of new complexity

- Nothing fundamentally challenging, but ...
- lots of issues to consider ...
- ... ideally before deployment.

Virtual use cases

- (see “your standard intro to VMs” for *hypervisor, guest, dom0, domU, image, Type-1, Type-2, para* vs *full,* ..); not talking about “OS compartments”, chOS, emulators etc – but concepts are similar. And “Virtualization” is an old concept...
- General “virtual” use cases:
  - Efficiency – pool underused machines
  - HA – live migration, “hide” interventions
  - Load balancing - “elastic” supply
  - Security: segregation, containment, honeypots
- HEP-specific:
  - Batch job segregation (CHEP – KIT/HTW)
  - Analysis environment, “VO appliance” (ATLAS/PANDA)
  - “Clouds”, User-supplied images (CERNVM)?

(interesting) Incident Investigation Impact (ideas)

- Prevention:
  - Restrictive network access != HEP/Grid ? (various callbacks, experiment monitoring etc)

Patching? Compliance checking?

- Detect: NIDS.
  - HyIDS – hypervisor as (N)IDS – interesting idea

–How to profile behaviour for unknown images?

- Forensics
  - Snapshot: *good* (but: who gets access?)
  - Unreliable local audit: *bad*
  - Yet another dimension: which VM was running?