

XCP Architecture Evolution Plan, Feb 2011

Dave Scott
System Architect

Talk outline

1. Future Hosts and Pools

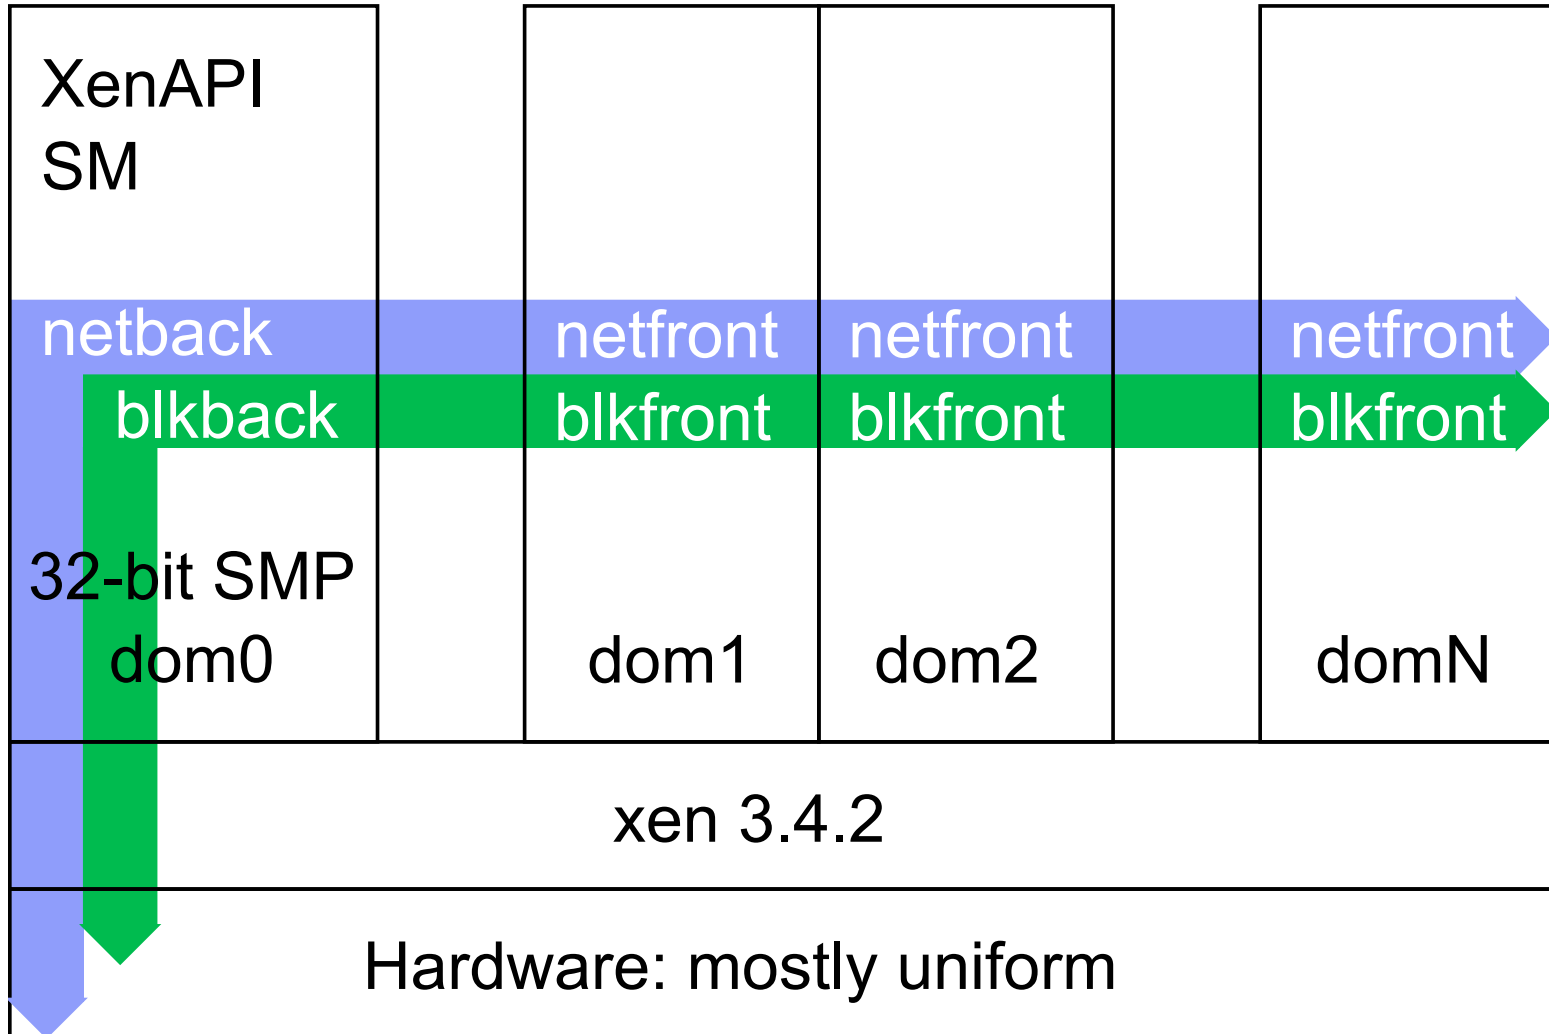
- Overview of how components and interfaces will evolve

2. The “evolution plan” itself

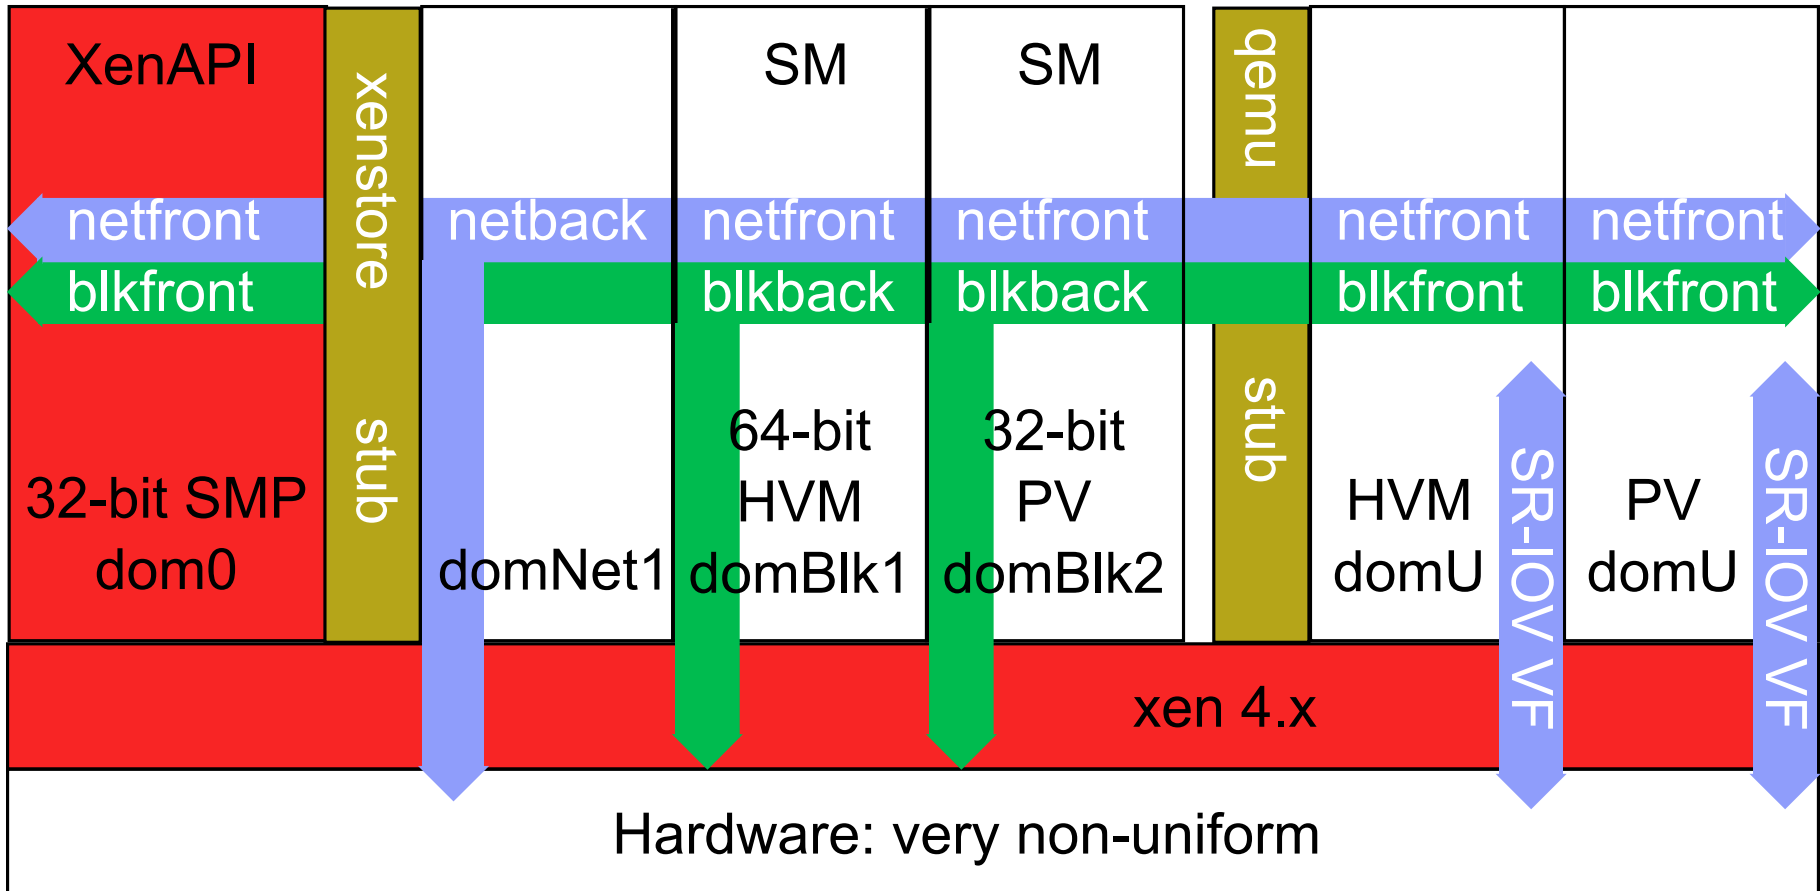
- Where it lives
- How it’s updated
- What’s on it

3. Q&A

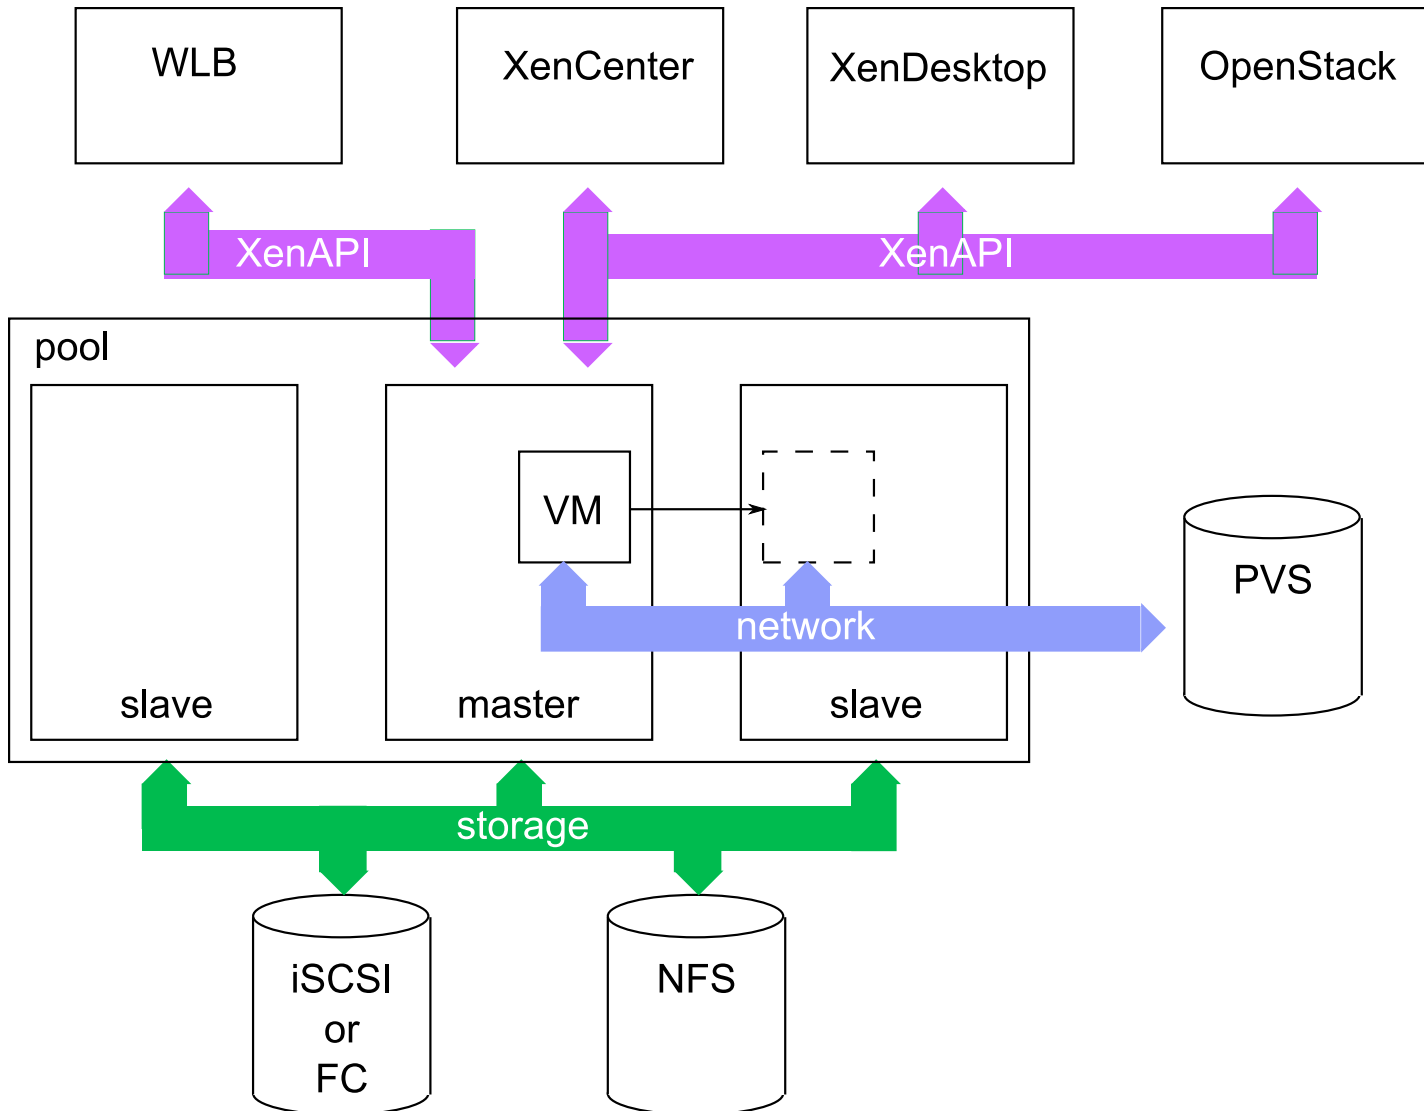
A Host - Today



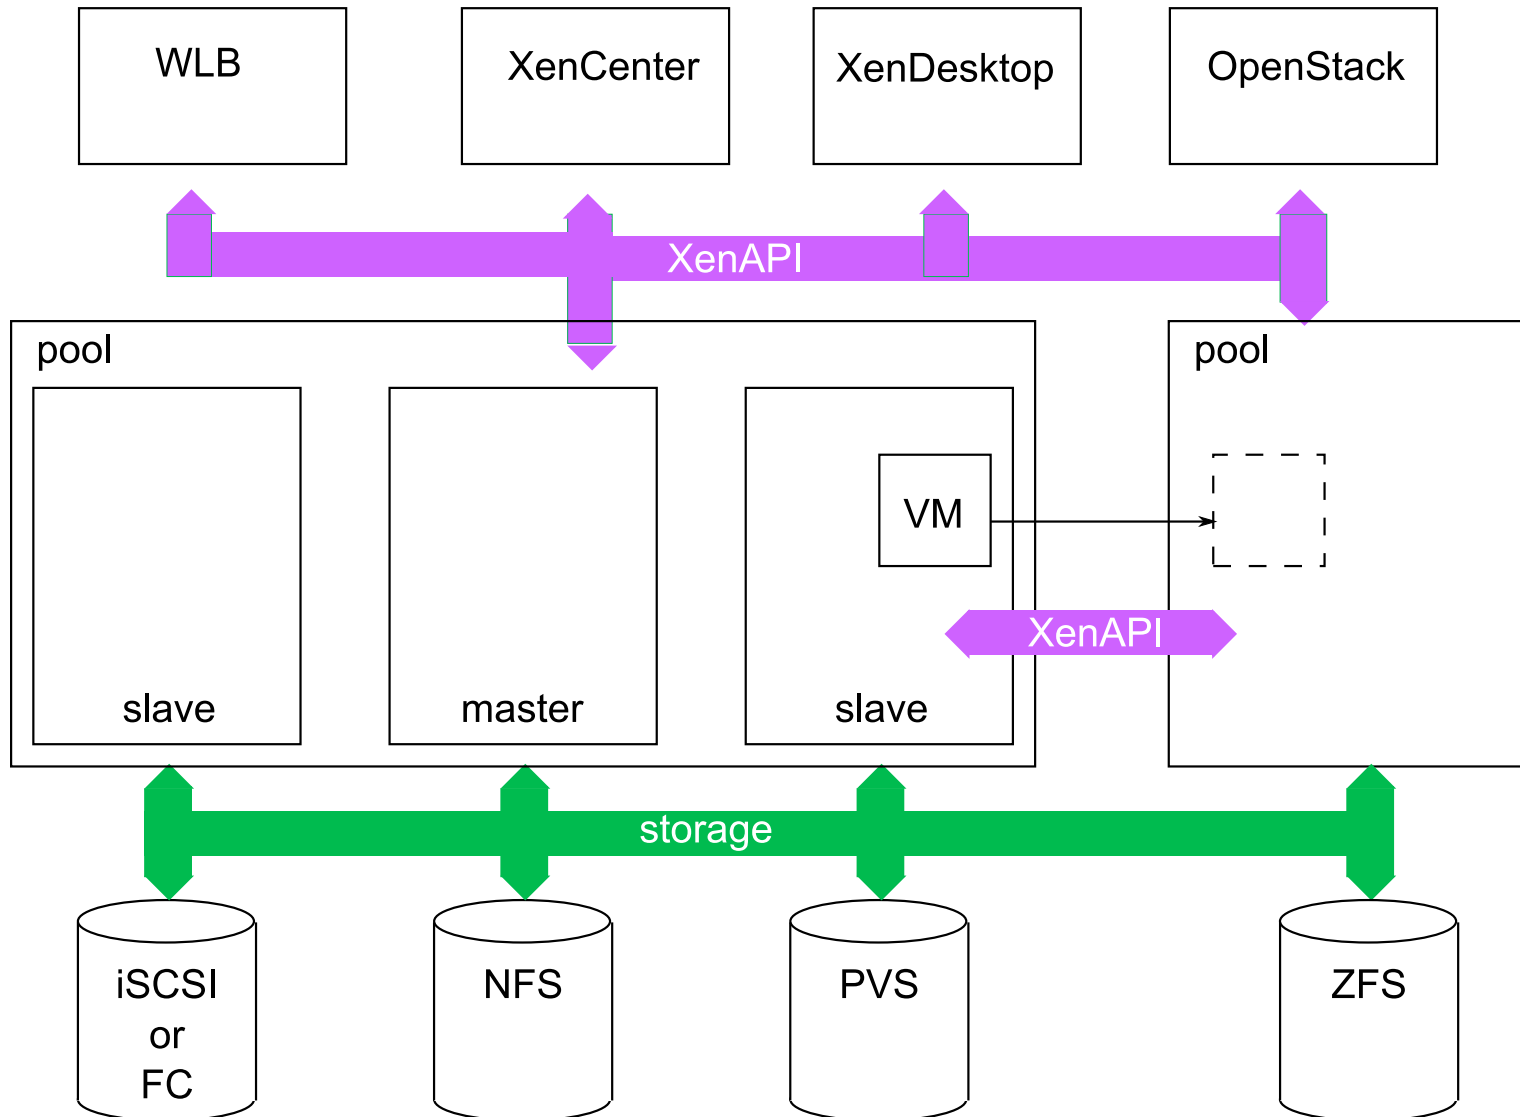
A Host - Tomorrow



A Pool - Today



A Pool - Tomorrow



The Evolution Plan


- On wiki:
 - <http://scale.ad.xensource.com/confluence/display/eng/XenServer+Architecture+Evolution+Plan>
- Items are Jira tickets: easy add/update
- Items are categorised
 - By theme: e.g. “cloud”
 - By (guess of) lead time
- Reviewed monthly by subsystem architects

Current Categories

- XenDesktop
- Cloud
- Appliances
- Performance and Scalability
- Robustness
- Supportability
- Tracking upstream vendors
- Code structure improvements

XenDesktop

Short-term




Long-term

- NUMA auto-affinity set
- Efficient VM image publishing
- Increase max dom0 memory
- IntelliCache for private image disks
- iSCSI/FC block storage for gold images
- Improved WLB power management
- Hypervisor swap
- Memory dedup and sharing

Cloud

Short-term




Long-term

- Cross-pool VM.migrate / VM.copy
- “Shard” the pool metadata db
- Storage motion for vhd storage
- Storage motion for non-vhd storage
- Optimise XenAPI for WAN links
- “stateless dom0”
- Cross-pool / Cloud support for WLB

Appliances

Short-term




Long-term

- v4v for non-IP inter-domain comms
- Detect pending updates for appliances
- Auto-update appliances (with rollback)
- Direct XenAPI client to VM guest agent comms

Performance and Scalability

Short-term




- Expose NUMA topology via XenAPI
- Stubdoms for intensive ops (eg VM export)
- Qemu stubdoms
- Scheduler2
- HVM container for 64-bit PV guests

Long-term

Robustness and Security

Short-term




Long-term

- Enable Xen Security Modules
- Pvgrub (or pygrub in helper domain)
- Snapshot/revert for dom0 fs
- SELinux in dom0 (+any other system domains)
- Driver domains
- Xenstored domain
- Measured boot (e.g. via Intel TXT)

Tracking Upstream


Short-term



- Update CPUID for Intel Sandy Bridge
- Xen 4.2
- Upgrade dom0 to CentOS 5.6
- Xapi <-> libxenlight
- Upgrade dom0 to CentOS 6.0

Long-term

Summary

- Architecture Evolution Plan helps us decide what to build, to satisfy future product/platform needs
- As platform customers, your input is needed!


Q&A