Libvirt presentation and perspectives

Daniel Veillard
veillard@redhat.com

Libvirt project goals

- Web site: libvirt.org
- Virtualization library: manage guest on one node
- Share the application stack between hypervisors
- Long term stability and compatibility of API and ABI
- Provide security and remote access “out of the box”
- Expand to management APIs (Node, Storage, Network)

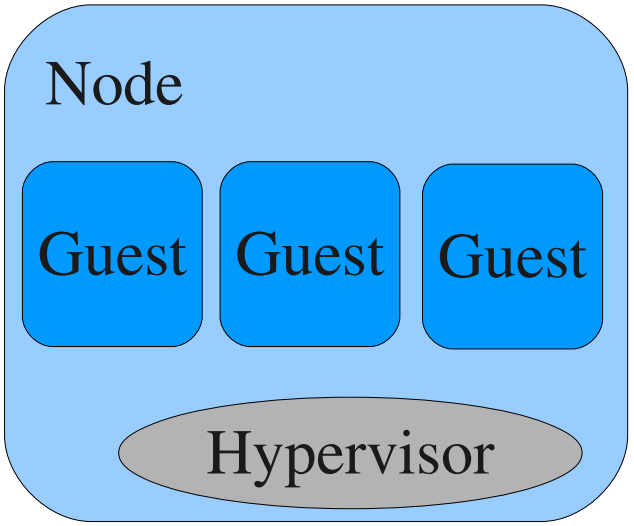

Limit duplication of efforts

Applications are costly to write and maintain !

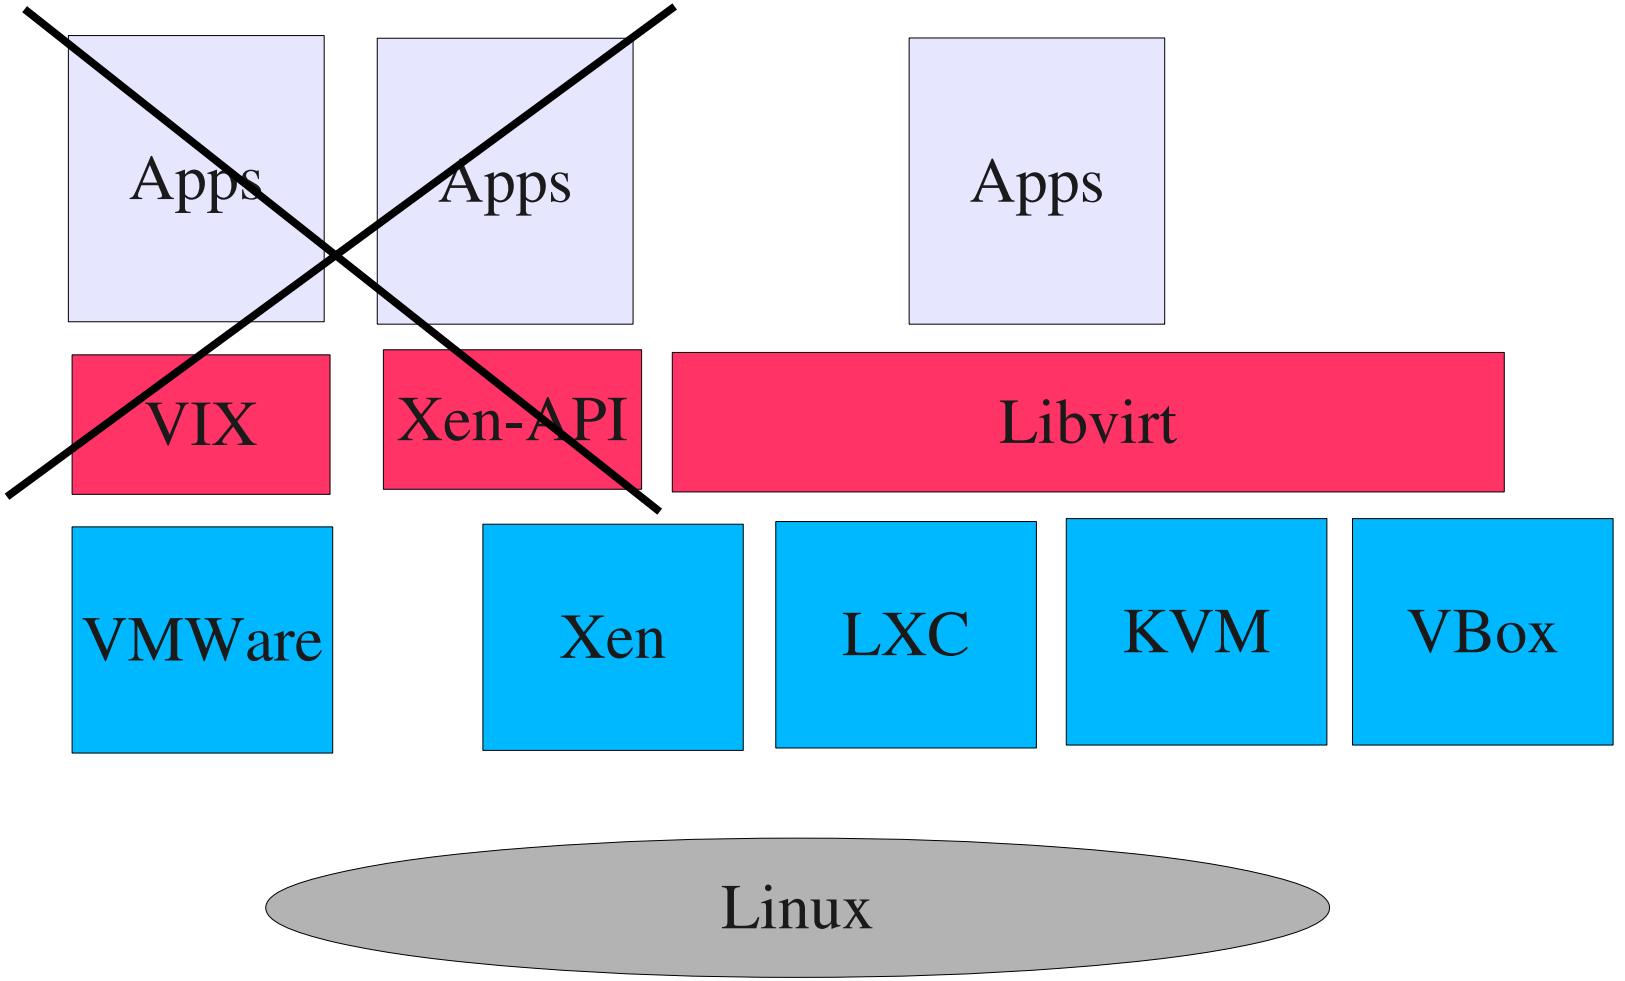

Project current status

- Started 2005
  - 20 commiters, 7 full time Red Hat persons
  - Active list libvir-list@redhat.com
  - A release every month
  - 200-300 git commits/month, 0.9.2
- Support for most hypervisors except Microsoft

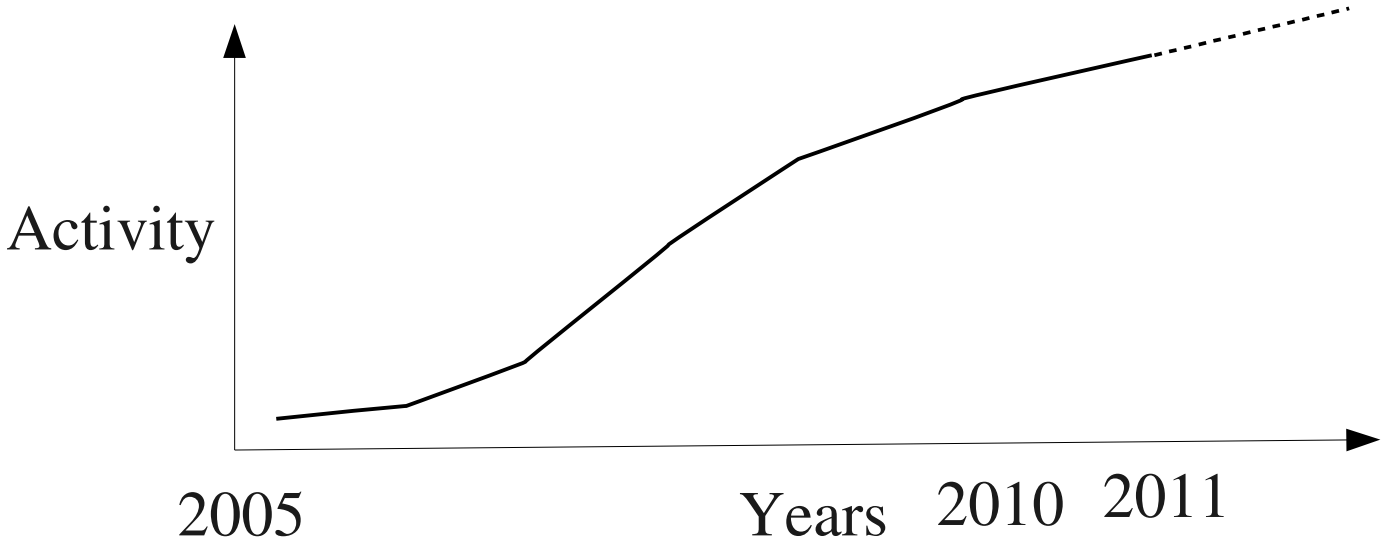

Architecture of libvirt

- Application links to the library
- The libvirt daemon talks to the hypervisor on the node
- Remote protocol is tunneled securely

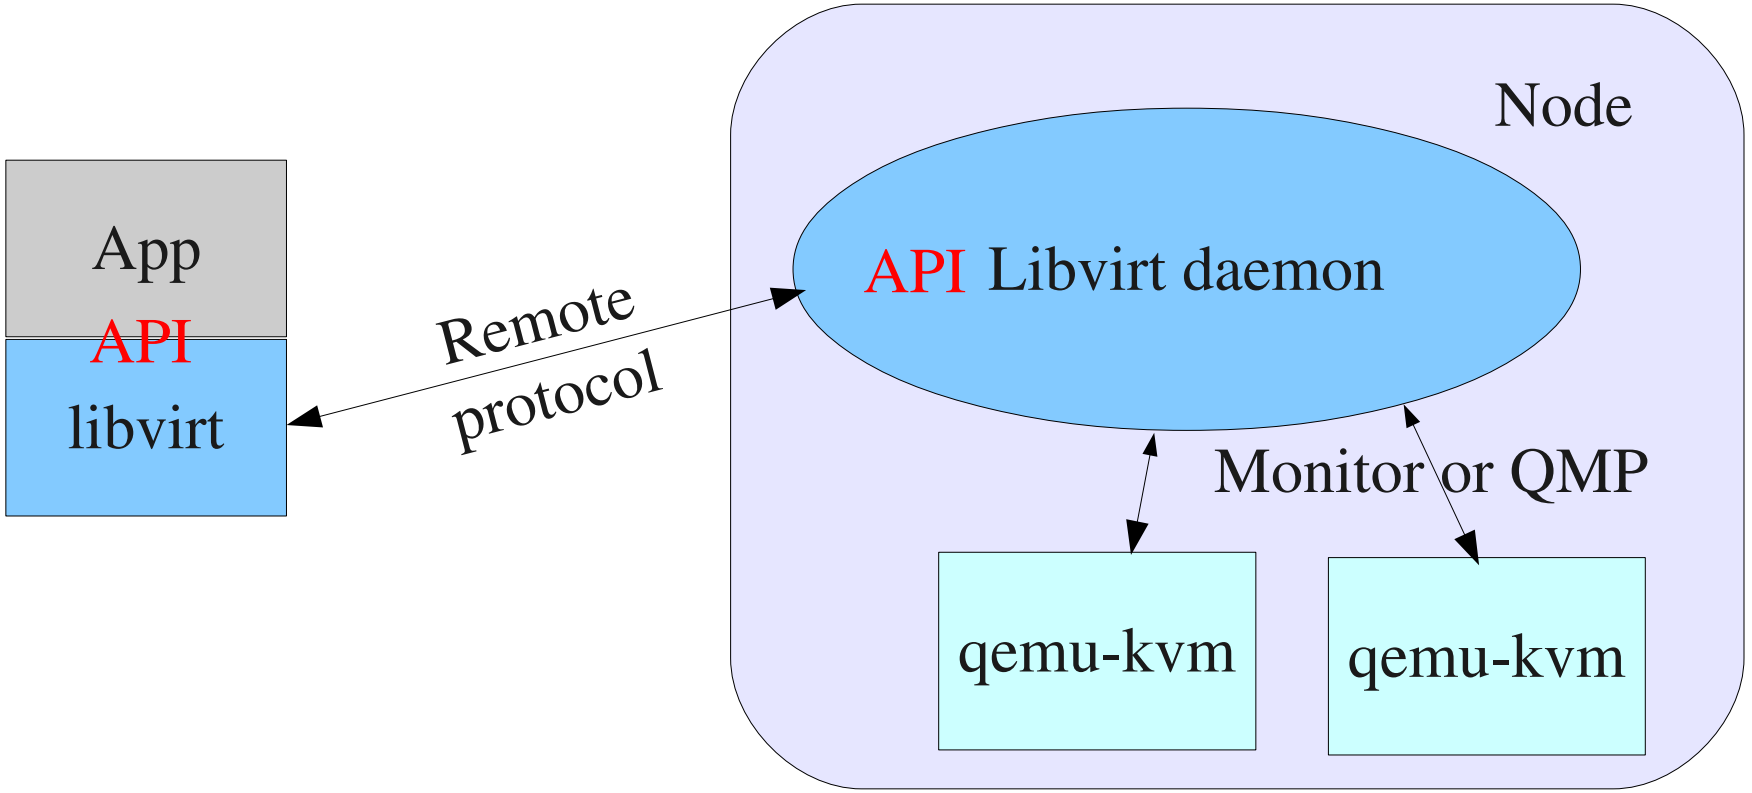

redhat

Current set of APIs (libvirt.h)

See the hypervisor support page for the full list

1) Domain state handling (save, restore, migration, core...)
2) Node and guests resource usage (memory, network, disk)
3) Security, audit and credential handling
4) Domain control (define, create, shutdown...)
5) Tuning (scheduler, memory, I/O, vcpu)

redhat

Current set of APIs (continued)

5) NUMA support (placement, topology, cells usage, pinning)
6) Dynamic or cold device attach and removal
7) Networking (virtual network, interfaces, filtering)
8) Storage handling (pools and volume)
9) Devices handling (enumeration, attach, detach, reset)
10) Asynchronous events callbacks

redhat

On the work and TODO

1) Transactions for interface APIs (0.9.2)
2) Virtual switches
3) Libxenlight and LXC driver improvements
4) Specific support for QEmu low level access (0.9.0)

1) Fine grained ACL for access control
2) Screenshot API, desktop integration

redhat

On the work and TODO (storage)

- Integration with lock managers
- Live snapshots
- Improvements on storage provisionning
- Live block migration

redhat

Classic libvirt applications

- Virt-manager (graphical GUI):
  - manage guests on a few hosts
  - Xen, Qemu/KVM support, LXC coming
- Virsh (CLI for libvirt)

- Libguestfs:
  - Read/modify guest disk images
  - Guestfish shell client
  - API with many bindings

redhat

Libvirt stack overview

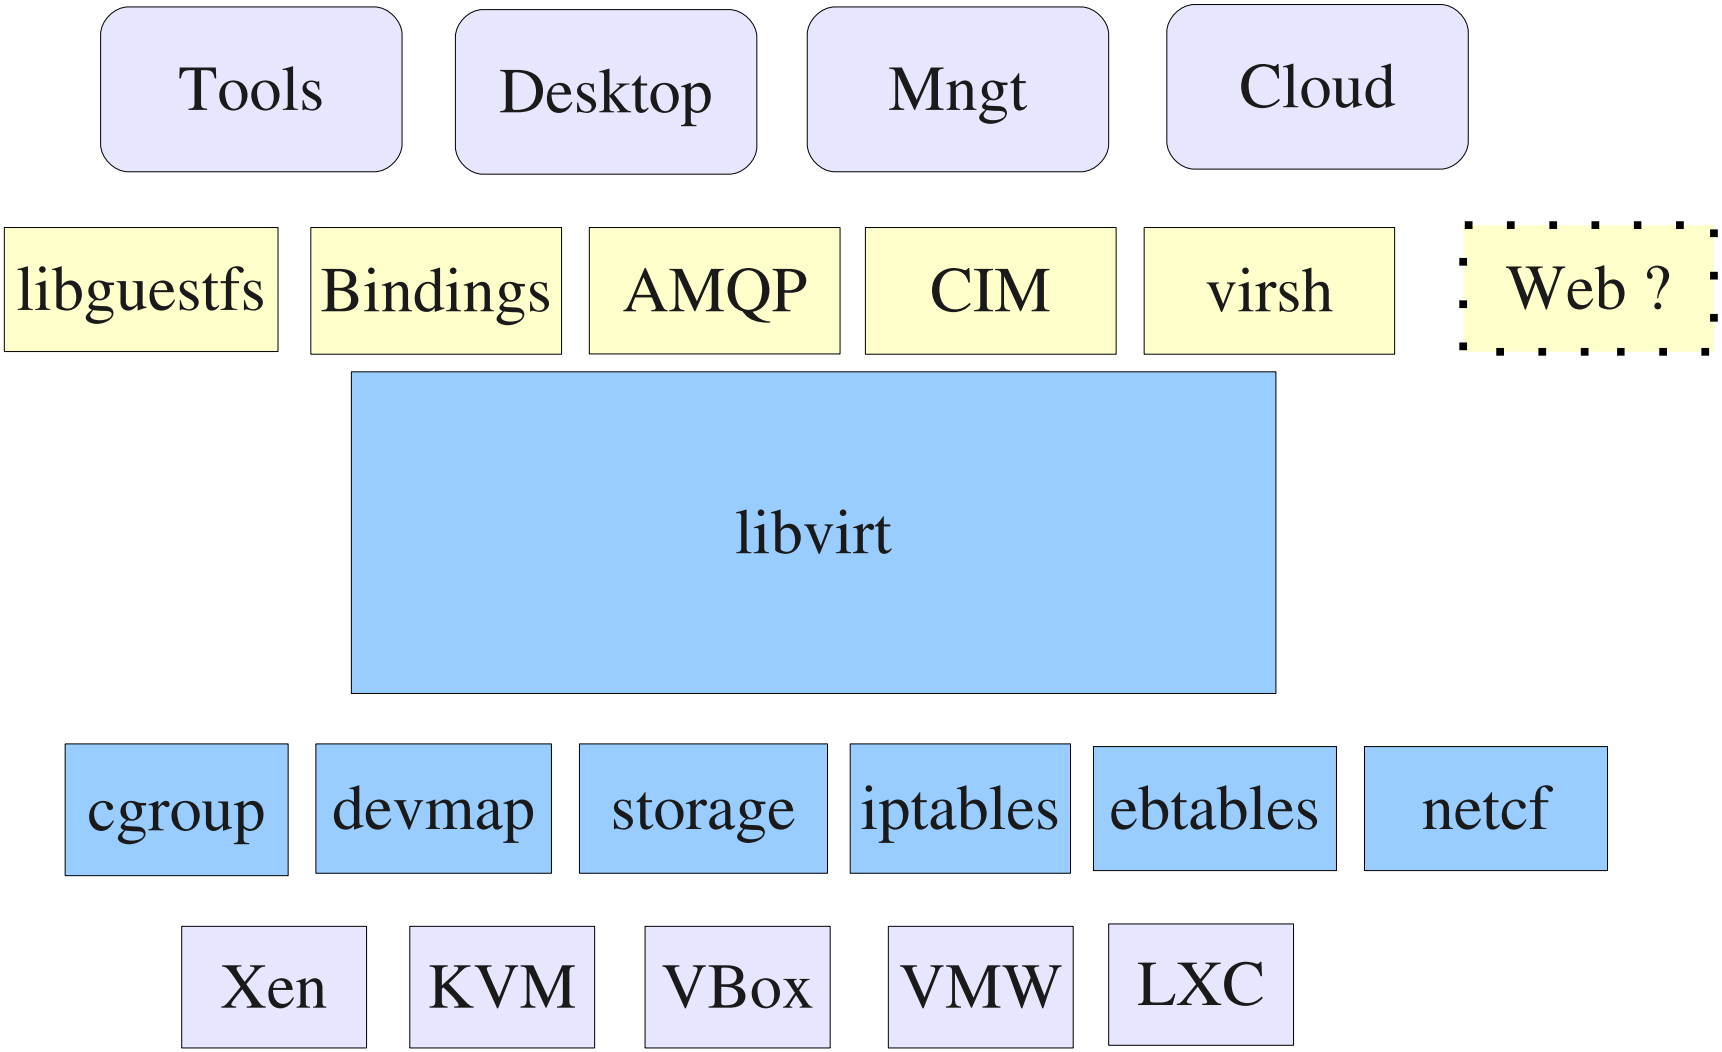

redhat

Prospective work (1.0.0 one day)

- Access to guest from the API (e.g. Matahari)
  - Optimize networking and storage handling
  - Better reboot support
  - General management tasks
- Access though Web/REST
  - Libvirt is not stateless, might be a challenge

redhat

Prospective work (continued)

- Higher level APIs
  - Define policies and use those for simpler APIs
  - Os specific knowledge
- Bridge with the Deltacloud APIs
  - Add a libvirt driver for Deltacloud
  - Simplify building cloud platform

redhat

Matahari

- Agents and APIs running in the guest:
  - Provide services for management
  - Core agent (gather info, network, services)
  - Framework for adding new API on QMF
- Using AMQP protocol
- Cross platform: Linux and Windows
- Transport (Standard IP, virtio-serial guests)
- https://github.com/matahari/matahari/wiki

redhat

Deltacloud (Cloud APIs)

- REST web based API to existing Clouds
  - Support for all major cloud providers
  - EC2, RHEV-M, Rackspace ...
  - Open to add more !
- Apache project
- Part of the project is now the Aeolus project
  - Provides “broker service”
  - Connect to multiple clouds
  - Creation, workflow policies
  - Access and permissions

redhat

Contacts: veillard@redhat.com

libvir-list@redhat.com

- Libvirt is mature
- It is still growing
- Feedback is important

http://libvirt.org/

Questions?