


VirtCloud: Virtualising Network for Grid Environments





David Antoš, Luděk Matyska,
Petr Holub, and Jiří Sitera

CESNET, z. s. p. o.
Prague, Czech Republic




Motivation—The MetaCenter Situation

- Czech national Grid infrastructure (<http://meta.cesnet.cz>)
- computation nodes in Brno, Pilsen, České Budějovice, and Prague connected with Czech NREN CESNET2 (10GE)




MetaCenter and Virtualisation

- virtualisation
 - increasing user control of the infrastructure
 - “dumb&passive” network no longer sufficient
 - network support necessary
 - VirtCloud design (not specific to CESNET2)
- goal architecture: “logical cloud over physical grid”
 - grid: distributed, heterogenous, various owners and administrators
 - cloud: single owner and administrator, on-demand services




Main Use Cases I

- user environment
 - user-provided OS image
 - user control over lifecycle of machines
 - even admin access for users
- privacy and security
 - mutual isolation of virtual clusters
 - ★ user protection
 - legacy insecure components
 - controlling inbound/outbound traffic
 - ★ attack/misuse prevention




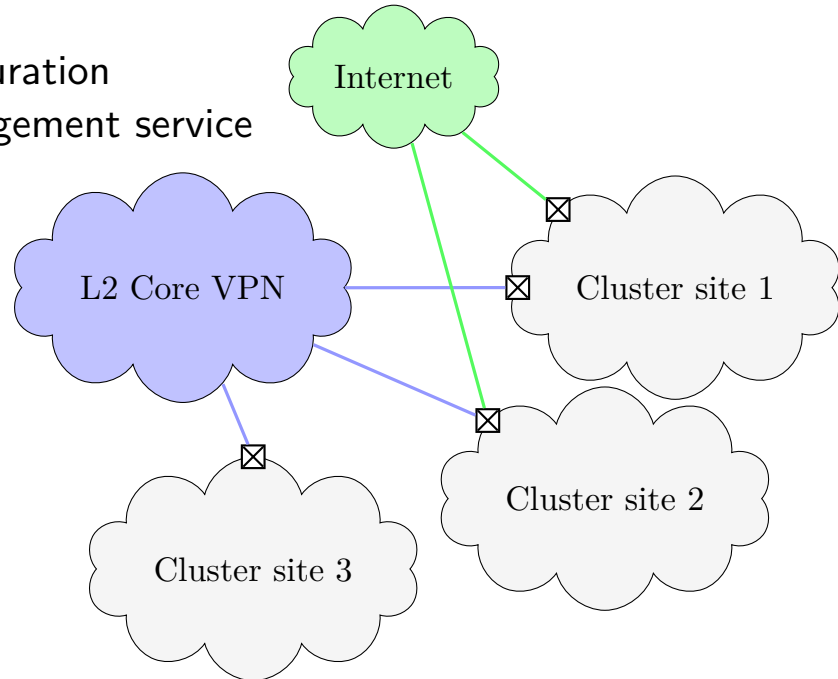
Main Use Cases II

- networking
 - limited amount of IPv4 addresses
 - hardcoded IP addresses in machine images
 - ★ allowing multiple instances of such clusters
 - attaching the cluster to user's network
 - required high performance over a state-wide network
 - ★ without reconfiguring the backbone




VirtCloud Architecture—Networking

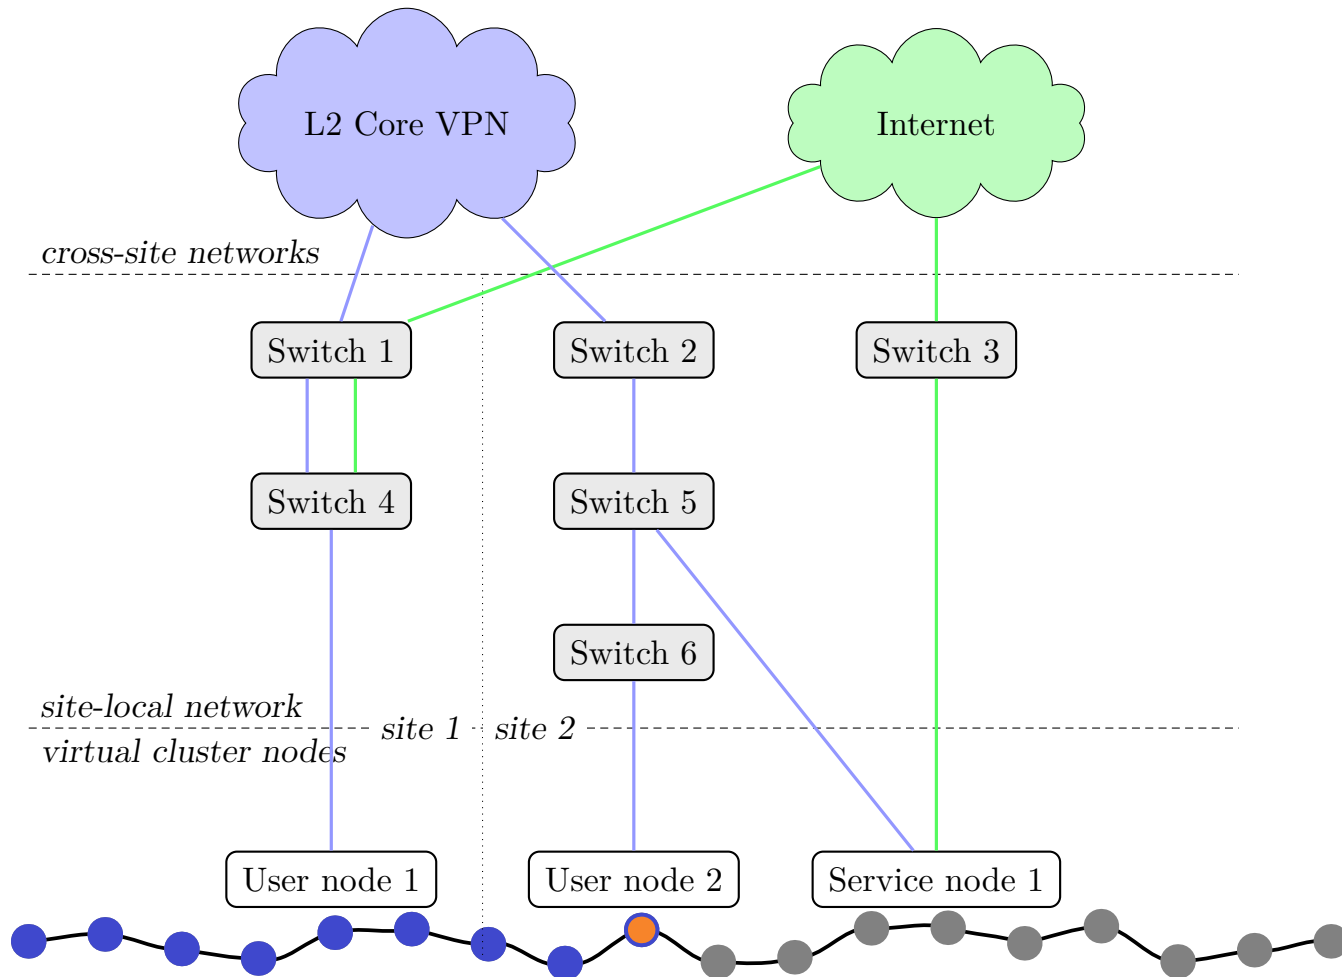
- L2 core network (over the state-wide backbone)
- cluster site network
- (physical) host configuration
- VLAN life cycle management service



VirtCloud Architecture—Auxiliary Services

- managing and booting virtual machines
 - image repository
 - distributing images to physical nodes
 - booting virtual machines
- user access to a closed cluster
 - to allow users to use the cluster
 - connecting the cluster to user's home network
- getting basic services inside the cluster
 - DNS, network attached storage, ...
 - preferably not via user's home network






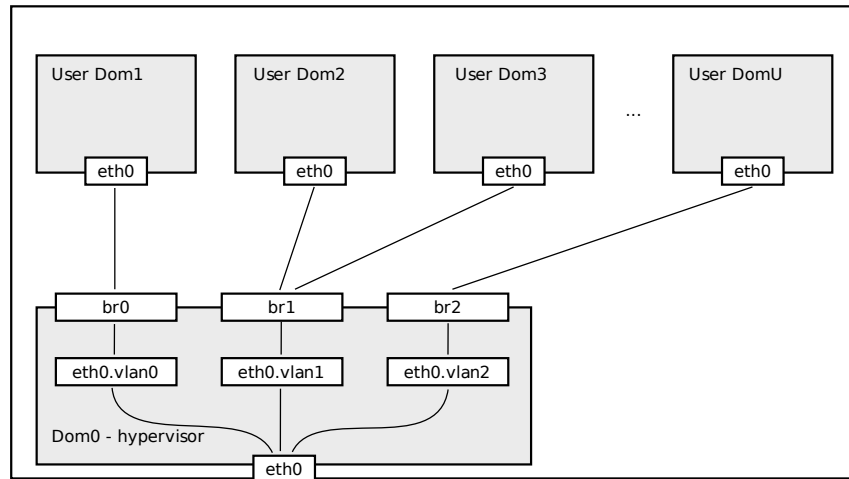
VirtCloud Implementation I

- core network—available technologies (only high-speed hardware-supported multipoint solutions)
 - Virtual Private LAN Service
 - Cisco Xponder
 - (IEEE 802.1ad (QinQ))
- site network: mix of Force10, HP, Cisco switches




VirtCloud Implementation II

- host configuration
 - bridging in Xen hypervisor
 - Layer 3 addressing is user-driven


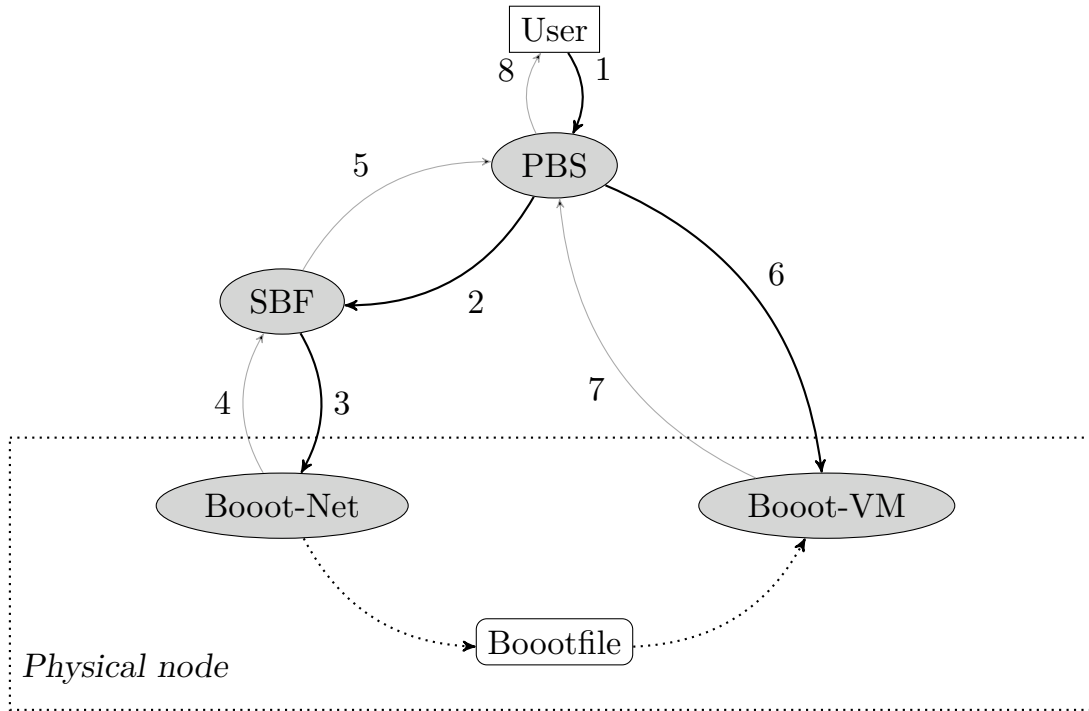


VirtCloud Implementation—Components

- job scheduling system
 - virtual cluster is a job for physical machines
 - PBS
- virtual network lifecycle management system
 - SBF
- managing host configuration and booting virtual machines
 - Boot-Net and Boot-VM




VirtCloud Implementation—Interactions




Access to (and from) Virtual Clusters

- implemented by a “service virtual machine”
 - booted with the cluster
 - user access from outside
 - ★ external tunnelling
 - ▷ OpenVPN (similar to Nimbus)
 - access to MetaCenter data/services
 - ★ attaching the service to the virtual cluster
 - access to external resources
 - ★ allowing the external traffic


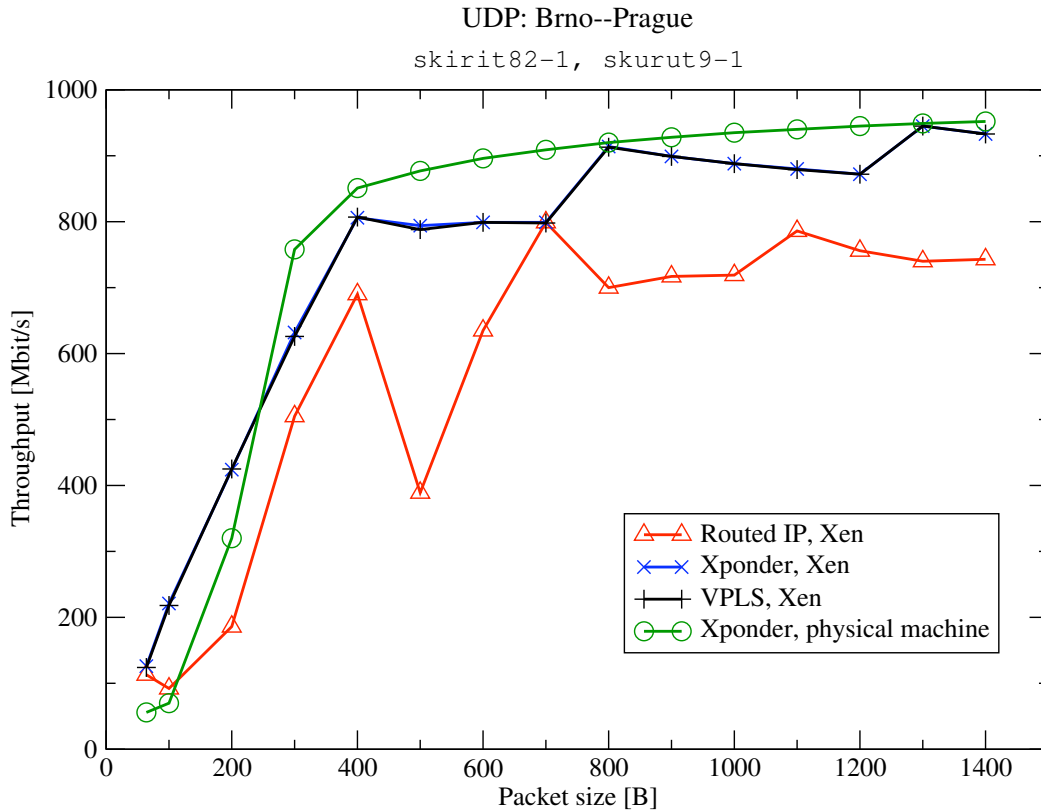


First VirtCloud Experiences


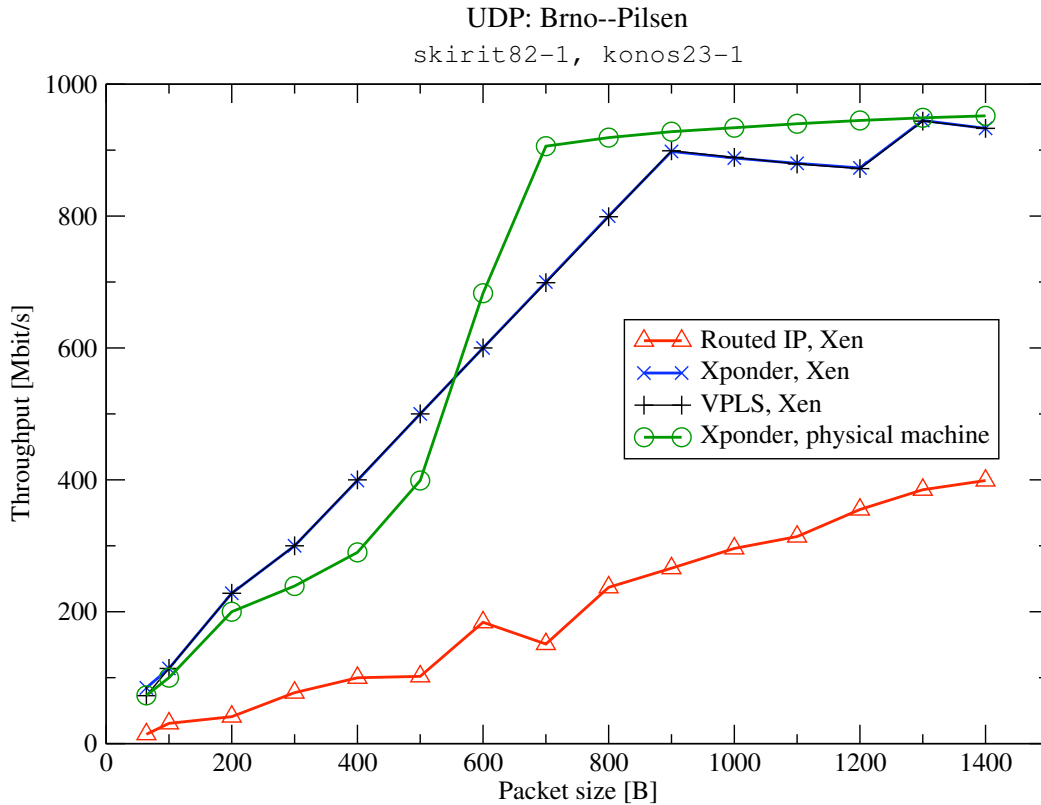
- interfering large-area networks has performance implications
- the performance must not be significantly worse than of native IP network
- to show feasibility of the concept
 - stability tests
 - throughput tests comparing
 - ★ Xponders in physical machines
 - ★ Xponders in Xen user domain
 - ★ VPLS in Xen user domain
 - ★ native IP connection
 - iperf, UDP bandwidth with $< 0.5\%$ loss



Results I



Results II



Conclusion

- architecture of VirtCloud, a system for internetworking dynamic virtual clusters over a large high performance network
- analysis, architecture, implementation
- evaluation: feasible
- future work
 - strategies of external resource access
 - integration of components
 - concepts of virtual machine migration
 - filesystems (home and scratch)

