


MetaCenter Virtual Networks




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

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




MetaCenter and Virtualisation

- Czech national Grid infrastructure (<http://meta.cesnet.cz>)
- computation nodes in Brno, Pilsen, and Prague
- connected with Czech NREN CESNET2 (10GE)
 - physical and virtual nodes have public IP addresses
- virtualisation
 - increasing user control of the infrastructure
 - “dumb&passive” network no longer sufficient
 - network support necessary
 - VirtCloud design (not specific to CESNET2)





Design Considerations—Summary

- high-performance virtual private network, state-wide
- dynamic virtual cluster creation
- encapsulation of virtual clusters
- no administrative privileges on the backbone
- controlled user access
- controlled access to external resources
- enabling migration of virtual machines
- multiple instances of identical L3 addresses
- optional cluster publishing




VirtCloud Architecture

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




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




VirtCloud Implementation II

- VLAN lifecycle manager (SBF, Slartibartfast)
 - allocates VLAN numbers
 - configures active elements
 - ★ has to understand the network topology,
 - ★ limitations of switches, ...
 - controlled by PBS
 - ★ virtual cluster is represented with a special job
 - cooperates with Magrathea (see poster no. 22)
 - ★ Xen domain configuration




Access to (and from) Virtual Clusters

- user access from outside
 - external tunnelling
 - ★ OpenVPN (similar to Nimbus)
 - ★ ssh tunnels
- 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
 - methods of publishing encapsulated clusters
 - scenarios of Layer 3 addressing
 - strategies of external resource access

