Grid y Computación
de Altas Prestaciones
GRyCAP

Carlos de Alfonso
Ignacio Blanquer

REALISTIC NETWORKING IN GENERIC MULTI-SITE CLOUD DEPLOYMENTS

Instituto de Instrumentación para Imagen Molecular (I3M).
Centro mixto CSIC - Universitat Politècnica de València - CIEMAT

UNIVERSITAT
POLITÈCNICA
DE VALÈNCIA

OUTLINE

- Introduction
  - Cloud Computing
  - Community Clouds and Networking
- EPFloater: an Endpoint Floater
- Network Virtualization and SDN
- Networking for Multi-site Federated Cloud Deployments
- Conclusions and Future Work

- Cloud computing
  - We all know about Cloud Computing
    - "on-demand self-service, broad network access, resource pooling, rapid elasticity, etc."
  - Cloud computing is based on virtualization
    - Virtual Machines, Storage, Database as a Service, etc.
    - We have tools to create such virtual resources
      - Hypervisors, Applications, SAN, etc.

- Community Clouds
  - Federation of resources from multiple organizations.
  - To provide similar services to any member coming from any of the organization.
- A set of services need to be agreed
  - Authorization, Virtual Machine Image distribution, Accounting, Common interfaces, Schedulers, etc.
- The users are able to run VMs in different sites

- A VM needs a network connection to be accessed
  - Public IP addresses are "expensive"
    - The sites have a limited amount of them and there is a large amount of VMs.
  - But not all the VMs need public routable IP addresses
    - e.g. in a cluster, the front-end needs a public IP address, but the working nodes do not need them.
- But the VMs need to communicate.

- Cloud Computing lacks from versatile Networking Virtualization tools
  - Private IP addresses
    - The VMs cannot communicate with VMs in other sites using the private IP addresses, as they do not travel out of the site.
  - Public IP addresses
    - Migrating VMs accross different domains or physical locations make that the VMs change their IP addresses.
      - According to the current model, routable IP addresses depend on the physical location of the servers.

EPFloater: Floating TCP/IP End Points

- Some VMs in a Cloud deployment need to be accessed from the Internet
  - We have solutions like "Floating IP Addresses"
    - Used in OpenStack and Amazon AWS.
- EPFloater manages Floating TCP/IP End Points
  - Generalization of Floating IP addresses
    - Consists in forwarding the traffict directed to a pair (Public IP:port) to other endpoint (Private IP:port).
      - A Floating IP forwards all the traffic from all the 65535 ports
    - A VM exposes few protocols (http, ssh, etc.)

EPFloater: Floating TCP/IP End Points

- Features
  - Using the lighweight mechanism of iptables
  - Includes intelligence to co-allocate ports and IP addresses
  - REST API

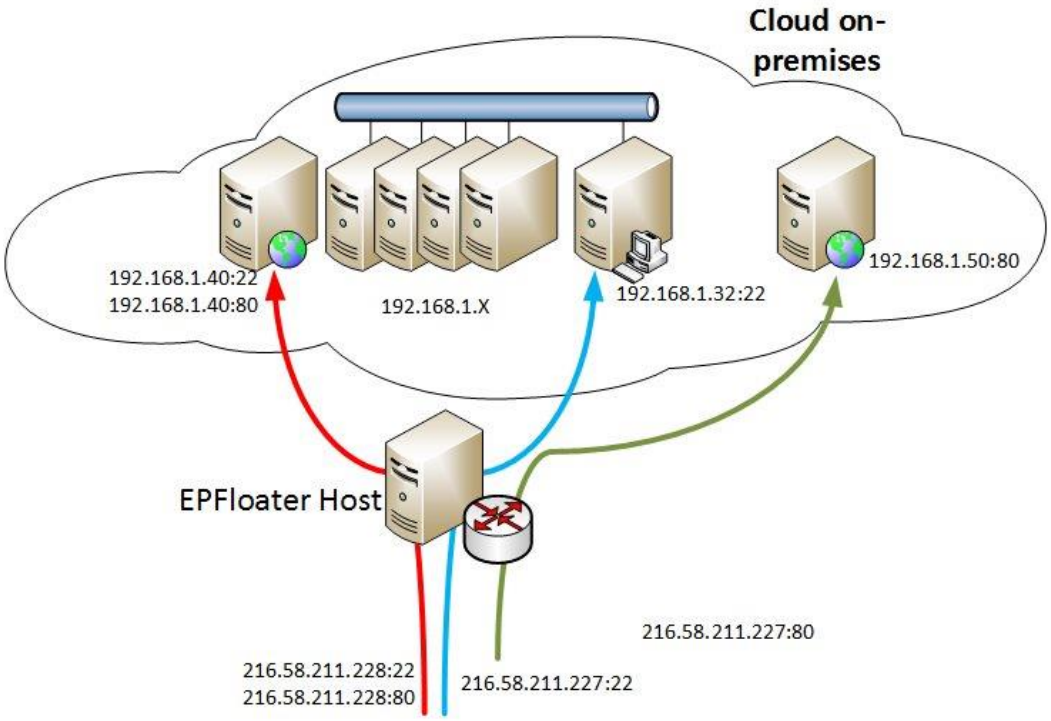

NETWORK VIRTUALIZATION AND SDN

- Virtual Networks
  - Logical networks that are decoupled from the underlying network hardware.
- Software-Defined Networks
  - The network control in switches and routers is decoupled from forwarding and is directly programmable.
    - This enables to program the network, instead of having to make a manual configuration through low-level interfaces.

- Overlay SDN-enabled Network as an approach to Network Virtualization
  - Uses the current network as a transport network.
  - We can dynamically reprogram the components to implement Virtual Networks for each tenant.

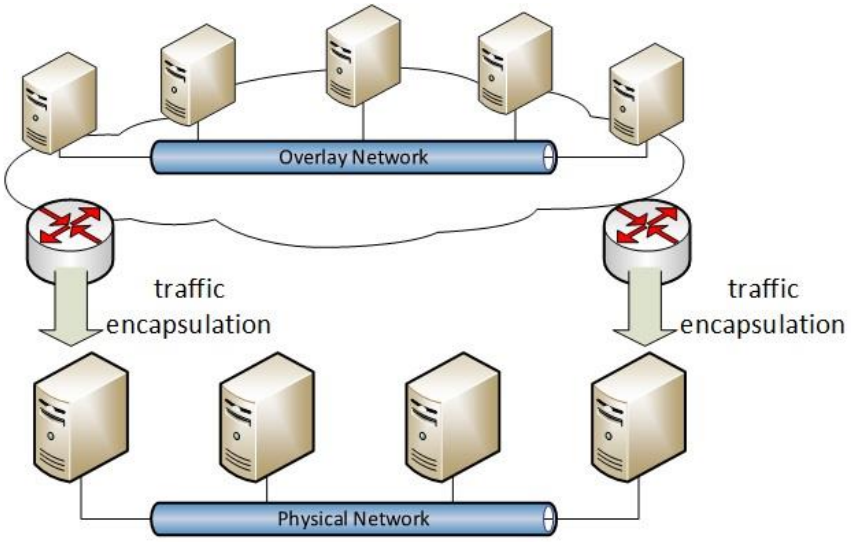

NETWORKING FEDERATED CLOUDS

- Usual scenario in Federated Clouds
  - Sites deploy their Clouds according to their requirements and internal features and needs.
    - Authentication, Cloud middleware, Network, etc.
  - Establishing mechanisms to integrate the sites.
- Some Federated Clouds have rules and mechanisms for other features than networking
  - In EGI FedCloud: VMCatcher for images, VOMS for authentication, rOCCI for interfaces, etc.

NETWORKING FEDERATED CLOUDS

- Connect the sites using a SDN-enabled network
  - Using boundary SDN-enabled switches to automate the control of the traffic between sites.
    - e.g. prevent that traffic from VMs is forwarded to other site if the destination is a VM in the same site.

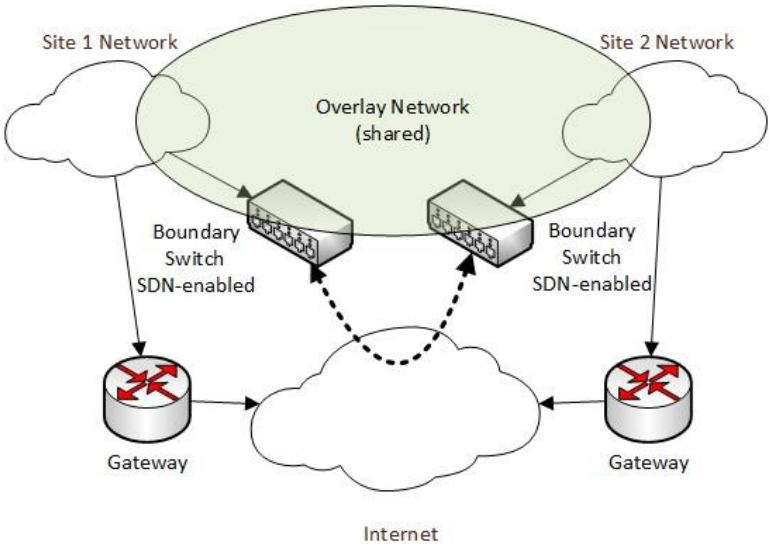

NETWORKING FEDERATED CLOUDS

- It is possible to create 3 kind of networks
  - Federated Networks (shared across sites)
  - Private Networks (isolated from other sites)
  - Wide Area Networks (to access to the Internet)
- It is possible to communicate between networks
  - Using gateways

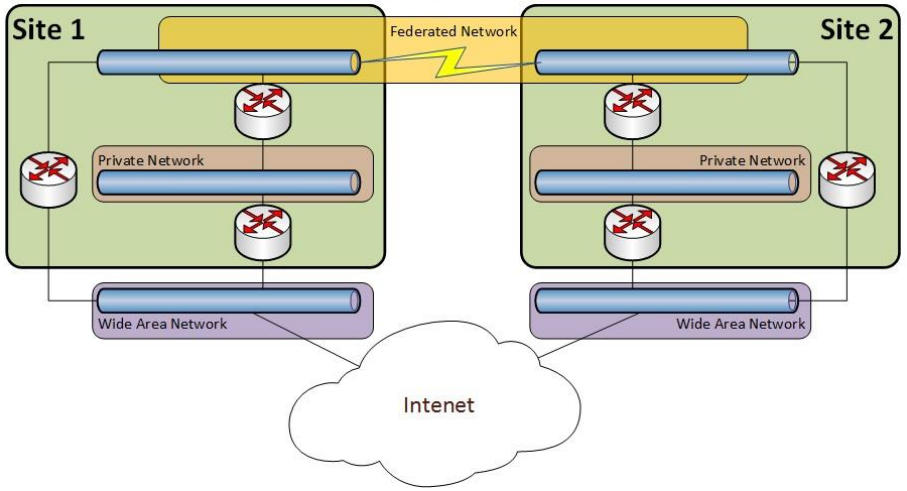

- Additional services
  - DHCP servers, Load Balancers, Firewall as a Service…
    - Each service can be provided on each kind of networks

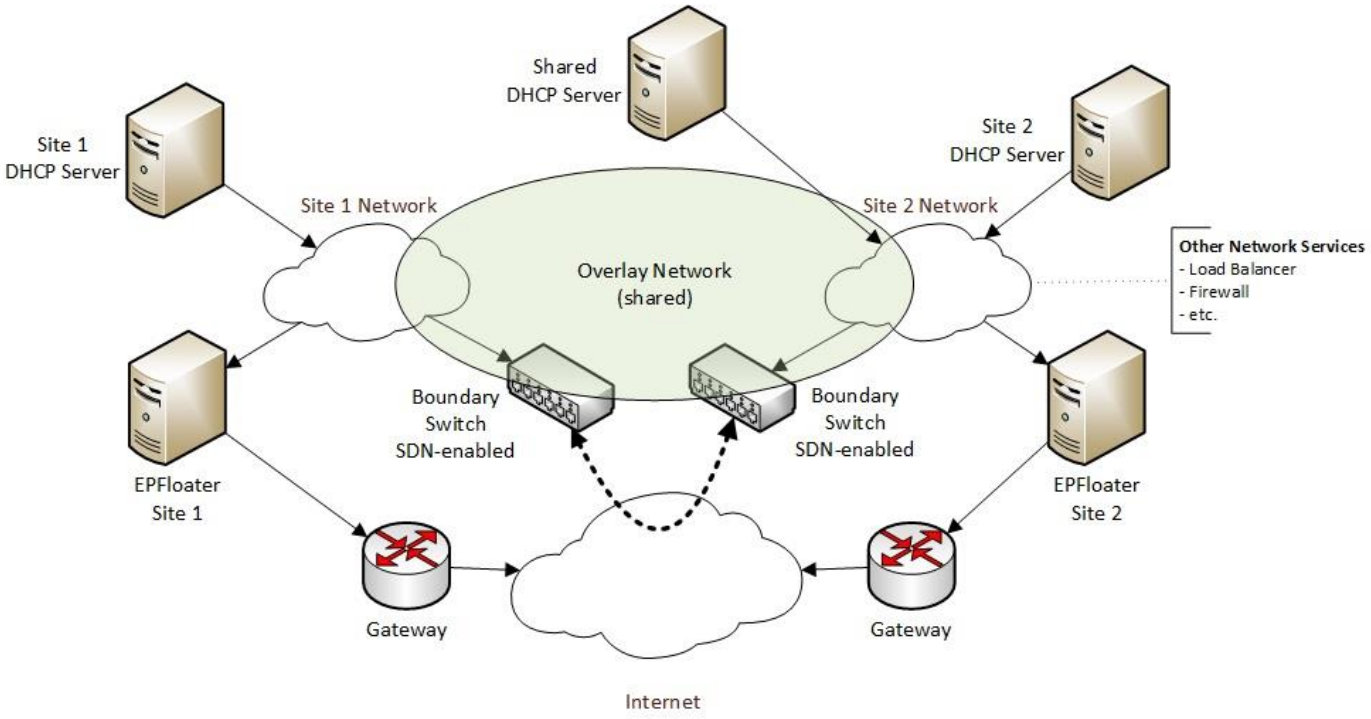

NETWORKING FEDERATED CLOUDS

- Features of this scheme
  - Enables shared networks between different sites.
    - Including DHCP servers
      - Using SDN we will create boundaries for the traffic and will forward to other sites only some DHCP traffic.
  - VM migration, keeping the public IP address
    - The traffic from IPs in the shared network will be forwarded only if needed.
    - EPFloater translates the public IPs into private IPs and so the traffic will be forwarded between sites if needed.
  - Load Balancers, Firewalls, etc. can be used in each kind of network independently.

CONCLUSIONS

- We have proposed a simple yet powerful scheme to enable realistic networking in federated multi-site cloud deployments.
- It is based in OpenVSwitch
  - Creates general boundary at a potential high network overload (due to protocols such as Spanning Tree).
- The future work will include one or more SDN controllers to automate tasks and program the boundary switches.