Stanford SDN-Based Private Cloud

Johan van Reijendam (jvanreij@stanford.edu)
Stanford University

Executive Summary

- The Web and its infrastructure continue to make phenomenal progress, allowing the creation and scaling of applications for millions of users.
- However, scientific computing groups on university campuses are unable to get the resources they need.
- Two disruptive technologies can bring a new paradigm to the campus.
    - Multi-Tenancy Virtualized Clusters
    - Software Defined Networking
- The Stanford proposal will bring both of these technologies to the campus, integrating compute clusters with a campus-wide SDN network.

Introduction

Networked computing has been very important for the research and education mission of universities but has failed to take advantage of new technologies.

- Scale-out cluster design
- Virtualization
- Multi-tenancy

Clouds

Google, Amazon, Yahoo!, and Facebook have pioneered scale-out computing and storage, achieving phenomenal scale with commodity pricing.

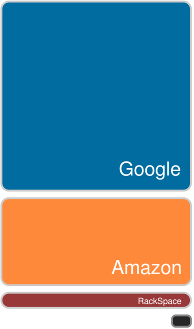

- Google: 12 self-build data centers, millions of servers
- Amazon: 7 data centers, over 454,400 servers
- RackSpace: 70,000+ servers
- Stanford (ITS): 6,000+ servers

Fast and Flexible Scaling

- The ability to self-service.
- No need to maintain a complex IT infrastructure.
- Lowered capital expenditures.
- Easy access to the best infrastructure.
- "Utility Supercomputing"

It now costs just pennies to get the necessary infrastructure up and running to get a venture off the ground.

Network Infrastructure Transformation

SDN/OpenFlow

- More control over the network infrastructure.
- Reduced capital and operational expenses.
- Customized services.
- New services, enabling more innovation.

SDN/OpenFlow is being embraced not just on campuses but also by large network owners and operators:

- Google OpenFlow WAN between data centers.
- NEC enterprise and data center deployments, reducing cost.
- VMWare acquisition of Nicira.
- Over 200 campus and research OpenFlow deployments.

What is Wrong

- Campus network infrastructure is out-dated compared with Web or cloud infrastructure.
- Researchers cannot get enough compute cycles and storage in a cost-effective way.
- Production networks built using closed and proprietary equipment.
- Proliferation of appliances (firewalls, load-balancers, ...)
- Different management interface for every device type.

A Typical University Infrastructure

- User requirements for computing, storage, and networking have been growing very fast.
- Universities have been adding compute, storage, and network infrastructure to meet demand.
- At the same time, they must keep the infrastructure secure and comply with regulations.

Stu: 15,723 Fac: 1,934 Staff: 10,680	300 Depts. 815 Bldgs.	Multiple 10Gb 14 Oper. Zones	Approx. 4500 APs	14 Pairs Juniper 5200 6 Pairs F5	Approx 50 Clusters 5 to 500 Nodes

Scientific Computing Community

- Universities have not been moving quickly enough.
- Groups typically acquire access to computing, storage, and networking resources via grants to pay for local dedicated resources.
- Access to remote resources is often limited by network bottlenecks, and shared supercomputers are often oversubscribed.
- Each group has its own, but insufficient, resources for its needs.
- Requirements are growing faster than can be met.

What is Wrong

Network Operations

Designs, deploys, and maintains campus-wide network infrastructure.

- Too much dependency on closed proprietary devices.
- Difficult or impossible to add features and customize services. Adding requested features takes too long and depends on one's relevance as a customer.
- Growing number of unique infrastructure devices.
- Performance bottlenecks introduced by middle boxes.
- Persistent network scaling problems.

  (Spanning-tree protocol, link redundancy)

Network Research

- Closed, proprietary, and production-critical nature of network makes it off-limits.
- A parallel infrastructure is unaffordable.
- No experimenting at scale.
- Limited to lab experiments and simulations.

Students

- Campus infrastructure has grown into a "production" infrastructure.
- Collaboration between campus IT and students has disappeared.
- Students have become passive users of the network.

Can write network and web applications, but cannot touch the infrastructure.

SDN-Based Private Cloud

An opportunity to bring a new paradigm to campus with virtualized scale-out clusters and SDN. Providing a more sustainable long term solution.

Using two enabling technologies:

- Multi-tenancy virtualized scale-out clusters
- Software Defined Networking (SDN/OpenFlow)

Proposed Private SDN Cloud Design

The proposed Stanford SDN-based Private Cloud will consist of several components:

1. Three virtualized computing clusters using SDN for inter-server and inter-VM networking, and OpenStack for VM orchestration.
2. Campus-wide sliceable/virtualized SDN backbone with approx. 10-15 switches.
3. Several sliceable/virtualized SDN edge networks.
4. SDN network control and management applications.

Logical View

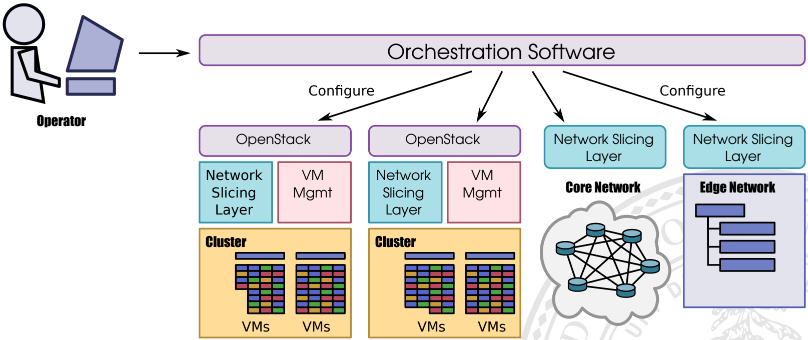

Figure : Logical View of the Infrastructure: Resource Orchestration

Research Groups - Radar Image Processing

- Raw radar data sets from archives around the globe to be reduced to individual radar images.
- 20-50 scenes per data set, each 500MB in size.
- Single-look complex images from scenes.
- Highly compute-intensive.

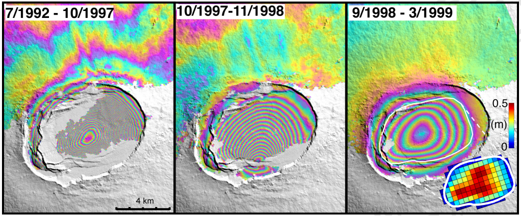

Radar Image Processing

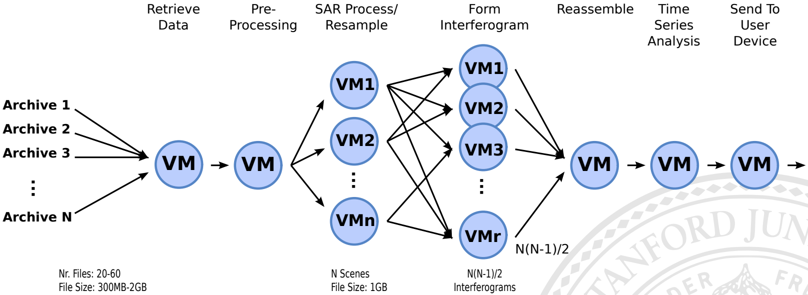

Figure : Radar Image Processing Re-engineered as Network of VMs

Network Operations

- Remove bottlenecks.
- Simplify management.
- Delegate control of segments and resources.
- Deploy network control and management as SDN applications.

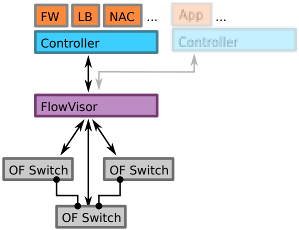

- Firewall
- Load balancing
- Guest access
- Bandwidth reservation

Private SDN Cloud Use Cases

Network Operations - Firewalling

- Cost of equipment and support.
- Mostly access-list type of policies.
- No application-layer filtering.
- Policies and objects per department.
- Must be in line with traffic.
- Creates bottleneck.
- Impacts more than a single user or department.

Network Operations - Load Balancing

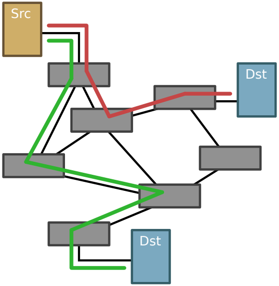

- Cost of equipment and support.
- Traffic bottleneck.

- Load-balancing is smart routing.
- With SDN, this function exists across the entire network.

Network Operations - Guest Access

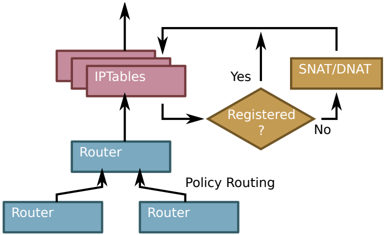

- Security through obscurity.
- IP address assigned based on registration.
- Requires policy-based routing to be configured across campus to direct traffic to a choke point.
- IPTables used to classify host MAC addresses or redirect to portal if not registered.

Network Research

- At-scale infrastructure to allow development and deployment of new ideas.
- Development of new tools to manage and debug this new infrastructure
- Testing and development on a production scale.

Components

- Private cloud will connect to the I2 100Gb innovation platform.
- SDN Core switches located in 10 ECHs across campus to provide full coverage.
- SDN edge switches located in buildings to connect departmental and research networks/buildings.
- Wireless connectivity included.

Stanford SDN Network Topology

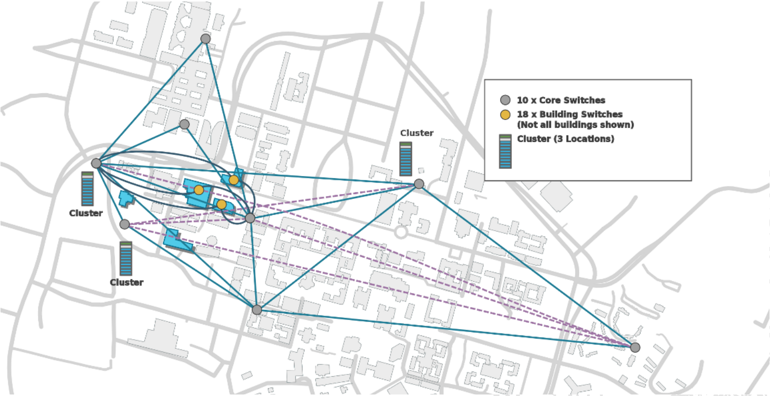

Stanford SDN Network Topology

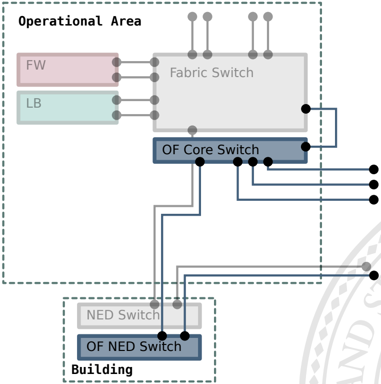

SDN/OpenFlow App Store

Without needing special expertise, users should be able to add functionality to their controllers.

- A central location for network apps.
- Enterprise Network Operators can download new features, install, and go.
- Anyone can post new apps.

A common interface between app and controller will be necessary.

Conclusion

The development of a SDN/OpenFlow-based private cloud on campus will

- Allow a campus to provide flexible and shared compute resources and infrastructure, tied to a highly programmable network.
- Make it possible to expand computing to off-campus public clouds.
- Enable the creation and deployment of new services on top of the network to the benefit of all campus users, not just researchers.

This NSF program provides university campuses an opportunity

- To gain more freedom and flexibility in how their network infrastructure is deployed and managed.
- To collaborate by inter-connecting campuses and extending these connections to public clouds.
- To develop and deploy the next generation campus network and services.
- To offer the scale, compute, storage, and network resources needed to continue future research.

Questions?