

University of Virginia



- Malathi Veeraghavan (MV), CC-NIE project PI
- Graduate students: Fatma Al-Ali, Shuoshuo (Shawn) Chen, Xiang Ji, Sourav Maji, Fabrice Mizero, Reza Rahimi, Xiaoyu (Sherry) Wang

CC-NIE project collaborators: Steve Emmerson, (UCAR), Ivan Seskar, Steve Decker, Joe Slezak (Rutgers U), Jerrold Robaidek and Dale Carder (U. Wisc)

Thanks to **NSF** for grants CNS-1116081, OCI-1127340, ACI-1340910, CNS-1405171, CNS- 1531065, ACI-0958998, and to **DOE** for grant DE-SC0011358.

Slides made for CC-NIE PI Meeting, Sept. 2015, Austin, TX, with some modifications



UVA research projects that use GENI

- Reliable multicast of filestreams
 - real-time weather data dissemination
- Large dataset transfers on rate-guaranteed VLANs
 - campus to campus
 - supercomputing center to campuses
- Traffic engineering
 - hybrid networking: best-effort and circuits/virtual circuits
- VNF for QoS mechanisms using Lagopus and OVS
 - policing and weighted fair queueing/scheduling
- Hadoop: assess potential for optical circuit switching in datacenter networks



Research project: Reliable multicast of filestreams

- Problem addressed: Current solution (LDM6) uses application-layer multicast to distribute real-time weather data from UCAR to 260 institutions
 - UCAR receives 11 GB/hr, but sends out 600 GB/hr
- Solution:
 - Use UVA's reliable file multicast transport protocol (FMTP)
 - Designed to run on top of OpenFlow Multicast
 - LDM7 implemented: FTMP + LDM6
 - LDM7 should save compute and link capacity
- Evaluation: using upto 80 VMs distributed across ExoGENI racks with 100 Mbps L2 multipoint topology to compare LDM6 and LDM7



Research project: Large dataset transfers on rate-guaranteed VLANs

- Problem addressed: even a small packet loss rate affects throughput of file transfers across high-speed WAN paths (large-RTT paths)
 - Cause: Without resource reservation, TCP sender has to dynamically estimate available bandwidth causing increases and decreases in rate
- Solution: Reserve resources on path; send at fixed rate
- Engineering:
 - Circuit TCP (TCP with congestion control disabled)
 - Linux traffic control (tc) to rate shape sending to match reserved path rate: Token Bucket Filter (TBF) vs Hierarchical Token Bucket (HTB)



Research project: Traffic engineering

- Problem addressed: Large, high-rate (alpha) flows fill up router buffers and add packet delays for real-time flows
- Solution:
 - Identify alpha flows from packet traffic within a network, and redirect to a separate queue on rate-guaranteed path
- Evaluation:
 - Use OVS and Lagopus; leverage Linux tc and DPDK QoS
 - Replay CAIDA traces, Skype traces, and iperf3 alpha flows
 - Does performance improve for real-time flows?



Research project: VNF for QoS mechanisms using Lagopus and OVS

- Problem addressed: At what scale can the data-plane networking functions of traffic policing and weighted fair queueing/scheduling be handled in software?
- Solution:
 - Use OVS and Lagopus software to emulate switches
 - Leverage Linux tc and Lagopus DPDK QoS functions
- Evaluation:
 - Use GENI VMs to create large numbers of flows with applications or CAIDA trace replay, and run Lagopus/OVS
 - What is the max. number of queues that can be supported?





UVA usage of GENI

Student	Resources used	Purpose	Positive experiences	Room for improvement
Fatma Al-Ali (fha6np@virginia.edu)	ProtoGENI, InstaGENI, ExoGENI	(i) Traffic-engr. (ii) Large transfers	(i) Resources (ii) Sudo (iii) Support	ProtoGENI: creating images InstaGENI: getting 1G between racks
Shawn Chen (sc7cq(@virginia.edu)	ExoGENI, InstaGENI Emulab	LDM7 FMTP	GENI provides "unlimited possibilities in designing and implementing tests." 80 VMs at four ExoGENI racks	Loss of connectivity; Insufficient VLAN IDs; Moving or parsing 2 TB worth of experimental data; Broadcast link stitching





UVA usage of GENI and DYNES

Student	Resources used	Purpose	Positive experiences	Room for improvement
Xiang Ji (xj4hm@virginia.edu)	DYNES ExoGENI	LDM6	Easy to use	GENI: slices going down and not easily renewable (need same set for repeated runs)
Reza Rahimi (gr5yf@virginia.edu)	InstaGENI Utah DDC	Lagopus for QoS VNF	Appreciates the great access to resources	Prevent accidents (such as reconfiguring address of ssh NIC); Reboots unsuccessful on multiple attempts
Xiaoyu Wang (xw5ce@virginia.edu)	ExoGENI	Hadoop jobs on hybrid DCN	Started recently	-



Summary

- Great to have such an extensive networking testbed. To say the least, it is an awesome resource and we can provide further details of our usage if the GENI financial planning effort needs it.
- After our experience with building a NSF-sponsored WAN tested called CHEETAH in the mid-2000s, the federation concept of GENI, and its support for HR to develop the tools and GUIs required for shared usage among (5000!) users is exactly what we should be doing

