SDN-based Virtual Machine Management for Research Testbeds

Richard Cziva
Internet2 Focused Technical Workshop
Miami, Florida, USA
1st April, 2015

Glasgow, Scotland, UK

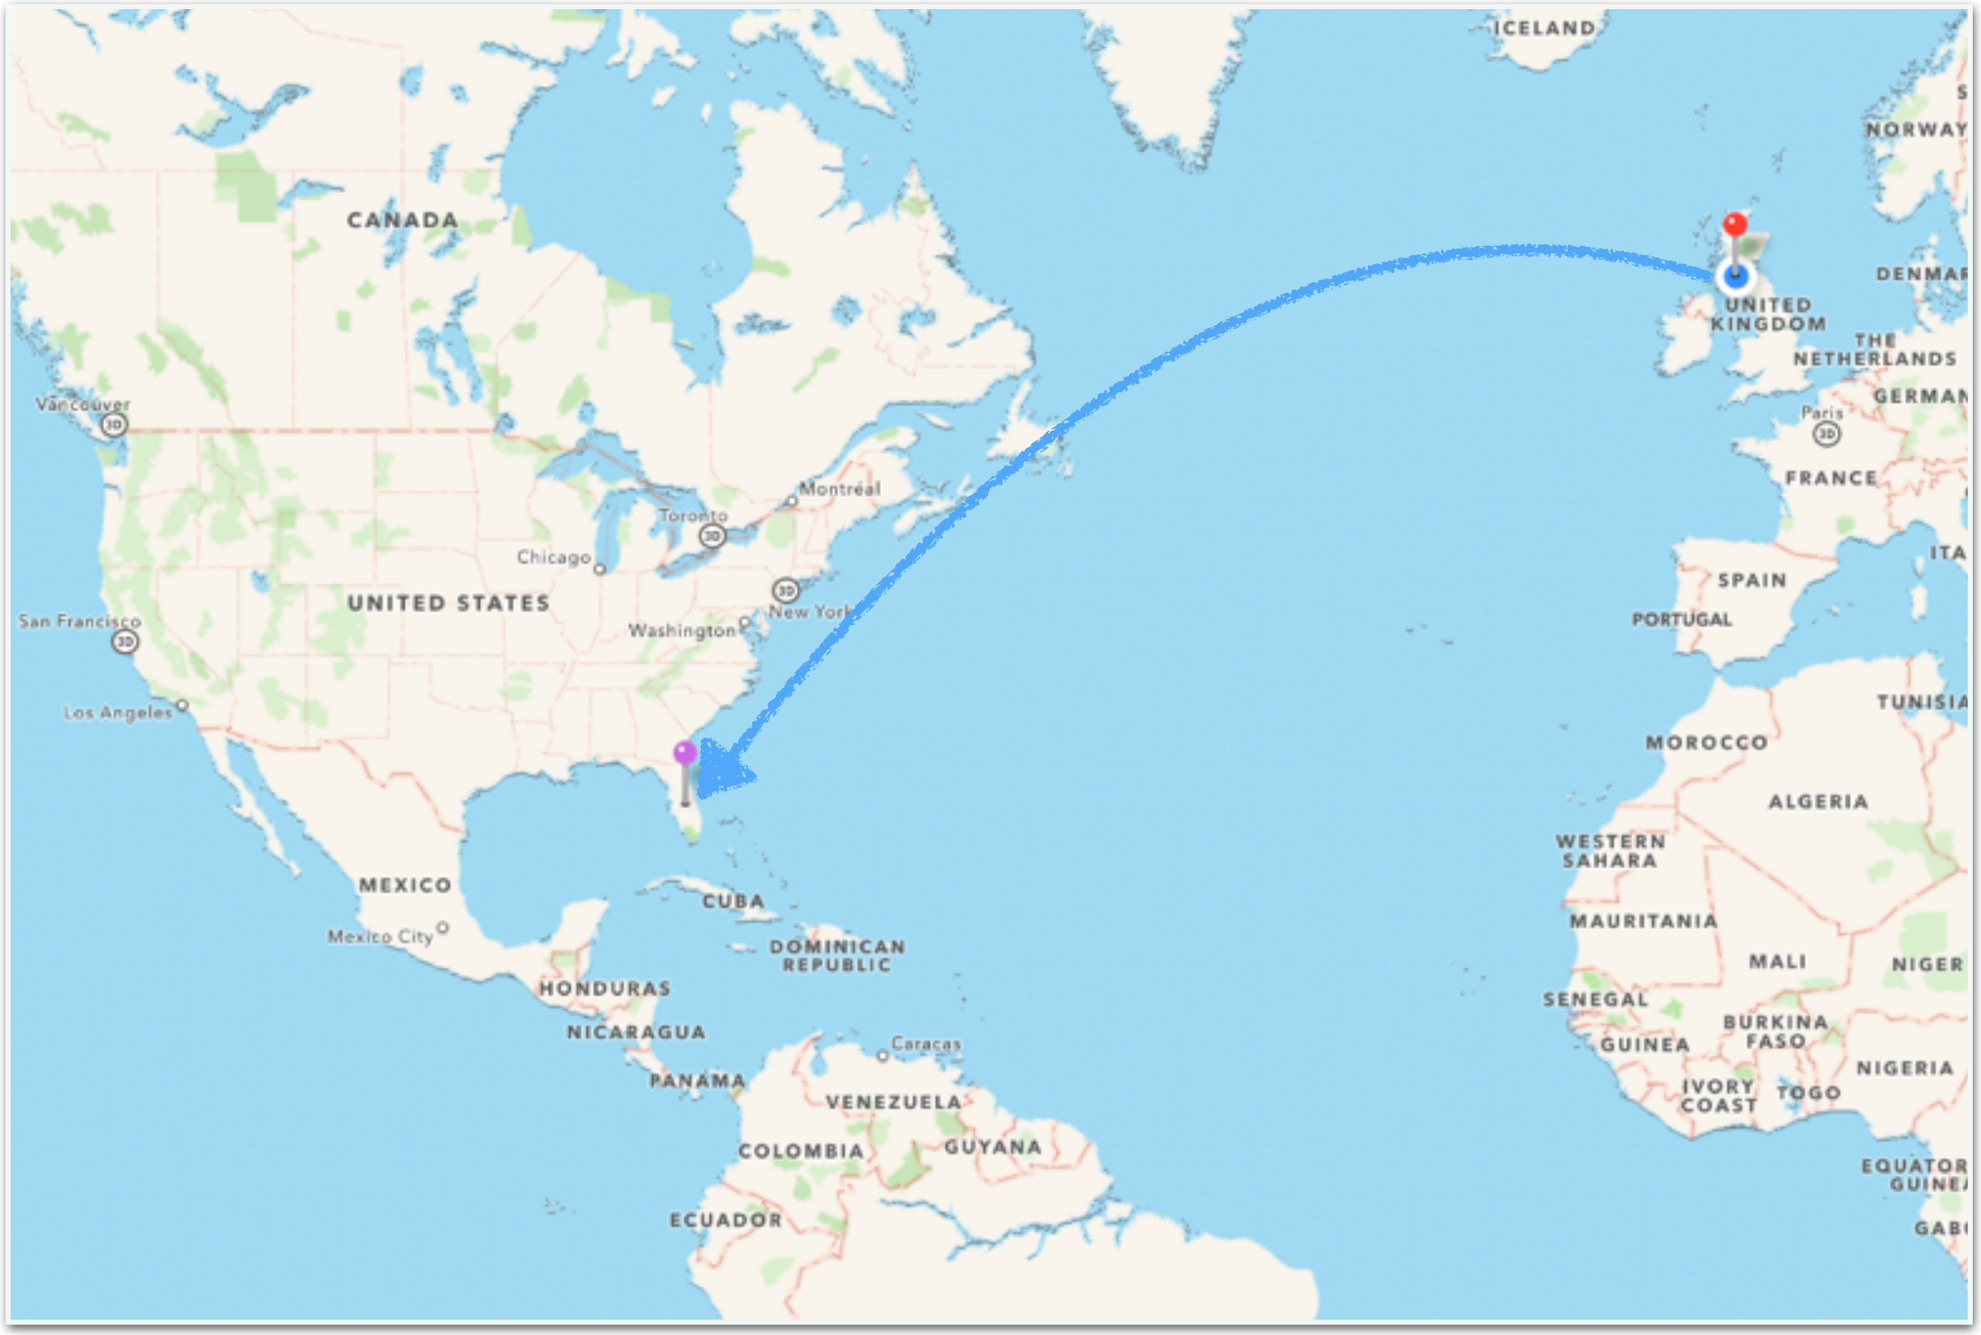

University of Glasgow

SDN Research in Glasgow

- Unified resource management for virtualized systems
- TCP parameter tuning using SDN
- Anomaly detection in critical infrastructures (e.g. Air Traffic Control Systems)
- Network Function Virtualization (Glasgow Network Functions)

IEEE CAMAD

- 20th IEEE Workshop on Computer-Aided Modeling and Design of Communication Links and Networks (CAMAD)

  - we organise a special session: "Service and Infrastructure Management for Cloud, Virtualized and Next Generation Networks”

- location: University of Surrey, UK

- conference: 7-9 September 2015

- submission: 1 May 2015

Motivation

In virtualized systems, **server** and **network** resources have disjoint control mechanisms.

Motivation

Server Resources

- VM
- Energy
- CPU
- Memory
- Hypervisors

Network Resources

- Switches
- Routing
- Middleboxes
- Policies

A **unified server-network control** mechanism is needed

Motivation

- The lack of inter-operation between various resource management schemes leads to
  - network-wide congestion
  - bandwidth bottlenecks
  - inefficient network usage
  - …

Our work

- We are building a unified control framework for virtualized systems that manages different types of resources (server, network, storage, etc)

    1. detects resource management issues by collecting and analysing measurement data

    2. suggests solutions / performs actions

- Initial work has focused on Cloud Data Centres

- First step: optimise network-wide link utilisation **and** VM allocation from a central point

S-CORE

- Scalable Communication Cost Reduction

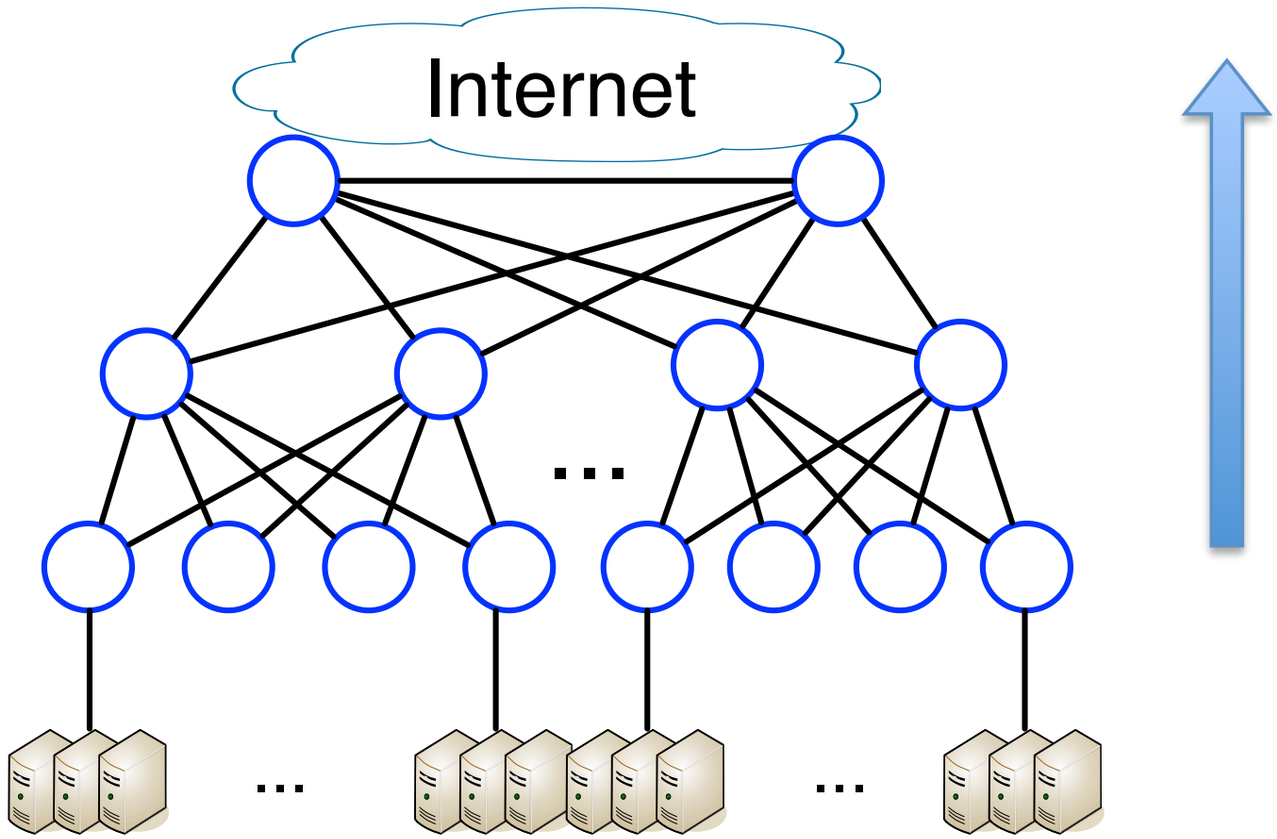

oversubscription ratio
higher link speed

- **Fung Po Tso, Konstantinos Oikonomou, Eleni Kavvadia, Dimitrios P. Pezaros**
- **Scalable Traffic-Aware Virtual Machine Management for Cloud Data Centers**
- **IEEE ICDCS 2014**

S-CORE

communication cost
for an allocations *A*

$$C(u,v) = \lambda(u,v) \sum_{i=1}^{\ell^{\mathcal{A}}(u,v)} c_i.$$

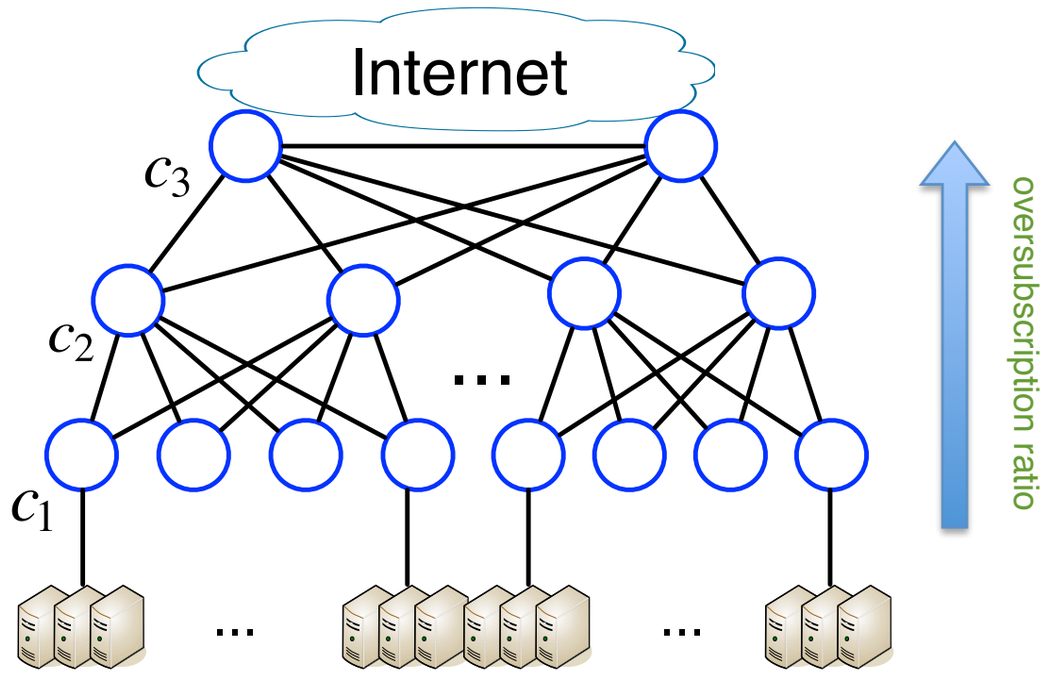

$\lambda(u,v)$ is the traffic load per time unit exchanged between VM u and VM v

link weight, c, can be set according to hierarchy, bandwidth, or policies but generally $c_1 < c_2 < c_3$

$l(u,v)$ communication level between VM u and VM v

S-CORE

Eventually,
overall communication cost

$$C^{\mathcal{A}} = \sum_{\forall u \in \mathbb{V}} \sum_{\forall v \in \mathbb{V}_u} \lambda(u,v) \sum_{i=1}^{\ell^{\mathcal{A}}(u,v)} c_i.$$

Thus, centralised optimal

$$C^{opt} \leq C^{\mathcal{A}}$$

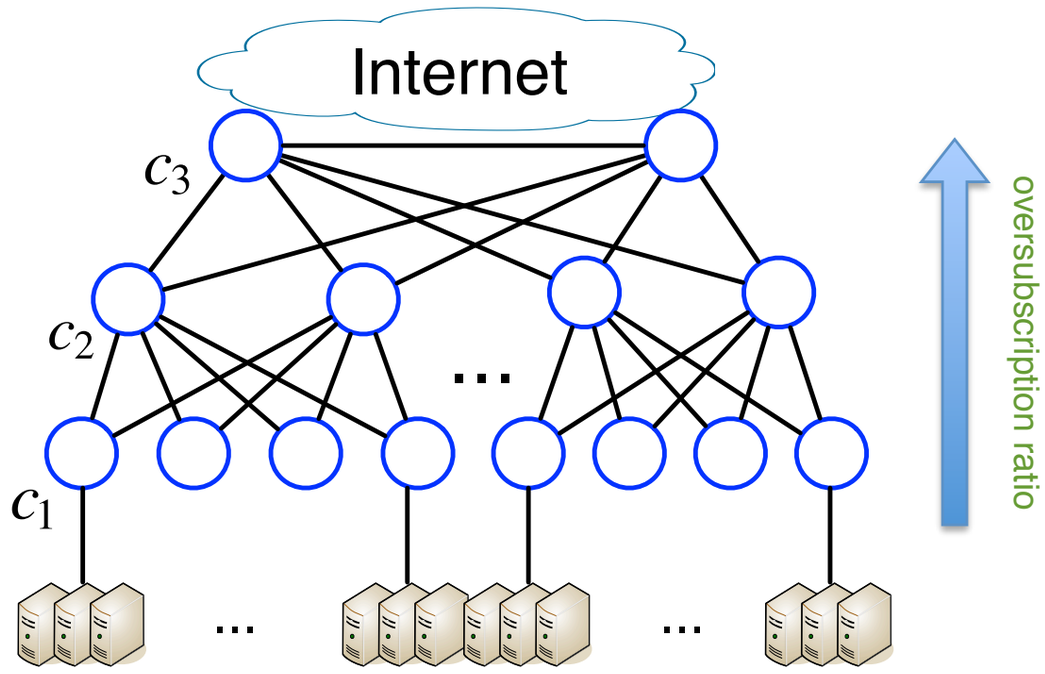

$\lambda(u,v)$ is the traffic load per time unit exchanged between VM u and VM v

link weight, c, can be set according to hierarchy, bandwidth, or policies but generally $c_1 < c_2 < c_3$

$l(u,v)$ communication level between VM u and VM v

OpenFlow

- A flow entry contains match rules, actions and **statistics**

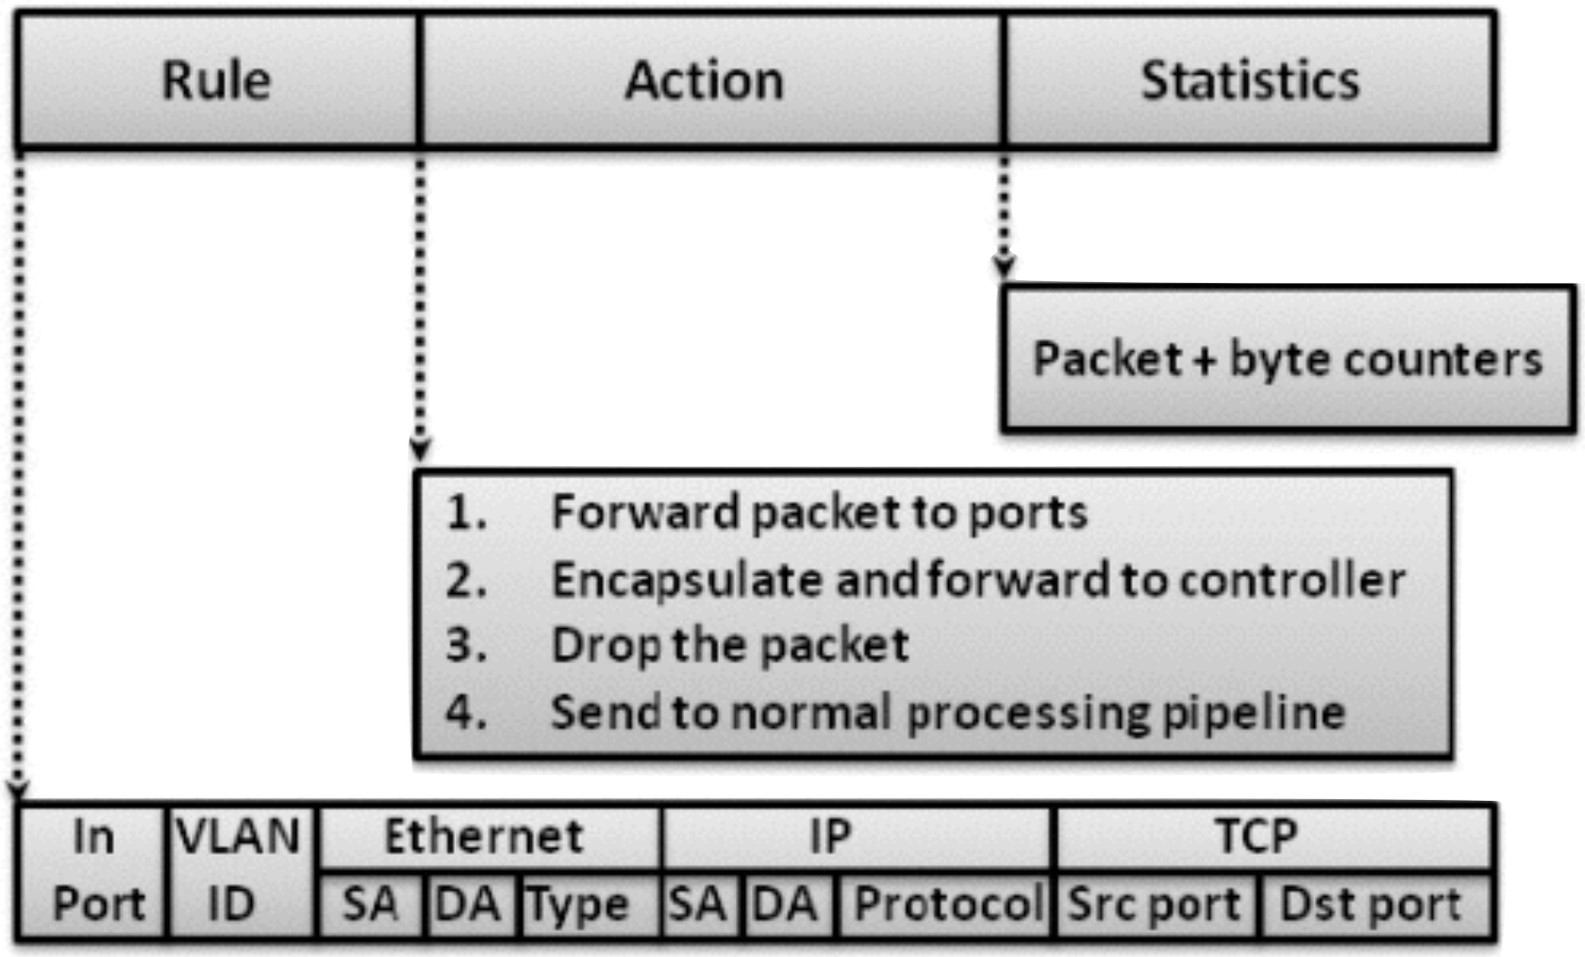

System design

- a central SDN controller module (Ryu)
  - collects flow statistics periodically (Statistics Request -> FlowStatsReceived)
    - managing topology, switches, hosts, link weights
    - orchestration of VM migration
- Hypervisors should support VM migration

Evaluation

Evaluation

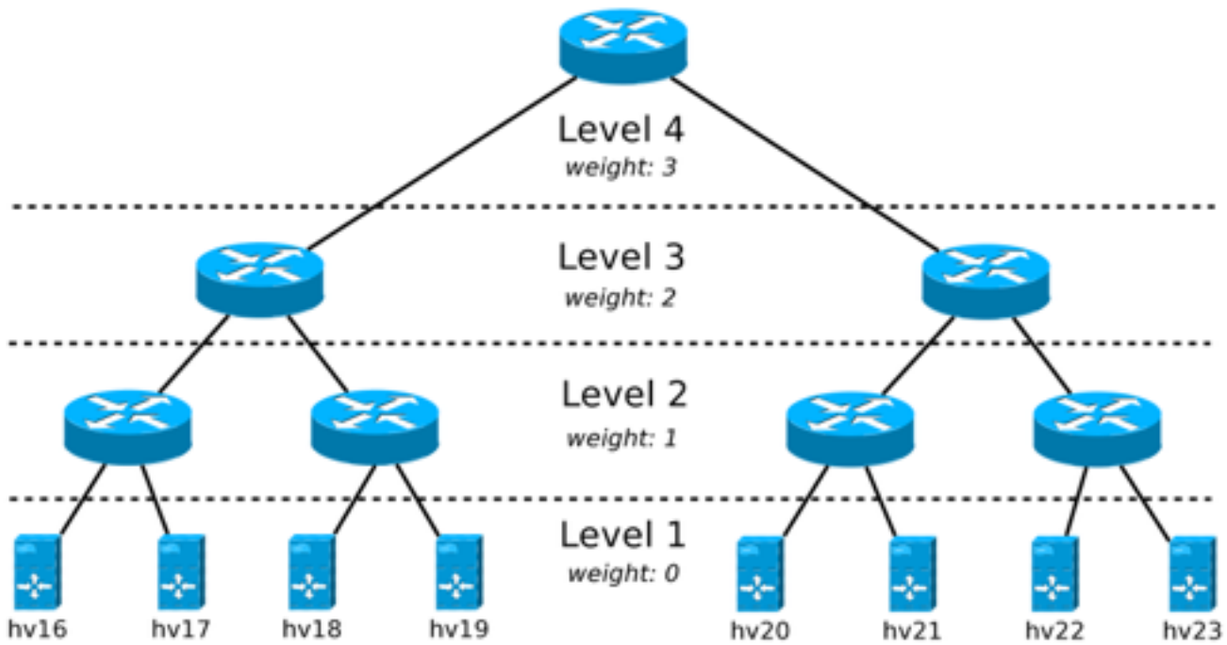

TABLE II. INITIAL TRAFFIC GENERATION IN OUR TEST SETUP.

Source VM	Source HV	Destination VM	Destination HV	Link cost
10.0.0.1	hv16	10.0.0.6	hv17	2
10.0.0.2	hv16	10.0.0.10	hv19	6
10.0.0.3	hv16	10.0.0.23	hv23	12
10.0.0.6	hv17	10.0.0.11	hv19	6
10.0.0.9	hv18	10.0.0.22	hv23	12
10.0.0.21	hv23	10.0.0.5	hv17	12

Evaluation

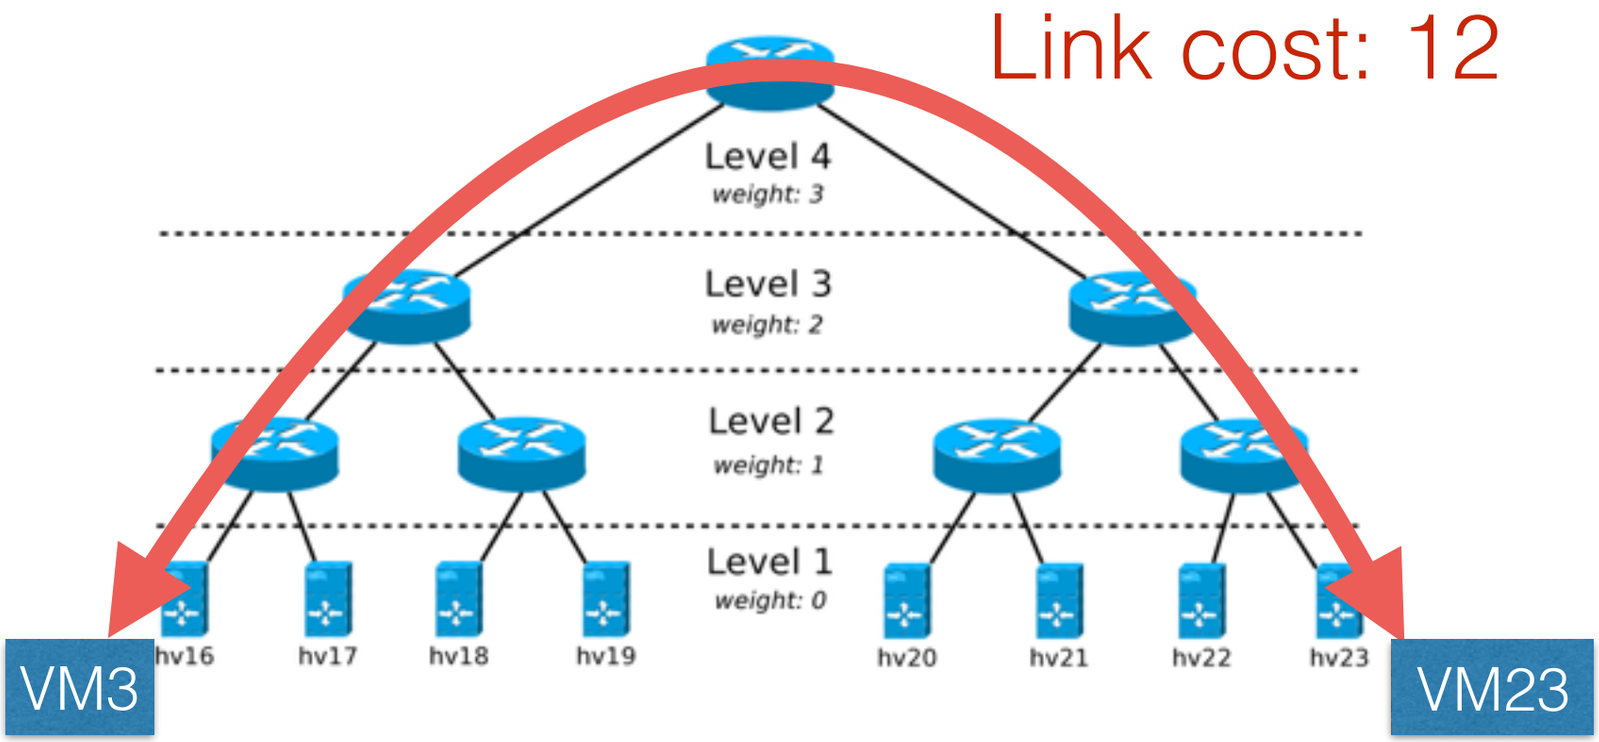

TABLE II. INITIAL TRAFFIC GENERATION IN OUR TEST SETUP.

Source VM	Source HV	Destination VM	Destination HV	Link cost
10.0.0.1	hv16	10.0.0.6	hv17	2
10.0.0.2	hv16	10.0.0.10	hv19	6
10.0.0.3	hv16	10.0.0.23	hv23	12
10.0.0.6	hv17	10.0.0.11	hv19	6
10.0.0.9	hv18	10.0.0.22	hv23	12
10.0.0.21	hv23	10.0.0.5	hv17	12

Experimental Results

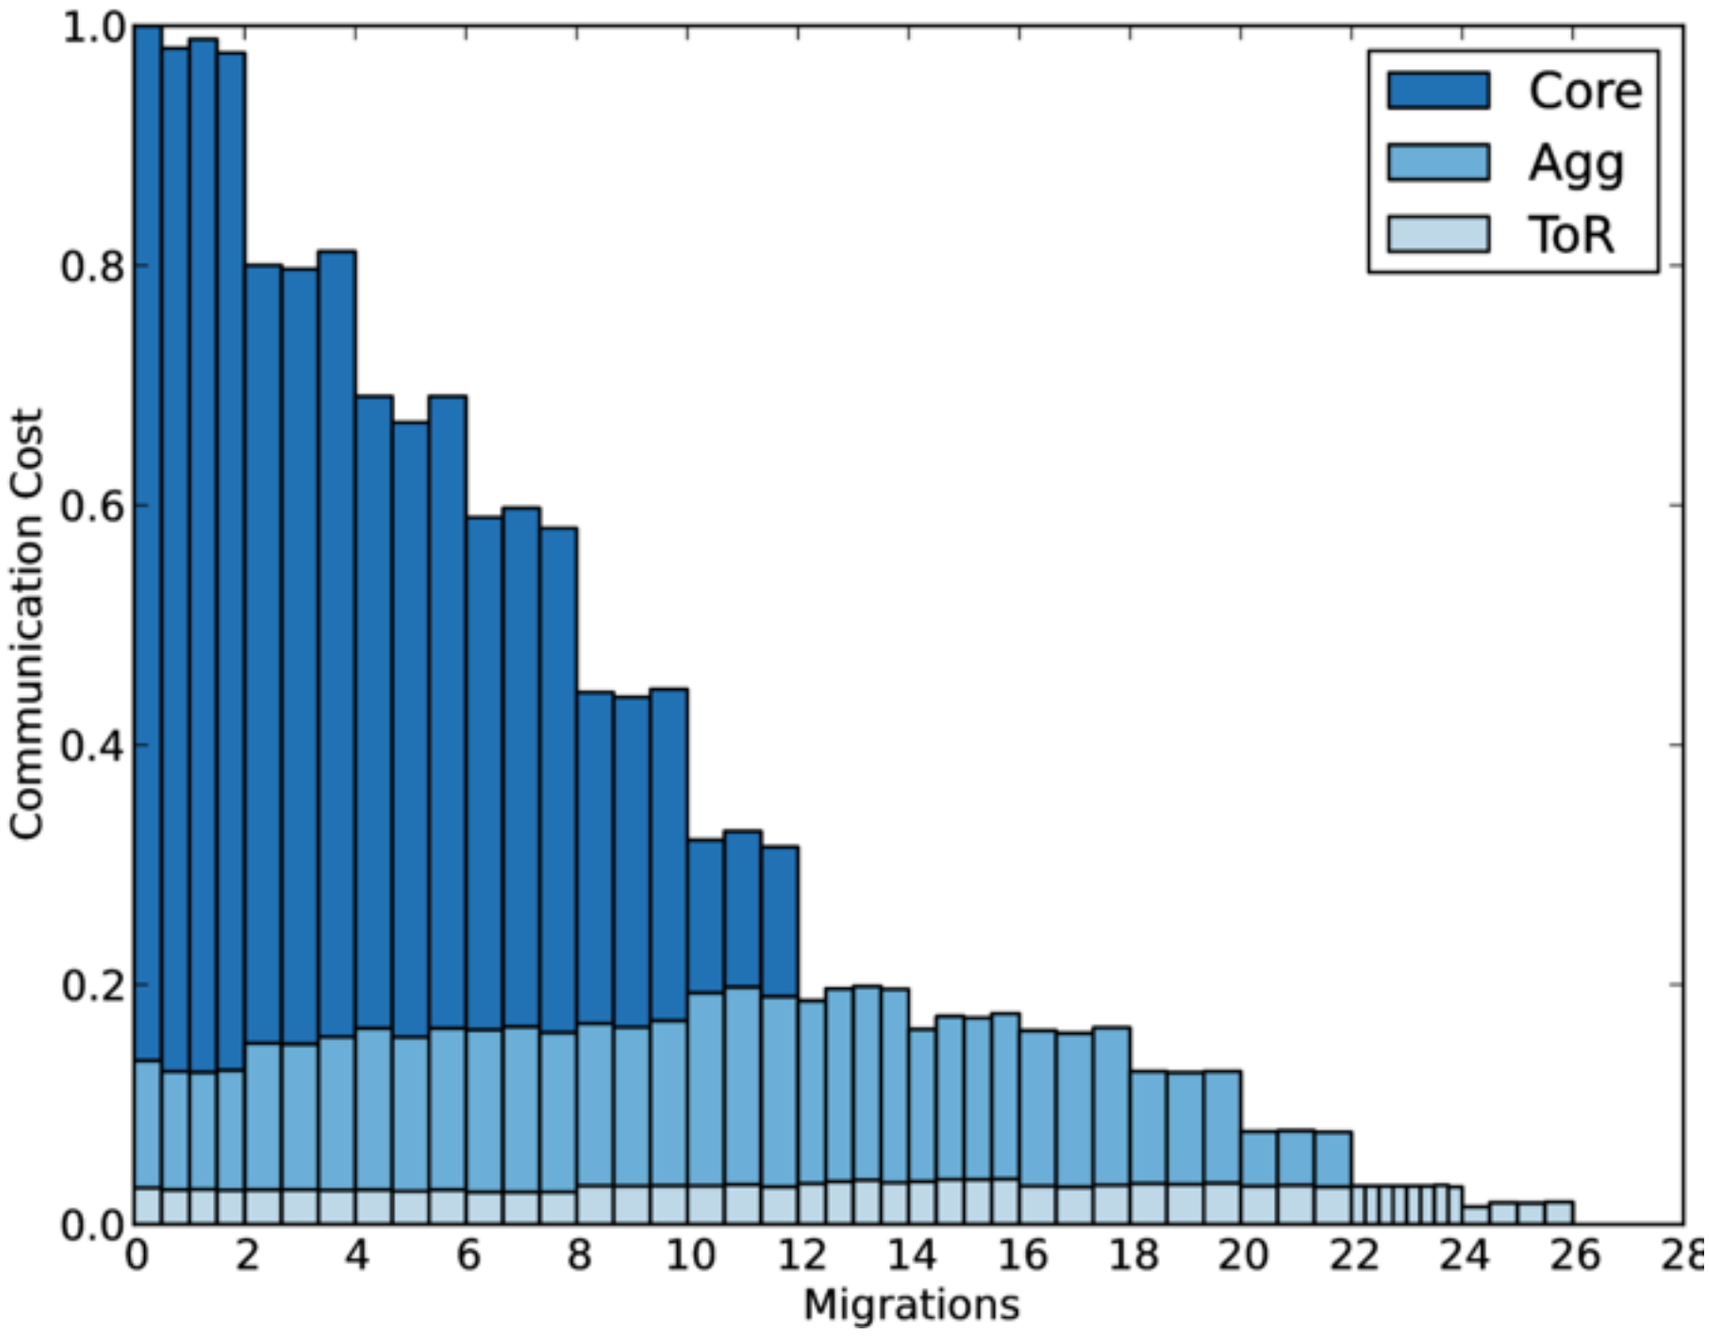

Experimental Results

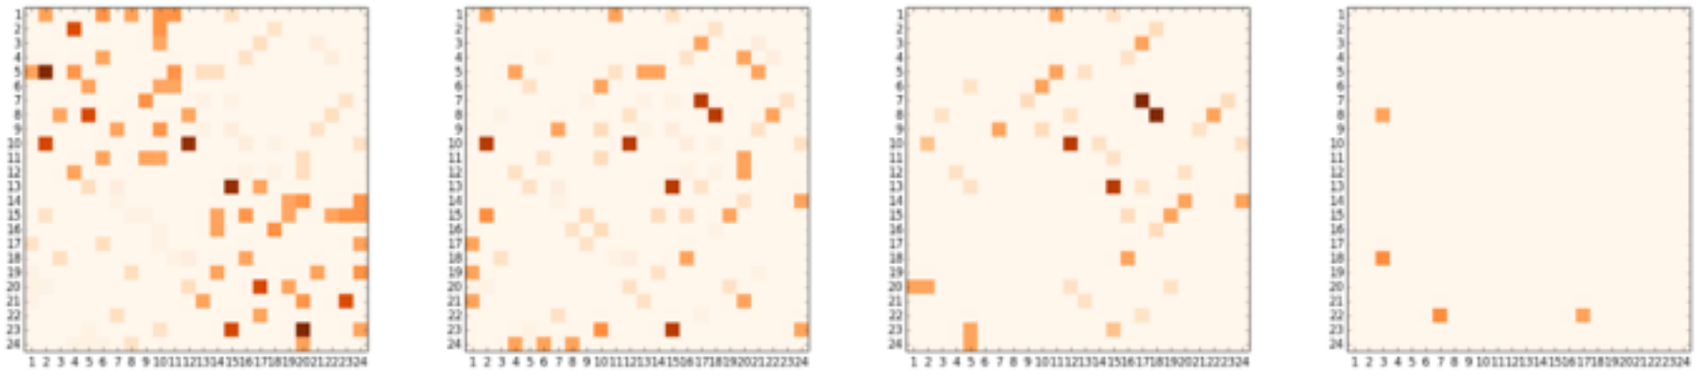

(a) Best fit - start (b) Best fit - 33% (c) Best fit - 66% (d) Best fit - stable

(e) Round Robin - start (f) Round Robin - 33% (g) Round Robin - 66% (h) Round Robin - stable

(i) Lookahead - start (j) Lookahead - 33% (k) Lookahead - 66% (l) Lookahead - stable

Research Testbeds

- Research and Education (R&E) Testbeds have started to utilize
  - Virtualized End-to-End Network Infrastructure
    - Virtual Machines
    - Virtual Storage
    - Multiple links between resources with various costs (speed) associated

Challenges

- R&E testbeds involve multiple administrative domains and multiple federation levels that raise
  - technical issues
  - operational issues
  - legal issues

Ongoing Research in Federation

- resource description
- resource registration
- resource access control
- service level agreements
- resource usage policies
- **unified resource management**
- resource lifecycle management
- legal framework
- operational procedures
- business frameworks

OpenCloud @ Internet2

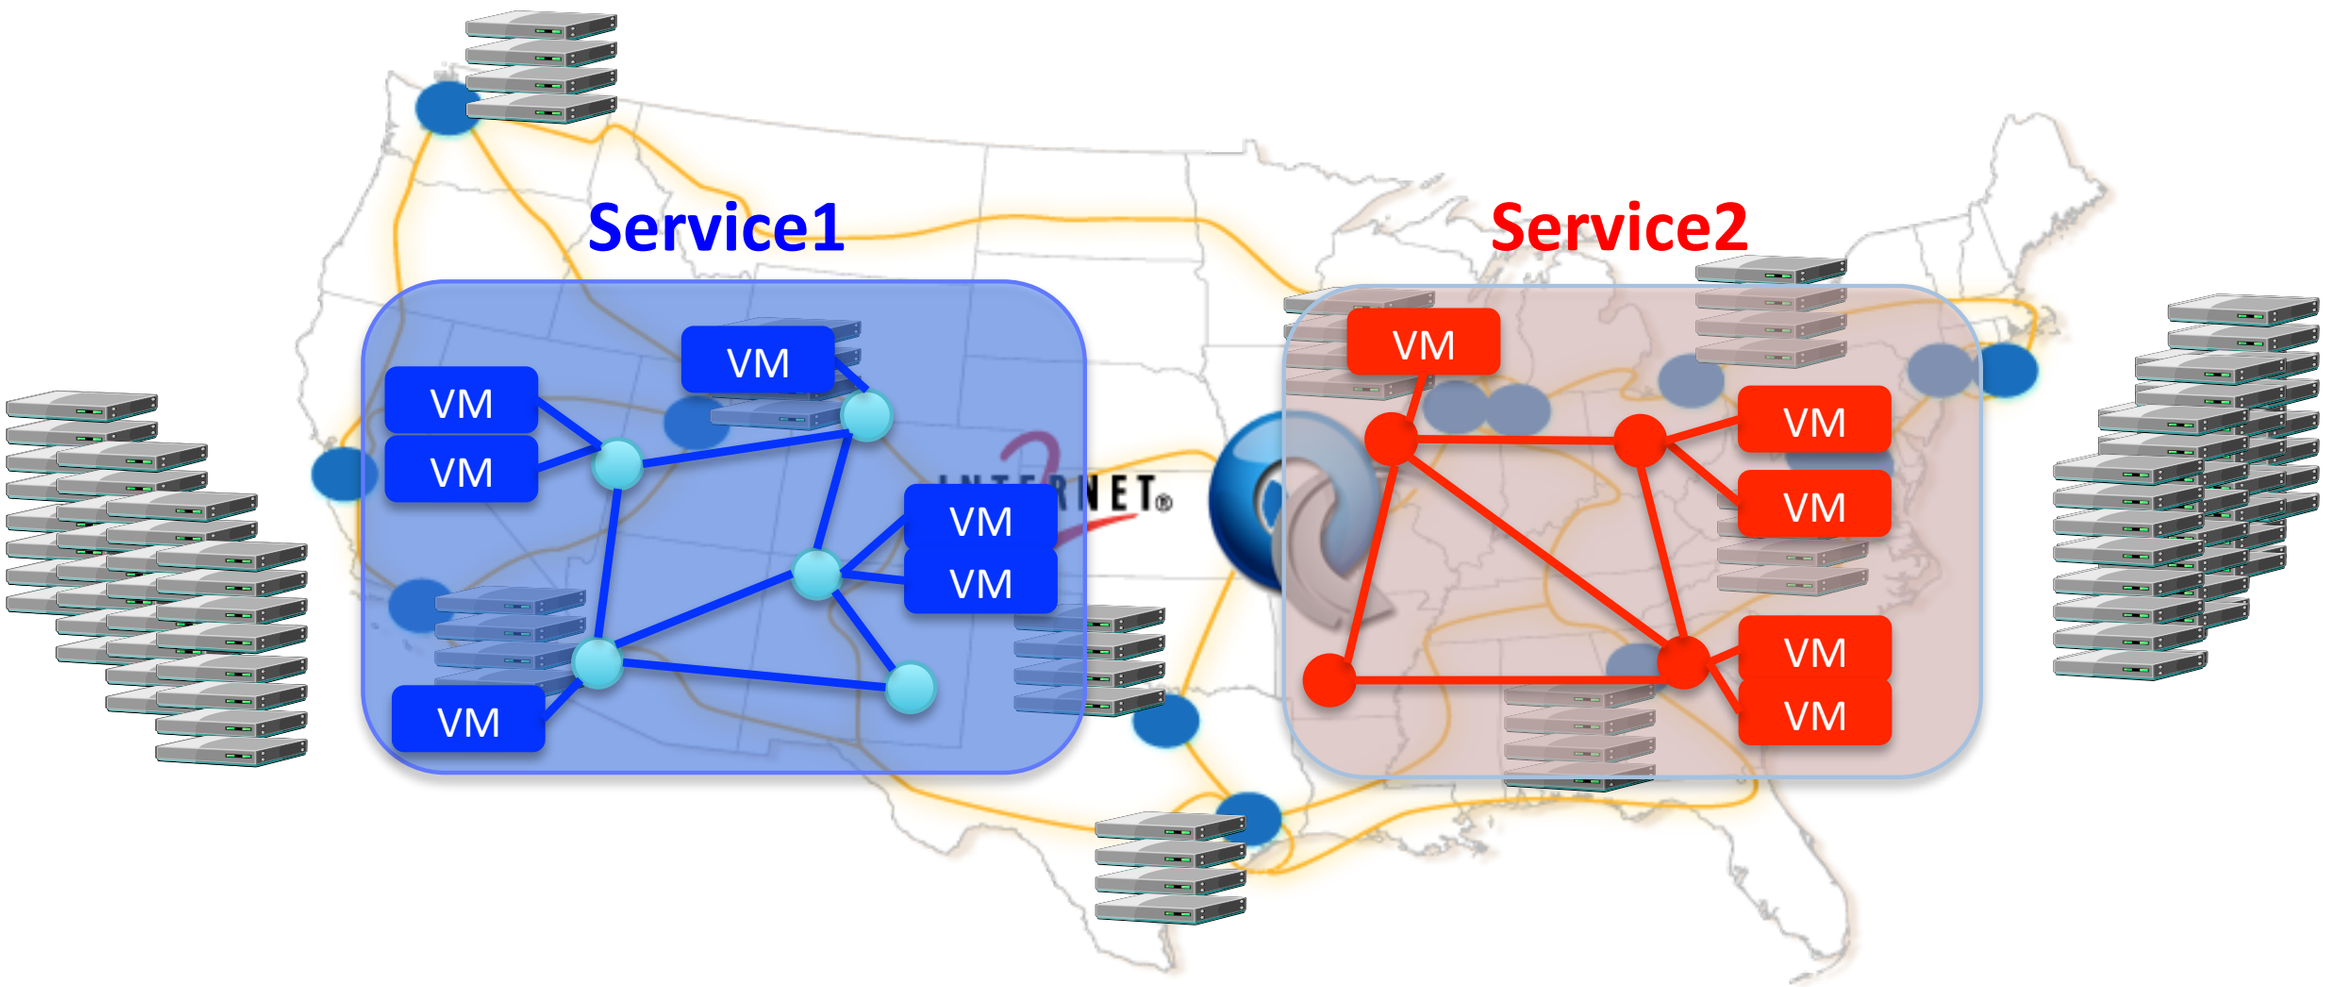

Copyright: Larry Peterson, Open Networking Lab

Conclusion

- We are building a unified control framework for virtualized systems (Clouds, R&E testbeds, etc) that manages server and network resources in a unified way

  1. detects resource management issues by collecting and analysing measurement data

  2. suggests solutions / performs actions

- **We are interested in your feedback and the applicability of such resource management scheme for R&E testbeds**

Thank you
for your attention!

Richard Cziva
http://richard.systems
r.cziva.1@research.gla.ac.uk

References

- Cziva, R., Stapleton, D., Tso, F. P., and Pezaros, D. P. (2014) SDN-based virtual machine management for cloud data centers. In: Third IEEE International Conference on Cloud Networking (IEEE CloudNet), 8-10 Oct 2014, Luxembourg [best paper award].
- White, K., Pezaros, D., and Johnson, C. (2014) Using Programmable Data Networks to Detect Critical Infrastructure Challenges. In: International Conference on Critical Information Infrastructures Security (CRITIS'14), October 13-15, 2014, Limassol, Cyprus.
- Tso, F. P., Oikonomou, K., Kavvadia, E., and Pezaros, D. P. (2014) Scalable traffic-aware virtual machine management for cloud data centers. In: IEEE International Conference on Distributed Computing Systems (IEEE ICDCS), 30 Jun - 3 Jul 2014, Madrid, Spain.
- Tso, F.P., and Pezaros, D.P. (2013) Improving data centre network utilisation using near-optimal traffic engineering. IEEE Transactions on Parallel and Distributed Systems (IEEE TPDS), 24 (6). pp. 1139-1148. ISSN 1045-9219. June 2013.
- Tso, F.P., Hamilton, G., Weber, R., Perkins, C., and Pezaros, D. (2013) Longer is better: exploiting path diversity in data center networks. In: IEEE International Conference on Distributed Computing Systems (IEEE ICDCS), 8-11 Jul 2013, Philadelphia, PA, USA.
- Tso, F.P., Hamilton, G., Oikonomou, K., and Pezaros, D. (2013) Implementing scalable, network-aware virtual machine migration for cloud data centers. In: IEEE International Conference on Cloud Computing (IEEE CLOUD), 27 Jun - 02 Jul 2013, Santa Clara, CA, USA.
- Jouet, S., and Pezaros, D. (2013) Measurement-Based TCP Parameter Tuning in Cloud Data Centers. In: IEEE International Conference on Network Protocols (IEEE ICNP), 7-11 October 2013, Gottingen, Germany