Software Defined Infrastructure: Challenges and Opportunities

Rajkumar Kettimuthu, Argonne National Laboratory

Bio:

Rajkumar Kettimuthu is a Computer Scientist in the Mathematics and Computer Science Division at Argonne National Laboratory, and a Senior Fellow in the Computation Institute at The University of Chicago and Argonne National Laboratory. His research is focused on software defined networking, high-speed transfer of large-scale data, and large-scale data analysis. He has co-authored several articles in the above-mentioned areas. He is a recipient of R&D 100 award. He is a Co-PI of DOE Software Defined Network Science Flows project. He led a demo at SC15 on using software defined networking techniques to do adaptive QoS for Science Flows. He is a senior member of both IEEE and ACM.

Background:

Significant work has been done in the past on dynamic provisioning of network resources. NSF funded DRAGON and DOE funded OSCARS project enable the provisioning of wide-area network resources. The tools developed by these projects enable one to provisioning bandwidth from the border router of one institution to the border router of another institution. Projects such as TeraPaths and LambdaStation (both funded by DOE) and DYNES (funded by NSF) attempt to extend the circuit all the way to the end hosts. But still the process to setup an end-to-end circuit involves complex authentication process and manual steps including obtaining a certificate from a well known certificate authority, emailing wide-area network administrators and getting approval from the user's site network administrator. Because of this tedious process, only very few applications make use of the network reservation capabilities. Software-defined networking can help address some of the issues but still a number of challenges remain.

Challenges and Opportunities:

- How do we do the resource slicing?
  - Coarse-grain versus fine-grained
  - Coarse-grain means hierarchical allocation i.e., slices are allocated to a collaboration or a service and the collaboration or service will in turn manage the allocation among its users.
- Authentication and authorization:
  - How do we simplify authentication and authorization?
    - Can we integrate with federated identity such as InCommon? How do we handle collaborations?
  - 'Connection tuple based enforcement' versus 'token based enforcement'
    - Some applications may not know the entire connection tuple at the time of reservation
      - Rather than forcing the user to provide the connection tuples at the time of the reservation, can we do token based enforcement
      - What are the implications of this to the transport and network protocols?
- Operating systems and Quality-of-Service:
  - In the current HPC environment, most common deployments have data transfer nodes and compute nodes mount and share a global shared parallel storage system, which is connected through a storage area network. How do we do end-to-end QoS for data transfers in such scenarios?
  - In such scenarios, it may not possible to ensure contention-free access to the storage system. And the data transfer nodes may not support scheduling CPU resources. The end result is that some portion of the reserved bandwidth on WAN and LAN may go unused when circumstances conspire to throttle the throughput of the hosts' connection. How do we ensure that this bandwidth is not effectively lost? Can SDI operating system dynamically alter the QoS parameters to give the unused bandwidth of a flow to other flows in priority order? Can we do this in a transparent manner i.e., if the flow underutilizing the reserved bandwidth is able to ramp up after a certain time, can we ensure that it can get bandwidth up to its reserved level without any noticeable delay?
  - How do we support flow-level QoS in SDI?