Submitter: Inder Monga, CTO, Energy Sciences Network (ESnet)
Contact: imonga@es.net

Towards network predictability – a missing ingredient

Challenge

The ‘separation of concerns’ in the internet data layers has worked well to enable the Internet bring data communication capabilities to billions of users and devices. This ‘best effort’ connectivity works for most applications and users, and most of the time. The challenge has been to build predictable application workflows, that leverage diverse wide-area resources of compute/storage/network, especially due to the wide variability and unpredictability of a shared network and I/O fabric that extends end-to-end. With the simple, but effective socket interface, applications do not have the ***levers*** to monitor, manage and change the level of performance their data flows get from the network. These levers might not be relevant for ‘mouse’ flows, that send low amounts of non-critical data but extremely relevant when large flows, long data transfers or time-critical transactions are needed over the shared infrastructure. The community tried to solve these objectives with various initiatives like Active Networks, Bandwidth-on-Demand, but none of those approaches has really solved this unmet need, mainly because predictability was not a principle that was front and center to the data movement design.

Vision

The vision is to build a self-correcting feedback loop between the application and the network, so higher level of predictability can be attributed to network- and data-intensive application interactions. Perfection, aka. high-precision predictability, might take more effort than needed, so the objective is to get much better at meeting the goals with reasonable error bars. Success can be attributed when within reasonable delta one can predict the time it is going to take a data transfer between two nodes on the network, assuming shared usage and varying network conditions. For certain applications, it can be imagined that ***predictability*** maybe the more important parameter than total throughput. The hypothesis is with more predictability, scientific workflows and other distributed workflows (like Superfacilities) can be easily automated requiring minimum manual human attention. In addition to science workflows, as the Internet starts being used to control critical facilities over this shared infrastructure, like self-driving automobiles, smart grids etc, predictability as a network value will apply to even a larger space.

Approach

To solve this problem, we need to bring fields of applications, systems, network control, performance management and build synergistic control loops of interaction between them. Research is needed to explore *closed vs. open feedback loops*, and determine the right, simple structure to apply to data movement. *Network analytics* as a field needs to be improved significantly to provide the kind of feedback application or data movement algorithms will be needed. In addition to the control and feedback aspects of the network, we need protocols that adapt well to such constraints – so exploring adaptive use of non-TCP protocols for data transfer will be important. *Building an integrated system* that ties data transfer mechanisms with network services is important – and those services need to have appropriate interfaces to share back information like service status. The feedback loop can be controlled with network programmability to adapt to changing conditions. This area is complex but would result significantly interesting research and impact to global, distributed DOE science.

Conclusion

After listening to many frustrations of infrastructure owners, scientists and science infrastructure providers, it is clear that predictability will greatly enhance automation and speed of scientific discovery. It is not a trivial problem, and requires cross-discipline collaboration and approaches,

Submitter: Inder Monga, CTO, Energy Sciences Network (ESnet)
Contact: imonga@es.net

integrating network data transfer protocols, multi-domain federation, easy intent-like interfaces to discover and access data like Named-data networking, and network analytics over performance measurement data.