

Use Cases for High Bandwidth Query and Control of Core Networks

draft-bernstein-alto-large-bandwidth-cases-00.txt

Greg Bernstein, Grotto Networking
Young Lee, Huawei


Motivation

- To enable the application layer to more fully and efficiently utilize the capabilities of networks on top of a flexible GMPLS/PCE infrastructure
- Road blocks:
 - Core networks typically consist of two or more layers, while applications are typically only know about the IP layer and above.
 - GMPLS routing exposes full network topology information which tends to be proprietary to a carrier or require specialized knowledge and techniques, e.g., RWA
 - GMPLS signaling interfaces are defined for either peer GMPLS nodes or via a user network interface (UNI)

Outline

- End System Aggregation Use Case
 - See draft
- Data Center to Data Center (DC-to-DC) Use Case
- Generic Optimization Problem
- Cross Stratum Interfaces

Data Center to Data Center Networking



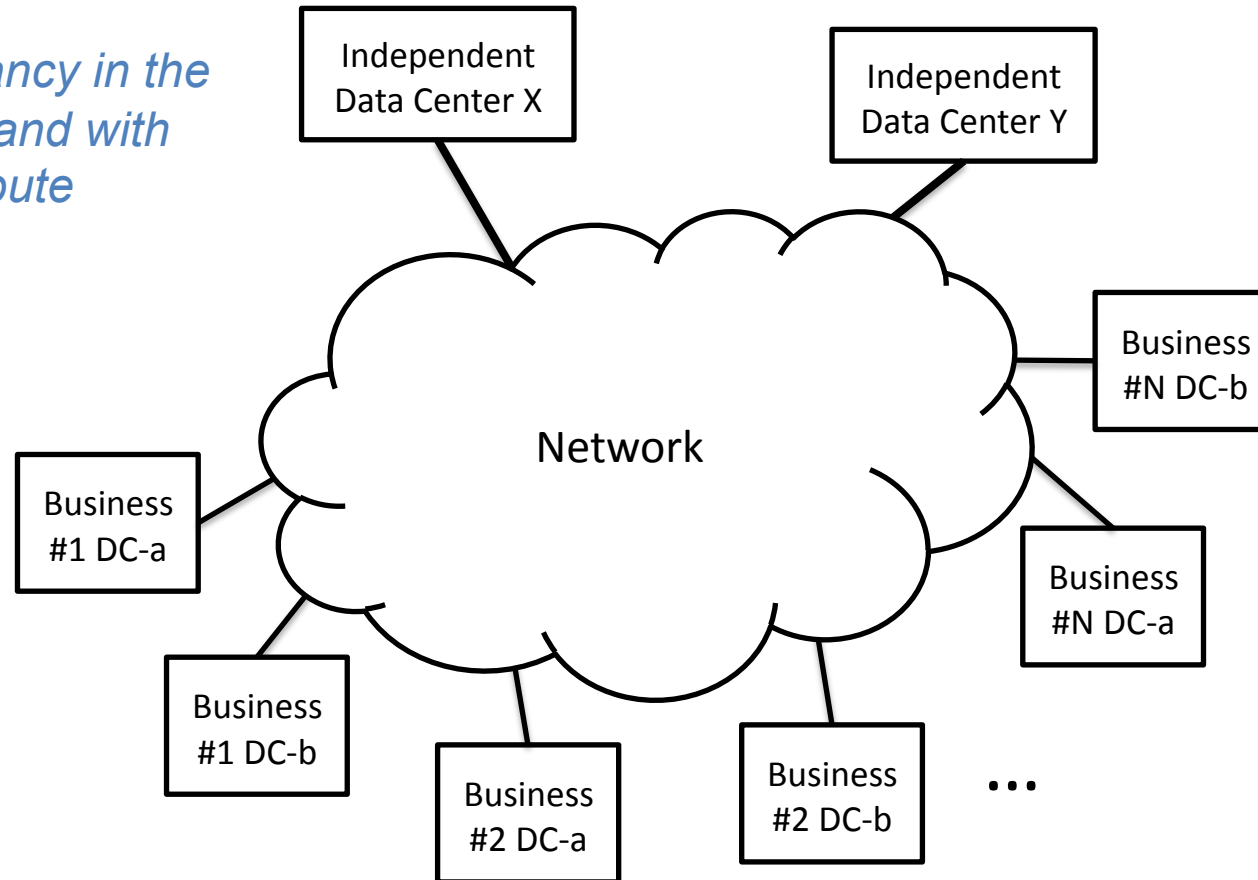
- For our purposes here we consider a data center any computation facility with **significant** access bandwidth to the network (this does not include relatively low bandwidth internet clients)

Cross Stratum Optimization Opportunities

- Compute Pricing
 - VM instances: Reserved, On Demand, Spot
 - Pricing can vary with demand (spot market) and time of day (e.g., due to changing power rates)
- Compute Constraints
 - Limits at local data centers, maintenance windows, ...
- Network Constraints
 - Path bandwidth and latency between sites at a given time
- Network Costs
 - Different providers, Congestion based pricing, etc...

Reliability and Recovery

Redundancy in the network and with the compute services



Reliability and Recovery

- Server Failure Recovery
 - within a data center → data center internal response
 - between data centers → network response to shift traffic to alternative data center
- Network Response
 - MPLS and GMPLS provide the ability to switch “flows” quickly minimizing downtime
- Network Failures
 - When network capacity is impacted due to link/node failures, network may be able to work with application to detour clients to other data centers or “shed load” in a controlled fashion.
- Coordinated Maintenance Operations

Cross Stratum Interfaces

Applications control/management doesn't want to deal with network details and networks don't want to always share them!

- Network Query → *Can base on ALTO work*
 - What amount of bandwidth might an application be able to get between two or more locations?
 - How much might that cost?
 - Now or some time in the future? Some applications are less time sensitive than others...
- Network Reservations
 - On demand or advanced reservations
 - Backup flow reservations
- Recovery
 - Activate pre-reserved recovery flows
 - On demand (perhaps higher cost) recovery action (optimize for cost after we have recovered)