Inter-domain QoS: MESCAL approach, functional architecture and solution options

David Griffin, *University College London*
dgriffin@ee.ucl.ac.uk

Motivation

- The Internet consists of a large collection of interconnected but *independently* operated networks
- ***Intra***-domain QoS problem is largely solved, if not yet widely deployed
- Biggest obstacle to ***inter***-domain QoS is suitable business agreements between providers in the absence of any central control or regulatory environment
- How can *end to end agreements* be negotiated, implemented and enforced?

Current Situation

- Intra-domain QoS mechanisms for traffic engineering and routing exist:
  - Intserv, Diffserv, MPLS, IP-based TE, over-provisioning
- QoS capabilities limited in scope to ingress-egress routers across a single domain
- Current inter-domain relationships are based on reachability only and are unable to support inter-domain QoS
  - Note: IETF is currently discussing mechanisms for signalling inter-domain LSPs, but a diffserv-based layer 3 solution would be more scalable for mass-market services

MESCAL project

Objective: to specify and validate scalable, incremental solutions to enable the flexible provisioning of inter-domain QoS across the Internet

- Industrial partners:
  - France Telecom R&D (Coordinator)
  - Thales Research Ltd (Partner)
  - Algonet SA (Partner)
- Academics:
  - University College London (Partner)
  - University of Surrey (Partner)

Overview of MESCAL approach

- No central point of control/decision making (no Internet God!)
  - interactions between interconnected peers only
  - cascaded hop by hop model
- Focus is on IP connectivity provider interactions at both *service layer* (pSLSs) and *network layer* (qBGP)
- Internal intra-domain means to achieve QoS (e.g. MPLS-TE, IP-TE, over-provisioning) do not impact on the interactions between providers
- 3 service options within the overall MESCAL solution
  - mass market vs VPN
  - loose vs hard QoS guarantees

MESCAL Focus from Business Perspectives

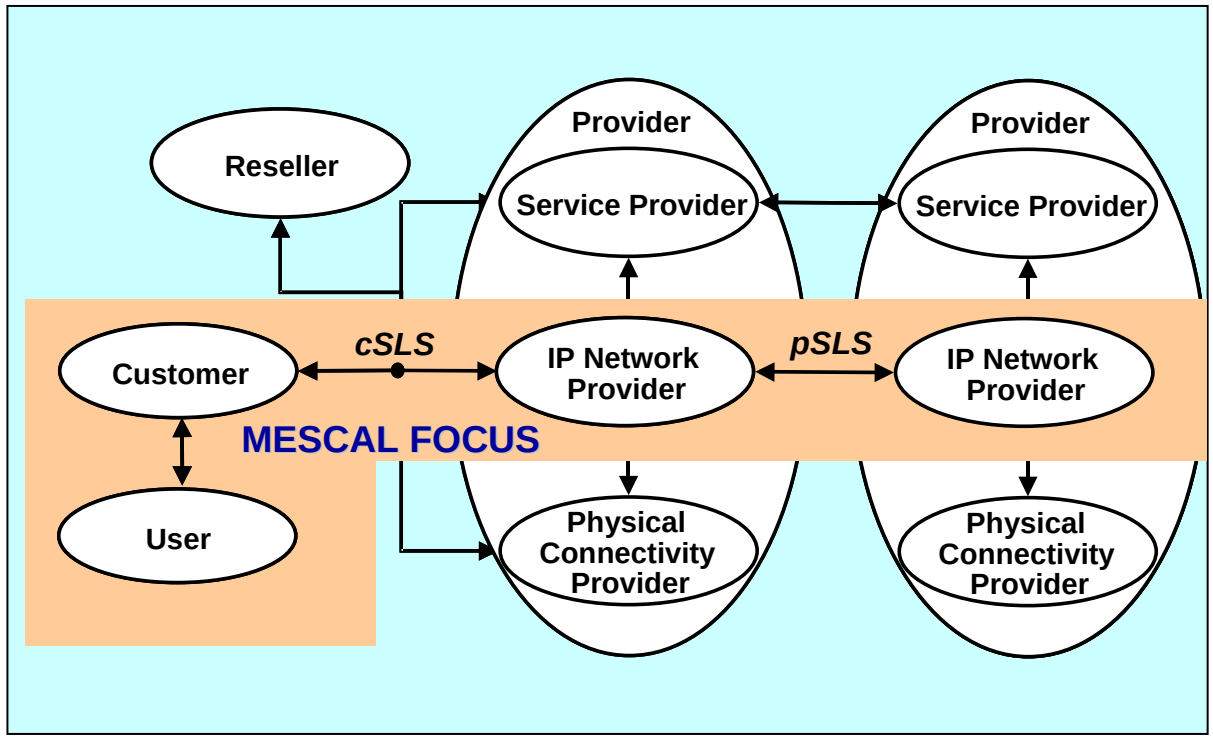

- MESCAL defines two types of service contracts:
  - **cSLS between customers and providers**
  - **pSLS between providers**

QoS Classes

- QoS-class (QC)
  - a basic *QoS transport capability* of a provider domain
  - *performance* attributes-value pairs: ordered set {delay, loss, jitter}
  - analogous to the IETF notion of Per-Domain Behaviour (PDB)

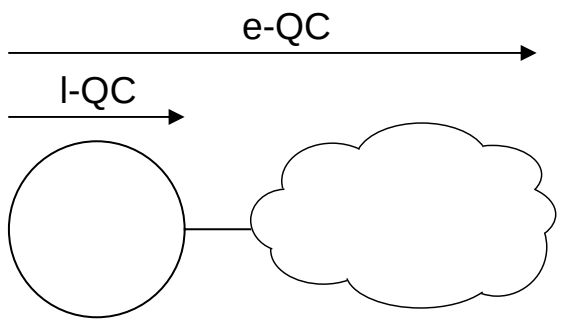

- local-QoS-class (l-QC)
  - a QC with the scope of a single provider
- extended-QoS-class (e-QC)
  - a QC which extends across the boundaries of multiple providers
- meta-QoS-class (meta-QC)
  - an abstract well-known QoS-class with standardised parameter values

Inter-domain QoS example

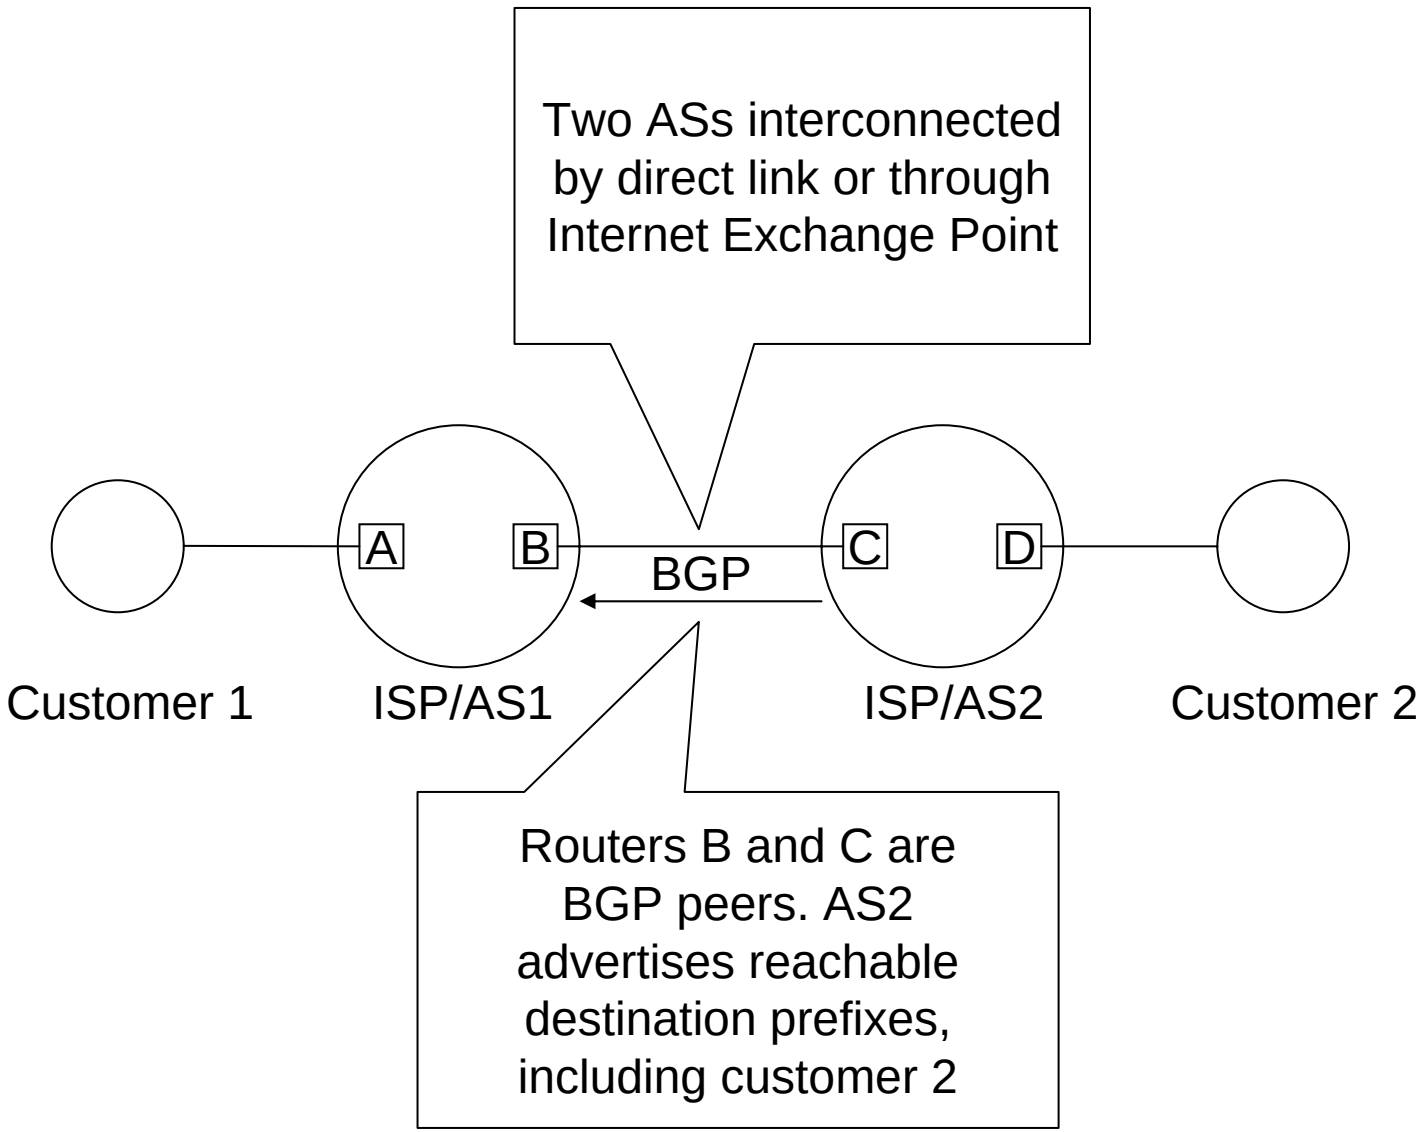

Inter-domain QoS example

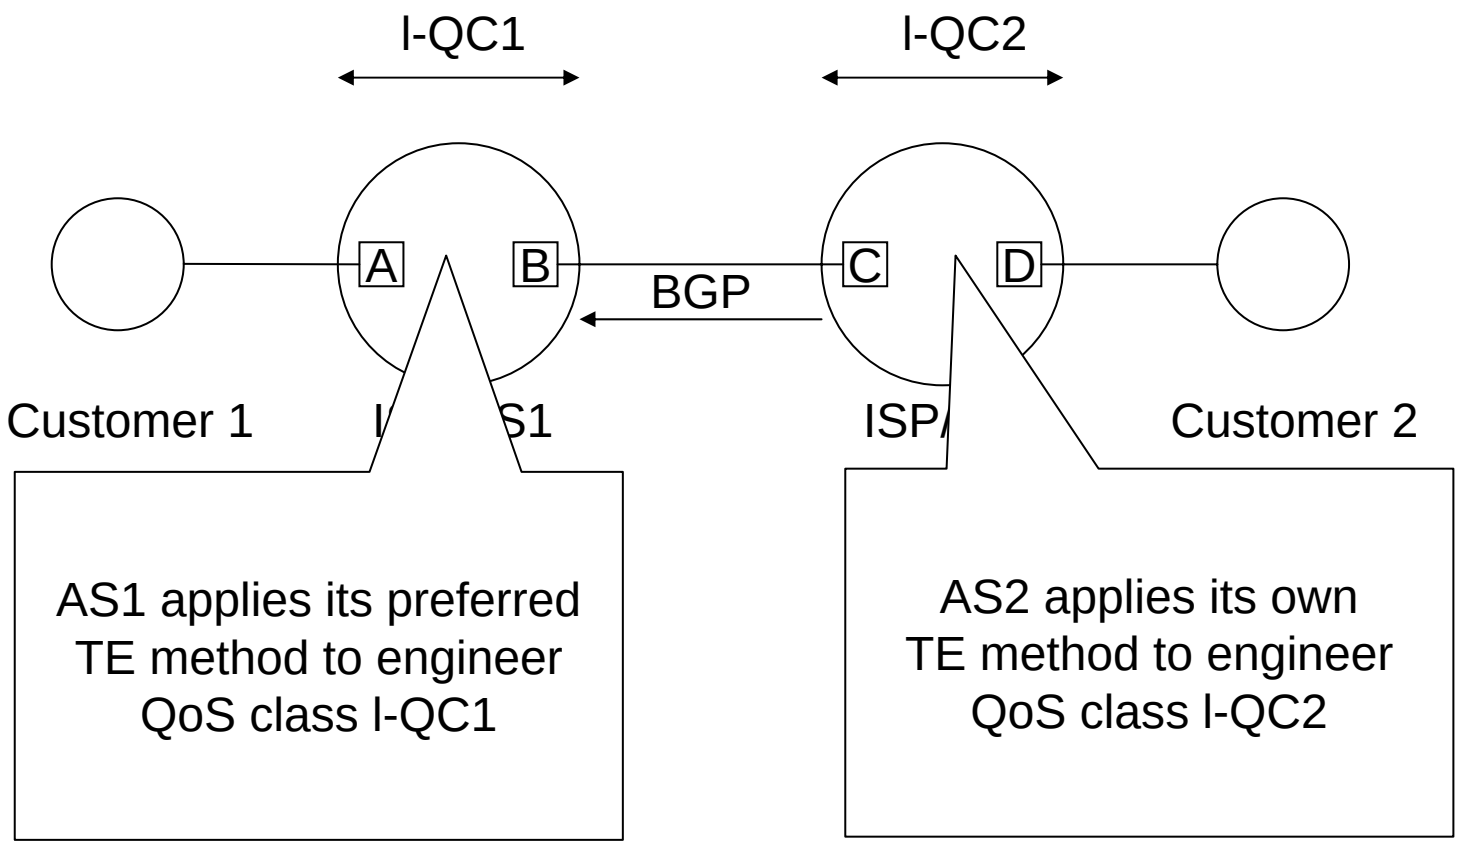

Inter-domain QoS example

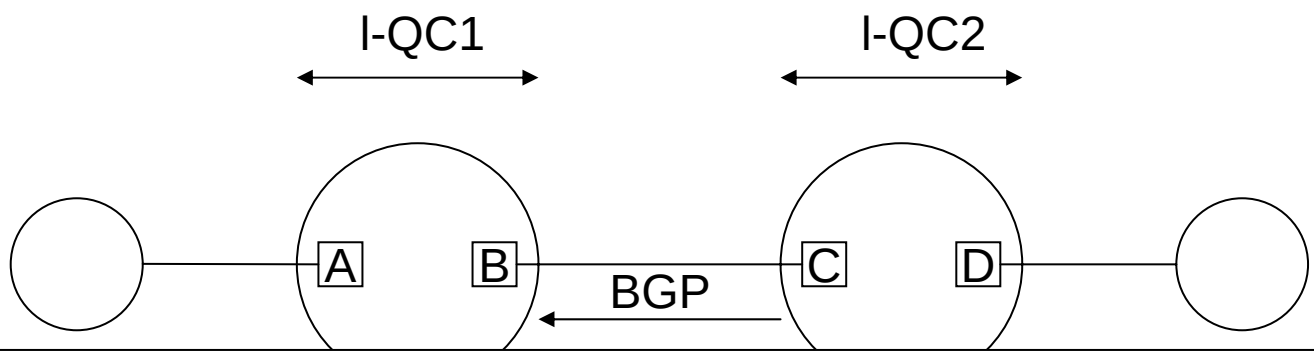

SLAs are required between ISPs and customers (or peer ISPs)
to use other than Best Effort QoS Classes: I-QC1 or I-QC2
- quantity of traffic
- topological scope
- performance parameters

Inter-domain QoS example

ISP1 is aware of ISP2's l-QC2 capability through, e.g. InterQoS marketplace.

According to its business objectives, customer requirements, ISP1 defines an *Inter-domain* QoS Class, e-QC1:

e-QC1 = l-QC1 *op* l-QC2

(*op*: e.g. *addition* for delay, *minimum* for throughput)

We call this a "QoS binding"

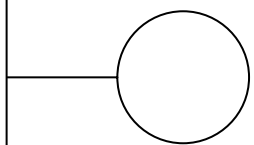

Customer 2

Inter-domain QoS example

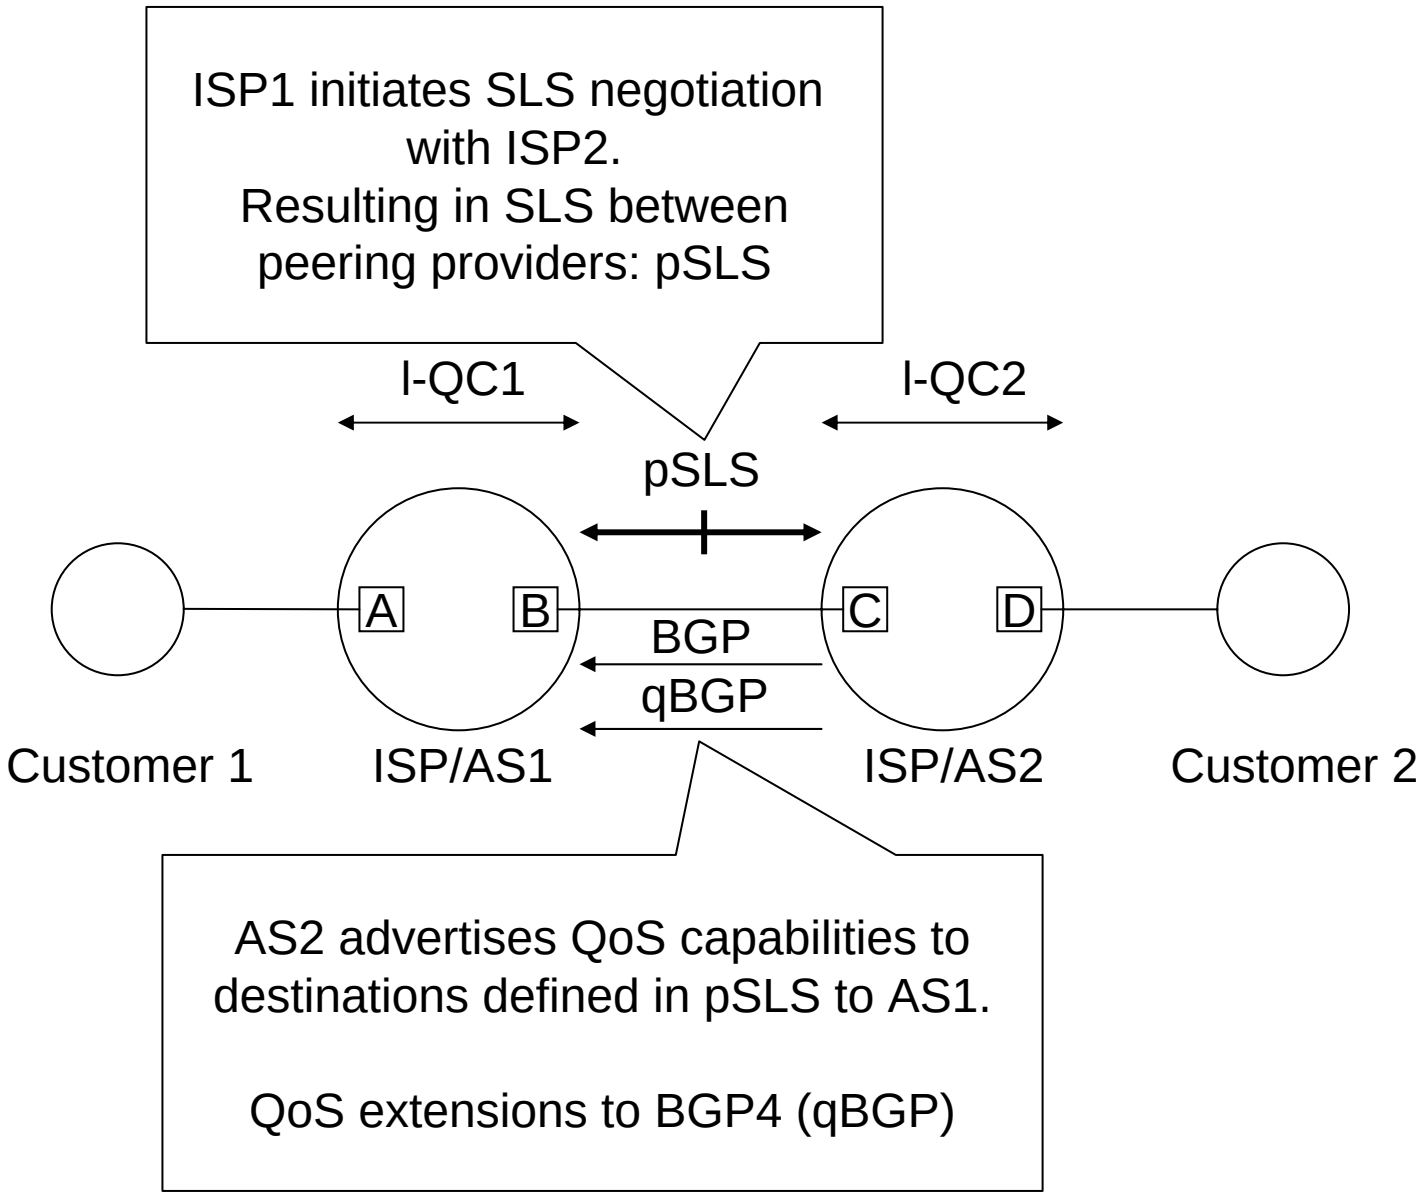

Inter-domain QoS example

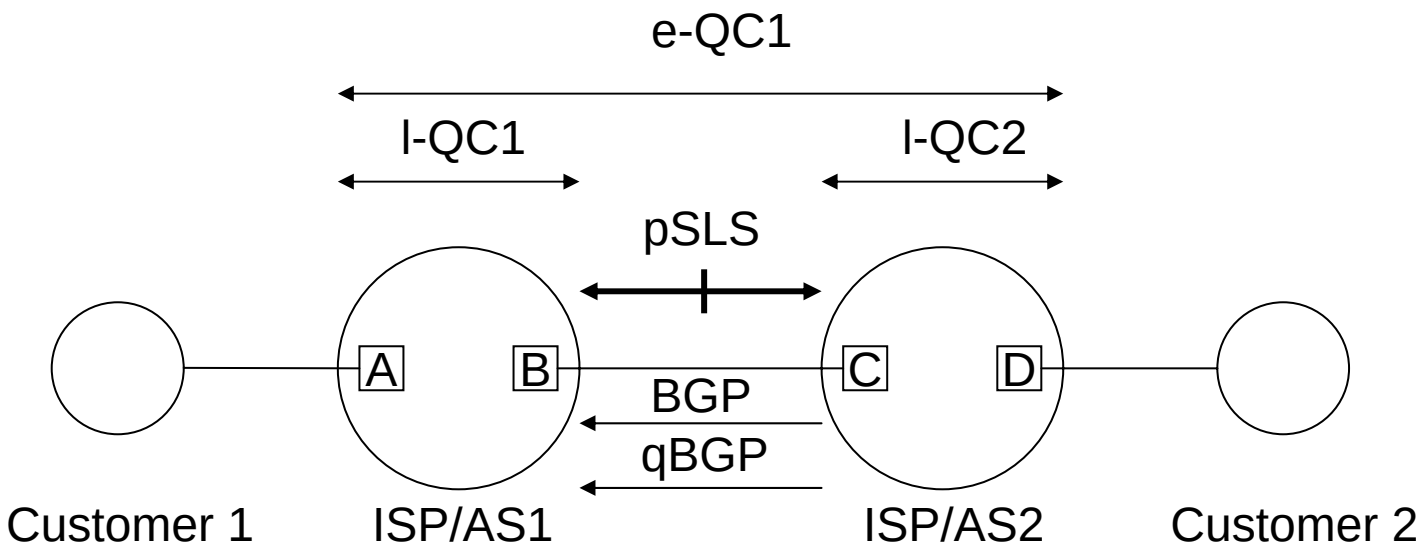

ISP1 is now in a position to offer inter-domain QoS Class e-QC1 to its customers in addition to intra-domain QoS Class l-QC1 and BE services

Cascaded and Source-based Approaches

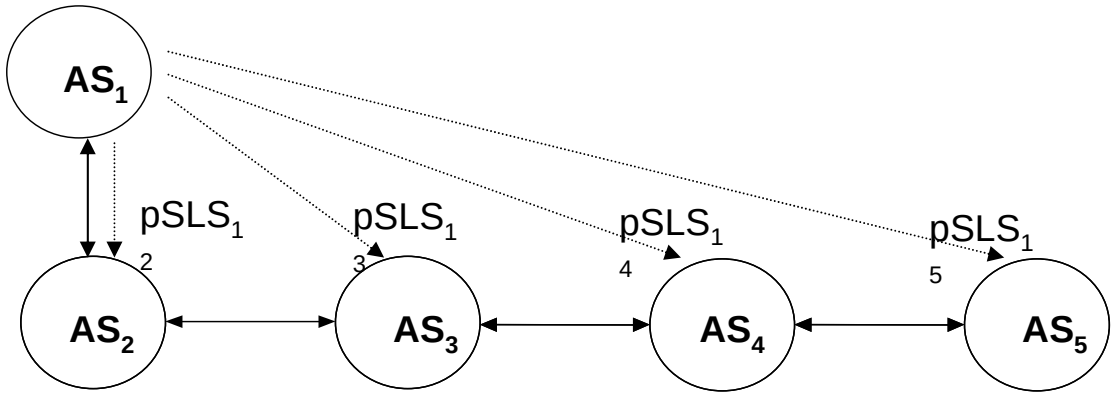

Source-based Approach

Originating ISP forms agreements with all ASs on the end-to-end path

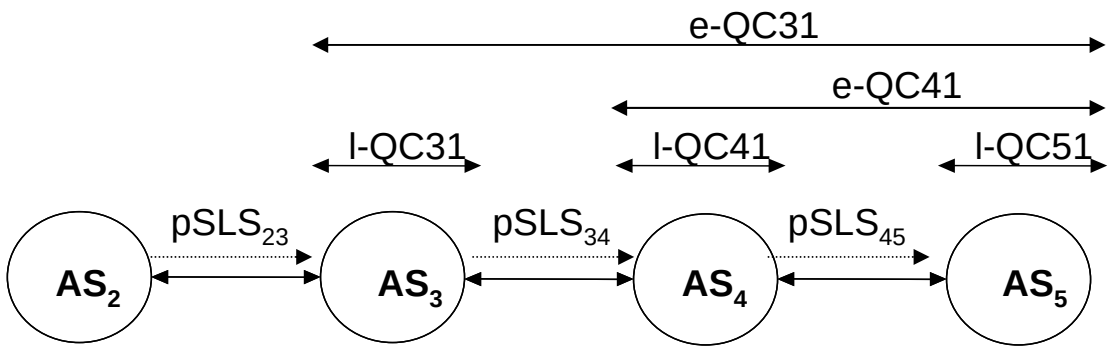

Cascaded Approach

ISPs form agreements only with immediate downstream peers.
A more scalable solution at the cost of some control of the path.

Solution Targets

customer types

residential → corporate

service guarantees as per cSLSs

service options	Loose	Statistical	Hard
end-to-end performance (delay, loss, litter)	Qualitative	Quantitative/ Qualitative	Quantitative
bandwidth guarantees	NO	YES (not per flow)	YES (per flow)
topological scope (reachable destinations)	NO	Any/Specific	Specific

Inter-domain QoS Solution Options

- Loose guarantees (solution option 1)
  - well known set of Meta-QoS-Classes
  - qualitative QoS guarantees/relative treatment (e.g. low-delay, low-loss for TCP traffic)
  - QoS bindings only take place within the same Meta-QoS-Class
  - pSLSs define aggregate bandwidth but do not restrict destinations (reachability is determined dynamically through qBGP)

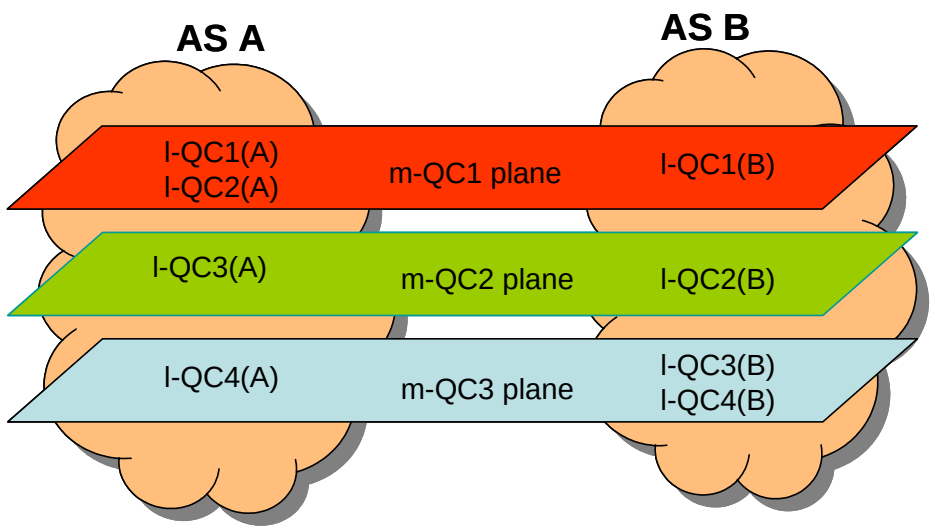

Inter-domain QoS Solution Options

- Statistical quantitative guarantees (solution option 2)
  - ISP offers l-QCs with well defined quality parameters (e.g. delay <= 100ms) to specific destinations
  - QoS bindings are not restricted: any l-QC may be bound to any l-QC (or e-QC) offered by downstream ISPs
  - pSLSs define aggregate bandwidth and specific set of destination prefixes (for the selected QoS binding)

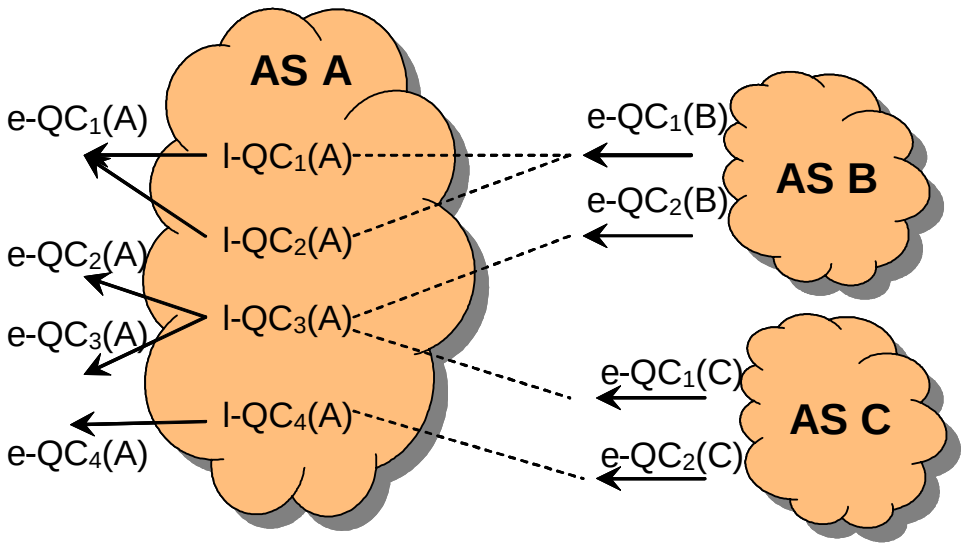

Inter-domain QoS Solution Options

- Hard guarantees (solution option 3)
  - built on solution option 1
  - explicit inter-domain MPLS LSP-TE tunnels define end-to-end path and resource reservation
  - Path Computation Protocol (PCP) runs between Path Computation Elements (PCEs)
  - path selection based on inter-domain reachability and QoS aggregate information learned via qBGP

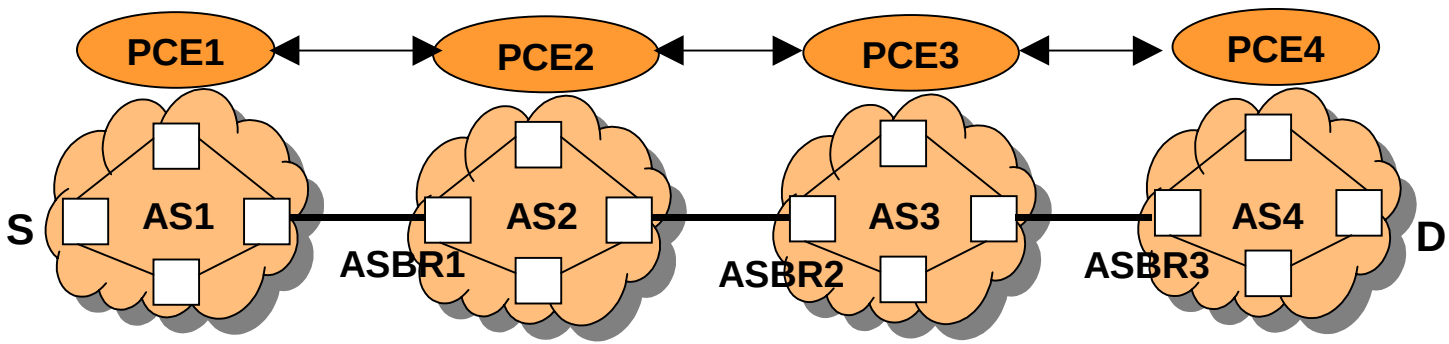

MESCAL functional architecture

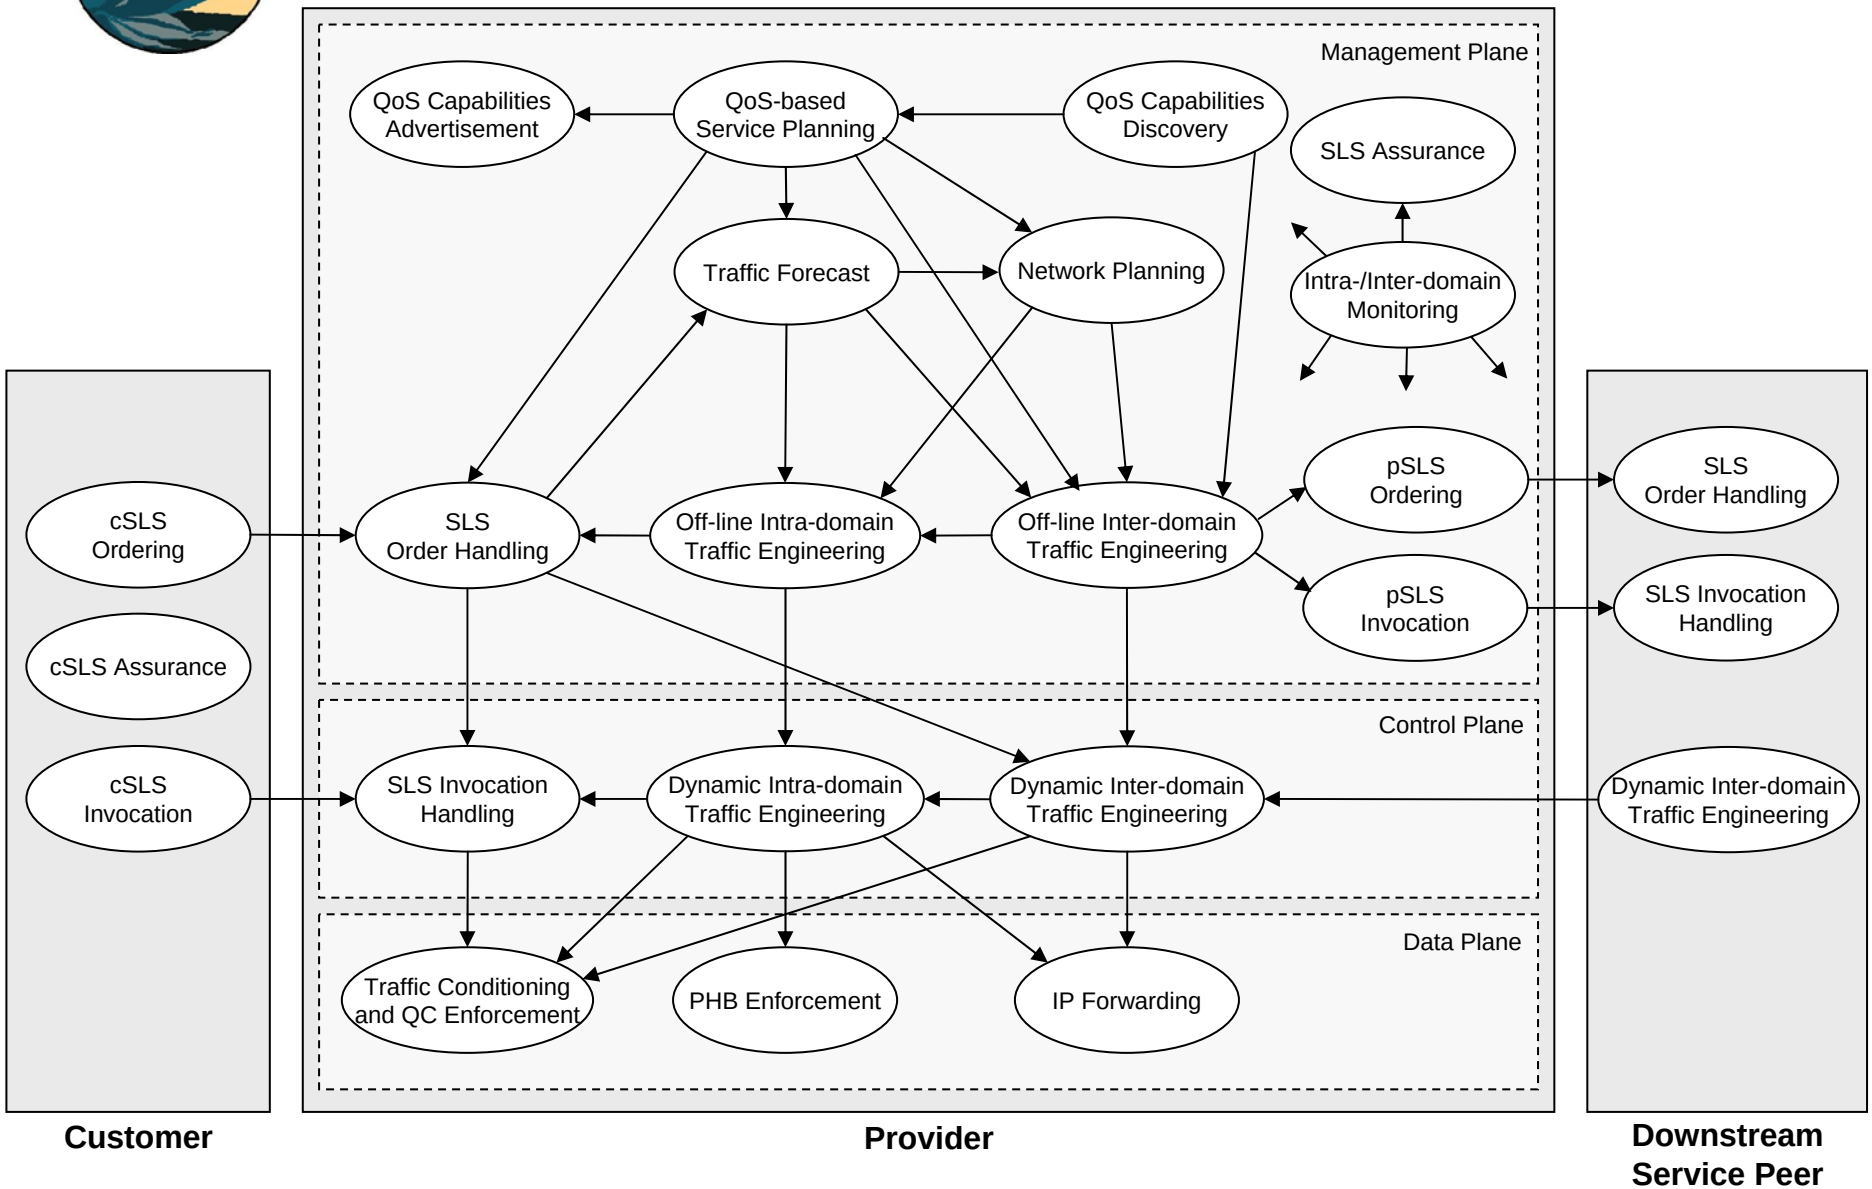

MESCAL key results

- Inter-domain QoS framework specifications
  - service model, functional architecture, solution options
- qBGP specifications
  - QoS enhancements to BGP protocol
  - QoS-based route selection process
- pSLS modelling, negotiation and provisioning
- Inter-domain, QoS-focussed TE algorithms
  - off-line and dynamic (from single AS's perspective)
  - for both uni- and multicast
- Intra-domain IP-based TE algorithms for QoS
- Admission control at subscription and invocation epochs
- Testbed prototypes and simulation models

http://www.mescal.org/

