Peering in an IP world
Technology Requirements

OPASTCO 2009 Technical & Marketing Symposium

steve ulrich - consulting systems engineer

sulrich@cisco.com

Internet structure

traditional assumption is that the Internet is based on a well ordered provider-client hierarchy

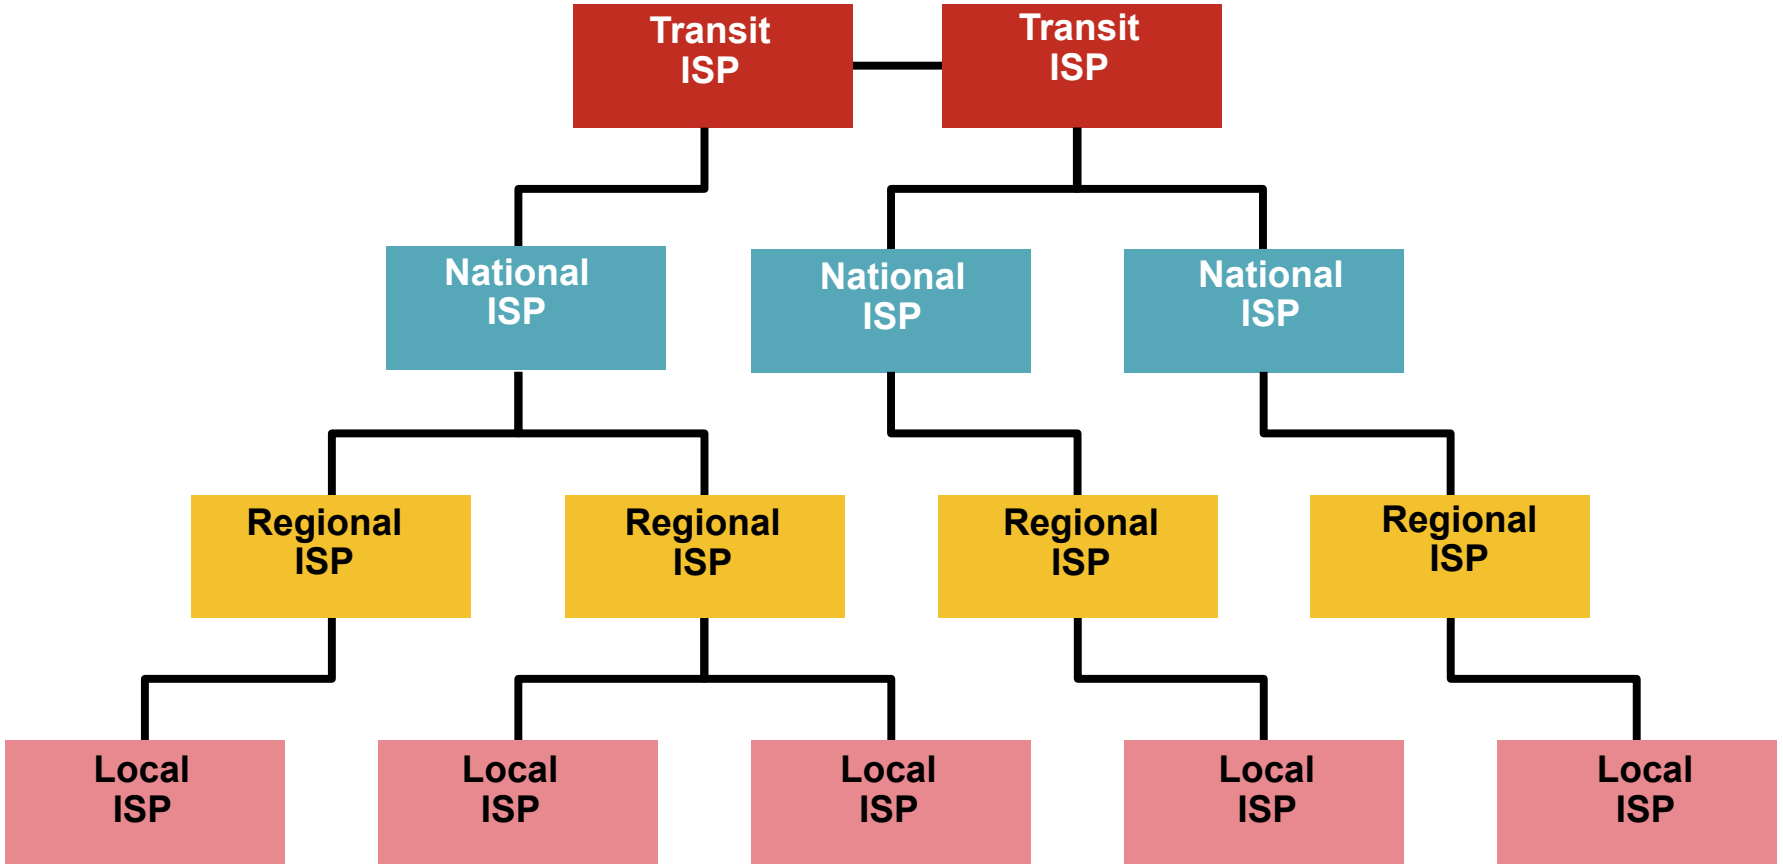

Internet structure

reality is not so ideal

- the Internet is an interconnection of ~ 30,000 (semi-) autonomous service providers
- there is no central coordination for the management of interconnections, services, and tariffs
- Internet peering ecosystem includes

  many policies / many services / one Internet

- unordered subset of interconnects
- driven by business requirements underpinned by performance
- non-disclosure and bi-lateral agreements
- peering is now considered a corporate asset and legal concern

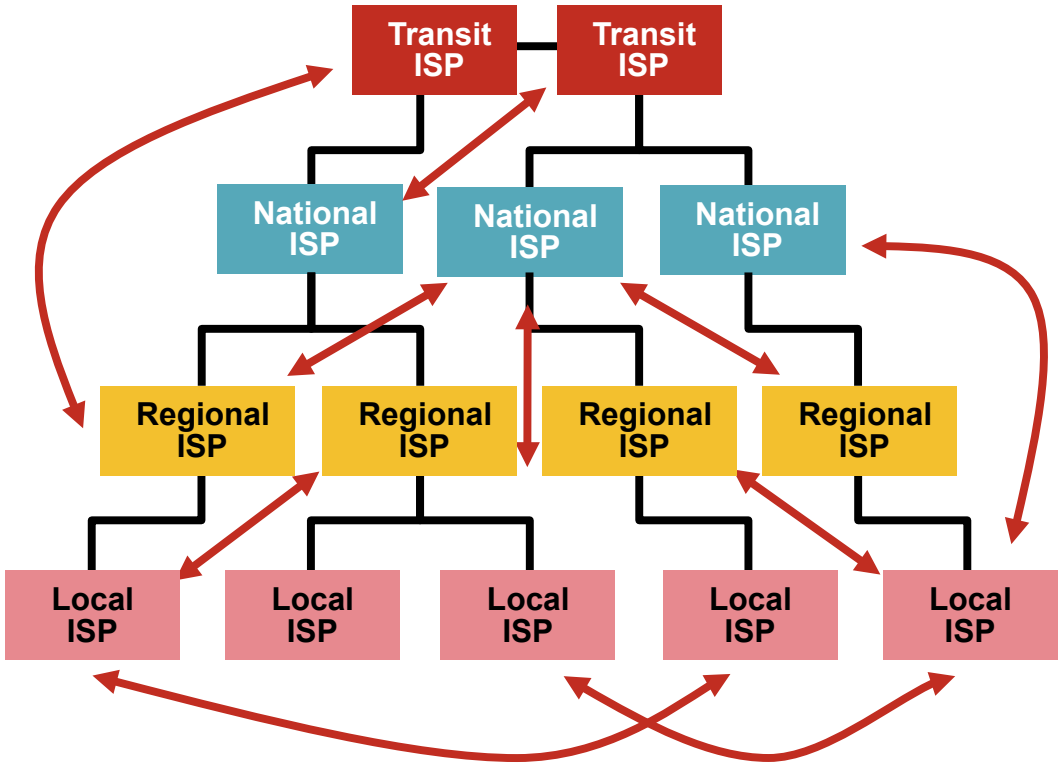

transit - definition

- transit is the business relationship where one ISP provides reach-ability to all destinations in its routing table to its customers
- transit provides connectivity to a superset of all destinations

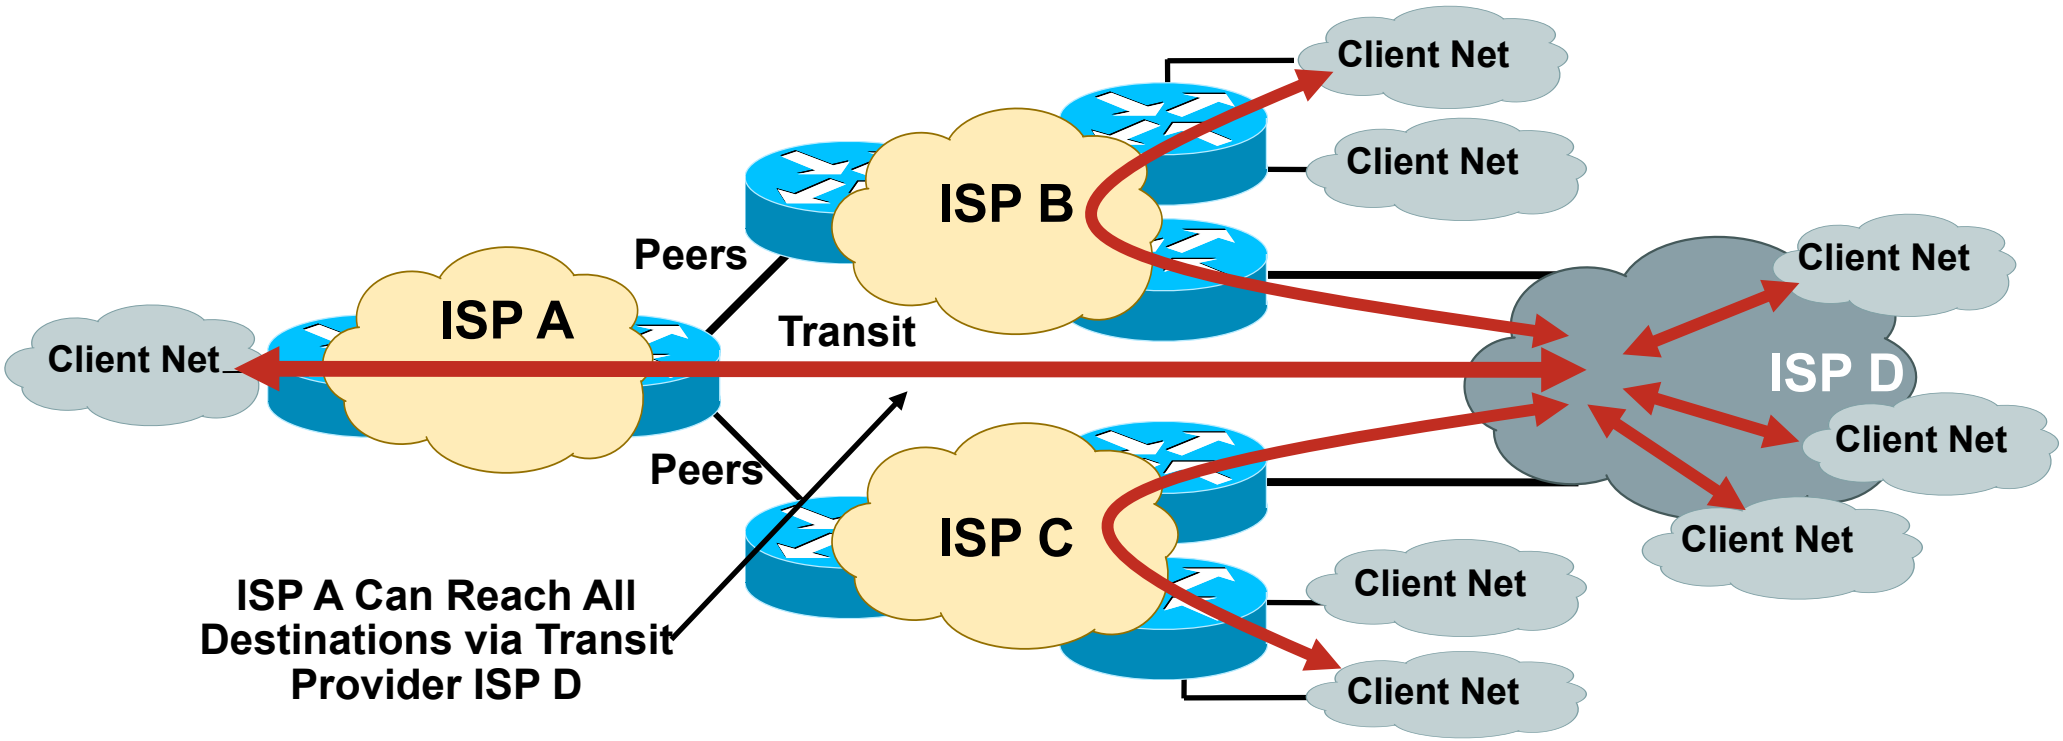

peering - definition

- peering is the business relationship where ISPs provide to each other reach-ability to each predefined portions of their routing table
- peering provides connectivity to a subset of a provider's customer destinations

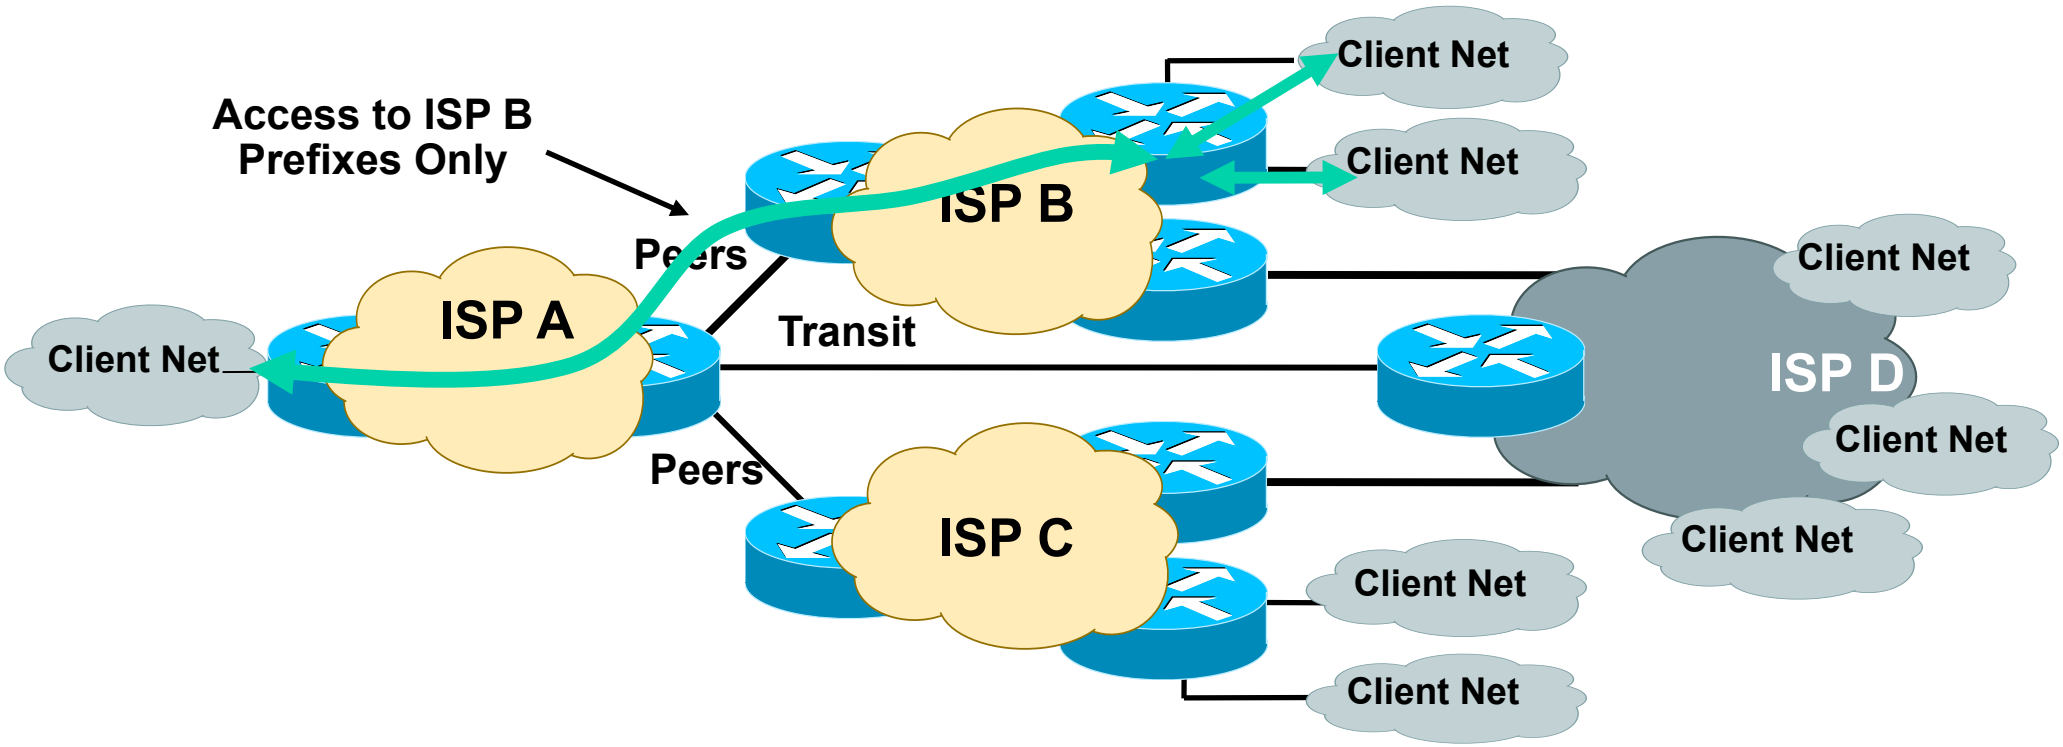

traditional Internet peering model

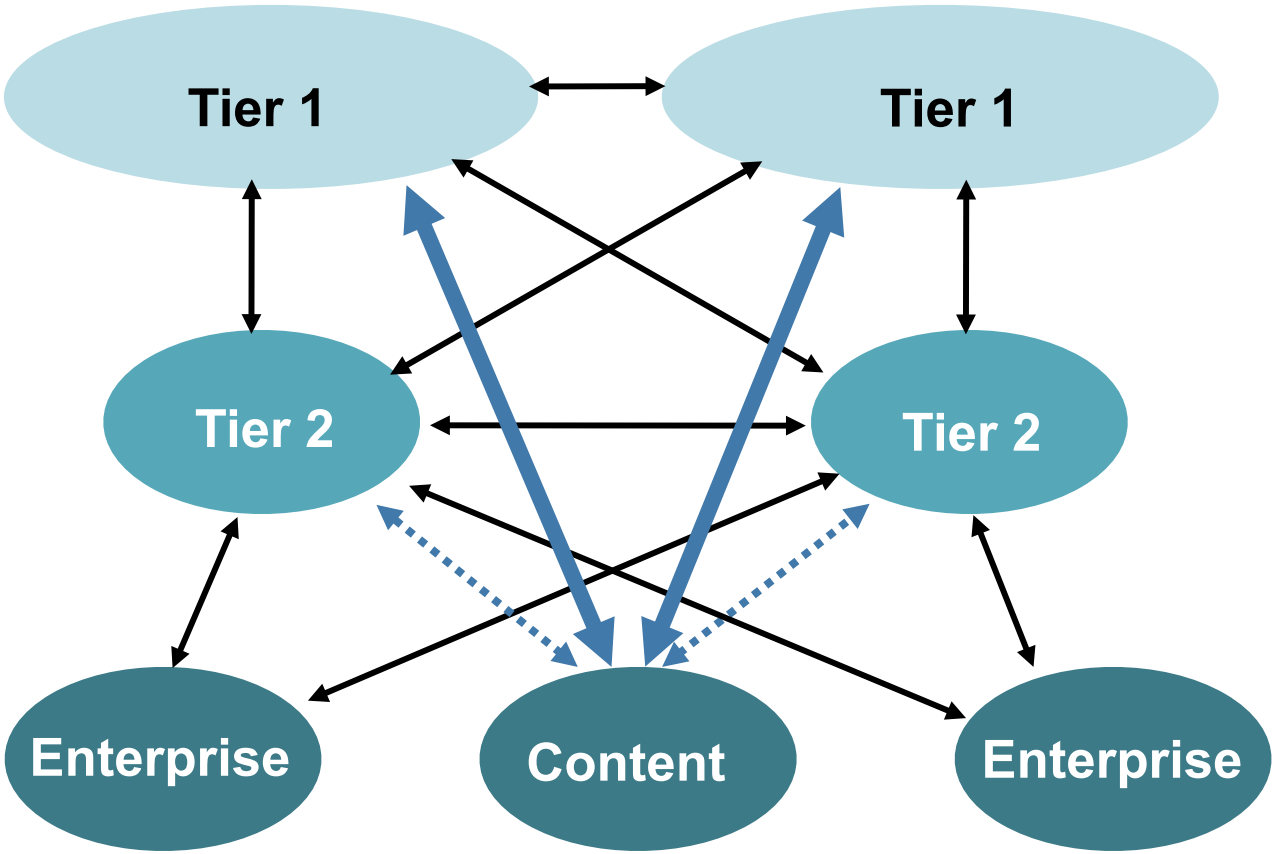

- tier 1 providers have access to the entire Internet (region) routing table solely through peering relationships
- tier 2 providers must buy some transit from tier 1 providers
- content providers buy transit (primarily from tier 1) to provide content

Internet peering evolution

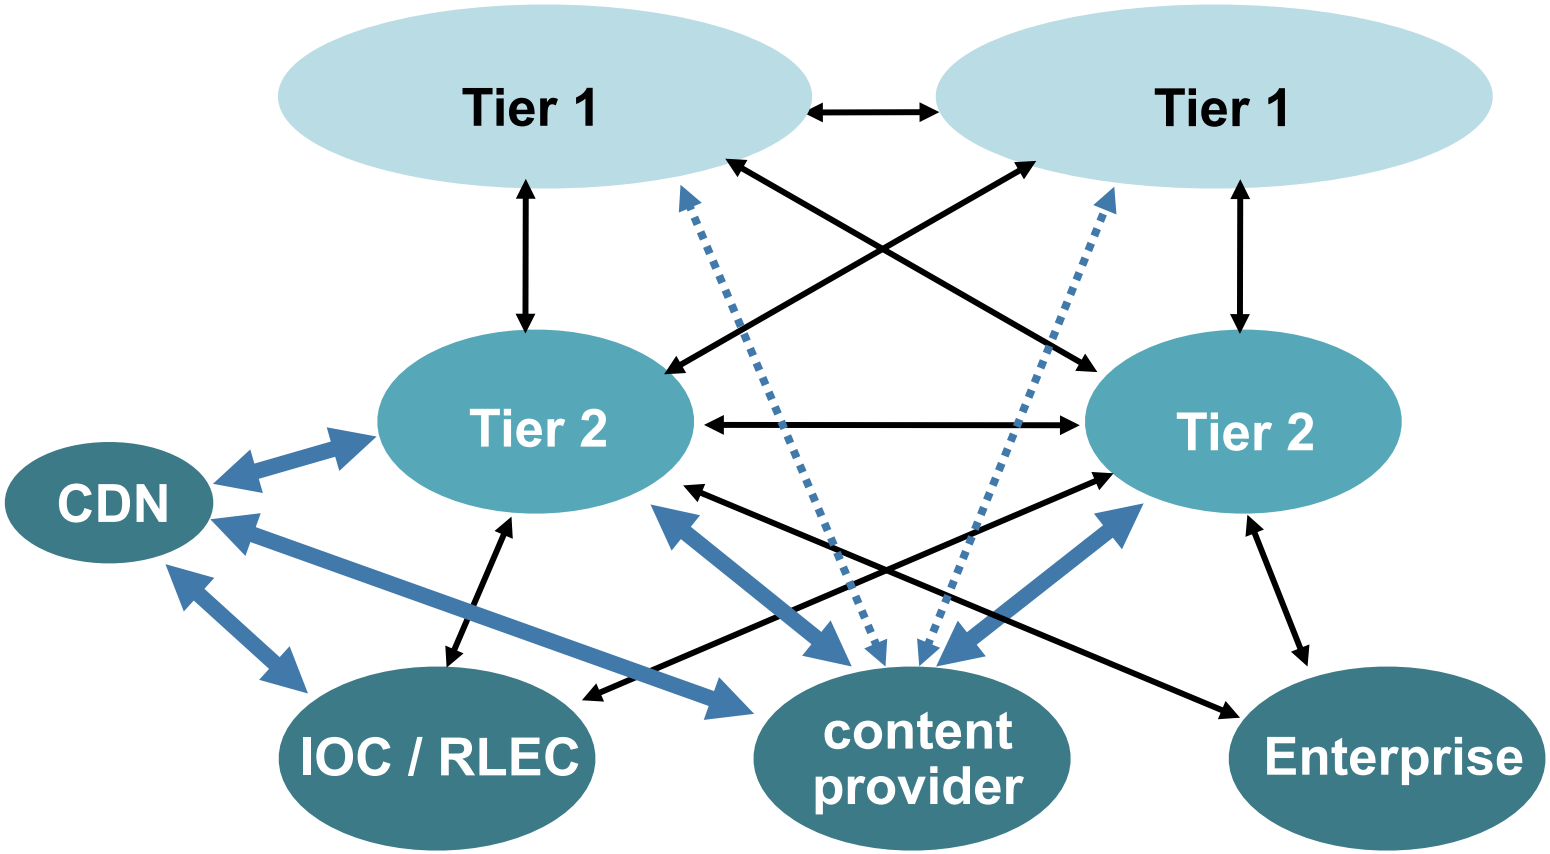

- tier 1 providers have access to the entire Internet (region) routing table solely through peering relationships
- tier 2 providers must buy some transit from tier 1 providers
- content providers peer (increasingly with tier 2) providing content directly onto the broadband networks

peering rationale

for the ISP

- commonly estimated, 10 - 20% of traffic can be peered away
- even under congestion, capacity can be upgraded and managed more effectively

for the content providers

- improve application performance, reduction in latency
- improvement in throughput
- CDNs as content providers ...
  - peering at NAPs or with ISPs improves burstability
  - backup for on-net servers
  - marketing - CDNs tout the number of interconnections they have to their customers

common to both

- reduce transit ISP service costs
- upgrades require less planning and costs
- greater control over routing and traffic load balancing

Internet peering interconnection

public / shared peering

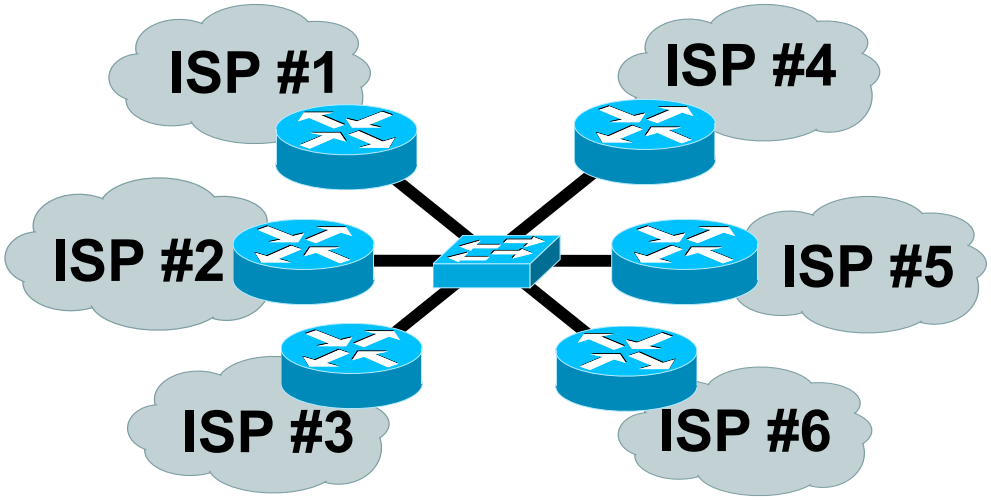

switched ethernet

private peering

ethernet / POS

- peering between equivalent sizes of service providers (e.g. tier 2 to tier 2)
  - shared cost private interconnection, equal traffic flows
  - “no cost peering”
- peering across exchange points
  - if convenient, of mutual benefit, technically feasible
- fee based peering
  - unequal traffic flows, “market position”

peering requirements / costs

- if you're not in an Internet exchange (IX) location already

  IX Colo / Power / transport to IX

- IX Port and/or cross-connect fees
- CapEx: routers, switches, optics, ports

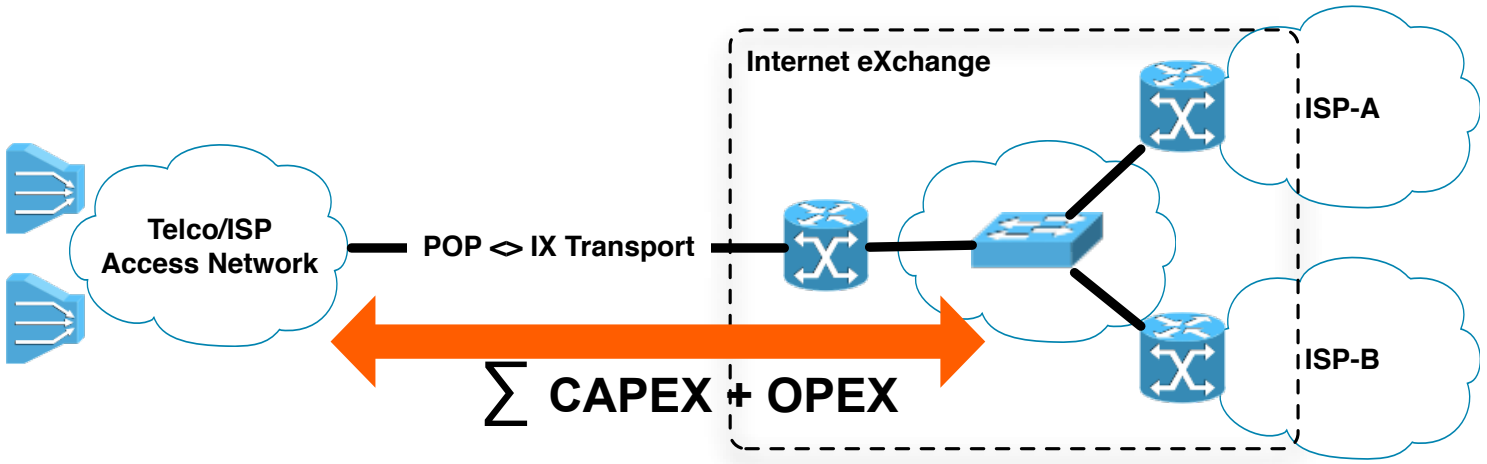

- OpEx: Network Engineers

peering technical requirements

- platform capacity

  increasingly oriented around 1GE and 10GE interfaces

- instrumentation

  what traffic is traversing your network and who's sending it to you? peering requires accounting on peering interconnect traffic based on its source, destination and their traversed AS path, grouping or service category

  this requires high performance and scalable NetFlow

- routing policy expression / peering policy enforcement

  announce what you need, to who you need to and apply policy to what you receive from your peers

  utilize hierarchical QoS policies on ingress and egress, to enforce peering policy. couple this to peering interfaces with QPPB

- advanced and high performance security measures

  ACL application at line rate - drop spoofed traffic and mitigate DDoS attacks

  automated and self-activated control-plane protection mechanisms

platform capacity

- highly Ethernet oriented connectivity
- GE increasingly minimum peering interface
- sub-rate 10GE peering quite common

  w/sub-rate provided via policing and/or QoS policy on a per-VLAN basis

- 10GE common on private connections or in peering fabrics
- requires line-rate application of various features

  ACLs - auto-generated ACLs of very large size (1000s of lines)common - require hardware based application

  QoS application must take place in hardware

  hardware based control-plane protection mechanisms

instrumentation - NetFlow

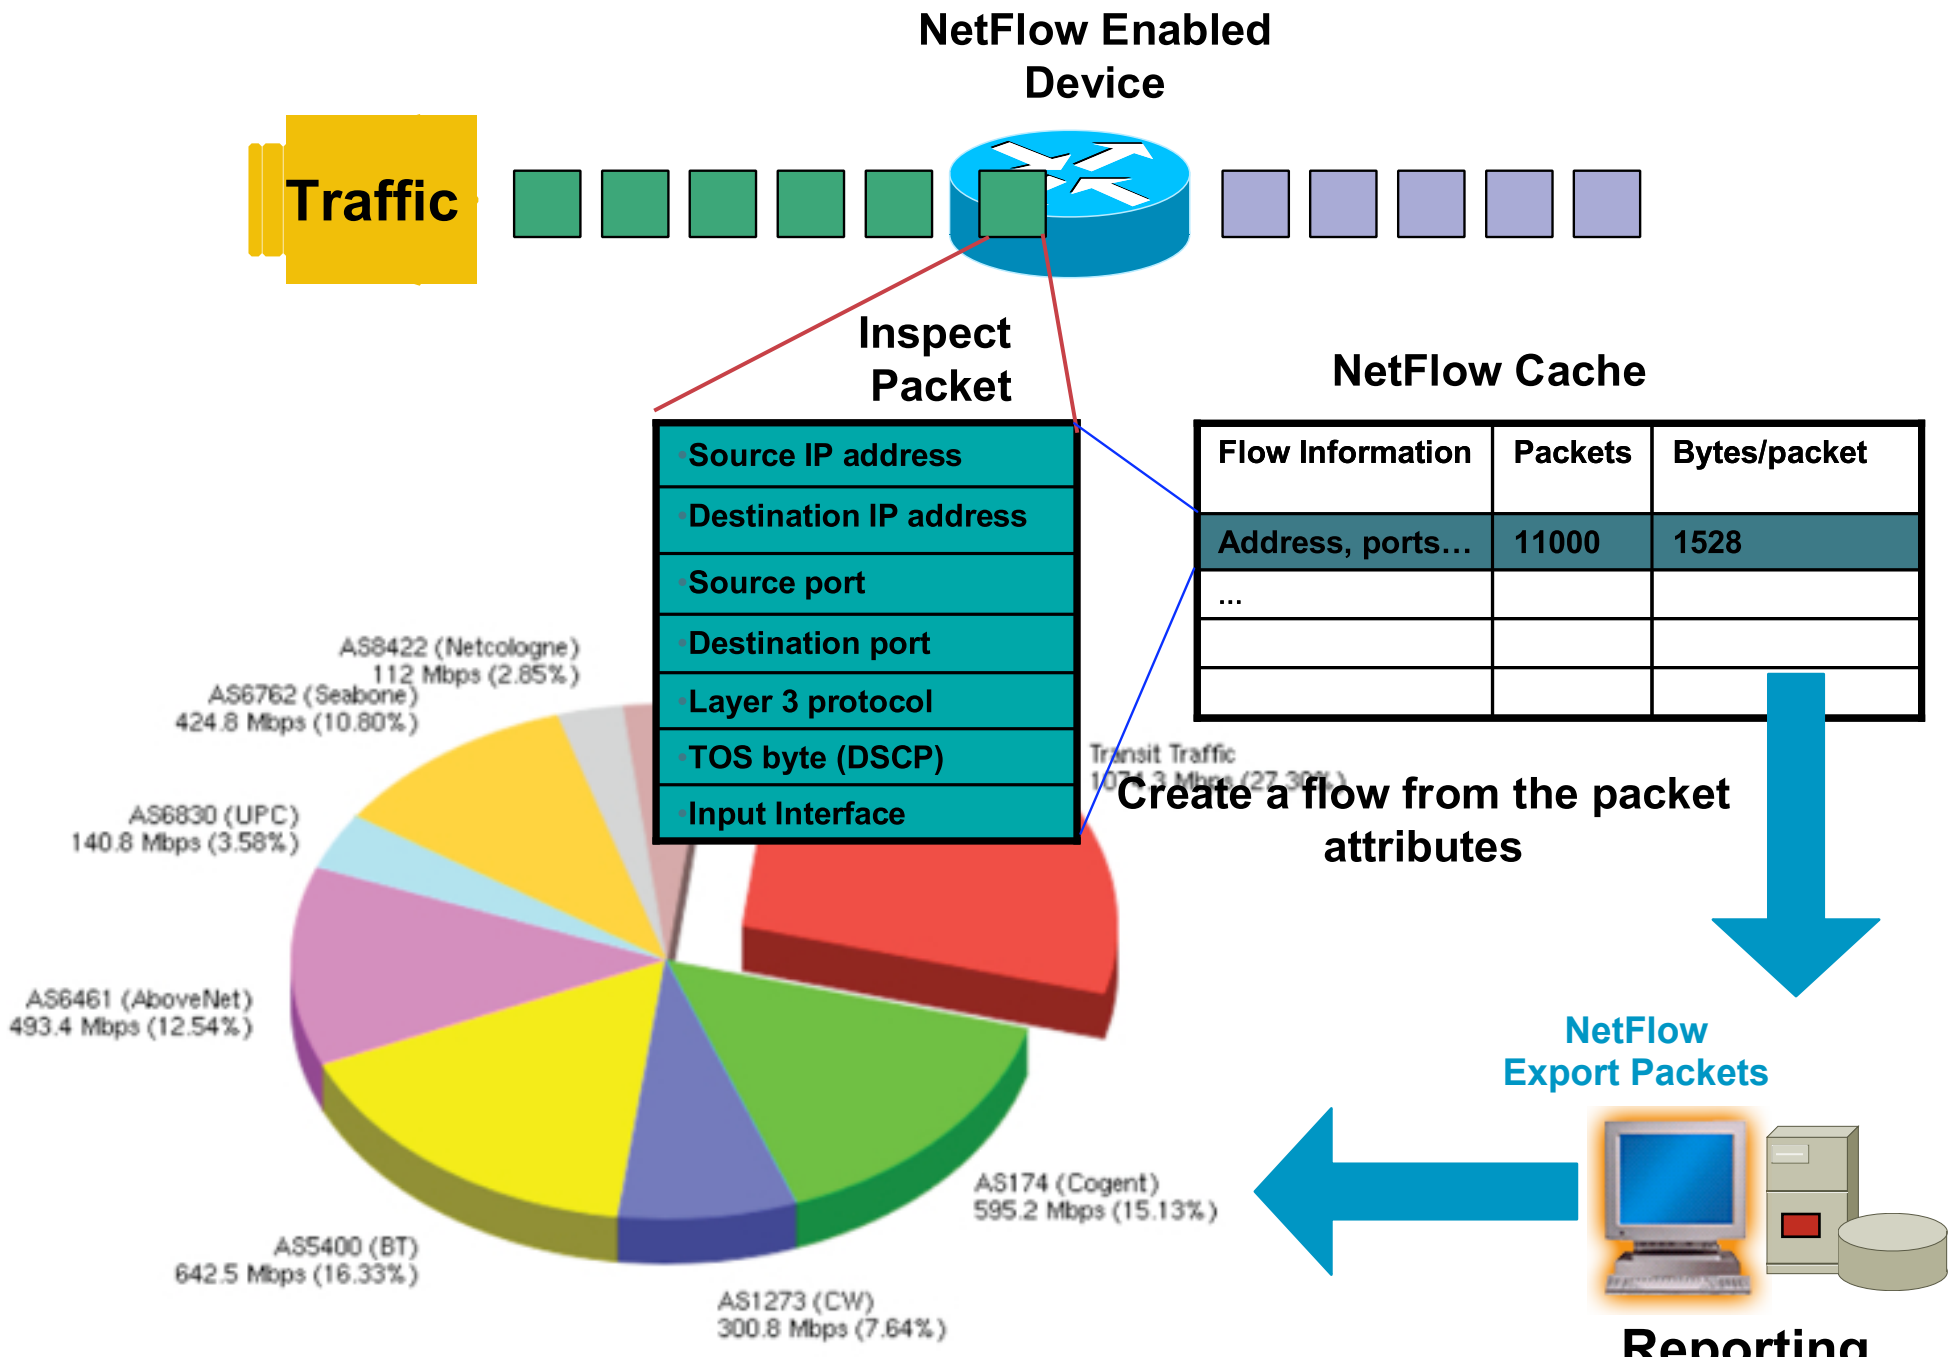

instrumentation - NetFlow - available info

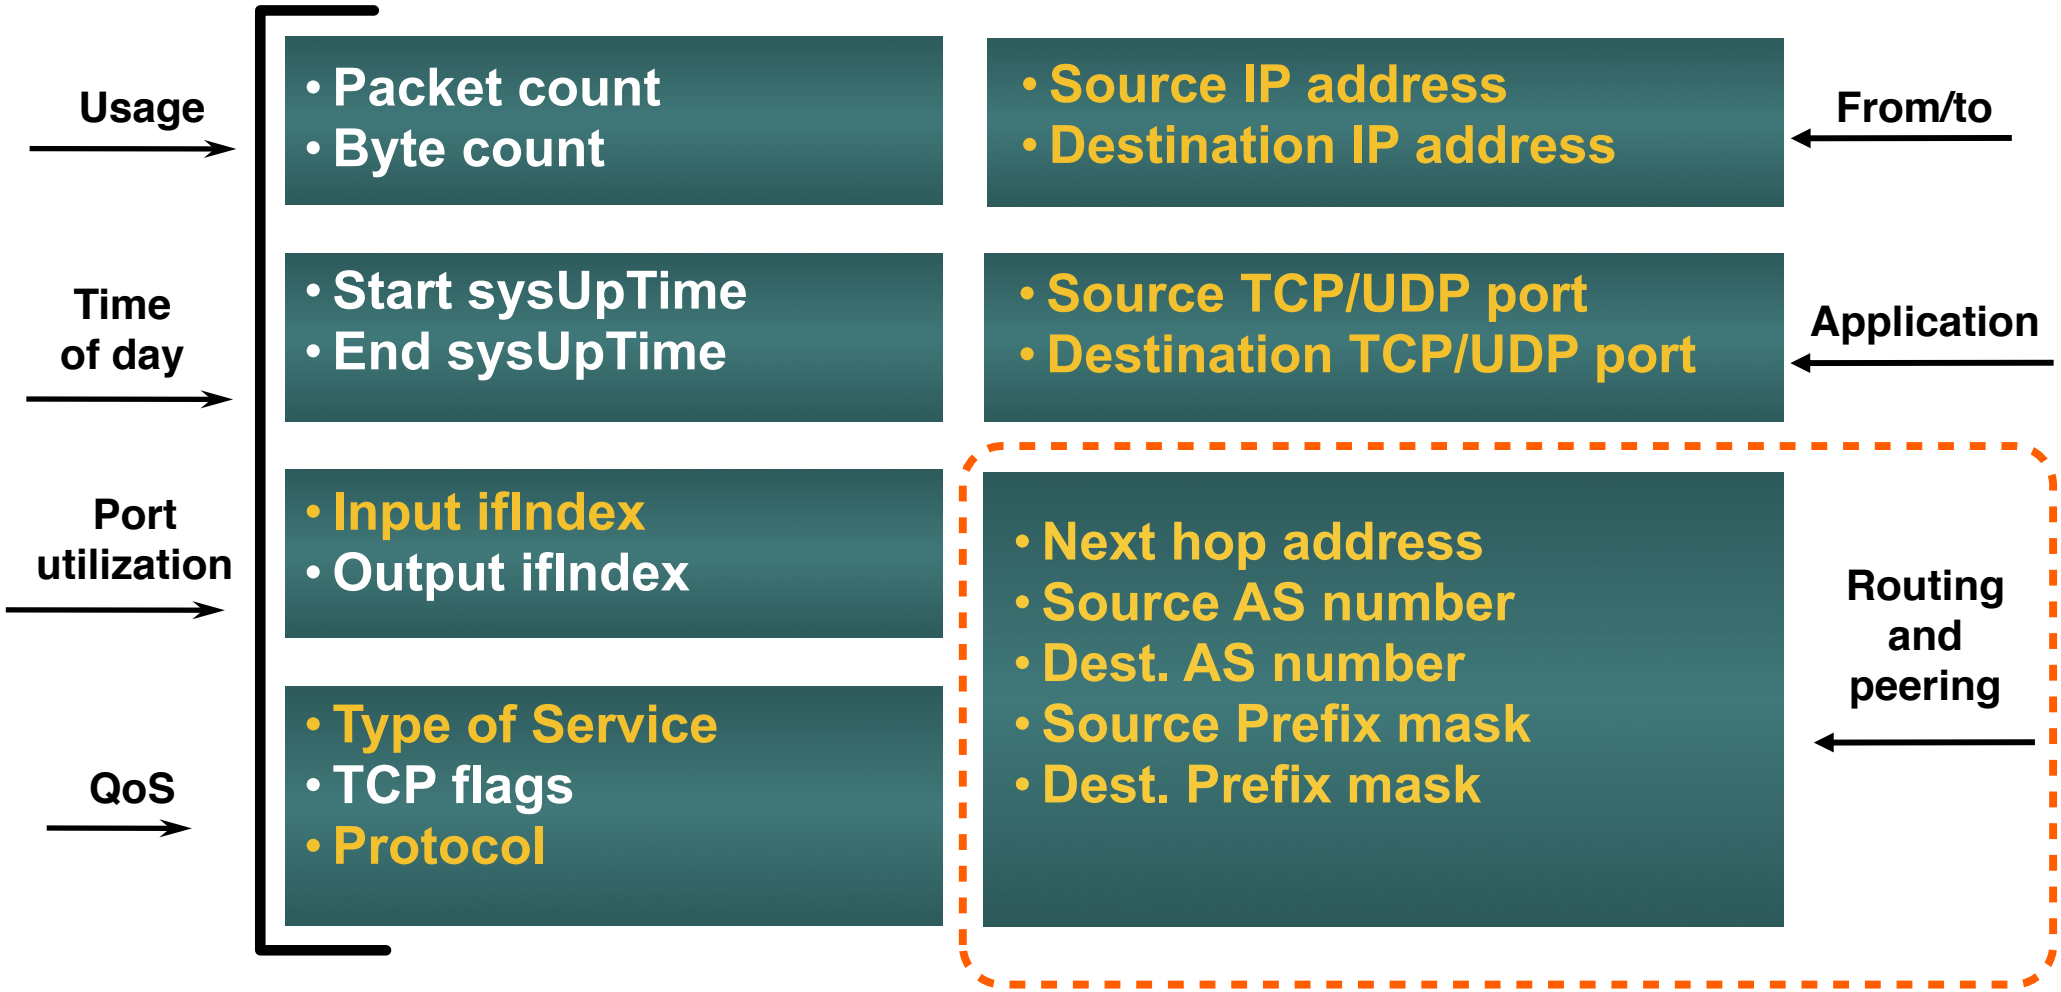

v5 used extensively today - v9 provides interesting and notable enhancements

instrumentation - NetFlow - reporting tools

Product Name	Primary Use	Comment	OS
Cflowd	Traffic Analysis	No longer supported	UNIX
Flow-tools	Collector Device	Scalable	UNIX
Flowd	Collector Device	Support V9	BSD, Linux
FlowScan	Reporting for Flow-Tools		UNIX
IPFlow	Traffic Analysis	Support V9, IPv4, IPv6, MPLS, SCTP, etc..	Linux, FreeBSD, Solaris
NetFlow Guide	Reporting Tools		BSD, Linux
NetFlow Monitor	Traffic Analysis	Supports V9	UNIX
Netmet	Collector Device	V5, support v9	Linux
NTOP	Security Monitoring		UNIX
Stager	Reporting for Flow-Tools		UNIX
Nfdump/nfsen	Traffic Analysis	Support V5 and v9	UNIX

note: there are many open source NetFlow reporting tools available

enforcing IP peering destination policy

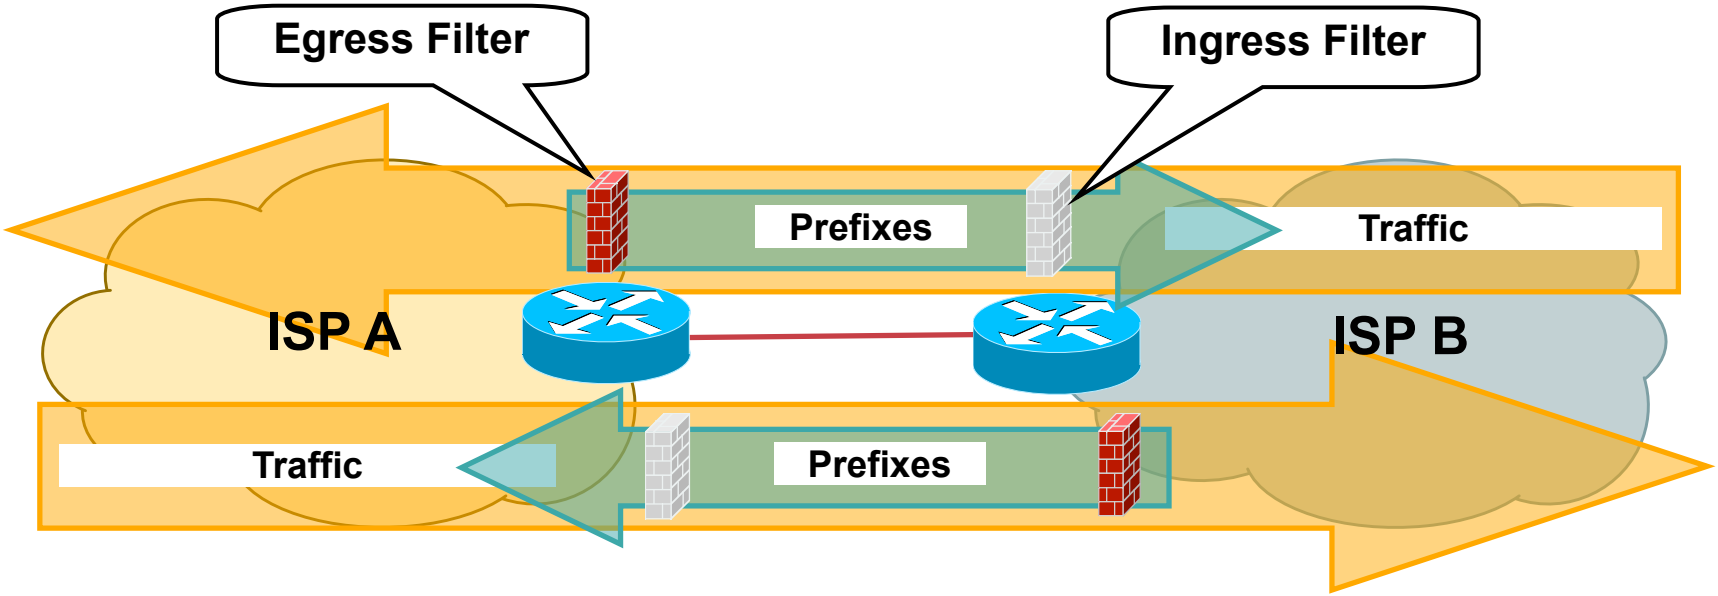

Guarded Trust

ISP A trust ISP B to send X prefixes from the Global Internet Route Table.

ISP B Creates a egress filter to insure only X prefixes are sent to ISP A.

ISP A creates a mirror image ingress filter to insure ISP B only sends X prefixes.

ISP A's ingress filter reinforces ISP B's egress filter.

enforcing IP peering source policy

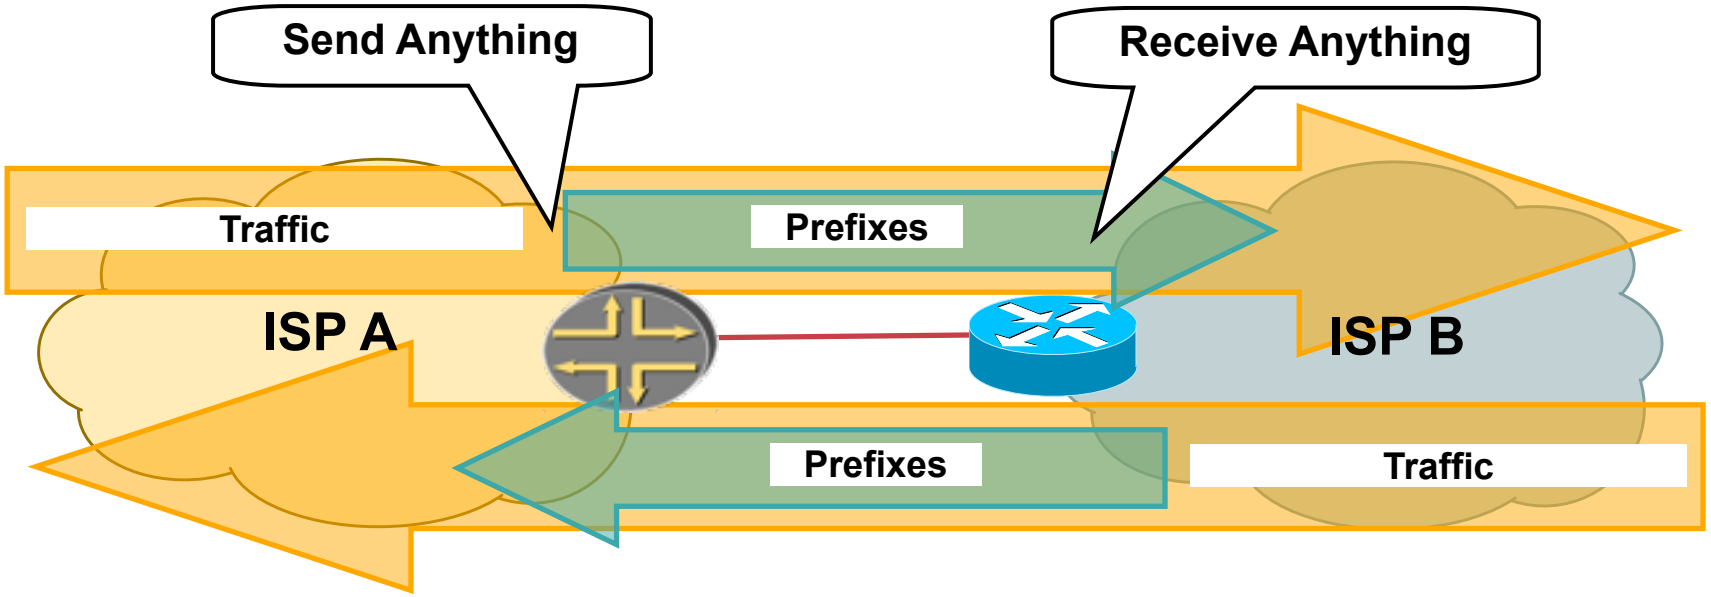

Enforcing Source Policy – requires a sophisticated tool kit

ISP A trust ISP B each other to send packets that match their peering agreement.

reality is that there is nothing to stop the ISPs from sending anything they want. **hence, traffic dumping**

tools like Netflow and BGP Policy Accounting are used identify abuses, data plane enforcement takes place with ACLs

security mechanisms

- layered control plane protection using multiple policers
  - L2 congestion control
  - line rate ACL filtering
  - control-plane session validation using pre-filter mechanisms
  - adjustable performance for trusted control plane session treatment
  - multiple queues to CPU

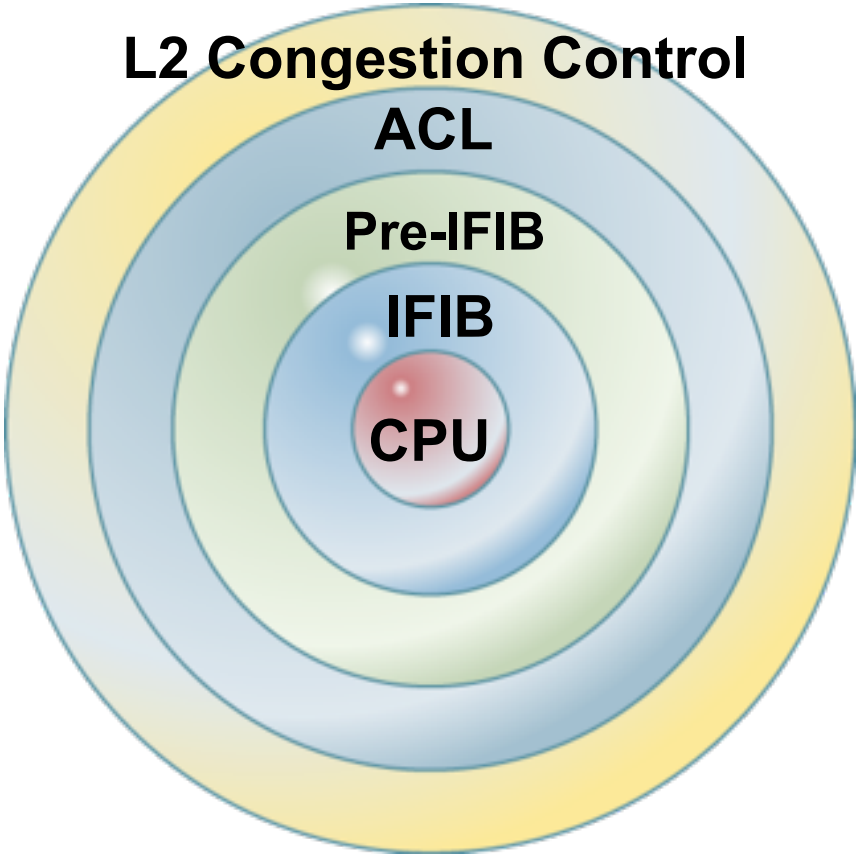

- support MD5 authentication for routing protocols
- defend against TCP finger printing
- priority in switch fabric ensures control traffic will never be dropped
- support GTSM RFC 3682 (formerly BTSH)

peering requirements summary

peering requirement	relevant technology	scale / notes / etc.
peering accounting	MAC and BGP policy accounting	support for line rate accounting
peering billing, flow monitoring	NetFlow	sampled NetFlow (better than 1:1000 sampling granularity)
peering bandwidth guarantee and traffic separation	Hierarchical QoS Per Vlan Policy	requires scalable line rate policy application within hardware
security, DDoS mitigation	uRPF	line-rate application
	Control-Plane Policing	requires an automated and self activated hardware based policer
	In and out ACL	high-ACE count - > 32K line rate application
peering policy enforcement	QPPB source and destination based FIB lookup for H-QoS classification and policing	support for line rate QoS classification application and QPPB binding
peering link fault detection and integrity	BFD detection	requires distributed BFD implementations with sub-second timer granularity
large network resilience techniques	fast convergence	BGP PIC, Fast IGP, IP FRR support for RIB/FIB scale from 1M - 2M routes


CISCO