Service Providers In a Global World

Design principles

Vijay Gill
vijay.gill@gmail.com
SANOG X
New Delhi, India
September 5, 2007

Disclaimer

Legalese: Please note the following

This talk is a summary of common practice in the industry, based on experience in the EU/US. **This is not in any way shape or form based on what anyone may be doing, may plan to do, or may have done in the past**.

This is intended as a discussion on best common practices that I've gathered over years in the industry

Agenda

Introduction

Fundamental Constraints

Differentiators

Design Principles

Conclusions

Slide 1

Hire RFP Engineers

“Oi! Vendor, run me up a backbone here, then”

Heidi Heiden's First Law:
When you want it bad, you get it bad,
and most people want it in the worst way.

Slide 2

Slide 3

PROFIT

For the Vendor

The Good Old Days

FRITCH (Frame Relay Interim Crutch)

3com cards with 2KB of packet memory.

IOS upgrades: Call Cisco. Start upgrade.

Routing updates stopped packet forwarding

OC12s built on protect side

Two words: DEC Gigaswitch

LS1010s that took 40 minutes to boot

TTM linecards

Pittsburgh construction contractors drop packets more than a cisco 2501 which is running "we poured coffee in the aui port" instead of IOS
-Faisal Jawdat

Today

No more fundamentally broken stuff

Routers are not running out of PPS*, obviating need for things like full-mesh ATM

Methodology is now understood

*Unless you are doing something stupid

Noble intentions should be checked periodically against results
-Warren Buffett,

Basic Axiom

All things come down to the fundamental problem:

- Map demand onto a set of lower level constraints

Constraints are at the lowest level

- Fiber topology
- Forwarding capacity
- Power & Space
- Follow The Money (FTM)

Everything thing else is an abstraction of the above constraints

Differentiators

- Everyone has the same constraints
  - No special routers available to others
  - Speed of light is constant (modulo fiber refractive index in your physical plant)
- So how do you differentiate yourself

  Latency

  Cost (note I did not use price for a reason)

  Open Networks

  Rich connectivity

  OSS/NMS

Latency - Dollar Based Forwarding

Latency Matters in Operational networks

- CDNs
- Latency Topology
  - Different forwarding classes (CBF)

Business Decision ($$-based Forwarding)

- Some traffic is more important than others

Traffic Topology changes very dynamically

Very fast control loops

- Human level visibility per transaction

Example

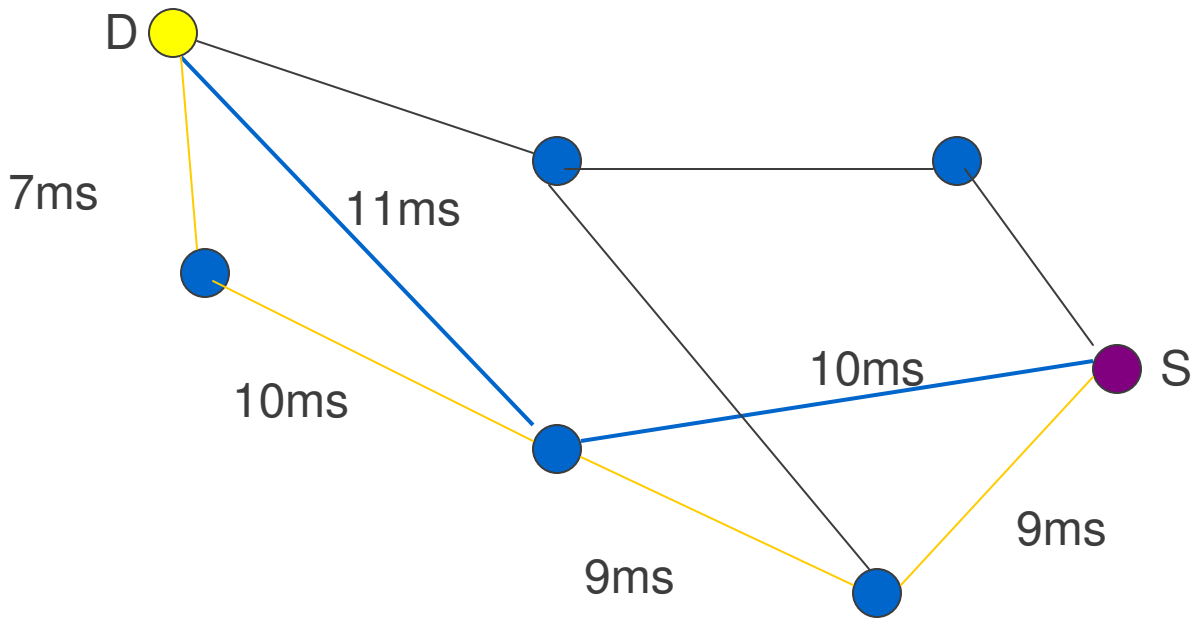

Cost

People that survive will be able to build a network at the lowest cost commensurate with their SLA – Sean Doran

Forgot to add – in a free market

So, what have we learnt

Cost - Efficiency

Based on models that are predicated on 50% or more margin on revenue

Those days are over

The killer App is bandwidth

Become the Walmart ™of the internet

Learn to survive based on 10% profit margins. Or lower.

OSS/NMS are competitive advantages

“it's the bandwidth, stupid”
-John McCalpin

Cost – This time it'll work for sure

And we react with great caution to suggestions that our poor businesses can be restored to satisfactory profitability by major capital expenditures. (The projections will be dazzling - the advocates will be sincere - but, in the end, major additional investment in a terrible industry usually is about as rewarding as struggling in quicksand.)

-Warren Buffet

Cost - Operational Issues

Fewer boxes

- We can barely manage what we have
- Ever growing mass of policy

Troubleshooting simplicity

- 3 AM on-call?

Forwarding complexity

- Actually running out of PPS now

Open Networks - Real Options

Theory of Real Options

- Market uncertainty is high
- Architecture that fosters experimentation at the edge creates potential for greater value than centralized administration
- Distributed structure promotes innovation
- Enables experimentation at low cost

Putting the intelligence in the applications

Spend capital on improving backbone or making application more resilient?

- Pick application
  - Spend can be leveraged across the entire Internet
  - Virtuous Cycle

Open Networks - Wait A Minute

•Whatever happened to “climb up the value chain”

•Decide what business you are in

•Moving up the value chain and providing ever higher-touch services are in direct conflict with providing low cost bandwidth

•Moving from bit transport into services is not the same business

•Picking both

- – Leaves you vulnerable to more narrowly focused competitors in each space
- • Providing services at a cost+ model and letting innovation happen at the edges will result in a virtuous cycle, resulting in every higher bandwidth usage

Rich Connectivity - Peering

Gaming companies are making the same mistakes as the content guys. They ALWAYS over-estimate the importance of the content and vastly underestimate the desire of users/people to communicate with each other and share…

-Joi Ito

The Internet is a network of networks

In and of itself, your network delivers some value

The real value is when your network connects to the other networks

Metcalfes Law, Reeds Law

Rich Connectivity - Peering

The word “peering” is heavily overloaded

- Used for BGP sessions, bilateral interconnects

Does not give insight into what relationship exists

Settlement Free Interconnection

- What people actually mean when they say “peering”

SFI is a business relationship, not technical

- You can peer with someone without it being Settlement Free

Peering Vs Transit

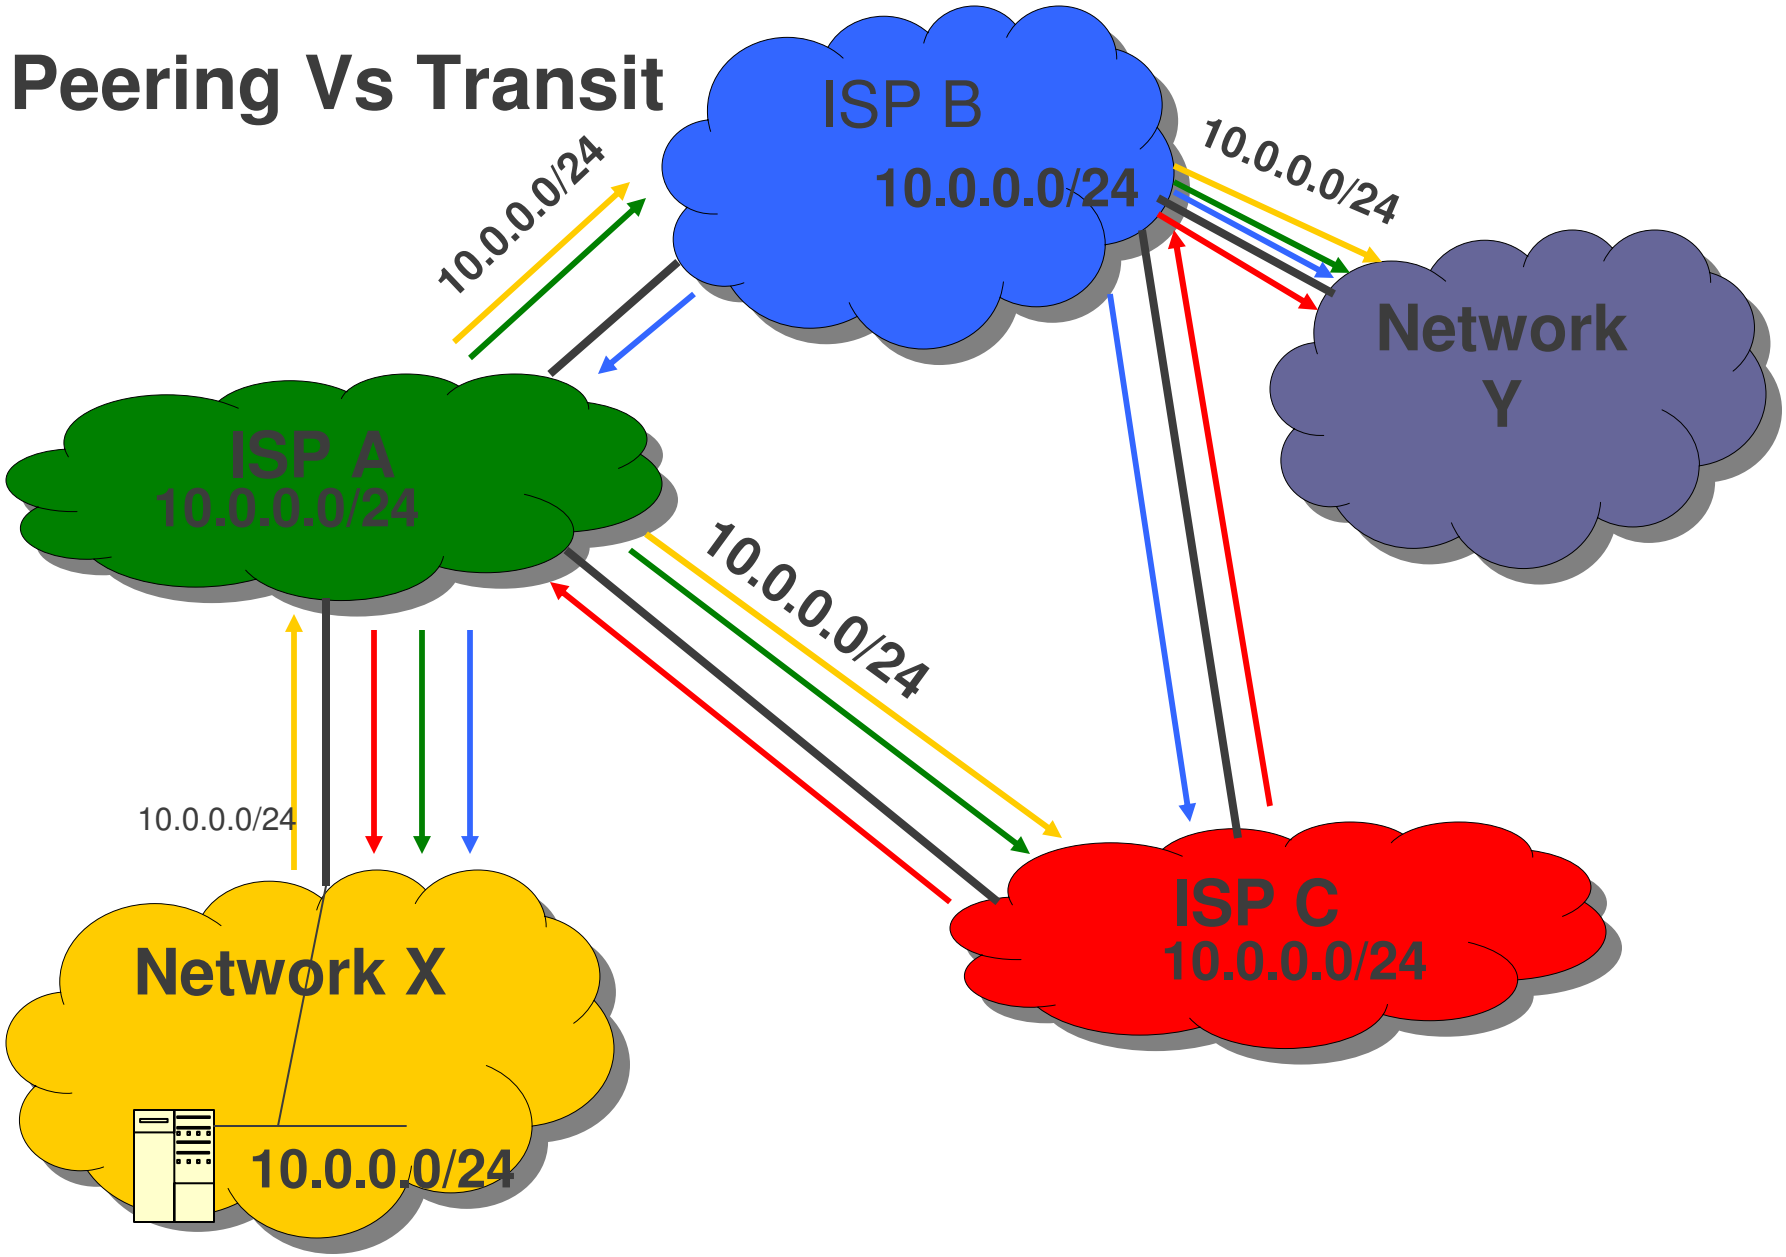

X is a transit customer of ISP *A, Y is a transit customer of B*. As a transit customer, *X* announces 10.0.0.0/24 to ISP *A*. *ISP A*, *B* and *C* peer with each other. Full Mesh at the Tier 1 level required

Donut Peering

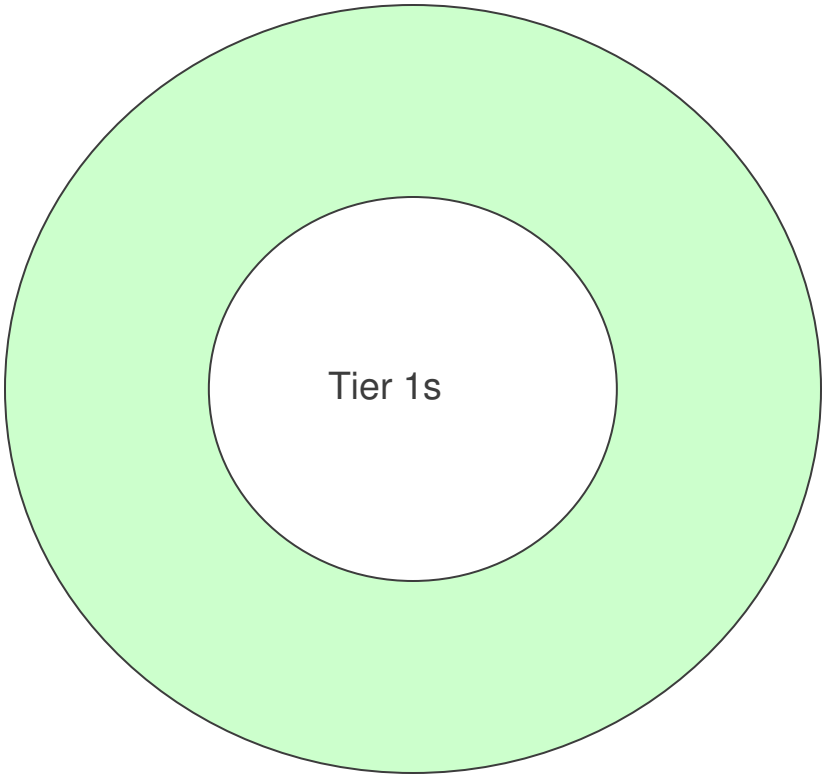

First coined by Bill Woodcock

Allows you to reduce cost while improving latency for customers

Peering is a key differentiator when answering RFPs for content customers

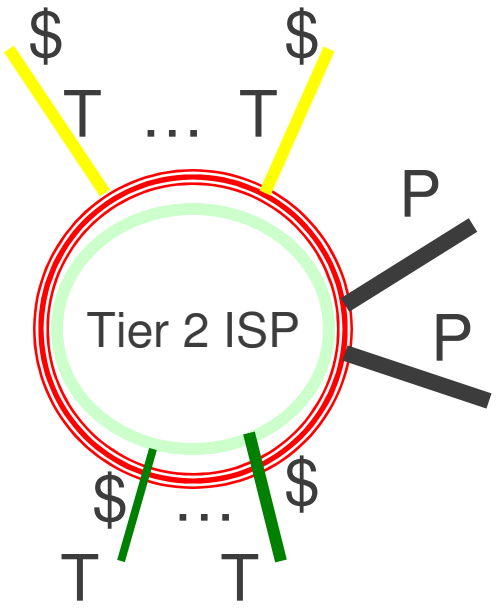

Connectivity End State

Purchase Transit to access some part of the Internet Region.

Peer in region to reduce transit fees.

Behavior: "Open" Peering or "Selective" Peering Policy
Active in Peering Forums

OSS/NMS

Database – One Source of Truth

Not the Network

Allows for faster provisioning

Consistency

Auditing

All these things give a better quality of service to the end user

Competitive advantage in reducing OPEX and SLA payouts due to error in configurations

Modern Design Principles

Complete life cycle in-house

- Architect, design, implement, and operate
- This is your core business
  - –Yet to see an ISP flourish in a free-market that did not hold this as core DNA

Design Criteria

- Diversity of component, paths, logical units
- No Single Point of Failure (SPOF)
  - –Redundant capacity to support peak load
- Routed topology
- System review to ensure performance & survivability goals

"Fault tolerant" is like "tier 1" - the companies that really are fault-tolerant aren't the ones going all over the place claiming to be.
-Robert Seastrom

It ain't the metal, it ain't the glass; it's the wetware.

-Tony Li

Design Principles – Operational Simplicity

Standards based (as much as possible)

Consistency

- Replicate design throughout all POPs
  - Boxes might be different sizes but nodal architecture is the same

"Choose the path ***more*** traveled"

- Problem with being on the cutting edge is that someone's got to bleed

Empower new employees with the ability to easily understand what has been built

- Push tasks into the NOC

"Do not be so proud of this technological terror you have constructed. The ability to criticize Star Wars is insignificant next to power of the Fans"

-Brandon David Short

Design Principles – Operational Simplicity

When to touch the network

- Routing policy based on simple performance metrics and cost (price/Mb)
- No fancy tricks – Traffic engineering based on using common policy across similar peers (e.g., free or transit) rather than exceptions

Achieved through engineering simplicity

- Focus on reducing OPEX
- It's not about building a network that's cool or is a challenge to engineer
  - "Make it simple, but no simpler"

Ask not what evolution can do for you, ask what you can do for evolution.

- Jimbo Kukla

Most Important Lessons

Invest in good software tools developers

Invest in processes and automation

Ruthlessly stamp out network authoritative thinking

Multiple Backbones

Application Aware

Dedicated Infrastructure used to implement each application – PSTN, FR, ATM, Private Line

Discourages Diversity

- Needs Large Market Demand before it is cost effective to go out and build the support infrastructure

ATM/SONET infrastructure

- More boxes, complex management issues
- Hard to upgrade bandwidth in sync for all backbones

Multi-Service Networks

- Simple IP core is the bulk of traffic
- Higher margin products can be leveraged across this core
- Requirement not to share a common control plane
- Use cheap boxes to build a transport fabric (mpls switch)
- Using L2VPN/CCC type techniques overlay multiple networks onto the common transport
- Possible cost reductions in the core fabric
  - Remove Longest prefix match logic/FIB tables/associated logic
  - Switch on fixed labels
  - Drive cost down on linecards
  - Remove logic for timing and sync (100gigE)

Common Backbone

Application Unaware

- Rapid innovation
  - Clean separation between transport, service, and application
  - Allows new applications to be constructed without modification to the transport fabric.

Less Complex (overall)

Alternative (in the real world)

One backbone (underlay) for pure low-cost IP

One High touch backbone for everything else

Mostly because we don't have cheap enough ports at the MPLS switch level yet and (and we won't)

In The Real World

IP traffic will dominate

Design for opaque bits

Walled Gardens ™ will fail (in a competitive market)

Q&A

"There is a difference between making something foolproof and reducing the number of fools"

-Bill Barns

"The venture will be profitable from day one"

-Michael Armstrong (announcing the Concert deal)