# Global Traffic Exchange among Internet Service Providers (ISPs)

OECD - Internet Traffic Exchange

Berlin, June 7, 2001

J. Scott Marcus,
Chief Technology Officer (CTO)

GENUITY

www.genuity.com

# Economics of Internet Interconnection

- Background
  - Internet interconnection in North America
  - International Internet interconnection
- Economic modeling of Internet backbone peering
  - The off-net pricing principle
  - Implications for international interconnection
- Scaling challenges to the peering system
- Cross-provider differentiated services, measurements and SLAs

# Internet Interconnection in North America

- Historical roots
- Peering and transit
- Shared and direct peering
- Shortest exit
- Hierarchical structure

# The Old NSFNET Backbone

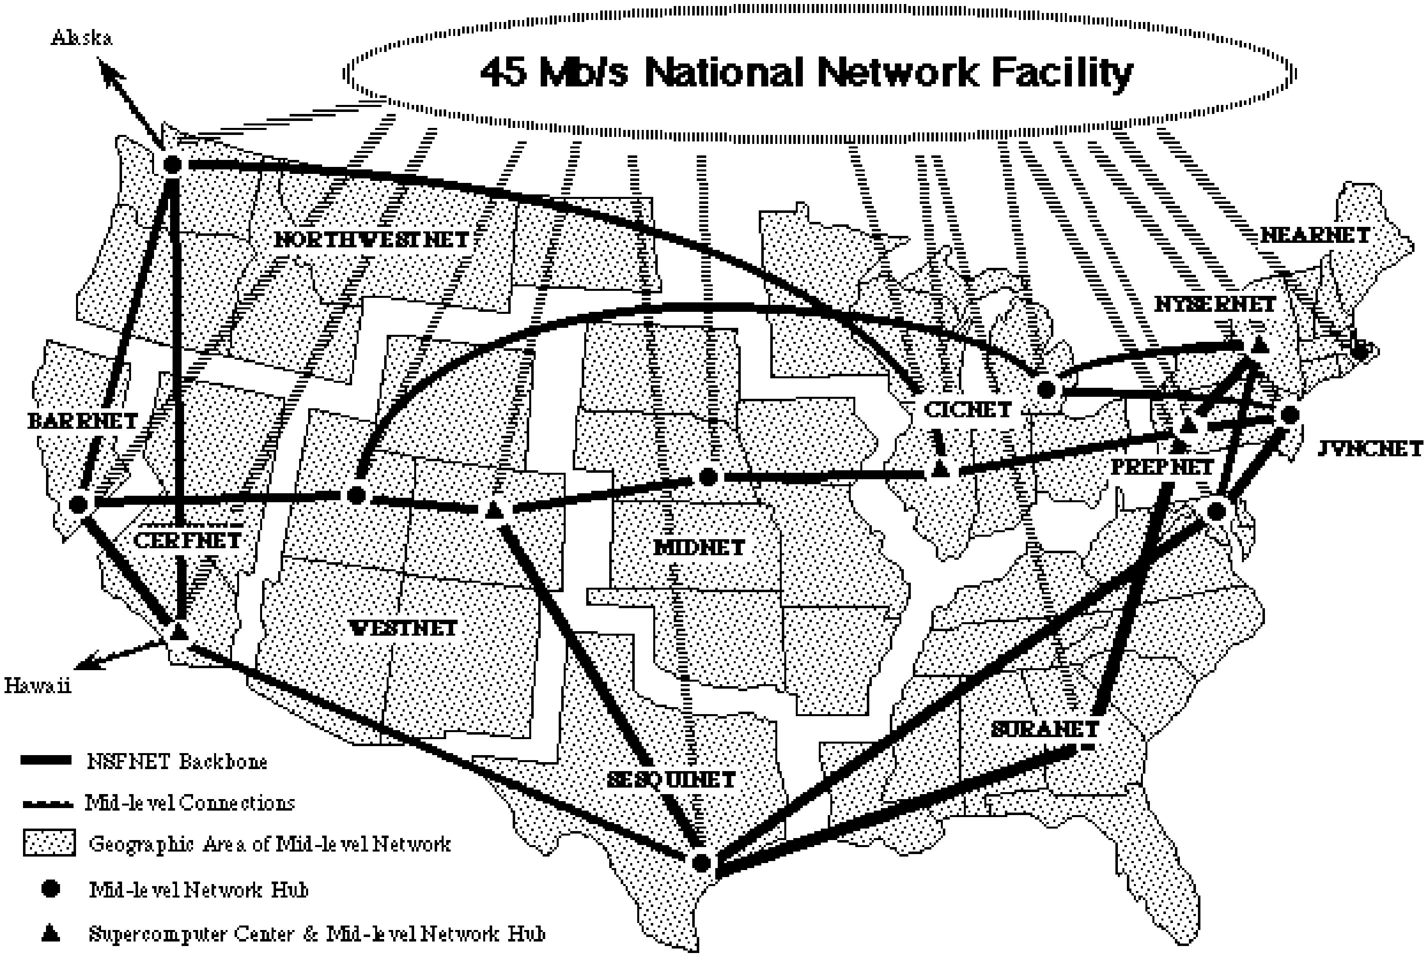

# Privatization of the Internet (NSF 93-52)

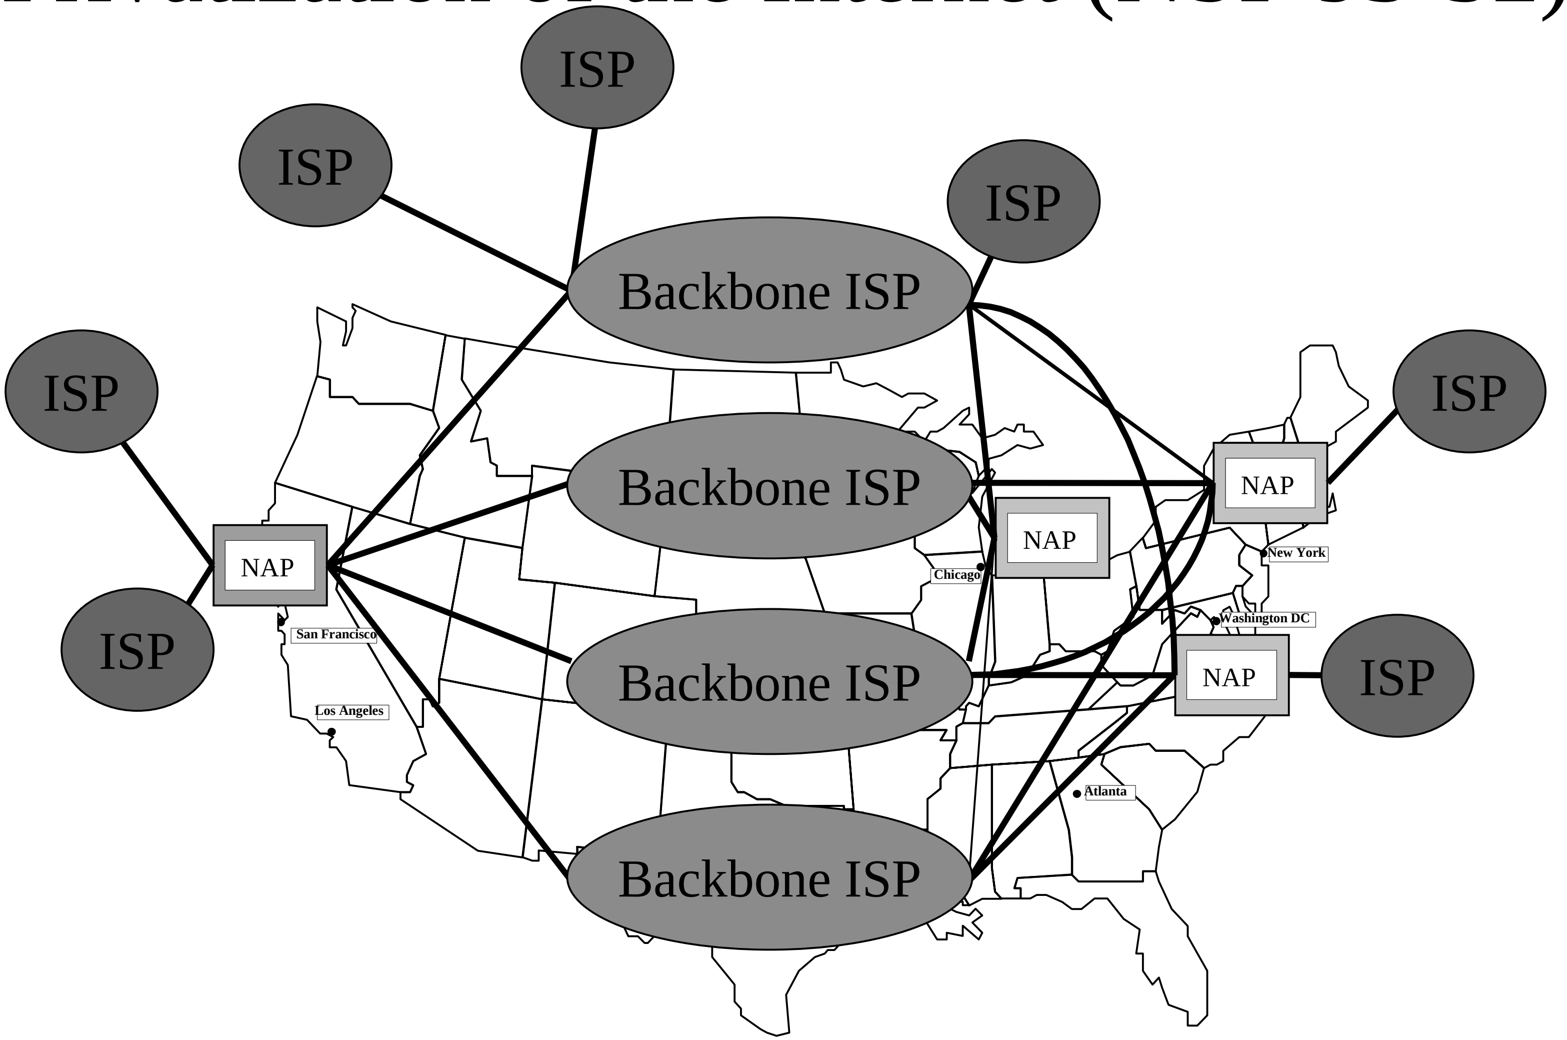

# Peering and Transit

- *Peering* is usually a bilateral business and technical arrangement, where two providers agree to accept traffic from one another, and from one another's customers (and thus from their customers' customers). Peering does not include the obligation to carry traffic to third parties.
- *Transit* is usually a bilateral business and technical arrangement, where one provider (the *transit provider*) agrees to carry traffic to third parties on behalf of another provider or an end user (the *customer*). In most cases, the transit provider carries traffic to and from its other customers, and to and from *every* destination on the Internet, as part of the transit arrangement.
- Peering thus offers a provider access only to a single provider's customers; transit, by contrast, usually provides access at a defined price to the entire Internet.
- Historically, peering has often been done on a bill-and-keep basis, without cash payments, where both parties perceive roughly equal exchange of value; however, there is often an element of barter.

# Shared and Direct Peering

- A few shared global traffic exchange points.
- Smaller domestic shared traffic exchange points for regional concentration and exchange of traffic.
- Direct traffic exchange carries most Internet backbone traffic. Even though shared traffic exchange points are losing market share, their traffic is likely to continue to grow in absolute terms.
- Carrier hotels and fiber interconnects - an emerging trend that seeks to provide the best of both worlds.
- Whether shared or direct, the prevailing pattern is *shortest exit* routing - the sending provider hands off traffic at the point most convenient to the sender.

# Shared and Direct Peering

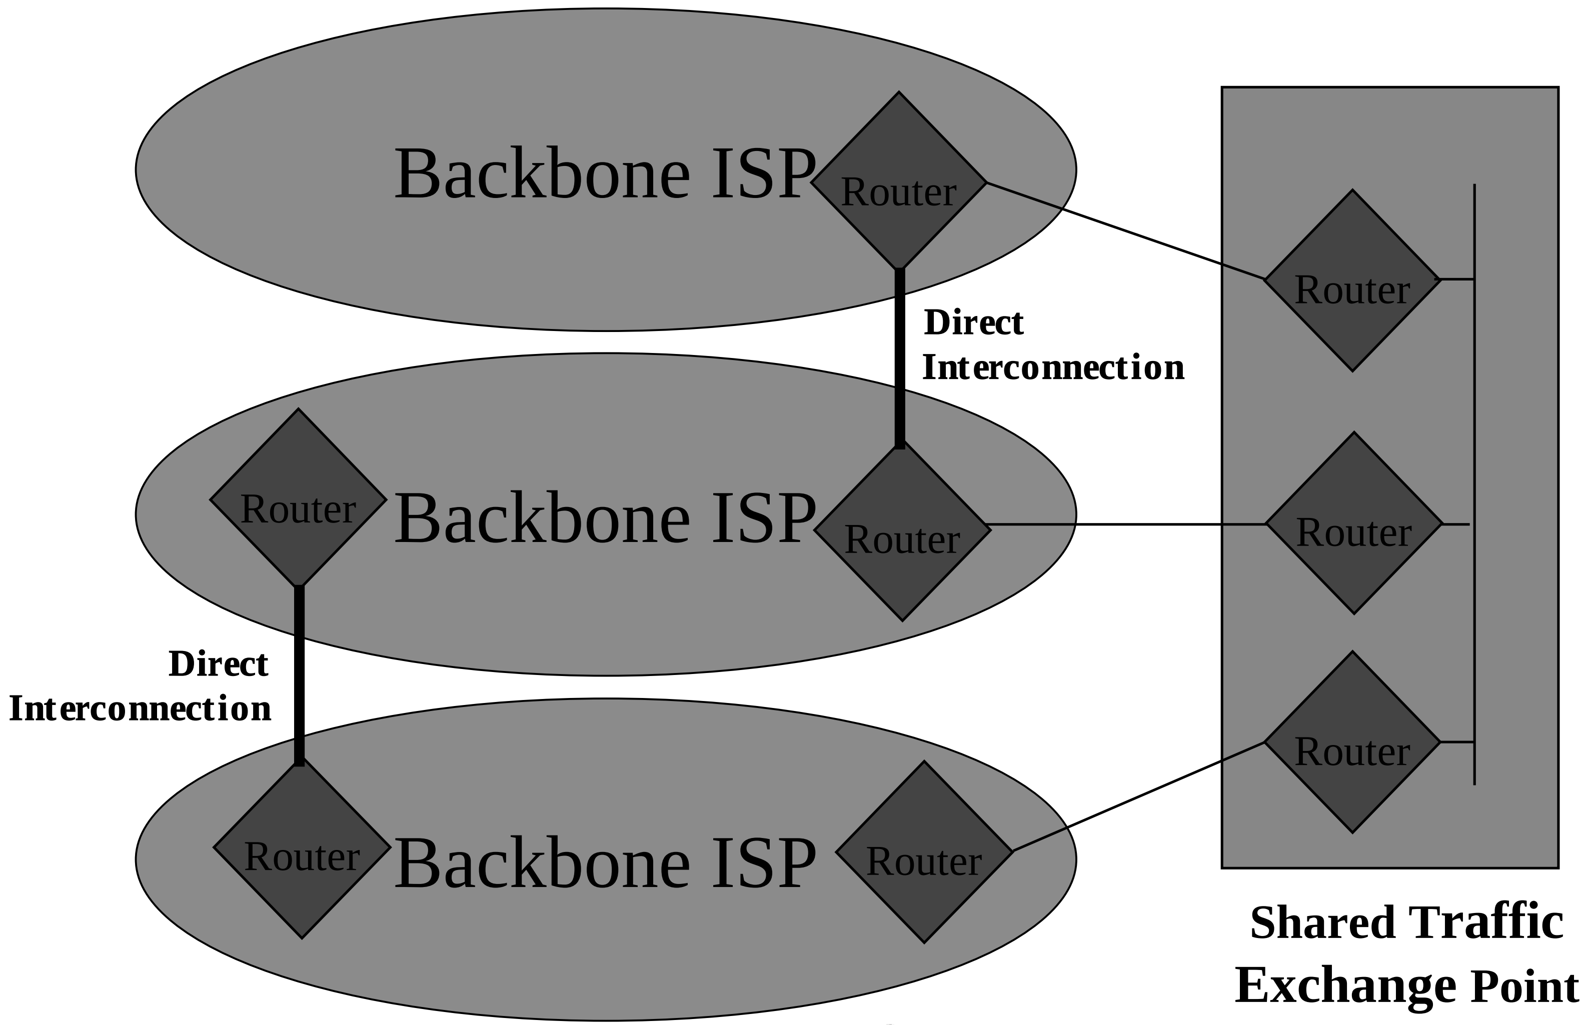

# Shortest Exit ("Hot Potato") Routing

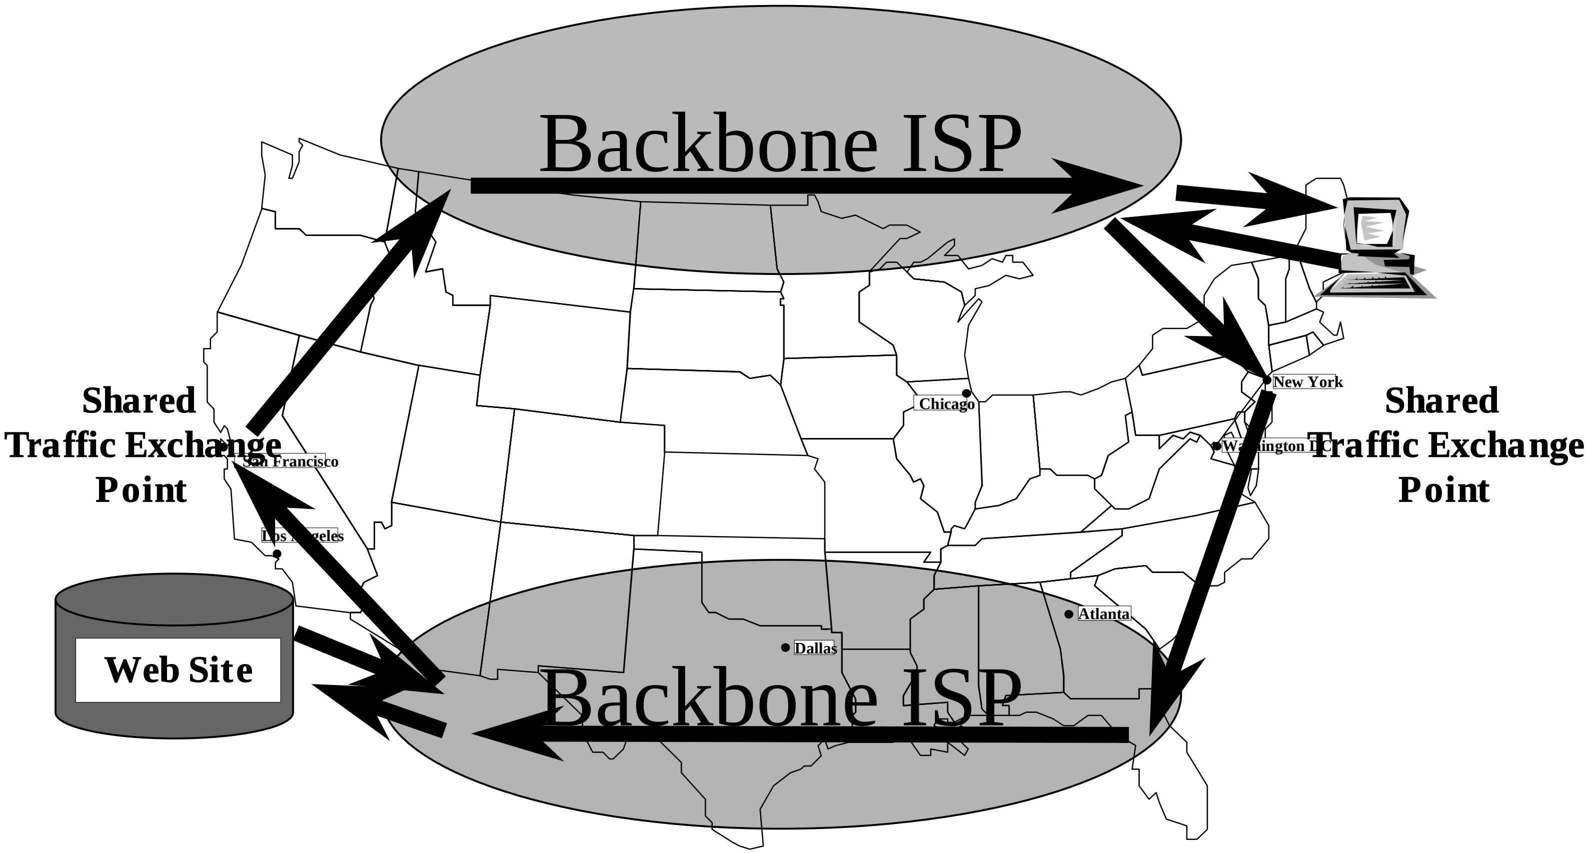

# Backbone and Secondary Peering

- Backbone peering interconnects providers that have no need for a transit relationship.  It provides the only interconnection between those providers.
- Secondary peering interconnects providers who would otherwise exchange traffic through some transit connection.
- Secondary peering is an economic optimization, reducing traffic over the transit connection.

# Fundamental Economics of a Local or Regional ISP

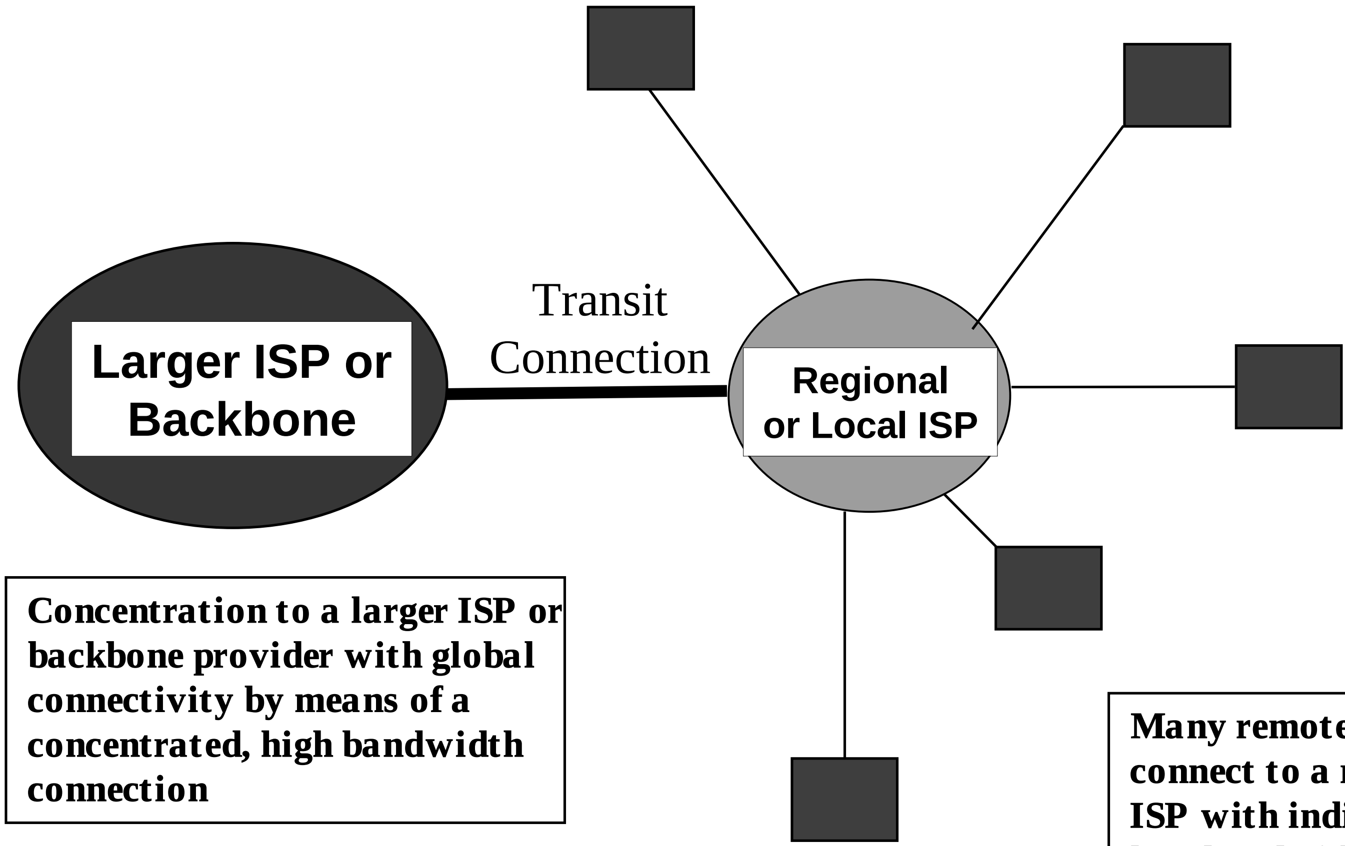


**Concentration to a larger ISP or backbone provider with global connectivity by means of a concentrated, high bandwidth connection**

**Many remote locations connect to a regional or local ISP with individual, low bandwidth connections**

# Secondary Peering

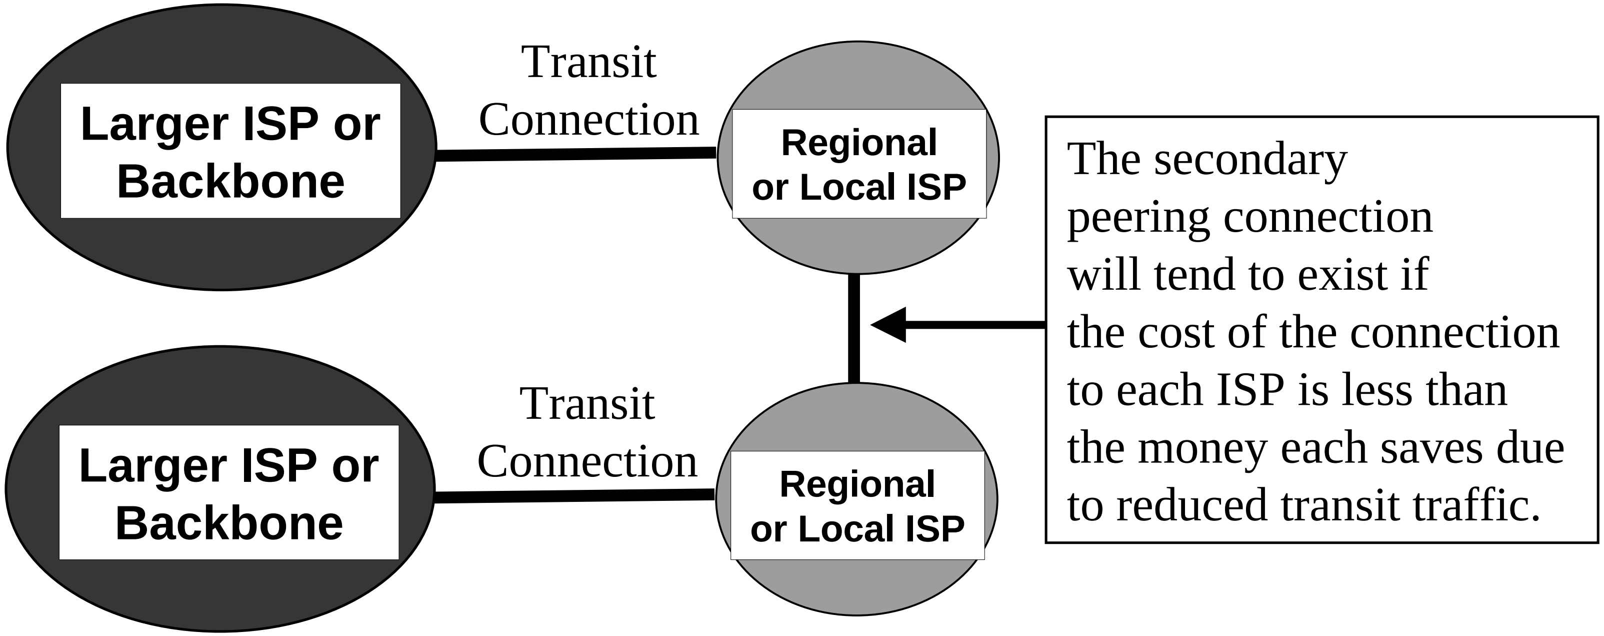

# A Hierarchical View of the Internet

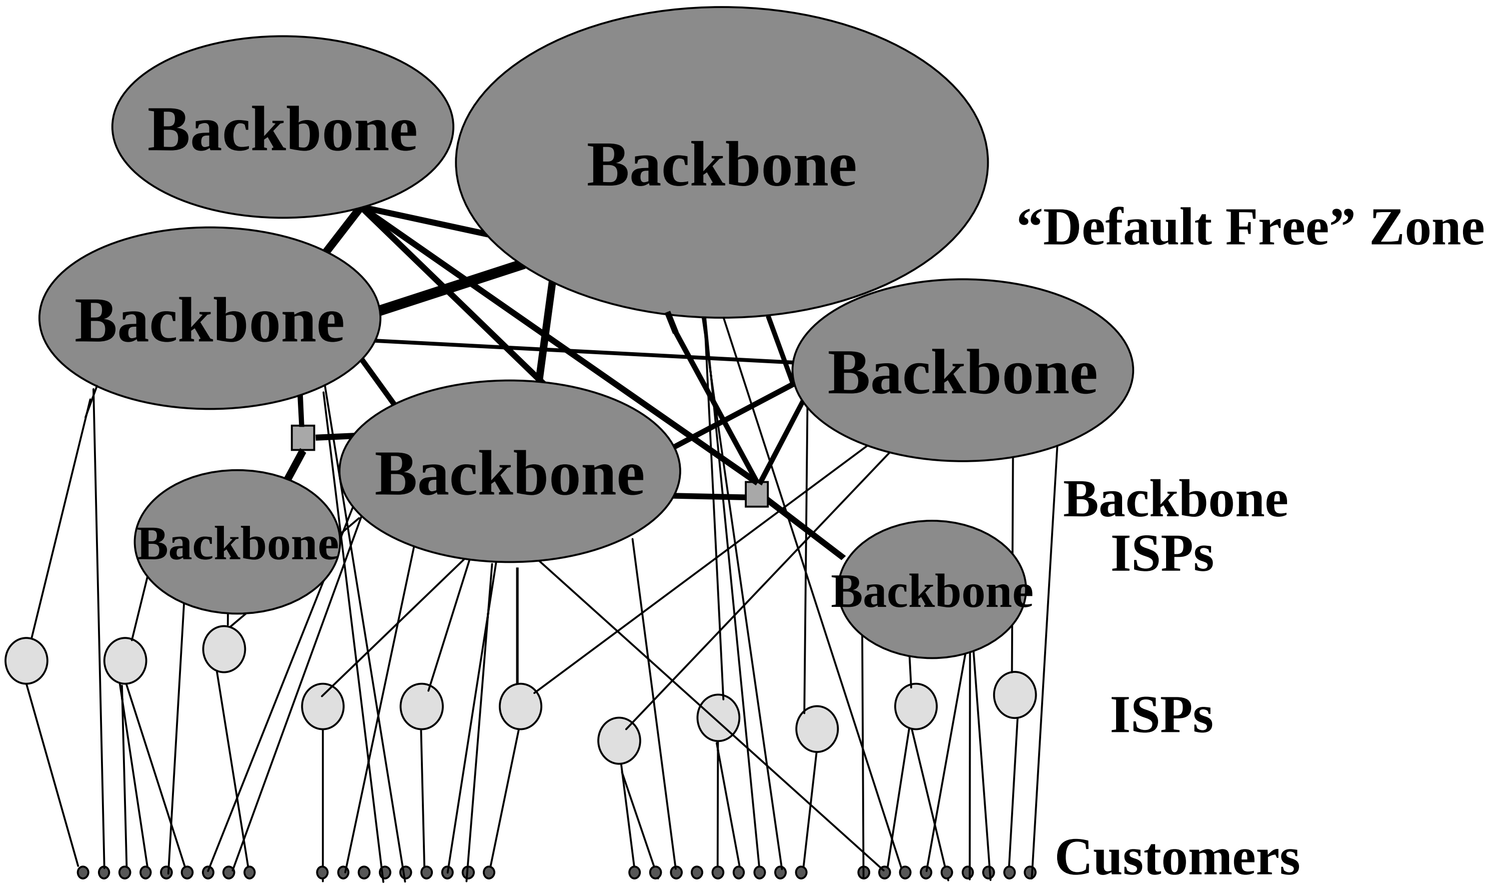

# A Hierarchical View of the Internet

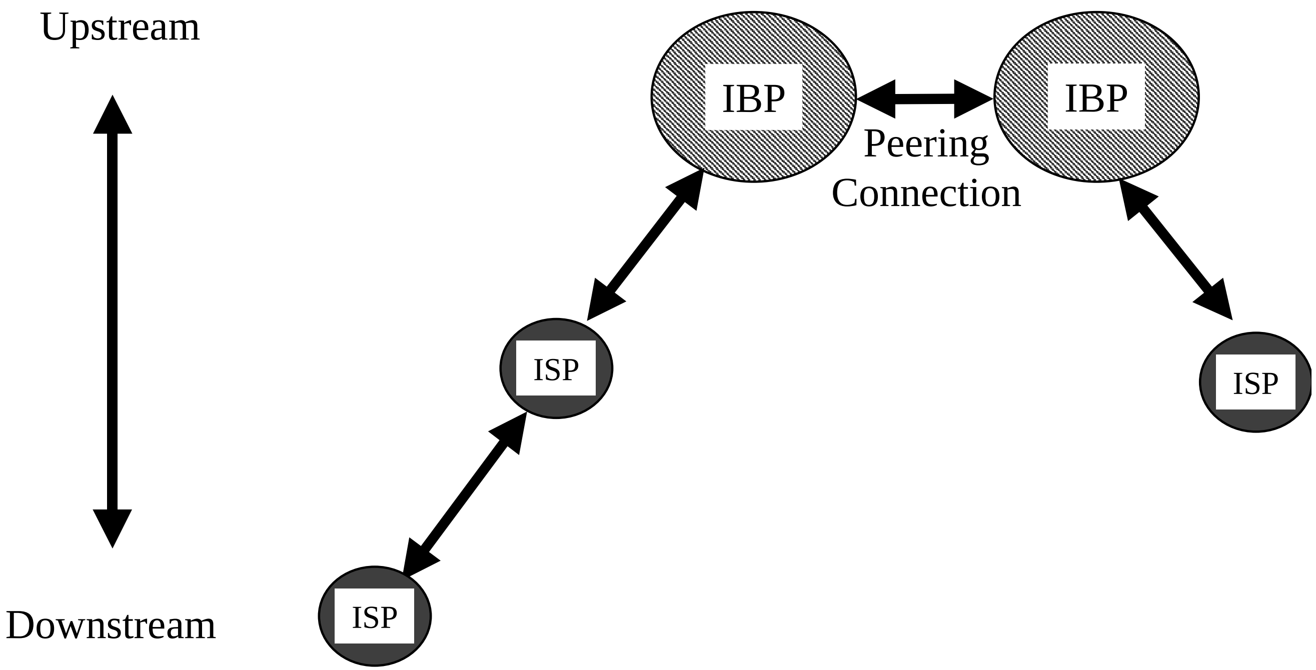


cf. Lixin Gao

# Negotiations for North American Backbone Internet Interconnection

- Typical US backbone interconnection guidelines
  - Bi-coastal US presence, with multiple potential points of interconnection
  - Significant transcontinental bandwidth
  - Consistent routes at all locations
  - Competent staff, professional 7 x 24 operation
  - Rough balance of ingress/egress traffic
  - Sufficient scale to justify transaction costs
- Where criteria are not met, a backbone may:
  - decline to exchange traffic, OR
  - expect cash or non-cash compensation in return
- Backbone providers negotiate traffic exchange terms and conditions on a case by case basis.

# Emerging Global Trends

- The traditional “hub and spokes” system
- Traffic exchange trends in Europe
- Global deployment
- International diffusion of users and content

# Traditional “Hub and Spokes” System

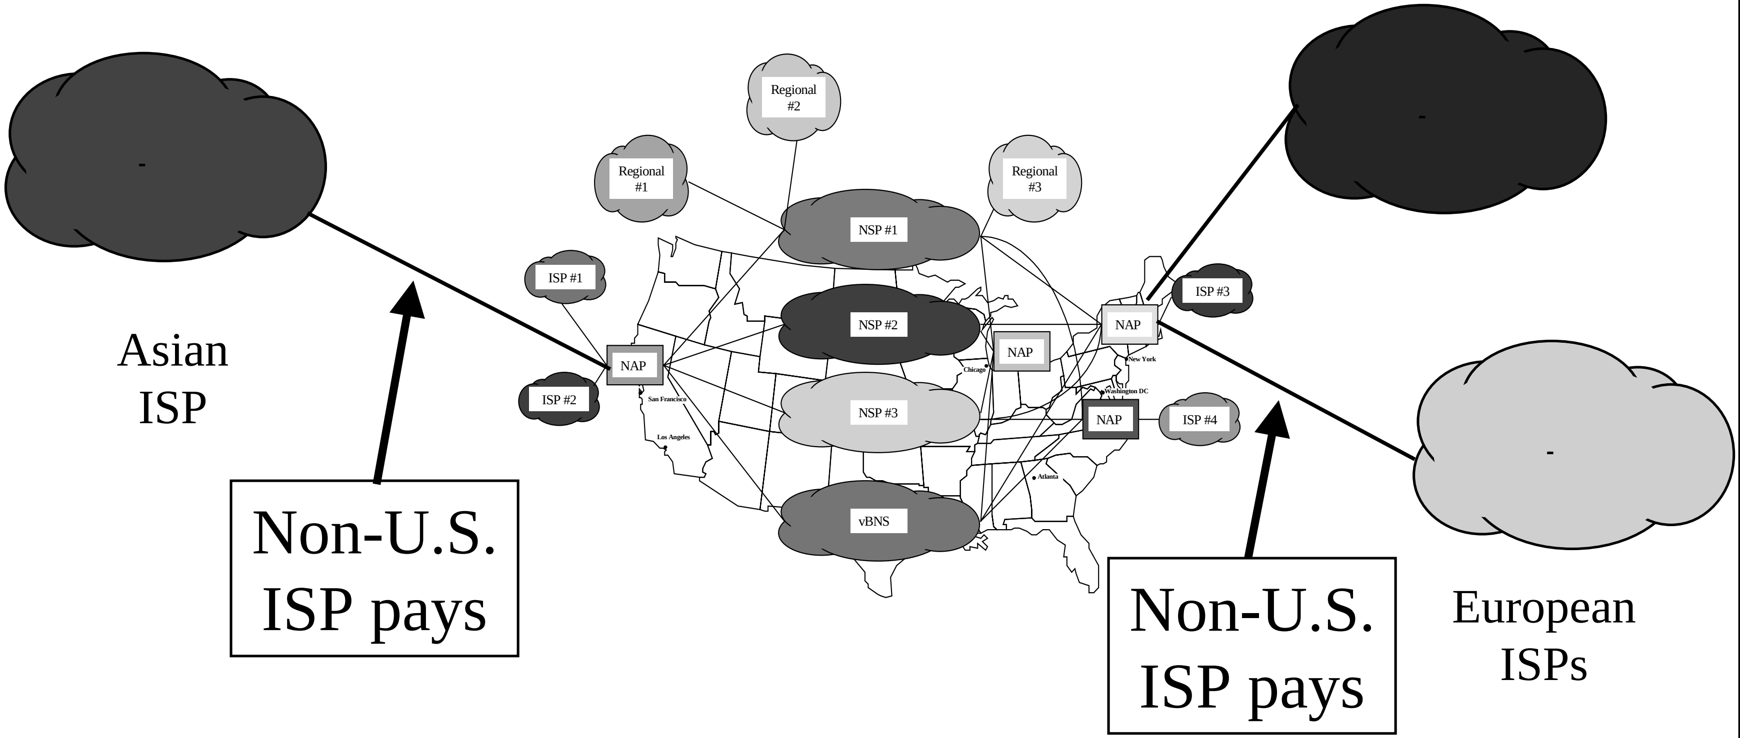

# Former “Hub and Spokes” in Europe

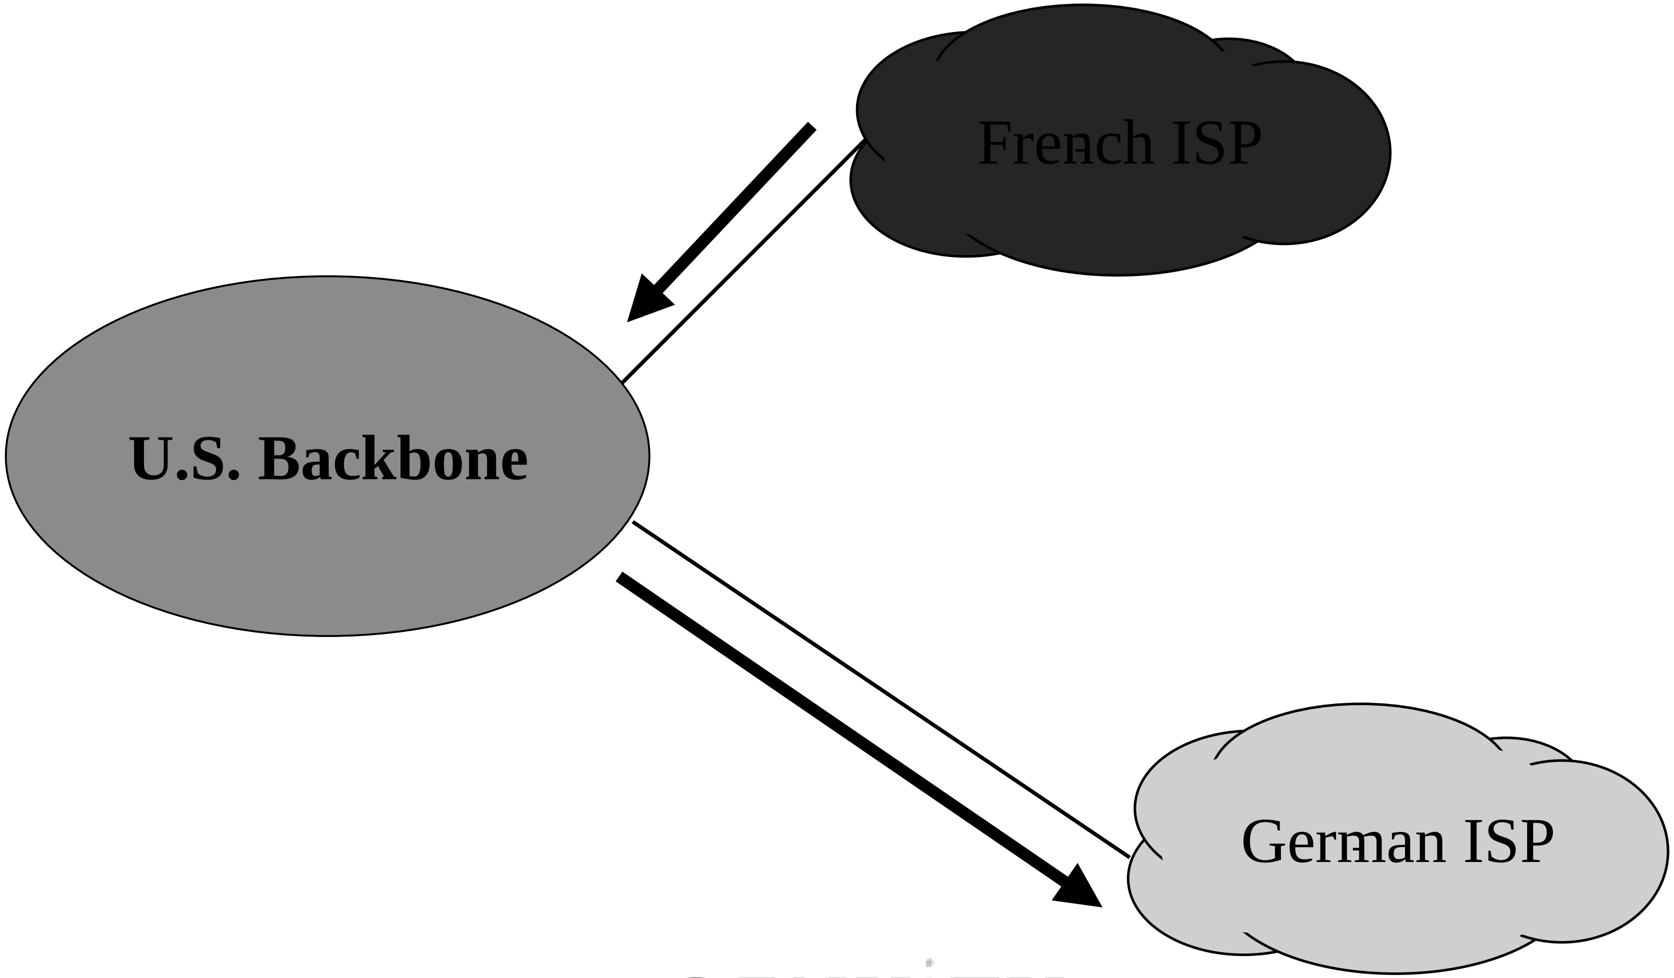

# Internet Traffic Exchange in Europe Today

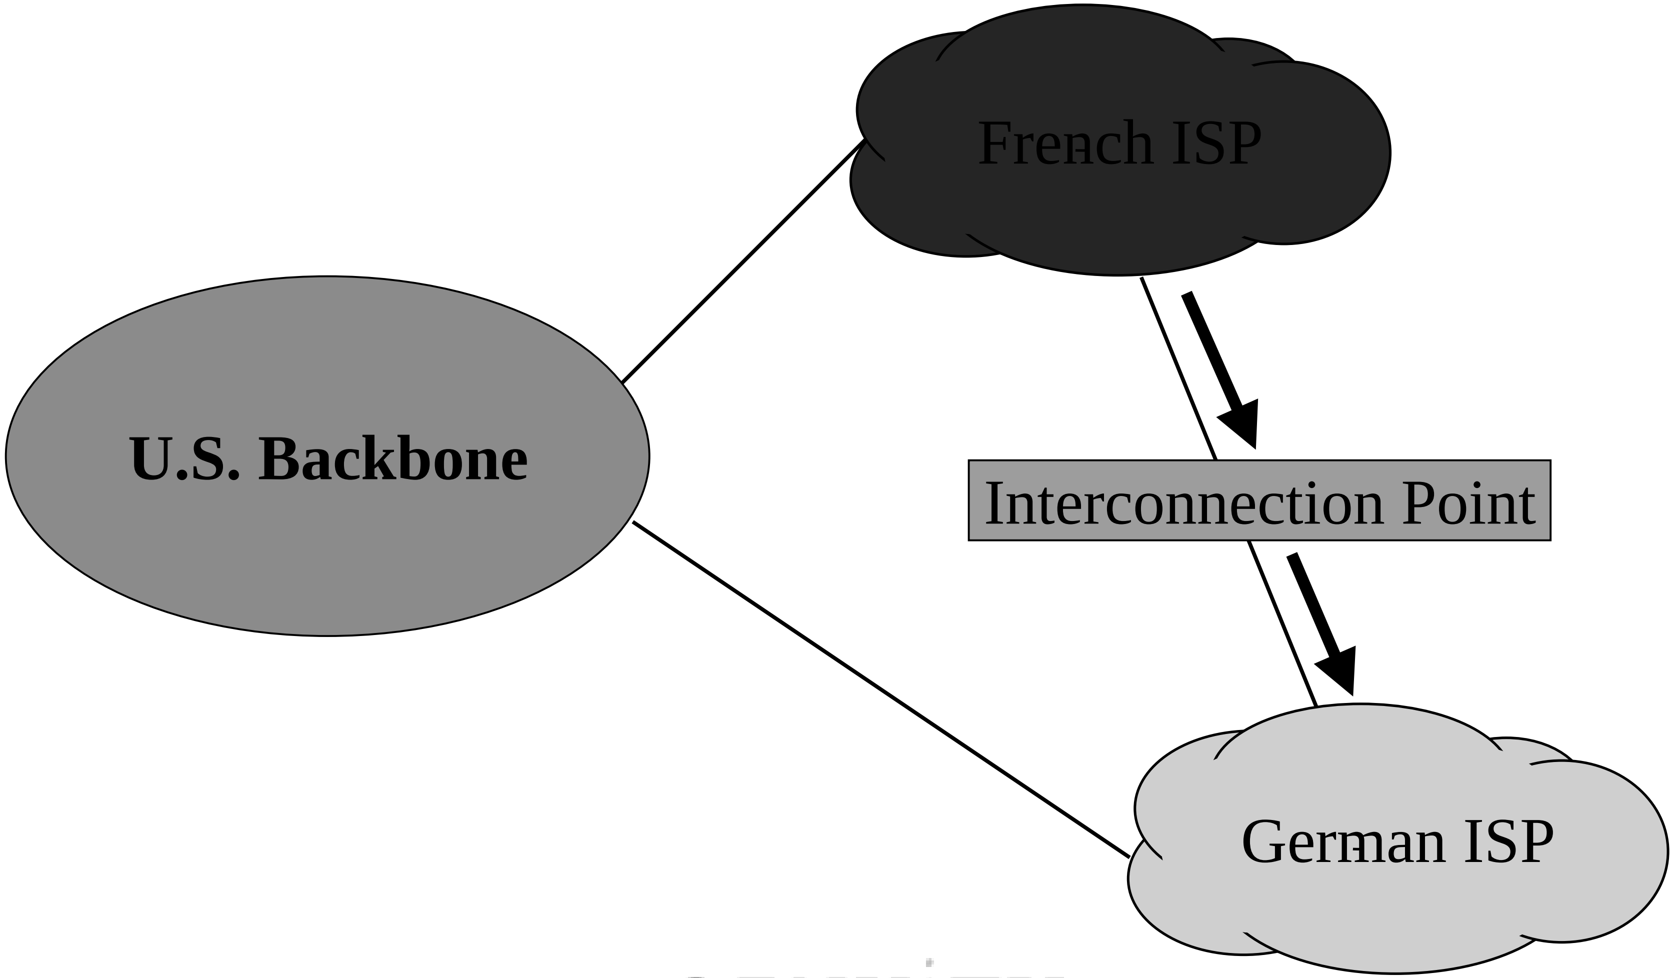

# Factors Driving European Evolution

- Decline in “street price” of circuits within Europe due to deregulation.
- Declining cost of transoceanic capacity.
- Increased number and density of customers and content (and caching).
- Improved number and distribution of shared peering points.
- Deregulation of European telecoms, and recognition of the need to minimize regulatory barriers to Internet growth.
- New transit services terminated in Europe and elsewhere.

# Global Internet Backbone Deployment

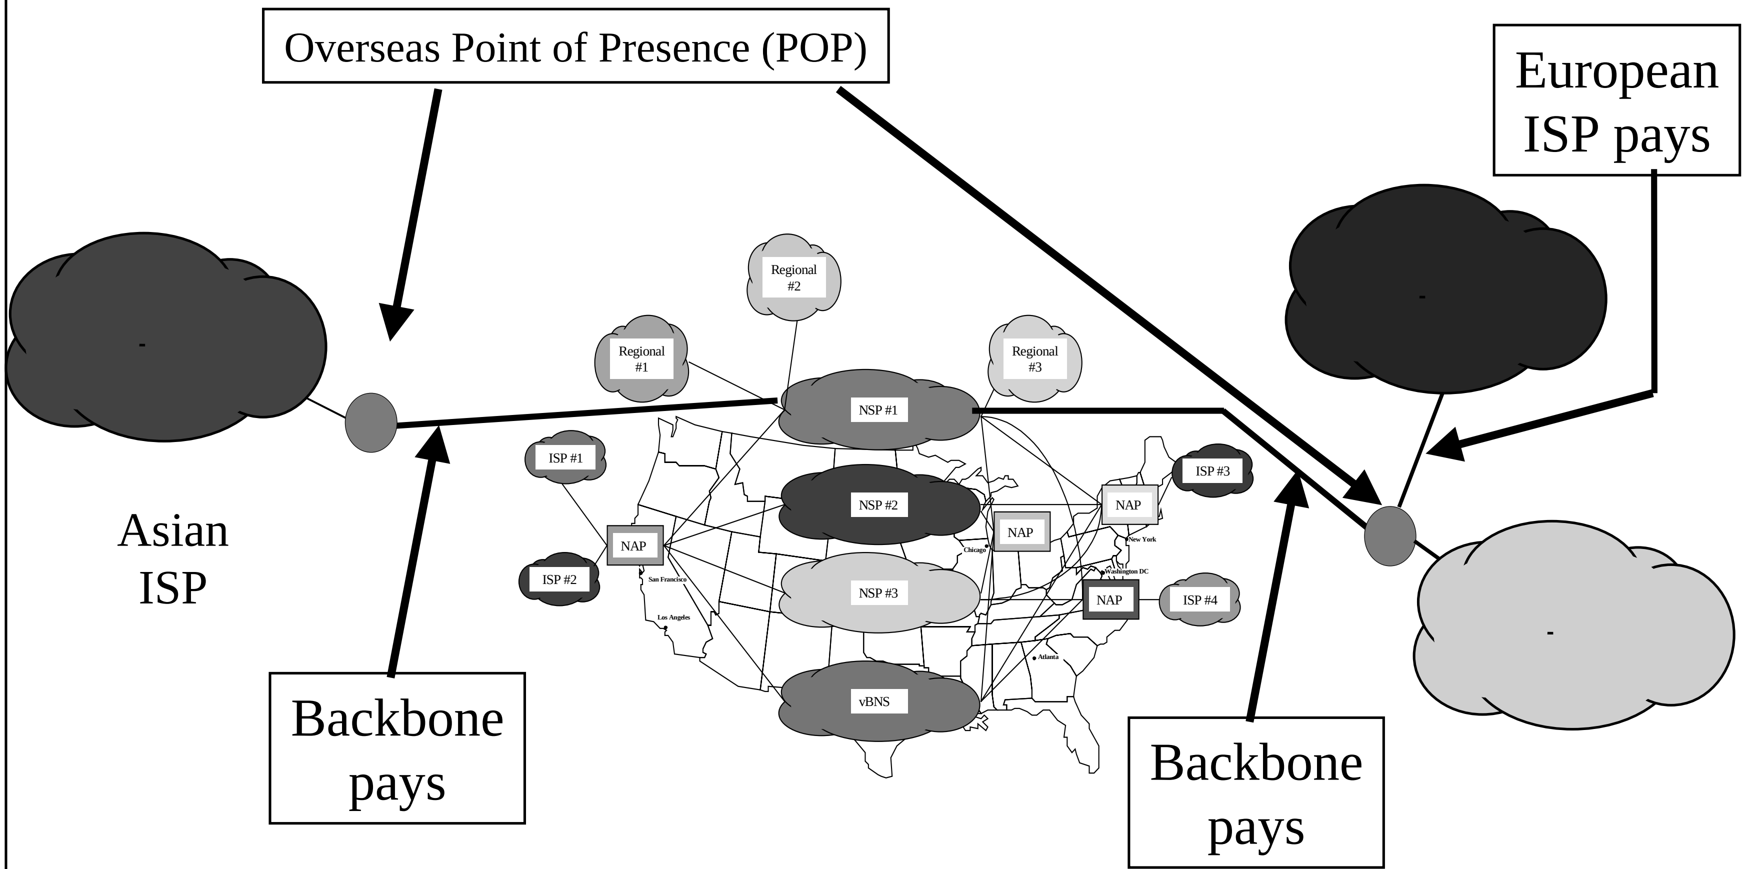

# International Diffusion

Number of Users Online Worldwide

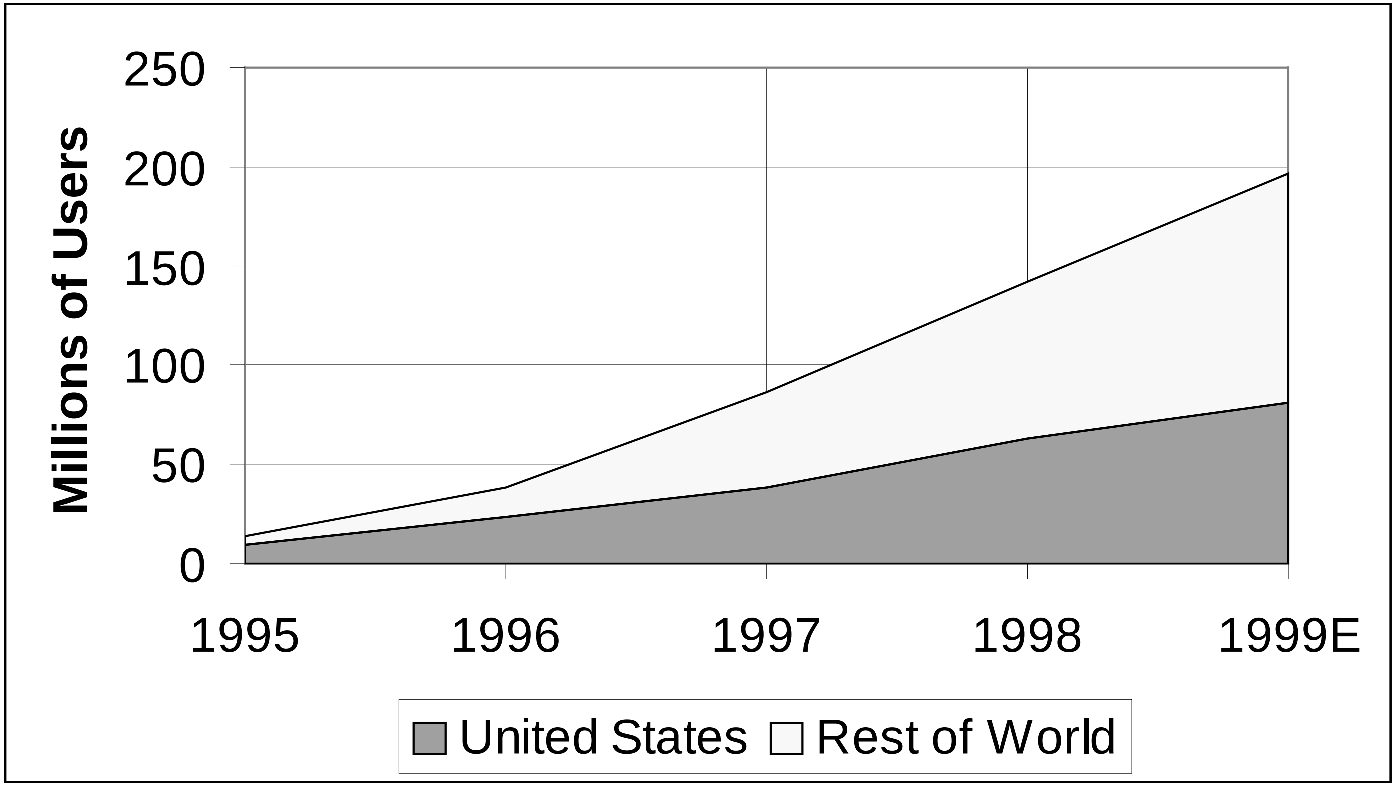


Source: IDC, Merrill-Lynch (Kende, FCC)

# Motivations for International Pressure for Cost Sharing (ICAIS)

- Mistaken perception that U.S.-based backbones discriminate against overseas providers in our interconnection policies.
- Dissatisfaction with allocation of transoceanic circuit costs, which often are fully carried by the non-U.S.-based provider.

# An Excerpt from Genuity's Guidelines

1. Presence at three or more Shared Interconnection Points listed above (two of which must be MAE-East ATM and MAE-West ATM), for Domestic ISPs; presence at two or more Shared Interconnection Points listed above for International ISPs.
2. For domestic ISPs, United States coast-to-coast nationwide backbone of at least 155Mbps.
3. Consistent route announcements at all exchange locations.
4. Experienced, professional Network Operations Center staffed 24x7.
5. Loose Source Record Route (LSRR) capability at core border routers on network.
6. For domestic ISPs, roughly balanced traffic.
7. A minimum Internet traffic exchange of 1 Mbps with Autonomous System 1.
8. Willingness to enter into a formal Internet Interconnection Agreement.

http://www.genuity.com/infrastructure/interconnection.htm

# Economic Modeling of Internet Backbone Peering

- Analytical framework
- The off-net pricing principle
- Implications for international interconnection

# Analytical Framework

- Define:

  $c_o$ as cost of origination
  $c_t$ as cost of termination
  a as an access charge levied on the sender

- Due to shortest exit, $c_t > c_o$
- Then

  cost for the originating network is $c_o + a$
  cost for the terminating network is $c_t - a$

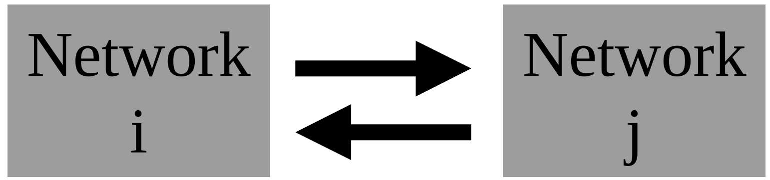

# Impact of Access Charges

- In a bill-and-keep system, access charges are zero.
- With equal access charges and symmetric traffic, net access charges are still zero.
- Providers will, however, view their *marginal* costs quite differently. (Laffont/Rey/Tirole)

# The Off-Net Pricing Principle

- Under a broad range of conditions, marginal price should equal marginal cost.
- Providers will tend to charge the same for on-net traffic as for off-net traffic, i.e. there are no strong incentives for price discrimination.

# Implications for International Interconnection

- Where a group of customers have distinguishable costs that are higher than those of other customers, the model predicts that marginal prices will be higher to exactly offset the higher costs.
- In the absence of market power, there is in fact no way for the ISP to absorb the cost difference without reflecting it in the price.
- This prediction is consistent with typical practices of US-based ISPs, both within the US and overseas.

# Scaling Issues

- AS Number growth
- IPv4 address growth
- Routing table growth

# Tony Bates's Data

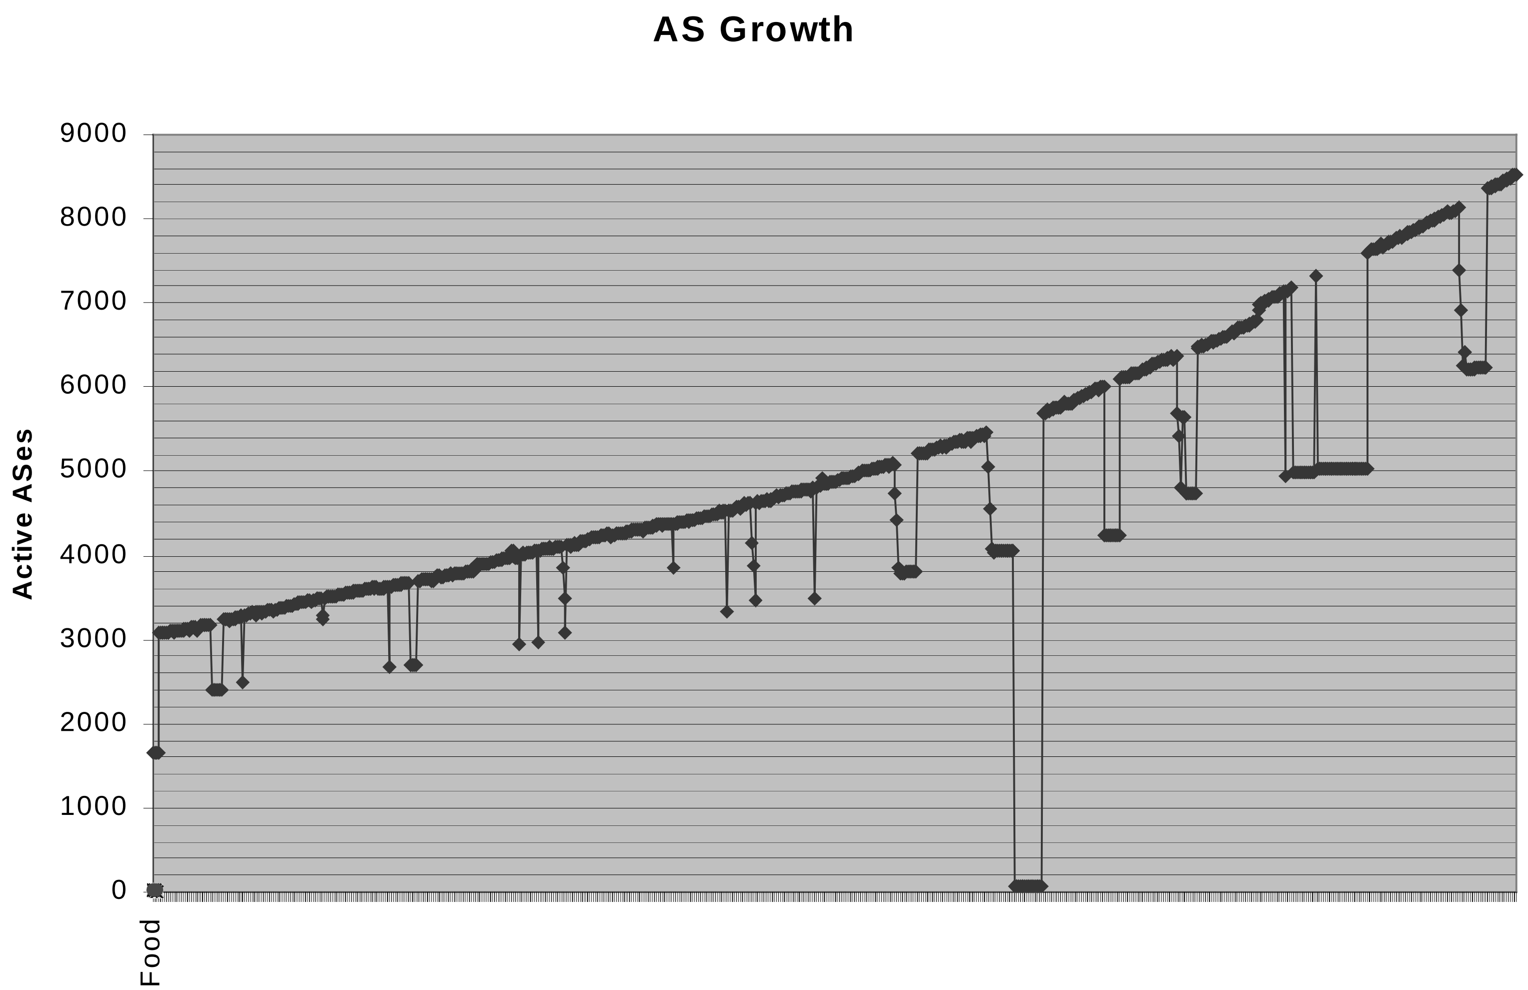


GENUITY
www.genuity.com

# Exponential Growth of Autonomous System (AS) numbers

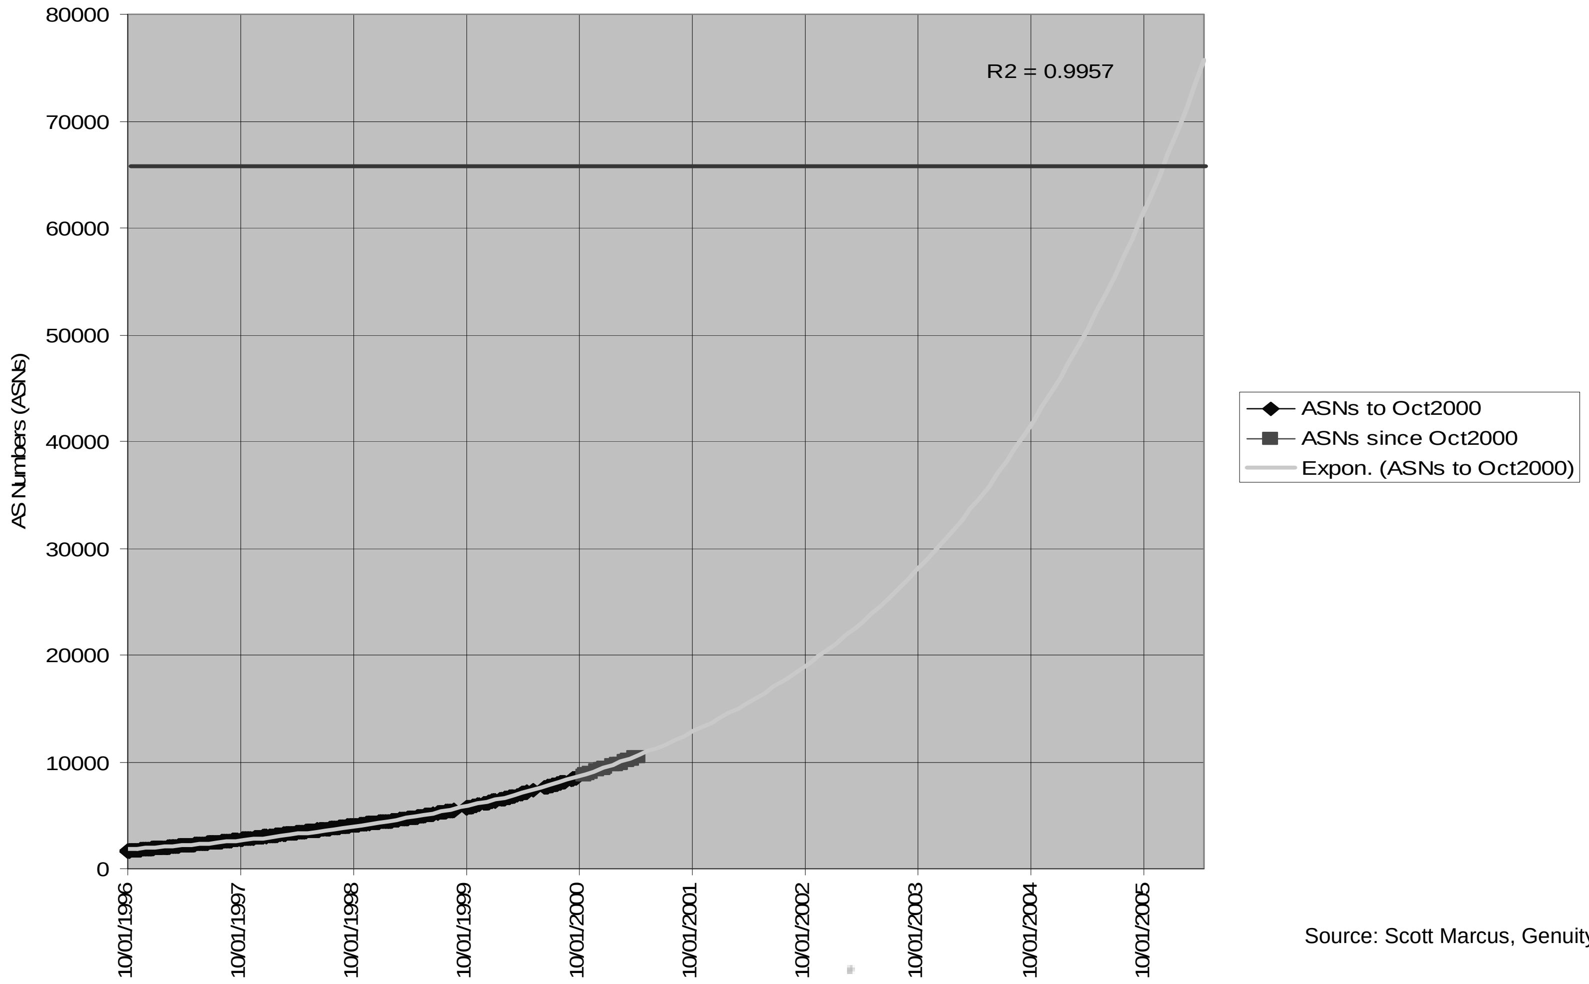


Source: Scott Marcus, Genuity

# Exponential vs Quadratic (Bates Data)

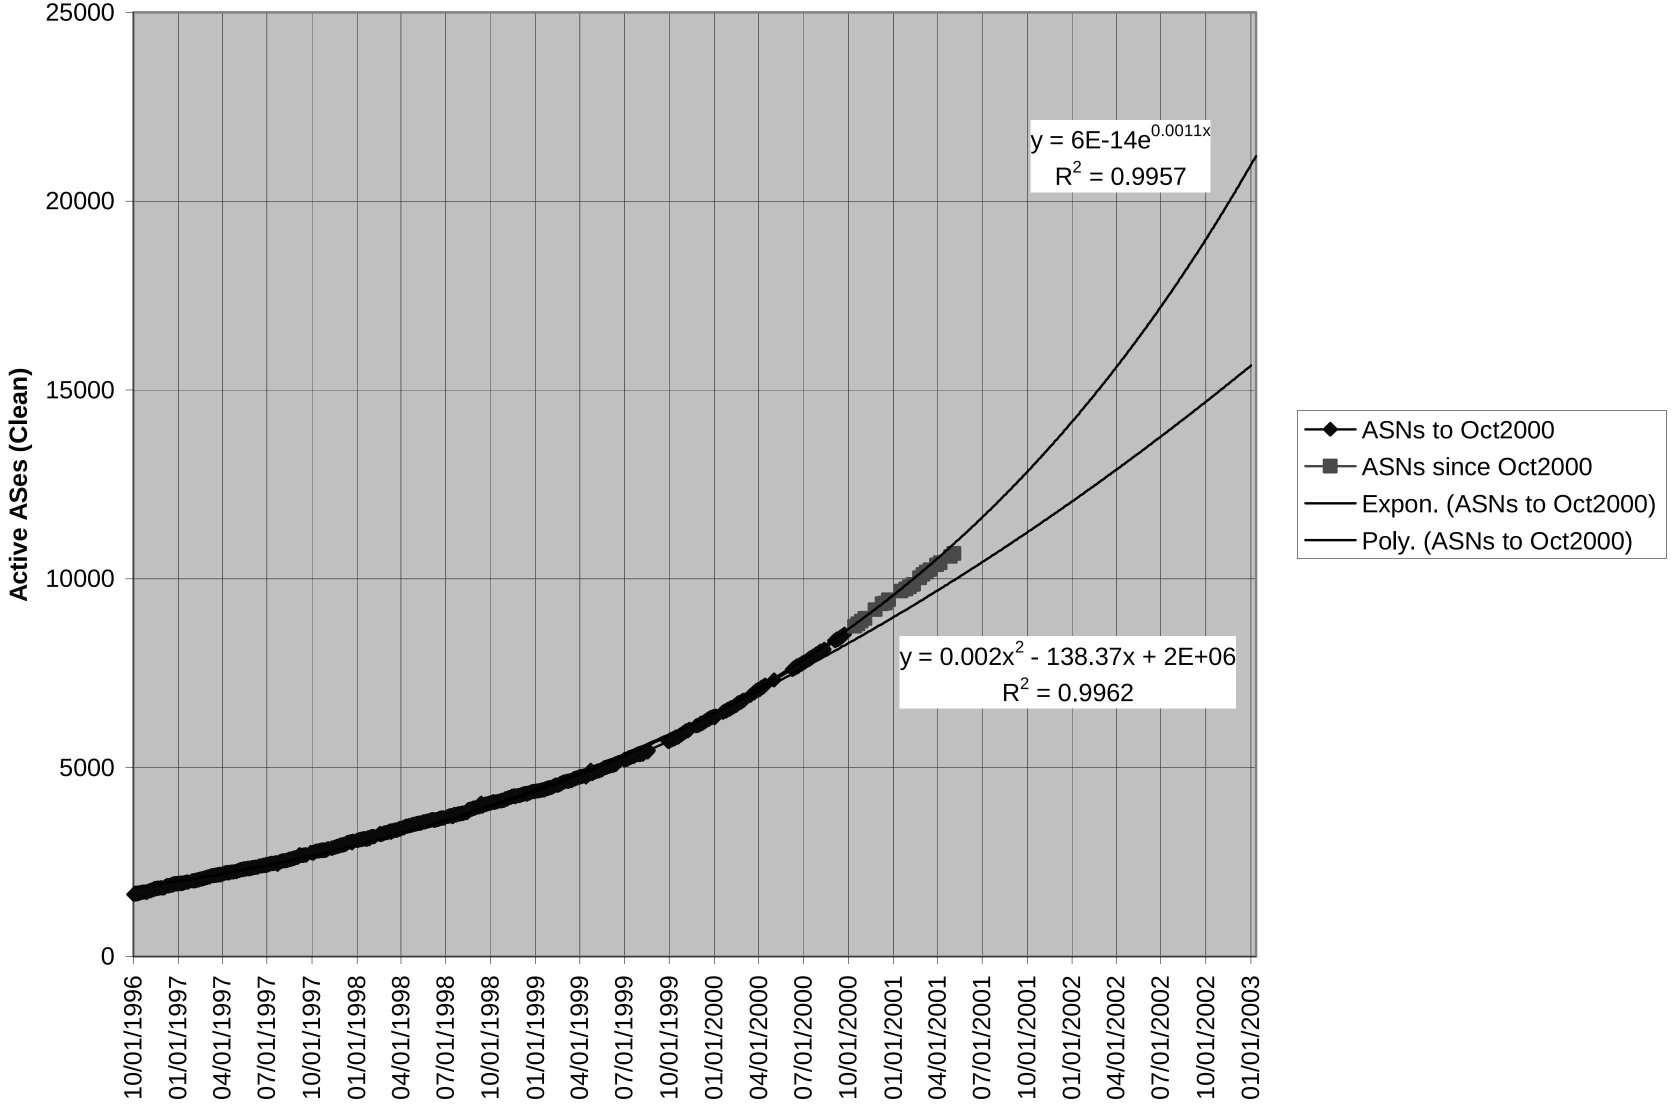

# Chicken Little was Wrong!

- This is far simpler to remedy than IPv4 address exhaustion, because
  - the solution need not impact end systems (hosts);
  - the solution need not impact DNS; and
  - the solution need not impact routers unless they speak BGP-4.
- Any solution is complicated by the need for backward compatibility and phased migration.
- Time until exhaustion is nonetheless sufficient to architect, design, implement and deploy solutions.
- Cisco and Juniper are reportedly well into implementation.

# The RIRs Recognize the Need for Forecasting

- Continuing need to further refine projections.
- Need for forward-looking proactive forecasting on a regular basis not only for AS numbers, but also for route table entries and IPv4/IPv6 addresses.
- Forecasting needs to incorporate allocation data from all three RIRs (APNIC, ARIN, RIPE NCC).
- Forecasting needs to be institutionalized by the RIRs themselves, with data readily available to independent researchers.

# The Team

- Assembled by ARIN
  - Frank Solensky    Gotham Networks
  - kc claffy    CAIDA
  - Scott Marcus    Genuity
- Active contributions and support by APNIC and RIPE NCC

# IPv4 Address Allocations

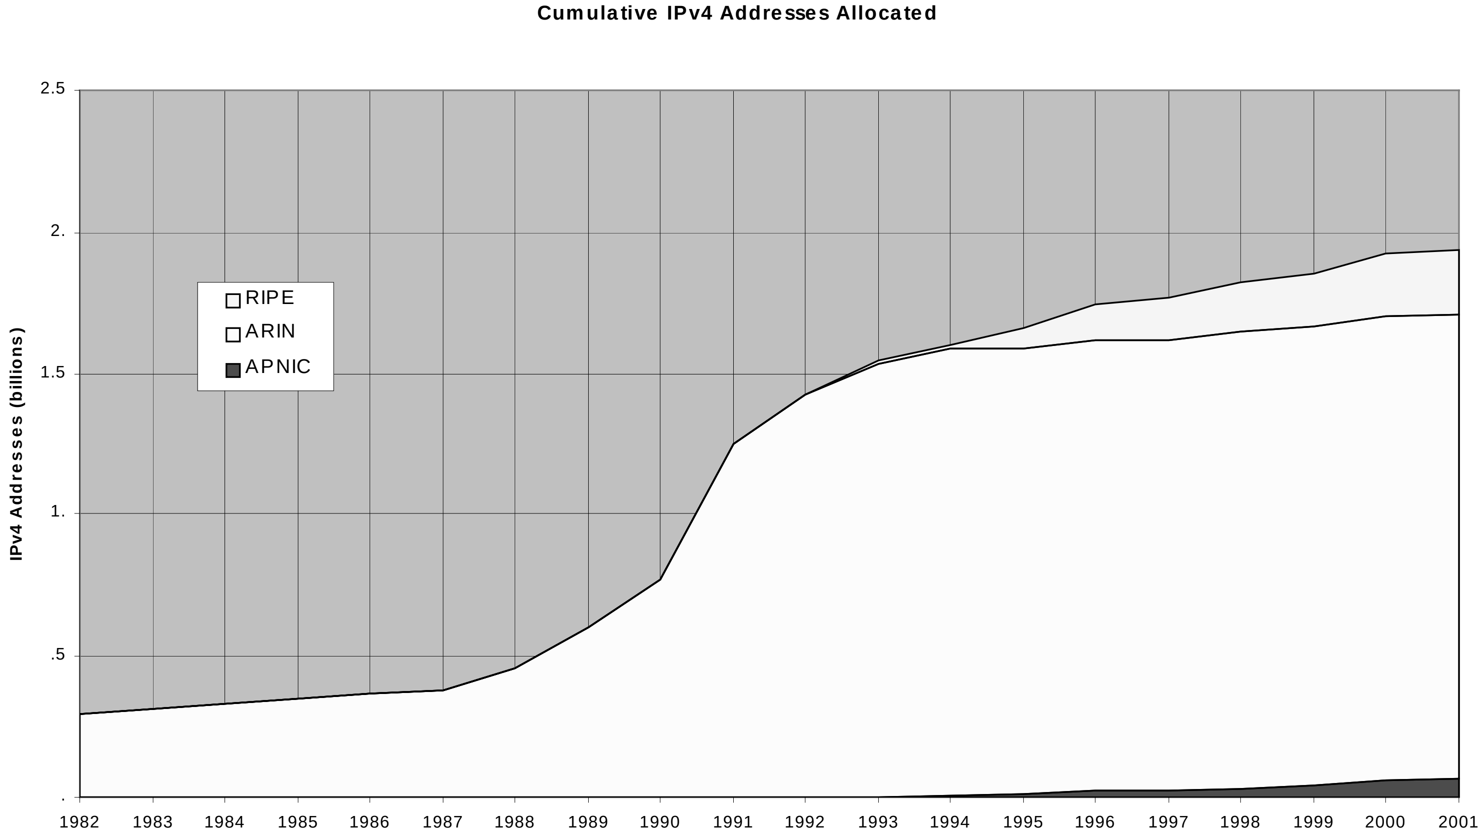


Source: Solensky/Marcus/Claffy

# Route Table Growth

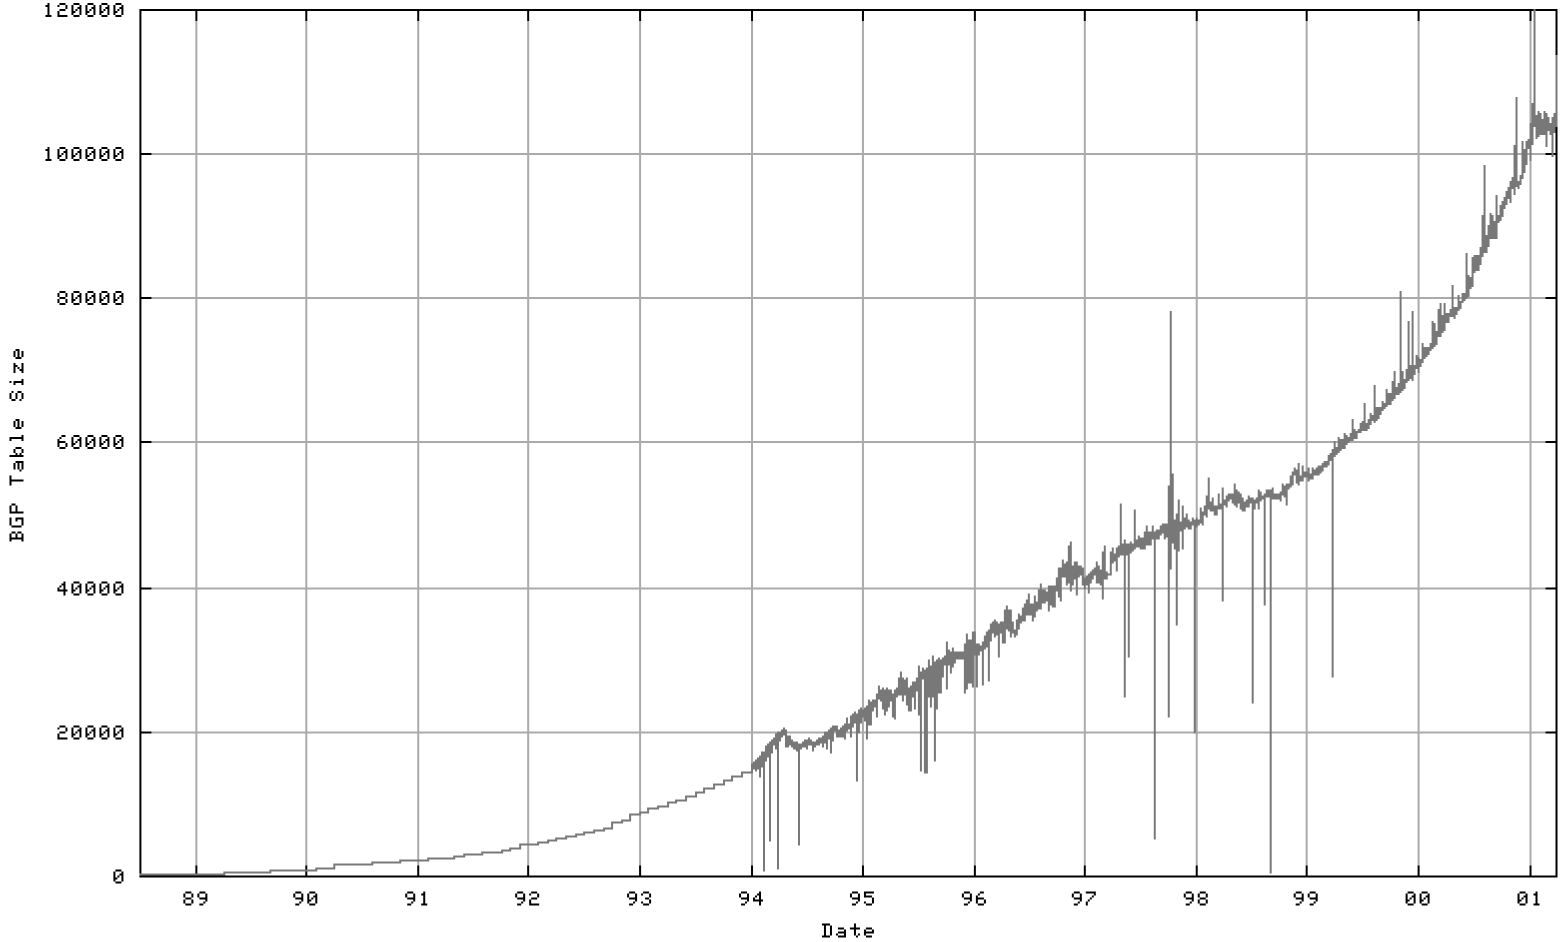


Source: Geoff Huston

# Applications Have Distinct Needs

- E-mail is tolerant of delay, but every bit of data must ultimately be delivered correctly.
- For real time WWW traffic, delay must be low, but moderate variability of delay is permissible.
- Real time Voice over IP (VoIP) and video are tolerant of modest loss, but intolerant of large and variable delay.

# Cross-Provider Differentiated Services

- Internet services today are best-efforts.
- Commoditized service, with marginal price equal to marginal cost.
- Differentiated services are of limited value today.
  - Best-efforts are *good* efforts, so little incremental advantage for differentiated services.
  - Customers have no demonstrated propensity to pay a premium for better CoS for data services.
  - Better CoS for real time voice and video still interesting.
  - CoS today is limited to a single provider’s network.

# Cross-Provider Measurements and SLAs

- Potentially valuable for best-efforts traffic; essential for Differentiated Services.
- Provider responsibility ends today at the peering interconnection point.
- No consistent measurement framework.
- No business arrangements to incent SLA adherence between peer networks.
- Information sharing of proprietary data is a difficult but not intractable problem.

# Summary

- The distinction between peering and transit is vital to an understanding of the underlying business models that drive global Internet connectivity.
- The business models that have evolved around peering, in North America and globally, are complex but rational.
- Peering business relationships have adapted quickly in response to economic or regulatory stimuli, and will continue to evolve.

# References

Cremer, Rey and Tirole, “Connectivity in the Commercial Internet”, presented at “Competition and Innovation in the Personal Computer Industry”, San Diego, April, 1999.

Marcus, *Designing Wide Area Networks and Internetworks: A Practical Guide*, Addison Wesley Longman, 1999. Chapter 14, which describes peering and transit, is available at http://www.genuity.com/about/executives/marcus_bio.htm

Huston, “Interconnection, Peering and Settlements”, http://www.isoc.org/inet99/proceedings/1e/1e_1.htm

Laffont and Tirole, *Competition in Telecommunications*, MIT Press, 2000. General analysis of the economics of telecommunications pricing.

Kende, “The Digital Handshake: Connecting Internet Backbones,” FCC, 2000.

Gao, “On Inferring Autonomous System Relationships in the Internet, IEEE GlobeCom 2000, San Francisco, November, 2000.

Milgrom, Mitchell, and Srinagesh, “Competitive Effects of Internet Peering Policies”, presented at the American Economics Association, January, 2001.

Laffont, Marcus, Rey and Tirole, “Internet Interconnection”, *American Economics Review*, May, 2001.