ATTACHMENT 2

Level 3 Proposal for Equitable Sharing of Backbone Costs between Peering Networks

A. Guiding Principles

Level 3 believes that there are four guiding principles for interconnection between two networks of comparable scale and reach:

1. **Equitable sharing of backbone cost burden**. Irrespective of the direction or type of traffic flows, the networks on average should bear a reasonably equal share of backbone costs in carrying traffic exchanged over interconnection points.

2. **Each network is responsible for its own first/last mile.** Neither party should use a peering agreement to generate fees to subsidize the local access portions of its network.

3. **Focus on quality service for each network's customers.** The agreement must ensure that both parties' customers receive high-quality services.

4. **Resilient, scalable and secure interconnection architecture.** The arrangement should include provisions that can quickly account for growth or changes in traffic, should be architected to minimize the impact of disruptions catastrophic events, and assure that customers' traffic remains secure.

B. Achieving the Guiding Principles

1. How to measure backbone cost burden

In order to achieve the first principle outlined above – equitable sharing of backbone cost burden – we must identify a reasonable measure of backbone costs. We will begin by looking at an illustrative traffic flow between a Level 3 content customer (e.g., an entertainment website) and an ISP's customer (e.g., a residential broadband subscriber), under the "hot potato" routing protocol that was the historical standard for peered Internet traffic (under "hot potato" routing each packet is delivered to the interconnection point that is closest to its source).

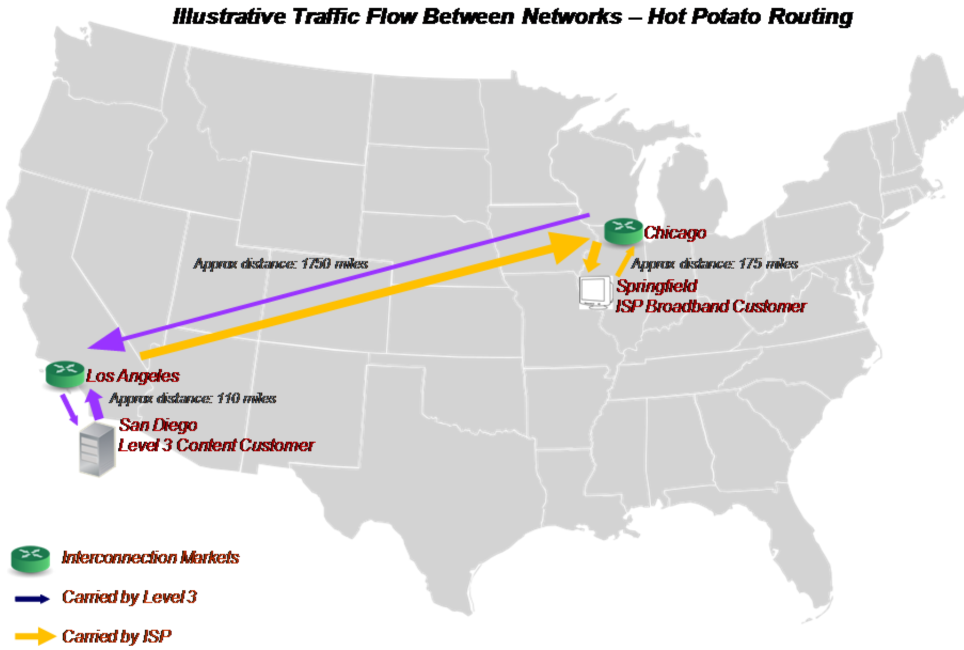

In this example, a residential broadband subscriber in Springfield requests a content file (e.g., a webpage, photo, or video clip) that is located at a content provider in San Diego. Level 3 and the residential ISP are peered in both Chicago and Los Angeles. As shown by the illustration above, using hot potato routing, the residential ISP hands off the request for content to Level 3 at the peering interconnection in Chicago, and Level 3 carries the request for content on its backbone between Chicago and Los Angeles. Level 3 then retrieves the requested content file from its customer in San Diego, carries that file to the peering interconnection in Los Angeles where it is handed off to the residential ISP, and the residential ISP carries the content on its backbone between Los Angeles and Chicago. Each network carries both the request for content and the content itself across the "middle mile" and "last mile" between the interconnection point and the specific customer location.

In this case, because the content contains more data than the request for that content, it is clear that the ISP would bear the larger backbone cost burden – the ISP would carry more traffic over the long distance between Los Angeles and Springfield than would Level 3 over the long distance between Chicago and San Diego.

To quantify the backbone cost burden, we should calculate and compare the number of **bit miles** carried on each of the networks. In this case, assume that the request for content represents 1 KB of traffic and the content itself represents 10 KB.

Type of traffic	City pair	Distance	Traffic	KB-miles	Carried by
Request for content	Springfield to Chicago	~175 mi	1 KB	175	ISP
Request for content	Chicago to LA	~1750 mi	1 KB	1750	Level 3

Request for content	LA to San Diego	~110 mi	1 KB	110	Level 3
Content	San Diego to LA	~110 mi	10 KB	1100	Level 3
Content	LA to Chicago	~1750 mi	10 KB	17,500	ISP
Content	Chicago to Springfield	~175 mi	10 KB	1750	ISP

So, KB-miles carried on the ISP's network would equal:

175 + 17,500 + 1750 = 19,425 KB-miles

And, KB-miles carried on Level 3's network would equal:

1750 + 110 + 1100 = 2,960 KB-miles

In this illustrative case, the ISP would bear a backbone cost burden that is 6.56 times that of Level 3 (19,425 KB-miles equals about 6.56 times 2,960 KB-miles).

Changing the routing protocol used by the parties, however, can significantly change the bit mileage on each network. If, instead of "hot potato" routing we use "cold potato routing," the ISP would carry the request for content between Chicago and Los Angeles, and Level 3 would carry the content itself between Los Angeles and Chicago.

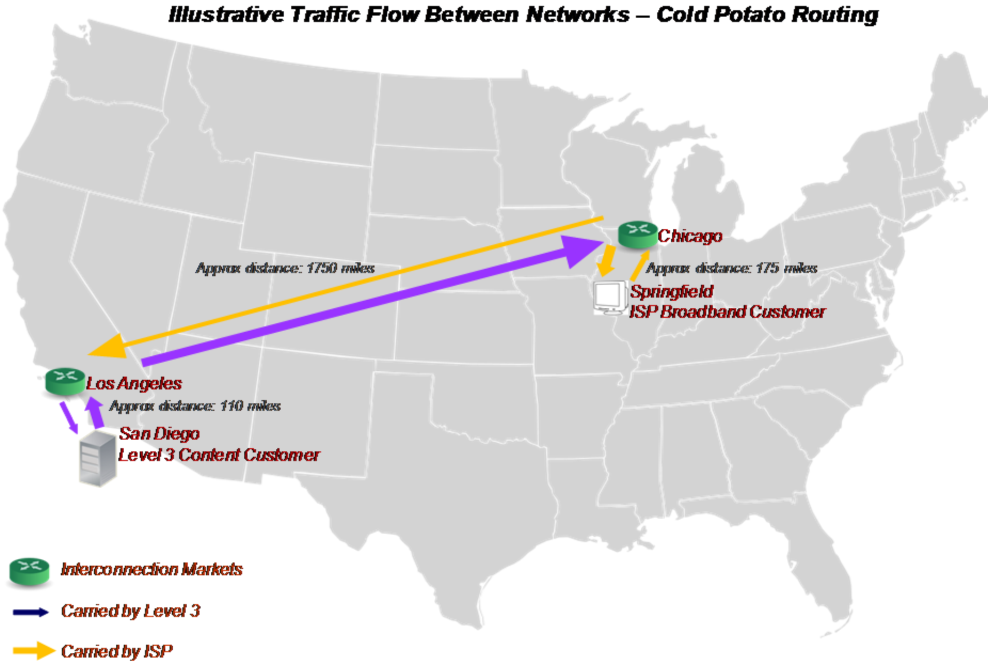

Type of traffic	City pair	Distance	Traffic	KB-miles	Carried by
Request for content	Springfield to Chicago	~175 mi	1 KB	175	ISP
Request for content	***Chicago to LA***	***~1750 mi***	***1 KB***	***1750***	***ISP***
Request for content	LA to San Diego	~110 mi	1 KB	110	Level 3
Content	San Diego to LA	~110 mi	10 KB	1100	Level 3
Content	***LA to Chicago***	***~1750 mi***	***10 KB***	***17,500***	***Level 3***
Content	Chicago to Springfield	~175 mi	10 KB	1750	ISP

Under cold-potato routing, KB-miles carried on the ISP's network would equal:

175 + 1750 + 1750 = 3,675 KB-miles

And, KB-miles carried on Level 3's network would equal:

110 + 1100 + 17,500 = 18,710 KB-miles

In the cold potato case, Level 3 would bear a backbone cost burden that is 5 times that of the ISP.

This type of bit mile analysis should be used to evaluate the total backbone cost burden borne by each network with respect to the traffic exchanged between the two networks. Bit miles are a much more relevant and manageable measure of backbone cost than send/receive ratios. Send/receive ratios vary for many reasons, and events outside of either party's control (new application development, customer and subscriber behavior, customer wins and losses by each network) can and do influence send/receive ratios greatly. Most importantly, *direction* of traffic does not drive backbone cost; if a network carries a bit on its backbone from Chicago to Los Angeles, the backbone costs are equivalent whether that bit was received from a peer in Chicago and delivered to a customer in Los Angeles or received from a customer in Chicago and delivered to the peer in Los Angeles. As the previous analysis demonstrates, changes to routing can substantially shift backbone cost burden between two networks independent of any changes to send/receive ratios.

Simple calculations using approximate average distances enable us to estimate total bit miles on each network under various traffic and routing scenarios. If the bit mileage of one party is significantly lower than the bit mileage of the other party in a peering relationship, some adjustment to routing practices or peering interconnection locations should be implemented to more equitably allocate backbone costs.

2. Estimating current and future backbone cost burden

Historically, peered networks have shouldered roughly equivalent backbone cost burdens. This may be seen in analysis of illustrative data representative of historical traffic patterns.

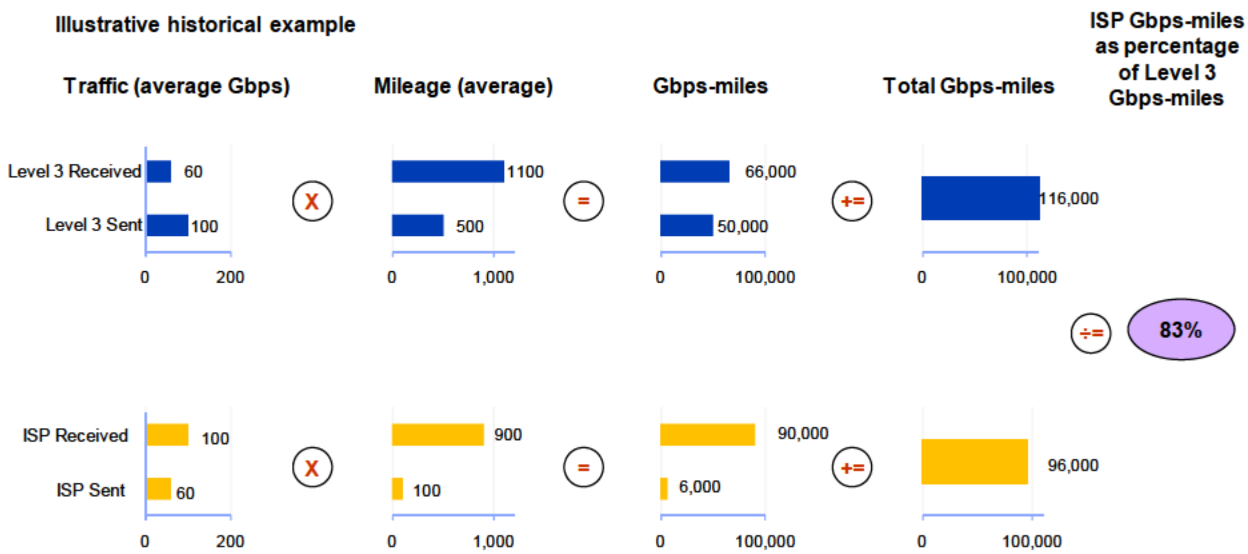

In this traffic flow, the residential ISP's bit mileage is 83% of Level 3's bit mileage for traffic delivered over the peering interconnection points. But historical traffic flows are changing – traffic is increasing rapidly and traffic patterns are changing as well. To account for this, we examined both recent traffic (similar to early 2011) and illustrative future traffic scenarios (see Appendix 1 for traffic scenario detail), and calculated bit mile equity under various routing scenarios. The routing scenarios include differentiated mileage assumptions for base traffic (mostly transit, like that in the illustrative traffic flow in section B1 above) and for CDN traffic, which is originated within Level 3's network and has a much greater degree of "localization," meaning it travels fewer miles in total (see Appendix 2 for mileage assumption detail).

If hot potato routing continues to be used, as traffic grows on the Level 3 network, a residential ISP would bear an increasing share of the backbone cost burden, as seen in the calculations below.

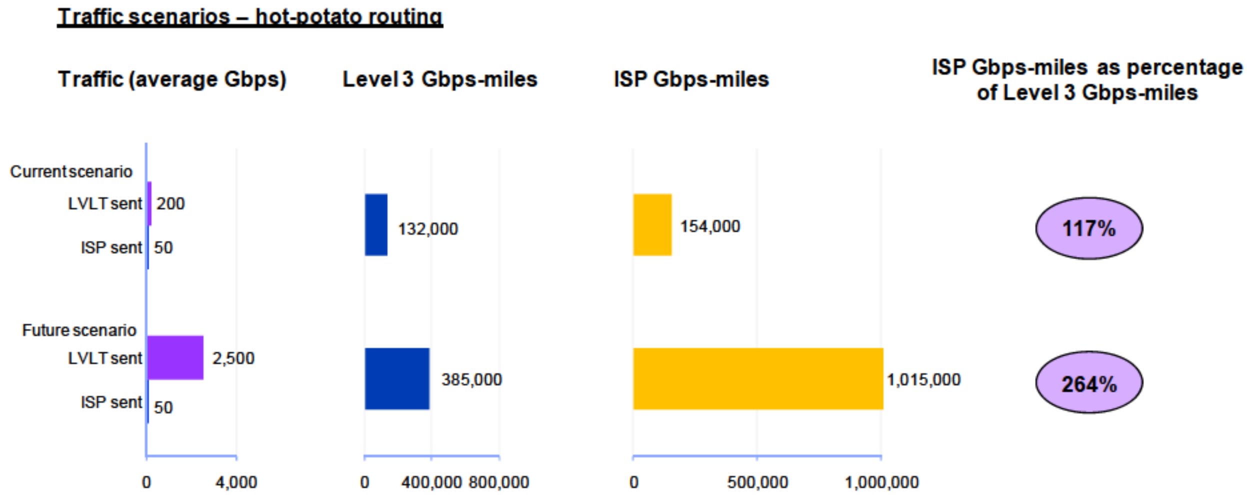

However, as in our example in section B1, cold potato routing between the two networks shifts the vast majority of the backbone cost burden onto Level 3.

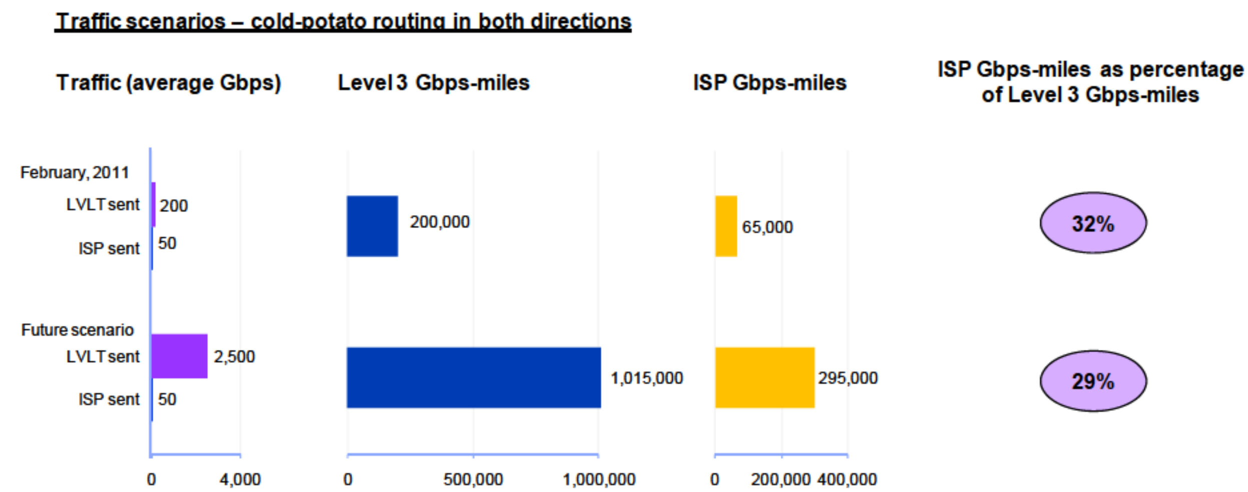

This analysis clearly demonstrates that equitable backbone cost burden can be achieved through simple routing changes alone, and that moving to full cold potato routing would (with these assumptions) disproportionately burden Level 3. Even if send/receive ratios hit 50 (sent by Level 3) to 1 (received by Level 3), Level 3 would (using cold potato routing) bear most of the backbone cost burden. This again shows that send/receive ratio is not the proxy for the relative backbone costs borne by each network.

C. Putting Equitable Backbone Cost Sharing into Practice

As indicated by the preceding analysis, in most cases equitable backbone cost sharing can be achieved through routing changes and without settlement, irrespective of asymmetry in the "send/receive ratio" for traffic between two networks. Therefore, traffic exchange agreements could achieve that guiding principle by incorporating two main elements:

a. Each network should cold potato route the traffic that originates from its customers to the interconnection point closest to the end destination to the extent required to achieve a reasonable equity in bit mileage between the two networks. A network that, over periods of time, is carrying less than 70% of the bit mileage of the other network should implement changes to routing to more equitably distribute bit mileage.

b. Traffic exchange would be settlement free as long as bit mileage is reasonably equitable.

As demonstrated in the analysis in section B2, full cold potato routing in both directions between the two networks does not equalize the bit mile balance between the two networks – it shifts the majority of the bit mile burden to Level 3, even under increasingly asymmetric send/receive traffic ratios. Were an unlikely future case to arise wherein the two networks were not able to equalize the bit mile balance with cold potato routing, the party with fewer bit miles could offer additional levers to further shift bit miles to its network, including one-way cold potato routing, deep interconnection (where the two networks would add more interconnection points closer to the party with more bit miles' customers, reducing the party's "middle mile" mileage), and, in the case where the party with fewer bit miles offers CDN services, deep caching (where the party would place CDN clusters within the other's network, enabling some content traffic to be originated closer to customers and therefore reducing "middle mile" mileage).

D. **Conclusion**

Interconnection between networks within the Internet is essential to providing quality customer service. Fair and equitable peering interconnection requires that each party bear a fair and equitable share of backbone costs. While "send/receive" ratios in the past were used as a proxy for cost allocation, shifting traffic patterns within the Internet compel – and innovative network intelligence make possible – a more precise measure of equity: bit mileage.

Appendix 1: Traffic Scenarios

All traffic in average Gbps	**Illustrative current scenario**		**Illustrative future scenario**	
Traffic type	**Level 3 sent/ ISP received**	**Level 3 received/ ISP sent**	**Level 3 sent/ ISP received**	**Level 3 received/ ISP sent**
Base traffic (mostly transit)	150	50	500	50
Incremental CDN	50	n/a	2000	n/a
Total	200	50	2500	50

Appendix 2: Mileage Assumptions

Hot-potato routing				
Type of traffic	**Network**	**Mileage for traffic sent by LVLT to ISP**	**Mileage for traffic sent by ISP to LVLT**	**Data source**
Base (primarily transit)	On LVLT network	500	1,100	Level 3 assumption based on actual mileage to certain peers
	On ISP network	900	100	Level 3 assumption based on understanding of certain peers' mileage
Incremental CDN traffic	On LVLT network	40	n/a	Level 3 assumption based on actual mileage to certain peers
	On ISP network	280	n/a	Level 3 assumption, some but not full localization of CDN

Cold-potato routing				
Type of traffic	**Network**	**Mileage for traffic sent by LVLT to ISP**	**Mileage for traffic sent by ISP to LVLT**	**Data source**
Base (primarily transit)	On LVLT network	1,100	500	Level 3 assumption, based on hot-potato assumptions above
	On ISP network	100	900	Level 3 assumption, based on hot-potato assumptions above
Incremental CDN	On LVLT network	220	n/a	Level 3 assumption, some but not full localization of CDN
	On ISP network	100	n/a	Level 3 assumption, based on hot-potato assumptions above