A Blockchain-backed Internet Segment Routing WAN (SR-WAN)

Whitepaper by NOIA Network

William B. Norton wbn@noia.network, Jonas Simanavicius jonas@noia.network

Abstract— *Barely a day goes by without a news report of some Internet infrastructure failure that impacted end-users across the Internet. There is often nothing affected users can do but wait for the storm to pass, and resume when connectivity is restored.*

Segment Routing has emerged from the IETF as a source routing RFC that can force network traffic to, and then through, a pre-defined sequence of relays called segment routers. Each segment router replaces the destination address with the address of the next segment to visit on towards the destination. This enables end-systems to force traffic along a network path that bypasses areas of network degradation.

This paper is divided into two halves. The first half presents a model for an inter-provider public Segment Routing WAN (SR-WAN) that provides hosts (Mac, PC, Linux) with additional paths through the Internet. This is done using segment routing across network paths shared by the others connected to the SR-WAN. Blockchain is used as the distributed ledger for providers and consumers of spare compute and bandwidth. The second half of the paper focuses on private SR-WANs, with commercially operated segment routers operating in well-connected Internet colocation centers. Both the public and the private SR-WANs share a common distributed database of segments available for use or purchase.

1. Internet Variability

The public Internet is composed of autonomous Internet Service Providers (ISPs), interconnected in arms-length distance relationships called "Internet Transit" and "Internet Peering." Both of these relationships provide connectivity to the public Internet, or portions of it, respectively. Importantly, packets are exchanged between these networks without consideration of the *quality* of the network paths enabled. As a result, routers will blindly forward traffic across lossy congested links.

Beyond occasional packet loss the public Internet is increasingly experiences latency fluctuations. To illustrate, one of the authors conducted inter-cloud measurements between 25 cloud instances across the AWS, Google, and Azure clouds in 2017. This study quantified the variability in packet delivery between the clouds, spotlighting latencies that deviated from the median latency measurement by more than 20%[1]. These types of latency anomalies appeared almost every day (see graph in Figure 1 for an example).

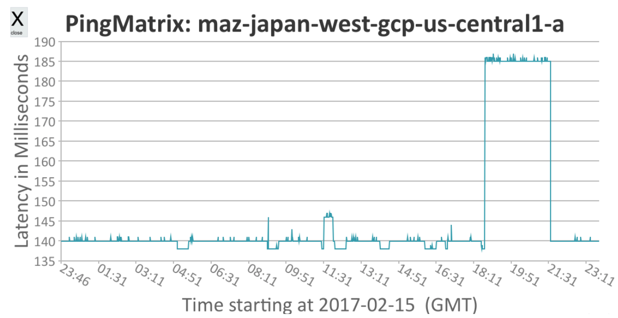

Figure 1 - Example of latency anomalies common on the Internet

Research with the cloud network operators identified the cause of this variability to be a MPLS traffic engineering technique called "auto-optimization." This is a router configuration instructing the MPLS underlayment to balance the router load evenly across the underlying transport links. When end-system traffic traverses these paths the latency can vary by up to 100ms for minutes or hours. This variability of course is detrimental to latency-sensitive applications, but also leads to variability in connection time and effective thruput.

Internet performance variability also comes from the Internet inter-domain routing system. The Internet Society reported 14,000 routing outages, leaks and hijacks in 2017. Most recently (June 2019), Allegheny Technologies Inc., a metals manufacturer, misconfigured its router leading to widespread partial and complete Internet connectivity failures.

All three of these examples led to Internet end-users routinely experience periods of poor performance for minutes or hours. When connectivity is lost or crippled, there is no general recourse but to complain to your provider and wait for the connectivity to recover, accepting that the public Internet has periods of intermittent network degradation.

Until now.

2. A Public Segment Routing System

A public segment routing system brings the wisdom of crowds and a technology called "Segment Routing" to the public Internet. The Segment Routing WAN (SR-WAN) continually measures and enumerates a list of alternative

[1] https://www.linkedin.com/pulse/internet-service-provided-cloud-premium-william-b-norton/

network paths through the public Internet using segment routers operated by the SR-WAN community of users. This “sharing economy” model enables connectivity for partially isolated users to leverage the better connectivity of other users. Network traffic is relayed through a path of participating segment routers (called a “Segment Router Path,” or a “SRPath”)

The block chain provides accounting for the community, rewarding utility tokens to those who relay the traffic for others, which can be applied when connectivity is problematic for the segment router operator.

2.1 Overview

The nodes of the SR-WAN operate segment routers. In the public SR-WAN, the segment router is software (open-source virtual router that runs in a container on your host.)

Each segment router packet contains the addresses of segment routers to visit before delivering the packet to the final destination. This source-routing system ”routes around” congested paths as shown in Figure 2.

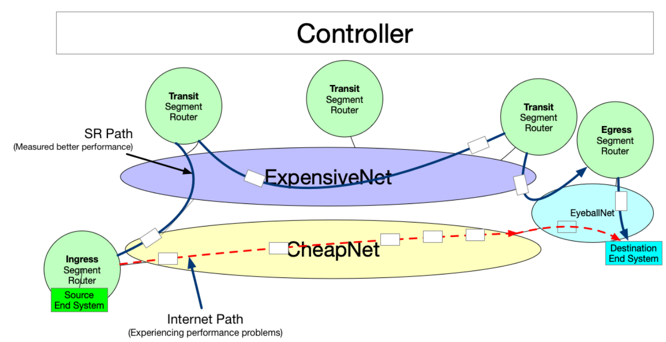

Figure 2 – Routing Internet traffic around a congested path

Much like MPLS is called a network *underlayment* acting as “Layer 2.5,” the SR-WAN can be thought of as a network *veneer* acting as “Layer 3.5.” Both plug into routers, and apply hardware to optimize either utilization of links or utilization of transit bandwidth respectively.

2.2 Components of the Public SR-WAN System

There are five main elements to the public SR-WAN system:

2.2.1 Ingress/Egress Segment Router

Traffic enters the segment routing system via an “Ingress Segment Router” and sent out to the final destination via an “Egress Segment Router.” For the purposes of this part of the paper, a “segment router” refers collectively to the segment router and support systems running in a docker on your host (PC,Mac,Linux).

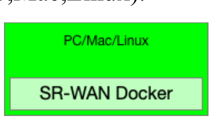

Figure 3 - Public SR-WAN Segment Router on the host

The docker provides an isolated environment for the segment routing software and supporting systems.

2.2.2 The SR-WAN Controller

The SR-WAN controller kick starts off the segment routing system with the assignment of segment routers into **segment router groups**.

Segment router groups scale to up to 100. One segment router in each group periodically communicates the group’s collective network state to the controller.

The controller converts the performance matrices into a database of available network segments and their recently measured network performance characteristics. This authoritative decentralized database is called the Distributed Transit Exchange (DITEX).

2.2.3 Segment Router Pulse

All segment routers are assigned to groups that “pulse” one another with their state information. This accomplishes two things:

First, it provides the payload to measure **one-way latency** between all segment routers in the group.

Secondly, each segment router can then construct a **full matrix of one-way latency measurements** which can be used to determine if there is a better path towards the destination through the segment router system.

Note that round trip measurements (i.e. ping) are too inaccurate as they conflate the performance of two distinct network paths, and are therefore not an accurate reflection of the performance of packets directed one-way through the segment routing system.

2.2.3 First and Last Hops

To match the destination addresses to nearby segment routers for ingress and egress, the segment routers and destination addresses are tagged to a geolocation:ASN tuple. This enables an Internet packet to be directed to an appropriate egress segment router when the SR-WAN provides a better performing path.

When an egress router receives the segment routing packet, it forwards the original packet on to the destination. During each transfer, each segment router calculates the one-way latency and thruput observed through the segment routing path. This path is called an **“extraordinary path”** and becomes part of the state information propagated across the group. This also allows the system to autonomously compare the calculated latency expected against the actual latency.

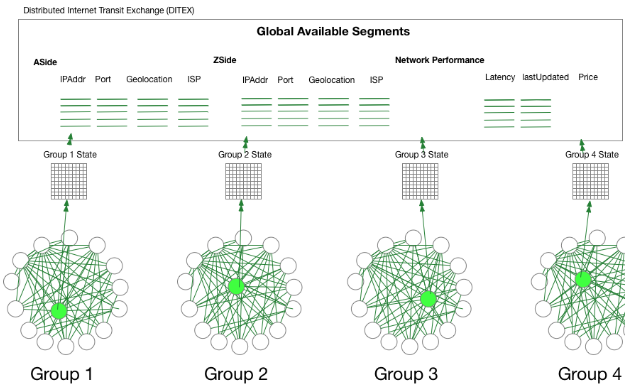

Figure 4 - The SR-WAN Controller assimilates segment performance measurements

The **DITEX** holds the SR-WAN segment state, and supports several uses. First, the DITEX provides data for group assignment. Secondly, the DITEX provides some network accounting data for validators and for smart contracts. Thirdly, the DITEX holds the data for path performance validation to identify "bad actor" and "poor performer" segment routers and quarantine them into groups. Fourth, the DITEX provides the global perspective required to optimize inter-group traffic across the global SR-WAN ecosystem.

2.2.5 Blockchain

Blockchain is the distributed ledger technology used to settle between suppliers and consumers of the collective resources of the SR-WAN. Utility tokens (aka "NOIA coins") are *earned* for traffic relayed for others, and *utilized* when they themselves wish to utilize the SR-WAN (i.e. its default Internet path is congested.)

Network data is continuously collected by the segment router and presented to the controller in the form of a 'ticket' along with group statistics. Validators verify the network statistics that accompany the claim for reward, and add the transaction to the blockchain so coins are transferred to the NOIA wallet of the bandwidth supplier as shown in Figure 5.

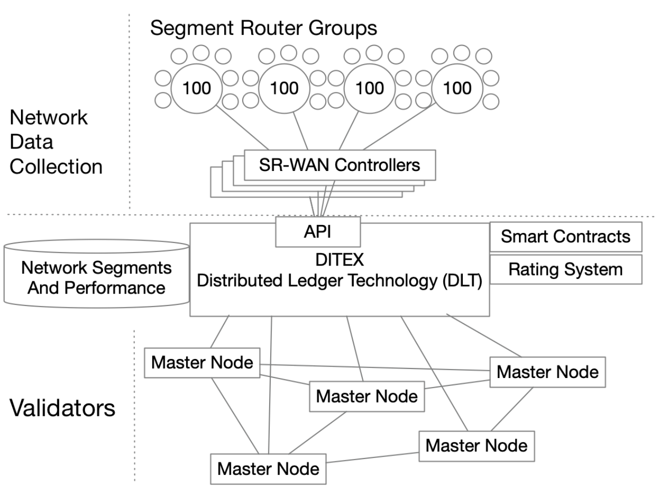

Figure 5 - Blockchain and Segment Routing System

The validators (Master Nodes) earn tokens for validating transactions in the DITEX.

There are two sources of tokens for rewarding validators. First, every transaction on the ledger generates a small transaction fee. This transaction fee pool is also subsidized systematically for the first three years. Transaction fees are allocated to a validator pool (25%) and a dividend pool (75%).

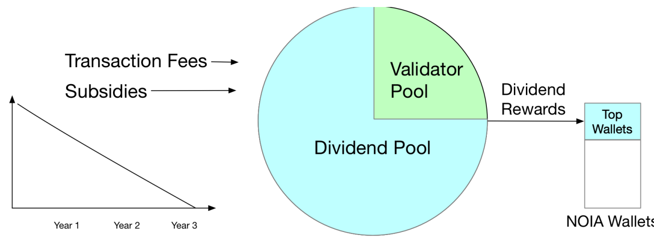

Figure 6 - Transaction Fees allocation

The dividend pool is rewarded to the top voted NOIA wallets. This approach provides an incentive for large token holders to get and share dividend yields, while maintaining a healthy token velocity. Coin holders can "vote" their coins with others to form groups to increase their chances of being part of one of the reward wallets and receive a portion of the dividend pool. These "votes" have no impact on the holders' wallet, but rather provide a means for smaller wallets to participate in the dividend pool reward.

For reference please see the NOIA Paper: "Economics of Decentralized Internet Transit Exchange: Utilization of Transit Capacity."

2.3 The Public SR-WAN Supports the Sharing Economy

There are many concepts, such as *collaborative consumption peer-to-peer markets* (i.e. BitTorrent), and *access-based consumption* (i.e. airBnB), which can all be put under the umbrella term of ***sharing economy***. The public SR-WAN brings forward a new example of a sharing economy that satisfies all seven criteria enumerated by Hawlitchek .et. al.[2]:

1. **Increasing utilization rates.** The Internet has spare capacity.
2. **Peer-to-peer principle.** Transactions are between buyer and seller.
3. **Existence of reimbursement.** Coins are rewarded for use of resources.
4. **No transfer of ownership.** The resource is used and then is available for others.
5. **Resource tangibility.** To participate, one shares part of their computer. The location of that computer makes it more or less attractive to the community of users.
6. **Leveraging of information systems.** Machine learning systems are applied to the data to optimize

[2] Florian Hawlitschek, Benedikt Notheisen, and Timm Teubner. The limits of trust-free systems: A literature review on blockchain technology and trust in the sharing economy. Electronic commerce research and applications, 29:50–63, 2018.

the routing system and settlement is done via modern block chain.

7. **Temporariness.** Access to infrastructure is granted for the short period of time.

The public SR-WAN demonstrates all of these characteristics and so meets the definition of a sharing economy. Network segments provided by a segment router are the commodities with intrinsic value, the seller grants access and the buyer consumes the resource, after which the buyer and seller settle with blockchain.

3. A Day in the Life of a SR-WAN Segment Router

To illustrate the system, let's follow the money for a 1GB use.

Assume each NOIA coin has a nominal price of $0.05 USD. Further assume that the cost to relay traffic through the public SR-WAN is $0.50 per GB. Each segment router earns a share of the $0.50, proportionate to the number of relays in the path[3] as shown in Figure 7.

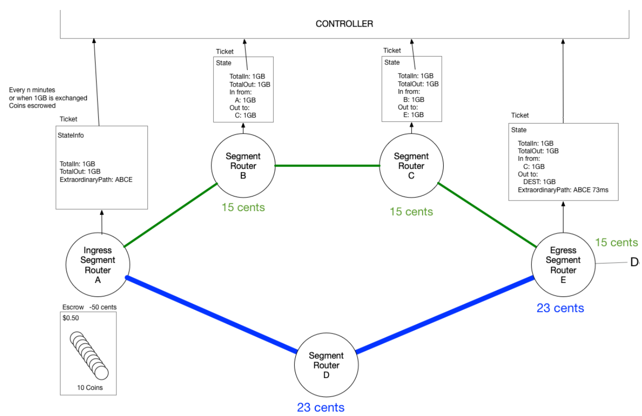

Figure 7 - Segment Routers earn tokens to the proportion of their participation in the path

After the software is downloaded and installed in the docker, there are six interactions.

3.1. Initialization

The segment router connects to the controller to receive its configuration (including geolocation:ISP tuple) and its group number. Coins are deposited in escrow for use when traffic is sent to the SR-WAN.

3.2. Pulse

Each segment router periodically (every 10 seconds for example) pulses their state information to all others in the group, so every member of the group knows each other's geolocation, ISP and one-way latency for traffic sent to them. This full matrix enables each segment router to calculate aggregate path latency that can be compared against the default Internet path latency.

[3] The allocation formula is a bit more complicated than this; each node earns proportionate to how many segment routers are in the path modulo a piece allocated to the operation of the shared infrastructure.

3.3. Detection

Most of the time, the system operates in pass-through mode, forwarding traffic on to the default gateway. The system looks up the geolocation:ISP tuple to find the best egress segment router for target destinations. A list of alternative routes is calculated and compared to default Internet routing performance.

3.4. Reroute through the Public SR-WAN

When congestion is detected in the form of packet loss or latency variance, etc., a segment router packet is constructed with the original packet as payload, and with the segment routers to visit listed in the segment router packet.

This segment router packet is sent to the first segment router in the path. Every router along the path forwards to the next segment router hop. Along the path, each segment router counts the amount of traffic sent and received by each other segment router.

3.5. Delivery of original packet to the destination

When the egress router receives the packet, it unpacks the original packet and sends it onto the destination. *Note that this packet originated elsewhere in the Internet, so may be filtered by ISPs implementing BCP-38. A variety of techniques can be used to address this.*

3.5. Blockchain settlement

After relaying 1GB, each segment router sends a claim 'ticket' to the controller with its state (including counter values like bootTime, totalTrafficIn, totalTrafficOut, and InOctets/OutOctets from each segment router). Validators review the group accounting data and transfer the coins from consumer's escrow to the provider's NOIA wallet.

The egress router also stores the performance statistics of these "exceptional paths" as part of its state.

Each segment router provides a portal (Web page) to provide feedback to the operator such as coins earned, coins applied, and traffic relayed. The dashboard also shows a live latency matrix of the SR-WAN from the perspective of the operator segment router.

For most users, the SR-WAN is a set it and forget it docker appliance. By participation they earn the right to use the segment router paths of others. NOIA Coins/Tokens ensure fair sharing.

4. SR-WAN Incentives

This public SR-WAN provides three powerful incentives.

1. It encourage operators to leave their segment routers running all the time to garner more coin,
2. It supports good networks. Segment routers on good networks attract more packets to relay and therefore earn more coin,
3. Its value grows as the network grows. As the network grows, the number and diversity of quality paths increases, and therefore increases the value derived by the community.

5. SR-WAN Leverages Market Forces

The SR-WAN leverages several current market forces:

5.1. Universality of fixed-price Internet Transit

Internet Transit provides connectivity to *any* endpoint on the Internet. The SR-WAN leverages this fact to reach the segment routers on other Internet transit services.

Since these services imbue customers with a fixed price for all the Internet traffic they can consume or produce (up to a limit), there is no additional cost to participate in the SR-WAN. Participation in the SR-WAN utilizes spare compute and bandwidth. And since the price of Internet Transit historically drops every year from about 20-30%[4], the cost of participation in the SR-WAN is expected to drop every year.

5.2 Acceptance of open-source distributed systems software

The world is shifting to open-source. In particular, the popularity and acceptance of open-source software today facilitates the acceptance of distributed systems such as this. Containers provide isolation between the client environment and the segment router. Frameworks like NodeJS ease the development of distributed systems, Express externalizes instrumentation as web pages and RESTful APIs. All of these open-source components are widely accepted and deployed today, providing the project with production-ready off-the-shelf building blocks to build a cooperative networked system.

5.3 Desire to participate in Crypto Currency

The rise in Bitcoin value to over $13,000[5] in 2019 has propelled cryptocurrency into the limelight and fueled speculative investment into digital assets. This SR-WAN automatically rewards NOIA coins for participation in the network, so provides a free and easy way to participate in the cryptocurrency market system. One can simply run some code and obtain crypto currency in return. While not the intent of the system, we acknowledge this strong allure.

5.4. Machine Learning is now accessible and applicable

The SR-WAN generates real-time network segment performance data that enables machine-learning systems to identify patterns in the data for use in dynamic routing.

Machine learning is utilized to automatically identify trends and optimally group to shift collectives of traffic along better segment-routed paths dynamically.

5.5 Emergence of Segment Routing from the IETF

About a dozen draft and full standards surrounding segment routing have emerged from the IETF, driven largely by Cisco Systems. Cisco enables routers to execute segment routing in hardware. While the public SR-WAN virtual router only leverages a small subset of features documented in the specifications, the marketing and technical support from major hardware vendors improves the chances of broader adoption of segment routing systems such as this.

5.6 The First Mover Advantage

The SR-WAN demonstrates network externality properties; the value of the SR-WAN is proportional to the number of nodes in the network and the quality of the paths they bring. Therefore, late entrants into the market will have to compete for users against an established and growing network that is generating coins of increasing value for them. This provides a sustainable competitive advantage for the first SR-WAN with a critical mass of segment routers.

8. SR-WAN Use Cases

There are several use cases envisioned for the SR-WAN.

8.1 Route Around Congestion

This is the normative case. On the public Internet, routers continue to send traffic along a congested path shown in red in Figure 8.

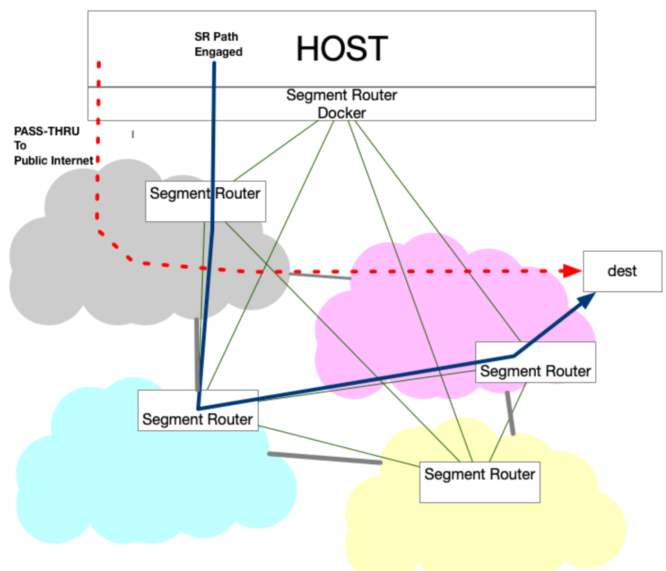

Figure 8 - Detection and bypassing congested network paths

As described earlier, the SR-WAN paths are constantly being compared and the SR-WAN path is used when better. When the regular Internet path returns to normalcy the default Internet path is once again preferred.

8.2 Buy cheap network bandwidth but use a better network

Since the Internet transit service provides access to <u>any</u> attached Internet device, there is nothing stopping one from simply forcing all of their traffic to a segment router sitting on an empirically better network. This traffic will be received by the segment router on the premium network and

[4] https://drpeering.net/FAQ/What-are-the-historical-transit-pricing-trends.php

[5] https://www.google.com/search?q=price+of+bitcoin&rlz=1C5CHFA_enUS763US763&oq=price+of+bitcoin&aqs=chrome..69i57j0l5.2711j1j8&sourceid=chrome&ie=UTF-8

forwarded onto its destination across the premium network path.

8.3 A Cost Effective Intranet

A cost-competitive alternative to MPLS can be constructed for small and medium-sized businesses over other peoples segment routers. This could deliver encrypted network segments between offices. Since Internet transit is always dropping in price, one can construct a private network solution at a fraction of the cost of current MPLS solutions, and the costs are expected to drop 20%-30% per year from there!

9. Improve Security-Control and Visibility

The future of cyber security is *adaptive* and *responsive*, something the SR-WAN is particularly well suited to do. State-sponsored cyber attacks leverage automated attack systems, so detection, mitigation and data collection systems must be automated as well. The SR-WAN system provides greater control and visibility than sending traffic over the public Internet alone.

Control. The SR-WAN is driven by a constant flow of performance data fed into an AI/machine learning system. This system dynamically re-*groups* segment routers based empirically on network performance reported by the receiving routers. This enables the re-routing of traffic dynamically based on current measurements and training models. This data collection and feedback mechanism enables the controller to adjust groupings based on measured one-way latency between ingress and egress systems.

Visibility. Each segment router provides the operator with visibility into the segment routing paths available. The segment routing instrumentation is a constantly updated web page run from within the segment router docker. Here the system shows uptime, coins earned, the traffic relayed, and a matrix of real-time performance data.

The matrix provides users with current network state as seen by the segment router, highlighting paths where improved traffic performance brings opportunities for better connectivity.

10. DITEX Premium Bandwidth Use Cases

"The most powerful asset in the digital age is data[6]." – Don Tascott

As introduced earlier, the Distributed Internet Transit Exchange (DITEX) holds the repository of available network segments and their performance characteristics, pre-populated from the segments available on the public SR-WAN.

Additional segments and paths can be listed on the DITEX with coin premiums associated with them. For example, commercial operators can sell premium segments with automated smart contracts. The blockchain is used for settlement between the buyer and seller, with a micro transaction fee allocated to the DITEX operations.

10.1 Use case: Monetize spare capacity

To illustrate a DITEX premium bandwidth use case, consider the Spread Network link between Chicago and New York. Spread Networks spent a lot of money to tunnel through the mountains to build a fiber path with the lowest possible latency. This exceptionally low latency path is expensive to lease, but is critical for automated trading during business hours. The demand curve for this circuit is shown in the graph, highlighting six hours of unused capacity before and after the workday.

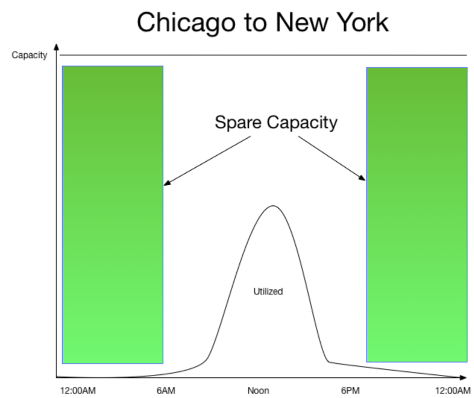

Figure 9 - Off-hours bandwidth available and monetized as network segments

This off-hours capacity can be monetized by simply running segment routers on either side of the link and listing them on the DITEX. With its empirically measured lower latency, this link will attract traffic, and its operator will receive coins for its use during those off-hours. The listing of the bandwidth can be anonymous and can be time-delineated, so there is no market cannibalization or impact on most-favored-nation-termed contracts.

10.2 Use case: Network-as-a-Service

Setting up a traditional network is both time and capital intensive. Shifting from networking *physical* routers to networking *virtual* routers only gets you half way there. The DITEX provides the missing piece: the directory of available resources that can be stitched together to create virtual network required. Since this virtual network rides on top of the Internet, it has maximal reach. Since it rides on the Internet it also inherits the path diversity across many networks.

10.3 Use case: Purchase Better Internet On-Demand

Some enterprises may want to leverage the path diversity during times of crisis, but not forward segment routing

[6] "How the blockchain is changing money and business," TEDSummit June 2016

traffic on behalf of others. The DITEX enables the purchase of coins for use if and when the segment routing path is used. This mode of operation is implemented by setting the price of the enterprise's segment links to infinity in the DITEX. This will in effect make these segments unavailable for others to use but still allow it to participate in the latency measurements so it knows when the SR-WAN provides an empirically better path. This configuration will generate no coins for the enterprise, but still allow the enterprise to participate in the group pulse and engage better paths with purchased coins when desired.

11. A Private Segment Routing WAN

The Private Segment Routing WAN (Private SR-WAN) leverages the same technology as the public SR-WAN but utilizes professionally operated segment routers in native IPv6 Segment Routers (SRv6), collocated in well-populated Internet colocation centers. The operators of the infrastructure purchase Internet Transit from ISPs that have peering with the regional eyeball networks to produce the lowest latency paths.

While these deployments are expensive, the performance characteristics are perfect for emerging latency sensitive and business critical edge applications. A single regional deployment can provide direct reach to 80% of the market within 25ms. This brings premium performance segments into the DITEX.

Deployments will match the eight interconnection regions across the U.S.: Seattle, Bay Area, Los Angeles, Chicago, Dallas, Ashburn, Newark, and Atlanta as shown in Figure 10.

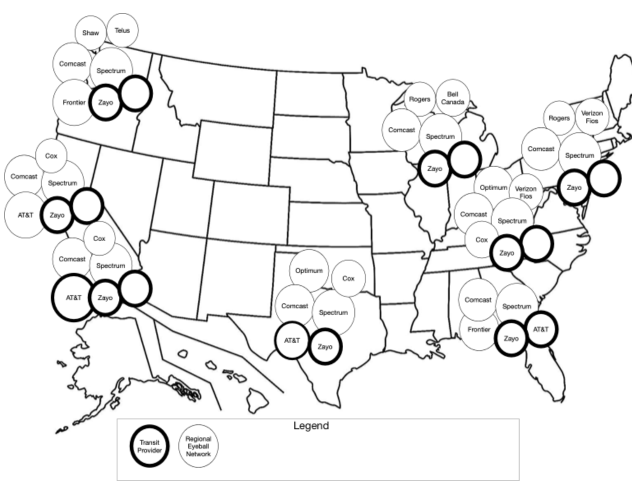

Figure 10 - Proposed US deployments (reach 80% of market within 25ms)

In Europe, to reach 80% of a region within 25 milliseconds one needs good connectivity to 5-8 eyeball networks. Here again, transit from well-peered ISPs can be used to reach these regional eyeballs in London, Amsterdam, Frankfurt, and Paris.

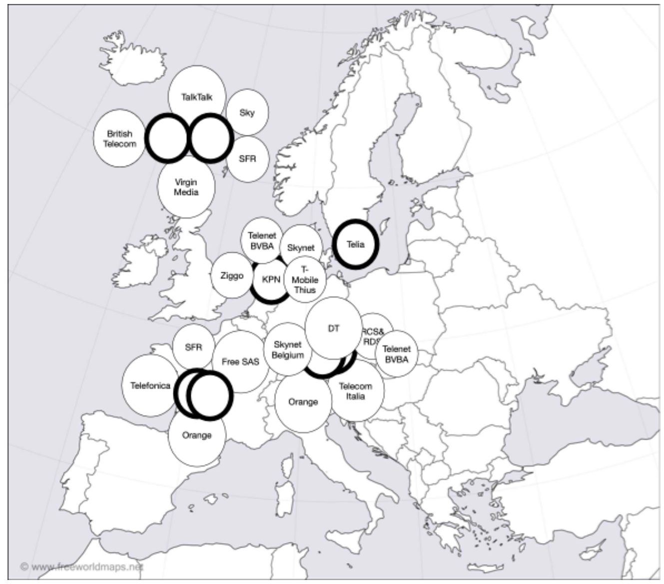

Figure 11 - Proposed European Deployments (reach 80% of market within 25ms)

Each deployment also leverages market competition for transit. Historically transit is 10-20% cheaper in a popular colocation center as compared with a corporate data center. Since the colocation centers house the transit providers core routers, the segment router operator is really only depending on and paying for its Internet traffic to travel across a single router backplane.

Each professional deployment provides multi-homed segment routers. This brings a richer selection of paths than the average public SR-WAN has access to. Peering at the dominant IXP enables peering for additional direct reach.

It is important to note that the cost of Internet transit services is borne by the segment router operator. There are usually volume discounts and minimum traffic volume commitments with Internet Transit purchases, so the professional operator will have spare capacity if they fail to meet their traffic commitments. In this case, the operator has purchased spare capacity and is compensated in NOIA coins.

While NOIA is prepared to roll out a private SR-WAN as a proof of concept, its preference is to engage with and coordinate other with interested parties.

12. Future Use Cases

Several future use cases are interesting to consider.

MPLS Supplement or Replacement. Current network operators can supplement MPLS offerings to enterprises with a lower cost private SR-WAN over the public Internet. This solution rides on top of a plethora of diverse Internet transit paths. This solution replaces partially utilized MPLS links with more efficiently utilized Internet bandwidth.

Network Last Mile. Many network operators do not have an inventory of last mile transport to bring customers onto their network. This segment routing system could be used to construct a public Internet on-ramp onto any physical network.

Service Level Agreement Token Escrow An ISP may escrow coins for customers to use in the event of congestion or packet loss across their network. This form of SLA goes beyond credits for downtime; it pays for tokens to bypass congestion or faults in the provider's network if at all possible.

Virtual Circuits. One could construct heuristics that spread the offered load across multiple SR-Paths with the same performance characteristics. By sending the traffic in round-robin pattern, the system learns more about the path characteristics and more users are rewarded with coin. In this way, last mile capacity can be aggregated into larger capacity bandwidth chunks.

Content Distribution Systems. A CDN can leverage the SR-WAN to populate caches deep in the last mile networks. Operators of the segment routers receive coin for the traffic exchanged, so caching provides one more reason why their segment routers receive traffic.

13. Conclusion

The Segment Routing WAN brings together buyers and sellers of spare Internet connectivity using rewards for participation and relaying, and a community block chain for settlement.

The public SR-WAN involves installing free open-source software in a docker container on a host system. Coins are earned and used by the host to gain access to a better Internet.

The private SR-WAN brings professional operators running hardware SRv6 nodes at well-populated colocation environments to enterprises seeking optimized connectivity for business-critical applications.