A CDN Without POPs?

TULIX SYSTEMS

A CDN Without POPs?

Why the conventional approach is flawed when evaluating a CDN for live streaming

When CDNs first came into existence, the internet was mostly static content, with media files being delivered via progressive download. Content was being consumed at an exponentially growing pace by users around the world, and CDNs sprung up as an answer to the need for global delivery and load distribution for high volumes of traffic.

Introduction

When explaining how our CDN functions to someone for the first time, whether to a potential customer or a seasoned industry veteran, the initial reaction is frequently a puzzled stare. Those with some familiarity with how CDNs function instantly steer the conversation to points-of-presence (POPs). Inevitably, we're asked how many POPs we have and where our edge servers are located, the two metrics that seem to be the most popular means of gauging the quality of a CDN at first glance. However, our unique CDN was built on an infrastructure with no POPs at all, and though we have added and are continuing to add servers in various locations, the core of the Tulix CDN is a centralized distribution system housed in our data centers in Atlanta, Georgia.

Our lack of POPs to list in conversations is often met with incredulity until we explain why we chose to forgo the traditional CDN model in building our infrastructure and delivery network. Even if our explanation isn't quite understood, a demonstration and test of stream delivery using our CDN is always enough to persuade even the most staunch doubters that it not only works, but shows consistently higher stability and quality than POP-based alternatives. Our decision to build a centralized distribution system was not an accident, but a very deliberate and systematic effort to build a highly scalable and global CDN, optimal for live stream delivery.

A Brief History of the CDN

When CDNs first came into existence, the internet was mostly static content, with media files being delivered via progressive download. Content was being consumed at an exponentially growing pace by users around the world, and CDNs sprung up as an answer to the need for global delivery and load distribution for high volumes of traffic. Locating edge servers in different countries and regions around the world made sense, as it reduced the risk of a single server becoming overloaded with traffic and ensured that users at greater distances from a file's point of origin could reliably download said file with limited interruption and stable streams.

CDNs grew both in number and in size as the demand for content continued to increase, and the POP-based model of delivery became the established standard. As live streaming technology advanced and became a viable method of delivering video and audio, the already-established CDNs simply adapted their existing infrastructure to distribute live content. Live streaming delivery, as it turns out, is uniquely different from the delivery of downloadable content. For one, as the name implies, live streaming is entirely in real time. Live streams are ingested from the source and immediately delivered to viewers around the world. Using the traditional model, the viewer is served the stream from an edge server (the CDNs POP) nearest to him or her, which receives the stream from the point of ingest.

Centralized Distribution System: The Origins of the Tulix CDN

When live streaming was still in its infancy and leagues away from seeing widespread adoption, we anticipated how much of an impact it would have on the consumption of content and were determined to devise a scalable solution to meet a demand we knew would be global

and unprecedented in volume. In 1996, we streamed live from Centennial Park in Atlanta during the Olympics. That same year, we built the first website for Red Lobster, later introducing the LobsterCam, a live stream of the lobster tank at the company's customer service center (which was only ever interrupted when the cable of the camera was cut in half by one of the tank's temporary residents).

From then on, we committed ourselves entirely to building a CDN that would overcome the challenges associated with live streaming, both from the hardware and the networking perspectives. Tulix created an infrastructure that was highly scalable and met the demands for global content delivery. We did this all from our data centers at 55 Marietta St., a premier carrier hotel in the heart of Atlanta, without adding POPs in other locations until recently.

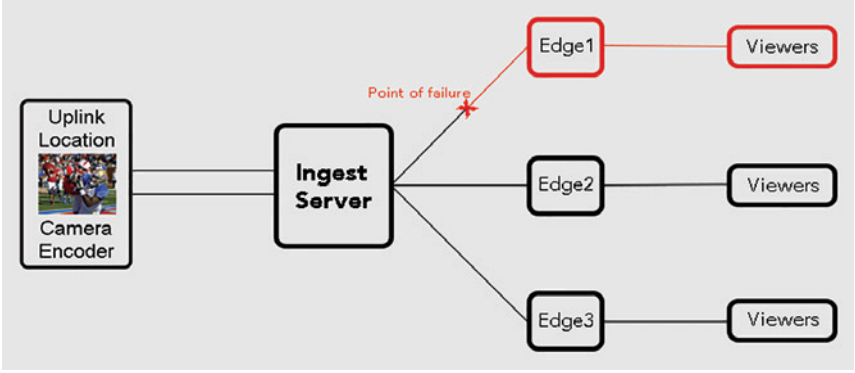

A connection problem from ingest to edge can cause widespread streaming outages.

Why Centralized Distribution?

As mentioned earlier, one of the key distinguishing elements of live streams is that they have to be delivered immediately from source to viewer. There is no caching, no progressive download, and any interruptions between the origin and destination of a live stream will be noticeable to the viewer. In the traditional model of POP-based delivery, this allows for 3 potential points of failure, at least from a networking standpoint. If any of the connections between a source and ingest, ingest and edge server, or edge server and viewer is poor, that viewer will receive a choppy stream in the best case or no stream at all in the worst.

Tulix's CDN uses two independent route-optimization systems to dynamically assess the connection of a viewer to our streaming

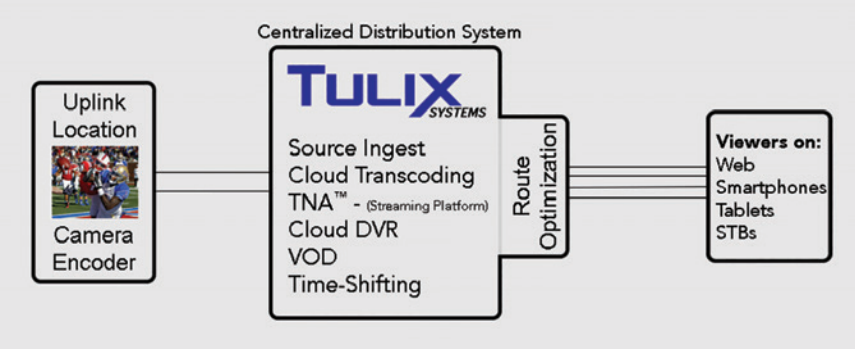

Tulix's CDN uses full redundancy and route-optimization along four of the largest backbones to guarantee stability and quality.

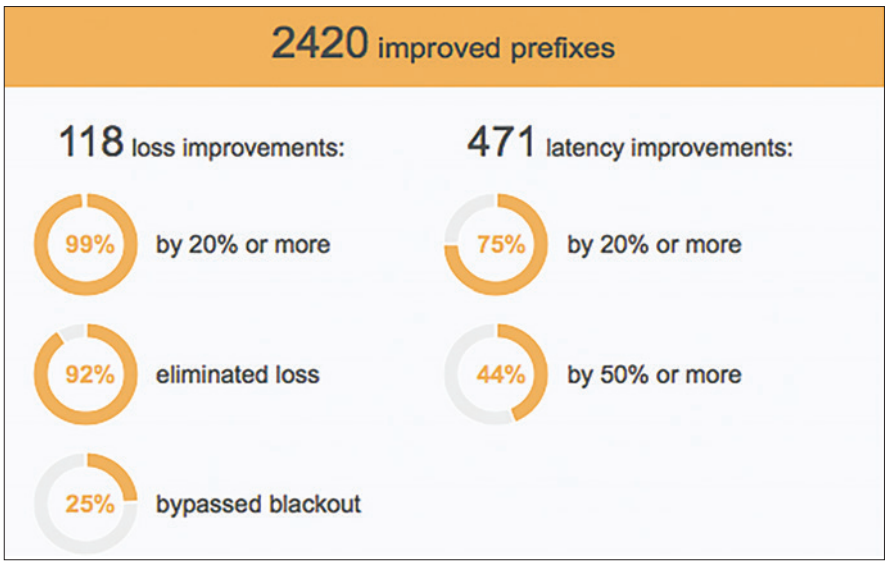

A generated report that highlights the effectiveness of route-optimization.

Most Improved Prefixes (Control)			
Prefix	Improvement (ms)	Latency (ms)	
		Old	New
62.47.0.0/16	457	593	136
101.58.0.0/18	442	545	103

Route-optimization can make the difference between an unwatchable and a smooth stream.

servers in Atlanta, and we are in the process of adding a third. The route optimizers choose between four of the largest internet backbones and Tier 1 providers to which our data centers are directly linked to determine the path to the viewer with the lowest latency and packet loss. These backbones were deliberately chosen to maximize global reach and stream stability. Being located in a carrier hotel at the largest internet exchange in the Southeastern United States means we have access to, in practical terms, unlimited bandwidth, and can add connections to the largest global networks in a matter of weeks. POP-based CDNs locate their servers in data centers they do not own, and thus have limited to no control over what networks they are connected to.

Owning our own data centers means we have complete control over our network architecture. We use 10Gbps streaming platforms and are currently upgrading to an 800Gbps slot per network infrastructure that will allow us to use 40Gbps systems in the near future. We use the fastest core routers and latest generation of supervisors available to maintain the highest level of performance. We have tested countless streaming servers and use hybrid streaming platforms that leverage the advantages of each for maximum scalability and quality. Everything in our data centers is fully redundant and optimized for performance in delivering live streamed content, a product of years of expertise and testing by our engineers.

This model is not a theoretical abstract, but the one by which we have delivered live streams to viewers around the world consistently with optimal performance since we began streaming. In recent tests of our new TNA hybrid streaming architecture, we achieved loading times of one second as far away as Pakistan. Like every other company, we have faced hurdles along the way as we built out our live streaming infrastructure, but our technical expertise and direct control of both hardware and network allowed us to build a robust and unparalleled CDN that excels in particular in the distribution of live streams.

POPs: Benefits and Drawbacks

Many CDNs do not own their own data centers, but instead rent space in those of other companies for their equipment. While a thorough vetting process might ensure that the data centers they choose meet acceptable standards in terms of network and reliability, the nature of the POP-based model puts those CDNs that rely on it entirely at risk. Each additional POP is a means of increasing the

#	Timestamp	Prefix	ASN	AS Name	Old			New			Type	Prob
					Provider	Loss %	Latency (avg, ms)	Provider	Loss %	Latency (avg, ms)		
1	02/06/14 04:15:19 PM	46.176.128.0/17	3329	Hellas OnL...	Cogent	0	222	Telia	0	177		
2	02/06/14 04:14:49 PM	202.46.52.0/24	2914	NTT Americ...	Cogent	28	225	Level3	0	223		
3	02/06/14 04:13:18 PM	173.85.0.0/19	5650	Frontier C...	NLAYER	0	113	Cogent	0	37		
4	02/06/14 04:11:02 PM	39.32.0.0/16	45595	Pakistan T...	Cogent	0	301	Telia	0	266		

A report demonstrating how Tulix's route-optimizers dynamically choose between networks to reduce latency and packet loss.

reach of content delivery and balancing traffic. However, there are a multitude of problems and failures that can arise at each POP that are not directly in the control of the CDN and can lead to outages and connectivity issues as a result.

This does not mean that there is no value in having a presence in different locations. Tulix has recently begun horizontally scaling our CDN by deploying nodes in Europe and is now adding additional ones in Japan, the Middle East, Australia, and the west coast of the United States. However, these are intended to enhance our existing system and act as a solution to the cases where streams cannot be delivered optimally via one of the four backbones we use. Initial requests for streams will come to our central streaming system in Atlanta, which will decide whether to deliver them to the requesting viewer via route optimization or nodes. In this scenario, the POPs are a complement or enhancement to our CDN, not the CDN itself.

Even if all of our international nodes were to fail, we would still have a fully functional CDN capable of global delivery to any number of viewers. All of the other services housed in our core data centers, including cloud transcoding, playout, and VOD, would remain fully operational. When POPs are down, traditional CDNs have to reroute traffic to other nodes, potentially resulting in instability, excessive load, and loss of connectivity. In many cases, they are also unable to directly resolve the issues causing the downtime, especially if these are a problem with the data center that is housing their servers. Streaming from a controlled environment allows Tulix's engineers and support staff to find the source of a problem quickly and resolve it immediately.

The Importance of a Controlled Environment

All of Tulix's services were created with maximum control in mind. By being in full control of the infrastructure from which our streams are delivered, we are able to perform extensive testing and continuously upgrade our technology to stay on the cutting edge. Tulix's data centers are the core of our centralized distribution system and were built with infrastructure redundancy in mind. The power system is fully redundant, with two independent generators that activate automatically in the event of a power outage. Stream stability can be guaranteed via dual ingest and dual outbound feeds, while VOD files are stored using RAID configuration to prevent data loss. The network is innately redundant based on the configuration of the route-optimization systems. Should connectivity to one of the four backbones it chooses between fail, the optimizers automatically reroute traffic to the stable networks.

These precautions extend not only to our CDN, but to all of the streaming-related services we provide. Rather than rely on third-party systems and risk extended downtimes or problems which we cannot address ourselves, we built our own solutions to meet the needs of our customers. Tulix's applications for OTT devices and set-top-boxes—as well as our transcoding, cloud-playout, and monetization systems—were all built by us to be able to provide the highest level of service and support to our customers.

The Race to the Bottom: CDN Pricing

It's no secret that the CDN market is as competitive as ever. Prices have been on a steady decline in past years and there are more and more companies trying to entice potential customers with rock-bottom

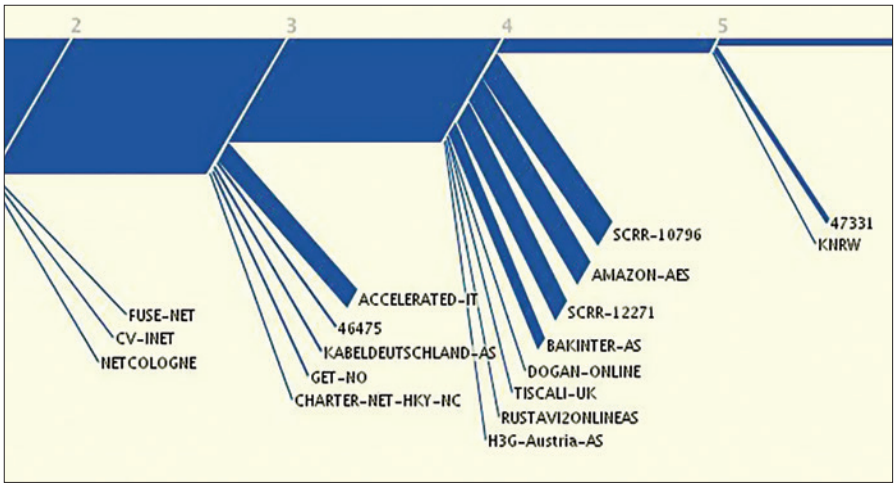

A generated graphic showing BGP hops and traffic volume to various connections.

bandwidth pricing. This change has been facilitated in part through the ease with which companies can now deploy servers in international locations without the need for the physical presence of engineers or support staff. While these companies may boast a large number of POPs, that number says almost nothing about the quality of the nodes used or the streaming expertise of the employees that operate them.

Broadcasters looking for the cheapest per-GB cost are not doing themselves, nor their customers, a favor. While on the surface prices seem like an apples to apples comparison, treating bandwidth as just another commodity for streaming ignores all of the factors that differentiate CDNs in terms of quality. The method of measuring bandwidth for streaming in terms of data transferred instead of connection (Kbps, Mbps, Gbps, etc.) is misleading itself and confusing to broadcasters who often do not understand what they are paying for, but this system of pricing has stuck around like the many other remnants of the traditional CDN model.

For many broadcasters and corporations, streaming has become either a core of their business or a vital organ that powers numerous other functions. While the urge to look for the cheapest price may be tempting, those looking for a CDN should ask themselves what the consequences would be should their streaming or experience widespread connectivity issues. Consistent quality and stability are crucial for live streaming, and sacrificing either for savings in the short term can prove to be a costly mistake in the future.

By using a centralized distribution system, Tulix maximizes both cost effectiveness and reliability, allowing us to compete with even the largest CDNs on pricing while offering superior stability and quality. Operating our CDN from a controlled environment means any service interruptions can be resolved immediately with minimal support intervention. Our expertise in streaming and our ability to develop custom solutions to meet the specific needs of our customers is what differentiates us from being a commodity vendor.

Trust, but Verify

The most important step in determining the quality of a CDN is testing it. While this may seem obvious at face value, the process for testing that many organizations follow frequently lacks rigor. All companies do their best to look good on paper, but the only way to be sure of a CDN's quality is to test it in a broad variety of situations. The test stream may work on an office connection in the United States, but will it work on a less optimal connection across the world? It may play fine on a new iPhone 5, but could someone watch it from a phone with an older version of Android? These (and many more) are questions many people may not ask when evaluating a CDN but may become all too apparent once viewer complaints start pouring in later.

We are no strangers to skepticism, particularly because some of the metrics people use to judge the quality of a CDN at first glance—such as location and number of POPs—belie the actual capabilities and performance of our CDN. Thus, we have always welcomed requests for testing as an opportunity to prove not just the viability but the superiority of our centralized distribution system as a means of content delivery. Testing is what has allowed us to fine-tune and perfect our CDN, so we embrace it as a means of both showing potential customers what we are capable of and as a way to satisfy our goal of continuously enhancing our delivery.