How does a multi-CDN strategy work for my business?
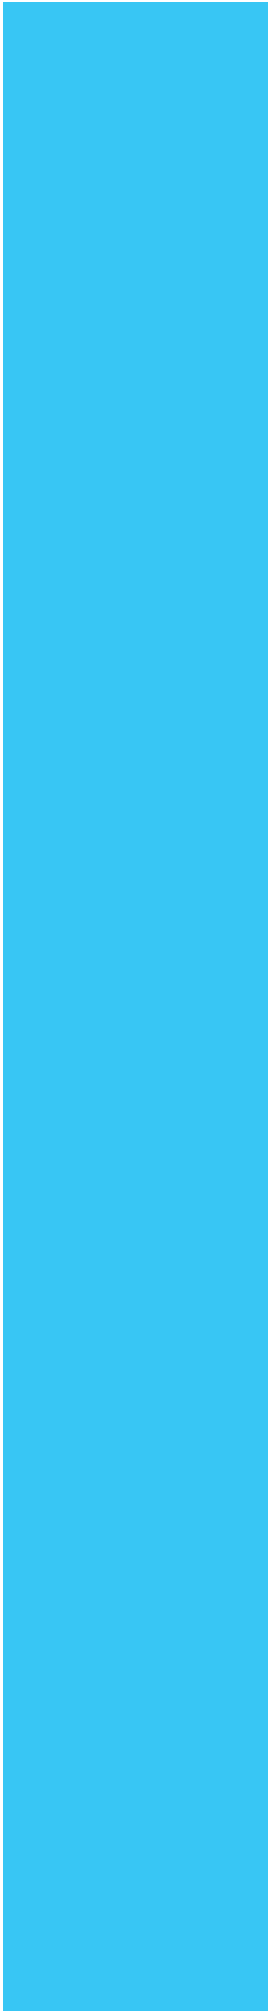
Understand the benefits a move to a multi-CDN environment will deliver to your business, your customers and the architectures involved.

LUMENSM

Executive summary

Relying on a single Content Delivery Network (CDN) exposes an online content provider to the risk of degraded performance and outages. Even the best CDN suffers from issues caused by failing hard drives, server quirks, bottlenecks and last mile congestion. A single CDN may be able to provide satisfactory or even very good service most of the time, but a single CDN will never 'always' be able to deliver exceptional performance in all regions or markets.

A multi-CDN strategy allows a content provider to share traffic between two or more providers in order to mitigate the risks of issues with any individual provider. Whether it be micro outages, peak-hour performance degradation or other issues, having alternatives allows content providers to protect their end-users against user-experience impacts caused by a single CDN.

Historically multi-CDN strategies were implemented with manual policies where human intervention was required to trigger a shift in traffic share between providers. Today, they are nearly all implemented dynamically with the use of policy-based automation algorithms to switch traffic. Policy-based automation has become prevalent as it allows for faster and more granular traffic-share shifts when the need arises and minimises the impact on the end-users.

LUMEN

Regardless of the approach taken the key is that a multi-CDN approach provides protection against chronic performance issues that can cause impacts to end-user experiences.

There are a small number of highly capable CDN providers in the market today that can be combined to create an effective multi-CDN architecture. There is also an array of tools that can be used to make a multi-CDN environment safe, easy to create, and easy to manage. Simply stated, if video delivery is core to your business, a multi-CDN approach is more than just a good idea, it's critical. The benefits of a multi-CDN environment far exceed simple vendor diversity. They include market-differentiating performance gains, faster reach to new regions, high-quality end-user experiences, and cost savings as the competition for your business heats up.

The mission of this document is to help organizations understand the reasons to move to a multi-CDN environment, the improvements that can result from this move, and the architectures involved. The best practices for making this type of move will also be discussed. It is our hope this white paper will provide a clear and concise path to a multi-CDN environment that circumvents pitfalls and gets you to the benefits of that environment as safely and efficiently as possible.

LUMEN℠

CDN performance issues

In this section, we will review examples of various types of CDN performance issues. Each of these issues can have significant impacts on end-user experiences.

Micro outages

Figure 1 below provides a typical representation of a micro outage when you have two CDNs in play. Just as usage is starting to ramp up, CDN 1 has an issue. The issue was localised, primarily impacting users of CDN 1, and it was not enough to significantly alter the median availability. However, as the chart shows during this short period, availability for CDN 1 dropped below 80 percent. This means that one in five requests for content would fail. For certain types of content, this kind of outage might be acceptable, but for pay-per-view, live sports or other premium content, this type of issue would quickly result in a high number of end-user complaints.

F1
Availability comparison of CDN 1 and CDN 2

The good news demonstrated in Figure 1 is that CDN 1 quickly recovered and was subsequently able to achieve an availability level that was on par with CDN 2.

What does this mean?

Although the duration of this event was short the impact during the event significantly affects end-user's ability to access content. With ever changing user patterns and demands it's important to understand what effect this has on end-users' perceptions as a provider of content.

A second point to note is that during these events, CDN 2 was unaffected. This is not to suggest that CDN 2 is superior to CDN 1. Issues that occur on CDN 1 are not connected to issues that occur on CDN 2. If a customer had been able to quickly shift traffic from CDN 1 to CDN 2 during these occurrences, they would have been able to mitigate impact on end-user experience.

LUMEN

Performance degradation

The second category of CDN performance issues to consider falls under performance degradation. These issues are subtler than the previous category of outages, which are characterised by error messages received by clients. In the event of performance degradation, content is received by clients more slowly and can result in the slowing of video start-up times, rebuffering or lower than average bit rates. The end-user is still able to receive content, but the quality of experience is lower than normal. The concern for content owners is the negative perception that a degraded user experience can generate. It has been demonstrated that user abandonment rises sharply with impacts to end-user experience, and as the range of choices in the over-the-top (OTT) video market becomes broader, users are not always prepared to tolerate low-quality video.

As seen with micro outages, impacts to CDN performance can be very localised. Figures 2 and 3 show a representation of the possible difference between CDN 1 and CDN 2 at a country level versus CDN 1 and 2 at a regional level. The difference can be quite marked further highlighting a consideration for multi-CDN.

F2
Country performance degradation example

F3
Regional performance degradation example

LUMEN℠

Background

The delivery of high-quality video over the Internet has an inherent reliance on a distributed architecture to achieve high definition / high bit-rate delivery to end-users. The CDNs that provide this functionality leverage DNS (the Domain Name System which is the part of the Internet that translates IP numbers into human-readable names) to redirect end-user requests for content to the CDNs edge servers, rather than to the content publisher's origin servers.

End-user redirect to CDN using DNS

As Internet-delivered video services move toward the dominant business model, and the high-quality, uninterrupted delivery of that video becomes mission critical to the success of those service providers, the risk of using a single CDN partner for delivery is something that many companies are no longer willing to accept. Also, empirical data shows that costs can be reduced while performance is improved by adding the right CDN partner to the mix.

Fortunately, there are now several acceptable options when it comes to global, Tier 1 CDN providers. Plus, the migration path to a multi-vendor CDN environment is well travelled, well understood, and can be quick and painless if you follow the trail blazed by those who have made the transition successfully. The two basic types of CDNs, those that own their own networks (such as Tier-1 network operators like Lumen) and those that don't ("PurePlay" CDNs), should also be taken into account. It is often good to have one of each in the mix, as their strengths tend to complement rather than duplicate each other.

LUMEN

The high level perspective

The four general considerations you must take into account when migrating from a single source to a multi-CDN environment are:

1. Workflow continuity
2. Traffic-distribution policy for business and performance
3. Performance measurement and telemetry
4. Traffic distribution tools.

Workflow continuity

An in-depth analysis of your current workflow, and the ability of a second CDN vendor to accommodate that workflow, is critical for a smooth and seamless transition to a multi-CDN environment. Identifying the "Must accommodate", "Should accommodate" and "Could or should change" elements of your workflow is the first step. Often, a change in the vendor ecosystem can be an excellent time to make tweaks to the workflow in order to increase efficiency or correct process shortcomings that you have previously identified.

To avoid duplication of effort and added workflow complexity in a multi-CDN environment, close attention needs to be paid to ubiquitous functionality. The CDN you integrate into existing workflows should have the flexibility to accommodate industry standard approaches to token authentication, content invalidation, geo-filtering and caching / origin-fill rules. You may find that your current vendor has done one of these in a proprietary manner, but all these functions have industry-standard best-practices equivalents that you can sync up between vendors.

Close partner-level interactions with the incumbent CDN, as well as prospective CDN partners, are key to a seamless transition to a multi-CDN strategy. Diligence in this phase will make the transition much easier, and any CDN vendor interested in winning a portion of your business will be willing to engage at a deep level. It is critical to understand and address workflow issues and design a proof-of-concept implementation plan to demonstrate each element of compatibility with your desired end-state workflow.

Traffic-distribution policy

Not all CDNs are equally strong in all areas. This truism underlies the core value of a well-considered multi-CDN strategy. Adding the right CDN to your ecosystem should increase performance and reliability. The key is making good decisions about the percentage and makeup of the traffic sent to each CDN. Decision criteria are pretty straightforward:

LUMEN℠

1. In what regions does each CDN perform best

A. What is the right level of regional granularity to consider? City? Region? Country?

2. Bearing the following in mind, what commitment levels make sense across all vendors?

A. Commitment drives unit costs.

B. To ensure maximum benefit CDNs need a minimum amount of customer traffic (>15%) in all markets / regions to keep caches warm in the event of a traffic spike or failover.
This better distribution ensures smoother transition and ramp up in the event of a problem and lessens the impact on the CDN provider and importantly the end-user.

3. Policy must accommodate regional strengths and support maximum aggregate performance.

4. Policy must move traffic in a closed-loop controlled manner to avoid traffic shifts in excess of 50 percent in a given region:

A. Cache-fill traffic can overload origin infrastructure if too much traffic shifts from one CDN to another.

B. Hands-on control and data analysis over several months should precede any attempts at full automation to establish safe boundaries for traffic movement.

Performance measurement and telemetry

The key to performance-based decisions is an impartial view of global and regional performance between CDNs. To determine CDN performance, vendor-neutral tools and application-appropriate analytics should be employed. Latency is a reasonable proxy for many performance elements, but it may not tell the whole story. Public domain tools like pingdom.com and webpagetest.org, and commercial tools from Cedexis to Conviva can provide you with deeper analytics specific to your application.

Third-party data snapshots from the Internet and from client-side telemetry will help you identify average performance as well as any specific time-of-day challenges your prospective vendors may be experiencing. Once you have that powerful data, you can create a traffic-distribution policy that maximises both performance and cost efficiencies. Real-time and historical reporting Application Protocol Interfaces (APIs) from your CDN vendors, or third-party analytics engines, can further enable efficient and appropriate traffic shifts at a high level using server-side data.

Client-side telemetry takes this approach to the next level, providing data that is a direct reflection of actual end-user experience. This telemetry provides average and discrete viewing bit rates, rebuffering percentages, fatal (restart) errors, and a host of other diagnostic data that can be critical to traffic-balancing decisions and troubleshooting.

Once you have all this performance data, it is important to understand how you can consume and utilise it to inform changes in traffic balance. KPIs, such as error-response codes, utilisation changes by region, rebuffering percentages and average throughput, coupled with the ability to set and alarm on specific thresholds, all determine how quickly and efficiently you can manage performance, traffic balances and customer experience / satisfaction in near real-time.

LUMEN℠

Traffic-distribution tools

Now we get down to the mechanics. What tools and approaches will give you the best performance, resilience and control in managing your multi-CDN environment? A hierarchy of regional splits (for performance), and then a split by percentage of users (for resilience and cache warming) is generally considered to be the best approach. Other approaches that require a process step to the publishing workflow, such as tagging specific content to specific CDNs, are not recommended because they limit your ability to make fast changes to traffic percentages given to each CDN. They can also induce large cache-fill volume if a CDN has to take over serving content that it has never cached.

The two broad categories of traffic management tools to consider are DNS (or DNS/HTTP) based traffic-management / load-balancing tools and client-embedded, traffic-management tools. Each of these comes with inherent advantages and limitations. Specific features and applicability will vary with each vendor. Companies often start with a DNS-based tool to establish and manage the multi-CDN environment for its simplicity, control and ease of use.

They then migrate to a client-side approach over time to augment or replace the DNS-based solution. The key to a successful migration from one to the other is to take the time to understand the behaviour of your end-users, CDNs and any content-specific nuances that affect performance.

DNS-based traffic management

A software-based DNS-load balancing system has many strengths. There are no hardware or service-contract costs. Since the members of your operational staff already know and use DNS to manage origins and the incumbent CDN, there is no steep learning curve for ease of use (usually a web-based Portal). You retain absolute and granular control over where your traffic is directed. There is no significant integration work other than to define and implement your traffic-management policies, and DNS itself is well understood and predictable. Below is a simplified view of DNS-based traffic-management flows.

A DNS-based system alone leaves you with load-balancing policies that are somewhat static. Without client-side telemetry, automated real-time traffic shifts can only be based on coarse server-side data, such as increases in HTTP response codes and traffic thresholds provided by the CDNs themselves. Changes to the distribution can always be made manually based on observed performance differences, and automated failover can be configured in the event of a catastrophic failure of one of your CDN vendors. However, overcoming real-time transient Internet conditions without additional client-side telemetry is difficult.

DNS-based traffic management can also suffer from the use of an end-user's DNS resolver to approximate their geographic location. Under some circumstances, a DNS resolver that serves a wide geographical subset of end-users can cause inaccurate identification of those users' locations. Advances in DNS resolution that include access to the IP of the end-user have reduced this shortcoming of DNS-based traffic management where the end-user's DNS resolver supports this feature. Unfortunately, this support is not ubiquitous at the moment.

Client-side telemetry

Client-side data can show you exactly what your end-users are experiencing globally, regionally and individually. By embedding a bit of code into your player, you can gain a lot of insight into end-user experiences and the underlying technical causes of impairments or sub-optimal viewing experiences. This data, when combined with a DNS-based load balancer, completes the picture, enabling granular real-time control, appropriate automation and precise traffic shifts to mitigate well-defined problems.

Embedded client-side traffic management

We should consider client-side load balancers as well. Usually, these use client-side telemetry and an outside-in view of CDN performance to make traffic-switching decisions at the individual client level. These network-aware clients can switch CDNs midstream in real time when conditions impair one CDN but not another. Client-side load balancers offer some benefits in reduction of re-buffering incidents and other factors that influence customer experience.

The downside to client-embedded technology of this nature is the integration process (modifications to the player itself) and iterative testing that is required before roll-out. Automated, end-user-oriented CDN switching can also make management of commitments to CDN vendors difficult, and it may impair or override any attempts to direct discretionary traffic to a specific CDN vendor for commercial or business policy reasons. Another important consideration is that not all adaptive-delivery protocols support CDN switching at the client level, so a client-side, traffic-control system cannot be ubiquitous across all devices.

Finally, commercially available client-side, traffic-management technologies may base pricing on monthly peak users or hours of content consumed through the client. This results in significant and unpredictable costs. Lumen offers its own in-stream switching capability via our multi-CDN Orchestrator, which addresses many of these challenges.

Conclusion

Sole sourcing of critical services or infrastructure has largely disappeared in all sectors of business. The historical dominance of one CDN vendor, and, until recently, the lack of credible competitors, has made the CDN delivery of video one of the last bastions of sole sourcing. Studies show that this no longer makes good sense. Most OTT Video companies agree that a multi-CDN environment is critical to their continued growth, competitive differentiation and ultimate success in this hypercompetitive quest for the video consumer's attention. Best practices suggest that good data, close CDN collaboration, and the combination of DNS-based control tools and client-side telemetry enable the optimal multi-CDN environment.

While the initial move to a multi-CDN environment may appear risky, in practice it is relatively straightforward, and the benefits are undeniable. The keys to successful implementation are choosing the right performance-measurement platform, the right traffic-control platform and policies, the right CDN partners, and joint planning with organisations experienced in multi-CDN environments. Follow these guidelines, and you will soon be enjoying the benefits of higher performance and availability, lower costs, better information, enhanced customer experience and the reduced operational overhead inherent in the right multi-CDN environment.

LUMEN