MAINSTREAMING®

Q1 | 23

Whitepaper

# Broadcast-grade streaming

How a world-class live video streaming delivery solution is designed, and why you should expect to simultaneously achieve the best possible viewer QoE, business ROI and environmental sustainability

Including **three case studies**

arqiva GEDI GRUPPO EDITORIALE DAZN

# Executive Summary

## A content delivery network (CDN) must be designed and operated in a particular way to achieve demanding technical and business KPIs, especially for live video.

The capacity of the network is the key to performance. If capacity does not exist, high performance cannot be achieved. Without capacity, viewers will suffer poor performance. When thinking about capacity, a fundamental principle based on simple physics is that capacity closer to the consumer is better than capacity that is further away.

The core responsibility of a CDN service provider working with live video delivery – the most performance sensitive content delivery use case – is to leverage **every possible pathway from the CDN Edge to the consumer to achieve perfect delivery performance.** This relies on intelligent software to manage content and streaming sessions, which requires industrial-grade real-time monitoring capabilities that can be acted upon by the CDN.

To leverage those pathways, they must be set up in advance ready to be called into action when needed, a critical network planning task. When an event occurs which reduces the capacity for video delivery at any point in the network – which could be a loss of capacity or a surge in demand – then the original network design and the exact rights of access to those other networks are the two key principles that will determine the ability to meet performance expectations.

If the demand is not fully manageable within the available capacity, then the capacity planning or capacity implementation was wrong. That said, we know that demand forecasting is not a perfect science, so it is easily possible for demand to over-subscribe the capacity. And even if demand is forecast accurately by a single D2C Streamer, in a multi-tenanted CDN environment the total demand from all tenants may cause a capacity overload, affecting everyone, especially live video streaming. **Unless the streaming capacity is guaranteed to the D2C Streamer, it is simply not guaranteed.**

Multi-CDN architectures have emerged, partly to address this problem, creating the ability to allocate streams to another CDN when there are performance issues. But this remains problematic if every alternative pathway is multi-tenanted and does not provide guaranteed capacity. Given that guaranteeing capacity in a multi-tenanted environment can be costly, there are cost-performance and control / no-control trade-offs to consider.

**The way to guarantee performance is to use a more private, managed content delivery platform or service.** This removes risks of poor performance by guaranteeing the capacity will always be available. **To achieve broadcast-grade streaming, it is important to have guaranteed capacity**, at least for a large proportion of the content delivered, and to have technology and service partners who are experts in video and focused on real-time performance management on a 24x7 basis.

**Based on innovative business models, a private CDN with guaranteed capacity can be very commercially attractive, particularly to D2C Streamers that routinely deliver content to their audiences in a well-defined geography**. Working with service partners with broadcast-grade video delivery expertise, advanced monitoring capabilities, and proactive service management methods can assure and continuously improve streaming performance to help a D2C Streamer **reach new heights of customer satisfaction and business profitability**.

The bottom line for D2C Streamers is that your **CDN partner really matters**. This whitepaper is for D2C business leaders, OTT product leaders, and video streaming engineers, and provides a detailed look at how MainStreaming delivers new capabilities for live video streaming distribution that can take your streaming service to the next level.

# Table of Contents

# What is broadcast-grade streaming?

As Broadcasters focus in on streaming-first strategies, streaming must achieve "broadcast-grade" standards. At MainStreaming, **broadcast-grade streaming** means **three** simple things.

1. Video delivery is of a **consistently high quality and low latency** that does not cause the viewer any dissatisfaction.

2. Reliably **scaling to many millions** of concurrent viewers without having problems with quality or latency.

3. A **predictable delivery cost** that enables a strong return on investment for the Streamer, even when viewership is unpredictable.

In this whitepaper, we will explain how these three business success factors can be simultaneously achieved. They are not mutually exclusive and there are four well-defined steps to take.

We will consider some industry leading organisations already turning to MainStreaming to help them achieve their own broadcast-grade streaming objectives. And we will describe how we work as a business to enable this to happen not just once, but minute after minute, and hour after hour.

**"Whether you are a subscription-driven business or ad-driven business, your customers' quality of experience will be a major driver or satisfaction, churn, and profitability"**

# Our Perspective

## Video streaming must continue advancing to reach broadcast-grade.

When streaming replaces existing dominant forms of video delivery like terrestrial, satellite, cable, and IPTV, we could require 80 times more streaming capacity than we have today. This number could be higher as the media industry introduces more advanced resolutions and viewing formats.

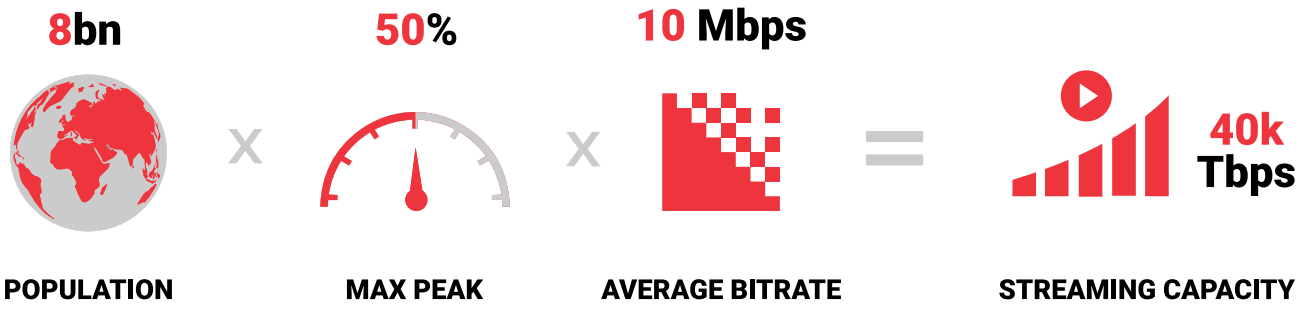


**Streaming will ultimately replace all the other existing technologies for two main reasons**. First, the finite frequency resource used for broadcasting is valuable for arguably more important use cases, like mobile connectivity and security. Second, streaming over the internet provides media companies with more delivery flexibility and new commercial opportunities, while providing consumers with more viewing flexibility.

But the full potential of streaming's success could be hampered by quality, scalability, and cost issues, plus valid concerns about streaming's sustainability credentials. Particularly for the delivery of live video, which the internet and ISP networks were not originally designed to support. As streaming generally grows, and especially as live video streaming grows, the whole media industry should look to new streaming models that have already been implemented by the pioneering D2C Streamers and are ready for adoption by the next wave of large-scale D2C Streamers.

**MainStreaming developed its intelligent media delivery platform (iMDP) exactly for this purpose**. This next-generation CDN solution, focused on energy-efficient Edge video delivery with innovative business models, enables D2C Streamers and Internet Service Providers (ISPs) to manage video delivery operations economically and to consistently deliver best quality. In other words, MainStreaming enables D2C Streamers to deliver "broadcast-grade streaming".

# Arqiva & MainStreaming partner to offer new streaming services

Broadcast infrastructure provider Arqiva and intelligent media delivery platform MainStreaming have forged a technology and services partnership, jointly to offer distribution services for the media streaming market.

The partners will explore how the combination of MainStreaming's CDN technology and broadcast-grade streaming experience with Arqiva's global media infrastructure and managed services capability can offer more scalable, flexible, and programming-centric content distribution services for the media streaming market.

According to the pair, with ever-growing viewer numbers on streaming services and the increasing strategic value of online audiences, the streaming needs of the biggest broadcasters and service providers are greater than ever. The combination of a large audience served, consistently high video quality, and low latency is the tough combination to get right hour after hour. As such, secure, scalable and cost-effective content distribution networks are vital.

They say that existing streaming distribution networks are not well suited to deliver either the quality of service required by service providers, or the quality of experience expected by audiences.

The growing carbon footprint of streaming services is also a concern for both providers and audiences. Arqiva and MainStreaming are coming together to address these issues and to challenge conventional approaches to content distribution.

“MainStreaming's technology makes true edge computing for the media industry a reality, and already delivers important benefits for our industry-leading customers," stated Antonio G. Corrado, CEO, MainStreaming. “We are excited to work together with Arqiva and the media industry to take advantage of our real-time, ultra-low latency, highly scalable streaming capabilities to deliver broadcast-grade streaming and also create new and exciting edge applications for video delivery."

“The streaming world is changing fast and navigating the commercial and technical issues has never been harder," added Clive White, CTO, Arqiva. “Arqiva and MainStreaming will be collaborating on a range of new capabilities and service offerings to meet these challenges with a view to optimising the customer experience and adding value to the biggest broadcasters in our core markets."

# Step 1: Deploy the right solution

The first step to achieve broadcast-grade streaming is to ensure the right solution is designed and deployed. Good planning and design are essential to the realisation of a broadcast-grade streaming platform.

## The QoE dependency on hardware

The overarching principle for high performance streaming is to focus on the customer experience, or QoE (Quality of Experience). Video is a specialised workload, and so the hardware and software – including the operating system kernel and the streaming analytics – must be optimised.

A first principle in achieving broadcast-grade streaming is that capacity cannot be created from nothing. And because streams require capacity, capacity must be in position for broadcast-grade streaming to be a reality. Capacity must therefore be planned and implemented all the way from the Origin to the millions of consumers. Ideally the full end-to-end CDN capacity should be guaranteed to the D2C Streamer. Any other approach risks congestion and interruptions to the streams. Servers, ports, switches, and cables all matter, and cannot be created from nothing.

A second principle is that this capacity must scale-out at the edge of the network, more than at the origination point, because millions of consumers may all watch the same content or at least some level of similar content. Therefore, the capacity should exist close to them, not far from them. Internet Service Providers (ISPs) know where to deploy CDN Edge Cache capacity based on their network congestion patterns. MainStreaming works with Content Providers and ISPs to deploy the primary locations and the inter-cache connectivity pathways that will be used as resources when making routing decisions.

After the capacity and location of the infrastructure are defined, then the design of the hardware itself becomes the next critical component.

For best performance, the processing hardware (i.e., servers) should work in a bare metal mode to avoid unwanted interrupts from routine background processing of operating systems and virtual machines. The CPUs in the servers should be tuned so they

are ready for simultaneous stream monitoring and video processing workloads. The CPU should also be tuned to achieve sustainability objectives, reducing power consumption as far as physically possible.

As a workload type, video needs to cater for legacy environments, including older TVs and set-top boxes (STBs). Sometimes, older equipment cannot handle the latest video streaming standards, like an HTTP redirect command. It is important to identify these situations and set up the streaming system to support these niche requirements, to assure quality of experience to these parts of the audience.

The transport hardware (i.e., network and switches) is a network design function, focused on connecting all the hardware together. Network architectures can be designed to have varying levels of performance and resilience. Networks are complex, somewhat like the human body. Network experts describe troubleshooting a network as similar to a doctor working with their patients – for the doctor, it is essential to see and understand the person to diagnose their problems. It is the same with networks. There are many conditions that can affect a network and its video streaming performance, so one must be close to the network to understand all of its dynamics.

To design a network to achieve high-performance objectives, MainStreaming engineers analyse optimal network paths, resource availability, and terms and conditions with connectivity providers. If a CDN provider does not study these subjects in detail, for the initial network design and then on an ongoing basis as environmental conditions change, then performance will be sub-optimal, and consumer QoE will not achieve all its goals. At the very cutting-edge of performance, the finest details always make the difference between good and great.

## The QoE dependency on software

Once hardware is in position to be utilised, the software layer becomes the next deciding factor for performance.

Video streaming software at the edge of the network has three primary purposes:

- To efficiently request video segments from the Origin platform
- To efficiently deliver video segments to the consumer devices
- To efficiently balance load between Edge servers to ensure the entire audience receives high quality video

MainStreaming took a simple yet technically innovative approach to achieve these objectives. Working backwards from the vision of world-class customer experience, the software was developed to be highly distributed, so that the entire platform would scale for the largest audiences. The software was developed to transfer streams or files from one Edge server to any other Edge in the network without adding any additional latency or disrupting the viewer's experience. The software was also developed to ask for content only once from the Origin.

The result is a set of autonomous Edge nodes, scalable to any number, that are widely distributed and share access to the same content. This core architectural design is how MainStreaming's CDN is able to reach broadcast-grade performance without slowing down under a heavy load. But the software is not only focused on consistent high-quality streaming – it is also about being able to stream intelligently. That is where monitoring comes in.

## Design in monitoring from the start

A common occurrence in network design is to plan first for the delivery of the content and plan afterwards for monitoring that delivery. However, this is not conducive to broadcast-grade performance, or even satisfactory QoE performance.
MainStreaming started its journey

differently. Delivery and monitoring were given equal importance in the design. This is significant, because monitoring millions of streams in real time is not easy, and the software and hardware architecture must be designed specifically to achieve this goal. By designing in monitoring from the beginning, MainStreaming prepared itself fully for a future of broadcast-grade streaming.

The path for OTT video from the Origin to the consumer is long and winding. There are many parts to the chain, which all change quickly in terms of capacity and availability. ISPs have frequent, often unannounced, maintenance windows. Servers fail. Audience demand unexpectedly exceeds expectations. The end-to-end delivery chain must withstand these shocks to the system. We cannot "fly blind" and expect to perform well in a highly reactive mode. The system must be as all-seeing as possible and we must strive to be proactive, and at least always prepared for failure.

With monitoring and streaming co-designed, it is possible to work with D2C Streamers to achieve specific goals, such as prioritising streams to different types of consumers or for different types of content. Often, this type of capability is tied to the client-side tools and the data they provide. But with the right server-side monitoring, good decision-making by D2C Streamers can be enabled by the CDN itself. Only with a real-time, scalable monitoring solution can this be achieved.So how does MainStreaming achieve high-performance monitoring alongside high-performance streaming?

**First**, the monitoring software is deployed on-board every Edge server. This makes each Edge server autonomous. Each Edge can independently determine to not accept any new streams, and it can decide to offload existing streams, or even just segments. This capability to make real-time capacity management decisions is essential to protecting the viewer's QoE. This approach to Edge-level monitoring means that monitoring scales as Edge streaming capacity scales. MainStreaming's solution does not reach an architectural limit for monitoring at scale.

**Second**, MainStreaming chose to work very close to the hardware

level because it is fundamental to performance. The software directly interfaces to every hardware component, like a NIC (Network Interface Card), and can push every hardware component to its limit. There is no virtual machine layer inside the server, so there are no interrupts to manage. These reduced layers between software and hardware mean more control, and more performance. They also mean more energy-efficient operations, limiting the processing power used by each Edge.

**Third**, MainStreaming tunes performance at the kernel level, where there is more system performance data available to be harvested (e.g., data about network, RAM, IO, storage, temperature, etc.). This gives MainStreaming more intelligence on which to make delivery decisions and gives D2C Streamers full governance of their service delivery. Through intensive analysis, MainStreaming has learned where to make kernel-level changes to improve performance specifically for video streaming and video monitoring workloads.

**Fourth**, specialised video workloads demand two operational priorities. The first is to collaborate very closely with hardware partners, to optimise the software for both streaming and monitoring. MainStreaming looks for ways to squeeze every piece of performance out of the available hardware, thus driving the highest levels of energy efficiency. The second is to provide a managed service to customers for monitoring applications. Experience and extreme attention to detail are required to correctly set the data capture and processing thresholds, so as not to overload the system which could impact streaming performance.

# Step 2: Optimise streaming operations

## The best Edge

Broadcast-grade streaming delivery is based on a fundamental principle, reimagining Occam's razor which states that entities should not be multiplied beyond necessity, meaning that the stream should be delivered to the consumer from an Edge cache that is in the nearest possible location to the consumer. This is easy to say, sometimes hard to do, and in practice does not mean that this is always the right answer.

In the final analysis, the best Edge cache to stream from is the Edge that can deliver the stream most effectively and efficiently. This means achieving the required business KPIs with minimal resource consumption. As a rule, the nearest Edge is the best Edge because of the simpler network connectivity between Edge and Consumer, which represents the most energy efficient pathway. But because this is not always the right answer, MainStreaming defines the best Edge as the Edge that best achieves the defined performance objectives.

Selecting the nearest Edge to the consumer requires awareness of the consumer's location. While DNS (Domain Name Server) Resolver IP is an industry-standard way to identify the location of the consumer, ISPs have more accurate ways to establish where the IP address is located for an individual consumer. ISPs know where the router or modem is physically located for their broadband customers but presenting this data back to a CDN that is outside the ISP network is not useful because the CDN often only has a very limited number of connection points to the ISP and these points are identified by a simple set of Resolver IP addresses. From this point on the ISP will deliver the content to the consumer.

An on-net CDN, however, can benefit from the ISP's detailed IP address information. The best method is eDNS0, which represents regional level information about consumer location. If an ISP has implemented this method (which would require them to invest in the appropriate infrastructure to do so) then it enables a CDN to select Edge Caches for stream delivery in a more precise way. If the Edges are inside the ISP network, then eDNS0 allows an Edge to be selected that is already close to the consumer.

Once the Edge Cache is selected, streaming initiates and will continue uninterrupted unless there is a change in the consumer's streaming performance. If there is a change it may be necessary to take evasive action and reroute the stream through another path (assuming it has been designed in, as previously mentioned). So, the ongoing task of the MainStreaming Edge at this point is to monitor, monitor, and monitor.

## Monitor, monitor, monitor

Monitoring is focused first and foremost on understanding the user experience. The monitoring application on every Edge will identify if its group of users, down to a sub-group level (e.g., HLS vs. DASH, fixed broadband vs. mobile), is being served poorly. These pre-defined performance thresholds are set inside the MainStreaming CDN Manager with direct input from the D2C Streamer. Thresholds can be changed in real-time and adapted for different types of content, events, and groups of consumers. Experience has taught us which configurations generally work best for different scenarios, and so we make choices about what we need to monitor to best understand performance and take actions.

Assuring QoE by the CDN is not a one-time activity. It requires continuous monitoring of stream performance, and that requires continuous monitoring of upstream Edge performance, on-board Edge performance, and downstream-Edge performance. MainStreaming uses a broad range of performance and system health parameters to proactively determine when QoE assurance thresholds will be broken, which is converted to a single measure of quality and used to proactively manage stream load on the Edge. If the quality measure breaches its threshold, then new streams and possibly existing streams are seamlessly diverted to a new best Edge, which may be in the same PoP, a new PoP, or even potentially on the MainStreaming public CDN service if a MainStreaming private CDN service is the primary platform delivering the content.

## The new best Edge

Before switching an existing stream to a new best Edge, we must consider the consumer experience. No one wants to be automatically disconnected and reconnected as this might completely disrupt the viewing experience. MainStreaming therefore implemented transparent Edge-switching. This means that a consumer's session is redirected to a new Edge, which has been selected for its better performance, without any disruption to their viewing experience. Without this, the customer experience of a switch would be poor, and only ABR ladder changes or an equally seamless CDN selection would otherwise save the viewing experience, which is not an ideal solution. The seamless transparency capability of Edge-switching is important for protecting the streaming customer's experience.

To make the choice to Edge-switch, MainStreaming uses all available data about the current streaming pathway and alternative streaming pathways. This "navigation processing" functionality is what we call the Orchestrator, and it is a core capability for broadcast-grade streaming. Alternative pathways need to exist if the customer experience is to be protected and maintained. If those pathways were not established during the deployment phases, then the customer will get what they get – a best efforts approach. This is where "good enough" has evolved in recent years to instead be "broadcast-grade" as streaming growth has accelerated.

Here we make another point about the fundamental importance of capacity. There may be an alternative pathway for a stream to take, but the question becomes: is there available and sufficient capacity for a stream to utilise? An alternative pathway without available capacity does not help. In the world of latency-sensitive and QoE-sensitive streaming, capacity matters.

The final decision to reroute a stream is rule-based. The rules are enhanced continuously based on analysis and experience, with changes agreed between MainStreaming, the ISP, and the D2C Streamer to avoid unintended consequences of changes. MainStreaming is typically dealing with outlier performance cases representing a small

percentage of the total audience, so we make controlled changes in software-based rules and alternative pathway configurations, rather than allowing an unsupervised AI to make changes. After many years of studying live streaming performance we know the best practices that need to be applied.

## The Smart Origin

There is a simple concept for CDNs interfacing to an Origin: the CDN should do everything possible to assure the Origin can deliver to the Edge on time. If the Origin delivers to the Edge on time, then the Edge has the best chance to deliver to the final viewer on time. When the Origin delivers late, the Edge is put under intense pressure to deliver perfectly in order that the Player has the chance to maintain a high quality of viewer experience.

MainStreaming has architected two important features that maximise the Origin's ability to deliver to the Edge on-time.

The first feature is called the Smart Origin. This software component can be configured to ensure the CDN asks for content only once from the Origin. Once the content is received into the CDN, it is distributed between Edges as required rather than having individual Edges or Intermediate Caches make their own requests to the Origin. Effectively, the Edge that ingested the content from the Origin becomes the new Origin for the rest of the CDN. Smart Origin therefore minimises the chance of overloading an Origin, it minimises the cache storage required across the CDN, it minimises energy consumption by reducing storage size in the CDN, and it also minimises egress cost from cloud-based Origins.

The second feature is called the Best Origin Selector, which brings important benefits particularly when working with multiple Cloud-based Origins. The Best Origin Selector feature monitors the egress performance of each Origin the CDN is connected to. If degradation is detected, the CDN will automatically switch to another Origin that is performing better. It is possible to continuously switch between Origins in real-time, which can be required when multiple CDNs are making different requests to the same Origins and placing heavy loads on

the system depending on their own content requests and their cache hit ratios. For MainStreaming customers using cloud-based Origins with multi-CDNs, this feature improves their Origin performance.

## Customer-focused optimisations

The pathway from the Origin to the consumer is long and winding. The Edge layer is uniquely placed to offer extra intelligence about streaming performance given its mid-point location along the streaming route.

MainStreaming Edges look upstream to the Origin and downstream to the consumer. Monitoring these two directions is a fundamental process as video segments are retrieved from the Origin and delivered to the consumer. With this visibility, and because we are obsessed with delighting our customers and our customers' customers, we have provided extra information to our customers to help them perform at their best.

In one example we implemented a report-back feature to our customer to tell them when a stream prioritisation decision could be made for a single household receiving multiple streams from our customer. Using this API-driven information exchange, our customer could then instruct the MainStreaming Edge to prioritise one stream over another, such as a live stream over a VOD download.

In another example we implemented a method of informing customers about an Origin time-out problem. Many Origins do not inform their customers if a stream fails to originate. The CDN, on the other hand, knows exactly when this happens because we fail to receive video segments as requested on behalf of the consumer. MainStreaming provided an automated alert to our customers of any Origin outage we could identify, which generally resulted in our customer restarting an encoder to continue delivering the streams.

# Step 3: Focus on continuous improvement

Improving video streaming is a complex human process. In other words, making improvements requires a lot of data input from multiple sources, a high level of expertise and experience, and a lot of analysis of the alternative solutions to identify the most effective and efficient way to improve. In the end, hardware and/or software must be updated somehow to improve future video streaming performance.

In addition, the area of real-time improvements – i.e., making changes to stream delivery during the consumer's viewing time - are limited by contractually agreed capacity availability. Alternative delivery paths can be taken if they exist. Contingency plans must be in place before streaming begins. If streaming fails, it is likely that the capacity plans need to be upgraded, requiring human intervention and approval.

For these reasons, machine learning and autonomous artificial intelligence are not at the heart of video streaming performance. Experience, planning, and design are the key ingredients.

At MainStreaming, we operate a highly customer-focused continuous improvement process. When we identify that consumers have not received perfect service our data scientists and video specialists analyze all the available data, from our own systems and from our customers (including other suppliers if possible). Some customers already invest significant internal effort to understand performance outliers and strive for perfect performance. Their detailed insights help us.

When we know what to change to make an improvement, we fix the software, hardware, and/or network set-up. We often also upgrade our monitoring solutions to better identify the root causes of similar problems in the future. Our overarching goal is to simplify solutions, so they withstand environmental impacts more efficiently. Simplifying technology is always the most difficult thing to do – the simplest solutions are the most elegant. And to make solutions elegant requires our teams to have vision, talent, a great work ethic, focus, and access to lots of information. These are the ingredients that enable MainStreaming to perform at the highest possible levels.

# Elemedia (Gruppo GEDI)- a customer case study


## Who is Elemedia

Elemedia is part of the GEDI Group, an Italian media conglomerate, formerly known as L'Espresso Group. Originally founded in 1955, it is based in Rome, Italy. With major publications including La Repubblica and La Stampa, and also national radio stations including Radio Deejay, Radio Capital, and M2o, the GEDI Group has over 5.5 million readers and listeners every day, making it one of the most important editorial businesses in Europe.

## Case study

Elemedia operates leading national radio broadcasters throughout Italy including Radio Deejay Radio Capital, and Radio m2o. Their streaming delivery is increasing all the time as listeners "tune in" from work and home. Streams are delivered direct to listeners via websites and Apps, but also via 3rdparty streaming platforms including aggregators (e.g., Radioplayer), Amazon Echos, Sonos, and Google Home.

Elemedia chose to work with MainStreaming to address two requirements. First, to continuously improve their streaming quality as listener count increases. Second, to find ways to understand stream consumption via the 3rd party platforms used by listeners.

As Paolo Ruzzier, Technical Director at Elemedia, says, "These two areas are critical for our business. We are delighted to have MainStreaming's customer-focused support to address these important broadcaster requirements and help us to continuously improve our listener experience and business performance."

# Step 4:
# Be easy to work with

MainStreaming was born as a service provider for D2C Streamers. Our core function is to take care of all the technological and operational details associated with delivering video from the point of origination to the point of delivery into the ISP networks. We manage all aspects of technical design, deployment, and maintenance to assure best-in-class streaming performance.

But we also acknowledge different types of customers have different requirements and varying levels of need for managed services. Some customers want to outsource 100% of their content delivery operations. Some customers are highly engaged in monitoring and making improvements. Some customers collaborate closely with us and ISPs to design network topologies. Some customers procure hardware and network connectivity themselves and we simply provide software and services. Some customers want to own and operate the entire platform themselves and rely on us only for technical support and maintenance.

Our view is that optimal performance is achieved by experts with a wide view of industry best practice, designing the right solution from the beginning, managing the day-to-day performance of the system, and leading continuous improvement efforts. Streaming is a complex, multi-faceted delivery process, operating in a complex networked environment. Laser focus by deeply committed experts is how to continue performing well day in and day out.

Whether we work with a D2C Streamer who wants to receive a fully managed service, or a Network Operator who wants to receive a partially managed service or a fully-supported platform, we bring our customer-centric attitude to ensure broadcast-grade performance.

# How to achieve best-in-class performance and ROI

Streaming performance is fundamentally determined by capacity availability, disciplined capacity management, and intelligent stream management. The goal is therefore to avoid congestion throughout the end-to-end chain. Streamers looking for broadcast-grade performance need the ultimate in capacity allocation and management. Private Edge capacity delivers this result.

At the same time Private Edge Services also deliver best-in-class costs. This can be counter-intuitive for many people because cloud services have focused on pay-per-use business models for years. Part of the reason for better costs is the close relationship with ISPs that a large streamer should have. The economics of internet delivery mean that certain costs remain in place when the relationships between Streamer and ISP are distant and managed via multi-tenant third party service providers. ISPs make money from selling connection capacity. But they must also invest significantly to create more connection capacity as streaming grows. And not only does this capacity cost, but it also consumes more energy. There are better ways to deliver video for the largest streamers who are responsible for the consumption of a disproportionate amount of internet bandwidth – back to Occam's razor.

That better way is the on-net private CDN model that is based on capacity, not consumption. This business model is fundamentally different for Streamers to consider and should be studied closely (see table below).

| Consideration | Consumption Model (per GB) | Capacity Model (per Gbps) |
|---|---|---|
| An increase in the Number of concurrent Viewers of the streaming service… | Will increase costs | Will increase costs |
| An increase in the Video Resolution (average Gbps) offered to the viewers… | Will increase costs | Will increase costs |
| An increase in Hours Viewed by a fixed number of viewers… | Will increase costs | Will **not** increase costs |
| The nature of the costs is… | **Variable**, and will increase in line with any form of consumption growth | **Fixed**, and will only increase if capacity must be expanded |
| Performance Guarantees… | Not typically offered because capacity is not guaranteed | Offered because capacity **is** guaranteed |
| The content delivery cost will economically scale… | Only if per-GB unit prices reduce | If per-Gbps capacity prices reduce **and/or** if capacity is utilised more |

Main Streaming models typically show a 30-50% cost saving for a customer moving from the consumption model to the capacity model. Why? Because a Streamer who delivers routinely – e.g., every day in a similar pattern like a national broadcaster, or frequently with higher capacity requirements like a live sports streamer – utilises capacity sufficiently to make it cost-effective.

An everyday analogy that illuminates the different models is how people pay for accommodation. Visiting a place for a night or a week would typically involve a hotel, at a high price per night for the irregular use and convenience. Living in a place for several months or even a few years, and remaining flexible because long-term plans are unclear, would typically involve renting a house or apartment because it is more cost-effective. Living in a place for many years because of a commitment to life in the area would likely involve buying a house or apartment because it is most cost-effective. There are different cost models and commitment levels in each scenario, but in principle the regular use and commitment to the area demands the most long-term economical solution.

Streamers have the same scenario. While streaming is being tested and is only a small part of the business it makes sense to "rent" delivery services on the variable cost consumption model. But when streaming is core to future business and must be leveraged for maximum profitability and economic return then it makes sense to "buy" delivery services or solutions on the capacity model.

A hybrid model is typically the best approach because a D2C Streamer generally has "run rate" delivery levels and "peaks". The run rate delivery can make excellent economic use of the fixed capacity and capacity-based model, while relying on a consumption-based model for peaks that do not make best economic use of fixed capacity.

**30-50%**

Cost saving for a customer moving from the consumption model to the capacity model.

# How to make streaming more sustainable

An intelligent media delivery platform must address our planet's need for sustainable solutions, otherwise it would not be intelligent enough.

Streaming is growing fast, and it is likely to be the dominant form of media consumption in the future. We must make every effort to reduce the energy consumption of our important content delivery networks, just like we must make every effort to reduce the CO2 emissions from our physical global product supply chains.

MainStreaming is a founding member of the Greening of Streaming, an initiative started to bring focus to this important subject. MainStreaming has two key focus areas to drive reductions in energy consumption, thereby reducing CO2e emissions. First, we architect content delivery networks for maximum efficiency. Second, we operate content delivery networks for maximum efficiency.

## Architecture

The architecture focus involves minimising the size and energy consumption of the content delivery network that is deployed. MainStreaming achieves this in three important ways.

- Ask once capability – the MainStreaming Edge requests content only once from the Origin. Once content is inside the Edge layer, it is distributed between Edges so that the backhaul bandwidth is reduced.
- Distributed Edge processing – instead of a hierarchical layered network of Edges and intermediate Edges, MainStreaming distributed its processing capabilities in a flat structure. Autonomous Edge processing and Edge intercommunication remove layers of hardware versus traditional designs.
- Video focus – video uses a lot of bandwidth and is the main consumer of internet capacity (and energy). But video is also a relatively simple form of data to deliver, having larger segment sizes than website delivery, for example. Hardware and software design can be optimized for video and use less energy as a result.

## Operations

The operational focus involves removing unnecessary workloads from the platform and ensuring the content delivery network does not consume unnecessary energy once it is deployed. MainStreaming addresses this need in two important ways.

- VOD assets, once inside the MainStreaming Edge layer, are distributed between Edges as required. But a file is only transferred and re-cached (i.e., stored, using energy) on a new Edge if it is clearly going to be watched multiple times based on popularity. Otherwise, the decision could be made to simply stream the file from the single Edge Cache that is holding the content to avoid a re-write of the data to storage.

- Edge caches are not used 100% of the time. Video viewing habits mean that most content is consumed in the evenings, with some significant but smaller midday peaks. But servers are always deployed, ready for the daily viewing audience, so how they consume energy while they are idle becomes the focus. MainStreaming works with its hardware partners to deliver up to 60% energy consumption reduction during times when the servers are not being used.

# The MainStreaming Vision

MainStreaming has built an advanced platform for delivering video to broadcast-grade standards with sustainability at its core. This platform - the MainStreaming Intelligent Media Delivery Platform – is the heart of the next-generation content delivery network architecture that delivers maximum performance and financial benefits to D2C Streamers and ISPs that are investing in streaming services. The world's leading D2C Sports Streamer is already benefitting from MainStreaming's technology and media-centric disciplines, as described in the following DAZN Case Study.

But this is not the end of the story. We are moving towards a world that requires at least an 80x capacity increase in streaming video delivery capacity. It could easily be a lot more than this if we consider augmented reality and virtual reality applications. We therefore must really focus on what it means to run a sustainable, broadcast-grade, streaming platform.

MainStreaming has a vision of highly scalable, highly efficient, high-performance, economically scalable video delivery based on continued developments in **five critical areas**:

## 1.Location data

We already use eDNS0 whenever possible to select the optimal Edge for delivering a stream. As eDNS0 zoning improves, we expect to be able to allocate streams to the best Edge to more granular levels, fine-tuning both performance and efficiency.

## 2.Deep edges

We already deploy Edges in multiple ISP PoPs (points of presence) deep inside ISP networks on a multi-ISP basis, working with leading D2C Streamers to optimise performance for their programming and their audiences. As streaming audiences grow, we expect to continue deploying Edges deeper into the networks, and to be able to use those Edges with more accurate location data.

## 3.Energy efficiency

We are proud of what we have achieved so far with reduced energy consumption from our different approach to CDN architecture and streaming operations and that all MainStreamin is carbon neutral. But this is only the beginning. Full server power-down capabilities must be considered, and we will continue working with our hardware partners to optimise processing methods.

## 4.Enhanced monitoring

MainStreaming already has a clear view of Edge delivery performance, with visibility upstream and downstream. As the last point at which video is processed before it enters a broadband or cellular network, the MainStreaming Edge makes important decisions to optimise performance. Integrating data directly from downstream platforms will ensure we have full knowledge of streaming performance before we make the final decision and deliver the video stream to the consumer.

## 5.Enhanced analysis

Expert insight and planning are the best ways to upgrade an existing content delivery network to improve performance. Every day we deepen our knowledge through data science, network design, and video streaming expertise. We codify knowledge into intelligent rules and algorithms and will continue researching how machine learning and artificial intelligence can support continuous improvement.

# #WeStreamTheFuture

MainStreaming has always had a vision of a world where streaming is the primary form of content delivery. It's why we set ourselves an ambitious mission in our tagline.

# The DAZN Case Study

## Introduction

DAZN has set a new standard for the media and sports industry, revolutionising the live sport viewing experience for fans around the world. In 2021, DAZN became the primary broadcaster for Serie A football in Italy - one of the largest sports streaming deals in history and a seminal tipping point in the migration of sports consumption to OTT. To assure the quality of delivery for the fans, DAZN pioneered a new technical and operational model for broadcast-grade live streaming at scale. DAZN Edge was born, powered by MainStreaming's specialised Private Edge CDN, and deployed with ISP collaboration, creating Italy's largest CDN.

## DAZN Edge, by MainStreaming

DAZN launched in Italy in 2018, with a minority package of Serie A rights. However, existing challenges with the local infrastructure and quality of OTT video delivery over the Italian networks convinced DAZN that it needed to formulate a strategy to not just enhance the viewing experience for Italy, but in all territories where network conditions are not ideal for live D2C streaming. This was the beginning of DAZN Edge.

In 2021, as the new Serie A rights window approached, DAZN's Distribution Technology Team had completed an intensive 2-year design, selection, and evaluation process to create a new broadcast-grade video streaming delivery platform, project named DAZN Edge. The requirements were clear: DAZN needed a solution that would enable it to collaborate with ISPs to deliver high quality streaming video to large national and regional audiences, and that would enable DAZN to directly control the speed and location of streaming capacity deployments for its audience. In addition, DAZN Edge needed to be quickly deployable on a global basis, be operated for DAZN as a managed service, and use a predictable cost model.

DAZN selected industry innovator MainStreaming and its commercially

available Private Edge CDN to underpin the DAZN Edge programme. MainStreaming won the competitive selection process for various reasons, including:

- Superior streaming performance when tested
- Engineering and operational agility to support DAZN's fast-paced business
- Relatively small hardware footprint for easy ISP deployments
- A specialised OTT video-centric operations model
- Years of experience operating a video-centric large-scale public CDN
- Innovative commercial models

MainStreaming's technical partner, Intel, played an important part in achieving the highest level of streaming performance by supporting the resource optimisation of MainStreaming's Intelligent Media Delivery Platform.

## Fast Deployment in Italy

In early 2021, the DAZN Edge programme was introduced to the rightsowners, Lega Serie A, the major Italian ISPs (TIM, Vodafone, WIND, Fastweb, Tiscali and EOLO), and the Italian regulator AGCOM. It was pitched as the specialised technical tool to overcome network congestion points and therefore assure the highest possible quality of experience for the fans. DAZN's business case for the Italian rights relied on supporting large, concurrent audiences. Delivering this in a streaming-only model was ambitious and collaboration with the ISPs was essential.

DAZN Edge was offered to the ISPs as a win-win solution, helping those without their own CDNs to deploy sufficient streaming capacity quickly and easily to improve the viewing experience of their own customers. For those with their own CDNs and the ability to integrate the DAZN App into their own video services, like TIM, it was a platform that delivered supplementary capacity quickly. DAZN Edge helped DAZN work with each ISP in the

way that worked best for them. The new DAZN Edge platform needed to be fully implemented from scratch to support the delivery of high-profile live video that was expected to require an amount of bandwidth significantly greater than total available public CDN interconnect capacity. In other words, the new DAZN Edge capacity and ability to deploy quickly inside ISP networks was necessary for DAZN to successfully stream Serie A matches. MainStreaming and DAZN worked with each of the partner ISPs to design the correct Private Edge CDN architecture for their networks, overcoming the fact that the ISP core networks were not optimised for latency-sensitive, large-scale live video delivery.

Within just 2 months, MainStreaming had designed and implemented Italy's largest dedicated CDN for DAZN. The DAZN Edge was deployed in tens of individual sites throughout Italy, more than doubling the capacity that had previously been available to DAZN, which was supplemented by DAZN's agreements with public CDNs and on TIM's own private CDN.

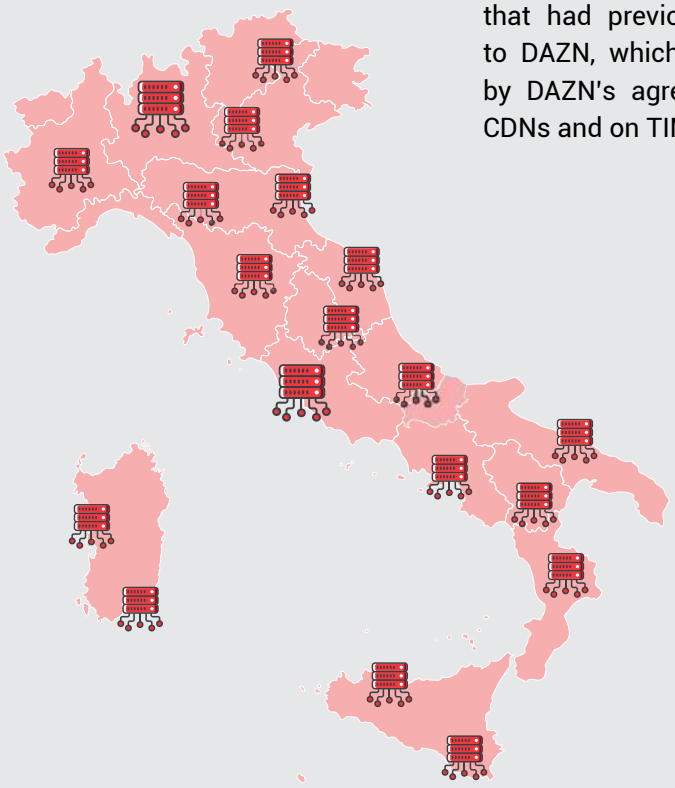

**Figure 2: MainStreaming Edge deployment by Region in Italy**

# Game-time

In August 2021, the focus turned to DAZN's streaming performance, with high levels of scrutiny to deliver a consistent high-quality viewing experience to football fans.

QoE (Quality of Experience) is a lead KPI for DAZN and operationally supported by MainStreaming, defined by a range of metrics including rebuffering ratio, bitrate consistency, latency, and start-up time. The deployment of DAZN Edge directly impacted all these KPIs in a positive way (see Figures 3, 4 and 5), driven by various factors:

- Distributed MainStreaming Edge architecture with intelligent video processing algorithms
- Dedicated server and port capacity for DAZN Edge that avoided congestion
- Strategic DAZN Edge locations inside the ISP networks that reduced core network congestion and delivered content to the fans from closer to their own location.

This distributed Private CDN Edge model performed especially well for bitrate consistency and latency, two KPIs that ranked highest in importance for the fan experience.

## Rebuffering

Fig.3

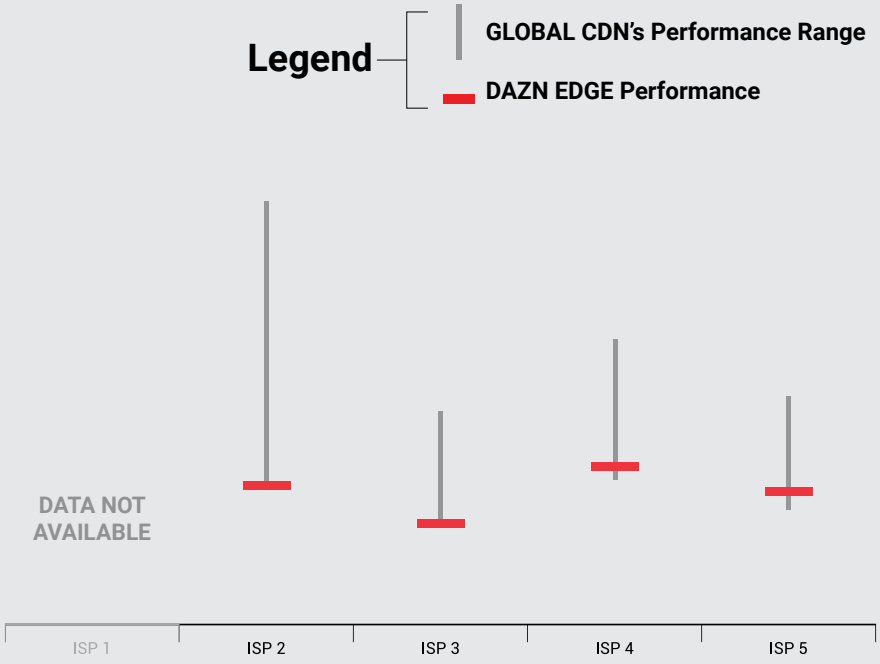

## Average bitrate

Fig.4

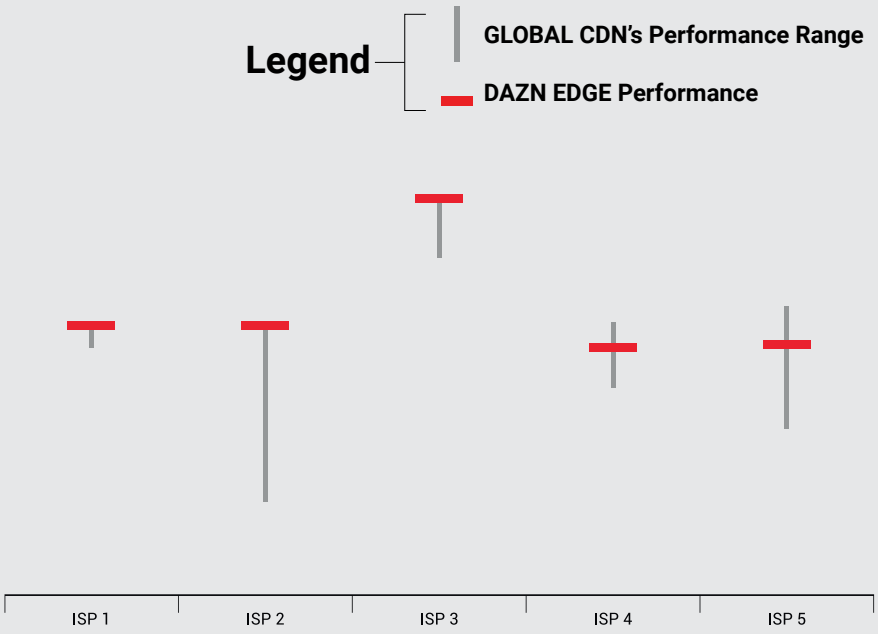


## Video start time

Fig.5

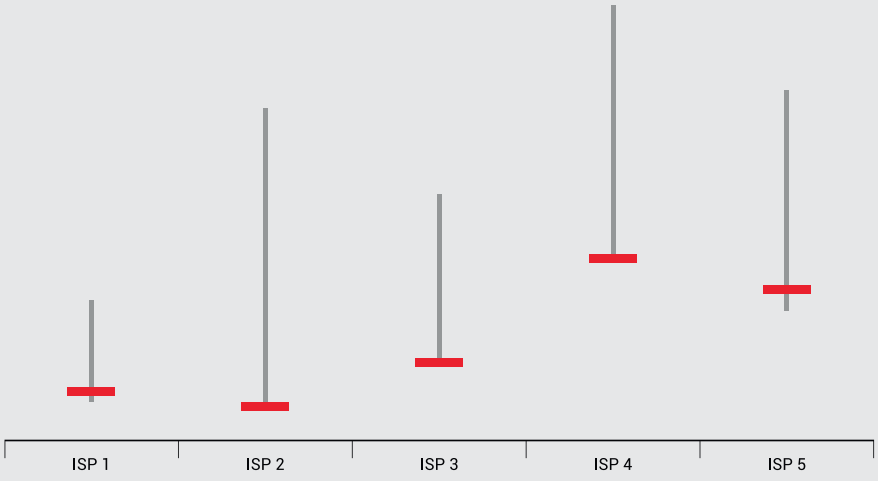


While theoretically the DAZN Edge's ability to remove congestion from ISP core networks was a key design principle, the effect needed to be observed in practice. Not only did the audience served directly by DAZN Edge observe improved results, but in the highest peak moments for the biggest matches when all types of CDN capacity was utilised, every fan received improved performance. The DAZN Edge's "total network decongestion effect" helped improve the performance of all content delivery networks utilised by DAZN in Italy and was a big win for DAZN and the ISPs.

# Post-game analysis

DAZN and MainStreaming have created a new streaming model in close collaboration with the leading ISPs. DAZN Edge is the result of a vision for broadcast-grade live streaming that is necessary for DAZN to succeed as a live-centric D2C Streamer.

As DAZN considers future service enhancements with 4K, HDR, immersive viewing experiences, additive business ventures like gaming and betting, securing rights in other markets, and expanding viewership in existing markets, the ability to cut through the internet's natural congestion and its natural speed of expansion is critical to success. DAZN Edge by MainStreaming creates a technical, operational, and commercial model that supports DAZN's path to becoming and remaining a world leading sports streaming service.

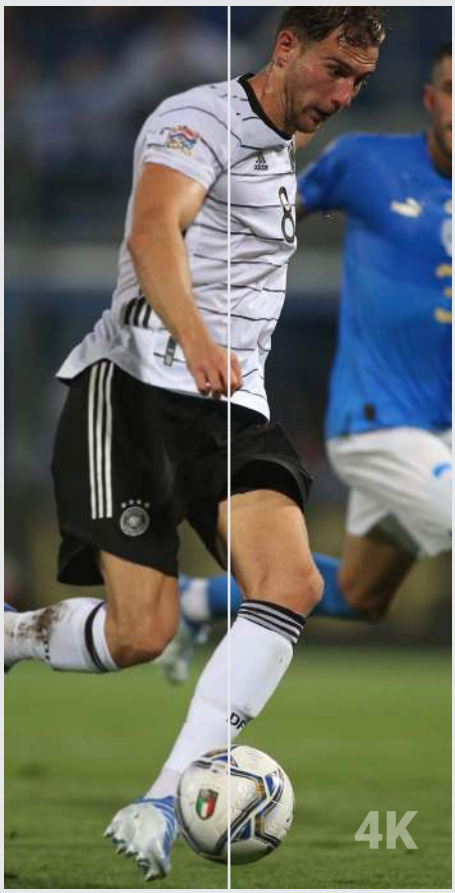

# Conclusion

OTT video streaming can reach broadcast-grade standards, and streaming can even become the new gold standard in content delivery.

Streaming is growing all the time. Even in the most mature markets, peak audiences could grow by 10x in the coming 5-10 years. 80x can be considered as a growth figure when we think about the whole world and the standard resolution changing from HD to UHD.

A sustainable, cost-effective, high-performance solution is required for this vision to become reality. And the delivery solution must also be manageable by Streamers that must first focus on monetizing their content creation, acquisition, and production. Distribution technology, which dominates the OTT conversation, should be left to the specialists who know how to manage and guarantee broadcast-grade delivery.

It's not clear how long it will take for streaming to become the dominant form of video delivery and consumption. It could be 4 years, when we reach the next FIFA World Cup, or it could be in 10-20 years as population demographics change. It is also important to remember that different countries are moving at different speeds according to their own needs.

**While the growth of streaming could be held back by poor quality, high distribution cost, and an inability to scale well for the biggest live events, actually there are solutions to all these problems already. They just need to be implemented to replace legacy designs and business models.** "MainStreaming is already delivering the new solution to leading Streamers like DAZN."

One day, broadcast-grade streaming is likely to be the normal state. For now, as OTT services expand rapidly, it is not. There are constant pressures to perform better. MainStreaming is focused on helping Streamers to deal with this pressure and deliver the right results.

***We thank you for reading this whitepaper, and would be delighted to hear from you with questions or comments.***

Please email us at
*marketing@mainstreaming.tv*
or contact us at this *URL* for more information or to share your feedback.

# The Author

**Paul Martin**
**Principal Consultant**

Paul has a Technology sector background in semiconductors, broadband, broadcast, and streaming, with deep experience in operations planning, change management, business strategy, and general management. This breadth and depth of knowledge creates a unique perspective on how the future of video delivery will evolve technologically and commercially to support growing demands from consumers.

Paul is also an OTT Contributing Editor for The Broadcast Bridge, a thought leadership publication with global readership from across the broadcasting and streaming industry.

He has published over 40 articles about OTT technology and business models, including subjects such as content personalisation, ultra-low latency delivery, video streaming infrastructure, and streaming's sustainability credentials.

# Contributors

**Giovanni Proscia**
**CTO - Technology Founder**

*15+ years of Tech experience in IT and Telecommunications*

**Philippe Tripodi**
**CPO - Product Founder**

*Entrepreneur and Innovation Expert*

**Sergio Carulli**
**CIO - Innovation**

*25+ years of Tech experience and Former IT Professor at Milan University*

**Marco Inzaghi**
**Head of Business Development**

*15+ years of Tech experience in IT and Telecommunications*

**Antonio G. Corrado**
**CEO & Founder**

*Entrepreneur and Tech Expert*

***“Live video streaming at scale is what we are born for„***

MAINSTREAMING®

# We Stream The Future!