

The Future of TCO

How Composable, AI-Native OTT Platforms
Drive Real Business Outcomes.

A report by:  merapar |  accedo



Content

1. Executive Summary	pg 3
2. From Spend to Outcomes: Redefining TCO in the OTT Era	pg 8
3. Lessons From D2C and E-Commerce: The Composability Advantage	pg 14
4. Why OTT is Lagging—And How to Catch Up?	pg 21
5. Laying the Groundwork for Composable, AI-Native OTT	pg 25
6. AI-Native Orchestration: The Next Frontier in Cost Optimization	pg 29
7. Why This Matters: The Future for OTT	pg 35

Most **OTT** services are **struggling** with mounting economic **pressure**.

Content acquisition costs are rising, churn is difficult to contain, and user expectations continue to evolve faster than many platforms can adapt. In this environment, media companies are forced to reassess how they allocate resources—not just to maintain service quality, but to drive sustainable growth.

This report is for technology and business leaders tasked with navigating this increasingly complex landscape. It introduces a new strategic lens on ‘Total Cost of Ownership’ (TCO); one that moves beyond static accounting models and toward a dynamic, outcome-driven framework. It draws on lessons from adjacent industries like D2C and e-commerce, where modular technology stacks and AI-native orchestration are already redefining how platforms scale, optimize, and compete.

The future of TCO will depend on a few core principles.

Composability as a strategic enabler

Traditional OTT platforms rely on tightly coupled systems that may be cost effective to acquire but expensive to maintain and slow to evolve. Composable platforms, by contrast, allow for modular assembly of capabilities, enabling faster iteration, greater flexibility, and reduced duplication. This architectural shift lowers integration costs and unlocks agility.

AI orchestration & agent automation

AI is becoming a foundational layer in platform design. Agent-based orchestration uses intelligent automation to optimize user journeys, triage operational issues, and personalize experiences in real time. This reduces manual effort, accelerates time-to-insight, and directly impacts business KPIs like churn, ARPU, and LTV.


TCO as a living and evolving system

Static TCO models, focused on CapEx and OpEx tradeoffs, fail to capture the full picture in today's dynamic OTT environment. A modern approach to TCO requires a continuous feedback loop between data, architecture, and experience. Platforms must become self-aware, adjusting operations in real time to meet business goals.

OTT services that **succeed** over the next five years will be the ones that **orchestrate** the best. The future of TCO lies not in **reducing** cost alone, but in maximizing the value of each dollar spent. That means designing platforms for **adaptability**, enabling AI to continuously **optimize** the service, and measuring **success** based on outcomes—not features.



What this report covers





TCO is no longer just a cost metric. It has become a lens through which to **design**, **operate**, and **scale** the entire business.

With the right **architecture** and orchestration **strategy**, OTT providers can reduce waste, **accelerate** growth, and unlock long-term **competitive** advantage.

From **spend** to **outcomes:** **Redefining** TCO in the OTT era.

In a slower-moving video streaming environment, treating TCO as a static metric made perfect sense. But OTT platforms today operate in a dynamic, high-stakes market where agility, speed, and adaptability are essential. The traditional model of TCO is no longer fit for purpose.

The old model

Build vs Buy


Historically, media companies made large infrastructure decisions based on multi-year ROI projections. Technology stacks were monolithic, and decisions around software, hardware, and operations were often siloed across departments.

TCO calculations focused on total outlay over time, including licensing, hosting, maintenance, and internal and external resourcing.

CapEx vs OpEx

This approach worked when platform change was incremental. But today's OTT environment is shaped by constant evolution—shifting consumer behavior, fragmented attention, and relentless competition.

A static cost model fails to capture the full picture, especially when the true cost of missed opportunity, slow time-to-market, or misaligned product decisions is so high.



Why static models fail today

The core flaw in traditional TCO thinking is its rigidity. It assumes a fixed set of inputs and outputs, ignoring the volatility of today's OTT landscape.

Technology evolves rapidly, and what is “efficient” today may be obsolete tomorrow.

User expectations shift, and platforms must respond in real time to retain engagement.

Competitive pressure intensifies, making slow iteration cycles a business risk.

As a result, platforms optimized for the lowest upfront cost often struggle with long-term adaptability. What's needed is a new way to think about TCO—not as a snapshot, but as a system in constant motion.

TCO as a strategic lever

Forward-looking OTT leaders are beginning to treat TCO as a lever for growth. This requires a shift in mindset: from managing spend to orchestrating outcomes. That shift rests on three pillars:

- 1 **Orchestration over optimization**
- 2 **Agility through modularity**
- 3 **Continuous optimization via data**

Orchestration over optimization

Cost efficiency is shifting the focus from simply cutting budgets to strategically aligning spending with performance. Orchestration, particularly when powered by AI, enables platforms to continually tune operations, workflows, and experiences to maximize return on investment.

Agility through modularity

Modular, composable architectures reduce the cost of change. Instead of rebuilding systems or forcing integration workarounds, teams can plug and play capabilities as needed. This reduces waste, enables faster experimentation, and insulates the platform from long-term technical debt.

Continuous optimization via data

Real-time telemetry, user behavior signals, and operational metrics make it possible to monitor and adjust performance continuously. Instead of relying on static roadmaps or quarterly reviews, platforms can use live data to steer investment decisions.

In this new model, TCO behaves more like a living system—**dynamic**, **responsive**, and **deeply connected to business outcomes**.

- When churn risk rises, orchestration agents can trigger retention workflows or UI changes.
- When a feature underperforms, analytics can reallocate attention to higher-impact areas.
- When operational costs spike, automation can identify and resolve root causes in real time.

With every system element connected, cost becomes a reflection of platform intelligence, agility, and alignment, not simply an input.

Linking cost to value

Perhaps most importantly, this modern approach to TCO connects spending to outcomes. Traditional models only measure input (what was spent); the new model also measures output (what was achieved). For OTT services, this could mean tracking:



- **Churn reduction per dollar spent on retention tooling.**
- **Customer lifetime value per dollar spent on personalization.**
- **Time-to-market acceleration per dollar spent on orchestration.**
- **Time to resolution of technical issues.**

These are not theoretical metrics—they are increasingly accessible in platforms that embrace composable design and AI-native orchestration.

So, why does this matter?

The shift from **static** to **dynamic** TCO models is a **direct response** to the **realities** of today's **OTT landscape**.

Service providers face growing pressure to do more with less: to innovate faster, retain users longer, and operate leaner. Relying on traditional cost frameworks in this environment proves not only inefficient but also a competitive risk.

Treating TCO as a living system enables OTT leaders to respond to volatility with precision. It connects investment to impact in real time, ensuring that spend is continuously justified by outcomes.

Without this shift, companies risk falling behind—trapped in rigid systems, unable to adapt quickly enough to market signals or user expectations.

The opportunity is clear: by **rethinking TCO** as a **strategic lever**—driven by **orchestration, modularity, and data**—OTT platforms can move beyond cost control and toward continuous value creation.

Lessons from **D2C** and **e-commerce**: The **composability** advantage.

While many OTT providers are still struggling to modernize aging architectures, adjacent industries—particularly direct-to-consumer (D2C) brands and e-commerce platforms—have already undergone a major transformation. Their advantage? A clear shift toward composable platforms, modular tech stacks, and AI-native orchestration. These industries provide a valuable blueprint for OTT providers looking to achieve agility, reduce cost, and accelerate growth.

How D2C brands rewrote the playbook

D2C companies rose to prominence by bypassing traditional retail channels and building direct relationships with consumers. Their success, however, comes from technology as much as it does from distribution.

D2C leaders were early to understand that modularity and flexibility are essential to staying close to customers.

Instead of building monolithic platforms, they adopted composable tech stacks: content systems, commerce engines, personalization tools, data layers, and engagement platforms—all loosely coupled and easily replaced. This allowed them to:

- Launch features and campaigns quickly.
- Swap out underperforming components without large-scale rebuilds.
- Align technical delivery with customer and business KPIs in real time.

More importantly, D2C brands focused on outcomes over features. Rather than chasing roadmap milestones, they tracked metrics like conversion, retention, and lifetime value. Their technology choices were governed by what moved those metrics, not by what added surface-level functionality.

This shift reshaped their cost structures. Instead of locking capital into large, fixed systems, they budgeted for modular pilots, iterative releases, and targeted optimizations. The result: reduced customer acquisition cost (CAC), higher retention, and greater ROI on digital spend.

E-commerce platforms and AI-native orchestration

The e-commerce sector, especially at the enterprise level, has gone even further by embedding orchestration and automation into the core of the user journey.

Take Amazon, for example. Its personalization engine isn't a siloed feature—it's an orchestrated layer that continuously adapts based on user behavior, inventory, location, and context.

Product recommendations, search results, promotional offers, and even delivery options are dynamically adjusted.

These optimizations are not manual—they're agent-driven, enabled by real-time data and AI. This architecture enables Amazon to:

- Reduce cart abandonment through personalized nudges.
- Optimize logistics and fulfillment for cost and speed.
- Continuously test and adapt user flows without major releases.

Shopify, though serving a very different segment of the market, is also a good example of how composability can scale. Its success lies in its modular app ecosystem, where thousands of businesses can extend their commerce stack with purpose-built tools—from inventory management to checkout optimization to AI-powered marketing. This approach gives merchants using the platform a competitive edge while insulating Shopify from the burden of feature bloat.

Nike: a case study in D2C transformation

Nike's shift to a D2C model is a textbook example of composability in action. Facing **margin pressure** and an **evolving retail landscape**, Nike chose to **prioritize direct relationships with customers over wholesale distribution**.

Behind the scenes, the company restructured its technology to support this shift:


- Migrate to modular, cloud-native services.
- Integrate AI-driven analytics to personalize content and offers.
- Embed dynamic orchestration across app, web, and retail touchpoints.

With this new stack, Nike was able to unify customer data, launch tailored experiences across markets, and react in near real time to changes in demand or behavior. The results included stronger brand loyalty, faster product cycles, and higher per-customer revenue—all while maintaining cost efficiency through automation and reuse.



Strategic outcomes: efficiency through modularity

Across these examples, a clear pattern emerges. Composability is not just about technical flexibility—it is an entire business strategy. The organizations that adopted it achieved measurable outcomes:




Reduced waste

By eliminating duplicate systems and unnecessary rework.



Faster iteration

Through decoupled services and agile experimentation.



Increased retention

Enabled by personalization, relevance, and responsiveness.

Lower total cost

Through efficient operations, and data-driven automation.

These outcomes were not achieved by adding more features, but by rethinking how digital platforms are designed and operated.

Implications for OTT

OTT services, especially those operating at scale, face many of the same challenges that e-commerce and D2C brands once did: fragmented customer journeys, rising churn, long release cycles, and unsustainable operational costs.

Adopting a composable approach can unlock the same advantages seen in D2C and e-commerce:

Ability to respond faster to changes in user behavior, where experiences are adjusted in real time.

Lower integration overhead, where components can be swapped or scaled with minimal disruption.

Smarter investment decisions, where cost is aligned with actual business outcomes.

AI-native optimization, where automation drives retention, engagement, and quality.

The gap is now strategic, rather than technical. OTT providers must stop thinking of their platform as a fixed asset and start treating it as a living system, one that evolves continuously based on data and demand.

So, what does this mean?

The lessons from D2C and e-commerce are clear: platforms that **prioritize modularity, data activation, and outcome-driven orchestration** gain an outsized advantage. They reduce cost not by cutting corners, but by removing inefficiencies and focusing resources where they matter most. OTT providers have an opportunity to leapfrog legacy limitations.

By adopting a composable, AI-native foundation, they can move faster, spend smarter, and deliver better experiences—all while building a more agile, scalable business.

The path forward isn't about copying features—it's about adopting the mindset that made these industries successful. With the right architecture in place, OTT services can operate with the same speed, intelligence, and efficiency that define today's most successful digital businesses.

Why **OTT** is **lagging**—and how to **catch** up?

While adjacent industries have modernized through modular platforms and AI-native orchestration, much of the OTT ecosystem remains stuck in an earlier phase of digital maturity. Technical debt, fragmented stacks, and legacy decision-making models are creating unnecessary cost, complexity, and ultimately customer churn.

To compete in today's environment, OTT providers must recognize what's holding them back—and take decisive steps to close the gap.



Fragmented architectures create drag

Most OTT platforms have grown through a patchwork of internal builds, vendor solutions, and one-off integrations. Over time, this leads to architectural sprawl: separate systems for video playback, user identity, analytics, personalization, content management, and monetization.

These components often operate in silos, with limited interoperability and redundant functionality. The result is a platform that's hard to evolve, expensive to maintain, and slow to adapt. Every new feature or integration adds incremental complexity. Teams spend more time managing dependencies than delivering value.

This fragmentation directly impacts cost. Engineering effort is duplicated across teams and markets. Operational overhead increases. **Time to market slows down—often measured in quarters rather than weeks. In contrast, composable platforms allow services to be swapped, scaled, or upgraded independently.** That kind of flexibility is difficult to achieve with tightly coupled legacy stacks.

Lack of unified data and real-time visibility

A modern OTT experience is only as good as the data that powers it. Yet many providers still lack a unified view of their users, content, and performance. Data is collected, but it's scattered across systems—billing platforms, CDNs, engagement tools, customer support systems—each holding a partial picture.

Without integrated data pipelines and real-time visibility, personalization becomes guesswork. Product decisions rely on lagging indicators or anecdotal feedback. Business teams can't easily connect user behavior with outcomes like churn, lifetime value (LTV), or upsell conversion.

This blind spot leads to missed opportunities: poor content recommendations, untargeted promotions, generic user journeys—all of which erode engagement and retention.

Other industries have already moved past this. E-commerce platforms, for example, operate on unified, real-time datasets that fuel personalized journeys, optimize spend, and adapt to user behavior in-flight. OTT needs the same level of data orchestration.

Over-investing in features, under-investing in outcomes

Feature delivery has long been the dominant operating model in OTT. Roadmaps are built around shipping new capabilities—offline viewing, search filters, multi-user profiles—without tying those features back to clear business results.

This mindset creates a surface-level sense of progress while masking deeper issues. Features are shipped, but churn doesn't improve. Cost goes up, but quality doesn't stabilize.

Users have more options, but not necessarily better experiences. The real question isn't what is being built—it's why. Without a framework to link technical delivery to measurable outcomes, OTT providers risk spending large portions of their budget on low-impact initiatives.

In contrast, outcome-driven organizations start with the KPI and work backward. If the goal is to increase LTV, they model the drivers: time to value, engagement depth, upsell readiness. Every product or operational decision is evaluated against its contribution to that metric. OTT providers need to adopt this same discipline.

Disconnected product and financial decision-making

One of the more persistent challenges in OTT is the disconnect between product roadmaps and financial accountability. Product teams are often measured by feature delivery or NPS, while CFOs focus on operating margins, churn rates, and ARPU. As a result, budget allocation is rarely guided by a shared understanding of value.

This disconnect creates friction and waste. Cost-reduction efforts focus on vendor negotiations instead of architectural improvements.

Monetization strategies are bolted on rather than designed into the user journey. Experimentation is limited by fear of destabilizing a fragile tech stack.

To catch up, OTT companies need to unify how they think about performance. Product, engineering, and finance leaders must align around shared KPIs—whether that's churn reduction, customer lifetime value, or quality of service. This alignment enables smarter investment, faster iteration, and clearer accountability.

So, how can you make the shift?

Fortunately, the **pain points** described above **are not permanent**. They are symptoms of an **outdated operating model**—and they are **solvable**.

The solution lies in composable OTT platforms: systems built from modular, interoperable services that can be orchestrated dynamically. In a composable architecture, each component—whether it's user identity, content curation, playback, or analytics—functions as an independent service.

This enables faster experimentation, localized deployments, and smarter reuse of functionality across markets and devices.

Critically, composability isn't just about how software is built—it's about how the entire platform is managed. The real unlock comes when composable architectures are paired with AI-native orchestration: agents that monitor, adapt, and optimize the user experience in real time.

Laying the groundwork for composable, AI-native OTT.

Understanding why OTT services must evolve is only part of the equation. The real challenge lies in the how. For many media companies, legacy infrastructure, siloed teams, and reactive product cycles make the shift to composable, AI-native architectures feel daunting. But transformation doesn't have to start with a full rebuild.

Instead, OTT services can take clear, practical steps to prepare their organizations for this shift—technically, operationally, and culturally. This section outlines the foundational work needed to ensure a successful transition toward a more agile, intelligent, and outcome-driven model.

Audit your current architecture for modularity gaps

The first step is to understand how modular your current platform really is by conducting a focused audit of your technology stack:

- Where are the tightly coupled dependencies?
- Which components (e.g. CMS, payments, personalization, analytics) are interchangeable, and which are locked in?
- How easily can new features be swapped in or out without triggering regression elsewhere?

The goal of this audit isn't just to catalog tools—it's to assess composability readiness.

Once complete, teams should identify priority areas for decoupling—starting with systems that are high-impact and high-friction.

Unify and enrich your organization's data layer

Composable architectures rely on shared, unified data to function intelligently. **If your**

customer, content, and behavioral data are spread across disconnected systems, AI-driven orchestration will struggle to deliver real value. The steps to take

in order to create a unified data layer include:

- Creating a centralized data strategy with clear ownership and governance.
- Investing in tooling that enables real-time access and enrichment of data across systems.
- Mapping key customer journeys and identifying where data signals are missing or underutilized.

This is not just about implementing machine learning models—it's about creating a solid foundation where every agent, module, and team can work from a shared, up-to-date view of the customer.

Establish a composability and orchestration mindset

Technology alone won't drive transformation.

Teams must learn to think in terms of outcomes, orchestration, and agility, rather than feature lists and static releases.

To shift your organization's internal mindset, you will have to:

- Redefine product KPIs around business outcomes (e.g. churn reduction, trial conversion, ARPU uplift).

- Encourage product and engineering leads to design services as reusable components.
- Break the cycle of quarterly roadmap "commitments" and shift toward goal-based planning.

Shifting mindsets is easier said than done. You are wise to support the process with ample training, incentives for cross-functional collaboration, and methods to ensure executive alignment on making the leap from control to composition.

Pilot agent-based workflows in targeted use cases

One of the most effective ways to build internal confidence in AI-native orchestration is to start small. Choose a contained but impactful use case—such as churn prediction, content recommendations, or automated QA—and deploy a pilot agent to handle it. Choose your pilot case by looking for areas where:

- Manual effort is currently high.
- Speed is critical to success.
- There is a direct line to a business KPI.

Run the pilot with clear success metrics, and make sure both technical and business stakeholders are involved in measuring impact. **These**

early wins not only prove value but help develop the operational muscle needed to manage agent-based systems.

Modernize governance to enable greater flexibility

Composable systems and AI agents thrive in environments where speed, experimentation, and continuous improvement are encouraged. **But many OTT companies still operate under outdated governance models built for slower, release-based cycles.**

That said, modernizing governance doesn't mean abandoning discipline—it means reframing it by:

- Approving modular experiments based on risk and impact, not size.
- Shifting compliance and QA toward observability and real-time monitoring.
- Empowering cross-functional teams to own outcome delivery, not just task completion.

This cultural shift is often the most difficult, but also the most transformative.

It allows technical progress to move at the speed of business needs. The move to composable, AI-native OTT requires a fundamental rethinking of platform construction, optimization, and evolution, extending beyond mere technology migration.

The five steps above offer a pragmatic way to begin the journey. Each one creates the conditions for AI orchestration to thrive, for costs to become more manageable, and for outcomes to take center stage.

AI-native orchestration: The next frontier in cost optimization.

The promise of composable architectures is flexibility—but flexibility alone doesn't create value. To unlock measurable business outcomes, OTT providers need a system that not only supports modularity, but also understands how to act on it.

That's the role of AI-native orchestration.

What is AI-native orchestration?

AI-native **orchestration** is the layer that sits atop a **composable** OTT platform. It connects services, activates **real-time data**, and **automates decisions** across the **end-to-end** user journey.

But unlike traditional rules-based workflows, this layer is powered by agents: autonomous AI systems trained to optimize for specific outcomes like retention, LTV, or operational efficiency.

In short, orchestration turns static roadmaps into living, adaptive systems. It enables the platform to respond continuously to changes in user behavior, content performance, and business priorities—without requiring manual intervention at every step.

Where composability offers structural efficiency, orchestration drives intelligent automation.



From workflows to agents: a new operating model

Traditional OTT operations rely on static workflows—manually defined triggers and business rules that govern everything from onboarding to recommendation logic. These workflows are often fragile, inflexible, and expensive to maintain across multiple markets or audience segments.

AI-native orchestration replaces this model with outcome-driven agents.

These agents operate like specialized digital employees: each trained to monitor a particular aspect of the user

experience, respond to signals in real time, and make decisions that optimize for a desired KPI. Examples of agents include:

- Retention agents that detect churn risk signals and trigger personalized interventions.
- Onboarding agents that adapt welcome flows based on user context and device.
- LTV optimization agents that test and refine upsell journeys continuously.

- Operational agents that triage platform issues, run diagnostics, or auto-resolve support tickets.

These agents don't just automate tasks—they learn and improve over time, creating a system that becomes smarter and more efficient the longer it runs.

Use cases across the OTT lifecycle

AI-native orchestration delivers measurable value across the full video streaming user journey. These are a few high-impact examples.

Reducing churn with real-time retention agents

Instead of waiting for post-churn surveys or lagging indicators, retention agents can monitor early signs of disengagement: declining watch time, skipped titles, payment friction. When risk increases, the agent can personalize interventions—a free trial extension, relevant content nudges, or an in-app prompt to re-engage. This moves churn management from reactive to proactive—and often prevents loss before it happens.

Increasing LTV through personalized journeys

Rather than offering the same upsell or recommendation to every user, LTV agents test variations based on cohort behavior, content preferences, and engagement timing. Over time, these agents learn which journeys yield the highest conversion—and optimize accordingly. This increases revenue per user while improving the user experience itself.

Reducing operational cost through automated QA and support

AI agents can monitor platform performance across devices, geographies, and user segments in real time. When errors occur—buffering spikes, failed logins, content mismatches—they can automatically log issues, assign priority, and trigger fixes or reroutes. In customer support, similar agents can auto-respond to common issues, escalating only when human input is truly needed. These capabilities reduce headcount dependency for repeatable tasks and improve service consistency.

The strategic impact on TCO

AI-native orchestration transforms cost from a static burden to a dynamic lever. It reduces overhead, avoids waste, and improves performance—while enabling faster iteration and higher responsiveness.

More importantly, it aligns technology effort directly with business outcomes. Instead of investing in broad feature sets and hoping for ROI, providers can target specific KPIs—churn, LTV, ARPU—and design their orchestration strategy to deliver them.

Here’s how orchestration directly improves the components of TCO:

TCO Driver	Traditional Approach	AI-Native Orchestration
Feature delivery	Manual and roadmap-driven	Outcome-driven and adaptive
Ops/support	Reactive and resource-heavy	Automated and self-healing
Personalization	Rules-based and generalized	Realtime, predictive and contextual
Churn management	Lagging and post-facto	Proactive, realtime interventions
Experimentation	Slow, human-run A/B testing	Agent run multivariate testing

The result is a system that can self-tune—not just to deliver services efficiently, but to evolve them continuously based on what’s working.

So, what comes next?

AI-native **orchestration** lays the foundation for a **smarter**, more **efficient** operating model. But the next step is what matters most: **aligning** these orchestration **capabilities** with your **business goals**.

This means moving beyond one-off automation use cases and starting to invest in outcomes—whether that's reducing churn, increasing LTV, lowering operational overhead, or improving service quality.

It requires treating KPIs not just as reporting metrics, but as design inputs that inform how agents are deployed, how platforms are composed, and how priorities are set.

The shift from features to outcomes doesn't happen overnight. But with composability and orchestration in place, the infrastructure is ready. What's needed now is a new approach to investment planning and roadmap design—one that focuses not on delivering more, but on delivering what moves the business forward.

Why this **matters**: The **future** for **OTT**.

The OTT industry stands at a critical inflection point. Cost pressures are intensifying, user expectations are rising, and the limits of traditional operating models are becoming harder to ignore. What once worked—monolithic platforms, linear roadmaps, feature-heavy development cycles—is no longer enough to compete, let alone grow.

This report has argued that the future of Total Cost of Ownership is not about spending less. It's about spending smarter.

That shift requires a new mindset. One where TCO is not a line item on a finance sheet, but a reflection of how well your architecture, your data, and your teams are aligned to deliver business outcomes. It means understanding cost not as a fixed burden, but as a dynamic system that can be optimized—continuously, intelligently, and in real time.

The takeaway is clear: composable, AI-native platforms give media companies the tools to make that shift.

- Composability creates the structural agility to evolve quickly.
- Orchestration makes that agility intelligent, outcome-focused, and scalable.
- AI agents transform optimization from a one-time task into an ongoing capability.

But the opportunity is time-sensitive.

Other industries—e-commerce, D2C, fintech—have already adopted modular stacks and AI-driven operations to reduce customer acquisition costs, increase loyalty, and drive profitability. Many OTT providers, in contrast, are still building roadmaps around features instead of results, or relying on teams to manually stitch together data, tools, and workflows that should already be connected.

The cost of inaction is more than technical debt. It's a missed opportunity: lower retention, slower time to market, and an inability to prove ROI on innovation spend.

Why this matters now is because the foundational change is already underway. The question is no longer if the industry will adopt composable, AI-native models—but who will do it first, and best.

Those who lead will see tangible results:

- Higher operational efficiency through intelligent automation.
- Faster experimentation and time-to-value through modular delivery.
- Better financial performance through outcome-driven orchestration.

Those who lag will continue to struggle with brittle systems, rising costs, and declining engagement.

The **future** of OTT will belong to companies that can **evolve** in real time. That starts by **rethinking** how **value** is **created**—and how **cost** is **controlled**. It starts by **transforming** your platform into a living system.

About the authors of this report



Merapar specializes in cloud and hybrid software engineering, focusing on media & entertainment while serving other sectors. We deliver scalable, (AI-powered) platforms that speed up time-to-market, simplify operations and cut costs, enabling clients to innovate with confidence

Website
merapar.com



Accedo is a global provider of award-winning software and services spanning the entire video streaming lifecycle - from strategy, UX design and engineering to quality assurance, delivery and ongoing service optimization. The company offers flexible engagement models ranging from SaaS to end-to-end managed services.

Website
accedo.tv