

Retention Dynamics:

Competitor Disruption Cohort Analysis

*Linking User Behavior, Acquisition Quality, and Customer Feedback
to Drive Targeted Retention Strategy*

Business Context: Acquisition Spike Driven by Competitor Disruption

CONTEXT



Competitor Content Disruption

A major content dispute caused a rival streaming platform to lose access to key programming, including high-demand live sports, creating a window for new subscriber acquisition.



Rapid User Acquisition (~280K subscribers)

Our platform saw a sudden influx of new users seeking alternative access. This was a significant short-term growth event with an uncertain long-term profile.



Core Business Uncertainty

Were these users genuine long-term subscribers, or event-driven opportunists who would leave once the dispute resolved and the competitor restored content?

THE CORE QUESTION

"The question wasn't just retention. It was whether this growth translated into durable long-term value."

~280K

New subscribers
acquired

~62%

(~174K subs)
Classified as
event-driven

~20.8%

(~58K subs)
Churned within
5 months

** All figures are synthetic approximations preserving original analytical patterns*

Analytical Framework: Three Independent Lenses on the Same Problem

APPROACH

To determine whether this was durable or transient growth, I designed a multi-layer analytical framework using three complementary approaches:

01 Behavioral Analysis



Cohort-based survival curves tracking month-over-month retention across the full acquired subscriber base

WHAT is happening

02 Segmentation Analysis



Retention decomposed by product tier, acquisition channel, and content engagement behavior

WHERE it's happening

03 Customer Feedback (LLM)



LLM classification pipeline applied to unstructured call transcripts to extract and quantify primary churn drivers

WHY it's happening

All three lenses independently converged on the same conclusion, which gave us the confidence to move from diagnosis to action.

Overall Retention: A Sharp Early Drop Signals Event-Driven Behavior

BEHAVIORAL

Survival Rate (%) of Subs Acquired



Key Pattern

Steepest drop occurs between M1 and M2, consistent with event-driven users disengaging once their immediate content need is fulfilled.

99.7%

Survival at Month 1
(~279K subs retained)

86.0%

Survival at Month 2
(~241K subs retained)

79.2%

Survival at Month 5
(~222K subs retained)


~20.8%

Total churn by Month 5
(~58K subs lost)

The M1 to M2 drop of ~13.7 percentage points (~38K users lost in a single month) is the critical signal. This is not gradual churn, it reflects a cohort that was never fully committed.

Product Tier: Core Package Churn Runs Nearly 2x Premium Tiers

SEGMENTATION



69.5%

Core Survival at M5
(~126K subs)

84.8%

Select Survival at M5
(~70K subs)

87.9%

Premium Survival at M5
(~56K subs)

85.4%


Elite Survival at M5
(~28K subs)

The ~30-point gap between Core and Premium suggests a content-product mismatch. These users didn't find sustained value, not just lower intent.

Acquisition Channel: Digital Channels Bring Lower-Intent Users

SEGMENTATION

Survival Rate by Acquisition Channel (%)



M5 Survival Gap

Assisted: 89.9% (~84K subs)
Digital: 78.8% (~196K subs)
Gap: ~11.1 pts

Why Assisted Performs Better

In-person acquisition creates higher friction that naturally qualifies intent.

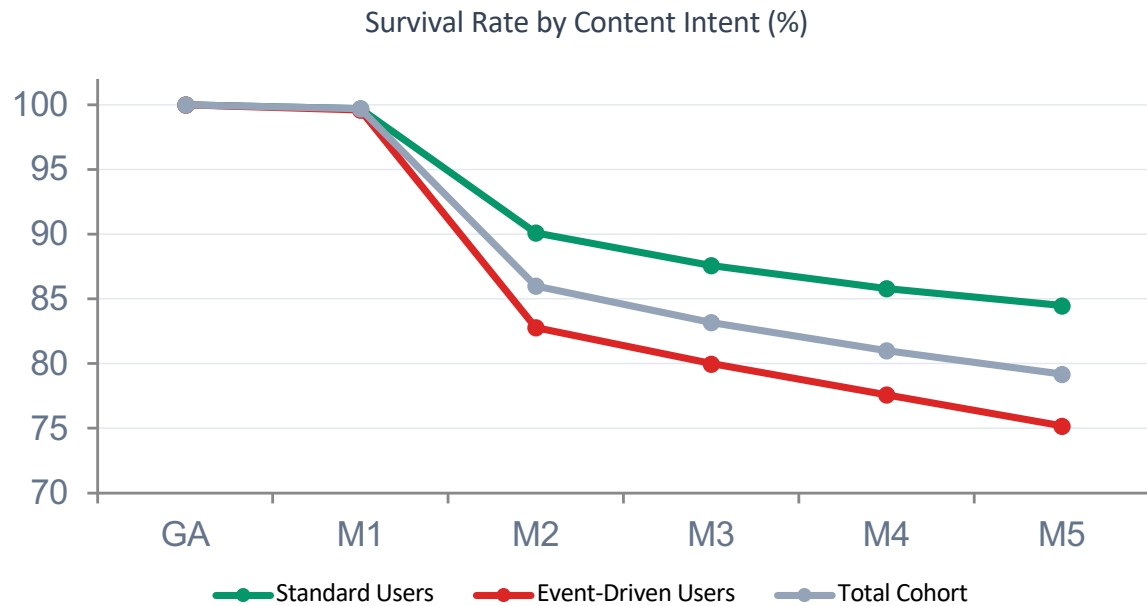
Why Digital Underperforms

Low-friction sign-ups attract opportunistic users responding to short-term events.

Optimizing for acquisition volume alone creates misleading growth signals. Channel quality is a reliable proxy for user intent and long-term value.

User Intent: Event-Driven Users Retain ~9 Points Below Standard Users

SEGMENTATION



Intent Segmentation

What is 'Event-Driven'?

Users whose engagement was concentrated on high-demand tentpole content (major live sporting events). Classified via content consumption patterns.

~62%

of cohort was event-driven (~174K subs)

75.2%

Event-driven M5 survival (~131K subs)

84.5%

Standard user M5 survival (~90K subs)

9.3pt


Intent-driven retention gap

Across all three cuts (package, channel, and intent), the same pattern emerges: churn is driven by low-intent, event-based users, not product failure.

Voice of Customer: LLM Analysis Confirms the Behavioral Story

LLM ANALYSIS

Primary Churn Drivers: LLM Categorization of Customer Feedback



Methodology

Unstructured call transcripts between subscriber and customer representative processed through an LLM classification pipeline. ~9.4K records across 5 categories. Results validated against independent customer research.

Validated Against Independent Customer Research

"I only signed up to access specific events. Once restored elsewhere, I had no reason to stay."

"The promotion ended and the price jumped. It wasn't worth it vs. alternatives."

"The app was difficult to navigate and had buffering issues on multiple devices."

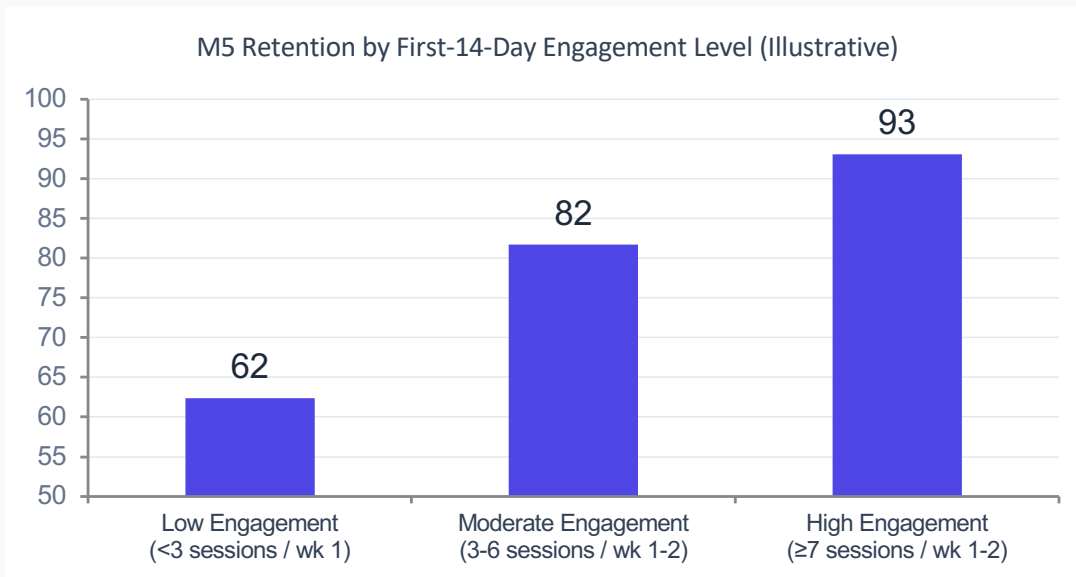
Representative synthetic quotes based on LLM-derived themes

Behavioral and qualitative signals are mutually reinforcing. Event-driven intent and pricing jointly explain the majority of early churn.

Predictive Signal: Early Engagement Thresholds Drive Long-Term Value

PREDICTIVE

These segments aren't just descriptive. They form the basis for predicting long-term value and targeting interventions. Early behavioral signals within the first 14 days strongly predict retention and estimated LTV:



THE PREDICTIVE SIGNAL

We identified a clear activation threshold: users reaching ≥ 7 sessions in the first 14 days retain at 93% vs. 62%, a 30+ point gap across $\sim 280K$ users. This becomes a decision boundary for targeting.

93.1%
High-engagement
M5 retention
($\sim 70K$ subs)

1.8x
LTV multiplier
vs. low-engage

14 days
Prediction
window

Implication: This signal was operationalized into onboarding interventions and lifecycle targeting, enabling proactive identification of at-risk users before churn intent crystallizes.

This shifts the question from "who churned?" to "who will churn?" Moving from retrospective diagnosis to a production-ready signal for targeting and intervention.

Synthesis: Three Independent Lenses, One Consistent Story

SYNTHESIS

"Churn was not driven by product failure. It was driven by misaligned acquisition quality and low user intent."

Product Tier

FINDING

Core package (~126K subs) retained at 69.5%, nearly 20 pts below Premium tier (~56K subs) at M5

INTERPRETATION

Product-content mismatch: lower-tier users didn't find sustained long-term value in the platform

Acquisition Channel

FINDING

Digital channel (~196K subs) retained at 78.8% vs. 89.9% for assisted (~84K subs), ~11pt gap

INTERPRETATION

Channel quality = intent quality: low-friction acquisition translates directly to lower commitment

User Intent

FINDING

Event-driven users (~174K subs, 62% of cohort) retained at 75.2% vs. 84.5% for standard (~106K subs)

INTERPRETATION

Transactional behavior: joined for a specific event, disengaged once that immediate need was met

The implication: stop optimizing for acquisition volume. Optimize for activation quality. Sustainable growth requires aligning acquisition strategy, product fit, and user expectations.

Business Impact: Quantifying the Cost of Event-Driven Churn

IMPACT

~\$150M

Est. Lost Cohort LTV

Using typical ARPU and observed tenure differences between segments, this early churn translates to roughly ~\$150M in lost cohort LTV.

~30%

CAC Efficiency Gap

Lower retained LTV per acquisition dollar in digital channels. Volume-optimized spend masks a significant quality gap.

~\$40M

Addressable Retention Value

A 3-pt retention improvement in high-intent digital segments yields ~\$40M in incremental LTV at current scale.

Channel Quality Gap: Acquisitions vs. Retained LTV (Illustrative)

Digital / Online



70% acq / 55% LTV

Assisted / In-Person



30% acq / 45% LTV

Acq. share

Retained LTV share

Executive Takeaway

This cohort looked like growth on the surface, but from a value perspective, it behaved more like short-term demand capture than durable customer acquisition.

Prioritized Action Plan: From Insight to Investment Decision

STRATEGY

Initiatives ranked by business impact, implementation effort, and evidence confidence. Highest-ROI actions identified for immediate deployment:




Initiative	Rationale	Impact	Effort	Timing	Confidence	Priority
Smarter Acquisition Targeting	~62% of joins were event-driven. Weighting spend toward multi-genre intent signals improves cohort quality.	High	Low	Near-term	High	#1
Pre-M2 Retention Intervention	M1>M2 is the critical churn window. Pre-emptive outreach intercepts intent before it crystallizes.	High	Med	Near-term	High	#2
App Onboarding UX Improvement	#1 exit theme (32%). Guided onboarding for platform migrants and clearer SVOD activation paths.	High	Med	Mid-term	Med	#3
Fix SVOD Auth Activation Flow	18% of churners cited content/auth friction. Streamlining app activation removes a preventable exit trigger.	Med	Low	Near-term	High	#4
Targeted Pricing Incentives	Effective for high-intent segments only. Must be constrained to users where incremental LTV justifies the cost.	Med	High	Selective	Med	#5

Trade-off note: Pricing incentives show strong short-term lift but are constrained to high-intent users where incremental LTV justifies the cost. Broad application erodes unit economics.

Strategic Research Roadmap: From Project to Portfolio

RESEARCH AGENDA

This analysis is the first thread in a broader research portfolio. Prioritized to maximize near-term decision impact while building longer-term modeling capability:

01 Usage-to-Retention Modeling 	02 Acquisition Quality Scoring 	03 Intervention Effectiveness 
STRATEGIC QUESTION <i>Which early engagement behaviors predict durable retention and long-term value?</i>	STRATEGIC QUESTION <i>Which acquisition channels and campaign types produce the highest-LTV users?</i>	STRATEGIC QUESTION <i>Which retention actions actually move the needle, and for whom?</i>
APPROACH Survival models with time-varying engagement features; threshold analysis on activation metrics	APPROACH Channel-level LTV modeling; CAC payback analysis by segment; intent scoring at acquisition	APPROACH Quasi-experimental methods (diff-in-diff, holdout groups); heterogeneous treatment effects
In progress (this case study) Ongoing	Next priority Q1-Q2	Planned Q2-Q3

Thread 1 informs acquisition targeting now. Thread 2 builds the scoring model to shift spend. Thread 3 closes the loop with causal measurement.

Limitations & Trade-offs: What We Can and Cannot Conclude

RIGOR

These findings carry important caveats that inform where results are actionable vs. directional:

Correlation vs. Causation

Survival curves describe retention patterns but do not isolate causality. Package, channel, and intent are correlated. A digital-acquired user is also more likely to be event-driven and on a lower tier.

LLM Classification Risk

LLM-based categorization carries inherent misclassification risk for ambiguous feedback. Mitigated by validating against independent customer research. Strong alignment increased confidence.

External Event Confounding

Churn timing is partly driven by when the competitor restored content, an exogenous event. This limits the generalizability of timing-based findings to non-disruption cohorts.

Retention Offer Attribution Bias

Users who received offers were pre-identified as high-risk, not randomly assigned. Measuring true lift required careful holdout group design to avoid selection bias.

These constraints inform Thread 3 of the research roadmap: causal measurement using holdout designs and quasi-experimental methods to isolate true treatment effects.

Applying This Framework in B2B SaaS

TRANSFERABILITY

The analytical principles demonstrated here (multi-method triangulation, intent segmentation, and LTV-driven prioritization) transfer directly to B2B product-led growth:

B2C Concept (This Case)		B2B SaaS Equivalent
Subscriber retention / churn	->	Account renewal / NRR / expansion
Content engagement patterns	->	Product usage & feature adoption stages
Event-driven acquisition spike	->	Campaign-driven or PLG trial surges
Product tier (Core / Premium)	->	License tier / feature adoption depth
Acquisition channel quality	->	GTM motion (PLG vs. Sales-led vs. CSM)
LLM-classified churn reasons	->	LLM-classified support tickets / NPS verbatims
Cohort survival curves	->	Time-to-value / onboarding completion analysis

Core transferable insight: In both B2C and B2B, acquisition volume can mask intent quality. The same multi-lens framework (behavioral signals, segmented retention, and unstructured feedback) identifies which accounts are durable and which are at risk.

Key Takeaways

01 Not All Growth Is Equal

Event-driven acquisition creates volume but not durable value. Retention metrics must be interpreted in light of how and why users entered, not just how many.

02 Behavioral + Qualitative = Confidence

Combining survival analysis with LLM-derived voice-of-customer data created mutually reinforcing signals. Neither alone would have been as complete or as actionable.

03 Intent Is Detectable at Acquisition

Channel type, product tier, and content engagement behavior are leading indicators of long-term retention. These can power proactive interventions, not just retrospective analysis.

04 Insight Without Action Is Incomplete

Findings directly informed targeted retention interventions and acquisition strategy adjustments, closing the loop between statistical analysis and business decisions.

"The core shift is from measuring churn to engineering retention. Once you can predict durable users early, you can change the system, not just react to it."