

**University of Bath**

School of Management

MSc in Entrepreneurship Management and Innovation

**Navigating the AI Transition through Incumbent Intrapreneurship:  
A Framework for Cross-Portfolio Workflow Automation Ventures  
in the Live Media Industry**

*Roy Hasson Collier*

Academic Year 2025/2026

*Submitted as part of the requirement for completing an MSc in  
Entrepreneurship and Management at the University of Bath*

# **Contents**

List of Figures

Acknowledgements

Abstract

## **Chapter 1 Introduction**

- 1.1 The category shift in live media
- 1.2 The strategic problem for incumbents
- 1.3 Research question and scope
- 1.4 Origin, motivation, and aspirations
- 1.5 Structure of the dissertation

## **Chapter 2 Exploratory Phase**

- 2.1 Literature Review
- 2.2 Customer and Competitor Overview

## **Chapter 3 Generative Phase**

- 3.1 Methodology
- 3.2 Findings: Cross-Interview Analysis
- 3.3 Discussion
- 3.4 The Framework

## **Chapter 4 Business Model and Financial Viability of the Focal Venture**

- 4.1 The focal venture
- 4.2 Business model
- 4.3 Organisational design
- 4.4 Financial model

## **Chapter 5 Conclusion and Reflection**

- 5.1 Answering the research question
- 5.2 Contribution
- 5.3 Reflection in action
- 5.4 Reflection on action
- 5.5 Limitations and further research

**References**

**Appendices**

Appendix A Interview Guide (Semi-Structured Protocol)

Appendix B Participant and Pseudonym Matrix

Appendix C Thematic Coding Tree

Appendix D Illustrative Financial Model Assumptions

## List of Figures

Figure 2.1 The integration gap in the live media value chain

Figure 2.2 The integration paradox in the literature: separation versus integration

Figure 2.3 Competitive positions and the unaddressed cross-layer gap

Figure 2.4 Two structural approaches to entering the category and their constraints

Figure 3.1 Four-pool source triangulation on a single research question

Figure 3.2 The integration paradox

Figure 3.3 The venture as a set of contingent design choices

Figure 4.1 Business Model Canvas for the focal venture, mapped to the integration paradox

Figure 4.2 Illustrative five-year financial shape — the transition J-curve

## **Acknowledgements**

I would like to thank my supervisor, Dr Stoyan Stoyanov, for his guidance and thoughtful feedback throughout this research. I am also grateful to the many industry professionals who generously shared their time and candid insights in the interviews; their perspectives are the foundation of this work. Above all, I thank my family for their unwavering support and encouragement, and for giving me the time and space to immerse myself in this research — I could not have completed it without them.

## **Abstract**

This dissertation examines how an established live media technology vendor should organise and execute an AI-enabled cross-portfolio workflow automation venture. Live media is undergoing a category shift in which competitive advantage is moving from the hardware layer that historically defined it to an AI-driven software layer that orchestrates the operational decisions connecting live production. Incumbent vendors, built for the hardware era, are ill-equipped to respond, yet hold complementary assets AI-native entrants cannot easily replicate. The central problem, termed the integration paradox, is that the venture's advantage depends on close integration with the incumbent's infrastructure and customers, while its survival depends on enough separation to operate at venture speed and be trusted as neutral.

Following an applied corporate-venture design, five design propositions were derived from the literature on corporate entrepreneurship, architectural disruption, and AI as a general-purpose technology, and tested through semi-structured interviews across four stakeholder pools: AI-native ventures, incumbent vendors, broadcast customers, and an industry-standards body. The analysis confirmed two propositions strongly, qualified one as contingent, recorded one as confirmed but contested, and reframed one. The principal finding is that, in this category, technology is a necessary condition of competing but not the source of durable advantage, which is instead human and tacit; the expected data moat does not hold in an asset-centric industry where the monetisable asset is content rather than operational data.

These findings are synthesised into a contingent framework organised around the integration paradox, and applied to a hypothetical focal venture through a Business Model Canvas and an illustrative financial model whose J-curve is the financial signature of the paradox. The study contributes a bounded divergence from the complementary-assets reading of incumbency and a transferable way of reasoning about incumbent-led ventures whose advantage and autonomy pull in opposite directions.

## **Chapter 1: Introduction**

### **1.1 The category shift in live media**

Live media is going through a category shift of structural significance. For three decades, the industry's competitive logic was organised around one value proposition: the reliable transport, encoding, and switching of live signals from acquisition to distribution. Hardware quality, signal integrity, and latency defined who won. The leading vendors built their businesses around mastering the physics of live signal transport — a mastery that is now a liability as much as an asset.

The new layer increasingly determining advantage is workflow automation: the AI-driven software coordinating what gets captured, how it moves through ingest and production, and how it reaches distribution. With AI this becomes closed-loop automation that detects conditions across the signal chain, decides on interventions, and executes them without humans in the loop. The economics shift from hardware box sales to platform pricing, the buyer from production technology to operational efficiency, and the competitive set widens beyond the traditional vendor list to AI-native software ventures that did not exist five years ago.

This shift is not incremental but architectural. The components of live production — cameras, encoders, switchers, distribution — are not disappearing; the intelligence layer connecting them is changing fundamentally, and incumbents whose structures and customer relationships were built for the hardware layer are ill-equipped to respond. Henderson and Clark (1990) explain why: firms whose competence lies in component-level expertise struggle most when the architecture connecting components changes, their existing knowledge becoming not merely irrelevant but actively misleading.

### **1.2 The strategic problem for incumbents**

Incumbent live media technology vendors face a specific, underexplored strategic problem: their operating models are built for a category that is no longer the locus of competitive advantage. The AI-native independent ventures emerging in workflow automation, among them companies focused on AI-driven production orchestration and story-centric newsroom management, have operating models calibrated to the new category from inception: software-native economics, iterative customer co-development, and products that improve as data accumulates.

Incumbents do, however, hold one structural advantage independent ventures cannot easily replicate. A cross-portfolio AI-enabled workflow automation venture, spanning acquisition

through distribution, depends on integration with the upstream and downstream systems incumbents already own, including signal-layer data AI-native entrants cannot access. The constraint is not technical. It is organisational. A cross-portfolio venture sits across product lines, competes with them for resources, and requires an operating model the incumbent does not yet have. The integration paradox, as this research terms it, is that realising the incumbent's structural advantage requires the venture to remain connected to the parent, while building the AI-native operating model that makes it competitive requires separation. This tension is unresolved in the existing literature on corporate entrepreneurship, incumbent disruption, and AI strategy.

### **1.3 Research question and scope**

The research question is *how should an established live media technology vendor organise and execute an AI-enabled cross-portfolio workflow automation venture.*

The output is a framework for AI-enabled intrapreneurial innovation in this category, illustrated through a hypothetical focal venture via the business model canvas and financial model. The framework is designed to be operationally specific, addressing the governance, operating-model, talent, data, and GTM questions a live media incumbent would actually need to resolve. The focal venture is generic and hypothetical rather than a specific company's roadmap, to protect commercial confidentiality and ensure the framework applies beyond any single vendor.

### **1.4 Origin, motivation, and aspirations**

This research emerged from direct engagement with the live media technology industry at a time when the category shift described above is already observable in practice. Practitioners are experiencing not a future disruption but a present one, and what is missing is not awareness of the shift but a framework for responding to it.

The academic literature on corporate entrepreneurship, incumbent disruption, and AI strategy addresses parts of the problem but has not engaged with the specific combination present in live media: a mature B2B technology category with hardware legacy, an architectural disruption driven by AI, and the integration paradox at the centre of the incumbent's strategic choices. The aspiration is twofold: to contribute to the literature on corporate entrepreneurship and incumbent response to AI, and to inform incumbents navigating this shift.

### **1.5 Structure of the dissertation**

Chapter 2 constitutes the exploratory phase. Section 2.1 presents the literature review across three bodies of theory, concluding with five design propositions that become the starting assumptions for the generative phase. Section 2.2 provides the customer and competitor overview: the workflow-automation category, the incumbent vendor landscape, and the AI-native ventures used as comparative reference cases.

Chapter 3 constitutes the generative phase: semi-structured interviews across four stakeholder pools, analysed thematically to test and refine the design propositions. Its findings inform the framework developed at the close of the chapter and, in turn, the business model canvas in Chapter 4, which is therefore a product of the synthesis rather than an instrument used during data collection.

Chapter 4 presents the business model summary: the nine canvas elements as they stand at the end of the generative phase, and the financial model across revenue, cost, working-capital, and investment components.

Chapter 5 provides discussion and reflection: an assessment of viability, a reflection on the research process and what remains unexamined, and the findings' contribution to academic understanding and practice.

## **Chapter 2: Exploratory Phase**

### **2.1 Literature Review**

#### **2.1.1 Introduction**

At the centre of this research sits a tension that existing literature has not yet resolved. Building a new venture inside an established firm requires, by most accounts, structural separation from the core: its own team, its own metrics, and protection from the parent's resource-allocation logic. But the venture this research investigates draws its competitive value precisely from not being separate: its advantage over independent entrants is access to the incumbent's infrastructure, data, customer relationships, and business context. Separate it fully and it becomes just another underfunded startup. Leave it embedded and the hardware business absorbs it. This is the integration paradox, and none of the three bodies of literature this review draws on has confronted it directly.

The first, on corporate entrepreneurship, offers the most developed account of how ventures survive or fail inside established firms, but was built largely around product-category ventures in manufacturing and consumer technology. It has not grappled with ventures whose value depends on remaining integrated with the parent's infrastructure rather than breaking from it. The review concludes by combining these three bodies into a set of design propositions that the generative phase will test.

### 2.1.2 Corporate Entrepreneurship and Intrapreneurship

The central puzzle is not why new ventures fail, but why established firms, with all their resources and market access, so consistently fail to launch them. Burgelman (1983) offers the most durable account of this puzzle. His fieldwork inside a large technology company led him to distinguish between two types of strategic behaviour: **induced**, where initiatives follow the firm's existing strategic logic and compete successfully for resources within it, and **autonomous**, where initiatives emerge from lower levels of the organisation in response to opportunities the firm has not formally recognised. The distinction matters because genuinely novel ventures almost always begin as autonomous initiatives: they do not fit the firm's current strategy, compete awkwardly with established business units, and depend for survival on a senior champion willing to translate their value into language the organisation can process. Without that translation, Burgelman found, autonomous initiatives tend to be starved or absorbed before they establish themselves.

A cross-portfolio AI-enabled workflow automation venture inside a live media incumbent fits only parts of this picture. It belongs to no existing product line, competes for talent and budget with hardware divisions whose revenue is immediate and measurable, and rests on software-native economics difficult to articulate in the performance metrics the parent uses to allocate resources. But Burgelman's model deserves scrutiny before it is applied uncritically. His fieldwork was conducted in a 1970s semiconductor company, where product-category ventures aimed at a standalone business unit independent of the parent. The venture this research concerns is different: *its competitive advantage depends on integration with the parent's existing product portfolio, not independence from it*. Burgelman's prescriptions for protecting autonomous initiatives, including dedicated resources and protected space, remain relevant, but his model does not fully account for the governance challenge of a venture that must be **separate enough** to develop an AI-native operating model and **connected enough** to exploit the incumbent's infrastructure. Industry bodies note this standards-based integration layer

remains underdeveloped relative to the pace of IP adoption (DPP, 2024) — part of what this research addresses (see Figure 2.1).




Figure 2.1: The integration paradox. Adapted from Burgelman (1983) and O'Reilly and Tushman (2008).

O'Reilly and Tushman (2008) formalise the structural argument through their concept of **organisational ambidexterity**. They define it as an organisation's ability to exploit its existing business while exploring new ones, and argue, backed by comparative research, that the two activities need to be structurally separated. Exploration units that share structures, processes, and incentives with exploitation units tend to be outcompeted for resources, because exploitation's returns are immediate while exploration's are delayed and uncertain. Their recommendation is structural separation with integration at senior-leadership level: its own team, metrics, and culture, but a leadership connection to the parent that protects its resource base and manages the interface between the two logics.

This prescription is not the complete picture. Birkinshaw and Gibson (2004) offer a complement: contextual ambidexterity holds that organisations can build the capacity for both exploitation and exploration within a single unit through behaviour, culture, and context, not structural separation alone. This applies directly to the integration paradox. Full separation

cannot resolve it precisely because the venture needs ongoing operational connection to the hardware business: its people must work with incumbent infrastructure teams, access signal-layer data, and maintain the customer relationships that give the venture its integration advantage. This is contextual ambidexterity at the level of the venture unit itself, where team members hold both AI-native development logic and incumbent-domain knowledge. Birkinshaw and Gibson's model therefore refines rather than contradicts the structural argument: separation from the core is necessary for organisational survival, but a contextually ambidextrous culture within the venture sustains the integration advantage. Raisch and Birkinshaw (2008) confirm that firm size and environment moderate ambidexterity outcomes, supporting a contingency reading: the balance between separation and integration varies by incumbent type, which the framework addresses directly.

What the intrapreneurship literature has not addressed is the combination of factors present in the live media case: a hardware-legacy B2B category, a venture whose competitive advantage depends on integration with the incumbent's own physical infrastructure and with the third-party tools, platforms, and interoperability standards that make up the broader live production ecosystem, and an architectural disruption driven by AI rather than by a new product category. How a venture can be structurally separated from the core and yet retain governed access to the infrastructure, data, and ecosystem connections that make it distinctive is a question the existing literature does not answer. The architectural disruption literature, addressed in the following section, explains why this challenge is particularly acute in the live media context.

### **2.1.3 Incumbent Response to Architectural and Platform-Level Technological Change**

Henderson and Clark (1990) studied the market for photolithographic alignment equipment, a narrow slice of semiconductor manufacturing, but the mechanism they identified has proved generative precisely because it does not depend on the specifics of that industry.

Their taxonomy distinguishes innovations **not just by whether components change, but by whether the relationships between components change**. Incremental innovation improves components without disturbing the architecture. Modular innovation changes a component while leaving its interfaces with the rest of the system intact. Radical innovation changes both. **Architectural innovation**, the category they introduce, changes the relationships between components while leaving the components themselves largely intact. This last type is the most dangerous for incumbents, for a reason not immediately obvious: from inside the firm it looks like incremental change. The components are familiar. The basic technologies have not shifted.

What has shifted is the logic that connects them, and incumbents' organisational structures, calibrated to the old logic, actively prevent them from seeing this.

The mechanism works through what Henderson and Clark call information filters: channels and routines that route information efficiently according to the existing architecture. When the architecture changes, the filters still route information, but by the old categories, so problems visible under the new logic fail to register and the firm carries on optimising for a configuration that is no longer competitively decisive.

The live media application of this argument is uncomfortably specific. The cameras, encoders, switches, and distribution systems that incumbents build and sell are not disappearing. **What is changing is the layer that connects and orchestrates them:** the workflow, the decision logic, the automation. An incumbent organised around signal encoding or hardware reliability will have R&D teams, product managers, and sales people whose knowledge and incentives are calibrated to the component layer. The workflow layer registers, if at all, as a software feature request for an integration team. Analysts observe that while IP adoption in live production has accelerated, automation of the workflows connecting these systems has not kept pace (DPP, 2024), and new venture formation has shifted toward software-layer orchestration even as major incumbents' announcements remain hardware-focused. The information-filter mechanism explains why.

Christensen (1997) addresses the demand side: where Henderson and Clark explain how incumbents fail to perceive architectural change, Christensen explains why, even perceiving it, they fail to invest. Disrupting technologies initially underperform on the metrics mainstream customers value, and an incumbent's best customers want better versions of what they have. The financial model reinforces this: the disrupting category offers lower margins initially, so allocating resources to it means withdrawing them from products that generate current profit. The rational, well-managed firm ignores the disruption until it is too late.

The theory has attracted criticism. Lepore (2014) argued the original cases were selectively presented and the theory cannot be falsified, since any outcome can be retrofitted into the disruption narrative. King and Baatartogtokh (2015) examined fifty-seven of Christensen's cases and found fewer than half satisfied the theory's own criteria. Danneels (2004) offers the most constructive response, operationalising the theory and identifying the conditions under which incumbent failure is and is not likely. This research draws on Christensen descriptively,

to account for the observed underinvestment in workflow automation among live media incumbents, not to predict which firms survive.

Teece (2007) argues that incumbent failure under technological change is not inevitable, introducing a concept directly significant here: complementary assets — the resources and capabilities a firm holds alongside its core technology that are required to capture value from an innovation. In live media, incumbents hold complementary assets AI-native entrants cannot easily replicate: an installed base of hardware across the signal chain from acquisition through distribution; long-standing relationships with broadcasters whose production environments are built around incumbent technology; proprietary operational data accumulated across years of live production; and the integration points across the full workflow that let an incumbent orchestrate automation across the whole signal chain rather than within one segment. On Teece's logic, these assets, appropriately leveraged, represent a potentially durable advantage in the category. The incumbent that combines its complementary-asset base with an AI-native operating model is, in principle, better positioned than an entrant building the same capability without that foundation.

The challenge is converting this potential into realised advantage. Teece identifies three dynamic capabilities that determine whether an incumbent can: sensing (identifying the opportunity before it is obvious), seizing (mobilising resources to act), and transforming (restructuring to sustain the new position). In live media, sensing is available to incumbents through their proximity to broadcaster pain points; seizing requires governance and operating-model decisions the existing organisation is not designed to make; and transforming requires structural and cultural changes that compete with the existing hardware business. The record on whether incumbents navigate this sequence is mixed, and Eisenhardt and Martin (2000) and Winter (2003) note that dynamic capabilities are better understood as a post-hoc description of what successful firms did than a guide to what firms should do. That gap between prescription and operational detail is one of the spaces this research fills. Critically, neither the architectural-innovation nor the disruption literature examines how incumbents' complementary assets interact with the governance and operating-model requirements of a cross-portfolio AI-native venture in a hardware-legacy category. That interaction is the focus of the following two sections.

#### **2.1.4 AI as a General-Purpose Technology and its Strategic Implications**

Brynjolfsson and McAfee (2014) argue that AI belongs to a class of technologies, alongside steam and electricity, whose significance spreads through gradual complementary adjustment rather than immediate displacement. What matters most here is not the aggregate claim but the adjustment mechanism. Firms that adopted electricity did not simply replace steam motors with electric ones and collect the gains; those came later, from redesigning factories around electricity's properties. The parallel for AI is that its value lies not in the technology itself but in the complementary organisational changes it makes possible, which take time and willingness to disrupt existing ways of working.

Brynjolfsson, Rock and Syverson (2021) formalise this as the productivity J-curve: firms that adopt AI without the necessary organisational adjustments see an initial decline in measured productivity before the gains materialise, because the disruption is immediate while the benefits of redesigned workflows are not. For live media incumbents this is a practical warning: industry analysis found early pilots within existing structures produced mixed results, gains offset by integration overhead and workflow disruption (DPP, 2024). Bolting AI onto existing workflows without redesigning them reproduces the J-curve in compressed form.

Agrawal, Gans and Goldfarb (2018) offer a more granular economic account. They argue that the most useful way to understand what AI does, across its many applications, is as a reduction in the cost of prediction: the conversion of available data into probabilistic outputs about uncertain states. When prediction becomes cheap, the value of what surrounds it rises: data as the raw input, human judgement on high-stakes decisions prediction informs but should not determine, and the systems that execute on predictions at scale. For live media, this reorients the question: not which AI model to use, but what operational data the incumbent can access, how it converts into workflow predictions, and what infrastructure executes them at the latency live production demands. An incumbent with years of signal quality records, ingest logs, and production workflow data has a potential input advantage over AI-native entrants who are building their training base from scratch, though whether that data is accessible, clean, and structured for machine learning use is a separate and less certain question.

Iansiti and Lakhani (2020) describe the organisational form that AI-native firms have converged on, which they call the **AI factory**: a combination of data pipelines, model development infrastructure, and automated decision and action systems that replaces the human-driven operational core of the business. Its economic properties differ from traditional operating models: costs do not scale proportionally with output, the system improves

automatically as data accumulates, and the same capability can be deployed across markets without proportional additional investment. A small but growing number of AI-native ventures in live media are, in their niches, early-form AI factories, their operating models calibrated to these economics from inception. A cross-portfolio venture inside a live media incumbent would need comparable properties while built inside a context designed for hardware product economics. Iansiti and Lakhani's framework was developed for large platform and pure-play AI firms; whether its prescriptions translate to an intrapreneurial venture within a hardware-legacy parent is a question this research is, in part, attempting to answer.

There is a genuine debate in this literature about whether AI amplifies incumbent advantages or creates space for new entrants. Agrawal and colleagues suggest that cheap prediction democratises capability previously available only to large firms. Acemoglu (2023), among others, argues that AI economics, particularly the data flywheel and scale advantages in model training, tend to favour large incumbents and platform companies, producing concentration rather than democratisation. Both dynamics probably operate at once: incumbents have infrastructure, data, and relationship advantages AI could amplify, while AI-native entrants have operating-model and speed advantages scale alone does not overcome. Which wins is a question for the generative phase.

### **2.1.5 Synthesis and Design Propositions**

The three bodies of literature do not simply add up to a single argument; they interact, and the interactions matter. The intrapreneurship literature identifies the conditions that suppress autonomous ventures — resource competition, cultural resistance, legitimacy deficits — but not why AI-enabled workflow automation specifically is hard to fund internally. The architectural-disruption literature explains why incumbents under-resource the disrupting category, but not the operating-model transformation AI requires even when resources are available. The AI-strategy literature describes AI-native operating models and their superior economics, but was not developed with hardware-legacy incumbents in mind and says little about building such a model from the inside. Each explains a different mechanism of failure, and the mechanisms reinforce one another: the architectural shift makes the category hard to perceive, intrapreneurship dynamics make it hard to fund, and the operating-model gap makes it hard to execute even when perception and funding are secured. The challenge is structurally layered, and not uniform across incumbents: a mid-sized vendor with software-curious leadership faces a different version of the problem than a large hardware company optimised

for margin protection, and the framework accounts for this variation rather than treating all incumbents as equivalent.

And yet incumbents are not without advantages. The combination of hardware ownership and SaaS or AI capability is, in principle, a structural moat AI-native entrants cannot easily replicate: the hardware layer generates the signal-domain data that feeds the AI layer; the customer relationship, built over years, provides access to production environments entrants are still trying to enter; and the integration points across acquisition, ingest, production, and distribution are already owned. *A pure-play software venture can automate within its slice of the workflow; only an incumbent can, in principle, automate across it.* These advantages are real but not permanent, and realising them is the integration paradox: the venture must remain connected to the parent to draw on them, yet separated enough to build the AI-native operating model that makes it competitive. The literature identifies both the need for separation and the value of integration without resolving how to achieve both at once. That is the gap this research addresses through the framework developed in the generative phase (see Figure 2.2).

Timing adds a further dimension. Christensen's disruption model and Agrawal and colleagues' data flywheel argument both imply that the window for incumbent response is finite. AI-native entrants improve as they accumulate data and refine their models. Industry interoperability standards, as they mature and adoption widens, gradually reduce the technical integration barriers that incumbents currently hold as structural advantages. The framework this research develops is therefore designed for a specific competitive moment, not as a timeless prescription.

From this synthesis, five design propositions are derived. These are working assumptions, not predictions. Each is grounded in the literature reviewed above. Each carries a sceptical counter-position, which the generative phase will use as an active test rather than a straw man.

**Design Proposition 1:** A cross-portfolio AI-enabled workflow automation venture inside a live media incumbent **requires structural separation from the core hardware business** to avoid being starved of resources or absorbed by the existing organisational logic, but must retain governed access to the incumbent's infrastructure, operational data, and customer relationships to offer something that independent AI-native ventures cannot. (Burgelman, 1983; O'Reilly and Tushman, 2008; Iansiti and Lakhani, 2020)

***Sceptical position:*** The counter-position is that structural separation in a mid-sized live media vendor is practically unachievable: the firm is too small for genuine ambidexterity, and any protected venture unit will eventually be raided for resources during a commercial downturn.

**Design Proposition 2:** Existing signal-domain technical competence is not transferable to AI-enabled workflow automation and is likely to actively distort product and go-to-market decisions if those decisions are shaped by people whose expertise is in the hardware layer. The venture requires **talent and knowledge profiles that are distinct** from the incumbent's existing engineering and product functions. (Henderson and Clark, 1990; Iansiti and Lakhani, 2020)

***Sceptical position:*** The counter-position is that domain knowledge of live production workflows is a significant advantage in designing AI-driven workflow automation, and that replacing hardware engineers with AI specialists would sacrifice exactly the contextual expertise that makes the venture's product credible to broadcasters.

**Design Proposition 3:** The primary design challenge for the focal venture is **operating model design rather than technology selection**. Building the venture around AI-native operating principles, including data flywheel economics, software-native unit economics, and iterative customer co-development, is more important than which AI technologies it uses. (Iansiti and Lakhani, 2020; Brynjolfsson and McAfee, 2014)

***Sceptical position:*** The counter-position is that in live media, where signal reliability and latency are non-negotiable, the technology capability constraint is primary. Operating model design cannot compensate for AI systems that are not production-grade, and getting the technology right is the prerequisite for everything else.

**Design Proposition 4:** The incumbent's years of **accumulated operational data represent a potential competitive advantage** that AI-native entrants cannot easily replicate in the short term. Whether this advantage is realisable depends on whether the venture can access, organise, and deploy that data within an AI-native operational structure, which is not guaranteed. (Agrawal, Gans and Goldfarb, 2018; Teece, 2007)

***Sceptical position:*** The counter-position is that legacy operational data from hardware-era production workflows may be poorly structured, inconsistently logged, and not representative of the AI-era workflows the venture is trying to automate, making it less

valuable than it appears and potentially a distraction from building the right data infrastructure from scratch.

**Design Proposition 5:** The structural advantages the incumbent holds in cross-portfolio workflow automation, including integration ownership, operational data, and existing customer relationships, are **time-limited**. AI-native entrants compound their learning advantages over time, and industry interoperability standards are gradually reducing the technical barriers to integration that incumbents currently benefit from. Acting within the current window matters. (Christensen, 1997; Agrawal, Gans and Goldfarb, 2018; Teece, 2007)

**Sceptical position:** The counter-position is that in B2B live media, customer switching cycles are so long and risk tolerance so low that the window is longer than the AI strategy literature implies, and that incumbents may have more time than this proposition suggests to build a credible response. The following section examines the competitive landscape of the live media workflow automation category, establishing the empirical context that informs the first iteration of the business model canvas.




Figure 2.2: Navigating the integration paradox — design propositions in sequence.  
Derived from synthesis of literature reviewed in sections 2.1.2–2.1.4.

## 2.2 Customer and Competitor Overview

### 2.2.1 The Live Media Workflow Automation Category

Live media workflow automation refers to the coordination, orchestration, and increasingly AI-driven execution of the operational decisions connecting live production from acquisition to distribution. It spans several layers: resource scheduling; ingest coordination and quality monitoring; production orchestration including rundown management, graphics triggering, and studio automation; and distribution routing and multi-platform publishing. What distinguishes it from the signal-transport layer that historically defined the category is that it operates on decisions rather than signals: not how reliably a signal moves from camera to encoder, but how intelligently the system decides what to capture, how to route it, and where to send it.

AI capability has crossed the threshold at which closed-loop automation — systems that detect conditions, decide on interventions, and execute without humans in the loop — is achievable in live production for a meaningful subset of decisions. And broadcaster cost pressure, intensified by multi-platform distribution and remote production, has created demand for operational efficiency that hardware-layer improvements alone cannot address (DPP, 2024). In Henderson and Clark's (1990) terms, this is an architectural shift rather than an incremental one: the components of live production are not disappearing, but the intelligence layer connecting and orchestrating them is changing fundamentally.

### 2.2.2 The Competitive Landscape: Two Structural Approaches

The live media workflow automation category is being entered from two structurally different directions simultaneously, and the difference matters analytically.

The first approach is **hardware-out**. Established vendors are extending into workflow automation by adding software and AI to hardware-anchored product lines. Harmonic introduced an AI orchestration service for live video workflows in April 2026 (Harmonic, 2026); Dalet launched Dalia at IBC 2025, an agentic AI layer embedded across its portfolio that can schedule ingest, trigger review workflows, and execute tasks end-to-end (Dalet, 2025). In each case AI is added to an existing platform rather than the platform being designed around AI from inception. The architecture of the incumbent's existing product bounds the automation; the workflow layer is, in Henderson and Clark's (1990) terms, a feature of the existing system rather than a reconfiguration of it.

The second approach is **software-in**. A smaller number of independent ventures have built workflow automation from inception, without hardware legacy, with AI as a foundational

design assumption. Two are used here as illustrative examples at different workflow layers. They are selected because they are publicly documented, active at industry events, and representative of distinct positions; they are not an exhaustive mapping of the AI-native landscape, which continues to evolve.

Nexog originated in telecoms network automation and has extended its AI-native platform into broadcast. Its trajectory is analytically significant: rather than a hardware company adding software, Nexog is a software-native, AI-driven platform from an adjacent domain entering live media, showing that barriers to entry at the operations layer are lower for software-native players than for hardware incumbents moving upward. Its AIOS platform covers resource scheduling, booking of equipment and personnel, multi-vendor network coordination, and transport stream management for live events and sports feeds (Nexog, 2025). The broadcast-specific G2 Orchestrator addresses the resource planning and network management decisions that precede and enable live production. Nexog's position sits at the operational layer immediately above the signal infrastructure, closest to the incumbent hardware vendor's territory.

Cuez operates at the editorial and production orchestration layer. Its platform puts the story rather than the broadcast schedule at the centre of production workflow, with AI agents handling repetitive tasks including headline generation, multichannel content adaptation, and rundown management within live production environments (Cuez, 2026). Cuez's open agentic AI framework lets broadcasters deploy their own fine-tuned or local models where data governance requires it, reflecting the editorial sensitivity of newsroom environments. The buyer is different from Nexog's: news directors and editorial leads rather than operations engineers and infrastructure managers.

These ventures are not competitors. They address different workflow layers, buyers, and operational problems, and represent different routes into the category: Nexog via telecoms network automation, Cuez from within the broadcast editorial context. What they share is the operating-model logic of AI-native ventures: built for software economics from inception, improving as they accumulate data and usage, with no architectural legacy constraining their product decisions. The distribution layer also has AI-native participants: cloud-native playout and FAST channel delivery platforms have emerged as genuinely software-native alternatives to traditional broadcast playout infrastructure, with AI-driven scheduling, ad insertion, and channel management at their core (Amagi, 2025). What no single AI-native venture currently

does, however, is span the full chain from acquisition and ingest through production orchestration and distribution within a single platform. Each occupies a position. None spans the whole. The cross-layer integration challenge remains unaddressed by any single independent player. This is the integration paradox in competitive form: the entities with full-chain infrastructure reach lack the AI-native operating model, and the entities with the AI-native operating model lack the infrastructure reach. Figure 2.3 maps these positions and makes the gap visible. Figure 2.4 compares the two structural approaches and their respective constraints.




Figure 2.3: Live media workflow automation — player positions by layer. Illustrative examples only.




Figure 2.4: Two structural approaches to live media workflow automation. Derived from section 2.2.

### 2.2.3 The Customer Response: Broadcasters Building Their Own

A third response, less visible in competitive analysis but equally significant here, is large broadcasters building bespoke workflow automation internally. Where the vendor landscape fails to provide an integrated solution, well-resourced broadcasters have in some cases invested in proprietary infrastructure tailored to their own production environments. This response is documented at the industry level through the IBC Accelerator programme, whose annual project challenges are proposed by broadcasters identifying technology gaps that the market has not yet filled.

The 2025 IBC Accelerator project on AI Agent Assistants for Live Production, which won Project of the Year, was proposed by broadcasters including BBC and ITN precisely because no commercial product provides an integrated agentic layer across the full live production stack (IBC, 2025). The proof-of-concept demonstrated a system of twelve coordinated AI agents covering audio, automation, graphics, rundown management, content discovery, and video enhancement, with participants including NBCUniversal, EVS, Google, Cuez, and Moments Lab. The project's significance for this research is not its technical output but its diagnostic value: the fact that a consortium of major broadcasters and technology vendors found it necessary to build a cross-vendor agentic system as an industry research project, rather than procuring it from a single vendor, confirms that no existing product solves the cross-portfolio, full-chain automation problem at scale.

The broadcaster build-your-own response is a second-best outcome for all parties, and is consistent with the demand-side mechanism Christensen (1997) identifies: when no commercial vendor adequately addresses the operational problem, well-resourced customers substitute internally rather than waiting for the market to catch up. It is expensive, slow to develop, and non-transferable: solutions built for one broadcaster's production environment typically cannot be productised or sold to others. It represents latent commercial demand for a solution that does not yet exist in the market.

#### **2.2.4 The Customer Problem**

The demand-side picture that emerges from the competitive landscape and the broadcaster response is coherent. Broadcasters need workflow automation that operates across the full signal chain, from acquisition coordination through ingest, production orchestration, and multi-platform distribution, with AI-native economics and the reliability live production demands. What is currently available does not fully address this need. Incumbent vendors offer deep integration across the signal chain but embed automation within architectures designed for the hardware era, producing workflow features rather than workflow platforms. AI-native independent ventures offer genuine AI-native operating models and fast iteration, but within single workflow layers and without the infrastructure integration access that cross-chain automation requires. Large broadcasters who build their own solutions gain a degree of integration but at significant cost and without the benefits of a productised, continuously improving platform.

Industry analysis confirms the gap between what the category requires and what current vendors provide is most acute at the integration layer: where automation decisions in one part of the workflow must propagate automatically to others (DPP, 2024). A broadcaster automating ingest decisions still requires manual handoffs to production scheduling. A broadcaster automating distribution routing still requires human intervention to coordinate with the ingest and production systems upstream. The manual effort remains concentrated precisely at the connection points between systems, which is where a cross-portfolio incumbent venture with integration access across the full signal chain would be positioned to operate. The incumbent's signal-layer infrastructure, operational data, and customer relationships are, in Teece's (2007) terms, complementary assets: they cannot easily be replicated by AI-native entrants and represent the potential basis for durable competitive advantage if combined with an AI-native operating model.

Whether this customer problem is experienced and described in the same terms by practitioners across the vendor, venture, and broadcaster communities is an empirical question. The generative phase of this research tests it directly.

## **Chapter 3: Generative Phase**

### **3.1 Methodology**

#### **3.1.1 Research problem and underlying assumptions**

This research asks how an established live media technology vendor should organise and execute an AI-enabled cross-portfolio workflow automation venture. The question is strategic and organisational rather than technological: it is concerned with governance, operating model design, and the conditions under which a structurally specific type of intrapreneurial venture can succeed inside a hardware-legacy incumbent. The literature review concluded with five design propositions — covering structural separation, architectural competence as liability, operating model primacy, the convertibility of complementary assets, and the time-bounded nature of the competitive window — which function in the generative phase as working hypotheses to be tested, refined, and where necessary revised against empirical evidence.

#### **3.1.2 Research approach and why qualitative interviews**

The generative phase is qualitative, exploratory, and interpretivist in character. The research is not testing a hypothesis with a representative sample or measuring variance across a defined population. It examines a complex strategic situation in which the mechanisms are partially understood, the relevant actors are few and senior, and the evidence that matters is contextual, tacit, and surfaces only through conversation.

An earlier research plan included a survey of live media professionals, dropped for reasons worth stating plainly. The questions the research must answer — whether structural separation protects an internal venture from resource reallocation, whether incumbents' operational data is genuinely convertible into AI advantage — cannot be answered with closed-response items; they concern mechanisms and conditions, not frequencies. A tick-box response from a senior broadcaster would reveal the position held without revealing why, what experience underlies it, or what might change it. More practically, the people whose perspectives matter most — CTOs, AI-native founders, heads of broadcast innovation — do not complete surveys. They have conversations, and those need to be real ones with time and thought invested.

Semi-structured interviews suit this: a structured core ensures comparability, while the freedom to follow up and probe allows the depth the research question requires. Participants were recruited through the researcher's professional network in the industry, supplemented by direct LinkedIn outreach to senior practitioners whose public profiles indicated direct relevance to the research question. Initial contact explained the research purpose; those who agreed were sent the participant information sheet at least forty-eight hours before the interview.

### **3.1.3 Sample design and rationale**

The first pool is AI-native ventures: founders or senior product leaders at companies that have built workflow automation capability from inception. This pool provides direct evidence on the operating model gap and on how incumbents look from the outside. Participants were also identified through public profiles at IBC and NAB and through DPP membership and event involvement.

The second pool is established incumbent vendors: senior product and technology leaders at live media companies navigating the decision about whether and how to build internal AI-enabled workflow capability. Participants were identified from publicly listed live media technology companies and through direct LinkedIn outreach to CPO, CTO, and VP-level contacts, with priority given to both hardware-origin and software-workflow incumbent types to allow comparison across the spectrum.

The third pool is broadcasters and production service companies: technology and operations decision-makers who buy live media technology and who have direct experience of evaluating, deploying, or building workflow automation tools. Participants were identified through their documented involvement in IBC Accelerator projects, DPP research initiatives, and the researcher's professional network.

The fourth pool is industry and standards bodies: organisations with cross-vendor and cross-broadcaster visibility. Participants were identified through their roles at independent analyst and standards bodies active in the live media technology category.

The target sample was eight to ten, and eight participants were secured across the four pools. The goal is thematic saturation across the stakeholder positions rather than statistical representativeness; the pools are four vantage points on one question, so convergence across them is the primary evidence. By the later interviews the dominant themes were recurring and the propositions stabilising. The venture pool comprises an infrastructure-layer and an

editorial-layer AI-native venture; the incumbent pool, one software-origin and one hardware-origin vendor, allowing comparison across incumbent types; the customer pool, a commercial-broadcaster transformation leader and two senior practitioners whose careers span broadcaster, cloud-infrastructure, and vendor roles; and the industry pool, an independent cross-market analyst body. Interviews were conducted in June 2026.

The adequacy of the sample is more defensibly judged by its information power than by a fixed participant count or a claim of numerical saturation (Malterud, Siersma and Guassora, 2016). Information power holds that the more relevant information a sample carries, the fewer participants are needed; it is assessed across five dimensions, each favouring this study. The aim is narrow, concerning one venture type in one category, which raises information power and lowers the participants required. Sample specificity is dense: participants were senior figures — CTOs, founders, principal product leaders — selected for direct experience of the problem under study, not representativeness. The study is strongly theory-informed, built on established work operationalised as five design propositions, so each interview is read against a developed scaffold rather than as raw discovery. Dialogue quality was high: in-depth, semi-structured exchanges with articulate expert respondents. And the analysis is cross-case, comparing vantage points across four pools rather than seeking depth within one. On each dimension the sample carries high information power, a more rigorous justification for its size than an appeal to saturation alone.

### **3.1.4 Data collection**

Each interview follows a semi-structured guide of twelve questions organised into five sections: warm-up and context; the category shift and its character; the integration paradox and organisational design; competitive dynamics and time horizon; and closing questions including an explicit invitation to challenge the framing. The full guide is provided as an appendix. Core questions are held constant across all interviews to allow cross-respondent comparison; follow-up probes vary by pool to surface the depth of perspective appropriate to each position.

Interviews are conducted by video call, scheduled for forty-five to sixty minutes, and recorded with the respondent's explicit consent confirmed verbally at the start of each recording. Transcripts are produced automatically during the call using AI transcription software and are available immediately after each interview. Initial impressions are written up immediately afterwards while the conversation is fresh. This early write-up is part of the analytical process, capturing what a transcript alone cannot convey — the hesitation before an answer, the

animation with which a respondent described frustration with incumbent processes, the gap between what a founder said they would do and what they actually did.

### **3.1.5 Analysis**

Interview data are analysed through thematic analysis across six phases: familiarisation, initial coding, theme identification, theme review, theme definition, and synthesis. Coding proceeds in two phases.


The first phase is deductive. The five design propositions from the literature review function as initial themes. Evidence from each interview is coded according to whether it confirms, refutes, or complicates each proposition. This structure is appropriate to a research design where the literature has produced specific testable propositions and the goal is to assess and refine them rather than discover them from scratch.

The second phase is inductive. Themes that emerge across interviews independently of the design propositions are identified and developed. This phase is included specifically to guard against the risk that deductive coding forecloses on evidence that does not fit the existing framework. One illustration of how this works in practice: the first interview surfaced a distinction between incumbents who are structurally disposed to create new ventures and those who are structural adopters — waiting for something to be proven before acquiring it. This distinction was not in the literature review and was not anticipated by the design propositions; it emerged from the conversation and entered the analysis through the inductive phase.

Cross-pool comparison is conducted for each theme: evidence from each of the four stakeholder pools is examined and compared. Convergence across pools strengthens the evidence base; divergence indicates that a proposition may need to be qualified or made contingent on context.

### **3.1.6 Triangulation and validity**

The design strengthens validity in two ways. Source triangulation across the four pools ensures no single perspective dominates: a claim recurring across the venture, incumbent, customer, and industry pools carries more weight than one from a single vantage point, and divergences are treated as findings in their own right. Theoretical triangulation comes from the three bodies of literature in Chapter 2, each a different lens on the interview evidence, reducing the risk that any single frame distorts the reading.



*Convergence across pools, rather than depth within any one, is the primary evidence.*

*Figure 3.1 Four-pool source triangulation on a single research question.*

### 3.1.7 Ethics

The research follows the University of Bath School of Management ethics policy for interview-based research with consenting adult professionals. All participants received an information sheet at least forty-eight hours beforehand, confirmed consent on the recording, and are pseudonymised by default; any direct quotations are reviewed by participants before publication; recordings are deleted within six months of submission. No vulnerable populations or sensitive personal data are involved beyond professional perspectives. The information sheet and consent form are provided as appendices.

### 3.1.8 Positionality and researcher reflexivity

The researcher is a current senior professional in the live media technology industry with direct experience of the strategic decisions examined in this research. This positionality has consequences that must be acknowledged clearly, because it is both the most significant advantage and the most significant risk in the research design.

The advantage is access. The researcher's professional network enabled interview recruitment at a level of seniority — with chief technology officers, product leaders, and venture founders — that would have been difficult to achieve from outside the industry. This directly improves the quality of evidence the research can produce.

The risk is confirmation bias: an active industry professional may unconsciously weight evidence that aligns with existing assumptions about how the industry works and what the right strategic answer is.

Several practices are built into the research design to manage this risk. The interview guide includes an explicit question at the close of every interview inviting the respondent to identify what the conventional analysis of the shift gets wrong — a deliberate invitation to challenge the researcher’s framing. The inductive coding phase is designed to surface evidence that the deductive framework would otherwise miss. One participant, the founder of an AI-native infrastructure venture, previously held a senior product role at the researcher’s employer before founding his company. All questions in that interview were framed at the strategic and generic level, no commercially sensitive territory was addressed, and the interview is treated with the same analytical discipline as every other in the sample.

The researcher’s positionality is therefore disclosed transparently rather than managed away. It is the condition that makes the research possible and the condition that requires the most careful methodological management.

### **3.1.9 Limitations**

The sample of eight is small in absolute terms and the findings cannot be generalised statistically; the design requires theoretical saturation across stakeholder positions, defended through information power above. Two limitations remain. With roughly two participants per pool, it would be overstated to claim full saturation of any single perspective in isolation, though cross-pool convergence mitigates this. And the supply side is more heavily represented than the demand side: though the customer perspective is captured through a broadcaster transformation leader and practitioners with broadcaster experience, more currently serving broadcaster operations leaders would have strengthened it. The framework offers design propositions for a specific incumbent type in a specific category; its applicability elsewhere is for future research to establish.

The sample is geographically concentrated in the United Kingdom, Western Europe, and North America, reflecting the researcher’s network and the distribution of decision-making authority in the industry. The framework’s applicability to Asia and the Middle East cannot be claimed from this sample, though live production’s operational dynamics are global enough that the core propositions are unlikely to be fundamentally geographically bounded.

The hypothetical focal venture used to illustrate the framework is deliberately generic rather than based on a named company's actual initiative. This protects commercial confidentiality and ensures the framework is not read as a business plan for a specific organisation. The trade-off is that the framework has not been empirically validated against a real implementation. That is a legitimate limitation and one that points toward the most productive direction for future research: a longitudinal study of an actual incumbent venture attempting to navigate the integration paradox in practice.

## 3.2 Findings: Cross-Interview Analysis

*Pseudonym key: V1 venture/infrastructure; V2 venture/editorial; I1 incumbent/software; I2 incumbent/hardware origin; C1 customer/broadcaster; C2 cross-ecosystem (media-org product leader, former vendor and cloud-platform); C3 cross-ecosystem (ex-broadcaster and cloud-platform AI specialist); A1 industry/standards body.*

### 3.2.1 Introduction and structure

This chapter presents the findings of the generative phase: a thematic analysis across the four stakeholder pools, organised around the five design propositions advanced in Chapter 2. The emphasis throughout is on convergence and divergence across the four pools: a claim recurring across perspectives is treated as more robust than one from a single vantage point, and disagreements between pools are reported as findings in their own right. Interpretation is reserved for Section 3.3.

### 3.2.2 Operating-model primacy (DP3)

Across all eight interviews, and from every pool, respondents located the decisive factor in the venture's operating model rather than its technology, treating technical capability as a baseline and the operating model as what distinguishes one venture from another. V1 contrasted a focused venture's iteration speed with an incumbent's as flying a fighter jet rather than a passenger aircraft. A1 distinguished vendors that sell features from those that sell business value. I1 described the lever as organisational design and hiring. C1 advised focusing on business-driving strategy rather than over-listening to feature requests, citing a vendor that cannibalised itself by being too reactive. I2 reduced it to the right mix of skills and mutual respect between domain experts and platform engineers. C2 urged vendors to refocus on the customer's problem; V2 insisted AI workflows must be owned by the business, not improvised; and C3 reduced it to a principle: AI is a tool whose value is realised through deliberate deployment to specific tasks, not through the technology itself.

DP3 was retained as stated, the most securely grounded of the five.

### 3.2.3 The data moat (DP4)

The proposition that an incumbent's operational data constitutes its primary convertible advantage was contradicted by every respondent who addressed it, most sharply by the source best placed to defend it. The hardware incumbent (I2), whose organisation holds more operational and signal data than any other participant, dismissed it outright: data was backward-looking and effectively worthless, while the customer relationship was vastly more important.

The software incumbent (I1) named brand, relationships, and domain excellence as the firm's assets and did not mention data. The infrastructure-venture founder (V1) argued that data-privacy constraints sharply limit the convertibility of any data advantage. The cross-ecosystem respondent (C2) warned that over-reliance on data synthesis undervalues human reasoning, characterising the language model as predicting the next word rather than reasoning.

V2 supplied the explanation the others only asserted: data is the moat in data-centric industries, but live media is not one. Its monetisable asset is the content itself — video, audio, and graphics, not metadata — so the operational data the proposition assumed valuable was not what the industry traded on.

DP4 was therefore not retained as stated. It was reframed: in live media the incumbent's convertible advantage lies in its relationships, reputation, and domain expertise, not in its operational data. The significance of this divergence is developed in the discussion.

#### **3.2.4 Structural separation (DP1)**

The proposition that an AI venture must be structurally separated while retaining integration access was supported, but separation proved one of several viable arrangements rather than a universal prescription. The software incumbent (I1) had implemented it in recognisable form: a distinct team built around an externally recruited entrepreneurial leader and a cloud-native engineering lead, also staffed with experienced internal personnel to preserve domain knowledge, and connected to the core through a deliberate product-management thread. The hardware incumbent (I2), by contrast, had not separated at all; its platform was built by a blended team held together by the right mix of skills and mutual respect, with the firm's broadcast heritage preserved at the core. His conclusion was explicit: the question was less about organisational structure than the right mix of skills.

Two further accounts sharpened the picture. V2 described a strong newsroom-software acquisition ruined when the acquirer layered the new product onto its legacy system rather than letting it supersede the old; I2 independently observed that a vendor acquiring an orchestration company loses its power, because absorption compromises the independence that made the firm effective. DP1 was therefore retained but qualified: structural separation with integration access is one valid model, blended integration is another, and the choice depends on whether the venture must escape the legacy business and on the incumbent's scale.

#### **3.2.5 Architectural competence as liability (DP2)**

The proposition that an incumbent's deep architectural competence becomes a liability when the category shifts received qualified support, and was the only proposition respondents actively argued against. The analyst (A1) and the software incumbent (I1) both described the commoditisation of code and the consequent vulnerability of vendors whose advantage rests on features. The broadcaster (C1) identified a financial mechanism beneath the architectural one: a hardware-origin vendor moving to subscriptions must accept reduced revenue before subscription revenue scales. The editorial-venture respondent (V2) added a scale dimension, citing a hardware incumbent's wholesale platform pivot as moving too early.

The contesting view came from the incumbents, and deserves reporting rather than dismissal. I2 argued that deep domain competence is not a liability but a transferable asset: the accumulated understanding of how to operate reliably, securely, and at scale in a live environment is precisely what an AI-native entrant lacks and cannot quickly acquire, so a vendor carrying that competence into the new category begins with an advantage. On this view the danger is not the competence but misreading which part to carry forward — the firm that mistakes attachment to a product form for the underlying capability is encumbered, while the firm that retains the capability and sheds the product form is strengthened. I1 added that the liability is real only where the organisation cannot separate the two: a deliberately structured unit can preserve the domain knowledge while escaping the legacy assumptions. The evidence therefore did not refute the proposition so much as qualify it: architectural competence is a liability when it manifests as attachment to a superseded product form, and an asset when it manifests as transferable domain capability. DP2 was accordingly retained as confirmed but not deterministic, its effect contingent on which form the competence takes.

### **3.2.6 The time-bounded window (DP5)**

The proposition that incumbents face a limited but real window was confirmed across all pools, with a consistent recalibration: respondents judged the window wider than the prevailing urgency implies. I1 judged market windows wider than commonly thought, while warning they are narrow at the individual anchor-customer level, where missing a procurement cycle can forfeit a relationship for years. I2 held that the industry has more time than it feels, locating the constraint in the slow human-adoption cycle within broadcasters. V2 judged that the world moves slower than it seems, characterising live media as an AI-adoption laggard. C3 reinforced this: an incumbent's relationship moat buffers it against moving slowly, so the genuine fast-

mover advantage lies less with established vendors than with smaller innovators and the fast-cycle creator-economy segment incumbents largely underserve.

The most useful refinement came from the cross-ecosystem respondent (C2), who distinguished two clocks: the product clock — the pace at which a finished AI product must reach the market — is relatively slow, but the capability clock — the pace at which an organisation must begin building AI competence — runs faster, because capability compounds. DP5 was retained, recalibrated along these lines.

### **3.2.7 Cross-cutting themes not anticipated by the propositions**

Several themes emerged consistently that the propositions had not anticipated. The most important was the integration layer as the industry's central unmet need, named in different language across all four pools. The cross-ecosystem respondent (C3) supplied the sharpest demand-side form: the most consistent message from buyers was a refusal to purchase further standalone tools and a preference for the capability built into systems they already operated. That the demand side named the industry's biggest gap as precisely the capability the focal venture is conceived to provide — and wanted it delivered through the incumbent's existing tools — is notable corroboration of the study's premise, and points to the incumbent's installed base as a natural channel for it.

A second was the primacy of internal productivity over product features as the near-term locus of AI's impact. The cross-ecosystem respondent (C3), who had worked directly on AI applications across the sector, located the mechanism: customer-facing and on-air generative uses were gated less by technical maturity than by copyright-litigation risk and craft-union sensitivity, so the value that shipped concentrated in back-office workflows such as localisation, transcription, and translation. A third was a consistently sober calibration across all eight respondents — a trough of disillusionment, investment running ahead of value, and a subsidised era of artificially low cost that may not persist. A fourth concerned competition: respondents argued against a winner-takes-all reading, holding that AI raised the baseline while leaving room for differentiation above it, so incumbents, entrants, and self-building customers would coexist.

### **3.2.8 Summary of proposition status**

The table below summarises the status of the five design propositions after the cross-interview analysis. It is a reader's anchor; the substantive evidence is in the preceding prose.

| Proposition                                         | Status                                     | Basis                                                                                                                                     |
|-----------------------------------------------------|--------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------|
| DP1 — Structural separation with integration access | Confirmed but contingent                   | Two valid models (separation vs blended integration); choice depends on whether the venture must escape the legacy business, and on scale |
| DP2 — Architectural competence as liability         | Confirmed but not deterministic; contested | A liability where it anchors a feature/hardware conception; the underlying domain competence can transfer and become an asset             |
| DP3 — Operating-model primacy                       | Strongly confirmed                         | Endorsed from every pool; the most securely grounded proposition                                                                          |
| DP4 — Operational data as moat                      | Not retained; reframed                     | Live media is asset-centric, not data-centric; the durable moat is human and tacit, not informational                                     |
| DP5 — Time-bounded window                           | Confirmed; recalibrated                    | Wider than the prevailing urgency implies; product clock slow, capability clock fast                                                      |

In summary, the cross-interview analysis confirmed two propositions strongly, reframed one, qualified one as contingent, and recorded one as confirmed but contested. Across these results a consistent pattern emerged in the evidence: respondents from every pool located the venture's prospects in its operating model, its relationships and domain judgment, and its organisational capability, rather than in the technology or the operational data the propositions had foregrounded. The integration layer recurred, unprompted, as the capability the industry most lacked. These results, and the connections between them, are interpreted in the discussion that follows.

### **3.3 Discussion**

This chapter interprets the findings reported in Section 3.2. The study asked how an established live media vendor should organise and execute an AI-enabled cross-portfolio workflow automation venture, testing five design propositions derived from the literature. The discussion draws the findings together into an answer, relates them to the prior work from which the propositions were built, considers alternative explanations, and states the limits of the interpretation. It does not re-walk the propositions individually or introduce new evidence; it works only from what Section 3.2 established.

#### **3.3.1 The organising argument: technology as precondition, not differentiator**

Read individually the findings appear as five separate results; read together they answered the research question with one underlying claim: the venture should be organised around the recognition that, in this category, technology is a necessary condition of competing but not the source of durable advantage, which is instead human and tacit. A venture without technical capability fails regardless of how it is organised, but as the technology becomes broadly available it ceases to be where differentiation lies. Each of the three most strongly supported findings expressed this claim: the operating-model finding located success in organisational design; the data-moat reversal dismissed the one asset assumed decisive because it was not human or tacit; and the time-window finding held that what had to be built early was capability, not a product.

This speaks directly to the general-purpose-technology literature the study drew on. Brynjolfsson and McAfee (2014) and Iansiti and Lakhani (2020) argue that a transformative technology reshapes competition, and the intuitive reading is that advantage accrues to whoever commands the best version of the technology or the most data to feed it. The findings qualify rather than contradict that view: consistent with Brynjolfsson, Rock and Syverson (2021), who show the returns to a general-purpose technology depend on complementary organisational change, the evidence here located advantage in the assets surrounding the technology, not the technology itself. As the technology became a shared input, advantage migrated back toward what it could not supply. This is also consistent with Agrawal, Gans and Goldfarb (2018): when prediction becomes cheap and ubiquitous, value shifts to its complements — judgment and the assets that contextualise it — which is what respondents described.

#### **3.3.2 A bounded divergence from complementary-assets theory**

The data-moat reversal is the study's **most consequential finding** and where it engages theory most directly. The proposition was built from the complementary-assets logic of Teece (2007): the assets an incumbent needs to commercialise an innovation are where its durable advantage lies. The data-moat reading is one instantiation of that logic, holding where the firm's monetisable asset is informational, as in adjacent data-centric industries. Yet every respondent in live media rejected it, most emphatically the incumbent holding the most data.

The reconciliation is that Teece's argument is not about data specifically but about whichever assets are complementary in a given industry's value structure. Live media is **asset-centric rather than data-centric**: its monetisable asset is content, and the operational data the proposition assumed valuable is neither what the industry trades on nor, by respondents' accounts, reliably usable. The complementary assets that matter are therefore human and tacit. The contribution is thus a bounded divergence from the data-moat reading, not a refutation of Teece: it specifies when the data-moat expectation holds and when it inverts. In doing so it aligns with the resource-based and dynamic-capabilities tradition (Eisenhardt and Martin, 2000; Teece, 2007), in which tacit, relational, hard-to-imitate resources, rather than codified data, are the foundation of sustained advantage. The evidence further suggested even the conditional data advantage that exists is unstable: one respondent observed that consolidation and senior-leadership churn can render a vendor's accumulated knowledge of a customer obsolete and even misleading, so the data asset not only fails to convert automatically but decays.

### 3.3.3 When competence is a liability and when it is an asset

The architectural-competence proposition, built from Henderson and Clark (1990) and Christensen (1997), holds that an incumbent's established competence becomes a liability when the basis of competition shifts. The findings confirmed this only conditionally: the incumbents argued, with evidence, that their competence was an asset that transferred. The disagreement is resolved by distinguishing two senses of competence. Where it denoted attachment to a superseded product form — a hardware conception, a feature-led proposition, a perpetual-licence model — the Henderson and Clark and Christensen prediction held. Where it denoted the underlying domain capability — operating reliably, securely, at scale in a live environment — that capability transferred and became an asset.

This refines the disruption account rather than simply endorsing it, as the critical literature anticipated. King and Baatartogtokh (2015) and Danneels (2004) note that disruption theory

under-specifies when incumbent assets become liabilities and when they remain advantages; the present finding supplies one such specification for this category. It also connects to the ambidexterity literature (O'Reilly and Tushman, 2008; Birkinshaw and Gibson, 2004): carrying forward domain capability while shedding attachment to the old product form is the simultaneous exploitation-and-exploration ambidexterity describes, here located at the level of competence rather than structure alone.

### 3.3.4 The integration paradox as the connecting structure




Figure 3.2 The integration paradox.

The findings not anticipated by the propositions — above all the integration layer as the industry's central unmet need — connect the foregoing into a single structure, the integration paradox. It was clearest in the acquisition evidence, where the acquired venture lost its effectiveness when absorbed into the parent rather than allowed to supersede the legacy product. This is the tension the corporate-venturing literature describes but does not fully resolve. Burgelman (1983) showed that autonomous internal ventures must be protected from the parent's resource-allocation logic, and Christensen (1997) likewise prescribed separation; yet both concerned ventures whose goal was to become independent businesses. The venture studied here is different: its advantage depends on remaining integrated with the parent's assets and customers, so separation cannot be the whole answer.

The integration paradox names this gap. The venture must be close enough to the incumbent to reach its complementary assets, yet independent enough that its orchestration is trusted as neutral and its pace is its own. This extends the ambidexterity tradition (O'Reilly and Tushman,

2008): where ambidexterity asks how a firm pursues exploitation and exploration at once, the integration paradox specifies a venture whose exploration depends on the parent's exploitation assets, so separation must be calibrated rather than applied wholesale. That a customer described the absence of this orchestration layer as a major gap helps ground the paradox in a demand-side need rather than a purely supply-side ambition.

### **3.3.5 Alternative explanations**

Before these interpretations are accepted, alternative explanations deserve consideration, the most important concerning the data-moat reversal. One reading is that respondents dismissed data not because it is structurally unimportant but because their firms have not yet learned to exploit it — a current capability gap rather than an enduring feature. A second is sample composition: participants were weighted toward product, strategy, and commercial roles rather than data-science roles, and a different pool might have valued data more highly. This weighting was deliberate and appropriate to the unit of analysis. The research question concerns how a venture should be organised and where its advantage lies — questions of organisational design and strategic orchestration resolved at business and product-leadership level, not as engineering problems. The choices to separate or blend, to lead with integration or features, and to treat data or relationships as the moat are made by the roles represented here; a data-science pool would speak to model performance but not to the organisational question. The risk is therefore narrow — a possible under-valuation of data's technical potential — not a misalignment between sample and question. A third is social desirability: incumbents may downplay data advantages they cannot realise, and ventures may do so because emphasising relationships suits their position.

These alternatives are acknowledged rather than dismissed, but two features of the evidence weigh against them. The reversal came most forcefully from the participant best placed to claim a data advantage — the incumbent holding the most operational data — the opposite of what self-interest would predict. And the convergence across the four structurally different pools, with their different incentives, makes a shared bias a less economical explanation than a shared perception of the industry's actual structure. The interpretation offered here is therefore advanced as the most defensible reading, not the only possible one, and the same caution applies to the competence and time-window findings, where the contested and recalibrated status of the propositions is itself reported as a finding rather than resolved by fiat.

### **3.3.6 Limitations of these interpretations**

Two limits should be stated. First, the human-and-tacit-asset argument concerns the present and near-term technology; several respondents were explicit that a more capable, still-affordable generative AI could shift the basis of advantage again, and the analysis does not assume the present settlement is permanent. Second, the reconciliations offered here are interpretive and drawn from a small, tightly scoped sample, advanced as well-grounded readings the convergence across pools makes defensible, not settled conclusions. Further research directions are taken up in the conclusion.

## 3.4 The Framework

### 3.4.1 From propositions to contingent design choices

The five propositions were advanced in Chapter 2 as independent claims, but the findings showed they interact and several are conditional rather than absolute. The framework therefore reconceives them not as a checklist but as a connected set of design choices, organised around the integration paradox and made in a definite order.




Figure 3.3 The venture as a set of contingent design choices.

### 3.4.2 The governing choice: structural form

The first, governing choice concerns the venture's structural relationship to the incumbent, following from the DP1 contingency. The decisive question is whether the venture must escape the incumbent's legacy architecture and revenue model to succeed. Where it must — where the new capability cannot be built on the existing technical base, or its subscription economics would be suppressed by the incumbent's perpetual-licence business — structural separation is indicated: a distinct unit with its own entrepreneurial leadership and engineering culture, connected to the parent through a narrow integration thread rather than absorbed. Where it need not — where the new platform is itself the firm's strategic core rather than a departure — blended integration becomes coherent, with domain capability preserved at the core rather than escaped. The two models are not hypothetical: the findings furnished an instance of each. One software incumbent separated structurally, building its next-generation platform in a distinct

team to escape the legacy product; one hardware incumbent deliberately did not, building with a blended team and preserving its broadcast heritage at the core, because that platform was the company's strategic future. The framework does not prefer one model; it identifies the condition — **the escape question** — that determines which applies.

This choice is conditioned by scale. For a large incumbent able to sustain both the existing business and the venture, either model is viable and the decision rests on the escape question above. For a smaller vendor, the findings indicated a standalone internal venture is often unaffordable, and wholesale transformation of the core, or acquisition, is the more realistic route. Where acquisition is chosen, the evidence was unambiguous on the failure mode to avoid: the acquired venture must run independently and, where necessary, supersede the legacy product, not be absorbed and stripped of the independence that gave it value.

### **3.4.3 The governing principles: how the chosen structure is run**

Once the structural form is set, four principles govern how the chosen structure is run. They are not independent prescriptions but follow in order from the study's central claim: because technology is a precondition rather than a differentiator, the operating model comes first; what that operating model must exploit is the reframed, human-and-tacit locus of advantage; realising that advantage requires disentangling the competence worth carrying forward from the product form worth shedding; and all of this must be timed against the two clocks the evidence distinguished.

The first is operating-model primacy. Competent technology is a precondition but does not differentiate, because it is increasingly available to every entrant. What differentiates is how the venture is organised to understand and serve a customer problem: its decision speed, customer proximity, team composition, and the metrics by which it allocates resources. The venture's design effort should therefore go to its operating model — hiring, team structure, iteration cadence — not to a technical edge that will not persist.

The second is the reframed locus of advantage. The venture should be designed to exploit the incumbent's durable, human and tacit assets — relationships, reputation for reliability, and domain judgment — because in this asset-centric industry these, not operational data, are what an AI-native entrant cannot quickly replicate. The practical implication is a governance one: the venture must retain privileged access to the incumbent's customer relationships and domain

experts — an argument for connection to the parent — while not mistaking the parent's accumulated data for the source of its advantage.

The third principle is the disentangling of competence. The incumbent's architectural heritage contains two things that must be separated: an attachment to a superseded product form, which is a liability, and a transferable domain capability — reliability, security, operation at scale — which is an asset. The venture's design must carry the latter forward while shedding the former. Operationally this is a talent and culture question: the venture needs people who embody the domain capability without being bound to the old product — the blended-team logic the hardware incumbent described, and why its platform was built from a broadcast mindset rather than retrofitted onto one.

The fourth is the two-clock discipline. The product clock and capability clock run at different speeds, and the venture should heed both. The finished product can reach the market on the slower product clock, but the organisation must begin building AI capability immediately on the faster capability clock, because that capability compounds and cannot be acquired late without disproportionate cost. An incumbent should therefore stand up the venture — team, tooling, early customers — well before it has a finished product, accepting an imperfect start for the compounding capability it builds.

#### **3.4.4 The organising purpose**

These choices serve the venture's reason to exist, identified by the demand side as the industry's central unmet need: the integration and orchestration layer connecting a broadcaster's fragmented systems. The framework's contribution is to make the tension explicit and convert it into a set of deliberate, ordered design decisions: the integration paradox, resolved not by choosing integration or separation but by calibrating both around the venture's purpose.

## **Chapter 4 — Business Model and Financial Viability of the Focal Venture**

This chapter translates the framework into a concrete venture design and tests its commercial logic. Where the framework establishes how an established live media vendor should organise an AI-enabled workflow automation venture, this chapter sets out what the venture is, how it creates and captures value, and whether it is financially viable. The focal venture is deliberately hypothetical and generic, both to protect commercial confidentiality and to keep the analysis at the level of strategic logic. The financial model presented is illustrative: it is designed to demonstrate the commercial logic and the financial signature of the integration paradox, not to forecast returns.

### **4.1 The focal venture**

The focal venture is an AI-enabled cross-portfolio workflow automation capability built inside an established hardware-origin live media vendor. Its value proposition is anchored in the strongest finding: the absence of a standard integration layer across the fragmented systems a broadcaster operates. The venture is that layer — an AI-triggered orchestration capability connecting the incumbent's portfolio and third-party systems, automating handoffs that currently require manual intervention. It is scoped to the enterprise and supply-chain layer, where AI adoption is real and return-driven, not the creative layer, where adoption is slow.

### **4.2 Business model**

The venture's business model is summarised as a Business Model Canvas (Osterwalder and Pigneur, 2010) in Figure 4.1, adapted here to map each block to the integration paradox: blocks shaded toward integration draw on the incumbent and are the source of advantage, while those shaded toward separation require the venture's autonomy and are the condition of survival. The canvas is thus a synthesis of the findings, not a neutral inventory.

**Figure 4.1 Business Model Canvas for the focal venture, mapped to the integration paradox**




*Figure 4.1 Business Model Canvas for the focal venture, mapped to the integration paradox.*

Three elements carry most of the model's logic. The value proposition is the integration and orchestration layer identified as the central unmet need, working with existing infrastructure rather than replacing it. The key resources weight toward the incumbent's human and tacit assets — relationships, reputation, domain judgment — not operational data. The channel exploits the incumbent's sales force and brand, lowering customer-acquisition cost; it is the most tangible expression of the complementary-asset advantage, realised only where the venture keeps integration access to the parent.

The revenue model is where the most consequential design choice sits. It is generated through a multi-tenant software-as-a-service (SaaS) model, a deliberate departure from the incumbent's perpetual hardware-licence heritage and a necessary one, since bespoke single-tenant builds are not scalable or profitable for the vendor. The pricing of that model, however, is not settled: the classical SaaS pattern of a fixed periodic subscription is, for an AI-enabled product, increasingly blended with consumption-based pricing in which the customer pays in proportion to usage — the volume of content processed, or the quantity of model tokens consumed. The venture therefore combines a subscription floor with a consumption-based component, a mixed model whose consequence for margin the financial analysis examines.

### 4.3 Organisational design

The venture's structure applies the framework's governing choice. Because the orchestration platform must escape the incumbent's legacy architecture and revenue model, the structural-separation branch is indicated: a distinct unit with an externally recruited entrepreneurial leader and cloud-native engineering alongside experienced internal staff, connected to the parent through a product-management thread — far enough from the core to move differently, close enough to draw on the complementary assets that constitute its advantage.

#### 4.4 Financial model

A distinction should be drawn between what the model evidences and what it merely illustrates. Its structure is grounded directly in the primary research: the shape of the transition trough was described by a broadcaster, who observed that hardware-origin vendors must accept short-term revenue loss for long-term sustainability; the cost composition, a largely fixed venture team alongside usage-scaling compute, follows from the operating model the findings established; the multi-tenant model rather than bespoke single-tenant builds was identified in interview as a commercial necessity; and the downside compute-cost risk reflects the editorial-venture respondent's warning about a subsidised phase of the market. Illustrative are the magnitudes — revenue figures, trough depth, year of inflection — chosen to demonstrate commercial shape rather than forecast returns. The model's analytical work lies in translating the qualitative findings into a financial logic, not projecting numbers.

The model projects the venture over five years: subscription revenue scales from the two tiers as adoption compounds, while the cost base is dominated by the separated venture team, usage-scaling compute, and a GTM cost reduced by the incumbent's channel. The single most important feature of the model is the transition trough: as customers migrate from legacy hardware and perpetual licences to subscriptions, the incumbent displaces a portion of its own revenue before the new revenue has scaled to replace it. A further cost risk, raised by the editorial-venture respondent, is that the present cost of AI compute reflects a subsidised phase of the market and may rise materially; the downside scenario reflects this.

The result is the characteristic J-curve shown in Figure 4.2. The operating result deteriorates through the first two years as the venture invests and the transition trough bites, reaches its low, then recovers to an inflection as multi-tenant subscription revenue scales past the largely fixed cost of the venture team. This shape is the **financial signature of the integration paradox**: the venture's advantage depends on integration with the incumbent's installed base and revenue, yet realising it requires deliberately sacrificing legacy revenue in the near term.




*Figure 4.2 Illustrative five-year financial shape — the transition J-curve. Figures are illustrative and rounded, intended to show the commercial shape, not forecast returns.*

#### 4.4.1 Pricing, tokens, and the margin tension

The financial model exposes a tension the shift toward consumption-based pricing makes newly important. On the cost side, an AI-enabled venture is already a token business: its dominant variable cost is compute, which scales with usage and is billed, in effect, per token processed. On the revenue side, a pure subscription recovers a fixed amount regardless of how much the customer consumes. Where those two sides diverge — consumption-based costs beneath subscription-based revenue — the margin is exposed to any movement in the unit cost of compute, since costs rise with usage while revenue does not.

Here the editorial-venture respondent's warning becomes financially concrete: his observation that today's AI compute cost reflects a subsidised phase is, in these terms, a claim about the future token price. If that price rises as the subsidy recedes, a venture priced on fixed subscription absorbs the increase into its margin, while one passing usage through on a consumption basis is insulated. The design choice between subscription and consumption pricing is therefore not merely a commercial preference but a decision about who bears the risk of a volatile, currently understated input cost. The model treats this as a central sensitivity: the base case assumes compute costs remain broadly stable, while the downside scenario reflects a material rise in the token price and the consequent margin compression under a subscription-weighted revenue model.

A further implication is competitive. Consumption-based pricing makes the cost of the service legible to the customer, and this transparency, while attractive to buyers, erodes the lock-in subscription models historically provided, reinforcing the commoditisation pressure identified

in the findings. The venture's defensibility therefore cannot rest on pricing opacity but on the same human and tacit assets the study identified as the durable moat — integration depth, domain judgment, and trusted relationships a transparent unit price does not capture.

#### **4.4.2 Working capital and investment**

The revenue and cost models are set out above; the working-capital and investment models complete the picture. The working-capital position is comparatively light. A multi-tenant SaaS model with subscription billing collects revenue alongside or ahead of delivery, so the venture carries little of the receivables burden a project-based or hardware-sales model would impose; its principal variable cost, compute, is paid in arrears against usage. The net effect is a modest requirement, dominated not by inventory or receivables but by the timing mismatch between largely fixed staff costs, due monthly, and subscription revenue, which builds gradually as the customer base compounds. That mismatch is most acute during the transition trough, where the venture funds a full team against still-thin revenue.

The investment model follows directly from that trough. The capital the incumbent must commit is determined not by the venture's eventual scale but by the depth and length of the trough preceding the inflection — the cumulative operating losses absorbed before subscription revenue overtakes the largely fixed cost of the team. In the illustrative model this is the single most important investment parameter, setting the minimum funding commitment the incumbent must make before the venture turns. The investment is therefore best understood not as a one-off build cost but as a sustained commitment to fund the trough to its far side, and viability, examined next, turns substantially on whether the incumbent has the patience and balance-sheet capacity to do so.

#### **4.5 Scenarios and viability**

Three scenarios frame the range of outcomes. The base case assumes the return-proven adoption pace indicated by the evidence. The upside assumes the integration gap proves acute and the incumbent moves decisively within the competitive window, shortening the trough. The downside assumes the window closes faster — because agentic coding lets customers rebuild commodity capability themselves, because budget pressure and cloud repatriation slow adoption, or because the token price rises as the compute subsidy recedes — deepening and lengthening the trough. The outcome is most sensitive to the adoption ramp and to pricing; the depth of the trough determines how much capital the incumbent must commit before the inflection.

The model supports a conditional judgment of viability. The venture is viable where the incumbent is large enough to sustain managed separation, funds the transition trough, leads with the integration value proposition rather than discrete features, and acts within the competitive window. It is not viable where the incumbent is sub-scale (transformation or acquisition is then better), where the capability is treated as a feature inside the core (the framework's liabilities then dominate), or where the window is missed. The financial analysis therefore returns to the central claim, the paradox is visible not only in the organisational design but in the shape of the venture's financial trajectory.

## Chapter 5 — Conclusion and Reflection

### 5.1 Answering the research question

This dissertation asked how an established live media vendor should organise and execute an AI-enabled cross-portfolio workflow automation venture. The research supports designing the venture around a single tension, **the integration paradox**: its advantage depends on close integration with the incumbent's complementary assets, while its survival depends on enough structural independence to move and be trusted on its own terms. The five propositions, tested across four pools, resolve into a connected set of contingent choices rather than a checklist: structural form is the governing choice, contingent on whether the venture must escape the legacy business and on scale, and the chosen structure is then run by the four principles in the framework — all serving the venture's reason to exist, the integration and orchestration layer the demand side identified as the industry's central unmet need.

### 5.2 Contribution

The study's contribution is empirical and theoretical at once. Empirically, it brings together venture, incumbent, customer, and industry perspectives on a question usually examined from one side, and finds a notable convergence: the durable advantage is human and tacit, while technology is a condition of competing rather than the basis of differentiation. Theoretically, it specifies a **bounded divergence from the complementary-assets reading** of incumbency: the expectation that operational data is the convertible moat holds in data-centric industries but inverts in asset-centric live media, where the monetisable asset is content and the durable advantage relational.

### 5.3 Reflection in action

My understanding of the problem changed more than once, and the present view is the product of those revisions rather than my starting assumptions. I began expecting operational data to be the incumbent's decisive asset and the time pressure to be acute; the evidence pushed me to reframe both, relocating the durable advantage in human and tacit assets and recalibrating the window as wider than the prevailing urgency implies. The business-model solution — the managed-separation venture organised around the integration-and-orchestration layer — therefore rests on a problem definition the research reshaped. It is valid but conditional: it holds where the incumbent is of sufficient scale, funds the transition trough, and leads with the integration value proposition. The most significant unresolved uncertainty is the pricing model,

where the shift from subscription toward consumption and token-based revenue is live and unsettled, and the venture's margin exposure greatest.

#### **5.4 Reflection on action**

The design served the question well in some respects and could have been steered differently in others. Testing propositions across four structurally different pools, rather than seeking depth within one, was right for an industry-wide question: the convergence gave the findings their weight and exposed the data-moat reversal most clearly. The discipline of letting the interviews contradict the propositions rather than defending them was the most valuable methodological lesson, turning a confirmatory exercise into a genuine reframing. Conducting the research while working in the industry was both advantage and risk: it gave access and fluency an outsider could not have had, but required deliberate steps to keep my priors from leading the analysis, including excluding a close contact from the pool. With hindsight I would have begun interviews earlier to allow a further iteration, and two aspects remain under-examined: the customer-side perspective is thinner than the vendor and venture ones, and the pricing question surfaced late.

#### **5.5 Limitations and further research**

The study's limits are those of a small, tightly scoped qualitative design: its claims rest on convergence across pools rather than depth within any one, and full saturation of any single perspective cannot be claimed. The framework is offered for a specific incumbent type in a specific category, and its transfer to others is a matter for future work. Two directions follow most naturally. The first is to test the framework against a live attempt to build such a venture, observing the contingent choices in practice. The second is to examine the pricing question the financial analysis raised, as consumption and token-based models displace subscription, and to trace what that shift does to the locus of advantage the study identified.

#### **References**

- Acemoglu, D., 2023. Automation and the future of work: challenges and opportunities. *AEA Papers and Proceedings*, 113, pp.66–72.
- Agrawal, A., Gans, J. and Goldfarb, A., 2018. *Prediction machines: the simple economics of artificial intelligence*. Boston: Harvard Business Review Press.
- Amagi, 2025. *Cloud solutions for broadcast, CTV and FAST*. [online] Available at: <https://www.amagi.com> [Accessed May 2026].

- Birkinshaw, J. and Gibson, C., 2004. Building ambidexterity into an organization. *MIT Sloan Management Review*, 45(4), pp.47–55.
- Brynjolfsson, E. and McAfee, A., 2014. *The second machine age: work, progress, and prosperity in a time of brilliant technologies*. New York: W. W. Norton & Company.
- Brynjolfsson, E., Rock, D. and Syverson, C., 2021. The productivity J-curve: how intangibles complement general purpose technologies. *American Economic Review: Insights*, 3(3), pp.333–350.
- Burgelman, R.A., 1983. A process model of internal corporate venturing in the diversified major firm. *Administrative Science Quarterly*, 28(2), pp.223–244.
- Christensen, C.M., 1997. *The innovator's dilemma: when new technologies cause great firms to fail*. Boston: Harvard Business School Press.
- Cuez, 2026. *Cuez brings four new innovations to NAB 2026: from story-centric newsroom to open AI agent framework*. [online] Available at: <https://cuez.app/blog/press-release-cuez-brings-four-new-innovations-to-nab-2026-from-story-centric-newsroom-to-open-ai-agent-framework/> [Accessed May 2026].
- Dalet, 2025. *Dalet unveils Dalia agentic AI media workflows at IBC2025*. [online] Available at: <https://dalet.com/news/dalia-agentic-ai-media-workflows-ibc2025/> [Accessed May 2026].
- Danneels, E., 2004. Disruptive technology reconsidered: a critique and research agenda. *Journal of Product Innovation Management*, 21(4), pp.246–258.
- DPP, 2024. Industry analysis reports on AI and workflow automation in live media production. Digital Production Partnership. [online] Available at: <https://www.thedpp.com> [Accessed May 2026].
- Eisenhardt, K.M. and Martin, J.A., 2000. Dynamic capabilities: what are they? *Strategic Management Journal*, 21(10–11), pp.1105–1121.
- Harmonic, 2026. *Harmonic introduces breakthrough AI orchestration service for live video workflows*. [online] Available at: <https://investor.harmonicinc.com/news-releases/news-release-details/harmonic-introduces-breakthrough-ai-orchestration-service-live/> [Accessed May 2026].

- Henderson, R.M. and Clark, K.B., 1990. Architectural innovation: the reconfiguration of existing product technologies and the failure of established firms. *Administrative Science Quarterly*, 35(1), pp.9–30.
- Iansiti, M. and Lakhani, K.R., 2020. *Competing in the age of AI: strategy and leadership when algorithms and networks run the world*. Boston: Harvard Business Review Press.
- IBC, 2025. *Accelerator project 2025: AI agent assistants for live production*. [online] Available at: <https://show.ibc.org/accelerator-project-ai-assistance-agents-live-production> [Accessed May 2026].
- King, A.A. and Baatartogtokh, B., 2015. How useful is the theory of disruptive innovation? *MIT Sloan Management Review*, 57(1), pp.77–90.
- Lepore, J., 2014. The disruption machine: what the gospel of innovation gets wrong. *The New Yorker*, 23 June. [online] Available at: <https://www.newyorker.com/magazine/2014/06/23/the-disruption-machine> [Accessed May 2026].
- Malterud, K., Siersma, V.D. and Guassora, A.D., 2016. Sample size in qualitative interview studies: guided by information power. *Qualitative Health Research*, 26(13), pp.1753–1760.
- Nexog, 2025. *NEXOG AIOS: cloud-native AI-driven operating system for media and network workflow*. [online] Available at: <https://nexog.com> [Accessed May 2026].
- O'Reilly, C.A. and Tushman, M.L., 2008. Ambidexterity as a dynamic capability: resolving the innovator's dilemma. *Research in Organizational Behavior*, 28, pp.185–206.
- Raisch, S. and Birkinshaw, J., 2008. Organizational ambidexterity: antecedents, outcomes, and moderators. *Journal of Management*, 34(3), pp.375–409.
- Teece, D.J., 2007. Explicating dynamic capabilities: the nature and microfoundations of (sustainable) enterprise performance. *Strategic Management Journal*, 28(13), pp.1319–1350.
- Winter, S.G., 2003. Understanding dynamic capabilities. *Strategic Management Journal*, 24(10), pp.991–995.

## Appendices

---

### Appendix A — Interview Guide (Semi-Structured Protocol)

The semi-structured guide below was used across all four stakeholder pools. Each interview opened with a short briefing and recorded consent, and closed with a request for onward referrals; these procedural scripts are summarised in the methodology (Section 3.1) and are not reproduced here. Pool-specific follow-up prompts (Venture, Incumbent, Customer, Industry) drew out each group's perspective on the same underlying questions, and the five design propositions (DP1–DP5) are noted against the questions that primarily tested them.

**Part 1 · Warm-up.** Q1: current role and its connection to live media workflow automation. Q2: the most significant shift in the industry right now (noting whether AI, workflow, or cost pressure is raised unprompted).

**Part 2 · The category shift.** Q3: how much of AI-driven workflow automation is a genuine structural shift versus hype (Venture — real buying decisions or pilots?; Customer — deciding differently today or still evaluating?; Industry — where is adoption real versus conceptual?). Q4: the most significant gap between what broadcasters need operationally and what the market provides (Venture — is that the gap you address?; Customer — solving it internally or waiting on vendors?; Industry — is it documented/raised by members?). Q5: what AI-native ventures do differently from established vendors, and where each is advantaged or constrained (Venture — could you have built this inside an incumbent?; Customer — what does each side do better in your evaluations?; Industry — can incumbents match entrants, or is the gap structural?).

**Part 3 · The integration paradox (analytical core).** Q6 (DP1): build the capability inside the existing organisation, as a separate venture unit, or some other arrangement — and why (Venture — what would an acquirer gain and risk?; Customer — does it matter to you where the capability sits?; Industry — have you seen incumbents attempt this, and what happened?). Q7 (DP4): what incumbents have — data, customer relationships, signal infrastructure — that an independent venture cannot easily replicate, and how much is actually convertible into advantage (Venture — would incumbent data genuinely improve your models?; Customer — genuine advantage or data-governance concern?; Industry — do interoperability standards help or hinder?). Q8 (DP3): how much of success or failure comes down to organisational

design versus technology versus leadership (Venture — the single hardest thing for an incumbent to replicate?; Customer — biggest barrier on a comparable transformation?; Industry — the usual proximate cause when incumbents fail?).

**Part 4 · Time and competitive dynamics.** Q9 (DP5): how fast the landscape is evolving and whether incumbents have a window to respond — and how long (all pools: what would close the window faster, and what would extend it?). Q10: the market in three to five years — who wins, who loses, and why (Venture — where is your own company?; Industry — what role for standards and interoperability?).

**Part 5 · Closing.** Q11: the one thing the conventional analysis of this shift gets wrong or misses (all pools; space left for an unprompted answer). Q12: recommendations of others to speak with.

---

## Appendix B — Participant and Pseudonym Matrix

All participants are pseudonymised by code. Named organisations are withheld in accordance with the pseudonymisation protocol; the role descriptions are deliberately general for the same reason.

| Code | Pool              | Role description                                                               | Date<br>(2026) | Approx.<br>duration |
|------|-------------------|--------------------------------------------------------------------------------|----------------|---------------------|
| V1   | Venture           | Founder, AI-native infrastructure-layer workflow venture                       | 4 June         | ~45 min             |
| V2   | Venture           | Leader, AI-native editorial-layer workflow venture                             | 19 June        | ~45 min             |
| I1   | Incumbent         | Product EVP, software-origin media workflow vendor                             | 8 June         | ~45 min             |
| I2   | Incumbent         | Product marketing director, hardware-origin live media vendor                  | 17 June        | ~45 min             |
| C1   | Customer          | Transformation leader, commercial broadcaster                                  | 9 June         | ~45 min             |
| C2   | Customer / bridge | Media-organisation product leader; ex-cloud-infrastructure, ex-hardware-vendor | 18 June        | ~45 min             |

| Code | Pool                 | Role description                                                            | Date<br>(2026) | Approx.<br>duration |
|------|----------------------|-----------------------------------------------------------------------------|----------------|---------------------|
| C3   | Customer /<br>bridge | Ex-broadcaster and ex-cloud-platform<br>media & entertainment AI specialist | 22 June        | ~45 min             |
| A1   | Industry             | Senior figure, independent cross-market<br>standards / analyst body         | 5 June         | ~45 min             |

*The four pools are venture (2), incumbent (2), customer (3, including two cross-ecosystem bridges), and industry (1) — eight participants. All interviews were conducted in June 2026 and ran approximately 45 minutes each.*

## Appendix C — Thematic Coding Tree

The analysis followed a theory-informed thematic approach. Raw interview material was coded into **first-order codes** (in participants’ own framing), aggregated into **second-order themes** (the analyst’s conceptual grouping), and consolidated into **aggregate dimensions** that map onto the design propositions and the emergent findings. The structure below sets out that three-tier aggregation; parenthetical codes (V1, I2, C3, etc.) indicate the pool/participant where a first-order code was most salient, though most codes recurred across pools, and cross-pool recurrence is what elevated a code to theme status.

Aggregate dimension 1 — Operating model is the locus of advantage (→ DP3)

| First-order codes (raw)                                                                                           | Second-order theme                                   |
|-------------------------------------------------------------------------------------------------------------------|------------------------------------------------------|
| “code is commoditised”; “everyone has the same models”;<br>“table stakes”; “AI is a tool, not the solution” (C3)  | Technology as baseline, not<br>differentiator        |
| “fighter jet vs passenger aircraft” (V1); “decision speed”;<br>“team composition”; “the right mix of skills” (I2) | Venture speed and structure<br>as the differentiator |
| “workflows must be owned, not improvised” (V2);<br>“organisational design and hiring” (I1)                        | Workflows owned by the<br>business                   |

Aggregate dimension 2 — The moat is human and tacit, not data (→ DP4, reframed)

| First-order codes (raw)                                                                        | Second-order theme                               |
|------------------------------------------------------------------------------------------------|--------------------------------------------------|
| “data is backward-looking”; “effectively worthless”;<br>“relationship far more important” (I2) | Data dismissed by the data-<br>richest incumbent |

| First-order codes (raw)                                                                                                               | Second-order theme                                       |
|---------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------|
| “brand and relationships are the asset” (I1); “deep trusted relationships”; “domain judgement”                                        | Relationships, reputation, domain judgement as the asset |
| “the monetisable asset is content” (V2); “data-privacy limits convertibility” (V1); “LLM predicts the next word, doesn’t reason” (C2) | Asset-centric vs data-centric mechanism                  |
| “consolidation and exec churn”; “knowledge goes stale”; “blindness” (C3)                                                              | Data advantage decays and can mislead                    |
| Aggregate dimension 3 — Structural form is contingent (→ DP1)                                                                         |                                                          |
| First-order codes (raw)                                                                                                               | Second-order theme                                       |
| “two separate teams”; “outside leader + cloud engineer + internal DNA”; “a thread connects the two” (I1)                              | Managed-separation model                                 |
| “not separate”; “blended team”; “preserve broadcast DNA”; “held by respect” (I2)                                                      | Blended-integration model                                |
| “ruined when absorbed into legacy” (V2); “must be allowed to supersede” (I2); “buy-and-build is fine, it’s what you do after”         | Acquisition-absorption failure mode                      |
| “internal-startup works for large orgs”; “small vendors need full transformation or acquisition” (A1)                                 | Scale contingency                                        |
| Aggregate dimension 4 — Architectural competence: liability or asset (→ DP2, contested)                                               |                                                          |
| First-order codes (raw)                                                                                                               | Second-order theme                                       |
| “feature-led vendors are vulnerable”; “code commoditised” (A1, I1); “moving too early” (V2)                                           | Competence as liability                                  |
| “broadcast DNA is an asset” (I2); “domain capability AI-native entrants lack”; “separable with the right structure” (I1)              | Competence as transferable asset                         |
| Aggregate dimension 5 — A time-bounded but uneven window (→ DP5, recalibrated)                                                        |                                                          |
| First-order codes (raw)                                                                                                               | Second-order theme                                       |
| “product clock slow, capability clock fast” (C2)                                                                                      | Capability clock vs product clock                        |
| “Avid vs Adobe — pressure, not lost sales” (C3); “windows narrow at the anchor-customer level” (I1)                                   | Relationship moat buffers incumbent slowness             |

| First-order codes (raw)                                                                                                     | Second-order theme                                       |
|-----------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------|
| “creator economy ~20% of revenue”; “faster adoption cycles” (C3); “slow human-adoption cycle in broadcast” (I2)             | Fast-mover advantage in adjacent segments                |
| Aggregate dimension 6 (emergent) — The integration layer is the central unmet need                                          |                                                          |
| First-order codes (raw)                                                                                                     | Second-order theme                                       |
| “no Zapier for broadcast” (C1); “rundown/automation collapse” (V2); “unifying ingest” (I1); “interoperability problem” (A1) | The gap named, in different words, across all four pools |
| “don’t make me buy another widget”; “build it into what I already use” (C3)                                                 | Demand-side preference for embedded capability           |
| Aggregate dimension 7 (emergent) — Productivity over product (sober AI calibration)                                         |                                                          |
| First-order codes (raw)                                                                                                     | Second-order theme                                       |
| “localisation, transcription, translation ship”; “litigation and union gating” (C3)                                         | Back-office ships, on-air/creative stalls                |
| “trough of disillusionment”; “investment ahead of value”; “subsidised era of cheap compute” (V2)                            | Hype calibration                                         |

## Appendix D — Illustrative Financial Model Assumptions

The financial model in Chapter 4 is **illustrative, not predictive**: its structure is grounded in the primary evidence, while the specific magnitudes are chosen to demonstrate the commercial shape rather than to forecast returns. This appendix sets out the underlying assumptions so the model can be inspected and challenged. (Values in \$’000 unless stated.)

### D.1 Pricing and subscription-tier assumptions

| Assumption                                                       | Value used<br>(illustrative) | Basis                         |
|------------------------------------------------------------------|------------------------------|-------------------------------|
| Anchor-customer subscription tier<br>(annual contract value)     | \$300k                       | Illustrative tier price point |
| Mid-market-customer subscription tier<br>(annual contract value) | \$75k                        | Illustrative tier price point |

| Assumption                                        | Value used<br>(illustrative) | Basis                                  |
|---------------------------------------------------|------------------------------|----------------------------------------|
| Net revenue retention (existing accounts)         | 1.12 (+12% p.a. expansion)   | Illustrative                           |
| Gross margin at scale (multi-tenant SaaS)         | 78%                          | Evidenced: multi-tenant SaaS necessity |
| Adoption ramp — new anchor customers per year     | 1, 2, 3, 3, 4 (Years 1–5)    | Illustrative                           |
| Adoption ramp — new mid-market customers per year | 2, 6, 12, 18, 25 (Years 1–5) | Illustrative                           |

#### D.2 Cost assumptions, including baseline token cost

| Assumption                                    | Value used (illustrative)                         | Basis                                   |
|-----------------------------------------------|---------------------------------------------------|-----------------------------------------|
| Separated venture team — annual cost (Year 1) | \$2,200k                                          | Evidenced: operating-model findings     |
| Team cost growth                              | +20% per year                                     | Illustrative                            |
| Cost of revenue                               | 22% of net revenue (78% gross margin)             | Evidenced: SaaS economics               |
| — of which, AI compute (token cost)           | ~4.8% of net revenue                              | Derived from the token assumption below |
| Baseline blended token cost                   | \$8 per 1M tokens                                 | Illustrative baseline                   |
| Token intensity                               | ~6M tokens per \$1k of net revenue                | Illustrative                            |
| Go-to-market cost                             | 18% of net revenue (reduced by incumbent channel) | Evidenced: incumbent channel advantage  |

The baseline token cost is the assumption the model is most sensitive to: compute is the dominant component of the 22% cost of revenue, and the downside scenario assumes this token cost rises materially as the current market subsidy on AI compute recedes.

#### D.3 Legacy revenue displacement (the transition trough)

The legacy-revenue displacement is modelled as **absolute dollar amounts** rather than a proxy percentage — the value of legacy hardware and perpetual-licence revenue lost as

customers migrate to subscription, per the broadcaster's account that hardware-origin vendors must accept short-term revenue loss before subscription revenue scales.

| Assumption (\$'000)      | Year 1 | Year 2 | Year 3 | Year 4 | Year 5 |
|--------------------------|--------|--------|--------|--------|--------|
| Displaced legacy revenue | 0      | -250   | -400   | -300   | -150   |

#### D.4 What is evidenced versus illustrative

1. **Evidenced** (from the primary research): the existence and shape of the transition trough; the cost composition (a largely fixed venture team alongside usage-scaling compute); the multi-tenant SaaS necessity and ~78% gross margin; the go-to-market advantage from the incumbent's channel; and the downside compute-cost risk.
2. **Illustrative** (chosen to demonstrate commercial shape, not forecast): all specific magnitudes — the subscription-tier prices, the adoption ramp, the team cost, the baseline token cost, and the absolute displacement values.

*The resulting five-year revenue, cost, and operating-result figures, and the J-curve they produce, are presented and discussed in Chapter 4.*