

Results Over Rhetoric: A Hands-On Ecosystem Guide for the AI Era

Pg Table of Contents

4 Executive Summary

7 Abstract

8 Part 1: Understanding Business Ecosystems

9 Chapter 1.1 – Why Ecosystems? The Strategic Shift in Value Creation

11 Chapter 1.2 – What Exactly is a Business Ecosystem, and How is it Different from Clusters?

13 Chapter 1.3 – Adding Value Through Multi Business and Multi Actor Ecosystems

16 Chapter 1.4 – From Technology to Business Design: Platforms vs. Ecosystems

18 Chapter 1.5 – Why Use Them? From Narrow to Broad; Physical to Digital; from Local to Digital

21 Chapter 1.6 – Not a Silver Bullet: Ecosystem's Limits and Failures

24 Cases Chapter 1: Taking traditional companies into the digital age through ecosystems: Opportunities for incumbents and their complementors

24 C1.1 – The Swedish Sydved Forestry Ecosystem

25 C1.2 – InSoil – Spanning Funding to Agri-Regeneration

26 C1.3 – Busroot – Manufacturing Analytics Platform

27 C1.4 – A Phygital Ecosystem: MAF and the Creation of the SHARE Ecosystem

28 C1.5 – From Fridges and Stoves to the Internet of Food

29 C1.6 – Linking Public Transport with Private Sector and Incentives: Velocia

30 C1.7 – Supporting Green Transition Through an Ecosystem Play: Qurator

31 C1.8 – Opportunities and Challenges in the Shift from Cars to an Electromobility Ecosystem

33 Part 2: Navigating, Designing, and Participating in Ecosystems

33 Chapter 2.1 – Strategic Choices: Join, Shape, or Orchestrate?

34 Chapter 2.2 – Roles in Ecosystems

37 Chapter 2.3 – Tools for Participation

39 Chapter 2.4 – Becoming a Plausible Orchestrator

40 Chapter 2.5 – Becoming a Good Partner or Complementor

44 Chapter 2.6 – SME Design and Readiness

45 Chapter 2.7 – Failure Patterns and Avoidance

47 Chapter 2.8 – “No Regrets” Steps

48 Cases Chapter 2: Ecosystems as drivers of focused value propositions

48 C2.1 – Growth Through Ecosystem Use in MasMovil and COSMOTE

49 C2.2 – KBC: intent-driven orchestration with Kate

50 C2.3 – DBS's POSB Ecosystem with “School Buddy”

51 C2.4. Ecosystem Support/Advisory: Synocus, FuzeQube, Evolution Ltd

52 Synocus

52 FuzeQube

53 Evolution Ltd

54 C2.5 – HDFC ERGO: From Insurance to Mindshare

55	<u>Case C2.6 – Kinaxis: turning supply-chain planning into an ecosystem</u>
56	<u>C2.7. Exnaton: Enabling SMEs to Orchestrate Energy Communities</u>
57	<u>C2.8. Katana MRP: Plug-in Ecosystem for Manufacturing SME</u>
57	<u>C2.9 – AI as an Ecosystem Accelerator: New Ways to Join, Shape, and Orchestrate</u>
59	<u>Part 3: From Goliath to David: Living Under the Shadow of Big Tech</u>
59	<u>Chapter 3.1 – Why Would You Want to Mess with the Big Tech Ecosystems?</u>
59	<u>Chapter 3.2 – When and Why Power Concentrates</u>
61	<u>Chapter 3.3 – Key Risks for Complementors and how EU regulation is trying to help</u>
63	<u>Chapter 3.4 – AI and Ecosystems: where power shifts, and what to do about it</u>
64	<u>Chapter 3.5 – Survival Strategies and Defensive Moves</u>
66	<u>Chapter 3.6 – The Role of Policy – Hope and Limitations</u>
67	<u>Chapter 3.7 – What does this Mean for SMEs?</u>
69	<u>Cases Chapter 3: Competing alongside big tech / building new services</u>
69	<u>C3.1. The Dilemmas of Smaller Organizations Working with BigTech: Life Under Amazon's Shadow for Airplay</u>
69	<u>C3.2. Shopify as an Engine and an Opportunity</u>
71	<u>C3.3. Nvidia and the growth of AI through ecosystems</u>
72	<u>C3.4. GAIA-X & Eurostack: Europe's Pushback on Platform Dominance</u>
74	<u>C3.5. Quibim: Radiology AI Ecosystem Partner</u>
74	<u>C3.6. TruBlo: modular content-trust rails instead of a mega-platform</u>
76	<u>C3.7. Philips HealthSuite & the Elusive Promise of Digital Health</u>
78	<u>References</u>
78	<u>Evolution Ltd Output</u>
79	<u>Further Readings</u>
80	<u>Links to Case Studies</u>

Michael G Jacobides (mjacobides@london.edu) is the Sir Donald Gordon Professor of Entrepreneurship & Innovation and Professor of Strategy at London Business School, and Evolution Ltd's Lead Advisor. He is Academic Advisor to BCG's Henderson Institute, member of the World Economic Forum's AI Governance Alliance, co-author of its White Paper on platforms and ecosystems, and consistently ranked as one of the 50 top management thinkers since 2019.

This report was written at the behest of and with funding from the Generalitat de Catalunya, drawing on Prof. Jacobides and Evolution Ltd's experience in researching and building ecosystems, and advising on ecosystem strategy.

© Michael G. Jacobides and Evolution Ltd, 2025.

Executive Summary

Why ecosystems now

- **Ecosystems ≠ clusters.** Clusters are geographic concentrations; ecosystems are **designed collaborations** organised by roles, rules and shared rails (identity, data, APIs, governance). Ecosystems can live inside clusters but often **span sectors and regions**.
- **Beyond Big Tech.** Ecosystem plays are not confined to hyperscalers: this report covers retail/venues (MAF/SHARE), banking (KBC/Kate), manufacturing SMEs (Katana), logistics/planning (Kinaxis), energy (Exnaton), mobility (Velocia), health (Philips), insurance (HDFC ERGO) and more.
- **AI shifts the front door.** Assistants and model services turn user **intent** into routable actions across partners, lowering search/coordination costs and moving power to whoever controls rails and ranking.
- **Roles–Rules–Data are the operating core.** Make participation terms legible and you reduce friction, attract better complements, and keep investment compounding.
- **Measurement = legitimacy.** Ecosystems scale when contribution is **provable** (closed-loop lift, time-to-resolution, churn avoided) and when **portability** is real.

What good looks like (design choices that travel)

- Start with a **minimum viable ecosystem** for one end-to-end job (quote → schedule → fulfil) using the fewest partners and shared rails.
- Publish **boundary resources**: consented identity, an event stream for intents, documented APIs/SDKs, and common metrics/attribution.
- Keep governance **light, bilateral, reversible** early; add councils/certification only after usage crosses visible thresholds.
- Tie economics to **verified outcomes**, not seats or clicks; publish the split so partners see how they win.
- Design for **multihoming** (your own and partners'), avoiding the “ecosystem” trap.

Implications for Catalonia (policy playbook)

- Build sector data spaces (health, industry, energy, tourism/retail) with consent, attestation and portability.
- Provide open rails: common identity and payments, shared event schemas, and public sandboxes so SMEs can ship assistant-callable tools.
- Fund SME ecosystem vouchers for integrations, verification and first pilots (not generic consulting).
- Issue lightweight governance templates (eligibility, SLAs, data rights, value split, design-for-exit) reusable across consortia.
- Launch an Interoperability & Measurement Scoreboard: time-to-first-API call, onboarding cycle time, share of transactions with closed-loop attribution.

- Use public procurement to reward interoperability and outcome-based pricing; require portability clauses.

Playbook for orchestrators (corporates & public operators)

- Pick one job you can own; keep the user surface simple and the rails strong.
- Treat interoperability as a product (versioned APIs, sandboxes, reference code, live status).
- Measure and share uplift; make ranking/eligibility changes auditable.
- Use AI as connective tissue—routing, anomaly detection, prediction—rather than lock-in theatre.
- Preserve reversibility: portable artefacts (logs, models, plans), deprecation windows, clear dispute paths.

Playbook for SMEs (complementors)

- Stay visible (maintain a direct channel) and portable (a second route to market/fulfilment).
- Be a clean plugin: expose “quote/schedule/fulfil”, document SLAs, and log outcomes you can bill on.
- Negotiate data and learning rights up front; keep adapters for a second assistant/platform.
- Price to outcomes where feasible; publish a one-pager pilot report (baseline vs uplift) to recruit partners.

Regulation as leverage (not a crutch)

- European rules (DMA/DSA/AI Act) improve contestability—transparency on ranking/defaults, fair access to rails, duties on high-impact AI.
- Use them to secure clarity in contracts (eligibility, rate limits, data rights), but let strategy set direction.

Sector illustrations (from the report's cases)

- **Retail/venues (MAF/SHARE):** outcome-measured retail media keeps brands, tenants and venues investing.
- **Banking (KBC/Kate):** assistants route intents to bank and partner services; rewards tied to solved tasks.
- **Manufacturing/logistics (Kinaxis):** a neutral planning core; complements build apps against a shared truth set.
- **Energy (Exnaton):** community energy rails let SMEs orchestrate local markets with transparent settlement.
- **SME commerce/production (Shopify, Katana):** merchant-first rails expand participation while keeping control closer to firms.
- **AI keystones (Nvidia):** layered rails that attract complements—illustrates power and dependency.

- **Europe's pushback (GAIA-X/Eurostack):** sovereignty via interoperability; success measured in cross-border portability and adoption.

Getting started with an ecosystem approach

- Next 90 days: Get a plan ready that sets your North Star, then: pick one job; run a narrow pilot with one anchor partner; track a single success metric.
- Next 12 months: draw on what you learnt; expand your remit; expand partners; publish rules; move pricing to outcomes; connect to data spaces; certify assurance partners.
- Throughout: maintain a second channel beyond the designed ecosystem; reserve engineering for adapters/tests; instrument dependencies (eligibility, latency, ranking).

Risks to avoid

- Big-bang launches without a north-star metric; bespoke integrations that decay; opaque value splits; exclusivities before traction.
- Egosystem drift: hoarding value and hiding rules cause partners to disengage and variety to collapse.


Abstract

Ecosystems—not standalone firms—are fast becoming the unit of competition. They are not the same as clusters: clusters are geographic concentrations of related actors; ecosystems are role- and rule-based collaborations organised around a user job and shared rails (identity, data, APIs, governance). Ecosystems may sit inside clusters, but they are not defined by geography and often cut across sectors and regions. AI and GenAI lower search and coordination costs, turning intent into action across many actors, while shifting control to whoever owns the rails for identity, data, ranking and rules. The report shows how well-designed ecosystems create measurable value (higher engagement, lower frictions, better economics) and why governance—eligibility, service levels, data rights, value-sharing and portability—decides who benefits. We also show that ecosystems are not just a Big Tech phenomenon: cases from retail/venues, banking, manufacturing SMEs, energy, mobility, health and public services demonstrate broad applicability. We translate these insights into a playbook apt for a European economy or region such as Catalunya: build sector data spaces and open rails, give SMEs cheap, trusted participation, and use light, reversible governance so partners keep options and keep investing.

Part 1: Understanding Business Ecosystems

Business (and digital) ecosystems have become all the rage. Prompted by the growth of Big Tech stalwarts such as Google, Amazon, Facebook and Microsoft, whose success is based on clever ecosystem plays, the business community has turned to ecosystems as the latest buzzword.

Excitement with ecosystems has gone beyond the tech world. In a review of all the corporate filings of large companies worldwide, researchers from BCG's Henderson Institute saw that the mention of the term ecosystem grew 13-fold from 2007 to 2016. In terms of research, following the publication of two influential articles in 2018, "business ecosystems" has become one of the most active areas of research in Strategic Management, with a 2024 report of the Strategic Management Society finding it was the most important new area of research.



But what is this all about? What defines business ecosystems, and why have they risen in importance? How do they relate to clusters, and how do they compare with other forms of organization? How do ecosystems relate to platforms? What are the key roles firms can play and what are the main varieties of ecosystems? This chapter provides an overview meant to help executives who are *not* in Big Tech figure out why they should care, and offers a number of case-studies that illustrate ecosystems on the smaller end – so as to bring them to life for companies in a broad set of sectors.


Section 1.1 explains why ecosystems have eclipsed supply chains; 1.2 defines their structure and clarifies how they differ from clusters; 1.3 shows how multi-business and multi-actor layers interlock; 1.4 separates the technical role of platforms from the managerial craft of ecosystem design; 1.5 details when and why firms should engage; and 1.6 warns where complexity and misaligned incentives derail even the boldest visions. Together, they arm readers with the vocabulary and frameworks needed for the pages ahead.

Chapter 1.1 – Why Ecosystems? The Strategic Shift in Value Creation

The organisation of value creation is shifting. For a long time most industries relied on linear “pipeline” models that marched resources through tightly specified stages toward a finished product. That logic is fraying. As markets fragment and product life cycles shorten, firms are pushed to deliver hyper-personalised, continuously updated solutions assembled from a widening menu of components. Linear chains struggle to keep pace because every adjustment reverberates across a fixed sequence of bilateral contracts, raising coordination costs and slowing response. Research on modularisation shows that when interfaces between modules are codified, old merits of vertical control diminish and a new organisational form, the business ecosystem, outperforms both integration and arm’s-length markets allowing many contributors to innovate in parallel while remaining loosely synchronised around a shared architecture.

Digital technology supplies the connective tissue that makes this alternative viable. Standardised APIs translate formerly bespoke hand-offs into machine-readable rules; elastic cloud services give even small firms on-demand scale; and edge analytics move intelligence to the periphery, enabling real-time feedback loops without central bottlenecks. Analysts trace these effects back to a simple mechanism: digitisation simplifies interfaces, lowers transaction costs, and fosters modularity, encouraging firms to open previously closed systems and invite third-party inventiveness (Baldwin, 2024).¹ What underpins ecosystems is that in many systemic solutions, there is a need to coordinate. While digitization offers a “glue” and opens up new realms of possibility, there is still a need for some organizing (or orchestrating) hand to put everything together. This is where ecosystems come in – as they represent webs of offerings provided by collaborating firms, working without the need of centralized ownership for a collective outcome. They thrive when there is a need for coordination – but where modularity allows for them to emerge as an organizing form.

¹ The evolution of the personal-computer architecture is often cited as proof. Once IBM’s proprietary design became an industry standard, an open market of clone makers emerged. Yet the fragmentation created new bottlenecks—performance improvements in one module were useless unless the rest of the system evolved in concert. Intel responded by establishing its Architecture Lab, using technical standards such as the PCI bus to coordinate outside innovators without owning them outright, thereby stepping into the role of ecosystem orchestrator.



Business ecosystems operate at two levels. First, the level of the product – or rather, the bundle that the customer sees, and second, the level of the organization – where old style ownership or use of supply chain gives way to interlocking organizations.

In terms of customer value proposition, the gravitational centre of this new logic is the user rather than the firm. Customers no longer consume isolated products; they expect seamless outcomes that traverse sectoral boundaries—booking a journey instead of a ticket, managing health rather than visiting a doctor. Ecosystem leaders therefore compete to orchestrate the total experience. Alibaba bundles retail, logistics, and finance around a single login; Uber extends ride-hailing into food, freight, and financial services; and Ping An links banking, property, tele-health, and insurance into one interface. In each case, the strategic advantage flows less from owning every asset than from aligning a constellation of partners around friction-free customer journeys. Ron Adner frames the resulting play as a shift “from constructing products to constructing value”; the task is to secure alignment across actors whose contributions are essential but outside direct control.

In terms of organization, ecosystems solve the classic make-or-buy dilemma: they preserve entrepreneurial autonomy for complementors yet achieve system-level alignment when multilateral complementarities make bilateral contracting unreliable. Intel’s stewardship of the PC standard, Apple’s curation of the App Store, and Visa’s rule-book for payment processors all demonstrate how a focal actor can wield “smart power” to shape direction, resolve conflicts, and evolve standards without full ownership. As we’ll show through examples in this report this isn’t only something Big Tech (and big scale) can achieve. Smaller number organizations can also benefit by rethinking their value add and learning from the cutting edge.

When ecosystems work, the pay-offs are striking. Research highlights three recurrent benefits: rapid access to specialised capabilities (“not all smart people work for you”), hyperscalable growth as new complements join through modular interfaces, and resilience through the ability to recombine modules as technology or demand shifts. Apple’s decision to open the iPhone to external developers is emblematic; the App Store’s explosion of innovation propelled iOS to global dominance within its first year and cemented Apple’s position as orchestrator without the

capital intensity of vertical integration. At a macro level, ecosystems sit midway on a spectrum of organisational forms, between the rigidity of integrated supply chains and the fluidity of open markets. They balance commitment and flexibility by letting orchestrators set high-level direction while letting complementors pursue their own experiments, creating option value the focal firm could never replicate internally.

None of this suggests that ecosystem strategy is a universal remedy—later chapters will address limits and failure modes—but it does explain why the concept has become central to competitiveness wherever modular technology meets systemic complexity. Pipelines struggle to meet the twin demands for personalisation and speed when improvement depends on chain of sequential approvals. Digital enablers render coordination costs low enough for distributed innovation to flourish, yet the systemic nature of many contemporary problems—energy transition, autonomous mobility, personalised healthcare—means ad-hoc markets are insufficient. Ecosystems square the circle: they mobilise diverse capabilities quickly, orchestrate them through digital governance and keep the user's holistic experience at the core.

Understanding this strategic shift is foundational for the analyses that follow. The remainder of Part 1 will formalise what ecosystems are (and are not), show how multi-actor and multi-business webs create new value, and clarify the boundary between platforms and ecosystems. Together, these chapters provide the conceptual toolkit managers need to navigate—and thrive in—the age of ecosystem competition.

Chapter 1.2 – What Exactly is a Business Ecosystem, and How is it Different from Clusters?

Business ecosystems are best understood as loosely coordinated constellations of independent yet interdependent actors, such as technology suppliers, complementors, service providers, data brokers, even regulators whose individual contributions plug into a shared architecture so that a unified value proposition can reach the end user. The architecture might be a set of technical standards, an API suite, a payments rail, or, more generally, a rule book that specifies how modules fit together, how data and money flow, and how disputes are resolved. Because no single firm owns the entire system, governance relies on norms, role definitions, and adaptive rules rather than bilateral contracts or vertical control. This structural view, often called the “ecosystem-as-structure” perspective, starts with the desired outcome and works backward to map the activities, actors, positions, and links required to make that outcome materialise.

That definition immediately sets ecosystems apart from traditional supply chains. In a supply chain, each link is bound by contractual hierarchy: component makers sell to assemblers, who sell to wholesalers, who sell to retailers, in a largely one-way flow. Complementarities are bilateral and governed by quantity, price, and quality terms. In an ecosystem, by contrast, complementarities are multilateral and supermodular: the value of a smartphone, for instance, rises with the availability of external apps, accessories, and services that no contract mandates yet all co-evolve around interface guidelines published by the platform owner. The resulting payoff is systemic rather than sequential. Once modular standards lower coordination costs,

firms no longer need common ownership to achieve joint performance; they can innovate in parallel while remaining loosely synchronised around the architecture.

Clusters, meanwhile, are a different, if loosely related animal. A cluster is a geographically bounded agglomeration (think the MedTech corridor around Minneapolis or Emilia-Romagna's packaging-machinery district) whose firms benefit from local knowledge spill-overs, specialised labour markets, and tacit norms built through face-to-face interaction. Regions such as Catalunya have distinguished themselves by building solid infrastructure that allowed their region to become an attractive place for firms that want to engage, say, in the creative sectors, ensuring that the region has growth momentum. The glue in these "local ecosystems" is proximity.

What makes clusters similar to business ecosystems is that there is some serious thought that goes into *orchestrating* them – here, through the work of authorities, be they European, National or, in the case of Catalunya, the local government, whose efforts have become the focus of a case-study by Harvard Business School. The mechanism is also similar in that the idea behind ecosystems is that they provide a mechanism that allows parties to add value to each other. A cluster means that the value of one firm (in say the creative space) increases in the presence of a value of another firm (say in technology), as they both might benefit from proximity, connection to the local universities, incubators etc. It's all about co-location.

Business ecosystems can be and increasingly are geographically dispersed. What binds them is architecture, not co-location. A Salesforce partner headquartered in Singapore can release a billing extension used by a German systems integrator and a Brazilian retailer within hours because all three conform to the same interface definitions in AppExchange. Spatial spill-overs, once the lifeblood of competitive advantage, are supplanted by digital interoperability. The latest scholarship on "innovation ecosystems" places clusters and supply chains on one side of a conceptual map and ecosystems on the other, noting that only the latter combine participant heterogeneity, non-hierarchical governance, and a coherent system-level output aimed at a recognisable audience.

Traditional clusters draw vitality from geographic proximity; ecosystems are bound instead by explicit digital links that make distance largely irrelevant. A machine-vision start-up in Helsinki can feed algorithms into a drone-mapping suite orchestrated from Sydney, confident that the interface contracts, data schemas, and revenue rules are identical whether the call originates in Lapland or New South Wales. Physical colocation may still confer advantages for tacit learning, but it is no longer a prerequisite for systemic collaboration. What binds the web is the intentional design of technical and relational touchpoints, not shared post codes.


The pattern shows up across sectors. Apple's App Store, Bosch's mobility-services hub, and the Salesforce partner marketplace differ in technology and industry context, yet each revolves around a focal orchestrator that curates standards and roles, a community of complementors that extend functionality, and users whose adoption feedback steers ongoing evolution. None behaves like a closed supply chain or a regional cluster. Apple does not contractually oblige developers to create navigation or payment apps; Bosch does not vertically integrate every electric-vehicle charger or ride-sharing fleet; Salesforce neither owns nor locally co-locates with its independent software vendors. Yet the combined offer is richer, more adaptive, and faster to

evolve than any linear chain could produce. Seen this way, business ecosystems may result from or be inspired by clusters, but they move beyond them; and they might be able to further sustain clusters, making them more adaptive.

Chapter 1.3 – Adding Value Through Multi-Business and Multi-Actor Ecosystems

The ecosystem lens becomes truly powerful when we recognise that two distinct but mutually reinforcing architectures—multi-business (or “experience”) ecosystems and multi-actor ecosystems—can be woven together to unlock value that no firm could generate alone. Both spring from the same digital foundations of modular interfaces and ubiquitous connectivity, yet they solve different problems and demand different design choices. Prescriptively, as the framework we’ve used in Evolution Ltd ecosystem projects, they are also different – one affects the choice of “where” to compete (and what products to bundle) and the other affects the choice of “how” to engage with other ecosystem partners.

ILLUSTRATIVE –
NOT EXHAUSTIVE



Multi-business ecosystems start with the customer journey and work backward. Their defining move is to wrap a constellation of goods and services around a focal need so that the presence of one element amplifies the usefulness of the others. Think of how an airline app adds carbon-offset options, ride-hailing, lounge access, and travel insurance: each component is optional, but the perceived convenience of having them one tap away nudges adoption across the bundle. Crucially, customer choice remains sovereign; the orchestrator wins by curating complementary options, not by forcing tie-in. That distinguishes a multi-business ecosystem from a traditional cross-sell or scope play. In the old model, firms pushed a wider catalogue because internal synergies lowered cost; the user value of A plus B was incidental. In the ecosystem model, the demand-side synergy is explicit: product A becomes more attractive because product B is easily discoverable, seamlessly integrated, and co-evolves through shared data and feedback loops.

Because the user experience is the binding agent, successful multi-business orchestrators obsess over interface coherence, unified identity, and frictionless hand-offs. Apple’s shift from devices to a services-heavy bundle illustrates the logic: iCloud, Apple Pay, Fitness+, and CarPlay live in different verticals, yet a common ID and payment rail make the whole feel like


an ambient service layer that follows the user from phone to watch to car. Loyalty is earned less through lock-in than through the cumulative convenience of staying inside the orbit.

Where multi-business ecosystems assemble experiences, multi-actor ecosystems organise production. They address the supply-side challenge that complex solutions rarely emerge from a lone firm's capabilities. Here the unit of analysis is not the bundle but the network of specialised participants—original-equipment manufacturers, software developers, integrators, data brokers—whose modules must interoperate in real time. Governance replaces ownership as the coordinating device. Instead of dictating specifications to captive suppliers, the orchestrator curates interface standards, contribution rules, and value-sharing protocols so that autonomous partners can innovate without derailing system coherence. Automotive over-the-air updates, smart-home compatibility programmes, and open banking APIs all show how a vertical ecosystem enables rapid recombination by standardising the “thin crossing points” where modules plug together.

A multi-actor ecosystem therefore differs from both vertical integration and conventional supply chains. Integration concentrates control inside one corporate boundary; supply chains coordinate through bilateral contracts. An ecosystem, by contrast, is multilateral and adaptive: members can join, exit, or reposition without renegotiating every other relationship, because the architecture insulates modules behind stable interfaces. That modular buoyancy is what lets a fitness-tracker maker replace its cloud-analytics provider overnight or a fintech swap identity-verification vendors with minimal disruption.²

Integration points between the two levels create flywheel effects. As the breadth of the experience ecosystem widens, more specialised partners see upside in joining, enriching the multi-actor layer. In turn, a deeper bench of partners allows the orchestrator to launch new user-facing propositions faster, tightening adoption loops and raising switching costs without resorting to hard lock-in. Network externalities, previously discussed in platform contexts, thus diffuse into the broader ecosystem fabric: complement variety makes the core more attractive, which raises demand and draws even more complements.

² Although conceptually distinct, the two architectures interlock. Most experience ecosystems rely on at least one multi-actor web to deliver the components of the bundle, while many vertical ecosystems exist precisely to feed into broader customer journeys. Consider an electric-mobility suite that promises “door-to-door carbon-neutral travel.” The user-facing bundle spans vehicle sharing, charging, parking, and payment. Under the hood, separate vertical ecosystems coordinate battery makers, charger OEMs, grid operators, and data aggregators. The orchestrator's art lies in aligning the economic incentives of upstream partners so that the downstream promise stays credible and evolves with customer expectations.




Source: BCG Henderson Institute analysis.

For managers, separating the layers clarifies strategic choices. A firm may decide to orchestrate the customer experience yet remain a partner, not an orchestrator, within several vertical ecosystems. Ride-hailing companies exemplify this split: they control the passenger-facing app but source mapping, payments, identity checks, and telematics from partner webs they do not lead. Conversely, a component specialist might orchestrate a tightly governed vertical ecosystem (say, a chipset reference design) while plugging its output into multiple branded experience bundles run by device makers. Decoupling the roles prevents the common pitfall of “ecosystem overreach,” where a firm tries to dominate every layer and ends up stifling the very complement innovation it needs.³

Finally, both architectures thrive on data, but the data flows differ. Experience ecosystems emphasise unified identity and context signals to personalise the front end; vertical ecosystems depend on technical telemetry and shared roadmaps to synchronise development. Successful orchestrators design data rights and APIs that respect these distinct purposes, ensuring that insights circulate where they add system value without eroding partner autonomy.

In sum, multi-business and multi-actor ecosystems are complementary blueprints for value creation. One curates holistic customer experiences; the other mobilises a diverse production base. Understanding their differences and the points where they intersect allows firms to craft roles that amplify their distinctive assets, avoid organisational overload, and tap the full combinatorial potential of digital modularity.

³ Distinguishing the two also sharpens the discussion around economies of scope. In a multi-business ecosystem, scope economies are demand-driven: value stems from complementarities perceived by the user. In a multi-actor ecosystem, scope economies are supply-driven: shared standards and reusable modules spread fixed costs across participants. Blending the two can generate compound benefits, but only if governance retains the flexibility for partners to capture a fair share. Heavy-handed orchestrators that rewrite rules unilaterally risk triggering defection spirals, as seen in several app-store controversies




Chapter 1.4 – From Technology to Business Design: Platforms vs. Ecosystems

The vocabulary of the digital economy often blurs two distinct constructs - *platforms* and *ecosystems*, but their purposes, mechanics, and managerial imperatives differ in fundamental ways. A platform is, at its core, an engineered stack of infrastructure and interface standards. Its job is technical: expose stable “ports” through which disparate actors can plug in code, data, or transactions without having to negotiate bespoke connections every time. When the iPhone launched, Apple’s iOS offered developers a common software development kit, payment rails, and security layers. That scaffold eliminated a thicket of bilateral integrations, encouraging thousands of apps to materialise at minimal marginal cost. The value of a platform therefore flows from modularity, economies of scope in innovation, and the indirect network effects that arise when distinct user groups—developers, advertisers, consumers—find one another across.

Platforms come in different forms, and their appeal can sometimes lead to confusion, as the term is used to describe a wide array of business models and technical architectures. As Cusumano, Gawer, and Yoffie (2019) point out, it is crucial to distinguish between **transaction platforms** and **innovation platforms**. Transaction platforms, such as Amazon Marketplace, Uber, or Airbnb, act primarily as intermediaries that facilitate exchanges between two or more parties, typically buyers and sellers or service providers and customers. Their value lies in reducing friction, aggregating supply and demand, and enabling network effects through scale. In contrast, **innovation platforms** like Microsoft Windows, Apple iOS, or Intel’s chip architecture, serve as technological foundations upon which complementary innovations can be developed. These platforms provide tools, standards, and access that enable third-party developers or firms to co-innovate and create value that extends the core platform’s functionality. While some firms, like Apple, span both types by enabling transactions on top of innovation infrastructures (i.e., app stores on mobile OS), the strategic challenges and governance mechanisms differ markedly between the two. Understanding this distinction helps

clarify not just what a platform is, but how it creates value, orchestrates participation, and scales over time



Source: Cusumano, Gawer, Yoffie (2019), *The Business of Platforms: Strategy in the Age of Digital Competition, Innovation, & Power*

A platform (Parker, Van Alstyne & Choudary, 2016) is typically a core technology or business infrastructure that connects different users, such as buyers and sellers in a marketplace, or developers and users in a software environment. Platforms work by enabling interactions and exchanges, often through a central interface controlled by a focal firm. An ecosystem, by contrast, is a broader structure. It involves a group of independent firms that work together to deliver a shared value proposition. What holds an ecosystem together is not necessarily a piece of technology, but a shared goal and a clear logic for how the parts fit together and complement each other.

An ecosystem, then, begins where the platform leaves off. It is a pattern of interdependence among autonomous but complementary organisations that jointly craft, deliver, and capture value for (and from) end users. The architecture may rest on a digital platform—indeed, most modern ecosystems do—but its essence is relational, not technical. Participants invest in co-specialised assets, coordinate road-maps, share knowledge, and align incentives so that the total solution is worth more than the sum of its modular parts. Apple's App Store is just one slice of a wider constellation that now encompasses media, health, mobility, and finance; those additions required not simply new APIs but negotiated governance about data rights, revenue splits, and curation rules. In other words, a platform can exist without an ecosystem (think of a data-sharing API that nobody builds on), but an ecosystem almost always leverages one or more platforms to do its work.

This distinction matters because the problems each structure solves are different. Platforms primarily remedy matching and innovation failures: they lower search and transaction costs, create predictable interfaces, and stimulate third-party experimentation by making it cheap to build complements that click into the core. Ecosystems address cooperation failures: situations where the benefit from a finished solution hinges on coordinated investments across several actors who lack a natural hierarchy. Electric-vehicle charging, for instance, requires hardware OEMs, grid operators, software integrators, and payment providers to move roughly in synchrony; no amount of well-documented APIs alone will persuade them to sink capital unless there is a governance layer that clarifies who does what, how surplus is divided, and how standards evolve.

Design, therefore, goes well beyond choosing open or closed APIs. Successful ecosystem architects choreograph a balance between autonomy and alignment. They must leave interfaces incomplete enough to spark external creativity, yet sufficiently stable to prevent fragmentation; they must let complementors capture enough value to keep contributing, yet retain levers such as branding, data access, certification to steer the whole. Recognising the boundary between platform and ecosystem also sharpens organisational choices. A firm may decide to open parts of its stack as a generic platform while tightly curating the ecosystem layers that create distinctive user experiences. Conversely, a specialist might ignore platform ownership entirely, opting instead to become a keystone complement whose presence is essential across many ecosystems (Iansiti & Levien, 2004).

Treating ecosystem design as a business problem, not just a technical one, helps leaders anticipate the darker side of interdependence. Platforms that are too open invite free-riding and quality erosion; governance that skews too heavily toward the hub breeds distributional grievances that deter vital complementors. Ecosystem health therefore hinges on continuous monitoring of both functional performance (is the joint solution getting to market?) and distributive fairness (is surplus shared in ways perceived as legitimate?). When those metrics drift, subtle recalibration of access rights, incentive schemes, or standard-setting forums is often more effective than unilateral diktat or wholesale redesign.

Chapter 1.5 – Why Use Them? From Narrow to Broad; Physical to Digital; from Local to Digital

Ecosystems are no longer a curiosity reserved for digital giants; they have become a pragmatic tool for any firm that wants to transform a focused product into a richer, stickier solution, expand its reach beyond traditional boundaries, and deepen the economic returns from its installed base. The appeal rests on three interlocking advantages: scope expansion, capability access, geographic transcendence and leveraging data.

First, ecosystems convert specialised offerings into comprehensive outcomes. A single product, even a sophisticated one, addresses only a fraction of the user's broader objective. By mobilising a constellation of complementors, the focal firm turns isolated value into an integrated journey. A utility that once sold smart meters alone can surround that device with demand-response software, home-energy dashboards, and solar-financing partners; the cumulative proposition lowers acquisition friction and raises switching costs without the orchestrator having to build every element itself. Because each complementor invests in its own module, the ecosystem scales its feature set faster than any vertically integrated rival and can flex to user feedback in near-real time.

Consider John Deere's evolution from selling tractors to orchestrating a full-stack agricultural ecosystem. While its core product remains machinery, Deere now integrates GPS-guided equipment, real-time field analytics, weather and soil data, remote diagnostics, and even financial services through its Operations Center platform. By partnering with agronomic advisors, software firms, and service providers, Deere offers farmers a bundled solution that improves yields, lowers input costs, and reduces downtime, without building every component itself. The result is a more valuable, sticky relationship: the tractor becomes a gateway to a broader system of tools and services tailored to each farmer's needs, scaling faster and adapting more flexibly than any standalone product or vertically integrated rival could manage.

Second, ecosystems unlock capabilities that would be prohibitively slow or costly to develop in-house. Industrial incumbents have long relied on bilateral contracts for specialised inputs, but bilateral deals cannot cope when multiple interdependent innovations must land simultaneously. Modular technical standards and clear governance lower transaction costs across the entire web, letting partners pursue parallel innovation while maintaining systemic coherence. Firms gain access to cutting-edge analytics, niche hardware, or local distribution channels without diluting focus on their core differentiation. The result is an accelerated innovation tempo and a higher return on invested capital, because external partners shoulder their own R&D and market risk.

Third, digital ecosystems dissolve geographic constraints. Physical businesses have traditionally expanded by planting new assets in each location; digital orchestration rewrites that logic. Cloud infrastructure, APIs, and federated data models allow a manufacturer in one region to interoperate with service providers half a world away, meeting customer needs that combine physical delivery with remote intelligence. A farm-equipment maker, for instance, can integrate satellite imagery from Europe, agronomic algorithms from North America, and local cooperative financing in Africa, none of which requires the orchestrator to own assets on those continents. This borderless reach is particularly valuable to small and mid-sized firms looking to punch above their weight.

Finally, another powerful motivation to embrace ecosystems is the ability to both gather more granular data from customers and turn that data into value through personalized, adaptive offerings. A firm operating alone typically sees only a sliver of customer behavior—limited to direct usage of its own product or service. But when it orchestrates an ecosystem, it gains privileged access to a broader canvas of data from adjacent activities, touchpoints, and connected devices. This richer dataset enables a deeper understanding of customer preferences, context, and evolving needs, especially when partners agree to share metadata or usage patterns through secure, consent-based protocols.

For instance, a mobility provider that operates a vehicle platform might only know where and when its cars are driven. But if it coordinates an ecosystem that includes insurance firms, charging station operators, in-vehicle entertainment apps, and fleet management tools, it can assemble a multidimensional view of driver behavior and vehicle health. This allows the orchestrator to offer tailored insurance premiums, dynamic maintenance scheduling, route optimization, or even personalized infotainment bundles, creating stickier value propositions while expanding monetization pathways. Importantly, this data loop not only enhances customer satisfaction but also improves the ecosystem itself: usage insights can guide future product design, prioritize feature development, and fine-tune the partner mix. In essence, ecosystems enable continuous learning at scale, where data fuels both differentiation and adaptability.

Ecosystems are becoming even more important as a result of the rapid development of AI and GenAI that are transforming our economy and the productive tissue, as we illustrate in case C1.7. They dramatically increase the ability to collaborate across boundaries; they offer new possibilities for intent-driven orchestration, where AI assistants translate user needs into routable requests across partners. This lowers participation costs, creates new roles for SMEs (as specialised complements, assurance layers or optimisers), and lets orchestrators compound value through shared data and feedback loops. We return to AI's implications for power and participation in §3.4 and point to concrete mechanisms throughout the cases in Parts 1 and 2.

Yet ecosystems are not a universal remedy; they outperform other governance models only under specific conditions. Modularity must be high enough that components can be separated and recombined without excessive engineering friction, while the overall solution must still require non-trivial coordination across actors. When products are tightly coupled and safety or regulatory concerns demand end-to-end control, vertical integration or a hierarchical supply chain remains superior. Conversely, if components are completely commoditised and users can assemble them without expert guidance, an open market may suffice. Ecosystems thrive in the middle ground: plenty of interchangeable modules, but a clear benefit from orchestrated coherence.

Two additional criteria sharpen the decision. First, is the underlying problem large or acute enough to justify the complexity premium that an ecosystem entails? Companies succeed when they remove a sizable customer friction (e.g. excess cost, delay, or uncertainty) or satisfy a new, unaddressed need. Second, can the focal firm credibly motivate partners to join? Without a compelling value blueprint, clear role definitions, and fair value-sharing, co-innovation and adoption risks will derail the effort before it reaches scale (see Adner, 2012, 2021).

When those hurdles are cleared, even asset-heavy industrial players can capture new layers of margin. By embedding their equipment inside a broader digital scaffold, they migrate from one-off sales to recurring revenue streams, harvest downstream data to refine their designs, and insert themselves deeper into customer workflows. Crucially, they do so without compromising their hard-won expertise in engineering and quality. Instead of competing head-to-head with tech entrants, they reframe the contest around the full system outcome (e.g. efficiency, reliability, sustainability) where their domain knowledge remains indispensable.

For smaller firms and complementors, ecosystems open secondary paths to growth. Joining a strong ecosystem grants instant visibility and trust, reduces customer acquisition costs, and provides a ready-made integration fabric. Contributors can multihome across several ecosystems to hedge dependency risk, specialise in a bottleneck niche to secure pricing power, or leverage orchestrator tools such as payment rails, identity services, analytics, to accelerate their own innovation. The contributor role may, in fact, deliver superior risk-adjusted returns: capital outlay is lower, strategic flexibility higher, and bargaining power can be substantial if the module is mission-critical.

Digital ecosystems also change the mindset required for success. Opportunities are pursued in terms of participation rather than ownership, and strategy becomes an exercise in aligning incentives across a network rather than optimising a single value chain. Governance, therefore, is integral to design. Rules for access, participation, and commitment must balance openness, which fuels variety and rapid uptake, with enough control to maintain quality and capture value. Overly restrictive schemes stifle complementor creativity; overly permissive ones invite free-riding and undermine trust.

Finally, the move from local to digital is not simply a shift in distribution channels but a redefinition of competitive terrain. Local incumbents can project their strengths globally by exposing APIs or data sets, while global entrants can localise offerings through partnerships rather than brick-and-mortar investment. Success in this arena demands a broader strategic aperture: scouting adjacent capabilities, monitoring regulatory shifts that affect data sharing, and continuously fine-tuning role expectations as technology and customer preferences evolve.

In short, firms use ecosystems to transcend the limits of product scope, organisational capability, and geography. When modularity is high and coordination matters, an ecosystem creates a bigger pie and helps each participant earn a larger, more defensible slice. The prize is a resilient, adaptive business model that can pivot with market turbulence, capture value across the life cycle of the solution, and engage customers in richer, more enduring relationships.

Chapter 1.6 – Not a Silver Bullet: Ecosystem's Limits and Failures

Ecosystems have moved from managerial novelty to strategic orthodoxy with remarkable speed, yet they come with costs, as they are complex to operate and require deft governance – let alone that some key organizations may decide that rather than building an ecosystem that might benefit the community it would want to keep it within its own boundaries. You should expect an ecosystem to be worth it when the complexity of customer needs exceeds what a single firm can deliver efficiently on its own—and when the rewards of coordination outweigh the frictions of collaboration. Ecosystems become especially attractive in fast-evolving contexts where customer value emerges from the interplay of multiple technologies, services, or regulatory environments, and where speed, optionality, and modularity are critical to success. They shine when no single player can credibly build or buy every component, yet a loosely aligned network can flexibly co-create and adapt at scale. In addition to that, ecosystems come with pathologies of their own.

First, the logic that makes them so attractive in some contexts can be self-defeating in others. The very features that power an ecosystem—modularity, multilateral interdependence, and optional participation—also create friction, overhead, and political risk. Whenever the number of moving parts rises faster than the benefits of coordination, complexity can dwarf reward, turning an elegant architectural vision into an unworkable tangle. That is why even firms with deep pockets and impeccable technical credentials sometimes stumble: it takes more than APIs and enthusiasm to keep dozens of semi-autonomous contributors moving in the same direction.

Complexity on its own is not fatal, but it magnifies every other weakness. A cautionary tale is GE's Predix initiative, which promised to be the universal operating system for industrial data yet faltered because the ecosystem it required (OEMs, analytics providers, service contractors, and plant operators) found neither a clear value proposition nor a credible path to monetisation. Partners hesitated to rewrite their road maps around a platform whose economics felt opaque; without their commitment, the critical mass of data and use-cases never materialised, reinforcing scepticism in a vicious spiral. The lesson is stark: orchestration only works when the economic upside for each actor is obvious and immediate enough to justify the switching costs.

Partner adoption hurdles often mask a deeper governance problem. If incentives are misaligned, no amount of brand charisma can keep an ecosystem afloat. Complementors decide on the margin whether to invest scarce resources; they compare what they will capture

against what they must contribute. Revenue-share formulas that appear generous in headquarters spreadsheets can look anaemic once development, integration, and ongoing support are factored in. Worse, fickle rule-making erodes trust. When orchestrators unilaterally re-scope APIs, alter ranking algorithms, or squeeze margins to hit quarterly targets, they signal that today's upside may be tomorrow's sunk cost. Predictable rules, transparent metrics, and credible dispute-resolution mechanisms are therefore as essential as any technical spec.

An equally common trap is the “ego-system”—a platform conceived primarily to showcase internal capabilities rather than to solve an external pain point. Vodafone's 360 service exemplified the risk: launched with fanfare, it bundled messaging, social feeds, and apps but never clinched meaningful third-party support. Developers saw little incremental demand, users found better alternatives elsewhere, and the initiative collapsed despite the parent's scale and marketing muscle. The moral is simple: ecosystems thrive on outside-in value. Starting with what customers and partners genuinely struggle with, then architecting participation terms around those struggles, beats any attempt to retrofit collaborators into a home-grown solution.

Governance missteps also manifest in the way value is measured and shared. Free-riding lurks whenever contributions are intangible, and overappropriation by the hub can chill innovation before it begins. Robust, system-level metrics such as churn reduction, basket expansion, service attach rate, create a shared scoreboard and reduce suspicion. Ping An's digital ventures allocate internal capital only after managers prove that ecosystem KPIs link to both partner success and customer outcomes; the discipline keeps everyone honest and focuses debate on the pie rather than the slice. By contrast, projects that trumpet vague ecosystem synergies without a quantifiable yard-stick almost always disappoint.

Openness is another double-edged sword. A broad interface encourages variety, speeds experimentation, and builds network effects, but too much permeability invites fragmentation, quality erosion, and litigation over intellectual property. The art lies in calibrated permeability, ie being open enough for distinctiveness to flourish, closed enough that system standards remain coherent. Apple's review process, Salesforce's AppExchange certification, and Visa's compliance audits all illustrate how rule-based gating can maintain trust while sustaining a high rate of external innovation.

Clarity of role is equally non-negotiable. Not every company can or should aspire to orchestrate. Orchestration demands a blend of assets that few possess simultaneously: user access, data gravity, brand legitimacy, technical infrastructure, and the balance-sheet stamina to subsidise early stages. Overreaching firms often dilute focus, alienate partners, and incur costs that never amortise. By contrast, a keystone or specialist complementor can earn robust returns by solving a unique bottleneck (e.g, security, localisation, analytics) across multiple ecosystems, without bearing the governance burden. Choosing the right seat at the table is therefore a strategic decision, not a matter of ego.

Before committing resources, managers should run a simple diagnostic. Is the customer problem large and unresolved enough to justify a coordinated answer? Can modular standards shrink transaction costs below the value created? Do we have assets or influence that make us a natural orchestrator, or are we better placed to complement someone else's vision? Have we defined transparent incentives and measurable outcomes that prospective partners will find credible? And finally, do we possess the organisational agility to evolve the rules as technology, regulation, and competitive dynamics shift? If any answer trends toward no, a more traditional alliance or vertical integration may deliver faster, surer returns.

Ecosystems remain a powerful mechanism for unlocking cross-boundary innovation, but they are not a strategic panacea. They magnify opportunity and risk in equal measure, rewarding careful design and punishing hubris. When they fail, it is rarely for lack of technology; rather, they collapse under the weight of misaligned incentives, governance gaps, and misplaced assumptions about partner commitment. Treating ecosystem building as a disciplined choice—anchored in user value, grounded in economics, and supported by explicit rules—turns a fashionable buzzword into a durable source of advantage.

Cases Chapter 1: Taking traditional companies into the digital age through ecosystems: Opportunities for incumbents and their complementors

C1.1 – The Swedish Sydved Forestry Ecosystem

Sydved's business begins in the old growth forests that blanket southern Sweden, but its competitive edge now lies in the invisible mesh of data and relationships it orchestrates above the tree line. For decades the company prospered by dispatching field foresters to thousands of small, family-owned woodlots, negotiating harvest volumes over coffee in farmhouse kitchens, and coordinating truckloads of logs through a well-oiled but primarily manual chain of buyers, hauliers, and mills. That high-touch model delivered impeccable local trust yet struggled to cope with a market that demanded faster quotes, tighter delivery windows, and verified sustainability credentials. Sydved's answer was not to discard its social capital but to graft a digital coordination layer onto the legacy network, creating a hybrid ecosystem that now sets the pace for Swedish timber logistics.

The pivot centred on a browser marketplace that matches standing timber with buyers in real time. Landowners upload plot data and preferences; buyers receive instant matches by species, grade and haulage radius; logistics partners see aggregated demand to optimise routes. Algorithms draw on satellite imagery, historical yield curves and live price feeds to propose thinning schedules and bid ranges, yet every digital offer still routes through the local forester, who validates conditions and handles face-to-face contracting. The platform mechanises the arithmetic and paperwork; humans keep the relationships.

Results came quickly. Quote lead times fell from weeks to hours; truck utilisation improved as dispatch stitched small lots into full loads; transparent bid histories gave landowners leverage they rarely enjoyed in one-to-one negotiations. Planning also changed: predictive harvest models began to factor soil moisture, projected stumpage and biodiversity buffers, turning static plans into quarterly living documents.

The model scaled. Freed from phone-tree admin, Sydved expanded to woodlots under 50 hectares, bundling micro-tracts into contiguous harvesting blocks that were previously uneconomic. As volumes rose, it negotiated better haulage tariffs and mill gate-slots, reinforcing a speed-and-cost flywheel.

Digitisation did trigger concerns. Veteran foresters feared algorithms would commoditise local knowledge, and some worried price transparency could erode cooperative norms. Management reframed the tools as augmentation—training foresters to use satellite overlays and dashboards—and shifted incentives from volume to a blended scorecard (transaction speed, landowner satisfaction, biodiversity compliance), signalling that advisory skill still anchors the model.

Governance stayed hybrid. Sydved kept control of algorithms and the data lake but opened APIs to accredited hauliers and independent silviculture advisers, with participation contingent on traceability tags, sustainable-forestry certifications and data-reciprocity rules: Enough openness to attract complements, enough discipline to avoid winner-takes-all dynamics.

The sector-level effect is practical: the velocity and transparency of each harvest window improve without bulldozing local culture. Landowners still shake hands with someone they

know; buyers still get a single point of accountability; dispatch still reacts to road bans—but the pieces now align faster and with clearer economics. Externally, digital log “passports” bolster due-diligence and climate-smart positioning; internally, every transaction feeds a data flywheel that sharpens matching and margin (see §1.5).

Sydvéd shows how incumbents can modernise a trust network by codifying roles and rules, not by centralising everything. The lesson is that as you consider how to engage with digital technology and shift from tradition-based webs to structured ecosystems, you can keep human brokers where they add confidence, standardise interfaces around them, and use shared data to coordinate complexity at scale.

C1.2 – InSoil – Spanning Funding to Agri-Regeneration

InSoil is a profit-seeking, ESMA-authorised platform created by the team behind HeavyFinance to solve a persistent gap: most small and mid-sized farms want to adopt regenerative practices, but cannot fund the transition on standard bank terms, and carbon markets have been too opaque to rely on. Its purpose is straightforward—to turn verified improvements in soil health into bankable cash flows that align the interests of farmers, investors and corporate buyers of carbon outcomes.

The keystone instrument is a “green loan”: zero-interest credit repaid from future carbon-credit royalties rather than fixed instalments. Farmers commit to a regenerative plan; InSoil’s monitoring stack quantifies additional soil carbon; verified tonnes are tokenised through a registry partner and purchased by a pre-committed buyer; part of the proceeds amortises the loan. Data, verification and cash movements are instrumented from the outset, so no party has to take performance on trust. The incentives line up: farmers access working capital without cash-flow strain, investors earn returns indexed to climate performance, and InSoil is paid for origination, servicing and a small share of carbon.

Early traction suggests product–market fit at scale. By mid-2025 the platform had channelled more than €80 million to over 4,000 farms, with roughly 800,000 hectares enrolled. Investor yields clustered around 12–14% while defaults on green loans stayed below 3%, supported by collateral and the carbon backstop. Dedicated facilities (around €100 million from a carbon buyer and €50 million from a European public investor) added balance-sheet heft so credits could be aggregated in volumes that matter for corporate Scope-3 programmes.

Operationally, the flywheel is data. Field sampling and remote sensing establish baselines; satellite indices track biomass; on-farm sensors feed moisture and nutrient signals into machine-learning models that forecast sequestration. Farmers receive practical agronomy prompts and issuance forecasts via a dashboard. A shared portal gives both farmers and investors clear sight of projected carbon issuance, amortisation profiles and peer benchmarks—transparency that has often been missing in carbon markets.

Governance and build choices are aimed at trust and scale. Operating under investment-services authorisation, InSoil applies standard investor protections such as segregated accounts and stress testing, and publishes portfolio-level carbon audits aligned with emerging European certification. The platform is API-first: new verification methodologies (for example, agroforestry or biochar) can plug in without rewiring the core, and third-party complements (e.g., weather-indexed insurance, invoice finance, cooperative group loans) attach where they add value. Internally, underwriting blends conventional credit metrics with agronomic risk

scores; staff incentives balance verified carbon outcomes with farmer satisfaction rather than sheer loan volume.

Strategically, InSoil chooses the role of keystone rather than would-be monopolist. It sets open, transparent participation rules and neutral rails instead of owning registries or dictating agronomy. That design is pragmatic: it lowers adoption risk for farmers and partners, crowds in specialist complements, and enables cross-border growth without heavy capital expenditure. Farmers keep their trusted advisers; banks can co-lend or refer; corporates gain traceable, verifiable tonnes; and a wider set of service providers can innovate on top of the same data backbone.

The case shows how an AI-enabled monitoring layer can act as shared infrastructure for an ecosystem, lowering participation costs and legitimising value-sharing. By aligning roles, rules and data, InSoil converts regenerative ambition into investable flows and demonstrates a workable route for regional scale-up of climate-positive agriculture.

C1.3 – Busroot – Manufacturing Analytics Platform

Most small and mid-sized factories still run blind: legacy machines hum without connectivity; cycle times are guessed; energy is metered monthly. Output Industries built Busroot to stitch those gaps closed with cloud-native plumbing a tiny workshop can afford. It is an ecosystem, not just a product: a lightweight gateway clips onto existing lines; low-cost sensors stream vibration, power and environmental data; a thin cloud layer turns signals into standard shop-floor metrics (OEE, scrap, energy intensity) in a browser cockpit. The architecture is vendor-agnostic: SMEs can mix sensors from different suppliers; Output Industries curates compatibility and keeps open APIs so third parties can add drivers and micro-apps without touching the core.

The modular design is the strategy- and as we mention, a foundation for ecosystem growth. SMEs digitise line-by-line rather than buying a monolithic MES; local integrators, often two- or three-person firms, earn fees wiring gateways and configuring dashboards; cloud providers supply security and elastic compute on usage. Everyone takes a slice, and no single actor dominates, avoiding the lock-in that makes big industrial IoT suites daunting. Output Industries orchestrates by setting device protocols, maintaining the analytics backbone and sharing revenue on transparent rules.

The evidence is encouraging. In the Made Smarter UK pilots, 23 of 25 factories kept Busroot post-trial; productivity rose c.10–15% as real-time alerts headed off cascading faults; energy dashboards cut power use by c.28% by exposing idle motors and inefficient changeovers; managers replaced whiteboard hunches with live bottleneck data that flowed straight into ERP. A textiles vignette makes it concrete: vibration sensors plus the cockpit exposed micro-stoppages tied to humidity spikes; adding a cheap dehumidifier and revising maintenance reduced downtime by a third within a month, financing the next wave of sensors—Busroot's pay-as-you-grow ethos in action.

Governance emphasises openness with certification. Any sensor vendor can join after security review and data-format adherence; integrators get sandboxes and training to earn credentials quickly; platform fees are modest and blended into device/subscription pricing, so the orchestrator's revenue scales with ecosystem adoption, not per-site lock-ins. Soft infrastructure

matters too: community forums for shared dashboard templates; a plugin marketplace (e.g., predictive spindle failure, recipe-based energy benchmarks); and hackathons that surface niche tools later rolled into core releases. Recognition has followed, and the roadmap adds AI anomaly detection, carbon-accounting modules and a partner API so financiers can underwrite equipment loans from live utilisation data, each step deepening interdependence across participants.

Beyond illustrating an ecosystem in action, this particular case study also aligns with our next Chapter, too. The take-away for Catalan firms is that we can treat Busroot as a template: design for incremental adoption; keep governance neutral and transparent; let local integrators capture value; and anchor claims in a single truth set of shop-floor metrics. That is how a mid-cap, not Big Tech, can orchestrate an industrial data ecosystem, and how regions can lift productivity by backing open interfaces, portability and SME-centric rails (see §2.1–2.3, §2.6).

C1.4 – A Phygital Ecosystem: MAF and the Creation of the SHARE Ecosystem

Majid Al Futtaim began as a regional champion in real estate and experiences—owning and operating malls, Carrefour hypermarkets, cinemas and leisure venues across MENA. The original business model was straightforward: curate attractive destinations, lease space to tenants, drive footfall, and collect rent. As consumer journeys became omnichannel, new opportunities for engagement arose. Tenants wanted measurable demand, not just traffic; brands wanted closed-loop proof that media moved product; and shoppers expected relevance across app, web and store. The Group's answer was to evolve its role from a pure landlord and start orchestrating a data-stitched ecosystem that aligns incentives across retailers, brands and venues.

The pivot started with SHARE, a single customer identity and consent framework spanning malls, Carrefour, cinemas, lifestyle brands and leisure. Instead of siloed loyalty programmes, shoppers now had one ID, one wallet and clear privacy controls. This identity layer made it possible to connect browsing, purchase and visit data, and to do so on permissioned terms. With that spine in place, Majid Al Futtaim invested in data engineering and AI: models to infer intent, predict basket lift and choose the right context for messages; governance to protect consent and keep sensitive data inside clean-room boundaries; and common measurement so every partner sees the same truth set.

On those rails, the Group then built a retail-media business that turns its physical and digital estate into addressable inventory. Brands can reach high-intent audiences on Carrefour's owned channels, extend that reach off-site, and reinforce it in-store and in-mall through smart screens mapped to shopper journeys, planned and reported as one campaign. Crucially, performance is tied to incremental sales, not vanity clicks. Attribution models distribute value beyond last-click to capture “halo” effects across a basket, so a screen impression that lifts category spend is recognised even when the final SKU differs. This makes media spend feel legitimate to tenants and brands alike and avoids the familiar tug-of-war over who gets credit.

The economics have moved from promising to material. As permissioned, location-aware campaigns reach shoppers roaming the malls, with creative shaped by purchase histories and predicted needs, retail media has emerged as a new profit pool: with a 60% margin, with clear runway as inventory expands across formats, categories and partner integrations. That profitability matters because it funds the flywheel: better targeting and measurement improve

tenant ROI; stronger tenant ROI attracts better brands and experiences; better experiences draw more footfall, which enriches the data and raises the value of the network for everyone.

The ecosystem now supports multiple on-ramps for complementors. Large CPGs plan full-funnel activity with closed-loop proof; SMEs buy self-serve flights targeted to look-alike cohorts they could never reach on their own; agencies and creative studios build formats timed to footfall and day-part; analytics partners publish models for optimal frequency and placement; fintech partners bundle media credits and dynamic offers into leases or supplier terms. Because IDs, consent, placements and reporting are standardised, new roles can flourish without bespoke data deals or one-off integrations.

AI is the quiet engine. Models infer intent from journeys (weekly stock-up versus discovery), allocate placements to minimise waste, predict lift by context and continuously adjust creative. In-store, computer vision and event streams can time messages to queue length or aisle congestion; online, generative tools cut creative turnaround so offers keep pace with real consumption cycles. The effect is a participation dividend: smaller tenants and local brands can credibly play alongside global advertisers because targeting, creative and measurement arrive as packaged services rather than bespoke projects.

Majid Al Futtaim's governance choices help keep the game fair as scale rises. Consent is explicit and portable; personally identifiable information stays separated from media planning; placement eligibility and ranking criteria are published; and outcomes are audited so partners are not asked to take performance on trust. Commercially, pricing is tied to demonstrable incrementality, not captive access to prime screens or app surfaces. Those rules are not altruistic; they are strategic. Partners stay and invest when they believe the scoreboard and the allocation of credit are sound.

The strategic significance extends beyond retail. This is a playbook for any regional orchestrator with a mix of physical and digital touchpoints: start by unifying identity and consent; invest in high-fidelity data and shared measurement; use AI to translate context into outcomes; and build a media and monetisation layer that pays for the rails while increasing partner returns. Done well, the landlord becomes an ecosystem coordinator whose economics depend on partner success, not just rent. The result is a network where value accrues to the orchestrator and to the complements that make the places worth visiting—turning footfall into a repeatable, high-margin profit stream and a defensible data asset that strengthens the whole.

C1.5 – From Fridges and Stoves to the Internet of Food

Haier's "Internet of Food" (IoF) reframes the kitchen as an ecosystem rather than a product aisle. After rebranding to Haier Smart Home, the firm set out to orchestrate end-to-end culinary journeys, from farms and cold-chain logistics to connected appliances, digital content and last-metre services, so the everyday act of cooking becomes a data-rich, repeatable experience rather than a string of disconnected purchases.

The operating model starts with scenarios, not SKUs. Instead of pushing ovens, Haier asks what people are trying to achieve—fast weeknight dinners, restaurant-grade dishes, healthier eating—and assembles the complementors each scenario needs. A flagship example bundled specialist farms, cold-chain distributors, QR-coded ingredient packs and a smart-oven programme into a one-tap, chef-guided meal that finishes at home. Success is measured in

repeat purchases, cross-device activation and the flow of behavioural data back into the Smart Home app.

Data sits at the centre of the flywheel. Appliances log usage; QR scans and voice commands enrich profiles; marketplace orders expand the catalogue of complements. Personalisation follows: recipes that match fridge contents, replenishment prompts timed to consumption and service nudges that drive after-sales revenue. The same data underwrites governance: partners who integrate APIs and share performance data gain visibility and analytics access; those who do not cede ground to more cooperative rivals—an open posture steered by clear consent flows. (See §1.5 on data flywheels.)

Ecosystem position varies by market. Where Haier can marshal dense complements and capture data at scale, it leans into a vertically integrated experience; where regulation, culture or rival platforms make dominance costly, it adopts a keystone role: Exposing robust APIs, embedding in others' stacks and using co-creation venues to probe niches with minimal capital. For smaller firms, the landscape is paradoxically welcoming: IoF lowers entry barriers by providing identity, device protocols, payment and logistics rails out of the box. The trade-off is dependency, mitigated by multihoming across major appliance ecosystems or by specialising in bottlenecks (for example, allergen tracking or niche dietary content). (See §2.1–2.2 on roles and governance.)

Value capture follows role clarity. Orchestrators monetise data and cross-sell, premium hardware makers differentiate through services, complementors choose between volume and niche authority, and logistics/payments earn transactional rents. Misaligned splits, exclusivity that throttles variety or opaque data policies can stall the web—hence the emphasis on transparent rules and portable interfaces. (See §2.3 on tools for participation.)

Three uncertainties shape the road ahead: whether standards like Matter unlock truly device-agnostic kitchens; how privacy/competition rules shift data portability; and whether consumers sustain subscription-based cooking and automated replenishment without fatigue. Whatever the outcome, the direction is clear: white-goods makers are competing to orchestrate daily rituals, harvest behavioural insight and partner with an expanding roster of digital and physical services. IoF shows how a legacy manufacturer can pivot from product innovation to scenario orchestration, with AI increasingly acting as the connective tissue that lowers participation costs, improves matching and keeps the experience learning over time. (See §3.4 on AI as an ecosystem accelerator.)

C1.6 – Linking Public Transport with Private Sector and Incentives: Velocia

Velocia is an incentive-driven mobility platform that uses rewards and a “mobility wallet” to shift everyday travel toward public transit, bikes and shared modes. Rather than add another trip planner, it layers behavioural incentives onto existing rails so riders earn as they hit simple goals, agencies see measurable uptake, and operators learn which nudges work where, all without heavy capex. Offerings include a pay-as-you-go rider incentive program, a wallet that lets users plan, pay and track across services, and an off-the-shelf “major events” module to move crowds smoothly to and from venues.

The ecosystem spans transit agencies, fare/payment providers, MaaS apps and card networks. Velocia's role is to make these actors legible to one another: it plugs into agency systems and

open-loop payments, integrates with ticketing and trip apps, and routes rewards to riders when verified behaviours occur. Partnerships with transit authorities (e.g., regional agencies in the U.S.), fare platforms such as Masabi, mobility apps like Transit, and networks such as Visa illustrate the connective tissue required to turn fragmented options into a coherent experience. The result is a brokered model in which each stakeholder sees how it wins—agencies in ridership and peak smoothing, payment and fare providers in throughput, riders in convenience and savings.

Incentives are the governance tool. By paying for specific, auditable actions—taking transit to a game, shifting a commute to off-peak, choosing active modes, Velocia aligns interests while generating data that improve targeting and service design. Embedded programs (for example, within regional fare apps) reduce friction, and the feedback loop from reward to behaviour to data to better targeting lowers acquisition costs and makes value sharing credible. For cities and operators, this turns “loyalty” from marketing into an operational lever: more predictable loads, fewer car trips for major events, and clearer signals on where to add capacity. For complementors, clear APIs and attribution mean rewards can be funded against verified outcomes rather than broad subsidies.

Strategically, the case shows how a small orchestrator can stitch a multi-actor system by standardising rewards, data and attribution. The platform does not try to own every surface; it makes participation cheap, proves lift, and lets agencies and partners keep their brands and economics, illustrating a pragmatic ecosystem pattern that other public-service domains can emulate.

C1.7 – Supporting Green Transition Through an Ecosystem Play: Qurator

Corporates racing toward net-zero targets discover that pledges are easy; selecting credible suppliers is not. Thousands of unfamiliar vendors offer carbon-cutting hardware, software and services, but the landscape is opaque and slow to navigate. Qurator, a Stockholm-based startup, tackles this by building a data-rich, two-sided platform that aims to be the neutral catalogue and due-diligence rail for B2B decarbonisation.

Rather than add another carbon-accounting tool, Qurator behaves like a roaming scout. A machine-learning engine continuously scrapes public and proprietary sources (professional networks, patent filings, VC portfolios and trade-fair catalogues) and classifies hundreds of thousands of companies across a deep sustainability-tech taxonomy. On top of that corpus sit two products: an open marketplace that makes suppliers discoverable, and a subscription “sustainability intelligence” portal that lets buyers express constraints (feedstock, CO₂e thresholds, regional rules) and receive shortlists in days instead of months. Vendors get profile pages and intent signals; buyers get triaged options and comparable evidence; over time, API feeds plug into procurement suites and carbon ERPs. Monetisation is balanced across both sides (credibility packages, scouting fees, data services) to keep incentives aligned and search costs falling.

The idea was de novo: sustainability managers were Googling, cold-calling and commissioning €100k landscape studies—yet still missing most relevant providers. The founding team, drawn from corporate innovation and data-science roles, shaped Qurator to collapse that waste and level the field for SMEs with strong tech but limited sales reach. Early capital came from

climate-focused VCs, InnoEnergy and EU innovation grants, providing both resources and policy goodwill.

Pilots suggest the approach works: a €20bn food group used the platform to screen 25 packaging alternatives and shortlist five against strict water-use and recyclability thresholds; a mid-size European utility found digital-twin software for grid optimisation; suppliers report higher-quality inbound leads at a fraction of trade-fair spend. Headwinds remain. Shifting subsidies and disclosure rules can pause corporate experimentation; “greenhushing” reduces the public signals that feed network effects; and general-purpose AI threatens to commoditise surface-level discovery. Qurator’s response is to double down on curated taxonomy depth, verification signals and neutrality such as vetting profiles, flagging hype and keeping buyer search behaviour separate from vendor analytics, so that ranking remains trusted and not simply pay-to-play.

The strategic significance is broader than a single startup. Qurator shows how a lean actor can orchestrate transaction-enabling infrastructure in a domain usually ceded to consultants or big platforms: stitch fragmented data into a shared map; standardise interfaces so newcomers can join; and price participation so value is created on both sides. It also underscores the limits: market opacity is both opportunity and risk; external policy swings can choke the very demand that powers data flywheels; and credibility depends on visibly neutral rules. The managerial takeaway is to treat discovery and due diligence as ecosystem services—open enough to attract complements, disciplined enough to sustain trust—so that decarbonisation spend moves faster, with better matches and fewer dead ends.

C1.8 – Opportunities and Challenges in the Shift from Cars to an Electromobility Ecosystem

Europe’s electrification push has delivered better batteries and denser fast-charging, yet the journey still feels stitched together. Drivers juggle apps, plugs and opaque tariffs because the problem is systemic: electrified mobility only works when energy, grid operations, charging hardware, vehicle software, data services and retail all interlock around user “jobs” from plug-in to bill-settlement and, ultimately, grid balancing (see §1.5).

The centre of gravity is contested. Utilities own the grid but optimise around monthly billing rather than real-time mobility. OEMs control the dashboard but lack energy-market DNA. Tech firms can mask fragmentation with routing, payment and optimisation, yet risk centralising data and discovery. Equipment makers ship smart boxes without user-level visibility. Oil & gas players control prime forecourt locations and fleets but must pivot to electrons without cannibalising cash flow. Each has assets a finished solution needs; none can deliver it alone.

Why hasn’t this “just happened”? Because electromobility mixes high co-innovation risk (every module must work in sync) with high adoption risk (users punish slow or unreliable journeys). Standards cured the worst of plug incompatibility, but the soft rails, ie identity, roaming tariffs, cybersecurity, fair value-sharing, are still patchwork, so capex rises while NPS stagnates. The fix is less hero tech than role clarity plus fair governance (Ch. 2): utilities expose time-of-use and flexibility as APIs that assistants and portals can call; OEMs standardise V2H/V2G interfaces and keep e-wallets open; tech hubs orchestrate intent (“route me to the cheapest green kWh now”) while staying neutral and transparent; device makers design for multihoming; forecourts become multimodal hubs combining fast-charge, retail and fleet services.

Policy can synchronise moves but cannot substitute for orchestration. Use regulation as **rails**: card readers and roaming transparency for public chargers; harmonised data-sharing across power and mobility; Data-Act-style portability in contracts and tenders (Ch. 3). Then make progress legible with a common scoreboard of user-centred KPIs: median plug-to-pay time, roaming success rate, renewable share per session, household bill impact from smart charging, fleet availability for peak-shaving. Visible metrics deter freeriding and crowd capital to what works.

For SMEs, the seam is software: roaming clearing, micro-forecasting, depot-energy optimisers, embedded insurance. The playbook is **plugin strategy**: API-first, multihome across hubs, price on verified savings, and insist on data-portability clauses (see §2.3, §2.6). For regions like Catalonia, the lesson echoes the report's core argument: ecosystems do not self-assemble; they must be continuously governed around a crisp user promise, fair value-sharing and interoperable rails. Hardware costs will keep falling; winners make the experience feel like one system end-to-end.

Part 2: Navigating, Designing, and Participating in Ecosystems

Part 2 translates the theory into action. Chapter 2.1 walks executives through the pivotal decision of whether to join, shape, or orchestrate an ecosystem. Chapter 2.2 disentangles the core roles (hub, keystone, complementor) and shows how value flows between them. Chapter 2.3 introduces practical mapping and risk-assessment tools that turn abstract webs into concrete blueprints for action. Chapter 2.4 sets out the tests a company must pass to be seen as a credible orchestrator, while Chapter 2.5 details the playbook for thriving as a partner or complementor. Chapter 2.6 focuses on SME readiness, from API design to governance discipline. Chapter 2.7 flags the most common failure traps and how to dodge them, and Chapter 2.8 closes with four “no-regrets” moves that strengthen ecosystem fitness regardless of the role ultimately chosen. We adopt a simple build–join roadmap (Figure 2.1) and a role-choice filter (Figure 2.2) to keep the analysis practical.

Chapter 2.1 – Strategic Choices: Join, Shape, or Orchestrate?

Firms exploring business ecosystems face an immediate fork in the road: should they participate as a complementor, try to shape emerging standards, or take on the demanding role of orchestrator? Making the right call hinges on a clear grasp of why the company wants to engage at all, what role best exploits its assets, and how it can integrate with partners at scale—the Why–What–How logic that underpins every successful ecosystem move.

For the vast majority of companies, entry begins as a complementor. Complementors plug focused pain points with modular offers: an AI-driven pricing engine for a marketplace, a payment widget for a mobility app, a diagnostics API for a smart-home hub. Because complements snap into pre-defined interfaces, development cycles are short, capital requirements modest, and the return horizon comparatively quick. The complementor path therefore suits firms that possess distinctive technology, data, or customer insight but lack the breadth, brand legitimacy, or balance-sheet strength to rally a constellation of partners on their own. The strategic question for these players is not whether to participate—opportunity cost virtually guarantees that—but which ecosystems to prioritise, how to differentiate amid rival modules, and how to secure access to the orchestrator’s distribution channels without becoming over-dependent. That evaluation starts with “Why”: Does the ecosystem offer a route to customers the firm cannot reach solo? Next comes “What”: Is its module sufficiently unique or hard to replicate to earn attractive economics? Finally, “How”: Can the firm adhere to the technical standards, compliance requirements, and roadmap cadence that the orchestrator demands while retaining room to evolve its own proposition?

Some firms advance beyond mere participation and seek to shape an ecosystem’s architecture. Shaping typically emerges in the formative stages of a domain: when standards are fluid, governance unsettled, and no single entity yet commands legitimacy. Here, technology suppliers, industry consortia, or coalitions of complementors collaborate to define data schemas, protocol layers, or compliance frameworks that eventually become hard wired into the system. Shapers trade positional authority for influence: they may not control end-user access, but by embedding their preferred interfaces or intellectual property into the scaffolding they lock in long-term advantage. The shaping role therefore appeals to actors with specialised expertise that others depend on, or with regulatory credibility that can convene rivals around a

common rulebook. The calculus again follows the Why–What–How sequence. Why invest in shaping rather than accepting an orchestrator's standards? Because early contribution can pre-empt rival architectures and secure licensing fees or switching costs later. What must be offered? A blueprint that addresses mutual pain points such as interoperability, security, data portability, better than competing proposals. How will shaping be executed? Typically through open governance bodies, shared reference implementations, and transparent road-mapping that reassure would-be adopters the standard will remain stable yet evolve with their needs.

At the apex sits orchestration: designing the architecture, governing access, and steering the ecosystem's evolution. Orchestrators define the technical interfaces, curate participation, and adjudicate conflicts; they also absorb the lion's share of upfront investment in infrastructure, brand building, and partner enablement. In return they capture outsized value through data aggregation, multi-sided fees, and platform lock-in. Yet orchestrating is expensive and risky. It demands a cocktail of assets that few incumbents or entrants possess simultaneously: broad customer access, a credible neutral brand that rivals are willing to trust, deep pockets to subsidise early adoption, and an organisational culture capable of balancing openness with control. The Why–What–How test here is unforgiving. Why should the market accept a new orchestrator when competing platforms already exist? What unique pain point or unserved customer journey justifies another hub? How will the orchestrator secure legitimacy—through exclusive data, regulatory endorsement, or an irresistible initial user base—and how will it fund the subsidy required to kick-start network effects?

Critically, the three choices are not mutually exclusive over time. A start-up might enter an ecosystem as a niche complementor, learn the standards, build relationships, and gradually evolve into a shaper as its module becomes indispensable. A corporate with strong domestic brand recognition could orchestrate at home yet act as a complementor abroad where it lacks legitimacy. Flexibility is baked into the Why–What–How logic: the answer can shift as capabilities deepen and the ecosystem landscape matures.

Deciding among join, shape, or orchestrate therefore begins with an unflinching audit of assets, ambitions, and risk tolerance. Firms should map the ecosystem's dependency graph—who controls user access, who owns critical interfaces, where bottlenecks reside—and locate their potential leverage points. They must quantify the investment required to ascend from complementor to shaper, or from shaper to orchestrator, and weigh it against probabilistic returns. Above all, they need clarity on integration and scaling: the technical hooks, governance commitments, and partner incentives that will translate strategic intent into repeatable execution.

By confronting the Why–What–How questions early, organisations avoid the twin perils of overreach—pursuing orchestration without the means to sustain it—or complacency, settling for a thin complementor role when shaping could lock in durable power. Strategic navigation of ecosystems is less about visionary rhetoric than disciplined choice: enter where you add irreplaceable value, shape only if you can set the rules, orchestrate only if you can pay the price and command the trust.

Chapter 2.2 – Roles in Ecosystems

Successful ecosystems thrive when each participant occupies a clearly defined role that fits the system's architecture and incentives. At the centre sits the orchestrator, the actor that frames

the customer promise, sets the technical interfaces and governance rules, and continuously refines the value blueprint as new opportunities and risks emerge (see Adner, 2012, 2021). Apple's stewardship of App Store guidelines, DBS's curation of its POSB "school-buddy" network and Majid Al Futtaim's data protocols for its SHARE platform illustrate how orchestration is less about owning every asset than about choreographing how others plug in and how the surplus is shared. Orchestrators earn the right to lead by combining privileged user access, brand legitimacy, and the balance-sheet capacity to subsidise early adoption; without those assets, attempts to dictate standards rarely attract follow-through.

Complementors anchor the ecosystem's variety and dynamism. They develop modular apps, services, hardware, or data feeds that enlarge the customer proposition without requiring the orchestrator to carry the cost or risk of innovation. A fintech that embeds instalment payment options in Carrefour checkout, a recipe-content creator that programmes Haier's smart oven, or a sensor start-up that pushes real-time diagnostics to an e-mobility routing engine all exemplify how complementors extend functionality while preserving their own identity and business models. Their strategic calculus turns on three questions: whether the ecosystem offers reach they cannot achieve alone, whether their capability is sufficiently differentiated to command rents, and whether integration overhead and policy volatility leave enough margin.

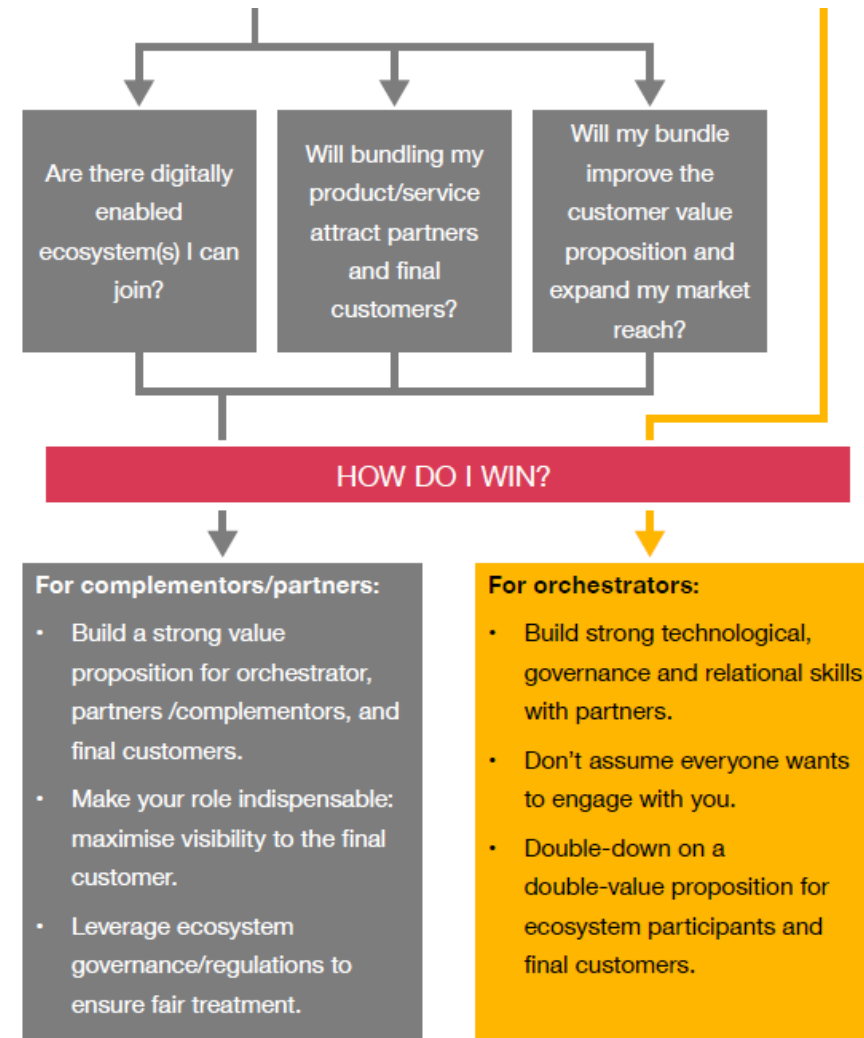
Keystones occupy the narrow set of positions that make the entire web resilient: identity providers, data-cleansing hubs, security certifiers, or critical middleware whose failure would cripple multiple flows at once. They are not necessarily large firms, but they stabilise the system by ensuring continuity and performance where fragmentation would be catastrophic. Stripe's payments rail in countless retail ecosystems, identity-verification services that underpin open banking APIs, or cloud observability tools in industrial IoT stacks show how keystones anchor trust and operational coherence. Their power is structural: by lowering systemic risk they become hard to displace, yet they must remain scrupulously neutral to maintain credibility with both orchestrator and complementors.

Users, finally, are not passive consumers but active contributors particularly in "experience-to-everything" models where behaviour data and feedback loops refine the offer in real time. Each grocery scan in SHARE, each home-energy preference logged by a connected thermostat, or each transit point earned through Velocia feed back into algorithms that reshape recommendations, price signals and product design. In turn, visible impact on the experience motivates deeper engagement, closing a virtuous circle that no orchestrator or complementor could simulate in isolation. Empowered users can also discipline governance: when rule changes threaten perceived fairness, defections or vocal backlash quickly erode network effects.

Healthy ecosystems therefore hinge on the alignment of these four roles. Orchestrators must offer complementors compelling economics and clear road-maps; complementors must respect interface integrity and data responsibilities; keystones must operate transparently; and users must perceive sustained value. Missteps in any link such as an orchestrator squeezing margins, a keystone abusing information asymmetry, a wave of low-quality complements propagate through the system and can trigger defection spirals. Conversely, when roles remain distinct yet mutually reinforcing, the ecosystem compounds innovation, loyalty and resilience faster than any vertically integrated firm could hope to match.

Figure 2.1: Deciding Participation and Role in an Ecosystem

Choosing your ecosystem strategy



Source: Michael Jacobides

Chapter 2.3 – Tools for Participation

Ecosystem participation is too complex to rely on instinct; firms need a systematic toolkit that shows where to play, what to watch, and how to prepare. The starting point is role-mapping. A concise map exposes the invisible dependencies that bind an ecosystem together: who controls end-user access, who owns data pipes, who supplies irreplaceable complements, and where bottlenecks could appear when demand spikes or standards change. Mapping begins with the user journey and works backward, tracing each activity through the actors, positions, and links required to make that journey seamless. The exercise quickly clarifies power asymmetries—perhaps a payment rail sits between every transaction, or a niche sensor vendor feeds critical telemetry to half the stack. By translating an abstract web of logos into a concrete dependency graph, role-mapping reveals leverage points for negotiation and highlights which relationships deserve contractual guardrails versus lighter, API-level agreements.

With dependencies visible, management can apply the Ecosystem-as-Structure lens to evaluate risk (Adner, 2012, 2017). This framework distinguishes two hazards that traditional strategy tools often blur: co-innovation risk and adoption risk. Co-innovation risk asks whether partners must advance in lockstep for the offering to work at all. A connected-car service, for example, only delights customers if automakers, telcos, mapping providers, and roadside infrastructure all deliver compatible upgrades on roughly the same schedule. Adoption risk, in contrast, probes whether end users will embrace the full bundle once the pieces are technically ready—an electric-vehicle ecosystem may be fully functional, yet drivers might hesitate without visible charging density or trust in battery longevity. By separating these risks, the lens helps managers decide where to invest persuasion capital. If co-innovation is the bottleneck, the firm might subsidize developer kits, share road-maps, or even underwrite partner prototypes to accelerate convergence. If adoption is the sticking point, pilot programs, warranties, or visible demonstrations may be more impactful than further engineering sprints.

The third tool in the participation toolkit is a capability audit that inventories the assets required to execute the role identified in the map and withstand the risks surfaced by the ecosystem lens. Four questions dominate. Does the firm control distinctive data that others need? Can it influence interface logic by owning or co-defining key APIs? Does it possess recognised authority to set or certify standards? And, crucially, does it reach users directly, either through a channel it owns or through persuasive brand gravity? The audit grades each dimension on both current strength and ease of improvement. A software scale-up may score high on data science but low on regulatory credibility; a hardware incumbent may enjoy brand trust yet lack agile release pipelines. By cross-referencing these scores with the dependency map, executives can spot gaps that threaten their chosen role. If user-access control is shaky, the firm might pursue distribution partnerships before committing R&D funds. If standards influence is weak, joining a consortium early, or hiring talent with protocol experience, can prevent downstream lockout.

Taken together, role-mapping, the Ecosystem-as-Structure risk assessment, and capability audits form a repeatable loop. The map highlights who matters; the lens clarifies what could go wrong; the audit tells whether the firm is equipped to handle it. Running the loop at regular intervals—say, quarterly for fast-moving digital domains—keeps strategy grounded in concrete interdependencies rather than aspirational slides. It also disciplines resource allocation: capital flows to filling capability gaps that the map and lens prove are critical, not to vanity projects that merely extend internal scope. Smaller firms, in particular, gain the confidence to negotiate fair terms when they can point to a clear contribution: exclusive data feed, compliance module, niche algorithm backed by a transparent assessment of ecosystem needs.

Perhaps most importantly, these tools prevent what might be called ecosystem myopia. Enthusiasm for joining a buzzy platform often obscures the reality that complementors bear disproportionate co-innovation costs, or that adoption hinges on a third party's marketing spend. By foregrounding dependencies, risks, and owned capabilities, the toolkit forces a sober conversation about timing and sequencing. A start-up may discover that it should first embed in a narrower set of use cases where adoption barriers are lower, gather usage data to harden its algorithm, and only then scale across a broader ecosystem once co-innovation hurdles drop. Conversely, an incumbent contemplating orchestrator status may realise that its current clout over interfaces is ephemeral and pursue a co-shaping role in standards bodies before unveiling a grand architecture.

In practice, applying these tools shifts ecosystem participation from an opportunistic gamble to a managed process. Role-mapping supplies the visual blueprint; the Ecosystem-as-Structure lens provides a diagnostic of interdependent timing; capability audits anchor ambitions in reality. Firms that institutionalise the combined toolkit make better decisions on partnership depth, governance structures, investment pacing, and exit options. Those that skip the discipline risk building impressive technology that languishes for lack of aligned complements, or burning cash subsidising partners when the real obstacle is end-user scepticism. The path to resilient participation starts not with code, but with a structured view of where value truly depends on others—and where the firm itself is uniquely positioned to shape the outcome.

Three distinct roles in multifactor ecosystems, each with upsides & costs



Roles in ecosystems	Role objectives
 Orchestrator	▪ Fulfils the end-user's central need & is the primary end-user contact point in the ecosystem ▪ Provides core (essential) products & services ▪ Arranges provision of further products and services from ecosystem's participants
 Partner	▪ Fulfils specific end-user needs from core business ▪ Clearly differentiates products, services, and brand ▪ Secures own end-user contact point
 Complementor	▪ Contributes to the fulfillment of a specific end-user need related to core business with a component ▪ Designs product and services for different ecosystems, products easily integrated by different orchestrators and partners

Chapter 2.4 – Becoming a Plausible Orchestrator

A credible orchestrator begins by framing a sharp, unmet problem and positioning the ecosystem as the most elegant way to solve it. That “north-star” proposition does two jobs at once: it attracts the first wave of partners by clarifying where new money lies, and it disciplines internal decision-making when tempting tangents appear. Leaders therefore launch not with a feature list but with a narrative—why the pain point matters, how modular collaboration beats vertical control, and what each participant stands to gain. When Apple introduced the App Store it spent as much time explaining the revenue split and review process as it did touting iPhone hardware; the clarity of both purpose and mechanics made enrolment a rational choice.

With the end-user promise nailed, the orchestrator designs a minimal viable ecosystem—just enough actors and modules to deliver the core experience and prove super-additive value. Superfluous domains wait outside the first release; every added module raises coordination costs and dilutes focus. A tight kernel also accelerates positive feedback loops: the shorter the distance from partner effort to visible customer impact, the faster complementary investment follows. Once momentum builds, the orchestrator can widen the circle, always sequencing additions in ways that extend the customer journey and deepen data insight rather than merely bulk up headcount.

Competitive advantage flows from moats that ecosystems alone can defend. Data gravity is the most obvious: usage telemetry, preference profiles, and performance benchmarks intensify with every transaction, making rival entry more expensive day by day. Interface entrenchment is the second: stable APIs, software development kits, certification regimes, and reference hardware create switching costs running through the developer’s codebase, the supplier’s tooling, and the user’s habits. Brand legitimacy is the third pillar; in a world where multiple ecosystems chase the same complements, the one perceived as most predictable and fair will attract higher-quality partners, compounding network effects and discouraging fragmentation. Moats therefore emerge from a blend of technology, process, and trust—not from secrecy or coercion.

Differentiation against competing ecosystems hinges on specific pain points rather than scale for its own sake. A retailer might build its moat around hyper-local delivery windows others struggle to match; a financial-data hub might focus on real-time reconciliation that incumbent networks clear only overnight. By solving a tangible bottleneck, the orchestrator secures indispensability even when broader functionality overlaps with rivals. Over time, the orchestrator layers adjacent services, but each extension must reinforce the original edge (turning proprietary logistics data into predictive inventory tools, or leveraging reconciliation speed to launch intraday lending) so the moat widens rather than drifts.

Trust is the orchestrator’s scarce currency and must be cultivated deliberately. Partners need confidence that the rules of the game will not shift when they become successful. Transparent roadmaps, published interface change logs, and dispute-resolution forums signal procedural fairness. Value-sharing formulas should be explicit, anchored in easily auditable metrics such as gross merchandise value or active-usage thresholds. When economic asymmetry is inevitable—for instance, a zero-marginal-cost digital service paired with capital-intensive hardware—the orchestrator can restore balance through tiered fees, marketing credits, or data-access privileges. Periodic ecosystem health reviews, shared openly, reinforce the message that mutual gain, not unilateral extraction, guides evolution.

Communication rhythms matter as much as content. At launch, orchestration is almost evangelical: concise manifestos, founder videos, and hackathons galvanise early adopters. As

the ecosystem matures, dialogue shifts to operating cadence—release notes, partner councils, road-map conferences, and quarterly performance dashboards. The orchestrator becomes less a preacher and more a choreographer, synchronising incremental innovations across hundreds of semi-autonomous actors. Crucially, communication channels remain two-way; partner feedback on pricing, certification bottlenecks, or user pain points feeds directly into governance tweaks, signalling responsiveness and cementing loyalty.

Finally, an orchestrator must internalise that competition occurs at two levels simultaneously: ecosystem versus ecosystem, and module versus module within each system. Winning the first battle means enabling healthy rivalry in the second. A vibrant marketplace of complements keeps users engaged and spurs continual improvement, which in turn enlarges the top-of-funnel for the orchestrator's core services. Attempts to suppress intra-ecosystem competition—by favouring an in-house app or cornering a complementary segment—often backfire, pushing high-value partners into rival ecosystems and eroding the differentiating variety that attracted customers in the first place.

In sum, becoming a plausible orchestrator is less about raw scale than about disciplined design: articulate a compelling purpose, launch with a focused kernel, deepen defensible moats around data and interfaces, differentiate through pain-point mastery, and build unambiguous rules that partners can trust. Do that, and the ecosystem's compounding loops of innovation and adoption will outpace any rival relying on size or proprietary control alone.

Chapter 2.5 – Becoming a Good Partner or Complementor

Ecosystem success is not reserved for giants; smaller and mid-sized companies can win by choosing disciplined participation plays that exploit their distinctive assets while neutralising power asymmetries with the hub. The first approach is the “plug-in” strategy: build a narrowly focused module—an algorithm, a payment microservice, a sensor—and design it for effortless multihoming so it snaps into several ecosystems without bespoke rewrites. A plug-in relies on technical excellence and speed; its value comes from reducing time-to-market for orchestrators that crave rapid feature expansion but cannot staff every niche. Revenue flows through usage-based fees or shared upside clauses that scale automatically with volume, keeping negotiation friction low.

A second path is collaborative bundling. Here, firms lacking direct consumer reach surround a high-engagement product with complementary services they co-design with the orchestrator. A fitness-equipment maker, for instance, can fold nutritional coaching and sleep-tracking into a single subscription that rides on the platform's identity and billing rails. The payoff is recurring income and richer telemetry, but the trade-off is partial cession of user access. Protection comes from embedding proprietary IP (patented biomechanics, exclusive content etc) that the hub cannot easily replicate and that users would miss if the bundle dissolved.

Hosting offers a third avenue. Instead of fighting to stand out in a crowded marketplace of apps, a supplier turns itself into micro-infrastructure for others: a regional data-cleansing service, a compliance sandbox, or a localisation layer that global platforms need to satisfy regulators. Hosting locks in demand through switching costs and regulatory inertia, turning what began as a point solution into a keystone that underwrites ecosystem resilience. Pricing combines flat subscription tiers with metered overages, ensuring predictable base revenue plus upside from growth.

Optimisation is the fourth play. Companies with deep operational know-how such as manufacturing analytics, energy balancing, last-mile routing, embed continuous-improvement algorithms inside the ecosystem's workflows. Their software mines live process data, flags inefficiencies, and triggers automated interventions. Fees are contingent on verified savings, aligning incentives and discouraging opportunistic haggling. Over time, optimisation partners gain privileged visibility into system performance, a data asset they can recycle into new offerings or leverage when renegotiating terms.

The fifth option, orchestration lite, is viable for regional incumbents with entrenched distribution. These firms curate a micro-ecosystem - say, an agri-tech stack serving a specific climate zone and then federate it upward into larger platforms via standard APIs. They retain authority within their niche while benefiting from the scale and reach of global hubs. Governance balances autonomy and compliance: local standards and brand posture domestically, strict adherence to umbrella rules on security, privacy, and revenue sharing internationally.

Across all five plays, bargaining power hinges on four assets: distinctive data, technical bottlenecks, regulatory licences, and trusted relationships. Firms should audit which of these they genuinely command and fashion their participation model accordingly. A start-up with a unique diagnostic dataset should negotiate data-share clauses that grant it derivative-work rights; a payments specialist with a coveted licence can demand equity options in exchange for integration exclusivity.

Risk management requires early clarity on exit options. Contracts should include data-portability rights, code-escrow provisions, and reasonable sunset periods for API deprecation. Small players lacking leverage can pool their negotiation through alliances or industry consortia, presenting a united front on royalties and interface changes.


Finally, success metrics must evolve beyond near-term revenue. Leading indicators such as module adoption rate, cross-ecosystem activation, latency improvements, and partner satisfaction scores show whether the chosen play is compounding or stalling. Quarterly ecosystem health reviews, tied to investment gates, prevent sunk-cost fallacies and keep strategy grounded in measurable traction.

In short, firms that understand their unique strengths, select the participation play that amplifies those strengths, and embed robust safeguards can thrive alongside hubs far larger than themselves. Ecosystems are not a zero-sum arena; with smart design, even modest contributors can capture durable value while helping the whole system flourish.

ROIC			
	<i>Return / profit potential</i>	<i>Capital Invested</i>	<i>Risk profile</i>
Orchestrator	High, but not guaranteed	High	High
Complementor	Limited, but substantial	Limited	Low / Medium

The Figure below provides a summary of our advice (drawing on projects that Evolution Ltd has undertaken) on how to make all the key choices that ecosystems entail. It helps provide the discipline that takes intuition and build it out so that it becomes actionable and so that we can empower executives to drive the ecosystem play- and measure their progress, too.

Figure 2.2: An Ecosystem Development Framework



Chapter 2.6 – SME Design and Readiness

Small and midsize enterprises rarely have the luxury of dictating ecosystem architecture, yet they do control how quickly and credibly they can plug into one. Three practical pathways dominate. The simplest is the plug-in approach, where an SME connects its niche solution to an established platform through pre-built connectors or marketplace apps. Katana MRP's one-click link with Shopify, which exposes real-time inventory to e-commerce orders, is a textbook example: zero new infrastructure, immediate reach to merchants, and a revenue model aligned with transaction growth. A second route is joint go-to-market, in which the SME and a larger orchestrator bundle their products into a single offer and push it through shared channels. This works best when the SME fills a visible gap - say, an AI scheduling module that lifts the conversion rate of a mobility platform, so the orchestrator has clear incentive to co-promote rather than subsume. The third path is deeper, API-based integration. Here the SME exposes or consumes programmatic interfaces so that workflows feel native inside the host environment; pay-per-use pricing and developer-first documentation make uptake frictionless for the orchestrator and its other partners.

Regardless of entry path, design principles decide whether the relationship scales. Rapid integrability is non-negotiable. That means RESTful or GraphQL APIs, authentication via standard OAuth, and an SDK or code samples in at least two mainstream languages. SMEs should assume the orchestrator's developer team will grant them no more than a half-day sanity check; if hand-holding is required, the slot will go to a rival whose library installs in minutes. Modularity is the second pillar. A complement that bundles six features behind one massive dependency tree will fail security review and slow release cadences; one that offers bite-sized endpoints for data ingestion, analytics, and visualisation separately lets the orchestrator compose only what the user journey needs today and extend later. Low onboarding friction rounds out the triad. Self-service sandboxes, automated SLA monitoring, and usage-based pricing that scales from free tier to enterprise assure product managers they can experiment without procurement red tape. In practice, every additional form or manual approval slashes adoption probabilities by double-digit percentages.

Yet technical elegance alone is insufficient; complementors must also align with orchestrator expectations while preserving strategic independence. Alignment starts with roadmap transparency. An SME should publish a quarterly summary of planned releases and deprecations and seek early feedback from the orchestrator's product leads. This reassures the hub that the complement will mature in step with evolving customer journeys. The flip side is guarding against dependence. Dual-core development (maintaining connectors to at least one alternative platform) keeps negotiation leverage and forces disciplined abstraction layers in the codebase. Revenue alignment is next. Usage-based revenue shares that mirror the orchestrator's own monetisation logic, or tiered pricing tied to active users, ensure the hub's growth directly benefits the complement and vice versa. Independence is maintained by blending these revenue streams with a stand-alone offering sold under the SME's own brand, so pivot options remain open.

Governance rituals reinforce the balance. Regular performance reviews anchored in jointly agreed key metrics (uptime, latency, conversion lifts) provide a factual basis to escalate issues without political friction. Clear escalation paths protect the SME from being blindsided by unilateral policy shifts, while also signalling accountability to the hub. For disputes, lightweight mediation clauses or predefined exit procedures keep legal costs in check and prevent souring relationships from paralysing the entire value chain.

Finally, cultural readiness often separates SMEs that thrive from those that stall. Engineering teams must treat external API changes as routine stimuli, not costly emergencies; agile sprints should include capacity for adaptation work. Customer-success staff need playbooks for two-step support flows, where first-line inquiries may come through the orchestrator's help desk before reaching the SME. Finance must be comfortable reconciling micro-transactions rather than monthly invoices. Leadership, for its part, has to embrace the paradox of being both collaborator and competitor in the same web—celebrating shared wins while continuously scanning the horizon for new niches where the SME can differentiate.

When plug-in connections, modular design, and balanced governance converge, an SME gains more than scale: it earns a reputation as a dependable, adaptable complementor. In an era where orchestrators face a cue of would-be partners, that reputation is the most valuable signalling asset an SME can accumulate, unlocking faster approvals, joint marketing opportunities, and privileged roadmap access. In short, ecosystem readiness is not a one-time integration project but an organisational capability—one that transforms the SME from a passive supplier into an agile, self-directed node in a larger, ever-evolving network.

Chapter 2.7 – Failure Patterns and Avoidance

Ecosystems promise outsized growth, but the same forces that create exponential upside can quickly turn on the orchestrator. Four recurring failure modes stand out. Understanding them early, and hard-wiring countermeasures into architecture, governance, and culture prevents months of re-work and millions in squandered goodwill.

The first and most important pitfall is the ego-system trap: designing around internal ambitions instead of external jobs to be done. Teams lavish resources on features that showcase proprietary technology, only to discover that complementors cannot build profitable services on top, or that end users barely notice the improvements. Symptoms appear in metrics long before revenue stalls: partner churn rises, API calls plateau, and net-promoter scores remain stubbornly flat. The cure is ruthless, ongoing validation of real customer frictions. Instead of anchoring strategy in a grand vision deck, orchestrators should embed objective-driven “problem sponsors” who own specific user outcomes and wield veto power over feature releases that drift from those outcomes. A public roadmap disciplined by customer pain points keeps the centre of gravity outside the corporate ego and inside the marketplace.

Opposite in spirit but equally lethal is over-centralisation. Fearful of fragmentation, orchestrators dictate detailed interface specifications, mandate exclusivity, or pre-empt promising niches with in-house products. Complementors respond by limiting investment, multihoming to rival ecosystems, or delaying critical updates. The platform then enters a negative spiral: fewer novel complements reduce user excitement, slowing adoption and further eroding partner ROI. Successful hubs strike a balance between coherence and creative latitude. They protect system integrity by certifying minimum performance and security standards, then step back, letting partners differentiate on experience and pricing. Selective modularisation such as tight control at safety-critical layers, openness elsewhere, creates guardrails without suffocating invention.


A subtler danger is governance ambiguity. Early in an ecosystem's life, informal promises and personal relationships fill gaps in contracts. As scale arrives, opacity around revenue splits, data rights, or algorithm-ranking criteria breeds suspicion. High-value partners hedge their bets;

new entrants balk at the cost of integration. Clarity is not only about legal documentation; it is a living discipline. Mature ecosystems publish value-sharing formulas tied to verifiable KPIs, version their terms alongside API changes, and convene routine councils where grievances surface before they metastasise. Equally important is a predefined escalation path (mediation, arbitration, or structured exit) so that inevitable conflicts have a constructive outlet instead of spilling into social media or regulatory complaints.

Even when technical and contractual scaffolding is sound, organisational resistance can derail momentum. Siloed functions inside the orchestrator (engineering, compliance, sales) optimise for their own KPIs, creating drag on partner enablement. A developer who waits weeks for security clearance or an account manager who cannot access unified partner data soon disengages. Internally, the remedy starts with cross-functional “ecosystem pods” that hold end-to-end accountability for partner success, from first contact to scaling. Externally, the orchestrator must equip complementors with self-service analytics and clear points of contact so they can glean insight from shared data rather than pleading for ad hoc reports. Cultural signals matter: leadership should celebrate wins where partners, not the core business, capture the lion’s share of new revenue. Such stories legitimize collaboration and dampen zero-sum mindsets.

Taken together, these four failure modes teach a common lesson: healthy ecosystems thrive on aligned autonomy. Orchestrators must locate the sweet spot between vision and humility, control and openness, promise and proof. That demands continuous investment in three capabilities. First, listening systems—user ethnography, partner advisory boards, telemetry feedback—ensure the problem space stays anchored in lived pain points. Second, adaptive governance codifies incentives and rights in language precise enough for scale yet flexible enough to evolve. Third, boundary-spanning teams dissolve internal walls, treating complementors as extensions of the core value proposition rather than downstream customers.

Ecosystem building is less a heroic sprint than a sustained choreographic exercise. Companies that avoid the ego-system trap, temper centralisation, codify value-sharing, and break their own silos position themselves not merely to launch a platform, but to nurture a living network that compounds value for every participant over time.



Chapter 2.8 – “No Regrets” Steps

Momentum in an ecosystem can stall for many reasons, but four moves deliver upside whatever role or governance model a company pursues. They need modest investment, build organisational muscle, and impose almost no lock-in—true “no-regrets” actions.

First, treat interoperability as a product. Every failed API call, opaque field and missing sandbox kills partner enthusiasm before commercial terms are discussed. A small cross-functional integration guild that owns docs, reference code and developer success metrics cuts months from onboarding and creates a positive selection effect. Versioned schemas, machine-readable OpenAPI specs and live status dashboards may look like plumbing, but they set the ceiling on how many parallel experiments the ecosystem can run. Lower friction gives strategy more choice.

Second, weave the ecosystem story into mainstream business development and marketing, not a side project. Partner value propositions should live in the same decks sales teams use with end customers; case studies should celebrate joint wins. This clarity deters the “egosystem” drift: if marketing cannot explain in one line why external actors matter, ambition has slid inward. Internally, align quotas and campaigns to ecosystem KPIs; externally, signal that the aim is growing the pie, not harvesting data.

Third, start governance light, bilateral and reversible. Early partnerships run on speed and trust, not 40-page frameworks. A concise memorandum covering data rights, branding and a simple revenue share is usually enough for a first release. Add councils, certification tiers or arbitration only after usage passes visible thresholds. This keeps legal overhead proportional to value and surfaces misalignment early: if a prospective partner wants exclusivity or complex indemnities before traction, take the hint.

Finally, choose battlefields where complementor success is visible and role boundaries are clear. High-performing ecosystems have living exemplars—apps earning seven figures, hardware add-ons with positive margins, analytics firms featured on stage. Such stories lower perceived risk for newcomers and help leaders justify resources. Clear demarcation also reduces channel conflict: when SMEs see that analytics partners earn telemetry fees while the orchestrator focuses on bundles, they know where to invest.

Taken together: interoperability lays the rails; narrative alignment sets commercial direction; light governance buys flexibility; visible partner wins create social proof. None of these fixes a role; they raise readiness, credibility and optionality. Whether the next turn is to orchestrate a niche, shape a standard or double down as a specialist complementor, the organisation will move faster and negotiate from strength. In a landscape where timing, trust and adaptability matter as much as assets, these are as close to guaranteed returns on strategic attention as one can get.

Cases Chapter 2: Ecosystems as drivers of focused value propositions

C2.1 – Growth Through Ecosystem Use in MasMovil and COSMOTE

Telecoms faced the same bind: connectivity was commoditising, price wars eroded margins, and portability kept churn high. MásMóvil (Spain) and COSMOTE (Greece) escaped by orchestrating daily-use, partner-delivered services on open rails, designed to lift touchpoints and reduce churn. Vodafone tried versions of the same idea but fell into classic ecosystem traps—the devil is in the design choices, not in the slogan.

MásMóvil set its scope ex-ante: move into high-frequency categories with low integration cost where it could reuse identity, billing and support (electricity first, then insurance). It partnered with agile specialists rather than incumbents, avoided in-house builds, and governed participation through clear APIs and revenue-shares tied to outcomes. One north-star metric, net avoided churn, gated every launch. The result was churn roughly half the Spanish market average and accretive economics even while undercutting on price.

COSMOTE started with payments as the spine. A wallet and debit product created daily data and trust, enabling partner-underwritten insurance and simple savings inside the self-care app. Roll-out was bottom-up: pilots had to show take-up and a positive NPS before national scale. Where complexity or culture argued against building, COSMOTE partnered instead of persisting with internal pet projects. The operator kept identity, billing and service quality, while letting complements own pricing, risk models and regulatory capital.

Vodafone, by contrast, shows how strong assets can be undermined by weak ecosystem design. The Vodafone 360 applications play (2009) launched with a thin partner network and an in-house experience that could not compete once Android and iOS app ecosystems surged. More broadly, Vodafone oscillated between internal and external plays, sending mixed signals to partners and spreading investment across initiatives that never compounded. Partner-selection criteria that looked sound on paper struggled in practice; several high-profile combinations suffered technology and culture mismatches; and expansion moved into arenas with little defensible advantage, yielding look-alike bundles without stickiness.

Three design choices explain the divergence. First, scope discipline: MásMóvil and COSMOTE entered categories with frequent use and clear customer jobs, where telecom rails (ID, billing, support) genuinely reduce friction; Vodafone spread into areas where its advantages were thin. Second, partner incentives: the winners picked complements whose success wasn't cannibalised by the telco, wrote simple rules (APIs, service levels, revenue-share) and measured outcomes everyone could live with; Vodafone's mixes of in-house and external efforts dulled incentives. Third, execution cadence: pilots, proof, then scale with a single metric (churn/NPS) deciding funding, versus broad initiatives without a unifying scoreboard.

The lesson is not that “ecosystems work for challengers” and “fail for incumbents.” It is that orchestration is a craft: choose domains where your rails create real advantage; keep complements' economics attractive; make rules legible; and measure the one outcome that actually pays back the bundle (churn, not vanity usage). Done this way, connectivity becomes the anchor for a wider, partner-led bundle that customers revisit—and stay for.

C2.2 – KBC: intent-driven orchestration with Kate

KBC's move from universal bank to ecosystem orchestrator rests on a simple premise: everyday actions reveal intent. Card taps, bill payments, address changes, travel bookings or a new payroll pattern each signal what a customer is trying to achieve. Kate, KBC's conversational assistant, listens for those consented signals, translates them into standardised "intents," and routes them to the right mix of KBC products and partner services—financing plus insurance plus a vetted repair in one flow, rather than separate, form-heavy journeys. KBC's investor materials now place AI and Kate at the core of digital sales and advice, reflecting how orchestration, not just distribution, has become a strategic spine.

The mechanics are straightforward. When behaviour suggests an opportunity (e.g., moving house, a car purchase, a large home expense) Kate proposes actions in natural language and calls the minimum set of partner endpoints (quote, schedule, fulfil) needed to get the job done. Partners plug in through published APIs and clear service-level and reporting rules; outcomes are measured on solved tasks, retained balances, claims paid on time and churn avoided, so the assistant does not devolve into a thin sales funnel. Industry coverage indicates Kate now handles hundreds of thousands of customer queries monthly and sends well over a million proactive nudges, illustrating the shift from static menus to intent routing at scale.

KBC has also rebuilt incentives around this orchestration. In September 2025 it upgraded Kate Coins from a traditional, brand-specific reward into a broader benefits layer: customers can earn and redeem more easily across KBC and a widening roster of partners (e.g., Colruyt, Kinopolis, HelloFresh, Telenet, Amazon, Pizza Hut, Q-Park, Bofrost). The change makes rewards portable across domains—retail, mobility, hospitality—so partners compete on verified outcomes rather than coupon breadth, and customers experience one coherent programme instead of siloed deals.

The transition mattered operationally too. As KBC introduced the new model, legacy partner deals attached to "old" Kate Coins expired on 31 August 2025, with a simpler, wider set of benefits rolling out the following week. That reset aligned incentives with the intent-routing design: fewer one-off promotions, more offers priced to measured lift and delivered through the same rails Kate already uses.

Signals from outside the bank corroborate the trajectory. Sia Partners' 2024 global benchmark ranked KBC Mobile the #1 banking app worldwide, citing the breadth of functionality and the virtual assistance layer as differentiators—useful external validation that the assistant-centred model is showing up in customer experience, not just in architecture diagrams.

Two design choices keep the ecosystem attractive to complements. First, neutral governance: documented eligibility and ranking criteria, auditable flows, explicit customer consent and the ability to revoke sharing by intent; no hard exclusivities where multiple high-trust providers are valuable. Second, programmable incentives: benefits attach to verified behaviours (safe driving, on-time bills, greener tariffs, helpful cross-sell) rather than generic segments, which lets both KBC and partners price rewards precisely and share uplift credibly. Combined, these features lower acquisition costs for smaller partners and make their participation repeatable without bespoke integrations.

There are risks to manage: an assistant can drift into self-preferencing, and a rich signal base can tempt over-targeting. The countermeasures are the same as those that made the system attractive in the first place: published rules, outcome-based measurement, conservative privacy

defaults and “design for exit” in partner contracts. Done that way, the assistant functions as shared infrastructure for an ecosystem rather than a one-sided gatekeeper.

The strategic lesson travels. Standardising intents and interfaces turns raw interaction data into routable requests that many firms can answer, while programmable rewards make participation economically rational for both global and local partners. That is how an incumbent can leverage its information advantage to orchestrate a wider network, converting signals into solved tasks, value for complements and durable economics for the hub.

C2.3 – DBS’s POSB Ecosystem with “School Buddy”

Smart Buddy is POSB/DBS’s school-based cashless payments and financial-literacy program, where students use a wearable or card linked to a parent app for allowances, spending controls, savings goals and (optionally) fitness tracking, accepted at school canteens/bookshops and selected merchants. It’s a solution for parents who can give flexible funding to their children, they can track their spending without them giving them a smartphone or cash or credit card.

This signifies POSB/DBS’s shift from traditional universal bank to a data-enabled ecosystem orchestrator, packaging identity, payments and consent into shared rails. POSB’s advantage was never exotic technology; it was proximity. As Singapore’s community bank within DBS, it sat at the junction of salaries, bill payments, school programmes and the neighbourhood merchant economy. That vantage point made a narrow retail bank model feel limiting. The shift was to use national digital rails—SingPass/MyInfo for instant KYC, PayNow and SGQR for ubiquitous payments—and a lightweight orchestration layer in the POSB/DBS app (with PayLah! as the wallet front door) to connect households, hawkers, schools and service providers around everyday tasks rather than products.

The design principle is that actions reveal intent. A salary credit plus a change of address signals a move; repeated canteen and stationery purchases suggest school-season needs; a run of taxi and food-delivery spends implies cash-flow pressure; a child’s canteen payments point to lunch routines and pocket-money limits. POSB wires these clues into the app experience so the bank proposes solutions at the right moment: pre-approved instalment plans when a large expense hits, bundled renter’s insurance alongside a deposit transfer, utilities switching when a new address is detected, school-linked savings rules when Smart Buddy activity spikes, merchant offers that match where the family actually shops. Partners join through documented APIs and clear reporting; the value proposition for them is straightforward—lower acquisition costs, closed-loop attribution and access to verified intent at the moment of need.

Two choices make the ecosystem work beyond payments. The first is to treat public infrastructure as shared plumbing rather than something to be replaced. Onboarding uses MyInfo; data portability flows through SGFinDex where the customer consents; payments default to PayNow and SGQR so even the smallest stall can participate without bespoke hardware. That approach lowers participation costs for complements and keeps the bank aligned with policy goals on inclusion and interoperability. The second is to package the ecosystem so SMEs and schools can plug in without a team of engineers. The Smart Buddy programme shows the pattern: a tap-to-pay watch and parent app bring cashless canteen payments, attendance and budgeting into one flow; the rails are open enough that schools choose their POS vendors and caterers, while POSB runs identity, settlement and the parent

interface. The result is a kid-sized financial system that quietly scales across thousands of micro-transactions and dozens of counterparties every day.

Evidence shows up in behaviour, not just downloads. Hawkers and small retailers kept using SGQR after incentive campaigns ended because the app experience removed friction from their daily flow and settlement was transparent. Schools extended Smart Buddy from canteen payments to bookshops and transport, because parents valued real-time controls and children adapted quickly. Merchant partners stayed because closed-loop reports linked offers to basket lift rather than clicks. Within the bank, service costs fell as routine tasks moved to chat and self-serve, while retention improved as the app became the default way to solve mundane problems that previously sat outside “banking”.

Governance keeps partners invested. Consent sits visibly with the customer; personally identifiable data are shielded from media planning and offer engines through clean-room patterns; partner eligibility, ranking criteria and service-level expectations are documented; and there are no hard exclusivities where competition benefits the user. Commercial terms emphasise verified incrementality rather than captive access to app real estate, so smaller brands can credibly participate alongside national chains. That mix—open standards at the edge, disciplined rules at the core—avoids the “egosystem” trap where a bank behaves like a tollbooth and gradually chokes variety.

AI shows up as connective tissue rather than a stunt. Models score likelihood of intent and choose when not to interrupt; natural-language flows in chat reduce abandonment; simple propensity and uplift models price rewards to behaviour rather than demographics, which is what makes small, hyper-local offers economical. In practice, this widens participation: a neighbourhood tuition centre, clinic or repair shop can expose one or two endpoints (quote, schedule, fulfil) and be discovered by the right families at the moment of need, with attribution handled by the platform instead of re-implementing CRM and analytics from scratch.

The strategic lesson is that a community bank can become an orchestrator by leaning into the rails it already shares with society and by packaging participation so that many small actors can join without bespoke deals. DBS's group-level transformation provides the architecture; POSB is its everyday, neighbourhood-scale illustration—turning national infrastructure into local ecosystems where value accrues to the orchestrator and to the complements that make daily life work better.

C2.4. Ecosystem Support/Advisory: Synocus, FuzeQube, Evolution Ltd

The growth of ecosystems has led to a set of new challenges. The factors underpinning their success or failure have only recently started to be understood and all stakeholders – from corporates to new ventures, governments and research institutions are slowly developing their ecosystem skills. Yet success or failure doesn't rely on knowing the buzzword- but rather, on being able to offer services rooted in an understanding of specific needs. While in the last few years all major consultants (Big Three, ie McKinsey, BCG and Bain, and Big Four ie EY, PwC, Deloitte and KPMG) or others like Accenture have all made ecosystem a key theme and topic for advisory, there is still considerable distance between professed expertise and ability to deliver results. This has led to the emergence of boutiques whose presence underlines the appetite for a more robust offering when navigating this complex landscape. Here are three examples.

Synocus

Synocus operates as an ecosystem integrator for regions and sectors that need many actors to move together but lack a shared operating model. The premise is practical: pick a concrete “job” that no single player can deliver—decarbonising a manufacturing basin, digitising port logistics, scaling circular materials—and then design the minimum rails that let anchors, SMEs, universities, financiers and agencies act in concert. The work starts with demand, not grants: anchor firms and public buyers pre-commit to a handful of use cases with measurable outcomes and timeframes. Around those, Synocus codifies roles, data exchanges and decision rights, creates a light governance forum with an agreed cadence, and seeds a shared KPI “scoreboard” that everyone sees, so progress and gaps are unambiguous.

What makes this more than facilitation is the emphasis on repeatable infrastructure. Synocus standardises the artefacts that usually bog projects down—eligibility rules for partners, reference contracts, data-sharing and IP clauses, identity/attestation for participants, and a neutral repository for evidence and learnings. That discipline lowers onboarding costs for SMEs and reduces political risk for public bodies. When AI is useful, it is used as connective tissue rather than theatre: data matching to surface qualified SMEs faster, simulation to test supply–demand changes before capital is committed, and simple assistants to keep multi-party coordination moving without meetings.

The payoffs tend to be visible within months: shorter time from pilot to scaled deployment; broader supplier participation because the rules are legible; and easier access to blended finance because outcomes are tracked on the same scoreboard that governs the programme. Crucially, Synocus avoids the “egosystem” trap by not trying to own the rails. It keeps governance neutral, documents change logs, and designs for multihoming so that solutions can plug into other regional or sectoral efforts. The strategic lesson echoes the report’s core argument on roles, rules and data: complex transformations are ecosystem problems, and they advance when an integrator turns policy intent and corporate ambition into a working, open operating system that others can join and reuse.

FuzeQube

FuzeQube addresses a stubborn gap in European innovation: many universities and labs generate strong IP, but tech transfer stalls because disclosure processes are uneven, market validation is slow, and it is hard to assemble experienced teams for spin-outs. The firm’s response is to treat tech transfer as a process-as-a-platform problem. It aligns incentives across the triple helix—academia, industry and government—by codifying how inventions are disclosed and assessed, how market fit is tested quickly, and how executive talent is matched to promising projects, then wrapping those steps in shared tooling and clear rules that multiple institutions can adopt.

Instead of each university rebuilding the stack, FuzeQube convenes consortia that share scarce resources—IP counsel, patentability screening, access to sector experts, pools of interim executives—under transparent cost- and value-sharing. The result is both cheaper and faster: researchers get a predictable path from lab note to protect/partner/park decisions; TTOs get comparable evidence they can act on; investors see standardised data rooms and governance; and potential corporate partners meet opportunities earlier, when co-development is still feasible. Where universities already collaborate in alliances, FuzeQube’s rails provide the missing operational mechanics, so platform effects—shared marketplaces of licensable IP, talent and pilot sites—actually emerge rather than remain a slide.

Team formation, a perennial bottleneck, is handled explicitly. The platform maintains a vetted bench of operators and subject-matter leaders, with incentives that can blend cash, success fees and equity, so spin-outs are not stranded for want of leadership at the moment they become viable. Simple AI tools support the flow but do not replace judgement: classifiers triage disclosures against prior art, matching engines shortlist mentors and executives, and dashboards surface comparable deals and likely partners; the objective is to cut search and coordination costs so humans can spend time where it matters.

For regions like Catalonia, the benefit is twofold. First, it raises the hit rate of institutional innovation by making the pathway legible and shared across institutions of different sizes, which is essential where smaller universities lack full in-house capability. Second, it crowds in complements—investors, corporates, alumni networks—because participation rules are consistent and outcomes are measured, not promised. The model's logic aligns with the report's themes: ecosystems form when roles are clear, rules are credible, and data make value visible. By turning tech transfer from bespoke craft into a common operating system, FuzeQube demonstrates that Europe's innovation problem is less a shortage of ideas than a shortage of coordination—and that the fix is an ecosystem design that lowers the cost of joining and raises the return on contributing.

Evolution Ltd

Evolution Ltd, which supported writing this report, is a boutique advisory that blends rigorous academic research with hands-on operator experience to help organisations design, govern, and scale digital ecosystems and AI-enabled services. The firm is deliberately vendor-agnostic and independence is part of its governance: no reseller focus, no captive build agenda. Rather, a focus on projects that are germane to the research and advisory focus of a group of highly experienced and skilled advisors, without the operational leverage that creates COI.

Its projects draw on its affiliates' experience and connections to global tech giants, startups, disruptors, entrepreneurs, and governments alike, and their capability is to engage key stakeholders in effective conversations and catalyse action. Its value add comes from moving beyond the buzzwords and applying a rigorous analysis to cutting-edge problems.

Evolution works with large corporates, governments, scale-ups and NGOs, and has supported leaders across finance, retail, telecoms, health, consumer, and mobility. Selected engagements include CEOs and leadership teams at KBC, Majid Al Futtaim, MásMóvil, HDFC ERGO, Lavazza, Philips, Mastercard, Imperial Brands and Match Group; collaboration with strategy firms (BCG, Strategy&) and peer integrators (Keystone, Synocus, FuzeQube); and venture partnerships with companies such as Velocia, Evolver.ai and Aegis Labs. Its frameworks and findings are widely cited and have appeared in leading outlets (e.g., HBR, SMR, CMR, *Financial Times*) and global forums such as the World Economic Forum.

Operating through a network of independent advisors, Evolution applies the same ecosystem logic it recommends: open rails, clear rules, transparent measurement. The result is design that partners can trust, economics that sustain participation, and options that keep clients in control as markets—and AI—shift.

C2.5 – HDFC ERGO: From Insurance to Mindshare

Inspired by insurance innovators like China's PingAn, the world's most valuable insurance brand, HDFC ERGO has recently built out an ecosystem play, as a result of a conscious project of strategic analysis and experimentation. What lie behind it was that it sought a creative way to increase the frequency of customer interactions, drive higher satisfaction and retention, and support better unit economics. The strategy is to turn insurance from an annual transaction into an ongoing service, raising engagement by solving adjacent jobs around core lines (motor, health, travel, home) with a web of partners.

In the customer's hand it looks simple. A cracked windscreen becomes a same-day repair via a nearby cashless garage; a fever at home turns into a tele-consultation, doorstep lab test and e-pharmacy fulfilment; a flight delay auto-triggers lounge access and instant settlement. Each flow blends cover, finance and third-party fulfilment so the customer experiences a solution rather than a policy number. That creates many small, useful moments across the year, keeping the brand salient and opening natural paths to add riders, upgrade from third-party to comprehensive, extend into home or travel, and renew with less price sensitivity.

Making this work required getting the information and data infrastructure right. HDFC ERGO built a unified, consented customer identity; a streaming event layer that detects intents in real time (damage, illness, travel disruption); a provider registry with clear eligibility and service-level rules; and an orchestration layer that calls the minimum partner endpoints (quote, schedule, fulfil) to complete the job. Without those rails, the ecosystem would feel bolted-on; with them, journeys are coherent, attribution is closed-loop, and both customers and partners can see what value was created.

The ecosystem play also required an understanding of the competitive landscape- the offerings available to the Indian market by different types of providers in different verticals, which helped pin down where the value add lie (the "outside in" analysis) and the understanding of the valuable bits that the company had to offer (the "inside out" work) that both needed to converge on an ecosystem play where the firm had a right to play and a right to win.

The economics show up on both sides of the P&L. On the loss side, guided fulfilment reduces leakage and severity (right provider, right parts, fraud caught by pattern checks) and prevention nudges lower claim frequency. On the revenue side, increased engagement lifts retention and product density per household, while outcome-priced bundles support better margins. Acquisition costs fall as the app becomes the default place to solve common problems; service costs fall as routine tasks move to self-serve and chat. Richer first-party data improves risk selection and pricing, strengthens reinsurance negotiations and makes marketing more precise.

This is an ecosystem, not a chatbot. The insurer keeps identity, adjudication and settlement; specialists keep clinical care, repair, logistics and last-mile. Roles and rules are explicit: published eligibility and ranking criteria for providers, documented service levels, audit trails for model decisions, conservative privacy defaults and "design-for-exit" so partners can leave without stranding customers mid-journey. Those choices attract better providers, deter self-preferencing and sustain participation at scale.

AI is the connective tissue rather than the headline. Computer vision validates damage, propensity models time offers only when useful, anomaly detection targets fraud without burdening honest customers, and routing models pick the best next step given cover, location

and capacity. The point is practical: AI lowers search and coordination costs so more complements can plug in and more customer intents can be fulfilled quickly.

The result is a franchise that behaves less like a commodity insurer and more like an orchestrator of everyday resilience. Customers get timely help and a reason to stay; partners get steady demand and fair credit; the insurer earns through lower combined ratios, higher lifetime value and new fee income from service orchestration. HDFC ERGO's journey shows that even in heavily regulated markets, incumbents can pivot from occasional insurer to pervasive life-assistant by systematically importing and localising ecosystem playbooks. The payoff is not only better margins but also durable mindshare—a strategic asset in a world where the battle for customers' attention is fought minute by minute, tap by tap.

Case C2.6 – Kinaxis: turning supply-chain planning into an ecosystem

Kinaxis began as a high-end planning tool; it now behaves like an orchestrator. The core move was to standardise how data, scenarios and decisions flow across a manufacturer's network—suppliers, contract plants, logistics partners, retailers—so planning is concurrent rather than sequenced. A single, shared model of demand, supply, capacity and constraints sits at the centre; partners plug into that model through documented interfaces to publish signals (orders, delays, quality alerts) and to consume what they need (forecasts, allocations, exception queues). The effect is to convert planning from a departmental exercise into a boundary resource others can build on (see §1.5).

Two design choices make this an ecosystem rather than a monolith. First, the core stays neutral and legible: versioned APIs, a common ontology for parts, sites and time buckets, and auditable scenario runs so everyone can see what changed and why. Second, Kinaxis invites complements. System integrators assemble vertical templates; niche ISVs publish add-ons—risk and ESG scores, warranty-return predictors, carbon and water intensity, supplier health checks; and customers' analytics teams ship their own models as plug-ins. Since data and scenarios sit in one place, these adds become apps against a shared truth set, not one-off integrations that decay.

The customer experience is a shift from firefighting to composition. When a contract plant goes down, the control tower does not just throw alerts; it proposes alternatives with quantified trade-offs—reroute through Plant B, expedite from Supplier C, accept a service-level dip for SKU family D—so cross-functional teams can agree the move in minutes. Execution signals loop back (order management, transport orchestration), closing the gap between plan and reality. For partners, this is a repeatable on-ramp: expose one or two endpoints (risk score, emission factor, lead-time forecast), meet service-level and transparency rules, and you're in front of global manufacturers without bespoke deals.

The economics follow. Inventory and expedite costs fall as decisions are made with the full system in view; cycle times contract because exception handling is codified; and the same rails support new KPIs (service, cost, carbon) without another data project. Critically, engagement rises: suppliers and logistics providers contribute data because they also get better visibility and fairer credit when their actions prevent a stock-out or cut emissions, mirroring the report's emphasis on measurement as legitimacy in Part 2.

AI is present but purposeful. Forecasting and constraint-solving models run “what-if” scenarios on live data; anomaly detectors surface bad master data before it cascades; natural-language

prompts generate first-pass scenarios for planners who then decide. The point is not novelty; it is lowering search and coordination costs so more actors can join the same decision loop—exactly the participation dynamic outlined in §3.4.

Governance keeps the flywheel honest. Partners are certified against security and data-handling rules; model changes are versioned with visible impact on plans; customers keep ownership of their data and can export artefacts (plans, allocations, audit logs) to avoid lock-in. That mix—open enough to attract complements, disciplined enough to sustain trust—prevents the “egcosystem” trap where a platform hoards advantage and partners disengage.

The strategic lesson is portable: treat planning as shared infrastructure. When a neutral core makes data and decisions legible, the whole network can co-innovate—manufacturers add services without ripping systems, SMEs sell into big programmes as specialised complements, and regulators and customers get credible, multi-metric proofs. That is how Kinaxis turned a planning product into an ecosystem play—and why roles, rules and data, not just algorithms, explain the performance gains.

C2.7. Exnaton: Enabling SMEs to Orchestrate Energy Communities

Running a local “energy community” used to demand bespoke IT and heroic spreadsheets. Exnaton, an ETH-spun Zurich SaaS, replaces that with a white-label platform any midsize utility or cooperative can deploy in weeks. It ingests 15-minute smart-meter data, matches surplus rooftop solar to nearby demand, prices trades under national rules, and issues clear, app-based bills—complete with CO₂ savings and peer-to-peer history. Residents sell excess power to neighbours at below-retail yet above feed-in rates; the utility earns a facilitation fee and keeps the grid stable.

The ecosystem design is simple and disciplined. Exnaton supplies the *rails*—API-ready billing, real-time settlement, user analytics—while utilities and energy cooperatives act as local orchestrators: they brand the app, set compliant price corridors and market the service. Prosumers connect inverters, EV chargers and batteries via standard protocols; households join with their existing retail account. Transparent governance (every trade and fee is logged and visualised) builds trust where utility-led innovation often struggles. Under the hood, microservices expose core tasks (data ingestion, tariff calculation, reporting) through documented APIs, and a Developer Portal lets complements plug in EV smart-charging, demand-response bids or gamified carbon dashboards without rewiring the core.

Traction is tangible. In three years the platform rolled out at 15+ utilities across Switzerland, Germany and Austria. A Bavarian pilot lifted community-traded solar from 20% to 55% of local demand in six months, cutting household bills by ~12% and shaving grid-export peaks; a Swiss cooperative layered dynamic EV tariffs so chargers preferentially absorb neighbourhood surplus and flatten evening spikes. Funding from Venture Kick, True Ventures and TotalEnergies’ accelerator, and regulatory tailwinds (e.g., Germany’s §14a revisions; Austria’s Renewable Expansion Act), signal a widening addressable market.

The strategic takeaway: orchestration is no longer the privilege of giants. By externalising complexity and meeting compliance out of the box, Exnaton lets “pipes-and-meters” SMEs curate credible local ecosystems that convert commodity electrons into loyalty-rich community services—while giving complementors and policymakers a neutral, modular operating layer to scale citizen energy, one neighbourhood at a time.

C2.8. Katana MRP: Plug-in Ecosystem for Manufacturing SME

When a craft chocolate maker or niche furniture studio tries to scale, the bottleneck is rarely marketing—it is the messy tangle of spreadsheets that sit between online orders and the shop floor. Katana MRP spotted this pain and built an ERP designed expressly for manufacturers with fewer than one hundred employees, bundling bill-of-materials, shop-floor control, and real-time inventory in a browser tab. Yet its masterstroke was not the core product; it was the decision to open the platform with Shopify-style APIs and a public Integration Marketplace so that outside tools could snap in as effortlessly as plugs into a socket.

The strategy came into its own once Katana shipped native commerce and finance connectors, making it an ecosystem enabler. A direct Shopify link lets D2C brands import SKUs, allocate stock by channel and push availability back to storefronts without CSV gymnastics; accounting integrations (Xero, QuickBooks and regional ledgers) automate COGS and payables; shipping partners (e.g., ShipStation/EasyPost) turn pick lists into labels with tracking. Add-ons handle barcode scanning, serial numbers for regulated devices, QC checkpoints and vendor lead-time forecasting. Each integration widens Katana's total addressable workflow while giving partners low-friction access to manufacturers who need automation but not heavyweight IT. By 2025, more than 1,500 companies across textiles, specialty foods, woodworking and hobby electronics were running on the stack, processing billions in order value that previously required enterprise ERPs.

The ecosystem design is deliberate. Katana keeps a neutral core with versioned APIs and a shared data model so complements build apps against a single source of truth rather than one-off integrations that decay. System integrators publish vertical templates; niche ISVs sell micro-apps (from batch genealogy to carbon footprints); implementation partners earn annuity fees on migrations and scanner roll-outs. Roles and rules are legible (documented interfaces, marketplace standards, transparent roadmaps) so value is shared without a hyperscale gatekeeper swallowing the margin.

Managerially, the case demonstrates how roles, rules and data enable SME participation at scale: plug into demand (Shopify and other channels), keep finance and logistics portable, and let specialist complements own narrow problems while orchestrating the flow in one cockpit. The result is a midsize SaaS from Tallinn turning isolated workshops into digitally fluent supply-chain nodes—proof that credible ecosystem leverage is attainable well below Big Tech scale.

C2.9 – AI as an Ecosystem Accelerator: New Ways to Join, Shape, and Orchestrate

One of the most important changes in today's world is the rapid growth of AI. While nobody really knows how the future will unfold, where there will be bubbles to burst and where impact of AI will be felt more acutely, we do know that significant changes are afoot. One of these changes relate to the way in which AI and ecosystems intertwine- and we expect ecosystems to be ever more important looking ahead.

In what ways will AI matter for ecosystems? In short, AI will transform how we organize, and redraw the opportunity landscape. It is important to remember that AI doesn't just add features; it changes who can participate and how coordination happens. In practice, three effects matter.

1) AI collapses search and coordination costs, widening the funnel of complements.

Matching the “right” partner has always been expensive: scanning markets, translating needs, testing fit. Machine-learning catalogues and LLM-based copilots now automate much of that grunt work. The result is more—and smaller—actors able to join. Qurator shows this dynamic in sustainability tech: its ML engine surfaces hundreds of niche vendors that would never make a corporate short-list, letting buyers trial options in days rather than months. Lower matching costs don’t merely speed deals; they create ecosystem variety, which raises user value and pulls in still more complements (see case study).

2) AI assistants become the ecosystem’s new “front door”.

Conversational agents standardise how users express intent (“renew policy”, “reprice inventory”, “find a charger”), turning intent into APIs that any qualified partner can answer. KBC’s Kate embodies this shift: partners plug skills into Kate’s intent catalogue; Kate Coin then rewards behaviour across banking, insurance, and merchant services. That design lowers a partner’s go-to-market friction and lets KBC orchestrate without owning every product—AI makes the rails intelligible and fair, so participation scales (see case study).

3) AI unlocks neutral data moats that sustain keystones.

Ecosystems thrive when a credible actor turns messy, multi-party data into shared, low-friction boundary resources—metrics, forecasts, attributions—that everyone trusts. AI/ML makes those resources accurate and timely. Consider the case studies we look at in this report:

MAF/SHARE uses models to compute “halo” contribution so tenants are paid for ecosystem value, not just their own till—partners stay because the rules feel legitimate.

Kinaxis lets SMEs publish AI apps (risk, carbon, warranty) against a single truth set; the core remains neutral while variety explodes through complements.

Exnaton transforms interval meter data into fair settlement and optimisation for citizen energy communities, letting small utilities orchestrate credible local webs.

Across these effects, AI creates new roles: the *intent router* (assistant owner), the *assurance layer* (safety/quality validators), and the *scenario lab* (simulation providers). For SMEs, that means more ways in—as a plugin answering an intent, as a model-powered optimizer riding a neutral data layer, or as an assurance service certifying other people’s AI.

The managerial takeaway is simple: treat AI as infrastructure for participation, not just as a product feature. Where you already have cases that lean on AI—KBC’s Kate, Kinaxis partner apps, MAF’s attribution models, Qurator’s ML scout, Busroot’s analytics roadmap, Exnaton’s optimisation—you can point to concrete mechanisms: *AI lowers entry barriers; AI makes rules legible; AI multiplies credible roles for smaller players*. That is how AI converts abstract “ecosystem talk” into repeatable growth paths (see case studies cited above).

Part 3: From Goliath to David: Living Under the Shadow of Big Tech

Chapter 3.1 – Why Would You Want to Mess with the Big Tech Ecosystems?

Working with hyperscalers is often unavoidable: they concentrate user attention, set de facto standards, and operate identity, payments, and discovery rails that smaller firms cannot cost-effectively replicate. The prize for engaging is threefold: distribution (instant reach), capability access (APIs, SDKs, data services) and credibility (borrowed trust from an incumbent interface). But there are equally strong reasons to challenge or constrain dependence: economic surplus tends to migrate to the hub (fees, ranking control, data capture), strategic options narrow, and product roadmaps become vulnerable to unilateral policy shifts. These are not hypotheticals; they are structural outcomes of the forces outlined in 3.2—network effects, control of touchpoints and data, multihoming frictions, and algorithmic curation. As we can see in the case of Amazon and X Fire, the temptation of leveraging growth is hard to say no to- and can help expand the boundaries of geographically constrained firms with a scalable product.

The practical question is not “engage or abstain,” but how to engage on your terms. The choice mirrors the Join / Shape / Orchestrate logic in Part 2:

- Join to learn quickly and monetise a module where the hub provides the cheapest route to customers.
- Shape when your capability (or licence) can be embedded in standards so that switching costs favour you over time.
- Orchestrate only in domains where you control a distinctive user relationship, critical data, or trust asset that others will credibly follow.

A disciplined entry test for Catalan SMEs: Necessity (do we need the hub to reach demand?), Differentiation (is our module hard to replace?), Bargaining power (what asset makes us non-fungible: data, licence, locality, brand?), and Reversibility (can we multi-home and exit with our data and users intact?). If any answer trends to “no,” design for participation with abstraction layers and optionality first, not lock-in.

Chapter 3.2 – When and Why Power Concentrates

Digital ecosystems promise fluid collaboration, but in consumer-facing arenas they frequently drift toward winner-takes-most outcomes. Four structural forces drive this concentration. First are strong network effects. Each new user makes the service more valuable for the next, raising the payoff from scale in a non-linear way. On a social network the value of connecting with friends multiplies as the graph fills in; on a marketplace, more buyers attract more sellers who in turn attract more buyers. Once a platform clears a critical mass threshold, rivals must offer an impossibly superior proposition to tempt users away, especially when the incumbent continually reinvests scale rents into better features.

Second is control of data and touchpoints. Consumer interactions create granular behavioural trails that power personalisation engines. The more data the orchestrator harvests, the sharper its recommendations, the higher its conversion rates, and the more data it collects—a self-reinforcing feedback loop. The orchestrator also owns the portals through which end users

engage: the app launcher, the voice assistant, the default search bar. Competing services must negotiate for visibility or pay for access to the very data they help generate.

Third, multi-homing frictions keep consumers and complementors anchored. Technical hurdles—unique APIs, identity tokens, proprietary file formats—raise switching costs. Contractual clauses restrict price differentials or prohibit parallel distribution. Behavioural inertia matters too: users resist abandoning familiar interfaces and accumulated reputation scores. Together these frictions create a “data and habit crust” that locks the ecosystem in place even when nominally open standards exist.

Fourth, algorithmic curation amplifies incumbency. Recommendation engines rank products, apps, or posts based on engagement metrics biased toward already popular options. Visibility begets clicks, clicks beget higher ranking, and a handful of winners snowball while the long tail languishes. Because those algorithms are opaque and proprietary, challengers cannot easily diagnose under-performance or optimise against the black box.

Hyperscalers leverage the interplay of these forces through three levers of dominance. Interface control comes first: by defining APIs, data schemas, and payment rails, Apple dictates how developers must build and monetise iOS apps; Amazon sets the bargaining terrain for marketplace sellers and fulfilment partners. Control over the rules lets the hub change commissions, surface its own substitutes, or throttle functionality with little notice. Complementors therefore face a lopsided dependency: they need the platform more than it needs any individual of them.

Next is reputation shaping. Platforms curate star ratings, feature lists, and search placement that steer user attention. A single tweak to ranking logic can swing millions in sales from one vendor to another. Because the orchestrator both operates the market and competes within it, its own offerings can quietly rise while rivals slip, all under the guise of “user relevance.”

Finally, switching costs solidify the “value-appropriation gravity well.” A consumer who has invested in app purchases, photo libraries, loyalty points, or smart-home devices faces real economic loss if she migrates. Developers tied to SDKs and in-app purchase mandates would need to rebuild back-end code and re-acquire users elsewhere. At scale these micro frictions aggregate into macro immobility, allowing the orchestrator to capture outsized margins without losing participants.

Power concentrates most acutely when three conditions coincide. The first is high indirect network effects—multiple sides of the market mutually reinforce adoption, as in app stores linking consumers and developers or ad networks linking advertisers and publishers. The second is data-rich engagement loops that personalise and thus differentiate the experience faster than competitors can imitate. The third is closed or mutable interfaces that gatekeep complementary innovation and extract rents through technical or policy edicts.

These dynamics do not doom every ecosystem to monopoly, but they skew the odds. Multi-homing can blunt scale advantages if standards remain open and switching costs stay low; differentiated verticals can sustain several coexisting hubs; regulators can mandate data portability or payments choice. Absent such counterweights, however, customer stickiness, algorithmic feedback, and governance asymmetry will pull surplus toward the centre as surely as gravity—leaving hyperscalers with disproportionate power over prices, innovation pathways, and ultimately the distribution of economic value.

Chapter 3.3 – Key Risks for Complementors and how EU regulation is trying to help

Complementors thrive when the platform they build on is predictable, contestable and fair. They suffer when rules change without notice, when the platform self-preferentially owns the shelf, when data or distribution are held hostage, or when ranking volatility makes demand fragile. AI intensifies each of these: assistants become new front doors; tool registries and API eligibility become the new chokepoints; and learning loops around usage data make the party that controls context hard to dislodge. The European regulatory package—**DMA, DSA, and the AI Act**—is meant to restore some balance by disciplining gatekeepers, forcing transparency and portability, and setting minimum safety and interoperability standards for AI. What follows is a pragmatic view of what those laws *do* for complementors, where enforcement has bitten, and what frictions remain.

The **Digital Markets Act (DMA)** targets the structural bargaining gap between very large “gatekeepers” and those who depend on them. Six firms—Alphabet, Amazon, Apple, ByteDance, Meta and Microsoft—have been designated, across 22 core platform services. That designation matters because the rulebook then bans self-preferencing in rankings, prevents anti-steering clauses, requires access to key interfaces and data on reasonable terms, and pushes open app-distribution pathways (choice screens, alternative app stores or payment routes). In 2025 we saw the first real teeth: the Commission imposed the first DMA fines (Apple and Meta) and issued preliminary findings against Alphabet and Apple in separate probes. That enforcement wave tells complementors two things: (i) steering and self-preferencing are not just “bad form”—they are illegal under the DMA; (ii) there is now a route to redress if “compliance” arrives in name only. Yet it also shows the limits: gatekeepers can try to comply while re-optimising fees and frictions (e.g., revised Store fees or “core technology” charges), so the *letter* of the DMA still needs continual policing to achieve its *spirit*. It also remains to be seen whether Big Tech doesn’t simply consider this to be a cost for doing business, emboldened by the support of the Trump administration that has turned everything into a geopolitical race.

The **Digital Services Act (DSA)** focuses on *system risk and transparency* for very large online platforms. For complementors, the value is indirect but real: clearer ad transparency and recommender-system duties make discovery less arbitrary and reduce dark-pattern “taxes” on user trust; data-access for vetted researchers can surface systemic bias; and “new feature” risk assessments constrain sudden design shifts that externalise harms. Importantly, 2024–25 enforcement has already changed behaviour: the Commission opened formal proceedings against X; and TikTok Lite’s “rewards” feature was withdrawn EU-wide after DSA pressure, setting a precedent that risky engagement mechanics can be stopped before they become the new normal. Current findings against Meta and TikTok point to a tougher line on minors’ protection and ad labelling—further signals that discovery and demand capture are no longer a cost-free optimisation exercise for the largest players. These steps won’t guarantee fair reach for every complement, but they *do* set guardrails on the most distortive tactics.

The **AI Act** recognises that the next set of chokepoints lives in assistants, orchestration layers and general-purpose models. Timelines matter: the Act entered into force in 2024; obligations for general-purpose AI (GPAI) providers began to apply on 2 August 2025, with extra duties for models deemed to pose *systemic risk* (e.g., evaluations, incident reporting, security measures and documentation). The Commission has also issued GPAI guidance and a Code of Practice

to reduce ambiguity for model providers and those building on them. For complementors, the practical upside is twofold: (i) a baseline of transparency and safety obligations on the models and assistants you depend on; (ii) a framework to insist on documentation and interface clarity when AI platforms become the new gatekeepers. The open question is pace and specificity: guidance is still evolving, systemic-risk designations will matter greatly, and interoperability at the *tool-API layer* remains more a market outcome than a legal mandate. There is considerable pushback as some of the regulations are seen as onerous, and as the EU worries about missing out in the AI race and its geopolitical standing.

What does this add up to for a firm building on others' rails? First, risks remain real: self-preferencing can move from search pages to assistant suggestions; ranking opacity can migrate from feeds to prompts; "fair" access can be eroded by pricing, latency budgets or eligibility criteria; and in AI, a platform can comply on paper but retain decisive behavioural levers (context windows, tool selection, safety filters) that tilt outcomes. Second, there is new leverage. The DMA supplies hard obligations you can reference in negotiations (anti-steering, access and ranking), with live precedents and fines; the DSA curbs the worst excesses around attention capture and gives regulators purchase over design choices that can arbitrarily punish complements; and the AI Act sets a compliance spine around models so "black-box" excuses carry less weight. Third, enforcement is happening, not just promised. Where the Commission has moved—DMA fines and statements of objections; DSA proceedings and binding commitments—it has already forced product and policy changes. The practical step for complementors is to *instrument and document harm* (steering blocks, ranking drops tied to policy changes, API access limits), then escalate with references to the relevant Article numbers; the law is finally written with your situation in mind.

There are, however, issues to watch. Gatekeepers will keep "complying while re-optimising"—for example, by restructuring fee stacks or redefining technical eligibility; these tactics are slower to adjudicate than a clear-cut anti-steering clause. In content governance, the DSA's push for safer defaults can be read as chilling or over-broad if enforcement leans too heavily on headline pressure rather than measured risk analysis. And in AI, the line between proportionate GPAI duties and barriers to European open-source or SME innovation is still being drawn; codes of practice reduce uncertainty, but they are not (yet) a safe harbour, and technical standards are catching up with the law. In short: the *direction* of travel favours contestability, but outcomes will depend on everyday implementation choices—by platforms, by regulators, and by complementors who use the new rules in contracts and escalation.

Two final points anchor this section to your broader argument. First, policy is a guardrail, not a substitute for strategy. Complementors should still multihome, keep their own CRM and measurement, and build an abstraction layer so that no one assistant, app store or model is an existential single point of failure (see Case C3.2, Shopify; 3.4 on AI & ecosystems). The regulations make that strategy more viable; they don't replace it. Second, Europe is adding portability to the plumbing. Although outside this chapter's core trio, the Data Act's *cloud-switching* duties (applicable from September 2025) backstop "design for exit" in contracts and public tenders, reinforcing the same optionality logic you promote elsewhere (see C3.3 on GAIA-X/Eurostack). Taken together, this is not a wall against global platforms; it is a disciplined attempt to keep choice, transparency and exit alive inside the ecosystems where complementors must operate.

Chapter 3.4 – AI and Ecosystems: where power shifts, and what to do about it

One of the most important changes in the world around us, though is the growth of AI. The question becomes, how does AI relate to ecosystems? The answer is, it increases their importance in a number of ways. First, AI itself is growing through ever-more-concentrated ecosystems. Big Tech firms are finding ways to absorb, ally with, engage and link with the new AI upstarts, as the figure below shows.

Companies don't provide independent goods and services; the AI segment is set in a number of partly overlapping ecosystems. Firms like Nvidia (see C3.3) have grown as a result of their ecosystem play on several levels. They have also created new ecosystems for those who are involved in AI – not just the major players, where Nvidia is aggressively creating webs of influence but also programmers and smaller organizations who could be part and parcel of the productive fabric of a highly trained, innovation-focused region such as Catalunya.

Furthermore, as we discussed in section 1.5, AI collapses search and coordination costs, widening the funnel of complements; more important, AI assistants become the ecosystem's new “front door” and they unlock data moats that sustain keystones. They amplify the opportunities firms like MAF or KBC are drawing on, and make it possible for them to develop new ecosystems. From a complementors perspective this also means that the opportunity space to connect grows, and with it grows the need to have a clearer ecosystem strategy.

There are some broader implications of the changes that AI induces, that we should expect will make ecosystems be even more important than before. AI does not simply automate tasks inside existing structures; it reshapes the structures themselves, and much of the action is in inter-dependent ecosystems that move beyond traditional sectoral boundaries. In ecosystems, power has typically accrued to hubs that control demand aggregation, access, and rules. AI introduces new control points—the user-facing assistant/interface, the orchestration layer that coordinates services across firms, and the data/feedback loops that compound learning. Control of any one of these allows a player to tilt the field without owning every asset in the stack (see Choudary, 2025).

The first shift is at the interface. If recommendations and actions flow through an assistant, the *default* becomes decisive: what the assistant proposes is what many customers accept. This is not just convenience; it is economics. We all tend to be much less active in choosing that theory would predict. As a result, we go for the default, allowing firm (dubbed “choice architects”) to benefit. In digital ecosystems, whoever sets the default prompt, ranking, or fulfilment path can reassign surplus—much as a dominant marketplace did with search placement. AI assistants compress that power into fewer clicks, sometimes into zero-click journeys, and this opens up a significant question of power and leverage.

The second shift is the rise of agentic orchestration. Instead of users stitching together steps across providers, AI agents call functions and APIs to achieve outcomes. This creates a new kind of “platform”: one that composes multiple complements at run-time. The orchestrator who controls which tools are callable, on what terms, with what telemetry, can favour some complements and degrade others—without ever appearing to “compete on the shelf.” The rules of participation thus migrate upward, from storefronts and app stores to assistant registries, tool schemas, and data contracts.

The third shift is about data. Traditional moats came from user scale and the network externalities that strengthen the orchestrators. AI moats emerge from task-specific feedback

and private contextual data. A smaller firm can hold a defensible position if it owns scarce signals (domain expertise, verified quality data, local trust channels) and can feed them into models safely. Conversely, complements that cannot take their own derivative learning with them become fungible and price-taker. The practical question is no longer whether to “join or abstain,” but how to join with optionality, securing rights to your data exhaust and proof of value creation.

These shifts raise both risks and opportunities. Risks echo earlier themes: hold-up through opaque ranking in assistants; policy drift when orchestration rules change; subtle self-preferencing via tool eligibility or context windows. Opportunities are equally real: *merchant-first rails* (e.g., infrastructure models that preserve brand, data, and multihoming) help smaller players shape rather than merely join big hubs; data spaces and standards reduce the cost of safe sharing; open interfaces let complements compose into outcomes without surrendering their crown-jewel data. (See Case C3.2 for an infrastructure-first counterweight and §3.4 for defensive moves.)

For Catalan SMEs, the practical playbook is consistent with the survival strategies we set out earlier, but AI both leads to new ecosystem options and intensifies the need for discipline for an ecosystem play. Build an abstraction layer that separates your product from any one assistant or model; maintain a side-channel to customers (support, community, CRM) so assistant-ranking shocks are survivable; and insist on data-rights schedules that protect derivative learning. Where you bring non-substitutable assets—regulated licences, local trust, proprietary datasets—tie them to measured P&L lift so removal is visibly value-destroying. For complements that cross assistant ecosystems, instrument eligibility and latency: when behaviour changes, you need evidence to escalate—or exit.

For the Generalitat, policy can be a springboard, not just a brake. Procurement should favour interoperable, merchant-first architectures; EDIHs and Catalonia Data Spaces can underwrite safe data-sharing; standard contract riders (API versioning, portability, audit trails on ranking in assistants) can be published for SMEs to append to platform agreements. The lesson of AI in ecosystems is simple but non-negotiable: ecosystems are built, not born. If the region wants to translate scientific and entrepreneurial strength into advantage, it must invest in the connective tissue that makes equitable orchestration possible—and in the governance that keeps it so.

Chapter 3.5 – Survival Strategies and Defensive Moves

Ecosystems create growth—but they also shift power toward whoever controls the front door and the rails (identity, data, ranking, rules). With AI assistants and orchestration layers mediating more journeys, the practical question is how a firm stays valuable and hard to sideline. The cases in this report point to a small set of moves that work in the wild. The aim is not to “beat” a hub, but to join on terms that preserve options, make your contribution visible, and keep you worth partnering with.

Own something others can’t copy fast. Pick a moat that matters to the ecosystem: regulated licences (health, finance), a high-trust local network (garages, clinics, integrators), or a scarce dataset tied to outcomes. KBC’s assistant is powerful because it sits on consented transaction signals; MAF’s retail media flies because it joins footfall, purchase and location. If removal of your asset obviously damages conversion, reliability or unit economics, you negotiate from strength.

Keep a direct line to your user. Even if discovery happens on a platform or through an assistant, maintain a channel you control—support communities, CRM, education content, newsletters. MAF turned identity (SHARE) into a durable spine; KBC did the same with Kate. A portable audience is your shock absorber when a hub changes rules.

Be easy to plug in—and ready to replace. The paradox of staying vital is designing for exit. Publish or adopt clean APIs; separate your product from any single model, app store or assistant; version changes and give partners deprecation windows. Katana thrives because complements build against a single source of truth; Kinaxis scaled by making scenarios legible across suppliers and logistics. If switching you out would strand users, hubs won't risk it; if switching is genuinely possible, more partners will join—raising the ecosystem's value and your relevance within it.

Instrument your dependency. Treat platforms like suppliers: monitor eligibility flags, ranking shifts, latency budgets and break-glass thresholds. When something changes, you need facts, not hunches. HDFC ERGO's shift to orchestrated, solved tasks worked because the data layer and provider registry were built first; without instrumentation, the experience would have felt bolted-on.

Make your contribution provable. Ecosystems endure when value splitting feels fair. Adopt closed-loop measurement that shows what your presence did to the scoreboard people care about: solved tasks, sales lift, time-to-resolution, retention, reduced loss or waste. MAF's incremental "halo" measurement keeps brands and tenants investing; Exnaton's transparent settlement keeps prosumers and utilities aligned. If you can prove uplift, you earn more latitude on placement, pricing and roadmap.

Multihome where it counts. Maintain at least one alternative route to market (another marketplace, partner, or direct channel) and one alternative route to fulfilment (second logistics or service network). Shopify's complementors survive turbulence because they also run D2C or additional marketplaces; POSB's community rails work because they sit on national standards (PayNow, SGQR, MyInfo) rather than a proprietary fork.

Use coalitions to close bargaining gaps. If your category faces a single dominant hub, act with peers. Shared taxonomies, data spaces and reference contracts lower onboarding costs and make fair access harder to deny. Energy communities (Exnaton) and automotive data spaces show how neutral artefacts crowd in complements without creating a new gatekeeper.

Price in optionality. Reserve a small, permanent slice of engineering for adapters, second SDKs, and test beds. It looks like overhead—until it is why you keep shipping when a hub changes policy. The SMEs that rode out commerce and advertising shocks were the ones with connectors ready.

Design incentives that reward verified behaviour. Programmable benefits tied to outcomes (safer driving, greener tariffs, on-time bills, loyalty to local merchants) pull more actors into the same loop. InSoil aligned farmers, verifiers and buyers by paying from measured sequestration; KBC uses benefits to make partners compete on lift, not coupons. When rewards are priced to proof, small players can win alongside giants.

Let AI lower costs to participate—not raise lock-in. Use assistants and recommender systems to translate messy intents into callable actions, but keep tool registries open and ranking criteria inspectable. The winning pattern we see is "AI as connective tissue": routing,

anomaly detection, prediction, and first-pass drafting that reduce search and coordination costs so **more** firms can credibly join the same journey.

Know which rulebook helps you. Contestability, transparency and portability obligations are finally catching up with assistant-era gatekeeping. Use them to secure clarity on ranking, interfaces and data rights when you escalate or contract—but don't outsource strategy to regulation. Your day-to-day leverage still comes from the moves above.

A simple workflow to apply tomorrow. Start from one customer job you keep losing (or fear losing) to a hub. Map the minimum partners needed to solve it end-to-end. Stand up the rails before the campaign: identity/consent, an event stream for intents, a thin orchestration layer that calls only “quote, schedule, fulfil.” Ship it to a narrow segment with a single success metric (e.g., churn avoided, minutes saved, verified lift). Publish the rules to partners, log changes, and expand only when the uplift is repeatable. That is exactly how the successful cases in this report moved from slideware to habit—and why they stayed relevant as the front door shifted from menus to models.

Chapter 3.6 – The Role of Policy – Hope and Limitations

Regulators have begun to pull hard on the few levers capable of counter-acting the winner-takes-most forces detailed earlier. Europe's Digital Markets Act stands at the centre of that effort. By designating “gatekeepers” and imposing bright-line duties—no self-preferencing, no compulsory in-app payment monopolies, no dark-pattern consent harvesting—the DMA tries to dull the blades that hyperscalers use to carve out unfair advantage. Equally important are the access mandates embedded in the text: business users may port data freely, edge devices must expose core hardware and software interfaces to rival service providers, and advertisers gain audit rights over platform metrics. These provisions recognise that without structural openness, data gravity and default positions will continue to cement incumbency.

Interoperability mandates push that logic even further. Messaging services must interconnect on basic text, image, and voice features; app stores must allow alternative catalogue fronts; cloud providers are pressed to publish migration toolkits for workload portability. Where negotiated interfaces fail, public authorities reserve the right to set technical standards directly. Transparency rules round out the toolbox: ranking algorithms, ad-delivery logic, and review moderation policies must be explainable in plain language, logged for audit, and open to third-party testing. Sunlight does not eliminate bias, but it raises its cost and gives rivals factual grounds on which to challenge discrimination.

Yet optimism should be tempered by the realities of enforcement. Regulations are imperfect, and while the EU has shown a globally unusual appetite to consider the impact of ecosystem power (e.g. in blocking the proposed merger between Booking.com and eTraveli) regulation has to be tested by courts and considered at the level of 27 different, fragmented jurisdictions in the EU. Pushback will not only be legal- but also political and perceptual. Gatekeepers lobby intensively, throwing financial and technical complexity into every consultation. Implementation guidance, codes of conduct, and delegated acts must wend their way through crowded regulatory calendars, and each delay preserves the status quo. Even after provisions enter force, small firms often struggle to recognise a violation, let alone finance litigation. A data-porting right is only as real as the legal budget available to compel compliance; an antitrust victory that arrives three years late seldom restores a closed business.

Recognising these limits, European policy makers have paired sticks with scaffolding for indigenous players. The EU's network of European Digital Innovation Hubs offers cloud credits, sandboxes, and advisory hours so that small manufacturers, agritech outfits, and creative studios can plug confidently into open standards rather than defaulting to a hyperscaler's proprietary stack. Horizon Europe, Digital Europe, and member-state recovery funds funnel grants into open-source middleware, privacy-preserving analytics, and edge-computing tool-chains that give domestic ecosystems a head start. Public procurement rules are increasingly written to favour solutions that expose interoperable APIs and guarantee data sovereignty, turning the state into a reference customer for modular alternatives.

Private-sector coalitions amplify this momentum. GAIA-X seeks to codify a framework of federated cloud services where identity, policy, and data-exchange protocols are agreed up-front, allowing multiple providers to appear as one logical infrastructure. Critics justly note slow progress and governance sprawl, yet the initiative signals that European industry no longer accepts dependency on a handful of foreign hyperscalers as inevitable. The emerging Eurostack coalition goes a step further, drafting a fully open runtime layer for AI and data applications that can sit atop any compliant cloud. Success is by no means assured, but the mere existence of these projects shapes behaviour: incumbents pre-emptively offer more open licensing and local data-residency options to head off regulatory or market push-back.

What, then, can policy realistically achieve? It cannot, on its own, reverse network effects or rewrite the economics of data-driven learning. But it can tilt incentives. By forcing gatekeepers to interoperate, to publish interfaces, and to justify ranking choices, it reduces the artificial frictions that keep users and complementors captive. By funding neutral infrastructure and mandating open standards in public tenders, it lowers entry costs for challengers and shrinks the distance between prototype and scale. The DMA and related measures thus act as guardrails and springboards simultaneously: restraining abuses while enabling local innovation ecosystems to form around alternative value propositions.

Still, hope must stay tethered to pragmatism. Rule-books grow stale; enforcement budgets lag; lobbying never sleeps. Policymakers need continuous market monitoring, rapid remedial powers, and alliances with civil-society watchdogs that can surface evidence in real time. SMEs require legal aid pools and collective redress mechanisms to make their rights tangible. And European initiatives like GAIA-X must resist bureaucratic drag, focusing relentlessly on shipping usable code and reference deployments.

In short, policy can slow the gravitational pull of digital mega-platforms and nurture space for regional and sector-specific ecosystems, but only if rules remain enforceable, support structures stay funded, and public institutions evolve as fast as the technology they aim to tame.

Chapter 3.7 – What does this Mean for SMEs?

Ecosystems can lower your cost to acquire customers and widen your reach—but only if you join on terms that keep you visible, portable and hard to replace. The winning SMEs in our examples did three simple things well: they solved one clear job end-to-end, they made it easy for partners to plug in, and they proved their contribution with numbers that others trusted.

Start with role clarity. Decide the specific job you solve in someone's day (repair the cracked screen; get a small loan linked to a verified outcome; match surplus energy to local demand). Everything else—partners, data, pricing—follows from that choice. Make joining you simple: expose one or two endpoints that partners can call (“quote”, “schedule”, “fulfil”), keep your data formats tidy, and publish the basic rules of engagement so nobody needs a bespoke deal to work with you. Then make your value provable: measure the lift you create (time saved, sales converted, losses avoided) and share those numbers with partners so they keep investing.

A 90-day plan that works in practice:

- Days 1–10: pick one job you can own; write down the minimum set of steps to solve it; name the single metric that proves success (minutes saved, churn avoided, verified uplift).
- Days 11–30: line up one anchor partner and one back-up; wire the basics—identity/consent, a small event stream so you can spot the moment to act, and the two endpoints partners need to call you.
- Days 31–60: run a narrow pilot for a narrow segment; ship fixes weekly; instrument everything (eligibility, latency, completion, the single success metric).
- Days 61–90: publish the rules to partners (what makes you rank, what service levels you honour, how you count outcomes); add one second channel so you're not captive; turn the pilot report into a one-pager you can reuse to recruit more partners.

Different SME types can use the same logic:

- If you make or sell products, plug into the channels your customers already use, keep finance and shipping portable, and capture uplift at the SKU or basket level so you can prove why you deserve premium placement.
- If you provide a specialist service (repairs, health, logistics), expose “quote/schedule/fulfil” and insist on clear, outcome-based reporting so you're rewarded for solved tasks, not just leads.
- If you build tech or analytics, be a clean plugin that answers a well-defined request; avoid dependence on a single platform; negotiate your rights to the learning you generate; and keep adapters ready for a second channel.

A few disciplines make the difference. Keep a direct line to your users (support community, newsletter, help centre) even if discovery happens elsewhere. Reserve a small, permanent slice of engineering for adapters and tests so a rule change upstream doesn't stop you shipping. Tie any rewards or discounts you offer to verified behaviour, not broad segments, so smaller partners—and you—can afford to play. And treat your data as a trust asset: clear consent, visible controls, and audit trails that make partners comfortable sharing theirs.

If you cannot commit to a narrow pilot, a single success metric and weekly iteration, consider dropping the section rather than keeping generalities. If you can, this section earns its keep by giving teams something they can run next quarter and expand the quarter after.

Cases Chapter 3: Competing alongside big tech / building new services

C3.1. The Dilemmas of Smaller Organizations Working with BigTech: Life Under Amazon's Shadow for Airplay

X Fire Paintball & Airsoft, a small U.S. retailer, joined Amazon Marketplace and used Fulfillment by Amazon (FBA) to reach customers far beyond its two local stores. Within a few years, online sales—largely via Amazon—accounted for the vast majority of revenue. The economics looked attractive: Amazon brought traffic, handled logistics, and helped X Fire win the Buy Box; the 15% commission felt like a fair price for growth.

The platform upside, however, came with escalating dependence. As X Fire's listings gained traction, Amazon began sourcing the same SKUs directly from suppliers and selling them itself—often advantaged in the Buy Box and sometimes pricing below manufacturers' minimum advertised price. For a small complementor, the host platform became a powerful competitor with superior data, bargaining leverage, and control over discovery. X Fire's near-term gains thus increased its exposure to a gatekeeper that could appropriate demand, compress margins, or disintermediate the seller altogether.

X Fire experimented with defensive moves that many SMEs will recognize: pressing suppliers not to sell directly to Amazon, negotiating return-rights if they did, and trying to organize peer retailers—efforts that yielded partial, fragile wins and constant renegotiation after supplier mergers. The firm ultimately declined Amazon's invitation to expand to Amazon.ca, fearing it would again “seed” the market only to face direct entry once demand was proven. These experiences capture the core calculus for small complementors: platforms can be the fastest route to scale, but success on the platform can trigger the very dynamics that erode a complementor's position.

What this illustrates for the report is twofold. First, the short-term benefits of platform participation (reach, convenience, credibility) are real and often decisive for SMEs—but they are not free. Data access, self-preferencing, Buy-Box design, and MAP enforcement are architectural choices that tilt outcomes toward the orchestrator. Second, these patterns explain why policy debates increasingly focus on curbing predatory practices (e.g., using third-party seller data to compete, steering to own retail, or imposing parity clauses). The X Fire case shows how even capable, well-reviewed complementors face a strategic dependency they cannot bargain away—and why regulatory guardrails and pro-complementor ecosystem design matter if we want platforms that foster durable SME participation rather than transient, high-churn growth.

C3.2. Shopify as an Engine and an Opportunity

Shopify illustrates how an infrastructure model can act as a more equitable partner for SMEs than gatekeeper marketplaces. Rather than aggregating consumer demand into a single “shelf,” Shopify provides the stack for merchants to run their *own* store and then attach sales channels (Amazon, social, etc.) on their terms. That design preserves the merchant's brand and customer relationship while making multihoming operationally feasible (catalogue, inventory and orders remain central). The company's public stance is explicit: prioritise merchant sovereignty—“own, don't rent”.

This architecture mitigates several risks discussed in Section 3. Because Shopify does not run a consumer marketplace that competes on the same shelf, the classic self-preferencing conflict is far less structural than on vertically integrated marketplaces that both set rules and sell against third-party sellers. Evidence on self-preferencing and Buy-Box-type concerns in large marketplaces underscores the contrast and helps explain why concentrating discovery and rule-setting in one hub tends to shift surplus to the centre.

Control over data and relationships is equally important. In the default configuration, merchants act as data controllers for customer data processed via their store, with Shopify generally acting as a processor/service provider (subject to carve-outs in specific additional services). Practically, this means merchants can maintain privacy notices, honour access/erasure/portability requests, and export operational data without renegotiating platform terms—exactly the portability and reversibility conditions we recommend for SMEs.

Shopify's complementor ecosystem (apps, themes, agencies) also aligns incentives relatively well for outsiders to invest. In 2021, Shopify dropped its App Store commission to 0% on the first US\$1m/year (15% thereafter), catalysing third-party tool-building for SMEs; in 2025 it tightened this to a lifetime US\$1m cap—a reminder that even “friendly” ecosystems evolve and rules change. This suggests that the rails are more equitable than a gatekeeper marketplace, yet prudent firms still hedge with abstraction layers and optional channels.

This also suggests that firms like Shopify are rational, trying to take advantage of the opportunities left by greedier large firms like Amazon, but still need to balance the value proposition of being more equitable towards the smaller complementors with the desire or need to capture value from the collaborative relationship. Complementors should focus on doing careful due diligence in deciding whom they ally with, what their incentives and policies are, and pick the relationships that make sense for them- to avoid tricky outcomes such as what awaited Airplay and its Amazon adventure.

In practical terms for Catalan SMEs, Shopify functions as a control point: run core brand, checkout, identity, CRM and data pipelines on your site; syndicate to marketplaces and social channels as optional demand taps; and keep evidence of the value you create (conversion, attach-rate, retention) in your own analytics. That configuration reduces algorithmic and fee shocks from any one channel, and—because the app ecosystem is broad—lets firms assemble capability without bespoke builds. The trade-off is that governance still matters: privacy settings, app permissions and data exports need to be managed proactively, and firms should assume that commercial terms (including developer rev-share and some data features) will change over time.

This case study suggests ecosystems *can* be designed to distribute power fairly (merchant-first infrastructure, multihoming by design, developer-friendly economics) but they don't appear by accident. They require conscious architectural choices, stable rules and continual tuning of incentives—exactly the orchestration work public policy and regional convenors can encourage. Shopify's model provides a concrete reference architecture for SMEs that want to join large demand hubs without surrendering the assets that make them non-fungible.

C3.3. Nvidia and the growth of AI through ecosystems

Nvidia's rise from a gaming-GPU specialist to one of the world's most valuable companies is, at heart, an ecosystem play. Rather than compete only on chips, Nvidia built and then relentlessly deepened a *full-stack* platform—hardware, interconnects, systems, networking, and, crucially, the software layers that bind developers, partners and customers into a mutually reinforcing orbit. Its annual report describes this explicitly: the company “delivers a full-stack computing platform” anchored in CUDA, surrounded by CUDA-X domain libraries and Nvidia AI Enterprise software; it now serves a developer community counted in the millions.

The flywheel starts with CUDA. By giving developers a stable programming model and a rich set of tuned libraries (cuDNN, TensorRT and the broader CUDA-X stack), Nvidia made its GPUs the default target for modern AI and accelerated computing. That lowers time-to-value for complementors—startups, ISVs, systems integrators and enterprise teams—because so much heavy lifting (kernels, compilers, drivers and reference containers) is pre-optimised and published through the NGC catalogue. In practice, this means new models and applications can be stood up quickly on on-prem clusters or any public cloud that exposes Nvidia acceleration.

Nvidia then reinforced the core with system-level control points that make very large AI workloads tractable. NVLink and NVSwitch turn racks of GPUs into a *single logical accelerator*, while DGX/NVL systems arrive as engineered building blocks that hyperscalers and enterprises can stitch together at data-centre scale. With Blackwell, the GB200/GB300 NVL72 cabinets provide unified high-bandwidth memory space and terabytes-per-second of GPU-to-GPU fabric—delivering step-changes in training and inference efficiency. This is as much ecosystem design as it is silicon: the more developers target this fabric, the more valuable it becomes for clouds, OEMs and end-users—and the harder it is for rival stacks to dislodge.

Networking rounds out the stack. The Mellanox acquisition pulled low-latency interconnect, DPUs and InfiniBand/Ethernet fabrics inside the platform boundary, so clusters scale predictably and software can assume high-performance communications. Spectrum-X (Ethernet for AI) and Quantum-class InfiniBand give Nvidia credible “north–south and east–west” stories, which matters because AI workloads are now limited as much by networking and orchestration as by FLOPS.

On top of the infrastructure, Nvidia is standardising *how* enterprises consume AI. AI Enterprise is positioned as an “operating system” for enterprise AI, and NIM microservices expose pre-optimised model endpoints behind stable, industry-standard APIs. That lowers switching costs for customers building apps, while still anchoring them to Nvidia's inferencing toolchain (TensorRT(-LLM), CUDA-X) and to any Nvidia-accelerated environment they choose—on-prem, cloud or edge. In parallel, the AI Foundry offering combines DGX Cloud capacity with NeMo tooling to help enterprises customise models and then export them as NIMs. These layers convert raw compute into repeatable, monetisable software and services, pulling more complementors into the fold.

The developer and partner side of the ecosystem is equally deliberate. NGC's ready-to-run containers, vertical SDKs (robotics, healthcare, AV) and a vast startup programme (Inception) reduce go-to-market friction and create a pipeline of third-party innovations that implicitly privilege Nvidia-ready deployment. For corporate buyers, this shows up as shorter pilots, simpler procurement (reference architectures with OEMs), and a talent pool already trained on CUDA-first workflows. The result is classic positive feedback: more apps and know-how attract more customers, which justifies more platform investment, which attracts more developers.

This ecosystem design has real economy-wide consequences. Nvidia's data-centre business has exploded as generative AI moved from POCs to production. Record fiscal-year results (and repeated record quarters) are not only a demand story; they reflect how the platform converts demand into usable capacity at scale—DGX/NVL systems, tuned software, and networking delivered as a coherent whole. That coherence is why hyperscalers and large clouds have been willing to pre-commit to vast deployments of Nvidia-based infrastructure, and why we see marquee orders to support frontier-model training.

Yet the same design raises governance questions central to Chapter 3 of this report. By controlling multiple layers—from developer tooling to interconnects and networking—Nvidia concentrates bargaining power at the platform core. CUDA's ubiquity creates “gravitational pull” for complementors; NIM's standard APIs make application logic portable across *Nvidia-accelerated* venues, but do little to encourage cross-platform portability beyond the Nvidia stack. The risk for complementors is dependency: pricing power, allocation in scarcity, and roadmap choices (for example, favouring certain form factors or features) can ripple through the value chain. Policymakers, meanwhile, confront a single-vendor chokepoint in critical compute, with spillovers for national AI programmes and industrial strategy. None of this happens by accident; it is the predictable outcome of ecosystem architecture that solves real problems at scale while centralising control points.

There is, however, a flip side that matters for regions and SMEs. Precisely because Nvidia reduces integration friction and performance risk, smaller players can now participate in advanced AI without building everything themselves. Startups can ship on NGC containers; local systems integrators can stand up RAG/agentic apps on AI Enterprise/NIM; manufacturers and city planners can leverage Omniverse and OpenUSD to design “physical-AI” systems and digital twins that interoperate across toolchains. In other words, a tightly orchestrated platform can expand the *addressable* ecosystem even as it strengthens the orchestrator. The policy questions become: how to ensure access (capacity, pricing, skills), how to preserve credible alternatives (where appropriate), and how to keep bottlenecks from muting downstream innovation.

Nvidia showcases both halves of our Chapter 3 thesis. First, complex innovations (frontier AI) *need* ecosystems: without common stacks, tuned libraries, proven interconnects and reference systems, most enterprises and public bodies could not absorb or scale the technology. Second, those ecosystems don't assemble themselves; the orchestrator invests heavily, curates standards, and uses product and programme design (DGX, NVL, CUDA-X, NIM, Inception) to align incentives. The result is growth that looks “inevitable” in hindsight—but is the product of a decade-plus of architectural choices. For complementors and policymakers, the lesson is not to fight ecosystems but to negotiate their terms: push for open interfaces where feasible, build local capability on top of the dominant stacks, and keep optionality alive to avoid single-point dependency as the technology (and pricing power) shift.

C3.4. GAIA-X & Eurostack: Europe's Pushback on Platform Dominance

Europe's answer to hyperscaler dominance has shifted from the dream of a single “European cloud” to a more pragmatic playbook: rules, rails and procurement that aim to keep markets open, portable and contestable while staying interoperable with the global tech economy. The political spark this cycle was the EU's cloud-security label (EUCS). In April 2024 a revised draft removed the most stringent “sovereignty” clauses—those that would have required immunity

from extra-EU law for top-tier certification—making it easier for U.S. providers to qualify and pushing the sovereignty debate away from certification formalities toward industrial policy and public buying power.

Against that backdrop, **GAIA-X** has settled into its role as the **rules-and-interfaces layer** for trusted data sharing. Its 25.05 Architecture and Trust Framework set out how identity, credentials, compliance “labels” and policy rules travel with data across firms and borders, with live demonstrations (e.g., at Hannover Messe 2025) showing that this is more than a paper standard. In essence, GAIA-X reframes “sovereignty” as verifiable attestations and portable contracts, not ownership of every server.

The proof points are sector data-spaces. In automotive, Catena-X has become a reference model for cross-company data exchange on traceability and compliance—and in June 2025 the Automotive Industry Action Group (AIAG) was named the North America hub for Catena-X, suggesting the method can travel when it solves real problems (carbon and parts provenance, recall efficiency). That kind of portability—European rules used abroad—matters more than badges if the goal is bargaining power and true interoperability.

Under the covers, the IPCEI-CIS programme is the supply-side complement: a Commission-backed, multi-country effort to build a multi-provider cloud-to-edge continuum so buyers can source capacity from different European providers without rewriting everything each time. It is deliberately not a “national cloud”; it is plumbing that should make switching and locality feasible where sectors or states need them.

The demand-side lever is the **Eurostack** initiative—a coalition arguing that sovereignty also needs procurement muscle. Its 2025 framework proposes European-made options being scored and, where competitive, preferred in public tenders, plus clear criteria for what counts as a sovereign provider. Whether or not hard quotas materialise, the centre of gravity has visibly moved toward what public buyers specify and purchase, not just how they certify.

Finally, Europe has added legal teeth to **portability**. From 12 September 2025, the EU Data Act applies, including cloud-switching and interoperability duties for providers of data-processing services. That flips “design for exit” from a negotiation point into a default legal expectation that regions, agencies and SMEs can cite directly in RFPs and contracts.

What this buys Europe is leverage and options. Trustable sharing via GAIA-X lowers the cost of cross-company collaboration; sector data-spaces like Catena-X show that standards can become operating reality; IPCEI-CIS promises local, multi-provider capacity where needed; Eurostack tries to anchor demand; and the Data Act gives users enforceable switching rights. Together, these are not a wall against global platforms but a federated operating model that can reduce hold-up risk and keep more value with European firms when they participate in global ecosystems.

Open questions remain. Will GAIA-X labels and proofs become “buy signals” in large tenders, or stay as showcases? Will IPCEI-CIS components reach the performance, developer-experience and price points needed so switching is more than a clause? Can procurement preferences advance without clashing with EU competition rules or practical cost constraints—especially during tight AI capacity cycles? And will the Data Act’s switching rights translate into low-friction exits once contracts, egress fees and integration effort are counted? These uncertainties are real, but the policy and technical rails are now present. For Catalonia, the practical move is to use them: adopt GAIA-X artefacts in pilots, write Data Act portability into

procurement and grants, and tap European providers where locality or multivendor optionality matters—while remaining plugged into global ecosystems where they add value.

C3.5. Quibim: Radiology AI Ecosystem Partner

Quibim, a Valencia-based imaging startup, shows how an SME can scale inside a regulated hospital stack without trying to replace it. The company builds disease-specific AI modules, wraps them in transparent compliance artefacts, and plugs them into the clouds and devices hospitals already use—co-evolving with Microsoft Azure and Philips rather than fighting them. That choice gives Quibim distribution and credibility while letting it keep control of its IP, roadmap and regulatory posture.

The product strategy is deliberately granular. Modules such as QPProstate, QPBrain and QPLiver target single pathologies, obtain CE marks and FDA 510(k) clearances, and appear inside existing PACS viewers as an extra toolbar—so radiologists see auto-segmented lesions and quantitative biomarkers without leaving their workflow. Modularity lowers validation burden and speeds multi-country approvals: if a regulator flags a feature, Quibim can revise that component without freezing the whole suite.

Scale comes from partnering on rails, not renting them. On the cloud side, Quibim containerises inference on Azure (Kubernetes + AD + encrypted Blob Storage), gaining elastic capacity for screening spikes while retaining control over binaries and de-identified flows. A ten-million-image research biobank feeds model improvement within GDPR/HIPAA boundaries; hospitals benefit from single-sign-on and existing Azure commercial terms. On the device side, Philips integrations push DICOM streams through edge accelerators for denoising before light feature maps go to Azure—keeping round-trips inside clinical reporting windows and preserving a vendor-agnostic API for future Siemens/GE additions.

Outcomes, not demos, anchor credibility. A £2.6 million NHS pilot rolled QPProstate across seven hospitals and ~3,000 patients, with early data showing a 10% lift in early-stage detection and ~5 minutes saved per study. European consortia use the platform to harmonise multi-vendor datasets for Alzheimer's biomarkers, and biopharma partners (e.g., Janssen) mine imaging endpoints for trials via QPDiscovery.

The operating model matches our “roles–rules–data” emphasis. Hyperscalers supply secure, elastic infrastructure; device vendors expose acquisition hooks; Quibim occupies a high-trust, high-expertise niche and makes its audit trails inspectable. Funding and compliance reinforce the loop: a recent ~\$50 million Series A underwrites ISO 27001/GDPR/HIPAA regimes and automates documentation packs with every model update, turning regulation from overhead into advantage. The lesson for European SMEs is portable: specialise where clinical risk and legal complexity deter generic AI, integrate cleanly with dominant clouds and devices, and use measurable clinical impact to earn bargaining power rather than seek to be the platform.

C3.6. TruBlo: modular content-trust rails instead of a mega-platform

Disinformation, deepfakes and invisible edits have eroded confidence in online material. Europe's TruBlo (“Trusted Content on Future Blockchains”) takes a pragmatic route: rather than build a single platform, the NGI programme seeded a federated ecosystem of ~45 startups, research groups and micro-SMEs, each tackling a slice of authenticity (provenance, reputation,

tamper evidence) and connecting through open standards so modules can be snapped into existing publisher and social-media workflows.

The structure was three competitive open calls (2021–2023) distributing about €4.2 m in small cheques (up to €200k) plus mentoring from anchors like Worldline, Deutsche Welle and ICCS. Grants came with obligations: open-source deliverables and API blueprints to ensure interoperability. The result was a toolbox—proof-of-location for photos, SBOM-style trackers to tag AI-generated images, graph analytics to map how false narratives propagate—that larger actors can compose as needed.

Two SME stories illustrate the approach. WordProof turned its funding into a timestamping plugin for WordPress/Shopify that embeds an immutable edit certificate in page footers—helpful as EU policy moved toward verifiable transparency. TruthNest integrated social-graph credibility scores into federated feeds (Mastodon/Matrix), adding a “signals layer” usually available only on Big Tech stacks. Both parlayed the programme’s visibility into commercial pilots with newsrooms and fact-checking NGOs.

Interoperability is the point, not a slogan. Shared JSON-LD vocabularies and W3C Verifiable Credentials mean a WordProof timestamp can flow into a TruthNest dashboard or a media-forensics toolchain; broadcasters can overlay provenance tags on archives; storage providers can attach integrity proofs to uploads; AI-filter startups can call the common reputation API for context. This “lego” philosophy reduces lock-in while letting SMEs specialise deeply.

By late 2024, early deployments signalled traction beyond demos: Catena-X (the GAIA-X automotive lighthouse) experimented with TruBlo’s proof-of-origin protocol for parts’ digital twins; the EU’s data-spaces programme referenced TruBlo modules as preferred building blocks for media data spaces; a Southeast Asian marketplace tested the timestamping SDK to combat counterfeit product photos. Each pilot fed code and documentation back into the commons, strengthening the next wave.

The programme’s broader contribution aligns with this report’s emphasis on roles, rules and data. For SMEs, TruBlo lowers exploration risk, supplies credible interfaces, and lends policy-endorsed legitimacy that opens doors with risk-averse buyers. For regulators, it demonstrates an alternative to top-down mandates: fund usable tools, seed standards, and let market selection operate—timely as the DSA and forthcoming Media Freedom Act demand verifiable transparency from platforms.

Three lessons travel: small, milestone-tied cheques beat monolithic moonshots; interoperability-first licensing and APIs keep the ecosystem compounding; and public visibility (registries, demo days, hackathons) turns niche code into options large buyers can adopt.

Managerial takeaway: content trust won’t be delivered by a single hub. The modular rails model—open interfaces, composable proofs, neutral governance—lets publishers, platforms and civic actors assemble fit-for-purpose pipelines while keeping room for competition. Regions aiming to catalyse similar ecosystems can copy the playbook: run open calls, require interoperable artefacts, pair SMEs with anchor mentors, and publish a public registry so complements can find and combine each other. TruBlo shows how to design an ecosystem that many can join—and how smart public money can accelerate it without creating a new gatekeeper.

C3.7. Philips HealthSuite & the Elusive Promise of Digital Health

Philips built HealthSuite to solve a familiar problem in healthcare: data and workflows are fragmented across monitors, scanners, home devices, hospital IT and outside services, so the “patient journey” breaks at every hand-off. HealthSuite acts as a cloud spine that normalises clinical and device data, exposes standard interfaces to partners and lets hospitals compose end-to-end care flows without ripping out what they already run. The strategic shift is clear: instead of selling stand-alone equipment, Philips orchestrates an ecosystem where many firms can plug into the same rails—aligning with this report’s emphasis on roles, rules and data.

The operating model is disciplined. HealthSuite abstracts identity, consent and data governance; converts device and clinical records into shared formats; and publishes documented APIs so hospitals, device makers, digital therapeutics and home-care providers can interoperate. Hospitals keep data ownership and clinical control; Philips supplies the “minimum viable infrastructure”: interoperability (e.g., FHIR/DICOM gateways), auditability, and a contractable uptime and security envelope. That stance lowers onboarding costs for complementors and makes participation legible to risk-averse buyers.

In practice the ecosystem shows up as solved tasks rather than new portals. A COPD patient discharged with a connected spirometer and patch streams signals to HealthSuite; anomaly detection flags deterioration; the care team receives a task in its existing workflow; a home-care partner is scheduled; medication adherence is nudged through the pharmacy’s app; and outcomes—symptom days avoided, readmission risk—are attributed across contributors. Imaging follows the same pattern: scans are harmonised, quantitative reports flow back into the EHR, and a specialised AI module marks subtle changes without forcing radiologists into a new viewer. What changes is not the badge on the device but the legible path data and actions follow across organisations.

Incentives are built into the rails. Partners see closed-loop attribution—alerts handled, visits avoided, time-to-diagnosis shortened—so they can price services to outcomes rather than seat licences. Hospitals get vendor choice and reversibility: if a module underperforms, it can be swapped without unpicking the whole stack. Philips earns on the platform and on advanced services, but the economics improve for everyone only when measurable outcomes do—mirroring the report’s argument that measurement is the basis of legitimacy.

AI is connective tissue, not theatre. Models triage signals, fill data gaps, and generate first-pass reads or care-plan suggestions, but always with provenance, versioning and audit trails. Because algorithms run against a shared, governed data layer, SMEs with high-trust point solutions can participate without negotiating bespoke integrations at every hospital. That is AI as a participation accelerator—the same logic developed elsewhere in the report—lowering search and coordination costs so more complements can credibly join.

Regulatory and sovereignty concerns are handled by design. HealthSuite supports data residency, consent and retention controls suited to GDPR and medical-device regimes; interfaces are documented and change-logged; and contracts specify portability so providers can export artefacts (studies, reports, audit logs) if they switch. This keeps the orchestrator from becoming a gatekeeper and aligns the platform with Europe’s push for contestability and trustworthy data spaces.

The lesson for readers is portable. HealthSuite shows how a legacy manufacturer can turn its installed base and domain knowledge into shared infrastructure: keep the core neutral and

inspectable, make contribution measurable, and let a broader set of partners attach with minimal friction. That is how an ecosystem grows around critical services like healthcare without collapsing into a single hub—and why roles, rules and data, supported by pragmatic AI, are the real levers of scale.

References

Evolution Ltd Output

- Jacobides MG. 2024. How Will AI Impact Your Business? From Productivity to Strategic Transformation, [White Paper](#), January, Evolution Ltd.
- Jacobides MG. 2022. A Playbook for Winning the Ecosystem Game, [White Paper](#), December, Evolution Ltd.
- Jacobides MG. 2021. What Drives and Defines Digital Platform Power? A Framework, with an Illustration of App Dynamics in the Apple Ecosystem, [White Paper](#), April, Evolution Ltd.
- Jacobides MG, Bruncko M, Langen R. 2020. Regulating Big Tech in Europe: Why, So What, and How Understanding Their Business Models and Ecosystems Can Make a Difference, [White Paper](#), December, Evolution Ltd.
- Jacobides MG, Brusoni S, Candelon F. 2021. The Evolutionary Dynamics of the Artificial Intelligence Ecosystem, [Academic Paper](#), May, London Business School.
- Jacobides MG, Cennamo C, Gawer A. 2023. Externalities and Complementarities in Platforms and Ecosystems: From Structural Solutions to Endogenous Failures, [White Paper](#), July, Evolution Ltd.
- Jacobides MG, Langen R, Pusnakovs N, Romanenkov Y. 2022. Enhancing Brands and Shaping Customer Experiences: From Partnerships to Experience Ecosystems, [Briefing Paper](#), May, Evolution Ltd.
- Jacobides MG, Levi S, Tas J. 2020. First Steps, Far to Go: The Promises and Pitfalls of Platforms and Digital Ecosystems in Healthcare, [White Paper](#), December, Evolution Ltd. & Philips.
- Jacobides MG, Lianos I. 2021. Ecosystems and Competition Law in Theory and Practice, [Journal Article](#), November, Industrial and Corporate Change.
- Jacobides MG, Lianos I. 2021. Regulating Platforms and Ecosystems: An Introduction, [Journal Article](#) (Introduction), October, Industrial and Corporate Change.
- Jacobides MG, Lianos I. 2021. Regulating Platforms and Ecosystems: Introduction to the Industrial and Corporate Change Special Issue, [Submitted Manuscript](#), July, London Business School & UCL.
- Jacobides MG, Ma MD, Trantopoulos K, Vassalos V. 2023. Gamification: How to Climb the Ladders and Avoid the Snakes, [White Paper](#), January, Evolution Ltd.
- Jacobides MG, MacDuffie JP, Tae J. 2023. Revisiting Disruption: Lessons from Automobile Transformation and Mobility Innovation, [White Paper](#), July, Evolution Ltd.
- Jacobides MG, Reeves M, Fuller J. 2019. The Myths and Realities of Business Ecosystems, [Digital Article](#), February, MIT Sloan Management Review.
- Jacobides MG, Schneider J, Trepte F. 2024. The E-Mobility Ecosystem: Who Will Win the Race for the Customer?, [Research Paper](#), November, Evolution Ltd. & Strategy&.
- Jacobides MG, Sundararajan A, Van Alstyne M. 2019. Platforms and Ecosystems: Enabling the Digital Economy, [Briefing Paper](#), February, World Economic Forum.
- Teng N, Jacobides MG. 2022. The Shape-Shifting Dynamics of a Digital Platform Entrant: How Grab Navigated and Changed the Industry Architecture of Mobility in Southeast Asia, [Academic Paper](#), February, London Business School.

Further Readings

Adner R. 2021. Ecosystem Disruption vs. Technology Disruption: A Value Architecture Perspective, Presentation, September, Tuck School of Business, Dartmouth College.

Adner R. 2017. Ecosystem as Structure: An Actionable Construct for Strategy, Journal Article, January, Journal of Management.

Adner R. 2012. *The Wide Lens: A New Strategy for Innovation*. Portfolio / Penguin.

Baldwin, C.Y. 2024, *Design Rules, Volume 2: How Technology Shapes Organizations*. MIT Press.

Caffarra C, Gawer A, Jacobides MG. 2024. Mapping Antitrust onto Digital Ecosystems, [Journal Article](#), April, CPI Antitrust Chronicle.

Cenamor J. 2021. Complementor Competitive Advantage: A Framework for Strategic Decisions, [Journal Article](#), January, Journal of Business Research.

Jacobides MG. 2022. How to Compete When Industries Digitize and Collide: An Ecosystem Development Framework, [Journal Article](#), March, California Management Review.

Jacobides MG. 2019. In the Ecosystem Economy, What's Your Strategy?, [Digital Article](#), September–October, Harvard Business Review.

Jacobides MG, Cennamo C, Gawer A. 2024. Externalities and Complementarities in Platforms and Ecosystems: From Structural Solutions to Endogenous Failures, [Journal Article](#), January, Research Policy.

Jacobides MG, Cennamo C, Gawer A. 2018. Towards a Theory of Ecosystems, [Journal Article](#), June, Strategic Management Journal.

Jacobides MG, Lianos I. 2021. Ecosystems and Competition Law in Theory and Practice, [Journal Article](#), November, Industrial and Corporate Change.

Jacobides MG, Reeves M, Fuller J. 2019. The Myths and Realities of Business Ecosystems, [Digital Article](#), February, MIT Sloan Management Review.

Jacobides MG, Schneider J, Trepte F. 2024. The E-Mobility Ecosystem: Who Will Win the Race for the Customer?, [Research Paper](#), November, Evolution Ltd. & Strategy&.

Pidun U, Reeves M, Schüssler M. 2020. How Do You Design a Business Ecosystem?, [Briefing Paper](#), January, BCG Henderson Institute.

Pidun U, Reeves M, Schüssler M. 2019. Do You Need a Business Ecosystem?, [Briefing Paper](#), May, BCG Henderson Institute.

Pidun U, Reeves M, Zoletnik B. 2021. How to Benefit From Business Ecosystems If You Are Not the Orchestrator?, [Article](#), September, BCG Henderson Institute.

Scott Morton FM. 2025. Digital Platform Regulation: Making Markets Work for People, [Monograph](#), October, Independent.

Thomas LDW, Autio E. 2020. Innovation Ecosystems, [Working Paper](#), January, SSRN.

Links to Case Studies

C1.1 – The Swedish Sydved Forestry Ecosystem

<https://forestry.trimble.com/sydved-selects-cfharvest-to-digitalize-its-forestry-harvesting-operations/>

C1.2 – InSoil – Spanning Funding to Agri-Regeneration

<https://insoil.com/>

C1.3 – Busroot – Manufacturing Analytics Platform

<https://www.output.industries/>

C1.4 – A Phygital Ecosystem: MAF and the Creation of the SHARE Ecosystem

<https://publishing.london.edu/cases/building-majid-al-futtaims-phygital-future-orchestrating-a-retail-ecosystem-compact-case/>

<https://www.majidalfuttaim.com/en/media-centre/press-releases/detail/2023/06/majid-al-futtaim-partners-with-tabby-to-introduce-flexible-payment-solutions>

C1.5 – From Fridges and Stoves to the Internet of Food

<https://publishing.london.edu/cases/from-products-to-experience-ecosystems-haiers-internet-of-food/>

C1.6 – Linking Public Transport with Private Sector and Incentives: Velocia

<https://velocia.io/rider-incentive-program/>

C1.7 – Supporting Green Transition Through an Ecosystem Play: Qurator

<https://www.qurator.com/>

C1.8 – Opportunities and Challenges in the Shift from Cars to an Electromobility Ecosystem

<https://www.evolutionltd.net/post/the-e-mobility-ecosystem-who-will-win-the-race-for-the-customer>

C2.1 – Growth Through Ecosystem Use in MasMovil and COSMOTE

<https://www.masmovil.es/>

<https://www.cosmote.gr/hub/>

C2.2 – KBC: intent-driven orchestration with Kate

<https://www.kbcbrussels.be/retail/en/products/payments/self-banking/on-your-smartphone/mobile/mobile-faq/Kate-FAQ.html>

C2.3 – DBS's POSB Ecosystem with “School Buddy”

https://www.dbs.com/newsroom/Worlds_first_in_school_wearable_tech_savings_and_payments_programme_POSB_Smart_Buddy_launches_in_19_schools

C2.4 – Ecosystem Support/Advisory: Synocus, FuzeQube, Evolution Ltd

— Synocus

<https://www.synocus.com/ecosystems/>

— FuzeQube

<https://www.fuzeqube.com/>

— Evolution Ltd

<https://www.evolutionltd.net/>

Case C2.6 – Kinaxis: turning supply-chain planning into an ecosystem

<https://www.kinaxis.com/en>

<https://supplychaindigital.com/news/kinaxis-ai-orchestration-for-resilient-supply-chains>

C2.7 – Exnaton: Enabling SMEs to Orchestrate Energy Communities

<https://www.exnaton.ai/en/renewable-energy-communities>

<https://www.exnaton.ai/en/energy-sharing>

C2.8 – Katana MRP: Plug-in Ecosystem for Manufacturing SME

<https://katanamrp.com/small-business-manufacturing-software/>

<https://katanamrp.com/blog/smart-manufacturing/>

C3.1 – The Dilemmas of Smaller Organizations Working with BigTech: Life Under Amazon's Shadow for Airplay

<https://www.hbs.edu/faculty/Pages/item.aspx?num=52247>

C3.2 – Shopify as an Engine and an Opportunity

<https://www.shopify.com/partners>

C3.3 – Nvidia and the growth of AI through ecosystems

C3.4 – GAIA-X & Eurostack: Europe's Pushback on Platform Dominance

<https://ainowinstitute.org/publications/x-european-digital-independence-building-the-eurostack>

C3.5 – Quibim: Radiology AI Ecosystem Partner

<https://quibim.com/>

C3.6 – TruBlo: modular content-trust rails instead of a mega-platform

<https://cordis.europa.eu/project/id/957228/results>

C3.7 – Philips HealthSuite & the Elusive Promise of Digital Health

<https://www.evolutionltd.net/post/digital-platforms-and-ecosystems-in-healthcare-promises-and-pitfalls>



Evolution Ltd is a boutique advisory that combines frontier research from world-class business academics and technologists with hands-on experience from senior executives to guide organizations in an increasingly complicated environment.

Evolution focuses on digital ecosystems, Artificial Intelligence and their impact on strategy and organization. Its independence and governance structure ensure rigor and bespoke solutions for its clients and inspire hands-on, [award-winning frameworks](#) that shape managerial practice.

Its clients and partners include large corporates, leading consultancies, governments and NGOs. Projects draw on its affiliates' vast experience and connections to global tech giants, startups, disruptors, entrepreneurs, and governments alike to engage key stakeholders in effective conversations and catalyse action.

www.evolutionltd.net



Generalitat
de Catalunya
**Departament
d'Empresa i Treball**

The Generalitat de Catalunya is the regional government and public administration of Catalonia (Spain). With its own powers in areas such as education, culture, research, and industry, the Generalitat promotes economic development, innovation, and the sustainability of Catalonia's productive sector.

This book reflects its commitment to advancing industrial transformation as a key driver of sustainable economic growth and a more cohesive society. One of the Generalitat's strategic instruments in this field is the Industry Observatory, established within the Directorate General for Industry under the first National Pact for Industry (2017–20). The Observatory plays a central role in monitoring the implementation and evolution of the successive National Pacts for Industry (2017–20, 2022–25, and the upcoming 2026–30), while also producing sectoral and forward-looking studies on industrial economics to inform policy and guide future action.

<https://empresa.gencat.cat/ca/inici/>